



Intra-operative Ultrasound to MRI Fusion with a Public Multimodal Discrete Registration Tool

Mattias P. Heinrich (✉)

Institute of Medical Informatics, University of Lübeck, Lübeck, Germany
heinrich@imi.uni-luebeck.de
<http://www.mpheinrich.de/>

Abstract. We present accurate results for multi-modal fusion of intra-operative 3D ultrasound and magnetic resonance imaging (MRI) using the publicly available and robust discrete registration approach *deeds*. After pre-processing the scans to have isotropic voxel sizes of 0.5 mm and a common coordinate system, we run both linear and deformable registration using the self-similarity context metric. We use default parameters that have previously been applied for multi-atlas fusion demonstrating the generalisation of the approach. Transformed landmark locations are obtained by either directly applying the nonlinear warp or fitting a rigid transform with six parameters. The two approaches yield average target registration errors of 1.88 mm and 1.67 mm respectively on the 22 training scans of the CuRIOUS challenge. Optimising the regularisation weight can further improve this to 1.62 mm (within 0.5 mm of the theoretical lower bound). Our findings demonstrate that in contrast to classification and segmentation tasks, multimodal registration can be appropriately handled without designing domain-specific algorithms and without any expert supervision.

1 Introduction and Related Work

Fusion of multimodal medical data is one of the most important application of image registration. In radiotherapy the registration of a computed tomography (CT) scan for dose-planning and an MRI for tumour and organs-at-risk delineation have to be aligned in a nonlinear fashion. In ultrasound guided neurosurgery an intra-operative 3D ultrasound (3DUS) has to be registered to a pre-treatment MRI to guide the surgeon during tumour resection. The correction of challenge on brainshift in intra-operative ultrasound (CuRIOUS) addresses the latter and provides a large training dataset of 22 clinical multimodal cases of 3T MRI T1w and T2 FLAIR scans as well as 3DUS after craniotomy but before dura opening with expert-labeled homologous anatomical landmarks as presented and described in detail in [1]. The main challenges that were discussed in previous

work revolve around a suitable way to define modality-invariant similarity metrics and fast and robust algorithm for finding the transform that optimises this measure. Popular choices of metrics for MRI-US registration include gradient-based correlation metrics [2], self-similarity descriptors [3] or advanced mutual information variants [4]. Secondly, a suitable optimisation framework has to be adapted to enable optimal performance on the given dataset. In this submission, we argue that excellent results can be achieved by employing off-the-shelf and publicly available algorithms that have been optimised for general medical image registration tasks, such as atlas-based segmentation propagation of abdominal CT, without any further domain specific adaption. In the next section, we will describe the employed method that is based on self-similarity context (SSC) descriptors [3] and the discrete optimisation framework *deeds* [5] that performed best in two MICCAI segmentation challenges Beyond the Cranial Vault (BVC) in 2015 [6] and Multimodal Whole-Heart Segmentation (MM-WHS) in 2017 [7] with the exact same default parameters and a subsequent label fusion step.

2 Method

Quantised **self-similarity context** descriptors (SSC) [3] are used to define a similarity metric. Instead of relying on direct intensity comparisons across scans, SSC aims to extract modality-invariant neighbourhood representations separately within each scan based on local self-similarities (normalised patch-distances). It naturally deals well with multi-modal alignment problems, enables contrast invariance, which is particularly beneficial for MRI scans, and focuses the alignment on image edges of the ultrasound.

The **dense displacement sampling** registration short *deeds* [5] is a discrete optimisation algorithm that aims to avoid local minima in the cost function. It is therefore in particular suitable for challenging image appearance often seen in intra-operative ultrasound. A dense displacement sampling covers a large range of potential displacements (capture range) and the combinatorial optimisation based on dynamic programming ensures plausible first-order B-spline transformations without unrealistic deformations on a specified control-point grid. A diffusion regularisation is used between edges that connect neighbouring displacement nodes and this graph is simplified to contain no loops (a minimum-spanning-tree) to simplify the optimisation. A symmetry constraint on the non-linear transform further increases the smoothness of deformations.

