

Multivariate Autoregressive Modelling of fMRI
time series

L. Harrison, W.D. Penny and K. Friston

Wellcome Department of Imaging Neuroscience,
University College London

Tel: 02078337472

Fax: 02078131420

E-mail: lharris,wpenny,karl@fil.ion.ucl.ac.uk

March 14, 2003

Abstract

We propose the use of Multivariate Autoregressive (MAR) models of fMRI time series to make inferences about functional integration within the human brain. The method is demonstrated with synthetic and real data showing how such models are able to characterise inter-regional dependence. We extend linear MAR models to accommodate nonlinear interactions to model top-down modulatory processes with bilinear terms. MAR models are time series models and thereby model temporal order within measured brain activity. A further benefit of the MAR approach is that connectivity maps may contain loops, yet exact inference can proceed within a linear framework. Model order selection and parameter estimation are implemented using Bayesian methods.

1 Introduction

Functional neuroimaging has been used to corroborate functional specialisation as a principle of organization in the human brain. However, disparate regions of the brain do not operate in isolation and more recently neuroimaging has been used to characterise the network properties of the brain under specific cognitive states [4] [8]. These studies address a complementary principle of organization, functional integration.

Functional MRI provides a unique opportunity to observe simultaneous recordings of activity throughout the brain evoked by cognitive and sensorimotor challenges. Each voxel within the brain is represented by a time series of neurophysiological activity that underlies the measured BOLD response. Given these multivariate, voxel-based time series, can we infer large-scale network behaviour among functionally specialised regions? In order to answer this question models are needed to describe the underlying connectivities implied by the data.

Effective connectivity is defined as the influence a neuron (or neuronal population) has on another. At the neuronal level this is equivalent to the effect pre-synaptic activity has on post-synaptic response, otherwise known as synaptic efficacy. Models of effective connectivity are designed to identify a suitable metric of influence among interconnected components (or regions of interest) in the brain. The notion of inferring influence from recorded data is, however, much more general. Consider the trajectory of an object as the result of external forces acting on it. These forces, which may be represented by equations of motion, determine the object's path. The equations of motion are an example of a model of the physical system. The observed data can be understood and analysed using this model. There are many different approaches to modeling a dynamic physical system, however, the motivation is the same for all: to identify operational principles responsible for generating the data.

There are two main approaches to modeling dynamic systems (e.g. physical bodies acted on by external forces or neuronal firing within a network), which can be used to understand spatial and/or temporal order within measured data, such as functional imaging data. These include equations of motion (as above), alternatively, we may model the systems behaviour by simply quantifying relationships within the measured data only. The first approach includes State-Space Models, e.g. used by the Kalman Filter, while the second includes simple regression analysis and convolution models (such as Volterra approach) to identify statistical dependencies, or patterns, within the data [18].

Both approaches have been used to measure effective connectivity among cortical regions from neuroimaging data starting with regression models [22, 14, 16, 17], then input-output models [13, 15] and later State-Space Models [7]. Regression techniques, such as Psychophysiological Interactions, are advantageous as they are easy to solve, yet may be used to approximate nonlinear modulatory interactions [14]. However, this is at the expense of ignoring temporal information, i.e. the history of an input (experimental task) or physiological variable (imaging data). This is important as interactions within the brain, whether over short or long distances, take time and are not instantaneous (which is implicit within regression models). Furthermore, the instantaneous state of any brain system that conforms to a dynamical system, will depend on the history of its input. Structural equation Modeling (SEM), as used by the neuroimaging community [22, 5] has similar problems ¹. Input-output models, such as the Volterra approach, model temporal effects in terms of an idealized response characterized by the kernels of the model [12]. A criticism of the Volterra approach is that it treats the system as a black box, meaning that it has no model of the internal mechanisms that may generate the data. State-Space Models account

¹There exist versions of SEM that do model dynamic information, see [9] for details of Dynamic Factor Analysis.

for correlations within the data by invoking state variables, whose dynamics generate the data. Recursive algorithms, such as the Kalman Filter, may be used to estimate these states through time, given the data. This approach was used to estimate variable regression coefficients between V1 and V5 activity in [7].

The MAR model fits into this classification scheme as one that models temporal effects across different variables (e.g. regions of interest), without using state variables. They characterise inter-regional dependencies within the data, specifically in terms of the historical influence one variable has on another. This is distinct from regression techniques that quantify instantaneous correlations, yet, similar to the Volterra model in that the relative influences, over time, are estimated. These considerations have motivated the investigation of MAR models, which may in some instances, be suitable for making inferences about functional integration in fMRI.

The paper is divided into 3 sections. First we describe the theory of MAR models, parameter estimation, model order selection and statistical inference. We have used a Bayesian technique for model order selection and parameter estimation which is described fully in [27]. Second, we test the method with synthetic data before modelling real neurophysiological data taken from an fMRI experiment addressing attentional modulation of cortical connectivity during a visual motion task [6]. The modulatory effect of one region upon the responses to other regions is a second order interaction which is precluded in linear models. To circumvent this we have introduced bilinear terms [14]. We assess the ability of bilinear

MAR models to capture top-down modulatory effects of the prefrontal cortex (PFC) and posterior parietal cortex (PPC) on motion sensitive regions in the dorsal visual pathway during attention. In the final section we discuss the advantages of MAR models, its use in spectral estimation and future developments

of the Bayesian approach used to estimate MAR parameters.

2 Theory

2.1 Multivariate Autoregressive Models

Given a univariate time series, its consecutive measurements contain information about the process that generated it. An attempt at describing this underlying order can be achieved by modelling the current value of the variable as a weighted linear sum of its previous values. This is an Autoregressive (AR) process and is a very simple, yet effective, approach to time series characterisation. The order of the model is the number of preceding observations used and the weights are the parameters of the model estimated from the data, that uniquely characterise the time series.

Multivariate Autoregressive models extend this approach to multiple time series so that the vector of current values of all variables is modelled as a linear sum of previous activities. Consider d time series generated from d variables (brain regions) within a system such as a functional network in the brain and where p is the order of the model. Here the scalar p denotes order, however, later we will use $p(\alpha)$ to mean the probability of α . A MAR(p) model predicts the next value in a d -dimensional time series, y_n as a linear combination of the p previous vector values

$$y_n = \sum_{i=1}^p y_{n-i} A(i) + e_n \quad (1)$$

where $y_n = [y_n(1), y_n(2), \dots, y_n(d)]$ is the n th sample of a d -dimensional time series, each $A(i)$ is a d -by- d matrix of coefficients (weights) and $e_n = [e_n(1), e_n(2), \dots, e_n(d)]$ is additive Gaussian noise with zero mean and covariance R .

We have assumed that the data mean has been subtracted from the time series.

The model can be written in the standard form of a multivariate linear regression model as follows

$$y_n = x_n W + e_n \quad (2)$$

where $x_n = [y_{n-1}, y_{n-2}, \dots, y_{n-p}]$ are the p previous multivariate time series samples and W is a $(p \times d)$ -by- d matrix of MAR coefficients (weights). There are therefore a total of $k = p \times d \times d$ MAR coefficients. If the n th rows of Y , X and E are y_n , x_n and e_n respectively and there are $n = 1..N$ samples then we can write

$$Y = XW + E \quad (3)$$

where Y is an $(N - p)$ -by- d matrix, X is an $(N - p)$ -by- $(p \times d)$ matrix and E is an $(N - p)$ -by- d matrix. The number of rows $N - p$ (rather than N) arises as samples at time points before p do not have sufficient preceding samples to allow prediction.

MAR models quantify the linear dependence of one region upon all others in the network. The weights in W can be interpreted as characterising the influence each region has upon it. Independence between a pair of regions results in a weight of zero while dependence is reflected in a non-zero magnitude.

A schematic representation of equation 3 is shown in Figure 1. Figure 1a shows the original d -dimensional time series (Y) modelled as a MAR process (XW) plus residual error (E).

W characterises the d -dimensional series as a network of connection strengths between all possible pairs of elements in the original series. Figure 1b is a schematic of W , which consists of p layers, one for each time lag used in the model. Each layer is a $d \times d$ matrix of weights (shown as the squares along the bottom of the figure, again, one for each time lag). The diagonal entries in these are ‘‘self-connections’’ and the off-diagonals are connections between regions. Any dependence among elements in the d -dimensional time series (brain

regions) is reflected in non-zero off-diagonal coefficients.

The use of MAR models for characterising networks of cortical activity associated with cognitive tasks allows one to quantify the dependence among all possible combinations of pairs of regions in the model. This way, connectivity architectures can be compared across different cognitive tasks. For example, the instruction to attend or not to a stimulus will induce plastic changes in connectivity, which may be quantified and compared using MAR models.

2.2 Nonlinear autoregressive models

Given a network coupling model of the brain we can think of two fundamentally different types of coupling; linear and nonlinear. The model discussed so far is linear. Linear systems are described by the principle of superposition, which is that inputs have additive effects upon the response that are independent of each other. However, if the inputs interact to produce a response, the response can no longer be described by a linear combination of the inputs. This is an example of a nonlinear interaction.

In [6] such nonlinear interactions have been modelled by making use of bilinear terms and is the approach adopted here. Specifically, to model a hypothesized interaction between variables $y_n(j)$ and $y_n(k)$ one can form the new variable

$$I_n(j, k) = y_n(j)y_n(k) \quad (4)$$

This is the bilinear variable. This is then orthogonalised with respect to the original time series. We then form an augmented MAR model with an extra time series and augmented connectivity matrices $\tilde{A}(i)$.

$$[y_n, I_n(j, k)] = \sum_{i=1}^p [y_{n-i}, I_{n-i}(j, k)]\tilde{A}(i) + e_n \quad (5)$$

The relevant entries in $\tilde{A}(i)$ then reflect modulatory influences eg. a change of the connection strength between $y(j)$ and other time series due to the influence

of $y(k)$.

It should be noted that each bilinear variable introduces only one of many possible sources of nonlinear behaviour into the model. The example above specifically models nonlinear interactions between $y_n(j)$ and $y_n(k)$, however other bilinear terms could involve, for instance, the time series $y_n(j)$ and the inputs ($u(t)$). The inclusion of these terms are guided by the hypothesis of interest eg. does time change the connectivity between earlier and later stages of processing in the dorsal visual pathway? Here $u(t)$ would model time.

2.3 Maximum Likelihood Estimation

Reformulating MAR models as standard multivariate linear regression models allows us to retain contact with the large body of statistical literature devoted to this subject; see eg. Box and Tiao (p. 423) [2].

The Maximum Likelihood (ML) solution (see eg. [29]) for the MAR coefficients is

$$\hat{W} = (X^T X)^{-1} X^T Y \quad (6)$$

The maximum likelihood noise covariance, S_{ML} , can be estimated as

$$S_{ML} = \frac{1}{N-k} (Y - X\hat{W})^T (Y - X\hat{W}) \quad (7)$$

where $k = p \times d \times d$. We define $\hat{w} = \text{vec}(\hat{W})$ where vec denotes the columns of \hat{W} being stacked on top of each other (for more on the vec notation, see [24]).

