Detecting Low-Frequency Functional Connectivity in fMRI Using a Self-Organizing Map (SOM) Algorithm

Scott J. Peltier,¹ Thad A. Polk,² and Douglas C. Noll³

¹Department of Applied Physics, University of Michigan, Ann Arbor, Michigan

²Department of Psychology, University of Michigan, Ann Arbor, Michigan

³Department of Biomedical Engineering, University of Michigan, Ann Arbor, Michigan

Abstract: Low-frequency oscillations (<0.08 Hz) have been detected in functional MRI studies, and appear to be synchronized between functionally related areas. A current challenge is to detect these patterns without using an external reference. Self-organizing maps (SOMs) offer a way to automatically group data without requiring a user-biased reference function or region of interest. Resting state functional MRI data was classified using a self-organizing map (SOM). Functional connectivity between the left and right motor cortices was detected in five subjects, and was comparable to results from a reference-based approach. SOMs are shown to be an attractive option in detecting functional connectivity using a model-free approach. *Hum. Brain Mapp.* 20:220–226, 2003. © 2003 Wiley-Liss, Inc.

Key words: functional connectivity; BOLD; model-free analysis; pattern recognition; clustering; functional MRI

INTRODUCTION

Recent studies in functional MRI have shown slowly varying fluctuations that are temporally correlated between functionally related areas. These low-frequency oscillations (<0.08 Hz) seem to be a general property of symmetric cortices, and have been shown to exist in the motor, auditory, visual, and sensorimotor systems, among others [Biswal et al., 1995; Cordes et al., 2000; Lowe et al., 1998]. Thus, these fluctuations agree with the concept of functional connectivity: a descriptive measure of spatio-temporal correlations between spatially distinct regions of cerebral cortex [Friston et al., 1993]. Several recent studies have shown

decreased low-frequency correlations for patients in pathological states (such as multiple sclerosis [Lowe et al., 2002] or cocaine use [Li et al., 2000]). Accordingly, low-frequency functional connectivity may be important as a potential indicator of regular neuronal activity within the brain.

However, it is a challenge to detect and quantify low-frequency spatio-temporal patterns in functional imaging data. The cross-correlation method in common use is sensitive to drifts in the data [Lowe and Russell, 1999], which are still present in low-frequency filtered data. There is also the obvious question of what reference waveform to use in a correlation analysis of resting-state data, where there is no external paradigm being presented. The use of investigator-defined regions of interest (ROIs) or “seed clusters” has been the primary method used in functional connectivity studies [Biswal et al., 1995; Hampson et al., 2002; Lowe et al., 1998; Peltier and Noll, 2002], in which the pixel time courses in a particular slice are correlated with the ROI reference waveform to form functional connectivity maps. This use of “seed clusters” is not an optimal way of detecting functional connectivity, in that it is (1) user-biased, and (2) not applicable in cases where pre-supposed ROIs are not known or for which a task activating the ROI is unknown.

Contract grant sponsor: NIH; Contract grant number: NS32756.

*Correspondence to: Scott J. Peltier, Emory University, Hospital Annex, 531 Asbury Circle, Suite N305, Atlanta, GA 30322-4600. E-mail: speltier@bme.emory.edu

Received for publication 31 July 2002; Accepted 14 August 2003
DOI 10.1002/hbm.10144

Several model-free approaches to fMRI analysis have been used recently, including principal component analysis (PCA) [Weaver et al., 1994] and fuzzy clustering [Golay et al., 1998]. However, each of these approaches has drawbacks that limit their application to investigating functional connectivity. PCA finds principal components that are orthogonal, and that contribute a large amount of variance to the data, but fMRI data can violate orthogonality assumptions [Le and Hu, 1995], and in functional connectivity, the signals of interest may be weak relative to other sources of noise (e.g., physiological noise). Fuzzy clustering groups data into clusters of related time courses, based on the membership of each time course in each cluster; but recently it has been shown that fuzziness is not necessary in fMRI clustering [Fischer and Hennig, 1999a], which allows implementation of computationally faster algorithms (such as k-means clustering or self-organizing maps).

