

Exploring Cortical Folding Pattern Variability Using Local Image Features

Rishi Rajalingham¹, Matthew Toews², D. Louis Collins³, and Tal Arbel¹

¹ Center for Intelligent Machines, McGill University

² Brigham and Women's Hospital, Harvard Medical School

³ Montreal Neurological Institute, McGill University

Abstract. The variability in cortical morphology across subjects makes it difficult to develop a general atlas of cortical sulci. In this paper, we present a data-driven technique for automatically learning cortical folding patterns from MR brain images. A local image feature-based model is learned using machine learning techniques, to describe brain images as a collection of independent, co-occurring, distinct, localized image features which may not be present in all subjects. The choice of feature type (SIFT, KLT, Harris-affine) is explored with regards to identifying cortical folding patterns while also uncovering their group-related variability across subjects. The model is built on lateral volume renderings from the ICBM dataset, and applied to hemisphere classification in order to identify patterns of lateralization based on each feature type.

1 Introduction

The cerebral cortex is characterized by complex folding patterns created by ridges called gyri, and fissures called sulci. The exploration and identification of these cortical folding patterns is important in the development of a general functional atlas of the cortex, as folds are linked to the division of functional areas in the brain. However, this task is difficult due to the fact that cortical structures may vary significantly in shape, size and appearance across subjects, or may not even be present in all subjects. This phenomenon is known as inter-subject variability, the variation in feature morphology across different subjects caused by inherent anatomic variability or pathology, and renders it difficult even for experts to identify folding patterns [1].

Uncovering cortical variability can be instrumental in the understanding of how various anatomical structures express themselves within and across subject groups; an appropriate quantification of variability can answer fundamental questions in medicine, such as which cortical folds are common or rare across a particular population. Quantifying anatomical variability at the group level, based for example on pathology, can give insight as to how similarities or differences in expression of particular folds are related to subsets of the population. This can then be used as a means of evaluating similarity of a new subject to established groups.

The conventional approach to examine and identify cortical folding patterns relies on manual labelling by experts, which is a tedious task that is subject to inter-rater variability. For this reason, a significant amount of work has been done on automatic, data-driven techniques to learn cortical folding patterns and study their variability using morphometric techniques on magnetic resonance imaging (MRI) data of the brain. Morphometry is the analysis of variation in form (appearance, shape or size) of objects. In this context, it aims to identify anatomical variability of cortical folds between subjects. Most automatic analyses rely firstly on warping each brain image to a common reference frame, using voxel intensities [2], landmark objects [3], cortex-specific features [4], to name a few, and computing statistical quantities based on this registration. These methods assume or even force a one-to-one correspondence between subject brains that may not generally exist. Anatomical variability in folding patterns can lead to imperfect registration, resulting in the averaging out of structural differences. Consider the case where a particular anatomical structure expresses itself in multiple, distinct types across a population, or only expresses itself in part of the population. This type of group-related variability cannot be uncovered by most morphometric techniques.

A number of methods address this issue by loosening the one-to-one correspondence assumption, by means of multiple atlases [5] or atlas stratification [6], or residual error components [7]. Feature based morphometry (FBM) is a probabilistic, parts-based model of images across subjects which effectively ignores the lack of a one-to-one correspondence in inter-subject registration, and rather attempts to identify local regions of images whose occurrences in groups is statistically significant [9]. This method has been applied to 2D MR slices to model anatomical variability [8], and volumetric brain images to discover group-related anatomical patterns [9]. It has had some success with surface renderings of the cortex [10], but was limited by the choice of local feature representation.

The main contribution of this paper is the introduction of an FBM technique for automatically learning new, unlabelled cortical folding patterns from a large set of subject brain images, attempting not only to cope with variability, but to uncover it in a completely data-driven, bottom-up fashion. The challenge lies in automatically identifying instances of the same folding pattern in different subjects, while not forcing correspondences (as they may not exist or vary significantly in some subjects). In particular, the goal is to automatically detect multiple distinct modes of appearance of particular cortical folds, in order to expose patterns of variability. This leads to several advantages over current approaches. Firstly, this permits the development of a bottom-up, parts-based description of cortical folds in healthy brains that copes with the inherent inter-subject variability. Secondly, uncovering this variability can also be helpful for numerous clinical applications by revealing which cortical folding patterns are common or rare across a subset of the population defined by pathology (e.g. schizophrenia, autism, some forms of epilepsy) versus the norm. The automatic, bottom-up nature of this system makes it possible to extract meaningful patterns, thus potentially leading to breakthroughs

in our understanding of neurological diseases, without requiring particular prior knowledge regarding the disease in question.