To recover the matrix \hat{W} we simply ‘un-stack’ the columns from the vector \hat{w} .

The ML parameter covariance matrix for \hat{w} is given by [20] (page 321)

$$\hat{\Sigma} = S_{ML} \otimes (X^T X)^{-1} \quad (8)$$

where \otimes denotes the Kronecker product (see, eg page 477 in Box and Tiao [2])

The optimal value of p can be chosen using a model order selection criterion such as the Minimum Description Length (MDL). See eg. [25].

2.4 Bayesian Estimation

It is also possible to estimate the MAR parameters and select the optimal model within a Bayesian framework [27]. This has been shown to give better model order selection and is the approach used in this paper. The maximum-likelihood solution is used to initialise the Bayesian scheme.

In what follows $\mathbf{N}(m, Q)$ is a multivariate Gaussian distribution with mean m and precision (inverse covariance) Q . Also, $\text{Ga}(b, c)$ is a Gamma distribution with parameters b and c . The gamma density has mean bc and variance b^2c . Finally, $\text{Wi}(s, B)$ denotes a Wishart density [2]. The Bayesian model uses the following prior distributions

$$\begin{aligned} p(W|p) &= \mathbf{N}(0, \alpha I) \\ p(\alpha|p) &= \text{Ga}(b, c) \\ p(\Lambda|p) &= |\Lambda|^{-(d+1)/2} \end{aligned} \tag{9}$$

where p is the order of the model, α is the precision of the Gaussian prior distribution from which weights are drawn and Λ is the noise precision matrix (inverse of R). In [27] it is shown that the corresponding posterior distributions are given by

$$\begin{aligned} p(W|Y, p) &= \mathbf{N}(\hat{W}_B, \hat{\Sigma}_B) \\ p(\alpha|Y, p) &= \text{Ga}(\hat{b}, \hat{c}) \\ p(\Lambda|Y, p) &= \text{Wi}(s, B) \end{aligned} \tag{10}$$

The parameters of the posteriors are updated in an iterative optimisation scheme described in the Appendix. Iteration stops when the ‘Bayesian evidence’ for model order p , $p(Y|p)$, is maximised. A formula for computing this is also provided in the appendix. Importantly, the evidence is also used as a model order selection criterion, that is, to select the optimal value of p .

2.5 Bayesian Inference

The Bayesian estimation procedures outlined above result in a posterior distribution for the MAR coefficients $P(W|Y, p)$. Bayesian inference can then take place using confidence intervals based on this posterior. See, for example, page 84 of Box and Tiao [2]. The posterior allows us to make inferences about the strength of a connection between two regions. Because this connectivity can be expressed over a number of time lags our inference is concerned with the vector of connection strengths, a , over all time lags. To make contact with classical (non-Bayesian) inference, we say that a connection is ‘significantly non-zero’ or simply ‘significant’ at level α if the zero vector lies outside the $1 - \alpha$ confidence region for a . This is shown schematically in Figure 2. We also refer to α as the ‘ p -value’ (see Appendix B).

3 Application

3.1 Synthetic data

To test the face validity of the method two sets of synthetic data were generated, which are shown in Figure 3. All time series were generated from known MAR(2) models. The known values were compared with estimates of model order and weights. Each data set contained six time series, the first being independent (Figure 3a) with all off diagonal MAR terms equal to zero while the second (Figure 3b) included two sets of three time series which were dependent within sets but independent between sets.

The upper right figure in Figure 3a shows the true MAR(2) structure from which time series were generated (lower left). The model has only non-zero diagonal terms and no covariance structure and thereby generates independent time series. Each series is essentially an AR(2) process. The simulated time

series were modelled as a $MAR(p)$ process, using Bayesian evidence to select the optimal order. This is shown in Figure 4a, demonstrating, as anticipated, an optimal model order of 2. Parameter estimates are shown for comparison on the right of Figure 3a. The matrix W is represented in two ways; the first (upper right) shows only the conditional means of the parameter estimates in the same format as the known $MAR(2)$ model. The general character of the known structure has been captured with dominant diagonal terms, however, off diagonal terms have many non-zero values. A more complete representation of the posterior distribution ($p(W|Y, p)$) is shown below which depicts the variance about the estimated means. Each plot within the matrix of graphs contains weight estimates at all time lags in the model for one connection. Zero is indicated and posterior distributions shown in relation to it. Those parameters which straddle zero, despite having a non-zero mean are not considered significant. All connections that are significantly non-zero are circled. The overall structure of the true parameters is reflected in the estimates with the exception of one connection (circled off diagonal coefficient).

The second set of synthetic data contained a mixture of dependence and independence and are shown in Figure 3b. The format is the same as Figure 3a. The known $MAR(2)$ model has two subgroups characterised by dependence within each subgroup (to the degree of the coefficients magnitude) and independence between subgroups (reflected in the zero coefficients). Modelling the time series (lower left) as a $MAR(p)$ and using Bayesian evidence (Figure 4a) the correct model order ($p = 2$) was identified. Parameter estimates (right) again reflect the known $MAR(2)$ structure.

The accuracy of model order selection using Bayesian evidence was generally stable however occasionally an incorrect order was calculated using smaller data sets (eg < 250). Increasing the number of data points to 500 produced robust and correct estimates in all cases. The stability of these results is worth noting

as estimates of the zero coefficients produced the most “false positives” (usually only one out of a possible thirty in the first data set and eighteen in the second) while “false negatives” occurred less frequently. Given these occasional discrepancies, the Bayesian framework for estimating model parameters was able to differentiate and quantify interdependence within a MAR process.

3.2 fMRI data

Attentional effects on the responsiveness of motion sensitive area V5 and PPC measured in electrophysiological and neuroimaging studies suggest attention is associated with changes in connectivity [1] [26]. In this paper we use data from a fMRI study investigating attentional modulation of connectivity within the dorsal visual pathways [6]. This provides a testbed for assessing how MAR models estimate changes in connectivity.

In brief, the experiment was performed on a 2T MRI scanner on several subjects. The visual stimulus involved random dots moving radially outwards at a fixed rate. Subjects were trained beforehand to detect changes in velocity of radial motion. Attentional set was manipulated by asking the subject to attend to changes in velocity or to just observe the motion. Both of these states were separated by periods of “fixation” where the screen was dark and only a fixation dot was visible. Each block ended with a “stationary” condition in which a static image of the previously moving dots was shown. Unknown to the subjects, the radial velocity remained constant throughout the experiment such that the only experimental manipulation being attentional set. Categorical comparisons using SPM(t) were used to identify changes in brain activity dependent on attentional set. This revealed activations throughout right and left hemispheres in the primary visual cortex V1/2 complex, visual motion region V5 and regions involved in the attentional network including posterior parietal cortex (PPC) and in the right prefrontal cortex (PFC).

Regions of interest (ROI) were defined with a diameter of 8mm centred around the most significant (< 0.05 , corrected) voxel and a representative time series was defined by the first eigenvariate of the region.

For details of the experimental design and acquisition see [6]. The time series of the right hemisphere regions, in one subject, are shown in Figure 4.

Inspecting the four time series reveals a number of characteristics worth noting. The series from the V1/2 complex show a dependence on the presentation of the moving image with a small difference between attention and non-attention. However, in the higher brain areas of PPC and PFC attentional set is the dominant influence, with a marked increase in activity during periods of attention. The relative influence each region has upon others, or indeed any non-additive interaction, is not obvious from visual inspection alone. Modelling the series as a MAR process provides a quantitative approach to these putative effects.

Three models were tested using the regions and bilinear terms shown below. Bilinear terms for interactions between V1/2 complex and PPC are written as $I_{v1,ppc}$, and regions V5 and PFC as $I_{v5,pfc}$. These time series were entered into bilinear MAR models. The interaction terms can be thought as “virtual” nodes in a network. Models 1 and 3 involved only right hemisphere PFC as no significant attention related activation in the left PFC was found.

- Model 1: V1/2, V5, PPC and PFC
- Model 2: V1/2, V5 and $I_{v1,ppc}$
- Model 3: V5, PPC and $I_{v5,pfc}$

The motivation for the first model was to ask a very general question: given time series of brain activity over the entire experiment, during all four states, is there, on average, a functional network connecting key regions in the visual and attentional systems? The second and third models were motivated by SEM

and Volterra analyses of the same data reported in [6] and [15] respectively. These different methodologies address the same issues of modulatory, or bilinear interactions, at different levels of the visual and attentional pathways and provide a convenient benchmark with which to validate the current approach. Models 2 and 3 therefore specifically address whether or not attentional influence from top-down regions could be mediated by second-order interactions. If there is no bilinear interaction among the regions then this should be reflected as zero-valued weights within a bilinear MAR

All pairwise connection strengths (MAR coefficients across all time lags) were tested separately across all time lags for significant differences from zero. This raises an important issue of how *a priori* knowledge of connectivity should determine the analytic strategy. By testing all connections we are essentially approaching the model with no prior knowledge. Alternatively a model led approach is directed by prior knowledge of connectivity, thereby only testing connections established by, for example, anatomical studies. Given the motivation of the first model, to test for average connectivities over all tasks, we used the former strategy.

The results of all connections that attained significance (p-values < 0.05) are shown for three subjects in Tables 1, 2 and 3. Connectivity maps have been used to illustrate the dependencies in Figure 6, 7 and 8 for models 1, 2 and 3 respectively. The width of the arrows are scaled to the p-values (thin for p-values between 0.05-0.001 and thick for values < 0.001). The optimal model order was selected using Bayesian evidence. Plots of which are shown for one subject and all three models in Figure 5b. The optimal model order was $p = 4$ for all

The first model characterises the network of connectivities on average over all attentional states. Table 1 reveals clear feedforward and backward connections within the network. All V1 to V5 connections reached very high significance

($p < 0.0004$) and for V5 to PPC connections this was true for two subjects ($p < 0.0009$). Top-down connections between PFC and PPC were demonstrated in all subjects ($p < 0.02$). A schematic of these results is shown in Figure 6 for one subject. Figure 6a is the weight matrix W shown as a 4×4 matrix with each element containing a representation of the posterior distribution ($p(W|Y, p)$) of weights for one connection at all time lags in the model, the same as in Figure 3. Given the hierarchical sequence of the regions the upper off-diagonal terms correspond to dependence ascending through the cortical hierarchy. For example the plots in position (1, 2) are the coefficients characterising the influence V1 has upon V5 at all time lags. In short, upper diagonal coefficients quantify the influence of forward connections. The lower off-diagonal terms complement these and characterise backward projections. The forward driving influence of V1 on V5, V5 on PPC and PPC on PFC is evident. Back projecting influences of PFC upon V5 and PPC are also shown. These are important observations, however they are limited in that they represent a linear characterisation that precludes attentional modulation.

The second and third models were designed to test whether attentional areas modulate coupling within the visual system. Bilinear terms $I_{v1,ppc}$ and $I_{v5,pfc}$ were included in Models 2 and 3 respectively. The results of model 2 are displayed in a similar fashion to model 1 in Table 2 and Figure 7.

