

Good Features for Reliable Registration in Multi-Atlas Segmentation

Fredrik Kahl^{1,2} Jennifer Alvé¹ Olof Enqvist¹ Frida Fejne¹
Johannes Ulén² Johan Fredriksson² Matilda Landgren² Viktor Larsson²

¹Department of Signals and Systems
Chalmers University of Technology, Sweden

²Centre for Mathematical Sciences
Lund University, Sweden

Abstract

This work presents a method for multi-organ segmentation in whole-body CT images based on a multi-atlas approach. A robust and efficient feature-based registration technique is developed which uses sparse organ specific features that are learnt based on their ability to register different organ types accurately. The best fitted feature points are used in RANSAC to estimate an affine transformation, followed by a thin plate spline refinement. This yields an accurate and reliable nonrigid transformation for each organ, which is independent of initialization and hence does not suffer from the local minima problem. Further, this is accomplished at a fraction of the time required by intensity-based methods. The technique is embedded into a standard multi-atlas framework using label transfer and fusion, followed by a random forest classifier which produces the data term for the final graph cut segmentation. For a majority of the classes our approach outperforms the competitors at the VISCERAL Anatomy Grand Challenge on segmentation at ISBI 2015.

1 Introduction

Segmentation is a key problem in medical image analysis, and may be used for numerous applications in medical research and clinical care. In this paper, a pipeline for the segmentation of whole-body CT images into 20 different organs is presented. The approach is based on multi-atlas segmentation, see [KSK⁺10, HKA⁺10, WSD⁺13] and the references therein, an approach which is known to produce state-of-the-art results for several segmentation tasks. The method requires pair-wise registrations from a set of atlas images to the unknown target image.

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>

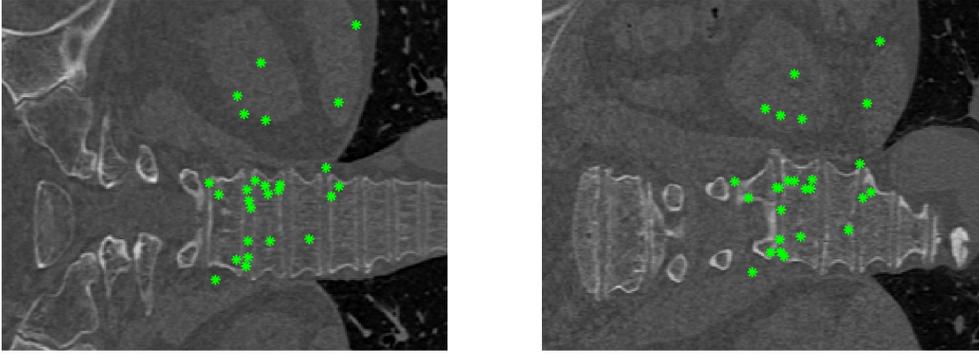


Figure 1: Two CT slices of a target (*left*) and atlas (*right*) with corresponding features after RANSAC for Lumbar Vertebra 1.

In principle, there are two different approaches to image registration, feature-based and intensity-based registration, see the surveys [KBG⁺11, SDP13]. Intensity-based methods are capable of producing accurate registrations but are sensitive to initialization and often slow. Feature-based methods are usually faster, but may risk failing due to many outlier correspondences between the images. Our approach is an adapted feature-based method that utilizes the speed of general feature-based methods while trying to eliminate the risk of establishing incorrect point-to-point correspondences between the images by identifying reliable feature points. We show that reliable organ localization can be computed using (i) robust optimization techniques and (ii) learned feature correspondences.

2 Proposed Solution

Our system segments each organ independently of each other using a multi-atlas approach. The pipeline has three steps:

1. Feature-based registration with RANSAC.
2. Label fusion with a random forest classifier.
3. Graph cut segmentation with a Potts model.

These steps will now be described in more detail.

