Group-wise registration of ultrasound to CT images of human vertebrae

Sean Gill[a], Parvin Mousavi[a], Gabor Fichtinger[a,b,c], David Pichora[c] and Purang Abolmaesumi[a,c,d]

[a] School of Computing, Queen's University, Kingston, ON, CANADA;
[b] Department of Mechanical Engineering, Queen's University, Kingston, ON, CANADA;
[c] Department of Surgery, Queen's University, Kingston, ON, CANADA
[d] Department of Electrical and Computer Engineering, Queen's University, Kingston, ON, CANADA;

ABSTRACT

Automatic registration of ultrasound (US) to computed tomography (CT) datasets is a challenge of considerable interest, particularly in orthopaedic and percutaneous interventions. We propose an algorithm for group-wise volume-to-volume registration of US to CT images of the lumbar spine. Each vertebra in CT is treated as a sub-volume and transformed individually. The sub-volumes are then reconstructed into a single volume. The algorithm dynamically combines simulated US reflections from the vertebrae surfaces and surrounding soft tissue in the reconstructed CT, with scaled CT data to simulate US images of the spine anatomy. The simulated US data is used to register preoperative CT data to intra-operative US images. Covariance Matrix Adaption – Evolution Strategy (CMA-ES) is utilized as the optimization strategy. The registration is tested using a phantom of the lumbar spine (L3-L5). Initial misalignments of up to 8 mm were registered with a mean target registration error of 1.87±0.73 mm for L3, 2.79±0.93 mm for L4, 1.72±0.70 mm for L5, and 2.08±0.55 mm across the entire volume. To select an appropriate optimization strategy, we performed a volume-to-volume registration of US to CT of the lumbar spine, allowing no relative motion between vertebrae. We compare the results of this registration using three optimization strategies: simplex, gradient descent and CMA-ES. CMA-ES was found to converge slower than gradient descent and simplex, but was more robust for rigid volume-to-volume registration for initial misalignments up to 20 mm.

Keywords: Intensity-based registration, image-guided therapy, multimodality registration, ultrasound.

1. INTRODUCTION

Spinal injections for back-pain management are carried out on a frequent basis. Currently, these procedures are performed under fluoroscopy or CT guidance in specialized interventional radiology facilities, and thus incur a major burden on the healthcare system. Another drawback with the current practice is patient and surgeon exposure to X-ray radiation. Furthermore, in cases where fluoroscopy is used, the surgeon accumulates significant exposure to radiation over time. The goal of this research is to design a spine intervention system with ultrasound image guidance, requiring only a single preoperative CT or MRI scan. This would greatly reduce the exposure of both the patient and the physician to ionizing radiation and allow the procedure to be performed outside of a specialized facility.

An essential component of such a system is the ability to reliably register an US scan of the spine with a preoperative CT volume. In this regard, many point- and surface-based registration methods have been proposed in the literature[1,2]. However, these techniques normally need manual intervention and require segmentation of US data which is time consuming and susceptible to errors. To avoid these problems, we chose to focus on intensity-based registration methods. In orthopaedic applications of US, the tissue occlusion caused by bone-soft tissue interfaces means there cannot be a direct correspondence between US intensities and CT intensities unless the CT data is modified to more accurately represent what is visible in US. Brendel et al. (2002) propose to define the bone surface in CT that is visible in an US image, with the similarity calculated as the US pixel intensities overlapping the surface[4,5]. The limitation of this method is that it requires *a prior* knowledge of the direction and orientation of the US probe. Penney et al. (2006) propose defining, for each pixel in CT and US, a probability that it is a bone edge[3]. These probability images are then

Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling,
edited by Michael I. Miga, Kenneth H. Wong, Proc. of SPIE Vol. 7261, 72611O

registered using normalized cross correlation. In order to have clinically relevant probabilities, a large dataset of CT and US images would have to be segmented. Wein et al. (2007, 2008) propose a registration method in which density information from CT data is used to iteratively simulate an US image throughout the registration process, thereby optimizing the simulation as the registration proceeds[6,7]. This has the benefit of not requiring any previous knowledge of the orientation of the US probe. This algorithm was designed for registration of soft tissue and does not account for the shadowing caused by an US beam reflecting at a soft tissue-to-bone interface. While the simulation simplifies the physics behind ultrasound beam propagation, it is accurate enough to bring the two modalities of CT and US closer together for registration purposes, and is also computationally efficient[4]. Shams et al. (2008) have built upon this algorithm to create a more realistic simulation for training physicians and technicians in the use of US imaging[8]. This technique requires preprocessing to create a scatter volume of the CT data using the Field II simulator[9,10], which remains computationally expensive.

