2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Visual Vein-Finding for Robotic IV Insertion

Reuben D. Brewer and J. Kenneth Salisbury

Abstract— **This paper presents a new algorithm for selecting the optimal needle insertion point in images of hand veins for robotic IV insertion. The 3D coordinates and orientation of the vein that the algorithm detects would eventually be fed to a robot for insertion of the IV needle. The goal of the algorithm is to identify venous bifurcations and determine an insertion point and approach angle for the needle in between their branches. The algorithm uses an annular tracking window that tracks along the veins and searches for bifurcations. We describe methods for centering the initial bifurcation estimates, error-checking, and positioning the needle exactly in between the bifurcation branches. We conclude with an experimental study of 50 subjects that shows a 32.4% success rate at detecting all bifurcations and a 82.6% success rate at finding at least one bifurcation in each image that contains bifurcations.**

I. Introduction

A. Background

Intravenous (IV) catheterization is a medical procedure wherein a flexible plastic tube, or catheter, is inserted into a vein for the delivery of medicinal fluids. The catheter initially surrounds a needle that punctures the wall of a vein so that the catheter can be slid off of the needle and into the vein, whereupon the needle is removed. Nearly 1 billion IV insertions take place in the United States annually, and 28% of those insertions fail on the first attempt in normal adults, with appreciably higher failure rates in children [1], [2]. Failed insertions commonly cause bruising and pain, but can also lead to long-term nerve damage and schlerosis of the veins.

Robotic IV insertion has been proposed as a possible solution to increase the insertion success rate through precise movement of the needle and enhanced sensory abilities [3], [4]. Towards this end, we are currently developing a 7-DOF robot, as shown in Figure 1, that can insert either under teleoperation or autonomously. Such a system could be used to treat people in remote or hostile locations where a human practitioner could not be present. Autonomous insertion could be used in hospitals to increase the success and through-put of phlebotomists, nurses, and anesthesiologists. While good mechanical design can provide precise needle motion, the robot must also include robust vein detection through various sensors to provide the target location and insertion trajectory for the needle.

One of the reasons for the high failure rate for human practitioners is the difficulty of locating veins precisely.

R. Brewer is with the Department of Mechanical Engineering, Stanford University, Stanford, CA, USA, rdbrewer@stanford.edu

J. Salisbury is with the Departments of Computer Science and Surgery, Stanford University, Stanford, CA, USA, jks@robotics.stanford.edu

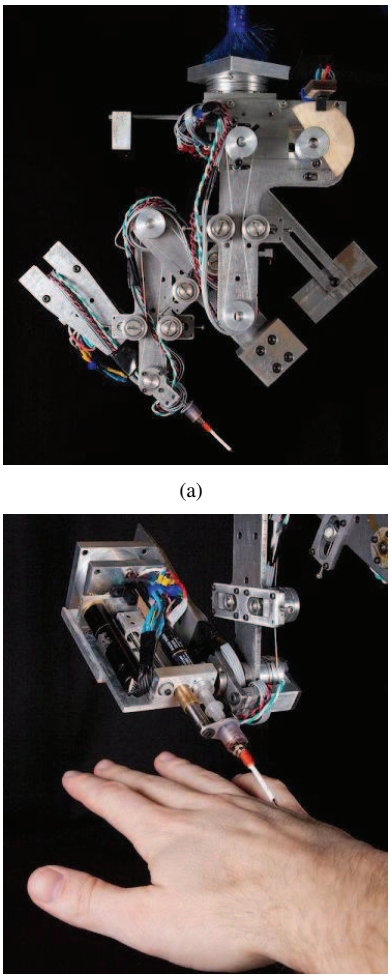

(a)

(b)

Fig. 1. 7 DOF robotic IV insertion system currently under development.