1. Feature-based registration with RANSAC. In order to register an atlas image to the target, a feature-based approach is used. Sparse features are extracted according to Svärm et al. [SEKO15], which uses a method similar to SIFT for feature detection and SURF for feature description. Typically around 8,000–10,000 features are extracted from a $512 \times 512 \times 800$ CT image, which takes less than 30s. Correspondences are obtained by matching a *subset* of the features in the atlas image to the features in the target. The matching is done with a symmetric neighbor approach, where each descriptor is matched to its nearest neighbor. The organ-specific subset of atlas features is determined as a pre-processing step in the following way. For each atlas image and organ, *golden transformations* are established to the other atlas images using the ground truth segmentations. Then, based on these transformations, one can check which features in the atlas image of interest that are most accurate, and rank the features accordingly. We have found empirically that using the top 300 best features for each organ provides robust and reliable registration. Standard RANSAC

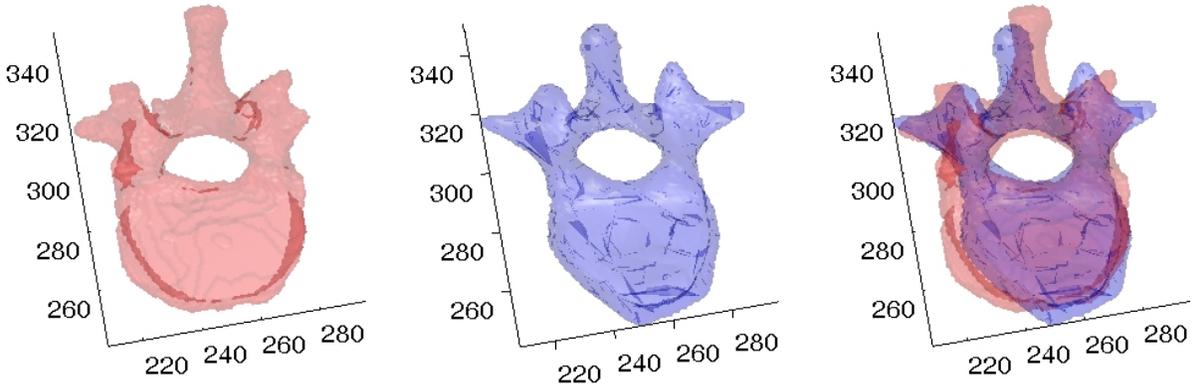


Figure 2: Example registration of Lumbar Vertebra 1. *Left*: the atlas ground truth mask. *Middle*: the TPS warped target mask. *Right*: The masks overlaid in the same coordinate system.

with the truncated ℓ_2 as cost function is used in order to remove outliers. The optimization is run 500,000 iterations and the truncation threshold was set to 30 mm. See Figure 1 for an example.

Finally, a coordinate transformation from the atlas to the target image is computed by applying thin plate splines (TPS) to the remaining correspondences, and thereafter used in order to transfer the labels of the atlas to the target image. The thin plate spline method proposed in [CR00] was used. One registration takes less than 10s in total. See Figure 2 for an example.

For several of the organs, the registrations are refined with a standard intensity-based method. More specifically, we used NiftyReg [ORS⁺01] which takes around 100–200s per registration.

2. Label fusion with a random forest classifier. The pairwise feature-based registrations give us a rough idea where the organ is located. In order to fuse the transferred labels, one can compute an average voxel map, denoted P . Hence, for voxel i , the map $P(i)$ gives a number between 0 and 1 which can (intuitively) be interpreted as the probability of voxel i belonging to the organ. For example, if half of the atlas images say that voxel i should be organ, then $P(i) = 0.5$.

The map P largely ignores the local appearance around the target organ and in order to improve the accuracy of the estimate, the map P along with a few other features is used to train a random forest classifier. We use Sherwood [CSK11] to train and evaluate large random forest instances efficiently. Given a target volume I and the map P , we begin by smoothing both using a Gaussian kernel with standard deviation $\sigma = 1$, resulting in two new volumes denoted I_s and P_s . Using cross validation we also determine a threshold level, τ , for P and construct a distance map D , where each voxel in D equals the (signed) distance to the boundary surface of the binary volume $P > \tau$, that is, the map P thresholded at τ . For each volume I we thus obtain 5 features per voxel i : $I(i)$, $I_s(i)$, $P(i)$, $P_s(i)$ and $P_t(i)$. The output of the classifier produces yet another map, denoted P_r , which is a refined estimate of the location of the organ. As previously, P_r can be interpreted as the probability for each voxel belonging to the organ of interest. See Figure 3 for an example.