In this paper, we propose an algorithm that extends the registration algorithm presented by Wein et al. (2007, 2008) to account for group-wise registration of vertebrae. The algorithm accounts for bone shadowing and the occlusion of tissue in the US image. We utilize the algorithm for rigid volume-to-volume registration and compare the results of three optimization strategies: simplex, gradient descent and Covariance Matrix Adaption – Evolution Strategy (CMA-ES). The registration algorithm applies a transform to the CT volume, simulates an US image from the transformed CT and then calculates an image similarity metric. This process iterates until convergence. We then present the algorithm for group-wise volume-to-volume registration. This approach allows free motion between the vertebrae in the CT data to account for the changes in patient posture between collection of the preoperative and intraoperative data. This algorithm is similar to the previous approach, but transforms each vertebra individually and reconstructs the sub-volumes into a single volume before calculating the simulated ultrasound. Initial and final target registration errors (TRE) are compared individually for each vertebra and across the entire volume. Finally, we discuss the strengths of the registration technique, sources of error in the computation and the direction of possible future work.

2. METHODS

A high-resolution CT volume (0.46 mm x 0.46 mm x 0.625 mm) and an US volume were acquired from a lumber spine (L3 to L5) phantom (Figure 1). The US volume was reconstructed from a freehand sweep using an SP10-60 wide-band linear-array transducer (General Electric Canada, Mississauga, ON, Canada) operating at 6.6 MHz with a depth of 5.5 cm. The probe was tracked using a Optotrack Certus Tracking System (Northern Digital Inc., Waterloo, ON, Canada) and calibrated using an N-wire US phantom[11]. All registrations were performed on a Dell Precision 690, with 2×2.33 GHz Intel Xeon Quad-core CPU and 16GB of RAM.

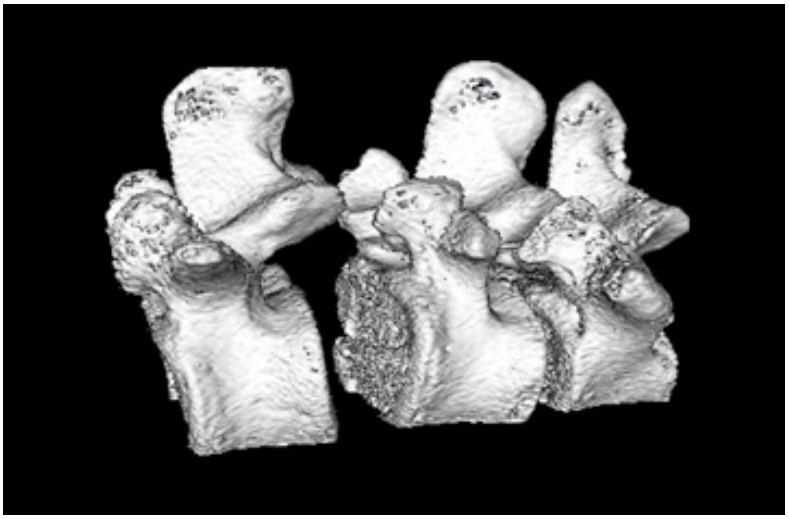

Figure 1. Surface model of the lumbar spine phantom. From left to right are the L3, L4 and L5 vertebrae.

2.1 Rigid Volume-to-Volume Registration

The US and CT volumes were initially aligned using the gold standard for registration calculated from fiducial markers implanted on the phantom. The CT volume was then perturbed by a random transform chosen from a uniform distribution of ±10 mm translation along each axis and ±10° rotation about each axis. The misaligned CT volumes were then registered back to the US volume using our US simulation-based registration algorithm. Accuracy was determined by the ability of the registration to recover to the fiducial-based gold standard and is reported as the mean target registration error (TRE) calculated as the misalignment of the four corners of the volume. One hundred such registrations were performed using our registration technique and a simplex optimization strategy. The registrations were repeated using gradient descent and Covariance Matrix Adaption – Evolution Strategy (CMA-ES)[12] as the optimizer.