Both left and right hemispheres were modelled separately and show consistent results in the left for the connection between $I_{v1,ppc}$ and V5 ($p < 0.05$).

This indicates that PPC changes the connection between V1 and V5. Only subject 1 reached significance on the right for this interaction. The bilinear effect is depicted in Figure 7 b and c, first (left figure) showing the bilinear variable as a “virtual” node and second (right figure) its implicit physiological interpretation. Given the posterior density of the connection strength, bilinear terms are seen to account for a significant component of the activity observed

within these regions.

The results of Model 3 are shown in Table 3 and Figure 8. Second order connections between $I_{v5,pfc}$ and PPC reach significance with p-values between 0.03 – 0.056 in two subjects. Figure 8b demonstrates a similar interaction as in Figure 7b, suggesting that attention may be mediated by PFC and modulation of the connection strength between V5 and PPC. That is, PFC changes how PPC responds to V5, the responsiveness being greater during attention than non-attention. This is a bilinear effect, similar to that found in model 2.

4 Discussion

We have proposed the use of MAR models for making inferences about functional integration using fMRI time series. One motivation for this is that the dominant model, used for making such inferences in the existing fMRI/PET literature, namely Structural Equation Modelling, as used in [22, 5], is not a time series model. Indeed, inferences are based solely on the instantaneous correlations between regions - if the time series were randomly permuted SEM would give the same results. Thus SEM throws away temporal information.

This deficiency is not shared by MAR models which are proper time series models. Further, MAR models may contain loops and self-connections yet parameter estimation can proceed in a purely linear framework ie. there is an analytic solution that can be found via linear algebra. In contradistinction, SEM models with loops require non-linear optimisation. The reason for this is that MAR models do not contain instantaneous connections. The between-region connectivity arises from connections between regions at different time-lags. Due to temporal persistence in the activity of each region (ie. the activities are similar from one sample to the next) this captures much the same effect - but in a computationally simpler manner.

Given that MAR models extract temporal information, how can this be interpreted? MAR models derive temporal information from the auto and cross-covariance function, which is used to estimate coefficients at different lags. The temporal profile of the coefficient estimates characterizes the temporal aspects of the dependencies. For example, the coefficients can alternate from positive to negative and decay with increasing lags, due to oscillatory interactions. As BOLD is a measurement of the haemodynamic response to neuronal processes, temporal information is smoothed, rendering the coefficients a summary of neuronal activity observed during the haemodynamic response. This may confound the interpretation of the exact timing of an interaction, however, general observations regarding the brain's response are possible. In Figure 6a the first and second coefficient of connectivity between PPC and PFC, both forward and backward, are similar. This suggests that on average, throughout the entire experiment, that there is an equivalent level of connectivity between the two regions, over a period of two time lags (approximately 6 seconds). The models in Figure 7 and 8 support the hypothesis of a modulatory interaction, modeled using a bilinear term, from PPC and PFC, whose effect becomes less prominent with time, i.e. up to four time lags (approximately 12 seconds).

In this paper we have used “off-the-shelf” MAR models in which every region is connected to every other region. Bayesian inferences about connections are then made on the basis of the estimated posterior distribution. This is in the spirit of how General Linear Models are used for characterising functional specialisation; all conceivable factors are placed in one large model and then different hypotheses are tested using t or F-contrasts [11]. We note that this approach is fundamentally different to the philosophy underlying SEM. In SEM, only a few connections are modelled and these are chosen on the basis of prior anatomical or functional knowledge. In cases where this knowledge is available this may be the preferred approach and, in the future, we envisage the use of

MAR models which are not fully connected (see Bayesian estimation below).

MAR models can be used for spectral estimation. In particular they enable parsimonious estimation of coherences (correlation at particular frequencies), partial coherences (the coherence between two time series after the effects of others have been taken into account), phase relationships [21] [23] and directed transfer functions [19]. MAR models have been used in this way to investigate functional integration from EEG and ECOG recordings [3]. This provides a link with a recent analysis of fMRI data [10] which looks for sets of voxels that are highly coherent. MAR models provide a parametric way of estimating this coherence, although in this paper we have reported the results in the time domain.

A further aspect of our “off-the-shelf” MAR models is that they capture only linear relationships between regions. Following Buchel and Friston [6], we have extended their capabilities by introducing bilinear terms. It is also possible to include further higher order terms, for instance, second-order interactions across different lags. Frequency domain characterisation of the resulting models would then allow us to report bi-spectra [28]. These describe the correlations between different frequencies which may be important for the study of functional integration [12].

A key aspect of our approach has been the use of a mature Bayesian estimation framework [27]. This has allowed us to select the optimal MAR model order. In future, the Bayesian approach could be greatly extended. In [27] we show how MAR coefficients can be placed into groups. For example, all of those connecting the same two regions could be placed in the same group. Different groups could then be associated with different prior precisions. Groups with prior means of zero and infinite prior precision in effect then specify the absence of a connection. In this way, we could design MAR models with sparse connectivities. More generally, prior means and precisions could be estimated

from data using a Bayesian framework that could be specified so as to include time series from multiple subjects. These further extensions will be the subject of subsequent papers.

A Bayesian estimation

Following the algorithm developed in [27], the parameters of the posterior distributions are updated iteratively as follows.

$$\begin{aligned}
\Lambda_D &= \hat{\Lambda} \otimes (X^T X) & (11) \\
\hat{\Sigma}_B &= (\Lambda_D + \hat{\alpha} I)^{-1} \\
\hat{W}_B &= \hat{\Sigma}_B \Lambda_D \hat{W} \\
\frac{1}{\hat{b}} &= \frac{1}{2} \hat{W}_B^T \hat{W}_B + \frac{1}{2} \text{Tr}(\hat{\Sigma}_B) + \frac{1}{b} \\
\hat{c} &= \frac{k}{2} + c \\
\hat{\alpha} &= \hat{b} \hat{c} \\
s &= N \\
B &= \frac{1}{2} (Y - X \hat{W}_B)^T (Y - X \hat{W}_B) \\
&\quad + \sum_n (I \otimes x_n) \hat{\Sigma}_B (I \otimes x_n)^T \\
\hat{\Lambda} &= s B^{-1}
\end{aligned}$$

The updates are initialised using the Maximum-Likelihood solution. Iteration terminates when the Bayesian evidence increase by less than 0.01%. The Bayesian evidence is computed as follows

$$\begin{aligned}
p(Y|p) &= \frac{N}{2} \log |B| - KL(p(W|p), p(W|Y, p)) & (12) \\
&\quad - KL(p(\alpha|p), p(\alpha|Y, p)) + \log \Gamma_d(N/2)
\end{aligned}$$

where $KL(p1, p2)$ denotes the Kullback-Liebler (KL) divergence between densities $p1$ and $p2$. Expressions for these are given in [27]. Essentially, the first term

in the above equation is an accuracy term and the KL terms act as a penalty for model complexity.

B Testing the significance of connections

The connectivity between two regions can be expressed over a number of time lags. Therefore, to see if the connectivity is significantly non-zero we make an inference about the vector of coefficients a , where each element of that vector is the value of a MAR coefficient at a different time lag. First we specify $(k \times k)$ ($k = p \times d \times d$) sparse matrix C such that

$$a = C^T w \quad (13)$$

returns the estimated weights for connections between the two regions of interest. For an MAR(p) model, this vector has p entries, one for each time-lag.

The probability distribution is given by $p(a) = \mathbf{N}(m, V)$ and is shown schematically in Figure 2. The mean and covariance are given by

$$\begin{aligned} m &= C^T \hat{w} \\ V &= C^T \hat{\Sigma}_B C \end{aligned} \quad (14)$$

where $\hat{w} = \text{vec}(\hat{W}_B)$ and $\hat{\Sigma}_B$ are the Bayesian estimates of the parameters of the posterior distribution of regression coefficients from the previous section. In fact, $p(a)$ is just that part of $p(w)$ that we are interested in.

The probability α that the zero vector lies on the $1 - \alpha$ confidence region for this distribution is then computed as follows. We first note that this probability is the same as the probability that the vector m lies on the edge of the $1 - \alpha$ region for the distribution $\mathbf{N}(0, V)$.

This latter probability can be computed by forming the test statistic

$$d = m^T V^{-1} m \quad (15)$$

which will be the sum of $r = \text{rank}(V)$ independent, squared Gaussian variables.

As such it has a χ^2 distribution

$$p(d) = \chi^2(r) \tag{16}$$

Acknowledgements

This work was funded by the Wellcome Trust

References

- [1] J.A. Assad and R.M. Maunsell. Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature*, 1995.
- [2] G.E.P Box and G.C. Tiao. *Bayesian Inference in Statistical Analysis*. John Wiley, 1992.
- [3] S.L. Bressler, M. Ding, and W. Yang. Investigation of cooperative cortical dynamics by multivariate autoregressive modeling of event-related local field potentials. *Neurocomputing*, 26-27:625–631, 1999.
- [4] C. Buchel and K.J. Friston. Characterizing functional integration. In R.S.J. Frackowiak, K.J. Friston, C.D. Frith, R.J. Dolan, and J.C. Mazziotta, editors, *Human Brain Function*, pages 127–140. Academic Press USA, 1997.
- [5] C. Buchel and K.J. Friston. Modulation of connectivity in visual pathways by attention: Cortical interactions evaluated with structural equation modelling and fMRI. *Cerebral Cortex*, 7:768–778, 1997.
- [6] C. Buchel and K.J. Friston. Modulation of connectivity in visual pathways by attention: Cortical interactions evaluated with structural equation modelling and fMRI. *Cerebral Cortex*, 7:768–778, 1997.

- [7] C. Buchel and K.J. Friston. Dynamic changes in effective connectivity characterized by variable parameter regression and kalman filtering. *Human Brain Mapping*, 6:403–408, 1998.
- [8] C. Buchel and K.J. Friston. Assessing interactions among neuronal systems using functional neuroimaging. *Neural Networks*, 13:871–882, 2000.
- [9] R. Cudeck. *Structural Equation Modeling: Present and Future*. 2002.
- [10] K. Muller et al. On multivariate spectral analysis of fmri time series. *Neuroimage*, 2001.
- [11] R.S.J. Frackowiak, K.J. Friston, C.D. Frith, R.J. Dolan, and J.C. Mazziotta, editors. *Human Brain Function*. Academic Press USA, 1997.
- [12] K.J. Friston. The labile brain.I. neuronal transients and nonlinear coupling. *Phil Trans R Soc London B*, 355:215–236, 2000.
- [13] K.J. Friston. Brain function, nonlinear coupling, and neuronal transients. *The Neuroscientist*, 2001.
- [14] K.J. Friston, C. Beuchel, G.R. Fink, J. Morris, E. Rolls, and R.J. Dolan. Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage*, 6:218–229, 1997.
- [15] K.J. Friston and C. Buchel. Attentional modulation of effective connectivity from V2 to V5/MT in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 97(13):7591–7596, 2000.
- [16] K.J. Friston, C.D. Frith, and R.S.J. Frackowiak. Time-dependent changes in effective connectivity measured with PET. *Human Brain Mapping*, 1:69–79, 1993.

