

An algorithm for optimal fusion of atlases with different labeling protocols

Juan Eugenio Iglesias^{a,*}, Mert Rory Sabuncu^b, Iman Aganj^b, Priyanka Bhatt^c, Christen Casillas^c, David Salat^b, Adam Boxer^c, Bruce Fischl^{b,d}, Koen Van Leemput^{b,e,f}

^a*Basque Center on Cognition, Brain and Language, Spain*

^b*Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, USA*

^c*Memory and Aging Center, University of California, San Francisco, USA*

^d*MIT Computer Science and Artificial Intelligence Laboratory (CSAIL)*

^e*Department of Applied Mathematics and Computer Science, Technical University of Denmark*

^f*Departments of Information and Computer Science and of Biomedical Engineering and Computational Science, Aalto University, Finland*

Abstract

In this paper we present a novel label fusion algorithm suited for scenarios in which different manual delineation protocols with potentially disparate structures have been used to annotate the training scans (hereafter referred to as “atlases”). Such scenarios arise when atlases have missing structures, when they have been labeled with different levels of detail, or when they have been taken from different heterogeneous databases. The proposed algorithm can be used to automatically label a novel scan with any of the protocols from the training data. Further, it enables us to generate new labels that are not present in any delineation protocol by defining intersections on the underlying labels. We first use probabilistic models of label fusion to generalize three popular label fusion techniques to the multi-protocol setting: majority voting, semi-locally weighted voting and STAPLE. Then, we identify some shortcomings of the generalized methods, namely the inability to produce meaningful posterior probabilities for the different labels (majority voting, semi-locally weighted voting) and to exploit the similarities between the atlases (all three methods). Finally, we propose a novel generative label fusion model that can overcome these drawbacks. We use the proposed method to combine four brain MRI datasets labeled with different protocols (with a total of 102 unique labeled structures) to produce segmentations of 148 brain regions. Using cross-validation, we show that the proposed algorithm outperforms the generalizations of majority voting, semi-locally weighted voting and STAPLE (mean Dice score 83%, vs. 77%, 80% and 79%, respectively). We also evaluated the proposed algorithm in an aging study, successfully reproducing some well-known results in cortical and subcortical structures.

Keywords: Segmentation, Label fusion

1. Introduction

Automatic segmentation of brain structures from MRI data makes it possible to carry out neuroimaging studies at larger scales than manual tracings would, since the latter are very time consuming to make. Moreover, automatic segmentation methods are also more repeatable and reliable than their manual counterparts. Brain MRI segmentation has been used in a number of applications, such as tractography[1], surgical planning[2] and studies of aging[3], brain development[4] and pathologies like Alzheimer’s disease[5].

One family of supervised segmentation techniques that has become popular in brain MRI is multi-atlas segmentation [6]. In conventional atlas-based segmentation, the grayscale image of the atlas is nonlinearly registered to the space of the test scan, and the resulting transform is then used to warp the corresponding labels, which pro-

vide an estimate of the segmentation. Since a single atlas is not sufficient to cover the whole spectrum of variability within a population, multi-atlas segmentation has emerged as a natural extension. Using multiple atlases, this family of techniques produces more accurate segmentations [7] by: (1) independently registering several atlases to the test scan; (2) using the resulting transforms to deform the corresponding label images; and (3) combining the registered label maps into a single estimate of the segmentation with a label fusion algorithm. Multi-atlas segmentation is becoming widespread for three reasons. First, the maturity of registration algorithms (e.g., ANTs/SyN[8], Elastix[9]) enable multi-atlas techniques to achieve very high performance. Second, the public availability of such methods makes multi-atlas segmentation easy to implement. And third, the relative computational cost associated with nonlinearly registering the atlases is quickly diminishing thanks to the rapid increase in processing power of computers.

The choice of label fusion method is critical for the performance of multi-atlas segmentation. Early algorithms

*Corresponding author

Email address: e.iglesias@bcmi.eu (Juan Eugenio Iglesias)

include *best atlas selection* [6] and *majority voting* [10]. The former estimates the segmentation as the labels of the atlas that is most similar to the test scan after registration. In this context, similarity can be measured with the same metrics that are typically used in image registration, such as cross-correlation, mutual information or sum of squared differences. Majority voting, on the other hand, operates at the voxel level by locally assigning the most frequent deformed atlas label at each spatial location - without considering the image intensity information. The performance of majority voting can be increased by an atlas selection process, in which only the deformed atlases that are most similar to the target scan are considered in the fusion[11, 12]

Later fusion methods compute the segmentation as a weighted combination of the labels of the registered atlases such that higher weights are given to more similar atlases. The weights can be global[13] or local[14, 15, 16, 17]. *Sabuncu et al.* [17] have shown that many of these multi-atlas methods can be written within a unified generative model. Another popular label fusion approach is STAPLE[18] and its variants[19, 20, 21, 22]; while this method was originally developed to combine multiple manual segmentations from different human raters, it is increasingly being applied in the context of multi-atlas label fusion.

All the aforementioned label fusion algorithms assume that all structures of interest are labeled in all atlases, which is a rather limiting constraint. Eliminating this requirement would have several practical implications:

- It would enable us to combine training scans from different datasets even if they have different sets of annotated structures. In turn, this would make it possible to take advantage of the increasing amount of heterogeneously labeled MRI data that are publicly available.
- It would also enable us to segment structures that are not included in any of the datasets, but defined as the intersection of labels. For instance, the intersection of the lateral postcentral region and the cerebral gray matter would define the primary somatosensory cortex.
- It would allow for the fusion of segmentations from different modalities with different field of views and resolution. For instance, it would be possible to combine standard resolution brain MRI (1 mm resolution) with high-resolution MRI with limited field of view or even histology or optical coherence tomography data.
- It would be useful if one were to manually relabel a subset of atlases to include finer structures in the annotations. For example, in a large dataset with the hippocampi already labeled, an expert rater can additionally delineate the hippocampal subfields – which

is extremely difficult and time consuming – in just a few cases. Traditional label fusion methods would only be able to use these few scans in the segmentation, having to disregard the information in all the scans in which the subfields are not labeled.

Despite the practical implications that a label fusion algorithm which allows for heterogeneously labeled atlases would have, this direction remains largely unexplored in the literature. To the best of our knowledge, only a particular case of label fusion with heterogeneous labels has been considered so far: the situation in which some of the labels are missing in some of the atlases. To tackle this problem, *Landman et al.*[23, 24, 25] propose an ad-hoc solution by modifying the STAPLE framework such that unlabeled voxels are ignored and the confusion matrix entries corresponding to the missing structures are fixed. *Commowick et al.*[26] propose ameliorating the effect of missing labels by adding a prior on the confusion matrices to the STAPLE algorithm that, when a label is missing, encourages higher a transition probability from that label to the background. However, such an approach treats as background all the voxels that have not been labeled with one of the foreground labels.

In this study, we present a family of probabilistic models for label fusion that make it possible to use atlases that have been annotated with different protocols. In our models, the atlases are assumed to have a hidden “fine” segmentation with all the structures present in the training data – including those defined by intersections of labels. The actual observed labels are assumed to have been obtained by collapsing groups of hidden fine labels into more general, coarse labels.

The contribution of this study is twofold:

- i We use probabilistic models of label fusion to extend three popular methods (majority voting, semi-locally weighted fusion and STAPLE) to the scenario of heterogeneously labeled atlases.
- ii We propose a new generative model for label fusion that can overcome the limitations of these generalizations – the inability to produce meaningful posteriors and to exploit the similarities between the atlases – and show that it outperforms the generalizations in experiments with four datasets.

The rest of this paper is organized as follows. In Section 2, we describe the general framework for label fusion with heterogeneously labeled atlases, propose the generalizations of the different methods, identify their disadvantages, and present a new fusion algorithm to address their shortcomings. In Section 3, we assess the performance of the different algorithms with experiments on four different datasets. Finally, Section 4 concludes the paper.

2. Methods

In this section, we first introduce the general framework and define the variables that we will use throughout

the paper (Section 2.1). Then, we present the generalizations of majority voting, semi-locally weighted voting and STAPLE (Sections 2.2-2.4), and identify their weaknesses in Section 2.5. Finally, we introduce a label fusion method that addresses these shortcomings in Section 2.6.

2.1. General framework

Throughout the remainder of this paper, we will assume that a test scan consisting of J voxels is to be segmented. We will use $\mathbf{y} = \{y_j, j = 1, \dots, J\}$ to refer to the image intensities, and $\mathbf{s} = \{s_j, j = 1, \dots, J\}$ to refer to its hidden, underlying segmentation. Let us also assume that a set of N atlases has been pre-registered to a test scan with a non-linear algorithm. Let $\{\mathbf{i}_n\}$ (where $\mathbf{i}_n = \{i_{nj}, j = 1, \dots, J\}$) be the observed image intensities of the N registered atlases, and let $\{\mathbf{l}_n\}$ (where $\mathbf{l}_n = \{l_{nj}, j = 1, \dots, J\}$) be the corresponding discrete labels, defined at the finest detail level. Their values range from 1 to L , the total number of fine labels.

