Sparse Patch-Based Label Fusion for Multi-Atlas Segmentation

Daoqiang Zhang[1,2], Qimiao Guo[1,2], Guorong Wu[1], and Dinggang Shen[1]

[1] Dept. of Radiology and BRIC, University of North Carolina at Chapel Hill, NC 27599
[2] Dept. of Computer Science and Engineering, Nanjing University of Aeronautics and Astronautics, Nanjing 210016, China
{dqzhang,qimiaoguo}@nuaa.edu.cn, {grwu,dgshen}@med.unc.edu

Abstract. Patch-based label fusion methods have shown great potential in multi-atlas segmentation. It is crucial for patch-based labeling methods to determine appropriate graphs and corresponding weights to better link patches in the input image with those in atlas images. Currently, two independent steps are performed, i.e., first constructing graphs based on the fixed image neighborhood and then computing weights based on the heat kernel for all patches in the neighborhood. In this paper, we first show that many existing label fusion methods can be unified into a graph-based framework, and then propose a novel method for simultaneously deriving both graph adjacency structure and graph weights based on the sparse representation, to perform multi-atlas segmentation. Our motivation is that each patch in the input image can be reconstructed by the sparse linear superposition of patches in the atlas images, and the reconstruction coefficients can be used to deduce both graph structure and weights simultaneously. Experimental results on segmenting brain anatomical structures from magnetic resonance images (MRI) show that our proposed method achieves significant improvements over previous patch-based methods, as well as other conventional label fusion methods.

1 Introduction

Automatic and accurate image segmentation is a critical step for many clinical studies in computational anatomy, including pathology detection and brain parcellation, etc. Recently, multi-atlas based segmentation methods have shown great success in segmenting brain into anatomical structures. There are two major steps in multi-atlas based segmentation, i.e., image registration and label fusion. In the first step, it registers each atlas image to the input image and further warps the corresponding label map by following the same estimated deformation field. Then, in the second step, it combines the multiple propagated labels from different atlases to obtain the final labels of the input image by some heuristics. We focus on label fusion in this study.

A number of label fusion strategies have been proposed for multi-atlas based segmentation. For example, majority voting (MV) and weighted MV are widely used in medical image segmentation. In [1], various weighting strategies were categorized into two groups, i.e., global weighted voting and local weighted voting, and it was

P.-T. Yap et al. (Eds.): MBIA 2012, LNCS 7509, pp. 94–102, 2012.

shown that the local weighted voting method outperforms the global one when segmenting high-contrast brain structures. Recently, besides weighting, some more advanced learning techniques have also been developed for further improving performance of label fusion. In [2], a probabilistic label fusion method was proposed to explicitly model the relationship between the atlas and the input image. In [3], a regression-based label fusion method was proposed to use the correlations between results from different atlases. In [4], a labeling confidence was estimated based on forward and backward k-nearest neighbor search to guide the subsequent sequential label fusion.

On the other hand, most of the above label fusion methods only consider one candidate on each atlas image when labeling a voxel in the input image, under the implicit assumption that the input and atlas images should be well registered, i.e., through a non-rigid registration. More recently, inspired by the success of non-local strategy and patch-based method in imaging applications, e.g., image denoising [5], patch-based label fusion that does not require any non-rigid registration was independently proposed in [6] and [7], respectively. Here, the main idea is to allow multiple candidates (usually in the neighborhood) on each atlas image and to aggregate them based on non-local mean. It is crucial for the patch-based label fusion method to determine the appropriate graphs and corresponding weights to better link patches in the input image with patches in the atlas images. At present, the graph construction processes are generally divided into two independent steps, i.e., first manually constructing graphs based on the fixed image neighborhood and then computing weights based on the heat kernel for all patches in the neighborhood.

In this paper, we show that many existing label fusion methods, including the patch-based method, can be unified into a graph-based framework, with differences in different definitions on the graph adjacency structure and graph weights. Furthermore, we propose a novel graph construction method for patch-based label fusion by using sparse representation. Here, our motivation is that each patch in the input image can be reconstructed by the sparse linear superposition of patches in the atlas images. Then, the reconstruction coefficients, which can be gotten by solving an l_1-norm regularized linear regression problem, will be used to deduce the graph adjacency structure and the graph weights simultaneously. It is worth noting that due to the use of l_1-norm, a lot of coefficients will have zero values, i.e., sparse patches are selected for subsequent label fusion. To the best of our knowledge, no previous works have used sparse representation for multi-atlas label fusion, although it has achieved great successes on a number of other applications such as face recognition [8]. We will apply our proposed sparse patch-based label fusion method for segmenting brain anatomical structures from magnetic resonance images (MRI).

