INTRA-OPERATIVE 3D ULTRASOUND AUGMENTATION

Stephen R. Aylward, Julien Jomier, Jean-Philippe Guyon, Susan Weeks

COMPUTER-AIDED DIAGNOSIS AND DISPLAY LAB
Department of Radiology, The University of North Carolina at Chapel Hill

ABSTRACT

We introduce an automated and accurate system for registering pre-operative 3D MR and CT images with intra-operative 3D ultrasound images based on the vessels visible in both. The clinical goal is to guide the radio-frequency ablation (RFA) of liver lesions using percutaneous ultrasound even when the lesions are not directly visible using ultrasound. The lesions locations and desired RFA sites are indicated on pre-operative images, and those markings are made to appear within the intra-operative 3D ultrasound images.

We present our current implementation, provide analyses of its components, and demonstrate its performance.

1. INTRODUCTION

We have developed an automated method for registering images based on the vasculature visible in them. This "vascular registration" method is extremely general. It is a multi-scale method, driven by the geometry of tubes, that can be applied across imaging modalities, to 2D and 3D images, and to any vascular organ. Vessels are generally densely distributed and move as their surrounding tissues move; therefore, vessels may provide better correspondence than landmarks or surfaces for tracking an organ's internal deformations. Our vascular registration method is fast and very accurate. Not only can this method be used, for example, to detect and track cancers by comparing lung CT scans over time; it can also be used to register pre-operative liver CT or MR images with intra-operative 3D ultrasound scans. The latter is the technical focus of this paper.

The goal of the work presented herein is to demonstrate an intra-operative vascular registration system that enables liver lesions, visible on pre-operative CT and yet not visible on intra-operative ultrasound, to be treated using percutaneous, ultrasound-guided, radio-frequency ablation (RFA). Traditionally, 40-50% of liver lesions seen on CT or MR cannot be treated using ultrasound guidance since they are not visible using ultrasound [3]. Via our 3D ultrasound augmentation system, pre-operative CT image markings such as lesion locations and desired RFA sites are automatically, quickly, and accurately transcribed into intra-operative 3D ultrasound images. Using this augmented intra-operative data, an interventional radiologist is able to accurately guide an RFA needle to any predefined site and thereby treat lesions that would otherwise require attempting RFA with laparoscopic guidance, performing an open surgical resection, or waiting until the lesion enlarges.

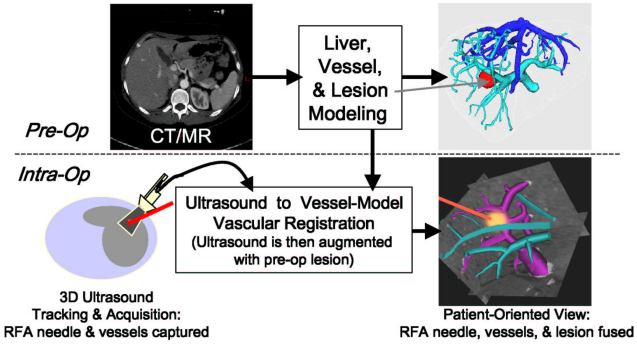

Fig. 1. Pre- and Intra-operative phases of the system.

The ultrasound augmentation system has pre- and intra-operative components (Fig. 1). In this paper we review related works and discuss our methods for pre-operatively modeling a liver's parenchyma, lesions, and vessels. Evaluations of the accuracies of these methods are interspersed. We then highlight the intra-operative components of our system: ultrasound tracking and image registration. We conclude with results that illustrate the operation and accuracy of the entire system.

2. BACKGROUND

Percutaneous RFA of hepatic lesions is performed by initially localizing the lesion, usually using ultrasound. An RFA needle is then advanced through the skin, into the lesion, and to a desired site. An electrical current is then sent through the needle and into the tumor. Needles can pro-

This system builds upon The National Library of Medicine's Insight Toolkit (ITK). This work is funded in part by an equipment grant from Dell Computer Corporation, the Whitaker Foundation (Aylward-RG-01-0341), the National Library of Medicine Insight Toolkit (Aylward-N01-LM-0-3501), and the National Cancer Institute (Bullitt-R01-CA67812)

duce up to a 7 cm burn which (with 1 cm margins for effectiveness) enables the treatment of a 5 cm spherical lesion with a single burn. For larger lesions, multiple overlapping burns are required. Pre-operative planning and ultimate needle positioning is of the utmost importance. If burns are poorly placed, then residual tumor will require repeat intervention. Our radiologists report seeking to position RFA needles within 5 mm of the intended treatment site. Other groups have stated similar tolerances [3].