These deformed labels $\{\mathbf{l}_n\}$ are not directly observed; instead, we have access to a different set of coarse labels $\{\mathbf{c}_n\}$ (where $\mathbf{c}_n = \{c_{nj}, j = 1, \dots, J\}$), which correspond to the actual manual delineations. The coarse labels $\{\mathbf{c}_n\}$ are obtained by collapsing the fine labels $\{\mathbf{l}_n\}$ into different groups of labels by means of a set of N deterministic, protocol-specific functions: $c_{nj} = f_n(l_{nj})$. A protocol function could, for instance, collapse the hippocampal subfields into a single hippocampal label. Having a separate f_n for each atlas enables us to combine different labeling protocols. Different protocol functions can collapse the same fine label into different coarse labels; for instance, orbital cortex could be collapsed into the cerebral cortex by one protocol and into the frontal lobe by another.

We will now generalize three existing models – majority voting, locally weighted voting and STAPLE – to the scenario with collapsed labels. In all cases, the original algorithm (i.e., without heterogeneously labeled atlases) is recovered when all protocol functions are bijective (i.e., no labels are collapsed).

2.2. Generalization of majority voting

As explained in [17], majority voting can be seen as the most likely labeling in a probabilistic model in which the segmentation s_j is sampled randomly from one of the N atlases, as indexed by a hidden discrete field $\mathbf{m} = \{m_j, j = 1, \dots, J\}$. Specifically, the value of the field at a certain voxel $m_j \in \{1, \dots, N\}$ indexes from which atlas we take the label at voxel j , i.e., $s_j = l_{m_j j}$. The field $\mathbf{m}$ follows a flat prior, i.e., $p(\mathbf{m}) \propto 1$. It is straightforward to show that, after marginalizing over $\mathbf{m}$, the distribution of the segmentation at a given location is equal to its relative frequency within the propagated labels of the atlases. Therefore, the most most likely segmentation is the mode of the propagated labels.

The graphical model for generalized majority voting is shown in Figure 1, and the corresponding equations




Figure 1: Graphical model for generalized majority voting. Shaded variables are observed.

1. $\{\mathbf{l}_n\} \sim p(\{\mathbf{l}_n\}) \propto 1$
2. $\mathbf{m} \sim p(\mathbf{m}) \propto 1$
3. $c_{nj} = f_n(l_{nj}), \forall n, j$
4. $s_j = l_{m_j j}, \forall j$

Table 1: Equations for the probabilistic model of generalized majority voting

in Table 1; note that we have assumed a flat prior over the hidden labels. The main difference with the original model is that the fine labels of the atlases are not available; instead, we observe their coarse labels – but still want to compute the segmentation at the fine level. To find the most likely segmentation within this model, we can operate at each spatial location j independently – since the model factorizes over voxels. The problem to solve is $\hat{s}_j = \operatorname{argmax}_{s_j} p(s_j | c_{\cdot j})$, where $c_{\cdot j}$ denotes all the coarse atlas labels at voxel j . The probability $p(s_j | c_{\cdot j})$ is given by:

$$p(s_j | c_{\cdot j}) = \sum_{m_j} p(m_j) p(s_j | m_j, c_{\cdot j}) \quad (1)$$

$$\begin{aligned} &= \sum_{m_j} p(m_j) \sum_{l_{m_j j}} p(s_j | l_{m_j j}) p(l_{m_j j} | c_{m_j j}) \\ &= \sum_{m_j} \frac{1}{N} \sum_{l_{m_j j}} \delta[l_{m_j j} = s_j] \frac{\delta[f_{m_j}(l_{m_j j}) = c_{m_j j}]}{\sum_{l=1}^L \delta[f_{m_j}(l) = c_{m_j j}]} \\ &= \frac{1}{N} \sum_{m_j=1}^N \frac{\delta[f_{m_j}(s_j) = c_{m_j j}]}{\sum_{l=1}^L \delta[f_{m_j}(l) = c_{m_j j}]}, \end{aligned} \quad (2)$$

where $\delta[\cdot]$ is Kronecker's delta. The interpretation of Equation 2 is simple: at each voxel j , every atlas equally spreads its vote over all the fine labels that are compatible with the coarse label $c_{m_j j}$. The segmentation is just the label that maximizes Equation 2 with respect to s_j .

2.3. Generalization of semi-locally weighted voting

Here we generalize the model from [17] to the scenario with collapsed labels. The graphical model is shown in Figure 2, and the corresponding equations in Table 2. The model for the labels is the same as for majority voting; however, now there is a Markov random field (MRF) prior with a predefined constant β on the membership field $\mathbf{m}$, in order to encourage spatial patterns of labels observed in the training data. In addition to the segmentation, the membership field now also generates the intensities of the




Figure 2: Graphical model for generalized semi-locally weighted voting. Shaded variables are observed.

- | |
|---|
| <ol style="list-style-type: none"> 1. $\{I_n\} \sim p(\{I_n\}) \propto 1$ 2. $\mathbf{m} \sim p(\mathbf{m}) \propto \prod_{j=1}^J \exp \left[\frac{\beta}{2} \sum_{j' \in \mathcal{H}_j} \delta[m_j = m_{j'}] \right]$ 3. $\sigma^2 \sim p(\sigma^2) \propto 1$ 4. $c_{nj} = f_n(l_{nj}), \forall n, j$ 5. $s_j = l_{m_j j}, \forall j$ 6. $y_j \sim p(y_j m_j, \sigma^2, i_{.j}) = \mathcal{N}(y_j; i_{m_j j}, \sigma^2)$ |
|---|

Table 2: Equations for the generative model of generalized semi-locally weighted voting. $\mathcal{H}_j$ represents the neighborhood of voxel j .

test scan at each voxel by selecting atlas m_j , taking its intensity $i_{m_j j}$, and corrupting it with Gaussian noise with variance σ^2 . The variance is independent of the spatial location. In [17], it is set to a predefined value; here, we estimate it from the data instead assuming a flat prior for it. The segmentation is formulated as:

$$\begin{aligned} \hat{\mathbf{s}} &= \underset{\mathbf{s}}{\operatorname{argmax}} p(\mathbf{s} | \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}) \\ &= \underset{\mathbf{s}}{\operatorname{argmax}} \int_{\sigma^2} p(\mathbf{s} | \sigma^2, \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}) p(\sigma^2 | \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}) d\sigma^2. \end{aligned} \quad (3)$$

$$\approx \underset{\mathbf{s}}{\operatorname{argmax}} p(\mathbf{s} | \hat{\sigma}^2, \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}), \quad (4)$$

where we have used the mode approximation for the integral. The point estimate $\hat{\sigma}^2$ of the variance is:

$$\hat{\sigma}^2 = \underset{\sigma^2}{\operatorname{argmax}} p(\sigma^2 | \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}). \quad (5)$$

We now describe algorithms to compute the most likely variance $\hat{\sigma}^2$ (Section 2.3.1) and subsequently the most likely segmentation (Section 2.3.2).

2.3.1. Computation of the most likely variance $\hat{\sigma}^2$.

We first note that, according to the model in Figure 2, the variance σ^2 is independent of the coarse labels $\{\mathbf{c}_n\}$ when the segmentation $\mathbf{s}$ is unknown. Therefore, the problem in Equation 5 can be rewritten as follows:

$$\begin{aligned} \hat{\sigma}^2 &= \underset{\sigma^2}{\operatorname{argmax}} p(\sigma^2 | \{\mathbf{i}_n\}, \{\mathbf{c}_n\}, \mathbf{y}) \\ &= \underset{\sigma^2}{\operatorname{argmax}} p(\sigma^2 | \{\mathbf{i}_n\}, \mathbf{y}) \\ &= \underset{\sigma^2}{\operatorname{argmax}} \log p(\mathbf{y} | \sigma^2, \{\mathbf{i}_n\}). \end{aligned} \quad (6)$$

Equation 6 requires marginalizing over $\mathbf{m}$, which leads to an intractable sum due to the MRF prior. Instead, we will use the variational expectation maximization (VEM) algorithm to estimate an approximate solution. Rather than optimizing Equation 6 directly, we maximize a lower bound J :

$$\begin{aligned} J(q(\mathbf{m}, \sigma^2)) &= \log p(\mathbf{y} | \sigma^2, \{\mathbf{i}_n\}) - KL[q(\mathbf{m}) || p(\mathbf{m} | \mathbf{y}, \sigma^2, \{\mathbf{i}_n\})] \quad (7) \\ &= H[q] + \sum_{\mathbf{m}} q(\mathbf{m}) \log p(\mathbf{m}, \mathbf{y} | \sigma^2, \{\mathbf{i}_n\}), \quad (8) \\ &\leq \log p(\mathbf{y} | \sigma^2, \{\mathbf{i}_n\}) \end{aligned}$$

where KL denotes the Kullback-Leibler divergence and H represents the entropy of a random variable. The inequality $J(q(\mathbf{m}, \sigma^2)) \leq \log p(\mathbf{y} | \sigma^2, \{\mathbf{i}_n\})$ holds because the KL divergence is non-negative. The distribution $q(\mathbf{m})$ represents an approximation to the posterior distribution of $\mathbf{m}$ given the observed intensities and the variance. This distribution is optimized over a class of restricted functions. The standard approximation, known as mean field approximation, is that q factorizes over voxels: $q(\mathbf{m}) = \prod_{j=1}^J q_j(m_j)$, where q_j is a categorical distribution over the indices of the atlases at voxel j .

