Functional Connectivity Mapping Using the Ferromagnetic Potts Spin Model

Larissa Stanberry,^{1*} Alejandro Murua,² and Dietmar Cordes³

¹Department of Statistics, University of Washington, Seattle, Washington

²Department of Mathematics and Statistics, University of Montreal, Montreal, Quebec, Canada

³Department of Radiology, University of Colorado Health Sciences Center, Denver, Colorado

Abstract: An unsupervised stochastic clustering method based on the ferromagnetic Potts spin model is introduced as a powerful tool to determine functionally connected regions. The method provides an intuitively simple approach to clustering and makes no assumptions of the number of clusters in the data or their underlying distribution. The performance of the method and its dependence on the intrinsic parameters (size of the neighborhood, form of the interaction term, etc.) is investigated on the simulated data and real fMRI data acquired during a conventional periodic finger tapping task. The merits of incorporating Euclidean information into the connectivity analysis are discussed. The ability of the Potts model clustering to uncover the hidden structure in the complex data is demonstrated through its application to the resting-state data to determine functional connectivity networks of the anterior and posterior cingulate cortices for the group of nine healthy male subjects. *Hum Brain Mapp* 29:422–440, 2008. © 2007 Wiley-Liss, Inc.

Key words: Potts model; functional connectivity; resting-state; fMRI; clustering

INTRODUCTION

Functionally connected regions of the brain exhibit synchronous blood flow fluctuations, which induce coherent low-frequency (<0.12 Hz) signal intensity changes in fMRI data. These coherent oscillations in different brain regions imply the existence of neural connections between them [Biswal et al., 1995; Cordes et al., 2000; Friston, 1995; Lowe et al., 1998; Stein et al., 2000; Xiong et al., 1999]. A variety of model-independent methods that aim to determine functional networks specific to a given task have been introduced in the literature and include independent com-

ponent analysis (ICA) [McKeown and Sejnowski, 1998], cluster analysis [Baumgartner et al., 1998; Cordes et al., 2002; Filzmoser et al., 1999; Goutte et al., 1999; Stanberry et al., 2003], and principal component analysis [Friston et al., 1993]. Most of the clustering methods (e.g. *K*-means, hierarchical clustering) require heavy assumptions such as a known distribution of data in clusters, or the expected number of clusters, while ICA produces spatially independent maps by optimizing a non-Gaussianity criterion (e.g. negentropy) that cannot be easily translated in terms of a functional connectivity measure.

Until now both model-dependent and model-independent methods have been largely used to determine periodic task-specific activations and functional connections. However, establishing the existence of functional connections in the resting-state brain or during continuous (steady-state) stimuli with any commonly used method has been a rather difficult task. When dealing with such complex fMRI data, researchers often resort to a seed voxel analysis [Biswal et al., 1995], where a functional network of the region of interest is defined simply by thresholding the correlation map

*Correspondence to: Larissa Stanberry, Department of Statistics, University of Washington, Box 354322, Seattle, WA 98195-4322. E-mail: lstan@u.washington.edu

Received for publication 11 January 2006; Revised 25 January 2007; Accepted 26 February 2007

DOI: 10.1002/hbm.20397

Published online 11 May 2007 in Wiley InterScience (www.interscience.wiley.com).

associated with the mean time course of the pre-selected seed voxels in that region. This method has been used to study functional connectivity of various brain regions during resting-state [Biswal et al., 1995; Cordes et al., 2000; Lowe et al., 1998, 2000; Stein et al., 2000] and continuous tasks [Hampson et al., 2002, 2004; Lowe et al., 2002]. Greicius et al. [2003] followed a similar framework in their study of functional connectivity of the cingulate cortex.

The major disadvantage of the seed voxel approach is that the resulting connectivity maps are conditionally dependent on the choice of the seed region. By the definition of the method, the brain region of interest is represented by a small number of carefully chosen voxels whose mean time courses serve as a reference function, and the resulting connectivity map is interpreted as being characteristic of the entire region of interest. However, to claim that the connectivity map obtained for a seed region reflects functional connections of the brain region from which the seed voxels were chosen, one needs to demonstrate that the mean time course of the seed region is indeed representative of the brain region of interest. In our studies [e.g. Stanberry et al., 2006], we observed that it is often the case that the structure of interest is comprised of several clusters, indicating inhomogeneity of the temporal responses in the selected region. Therefore, seed-voxel connectivity maps can potentially be inaccurate as they could omit the connections from parts of the structure underrepresented in the seed region. Also, the seed voxel approach does not allow for the automatic identification of the underlying functional dependencies, and, to get a full picture of the induced connections, one needs to sift through the entire set of brain regions, which is rather tedious and time consuming.

In this paper we adapt a stochastic clustering method based on the ferromagnetic Potts model to determine functional connectivity networks. The method was originally developed by Blatt et al. [1996] and its physical aspects and performance were rigorously studied in successive publications [Blatt et al., 1997; Domany, 1999; Wiseman et al., 1998]. The Potts model is native of statistical physics [Wu, 1982] and models an inhomogeneous ferromagnetic system where each data point is viewed as a marked vertex on the graph. The mark is a cluster label, or spin, associated with the vertex. The configuration of the system is defined by the interactions between the vertices and a parameter, T , called temperature. At low temperatures, all labels are identical (spins are aligned), which is equivalent to the presence of a single cluster. As temperature rises the single cluster starts to split as the interactions between weakly coupled voxels vanish. In the statistical mechanics literature, the splitting corresponds to changes in entropy, while from a clustering viewpoint it is equivalent to the partition of data into distinct groups containing objects with similar characteristics. As the temperature continues to climb, the remaining clusters split into a multitude of smaller clusters. This implicit behavior of the Potts model can be explained through its connection with the random clusters model [Sokal, 1996; Murua et al., On Potts model

clustering, kernel K -means and density estimation, J Comput Graph Stat, under review].

One of the main advantages of the Potts model clustering is that the number of clusters and cluster memberships are determined by consensus. That is, several possible and likely clustering structures derived from the data are combined to form cluster groups using the associated consensus graph (the definition is given later); the number of clusters in the data is equal to the number of connected components of the consensus graph.

The superparamagnetic clustering method, as it was dubbed in the original paper, has been successfully applied to different research problems [Agrawal and Domany, 2003; Domany, 2003; Domany et al., 1999; Einav et al., 2005; Getz et al., 2000a; Kullmann et al., 2000; Ott et al., 2004, 2005; Quiroga et al., 2004] and various extensions of the method were proposed, including the iterative coupled two-way clustering [Getz et al., 2000b; Getz and Domany, 2003], the modified Hamiltonian and an improved version of the cluster update algorithm for the image segmentation problem [von Ferber and Woergoetter, 2000], a sequential extension for inhomogeneous clusters [Ott et al., 2004, 2005], and a spin-glass Hamiltonian for complex networks data [Reichardt and Bornholdt, 2004]. Murua et al. (On Potts model clustering, kernel K -means and density estimation, J Comput Graph Stat, under review) related the Potts model based clustering to kernel-based clustering methods and density estimation and introduced yet another penalized Hamiltonian to favor clusters of equal sizes.

The remainder of the paper is structured as follows. In the next section we briefly describe the Potts model clustering method. The performance of the method is studied on the simulated data and data acquired from a conventional on-off finger tapping task. The method is applied to the resting-state data to detect functional networks of the posterior and anterior cingulate cortices. The discussion and conclusions are presented in the last two sections.

