

Ultrasound Decompression for Large Field-of-View Reconstructions

C. Schulte zu Berge¹, M. Salehi^{1,2}, F. Bender³, and W. Wein¹

¹ImFusion GmbH, Munich, Germany

²Chair for Computer Aided Medical Procedures, Technische Universität München, Munich, Germany

³piur Imaging GmbH, Vienna, Austria

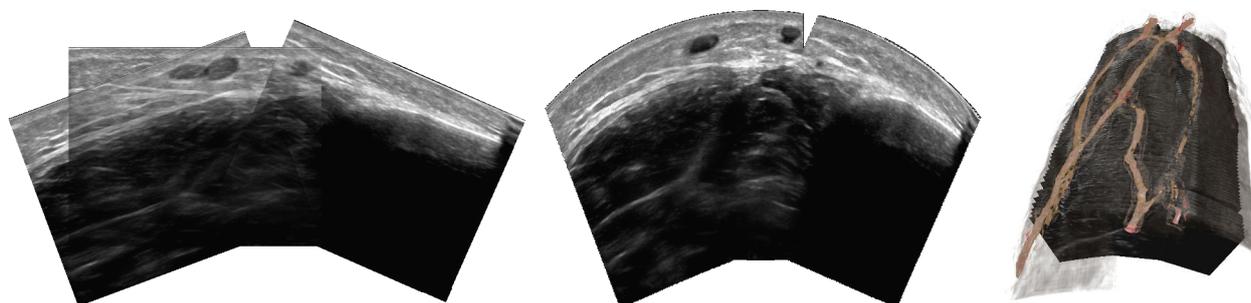


Figure 1: Application of our technique to vascular ultrasound. Due to tissue compression artifacts the original acquisition (left) exhibits poor alignment of important anatomical features. For instance, the same vessel is shown twice at different locations. Our decompression and image stitching technique estimates the original shape of the anatomy and provides a plausible extended FOV reconstruction of the anatomy (middle). This can for instance be used for showing the whole vessel tree (right) providing important spatial context.

Abstract

Tracked medical ultrasound allows for cost-effective and radiation-free imaging of anatomy featuring a very high spatial resolution. To overcome the limitations of the small field-of-view, sonographers can acquire multiple adjacent sweeps and compound them into a single volumetric representation. However, due to the inherent local and non-uniform compression of the underlying anatomy (caused by the ultrasound probe) the adjacent sweeps often exhibit poor alignment and discontinuities. We propose a novel decompression model to compensate for probe pressure related artifacts. It incorporates domain knowledge of the global acquisition pattern for regularization and allows for seamless stitching of multiple overlapping 3D freehand sweeps into one volume. The resulting extended field-of-view visualization provides the clinician with spatial context so that the relationship between individual features are easier to understand. Our experiments show that the resulting extended field-of-view reconstructions have a superior image quality in terms of alignment and continuity of the visible anatomy compared to the original acquisitions. Comparison to ground truth MRI data demonstrates the plausibility of our non-rigid decompression model.

1. Introduction

Traditional 2D ultrasound is an established imaging modality in today's clinical practice. Its cost-effectiveness, non-invasiveness, and flexibility make it attractive for both diagnostic and interventional applications. Different approaches have been introduced to extend ultrasound imaging to the third dimension. Matrix array probes acquire a full 3D volume at once and are established in clinical practice for certain applications in spite of being rather expensive, bulky and providing only a limited field-of-view. In the recent past, 3D freehand ultrasound where a 2D probe is equipped with 6-DOF tracking hardware has more and more proven its poten-

tial as a lightweight alternative. State-of-the-art ultrasound spatial compounding techniques allow for accurate and high-quality volumetric reconstructions from the originally irregularly sampled 2D+t data [SLT*07].

However, since optimal acoustic coupling requires continuous contact with the skin surface the sonographer is required to apply a certain force during the acquisition to achieve a satisfying image quality. The pressure exerted by the transducer leads to a non-rigid deformation and compression of the underlying anatomy. As a result, 3D freehand ultrasound acquisitions are typically restricted to either angular-swept motion, or linear trajectories perpendicular

