ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage xx (2006) xxx – xxx

Learning functional structure from fMR images

Xuebin Zheng[a] and Jagath C. Rajapakse[a,b,*]

[a]BioInformatics Research Center, School of Computer Engineering Nanyang Technological University, Singapore
[b]Biological Engineering Division, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

Received 6 June 2005; revised 26 January 2006; accepted 30 January 2006

We propose a novel method using Bayesian networks to learn the structure of effective connectivity among brain regions involved in a functional MR experiment. The approach is exploratory in the sense that it does not require an a priori model as in the earlier approaches, such as the Structural Equation Modeling or Dynamic Causal Modeling, which can only affirm or refute the connectivity of a previously known anatomical model or a hypothesized model. The conditional probabilities that render the interactions among brain regions in Bayesian networks represent the connectivity in the complete statistical sense. The present method is applicable even when the number of regions involved in the cognitive network is large or unknown. We demonstrate the present approach by using synthetic data and fMRI data collected in silent word reading and counting Stroop tasks.

Keywords: Bayesian networks; Effective connectivity; Functional MRI; Interference counting; Silent reading; Stroop task

Introduction

With the rapid development of medical imaging techniques, researchers are now able to obtain a multifaceted view of brain function and anatomy (Maurer and Fitzpatrik, 1993). Functional brain imaging represents a range of measurement techniques, which extract quantitative information about physiological function and provide functional maps showing which regions are specialized for different sensory or cognitive functions (Maintz and Viergever, 1996). Although many researchers have attempted to identify the individual brain areas involved in various cognitive tasks, holistic views of effective connectivity of higher-order functions have not been investigated thoroughly. More recently, functional integration studies describing how functionally specialized areas interact and how these interactions lead the brain to perform a specific task have become one of the hot topics in brain mapping research (Penny et al., 2004a). In this paper, we present an exploratory approach to determine effective connectivity among brain regions from fMRI data based on Bayesian graphical models where interactions among the regions are represented by conditional probabilities.

Presently, the information about neural interactions is often extracted by decomposing interregional covariances among activations. Structural Equation Modeling (SEM) has been the most commonly used method to analyze the effective connectivity among brain regions. McIntosh and Gonzalez-Lima (1994) first described SEM and applied for network analysis of vision tasks using PET. Other researchers (McIntosh et al., 1994; Krause et al., 1999; Nyberg et al., 1996; Bavelier et al., 2000; Honey et al., 2002; Nezafat et al., 2001; McKiernan et al., 2001; Petersson et al., 2000; Buchel and Friston, 1997) have later used SEM for the analysis of networks of brain regions involved in sensory or cognitive tasks. Bullmore et al. showed how to search for the best fitting covariance model of connectivity from fMRI data by using SEM (Bullmore et al., 2000). Mechelli et al. (2002) constructed a multisubject network based on SEM to illustrate the differences in connectivity among subjects. The covariances between the brain regions in SEM describe the behavior of a neural system only in the second-order statistical sense, whereas the conditional probability densities (CPDs) characterizing graphical models describe the behavior of a network in the complete statistical sense.

Dynamic Causal Modeling (DCM) was introduced by Friston (2003) to model functional interactions at the neuronal level and comprises a bilinear model for neurodynamics and an extended balloon model for hemodynamics. DCM has shown to be a potential model for making inferences about the temporal changes of effective connectivity from fMRI data (Penny et al., 2004a,b; Friston, 2003). DCM models interactions at the neuronal rather than the hemodynamic level (Penny et al., 2004a), which is more useful in analyzing the temporal interactions between brain regions. Granger causality mapping (GCM), a linear method developed for modeling time-resolved fMR time-series, investigates effective connectivity among activated brain areas by using a vector autoregressive (VAR) model (Goebel et al., 2003). The connectivity is computed by evaluating interactions between a current voxel and a reference voxel and introducing a statistical framework for distinguishing different types of interactions.

* Corresponding author. School of Computer Engineering, Nanyang Technological University, Blk. N4, 50 Nanyang Avenue 639798, Singapore.
E-mail address: asjagath@ntu.edu.sg (J.C. Rajapakse).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2006.01.031

Granger causality mapping renders a voxel-wise connectivity analysis, whereas the present approach is region-wise and seeks for a global representation of a neural system.

The existing methods of connectivity analysis, such as SEM, DCM, and GCM, are confirmatory in the sense that they need a prior connectivity model to begin with. The prior models are often under anatomical constraints and complicated by the fact that many of them have been obtained in the studies of monkeys. And it is not always certain which areas are to be included in the study, especially if the brain regions are involved in functions unique to human, such as language and cognition (Bullmore et al., 2000). Our method based on Bayesian networks allows extraction of the connectivity among brain regions from functional MRI data in an exploratory manner. Bayesian network modeling is widely applicable for compactly representing the joint probability distribution over a set of random variables (Jordan, 1999). In our functional brain networks, the nodes represent the activated brain regions and a connection between two regions represents an interaction between them. The Maximum A Posteriori (MAP) estimation of the structure of the functional network is derived from fMRI data to maximize the Bayesian Information Criterion (BIC) by using a greedy search algorithm.

A synthetic fMRI data set was used to test the feasibility and robustness of the proposed method. The method was further demonstrated by exploring the functional structure from fMRI data obtained in two experiments: a silent word reading task and a counting Stroop task. The network derived for the reading task was compared with the previous literature. The neural systems derived for neutral and interference counting Stroop tasks performed by normal control subjects were used to infer the differences of the performances in the two tasks. The results obtained in the two real fMRI data were consistent with the previous literature and hypotheses, validating the present approach.

Method

Neural systems modeling with Bayesian networks

A Bayesian network, a specific graphical model that utilizes Bayes' rule for inference, consists of a graph structure and a set of parameters indicating the path coefficients. The graph structure S is a *directed acyclic graph* (DAG) that encodes a set of conditional independence assertions about the variables at nodes. The parameters are represented by conditional probability distributions (CPDs) defining the probabilities of the nodes given their parent nodes.

Fig. 1 shows an example of a Bayesian network, representing a neural system consisting of five brain regions; $\{r_i: i = 1, 2, \ldots 5\}$ denotes the set of brain regions activated during the task where r_i represents the ith brain region and x_i denotes the activation of the region; the set of the directed arcs and the conditional probabilities $\{p(x_i|x_j): i, j = 1, 2 \ldots 5; i \neq j\}$ characterize the functional connectivity among the brain regions, in the neural system. The brain regions are presumed to collectively and interactively perform the sensory or cognitive task in the fMRI experiment.

Consider a neural system consisting of a set of n brain regions $R = \{r_i; 1, 2, \ldots, n\}$ that is capable of collectively performing a particular sensory or cognitive task. The activation of a brain region r_i is represented by the average of the time courses of

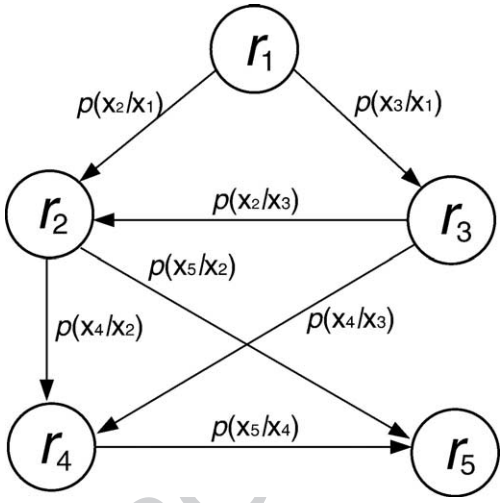

Fig. 1. Illustration of a neural system represented by a Bayesian network: the set of five activated brain regions $\{r_i: i = 1, 2, \ldots, 5\}$ is represented by the nodes, and the conditional probabilities among them, $\{p(x_i|x_j): i, j = 1, 2, \ldots, 5; i \neq j\}$, represent the interactions.

hemodynamic responses of the neurons in the region. Suppose that the average of the time-series responses of the activated brain region is x_i. The fMRI experiment is represented by the data set containing activations of all activated brain regions: $x = \{x_i: 1, 2, \ldots n\}$. From the chain rule of probability, the likelihood of the activation of the neural system is given by:

$$p(x) = \prod_{i=1}^{n} p(x_i|x_1,..,x_{i-1}) \quad (1)$$

where $p(x)$ indicates the joint probability of the activations of all brain regions in the neural system and defines the likelihood of the function of the neural system. For each variable x_i, let $a_i \subseteq \{x_j; j = 1, 2, \ldots, n, i \neq j\}$ be a set of parent nodes of x_i that renders x_i and its ancestors conditionally independent. That is,

$$p(x_i|x_1, x_2,\ldots, x_{i-1}) = p(x_i|a_i, \theta_i) \quad (2)$$

where θ_i denotes the parameters of the distribution.

Then, a Bayesian network representing the joint probability of the activation of all brain regions, i.e., of the whole brain system, can be written as:

$$p(x) = \prod_{i=1}^{n} p(x_i|a_i, \theta_i) \quad (3)$$

where θ_i indicates the parameters of the CPDs, involving brain region r_i and its parent nodes in a_i. Let $\theta = \{\theta_{i,j}: i, j = 1, 2, \ldots n; i \neq j\}$ denotes the set of parameters of the whole neural system. We presume that all CPDs in the graphical model carry the same form.

For two activated regions r_1 and r_2, the interaction or the influence from region r_1 to r_2 is indicated by the conditional probability $p(x_2|x_1)$, and the influence from r_2 to r_1 is $p(x_1|x_2)$. Since the activities of r_1 and r_2 are not independent, the distribution of x_1 will be affected when x_2 is given, and vice versa. Thus, the interactions of two linked nodes are bi-directional in a Bayesian network. One of the biggest advantages for choosing Bayesian networks is that they have the bi-directional message passing architecture and can be learned in an unsupervised manner from data.

