

Deriving Anatomical Context from 4D Ultrasound

M. Müller^{1,2}, L. E. S. Helljesen³, R. Prevost¹, I. Viola^{4,5}, K. Nylund^{3,6}, O. H. Gilja^{3,6}, N. Navab² and W. Wein¹

¹ImFusion GmbH, München, Germany

²Chair for Computer Aided Medical Procedures (CAMP), Technische Universität München, Germany

³National Centre for Ultrasound in Gastroenterology, Haukeland University Hospital, Norway

⁴Vienna University of Technology, Austria

⁵Department of Informatics, University of Bergen, Norway

⁶Department of Clinical Medicine, University of Bergen, Norway

Abstract

Real-time three-dimensional (also known as 4D) ultrasound imaging using matrix array probes has the potential to create large-volume information of entire organs such as the liver without external tracking hardware. This information can in turn be placed into the context of a CT or MRI scan of the same patient. However for such an approach many image processing challenges need to be overcome and sources of error addressed, including reconstruction drift, anatomical deformations, varying appearance of anatomy, and imaging artifacts. In this work, we present a fully automatic system including robust image-based ultrasound tracking, a novel learning-based global initialization of the anatomical context, and joint mono- and multi-modal registration. In an evaluation on 4D US sequences and MRI scans of eight volunteers we achieve automatic reconstruction and registration without any user interaction, assess the registration errors based on physician-defined landmarks, and demonstrate real-time tracking of free-breathing sequences.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement—Registration

1. Introduction

With native 3D ultrasound probes more commonly available, methods which extend the long known panoramic imaging approaches into the third dimension are emerging. Such extended field-of-view (FOV) volumetric data can then be superimposed with other pre-operative or diagnostic imaging data. This may be particularly useful for the diagnosis, treatment planning, interventional navigation, and tumor follow-up such as lesions in liver and kidney; however a multitude of other applications would benefit from the widened anatomical context of 3D ultrasound data. Generally speaking, the capabilities of medical ultrasound shall be extended to allow for reproducible, large-volume tomographic imaging, while preserving its unique advantages as interactive real-time imaging modality.

Regarding image-based 3D ultrasound tracking for achieving this, the main challenge lies in compensating registration errors between the overlapping volumes which can lead to significant drift distorting the shape of the final ex-

tended FOV volume. Using multi-modal registration to CT or MRI data which may be available in some clinical application scenarios, remains a difficult problem to solve with automatic algorithms.

2. Related Work

For image-based tracking of 3D ultrasound volumes, a simultaneous optimization of all volumes is suggested in [WWN07] to reduce drift, together with designated multi-variate similarity measures. However, for the scenario here with a large number of volumes streaming in real-time, such an approach is both computationally prohibitive and would not fully address the drift problem as only some of the consecutive volumes overlap. In [ØWU*12] a consecutive tracking technique also taking previously registered volumes into account to reduce drift is presented (we adapt a similar approach in this work).

Additional MRI or CT data can serve as anatomical reference to improve the ultrasound reconstruction, as suggested

in [KWN09] along with preliminary results on phantom data and freehand ultrasound volumes. The required multi-modal registration of ultrasound is still a complex research topic though. Common image-based approaches try to simulate ultrasound from CT during registration to ease comparison [WBK*08, XLZG13]. Such an approach is not directly applicable on MRI because of the missing physical link of the image intensities. Instead [RPMA01] and [WLF*13] have shown that comparing the ultrasound to the intensity and gradient of the MRI simultaneously produces accurate and stable results for the brain.

In this work, we present a novel system for purely image-based abdominal US-MRI fusion based on a robust extended FOV reconstruction, high-fidelity median compounding, a machine learning approach to determine the initial pose, and the LC² metric [WLF*13] for multi-modal registration. Besides, to our knowledge for the first time, we offer a quantitative evaluation of reconstruction drift and registration errors on in-vivo liver data, while previous assessment of related techniques was based only on in-vitro error analysis.

3. Method

An overview of the system is shown in Fig. 1. After the acquisition of a breath-hold sequence, our automatic reconstruction consisting of registration and compounding (Section 3.1) creates the extended FOV volume of the liver. If an MRI scan is available, it is used to place the ultrasound volume into its anatomical context by means of an initialization step (Section 3.2) followed by multi-modal registration (Section 3.3) and optional refinement (Section 3.4). Last but not least, during live acquisition the aforementioned steps may be used to provide the internal positioning with respect to the anatomy from the breath-hold ultrasound volume and optionally the MRI data, by means of image-based real-time tracking as demonstrated in Section 4.5. This could for example guide the clinician to relevant structures such as vascular trees or possible pre-segmented liver pathologies.

