System for Deep Venous Thrombosis Detection Using Objective Compression Measures

Julian Guerrero*, *Student Member, IEEE,* Septimiu E. Salcudean, *Fellow, IEEE,* James A. McEwen, *Member, IEEE,* Bassam A. Masri, and Savvas Nicolaou

Abstract—A system for objective vessel compression assessment for deep venous thrombosis characterization using ultrasound image data and a sensorized ultrasound probe is presented. Two new objective measures calculated from applied force and transverse vessel area are also presented and used to describe vessel compressibility. A modified star-Kalman algorithm is used for feature detection in acquired ultrasound images, and objective measures of vessel compressibility are calculated from the detected features and acquired force and location data from the sensorized probe. A three-dimensional shape model of the examined vessel that includes compressibility measures mapped as colors to its surface is presented on the user interface, as well as a virtual representation of the image plane.

The compressibility measures were validated using expert segmentation of healthy and diseased vessels and compared using paired t-tests, which showed a significant difference between healthy and diseased cases for both measures. 100% sensitivity and specificity were obtained for both measures. The system was implemented in real-time (16 Hz) and evaluated using a tissue phantom and on healthy human subjects. Sensitivity was 100% and 60%, while specificity was 97% for both measures when implemented. The initial results for the system and its components are promising.

Index Terms—Deep venous thrombosis, Kalman filter-based segmentation, sensorized screening system, ultrasound image segmentation.

I. INTRODUCTION

SCREENING for deep venous thrombosis (DVT) currently forms an integral part of many postoperative procedures in modern health care facilities. Patients at high risk as well as symptomatic patients are routinely scanned in search of blood clots or thrombi within the deep veins. These thrombi may occlude venous flow, or break off from the vessel wall and cause a possibly fatal pulmonary embolism.

It has been reported that there are from 170 000 to as many as 260 000 DVT patients diagnosed and treated each year in the United States [1], and the mortality rate of DVT if left untreated is approximately 2.5% [2]. The annual incidence of a first episode of clinically suspected venous thrombosis has been estimated at 2%–4% in the general population [3].

The primary approach to DVT diagnosis is considered to be the compression ultrasound (CUS) examination [4]. Even though thrombi and blood have similar echogenicity, B-mode ultrasound images can be used to determine the absence of a thrombus in a vein. Indeed, when the examiner applies gentle ultrasound transducer pressure, a vein collapses unless there is a thrombus inside [5], [6]. In a CUS examination, the deep veins are scanned every 2–5 cm from the common femoral vein near the groin down to the deep calf veins, to the extent possible. From side-by-side ultrasound images depicting veins in both compressed and uncompressed states, the locations of thrombi can be inferred and recorded on a paper diagram.

CUS examinations are highly operator dependent [7], [8]. Screenings may take in excess of 40 min and repeat scans are often required ([2] reports repeat scans in 76% of cases). A system that can objectively characterize thrombi in a single examination would, therefore, be beneficial.

A "single ultrasound" examination strategy for DVT detection has been suggested [9]–[11]. The most recent methodology presented by Schellong *et al.* [10] includes the strict standardization of the ultrasound examination protocol and a sound training of the sonographer, with a low reported technical failure rate. The authors argue that by making the examination more objective (standardization of the examination) and by reducing user variability (sound ultrasonographer training), the diagnostic workup of patients can be reduced to a single examination.

Prior work has been targeted at characterizing thrombi using ultrasound elasticity imaging. In [12], off-line frame-to-frame motion estimation was performed on several hundred images using speckle tracking to produce displacement and strain images. Results indicate that the stiffer clots can be distinguished from surrounding tissue with this method. The method requires significant computational resources and interpretation of strain data. Its implementation in a manner that augments the conventional CUS examination without introducing delays that affect the "real-time" nature of ultrasound examinations requires further research and has not been demonstrated. Other strain imaging techinques [13], [14] have similar drawbacks.

This paper presents a new CUS approach to the diagnosis of DVT. The approach aims to make the CUS examination less operator dependent and faster in two ways: by obtaining an objective measure of venous compressibility and by recording three-dimensional (3-D) venous compressibility maps automatically.

An objective measure of vein compressibility is achieved by segmenting the vein in cross-sectional ultrasound images, and computing the vein cross-sectional area during compression, while at the same time measuring the force applied by the ultrasound transducer. This is implemented in our system in real-time, with frame-rates of 16 Hz and better, by assuming an elliptical venous cross section and using an extended Kalman filter to extract the ellipse parameters. The segmented venous cross sections are registered relative to each other in 3-D space by using an electromagnetic position sensor attached to the ultrasound transducer. The forces applied by the ultrasound transducer are also measured. The vein cross-sectional areas and transducer forces are used to define objective venous compressibility measures. The proposed system provides a 3-D venous compressibility map based on measurements and displays this to the examiner in an intuitive manner and in real-time. Its potential is demonstrated in experiments with phantoms and healthy subjects.

An objective CUS method with spatial localization would be suitable as a DVT screening system. Unlike the strain-imaging approaches surveyed above, the system proposed in this paper makes minimal changes to conventional CUS examinations, requires little additional instrumentation, and is implementable in real-time. It can be viewed as a simplified version of an elastography system, with only one single *lumped* stiffness being computed and displayed. The tradeoff involved is giving up on elasticity imaging detail for simplicity of implementation and real-time performance.

The paper is organized as follows. The compression assessment criteria for DVT are presented in Section II, Section III presents the instrumentation of the ultrasound transducer and DVT system, Section IV presents the examination protocol, including the approach taken for vein segmentation and the examiner computer interface, while Section V presents experimental results. A discussion and future work ideas follow in Section VI, with conclusions in Section VII.

II. VESSEL COMPRESSION ASSESSMENT CRITERIA

A CUS examination, as outlined in Section I, is a user-dependent procedure where an examiner not only interprets ultrasound images of vessel cross sections, but also takes into account the location of the ultrasound probe and the force applied thereat, in order to reconstruct a mental image of a patient's vasculature and possible thrombi. Broken down, the mental reconstruction task focuses primarily on determining vessel area at a given location, while applying varying amounts of pressure. An assessment of vessel compressibility is obtained from analysis of this information. Therefore, in order to develop an objective measure of DVT based on CUS, the *where*, *how hard*, and *vessel area* must be quantified, and an appropriate analysis of the data performed.

Two new distinct measures to indicate the possibility of DVT for an examined vessel segment have been developed and are presented in this paper. These measures follow from a quantitative evaluation of conventional CUS examinations, where sensor measurements and feature extraction results have replaced the examiner's interpretation, as will be described below.

The first DVT likelihood measure is called the *Transverse Area Ratio* (TAR), and is defined as the ratio of the minimum ($A_{\min}$) to the maximum ($A_{\max}$) detected transverse vessel area for a specific segment, or

$$\text{TAR} = \frac{A_{\min}}{A_{\max}}. \tag{1}$$

The TAR indicates how much the vessel area decreases under compression, as a percentage of the original transverse vessel area, and is calculated from several images. A large TAR ($\sim$100%) would, therefore, indicate an incompressible vein segment and the possibility of DVT, while a small TAR ($\sim$0%) would indicate a normal compressible vein.