Learning the structure

The structure learning refers to the learning of the topology of the functional network with respect to the parameterization used. We attempt to learn the structure of the neural system from functional MRI data by taking a Bayesian approach considering the probability distributions over the parameters or models. This allows the determination of the confidence of one's estimate and the usage of predictive techniques such as Bayesian model averaging (Murphy, 2002). The present model is a directed model, as referred to the Bayesian networks, where all the nodes are fully observed and the interactions are presumed to be Gaussian.

We attempt to obtain the Maximum A Posteriori (MAP) estimation of the structure, $\hat{S}$, given all the data set:

$$\hat{S} = \max_S p(S|D) \tag{4}$$

where from Bayes theorem,

$$p(S|D) = \frac{p(D|S)p(S)}{p(D)}; \tag{5}$$

as the denominator does not depend on S, only the numerator is needed to be maximized. $p(S)$ is assumed to have a uniform prior over the structures (Heckerman and Geiger, 1995), and, to compute $p(D|S)$, the Bayesian approach averages over all possible parameters, weighing each by their posterior probability:

$$p(D|S) = \int p(D|S,\theta)p(\theta|S)d\theta. \tag{6}$$

For large samples, the term $p(D|S, \theta)p(\theta|S)$ is reasonably approximated as a multivariate Gaussian (Kass and Raftery, 1995). In addition, approximating the mean of the Gaussian with the maximum likelihood (ML) estimates of θ and ignoring the terms that do not depend on the data set size N, we obtain the Bayesian Information Criterion (BIC), indicating the fitness of the graph to the data:

$$\mathrm{BIC}(\theta) = \log\left\{p\left(D|\hat{\theta}\right\} - 0.5l\ \log\{N\}\right. \tag{7}$$

where $\hat{\theta}$ is the ML estimate of the parameters and l is the number of free parameters of the model. The present approach assigns a score to each candidate graphical model, which measures how well the graphical model describes the data set D (Margaritis, 2003) and yields the best fit model by optimizing the BIC score.

There are two different approaches for learning the structure of the network: constraint-based approach and search-and-score approach (Jordan, 1999). The constraint-based approach begins with a fully connected graph and removes edges in a sequential manner if certain conditional independencies are absent in the data. This approach has the disadvantage of repeated independence tests, leading to a loss of statistical power. The more popular search-and-score approach searches through the space of possible DAGs and returns either the best one or a sample of the best models by using a fitness score (Murphy, 2004). Since the number of DAGs is super-exponential of the number of nodes, an exhaustive search in the space is impractical. So, either a local search algorithm, such as greedy hill climbing, or a global search algorithm, such as Markov Chain Monte Carlo (MCMC) method (Wesley, 1994), should be employed. We used the Metropolis–Hastings (MH) algorithm (Wesley, 1994), an MCMC algorithm, to search the space of DAGs to find the optimal structure of the network.

Experiments and results

In this section, we illustrate our technique with experiments on a synthetic data set and two fMRI data sets obtained from the fMRI Data Center, Dartmouth College (fMRIDC): a silent word reading task (access number: 2-2000-11189) and a counting Stroop task (access number: 2-2000-1123B). We tested our method on a synthetic data set for robustness and compared the results with the SEM approach. The structures of the neural systems involved in the two tasks were derived, and their validity was investigated with the help of the past literature and known hypotheses.

Synthetic data

Synthetic fMRI data sets were generated to test the feasibility and robustness of the proposed method for detecting the underlying neural system.

Data generation and simulation

A neural system was simulated with synthetic time-series where interactions among the brain regions are represented by linear coefficients. Suppose that the activities had zero mean Gaussian variates with an $n \times n$ covariance matrix $\mathbf{\Sigma}$, i.e., $N(x; 0, \mathbf{\Sigma})$. Regression equations describe how the activity of one region is related to the activity of the other regions with a set of linear coefficients:

$$\boldsymbol{x}_t = \mathbf{M}_{\boldsymbol{x}_t} + e_t \tag{8}$$

where $\boldsymbol{x}_t$ denotes the vector of activations of the regions at time t and e_t is the zero mean Gaussian innovation. Matrix $\mathbf{M} = \{m_{ij}\}_{n\times n}$ is formed by the predicted interactions among regions. By subtracting $\mathbf{M}_{\boldsymbol{x}_t}$ from both sides of the regression equation and multiplying by $(\mathbf{I} - \mathbf{M})^{-1}$, where $\mathbf{I}$ is an $n \times n$ identity matrix, the equation becomes:

$$\boldsymbol{x}_t = (\mathbf{I} - \mathbf{M})^{-1} e_t. \tag{9}$$

Eq. (9) can be used to generate synthetic data from a known model given by $\mathbf{M}$. The Gaussian variates e_t was randomly generated and then pre-multiplied by $(\mathbf{I} - \mathbf{M})^{-1}$. This approach was repeated for each t to obtain the time-series.

All synthetic time-series were simulated to have 300 time points, and the data set was generated based on the following parameters: the structure was the same as in Fig. 1; the nonzero elements of the linear coefficient matrix $\mathbf{M}$ were $m_{21} = 1.1$, $m_{23} = 0.6$, $m_{31} = 0.8$, $m_{42} = 1.3$, $m_{43} = 1.1$, $m_{52} = 0.9$, and $m_{54} = 1.2$. We used the present method to derive the functional structure from the synthetic data set.

Robustness

The synthetic data set was corrupted by adding random Gaussian noise (Signal/Noise = 1.0) at randomly selected time points for each time-series to test the robustness of our method. The percentage of corrupted time points was varied from 10% to 60% in steps of 10%.

We used a likelihood ratio (LR) measure to assess the matching between the learned structure and the known structure as for a given specificity, no other test renders a higher sensitivity (Penny et al., 2004a). If $p(x|\theta, \hat{S})$ and $p(x|\theta, S)$ are the likelihoods of the

estimated structure $\hat{S}$ and the actual structure S, then the log of the likelihood ratio is given by

$$\log R = \log p\left(x|\theta,\hat{S}\right) - \log p(x|\theta,S). \quad (10)$$

Under the null hypothesis that the models are identical, and for large t, $-2 \log R$ is distributed as a χ^2 variable having degrees of freedom equal to the difference in number of parameters between the models. The results of fitness of our method at various amount of noise are shown in Fig. 2. The values of log LRs were scaled between 0 and 1 for better display. The results were stable until 40% of the data points were corrupted by random noise.

Comparison with SEM

The SEM approach proposed by Bullmore et al. (2000) was used to derive the neural systems generated by the synthetic data sets, and the performances were compared with our technique with Bayesian networks. Several synthetic data sets were generated to simulate brain systems with different number of regions, $n = 3, 4, \ldots, 15$, as illustrated in Fig. 1. The log-likelihood ratios against the number of brain regions are shown in Fig. 3.

As seen, our technique with Bayesian networks derived the neural systems closer to the ground truth on all randomly generated synthetic data sets. In the case of synthetic network with 13 regions, the estimated structure did not match well with the actual structure, indicating that the algorithm might have fallen into a local minimum during searching. As the number of regions in the neural system increases, the probability of the structure falling into the local minimum becomes higher.

Silent reading task

Data

The fMRI data used in this experiment consist of six subjects (five males, one female), aged between 20 and 34, with English as the first language. The experiment consisted of a 3 × 2 factorial design, three frequencies of presentation: 20, 40, and 60 words per minute, and for each, words and pseudowords presentations alternated with a resting condition. The task involved silent reading of words and pseudowords as soon as they appeared on the screen; the resting condition involved fixating to a cross in the middle of the screen. Each subject was presented with 105 words and 105 pseudowords. Stimuli were composed of four, five, or six letters and were presented in 12 blocks. Each block lasted 21 s and was followed by a resting period of 16 s. Data for each subject contain 360 volume images with a repetition time (TR) of 3.15 s/volume. For more details of the experiment, the reader is kindly referred to Mechelli et al. (2000).

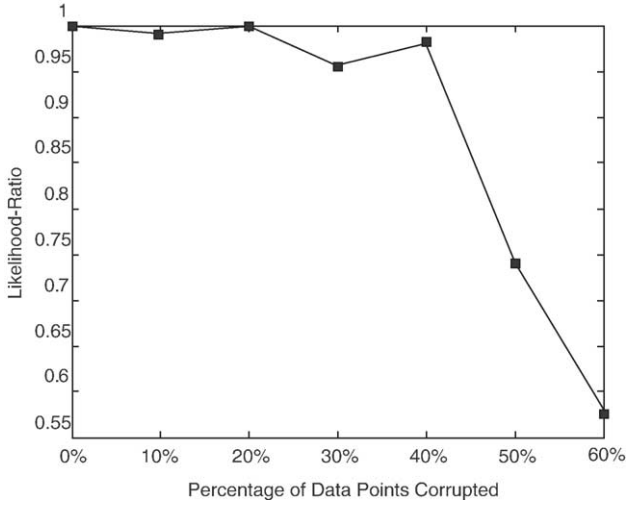

Fig. 2. Illustration of the robustness of the proposed method for deriving neural systems: the log-likelihood ratios of prediction versus the percentage of number of data points corrupted by random noise.

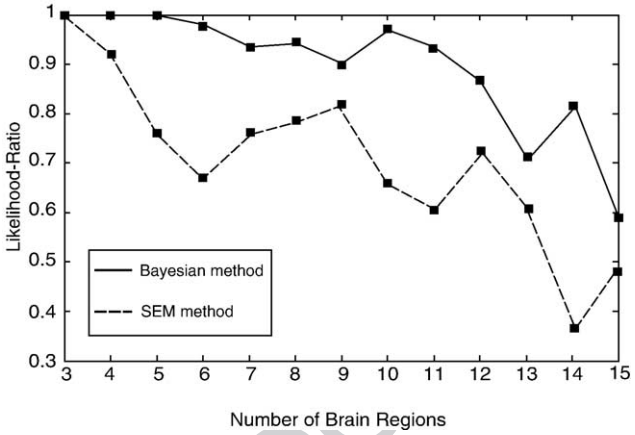

Fig. 3. The comparison of performances in deriving the functional structures of neural systems, by the SEM method and the present approach: the log-likelihoods are shown against the number of brain regions.

