Backward-Warping Ultrasound Reconstruction for Improving Diagnostic Value and Registration

Wolfgang Wein¹, Fabian Pache¹, Barbara Röper², and Nassir Navab¹

¹ Chair for Computer Aided Medical Procedures (CAMP), TU Munich
Boltzmannstr. 3, 85748 Garching, Germany
{wein, pache, navab}@cs.tum.edu

² Clinic and Policlinic of Radiation Oncology,
Klinikum Rechts der Isar, TU Munich, Germany
barbara.roeper@lrz.tum.de

Abstract. Freehand 3D ultrasound systems acquire sets of B-Mode ultrasound images tagged with position information obtained by a tracking device. For both further processing and clinical use of these ultrasound slice images scattered in space, it is often required to reconstruct them into 3D-rectilinear grid arrays. We propose new efficient methods for this so-called ultrasound spatial compounding using a backward-warping paradigm. They allow to establish 3D-volumes from any scattered freehand ultrasound data with superior quality / speed properties with respect to existing methods. In addition, arbitrary MPR slices can be reconstructed directly from the freehand ultrasound slice set, without the need of an extra volumetric reconstruction step. We qualitatively assess the reconstruction quality and quantitatively compare our compounding method to other algorithms using ultrasound data of the neck and liver. The usefulness of direct MPR reconstruction for multimodal image registration is demonstrated as well.

1 Introduction

As a cost effective and non-invasive method, two-dimensional ultrasound is the most widely used medical imaging modality. As it depicts only two-dimensional planar images from the anatomy, a lot of research has been carried out to transform it to a three-dimensional modality [1]. Nowadays there are a number of 3D-ultrasound systems available, using either a mechanical wobbling setup or native 2D matrix arrays. They establish a fan of ultrasound planes, which in turn can be scan-line converted to yield a cartesian 3D volume. These volumes are however very limited in their size. Freehand ultrasound imaging uses mainly position sensing to record simultaneously the 2D ultrasound images and their position and orientation in space. This allows to establish the spatial relation from arbitrary movement of the ultrasound transducer, potentially covering a much larger area of the human anatomy. In order to precisely know the transformation in space of each image plane based on the measurement of e.g. a magnetic tracking sensor or optical tracking target attached to the ultrasound probe, a

careful calibration is necessary [2]. Using the three-dimensional characteristics, in particular any-plane views, has many medical advantages, including

- independency from examiner and, to some extent, used probe positions
- freedom to display pathological process in any angle, e.g. along and perpendicular to its main axis for visualizing its full extend, or planes that focus on relations to relevant neighbouring normal tissue structures
- possibility to visualize planes parallel to the skin, that cannot directly be derived from diagnostic sweeps in B-mode
- upvalue ultrasonography into a comparable line with other sectional imaging modalities that allow for free choice of plane at acquisition (e.g. MRT) or reconstruction (e.g. multi-slice CT)

To exploit them, a reformatting either into a cartesian volume or a plane arbitrarily located in space is necessary. This process is denoted *Spatial Compounding*, and there are two distinct approaches to it, as already pointed out by [3].

A footprint of each of the B-mode images scattered in space can be created in the initially empty 3D volume. If an additional volume channel is used, an averaging can be achieved where several ultrasound planes intersect the same voxels, otherwise often the maximum is used as final intensity [4]. This so-called forward-warping is computationally efficient, while it has the potential to cause gaps in the reconstruction. It can also be used to directly create Multi-Planar Reconstructions (MPRs) by assuming a constant elevational thickness of each ultrasound image [5]. For this purpose, a polygon depicting the intersection of each ultrasound image with the desired plane is drawn onto the screen with hardware-accelerated texture-mapping.

A backward-warping strategy would traverse the target plane or volume, for each grid point identify the relevant original ultrasound information and accumulate the voxel intensity using e.g. distance-weighted interpolation or other merging schemes. We will present an algorithm implementing this approach efficiently, despite the obviously high computational effort. A simplified approach is taken in [6], where a continuous probe motion without any ultrasound plane intersections is assumed, henceforth each voxel intensity is weighted from the two neighboring ultrasound slices, using the probe trajectory rather than the perpendicular projections.