The second DVT likelihood measure is a *vessel stiffness*, which, given the availability of transverse images of veins, and of applied transducer forces, can be defined as the relative change in vessel area divided by the relative change in applied transducer force. A linear fit is obtained to data points defined by the normalized vessel area (A) at a given normalized compression force (F). Several such points are obtained for each compression cycle (when the examiner presses and releases the transducer at one location) and normalized to the maximum area and force values for that cycle. The vessel stiffness measure is the slope m of the fitted line

$$A = mF + b \tag{2}$$

where the y-axis intercept b is not used. A value $m \simeq 0$ indicates venous incompressibility and possible DVT, while a healthy vein would generate an $m \simeq -1$.

It is expected that the minimum transverse area will be obtained when maximum force is applied and that the maximum transverse area will be obtained when the minmum force is applied to the patient's limb. The area is determined, in the case of the system presented herein, by a modified star-Kalman algorithm that detects the vessel contour in the ultrasound image, and approximates the vessel area using an elliptical model, as outlined in Section IV-A. Force measurements are obtained through sensors as outlined in Section III.

Both of these measures can be calculated for various vessel locations, in this manner constructing an objective vessel compressibility map, and characterizing an examined vessel.

III. SYSTEM INSTRUMENTATION

A sensorized hand-held ultrasound probe was developed in order to obtain the necessary measurements. The probe includes a 6-degree-of-freedom (DOF) force/torque sensor, and a 6-DOF location sensor.

Two aluminum shells were constructed to surround a linear 9–4 MHz ultrasound probe. The inner shell is fixed to the probe, while the outer shell is connected to the inner shell through the force/torque sensor (Nano25, ATI Industrial Automation, Inc.) at the rear, as shown in Fig. 1(a), (b). The examiner can grasp and manipulate the ultrasound probe in a regular manner, and all applied forces and torques are acquired.

The full-scale loads of the force-torque sensor are rated as 125 N for the X and Y axes, 500 N for the Z axis, and a torque of 3 N-m for all three axes. The data is acquired by a 16-bit data acquisition card (PCI-6034E, National Instruments Corp.).
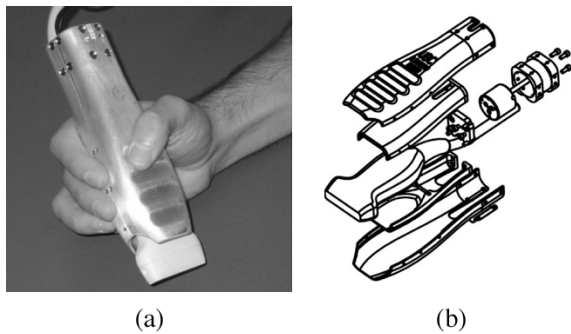
Fig. 1. Sensorized ultrasound probe used in the DVT screening system shown, with (a) examiner grasping the outer shell which surrounds the ultrasound probe itself, (b) an inner shell and a force/torque sensor.

Both shells combined add approximately 70 grams to the weight of the ultrasound probe, and the size of the ultrasound probe increases by a few millimeters.

The location of the probe is obtained by attaching an electromagnetic sensor (PCIBird, Ascension Technology Corp., Burlington, VT) to the rear of the outer shell through a 5.33″ plexiglas rod. Since the sensor is maintained in rigid relationship to the ultrasound probe and, therefore, the ultrasound image plane, the location of the image plane can be calculated using a calibrated homogeneous transformation and the location measurements. The extracted two-dimensional contours can, therefore, also be placed in a 3-D reference frame.

The static resolution of the sensor with no interference is specified as 0.5 mm for position measurements and $0.1°$ for angular measurements, with an accuracy of 1.8 mm and $0.5°$ for position and orientation, respectively. Authors have reported the successful use of these sensors in imaging applications [15], such as in 3-D vascular ultrasound of the carotid artery [16], with low reported errors ($\sim$1%) [17]. Additionally, [18] reports that the precision of an Ascension Technology electromagnetic sensor attached to 140 g of aluminum did not change for a low measurement rate (26.5 Hz). The measurement rate for our system was set to 26.5 Hz.

The complete DVT screening system has been implemented on a PC-based ultrasound machine (RP 500, Ultrasonix Medical Corp., dual-processor 1.7-GHz Intel Pentium IV), and image acquisition is done using the provided Ultrasonix libraries, which allow direct access to memory and image data, at the native frame rate. Therefore, no video acquisition hardware is necessary.

IV. EXAMINATION PROTOCOL

An examination protocol was developed for the DVT screening system, based on the CUS protocol. It is assumed that location and force values are available for each image, as well as transverse vessel contours and corresponding areas.

The user initially scans the vessel segment of interest while applying as little force as possible in order to obtain uncompressed vessel contours, which are used to build a rough 3-D model of the vessel segment as described in Section IV-B. If the measured force value is too large (>2 N)[1] the detected contour will not be used. As the user sweeps over the region of interest with the probe, contour data are obtained at about 8 mm apart and the contours are used to build the model. Conventional CUS examination protocols suggest assessing veins for compression every 1 to 5 cm [19], depending on the region being scanned.

Once a model has been constructed, the user returns the probe to the approximate locations where a contour was acquired, and performs a compression exam. During this exam, several ultrasound images are obtained while the user keeps the ultrasound probe at one location and presses down. Once a minimum force threshold (2 N) has been passed, compression data consisting of detected contours, transverse areas, and force and location measurements, is stored. All compression data must be taken close to an uncompressed vessel contour, as is verified by the location measurements. A compression assessment for the examined segment is then calculated as specified by the criteria in Section II. The area and force data obtained from uncompressed model measurements are also included in the compression assessment, to establish the uncompressed state of the vessel. Compression data is then mapped using a color code to the corresponding 3-D model location, as described in Section IV-B. This compression exam procedure is then repeated for the next location and so on, until all model contours show compression data indicating the complete vessel segment has been scanned. It is also possible for the model and compression exam procedures to be done in a single step, but will require that accurate force thresholds be determined from patient data.

A compression exam is inadequate if insufficient compression data (<5 sets) is gathered or insufficient force (<10 N) is applied at one location. The examiner must then press and release again, repeating the compression assessment at that given location.

A. Segmentation of Veins in Cross-Sectional US Images

Feature detection in ultrasound images is a difficult task because of the presence of speckle and because boundaries are not clearly delineated. Methods that incorporate knowledge about the shape or geometry of the desired feature may, therefore, improve detection.

Among the most prominent such methods are deformable contours or snakes [20]. Snakes and methods derived from them have the drawback that they require careful initialization, are difficult to implement in real-time, and often require careful tuning of parameters for convergence.

