Real-Time Retinal Vessel Mapping and Localization for Intraocular Surgery

Brian C. Becker and **Cameron N. Riviere**
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA

Abstract

Computer-aided intraocular surgery requires precise, real-time knowledge of the vasculature during retinal procedures such as laser photocoagulation or vessel cannulation. Because vitreoretinal surgeons manipulate retinal structures on the back of the eye through ports in the sclera, voluntary and involuntary tool motion rotates the eye in the socket and causes movement to the microscope view of the retina. The dynamic nature of the surgical workspace during intraocular surgery makes mapping, tracking, and localizing vasculature in real time a challenge. We present an approach that both maps and localizes retinal vessels by temporally fusing and registering individual-frame vessel detections. On video of porcine and human retina, we demonstrate real-time performance, rapid convergence, and robustness to variable illumination and tool occlusion.

I. Introduction

VɪTREORETINAL surgery is often regarded as particularly demanding due to the extraordinary precision required to manipulate the small, delicate retinal structures, the confounding influence of physiological tremor on the surgeon's micromanipulation ability, and the challenging nature of the surgical access [1], [2]. Routine procedures such as membrane peeling require the surgeon to manipulate anatomy less than 10 μm thick [3–5], and laser photocoagulation operations benefit from accurate placement of laser burns [6], [7]. Promising new procedures such as vessel cannulation necessitate precise and exacting micromanipulation to inject anticoagulants into veins less than 100 μm in diameter [8], [9].

To address micromanipulation challenges in retinal surgical procedures, a variety of assistive robots have been proposed. Master/slave robots developed for eye surgery include the JPL Robot Assisted MicroSurgery (RAMS) system [10], the Japanese ocular robot of Ueta et al. [11], and the multi-arm stabilizing micromanipulator of Wei et al. [12]. Retinal surgery with the da Vinci master/slave robot has been investigated [13] and led to the design of a Hexapod micropositioner accessory for the da Vinci end-effector [14]. The Johns Hopkins SteadyHand Eye-Robot [15] shares control with an operator who applies force to the instrument while it is simultaneously held by the robot arm. A unique MEMS pneumatic actuator called the Microhand allows grasping and manipulation of the retina [16]. The Microbots of Dogangil et al. aim to deliver drugs directly to the retinal vasculature via magnetic navigation [17]. A lightweight micromanipulator developed in our lab, Micron, is fully handheld and has been used with vision-based control to aid retinal surgical procedures [18], [19].

While classic robot control can provide general behaviors such as motion scaling, velocity limiting, and force regulation, more specific and intelligent behaviors require knowledge of the anatomy. Vision-based control combines visual information of the anatomy with robotic control to enforce tip constraints, or virtual fixtures, which enact task-specific behaviors and provide guidance to the surgeon during procedure [20], [21]. In retinal vessel cannulation, knowledge of the vessel location relative to the instrument tip can aid robotic behavior and more effectively help guide the robot during injections into the vessel [22], [23]. During retinal laser photocoagulation, placing bums on retinal vessels should be avoided as this can occlude the vein, possibly causing vitreous hemorrhage [7]. However, existing methods for vessel detection or retinal registration are not suited to real-time operation, preferring accuracy over speed for offline use, and do not handle constraints required for intraocular surgery, such as robustness to tool occlusion.

In this paper, we propose an approach to map and localize the vasculature of the retina in real time that is robust to tool occlusions and variable illumination conditions, for use in intraocular surgery. In Section II, related work in simultaneous localization and mapping (SLAM) and in retinal registration is described. Section III describes our approach of using the fast retinal vessel detection of [24] for feature extraction, an occupancy grid for mapping [25], and iterative closest point (ICP) [26] for localization. In Section IV, we evaluate our approach on videos of an *in vitro* eyeball phantom, *ex vivo* porcine retina, and *in vivo* human retina (see Fig. 1). Section V concludes with a discussion and future work.

II. Related Work

A wealth of published work related to localization and mapping of retinal vessels exists and can be grouped into three general categories: vessel detection, retinal registration, and the more general robotic approach of simultaneous localization and mapping (SLAM).

