Unsupervised Spatiotemporal Analysis of FMRI Data Using Graph-Based Visualizations of Self-Organizing Maps

Santosh B. Katwal*, *Member, IEEE*, John C. Gore, René Marois, and Baxter P. Rogers

Abstract—We present novel graph-based visualizations of self-organizing maps for unsupervised functional magnetic resonance imaging (fMRI) analysis. A self-organizing map is an artificial neural network model that transforms high-dimensional data into a low-dimensional (often a 2-D) map using unsupervised learning. However, a postprocessing scheme is necessary to correctly interpret similarity between neighboring node prototypes (feature vectors) on the output map and delineate clusters and features of interest in the data. In this paper, we used graph-based visualizations to capture fMRI data features based upon 1) the distribution of data across the receptive fields of the prototypes (density-based connectivity); and 2) temporal similarities (correlations) between the prototypes (correlation-based connectivity). We applied this approach to identify task-related brain areas in an fMRI reaction time experiment involving a visuo-manual response task, and we correlated the time-to-peak of the fMRI responses in these areas with reaction time. Visualization of self-organizing maps outperformed independent component analysis and voxelwise univariate linear regression analysis in identifying and classifying relevant brain regions. We conclude that the graph-based visualizations of self-organizing maps help in advanced visualization of cluster boundaries in fMRI data enabling the separation of regions with small differences in the timings of their brain responses.

Index Terms—Functional MRI (fMRI), SOM visualization, reaction time, self-organizing map.

I. INTRODUCTION

FUNCTIONAL magnetic resonance imaging (fMRI) is a noninvasive imaging technique that has emerged as a powerful tool to identify brain regions involved in cognitive processes. FMRI offers spatial and temporal resolutions adequate to measure the location, amplitude, and timing of brain activity. However, the low contrast to noise and generally poor resolutions of fMRI data present challenges in accurately mapping task-related brain regions.

The analysis of fMRI data aims to correctly identify the regions in brain that are activated during the course of a cognitive task. Several methods for analyzing fMRI data have been reported in the literature. These methods can be broadly classified into two categories: hypothesis-driven and data-driven methods. Statistical parametric mapping (SPM) based on the general linear model (GLM) is a commonly used hypothesis-driven method that assumes a simple parametric linear model for signals with a specific noise structure and uses voxel-based linear regression analysis [1]. The success of SPM for fMRI analysis is due mainly to the simplicity of the approach in principle and application. More complex parametric models based on Markov random field models [2], [3] and hidden Markov models [4] have also been used in fMRI analysis. These methods require information about stimulus timings and assumption about the shape and timing of the hemodynamic response and are more suitable for simple task paradigms. A typical SPM analysis performs linear convolution of an assumed hemodynamic response function (HRF) with the deterministic stimulus timing function to construct reference functions. These modeling assumptions and the deterministic character assigned to the stimulus timing function may be too restrictive to capture the broad range of possible brain activation patterns in space and time and across subjects. SPM performs voxel-by-voxel analysis which is massively univariate. Due to spatial coherence and temporal autocorrelation between brain voxels [5], a multivariate approach may be more appropriate than the voxel-by-voxel approach for fMRI analysis.

Data-driven methods follow multivariate approach for exploratory fMRI analysis. The most commonly used data-driven techniques for fMRI analysis include principal component analysis (PCA) [6], [7], independent component analysis (ICA) [8], [9], and data clustering [11]–[23]. Several data clustering algorithms have been used including K-means clustering [11], fuzzy clustering [12]–[14], hierarchical clustering [15], [16], and self-organizing maps (SOMs) [17]–[23]. The data-driven methods make few or no assumptions about HRF shape and do not require *a priori* knowledge about stimulus timings. However, each of these methods has its own shortcomings. PCA assumes that each component is mutually orthogonal and considers only second order statistics (variance). However, second order statistics

is not sufficient to characterize noisy fMRI datasets, and mutual orthogonality between components may not be applicable to fMRI data [7]. ICA works with higher order statistics (entropy or non-Gaussianity) to separate maximally independent sources from fMRI data. However, it makes strong assumptions about independence between spatial or temporal components that may result in biased decomposition [9], [10]. Additionally, the components are unranked and sometimes difficult to interpret. The K-means algorithm is constrained by the assumption that the clusters are spherically symmetric which is not applicable in the context of fMRI. K-means clustering results depend largely on the initial cluster assignment and the number of clusters, which must be defined *a priori* [11]. The results from fuzzy clustering also depend on the initial cluster assignment and fuzzy factor [20]. Hierarchical clustering overcomes the limitations of K-means and fuzzy clustering but suffers from computational complexity and inability to adjust once a merge or split decision has been made [11].

An SOM overcomes these limitations by topology-preserving unsupervised mapping of the high-dimensional data onto a low-dimensional (often a 2-D) lattice of nodes (neurons). Kohonen's SOM is an artificial neural network model that uses unsupervised learning to reveal underlying structures of the data [24], [25]. The SOM has been successfully used as a data-driven technique in fMRI for detecting brain activity and functional connectivity in an unsupervised way [21], [22].

The SOM transforms data through an adaptive vector quantization that ensures an orderly arrangement (map) of the prototypes (feature vectors) in the output space [25]. This helps in visualization of natural clusters in the data. However, a post-processing scheme is required to capture cluster boundaries in the map. The learned map has several components that may represent neighborhood relations and connectivity between prototypes in the data space. These include sizes of the receptive fields, local density distribution or distance (or similarity) across node prototypes. These components have been exploited for visualization of data clusters via several visualization schemes such as U-matrix [27] or its variants [28], [29], visualization-induced SOM [30], or double SOM [31]. A more comprehensive review on various visualization schemes of the SOM can be found in [26], [32], and [33].

FMRI data typically comprise a large noisy dataset where the magnitude of detectable signals could be very low and signals of interest may be confined to a few voxels in the high-dimensional image space. This makes the task of delineating important features and fine structures of data in fMRI data more difficult. In some previous fMRI studies, a two-stage clustering approach has been employed where the SOM is accompanied by a second-stage of clustering for effective delineation of fMRI signal patterns. The first stage constitutes the SOM with sufficient number of output nodes to ensure flexibility and good quality of mapping of the data. In the second stage, similar and redundant SOM output nodes are merged. To this end, fuzzy c-means [20], hierarchical clustering [22] or special node merging criteria employing least square distance [21] or reproducibility of the fMRI data across epochs [23] have been used to merge similar SOM nodes in fMRI analysis.

The two-stage clustering techniques have made use of temporal and spatial relationships between node prototypes in the SOM for cluster delineation. However, a critical measure of data topology provided by the distribution of data across the SOM nodes has been overlooked. The local density distribution is indicative of how the data are distributed within the receptive fields of node prototypes with respect to their neighbors. The inclusion of data topology in the visualization of the SOM provides important clues for delineation of fine cluster structures especially in high-dimensional, large and noisy datasets [26]. The data topology can be visualized by rendering of a topology-representing graph over SOM lattice.

In this study, we use a combination of two graph-based visualization techniques that incorporate 1) local density distribution across SOM prototypes; and 2) local similarities (correlations) between the prototypes. The combined visualization effectively captures cluster boundaries and delineates detailed connectivity structures of the meaningful data [34], [35]. We show that this approach improves detection of signals in fMRI data, including effective separation of regions that have small differences in the timings of their responses. We use this approach on a simulated dataset to identify and classify activated regions based on timings of the signals they constitute. We also use this approach to study the relationship between fMRI signal timing and reaction time (RT) during a visuo-motor task. We extract average fMRI signals from the regions identified via the SOM and compute time-to-peak using an inverse logit (IL) model of the hemodynamic response. We compare the SOM visualization method with ICA, another data-driven technique, and GLM-based multiple regression analysis using SPM. We fit a linear mixed-effects model on the time-to-peak measures and compare the slopes and intercepts of the fits to compare the timing accuracy of measurements obtained via the SOM, ICA, and GLM.

II. RELATED THEORY

A. SOM

Kohonen's SOM is a two-layer feedforward artificial neural network model that maps high-dimensional input data into a set of nodes arranged in a low-dimensional (often a 2-D) rigid lattice using unsupervised learning [24], [25]. A weight vector of the same dimension as the input vector is associated with each SOM node. The SOM algorithm consists of a series of training steps that tune the weight vectors of the nodes to the input vectors. At the end of training, this process generates a learned SOM with each node having an associated weight vector or prototype. Each prototype is associated with a different set of input vectors. This transformation helps in effective visualization and abstraction of high-dimensional data for exploratory data analysis.

B. The SOM Algorithm

The SOM algorithm consists of two major steps: 1) determining the winner node; and 2) updating the weight vectors associated with the winner node and some of its neighboring nodes. Prior to training, the weight vectors associated with each node

of the map are suitably initialized. Random initialization often suffices. For a profitable initialization, the vectors are sampled evenly from the subspace spanned by the two largest principal component eigenvectors [25].

