



HHS Public Access

Author manuscript

Med Image Anal. Author manuscript; available in PMC 2018 February 01.

Published in final edited form as:

Med Image Anal. 2017 February ; 36: 162–171. doi:10.1016/j.media.2016.11.005.

Progressive Multi-Atlas Label Fusion by Dictionary Evolution

Yantao Song^{a,b}, Guorong Wu^b, Khosro Bahrami^b, Quansen Sun^a, and Dinggang Shen^{b,c,*}

^aSchool of Computer Science & Technology, Nanjing University of Science and Technology, Nanjing, Jiangsu 210094 China

^bDepartment of Radiology and BRIC, University of North Carolina at Chapel Hill, NC 27599 USA

^cDepartment of Brain and Cognitive Engineering, Korea University, Seoul 02841, Republic of Korea

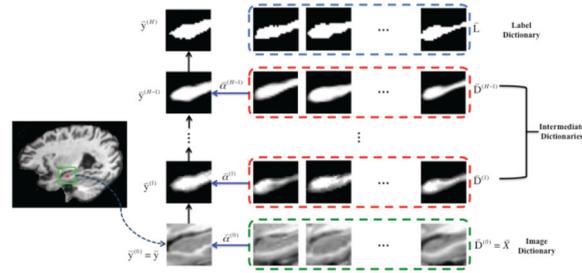
Abstract

Accurate segmentation of anatomical structures in medical images is important in recent imaging based studies. In the past years, multi-atlas patch-based label fusion methods have achieved a great success in medical image segmentation. In these methods, the appearance of each input image patch is first represented by an *atlas patch dictionary* (in the image domain), and then the latent label of the input image patch is predicted by applying the estimated representation coefficients to the corresponding anatomical labels of the atlas patches in the *atlas label dictionary* (in the label domain). However, due to the generally large gap between the patch appearance in the *image domain* and the patch structure in the *label domain*, the estimated (patch) representation coefficients from the image domain may not be optimal for the final label fusion, thus reducing the labeling accuracy. To address this issue, we propose a novel label fusion framework to seek for the suitable label fusion weights by progressively constructing a dynamic dictionary in a layer-by-layer manner, where the intermediate dictionaries act as a sequence of guidance to steer the transition of (patch) representation coefficients from the image domain to the label domain. Our proposed multi-layer label fusion framework is flexible enough to be applied to the existing labeling methods for improving their label fusion performance, i.e., by extending their single-layer *static* dictionary to the multi-layer *dynamic* dictionary. The experimental results show that our proposed progressive label fusion method achieves more accurate hippocampal segmentation results for the ADNI dataset, compared to the counterpart methods using only the single-layer *static* dictionary.

Graphical Abstract

*Corresponding author. dgshen@med.unc.edu (D. Shen).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Keywords

Brain MRI; sparse representation; hippocampus; label fusion; multi-atlas

1. Introduction

Magnetic resonance imaging (MRI) is an advanced medical imaging technique, which plays an essential role in neuroscience research and clinical studies. However, due to the large amount of MRI data produced every day, it is time-consuming and expensive to process medical images manually. Therefore, automated and accurate segmentation is in high demand in existing imaging-based studies, in order to *either* discover group differences between individual subjects *or* quantify subtle changes over time. For instance, the hippocampus is known as an important structure related to Alzheimer's disease, epilepsy, and schizophrenia (Devanand et al., 2007; Dickerson et al., 2001; Holland et al., 2012; Van Leemput et al., 2009). Therefore, automated and accurate segmentation of the hippocampus is critical.

However, since anatomical structures (i.e., hippocampus) vary significantly across individuals, the prior knowledge of shape and appearance learned from a certain template is often not sufficient for guiding the segmentation of target anatomical structures. To this end, multi-atlas based segmentation methods have been recently developed and achieved great success by letting the target labels on the target image follow the consensus of labels of multiple atlases with similar local image appearance. Generally, with more atlases, higher segmentation accuracy can be achieved by reducing the variations between the target and atlas images.

To do the segmentation, followed by registering atlas images to the target image, the latent anatomical label on each target image point can be determined by a certain label fusion strategy, such as majority voting (MV) Heckemann et al. (2006); Rohlfing et al. (2005). Majority voting is a classical label fusion method, which simply chooses the label with the highest vote as the final label. To improve the labeling accuracy, local weighted voting (LWV) was also proposed by replacing the *hard* voting (considering only the label information) with *soft* voting which is proportional to the patch-wise appearance similarity Sabuncu et al. (2010).

Apparently, the above *point-wise* label fusion strategies are highly dependent on the accuracy of image registration. To address the potential issue of inaccurate registration,

many *patch-based* label fusion methods have been proposed in recent years Artachevarria et al. (2009); Song et al. (2015); Wang et al. (2011); Yan et al. (2013); Zhang et al. (2011). In these methods, the main assumption is that, if two image patches have similar appearance, they should bear the same anatomical label. The typical patch-based label fusion methods include *nonlocal patch-based labeling* (NPBL) Coupé et al. (2011); Rousseau et al. (2011) and *sparse patch-based labeling* (SPBL) Tong et al. (2013); Zhang et al. (2012a,c). Note that all patch-based label fusion methods collect candidate atlas patches in a search neighborhood across all registered atlas images. The weights used for label fusion in NPBL are proportional to the decayed patch-wise similarities penalized by the exponential function. Inspired by the discriminative power of sparse representation Tibshirani (2011); Zhang et al. (2012b,a), the SPBL method has been proposed to introduce sparsity into the optimization of the weighting vector at each image point. Since the sparsity constraint enforces many zero elements in the weighting vector, the SPBL method can reduce the risk of taking incorrect or ambiguous patches and finally use only a small number of well-matched patches for labeling. More advanced methods can be found in Wang et al. (2013); Wu et al. (2014), where the pairwise dependency between atlas patches is further modeled to avoid the repeating label fusion error by similar atlas patches.

In all of the above state-of-the-art methods, the weights are exclusively optimized in terms of patch-wise image appearance. The computed weighting vector is regarded as an *appearance representation profile* and then directly used to determine the (binary) labels for the target image. Despite its simplicity and effectiveness, there is no evidence showing that such weights are domain-invariant, i.e., the optimized weights derived from the best image patch presentation may be not necessarily optimal for label fusion. Fig. 1 demonstrates the significant gap of representation profiles estimated in the *image domain* using appearance information (left) and in the *label domain* using label information (which is assumed to be known for the unseen target image). It is clear that there is no guarantee for the current state-of-the-art label fusion methods to achieve the optimal labeling results by directly applying the *appearance-based* representation profile for label fusion.

To address this issue, we propose a novel label propagation framework to progressively convert the *representation profile* from the image domain to the *optimal weighting vector* (for label fusion) in the label domain by constructing a set of intermediate dictionaries to bridge the image domain and the label domain. Such intermediate dictionaries provide a sequence of guidance to steer the estimation of the *appearance representation profile* to the *optimal weighting vector* for label fusion.

Specifically, in the training stage for each target image patch, the initial-layer dictionary consists of the original atlas image patches (in the image domain), similar to the most of traditional label fusion methods. Since each atlas image patch has its corresponding label patch, it is straightforward to build the label patch dictionary (in the label domain) by arranging the corresponding label patches with the same order as the original atlas image patches in the initial-layer dictionary. To remedy the large transitions from the image domain to the label domain, we first apply a label fusion technique (e.g., NPBL or SPBL) to obtain the *representation profile* for each atlas image patch in the initial-layer dictionary, while regarding all other instances in the initial-layer dictionary as the atlas image patches. Then,

we compute a *label probability patch* by applying the obtained *representation profile* to the respective atlas label patches. By repeating the above leave-one-out label fusion procedure to all the patches in the initial-layer dictionary, we can construct the first-layer intermediate dictionary. Similarly, we can construct the subsequent intermediate dictionaries, as shown in Fig. 2. In the end, we can construct a sequence of intermediate dictionaries, where the label probability patches become sharper and sharper, close to the binary shape of the corresponding atlas label patches.

In the testing stage, given the learned multi-layer dictionary at each target image location, the final weights for voting the label are also estimated in a progressive way. Specifically, starting from the initial layer, we gradually refine the label fusion weights by alternating the following two steps. *First*, we compute the *representation profile* of the target image patch by using the patch dictionary in the current layer. *Second*, we refine the label probability map within the target image patch by applying the latest *representation profile* to the binary atlas label patches, and then use the obtained new probability patch as the new target image in the next layer of the label estimation. In this way, we can obtain more and more accurate weights to determine the anatomical label for the original target image, with the guidance of the intermediate dictionary at each layer.

The contributions of our proposed method include: (1) Since we harness the multi-layer dictionary to remedy the gap between patch appearances and anatomical labels, our label fusion essentially seeks the best label fusion weights, instead of just patch-wise representation; (2) The sequence of built intermediate dictionaries allows us using *not only* appearance features *but also* structural context information Tu and Bai (2010) to significantly improve the robustness in patch representation; (3) Our proposed progressive patch representation by a multi-layer dictionary is general enough to be integrated with many conventional patch-based segmentation methods for improving their performances. Our proposed method has been evaluated in the segmentation of the hippocampus from elderly brain MR images in the ADNI dataset. More accurate segmentation results have been achieved, compared to the state-of-the-art methods, i.e., NPBL and SPBL.

The remainder of the paper is organized as follows. In Section 2, we present a short introduction of the existing hippocampus labeling methods in the literature, followed by a detailed description of our progressive multilayer label fusion method. Section 3 shows the results of our proposed method in hippocampus segmentation. Then, in Section 4, we give the discussion of the proposed method along with the future work. Finally, we conclude our work in Section 5.