The key technologies of intra-operative ultrasound augmentation are vascular modeling and vascular registration. There are three basic approaches to vessel modeling: centerline modeling and radii estimation, spatial filtering, and voxel labeling. Our philosophy is that centerline modeling should be the basis of vessel modeling. Centerlines of tubular objects such as vessels can be estimated by integrating over a larger image extent than is used by most spatial filtering, voxel labeling, and edge detection methods. Using a larger extent provides insenitivity to noise, stabilizes subsequent radii estimation, and enables the extraction of small, faint vessels. A centerline and radii representation is required by our vessel-model-to-image registration method.

The three general approaches to image registration are: image-image, feature-feature, and feature-image matching. Image-image methods are exemplified by the mutual information metric. Feature-feature methods include iterative closest point and landmark methods. Feature-image methods are the focus of Petra van den Elsen's work and our own work [1]. For ultrasound to MR/CT registration, a feature-feature image registration method is proposed in [4]. Vessel voxels are identified in both images and the distance between those sets of voxels is minimized. In [3] a deformable surface method is used to align ultrasound images and quantify liver motion due to breathing.

Feature-image registration methods have great potential for intra-operative guidance. Time can be spent to develop accurate models that can be repeatedly and rapidly matched with intra-operative images to account for patient movement, e.g., respiration. The details of our system's pre- and intra-operative phases are presented next.

3. PRE-OPERATIVE MODELING

The pre-operative phase produces models of the liver, its lesions, and its vessels. These segmentations are only performed on the pre-operative data. Emphasis is placed on accurate vessel modeling since those models drive the intra-operative registration process.

To demonstrate and quantify these methods for this paper, we use CT scans from potential donors for living related liver transplantation and from a home-made phantom. The phantom is a gallon tub of gelatine in which Soba noodles are suspended to simulate vessels and small dough balls are used to represent lesions. The donors and the phantom are imaged using a Siemens Somatom Plus CT scanner to create volumes having 1x1x3 mm voxels. We use B-spline interpolation across the slices to construct isotropic voxels. For the potential donors, a scan is acquired 30 seconds after a contrast bolus to highlight the portal venous system in the liver. A second scan is acquired 30 seconds after the first to image the hepatic venous system in the liver.

Livers and lesions are modeled using a semi-automated method. In CT data, the parameters of the following processes are fixed for all patients. (1) Voxels spatially connected to a specified starting point and having intensities within specified thresholds are identified. Simple heuristics are used for automated seed-point selection. (2) The connected component is eroded and then dilated to remove pertrusions. (3) The main component is then dilated and then eroded to fill-in small holes. The result is smoothed mask of the main component.

To determine the speed and accuracy of our segmentation system we CT scanned three surgical gloves and three balloons that contained known amounts of water. The semi-automated system required approximately one minute per item and estimated volumes within 6% of ideal. Hand contouring required 10 minutes per item and estimated volumes within 11% of ideal. Additionally, we have used this method to model the livers of over 70 potential donors. For some data it is necessary to edit the segmentations; however, the time required and the inter- and intra-user variability for such editing is minimal.

Our vessel segmentation method operates by performing a multi-scale extraction of the centerline of a vessel and then estimating the radius of the vessel about that centerline. See [2]. In summary, the method extracts the representation of a vessel in three steps:

1. Seed points for initiating the vessel extraction process are specified automatically. We use the liver model (defined above) to delineate a region of interest. The 0.1% brightest voxels in that region are used as seed points. An initial scale/radius of 3 mm is used to initiate the multi-scale centerline traveral process at these seeds. The method's performance is not statistically significantly dependent on the seed location along a vessel or the initial scale.

2. From the seed, a conjugate gradient ascent is used to reach the local intensity/height ridge that is the centerline of the tube. The automated traversal of the ridge progresses as follows: (a) The plane normal to the ridge is approximated by the eigenvectors of the scaled Hessian and that plane is shifted one-fifth of a voxel along the ridge's approximate tanget direction. Assuming the centerline varies smoothly, the ridge will pass through that shifted normal plane. (b) The local maximum in intensity in that plan is located. This is the next ridge/centerline point. (c) At fixed intervals during this traversal, the radius of the tube is estimated by find-

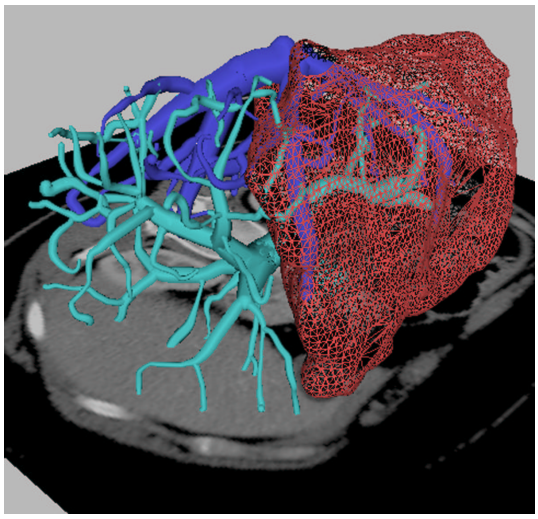

Fig. 2. Models formed pre-operatively. Portal and hepatic veins and left lobe of liver are shown.