VEM alternates between an expectation (E) step and a maximization (M) step. In the E step, we maximize the bound J with respect to $q(\mathbf{m})$. In the M step, we maximize J with respect to the model parameters – in this case, the variance σ^2 . In the **E step**, it is convenient to work with Equation 7: maximizing J is equivalent to minimizing the KL divergence, which yields the following update:

$$q_j(m_j) \propto p(y_j | i_{m_j j}, \sigma^2) \exp \left[\beta \sum_{j' \in \mathcal{H}_j} q_{j'}(m_{j'}) \right], \quad (9)$$

which can be solved with fixed point iterations, normalizing q_j after each step.

In the **M step**, it is more convenient to work with Equation 8, since the entropy term can be disregarded. The maximization yields the following update:

$$\sigma^2 = \frac{1}{J} \sum_{j=1}^J \sum_{m_j=1}^N q_j(m_j) (y_j - i_{m_j j})^2. \quad (10)$$

The VEM algorithm typically converges in a few (5-6) iterations. Note that, if we set $\beta = 0$ in the model, we recover the standard EM algorithm [27].

2.3.2. Computation of the most likely segmentation $\hat{\mathbf{s}}$.

Given $\hat{\sigma}^2$, computing the segmentation $\hat{\mathbf{s}}$ still requires evaluating an intractable sum over $\mathbf{m}$. However, since $q(\mathbf{m})$ minimizes the KL divergence with $p(\mathbf{m} | \mathbf{y}, \sigma^2, \{\mathbf{i}_n\})$, we can approximate the problem in Equation 4 with:

$$\begin{aligned} \hat{\mathbf{s}} &= \underset{\mathbf{s}}{\operatorname{argmax}} \sum_{\mathbf{m}} p(\mathbf{s} | \mathbf{m}, \{\mathbf{c}_n\}) p(\mathbf{m} | \mathbf{y}, \{\mathbf{i}_n\}, \hat{\sigma}^2) \\ &\approx \underset{\mathbf{s}}{\operatorname{argmax}} \prod_{j=1}^J \sum_{m_j=1}^N q_j(m_j) p(s_j | m_j, c_{.j}). \end{aligned}$$




Figure 3: Graphical models for generalized STAPLE. (a) Expanded version. (b) Compact version. Shaded variables are observed.

Therefore, the optimal segmentation can be computed voxel by voxel as:

$$\hat{s}_j \approx \operatorname{argmax}_{s_j} \sum_{m_j=1}^N q_j(m_j) p(s_j | c_{\cdot j}, m_j),$$

which is almost identical to the right hand side of Equation 1, with the difference that the term $p(m_j)$ has been replaced by $q_j(m_j)$. It is straightforward to show that the approximate posterior label probabilities are given by:

$$p(s_j | y_j, c_{\cdot j}, i_{\cdot j}, \hat{\sigma}^2) \approx \sum_{m_j=1}^N q_j(m_j) \frac{\delta[f_{m_j}(s_j) = c_{m_j j}]}{\sum_{l_{m_j j}} \delta[f_{m_j}(l_{m_j j}) = c_{m_j j}]} \quad (11)$$

This expression is similar to the equation for the label posteriors of generalized majority voting (Equation 2). The difference is that the constant term $1/N$ has been replaced by the approximate membership posteriors q_j . This term depends on the image intensities, such that the contribution is higher for the atlases that are semi-locally more similar to the novel scan. The vote of each atlas is still spread equally among the fine labels they are compatible with the coarse label at each voxel.

2.4. Generalization of STAPLE

The generative model of STAPLE is as follows: the hidden segmentation is generated by a prior $p(\mathbf{s}) = \prod_{j=1}^J p(s_j)$, such that $p(s_j)$ is a categorical distribution that reflects the prior frequencies of the classes. In our scenario, we used a flat prior (i.e., $p(s_j) \propto 1$); preliminary experiments showed that the prior used in the original STAPLE algorithm – in which the frequencies are proportional to the volumes of the structures – considerably decreased the segmentation accuracy in small structures. Given the segmentation, the (deformed) atlas labels are assumed to be independent corrupted observations of the hidden ground truth segmentation. Each atlas has a corresponding confusion matrix through which its labels are generated, whereas the atlas intensities are disregarded in the fusion. The STAPLE algorithm first computes point estimates for the confusion matrices in light of the observed data, and then uses them to estimate the segmentation.

The graphical model for generalized STAPLE is shown in Figure 3a. Each atlas is characterized by a confusion matrix Ω^n . The element Ω_{ls}^n corresponds to the probability that the true segmentation s is observed as label l by atlas n . The fine labels of the atlases are collapsed into the

coarse, observed labels through the protocol functions f_n . For our purposes, it is actually more convenient to work with a more compact version of this model, as shown in Figure 3b. We assume that the collapsed labels are generated directly by confusion matrices Θ^n : Θ_{cs}^n corresponds to the probability that the true segmentation s is observed as coarse label c by atlas n . The relation between Ω^n and Θ^n is simple: $\Theta_{cs}^n = \sum_{\{l | f_n(l)=c\}} \Omega_{ls}^n$. Note that the matrices $\{\Theta^n\}$ are not necessarily square, whereas the $\{\Omega^n\}$ are. Moreover, the matrices $\{\Theta^n\}$ will in general have different number of rows.

As in the original STAPLE algorithm, we compute the maximum likelihood estimates of the confusion matrices as follows:

$$\begin{aligned} \{\hat{\Theta}^n\} &= \operatorname{argmax}_{\{\Theta^n\}} \log p(\{\mathbf{c}_n\} | \{\Theta^n\}) \\ &= \operatorname{argmax}_{\{\Theta^n\}} \log \sum_{\mathbf{s}} p(\{\mathbf{c}_n\} | \mathbf{s}, \{\Theta^n\}) \\ &= \operatorname{argmax}_{\{\Theta^n\}} \log \prod_{j=1}^J \sum_{s_j=1}^L \prod_{n=1}^N p(c_{nj} | s_j, \Theta^n) \\ &= \operatorname{argmax}_{\{\Theta^n\}} \sum_{j=1}^J \log \left[\sum_{s_j=1}^L \prod_{n=1}^N p(c_{nj} | s_j, \Theta^n) \right]. \quad (12) \end{aligned}$$

Equation 12 can be iteratively optimized with EM. In the **E step**, we compute $\tilde{W}_j^s$, a soft classification of voxel j :

$$\tilde{W}_j^s = \frac{\prod_{n=1}^N p(c_{nj} | s, \tilde{\Theta}^n)}{\sum_{s'=1}^L \prod_{n=1}^N p(c_{nj} | s', \tilde{\Theta}^n)} = \frac{\prod_{n=1}^N \tilde{\Theta}_{c_{nj}s}^n}{\sum_{s'=1}^L \prod_{n=1}^N \tilde{\Theta}_{c_{nj}s'}^n} \quad (13)$$

In the **M step**, we update the confusion matrices:

$$\Theta_{cs}^n = \frac{\sum_{j=1}^J \tilde{W}_j^s \delta[c_{nj} = c]}{\sum_{j=1}^J \tilde{W}_j^s} \quad (14)$$

Once the EM algorithm has converged, the approximate label posteriors are given by:

$$p(s_j | c_{\cdot j}) \approx p(s_j | c_{\cdot j}, \{\hat{\Theta}^n\}) = \tilde{W}_j^{s_j}, \quad (15)$$

and the discrete segmentation is just:

$$\hat{s}_j \approx \operatorname{argmax}_s \tilde{W}_j^s.$$

2.5. Shortcomings of the generalized methods

The presented generalizations of majority voting, semi-locally weighted voting and STAPLE suffer from several limitations. Generalized majority voting inherits from its parent the inability to exploit the information in the deformed atlas intensities. In addition, this method is unable to produce realistic label posteriors (soft segmentations) in many common multi-protocol scenarios due to the spreading of the votes across all the compatible fine labels at each




Figure 4: Example to illustrate the shortcomings of generalized majority voting. (a) Labels for a registered atlas with labels for the whole cortex and for a number of subcortical structures. (b) Labels for a registered atlas with labels for the hippocampal subfields. (c) Fusion of (a) and (b) with generalized majority voting. (d) Posterior probability map for the cortex; note that the values are close to 0.5 (rather than 1) in the cortex, and that a large amount of probability mass is distributed all across the image in non-cortical areas (other than the hippocampus). (e) Posterior probability map for the amygdala; similar observations as for (d) can be made.

location. For example, let’s consider the problem of fusing the segmentation in Figure 4a – which includes the cortex as a whole and most subcortical structures – with the segmentation in Figure 4b – which includes the hippocampal subfields. Since the background label of the subfield data is compatible with all non-hippocampal labels, this scan would contribute a flat map of non-zero probability for all non-hippocampal structures all over the image domain. Even though the hard segmentation (Figure 4c) might still be meaningful, the label posteriors given by the method (Figures 4d and 4e) are not realistic. For instance, the posterior for the cortex is close to 0.5 (rather than 1) around this structure in Figure 4d, and the same thing happens with the amygdala in Figure 4e.

