



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Effects of registration regularization and atlas sharpness on segmentation accuracy

B.T. Thomas Yeo^{a,*}, Mert R. Sabuncu^{a,1}, Rahul Desikan^b, Bruce Fischl^{a,c,d,2}, Polina Golland^{a,2}^a Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA^b Department of Anatomy and Neurobiology, Boston University School of Medicine, Boston, MA, USA^c Department of Radiology, Harvard Medical School, Charlestown, MA, USA^d Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

Article history:

Received 31 January 2008

Received in revised form 9 May 2008

Accepted 10 June 2008

Available online 19 June 2008

Keywords:

Generative model

Registration

Segmentation

Parcellation

Multiple atlases

Markov random field

Regularization

ABSTRACT

In non-rigid registration, the tradeoff between warp regularization and image fidelity is typically determined empirically. In atlas-based segmentation, this leads to a probabilistic atlas of arbitrary sharpness: weak regularization results in well-aligned training images and a sharp atlas; strong regularization yields a “blurry” atlas.

In this paper, we employ a generative model for the joint registration and segmentation of images. The atlas construction process arises naturally as estimation of the model parameters. This framework allows the computation of unbiased atlases from manually labeled data at various degrees of “sharpness”, as well as the joint registration and segmentation of a novel brain in a *consistent* manner.

We study the effects of the tradeoff of atlas sharpness and warp smoothness in the context of cortical surface parcellation. This is an important question because of the increasingly availability of atlases in public databases, and the development of registration algorithms separate from the atlas construction process. We find that the optimal segmentation (parcellation) corresponds to a unique balance of atlas sharpness and warp regularization, yielding statistically significant improvements over the FreeSurfer parcellation algorithm. Furthermore, we conclude that one can simply use a single atlas computed at an optimal sharpness for the registration–segmentation of a new subject with a pre-determined, fixed, optimal warp constraint. The optimal atlas sharpness and warp smoothness can be determined by probing the segmentation performance on available training data. Our experiments also suggest that segmentation accuracy is *tolerant* up to a small mismatch between atlas sharpness and warp smoothness.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

In this work, we propose a generative model for the joint registration and parcellation of cortical surfaces. We formulate the atlas construction process as estimation of the generative model parameters. This provides a consistent framework for constructing the atlas as well as the registration and segmentation of a novel subject. We explore the effects of registration regularization on the atlas construction and the segmentation of a new image. We conclude that optimal segmentation corresponds to a unique balance of atlas sharpness and warp regularization. However, the segmentation accuracy is *tolerant* up to a small mismatch between atlas sharpness and warp smoothness.

1.1. Probabilistic atlases

Probabilistic atlases are powerful tools in segmentation (Ashburner and Friston, 2005; Collins et al., 1999; Desikan et al., 2006; Fischl et al., 2002, 2004; Pohl et al., 2006; Van Leemput et al., 1999). The simplest probabilistic segmentation atlas provides only the prior probability of labels at a spatial position and no information about the expected appearance of the image. One uses this type of atlas to initialize and to guide the segmentation of a new image (Van Leemput et al., 1999). More sophisticated algorithms deform the atlas to the new image and transfer label probabilities from the atlas to the image (Ashburner and Friston, 2005, 2006). These label probabilities are then combined with the intensity of the new image to produce the final segmentation. We note that the phrase “image intensities” is used in a generic sense to indicate all image-derived features, such as MR intensity or cortical geometry. In this approach, the relationship between labels and image intensities is estimated in the new image and not from the training images used to compute the atlas.

More complex probabilistic atlases provide statistics on the relationship between the segmentation labels and the image

* Corresponding author. Tel.: +1 617 253 4413.

E-mail addresses: ythomas@csail.mit.edu (B.T. Thomas Yeo), msabuncu@csail.mit.edu (M.R. Sabuncu), rahul@nmr.mgh.harvard.edu (R. Desikan), fischl@nmr.mgh.harvard.edu (B. Fischl), polina@csail.mit.edu (P. Golland).¹ Thomas Yeo and Mert Sabuncu contributed equally to this work.² Bruce Fischl and Polina Golland contributed equally to this work.

intensities (Desikan et al., 2006; Evans et al., 1993; Fischl et al., 2002, 2004; Pohl et al., 2006). The intensity model in the atlas is used to bring the atlas and new image into the same space. The label probabilities and intensity model are then used to segment the new image. Such a registration process can be done sequentially (Desikan et al., 2006; Fischl et al., 2002, 2004) or jointly (Pohl et al., 2006). Since this complex atlas contains more information than the simpler atlas, it can conceivably be a more powerful tool for segmentation. However, extra care is needed if the modalities of the new image and training images are different.

In this work, we focus on probabilistic segmentation atlases that model both labels and intensities. An initial step in probabilistic atlas computation is the spatial normalization of the training images. The features used for co-registering images are usually derived from the images themselves (Allasonnière et al., 2007; Bhatia et al., 2004; Fischl et al., 1999; Guimond et al., 2000; Joshi et al., 2004; Mazziotta et al., 1995; Paus et al., 1996; Twining et al., 2005) or from the segmentation labels (De Craene et al., 2004; Lorenzen et al., 2006; Van Leemput, 2006; Weisenfeld and Warfield, 2007; Yeo et al., 2007). After normalization to a common space, we can compute the variability in intensity values (and/or labels) in the atlas space. Our approach uses both labels and intensities to co-register the training images. The need to utilize both labels and intensities arises naturally from the proposed generative model.

Joint registration–segmentation algorithms are generally more effective than sequential registration–segmentation as registration and segmentation benefit from additional knowledge of each other (Ashburner and Friston, 2005; Pohl et al., 2006; Wyatt and Noble, 2003; Xiaohua et al., 2004, 2005; Yezzi et al., 2003). In our case, the requirement to jointly register and segment is also a natural consequence of the generative model.

1.2. Warp regularization

Spatial normalization of the training images can be achieved with different registration algorithms that vary in the flexibility of warps they allow. Both low-dimensional warps, e.g., affine (Pohl et al., 2006), and richer, more flexible warps can be employed, e.g., splines (Bhatia et al., 2004; Twining et al., 2005; Weisenfeld and Warfield, 2007), dense displacement fields (De Craene et al., 2004; Guimond et al., 2000; Fischl et al., 1999, 2004; Van Leemput, 2006) and velocity fields (Joshi et al., 2004; Lorenzen et al., 2006). More restricted warps yield blurrier atlases that capture inter-subject variability of structures, enlarging the basin of attraction for the registration of a new subject. However, label prediction accuracy is limited by the sharpness of the atlas. Recently, it has been shown that combining information on the warp flexibility and the residual image of the registration process improve the classification accuracy of schizophrenic and normal subjects (Makrogiannis et al., 2006). Finding the optimal warp regularization tradeoff has received attention in recent years. Twining et al. (2005) and Van Leemput (2006) propose frameworks to find the least complex models that explain the image intensity and segmentation labels of the training images.

In contrast, we explicitly explore the relationship between the warp flexibility and the atlas sharpness in the context of atlas-based segmentation. Rather than finding one optimal warp regularization when computing the atlas, we compute atlases with varying degrees of warp regularization and thus obtain atlases of varying degrees of sharpness. We use these to study the effect of atlas sharpness and warp regularization on segmentation accuracy. In particular, we compare three specific approaches: (1) progressive (i.e., using increasingly flexible warps) registration–segmentation of a new brain with increasingly sharp atlases; (2) progressive registration with an atlas of a particular sharpness; (3) registration with an atlas of a particular sharpness using a pre-determined, fixed constraint on the warp.

1.3. Cortical parcellation

We investigate the question of regularization in the context of automatic parcellation of cortical surfaces. Automatic labeling of surface models of the cerebral cortex is important for identifying regions of interest for clinical, functional and structural studies (Desikan et al., 2006; Rivière et al., 2002). In particular, it has been shown that cortical surface registration can increase the power of functional alignment or activation consistency (Fischl et al., 1999), as well as the alignment of cytoarchitectonics (Fischl et al., 2007). Recent efforts range from identification of sulcal or gyral ridge lines (Tao et al., 2002; Tu et al., 2007) to segmentation of sulcal or gyral basins (Desikan et al., 2006; Fischl et al., 2004; Klein and Hirsch, 2005; Lohmann and Von Cramon, 2000; Mangin et al., 1994; Rettmann et al., 2002; Rivière et al., 2002). Similar to these prior studies, we are interested in the parcellation of the entire cortical surface, i.e., the automated labeling of each point of the surface of a new subject.

Because the local geometries of different sulci and gyri might be similar, learning the geometries of the different sulcal and gyral labels using information from a new image only is difficult. Instead, we will follow the approach of (Fischl et al., 2004; Desikan et al., 2006) and utilize an atlas that encodes both the spatial distributions of the labels and their spatially varying geometries.