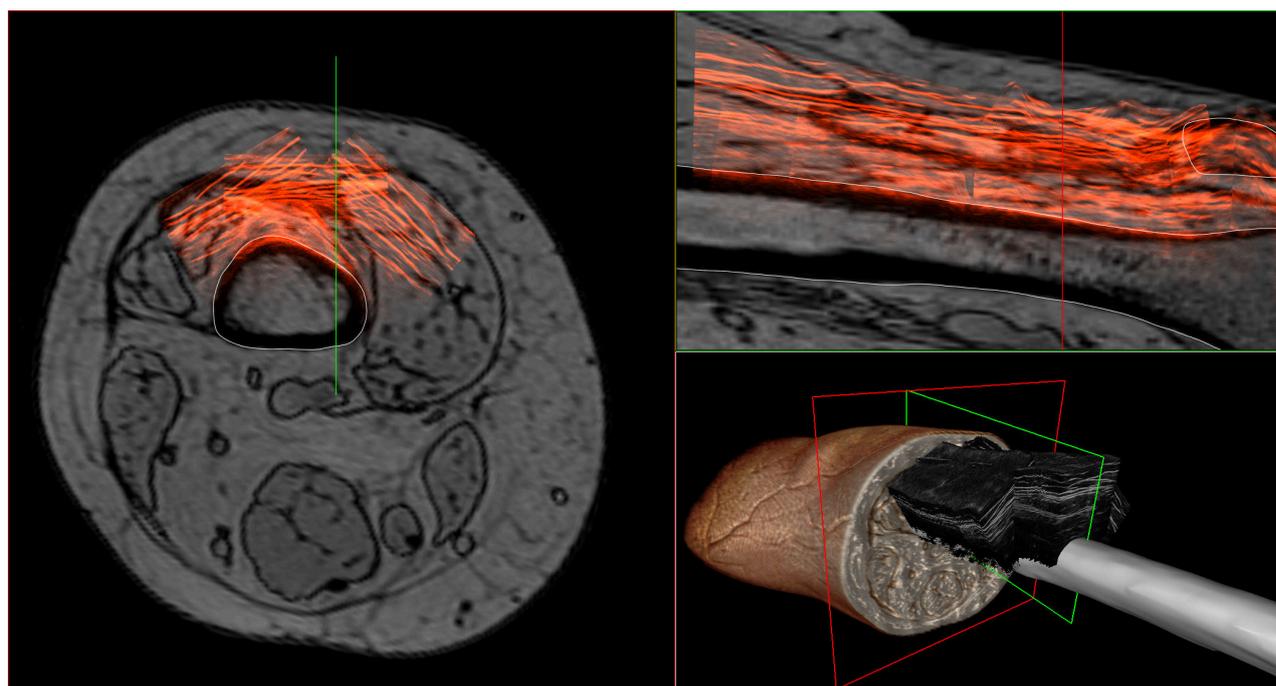


Figure 2: Illustration of a use case for our technique. Multi-modal rendering of MRI and 3D freehand ultrasound acquisitions of the human femur. Each of the three ultrasound sweeps has been co-registered to the MRI based on the bone surface. However, due to the applied probe pressure the underlying anatomy has been compressed and deformed. Thus, the features above the bone surface shown in both MRI and US do not align. Neither do they align between the individual ultrasound sweeps.

to the long side of the transducer so that adjacent frames exhibit similar compression artifacts and the resulting 3D reconstruction is formed of continuous anatomical features. While there are theoretically no limits on the length of such trajectories, such acquisitions suffer from a limited vertical field-of-view (FOV) which is bound by the width and lateral opening angle of the used transducer. As soon as the target anatomy can not be covered by a single sweep, the sonographer has to resort to acquiring multiple parallel sweeps. However, the image content of the overlapping areas will not align due to the inhomogeneous compression artifacts (cf. Figure 2).

Traditional image-based deformable registration techniques [CMG17] might be able to compensate for visual discontinuities, however, they usually lack the domain-specific regularization so that the results do not necessarily represent the actual shape of the anatomy. We propose a decompression model to estimate the original shape of the underlying anatomy. It solely relies on the acquired ultrasound data as input but at the same time takes domain knowledge into account. Thereby, we can generate high-quality extended FOV reconstructions of multiple ultrasound sweeps, which provide a realistic spatial context for visual computing applications.

2. Related Work

A general overview of the ultrasound processing and visualization pipeline can be found in [BSH*12]. To our knowledge there is no prior work that uses a decompression model to reconstruct the original shape of a set of 3D ultrasound acquisitions without using any

additional prior knowledge. However, there are several topics that are related to our work:

Ultrasound mosaicking describes techniques to seamlessly combine several 2D or 3D ultrasound images to generate an extended field-of-view and/or resolve incomplete information in acoustic shadows by scanning the anatomy from different angles. Early works such as [GTP*03] or [PR06] derive the global alignment of multiple images from a sequence of pairwise registrations. Wachinger et al. propose to use global optimization schemes, which are usually better suited to solve the underlying complex non-linear registration problem [WWN07]. However, all these methods do only apply a rigid registration or use a very local block-based model that does not consider any compression artifacts from the acquisition process.

Different approaches have been proposed for the *detection of ultrasound compression*. Elastography imaging [GDFT13] provides different means of estimating the local stiffness and other mechanical properties of the anatomy by using special scanning protocols. In a very recent work Virga et al. investigate the usage of pressure sensors of robotic ultrasound for both estimating and correcting the induced tissue compression [VGB*18].

Once information on the tissue properties has been acquired, *bio-mechanical models* allow for detailed prediction of deformations. However, they come with high computational costs and a large number of parameters. Common approaches include finite element methods [BHN01] as well as 3D mesh-based methods [PM15].

Dahmani et al. recently proposed to use Mutual Information to derive a more lightweight deformation model [JD17].

The majority of the existing work on the topic of ultrasound pressure compensation tries to find a physically correct estimation of the deformation so that measurements (e.g. length, volume) can be performed on the reconstructed volume. However, we argue that physical measurements usually happen on a very focused region of interest that fits into the field-of-view of a single acquisition and therefore they are not the main use case of such techniques. Instead, the main benefit of extended FOV imaging is that it provides the clinician with spatial context of the whole anatomy so that relationships between individual features are easier to understand [KCK*03]. Furthermore, the increased FOV may also allow ultrasound to have a more similar appearance to tomographic imaging modalities such as CT/MRI so that findings can be related more easily [HSK*03].

Therefore, in this work we do not focus on providing fully physically accurate reconstructions of the original anatomy that would allow for measurements. Solving this highly challenging inverse problem would require sophisticated biomechanical models of the underlying anatomy and the exerted forces, and might even be impossible without any additional input knowledge. Instead, we target a mathematical model that is simple enough to be applicable in real-time visual computing scenarios yet is capable of yielding plausible visual results to enable the clinician to comprehend correct spatial relationships.

3. Methods

The proposed decompression model was developed for ultrasound acquisitions using straight linear array transducers where a set of adjacent 3D freehand sweeps were acquired in order to gain an extended field of view. Our goal was to keep the number of parameters as minimal as possible and the overall model simple enough so that it can be evaluated in real-time. Furthermore, it should not rely on any additional input apart from the set of 3D freehand ultrasound sweeps. We assume the sweeps to be co-registered which is implicitly the case if the tracking hardware is well calibrated and the imaging subject does not move during the acquisition.