Soft segmentations are desirable for two reasons: first, they are useful to estimate the uncertainty in the boundary locations; and second, they allow to compute more accurate estimates of the volumes of the different structures using expectations. If V_l is the volume of structure l in voxels [28]:

$$E[V_l] = E \left[\sum_{j=1}^J \delta(s_j = l) \right] = \sum_{j=1}^J E[\delta(s_j = l)] = \sum_{j=1}^J p_j^{post}(l),$$

where p_j^{post} is the posterior label probability at voxel j , given by Equation 2 (majority voting), 11 (semi-locally weighted voting) or 15 (STAPLE). If we computed volumes from the posteriors in Figure 4d or 4e, the estimates would clearly not be accurate.

In contrast to generalized majority voting, generalized semi-locally weighted voting uses the image intensities to estimate the contributions from the different atlases; however, it still suffers from the shortcoming that it spreads the votes across all the fine labels and cannot generate meaningful posteriors. STAPLE, on the other hand, can produce realistic posteriors thanks to the nature of its generative model, but it shares with majority voting the limitation that it does not consider the image intensities in the fusion. Moreover, STAPLE was originally conceived as a method to merge manual segmentations, and generally performs poorly in the label fusion step of multi-atlas segmentation, where it is often outperformed even by ma-

jority voting (see for instance [29]).

Even though some of these shortcomings could be addressed by newer versions of the algorithms (e.g., an extension of STAPLE using image intensities was presented in [20]), none of the generalized methods supports exchange of information *between* the atlases during the fusion. In the standard case where all atlases have all possible labels, this is not necessary, since there is nothing they can learn from one another. However, in the multi-protocol scenario, the exchange of information between atlases with missing labels can improve the segmentation.

For instance, let us assume that we have three registered atlases $\{\mathbf{i}_1, \mathbf{c}_1\}$, $\{\mathbf{i}_2, \mathbf{c}_2\}$, $\{\mathbf{i}_3, \mathbf{c}_3\}$ and that, at a given voxel j , their deformed coarse labels are $c_{1j} = c'$, $c_{2j} = c''$ and $c_{3j} = c'''$, respectively. Let us further assume that c' and c'' correspond to fine labels l' and l'' without any ambiguity, i.e., $f_1(l) = c' \iff l = l'$ and $f_2(l) = c'' \iff l = l''$, but that c''' is the result of collapsing l' and l'' , i.e., $f_3(l') = f_3(l'') = c'''$. In that case, if $i_{3j} \approx i_{1j}$, we would expect the hidden fine label l_{3j} to be l' , whereas if $i_{3j} \approx i_{2j}$, we would expect it to be l'' instead.

To address the described issues, we present in Section 2.6 below a label fusion method specifically designed for multi-protocol scenarios.

2.6. Proposed fusion method

Figure 5 shows the graphical model of the proposed framework, which we coin “multi-protocol label fusion” (MPLF). Table 3 displays the corresponding equations. Essentially, we are assuming that both the registered atlases and the test scan are generated by a (latent) statistical atlas of labels (α^L) and intensities (α^I) in the space of the test scan. The assumption that the atlases were generated by the same statistical atlas is what allows the model to integrate the information across the atlases. Specifically, the vector $\alpha_j^L = \{\alpha_{jl}^L, l = 1, \dots, L\}$ represents the a priori probability of observing the different labels (at the fine detail level) at voxel j . The vector $\alpha_j^I = \{\alpha_{jl}^I, l = 1, \dots, L\}$ stores, for each fine label, the mean of a Gaussian that models the distribution of intensities at voxel j conditioned on that label. All the Gaussians share a predefined variance σ^2 . The voxels are assumed to




Figure 5: Graphical models for MPLF, the proposed fusion method. Shaded variables are observed.

1. $\alpha^{\mathcal{L}} \sim p(\alpha^{\mathcal{L}}) = \prod_{j=1}^J \text{Dir}[\alpha_j^{\mathcal{L}}; (1 + \epsilon)\mathbf{1}]$
2. $l_{nj} \sim \alpha_j^{\mathcal{L}}, \forall n, j$
3. $c_{nj} = f_n(l_{nj}), \forall n, j$
4. $s_j \sim \alpha_j^{\mathcal{L}}, \forall j$
5. $\alpha^{\mathcal{I}} \sim p(\alpha^{\mathcal{I}}) = \prod_{j=1}^J \prod_{l=1}^L \mathcal{N}[\alpha_{jl}^{\mathcal{I}}; \mu_0, \frac{\sigma^2}{\epsilon}]$
6. $i_{nj} \sim \mathcal{N}(\alpha_{jl_{nj}}^{\mathcal{I}}, \sigma^2), \forall n, j$
7. $y_j \sim \mathcal{N}(\alpha_{js_j}^{\mathcal{I}}, \sigma^2), \forall j$

Table 3: Equations for the generative model of Figure 5. $\text{Dir}[\cdot]$ represents the Dirichlet distribution, and $\mathbf{1}$ is the all one vector.

be independent of one another. As in the previous sections, the fine labels of the atlases are hidden, and we only have access to corresponding coarse segmentations $\{c_n\}$.

The model is completed with priors for $\alpha^{\mathcal{L}}$ and $\alpha^{\mathcal{I}}$, which we assume factorize over voxels as well. We use conjugate priors: for $\alpha_j^{\mathcal{L}}$, we assume a Dirichlet distribution with concentration parameter $1 + \epsilon$, i.e., we assume that we have ϵ prior observations for each class at each voxel. For $\alpha_{jl}^{\mathcal{I}}$ we assume a Gaussian distribution with mean μ_0 and variance σ^2/ϵ , which is equivalent to having ϵ priors observations with sample mean μ_0 . In practice, ϵ is small and the only objective of these priors is to ensure the numerical stability of the algorithm.

Exact inference within this model would require marginalizing over the model parameters, i.e., the statistical atlas $(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}})$, which leads to an intractable integral. We use the assumption that the posterior distribution of the model parameters is heavily peaked to approximate:

$$p(\mathbf{s}|\mathbf{y}, \{\mathbf{i}_n\}, \{\mathbf{c}_n\}) \simeq p(\mathbf{s}|\hat{\alpha}^{\mathcal{I}}, \hat{\alpha}^{\mathcal{L}}, \mathbf{y}, \{\mathbf{i}_n\}, \{\mathbf{c}_n\}), \quad (16)$$

where the point estimates $\hat{\alpha}^{\mathcal{I}}$ and $\hat{\alpha}^{\mathcal{L}}$ are given by:

$$(\hat{\alpha}^{\mathcal{L}}, \hat{\alpha}^{\mathcal{I}}) = \underset{\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}}{\text{argmax}} p(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}|\mathbf{y}, \{\mathbf{i}_n\}, \{\mathbf{c}_n\}). \quad (17)$$

To compute the point estimates, it is convenient to notice that the test image can be considered an extra atlas, such that:

$$\begin{aligned} \mathbf{i}_{n+1} &= \mathbf{y}, \\ \mathbf{l}_{n+1} &= \mathbf{s}, \\ \mathbf{c}_{n+1} &= \mathbf{1}, \\ f_{n+1}(s) &= 1, \forall s, \end{aligned}$$

