AUTOMATICALLY LEARNING CORTICAL FOLDING PATTERNS

Matthew Toews[1], D. Louis Collins[2] and Tal Arbel[1]

[1]Centre for Intelligent Machines, McGill University
{mtoews,arbel}@cim.mcgill.ca

[2]Montreal Neurological Institute, McGill University
louis@bic.mni.mcgill.ca

ABSTRACT

A data-driven technique is presented for automatically learning cortical folding patterns from MR brain images of different subjects. Cortical patterns are represented in terms of generic scale-invariant image features. Learning automatically identifies a set of features that occur with statistical regularity in appearance and geometry from a large set of MR volume renderings, based on a predescribed anatomical region of interest. A filtering technique is presented for distinguishing between valid cortical features and those likely to arise from incorrect correspondences, based on feature geometry. Expert validation of 100 feature instances shows that 77% correctly identify the same underlying cortical structure in different brains despite high inter-subject variability, and filtering improves the ability to identify the most meaningful patterns.

1. INTRODUCTION

The task of identifying and describing folding patterns in the cortex is of high interest, as cortical folds are closely related to functional divisions within the brain. Automatically identifying cortical folding patterns in magnetic resonance (MR) images of different subjects has proven a challenge, however, due to inter-subject variability. A specific fold may vary significantly in shape, exhibit multiple distinct morphologies, or may not exist in all subjects. Although primary folding patterns such as the central sulcus or lateral fissure are identifiable in virtually all adult brains, secondary or tertiary folding patterns are difficult to reliably identify in all subjects, even for human experts [1].

A variety of different techniques are used to identify or analyze cortical folding. Manual labeling is widely used [2], however it is prone to inter-rater variation and tedious, particularly for the large datasets required to observe the full range of cortical variability. Images of different subjects can be registered or aligned within a common geometrical reference frame or atlas for analysis [3], and cortex-specific image features such as sulci can be used to improve alignment [5, 6], however a one-to-one mapping between subjects is difficult to determine and may not generally exist due to inter-subject variation [4, 6]. Machine learning can be used to reproduce expert sulcal labelings from training examples [7, 8]. Automatically identifying new, unlabeled cortical folding patterns that may only be present in a subset of subjects remains an open research challenge.

The contribution of this paper is a technique for automatically learning new, unlabeled cortical folding patterns from a large set of subject images, based on a predescribed anatomical region of interest. The primary challenge is in reliably identifying instances of the same folding pattern in different subjects, as a specific folding pattern may not be present in all subjects or may appear similar to other unrelated patterns on a local scale. Cortical appearance is represented as a generative probabilistic model based on a collection of distinctive, localized scale-invariant image features [10, 11]. Machine learning is applied to automatically identify clusters of features that arise from the same underlying cortical structure in different training images [9]. Coping with inter-subject cortical variability remains a challenge, however, and many learned features arise from invalid, coincidental correspondences between unrelated cortical structure. This paper develops a novel measure to quantify the likelihood of invalid feature correspondence, based on feature geometry, which can be used to filter or discard such features from the model.

The remainder of this paper is as follows. Section 2 reviews appearance modeling from scale-invariant features. Section 3 presents a novel measure to quantify the tendency of features to result in incorrect correspondence between unrelated cortical patterns. Section 4 presents experiments modeling 196 lateral volume renderings of the International Consortium of Brain Mapping (ICBM152) data set [12]. Expert validation of a subset of 10 model features shows that 77% of feature occurrences represent the same underlying cortical structure, and filtering based on feature geometry significantly reduces the number of invalid features.

2. SCALE-INVARIANT FEATURE MODELING

This paper proposes modeling the cortical surface using local invariant features [10, 11]; distinctive, generic image patterns that can be robustly extracted and used to establish correspondences between a wide variety of different image types. While special-purpose features corresponding to known cortical structures can be used, e.g. specific sulci [13, 7], generic

features are useful when the features of interest are not known a priori, e.g. generic cortical patterns. A scale-invariant feature is an oriented image region, characterized geometrically in terms of location x, orientation θ and scale σ within the image. Although scale-invariance is useful for establishing correspondences between images of different resolution, its primary significance in brain modeling is in determining the characteristic scale of an image pattern. In cortical images for instance, feature scales are typically related to the width of anatomical structures in the image such as gyri or sulci, see Figure 1.