Many powerful local features have been developed in the field of computer vision, particularly for application areas such as object detection or classification. In this paper, we wish to explore the application of established features from the computer vision literature to this domain for the first time. In particular, we explore the power of different specifically chosen local image features to represent cortical folding patterns. Using machine learning techniques, a feature-based model is built to describe brain images as a collage of independent, co-occurring, distinct, localized image features which may not be present in all subjects. This approach is completely bottom-up, which does not restrict the discovery of variability to explicitly extracted sulcal structures, but rather automatically exposes patterns in the image independently of knowledge from particular anatomical structures. Three well-established feature types are explored (SIFT, KLT, Harris-affine), and their usefulness in learning cortical folding patterns is reported. By comparing them, the importance of selecting context-specific descriptions is emphasized. We also apply this model to hemisphere classification, in order to identify patterns of lateralization, and identify the usefulness of different feature types for this task.

The remainder of this paper is organized as follows. Section 2 describes the various local feature types used and reviews the modelling algorithm. Section 3 presents experimental results and discussions on modelling 196 lateral volume renderings of the International Consortium of Brain Mapping (ICBM) [11] data set for the tasks of identifying folding patterns and classifying hemispheres.

2 Feature-Based Modelling

2.1 Local Image Features

Local image features are distinct image patterns that can be detected automatically and robustly based on a saliency criterion. Depending on the saliency criterion of choice, an extracted local feature $f = (g, a)$ can be characterized geometrically by parameters including position, scale, orientation and eccentricity $g = (x, \sigma, \theta, e)$ of the local region. Its appearance is described by a 128-dimension vector a , storing gradients computed within the local region.

We use three different types of local features, and compare their performance in accurately describing cortical structures. We note that their relevance to the context plays a significant role in the usefulness of the model, both in describing brains as a collage, and in identifying class-distinctive features.

1. **SIFT:** Lowe’s Scale invariant feature transform (SIFT) [12] evaluates saliency in the Gaussian scale-space. Although Gaussian scale-space maxima is effective for a wide range of applications, because it is generic enough to extract local features when there is no prior on the features of interest, it has had limited success with images of the cortex. SIFT generally extracts salient blob-like features, corresponding in this context to specular gyral reflections. See Fig 1(a) for examples.




Fig. 1. Examples of extracted features of each type, (a) SIFT, (b) KLT (c) Harris-affine

2. **KLT:** The Kanade-Lucas-Tomasi (KLT) [13] feature tracker defines saliency of local image patches using the Harris corner measure with uniform weighting function. Corner detectors are relevant for the cortex, as sulcal junctions and intersections are folding patterns of interest. See Fig 1(b) for examples.
3. **Harris-affine detector:** Mikolajczyk and Schmid’s affine invariant key-point detector [14] uses the Harris-Laplace transform for initial region point localization, and is robust to affine transformations by describing affine regions. The scale-invariant, elliptical affine regions are useful in describing elongated structures of interest, e.g. long gyri/sulci, of the cortex. See Fig 1(c) for examples.

2.2 Learning a Model

The subject images are assumed to be aligned after a pre-processing normalization step consisting of a 12-parameter affine image alignment, rather than a non-linear registration to a common reference frame. Local image features are extracted and stored for each subject image.

Learning a model effectively consists of identifying a set of model features $\{m_i\}$, consisting of cluster centroids of the very large set of extracted local image features $\{f_i\}$ [9] by identifying for each f_i , others which are similar in geometry and appearance. In fact, two different clusters G_i, A_i are constructed based on similarity in geometry and appearance respectively as follows:

$$G_i = \{f_j : \frac{\|x_i - x_j\|}{\sigma_i} < \varepsilon_x, \log|\frac{\sigma_j}{\sigma_i}| < \varepsilon_\sigma, \log|\frac{\theta_j}{\theta_i}| < \varepsilon_\theta, \frac{|e_i - e_j|}{e_i} < \varepsilon_e\}, \quad (1)$$

$$A_i = \{f_j : \|a_i - a_j\| < \varepsilon_{a_i}\} \text{ where } \varepsilon_{a_i} = \sup\{\varepsilon_a > 0 : \frac{A_i(\varepsilon_a) \cap G_i}{A_i(\varepsilon_a) \cap \tilde{G}_i} > 1\}. \quad (2)$$

Given A_i, G_i , the model cluster associated with f_i is $M_i = A_i \cap G_i$, i.e. corresponding to similarity in both appearance and geometry. Features with strong appearance similarity are allowed to cluster together within a lenient geometrical window. Moreover, given similar geometry, features with significant appearance dissimilarity aren’t forced to cluster together, thus allowing for multiple appearances modes at any particular location on the cortex.