- [17] K.J. Friston, L.G. Ungerleider, P. Jezzard, and R. Turner. Characterizing modulatory interactions between v1 and v2 in human cortex: A new treatment of functional MRI data. *Human Brain Mapping*, 2:211–224, 1995.
- [18] J.N Juang. *Identification and control of mechanical systems*. Cambridge University Press, 2001.
- [19] M. Kaminski, K. Blinowska, and W. Szelenberger. Topographic Analysis of coherence and propagation of EEG activity during sleep and wakefulness. *Electroencephalography and Clinical Neurophysiology*, 102:216–227, 1997.
- [20] J.R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. John Wiley, 1997.
- [21] S.L. Marple. *Digital spectral analysis with applications*. Prentice-Hall, 1987.
- [22] A.R McIntosh, C.L Grady, L.G Underleider, J.V Haxby, S.I Rapoport, and B. Horwitz. Network analysis of cortical visual pathways mapped with pet. *Journal of Neuroscience*, 1994.
- [23] M.J.Cassidy and P.Brown. Stationary and non-stationary autoregressive models for electrophysiological signal analysis and functional coupling studies. *preprint*, 2002.
- [24] R.J. Muirhead. *Aspects of Multivariate Statistical Theory*. John Wiley, 1982.
- [25] A. Neumaier and T. Schneider. Estimation of parameters and eigenmodes of multivariate autoregressive models. *Submitted to ACM Trans. Math. Softw.*, 2000.
- [26] K.M. O’Craven and R.L. Savoy. Voluntary attention can modulate fmri activity in human mt/mst. *Investigational Ophthalmological Vision Science*, 1995.

- [27] W.D. Penny and S.J. Roberts. Bayesian multivariate autoregressive models with structured priors. *IEE Proc.-Vis Image Signal Process*, 149(1):33–41, 2002.
- [28] M.B Priestley. *Nonlinear and non-stationary time series analysis*. Harcourt Brace Jovanovich, 1988.
- [29] S. Weisberg. *Applied Linear Regression*. John Wiley, 1980.

Table Captions

1. P-values for testing individual connection strengths across all time lags for Model 1 for all three subjects, right hemisphere only. The table is divided in two with the upper part showing values for ascending connections (bottom-up) and the lower part descending connections (top-down) within the visual and attentional systems. The most significant coefficients are shown schematically for subject 1 in Figure 6. (* for p-values between 0.05-0.001, ** for p-values < 0.001 , – for estimates not significantly different from zero ($p > 0.05$)).
2. P-values for connection strengths across all time lags for Model 2 for all three subjects. Left and right hemispheres were modelled separately. The table includes top-down connections in both hemispheres. The bilinear term $I_{v1,ppc}$ introduces second order interactions among PPC, V1 and V5. There is a consistent dependence between $I_{v1,ppc}$ and V5 in the left hemisphere for all three subjects. Diagrams of the weight matrix W and connectivity maps are shown in Figure 7.
3. Similar layout for Model 3, right hemisphere only. Second order interactions are modelled between $I_{v5,pfc}$ and PPC and show significant coefficients in two of the three subjects. Weight matrix and connectivity maps are shown in Figure 8.

Table 1:

Direction	regions	subject 1	subject 2	subject 3
Bottom-up	v1-v5	(0.0000**)	(0.0004**)	(0.0000**)
	v1-ppc	(0.0000**)	–	–
	v1-pfc	–	(0.007*)	–
	v5-ppc	(0.0002**)	(0.0009**)	–
	ppc-pfc	(0.0182*)	–	–
Top-down	pfc-v1	–	–	(0.0045*)
	pfc-v5	(0.0005**)	–	–
	pfc-ppc	(0.0197*)	(0.0017*)	(0.0001**)

Table 2:

Direction	regions	subject 1	subject 2	subject 3
Top-down	$I_{v1,ppc-v5}$ right	(0.0205*)	–	–
	$I_{v1,ppc-v5}$ left	(0.0384*)	(0.0568)	(0.0474*)

Table 3:

Direction	regions	subject 1	subject 2	subject 3
Top-down	$I_{v5,pfc-ppc}$	(0.056)	(0.03*)	–

Figure Captions

Figure 1: A schematic of equation 3 (main text). (a) The original d -dimensional time series Y is modelled as a MAR process (XW) plus residual error (E). (b) W is a matrix containing all the weights which characterise the interactions among the elements of d . It consists of p layers (one for each time lag used in the model), each layer containing a $d \times d$ matrix of weights. The p layers of W have been placed in sequential order at the bottom of the figure. The diagonal elements are “self-connections” while the off diagonals reflect the dependence between different variables in the original d -dimensional series. W can be used to compare network properties associated with different cognitive tasks.

Figure 2: For a MAR(2) model the vector of connection strengths, a , between two regions consists of two values, $a(1)$ and $a(2)$. The probability distribution over a can be computed from the posterior distribution of MAR coefficients as shown in Appendix B and is given by $p(a) = N(m, V)$. Connectivity between two regions is then deemed significant at level $alpha$ if the zero-vector lies on the $1 - \alpha$ confidence region. The figure shows an example $1 - \alpha$ confidence region for a MAR(2) model.

Figure 3: Two synthetic data sets are generated from known MAR(2) models (a and b). Both 3a and 3b have the same layout with the known MAR(2) model (upper left) from which time series are generated (sample in lower left). See Figure 4a for a plot of the Bayesian evidence verses model order identifying a model order of $p = 2$. Parameter estimates are shown on the right in two formats. The first (upper right) is a plot of estimated means of weighting parameters for comparison with the known MAR model. Below this is a more comprehensive representation of the posterior distribution ($p(W|Y, p)$) of the estimates. There are 6×6 plots with mean and variance (two standard deviations) shown at all time lags (2) in relation to zero. Distributions that do not straddle zero are significantly non-zero. Figures a and b differ in that the known MAR(2) models used in 3a produce largely independent series whereas 3b generates mixed dependences between series. Estimates that are significantly different (< 0.05) from zero across both time lags are circled and reveal the same MAR(2) underlying process that generated the series originally.

Figure 4: These are the representative time series of regions V1/2 complex, V5 and PPC and PFC from one subject, in the right hemisphere. All plots have the same axes of activity (adjusted to zero mean and unit variance) vs scan number (360 in total). The experiment consisted of four conditions in four blocks of 90 scans. Periods of “attention” and “non-attention” were separated by a “fixation” interval where the screen was dark and the subject fixated on a central cross, and each block ended with a “stationary” condition where the screen contained a freeze frame of the previously moving dots. Epochs of each task are indicated by the background grayscale (see key) of each series. Visually evoked activity is dominant in the lower regions of the V1/2 complex whereas attentional set becomes the prevalent influence in higher regions PPC and PFC.

Figure 5: Model order selection using the Bayesian evidence versus model order p . Figure 5a contains two plots calculated for the synthetic data described in Figure 3. In both data sets (independent and mixed dependence) the optimal order was identified as $p = 2$. Figure 5b shows three plots generated from one subject for all three models of the real data described in Figure 5. The optimal order was $p = 4$ for each subject.

Figure 6: Model 1 including regions V1/2 complex, V5, PPC and PFC. (a) W from a MAR(4) model, with each region indicated on the rows and columns of the diagram. Each element within the matrix of plots contains means and variances of parameter estimates at all time lags for that connection (same as Figure 3). Diagonals are “self connections” and given the order of variables the *upper* off-diagonal coefficients represent ascending, feed forward influence and can be interpreted as characterising driving connections through the cortical hierarchies. The *lower* off-diagonals complement these and characterise the influence higher regions have upon lower, estimating back projecting activity. All significant connections are circled. (b) Connectivity map of all connection strengths that were significantly different from zero across all time lags for model 1. Arrow width is scaled to the p-value of the connection strength estimated in W (thin for p-values between 0.05-0.001 and thick for < 0.001).

Figure 7: Model 2 including regions V1/2 complex, V5 and bilinear term $I_{v1,ppc}$ for 1 subject (right hemisphere). (a) W is shown in the same layout as Figure 6. The weight of interest is between $I_{v1,ppc}$ and V5 as it describes a second order interaction characterising the modulatory role upon downstream processing (upon the connection strength between V1-V5). These coefficients were significantly non-zero (circled). A diagram of the model including the bilinear term as a “virtual” node is shown in (b) and the physiological interpretation of this is shown in (c) where $I_{v1,ppc}$ is not shown but is represented implicitly (the thin arrow represents a p-value of between 0.05-0.001).

Figure 8: Model 3 including regions V5, PPC and bilinear term $I_{v5,pfc}$ for 1 subject (right hemisphere). The significance of the connection between $I_{v5,pfc}$ -PPC is tested as in previous models and connectivity maps are shown as in Figure 6. The model supports the notion that PFC plays a modulatory role during attention upon the connectivity between V5 and PPC.

