# Toward Monocular Camera-Guided Retinal Vein Cannulation with an Actively Stabilized Handheld Robot

Shohin Mukherjee[1], Sungwook Yang[2], Robert A. MacLachlan[1], Louis A. Lobes, Jr.[3], Joseph N. Martel[3] and Cameron N. Riviere[1]

***Abstract*— In this paper we describe work towards retinal vessel cannulation using an actively stabilized handheld robot, guided by monocular vision. We employ a previously developed monocular camera based surface reconstruction method using automated laser beam scanning over the retina. We use the reconstructed plane to find a coordinate transform between the 2D image plane coordinate system and the global 3D frame. Within a hemispherical region around the target, we use motion scaling for higher precision. The contribution of this work is the homography matrix estimation using monocular vision and application of the previously developed laser surface reconstruction to Micron guided vein cannulation. Experiments are conducted in a wet eye phantom to show the higher accuracy of the surface reconstruction as compared to standard stereo reconstruction. Further, experiments to show the increased surgical accuracy due to motion scaling are also carried out.**

## I. INTRODUCTION

Retinal cannulation refers to the direct delivery of drugs into the veins of the retina. Cannulation has the potential to be useful in the treatment of diseases such as retinal vein occlusion (RVO), which occurs when a clot obstructs blood flow in a central (CRVO) or branch (BRVO) vein of the eye. RVO is the second most common retinal vascular disease after diabetic retinopathy and affects an estimated 16.4 million adults worldwide [1]. Though RVO has no proven effective treatment, a promising experimental procedure called retinal endovascular surgery (REVS) exists. REVS involves cannulating the vein and directly injecting clot-dissolving plasminogen activator (t-PA) [2].

There have been attempts to employ robot-aided micromanipulation to retinal surgery. The Steady Hand robot developed at JHU receives input in the form of force/torque and allows voluntary motion while suppressing tremor through selective compliance [3], [4]. However, it suffers from the drawback that it cannot handle large movements of the patient's eye, which are common in ophthalmic surgery since the patient is usually sedated rather than anesthetized. In [5], an apparatus consisting of a micropipette, micromanipulator and positioner mounted to a base attached to a wrist rest was used to cannulate a branch retinal vein in order to treat CRVO [6].

Micron is an actively stabilized handheld surgical robot, that compensates for the surgeon's physiological hand tremor. The current version of Micron, as introduced in [7] is a 6 DOF system with a Gough-Stewart platform. The actuators are piezoelectric linear motors. The device has two sets of infrared LEDs, three of which are mounted on the moving platform and three are fixed to the handle. The LEDs are optically tracked by a custom built tracking system: ASAP(Apparatus to Sense Accuracy of Position), which uses two position sensitive detectors to triangulate the position of each LED in space [8]. Figure 1 shows the exploded view of the latest version of Micron.

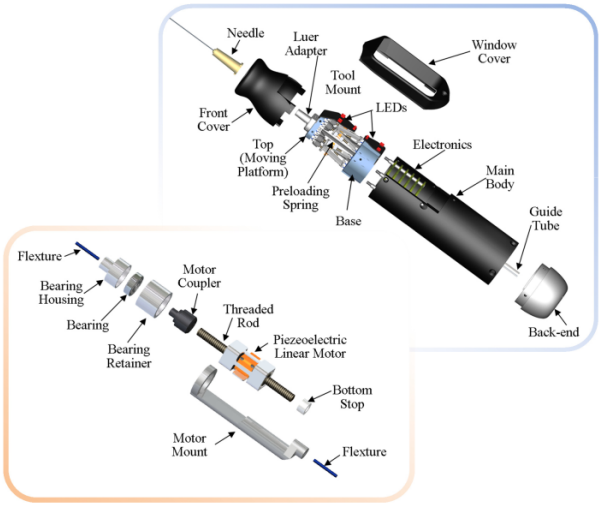


Fig. 1. Exploded view of Micron.