In the context of cortical parcellation, our experiments suggest that the optimal parcellation in all three registration schemes for exploring atlas sharpness and warp regularization corresponds to a unique balance of atlas sharpness and warp regularization. Furthermore, the optimal parameters yield statistically significant improvements over the FreeSurfer parcellation algorithm (Fischl et al., 2004). There is no difference in the optimal parcellation accuracy achieved by the three schemes.

While one might expect that outlier images require weaker regularization to warp closer to the population “average” to achieve better segmentation accuracy, our experiments show that it is sufficient to use a single atlas computed at an optimal sharpness for the registration–segmentation of a new subject with a pre-determined, fixed, optimal warp constraint.

The optimal atlas sharpness and warp smoothness can be determined by probing the segmentation performance on available training data. We find that the optimal warp smoothness for the new subject should be the same as the warp smoothness used for creating the atlas. While this might be an obvious consequence of our explicit generative model, in practice, it is common to use or develop a registration algorithm separate from atlas construction. This is especially true with the increasing availability of atlases in publicly available databases, such as MNI305 (Evans et al., 1993). It is therefore important to determine whether using different warp smoothness and atlas sharpness is necessarily detrimental to segmentation in practice. Our experiments suggest that segmentation accuracy is tolerant up to a small mismatch between atlas sharpness and warp smoothness.

A preliminary version of this work was published at the International Conference on Medical Image Computing and Computer Assisted Intervention (Yeo et al., 2007). This article expands the conference paper with a more detailed theoretical development and more extensive experimental work. In particular, the theory for the atlas construction process was only briefly discussed in (Yeo et al., 2007), but is covered in depth in this paper.

Our contributions are as follows:

- (1) We propose a generative model for the joint registration and segmentation of images. The atlas construction process is formulated as a parameter estimation problem. This provides a consistent framework for both estimating the atlas, as well as for the registration and segmentation of a new image.

- (2) We explore the space of atlas sharpness and warp regularization for registering and parcellating cortical surfaces. We find that the optimal parcellation corresponds to a unique balance of atlas sharpness and warp regularization. The robustness of the optimal atlas sharpness and warp regularization across trials suggests that we can use the same parameters for future cortical surface parcellation.
- (3) We show improved parcellation results over the state-of-the-art FreeSurfer parcellation algorithm (Fischl et al., 2004).

In the next section, we introduce the generative model, describe the resulting atlas estimation process, and the registration and segmentation of a novel image. Section 3 introduces the cortical surface parcellation problem and describes further modeling assumptions made in the context of this problem. We present experimental results in Section 4.

2. Theory

Given a training set of N images $I_{1:N} = \{I_1, \dots, I_N\}$ with label maps $L_{1:N} = \{L_1, \dots, L_N\}$, joint registration–segmentation aims to infer the registration parameters R and segmentation labels L of a new image I . To achieve this goal, we first learn the parameters of the generative model (Fig. 1) from the training images. These parameters correspond to the atlas A and smoothness parameter S . Here, we assume the generative model for the warp R is parameterized by S . According to the generative model, the parameters R and A are independent. However, conditioned on the observed image, R and A become dependent. Estimating the parameters S and A involves co-registration of the training images into a common space. We emphasize that our co-registration process uses both the labels and image intensities of the training images, as described in the next section.

2.1. Generative model for registration and segmentation

We consider the generative model of Fig. 1. L' is a label map in the atlas space generated probabilistically by the atlas A . For example, L' could be the tissue type at each MRI pixel, generated by a tissue probability map. Given the label map, the atlas then generates image I' . For example, at each pixel, we can generate an MR intensity conditioned on the tissue type and spatial location. Finally, we generate a random warp field R controlled by the smoothness parameter S . For example, we can generate a random displacement field, with neighboring displacements encouraged to be “close” by setting S to be the spacing of spline control points or the penalty in a cost function that discourages large or discontinuous deformation fields. The random warp R is then applied to the label map L' and image I' to create the observed label map L and observed image I , i.e., $I(R(x)) = I'(x)$




Fig. 1. Generative model for registration and segmentation. A is an atlas used to generate the label map L' in some universal atlas space. The atlas A and label map L' generate image I' . S is the smoothness parameter that generates random warp field R . This warp is then applied to the label map L' and image I' to create the label map L and the image I . We assume the label map L is available for the training images, but not for the test image. The image I is observed in both training and test cases.

and $L(R(x)) = L'(x)$. Thus a location x in the atlas space is mapped to a location $R(x)$ in the native (or raw) image space. We defer a detailed instantiation of the model to Section 3.

During co-registration, a small value of smoothness parameter S leads to less constrained warps, resulting in better alignment of the training images.³ This results in a sharper atlas. On the other hand, a larger smoothness parameter yields more regularized warps and a blurrier atlas.

2.2. Atlas building: estimating parameters of generative model

To estimate the parameters of the generative model, we maximize the likelihood of the observed images $I_{1:N}$ and $L_{1:N}$ over the values of the non-random smoothness parameter S and atlas A .

$$(S^*, A^*) = \operatorname{argmax}_{S, A} \log p(I_{1:N}, L_{1:N}; S, A). \quad (1)$$

$$= \operatorname{argmax}_{S, A} \log \int p(I_{1:N}, L_{1:N}, R_{1:N}; S, A) dR_{1:N}. \quad (2)$$

Here, $p(a; b)$ indicates the probability of random variable a parameterized by a non-random parameter b while $p(a|b)$ indicates the probability of random variable a conditioned on a random variable b .

In this case, we need to marginalize over the registration warps $R_{1:N}$, which is difficult. Previously demonstrated methods for atlas construction use various approximations to evaluate the integral. In this paper, we adopt the standard approximation that replaces the integral with the value of the integrand estimated at the maximum likelihood estimate of deformation. In contrast, (Richard et al., 2007) uses a sampling method while (Van Leemput, 2006) uses the Laplace approximation, which essentially models the distribution to be a Gaussian centered at the maximum likelihood estimate. It is unclear that these more complex methods, while theoretically more sound, lead to (practically) better approximation. Based on this approximation, we seek

$$(S^*, A^*, R_{1:N}^*) = \operatorname{argmax}_{S, A, R_{1:N}} \log p(I_{1:N}, L_{1:N}, R_{1:N}; S, A). \quad (3)$$

$$= \operatorname{argmax}_{S, A, R_{1:N}} \sum_{n=1}^N \log p(I_n, L_n, R_n; S, A). \quad (4)$$

$$= \operatorname{argmax}_{S, A, R_{1:N}} \sum_{n=1}^N [\log p(R_n; S) + \log p(I_n, L_n | R_n; A)]. \quad (5)$$

The second equality comes from the fact that the training images are independent of each other given the atlas A and smoothness parameter S . The last equality is implied by the independence relations specified by the graphical model in Fig. 1.

Optimizing the above expression yields the atlas A^* and smoothness parameter S^* . As mentioned before, a smaller value of the smoothness parameter S results in a sharper atlas. Since we are interested in how atlas sharpness affects segmentation accuracy, instead of estimating one single optimal S^* , we construct a series of atlases corresponding to different values of the regularization parameter S . In particular, we discretize S into a finite set $\{S_k\} = \{S_1 > S_2 > \dots > S_K\}$. For each value S_k , we seek the optimal atlas and set of deformations:

$$(A^*, R_{1:N}^*) = \operatorname{argmax}_{A, R_{1:N}} \sum_{n=1}^N [\log p(R_n; S_k) + \log p(I_n, L_n | R_n; A)]. \quad (6)$$

³ By a “better alignment of images”, we mean that the warped images look more similar, i.e., the similarity measure is improved. However, an improved similarity measure does not necessarily imply deformations with better label alignment. In fact, we show in the paper that the best segmentation occurs when warps are not overly flexible.

We shall refer to the atlas computed using a particular S_k as $A_{\alpha=S_k}$. We use alternating optimization to maximize Eq. (6). In each step, we fix the set of registration warps $R_{1:N}$ and estimate the atlas $A_{\alpha=S_k}$:

$$A_{S_k}^* = \operatorname{argmax}_{A_{S_k}} \sum_{n=1}^N \log p(I_n, L_n | R_n; A_{S_k}). \quad (7)$$

We then fix the atlas $A_{\alpha=S_k}$, and optimize the registration warps $R_{1:N}$ by optimizing each warp independently of others:

$$R_n^* = \operatorname{argmax}_{R_n} \log p(R_n; S_k) + \log p(I_n, L_n | R_n; A_{S_k}), \quad (8)$$

$$\forall n = 1, \dots, N.$$

This process is repeated until convergence. Convergence is guaranteed since this is essentially a coordinate-ascent procedure operating on a bounded function.