Figure 1a

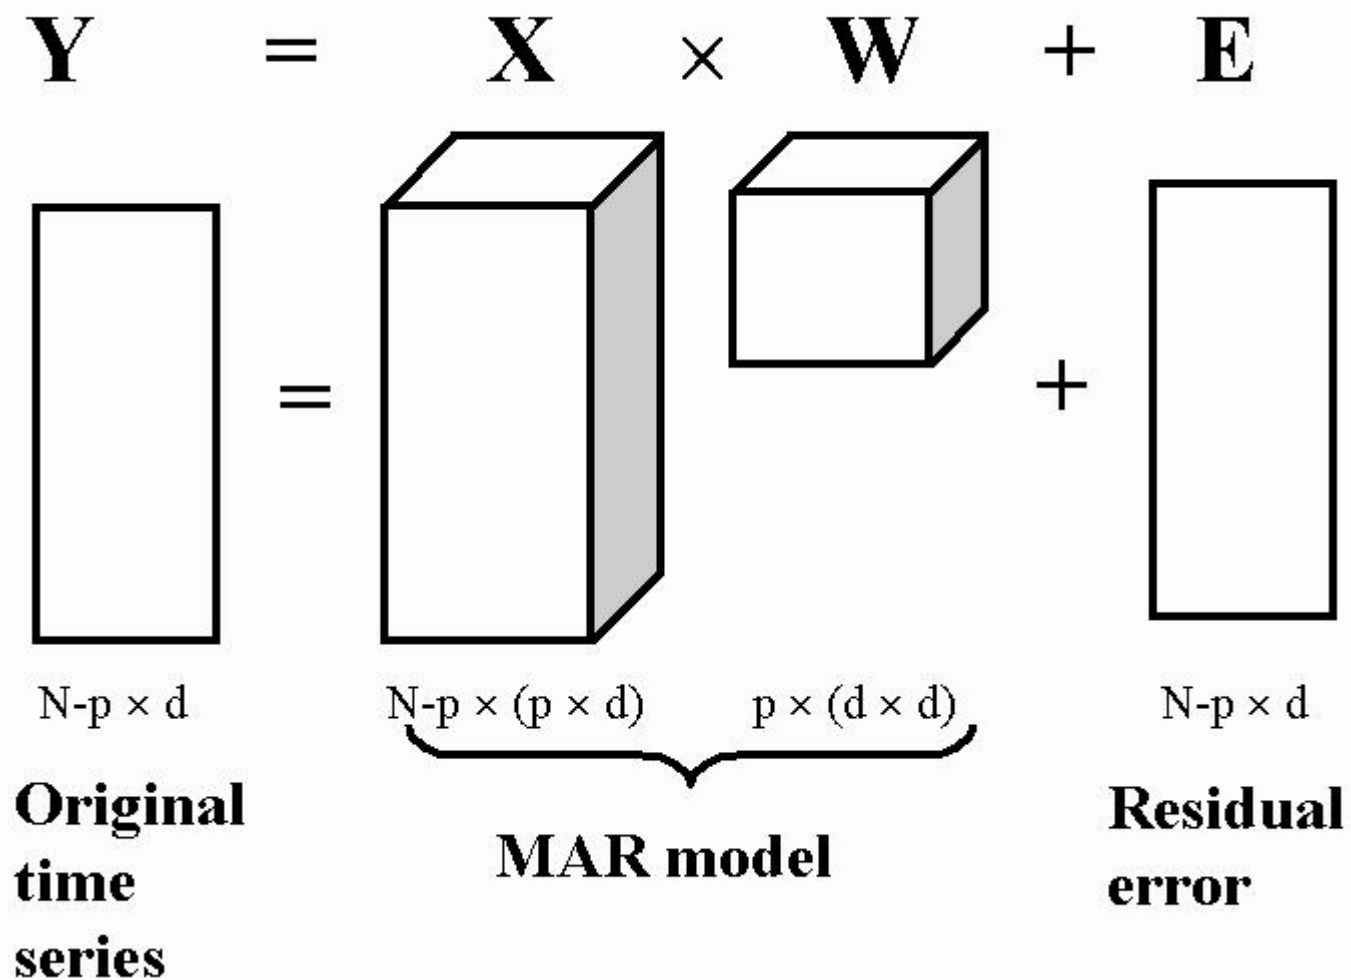




Figure 1b

$$\mathbf{W} =$$




Time lag
1



$$d \times d$$


2



$$d \times d$$

...

p



$$d \times d$$

Figure 2






Figure 3a Synthetic data - independent time series


Known parameters



Estimated parameters



Generated time series



Mean and variance of weights (W)

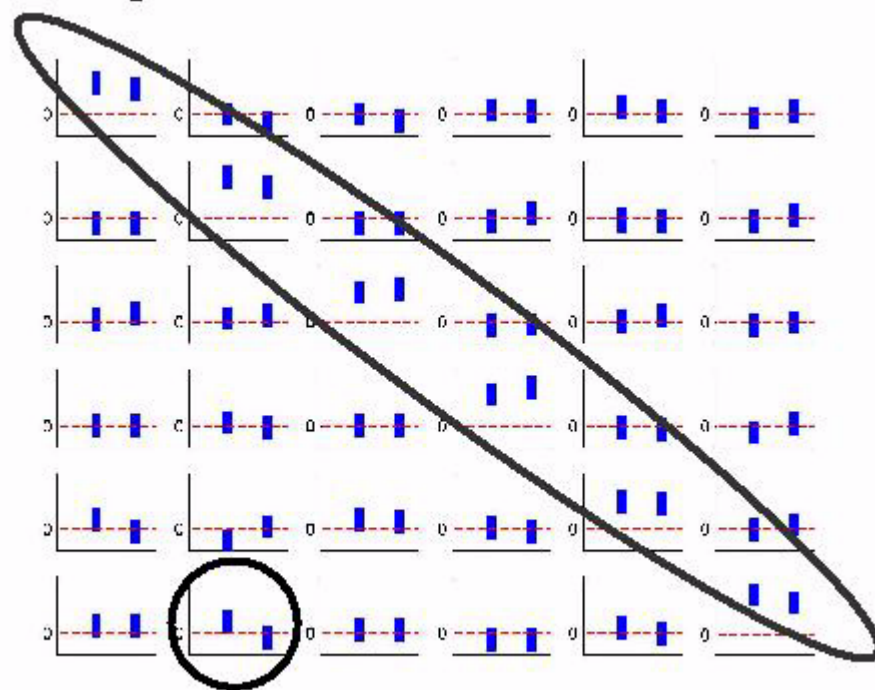



Figure 3b

Synthetic data - mixed dependence

Known parameters

Estimated parameters



Generated time series

Mean and variance of weights (\mathbf{W})





Figure 4


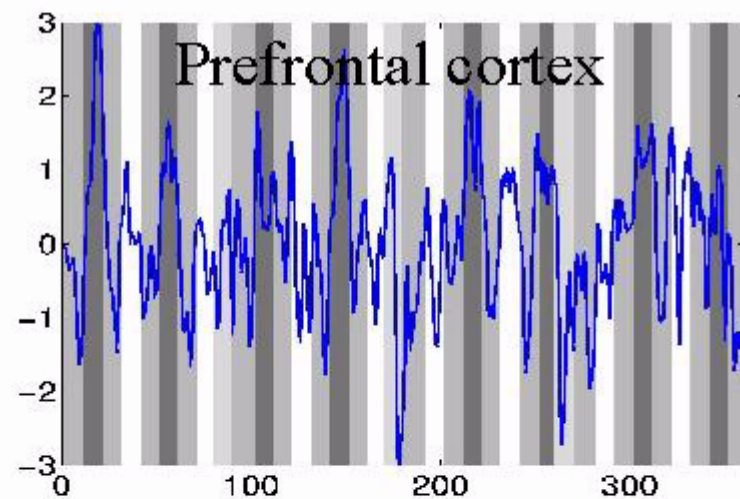
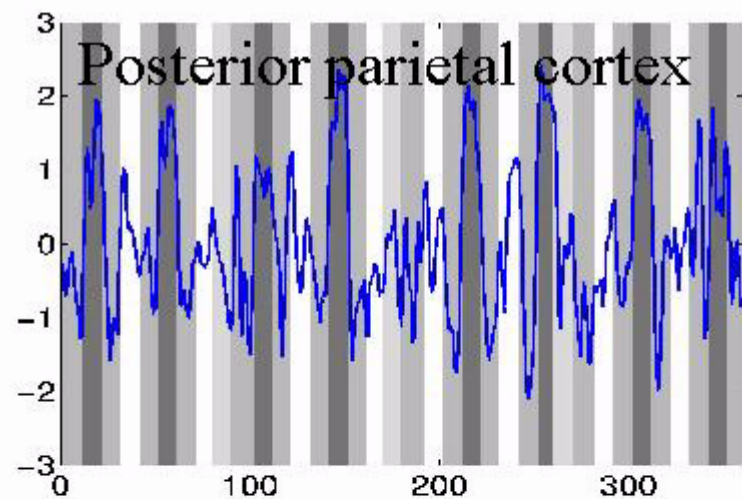