3.1. 3D Ultrasound Reconstruction

Instead of a pair-wise image registration on the two most recent frames of ultrasound volumes streaming in real-time, we use a number of m preceding volumes as proposed in [ØWU*12]. We extend this approach by comparing the image data with normalized cross-correlation (NCC) rather than squared intensity differences, which yields higher robustness due to its invariance to brightness and contrast changes. We define each registration update as:

$$\arg \max_{T_k} \sum_{i=1}^m \left(\omega(k, k-i) \sum_{x \in \Omega} \frac{(S_{k-i}(T_{k-1} \cdot x) - \bar{S}_{k-i})(S_k(T_k \cdot x) - \bar{S}_k)}{\sigma_{S_k} \sigma_{S_{k-i}}} \right) \quad (1)$$

Here $S = (S_1, \dots, S_n)$ are the available ultrasound sweep volumes and k the volume frames $\in [2 \dots n]$. \bar{S}_k denotes the

mean intensity value and σ_{S_k} the intensity variance of the ultrasound volume S_k . The weighting $\omega(k, k-i)$ is defined as the relative overlap of the volumes S_k and S_{k-i} , consequently volumes with small overlap contribute less to the final similarity value. The linear transformation matrix $T_k \in [T_1, \dots, T_n]$ describes the relative transformation from the first volume S_1 to S_k (with $T_1 = I_4$). We use 6 degrees-of-freedom (DOF) rigid parameters, since 9 or 12 DOF affine transformations accumulate much stronger drift errors and distort the final volume data. Those transformation parameters are adjusted with a non-linear optimization algorithm based on quadratic approximations [Pow09] for fast local convergence. Implementation on the GPU allows to fully parallelize the inner sum in (1), creating one thread per voxel.

For creating a compounded extended FOV volume from the registered data, each target voxel value is chosen as the median of all original volume intensities at the respective location. As opposed to a simple averaging, this better preserves small vascular structures which are not perfectly registered due to pulsation or other deformations, because the majority of overlapping voxel intensities determine the outcome (see Fig. 2). For each volume slice, all volume intersections are interpolated and sorted in parallel on the GPU, resulting in negligible overall computation time of this step.

3.2. Anatomical Context Initialization

With the extended FOV US volume in place, we need to find a coarse but robust estimate of its geometric transformation to the MRI scan, that will serve as an initialization for a subsequent more elaborate registration algorithm. As in [KZW*12], such an alignment is found using the diaphragm (a thin muscle between the lungs and the liver), which is clearly visible in both modalities (see Fig. 3). To detect it, we learn the local appearance using the popular Random Forests framework [CS13]. One forest is learnt for each modality; in both cases, the features used are the intensity, the gradients and the laplacian of the pixel and its 7³ neighborhood after smoothing at three different scales. An example of such response maps are also shown in Fig. 3. The registration is then found as the rigid transformation that minimizes the sum of the Euclidean distances of all the detected points in the MR image (*i.e.* voxels with a response greater than 0.5) to the US probability map, using a global optimizer [KA06]. The asymmetry in this definition comes from the fact that diaphragm prediction is usually more accurate in MR than in US images, in which some other structures like the heart might have a similar appearance. The whole process takes a few seconds on a standard desktop computer. Compared to previous work, this represents a straightforward and unified approach for the two modalities, while [KZW*12] had to solve both a graph-cut problem with heuristic features for the diaphragm detection in the US image and a model-based segmentation (which also required an initialization) for the MRI image. Besides, the extended FOV volume available in our workflow allows

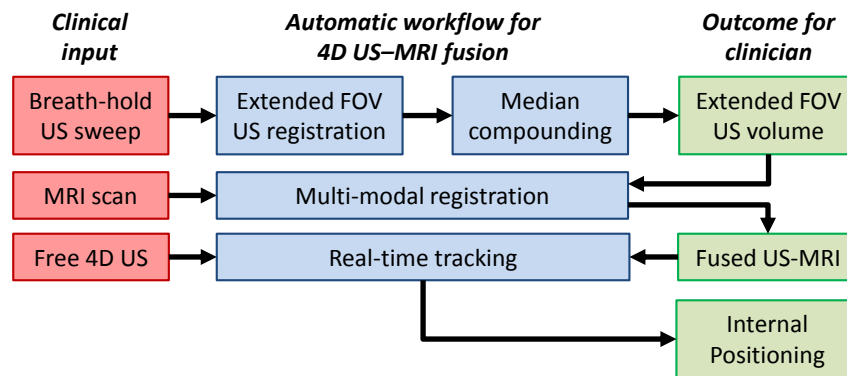


Figure 1: Overview of the proposed system.




