Detecting Discriminative Functional MRI Activation Patterns Using Space Filling Curves

D. Kontos[1,*], V. Megalooikonomou[1,*], N. Ghubade[1], C. Faloutsos[2]
[1] Department of Computer and Information Sciences, Temple University, Philadelphia, PA
[2] Department. of Computer Science, Carnegie Mellon University, Pittsburgh, PA

Abstract— **We propose a novel approach for detecting discriminative patterns of functional MRI (fMRI) activation that are associated with non-spatial clinical variables (e.g. disease). The main idea is to map the 3D volumes to 1D following the traversal of the Hilbert space-filling curve, which has been shown to exhibit optimality in preserving the locality of the voxels after the domain transformation. We apply statistical tests of significance on groups of points, i.e., segments of the transformed domain, to detect discriminative patterns. To discover discriminative areas, we project these patterns to the initial 3D space by following the inverse mapping of the transformation. As a case study, we analyze an fMRI data set obtained from a study that explores neuroanatomical correlates of semantic processing in Alzheimer's disease. We seek to discover activation regions that discriminate controls from patients. We evaluate the results by presenting classification experiments that utilize information extracted from these regions. The discovered areas identified through back-projection are within the medial temporal lobe being consistent with prior findings. The overall classification accuracy ranged from 81% up to 100% for certain experimental settings. The proposed approach has great potential for elucidating structure-function relationships and can be valuable to human brain mapping.**

Keywords— **activation patterns, classification, fMRI, Hilbert space-filling curve, linear mapping.**

I. Introduction

The detection of relationships between human brain structures and brain functions (i.e., human brain mapping) has been recognized as one of the main goals of the Human Brain Project [1]. Several approaches have been used in this problem domain [2]. One of the approaches used in functional brain mapping is to seek associations between brain activation patterns and tasks performed. A current obstacle in this type of analysis is the lack of methods to automatically classify such patterns (i.e., activation regions) and quantitatively measure levels of their similarity. In this paper we propose a technique for detecting and classifying functional Magnetic Resonance Imaging (fMRI) activation patterns. More specifically, we seek to discover brain activation patterns that are associated with a particular disease.

* This work was supported in part by the NSF (IIS-0083423) and the Pennsylvania Department of Health. Funding parties specifically disclaim responsibility for any analyses, interpretations and conclusions.

One of the most common approaches currently in use for analyzing fMRI activation is *statistical parametric mapping* (SPM)[3]. SPM analyzes each voxel's changes independently and ascertains their significance by means of statistical tests. This can be computationally expensive for large volumes, since each voxel is being considered independently. A problem associated with this is the multiple comparison problem, which occurs when computing a statistic for many pairwise tests. Another approach for classifying fMRI activation is to model (estimate) the underlying activation distributions utilizing parametric, non-parametric or semi-parametric techniques [4]. The main problem of these techniques is that real data are not accurately modeled using a simple mixture of Gaussian components, which is the approach that most of these approaches follow. Real data usually correspond to highly non-uniform distributions and estimating those in a multidimensional space is a difficult task to accomplish [5]. The additional dimensions represent more degrees of freedom in modeling the underlying activation distributions. This can easily mislead the estimation process to converge in local minima on the error surface of searching the best distribution hypothesis that can fit the observed activation data.

To overcome these difficulties, we propose a technique that is based on a linear mapping of the multidimensional space. Several functions have been proposed in the literature for this domain transformation. The z-ordering is based on interleaving bits from coordinates [6] and can be improved using Gray coding on the interleaved bits [7]. Here, we use the Hilbert space-filling curve [8], which has been shown to be optimal in preserving the locality of the voxels after the domain transformation, compared to the other functions proposed in the literature for this domain transformation [9]. The Hilbert space-filling curve has been extensively used for improving indexing in multidimensional databases [10] and similarity searches in spatial databases [11]. To avoid the multiple comparison problem, we apply statistical tests of significance on groups of voxels (bins) in the transformed domain. It has been shown that reproducing fMRI activations at a regional level is more robust across sites, subjects, and techniques than considering only the most significant voxels [12,13]. Here we compare patterns across activation maps taking into account their spatial extent instead of just the location of the most significant peaks. To identify the areas defined by the indicated discriminative bins in the linear domain, we project them back to the initial 3D space following the inverse mapping procedure. Finally, to evaluate the discriminatory significance of the detected