In other words, the test image is an additional atlas (index $N + 1$) with a constant coarse label 1, which is compatible with all the labels at the fine level. Then, we can use Bayes's rule to rewrite the problem in Equation 17 as:

$$\begin{aligned} (\hat{\alpha}^{\mathcal{L}}, \hat{\alpha}^{\mathcal{I}}) &= \underset{\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}}{\text{argmax}} p(\{\mathbf{i}_n\}, \{\mathbf{c}_n\}|\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}) p(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}) \\ &= \underset{\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}}{\text{argmax}} \log p(\alpha^{\mathcal{L}}) + \log p(\alpha^{\mathcal{I}}) + \sum_{n=1}^{N+1} \log p(\mathbf{i}_n, \mathbf{c}_n|\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}) \\ &= \underset{\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}}{\text{argmax}} \sum_{j=1}^J \log p(\alpha_j^{\mathcal{L}}) + \sum_{j=1}^J \sum_{l=1}^L \log p(\alpha_{jl}^{\mathcal{I}}) + \dots \\ &\quad \dots + \sum_{n=1}^{N+1} \sum_{j=1}^J \log \sum_{l_{nj}=1}^L p(i_{nj}|\alpha_j^{\mathcal{I}}, l_{nj}) p(c_{nj}|l_{nj}) p(l_{nj}|\alpha_j^{\mathcal{L}}), \end{aligned} \quad (18)$$

where

$$\begin{aligned} p(\alpha_j^{\mathcal{L}}) &= \frac{1}{Z_\epsilon} \prod_{l=1}^L (\alpha_{jl}^{\mathcal{L}})^\epsilon, \\ p(\alpha_{jl}^{\mathcal{I}}) &= \mathcal{N}(\alpha_{jl}^{\mathcal{I}}; \mu_0, \sigma^2/\epsilon), \\ p(i_{nj}|\alpha_j^{\mathcal{I}}, l_{nj}) &= \mathcal{N}(i_{nj}; \alpha_{jl_{nj}}^{\mathcal{I}}, \sigma^2), \\ p(c_{nj}|l_{nj}) &= \delta[f_n(l_{nj}) = c_{nj}], \\ p(l_{nj}|\alpha_j^{\mathcal{L}}) &= \alpha_{jl_{nj}}^{\mathcal{L}}. \end{aligned}$$

To solve the optimization problem of Equation 18, we use an EM algorithm: we iteratively build a lower bound to the objective function of Equation 18 that touches it at the current estimate of $(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}})$ (E step), and then optimize this bound with respect to $(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}})$ (M step).

In the **E step**, we make a soft assignment of each label l at the fine level of detail to each voxel j in each atlas n :

$$\tilde{W}_{nj}^l = \frac{\mathcal{N}(i_{nj}; \alpha_{jl}^{\mathcal{I}}, \sigma^2) \alpha_{jl}^{\mathcal{L}} \delta[c_{nj} = f_n(l)]}{\sum_{l'} \mathcal{N}(i_{nj}; \alpha_{jl'}^{\mathcal{I}}, \sigma^2) \alpha_{jl'}^{\mathcal{L}} \delta[c_{nj} = f_n(l')]} \quad (19)$$

Note that, if atlas n is labeled at the fine level already, then $(f_n)^{-1}(c_{nj})$ is unique and $\tilde{W}_{nj}^l = \delta[(f_n)^{-1}(c_{nj}) = l]$. These soft assignments are used to form the lower bound:

$$\begin{aligned} Q(\alpha^{\mathcal{L}}, \alpha^{\mathcal{I}}; \tilde{\alpha}^{\mathcal{L}}, \tilde{\alpha}^{\mathcal{I}}) &= -J \log Z_\epsilon - \sum_{n=1}^{N+1} \sum_{j=1}^J \sum_{l=1}^L \tilde{W}_{nj}^l \log \tilde{W}_{nj}^l + \dots \\ &\quad \dots + \sum_{j=1}^J \sum_{l=1}^L \left[\epsilon \log \alpha_{jl}^{\mathcal{L}} + \log \mathcal{N}(\alpha_{jl}^{\mathcal{I}}; \mu_0, \sigma^2/\epsilon) \right] + \dots \\ &\quad \dots + \sum_{n=1}^{N+1} \sum_{j=1}^J \sum_{l=1}^L \tilde{W}_{nj}^l \log \left[\mathcal{N}(i_{nj}; \alpha_{jl}^{\mathcal{I}}, \sigma^2) \delta[f_n(l) = c_{nj}] \alpha_{jl}^{\mathcal{L}} \right], \end{aligned} \quad (20)$$

where $\tilde{\alpha}^{\mathcal{L}}$ and $\tilde{\alpha}^{\mathcal{I}}$ are the current estimates of the parameters.

The **M step** updates can be derived as:

$$\alpha_{jl}^{\mathcal{I}} = \frac{\epsilon \mu_0 + \sum_{n=1}^{N+1} \tilde{W}_{nj}^l i_{nj}}{\epsilon + \sum_{n'=1}^{N+1} \tilde{W}_{n'j}^l}, \quad (21)$$

$$\alpha_{jl}^{\mathcal{L}} = \frac{\epsilon + \sum_{n=1}^{N+1} \tilde{W}_{nj}^l}{\epsilon L + N + 1}. \quad (22)$$

Once the algorithm has converged, we can substitute the point estimates $\hat{\alpha}^{\mathcal{L}}$ and $\hat{\alpha}^{\mathcal{T}}$ back into Equation 16 and use Bayes's rule to compute the segmentation. It is straightforward to show that the approximate posterior label probabilities of the voxels – which are independent of each other – are given by:

$$p(s_j | \hat{\alpha}_j^{\mathcal{L}}, \hat{\alpha}_j^{\mathcal{T}}, y_j, i_{\cdot,j}, c_{\cdot,j}) = \tilde{W}_{N+1,j}^{s_j}. \quad (23)$$

Therefore, the optimal discrete segmentation is:

$$\hat{s}_j \approx \underset{s}{\operatorname{argmax}} \tilde{W}_{N+1,j}^s, \quad (24)$$

and the expectation of the volume of label l is (in voxels):

$$E[V_l] = \sum_{j=1}^J \tilde{W}_{N+1,j}^l. \quad (25)$$

3. Experiments and results

3.1. MRI data

We used four different datasets of manually labeled T1 MRI scans in this study (see sample slices in Figure 6):

- FreeSurfer dataset: 39 T1-weighted, 1 mm isotropic scans with 36 cortical and subcortical labels (see delineation protocol in [30]). The cerebral cortex and white matter are considered single entities. We note that these are the subjects that were used to train the probabilistic atlas in FreeSurfer [31].
- Brainstem dataset: ten T1-weighted, 1 mm isotropic scans with manual labels for the medulla oblongata, pons and midbrain, i.e., the substructures of the brainstem. The delineation protocol is described in [32].
- Winterburn dataset: five 0.6 mm isotropic scans¹ with annotations of the hippocampal subfields – subiculum, CA1, CA23, CA4 and molecular layer. The acquisition and manual delineation of the data are described in [33]. The dataset includes T1-weighted and T2-weighted scans; only the T1-weighted volumes were used here.
- Hammers dataset²: 20 T1-weighted, 1 mm isotropic scans with 67 labels of cortical and subcortical structures. The labels of the cortical structures do not separate white from gray matter, and there is a single cerebellar label, which groups its gray and white matter. Further details on the dataset and on the manual labeling protocol can be found in [34] and [35].

The scans from all four datasets have fields of view covering the whole brain. Additional details on the acquisition can be obtained from the corresponding publications.

¹In the original publication they work with 0.3 mm upsampled data, but the native resolution is 0.6 mm.

²www.brain-development.org. ©Copyright Imperial College of Science, Technology and Medicine 2007. All rights reserved.

3.2. Definition of fine labels and protocol functions

In the fusion, we defined a set of 148 labels given by all the possible intersections of regions defined in the four datasets; note that this process generates more regions than the total number of unique labels in the four datasets (102). Specifically, we defined two labels for each lobe in the Hammers dataset, one for the corresponding white matter and one for the corresponding cortex, based on their intersections with the gray and white matter in the FreeSurfer dataset. We also defined a number of labels to cope with differences in labeling protocols of the same structure; even though the protocols are in general very similar to each other, there are two exceptions. First, the Hammers hippocampus does not include the tail, while the FreeSurfer and Winterburn protocols do. To cope with this, we split up each of the five hippocampal subfields into an anterior (head/body) and a posterior (tail) region; this yields a total of 20 hippocampal labels – 10 per side. Second, the midbrain of the Brainstem dataset coincides with the superior part of the brainstem in the Hammers dataset, but extends further in the superior direction than the brainstem in FreeSurfer. To model this difference, we split the midbrain into an inferior and a superior part; the latter is further split into left and right. The midbrain label in the Brainstem dataset and the brainstem in the Hammers dataset include all three regions, while the FreeSurfer dataset considers the inferior region part of the brainstem and the superior regions part of the left and right diencephala, respectively.

Given these region definitions, we specified four unique protocol functions f (one per dataset, as illustrated in Figure 7):

- FreeSurfer dataset: the protocol collapses all cortical structures into two generic cortex labels (left and right), all white matter structures (including the left and right corpus callosum, only defined in the Hammers dataset) into two white matter regions, all the hippocampal subfields into whole hippocampi, and the brainstem labels into a whole brainstem region – except for the superior midbrain regions, which are collapsed together with the diencephala.
- Brainstem dataset: the protocol collapses all the non-brainstem regions into a single background label.
- Winterburn dataset: the protocol collapses all the non-hippocampal regions into a generic background label.
- Hammers dataset: the protocol collapses the white and gray matter labels of each lobe into a single label, all the brainstem and hippocampal regions into two generic labels, and the white matter and cortex of the cerebellum into a single structure.




Figure 6: Sample sagittal slices of the four datasets used in this study, with manual annotations overlaid: (a) FreeSurfer, (b) Brainstem, (c) Winterburn, (d) Hammers dataset. We only show a region of interest around the labels in (b,c) – the scans cover the whole brain.