3.1. Sweep Trajectory Analysis

Due to the nature of the underlying decompression model, our method works best if all ultrasound frames share roughly the same orientation (i.e. the maximum angle between the normal vectors of two frames is less than 45°). Furthermore, the trajectory of each sweep should be formed of a single linear motion into one direction. Since most ultrasound compounding techniques have similar limitations, these are rather common restrictions and most sonographers acquire sweeps of this kind anyway.

Nevertheless, we employ an optional pre-processing step where we analyze the sweep trajectory to identify common deviations, such as

- **Loops:** frames where the forward velocity is negative. They usually occur when the sonographer has lost track of the target feature and therefore moves the probe back until the target is back in the focus region

- **Strafing motion:** frames where the sideways velocity is larger than the forward velocity. They usually occur when the target anatomy is no longer in the center of the field-of-view and the sonographer therefore performs a sideways motion with the probe.

If our sweep trajectory analysis detects such events it will remove duplicate frames and optionally also split the sweep into separate sub sweeps if needed.

3.2. Sweep Decompression Model

We assume a simplified deformation model, where the compression only happens in-plane and in the direction of the transducer. This tissue compression can be described by two main features:

- **Compression amount:** the total amount of downwards displacement of the skin surface in the direction of the applied pressure.
- **Stiffness distribution:** based on the local tissue properties the displacement amount decreases with increasing depth. At a certain depth no more compression takes place.

We derive our decompression model from these observations and will introduce two parameters corresponding to these features. Additionally, we use the fact that the top parts of all ultrasound frames build one continuous (skin) surface as a regularization. Our model approximates this shape with a set of local circle segments fitted to the top edges of each ultrasound frame. The top edge of each frame will be displaced to reside on its corresponding circle segment while the bottom edge of each frame will remain undisplaced. In between we model non-linear distribution of compression along depth using an exponential decay factor.

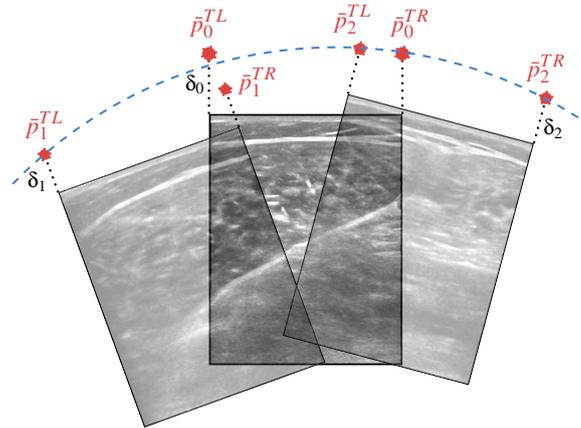


Figure 3: Illustration of the circle fitting method for a single plane P . For each frame i , we project the top left and top right corners (red points) displaced by the vertical displacement parameter δ onto P . We then fit a weighted circle to these points where the weighting factor (here illustrated as point size) is defined by the distance to P .

3.2.1. Circle Fitting

The circle fitting is done for each ultrasound frame independently. Every frame defines a plane P in which we fit a weighted circle defining the local shape of the skin surface. Since by definition all ultrasound top edges are part of the skin surface, the circle's support points are determined by the top corners of the surrounding ultrasound frames while the weighting is given by the distance to the plane. In order to model the *compression amount* we introduce vertical displacement δ as a parameter. The vertical displacement is added to the top edge of each ultrasound frame prior to the circle fitting. It is possible to define one global δ for all frames. However, our experiments showed that defining a separate δ_s for each acquired sweep s yields superior reconstruction results since the applied pressure may vary between sweeps and also the underlying tissue may exhibit a different compressibility all-together.

Given a plane P , we first compute for each frame i the distance d_i of the top center point to P (for the frame from which P originates, the distance naturally is 0). This distance provides us with the weighting factor, which we define by normalizing the distance w.r.t. a maximum distance d_{\max} that describes the smoothness of the reconstructed surface:

$$w_i = \frac{d_{\max} - |d_i|}{d_{\max}}. \quad (1)$$

Our experiments show that using the image width as maximum distance yields good results.

For each frame i we displace both the top left and top right corner by δ and project the resulting point onto P . The projected points \bar{p}_i^{TL} and \bar{p}_i^{TR} will serve as support points using w_i as weights. Thus, the set of support points for the circle fitted to plane P is given by

$$S_P = \bigcup_i \left\{ \left(\bar{p}_i^{TL}, w_i \right), \left(\bar{p}_i^{TR}, w_i \right) \right\}, \quad (2)$$

where i iterates over all frames, as illustrated in Figure 3. We use the Levenberg-Marquardt algorithm for weighted least squares fitting of the circle to S_P .

Even though the presented circle fitting is a rather simple local model defining the surface shape of each frame individually, the used weighting scheme is symmetric for parallel frames. As long as the frames of all sweeps share roughly the same orientation, the set of circles forms a smooth and homogeneous reconstruction of the skin surface shape without requiring an expensive global regularization step (cf. Figure 4).