In [9] vessel cannulation using Micron was attempted. However, the system suffered from several challenges due to which its application in an intact eye was not feasible. For example, the micropipette tip was painted for tracking which was not robust. However the primary shortcoming of the system was that it used stereo vision to reconstruct the retinal surface and to convert the image space coordinates into the global ASAP coordinates. Stereo reconstruction does not work in an intact eye during vitreoretinal surgery because of the nonlinear optics of the eye. The optical path includes the cornea, the lens and saline solution introduced during vitrectomy. The eye is also covered with a contact lens to provide a wide-angle view during operation. All this

*This work was supported in part by the National Institutes of Health (grant no. R01 EB000526).

S. Mukherjee, R. A. MacLachlan and C. N. Riviere are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: shohin@cmu.edu; robmacl@ieee.org; camr@ri.cmu.edu).

S. Yang is with the Center for BioMicrosystems, Korea Institute of Science and Technology, Seoul 136-791, Korea (e-mail: swyang@kist.re.kr)

L. A. Lobes, Jr. and J. N. Martel are with the Department of Ophthalmology, University of Pittsburgh, Pittsburgh, PA 15213 USA (e-mail: lalobes@gmail.com; marteljn@upmc.edu).

leads to considerable optical distortion. The accuracy of stereo reconstruction depends greatly on the accuracy of the camera calibration. Most camera calibration methods assume a classical perspective camera model in a single medium as air, an assumption which does not hold in an intact eye.

In [10], a new method for retinal surface reconstruction using monocular vision was introduced. It relied on automatic circular scanning of the retinal surface using a surgical laser attached to Micron and geometric analysis of the projected elliptical beam pattern. The method was tested in various conditions and its feasibility was established in a realistic eye model which incorporated optical distortion by lenses. The reconstruction was then used to carry out laser photocoagulation with Micron. The 3-DOF motion of the tool tip was decoupled into the 2-DOF planar motion parallel to the retinal surface and the 1-DOF motion along the axis of the tool. The decoupled 2-DOF motion was then controlled via image-based servoing, to place the laser beam onto a target position using a monocular camera. The 1-DOF axial motion was controlled to maintain a constant standoff distance from the estimated surface. However, visual servoing is challenging in cannulation. Instead of tracking a laser spot, a needle tip needs to be tracked which is unreliable in a real eye. There have been some tool tracking algorithms developed [11]–[13]. However, the cannulation needle is very thin and accurate detection of the tool tip is hard. In this paper, we apply the new surface reconstruction technique to develop a vessel cannulation technique that can be used in an intact eye.

## II. Methods

### A. Tool and eye phantom

In order to use the surface reconstruction using laser as described in II-C, it was essential to have a tool that has both the laser probe and the cannulation needle. The eye phantom consisted of a network of blood vessels printed on a surface with water on top, to serve as saline that replaces the vitreous humor during vitrectomy. The tool and the phantom are shown in Figure 2.

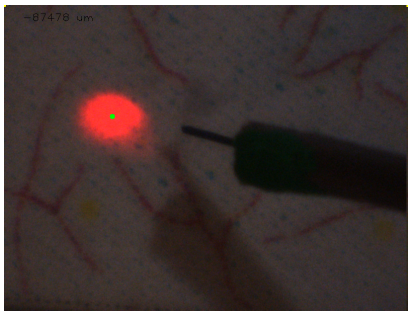

Fig. 2. Tool and eye phantom used.

### B. Tool tip calibration

The surface reconstruction described in [10] is very sensitive to the tool tip offset, which is the offset of the tool tip from the origin of Micron's coordinate system. In order to calibrate the tool tip offset, we use an approach developed in [14]. An optimizer is used to minimize the error between the projected and detected tip positions in both the cameras, while scanning the tool over the surface. The projected tip positions are obtained by multiplying the camera projection matrices with the 3D ASAP coordinates of the tool tip. The camera projection matrices are simultaneously calibrated during the optimization. The tip offset is modeled in terms of angular variations $\theta_x$ and $\theta_y$, and the length of the tool $l_0$, from the origin of Micron's coordinate system.. The optimization minimized the following objective function:

$$\min \sum_{i=1}^{n} [(p^i_{ProjL} - p^i_{DetL})^2 + (p^i_{ProjR} - p^i_{DetR})^2] \quad (1)$$

where $p^i_{ProjL}$ and $p^i_{ProjR}$ are the projected tip positions for the left and right cameras and $p^i_{DetL}$ and $p^i_{DetR}$ are the detected tip positions in the left and right cameras.

