Published in final edited form as:

Int J Comput Assist Radiol Surg. 2015 November ; 10(11): 1777–1792. doi:10.1007/s11548-015-1210-5.

Toward a generic real-time compression correction framework for tracked ultrasound

Thomas S. Pheiffer^{1,2} and Michael I. Miga^{1,3,4}

¹ Department of Biomedical Engineering, Vanderbilt University, 5824 Stevenson Center, Nashville, TN 37232, USA

² Present Address: Siemens Corporation, Corporate Technology, Princeton, NJ, USA

³ Department of Neurological Surgery, Vanderbilt University Medical Center, Nashville, TN, USA

⁴ Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center, Nashville, TN, USA

Abstract

Purpose—Tissue compression during ultrasound imaging leads to error in the location and geometry of subsurface targets during soft tissue interventions. We present a novel compression correction method, which models a generic block of tissue and its subsurface tissue displacements resulting from application of a probe to the tissue surface. The advantages of the new method are that it can be realized independent of preoperative imaging data and is capable of near-video framerate compression compensation for real-time guidance.

Methods—The block model is calibrated to the tip of any tracked ultrasound probe. Intraoperative digitization of the tissue surface is used to measure the depth of compression and provide boundary conditions to the biomechanical model of the tissue. The tissue displacement field solution of the model is inverted to nonrigidly transform the ultrasound images to an estimation of the tissue geometry prior to compression. This method was compared to a previously developed method using a patient-specific model and within the context of simulation, phantom, and clinical data.

Results—Experimental results with gel phantoms demonstrated that the proposed generic method reduced the mock tumor margin modified Hausdorff distance (MHD) from 5.0 ± 1.6 to 2.1 ± 0.7 mm and reduced mock tumor centroid alignment error from 7.6 ± 2.6 to 2.6 ± 1.1 mm. The method was applied to a clinical case and reduced the in vivo tumor margin MHD error from 5.4 ± 0.1 to 2.9 ± 0.1 mm, and the centroid alignment error from 7.2 ± 0.2 to 3.8 ± 0.4 mm.

Conclusions—The correction method was found to effectively improve alignment of ultrasound and tomographic images and was more efficient compared to a previously proposed correction.

Thomas S. Pheiffer thomas.s.pheiffer@vanderbilt.edu.

Conflict of interest Thomas Pheiffer and Michael Miga declare that they have no conflict of interest.

Ethical standard All procedures followed were in accordance with the ethical standards of the responsible committee on human experimentation (institutional and national) and with the Helsinki Declaration of 1975, as revised in 2008 (5).

Keywords

Ultrasound; Registration; Compression; Finite element method; Biomechanics

Introduction

Ultrasound is ubiquitous as an interventional imaging modality and is commonly used to assess the location and geometry of disease intraoperatively. An inherent problem with this role is the shape distortion of visualized tissue structures introduced by the probe pressure exerted. It is widely recognized that relatively large tissue compression can occur in soft tissue anatomy, e.g., the liver or breast. As a result, compression can obfuscate geometrical and locational measurements of subsurface targets such as tumors. This is particularly a problem for image-guided interventions, which rely upon tracked ultrasound to provide intraoperative spatial measurements of structures taken during an intervention and then compared to their co-registered preoperative imaging data counterparts. Nonrigid tissue compression is a primary cause of misalignment and shape distortion with these other sources of information. As image-guided navigation strategies in soft tissue environments continue to be developed, methods of correcting the tissue deformation from routine ultrasound imaging are necessary in order to ensure that all of these data are in a consistent spatial arrangement.

There are several methods described in the literature for performing compression correction. A common approach is to utilize the intensity information in the ultrasound images to perform a nonrigid intensity-based registration with positional tracking of compressed images over a range of compression states [1,2]. One drawback of this method is that it requires a stream of ultrasound images, and intensity-based registration for ultrasound is a challenging task in practice. For example, in [1], they demonstrated a method to correct for compression using correlation of a stream of radio frequency (RF) or amplitude frames, and although the method performed well in a phantom dataset, the authors noted its reliance on good image quality as well as the possibility of correction drift when compression estimates are accumulated across a large sequence of images. Another method of correction is to use a mechanical model of the tissue in order to estimate the subsurface tissue displacements caused by the interaction of the probe with the tissue surface. One group proposed using a force measurement apparatus to provide force boundary conditions to a tissue model [3,4], although force boundary conditions require some prior estimate of absolute material properties for the tissue. We recently proposed an alternative method which utilizes a biomechanical model-based correction which is driven by displacement boundary conditions provided by the position of a tracked ultrasound probe within a co-registered patient-specific organ surface from preoperative tomograms [5]. This method was shown to reduce ultrasound compressional error of nearly 1cm to approximately 2–3mm.

There is a subset of image-guided procedures for which preoperative tomographic image volumes are not commonly acquired, or the volumes are acquired with the patient in a much different presentation than the operative state. This can be the case in open liver surgery, for example, in which there is often significant manipulation of the organ by the surgeon

leading up to the surgical presentation of the tissue. Therefore, a method of compression compensation that does not rely on a preoperative model would be more valuable. In addition, it is often the case that subsurface structures may be necessary for enhancing image-to-physical registration, and it is easily seen that there are implications if subsurface deformation is not addressed in registration frameworks. Provided with at least some form of intraoperative measurement of compression, subsurface structures could be uncompressed to give true shapes in physical space. These true subsurface shapes could then be used in combination with surface information to compute a combined image-to-physical registration. An example of this would be a registration framework that used a surface point cloud from a laser measurement device and subsurface structures like a tumor [6] or perhaps blood vessels [7]. With these possibilities in mind, our initial goal was to create a compression correction method, which utilizes a generic approach that is independent of tomographic imaging and requires no registration to a preoperative surface. This correction is intended to produce ultrasound images in which the tissue structures can be rendered in their uncompressed biomechanical state so as to provide more accurate shape measurements or as a source of intraoperative data for geometric comparisons. The second goal was to compare the new method with the method previously described in [5] by deployment in phantoms and clinical data.

Methods

The compression compensation method is one step within a pipeline for image guidance using tracked ultrasound. A brief description of the pipeline is displayed in Fig. 1 to provide context for the new method and to indicate how it differs from the previous correction method. In summary, the input data to the generic method are the original ultrasound images showing compressed tissue, a calibrated and tracked ultrasound probe, a surface from some other intraoperatively tracked instrumentation, and a biomechanical tissue model that is used to warp the ultrasound data to show the tissue in an uncompressed state. The previously reported patient-specific method also requires an additional registration to preoperative imaging to perform this correction. Each of the correction steps for workflow in Fig. 1 is described in more detail below.