Figure 2: Mean (left) versus median compounding on patient 2.

for much more reliable and less ambiguous surface matching since a significantly wider portion of the diaphragm can be used.

3.3. Multi-modal Registration

Now that the diaphragm is lined up, the US-MRI registration is refined using a method based on the LC^2 similarity metric [WLF*13]. We use it with a three-dimensional neighborhood around every voxel, such that a least-squares fit of the linear intensity mapping between US, MRI intensity and MRI gradient magnitude is established for every region within the volumes. This allows for a robust assessment of anatomical structure alignment despite strong modality-specific differences, because the LC^2 formulation is locally invariant with regard to whether ultrasound intensities correspond to structures in MRI (as most vessels), or to interfaces/gradients between them (as other particular hepatic vasculature with echogenic borders or the diaphragm). First, the rigid pose parameters are optimized using a global stochastic optimization strategy [KA06] within bounds centered around the initialization. This is followed by local optimization [Pow09] of a parametric cubic spline based deformable model. Alternatively, a refinement using a global

affine transformation, i.e. containing non-uniform scaling and shearing parameters, may be used as well.

3.4. Joint Registration Refinement

With the MRI in place, it can be used as an anatomic reference in order to try to reduce any remaining drift caused by the reconstruction in Section 3.1. This is achieved by means of multi-modal registration of all original ultrasound volumes to MRI, either by optimizing all transformations, or the first and last, interpolating between them to correct for drift errors. Now the sum of multi-modal LC^2 similarity to MRI and mono-modal NCC similarity to the neighboring volumes is used as optimization cost.

If the sweep was recorded in a continuous motion and does not contain rapid changes in direction, we can assume that the overall shape of the sweep was reconstructed correctly (confirmed by the results of [ØWU*12]), i.e. the anatomical structures are not completely distorted. Otherwise the global registration would most likely have failed anyway because the structures did not match with the MRI. From this, we can conclude that the relative transformation between volumes should only have minor errors and that the




Figure 3: Diaphragm detection in both MR (top) and US (bottom) images using random forests. From left to right: 2D slices of the original image, corresponding response map, 3D volume rendering of the detection.




Figure 4: Illustration of the transformation interpolation scheme used in the refinement step.

largest absolute error is accumulated in the last frame. Therefore the drift error may be compensated by merely changing the transformation of the last frame with the MRI as reference and distribute the transformation over the whole sweep. However, because we can not assume that the global registration aligned the first frame perfectly and just the last frames are misaligned due to drift, our proposed refinement transforms the first and last frame and distribute this motion over the sweep as illustrated in Fig. 4. This can be imagined as moving both ends of a spring independently. In order to distribute the motion uniformly across all frames, it needs to be split up into smaller increments and applied successively, using one out of a variety of known formulas for interpolating transformation matrices. Here, we use a method proposed in [WL13] for correcting drift on motion estimated from successive cone-beam CT X-Ray imaging.

After distributing the motion to all frames, the LC^2 similarity between the MRI and the new positioned frames is accumulated and the resulting value serves as an overall similarity for optimization. Since there is usually a significant overlap between consecutive frames it is also possible to skip frames for the computation of the similarity (not for the transformation distribution though) and therefore speed-

ing up the overall computation while yielding similar results. Utilizing this factor makes this refinement step computationally efficient since the parameter space only consists of two linear transformations.