3.2.2. Computation of the Displacement Field

Once the outer shape of the skin surface has been determined by the circle fitting, we can use it as the base for the decompression displacement field. In our model we assume that the top edge of each ultrasound frame is displaced in upwards direction to match the curvature of the local circle segment. Given the general circle equation $(p_x - c_x)^2 + (p_y - c_y)^2 = r^2$ for a point $p = (p_x, p_y)$, circle center $c = (c_x, c_y)$ and circle radius r , we can compute the location of the top circle edge t_p for each pixel p as

$$t_p = \sqrt{r^2 - (p_x - c_x)^2} + c_y. \quad (3)$$

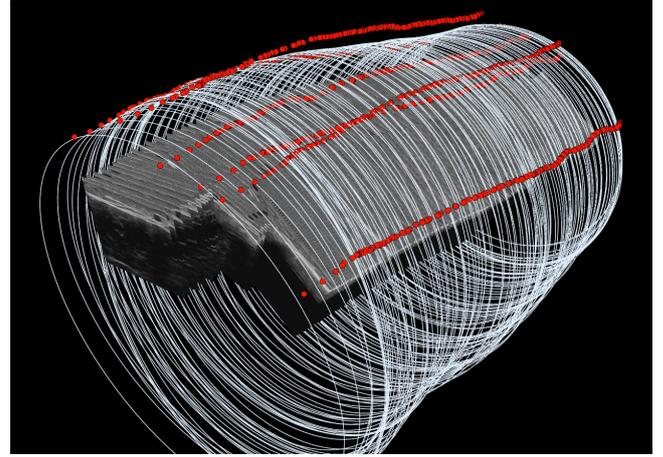


Figure 4: Illustration of the circle fitting applied to a set of 3 sweeps: Similar to Figure 3, the red points indicate the support points \bar{p}^{TL} and \bar{p}^{TR} generated from the top left and top right corners. Due to the symmetric weighting scheme the fitted circles form a smooth surface. The ultrasound sweeps are shown in their original compressed form.

We model the bottom edge of each ultrasound frame to have zero displacement by default. As observed in earlier there is a non-linear decline in compression with increasing depth. We model this in an exponential fashion by introducing the stiffness s as a second parameter. With h being the height of the ultrasound frame and $p_y \in [0, h]$ the vertical position (inverse depth) of the pixel, the final vertical position p'_y for point p is given by

$$p'_y = p_y + \left(\sqrt{r^2 - (p_x - c_x)^2} + c_y - p_y \right) \cdot \left(\frac{p_y}{h \cdot t_p} \right)^s. \quad (4)$$

The stiffness parameter s allows us to model the distribution of the compression along the depth of the image:

- $s = 1$: the amount of displacement is interpolated linearly between the top edge (maximum displacement) and the bottom edge (zero displacement).
- $s > 1$: the displacement will concentrate on the top part of the image, while the bottom edge remains undistorted.
- $s < 1$: the bottom edge will start to be displaced as well. Moving $s \rightarrow 0$ will approach a constant displacement modeling that the original tissue compression did occur below the image.

Since both vertical displacement δ and stiffness s relate to very intuitive terms it is a viable approach to have these parameters to be user-controlled and/or given by application-specific presets. However, if there is sufficient overlap between the individual sweeps, we can optionally use image-based similarity metrics to automatically optimize the parameters. Therefore, we compute the Local Normalized Cross-Correlation (LNCC) with a patch size of 9x9 pixels on the overlapping regions. To speed up the computation we limit the number of compared frames and distribute the sparse samples equally over the length of the reference sweep. Our experiments show that this automatic parameter optimization is particularly ef-