Scale-invariant features alone cannot be reliably matched between cortical images of different subjects due to inter-subject variability, however probabilistic modeling can be used to identify a set of features which exhibit similar geometry and appearance over a population of subjects. The model of [9] is considered here, consisting of a set of features $\{m_i\}$ within a common reference frame o. A model feature is denoted as $m_i = \{m_i^b, m_i^g, m_i^a\}$ representing the occurrence, geometry and appearance of a scale-invariant feature within an image, respectively. Occurrence m_i^b is a binary random variable representing feature presence/absence. Geometry $m_i^g = \{x_i, \theta_i, \sigma_i\}$ is an oriented region in $\mathbf{R}^N$ image space, represented by N-parameter image location x_i, an $N(N-1)/2$ parameter orientation θ_i and a scale σ_i. Appearance m_i^a represents the image content within region m_i^g, and can generally be parameterized in a number of ways, e.g. principal components. The reference frame is denoted as $o = \{o^b, o^g\}$ representing the occurrence and geometry of a common structure with respect to which the geometric variability of features of m_i can be measured, where o^b and o^g are parameterized in the same manner as m_i^b and m_i^g.

The model is formulated probabilistically as the conditional posterior probability of the reference frame o given a set of model features $\{m_i\}$:

$$p(o|\{m_i\}) = \frac{p(o)p(\{m_i\}|o)}{p(\{m_i\})} = \frac{p(o)\prod_i^M p(m_i|o)}{p(\{m_i\})}, \quad (1)$$

by applying Bayes rule and assuming conditional independence of model features given o. The assumption of conditional independence states that knowing o, all remaining feature variability can be quantified locally on an individual feature basis by factor $p(m_i|o)$, which can be expressed as:

$$\begin{aligned} p(m_i|o) &= p(m_i^a, m_i^b|o)p(m_i^g|o) \\ &= p(m_i^a|m_i^b)p(m_i^b|o^b)p(m_i^g|o^b, o^g), \quad (2) \end{aligned}$$

under the assumptions that feature appearance and occurrence are conditionally independent of feature geometry given o, and that feature appearance and the reference frame are conditionally independent given feature occurrence. Here, $p(m_i^a|m_i^b)$ is a density over feature appearance given feature occurrence, $p(m_i^b|o^b)$ is the probability of feature occurrence

given reference frame occurrence, and $p(m_i^g|o^b, o^g)$ represents the residual error of a linear transform from feature to reference frame geometry given reference frame occurrence.
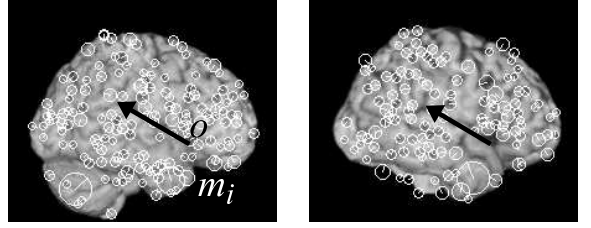

Fig. 1. A subset of features m_i (white circles) and the reference frame o (black arrow) in volume renderings of MR brain images from two different subjects. The reference frame is predefined according to a region of interest, here the lateral fissure, and can be automatically identified when fitting the model to new images. Note the high variability of cortical appearance from one subject to the next.

3. GEOMETRY AND MATCHING ERROR

Model learning involves labeling the reference frame o and extracting features in images, then automatically identifying clusters of features in different images which are similar in terms of their appearance and geometry with respect to o. Each such cluster is a model feature m_i, and represents instances of the same underlying cortical folding pattern in different subjects. The set of model features can be investigated to identify the folding patterns that occur most frequently or are most representative of a cortical region of interest.

The difficulty is that learning is susceptible to incorrectly clustering or falsely matching features arising from anatomically unrelated cortical folds. As shown in Figure 2, given intrinsic error in the reference frame scale $d\sigma$ (multiplicative) and orientation $d\theta$, the number of such potential feature matches is a function of 1) the distance r_i between the feature and the reference frame and 2) the feature scale σ_i. In 2D volume renderings, for instance, number of potentially geometrically consistent yet false matches is:

$$FalseMatches \approx \frac{A_{arc}}{A_{m_i}} = \left(\frac{r_i}{\sigma_i}\right)^2 \frac{d\theta(d\sigma^2 - \frac{1}{d\sigma^2})}{2\pi}. \quad (3)$$

Thus, the smaller the feature scale σ_i or the larger the distance r_i from the reference frame, the larger the area A_{arc} in which geometrically similar yet false matches could arise. This is particularly problematic for features arising from cortical folds, whose scales are typically small relative to the size of the cortical surface and whose appearances may be ambiguous. We thus propose the distance-scale ratio $\frac{r_i}{\sigma_i}$ as a measure of the potential of a feature m_i to match unrelated cortical structures, which can be used to discard or filter such invalid model features.
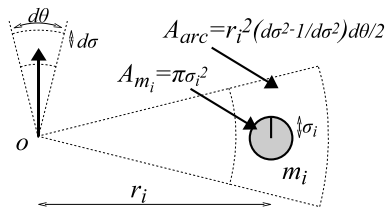
Fig. 2. The number of geometrically consistent yet false matches for a given feature m_i can be expressed as the ratio of the areas A_{arc} and A_{m_i}.

