Vision-Based Control of a Handheld Micromanipulator for Robot-Assisted Retinal Surgery

THESIS PROPOSAL

BRIAN C. BECKER
briancbecker@cmu.edu

THESIS COMMITTEE

Cameron N. Riviere (Chair)
George A. Kantor
George D. Stetten
Gregory D. Hager (Johns Hopkins University)

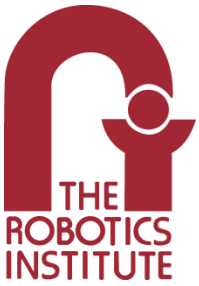

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, US, 15213

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

ABSTRACT

Surgeons increasingly need to perform complex operations on extremely small anatomy. Many promising new surgeries are effective, but difficult or impossible to perform because humans lack the extraordinary control required at sub-mm scales. Using micromanipulators, surgeons gain better positioning accuracy and additional dexterity as the instrument smoothes tremor and scales hand motions. While these aids are advantageous, they do not actively consider the goals or intentions of the operator and thus cannot provide context-specific behaviors, such as motion scaling around anatomical targets, prevention of unwanted contact with pre-defined tissue areas, and other helpful task-dependent actions.

This thesis explores the fusion of visual information with micromanipulator control and builds a framework of task-specific behaviors that respond synergistically with surgeon's intentions and motions throughout surgical procedures. By exploiting real-time microscope view observations, a-priori knowledge of surgical procedures, and pre-operative data used by the surgeon while preparing for the surgery, we hypothesize that the micromanipulator can better understand the goals of a given procedure and deploy individualized aids in addition to tremor suppression to further help the surgeon. Specifically, we propose a vision-based control framework of modular virtual fixtures for handheld micromanipulator robots. Virtual fixtures include constraints such as "maintain tip position", "avoid these areas", "follow a trajectory", and "keep an orientation" whose parameters are derived from visual information, either pre-operatively or in real-time, and are enforced by the control system. Combining individual modules allows for complex task-specific behaviors that monitor the surgeon's actions relative to the anatomy and react appropriately to cooperatively accomplish the surgical procedure.

Particular focus is given to vitreoretinal surgery as a testbed for vision-based control because several new and promising surgical techniques in the eye depend on fine manipulations of delicate retinal structures. Preliminary experiments with Micron, the micromanipulator developed in our lab, demonstrate that vision-based control can improve accuracy and increase usability for difficult retinal operations, such as laser photocoagulation and vessel cannulation. An initial framework for virtual fixtures has been developed and shown to significantly reduce error in synthetic tests if the structure of the surgeon's motions is known. Proposed work includes formalizing the virtual fixtures framework, incorporating elements from model predictive control, improving 3D vision imaging of retinal structures, and conducting experiments with an experienced retinal surgeon. Results from experiments with *ex vivo* phantoms and *in vivo* tissue for selected retinal surgical procedures will validate our approach.

TABLE OF CONTENTS

1 INTRODUCTION

Microsurgery is hard, even for experienced surgeons. As surgeons attempt to manipulate small anatomy in the sub-mm range, human control of surgical instruments becomes degraded [88]. Retinal vasculature is often less than 100 μm and membranes such as the Internal Limiting Membrane (ILM) in the eye are only tens of microns thick [21]. New surgical procedures depend on safely manipulating such small structures [2, 8, 23, 47, 87, 103, 104, 115]. For operations on extremely small anatomy, physiological tremor becomes a serious challenge. Peak-to-peak tremor of vitreoretinal surgeons can exceed 100 μm [105], creating potentially dangerous unintentional movements and loss of dexterity which can lead to the inability to successfully complete surgical procedures.

A popular solution to the problem of manipulation at small scales is to use robotic aids, such as micromanipulators [43]. Micromanipulators provide the surgeon a method to more precisely control μm-level movements. Usually, micromanipulators feature a number of useful behaviors, including tremor suppression, motion scaling, and limits on movement speed or operating area. Tremor suppression acts as a filter to reduce involuntary, physiological human movements that usually occur around 8-20 Hz [51] while allowing lower-frequency voluntary motion. Motion scaling increases movement precision by reducing all movements made by the operator by a fixed amount, so large motions become identical, but smaller motions. For safety, often the robot enforces a maximum velocity on the tip of the instrument and can prevent the surgeon from straying too far away from the center of the workspace. These features are generally either always active (i.e. tremor suppression) or activated manually by the operator. As evidenced by the widespread use of surgical robots such as the da Vinci by Intuitive Surgical and the published work on micromanipulators [17, 31, 80, 94, 110], these tools represent an effective way to tackle delicate surgeries.

For the purposes of this thesis, we define a taxonomy of micromanipulators consisting of three classes: master/slave, cooperative, and handheld. Master/slave micromanipulators separate the human input into one subsystem and the instrument actuation into another (often remotely located) subsystem. Surgical Intuitive's da Vinci robot [17] is perhaps the most popular master/slave configuration and has been outfitted with a prototype manipulator for intraocular microsurgery [82]. The micromanipulators designed by Hunter et al. [54], Jensen et al. [56] and Ueta et al. [110] are more examples of a master-slave robots designed specifically for retinal microsurgery. Because there is no mechanical linkage between the input and output, the available workspace is large and control is usually quite versatile as arbitrary tremor suppression filters and behaviors can modify the signals on the drive-by-wire connection. However, the additional pathways from input to control introduce extra latency and remove physical feedback, such as the ability to sense forces applied to the tissue. Because of the equipment required for the two separate subsystems and the complexity of synchronizing them, master/slave configurations are often more expensive than other micromanipulators.

The class of cooperative micromanipulators involves robotic arms and linkages that the human can grasp and where control is jointly shared between the human and the robot. Johns Hopkins University's Steady-Hand robot [80, 111] is a classic example of a robot arm for microsurgery that

is operated cooperatively by the human and robot. The human guides the micromanipulator while the robot acts as a mechanical low-pass filter to smooth movement. Hard constraints are particularly easy to enforce with non-backdrivable linkages. Because of the inherent stiffness in the mechanical linkages, movement is often sluggish and pressures exerted on tissue by the tip are not felt by the operator. More precise movements are usually achieved with virtual fixtures by commanding scaled velocities based on the forces applied to the micromanipulator [16]. A number of virtual fixtures frameworks have been developed for the SteadyHand [1, 16, 58, 68] based on the forces applied to the cooperatively controlled handle, providing even more fine-grained control to the operator.

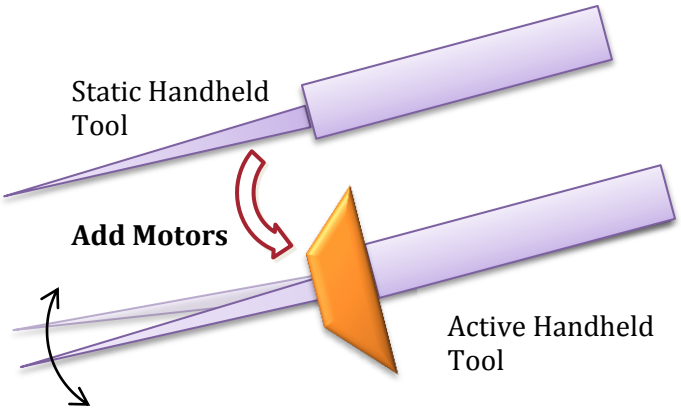

Fig. 2. Handheld instruments contain actuators that move the tip relative to the handle.

The micromanipulator class of most interest to the focus of this thesis is the handheld micromanipulator proposed by Riviere et al. [93, 96]. In a handheld micromanipulator, the entire instrument, including actuation, is contained in the shaft of the instrument, and control of the instrument tip happens by offsetting the tip relative to the current handheld position (see Fig. 2). The Micron robot [75, 94] developed in our lab at Carnegie Mellon University is the first of its kind, a fully handheld micromanipulator not much larger than a typical un-actuated surgical tool. Behaviors are accomplished by moving the tip relative to the hand motion via actuators between the handle and the tip. For instance, high-speed control performs tremor suppression by rapidly moving the instrument opposite to hand tremor. While remaining very lightweight and flexible, handheld micromanipulators have a limited range of motion and require high rate sensing and control to actively compensate for hand motion in addition to enforcing behaviors. However, there are a number of key advantages to this design including familiarity, reliability, and inexpensiveness. Handheld designs are familiar to surgeons, who can leverage their extensive experience with handheld instruments operated under traditional surgical-grade microscopes to avoid lengthy training processes. If the instrument stops working properly, a surgeon can simply turn it off, and complete the surgery as usual with the inert handheld instrument. This increases safety as the procedure can always be performed no differently than a typical, unaided surgery with regular tools. Likewise in case

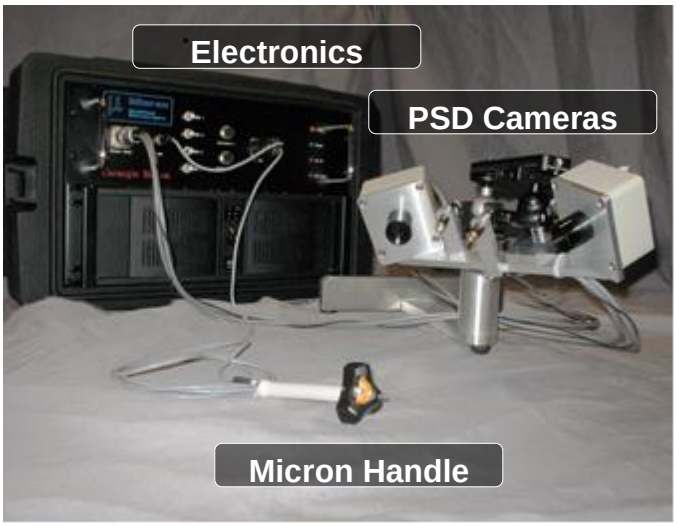

Fig. 1. Micron system with electronics connecting the PSD trackers and the actuated handheld tool.

of device malfunction, a surgeon can easily and rapidly move the tool's small form factor away from any tissue. Finally, handheld micromanipulators are often comparatively inexpensive. All proposed research will be performed with Micron, a tool with demonstrated benefits including tremor compensation, motion-scaling, and other aids shown in experimental studies with surgeon subjects [11, 14, 30].

1.1 MICRON: HANDHELD MICROMANIPULATOR

Handheld micromanipulators present a challenging engineering problem, requiring high control bandwidth, low error position measurements, and sufficient range-of-motion to complete the desired micromanipulation tasks. To date, Micron and its descendent design ITrem [63], are the only handheld micromanipulators designs that have been built and reported. We briefly describe the Micron system, shown in Fig. 1.

1.1.1 HANDPIECE

Micron, pictured in Fig. 3, is a fully handheld micromanipulator that self-actuates via Thunder® TH-10R piezoelectric actuators (Face International Corp., Norfolk, Virginia, USA) located between the handle grip and the tip of the instrument. Arranged in a 3-pointed star configuration, the piezoelectric actuators can flex in and out individually to give Micron a 3DOF range of movement that is approximately 0.5 mm axially and 1.8 mm transversely, depending on the tool mounted to Micron's base. See Fig. 4 for the arrangement of motors and a diagram of actuator operation. Actuators are driven with charge control at 2 kHz in a closed loop feedback control system that uses the position measurement system described in the following section. Control wires, laser fiber optics, and surgical instrumentation fits down the hollow center of the shaft and optionally to the tip of the instrument via flexible linkages.

1.1.2 POSITION MEASUREMENT

Low-latency, high-bandwidth positioning information is obtained from custom optical tracking hardware named ASAP [76]. Two position sensitive detectors (PSDs) at 60° track LEDs mounted to Micron within a 4 cm^3 workspace at a rate of 2 kHz, with measurement accuracies of <10 μm RMS at the tip of the instrument. The 3D positions of three LEDs on the actuated shaft holding the end-effector are reconstructed and used to calculate [53] the 6DOF position tracking of Micron's tip, as shown in Fig. 3. An additional LED attached to the handle provides hand movement information.

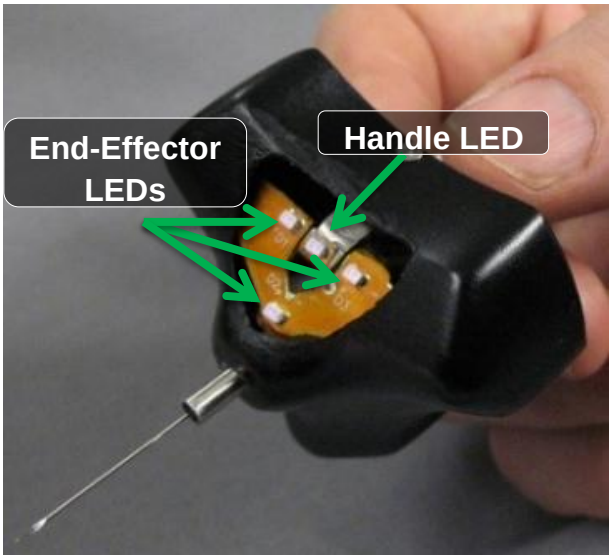

Fig. 3. Micron 6DOF sensing of the tip pose and handle pose with LEDs tracked by Position Sensitive Detectors (PSDs).

1.1.3 TREMOR SUPPRESSION

Tremor compensation can be achieved by inserting filters between the hand motion and the drive mechanism to reduce the higher-frequency tremor while retaining the operator's

gross lower-frequency movements. A 2nd order, low-pass Butterworth filter based on the handle motion is used to implement tremor reduction. Motion scaling, often used in conjunction with tremor suppression, scales hand movements so the tip only moves a set fraction of the hand movement [14]. For example, a scale factor of ½ would transform hand movement of 2 mm to 1-mm movement at the robotic tip. Recent developments in tremor compensation include a low-pass shelving filter [6], which acts as a hybrid tremor suppression filter that affords relative motion scaling and reports an additional 30-50% reduction in tremor compared to the low-pass only filter.

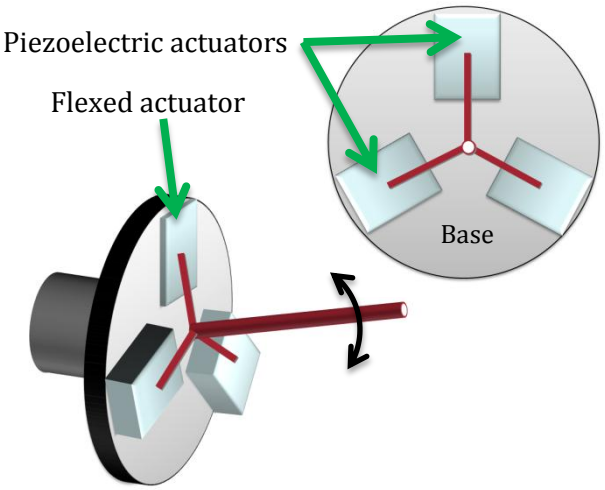

Fig. 4. Micron actuation with piezeoelectric bender motors mounted in a 3-pointed star configuration. Tip movement is caused by flexing the bender motors.

1.1.4 Microscope & Cameras

The operator uses a Zeiss® OPMI® 1 microscope to view the workspace under 25x magnification. Two Flea2® cameras (Point Grey Research, Richmond, B.C., Canada) are mounted to the eyepiece, acquiring 30-Hz stereo video of the 2x3 mm to 6x9-mm workspace at a resolution of 800x600. The real-time video is displayed to the operator on a 22" 3D computer display (Trimon ZM-M220W, Zalman Tech Co., Ltd., Seoul, Korea). The computer display enables the system to easily overlay extra information on the real-time video, such as target locations or depth cues, to aid the operator.

1.2 Problem Statement

Unfortunately, most existing techniques utilized in robotic micromanipulators such as tremor suppression and motion scaling lack any sort of understanding of the surgeon's intentions, plans, or goals. Used as black-boxes, each behavior operates blindly upon activation in the same manner regardless of what the surgeon needs or the situation requires. It is the surgeon's sole responsibility to turn behaviors on and off as well as tweak parameters to keep the behaviors at maximum helpfulness during the entire procedure. For instance, a robotic system may have a motion-scaling knob defining the scaling parameter that the operator must manually change when nearing the targeted tissue or when transitioning between steps in a procedure. In fact, it is possible that a behavior designed for aid may become a hindrance if ill-tuned or accidentally applied inopportunely. The fundamental assumption underlying this thesis is that further aids to surgeons necessitate at least basic domain knowledge of surgical procedures coupled with real-time monitoring and at least rudimentary understanding of a surgeon's actions and goals.

Much like an auto-pilot system eliminates the need for constant course corrections and tweaking many different knobs by understanding the aviator's flight plan and the current situation via instrumentation, our goal is to provide task-specific automation behaviors by understanding the surgeon's *operation plan* and the current state of the procedure *via visual feedback*. Just as an auto-pilot does not eliminate the need for trained and experienced pilots, neither should a robotic surgical system supplant surgeons. Instead, it will only provide them with methods to reduce

tedious repetition and automate mundane procedures so surgeons are free to exploit their extensive judgment and experience without having to worry about low-level performance.

A common approach is to use visual information and track the surgeon's actions relative to anatomy during the procedure. A significant body of research is dedicated to related algorithms such as tracking surgical instruments [89], understanding surgical actions [3, 91], and integrating and evaluating such technologies in robotic systems [11, 68, 80]. While all are necessary for a comprehensive robotic system that is able to synergistically aid surgeons during operations, one particularly promising avenue of research is virtual fixtures. Virtual fixtures (VF) enforce constraints on the behavior of the instrument by software controlled hardware. The motivation behind virtual fixtures is that if the exact behavior the operator wants to follow is known, inaccuracies can be eliminated by introducing an artificial fixture that constrains the user to desired behavior. For instance, if you know an artist wants to draw a straight line or a particular shape, a ruler or stencil improves performance by more precisely controlling the drawing instrument than the unaided operator could otherwise manage. For domains such as drafting where the work consists of many simple primitives, extensive use of fixtures is beneficial. Some surgical procedures lend themselves well to compositions of primitive movements, a fact exploited by researchers to use virtual fixtures to show beneficial performance improvements in specific surgeries [16, 66-68].

This thesis focuses on combining virtual fixtures, tremor suppression, motion scaling, and other behaviors into a unified vision-based control framework for the unique capabilities and challenges of handheld micromanipulators. While keeping the theoretical aspects of our work applicable in the general surgical mindset, this thesis will focus on applications to retinal surgery, often considered one of the most demanding areas of microsurgery because of the extremely small structures that must be routinely manipulated [97]. Using domain knowledge coupled with real-time vision analysis of the surgeon's actions during the procedure, the proposed theoretical framework and practical system implementation will provide task-specific behavioral aids to the surgeon. Work in necessary prerequisite research including model based control and 3D visual sensing will enable the application of the proposed framework to targeted retinal surgeries.

1.3 Thesis Statement

We propose three hypotheses which this thesis will investigate:

1. Micromanipulation tasks can be decomposed into steps, during which applying tip position constraints (virtual fixtures) with a handheld tool increase operator precision and safety.
2. Real-time visual analysis of anatomy and the surgeon's movements provides the context that allows the controller to select and customize the virtual fixtures.
3. Vision-based virtual fixtures aid surgeons during microsurgical operations.

Summarizing these hypotheses yields our thesis statement:

> *Fusing vision information into the control system of a handheld micromanipulator enables a rich set of task-specific behaviors known as virtual fixtures that can be composited and applied semi-automatically to reduce error and increase performance during microsurgery, such as vitreoretinal operations.*

2 PROBLEM DOMAIN: VITRORETINAL MICROSURGERY

The theoretical aspects of handheld micromanipulation presented in this thesis are kept generalizable to many different types of micromanipulation, ranging from microsurgery to industrial processes to biomedical experimentation. However, it is neither practical nor advantageous to design and conduct experimentation covering such a wide variety of manipulation tasks. Instead, this thesis focuses on microsurgery, and in particular retinal microsurgery. Although many areas of surgery are challenging from a manipulation viewpoint, vitroretinal surgery is considered to be one of the most demanding fields of microsurgery [52]. Thus we propose to use retinal surgery as the practical experimental testbed for the more general theoretical frameworks proposed.

This section highlights some of the common procedures in the retina and a few promising techniques developed recently. The purpose of this chapter is two-fold. First, it serves to motivate the continued development of handheld micromanipulators (see Fig. 8) and drive the design requirements for the framework of beneficial micromanipulator behaviors. Second, it provides important background information for the experiments proposed to validate the proposed thesis work. We begin by overviewing the structure of the human eye and describing the general group of retinal surgeries known as pars plana vitrectomies. We then overview some of the specific surgeries that are performed during a vitrectomy.