Since several f_j (specifically those in M_i) will generate quasi-identical model clusters, the cluster set $\{M_i\}$ is pruned to remove this redundancy. Pruning consists of automatically removing, for all i , clusters M_j such that $|M_j| < |M_i|$ and $|M_i \cap M_j| > \varepsilon_M |M_j|$, where ε_M is an empirically determined pruning ratio. Additionally, degenerate clusters consisting of few features are also deleted. Finally, the set of model features is achieved by computing the centroids of model clusters $m_i = \langle M_i \rangle$.

The model features $\{m_i\}$ give insight to inherent group related variability of anatomical structures, as any particular feature is not necessarily found in all images. Rather, each subject’s brain is modelled as an independent mosaic of local features, similar to multi-atlas approaches.

Furthermore, our feature-based modelling can automatically uncover this variability with a secondary clustering/pruning step on the set of model features. Features m_{i_k} , $k = 1, 2, \dots, N$ with similar geometry (as defined in Eq. 1) are grouped together if they are also similar in appearance (as defined in Eq. 2, using a global ε_a), and averaged: $m_i = \langle \{m_{i_k}\}_k \rangle$. This removes multiple model feature representations of the same or similar image feature, and ideally reduces the number of model features at any particular location to a few distinct ones. In the context of sulci, these are representative of group-related variability of the cortex, corresponding to possible multiple distinct modes of appearance of particular cortical folds arising across a population.

3 Experimental Results and Discussion

3.1 Identifying Cortical Folds

Our first goal is to explore the cortical variability over a population of healthy subjects. For this reason, we chose to extract volumetric brain image data from the ICBM dataset [11]. It is common to first perform a continuous deformation, such as a spherical [16] or conformal [17] mapping, of the extracted cortical surface prior to looking at cerebral sulci. Rather than projecting the cortex onto a sphere, we chose to use a lateral projection. While this mapping is not a homeomorphism, it preserves the appearance of folding patterns with minimal perspective distortion within the region of interest.

Using the MRICro software after applying the brain extraction tool, lateral renderings of the cortical surface, i.e. intensities in 2D lateral volume renderings defined by perspective projection, were outputted in image format. 98 unique subjects are used, mirroring hemispheres to obtain 196 images of resolution 217x181. In lateral views, the voxel resolution is 1mm isotropic, which is more than sufficient to visualize gyral/sulcal patterns in the laterally oriented surfaces.

Local features (KLT, SIFT, and Harris-affine) are automatically extracted from these images. KLT features are quite rare, averaging 50 per image, while SIFT and Harris-affine detectors can extract 400-500 features per image. Foreshortening along the edges of the brain due to projection results in lower resolution on the edges of the brain, therefore these features are disregarded. Additionally, all features are filtered by scale to remove extractions of small




Fig. 2. Frequency of occurrence plot, illustrating the most frequent model features for each type

specular reflections and large structures, as the features of interest are at a characteristic scale determined by gyri/sulci dimensions. A feature-based model is built, as described in Section 2.2, separately for each feature type.

Figure 2 illustrates the frequency of model feature occurrences per type. The relatively small SIFT model size seems to imply that this generic feature does not have the matching capabilities of both KLT and Harris-affine features. Note that the most frequent features of each type correspond to various anatomical structures with low variability, tuned to the feature saliency criteria. As seen in Figure 2, the KLT corner detector reliably identifies the strong corner arising due to the cerebellum (shown in red), while the Harris-affine detector seems to extract the slightly more elongated tri-corner of the pars triangularis in Broca’s area (shown in green). In general, SIFT detects specular reflections off gyri, but its most frequent feature coincides with a fold near the mid temporal sulcus (shown in blue).