Figure 5a

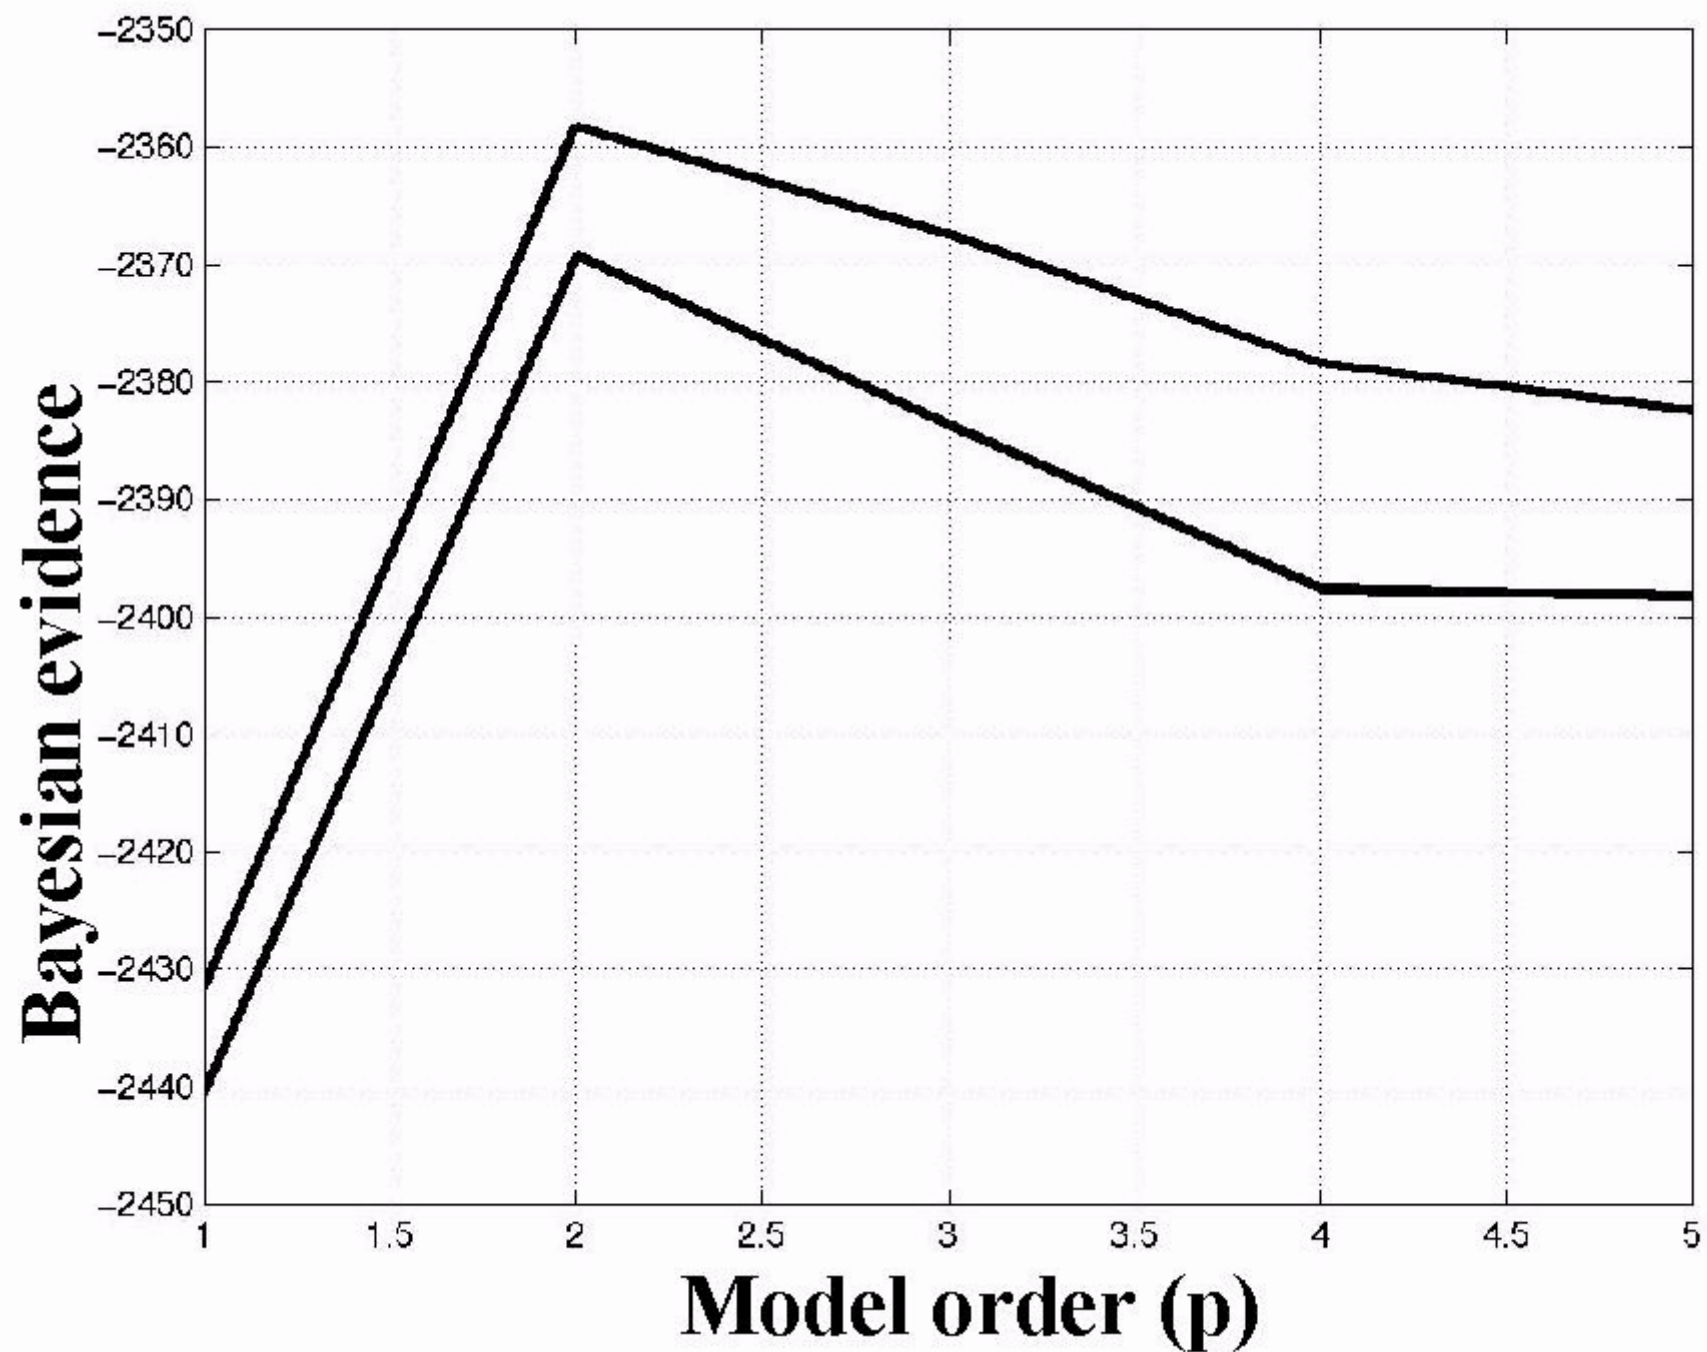


Figure 5b

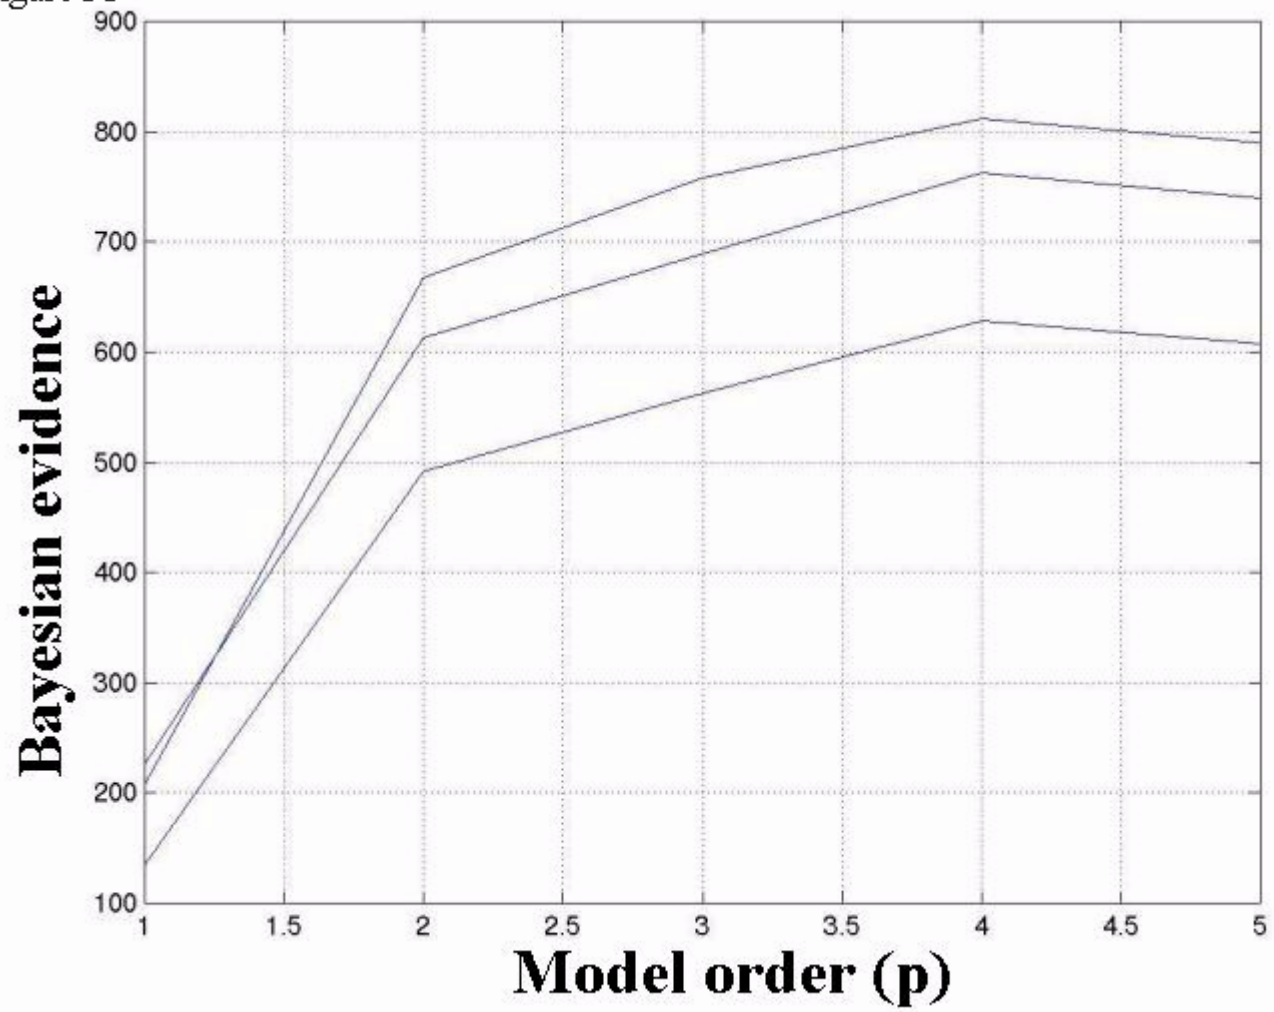


Figure 6a

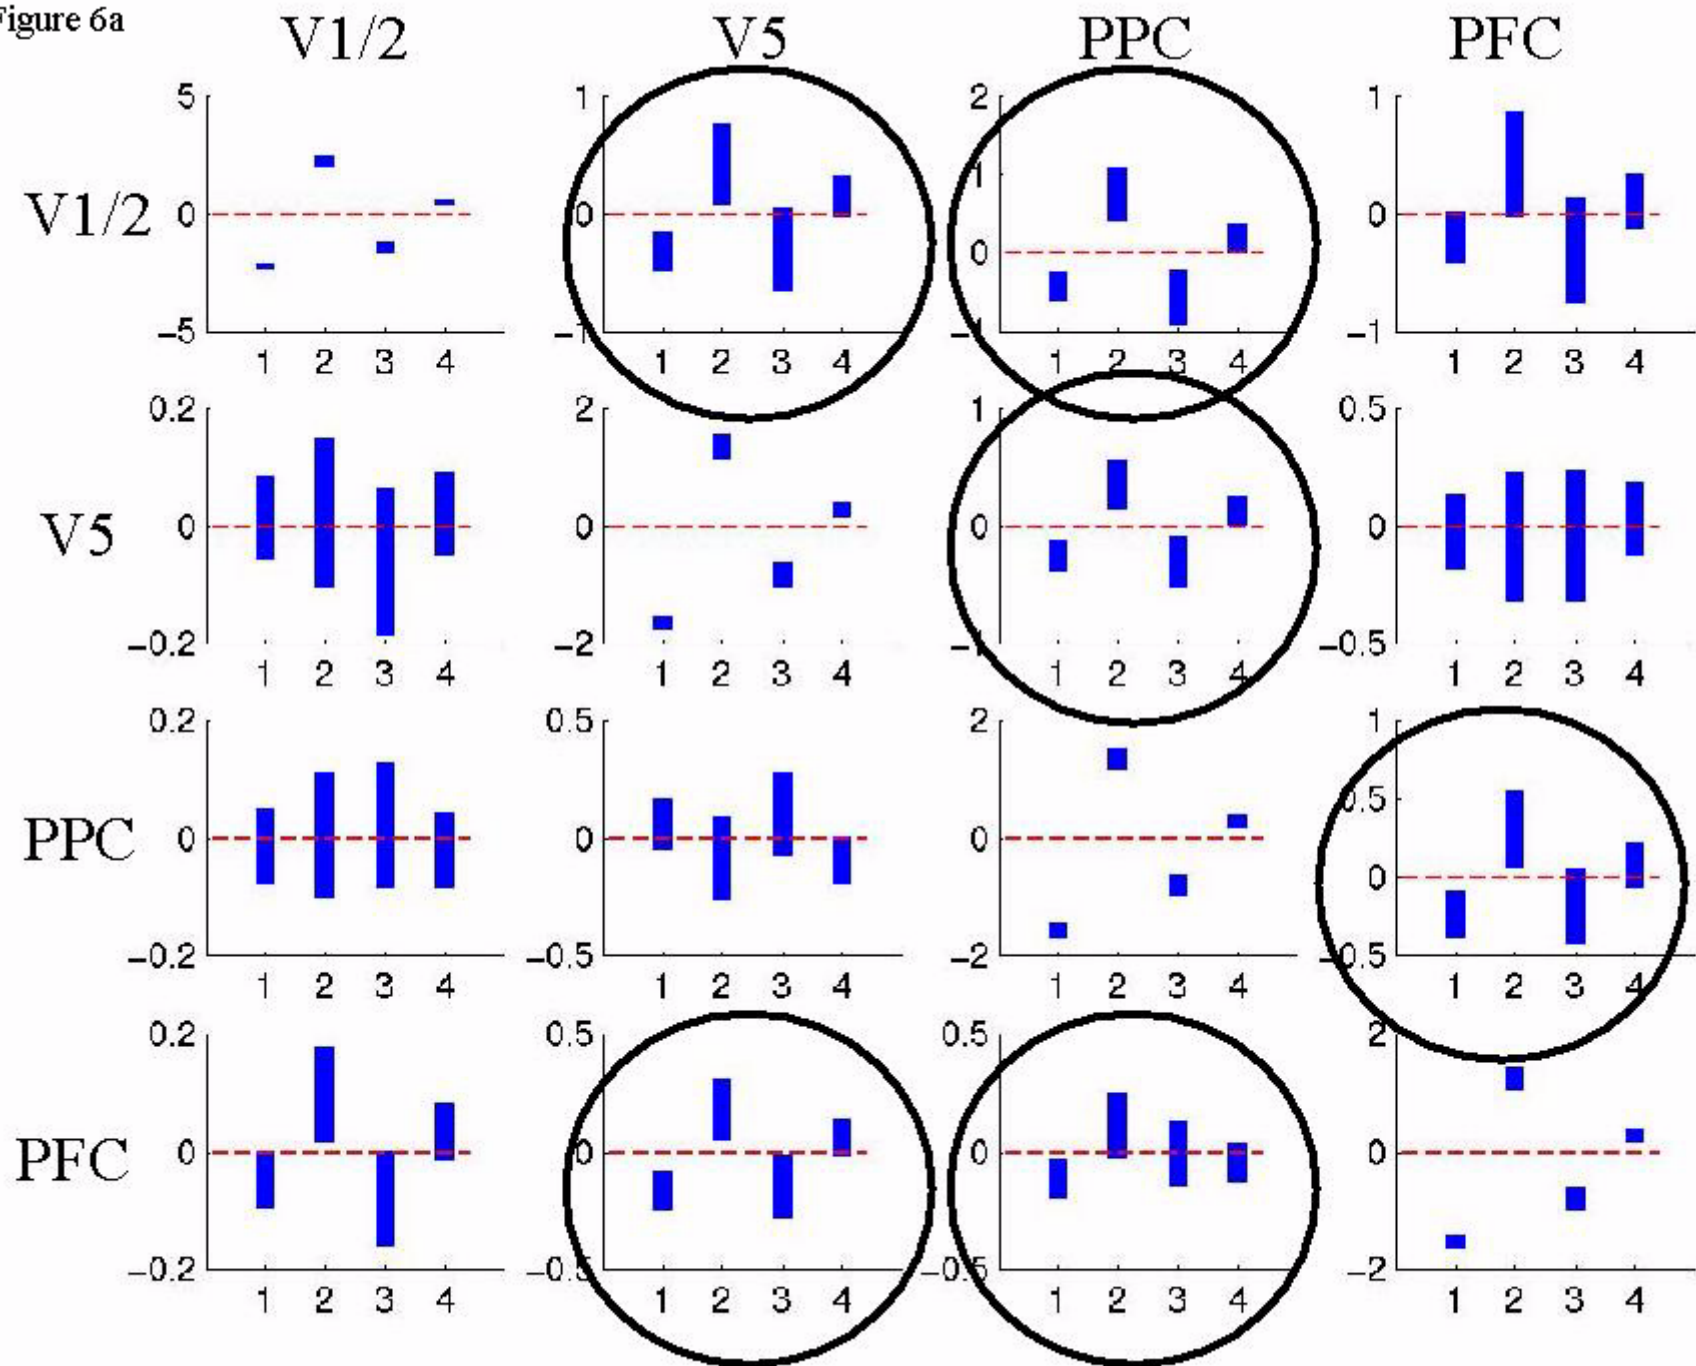


Figure 6b




Figure 7a