### C. Surface reconstruction

Since the control of Micron happens in the global coordinates of ASAP, there is a need to model the retina in the ASAP coordinates. Due to the unreliability of stereo reconstruction in an intact eye [9], Yang et al. [10] developed a surface reconstruction method, which uses a surgical laser scanning its aiming beam over the retinal surface. The region of interest of the retina is approximated as a plane. A circular scan is carried out by Micron over the retina, around a pivot point. The pivot point is the RCM (remote center of motion) in vitreoretinal surgery. The RCM in vitreoretinal surgery is the point of insertion of the tool into the eyeball. The circular pattern projects an ellipse on the plane to be estimated. The beam locations on the image plane are detected using a beam detection algorithm, which uses a sequence of image processing techniques like thresholding, morphological operations and contour detection. An ellipse is fitted through these beam points after outliers have been removed using RANSAC. The plane to be estimated can be described by tilt of the plane initially perpendicular to the axis of the cone. From Figure 3, $\theta_{plane}$ and $d_{plane}$ are estimated by (2) and (5). In the following equations, $m_a$ and $m_b$ are the major and minor axis lengths respectively of the fitted ellipse and $\boldsymbol{v_{cone}}$ is the vector along the axis of the cone. $s_{cam}$ is the image scale in $\mu m/pixel$ .

$$\theta_{plane} = \pm \arcsin \sqrt{\frac{1 - \frac{1}{\gamma^2}}{1 + \frac{1}{\tan^2 \theta_{cone}}}} \quad (2)$$

$$\gamma = \frac{m_b}{m_a} \quad (3)$$

$$\tan \theta_{cone} = \frac{h_{RCM}}{r_{cone}} \quad (4)$$

$$d_{plane} = \frac{\cos \theta_{plane} (\tan^2 \theta_{cone} - \tan^2 \theta_{plane})}{\tan \theta_{cone}} m_a s_{cam} \quad (5)$$

$$\boldsymbol{P_{plane}} = \boldsymbol{P_{RCM}} - d_{plane}\boldsymbol{v_{cone}} \tag{6}$$

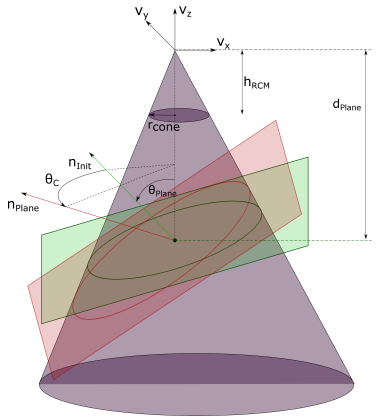


Fig. 3. Geometric analysis and parameters of ellipse formed on the plane from the circular laser scan [10].

The angle $\theta_{plane}$ does not uniquely describe the tilted plane, since any arbitrary vector orthogonal to the axis of the cone can be taken as the axis of rotation. Thus, rotating the true normal about the axis of the cone by any arbitrary angle leads to an ellipse of the same dimensions. Here, by true normal we mean the normal of the surface we are trying to estimate.

An initial normal vector, as explained in [10], can be set by rotating the axis of the cone $v_z$ about the y-axis of the cone coordinates by $\theta_{plane}$.

$$\boldsymbol{n_{init}} = \boldsymbol{R}(\boldsymbol{v_y}, \theta_{plane})\boldsymbol{v_z} \tag{7}$$

The true surface normal can be considered as a rotation of an initial normal vector $n_{init}$ about the axis of the cone $v_z$ by an angle $\theta_c$. In order to determine $\theta_c$, an optimizer is used to match the relative orientation of the 2D beam positions in the image with the 3D beam positions obtained by finding the intersection of the ray from the RCM to tip position of Micron with the estimated plane.