We can think of Eq. (8) as the atlas co-registration objective function by treating $\log p(I_n, L_n | R_n; A_{S_k})$ as the data fidelity function and $\log p(R_n; S_k)$ as the regularization term. We effectively register each image independently to the atlas $A_{\alpha=S_k}$. This iterated process of updating the atlas is similar to (Joshi et al., 2004), except we include training labels in the co-registration. We will show a concrete instantiation of this formulation in Section 3.

In practice, we first create atlas A_∞ based on a simple rigid-body registration and use it to initialize the atlas A_{S_1} , where S_1 is large enough such that the resultant warp is almost rigid. We then use atlas A_{S_1} to initialize the atlas A_{S_2} where $S_1 > S_2$, and so on. The result is a set of atlases $\{A_\alpha\} = \{A_{S_1} \dots A_{S_K}\}$. With enough samples, the finite set $\{S_k, A_{\alpha=S_k}\}$ accurately represents the underlying continuous space of atlases at different levels of sharpness.

2.3. Registration and segmentation of a new image

Given an atlas A and smoothness parameter S , the registration and segmentation of a new image I can be computed using a maximum-a-posteriori (MAP) estimate. This involves finding the mode of the posterior distribution of the new image labels L and registration R given the observed image I and atlas A and smoothness parameter S :

$$(L^*, R^*) = \operatorname{argmax}_{L, R} \log p(L, R | I; S, A). \quad (9)$$

$$= \operatorname{argmax}_{L, R} \log p(I, L, R; S, A). \quad (10)$$

$$= \operatorname{argmax}_{L, R} \log p(R; S) + \log p(I, L | R; A). \quad (11)$$

The second equality follows from the definition of the conditional probability. The last equality follows from the independence relations specified by the graphical model in Fig. 1. In prior work, this MAP approach is favored by some (Wyatt and Noble, 2003; Xiaohua et al., 2004, 2005), while others suggest estimating the MAP solution for the registration warp alone (Ashburner and Friston, 2005; Pohl et al., 2006):

$$R^* = \operatorname{argmax}_R p(I, R; S, A). \quad (12)$$

$$= \operatorname{argmax}_R \log \sum_L p(L, I, R; S, A), \quad (13)$$

and recovering the segmentation labels as a final step after recovering the optimal registration R^* . Prior work in joint registration and segmentation did not consider atlas construction in the same framework (Ashburner and Friston, 2005; Pohl et al., 2006; Wyatt and Noble, 2003; Xiaohua et al., 2004, 2005). Furthermore, S is usually set by an expert rather than estimated from the data.

We previously reported results based on the latter two-step approach (Yeo et al., 2007). In this version, we use the former MAP framework since it requires fewer assumptions for practical opti-

mization. As we show in Section 3, optimizing Eq. (11) using a soft-MAP coordinate-ascent approach using the Mean Field approximation (Jaakkola, 2000; Kapur et al., 1998) results in the same update rule used by our previously demonstrated method (Yeo et al., 2007).

The differences between the new image registration specified by Eq. (11) and the atlas co-registration objective function in Eq. (8) comes from the unavoidable fact that for the new image, the label map is not observed. To optimize Eq. (11), we use a coordinate-ascent scheme. In step t , we fix the registration parameters $R^{(t)}$ and estimate the labels $L^{(t+1)}$:

$$L^{(t+1)} = \operatorname{argmax}_L \log p(R^{(t)}; S) + \log p(I, L | R^{(t)}; A). \quad (14)$$

$$= \operatorname{argmax}_L \log p(I, L | R^{(t)}; A). \quad (15)$$

$$= \operatorname{argmax}_L \log p(L | I, R^{(t)}; A). \quad (16)$$

Eq. (16) optimizes the log posterior probability of the label map L given the image I , atlas A and current estimate of the registration parameters $R^{(t)}$. Next, we fix the label estimate $L^{(t+1)}$ and re-estimate the registration parameters $R^{(t+1)}$:

$$R^{(t+1)} = \operatorname{argmax}_R \log p(R; S) + \log p(I, L^{(t+1)} | R; A). \quad (17)$$

We can think of Eq. (17) as the new image registration objective function by treating $\log p(I, L^{(t+1)} | R; A)$ as the data fidelity term and $\log p(R; S)$ as the regularization term.

To maintain generality, we allow the use of a smoothness parameter S_k for the registration and segmentation of a new subject with an atlas A_α where S_k may not be equal to α . In other words, we can, for instance, compute an atlas using affine transformations, while using a flexible deformation model for the registration of a new subject. Strictly speaking, S_k should theoretically be equal to α in Eq. (11) from the perspective of probabilistic inference. Here, we examine whether using different warp smoothness and atlas sharpness is necessarily detrimental to segmentation in practice.

More specifically, we investigate three strategies for exploring the space of atlas sharpness and warp smoothness to register and segment a new image as illustrated in Fig. 2.

- (1) *Multiple atlas, multiple smoothness (MAMS)*: A multiscale approach where we optimize Eq. (16) and (17) w.r.t. R and L with the blurry atlas A_{S_1} and warp regularization S_1 . The resulting alignment is used to initialize the registration with a sharper atlas A_{S_2} and a correspondingly flexible warp regularization S_2 , and so on (Fig. 2a).
- (2) *Single atlas, multiple smoothness (SAMS)*: We use a fixed atlas of a fixed sharpness scale A_{S_k} to compute R and L according to Eq. (16) and (17) using a progressively decreasing warp smoothness S (Fig. 2b).
- (3) *Single atlas, single smoothness (SASS)*: We optimize Eq. (16) and (17) w.r.t. R and L with a fixed atlas A_{S_k} and warp regularization S_m , where S_k might not be equal to S_m (Fig. 2c).

So far, the derivations have been general without any assumption about the atlases A_α , the prior $p(R; S_k)$ or the image-segmentation fidelity $p(I, L | R; A_\alpha)$. In the next section, we instantiate this approach for a concrete example of cortical surface registration and parcellation.

3. Cortical surface parcellation

We now demonstrate the framework developed in the previous section for the joint registration and parcellation of surface models

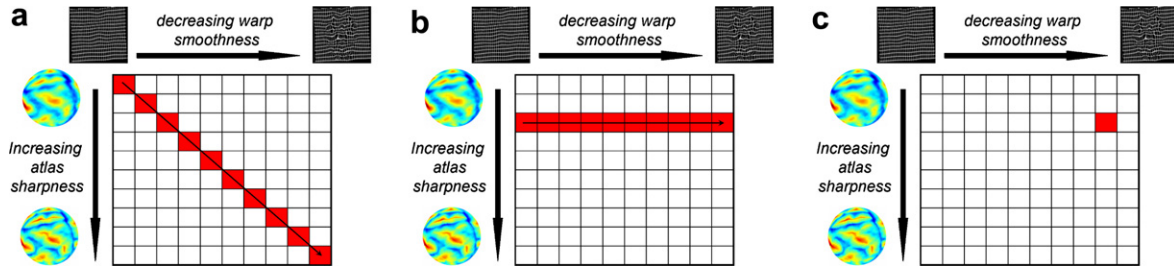


Fig. 2. Strategies for exploring space of atlas sharpness and warp smoothness of a new image: (a) multiple atlases, multiple smoothness (MAMS); (b) single atlas, multiple smoothness (SAMS); and (c) single atlas, single smoothness (SASS).

of the cerebral cortex. These surfaces are represented by triangular meshes with a spherical coordinate system that minimizes metric distortion (Dale et al., 1999; Fischl et al., 1999). The aim is to register an unlabeled cortical surface to a set of manually labeled surfaces and to classify each vertex of the triangular mesh in terms of the anatomical units. To construct the generative model for this problem, we need to define the prior on the registration parameters $p(R; S)$ and the model for label and image generation $p(I, L|R; A_x)$. In general, our model decisions were inspired by previous work on cortical surface parcellation (Fischl et al., 2004).

3.1. Generative model for registration and segmentation

We model the warp regularization with an MRF parameterized by S :

$$p(R; S) \triangleq \frac{F(R)}{Z_1(S)} \exp \left\{ -S \left[\sum_i \sum_{j \in \mathcal{N}_i} \left(\frac{d_{ij}^R - d_{ij}^0}{d_{ij}^0} \right)^2 \right] \right\}, \quad (18)$$

where d_{ij}^R is the distance between vertices i and j under registration R , d_{ij}^0 is the original distance, \mathcal{N}_i is a neighborhood of vertex i and $Z_1(S)$ is the partition function. Our regularization penalizes percentage metric distortion weighted by a scalar S that reflects the amount of smoothness (rigidity) of the final warp. We choose a percentage metric distortion instead of an absolute metric distortion (Fischl et al., 1999) to ensure that the tradeoff between the regularization and the data fidelity term is invariant to scale in our multi-resolution framework. This is especially important in our work since we explore a large range of warp smoothness parameter S .