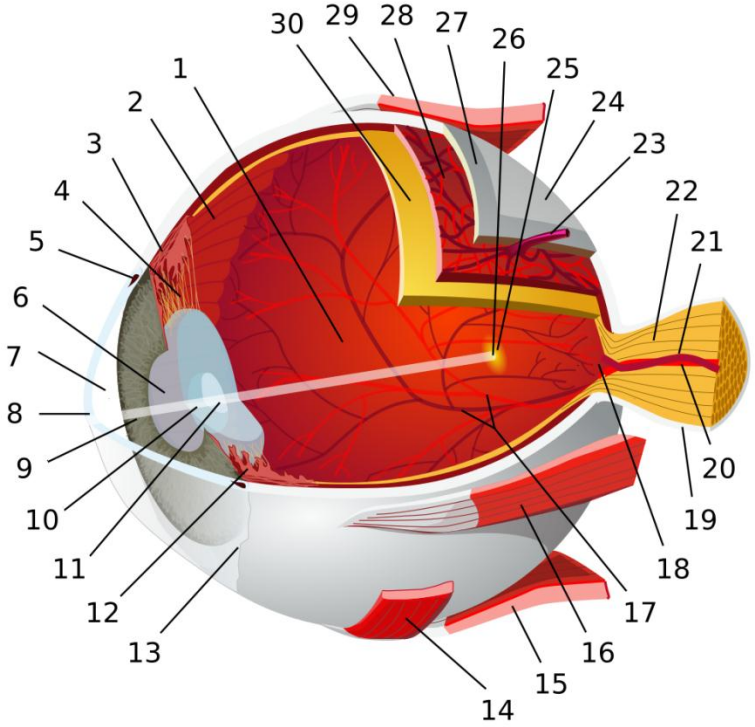

Fig. 5. Diagram of the human eye (credit to chabacano, used under Creative Commons Attribution-Share Alike 3.0 license). 1:posterior vestibule 2:ora serrata 3:ciliary muscle 4:ciliary zonules 5:canal of Schlemm 6:pupil 7:anterior chamber 8:cornea 9:iris 10:lens cortex 11:lens nucleus 12:ciliary process 13:conjunctiva 14:inferior oblique muscle 15:inferior rectus muscle 16:medial rectus muscle 17:retinal arteries and veins 18:optic disc 19:dura mater 20:central retinal artery 21:central retinal vein 22:optic nerve 23:vorticose vein 24:bulbar sheath 25:macula 26:fovea 27:sclera 28:choroid 29:superior rectus muscle 30:retina.

2.1 ANATOMY OF THE EYE

The human eye is a complex organ with many subsystems, as depicted in Fig. 5. The entire eyeball is held inside the sclera (27), a strong, thin hollow sphere that protects the eye. Light enters the clear cornea (8) through the pupil (6) in the iris (9) and is focused by the lens (11). It passes through a clear, jelly-like substance known as vitreous humor (also referred to as the vitreous body or just

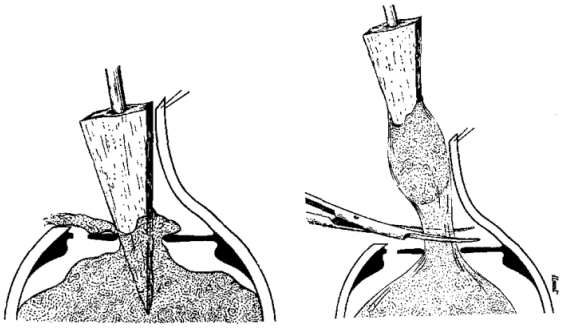

Fig. 6. Example of the "open sky" approach to removing vitreous, where a "toilet shape" U cut is made. A sponge is used to draw the vitreous, which is subsequently cut away. Photo credit to [73].

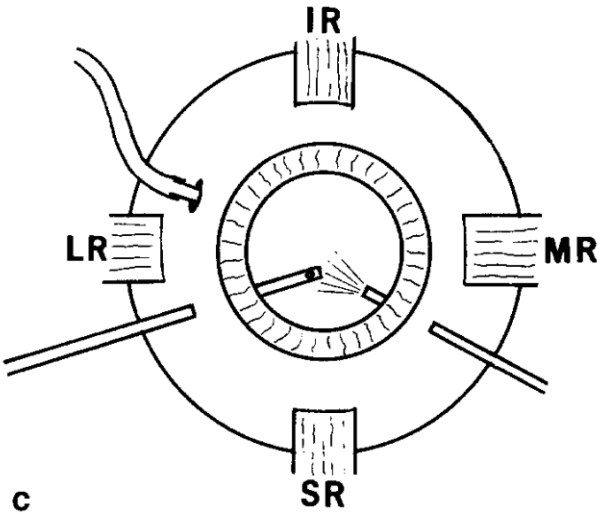

Fig. 7. Pars plana vitrectomy with three ports through the sclera for infusion, light, and an active tool. Photo credit to [73].

vitreous) that keeps the eye inflated. After passing through the vitreous, light hits the nerves in the retina (30), which transmits the light information to the brain via the optic nerve (20). The retina is approximately 130 μm to 550 μm thick and is nourished by various vasculature (17, 20, 21). The retina is separated from the vitreous by membranes such as the internal limiting membrane. Various problems occurring in the retina require surgical intervention, as detailed in subsequent sections.

2.2 Pars Plana Vitrectomy

Surgical access to the retina is difficult because of its placement behind the sclera and vitreous. In the late 1960s, Kasner developed an "open sky" vitrectomy technique for cataract surgery [59]. His solution was to cut through the sclera and remove the vitreous with a procedure known as a vitrectomy. The term vitrectomy literally means to excise the vitreous. During a vitrectomy, a "U" was cut into the scleara, the cornea temporarily removed, and the vitreous excised with sponge and scissors (see Fig. 6). Macherner [73] is credited with popularizing the general technique known as pars plana vitrectomy, which provides much less intrusive access to the retina. He extended Kasner's technique by replacing the "open sky" approach of making a large U-shaped incision in the sclera with several small ports in the pars plana section of the sclera through which instruments could be inserted. Furthermore, he developed a vitreous cutter that could be inserted through the port, cutting and removed the vitreous via suction pumps. Generally, a pars plana vitrectomy is performed as shown in Fig. 7, with three ports for illumination, an active tool (such as vitreous cutter, forceps, etc.), and infusion to maintain proper ocular pressure. Throughout the 70s, the pars plana vitrectomy procedure was used to treat previously incurable problems, including non-clearing vitreous hemorrhages, vitreous traction causing retinal detachment, and giant retinal tears. Today, one may view the pars plana vitrectomy (usually referred to as just a vitrectomy) as a general procedure, which may have many different goals and sub-procedures. Common retinal problems and associated sub-procedures performed during treatment are listed in Table 1.

Problem	Laser Photocoagulation	Vessel Cannulation	Membrane Peeling	Arteriovenous Sheathotomy
Diabetic Retinopathy	X			
Retinal Detachment	X		X	
Epiretinal Membrane			X	
Macular Hole			X	
Vessel Occlusion	X	X		X
Foreign Body Removal	X			
Macular Degeneration	X			

Table 1: Various problems occurring in the retina and corresponding solutions that are used by vitreoretinal surgeons.

We choose this subset of procedures because they are known to be particularly challenging. In the following sections, we discuss each vitrectomy treatment option briefly and describe micromanipulator behaviors that could be applicable to the procedure.

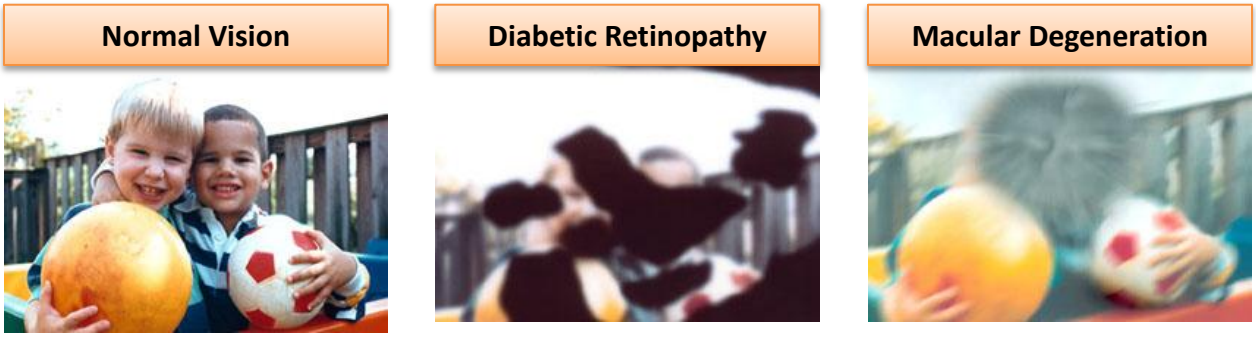

Fig. 8. Motivation for this work derives from retinal conditions and their effect on eyesight. Laser photocoagulation treats retinal ailments such as diabetic retinopathy and macular degeneration, which affect 4.1 million and 1.8 million Americans, respectively. Photo credit to NIH.

2.3 LASER PHOTOCOAGULATION

Laser photocoagulation of the retina is a common adjunct in pars plana vitrectomy surgery for common diseases such as diabetic retinopathy [104], retinal detachment [47], macular edema [8, 103], branch vein occlusion [87], and intraocular foreign body removal [2]. After the anatomical goals of clearing the media and successfully re-opposing the retina to the pigment epithelium, three common laser procedures are used. Panretinal photocoagulation lowers the production of VEGF by the ischemic retinal cells and decreases the potential for regrowth of neovascular tissue. Peripheral panretinal photocoagulation in retinal detachment seals retinal breaks and stimulates fibrous metaplasia which prevents retinal breaks and retinotomies from leading to re-detachment of the retina. Focal laser patterns are used to surround retinal breaks iatrogenically created to drain subretinal fluid, to treat accidental retinal breaks occurring when scar tissue is being peeled, and to surround traumatic retinal breaks prior to removing intraocular foreign bodies.

Accuracy in laser photocoagulation is important for optimal clinical results as inadvertent photocoagulation of a retinal vein can cause occlusion of the vein, possibly leading to vitreous hemorrhage [55]. Laser applications frequently are applied within one millimeter of the foveola,

requiring careful avoidance of unintended targets such as the optic nerve and the fovea to avoid permanent central vision loss [42]. Furthermore, accidental coagulation of the macular venule or arteriole can cause foveal ischaemia or intraretinal fibrosis [64]. For best results in previously photocoagulated retinas, previous burn locations are also to be avoided [101]. Research efforts to improve positioning of burns have yielded automated approaches [78], but system ergonomics and complexity have prevented clinical adaptation [18].

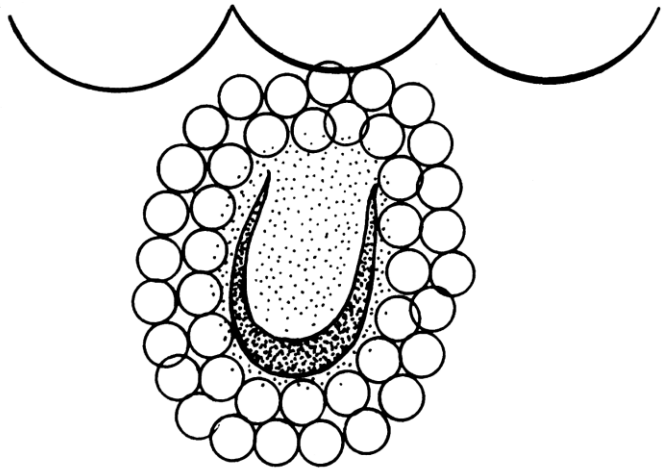

Fig 9. Placement of laser burns around a retinal break. Similar ring patterns are used around macular holes and other retinal work. Photo credit [86].

2.3.1 Treatment Procedure

Laser photocoagulation is used in a wide variety of treatments, but can be grouped into two typical scenarios. To seal retinal breaks, two or three rings of laser burns as seen in Fig 9 are applied around the surgery site [86]. For pan-retinal photocoagulation, a grid of hundreds to thousands of laser burns are applied to the retina. Burn sizes are typically around 200-400 μm in diameter with laser pulse durations of 10-100 ms. Currently, laser pulses are applied at a constant rate of 0.5 – 2 Hz when the surgeon depresses the foot pedal.

2.3.2 Behavioral Aids

- **Select Target**: Select the best target based on closest targets and hand motion.
- **Maintain Orientation**: Keep the tip (and thus the laser) pointed at a targeted burn location to increase accuracy of burns.
- **Fire Laser**: Automatically fire laser upon target acquisition.
- **Avoid Zones**: Prevent the laser from burning sensitive anatomy such as vessels.
- **Maintain Distance**: Keep the tip of the fiber optic within a specified range of the retina to prevent too hot or too cold burns.

2.4 Vessel Cannulation

New microsurgical procedures in the eye are yielding encouraging results in the treatment of diseases such as retinal vein occlusion (RVO). RVO occurs when a clot obstructs blood flow in a central or branch vein of the eye (CRVO and BRVO, respectively). As the second most common retinal vascular disease, RVO affects an estimated 16.4 million adults worldwide [107] and has no proven effective treatment [83]. A promising experimental surgical procedure is retinal endovascular surgery (REVS), which involves cannulating the vein and directly injecting clot-dissolving plasminogen activator (t-PA) to remove obstructions [23]. In a test of 28 eyes with CRVO, 54% recovered more than 2 lines of visual acuity within 3 months [117], although follow up studies have reported mixed results [23, 39].

In vivo retinal vessel cannulations were first reported with cats and rabbits using a micromanipulator to fix the micropipette and damp hand movements [4]. Although subsequent research has demonstrated successful cannulations in a variety of animal and human models, there are difficulties involved in manipulating such tiny vessels. It has been noted that larger vessels are easier to puncture than smaller vessels more distal to the optic nerve [109]. In experiments with human cadaver eye models, Tang and Han recorded that only 10 out of 18 attempted cannulations were successful and commented that manipulating vessels with currently available instrumentation is generally traumatic [108]. Because retinal vessels are so small (less than 200 μm), it is common to cannulate larger, nearby vessels, which is thought to lessen efficacy [116]. Overall, the procedure is considered difficult to perform [57] primarily because of limited dexterity while keeping the cannula in the vessel [116].

2.4.1 Treatment Procedure

Before cannulation, intraocular pressure is lowered to about 5 mm Hg, allowing vessels to dilate; following vessel puncture, the pressure should be increased to lessen bleeding [116]. Once the vessel has been punctured, 3.4 ml of rt-PA with a concentration of 200 μg/ml is injected into a cannulated vessel towards the optic nerve head [52]. The infusion process takes several minutes, during which the cannula should remain in the vessel.

2.4.2 Behavioral Aids

- **Steady Approach**: Since tremor amplitude can be larger than the width of the vessel, removal of tremor is important (indeed, [46], [4], and [116] use a micromanipulator during the procedure).
- **Motion Scaling**: To increase precision with <100 μm vessels, motion around the vessel can be scaled. Anisotropic scaling in the direction of the vein allows unimpeded thrusting puncture movements.
- **Maintain Position**: During infusion, the position of the cannula should be kept as motionless as possible to prevent trauma to the vessel or accidental dislodging of the cannula.
- **Prevention Zones:** Keep the tip of the instrument away from the retina to prevent tears or tissue trauma.

2.5 Membrane Peeling

The purpose of membrane peeling is to remove either the epiretinal or internal limiting membranes on the top of the retina. Epiretinal membranes are thin fibrous layers that form on the retina gradually over time or in conjunction with other retinal diseases. They obscure vision and are often referred to as a “macular pucker” because as they develop, they can distort and pucker the surface of the retina, causing straight lines to appear wavy. Removing the eipiretinal membrane is necessary to restore visual acuity.

The internal limiting membrane (ILM), a very thin layer that separates the vitreous from the retina, and is often removed during the treatment of macular holes. Macular holes occur when the vitreous pulls away and tears the retina, a condition that is common as the eye ages [86]. Kelly and Wendell first demonstrated using vitrectomy procedures to relieve pressure from the vitreous on the retina

[60]. To facilitate closure of the macular hole, some researchers propose removal of the ILM around the hole to alleviate tangential pressure from the surrounding retina [92].

2.5.1 TREATMENT PROCEDURE

Removing the epiretinal membrane or ILM requires a technique known as membrane peeling (shown in Fig. 10). A microvitreoretinal (MVR) or other blade is used to puncture the membrane. A pick or ILM elevator is then inserted just under the membrane to lift it up. This is very challenging work and sometimes takes minutes' worth of attempts [86]. Once the membrane has been lifted with the stripping motion, the pick is advanced. The surgeon tunnels between the membrane and the retina, periodically lifting the membrane off the retina. Forceps can be used to remove the peeled membrane. This technique is a technically challenging procedure, and unsuccessful peeling can result in poor visual outcome [106]. One study showed that up to 50% of patients exhibited inadvertent injury and defects in the nerve fiber layer after performing ILM [49].

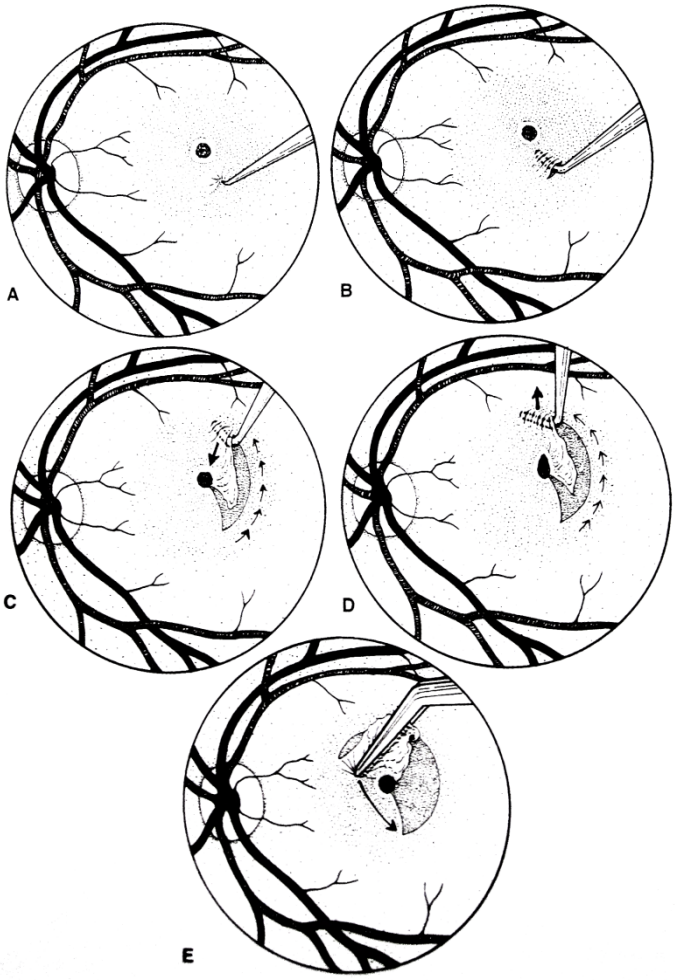

Fig. 10. Internal Limiting Membrane Peeling (ILM) around a macular hole. Photo credit [86].

2.5.2 BEHAVIORAL AIDS

- **Downward Motion Limits**: Preventing fast or extraneous motion in the downward direction could avoid accidental trauma to the retina.
- **Motion Scaling:** Increasing precision in the Z motion could reduce the number of tries it takes surgeons to insert the pick between the ILM and the retina.
- **Follow Curved Path:** Once the pick has been inserted, help with the fine tunneling motions (advance in a curve, lift up slightly) could also reduce accidental injury to the retina.

2.6 ARTERIOVENOUS SHEATHOTOMY

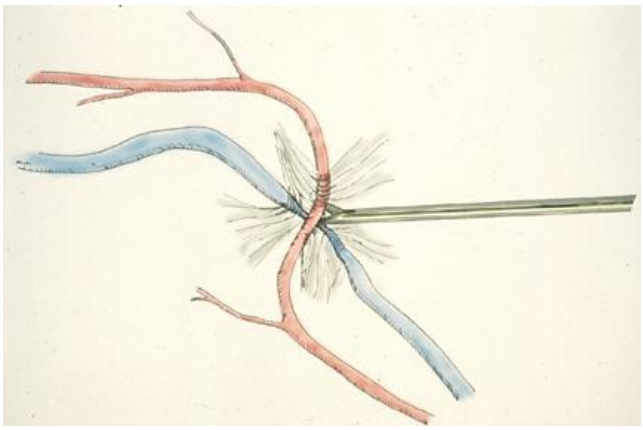

Fig. 11. Demonstration of performing an arteriovenous sheathotomy by separating the sheath that holds the vein and artery together. Photo credit [45].