Given $k$ points on the trajectory, the objective function in (8) was defined as in [10].

$$F = \min_{\theta_c \in [-\pi, +\pi]} \sum_{i=1}^{k} (\theta^i_{beam} - \theta^i_{proj})^2. \tag{8}$$

where the phase angles of the $i$th data point with respect to the major axis of the ellipse obtained from the beam trajectory on the image and the one obtained from the projected points on the estimated plane are $\theta^i_{beam}$ and $\theta^i_{proj}$ respectively. Once we have the $\theta_c$ that minimizes the given objective function, we can find the true normal by (9).

$$\boldsymbol{n_{plane}} = \boldsymbol{R}(\boldsymbol{v_z}, \theta_c)\boldsymbol{n_{init}} \tag{9}$$

The plane point $P_{plane}$ given by (6) and the plane normal given by (9) completely define the reconstructed retinal plane.

### D. Homography matrix estimation

In order to convert a cannulation target in the image space to the ASAP coordinate plane, we utilize the planar assumption of the region of interest. We then find a homography matrix between the retina and the image plane. This is done by obtaining a Direct Linear Transform (DLT) [15] between the detected 2D beam positions in the image and the corresponding expected 3D beam positions in the ASAP frame, obtained by finding the intersection of the ray from the RCM to the tip, with the estimated plane from the cone beam reconstruction. These are the same correspondences that were used in the surface reconstruction. Since the distance between the microscope and the retina is large compared to the dimensions of the region of interest, the angle of incidence for all points on the region of interest is approximately the same. Therefore, a linear relationship exists between the two coordinate systems as long as the optics of the medium remain constant. If the surface is reconstructed just before the cannulation operation, this approximation can be made.

Therefore,

$$\begin{bmatrix} X_{ASAP} \\ Y_{ASAP} \end{bmatrix} = \boldsymbol{H} \begin{bmatrix} x_{camera} \\ y_{camera} \end{bmatrix} \tag{10}$$

$$Z_{ASAP} = (-aX_{ASAP} - bY_{ASAP} - 1)/c \tag{11}$$

where $[X_{ASAP}, Y_{ASAP}, Z_{ASAP}]$ is the 3D world coordinate, $[x_{camera}, y_{camera}]$ is the 2D image space pixel coordinate, $[a, b, c]$ are the geometric plane parameters and $H$ is the homography matrix.

### E. Surface tracking with EyeSLAM

During vitreoretinal surgery, the surgeon manipulates the eye with a surgical tool to explore the region of interest. The patient may also move the eye. Because of this voluntary and involuntary motion of the eyeball, there is movement in the microscopic view of the retina. Therefore, any goal position defined in the image coordinates needs to be updated to account for this movement. For this purpose, Becker et. al developed EyeSLAM - a real time vessel mapping and localization algorithm [16] [17]. We use EyeSLAM to track the 2D motion described by 3 degrees of freedom ($[t_x, t_y, \theta]$) of the retinal plane and update the image-space goal position with a rigid transformation. This updated goal position is then mapped to ASAP coordinates using the homography matrix. Figure 4 shows EyeSLAM working in a wet eye phantom. EyeSLAM always runs in the background to track eyeball motion. The cone beam surface reconstruction and the homography matrix calculation is done once, after which motion scaling is engaged around the target point in the ASAP coordinates obtained by converting the image space target with the computed homography matrix.

### F. Motion scaling

Once the goal position for cannulation has been established in the ASAP coordinates, we define a hemispherical region around it. When the tool tip is outside this region, Micron operates in the normal tremor-cancellation mode

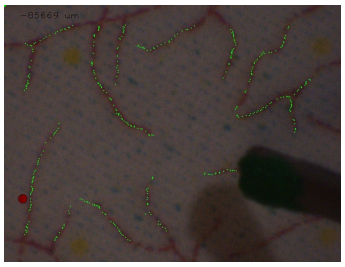

Fig. 4. Vessel tracking with EyeSLAM in a wet eye phantom. The green dots correspond to the current location of vessels.

which employs a low-pass filter that serves to attenuate high-frequency motions due to hand tremor [18]. Inside this region, motion scaling is engaged.

$$\boldsymbol{P_T} = \boldsymbol{P_R} + \lambda(\boldsymbol{F_L}(\boldsymbol{P_N}) - \boldsymbol{P_R}) \tag{12}$$

