

Orientation-Driven Ultrasound Compounding Using Uncertainty Information

Christian Schulte zu Berge¹, Ankur Kapoor², and Nassir Navab¹

¹ Chair for Computer Aided Medical Procedures
Technische Universität München, Munich, Germany
christian.szb@in.tum.de, nassir.navab@tum.de

² Imaging and Computer Vision,
Siemens Corporate Technology, Princeton, NJ, USA
ankur.kapoor@siemens.com

Abstract. Compounding 2D ultrasound sweeps into 3D volumes is, due to its cost- and time-efficiency, of great clinical significance in both diagnostic and interventional imaging. However, today's algorithms restrict the sweeps to have homogeneous pressure and a linear trajectory, which limits their use in clinical applications such as breast or musculoskeletal ultrasound where artifacts occur due to soft and uneven surfaces. In this work, we present two techniques to resolve those restrictions by using an orientation-driven approach, first compensating for probe pressure changes and then resolving ambiguities in regions, where multiple ultrasound frames from different acoustic windows overlap. After clustering incoming frames by orientation, we determine the final voxel intensities based on per-pixel uncertainty information. Qualitative and quantitative evaluation of our methods shows that these techniques provide reconstructions of superior quality for ultrasound sweeps of inhomogeneous pressure and twisted trajectories. Furthermore, we propose optimizations in the implementation of these techniques towards real-time applications, interactively updating and refining the reconstructed volume.

1 Introduction

Ultrasound spatial compounding is the reconstruction of 3D volumes from 2D ultrasound sweeps and has the potential to replace or extend current standard clinical procedures for several applications, such as breast cancer diagnosis and musculoskeletal (MSK) applications. Here, X-Ray does not only have the drawback of using ionizing radiation but also shows weak tissue contrast. MR imaging is rather slow, expensive and additionally restricts the patient to be in a position that might not be suited well for diagnostics. In contrast, ultrasound is relatively low-cost, portable, real-time capable and offers good soft tissue contrast.

Recent advances in tracking calibration and compounding algorithms have led to a significant increase in image quality of ultrasound compounding: Current ultrasound probe calibration methods achieve millimeter tracking-accuracy and the various spatial compounding algorithms offer different tradeoffs between algorithm complexity and quality of the compounded volume. As a consequence,

ultrasound compounding is making its way into commercial products shaping the term 3D freehand ultrasound.

So far however, the term 3D freehand ultrasound promises more than it actually can deliver, since current methods implicitly assume constraints such as constant probe pressure and/or constant motion of the ultrasound transducer along a linear path. While this may be negligible for applications such as carotid ultrasound where the anatomy is easily accessible, breast and MSK applications have highly curved surfaces requiring reconstruction of twisted sweep trajectories.

Curved sweeps lead to a challenging issue during the spatial compounding process because some of the acquired ultrasound frames may overlap with each other. Due to the dynamics and high complexity of the ultrasound image formation being dependent on incident angle, probe pressure and patient positioning, ultrasound may yield different information (i.e. image intensities) for the same point within the anatomy if scanned from different perspectives or at different times. Our orientation-driven methods handle these ambiguities to result in more accurate 3D reconstructions than current state-of-the art methods.

2 Related Work

In [1] Solberg et al. provide an overview on different 3D ultrasound compounding techniques and identify three different classes of algorithms:

Pixel-based methods traverse the pixels in each 2D ultrasound frame, transform the pixel location into voxel coordinates and write the pixel's intensity information into the initially empty volume. Since multiple pixels might contribute to a single voxel, the final voxel value may be determined by averaging or using the maximum intensity of all contributing pixels.

Voxel-based methods work the other way around by traversing the voxel grid of the target volume and are thus also referred to as backward-warping methods. For each voxel, they compute the corresponding pixels of the nearby ultrasound frames and use a weighting function based on intensity and/or distance to determine the final voxel value. Wein et al. show in [2] that voxel-based methods yield superior quality and smaller computation time than pixel-based methods. Furthermore, backward-warping algorithms can easily be used to compute multi-planar reconstructions (MPR) from the original ultrasound images without computing the reconstructed volume before.

Finally, *function-based methods* estimate the coefficients for a set of locally supported basis functions to approximate the input data. These functions are then evaluated on the voxel grid to reconstruct the compounded volume [3,4]. Klein et al. [5] propose to use radio frequency (RF) data instead of reconstructed B-mode images and a finite mixture model to obtain reconstructions of higher quality and address the view-dependency of ultrasound. While these methods yield 3D ultrasound reconstructions of very high quality, they are currently not feasible for clinical practice due to being computationally expensive.