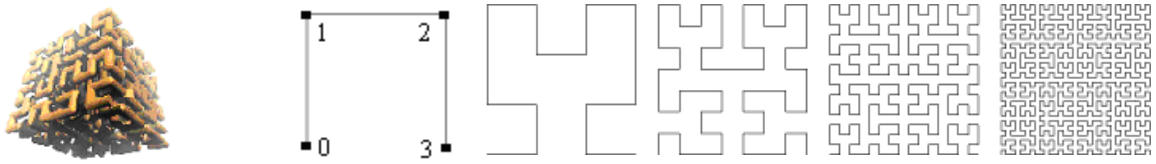

(a) (b)

Fig. 1. (a) An example of the 3D traversal of the curve and (b) the fractal like progress of the Hilbert space-filling curve in the 2D after a few iterations.

patterns, we develop a classification model based on Neural Networks that utilizes information extracted from the indicated regions to provide prediction and diagnosis. Here the goal is, given an fMRI image of a new subject, to determine the group to which it belongs, i.e., control versus various disease states such as Alzheimer's disease (AD).

II. Methodology

We seek to discover regions that provide significant discriminative information with respect to non-spatial properties, such as class membership (controls vs. patients). In the discussion that follows we present the proposed method for a two-class problem although it can be easily extended to more than two classes. In order to evaluate the method we construct features (attributes) that we use to develop and train a classification model for prediction and diagnosis. The method is applied on activation contrast maps that are the output of SPM (operating on individual subjects independently). SPM creates 3D activation maps of contrast and statistical significance values for pairs of conditions. We replace the typical "second level" of SPM analysis (group model).

Our method is based on a linear mapping of the 3D space using space filling curves. Informally, a space filling curve defines a continuous path in a multidimensional grid, visiting each point exactly once and never crossing itself. A desired property is to preserve the distances, in other words the spatial locality of the initial domain. We propose using the Hilbert space-filling curve which, as mentioned earlier, has been shown to be optimal in preserving the clustering properties of the initial multidimensional points, compared to other functions that have been used for such a domain transformation [9]. The main idea of the traversal followed by the Hilbert curve in the 3D space is shown in Fig. 1 (a). For clarification purposes, a clearer demonstration of the fractal-like traversal followed by the curve is shown in Fig. 1. (b) for the case of 2D. The procedure to derive higher orders of the curve, such as the 3D curve that we propose to use, is based on rotating and reflecting the curve at vertex 0 and vertex 3. The traversal proceeds in a recursive manner, following the same rotation and reflection pattern at each vertex of the basic curve. A detailed algorithm on how to grow the curve can be found in [14].

We apply the 3D Hilbert space-filling curve to all the 3D fMRI activation contrast maps of our dataset, mapping their voxels in the linear domain. We continue our statistical analysis in the linear domain. To reduce the dimensionality of the data and the effect of the multiple comparison problem, we employ an intermediate step of binning. For each subject we group together K consecutive points in the 1D and use as a representative attribute of the corresponding bin their mean value, V_{mean}. These representative values are then used in further statistical processing. The process of binning introduces a static partitioning of the space and smoothing of the corresponding voxel values. Furthermore, viewing K as a product of $K=k*k*k$ approximates a $k*k*k$ neighborhood in the initial multidimensional space, since the Hilbert curve preserves locality. Depending on the size of k the proposed approach has the ability to introduce a significant dimensionality reduction.

In order to detect discriminative activation patterns, we apply statistical tests of significance on the voxel bins. More specifically, for each separate bin we compare statistically the V_{mean} measurements (attributes) of all the samples belonging to one class to the corresponding V_{mean} measurements of the contrasting class. In other words, we estimate whether the V_{mean} distributions for each bin differ substantially in the separate classes. For this purpose several statistical tests can be applied. For example, this can be estimated using parametric (e.g. t-test) or non-parametric tests (e.g. Wilcoxon rank sum) [15]. Also, the Pearson correlation coefficient [15] between the class label (considered as a binary numeric value) and the attribute value V_{mean} for each sample bin can be computed (in this case an attribute is considered significant if the correlation coefficient is larger than a pre-determined threshold).

The voxel activation of a specific bin is considered to differ substantially class-wise when the statistical significance of the divergence is above a certain predefined threshold θ. Divergence is estimated using one of the previously mentioned statistical tests on the distributions of the V_{mean} measurements. This statistical significance is assigned to all the voxels of the corresponding bin. A discriminatory spatial activation pattern is constructed by back-projecting the indicated significant groups of voxels, forming areas in 3D.