In (12), $\boldsymbol{P_T}$ is the scaled goal position being sent to Micron's position controller, $\boldsymbol{F_L}(\boldsymbol{P_N})$ is the tremor-compensated null position, $\boldsymbol{P_R}$ is a reference point and $\lambda$ is the scale factor. Two choices for the reference point are logical. One possible choice is the point of entry of the tool tip into the region. In this case, there is a smooth transition from tremor cancellation mode to motion scaling mode. However, while exiting the region, if the point of exit is not the same as the point of entry, there is a jerk during the transition. A more obvious choice is to choose the intended cannulation goal position as $P_R$. However, in case of a constant scaling factor, there is a jerk during transition while both entering and exiting the motion scaling region. This is because there is a sudden transition from the scaled goal position to the null position of the manipulator, which is the default goal position in the absence of motion scaling.

In [19], in order to carry out a graceful transition between motion scaling and tremor compensation, the scale was decreased exponentially over time, which also provided robustness to intermittently noisy distance measurements due to unreliable tip tracking. However, since all our control happens in the much more reliable ASAP frame, and distances are measured in the global ASAP coordinates instead of the image coordinates, we simply change the scale factor linearly as a function of distance from the target, as given by (13).

$$\lambda(r) = \frac{1-\lambda_0}{R} r + \lambda_0 \tag{13}$$

$R$ is the distance at which motion scaling is engaged and $r$ is the current distance from the target. A scale factor of 0 corresponds to a hard virtual fixture and a scale factor of 1 corresponds to normal tremor compensation and no motion scaling. $\lambda_0$ corresponds to the maximum motion scaling at $r = 0$.

## III. Results

### A. Accuracy of coordinate transform

In order to measure the accuracy of the surface reconstruction and the image space to ASAP coordinate transform with the homography matrix, a grid of targets was defined on the retinal plane in the image space. The Micron tip was then commanded to servo to each of these targets, while maintaining a clearance of 500 $\mu$m in the cone beam case and 1500 $\mu$m in the stereo case. It was observed that without a large enough clearance, the needle tip would hit the surface of the retina phantom, due to the extremely inaccurate stereo reconstruction. The cone beam reconstruction is much more accurate, therefore a smaller clearance was sufficient for the experiments. A tip-detection algorithm was used to detect the final tip position in the image. Since there was no ground truth available for the position of the tool tip with respect to the surface along the surface normal dimension, the error was defined in terms of 2D pixel distance between the detected tip position and the target point. Figure 5 shows the target points and the detected tip positions for the cone beam reconstruction and stereo reconstruction respectively in a dry eye phantom. Figure 7 shows the plot of pixel error with the corresponding target. Figures 6 and 8 show the same for a wet eye phantom.

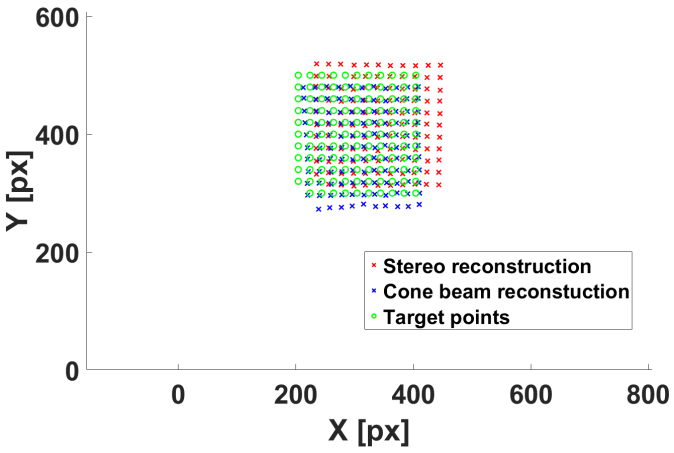


Fig. 5. Map of target points for cannulation and the corresponding tool tip location in image coordinates for cone beam reconstruction and stereo reconstruction, in a dry eye phantom.