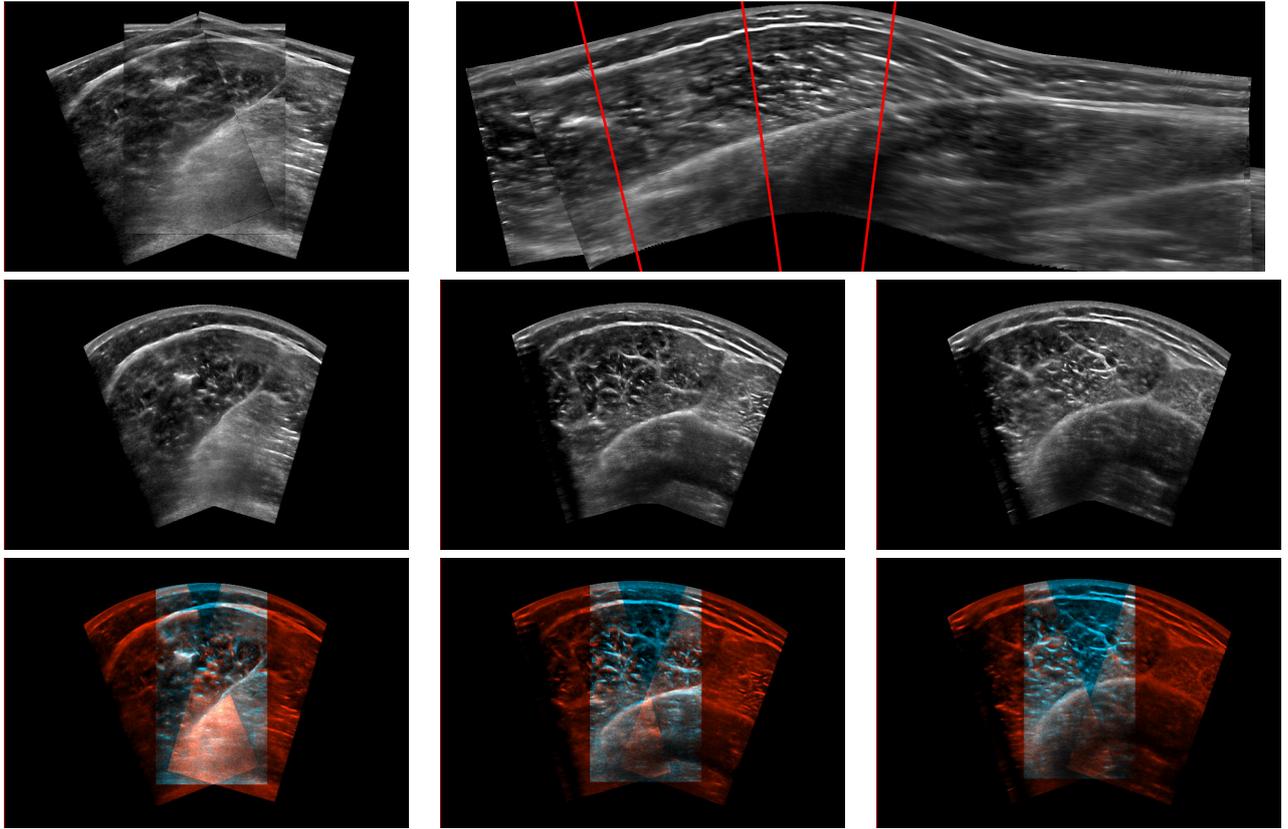


Figure 5: Sweep decomposition results for the MSK data set. The top left image shows the original configuration of the 3 acquired sweeps on the left. The top right image shows the sagittal reconstruction including indicators of the exemplary axial slices shown below. The middle row shows axial MPRs through the data set after decomposition and stitching. The bottom row shows a special color blending mode that allows for assessing the alignment differences.

fective when we use distinct values of δ and s for each acquired sweep.

3.3. Volume Reconstruction

Our decompression technique can be integrated into most state-of-the-art algorithms for ultrasound compounding (including both forward-warping and backward-warping methods as described in [SLT*07]) in order to generate a volumetric representation from the irregularly sampled 3D freehand ultrasound data. Since the model is solely working in-plane for each frame, the logic for transforming between the sweep geometry and the rectilinear target grid can be kept. The only required modification is to consider the decompression offset during the intensity lookup. This is a simple application of Equation 4 yielding an offset for the lookup coordinates which can be easily implemented in either CPU and GPU code.

After compounding each ultrasound sweep individually, they are stitched together to yield the final pixel/voxel value. To avoid visible seams we do not use traditional global alpha blending or median compounding but use a distance-based alpha blending scheme instead. For each ultrasound intensity I we determine the distance d_I to the center of the closest ultrasound frame. This can be easily

computed as a byproduct of the compounding step. During stitching we define the local (per-pixel) blending factor $\alpha_{B,F}$ for given background and foreground intensities B and F with distance values d_B and d_F as

$$\alpha_{B,F} = \frac{d_F - d_B}{2 \cdot \max\{d_B, d_F\}} + 0.5. \quad (5)$$

The resulting pixel intensity is then determined by standard alpha compositing.

4. Implementation

We implemented the presented technique in the ImFusion SDK, a commercial framework for medical imaging and accelerated computing. The computation of the circle support points as well as the circle fitting is performed on the CPU. All further processing steps are handled entirely on the GPU using OpenGL 4.3. We store the per-frame deformation parameters in Shader Storage Buffer Objects so that the local displacement can be computed directly during the spatial compounding, for which we implemented a standard backward compounding scheme using inverse-distance weighting.

The test system equipped with an Intel i7-6700K CPU @ 4 GHz

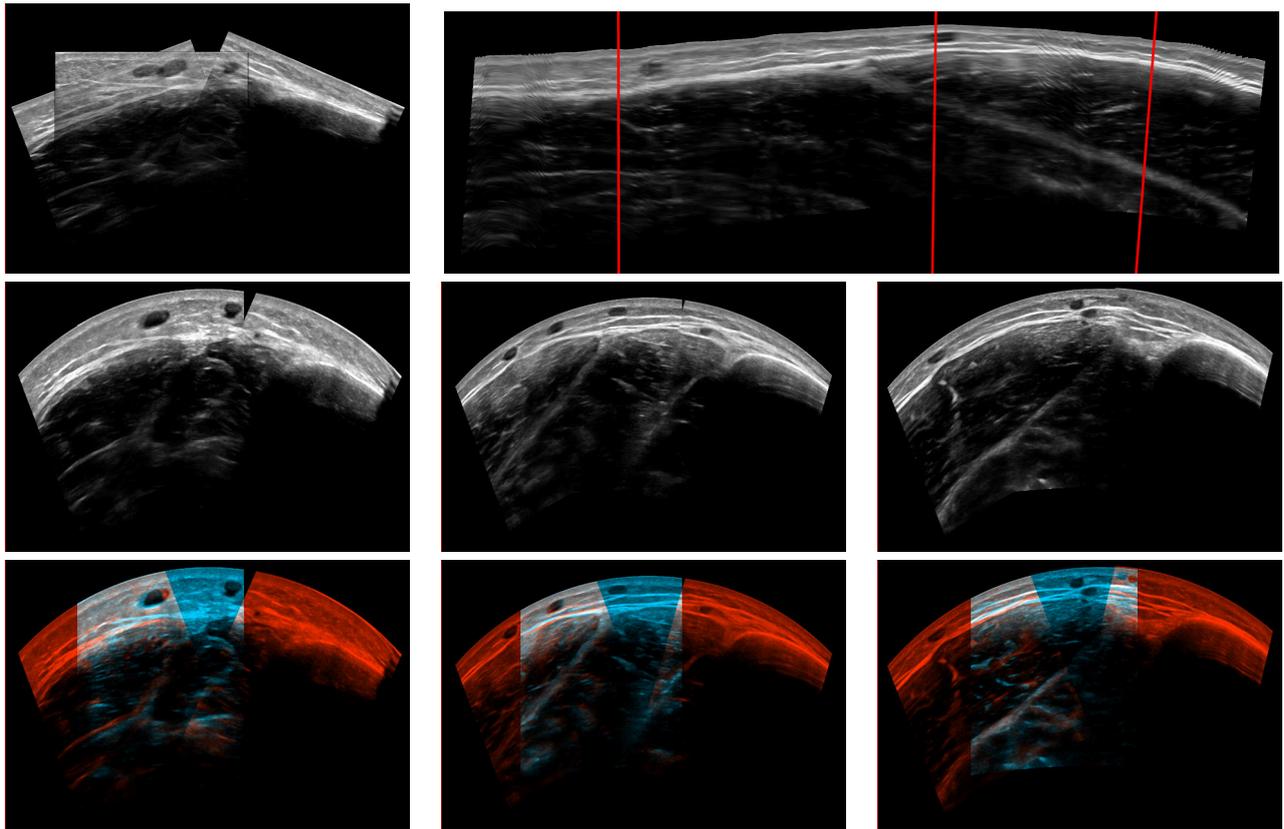


Figure 6: Sweep decomposition results for the vascular data set. The top left image shows the original configuration of the 3 acquired sweeps on the left. The top right image shows the sagittal reconstruction including indicators of the exemplary axial slices shown below. The middle row shows axial MPRs through the data set after decomposition and stitching. The bottom row shows a special color blending mode that allows for assessing the alignment differences.