One hypothesis in vessel occlusion is the adventitial sheath that binds arteries and veins together at ateriovenous crossings restricts blood flow, causing turbulence, which may lead to thrombosis. This is supported by the fact that the arteriole crosses over the vein between 70-99% of the time when BRVO occurs at arteriovenous crossings [114, 120]. It is thought that by removing the sheath, the vein will decompress, allowing for vessel recannulation. Osterloh and Charles first succeeded in showing surgical decompression via arteriovenous sheathonomy [85]; as a result, researchers have corroborated a variety of similar procedures [62, 84]. However, other experiments have reported failure to surgically separate the vessels because of the difficulty of the operation [15, 37, 38].

2.6.1 Treatment Procedure

As depicted in Fig. 11, an arteriovenous sheathotomy procedure aims to separate the artery from the underlying retinal vein by severing the sheath that connects the two. A bent microvitreoretinal (MVR) blade or scissors are used gently cut the sheath and pry the vein and artery apart from each other. Complications include nerve fiber layer defects, hemorrhage, and retinal detachment. Interesting, some research has indicated that simply freeing the vessels from the retina might be enough to restore blood flow [38].

2.6.2 Behavioral Aids

- **Tremor Compensation:** Because of the extremely delicate nature of the operation and the many reported failures and hemorrhages, reducing tremor is essential.
- **Avoidance Zones:** Setting limits on tip movements too far into the retina could prevent tearing or scratching the retina.
- **Motion Scaling:** Reducing movements around the critical junction of the vein and artery would increase surgeon precision in manipulation.
- **Velocity Limiting:** Preventing fast tip movements could help avoid accidental hemorrhage of the blood vessels.

2.7 Discussion

We have analyzed four vitreoretinal surgeries, their treatment procedures, and how a handheld micromanipulator could be beneficial during the operation. While Micron is applicable to a wide range of micromanipulation tasks, we chose these surgeries because of the precision required to perform the operation, making them ideal candidates for robotic aid. Because of the prevalence of laser photocoagulation and the skill required by cannulation, we focus on these two surgical techniques for the experimental sections of this thesis.

3 PRELIMINARY WORK

In this section, we document preliminary work in previous experimental approaches. In particular, we focus on investigating the feasibility of vision-controlled micromanipulation with two specific procedures: laser photocoagulation and vessel cannulation. Work on retinal laser photocoagulation demonstrates the ability to rapidly orient and fire at anatomical targets tracked via vision; while the retinal vessel cannulation research demonstrates how vision can be used modify motion scaling behavior in real-time. Experimental results from this preliminary work validates that a handheld micromanipulator can increase performance of the procedure.

3.1 RETINAL LASER PHOTOCOAGULATION

A recently introduced laser photocoagulation approach by Blumenkranz et al. utilizes a systems of mirrors mounted on a two axis galvanometric scanner attached to a slit lamp which deflects the laser beam in order to apply patterns of up to 50 pulses rapidly at a single command from the foot pedal [18]. Patterns demonstrated by this system include circular arcs, lines, and rectangular and circular grids. The semi-automatic application of laser spots has the potential for significant reduction of treatment time, although the system could not always avoid previous laser burns when applying a new pattern [101]. This approach has been commercialized as the Pascal Photocoagulator [81, 101]. However, it is designed for office use rather than for the operating room.

With the goal to extend the benefits of semi-automated systems to the realm of intraocular surgical application, we present the initial phases in the development of an assisted intraocular laser system that will increase the speed and accuracy of the placement of laser burns. It avoids accidental burns in dangerous areas such as the macula and optic nerve and blocks firing of the laser when the distance between the target tissue and the laser is too small or too great. The robotic platform of this system for laser photocoagulation is Micron (see for Fig. 12 prototype), a fully handheld active micromanipulator that has been reported previously [13, 94]. Micron uses frequency-multiplexed optical tracking to sense its own motion in six degrees of freedom (6DOF) and control a 3DOF parallel manipulator built into the handle [76]. By deflecting the tool tip or end-effector in real-time, Micron can compensate for undesired motion such as physiological hand tremor [95], or actively guide the tip toward a known target [13], such as a desired laser burn location.

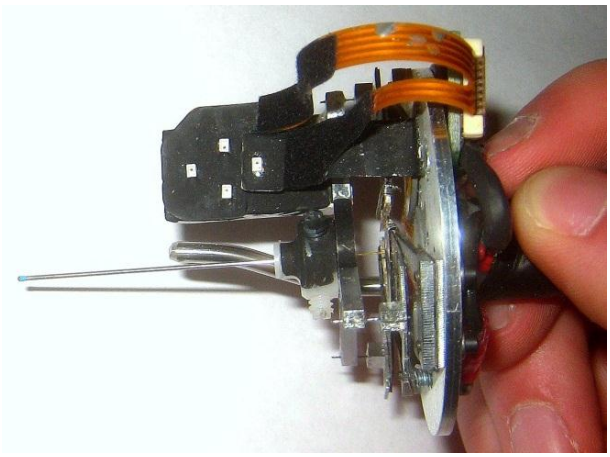

Fig. 12. Close-up view of Micron in the hand of the user. The fiber-optic cable threads through the hollow center of the instrument, running down the actuated metal shaft to the tip. Four white LEDs act as optical trackers to provide micron-accurate tracking of the tip of Micron.

3.1.1 SYSTEM SETUP

An Iridex® Iriderm Diolite 532 Laser is attached to Micron, and the optical fiber from an Iridex® 20-gauge standard straight EndoProbe handpiece runs through the shaft of the instrument to the end-effector. The setup is shown in Fig. 13. When fired, the laser optic setup yields a laser spot size that is 200-400 μm in diameter. Planning software to preoperatively

specify the location of each burn was written in LabVIEW® (National Instruments Corp., Austin, Texas, USA). A retinal image captured through the microscope at any time before the start of photocoagulation is used as the background upon which the desired pattern can be planned. The operator selects predefined patterns such as rectangular grids, circles, arcs, and ovals to place on the retinal image. A complete set of burn locations can be specified by moving, scaling, and stretching these primitive shapes. The whole pattern of target burn locations consisting of individual targets can be loaded into the control software that runs during the operation. For the remainder of this paper, it is assumed that all 2D positions in the image are measured in units of pixels and are defined in homogenous coordinates.

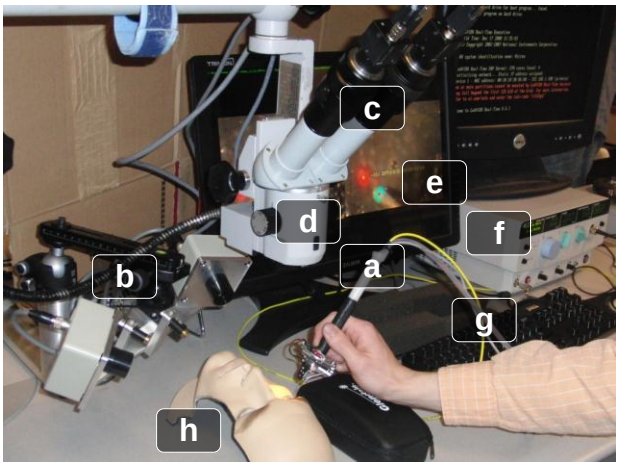

Fig. 13. The setup for laser photocoagulation with Micron, an active handheld micromanipulator. Micron (a) is shown held in the hand of the operator. The position of Micron is measured optically in 3D by the two ASAP PSDs (b) and stereoscopically in 2D via two cameras (c) mounted to the microscope (d). Real-time video with overlaid information is displayed on the 3D computer monitor (e). The laser (f) connects to Micron (a) via fiber optic (g). Targets are placed under the hollowed out eye of the face model (h).

3.1.2 Tracking of Burn Locations and Tip Positions

Registration of the preoperative image to the real-time video aligns the burn locations to the current frame, compensating for any movement of the eye. A number of possible algorithms can be used [48, 77, 121], but a straightforward feature-based approach of aligning interest points between the current frame of the video and the preoperative image works well for feature-rich images. Interest points are detected using Speeded Up Robust Features (SURF) [10] and aligned with a planar homography that is estimated using the standard computer-vision RANSAC algorithm [40]. Thus, assuming an approximately planar scene, registration can be maintained throughout the procedure, keeping burn locations consistent with their initial placement on the preoperative image.

In addition to registering the target burn locations, the vision system also tracks the 2D position of the color-coded tip of Micron in real time in the left and right video views, respectively. A color-based algorithm [118] tracks the centroid of a blue patch painted on the end of the tip of Micron. A red finder beam generated by the laser while in treatment mode provides an aiming guide to the operator between firings of the laser. The aiming beam is important because it shows the location of the burn before the laser is activated, allowing for continuous targeting. The vision system also tracks the centroid of the aiming beam in both left and right views.

Finally, the vision system is responsible for removing noisy images. During the execution of the procedure, video frames captured while the laser is firing are automatically detected and removed to avoid low contrast images and blooming effects caused by the intensity of the laser. The popular Intel® OpenCV library is used to implement the computer vision techniques.

3.1.3 Coordinate System Calibration

Only the vision system can sense the targets, the aiming beam, and the relative error between them. However, control occurs in the 3D space of ASAP by setting a 3D goal position and using a PID controller to reach that goal. This separation of sensing and control necessitates transformations between the image and the ASAP coordinate systems. This is possible by measuring the only variable that is observable in both coordinate systems and using it to calculate the transformations. The tip position of Micron is measured in three places: the left image p_L, the right image p_R, and in 3D space P. Measurements of p_L, and p_R are obtained from color trackers while P is measured by the PSD optical trackers in ASAP sensing the pulsed LEDs on the shaft of Micron.

The coordinate system transformation takes the form $p_c = M_c P \;\; \forall\, c \in (L, R)$, where P is the 3D homogenous coordinate of a point in space, p_c is the 2D homogenous coordinate of P imaged in camera $c \in (L, R)$, and M_c is a 3×4 camera projective matrix. A preoperative 30 s calibration is used to measure the 3D tip position P and the corresponding 2D tip positions p_L and p_R, which are used to estimate the transformations M_L and M_R using the Direct Linear Transform [50].

Because this estimation is performed with precision in the micrometer range, even very slight shifts or movements in the camera or PSD arrangement can introduce significant errors into the transformations, as can nonlinearities of the optical tracking sensors caused by axial rotations of the tool. To rectify these errors, Micron employs an online recursive least squares approach that minimizes the error of the transformation estimate for each frame during the entire operation, refining and adapting the coordinate system transformations over time for camera c:

$$M_c \leftarrow M_c + \eta(p_c - M_c P)P^T \tag{1}$$

where $(p_c - M_c P)P^T$ is the error of the transformation and η governs how closely the online calibration algorithm adapts the transformation to the measured data. This adaptive calibration is important for maintaining an accurate transformation between coordinate systems throughout the procedure.

Using these transformations, we define two mappings between the coordinate spaces: image projection $\phi(P, M_c) \rightarrow p_c$ and stereo triangulation $\Phi(p_L, p_R, M_L, M_R) \rightarrow P$. The first mapping ϕ is a projection that maps a point P in 3D space, defined by the coordinate system of ASAP, to a 2D point p_c on an image seen by camera c. The second mapping Φ involves solving for the intersection of two rays in space, and triangulates a pair of 2D points p_L and p_R, seen by the left and right cameras, to find the most likely 3D point P that corresponds to the observations $p_L = \phi(P, M_L)$ and $p_R = \phi(P, M_R)$.

3.1.4 Initial Control Mechanism

One possible approach [13] measures the error directly from the images as the relative distance between the finder beam (where the laser is currently aiming) and the closest aligned target i (where the laser should be aiming), transforming the error into the 3D ASAP coordinate system:

$$e = \Phi(l_L, l_R, M_L, M_R) - \Phi(a_L, a_R, M_L, M_R) \tag{2}$$

Where l_L, l_R and a_L, a_R are the burn target and laser aiming beam tracked locations in the left and right images, respectively. M_L, M_R are the camera calibrated project matricies that enable the

mapping Φ to project 2D positions in the left and right images to 3D points in ASAP space. A control signal can be directly generated by using the error as an input to a PID controller. Since both the aiming beam and target are on a locally planar surface, this control signal drives the aiming beam to the target without affecting the distance of the instrument relative to the surface. Once the target location has been acquired by the aiming beam, the laser can be fired. While yielding very accurate results, control signals can only be updated at the camera capture rate of 30 Hz, resulting in slow convergence and laser burn rates less than 0.25 Hz.

Instead of waiting for a new frame from the camera to directly measure the error between the target and the aiming beam, an alternate method is to reconstruct the model of the targets in the 3D ASAP coordinate system and use a model of Micron to update control signals between camera frames. Since the 3D tip position and pose is measured by ASAP at 2 kHz, control signals to point the tip toward the target can be generated much faster, leading to quicker convergence rates and a higher overall burn rate. Since the tip of Micron generally moves more rapidly than the retinal surface, the 3D model of burn targets can be updated at 30 Hz and still remain largely valid in between camera frames. However, as this approach depends on good models of the targets and Micron, accuracy of burns is expected to be lower than the approach of [13].

3.1.5 Surface Reconstruction

Developing a model of the targets can be accomplished by reconstructing the surface of the retina in 3D. Noting that that burn locations are applied on the surface of the retina, the 3D surface can be reconstructed using dense stereo algorithms operating on the left and right images with each new set of frames [102]. However, assuming a non-deformable surface, a more relevant approach is to use structured light [9] since Micron is equipped with a red aiming beam laser. Because of the simplicity of the retinal surface being reconstructed (i.e., no discontinuities, no large spectral reflections, no opacity, no occlusions), complex coding of the structured light is unnecessary. By sweeping Micron back and forth above the surface and observing the intersection of the aiming beam with the surface in both images, the 3D structure of the surface S can be calculated by transforming each of the aiming beam points a_L and a_R to a 3D point P_s such that each point P_s belongs to surface S:

$$P_s = \Phi(a_L, a_R, M_L, M_R) \mid P_s \in S \tag{3}$$

While one could use any model for the surface S (plane, quadratic, splines, etc), our system uses a planar representation by fitting the observed 3D points $P_s \in S$ to a plane using a least squares algorithm. The planar assumption works for our experiments in vitro; however, a higher-order model such as a quadratic surface fit might be more suited for testing in vivo, in which the retina adheres to the curved shape of the eye. Only an initial 5-15 s calibration routine is necessary to collect enough correspondences from the structured light to reconstruct the planar surface. If desired, further refinement at time t of the surface S can be calculated iteratively during the procedure to yield a time-varying surface S^t. Once the surface has been reconstructed, the burn locations from the left and right images L_L and L_R are projected onto the reconstructed surface S to get the 3D burn locations:

$$l_{3D} = proj(\Phi(l_L, l_R, M_L, M_R), S) \in L_{3D} \quad \forall\, l_L \in L_L, l_R \in L_R \tag{4}$$

where, in the planar case, the projection maps the 3D burn location to the plane. The purpose of this projection is to reduce noise in 3D target calculations, since the Z-component of the reconstructed point is subject to the most uncertainty.

3.1.6 Control System

Mimicking the standard surgical procedure of burning target locations in sequence, the controller selects the nearest target within range, deflects the tip to aim at it, fires the laser while locking the aiming solution, and only moves on to the next nearest target location after the current burn is completed. If the next nearest target is not within range or all preoperatively specified burns have been applied, the Micron tip goes to its neutral position. As the operator executes a fly-over maneuver with the instrument, Micron can quickly flick out as it goes by to burn passing targets. With this semiautomated method, the operator performs the gross motion by sweeping Micron over clusters of targets and lets the control system handle the exact positioning of the tip and timing of the laser activation.

The controller uses the tip position P, instrument rotation R, and targets L_{3D}, all of which are known in the 3D ASAP space. The tip position and instrument rotation are sensed by ASAP at 2 kHz, while the 3D reconstructed targets are updated from the cameras at 30 Hz. This disparity in update rates is acceptable because the targets move slowly, if at all, during the procedure.

Selection of the current target is done by choosing the yet-untreated target that requires the least amount of movement from Micron's neutral position. The neutral position is defined by all actuators being at the zero position, thereby allowing for the greatest movement in any arbitrary direction; functionally, this state is equivalent to Micron in the off state. Since the arrangement of actuators defines a pivot point near the base of the handle, conceptually this optimization of least movement finds the target location that requires the minimum rotation of the tool about the pivot.

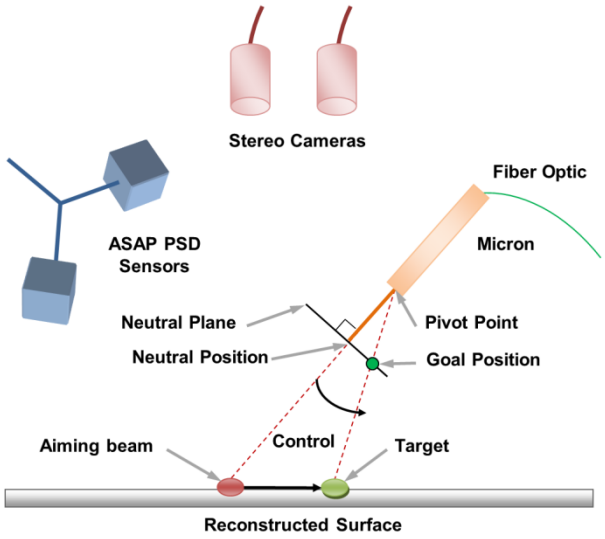

Fig. 14. Micron laser setup. The Micron handle attaches to the laser via a fiber optic cable. Cameras track the finder beam and target, while ASAP optically tracks the tip and pose of Micron. The goal location is determined by intersecting the ray connecting the pivot point and the target with the neutral plane. The piezoelectric actuators then deflect the tip from the neutral position to the goal position and fire the laser to burn the target.

Micron uses a PID controller to reach specified 3D goal positions, so a goal position that causes the aiming beam to hit the target must be calculated. This calculation forms the core of the control system and is highly dependent on accurate estimated transformations and target models. While any goal position on the line connecting the pivot position with the target would result in the laser aiming at the target, but the best 3D goal position for Micron is the one that maximizes the available range of motion to

account for tremor and the gross sweeping movements of the operator.

Since the workspace of Micron is shaped like a disc centered at the neutral position with its thickness tapering off towards the edges, the maximum transverse motion can be achieved on the plane normal to the shaft of the instrument and intersecting the neutral position. Because the axial range of motion is small and actuation in this direction greatly reduces available transverse movement, the controller leaves the job of depth stabilization to the operator. With this maximal maneuverability constraint, goal positions are restricted to this neutral plane N_P, yielding a single goal position P_G for any combination of tip position P, instrument rotation R, and target location l_{3D}. The 3D goal P_G is calculated at 1 kHz by intersecting the ray connecting the pivot point and the selected 3D target with the neutral plane:

$$P_G = intersect(N_P, l_{3D} - P_V) \tag{5}$$

The pivot point is defined by the tip position offset by the length of the shaft D and rotated by the pose R of the instrument in space: $P_V = P - R^{-1}D$. Using the PID controller to directly reach the goal point P_G, the system exhibits an underdamped behavior that has a tendency to overshoot the goal. Instead, a 30-ms minimum jerk trajectory is planned to the goal position, which results in much better tracking and significantly less overshoot. Once the goal position is reached, the laser is activated. The goal position is adjusted at 1 kHz to eliminate tremor or other motion until the laser has finished firing, at which point the tip returns to the neutral position and another untreated target is selected. The procedure terminates when all targets have been treated. See Fig. 14 for a graphical representation of the control system.

Small errors in the model of Micron or the calibration can cause large errors during the 3D reconstruction, so an additional 5-15 s calibration procedure is executed after the surface reconstruction. Errors in the calculated intersection of the aiming beam with the reconstructed surface and the observed location of the aiming beam are recorded during this calibration. The control system then adjusts the surface position with the measured mean error to better align the reconstructed surface with the observations. This is usually then valid for the remainder of the procedure, provided the operator does not significantly rotate Micron.

Several safety measures are designed to prevent firing errors and misplaced burns. The first safety check limits the maximum distance the Micron tip can move when deflecting to aim at a target. A conservative threshold is used to allow enough range of motion to accommodate the user's tremor or gross hand movements during targeting and burning. The second safety check allows activation of the laser only after targeting is complete and if enough reserve manipulator range is available. In the case where Micron is close to the limits of its movement, it aborts and returns to the neutral position to avoid being unable to continue targeting the burn location while the laser is activated. The third check ensures proper Z-distance between the tip and target before firing the laser, because if the laser is too close or too far away, the tissue will burn unevenly or not at all. If at any point Micron decides the target is unreachable, the tip is returned to the neutral position and Micron begins the target selection process again.

3.1.7 EXPERIMENTAL PROCEDURE

For ease of initial development and testing, experiments were performed on pictures of a retina printed on a color, high-resolution printer. The yellow/orange tones of the ink on the paper served as a good absorptive material for the laser, yielding distinct, solid black burns. The paper slide was placed under a formed face mask with the eyes hollowed out to give a more realistic operating environment. The Micron shaft is then inserted through the eye slot during the procedure to place laser burns on the paper. The surgeon's hand rested on a block next to the face mask; Micron was not braced against anything.