Recently, model-free analysis using self-organizing maps (SOMs) has been applied to functional MRI [Fischer and Hennig, 1999b; Ngan and Hu, 1999], and has shown promise in detecting activation patterns related to performing a cognitive task. We sought to assess whether the use of self-organizing maps (SOMs) can be extended to the detection of resting state functional connectivity. This work looks at detecting functional connectivity patterns in resting state fMRI data using self-organizing maps, and compares the results to those found using seed clusters, a standard reference-based approach. We demonstrate that the use of self-organizing maps offers an attractive alternative as a model-free analysis method for detecting functional connectivity.

METHODS

Data Acquisition

A series of fMRI experiments were performed on a 1.5 T Signa-LX scanner (GE Medical Systems, Milwaukee, WI) using a spiral pulse sequence. The sequence implemented a spectrospatial excitation pulse, and acquired 400 images. Pulse sequence parameters were TR/TE/FA/FOV of 520 msec/48.66 msec/45°/20 cm. Five 5-mm-thick axial slices were acquired in each TR, with an in-plane resolution of 3.57 × 3.57 mm.

Six subjects were studied under conditions of activation and rest. Four distinct sets of data, two task activation sets and two resting state sets, were acquired for each subject. To minimize the effect of possible attentional and/or physiological changes on the data, one set of resting data was acquired before and one set after the two sets of activation data for each subject, with minimal delay between each set. A sequential finger-tapping motor paradigm (20.8-sec fixation, 20.8-sec task, 5 repeats) was implemented for the activation data. The paradigm cues were visually presented to the subjects using IFIS (Integrated Functional Imaging System, Psychology Software Tools, Pittsburgh, PA). Resting state data was acquired while the subjects were inactive (lying still, with fixation cross being presented), and matched to the duration of the activation data (208 sec total).

The cardiac rhythm of the subjects was recorded during all runs using a pulse oximeter.

Post-Processing

Motion detection was performed on all data in order to detect gross head movement. This analysis was performed using 3-D rigid-body registration in AIR (Automated Image Registration) [Woods et al., 1998]. Following prior work [Cordes et al., 2000; Lowe et al., 1998], a cutoff of 0.4 mm was used for displacement in x , y , or z . One subject exceeded this motion threshold and was excluded from the study. The data of the remaining five subjects was then analyzed using the non motion corrected data, to avoid the introduction of artifactual spatio-temporal correlation through the motion correction process [Lowe et al., 1998].

A method of systematic noise removal was employed on all data following acquisition, as follows. Linear trends were removed from the data, to eliminate the effect of gross signal drifts, which could be due to scanner instabilities and/or gross physiological changes in the subject. Physiological noise variations in the data due to the cardiac rhythms were then removed using the regression analysis method proposed by Hu et al. [1995]. This approach removes the effects of the first and second order harmonics of the externally collected physiological waveforms. (Furthermore, the primary cardiac harmonic is removed when the low-pass filter is applied.) In addition, the time course of the highest varying voxel in the sagittal sinus was used as a regressor [Lund and Hanson, 2001], in order to remove any cardiac effects that were not in synchrony with the external cardiac measurement. The noise-corrected data was used for all subsequent analyses.

For the resting state data, the functional data was first low-pass filtered, by convolving the time courses with a rect filter with a cutoff frequency of 0.08 Hz. This cutoff avoids aliasing of the unwanted respiration and primary cardiac harmonic into the region of interest, while keeping the frequencies shown to contribute to functional connectivity [Biswal et al., 1995; Cordes et al., 2000; Lowe et al., 1998]. Each time course was then normalized, by subtracting its mean and dividing by its standard deviation, to give time courses with zero mean and unit variance. This was done to avoid convergence of the SOM algorithm based on the mean value of the time courses.