A real-time method for identifying the carotid artery contour in cross-sectional ultrasound images was presented in [21] by one of the authors. By "real-time" for vessel contour identification and tracking, we mean a method that does not interfere with operator hand motion during conventional CUS exams. It is known from biomechanics studies that voluntary hand motions do not exceed frequencies of 7–8 Hz. In CUS, much lower frequency components are expected. Nevertheless, the update frequency of 16 Hz would provide tracking of reasonably fast operator hand motion and will generate many data points during a vein compression cycle. Using a modified star algorithm [22]

---

[1]Thresholds have been defined on a trial and error basis, and are initial estimates, pending data from a clinical study.
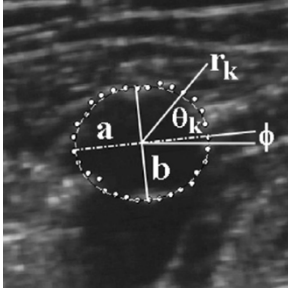
Fig. 2. Ellipse model for transverse vessel area. A vessel contour in an ultrasound image can be approximated by using an elliptical model with parameters a, b and ϕ. Each contour point can be described in polar coordinates by $r_k = f(a, b, \phi, \theta_k)$.
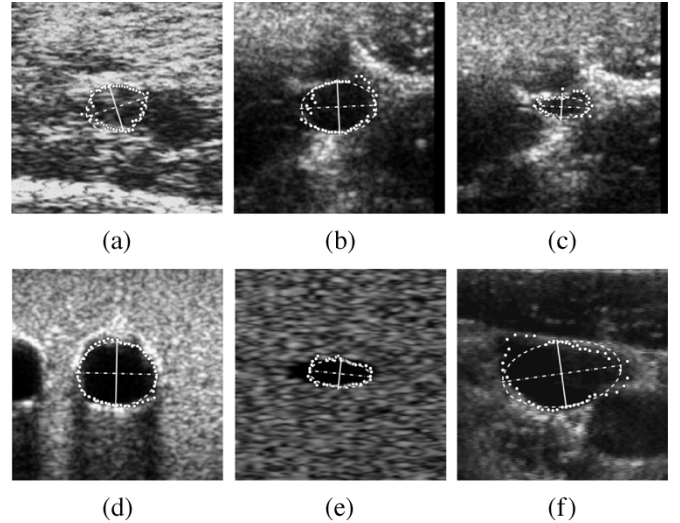

Fig. 3. Examples of contour and parameter estimation. (a), (b), (c), (f) Human, (d) phantom, and (e) simulated data with detected contours (white circles). The estimated semi-major axis a and semi-minor axis b are presented as a solid and dashed line, respectively, oriented according to the estimated ϕ.

and a Kalman filter, the algorithm in [21] assumes that the underlying vessel contour is a circle. Starting with a seed point inside the carotid artery, intensity data is detected along radii distributed uniformly in a star shape. The most probable location of the vessel lumen is detected based on a probabilistic edge detection function. The distance from the seed point to the detected vessel lumen is treated as the measurement in a spatial Kalman filter having the radius as a function of radius angle as a state. No model parameters need to be identified in this algorithm, as the Kalman filter dynamics simply state that the radius is a constant as a function of radius angle. This algorithm has been shown to perform well in detecting features even with echo drop-outs and shadowing artifacts.

It was found that the algorithm presented in [21], with a circular underlying model, does not perform well in identifying compressed veins. Indeed, because the radial contour detection algorithm fails to find the correct vessel edge when the vessel is under compression, the transverse vessel area can be underestimated and the vessel can be lost as a feature.

We, therefore, proposed an elliptical model to approximate the contour of the vein and the vessel area in a transverse ultrasound image [23]–[25]. In polar coordinates centered at the center of the ellipse, the radius r is given as a function of the radius angle θ by

$$r = \frac{ab}{\sqrt{b^2 \cos(\theta - \phi)^2 + a^2 \sin(\theta - \phi)^2}} \quad (3)$$

where the parameters ϕ, a and b are the angle and the length of the semi-major axis, and the length of the semi-minor axis, respectively, as shown in Fig. 2. While the ellipse model is a reasonable one, knowledge of the ellipse parameters for each segmented feature are required to describe the ellipse radius length as a function of the radius angle. Alternatively, the ellipse parameters can be included in an extended Kalman filter and are estimated from measurement data, as is commonly done [26]. This is the approach taken by the authors, and is described in more detail in Appendix A.

N radii are cast with equal angular spacing, from a *seed point* (x_c, y_c) previously selected inside the contour of the desired vessel. A one-dimensional (1-D) edge detector [21] applied onto each radius provides a measurement of the distance between the seed point and the contour. Based on these measurements, the

extended Kalman filter provides a sequence $\hat{r}_k$ of N points corresponding to the vessel contour and an estimate $\{\hat{a}, \hat{b}, \hat{\phi}\}$ of the elliptical model parameters. These are used to compute an approximate value $\pi \cdot \hat{a}\hat{b}$ of the cross-sectional area of the vessel. This results in a fast implementation of the DVT measures from Section II and a compression assessment that is almost immediately available.

Only the initial seed point (x_{c_0}, y_{c_0}) employed in the spatial Kalman filter mentioned above is entered manually by the operator using a mouse click in the interior of the image of a vein. A seed tracking method following the approach presented in [21] developed to accommodate feature motion calculates the seed point for subsequent frames. The vein seed point (x_{c_i}, y_{c_i}) is estimated through successive image frames i by a temporal Kalman filter using measured location data, where it is assumed that the vein center moves with constant velocity from frame to frame.

In many cases, [see Figs. 2 and 3(a)–(e)], the compressed veins can be accurately approximated by ellipses. In others, this approximation has significant errors. The underlying elliptical shape is used to approximate the vessel in order to provide a likely search region for the edge-detection algorithms to locate the boundary. This model provides a convenient way to capture global information about the feature to be extracted (the vessel contour) and gives coherence to the segmentation, but it is not crucial for it to provide small approximation errors. Indeed, Fig. 3(f) shows an example of detected vessel walls that are not elliptical.

The ellipsoidal shape was selected because it is a smooth curve simply parametrized by a differential equation with few parameters, and captures the essential feature of vein collapse—a shape with changing aspect ratio and orientation. In addition, it provides a quick way to compute the approximate area enclosed in order to provide a global fit error. Indeed, the mean error between expert segmented area and the detected ellipse area is small (~5%), as reported in Section V-B.
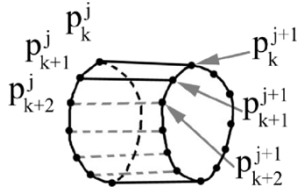
Fig. 4. Model mesh. Three-dimensional vessel model is constructed by creating polygons using detected contour points as vertices.

B. Model Construction

All valid detected contours acquired under no compression are used to reconstruct a 3-D model of the examined vessel. The coordinates of the detected contour points from an image are described in a global reference frame by using the location measurements.

The model building procedure begins once two contours have been acquired that satisfy the aforementioned restrictions (contours ~ 8 mm apart, <2 N applied force). Polygons are then created between the contours whose vertices are defined by

$$\left\{ \boldsymbol{p}_k^j, \boldsymbol{p}_{k+1}^j, \boldsymbol{p}_{k+1}^{j+1}, \boldsymbol{p}_k^{j+1} \right\}, \text{ for } \begin{matrix} j = 1 : N_1 - 1 \\ k = 1 : N_2 \end{matrix} \qquad (4)$$

where $\boldsymbol{p}_k^j$ denotes the kth index point on the jth contour, N_1 is the maximum number of contours allowed in a model, and N_2 is the number of points per contour. We also define $\boldsymbol{p}_{N_2+1}^j = \boldsymbol{p}_1^j$ in order to generate the last polygon for a pair of contours. A vessel model is represented in Fig. 4. Differences in shape between contours are small; therefore, k indexed points remain close to each other and large errors, e.g., due to rotation about the vessel axis, are not expected and have not been encountered.