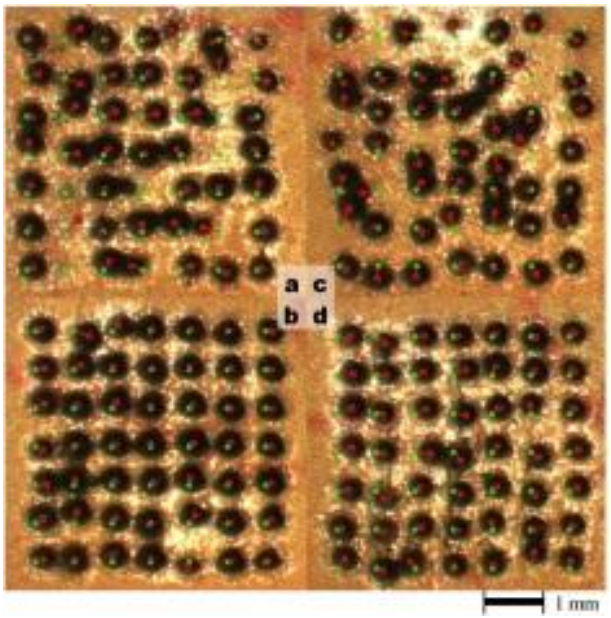

Fig. 15. Results from paper slide trials. Green solid dots indicated preoperatively specified burn locations, and red hollow circles indicate the centroid of the burn. (a) Top left: 0.5 Hz unaided. (b) Bottom left: corresponding aided trial from 0.5 Hz trial. (c) Top right: 1.0 Hz unaided. (d) Bottom right: corresponding aided from 1.0 Hz trial.

A 7x7 grid with the preoperative dot locations placed approximately 650 μm apart was selected as a good test pattern. Two different cases were tested during the experiment: unaided and aided. During the unaided case, the laser fires at a fixed frequency when the pedal is depressed. In this scenario, Micron is switched off, and it is the operator's responsibility to move the tip of the instrument to accurately place the burns in time with the laser repeat rate.

During the aided case, Micron is turned on and actively helps the user. The operator is responsible only for the gross movement, while Micron handles the precise targeting and firing of the laser. As with the unaided case, for safety reasons, the laser fires only when the pedal is depressed; however, the laser firing mechanism is under software control instead of being pulsed at a fixed frequency. This allows Micron to operate as fast or as slow as the operator feels comfortable moving the instrument.

During the paper trials, the individual laser pulses were set for a duration of 45 ms and a power of 3.0 W. The laser repeat rate for the unaided case was set to 0.5 Hz in the first set and 1.0 Hz in the second set of experiments. During the 0.5 Hz set, the surgeon was instructed to go slowly and steadily during the aided case. In the 1.0 Hz set, the surgeon was asked to proceed as fast as he felt comfortable during the aided task.

Tests with porcine retina in vitro were performed by the surgeon in a similar fashion to the paper slide procedure. Excised pig eyes were refrigerated until used and dissected immediately before the experiment. The retina was lifted out, placed on black felt, and smoothed out to form an even thickness. The prepared retina was mounted under the microscope (without the face model for simplicity). Preparation and mounting were done immediately before experimentation, to avoid drying of the tissue. The same 7x7 grid with 650 μm spacing was used with both unaided and aided scenarios. Laser pulses were 2.2 W for 27 ms, causing 200-300 μm diameter burns that appear on the retina as milky white spots. Because the surgeon was more familiar with the procedure by this

time and the aided cases were averaging near 2.0 Hz for the paper slide trials, we fixed the laser repeat rate at 2.0 Hz for a more equitable mean error comparison during the porcine experiments.

3.1.8 Effects of Visualization and Tool Ergonomics

Two potentially significant differences exist between the setup used to perform the paper slide and porcine tissue experiments and the setup a surgeon typically uses: the visualization of the workspace and the ergonomics of the tool. We devised two experiments to test each of these factors to determine the potential impact on accuracy of burn placement.

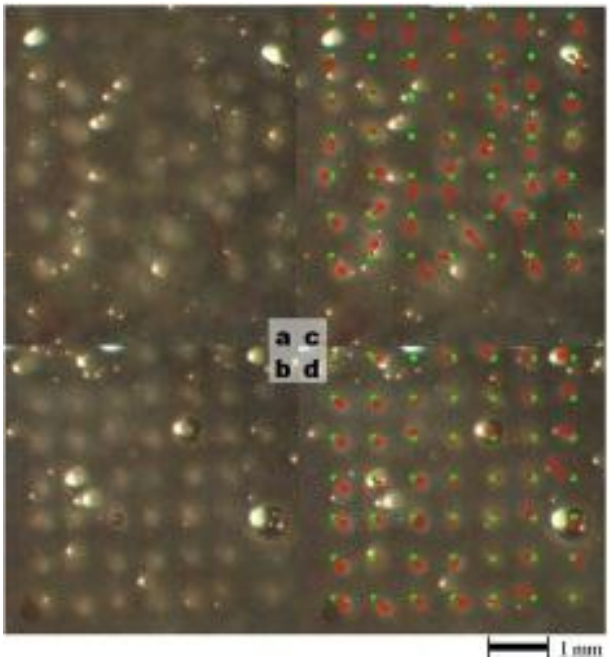

Fig. 16. Results from porcine retina trials. Green solid dots indicated preoperatively specified burn locations, and red hollow circles indicate the centroid of the burn. (a) Top left: 2.0 Hz unaided. (b) Bottom left: corresponding aided trial from 2.0 Hz trial. (c) Top right: 2.0 Hz unaided with errors depicted. (d) Bottom right: corresponding aided trial with errors depicted.

While the computer monitor can display informative overlays, such as the targets in the video stream, the quality is not as high as viewing the workspace directly with the microscope. Specifically, the cameras capture images with a lower dynamic range and at lower resolution than the human eye. Additionally, a 60-ms lag is introduced in order to capture, process, and display the real-time video. Of particular interest is how these factors affect performance. To investigate this question, we executed an additional test similar to those mentioned earlier with a 7x7 grid, except that the grid targets were printed directly onto the paper slides and the surgeon used the normal microscope eyepieces instead of the cameras and video. Only the unaided trials with an inert Micron could be run under such conditions. The surgeon executed one trial at the 0.5 Hz repeat rate, and one at the 1.0 Hz repeat rate.

The other potential factor we investigated was the ergonomics of the tool. The IRIDEX EndoProbe standard straight 20-gauge handpiece typically used in practice is significantly lighter and thinner than Micron. To determine if the ergonomics of Micron were significantly impacting performance, we ran the same two experiments again at the 0.5 and 1.0 Hz repeat rates with the IRIDEX EndoProbe instead of an inert Micron.

3.1.9 Experimental Protocol

Experiments were performed by a vitreoretinal surgeon with more than 20 years of experience. For each set of experiments, eight trials were performed sequentially by the surgeon during a one-hour period, alternating between unaided and aided. Roughly one minute of rest was provided between the interleaved unaided and aided trials. Video of each trial was recorded, along with the target locations on the preoperative image. From this information, the mean error and duration for each trial were measured. Error was measured by taking the final frame of the video sequence and hand-

marking the centroid of each burn. Individual error was calculated as the distance between the centroid of the burn and the target. Nearest-neighbor matching was used to match burns to targets. To avoid spurious matches, errors greater than the spacing between the target burn locations were discarded and noted as missed targets. Thus mean error was calculated only on the subset of targets where the burn location was in the neighborhood of the target. Statistical significance was determined by calculating p-values assuming a two-tailed test.

3.1.10 RESULTS

Fig. 15 and Fig. 16 present sample trials of the aided and unaided cases for both paper slide and porcine retina experiments. Mean error and mean duration for all the sets of trials are shown in Fig. 17 and Fig. 18. P-values for the mean error of all trials were less than 0.005 for all three sets of experiments. P-values for the trial durations were less than 0.005 only for the 0.5 and 1.0 Hz sets of experiments. Table 2 lists the mean errors for the three different sets of experiments, along with the overall reduction in mean error between the unaided and aided cases.

Trial/error (Hz)	Unaied (μm)	Aided (μm)	Reduction (%)
0.5	149	125	16.1
1.0	166	129	22.3
2.0	203	123	39.4

Table 2: Reduction in Overall Mean Error

In the investigation of the effects of different visualization systems, the mean error with the microscope view, when compared with the cameras and computer display view, was 58% lower at the 0.5 Hz repeat rate and 60% lower at the 1.0 Hz repeat rate. When running the same experiments with the normal probe instead of the inert Micron, the surgeon achieved error reductions of 72% and 68% in the 0.5 and 1.0 Hz cases, respectively.

3.1.11 DISCUSSION

The results presented demonstrate the general feasibility of an active handheld surgical device to decrease spot placement error and procedure duration in intraocular laser retinal photocoagulation. In both the paper slide and porcine retina trials, Micron significantly reduced the positioning error. The mean error for the unaided cases shows an interesting, but expected trend. At lower repeat rates, the surgeon has enough time to move and accurately place laser burns. At higher laser repeat rates, the mean error increases as the surgeon's ability to precisely control the placement of burns is exceeded. The position error for the aided cases is generally consistent across experiments, indicating that Micron can handle precision targeting even with fast gross movements of the instrument. Mean duration for the procedure was consistently shorter for the aided cases.

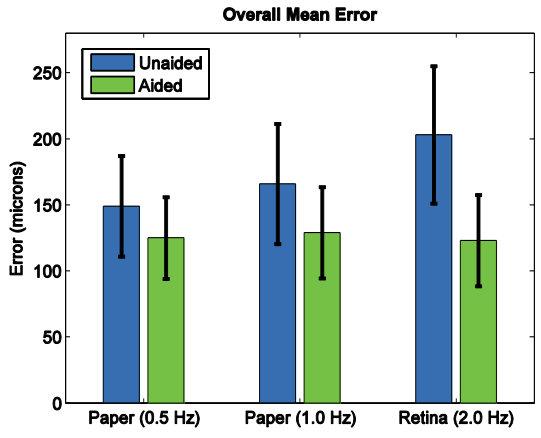

Fig. 17. Mean error bar chart with error bars representing the standard deviation for all three trials with laser repeat set to 0.5, 1.0, 2.0 Hz from left to right. Each scenario consisted of four unaided trials and four aided trials interleaved.

Excepting the 0.5 Hz case, in which the surgeon was instructed to execute the procedure slowly and methodically, mean duration was similar for all aided cases. Although the present demonstration is limited to simplified conditions in vitro, the experimental results show that a surgeon using the current prototype of Micron can perform the photocoagulation procedure with an accuracy improvement of 15% or more, compared to the control case, and up to four times faster, depending on the laser repeat rate selected.

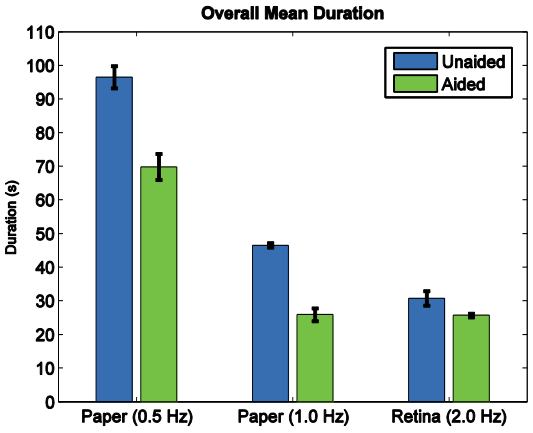

Fig. 18. Mean duration bar chart with error bars representing the standard deviation for all three trials with laser repeat set to 0.5, 1.0, 2.0 Hz from left to right. Each set consisted of four unaided trials and four aided trials interleaved. The apparent discrepancy in the 0.5 Hz case is because the surgeon was instructed to take his time during the aided trials.

Furthermore, this approach has three attractive side effects. First, because targets are specified a priori and Micron incorporates safety features to prevent firing at non-target locations, accidental extraneous movements do not risk laser burns being applied inadvertently to dangerous areas, which is risk present if the laser is in repeat mode and the surgeon experiences unintentional movement. Second, the surgeon is relieved of explicitly targeting burn locations and avoiding critical areas such as the fovea or vasculature. Using our system, the surgeon is responsible for only the gross movements of the instrument while the robotic assistance performs the precise targeting and firing behaviors. This level of automation reduces the surgeon's cognitive workload, which is generally thought of as a worthwhile goal [24]. Third, Micron's pose information combined with the 3D reconstruction of the surface enables a safety mechanism that only fires the laser when the distance from tip to retina is within a pre-specified range. This prevents ineffectual burns when the tip is too far away from the surface and "too hot" laser burns when the probe moves too close to the retina for the specified duration and power of the laser. Avoiding burns that are "too hot" has the desirable effect of lowering the risk of unintended choroidal and retinal neovascularization as a late complication of the laser surgery [100].

The present study is a proof of concept; substantial further work is necessary in order to bring the technique to clinical feasibility. The comparisons dealing with visualization and instrument ergonomics show that both have an effect on performance. Viewing the workspace via cameras and computer monitor led to degradation in performance compared to viewing directly through the stereo operating microscope. How much of this degradation could be eliminated by training is not known. In any case, a system that injects graphics directly into the optics of the microscope would likely alleviate the problem. As for ergonomic factors, the 11% decrease in unaided accuracy with Micron when compared with the IRIDEX EndoProbe confirms that the size and weight of the existing prototype, although small enough to be usable, nonetheless measurably degrade performance. To minimize this effect, a smaller and lighter instrument is presently under development. Future research must also include techniques for proper control in the presence of the fulcrum imposed by the entry point at the sclerotomy, and suitable calibration for accurate tracking and control of Micron when viewed through the cornea and lens of the eye. Because pre-

operatively selecting burn locations can be a tedious procedure, a semi-automatic method that can suggest burn locations from learned examples of previous surgeons' patterns on the retina could further expedite the operation. For extended procedures, the capability to re-image and re-register views of the eye to plan or re-plan new burn patterns intraoperatively also could be added and may be beneficial. Future evaluation in vivo in animal models will undoubtedly clarify the need for additional refinements.

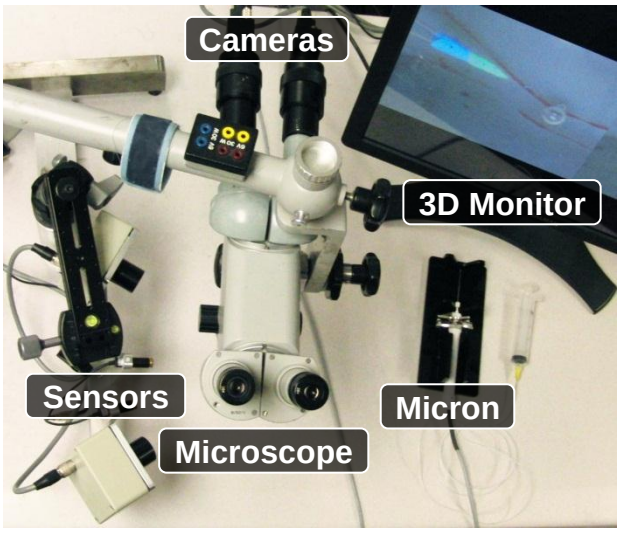

Fig. 19. Cannulation setup with Micron micromanipulator, microscope, stereo cameras, and PSD optical sensors

3.2 Retinal Vessel Cannulation

Although successful cannulations in a variety of animal and human models have been reported, there are difficulties involved in manipulating such tiny vessels. These reports are not surprising considering that physiological tremor of vitreoretinal surgeons has been measured at over 100 μm in amplitude [105], which is greater than the typical 40-120 μm diameter of retinal vessels. Recognizing the need for reduced tremor and increased precision, [41] describes cooperative control between a surgeon and a Steady-Hand robot for *in vivo* cannulation of 80 μm vessels in the chorioallantoic membrane of chicken embryos. Though the approach did not improve the cannulation success rate, the experiment did show improved ability to maintain the cannula in the vessel. We propose a vessel cannulation system based on a handheld robot, "Micron" [74], using vision-based control. Preliminary results *ex vivo*, performed by an experienced retinal surgeon, validate the system.

3.2.1 System Setup

As shown in Fig. 19 and Fig. 20, the surgeon uses Micron in a 2x3 mm workspace under a Zeiss OPMI® 1 microscope at 25X magnification. An optical bridge splits the microscope view between the microscope eyepieces and two Flea®2 cameras (Point Grey Research, Richmond, BC, Canada). Mounted as a stereo pair, the cameras capture 800x600 resolution video at 30 Hz. For cannulation procedures, TIP1TW1 micropipettes (World Precision Instruments, Saratoga, FL, US) with an inner diameter (ID) of 1 μm are secured to the output plate of Micron. Tubing threads through the hollow center of Micron's shaft and connects to a 20 mL syringe, which can be depressed to force gas or fluid through the micropipette and into the vessel.

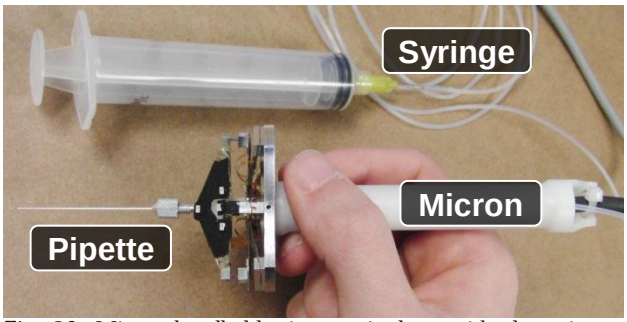

Fig. 20. Micron handheld micromanipulator with glass pipette attached with tubing to a 20 mL syringe.

3.2.2 Vision System

The stereo cameras are responsible for tracking both the tip of the micropipette and the vessels in the image to ascertain the proximity of the micropipette to the vessel as depicted in Fig. 21.

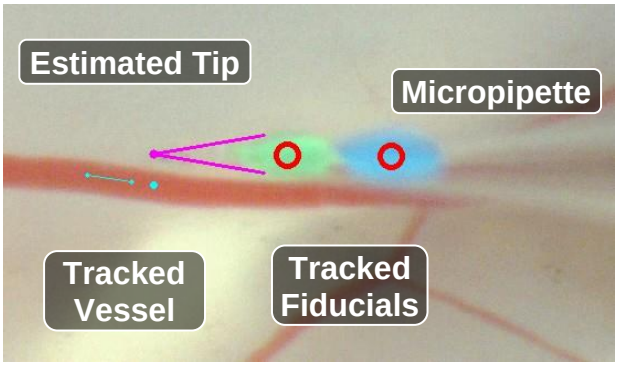

Fig. 21. Stereo cameras track 3D positions of green and blue fiducials on the micropipette to estimate the tip position. The vision system also detects vessels in the image to calculate the center of the vessel nearest the micropipette tip.

The relative distance is used by the control system, which adjusts the scale factor that determines how much hand motion is transferred from the handle to the tip of the instrument.

Since the micropipettes are formed from clear glass and the tip tapers off to a nearly invisible point, directly tracking the tip with vision poses a significant challenge. As an effective alternative, two unique colored fiducials are painted in stripes around the micropipette near the tip. Fast color segmentation [22] tracks the centroids of the fiducials, and the true micropipette tip position is estimated with interpolation. A quick 30 s calibration period registers the camera views with the world coordinate system as defined by the PSDs measuring the 6-DOF pose of the instrument; registration is maintained by an online adaptive least squares calibration procedure that adjusts the camera coordinate system to the world frame with each new camera image.

Vessels are similarly tracked with a Gaussian Mixture Model of learned vessel colors [118]. By searching locally around the estimated position of the micropipette tip, the XY center of the nearest vessel is located. Because vessels appear as locally invariant structures, i.e., one section of the vessel does not look significantly different from another section, correspondences for 3D reconstruction of the vessel are difficult to find. Instead, a global estimate of the mean depth Z of all blood vessels in view is calculated by running a simplified Lucas-Kanade tracker [6], which estimates the translation between the left and right camera views of the blood vessels. The 3D distance between the micropipette tip and the center of the nearest vessel as calculated by the stereo cameras is then given to the control system.

3.2.3 Tremor Compensation and Motion Scaling

Two key components form basis of the control system: tremor compensation and motion scaling. Hand tremor is compensated by commanding the micromanipulator tip $P_T \in \mathbb{R}^3$ to the the null position $P_N \in \mathbb{R}^3$ that has first been filtered by a lowpass filter $\boldsymbol{F_L}$. Since the null position is mechanically tied to the handle, the micropipette tip output is the smoothed input of the operator's hand movements:

$$P_T = \boldsymbol{F_L}(P_N) \tag{6}$$

Since physiological hand tremor generally lies in a frequency band of 8-12 Hz, a first-order 1.5 Hz lowpass Butterworth filter was selected for $\boldsymbol{F_L}$, which greatly attenuates tremor while only introducing a minor lag and settling time into the eye-hand coordination feedback loop.