A. Vessel Detection

Vessel detection is the process of extracting vasculature in retinal imagery and often includes calculating the center lines, width, and orientation of vessels. One set of methods uses local color and intensity information to classify the image on a per-pixel basis [27–30]. Another popular approach is to search across the image for vessel-like structures using matched filters at various locations, scales, and orientations [24], [31], [32]. Other algorithms use a bank of Gabor wavelets to do a pixel-wise classification of the image [33–35]. However, most focus on offline analysis of low-magnification, wide-area images such as fundus images where accuracy is prioritized over speed. With the exception of speed-focused algorithms such as [24], [30], [35] and other hardware-accelerated methods [36], [37], most algorithms require more than 1 s to run, which is insufficiently fast to benefit robotic control loops.

One notable exception is the rapid exploratory algorithm of Can et al. [24] that traces the vasculature, yielding a monotonically improving set of partial results suitable for real-time deployment at 30 Hz. Can et al. [24] achieves high-speed vessel detection through a very fast sparse initialization followed by a tracing algorithm. First, a fast search for vessel points along a coarse grid is performed to initialize a set of seed points on vessels. Each seed point, or detected candidate vessel, is then explored in both directions along the vessel with an approximate and discretized matched filter. At each iteration, the best fit for location, orientation, and width of the vessel center line is estimated through the evaluation of several matched filters. Using orientation estimates to initialize the next iteration, the network of vessels in the image is traced without having to evaluate areas lacking vessels. Because only a small fraction of the total number of pixels in the image is ever processed, most of the

vessels can be detected very rapidly. However, the vessel detections of [24] are noisier and less complete than other, more computationally expensive methods.

B. Retinal Image Registration

Numerous approaches to registering, mosaicking, and tracking exist to take a sequence of retinal imagery and calculate relative motion between images. In general, approaches match one or more of several features between images: key points, vasculature landmarks, or vasculature trees. Key point algorithms use image feature descriptors such as SIFT [38] to find and match unique points between retinal images [39–42]. Vasculature landmark matching algorithms find distinctive points based on vessel networks, such as vein crossings or bifurcations, and match custom descriptors across multiple images [43–45]. Other approaches augment or eschew key points and use the shape of extracted vessels to match vasculature trees [46], [47].

Methods that depend on local key point features [39], [40], [42], [48] often result in poor tracking at high magnification because of the lack of texture on the retina and the non-distinctive nature of individual points on the veins. With optimization, the algorithms of [43], [45] could be run in real time on modem hardware; however, they only use sparse retinal vessel landmarks, which are relatively few or non-existent at high magnifications. More importantly, they only perform localization and do not build a map of the vasculature. Also, all of these approaches suffer from interference caused by the instruments, which both occlude existing features in the image and create new, spurious features on the moving shaft. Our approach is most similar to Stewart et al. [44] which uses a robust, dual-boot ICP algorithm to register vasculature landmarks (bifurcations and crossovers) and trees (centerlines of the vessels). Stewart et al. achieves very accurate results, but our algorithm is 100X faster and handles occlusions, dynamic lighting conditions, and occlusion while yielding temporally consistent map.

C. Simultaneous Localization and Mapping (SLAM)

The problem this paper addresses is similar to a core problem addressed in robotics: simultaneous localization and mapping, or SLAM. In SLAM, a robot with imprecise, noisy localization (e.g., odometry) explores an unknown environment with local sensors (such as a laser range-finder) with the goal of building a global map and localizing itself relative to this map [49]. Using a probabilistic formulation, SLAM optimizes a joint probability over the map and the localization to simultaneously solve for the true positions of the robot and global environmental features. Early solutions such as the Extended Kalman Filter (EKF) scaled poorly and did not handle ambiguous landmark associations well [49]. Recent particle filter approaches such as FastSLAM are faster and more robust [50]. With the introduction of occupancy grids, which discretize the map and maintain a grid of probabilities representing whether each cell is occupied, SLAM algorithms scale more effectively [25].

Comparing SLAM to our problem, the task of building a temporally consistent map of vasculature and localizing the current observation of vessels to this map exhibits many similarities. However, most implementations of SLAM are tailored to space-carving sensors such as laser range-finders instead of over-head sensors and assume a reasonably good robotic motion model, both of which are poor assumptions in the problem of retinal localization and mapping.