Detection of activation

All functional images of the subjects were realigned, coregistered, normalized, and smoothed as the preprocessing steps. The design matrix, convolved with a synthetic hemodynamic response function (HRF), was used as the reference waveform for each time-series and then estimated the parameters of the linear model. The time-series were high-pass-filtered using a set of discrete Cosine basis functions with a cutoff period of 156 s and low-pass-filtered using a symmetric HRF as the smoothing kernel to condition the temporal autocorrelations (see Mechelli et al., 2000 for details).

The regions showing increased activity during reading for both words and pseudowords were identified by statistically comparing the fMRI signal while reading relative to the rest condition. The changes in the blood oxygenation level dependent (BOLD) contrast, associated with the performance of the reading task, were assessed on a voxel-by-voxel basis by using the general linear model (Friston et al., 1995) and the theory of Gaussian fields (Worsley and Friston, 1995). This analysis pipeline thus uses multivariate regression analysis and corrects for temporal and spatial autocorrelations of the fMRI data. Group analyses were performed using a fixed-effect analysis (FFX) (Friston et al., 1999). Significant hemodynamic changes for each contrast were assessed using the t statistical parametric maps, and the results were reported by giving the t values; and the statistical inferences were made at $P < 0.05$ corrected for multiple comparisons by using Family-wise Error Rate (FWER) (Worsley et al., 1996, 2004).

We used SPM2 (Friston et al., 1995) for the above analysis-preprocessing and identification of significantly activated regions. Talairach daemon database (Lancaster et al., 2004) and the co-planar stereotaxic atlas (Talairach and Tournoux, 1988) were used to assist the specification of the activated regions in Talairach coordinates. The Montreal Neurological Institute (MNI) coordi-

t1.1 Table 1
t1.2 Significantly activated regions during the reading condition relative to the rest condition are shown in 3D MNI coordinates with t statistics

t1.3 Brain regions (Brodmann areas)	Coordinates	t value
t1.4 Left extrastriate cortex (LEC: BA18, BA19)	(−16, −98, 6)	17.19
t1.5 Right extrastriate cortex (REC: BA18, BA19)	(16, −99, −6)	17.19
t1.6 Left superior parietal lobule (LSPL: BA7)	(−28, −60, 56)	7.65
t1.7 Right superior parietal lobule (RSPL: BA7)	(24, −58, 54)	7.53
t1.8 Left middle temporal cortex (LMTC: BA21, BA22)	(−50, −52, 8)	6.51
t1.9 Right middle temporal cortex (RMTC: BA21, BA22)	(58, −46, 8)	8.13
t1.10 Left inferior frontal gyrus (LIFG: BA44, BA45)	(−40, 12, 28)	7.33
t1.11 Right inferior frontal gyrus (RIFG: BA44, BA45)	(40, 8, 30)	7.46
t1.12 Left middle frontal gyrus (LMFG: BA46, BA9)	(−48, 36, 6)	6.68
t1.13 Right middle frontal gyrus (RMFG: BA46, BA9)	(40, 38, −8)	6.50

t1.14 Statistical inferences were made at $P < 0.05$ corrected for multiple comparisons by using FWER.

nates given by SPM2 were converted to the corresponding Talairach coordinates by using the technique described by Brett (2002). Table 1 and Fig. 4 show the activations found during the silent word reading task. The activations were found in bilateral extrastriate cortices, superior parietal lobes, middle temporal cortices, inferior frontal sulci, and middle frontal cortices, and the cerebellum.

Derivation of neural system

The time courses of significantly activated regions were extracted by taking the averages of the time-series at the peak-activated voxels and its neighbors at the cluster level for all subjects. All extracted time-series representing activated regions were formed into a matrix as the input to learn the structure of the neural system. The Metropolis–Hastings algorithm was used to search the space of all DAGs, with the Bayesian Information Criterion (BIC) as the score function to find the optimal model. The software package, *Bayes Net Toolbox*, written by Murphy (2004) was used for structure learning. Fig. 5 shows the posterior probability of the DAGs, assuming a uniform structural prior, and each point in the horizontal axis, representing a possible graph structure; the structure with the highest score was chosen to

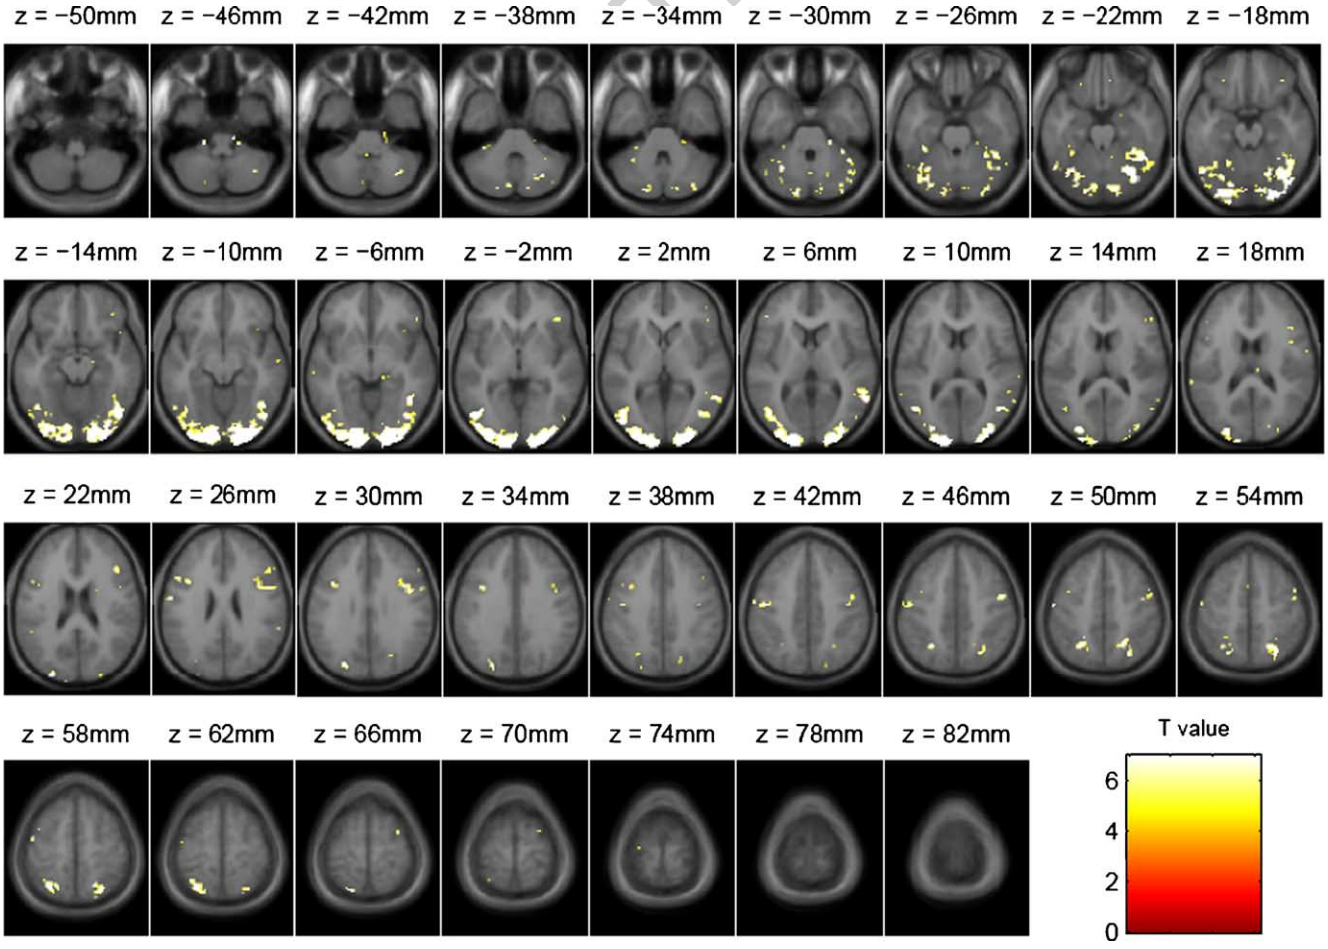

Fig. 4. Significantly activated brain regions obtained in the group study (using the fixed-effect analysis) of the silent reading task.

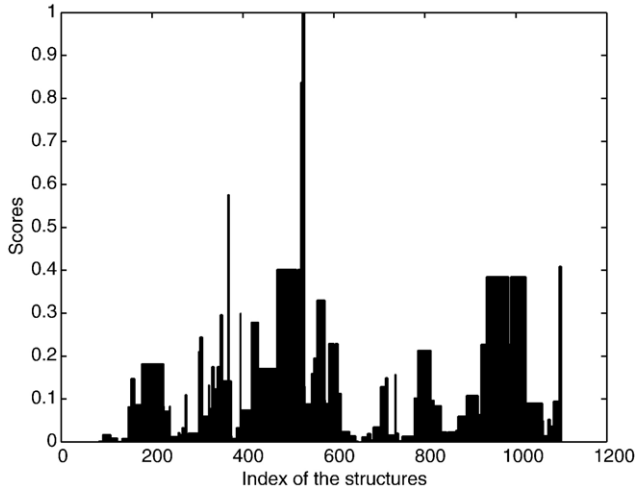

Fig. 5. The posterior scores of the possible DAGs derived from the Metropolis–Hastings algorithm, assuming a uniform prior for the structures.

represent the network of this particular task. Fig. 6 shows the acceptance ratio versus the number of the iteration steps as a crude convergence diagnostic during the search for the optimal structure. The network which had the highest BIC score is shown in Fig. 7.