SOM Algorithm

The data were analyzed using the self-organizing map (SOM) algorithm developed by Kohonen [1995]. This algorithm produces a predetermined number of exemplar time courses that represent the probability density function of the underlying data. This process is done iteratively. First, the exemplar matrix is initialized to random noise. Then, in one iteration, every voxel time course is compared to all the exemplar time courses, and the minimum distance is calculated, using a least squares metric:

Figure 1.

A: Normalized mean square error (NMSQE) vs. iteration number for a typical subject. Results are normalized by the MSQE of the last iteration. **B:** Voxel patterns found in one supercluster at the first (**top**) and last (**bottom**) iteration, overlaid on the corresponding structural image.

$$\|x - m_c\| = \min_j (\|x - m_i\|) \quad i = 1, \dots, N \quad (1)$$

where x is the time course of the data voxel under consideration, m_i denotes the time course of exemplar i , and m_c is the time course of the closest exemplar c . The exemplars are then updated at each iteration using:

$$m_i(t+1) = m_i(t) + h_{ci}(t) * [x - m_i(t)] \quad (2)$$

where t is the current iteration number, and $h_{ci}(t)$ is a (time-dependent) neighborhood function that controls how many neighboring exemplars in addition to the closest exemplar are also updated, and to what degree. As the iterations progress, the neighborhood function shrinks the neighborhood; so while initially, the exemplar map receives global ordering, at the end, only individual nodes are updated. The above steps are repeated for all voxels until convergence is reached, and further iterations produce no change in node

assignments. The final exemplar map will have a topologically ordered feature map that represents the underlying probability density function of the data with minimal error [Kohonen, 1995].

Implementation

For this study, SOM's consisting of 100 exemplars arranged in a 10×10 2-D grid were implemented. This gives a hundred exemplar time courses, and in general seems ample enough size for classifying four to five possible fMRI cluster types (activation, cardiac/respiratory rhythms, head motion, noise, functional connectivity, among the possibilities). The SOM algorithm defined by equations (1) and (2) was employed, with the neighborhood contraction rate implemented as a shrinking Gaussian neighborhood function, dependent on the iteration number (t):

$$h_{ci}(t) = \alpha * \exp(-\|r_i - r_c\|^2 / (2 * \sigma(t)^2)) \quad (3)$$

Figure 2.

A: Resultant exemplar time course matrix after applying the SOM algorithm to a typical low-frequency filtered fMRI data set. The colored time courses correspond to one supercluster. **B:** Voxel patterns corresponding to the supercluster in A, overlaid on the corresponding anatomic image.

Figure 3.

Comparison of the significant patterns found in the motor areas of five subjects in a task activation study (**top**) with the corresponding resting state functional connectivity studies using (**middle**) seed correlation analysis and (**bottom**) self-organizing map (SOM)

analysis. The activation study is thresholded at ($r > 0.4$) for task, and the seed correlation study at ($r > 0.35$) for low-frequency reference waveform. All resting state maps were contiguity thresholded at >3 voxels, for viewing purposes.

where α is a learning rate that controls how fast the exemplars change, set to be 0.1; r_i and r_c are the coordinates (in the two-dimensional exemplar matrix) of exemplars i and c ; and $\sigma(t)$ is the FWHM of the Gaussian function, initially set at seven nodes to give global topographical ordering, but then decreasing by 5% at each iteration to switch to local ordering. At each iteration step, all time courses of interest (the time courses of all voxels that occur within the brain in each slice) were compared to the exemplar map as described above. The number of total iterations was set at 100 in order to establish convergence of the algorithm.

Analysis

For each data set, the SOM algorithm was used to generate a set of exemplars. In order to reduce the amount of resultant data, the 100 exemplars in the final exemplar matrix were then grouped into superclusters, to examine the principal groupings of the exemplars. Superclusters were formed by finding the minimum least-squares distance between each exemplar time course and those of its immediate neighbors, placing the closest time courses into the same cluster, and repeating the procedure, until the entire exemplar matrix was grouped into 24 superclusters. (There is no neighborhood function as when using the SOM algorithm, as the exemplars are already topologically ordered.) These superclusters were then examined for significant spatial patterns in the motor-related areas.

