Surface-Based Spatial Pyramid Matching of Cortical Regions for Analysis of Cognitive Performance

Kristen M Campbell[1][0000−0001−8196−4423], Jeffrey S Anderson[2][0000−0002−9669−3846], and P Thomas Fletcher[3]

[1] Scientific Computing & Imaging Institute, University of Utah, Salt Lake City, UT, kris@sci.utah.edu

[2] Department of Radiology & Imaging Sciences, University of Utah, Salt Lake City, UT, andersonjeffs@gmail.com

[3] University of Virginia, Charlottesville, VA

Abstract. We propose a method to analyze the relationship between the shape of functional regions of the cortex and cognitive measures, such as reading ability and vocabulary knowledge. Functional regions on the cortical surface can vary not only in size and shape but also in topology and position relative to neighboring regions. Standard diffeomorphism-based shape analysis tools do not work well here because diffeomorphisms are unable to capture these topological differences, which include region splitting and merging across subjects. State-of-the-art cortical surface shape analyses compute derived regional properties (scalars), such as regional volume, cortical thickness, curvature, and gyrification index. However, these methods cannot compare the full extent of topological or shape differences in cortical regions. We propose icosahedral spatial pyramid matching (ISPM) of region borders computed on the surface of a sphere to capture this variation in regional topology, position, and shape. We then analyze how this variation corresponds to measures of cognitive performance. We compare our method to other approaches and find that it is indeed informative to consider aspects of shape beyond the standard approaches. Analysis is performed using a subset of 27 test/retest subjects from the Human Connectome Project in order to understand both the effectiveness and reproducibility of this method.

1 Introduction

Much work has been done to understand what parts of the brain are responsible for which functions. This localization of function is important for fundamental understanding of how the brain works and for neurosurgical tasks such as planning regions to target or avoid during the removal of tumors [2]. While localization of function is generally consistent across individuals, there is variability in the shape of functional regions.

Glasser et al. showed in [3] that certain cortical regions can have multiple *topological* variants that are present in a significant percentage of the population.

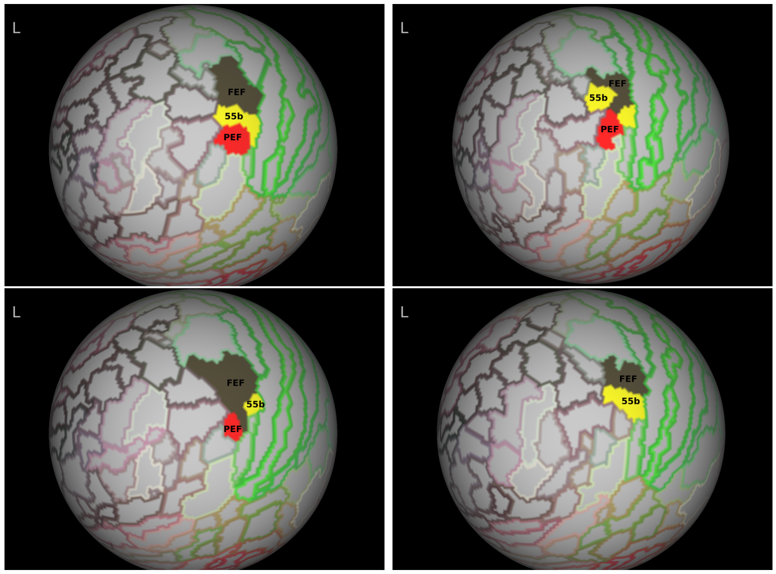

Fig. 1: Four examples of the variety of 55b, FEF, and PEF configurations. Each example is the cortical surface of a subject's left hemisphere projected onto a sphere.

In particular, they highlight the 55b region, a lightly myelinated area of the premotor cortex that separates the frontal eye field (FEF) and premotor eye field (PEF). However, there is a consistent variant in which 55b is split into two disjoint pieces, with FEF in between the split sections. A second variant exists in which 55b can be shifted relative to FEF, so that FEF is now between 55b and PEF. See Figure 1 for examples of the variety in 55b, FEF, and PEF configurations. Furthermore, it is possible that a region is missing in a subject. As an example, PEF is missing in the bottom-right subject in Figure 1.