The left hemisphere has been the focus of the analysis of the neural correlates of reading tasks. Since some language tasks such as those involving different languages, English-knowing bilinguals, literate versus illiterate, etc., show activation in both hemispheres (Kim et al., 1997; Tan et al., 2000; Petersson et al., 2000), we included all the activated regions of the cortex and explored all possible connections among all the brain regions.

The extrastriate cortex (EC: BA18, BA19) in the visual cortex plays the role of visual representation in word processing (Kolb and Whishaw, 1996). The connection from the extrastriate cortex to superior parietal lobe (SPL: BA7) forms the dorsal stream of visual analysis, performing the perception of visual word form. As seen in Fig. 7, the connections from EC to SPL are found in both hemispheres (LEC → LSPL and REC → RSPL). Meanwhile, the connections from the EC to prefrontal cortex including middle frontal gyrus (MFG: BA46, BA9) and inferior frontal gyrus (IFG: BA44, BA45) represent the information flow for the processing of semantic analysis and decision (LEC → LMFG, LEC → RMFG, LEC → RIFG, REC → RMFG, REC → LMFG, and REC → LIFG) (Bullmore et al., 2000). Furthermore, the connections between EC and middle temporal cortex (MTC: BA21, BA22), associated with the retaining and recalling of words from the memory (Kolb and Whishaw, 1996), are found in both hemisphere with reversed directions (REC → RMTC, LMTC → LEC); the reversed direction may be due to the bi-directional characteristic of the connectivity, represented by the Bayesian networks. In addition, a homologous interhemispheric connection between the ECs of both sides (REC → LEC) is found, which may be due to the transcallosal inferences between two hemispheres (McIntosh et al., 1994).

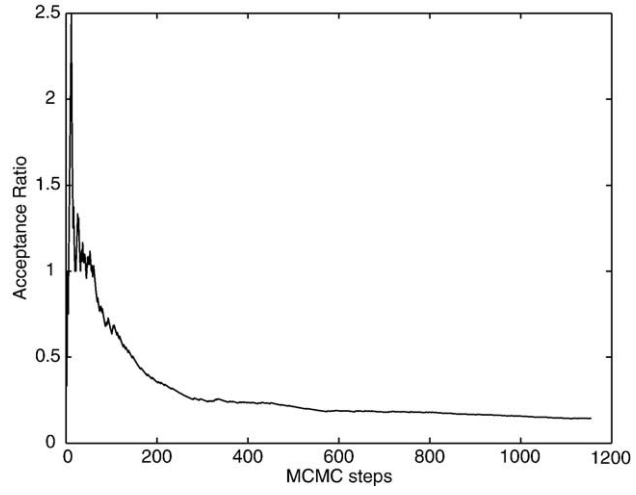

Fig. 6. The acceptance ratio versus the number of MCMC steps in finding the optimal structure of the neural system.

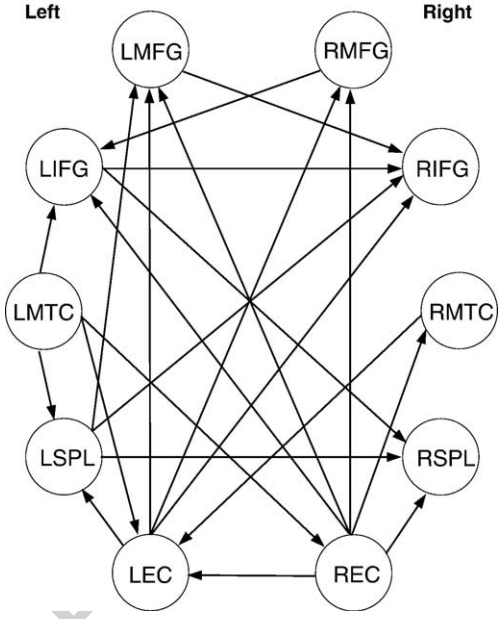

Fig. 7. The neural system learned from fMRI data of the silent reading task. L(R)EC: left (right) extrastriate cortex, L(R)SPL: left (right) superior parietal lobe, L(R)MTC: left (right) middle temporal cortex, L(R)IFG: left (right) inferior frontal gyrus, L(R)MFG: left (right) middle frontal gyrus.

The parietal lobe generally performs the integration of sensory information for the control of movement. In particular, the superior parietal lobe (SPL: BA7) plays the role of visual analysis and mainly makes efferent connections to the prefrontal cortex including MFG and IFG, providing more elaborate information (LSPL → LMFG, LSPL → RIFG) (Kolb and Whishaw, 1996). A homologous interhemispheric connection is also found between the SPLs (LSPL → RSPL). As seen in Fig. 7, the functional links from EC via SPL to prefrontal cortex form the dorsal visual pathway of language processing (LEC → LSPL → LMFG, LEC → LSPL → RIFG) (McIntosh et al., 1994).

The temporal lobes are involved in understanding and processing language, intermediate and long-term memory, complex memories, the retrieval of language or words, and emotional responses (BrainPlace.com). The middle temporal cortex (MTC: BA21, BA22) involved in our model is the general association cortex that integrates the input from the lower level auditory and visual areas for retaining in the memory. In particular, the posterior aspect of the left middle temporal cortex, which is also called the Wernicke's area, is involved in storing the visual word forms and processing lexical–semantic information (Fiebach et al., 2002). It

is supposed to have connections with LSPL for movement control (LMTC → LSPL), with the prefrontal cortex for semantic phonologic retrieval and semantic processing (LMTC → RIFG, LMTC → LMFG) and with EC for memory retention (LMTC → LEC, LMTC → REC) (Price, 2000; Hampson et al., 2002; Horwitz and Braun, 2003).

The MFG is involved in tasks that require executive control, such as the selection of behavior based on short-term memory (Krause et al., 1999). It receives inputs from the posterior parietal and superior temporal sulci. The IFG is most active for phonemic decisions and receives inputs from temporal lobes and parietal lobes (Price, 2000; BrainPlace.com). As seen in Fig. 7, except for the connections that have been mentioned above, there are interhemispheric connections between the prefrontal regions, including the interconnection between the homologous regions of IFG (LMFG → RIFG, RMFG → LIFG, LIFG → RIFG), which may be involved in semantic processing during inner speech (Bullmore et al., 2000).

In the derived interhemispheric language network, the left hemisphere showed dominant pathways, which is consistent with the traditional language network (left-hemispheric). Although the activations have been symmetrically distributed in both hemispheres, the right hemisphere activations may be due to the transcallosal influence of the left. This hypothesis is supported by the fact that there are more connections between the regions in the left hemisphere and the regions in the right hemisphere receive only results of processing in the left regions.

The connections in our model that are consistent with the previous literature are given in Table 2. Due to the fact that the specific networks for each cognitive task are different even though the tasks are very similar (e.g., different presenting rate, different words, or different block design in reading tasks), the existing literature can only be used as a general reference to an existing connection. The connectivity pattern derived from our method is consistent with the information flow in the silent reading task as evidenced by the literature, but the connections without a corresponding reference cannot be corroborated.

Table 2 t2.1
The list of the connections between the activated brain regions, found to be involved in the silent reading task, which had been previously verified in other language-based tasks t2.2

Connection	Functional description	Relative reference	
LEC → LSPL	Perception of visual word form	(Horwitz et al., 1998)	t2.4
REC → RSPL	Perception of visual word form	(McIntosh et al., 1994)	t2.5
LEC → LMFG	Semantic decision and analysis	(Krause et al., 1999)	t2.6
		(Bullmore et al., 2000)	t2.7
REC → RMFG	Semantic decision and analysis	(Krause et al., 1999)	t2.8
REC → LEC	Homologous interconnection	(McIntosh and Gonzalez-Lima, 1994)	t2.9
		(McIntosh et al., 1994)	t2.10
		(Krause et al., 1999)	t2.11
LSPL → LMFG	Executive control	(Honey et al., 2002)	t2.12
LSPL → RIFG	Phonemic decisions	(Honey et al., 2002)	t2.13
LSPL → RSPL	Homologous interconnection	(Honey et al., 2002)	t2.14
LMTC → LSPL	Semantic processing	(Price, 2000)	t2.15
		(Horwitz et al., 1998)	t2.16
LMTC → LIFG	Semantic phonologic retrieval	(McKiernan et al., 2001)	t2.17
		(Matsumoto et al., 2004)	t2.18
		(Hampson et al., 2002)	t2.19
		(Mechelli et al., 2002)	t2.20
LMTC → LEC	Memory retention	(Nyberg et al., 1996)	t2.21
LMTC → REC	Memory retention	(McIntosh et al., 1994)	t2.22
LMFG → RIFG	Inner speech production	(Krause et al., 1999)	t2.23
		(Nyberg et al., 1996)	t2.24
		(Petersson et al., 2000)	t2.25
LIFG → RIFG	Homologous interconnection	(Honey et al., 2002)	t2.26

Interference counting task

Data

Functional MRI data used in this experiment were obtained from a counting Stroop task testing the cognitive interference that occurs when processing of one stimulus feature impedes the simultaneous processing of a second stimulus attribute (Bush et al., 1998). Data were collected by Tamm et al. (2002) to investigate the performance of females with fragile X-syndrome on the cognitive interference task compared to a healthy control group. The participants included 14 females with fragile X-syndrome and 14 age-matched healthy control females without the fragile X mutation, ranging in age from 10 to 22 (mean age 15.43). The task consisted of 12 alternating experimental (interference) and controlled (neutral) conditions with the rest condition. For both conditions, the subjects were instructed to press the button that corresponded to the number of words appearing on the screen. During the neutral counting task, the word "fish" was presented 1, 2, 3, or 4 times on the screen (15 trials) and during the interference counting task, the words "one", "two", "three", and "four" were presented 1, 2, 3, or 4 times on the screen (15 trials). Stimuli were presented for 1350 ms at a rate of one every 2 s (TR) for a total of 180 trials (90 experimental, 90 control). For more details of the experiment, the reader is referred to Tamm et al. (2002). Our method is demonstrated using the data collected only on the control group.