III. Methods

Our goal is to design an algorithm that maps and localizes retinal vessels by merging retinal vessel detection with retinal image registration and taking advantage of temporal

information as seen in SLAM approaches. A fusion of methods is needed: fast retinal vessel detections algorithms are noisy, incomplete, and do not handle occlusions [24]; retinal image registration methods that do build vasculature maps are orders of magnitude slower than required for real-time robotic guidance [46], [47]; and SLAM algorithms are not designed or tuned for application in intraocular surgeries. Fig. 2 shows a new algorithm that incorporates aspects of [24], [44], [49] to perform 30 Hz vasculature mapping and localization of retinal video using rapidly-detected vessels as features, an occupancy grid for mapping, and iterative closest point (ICP) for localization to robustly handle noise, tool occlusions, and variable illumination.

A. Problem Definition

Given an series of input video frames $I = [I_0, I_1, \ldots, I_T]$ over a discretized time period $t \in [0, 1, \ldots, T]$, the algorithm should output a global map in the form of N vasculature points $M = [M^0, M^1, \ldots, M^N]$ and the corresponding camera viewpoint locations $L = [L_0, L_1, \ldots, L_T]$ of the input video frames in the map. At time t, we parameterize the ith vasculature point as a 2D location $M_t^i = \left[x_t^i, y_t^i \right]$. Because typical retinal surgeries have high magnification (often the view is only a few mm^2 of a 25 m diameter eye), we approximate the global map as a plane planar section of retina. Similar to many other approaches to retinal registration, we assume an affine camera with a viewpoint at time t as a 2D translation and rotation $L_t = [x_t, y_t, \theta_t]$ from the initial position at $t = 0$. As seen from our results, this 3-DOF motion model is sufficient even with fair large field of views of the retina. Observations of vessels in the camera at time t are denoted by Z_t. The remainder of this section describes the steps that take these video inputs and return these mapping and localization outputs.

B. Feature Extraction

Finding features for matching at high magnification is difficult. Traditional key point detectors such as SIFT [38] or SURF [51] fail to find distinctive points on the textureless retina. Likewise, enough vasculature landmarks such as crossovers or bifurcations may not be present in sufficient numbers to function as good features with high magnification. We instead use many anonymous feature points extracted from vessel detection algorithms. These features cannot be matched individually with a local feature descriptor, but can instead be matched as a group based on the shape of the vasculature network of vessels. We use the highly-efficient, but noisy vessel tracing algorithm of Can et al. [24] to detect vessels, which form the anonymous points (see Fig. 3).

To cull spurious detections on vessel-looking structures such as the tip of the instrument or light-pipe, each vessel point must pass a color test and bloom proximity test. The color test rejects pixels that are too dark or insufficiently red, while the bloom proximity test rejects vessels points that are too close to large white specular blooms in the image. These two simple tests reject many false positives in the detection stage and yield the current observation Z_t as 2D points.

C. Mapping

A global map that holds the current best estimate of all the observed vasculature is maintained using an occupancy grid O_t, which discretizes the map into pixel-sized cells (see Fig. 3). Each pixel in the occupancy grid represents the probability that a vessel occupies that particular spatial location. At each time instance t, the current observations Z_t are transformed to the map with the best estimate of the location L_t^{-1} and used to update the probabilities in the occupancy map by adding a Gaussian around each detected vessel point Z_t^i as each observation increases the probability that a vessel exists at the detected point. The

occupancy grid has a maximum value to prevent unbounded evidence from accumulating. A global decay function decreases the probability of all grid cells, allowing vessels that have not been detected to vanish after some time. While it might be more robust to explicitly consider deformation instead of a decay to let the map react to changes, tool/tissue interaction is non-rigid and difficult to model, especially in real-time.

The formulation of the occupancy map reasons about uncertainty over time, smoothing noise and handling occlusions and deformations. A final map containing the centerlines of the most probable vessels is constructed by skeletonizing the occupancy grid, which is approximated by thresholding, computing the distance transform, calculating the Laplacian, and thresholding again. This yields a map of 2D points M_t of vessels in the occupancy grid (see Fig. 3).