The presence of these and other topological variations in a large subset of individuals leads to the question of whether these variants are related to any sort of change in cognitive or social ability, psychological well-being, or personality. To answer this question, we first need a statistical shape analysis method that can (1) compare shapes with differing topology and (2) handle the possibility of missing shapes, that is, the method allows the null set to be a valid shape.

Existing shape correspondence methods require shapes to be prealigned and need to find a one-to-one mapping for each point around the shape [10]. Diffeomorphic methods such as large deformation diffeomorphic metric mapping (LDDMM) [1] and spherical demons [12] assume that a diffeomorphism exists that can deform one shape into the other. Neither shape correspondence nor diffeomorphic methods can compare one region to a region split into two separate pieces.

Other shape analysis methods allow topologically different shapes, but do not allow missing shapes. Leventon et al. [6] represent shapes as level sets of the signed distance from the shape boundary. This method does not have a unique representation for a shape and cannot handle the null set. Another approach is the Gromov-Hausdorff metric [9], which provides a valid distance metric between shapes and allows multiple topologies, but does not allow for a shape to be the null set.

Another metric that we could use is the Jaccard distance that measures the difference between the size of union and the size of intersection of two sets. This metric allows one of the shapes to be the null set, but it does not differentiate between nonoverlapping sets that are near versus far away from each other. Both of these cases will have the same distance of 0.

We propose a new distance metric on shapes defined as regions on a sphere, called icosahedral spatial pyramid matching (ISPM). This metric is an extension of spatial pyramid matching (SPM) [5, 13] defined for images. ISPM can measure distances between shapes that are not topologically equivalent and differences in configurations of shapes relative to each other. Our goal is to analyze relationships between the shapes of cortical regions and cognitive measures. However, standard correlation or regression approaches are not applicable, since our shapes do not live in a Euclidean vector space, but rather a non-trivial metric space defined through the ISPM metric. Instead, we use a statistic called distance correlation [11], which measures the dependency between two random values sampled from metric spaces.

2 Methods

ISPM computes spatial pyramid matching (SPM) on a mesh of a spherical surface using labeled borders of parcellated regions as features. We will review d-dimensional image space SPM here and then describe how we adapt the computation to an icosahedral mesh of a spherical surface. Once we can compute distances between shape complexes on a sphere, we show how to use distance correlation in order to find relationships between shape and cognitive performance measures.

2.1 Spatial Pyramid Matching

SPM [5] is a measure of similarity of features between two images, A and B, using the pyramid match kernel (PMK) from [4] while preserving spatial information on a rectangular d-dimensional image grid. SPM starts by separating each image into C channels, one channel per feature type, where each channel contains the spatial locations where that feature type is found in the image.

We then construct a pyramid of L resolution levels, where the finest resolution is level L, and the number of bins in a level, D, decreases by a factor of 2 in each direction at each level.

For each resolution level, l, in the pyramid, a spatial histogram is computed for each channel, c, that counts how many times that feature appears within the spatial extent covered by each bin. The amount a feature "matches" at a particular level of resolution can be measured by the intersection of the histograms $I(H^l_{A_c}, H^l_{B_c})$ for A_c and B_c. Histogram intersection is used here because it has the useful property of being positive-definite, as shown in [4]. The intersection of a channel histogram at a particular level is

$$\mathrm{I}^l_c(H^l_{A_c}, H^l_{B_c}) = \sum_{i=1}^{D} \min(H^l_{A_c}(i), H^l_{B_c}(i)). \tag{1}$$

To avoid counting the same contribution at more than one level, we look at the number of new matches, N^l, at each level:

$$N^l_c(A_c, B_c) = \mathrm{I}^l_c(H^l_{A_c}, H^l_{B_c}) - \mathrm{I}^{l+1}_c(H^{l+1}_{A_c}, H^{l+1}_{B_c}), \tag{2}$$

$$N^L_c(A_c, B_c) = \mathrm{I}^L_c(H^L_{A_c}, H^L_{B_c}). \tag{3}$$

The pyramid match kernel is the weighted combination of these new matches across all levels:

$$\kappa_c(A_c, B_c) = \sum_{l=0}^{L} \frac{1}{2^{L-l}} N^l_c(A_c, B_c). \tag{4}$$