Figure 7: Protocol functions: (a) sagittal slice of a segmentation at the fine label level; (b) effect of applying protocol function corresponding to FreeSurfer dataset; (c) Brainstem dataset; (d) Winterburn dataset; (e) Hammers dataset. It is the goal of this study to produce segmentations such as (a) by merging segmentations like (b-e).

3.3. Experimental setup

All 74 brain scans were resampled to 1 mm isotropic resolution, skull stripped, bias field corrected and intensity normalized with FreeSurfer[31]. The intensity normalization is necessary because it enables us to directly compare image intensities (even across datasets), which is a critical assumption in semi-locally weighted fusion and MPLF. The scans were then pairwise registered with the software package ANTS [8]). We used the default parameters for the deformation model (SyN[0.25]), number of iterations ($30 \times 90 \times 20$) and cost function (neighborhood cross correlation). However, we increased the radius of the window of the cost function from 4 to 6, and the regularization parameter from 3 to 4; these values better coped with the differences in intensity profile between datasets that cannot be completely eliminated by the intensity normalization step.

The four proposed label fusion schemes were used to segment the scans in a leave-one-out manner, where the test subject was left out from the atlas set $\{\mathbf{i}_n, \mathbf{l}_n\}$. The parameter setting and initialization was as follows. In semi-locally weighted fusion, we initialized $\sigma^2 = 100$, and set $\beta = 0.75$ – as in[17]. In STAPLE, we initialized each column of each Θ^n such that 0.95 of the probability mass

is equally distributed among the coarse labels that are compatible with the fine label corresponding to that column; the remaining 0.05 is equally spread across the non-compatible coarse labels. In MPLF, we set $\epsilon = 10^{-6}$, $\mu_0 = 65$ (which is the typical intensity of gray matter after normalization) and $\sigma^2 = 100$ (again, inspired by[17]).

We evaluated the label fusion methods in two different ways: directly with Dice scores, and indirectly through an aging experiment. First, we computed Dice scores between the manual delineations and the protocol-transformed, automated segmentations produced by the different label fusion methods, i.e., we compared $\mathbf{c}_{tj}$ with $f_t(\hat{\mathbf{s}}_j)$ (where we use the subscript t to refer to the test scan) for all four datasets. In addition, we also evaluated the performance of MPLF against the number of training scans. In this experiment, we used the FreeSurfer dataset, which has the most scans and therefore provides the largest range for the analysis. We computed segmentations of the FreeSurfer scans using N_{fs} other randomly selected FreeSurfer scans (“fs” stands for FreeSurfer) and a pool of (randomly selected) scans from the other datasets, such that the proportion of atlases from each of the datasets in the pool was approximately constant (see Table 4). This ensures that the performance depends on the number of atlases,

FreeSurfer (N_{fs})	1	3	5	8	16	23	31	37
Brainstem	1	1	2	2	4	6	8	10
Winterburn	1	1	1	1	2	3	4	5
Hammers	1	2	3	4	8	12	16	20

Table 4: Number of atlases from each dataset in the training pool for the experiment testing the performance of MPLF against the number of training scans.

Dataset	Maj. vot.	SL weight.	G-STAPLE	MPLF
FreeSurfer	80.6	82.6	82.6	86.5
Brainstem	85.2	85.6	74.2	86.9
Winterburn	43.1	46.8	36.4	50.5
Hammers	75.1	78.6	78.2	81.6
Combined	77.0	79.8	79.1	83.1

Table 5: Mean Dice score (in %) across all structures within each dataset, as well as for all structures from all datasets combined. “Maj. vot.” represents generalized majority voting, “SL weight.” represents generalized semi-locally weighted fusion, and “G-STAPLE” represents generalized STAPLE.

rather than the proportion of scans from the FreeSurfer dataset in the pool of training images. The experiment was repeated 10 times for each N_{fs} with different random selections of scans except for the cases $N_{fs} = 1$, where we used all 38 left-out scans in the dataset, and $N_{fs} = 38$, where there is only one possible combination scans (since we are using all the available atlases).

In a second set of experiments, we evaluated MPLF indirectly by analyzing the effect of age on the median thickness of each cortical region, as well as on the volume of the different subcortical structures. This analysis was carried at the fine level of label detail, so we measured volumes and thicknesses of structures that were not defined in any of the training datasets. Both the thicknesses and the volumes were computed from the soft segmentations: the volumes were computed with Equation 25, and the thicknesses were estimated from the label posteriors (Equation 23) using the algorithm described in [36]. This technique estimates the thickness at a point of interest by minimizing the line integral over the probability map of the gray matter on line segments centered at that point. For each cortical region, we took the median (a robust estimate) of the thicknesses given by this method for the voxels belonging to that region, as estimated by the hard segmentation from Equation 24. Note that both the estimation of the volume and of the cortical thickness from soft segmentations require faithful posteriors; probability maps like the ones showed in Figure 4c or 4c would yield unrealistic estimates.

The aging experiment was performed on the FreeSurfer dataset, which has the most subjects and the widest age range (53.3 ± 23.3 years). For each subcortical brain structure, we first fitted a generalized linear model (GLM) predicting its volume as a linear combination of the age of the subject, his intracranial volume (as estimated by FreeSurfer) and a bias. Then, for each cortical structure, we fitted a




Figure 8: Mean Dice score (in %) for the FreeSurfer dataset as a function of the number of training scans of the FreeSurfer dataset N_{fs} . The pool of training scans also includes atlases from the other three datasets, such that the ratio of scans from the different datasets is approximately constant (see Table 4). The Dice scores are averaged over all structures and ten random selections of scans (except for $N_{fs} = 1$, where we used all 38 left-out scans, and $N_{fs} = 38$, where there is only one possible combination of atlases).

GLM predicting its median thickness as a linear combination of the age of the subject and a bias. Finally, a statistical t-test was used to assess whether the coefficient related to age in each GLM was significantly different from zero. In order to increase the power of the analysis in the subcortical structures, we left-right averaged their volumes – for the median cortical thickness this is not as advantageous, since it is a robust estimate already.

3.4. Results

3.4.1. Direct validation: Dice scores

Table 5 shows the mean Dice scores across the structures defined within each dataset. Generalized STAPLE produces very variable results, yielding excellent segmentation for some structures but poor outputs for others (e.g., brainstem and Winterburn datasets). On average, it outperforms generalized majority voting by 2% Dice. Generalized semi-locally weighted voting takes advantage of the image intensities of the deformed atlases to yield an average Dice 1% higher than that of generalized STAPLE. MPFL clearly outperforms all the other methods by communicating information between the atlases: its average Dice is 3% higher than that of the second best method (generalized semi-locally weighted voting). Figure 8 shows the mean Dice score produced by MPFL in the FreeSurfer dataset as a function of the number of training scans. The plot shows that MPFL only requires 3 atlases to yield Dice scores similar to those of generalized semi-locally weighted




Figure 9: Box plots for the Dice scores corresponding to the four datasets. The central mark is the median, the box spans from the first to the third quartile, and the whiskers span the extreme data points not considered outliers (which are marked with red crosses).




Figure 10: (a) Sagittal slice of a scan from the FreeSurfer dataset. (b) Corresponding manual segmentation. (c) Automated segmentation with generalized majority voting. (d) Generalized semi-locally weighted segmentation. (e) Generalized STAPLE. (f) MPLF. (g) 3D rendering of the pial surface using MPLF. (h) Rendering of the white matter surface.

voting with 38 atlases. The performance of MPFL saturates at approximately 30 scans.

Figure 9 shows box plots for the Dice scores between the manual and automated segmentations for each structure in the four datasets. In the FreeSurfer dataset, generalized majority voting performs satisfactorily for most structures except for the cortex, which is difficult to register. Semi-local weighting provides a boost in the performance for some of the structures, particularly the caudate and the cortex, thanks to the use of image intensity information in the fusion. Generalized STAPLE produces very variable results: for some structures it outperforms semi-locally weighted fusion (diencephalon, amygdala, pallidum, putamen), but for others it produces rather poor results (most notably the cortex). MPLF successfully combines all the training data to produce the highest Dice scores for every structure other than the cerebellum.

In absolute terms, the results from MPLF in subcortical structures are slightly worse than those reported in the literature by state-of-the-art fusion algorithms working in single-protocol settings (e.g., [17]). This is not caused by shortcomings of our algorithm, but rather by the fact that it produces a segmentation that “averages” subtle differences in labeling protocols that are not explicitly modeled in our fusion framework. Therefore, the automatic segmentation does not necessarily agree perfectly with any of the individual protocols, and is penalized when computing Dice scores against the ground truth segmentations of the

individual datasets. For example, this is the reason – in addition to differences in registrations – for the discrepancy between our results for generalized majority voting in the Hammers dataset and the results reported in [10] using the same brain scans.