For each new acquired contour $\boldsymbol{p}^N$, where $2 \leq N \leq N_1$, the distances between the seed points of each existing contour $(\boldsymbol{p}^j, j = 1 : N - 1)$ and the new contour are determined. If a new contour p^N is at least 3 mm from the closest existing contour $\boldsymbol{p}^{N-l}$, new polygons are created linking it to this contour. Otherwise, points in the new p^N and existing closest contour $\boldsymbol{p}^{N-l}$ are averaged $(\hat{\boldsymbol{p}}_k^{N-l}) = (\boldsymbol{p}_k^N + \boldsymbol{p}_k^{N-l})/2, k = 1 : N_2)$ and a new set of polygons is not created. In this manner, the user can refine and smooth out the constructed model.

Compression assessment results are assigned directly to the contour points that comprise the vessel model. This is done by assigning a color value as an extended property to each point. The color for a certain TAR/slope value is determined by a colormap or lookup table; color or greyscale can be used. The rendering and visualization package (The Visualization Toolkit (VTK), Kitware Inc. [27]) then takes this property value and the known polygon configurations, to map the corresponding colors to the surface of the appropriate polygons. Colors are blended by VTK as necessary. Other steps in the rendering pipeline, such as triangulation, are done internally by VTK.

Initially, the points are assigned a "neutral" value, indicating that no compression data is available for that section. Once a compression examination has been performed and TAR/slope values are available, the model contour closest to the compression data is found. This is done by calculating the average seed points from compressed contours, and finding the closest model seed point. The corresponding TAR/slope value is then assigned to all points on that closest model contour.
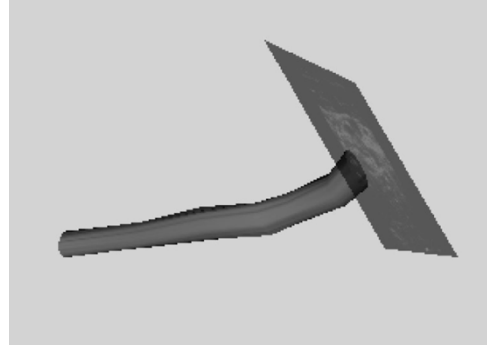

Fig. 5. Screen shot of the data display for ultrasound screening system, showing a "neutral" 3-D vessel model and the ultrasound image plane.
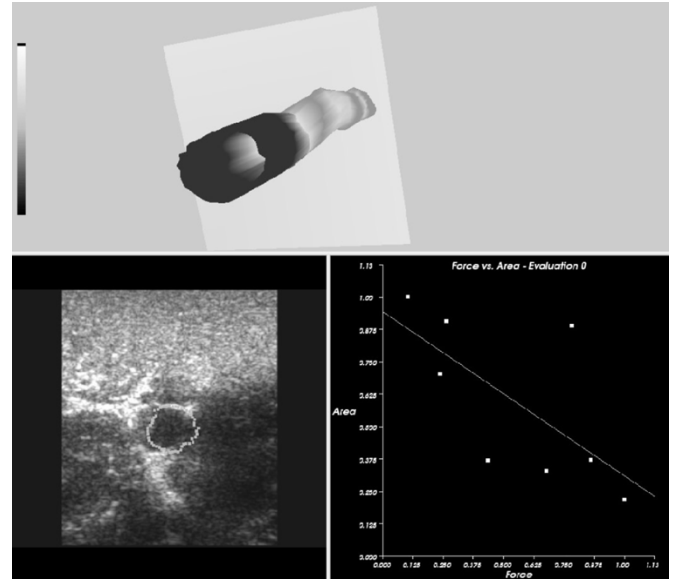

Fig. 6. Screen shot showing constructed model (top), with ultrasound data and detected contour (bottom left) and compression assessment (bottom right) for area indicated by virtual plane on the constructed model. Data from a healthy femoral vein from volunteer; dark indicates missing data.

C. Data Display

The 3-D model is displayed in real-time with a representation of the image plane in correct perspective using the location data. The scene can be viewed from any point of view, and a conventional ultrasound image display can also be presented to the examiner. Once available, compression data is mapped to the surface of each vessel segment making incompressible vessel segments easy to identify. The implemented system uses green for compressible vessels and red for incompressible vessels, although other color schemes or greyscale can be used, as is done for Figs. 5, 6, and 8(b). Additional data, such as examination adequacy and missing data, are also easily mapped in the same way, e.g., using different colors. In this manner, a summary of the venous examination is presented on the model itself.

A vessel model with no compression data, along with the image plane with the current ultrasound image data displayed on the surface is shown in Fig. 5. In Fig. 6, a user interface displays the model, the image with detected contour, and the compression assessment for the model location indicated by a virtual plane.

TABLE I
SLOPE VALUES FOR HEALTHY AND DISEASED VESSELS

	Healthy (H)	Diseased (D)
Mean (μ)	-1.18	-0.22
St. D.	0.15	0.19

V. EXPERIMENTAL RESULTS

A. Validation of Vessel Compressibility Criteria

Initial validation for the TAR and slope criteria was done with data from instructional videos [28] and videos of conventional compression examinations. While force measurements were not available for this validation, compression-release cycles were easily identified in the video, as well as the maximum and minimum transverse vessel area, making it possible to infer which image frames corresponded to those of minimum and maximum applied force.

Examples of incompressible vessels $(n = 12)$ and of healthy vessels $(n = 12)$ were identified and processed. Arterial segments were also included in this validation, and considered as incompressible vessel segments. The 8 to 20 images for each case were segmented manually (M) by the authors.

The mean value for the TAR for healthy veins (M) was 0% [standard deviation (SD) 0%], which indicates that the vessel was completely compressed, while for diseased vessels (M) the values was 71% (SD 13%), with a maximum value of 87% and a minimum value of 47%. The mean values for healthy and diseased veins are significantly different (p−value = 0, double-sided t-test, 99% CI). Sensitivity and specificity were 100% for a range of TAR cutoff values [1%, 45%].

The results for the manually obtained slope values are presented in Table I. Note that the mean value for the slope of a healthy vein μ_H is nearly -1.2, while the mean slope for a diseased vessel μ_D is approximately -0.2, which correspond to the expected values and are also significantly different (p−value = 0, double-sided t-test, 99% CI). Sensitivity and specificity were 100% for a range of slope cutoff values $[-0.8, -0.6]$.

A subset of this data was segmented using the feature extraction algorithm (E). The average value for the TAR for healthy veins $(n = 10)$ from the extracted contour (E) was 13% (SD 8%), while the mean TAR for uncompressible vessels $(n = 3)$ was 59% (SD 4%). These values still show a statistically significant difference (p−value = 0, double-sided t-test, 99% CI), despite the low sample number, as well as sensitivity and specificity of 100% for a range of cutoff values [30%, 50%]. The TAR for healthy vessels is larger when using the feature extraction algorithm because the algorithm cannot detect an ellipse with zero area.

Fig. 7 shows typical examples of transverse areas of diseased and healthy vessels and ellipse area computation.