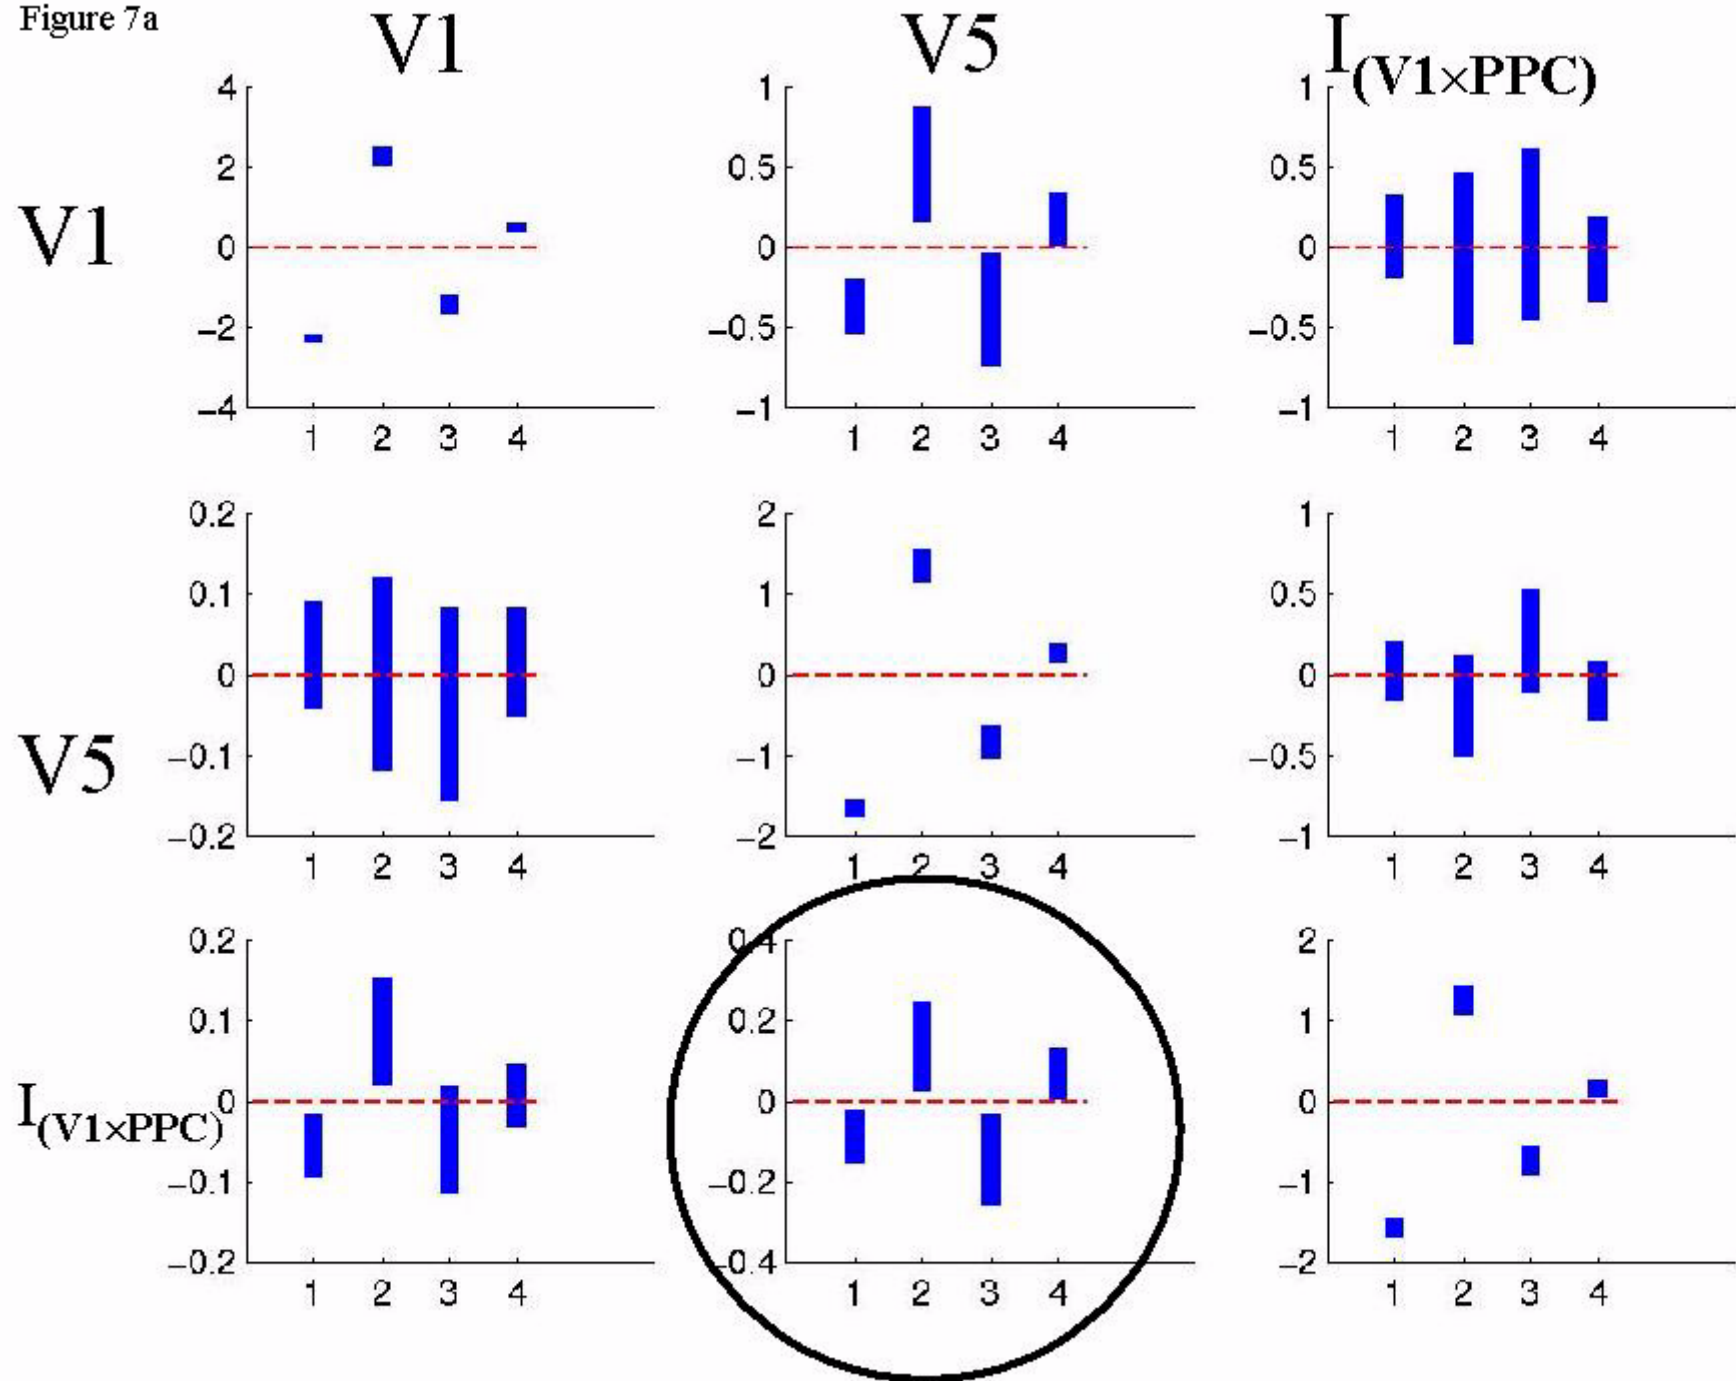
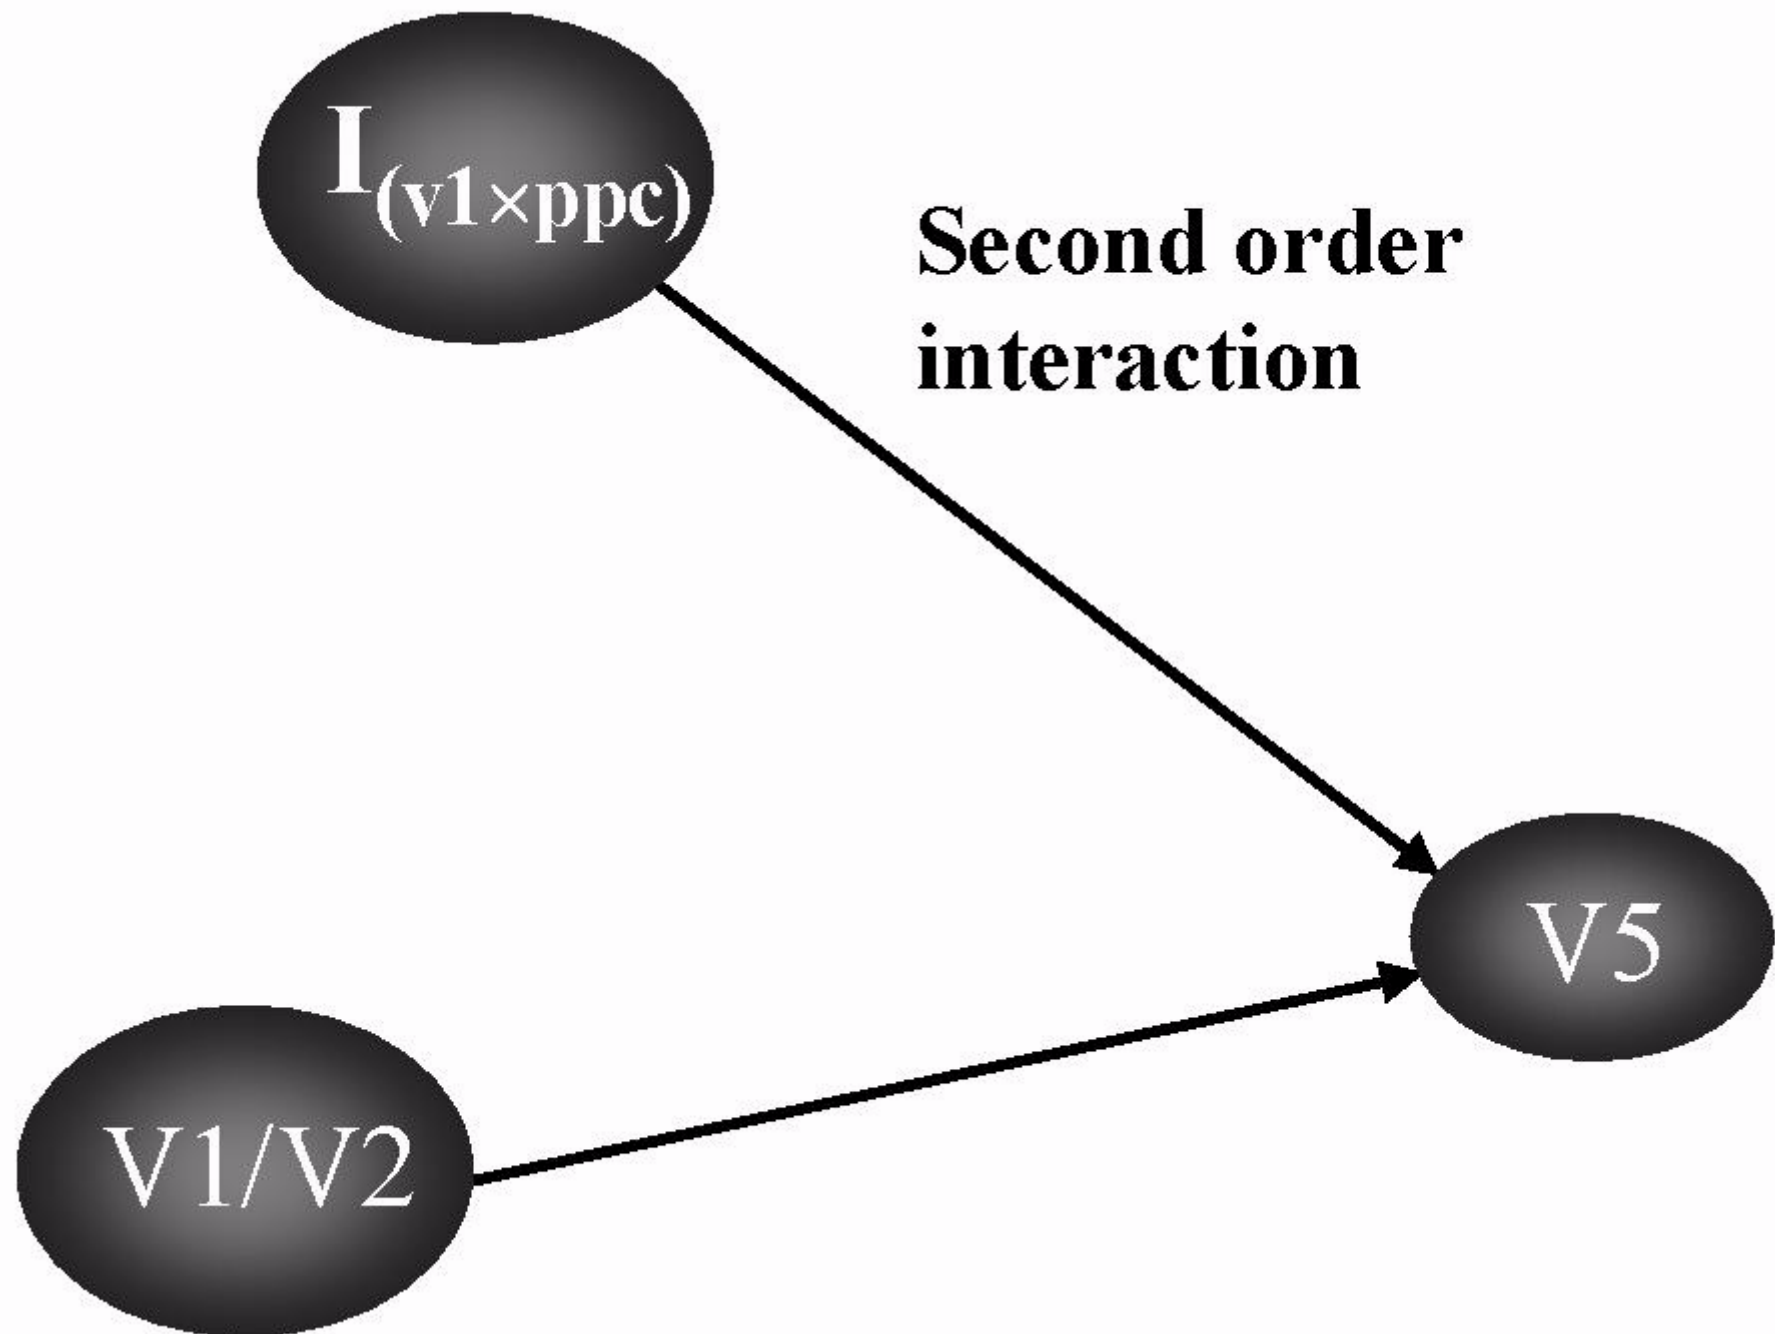




Figure 7b



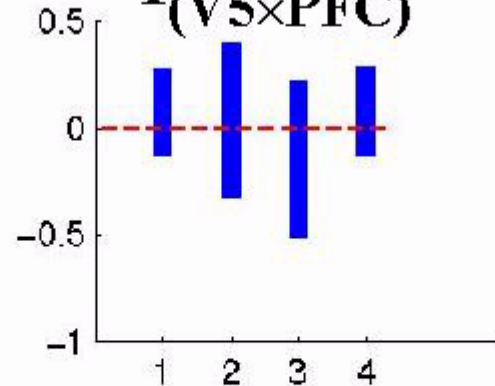
V5



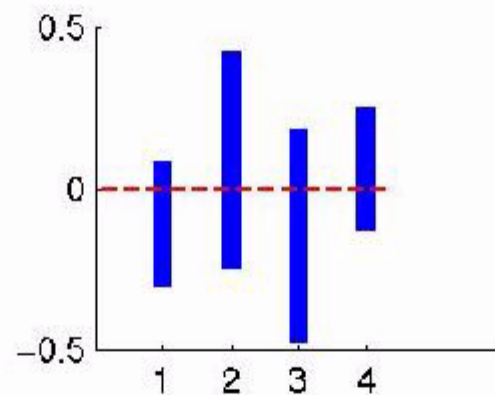