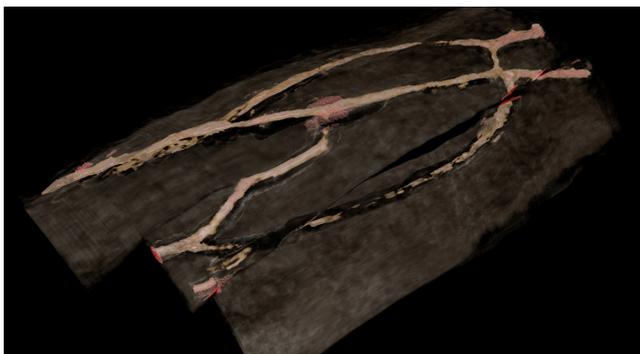


Figure 7: Exemplary use case for the results of the proposed technique: The stitched 3D reconstruction of the decompressed sweeps from Figure 6 can be used to extract a topologically correct vessel tree, which is shown here using traditional direct volume rendering.

and an nVidia GeForce GTX 970 was capable of maintaining interactive frame rates for all data sets. The most expensive part of the algorithm is the circle fitting, which takes an average of 48ms for

a representative data set consisting of 3 sweeps with a total of 608 ultrasound frames. The full processing time including computation of the final displacement field and backward compounding of an axial MPR in a 1920x1080 pixel viewport takes about 73ms. However, the circle fitting is only needed to be recomputed when the vertical displacement parameter has changed (e.g. during parameter optimization). Once the decomposition parameters are fixed, the overhead for applying the displacement during rendering is negligible since the displacement data only needs to be uploaded once. Therefore, when browsing the MPRs of the data set after parameter optimization, our system easily exceeds 30 fps.

5. Evaluation/Results

The majority of the used data sets were acquired using the *piur tUS* system for 3D freehand ultrasound (piur Imaging GmbH). It extends off-the-shelf medical ultrasound devices with a frame grabber to capture the image data and electromagnetic 6-DOF tracking hardware to gather the 3D pose of each frame. Calibration of the tracking system and the ultrasound probe was performed with [WK08].

5.1. Qualitative Results

We applied our technique to ultrasound sweeps of different anatomy using individual vertical displacement and stiffness parameters for each sweep. All parameters were determined in a fully automatic fashion using the similarity metric-based optimization scheme on 50 reference planes equidistantly spread over the trajectory. Thus, no user interaction was necessary. The whole computation time to yield the shown results was less than 15 seconds.

Figure 5 shows a musculoskeletal (MSK) data set consisting of three adjacent sweeps featuring approximately 50% overlap. Since the images show mainly a single muscle layer, the original compressed reconstruction (top left image) already exhibits very good continuity of the lower anatomical structures. However, the top skin and fat layers show rather poor alignment. After decompression all main features are well aligned. The difference images in the bottom row show only small errors. They can partially be explained by the fact that the sweeps have different sonification angles (the ultrasound point-spread function is angle-dependent).

Figure 6 shows a vascular data set of the leg also consisting of three adjacent sweeps. The left and the center sweep share roughly 60% overlap while the right and center sweep have roughly 25% overlap. The original reconstruction (top left image) shows poor alignment of the anatomical structures both in the upper and lower parts of the image. This is particularly problematic in case of the vessels since the same vessel is shown twice at different locations meaning that the depicted vessel topology is wrong. After decompression, vessels as well as the muscle layers in the deeper parts of the image are well aligned and each vessel is shown only once. This enables us for instance to reconstruct a full 3D volume of all sweeps combined, and extract a segmentation of the full vessel tree. This can then be shown using direct volume rendering to provide the clinician with a spatial overview of the anatomy (cf. Figure 7).

5.2. Parameter Study

To gain a better understanding of the two parameters, we conducted a small parameter study on the vascular ultrasound data set. Figure 8 shows the same ultrasound frame with varying combinations of δ and s , added annotations of the fitted circle, and a ruler overlay. It is clearly visible how the vertical displacement δ pulls the top edge of the frame up and how the stiffness parameter s affects the non-linear distribution of the image content.

5.3. Ground Truth Comparison

To evaluate the accuracy of our method w.r.t. the natural uncompressed shape of tissue, we compared the decompressed ultrasound sweeps with a co-registered MRI volume serving as ground truth. The MRI data was acquired from the upper leg of a subject shortly before recording the tracked ultrasound sweeps. The two modalities were registered to each other based on the bone surface of the femur as described in [SPM* 17] (cf. left column of Figure 9).