METHODS

Potts Model Clustering

Consider a data set containing N voxels $v_1, \dots, v_N$. We define the functional distance, f_{ij} , between two voxels v_i and v_j as

$$f_{ij} = 1 - \text{corr}(v_i, v_j) \quad (1)$$

where the correlation coefficient, $\text{corr}(v_i, v_j)$, measures the similarity of the acquired signals. Define the interaction, k_{ij} , between voxels v_i and v_j as

$$k_{ij} = \frac{1}{\hat{K}} \exp\left(-\frac{f_{ij}^2}{2a^2}\right) \quad (2)$$

where $\hat{K}$ is the average number of neighbors per voxel and a is the average nearest neighbor functional distance. Note

that the interaction term (2) is in the form of an isotropic exponential kernel. Next, let $S = \{s_1, \dots, s_N\}$ denote the labels or, equivalently, the cluster assignments of the data points. The probability density of the label assignments at a given temperature T , a purely abstract parameter, is given by the Boltzmann distribution

$$P(S; T) = \frac{1}{Z} \exp\left(-\frac{H(S)}{T}\right) \quad (3)$$

where Z is the normalizing constant and the Hamiltonian, $H(S)$, defines the energy of the system

$$H(S) = \sum_{(i,j)} k_{ij} (1 - \delta_{s_i s_j}) \quad (4)$$

The summation in (4) is over all *neighboring* sites v_i and v_j , and $\delta_{s_i s_j}$ is the Kronecker's delta, e.g. $\delta_{s_i s_j} = 1$ if $s_i = s_j$ and $\delta_{s_i s_j} = 0$ otherwise. The Hamiltonian is defined by the between-cluster interactions, whereas pairs of voxels belonging to the same cluster ($\delta_{s_i s_j} = 1$) do not contribute to it. Potts model clustering optimizes $P(S; T)$, which can be seen as a clustering figure-of-merit that penalizes the assignment of different labels to neighboring voxels.

Graph Structure

The graph, whose vertices are voxels of the fMRI image and whose edges correspond to the 3D Euclidean spatial neighborhood system, defines a natural graph structure for such data. For large data sets, in the absence of a natural graph structure, or when a different similarity measure is more appropriate, one may define a K -nearest neighbor graph where an object A is said to be a neighbor of an object B , if and only if A is among the K -nearest neighbors of B .

In context of the functional connectivity analysis, the graph structure with edges determined by the functional rather than Euclidean distances is more meaningful as it allows for links (i.e. edges) between disjoint brain regions. Pairwise interactions (2) can be thought of as weights thus transforming data into a weighted graph. Taking into consideration the size of a typical fMRI data set, it is computationally advantageous and more rational to construct a K -nearest neighbor graph rather than a complete one. The construction of the K -nearest neighbor graph for N data points requires $O(KN \log N)$ operations [Vaidya, 1989], while Potts model clustering needs to track NK interacting pairs of voxels. Clearly, large values of K will increase the computing time, however, small values of K might affect the algorithm ability to accurately recover the underlying structure. Later in the paper we explore the performance of the method as a function of the neighborhood size and give recommendations for the choice of the constant.

Distribution of Label Assignments

To find the likely label assignment S under the Boltzmann distribution at a given temperature T we use the Swendsen–Wang [Swendsen and Wang, 1987] algorithm. This is an efficient Markov Chain Monte Carlo technique specifically developed for the Potts model. We define $p_{ij} = 1 - \exp(-k_{ij}/T)$ and rewrite the Potts model density (3) as

$$P(S) = Z^{-1} \exp\left(-\frac{1}{T} \sum_{(i,j)} k_{ij} (1 - \delta_{s_i s_j})\right) = Z^{-1} \prod_{(i,j)} ((1 - p_{ij}) + p_{ij} \delta_{s_i s_j}) \quad (5)$$

Applying the identity $x + y = \sum_{b=0}^1 x \delta_{b,0} + y \delta_{b,1}$ to the last formula [Sokal, 1996] we can express the normalization constant Z as

$$Z = \sum_{\{S\}} \sum_{\{b_{ij}\}} \prod_{(i,j)} ((1 - p_{ij})(1 - b_{ij}) + p_{ij} b_{ij} \delta_{s_i s_j}) \quad (6)$$

where a Bernoulli variable b_{ij} defines the bond between voxels v_i and v_j , and may take the value of one whenever voxels v_i and v_j are neighbors and belong to the same cluster, and zero otherwise. The collection $B = \{b_{ij}\}$ defines the bond configuration of the system. Completely connected graphs formed by the unity bonds are subgraphs of the original data graph and referred to as connected components of the original graph.

The joint density of the label assignments and bond configuration, known as the Fortuin–Kasteleyn–Swendsen–Wang model, is

$$p(S, B) = Z^{-1} \prod_{(i,j)} ((1 - p_{ij})(1 - b_{ij}) + p_{ij} b_{ij} \delta_{s_i s_j}) \quad (7)$$

It can be shown [Sokal, 1996] that the marginal over all possible bond assignments is exactly the Potts model, and the marginal over all possible cluster assignments is the *random-cluster* model.

The samples from the Potts model are generated with a Gibbs sampler [Swendsen and Wang, 1987] making use of the conditional distributions for the model in (7). It starts with randomly chosen label assignments (e.g. assigning all voxels to the same cluster) and proceeds by alternating between the two steps:

- Given a label configuration, S , independently for each bond, set $b_{ij} = 1$ with probability $p_{ij} = 1 - \exp(-k_{ij}/T)$, if voxels v_i and v_j are neighbors and have the same label. Otherwise, set b_{ij} to zero.
- Given a bond configuration, B , assign the same label to all voxels in the same connected component. The labels are chosen independently of the other components and uniformly at random among all admissible labels.

Clustering by Consensus

At any given temperature T , M simulated samples are collected giving rise to M possible data partitions $\{S_1, \dots, S_M\}$. Under the Potts model, the probability of two given voxels having the same label assignment is estimated by

$$\hat{\delta}_{ij} = \frac{1}{M} \sum_{k=1}^M \delta_{ij}^{S_k} \quad (8)$$

where $\delta_{ij}^{S_k}$ takes the value of 1 if two voxels have the same label assignment under the partition S_k . The final cluster assignment is determined by consensus, i.e., two voxels are assigned to the same cluster if the estimated probability exceeds a specified threshold, e.g. 0.5. The probabilities define the edges of a new graph called the consensus graph. In this graph the edge between two voxels exists, if and only if $\hat{\delta}_{ij}^{S_k}$ exceeds the prespecified threshold. The connected components of the consensus graph are exactly the distinct clusters formed by consensus. By construction, clusters are specific to the given temperature T .

Identifying the Stable Phase

In general, according to the Potts model we would expect to observe three distinct phases of the clustering structure depending on the temperature value. During the first phase corresponding to low temperatures, all points form a single cluster. When the system is in the second phase, clusters of strongly interacting voxels are still present, but weak interactions have already been dissolved by the rising temperature. During the third phase corresponding to high temperatures, clusters of strongly interacting voxels disintegrate and the system is comprised of multiple small size clusters. Obviously, the phase of interest with regard to clustering is the second phase.

To properly determine the range of temperatures defining the second phase of the system, Blatt et al. [1997] suggested monitoring the changes in the system “magnetization” as a function of temperature, which in statistical terms corresponds to the variance of the largest cluster size. However, it is not always the case that the largest cluster must split before others do. Therefore, it is more meaningful to monitor the size changes of two to three largest clusters. In this paper we study changes of the standard deviations trajectories of the two to four largest clusters to determine the critical temperatures where the phase transitions occur.

Comparing Partitions

To monitor the stability of the clustering structure at any given phase, one could evaluate the agreement between cluster assignments obtained by the model and the true clustering structure (e.g. known in simulation studies, or from previous investigations). When the ground truth is

unknown, which is usually the case, one could compare partitions at neighboring temperatures using some similarity statistic. Relating to the three different phases of the clustering structure one would intuitively expect the partitions to be stable during the second phase and differ substantially when the phase transitions take place.

In this paper we used the Hubert and Arabie’s [1985] adjusted Rand index as a statistical measure of correspondence between two partitions, which is based on the cluster assignments of the pairs of data points.