4. Evaluation

The abdomen of eight volunteers was scanned with MRI T1 Dixon sequences on a Siemens Magnetom Avanto 1.5T scanner, resulting in in-plane resolution of $1.23 \times 1.23mm$ and out-of-plane spacing of $3mm$. 4D ultrasound was then acquired in all subjects using a GE Vivid E9 machine with a 4V matrix probe, including both breath-hold sweeps across as much of the liver as possible, and free-breathing sequences with steady probe. The cartesian 4D US data comprising 30 – 60 volumes per recording had a resolution of $0.64mm$ in axial direction, and $0.93mm$ in lateral and elevational direction, respectively. Ultrasound reconstruction, median compounding, anatomical context initialization and registration to MRI were computed as described above. The diaphragm classifiers have been learnt using a leave-one-out strategy for the US images (due to the limited number of images available) and using a completely different set of patients for the MR images. For evaluating the overall registration errors, a physician selected 5 – 7 anatomical landmarks in both the compounded US and MRI volumes. These included the bifurcation of inferior vena cava and the right, middle and left hepatic vein, portal vein, the neck of the gallbladder, the upper pole or hilum of the kidney, and further vessel bifurcations which could reliably be located in 3D space. Due to the coarse MRI spacing of $3mm$ and blurred appearance of major bifurcations, the landmark localization error is estimated as 2 – $7mm$.




Figure 5: Registration results on patient 8. Top row: Axial slice of MRI (left) and extended FOV US after rigid registration. Middle row: Color overlay of axial and coronal slice after rigid registration. Bottom row: Overlay after deformable registration, please note the precise superimposition of liver vasculature.

Patient	1	2	3	4	5	6	7	8
Rigid	4.8	8.2	11.6	25.1	17.0	8.8	12.3	7.4
Affine	4.7	7.8	11.5	24.0	15.1	9.0	11.6	7.3
Deformable	5.3	6.2	7.8	12.1	22.4	16.1	10.4	6.5
Landmarks	4.5	5.5	3.1	9.7	9.3	5.2	4.0	5.1

Table 1: Landmark distance errors in mm after multi-modal registration.

4.1. Registration Accuracy

Our automatic registration pipeline succeeds without user intervention in all cases. Tab. 1 lists the average landmark distance after rigid, affine and deformable registration on all subjects. In addition, the error after rigid registration of the landmarks themselves is shown, asserting that strong defor-

mations between the modalities exist. Fig. 5 shows the results on Patient 8. After deformable registration, the liver vasculature lines up significantly better. Nevertheless the errors do not decrease as much, because it is not possible to reliably place landmarks along the smaller vessels in the liver gland, which are mostly affected by the non-linear matching.

m	Patient 1		Patient 2		Patient 3		Patient 4		Patient 5		Patient 6		Patient 7		Patient 8	
	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr
1	3,19	1,17	10,22	8,24	5,35	2,94	4,14	0,57	1,65	2,84	7,38	3,95	3,80	3,53	4,06	1,96
2	1,16	0,39	3,63	2,12	1,16	2,44	1,06	0,71	1,35	2,23	4,21	1,37	2,62	1,37	4,18	2,11
3	3,85	1,34	6,78	9,56	3,03	5,41	3,00	1,91	1,15	1,35	4,26	1,97	3,47	1,75	4,81	3,28

Table 2: Evaluation of drift errors with respect to number of volumes m used in reconstruction. Translation errors Δt are in mm and rotation errors Δr in degrees. The best results for every patient are highlighted in bold font.

	Patient							
	1	2	3	4	5	6	7	8
Recon.	23s	25s	37s	11s	17s	67s	40s	57s
Init.	17s	19s	15s	21s	15s	20s	23s	15s
Reg.	34s	20s	25s	25s	20s	26s	37s	45s
Total	74s	64s	77s	57s	52s	113s	100s	117s

Table 4: Computation time of the pipeline steps.

	Patient							
	1	2	3	4	5	6	7	8
Affine	2.5s	2.7s	2.5s	2.8s	2.3s	3.2s	3.0s	2.1s
Interp.	83s	34s	123s	86s	57s	56s	36s	89s
Simult.				>1h				
Deform.	120s	98s	92s	150s	215s	164s	113s	160s

Table 5: Computation times of the refinement.

4.2. Reconstruction Drift

In order to assess possible drift accumulation in the ultrasound reconstruction, we have compared the results by using both expansion directions, as well as a varying number of frames m used in the multi-volume approach (Section 3.1). After the forward reconstruction, we register all volumes backward and compare relative translation and rotation of the first frame. The results are depicted in Tab. 2. Except on Patient 8, using $m = 2$ frames instead of one reduces the reconstruction drift significantly. More frames not only increase the computational cost, but also yield higher drift again, probably because the averaging of non-linear motion within successive frames suppresses the optimal expansion of the rigid probe motion.

