

Multi-modal Multi-Atlas Segmentation using Discrete Optimisation and Self-Similarities

Mattias P. Heinrich Oskar Maier Heinz Handels
heinrich@imi.uni-luebeck.de

Institute of Medical Informatics, University of Lübeck, Germany

Abstract

This work presents the application of a discrete medical image registration framework to multi-organ segmentation in different modalities. The algorithm works completely automatically and does not have to be tuned specifically for different datasets. A robust similarity measure, using the local self-similarity context (SSC), is employed and shown to outperform other commonly used metrics. Both affine and deformable registration are driven by a dense displacement sampling (deeds) strategy. The smoothness of displacements is enforced by inference on a Markov random field (MRF), using a tree approximation for computational efficiency. Consensus segmentations for unseen test images of the VISCERAL Anatomy 3 data are found by majority voting.

1 Introduction

Organ segmentations are an important processing step in medical image analysis, e.g. for image-guided interventions, radiotherapy, or improved radiological diagnostics. General solutions are preferable over organ specific models for large scale image processing. Machine learning approaches, in particular the popular random decision forests (RDF), have been recently used for multi-organ localisation [CSB09] and segmentation [GPKC12], yet for more challenging modalities (e.g. structural MRI) they have had limited success. This is partly due to the inhomogeneous intensity variations within and across MR scans. Registration-based multi-atlas segmentation can provide more robustness by using contrast-invariant similarity measures to guide the alignment of atlas to patient data. Here, we propose to employ a discrete registration model, which can capture large deformations to accurately segment volumes with large differences in patient anatomy and geometry. Combined with a robust multi-modal similarity metric (self-similarity context) it can be applied to registering both CT and MRI scans reliably. The method is briefly reviewed in the next section

Copyright © by the paper's authors. Copying permitted only for private and academic purposes.

In: O. Goksel (ed.): Proceedings of the VISCERAL Anatomy Grand Challenge at the 2015 IEEE International Symposium on Biomedical Imaging (ISBI), New York, NY, Apr 16th, 2015 published at <http://ceur-ws.org>

and the experimental setting detailed thereafter. The results on both training and test datasets are discussed in Sec. 4 and compared to some state-of-the-art approaches.

2 Method

Discrete optimisation can capture large motions by defining an appropriate range of displacements \mathbf{u} . It enables a flexible choice of different similarity terms, since no derivative is required. We use the framework presented in [HJBS13], which defines a graphical model with nodes $p \in \mathcal{V}$ (with spatial location \mathbf{x}_p) that correspond to control points in a uniform B-spline grid. For each node, the hidden labels f_p (from a large quantised set \mathcal{L}) are defined as potential 3D displacements $f_p = \mathbf{u}_p = \{u_p, v_p, w_p\}$ between a control point p in the fixed image F and moving image M . Edges between nodes used for inference of the pair-wise regularisation costs $R(f_p, f_q)$ ($p, q \in \mathcal{E}$) are modelled by a minimum spanning tree (MST) for computational efficiency. The displacement field is regularised using the squared differences of the displacements of neighbouring control points:

$$R(f_p, f_q) = \sum_{(p,q) \in \mathcal{E}} \frac{\|\mathbf{u}_p - \mathbf{u}_q\|^2}{\|\mathbf{x}_p - \mathbf{x}_q\|} \quad (1)$$

For the image similarity (data term) self-similarity descriptors are used [HJP⁺13]. The self-similarity context is based on local patch distances within each image and invariant to contrast change, robust to noise and modality independent. The dissimilarity metric D , the L_1 norm between 64 bit binary descriptor representations SSC_F (for fixed image) and SSC_M (for moving image) at two locations \mathbf{x} and $\mathbf{x} + \mathbf{u}$, can be efficiently calculated in the Hamming space:

$$D(\mathbf{x}_p, \mathbf{u}_p) = 1/|\mathcal{P}| \sum_{y \in \mathcal{P}} \Xi\{SSC_F(\mathbf{x}_p + y) \oplus SSC_M(\mathbf{x}_p + \mathbf{u}_p + y)\} \quad (2)$$

where \oplus defines an exclusive OR, Ξ a population count and $y \in \mathcal{P}$ the local patch coordinates. The combined energy function with regularisation parameter α becomes: $E(f) = \sum_{p \in \mathcal{V}} D(f_p) + \alpha \sum_{(p,q) \in \mathcal{E}} R(f_p, f_q)$. Belief propagation [FH06] on the MST (our relaxed graphical) is employed to find the global minimum without iterations in only two passes.