To compensate for probe pressure changes, Treece et al. use an image-based non-rigid registration technique [6]. By computing the line-wise maximum

normalized correlation between two adjacent B-mode images and applying a monotonicity constraint they estimate the deformation in depth introduced by the probe pressure. To avoid drift in the registration they constrain the registration results to the tracking. However, regularization will fail in case of inaccurate calibration of the ultrasound probe, especially in the error-sensitive rotational part.

3 Methods

3.1 Inter-frame Registration and Pressure Compensation

To correct for errors and inaccuracies in the tracking data (e.g. due to inaccurate calibration or patient movement), as well as for artifacts due to probe pressure changes, we propose an orientation-driven inter-frame registration technique:

Similar to Treece et al. [6], we perform an intensity-based registration between adjacent ultrasound frames. Using a simple and thus real-time capable pixel-wise uphill search evaluating the SSD, each ultrasound frame is registered to its surrounding frames independently in terms of in-plane translation and in-plane rotation. However, we perform the regularization by registering each ultrasound frame to a window W of surrounding frames. This ensures to compensate for drift independently of the tracking calibration accuracy.

Since the correlation between two ultrasound frames does not only depend on their proximity but also on their orientation to each other [7], we determine the weights for the frames in W by a combination of a Gaussian kernel of size N and a term C , which describes the orientation-based correlation between two images. For a given reference patch P and equally sized moving patch P' the windowed SSD (wSSD) is given by

$$\text{wSSD}_{P,P',N}(i) = \sum_{\substack{p \in P, \\ p' \in P'}} \sum_{n=-N}^N C(i, i+n) \cdot e^{-\frac{n^2}{2\sigma^2}} \cdot (I_i(p) - I_{i+n}(p'))^2 \quad (1)$$

where i is the index of the reference frame and $I_i(p)$ denotes the image intensity of ultrasound frame i at the position p . The correlation term $C(i, j)$ for frames i and j is defined by the cosine distance of their normals n_i, n_j to model the decreasing correlation between frames of increasing orientation difference:

$$C(i, j) := 1 - \frac{2}{\pi} \cdot \text{acos} \left(\frac{n_i \cdot n_j}{\|n_i\| \|n_j\|} \right) \quad (2)$$

To compensate for probe pressure artifacts, our method applies the above inter-frame registration technique not only to a single patch, but to a grid of independent patches of $1\text{cm} \times 1\text{cm}$ size. Since we expect the deformation to be orthogonal to the skin surface, our model allows free in-plane movement of the patches to be flexible enough to allow both linear and curvilinear probes. After computing the transformation for each patch as above, we set the transformation of the central patch as rigid part and the difference to the other patches as deformation field. The results can be seen in Fig. 1.




Fig. 1. Reconstruction of an abdominal phantom scan with probe pressure changes: (a) MPR through the compounded volume without applying our pressure compensation technique; (b) the same MPR through the compounded volume with pressure compensation applied

3.2 Compounding of Non-homogeneous Sweeps

In a tortuous acquisition sweep parts of the ultrasound frames overlap and may show different information for the same anatomy depending on the viewing angle. Here, classical ultrasound compounding techniques with averaging or distance-based weighting fail in correctly reconstructing such regions:

Given a set of ultrasound frames from different angles that all intersect near our target voxel to reconstruct as depicted in Fig. 2. Standard compounding algorithms such as [2] take the closest pixels in each ultrasound frame and determine the final voxel intensity based on a weighting function usually preferring closer pixel over pixels being farther away, hence the closest ultrasound frame has the highest influence. If we now consider a neighbor voxel, the closest frame may have a completely different orientation and thus show different information (due to the view dependency of ultrasound). This yields to artifacts in the compounded volume as depicted in Fig. 3a.

Furthermore, distance-based weighting can lead to incorrect reconstruction since the distance of the frame to the voxel has no correlation with the amount of information present in this pixel (i.e. level of uncertainty/noise). For instance, it may ignore a pixel being farther away but having low uncertainty and instead prefer a high uncertainty pixel (i.e. noise) because it is closer to the voxel.