2. Methods

2.1. Overview

In general, multi-atlas patch-based segmentation aims to determine the label of each point in the target image T by using a set of N registered atlas images I_s and label images $L_s (s = 1, \dots, N)$. For each voxel $v \in \Omega_T$ in the target image, $P_T(v)$ is the target patch centered at point v . The candidate atlas patches are recruited within a small neighborhood n_v across all N registered atlas images. We assume there are K candidate atlas patches $P_s(u)$ centered at

point $u(u \in n_v)$ after patch pre-selection Coupé et al. (2011). Note that patch pre-selection is often used in a searching neighborhood to speed up the label fusion and also preclude the most dissimilar patches. Generally, patch pre-selection provides a balance between the segmentation accuracy and computational time. In our implementation, we follow the criteria of pre-selection in Coupé et al. (2011) by computing the well-known structural similarity measure (SSIM) Wang et al. (2004).

Then, we vectorize each atlas patch $P_s(u)$ into an M -length column vector $\vec{x}_s(u)$ to form a patch dictionary $\mathbf{X} = [\vec{x}_s(u)]_{s=1, \dots, N, u \in n_v}$. Since each atlas patch has label information, it is straightforward to construct a corresponding label patch dictionary $\mathbf{L} = [\vec{l}_s(u)]_{s=1, \dots, N, u \in n_v}$, where $\vec{l}_s(u)$ is a binary vector of $\{0, 1\}^M$ representing the particular label (with '0' denoting background and '1' denoting the hippocampus in this paper) which is associated with $\vec{x}_s(u)$ of atlas patch $P_s(u)$. So the latent label for point v can be estimated through the interaction between the target patch $P_T(v)$ and all candidate atlas patches $P_s(u)$. Given the weight $\alpha_s(u, v)$ for the pair of $P_T(v)$ and $P_s(u)$, the predicted label vector $\vec{\phi}(v)$ for the target point v can be estimated via

$$\vec{\phi}(v) = \frac{\sum_{s=1}^N \sum_{u \in n_v} \alpha_s(u, v) \vec{l}_s(u)}{\sum_{s=1}^N \sum_{u \in n_v} \alpha_s(u, v)} \quad (1)$$

where $\vec{\phi}(v)$ is a vector of label probability, with each element in $\vec{\phi}(v)$ associated with one point in the target patch $P_T(v)$. Then, it is straightforward to obtain hippocampus segmentation within the target patch $P_T(v)$ by binarizing the probability vector $\vec{\phi}(v)$ by a threshold of 0.5.

In the following, we first introduce the calculation of weight $\alpha_s(u, v)$ by the conventional patch-wise labeling methods in Section 2.2. Then, we present our multi-layer framework in Section 2.3, including the construction of the multi-layer dictionary.

2.2. Conventional patch-based labeling

The principle of the conventional NPBL method is originated from the non-local strategy, which was widely used in computer vision, i.e., for image denoising Buades et al. (2005) and super-resolution Protter et al. (2009). The applications of NPBL in medical images can also be found in Awate and Whitaker (2006); Manjón et al. (2010). In the NPBL method, each candidate atlas patch $P_s(u)$ contributes to the final label fusion. And the weight of each candidate patch is calculated by a patch-wise similarity between the target patch $P_T(v)$ and the candidate patch $P_s(u)$, defined as

$$\alpha_s(u, v) = \exp(-\|P_T(v) - P_s(u)\|^2 / 2\sigma^2) \quad (2)$$

where σ controls the strength of exponential penalty on patch-wise difference.

In contrast to the NPBL method, the SPBL method assumes that the representation should be sparse, which means the optimal representation should use a smaller number of atlas

patches. Given the patch dictionary \mathbf{X} , sparse representation aims to find a sparse linear combination of patches in the dictionary \mathbf{X} for the best representation of the target patch P_T (\mathbf{v}). For simplicity, we denote the target patch $P_T(\mathbf{v})$ by \vec{y} and consider the weights associated with the patch dictionary \mathbf{X} forming a weighting vector $\vec{\alpha}$. Thus, the estimation of weighting vector $\vec{\alpha}$ can be formulated as a problem of minimizing the reconstruction error in the following energy function:

$$\operatorname{argmin}_{\vec{\alpha}} \|\vec{y} - \mathbf{X}\vec{\alpha}\|_2^2 + \lambda \|\vec{\alpha}\|_1 \quad (3)$$

In this objective function, the first term measures the reconstruction error, and the second term uses the l_1 -norm constraint on the weighting vector $\vec{\alpha}$. The regularization parameter λ controls the strength of sparsity on $\vec{\alpha}$.

2.3. Proposed method

2.3.1. Multi-layer labeling—The limitation of conventional multi-atlas patch-based segmentation methods is that *they assume high correlation between original image patches and label patches*. However, such an assumption may not be held in medical images which often have one-to-many mapping between appearance patches and label patches. To mitigate the inconsistency between the appearance patches (continuous) and the label patches (binary), we hypothesize a pathway of patch evolution that can smoothly transit from the image appearance to the corresponding labels. Our solution is to progressively construct a dynamic dictionary in a layer-by-layer manner, where a sequence of intermediate dictionaries can be constructed to gradually guide the transition from the *appearance representation profile* in the image domain to the *optimal weighting vector* (label fusion) in the label domain. As shown in both ends of Fig. 3, a set of atlas appearance patches (in left end) and their associated label patches (in right end) are extracted to form the image and label dictionaries, respectively. The constructed intermediate patch dictionaries (displayed in the middle of Fig. 3) describe the transition path, in terms of (from fuzzy to sharp) label probability maps. In this way, we can formulate the single-layer estimation of the *representation profile* to the multi-layer optimization of the *weighting vector* (for label fusion) by progressively refining the *weighting vector* with the evolution of the intermediate dictionaries from the image appearance to binary labels. The construction of the multi-layer dictionary and also the progressive label fusion are detailed below.

2.3.2. Multi-layer patch dictionary construction—In the training stage, to construct multi-layer patch dictionaries, we use the original intensity patch dictionary \mathbf{X} to form the initial layer $\mathbf{D}^{(0)} = \mathbf{X}$ as shown in the bottom of Fig. 2, i.e., $\mathbf{D}^{(0)} = [\vec{d}_k^{(0)}]$, where $\vec{d}_k^{(0)} = \vec{x}_k$. Note that we here use k to refer the k -th patch in the dictionary. From the first layer $\mathbf{D}^{(1)}$, we iteratively construct the intermediate dictionary $\mathbf{D}^{(h)} (h = 1, \dots, H-1)$ by alternating the following three steps.

- *First, starting from $h = 1$, for each instance $\vec{d}_k^{(h-1)}$ in the dictionary $\mathbf{D}^{(h-1)}$, we use all the other instances*

$\vec{d}_j^{(h-1)}$ ($j \neq k$) in $\mathbf{D}^{(h-1)}$ to represent the underlying instance $\vec{d}_k^{(h-1)}$. In such case, all the other instances in $\mathbf{D}^{(h-1)}$ form the instance-specific dictionary $\mathbf{B}_k^{(h)} = [\vec{d}_j^{(h-1)}]_{j=1, \dots, K; j \neq k}$, where $\mathbf{B}_k^{(h)}$ has $K-1$ column vectors. It should be mentioned that we can obtain the patch representation profile $\vec{\beta}_k^{(h)}$ for $\vec{d}_k^{(h-1)}$ via any fusion strategy, e.g., non-local mean in Eq.(2) or sparse representation technique in Eq.(3). Note that $\vec{\beta}_k^{(h)}$ is the column vector of length $K-1$.

Second, since each atom in $\mathbf{B}_k^{(h)}$ is associated with one label vector in \mathbf{L} , we can build a surrogate atlas-domain patch dictionary $\mathbf{L}_k = [\vec{l}_j]_{j=1, \dots, K; j \neq k}$ by arranging the atlas patches with the same order in $\mathbf{B}_k^{(h)}$. Then, we can compute the label patch $\vec{p}_k^{(h)}$ by using $\vec{p}_k^{(h)} = \mathbf{L}_k \vec{\beta}_k^{(h)}$.

Algorithm 1

Learning the intermediate dictionaries

Input: Intensity atlas patch dictionary $\mathbf{D}^{(0)} = [\vec{d}_k^{(0)}]$, label atlas patch dictionary $\mathbf{L} = [\vec{l}_k]$, patch number K , and layer number H .

Initialize: $h = 1$

While $h < H$

For the k -th instance **do**

1. $\mathbf{B}_k^{(h)} = \{ \vec{d}_1^{(h-1)}, \dots, \vec{d}_{k-1}^{(h-1)}, \vec{d}_{k+1}^{(h-1)}, \dots, \vec{d}_K^{(h-1)} \}$;

2. $\mathbf{L}_k = \{ \vec{l}_1, \dots, \vec{l}_{k-1}, \vec{l}_{k+1}, \dots, \vec{l}_K \}$;

3. Compute $\vec{\beta}_k^{(h)}$ for $\vec{d}_k^{(h-1)}$ by $\mathbf{B}_k^{(h)}$ according to Eq.(2) or Eq.(3);

4. $\vec{p}_k^{(h)} = \mathbf{L}_k \vec{\beta}_k^{(h)}$;

5. $\vec{d}_k^{(h)} = \vec{p}_k^{(h)}$.

End for

$$\mathbf{D}^{(h)} = [\vec{d}_k^{(h)}], \text{ and } h = h + 1.$$

End While

Output: The learned intermediate dictionaries $\{\mathbf{D}^{(0)}, \dots, \mathbf{D}^{(H-1)}\}$.

- *Third*, by repeating the above two steps for all instances $\vec{d}_k^{(h-1)}$, we can evolve the intermediate patch dictionary $\mathbf{D}^{(h-1)}$ to $\mathbf{D}^{(h)}$ by letting $\mathbf{D}^{(h)} = [\vec{d}_k^{(h)}]$, where $\vec{d}_k^{(h)} = \vec{p}_k^{(h)}$.