Detection of activation

We explore the networks involved in the neutral and interference counting tasks by normal controls and attempt to make inferences on the differences of the performances of the two tasks (Fig. 8).

The preprocessed functional images of the subjects were provided by the fMRIDC; images were reconstructed by using Inverse Fourier Transform from each of the 225 time points into 64 × 64 × 18 image matrices and voxel size of 3.75 × 3.75 × 7 mm^3. Using SPM2, the images were motion-corrected again to reduce the artifacts (Friston et al., 1996) and the regions showing significant activation during counting relative to the rest condition were detected using the fixed-effect analysis. The statistical inferences were made at $P < 0.05$ corrected for multiple comparisons by using the Family-wise Error Rate (FWER). Table 3 and Fig. 3 show significant activations of the control group in this experiment. The activations were found in both neutral and interference conditions in right superior parietal lobe (RSPL), left inferior parietal lobe (LIPL), anterior frontal gyrus (AFG), right lateral middle frontal gyrus (RLMFG), medial middle frontal gyrus (MMFG), ventral inferior frontal gyrus (VIFG), primary motor area (PMA), supplementary motor area (SMA) and anterior cingulate cortex (ACC). The left superior parietal lobe (LSPL) and left lateral middle frontal gyrus (LLMFG) were significantly activated only in

the interference task. On the other hand, activation was seen on either side of the lateral inferior frontal gyrus (LIFG) for both tasks (left for the neutral task and right for the interference task). Thus, despite the similar activation in the medial cortex (including ACC, SMA, VIFG, and AFG), the left hemisphere showed more activations in the interference counting task. Although proper motion correction was performed on the data, the crescentic frontal activations (AFG) in Fig. 3 may look like motion artifact (Bullmore et al., 1999; Field et al., 2000; Friston et al., 1996; Gavrilescu et al., 2004).

Derivation of neural system

The time courses of significantly activated brain regions were extracted by taking the averages of the time-series at peak-activated voxels and neighboring voxels at the cluster level for all subjects. The extracted time courses were then used as the input data for learning the structure of the neural system, by using a search-and-score method, similar to the silent reading word task. The networks which had the highest BIC scores for the two tasks are shown in Fig. 3; here onwards, we refer them as "neutral network" and "interference network", respectively.

The similar activation seen in the medial cortices for both conditions may indicate that the function of counting is mainly processed by the medial areas especially in the anterior cingulate cortex (ACC: BA24), which had been shown to be playing an essential role in counting Stroop (Hayward et al., 2004; Shin et al., 2001; Bush et al., 1998). Thus, the different activation in the lateral cortices between the two conditions may reflect the effects of "interference"; more activation in the language areas in the left hemisphere was found in the interference counting task. This is due to the fact that the subjects had been distracted by the meaning of the words being counted in the interference counting task.

The ACC is engaged during the Stroop task in order to resolve competing streams of information in the selection of sensory inputs and responses (Bush et al., 1998). The effects are reflected in the interference network by the connections, ACC → LLMFG (BA9) and ACC → RLMFG (BA9), for resolving interference effects, and ACC → LLIFG (BA44), for phonemic decisions. The absence of connections from ACC to the left hemisphere in the neutral task shows more involvement of the semantic processing and decision making in the interference network (Fig. 9).

The LMFG (BA9) is involved in tasks that require executive control and selection of behavior based on the short-term memory and receives inputs from the posterior parietal region (Price, 2000; BrainPlace.com). In this experiment, this region is involved in processing Stroop-related conflict and resolving interference effects (Tamm et al., 2002). The LLMFG in the interference network is connected to the LLIFG (BA44), and the RLMFG in the

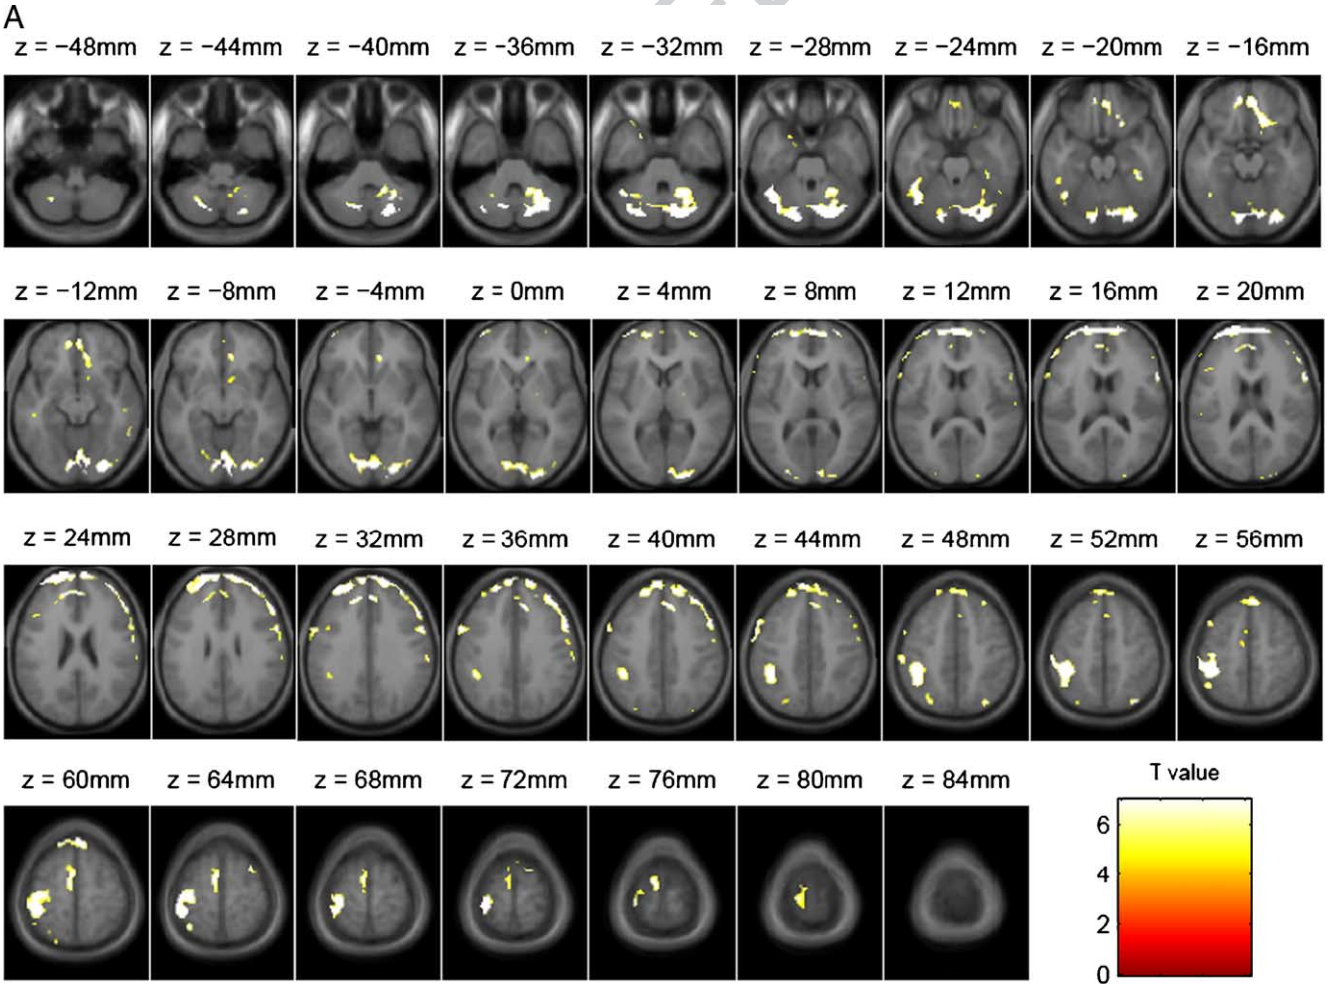

Fig. 8. Brain regions showing significant activation in counting Stroop tasks relative to the rest condition: (A) the neutral counting and (B) the interference counting. Statistical inferences were made at $P < 0.05$ corrected for multiple comparisons using FWER.