2 Graph-Based Framework for Label Fusion

We follow the notations in [7]. Let A be an anatomy textbook (or atlas) containing a set of atlas images and label maps, denoted as $\{(I_i, L_i), i = 1, \dots, n\}$. Given an input image I, we construct a weighted graph G_i between voxels $\boldsymbol{x}$ of input image I and

voxels $\boldsymbol{y}$ of each atlas image I_i, along with corresponding weight $w_i(\boldsymbol{x}, \boldsymbol{y})$, for $(\boldsymbol{x}, \boldsymbol{y}) \in \Omega^2$ where Ω denotes the image domain. Once we have obtained the graph weights, we can perform label fusion for voxel $\boldsymbol{x}$ of input image I as below:

$$L(\boldsymbol{x}) = \frac{\sum_{i=1}^{n} \sum_{\boldsymbol{y} \in \Omega} w_i(\boldsymbol{x}, \boldsymbol{y}) L_i(\boldsymbol{y})}{\sum_{i=1}^{n} \sum_{\boldsymbol{y} \in \Omega} w_i(\boldsymbol{x}, \boldsymbol{y})}, \forall \boldsymbol{x} \in \Omega. \tag{1}$$

We will show that many existing label fusion methods, e.g., majority voting (MV), and patch-based method (PBM), can be derived from Eq. 1, by using different definitions for graph weights.

1) Majority Voting (MV): We can get the label fusion rule for majority voting from Eq. 1, if we define the following graph weights:

$$w_i^{MV}(\boldsymbol{x}, \boldsymbol{y}) = \begin{cases} 1, & \forall\, (\boldsymbol{x}, \boldsymbol{y}) \in \Omega^2 \text{ and } \boldsymbol{x} = \boldsymbol{y}; \\ 0, & otherwise. \end{cases} \tag{2}$$

2) Patch-Based Method (PBM): We can get the label fusion rule for patch-based method from Eq. 1, if we define the following graph weights:

$$w_i^{PBM}(\boldsymbol{x}, \boldsymbol{y}) = \begin{cases} g\left(\dfrac{\sum_{\boldsymbol{x}' \in P_I(\boldsymbol{x}),\, \boldsymbol{y}' \in P_{I_i}(\boldsymbol{y})} \left(I(\boldsymbol{x}') - I_i(\boldsymbol{y}') \right)^2}{h_p} \right), \\ \qquad \forall\, (\boldsymbol{x}, \boldsymbol{y}) \in \Omega^2 \text{ and } \boldsymbol{y} \in N_i(\boldsymbol{x}); \\ 0, \qquad otherwise. \end{cases} \tag{3}$$

Where $N_i(\boldsymbol{x})$ denotes the neighborhood of voxels $\boldsymbol{x}$ in the atlas image I_i; $P_I(\boldsymbol{x})$ and $P_{I_i}(\boldsymbol{y})$ denote the patch centered at voxel $\boldsymbol{x}$ in the input image I and the patch centered at voxel $\boldsymbol{y}$ in the atlas image I_i; g is a smoothing kernel function (in PBM, the heat kernel $g(z) = e^{-z}$ was used), and h_p is the kernel width.

It is worth noting that the constructed graph and its corresponding weights essentially determine the label fusion rule of a certain method, which inspires us to seek other graph construction ways in order for devising new label fusion methods.

3 Sparse Patch-Based Method

Inspired from the discriminating power of sparse representation which has been validated on other tasks such as face recognition, we propose to reconstruct each patch in the input image by the sparse linear superposition of patches in the atlas images. Then, the reconstruction coefficients of those patches are used to deduce the graph adjacency structure and the graph weights simultaneously.

Following the notations in Section 2, given a set of atlas images and label maps $\{(I_i, L_i), i = 1, \ldots, n\}$, for each voxel $\boldsymbol{x}$ of input image I and each voxel $\boldsymbol{y} \in N_i(\boldsymbol{x})$ in atlas image I_i, we want to automatically optimize the weight $w_i(\boldsymbol{x}, \boldsymbol{y})$ based on sparse representation.