D. Localization and Motion Model

To localize eye-ball motion (which is mathematically identical to localizing camera motion), a 3-DOF planar motion model is chosen. The problem of localization is then to estimate the 2D translation and rotation L_t between the current observations Z_t and the map M_t, both of which are represented by an un-ordered, anonymous set of 2D points. Iterative closest point (ICP) is used to find point correspondences and calculate the transformation. To guarantee real-time performance, Z_t and M_t are randomly sampled to have a maximum of 500 points. To prevent spurious detections from causing large mismatches and adversely affecting the solution, candidate correspondences are only used to estimate the transform if their distance is under some threshold. Horn's quatemion-based method [52] is used to estimate the rigid transform instead of an affine or similarity transform because scale and shear are negligible. Incomplete vessel detections at each frame are noisy, so the final ICP estimation of the localization is smoothed using a constant-velocity Kalman filter, yielding the localization L_t. The occupancy map is then updated with the newly registered vessel points Z_t to close the loop on the algorithm.

IV. Evaluation

We have evaluated the proposed algorithm on a variety of videos of recorded *in vitro* eye phantom, *ex vivo* porcine retina, and *in vivo* human procedures.

A. Setup and Timing Performance

For ease of robotic testing in our lab, color video recorded of an eyeball phantom or porcine retina is captured at 30 Hz with a resolution of 800×608 at a variety of high magnifications (10-25X). Each frame is converted to grayscale by selecting the green channel, a common practice in many vessel detection algorithms. For efficiency, the image is scaled to half-size. On a fast, modem computer (Intel i7-2600K), our algorithm implemented in C++ with OpenCV [53] runs at 30-40 Hz, including all vessel feature detection, occupancy grid mapping, and ICP localization run in a single thread. Fig. 1 shows the proposed algorithm output on an *in vivo* human procedure.

B. Initialization and Convergence Results

Initialization is fast, requiring less than a second to start building the map and only a few seconds to build a full map. Fig. 4 shows the output of the algorithm running on a 48 s clip of a surgeon tracing a vein in an eyeball phantom. Within half a second, the map for visible areas is initialized. Because the light-pipe only illuminates portions of the eye at a time, the map is built as new vessels become visible. The proposed algorithm is able to handle the occlusion of the vessels by the instrument shaft and light-pipe. Notice the algorithm has

been able to accurately map and correctly localize even though the entire view has moved, rotated, and been occluded by the tool under variable illumination.

C. Intermediate Detection and Occupancy Grid Results

Fig. 3 shows the intermediate steps of the algorithm on *ex vivo* porcine retina in an eyeball phantom filled with saline. The view is through a vitrectomy lens and illumination is provided by the surgeon solely through a light-pipe. Incomplete vessel detections are merged over time into the occupancy grid to form a full and accurate map after a few seconds. Fig. 3(c) shows that the localization is still maintained after movement of the eyeball by the surgeon.

D. Comparison to Vessel Detection Algorithms

Fig. 5 demonstrates why current vessel tracing algorithms are insufficient for robotic guidance in real time surgical environments. We compare to three existing vessel tracing algorithms on video of human retina *in vivo* taken in during a membrane peel. Our proposed approach provides a more complete output than any of the other methods. In particular, as examination of Figs. 5(a) and 5(b) shows, our algorithm learns over time to ignore the tool shaft, exhibiting more robustness to occlusion and illumination. Overall, the temporal fusion of the proposed approach increases coverage and consistency.

V. Discussion

We have presented a new algorithm for retinal mapping and localization that operates in real time at 30 Hz. Designed to handle the dynamic environment of high-magnification, variable illumination retinal surgery, our approach converges quickly and is robust to occlusion. In comparisons on retinal video, it has proven to be an effective method to temporally smooth vessel detections and build a comprehensive map of the vasculature. Shortcomings to the algorithm include some lag when smoothing jitter with the Kalman filter and loss of tracking in the case of large, sudden movements. Future improvements should include robust vessel detection, ICP, or scan-matching approaches. More effective handling of uncertainty during localization and advanced motion models (such as spherical) would also be beneficial. Finally, future work includes quantitative comparisons to existing methods.

Acknowledgments

References

[1]. Holtz, F.; Spaide, RF. Medical Retina. Springer; Berlin: 2007.