$F(\cdot)$ ensures invertibility. It is zero if any triangle is folded by warp R and one otherwise. We represent the warp R as a displacement field on the sphere. Therefore, a term like $F(R)$ is necessary. One could replace $F(R)$ with a more sophisticated invertibility prior (e.g., Ashburner et al., 1999; Nielsen et al., 2002) or restrict the space of feasible warps to be the space of diffeomorphisms (Joshi et al., 2004; Glaunès et al., 2004).

We first decompose the label and image prior:

$$p(I, L|R; A_x) = p(L|R; A_x)p(I|L, R; A_x), \quad (19)$$

and impose an MRF prior on the parcellation labels

$$p(L|R; A_x) \triangleq \frac{1}{Z_2(A_x)} \exp \left\{ \sum_i U_i(A_x)L(R(x_i)) + \sum_i \sum_{j \in \mathcal{N}_i} L^T(R(x_i))V_{ij}(A_x)L(R(x_j)) \right\}. \quad (20)$$

Here, we use the vectorized MRF notation of (Kapur et al., 1998). Assuming the total number of labels is M , $L(R(x_i))$ is a column vector of size $M \times 1$ that sums to 1. Each component of $L(R(x_i))$ is an indicator variable. In particular, suppose the image has label m at location $R(x_i)$, then $L(R(x_i))$ is zero for all entries, except for the m -th entry, which is equal to 1. It is a common practice to relax the con-

straint, so that L still sums to 1 but the entries might take on fractional values to indicate uncertainty in the segmentation results. $U_i(A_x)$ is a $1 \times M$ local potential vector that captures the frequency of label $L(R(x_i))$ at vertex i of the atlas mesh. The $M \times M$ compatibility matrix $V_{ij}(A_x)$ reflects the likelihood of two labels at vertices i and j being neighbors. $Z_2(A_x)$ is the partition function dependent on the atlas A_x , ensuring the given MRF is a valid probability distribution.

We further assume that the noise added to the mean image intensity at each vertex location is independent, given the label at that location.

$$p(I|L, R; A_x) \triangleq \prod_i W_i(I(R(x_i)); A_x)L(R(x_i)), \quad (21)$$

where $W_i(I(R(x_i)); A_x)$ is a $1 \times M$ observation potential vector defined at each atlas vertex i . The m -th entry corresponds to the likelihood of observing a particular image intensity or vectors of intensity (in this case, the local surface geometries) at location $R(x_i)$ given a particular label m . We assume W_i follows a Gaussian distribution, e.g., given that the post-central gyrus is at location $R(x_i)$ of the image, we expect the local curvature and/or sulcal depth $I(R(x_i))$ to follow a Gaussian distribution whose parameters are estimated from the training images.

The collection of model components $\{U_i, V_{ij}, W_i\}$ define the atlas A_x . Eq. (20) defines an isotropic prior on the labels, which is a simpler model than that used by modern approaches. The FreeSurfer parcellation algorithm uses a spatially varying and anisotropic MRF model (Fischl et al., 2004), whose parameters change dynamically with the geometries of the subject being segmented. An anisotropic MRF improves the parcellation accuracy, because the boundaries of certain gyral regions are predicted robustly by the variation in curvature. For example, the boundary between the pre-central and post-central gyrus is the central sulcus. Along the boundary, there is high curvature, while across the boundary, the curvature drops off sharply.

We made the explicit choice of warping an image (interpolating an image) to the atlas space. The alternative of warping the atlas (interpolating the atlas) to image space would require us to interpolate the MRF, which is non-trivial. Interpolating U would result in a change in the partition function, which is exponentially hard to compute. In addition, if we were to use the dynamic model of FreeSurfer, since we have made the choice of warping the subject, this would mean that the model parameters and hence the partition function of the MRF would change during the registration step. FreeSurfer does not have this problem, because it does not perform joint registration and segmentation. Therefore, we assume V to be spatially stationary and isotropic and drop the subscripts i, j .

3.2. Atlas building: estimating parameters of generative model

Substituting Eq. (18), (20) and (21) into the atlas co-registration objective function in Eq. (8), we obtain:

$$\begin{aligned}
R_n^* &= \operatorname{argmax}_{R_n} \log p(R_n; S_k) + \log p(I_n, L_n | R_n; A_{S_k}) \\
&= \operatorname{argmax}_{R_n} \log F(R_n) - S_k \sum_i \sum_{j \in \mathcal{N}_i} \left(\frac{d_{ij}^R - d_{ij}^0}{d_{ij}^0} \right)^2 \\
&\quad + \sum_i U_i(A_x) L(R_n(x_i)) + \sum_i \sum_{j \in \mathcal{N}_i} L^T(R_n(x_i)) V(A_x) L(R_n(x_j)) \\
&\quad + \sum_i W_i(I(R_n(x_i)); A_x) L(R_n(x_i)) + \text{const.} \tag{22}
\end{aligned}$$

where the first term prevents folding triangles, the second term penalizes metric distortion, the third and fourth terms are the Markov prior on the labels and the last term is the likelihood of the surface geometries given the segmentation labels of the n -th training image.

We can then fix R_n^* and estimate the atlas parameters $A_x = \{U_i, V, W_i\}$ using Eq. (7). In practice, we use the naive approach of frequency counts (Fischl et al., 2004) to estimate the clique potentials U, V and the maximum likelihood estimation of the Gaussian likelihood terms W . Appendix B provides the implementation details.

3.3. Registration and segmentation of a new image

Similarly, we substitute Eq. (18), (20) and (21) into the update rules for the new subject segmentation in Eq. (16) and registration in Eq. (17).

Warping the subject to the atlas, optimization in Eq. (16) with a fixed $R^{(t)}$ involves estimating the segmentation labels at positions $R^{(t)}(\{x_i\})$ of the subject, where $\{x_i\}$ are vertices of the atlas mesh. We will denote this segmentation estimate by $\hat{L}^{(t+1)}$. Eq. (16) becomes

$$\hat{L}^{(t+1)} = \operatorname{argmax}_{\hat{L}} \log p(\hat{L} | I(R^{(t)}(\{x_i\})); A_x). \tag{23}$$

Even with fixed $R^{(t)}$, solving the MAP segmentation Eq. (23) is NP-hard. We adopt the mean field approximation (Jaakkola, 2000; Kapur et al., 1998). We then use the complete approximate distribution provided by the mean field solver in optimizing Eq. (17). This approximation effectively creates a soft segmentation estimate $\hat{L}_i^{(t+1)}$ at each location $R^{(t)}(x_i)$ of the new subject. $\hat{L}_i^{(t+1)}$ is a column vector of size $M \times 1$. The m -th component of $\hat{L}_i^{(t+1)}$ is the probability of finding label m at location $R^{(t)}(x_i)$ of the new subject. To estimate the label $\hat{L}(x)$ at an arbitrary location x in the subject space, we interpolate from $\hat{L}^{(t+1)}$ defined at $R^{(t)}(\{x_i\})$ onto location x .

This optimization procedure leads to fewer local minima since it avoids commitment to the initial estimates obtained through hard thresholding that might be very far from a good solution if the new image is originally poorly aligned to the atlas. Appendix A provides an outline for computing $\hat{L}^{(t+1)}$ via the mean field approximation. Warping the subject to the atlas, Eq. (17) becomes:

$$\begin{aligned}
R^{(t+1)} &= \operatorname{argmax}_R \log F(R) - S_k \sum_i \sum_{j \in \mathcal{N}_i} \left(\frac{d_{ij}^R - d_{ij}^0}{d_{ij}^0} \right)^2 \\
&\quad + \sum_i U_i(A_x) \hat{L}(R(x_i)) + \sum_i \sum_{j \in \mathcal{N}_i} \hat{L}^T(R(x_i)) V(A_x) \hat{L}(R(x_j)) \\
&\quad + \sum_i W_i(I(R(x_i)); A_x) \hat{L}(R(x_i)) + \text{const.} \tag{24}
\end{aligned}$$

Further implementation details can be found in Appendix B.

4. Experiments and discussion

We consider 39 surfaces that represent the gray-white matter interface of the left and right hemispheres automatically seg-

mented from 3D MRI using FreeSurfer (Dale et al., 1999). This data set exhibits significant anatomic variability since it contains young, middle-aged, elderly subjects and Alzheimer's patients. The surfaces are topologically corrected (Fischl et al., 2001; Ségonne et al., 2007) and a spherical coordinate system is established by minimizing metric distortion (Fischl et al., 1999). For each hemisphere, the 39 cortical surfaces are first rigidly aligned on the sphere, which corresponds to rotation only. The surfaces are manually parcellated by a neuroanatomical expert into 35 labels (Desikan et al., 2006). Fig. 3 illustrates the manual parcellation for one subject. A complete list of the parcellation units is included in Table 1.