2.2 Group-wise Registration

Next, the registration process was extended to simultaneously register three independently transformed lumbar vertebrae to a single US volume. CMA-ES was used as the optimization strategy to calculate the 18 rigid-body registration parameters: individual rotations and translations for each vertebra piece. The volume was cut into three sub-volumes, each containing a single vertebra. There was some overlap between sub-volumes, as the vertebrae overlap. Any pixels in the sub-volumes corresponding to bone from an external vertebra were masked. First, the entire volume was transformed, similar to Section 2.1, using the same set of random transformations; next, the individual vertebra sub-volumes were randomly transformed using a uniform distribution of ±5 mm translation along each axis and ±5° rotation about each axis. The former random transformation was applied to initially misalign the entire CT volume, while the latter was applied to simulate the deformation in the spine anatomy which could happen between pre-operative imaging and intra-operative intervention phases. After the individual transforms were applied, the pieces were stitched into a single volume. The maximum intensity was taken of any overlapping voxels across the sub-volumes, thus preserving bone structure in the final volume. Gaps were filled with a default value that approximates the intensity of soft tissue in CT. The US simulation was then applied to the reconstructed volume and the similarity with the US volume was calculated. It is important that the pieces be stitched together before the simulation is applied, as shadowing from the bone may occlude tissue from another piece. The TRE was calculated and reported individually for each piece along with the TRE for the entire volume.

2.3 US Simulation

Our simulation builds upon that presented for a linear probe in Wein et al. (2007, 2008)[6,7]. Their simulation was designed for registration of soft tissue and does not directly translate to simulation of US images containing bone. This is due to the full reflection of the US beam at a soft-tissue/bone interface and occlusion of all tissue in the US image past the interface. We have made several modifications to the simulation algorithm to account for this. A brief overview of the Wein et al. (2007, 2008)[6] algorithm, for the case of a linear probe, is presented, along with an explanation of our additions. The simulation can be divided into three calculations: Simulation of the US reflection from CT, mapping of the CT values to those found in US, and weighting of the simulated reflection, mapped CT and a bias to best represent the US image. A diagram of the simulation procedure with example images is shown in Figure 2.

In calculating the simulated US reflections, we model an ultrasound beam as a ray passing through each column in the CT image. Each ray begins with an intensity, I, of 1 and the reflection and transmission are calculated at each pixel using the following equations:

$$\Delta r(x,y,d) = \left(d^T \nabla\mu(x,y)\right) \frac{|\nabla\mu(x,y)|}{\left(2\mu(x,y)\right)^2}\,, \tag{1}$$

$$\Delta t(x,y) = 1 - \left(\frac{|\nabla\mu(x,y)|}{\left(2\mu(x,y)\right)}\right)^2, \tag{2}$$

$$r(x,y) = I(x,y-1)\Delta r(x,y,d)\,, \tag{3}$$

$$I(x,y) = \begin{cases} I(x,y-1)\Delta t(x,y)\,, & |\nabla\mu(x,y)| < \tau \\ 0\,, & |\nabla\mu(x,y)| \geq \tau \end{cases}, \tag{4}$$

where d is the direction of the US beam, μ is the intensity of the CT image, Δr is reflection coefficient, r is the simulated reflection intensity, Δt is the transmission coefficient, τ is the threshold for full reflection and I is the intensity of our simulated US beam. The equations are presented as two-dimensional for the sake of simplicity, but can be easily extended to three dimensions. Additionally, we apply the condition where any gradient value greater than a set threshold (450 h.u. in our simulations) causes full reflection of the US beam intensity at that point, setting the incoming US beam intensity for all subsequent points on the scan line to zero. The threshold value is the approximate gradient at soft tissue to bone interface in the CT images of our phantom and thus, simulates occlusion below this interface. To amplify the impact of the intensities of small reflections, a log-compression is applied to the simulated reflection image,

$$r(x,y) = \frac{\log(1+ar(x,y))}{\log(1+a)}, \tag{5}$$

To map the CT intensities to values closer to those corresponding to the tissues in the US data, we approximate the curve presented in Wein et al. (2007) [5] and use the following equation,