The proposed procedure of progressively constructing multi-layer dictionary \mathbf{D} is briefly summarized in Algorithm 1.

2.3.3. Progressive label fusion via multi-layer dictionary—In the testing stage, given the multi-layer dictionary \mathbf{D} , the weighting vector $\vec{\alpha}$ can be gradually refined from $\vec{\alpha}^{(0)}$ at the initial layer (which is optimal only for the patch appearance representation) to $\vec{\alpha}^{(H-1)}$ in the last layer (which is optimal for the label fusion). Hereafter, we call the weighting vector $\vec{\alpha}^{(0)}$ as the *appearance representation profile*.

Algorithm 2

Multi-layer label fusion

Input: The testing image, multi-layer dictionary \mathbf{D} , and label dictionary \mathbf{L} .

Initialize: $h = 0$, $\vec{y}^{(0)} = \vec{y}$.

For each voxel v **do**

While $h < H$

1. Compute $\vec{\alpha}^{(h)}$ for $\vec{y}^{(h)}$ by $\mathbf{D}^{(h)}$ according to Eq.(2) or Eq.(3);
2. $\vec{y}^{(h+1)} = \mathbf{L}\vec{\alpha}^{(h)}$.

End while

End for

Binarize the probability map $\vec{y}^{(H)}$.

Output: The segmentation result (binary label map) of the testing image.

In the initial layer, we use the original intensity patch dictionary $\mathbf{D}^{(0)}$ to present the target image patch $\vec{y}^{(0)} = \vec{y}$ located at v and thus obtain the representation profile $\vec{\alpha}^{(0)}$ of the initial layer. Conventional label fusion methods stop here and then vote for the label via the weights in $\vec{\alpha}^{(0)}$. Instead, our progressive label fusion method computes the label probability patch $\vec{y}^{(1)}$ by letting $\vec{y}^{(1)} = \mathbf{L}\vec{\alpha}^{(0)}$. It is worth noting that the intensity target image vector $\vec{y}^{(0)}$ now turns to a label-probability vector after going through the initial layer. Then, we iteratively refine the probability patch vector $\vec{y}^{(1)}$ until it reaches the label dictionary. Specifically, we use $\vec{y}^{(1)}$ as the new target patch and further represent it by the intermediate dictionary $\mathbf{D}^{(1)}$ in the same layer, thus obtaining the new label probability patch $\vec{y}^{(2)}$. Next, we continue to represent $\vec{y}^{(2)}$ in the second layer using the intermediate dictionary $\mathbf{D}^{(2)}$. By repeating the same procedure until reaching the last layer, the estimated probability patch $\vec{y}^{(h)}$ becomes sharper and sharper, as shown in Fig. 4. Accordingly, the representation profile $\vec{\alpha}^{(H-1)}$ can be regarded as the best weighting vector, which can be finally used to determine the latent label on the target image point v . The algorithm of the proposed multi-layer label fusion is summarized in Algorithm 2.

3. Experiments

3.1. Dataset

In our experiments, we randomly select 64 3T MRI and 240 1.5T MRI from the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset (www.adni-info.org), which include three types of subjects, i.e., normal control (NC), MCI (Mild Cognitive Impairment) and Alzheimer's disease (AD). The image resolution is resampled to $1 \times 1 \times 1 \text{ mm}^3$. The detailed subject information is shown in Table. 1.

The following three pre-processing steps were performed to all subject images. First, we removed the skull by a learning based meta-algorithm Shi et al. (2012). Second, we used N4-based bias field correction Tustison et al. (2010). Third, we applied intensity standardization to normalize the intensity range Madabhushi and Udupa (2006). Both left and right hippocampi have been manually labeled for each subject. In our experiments, we regard those manual segmentations as ground-truth. To label the target subject, we first aligned all the atlas images to the underlying target subject. In order to improve the computational efficiency, both atlas selection and patch selection strategies were applied, where the selection criterion was based on image intensity similarity.

To evaluate the performance, we use both the Dice ratio and sensitivity to calculate the accuracy and stability by comparing the automated segmentation results with the manual ground-truth. The Dice ratio is defined to measure the overlap between regions R_1 and R_2 via

$$\text{Dice}(R_1, R_2) = 2 \times \frac{|R_1 \cap R_2|}{|R_1| + |R_2|} \quad (4)$$

Sensitivity is a statistical measure of the segmentation results, which is defined as

$$\text{Sens} = \frac{|R_1 \cap R_2|}{|R_1|} \quad (5)$$

where $|\cdot|$ denotes for the regional volume, and R_1 and R_2 represent the estimated and ground-truth labels, respectively.

The mean absolute surface distance (MASD) is a symmetric border positioning measure that gives the mean minimal distance between two surfaces Sluimer et al. (2005). Given the two surfaces S_1 (from the automated segmentation result) and S_2 (from the ground-truth segmentation), let $\bar{d}_{\min}(S_1, S_2)$ be the average minimal distance from all points on surface S_1 to surface S_2 , and then MASD can be defined as

$$\text{MASD}(S_1, S_2) = \frac{1}{2} (\bar{d}_{\min}(S_1, S_2) + \bar{d}_{\min}(S_2, S_1)) \quad (6)$$

It should be mentioned that all testing images were evaluated in a leave-one-out manner. Specifically, in each leave-one-out case, we use FLIRT in the FSL toolbox Smith et al. (2004) with 12 degrees of freedom and the search range of ± 20 in all directions, for the initial linear registration. For deformable registration, we use the diffeomorphic Demons method Vercauteren et al. (2009) with the smoothing kernel size of 2.0, and iteration numbers of 15, 10, and 5 in the low, middle, and high resolutions, respectively.

3.2. Parameters setting

3.2.1. Parameter selection—In the following experiments, we fix the patch size as $5 \times 5 \times 5 \text{ mm}^3$ and the search window as $5 \times 5 \times 5 \text{ mm}^3$ for constructing the dictionaries. In the non-local mean method, the penalty strength σ is set to 0.5, while, in the sparse patch-based label fusion method, the sparse constraint λ is set to 0.1. In all experiments, we discard the searched candidate atlas patches if the pre-selection threshold on the SSIM measurement is less than 0.9.

3.2.2. The influence of atlas number—In order to test the effect of the atlas number, we evaluate the labeling accuracy on 64 3T MRI by using a different number of atlases. For simplicity, we fix the number of the dictionary layer to 4. As shown in Fig. 5, the average Dice ratio keeps increasing as the number of atlases increases. When the atlas number reaches beyond 15, the improvement of the Dice ratio becomes marginal. Considering the computational time, we thus use 15 atlases throughout all the following experiments.

3.2.3. The influence of the number of layer—In this experiment, we use the same dataset as in Section 3.2.2 to investigate the influence of the number of layers on the performance of our proposed method. The evolution curve of the average Dice ratio with respect to the number of dictionary layers is shown in Fig. 6. It can be observed that the improvement of our progressive label fusion method is obvious after using the initial layer (corresponding to the baseline methods), i.e., more than 1% improvement of the Dice ratio over both the non-local and SPBL methods. Our progressive label fusion framework gradually converges after using three layers. Considering the computation time, we use four layers ($H = 4$) in the following experiments.

3.3. Entropy of multi-layer dictionary

We use entropy as a measure of the quality for evaluating our constructed multi-layer dictionary. Note that label map with simple constant values has lower entropy, while the probability map with complex intensity values has higher entropy. Actually, minimum entropy has been used as a good criterion for the design of classification/segmentation methods Palubinskas et al. (1998) Viola and Wells III (1997).

Fig. 7 is an example of the evolution of entropy in the multi-layer dictionary. For a testing patch (i.e., a green solid block in the (top-left) original image) centered at a (red point) voxel, a set of similar patches (i.e., 24 patches used in this figure) are searched (i.e., in the blue dash box), along with their corresponding known label patches (i.e., in the green dash box). By construction of a multi-layer patch dictionary as described in Section 2.3.2, we can obtain the intermediate dictionaries (i.e., those in the red dash box). The tables in the right

panel of this figure provide the corresponding entropies for the patches in the blue, red, and green dash boxes.

Each element in the right tables of Fig. 7 shows the entropy of the corresponding patch in the multi-layer dictionary (i.e., blue and red dash boxes) as well as in the ground-truth label patches (i.e., green dash box). Consistent with our above description, the label patches have lower entropy (as shown at the bottom of the right figure). On the contrary, the intensity image patches have higher entropy (as shown at the top of the right figure), because of their relatively uniform intensity distribution. The right tables also show the entropy of the patch element in each layer of the multi-layer dictionary (i.e., layer 0 to layer 3, as shown from the middle top to middle bottom of the right figure). Since entropy is related to the minimum number of bits required to code the intensity distribution, the predictability is closely related to the value of entropy Mangin (2000). That is, a predictable variable has low entropy, while an unpredictable variable has high entropy. As shown in Fig. 7, from layer 0 to layer 3, the entropy is decreased for most patches, thus, becoming easier to predict. As a result, with the increase of the number of layers, the intermediate dictionary becomes sharper and sharper, which finally becomes closer to the binary label maps.

3.4. Labeling results for 3T MR images

Table. 2 and Table. 3 show the mean and standard deviation of the Dice ratio on the hippocampus (left, right and overall) in linear and deformable registration scenarios, respectively. Compared to the baseline *non-local* and *SPBL* methods with single-layer dictionary, our progressive label fusion framework can improve the labeling accuracy with more than 1% increase of the Dice ratio. The maximum improvement is 1.8% between the conventional non-local method and our progressive non-local method in the case of linear image registration. Also, the improvements over the baseline methods in terms of the overall Dice ratio are statistically significant with a *p*-value less than 0.05 by using paired *t*-test. Significant improvements are indicated with ‘*’ in Table. 2 and Table. 3.