B. Contour Detection Results

Vessel segmentation was evaluated by comparing detected contours $(n = 3810)$ on simulated ultrasound images (Field II [29]) with known elliptical features, and contours $(n = 1998)$ from real ultrasound images including cartoid and saphenous veins and arteries, and popliteal veins. Seed points were chosen
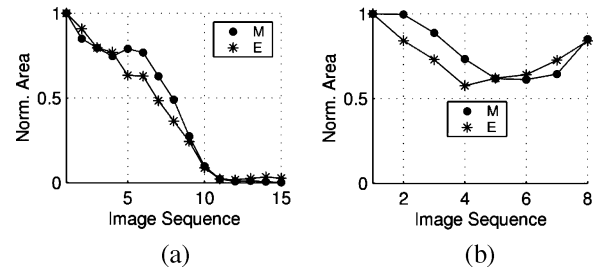

Fig. 7. TAR validation data. (a) Typical example of area versus applied force plot for a healthy vessel. (b) Typical example for a diseased vessel from manually extracted data (M) and our algorithm (E).

TABLE II
MEAN ERRORS (μ) AND STANDARD DEVIATIONS (σ) FOR VESSEL
SEGMENTATION ON SIMULATED AND REAL IMAGES

	a (%)	b (%)	ϕ (°)	FArea (%)	EArea (%)	MEr (pix)	RMSEr (pix)
Sim, μ	0.98	-5.37	-1.06	-3.11	2.66	-0.07	7.22
Sim, σ	8.74	37.8	18.9	29.7	23.8	3.46	3.64
Real, μ	NA	NA	NA	NA	-6.54	1.80	9.55
Real, σ	NA	NA	NA	NA	29.8	4.73	3.84

TABLE III
EXECUTION TIME (MS)—DVT SYSTEM

	Read Sensors	Contour/ Tracking	Model/ Comp.	Render/ Display	Total Time
Mean	1.2	23.4	10.0	10.7	62.1
Max	22.8	96.8	4,281	36.4	4,342
Min	0.2	0.01	0.01	1.2	11.9

in an area extending radially from the center of the feature to one quarter of the distance to the expert tracing. Mean (MEr) and rms (RMSEr) radial distance errors between boundaries [30] were calculated in pixels between expert tracings and detected contours, while mean errors between estimated ellipse parameters and known parameters for simulated images were generated. Transverse areas from estimated ellipse parameters were compared to known "true" areas (FArea) for simulated images, and to expert areas (EArea) for all images, and mean errors were obtained. Mean errors (μ) and standard deviations (σ) are shown in Table II.

C. System Evaluation

The system performs in real-time, including the contour extraction and tracking, the model building procedure, as well as the gathering and display of vessel compression data. Typical execution time values (in milliseconds) are presented in Table III. Data was obtained from scanning one vessel ($n = 2625$ frames processed) following the presented protocol. The current implementation of the system operates on average at about 16 Hz. Maximum execution time occurs during the model building and compression exam procedure, after all compression data has been acquired for a single location, and the compressibility measures are being calculated and acquired data is written to disk. In this case, this occurs in 3 of 2625 cycles (3 compression exams for this model).
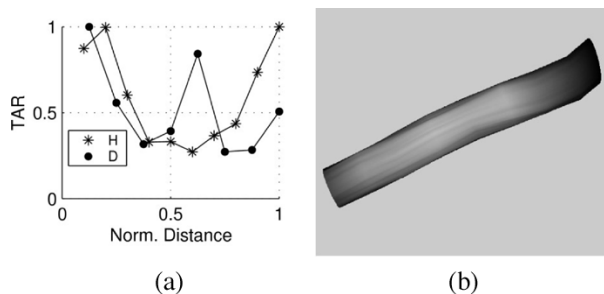
Fig. 8. Example of experimental phantom data. Typical TAR measures for a healthy (H) and diseased (D) vessel are presented (a) along the length of a phantom vessel. An example of a 3-D vessel model (b) with compression data mapped to it's surface as grayscale values (light is compressible).
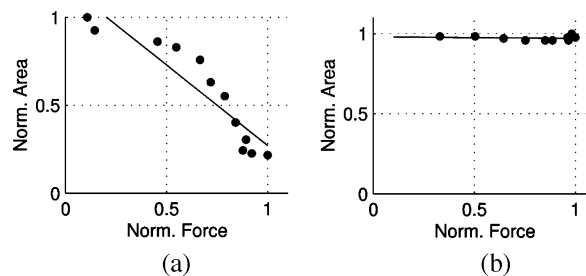

Fig. 9. Typical human compression data consisting of normalized transverse area versus normalized force is presented. (a) Vein transverse area almost disappears as force is applied, while (b) artery transverse area does not change over a range of applied forces.
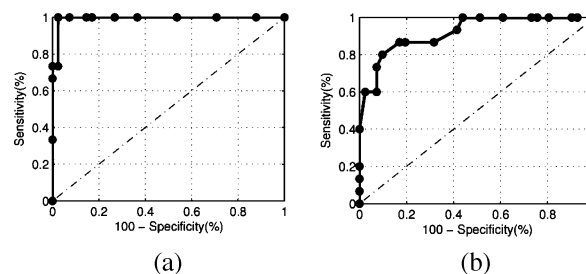

Fig. 10. ROC curves for (a) TAR and (b) slope venous compressibility criterion.

It is known that high conductivity, low permeability metals, such as the aluminum ultrasound transducer shell, can introduce errors in magnetic tracking systems due to eddy currents [18]. Alternatively, by reducing the sensor measurement rate, the effect of the eddy currents is minimized and location errors can be reduced. Additionally, the sensor manufacturer claims high immunity to distortion caused by residual eddy currents. Therefore, an extensive quantitative evaluation of location accuracy was not deemed necessary. The location accuracy was tested using a qualitative evaluation on known phantoms and was found to be sufficient for it not to interfere with the prototype demonstrated in this paper.

The complete system has also undergone several evaluations, described as follows.

1) Phantom Experiments: For developing and testing the system, a physical simulator or phantom was constructed using either polyvinyl-alcohol (PVA) cryogel [31] vessels, with inner and outer diameters approximately 8 and 10 mm, respectively, or vessels made from 6 to 8 coats of liquid latex (Coast Fiber-Tek Productions Ltd., Burnaby, BC, Canada) forming a 0.4 mm thick, 10 mm in diameter vessel. PVA has been shown to have similar physical properties as porcine vessels. These vessels were fixed within a Plexiglas container, and the interior of each vessel was accessible from the outside of the container. The tissue mimic surrounding the vessels was made using type B gelatin from bovine skin (Sigma-Aldrich, Inc.), following a previously presented method [32]. A solution of 13% gelatin by weight was prepared in water and heated to 50 °C while stirring. Afterwards 3% 50 μm cellulose by weight was added for appropriate acoustic attenuation. The mixture was cooled to approximately 30 °C and poured into the plexiglas container, surrounding the previously prepared vessels. This phantom material proved to be resilient, and of good image quality.

The phantom vessels were scanned according to the examination protocol. Fig. 8(a) presents the TAR versus phantom location for the "healthy" (H) and "diseased" (D) case. The location of the thrombus phantom is clearly evident from the high TAR value in this example.

Fig. 8(b) shows a 3-D vein model generated by scanning the healthy vein phantom. The compression information is mapped to the surface of the model as grayscale values. The midsection

of the vessel corresponds to a healthy compressible vessel while both ends, attached to fixtures in the phantom, show less compressibility. This is also observed in Fig. 8(a).

2) Human Experiments: The system underwent initial testing on healthy volunteers ($n = 3$), where several vessel segments belonging to the deep venous system were scanned in each case (total $n = 10$).

Typical results of a compression exam on a femoral vein are presented in Fig. 9(a). The results show a line fit to the data with slope near -1 and a small TAR value indicating the absence of DVT as expected.