Prior to the deformable registration, a block-matching based linear registration using also the SSC metric is employed as detailed in [HPSH14].¹

3 Experiments

The deformations between different anatomies make a large number of degrees of freedoms necessary. As pre-processing the images are resampled to an isotropic resolution of $1.8 \times 1.8 \times 1.8$ mm³ and padded or cropped to have same dimensions. For the affine pre-registration, three scales of control-point grids with spacings of [9, 8, 7] voxels are used. The displacement label space is defined by two parameters: number of steps l_{\max} and quantisation step q , which together define the label space $\mathcal{L} = q \cdot \{0, \pm 1, \dots, \pm l_{\max}\}^3$ voxels. We used $l_{\max} = [6, 5, 4]$ and $q = [5, 4, 3]$ voxels. For the deformable registration four scale levels with spacings of [8, 7, 6, 5], numbers of steps of $l_{\max} = [6, 5, 4, 3]$ and quantisations of $q = [4, 3, 2, 1]$ voxels were used. The number of random samples and the regularisation weight were left at their default parameters 50 and 2. Inverse consistent is improved by employing a symmetric calculation of deformations (see [HJP⁺13]).

To asses the impact of the similarity metric, we additionally performed experiments using mutual information (MI) and normalised gradient fields (NGF) [HM05]. For a more detailed comparison of the optimisation, we also applied the popular continuous-optimisation based framework NiftyReg [MRT⁺10] (which uses a B-spline parameterisation) with an affine initialisation [ORPA00].

¹Our software is publicly available for download at www.mpheinrich.de (deedsRegSSC)

Table 1: Experimental results for training dataset of VISCERAL Anatomy 3 challenge. Dice volume overlap for the 7 most common organs (psoas major muscles are abbreviated by pmm) in abdominal and thorax scans when using majority voting. The results of [GSG14] are from a different subset of the challenge (Anatomy 2), so there are not directly comparable.

method	liver	spleen	bladder	r kidney	l kidney	r pmm	l pmm	avg
deeds+SSC CT-CT	0.92	0.84	0.82	0.91	0.91	0.85	0.84	0.872
deeds+MI MR-CT	0.77	0.66	0.16	0.52	0.85	0.69	0.64	0.610
deeds+NGF MR-CT	0.77	0.74	0.31	0.55	0.86	0.75	0.73	0.673
deeds+SSC MR-CT	0.82	0.78	0.44	0.62	0.88	0.80	0.79	0.732
NiftyReg+MI MR-MR	0.81	0.79	0.05	0.58	0.77	0.52	0.36	0.554
[GSG14] MR-MR	0.83	0.66	0.21	0.88	0.85	0.64		0.677
deeds+SSC MR-MR	0.80	0.82	0.63	0.55	0.88	0.79	0.76	0.744
proposed Test MR-MR	0.79	0.71	0.36	0.78	0.83	0.76	0.78	0.714

4 Results

Our results are summarised in Table 1 for a subset of 10 training scans of the contrast enhanced (ce) abdominal MRI modality (or thorax/abdominal ceCT) and a leave-one-out validation. It can be seen that MRI segmentation is substantially more challenging yielding average results of Dice overlap for 7 organs of at most 0.744, while the results for the same setting for CT scans are ≈ 0.13 higher. Either of the two compared discrete optimisation strategies, by Gass et al. [GSG14] and our framework [HJBS13], outperforms the continuous optimisation approach of [MRT⁺10]. Using SSC as similarity metric improves the segmentation by 0.12 compared to MI and by 0.06 compared to NGF within the same framework. The multi-modal segmentation, for which we used MRI scans as fixed and CT scans as moving atlas scans, shows nearly identical accuracy to using same modality priors. This is an interesting finding, which could be employed for generating synthetic CT scans from MRI scans, e.g. for MR-PET reconstruction [HSS⁺08]. Due to time limitations only preliminary results for the hidden test datasets could be computed (last row of Table 1), for which we employed only three atlas scans each. We anticipate further improvements for our final results, which will subsequently be published on the VISCERAL leaderboard. The run-time of our algorithm on the virtual machine was on average 4 minutes per registration, which can be reduced with an optimised CPU implementation to less than a minute.