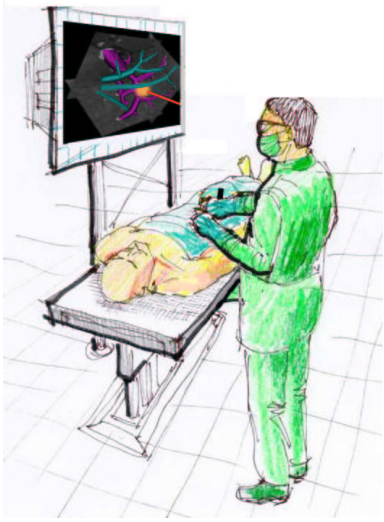

Fig. 3. The clinical setup being deployed.

ing the maximum through scale of a medialness function (see the next step).

3. Radius estimation is stabilized by the centerline. Medialness functions respond maximally when applied at the center of an object and at a scale proportional to the object's width. Centered at each point along a centerline, we search for the local maximum of medialness through scale with respect to the maximal scale at the previous centerline point.

We have conducted extensive tests to quantify the accuracy, speed, and automation (reliance on initial parameters) of the centerline and radii modeling methods [2]. For 200 randomized simulated vessel extractions in a simulated image with noise characteristics similar to those of CT, average time to extract 20 voxels of centerline was about 0.3 seconds, centerline representation average error was less than one voxel, maximum error was less than two voxel, 90% of the centerline points were within one voxel of ideal, and an optimally difficult branch point was passed nearly 70% of the time. Regarding automation, the centerline extractions accuracy was not statistically significantly affected by its starting parameter values (seed scale and seed location). For the radii estimations, for 200 random extractions of a small, tortuous vessel from an MRA image, no two corresponding radius estimates differed by more than one-fifth of a voxel. Illustrations of typical results from liver and vessel modeling are provided in Fig. 2. These liver and vessel models were generated as part of our living related donor liver transplant planning research.

4. INTRA-OPERATIVE PROCESSING

The proposed clincal setup of the system is illustrated in Fig. 3. The intra-operative methods seek to balance speed and accuracy. Additionally, the vascular registration method inherently provides a quantification of the quality of its results so that registration failures can be automatically detected, reported to the clinician, and retried; thereby minimizing risk to the patient.

The ultrasound system is a Voluson 530D by Medison, Inc. Volumes are recorded by holding the probe stationary for 7 seconds while the transducer internal to the probe is automatically swept. The volumes are resampled to cartesian 0.5x0.5x0.5 mm voxels. Standard mode (non-doppler) scans are acquired. Our goal is to process these data in under 7 seconds - registration initialization is critical.

To initialize the vascular registration method we use a calibrated magnetic tracker to record the ultrasound position in the operating room. Vascular image registration is required to account for the 10+ mm of respiration-induced liver movement [3].

Our vascular registration method must account for vascular network changes and non-rigid deformations between pre- and post-operative data: (1) The number, lengths, and radii of vessels visible in the images may differ and hence the vascular network will appear to change between images. (2) Images may only partially overlap. (3) Because of patient movement and surgical procedures, groups of vessels within and across organs may undergo non-rigid deformations with respect to location, path, and radius.

Our method is a rigid registration technique. It is formulated as a transformation of point x in a source image into the coordinate space y of a target image. A rigid transformation occurs as a rotation matrix multiplication plus a translation $y = xR + o$ where R is a Euler matrix parameterized by α, β, and γ as rotations about the z, y, and x-axes respectively and where offsets $o = [o_x, o_y, o_z]$. Extensive

analyses and visualizations of the parameter space of our registration metric are in [1].