3 Implementation

We used *c3d*, which is a general purpose medical image processing command-line tool and can be found at <http://www.itksnap.org/pmwiki/pmwiki.php?n=Downloads.C3D> to resample the provided nifti files of 3DUS and T2 FLAIR

into a common reference frame and to isotropic voxel sizes of 0.5 mm^3 . The command used for the MRI FLAIR and 3DUS data respectively therefore was:

```
c3d *folder*/Case1-US-before.nii.gz *folder*/Case1-FLAIR.nii.gz
  -reslice-identity -resample-mm 0.5x0.5x0.5mm -o Case1-MRI_in_US.nii.gz
c3d *folder*/Case1-US-before.nii.gz -resample-mm 0.5x0.5x0.5mm -o Case1-US.nii.gz
```

We then used both linear and deformable parts of the *deeds* framework as downloaded from <https://github.com/mattiaspaul/deedsBCV/> with default settings. These include an linear pre-registration that performs a block-matching on four scale levels and estimates a rigid transformation using:

```
linearBCV -F Case1-US.nii.gz -M Case1-MRI_in_US.nii.gz -R 1 -O affine1
```

In order to be able to apply the estimated linear and nonlinear transformations to manual landmark positions, the algorithm requires segmentation masks. In our case, these will represent landmarks as 3D spheres. After generating two text files with a custom python implementation¹. The landmark segmentations can be easily generated using *c3d* as follows:

```
python landmarks_split_txt.py --inputtag *folder*/Case1-MRI-beforeUS.tag --savetxt Case1_lm
c3d Case1-MRI_in_US.nii.gz -scale 0 -landmarks-to-spheres Case1_lm_mri.txt 1
  -o Case1-MRI-landmarks.nii.gz
```

The transformation matrix is fed into the deformable part of *deeds* using the following (default) parameters: number of displacement steps $l_{\max} = [8, 7, 6, 5, 4]$, quantisation/stride $q = [5, 4, 3, 2, 1]$, and B-spline grid spacings of $[8, 7, 6, 5, 4]$ voxels. A default weighting of $\alpha = 1.6$ between the SSC-similarity and the diffusion regularisation in *deeds* was used. An example command to run a registration is as follows:

```
deedsBCV -F Case1-MRI_in_US.nii.gz -M Case1-US.nii.gz -O Case1-deeds
  -S Case1-US-landmarks.nii.gz -A affine1_matrix.txt
```

The computation times are approx. 5s for linear alignment and 20s for deformable registration on a mobile dual-core CPU based on the efficient OpenMP implementation.

After applying the combined transformations to the 3D landmark spheres, their spatial (voxel) coordinates are extracted by calculating the centre of mass in python using² and the following command, which stores them in a text file and can directly calculate the mTRE when provided with the target landmarks:

```
python landmarks_centre_mass.py --inputnii Case1-MRI-landmarks.nii.gz
  --movingnii Case1-deeds_deformed_seg.nii.gz --savetxt Case1-results
```

In the next sections the results are presented both visually and numerically and their implications are discussed.

4 Results and Discussion

All experiments were run with same settings on the 22 training scans of the challenge and evaluated using all manual landmarks as provided by the organisers. The algorithms are fully automatic and require no manual initialisation.

¹ <https://gist.github.com/mattiaspaul/56a49fa792ef6f143e56699a06067712>.

² <https://gist.github.com/mattiaspaul/f4183f525b1c6c65e71ad23298d6436e>.

Table 1. Numerical results of accuracy of multimodal registration evaluated with manual landmarks in mm using the *deeds* algorithm in three different settings. First, only the linear part is considered, which yields an mTRE of 1.88 ± 0.53 mm. Second, a nonlinear transform is estimated in addition, yielding a slightly higher error of 1.92 ± 0.60 mm. But when finally fitting another rigid transform to the nonlinear result, the best mTRE of 1.67 ± 0.54 mm is reached.