The second key component used by the control system is motion scaling, which gives the surgeon increased precision. Motion scaling takes the handle motion as calculated by the change in null position ΔP_N and transfers only a fraction $1/s \in \mathbb{R}$ of the movement to the micromanipulator tip ΔP_T:

$$\Delta P_T = \frac{1}{s}\Delta P_N \tag{7}$$

For example, if the scale factor is 2 and the surgeon moves the handle 100 μm, the tip only moves 50 μm. Any uncompensated tremor, drift, or deliberate movement is scaled down, granting the surgeon the additional comfort and safety of using larger movements to manipulate tiny anatomy. A scale factor of unity is equivalent to running no motion scaling (i.e., Micron turned off) while a scale factor of ∞ corresponds to fixing the tip in space (i.e., disregarding all operator's movement).

One important limitation of motion scaling is the decreased range of motion. During tremor compensation, the lowpass filter gain rolls off to unity at 0 Hz so the manipulator range of motion is only used temporarily to smooth out sudden jerks and/or tremor. However, since motion scaling is applied all the way down to 0 Hz, the effective range of motion of the device is reduced by a factor s. For this reason, tremor compensation can be performed everywhere, whereas motion scaling can only be applied in a limited volume, generally near the vessel where it is most effective in aiding the cannulation procedure.

Additionally, because the 0.8 mm axial range of Micron is significantly smaller than the 3.0 mm transverse range, motion scaling is applied anisotropically. While tremor compensation is applied in all directions at all times, motion scaling is only activated as discussed earlier in the vicinity of the vessel and only in the transverse directions of the tool. In addition to preserving manipulator range along the axis of the tool where the range of motion is most limited, anisotropic scaling has an additional benefit. Since the surgeon approaches the vessel along a roughly parallel trajectory, unity motion scaling axially allows for quick thrusts into the vessel while motion scaling in the transverse direction increases the ability of the surgeon to keep the micropipette centered on the vessel.

3.2.4 Control System

The control system combines tremor compensation with anisotropic motion scaling to aid the surgeon in cannulation procedures. Tremor compensation is applied at all times for smooth movements. Motion scaling is activated only when the micropipette is closer than a certain threshold to the vessel, as measured by the vision system. A threshold value of 500 μm was selected empirically. Although motion scaling could be implemented with a velocity controller, estimation of the velocities from measured positions is noisy. A preferable approach in our system is to select some reference point $P_R \in \mathbb{R}^3$ for motion scaling, measure the offset between the tremor-compensated null position $\boldsymbol{F_L}(P_N)$ and the reference point P_R, and use a position controller to drive the micromanipulator tip P_T to the measured offset position, scaled by $1/s$. Anisotropic motion scaling is achieved by taking into account the rotation $R \in \mathbb{R}^{3\times3}$ of the tool, and introducing individual scale factors $s_X, s_Y, s_Z \in [1, \infty]$:

$$P_T = S(R\boldsymbol{F_L}(P_N) - RP_R) + RP_R \quad \text{where } S = \begin{bmatrix} 1/s_X & 0 & 0 \\ 0 & 1/s_Y & 0 \\ 0 & 0 & 1/s_Z \end{bmatrix} \tag{8}$$

The reference point P_R is selected as the tremor compensated null position at the time of activating motion scaling $P_R = \boldsymbol{F_L}(P_N)$ and held constant until scaling is deactivated.

As the operator moves Micron during the procedure, the motion scaling displaces the tip away from the null position, reducing the range of motion. At some point, the micromanipulator tip must return to the null position, otherwise the actuators will saturate and Micron will be unable to provide any assistance. Ideally, we would prefer sufficient range of motion to complete the cannulation and withdraw from the vessel before returning the tip to the null position, thus maximizing the range for other operations.

Turning off motion scaling is done when the vision system senses the tip has left the vicinity of the vessel (i.e., the distance between them exceeds a certain threshold). However, directly commanding the micromanipulator tip to the null position causes a rapid and unpleasant twitch in the tool tip. Furthermore, transient tracking errors may falsely trigger turning off the motion scaling. Since tracking is not entirely reliable, the problem of determining when and how to return to pure tremor compensation is exacerbated.

The proposed solution involves a graceful transition from motion scaling to tremor compensation while at the same time gradually re-centering the manipulator at the null position. To account for noisy tracking, the scale is decreased exponentially over time, providing robustness to intermittently noisy or incorrect distance measurements. Denoting t as the number of successive time steps the vision system has detected the tip outside the vicinity of the vessel, the scale factor for each direction is reduced at each time step:

$$s_D \leftarrow s_D - \frac{\alpha}{(1+\gamma)^t} \; \forall D = (X, Y, Z) \tag{9}$$

where α specifies how much to reduce the scale at each time-step. An exponential discount factor γ encodes a sense of how confidence over time increases the convergence rate. Thus, several intermittent noisy measurements will only decrease motion scaling slightly. In contrast, a sequence of measurements indicating that motion scaling is no longer needed signifies a much higher confidence and effects a higher convergence rate back to tremor compensation. Equation (4) terminates for each direction when the scale factor is unity, at which point the manipulator tip has returned to the null position. Alternatively, (4) terminates if the tip is re-detected within the vicinity of the vessel, in which case the reference point P_R is recalculated as:

$$P_R = (S - I)^{-1}\left(SRP_T - R\boldsymbol{F_L}(P_N)\right) \tag{10}$$

and the scale factor is reset to avoid discontinuous jumps in tip position due to changing scale factors. Thus, the control system accomplishes a gentle transition from motion scaling to tremor compensation in a way that robustly deals with noise. Empirical testing indicates values of $\alpha = 3 * 10^{-5}$ and $\gamma = 1 * 10^{-4}$ are effective for a smooth 4–8 s transition.

3.2.5 Experiments & Results

Cannulation experiments were performed ex vivo by an experienced retinal surgeon under two scenarios: Aided (with the assistance of Micron) and Unaided (freehand, with Micron powered off). Porcine eyes were used as the animal model for all experiments. Sections of the back eye around the optic nerve were removed with the retina and vessels still attached. Vessel diameters were measured before each cannulation attempt, and total cannulation duration was recorded. Total

duration excludes time spent clearing any bubbles in the vitreous caused by failed injections. Injections were performed by depressing a 20 mL syringe of air connected to the micropipette. The surgeon performed the insertion of the cannula into the vessel and orally indicated when to depress the syringe. Multiple injection attempts were allowed per trial. Successful cannulation was defined as air entering the vessel and displacing the blood; a cannulation was recorded as a failure when the vessel became too damaged to continue. Trials alternated between several Unaided attempts and several Aided attempts to limit ordering effects.

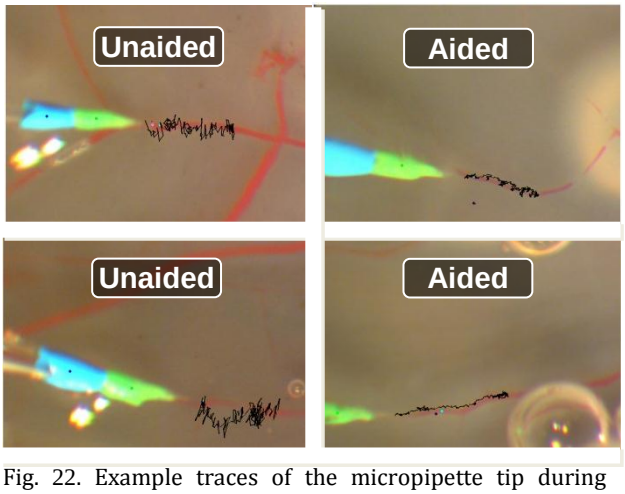

Fig. 22. Example traces of the micropipette tip during cannulation overlaid onto microscope images, showing tremor during the procedure without Micron (Unaided) and with Micron (Aided).

Trials	Total Experiments	Successful Experiments	Successful Durations (s)
Unaided	7	2 (29%)	76±31
Aided	8	5 (63%)	72±19

Table 3. Cannulation results for 40-60 µm diameter vessels

As seen in Table 3, for vessels with diameters between 40 and 60 µm, the cannulation success rate was higher with similar durations in successful trials. Vessels larger than 60 µm exhibit more equal success rates. Visually, Fig. 23 presents the paths inscribed by the micropipette tip during cannulation as traces on the image for both the Aided and Unaided cases. From these images, it can be seen that tremor is lessened and the surgeon can better track the vessel in the Aided case, thus reducing trauma to the vessel and surrounding tissue.

3.2.6 Discussion

In these preliminary surgeon experiments, the robot aid of tremor compensation with motion scaling applied by the active handheld micromanipulator increased the success rate of small vessels while maintaining similar procedure duration. Better tracking algorithms could further increase effectiveness.

Two months later, we performed additional testing with both *ex vivo* porcine eyes

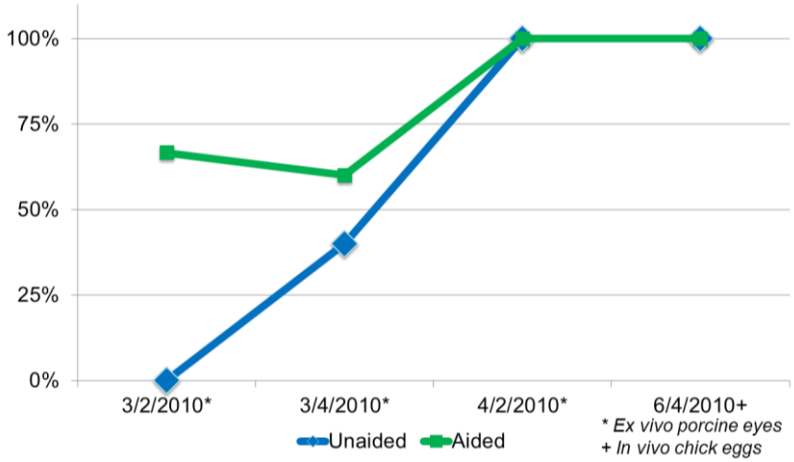

Fig 23. Learning curve of cannulation success for 40-60 µm vessels over four sessions spaced over a 3 month period.

and *in vivo* in chick embryo chorioallantoic membrane [65] to examine the effect of task learning curves as the surgeon becomes more proficient. As shown by the Aided and Unaided test scenario success rates plotted over time in Fig 22, we discovered the surgeon eventually became equally proficient under either scenario. It is clear, however, that the Aided case is useful to ease the learning curve. Further testing is necessary to examine how the learning curve generalizes to more realistic scenarios and other surgeons.

4 APPROACH: POSITION-BASED VIRTUAL FIXTURES

From preliminary work and experimenting with different behaviors for the retinal laser photocoagulation and retinal vessel cannulation, two conclusions can be reached. First, the overall approach of utilizing vision-based, task-specific behaviors for retinal procedures is helpful. Second, the theoretical aspects of the previous work were combined in an ad-hoc manner, whereas a more general, unified, and extensible approach would be desirable. This section describes the central approach proposed in the thesis: the development of a theoretical framework for position-based virtual fixtures derived from the vision system and specifically designed for handheld instruments. Fig. 24 depicts our experimental setup.

4.1 RELATED WORK

Virtual fixtures, in contrast to tremor compensation or motion scaling that operate in the general case, instead aim to improve specific motions or tasks. Much like a ruler 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 [98] as a method to overlay abstract sensory information onto a force-reflecting master workspace, virtual fixtures were originally developed to address latency in teleoperation tasks. They were also found to reduce cognitive load and increase precision and performance [99].

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. [44] presented an optimization technique for commanding multi-linkage robotic arms via a joystick, while keeping joints within their limits, maintaining orientations of the end-effector, and imposing task-space constraints on positions. Kapoor et al. [58] extends Funda's work in constrained optimization to jointly handle simultaneous primitives such as maintain a direction, rotate around a line, stay above a plane, etc. with a weighted optimization. Davies et al. [32] demonstrate the ACROBOT (Active Constraint Robot) with virtual fixtures preventing the surgical cutter from entering pre-defined avoidance zones. Everett et al. [36] showed decreased operation time and operator fatigue in tele-manipulating 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. [61] decomposes surgical tasks into basic system states where the surgeon activates primitive actions with a foot pedal to execute position, orient, approach, insert, etc. subroutines. Bettini et. al [16] introduces 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 behaviors such as trajectory following, positioning, and volumetric restrictions. Hard virtual fixtures disregard any forces in non-preferred directions to prevent deviations from the virtual fixtures, while varying degrees of soft virtual fixtures allow a mixture between human commands and the active virtual fixture. If the robot is non-backdriveable (such as the JHU SteadyHand [80]), strict adherence to the virtual fixture can be enforced by ensuring velocity components are zero in directions that move the tip away from the 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 [33, 66], follow trajectories in tight anatomies [67], and aid in simulated surgical procedures [67, 68].

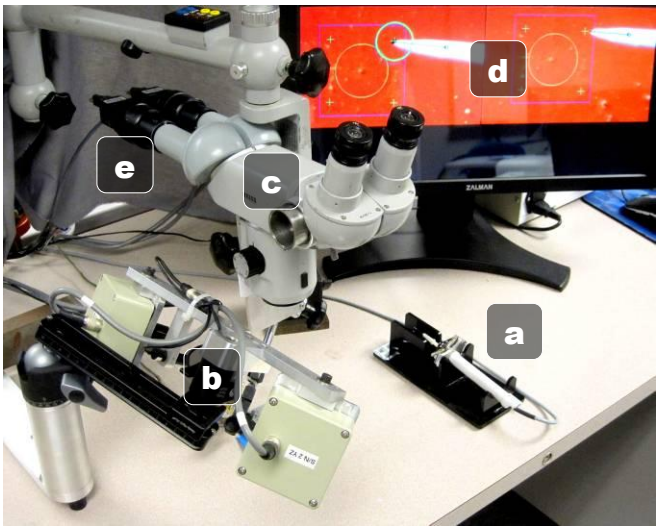

Fig. 24. System setup with (a) Micron, (b) ASAP position sensors, (c) surgical microscope, (d) 3D monitor, and (e) stereo cameras.

Although a significant amount of research has been accomplished in applying virtual fixtures to various robotic platforms and micromanipulation problems, much of the work focuses on master/slave or cooperative micromanipulators. In most formulations of virtual fixtures, the user manipulates a robot arm that is attached to the instrument directly or remotely via teleoperation. Forces or velocities on the robot arm are transformed to velocity commands at the instrument tip, which are shaped first in software. 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 positions; thus, the input to the virtual fixtures must be handle (or hand) motion. Furthermore, because the tool is handheld, the hand movements actually correspond to meaningful gestures relative to the tissue, unlike the arbitrary positioning of master/slave configurations. This fundamental difference, the use of handle position instead of force or master position as the control input, necessitates the development of a different formulation of virtual fixtures specifically designed for this class of handheld instruments. As will be seen later, the development of new control laws is advantageous because motion scaling can be derived quite naturally.

4.2 Formal 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 subspace. In the case of hard virtual fixtures, the tip of Micron should always lie on the subspace; in the case of soft virtual fixtures, the error between the hand motion and the virtual fixture should be scaled by the specified amount. In both cases, tremor compensation should smooth tip movement.

We denote the tip of Micron in 3D space as $P_T \in \mathbb{R}^3$. As Micron actuates, the tip will move independently of the hand (or handle) motion. We define the null position as the 3D tip position $P_N \in \mathbb{R}^3$ under the assumption that Micron is off; in another words P_N exactly reflects the hand motion. One can think of P_N as being mechanically tied to the handle, and unless Micron actuates its piezoelectric motors, $P_T = P_N$. 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 of the operator's intentions.

4.3 Point Virtual Fixtures

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

$$P_T = P_G \tag{11}$$

and use approximated linear inverse kinematics with a PID 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 <200 ms if the goal point is within Micron's current range of motion. Once the point virtual fixture is activated, Micron uses the PID controller to maintain the tip on the goal at a control frequency of 1 kHz. Moving the virtual fixture, for instance in response to shifting anatomy, is possible by setting a new goal point P_G. To avoid high-frequency oscillations caused by PID overshoot, large changes in P_G should be smoothed with either a low-pass filter or a minimum jerk trajectory.

4.4 Higher-Order Subspace Virtual Fixtures

Higher order subspaces can be built on top of the point virtual fixture to obtain more interesting virtual fixtures. Each additional level adds a degree of freedom to the tip motion. For instance, the point virtual fixture restricts all motion. However, the line virtual fixture allows the tip to freely travel along a line while restricting motion orthogonal to the line. Likewise, we can define a hierarchy of virtual fixtures as point, curve, surface, and volume (see Fig. 25). 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 [14], scanning tissue with OCT [7]
- **Curve** (1 DOF): Following path for laser oblation, 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 [32]

Each virtual fixture in the hierarchy collapses to a special case of the one below it. 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 right point on the higher-order virtual

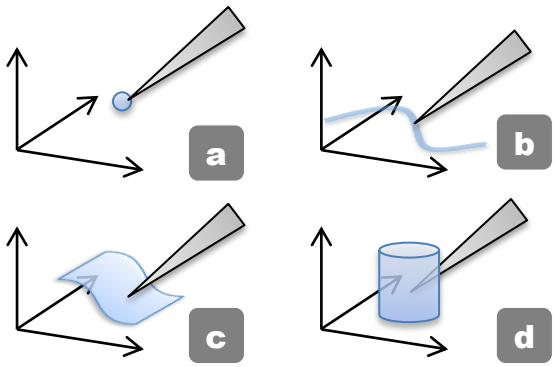

Fig. 25. Virtual fixtures constraining the tip to a subspace with increasing degrees of freedom: (a) point, (b) curve, (c) surface, and (d) volume.

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 optimal 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:

$$P_G = \boldsymbol{M_O}(V, P_N) \tag{12}$$

For the simple geometric structures used in this paper (lines, planes, circles, cylinders, etc.), analytic solutions for $\boldsymbol{M_O}$ exist. For curves and surfaces described by more complex models such as NURBS, interested readers should look to numerical methods [72].

4.5 Tremor Suppression

In point virtual fixtures, all degrees of freedom are proscribed by the fixture. However, for higher-order virtual fixtures, tremor parallel to the subspace is not affected by the orthogonal projection $\boldsymbol{M_O}$. In fact, tremor is evident in all 3 DOF for volume virtual fixtures. One easy solution is to use a tremor suppression filter $\boldsymbol{F_T}: \mathbb{R}^{\boldsymbol{n}} \rightarrow \mathbb{R}^{\boldsymbol{n}}$ on the null position P_N before it is orthogonally projected to the virtual fixture.

$$P_G = \boldsymbol{M_O}(V, \boldsymbol{F_T}(P_N)) \tag{13}$$

Alternatively, the tip position can be re-parameterizing on the subspace represented by the virtual fixture with a lower degree of freedom. In this case, we redefine the mapping $\boldsymbol{M_O^n}: \mathbb{R}^3 \rightarrow \mathbb{R}^n$ 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 single dimensional tremor suppression filter $\boldsymbol{F_T}$. For an n DOF virtual fixture, we can calculate the goal point as:

$$P_G = \boldsymbol{F_T}\left(\boldsymbol{M_O^n}(V, P_N)\right) \tag{14}$$

For experiments in this paper, we use Eq. 3 for volume and plane virtual fixtures and Eq. 4 for curves to prevent any nonlinearities in the orthogonal projection from affecting the tremor suppression. 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 with a low-pass filter, Kalman Filter, or some other method.

4.6 Motion Scaling

So far we have been describing hard fixtures where the tip position cannot deviate from the constraint imposed by the virtual fixture. However, most virtual fixture derivations introduce soft virtual fixtures that share control between the virtual fixture and the operator. If hard virtual fixtures are analogous to an unyielding metal ruler, then soft virtual fixtures can be likened to a yielding rubber ruler that helps draw straight 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 actuates to the tip of the instrument.

$$P_T = (1 - \lambda)P_G + \lambda P_N \tag{15}$$

In essence, λ functions as a weighted average of the goal and null position. $\lambda = 0$ corresponds to a hard virtual fixture, and $\lambda = 1$ disables virtual fixtures entirely. However, for values of λ between 0 and 1, Eq. 5 can be directly manipulated into a motion scaling paradigm:

$$P_T = P_G + \lambda(P_N - P_G) \tag{16}$$

$$e = P_N - P_G \tag{17}$$

$$P_T = P_G + \lambda e \tag{18}$$

Thus, the parameter λ is seen to be the scaling factor on the error e between the hand position P_N and the goal point P_G on the virtual fixture. For example, if $\lambda = 1/2$, then all hand motions that deviate from the virtual fixture will be scaled by one half.