Preoperative imaging and patient model

Image-guided interventions often begin with acquisition of high-resolution CT or MR image volumes prior to the procedure. The patient-specific compression correction method utilized a patient model created from these images. However, these data are unnecessary for the proposed generic method. The specific details of the model creation for the phantom and clinical studies are described later in the respective experimental sections.

Intraoperative data collection

The intraoperative data consist of ultrasound images, which are captured onto a host computer, along with position and orientation information from an optically tracked target that is attached to the ultrasound probe. After a calibration procedure (described in the phantom experiments), the pose of the ultrasound images is therefore known in the

coordinate system of the tracking device. In addition, other tracked devices are present which also provide physical measurements in the same coordinate frame, such as a tracked stylus or laser range finder devices.

Compression correction

The goal of this work was to reduce the localization error of soft tissue subsurface structures as visualized by ultrasound imaging arising from soft tissue deformations exerted by the probe itself. We recently presented a method in [5] utilizing both probe tracking information in combination with a co-registered patient model in order to estimate the compression depth of the probe into the tissue during insonation and then to use that depth to correct the tracked ultrasound image poses using a biomechanical model-based approach. The novel method that we now propose does not use a patient-specific model derived from preoperative imaging, but instead uses a generic model to drive the correction as shown in Fig. 2. In order to rigorously compare the two methods, a summary of the patient-specific method is described in the next subsection, followed by a description of our novel generic method.

Patient-specific correction

The patient-specific compression correction method presented in [5] made several key assumptions in order to compute a reasonable correction. The primary assumption was that intraoperative ultrasound data could be aligned to the preoperative imaging with an initial rigid registration. The second assumption was that the tissue presentation in terms of mechanical state in the preoperative imaging was similar to the intraoperative presentation, in the absence of tissue manipulation by the ultrasound probe. These two assumptions led to the next assumption, which was that in the presence of tissue compression by the ultrasound probe, the tip of the probe would be some distance below the surface of the co-registered patient model depending on the magnitude of compression.

The strategy in [5] was to use the position of the probe tip within the tissue model in order to estimate the tissue deformation, resulting from application of the probe and then use these data to correct every acquired ultrasound image. The geometry of the probe tip was constructed by scanning it with the LRS, thus providing a digital probe surface model. The probe was tracked in physical space using a tracked target attached to the probe. The next assumption was that during imaging the user would press the probe only in the depth direction of the ultrasound plane into the tissue, with no lateral or out-of-plane movement (no dragging effects). This assumption was made to simplify the next step of the correction, which was to assign boundary conditions to the biomechanical tissue model using the pose of the probe geometry. Assuming purely depth compression, the model surface nodes directly above the digital probe surface were assigned Dirichlet boundary conditions corresponding to the compression vectors calculated from the patient surface to the probe surface. The rest of the patient model was assigned a set of initial boundary conditions based upon the surgical plan and prior knowledge of the patient presentation. In the case of the liver phantom, the inferior surface was fixed as it was on a rigid conforming base, and the superior surface was set to stress-free boundary conditions and allowed to freely deform. In the case of the clinical brain case, the mesh nodes corresponding to the craniotomy region were set as stress-free, the base of the brain was set as fixed, and the rest of the brain nodes

were set to have zero displacement in the normal direction but stress-free in the tangential directions to allow for slipping along the skull. These boundary conditions have been used in brain models previously for model-updated image-guided surgery [8,9].

After the assignment of boundary conditions, the model was solved for 3D displacements over the entire mesh to estimate the probe-deformed state of the tissue. The model used in both methodologies was the standard 3D Navier–Cauchy equations for the displacement field:

$$\frac{E}{2(1+\nu)}\nabla^2 u + \frac{E}{2(1+\nu)(1-2\nu)}\nabla(\nabla \cdot u) = 0 \quad (1)$$

where E is Young's modulus, ν is Poisson's ratio, and u is the 3D displacement vector at a point in the tissue. The partial differential equation is solved within a finite element method framework using the Galerkin weighted residual technique with linear basis functions. The system of equations that solves for the displacement vectors at every node in the mesh can be written as:

$$[K]\{u\} = \{f\} \quad (2)$$

where K is the global stiffness matrix, u is the vector of nodal displacements, and f contains the contributions of boundary conditions. For each ultrasound image to be corrected, this system of equations is constructed and solved for the nodal displacements, which satisfy static equilibrium for the supplied boundary conditions. These displacements are then reversed and interpolated onto the tracked ultrasound data, which was then deformed with this 3D displacement field to an estimate of its state in the absence of compression. We should note that there are important implications to the nature of this patient-specific computation with respect to encumbrance that will be discussed in comparison with our generic model in the next section.

Proposed generic correction

The first difference between the generic correction and the patient-specific correction is that instead of a patient-specific mesh constructed from preoperative imaging and registered to intraoperative space, the generic method instead uses a pre-constructed block mesh (see Fig. 2), which is calibrated to follow the tip of the tracked ultrasound probe. The most important consequence of this framework is that the generic method only requires a sparse intraoperative measurement of tissue compression in order to provide a model correction, rather than a registration to preoperative imaging. This could be either provided by having separate digitization of the surface in physical space (e.g., a laser range scan of the surface of interest) or would need a trigger to track ultrasound position once in contact with the tissue. In this work, we have chosen the former rather than latter methodology. Lastly, we should note that a pre-computed mesh in this instance is possible and offers distinct computational advantages that are described later below.

The block mesh calibration procedure simply requires the alignment of the top of the ultrasound image with the center of one side of the mesh, and of the image plane itself with

the plane through the center of the block. The pose of the generic block mesh thus is defined by the same tracking information which defines the pose of the ultrasound image, i.e., the optically tracked attached target in this case. In this realization, the general strategy is to acquire intraoperative measurements of the undeformed tissue surface using an LRS and then use that surface in conjunction with the location of the ultrasound probe to estimate the depth to which the tissue was compressed. This depth is computed by casting rays down from each point of the LRS cloud in the depth direction and finding the average length of the ray segments which intersect with the tracked probe tip surface. This depth is then used to assign Dirichlet boundary conditions to the block mesh in a similar manner as the patient-specific correction, although with a slight difference. The initial boundary conditions for the mesh in this method were assigned such that the far-field face of the block was fixed, and the superior and side surfaces were stress-free. This is one of the key assumptions of this work; far-field deformations are near zero. This certainly represents a potential source of error depending on the angle of insonation, the location on the organ, and the size of the organ. One way to mitigate this is to pre-compute multiple tissue blocks with differing far-field lengths (see Fig. 3). Similar to how ultrasound transducers can have different fields of view that can be selected by the user, we envision that a user could select a correction far-field assumption based on in vivo conditions at the time of interrogation. This is a central reason why we examine the effects of differing far-field assumptions in the experiments described further below.