We perform cross-validation by leaving out subjects 1 through 10 in the atlas construction, followed by the joint registration–segmentation of the left-out subjects. We repeat with subjects 11 through 20, 21 through 30 and finally 31 through 39. We select S to be the set $\{100, 50, 25, 12.5, 8, 5, 2.5, 1, 0.5, 0.25, 0.1, 0.05, 0.01\}$. We find that in practice, $S = 100$ corresponds to allowing minimal metric distortion and $S = 0.01$ corresponds to allowing almost any distortion. The intervals in the set S are chosen so that each decrease in the value of S roughly corresponds to an average of 1mm increased displacement in registration.

Since we treat the subject mesh as the moving image, both registration and parcellation are performed on the fixed atlas mesh. The segmentation is interpolated from the atlas mesh onto the subject mesh to obtain the final segmentation. We compute segmentation quality by comparing this final segmentation with the “ground truth” manual parcellation.

To speed up the algorithm, we construct the atlas on a sub-divided icosahedron mesh with about 40k vertices. Typically, each subject mesh has more than 100k vertices. The segmentation labels inferred on the low resolution atlas mesh are therefore computed on a coarser grid than that of the manual parcellation. Yet, as we discuss in the remainder of this section, the proposed implementation on average outperforms the FreeSurfer parcellation algorithm (Fischl et al., 2004).

Despite working with a lower resolution icosahedron mesh, registration at each smoothness scale still takes between 20 min to 1.5 hrs per subject per atlas. Registration with weakly constrained warps ($S \leq 0.1$) requires more time because of the need to preserve the invertibility of the warps. The entire set of experiments took approximately 3 weeks to complete on a computing cluster, using up to 80 nodes in parallel.

4.1. Exploration of smoothness S and atlas A_x

In this section, we discuss experimental results for the exploration of the smoothness parameter S and atlases A_x . We measure the segmentation accuracy using the Dice coefficient, defined as the ratio of cortical surface area with correct labels to the total surface area, averaged over the test set.

Fig. 4 shows the segmentation accuracy for SAMS ($A_x = A_1$) and MAMS as we vary S . The average Dice peaks at approximately $S = 1$ for all cross-validation trials, although individual variation exists, as shown in Fig. 5.⁴ For a particular value of S , outliers warp more because the tradeoff between the data fidelity and regularization is skewed towards the former. However, it is surprising to find that the optimal S is mostly constant across subjects. It also appears that peaks of the segmentation accuracy plots are relatively broad, implying that good parcellation results can be obtained for a range of S between 1 and 2.5.

For large S (highly constrained warps), MAMS consistently outperforms SAMS. Because the surfaces are initially misaligned, using

⁴ The optimal value of $S = 1$ is coincidental in the sense that it depends on the unit chosen for metric distortion.




Fig. 3. Example of manual parcellation shown on a partially inflated cortical surface. In our data set, the neuroanatomist preferred gyral labels to sulcal labels. There are also regions where sulci and gyri are grouped together as one label, such as the superior and inferior parietal complexes.

Table 1
List of parcellation structures

1. Sylvian fissure/unknown	2. Bank of the superior temporal sulcus
3. Caudal anterior cingulate	4. Caudal middle frontal gyrus
5. Corpus callosum	6. Cuneus
7. Entorhinal	8. Fusiform gyrus
9. Inferior parietal complex	10. Inferior temporal gyrus
11. Isthmus cingulate	12. Lateral occipital
13. Lateral orbito frontal	14. Lingual
15. Medial orbito frontal	16. Middle temporal gyrus
17. Parahippocampal	18. Paracentral
19. Parsopercularis	20. Parsorbitalis
21. Parstriangularis	22. Peri-calcarine
23. Post-central gyrus	24. Posterior cingulate
25. Pre-central gyrus	26. Pre-cuneus
27. Rostral anterior cingulate	28. Rostral middle frontal
29. Superior frontal gyrus	30. Superior parietal complex
31. Superior temporal gyrus	32. Supramarginal
33. Frontal pole	34. Temporal pole
35. Transverse temporal	

a sharp atlas (in the case of SAMS, atlas A_1) results in poor segmentation accuracy due to a mismatch between the image features and the atlas statistics. As we decrease the smoothness S , SAMS allows for more flexible warps towards the population average than MAMS since it uses a sharper atlas. The similarity measure is therefore given higher weight to overcome the regularization. This results in better segmentation accuracy than MAMS. Eventually, SAMS and MAMS reach similar maximal values at the same optimal smoothness S . Beyond the optimal S , however, both MAMS

and SAMS exhibit degraded performance. This is probably due to overfitting and local optima created by more flexible warps.

We also examine the Dice measure of each parcellation structure as a function of the warp smoothness S . In general, the curves are noisier, but follow those of Fig. 4. Fig. 6a shows a typical curve that peaks at $S = 1$, while Fig. 6b shows a curve that peaks at $S = 5$. However, in general, for most structures, the optimal smoothness S occurs at approximately $S = 1$ (Fig. 7), which is not surprising since for most subjects, the optimal overall Dice (computed over the entire surface) also occurs at $S = 1$ (Fig. 5).

We now explore the effects of both warp smoothness S and atlas sharpness α on parcellation accuracy. Fig. 8 shows a plot of Dice averaged over all 39 subjects. The performances of MAMS, SAMS and SASS at ($S = 1, \alpha = 1$) are indistinguishable. As an example, we see that for both hemispheres, SAMS with $\alpha = 0.01$ (green line) starts off well, but eventually overfits with a worse peak at $S = 1$ ($p < 10^{-5}$ for one-sided paired-sampled t -test, statistically significant even when corrected for multiple comparisons). Similarly for SASS, the best values of α and S are both equal to 1. We also show SASS with $\alpha = 0.01$ and $S = 1$ in Fig. 8. In general, there is no statistical difference between MAMS, SAMS or SASS at their optima: $\alpha = 1, S = 1$.

Originally, MAMS and SAMS were introduced to reduce local optima in methods, such as SASS. It is therefore surprising that the performance of all three methods is comparable. While using the correct smoothness and atlas sharpness is important, our “annealing” process of gradual reduction of smoothness (MAMS and SAMS) does not seem to increase the extent of the basins of

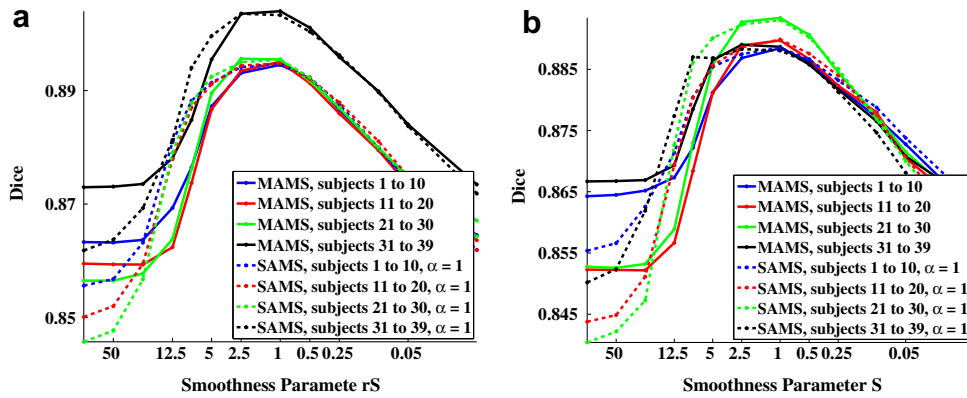


Fig. 4. Parcellation accuracy as a function of warp smoothness. S is plotted on a log scale: (a) left hemisphere; and (b) right hemisphere.




Fig. 5. Histogram of optimal warp smoothness S across subjects (MAMS): (a) left hemisphere (b) right hemisphere.

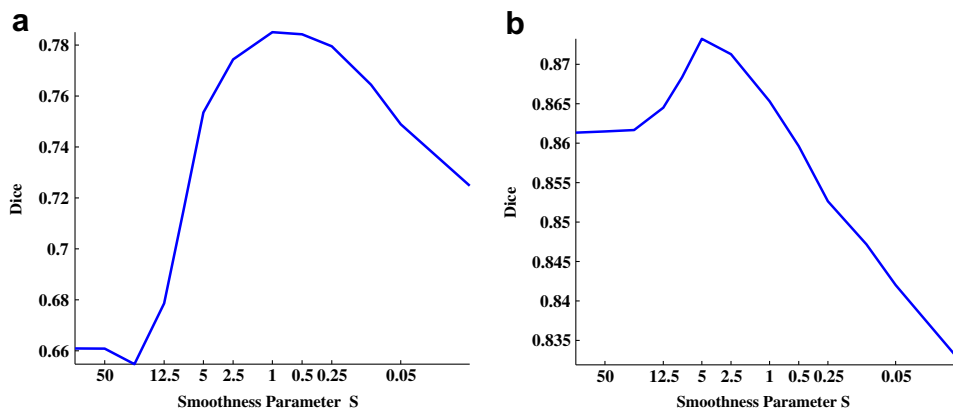


Fig. 6. (a) Typical plot of Dice against smoothness S . (b) A noisy plot of Dice against smoothness S . (a) Right inferior temporal gyrus. (b) Right temporal pole.

attraction in the context of cortical parcellation. One possible reason is that on the cortical surfaces, two adjacent sulci might appear quite similar locally. Smoothing these features might not necessarily remove the local minima induced by such similarity. Incorporating multiscale features (Nain et al., 2007; Yu et al., 2007a,b) with multiple smoothness offers a promising approach for avoiding such local optimal issues on the cortex.