Practitioners often have great difficulty seeing or feeling small veins, as found in women and children, or veins that lie beneath a layer of fat, as found in children and obese patients. However, detection of the veins through various sensors could allow a robot to target veins that human practitioners otherwise could not localize. In [3], robotic palpation, or tactile sensing, was used to locate veins by instrumenting a probe with a force sensor and examining changes in tissue stiffness across the arm. However, the pressure of palpation may roll the vein away from the probe, thereby skewing the position information. Further, palpation suffers from an inability to map more than a small segment of vein or provide real-time tracking of the vein in case the vein/arm moves during insertion. An alternative is the use of

infrared imaging, as in [4]. The deoxygenated haemoglobin in veins absorbs infrared light more than surrounding tissues, so veins appear as dark on a light background under infrared light [5]. Such imaging systems are commercially available for assisting human practitioners and can be readily adapted to provide images of veins to a robot [6], [7]. However, a topic of ongoing research is how to provide a robot with the ability to analyze and interpret these vein images so as to select an optimal insertion location.

The conventional wisdom of many practitioners is that inserting at a venous bifurcation minimizes the vein rolling as opposed to inserting along the midsection of a single vein because bifurcations are tethered by more connective tissue [8]. For this reason, our algorithm searches the venous network for bifurcations as the optimal insertion points. Since the dorsum of the hand contains far more bifurcations than does the forearm or antecubital fossa (inside of the elbow), we are restricting our insertions to the hand. Concentrating on the hand has the additional benefits of allowing for easier mechanical access to the insertion point and the ability to use simple transillumination of the veins, whereas the forearm is too thick for transillumination.

B. Related Work

Bifurcation detection has been studied widely for images of retinal arteries. In [9], the entire image is searched using the assumption that bifurcations will occur at areas of high variation in Sobel edge direction. Similarly, in [10] and [11] a course grid of seed-points that looks for local gray-scale minima between oppositely-signed edges is applied across the entire image. Recursive tracing proceeds from these seed-points by moving in the direction that best fits the vein model of a gray-scale minimum between antiparallel edges. However, searching the entire image can generate many false positives, especially if the image is noisy or has artifacts. In the case of our particular vein images, the variation in gray-scale intensity and contrast of the veins leads to poor edge detection such that the edges cannot be followed reliably. An alternative to searching the entire image is to trace along the arteries from a known landmark, hopefully reducing the number of false-positives. In [12], Canny edge detection is performed, and tracing occurs along the detected edges, starting from the optic disk. In [13], a steerable Gaussian filter is used to trace from the optic disk. This tracing assumes thin, well-defined vessels with consistent cross-sectional profiles. Unfortunately, the vessels in our images exhibit cross-sectional intensity profiles of varying size, noisiness, and degree of saturation such that looking at their cross-sectional profile does not give much reliable information.

C. Overview

In this paper, we present a new algorithm for finding the position and orientation of venous bifurcations in an image. We first detail a method for locating the wrist veins as a reliable starting point for vessel tracking. We then present a new annular tracking window that examines the macro-structure of the vein instead of a small slice, allowing for tracking noisy, variable veins with erratic edges, as well as detecting bifurcations during the tracking process. We describe a method for centering an initial bifurcation estimate on the true center of the bifurcation by using a set of concentric annular windows. We detail methods for error-checking and finding the actual point of insertion in between the branches of a bifurcation. Finally, we discuss the experimental results of applying our bifurcation detector to a set of sample images.

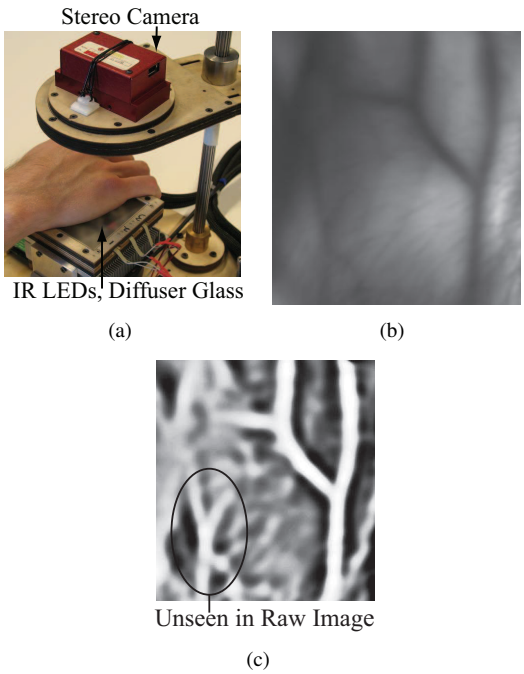

Fig. 2. (a) Imaging system. (b) Raw image of veins. Veins appear as dark on a light background. (c) Enhanced image. Veins now appear as bright on a dark background.