More precisely, consider two clustering partitions S and S' of the data set containing N points. Define N_{ij} to be the number of points that belong to the i -th cluster of S and to the j -th cluster of S' ; N_{ig} (N_{gj}) will denote the number of points in i -th cluster of S (j -th cluster of S'). Define A to be the number of pairs placed in the same cluster in both clusterings, B to be the number of pairs placed in the same cluster in S , but in different clusters in S' , C to be the number of pairs placed in the same cluster in S' , but in different clusters in S , and D to be the number of pairs placed in different clusters in both clusterings. Clearly, A and D can be interpreted as agreements between two partitions, whereas B and C reflect disagreements between them. The original Rand index [Rand, 1971] is a nonnegative similarity measure, which attains 1 when the two partitions are identical,

$$\text{Rand. Index} = (A + D) / \binom{N}{2} = \left(\binom{N}{2} + 2 \sum_{i,j} \binom{N_{ij}}{2} - \left(\sum_i \binom{N_{ig}}{2} + \sum_j \binom{N_{gj}}{2} \right) \right) / \binom{N}{2} \quad (9)$$

The Rand index has a straightforward interpretation as the probability of agreement between the two partitions. The index, however, is not corrected for randomness as it does not take on a constant value for two completely random partitions. Given the number of clusters and objects in each of the two partitions S and S' , the simplest null model for randomness assumes that the $I \times J$ contingency table is based on the generalized hypergeometric distribution. Under the null hypothesis, the expected number of object pairs placed in the same cluster under two partitions is [Hubert and Arabie, 1985]

$$E \left(\sum_{i,j} \binom{N_{ij}}{2} \right) = \sum_i \binom{N_{ig}}{2} \sum_j \binom{N_{gj}}{2} / \binom{N}{2} \quad (10)$$

From here, the expected value of the Rand index is easy to compute. Correcting the index (9) for randomness

$$\frac{\text{Index} - E(\text{Index})}{\max(\text{Index}) - E(\text{Index})} \quad (11)$$

we obtain the expression for the adjusted Rand index

Figure 1.

Left: Two symmetric anatomical regions selected for the simulation experiment; right: artificial signal profiles obtained by convolving box-car designs with 60- (top) and 20-s (bottom) periods with the hemodynamic response function. The units on the horizontal axes are in timeframes of TR 400 ms.

$$\frac{\sum_{ij} \binom{N_{ij}}{2} - \left(\sum_i \binom{N_{ik}}{2} \sum_j \binom{N_{jl}}{2} \right)}{\binom{N}{2}} \quad (12)$$

$$\frac{\frac{1}{2} \left(\sum_i \binom{N_{ik}}{2} + \sum_j \binom{N_{jl}}{2} \right) - \left(\sum_i \binom{N_{ik}}{2} \sum_j \binom{N_{jl}}{2} \right)}{\binom{N}{2}}$$

It is easy to see that the adjusted index is bounded by 1 and has expectation 0 under the null hypothesis of randomness. It achieves its maximum when the two partitions are identical and the higher the value of the index the more alike the two partitions are.

Milligan and Cooper [1986] have conducted a comprehensive study investigating the effect of the density factor, dimensionality, and sample size on the performance of five different indices used to recover the true cluster structure. They found that the adjusted Rand index by Hubert and Arabie [1985] had superior performance in the null case, when no cluster structure exists in the data. Specifically, the mean value of the index in the null case was essentially zero, correctly indicating absence of the underlying grouping structure in the data. Different density conditions and unequal cluster sizes have shown no effect on the mean index value and only a mild influence on its variability. Other measures including the Rand index [Rand, 1971], Jaccard index [Downton and Brennan, 1980], Fowlkes and Mallows index [Fowlkes and Mallows, 1983], and the Morey and Agresti [Morey and Agresti, 1984] adjusted Rand index had positive mean values in the null case and much larger variability when compared with the Hubert and Arabie's measure. In case of the nonoverlapping distinct groups in the data, the adjusted Rand index correctly indicated near perfect recovery of the clustering structure.

EXPERIMENTS

All structural and functional MR images used in this research were collected in accordance with institutional

regulations (IRB approval) on a commercial 1.5T MR scanner (General Electric, Waukesha, WI) equipped with echo-speed gradients and a standard birdcage head coil.

Simulated Data

The simulated data set was composed of two circular clusters, each containing 37 voxels (Fig. 1, left), with distinct signal profiles resembling an on-off design with periods 20 and 60 s, respectively. The artificial signal templates for two clusters were obtained by convolving a box-car function with a canonical hemodynamic response function as implemented in SPM2 (Wellcome Department of Imaging Neuroscience, London, UK) (Fig. 1, right panels). The noise component was estimated from the resting state fMRI data acquired with parameters FOV 24 cm × 24 cm, BW ±62.5 KHz, TR 400 ms, Flip 50°, slice thickness 7 mm, 64 × 64 resolution, 750 time-points. The data were checked for subject motion using AFNI (Robert Cox, NIH) and band-pass filtered to retain frequencies in the 0.02–0.12 Hz range. As a result, low-frequency noise, cardiac and respiratory artifacts were effectively eliminated from the data acquired with a high sampling rate. Following the approach of Beckmann and Smith [2004] we forgo the autocorrelation structure and assume the background noise to be normally distributed. The preprocessed resting-state data were used to estimate the noise parameters voxel-wise in the two symmetric regions chosen in the motor cortex (Fig. 1, left). Simulated Gaussian noise was then added to the artificial signal to complete the simulations. The signal and noise content was varied to produce data sets with a signal-to-noise ratio (SNR) ranging from about 0.06 to 1.

The Potts model clustering algorithm was applied to the simulated data sets. For every temperature, 1,000 sweeps of the Swendsen–Wang algorithm were run. The initial 500 sweeps were discarded from further analysis.

fMRI: Motor Task

The motor paradigm consisted of 30 identical cycles of finger tapping and rest with a period of 10 s (5 s on, 5 s off). Scanning was performed with the following EPI parameters: four slices, FOV 24 cm × 24 cm, BW ±62.5 KHz, TR 400 ms, Flip 50°, slice thickness 7 mm, gap 2 mm, 64 × 64 resolution, 750 time-points.

fMRI: Resting State

Resting-state data were collected on nine 30–45-year-old healthy men. Subjects were instructed to lay still with their eyes closed and refrain from any cognitive activity. Scanning was performed with the following acquisition parameters: 20 axial slices, FOV 24 cm × 24 cm, BW ±62.5 kHz, TR 2,000 ms, Flip 82°, slice thickness 6 mm, gap 1 mm, resolution 64 × 64, 186 time-points. The cardiac and respiratory rates were digitally acquired with a sampling fre-

quency of 100 Hz using a pulse-oximeter and a respiratory belt, respectively.

fMRI: Image Preprocessing

In each data set the motion was assessed using a 3-D registration algorithm in AFNI (Robert Cox, NIH). Voxels with an SNR value in the first decile were discarded to eliminate voxels outside the brain as well as noisy time courses associated with pulsating arteries and some portion of the cerebral spinal fluid. For each individual, the time series were normalized to have a mean of zero and standard deviation of 1. Since functionally connected regions of the brain are characterized by synchronous signal changes in the low-frequency range (0.02–0.12 Hz), the cross-correlation coefficients specific to this frequency

band were used in the analysis. To further reduce the size of the data set only voxels having low-frequency specific correlations of at least 0.3 with five or more voxels were retained. For the resting-state data potential aliasing of the cardiac and respiratory rates was evaluated from digitally sampled physiological data.

The Potts model clustering algorithm was applied to the reduced data sets. The size of the neighborhood K was set to 10. For every temperature, 1,000 sweeps of the Swendsen–Wang algorithm were run and the last 500 sweeps were retained for further analysis.

Group Functional Connectivity Maps

To obtain group connectivity maps for the resting state data associated with anterior and posterior cingulate corti-

Figure 2.

Summary statistics as a function of temperature T (on the log scale) for SNR = 0.2 (left column), 0.25 (middle column), and 0.65 (right column). The peaks of the standard deviation trajectories (first and second rows) correspond to critical temperatures associated with significant changes in the clustering structure. As temperature grows the number of clusters steadily

increases (third row). The adjusted Rand index (bottom row) shows the existence of the stable phase during which the clustering structure is perfectly uncovered for SNR > 0.2. For the lowest SNR only the cluster containing low-frequency signals is identified, while the high-frequency signals form multiple independent groups.

ces, subject-specific cluster maps were coregistered to the standardized brain using the FSL software (Image Analysis Group, FMRIB, Oxford, UK). The spatial extent of the posterior cingulate and anterior cingulate cortices was defined by warping them from the Automated Anatomical Labeling system [Tzourio-Mazoyer et al., 2002] to the individual brain using FSL Flirt. Cluster memberships of all the voxels belonging to the anterior cingulate cortex were identified for each subject. Subject-specific indicator maps were constructed to have a value of 1 at voxels that belong to any of the clusters comprising the anterior cingulate cortex, and zero, otherwise. Group connectivity maps for the anterior cingulate cortex were constructed by averaging individual indicator maps across subjects on a voxel by voxel basis. Group maps reflecting functional connections of the posterior cingulate cortex were constructed in the same fashion. The group maps were smoothed with a 5-mm smoothing window using the SPM2 package (Wellcome Department of Cognitive Neurology). Functional connectivity maps were overlaid on the standardized brain using MRICro software (Chris Rorden, University of Nottingham).