The exemplar map was examined after each iteration to verify convergence of the algorithm. The mean squared

error (MSQE) between the data time courses and the representative exemplar time courses was calculated. Normalized mean square error (NMSQE) was calculated by dividing the MSQE at each iteration by the MSQE at the final iteration. The behavior of the NMSQE with iteration number was used to gauge the rate of convergence and reduction in error as the SOM analysis progressed.

For comparison, task activation correlation maps and functional connectivity correlation maps generated using the seed cluster method were formed, as follows. First, every time course of the task activation data set was correlated with the motor task reference waveform (square wave with

TABLE I. Similarity measure for the seed cluster and SOM analysis methods*

Subject no.	Similarity measure	
	Seed cluster	SOM
1	0.5530 (0.0134)	0.6231 (0.0496)
2	0.5483 (0.0301)	0.5808 (0.0641)
3	0.5238 (0.0183)	0.5428 (0.0413)
4	0.5659 (0.0256)	0.5959 (0.0783)
5	0.5300 (0.0271)	0.5021 (0.0864)

*Values are mean (SD) of the average correlation value between all voxel timecourses in the motor-related cluster and the mean time-course of that cluster. The seed cluster connectivity maps were thresholded at ($r > 0.35$).

Figure 4.

The largest six supercluster patterns found for a typical subject. The patterns were contiguity thresholded at >3 voxels, and arranged in descending number of voxels, from **A–F**.

5.2 sec of hemodynamic offset) to form activation correlation maps. Seed clusters were then selected for each slice by identifying the block of four voxels containing the highest average correlation coefficient. This method of using seed clusters has been implemented in other functional connectivity studies [Cordes et al., 2000; Hampson et al., 2002; Lowe et al., 1998]. The corresponding seed cluster time courses in the filtered resting state data were averaged together to create a single low-frequency reference waveform. This low-frequency reference waveform was cross-correlated with all other low-pass filtered voxels to form functional connectivity correlation maps. This process was performed for all slices in all sets of resting state data.

The spatial patterns generated by the seed cluster and SOM methods were compared for significant functional connectivity patterns in the motor-related areas. The reference waveform for each method (seed cluster reference waveform, SOM node exemplar) was used to investigate the correlation strength between each reference waveform and the time courses of the voxels belonging to its associated pattern.

RESULTS

Convergence in the SOM algorithm was evidenced in the MSQE analysis of individual data sets. Figure 1A shows the MSQE vs. iteration number for a typical data set. The MSQE decreases until after ~ 60 iterations it reaches a stable level. It can be seen that the initial node assignment at iteration one results in approximately $1.75\times$ the relative error as compared to the final configuration at iteration 100. The associated spatial pattern of one of the superclusters is shown in Figure 1B, after the initial and final iteration. The normalized MSQE at iteration one is high, and the supercluster's voxels are widespread, corresponding to no discernable pattern. At the final iteration, however, it is seen that the exemplar map has self-organized, so that the supercluster pattern is seen to be localized in the motor-related areas.

The SOM algorithm was effective in classifying the resting state data. Figure 2 displays the results for a typical data set. The exemplar matrix in Figure 2A represents the data with 100 exemplar time courses. It can be seen that the exemplar time courses are topologically ordered across the matrix, with sim-

ilar time courses being grouped together. From this exemplar matrix, superclusters of exemplar time courses are examined for functional connectivity patterns. Figure 2B shows the pattern identified by the colored supercluster in Figure 2A. The pattern clearly has spatial groupings in the left and right motor cortex, as well as the supplementary motor area.

Figure 3 shows typical results of the SOM algorithm in detecting functional connectivity in the motor cortex across all subjects, compared to the results found by the seed correlation method. It is seen that the SOM algorithm identifies similar connectivity patterns as compared to the correlation method, but with no use of external reference, with patterns being found in the motor-related areas in all subjects.