Case	#1	#2	#3	#4	#5	#6	#7	#8	#12	#13	#14	stddev
Before	1.86	5.75	9.63	2.98	12.20	3.34	1.88	2.65	19.76	4.71	3.03	4.29
Linear	1.88	2.38	1.29	1.31	1.87	1.86	1.58	2.66	1.43	3.47	1.33	0.53
Nonlinear	2.60	2.58	2.29	1.35	2.12	2.30	1.63	2.95	1.21	1.70	1.91	0.60
NL+fit	1.45	2.21	2.00	1.20	1.64	1.89	1.46	3.22	1.14	1.18	1.18	0.54
Case	#15	#16	#17	#18	#19	#21	#23	#24	#25	#26	#27	avg
Before	3.37	3.41	6.41	3.66	3.16	4.46	7.05	1.13	10.10	2.93	5.86	5.42
Linear	2.32	1.41	1.78	1.23	2.12	1.90	1.59	1.57	3.21	1.60	1.58	1.88
Nonlinear	1.97	1.73	2.20	1.27	2.29	1.50	1.33	2.54	1.29	1.23	2.30	1.92
NL+fit	1.89	1.50	1.78	1.03	2.05	1.30	1.21	2.69	1.45	1.31	1.86	1.67

We confirmed that the original error (before registration) was approx. 5.4 mm as mentioned in [1]. The numerical results are presented in Table 1 and using distribution plots in Fig. 2. A clear advantage over the initial error can be seen from 5.42 mm to 1.88 mm when using a linear transform only. We were also interested in exploring whether the nonlinear part of the registration may provide a better alignment despite the fact that the ultrasound images were acquired




Fig. 1. Visual example of US-MRI registration for #3 of the training dataset. The top row shows a colour overlay (US in jet) on top of the original MRI. The bottom row demonstrates a clearly improved alignment when applying the automatically estimated linear transform to the MRI scan.

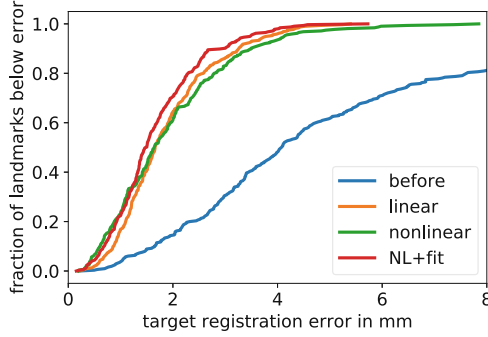


Fig. 2. Cumulative distribution of landmark errors sorted in ascending order. All variants of the automated multimodal discrete registration decrease the landmark error and improve image alignment.

before opening the dura. This is not directly the case as the result slightly deteriorate to 1.92 mm on average. However, when fitting again a rigid transform to the nonlinearly displaced landmark correspondences a mTRE of 1.67 mm. The fitting has been carried out using the technique described in [8]. This indicates that the more flexible deformable registration can improve the match of certain landmarks, but is also less robust in areas of limited contrast. Therefore, the following restriction to a rigid transform, which reduces the influence of outliers, improves the overall outcome.

We further noted that aligning MRI to ultrasound is slightly more accurate than in reverse order. Since, the nonlinear part of *deeds* is already symmetric this discrepancy could be alleviated by using the approach of [9]. Furthermore, the regularisation parameter could be further optimised from $\alpha = 1.6$ to $\alpha = 0.4$ yielding a modest improvement to 1.62 mm mTRE. A visual example of the registration and multi-modal fusion outcome is shown in Fig. 1.