Furthermore, we calculate the surface distance between the ground-truth hippocampus and the estimated hippocampus (left and right). The SPBL method is used as the example to demonstrate the evolution of the surface distance during the progressive label fusion in Fig. 8. According to the color bar shown in the right side of Fig. 8, the surface distances keep decreasing with the increase of the number of intermediate layers. Table. 4 shows the corresponding surface distances on the whole hippocampus in Fig. 8. In accordance with Fig. 8, by increasing the number of intermediate layers, the mean surface distance decreases gradually. When $H=1$, which corresponds to the conventional one-layer SPBL method, the maximum distance is 2.83mm, and it significantly decreases to 1.22mm at $H=4$ by our method.

3.5. Labeling results for 1.5T MR images

We have also applied our multi-layer method on the 240 1.5T MRI dataset. Specifically, we test three datasets, including normal control (NC), mild cognitive impairment (MCI) and Alzheimers disease (AD), and each dataset has 80 cases as shown in Table. 1. For further evaluation, we compute the mean and standard deviation of the hippocampus volumes of the

manual ground truths for each dataset in Table. 5. The table shows significant volume differences among the three datasets. Specifically, compared to the NC data, the volumes of AD and MCI are smaller; in particular, AD has much smaller volume than NC, and also AD has the largest volume variation.

Moreover, the evaluation measures (i.e., Dice ratio, sensitivity and MASD) are obtained using different combined strategies (i.e., non-local and SPBL under linear or deformable registration) on the 1.5T MRI dataset (for AD, MCI and NC), as summarized in Fig. 9, Fig. 10 and Fig. 11, respectively. Obviously, our progressive methods consistently outperform all other methods in all these three evaluation measures, regardless of AD, MCI or NC (with larger improvement for the NC subjects). The best result is achieved for our progressive SPBL method in the case of using deformable registration. As shown in Fig. 9, compared to the baseline methods, our progressive label fusion methods can improve the labeling accuracy for more than 1% Dice ratio, in the cases of both linear and deformable registrations. According to Fig. 11, the mean surface distance is also getting smaller.

In order to further testify to the effectiveness of our proposed method, we also conduct experiments on the mixed datasets of NC, MCI and AD. Table. 6 shows the average Dice ratios by the four methods in both linear and deformable registration cases. Obviously, our progressive methods consistently outperform the conventional methods. Specifically, compared to the conventional non-local method, our progressive method gains improvements of 1.3% and 1.5% for the linear and deformable registration cases, respectively.

3.6. Computational complexity

As illustrated in Algorithm 1 and Algorithm 2, our proposed multilayer label fusion method takes major computational cost in two stages, including (1) dictionary learning and (2) progressive multi-layer label fusion. By analyzing the proposed method, we found three major factors determine the computational cost in including two stages, including the number of layers H , the number of atlas patches K , and the size of each patch M .

Based on these factors, the computational complexities of the dictionary learning and the multi-layer label fusion stages are $O(HK^2M)$ and $O(HM)$, respectively. Therefore, it costs $O(HK^2M)$ for each voxel. Considering the computational time, we have used various strategies to speed up our algorithm, including parallel programming and patch pre-selection.

4. Discussion

4.1. Application to volumetric labeling

Algorithms for automatic segmentation of brain structures are important for brain analysis. Hippocampal volume has been found as an early biomarker for AD, but it is hindered by various limitations of manual segmentation. Although many methods have been proposed for hippocampal segmentation in MR images, they often directly determine the label according to the weights computed from the appearance representation of the testing image. However, the intensity images and label maps are in two different domains, and there is no

evidence to show a direct connection between intensity images and label maps. Thus, in Fig. 2, we construct a bridge to connect these two domains. As shown in Fig. 3, with the help of this built "bridge" (multi-layer dictionary), the probability patches can become sharper and sharper, and thus the intensity image domain and the label domain can be better connected. In this way, the representation coefficients can be iteratively computed from the built dictionaries, and then the latent label can be predicted by applying the estimated coefficients to the atlas labels in a layer-by-layer manner. The proposed method shows its potential in volumetrical labeling. The optimized dictionary also makes our labeling method more robust to registration errors.

4.2. Segmentation quality

The results presented in this work demonstrate that our proposed multilayer label fusion method is well-suited for automated hippocampus segmentation. First, we have tested the influences of both the atlas and layer number. In Fig. 5, we found that further increasing the number of training images will not improve the performance significantly. Considering the computational cost, we select 15 as the atlas number in all experiments of this paper. Also, the multi-layer dictionary learning makes the probability maps much closer to the label maps. Thus, it is also important to choose a suitable layer number. Similar to the atlas number, the use of more layers can lead to better results. As shown in Fig. 6, the use of 4 layers produces the best results for all experiments, and thus we use 4 layers in this paper. To further validate the effectiveness of the multi-layer dictionary learning, we analyze each built dictionary quantitatively by entropy. It is found that entropy is decreased with the increase of the number of layers, which further demonstrates the significance of our multi-layer dictionary learning.

As pointed out in Table. 2 and Table. 3, using the 3T MRI, the Dice ratio can exceed 0.85, showing good overlap with manual segmentations. Also, the surface of automated segmentation is very close to the surface of ground-truth segmentation. Besides, with the 1.5T MRI, our proposed progressive methods also outperform all 1-layer methods, in terms of Dice ratio, sensitivity, and MASD.

4.3. Comparison to other methods

In our method, we have applied either the NPBL or SPBL as a basic label fusion algorithm to obtain the *representation profile* at each layer. Compared to the original (single-layer) NPBL and SPBL methods, our proposed method is consistently better, with statistically significant improvements on both the Dice ratio and surface distance. The direct visual comparison on the large image dataset also shows better results with our method, compared to the two baseline methods.

The performance of our progressive multi-layer label fusion scheme is further compared with the performances of other two single-layer algorithms (i.e., non-local and SPBL) on both the 3T and 1.5T MRI scans. For both cases of linear and deformable registrations, our method outperforms these baseline methods in hippocampal segmentation. On the other hand, the use of deformable registration often gives better results than the use of linear registration, for all methods (including our method and conventional methods).

4.4. Limitation and future work

The main limitation of our proposed method is the time complexity, for the case of using a large number of layers or a large number of selected atlases. In such case, for each testing patch centered at each voxel, a set of patches needs to be selected to estimate the testing patch in a leave-one-out manner, for creating a multi-layer dictionary. Since this procedure needs to be online, it further increases the computational time for our proposed method.

In this paper, we just focus on segmentation of the hippocampus. With the shown benefit of using a multi-layer dictionary for image labeling, it will be interesting to also apply our method on the segmentation of other organs in the future.

5. Conclusion

In this paper, we proposed a progressive label fusion framework for multiatlas segmentation by dictionary evolution. Different from the conventional label fusion methods, in our proposed method, we constructed a sequence of intermediate dictionaries in a multi-layer manner to progressively optimize the weights for label fusion. In this way, our proposed method seeks for the representation weights that can be progressively improved for final label fusion, instead of employing only the original intensity patches based representation weights, which are used in the conventional methods. To this regard, our proposed *multi-layer* method can also be considered as an extension of the conventional *single-layer* methods. We have applied our novel label fusion method to hippocampus segmentation in both 1.5T and 3T MRI brain images from the ADNI dataset. Compared to the *state-of-the-art* counterpart label fusion methods with a *single-layer* dictionary, our proposed *multi-layer* label fusion method achieves better performances.

References

- Artaechevarria X, Munoz-Barrutia A, Ortiz-de Solorzano C. Combination strategies in multi-atlas image segmentation: Application to brain mr data. *Medical Imaging, IEEE Transactions on*. 2009 Aug; 28(8):1266–1277.
- Awate S, Whitaker R. Unsupervised, information-theoretic, adaptive image filtering for image restoration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2006 Mar; 28(3): 364–376.
- Buades A, Coll B, Morel J-M. A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation*. 2005; 4(2):490–530.
- Coupé P, Manjón JV, Fonov V, Pruessner J, Robles M, Collins DL. Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage*. 2011; 54(2): 940–954. [PubMed: 20851199]
- Devanand D, Pradhaban G, Liu X, Khandji A, De Santi S, Segal S, Rusinek H, Pelton G, Honig L, Mayeux R, et al. Hippocampal and entorhinal atrophy in mild cognitive impairment prediction of alzheimer disease. *Neurology*. 2007; 68(11):828–836. [PubMed: 17353470]
- Dickerson BC, Goncharova I, Sullivan M, Forchetti C, Wilson R, Bennett D, Beckett L, et al. Mri-derived entorhinal and hippocampal atrophy in incipient and very mild alzheimers disease. *Neurobiology of aging*. 2001; 22(5):747–754. [PubMed: 11705634]
- Heckemann RA, Hajnal JV, Aljabar P, Rueckert D, Hammers A. Automatic anatomical brain mri segmentation combining label propagation and decision fusion. *NeuroImage*. 2006; 33(1):115–126. [PubMed: 16860573]