Denote $A_{\boldsymbol{y}}^{i} \triangleq col\left(\left\{I_i(\boldsymbol{y}')|\ \boldsymbol{y}' \epsilon P_{I_i}(\boldsymbol{y})\right\}\right)$ as a patch vector corresponding to $P_{I_i}(\boldsymbol{y})$, and $b_{\boldsymbol{x}} \triangleq col(\{I(\boldsymbol{x}')|\boldsymbol{x}' \epsilon P_I(\boldsymbol{x})\})$ as a patch vector corresponding to $P_I(\boldsymbol{x})$, where col is an operator which aligns all elements in a set into a column vector. Then, for each $\boldsymbol{x} \in \Omega$, we optimize the following objective function to get the graph weights $w_i(\boldsymbol{x}, \boldsymbol{y})$, for each $\boldsymbol{y} \in N_i(\boldsymbol{x})$ and $i = 1, \ldots, n$:

$$\min_{\{w_i(\boldsymbol{x},\boldsymbol{y})\}} \frac{1}{2}\left\|\sum_{i=1}^{n}\sum\nolimits_{\boldsymbol{y}\in N_i(\boldsymbol{x})} A_{\boldsymbol{y}}^{i} w_i(\boldsymbol{x},\boldsymbol{y}) - b_{\boldsymbol{x}}\right\|_2^2 + \lambda\sum_{i=1}^{n}\sum\nolimits_{\boldsymbol{y}\in N_i(\boldsymbol{x})} |w_i(\boldsymbol{x},\boldsymbol{y})|. \tag{4}$$

The intuition of Eq. 4 is to sparsely reconstruct the patch vector $b_{\boldsymbol{x}}$ in the input image using by patch vectors $A_{\boldsymbol{y}}^{i}$ in the atlas images. Specifically, the function of the first term of Eq. 4 is to minimize the reconstruction error, while the second term, which is equivalent to the l_1-norm, requires a sparse solution. The regularization parameter λ balances the relative contributions of the two terms and also controls the 'sparsity' of the linear model.

Furthermore, Eq. 4 can be simplified into the following equivalent form:

$$\min_{w_{\boldsymbol{x}}} \frac{1}{2}\|Aw_{\boldsymbol{x}} - b_{\boldsymbol{x}}\|_2^2 + \lambda\|w_{\boldsymbol{x}}\|_1. \tag{5}$$

where $w_{\boldsymbol{x}} \triangleq col(\{w_i(\boldsymbol{x},\boldsymbol{y})|\boldsymbol{y}\epsilon N_i(\boldsymbol{x}), i\epsilon\{1,2,\ldots,n\}\})$, and A is defined as $A \triangleq row\left(\left\{A_{\boldsymbol{y}}^{i}|\boldsymbol{y}\epsilon N_i(\boldsymbol{x}), i\epsilon\{1,2,\ldots,n\}\right\}\right)$ where row is an operator which aligns all elements in a set into a (block) row vector. Note that since $A_{\boldsymbol{y}}^{i}$ is a column vector, aligning $A_{\boldsymbol{y}}^{i}$ with the operator row will lead to a matrix A. Eq. 5 is a standard l_1-norm optimization problem (with the same form as the Lasso method [9]), which can be efficiently solved by a number of existing optimization software [10]. It is worth nothing that a non-negative constraint is required on weight $w_i(\boldsymbol{x}, \boldsymbol{y})$ in this study.

Once we have obtained the graph weights $w_i(\boldsymbol{x}, \boldsymbol{y})$ by solving Eq. 4 (or 5), we can perform the label fusion process using Eq. 1, to get the propagated label $L(\boldsymbol{x})$ for each voxel $\boldsymbol{x}$ in the input image I. Note that generally $L(\boldsymbol{x})$ takes a continuous value, e.g., a value between 0 and 1 for binary segmentation (i.e., two classes with labels 0 and 1). We focus on binary segmentation in this paper. To get the final label, we use the following equation, for each $\boldsymbol{x} \in \Omega$:

$$\Gamma(\boldsymbol{x}) = \begin{cases} 1, & L(\boldsymbol{x}) \geq 0.5; \\ 0, & otherwise. \end{cases} \tag{6}$$

4 Experiments

In this section, we evaluate the performance of our proposed sparse-patch based method (SPBM) on segmenting brain anatomical structures, i.e., region-of-interests (ROIs), from the MR brain images of NA0-NIREP database [11]. We compare our

method with other label fusion methods, including majority voting (MV), local weighted voting (LWV), STAPLE [12], and patch-based method (PBM) [6, 7].