2 Methods

For every discrete position $\mathbf{x}_i \in V$ in the reconstruction volume V , we need to compute a set of tuples $A_i = \{(d, y)\}$ where d is the distance of an original ultrasound data point to $\mathbf{x}_i$ and y its intensity value.

$$\forall \mathbf{x}_i \in V : A_i = \{(d, y) | d < D; d = \|\mathbf{p} - \mathbf{x}_i\|\} \quad (1)$$

$$\mathbf{p} = H_j(u, v, 0, 1)^T; \quad y = Y_j(u, v) \quad (2)$$

Here, D is the maximum distance at which points should be taken into account for accumulation of the voxel intensity. The homogenous 4x4 transformation

matrix H_j for a particular ultrasound slice j maps a 2D point $(u, v)^T$ of the slice plane into a point $\mathbf{p}$ in 3D. The corresponding ultrasound image intensity is denoted $Y_j(u, v)$. The main effort is now to determine the set A_i for each voxel from the whole ultrasound information available. However, we can restrict ourselves to points $\mathbf{p}$ originating from slices whose perpendicular distance to $\mathbf{x}_i$ is smaller than D :

$$S_i = \{j | d_{i,j} < D; d_{i,j} = (0, 0, 1, 0) H_j^{-1} \mathbf{x}_i\} \quad (3)$$

2.1 Fast Slice Selection

In the following we devise an efficient means to successively compute S_i for all voxels, using the following facts:

- We can traverse the volume in a way such that the distance of two successive voxels is always $\|\mathbf{x}_{i+1} - \mathbf{x}_i\| \leq k$, where k is the maximum voxel spacing. Figure 1 illustrates a simple possible scheme.
- If the distance of ultrasound slice j to $\mathbf{x}_i$ is $d_{i,j}$, then the distance of the same slice to voxel $\mathbf{x}_{i+1}$ can not be smaller than $d_{i,j} - k$.
- Hence this mentioned slice j can only be required for an voxel index $i + \lfloor d_{i,j}/k \rfloor$ or later.

We use a rotation queue with $\lceil k/D_V \rceil$ elements (D_V being the volume diagonal), which is equivalent to the maximum distance of a slice contributing to the reconstruction, from a particular voxel. Each element contains a set of slice indices, at the beginning all slices are in the head element of the queue. For every voxel $\mathbf{x}_i$, all slices in the queue head are removed, their distance d is computed and they are reinserted into the rotation queue corresponding to that distance. If $d < D$, then the slice is added to S_i and considered for accumulation of the voxel intensities. For the next voxel $\mathbf{x}_{i+1}$, the rotation queue head is advanced.

Fig. 1. Volume traversal scheme for advancing only one voxel at a time

This will allow us to implement efficient backward-warping compounding methods to reconstruct three-dimensional volumes of ultrasound information. In order to create online Multi-Planar Reconstructions (MPRs) directly from

the original data, we will just consider a reconstruction volume with a single slice, arbitrarily located in space.

2.2 Intensity Accumulation

For a voxel $\mathbf{x}_i$, all pixels on each slice $\in S_i$ that satisfy $d < D$ are added to the set of distance-intensity tuples A_i defined in equation 1. For a given set $A = \{(d_i, y_i)\}$, the reconstructed voxel intensity can be any weighting or selection function $f(A)$. We considered the following functions in our work:

Inverse Distance Weighting. Originally defined in [7], it assures that the reconstructed intensity approximates the original data values for $d \rightarrow 0$ ($\mu > 1$ being a smoothness parameter):

$$f(A) = \sum_{i=1}^n y_i \frac{d_i^{-\mu}}{\sum_{j=1}^n d_j^{-\mu}} \quad (4)$$

Gaussian Weighting. A 3D Gaussian kernel with size σ can be used to weight the data values as well:

$$f(A) = \frac{\sum_{i=1}^n y_i e^{-d_i^2/\sigma^2}}{\sum_{i=1}^n e^{-d_i^2/\sigma^2}} \quad (5)$$