Figure 3 illustrates examples of model features for each type: occurrence matches are shown for two different, typical SIFT (S_1, S_2), four KLT (K_1, K_2, K_3, K_4), and four Harris-affine (H_1, H_2, H_3, H_4) model features across nine arbitrary subjects. The KLT tracker and Harris-affine detector are able to identify plausible cortical structures much better than the generic SIFT. KLT accurately detects strong corners arising from sulcal junctions, while the complementary Harris-affine features often coincide with elongated structures such as linear sulci and gyri edges. Although expert validation is needed to claim that these learned features correspond to valid semantic structures of the cortex, current results are encouraging for the context-specific features (KLT, Harris-affine), especially considering that matches are made despite significant anatomical variability. For example, H_1 is found at varying geometrical positions (see Figure 3), but seems to consistently correspond to the superior temporal gyrus. Similarly, H_4 accurately captures the length of the lateral sulcus and superior temporal sulcus, despite the variability in appearance, orientation and scale of these structures.

Additionally, the system automatically identifies examples of local regions of the cortex with at least two distinct modes of appearance, using the secondary

pruning step described in Section 2.2 based on a KLT model. Figure 4 shows three examples of such regions, with the corresponding appearance modes shown as image gradients. The distinct modes could correspond to the presence/absence of bridges or links between neighbouring gyri in parts of the population. It is noteworthy that, under inter-subject correspondence assumptions, these folding patterns would have been merged, and the underlying group-related variability would be lost in the registration.

3.2 Hemisphere Classification

In order to validate that the features we are examining do actually correspond to meaningful anatomical patterns, we now look at how they can be used to extract group-based differences. For the purposes of this paper, we explore the example toy application of hemisphere classification based on cortical folding patterns. The goal is not to establish an optimal strategy for hemisphere classification. Indeed, simpler methods can be successfully employed to address this particular task. Instead, the aim is to explore how well the particular features chosen can be used to automatically identify patterns of sulcal folds distinctive to a hemisphere, possibly indicative of lateralization of the brain.

Hemisphere classification is done for 98 subjects using 20-fold cross validation. Subject images (two hemispheres per subject) are either both in the training set or both in the test set, to avoid errors due to possible subject-specific feature correlation. Otherwise, training and test sets are randomly created. In Table 1, hemisphere classification error rates are shown for models based on three local feature types (KLT, SIFT, Harris-affine) using a support vector machine (SVM). Although our error rates are not comparable to other methods [15] which are specifically tailored to this problem, these results reveal the potential of classification based on various image features extracted from lateral renderings in a completely bottom-up fashion, without prior information about the problem context. In particular, tailoring features to the context seems to increase the performance of a classification model.

Table 1. 20-fold Cross Validation Error

-	KLT	SIFT	Harris-affine
SVM Error	0.24	0.45	0.33

While classification results are promising, our main goal is to explore how class-distinctive model features can be established automatically. Figure 5 shows the statistically significant class-distinctive model features, along with the corresponding contingency tables and significance values (note $p \ll 0.05$). The results of hemisphere classification indicate that it is likely not particular cortical folds that are most hemisphere-specific, but rather the overall geometry. The most distinguishing KLT model features, as shown in Fig 5(a), are corners that occur on and around the cerebellum, indicating that the size and shape of the temporal




Fig. 3. Each grouping of nine images corresponds to occurrence matches of a particular model feature. Two different SIFT S_1, S_2 (red circles), four KLT K_1, K_2, K_3, K_4 (red squares), and four Harris-affine H_1, H_2, H_3, H_4 (yellow ellipses) features are shown.

lobe, which can obstruct this feature, is related to lateralization. This geometric warp between hemispheres is well studied [18]. However, sulcal folds can also be hemisphere-specific: the most distinguishing Harris-affine model features seem to occur on the pars triangularis of the inferior frontal gyrus (see Fig 5(b)), which is known to be larger on the left hemisphere [19].





Fig. 4. Automatically found multiple distinct modes of appearances: Three local regions are shown on a brain template, with corresponding appearance modes shown as image gradients of folds



	m_i	\bar{m}_i
Right	12	86
Left	37	61

(c)

$$p = 2.7783 \times 10^{-5}$$

	m_i	\bar{m}_i
Right	65	33
Left	89	9

(d)

$$p = 2.1535 \times 10^{-5}$$

Fig. 5. Significant class distinctive model features (a) found on the cerebellum (KLT) and (b) near the pars triangularis of Broca's area (Harris). Corresponding contingency tables and p-values are shown in (c), (d).