III. Results

Our dataset consisted of 3D activation maps of 9 controls and 9 Alzheimer's disease patients on a category-exemplar word pair. The task consisted of an auditory presentation of word pairs (categories and possible exemplars) requiring a semantic decision (match-mismatch) [16]. Prior to the application of the proposed technique, we

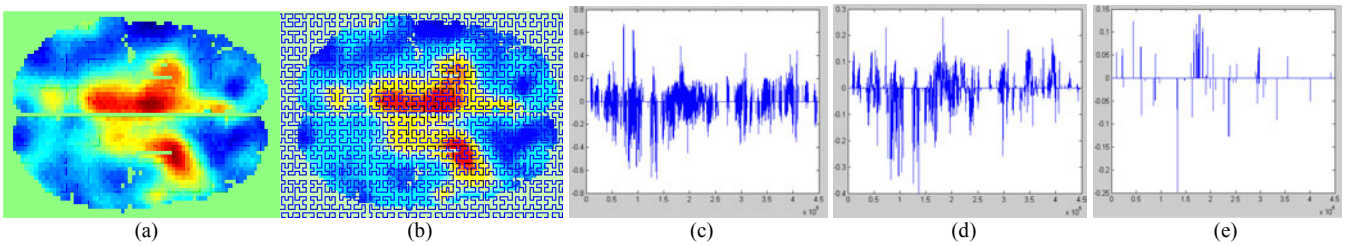

(a) (b) (c) (d) (e)

Fig. 2. Steps of the proposed technique on the AD fMRI contrast map dataset: (a) A 2D slice of the 3D fMRI brain activation map and (b) the traversal of the Hilbert space-filling curve through it. (c) The result of the linear mapping of a 3D fMRI scan and (d) the effect of binning by representing each bin with its V_{mean} measurement. (e) The discriminative voxels as indicated after applying the t-test with θ=0.05 in the transformed domain.

TABLE I

k	Classification Accuracies for **t-test**					
	***θ*=0.05**			***θ*=0.01**		
	Controls	Patients	**Overall**	Controls	Patients	**Overall**
2	0.9444	0.9889	0.9667	1	1	1
3	0.8667	1	0.9333	0.9556	1	0.9778
4	0.9111	0.9889	0.9500	0.9667	1	0.9833
5	0.8111	0.9778	0.8944	0.8667	0.9667	0.9167
6	0.7889	1	0.8944	0.7556	0.9778	0.8667
7	0.8889	0.9444	0.9167	0.8778	0.9222	0.9000
8	0.9000	0.9222	0.9111	0.8556	0.7778	0.8167
9	0.8556	0.9556	0.9056	0.8889	0.7444	0.8167
10	0.8000	0.9778	0.8889	0.7889	0.8444	0.8167

performed spatial normalization, i.e., registration of the scans to a standard template brain using the anatomical reference images. This was carried out in SPM99 [3]. Each subject's task-related activation was analyzed individually versus the subject's rest condition, resulting in individual contrast maps giving a measurement of fMRI signal change at each voxel. First, we removed the effect of the background noise by subtracting the signal value measured in representative background voxels from all the voxels of the 3D volume. Finally, we masked the data using a binary mask extracted from the T1 canonical atlas that was used as the template for the registration. Only signal within the binary mask was included in the analysis.

After these preprocessing steps we first applied the linear transformation imposed by the 3D Hilbert space-filling curve to all the sample 3D contrast maps (9 controls and 9 patients). For visualization purposes only, we illustrate the main idea for the traversal of the Hilbert space-filling curve on a 2D slice of the activation maps in Fig. 2. (a) and (b), although the 3D curve was applied on the fMRI volumes, providing a direct 3D to 1D linearization. To reduce the dimensionality of the dataset in the 1D space we experimented with a range of values for the bin size K, from K=2*2*2 (k=2) to K=10*10*10 (k=10). We used the t-test for detecting discriminative spatial patterns and experimented with significance thresholds θ=0.05 and θ=0.01. Fig. 2 shows the results of binning (d) and the discriminative bins detected by the application of the t-test (e) on the 1D transformation of a patient sample (c).