4. EXPERIMENTS

Experiments model the cortical surface in lateral volume renderings of the brain, using the lateral fissure as the reference frame as shown in Figure 1. The lateral fissure is an obvious structure that can be labeled with minimal expert knowledge. Using this reference frame, modeling is expected to be most effective in cortical regions near the origin where the ratio $\frac{r_i}{\sigma_i}$ is minimal, e.g. the inferior frontal gyrus or Broca's area. In general, reference frames can be defined throughout the cortex in order to investigate specific regions of interest. Lateral volume renderings are generated from MR volumes of the ICBM152 dataset using MRIcro [14], after applying the brain extraction tool [15] to strip the scalp. 98 unique subjects are used, mirroring hemispheres to obtain 196 images of resolution 217x181. Features are automatically extracted and represented using the scale-invariant feature transform (SIFT) technique [10], resulting in 500-600 features per image. A number of different feature types could be used, the SIFT technique is currently popular as it has been shown to outperform other techniques in terms of detection repeatability [11].

The set of learned cortical folding patterns is large ($\approx 10K$), and it is of interest to know which represent valid instances of the same neuroanatomical structure in different subjects. Exhaustively validating all model features is unreasonable, as most occur rarely in small numbers of subjects and are statistically insignificant. A number of model features do represent stable cortical patterns which occur in numerous subjects, however, and expert validation here is performed on a limited set of 10 of the 100 most frequently occurring model features. In selecting model features to validate, spurious features in the fringes of the image are disregarded, as are large stable features corresponding to cerebral lobes. For each of the 10 model features, a neuroanatomical expert is presented with 10 different occurrences automatically identified in different subjects. The expert first determines which underlying cortical structure is associated with the majority of the 10 feature instances, and then labels each individual instance as either correct (i.e. properly identified structure), incorrect or unsure. A histogram of results is shown in Figure 3, and Figure 4 illustrates occurrences of three different model features in 3 different subjects.
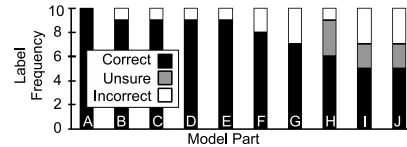

Fig. 3. A histogram of expert labeling results for 10 different model features A-J, based on 10 occurrences of each feature in different images. Of 100 occurrences, 77 are labeled as correct, 7 as unsure and 16 as incorrect by a neuroanatomical expert. Incorrect occurrences typically arise in similar yet incorrect nearby structure. Note how model inaccuracy is associated with human rater uncertainty (e.g. H, I and J).
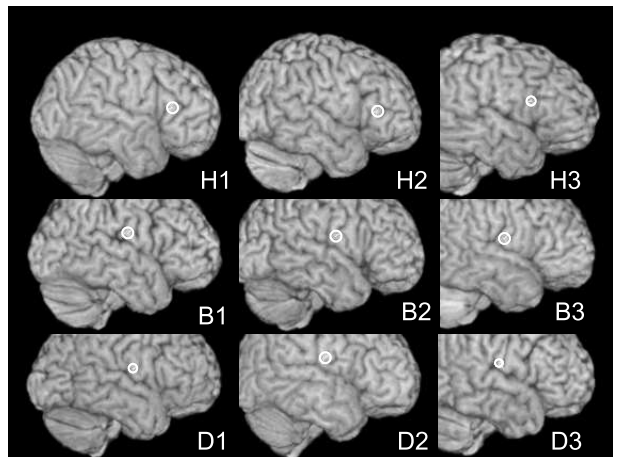

Fig. 4. Occurrences of model features B, D and H in 3 different subjects. Cortical patterns identified generally correspond to gyral bulges. H occurs anterior to the ascending sulcus in the pars triangularis of Broca's area, in 24% of subjects. Note that the image H3 was labeled an incorrect instance. B and D represent two distinct appearance variations of a similar section of the postcentral gyrus, and occur in 58% and 62% of subjects, respectively. Note that model features are not mutually exclusive and can co-occur in the same subject.