To further quantify the motor-related clusters found by the seed cluster and SOM analysis methods, a similarity measure was formed. The average time course of each cluster was correlated with the time course of each voxel contained within the cluster. The average correlation coefficient over all voxels within the cluster was then used as the similarity measure. The results, shown in Table I, show that the seed cluster method and the SOM analysis are both strongly correlated with their average waveforms. It is also seen that the amount of correlation is not significantly different between the methods, indicating the SOM analysis finds clusters that are as self-correlated as the seed cluster method. The standard deviation is higher in the SOM case, which may be due to the inclusion of a wider range of voxel time courses, since the SOM algorithm to iteratively cluster all the data, as opposed to the seed cluster method. The results for the seed cluster method depend on the correlation threshold used, which also changes the size of the resultant connectivity patterns; but no significant differences in the similarity measure were found when the pattern size within the motor cortex was matched between the two methods.

Figure 4 displays the largest superclusters for the SOM exemplar matrix from Figure 2. The previously identified motor functional connectivity pattern is seen to be supercluster B. Supercluster A seems to be vascular in origin, with the cluster located around the lower sagittal sinus. Supercluster C may be head motion, as the cluster is located along the edge of the head. Clusters D, E, and F are possible connectivity clusters, with separate clusters found in spatially distinct regions.

A

B

C

Figure 5.

Total squared error between a typical data set and the exemplar matrix found by the SOM algorithm at each iteration, for different sets of parameters (number of exemplar nodes, learning rate [LR], or neighborhood contraction rate [NCR]). The default param-

eters (marked by an asterisk), which were implemented on the data, were 100 exemplar nodes, a learning rate of 0.1, and a neighborhood contraction rate of 0.05.

DISCUSSION

A self-organizing map (SOM) algorithm was applied to resting-state functional MRI data. The algorithm was able to successfully cluster the underlying data and capture significant spatio-temporal features. In particular, low-frequency functional connectivity associated with the motor cortex was detected, with no use of external reference or bias. By using a model-free approach to analyze fMRI data based on the signal alone, SOM algorithms offer an attractive analysis alternative for detecting resting state low-frequency functional connectivity without user bias.

The SOM algorithm operates by using the data to form exemplars that represent the data set, which offers advantages over the commonly used seed cluster method of examining functional connectivity. First, there is no need to collect an activation data set, so scan time is reduced. Secondly, the entire data set is classified in the SOM analysis, so that all voxels are grouped to similar voxels. This is in contrast to the seed cluster approach, which yields a correlation map that is only relative to the chosen seed cluster. Thus, the SOM analysis permits investigation of all interesting clusters in the data, as shown in Figure 4, so that even when pre-supposed ROIs are not known, the SOM analysis may be implemented.

Further optimization of the implemented SOM analysis can be performed. The convergence results (Fig. 1) show that a running analysis of MSQE can be implemented to automatically stop the algorithm when no further appreciable changes in MSQE or time course classification occur (i.e., all data time courses continue to map to the same winning exemplar time course). Also, by incorporating anatomically defined landmarks, automatic selection of superclusters can be employed instead of user selection, further removing user bias and/or error.

It may also be possible to further optimize the SOM parameters. Figure 5 shows the effect of changing the learning rate, number of nodes, or neighborhood contraction rate. It is seen that the parameters implemented in this study (marked by an asterisk in Fig. 5) generally attain the lowest total squared error, given practical considerations. In Figure 5A, the use of 100 nodes results in lower error than that found using a smaller matrix, while not sacrificing the extra processing time that using a larger matrix would entail. Larger matrix sizes would also entail a greater need for superclustering. The learning rate of 0.1 has a slightly higher final error than when using a learning rate of 0.05 (Fig. 5B). However, it can be seen that the learning rate of 0.025 results in an even higher final error, probably because the learning rate is so slow that even when the neighborhood has contracted to give local ordering, the exemplars have not been globally optimized. Thus, using a slightly higher learning rate could be beneficial, as different underlying data may require a slightly higher learning rate to order the data to the same degree. Finally, the different values of the neighborhood contraction rate result in almost the same error (Fig. 5C), so that minimal error can be sacrificed for a reduction in convergence time. Since there is likely an interaction between the learning rate and neighborhood contraction rate, a multi-dimensional parameter optimization is probably warranted.