RESULTS

Simulated Data

Algorithm illustration

To investigate the changes in the clustering structure as a function of temperature and signal-to-noise ratio characteristics we simulated three data sets as described earlier with SNR = 0.2, 0.25, 0.65. Figure 2 shows the standard deviation trajectories of the size of the first (top row) and the second (second row) largest clusters, number of clusters (third row), and adjusted Rand index of the cluster assignments (bottom row) as a function of temperature for three different signal-to-noise ratios SNR = 0.2 (left column), 0.25 (center column), and 0.65 (right column). The adjusted Rand index here measures the similarity between the obtained clusterings and ground truth at each temperature. The temperature parameter is shown on the log (base e) scale and will be further referred to as $\log T$.

At low temperatures voxels form a single cluster for all three signal strengths. As temperature rises, the cluster begins to split causing variations in the size of the clusters and an abrupt decrease of the adjusted Rand index. The changes in the standard deviations trajectories for the two largest clusters occur at the same temperature, since the reduction in the size of the initial frozen cluster is equivalent to formation of another cluster(s). The temperature, at which changes in the system take place, varies for signals with different SNR. For SNR = 0.2, the frozen cluster immediately dissolves into 35 clusters at $\log T = -4.6$, with the largest cluster containing 36 voxels with low-frequency time courses (Fig. 1, right top panel). The remaining clusters contain at most two voxels with high-frequency sig-

nals (Fig. 1, right bottom panel). The clustering at $\log T = -4.6$ is closest to the ground truth classification that so far has been observed, which is manifested by the jump in the value of the adjusted Rand index to 0.43. As temperature increases to $\log T = -4.45$ the interactions between the low-frequency signals vanish and the system crumbles into 63 clusters. The system continues to dissolve, and at $\log T > -1.4$ we observe a complete chaos with each voxel forming an independent cluster.

For the SNR = 0.25, standard deviations of the two largest clusters achieve their maximum at $\log T = -5.35$, while voxels are still being frozen together, indicating increasing instability of the system and forecasting the upcoming changes in the grouping structure. The first reorganization of the system takes place at $\log T = -5.65$ when two clusters coinciding with the true underlying structure of the data are formed (the adjusted Rand index is 1). However, the formation is not quite stable yet, and it is not until $\log T = -5.2$ that bonds between the two clusters vanish and we observe a stable plateau, during which voxels form two clusters coinciding with the true grouping structure of the data. As interactions within clusters weaken (see local maxima of the standard deviation trajectories at $\log T = -4.9$) the system splits into nine clusters at $\log T = -4.75$, with a subsequent drop of the adjusted Rand index to 0.82. The detrimental effect of rising temperature on the system structure continues as the number of clusters steadily increases to 74 singletons.

The evolution of the system at SNR = 0.65 is similar to that of SNR = 0.25. We would like to point out the following differences. First of all, stronger interactions between the voxels require higher temperature to be broken down. First signs of the approaching changes appear at $\log T = -5.5$ (compare to $\log T = -5.65$ for SNR = 0.25) and the stable clusters are formed at $\log T = -5.05$ (compare to $\log T = -5.2$ for SNR = 0.25). Another important feature is the prolonged duration of the stable phase, which in this case spans the interval of $\log T$ between -5.05 and -3.85 and is due to stronger interactions between the simulated time courses with SNR = 0.65. During the stable phase, both standard deviation trajectories steadily decline and spikes at $\log T = -3.7$ correspond to another change in the system structure when a single voxel was separated from the cluster containing high-frequency signals. This cluster continues to dissolve and at $\log T = -3.1$ the system experiences yet another drastic change in its organization, when the high-frequency cluster completely dissolves, while the low-frequency cluster remains largely intact due to stronger correlations between the signals that comprise it. The remaining low-frequency cluster finally dissipates at $\log T = -2.9$ into a multitude of clusters containing predominantly a single voxel.

Effect of K on algorithm performance

In general, the size of the neighborhood in imaging applications is dictated by the intrinsic characteristics of

Figure 3.

Left: Adjusted Rand index as a function of the neighborhood size, K , for the simulated data with $\text{SNR} = 0.2$ (dotted line), 0.25 (dot-dash line), and 0.65 (solid line); right: the duration of the stable phase for $\text{SNR} = 0.25$ and 0.65 as a function of K ; the duration was defined as the length of the $\log T$ interval spanning the stable phase. Extreme values of K have a negative impact on

the algorithm performance even for high SNR data. The range of K for which the algorithm performs best decreases as SNR decreases. The duration of the stable phase is longer for data with highest SNR, whereas for the $\text{SNR} = 0.25$ the transitions happen almost momentarily (except for $K = 10$); this is also the case for $\text{SNR} = 0.2$, so the null line was omitted.

the image. For example, a four- or a eight-neighbor systems are naturally used for the regular rectangular lattice of sites and a six-neighbor system is common for the regular hexagonal lattice. In the current paper we operate in a high dimensional space and, therefore, the choice of the size of the neighborhood is not entirely obvious. Choosing the right value of K is similar to the problem of choosing the proper bandwidth in the kernel density estimation, i.e. too large values of K tend to conceal fine clustering structures in the data, while too small values of K would lead to spurious groupings. There is no strict rule for the proper choice of K and one should decide on it using data at hand.

To illustrate the effect of K on the accuracy of clustering we classified simulated data for various values of K in the definition of the K -nearest neighbor graph. Figure 3 shows (left) the values of the adjusted Rand index measuring the similarity between the recovered clustering and the ground truth and (right) the duration of the stable phase, defined as the length of the $\log T$ interval corresponding to the stable phase, as functions of K for $\text{SNR} = 0.2$ (dotted line), 0.25 (dot-dash line), and 0.65 (solid line). For $\text{SNR} = 0.65$ the algorithm performance is very solid and for the graph spanned by five or more nearest neighbors the Potts model clustering correctly classifies all data points. When the graph structure is defined by a minimum spanning

tree ($K = 1$), the method correctly uncovers the membership of about 74% of voxels, but the remaining voxels form nine small-size clusters, which causes a drop in the value of the adjusted Rand index. When the size of the neighborhood grows to 40 and becomes larger than the size of the clusters present in the data, the stable phase vanishes due to the increased interaction range, which does not anymore suppress the interactions between the noisier high-frequency voxels when compared with the interactions between the low-frequency signals, and the frozen cluster is seen to be instantly dissolved by the rising temperature.

At lower $\text{SNR} = 0.25$ the behavior of the system is quite different. The reduced interaction range a in the minimal spanning tree graph ($K = 1$) and a high noise content of the observed signals weaken the interaction strengths within clusters, thus, obscuring the existence of the high density regions. Consequently, the initial phase where all voxels are frozen into a single cluster does not exist and the system is completely disordered even at low temperatures. For $K = 5$ and $K = 7$ the increased interaction range elevates the interaction strength between the voxels, and at null temperature we find not one but two large clusters, which closely resemble the true structure of the system (the adjusted Rand index is about 0.84 with 92% of voxels being classified correctly). These clusters are momentarily

Figure 4.

Adjusted Rand index as a function of SNR to compare the performance of the Potts model clustering (solid line) to that of the hierarchical clustering with dendrogram sharpening (dotted line). The hierarchical clustering based method was more sensitive at SNR = 0.2 when compared with the Potts model based clustering with the minimum cluster size set to one (left). The Potts

model clustering with the minimum allowable cluster size set to five (right) outperforms the hierarchical method for SNR = 0.13 and gives reliable outcome for SNR = 0.2 by correctly identifying cluster membership of 92% of voxels. Both methods proved to be effective in recovering the clustering structure for SNR > 0.25.

dissolved by rising temperature for $K = 7$ but survive longer for $K = 10$ (the log T interval spanning the stable phase is $[-5.29, -4.61]$, see also Fig. 3, right) before the system plunges into chaos. For $15 \leq K < 30$ the frozen cluster formed at low temperatures dissolves into a cluster containing low-frequency voxels and a number of small clusters with high-frequency signals. For larger $K \geq 30$, the stable phase vanishes and the system goes from the frozen into chaotic state in a single jump.