- Holland D, Desikan RS, Dale AM, McEvoy LK, et al. Rates of decline in alzheimer disease decrease with age. *PloS one*. 2012; 7(8):e42325. [PubMed: 22876315]
- Madabhushi A, Udupa JK. New methods of mr image intensity standardization via generalized scale. *Medical Physics*. 2006; 33(9):3426–3434. [PubMed: 1702239]
- Mangin, J-F. Entropy minimization for automatic correction of intensity nonuniformity; *Mathematical Methods in Biomedical Image Analysis*, 2000. Proceedings. IEEE Workshop on. IEEE; 2000. p. 162-169.
- Manjón JV, Coupé P, Martí-Bonmatí L, Collins DL, Robles M. Adaptive non-local means denoising of mr images with spatially varying noise levels. *Journal of Magnetic Resonance Imaging*. 2010; 31(1):192–203. [PubMed: 20027588]
- Palubinskas, G.; Descombes, X.; Kruggel, F. An unsupervised clustering method using the entropy minimization; *Pattern Recognition*, 1998. Proceedings. Fourteenth International Conference on. Vol. 2. IEEE; 1998. p. 1816-1818.
- Protter M, Elad M, Takeda H, Milanfar P. Generalizing the nonlocal-means to super-resolution reconstruction. *Image Processing, IEEE Transactions on*. 2009; 18(1):36–51.
- Rohlfing, T.; Brandt, R.; Menzel, R.; Russakoff, DB.; Maurer, CR, Jr. *Handbook of Biomedical Image Analysis*. Springer; 2005. Quo vadis, atlas-based segmentation?; p. 435-486.
- Rousseau F, Habas PA, Studholme C. A supervised patch-based approach for human brain labeling. *Medical Imaging, IEEE Transactions on*. 2011; 30(10):1852–1862.
- Sabuncu MR, Yeo BT, Van Leemput K, Fischl B, Golland P. A generative model for image segmentation based on label fusion. *Medical Imaging, IEEE Transactions on*. 2010; 29(10):1714–1729.
- Shi F, Wang L, Dai Y, Gilmore JH, Lin W, Shen D. Label: pediatric brain extraction using learning-based meta-algorithm. *Neuroimage*. 2012; 62(3):1975–1986. [PubMed: 22634859]
- Sluimer I, Prokop M, Van Ginneken B. Toward automated segmentation of the pathological lung in ct. *IEEE transactions on medical imaging*. 2005; 24(8):1025–1038. [PubMed: 16092334]
- Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, et al. Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage*. 2004; 23:S208–S219. [PubMed: 15501092]
- Song, Y.; Wu, G.; Sun, Q.; Bahrami, K.; Li, C.; Shen, D. *Medical Image Computing and Computer-Assisted Intervention- MICCAI 2015*. Springer; 2015. Progressive label fusion framework for multi-atlas segmentation by dictionary evolution; p. 190-197.
- Tibshirani R. Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2011; 73(3):273–282.
- Tong T, Wolz R, Coupé P, Hajnal JV, Rueckert D, Initiative ADN, et al. Segmentation of mr images via discriminative dictionary learning and sparse coding: Application to hippocampus labeling. *NeuroImage*. 2013; 76:11–23. [PubMed: 23523774]
- Tu Z, Bai X. Auto-context and its application to high-level vision tasks and 3d brain image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2010; 32(10):1744–1757.
- Tustison NJ, Avants BB, Cook P, Zheng Y, Egan A, Yushkevich P, Gee JC, et al. N4itk: improved n3 bias correction. *Medical Imaging, IEEE Transactions on*. 2010; 29(6):1310–1320.
- Van Leemput K, Bakkour A, Benner T, Wiggins G, Wald LL, Augustinack J, Dickerson BC, Golland P, Fischl B. Automated segmentation of hippocampal subfields from ultra-high resolution in vivo mri. *Hippocampus*. 2009; 19(6):549–557. [PubMed: 19405131]
- Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*. 2009; 45(1):S61–S72. [PubMed: 19041946]
- Viola P, Wells WM III. Alignment by maximization of mutual information. *International journal of computer vision*. 1997; 24(2):137–154.
- Wang, H.; Suh, JW.; Das, S.; Pluta, J.; Altinay, M.; Yushkevich, P. Regression-based label fusion for multi-atlas segmentation; *Computer Vision and Pattern Recognition (CVPR)*, 2011 IEEE Conference on. IEEE; 2011. p. 1113-1120.

- Wang H, Suh JW, Das SR, Pluta JB, Craige C, Yushkevich P, et al. Multi-atlas segmentation with joint label fusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on.* 2013; 35(3):611–623.
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on.* 2004; 13(4):600–612.
- Wu G, Wang Q, Zhang D, Nie F, Huang H, Shen D. A generative probability model of joint label fusion for multi-atlas based brain segmentation. *Medical image analysis.* 2014; 18(6):881–890. [PubMed: 24315359]
- Yan, Z.; Zhang, S.; Liu, X.; Metaxas, DN.; Montillo, A. Accurate segmentation of brain images into 34 structures combining a non-stationary adaptive statistical atlas and a multi-atlas with applications to alzheimer’s disease; *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on.* IEEE; 2013. p. 1202-1205.
- Zhang, D.; Guo, Q.; Wu, G.; Shen, D. *Multimodal Brain Image Analysis.* Springer; 2012a. Sparse patch-based label fusion for multi-atlas segmentation; p. 94-102.
- Zhang, S.; Zhan, Y.; Dewan, M.; Huang, J.; Metaxas, D.; Zhou, X. Deformable segmentation via sparse shape representation. In: Fichtinger, G.; Martel, A.; Peters, T., editors. *Medical Image Computing and Computer- Assisted Intervention.* Vol. 6892 of *Lecture Notes in Computer Science.* Berlin Heidelberg: Springer; 2011. p. 451-458.
- Zhang S, Zhan Y, Dewan M, Huang J, Metaxas DN, Zhou XS. Towards robust and effective shape modeling: Sparse shape composition. *Medical image analysis.* 2012b; 16(1):265–277. [PubMed: 21963296]
- Zhang S, Zhan Y, Metaxas DN. Deformable segmentation via sparse representation and dictionary learning. *Medical Image Analysis.* 2012c; 16(7):1385–1396. [PubMed: 22959839]

Highlights

- A progressive multi-atlas label fusion method by deep dictionary evolution is proposed.
- A sequence of intermediate dictionaries was constructed to progressively optimize the weights for label fusion.
- As an extension of the conventional single-layer methods by improving their label fusion performance.

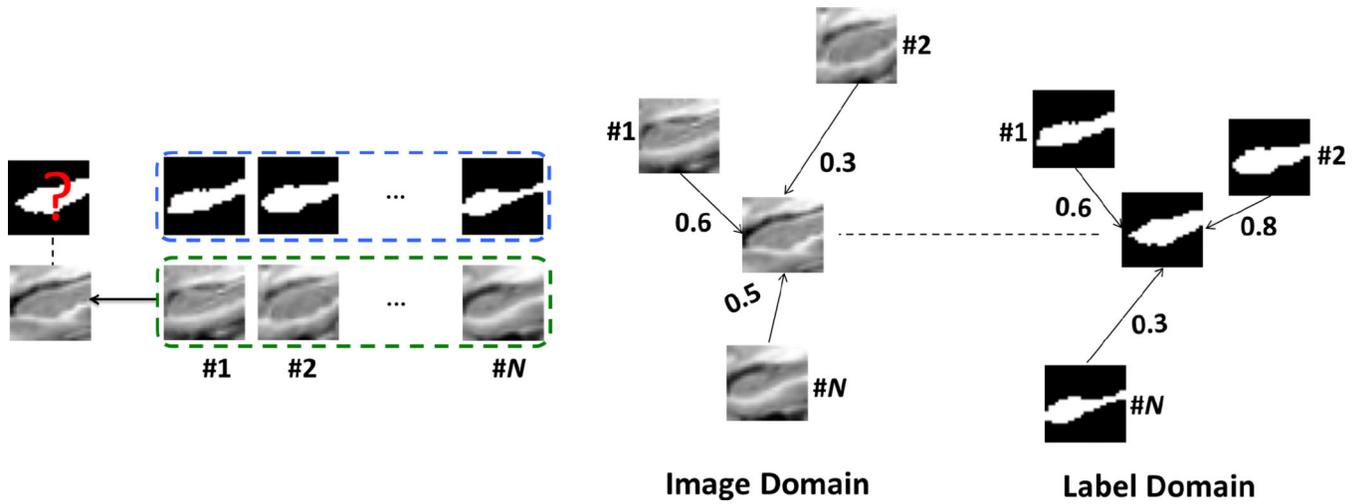


Figure 1. Demonstration of the significant gap of representation profiles computed in the image domain using image appearance and in the label domain using label information.

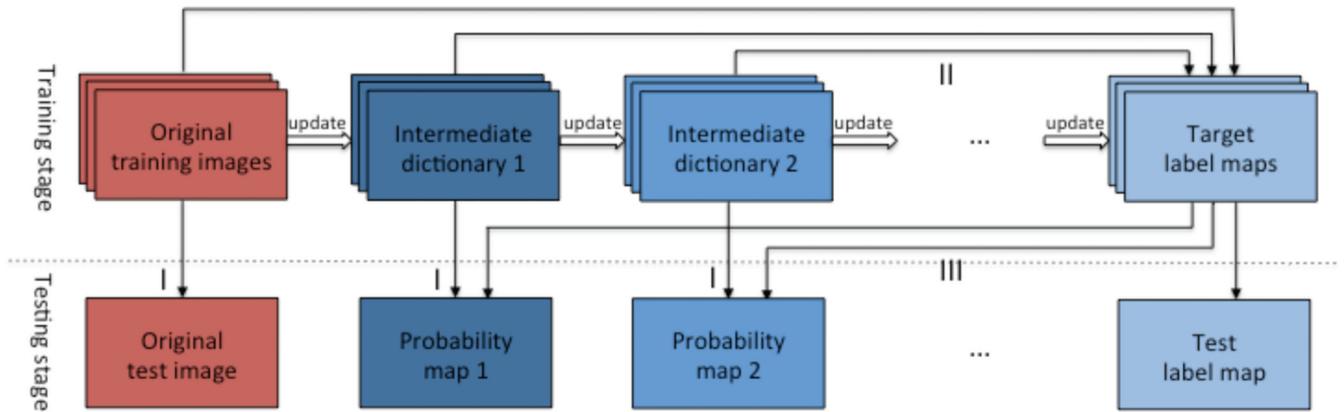


Figure 2. The progressive label fusion framework. To overcome the large gap between the two different dictionaries, i.e., from the original images (red) and label maps (blue), we use a set of intermediate dictionaries to gradually encode the transition from the *representation profile* in the original image domain to the *optimal weighting vector* (for label fusion) in the final label domain. Please see text for more details of this figure.