The spatial pyramid matching is the normalized combination of these kernels across the different channels:

$$K_S(A, B) = \frac{\sum_{c=1}^{C} \kappa_c(A_c, B_c)}{\sqrt{\sum_{c=1}^{C} \kappa_c(A_c, A_c) * \sum_{c=1}^{C} \kappa_c(B_c, B_c)}}. \tag{5}$$

Because histogram intersection is used, and κ_c sums these intersections with a weight that is decreasing as the grid coarsens, the pyramid match kernel is positive-definite, as shown in [4]. The positive-definiteness continues to be maintained when summing across channels, and thus K_S is a Mercer kernel since κ_c is one. Since $K_S(A, A) = 1$ for all images, A, we convert this similarity kernel to a distance metric, $d_S(A, B)$, scaling by 0.5 to keep $0 \leq d_S(A, B) \leq 1$ as follows:

$$\begin{aligned} d_S(A, B) &= \frac{1}{2}(K_S(A, A) + K_S(B, B) - 2K_S(A, B)) \\ &= 1 - K_S(A, B). \end{aligned} \tag{6}$$

2.2 Icosahedral Spatial Pyramid Matching

We adapt SPM from a d-dimensional image grid to an icosahedral mesh approximation of a spherical surface so that we can compute the similarity of cortical surface features. One way to produce a regular mesh of a spherical surface is to start with an icosahedron, subdivide each triangle face into 4 smaller triangles, project the new vertices onto the sphere and repeat until the desired level of approximation of the sphere has been achieved.

We use this process in reverse to generate our sequence of pyramid levels. We start with triangles on the finest mesh, merging them with the 3 neighboring triangles that share a common edge into larger triangles. Whereas the SPM pyramid reduces the number of bin cells by a factor of 2^d at each level, ISPM reduces the number of bin faces by a factor of 4 at each level and the number of vertices by a factor of 2. We implemented this pyramid generation as a histogram resampling method added to the Connectome Workbench tool [8].

Once the pyramid is constructed, we proceed as for SPM, computing the histogram at each level by counting the number of features of a particular type that fall within a spatial bin and then summing across feature channels. We use the same weighting scheme as above to combine across pyramid levels. Since we are using a decreasing weighting scheme and summing positive-definite kernels, we preserve the Mercer kernel property, and ISPM is a distance metric.

We use the icosohedral mesh instead of a cuboid or uv mesh so that each bin in each spatial histogram has nearly equal area, and so that we have a simple way to generate the pyramid of histograms. If different features like structure tensors [13] are used instead of the region membership used here, then care must be taken to rotate the features appropriately as they are moved around on the surface of the sphere.

2.3 Distance Correlation

We want to look at whether there is a relationship between the shape of functional regions and cognitive measures, but we only have distances between shape complexes, not an individual measurement of shape. Therefore, we ask the question: *If two subjects have similar (or dissimilar) shapes of functional regions, do they also have similar (or dissimilar) cognitive measures?* We address this question by using the distance correlation, dCor [11], between two distance matrices, which tests the joint independence of two random variables of arbitrary dimension. The dCor is zero when the random variables are independent and otherwise ranges between 0 and 1. The dCor t test has a null hypothesis that the two variables are independent (i.e., dCor = 0). A small p-value of this test rejects this null hypothesis and indicates that the dCor is significantly different from 0, with dCor representing the amount of dependence.

Lyons [7] extended the theory of distance correlation to metric spaces and showed that it is necessary and sufficient for the metric space to be of the strong negative type to test independence. Because we were careful to ensure ISPM is a distance metric, we are able to use dCor to test independence. This property also holds for the Jaccard distance metric we evaluate in the results section.

To compare cognitive measures to the ISPM distances, we compute the pair-wise distance between each subject's value of that cognitive measure. To analyze k cognitive measures concurrently, we can construct a k-vector of the measures and compute the pair-wise Euclidean distances between these vectors. One of the many attractive properties of the distance correlation is that it is scale-invariant.