$$p(x,y) = 1.36\mu(x,y) - 1429. \tag{6}$$

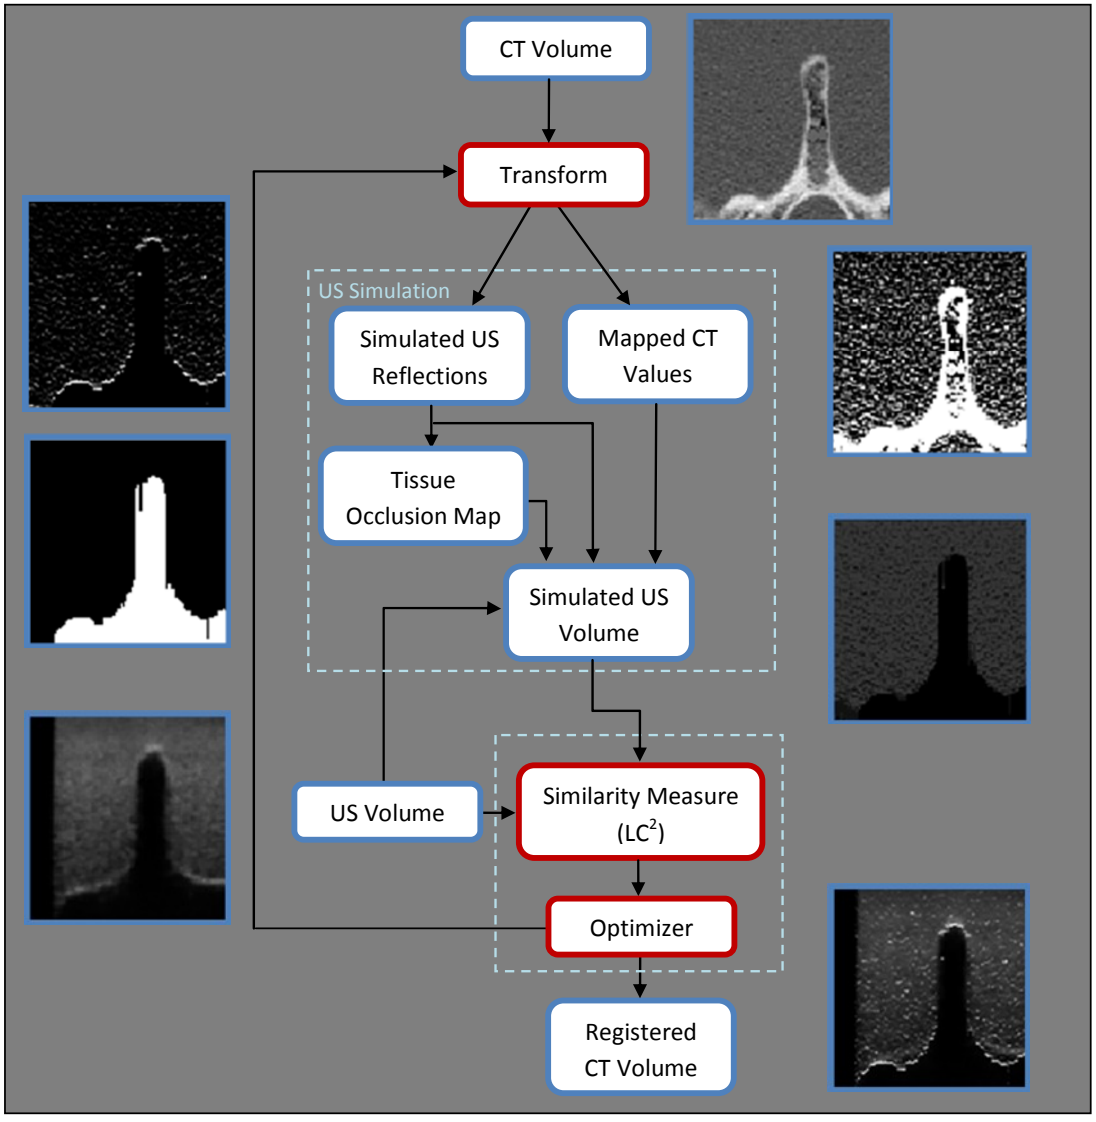

Figure 2. Workflow for US to CT Registration.

The final step of the US simulation is the weighting of the simulated US reflection, the mapped CT and a bias term. Assuming overlapping pixels in the US and CT datasets correspond, we use a least-squares optimization method to calculate the weights, such that the simulation values best correspond to the real US image. The final simulation is calculated as,

$$f(x,y) = \begin{cases} \alpha p(x,y) + \beta r(x,y) + \gamma\,, & I(x,y) > 0 \\ 0\,, & I(x,y) = 0 \end{cases}\,, \tag{7}$$

where f is the simulated US image and α, β, γ are the weights for their respective images. In our algorithm, the weights are calculated across the entire volume. We do not include any pixels that are occluded in the simulation as part of the weight calculation. This emphasizes the alignment of bone between the two modalities in the registration process. When the bone is misaligned, pixels in the US image that are occluded are included in the weight calculation, reducing the accuracy of the simulation. When the bone is properly aligned, only US pixels corresponding to soft tissue are included in the weight optimization (Figure 3). Occluded pixels are identified as any pixel where the intensity of the incoming simulated US beam is zero. In the final US simulation, all occluded pixels are set to zero.

2.2 Similarity Metric

The US simulation is updated at each step of the registration. Similarity between the actual US image and simulated US image is calculated using the Linear Correlation of Linear Combination (LC2) metric suggested by Wein et al. (2007)[6] and derived by incorporating the US and simulation values into a Correlation Ratio framework[13],

$$LC^2 = \frac{\sum(U(x,y) - f(x,y))^2}{N \times Var(U)}\,, \tag{8}$$

where N is the number of overlapping pixels between the US and CT images, and U is the actual ultrasound image intensity.