We ran two experiments for defining the parameters of the decompression model:

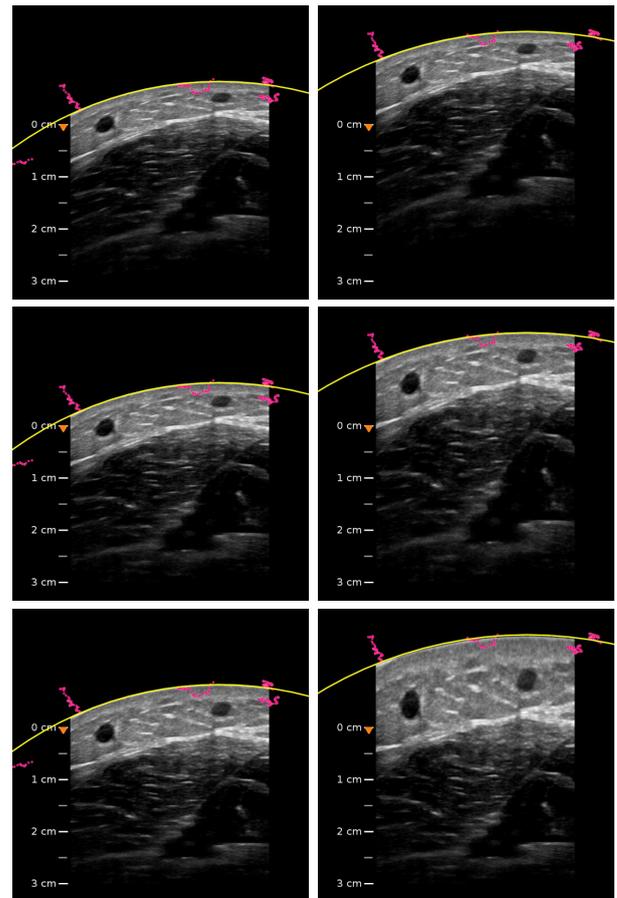


Figure 8: Influence of the two parameters on the decompression: Left column: $\delta = 5\text{mm}$, right column: $\delta = 15\text{mm}$. Top row: $s = 0.4$, middle row: $s = 1.0$, bottom row: $s = 2.0$. The yellow line indicates the circle segment, the red dots show the support points used for circle fitting.

- (a) In the first experiment we wanted to investigate how well the technique performs with manually selected parameters (since automatic parameter optimization is prone to locking into local minima). Therefore, we manually chose optimal values for both vertical displacement δ and stiffness s solely based on the feature alignment in the ultrasound images without considering the MRI volume.
- (b) In the second experiment, we extracted the skin surface from the MRI volume and use it as input for the circle fitting: Instead of using the top corners of the ultrasound frames for the support points we fit each circle to the intersection of its plane with the skin mesh. Then, the stiffness parameter s was automatically optimized by maximizing the image similarity in the overlapping section of the ultrasound.

The results are shown in Figure 9. The bone surface was used for co-registration. Thus, the original compressed configuration (left column) naturally has the best alignment of this feature. How-

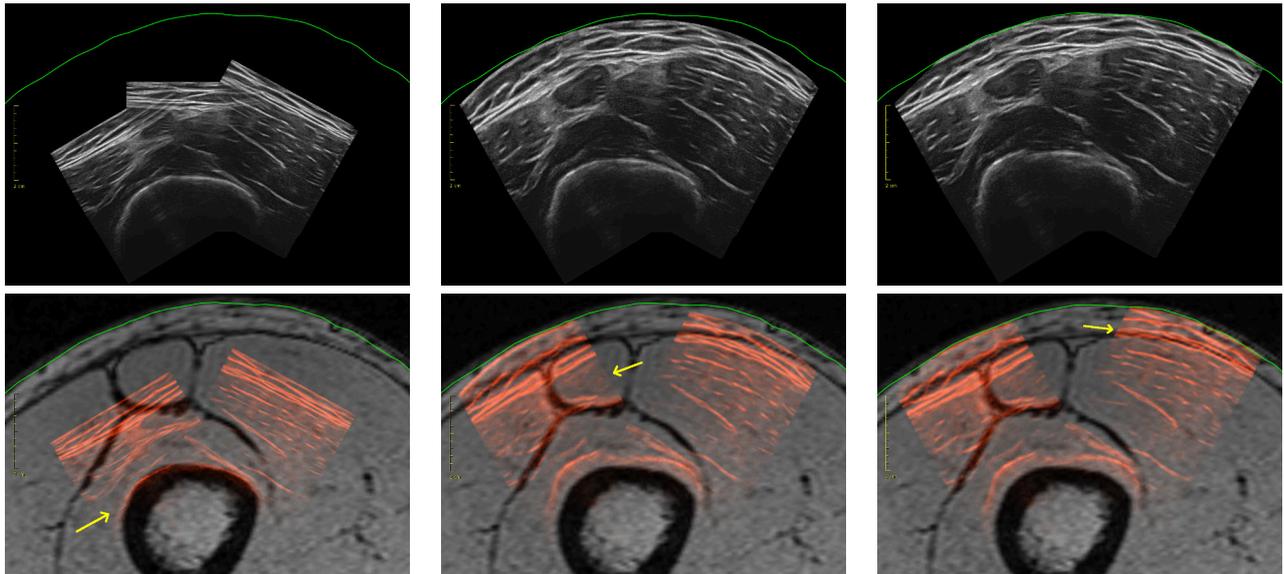


Figure 9: Comparison of sweep decomposition results with ground truth data. Top row images show the stitched ultrasound frames and the ground truth skin surface (green line). Bottom row images additionally show the ground truth MRI data fused with the US frames (the middle sweep is removed for the sake of visualization) to assess the alignment of individual anatomical features. Left column shows the original ultrasound data; middle and right columns show the results of experiment (a) and (b), respectively.

ever, all other anatomical features of the fat and muscle layers are severely misaligned.

The center column shows the results of experiment (a). The most notable finding is that even though the decompression parameters were chosen completely independent of the MRI, the reconstructed skin surface closely matches the ground truth (green line). Furthermore, the muscle borders in the center of the ultrasound frame (arrow) align well. However, the alignment of the fat-muscle interface at the top could be improved.

The right column shows the result of experiment (b). As expected, the top part of the ultrasound (arrow) now aligns almost perfectly with the MRI.