REFERENCES

- Biswal B, Yetkin F, Haughton V, Hyde J (1995): Functional connectivity in the motor cortex of resting human brain using echoplanar MRI. *Magn Reson Med* 34:537–541.
- Cordes D, Haughton V, Arfanakis K, Wendt G, Turski P, Moritz C, Quigley M, Meyerand, ME (2000): Mapping functionally related regions of brain with functional connectivity MR imaging. *Am J Neuroradiol* 21:1636–1644.
- Fischer H, Hennig J (1999a): Is fuzziness useful in fMRI clustering? In: *Proceedings of the International Society of Magnetic Resonance in Medicine, 7th Annual Meeting*. Philadelphia, PA. p 1720.
- Fischer H, Hennig J (1999b): Neural-network based analysis of MR time series. *Magn Reson Med* 41:124–131.
- Friston KJ, Frith C, Liddle P, Frickowiak R (1993): Functional connectivity: the principal components analysis of large (PET) data sets. *J Cereb Blood Flow Metab* 13:5–14.
- Golay X, Kollias S, Stoll G, Meier D, Valavanis A, Boesiger P (1998): A new correlation-based fuzzy logic clustering algorithm for fMRI. *Magn Reson Med* 40:249–260.
- Hampson M, Peterson BS, Skudlarski P, Gatenby JC, Gore JC (2002): Detection of functional connectivity using temporal correlations in MR images. *Hum Brain Mapp* 15:247–262.
- Hu X, Le TH, Parrish T, Erhard P (1995): Retrospective estimation and correction of physiological fluctuation in functional MRI. *Magn Reson Med* 34:201–212.
- Kohonen T. 1995. *Self-organizing maps*. New York: Springer Verlag. 362 p.
- Le TH, Hu X (1995): Potential pitfalls of principal component analysis in fMRI. In: *Proceedings of the Magnetic Resonance in Medicine, 3rd Annual Meeting*. Nice, France. p 820.
- Li SJ, Biswal B, Li Z, Risinger R, Rainey C, Cho J, Salmeron B, Stein E (2000): Cocaine administration decreases functional connectivity in human primary visual and motor cortex as detected by functional MRI. *Magn Reson Med* 43:45–51.
- Lowe MJ, Russell DP (1999): Treatment of baseline drifts in fMRI time series analysis. *J Comp Asst Tomogr* 23:463–473.
- Lowe MJ, Mock B, Sorenson, JA (1998): Functional connectivity in single and multislice echoplanar imaging using resting state fluctuations. *Neuroimage* 7:119–132.
- Lowe MJ, Phillips MD, Mattson D, Dziedzic M, Matthews, VP (2002): Multiple sclerosis: low-frequency temporal blood oxygen level-dependent fluctuations indicate reduced functional connectivity-initial results. *Radiol* 224:184–192.
- Lund TE, Hanson LG (2001): Physiological noise reduction in fMRI using vessel time-series as covariates in a general linear model. In: *Proceedings of the International Society of Magnetic Resonance in Medicine, 9th Annual Meeting*. Glasgow, UK. p 22.
- Ngan SC, Hu X (1999): Analysis of functional magnetic resonance imaging data using self-organizing mapping with spatial connectivity. *Magn Reson Med* 41:939–946.
- Peltier SJ, Noll DC (2002): T_2^* dependence of low frequency functional connectivity. *Neuroimage* 16:985–992.
- Weaver JB, Saykin AJ, Burr RB, Riordan H, Maerlender A (1994): Principal components analysis of functional MRI of memory. In: *Proceedings of the Society of Magnetic Resonance, 2nd Annual Meeting*. San Francisco, CA. p 808.
- Woods R, Grafton S, Holmes C, Cherry S, Mazziotta J (1998) Automated image registration: I. General methods and intrasubject, intramodality validation. *J Comp Asst Tomogr* 22:141–154.