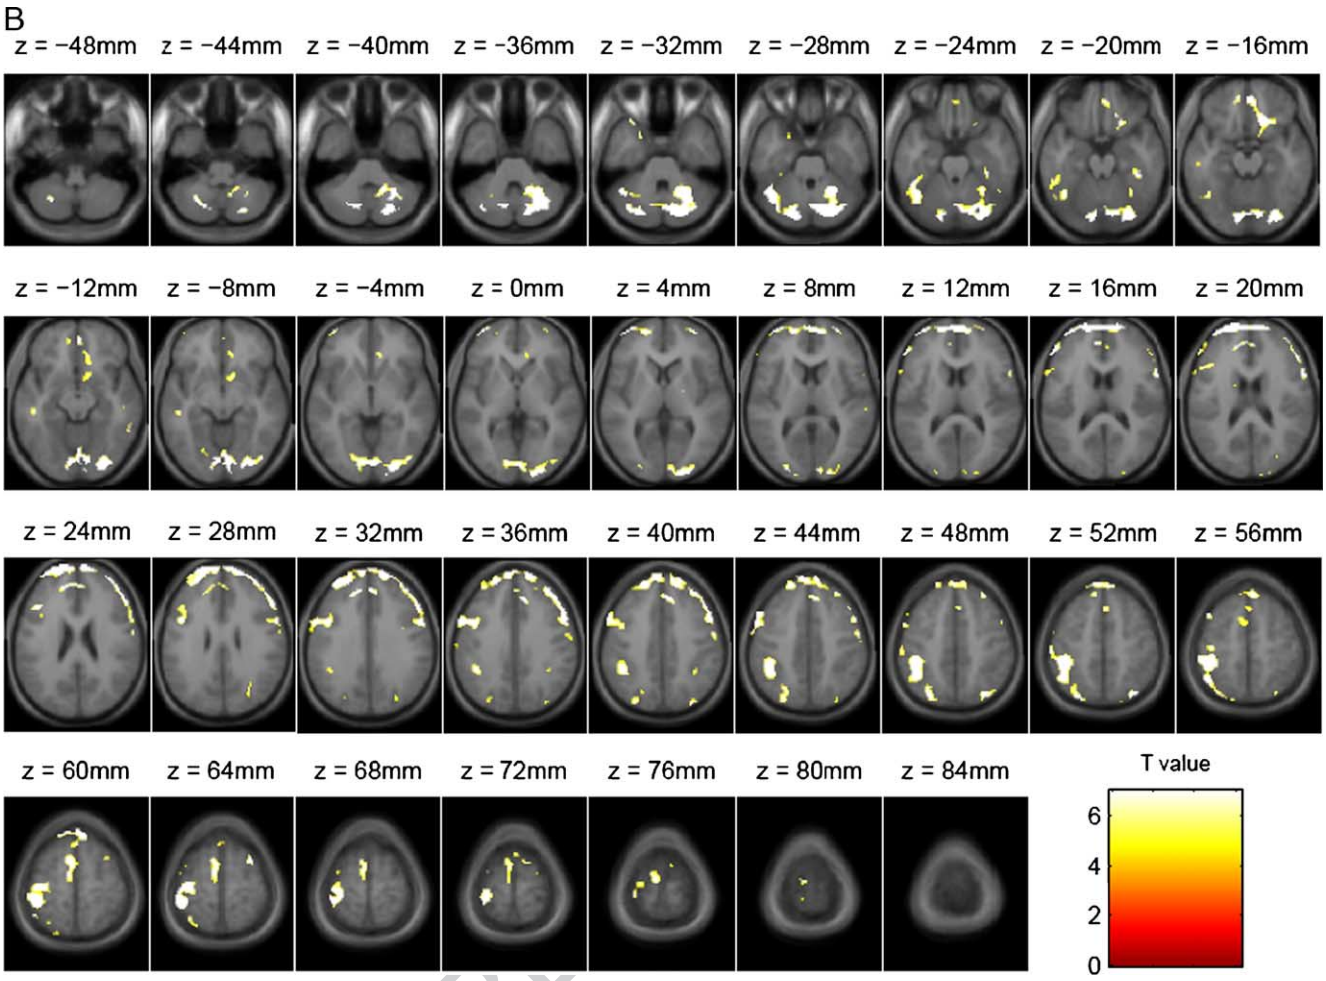

Fig. 8 (*continued*).

neutral network has been connected to the RLIFG for executive controls. The MMFG (BA8) is believed to play an important role in the control of eye movements (Faw, 2002). The common connections found for both tasks are MMFG → RLMFG, MMFG → RLIFG, and MMFG → LLIFG. The connection MMFG → SMA in the neutral network is absent in the interference network, while the connections MMFG → RSPL, MMFG → ACC, and MMFG → LPMA in the interference network are absent in the neutral network. The difference may be due to the different demands of concentration needed by the tasks.

Table 3
The results of the analysis of the activation patterns of the control group performing the counting Stroop task: significantly activated regions during the counting tasks relative to the rest condition are shown in 3D MNI coordinates with the significance values given by *t* statistics

Brain regions	Neutral counting		Interference counting	
	Coordinates	*t* value	Coordinates	*t* value
Left superior parietal lobe (LSPL: BA7)			(−28,−74, 50)	6.42
Right superior parietal lobe (RSPL: BA7)	(32, −72, 50)	4.98	(32, −72, 50)	6.10
Left inferior parietal lobe (LIPL: BA40)	(−42, −38, 58)	10.97	(−42, −38, 60)	11.00
Anterior frontal gyrus (AFG: BA10)	(2, 64, 14)	10.52	(4, 64, 14)	10.86
Left lateral middle frontal gyrus (LLMFG: BA9)			(−54, 16, 44)	5.14
Right lateral middle frontal gyrus (RLMFG: BA9)	(54, 12, 38)	5.09	(54, 12, 38)	6.32
Medial middle frontal gyrus (MMFG: BA8)	(6, 34, 40)	8.47	(6, 34, 40)	8.64
Left lateral inferior frontal gyrus (LLIFG: BA44)			(−56, 8, 34)	5.44
Right lateral inferior frontal gyrus (RLIFG: BA44)	(56, 8, 34)	5.06		
Ventral inferior frontal gyrus (VIFG: BA47)	(16, 26, −16)	5.36	(16, 26, −16)	5.98
Supplementary motor area (SMA: BA6)	(−6, −4, 64)	6.36	(−6, −4, 66)	6.29
Left primary motor area (LPMA: BA4)	(−32, −26, 68)	6.98	(−34, −24, 66)	6.64
Anterior cingulate cortex (ACC: BA24)	(10, 36, −8)	5.17	(10, 34, −8)	4.77

Statistical inferences were made at $P < 0.05$ corrected for multiple comparisons by using FWER.

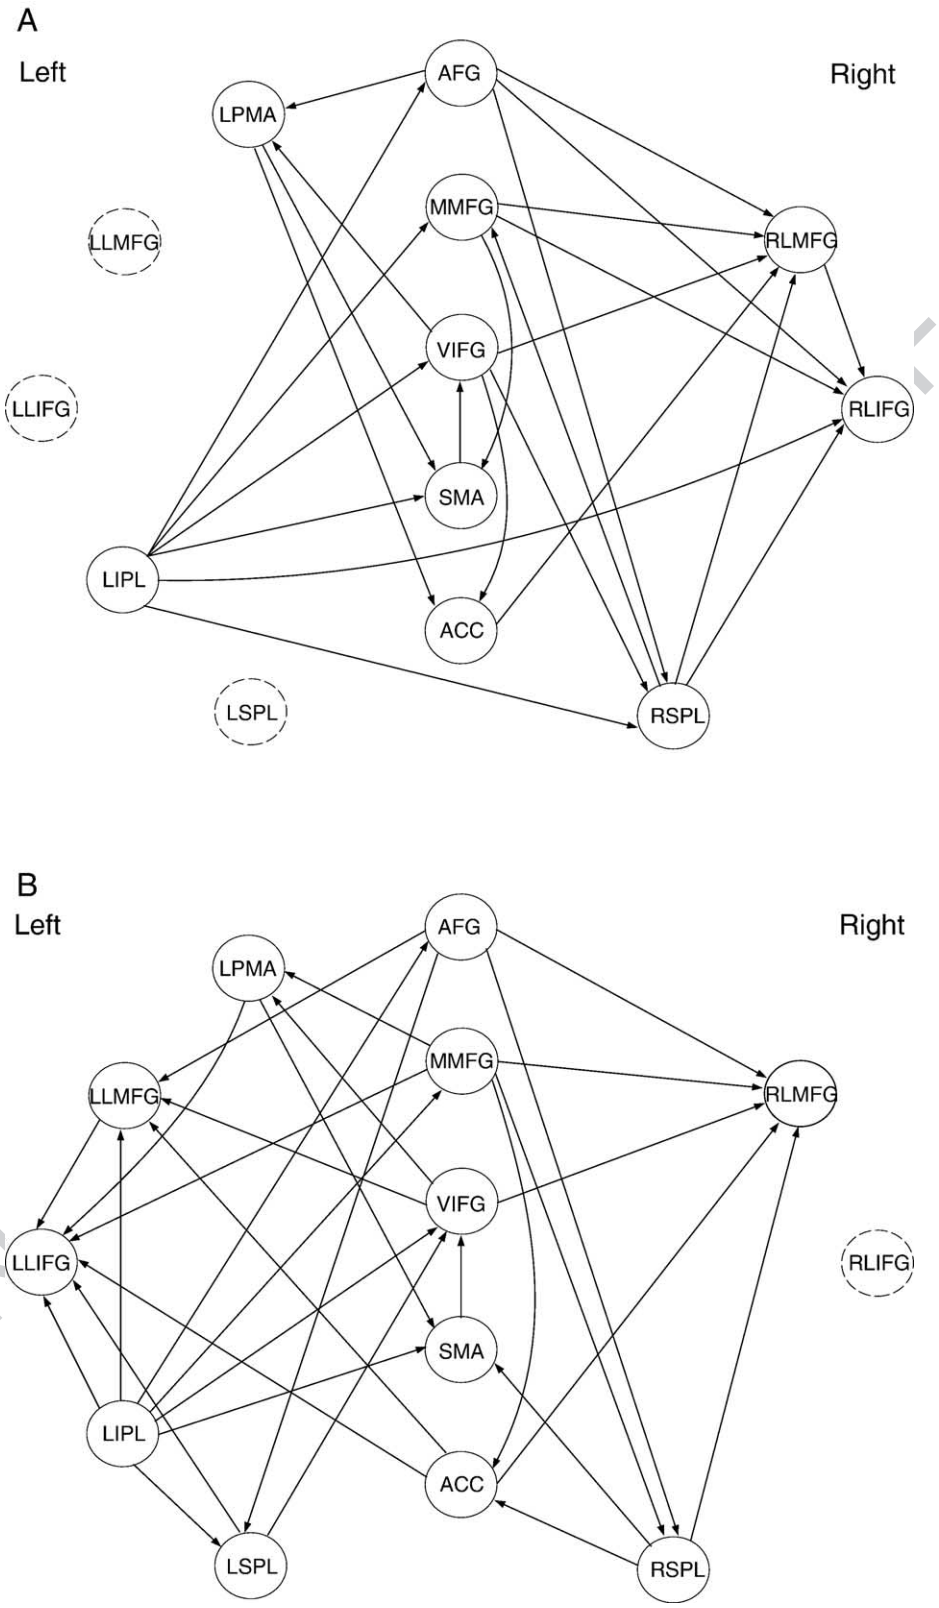

Fig. 9. Structures learned from the data for (A) the neutral counting task and (B) the interference counting task. A dotted circle indicates that the region is not significantly activated in the particular task.