An arterial segment was also examined, representing an incompressible vessel. The results for one compression-release cycle are presented in Fig. 9(b) and show a high TAR value and slope value near 0, clearly, indicating an incompressible vessel, as expected.

Additional TAR and slope data generated by the system from healthy volunteers was used to calculate sensitivity and specificity values for each individual cross section. Veins were considered "healthy" ($n = 41$) and arteries were considered "diseased" ($n = 15$). Receiver operating characteristic (ROC) curves were generated for different cutoff values, and presented in Fig. 10. Optimal cutoff value for the TAR as determined by the likelihood ratio (LR) was 55% with a sensitivity of 100% and specificity of 97%, while the optimal cutoff value for the slope was -0.2 with a sensitivity of 60% and specificity of 97%.

No diseased patients were included because of ethical considerations, but clinical studies including these patients will soon be underway.

VI. Discussion

Our measurement-based DVT screening system mimics a conventional CUS examination and, therefore, comprises several subsystems that perform tasks that a human examiner would traditionally perform. Additionally, a sensorized probe has been developed to permit the acquisition of necessary data, and finally two DVT measures are calculated to determine the degree of compressibility and, therefore, the possibility of DVT of an examined venous segment.

With respect to the developed measures, the large difference in TAR values from Section V-A provide a good indication of the usefulness of the proposed measure. It is also clear that lines fit to this data will produce slope values that can be used to readily distinguish diseased vessels from healthy vessels. These consistent results suggest that DVT can be well characterized by these measures. Even for the low sample sizes used, the measures for healthy and diseased vessels are significantly different. 100% sensitivity and specificity for both criteria obtained from manually segmented data further validate the measures. Moreover, both of the ROC curves for the TAR and slope criteria indicate that these measures are quite useful (large area left of diagonal) when implemented. It is believed that the lower sensitivity for the slope criterion is due to the presence of outliers created from current inability to detect zero area, and clearly indicates that the implementation is not matching its potential.

The value of these measures are dependent on the accuracy of the measurements, specifically on the calculated areas. While the mean errors for detected ellipse parameters are very small, the standard deviations are relatively large. The reported mean error of the critical measure, the detected transverse area when compared to expert segmented area is only about 5%, with standard deviation about 20%. Even so, given the very large difference between healthy and diseased values for both the TAR and slope measures obtained from manual segmentation (also shown by the 100% sensitivity and specificity for cutoff values of up to 40%), it is expected that these measures would still be quite useful even if mean area errors were on the order of 10–20%.

It is expected that a consistent over- or underestimation of transverse vessel areas will not negatively affect the compressibility assessment, while overestimation of minimum areas coupled with underestimation of maximum areas will have a minimal effect. On the other hand, overestimation of larger (maximum) areas combined with underestimation of smaller (minimum) areas could reduce the sensitivity of the developed metrics. However, the authors have observed that the current algorithm tends to overestimate smaller (minimum) areas.

The use of the additional force and location data provide the necessary information to perform an adequate and complete compression examination without which the compression assessment could not be made. The sensorized probe allows this additional data to be readily gathered, with minimal changes to existing hardware, and without sacrificing the portability of modern ultrasound systems. Because of reported high measurement accuracy and low errors in systems similar to ours presented in Section III, a quantitative evaluation of location accuracy was not deemed necessary. Qualitative results indicate location accuracy is adequate for this application.

The DVT system displays the compressibility and general shape of vessels, as well as the relative location of a thrombus, if found, and archives this data for future reference. The vessel model can be used to provide a spatial reference to the examiner and for displaying the results of the examination.

In use, compressible vessels were identified as such, and incompressible vessels such as the examined arteries provided an example of a true positive result. Additionally, it was shown that the phantom provides an adequate platform for the development and testing of the current system.

Currently patient movement cannot be accounted for by the system. This and others, such as arterial pulse, are possible sources of error. Errors are also expected from the location measurements, but these are small enough not to interfere with the prototype presented in this paper.

With respect to future work, many improvements are still possible, as well as in depth testing to validate the system.

Improvements to the feature detection algorithm are underway in order to adequately detect a completely compressed vessel. Since the feature disappears from an image under compression, it cannot be detected. This is why a noticeably larger average TAR for healthy vessels in Section V-A was found when using the segmentation algorithm, compared to a TAR of 0% for expert-segmented data. Also, bifurcating vessels cannot be tracked with the current system.

The use of other imaging modalities for characterizing the scanned vessel, such as Doppler and flow information, and elastography as this method becomes feasible for real-time applications, are also being explored.

Additionally, similar tissue characterization applications using the presented sensorized probe are easy to envision.

Extensive clinical and laboratory testing on healthy and diseased subjects is planned for validating the system. The phantom must be validated as well, by insuring that compression versus area data and ultrasound images obtained from the phantom are consistent with data obtained from humans.

An interesting extension of this system is as a platform for the real-time evaluation of elastic properties of different anatomical structures based on feature extraction and the sensorized probe.

VII. Conclusion

An experimental system and interface for the screening of DVT has been presented. Two numerical measures for the likelihood of DVT are determined by the system using automatic area computation, force sensing and sensor location. These measures along with the sensor data are presented in a 3-D display to the examiner, as well as providing a detailed record of the screening.

The system was evaluated on a custom made phantom as well as on healthy human subjects, with promising results. The area detection system and the area-based DVT detection system using the Transverse Area Ratio and slope criterion as objective characterizations of venous compression were also evaluated on videos including DVT cases and performed as expected.

APPENDIX A
VESSEL SEGMENTATION ALGORITHM

The proposed modified star-Kalman algorithm functions as follows. A *seed point* (x_c, y_c) is selected inside the desired vessel on a transverse ultrasound image. N angularly equispaced radii are then projected from the seed point to a maximum search distance r_{max}, which must be larger than the feature to be detected. The radii are sequentially scanned from $k = 1$ to N. Each radius is filtered by a median filter to improve the edge detection [23]. An estimate of the state vector $\boldsymbol{x}_k = [a_k, b_k, \phi_k]^\top$ is obtained from an extended Kalman filter, and the detected edge r_k is the output of a function evaluated using the estimated parameters at each k.

If the ellipse parameters are known we can describe the radius length r of an ellipse centered at a seed point as the output of the function (3) using the radius angle θ and ellipse parameters a, b and ϕ. If the ellipse parameters are *not* known they can be estimated using the framework of nonlinear state estimation of an extended Kalman filter [26]. The nonlinear system to which the extended Kalman filter is applied is the following:

$$\begin{cases} \boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \boldsymbol{\varsigma}_k \\ r_k = \boldsymbol{C}(\boldsymbol{x}_k) + \eta_k \end{cases} \qquad \text{(A-1)}$$

where

$$\boldsymbol{C}(\boldsymbol{x}_k) = \frac{a_k b_k}{\sqrt{b_k^2 \cos^2(\theta_k - \phi_k) + a_k^2 \sin^2(\theta_k - \phi_k)}} \qquad \text{(A-2)}$$

the state is $\boldsymbol{x}_k = [a_k, b_k, \theta_k]^\top$; k is a radius angle index, $\theta_k = 2\pi k/N$, and $\boldsymbol{\varsigma}_k$ and η_k are assumed to be sequences of white, zero-mean, Gaussian process and measurement noise, respectively, with covariances $\boldsymbol{S}_k$ and $\boldsymbol{R}_k$, respectively. The output r_k is the radius length from the seed point to the vessel boundary at angle θ_k and a function of the parameters a_k, b_k, and ϕ_k.