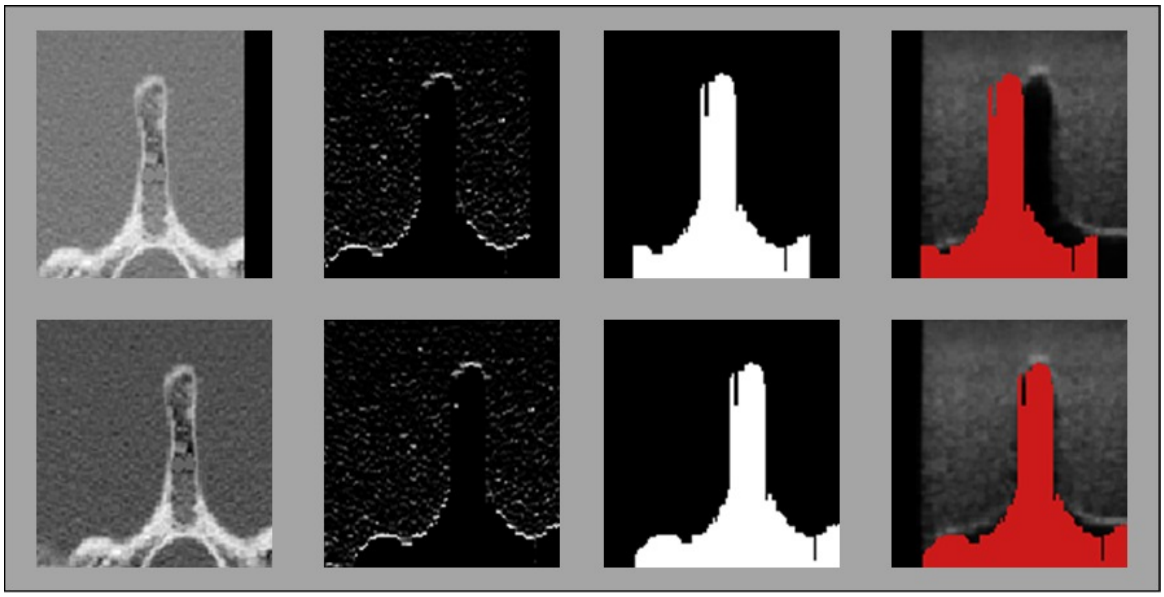

Figure 3. Example of pixels excluded in weight calculation for misaligned (Top) and aligned (Bottom) slices. Displayed are the CT slice (Left), simulated US reflections (Center Left), occluded pixel mask (Center Right) and the US slice overlaid with the occluded pixel mask (Right). Red pixels are not included in the weight calculation.

3. RESULTS AND DISCUSSION

Examples of misaligned and registered volumes are shown in Figures 4 and 5. Results for rigid volume-to-volume registration using simplex, gradient decent and CMA-ES optimization are displayed in Figure 6. Simplex registration resulted in a mean TRE of 3.39 mm and a standard deviation of 3.80 mm. For gradient decent optimization, the mean TRE was 3.27 mm with standard deviation of 3.48 mm. Using CMA-ES optimization resulted in a mean TRE of 2.26 mm and standard deviation of 0.50 mm. CMA-ES proved to be a more robust optimization strategy and was able to reliably register misalignments up to an initial TRE of 20 mm. Simplex and gradient descent optimization began to fail when initial misalignment approached 14 mm and 8 mm, respectively. CMA-ES was the most computationally expensive with a mean runtime of 258 seconds, while gradient descent was slightly faster with a mean runtime of 246 seconds. Simplex was the fastest, with a mean of runtime of 66 seconds.