In cortical structures, MPLF shares the limitation of other label fusion algorithms (and registration based segmentation methods in general), that volumetric registration of the cerebral cortex is extremely difficult. The cortical segmentation could possibly be improved by replacing ANTs with a registration algorithm specifically designed for the cortex, e.g., [37].

The same trends that were observed in the FreeSurfer dataset apply to the Brainstem, Winterburn and Hammers datasets. Semi-locally weighted voting provides a small improvement over majority voting, STAPLE yields very variable results (rather poor, in some cases) and MPLF outperforms the other three for all the structures of interest, other than CA1 in the hippocampus and the cerebellum (and some lobes/gyri) in the Hammers dataset. In absolute terms, the Dice scores are high in the Brainstem dataset – though difficult to place in a global context due to the lack of midbrain, pons and medulla segmentation algorithms in the literature. On the other hand, they are low for the hippocampal subfields, due to their thin shapes (all but CA4) and the insufficient resolution to segment them accurately. In the Hammers dataset, Dice scores of the subcortical structures are one notch below the results

from the FreeSurfer dataset; this is possibly due to the fact that the averaging in labeling protocols is skewed towards the FreeSurfer data due to the larger presence in the training dataset – twice as many.

Figure 10 shows segmentations and 3D renderings of a sample test scan from the FreeSurfer dataset. The automated segmentations in Figure 10c-f are much richer than the manual labels in Figure 10b. For instance, the parcelated pial and white matter surfaces (Figure 10g-h) could not have been generated using any of the training datasets independently). Since they do not consider image intensities in the fusion, the generalizations of majority voting and STAPLE do not correctly segment the cortex, which is difficult to register. They also oversegment structures such as the pallidum. Generalized semi-locally weighted fusion ameliorates these problems through the use of intensities – particularly the cortical segmentation. However, it is outperformed by MPLF, which produces more accurate segmentations for the thalamus, pallidum, putamen and choroid plexus – while providing meaningful estimates of the label posteriors.

3.4.2. Indirect validation: aging study

The association between morphometric measurements and age (Table 6) shows strong consistency with prior work using other procedures. For example, we found strong associations between age and cortical thickness in frontal lobe regions including the superior frontal and precentral gyrus, weaker association with medial temporal regions, and strong reductions in volume of the thalamus with more moderate reductions in volume of the hippocampus [3, 38]. Even though the resolution of the scans was not sufficient to clearly distinguish among hippocampal subfields, we observed a larger effect of aging on CA1 and CA4-DG, consistent with prior work [39]. The results on the brainstem also showed strong agreement with prior studies: the volume of the midbrain steadily decreases with age, the pons is spared, and the medulla suffers from minimal atrophy [40].

4. Conclusion and discussion

In this paper, we have generalized three popular label fusion methods to scenarios where the atlases have been manually traced with different protocols. We have discussed the limitations of the generalized methods and proposed MPLF, an alternative algorithm that was shown to outperform them: the average Dice score on four datasets improved between 3% and 6% with respect to the generalizations. The core of both the generalized methods and MPLF is the definition of protocol functions that group sets of fine labels (hidden) into coarser labels (observed). Adding the coarse labels to the corresponding graphical models, the generalization of the three existing methods to the multi-protocol scenario is straightforward.

The extensions of majority voting, semi-locally weighted fusion and STAPLE are easy to implement and require

very few changes with respect to the original algorithms. On the other hand, MPLF is computationally more expensive, since it requires optimizing a function with the EM algorithm at each voxel. In any case, the computational cost of the fusion (approximately one hour in our experiments) is small compared with the cost of nonlinearly registering the atlases (approximately one and a half hours per atlas, tens of hours in total), and the algorithm can be easily parallelized if necessary – since the voxels of the input scan are processed independently.

Future work will follow three directions. First, we will generalize newer, more sophisticated label fusion algorithms and compare them with MPLF. In particular, it will be interesting to consider extensions of STAPLE that support soft labels and spatially varying performance parameters. Second, we will consider the possibility of placing smoothness priors on the intensities and label probabilities of the statistical atlas in MPLF, as well as on the segmentation of the test scan. Even though the automated segmentations were accurate and smooth in our experiments, smoothness constraints might be important when the number of atlases is not as high as in this study. And third, we will generalize MPLF to cross-modality scenarios, which will introduce the capability to handle microscopic images (e.g., BigBrain³ [41]) or optical coherence tomography [42], in order to model with very fine detail brain areas that are not visible with MRI.

As the amount of publicly available, heterogeneously labeled data continues to grow, we believe that segmentation methods that can cope with different protocols – such as the one we have described – will become increasingly important.

Acknowledgements

Support for this research was provided in part by the National Center for Research Resources (U24 RR021382, P41-RR14075, 1KL2RR025757-01), the National Institute for Biomedical Imaging and Bioengineering (P41-EB015896, R01EB006758, R01EB013565, 1K25EB013649-01), the National Institute on Aging (AG022381, 5R01AG008122-22, R01 AG016495-11), the National Center for Alternative Medicine (RC1 AT005728-01), the national Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963, R01NS083534), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. This research was also funded by TEKES (ComBrain), Harvard Catalyst, and financial contributions from Harvard and affiliations. JEI was supported by the

³<https://bigbrain.loris.ca>

(a)		(b)		
Structure	$-\log_{10} p$	Cortical region	$-\log_{10} p$ (left side)	$-\log_{10} p$ (right side)
Lat.Ventricle	7.51	Ant. temp. lobe medial	4.14	4.97
Cerebellum WM	0.43	Ant. temp. lobe lateral	8.30	7.91
Cerebellum CT	2.37	Parahip. & ambient gyri	0.60	1.66
Thalamus	13.70	Sup. temp. gyrus	9.26	11.21
Caudate	0.72	Middle & inf. temp. gyri	4.15	4.86
Putamen	6.8	Fusiform gyrus	1.85	2.94
Pallidum	6.56	Insula	8.90	9.49
Whole hippo.	3.72	Lat. occipital lobe	4.41	4.92
Amygdala	3.98	Gyrus cinguli ant. part	7.04	8.10
Accumbens	3.99	Gyrus cinguli post. part	9.10	9.84
Diencephalon	10.27	Middle front. gyrus	11.69	10.94
3rd ventricle	7.10	Post. temp. lobe	7.25	8.69
4th ventricle	0.59	Inferolat. pariet. lobe	12.71	10.61
Corp.Call.	0.29	Precentral gyrus	11.30	10.37
Whole brainst.	0.70	Gyrus rectus	1.367	0.84
Cerebral WM	1.13	Orbitofrontal gyri	8.82	11.93
CA1	2.84	Inf. front. gyrus	11.73	12.73
CA23	2.42	Sup. front. gyrus	11.87	13.03
CA4-DG	3.45	Postcentral gyrus	8.44	8.35
Subiculum	1.38	Sup. parietal gyrus	6.63	8.20
Molec. Layer	2.57	Lingual gyrus	6.10	6.34
Midbrain	6.14	Cuneus	5.33	7.54
Pons	0.25			
Medulla Obl.	2.68			

Table 6: Effect of age: p-values (given as $-\log_{10} p$) for t-test on significance of slope in the generalized linear model. (a) Volume of subcortical structures (left-right averaged). (b) Cortical thickness. Significant p values (below 0.05 after Bonferroni correction) are in bold. The threshold for significance at $\alpha = 0.05$ is, after Bonferroni correction, $-\log_{10} p \approx 2.65$.

Guipuzcoa Fellows program. In addition, BF has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. BF’s interests were reviewed and are managed by Massachusetts General Hospital and Partners HealthCare in accordance with their conflict of interest policies.