4.1 Dataset and Experimental Settings

The NA0-NIREP dataset [11] used for evaluation consists of 16 annotated MR images, including 8 normal adult males and 8 females. The 16 MR images have been manually segmented into 32 ROIs. The MR images were obtained in a General Electric Signa scanner operating at 1.5 Tesla, using the following protocol: SPGR/50, TR 24, TE 7, NEX 1 matrix 256 × 192, FOV 24 cm.

For evaluation, a random sampling of approximate two thirds (i.e., 11) of the total 16 MR images in NA0-NIREP dataset are used as atlases and the rest 5 images are used as input images. 10 independent runs are performed and the averaged results are reported. We segment the input images using different multi-atlas label fusion methods, and then measure their performances using the Dice overlap, defined as $Dice(S_a, S_b) = 2|S_a \cap S_b|/(|S_a| + |S_b|)$, where the symbol $\cap$ denotes the overlapping voxels between the two segmentations, and $|S|$ denotes the number of voxels of the corresponding segmentation.

In LWV, PBM and our SPBM, there is a common parameter, i.e., the size of neighborhood region N or N_i. In our experiments, a neighborhood size of $5 \times 5 \times 5$ is used for all the three algorithms. On the other hand, both PBM and SPBM have another common parameter, i.e., the size of patches P_I and P_{I_i}. In our experiments, following [7], a small size of patches (i.e., $3\times 3 \times 3$ voxels) is used for both methods. Besides, both LWV and PBM adopt the heat kernel, and the corresponding kernel widths h_w and h_p are estimated based on the minimal distances between region $N(\boldsymbol{x})$ and regions $N_i(\boldsymbol{x})$ $(i = 1, \ldots, n)$, and between patch $P_I(\boldsymbol{x})$ and patches $P_{I_i}(\boldsymbol{y})$ $(\boldsymbol{y} \in N_i(\boldsymbol{x})$ and $i = 1, \ldots, n)$, for LWV and PBM, respectively. Similar strategy has also been used in [6]. We use the SLEP software [10] to solve the l_1-norm optimization in Eqs. 4-5, and the parameter λ is set as $0.01\lambda_{max}$, where $\lambda_{max}$is automatically computed by the program, and it denotes the maximal value of λ, above which shall lead to the zero solution.

4.2 Comparison on Segmentation Results

We first compare the segmentation results of different multi-atlas label fusion algorithms on NA0-NIREP. It is worth noting that we focus on binary segmentation, and thus in our experiments for each method we perform 32 independent binary segmentations corresponding to 32 ROIs. At each segmentation corresponding to a certain ROI, we set the (ground truth) label of a voxel as 1 if it belongs to the ROI, and 0 otherwise. Table 1 gives the segmentation results measured by Dice overlap using five different methods on different ROIs of brain.

As can be seen from Table 1, our SPBM method achieves the best performance among all five methods on segmenting nearly all brain ROIs. Specifically, it outperforms MV, STAPLE and LWV on all 32 brain ROIs, and PBM on 30 of 32 ROIs. We

also compute the averaged Dice overlap of different methods across 32 ROIs, and the results are 53.6%, 56.9%, 68.0%, 73.6% and 75.6%, for MV, STAPLE, LWV, PBM and SPBM, respectively. These results show that in average our method achieves 2.0 percent improvement over the conventional PBM, and both patch-based method (including PBM and SPBM) significantly outperform the other label fusion methods, which validate the advantage of using the one-to-many correspondence strategy in the patch-based methods over the use of conventional one-to-one correspondence strategy in the non-patch-based methods. Tables 1 also indicates that by considering the image information through local weighting, LWV achieves a big improvement than MV and STAPLE which do not use image information in label fusion, but is still much inferior to the patch-based methods (including both PBM and SPBM) due to the use of one-to-one correspondence.

Tables 1 also show that there exist great differences on performances of five methods across different ROIs. For example, the differences between results on L insula gyrus (one of the best segmented ROI) and those on L postcentral gyrus (one of the worst segmented ROI) are 30.1%, 22.6%, 22.4%, 21.5% and 21.8%, for MV, STAPLE, LWV, PBM and SPBM, respectively. Clearly, these results show that segmenting the latter brain structure is more challenging than segmenting the former brain structure.

Finally, in Fig. 1 we visually plot the segmentation results of five methods on segmenting *L insula gyrus*. For comparison, we also show the original image and corresponding ground truth. As can be seen from Fig. 1, our SPBM method achieves the best visual quality of segmentation results among the five methods.