4 Conclusion

The contribution of this paper is a feature-based model of the cortex which aims to automatically uncover unlabelled cortical folding patterns from a large set of subject brain images. By describing brain images as a collection of independent, co-occurring, distinct, localized image features, the feature-based model is able to identify instances of the same folding pattern in different subjects while not forcing correspondences. Consequently, the model can automatically uncover group-related variability as well as multiple distinct modes of appearance of particular cortical folds. This will be helpful in complementing the laborious,

manual studies by neuroanatomists is developing an understanding of the the cortical variability of healthy brains, whose goal is to develop a general functional atlas of the brain. Uncovering this variability can be helpful for numerous clinical applications, by giving insight into which cortical folding patterns are common or rare across a subset of the population defined by pathology, for example. For instance, looking at brains of subjects with schizophrenia or autism may reveal differences in expression of particular folds with respect to the norm. This could be potentially instrumental in permitting us to learn about diseases where such variability is difficult to establish. As well, it could then be used as a means of evaluating similarity of a new subject to these pathological groups.

Variou feature types (SIFT, KLT, Harris-affine) were explored in this work, and resulting models were also applied for hemisphere classification, exposing the power of different features and classifiers for this task. The class-distinctive features that were discovered may be indicative of brain lateralization. However, expert validation on the learned features is a necessary future step to confirm this. Interesting future work will involve applying this framework to subjects with schizophrenia or Alzheimer's, in order to analyze the class-distinctive features for this group. Future work could also involve the conjunction of features optimized for the task of cortical modelling and classification.

References

1. Ono, M., et al.: Atlas of the Cerebral Sulci. Thieme Medical (1990)
2. Toga, A.W., et al.: Probabilistic approaches for atlas normal and disease-specific brain variability. *Anat. Embryol.*, 267–282 (2001)
3. Mangin, J., et al.: Object-Based Morphometry of the Cerebral Cortex. *IEEE TMI* (2004)
4. Mangin, J. et al.: A framework to study the cortical folding patterns. *NeuroImage* (2004)
5. Klein, A. et al.: Mindboggle: a scatterbrained approach to automate brain labeling. *NeuroImage* (2005)
6. Blezek, D.J., Miller, J.V.: Atlas stratification. In: Larsen, R., Nielsen, M., Sparring, J. (eds.) *MICCAI 2006*. LNCS, vol. 4190, pp. 712–719. Springer, Heidelberg (2006)
7. Baloch, S., et al.: An anatomical equivalence class based joint transformation-residual descriptor for morphological analysis. In: Karssemeijer, N., Lelieveldt, B. (eds.) *IPMI 2007*. LNCS, vol. 4584, pp. 594–606. Springer, Heidelberg (2007)
8. Toews, M., Arbel, T.: A Statistical Parts-based Appearance Model of Anatomical Variability. In: *IEEE TMI*, pp. 497–508 (2007)
9. Toews, M. et al.: Feature-Based Morphometry: Discovering Group-related Anatomical Patterns. *NeuroImage* (2009)
10. Toews, M., et al.: Automatically Learning Cortical Folding Patterns. In: *IEEE ISBI*, pp. 1330–1333 (2009)
11. 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), 1293–1322 (2001)
12. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* 60(2), 91–110 (2004)

13. Tomasi, C., Shi, J.: Good Features to Track. In: CVPR, pp. 593–600 (1994)
14. Mikolajczyk, K., et al.: Scale and affine invariant interest point detectors. *IJCV* 60(1), 63–86 (2004)
15. Duchesnay, E., et al.: Classification from cortical folding patterns. *IEEE TMI* 26(4), 553–565 (2007)
16. Xu, C., et al.: A Spherical Map for Cortical Geometry. In: Proc. Int. Conf. Functional Mapping Human Brain, pp. 73–74 (1998)
17. Angenent, S., et al.: On the Laplace-Beltrami operator and brain surface flattening. *IEEE Trans. Med. Imag.* 18(8), 700–711 (1999)
18. Lyttelton, O., et al.: Positional and Surface Area Asymmetry of the Human Cerebral Cortex explored through automated surface-based analysis. *Neuroimage* (2009)
19. Foundas, A.L., et al.: Pars triangularis asymmetry and language dominance. *Proc. Natl. Acad. Sci. USA* 93(2), 719–722 (1996)