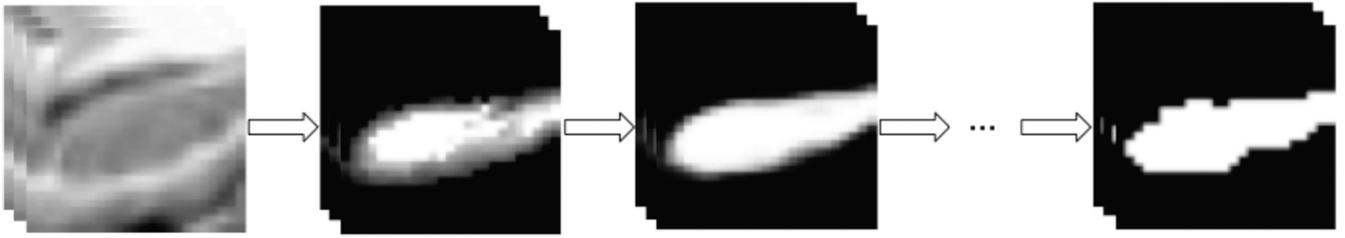


Figure 3. The evolution of the constructed multi-layer dictionary from the *continuous* appearance patches to the *binary* label patches.

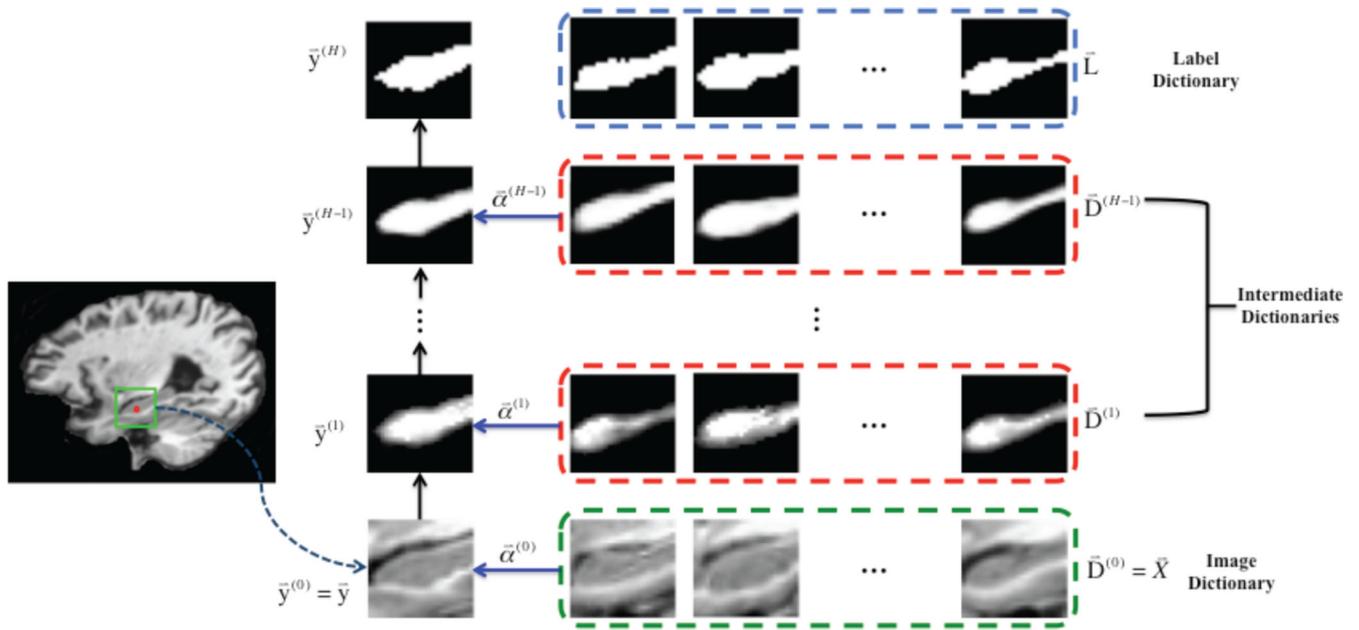


Figure 4.

The framework of the proposed method. Conventional methods estimate the *patch representation profile* using only the image dictionary (i.e., green dash box), and then directly apply it to the label dictionary (blue dash box) for label fusion. To address the large gap between these two different dictionaries (image and label), our method uses a set of intermediate dictionaries (red dash boxes) to gradually guide the estimation of the *representation profile* from the image domain to the label domain. In the application stage, we sequentially go through these intermediate dictionaries and obtain the final binary label.

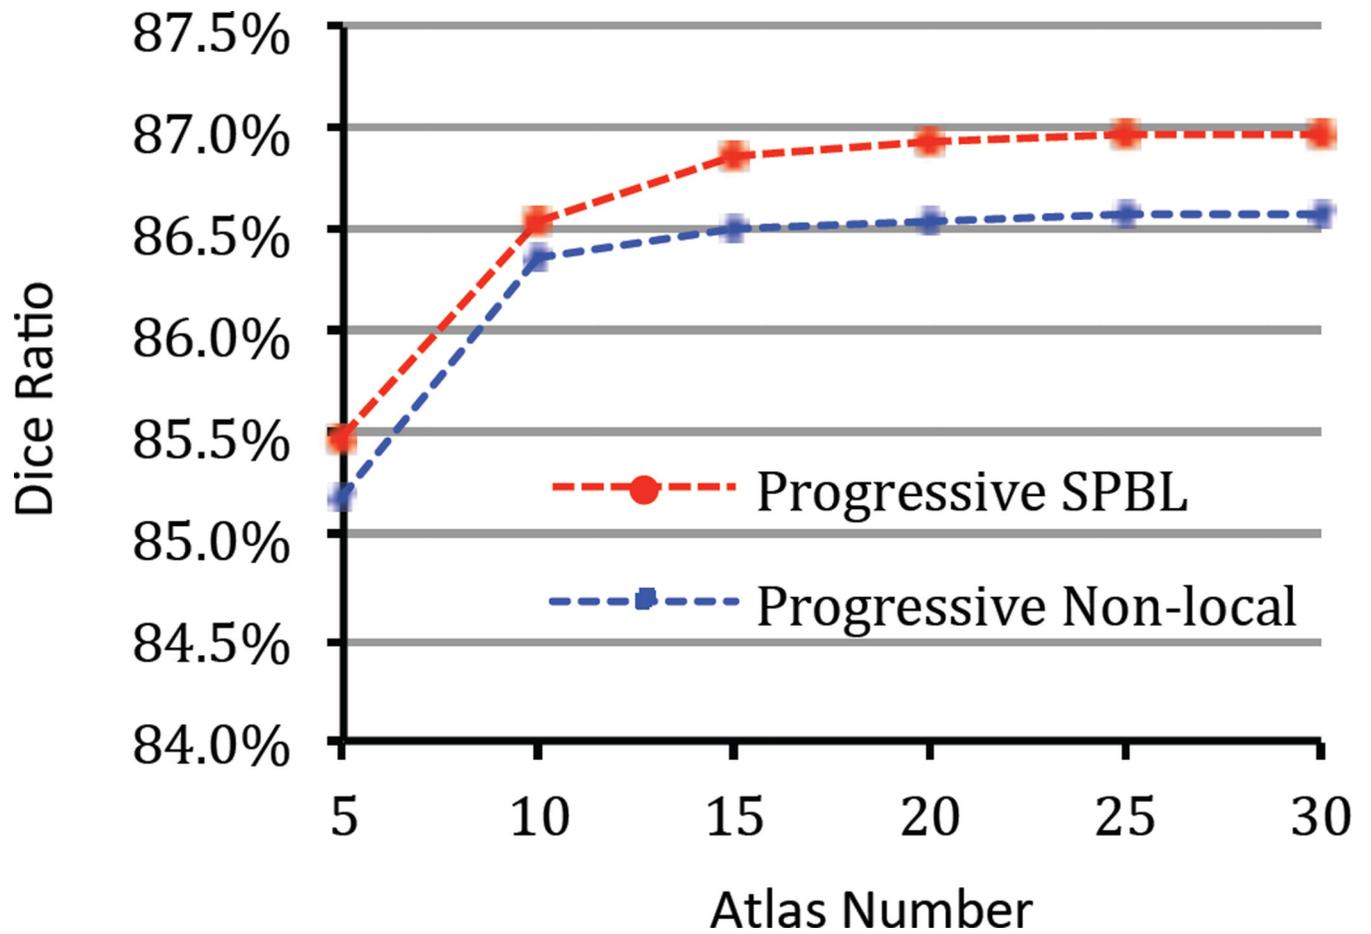


Figure 5. The evolution curves of the Dice ratio *with respect to* the increase of the number of atlas.