To evaluate the predictive power and the association of the indicated patterns with the disease, we proceeded with classification experiments. As a classification model we used Neural Networks. To avoid overfitting due to a small training dataset we applied one-layer perceptron networks trained by the Pocket algorithm [17]. As inputs to the classifier we used the attributes V_{mean} corresponding to the discriminative bins. The output was a binary class label indicating the class of the samples (control vs. patient). To improve the neural network training, prior to feeding to the network, the attributes were standardized to have zero mean and unit standard deviation. The leave-one-out approach was employed to evaluate classification performance on unseen data. Taking into consideration the stochastic nature of the Pocket algorithm, we repeated the process of training and testing the model in each of the leave-one-out loops 10 times and averaged the percentage of the correct predictions to obtain the reported accuracy.

The indicated regions forming discriminative spatial patterns that correspond to the best classification results for each θ selection are shown in Fig. 3 after being overlaid on the T1 anatomical template. The significance of each region is annotated using a color coding.

Table I shows the classification results obtained for control and patient samples separately as well as the overall classification accuracy for each experiment. Very good classification was achieved, reaching even 100% for specific experimental settings. These results support the argument that the regions discovered by the proposed approach are indeed associated with AD, thus providing significant discriminative information.

Fig. 3 shows that the majority of the significant regions determined by the proposed approach that could discriminate Alzheimer patients from controls were within the medial temporal lobe. This is consistent with other findings [18]. The neuropathology of early AD is relatively diffuse with atrophy in widespread cortical and subcortical areas, including the medial temporal lobes and temporal parietal and frontal cortical regions. On functional neuroimaging studies (fMRI and PET) patients with very early AD manifest as Mild Cognitive Impairment (MCI) often show compensatory activations outside of areas typically used by healthy elderly controls [18]. This is thought to represent the brain's recruitment of proximal and possibly distal neural units to maintain performance in the face of progressive pathology. Therefore, our findings of multiple distributed regions that differentiate patients and controls may be consistent with a distributed reorganization of networks subserving the semantic memory task [16].

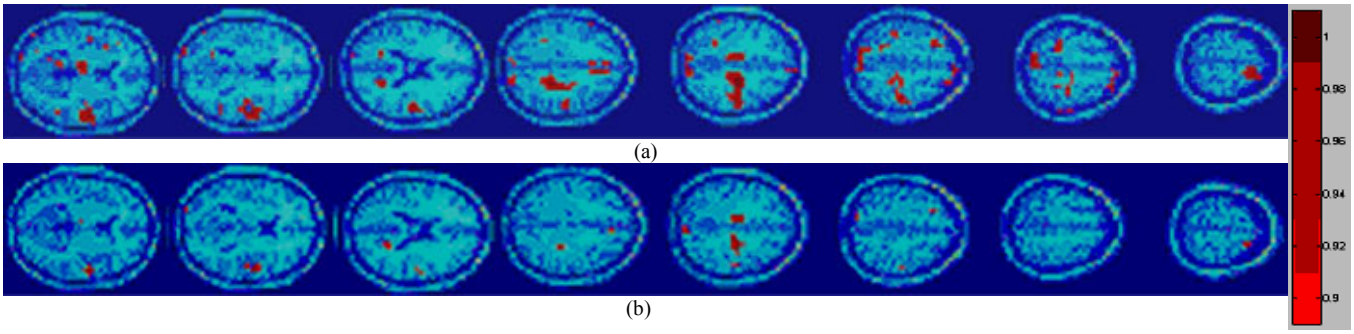

(a)

(b)

Fig. 3. Transaxial slices of the T1 canonical atlas showing the areas discovered by the proposed method when applied with significance threshold (a) $\theta = 0.05$, (b) $\theta=0.01$. The colorbar shows the significance.

IV. Conclusions

We proposed and evaluated methods for the analysis of brain activation scans potentially suitable for the effective discovery of spatial activation patterns that are discriminative among different groups of subjects. The methods are applied on activation maps that are the output of SPM (operating on individual subjects independently). We replace the typical "second level" of SPM analysis (group model). The proposed approach is based on a mapping of the 3D space to the 1D space introduced by the traversal of the Hilbert space-filling curve. We applied statistical tests of significance in the linear domain to detect discriminative activation patterns. The proposed technique considers groups of voxels (spatial sub-domains) and effectively reduces the computational cost introduced by repeated statistical tests (multiple comparison problem). It is also more robust than methods that perform voxel-wise analysis which are more prone to registration errors and variability of individual voxel values across runs, subjects, and analysis techniques. This is due to the fact that it seeks for activation patterns on a regional level by introducing the intermediate step of voxel binning (static partitioning of the space). Work in progress includes experiments aimed to combining the regions discovered when using different statistical tests. Finally, we plan to explore the ability of the proposed technique as a tool for mining spatial patterns and associations related to other diseases and disorders in the framework of the human brain structure and function.