Many features identified through learning are invalid, and hamper the ability to analyze meaningful cortical folding patterns. The distance-scale ratio $\frac{r_i}{\sigma_i}$ presented in Section 3 is used here to filter model features due to invalid correspondences, and produce a more meaningful set of cortical folding patterns. To investigate the effect of filtering, the 10 most commonly occurring but invalid model features are identified, which clearly represent correspondences between unrelated cortical structure. All model features m_i such that $\frac{r_i}{\sigma_i} > 20$ are then filtered/removed, 20 is an empirically-determined threshold. Figure 5 illustrates the result of filtering, where the number of features is reduced from $\approx 10K$ to $\approx 2K$, the 10 most common invalid features are removed, while 9 of 10 validated features are preserved.
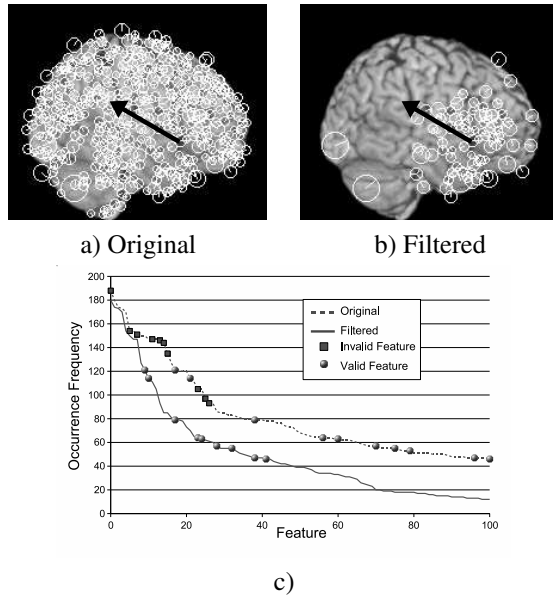
a) Original   b) Filtered

c)

Fig. 5. Images a) and b) show original and filtered features for a single image. Graph c) plots the 100 most frequently occurring model features, in descending order of occurrence frequency, for original and filtered feature sets. Note that filtering removes a large number of smaller/distant features, including the 10 most frequent incorrect model features, but preserves features which are larger/closer to the reference frame, including 9 of 10 validated features.

5. DISCUSSION

A technique is proposed for automatically learning new, unlabeled cortical folding patterns associated with a predescribed anatomical region of interest. A set distinct cortical patterns is learned from 196 lateral volume renderings, and expert validation of 10 different cortical patterns, each occurring in 10 in different subjects, shows 77% accuracy. Other interesting cortical patterns were identified, but could not be included in this paper due to the time required for expert validation. The ratio of the feature distance to the reference frame vs. feature scale $\frac{r_i}{\sigma_i}$ is shown to reflect the likelihood of false correspondence, and is used to effectively discard invalid model features arising from unrelated cortical structure.

Experimental results based on volume renderings are encouraging. Modeling could be extended to identify patterns from alternative data representations used in cortical analysis, e.g. surface representations [16, 4]. A variety of local feature types other than SIFT exist, and could be incorporated to identify a wider variety of cortical characteristics. We are currently investigating links between learned model features and activation regions within the brain, and the identification of clusters of similar folding patterns.

6. REFERENCES

[1] Ono, M. et al.: Atlas of the Cerebral Sulci. Thieme Medical, New York (1990)

[2] Naidich, T.P. et al.: Anatomic relationships along the low-middle convexity: Part I-Normal specimens and magnetic resonance imaging. Neurosurgery **36**(3) (1995) 517–532

[3] Fischl, B. et al.: High-resolution inter-subject averaging and a coordinate system for the cortical surface. Human Brain Mapping **8** (1999) 272–284

[4] Tosun, D. et al.: A geometry-driven optical flow warping for spatial normalization of cortical surfaces. IEEE TMI **27**(12) (2008) 1739–1753

[5] Hellier, P. et al.: Coupling dense and landmark-based approaches for non rigid registration. IEEE TMI **22**(2) (2003) 217–227

[6] Mangin, J. et al.: Object-based morphometry of the cerebral cortex. IEEE TMI **23**(8) (2004) 968–983

[7] Riviere, D. et al.: Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. MIA **6** (2002) 77–92

[8] Zheng, S. et al.: A learning based algorithm for automated extraction of the cortical sulci. In: MICCAI. (2006) 241–248

[9] Toews, M. et al.: A statistical parts-based appearance model of anatomical variability. IEEE TMI **26**(4) (2007) 497–508

[10] Lowe, D.G.: Distinctive image features from scale-invariant keypoints. IJCV **60**(2) (2004) 91–110

[11] Mikolajczyk, K. et al.: Scale and affine invariant interest point detectors. IJCV **60**(1) (2004) 63–86

[12] Mazziotta, J. et al.: A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (ICBM). Philos Trans R Soc Lond B Biol Sci **356**(1412) (2001) 1293–1322

[13] Cachier, P. et al.: Multisubject non-rigid registration of brain MRI using intensity and geometric features. In: MICCAI. (2001) 734–742

[14] Rorden, C.: MRIcro. (www.sph.sc.edu/comd/rorden/mricro.html)

[15] Smith, S.: Fast robust automated brain extraction. Human Brain Mapping **17** (2002) 143–155

[16] Fischl, B. et al.: Cortical surface-based analysis ii: Inflation, flattening, and a surface-based coordinate system. NeuroImage **9**(2) (1999) 195–207