Reducing SNR to 0.2 we observe small changes in the behavior of the system. For $K < 5$ a completely disordered system exists even at low temperatures and the frozen state of the system is not observed until $K = 7$. For $K = 5$ at the null limit there exists a number of predominantly single-voxel clusters formed by the high-frequency signals and a large cluster containing all the low-frequency signals (the adjusted Rand index is 0.38), which instantly dissolves by rising temperature. The observed duration of the stable phase for the SNR = 0.2 is essentially zero for the entire range of K values and was not plotted on Figure 3 (right).

Comparison with other clustering methods

We used the simulated data to compare the Potts model based clustering to hierarchical clustering empowered by the dendrogram sharpening previously used in fMRI connectivity analysis [Stanberry et al., 2003, 2006]. The two

clustering methods are similar in that they do not require a priori assumptions of the number or nature of clusters present in the data and, therefore, the comparison should be fair. Figure 4 (left) shows the performance of the two methods as a function of SNR. Both methods were found to be equally effective in distinguishing between the different signal profiles for an SNR of 0.25 and higher. The dendrogram based method appears to be more sensitive correctly classifying voxels with an SNR as low as 0.2, while the Potts model clustering at SNR = 0.2 preserves the low-frequency cluster and splits the high-frequency one. This seemingly superior performance of the hierarchical algorithm is explained by its definition. The method, by default, classifies all the voxels by assigning them to the closest cluster core identified in the data set reduced by sharpening. Thus, the number of clusters in the data is strictly less than the number of data points.

As we have seen in the Potts model based method, clusters formed by signals with low SNR are very unstable and melt into a multitude of single-point clusters, whereas clusters with high SNRs are preserved. Single-voxel clusters often do not have a meaningful interpretation, i.e. their occurrence in fMRI data is hardly justifiable from the biological standpoint. To address this issue we impose a lower bound on the size of the final clusters. This modification also makes the performance of the algorithm during the *chaotic* phase to be alike to the single linkage method with the number of clusters determined as a function of

Figure 5.

the minimum allowable cluster size. Restricting any cluster to contain at least five voxels, Potts model clustering outperforms the hierarchical method for $\text{SNR} = 0.13$ in terms of the adjusted Rand index and gives reliable outcome for $\text{SNR} = 0.2$ by correctly identifying cluster membership of 92% of voxels while misclassifying only a few high-frequency signals.

fMRI Data

Motor paradigm

We applied the Potts model clustering to the motor task data twice, first imposing no restriction on the size of the smallest cluster, and then reclustering the data with the size of the smallest cluster set to 50. Figure 5 shows changes in the size of the three largest clusters (top plot), the number of clusters in the data (middle plot), and the adjusted Rand index comparing partitions at neighboring temperatures as a function of temperature (on the log scale); the minimum allowed cluster size was one. Initially

Figure 5.

Top: The standard deviation trajectories of the three largest clusters; center: number of clusters; and bottom: adjusted Rand index as functions of temperature (log scale) for the motor task experiment (the minimum cluster size set to one). The adjusted Rand index measures the similarity of cluster assignments at given and preceding temperatures. Phase transitions are identified by the peaks of standard deviation trajectories and abrupt changes of the adjusted Rand index. The stable phase spans the log T interval of $[-6.25, -5.95]$, which is reflected by the similarity of neighboring partitions in this temperature range (bottom plot). Relatively high rate of the task presentation (~ 0.1 Hz) suppresses the signal changes, masking the transition to the chaotic state, so the number of clusters increases monotonically with rising temperatures. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 6.

Clustering results for the motor-task experiment determined at $\log T = -5.95$; minimum cluster size was set to one. Clusters were color-coded and overlaid on high-resolution anatomical images. The largest cluster (red) corresponds to the motor cor-

tex and supplementary motor area engaged in the motor task; the remaining clusters predominantly contain a single voxel. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 7.

Top: The standard deviation trajectories of the three largest clusters; center: number of clusters; and bottom: adjusted Rand index as functions of temperature (log scale) for the motor task experiment (the minimum cluster size set to 50). The adjusted Rand index measures the similarity between the cluster assignments at the given and preceding temperatures. The changes in the clustering structure are now more apparent and take place at higher temperatures. During the stable phase, which is spanned by a $\log T$ interval of $[-1.49, -0.86]$ and is twice as wide as the one observed in Figure 5, cluster assignments remain robust (the adjusted Rand index is about 0.8). Since the number of clusters present in the system is bounded above the number of clusters and the similarity between the neighboring partitions stabilize during the chaotic phase ($\log T > -0.6$). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 8.

Top: Clustering results for the stable phase at $\log T = -0.98$ (minimum cluster size is 50) identify the largest cluster (red) corresponding to the motor cortex and supplementary motor area. The remaining clusters could be attributed to the motion component and vascular contributions; bottom: disordered state at $\log T = -0.86$; the imposed cluster size restriction accentuates strong dependencies between voxels in the immediate Euclidean neighborhood. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 8.

Figure 9.

Fourier and wavelet spectra (using the continuous Morlet transform) for the three largest clusters identified for the periodic motor task (minimum cluster size is 50). The largest contribution to the mean time course of the first largest cluster, corresponding to the motor cortex and supplementary motor area, comes from the paradigm frequency of 0.1 Hz (dashed lines on Fourier spectra plots). The contributions are persistent through

the task and mimic the paradigm design [see the high periodic values (dark red spots) at 0.1 Hz on the wavelet spectrum]. This frequency contributes to the mean time courses of the other two clusters, but is not dominant. Note that respiratory artifacts were accurately localized (0.4 Hz oscillations in the second row) into a separate cluster. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

$\log T = -6.25$ precede the meltdown of the system at $\log T = -6.1$ into about 500 clusters containing at most two voxels, except one 450-voxel cluster. The content of the largest cluster remains approximately the same as $\log T$ increases to -5.95 , but it literally crumbles as noisy voxels separate from it to form independent groups as temperature rises. Consequently, the number of clusters in the system steadily increases and the similarity index drops. The largest cluster (red) identified at $\log T = -5.95$ corresponds to the motor cortex and supplementary motor area engaged in the motor task, while the remaining clusters predominantly contain a single voxel (Fig. 6, the clusters were color-coded to agree with the previous figure and overlaid on high-resolution anatomical images).

Imposing the restriction on the size of the smallest cluster to contain at least 50 voxels, consistently with our expectations we observe (Fig. 7, top plot) more pronounced changes in the standard deviation trajectories of the first two largest clusters at $\log T = -1.61$ and at $\log T = -0.92$, forecasting upcoming changes in the clustering structure. At $\log T = -1.49$, the system splits into four clusters (Fig. 7, middle plot), causing the immediate drop in the adjusted Rand index (Fig. 7, bottom plot). During the stable phase of the system the size of the clusters and their content remain approximately unchanged, which is reflected by the plateau of the standard deviation trajectories and relatively high values of the adjusted Rand index, which are on the order of 0.8 being equivalent to the change in cluster assignments for less

Figure 10.

Standard deviation trajectories of the three largest clusters for the resting state data vary across subjects both in shape and magnitude. The low signal strength reduces the duration of the stable phase, which, nevertheless, is clearly visible. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

than 6% of voxels. The transition from the stable to chaotic state occurs immediately after the second peak of the standard deviation trajectories ($\log T = -0.86$) as the system splits into 10 clusters each containing 50–200 voxels. This change in the system structure is accompanied by the sharp decrease of the adjusted Rand index (<0.2 , corresponding to the change of cluster assignments of 35% of voxels). As temperature increases, the system gradually dissolves. However, the minimum cluster size condition restricts the number of clusters present in the system and, therefore, for higher temperatures ($\log T > -0.6$) the number of clusters stabilizes and the similarity between the neighboring partitions increases.