Nevertheless, after assignment of boundary conditions, the model is solved for 3D displacements in the block of tissue, and the displacements are reversed and interpolated onto the ultrasound data to perform the correction. The model construction is governed by the same constitutive equations given by (1) and (2). However, there are several advantages that the generic correction offers compared to the patient-specific model. With respect to the patient-specific method, the global stiffness matrix, K , would need to be reconstructed whenever different regions of the organ boundary were engaged for imaging, i.e., with the application of Dirichlet boundary conditions at different nodes based on localization, the stiffness matrix would need to be altered, such as from a displacement to a stress-free condition or vice versa. While all equations for every boundary node could be stored to prevent repeating element assembly, preconditioners would likely need to be recomputed followed by iterative matrix solutions. This would be an expensive process, especially if trying to achieve framerate updates. One strategy would be to use force-based boundary conditions, which would allow for a great deal of pre-computation but would require accurate measurement of applied force as well as material properties of the tissue [3,4]. With respect to the generic correction, however, the type of boundary condition assigned to each boundary node will always remain the same, as each correction proceeds by merely altering the magnitude of the displacement boundary conditions on the top of the block mesh. Thus, it is possible to pre-compute K and reuse it for each correction whenever f is updated in a simple matrix multiplication:

$$\{u\} = [K]^{-1} \{f\} \quad (3)$$

Another property of the generic method offers a further computational speedup. In order to correct the ultrasound data, only the model solution at a plane of the mesh which corresponds to the ultrasound image plane is actually needed. The computations solving for the rest of the mesh node displacements are not essential, but only their influence on the localization information within the slice itself is needed. This makes it desirable to somehow eliminate the computational burden of those nodes from the system of equations during surgery. This can be accomplished through the method of condensation, which results in a smaller system of equations that can be solved much more rapidly [10]. The first step in this process is to carefully arrange the ordering of the mesh node indices to ensure that the first N equations belong to the nodes lying on the ultrasound plane, as well as any nodes on the top surface which are assigned varying amounts of compression boundary conditions. Assuming this ordering, the equation from (2) can be rewritten as a block matrix system where the subscripts p and a indicate the plane nodes and all other nodes, respectively:

$$\begin{bmatrix} K_{pp} & K_{pa} \\ K_{ap} & K_{aa} \end{bmatrix} \begin{Bmatrix} \mathbf{u}_p \\ \mathbf{u}_a \end{Bmatrix} = \begin{Bmatrix} \mathbf{f}_p \\ \mathbf{f}_a \end{Bmatrix} \quad (4)$$

The block matrix system in (4) can be rearranged to a form involving only the displacement solution of the plane nodes, $\mathbf{u}_p$:

$$[\check{K}_{pp}] [\mathbf{u}_p] = \{\check{\mathbf{f}}_s\} \quad (5)$$

where

$$[\check{K}_{pp}] = [K_{aa}] - [K_{pa}] [K_{aa}]^{-1} [K_{ap}] \quad (6)$$

$$\{\check{\mathbf{f}}_s\} = \{\mathbf{f}_s\} - [K_{pa}] [K_{aa}]^{-1} \{\mathbf{f}_a\} \quad (7)$$

The modified stiffness matrix given in (6) represents a transfer of the displacements from all nonplane nodes to the plane nodes which are the primary concern and maintains the volumetric nature of the model. Using this stiffness matrix offers significant computational benefit because it is a fraction of the size of the full K matrix. It can be similarly pre-computed and stored for very fast solutions of the $\mathbf{u}_p$ vector of plane node displacements. In addition, given the careful ordering of the node indices explained above and the assignment of initial boundary conditions, it will also be the case that all nodes in the $\mathbf{f}_a$ vector will always be assigned either zero stress or zero displacement boundary conditions. Given (7), this implies that changes in the compression depth during imaging will result in simple reassignment of the values in $\check{\mathbf{f}}_s$:

$$\{\check{\mathbf{f}}_s\} \Big|_{\{\mathbf{f}_a\}=0} = \{\mathbf{f}_s\} \quad (8)$$

Given the pre-computation of the modified stiffness matrix in (6) and the speed of assigning new values in (8), the generic method offers a very large speed increase compared to the

patient-specific method and can potentially be performed at near real-time frame rates. Both correction methods were implemented in MATLAB on an Intel Core 2 Quad CPU at 2.4GHz with 4GB of RAM.

Experimental validation

Simulations

Several simulations were performed to examine the sensitivity of the generic correction method to various factors. The first simulation performed was to analyze the effect of the finite element mesh resolution on the model correction. This simulation consisted of constructing equivalent $10 \times 10 \times 5$ cm block meshes with a tetrahedral element edge length ranging from 2 to 10 mm. The number of nodes in the meshes ranged from 53,018 to 634 between the 2- and 10-mm resolution, respectively. Three simulated tumors were created with diameters of 10, 20, and 40 mm and placed separately in an instance of each mesh. A 10-mm surface compression was then simulated for each mesh, and the model solution was interpolated to the tumor nodes for comparison of the effects of the mesh resolution on the correction strategy. The comparison was performed by utilizing the most finely resolved mesh (2-mm edge length) as the ground-truth solution, with each subsequent model solution from the coarser meshes being compared to the ground-truth solution in terms of the difference in final tumor position.

The second simulation performed was aimed at determining the potential effects of the following variables upon the correction: (1) the block mesh size, (2) the tumor stiffness, and (3) the tumor size. Three block meshes were created with dimensions of $10 \times 10 \times 5$ cm, $10 \times 10 \times 10$ cm, and $10 \times 10 \times 15$ cm constructed with 5-mm edge length. Three simulated tumors were created with diameters of 10, 20, and 40 mm and each placed at a 3 cm depth (half the maximum depth of the linear array probe used for the phantom and clinical data in this work) in an instance of each block mesh described above. To illustrate, the 20-mm tumor is shown in the meshes of different sizes in Fig. 3. The tumors were assigned stiffness values of 1:1, 10:1, and 30:1 compared to the rest of the tissue block, resulting in 27 meshes (three mesh sizes, three tumor sizes, and three stiffness ratios). Each mesh was then subjected to surface compression ranging from 0 to 10 mm. For each state of compression, the model-deformed tumor surfaces were compared to the uncompressed tumor surfaces to illustrate the effect of tumor stiffness and size on the model solution, which would in turn affect the correction.