5 Conclusion and Outlook

In summary, we have demonstrated that the general purpose, publicly available, discrete registration toolbox *deeds* provides excellent accuracies of 1.62 mm for a challenging ultrasound to MRI brain registration. The method relies on no training data, but potentially the widely applicable self-similarity descriptors could be replaced by a learning-based approach that relies on known correspondences in training cf. [10] or [11]. However, the impact will probably be more pronounced when considering scans with more brain-shift.

A further interesting research direction would be to only learn the spatial layout used for self-similarity distance computations by means of deformable convolutions. These have been successfully applied to registration and segmentation tasks with few labelled datasets [12]. Moreover, the computation time of the algorithm could be drastically reduced (to subsecond runtimes) by performing the similarity and regularisation calculations on a GPU, which we have

already demonstrated for parts of the algorithm in [13] and we intend to complete this for the whole algorithm in the near future.

Acknowledgements. We would like to thank the CuRIOUS 2018 organisers for providing this new multimodal dataset to the public.

References

1. Xiao, Y., Fortin, M., Unsgård, G., Rivaz, H., Reinertsen, I.: Retrospective evaluation of cerebral tumors (RESECT): a clinical database of pre-operative MRI and intra-operative ultrasound in low-grade glioma surgeries. *Med. Phys.* **44**(7), 3875–3882 (2017)
2. Fuerst, B., Wein, W., Müller, M., Navab, N.: Automatic ultrasound-MRI registration for neurosurgery using the 2D and 3D LC2 metric. *Med. Image Anal.* **18**(8), 1312–1319 (2014)
3. Heinrich, M.P., Jenkinson, M., Papież, B.W., Brady, S.M., Schnabel, J.A.: Towards realtime multimodal fusion for image-guided interventions using self-similarities. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013. LNCS, vol. 8149, pp. 187–194. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40811-3_24
4. Rivaz, H., Karimaghloo, Z., Collins, D.L.: Self-similarity weighted mutual information: a new nonrigid image registration metric. *Med. Image Anal.* **18**(2), 343–358 (2014)
5. Heinrich, M., Jenkinson, M., Brady, M., Schnabel, J.: MRF-based deformable registration and ventilation estimation of lung CT. *IEEE Trans. Med. Imag.* **32**(7), 1239–1248 (2013)
6. Xu, Z., et al.: Evaluation of six registration methods for the human abdomen on clinically acquired CT. *IEEE Trans. Biomed. Eng.* **63**(8), 1563–1572 (2016)
7. Heinrich, M.P., Oster, J.: MRI whole heart segmentation using discrete nonlinear registration and fast non-local fusion. In: Pop, M. (ed.) STACOM 2017. LNCS, vol. 10663, pp. 233–241. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75541-0_25
8. Arun, K.S., Huang, T.S., Blostein, S.D.: Least-squares fitting of two 3-D point sets. *IEEE Trans. Pattern Anal. Mach. Intell.* **5**, 698–700 (1987)
9. Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. *J. Med. Imaging* **1**(2), 024003 (2014)
10. Masci, J., Bronstein, M.M., Bronstein, A.M., Schmidhuber, J.: Multimodal similarity-preserving hashing. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(4), 824–830 (2014)
11. Blendowski, M., Heinrich, M.P.: 3D-CNNs for deep binary descriptor learning in medical volume data. In: Maier, A., Deserno, T., Handels, H., Maier-Hein, K., Palm, C., Tolxdorff, T. (eds.) *Bildverarbeitung für die Medizin 2018*. I, pp. 23–28. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-662-56537-7_19
12. Heinrich, M.P., Oktay, O., Bouteldja, N.: OBELISK - one kernel to solve nearly everything: unified 3D binary convolutions for image analysis. In: *Medical Imaging with Deep Learning* (2018)
13. Ha, I.Y., Wilms, M., Handels, H., Heinrich, M.: Model-based sparse-to-dense image registration for realtime respiratory motion estimation in image-guided interventions. *IEEE Trans. Biomed. Eng.* (2018). <https://doi.org/10.1109/TBME.2018.2837387>