4.3 Comparison on Graph Weights

In this experiment, we compare the graph weights gotten by different graph-based label fusion methods. Because the graphs constructed by MV and LWV are very simple (i.e., one-to-one correspondence), we only compare the graph weights constructed from two patch-based methods, i.e., PBM and SPBM. To this end, in Fig. 2, we plot the graph weights constructed by both PBM and SPBM methods, for a 'positive' input patch (i.e., centered voxel has a label of 1) on segmenting *L insula gyrus*. In Fig. 2, red lines denote the weights for 'positive' patches from 11 atlas images, while blue lines denote those for 'negative' patches (i.e., centered voxel has a label of 0), and the green lines denote the separation lines between different atlas images.

As can be seen from Fig. 2, PBM obtains a dense graph even after the pre-selection of patches. This is because in some (especially smooth) regions, there may exist a number of adjacent patches that have high similarities to the input patch, and thus it is difficult to get a sparse graph through threshold on similarities. In contrast, SPBM get a very sparse graph by explicitly imposing the 'sparsity' constraint into the objective function. On the other hand, Fig. 2(b) indicates that SPBM also achieves a discriminating power by using sparse representation, which has been validated on other tasks, e.g., face recognition [8]. Specifically, Fig. 2(b) shows that the graph weights of SPBM for the current input patch are dominated by 'positive' patches. This

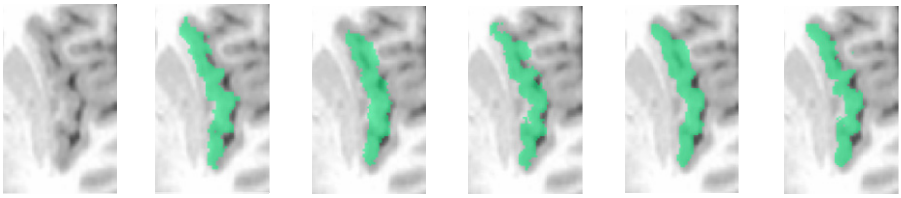

Fig. 1. Visual views on segmentation results of five different label fusion algorithms on *L insula gyrus*

spontaneously-emerged discriminating power is very helpful for subsequent labeling, as shown in above results.

Table 1. Segmentation results on different ROIs of brain, measured by Dice overlap (%) using five different algorithms. Results are presented in (left hemisphere - right hemisphere).

Brain ROIs	MV	STAPLE	LWV	PBM	SPBM
Occipital lobe	52.1 - 55.2	58.5 - 61.7	69.3 - 74.9	73.6 - 79.6	**77.3 - 81.2**
Cingulated gyrus	59.4 - 61.0	62.5 - 64.7	70.0 - 71.5	72.9 - 75.5	**75.7 - 77.0**
Insula gyrus	65.6 - 69.0	64.0 - 67.1	74.3 - 76.3	77.8 - 80.3	**82.3 - 84.0**
Temporal pole	59.7 - 65.4	62.0 - 65.9	73.2 - 78.0	76.6 - 81.6	**78.6 - 81.9**
Superior temporal gyrus	48.0 - 43.8	52.5 - 52.9	62.6 - 64.3	67.2 - 70.0	**73.1 - 72.2**
Infero emporal region	59.9 - 59.8	60.5 - 59.8	73.3 - 76.2	76.8 - 80.7	**80.1 - 80.9**
Parahippocampal gyrus	64.6 - 70.4	61.2 - 63.7	73.4 - 77.5	77.3 - 82.1	**79.9 - 82.5**
Frontal pole	62.8 - 62.5	63.6 - 58.9	75.3 - 74.7	**81.0 - 81.6**	80.6 - 79.6
Superior frontal gyrus	52.6 - 53.0	55.7 - 55.1	69.7 - 69.5	74.4 - 74.6	**76.0 - 75.3**
Middle frontal gyrus	53.4 - 48.7	55.3 - 55.1	68.2 - 65.9	72.7 - 70.7	**74.9 - 71.7**
Inferior gyrus	47.2 - 44.8	49.7 - 53.6	64.1 - 62.8	69.9 - 67.8	**71.7 - 69.5**
Orbital frontal gyrus	61.9 - 61.6	62.4 - 61.6	74.6 - 73.3	79.9 - 77.4	**81.3 - 78.1**
Precentral gyrus	39.1 - 37.7	52.1 - 48.4	61.9 - 62.3	66.1 - 67.6	**70.0 - 69.6**
Superior parietal lobule	42.5 - 46.2	49.8 - 51.9	62.4 - 64.2	67.3 - 69.4	**69.8 - 72.2**
Inferior parietal lobule	49.7 - 50.6	53.4 - 54.8	64.3 - 67.8	68.9 - 73.0	**72.4 - 73.3**
Postcentral gyrus	35.5 - 30.5	41.4 - 41.8	51.9 - 55.6	56.3 - 63.4	**60.5 - 65.7**