The fact that the optimal smoothness parameter S^* corresponds to the optimal atlas sharpness parameter α^* is not surprising. According to the graphical model in Fig. 1 and as mentioned in the derivations in Section 2.2 and Section 2.3, theoretically, we

do expect $S^* = \alpha^*$. Learning this optimal S^* in the atlas construction process is a future avenue of investigation.

Alternatively, we can also determine the best S and A_z for a new image registration by optimizing the objective function in Eq. (11). Unfortunately, there are technical difficulties in doing this. First, we notice that the objective function in Eq. (11) increases with decreasing S . This model contains no Occam's razor regularization terms that penalize overfitting due to flexible warps. This is mainly because Eq. (11) omits certain difficult-to-compute normalization terms, such as the partition function that depends on U , V and W and thus dependent on S and α . These terms are ignored for fixed




Fig. 7. Histogram of optimal S across structures (MAMS): (a) left hemisphere; and (b) right hemisphere.

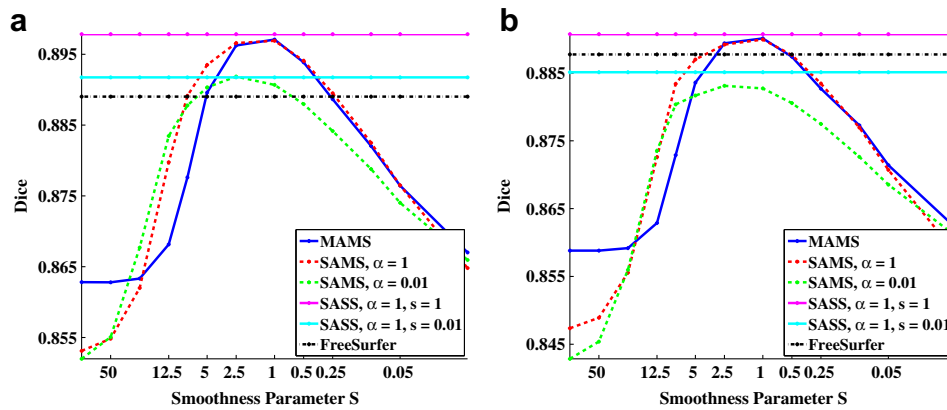


Fig. 8. Overall Dice versus smoothness. S is plotted on a log scale: (a) left hemisphere; and (b) right hemisphere.

values of S and α . We can use various approximation strategies to compute the normalization terms. But it is not clear whether these approximations yield sufficient accuracy to determine the optimal values for S and α . On the other hand, empirically we find that $S = \alpha = 1$ consistently yields the optimal (or very close to the optimal) segmentation performance. This suggests that one can probe the training data using a MAMS-type strategy to determine the optimal values of warp smoothness and atlas sharpness, and then use the SASS strategy for future registration and segmentation of new subjects. Furthermore, our experiments suggest that segmentation accuracy is tolerant up to a small mismatch between atlas sharpness and warp smoothness.

Further work will involve the application of our framework to other data sets (including volumetric data) and experiments with other models of data fidelity and regularization. We expect the optimal smoothness and atlas sharpness to be different but it would be interesting to verify if these values are consistent across subjects. It is possible that in the volumetric case where the intensity features are more predictive of labels, a MAMS-type strategy may outperform SASS, particularly in cases where the anatomy is dramatically different (e.g., ventricles in Alzheimer patients).

In this work, the optimal smoothness parameter found is global to the entire surface. Previous work demonstrated the advantage of using spatially varying smoothness parameters (Commowick et al., 2005). Unfortunately, the approach of exhaustive search presented here cannot be directly applied, since it would involve searching over a much larger space of parameters.

4.2. Comparison with FreeSurfer

We now compare the performance of our algorithm with the FreeSurfer parcellation algorithm, described in (Fischl et al., 2004) and extensively validated in (Desikan et al., 2006).

It is unclear which algorithm has a theoretical advantage. The FreeSurfer algorithm is essentially a “single atlas, single smoothness” (SASS) method that uses a sequential registration–segmentation approach and a more complicated anisotropic MRF model that has been specifically designed and fine-tuned for the surface parcellation application. Our model lacks the anisotropic MRF and introducing it would further improve its performance. On the other hand, FreeSurfer uses iterated conditional modes (Besag, 1986) to solve the MRF, while we use the more powerful mean field approximation (Jaakkola, 2000; Kapur et al., 1998). FreeSurfer also treats the subject mesh as a fixed image and the parcellation is performed directly on the subject mesh. Therefore, unlike our approach, no interpolation is necessary to obtain the final segmentation.

As shown in Fig. 8, the optimal performances of MAMS, SAMS and SASS are statistically significantly better (even when corrected for multiple comparisons) than the FreeSurfer, with p -value 1×10^{-8} (SASS) for the left hemisphere and 8×10^{-4} (SASS) for the right hemisphere.

Because Dice computed over the entire surface can be deceiving by suppressing small structures, we show in Fig. 9 the percentage improvement of SASS over FreeSurfer on inflated cortical surfaces. Qualitatively, we see that SASS performs better than FreeSurfer




Fig. 9. Percentage improvement of SASS over FreeSurfer. The boundaries between parcellation regions are set to reddish-brown so that the different regions are more visible: (a) left lateral; (b) right lateral; (c) left medial; and (d) right medial. Note that none of the blue-colored structures achieves statistically better segmentation accuracy using Free Surfer than SASS (see text).

since there appears more orange-red regions than blue regions. The fact that the colorbar has significantly higher positive values than negative values indicates that there are parcellation regions with significantly greater improvements compared with regions that suffer significant penalties.

More quantitatively, Figs. 10 and 11 display the average Dice per structure for FreeSurfer, MAMS, SAMS and SASS at ($S = 1, \alpha = 1$) for the left and right hemispheres, respectively. Standard errors of the mean are displayed as red bars. The numbering of the structures correspond to Table 1. The structures with the worst




Fig. 10. Structure-specific parcellation accuracy for the left hemisphere. First column (dark blue) corresponds to FreeSurfer. Second (light blue), third (yellow) and fourth (brown) columns correspond to MAMS, SAMS and SASS respectively. ($S = 1, \alpha = 1$). * indicates structures where SASS shows statistically significant improvement over FreeSurfer. There is no structure that becomes worse. (a) Left hemi structures: 1–9; (b) left hemi structures: 10–18; (c) left hemi structures: 19–27; and (d) left hemi structures: 28–35.




Fig. 11. Structure-specific parcellation accuracy for the right hemisphere. First column (dark blue) corresponds to FreeSurfer. Second (light blue), third (yellow) and fourth (brown) columns correspond to MAMS, SAMS and SASS respectively. ($S = 1, \alpha = 1$). * indicates structures where SASS shows statistically significant improvement over FreeSurfer. There is no structure that becomes worse. (a) Right hemi structures: 1–9; (b) right hemi structures: 10–18; (c) right hemi structures: 19–27; and (d) right hemi structures: 28–35.

Dice are the frontal pole, corpus callosum and entorhinal cortex. These structures are small and relatively poorly defined by the underlying cortical geometry. For example, the entorhinal cortex is partially defined by the rhinal sulcus, a tiny sulcus that is only visible on the pial surface. On the other hand, the corpus callosum is mapped from the white matter volume onto the cortical surface. Its boundary is thus defined by the surrounding structures, rather than by the cortical geometry.

For each structure, we perform a one-sided paired-sampled t -test between SASS and FreeSurfer, where each subject is considered a sample. We use the false discovery rate (FDR) (Benjamini and Hochberg, 1995) to correct for multiple comparisons. In the left hemisphere, SASS achieves statistically significant improvement over FreeSurfer for 17 structures ($FDR < 0.05$), while the remaining structures yield no statistical difference. In the right hemisphere, SASS achieves improvement for 11 structures ($FDR < 0.05$), while the remaining structures yield no statistical difference. The p -values for the left and right hemispheres are pooled together for the false discovery rate analysis.

A major factor influencing the accuracy of the automatic parcellation is the manual segmentation. In well-defined regions such as pre- and post-central gyri, accuracy is more than 90% and within the range of inter-rater variability. In the ambiguous regions, such as the frontal pole, inconsistent manual segmentation leads to less consistent training and worse segmentation performance.