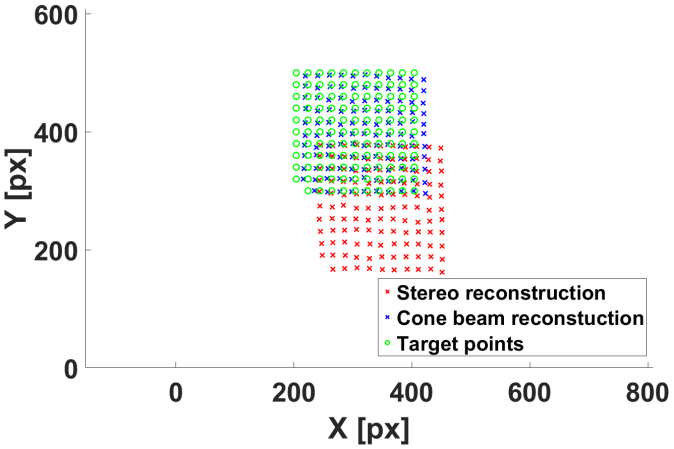


Fig. 6. Map of target points for cannulation and the corresponding tool tip location in image coordinates for cone beam reconstruction and stereo reconstruction, in a wet eye phantom.

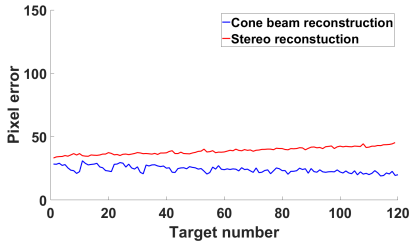


Fig. 7. 2D Pixel error between target point and detected tip position for cone beam reconstruction and stereo reconstruction, in a dry eye phantom.

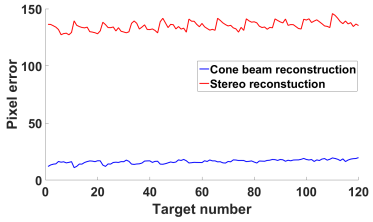


Fig. 8. 2D Pixel error between target point and detected tip position for cone beam reconstruction and stereo reconstruction, in a wet eye phantom.

The accuracy of the surface reconstruction was compared to a standard stereo algorithm which was impelemented as follows. The projection matrices for the cameras were calibrated as in [15]. The laser probe was swept over the retinal surface while detecting the 2D tip positions $p_c^i \in \mathbb{R}^2, \forall c \in \{L, R\}$ in the left and right images, which were matched with the corresponding 3D ASAP tip positions $P_{ASAP}^i \in \mathbb{R}^3$. Thus we obtain camera projection matrices $M_c \in \mathbb{R}^{3X4}, \forall c \in \{L, R\}$ using DLT, which are then used for triangulating a goal position in 3D.

$$\boldsymbol{p_{tip} = M_c P_{tip}^{ASAP}} \tag{14}$$

TABLE I

AVERAGE PIXEL ERROR.

| | Stereo | Cone beam |
|---|---|---|
| Air | 38.9 px | 24.1 px |
| Water | 135.2 px | 16.3 px |

TABLE II

STANDARD DEVIATION OF PIXEL ERROR.

| | Stereo | Cone beam |
|---|---|---|
| Air | 2.8 px | 2.6 px |
| Water | 3.9 px | 1.7 px |

Table I shows the average pixel error between the target and the detected tool tip. It was observed that the average pixel error between the target and the detected tip positions was similar in a dry eye phantom for the two cases. However in a wet eye phantom, the pixel error in the cone beam reconstruction case was about 12% of the error in the stereo reconstruction case.

### *B. Motion scaling*

In order to demonstrate the automatic activation of motion scaling close to the target, the approach of the vessel during cannulation was simulated in the wet eye phantom. The cannulation needle was attached to the end of the laser probe, and the tool tip calibration as explained in II-B was re-run to account for the needle length. The motion was scaled about the target point and the scale was linearly decreased from 1 to $\lambda_0 = 0.3$ as the distance between the target and the needle tip decreased from R = 4 mm to 0 mm. Figures 9, 10 and 11 show the $X$,$Y$ and $Z$ positions of the tool tip in the ASAP coordinates for three conditions: no tremor cancellation, normal tremor cancellation and motion scaling with tremor cancellation.