5 Conclusion

In this paper we proposed a novel sparse patch-based method (SPBM) for multi-atlas label fusion, based on sparse representation. We also show that the patch-based methods as well as many existing methods can be unified into a graph-based label fusion framework, with different ways for graph construction. Different from the conventional patch-based method (PBM) which constructs a graph and computes weights based on similarities between patches in the input image and the atlas images, our method automatically obtains the graph and weights by solving l_1 optimization. Our experiments on segmenting brain anatomical structures on NA0-NIREP dataset show

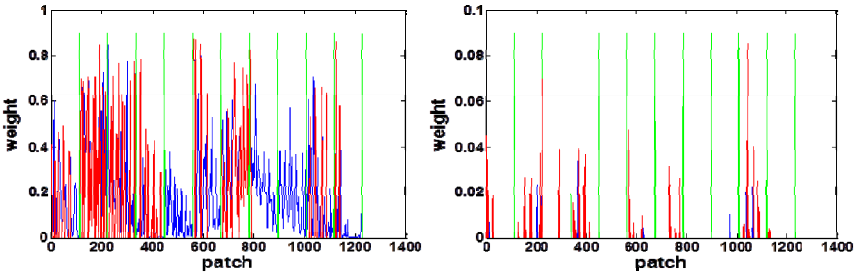

(a) PBM (b) SPBM

Fig. 2. Comparison of the graph weights constructed by PBM (a) and SPBM (b)

that SPBM achieves up to 5.9 (and 2.0 in average) percent improvements over PBM, and significantly outperforms other label fusion methods.

References

1. Artaechevarria, X., Munoz-Barrutia, A., Ortiz-de-Solorzano, C.: Combination strategies in multi-atlas image segmentation: application to brain MR data. IEEE Transactions on Medical Imaging 28, 1266–1277 (2009)
2. Sabuncu, M.R., Yeo, B.T., Van Leemput, K., Fischl, B., Golland, P.: A generative model for image segmentation based on label fusion. IEEE Transactions on Medical Imaging 29, 1714–1729 (2010)
3. Wang, H., Suh, J.W., Das, S.R., Pluta, J., Altinay, M., Yushkevich, P.: Regression-based label fusion for multi-atlas segmentation. In: CVPR, pp. 1113–1120 (2011)
4. Zhang, D., Wu, G., Jia, H., Shen, D.: Confidence-Guided Sequential Label Fusion for Multi-atlas Based Segmentation. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part III. LNCS, vol. 6893, pp. 643–650. Springer, Heidelberg (2011)
5. Coupe, P., Yger, P., Barillot, C.: Fast non local means denoising for 3D MR images. Med. Image Comput. Comput. Assist. Interv. 9, 33–40 (2006)
6. Coupe, P., Manjon, J.V., Fonov, V., Pruessner, J., Robles, M., Collins, D.L.: Patch-based segmentation using expert priors: application to hippocampus and ventricle segmentation. NeuroImage 54, 940–954 (2011)
7. Rousseau, F., Habas, P.A., Studholme, C.: A supervised patch-based approach for human brain labeling. IEEE Transactions on Medical Imaging 30, 1852–1862 (2011)
8. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. IEEE Transactions on Pattern Analysis and Machine Intelligence 31, 210–227 (2009)
9. Tibshirani, R.: Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B 58, 267–288 (1996)
10. Liu, J., Ji, S., Ye, J.: SLEP: Sparse learning with efficient projections. Arizona State University (2009)

11. Christensen, G.E., Geng, X., Kuhl, J.G., Bruss, J., Grabowski, T.J., Pirwani, I.A., Vannier, M.W., Allen, J.S., Damasio, H.: Introduction to the Non-rigid Image Registration Evaluation Project (NIREP). In: Pluim, J.P.W., Likar, B., Gerritsen, F.A. (eds.) WBIR 2006. LNCS, vol. 4057, pp. 128–135. Springer, Heidelberg (2006)
12. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. IEEE Transactions on Medical Imaging 23, 903–921 (2004)