The LIFG is mostly active for phonemic decisions and receives inputs from parietal lobes (Price, 2000; BrainPlace.com). In Fig. 3, the LIFG in both networks has no output connection to other regions. The VIFG (BA47), including orbitofrontal cortex, plays a specific role in controlling voluntary goal-directed behavior (Tamm et al., 2002). The common connection for both tasks,

VIFG → LPMA (BA4), stores the voluntary activities involved (Faw, 2002; Wu et al., 2004). The connection VIFG → LLMFG (BA9), found only in the interference network, is related to the specific function of LMFG involved in processing Stroop-related conflict.

The AFG (BA10) is believed to play a part in strategic processes involved in memory retrieval and executive function (Faw, 2002). The connections from AFG to other regions that are common in both networks include: AFG → SPL (BA7) and AFG → LMFG (BA9). The connections AFG → LPMA (BA4) and AFG → RLIFG (BA44) are present in the neutral network but absent in the interference network. The SMA is believed to play a role in the planning of complex and coordinated movements (Kolb and Whishaw, 1996). A connection from SMA to ventral inferior frontal gyrus was found in both networks. The PMA is treated as the storage of motor patterns and voluntary activities and is involved in the expressive language of lips and tongue areas and writing and sign language of hand and arm areas (Faw, 2002). The connection LPMA → SMA is common for both tasks, indicating the voluntary movements involved in counting task (Wu et al., 2004).

The parietal lobe generally performs the function of processing and discriminating of sensory inputs (Kolb and Whishaw, 1996). The activation in LIPL or supramarginal gyrus (BA40) observed in this experiment has been linked to memories of visual word forms of the language system and is likely to be associated with arithmetic computing (counting) and language processing (reading). As seen in Fig. 3, the LIPLs in both networks send the representations of the inputs to the medial regions, AFG, MMFG, VIFG, and SMA, which are mainly involved in the counting function. The differences are seen as the extra activations in the language areas of the interference network: the connections from LIPL to LLIFG (BA44) and LSPL (BA7); as well as the connection for processing Stroop-related conflict and resolving interference effects: LIPL to LLMFG (BA9). The connections from LIPL to RSPL (BA7) and RLIFG (BA44) are seen only in the neutral network; this may account for a compensation function for the absence of language pathways present in the interference network and is likely to be involved in the visualization of symbols instead of reading, i.e., "automatic speech", where the right hemisphere is subserving residual aphasia speech (Vanlancker-Sidtis et al., 2003). The LSPL (BA7) was activated only in the interference counting task and has connections to the regions, LLIFG and VIFG (BA47); the RSPL was activated in both tasks and connected to the RLMFG (BA9), while the connections from RSPL to SMA (BA6) and ACC (BA24) are found only in the interference network.

In summary, the structures involved in both tasks are mostly common, and the differences are mainly due to the specific language areas activated in the interference counting task. Connections present only in the interference network (such as LIPL → LLMFG → LLIFG) are part of the language pathway, thus performing phonetic and semantic analysis and decision. Meanwhile, connections found only in the neutral network (such as LIPL → RLIFG) may perform a compensatory function for the non-activated functions corresponding to the connections, LIPL → LLIFG and LIPL → SPL, present in the interference network. In addition, since the interactions between two regions were allowed to be bi-directional, some connections are seen reversed between the two networks such as MMFG → RSPL in the interference network versus RSPL → MMFG in the neutral network.

Discussion

Earlier approaches to neural systems analysis, such as SEM, DCM, and GCM, are confirmatory; a researcher is more likely to use them to determine whether a previously known or hypothesized neural system model is valid rather than to "find" a suitable model from the data (Maruyama, 1989). The structures of those models were constrained by the prior models derived from previous studies or by anatomical constraints, although the exact model for the experiment under consideration is often unknown. Our method investigated the use of Bayesian networks to learn large or unexplored cognitive networks from fMRI data by assuming that the basis of such networks does not have proper prior models.

In SEM, effective connectivity was explored using path coefficients indicating the covariances among regions (Bullmore et al., 2000). The present approach uses conditional probability densities in graphical models to determine the structure of a functional network. In contrast to the second-order models, such as SEM, the connections between the regions in the present approach were derived by considering CPDs describing the behavior of a network in the complete statistical sense, which renders more information about the effective connectivity. The results on the synthetic data showed that the Bayesian networks can better fit the functional imaging data than the covariance-based models. The connectivity analysis by GCM is voxel-wise; in contrast, our approach is region-wise and seeks for a global representation of a neural system. Both DCM and the present approach make inferences about the connectivity of the network in the Bayesian framework, therefore, there are no limits on the number of connections that can be modeled without an overfitting problem. However, the DCM analyzes interactions at the neuronal rather than the hemodynamic level, which is more useful in analyzing the temporal interactions among brain regions. Instead, our approach focuses on exploring the static structure of interactions of the neural systems.

The complexity of the brain makes it difficult to be explored, especially in higher cognitive tasks; the analysis of functional integration (functional connectivity and effective connectivity) is still far from full understanding. The proposed method of exploring global neural systems from functional imaging data provides an alternate method to study brain function in terms of networks. The networks derived from our method for silent reading and Stroop tasks were consistent with the literature, providing a partial validation of our approach though the gold standard of the networks of the tasks considered is unavailable. In the silent reading task, the network demonstrated that the dominance of language processing in the left hemisphere and the regions in the right hemisphere receives the effects of processing from the left hemisphere. The interference network derived showed the involvement of language areas in the interference counting task compared to the neutral counting task.

The structure of the present functional brain network was determined from the data by the present method in a completely exploratory manner. As seen in the experiments with synthetic data, the method was robust to random noise and outperformed SEM in determining the structure. The MCMC algorithm searches the DAG space and returns a sample of structures after search-and-score learning. We choose the structure with the highest score as it matches the data the best. This may not always be the best choice because of possible local minima. The simulations with synthetic data showed that, as the number of region increases, the search has a

higher rate of falling into local minima. However, this problem can be mitigated even if partial a priori knowledge of the regions of activation or their connectivity is available. A compromise between confirmatory and exploratory approaches might be more appropriate for analysis of brain connectivity.

As illustrated in the experiment of Stroop task, the present method offers the feasibility of comparing the differences how brain regions interact in realizing the different tasks. This could be extended to differentiate the performance of patients and healthy participants performing the same cognitive tasks and explore disconnectivity hypotheses in brain disease. A major advantage of Bayesian networks might be its ability to infer network function in the case of brain disorders as inferencing is a strength of the graphical models. The main objective of the present work is to determine the existence of significant interactions among brain regions. Estimating the strengths of these interactions and exploring the behavior of such networks due to an abnormal event such as a stroke remain as future work.

Uncited references

Geiger, 1994
Genovese et al., 2002
He et al., 2003
SPM99
Zoubin, 2004

Acknowledgments

The authors wish to thank Dr. Susanta Mukhopadhyay, Dr. Wang Yang, and anonymous reviewers for their comments which greatly improved the quality of the manuscript. This work was supported by a grant to J.C. Rajapake, jointly by the Ministry of Education (MOE) and the Agency of Science, Technology and Research (A*Star), Singapore. The data sets were supported by fMRIDC (The fMRI Data Center, Dartmouth College, http://www.fmridc.org, accession number: 2-2000-11189 and 2-2001-1123b).

References

Bavelier, D., Tomann, A., Hutton, C., Mitchell, T., Corina, D., Liu, G., Neville, H., 2000. Visual attention to the periphery is enhanced in congenitally deaf individuals. J. Neurosci. 20 (RC93), 1–6.

BrainPlace.com, Brain SPECT Information and Resources. Brain Systems, Functions and Problems. http://www.brainplace.com/bp/brainsystem/.

Brett, M., 2002. The MNI brain and the Talairach atlas. http://www.mrc-cbu.cam.ac.uk/Imaging/Common/mnispace.shtml.

Buchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. Cereb. Cortex 7, 768–778.

Bullmore, E., Brammer, M., Rabe-Hesketh, S., Curtis, V., Morris, R., Williams, S., Sharma, T., McGuire, P., 1999. Methods for diagnosis and treatment of stimulus-correlated motion in generic brain activation studies using fMRI. Hum. Brain Mapp. 7 (1), 38–48.

Bullmore, E., Horwitz, B., Honey, G., Brammer, M., Williams, S., Sharma, T., 2000. How good is good enough in path analysis of fMRI data? NeuroImage 11, 289–301.

Bush, G., Whalen, P.J., Rosen, B.R., Jenike, M.A., McInerney, S.C., Rauch, S.L., 1998. The counting Stroop: an interference task specialized for functional neuroimaging validation study with functional MRI. Hum. Brain Mapp. 6, 270–282.

Faw, B., 2002. Pre-frontal executive committee for perception, working memory, attention, long-term memory, motor control, and thinking: a tutorial review. Conscious. Cogn. 12, 83–139.