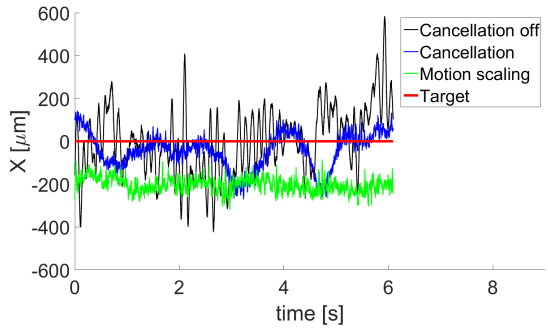


Fig. 9. $X$ position of the needle tip with time while holding steady at the cannulation target in the wet eye phantom, in ASAP coordinates.

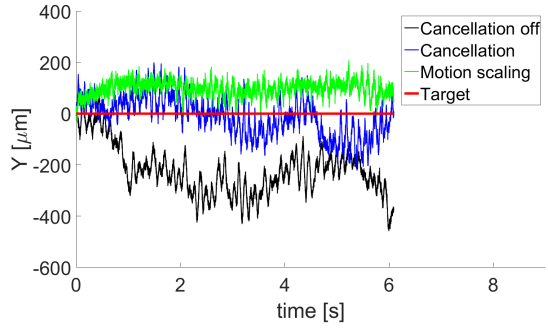


Fig. 10. $Y$ position of the needle tip with time while holding steady at the cannulation target in the wet eye phantom, in ASAP coordinates.

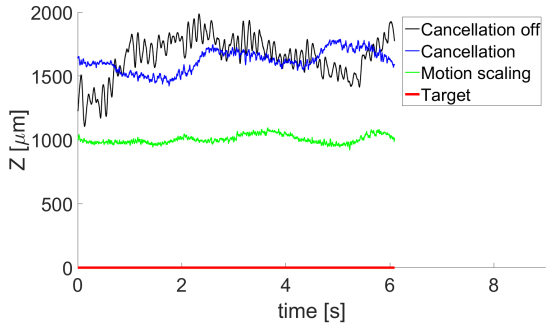


Fig. 11. $Z$ position of the needle tip with time while holding steady at the cannulation target in the wet eye phantom, in ASAP coordinates.

TABLE III

MEAN ERROR AND STANDARD DEVIATION OF ERRORS IN ASAP COORDINATES.

| | $\mu_{net}$ [mm] | $\sigma_{net}$ [mm] | $\sigma_x$ [mm] | $\sigma_y$ [mm] | $\sigma_z$ [mm] |
|---|---|---|---|---|---|
| Cancellation off | 1.68 | 0.176 | 0.153 | 0.097 | 0.165 |
| Cancellation on | 1.62 | 0.086 | 0.082 | 0.074 | 0.085 |
| Motion scaling | 1.03 | 0.031 | 0.034 | 0.030 | 0.033 |

Table III shows the mean and standard deviation of 3D errors measured in ASAP coordinates. Error standard deviation is the most meaningful metric here. Mean error or bias is primarily a function of the operator's visual perception, and is not affected by the technique presented here, except to the extent that cancellation of higher frequency error may assist the operator to correctly perceive mean positions. With motion scaling engaged, the standard deviation in error was about 36% of the standard deviation with pure cancellation and 18% of the error with no cancellation. The duration for the experiments was chosen based on our prior experience of using Micron in simulated surgery settings [18] and has no influence on the proposed method.

## IV. DISCUSSION

In this work, image guided retinal vessel cannulation was proposed. Challenges faced in our previous attempts at cannulation were successfully solved using monocular vision based surface reconstruction and subsequent coordinate mapping. The error in the coordinate mapping was about 12% of the error from a standard stereo-based triangulation approach, in a realistic wet eye phantom. The automated motion scaling methodology that relies on this coordinate mapping led to increased precision as indicated by the 64% reduction in the standard deviation of error, compared to pure tremor cancellation. As can be seen in Figures 9-11, the high-frequency oscillations due to hand tremor are clearly large when tremor compensation is turned off. It can also be seen that the low-frequency oscillations are much lesser in amplitude when motion scaling is active as compared to filter-based tremor compensation. This is expected since any high-frequency deviation which passes through the low-pass shelving filter is scaled down. Importantly, as evident from Figure 11, the error in the $Z$ axis, which was roughly perpendicular to the plane, was less in the motion-scaling case, showing that it was easier to keep the cannulation needle closer to the target without hitting the retinal plane when motion scaling was engaged. Retinal vessels are around $100\mu m$ in diameter. As reported in Table III., the standard deviation in error in our approach is lesser than $100\mu m$. In contrast, in the unaided case, the error standard deviation is over $100\mu m$. This clearly establishes the advantage of using the methods developed here for retinal vessel cannulation.