Our orientation-driven ultrasound compounding technique tackles these issues by exploiting additional uncertainty information using a two-step approach. Our method assumes that for each ultrasound pixel with intensity I_i , we also have an uncertainty value u_i that we later use for weighting the image intensities. While the actual method is independent from it, we use for our implementation the attenuation maps proposed by Karamalis et al. [8]. Even though they model ultrasound physics only to a limited amount, their attenuation maps can be interpreted as uncertainty information.

Clustering of the Ultrasound Sweep by Direction: As a first step, we perform a hierarchical clustering to identify tortuous sweep trajectories and regions of overlapping ultrasound frames. This partitions the ultrasound sweep




Fig. 2. Illustration of artifacts occurring in distance-weighted compounded regions where multiple ultrasound frames from different angles intersect. The intensity of voxel v_1 will be mainly influenced by the information of frame f_3 while the intensity of neighbour voxel v_2 will be mainly influenced by the information in frame f_1 . Since the frames travel through different acoustic windows, the information at this spatial location may significantly differ.




Fig. 3. Effect on the clustering of ultrasound frames by normal direction: (a) shows a compounding of a twisted ultrasound sweep with artifacts caused by the filtering based on the distance to the voxel. (b) shows a compounding of the same sweep with our clustering technique applied.

trajectory into parts where the frames have homogeneous orientation without requiring us to predefine the number of clusters. We apply an average group linkage algorithm using cosine distance to the normals of the ultrasound frames Eq. (2). This yields a set of sub-sweeps meeting the usual restriction of being contiguous and uniformly oriented.

A backward-warping algorithm then compounds each cluster c into a 3D volume applying our pressure compensation method as discussed in Section 3.1. Since the ultrasound frames of each cluster are guaranteed to have the same orientation and are thus travelling through the same acoustic window, we can safely assume the distributions of uncertainty within the frame to be homogeneous within nearby frames. We compute the intensity for voxel x as

$$I_c(x) = \frac{\sum_{i \in S} I_i \cdot d_i^{-\mu}}{\sum_{i \in S} d_i^{-\mu}} \quad (3)$$

where S is the set of frame pixels close to the compounded voxel x , d_i the Euclidean distance of pixel i to the compounded voxel and $\mu > 1$ a smoothness parameter ensuring that $I_c(x)$ approximates the original data for $d_i \rightarrow 0$ [9]. Furthermore, we propagate the uncertainty to the 3D volume using the same weighting:

$$U_c(x) = \frac{\sum_{i \in S} u_i \cdot d_i^{-\mu}}{\sum_{i \in S} d_i^{-\mu}} \quad (4)$$

Uncertainty-Based Fusion of the Compounded Clusters: Since ultrasound image formation is a highly non-linear process and the pixel-based uncertainty values u_i are relative to the image content and thus not necessarily comparable between different frames, we perform the uncertainty-based fusion in a second step to avoid artifacts such as the ones depicted in Fig. 3a. In this second step our method fuses clusters into the final 3D volume based on the propagated uncertainty values. Let C be the set of clusters, then the final intensity I at voxel x is given by

$$I(x) = \frac{\sum_{c \in C} (1 - U_c(x)) I_c(x)}{\sum_{c \in C} 1 - U_c(x)} \quad (5)$$

4 Implementation

Our implementation of orientation-driven ultrasound compounding employs several optimizations to allow real-time applications such as an interactive update and refinement of the compounded volume: The regularized inter-frame registration needs only a limited number of frames lookahead (i.e. size of the regularization window) and can hence be performed on-line as well as the clustering by orientation, which simply starts a new cluster as soon as the cosine distance is beyond the threshold.

Our incremental compounding adapts the two-step compounding of multiple clusters to an in-place algorithm. Instead of reconstructing a separate volume for each cluster, we use a single volume as accumulation buffer. The reconstructed voxels of each cluster can be incrementally added by rewriting equation 5 to a recurrence scheme, to gain a significantly lower complexity and memory footprint: Given the voxel intensity I_{i-1} and uncertainty U_{i-1} of the previous runs and I_c, U_c of the current run, we define the new intensity I_i and uncertainty U_i as:

$$I_i = \frac{U_{i-1} I_{i-1} + (1 - U_c) I_c}{U_{i-1} + (1 - U_c)}, \quad U_i = U_{i-1} + (1 - U_c) \quad (6)$$

The incremental compounding technique is further accelerated by using an intermediate lookup structure for the backward-warping: Each ultrasound frame is sampled into a lower resolution brick structure using a scanline voxelization technique as an efficient sampling method. Similar to scanline rasterization in Computer Graphics, we compute the coordinates of the four corners of the frame,