The training expands over several iterations and is based on competitive learning. In each iteration, a vector $\boldsymbol{x} = [x_1, x_2, \ldots, x_n]^T \in \mathbb{R}^n$ (where n is the length of fMRI data) from the input space is compared with the weight vectors of the nodes $\boldsymbol{m}_i = [m_{i1}, m_{i2}, \ldots, m_{in}]^T \in \mathbb{R}^n$ (where $i = 1, 2, \ldots, N$; N being the total number of nodes) to determine the winner node, often referred to as the best matching unit (BMU). The BMU refers to the node whose weight vector is the closest match of the input vector based upon a similarity metric. The most commonly used metric is the Euclidean metric:

$$\|\boldsymbol{x} - \boldsymbol{m}_c\| = \min_i \{\|\boldsymbol{x} - \boldsymbol{m}_i\|\} \quad i = 1, \ldots, N \quad (1)$$

where $\| \ \|$ represents the Euclidean norm, $\boldsymbol{x}$ is the vector under consideration, $\boldsymbol{m}_i$ denotes the weight of the i^{th} node on the map, and $\boldsymbol{m}_c$ represents the weight of the BMU. Once the BMU is determined, the weight vectors associated with the BMU and some of its neighbors in the map are updated to make them more similar to the input vector

$$\boldsymbol{m}_i(t+1) = \boldsymbol{m}_i(t) + h_{ci}(t)[\boldsymbol{x}(t) - \boldsymbol{m}_i(t)] \quad (2)$$

where t is the current iteration number; $h_{ci}(t)$ is defined as the neighborhood kernel that controls the number of neighboring nodes to be updated and the rate of update in each iteration. The magnitude of this update decreases with time (iteration) and for nodes farther away from the BMU with a suitable kernel. In general, the neighborhood kernel takes the form of a Gaussian function:

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

where r_i and r_c are spatial coordinates of the i^{th} node and the winner node, respectively, in the output space; σ is the full-width at half-maximum (FWHM) of the Gaussian kernel that determines the neighboring nodes to be updated. α denotes the learning rate that controls how fast the weights get updated. Both σ and α decrease monotonically with the increase in the learning iteration t.

C. Graph-Based Visualizations of the SOM

We delineated clusters in the output map using a combination of two graph-based visualization techniques that incorporated 1) local density distribution across SOM prototypes (density-based connectivity); and 2) local similarities (correlations) between the prototypes (correlation-based connectivity) [34].

1) Density-Based Connectivity Visualization, CONNDDvis: The density distribution refers to the number of input vectors (BOLD signals in fMRI) included in the receptive fields of the SOM prototypes. The local density distribution within the receptive fields represents the distribution of data with respect to neighboring prototypes and can be visualized with a topology representing weighted Delaunay graph. It is realized by rendering of the connectivity matrix, CONNDD, over the SOM lattice [26]. The existence of an edge between two prototypes

$\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ on the graph indicates that they are neighbors in the input data space and the weight of the connection between them gives its connectivity strength:

$$\text{CONNDD}(i,j) = |\text{RF}_{ij}| + |\text{RF}_{ji}| \quad i, j = 1, 2, \ldots, N \quad (4)$$

where $|\text{RF}_{ij}|$ denotes the number of input vectors in the receptive field of prototype $\boldsymbol{m}_i$ for which $\boldsymbol{m}_j$ is the second BMU ($\boldsymbol{m}_i$ being the first BMU). CONNDD provides finer density distribution by showing the connection between adjacent prototypes in terms of the number of data points for which these prototypes are the top two BMUs. The connectivity strengths can be normalized to 1 by dividing each value with the mean of the strongest connection of each prototype [26]. Any stronger connections will be assigned a value of 1. Outliers can be excluded by discarding connectivity strength smaller than a threshold. One heuristic for the threshold is the mean connectivity strength of the fourth strongest connection (rank 4) of each prototype when the mean strength falls sharply [26]. The normalized connectivity strengths can be effectively visualized in grayscale and binned widths where darker and thicker edges represent stronger connections. The binarized CONNDD is equivalent to the induced Delaunay graph and when rendered over the SOM lattice, it intersects Delaunay triangulation with the data manifold. This captures cluster boundaries and makes the discontinuities with the dataset visible. We refer to the visualization of CONNDD on the SOM lattice by CONNDDvis, analogous to CONNvis in [26]. Previous works have demonstrated that CONNDDvis helps in the extraction of details in the data structure when the data vectors outnumber the SOM prototypes [26].

2) Correlation-Based Connectivity Visualization, CONNC-Cvis: The correlation coefficient matrix, CONNCC, which includes temporal similarities (correlation coefficients) of neighboring prototypes, can be visualized graphically to display local similarities in the prototypes. The weight on the edges between two prototypes gives the measure of their similarity. The visualization of CONNCC on the SOM lattice is termed as CONNCCvis, similar to DISTvis in [26] where local Euclidean distances or dissimilarities were visualized. We used correlation coefficients between prototypes as measures of their local similarities. CONNCCvis highlights prototypes from the output map that have high temporal similarities in their neighborhood. This helps to distinguish signals from noise in fMRI data visualization.

3) Combined Connectivity Visualization, CONNDDCCvis: The visualization obtained from CONNDDvis and CONNC-Cvis can be merged to obtain a combined connectivity visualization that emphasizes delineation of connectivity structures of prototypes representing task-related signals. We refer to this visualization as CONNDDCCvis, which combines CONNDDvis and CONNCCvis by multiplying (element-by-element) the normalized density-based connectivity matrix, CONNDD, with the normalized correlation-based connectivity matrix, CONNCC

$$\text{CONNDDCC}(i,j) = \text{CONNDD}(i,j) \times \text{CONNCC}(i,j). \quad (5)$$

CONNDDCC denotes overall connectivity strength between $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ and includes both density-based connectivity and correlation-based connectivity between prototypes. The
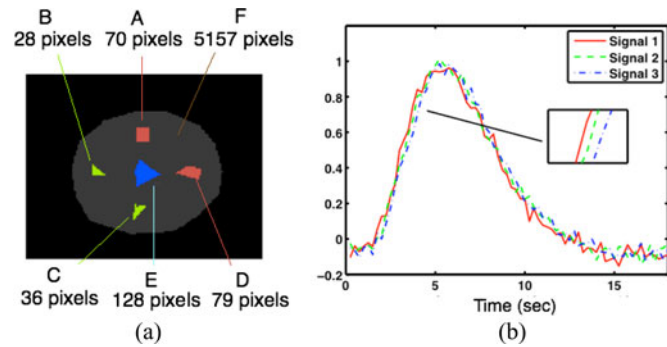
Fig. 1. (a) Synthetic brain slice (128×128) comprising five activated clusters (A–E). The clusters were divided into three groups (indicated in different colors) based on delay between the signals (SNR = 6, TR = 250 ms) they constituted. The signals from red clusters led the signals from green clusters by 100 ms and from blue by 200 ms. F represents gray or white matter constituting noise. (b) Average simulated BOLD responses from red (A and D), green (B and C), and blue (E) clusters in the respective colors after averaging across trials.

resulting visualization suppresses the visualization of noise and delineates detailed connectivity structures of correlated signals [34].

III. MATERIALS AND METHODS

A. Simulated Data

To illustrate the application of the method, we created synthetic fMRI dataset with timing variability. Synthetic BOLD signals simulating an event-related acquisition were created with a stimulus time series comprising a 2-s box car-shaped stimulus followed by a 16 s rest repeated over 17 trials with a total runtime of 306 s. The stimulus time series was convolved with a canonical hemodynamic response based on a gamma variate function with a 100 Hz sampling frequency. It was downsampled to create the synthetic BOLD signal with a 250 ms repetition time (TR). Gaussian noise was added to the resulting signal such that the signal-to-noise ratio (SNR) was 6. The SNR was calculated as the ratio of the amplitude of the signal to the standard deviation of the noise.

Three sets of simulated stimuli, shifted in time to introduce known delays between the simulated BOLD signals, were used. The first set led the second by 100 ms and the third by 200 ms. A synthetic brain slice (128×128) was created to simulate the spatial distribution of activated regions on brain [see Fig. 1(a)]. The first set of simulated BOLD signals was assigned to the red regions (A and D). Regions B and C (green) constituted signals from the second set, and blue pixels were assigned the signals from the third set. Region F represented the gray or white matter in brain and consisted of noise only. BOLD responses from these regions are shown in Fig. 1(b) after averaging across trials.

B. Real FMRI Data

1) Experimental Design: We used the fMRI dataset originally collected for a study of the unified attentional bottleneck [36]. The original experiment was composed of trials that comprised either an auditory-vocal (AV) task, a visual-manual (VM) task, or both tasks simultaneously (dual tasks, DT). Nine
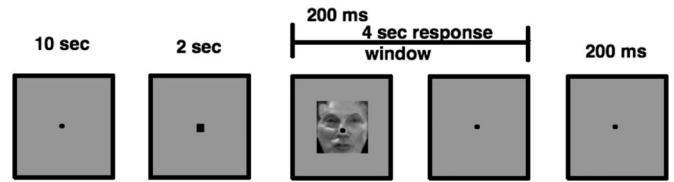

Fig. 2. Experimental design of the face discrimination task comprising visual stimulus-manual response trial (adapted from [36]).

functional runs with six VM trials, six AV, and six DT trials, presented randomly, were acquired. For this work, we extracted only the VM trials while excluding other tasks from our analysis. The VM face discrimination task used one of three male faces (6.4° of visual angle). Each trial comprised 12 s of a black fixation marker (0.1° of visual angle) followed by the stimuli presentation for 200 ms. The last 2 s of the fixation included an enlargement of the fixation point to alert the subjects about the following stimulus presentation (see Fig. 2). With the stimulus presentation, a 4 s response interval was initiated, which was followed by a 200 ms postresponse period before the start of the next trial. Responses were made with right index, middle, and ring fingers for the respective faces.