Thus, we do not need to center or scale the variables prior to computing the distance correlation. Instead, *dcor.ttest* from the *energy* package in R performs a U-centering of each of the distance matrices and then provides a bias-corrected dCor and p-value. After bias correction, the dCor values are shifted, and may fall slightly below 0.

3 Results

To gain insight into whether the shape and topological variability of cortical regions is related to cognitive performance, we looked at the distance correlation between the ISPM of cortical regions and several cognitive measures. We also compared the ISPM distance metric performance to that of a more commonly used metric, the Jaccard index (JI). Finally, we looked at whether the ISPM metric is more informative than using measures like cortical thickness or surface area of cortical regions.

Because we were specifically interested in looking at topological variability of region configuration, we needed to start with a cortical region parcellation that allows for topological variation. The Glasser 2016 multi modal parcellation (MMP) finds high-quality subject-specific parcellations using an areal classifier that defines 180 regions per hemisphere based on features such as myelin maps, resting state fMRI, and cortical folding patterns. Regions are constrained to be in proximity to the group parcellation, but are not constrained to be topologically equivalent. The Glasser parcellations are publicly available for 27 subjects who are in the Test/ReTest subgroup of the Human Connectome Project (HCP) Young Adult dataset. Each subject was scanned twice and the subject scans were processed independently through the HCP pipelines. The cortical surfaces of each hemisphere were inflated onto a sphere where each vertex in channel, c, is labelled 1 if that vertex followed the border of a parcellation region c. We computed ISPM for each hemisphere and then summed them to get a total distance between subjects. We compared dCor and p-values for the Test and ReTest data sets to get a sense of the consistency of these results.

We chose four cognitive measures for analysis. The picture vocabulary test (PVT), reading test (Read), processing speed (Speed), and fluid cognition composite score (Fluid). PVT and Read were chosen to evaluate relationships between language-related areas and language ability, whereas Speed and Fluid where chosen since they are representative of more global relationships between all brain regions and cognitive ability. We used the age-adjusted version of cognitive measures whenever possible. All p-values shown were corrected for multiple comparisons with FDR.

We started by looking at the relationship between the ISPM of the cortical regions and each of these cognitive measures separately and then combined into a 4-vector (All). As shown in Table 1, we see that the distance correlation of 0.21 between ISPM and fluid cognition is consistently significant, but there is more variability in Test/ReTest in reading and processing speed.

We compared our results to a more standard metric for the comparison of segmentations, the Jaccard index, $d_J(A,B) = \frac{|A \cup B| - |A \cap B|}{|A \cup B|}$. Although JI is often used to compare segmentation results, this is the first time it has been used in shape analysis of cortical regions. We see in Table 1 that although JI did correlate significantly with fluid cognition, the distance correlation between JI and fluid cognition decreased to 0.15 for Test, 0.18 for ReTest. Other results were less significant or not significant compared to ISPM. As we expect, JI and ISPM were strongly related where the dCor of JI and ISPM was 0.905 with a p-value $<$ 2.2e-16.

We computed the surface area and mean cortical thickness of each of the 360 regions. We then compared the distances between each subject's 360-d thickness vector and the same cognitive measures as above. We repeated this comparison for the 360-d surface area vector. As you can see in Table 1, there were some significant correlations between surface area and fluid cognition and also correlations with the four combined measures for the ReTest dataset. The correlations were smaller and less consistent than the JI and ISPM results, with the exceptions of a significant correlation between surface area and Read in the ReTest data set, and a stronger correlation of the surface area vs all four combined measures than JI and ISPM had for the ReTest dataset. The dCor of mean cortical thickness and ISPM was 0.275 with a p-value of 6.28e-4, while the dCor of surface area and ISPM was 0.445 with a p-value $<$ 2.2e-16. These results demonstrate that considering aspects of shape beyond cortical thickness and surface area is informative. ISPM appears to be more discriminative than JI in this regard.

4 Conclusion

We propose a novel method of analyzing shapes on a spherical surface using two distance metrics, the Jaccard index and our icosahedral spatial pyramid matching metric. This measure is particularly useful to capture topological shape variants such as region splitting, region location changes relative to neighbors, and missing regions. Even though we only had 27 subjects, our distance correlation analysis finds statistically significant dependencies between the shape of functional cortical regions and fluid intelligence.