6. Discussion and Conclusion

We have presented a geometric decompression model and volumetric reconstruction scheme that allows to create arbitrarily large 3D ultrasound volumes composed of multiple overlapping freehand sweeps. It is based on the assumption that all scans are located on a convex portion of the skin surface and consists of two parameters per sweep, denoting the overall decompression at the surface, and the distribution of it within the body, respectively. The evaluation results indicate that the proposed technique is capable of generating realistic reconstructions of the original anatomy, which can then be used to support the user with an extended spatial context.

Our model can be improved and extended in a straight-forward fashion for various clinical applications and organs. We are currently experimenting with means to support additional transducer types such as curved linear array and sector probes. This requires adjusting the displacement vectors according to the frame geometry

as well as compensating for the non-uniform amount of displacement due to the curved shape of the transducer surface. Further directions of future work are the automatic classification of vessels, which may be used as weighting scheme during registration if an improved reconstruction of vascular structures is desired. Likewise, when scanning limbs with a bone prominently featured within the ultrasound images, a bone surface detection as e.g. in [SPM*17] can be used to define the fixed portion of the image where a decompression starts.

The largest simplifying assumption in our model is that we currently only decompress image content vertically per ultrasound frame. However, real anatomy may move sideways and out-of-plane when pushed with an ultrasound transducer; this holds especially for imaging applications where fat and muscle layers with multiple fascia are involved, which tend to slide past each other. Similar to above, some of this motion may be recovered through organ-specific assumptions and classification of anatomical structures. In addition, systematic measurement of tissue deformation with increasing probe pressure [VGB*18] may be used to derive a realistic deformation model, possibly as input to complete biomechanical modeling. Afterwards, parameter reduction techniques such as PCA may yield few significant modes to be optimized with the same method as presented here.

Acknowledgments

We thank Stryker Leibing GmbH & Co. KG, Freiburg, Germany for providing the co-registered MRI-ultrasound data sets used for the ground truth evaluation.

References

- [BHN01] BURCHER M. R., HAN L., NOBLE J. A.: Deformation correction in ultrasound images using contact force measurements. In *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA 2001)* (2001), pp. 63–70. 2
- [BSH*12] BIRKELAND Å., SOLTÉSZOVÁ V., HÖNIGMANN D., GILJA O. H., BREKKE S., ROPINSKI T., VIOLA I.: The ultrasound visualization pipeline - A survey. *CoRR abs/1206.3975* (2012). 2
- [CMG17] CHE C., MATHAI T. S., GALEOTTI J.: Ultrasound registration: A review. *Methods 115* (2017), 128 – 143. *Image Processing for Biologists*. 2
- [GDFT13] GENNISSON J.-L., DEFFIEUX T., FINK M., TANTER M.: Ultrasound elastography: Principles and techniques. *Diagnostic and Interventional Imaging 94*, 5 (2013), 487 – 495. *Ultrasound elastography*. 2
- [GTP*03] GEE A. H., TREECE G. M., PRAGER R. W., CASH C. J. C., BERMAN L.: Rapid registration for wide field of view freehand three-dimensional ultrasound. *IEEE Transactions on Medical Imaging 22*, 11 (Nov 2003), 1344–1357. 2
- [HSK*03] HENRICH W., SCHMIDER A., KJOS S., TUTSCHEK B., DUDENHAUSEN J. W.: Advantages of and applications for extended field-of-view ultrasound in obstetrics. *Archives of Gynecology and Obstetrics 268*, 2 (Jun 2003), 121–127. 3
- [JD17] JAWAD DAHMANI YVAN PETIT C. L.: Model-based correction of ultrasound image deformations due to probe pressure, 2017. 3
- [KCK*03] KIM S. H., CHOI B. I., KIM K. W., LEE K. H., HAN J. K.: Extended field-of-view sonography: advantages in abdominal applications. *Journal of ultrasound in medicine : official journal of the American Institute of Ultrasound in Medicine 22*, 4 (April 2003), 385–94. 3
- [PM15] PHEIFFER T. S., MIGA M. I.: Toward a generic real-time compression correction framework for tracked ultrasound. *International Journal of Computer Assisted Radiology and Surgery 10*, 11 (Nov 2015), 1777–1792. 2
- [PR06] POON T. C., ROHLING R. N.: Three-dimensional extended field-of-view ultrasound. *Ultrasound in Medicine & Biology 32*, 3 (2006), 357 – 369. 2
- [SLT*07] SOLBERG O. V., LINDSETH F., TORP H., BLAKE R. E., HERNES T. A. N.: Freehand 3d ultrasound reconstruction algorithms - a review. *Ultrasound in Medicine & Biology 33*, 7 (2007), 991 – 1009. 1, 5
- [SPM*17] SALEHI M., PREVOST R., MOCTEZUMA J.-L., NAVAB N., WEIN W.: Precise ultrasound bone registration with learning-based segmentation and speed of sound calibration. In *MICCAI (2017)*, Springer, pp. 682–690. 7, 8
- [VGB*18] VIRGA S., GÖBL R., BAUST M., NAVAB N., HENNERSPERGER C.: Use the force: deformation correction in robotic 3d ultrasound. *International Journal of Computer Assisted Radiology and Surgery 13*, 5 (May 2018), 619–627. 2, 8
- [WK08] WEIN W., KHAMENE A.: Image-based method for in-vivo freehand ultrasound calibration. *Proc. SPIE, Medical Imaging 2008 6920* (2008). 6
- [WWN07] WACHINGER C., WEIN W., NAVAB N.: Three-dimensional ultrasound mosaicing. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007* (Berlin, Heidelberg, 2007), Ayache N., Ourselin S., Maeder A., (Eds.), Springer Berlin Heidelberg, pp. 327–335. 2