2) Data Acquisition: The data were collected after approval from the institutional review board at Vanderbilt University [36]. Twelve healthy volunteers (five female; aged 21–33) participated. Functional (T2*) images (TR = 1200 ms, TE = 35 ms, flip angle = 70°, FOV = 220 mm; 64×64 matrix) consisting of 20 slices (4.5 mm thick, 0.5 mm skip) were acquired parallel to the anterior commissure–posterior commissure line on a Philips Achieva 3T Intera MR scanner. Stimuli were presented on a liquid crystal display panel and back-projected onto a screen in front of the scanner. Manual responses were collected with five-key keypads (Rowland Institute of Science). RT was recorded for each trial. Stimulus presentation was synchronized with fMRI volume acquisition with the interleaved design such that the intervening trials began 600 ms (1/2 TR) after the scanner pulse. There were in total nine runs. One subject completed only seven runs due to discomfort.

3) Data Preprocessing: Before analyzing the data for voxel selection, we ran them through a preprocessing pipeline that included motion correction and slice-time correction. We then extracted VM trials (a total of 54) from all runs for a given subject, and using the K-means algorithm, divided the trials based on RTs into two groups (clusters): slow RT and fast RT. If one cluster ended up having just a few trials, we discarded it as outliers and divided the remaining trials into two groups. We concatenated fast and slow RT trials to choose voxels from the combined signals. The concatenated time series was temporally filtered with a 120 s (0.0083 Hz) high-pass filter including detrending to remove low-frequency drifts and linear trends in the data.

C. Voxel Selection

1) SOM: The preprocessed signals were used as inputs to the SOM algorithm to detect voxels responding to the task. The total number of nodes N, initial learning rate α, and number of

iterations for the SOM algorithm were chosen from the test for convergence procedure used in [21]. This would optimize the parameters to achieve the best possible results with minimum iteration and time. An initial learning rate of 0.1 produced the least total squared error between the input data vectors and the representative node prototype vectors for the total number of nodes of 100 (arranged in a 10×10, 2-D lattice grid format) and 100 iterations.

We initialized the weight vectors associated with the nodes with the first two principal components of the input data from the brain region. The winner node (BMU) was selected using the lagged correlation coefficient metric:

$$corr(\boldsymbol{x}, \boldsymbol{m}_c) = \max_i \left\{ lagcorr(\boldsymbol{x}, \boldsymbol{m}_i) \right\}, \quad i = 1, \ldots, N \tag{6}$$

where $lagcorr(\boldsymbol{x}, \boldsymbol{m}_i)$ denotes the lagged correlation coefficient between the input $\boldsymbol{x}$ and the weight vector of the i^{th} node, $\boldsymbol{m}_i$. We used lag-1 correlation in the simulation so it would improve the sensitivity to timing differences between signals. Lag-0 versus lag-1 did not make a difference in the real fMRI data so we used lag-0. The initial value of the FWHM of the Gaussian kernel (σ) in the neighborhood function was set to be equal to the radius of the lattice, equal to seven nodes for the 10×10 lattice [21], [22]. The learning rate α was set to an initial value of 0.1. Both learning rate (α) and FWHM of the Gaussian kernel (σ) were decreased exponentially with the increase in the learning iteration. The training resulted in a 10×10 map of output nodes with a prototype and a voxel map for each node. Using our visualization scheme, we delineated clusters on the map. For the fMRI experiment, the clusters whose voxels mapped to the visual and motor regions were chosen for further analysis.

2) ICA: We used GIFT (Medical Image Analysis Lab), which implements spatial ICA, to extract task-related signals from the fMRI dataset. We ran the ICA for each individual subject using the FastICA algorithm incorporated in GIFT. The number of components was determined by the minimum description length principle. We examined the spatial map for each component and selected the one whose voxel map included the visual and motor regions. Brain signals from the regions were extracted using a suitable threshold to match the number of voxels obtained from the SOM.

3) GLM: We performed GLM-based multiple regression analysis using statistical parametric mapping (SPM8— http://www.fil.ion.ucl.ac.uk/spm/software). The regressor was constructed by convolving the event related stimulus time series with a canonical hemodynamic response based on gamma variate functions. GLM was fitted to the response using SPM8 and regression parameter was estimated. A suitable threshold for the t statistic was chosen so the number of voxels from visual and the motor regions matched with other methods.

D. Measuring Timing Differences Using an IL Model

Inverse logit functions were used to model the hemodynamic response. We fitted the model on average BOLD responses from visual and left motor regions for two RT groups and estimated their timing parameters.

To model the HRF, a superposition of three IL (sigmoid) functions, $L(x) = (1 + e^{-x})^{-1}$, was used. The first function modeled the rise after activation, the second modeled the subsequent fall and undershoot, and the third function modeled the stabilization or return to baseline [37]. The model of the HRF was given by:

$$h(t \mid \theta) = \alpha_1 L((t - T_1)/D_1) + \alpha_2 L((t - T_2)/D_2)$$
$$+ \alpha_3 L((t - T_3)/D_3). \tag{7}$$

Each function had three variable parameters representing the amplitude, position, and slope of the response. The α_i parameter controlled the amplitude and direction of the curve, T_i controlled the position, and D_i controlled the angle of the slope of the curve. We constrained the values of α_2 and α_3 (so that the fitted response begins at zero at the time point $t = 0$ and ends at magnitude 0) and used a four parameter model where the position of each function and the total amplitude were allowed to vary. We used following constraints for the amplitude:

$$\alpha_2 = \alpha_1 (L((-T_3)/D_3) - L((-T_1)/D_1))$$
$$/(L((-T_3)/D_3) + L((-T_2)/D_2)). \tag{8}$$

And

$$\alpha_3 = |\alpha_2| - |\alpha_1|. \tag{9}$$

We used a gradient descent solution to fit the model and used the parameter estimation procedure described in [37] to calculate the height H, time-to-peak T, and FWHM W from fitted HRF estimates. The difference in time-to-peak (TTPD) between average motor and visual responses was used to measure RTs.

IV. RESULTS

A. Simulated Data Analysis

Signals from synthetic brain regions, denoted by the inner gray area comprising five activated regions (A–E) in Fig. 1(a), were fed as input to the SOM algorithm. There were 5498 total brain pixels with 341 ($\sim$6%) that constituted synthetic BOLD signals and rest noise. Fig. 3 shows the 10×10 SOM output map showing traces of prototypes (average across trials) of the nodes. The prototypes corresponding to task are on the top right-hand corner of the map. We then applied the graphical visualization scheme to delineate clusters on the output map. Fig. 4(a) shows CONNDDvis obtained by draping the local density distribution matrix CONNDD over the SOM lattice. The strength of the connection between prototypes is indicated in grayscale and binned width of the lines where darker and wider lines represent stronger connections. Outliers and connections involving noise were removed by discarding connectivity strength smaller than the mean connectivity strength of the fifth strongest connection of each prototype when the mean strength fell sharply. A number of connections appeared between the nodes including three distinct clusters at the top right-hand corner of the map. Fig. 4(b) shows CONNCCvis obtained by draping the CONNCC matrix over the SOM lattice. The prototypes at the top right-hand corner, possibly including task-related signals,
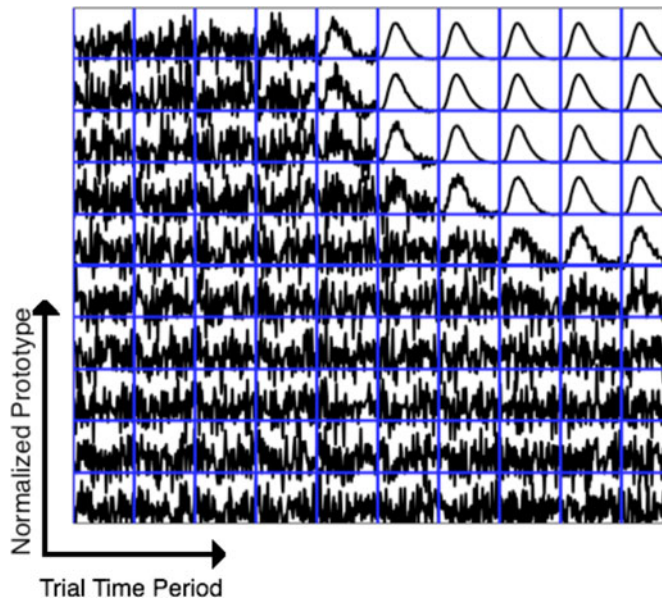
Fig. 3. 10×10 matrix of SOM output node prototypes (averaged across trials) from the simulated dataset. The prototypes corresponding to task are on the right-hand corner of the map.
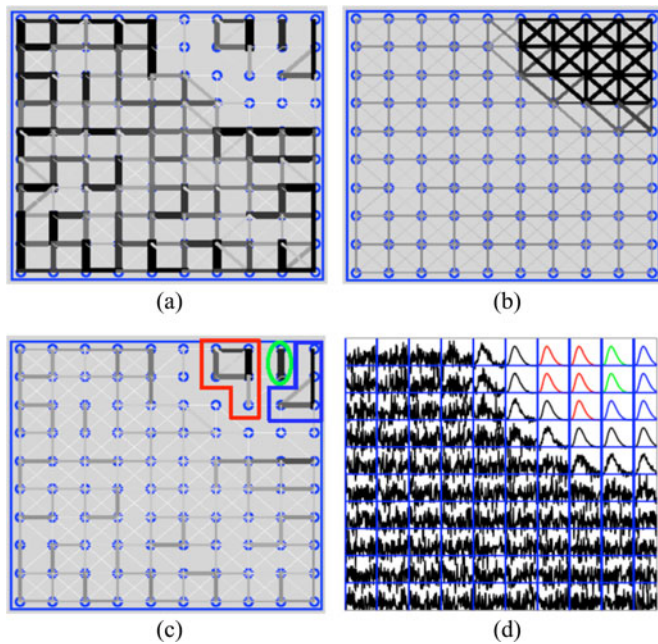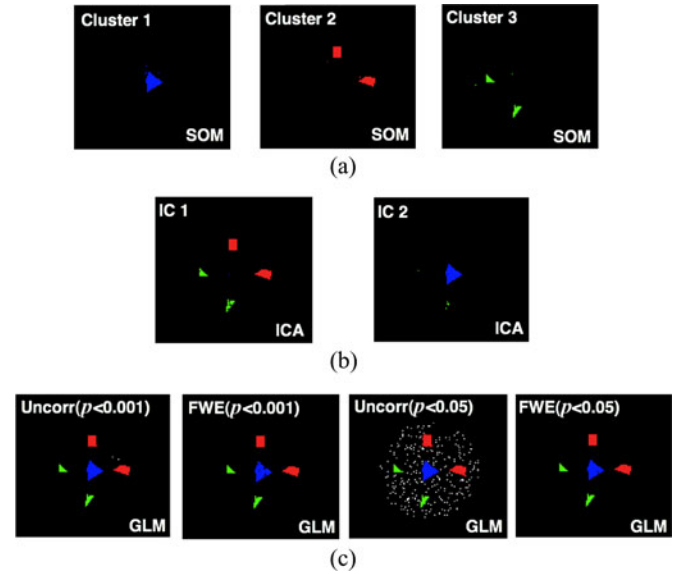