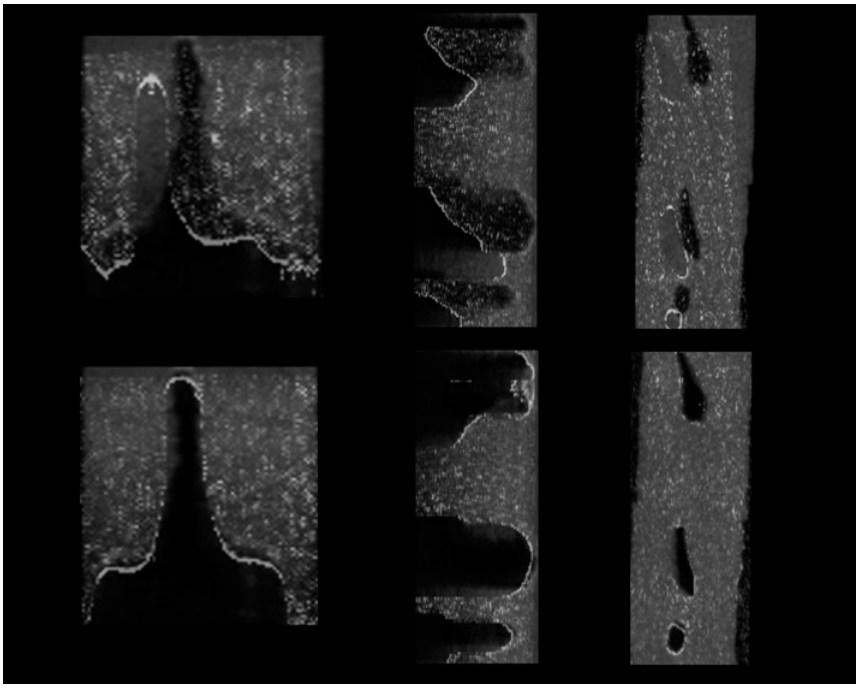

Figure 4. Transverse (Left), Sagittal (Center) and Coronal (Right) slices of the original US volume overlaid with the simulated US reflections of a misaligned CT volume (Top) and the volume after rigid volume to volume registration (Bottom).

As can be seen in Figure 6, no registration was able to achieve a final misalignment of less than 1 mm and all the misaligned registration converged to roughly the same position. This is most likely due to the US probe calibration error (~ 0.6 mm) and the gold standard registration error (0.62 mm). Another possible source of error is introduced when we reconstruct the US slices into a single volume. We must assume that the angle of the US probe is directly down along the y-axis of the volume and not along the actual plane of the US image that was acquired. However, we believe that the error introduced by this assumption is minimal. The data was collected with the probe in a vertical position and the collection was dense enough, ~ 2 frames/mm, that the direction of the US signal across the entire volume should be very close to vertical.

The results for the group-wise registration can be found in Figure 7, with individual results for each vertebra (Top Left, Top Right and Bottom Left) and for the overall volume (Bottom Right). Registration resulted in a mean TRE of 2.21 ±1.20 mm for L3, 3.28 ±1.10 mm for L4, 2.04 ±1.27 mm for L5, and 2.58 ±0.78 mm across the entire volume. The capture range of the registration was much lower than rigid volume-to-volume registration and the algorithm was able to reliably register vertebrae with an initial misalignment of 8 mm. Registrations with initial TRE of less than 8 mm resulted in a mean TRE of 1.87 ±0.73 mm for L3, 2.79 ±0.93 mm for L4, 1.72 ±0.70 mm for L5, and 2.08 ±0.55 mm across the entire volume.

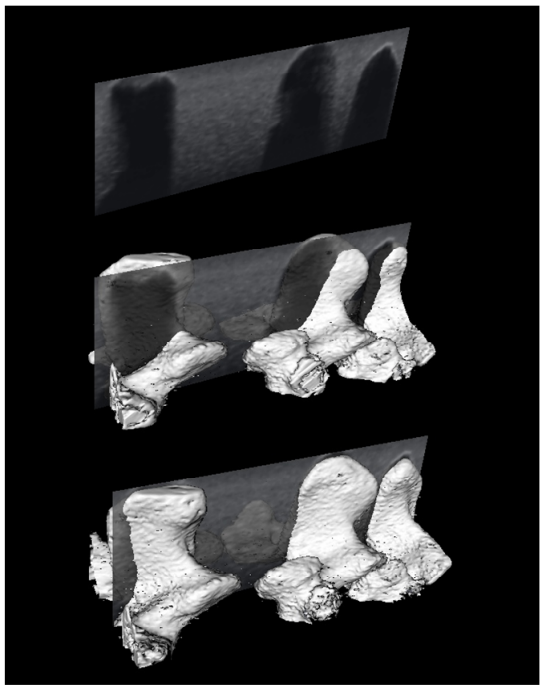

Figure 5. A sagittal slice from the US volume (Top) overlaid on the extracted surface of an example misaligned (Middle) and registered (Bottom) CT volume, using rigid volume-to-volume registration.

As the same dataset was used in the rigid volume-to-volume registration, the same sources of error apply. In addition, the largest difficulties presented by the algorithm are found in the overlapping of bone, causing occlusion and the fact we do not penalize collision of bone segments. This seems to be a problem in the registration of L4 and L5, where one vertebra would align with the US and the other would collide with the first. The type of alignment does not cause a significant decrease in the similarity metric, making it difficult for the algorithm to identify this as a poor registration.