3. Graph cut segmentation with a Potts model. The random forest classifier generates a new estimate $P_r(i)$ for each voxel i , but each decision in the classifier is taken independently of the output of the neighboring voxels of i . This has a tendency of producing noisy boundary estimates and therefore we will regularize the solution by using a standard Potts model. The final solution can then be computed with graph-cuts.

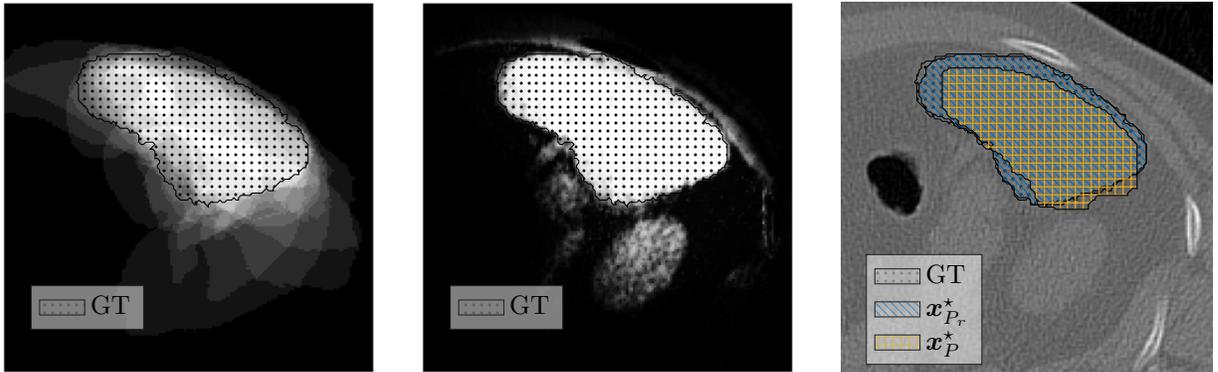


Figure 3: Example of the resulting probability estimates and segmentation of the spleen for one CT slice; in each image the ground truth (GT) is indicated. *Left*: the initial probability, P . *Middle*: the probability given by random forest, P_r . *Right*: the resulting segmentation \mathbf{x}_P^* using P and $\mathbf{x}_{P_r}^*$ using P_r overlaid on the original image.

The Potts model penalizes neighboring voxels if they take different labels. Let x_i be a Boolean indicator variable for voxel i , i.e., $x_i \in \{0, 1\}$. Then, for two neighboring voxels x_i and x_j , the cost should be zero if $x_i = x_j$ and λ otherwise, where λ is a positive scalar. This cost can compactly be written as $\lambda x_i(1 - x_j)$. Further, the data cost for voxel i is set to take value $1/2 - P_r(i)$ if $x_i = 1$ and zero otherwise. This favors voxels with probabilities in the interval $[0.5, 1]$ to be foreground and voxels with $[0, 0.5]$ to be background.

In summary the final segmentation, \mathbf{x}^* , is given by the solution to the optimization problem:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \{0,1\}^n} \sum_{i=1}^n x_i \left(\frac{1}{2} - P_r(i) \right) + \lambda \sum_{i=1}^n \sum_{j \in \mathcal{N}(i)} \mu_{ij} x_i (1 - x_j), \quad (1)$$

where λ is a regularization weight and μ_{ij} compensates for anisotropic resolution. For all organs we use a 6-connected neighborhood \mathcal{N} . In order to save memory and speed-up calculations we only process a volume around the zero level of the distance map D with a 20 voxels margin. The function in (1) is submodular and is minimized efficiently using the graph-cut implementation of [JSH12].

3 Experimental Results

All the tuning parameters in our system have been set by leave-one out cross validation on the first 15 of the 20 whole-body CT images available in the VISCERAL challenge. The 5 remaining images have been used to validate the performance of the random forest classifier and the graph-cut segmentation. In the training phase, the first 15 images have served as the atlas set, while in the final version all 20 images are utilized in the atlas.

Our system has been evaluated on a test set of 10 whole-body CT images by the organizers of the VISCERAL Anatomy Grand Challenge at ISBI 2015. Note that this test set is only available to the organizers. The final results are given in Table 1 together with the best competitors to date:

- *CMIV* - “Center for Medical Image Science and Visualization, Linköping University”,
- *HES-SO* - “University of Applied Sciences Western Switzerland” and
- *SIAT* - “Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences”.