[2]. Singh SPN, Riviere CN. Physiological tremor amplitude during retinal microsurgery. Proc. IEEE Northeast Bioeng. Conf. 2002:171–172.

[3]. Brooks HL. Macular hole surgery with and without internal limiting membrane peeling. Ophthalmology. 2000; vol. 107(no. 10):1939–1948. [PubMed: 11013203]

[4]. Packer, AJ. Manual of Retinal Surgery. Butterworth-Heinemann Medical; 2001.

[5]. Rice TA. Internal limiting membrane removal in surgery for full-thickness macular holes. Am. J. Ophthalmol. 1999; vol. 129(no. 6):125–146.

[6]. Frank RN. Retinal laser photocoagulation: Benefits and risks. Vision Research. 1980; vol. 20(no. 12):1073–1081. [PubMed: 7196642]

[7]. Infeld DA, O'Shea JG. Diabetic retinopathy. Postgrad. Med. J. 1998; vol. 74(no. 869):129–133. [PubMed: 9640436]

[8]. Weiss JN, Bynoe LA. Injection of tissue plasminogen activator into a branch retinal vein in eyes with central retinal vein occlusion. Ophthalmology. 2001; vol. 108(no. 12):2249–2257. [PubMed: 11733266]

[9]. Bynoe LA, Hutchins RK, Lazarus HS, Friedberg MA. Retinal endovascular surgery for central retinal vein occlusion: Initial experience of four surgeons. Retina. 2005; vol. 25(no. 5):625–632. [PubMed: 16077361]

[10]. Das H, Zak H, Johnson J, Crouch J, Frambach D. Evaluation of a telerobotic system to assist surgeons in microsurgery. Comput. Aided Surg. 1999; vol. 4(no. 1):15–25. [PubMed: 10417827]

[11]. Ueta T, Yamaguchi Y, Shirakawa Y, Nakano T, Ideta R, Noda Y, Morita A, Mochizuki R, Sugita N, Mitsuishi M. Robot-assisted vitreoretinal surgery: Development of a prototype and feasibility studies in an animal model. Ophthalmology. 2009; vol. 116(no. 8):1538–1543. [PubMed: 19545902]

[12]. Wei W, Goldman RE, Fine HF, Chang S, Simaan N. Performance evaluation for multi-arm manipulation of hollow suspended organs. IEEE Trans. Robot. 2009; vol. 25(no. 1):147–157.

[13]. Bourla DH, Hubschman JP, Culjat M, Tsirbas A, Gupta A, Schwartz SD. Feasibility study of intraocular robotic surgery with the da Vinci surgical system. Retina. 2008; vol. 28(no. 1):154–158. [PubMed: 18185154]

[14]. Mulgaonkar AP, Hubschman JP, Bourges JL, Jordan BL, Cham C, Wilson JT, Tsao TC, Culjat MO. A prototype surgical manipulator for robotic intraocular micro surgery. Stud. Health Technol. Inform. 2009; vol. 142(no. 1):215–217. [PubMed: 19377152]

[15]. Uneri A, Balicki M, Handa J, Gehlbach P, Taylor RH, Iordachita I. New Steady-Hand Eye Robot with micro-force sensing for vitreoretinal surgery. Proc. IEEE Int. Conf. Biomed. Robot. Biomechatron. 2010:814–819.

[16]. Hubschman JP, Bourges JL, Choi W, Mozayan A, Tsirbas A, Kim CJ, Schwartz SD. The Microhand: A new concept of micro-forceps for ocular robotic surgery. Eye. 2009; vol. 24(no. 2):364–367. [PubMed: 19300461]

[17]. Dogangil G, Ergeneman O, Abbott JJ, Pané S, Hall H, Muntwyler S, Nelson BJ. Toward targeted retinal drug delivery with wireless magnetic microrobots. Proc. IEEE Intl. Conf. Intell. Robot. Syst. 2008:1921–1926.

[18]. MacLachlan RA, Becker BC, Cuevas Tabares J, Podnar GW, Lobes LA Jr. Riviere CN. Micron: An actively stabilized handheld tool for microsurgery. IEEE Trans. Robot. 2012; vol. 28(no. 1): 195–212. [PubMed: 23028266]

[19]. Becker BC, MacLachlan RA, Lobes LA Jr. Hager G, Riviere C. Vision-based control of a handheld surgical micromanipulator with virtual fixtures. IEEE Trans. Robot. submitted.

[20]. Rosenberg LB. Virtual fixtures: Perceptual tools for telerobotic manipulation. IEEE Virt. Reality Ann. Int. Symp. 1993:76–82.

[21]. Bettini A, Marayong P, Lang S, Okamura AM, Hager GD. Vision-assisted control for manipulation using virtual fixtures. IEEE Trans. Robot. Dec.2004 vol. 20(no. 6):953–966.

[22]. Lin HC, Mills K, Kazanzides P, Hager GD, Marayong P, Okamura AM, Karam R. Portability and applicability of virtual fixtures across medical and manufacturing tasks. Proc. IEEE Int. Conf. Robot. Autom. 2006:225–231.

[23]. Becker BC, Voros S, Lobes LA Jr. Handa JT, Hager GD, Riviere CN. Retinal vessel cannulation with an image-guided handheld robot. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. 2010:5420–5423.

[24]. Can A, Shen H, Turner JN, Tanenbaum HL, Roysam B. Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. IEEE Trans. Inform. Technol. Biomed. 1999; vol. 3(no. 2):125–138.

[25]. Hahnel D, Burgard W, Fox D, Thrun S. An efficient FastSLAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements. Proc. IEEE Intl. Conf. Intell. Robot. Syst. 2003; vol. 1:206–211.

[26]. Besl PJ, McKay ND. A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Mach. Intell. 1992; vol. 14(no. 2):239–256.

[27]. Chanwimaluang T. An efficient blood vessel detection algorithm for retinal images using local entropy thresholding. Proc. Int. Symp. Circuits Syst. 2003; vol. 5:21–24.

[28]. Mojon D. Adaptive local thresholding by verification-based multithreshold probing with application to vessel detection in retinal images. IEEE Trans. Pattern Anal. Mach. Intell. Jan. 2003 vol. 25(no. 1):131–137.

[29]. Staal J, Abràmoff MD, Niemeijer M, Viergever MA, van Ginneken B. Ridge-based vessel segmentation in color images of the retina. IEEE Trans. Med. Imag. Apr.2004 vol. 23(no. 4): 501–9.

[30]. Alonso-Montes C, Vilariño DL, Dudek P, Penedo MG. Fast retinal vessel tree extraction: A pixel parallel approach. International Journal of Circuit Theory and Applications. Jul.2008 vol. 36(no. 5-6):641–651.

[31]. Chaudhuri S, Chatterjee S, Katz N, Nelson M, Goldbaum M. Detection of blood vessels in retinal images using two-dimensional matched filters. IEEE Trans. Med. Imag. Jan.1989 vol. 8(no. 3): 263–269.

[32]. Sofka M, Stewart CV. Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures. IEEE Trans. Med. Imag. Dec.2006 vol. 25(no. 12):1531–1546.

[33]. Soares JVB, Leandro JJG, Cesar RM, Jelinek HF, Cree MJ. Retinal vessel segmentation using the 2-D Gabor wavelet and supervised classification. IEEE Trans. Med. Imag. Sep.2006 vol. 25(no. 9):1214–1222.

[34]. Lupascu CA, Tegolo D, Trucco E. FABC: Retinal vessel segmentation using AdaBoost. IEEE Trans. Inform. Technol. Biomed. Sep.2010 vol. 14(no. 5):1267–1274.

[35]. Bankhead P, Scholfield CN, McGeown JG, Curtis TM. Fast retinal vessel detection and measurement using wavelets and edge location refinement. PLOS ONE. Jan.2012 vol. 7(no. 3)

[36]. Perfetti R, Ricci E, Casali D, Costantini G. Cellular neural networks with virtual template expansion for retinal vessel segmentation. IEEE Transactions on Circuits and Systems II: Express Briefs. Feb.2007 vol. 54(no. 2):141–145.

[37]. Nieto A, Brea VM, Vilarino DL. FPGA-accelerated retinal vessel-tree extraction. International Conference on Field Programmable Logic and Applications. 2009:485–488.

[38]. Lowe DG. Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. 2004; vol. 60(no. 2):91–110.

[39]. Cattin P, Bay H, Van Gool L, Székely G. Retina mosaicing using local features. Medical Image Computing and Computer-Assisted Intervention. 2006; vol. 4191:185–192. [PubMed: 17354771]

[40]. Wang Y, Shen J, Liao W, Zhou L. Automatic fundus images mosaic based on SIFT feature. Journal of Image and Graphics. 2011; vol. 6(no. 4):533–537.

[41]. Chen J, Smith R, Tian J, Laine AF. A novel registration method for retinal images based on local features. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. 2008; vol. 2008:2242–2245.

[42]. Becker BC, Yang S, MacLachlan RA, Riviere CN. Towards vision-based control of a handheld micromanipulator for retinal cannulation in an eyeball phantom. Proc. IEEE Int. Conf. Biomed. Robot. Biomechatron. 2012:44–49.

[43]. Becker DE, Can A, Turner JN, Tanenbaum HL, Roysam B. Image processing algorithms for retinal montage synthesis, mapping, and real-time location determination. IEEE Trans. Biomed. Eng. 1998; vol. 45(no. 1):105–118. [PubMed: 9444845]

[44]. Can A, Stewart CV, Roysam B, Tanenbaum HL. A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina. IEEE Trans. Pattern Anal. Mach. Intell. Mar.2002 vol. 24(no. 3):347–364.

[45]. Broehan AM, Rudolph T, Amstutz CA, Kowal JH. Realtime multimodal retinal image registration for a computer-assisted laser photocoagulation system. IEEE Trans. Biomed. Eng. Oct.2011 vol. 58(no. 10):2816–2824. [PubMed: 21689999]

[46]. Stewart CV, Tsai C-L, Roysam B. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. IEEE Trans. Med. Imag. Nov.2003 vol. 22(no. 11): 1379–1394.

[47]. Chanwimaluang T, Fan G, Fransen SR. Hybrid retinal image registration. IEEE Trans. Inform. Technol. Biomed. Jan.2006 vol. 10(no. 1):129–142.

[48]. Chen J, Tian J, Lee N, Zheng J, Smith RT, Laine AF. A partial intensity invariant feature descriptor for multimodal retinal image registration. IEEE Trans. Biomed. Eng. Jul.2010 vol. 57(no. 7):1707–1718. [PubMed: 20176538]

[49]. Thrun, S. Simultaneous Localization and Mapping. In: Jefferies, M.; Yeap, W-K., editors. Robotics and Cognitive Approaches to Spatial Mapping. Vol. vol. 38. Springer; Berlin / Heidelberg: 2008. p. 13-41.

[50]. Montemerlo M, Thrun S, Koller D, Wegbreit B. FastSLAM: A factored solution to the simultaneous localization and mapping problem. Proc. Nat. Conf. Artificial Intell. 2002:593–598.

[51]. Bay H, Ess A, Tuytelaars T, van Gool L. SURF: Speeded Up Robust Features. Comp. Vis. Img. Understand. 2008; vol. 110(no. 3):346–359.

[52]. Horn BKP. Closed-form solution of absolute orientation using unit quaternions. J. Opt. Soc. Am. A. 1987; vol. 4(no. 4):629–642.

[53]. OpenCV Library. [Online]. Available: http://opencv.willowgarage.com/
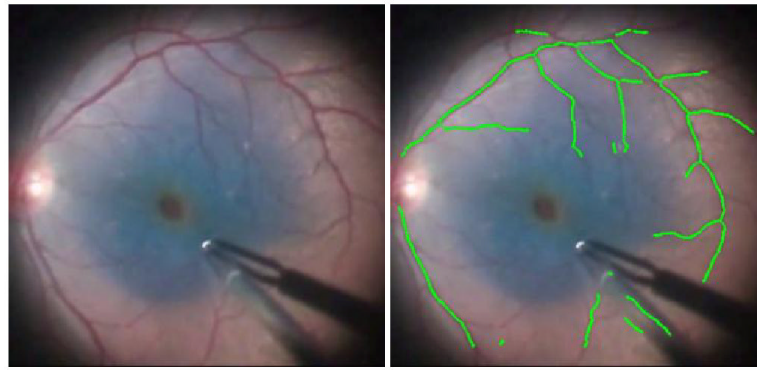
Fig. 1.
The proposed mapping and localization algorithm for retinal vasculature running in real time on recorded *in vivo* human retina during a retinal peeling with blue die. Video source: http://youtu.be/CTnavOgDsXA
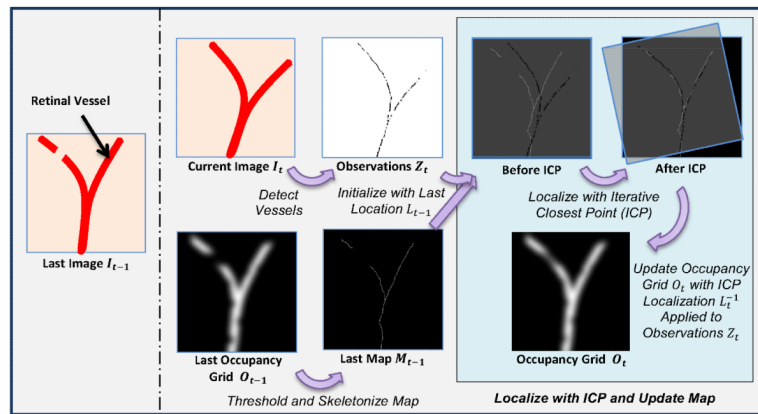
Fig. 2.
Block diagram detailing the steps in our proposed algorithm that maps and localizes retinal vessels during intraocular surgery. Retinal vessels are detected III each frame, localized to a skeletonized map of the occupancy grid with iterative closest point (ICP), and the map probabilities are updated.
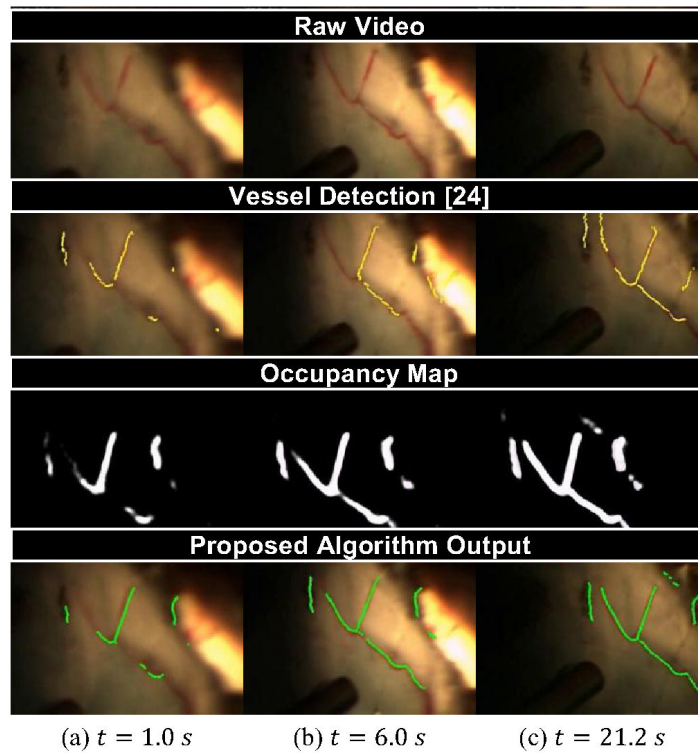
(a) $t = 1.0\ s$ (b) $t = 6.0\ s$ (c) $t = 21.2\ s$

Fig. 3.
Snapshots of (a) the raw video, (b) feature extraction, (c) mapping, and (d) localization of the proposed algorithm at various times during 21 s video of porcine retina. Notice that our algorithm builds up the full vessel network even with incomplete frame-to-frame detections from [24].

Fig. 4.
Output of the proposed algorithm at various points during a 48 s video of an eye phantom filled with saline and illuminated with a light-pipe.

(a) $t = 1.0\ s$ (b) $t = 5.7\ s$ (c) $t = 19.9\ s$

Fig. 5.
Comparison of our approach (30 Hz) to the vessel detection algorithms of Can et al. [24] (80 Hz), Sofka et al. [32] (0.3 Hz), and Bankhead et al. [35] (8 Hz) on a 20 s video of human retina during an *in vivo* membrane peel. Source: http://youtu.be/_naooJFuxPI