Figure 8 (top plot) shows the clustering results for the stable phase at $\log T = -0.98$; the clusters were color-coded to match with color map on Figure 7 and overlaid on high-resolution anatomical images. The large red cluster corresponds to the motor cortex and supplementary motor area engaged in the motor task. The remaining clusters could be attributed to a motion component or vascular contributions. Frequency spectra and wavelet power spectra of the mean time courses of the three largest clusters (see Fig. 9) show a strong influence of the paradigm frequency (0.1 Hz) to the time courses of all, but the third, clusters. As expected, the power of the task frequency

closely mimics the task design in the motor areas (Fig. 9, top right). Respiratory artifacts are noticeable in the second largest cluster (Fig. 9, second row) as the periodic spikes at around 0.4 Hz. The cluster partition at $\log T = -0.86$ (see Fig. 7) shows that at higher temperatures the red cluster associated with the motor task splits into multiple clusters of considerably smaller sizes.

Resting-state connectivity

Figure 10 shows the standard deviations trajectories of the three largest cluster sizes for nine subjects under resting-state conditions. The lower limit for the cluster size was set to 50. The standard deviation profiles are similar to those observed in the previous experiment. The trajectories vary across subjects both in shape and magnitude. The range of temperatures spanning the stable clustering phase is much shorter for the resting state data than for the motor paradigm data. The temperature scale is approximately the same, suggesting the equivalent interaction strength among all nine data sets.

The temperature corresponding to the last peak in the variation of the size of the second largest cluster was used to determine the consensus clustering structure for each subject. Group connectivity maps show that at rest the *anterior*

Figure 11.

Group functional connectivity maps of (obtained by Potts model clustering) of the anterior cingulate cortex show its connections to the inferior, superior, and middle frontal orbital cortex, parahippocampal gyrus, fusiform gyrus, right hippocampus, middle

temporal gyrus, thalamus, posterior cingulate cortex, superior medial frontal cortex, and left postcentral gyrus in the resting-state data. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

cingulate cortex is functionally connected to the bilateral inferior, superior and middle frontal orbital cortex, superior medial frontal cortex, bilateral parahippocampal gyrus, bilateral fusiform gyrus, right hippocampus, bilateral middle temporal gyrus, thalamus, posterior cingulum, and left postcentral gyrus (see Fig. 11); whereas functional connections of the *posterior* cingulate cortex (see Fig. 12) extend to the thalamus, anterior cingulate cortex, left postcentral gyrus, bilateral

inferior frontal gyrus, hippocampus, orbital frontal gyrus, left fusiform gyrus, and right parahippocampal gyrus. Figure 13 shows the group spectral density estimates and changes in the frequency content over time for the mean time courses of the anterior and posterior cingulate cortices. The group estimates were obtained by averaging frequency histograms and time-frequency spectra across individuals. The spectra look fairly similar for the two structures. The

Figure 12.

Group functional connectivity maps (obtained by Potts model clustering) of the posterior cingulate cortex during rest include the thalamus, anterior cingulate cortex, left postcentral gyrus, bilateral inferior frontal gyrus, hippocampus, orbital frontal, left

fusiform, and right parahippocampal gyri. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 13.

Group Fourier and wavelet spectra (using the continuous Morlet transform) for the anterior cingulate cortex (ACC, top row) and posterior cingulate cortex (PCC, bottom row) obtained by

averaging frequency and time-frequency spectra of the cortices across the nine subjects. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 14.

Clustering results for the periodic motor task obtained by extending the K -nearest neighbor graph to include edges with immediate spatial neighbors in six orthogonal directions are sim-

ilar to those observed with a K -nearest neighbor graph based only on functional distances. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

dominant component of the oscillations corresponds to the low-frequency range below 0.12 Hz, and the magnitude of these oscillations varies with time, exhibiting a strong increase closer to the end of the task.

Incorporating spatial information in connectivity analysis

We explored the possibility to incorporate spatial information from the Euclidean neighborhood via extending the K -nearest neighbor graph to include edges of the six immediate (in 3-D) neighboring voxels. That is for any two voxels identified as spatial neighbors in a conventional Euclidean sense, but not as K -nearest neighbors, the edge was added to the existing K -nearest neighbor graph and the corresponding interaction term was computed as in (2). The standard deviation trajectories were rather similar to those observed previously slightly differing only in the temperature scale. The quality of the obtained connectivity maps (see Fig. 14) were also alike to those described earlier with the biggest identified cluster corresponding to the motor activity.

DISCUSSION

We have demonstrated the use of the Potts model based clustering as an automatic tool to determine functional connectivity networks during periodic stimuli and in resting-state data. The rationale for the method stems from the ideas of statistical mechanics and information theory. The method has a very intuitive physical interpretation and gives the solution in terms of a probability distribution of the possible cluster assignments, which is then used to determine a unique group labeling. Another advantage of the Potts model based method is its stochastic formulation, which allows it to be incorporated into hierarchical probabilistic models.

In contrast to the seed voxel based connectivity approach the method is totally automatic, i.e., there is no need for expert intervention to construct the connectivity maps. Moreover, it is exhaustive in the sense that there is no omission of seemingly connected brain regions. The Potts model clustering is a kernel-based method [see Murua et al. (On Potts model clustering, kernel K -means and density estimation, *J Comput Graph Stat*, under review) for connection between the Potts model clustering and weighted kernel K -means] and, as such, shares the flexibility and good performance of kernel-based clustering methods including kernel K -means [Girolami, 2002] and the multiway normalized cut [Meila and Xu, 2003]. However, in contrast to the K -means algorithm one makes no assumption on the distribution of the data or the expected number of clusters.

Commonly to other clustering techniques the method uses the dissimilarity measure between the objects. In addition, the interaction kernel and the number of nearest neighbors K to create a graph structure are required to be specified. Once the samples from the underlying probability model are acquired it remains to identify the range of temperatures corresponding to the stable partitioning

phase and determine the final cluster assignment. We explored the performance of the method and its sensitivity to various parameters on simulated data and recapitulate the main findings later.

Effect of Parameters on Method Performance

The choice of the similarity measure

The Potts model clustering is based on the dissimilarity measure and avoids the explicit modeling of the observed signals. In this paper, the interactions were computed from the low frequency specific correlations. A paradigm design, stimuli administration, and physiological processes might induce spurious correlations that could potentially lead to erroneous connectivity maps. The problem can be circumvented by, for example, using first- or higher order partial correlations [Jobson, 1992] to account for potential confounders, or by defining the interactions between the feature vectors derived from the data rather than raw observations [Goutte et al., 2001]. The use of the correlation coefficient in this study was justified by the preprocessing of the data, absence of the task-related confounders in the resting-state data, and restricted spatial coverage of the brain volume for the motor task experiment. It is of interest to note that although the paradigm frequency induces notable signal changes not just in the motor area but in voxels attributed to various artifacts (Fig. 9), the lower ends of the frequency spectra were substantially different and, as a result, voxels were assigned to separate clusters.

The effect of the form of the interaction kernel

The effect of the form of the interaction kernel on algorithm performance has been experimentally shown to be negligible by previous researchers [Wiseman et al., 1998]. In particular, the standard deviation trajectories for the power law interactions, interactions specified by the Gaussian kernel, and interactions specified by the exponential kernel were found to be similar and differ only in the range of temperatures where the changes in the system occur. The irrelevance of the interaction form can also be seen through the connection of the Potts model based clustering to the kernel density estimation problem [Silverman, 1986; Murua et al., On Potts model clustering, kernel K -means and density estimation, *J Comput Graph Stat*, under review), for which it has been established both theoretically and experimentally [Epanechnikov, 1969] that the choice of the kernel has negligible effect on the performance of the estimation method.

Effect of K

The Potts model was originally considered on a system of spins confined to the regular lattice [Wu, 1982]. In this work, we did not view data as acquired on a regular lattice corresponding to the 3D representation of the brain;

rather, we “mapped” the data into a functional space where the distances between voxels were defined by the similarity of the observed time courses. The transformed data were not forming a regular lattice anymore and the definition of the neighborhood structure was not straightforward. One of the possible choices in this situation was to identify neighbors using the K -neighborhood criterion. Choosing K too big or too small would either smooth out the important features in the data or make its structure appear too noisy. The effect of K was not found to be crucial to the performance of the clustering procedure for a many-cluster model with noise [Wiseman et al., 1998], however, the size of the neighborhood considered in that paper did not vary considerably ($K = 6, 12, 14, 18$, and the Voronoi cell structure were studied), the contrast to noise ratio of the image, though not specified, appeared to be reasonably high, and the performance of the algorithm was measured in terms of accuracy, a measure that is less conservative than the adjusted Rand index.