Fig. 5. (a) The graph-based visualizations of the SOM successfully detected three clusters based on the delay groups: 0, 100, and 200 ms. (a) The blue cluster in Fig. 4(d) included voxels shown in left (Cluster 1), red cluster included voxels shown in middle (Cluster 2), and green cluster included voxels shown in right (Cluster 3). (b) ICA detected two components that could not distinguish the delay groups separately, and (c) GLM analysis with uncorrected and FWE corrected p values ($p < 0.001$ and $p < 0.05$) could not distinguish the delay groups either.

Fig. 4. (a) Density-based connectivity visualization (CONNDDvis): Visualization of node connectivity based on local density distribution on the 10×10 SOM lattice. Connectivity is interpreted in grayscale where darker and wider lines mean strong connections. (b) Correlation-based connectivity visualization (CONNCCvis): Visualization of connectivity based on local correlation (correlation coefficient between the neighboring prototypes) on the SOM lattice. (c) Combined connectivity visualization (CONNDDCCvis): Visualization of connectivity based on local density distribution and local correlation on the SOM lattice. Three clusters of nodes were identified as indicated in red, green, and blue. (d) The output map showing traces of prototypes in different colors for three different clusters.

showed high correlations. Fig. 4(c) shows CONNDDCCvis obtained by draping the density and correlation combined matrix CONNDDCC over the lattice. The three distinct clusters at the top right-hand corner of the map became dominant while other connections including noise were suppressed. The same clusters are shown with the node prototypes in Fig. 4(d) and are mapped to the image space in Fig. 5(a).

The pixels included in the prototypes indicated in blue in Fig. 4(c) are shown in Fig. 5(a), left. These represented blue pixels in the image space. Similarly, the prototypes shown in red in Fig. 4(c) included red pixels [see Fig. 5(a), middle] and green prototypes included green pixels as shown in Fig. 5(a), right. This shows the ability of the graph-based visualization schemes of the SOM in 1) detecting voxels responding to the task; and 2) distinguishing voxels on the basis of timings of the signals they constituted.

We compared the performance of the SOM with ICA and GLM. The ICA produced two independent task-related components from the simulated dataset but could not distinguish the three timing groups [see Fig. 5(b)]. The results did not change even when the number of components to be determined was set to more than 3. GLM with uncorrected or FWE corrected p values also detected the activated voxels but could not distinguish the delays in them [see Fig. 5(c)]. In summary, the SOM with graph-based visualizations provided highest sensitivity in detecting voxels responding to tasks and distinguishing them based on timings of the corresponding signals in the simulated data.

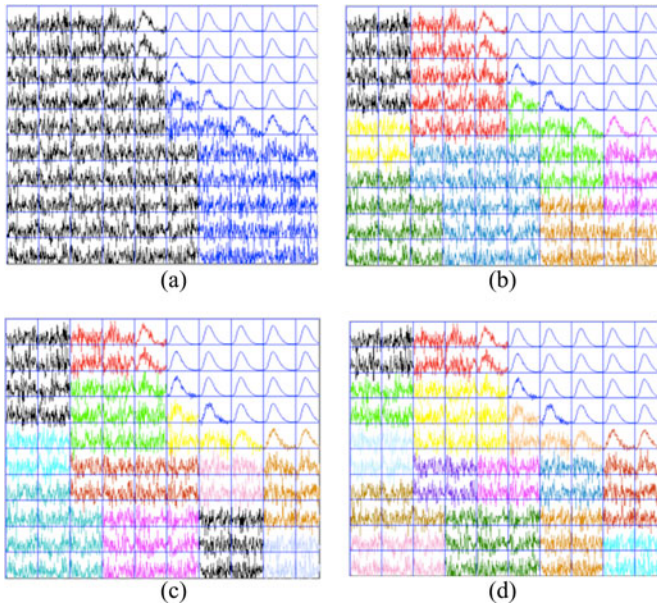We compared the performance of the visualization technique with the two-stage clustering approach where SOM prototypes

Fig. 6.   Clustering of SOM prototypes using hierarchical clustering with different "optimal" partitions (Np). The colors represent clusters. The prototypes corresponding to signals at the top right-hand corner all belonged to one cluster irrespective of the number of partitions. This indicates the inability of the two-stage clustering approach using hierarchical clustering of the SOM to distinguish delays in the signals. (a) $N_p = 2$. (b) $N_p = 9$. (c) $N_p = 13$. (d) $N_p = 16$.
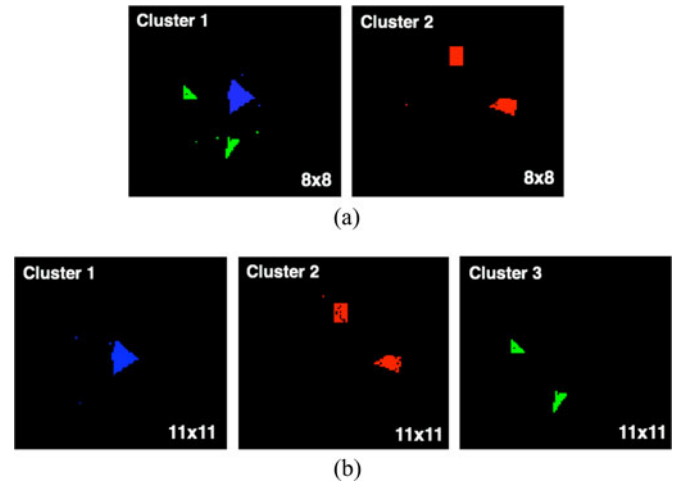

(a)

(b)

Fig. 7.   Effect of map size (a) 8×8 map size could not distinguish the three delay groups: 0, 100, and 200 ms separately. (b) 11×11 size successfully delineated the three delay groups in three distinct clusters.
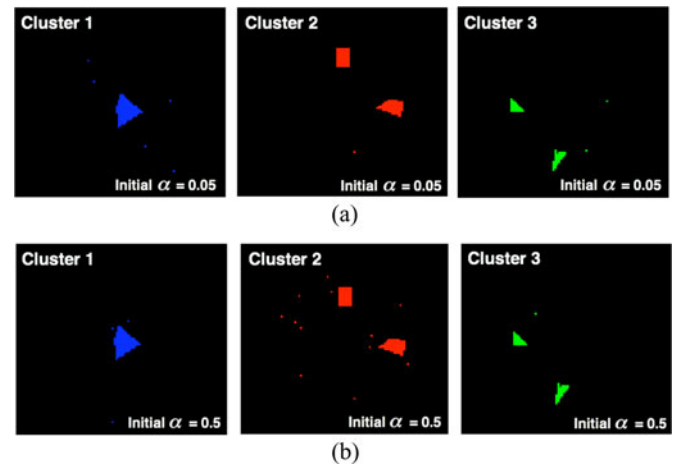

(a)

(b)

Fig. 8.   Effect of initial learning rate. (a) Slower initial learning rate, $\alpha = 0.05$, which was half the original, could still classify the three delay groups: 0, 100, and 200 ms. (b) Faster learning rate, $\alpha = 0.5$, could also delineate the three delay groups. However, there were more false positives in the results.

from the first stage were clustered using hierarchical clustering in the second stage [22]. The quality of the partition was assessed by the within-class inertia measure [11]. We clustered the SOM prototypes (in Fig. 3) with the hierarchical clustering approach using optimal partitions that resulted in 2, 9, 13, and 16 optimal numbers of clusters (see Fig. 6). None of these partitions were able to distinguish the delay groups in the signals. The prototypes related to signals belonged to one cluster in all partitions. We increased the cluster number further, but the signal-related prototypes always remained in one cluster.