Nearest Sample. Here, the data value closest to the considered voxel is directly accepted as reconstructed intensity:

$$f(A) = y_i | d_i = \min\{d_i\} \quad (6)$$

Weighted Median. Using the median has the advantage that one of the original intensities is chosen, which would be the centermost one from the sorted intensities $[y_i]$. In order to incorporate the distances, we "stretch" the sorted list with their respective inverse linear weightings ¹

$$\begin{aligned} \forall k \in [1 \dots n] : y_{k+1} &\geq y_k, w_k = 1 - \frac{d_k}{D} \\ f(A) = y_i \left| \sum_{k=1}^{i-1} w_k \right. &\leq \frac{\sum_k w_k}{2} \leq \sum_{k=1}^i w_k \end{aligned} \quad (7)$$

3 Results

3.1 Computation Time

We implemented a forward-warping compounding algorithm for comparison purposes, which averages all ultrasound pixel hits onto reconstruction volume voxels, as described in [4]. The following table compares the impact of volume and slice resolution on both forward and backward compounding.

¹ Linear mapping $[0, D] \rightarrow [1, 0]$, pixel beyond distance D are disregarded and therefore have weight 0.

Fig. 2. Longitudinal MPR (a) from an axial freehand sweep of a right neck created using backward compounding with weighted median. Anatomical details, from superficial to deep structures: skin, subcutaneous tissue, platysma, sternocleidomastoid muscle, lymph nodes (left: normal, center+right: malignant), internal jugular vein with physiological pulsation, marginal section of carotid artery with arteriosclerotic plaque. Enlarged details of images interpolated from compounded volume (b,c,e,f) and online MPR (d).

	forward		backward multiple		backward single	
D_V	$D_S = 256$	$D_S = 454$	$D_S = 256$	$D_S = 454$	$D_S = 256$	$D_S = 454$
16	517s	1687s	226s	401s	119s	295s
134	538s	1915s	712s	942s	442s	510s

D_V denotes the number of million voxels of the reconstruction volume, D_S is the width and height of the input ultrasound slices. The times were taken using a set of 1024 slices, on an AMD64 3200+ with 1GB RAM. While forward compounding is largely unaffected by the number of voxels, the amount of pixels to be processed linearly affect the computation time. Backward compounding reacts to changes in both input and target data, however the increase depending on the size of the slices can be largely eliminated if only a single pixel of each slice is required for each voxel as in a nearest neighbor accumulation.

3.2 Qualitative and Quantitative Comparison

Figure 2(a) shows one reslice approximately orthogonal to the original slices. 200 slices were processed using different accumulation schemes. Mean-Squared-Error

Fig. 3. The MPR Viewport was frozen in place at the spatial location of the original slice (d). The individual contributions of the right and left sequence are (a) and (b). The SMD (Squared Mean Differences) of the reconstruction drops significantly when both sequences are combined in (c).

(MSE) values were calculated according to the leave-one-out strategy for the whole reconstruction volumes, as described in [6]. Each method has its own merits for specific applications. The more homogenous appearance of the gaussian and smoothing kernels are favourable for volume rendering and segmentation. On the other hand, methods that retain original intensities, median and nearest sample produce more diagnostically relevant images than methods recomputing the values. Nearest Neighbor is the fastest but also the most unforgiving on noisy datasources and jittery tracking information. The online MPR reconstruction (fig. 2(d)) provides the sharpest and detailed images, as one data resampling step is skipped.

3.3 Interleaved 3D-Ultrasound Data

Using the online backward-warping MPR reconstruction it is possible to fuse multiple freehand ultrasound sweeps regardless of their relative spatial positions. The samples presented here were created using three sequences. A selected slice of the center sequence was recreated using data from sequences to the right and left, see Figure 3. Reconstruction using backward compounding with weighted median accumulation of one 256×256 slice from sweeps with a total of over 1000 slices takes about 1 second on an AMD64 3200+ using high quality settings.