5. Conclusions

In this paper, we proposed a generative model for the joint registration and segmentation of images. The atlas construction pro-

cess is formulated as estimation of the graphical model parameters. The framework incorporates consistent atlas construction, multiple atlases of varying sharpness and MRF priors on both registration warps and segmentation labels. We show that atlas sharpness and warp regularization are important factors in segmentation and that the optimal smoothness parameters are stable across subjects in the context of cortical parcellation. The experiments imply that the optimal atlas sharpness and warp smoothness can be determined by cross-validation. Furthermore, segmentation accuracy is tolerant up to a small mismatch between atlas sharpness and warp smoothness. With the proper choice of atlas sharpness and warp regularization, even with a less complex MRF model, the joint registration–segmentation framework achieves better segmentation accuracy than the state-of-the-art FreeSurfer algorithm (Desikan et al., 2006; Fischl et al., 2004).

There are various promising directions for future work. We believe that incorporating multiscale features (Nain et al., 2007; Yu et al., 2007a,b) can further improve the registration and parcellation of the cortical surfaces. Learning the optimal spatially varying smoothness parameters in the training set is also a problem we are currently working on. Finally, we plan to apply this framework to other problems, including volumetric segmentation.

Acknowledgements

The authors would like to thank Serdar Balci, Wanmei Ou, Kilian Pohl and Lilla Zollei for helpful conversations. They would like to especially thank Koen Van Leemput for illuminating discussions on atlas construction and other issues in registration. Support for this research is provided in part by the National Alliance for

Medical Image Analysis (NIH NIBIB NIMIC U54-EB005149), the Neuroimaging Analysis Center (NIT CRR NAC P41-RR13218), the Morphometry Biomedical Informatics Research Network (NIH NCRR mBIRN U24-RR021382), the NIH NINDS R01-NS051826 grant, the NSF CAREER 0642971 grant, National Center for Research Resources (P41-RR14075, R01 RR16594-01A1 and the NCRR BIRN Morphometric Project BIRN002, U24 RR021382), the National Institute for Biomedical Imaging and Bioengineering (R01 EB001550, R01EB006758), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01) as well as the Mental Illness and Neuroscience Discovery (MIND) Institute. B.T. Thomas Yeo is funded by the Agency for Science, Technology and Research, Singapore.

Appendix A. Mean field derivation outline

The mean field approximation uses a variational formulation (Jaakkola, 2000), where we seek to minimize the KL-divergence (denoted by $D(\cdot||\cdot)$) between $q(\hat{L}) = \prod_i \mathbf{b}_i(\hat{L}_i)$ and $p(\hat{L}|I(R^t)(\{x_i\}); A_x)$:

$$\{\mathbf{b}_i^*\}_{\{\mathbf{b}_i\}} = \operatorname{argmin}_{\{\mathbf{b}_i\}} D(q(\hat{L})||p(\hat{L}|I(R^t)(\{x_i\}); A_x)). \quad (\text{A.1})$$

This results in a fixed-point iterative solution. Since this is a fairly standard derivation (Jaakkola, 2000), we only provide the final update:

$$\begin{aligned} \mathbf{b}_i(m) &\propto \exp \{U_i(m) + \log p(I(R^t)(x_i))\} \hat{L}_i \\ &= m; A_x) + \sum_{j \in \mathcal{V}_i} \sum_{\hat{L}_j} \mathbf{b}_j(\hat{L}_j) [V(m, \hat{L}_j) + V(\hat{L}_j, m)], \end{aligned} \quad (\text{A.2})$$

where \mathbf{b}_i is normalized to be a valid probability mass function.

Appendix B. Implementation

We now present some implementation details for completeness. We first discuss the estimation of the atlas A_x defined by $\{U_i, V, W_i\}$ in Eq. (20) and (21) from the maximum likelihood function Eq. (7). In our model, estimating U_i and V is hard in practice, since evaluating Eq. (7) requires computing the NP-hard partition function. Instead, we use frequency counts to estimate the clique potentials, similar to FreeSurfer Fischl et al. (2004).

- In our implementation, the singleton potential U_i is a row vector of length M and L_n is a column indicator vector. We set

$$U_i = \log \frac{1}{N} \sum_n L_n^T(R_n(x_i)), \quad (\text{A.3})$$

where $(\cdot)^T$ indicates transpose.

- The pairwise potential V is a $M \times M$ matrix. Following (Fischl et al., 2004), we set

$$V = \log \frac{1}{2NE} \sum_n \sum_i \sum_{j \in \mathcal{V}_i} L_n(R_n(x_i)) L_n^T(R_n(x_j)), \quad (\text{A.4})$$

where E is the number of edges in the atlas mesh. More rigorous methods of optimizing the clique potentials through iterative proportional fitting (Jirousek and Preucil, 1995) would further improve the clique potential estimates.

- The likelihood potential W_i is a row vector of length M defined at each vertex. The m -th entry of W_i corresponds to the likelihood of observing a particular image intensity or vectors of intensity (in our case, the local surface geometries) at location $R(x_i)$ given a particular label m . While we might be observing multiple geometric features at any vertex, the likelihood of these features is combined into the row vector W_i . We use

maximum likelihood estimates of the mean and variance of the Gaussian distribution of cortical geometries conditioned on the segmentation labels to parameterize this distribution. In this work, we use the mean curvature of the original cortical surface and average convexity, which is a measure of sulcal depth (Fischl et al., 1999), as intensity features. At spatial locations where there is no training data for a particular label, it is unclear what the value of the entry in W should be since it is spatially varying. We simply assume a mean of zero and a large variance, essentially being agnostic to the value of intensity we observe. A more sophisticated method would involve the use of priors on the atlas parameters, so that the atlas parameters become random. In that case, when there is no observations, the maximum likelihood estimates of the atlas parameters become the priors.

Secondly, we discuss the registration of a training image to an atlas (Eq. (22)) and the new image registration (Eq. (24)).

- The registration warp R is a map from a 2-Sphere to a 2-Sphere. We represent R as a dense displacement field. In particular, each point x_i has an associated displacement vector u_i tangent to the point x_i on the sphere. $R(x_i)$ maps x_i to $x_i + u_i$ normalized to be a point on the sphere.
- To interpolate (for example the mean curvature of a cortical surface onto $R(x_i)$), we first find the intersection between the vector $R(x_i)$ and each planar triangle of the spherically mapped cortical surface. We then use barycentric interpolation to interpolate the values at the vertices of the mesh onto $R(x_i)$.
- The above two bullets completely specify the computation of the atlas co-registration objective function Eq. (22) and the new subject registration function Eq. (24). This allows us to compute the gradients of the objective function via the chain rule.
- We use conjugate gradient ascent with parabolic line search (Press, 1992) on a coarse-to-fine grid. The coarse-to-fine grid comes from the representation of the atlas as a subdivided icosahedron mesh.
- The final segmentation is obtained by selecting for each vertex the label with the highest posterior probability.
- To satisfy $F(R)$, the regularization that induces invertibility, we ensure that no step in the line search results in folded triangles. Unfortunately, in practice, this results in many small steps. It is much more efficient to perform the line search without considering $F(R)$, and then unfold the triangles using $F(R)$ after the line search. In general, we find that after unfolding, the objective function is still better than the previous iteration. This unfolding process can be expensive for small smoothness parameter S ($S \leq 0.1$), resulting in long run times of about 1.5 hrs per subject per atlas for $S \leq 0.1$.

References

- Allasonnière, S., Amit, Y., Trounev, A., 2007. Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society B* 68, 3–29.
- Ashburner, J., Friston, K., 2005. Unified segmentation. *NeuroImage* 26, 839–851.
- Ashburner, J., Andersson, J., Friston, K., 1999. High-dimensional image registration using symmetric priors. *NeuroImage* 9, 619–628.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of Royal Statistical Society B* 57 (1), 289–300.
- Besag, J., 1986. On the statistical analysis of dirty pictures. *Journal of Royal Statistical Society B* 48 (3), 259–302.
- Bhatia, K.K., Hajnal, J.V., Puri, B.K., Edwards, A.D., Rueckert, D., 2004. Consistent groupwise non-rigid registration for atlas construction. In: *International Symposium on Biomedical Imaging*.
- Collins, D.L., Zijdenbos, A.P., Baaré, W.F.C., Evans, A.C., 1999. ANIMAL + INSECT: improved cortical structure segmentation. *Information Processing in Medical Imaging* 1613, 210–223.