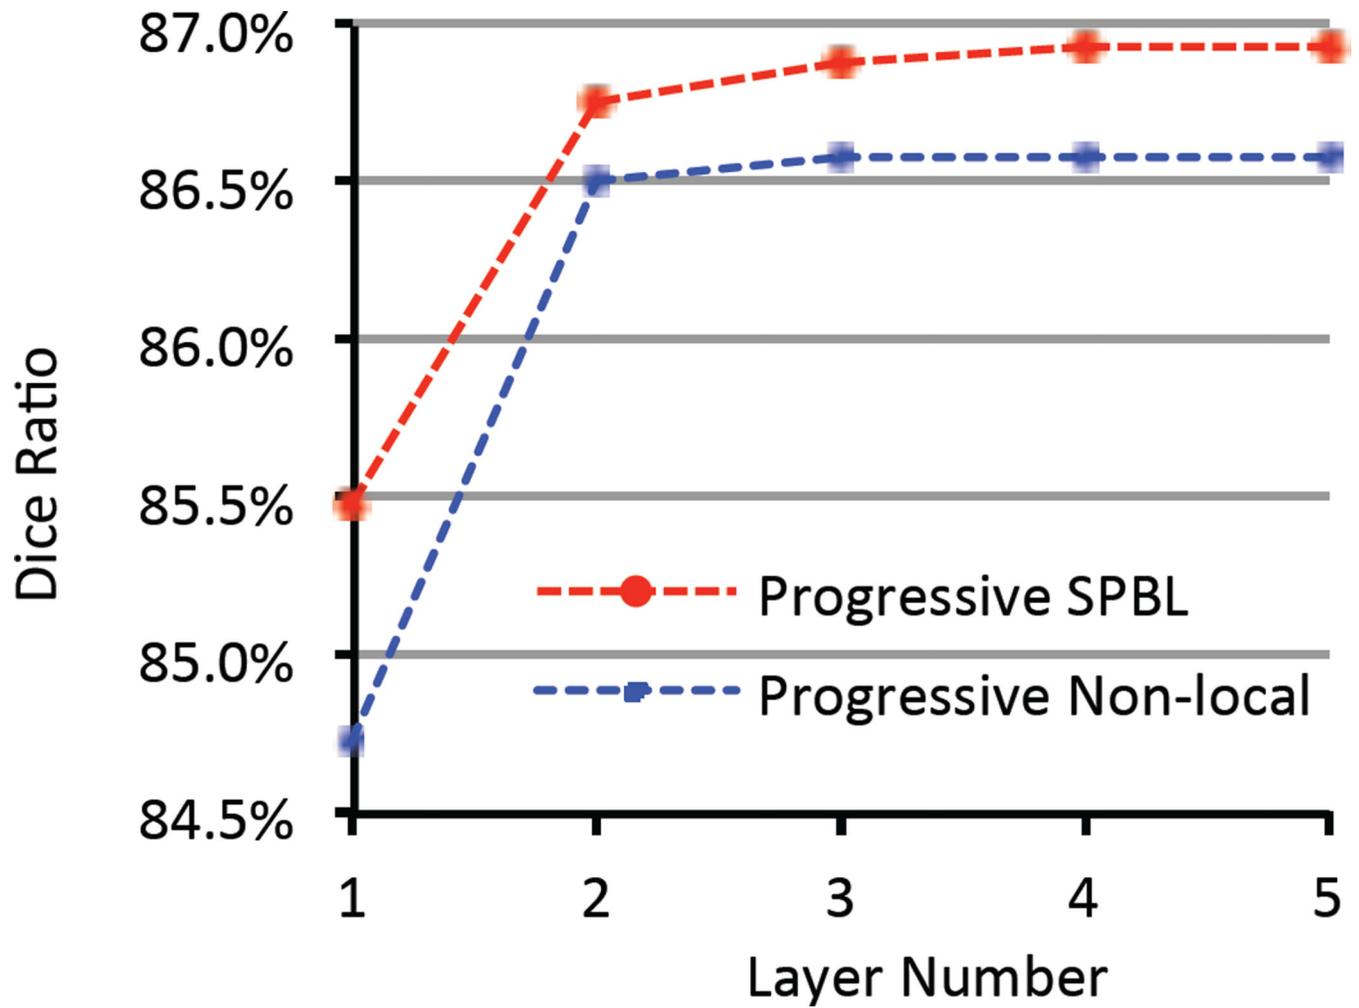


Figure 6. The evolution curves of the Dice ratio *with respect to* the increase of the number of layer.

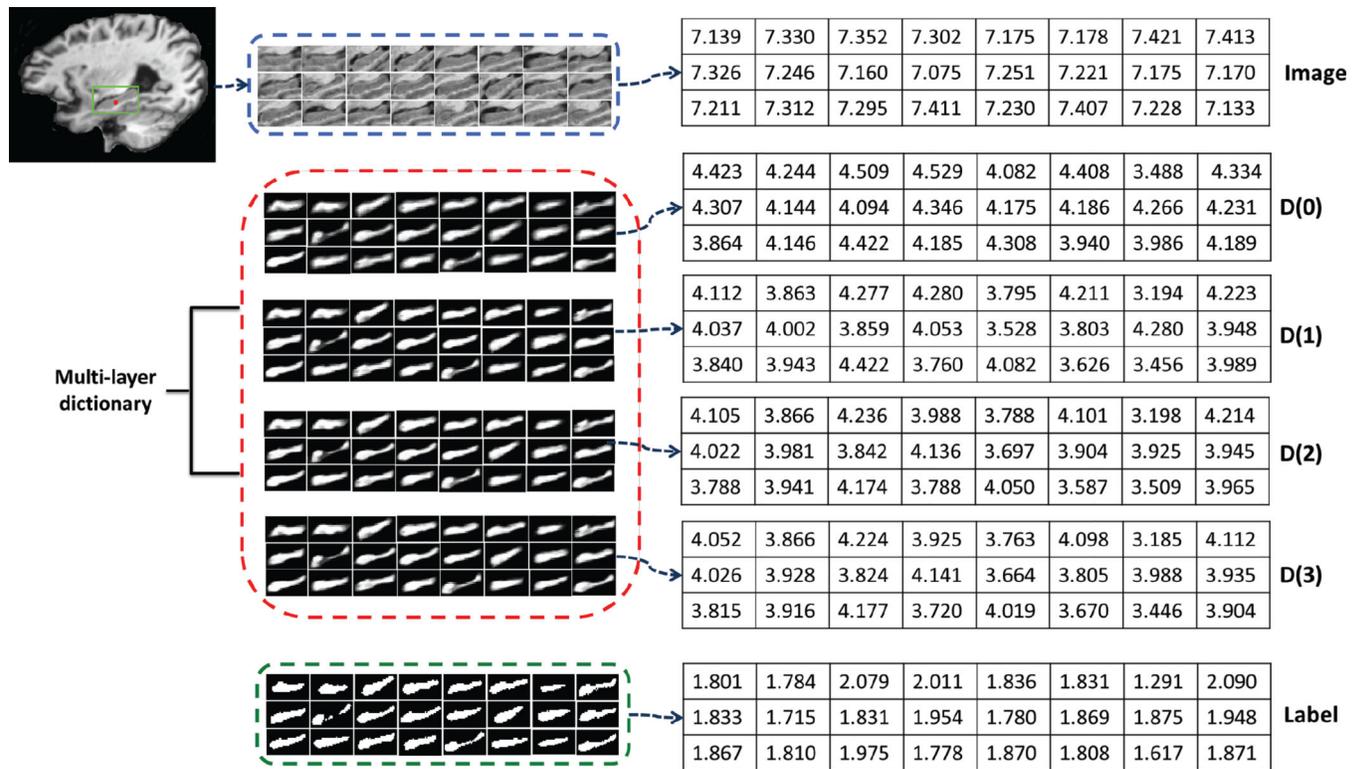


Figure 7.
Multi-layer dictionary and calculated entropy.

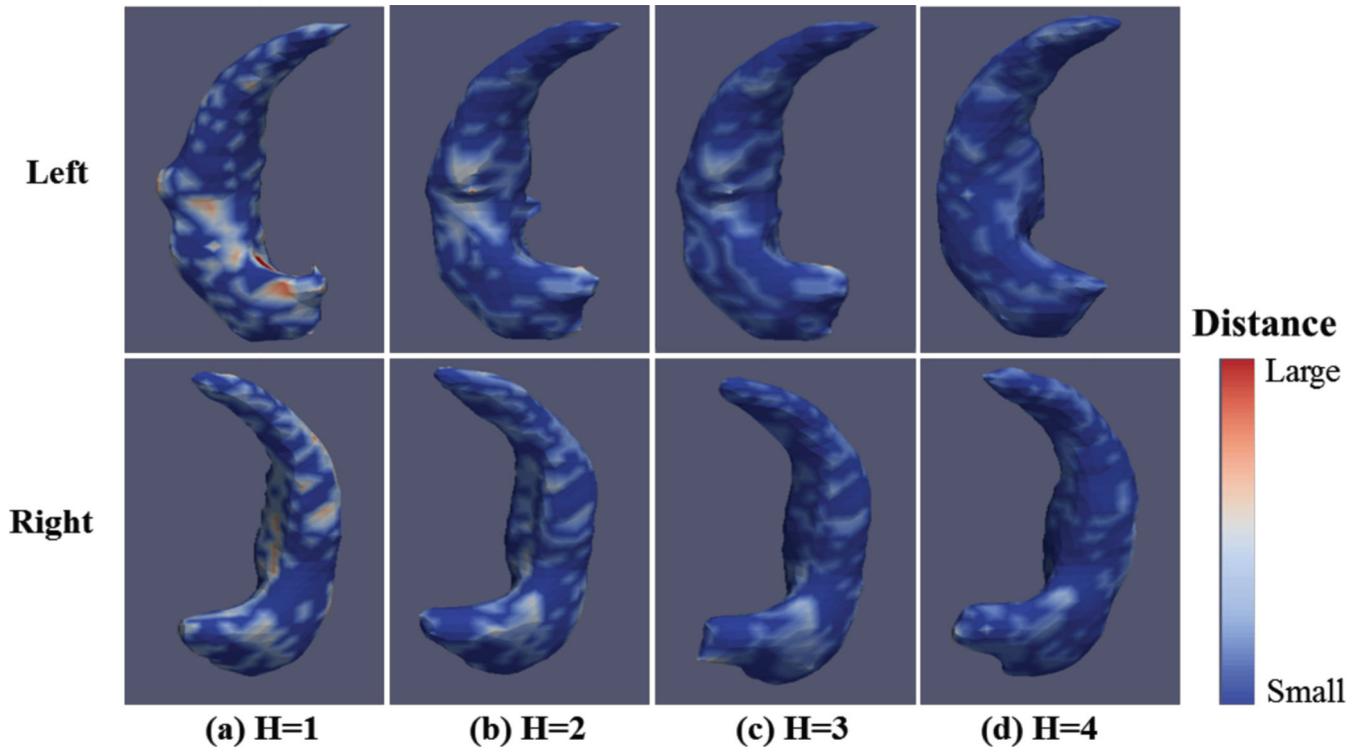


Figure 8. The evolution of surface distances between automatic segmentations and manual ground-truth segmentations from the initial layer (left) to the last layer (right).