The last simulation performed was aimed at examining the sensitivity of the generic correction to varying tumor locations within the tissue. An organ-like finite element mesh was first constructed from a CT image volume of a liver phantom. Three 20-mm-diameter spheres were manually inserted into the mesh at different locations to act as simulated tumors. The three locations are shown in Fig. 4 and were chosen to demonstrate that the generic correction method improves subsurface target localization regardless of the local organ geometry involved. In the case of each tumor, a surface compression of 1 cm was simulated using the liver mesh to compress the tumor, and then the generic correction method was used to correct the deformed tumor using a $10 \times 10 \times 10$ cm block mesh. The

generic corrected tumors were then each compared to the original uncompressed spheres in terms of boundary node error and tumor centroid error.

Phantom experiments

A compliant gel phantom was constructed by mixing 7% by mass polyvinyl alcohol (PVA) in water with 10% by volume glycerol. A small amount was poured into a tumor mold and subjected to four freeze–thaw cycles, in which the gel was frozen at -40°C for 12h and then thawed at room temperature for 12 additional hours. The first freeze–thaw cycle produces a gel with a tissue-like consistency, and each additional cycle results in an increasingly stiffer material. The phantom tumor was then suspended by wire in a larger anthropomorphic liver mold with PVA mixture and subjected to one additional freeze–thaw cycle. This resulted in a soft tissue phantom containing a stiffer tumor. The completed phantom was fixed to a rigid base containing fiducials, which were used to initialize the image-to-physical registration.

Phantom tomograms were acquired in this study in order to compare the patient-specific and generic correction methods to a ground-truth CT image set. CT image volumes were acquired for the phantoms using a clinical CT machine at $512 \times 512 \times 422$ with 0.6-mm isotropic voxels. The bulk phantom and tumor were segmented using intensity thresholding in Analyze 9.0 (Mayo Clinic, Rochester, MN). Isosurfaces were generated from the segmentations using the marching cubes algorithm and smoothed with a Laplacian filter. A patient-specific finite element mesh with tetrahedral elements was created from the smoothed phantom isosurface using custom-built mesh generation software [11].

The phantom fiducial markers were localized in physical space with a tracked pointer, and an initial rigid point-based registration to the CT images was performed [12]. An LRS scan of the liver surface was acquired, and an iterative closest point (ICP) registration was performed of the tracked LRS point cloud to the CT surface in order to refine the registration [13]. This alignment was used to perform the patient-specific correction method and served as the gold standard validation for the proposed generic correction method.

An Acuson Antares ultrasound machine (Siemens Inc., Munich, Germany) with a VFX13-5 linear array probe at 10MHz was used to acquire all ultrasound images in this study. The machine was used to collect both B-mode and strain images with the eSie Touch elasticity software from the manufacturer in order to illustrate the general applicability of the correction method to all forms of ultrasound data. The ultrasound images were tracked in 3D space by synchronizing each image with the pose detected by a Polaris Spectra optical tracking system (Northern Digital, Waterloo, ON, Canada) for a passive rigid body attached to the ultrasound probe. The tracked ultrasound probe was calibrated using the method described by Muratore and Galloway [14] such that all pixels in each image were associated with a 3D pose.

In addition to the ultrasound data collected above, the other intraoperative tools used in this study were a tracked pointer and LRS [15]. The pointer was used to digitize point fiducials such as beads on the phantom base and craniofacial landmarks on the patient. These points were used to initialize a surface-based registration of the dense point clouds from the LRS to the preoperative patient-specific model in the case of the phantoms.

Tracked B-mode and strain images were acquired of the embedded tumor, and the transformation matrix from the ICP registration was used to automatically align all tracked ultrasound images with the CT data. A total of 178 B-mode and 83 strain images were collected of the tumor. The tracking and registration transformations were then applied to the digital probe surface in order to generate boundary conditions for the two correction methods as described previously.

After the generation of boundary conditions, the patient-specific and generic correction methods were applied to each ultrasound image. This resulted in a collection of uncorrected, patient-specific corrected, and generic corrected images. With respect to the generic correction, a $10 \times 10 \times 10$ cm block mesh was used. For both methods, the meshes were assigned a tumor Young's modulus ratio of 1:1 with Poisson's ratio at 0.49 because PVA is known to be nearly incompressible. Each population of images was compared to the baseline CT images in terms of tumor geometry in order to evaluate the corrections. The tumor borders in each B-mode and strain image were segmented semiautomatically using the Livewire technique, and for each ultrasound image, the CT volume was re-sliced to provide a coplanar CT slice and tumor contour corresponding to the 3D pose of the co-registered ultrasound slice. The tumor borders from ultrasound and CT were then compared in terms of modified Hausdorff distance (MHD) and centroid distance between the two contours [16]. The MHD and centroid error metrics were computed prior to and after each correction and were the primary metrics in determining the efficacy of the methods.

Clinical case

The patient-specific and the generic correction methods were deployed in a preliminary clinical case. The clinical dataset consisted of a meningioma patient at Vanderbilt Medical Center. Informed written consent was obtained from the patient prior to the study with the approval of our institutional review board. The preoperative MR volume was segmented to produce brain and tumor surfaces, which were used to create a patient-specific model in the same manner as the phantom data. With respect to the generic correction, a $10 \times 10 \times 10$ cm block of tissue was used. The tumor in this case was a meningioma located superficially on the left side. In this case, the tumor and brain were assigned a 1:1 stiffness ratio and Poisson's ratio of 0.45 [17]. Alignment of the intraoperative tracked ultrasound images to the MR was performed by scanning the face of the patient with the LRS and performing an ICP registration between the LRS face point cloud and the MR patient model. There was no LRS cloud of the brain surface available in the case of this specific patient, and so a random sampling of the MR model surface was used to simulate LRS data in that case. Tracked B-mode images were obtained immediately after the craniotomy. Both correction methods were then applied to the ultrasound data and compared to the co-aligned MR tumor borders in terms of the coplanar contour MHD and centroid error.

Computational efficiency

In order to provide an estimate of the computational speed offered by the generic correction framework, the time to compute the generic model solution for an image slice in the phantom B-mode dataset was recorded in the case of a $10 \times 10 \times 5$ cm block mesh with 5-mm edge length consisting of 4042 nodes and 19,672 tetrahedral elements. When taking into

account our condensation approach, the number of nodes and elements used in the generic correction was reduced to 697 and 2698, respectively. The analogous correction using the patient-specific correction was computed using a mesh with 5-mm edge length consisting of 10,989 nodes and 55,165 tetrahedral elements. The difference in the numbers of nodes and elements in the meshes having similar edge length corresponds to the difference between the volume of the full patient-specific organ versus the volume of the block of tissue in the generic method.

Results

Simulations

The results of the mesh resolution simulations are shown below in Fig. 5. This figure displays how the model solutions at varying mesh resolutions changes compared to the solution to the high-resolution mesh using a 2-mm element edge length. The general trend in each case was that as the mesh becomes coarser, the interpolated model solution deviates from the solution obtained from the more finely resolved mesh, especially above an edge length of 7mm.

The results of the second simulation testing the effects of mesh size, tumor stiffness, and tumor size are shown below in Fig. 6. Each graph shows that as the applied surface compression increases, the tumor boundary nodes displace correspondingly. We see that compression is communicated to the tumor boundary displacement more effectively as the tumors decrease in stiffness with larger distortions experienced by larger tumors (which is expected, i.e., softer materials would experience larger shape distortion). Looking across depth block sizes, we see that the tumor boundary experiences more displacement with increasing block depth. This is expected as the far-field fixed displacement will inhibit internal block motion less in larger depth blocks. However, looking at the top subfigures in Fig. 6, we do see more pronounced effects from stiffness at large compressions with larger tumors.

The results of the last simulation testing the effects of tumor location are shown in Fig. 7. This figure shows the tumor boundary node error (for 50 boundary nodes) prior to correction and after application of the generic method. Prior to correction, the centroid error for the compressed tumor in location 1, 2, and 3 designated in Fig. 4 was 6.7, 5.0, and 2.0mm, respectively. After correction, the centroid error in location 1, 2, and 3 was 1.6, 1.8, and 0.8mm, respectively.

Phantom experiments

The patient-specific and generic model corrections were deployed in the liver phantom, and an example of the correction process applied to a tracked ultrasound slice is shown in Fig. 8. Qualitatively, there was a clear improvement to the alignment between ultrasound and co-registered tomograms in the phantom experiments. In addition, the ultrasound contours corrected with the generic model method were very geometrically similar to the ultrasound contours corrected with the patient-specific method. We should note that difference in contour and green rendered object in Fig. 8b, e represents the error in object localization and

shape if no correction is performed in the case of patient-specific, and generic model respectively.

The quantitative results of the phantom experiments in Fig. 9 show the MHD and coplanar-centroid distances as error metrics in comparing the ultrasound tumor borders with the co-registered CT borders, for both B-mode and strain images. The B-mode MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.0 ± 1.6 , 1.9 ± 0.6 , and 2.1 ± 0.7 mm, respectively. A Wilcoxon signed rank test was computed for the null hypothesis that the median difference between the error metrics was zero. It was found that there was a statistically significant difference between each of the image populations using this metric ($p < 0.01$). The B-mode centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 7.6 ± 2.6 , 2.0 ± 0.9 , and 2.6 ± 1.1 mm, respectively. The Wilcoxon test again found the three image populations to be significantly different from one another based on this metric ($p < 0.01$).

With respect to the strain images, the strain MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.6 ± 1.1 mm, 2.0 ± 0.5 mm, and 2.2 ± 0.5 mm, respectively. The Wilcoxon test found all three image populations to be statistically different using the MHD metric ($p < 0.01$). The strain centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 8.0 ± 1.6 mm, 3.0 ± 0.9 mm, and 3.3 ± 1.1 mm, respectively. The Wilcoxon test also found all three image populations to be statistically different using the centroid error metric ($p < 0.01$).

Clinical case

The patient-specific and generic model corrections were deployed in the clinical case, and an example result is shown in Fig. 10. The quantitative results of the clinical case in Fig. 11 show the MHD and coplanar-centroid distances as error metrics in comparing the ultrasound tumor borders with the co-registered CT borders, for both B-mode and strain images. The B-mode MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.4 ± 0.1 , 2.6 ± 0.1 , and 2.9 ± 0.1 mm, respectively. A Wilcoxon signed rank test was computed for the null hypothesis that the median difference between the error metrics was zero. It was found that there was a statistically significant difference between each of the image populations using this metric ($p < 0.01$). The B-mode centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 7.2 ± 0.2 mm, 3.5 ± 0.4 mm, and 3.8 ± 0.4 mm, respectively. The Wilcoxon test again found the three image populations to be significantly different from one another based on this metric ($p < 0.01$).

Computational efficiency

A breakdown of the various computational costs in terms of execution time for the patient-specific correction and generic correction is given in Table 1. In the case of the patient-specific method, the mesh is created from preoperative imaging, which typically requires at least 30 min, assuming that some manual oversight of the image segmentation is required. Creation of the mesh from the segmentation mask takes at least 5 min. In terms of actual intraoperative expense, the construction of the $32,967 \times 32,967$ stiffness matrix K and the

solution of Eq. (2) were conducted together in approximately 50 s, and this represented the vast majority of the total intraoperative computation time of 52.5 s.

In the case of the generic method, the modified stiffness matrix given by (6) was pre-computed for the block mesh prior to collection of ultrasound data. There were only 697 nodes in the ultrasound plane region of the mesh, and so the condensed stiffness matrix in this case contained 2091×2091 entries. For each model correction, the vector given by (7) was modified with the detected compression vectors, and (5) was solved for the plane node displacements. The inverse of the modified stiffness matrix in (6) was stored, and the solution time for the model was approximately 10 ms. The overall intraoperative computation time was approximately 80 ms.

Discussion

Simulations

The mesh resolution sensitivity study in Fig. 5 shows that in all but the largest tumors, there is less than an approximate ~5% difference between mesh solutions when the element edge length is approximately 5mm. This would suggest that in domains with large heterogeneities more care and finer resolution may be needed but as a routine method, 5-mm edge lengths may be sufficient.

The second simulation study showed how the block mesh size, tumor size, and tumor stiffness affect the model solution. The first observation is that the size of the block mesh did not affect the solution at the tumor nodes until the size was much larger than the depth at which the tumor was located (recall the tumor was placed at 3cm to be in the center of a 6cm US image). The solutions at the tumor nodes were similar when utilizing block sizes of $10 \times 10 \times 5$ cm and $10 \times 10 \times 10$ cm, but at the $10 \times 10 \times 15$ cm size, the tumor solutions tended to become less variable across different tumor sizes and stiffness. One trend to note from Fig. 6 is that there is very little impact on the model solution from the size of a tumor or its stiffness ratio at low levels of surface compression. However, as the surface compression becomes quite large, there is a divergence in the solutions on the basis of both tumor size (illustrated by the displacement magnitudes in each graph) and the stiffness ratio (the vertical axis on each graph). Larger tumor size resulted in greater overall tumor boundary deformation, which was expected because a larger tumor diameter implies that a greater proportion of tumor nodes were closer to the surface deformation, since all three tumors were placed at the same tissue depth. It is especially worth observing that the importance of tumor stiffness increased with increasing tumor size. In the case of the 10-mm-diameter tumor at the maximum surface displacement of 10mm, the difference in mean tumor displacements when using the 1:1 and 30:1 stiffness ratios varied by approximately 1mm. However, in the case of the 40-mm-diameter tumor, the difference in mean tumor displacements when using the 1:1 and 30:1 stiffness ratios varied by approximately 3mm for the same surface displacement magnitude. These simulations indicate that although in many cases the tumor geometry and material properties do not greatly impact the model solution, these variables can become important when the target is a large tumor with a much different stiffness from the surrounding normal tissue. This becomes somewhat of a limitation for the generic method, as with the generic method the mesh is pre-computed based on a

homogeneous domain. There may be some superposition and weighted combinatorial possibilities that could account for stiffness variations that would entail more pre-computation but would likely reduce these errors. While this is outside the scope of this paper, it does provide some intriguing directions for the work.

The last simulation showed the effect of tumor location within the organ on the correction. The main observation from Fig. 7 is the trend in tumor boundary error reduction after application of the generic correction compared to the uncorrected error. These simulations predict very well the subsequent results observed in the phantom and clinical experiments. It is also interesting to note the consistent behavior of the correction regardless of the tumor location and the local organ geometry. For instance, the median boundary error for the tumor in location 1 was reduced from a median value of 6.3–1.7mm, over a 70% reduction. In the case of the tumor in location 3, which was located deep in a thick region of the organ and thus was not deformed as much, the median uncorrected error was only 1.8mm and was still reduced to 0.8mm, over a 50% reduction. This finding is significant because it demonstrates that even with varying deformation scenarios due to tumor location, the generic correction method can improve tumor localization accuracy compared to no correction.

Phantom experiments

The results of the compression correction methods shown in Fig. 9 clearly demonstrate the improvement offered by both the patient-specific and generic methods to the alignment between ultrasound and co-registered tomograms in the phantom experiments. The MHD error metric showed a significant decrease in misalignment after application of both methods, but with a greater decrease for the patient-specific method. The centroid distance error metric showed an even clearer improvement after applying the two model-based corrections compared to the uncorrected ultrasound images. Interestingly, a similar trend was noted with respect to the corrections, in that the patient-specific method resulted in a greater decrease in error overall compared to the generic model correction. However, although the correction for the patient-specific method was considered significantly better than the generic method according to the Wilcoxon test, the mean difference between the resultant errors for these correction frameworks was submillimetric for both the B-mode and strain images. This result is important because it indicates that the generic model correction performs nearly as well as the patient-specific method, making it a reasonable alternative for cases in which a patient-specific model may not be available or is too cumbersome computationally.

Clinical case

The results of the clinical case in Fig. 11 showed a clear improvement in alignment between the tracked B-mode and MR tumor borders after both of the correction methods. The same trend from the phantom dataset was noted in this case, which was that the patient-specific model correction yielded a slightly greater reduction in error than the generic model correction. However, as with the phantom data, the difference in the mean error for both metrics was submillimetric in comparing the two corrections. This reinforces the idea that the generic model correction could be used to perform a comparable compression correction

in the absence of a patient-specific model from preoperative imaging or when computational speed is paramount.

Computational efficiency

It was found that the patient-specific method on average needed approximately 50 s to provide a compression correction update to each individual ultrasound frame during freehand movement of the probe. This long computation time was primarily a consequence of the need to re-assign boundary condition types to surface nodes in the patient-specific mesh as the probe was moved around the tissue, thus necessitating a full reconstruction of the stiffness matrix K for each correction. The stiffness matrix for the patient-specific mesh was much larger than the mesh in the generic correction due to the greater number of nodes needed to represent the full patient organ, thus leading to a longer solution time as well. We should note that with some clever storage strategies, the potential to nearly eliminate the 40 s element assembly construction is likely, but one would still be faced with an approximate 12.5 s process. Regardless, these correction times still remove one of the primary advantages of ultrasound as an interventional imaging modality, which is its real-time data acquisition. There is clearly motivation to provide both a corrected image while maintaining a high frame rate.

By contrast, the generic correction method was shown to provide a model solution in approximately 10 ms using the condensation method to only solve for the mesh nodes in the immediate vicinity of the ultrasound plane. This essentially represents the removal of the primary computational bottleneck from the patient-specific correction, which was the 52.5/12.5 s required for construction and solution of the stiffness matrix for a large organ-shaped mesh. The rest of the intraoperative steps consisted of determining boundary conditions prior to the model solution and then interpolating the model solution to the ultrasound data. The other steps combine with the model solution to give a total intraoperative correction time of approximately 80 ms, which is nearly real time at 12.5 frames per second. In addition, this work was implemented on only a single CPU and could easily be further improved by the use of GPU programming. In closing, this work demonstrates that although there is a modest reduction in the accuracy of the solution provided by the generic correction versus the patient-specific method that is statistically significant, the dramatic computational benefit provided by the former at a cost of minor inaccuracy cannot be discounted.

Limitations

The proposed generic model correction shares many of the same limitations of the patient-specific correction enumerated in [5]. For example, the generic correction is still subject to several sources of propagating error in the image guidance workflow. It heavily relies upon the optical tracking system, which imparts an inherent error to each measurement made with the device, including the surface digitization using a tracked pointer or LRS, as well as the tracked ultrasound data. It also retains the assumption from the patient-specific method that the user applies the probe purely in the depth direction for each image acquisition. This again simplified the creation of boundary conditions for the model, which is a challenge shared by both methods. In addition to the accuracy of the boundary conditions, the

geometry of the mesh itself was likely the primary cause for the difference in error observed between the generic model correction and the patient-specific correction. A block of tissue is clearly a very simplistic representation of most anatomical structures on which this method would be used. The size of the block mesh also needs to be chosen before the correction can occur. Although a $10 \times 10 \times 10$ cm cube was used for each generic model correction for the phantoms and clinical case in this study, it would be possible to pre-construct block meshes of various sizes as shown in Fig. 3 based on observable knowledge of the anatomy of interest that could be selected during the procedure. It would be fairly trivial to exchange various pre-constructed block meshes of different depths intraoperatively. The most computationally expensive operation during the procedure would be computing the compression depth using the LRS or pointer point cloud and then performing the interpolation of the solution to the ultrasound data. It would conceivably be possible to perform all of the operations in Table 1 at a real-time frame rate if efficiently implemented and with the support of GPU programming. Going a bit further, it might be possible to pre-compute the full displacement field on the slice itself, although this raises questions as to how to handle applications of the probe in compression that are not uniform, i.e., cases where the probe face is not approximately parallel with the intraoperative organ surface. More than likely, a lookup table with an interpolation and superposition scheme could be used intraoperatively to select the solution combination corresponding to the correct compression magnitude. This would likely result in an additional speedup compared to our current results although validation of this approach would be needed.