Fiebach, C.J., Friederici, A.D., Mu¡§ller, K., Cramon, D.Y.V., 2002. fMRI evidence for dual routes to the mental lexicon in visual word recognition. J. Cogn. Neurosci. 14 (1), 11–23.

Field, A.S., Yen, Y.F., Burdette, J.H., Elster, A.D., 2000. False cerebral activation on BOLD functional MR images: study of low-amplitude motion weakly correlated to stimulus. Am. J. Neuroradiol. 21 (8), 1388–1396.

fMRIDC, Dartmouth College. The fMRI Data Center. http://www.fmridc.org.

Friston, K.J., 2003. Dynamic causal modelling. NeuroImage 19, 1273–1302.

Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. Hum. Brain Mapp. 2, 189–210.

Friston, K.J., Williams, S., Howard, R., Frackowiak, R.S., Turner, R., 1996. Movement-related effects in fMRI time-series. Magn. Reson. Med. 35 (3), 346–355.

Friston, K.J., Holmes, A.P., Worsley, K.J., 1999. How many subjects constitute a study? NeuroImage 10, 1–5.

Gavrilescu, M., Stuart, G.W., Waites, A., Jackson, G., Svalbe, I.D., Egan, G.F., 2004. Changes in effective connectivity models in the presence of task-correlated motion: an fMRI study. Hum. Brain Mapp. 21 (2), 49–63.

Geiger, D., 1994. Learning Gaussian Networks. Technical Report, Microsoft Research, Advanced Technology Division, Microsoft Corporation MSR-TR-94-10.

Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. NeuroImage 15, 870–878.

Goebel, R., Roebroecka, A., Kimb, D., Formisanoa, E., 2003. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magn. Reson. Imaging 21, 1251–1261.

Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., Gore, J.C., 2002. Detection of functional connectivity using temporal correlations in MR images. Hum. Brain Mapp. 15, 247–262.

Hayward, G., Goodwin, G.M., Harmer, C.J., McCharley, R.W., 2004. The role of the anterior cingulate cortex in the counting Stroop task. Exp. Brain Res. 154, 355–358.

He, A.G., Tan, L.H., Tang, Y., James, G.A., Wright, P., Eckert, M.A., Fox, P.T., Liu, Y., 2003. Modulation of neural connectivity during tongue movement and reading. Hum. Brain Mapp. 18, 222–232.

Heckerman, D., Geiger, D., November 1995. Likelihoods and parameter priors for Bayesian networks. Technical Report MSR-TR-95-54, MicroSoft Research, Redmond, WA.

Honey, G.D., Fu, C.H.Y., Kim, J., Brammer, M.J., Groudace, T.J., Suckling, J., Pich, E.M., William, S.C.R., Bullmore, E.T., 2002. Effects of verbal working memory load on corticocortial connectivity modeled by path analysis of functional magnetic resonance imaging data. NeuroImage 17, 573–582.

Horwitz, B., Braun, A.R., 2003. Brain network interactions in auditory, visual and linguistic processing. Brain Lang.

Horwitz, B., Rumsey, J.M., Donohue, B.C., 1998. Functional connectivity of the angular gyrus in normal reading and dyslexia. Neurobiology 95, 8939–8944.

Jordan, M.I., 1999. Learning in Graphical Models. The MIT Press.

Kass, R., Raftery, A., 1995. Bayes factors. J. Am. Stat. Assoc. 90, 773–795.

Kim, K.H.S., Relkin, N.R., Lee, K., Hirsch, J., 1997. Distinct cortical areas associated with native and second languages. Nature, 171–174.

Kolb, B., Whishaw, I.Q., 1996. Fundamental of Human Neuropsychology. W. H. Freman and Company.

Krause, B.J., Horwitz, B., Taylor, J.G., Schmidt, D., Mottaghy, F.M., Herzog, H., Halsband, U., Muller-Gartner, H.W., 1999. Network analysis in episodia encoding and retrieval of word-pair associates: a PET study. J. Neurosci. 11, 3293–3301.

Lancaster, J.L., Fox, P.T., Mikiten, S., Rainey, L., 2004. Talairach Deamon Database. http://ric.uthscsa.edu/projects/talairachdaemon.html.

Maintz, J.B., Viergever, M.A., 1996. An overview of medical image registration methods. Imaging Science Department, Imaging Center Utrecht, Technical Report, UU-CS-1998-22.

Margaritis, D., 2003. Learning Bayesian network model structure from data. Thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh.

Maruyama, G.M., 1989. Basics of Structural Equations Modeling. Comput. John Wiley and Sons, Inc.

Matsumoto, R., Nair, D.R., LaPresto, E., Najm, I., Bingaman, W., Shibasaki, H., Luders, H.O., 2004. Functional connectivity in the human language system: a cortico-cortical evoked potential study. Brain 127, 2316–2330.

Maurer, C.R., Fitzpatrik, J., 1993. A review of medical image registration. Am. Assoc. Neurol., 17–44.

McIntosh, A.R., Gonzalez-Lima, F., 1994. Structural equation modeling and its application to network analysis in functional brain imaging. Hum. Brain Mapp. 2, 2–22.

McIntosh, A.R., Grady, C.L., Ungerieider, L.G., Haxby, J.V., Rapoport, S.I., Horwitz, B., 1994. Network analysis of cortical visual pathways mapped with PET. J. Neurosci. 14 (2), 655–666.

McKiernan, K.A., Conant, L.L., Chen, A., Binder, J.R., 2001. Development and cross-validation of a model of linguistic processing using neural network and path analyses with FMRI data. NeuroImage 13 (6), 2.

Mechelli, A., Friston, K.J., Price, C.J., 2000. The effects of presentation rate during word and pseudoword read: a comparison of PET and fMRI. Cogn. Neurosci. 12, 145–156.

Mechelli, A., Penny, W.D., Price, C.J., Gitelman, D.R., Friston, K.J., 2002. Effective connectivity and intersubject variability: using a multisubject network to test differences and commonalities. NeuroImage 17, 1459–1469.

Murphy, K.P., 2002. Dynamic Bayesian networks: representation, inference and learning. PhD thesis, UC Berkeley, Computer Science Division, http://www.ai.mit.edu/murphyk/Thesis/thesis.html.

Murphy, K.P., 2004. The Bayes net toolbox for Matlab. http://www.ai.mit.edu/murphyk/Software/BNTsage.html.

Nezafat, R., Shadmehr, R., Holcomb, H.H., 2001. Long-term adaptation to dynamics of reading movement: a PET study. Exp. Brain Res. 140, 66–76.

Nyberg, L., McIntosh, A.R., Cabeza, C., Nilsson, L., Houle, S., Habib, R., Tulving, E., 1996. Network analysis of positron emission tomography regional cerebral blood flow data: ensemble inhibition during episodic memory retrieval. J. Neurosci. 16 (11), 3753–3759.

Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004a. Modelling functional integration: a comparison of structural equation and dynamic causal models. NeuroImage 23, s264–s274.

Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004b. Comparing dynamic causal models. NeuroImage 22, 1157–1172.

Petersson, K.L., Reis, A., Askelof, S., Castro-Caldas, A., Ingvar, M., 2000. Language processing modulated by literacy: a network analysis of verbal repetition in literate and illiterate subjects. Congnitive Neuroscience 12 (3), 364–382.

Price, C.J., 2000. The anatomy of language: contributions from functional neuroimaging. J. Anat. 179, 335–359.

Shin, L.M., Whalen, P.J., Pitman, R.K., Bush, G., Macklin, M.L., Lasko, N.B., Orr, S.P., McInerney, S.C., Rauch, S.L., 2001. An fMRI study of anterior cingulate function in posttraumatic stress disorder. Soc. Biol. Psychiatry 0006-3223 (01) (01215–X).

SPM99, Wellcome Department of Cognitive Neurology, London, UK. Statistical Parametric Mapping. http://www.fil.ion.ucl.ac.uk/spm/.

Talairach, J., Tournoux, P., 1988. Co-planar Stereotaxic Atlas of the Human Brain. Thieme Medical Publishers, Inc.

Tamm, L., Menon, V., Johnston, C.K., Hessl, D.R., Reiss, A.L., 2002. fMRI study of cognitive interference processing in females with fragile X syndrome. J. Cogn. Neurosci. 14 (2), 160–171.

Tan, L.H., Spinks, J.A., Gao, J., Liu, H., Perfetti, C.A., Xiong, J.H., Stofer, K.A., Pu, Y., Liu, Y., Fox, P.T., 2000. Brain activation in the processing of Chinese characters and words: a functional MRI study. Hum. Brain Mapp. 10, 16–27.

Vanlancker-Sidtis, D., McIntosh, A.R., Grafton, S., 2003. PET activation studies comparing two speech tasks widely used in surgical mapping. Brain Lang. 85, 245–261.

Wesley, J.B., 1994. Statistical Analysis for Engineers and Scientists. The McGraw-Hill, Inc.

Worsley, K.J., Friston, K.J., 1995. Analysis of fMRI time-series revisited-again. NeuroImage 2, 173–181.

Worsley, K.J., Marrett, S., Neelin, P., Vandal, A.C., Friston, K.J., Evans, A.C., 1996. A united statistical approach for determining significant signals in images of cerebral activation. Hum. Brain Mapp. 4, 58–73.

Worsley, K.J., Taylor, J.E., Tomaiuolo, F., Lerch, J., 2004. Unified univariate and multivariate random field theory. NeuroImage 23, s189–s195.

Wu, T., Kansaku, K., Hallett, M., 2004. How self-initiated memorized movements become automatic: a functional MRI study. Neurophysiology 91, 1690–1698.

Zoubin, G., 2004. Graphical models: parameter learning. Gatsby Computational Neuroscience Unit, University College London, http://www.gatsby.ucl.ac.uk/zoubin/.