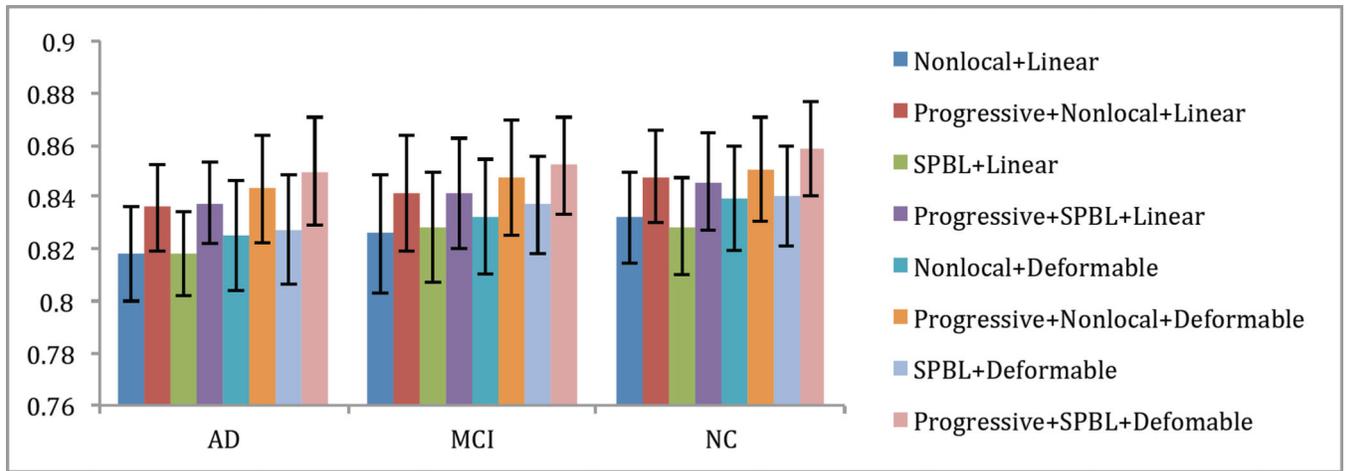


Figure 9.
Dice ratios on 1.5T MRI of AD, NC and MCI subjects.

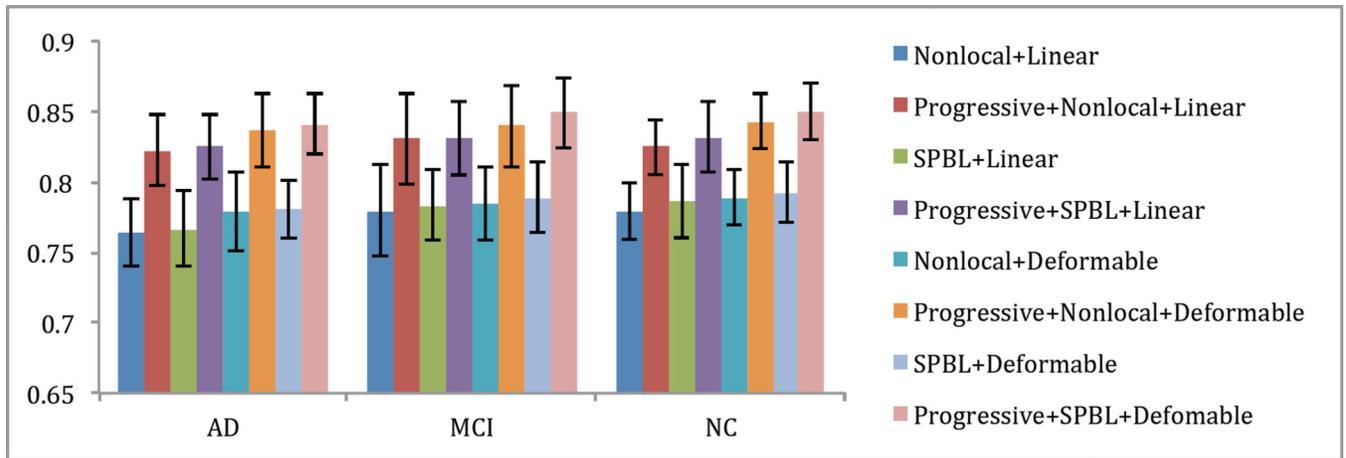


Figure 10.
Sensitivity on 1.5T MRI of AD, NC and MCI subjects.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

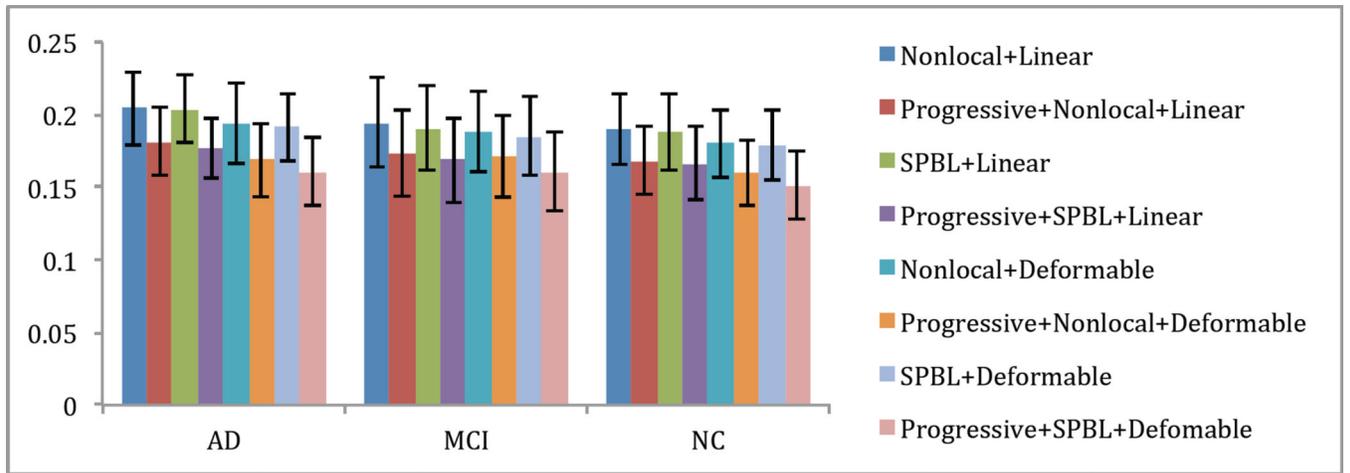


Figure 11.
MASD on 1.5T MRI of AD, NC and MCI subjects.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

Subject information of selected ADNI data used in experiment.

	NC	MCI	AD
3.0 T	20	32	14
1.5 T	80	80	80

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

The mean and standard deviation of the Dice ratio (in %) in labeling the 3T-MRI hippocampus for the case of *linear* registration.

Method	Left	Right	Overall
Conventional Non-local	85.1±5.7	84.3±5.1	84.7±4.2
*Progressive Non-local	86.8±4.5	86.2±5.1	86.5±3.7
Conventional SPBL	85.8±4.5	85.1±4.8	85.5±3.7
Progressive SPBL	87.1±3.2	86.7±5.3	86.9±3.3

Table 3

The mean and standard deviation of the Dice ratio (in %) in labeling the 3T-MRI hippocampus for the case of *deformable* registration

Method	Left	Right	Overall
Conventional Non-local	86.8±4.9	86.6±2.9	86.7±3.2
*Progressive Non-local	87.9±4.0	88.1±3.2	88.0±3.0
Conventional SPBL	87.2±3.6	87.1±3.3	87.2±2.9
*Progressive SPBL	88.2±3.6	88.5±3.1	88.3±2.8

Table 4

The surface distance (unit: mm) on hippocampus labeling between automatic segmentations and manual ground-truth segmentations with different number of intermediate layers

Number of Layers	<i>H=1</i>	<i>H=2</i>	<i>H=3</i>	<i>H=4</i>
Maximum Distance	2.83	2.00	1.52	1.22
Mean Distance	0.27	0.22	0.19	0.18

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 5

Comparison of mean volumes of hippocampus of three data types (AD, MCI and NC) (unit: mm^3).

Data	Left	Right	Overall
AD	1627±395	1574±349	3201±711
MCI	1864±365	1790±355	3654±688
NC	2133±302	2080±277	4213±546

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6

The mean and standard deviation of Dice ratio (in %) in labeling 1.5T-MRI hippocampi.

Method	Linear	Deformable
Conventional Non-local	81.8±3.3	83.4±2.6
Progressive Non-local	83.1±3.1	84.9±2.5
Conventional SPBL	81.9±3.2	83.7±2.6
Progressive SPBL	83.3±3.0	85.1±2.5

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript