

# VoteNet: A Deep Learning Label Fusion Method for Multi-Atlas Segmentation

Zhipeng Ding<sup>1</sup>, Xu Han<sup>1</sup>, Marc Niethammer<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, UNC Chapel Hill, USA

<sup>2</sup>Biomedical Research Imaging Center, UNC Chapel Hill, USA

**Abstract.** Deep learning (DL) approaches are state-of-the-art for many medical image segmentation tasks. They offer a number of advantages: they can be trained for specific tasks, computations are fast at test time, and segmentation quality is typically high. In contrast, previously popular multi-atlas segmentation (MAS) methods are relatively slow (as they rely on costly registrations) and even though sophisticated label fusion strategies have been proposed, DL approaches generally outperform MAS. In this work, we propose a DL-based label fusion strategy (VoteNet) which locally selects a set of reliable atlases whose labels are then fused via plurality voting. Experiments on 3D brain MRI data show that by selecting a good initial atlas set MAS with VoteNet significantly outperforms a number of other label fusion strategies as well as a direct DL segmentation approach. We also provide an experimental analysis of the upper performance bound achievable by our method. While unlikely achievable in practice, this bound suggests room for further performance improvements. Lastly, to address the runtime disadvantage of standard MAS, all our results make use of a fast DL registration approach.

## 1 Introduction

Image segmentation, i.e. giving pixels or voxels in an image meaningful labels, is an important medical image analysis task [7]. While a large number of segmentation approaches exist, deep convolutional neural networks (CNNs) have shown remarkable segmentation performance and are considered state-of-the-art [4,10]. This was not always the case. Prior to the dominance of CNNs, multi-atlas segmentation (MAS) has been highly popular and successful for medical image segmentation [7]. MAS approaches rely on a set of previously labeled atlas images. These images are then registered to an unlabeled target image and their associated labels are used to infer the labeling of the target image via label fusion. Hence, MAS performance relies on high-quality registrations or advanced label fusion methods. MAS is slow as it requires computationally costly registrations, but provides good spatial consistency via the given atlas segmentations.

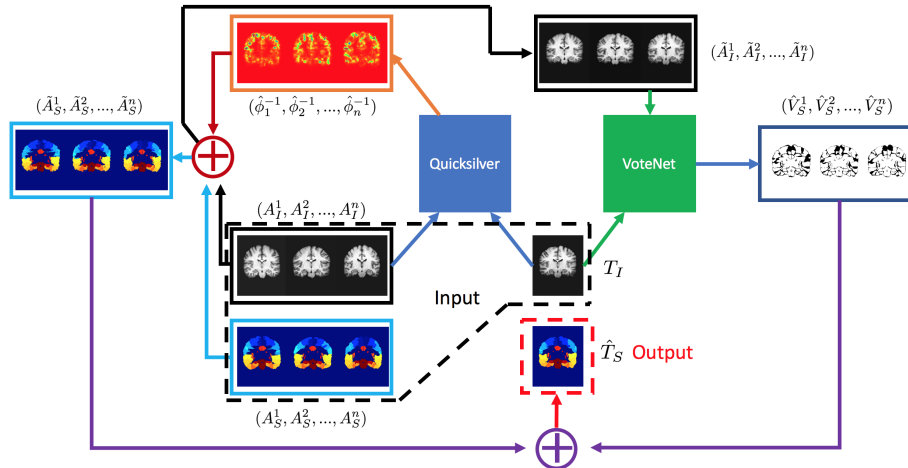
In contrast, CNN approaches use sophisticated network architectures, with parameters trained on large sets of labeled images. A popular architecture for medical image segmentation is the U-Net [4]. For DL approaches, the majority

of computational cost is spent during training. Hence, these approaches are fast at test time and have shown excellent segmentation performance for medical images [4,10]. However, as image labels are not directly spatially transformed, spatial consistency is only indirectly encouraged during training and DL approaches may miss or add undesired structures. Furthermore, large numbers of labeled images are desirable for training, but may not always be available.

Conceptually, MAS is attractive as it provides a direct and intuitive way to specify and obtain segmentations via a set of labeled atlases. Given atlases that are similar to the target image to be segmented, it is plausible that good registrations can be achieved, that atlas labels can be transferred well, and consequentially that high-quality segmentations can be obtained. However, it is a-priori unclear which atlases should be used to estimate the segmentation as not all atlases will align well via registration. Label fusion strategies aim at addressing the resulting spatial inconsistencies between the atlas labels warped to the target image space. Approaches include, majority and plurality voting [6,5], global weighted voting [1], and local weighted voting strategies [13]. Statistical modeling approaches have also been proposed [14] and patch-based approaches directly aim to compensate for local registration errors [2]. Most recently, machine learning [12] and deep learning [15] label fusion methods have been proposed. These methods all assume that all atlases might contribute to the labeling decision. Instead, we propose making decisions only based on atlases considered trustworthy. In contrast to global [11] and patch-based atlas selection [9], our approach *locally* predicts the set of trustworthy atlases voxel by voxel. Experimentally, we show that this strategy, even combined with simple plurality voting, has excellent segmentation performance on par or slightly outperforming a U-Net and significantly outperforming other label fusion strategies. All our results make use of fast DL-based deformable image registration, thereby resulting in a MAS approach which is fast and accurate while providing more direct control over spatial label consistency.

**Contributions.** (1) *New label fusion method:* We propose a DL label fusion method (VoteNet), which locally identifies sets of trustworthy atlases. (2) *Fast implementation:* All our results are based on a fast DL registration approach (Quicksilver). (3) *Performance upper bound:* We experimentally assess the best possible performance achievable with our approach and illustrate that there is a large margin for improvements. (4) *Comprehensive experimental comparison:* We compare to a variety of other label fusion strategies and a U-Net for 3D brain segmentation. Our approach performs consistently best.

**Organization.** Sec. 2 describes our proposed MAS framework in detail. Sec. 3 discusses experimental details and results on the LPBA40 dataset. Sec. 4 concludes the paper with a summary and an outlook on future work.



**Fig. 1:** Our MAS framework. Atlas images and target image are input into Quicksilver to predict transformation maps  $\hat{\phi}^{-1}$  from atlas to target space. The atlas images and labels are then warped to target space by  $\hat{\phi}^{-1}$ . VoteNet uses the warped atlas images and the target image to predict a binary mask  $\hat{V}_S$  for each atlas indicating locally if an atlas should be used. Final label fusion is done via plurality voting using only the atlases which should be used according to the VoteNet prediction. See Eq. 2.


## 2 Methodology

Before discussing our proposed label fusion method, we first describe multi-atlas segmentation. Let  $T_I$  be the target image to be segmented and  $A^1 = (A_I^1, A_S^1)$ ,  $A^2 = (A_I^2, A_S^2)$ , ...,  $A^n = (A_I^n, A_S^n)$  be  $n$  atlas images and their corresponding manual segmentations. Assume there is a reliable deformable image registration method that warps all atlases into the space of the target image  $T_I$ , i.e.  $\tilde{A}^1 = (\tilde{A}_I^1, \tilde{A}_S^1)$ ,  $\tilde{A}^2 = (\tilde{A}_I^2, \tilde{A}_S^2)$ , ...,  $\tilde{A}^n = (\tilde{A}_I^n, \tilde{A}_S^n)$ . Each  $\tilde{A}_S^i$  is now a candidate segmentation for  $T_I$ . Finally, a label fusion method  $\mathcal{F}$  combines all the candidate segmentations to produce the final segmentation  $\hat{T}_S$  for  $T_I$ , i.e.,

$$\hat{T}_S = \mathcal{F}(\tilde{A}^1, \tilde{A}^2, \dots, \tilde{A}^n, T_I). \quad (1)$$

Our framework uses two deep convolutional neural networks (CNNs) for MAS: the Quicksilver registration network to compute the spatial transformation to target image space and our label fusion network (VoteNet). While other registration approaches could be used, a DL approach greatly speeds-up the typically slow registrations for MAS, when based on numerical optimization. By using a DL approach we also demonstrate that it integrates well with MAS.

Fig. 1 illustrates our approach. Quicksilver and VoteNet are discussed below. **Quicksilver:** Quicksilver [16] uses the target image  $T_I$  and the atlas images  $(A_I^1, A_I^2, \dots, A_I^n)$  to predict the deformation maps  $(\hat{\phi}_1^{-1}, \hat{\phi}_2^{-1}, \dots, \hat{\phi}_n^{-1})$  which are used to generate warped atlas images  $(\tilde{A}_I^1, \tilde{A}_I^2, \dots, \tilde{A}_I^n)$  and their corresponding labels  $(\tilde{A}_S^1, \tilde{A}_S^2, \dots, \tilde{A}_S^n)$ . Using a DL registration approach such as Quicksilver



**Fig. 2:** VoteNet architecture. We use a U-Net structure with skip connections and batch normalization. Numbers indicate feature numbers (channels) in each convolutional layer. Inputs are target image and a warped atlas image. Output is a binary mask (after thresholding the sigmoid output at 0.5) to select trustworthy atlas voxels.

speeds-up pairwise registrations by at least an order of magnitude [16] compared to numerical optimization. Since registrations are the computational bottleneck of MAS, similar speed-ups can be obtained. E.g., MAS with our approach and 17 atlases requires only 15 mins on an NVIDIA GTX1080Ti. Experiments (Sec. 3) show that Quicksilver yields good results when combined with MAS (Tab. 1).

**VoteNet:** Given the warped atlas images ( $\tilde{A}_I^1, \tilde{A}_I^2, \dots, \tilde{A}_I^n$ ) and the target image, VoteNet independently predicts binary masks ( $\hat{V}_S^1, \hat{V}_S^2, \dots, \hat{V}_S^n$ ) for each warped atlas image, locally indicating if a warped atlas should be considered for the final labeling decision of the target image. In effect, VoteNet predicts for each spatial location the *set of trustworthy atlases*; all atlases which should likely not be used are discarded. Hence, VoteNet, implements a form of locally adaptive statistical trimming. Once the set of trustworthy atlases has been determined, their associated labels can be fused with any chosen label fusion strategy. For simplicity we use plurality voting. Our VoteNet strategy shifts the notion of a plurality to a *plurality of trusted atlases*. We define *trusted plurality voting* as

$$\hat{T}_S(x) = \arg \max_{l \in \Omega} \sum_{i=1}^n \mathbb{1}[\hat{V}_S^i(x) \odot \tilde{A}_S^i(x) = l], \quad (2)$$

where  $l \in \Omega = \{0, \dots, N\}$  is the set of labels ( $N$  structures; 0 indicating background),  $\mathbb{1}[\cdot]$  is the indicator function and  $x$  denotes a voxel position. We define  $\hat{V}_S^i(x) \odot \tilde{A}_S^i(x) := \tilde{A}_S^i(x) \in \Omega$  if  $\hat{V}_S^i(x) = 1$ ;  $\hat{V}_S^i(x) \odot \tilde{A}_S^i(x) := -1$  if  $\hat{V}_S^i(x) = 0$ .

**VoteNet training:** Fig. 2 shows the VoteNet architecture, which is based on the 3D U-Net [4] (but uses the target image and a warped atlas image as inputs). VoteNet processes an image patch-wisely, with a patch size of  $72 \times 72 \times 72$  from the target image  $T_I$  and a warped atlas image  $\tilde{A}_I$  at the same position, where the  $40 \times 40 \times 40$  patch center is used to tile the volume. To train VoteNet, we use 20

images from LONI Probabilistic Brain Atlas dataset (LPBA40)<sup>1</sup>; we randomly select 17 images and their labels as atlases and use Quicksilver to register all 17 atlases to the 20 images excluding themselves. This results in 323 ( $17 \times 16 + 17 \times 3$ ) pairwise registrations. Given the manual segmentation labels of these 20 images, we determine at which location the warped atlas labels agree (1) and where they disagree (0). These agreement/disagreement labels are our training labels for VoteNet, trained via a binary cross entropy loss (Fig. 2) in PyTorch. As VoteNet produces continuous outputs,  $P_i$ , (in  $(0, 1)$  via a sigmoid) we threshold at 0.5 at test time, i.e., the local prediction is  $\hat{V}_S^i = 0$  if  $P_i < 0.5$  and  $\hat{V}_S^i = 1$  otherwise. We train using ADAM over 300 epochs with a multi-step learning rate. The initial learning rate is 0.001; reduced to 0.0001 after 150 epochs; and finally to 0.00001 after 250 epochs. Training image patches were randomly extracted so that 0 labels account for no less than 5% of the entire patch volume. Training requires  $\approx 24$ h. The prediction of a single atlas mask takes  $< 20$ s.

### 3 Experimental Results and Discussion

We use LPBA40 for evaluation. We use two-fold cross-validation, i.e. the dataset is randomly divided into two non-overlapping subsets of equal size. One set is used for training (Sec. 2) the other for testing for each of the two cross-validation experiments. The results below are averaged over the cross-validation folds.

**Benchmark methods:** We compare against plurality voting (PV) [6], majority voting (MV), simultaneous truth and performance level estimation (STAPLE) [14], multi-atlas based multi-image segmentation (MAMBIS)<sup>2</sup> [8], joint label fusion (JLF) [13], patch-based label fusion (PB) [2] and a U-Net [4]. PV locally assigns the most frequent segmentation label among the atlases. MV assigns a label only if more than half of the atlases ( $\geq 9$ ) agree. STAPLE uses a statistical model to estimate a true hidden segmentation based on an optimal weighting of the segmentations. MAMBIS puts atlases in a tree structure to consider their correlations for concurrent alignment. JLF regards label fusion as an optimization problem and minimizes the total expectation of labeling errors. PB searches in a neighborhood to reduce registration errors and utilizes patch intensity and label information within a Bayesian label fusion framework. We include a U-Net for a direct comparison to a popular DL image segmentation approach. We also create an oracle label fusion strategy, which has access to the true label during local atlas selection. This allows establishing upper performance bounds.

**Oracle label fusion:** MAS depends on the interplay of image registration and label fusion. Conceivably, given a high-quality registration, one should

<sup>1</sup> LPBA40 contains 40 3D brain MRIs with 56 manually segmented structures. Preprocessing includes affine registration to the MNI152 atlas and histogram equalization.

<sup>2</sup> MAMBIS uses Diffeomorphic Demons for registration and hence results are not directly comparable to ours; we use Quicksilver for all other label fusion methods.


Method	Avg. Surf. Dist. (mm)	Surf. Dice (%)	Hausdorff Dist. (mm)	95% Max. Dist. (mm)	Volume Dice (%)
Oracle(1)	0.04 ± 0.01	99.06 ± 0.23	3.64 ± 0.32	0.18 ± 0.08	98.99 ± 0.19
Oracle(3)	0.17 ± 0.03	95.91 ± 0.66	5.80 ± 0.43	1.17 ± 0.16	96.45 ± 0.46
Oracle(6)	0.45 ± 0.05	88.97 ± 1.17	8.00 ± 0.51	2.46 ± 0.22	91.53 ± 0.76
Oracle(9)	0.86 ± 0.07	79.11 ± 1.49	9.87 ± 0.52	3.86 ± 0.28	84.79 ± 0.95
PV	1.19 ± 0.09	70.11 ± 2.02	<b>9.97 ± 0.53</b>	4.22 ± 0.29	77.47 ± 1.18
MV	1.26 ± 0.09	68.88 ± 2.03	10.34 ± 0.52	4.44 ± 0.29	77.04 ± 1.17
STAPLE	2.90 ± 0.29	62.54 ± 1.99	21.77 ± 0.98	8.88 ± 0.73	73.83 ± 1.13
MABMIS	1.24 ± 0.09	68.66 ± 2.07	10.07 ± 0.52	4.36 ± 0.28	77.89 ± 1.22
JLF	1.15 ± 0.08	72.44 ± 1.73	10.08 ± 0.54	4.30 ± 0.29	79.50 ± 1.12
PB	1.31 ± 0.09	67.00 ± 2.01	10.28 ± 0.50	4.54 ± 0.30	77.60 ± 1.16
U-Net	1.10 ± 0.12	74.69 ± 2.15	13.79 ± 3.86	4.20 ± 0.49	80.46 ± 1.29
VoteNet	1.02 ± 0.08	75.32 ± 1.90	10.89 ± 0.55	3.84 ± 0.28	80.55 ± 1.13
VoteNet + U-Net	<b>0.99 ± 0.08</b>	<b>76.23 ± 1.83</b>	11.67 ± 1.48	<b>3.78 ± 0.28</b>	<b>80.75 ± 1.11</b>

**Table 1:** Evaluation metrics for LPBA40 segmentation performance. **Avg. Surf. Dist.:** symmetric average surface distance in mm between each segmentation label and the true segmentation. **Surf. Dice:** Dice score of segmented label surface and true label surface at a tolerance of 1 mm. **Hausdorff. Dist.:** Hausdorff distance between segmented label volume and true label volume. **95% Max. Dist.:** 95 percentile of the maximum symmetric distance between segmented label volume and true label volume. **Volume Dice:** Average Dice score over segmented labels (excluding the background). We use a Mann-Whitney U-test to check for significant differences to VoteNet+U-Net. We use a significance level of 0.05 and the Benjamini/Hochberg correction [3] for multiple comparisons with a false discovery rate of 0.05. Results are highlighted in green if Vote+U-Net performs significantly better than the corresponding method. VoteNet is always better than U-Net and VoteNet+U-Net achieves the best performance.


be able to obtain a high-quality segmentation. To assess how well an ideal label fusion strategy could work, we investigate the behavior of an oracle label fusion method following our Quicksilver atlas to target image registrations. Specifically, Oracle( $n$ ) assigns the correct label to a voxel if at least  $n$  warped atlases (out of our 17) correctly label this voxel; otherwise the background label (0) is assigned.

**Results:** We use five measures to evaluate segmentations: average surface distance, average surface Dice score (i.e., a surface element is considered overlapping if it is within a certain distance ( $\leq 1$ mm) to the other surface), Hausdorff distance, 95% maximum surface distance, and average volume Dice score.

*Oracle results:* Tab. 1 shows that Dice scores of Oracle(1) are close to 100%, indicating that at least one warped atlas label image locally agrees with the manual segmentation. Even Oracle(9) (where the correct label is only assigned if at least 9 of the 17 atlases agree on this labeling) results in Dice scores higher than state-of-the-art approaches. In contrast, MV is significantly worse than Oracle(9). Note that all labels (excluding background) of Oracle(9) are also contained in MV. Hence, MV contains *incorrect* labels for which the majority of atlases *agree*. Therefore, if VoteNet can identify good subsets of these atlases, a good segmentation should be achievable by majority or plurality voting.



**Fig. 3:** One example of the output of VoteNet and its improvement over plurality voting. **Left two:**  $V_S$ : difference between true target image label and warped atlas label for one atlas;  $\hat{V}_S$ : predicted difference from VoteNet. Black indicates that labels are different, white indicates agreement. VoteNet prediction captures most of the label differences. **Right two:**  $P_P$ : recall map (i.e., probability of true positives out of all atlases) for plurality voting.  $P_V$ : recall map of VoteNet. VoteNet greatly improves recall, because it filters out true negatives so that the final voting is more accurate.



**Fig. 4:** 3D rendering of segmented results. Compared with the manual segmentation (GS) U-Net results show shape inconsistencies, highlighted by the red arrows. VoteNet and VoteNet+U-Net retain spatial consistency. Plurality voting (PV) results smooth out cortical structures, but these are well maintained by VoteNet and VoteNet+U-Net.

*Label fusion results:* Tab. 1 shows that VoteNet greatly improves performance over MV/PV and significantly outperforms all other evaluated label fusion strategies on most measures. These results also illustrate that VoteNet successfully locally eliminates atlases that would otherwise have tipped the results to incorrect PV label assignments. Further, we observed that there are some voxels ( $\approx 7.5\%$  inside the brain) that are not assigned any labels by VoteNet (i.e., locally all warped atlases are rejected). We therefore propose a combined VoteNet + U-Net strategy which fills in missing voxels via the U-Net segmentation. This strategy outperforms both VoteNet and U-Net. There is still a large gap between our VoteNet and the Oracle results. Hence, there is significant room for future improvement. Fig. 3 illustrates the performance of VoteNet. The predicted binary mask  $\hat{V}_S$  is close to the ground truth binary mask  $V_S$ , indicating that VoteNet captures most areas of poor label alignments for a given atlas image. In fact, VoteNet achieves a volume Dice score of 0.86 on local atlas

selection. Fig. 3(right) also shows that by only retaining locally trustworthy atlases the percentage of true positives (after VoteNet atlas selection) over all atlases grows significantly. Consequentially, subsequent plurality voting better predicts the correct labels.

*U-Net results:* U-Net results are generally good with respect to the volumetric Dice scores. However, as indicated by the surface measures (in particular, Hausdorff and 95% maximum surface distance), shapes of segmented structures may locally be distorted, as shape constraints are not straightforward to integrate into a CNN. This drawback is much less present in MAS, as a good deformable image registration method will preserve local structure and topology in target image space (based on transformation smoothness). Fig. 4 illustrates this effect. As highlighted by the red arrows, U-Net results often show inconsistent shapes, while VoteNet and VoteNet+U-Net produce shapes more consistent with the manual segmentations. Furthermore, our VoteNet and VoteNet+U-Net retain the cortical foldings, while PV tends to flatten them. This indicates that our proposed approach indeed complements a label fusion method such as PV well.

## 4 Conclusion and Future Work

We presented a new label fusion method (VoteNet) which helps *locally* select the most trustworthy atlases. With VoteNet, we achieve state-of-the-art segmentation performance, even surpassing a deep network (U-Net) while maintaining spatial shape consistency. We also provided an empirical analysis of best possible achievable performance of our approach, indicating that there is still substantial room for further performance improvements. In particular, it would be interesting to combine VoteNet with more advanced label fusion strategies than plurality voting. As such strategies have shown improved performance for MAS, it is conceivable that they could also further improve our approach, for example, by leveraging local image information to assess atlas to target image similarity. It would also be valuable to explore more advanced network architectures as well as end-to-end formulations integrating the registration network.

**Acknowledgements.** Research reported in this work was supported by the National Institutes of Health (NIH) and the National Science Foundation (NSF) under award numbers NSF EECS-1711776 and NIH 1R01AR072013. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or the NSF.

## References

1. Artaechevarria, X., Muñoz-Barrutia, A., Ortiz-de Solórzano, C.: Efficient classifier generation and weighted voting for atlas-based segmentation: Two small steps faster and closer to the combination oracle. In: SPIE. vol. 6914 (2008)



2. Bai, W., Shi, W., O'regan, D.P., Tong, T., Wang, H., Jamil-Copley, S., Peters, N.S., Rueckert, D.: A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: application to cardiac MR images. *TMI* 32(7), 1302–1315 (2013)
3. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57(1), 289–300 (1995)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: *MICCAI*. pp. 424–432 (2016)
5. Hansen, L., Salamon, P.: Neural network ensembles. *PAMI* (10), 993– (1990)
6. Heckemann, R.A., Hajnal, J.V., Aljabar, P., Rueckert, D., Hammers, A.: Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33(1), 115–126 (2006)
7. Iglesias, J.E., Sabuncu, M.R.: Multi-atlas segmentation of biomedical images: a survey. *MEDIA* 24(1), 205–219 (2015)
8. Jia, H., Yap, P.T., Shen, D.: Iterative multi-atlas-based multi-image segmentation with tree-based registration. *NeuroImage* 59(1), 422–430 (2012)
9. Konukoglu, E., Glocker, B., Zikic, D., Criminisi, A.: Neighbourhood approximation using randomized forests. *MEDIA* 17(7), 790–804 (2013)
10. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR*. pp. 3431–3440 (2015)
11. Sanroma, G., Wu, G., Gao, Y., Shen, D.: Learning to rank atlases for multiple-atlas segmentation. *TMI* 33(10), 1939–1953 (2014)
12. Wang, H., Cao, Y., Syeda-Mahmood, T.: Multi-atlas segmentation with learning-based label fusion. In: *MLMI*. pp. 256–263 (2014)
13. Wang, H., Suh, J.W., Das, S.R., Pluta, J.B., Craige, C., Yushkevich, P.A.: Multi-atlas segmentation with joint label fusion. *PAMI* 35(3), 611–623 (2013)
14. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *TMI* 23(7), 903 (2004)
15. Yang, H., Sun, J., Li, H., Wang, L., Xu, Z.: Neural multi-atlas label fusion: Application to cardiac MR images. *MEDIA* 49, 60–75 (2018)
16. Yang, X., Kwitt, R., Niethammer, M.: Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage* 158, 378–396 (2017)