PPC



$I_{(V5 \times PPC)}$



PPC



$I_{(V5 \times PPC)}$


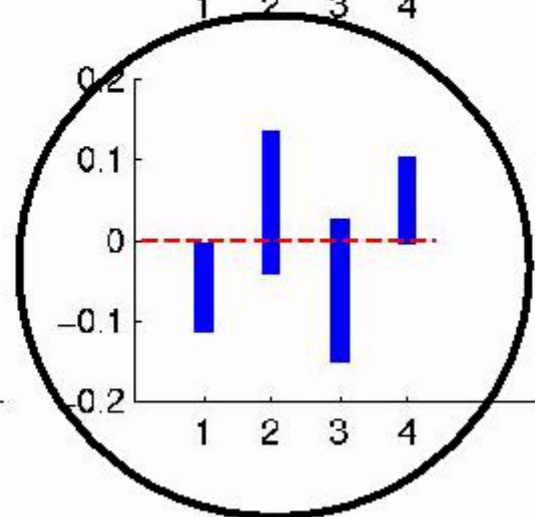
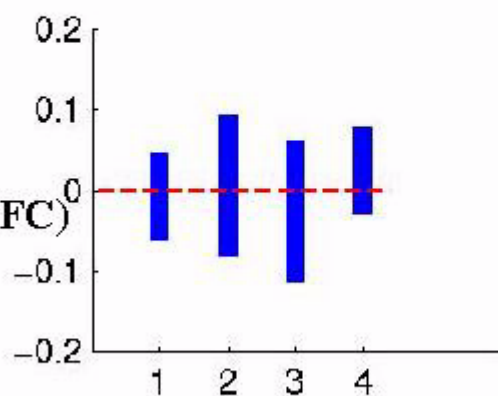


Figure 8b




Figure 8c