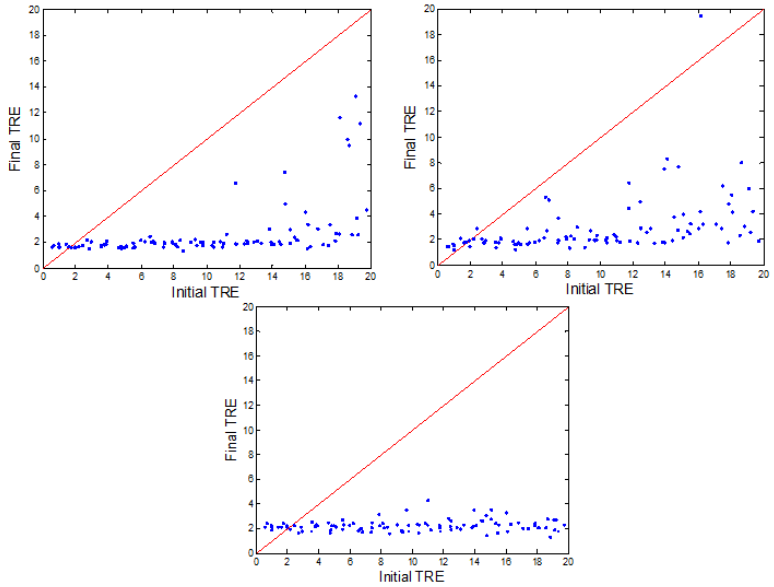

Figure 6. Results for volume-to-volume registration using simplex (Top Left), gradient descent (Top Right), and CMA-ES (Bottom Center) as the optimization strategy. Final target registration error (TRE) is plotted against the initial TRE.

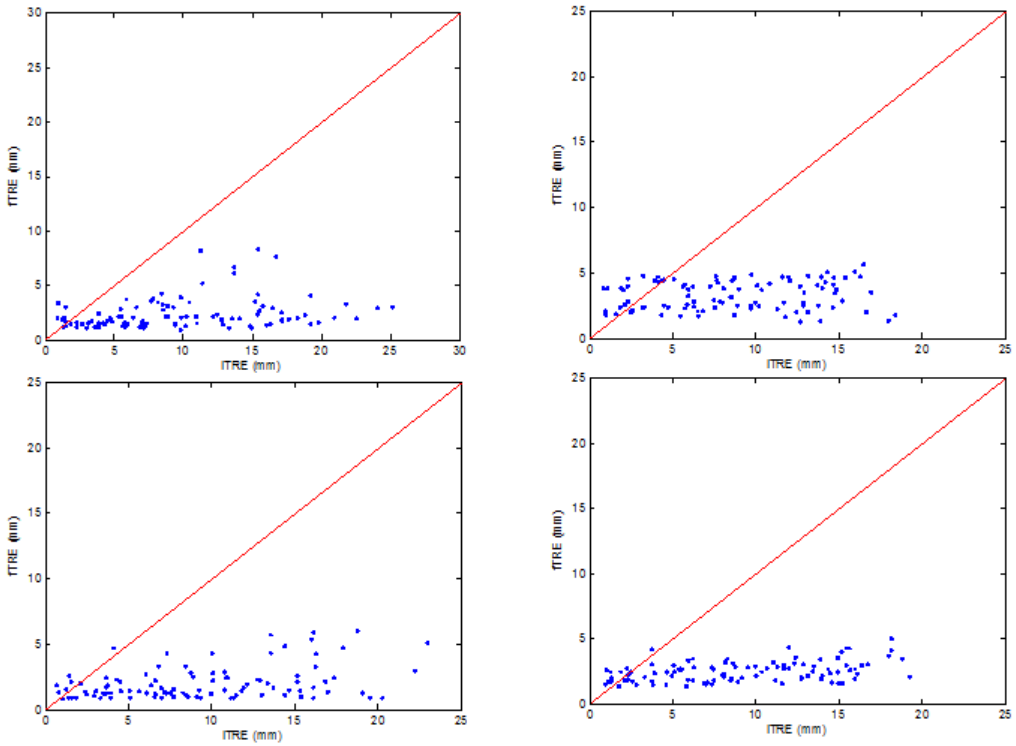

Figure 7. Results for group-wise volume-to-volume registration using CMA-ES as the optimization strategy. Final target registration error (TRE) is plotted against initial TRE for L3 (Top Left), L4 (Top Right), L5 (Bottom Left), and for the entire volume (Bottom Right).