5 Conclusion

We have demonstrated that deformable registration using discrete optimisation enables accurate automatic MRI organ segmentation. Choosing both a robust similarity metric and optimisation strategy has been found to be important for achieving high overlap. Local similarity-weighted atlas performance estimation and advanced label fusion [AL13] may further improve the results. While machine learning techniques alone may not achieve the same accuracy as registration-based approaches for MRI segmentation, the combination of both can boost the performance. In initial experiments, we found that an RDF trained with both atlas-based priors and intensity features [MWG⁺15] improves the segmentation overlap of liver, spleen and kidneys by ≈ 0.06 .

References

- [AL13] Andrew J Asman and Bennett A Landman. Non-local statistical label fusion for multi-atlas segmentation. *Medical Image Analysis*, 17(2):194–208, 2013.

- [CSB09] Antonio Criminisi, Jamie Shotton, and Stefano Bucciarelli. Decision forests with long-range spatial context for organ localization in CT volumes. *MICCAI workshop on Probabilistic Models for Medical Image Analysis*, pages 69–80, 2009.
- [FH06] Pedro Felzenszwalb and Daniel Huttenlocher. Efficient belief propagation for early vision. *International Journal of Computer Vision*, 70:41–54, 2006.
- [GPKC12] Ben Glocker, Olivier Pauly, Ender Konukoglu, and Antonio Criminisi. Joint classification-regression forests for spatially structured multi-object segmentation. In *ECCV 2012*, pages 870–881. Springer, 2012.
- [GSG14] Tobias Gass, Gabor Szekely, and Orcun Goksel. Multi-atlas segmentation and landmark localization in images with large field of view. In *Medical Computer Vision: Algorithms for Big Data*, pages 171–180. Springer, 2014.
- [HJBS13] Mattias P. Heinrich, Mark Jenkinson, Sir Michael Brady, and Julia A. Schnabel. MRF-based deformable registration and ventilation estimation of lung CT. *IEEE Transactions on Medical Imaging*, 32(7):1239–1248, 2013.
- [HJP⁺13] Mattias P. Heinrich, Mark Jenkinson, Bartłomiej W. Papież, Sir Michael Brady, and Julia A. Schnabel. Towards realtime multimodal fusion for image-guided interventions using self-similarities. In *MICCAI, LNCS*, pages 187–194. Springer, 2013.
- [HM05] Eldad Haber and Jan Modersitzki. Beyond mutual information: A simple and robust alternative. In *Bildverarbeitung für die Medizin 2005*, pages 350–354. Springer, 2005.
- [HPSH14] Mattias P. Heinrich, Bartłomiej W. Papież, Julia A. Schnabel, and Heinz Handels. Multispectral image registration based on local canonical correlation analysis. In *MICCAI, LNCS*, pages 202–209. Springer, 2014.
- [HSS⁺08] Matthias Hofmann, Florian Steinke, Verena Scheel, Guillaume Charpiat, Jason Farguhar, Philip Aschoff, Sir Michael Brady, Bernhard Schölkopf, and Bernd J. Pichler. MRI-based attenuation correction for PET/MRI: a novel approach combining pattern recognition and atlas registration. *Journal of Nuclear Medicine*, 49(11):1875–1883, 2008.
- [MRT⁺10] Marc Modat, Gerard R Ridgway, Zeike A Taylor, Manja Lehmann, Josephine Barnes, David J Hawkes, Nick C Fox, and Sébastien Ourselin. Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, 98(3):278–284, 2010.
- [MWG⁺15] Oskar Maier, Matthias Wilms, Janina von der Gablentz, Ulrike M Krämer, Thomas F Münte, and Heinz Handels. Extra tree forests for sub-acute ischemic stroke lesion segmentation in MR sequences. *Journal of Neuroscience Methods*, 240:89–100, 2015.
- [ORPA00] Sébastien Ourselin, Alexis Roche, Sylvain Prima, and Nicholas Ayache. Block matching: A general framework to improve robustness of rigid registration of medical images. In *MICCAI 2000*, pages 557–566. Springer, 2000.