Future work will include testing the developed technique in varying lighting conditions with different phantoms. It will also include implementation with a microneedle-equipped end-effector, and testing in animal retinal tissue ex vivo and in vivo.

## REFERENCES


[1] S. Rogers, R. L. McIntosh, N. Cheung, L. Lim, J. J. Wang, P. Mitchell, J. W. Kowalski, H. Nguyen, T. Y. Wong, I. E. D. Consortium *et al.*, "The prevalence of retinal vein occlusion: pooled data from population studies from the United States, Europe, Asia, and Australia," *Ophthalmology*, vol. 117, no. 2, pp. 313–319, 2010.

[2] 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, no. 5, pp. 625–632, 2005.

[3] R. Taylor, P. Jensen, L. Whitcomb, A. Barnes, R. Kumar, D. Stoianovici, P. Gupta, Z. Wang, E. deJuan, and L. Kavoussi, "A steady-hand robotic system for microsurgical augmentation," *International Journal of Robotics Research*, vol. 18, no. 12, pp. 1201–1210, 1999.

[4] B. Mitchell, J. Koo, I. 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 *IEEE International Conference on Robotics and Automation*, 2007, pp. 623–629.

[5] 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, no. 12, pp. 2249 – 2257, 2001.

[6] J. N. Weiss, "Apparatus and method for cannulating retinal blood vessels," Jul. 12 2005, US Patent 6,916,000.

[7] S. Yang, R. A. MacLachlan, and C. N. Riviere, "Manipulator design and operation of a six-degree-of-freedom handheld tremor-canceling microsurgical instrument," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 2, pp. 761–772, 2015.

[8] R. A. MacLachlan and C. N. Riviere, "High-speed microscale optical tracking using digital frequency-domain multiplexing," *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 6, pp. 1991–2001, 2009.

[9] B. C. Becker, R. A. MacLachlan, L. A. Lobes, G. D. Hager, and C. N. Riviere, "Vision-based control of a handheld surgical micromanipulator with virtual fixtures," *IEEE Transactions on Robotics*, vol. 29, no. 3, pp. 674–683, June 2013.

[10] S. Yang, "Handheld micromanipulator for robot-assisted microsurgery," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2015.

[11] R. Richa, M. Balicki, E. Meisner, R. Sznitman, R. Taylor, and G. Hager, "Visual tracking of surgical tools for proximity detection in retinal surgery," in *International Conference on Information Processing in Computer-Assisted Interventions*. Springer, 2011, pp. 55–66.

[12] R. Richa, M. Balicki, R. Sznitman, E. Meisner, R. Taylor, and G. Hager, "Vision-based proximity detection in retinal surgery," *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 8, pp. 2291–2301, 2012.

[13] R. Sznitman, R. Richa, R. H. Taylor, B. Jedynak, and G. D. Hager, "Unified detection and tracking of instruments during retinal microsurgery," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 5, pp. 1263–1273, 2013.

[14] S. Yang, R. A. MacLachlan, and C. N. Riviere, "Toward automated intraocular laser surgery using a handheld micromanipulator," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sept 2014, pp. 1302–1307.

[15] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.

[16] B. C. Becker and C. N. Riviere, "Real-time retinal vessel mapping and localization for intraocular surgery," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 5360–5365.

[17] D. Braun, S. Yang, J. N. Martel, C. N. Riviere, and B. C. Becker, "Real-time localization and mapping of retinal vessels during intraocular microsurgery," *Int. J. Med. Robot.,* submitted 2016.

[18] R. A. MacLachlan, B. C. Becker, J. C. Tabarés, G. W. Podnar, L. A. L. Jr., and C. N. Riviere, "Micron: An actively stabilized handheld tool for microsurgery," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 195–212, Feb 2012.

[19] B. Becker, "Vision-based control of a handheld micromanipulator for robot-assisted retinal surgery," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, September 2012.