4. CONCLUSION AND FUTURE WORK

In this work, we presented an extended version of registration technique presented by Wein et al. (2007)[6] designed for registration of vertebrae in US and CT data sets. Using tissue density from CT intensities, we simulated US images and automatically identified visible bone surfaces and occluded tissue. We tested the registration by randomly misaligning the CT volume within the intervals of ±10 mm translation along each axis and ±10° rotation about each axis. We compared the results of the registration using three optimization strategies: simplex, gradient descent and CMA-ES. Our algorithm was able to successfully register initial misalignments of up to 20 mm. We found CMA-ES optimization to be more robust than both simplex and gradient descent, though it was less computationally efficient than simplex. Additionally, we presented a novel group-wise registration method incorporating the US simulation. This registration experiment randomly transformed the CT volumes containing individual vertebra and subsequently reconstructed the sub-volumes to a single volume. US simulation was then applied to the volume, which was then used to register the CT dataset to the US volume. We chose to use CMA-ES as the optimization strategy since CMA-ES was found to be the most robust in our rigid volume-to-volume registration. This registration required a better initial alignment, but was still able to reliably register misalignments in individual vertebrae of up to an initial TRE of 8 mm. The middle vertebra (L4) was the source of the largest error in the registration. This is likely due to the occlusion of tissue at the overlap of the vertebrae and the lack of penalty for collision of bone in the sub-volumes.

In our future work, we plan to explore methods of improving the registration runtime. As suggested in Shams et al. (2008) the nature of this type of simulation lends itself well to multithreading and could see great improvements in computation time with a GPU-based implementation[8]. We are currently working on combining a biomedical model of the lumbar spine with the registration algorithm to constrain the relative motions of the vertebrae to biologically realistic orientations.

REFERENCES

[1] Muratore, D.M. , Russ, J.H., Dawant, B.M. and Galloway, R.L., “Three-Dimensional Image Registration of Phantom Vertebrae for Image-Guided Surgery: A Preliminary Study,” Computer Aided Surgery, 7, 342-352, 2002.

[2] Moghari, M.H. and Abolmaesumi, P., “A Novel Incremental Technique for Ultrasound tact Bone Surface Registration Using Unscented Kalman Filtering,” Proc. MICCAI 3750, 197–204, 2005.

[3] Penney, G.P., Barratt, D.C., Chan, C.S.K., Slomczykowski, M., Carter, T.J., Edwards, P.J. and Hawkes, D.J., "Cadaver validation of intensity-based ultrasound to CT registration," Medical Image Analysis, 10, 385-395, 2006.

[4] Brendel, B., Winter, S., Rick, A., Stockheim, M. and Ermert, H., "Registration of 3D CT and Ultrasound Datasets of the Spine Using Bone Structures," Computer Aided Surgery, 7, 146-155, 2002.

[5] Brendel, B., Siepermann, J., Winter, S. and Ermert, H., “In vivo evaluation and in vitro accuracy measurements for an ultrasound-CT registration algorithm”, Proc. CARS 1281, 583-588, 2005.

[6] Wein, W., Khamene, A., Clevert, D., Kutter, O., and Navab, N., "Simulation and fully automatic multimodal registration of medical ultrasound," Proc. MICCAI 4791, 136–143, 2007.

[7] Wein, W., Brunke, S., Khamene, A., Callstrom, M. and Navab, N., “Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention,” Medical Image Analysis 12, 577-585, 2008.

[8] Shams, R., Hartley, R. and Navab, N., “Real-time Simulation of Medical Ultrasound from CT Images,” Proc. MICCAI 5242, 734–741, 2008.

[9] Jensen, J.A., “ Field: A program for simulating ultrasound systems,” Proc. 10th Nordic-Baltic Conference on Biomedical Imaging Published in Medical & Biological Engineering & Computing, 351–353, 1996.

[10] Jensen, J.A. and Svendsen, N. B., “Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers”, IEEE Trans. Ultrason., Ferroelec., Freq. Contr., 39, 262-267, 1992.

[11] Chen, T.K., Thurston, A.D., Moghari, M.H., Ellis, R.E. and Abolmaesumi, P., "A real-time ultrasound calibration system with automatic accuracy control and incorporation of ultrasound section thickness", Proc. SPIE 6918, 69182A, 2008.

[12] Hansen, N. and Ostermeier, A., "Completely Derandomized Self-Adaptation in Evolution Strategies", Evolutionary Computation, 9(2), 159-195, 2001.

[13] Roche, A., Pennec, X., Malandain, G. and Ayache, N., “Rigid registration of 3D ultrasound with MR images: a new approach combining intensity and gradient information,” IEEE Trans. Med. Imag. 20, 1038–1049, 2001.