Fig. 4. Illustration of the scanline voxelization scheme (for simplicity in 2D): Starting at the bottom-most scanline around v_0 , we compute the left-most and right-most voxel covered by the ultrasound frame. Using the slope of the edges $E_{02} = E_{13}$ and $E_{01} = E_{23}$ we can incrementally compute the start- and end-voxel for the next scanline by simple additions. The increments have to be changed when advancing beyond v_1 resp. v_2 (advancing to middle/top part).

define a scanline axis and sort the corners along the axis. Exploiting the rectangularness of the image, we can compute the increments (i.e. slopes) along the other two axes for one step along the scanline axis. Iterating brick-wise along the scanline axis, we can compute all bricks touched by the frame using simple additions as depicted in Fig. 4. The brick structure can then be used to accelerate the lookup of all ultrasound frames close to a voxel.

5 Evaluation and Results

To evaluate our methods, we used an ACUSON S2000TM ultrasound machine equipped with an Acuson 9L4 linear transducer and Ascension trakSTARTM2 electromagnetic tracking hardware being calibrated as described in [10].

To confirm the physically correct reconstruction of anatomy, we acquired ultrasound sweeps of an abdominal phantom including a tumor target of spherical shape as depicted in Fig. 1. We computed 50 MPRs of arbitrary orientation through the target and compared the maximum diameter with measurements acquired from CT: The reconstructed ultrasound volume yielded an average target diameter of 14.63 ± 0.48 mm compared to 14.5 ± 0.84 mm in CT. Since the target is positioned relatively close to the surface, it can be scanned from different directions and is prone to deformation, hence being a relevant scenario to evaluate our method.

The effects of our inter-frame registration and pressure compensation technique can be observed in Fig. 1 showing the reconstructions of an ultrasound sweep through the abdominal phantom: Due to the probe pressure changes the

MPR through the reference volume (a) shows deformation of the originally round target. Our techniques restore the original shape, seen in (b) showing the same MPR through the volume compounded with pressure compensation.

Figure 3 shows the effect of our clustering technique when reconstructing a twisted ultrasound sweep of human shoulder. Due to the overlapping frames the baseline compounding in (a) shows artifacts because the closest frames for neighboring voxels may be acquired from different angles. The reconstruction in (b) uses our clustering technique to avoid overlapping frames and the occurring artifacts and additionally exploits uncertainty information when fusing the clusters so that unreliable intensities do not bias the final result.

Table 1. NCC and log-scale SNR in the overlapping region after registering the two compounded volumes of two sweeps with perpendicular trajectories of the same anatomy

	Baseline [2]		Our technique	
	NCC	SNR _{dB}	NCC	SNR _{dB}
Phantom / constant pressure	0.90	19.39	0.94	23.16
Phantom / pressure changes	0.81	13.02	0.94	22.47
In-vivo leg / constant changes	0.72	9.21	0.76	11.69
In-vivo leg / pressure changes	0.67	8.53	0.75	11.03

For quantitative evaluation we acquired pairs of overlapping sweeps with perpendicular main trajectory of both phantom and in-vivo data. After compounding the sweeps into separate 3D volumes using our methods, we applied a 3D-3D rigid registration using the tracking data as initialization. Expecting our techniques to yield better matching volumes, we compared their differences in the overlapping region with the baseline method (standard backward-compounding and no pressure compensation as described in [2]). With the average of the two volumes as expected result for a correct reconstruction, we quantify their difference in Normalized Cross-Correlation (NCC) and log-scale Signal to Noise Ratio (SNR_{dB}), for which define the signal as average of the volumes and the noise as RMS of the differences (Table 1). The sweeps with pressure changes show a significant improvement in terms of increase in both NCC and SNR_{dB} when our technique is applied. Furthermore, when comparing constant pressure with pressure changes, our technique shows significantly less drop of the measures. The slight improvements for the sweeps acquired with constant pressure are mainly due to the inter-frame registration correcting for the tracking error. Since the sweeps are acquired with perpendicular trajectories and the volumes therefore show different interpretations of the underlying data, no algorithm yields a perfect match. Furthermore, the in-vivo sweeps are expected to have lower similarity since they show by far less homogeneous anatomy. Figure 5 shows the difference images for the second phantom sweep.