Our experiments on simulated data indicate a notable effect of the extreme values of K on the algorithm performance, which varies with SNR of the data. For high SNRs, when clusters are clearly distinguishable, the influence of K is negligible as the algorithm perfectly uncovers the structure for a large range of values. However, extremely small values of K produce completely disordered (noisy) systems, whereas too large values of the parameter inflate the interaction strength even between points from different clusters completely annihilating the stable phase. These effects are amplified by the deteriorating signal strength since interactions grow rapidly as dissimilarity between the voxels increases. Hence, the range of K where the algorithm performs well shortens as signal content weakens. Based on our experimental results, we settled for $K = 10$ for fMRI data collected at 1.5 T.

Comparison with other clustering methods

Previous studies have shown [Blatt et al., 1997; Murua et al., On Potts model clustering, kernel K -means and density estimation, J Comput Graph Stat, under review]) superior performance of the Potts model clustering for real and simulated data in comparison with other clustering techniques including the hierarchical single linkage, complete linkage, minimal spanning tree, directed graph method [Fukunaga, 1990], etc. Our study shows that a minimum cluster size restriction improves the results of the method, which performs better or as good as a hierarchical clustering with dendrogram sharpening for a wide range of SNRs. The analysis of the fMRI motor task data with and without a cluster size restriction demonstrated other advantages of this condition including the increased duration of the stable phase and more apparent phase transitions.

Mapping Connectivity in Functional MRI Data

We have demonstrated how the Potts model clustering method can be applied to the functional connectivity prob-

lem using data from a conventional box-design motor task and a more complex resting-state condition. Following the changes in the data clustering structure over the extended temperature range we observed a clear transition from the frozen to the stable state. The observed steady monotone increase in the number of identified clusters (see Fig. 5) is due to a relatively high frequency of the on-off task design (0.1 Hz), which subdues the contrast between the “on” and “off” images [Hathout et al., 1999].

By imposing the lower limit on the size of the clusters, we force the system to remain in any present state for a longer range of temperatures, and higher temperatures were required for the bonds to be broken. Consequently, the phase transitions were taking place at higher temperatures, the stable phase became more pronounced, and the range of the temperatures spanning the stable phase increased. The restriction on the size of the smallest cluster is equivalent to the condition on the largest number of clusters to be present in the system at any given temperature. Therefore, once the transition to the chaotic phase occurs, clusters containing only a few voxels are forced to join together, which explains why the number of clusters in the system and the similarity index of neighboring assignments are stabilized at high temperatures (Fig. 7, middle and bottom plots).

Relative robustness of the clustering solutions within the stable phase allows one to select a partition corresponding to any log T within the stable phase interval as a final clustering solution. Alternatively, one could construct an augmented partition for the stable phase by estimating the proportion of times each pair of voxel was assigned to the same cluster and then thresholding it to get a unique labeling.

Resting state connectivity

The resulting functional connectivity maps for the anterior and posterior cingulate cortices during rest are consistent with earlier findings by Greicius et al. [2003] obtained with a seed voxel approach. In particular, Greicius et al. [2003] showed the connectivity of the *posterior* cingulate cortex with the medial prefrontal cortex, and found the ventral *anterior* cingulate cortex to be connected to the posterior cingulate and medial prefrontal cortex, including orbital frontal cortex and hypothalamus. However, there are some notable differences between the connectivity maps of Greicius et al. [2003] and ours. Our findings show the posterior cingulate cortex to be connected to both the right parahippocampal gyrus and the bilateral hippocampus, whereas Greicius et al. [2003] only found a connection to the left parahippocampal gyrus. We also found the posterior and anterior cingulate cortices to be connected to the left and bilateral fusiform gyrus, respectively. These differences could be in part due to the reduction of the signals from the entire brain structure (anterior or posterior cingulate cortices) to the sole mean time course of the region in the seed-voxel approach implemented in Greicius et al. [2003].

Incorporating spatial information into analysis

Conceptually, it would be beneficial to account for the Euclidean neighborhood system in the analysis, provided it improves the results. Spatial information can be incorporated in a number of ways, the simplest being to consider only a nearest-neighbor Euclidean system not regarding the similarity of the observed signals. With regard to the functional connectivity problem, a simple presumption of the conventional Euclidean neighborhood structure is fallacious considering complicated folding patterns of gray matter, though one could try and solve the clustering problem on the flattened cortical surface where the assumption of the Euclidean spatial dependences seems more appropriate (note that this approach would exclude structures deep inside the brain from the analysis). Another way to account for the natural 3D structure is, as it was done in the paper, to augment the existing K -nearest neighbor graph structure with edges defined by spatial proximities. Alternatively, one could redefine the distance measure as a weighted combination of the Euclidean distances and the similarity measure of interest or update the interaction term to include morphometric characteristics of the brain structures.

Although, our experiments did not indicate the effect of the combined neighborhood structure on clustering results, it would be of general interest to pursue this venue and further explore the use of spatial information with the goal to increase the sensitivity and accuracy of the clustering method.

CONCLUSIONS

We introduced an unsupervised stochastic clustering method based on the ferromagnetic Potts model, which provides a powerful alternative to study coherent signal changes and can be effectively used as an automatic tool to determine functional connectivity networks under various complex conditions such as the resting-state and continuous stimuli. We demonstrated the specifics of the method by applying it to the simulated data and a conventional on-off finger tapping paradigm. The method shown in the literature to be insensitive to the choice of the interaction kernel was found to give a robust performance for data with SNR as low as 0.2. The ability of the method to recover a hidden clustering structure depends on the size of the neighborhood and one should avoid using the extreme values of K . Our experiments showed the stability of the resulting clustering solutions in the sense that there was little variation of the cluster assignments within the stable phase. We observed that a lower bound on the size of the smallest cluster slows the changes in the clustering structure, increases the duration of the stable phase, and improves the clustering results for low SNR data. Spatial information was not found to enhance the clustering solutions, though additional research on this topic is encouraged. Functional connections of the anterior and posterior cingulate cortices for a group of nine healthy subjects dur-

ing rest corroborate previous findings, therefore, attesting to the ability of the Potts model clustering to operate on the complex data sets.

Potts model based method provides an elegant way for functional connectivity mapping since it makes no assumption with regard to the underlying grouping structure of the data and can *simultaneously* identify functional networks present in the data. To our knowledge, Potts model clustering and ICA [Beckman and Smith, 2005; Cordes et al., 1999] are so far the only two methods reported in the literature that are capable of simultaneous identification of connectivity networks in resting-state data.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Tilmann Gneiting at the University of Washington for his valuable suggestions and the anonymous referees whose comments helped to enhance the quality of the manuscript. This research was in part supported by a special multidisciplinary center grant P50-33812 from the National Institute of Child Health and Human Development (NIHCD) and by a fellowship from the National Science Foundation program for Vertical Integration of Research and Education in the Mathematical Sciences (VIGRE).

REFERENCES

- Agrawal H, Domany E (2003): Potts ferromagnets on coexpressed gene networks: Identifying maximally stable partitions. *Phys Rev Let* 90:158102.
- Baumgartner R, Windischberger C, Moser E (1998): Quantification in functional magnetic resonance imaging: Fuzzy clustering vs. correlation analysis. *Magn Reson Imaging* 16:115–125.
- Beckmann CF, Smith SM (2005): Tensorial extensions of independent component analysis for group fMRI data analysis. *Neuroimage* 25:294–311.
- Biswal B, Yetkin F, Haughton V, Hyde J (1995): Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34:537–541.
- Blatt M, Wiseman S, Domany E (1996): Superparamagnetic clustering of data. *Phys Rev Let* 76:3251–3254.
- Blatt M, Wiseman S, Domany E (1997): Data clustering using a model granular magnet. *Neural Comput* 9:1805–1842.
- Cordes D, Carew J, Eghbalnia H, Meyerand E, Quigley M, Arfanakis K, Assadi A, Turski P, Haughton V (1999): Resting-state functional connectivity study using independent component analysis. In: *Proceedings International Society for Magnetic Resonance in Medicine, 7th Annual Meeting*, Philadelphia.
- Cordes D, Haughton V, Arfanakis K, Wendt G, Turski P, Moritz C, Quigley M, Meyerand M (2000): Mapping functionally related regions of the brain with functional connectivity MR imaging. *Am J Neuroradiol* 21:1636–1644.
- Cordes D, Haughton V, Carew J, Arfanakis K, Maravilla K (2002): Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn Reson Imaging* 20:305–317.
- Domany E (1999): Superparamagnetic clustering of data—The definitive solution of an ill-posed problem. *Phys A* 263:158–169.
- Domany E (2003): Cluster analysis of gene expression data. *J Stat Phys* 110:1117.