Although the learning parameters and size of the SOM were optimized using the test for convergence procedure described in [21], we tweaked these parameters to see their effects on the result. Fig. 7 shows the effect of map size. For the same set of delays (0, 100, and 200 ms) and SNR level 6, a smaller map size (8×8) delineated only two groups of delays in the signals [see Fig. 7(a)]. Map size 9×9 (not shown here) could still delineate the three delay groups. A larger map size (11×11) could distinguish the delays resulting in three clusters as shown in Fig. 7(b), however, took longer to converge due to increase in the number of output nodes. We also tested the effect of the learning rate (see Fig. 8). For slower initial learning rate ($\alpha = 0.05$, half the original), the three delay groups were still clearly delineated. These clusters became apparent for faster learning rate ($\alpha = 0.5$) as well. However, there were more false positives in the results than those at $\alpha = 0.1$ [see Fig. 5(a)]. The initial learning rate did not have much impact on the results when varied from 0.05 to 0.5. However, it should be remembered that if the learning rate is too slow or too fast than the optimum, the results may not converge or in this study, it may fail to clearly delineate the clusters.

We tested the robustness of the method by changing the delays and SNR of the signals. For the same SNR level 6 but shorter delays (0, 50, and 100 ms), the SOM visualization could not distinguish the delays (not shown here) while longer delays (0, 200, and 400 ms) resulted in distinct clusters as expected (see Fig. 9). However, ICA still resulted in only two components, and GLM could not distinguish these longer delays in the signals [see Fig. 9]. For lower SNR of level 4, SOM visualization gave only two clusters for delays: 0, 100, and 200 ms [see Fig. 10]. However, longer delays (0, 200, and 400 ms) were clearly separable even at SNR level 4. When SNR increased beyond 6, even shorter delays started to become distinct. Fig. 11 shows results for delays 0, 50, and 100 ms between signals at SNR level 7. While SOM visualization could delineate two clusters, ICA and GLM could not distinguish the delays at all. In summary, our results showed that longer delays and higher SNR were favorable for clear delineation of timing differences in signals.
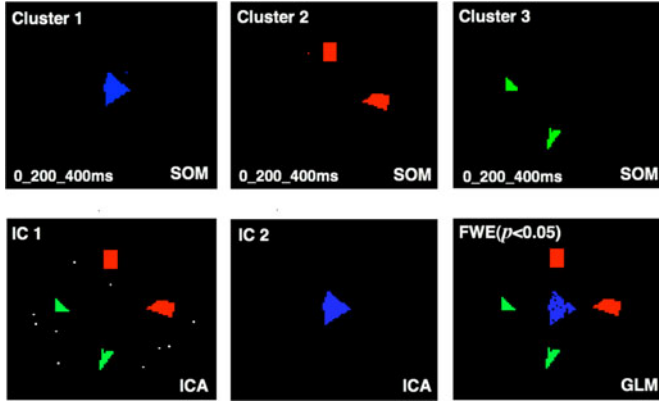
Fig. 9. Effect of delays. Larger delays between the signals (0, 200, and 400 ms) at the same SNR = 6 improved the clustering by resulting in fewer false positives than at shorter delays [0, 100, and 200 ms—Fig 5(a)]. However, the results from ICA and GLM were not much different.
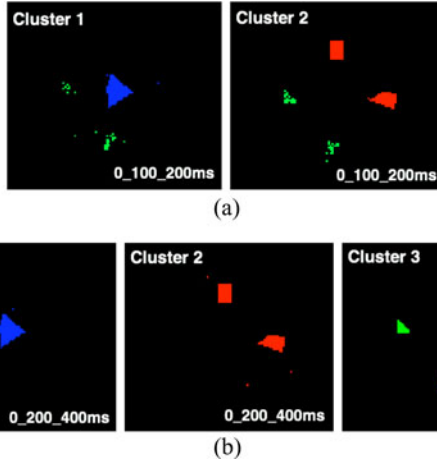

Fig. 10. Effect of SNR. (a) At a lower SNR = 4, the resulting clusters could not distinguish 0, 100, and 200 ms from each other. (b) For larger delays (0, 200, and 400 ms), three distinct clusters resulted for the delay groups.

Visualization of the SOM was the most sensitive and capable of delineating small timing differences compared to ICA or GLM.

B. FMRI Data Analysis: RT Study

The extracted functional time series data from brain regions was fed as input to the SOM algorithm. The density-based connectivity (CONNDDvis) showed a number of clusters at different regions of the map [see Fig. 12(a)] for a subject. The correlation-based connectivity showed strong and weak connections on the map as well [see Fig. 12(b)]. The combined connectivity delineated clusters that could be relevant to our analysis [see Fig. 12(c)]. A number of (about 12) clusters became apparent on the map. The clusters indicated in blue and red included prototypes whose voxels mapped to visual and left motor regions, respectively, in Fig. 13(a). Average signals from these regions were extracted and split into fast and slow RT groups/trials. The IL hemodynamic model was fit on the average signals for both groups. Fig. 14 shows average signals from (a) visual and (b) motor regions for the subject and their corresponding model fits. The RT difference is quite evident from
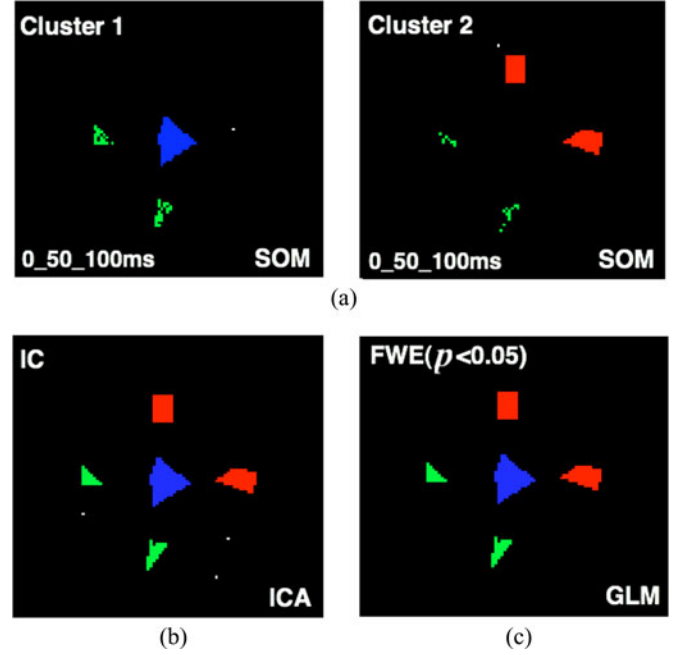

Fig. 11. Effect of SNR. (a) At a higher SNR = 7, two clusters resulted for delay groups: 0, 50, and 100 ms. (b) ICA and (c) GLM could not distinguish the delays at all.

the motor signals [see Fig. 14(b)]. The estimated HRFs for the two regions and RT groups for the same subject are shown in Fig. 15. Timing differences between visual and motor functions and slow and fast RT functions are evident. We estimated time-to-peak parameters of each signal. The TTPDs between visual and motor signals for the two RT groups were compared with the mean of the recorded RTs. The TTPDs increased (from 1.52 to 1.8 s) with the increase in RT (from 0.8 to 1.19 s). This is indicative of the ability of average signals from visual and left motor cortices to reflect the RT differences.

We also chose voxels using ICA [see Fig. 13(b)] and GLM [see Fig. 13(c)] and computed the time-to-peak measures from average responses for slow and fast timing groups. The TTPD between motor and visual cortices for both timing groups were compared with the corresponding RTs. Fig. 16 shows relationships between the computed TTPD measures and the mean RTs for 12 subjects for all three voxel selection techniques. The linear relationships between TTPD and RT were strongest for the SOM [see Fig. 16(c)] as indicated by the linear mixed-effects model (dark line, $p = 0.007$). For voxels selected by ICA, the TTPD increased with RT in some subjects while decreased for others [see Fig. 16(b)]. The mixed effects model indicated weaker linear relationship (dark line, $p = 0.06$). For GLM [see Fig. 16(a)], the linear relationship was not apparent (dark line, $p = 0.82$).