4.7 Generalized Control Law

To summarize the control law that incorporates virtual fixtures, tremor suppression, and motion scaling into the Micron framework, we assume for simplicity's sake that tremor suppression on the null position is sufficient and re-parameterization is not necessary. The generalized control law then becomes:

$$P_G = \boldsymbol{M_O}(V, \boldsymbol{F_T}(P_N)) \tag{19}$$

$$e = \boldsymbol{F_T}(P_N) - P_G \tag{20}$$

$$P_T = P_G + \lambda e \tag{21}$$

Eq. 9 selects the goal point on the virtual fixture closest to the tremor suppressed null position. Eq. 10 then calculates the error between these two points. Finally, Eq. 11 drives the tip to either the virtual fixture or, if λ is non-zero, scales the error to achieve motion scaling about the virtual fixture.

4.8 Composition of Virtual Fixtures

The primitive virtual fixtures of point, curve, surface, and volume can be easily combined into compound virtual fixtures. 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 $\boldsymbol{M_O}$ 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. One solution to prevent the tip from oscillating wildly between virtual fixtures is to disable all fixtures except the current one until the tip nears the edge. Only when the tip reaches the boundary between fixtures can Micron transition from one virtual fixture to another. This approach allows the construction of stable compound virtual fixtures.

4.9 SYNTHETIC TESTS

Several experiments were performed with Micron to validate the presented virtual fixture formulation. We use a similar experimental setup and series of tests as [30]. As depicted in Fig. 26, virtual fixtures are derived from a target laser etched in to rubber. The target includes four crosses arranged as corners of a 600 μm square with a 500 μm circle in the middle. Four tasks were considered:

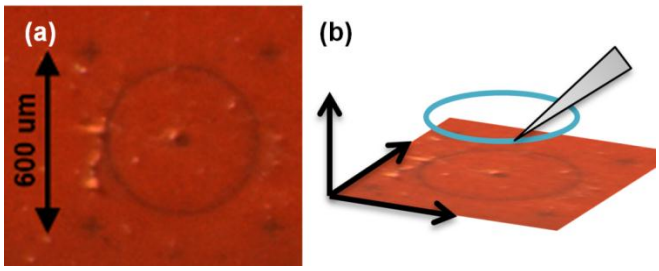

Fig. 26. (a) Laser etched target onto rubber surface (b) Generating a circle virtual fixture from the tracked target.

- **Hold Still**: Hold the tip of the instrument 500 μm above the top right cross for 60 s. A point fixture was used.
- **Circle Tracing**: Trace a circle with a 500 μm offset from rubber surface three times. A 3D curve fixture in the shape of a circle was used.
- **Move and Hold**: Move the tip to each cross sequentially and hold at each cross for 3 s. Four plane segments oriented vertically and connecting each of the four crosses 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 vertical the offset from the rubber surface.
- **Volume Restriction**: Allow free movement inside the volume while preventing the tip of the instrument from leaving the volume. A cylindrical volume was used with the circle serving as the perimeter constraint and the rubber surface serving as the bottom constraint.

All experiments were performed by a single individual familiar with Micron but without surgical experience. Experiments were performed randomly to alleviate ordering effects. In the first three experiments, four different scenarios were tested three times.

- **Unaided**: Micron was turned off.
- **Aided with Shelving Filter**: Micron was turned on with the state-of-the-art tremor suppression and relative motion scaling algorithm from [30].
- **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 and using the virtual fixtures without motion scaling ($\lambda = 0$).

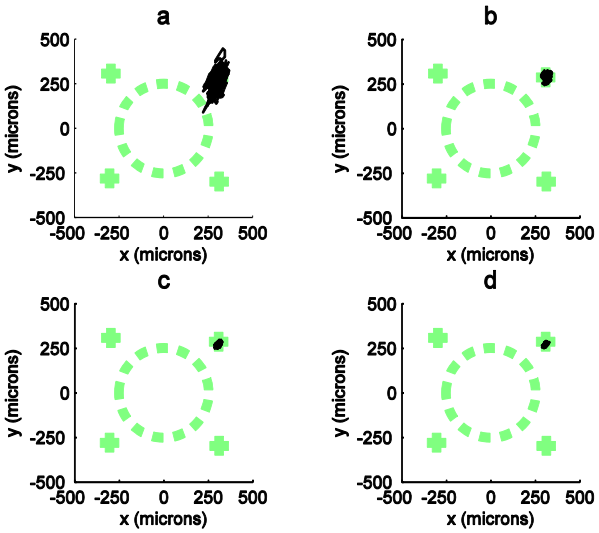

Fig. 27. Hold still results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.

All error was measured as the distance between the tip position measured by the ASAP optical trackers and the closest point on the virtual fixture as measured by the stereo cameras. Error was recorded at ~2 kHz. Because hard virtual fixturing attempts to remove all error between the tip and the virtual fixture, it is imperative to provide the operator with some visual feedback that lets the operator keep the gross positioning of the handle near the virtual fixture. Otherwise, the operator will unknowingly drift away from the virtual fixture error eliminating controller is effectively disabling the eye-hand feedback loop. To remedy this, we provide the operator an on-screen visual cue that provides 3D information as to the location of the unseen null position P_N. In all tests, the operator used these visual cues to maintain correct positioning.

Fig. 30 displays mean RMS error with standard deviation across all three trials for each task and scenario combination, demonstrating that virtual fixtures are effective in reducing error. Fig. 28, Fig. 29, and Fig. 27 show one trial of each task and scenario with a trace of the 2D tip position overlaid in black on the target represented by dashed light green lines. Maximum error for each set of trials is presented in Table 4. It is interesting to note that while the Move & Hold has much lower overall error, it is an artifact of the vertical orientation of the fixtures. Since the fixture is not restricting Z-movement, the controller is absorbing the error normally caused by poor depth perception. In volume restriction tests, hand motion exceeding 500 μm was successfully restricted to the volume with only a max error of 25 μm.

	Unaided	**Aided**		
Task	Micron Off (μm)	Shelving Filter (μm)	Soft Fixtures (μm)	Hard Fixtures (μm)
Hold Still	489	253	165	67
Circle Tracing	446	312	100	171
Move & Hold	246	172	52	109

Table 4. Max error in positioning tasks

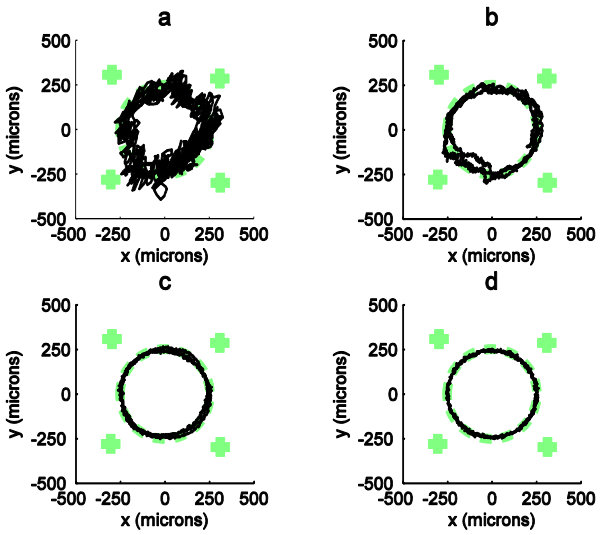

Fig. 29. Circle tracing results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.

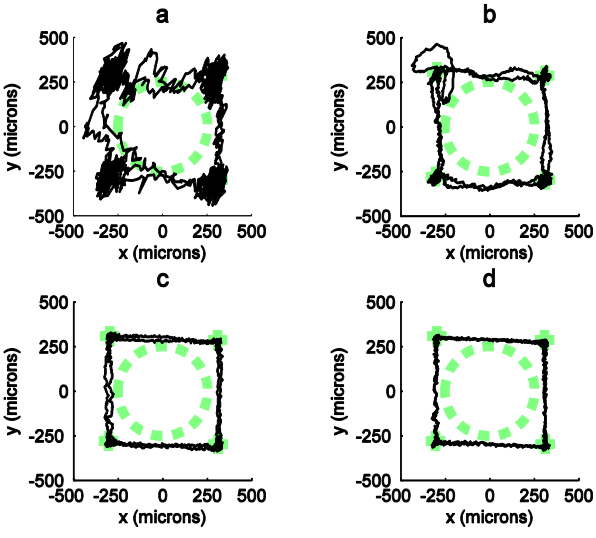

Fig. 28. Move & hold results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.

4.10 DISCUSSION

In this section, we have presented a derivation of virtual fixtures for handheld micromanipulators that depends not on forces applied to a robot arm, but hand motion at the handle of the instrument. The virtual fixtures are generated from stereo cameras in real-time and integrate tremor suppression and motion scaling. Using Micron as a test platform, virtual fixtures have been validated with medically relevant artificial tests and have been shown to outperform state-of-the-art tremor compensation in reducing the positioning error of the micromanipulator.

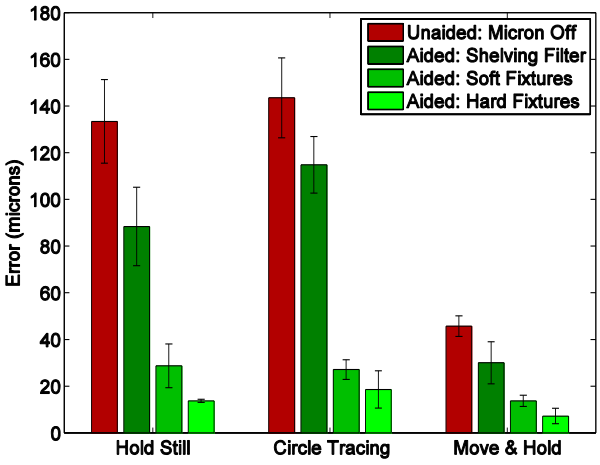

Fig. 30. Mean RMS error across three trials of each combination of task and scenario, with error bars representing the standard deviation.

While defining virtual fixtures based on the null position of a handheld micromanipulator introduces difficulties relating to degenerate solutions, this formulation offers a number of advantages that simplifies implementation compared to force-based virtual fixtures. No special treatment is needed when maintaining a point virtual fixture as it is in [16]. Implementation of soft virtual fixtures under our formulation directly results in intuitive motion scaling. On the other hand, the handheld nature of Micron means the system depends on reasonable gross hand motion so that the actuated tip remains in the instrument's range of motion. This can be viewed as a disadvantage because Micron cannot prevent the surgeon from violating the virtual fixture; however, a concomitant advantage is that the surgeon can override Micron in emergency situations.

5 PROPOSED RESEARCH

Formalizing the previous virtual fixtures work into an extensible, unified framework of compositable primitives for use in surgical procedures is central to the proposed work. Supporting research falls into three categories: model predictive control, 3D vision imaging, and practical implementation with task-specific behaviors that target individual surgical procedures. The proposed work will be evaluated in the two different retinal surgical procedures tested previously: retinal laser photocoagulation and retinal vessel cannulation. A schedule detailing the timeline of proposed work is also presented so as to finish the proposed work reasonably in about a year's timeframe.

5.1 CORE THESIS COMPONENTS

With the goal of enduing the micromanipulator with knowledge of the procedure and the surgeon's operation plan so as to provide task-specific aids, we propose research into four key areas. The first area of research is to formalize and extend the scope of the virtual fixtures framework developed in the preliminary work. Second, we propose to implement a model predictive control law that models system latency and predicts hand tremor to decrease high-frequency noise caused by high velocity hand movement. The third proposed work area is to use stereo cameras mounted to the microscope for improved tracking the micromanipulator's tip and retinal anatomy in 3D space. The fourth and final core component of the thesis is to apply the theoretical work in virtual fixtures framework, model predictive control, and 3D vision imaging to develop a prototype surgical system that targets retinal laser photocoagulation and retinal vessel cannulation procedures.

5.1.1 VIRTUAL FIXTURES FORMALIZATION

Preliminary work shows position-based virtual fixtures are helpful in significantly improving positioning accuracy in synthetic, yet relevant primitive tasks such as path tracing. However, further work must be accomplished before the virtual fixture framework is readily applicable to complex surgical techniques. We propose to use a more general spline parameterization for fixtures, add orientation fixtures, and develop a cleaner compositing system for combining simple fixtures into complex ones.

Spline Parameterization of Higher Order Virtual Fixtures

The major limitation of the proposed virtual fixtures framework is the inability to use a generalized parameterization for constructing primitive virtual fixtures. Only circles and lines are currently supported, and their parameterization based on the visual information is hard-coded into the system. Certainly this level of micro-managing the virtual fixture representation is tedious. Ideally, the surgeon should be able to draw arbitrary paths to follow or boundaries not to cross; furthermore, if the surgeon specifies anatomy to avoid or track, the virtual fixtures should be automatically generated from vision information. This argues for a flexible and dynamically generated representation for virtual fixtures. One possible solution, often used in graphics, is to approximate complex curves and surfaces with many primitive shapes such as line segments or triangles. Better quality approximations require increasingly large numbers of these primitives, yielding a fairly unwieldy representation.

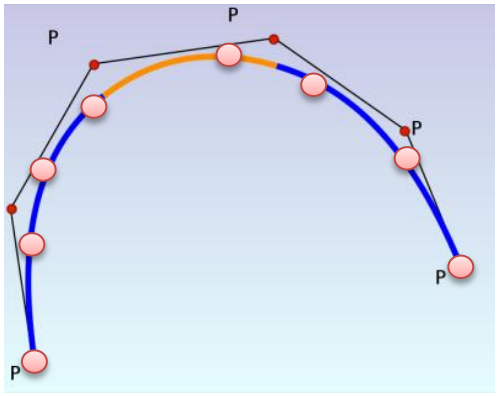

Fig. 32. Splines can represent arbitrary curves and can be fit to measured data points. Photo credit Gary R. Osgood.

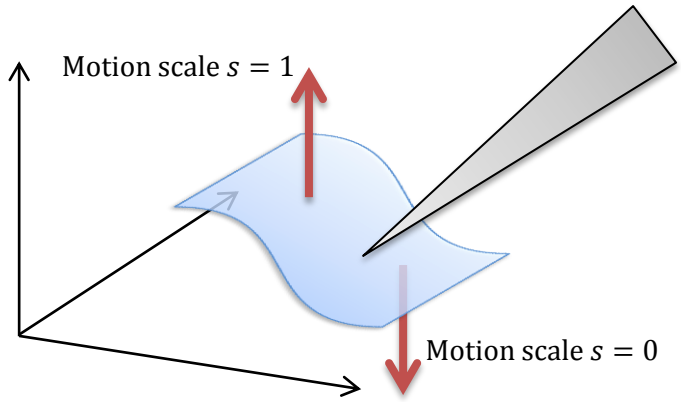

Fig. 31. Volumetric motion-scaling by selecting different motion scaling parameters for each side of a surface.

Instead, we hypothesize that a spline-parameterization of higher order virtual fixtures will be an effective and easy-to-use representation. A spline is a curve defined by piecewise polynomials where control points define the transition between the polynomials, as shown in Fig. 21. In particular, B-splines are a popular spline representation that use basis functions to define internal knot points that the curve must pass through. Complex and smooth trajectories can be created with relatively few parameters and have a number of advantageous features. B-splines can be fit to measured data [71, 113], can be easily extended to surfaces [5, 29, 35], and are already used in a number of surgical approaches [32, 33, 66, 67].

An important requirement for all virtual fixtures in our framework is the ability to do orthogonal mapping (often referred to as point projection) to find the closest point on the fixture, which is made possible by a wide variety of numerical algorithms [25-27, 70, 72]. Volume constraints can be enforced by B-spline surfaces with side-dependent motion-scaling (e.g. motion scale that changes depending on whether the tip is above or below the surface fixture as in Fig. 31). Thus, when crossing the surface into the area defined by separating surface fixture, the tip motion can be stopped by a hard virtual fixture or motion-scaled by a soft virtual fixture. Constructing multiple of these surfaces allows for arbitrary shaped volumes, although care needs to be taken during construction to prevent incompatible settings.

Orientation Virtual Fixtures

While tip position fixtures are available in the virtual fixtures framework, an existing omission is control of the instrument orientation. Orientation-based virtual fixtures are particularly important for laser-based applications where the surgeon is less interested in tip position and more interested on where the laser beam hits the anatomy. Perhaps the best way to formulate such a problem is with projective geometry. By controlling the tip position and orientation, the laser beam can be treated as a ray in space. To enforce an orientation virtual fixture that keeps the laser beam incident on a particular position, the ray representing the laser beam must intersect the Micron pivot position, the tip, and the desired target on the anatomy as shown in Fig. 33. This orientation problem is under-constrained because both the pivot point and the tip can be changed. For maximum effectiveness, additional constraints are necessary, such as on the tip position or minimum actuation effort to solve for the least-squares solution.

Unfortunately, several non-idealities are clearly evident in the system. The tip may not be straight, or the fiber optic sheathed inside the tip may not be straight. The location of the pivot point is currently treated as a rigid transform; however, it is dependent on mechanical flextures which may be difficult to model. Introducing a skew angle at the tip of the instrument, creating a lookup table for pivot positions at various angles, and performing other calibrations of the system non-idealities is likely to improve estimates, although the calculated orientations will likely only be good enough for first order approximations. Feedback control on the tracked laser beam will probably still be required for the most accurate orientation virtual fixture. Unfortunately, feedback control on the tracked laser beam depends on the video cameras, which have a much slower update rate and higher latencies. Incorporating the vision information into the feedback loop will be left for the more application-driven task-specific behaviors if needed.

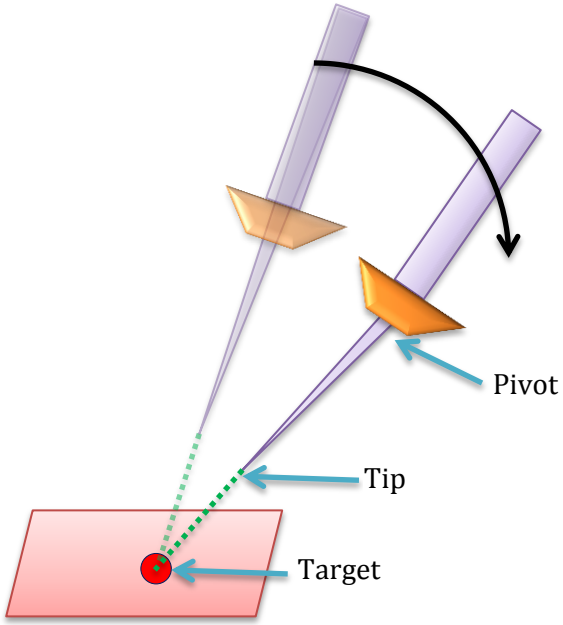

Fig. 33. Orientation virtual fixtures constrain the tool to be oriented towards a point.

Composition of Virtual Fixtures

An area of weakness in the current formulation is a neat, compact, and efficient method for combining primitive virtual fixtures to encompass more complex surgical movements. For clarity, we define composition of virtual fixtures as two distinct classes: temporal compositions and spatial compositions. Temporal compositions model how virtual fixtures change over time, whether in response to changing anatomy or changes in the surgeon's current plan, i.e. transitioning from one step in a procedure to another. An important consideration is how to enforce the virtual fixtures in the presence of temporal uncertainty due to the way virtual fixtures are derived from noisy vision data. Spatial compositions occur when multiple virtual fixtures exist simultaneously but are separated spatially. Examples of this include compound virtual fixtures such as boxes constructed from planes and compositions of position and orientation virtual fixtures. Another important consideration is the priority of each virtual fixture. During operation, it may become impossible to satisfy all of the fixtures' constraints. A form of management will be necessary to choose which virtual fixtures have priority in cases of conflict. We propose to extend the virtual fixtures framework to allow easy combinations of prioritized virtual fixtures.

5.1.2 Model Predictive Control

Experiments with virtual fixtures has shown that as we increasingly drive down positioning error, the latency in our control effort is becoming one of the largest contributors to positioning inaccuracy.