Another assumption retained in the generic correction method was the assignment of material properties to the finite element mesh. Accurate intraoperative measurement of tissue mechanical properties is very challenging in practice. The approach taken in this work was to assume the mesh was composed of a single homogenous tissue type. Under this assumption, the biomechanical model solution would become less accurate with increasing contrast in tissue stiffness. It should be noted that only the relative stiffness values would affect the solution of the model in either the patient-specific or the generic model correction because only Dirichlet boundary conditions drive the solution. Absolute values for Young's modulus would only affect the solution if force conditions were integrated into the approach. Similar to the previous comments regarding pre-computing the displacement field itself, it is likely that a pre-computational strategy that uses superposition and combinatorial approaches intraoperatively could be used to account for stiffness differences driven by perhaps strain-imaging data which is an intriguing direction for future work but outside of the scope of this paper.

Although it is true that the assumptions and simplifications made in this work could be seen as defeating the purpose of a patient-specific biomechanical model, these simplifications can be interpreted as methodologies that allow for improved fidelity measurements within the context of soft tissue image-guided environments. More specifically, when one considers the amount of tissue deformation that is experienced during the presentation of open or laparoscopic liver resection as an example, the generic framework allows for some re-establishment of measurement fidelity of subsurface structures acquired by ultrasound imaging. Ultimately, as guidance systems continue toward the full employment of nonrigid approaches for registration, the need for accurate localization of tissue structure within a

consistent patient space that is workflow sensitive will be of high importance. This balance of accuracy and workflow within the context of sparse data will remain of paramount concern in these new paradigms.

A final limitation of both the patient-specific and generic methods is that each requires a surface acquisition during the procedure, which in our setup was provided by LRS. It is generally understood that adding hardware requirements to a procedure would limit its adoption. However, there are four aspects to this to consider. The first is that in certain procedures with very high accuracy requirements, such as neurosurgery, the added benefits of an enhanced spatial understanding of the interventional target may justify the hardware burden. The second aspect is that there are other methods besides an LRS device, which can be used to provide the necessary surface measurement. For example, some have proposed to use a tracked ultrasound probe itself as a surface digitizer by swabbing the surface [18], and others have proposed to use stereoscopic surgical cameras to construct the organ surface [19]. A third consideration is that standardized surface acquisition is likely on the horizon for image-guided surgery systems. For example, surface registration via swabbing is already a standard feature on many commercial image-guided platforms. The last point to consider is that it is possible to trigger this correction using contact conditions by monitoring the ultrasound imaging data itself as it comes into contact with the tissue. This has its own challenges but is certainly worth considering. Nevertheless, much of the equipment utilized in this work is already routinely available in many operating rooms or can be modified to acquire appropriate data (e.g., stereo-pair is a standard feature in surgical microscopes). The method we used for digitizing the surface is just one embodiment of a more general concept.

The overall result of this work is that information in tracked ultrasound data can be corrected in near real time, provided that a measure of tissue compression is available intraoperatively. The immediate benefits are obvious in providing the clinician with more accurate size and position measurements of subsurface targets. This is important in a wide variety of procedures and anatomy, such as determining resection or ablation margins. Additionally, there are implications for more speculative work using subsurface information for enhanced registration [20]. An analysis of the effects of integrating corrected and uncorrected ultrasound data in a registration framework would be very interesting and also awaits further study.

Conclusions

In this work, we proposed and validated a novel method for correcting tissue compression error exerted by an ultrasound probe. Our novel generic tissue model was used to estimate physical tissue deformation as a result of pressing the tracked probe into the tissue surface and also compared to a previous compensation framework as well as to gold standard intraoperative CT imaging measurements. The experimental results indicate that the generic model correction method provides significantly improved intraoperative localization data. This is particularly important when patient-specific models may not be available from preoperative imaging or when minimizing computational encumbrance is important.

Acknowledgments

This work was supported in part by the National Institutes of Health award R01 NS049251 of the National Institute for Neurological Disorders and Stroke, and by the National Institutes of Health award R01 CA162477 from the National Cancer Institute.

References

1. Treece GM, Gee AH, Prager RW. RF and amplitude-based probe pressure correction for 3D ultrasound. *Ultrasound Med Biol*. 2005; 31(4):493–503. doi:10.1016/j.ultrasmedbio.2004.12.018. [PubMed: 15831328]
2. Xiao G, Brady JM, Noble JA, Burcher M, English R. Non-rigid registration of 3-D free-hand ultrasound images of the breast. *IEEE Trans Med Imaging*. 2002; 21(4):405–412. doi:10.1109/TMI.2002.1000264. [PubMed: 12022628]
3. Burcher MR, Lianghao H, Noble JA. Deformation correction in ultrasound images using contact force measurements. *IEEE workshop on mathematical methods in biomedical image analysis (MMBIA)*. 2001:63–70. 2001 doi:10.1109/mmbia.2001.991700.
4. Sun, S-Y. Thesis S.M.—Massachusetts Institute of Technology Dept. of Electrical Engineering and Computer Science 2010. Massachusetts Institute of Technology; Cambridge: 2010. Deformation correction in ultrasound imaging in an elastography framework..
5. Pheiffer TS, Thompson RC, Rucker DC, Simpson AL, Miga MI. Model-based correction of tissue compression for tracked ultrasound in soft tissue image-guided surgery. *Ultrasound Med Biol*. 2014 doi:10.1016/j.ultrasmedbio.2013.11.003.
6. Rucker DC, Wu Y, Clements LW, Ondrake JE, Pheiffer TS, Simpson AL, Jarnagin WR, Miga MI. A mechanics-based nonrigid registration method for liver surgery using sparse intra-operative data. *IEEE Trans Med Imaging*. 2013; 33(1):147–158. doi:10.1109/tmi.2013.2283016. [PubMed: 24107926]
7. Lange T, Papenberg N, Heldmann S, Modersitzki J, Fischer B, Lamecker H, Schlag PM. 3D ultrasound-CT registration of the liver using combined landmark-intensity information. *Int J Comput Assist Radiol Surg*. 2009; 4(1):79–88. doi:10.1007/s11548-008-0270-1. [PubMed: 20033605]
8. Chen I, Ong RE, Simpson AL, Sun K, Thompson RC, Miga MI. Integrating retraction modeling into an atlas-based framework for brain shift prediction. *IEEE Trans Biomed Eng*. 2013; 60(12):3494–3504. doi:10.1109/Tbme.2013.2272658. [PubMed: 23864146]
9. Coffey AM, Miga MI, Chen I, Thompson RC. Toward a preoperative planning tool for brain tumor resection therapies. *Int J Comput Assist Radiol Surg*. 2013; 8(1):87–97. doi:10.1007/s11548-012-0693-6. [PubMed: 22622877]
10. Bro-Nielsen M. Finite element modeling in surgery simulation. *Proc IEEE*. 1998; 86(3):490–503. doi:10.1109/5.662874.
11. Sullivan JM, Charron G, Paulsen KD. A three-dimensional mesh generator for arbitrary multiple material domains. *Finite Elem Anal Des*. 1997; 25(3–4):219–241. doi:10.1016/S0168-874x(96)00027-3.
12. Arun KS, Huang TS, Blostein SD. Least-squares fitting of 2 3-D point sets. *IEEE Trans Pattern Anal Mach Intell*. 1987; 9(5):699–700.
13. Besl PJ, McKay ND. A method for registration of 3-D shapes. *IEEE Trans Pattern Anal Mach Intell*. 1992; 14(2):239–256.
14. Muratore DM, Galloway RL. Beam calibration without a phantom for creating a 3-D freehand ultrasound system. *Ultrasound Med Biol*. 2001; 27(11):1557–1566. [PubMed: 11750755]
15. Pheiffer TS, Simpson AL, Lennon B, Thompson RC, Miga MI. Design and evaluation of an optically-tracked single-CCD laser range scanner. *Med Phys*. 2012; 39(2):636–642. doi:10.1118/1.3675397. [PubMed: 22320772]
16. Dubuisson MP, Jain AK. A modified Hausdorff distance for object matching. *Proceedings of the 12th IAPR international conference on pattern recognition, vol 1—conference A: computer vision and image processing*. Oct 9–13.1994 561:566–568. 1994. doi:10.1109/icpr.1994.576361.