References

- [1] A. Yendiki, P. Panneck, P. Srinivasan, A. Stevens, L. Zöllei, J. Augustinack, R. Wang, D. Salat, S. Ehrlich, T. Behrens, S. Jbabdi, R. Gollub, B. Fischl, Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy, *Frontiers in neuroinformatics* 5 (23) (2011) 1–12.
- [2] H. E. Cline, W. E. Lorensen, R. Kikinis, F. Jolesz, Three-dimensional segmentation of MR images of the head using probability and connectivity, *Journal of computer assisted tomography* 14 (6) (1990) 1037–1045.
- [3] K. B. Walhovd, A. M. Fjell, I. Reinvang, A. Lundervold, A. M. Dale, D. E. Eilertsen, B. T. Quinn, D. Salat, N. Makris, B. Fischl, Effects of age on volumes of cortex, white matter and subcortical structures, *Neurobiology of aging* 26 (9) (2005) 1261–1270.
- [4] R. C. Knickmeyer, S. Gouttard, C. Kang, D. Evans, K. Wilber, J. K. Smith, R. M. Hamer, W. Lin, G. Gerig, J. H. Gilmore, A structural MRI study of human brain development from birth to 2 years, *The Journal of Neuroscience* 28 (47) (2008) 12176–12182.
- [5] L. De Jong, K. Van der Hiele, I. Veer, J. Houwing, R. Westendorp, E. Bollen, P. De Bruin, H. Middelkoop, M. Van Buchem, J. Van Der Grond, Strongly reduced volumes of putamen and thalamus in alzheimer’s disease: an MRI study, *Brain* 131 (12) (2008) 3277–3285.
- [6] T. Rohlfing, R. Brandt, R. Menzel, C. R. Maurer Jr, Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains, *NeuroImage* 21 (4) (2004) 1428–1442.
- [7] S. Awate, R. Whitaker, Multiatlas segmentation as nonparametric regression., *IEEE transactions on medical imaging* 33 (9) (2014) 1803–1817.
- [8] B. B. Avants, C. L. Epstein, M. Grossman, J. C. Gee, Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain, *Medical image analysis* 12 (1) (2008) 26–41.
- [9] S. Klein, M. Staring, K. Murphy, M. A. Viergever, J. P. Pluim, Elastix: a toolbox for intensity-based medical image registration, *Medical Imaging, IEEE Transactions on* 29 (1) (2010) 196–205.
- [10] R. Heckemann, J. Hajnal, P. Aljabar, D. Rueckert, A. Hammers, Automatic anatomical brain MRI segmentation combining label propagation and decision fusion, *NeuroImage* 33 (1) (2006) 115–126.
- [11] P. Aljabar, R. A. Heckemann, A. Hammers, J. V. Hajnal, D. Rueckert, Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy, *Neuroimage* 46 (3) (2009) 726–738.
- [12] A. K. H. Duc, M. Modat, K. K. Leung, M. J. Cardoso, J. Barnes, T. Kadir, S. Ourselin, A. D. N. Initiative, et al., Using manifold learning for atlas selection in multi-atlas segmentation, *PloS one*

- 8 (8) (2013) e70059.
- [13] X. Artaechevarria, A. Muñoz-Barrutia, C. Ortiz-de Solorzano, Efficient classifier generation and weighted voting for atlas-based segmentation: Two small steps faster and closer to the combination oracle, in: Medical Imaging, International Society for Optics and Photonics, 2008, pp. 69141W–69141W.
 - [14] I. Isgum, M. Staring, A. Rutten, M. Prokop, M. Viergever, B. van Ginneken, Multi-atlas-based segmentation with local decision fusion – application to cardiac and aortic segmentation in CT scans, *IEEE Trans.Med.Im.* 28 (7) (2009) 1000–1010.
 - [15] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, D. L. Collins, Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation, *NeuroImage* 54 (2) (2011) 940–954.
 - [16] H. Wang, J. W. Suh, S. R. Das, J. B. Pluta, C. Craige, P. A. Yushkevich, Multi-atlas segmentation with joint label fusion, *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 35 (3) (2013) 611–623.
 - [17] M. Sabuncu, B. Yeo, K. Van Leemput, B. Fischl, P. Golland, A generative model for image segmentation based on label fusion, *IEEE Trans. Med. Im.* 29 (2010) 1714–1729.
 - [18] S. K. Warfield, K. H. Zou, W. M. Wells, Simultaneous truth and performance level estimation (staple): an algorithm for the validation of image segmentation, *Medical Imaging, IEEE Transactions on* 23 (7) (2004) 903–921.
 - [19] A. J. Asman, B. A. Landman, Formulating spatially varying performance in the statistical fusion framework, *Medical Imaging, IEEE Transactions on* 31 (6) (2012) 1326–1336.
 - [20] A. J. Asman, B. A. Landman, Non-local statistical label fusion for multi-atlas segmentation, *Medical image analysis* 17 (2) (2013) 194–208.
 - [21] M. Jorge Cardoso, K. Leung, M. Modat, S. Keihaninejad, D. Cash, J. Barnes, N. C. Fox, S. Ourselin, Steps: Similarity and truth estimation for propagated segmentations and its application to hippocampal segmentation and brain parcellation, *Medical image analysis* 17 (6) (2013) 671–684.
 - [22] A. Akhondi-Asl, S. Warfield, Simultaneous truth and performance level estimation through fusion of probabilistic segmentations., *IEEE transactions on medical imaging* 32 (10) (2013) 1840.
 - [23] B. Landman, J. Bogovic, J. Prince, Efficient anatomical labeling by statistical recombination of partially label datasets, in: *Proc. of ISMRM*, 2009, p. 269.
 - [24] B. A. Landman, J. A. Bogovic, J. L. Prince, Simultaneous truth and performance level estimation with incomplete, overcomplete, and ancillary data, in: *SPIE Medical Imaging, International Society for Optics and Photonics*, 2010, pp. 76231N–76231N.
 - [25] B. A. Landman, A. J. Asman, A. G. Scoggins, J. A. Bogovic, F. Xing, J. L. Prince, Robust statistical fusion of image labels, *Medical Imaging, IEEE Transactions on* 31 (2) (2012) 512–522.
 - [26] O. Commowick, A. Akhondi-Asl, S. K. Warfield, Estimating a reference standard segmentation with spatially varying performance parameters: Local map staple, *Medical Imaging, IEEE Transactions on* 31 (8) (2012) 1593–1606.
 - [27] A. P. Dempster, N. M. Laird, D. B. Rubin, et al., Maximum likelihood from incomplete data via the em algorithm, *J. R. Stat. Soc.* 39 (1) (1977) 1–38.
 - [28] J. E. Iglesias, M. R. Sabuncu, K. Van Leemput, Improved inference in bayesian segmentation using monte carlo sampling: Application to hippocampal subfield volumetry, *Medical image analysis* 17 (7) (2013) 766–778.
 - [29] J. Iglesias, M. Sabuncu, K. Van Leemput, A generative model for multi-atlas segmentation across modalities, in: *9th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2012, pp. 888–891.
 - [30] V. Caviness Jr, P. Filipek, D. Kennedy, Magnetic resonance technology in human brain science: blueprint for a program based upon morphometry, *Brain Dev.* 11 (1) (1989) 1–13.
 - [31] B. Fischl, D. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, A. Dale, Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain, *Neuron* 33 (2002) 341–355.
 - [32] J. Iglesias, K. Van Leemput, P. Bhatt, C. Casillas, S. Dutt, A. Boxer, B. Fischl, Bayesian segmentation of brainstem structures in MRI, manuscript submitted for publication.
 - [33] J. L. Winterburn, J. C. Pruessner, S. Chavez, M. M. Schira, N. J. Lobaugh, A. N. Voineskos, M. M. Chakravarty, A novel in vivo atlas of human hippocampal subfields using high-resolution 3T magnetic resonance imaging, *Neuroimage* 74 (2013) 254–265.
 - [34] A. Hammers, R. Allom, M. J. Koepp, S. L. Free, R. Myers, L. Lemieux, T. N. Mitchell, D. J. Brooks, J. S. Duncan, Three-dimensional maximum probability atlas of the human brain, with particular reference to the temporal lobe, *Human brain mapping* 19 (4) (2003) 224–247.
 - [35] I. S. Gousias, D. Rueckert, R. A. Heckemann, L. E. Dyet, J. P. Boardman, A. D. Edwards, A. Hammers, Automatic segmentation of brain MRIs of 2-year-olds into 83 regions of interest, *Neuroimage* 40 (2) (2008) 672–684.
 - [36] I. Aganj, G. Sapiro, N. Parikshak, S. K. Madsen, P. M. Thompson, Measurement of cortical thickness from MRI by minimum line integrals on soft-classified tissue, *Hum. brain mapp.* 30 (10) (2009) 3188–3199.
 - [37] G. Postelnicu, L. Zollei, B. Fischl, Combined volumetric and surface registration, *Medical Imaging, IEEE Transactions on* 28 (4) (2009) 508–522.
 - [38] D. H. Salat, R. L. Buckner, A. Z. Snyder, D. N. Greve, R. S. Desikan, E. Busa, J. C. Morris, A. M. Dale, B. Fischl, Thinning of the cerebral cortex in aging, *Cerebral cortex* 14 (7) (2004) 721–730.
 - [39] S. G. Mueller, M. W. Weiner, Selective effect of age, apo e4, and alzheimer’s disease on hippocampal subfields, *Hippocampus* 19 (6) (2009) 558–564.
 - [40] R. Raininko, T. Autti, S.-L. Vanhanen, A. Ylikoski, T. Erkinjuntti, P. Santavuori, The normal brain stem from infancy to old age, *Neurorad* 36 (5) (1994) 364–368.
 - [41] K. Amunts, C. Lepage, L. Borgeat, H. Mohlberg, T. Dickscheid, M.-É. Rousseau, S. Bludau, P.-L. Bazin, L. B. Lewis, A.-M. Oros-Peusquens, et al., Bigbrain: an ultrahigh-resolution 3d human brain model, *Science* 340 (6139) (2013) 1472–1475.
 - [42] C. Magnain, J. C. Augustinack, M. Reuter, C. Wachinger, M. P. Frosch, T. Ragan, T. Akkin, V. J. Wedeen, D. A. Boas, B. Fischl, Blockface histology with optical coherence tomography: A comparison with nissl staining, *NeuroImage* 84 (2014) 524–533.