II. Image Acquisition

A. Imaging System

Our imaging system is a low-cost approximation of commercially-available infrared imaging systems such as [6] and [7]. Our setup transilluminates the hand from the palm side and views the illuminated venous structure from the dorsal side, as shown in Figure 2(a). The light source consists of 5 high-power Osram SFH-4730 infrared LEDs ($\lambda_{peak} = 850nm$) spread out across the palm and covered with a diffusing glass to provide even lighting. It should be noted that according to [5] and [14], the light-blocking capacity of deoxyhaemoglobin (the deoxygenated blood in veins) is highest at a wavelength of 760nm and is reduced by as much as 50% from its peak value at 850nm. Further, the light-blocking capacity of oxyhaemoglobin (the oxygenated blood in arteries) is much lower at 760nm than at 850nm. Therefore, the contrast between veins and the surrounding tissues, and in particular arteries, is maximal at 760nm, resulting in more readily-identifiable veins in the images and less ghosting effects from faint arteries. However, the available

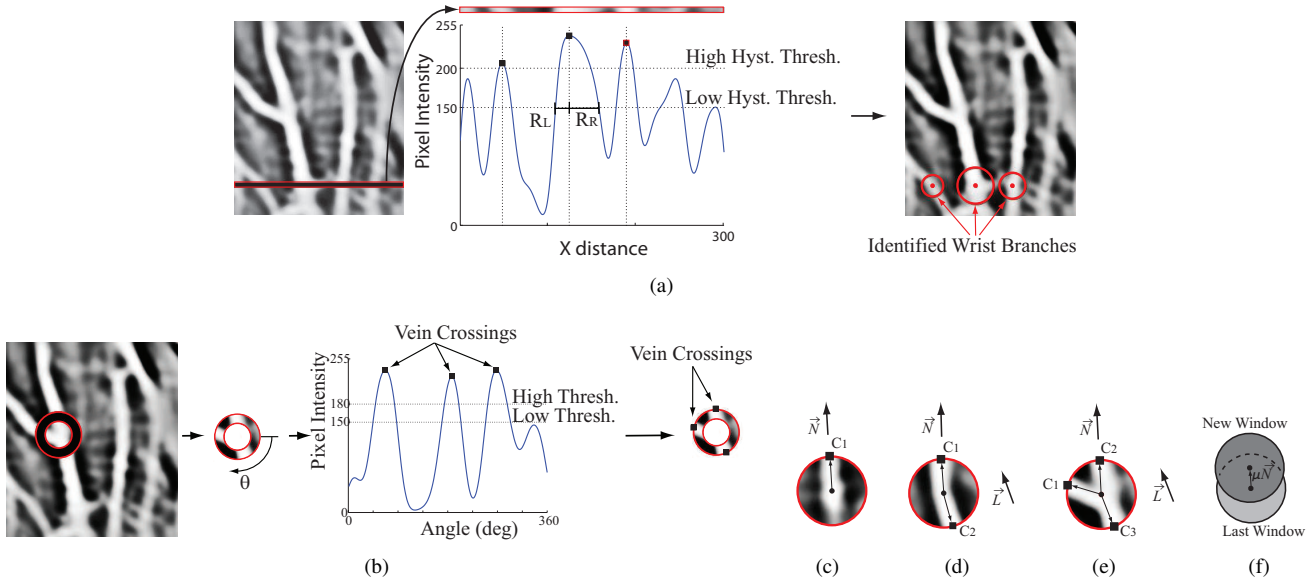

Fig. 3. (a) Identifying wrist branches with histogram at bottom of image. (b) Polar histogram of annular tracking window. (c) 1 Vein crossing represents a faded, single vein. (d) 2 Vein crossings represent a solid, single vein. (e) 3 Vein crossings represent a bifurcation. (f) Incremental movement of tracking window.

light source for a wavelength of 850nm was considerably more practical. The camera is a Videre STOC stereo camera (6 cm baseline) with the IR filter removed and provides 640x480 monochrome images, as shown in Figure 2(b). The stereo camera allows us to calculate the 3D position and orientation of the vein that we detect in a 2D image.

B. Image Enhancement

The veins are often difficult to identify in the raw images due to low contrast and fading of some vein sections, as seen in Figure 2(b). For this reason, we enhance the images to increase the vein vs. non-vein contrast and normalize the appearance of the veins so that all sections appear similar. We use Laplace of Gaussian (LOG) filtering (σ = 8, window width of 5σ, determined empirically) to accentuate the vein macrostructure over the background and histogram equalization to improve the overall contrast in the image. Figure 2(c) shows the enhanced version of the raw image seen in Figure 2(b).

III. Finding Initial Bifurcation Estimates

A. Identifying Wrist Veins

Our bifurcation finder operates by finding the most prominent veins at the base of the wrist and tracking along those veins with a window that looks for bifurcations. The wrist veins are composed of the basilic, cephalic, and their branches, such as the accessory cephalic vein. The desired bifurcations occur in the dorsal metacarpal veins in the top of the hand. Whereas the venous network in the hand varies greatly between individuals, the location of wrist veins is fairly consistent, providing an ideal starting location for tracking. As in [12] and [13], we track along the veins from a known location instead of searching the entire image because a bifurcation that we detect while tracking along a vein is more likely to be real than a bifurcation that we detect at a random point in the image. To identify the wrist veins, we take a horizontal sample strip (height = 8 pixels, determined empirically) across the bottom of the image and examine the smoothed histogram of gray-scale intensity, as shown in Figure 3(a). After collapsing the strip to a single pixel-width height by taking the median in the vertical direction, we apply gaussian filtering in the horizontal direction and hysteresis-thresholding to mitigate the noise in the histogram. Since veins appear as bright against a dark background, they are identified on the histogram as local maxima. Each identified wrist vein becomes a starting point for tracking. We use the histogram to size each tracking window to be twice the width of the wrist vein that it will track. The tracking window is larger than the vein being tracked so that the tracking window is never entirely inside the vein.

B. Tracking Veins

Tracking along the wrist veins employs a steerable window that uses information about the macrostructure of the vein section contained within the window to determine how to move along the vein. The tracking window consists of an annulus (thickness = 8 pixels, determined empirically) that is centered on a section of the vein. A polar histogram of gray-scale intensity in the annular sample describes the shape of the vein section contained within the window. As before, the veins are identified in the smoothed histogram as local maxima. After collapsing the annulus to a single pixel-width thickness by taking the median in the radial direction, we apply gaussian smoothing in the angular direction and hysteresis-thresholding to reduce noise in the histogram. Figure 3(b) shows the tracking window in detail.

Figures 3(c) - 3(f) show how the number of vein crossings detected in the tracking window conveys information about

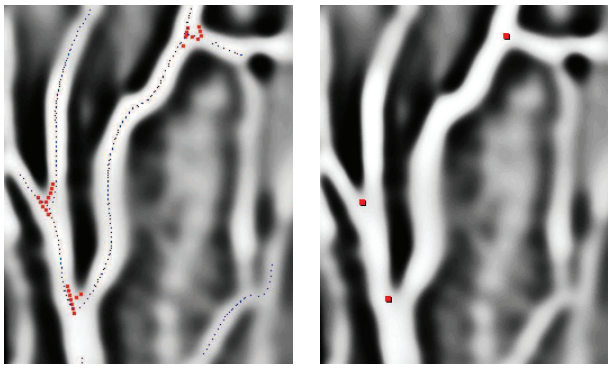

(a) (b)

Fig. 4. Tracking along wrist veins. (a) Un-clustered initial bifurcation estimates. Small black dots show the path of the tracking window, and large red dots show bifurcation estimates. (b) Clustered estimates of bifurcations (shown as large red dots).

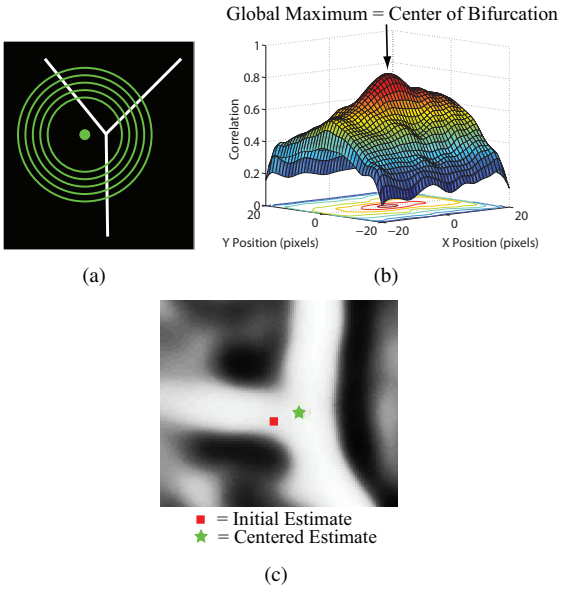

Fig. 5. (a) Concentric annular samples near a bifurcation. (b) Correlation in the set of annular samples in the vicinity of the bifurcation. (c) Example of centering algorithm on a real bifurcation.

the section of vein within the window and how the tracking window should move along the vein. We assume that one vein crossing (C_1) indicates that the tracking window is centered on a faded, single section of vein, two crossings (C_1, C_2) indicate that the tracking window is centered on a solid, single section of vein, and three vein crossings (C_1, C_2, C_3) indicate that the tracking window is centered on a vein near a bifurcation. In all three cases, the tracking window steps a small amount μ (5 pixels, determined empirically) in the new direction $\overrightarrow{N}$ that is computed as the vector between the window center and the vein crossing most closely aligned with the previous direction of movement $\overrightarrow{L}$. The degenerate cases are zero vein crossings, which indicates a lack of information, and more than three crossings, which indicates noise. In either degenerate case, the tracking window steps in the last known direction of movement. In the case of a detected bifurcation, the tracking window continues tracking along the branch most closely aligned with the last known direction of movement until it reaches an edge of the image or stalls, whereupon it returns to track along the alternate (third) branch of the bifurcation. Figure 4(a) shows an example of vein tracking and initial bifurcation detection. Due to the thickness of the veins, there are many points in the vicinity of a single bifurcation that show three vein crossings, or a detected bifurcation. Clustering all of the bifurcation estimates based on distance from each other reliably thins this mass of estimates to one estimate per actual bifurcation detected, as shown in Figure 4(b).

IV. Refining and Error-Checking Bifurcation Estimates

A. Centering the Bifurcation Estimate

The initial bifurcation estimate often does not coincide with the true center of the bifurcation due to the asymmetric distribution of initial guesses that were eventually clustered. Thus, we need an algorithm for finding the true center of the bifurcation given a moderately-close initial estimate. We have devised a solution to this problem by looking at the correlation in a set of concentric annular samples taken at different diameters, as shown in Figure 5(a). These samples are resized versions of the tracking window described above. If we position this set of annuli at each pixel in the vicinity of the bifurcation and examine the correlation, we find that the global maximum of the correlation corresponds to the true center for bifurcations with 3 branches. Since there are more than two samples, we compute the correlation as the sum of correlations between adjacent rings. Figure 5(b) shows the correlation plot for the example bifurcation in Figure 5(a). Figure 5(c) shows an example of the centering algorithm on a real bifurcation. Note that the initial estimate is appreciably off-center but that the centered estimate appears directly at the center of the bifurcation. For brevity, we omit the derivation of this centering algorithm.

B. Error-Checking

We employ two simple methods of error-checking to detect false-positive bifurcations. In a fashion similar to the centering algorithm, we examine concentric annular samples of each bifurcation over 35 different diameters and count the number of diameters that elicit the characteristic three vein crossings of a true bifurcation. If our estimate is a true bifurcation, then it looks like a bifurcation over many scales, whereas false-positives appear like bifurcations over only a few diameters. In practice, false-positives almost uniformly return a low number of counts (below 10), whereas true bifurcations return a much higher number (typically above 25). A final, simple error-check is to ensure that each bifurcation is centered on a bright spot (vein). Figure 6(a) shows examples of both error-checks. Note that the false positive is centered on a dark spot (non-vein) and only shows the three vein crossings for the particular diameter shown, whereas the true bifurcation is centered on a bright spot and shows the three vein crossings over most diameters.

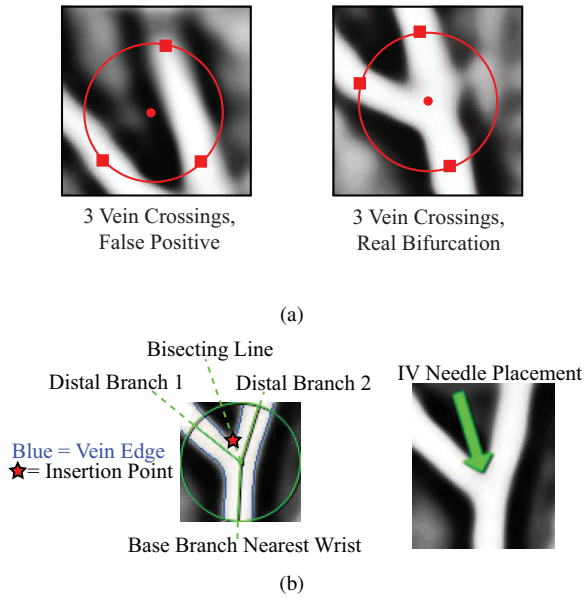

Fig. 6. (a) Error-checking for false-positives. (b) Finding the insertion point, angle.

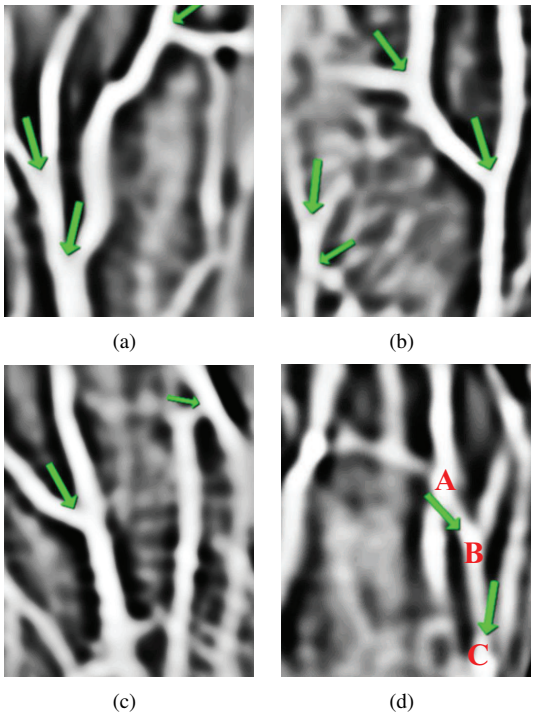

Fig. 7. (a,b,c) Examples of correctly identified and characterized bifurcations. The arrows show the desired position and orientation of the IV needle. (d) Failure modes. Point A is a false negative, point B is a correctly identified but misplaced bifurcation, and point C is a false positive that arose because two separate veins appeared to be a bifurcation.

C. Finding the Insertion Point

The desired insertion point at a bifurcation is not the center of the bifurcation, but, rather, along the edge in between the two branches of the bifurcation furthest from the wrist, as shown in Figure 6(b). To find this point, we use Canny edge detection to compute the edges of the veins in the vicinity of the bifurcation. This edge detection uses locally determined threshold values because the intensity of the veins varies somewhat over the image, precluding the usage of global threshold values. Once the local vein edges are found, we sample along the line that bisects the two distal branches of the bifurcation until we intersect with a vein edge. This intersection at the edge of the vein is the desired insertion point, and the bisecting line is the desired direction for the IV needle. It is acceptable for the bisecting line and needle not to align exactly with the branch nearest the wrist because the needle will not enter the vein far enough to hit the back wall, and the catheter is sufficiently flexible to deflect safely away from the back wall and into the lumen of the base vein branch.

D. Weighting Multiple Bifurcations

If multiple bifurcations are detected in the same image, we must be able to decide which is optimal for insertion. Towards this end, we examine each bifurcation and compute a weighted score based on a variety of desirable traits. These traits include the correlation found during the centering process, orientation with respect to the wrist direction, distance from the center of the hand, and average gray-scale intensity in the vicinity of the bifurcation. The correlation and intensity metrics address our confidence in the authenticity of the bifurcation. The orientation and distance metrics address the practical concerns that we wish to insert the needle in the direction of the wrist and near the center of the hand, as is standard practice among practitioners. While these metrics make sense from an engineering perspective, it will be necessary in the future to determine additional clinical criteria used by practitioners to select amongst bifurcations and implement the same weighting in our algorithm.

V. Experimental Methods and Results

We developed our bifurcation-finding algorithm on a set of thirteen images that provided for iterative refinement. We determined algorithm parameters such as LOG filter σ, window step size μ, and annular thickness empirically in these 13 images. To test the general applicability of the algorithm, we applied it to fifty images that were unseen during development. The test population was 72% male, had an age range of 22-56 (μ = 31.0, σ = 9.1), and was 72% caucasian and 28% asian. To compare the performance of our algorithm with that of human practitioners, we provided four medical doctors with the same images and asked them to identify all bifurcations, without preference to size or desirability. Our algorithm successfully detected and characterized 32.4% of the bifurcations identified by the doctors with a false positive rate of 11.6%. However, our algorithm correctly found and characterized at least one bifurcation in 82.6% of the images that contained bifurcations, which comprised 92% of all images (4 images did not contain any bifurcations). Assuming that it is clinically sufficient to find only one (instead of every) bifurcation per hand for the robot to insert an IV, our vein-finding algorithm would allow us to

insert on the majority (82.6%*92% = 76%) of individuals. Figures 7(a-c) show examples of successfully detected and characterized bifurcations, and Figure 7(d) shows several failure cases. Although the algorithm was not tested on children, we expect future pediatric tests to be successful due to the algorithm's ability to auto-size the tracking window to differently-sized veins. It should be noted that the large width of the tracking window creates margins at the edges of the image that are unsearchable by the algorithm, whereas humans can search the entire image. In fact, many of the bifurcations that the doctors identified and our algorithm did not were near the edges of the image/hand where the algorithm could not search. These "missed" bifurcations are not as important as others because it would be inconvenient to insert an IV at the edge of the hand.

VI. Discussion and Conclusion

In this paper we developed a new visual vein-finding technique for use in a robotic IV system, as shown in Figure 1. In such a system, our algorithm would analyze infrared images of the hand to locate suitable insertion points at bifurcations and provide the robot with the 3D position and orientation of the IV needle for insertion. Specific to our vein-finding algorithm, we discussed a new method for tracking veins that uses an annular tracking window to examine the macro-scale structure of local sections of vein. We then discussed a method for centering our initial rough bifurcation estimate by maximizing the correlation in a set of concentric annular samples near the bifurcation. We detailed a method for error-checking bifurcation estimates and positioning the IV needle exactly in between the branches of a bifurcation. Finally, we tested the performance of our algorithm on a set of 50 images and were able to correctly find and characterize 32.4% of all bifurcations with a false positive rate of 11.6% and find at least one bifurcation in 82.6% of images that contained bifurcations.

The experimental results have revealed several insights as to the limitations of our algorithm. The algorithm has particular difficulty in detecting closely-spaced bifurcations because they appear as a single, noisy bifurcation rather than distinct bifurcations. A limitation of the centering algorithm is the assumption that bifurcations consist of 3 straight branches. In some cases, there are more than 3 branches, or the branches are sufficiently curved to break the centering process. A further problem is that each tracking window maintains a constant size as it tracks along a wrist vein, even though the veins periodically change diameter. This can lead to either too much noise for an oversized tracking window or too little information for an undersized tracking window. Dynamic resizing of the tracking window is a potential solution to this problem.

Beyond addressing these limitations, future work will focus on using depth information from the stereo images to calculate the desired world-frame position and orientation (roll, pitch, and yaw) of the IV needle based on the selected bifurcation in a particular image. The robot will use this calculated configuration to generate a trajectory for insertion. Real-time visual tracking of the selected insertion point will continue to update the desired insertion trajectory for the robot throughout the insertion. We will also investigate the effect of light scattering on the accuracy of depth information in the stereo infrared images, as well as possible benefits of switching between infrared and visible-light imaging so as to switch between internal and external views of the hand, respectively.

VII. Acknowledgments

R. Brewer was supported in part by an ASEE NDSEG Fellowship and a NSF Graduate Research Fellowship. We are thankful to Dr. Greg Hager and Dr. Jana Kosecka for their consultation in this work.

References

[1] R. Lenhardt, "Local warming and insertion of peripheral venous cannulas: single blinded prospective randomised controlled trial and single blinded randomised crossover trial," *British Medical Journal*, vol. 325, no. 7, pp. 409–412, 2002.

[2] R. Lininger, "Pediatric peripheral iv insertion success rates," *Pediatric Nursing*, vol. 29, no. 5, pp. 351–355, 2003.

[3] A. Zivanovic, "The development of a haptic robot to take blood samples from the forearm," *MICCAI*, p. 614620, 2001.

[4] V. Paquit, "Near-infrared imaging and structured light ranging for automatic catheter insertion," *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, vol. SPIE 6141, p. 19, 2006.

[5] D. Delpy, "Quantification in tissue near-infrared spectroscopy," *Philos Trans. of the Royal Soc. B: Biol Sci.*, vol. 352, pp. 649–659, 1997.

[6] InfraRed Imaging Systems, http://www.irimagesys.com/.

[7] Luminetx Corporation, http://www.luminetx.com/.

[8] J. Payne, "Alternative techniques for venipuncture," *The American Journal of Nursing*, vol. 72, no. 4, pp. 702–703, 1972.

[9] D. Becker, "Image processing algorithms for retinal montage synthesis, mapping, and real-time location determination," *IEEE Trans. on Biomed. Eng.*, vol. 45, no. 1, pp. 105–118, 1998.

[10] C. Tsai, "Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images," *IEEE Transactions on Information Technology in Biomedicine*, vol. 8, no. 2, pp. 122–130, 2004.

[11] A. Can, "Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms," *IEEE Transactions on Information Technology in Biomedicine*, vol. 3, no. 2, pp. 125–138, 1999.

[12] L. Gagnon, "Procedure to detect anatomical structures in optical fundus images," *Proc. of Conf. Medical Imaging 2001 : Image Processing*, pp. 1218–1225, 2001.

[13] B. Kochner, "Course tracking and contour extraction of retinal vessels from color fundus photographs: Most efficient use of steerable filters for model based image analysis," *SPIE Conference on Image Processing*, vol. 3338, pp. 755–761, 1998.

[14] I. Meglinski, "Quantitative assessment of skin layers absorption and skin reflectance spectra simulation in the visible and near-infrared spectral regions," *Physiol. Meas.*, vol. 23, pp. 741–753, 2002.