C. Performance Comparison

We fit a linear mixed-effects model on the TTPD measures and computed slope and intercept of the fit for each voxel selection method. The t statistic for the slope of the model was calculated to compare the precision of each method in terms of the standardized effect size. The precision compares the strength of
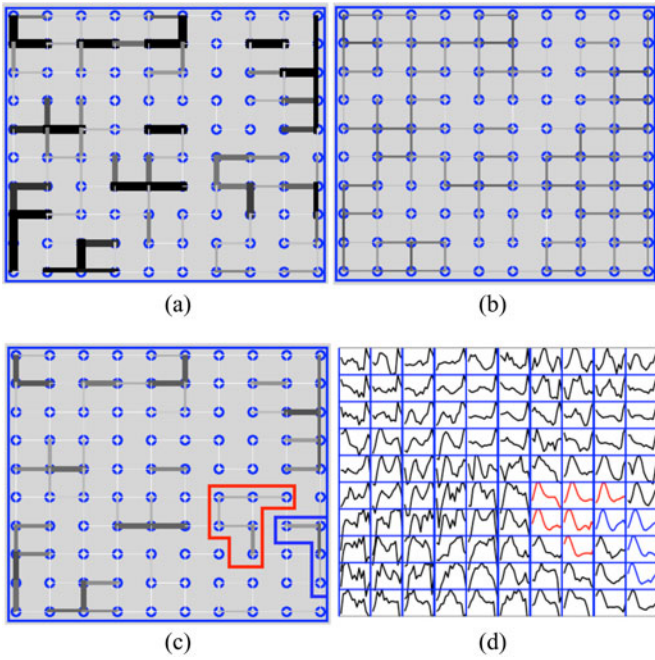
Fig. 12. (a) CONNDDvis: Density-based connectivity visualization from a subject performing the visual stimulus-manual response (VM) task. (b) CON-NCCvis: Correlation-based connectivity visualization: (c) CONNDDCCvis: Combined connectivity visualization. A number of node clusters were identified including those shown in red and blue. (d) The same clusters shown on the output map with traces of prototypes. The prototypes indicated in blue and red contained voxels in visual and motor regions, respectively [see Fig. 13(a)].
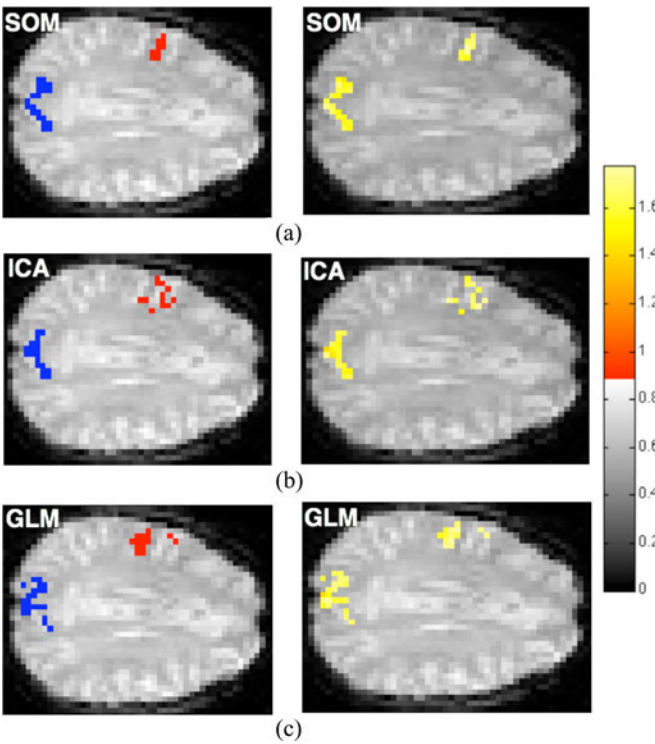

Fig. 13. Voxels identified by (a) SOM (Voxel count: Visual—21, Motor—16), (b) ICA (Voxel count: Visual—18, Motor—12), and (c) GLM (Voxel count: Visual—22, Motor –14) for the subject. The same voxels are shown with the colormap (based on correlation with the reference function) on the right.
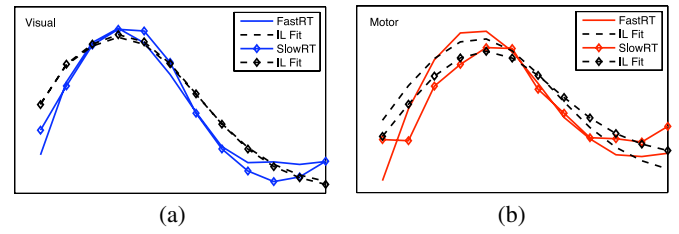

Fig. 14. Average signals from (a) visual and (b) motor regions (chosen via SOM) for the subject for fast and slow RT groups. The broken lines are average fits from the IL model. The timing difference between fast and slow RT groups is evident from the motor signals.
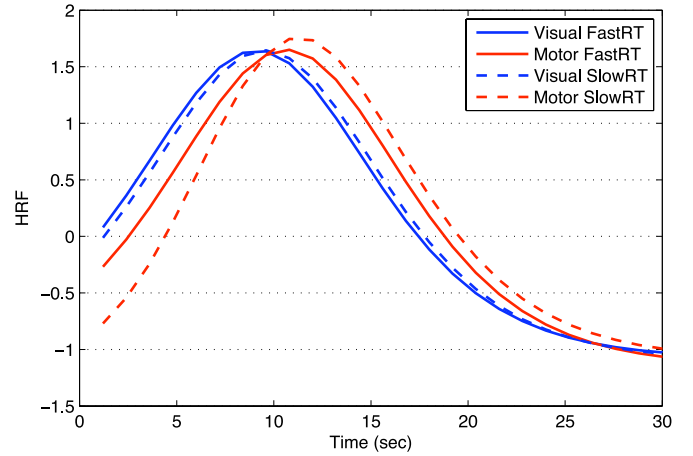

Fig. 15. Estimated HRFs from visual (blue) and motor (red) regions (chosen via SOM) for two RT groups for the subject. Motor–Visual (solid curves) time-to-peak difference: $TTPD_1 = 1.52$ s. Motor–Visual (broken curves) time-to-peak difference: $TTPD_2 = 1.8$ s. Actual mean RTs were: $RT_1 = 0.80$ s and $RT_2 = 1.19$ s.
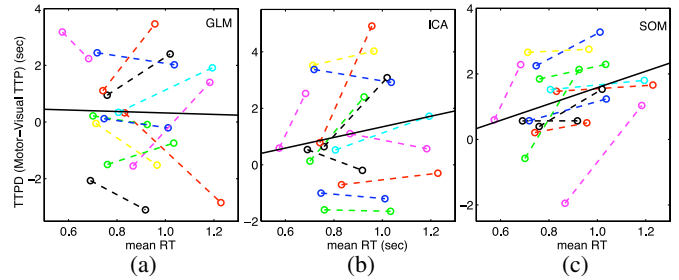

Fig. 16. TTPD between average motor and average visual cortex signals obtained via (a) GLM, (b) ICA, and (c) SOM versus mean RTs from 12 subjects. TTPD increased linearly with RT in most of the subjects for voxels selected by the SOM as indicated by the linear mixed-effects model (dark line, $p = 0.007$). For voxels selected by ICA, the TTPD increased with RT in some subjects while decreased for others (dark line, $p = 0.06$). For GLM, a linear relationship was not apparent (dark line, $p = 0.82$).

the linear relationship between TTPD and RT against the amount of variability in the data. Similarly, the t statistic for the intercept was measured to compare the bias of the fits. Fig. 17 shows (a) precision and (b) bias of the fits on the measures from the SOM, ICA, and GLM. The error bars indicate 95% confidence intervals from 2000 bootstrap samples using the case-resampling bootstrap. Signals identified by the SOM produced the highest precision of all. No bias was evident statistically for any of the methods [see Fig. 17(b)].
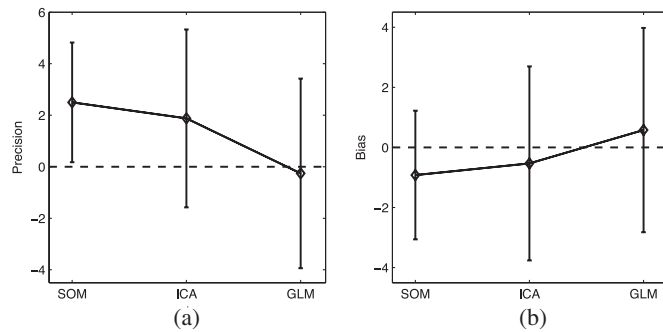
Fig. 17. Performance comparison of the voxel selection methods using results from linear mixed-effects modeling. (a) Precision was measured with the t statistic for the slope (linear term of the fit) describing the relationship to RT. (b) Bias was measured with the t statistic for the intercept of the fit. The error bars indicate 95% confidence intervals from 2000 bootstrap samples generated by case-resampling bootstrap.