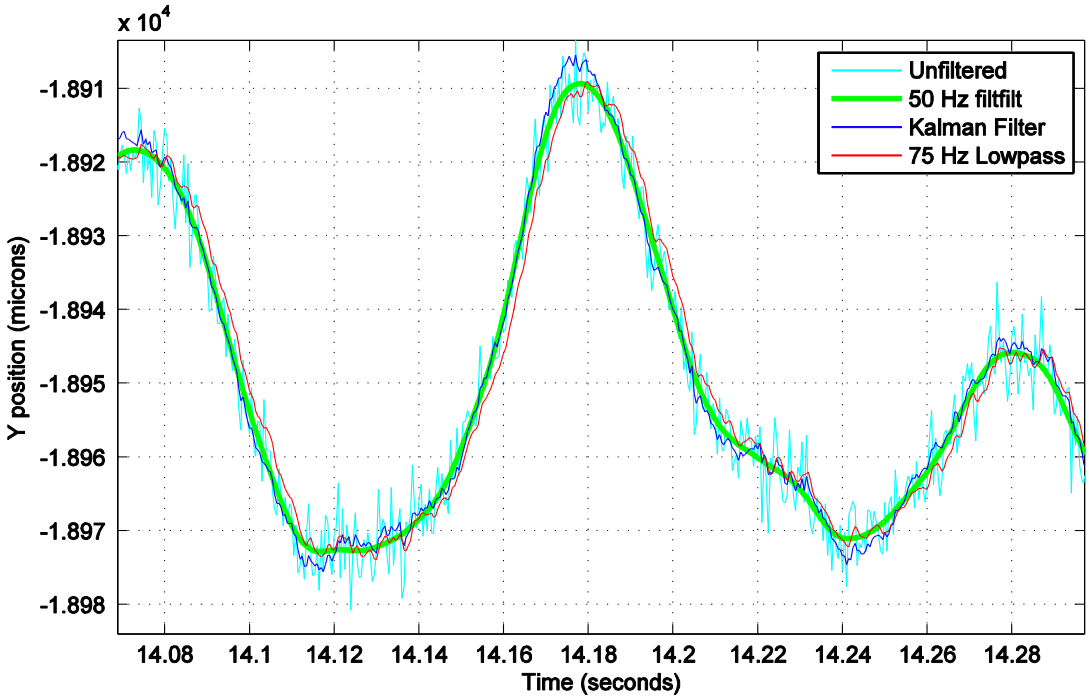

Fig. 34. Comparison of an Extended Kalman Filter and a low-pass filter on recorded tremor data for state estimation of position. The true position is estimated as a 50 Hz forward/backward (filtfilt) filter operation.

State Estimation

Good state estimation is useful for reducing high frequency noise and for accurate predictions into the future. One of the key observations in reducing noise is that some behaviors are physically impossible, i.e. extremely rapid changes in position are simply not possible. In another words, there is a model that describes the physical behavior under observation. Likewise, there are physical constraints on the measurements, i.e. the entire tool must rigidly rotate and translate. These additional constraints on the measurement can be employed to reduce noise. A related problem is estimating states of the system that are not directly measured. For instance, Micron measures only positions; however, in many situations, an estimate of the velocity or acceleration is useful. Numerical differentiation is possible, but is extremely noisy.

Similar to [34], we propose to use the Extended Kalman Filter (EKF) to simultaneously reduce noise and provide state estimates for the Micron system. The available measurements are the positions of four LED lights on Micron; three are mechanically located on the actuated shift/tip and one is located on the handle. Sensed by two PSDs, stereo reconstruction is used to obtain 3D positions of each of the four LED lights. We propose to use these 3D positions in the measurement model, although it may be worth investigating using the raw PSD data as the observations and the stereo projection as the measurement model. The desired output of the EKF is the full 6-DOF pose of the tip and handle (i.e. translation and rotation), and their derivatives (i.e. translational velocity/acceleration and angular velocity). The input is the poses estimated from the 3D measurements using the three coplanar point pose estimation by Horn's method [53]. Good state estimation benefits both low-level and high-level control systems. Fig. 34 shows a comparison of an EKF vs. a low-pass filter for estimating the true position of the tip of Micron by the proposed

system. Notice that the EKF overshoots by a few microns during rapid velocity changes, but does not incur the several millisecond time lag that a low-pass filter exhibits.

Hand Tremor Feed Forward Control

With virtual fixtures, it is theoretically possible to eliminate all positioning error and keep the tip of Micron perfectly located on the fixture. However, practically there are still sources of inaccuracies in the system that prevent perfect positioning. Investigation into the sources of these inaccuracies revealed that hand tremor is one of the largest contributing causes of positioning inaccuracy. While the majority of tremor is suppressed, a small part of the tremor is still evident in the output, causing a high frequency jitter motion around the virtual fixture. Upon further analysis, the error was discovered to be highly correlated with the velocity of the hand motion. Because there is a 3-6 ms latency in actuation and the feedback control is based on the position of the tip of the instrument, high velocities cause the goal position calculated to cancel tremor to be out of date by the time the actuation is accomplished.

We propose to develop models of the hand tremor, predict the motion of the tremor 3-6 ms into the future, and feed forward the results into the control system to more effectively suppress the tremor. Using the state estimation from the EKF, tip velocities will be available. By fitting models to the hand velocity and feeding forward the 3-6 ms velocity-compensated goal positions into the control system should provide more accurate positioning and better tremor suppression. Tremor models such as FLC [93, 112], autoregressive [12, 119], and Kalman filters [19, 20] will be investigated.

5.1.3 3D Vision Imaging

In order to provide the surgeon task-specific behaviors and assistive virtual fixtures derived from anatomy, the system must be able to quickly and accurately track the instrument tip and relevant anatomy in 3D. To avoid broadening the scope to all surgical tracking, we focus on the previously specified problem domain of retinal laser photocoagulation and vessel cannulation. In these surgeries, the necessary 3D vision imaging work can be broken down into several components: vessel localization and anatomy tracking.

Vessel Localization

In both domains of interest, vessel identification and localization is of great importance. In laser photocoagulation, accidentally hitting a vessel with the laser can result in blindness. In vessel cannulation, accurate localization is required to aid the surgeon in piercing the vessel and injecting medicine. Color-based models such as [118] are effective at locating the 2D locations of the vessels; however, the 3D depth of the vessels is more difficult to estimate. A full 3D representation of the vessel at each point in the image is essential to providing task-specific behaviors. Because vessels are usually smooth and textureless, neither typical stereo matching algorithms nor descriptor-based matching works well. Although existing approaches exist [69, 79], they work with holistic views of the entire retina such as fundus images. The proposed work will attempt to apply these approaches to smaller regions of the retina at higher magnifications. Conceptually, such algorithms should yield results as seen in Fig. 35.

Anatomy Tracking

Virtual fixtures are usually defined relative to some anatomy. For example, one virtual fixture might provide motion scaling around a vessel and another one might prevent the tip from touching the retina. Therefore, one important vision task is to track anatomy of interest to accurately generate the virtual fixtures. Eye movement is normal during retinal surgery as the surgeon will often rotate the entire eye in its socket via the schlerotomy. Similar to tracking camera movement in a typical computer vision application, knowing when the scene is moving is useful for maintaining correctly generated virtual fixtures and developing feed forward models for the control system. Likewise, tracking anatomy to target (for instance, with the laser beam in photocoagulation) or to avoid (such as the retina in case of cannulation) can influence the success of the surgery and limit collateral damage. A wide variety of computer vision algorithms could be applied to this problem, including tracking keypoint descriptors such as SURF [5]. Another option is template trackers [6] that can be initialized around anatomy of interest and iteratively located in each frame. Full reconstruction of the retina in 3D has been reported [28] as have various surface parameterizations of stereo reconstruction [29, 90]. Such algorithms will be applied and modified as necessary to track anatomy and construct meaningful virtual fixtures.

5.1.4 Task-Specific Behaviors

The final core component is to combine all the theoretical aspects of the proposed work and apply it to realistic surgical scenarios. Specifically, we propose to develop task-specific behaviors for two target retinal procedures: laser photocoagulation and vessel cannulation. In both instances, a large part of the development of the task-specific behaviors is not only what behaviors to select, but when to activate them and how to transition between behaviors. Much of this research will require working closely with the surgeon on the project to determine which aids are most effective.

Retinal Laser Photocoagulation

In laser procedures, there are several behaviors that are useful to the surgeon. Tremor compensation serves as both gross positioning aid as well as a psychological aid to instill trust into

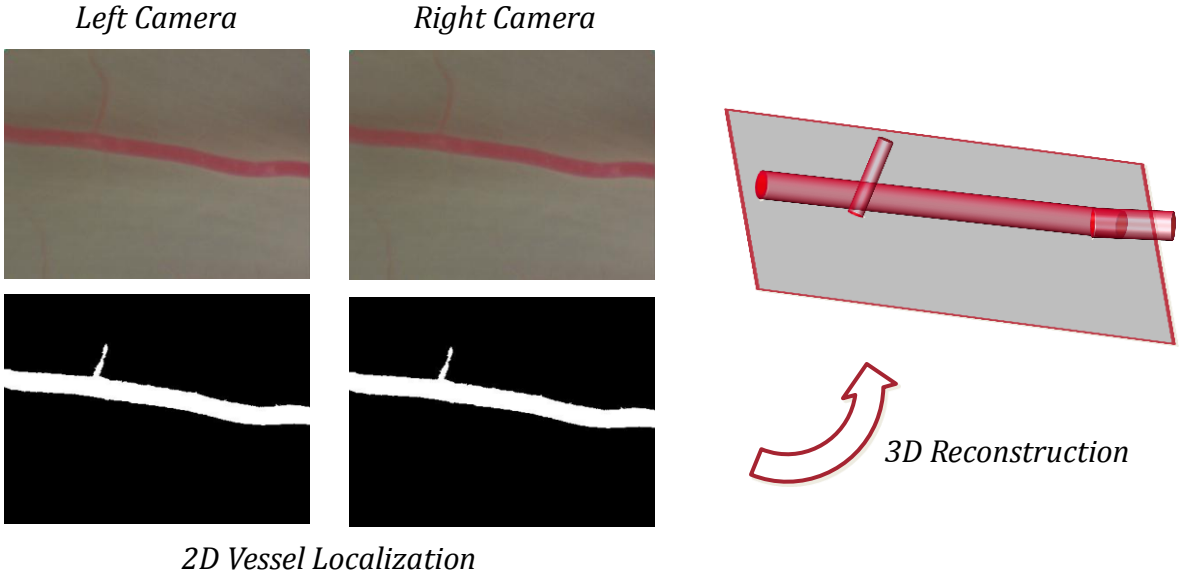

Fig. 35. Desired vessel reconstruction should use stereo matching to locate the image in 2D and reconstruct it in 3D.

the surgeon. Orientation virtual fixtures are particularly useful to burn designed targets with the laser, as are curved trajectories to move between orientations. An important constraint to maintain is the distance from the tip to the target location to ensure the laser burn is effectual, yet not too hot. Such constraints and others must be considered to build the correct set of virtual fixtures and transition between them during the operation.

Retinal Vessel Cannulation

In cannulation procedures, fine positioning is extremely important as vessels can be less than 50 μm in diameter. The approach of the tip to the vein needs to be smooth, the actual insertion needs to be free from axial movement, and during infusion the tip must be as motionless as possible. Equally important is preventing the tip from accidentally coming into contact or gouging the retina on which the vessel lies. A number of virtual fixtures can be used for this procedure. First, if the retina surface can be accurately located, volume constraints can prevent the needle from touching the retina. Second, a curve virtual fixture can enforce the tip to follow a smooth trajectory to the vein. Third, once the tip punctures the vessel, a point virtual fixture can keep the needle in the vessel for the infusion.

5.2 EVALUATION

To validate the effectiveness of the proposed research, a number of tests will be performed with novices and surgeons. For each of the three hypothesis proposed earlier, we present a set of experiments for validation.

1. *Micromanipulation tasks can be decomposed into steps, during which applying tip position constraints (virtual fixtures) with a handheld tool increase operator precision and safety.*

The theoretical aspects of the work, such as the virtual fixtures framework including the model predictive control, will be tested by novices and surgeons in synthetic tests such as hold still and curve tracing. In these experiments, we assume a-priori knowledge of the behavior the surgeon is performing and enforce virtual fixtures that aid the behavior. For more applicable tests, synthetic touching and pointing tasks will be performed to simulate vessel cannulation and laser photocoagulation tasks. Accuracy in these tasks will be measured with varying degrees of micromanipulator control, ranging from none to full. Safety will be measured by defining avoidance zones and measuring how well the virtual fixtures framework keeps the tip inside these zones.

2. *Real-time visual analysis of anatomy and the surgeon's movements provides the context that allows the controller to select and customize the virtual fixtures.*

The 3D vision imaging will be validated experimentally in so much as the reconstructions seem internally consistent (i.e. physical contact results in contact in the tracked 3D models). Other imaging analysis will be evaluated both offline and online. For instance, alignment accuracy between pre-operative images and video images will be reported. To determine if the task-specific behaviors activate properly based on the context, tests under a variety of operating conditions will be simulated and logs of the behaviors that activate (and the chosen parameters) will be analyzed.

3. Vision-based virtual fixtures aid surgeons during microsurgical operations.

The task-specific behaviors represent a more practical application of the proposed work and will be tested by at least one surgeon in surgical procedures with *ex vivo* or *in vivo* animal models. For retinal laser photocoagulation, paper slides with printed patterns or excised porcine retina will be used as the model. For retinal vessel cannulation, excised porcine eyes or chicken eggs [65] will be used as the animal model. Tests with surgeons on live animal models, such as *in vivo* rabbits, are planned as most applied test of the thesis concepts in real microsurgical conditions. We will analyze the resulting data for positioning accuracy, operation time, surgeon's cognitive load, and operation success rates.

5.3 EQUIPMENT BUDGET

In order to complete the proposed work with optimal results, additional hardware would be helpful. Chiefly, higher resolution and framerate cameras would increase accuracy in tracking the instrument tip and anatomy, yielding more stable virtual fixtures. To handle the increased data rates from the cameras, a new multi-core server-grade computer is necessary. Finally, to provide visual cues to the user, various lenses and displays are needed to inject images directly into the microscope. A tentative budget is shown in the table below.

Item	**Price**	**Date**	**Justification**
GTX 470 Graphics Card	$300	Nov '10	For fast GPU image processing
(2X) Flea3 USB 3.0 Cameras	$2000	Sum '11	High resolution (1280x1024) cameras at high speeds (up to 150 Hz) for 3D sensing.
Dual Opteron 6128 Computer	$1700	Sum '11	16 core server grade computer to process visual data and handle composite fixtures
Display, lenses, & other optics	$1000	Sum '11	For image injection into the microscope

5.4 WORK SCHEDULE

We propose a year-long timeline to complete the proposed work. Development of the virtual fixtures framework and preliminary model predictive control work has already been accomplished. During Spring 11', the thesis proposal will be given and the model predictive control will be finished. Summer '11 primarily focuses on formalizing the virtual fixtures framework and developing the 3D vision. Fall '11 consists of applying the new virtual fixtures framework to the task of vessel cannulation and laser photocoagulation and improving the vision sensing. Spring '12 is reserved for final experimentation, dissertation writing, defense, and revisions.

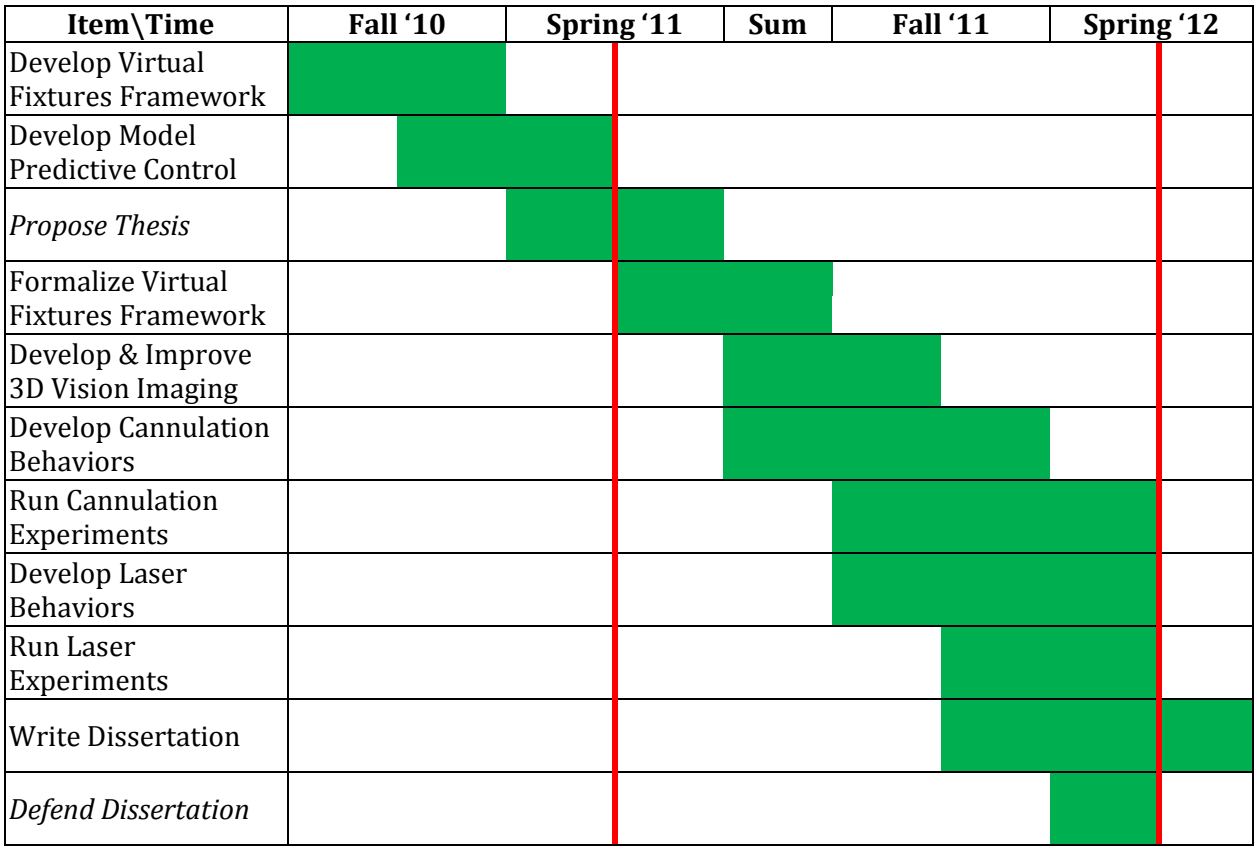

6 CONCLUSION

In this thesis proposal, we have reviewed some of the challenges in micromanipulator behavioral aids and identified the central hypothesis that the development of more effective behaviors can be achieved by the fusion of vision and control to provide real-time, contextual aids to the surgeon. We have presented preliminary work in retinal laser photocoagulation and vessel cannulation, motivating the development of a more general and extensible a framework that incorporates tremor suppression, motion-scaling, and virtual fixtures. We have proposed four core areas of research for the thesis: formalization of the virtual fixtures framework, development of model predictive control, 3D imaging of anatomy, and integration with task-specific behaviors. Evaluations will be performed on *ex vivo* and *in vivo* tissue, demonstrating the applicability of the research to microsurgical operations.

6.1 EXPECTED CONTRIBUTIONS

As a result of proposed research, we expect the main contributions of the thesis will be:

- The formalization of a general and extensible position-based framework for handheld microsurgical instruments that incorporates tremor suppression, motion scaling, and virtual fixtures.
- The development of model predictive control and tremor models to explicitly address latencies in the control systems of handheld micromanipulators.
- The integration of control and vision algorithms into a unified system for vitreoretinal microsurgery.

6.2 FUTURE DIRECTIONS

The proposed work makes a progressive stride towards the goal of synergistic behavior aids that attempt to understand the surgeon's operation plans, but this only represents a small step in the direction of intelligent aids for the surgeon. Future work should include investigating how the proposed virtual fixtures framework can be applied to a wider range of surgeries. New sensor modalities such as OCT should be investigated to allow sub-micron accuracy in 3D tissue reconstruction [7]. Hardware improvements to increase the range-of-motion, increase slew rates, and decrease latency would be of great benefit in many surgical applications. Another future area of great interest is force sensing, where Micron could not only sense, but control how much force is being applied to delicate tissue in order to prevent accidental damage. Ideally, research in these and related directions will yield robotic platforms that couple surgeons' expert knowledge with the positioning accuracy of micromanipulators to make the operating room a less scary place.

7 WORKS CITED

[1] J. J. Abbott, G. D. Hager, and A. M. Okamura, "Steady-hand teleoperation with virtual fixtures," in *Robot and Human Interaction*, 2003, pp. 145-151.

[2] H. Ahmadieh, H. Sajjadi, M. Azarmina, M. Soheilian, and N. Baharivand, "Surgical management of intraretinal foreign bodies," *Retina,* vol. 14, pp. 397-403, 1994.

[3] N. Ahmidi, G. Hager, L. Ishii, G. Fichtinger, G. Gallia, and M. Ishii, "Surgical Task and Skill Classification from Eye Tracking and Tool Motion in Minimally Invasive Surgery," *MICCAI,* pp. 295-302, 2010.

[4] B. Allf and E. de Juan, Jr., "In vivo cannulation of retinal vessels," *Graefes Archive for Clinical and Experimental Ophthalmology,* vol. 225, pp. 221-225, 1987.

[5] G. Ateshian, "A B-spline least-squares surface-fitting method for articular surfaces of diarthrodial joints," *Trans Am. So. Mech. Engr. J. of BioMech Engr. ,* vol. 115, pp. 366-366, 1993.