4.3. Joint Registration

As a possible means to overcome drift problems, we have applied the joint mono- and multi-modal registration using the following transformation configurations: 1) pose of first and last volume with transformation interpolation of the intermediate volumes, 2) rigid pose of all volumes. The average pose deviations in translation and rotation are depicted in Tab 3. Given that we have limited knowledge (due to the high landmark localization error) regarding how much overall non-linear motion is present between US and MRI (due to probe pressure deformations, different patient position etc.), these values represent an upper worst-case bound regarding the drift we may expect in our image-based tracking approach. Fig. 6 illustrates that the alignment does seem to improve slightly at key anatomical locations.

4.4. Performance

The overall computational time of the algorithm can be split into the distinct steps of ultrasound reconstruction, initialization, multi-modal registration and refinement, with the time for the refinement depending on the used scheme. Tab. 4 show the benchmarks for the overall pipeline without the refinement step, and Tab. 5 the values for different refinement steps individually. The latter table also contains the computation time for affine and deformable MRI-US registration, since this can be categorized as a refinement step and is always executed after rigid registration. Our test system is equipped with an Intel Core i7-3770 processor and a Nvidia GTX Titan graphics board. The ultrasound reconstruction time heavily depends on the number of frames in the sweep (which ranges from 26-58) and the resolution of the volume. The global registration and the context initialization depend mainly on the overall extent of the sweep because the compounding volumes will be larger. The first-last interpolation method was computed using the similarity of ≈ 10 volumes, offering a reasonable trade-off between accuracy and performance. Due to the very high computational cost, a simultaneous registration of the complete sweep is currently not feasible for clinical applications. Since the simultaneous method yields only minor improved results compared to the first-last interpolation method for continuous linear sweeps, it is only advisable for sweeps where the other refinements fail, for example, due to lack of overlap.

4.5. Tracking of Arbitrary Sequences

With reconstruction and multi-modal registration having established the link between large-FOV US and MRI, we can track real-time 3D US volumes, including free breathing sequences. To that end, we use a straightforward registration with an affine transformation model (12 DOF) and the NCC similarity measure. This allows for ≈ 500 evaluations per second using the full US volume size, hence updating the tracking several times per second. In sequences from every subject, the superimposition of the diaphragm and major vessels are properly kept. At the beginning we expect the sonographer to idle the probe for a few seconds with the patient in the same respiration state as in the US compound volume. During this time, the initialization of the first rigid pose is computed using global optimization with bounding conditions derived from the extent of the compounded volume.




Figure 6: Coronal slice of the Patient 8 dataset. Left: Reconstruction without correction, Middle: MRI volume, Right: Reconstruction with "first-last interpolation" correction. The tips of the arrows are located on the same world space coordinates.

mode	Patient 1		Patient 2		Patient 3		Patient 4		Patient 5		Patient 6		Patient 7		Patient 8	
	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr	Δt	Δr
1)	8,58	7,26	4,09	2,78	3,43	2,28	28,8	9,1	9,75	6,66	2,37	1,94	3,28	1,86	5,01	3,2
2)	2,59	8,76	7,6	5,59	15,56	31,16	0,5	3	0,29	3,36	2,55	12,79	2,6	12,8	3,04	16,7

Table 3: Translation distances Δt are in mm and rotation Δr in degrees.

Since both a quantitative evaluation as well as visual assessment in a non-animated print medium of such live data is difficult, we offer a different test of the consistency of the tracking, which is depicted in Fig. 7: Two different ultrasound acquisitions on the same patient are used; the first, shown in Fig. 7(left), is reconstructed using the pipeline introduced in Section 3. We then apply the real-time tracking on the second acquisition using the first sweep as reference. Even though the latter does not use the multi-modal MRI data, the superimposition shown in Fig. 7(middle) is visually convincing. The right image in Fig. 7 depicts a color overlay of both ultrasound sweeps, which confirms that the resulting geometry is approximately identical.

5. Conclusion

We have established a complete pipeline for processing real-time 3D ultrasound to allow for whole organ imaging and multi-modal fusion, without any required user interaction. Evaluation using physician-defined point correspondences reveal inconsistencies on the scale of $\approx 1\text{cm}$, however the localization error is estimated within the same order of magnitude. The entire computation time comprises a few minutes, depending on the actual method chosen for fine-tuning of the multi-modal registration. Besides, we have demonstrated that real-time tracking of free-breathing sequences is possible as well.