The estimate $\hat{\boldsymbol{x}}_{k|k}$ of the state $\hat{\boldsymbol{x}}_k$ is obtained from

$$\hat{\boldsymbol{x}}_{k|k} = \hat{\boldsymbol{x}}_{k|k-1} + \boldsymbol{G}_k \left(r_k - \boldsymbol{C}(\hat{\boldsymbol{x}}_{k|k-1}) \right) \qquad \text{(A-3)}$$

where $\boldsymbol{G}_k$ is the Kalman gain [26], and $\hat{\boldsymbol{x}}_{k|k-1} = \hat{\boldsymbol{x}}_{k-1|k-1}$. The estimated edge is generated by evaluating $\hat{r}_k = \boldsymbol{C}(\hat{\boldsymbol{x}}_{k|k})$. A system similar to (A-1) that encompasses the ellipse shape as *a priori* information in the dynamics can also be derived.

A 1-D edge detector provides the measurement residual $z_k = r_k - \boldsymbol{C}(\hat{\boldsymbol{x}}_{k|k-1})$, which is obtained by processing the brightness values along the radius r_k emanating at angle θ_k. M number of candidate points ρ_i are selected based on the results of an edge detection function, and a probability distribution function describing the actual location of the edge is constructed from these points for each radius, as in [21].

The estimation procedure is performed by consecutively traversing the contour several times, typically three, and the original seed point is used throughout. Initial values $\boldsymbol{x}_0 = \boldsymbol{x}_{0|0}$ are provided once, and data obtained from going around the contour once are used in consecutive traversals.

The vein edge locations are described by the generated output $\hat{r}_k$. The estimated ellipse parameters are defined as $\hat{a} = \hat{b}_N$, $\hat{b} = \hat{b}_N$, and $\hat{\phi} = \hat{\phi}_N$, and are used in place of a, b, and ϕ in (3) to reconstruct an ellipse centered at (x_c, y_c).

The extracted contour $\hat{r}_k$ and the estimated contour based on the ellipse parameters are compared to each other, at additional computation expense, to insure contour smoothness and consistency. The root mean squared (rms) radial distance between boundaries [30] is used as an error measure, and is computed by measuring the distance between the generated points $\hat{r}_k$ and the corresponding points on the generated ellipse. A data fit is deemed invalid if this error is larger than a predetermined threshold, typically

$$\text{error threshold} = 1.5 \cdot \sqrt{r_{max}} \qquad \text{(A-4)}$$

as determined by trial and error. The threshold must be increased in order to detect less elliptically shaped features. If a data fit is invalid, the contour detection can be repeated using a smaller r_{max} until an predetermined absolute minumum search area is reached. If data is still invalid, the contour detection has failed.

Typical parameter values and initial conditions are $N = 50$, $r_{max} = 80$, $\boldsymbol{S}_k = \text{diag}(2, 2, 0.1)$, $\boldsymbol{R}_k = 20$, $M = 5$, and $\hat{\boldsymbol{x}}_{0|0} = [40\ 40\ 0]^\top$.

REFERENCES

[1] F. A. Anderson *et al.*, "A population-based perspective of the hospital incidence and case-fatality rates of deep vein thrombosis and pulmonary embolism. The Worcester DVT study," *Arch. Intern. Med.*, vol. 151, pp. 933–938, 1991.
[2] N. Perone *et al.*, "Comparison of four strategies for diagnosing deep vein thrombosis: a cost-effectiveness analysis," *Am. J. Med.*, vol. 110, no. 1, pp. 33–40, 2001.
[3] M. C. H. Janssen *et al.*, "Diagnosis of deep vein thrombosis, an overview," *Netherlands J. Med.*, vol. 48, pp. 109–121, 1996.
[4] M. Atri *et al.*, "Accuracy of sonography in the evaluation of calf deep vein thrombosis in both postoperative surveillance and symptomatic patients," *Am. J. Roentgenol.*, vol. 166, pp. 1361–1367, 1996.
[5] J. J. Cronan *et al.*, "Deep venous thrombosis: US assessment using vein compression," *Radiology*, vol. 162, pp. 191–194, 1987.
[6] ——, "Lower-extremity deep venous thrombosis: further experience with and refinements of US assessment," *Radiology*, vol. 168, pp. 101–107, 1988.
[7] B. W. Frazee, E. R. Snoey, and A. Levitt, "Emergency department compression ultrasound to diagnose proximal deep vein thrombosis," *J Emergency Med.*, vol. 20, no. 2, pp. 107–111, 2001.
[8] American Thoracic Society, "The diagnostic approach to acute venous thromboembolism—clinical practice guideline," *Am. J. Respir. Crit. Care Med.*, vol. 160, pp. 1043–1066, 1999.
[9] J. Cornuz, S. D. Pearson, and J. F. Polak, "Deep venous thrombosis: complete lower extremity venous US evaluation in patients without known risk factors—outcome study," *Radiology*, vol. 211, pp. 637–641, 1999.
[10] S. M. Schellong *et al.*, "Complete compression ultrasonography of the leg veins as a single test for the diagnosis of deep vein thrombosis," *Thrombosis & Haemostasis*, vol. 89, no. 2, pp. 228–234, 2003.
[11] B. Wolf, D. M. Nichols, and J. L. Duncan, "Safety of a single duplex scan to exclude deep venous thrombosis," *Br. J. Surg.*, vol. 87, no. 11, pp. 1525–1528, 2000.
[12] S. Y. Emelianov *et al.*, "Triplex ultrasound: elasticity imaging to age deep venous thrombosis," *Ultrasound Med. Biol.*, vol. 28, no. 6, pp. 757–767, 2002.
[13] J. J. Mai, J. K. Tsuo, C. Pellot-Barakat, W. J. Hornof, C. Kargel, and M. F. Insana, "Vascular compliance using elasticity imaging," in *Proc. Ultrasonics Symp.*, 2001, pp. 1577–1580.
[14] J. C. Berrios and P. C. Pederson, "Deep vein thrombosis detection," in *Proc. Ultrasonics Symp.*, 1990, pp. 1567–1570.

[15] A. Fenster and D. B. Downey, "3-D Ultrasound imaging: a review," *IEEE Eng. Med. Biol. Mag.*, vol. 5, pp. 41–51, Nov./Dec. 1997.

[16] A. Fenster, D. Lee, S. Sherebrin, R. Rankin, and D. B. Downey, "Three-dimensional ultrasound imaging of the vasculature," *Ultrasonics*, vol. 36, pp. 629–633, 1998.

[17] S. W. Hughes, T. J. D'Arcy, D. J. Maxwell, W. Chiu, A. Milner, J. E. Saunders, and R. J. Sheppard, "Volume estimation from multiplanar 2D ultrasound images using a remote electromagnetic position and orientation sensor," *Ultrasound Med. Biol.*, vol. 22, no. 5, pp. 561–572, 1996.

[18] D. F. Leotta, P. R. Detmer, and R. W. Martin, "Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging," *Ultrasound Med. Biol.*, vol. 23, no. 4, pp. 597–609, 1997.