[6] S. Baker and I. Matthews, "Lucas-kanade 20 years on: a unifying framework," *Intl J of Computer Vision,* vol. 56, pp. 221-255, 2004.

[7] M. Balicki, J. H. Han, I. Iordachita, P. Gehlbach, J. Handa, R. Taylor, and J. Kang, "Single fiber optical coherence tomography microsurgical instruments for computer and robot-assisted retinal surgery," *Proc. 12th Annu. Conf. Med. Img. Comp. Aid. Int. (MICCAI),* pp. 108-115, 2009.

[8] F. Bandello, P. Lanzetta, and U. Menchini, "When and how to do a grid laser for diabetic macular edema," *Doc. Ophthalmologica,* vol. 97, pp. 415-419, 1999.

[9] J. Batlle, E. Mouaddib, and J. Salvi, "Recent progress in coded structured light as a technique to solve the correspondence problem: a survey," *Pattern Recognition,* vol. 31, p. 977, 1998.

[10] H. Bay, A. Ess, T. Tuytelaars, and L. van Gool, "SURF: Speeded Up Robust Features," *Proc. IEEE Computer Vision and Pattern Recognition,* vol. 110, pp. 346-359, 2008.

[11] B. Becker, R. MacLachlan, L. Lobes Jr, and C. Riviere, "Semiautomated intraocular laser surgery using handheld instruments," *Lasers Sur. Med.,* vol. 42, pp. 264-273, 2010.

[12] B. Becker, H. Tummala, and C. Riviere, "Autoregressive modeling of physiological tremor under microsurgical conditions," in *EMBC*, 2008, pp. 1948-1951.

[13] 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.*, Kobe, Japan, 2009, pp. 339 - 344.

[14] B. C. Becker, S. Voros, J. Louis A. Lobes, 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] F. Becquet, J. Le Rouic, X. Zanlonghi, P. Peronnet, E. Hermouet-Leclair, C. Pousset-Decre, and D. Ducournau, "Efficiency of surgical treatment for chronic macular edema due to branch retinal vein occlusion," *Journal français d'ophtalmologie,* vol. 26, p. 570, 2003.

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

[17] J. Binder and W. Kramer, "Robotically assisted laparoscopic radical prostatectomy," *BJU International,* vol. 87, pp. 408-410, 2001.

[18] M. S. Blumenkranz, D. Yellachich, D. E. Andersen, M. W. Wiltberger, D. Mordaunt, G. R. Marcellino, and D. Palanker, "Semiautomated pattern scanning laser for retinal photocoagulation," *Retina,* vol. 26, pp. 370-376, 2006.

[19] A. Bo, P. Poignet, and C. Geny, "Filtering voluntary motion for pathological tremor compensation," in *Proc. IEEE Intl. Conf. Int. Rob. and Sys.*, 2009, pp. 55-60.

[20] A. Bo, P. Poignet, F. Widjaja, and W. Ang, "Online pathological tremor characterization using extended Kalman filtering," in *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2008, pp. 1753-1756.

[21] H. L. Brooks, "Macular hole surgery with and without internal limiting membrane peeling," *Ophthalmology,* vol. 107, pp. 1939-1948, 2000.
[22] J. Bruce, T. Balch, and M. Veloso, "Fast and inexpensive color image segmentation for interactive robots," 2000.
[23] 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.
[24] C. Carswell, D. Clarke, and W. B. Seales, "Assessing mental workload during laparoscopic surgery," *Surgical Innovation,* vol. 12, pp. 80-90, 2005.
[25] X. Chen, G. Xu, J. Yong, G. Wang, and J. Paul, "Computing the minimum distance between a point and a clamped B-spline surface," *Graphical Models,* vol. 71, pp. 107-112, 2009.
[26] 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.
[27] X. Chen, Y. Zhou, Z. Shu, and H. Su, "Improved algebraic algorithm on point projection for Bezier curves," in *IMSCCS*, 2007, pp. 158-163.
[28] T. Choe, I. Cohen, and G. Medioni, "3-D shape reconstruction of retinal fundus," CVPR, 2006, pp. 2277-2284.
[29] J. Corso, N. Ramey, and G. Hager, "Stereo-Based Direct Surface Tracking with Deformable Parametric Models," Johns Hopkins University2003.
[30] J. Cuevas Tabarés, R. A. MacLachlan, C. A. Ettensohn, and C. N. Riviere, "Cell Micromanipulation with an Active Handheld Micromanipulator," in *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2010, pp. 4363-4366.
[31] 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.
[32] 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.
[33] M. Dewan, P. Marayong, A. Okamura, and G. Hager, "Vision-based assistance for ophthalmic micro-surgery," *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2004,* pp. 49-57, 2004.
[34] K. Dorfmüller-Ulhaas, "Robust optical user motion tracking using a Kalman filter," in *Proc. ACM Symp. on Virt. Real. Soft. Tech.*, 2003.
[35] M. Eck and H. Hoppe, "Automatic reconstruction of B-spline surfaces of arbitrary topological type," in *SIGGRAPH*, 1996, pp. 325-334.
[36] S. E. Everett, Y. Isoda, R. V. Dubey, and C. Dumont, "Vision-based end-effector alignment assistance for teleoperation," 1999, pp. 543-549 vol. 1.
[37] N. Feltgen, C. Auw-Haedrich, R. Buchen, and L. L. Hansen, "Arteriovenous dissection in a living human eye: clinicopathologic correlation," *Archives of ophthalmology,* vol. 123, p. 571, 2005.
[38] N. Feltgen, J. Herrmann, H. Agostini, A. Sammain, and L. L. Hansen, "Arterio-venous dissection after isovolaemic haemodilution in branch retinal vein occlusion: a non-randomised prospective study," *Graefe's archive for clinical and experimental ophthalmology,* vol. 244, pp. 829-835, 2006.
[39] N. Feltgen, B. Junker, H. Agostini, and L. L. Hansen, "Retinal Endovascular Lysis in Ischemic Central Retinal Vein Occlusion:: One-Year Results of a Pilot Study," *Ophthalmology,* vol. 114, pp. 716-723. e1, 2007.
[40] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM,* vol. 24, pp. 381-395, 1981.

[41] I. Fleming, M. Balicki, J. Koo, I. Iordachita, B. Mitchell, J. Handa, G. Hager, and R. Taylor, "Cooperative robot assistant for retinal microsurgery," *MICCAI 2008,* vol. 5242, pp. 543-550, 2008.
[42] R. N. Frank, "Retinal laser photocoagulation: benefits and risks," *Vision Research,* vol. 20, pp. 1073-1081, 1980.
[43] K. Frischmann, "Micromanipulator," US Patent, 1956.
[44] J. Funda, R. H. Taylor, B. Eldridge, S. Gomory, and K. G. Gruben, "Constrained Cartesian motion control for teleoperated surgical robots," *IEEE Transactions Robotics and Automation,* vol. 12, pp. 453-465, 1996.
[45] S. J. Garg and J. Vander, "Surgery for Retinal Venous Occlusions," *Review of Opthalmology,* 2003.
[46] M. R. Glucksberg, R. Dunn, and C. P. Giebs, "In vivo micropuncture of retinal vessels," *Graefe's archive for clinical and experimental ophthalmology,* vol. 231, pp. 405-407, 1993.
[47] B. Gupta, M. Elagouz, D. McHugh, V. Chong, and S. Sivaprasad, "Micropulse diode laser photocoagulation for central serous chorio-retinopathy," *Clinical & Experimental Ophthalmology,* vol. 37, pp. 801-805, 2009.
[48] J. V. Hajnal, D. J. Hawkes, and D. L. G. Hill, *Medical image registration*: CRC, 2001.
[49] C. Haritoglou, C. A. Gass, M. Schaumberger, A. Gandorfer, M. W. Ulbig, and A. Kampik, "Long-term follow-up after macular hole surgery with internal limiting membrane peeling," *American journal of ophthalmology,* vol. 134, pp. 661-666, 2002.
[50] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge, UK: Cambridge University Press, 2003.
[51] R. Harwell and R. Ferguson, "Physiologic tremor and microsurgery," *Microsurgery,* vol. 4, pp. 187-192, 1983.
[52] F. Holtz and R. F. Spaide, *Medical retina*: Springer, 2007.
[53] B. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. of Am. A,* vol. 4, pp. 629-642, 1987.
[54] I. W. Hunter, T. D. Doukoglou, S. R. Lafontaine, P. G. Charette, L. A. Jones, M. A. Sagar, G. D. Mallinson, and P. J. Hunter, "A teleoperated microsurgical robot and associated virtual environment for eye surgery," *Presence,* vol. 2, pp. 265-280, 1993.
[55] D. A. Infeld and J. G. O'Shea, "Diabetic retinopathy," *Postgrad. Med. J.,* vol. 74, pp. 129-133, 1998.
[56] P. S. Jensen, K. W. Grace, R. Attariwala, J. E. Colgate, and M. R. Glucksberg, "Toward robot-assisted vascular microsurgery in the retina," *Graefe's Archive for Clinical and Experimental Ophthalmology,* vol. 235, pp. 696-701, 1997.
[57] A. M. Joussen, T. W. Gardner, and B. Kirchhof, *Retinal vascular disease*: Springer Verlag, 2007.
[58] 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.
[59] D. Kasner, G. Miller, W. Taylor, R. Sever, and W. Norton, "Surgical treatment of amyloidosis of the vitreous," *Trans Am Acad Ophthal Otolaryngol,* vol. 72, p. 410, 1968.
[60] N. E. Kelly and R. T. Wendel, "Vitreous surgery for idiopathic macular holes: results of a pilot study," *Archives of ophthalmology,* vol. 109, p. 654, 1991.
[61] R. Kumar, G. Hager, A. Barnes, P. Jensen, and R. Taylor, "An augmentation system for fine manipulation," in *Medical Image Computing and Computer Assisted Intervention*, 2000.
[62] R. R. Lakhanpal, M. Javaheri, H. Ruiz-Garcia, and J. DE, "Transvitreal limited arteriovenous-crossing manipulation without vitrectomy for complicated branch retinal vein occlusion using 25-gauge instrumentation," *Retina,* vol. 25, p. 272, 2005.

[63] W. T. Latt, U. X. Tan, C. Y. Shee, C. N. Riviere, and W. T. Ang, "Compact Sensing Design of a Handheld Active Tremor Compensation Instrument," *IEEE Sensors Journal,* vol. 9, pp. 1864-1871, 2009.
[64] P. Leaver and C. Williams, "Argon laser photocoagulation in the treatment of central serous retinopathy," *The British Journal of Ophthalmology,* vol. 63, p. 674, 1979.
[65] T. Leng, J. M. Miller, K. V. Bilbao, D. V. Palanker, P. Huie, and M. S. Blumenkranz, "The chick chorioallantoic membrane as a model tissue for surgical retinal research and simulation," *Retina,* vol. 24, pp. 427-434, 2004.
[66] 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.
[67] M. Li, A. Kapoor, and R. Taylor, "Telerobotic control by virtual fixtures for surgical applications," *Advances in Telerobotics,* vol. 31, pp. 381-401, 2007.
[68] 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.
[69] D. Liu, N. Wood, X. Xu, N. Witt, A. Hughes, and T. SAMcG, "3D Reconstruction of the Retinal Arterial Tree Using Subject-Specific Fundus Images," *Advances in Computational Vision and Medical Image Processing,* pp. 187-201, 2009.
[70] X. Liu, L. Yang, J. Yong, H. Gu, and J. Sun, "A torus patch approximation approach for point projection on surfaces," *Computer Aided Geometric Design,* vol. 26, pp. 593-598, 2009.
[71] W. Ma and J. Kruth, "Parameterization of randomly measured points for least squares fitting of B-spline curves and surfaces," *Computer-Aided Design,* vol. 27, pp. 663-675, 1995.
[72] 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.
[73] R. Macherner, "The development of pars plana vitrectomy: a personal account," *Graefe's archive for clinical and experimental ophthalmology,* vol. 233, pp. 453-468, 1995.
[74] R. A. MacLachlan, B. C. Becker, J. Cuevas Tabarés, L. A. Lobes, Jr., and C. N. Riviere, "Micron: an actively stabilized handheld tool for microsurgery," *IEEE/ASME Transactions on Mechatronics,* submitted for publication.
[75] R. A. MacLachlan, B. C. Becker, J. C. Tabares, G. W. Podnar, J. Louis A. Lobes, and C. N. Riviere, "Micron: an Actively Stabalized Handheld Tool for Microsurgery," *IEEE Trans. Robot.,* submitted.
[76] 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.
[77] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical image analysis,* vol. 2, pp. 1-36, 1998.
[78] M. S. Markow, Y. Yang, A. J. Welch, H. G. Rylander, III, and W. S. Weinberg, "An automated laser system for eye surgery," *IEEE Eng Med Biol Mag,* vol. 8, pp. 24-29, 1989.
[79] M. Martínez-Pérez, A. Espinosa-Romero, U. IIMAS, M. FM, and M. UADY, "3D reconstruction of retinal blood vessels from two views," in *4th Indian Conference on Computer Vision and Graphics*, 2004, pp. 258-263.
[80] 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.
[81] D. Modi, P. Chiranand, and L. Akduman, "Efficacy of patterned scan laser in treatment of macular edema and retinal neovascularization," *Clin Ophthalmol.,* vol. 3, pp. 465–470, 2009.
[82] A. P. Mulgaonkar, J. P. Hubschman, J. L. Bourges, B. L. Jordan, C. Cham, J. T. Wilson, T. C. Tsao, and M. O. Culjat, "A Prototype Surgical Manipulator for Robotic Intraocular Micro Surgery," *Studies in Health Technology and Informatics,* vol. 142, pp. 215-217, 2009.

[83] B. Nilufer and B. Cosar, "Surgical treatment of central retinal vein occlusion," *Acta Ophthalmologica,* vol. 86, pp. 245-252, 2008.
[84] M. E. Opremcak and R. A. Bruce, "Surgical decompression of branch retinal vein occlusion via arteriovenous crossing sheathotomy: a prospective review of 15 cases," *Retina,* vol. 19, p. 1, 1999.
[85] M. D. Osterloh and S. Charles, "Surgical decompression of branch retinal vein occlusions," *Archives of ophthalmology,* vol. 106, pp. 1469-71, 1988.
[86] A. J. Packer, *Manual of Retinal Surgery*: Butterworth-Heinemann Medical, 2001.
[87] M. B. Parodi and F. Bandello, "Branch retinal vein occlusion: classification and treatment," *Ophthalmologica,* vol. 223, pp. 298-305, 2009.
[88] M. Patkin, "Erogonomics applied to the practice of microsurgery," *Australian and New Zealand Journal of Surgery,* vol. 47, pp. 320-329, 1977.
[89] Z. Pezzementi, S. Voros, and G. Hager, "Articulated object tracking by rendering consistent appearance parts," 2009, pp. 3940-3947.
[90] N. Ramey, J. Corso, W. Lau, D. Burschka, and G. Hager, "Real-time 3D surface tracking and its applications," in *CVPR*, 2005, p. 34.
[91] C. Reiley, H. Lin, D. Yuh, and G. Hager, "Review of methods for objective surgical skill evaluation," *Surgical Endoscopy,* pp. 1-11, 2010.
[92] T. Rice, "Internal limiting membrane removal in surgery for full-thickness macular holes," *Macular hole: pathogenesis, diagnosis, and treatment. Woburn, MA: Butterworth-Heinemann,* pp. 125–146, 1999.
[93] C. Riviere, R. Rader, and N. Thakor, "Adaptive cancelling of physiological tremor for improved precision in microsurgery," *IEEE Trans. on Biomedical Engineering,* vol. 45, pp. 839-846, 1998.
[94] 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.
[95] C. N. Riviere, J. Gangloff, and M. de Mathelin, "Robotic compensation of biological motion to enhance surgical accuracy," *Proc. of the IEEE,* vol. 94, pp. 1705-1716, 2006.
[96] C. N. Riviere and N. V. Thakor, "Modeling and canceling tremor in human-machine interfaces," *IEEE Eng. Med. Biol. Mag,* vol. 15, pp. 29-36, 1996.
[97] S. Rizzo, F. Patelli, and D. Chow, *Vitreo-Retinal Surgery: Progress III*: Springer, 2008.
[98] L. B. Rosenberg, "The use of virtual fixtures as perceptual overlays to enhance operator performance in remote environments," USAF Armstrong Laboratory1993.
[99] L. B. Rosenberg, "Virtual fixtures: perceptual tools for telerobotic manipulation," in *IEEE Virtual Reality Annual International Symposium*, 1993, pp. 76-82.
[100] R. Sandra, W. Joan, and O. Semin, "Review of the ocular angiogenesis animal models," *Semin Ophthalmol,* vol. 24, pp. 52-61, 2009.
[101] C. Sanghvi, R. McLauchlan, C. Delgado, L. Young, S. J. Charles, G. Marcellino, and P. E. Stanga, "Initial experience with the Pascal photocoagulator: a pilot study of 75 procedures," *British Journal of Ophthalmology,* vol. 92, p. 1061, 2008.
[102] D. Scharstein and R. Szeliski, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," *International Journal of Computer Vision,* vol. 47, pp. 7-42, 2002.
[103] I. U. Scott, M. S. Ip, P. C. VanVeldhuisen, N. L. Oden, B. A. Blodi, M. Fisher, C. K. Chan, V. H. Gonzalez, L. J. Singerman, and M. Tolentino, "A randomized trial comparing the efficacy and safety of intravitreal triamcinolone with observation to treat vision loss associated with macular edema secondary to central retinal vein occlusion: the Standard Care vs Corticosteroid for Retinal Vein Occlusion (SCORE) study report 6," *Archives of ophthalmology,* vol. 127, pp. 1115-28, 2009.

[104] C. P. Shah, "A randomized trial comparing intravitreal triamcinolone acetonide and focal/grid photocoagulation for diabetic macular edema," *Evidence-Based Ophthalmology,* vol. 10, p. 29, 2009.
[105] S. P. N. Singh and C. N. Riviere, "Physiological tremor amplitude during retinal microsurgery," in *Proc. Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2002, pp. 171-172.
[106] W. E. Smiddy, W. Feuer, and G. Cordahi, "Internal limiting membrane peeling in macular hole surgery," *Ophthalmology,* vol. 108, pp. 1471-1476, 2001.
[107] R. Sophie, L. M. Rachel, C. Ning, L. Lyndell, W. Jie Jin, M. Paul, W. K. Jonathan, N. Hiep, and Y. W. Tien, "The prevalence of retinal vein occlusion: pooled data from population studies from the United States, Europe, Asia, and Australia," *Ophthalmology,* vol. 117, pp. 313-319.e1, 2010.
[108] W. M. Tang and D. P. Han, "A study of surgical approaches to retinal vascular occlusions," *Archives of Ophthalmology,* vol. 118, pp. 138-143, 2000.
[109] M. K. Tsilimbaris, E. S. Lit, and D. J. D'Amico, "Retinal microvascular surgery: a feasibility study," *Invest Ophthalmol Vis Sci* vol. 45, pp. 1963-68, 2004.
[110] 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.
[111] 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 *3rd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics*, pp. 814-819.
[112] K. Veluvolu, W. Latt, and W. Ang, "Double adaptive bandlimited multiple Fourier linear combiner for real-time estimation/filtering of physiological tremor," *J of Bio. Sig. Proc. Con.,* vol. 5, pp. 37-44, 2010.
[113] W. Wang, H. Pottmann, and Y. Liu, "Fitting B-spline curves to point clouds by curvature-based squared distance minimization," *ACM Transactions on Graphics (TOG),* vol. 25, pp. 214-238, 2006.
[114] D. Weinberg, D. Dodwell, and S. Fern, "Anatomy of arteriovenous crossings in branch retinal vein occlusion," *American journal of ophthalmology,* vol. 109, pp. 298-302, 1990.
[115] J. N. Weiss, "Retinal surgery for treatment of central retinal vein occlusion," *Ophthalmic surgery and lasers,* vol. 31, pp. 162-165, 2000.
[116] J. N. Weiss, "Apparatus and method for cannulating retinal blood vessels," USA Patent 6,402,734, 2002.
[117] 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.
[118] M.-H. Yang and N. Ahuja, "Gaussian mixture model for human skin color and its applications in image and video databases," *Proc. SPIE,* vol. 3656, pp. 458-466, 1998.
[119] J. Zhang and F. Chu, "Real-time modeling and prediction of physiological hand tremor," in *ICASSP*, 2005.
[120] J. Zhao, S. M. Sastry, R. D. Sperduto, E. Chew, and N. Remaley, "Arteriovenous crossing patterns in branch retinal vein occlusion," *Ophthalmology,* vol. 100, p. 423, 1993.
[121] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing,* vol. 21, pp. 977-1000, 2003.