Fig. 5. Illustration of our evaluation method: First two images show the MPRs for each perpendicular sweep, the third one shows the squared difference of their intensities after 3D-3D rigid registration. (a) traditional backward-compounding fails to align the different structures; (b) our technique yields alignment of all structures. The quantitative results are shown in Table 1.

6 Discussion and Conclusion

In this work, we presented a novel orientation-driven approach to allow 3D free-hand ultrasound for a broader range of clinical applications. Typical acquisition sweeps in breast or musculoskeletal (MSK) ultrasound have pressure changes, back- and forth or twisting motion, which are not handled well by current state-of-the-art methods yielding artifacts for regions where the frames overlap. We cluster the ultrasound frames based on orientation and proximity and thereby guarantee that no frames in a cluster overlap. We further use per-pixel uncertainty information when fusing the clusters into the compounded volume, which yields more accurate reconstructions in places where we have information from different acoustic windows, because intensities from uncertain regions do not affect reliable intensities. Our method for probe pressure compensation uses a similar inter-frame registration approach as [6] but also incorporates the orientation of the frames to each other and uses a regularization independent from the tracking calibration quality.

Since the evaluation of our methods shows very good results for the reconstruction of non-homogeneous ultrasound sweeps, the question arises why our two-step compounding approach of first clustering by frame orientation and then fusing based on uncertainty information is superior to a classical one-step approach. We assume that the low signal-to-noise ratio of ultrasound, its high view-dependency and thus very limited consistency in time and movement sets the main challenge when compounding non-homogeneous ultrasound sweeps. Our orientation-driven two-step compounding technique introduces an additional interpolation step and thus compensates better for those highly non-linear effects. By exploiting uncertainty information we ensure that this additional interpolation does not impair the final image quality but even improves the result in regions where we have inconsistent image information from different acoustic windows. Also other applications [11] have shown the benefit of this uncertainty based approach to information fusion.

3D freehand ultrasound has a wide band of applications in both diagnostic and interventional imaging. Our work allows high quality reconstructions also in applications such as breast or MSK where soft and uneven surfaces lead to sweeps of non-homogeneous pressure and non-linear trajectory. Our implementation shows optimizations to stream-line our methods to allow real-time applications where the compounded volume gets updated and refined interactively during the acquisition, providing the clinician with direct feedback. Hence, we believe that our orientation-driven methods will have a significant impact in bringing ultrasound compounding further to clinical applications.

References

1. 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), 991–1009 (2007)
2. Wein, W., Pache, F., Röper, B., Navab, N.: Backward-warping ultrasound reconstruction for improving diagnostic value and registration. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) *MICCAI 2006*. LNCS, vol. 4191, pp. 750–757. Springer, Heidelberg (2006)
3. Rohling, R., Gee, A., Berman, L.: A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis* 3(4), 339–359 (1999)
4. Sanches, J.M., Marques, J.S.: A multiscale algorithm for three-dimensional freehand ultrasound. *Ultrasound in Medicine & Biology* 28(8), 1029–1040 (2002)
5. Klein, T., Hansson, M., Navab, N.: Modeling of multi-view 3d freehand radio frequency ultrasound. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part I*. LNCS, vol. 7510, pp. 422–429. Springer, Heidelberg (2012)
6. Treece, G., Prager, R., Gee, A., Berman, L.: Correction of probe pressure artifacts in freehand 3d ultrasound. *Medical Image Analysis* 6(3), 199–214 (2002), Special Issue on Medical Image Computing and Computer-Assisted Intervention 2001
7. Housden, R.J., Gee, A.H., Treece, G.M., Prager, R.W.: Sensorless reconstruction of unconstrained freehand 3d ultrasound data. *Ultrasound in Medicine & Biology* 33(3), 408–419 (2007)
8. Karamalis, A., Wein, W., Klein, T., Navab, N.: Ultrasound confidence maps using random walks. *Medical Image Analysis* 16(6), 1101–1112 (2012)
9. Shepard, D.: A two-dimensional interpolation function for irregularly-spaced data. In: *Proceedings of the 1968 23rd ACM National Conference*, ACM 1968, pp. 517–524. ACM, New York (1968)
10. Wein, W., Khamene, A.: Image-based method for in-vivo freehand ultrasound calibration. In: *SPIE Medical Imaging 2008*, San Diego (February 2008)
11. Comaniciu, D., Zhou, X.S., Krishnan, S.: Robust real-time myocardial border tracking for echocardiography: an information fusion approach. *IEEE Transactions on Medical Imaging* 23(7), 849–860 (2004)