V. DISCUSSION AND CONCLUSION

In this paper, we proposed a data-driven method for fMRI analysis using a novel graph-based visualization scheme for self-organizing maps (SOMs). The visualization scheme was formulated using two metrics of SOM node connectivity based on 1) density distribution across nodes; and 2) correlation between the node prototypes of the SOM output. The correlation-based connectivity (CONNCC) considered relationships between the neighboring prototypes in terms of their correlation. This helped to segregate signal from noise, because prototypes containing noise were less likely to be correlated. If the signals were highly correlated and had small timing differences, CONCC might not reveal the differences because the within-cluster similarity and between-cluster similarity among these signals would be comparable. Hence, all signals would seem to fall in one cluster. The density-based connectivity (CONNDD) on the other hand did not directly use correlation between prototypes. It considered relationships between the neighboring prototypes in terms of data distribution across the prototypes. The number of data vectors (fMRI time series) for which the prototypes were the best two matching prototypes accounted for connectivity between the neighboring prototypes in CONNDD. Separation between clusters could be seen when the neighboring prototypes did not have data distribution. This finer density distribution visualization helped to distinguish small timing differences within the signals. This visualization scheme was applied on simulated and fMRI datasets in this study to identify activated voxels and classify them based on differences in timings of their corresponding signals. The SOM applied in conjunction with the graph-based visualization scheme outperformed ICA and GLM by providing highest sensitivity in classifying regions based on the timing of their responses. On the real fMRI dataset, we fit the average signals using IL functions to model the hemodynamic response and measured TTPD measures to compare to RTs. Signals identified by the SOM suggested a statistically strong linear relationship between RTs and TTPDs.

A. Significance of Graph-Based Visualizations of SOMs in FMRI-Based Timing Studies

We used three methods for voxel selection in this study: GLM, ICA, and SOM. GLM was a univariate method that followed a hypothesis-driven approach whereas ICA and SOM were multivariate methods that followed a data-driven approach.

GLM followed voxel-based linear regression analysis based on the GLM using SPM. The modeling assumptions and deterministic characteristics of the reference function in GLM analysis may have been too restrictive to capture the full range of stimulus-driven BOLD transitions [22]. Additionally, the voxel-by-voxel hypothesis testing approach in SPM was massively univariate, not making full use of the neighborhood relationships or intervoxel dependencies that could improve sensitivity. This may have resulted in lower sensitivity of GLM in identifying and classifying task-evoked voxel responses based on their timings.

For ICA, we used spatial ICA that determined maximally independent components from the data by maximizing non-Gaussianity. This ensured segregation of task-related signals from other nonrelevant signals and noise. Although data driven, ICA made a strong assumption of independence between components which may have resulted in biased decomposition and decreased ability to detect task-related signals [10], [38]. This may have resulted in lower sensitivity of ICA as shown by the results (see Fig. 16). Our implementation of SOM used neighborhood correlation to group data based on similarities in their temporal patterns. The neighborhood function essentially made use of intervoxel relationships that increased its sensitivity. The graph-based visualization scheme delineated small structures of clusters in the data that could distinguish early BOLD signals from late signals. This was demonstrated by our results from simulation. This may have helped in reducing interference from late signals coming through draining veins in the real fMRI experiment.

A weakness of the graph-based visualizations of the SOM is that connections between nonneighboring SOM prototypes can obscure the visualization, especially when the dataset is noisy like a typical fMRI dataset. Using a suitable threshold to discard noise and outliers is important for clear distinction and visualization of meaningful clusters.

B. Time-to-Peak From IL Fits for Timing Studies

A comparative study of the hemodynamic response models conducted in [37] suggested that the IL model was immune to a large degree of model misspecification providing the least amount of bias and confusion between the response parameters than other models such as canonical gamma and finite impulse response. In [39], hemodynamic response timing parameters were investigated concluding that the time-to-peak estimate was stable across separate datasets for the same region within a subject and was a reliable measure of the hemodynamic response.

Our estimates of TTPDs parameters obtained from IL fits did not match the corresponding mean RTs in absolute terms [see Fig. 15] although their linear relationship was evident [see Fig. 16(c)]. This was understandable as the HRF is a complex,

nonlinear function of the neuronal or vascular changes. The HRF model is limited in terms of its statistical accuracy for accurate recovery of the true response parameters (time-to-peak, width and height) [37].

We observed large variation in TTPD measurements in some subjects across methods [see Fig. 16]. This could be due to small difference between slow and fast mean RTs for these subjects. In this case, even a slight difference in the voxels selected by the methods could result in large difference in timing measurements. Also, inclusion of voxels contributing to late signals from sources such as draining veins can introduce bias to timing measurements. Our approach of voxel selection using graph-based visualizations of the SOM was more sensitive to timing differences than other methods. It might have been able to avoid voxels contributing to late signals thereby reducing their interference on these timing measurements.

A major source of bias may arise in fMRI-based timing studies in the form of hemodynamic variability. The hemodynamic shape may vary across brain regions and subjects [40], [41] which could give misleading inferences on the actual delay. In this case, modulation of the delay (in our case RT differences) by experimental demands and cognitive context can be employed to rule out hemodynamics as the cause of the results and validate the measured timing differences. The linear relationship between TTPD measures and RTs, especially with signals identified by the SOM in our study, is compelling evidence that these measurements are not source of hemodynamic artifacts.

In summary, with careful consideration of experimental design, graph-based visualizations of SOMs can be used to identify spatiotemporal patterns of activated brain regions and measure relative timings of brain activities using fMRI.

ACKNOWLEDGMENT

REFERENCES

[1] K. J. Friston, A. P. Holmes, K. J. Worsley, J. –P. Poline, C. D. Frith, and R. S. J. Frackowiak, "Statistical parametric maps in functional imaging: A general linear approach," *Human Brain Mapping*, vol. 2, no. 4, pp. 189–210, 1994.

[2] X. Descombes, F. Kruggel, and D. Y. von Cramon, "Spatio-temporal fMRI analysis using Markov random fields," *IEEE Trans. Med. Imag.*, vol. 17, no. 6, pp. 1028–1039, Dec. 1998.

[3] M. Svensén, F. Kruggel, and D. Y. von Cramen, "Probabilistic modeling of single trial fMRI data," *IEEE Trans. Med. Imag.*, vol. 19, no. 1, pp. 25–36, Jan. 2000.

[4] S. Faisan, L. Thoraval, J. P. Armspach, J. R. Foucher, M. N. Metz-Lutz, and F. Heitz, "Hidden Markov event sequence models: Toward unsupervised functional MRI brain," *Acad. Radiol.*, vol. 12, no. 1, pp. 25–36, Jan. 2005.

[5] E. Zarahn, G. K. Aguirre, and M. D'Esposito, "Empirical analyses of BOLD fMRI statistics," *Neuroimage*, vol. 5, no. 3, pp. 179–197, Apr. 1997.

[6] L. K. Hansen, J. Larsen, F. A. Nielsen, S. C. Strother, E. Rostrup, R. Savoy, N. Lange, J. Sidtis, C. Svarer, and O. B. Paulson, "Generalizable patterns in neuroimaging: How many principal components," *Neuroimage*, vol. 9, no. 5, pp. 534–544, May 1999.

[7] T. H. Le and X. Hu, "Potential pitfalls of principal component analysis in fMRI," presented at the Int. Soc. Mag. Reson. Med. 3, Nice, France, 1995, p. 820.

[8] V. D. Calhoun, T. Adali, G. D. Pearlson, and J. J. Pekar, "Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms," *Human Brain Mapping*, vol. 13, no. 1, pp. 43–53, May 2001.

[9] M. J. McKeown, S. Makeig, G. G. Brown, T. P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski, "Analysis of fMRI data by blind separation into independent spatial components," *Human Brain Mapping*, vol. 6, no. 3, pp. 160–188, 1998.

[10] I. Daubechies, E. Roussos, S. Takerkart, M. Benharrosh, C. Golden, K. D. Ardenne, W. Richter, J. D. Cohen, and J. Haxby, "Independent component analysis for brain fMRI does not select for independence," *Proc. Nat. Acad. Sci.*, vol. 106, no. 26, pp. 10415–10422, Jun. 2009.

[11] C. Goutte, P. Toft, E. Rostrup, F. A. Nielsen, and L.K. Hansen, "On clustering fMRI time series," *Neuroimage*, vol. 9, no. 3, pp. 298–310, Mar. 1999.

[12] G. Scarth, M. McIntrye, B. Wowk, and R. L. Somorjai, "Detection novelty in functional images using fuzzy clustering," presented at the Proc. SMR 3rd Annu. Meeting, Nice, France, 1995, p. 238.

[13] L. He and I. R. Greenshields, "An MRF spatial fuzzy clustering method for fMRI SPMs," *Biomed. Signal Process. Control*, vol. 3, no. 4, pp. 327–333, Oct. 2008.

[14] X. Ding, J. Tkach, P. Ruggieri, and T. Masaryk, "Analysis of time course functional MRI data with clustering method without the use of reference signal," presented at the Proc. Int. Soc. Mag. Reson. Med. 17, San Francisco, CA, USA, 1995, p. 630.

[15] H. Cordes, V. Haughton, J. D. Carew, K. Arfanakis, and K. Maravilla, "Hierarchical clustering to measure connectivity in fMRI resting state-data," *Magn. Reson. Imag.*, vol. 20, no. 4, pp. 305–317, May 2002.

[16] H. Chen, H. Yuan, D. Yao, L. Chen, and W. Chen, "An integrated neighborhood correlation and hierarchical clustering approach of functional MRI," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 3, pp. 452–458, Mar. 2006.