Acknowledgment

The authors would like to thank A. Saykin for providing the fMRI data set and clinical expertise.

References

[1] S.H. Koslow, M.F. Huerta, "NeuroInformatics: An Overview of the Human Brain project," Mahway, NJ, Lawrence Erlbaum, 1997.

[2] V. Megalooikonomou, J. Ford, L. Shen, F. Makedon, F. Saykin, "Data mining in brain imaging," Statistical Methods in Medical Research, vol. 9, no. 4, pp. 359-394, 2000.

[3] K.J. Friston, A.P. Holmes, K.J. Worsley, J.P. Poline, CD. Frith, R.S.J. Frackowiak, "Statistical parametric maps in functional imaging: a general linear approach," Human Brain Mapping, 2: 189–210, 1995

[4] A. Lazarevic, D. Pokrajac, V. Megalooikonomou, Z. Obradovic, "Distinguishing Among 3-D Distributions for Brain Image Data Classification," in Proceedings of the 4th Intl. Conference on Neural Networks and Expert Systems in Medicine and Healthcare, Milos Island, Greece, pp. 389-396, 2001.

[5] V. Megalooikonomou, D. Pokrajac, A. Lazarevic, and Z. Obradovic, "Effective classification of 3-D image data using partitioning methods", in Proceedings of the SPIE 14th Annual Symposium in Electronic Imaging: Conference on Visualization and Data Analysis San Jose, CA, pp. 62-73, Jan. 2002.

[6] A. Orenstein, "Spatial query processing in an object oriented database system," In SIGMOD, Washington, D.C., pp. 326-236, 1986.

[7] C. Faloutsos, "Multiattribute hashing using gray codes," Proceedings of the ACM SIGMOD International Conference on Management of Data, pages 227—238, 1986.

[8] D. Hilbert. "Uber die stetige Abbildung einer Linie auf Flachenstuck," Math. Annln., 38:459-460, 1891.

[9] B. Moon, H.V. Jagadish, C. Faloutsos, J.H. Saltz." Analysis of the Clustering Properties of the Hilbert Space-Filling Curve," IEEE Transactions on Knowledge and Data Engineering, 13(1):124--141, Feb. 2001.

[10] C. Faloutsos, S. Roseman. "Fractals for secondary key retrieval," In Proc. of the ACM SIGACT/ SIGMOD Symposium on Principles of Database Systems, March 1989.

[11] G.R. Hjaltason, H. Samet. "Ranking in spatial databases," In Advances in Spatial Databases, Fourth International Symposium (SSD), pages 83--95, 1995.

[12] B. Casey, J. Cohen, K. O'Craven, et al. Reproducibility of fMRI results across four institutions using a spatial working memory task. NeuroImage, 8(3):249-261,October 1998.

[13] C. Tegeler, S. Strother, J. Anderson, and S. Kim. Reproducibility of BOLD-based functional MRI obtained at 4 T. Human Brain Mapping, 7(4):267{283, 1999.

[14] J. G. Griffiths, "An algorithm for displaying a class of space-filling curves," Software—Practice & Experience, v.16 n.5, p.403-411, May 1986.

[15] J.L. Devore, "Probability and Statistics for Engineering and the Sciences," 5th ed., International Thomson Publishing Company, Belmont, 2000.

[16] A.J. Saykin, L.A. Flashman, S.A. Frutiger, S.C. Johnson, A.C. Mamourian, C.H.Moritz, J.R. O'Jile, H.J. Riordan, R.B. Santulli, C.A. Smith, J.B. Weaver, "Neuroanatomic substrates of semantic memory impairment in Alzheimer's disease: Patterns of functional MRI activation," Journal of the International Neuropsychological Society, 5, 377-392,1999.

[17] S.I. Gallant, "Perceptron-Based Learning Algorithms," in IEEE Transactions on Neural Networks, vol. 1, no. 2, pp. 179-191, June 1990.

[18] A.J. Saykin, H.A. Wishart, "Mild cognitive impairment: Conceptual issues and structural and functional brain correlates," Seminars in Clinical Neuropsychiatry, vol. 8, pp. 12-30, 2003.