- Commowick, O., Stefanescu, R., Fillard, P., Arsigny, V., Ayache, N., Pennec, X., Malandain, G., 2005. Incorporating statistical measures of anatomical variability in atlas-to-subject registration for conformal brain radiotherapy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, LNCS 3750, pp. 927–934.
- Dale, A., Fischl, B., Sereno, M., 1999. Cortical surface-based analysis I: segmentation and surface reconstruction. *NeuroImage* 9 (2), 179–194.
- De Craene, M., du Bois d'Aische, A., Macq, B., Warfield, S., 2004. Multi-subject registration for unbiased statistical atlas construction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, LNCS 3216, pp. 655–662.
- Desikan, R., S'egonne, F., Fischl, B., Quinn, B., Dickerson, B., Blacker, D., Buckner, R., Dale, A., Maguire, P., Hyman, B., Albert, M., Killiany, R., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31, 968–980.
- Evans, A.C., Collins, D.L., Mills, S.R., Brown, E.D., Kelly, R.L., Peters, T.M., 1993. 3D statistical neuroanatomical models from 305 MRI volumes. Nuclear Science Symposium and Medical Imaging Conference, IEEE Conference Record 3 (31), 1813–1817.
- Fischl, B., Sereno, M., Dale, A., 1999. Cortical surface-based analysis II: inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9 (2), 195–207.
- Fischl, B., Sereno, M., Tootell, R., Dale, A., 1999. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* 8 (4), 272–284.
- Fischl, B., Liu, A., Dale, A., 2001. Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Transactions on Medical Imaging* 20 (1), 70–80.
- Fischl, B., Salat, D., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klavness, S., Montillo, A., Makris, N., Rosen, B., Dale, A., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33 (3), 341–355.
- Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D., Busa, E., Seidman, L.J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., Dale, A., 2004. Automatically parcellating the human cerebral cortex. *Cerebral Cortex* 14, 11–22.
- Fischl, B., Rajendran, N., Busa, E., Augustinack, J., Hinds, O., Yeo, B.T.T., Mohlberg, H., Amunts, K., Zilles, K., 2007. Cortical folding patterns and predicting cytoarchitecture. *Cerebral Cortex*.
- Glaunès, J., Vaillant, M., Miller, M., 2004. Landmark matching via large deformation diffeomorphisms on the sphere. *Journal of Mathematical Imaging and Vision* 20, 179–200.
- Guimond, A., Meunier, J., Thirion, J.-P., 2000. Average brain models: a convergence study. *Computer Vision and Image Understanding* 77 (2), 192–210.
- Heckemann, R.A., Hajnal, J., Aljabar, P., Rueckert, D., Hammers, A., 2006. Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33 (1), 115–126.
- Jaakkola, T., 2000. Tutorial on variational approximation methods. In: *Advanced Mean Field Methods: Theory and Practice*. MIT Press.
- Jirousek, R., Preucil, S., 1995. On the effective implementation of the iterative proportional fitting procedure. *Computational Statistics and Data Analysis* 19, 177–189.
- Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23, 151–160.
- Kapur, T., Grimson, W.E.L., Kikinis, R., Wells, W., 1998. Enhanced spatial priors for segmentation of magnetic resonance imagery. In: International Conference on Medical Image Computing and Computer Aided Intervention, pp. 457–468.
- Klein, A., Hirsch, J., 2005. Mindboggle: a scatterbrained approach to automate brain labeling. *NeuroImage* (24), 261–280.
- Lohmann, G., Von Cramon, D.Y., 2000. Automatic labelling of the human cortical surface using sulcal basins. *Medical Image Analysis* 4, 179–188.
- Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., Joshi, S., 2006. Multi-modal image set registration and atlas formation. *Medical Image Analysis* 10 (3), 440–451.
- Makrogiannis, S., Verma, R., Karacali, B., Davatzikos, C., 2006. A joint transformation and residual image descriptor for morphometric image analysis using an equivalence class formulation. *Mathematical Methods in Biomedical Image Analysis*.
- Mangin, J.-F., Frouin, V., Bloch, I., Régis, J., Lopez-Krahe, J., 1994. Automatic construction of an attributed relational graph representing the cortex topography using homotopic transformations. *Proceedings of SPIE* 2299, 110–121.
- Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., 1995. A probabilistic atlas of the human brain: theory and rationale for its development the international consortium for brain mapping (ICBM). *NeuroImage* 2 (2), 89–101.
- Nain, D., Haker, S., Bobick, A., Tannenbaum, A., 2007. Multiscale 3D shape representation and segmentation using spherical wavelets. *IEEE Transactions on Medical Imaging* 26 (4), 598–618.
- Nielsen, M., Johansen, P., Jackson, A.D., Laurup, B., 2002. Brownian warps: a least committed prior for non-rigid registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, LNCS 2489, pp. 557–564.
- Paus, T., Tomaiuolo, F., Otaky, N., MacDonald, D., Petrides, M., Atlas, J., Morris, R., Evans, A., 1996. Human cingulate and paracingulate sulci: pattern, variability, asymmetry, and probabilistic map. *Cerebral Cortex* 6, 207–214.
- Pohl, K., Fisher, J., Grimson, W.E.L., Kikinis, R., Wells, W., 2006. A Bayesian model for joint segmentation and registration. *NeuroImage* 31, 228–239.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, Cambridge, UK.
- Rettmann, M., Han, X., Xu, C., Prince, J., 2002. Automated sulcal segmentation using watersheds on the cortical surface. *NeuroImage* 15, 329–344.
- Richard, F., Samson, A., Cuenod, C., 2007. A SAEM algorithm for the estimation of template and deformation parameters in medical image sequences. Preprint MAP5, Nov.
- Rivière, D., Mangin, J.-F., Papadopoulos-Orfanos, D., Martinez, J.-M., Frouin, V., Régis, J., 2002. Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. *Medical Image Analysis* 6, 77–92.
- Ségonne, F., Pacheco, J., Fischl, B., 2007. Geometrically accurate topology-correction of cortical surfaces using non-separating loops. *IEEE Transactions on Medical Imaging* 26 (4), 518–529.
- Tao, X., Prince, J., Davatzikos, C., 2002. Using a statistical shape model to extract sulcal curves on the outer cortex of the human brain. *Transactions on Medical Imaging* 21, 513–524.
- Tu, Z., Zheng, S., Yuille, A., Reiss, A., Dutton, R., Lee, A., Galaburda, A.M., Dinov, I., Thompson, P., Toga, A., 2007. Automated extraction of the cortical sulci based on a supervised learning approach. *IEEE Transactions on Medical Imaging* 26 (4), 541–552.
- Twining, C., Cootes, T., Marsland, S., Petrovic, V., Schestowitz, R., Taylor, C., 2005. A unified information theoretic approach to groupwise non-rigid registration and model building. *Information Processing in Medical Imaging* 9, 1–14.
- Van Leemput, K., 2006. Probabilistic brain atlas encoding using Bayesian inference. In: International Conference on Medical Image Computing and Computer Aided Intervention, pp. 704–711.
- Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging* 18 (10), 897–908.
- Weisenfeld, N., Warfield, S., 2007. Simultaneous alignment and central tendency estimation for brain atlas construction. In: *Workshop on Statistical Registration: Pair-wise and Group-wise Alignment and Atlas Formation*.
- Wyatt, P., Noble, A., 2003. MAP MRF joint segmentation and registration. *Medical Image Analysis* 7 (4), 539–552.
- Xiaohua, C., Brady, M., Rueckert, D., 2004. Simultaneous segmentation and registration for medical image. In: International Conference on Medical Image Computing and Computer Aided Intervention, 3216, pp. 663–670.
- Xiaohua, C., Brady, M., Lo, J., Moore, N., 2005. Simultaneous segmentation and registration of contrast-enhanced breast MRI. *Information Processing in Medical Imaging* 3565, 126–137.
- Yeo, B.T.T., Sabuncu, M., Mohlberg, H., Amunts, K., Zilles, K., Golland, P., Fischl, B., 2007. What data to co-register for computing atlases. In: *Proceedings of the International Conference on Computer Vision, IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis*.
- Yeo, B.T.T., Sabuncu, M., Desikan, R., Fischl, B., Golland, P., 2007. Effects of registration regularization and atlas sharpness on segmentation accuracy. In: International Conference on Medical Image Computing and Computer Aided Intervention, pp. 683–691.
- Yezzi, A., Zollei, L., Kapur, T., 2003. A variational framework for integrating segmentation and registration through active contours. *Medical Image Analysis* 7 (2), 171–185.
- Yu, P., Grant, E., Qi, Y., Han, X., Ségonne, F., Pienaar, R., Busa, E., Pacheco, J., Makris, N., Buckner, R., Golland, P., Fischl, B., 2007a. Cortical surface shape analysis based on spherical wavelets. *IEEE Transaction on Medical Imaging* 26 (4), 582–597.
- Yu, P., Yeo, B.T.T., Grant, E., Fischl, B., Golland, P., 2007b. Cortical folding development study based on over-complete spherical wavelets. In: *Proceedings of MMBIA: IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis*.