- Domany E, Blatt M, Gdalyahu Y, Weinshall D (1999): Superparamagnetic clustering of data: Application to computer vision. *Comp Phys Commun* 121/122:5–12.
- Downton M, Brennan T (1980): Comparing classifications: An evaluation of several coefficients of partition agreement. Paper presented at the meeting of the Classification Society, Boulder, CO.
- Einav U, Tabach Y, Getz G, Yitzhaky A, Ozbek U, Amarglio N, Izraeli S, Rechavi G, Domany E (2005): Gene expression analysis reveals a strong signature of an interferon-induced pathway in childhood lymphoblastic leukemia as well as in breast and ovarian cancer. *Oncogene* 24:6367–6375.
- Epanechnikov VA (1969): Non-parametric estimation of a multivariate probability density. *Theory Prob Appl* 14:153–158.
- Filzmoser P, Baumgartner R, Moser E (1999): A hierarchical clustering method for analyzing functional MR images. *Magn Reson Imaging* 17:817–816.
- Fowlkes EB, Mallows CL (1983): A method for comparing two hierarchical clusterings, (with comments and rejoinder). *J Am Stat Assoc* 78:553–584.
- Friston KJ (1995): Functional and effective connectivity in neuroimaging: A synthesis. *Hum Brain Mapp* 2:56–78.
- Friston KJ, Frith C, Liddle FP, Frackowiak RSJ (1993): Functional connectivity: The principal-component analysis of large (PET) data sets. *J Cereb Blood Metab* 13:5–14.
- Fukunaga K (1990): *Introduction to Statistical Pattern Recognition*. San Diego: Academic Press.
- Getz G, Domany E (2003): Coupled two-way clustering server. *Bioinformatics* 19:1153–1154.
- Getz G, Levine E, Domany E, Zhang M (2000a): Super-paramagnetic clustering of yeast gene expression profiles. *Phys A* 279:457–464.
- Getz G, Levine E, Domany E (2000b): Coupled two-way clustering analysis of gene microarray data. *Proc Natl Acad Sci USA* 97:12079–12084.
- Girolami M (2002): Mercer kernel based clustering in feature space. *IEEE Trans Neur Net* 13:669–688.
- Goutte C, Toft P, Rostrup E, Nielsen FA, Hansen, LK (1999): On clustering fMRI time series. *Neuroimage* 9:298–310.
- Goutte C, Hansen LK, Liptrot MK, Rostrup E (2001): Feature-space clustering for fMRI meta-analysis. *Hum Brain Mapp* 13:165–183.
- Greicius M, Krasnow B, Reiss A, Menon V (2003): Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proc Natl Acad Sci USA* 100:253–258.
- Hampson M, Peterson B, Skudlarski P, Gatenby J, Gore J (2002): Detection of functional connectivity using temporal correlations in MR images. *Hum Brain Mapp* 15:247.
- Hampson M, Olson I, Leung H, Skudlarski P, Gore J (2004): Changes in functional connectivity of human MT/V5 with visual motion input. *Neuroreport* 15:1315–1319.
- Hathout GM, Copi RK, Bandettini P, Gambhir SS (1999): The lag of cerebral hemodynamics with rapidly alternating periodic stimulation: Modeling for functional MRI. *Magn Reson Imaging* 17:9–20.
- Hubert L, Arabie P (1985): Comparing partitions. *J Classif* 2:193–218.
- Jobson JD (1992): *Applied Multivariate Data Analysis*, Vol. 2. New York: Springer-Verlag.
- Kullmann L, Kertesz J, Mantegna RN (2000): Identification of clusters of companies in stock indices via Potts super-paramagnetic transitions. *Phys A* 287:412–419.
- Lowe M, Mock B, Sorenson J (1998): Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *Neuroimage* 7:119–132.
- Lowe M, Dzemidzic M, Lurito J, Mathew V, Phillips M (2000): Correlations in low-frequency BOLD fluctuations reflect cortico-cortical connections. *Neuroimage* 12:582–587.
- Lowe M, Phillips M, Lurito J, Mattson D, Dzemidzic M, Mathews V (2002): Multiple sclerosis: Low-frequency temporal blood oxygen level-dependent fluctuations indicate reduced functional connectivity—Initial results. *Radiology* 224:184–192.
- McKeown MJ, Sejnowski TJ (1998): Independent component analysis of fMRI data: Examining the assumptions. *Hum Brain Mapp* 6:368–372.
- Meila M, Xu L (2003): *Multiway cuts and spectral clustering*, Department of Statistics, University of Washington: Seattle, WA. Technical report No 442.
- Milligan G, Cooper M (1986): A study of the comparability of external criteria for hierarchical cluster analysis. *Multivariate Behav Res* 21:441–458.
- Morey L, Agresti A (1984): The measurement of classification agreement: An adjustment to the Rand statistic for chance agreement. *Educ Psychol Meas* 44:33–37.
- Ott T, Kern A, Schuffenhauer A, Popov M, Acklin P, Jacoby E, Stoop R (2004): Sequential superparamagnetic clustering for unbiased classification of high-dimensional chemical data. *J Chem Inf Comput Sci* 44:1358–1364.
- Ott T, Kern A, Steeb W, Stoop R (2005): Sequential clustering: tracking down the most natural clusters. *J Stat Mech* 11:P11014.
- Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004): Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16:1661–1687.
- Rand WM (1971): Objective criteria for the evaluation of clustering methods. *J Am Stat Assoc* 66:846–850.
- Reichardt J, Bornholdt S (2004): Detecting fuzzy communities in complex networks with a Potts model. *Phys Rev Lett* 93: 218701.
- Silverman BW (1986): *Density estimation for statistics and data analysis*, London, New York, Chapman and Hall.
- Sokal AD (1996): Monte-Carlo methods in statistical mechanics: Foundations and new algorithms. In: *Lectures at the Cargese Summer School on “Functional Integration: Basics and Applications.”*
- Stanberry L, Nandy R, Cordes D (2003): Cluster analysis of fMRI Data using dendrogram sharpening. *Hum Brain Mapp* 20:201–219.
- Stanberry L, Richards T, Berninger VW, Nandy R, Aylward E, Maravilla K, Cordes D (2006): Low frequency signal changes reflect differences in functional connectivity between good readers and dyslexics during continuous phoneme mapping. *Magn Reson Imaging* 24:217–229.
- Stein T, Moritz C, Quigley M, Cordes D, Haughton V, Meyerand E (2000): Functional connectivity in the thalamus and hippocampus studied with functional MR imaging. *Am J Neuroradiol* 21:1397–1401.
- Swendsen RH, Wang JS (1987): Nonuniversal critical dynamics in Monte Carlo simulations. *Phys Rev Lett* 58:86–88.
- Tzourio-Mazoyer N, Landeau B, Papathanassiou DF, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002): Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15(1): 273–289.
- Vaidya PM (1989): An $O(n \log n)$ algorithm for the all-nearest-neighbors problem. *Disc Comp Geom* 4:101–115.
- von Ferber C, Worgoetter F (2000): Cluster update algorithm and recognition. *Phys Rev E* 62(Part A):1461–1464.
- Wiseman S, Blatt M, Domany E (1998): Superparamagnetic clustering of data. *Phys Rev E* 57:3767–3783.
- Wu FY (1982): The Potts model. *Rev Mod Phys* 54:235–268.
- Xiong J, Parsons L, Gao J, Fox P (1999): Interregional connectivity to primary motor cortex revealed using MRI resting-state images. *Hum Brain Mapp* 8:151–156.