[17] A. Wismüller, A. Meyer-Bäse, O. Lange, D. Auer, M. F. Reiser, and D. Sumners, "Model-free functional MRI analysis based on unsupervised learning," *J. Biomed. Inf.*, vol. 37, no. 1, pp. 10–18, Feb. 2004.

[18] H. Fischer and J. Hennig, "Neural network-based analysis of MR time series," *Magn. Reson. Med.*, vol. 41, no. 1, pp. 124–131, Jan. 1999.

[19] S. -C. Ngan and X. Hu, "Analysis of functional magnetic resonance imaging data using self-organizing map with spatial connectivity," *Magn. Reson. Med.*, vol. 41, no. 5, pp. 939–946, May 1999.

[20] K. H. Chuang, M. H. Chiu, C.C. Lin, and J. H. Chen, "Model-free functional MRI analysis using Kohonen clustering neural network and fuzzy c-means," *IEEE Trans. Med. Imag.*, vol. 28, no. 12, pp. 1117–1128, Dec. 1999.

[21] S. J. Peltier, T. A. Polk, and D. C. Noll, "Detecting low-frequency functional connectivity in fMRI using a self-organizing map (SOM) algorithm," *Human Brain Mapping*, vol. 20, no. 4, pp. 220–226, Aug. 2003.

[22] W. Liao, H. Chen, Q. Yang, and X. Lei, "Analysis of fMRI data using improved self-organizing map and spatio-temporal metric hierarchical clustering," *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1472–1483, Oct. 2008.

[23] S. -C. Ngan, E. S. Yacoub, W. F. Auffermann, and X. Hu, "Node merging in Kohonen's self-organizing mapping of fMRI data," *Artif. Intell. Med.*, vol. 25, no. 1, pp. 19–33, May 2002.

[24] T. Kohonen, "The self-organizing map," *Proc. IEEE*, vol. 78, no. 9, pp. 1464–1480, Sep. 1990.

[25] T. Kohonen and O. Simula, "Engineering applications of the self-organizing map," *Proc. IEEE*, vol. 84, no. 10, pp. 1358–1384, Oct. 1996.

[26] K. Taşdemir, "Graph based representations of density distribution and distances for self-organizing maps," *IEEE Trans. Neural Netw.*, vol. 21, no. 3, pp. 520–526, Mar. 2010.

[27] A. Ultsch, "Self-organizing neural networks for visualization and classification," in *Information and Classification-Concepts Methods and Applications*, O. B. Lausen and R. Klar, Eds.   Berlin, Germany: Springer-Verlag, 1993, pp. 307–313.

[28] A. Ultsch, "Maps for the visualization of high-dimensional data spaces," in *Proc. 4th Workshop Self-Org. Maps*, 2003, vol. 3, pp. 225–230.

[29] M. Kraaijveld, J. Mao, and A. Jain, "A nonlinear projection method based on Kohonen's topology preserving maps," *IEEE Trans. Neural Netw*, vol. 6, no. 3, pp. 548–559, May 1995.

[30] H. Yin, "ViSOM-a novel method for multivariate data projection and structure visualization," *IEEE Trans. Neural Netw*, vol. 13, no. 1, pp. 237–243, Jan. 2002.

[31] M.-C. Su and H.-T. Chang, "A new model of self-organizing neural networks and its applications," *IEEE Trans. Neural Netw.*, vol. 12, no. 1, pp. 153–158, Jan. 2001.

[32] K. Taşdemir and M. Erzsébet, "Exploiting data topology in visualization and clustering of self-organizing maps," *IEEE Trans. Neural Netw.*, vol. 20, no. 4, pp. 549–562, Apr. 2009.

[33] J. Vesanto, "SOM-based data visualization methods," *Intell. Data Anal.*, vol. 3, no. 2, pp. 111–126, Aug. 1999.

[34] S. B. Katwal, J. C. Gore, and B. P. Rogers, "Analyzing fMRI data with graph-based visualizations of self-organizing maps," in *Proc. IEEE Int. Symp. Biomed. Imag.*, Mar./Apr. 2011, pp. 1577–1580.

[35] S. B. Katwal, J. C. Gore, J. C. Gatenby, and B. P. Rogers, "Measuring relative timings of brain activities using fMRI," *Neuroimage*, vol. 66, no. 1, pp. 436–448, Feb. 2013.

[36] M. N. Tombu, C. L. Asplund, P. E. Dux, D. Godwin, J. W. Martin, and R. Marois, "A unified attentional bottleneck in the human brain," *Proc. Nat. Acad. Sci. USA*, vol. 108, no. 33, pp. 13426–13431, Aug. 2011.

[37] M. A. Lindquist, J. M. Loh, L. Y. Atlas, and T. D. Wager, "Modeling the hemodynamic response function in fMRI: Efficiency, bias and mismodeling," *Neuroimage*, vol. 45, no. 1, pp. S187–S198, Mar. 2009.

[38] A. Meyer-Bäse, O. Lange, and A. Wismüller, "Comparison of two exploratory data analysis methods for fMRI: Unsupervised clustering versus independent component analysis," *IEEE Trans. Inf. Technol. Biomed.*, vol. 8, no. 3, pp. 387–398, Sep. 2004.

[39] F. M. Miezin, L. Maccotta, J. M. Ollinger, S. E. Petersen, and R. L. Buckner, "Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing," *Neuroimage*, vol. 11, no. 6, pp. 735–759, Jun. 2000.

[40] G. K. Aguirre, E. Zarahn, and M. D'esposito, "The variability of human, BOLD hemodynamic responses," *Neuroimage*, vol. 8, no. 4, pp. 360–369, Nov. 1998.

[41] D. A. Handwerker, J. M. Ollinger, and M. D'Esposito, "Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses," *Neuroimage*, vol. 21, no. 4, pp. 1639–1651, Apr. 2004.

John C. Gore received the Ph.D. degree in physics from the University of London, London, U.K., in 1976. He also received the degree in law.

He holds the Hertha Ramsey Cress Chair in Medicine and is a University Professor of Radiology and Radiological Sciences, Biomedical Engineering, Physics and Astronomy, and Molecular Physiology and Biophysics at Vanderbilt University, Nashville, TN, USA. He founded the pioneering MRI research program at Hammersmith Hospital in the U.K. in the late 1970s prior to establishing and directing the MRI research program at Yale University from 1982 to 2002. Since 2002, he has served as the Founding Director of the Vanderbilt University Institute of Imaging Science, a comprehensive, trans-institutional center that is engaged in multimodal research for biomedical applications. He has published more than 500 original papers and contributions within the medical imaging field. His research interests include the development and application of imaging methods for understanding tissue physiology and structure, molecular imaging and functional brain imaging.

Dr. Gore is a member of the National Academy of Engineering and an elected Fellow of the American Association for the Advancement of Science, the American Institute of Medical and Biological Engineering, the International Society for Magnetic Resonance in Medicine (ISMRM), and the Institute of Physics (U.K.). In 2004, he received the Gold Medal of the ISMRM for his contributions to the field of magnetic resonance imaging. He has served twice as a Trustee of the ISMRM and is an Editor-in-Chief of the journal *Magnetic Resonance Imaging* and a member of the National Advisory Council on Biomedical Imaging and Bioengineering for NIBIB.

René Marois received the Ph.D. degree in neuroscience from Yale University, New Haven, CT, USA, in 1996.

He is a Professor of psychology and an Associate Professor of radiology and radiological sciences at Vanderbilt University, Nashville, TN, USA. His research interests include human attention, working memory, and awareness. His research interest also includes psychophysical and functional neuroimaging approaches.

Santosh B. Katwal (M'05) received the B.S. degree from Tribhuvan University, Kathmandu, Nepal, the M.S. degree from the University of Tennessee, Knoxville, USA, and the Ph.D. degree from Vanderbilt University, Nashville, TN, USA, in 2012, all in electrical engineering.

He is currently a Research Staff Member at the University of Chicago, Chicago, IL, USA, where he is working on developing computer-aided diagnosis methods for breast cancer based on mammograms. From 2007 to 2012, he was a Research Assistant with the Institute of Imaging Science (VUIIS), Vanderbilt University. His research interests include functional brain mapping, machine learning, statistical modeling and the development and application of algorithms for image and high-dimensional data analysis.

Dr. Katwal has been a member of the International Society for Magnetic Resonance in Medicine (ISMRM) since 2008.

Baxter P. Rogers received the B.S. degree in physics from Furman University, Greenville, SC, USA, in 1998, and the M.S. and Ph.D. degrees in medical physics from the University of Wisconsin-Madison, Madison, USA, in 2001 and 2004. From 2004 to 2006, he took a postdoctoral fellowship in neuroimaging at the Vanderbilt University Institute of Imaging Science in Nashville, TN, USA.

He joined the Institute of Imaging Science as a Core Faculty Member in 2006, and is currently appointed at the Vanderbilt University Medical Center as a Research Associate Professor of radiology and radiological sciences, with secondary appointments in biomedical engineering and psychiatry. His research interest includes the development of neuroimaging methods and their application to the study of neurological and mental disorders.

Dr. Rogers has been a member of the International Society for Magnetic Resonance in Medicine since 2000, and the Organization for Human Brain Mapping since 2002. He also currently serves as a Review Editor for *Frontiers in Neuroscience*.