[19] B. B. Tempkin, *Ultrasound Scanning: Principles and Protocols*. Philadelphia, PA: Saunders, 1993.

[20] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: active contour models," *Int. J. Comput. Vis.*, pp. 321–331, 1988.

[21] P. Abolmaesumi, S. E. Salcudean, W. H. Zhu, M. R. Sirouspour, and S. P. DiMaio, "Image-guided control of a robot for medical ultrasound," *IEEE Trans. Robot. Automat.*, vol. 18, no. 1, Feb. 2002.

[22] N. Friedland and D. Adam, "Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing," *IEEE Trans. Med. Imag.*, vol. 8, no. 4, pp. 344–353, Aug. 1989.

[23] J. Guerrero, "Improved interface for deep venous thrombosis screening using B-mode ultrasound," Master's thesis, Univ. British Columbia, Vancouver, BC, Canada, 2002.

[24] J. Guerrero, S. E. Salcudean, J. A. McEwen, B. A. Masri, and S. Nicolaou, "Measurement-based deep venous thrombosis screening system," in *6th Annual International Conference on Medical Image Computing & Computer Assisted Intervention (MICCAI)*, Nov. 2003, pp. 214–221.

[25] ——, "Deep venous thrombosis screening system using numerical measures," in *Proc. 25th Annu. Int. Conf. IEEE EMBS*, Sep. 2003, pp. 894–897.

[26] Y. Bar-Shaloom and T. E. Fortmann, *Tracking and Data Association*. New York: Academic., 1988.

[27] The Visualization Toolkit [Online]. Available: http://www.vtk.org/ 2005

[28] S. R. Talbot, *Venous Imaging Techniques With Stephen R. Talbot, RVT*. Pasadena, CA: Davies, 1993.

[29] J. A. Jensen, "Field: a program for simulating ultrasound systems," *Med. Biol. Eng. Comput.*, vol. 34, pt. 1, pp. 351–353, 1996, suppl. 1.

[30] V. Chalana and Y. Kim, "A methodology for evaluation of boundary detection algorithms on medical images," *IEEE Trans. Med. Imag.*, vol. 16, no. 5, pp. 642–652, Oct. 1997.

[31] K. C. Chu and B. K. Rutt, "Polyvinyl alcohol cryogel: an ideal phantom material for MR studies of arterial flow and elasticity," *Magn. Reson. Med.*, vol. 37, pp. 314–319, 1997.

[32] T. J. Hall, M. Bilgen, M. F. Insana, and T. A. Krouskop, "Phantom materials for elastography," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 44, no. 6, pp. 1355–1365, Nov. 1997.

Julian Guerrero (S'00) received the B.Eng degree in electrical engineering from the Instituto Tecnológico y de Estudios Superiores de Monterrey, Cuernavaca, Mexico, in 1999, and the M.A.Sc. degree from the University of British Columbia, Vancouver, BC, Canada, in 2002, under the co-supervision of S. Salucudean and J. McEwen. He is currently working toward the Ph.D. degree in the Department of Electrical Engineering at the University of British Columbia.

His research interests are in image segmentation for medical applications and diagnosis.

Septimiu (Tim) E. Salcudean (S'78–M'79–SM'03–F'05) received the B. Eng and M.Eng. degrees from McGill University, Montreal, QC, Canada, and the Ph.D. degree from the University of California at Berkeley, all in electrical engineering.

From 1986 to 1989, he was a Research Staff Member in the robotics group at the IBM T. J. Watson Research Center, Yorktown Heights, NY. He then joined the Department of Electrical and Computer Engineering at the University of British Columbia, Vancouver, BC. Canada, where he is now a Professor and holds a Canada Research Chair. He spent one year at the ONERA-CERT (aerospace-controls laboratory), Toulouse, France, in 1996–1997, where he held a Killam Research Fellowship, and six months, during 2005, in the medical robotics group (GMCAO) at CNRS, Grenoble, France.

Prof. Salcudean is interested in haptic interfaces, teleoperation and virtual environments. He is pursuing applications to medical diagnosis and intervention and to the control of heavy-duty hydraulic machines such as excavators. He has been a co-organizer of several symposia on haptic interfaces and a Technical and Senior Editor of the IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION. He is a Fellow of the Canadian Academy of Engineering.

James A. McEwen (S'74–M'75) received the B.A.Sc. and Ph.D. degrees in electrical engineering (biomedical) from the University of British Columbia, Vancouver, BC. Canada, in 1971 and 1975, respectively.

He is currently Adjunct Professor, Department of Electrical and Computer Engineering, and Adjunct Professor of Orthopaedics, Faculty of Medicine, at the University of British Columbia. He is also Adjunct Professor in the School of Engineering Science at Simon Fraser University, Burnaby, BC, Canada. He holds numerous patents for medical devices including for automatic tourniquet systems for surgery, which he invented and developed, and for which he received the $100 000 Principal Award for Innovation in Canada from the Ernest C. Manning Foundation in 1997. He founded and served as Director of the Biomedical Engineering Department at the Vancouver Hospital Health Sciences Centre, from 1975–1990, and was one of the founders, and remains a director, of the Medical Device Development Centre in Vancouver. His main interests are in the development and evaluation of need-oriented medical technology in order to improve the quality of diagnosis and treatment, and control the costs of health care.

Dr. McEwen is a Registered Professional Engineer, a Certified Clinical Engineer, and a member of the Canadian Medical and Biomedical Engineering Society.

Bassam A. Masri received the B.Sc. and Ph.D. degrees in mechanical engineering from the Technion, Haifa, Israel, in 1988 and the University of Sydney, Sydney, Australia, in 1998, respectively.

He is Associate Professor of Orthopaedics at the University of British Columbia (UBC), Vancouver, BC, Canada, and is the Academic Head of the Division of Lower Limb Reconstruction and Oncology at UBC. He is also the Head of the Division of Reconstructive Orthopaedics at VGH. His area of practice is restricted to Adult Hip and Knee Reconstruction and Replacement and to Musculoskeletal Oncology through affiliation with the BC Cancer Agency. He is the Head of the Musculoskeletal Tumor Group for the Surgical Oncology Network at the BC Cancer Agency. He has authored numerous research and review articles and book chapters, and has lectured extensively throughout North America and Overseas.

He is a member of the Executive and Board of the Canadian Orthopaedic Association, a fellow of the American Academy of Orthopaedic Surgeons, a fellow of the American Orthopaedic Association, a member of the Knee Society and the Hip Society.

Savvas Nicolaou received the B.Sc. degree in biological sciences from Rutgers University, New Brunswick, NJ, in 1988, the M.D. degree from the Medical School, the University of Toronto, Toronto, ON, Canada, in 1992, and specialization in radiology from the University of British Columbia, Vancouver, BC, Canada.

He is an Assistant Professor of Radiology at the University of British Columbia. He is the Head of Undergraduate Teaching (VGH Radiology) for Medical Students and the Head of the Radiology Faculty Mentoring Program. His areas of research include musculoskeletal, emergency, trauma, cadiovascular, and space medicine.

Dr. Nicolaou is a member of the Golden Key National Honor Society, a member of the Canadian Medical Association, a member of the Canadian Association of Radiologists, and a member of the American Society of Emergency Radiology. He has authored numerous publications, including research and review articles and book chapters, and also has lectured throughout North America.