17. Dumpuri P, Thompson RC, Dawant BM, Cao A, Miga MI. An atlas-based method to compensate for brain shift: preliminary results. *Med Image Anal.* 2007; 11(2):128–145. doi:10.1016/j.media.2006.11.002. [PubMed: 17336133]
18. Miller, KE.; Ondrake, JE.; Pheiffer, TS.; Simpson, AL.; Miga, MI. Utilizing ultrasound as a surface digitization tool in image guided liver surgery.. In: Kenneth, HW., editor. David RH III. SPIE; 2012. p. 83163D
19. Kumar AN, Miga MI, Pheiffer TS, Chambless LB, Thompson RC, Dawant BM. Persistent and automatic intraoperative 3D digitization of surfaces under dynamic magnifications of an operating microscope. *Med Image Anal.* 2015; 19(1):30–45. doi:10.1016/j.media.2014.07.004. [PubMed: 25189364]
20. Lange T, Eulenstein S, Hunerbein M, Lamecker H, Schlag PM. Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery. *Medical image computing and computer-assisted intervention-Miccai 2004, Pt 2, proceedings.* 2004; 3217:534–541.

Fig. 1.

Overview of the clinical workflow for the proposed method. The ultrasound data are calibrated to a tracked target attached to the probe (a), which is tracked intraoperatively (b) with a tracking system. A surface digitization tool such as a laser range scanner (LRS) is also tracked in a consistent coordinate system, and in the patient-specific method, this surface information is registered to the preoperative data (c), which is an optional step in the proposed generic correction method (denoted by the *dashed box*). The surface is also used to calculate the compressive depth for the model to correct the deformed ultrasound data (d), and it is this step that is the key difference between the two compared methods (denoted by the *solid box*)

Fig. 2.

Example of a mesh used for the patient-specific model correction method **(a)**, constructed from preoperative imaging and aligned to the ultrasound data using intraoperative registration methods, and an example of a block mesh for the generic model correction method **(b)**, which is pre-aligned to the ultrasound data by performing a calibration to the ultrasound probe

Fig. 3.

Example of the meshes used in the simulation studies with sizes of $10 \times 10 \times 5\text{cm}$ (a), $10 \times 10 \times 10\text{cm}$ (b) and $10 \times 10 \times 15\text{cm}$ (c). In this case, a simulated 20-mm-diameter tumor is shown at the equivalent 3cm depth in each mesh

Fig. 4. Simulated tumor locations, showing a tumor in the middle of the liver (1), in the small left lobe (2), and in the larger right lobe (3)

Fig. 5.

Effects of mesh resolution on three sizes of a tumor after a simulated 10-mm surface compression. The tumor node error is defined relative to the result of the solution of a mesh with 2-mm edge length resolution

Fig. 6. Effects of block mesh depth, tumor stiffness, and tumor size upon the model-predicted tumor border deformation under varying amounts of surface compression (*colorbar* is in millimeters)

Fig. 7. Effects of tumor location on generic block mesh correction improvement in three simulated tumors (number of tumor boundary nodes $N = 50$)

Fig. 8. Example of B-mode image slice correction with the patient-specific model correction (**a–c**) and generic model correction (**d–f**). **a, d** The co-registered LRS point cloud on the respective mesh, **b, e** the tracked probe surface and the misalignment between the ultrasound tumor border with the CT tumor, and **c, f** the corrected ultrasound image

Fig. 9.

Alignment error results for the B-mode (**a**, **b**) and strain imaging (**c**, **d**) modalities for the organ-like phantom ($n = 178$ for B-mode, and $n = 83$ for strain). The position of tumor borders in each modality was evaluated in terms of MHD to the co-aligned CT borders (**a**, **c**), as well as the distance between the centroid of the ultrasound tumor with the coplanar CT tumor border (**b**, **d**). The edges of the *boxes* are the 25th and 75th percentiles, and the *whiskers* extend to the most extreme data points not considered as outliers

Fig. 10.

Example of the patient images of a superficial tumor visible in MR and B-mode (a) and the patient-specific correction (b, c) and generic model correction (d, e) applied to a B-mode image slice for the clinical case. b, d The tracked probe surface and the misalignment between the ultrasound tumor border with the CT tumor, and c, e the corrected ultrasound image

Fig. 11.

Alignment error results for the clinical case ($n = 118$ B-mode images). The position of tumor borders was evaluated in terms of MHD to the co-aligned MR borders (a), as well as the distance between the centroid of the ultrasound tumor with the coplanar MR tumor border (b). The edges of the *boxes* are the 25th and 75th percentiles, and the *whiskers* extend to the most extreme data points not considered as outliers

Table 1

Approximate execution time for each step in the patient-specific and the generic correction pipelines

	Patient-specific correction	Generic correction
<i>Preoperative phase</i>		
Image segmentation	30 min	–
Mesh construction	5 min	2 min
Mesh calibration	N/A	2 min
Stiffness matrix pre-construction	N/A	1 min
Total	35 min	5 min
<i>Intraoperative phase</i>		
Boundary condition determination	50 ms	50 ms
Stiffness matrix construction	40 s	N/A
Model solve	12 s	10 ms
Ultrasound tumor interpolation	0.4 s	20 ms
Total	52.5 s	80 ms

Times were determined using a single thread of an Intel Core2 Quad CPU at 2.4 GHz