The registration metric quantifies how well a rotation matrix and offset vector align two vascular images. The metric is based on the fact that vessel centerlines are scaled intensity ridges in the image; therefore, when two vascular images are aligned, the centerline points in one will map to bright points in the other, thereby maximizing our metric, a weighted sum of the scaled intensities of the target image at the transformed points:

$$\mathrm{F}(R,o) = \frac{1}{\sum_{j=1}^{n} w_j} \sum_{i=1}^{n} w_i I_{\kappa r_i}(x_i R + o) \qquad (1)$$

This metric is controlled by the sampling (n) of the centerline points (x_i), the scaling (κr_i = the scale of measures is proportional to the radius of the sample) of the image data (I), and the weighting (w_i) of the centerline points. When these values are fixed, failed registrations are clearly indicated by a poor final metric value. [1] In summary:

Sampling: We use a coarse-to-fine optimization strategy that progresses from n =one-twentieth to n =one-tenth of the centerline points extracted during height-ridge traversal. Any points whose medialness (vessel extraction step 3) is less than 0.2 is rejected for use in the metric.

Scaling: By default, the local scaling of the image is equal to the radius of the tube $(\kappa = 1)$. A coarse-to-fine strategy could be implemented using κ so that an initially large, smooth capture region is associated with each sample.

Weighting: Sample points are weighted based on their radius. Points with a small radius are affected by image noise and therefore are demoted. This weighting can be modified during coarse-to-fine registration.

Because of intensity irregularities along a centerline, the derivatives of this metric may induce shifts along a vessel; ideally shifts are only produced normal to a vessel, e.g., horizontal tubes should be limited to inducing vertical shifts. To implement this, at each centerline point, the gradients influence is limited to the direction normal to the tube in the original data. Additionally, we adjust for the principal orientation of the tubes in a network so that the transformation parameters gradients do not have an orientation bias, i.e., if most tubes are horizontal, the system should not be unduly biased towards vertical transformations. Lastly, we explicitly solve for $do, d\alpha, d\beta$, and $d\gamma$ during optimization.

To test registration accuracy, we generated 100 random transformations that mis-aligned a set of portal-phase vascular models and their associated hepatic-phase images up to ± 10 voxels (± 1.25 cm) and ± 0.1 radians (± 5.73 degrees). We then applied our optimization strategy and compared each of the final registration parameters" values with the mean final registration parameters" values. Mean offset error was < 0.3 mm, max offset error was < 1.3 mm in any direction, mean rotation error was < 0.01 radians, and max rotation error was < 0.05 radians in any direction.

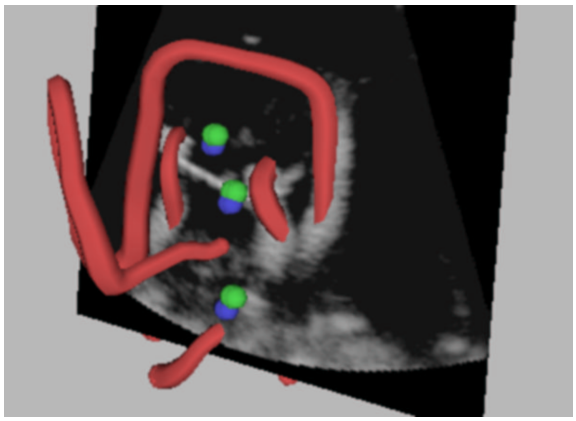

Fig. 4. CT-base noodle models registered with ultrasound data. Balls from CT and ultrasound data overlap - centers differ by less than 2.3 mm.

We tested the full augmentation system using the home-made phantom. The pre- and intra-operative methods were applied as presented. We then calculated the centers of mass of three 5x5x5 mm dough balls in the CT and ultrasound data. We then transformed the CT dough balls into the ultrasound data using the transformation defined via tracked ultrasound and vascular registration. Fig. 4 illustrates the results. The mean distance between any two corresponding centers of mass was 2.3 mm and the maximum distance was < 2.9 mm. This error is well within our desired 5 mm margin of accuracy. Our work is now focusing on integrating this system using ITK (the NLM's Insight toolkit), conducting performance studies in the lab, and then transitioned the system to the clinical for research. See http://caddlab.rad.unc.edu.

5. REFERENCES

[1] S. R. Aylward, S. Weeks, and E. Bullitt, "Analysis of the parameter space of a metric for registering 3d vascular images," in *MICCAI*, Utretch, The Netherlands, 2001

[2] S. R. Aylward and E. Bullitt, "Initialization, Noise, Singularities, and Scale in Height Ridge Traversal for Tubular Object Extraction" *Accepted IEEE Transactions on Medical Imaging*, 2002

[3] A. P. King, J. Blackall, G. Penney, D. Hawkes, "Tracking Liver Motion Using 2D Ultrasound and a Surface Based Statistical Shape Model," in *MMBIA*, Hawaii, 2001

[4] B. P. Porter, D. Rubens, J. Strang, J. Smith, S. Totterman, K. Parker, "3D Registration and Fusion of Ultrasound and MRI Using Major Vessels as Fiducial Markers," *IEEE Transactions on Medical Imaging*, 20(4) 2001