We believe that the presented methods pave the way towards fully automatic multi-modal integration of 3D ultrasound with other image data, without the need to use external tracking hardware. This should further strengthen the emerging role of ultrasound as quasi-tomographic imaging modality, paired with its flexibility allowing for intra-operative use.

The clinical viability of the proposed methods could be

better assessed with more accurate expert-defined measurements. Advanced high-resolution MRI sequences may be used in future studies to address this problem, as well as different approaches of generating ground truth than just landmarks, such as the definition of organ contours or unique hepatic vessel structure measurements. While a major advantage of the proposed methods is the general applicability independent of the examined organ (the learning-based initialization may be applied to any other structures rather than the diaphragm), better results will always be obtained with organ-specific tweaks. In particular, the deformable registration could certainly be improved by considering biomechanical liver models. The real-time tracking can be extended to allow for on-the-fly creation of a 4D respiration model. Last but not least, the advantages for actual clinical applications such as accurate liver lesion diagnosis, treatment and follow-up, have to be proven in designated clinical studies.

Acknowledgments

We would like to express our appreciation to the MedViz Network Bergen (medviz.uib.no) and GE Vingmed Ultrasound (www.gehealthcare.com) for providing the MRI and US datasets from the study on volunteering participants. The involvement of Ivan Viola was supported by the Vienna Science and Technology Fund (WWTF) through project VRG11-010 and by EC Marie Curie Career Integration Grant through project PCIG13-GA-2013-618680.

References

- [CS13] CRIMINISI A., SHOTTON J.: *Decision forests for computer vision and medical image analysis*. Springer, 2013. 2
- [KA06] KAELO P., ALI M.: Some variants of the controlled ran-




Figure 7: Comparison of sweep 1 reconstructed with MRI as in Section 3 (left) and sweep 2 reconstructed using only another previously compounded sweep (middle), difference between the two (right, 1 in red, 2 in green).

dom search algorithm for global optimization. *Journal of optimization theory and applications* 130, 2 (2006), 253–264. 2, 3

- [KWN09] KUTTER O., WEIN W., NAVAB N.: Multi-Modal Registration Based Ultrasound Mosaicing. *Medical Image Computing and Computer-Assisted Intervention 2009* (2009), 763–770. 2
- [KZW*12] KADOURY S., ZAGORCHEV L., WOOD B. J., VENKATESAN A., WEESE J., JAGO J., KRUECKER J.: A model-based registration approach of preoperative MRI with 3D ultrasound of the liver for interventional guidance procedures. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on* (2012), IEEE, pp. 952–955. 2
- [ØWU*12] ØYE O., WEIN W., ULVANG D., MATRE K., VIOLA I.: Real Time Image-Based Tracking of 4D Ultrasound Data. *Medical Image Computing and Computer-Assisted Intervention 2012* (2012), 447–454. 1, 2, 3
- [Pow09] POWELL M. J.: The BOBYQA Algorithm for Bound Constrained Optimization without Derivatives. *Cambridge Report NA2009/06, University of Cambridge* (2009). 2, 3
- [RPMA01] ROCHE A., PENNEC X., MALANDAIN G., AYACHE N.: Rigid Registration of 3-D Ultrasound with MR Images: A New Approach Combining Intensity and Gradient Information. *Medical Imaging, IEEE Transactions on* 20, 10 (2001), 1038–1049. 2
- [WBK*08] WEIN W., BRUNKE S., KHAMENE A., CALLSTROM M., NAVAB N.: Automatic CT-Ultrasound Registration for Diagnostic Imaging and Image-Guided Intervention. *Medical Image Analysis* 12, 5 (2008), 577. 2
- [WL13] WEIN W., LADIKOS A.: Towards general motion recovery in cone-beam computed tomography. In *Fully3D 2013 Proceedings* (June 2013). 4
- [WLF*13] WEIN W., LADIKOS A., FUERST B., SHAH A., SHARMA K., NAVAB N.: Global registration of ultrasound to MRI using the LC2 metric for enabling neurosurgical guidance. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*. Springer, 2013, pp. 34–41. 2, 3
- [WWN07] WACHINGER C., WEIN W., NAVAB N.: Three-dimensional ultrasound mosaicing. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007*. Springer, 2007, pp. 327–335. 1
- [XLZG13] XU L., LIU J., ZHAN W., GU L.: A novel algorithm for CT-ultrasound registration. In *Point-of-Care Healthcare Technologies (PHT), 2013 IEEE* (2013), IEEE, pp. 101–104. 2