<i>Organ</i>	<i>Ours</i>	<i>CMIV</i>	<i>HES-SO</i>	<i>SIAT</i>
Left Kidney	0.934	0.896	0.784	-
Right Kidney	0.915	0.796	0.790	-
Spleen	0.870	0.910	0.703	0.874
Liver	0.921	0.936	0.866	0.923
Left Lung	0.972	0.961	0.972	0.952
Right Lung	0.975	0.970	0.975	0.957
Unirary Bladder	0.763	0.713	0.698	-
Muscle Body of Left Rectus Abdominis	0.746	-	0.551	-
Muscle Body of Right Rectus Abdominis	0.679	-	0.519	-
Lumbar Vertebra 1	0.775	-	0.718	-
Thyroid	0.424	-	0.549	-
Pancreas	0.383	-	0.408	-
Left Psoas Major Muscle	0.861	0.828	0.806	-
Right Psoas Major Muscle	0.847	0.817	0.787	-
Gallbladder	0.190	-	0.276	-
Sternum	0.847	-	0.761	-
Aorta	0.830	-	0.753	-
Trachea	0.931	-	0.92	-
Left Adrenal Gland	0.282	-	0.373	-
Right Adrenal Gland	0.220	-	0.355	-
Average	0.718	-	0.678	-

Table 1: Final results measured in DICE metric for whole-body CT images. Our approach gives the best results for 13 out of the 20 organs. Here ‘-’ means that no segmentation was provided.

4 Conclusions

We have demonstrated that by using a feature-based approach to multi-atlas segmentation, it is possible to reliably locate and segment organs in whole-body CT images with state-of-the-art results. Still, there is room for improvement. For example, there is no guarantee that the system produces a valid organ shape. We are currently working on ways to directly incorporate such shape priors in the framework. Further, the speed of the system can be improved, for example, by circumventing the need to perform 20 pairwise registrations for every new target image [ANEK15].

References

- [ANEK15] J. Alvé, A. Norlén, O. Enqvist, and F. Kahl. Überatlas: Robust speed-up of feature-based registration and multi-atlas segmentation. *Scandinavian Conference on Image Analysis*, 2015. To appear.
- [CR00] H. Chui and A. Rangarajan. A new algorithm for non-rigid point matching. *IEEE Conference on Computer Vision and Pattern Recognition*, 2:44–51, 2000.
- [CSK11] A. Criminisi, J. Shotton, and E. Konukoglu. Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Foundations and Trends® in Computer Graphics and Vision*, 7:81–227, 2011.
- [HKA⁺10] RA. Heckemann, S. Keihaninejad, P. Aljabar, D. Rueckert, JV. Hajnal, and A. Hamers. Improving intersubject image registration using tissue-class information benefits

- robustness and accuracy of multi-atlas based anatomical segmentation. *NeuroImage*, 51(1):221–227, 2010.
- [JSH12] O. Jamriška, D. Sýkora, and A. Hornung. Cache-efficient graph cuts on structured grids. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3673–3680, 2012.
- [KBG⁺11] F. Khalifa, GM. Beache, G. Gimel’farb, JS. Suri, and AS. El-Baz. State-of-the-art medical image registration methodologies: A survey. In *Multi Modality State-of-the-Art Medical Image Segmentation and Registration Methodologies*, pages 235 – 280. Springer, 2011.
- [KSK⁺10] HA. Kirisli, M. Schaap, S. Klein, LA. Neefjes, AC. Weustink, T. van Walsum, and WJ. Niessen. Fully automatic cardiac segmentation from 3d cta data: a multi-atlas based approach. *Proceedings of SPIE*, 2010.
- [ORS⁺01] S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache. Reconstructing a 3D structure from serial histological sections. *Image and vision computing*, 19(1):25–31, 2001.
- [SDP13] A. Sotiras, C. Davatzikos, and N. Paragios. Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging*, 32(7):1153–1190, 2013.
- [SEKO15] L. Svärm, O. Enqvist, F. Kahl, and M. Oskarsson. Improving robustness for inter-subject medical image registration using a feature-based approach. *International Symposium on Biomedical Imaging*, 2015.
- [WSD⁺13] H. Wang, JW. Suh, SR. Das, J. Pluta, C. Craige, and PA. Yushkevich. Multi-atlas segmentation with joint label fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(3):611–623, 2013.