It would be interesting to perform a more extensive statistical analysis involving the Glasser MMP individual parcellations for all HCP subjects when this data become publicly available. Also, our ISPM method provides an interesting distance metric on shapes and complexes of shapes that can open up avenues for other types of statistical analysis involving shape metrics, e.g., Fréchet means.

Acknowledgements Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

Table 1: Distance correlation between average cortical thickness (Thick), surface area (Area), Jaccard index (JI) or ISPM and cognitive measures.
* p-value < 0.05. ** p-value < 0.01. *** p-value < 0.001.

Test Name	Test dCor	Test P-Value	ReTest dCor	ReTest P-Value
Thick vs PVT	-0.042	1	0.008	1
Thick vs Read	-0.003	1	0.058	0.573
Thick vs Speed	0.109	0.279	0.024	0.958
Thick vs Fluid	0.053	0.836	0.107	0.314
Thick vs All	0.043	0.836	0.066	0.573
Area vs PVT	-0.012	1	0.053	0.392
Area vs Read	0.001	1	0.135	**0.029***
Area vs Speed	0.048	0.732	0.088	0.161
Area vs Fluid	0.099	0.431	0.165	**0.00967****
Area vs All	0.060	0.732	0.162	**0.00967****
JI vs PVT	0.037	0.575	0.081	0.276
JI vs Read	0.099	0.106	0.070	0.294
JI vs Speed	0.102	0.106	0.022	0.787
JI vs Fluid	0.154	**0.0156***	0.179	**6.95e-3****
JI vs All	0.182	**5.61e-3****	0.134	**0.0436***
ISPM vs PVT	0.071	0.227	0.086	0.229
ISPM vs Read	0.128	**0.039***	0.042	0.509
ISPM vs Speed	0.117	**0.0494***	0.044	0.509
ISPM vs Fluid	0.212	**3.26e-4*****	0.211	**7.37e-4*****
ISPM vs All	0.239	**7.36e-5*****	0.144	**0.0267***

References

1. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International journal of computer vision **61**(2), 139–157 (2005)
2. Duffau, H.: A two-level model of interindividual anatomo-functional variability of the brain and its implications for neurosurgery. Cortex **86**, 303–313 (2017)
3. Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., et al.: A multi-modal parcellation of human cerebral cortex. Nature **536**(7615), 171–178 (2016)
4. Grauman, K., Darrell, T.: The pyramid match kernel: Discriminative classification with sets of image features. In: Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on. vol. 2, pp. 1458–1465. IEEE (2005)
5. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 2169–2178. IEEE (2006)
6. Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: 5th IEEE EMBS International Summer School on Biomedical Imaging. IEEE (2002)

7. Lyons, R.: Distance covariance in metric spaces. The Annals of Probability **41**(5), 3284–3305 (2013)
8. Marcus, D., Harwell, J., Olsen, T., Hodge, M., Glasser, M., Prior, F., Jenkinson, M., Laumann, T., Curtiss, S., Van Essen, D.: Informatics and data mining tools and strategies for the human connectome project. Frontiers in neuroinformatics **5**, 4 (2011)
9. Mémoli, F.: Gromov–Wasserstein distances and the metric approach to object matching. Foundations of Computational Mathematics **11**(4), 417–487 (2011)
10. Oguz, I., Cates, J., Datar, M., Paniagua, B., Fletcher, T., Vachet, C., Styner, M., Whitaker, R.: Entropy-based particle correspondence for shape populations. International journal of computer assisted radiology and surgery **11**(7), 1221–1232 (2016)
11. Székely, G.J., Rizzo, M.L., Bakirov, N.K.: Measuring and testing dependence by correlation of distances. The annals of statistics **35**(6), 2769–2794 (2007)
12. Yeo, B.T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: Fast surface registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 745–753. Springer (2008)
13. Zhu, P., Awate, S.P., Gerber, S., Whitaker, R.: Fast shape-based nearest-neighbor search for brain mris using hierarchical feature matching. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 484–491. Springer (2011)