3.4 Improving Registration

For multimodal registration of freehand ultrasound images with a CT scan, we had developed automatic image-based registration techniques in previous work [8]. There, a set of axial images from a continuous caudo-cranial sweep along the neck is selected and used for registration. Using our compounding methods, arbitrary planes of ultrasound information can be considered in addition. Figure 4 shows two original transversal images from a freehand ultrasound sequence along the jugular vessels, as well as two additional MPR-reconstructed slice images.

Fig. 4. Examples for spatially related ultrasound images from the neck: Original axial slices (a+b) and longitudinal MPRs (c+d) using our reconstruction technique

One could argue that the examiner should rotate the ultrasound probe after the continuous caudo-cranial motion to acquire longitudinal images supporting the registration. This however will introduce significant deformation errors due to tissue compression distributed differently on the patient's skin, as well as motion induced by the blood pulsating through the vessels. For a rigid registration, it is hence preferred to use additional planes derived directly from the data of the original continuous probe motion. This is done using our MPR reconstruction method, which allows in addition to create planes of information parallel to the skin surface within the body, which cannot be acquired by ultrasound itself.

The robustness of automatic rigid CT-Ultrasound registration on such a sequence improved significantly when oblique reconstruction planes were considered for registration. The standard deviation of the Target Registration Error on a lymph node decreased from 3.2mm to 1.5mm for a random displacement study for the sequence depicted in figure 4, when two MPRs were considered in addition to 5 slices from the original sweep (only two of which are shown in the figure for better spatial impression).

4 Conclusion

We have developed new methods for spatial compounding of freehand ultrasound data, using a backward-warping approach which collects the scattered image information for each voxel in an efficient manner. They allow to perform reconstructions with superior quality and smaller computation time compared to forward-projection techniques known from literature, while a choice of smoothness and continuity versus retaining original image characteristics can be made using different accumulation functions. These algorithms can also be used to compute Multi-Planar Reconstructions (MPRs) in real-time from the original data. This further increases the image quality, as an extra interpolation step is saved that would be necessary when rendering MPRs from previously compounded 3D-volumes. As a result, more detailed diagnostic information can be

gathered and visualized not only for the person performing the ultrasound examination, but also for demonstration to colleagues at a later point of time without the presence of the patient necessary. Pathological changes in tissue texture and their relation to anatomical landmarks can be demonstrated in an optimized plane without haste, which gives the possibility to combine the advantages of ultrasonography (high spatial resolution and tissue contrast depending on the frequency of the ultrasound device) with the advantages of sectional imaging in any plane. Furthermore, multimodal image-registration can be improved by adding oblique ultrasound information in addition to original image data. Finally, the online MPR algorithm can yield high-quality real-time visualization of oblique slices for freehand ultrasound systems.

References

1. Fenster, A., Downey, D.B., Cardinal, H.N.: Three-dimensional ultrasound imaging. *Physics in Medicine and Biology* **46**(5) (2001) R67–R99
2. Mercier, L., Langø, T., Lindseth, F., Collins, D.: A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Med. & Biol.* **31** (2005) 449–471
3. Barry, C., Allott, C., John, N., Mellor, P., Arundel, P., Thomson, D., Waterton, J.: Three-dimensional freehand ultrasound: Image reconstruction and volume analysis. *Ultrasound in Med. & Biol.* **23** (1997) 1209–1224
4. Rohling, R., Gee, A., Berman, L.: Automatic registration of 3-D ultrasound images. Technical report, Cambridge University Engineering Department (1997)
5. Prager, R., Gee, A., Berman, L.: Stradx: Real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis* **3** (1999) 129–140
6. Coupe, P., Hellier, P., Azzabou, N., Barillot, C.: 3d freehand ultrasound reconstruction based on probe trajectory. In: *MICCAI Proceedings*. (2005) 597–604
7. Shepard, D.: A two-dimensional interpolation function for irregularly spaced data. In: *Proc. 23 Nat. Conf. ACM*. (1968) 517–524
8. Wein, W., Röper, B., Navab, N.: Automatic registration and fusion of ultrasound with CT for radiotherapy. In: *MICCAI 2005 Proceedings*. *Lecture Notes in Computer Science*, Springer (2005)