

Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment

Mingzhou Ding, Steven L. Bressler, Weiming Yang, Hualou Liang

Center for Complex Systems and Brain Sciences, Florida Atlantic University, Boca Raton, FL 33431, USA

Received: 20 August 1999 / Accepted in revised form: 17 December 1999

Abstract. In this article we consider the application of parametric spectral analysis to multichannel event-related potentials (ERPs) during cognitive experiments. We show that with proper data preprocessing, Adaptive MultiVariate AutoRegressive (AMVAR) modeling is an effective technique for dealing with nonstationary ERP time series. We propose a bootstrap procedure to assess the variability in the estimated spectral quantities. Finally, we apply AMVAR spectral analysis to a visuomotor integration task, revealing rapidly changing cortical dynamics during different stages of task processing.

Even with the help of the most advanced windowing (tapering) techniques, nonparametric analysis produces severely biased spectral estimates for these short lengths (Granger and Hughes 1968; Jenkins and Watts 1968; Percival and Walden 1993; Muthuswamy and Thakor 1998).

A potential solution to both problems is provided by parametric spectral analysis in which MultiVariate AutoRegressive (MVAR) time series models are adaptively extracted from the data and become the basis for deriving spectral quantities. The theoretical foundation for this solution lies in treatment of the time series in a short time window as being generated by an underlying (approximately) stationary stochastic process. To make use of this formulation, we rely on the standard design of behavioral neurophysiology experiments in which a single task is often repeated hundreds, even thousands, of times (trials). Viewed from this perspective, the collection of ERPs from these trials is treated as an ensemble of realizations of a nonstationary stochastic process with locally stationary segments. As such, the ERPs from different trials provide statistical samples for reliably estimating time series models over time intervals nearly as brief as that spanned by the autoregressive model.

The objective of this work is to demonstrate the plausibility and utility of this approach to parametric analysis both theoretically and experimentally. Previous work (Florian and Pfurtscheller 1995; Schack and Krause 1995) has applied model-based spectral analysis to short time segments of EEG recordings. Our main contributions lie in the following three areas. First, we identify a set of preprocessing steps, including removal of the averaged event-related potentials (ensemble means) from the ensemble of single-trial ERPs, that prove to be essential for the meaningful estimation of the time series models. Second, we design a set of tests, including the residual whiteness test, to show that the fitted models adequately represent the statistical properties of the ERP data. Third, we devise a variability assessment strategy based on the ideas of bootstrapping that allows statistical testing of significance to be carried out on the derived spectral quantities.

1 Introduction

Spectral analysis is a mature methodology that has been successfully applied to nearly every field of science and technology (Jenkins and Watts 1968; Bendat and Piersol 1986; Percival and Walden 1993). We sought to apply spectral analysis to multichannel event-related potentials (ERPs) from the cerebral cortex to understand how different cortical areas work together during cognitive processing. We discovered, however, that two fundamental aspects of this problem present serious challenges to the application of spectral analysis. First, because different cortical areas must be treated as belonging to the same integrated system to understand their interactions, spectral analysis of cortical ERPs must be fully multivariate. Second, due to the rapidly changing (nonstationary) nature of neural activity related to cognitive processing, one must examine ERPs in brief time intervals on the order of 40–80 ms (or possibly even shorter). For the 200-Hz sampling rate used in this study, this length corresponds to a data string of only 8–16 points. The traditional nonparametric approach to spectral analysis, based on discrete Fourier transform (DFT), is rendered ineffective by such short data lengths.

From the experience gained in this investigation, we propose the following general approach to the spectral analysis of ERPs during a cognitive task:

1. Cover the entire task time course with highly overlapped time windows short enough that the process in each window can be treated as locally stationary.
2. In each time window, derive a linear stochastic MVAR model of the process by fitting the ERP data (after suitable preprocessing) from an ensemble of trials.
3. Derive spectral quantities such as power and coherence from the model parameters. The evolution of these quantities, obtained by adaptively fitting models over successive windows, yields a very finely resolved depiction of cortical dynamics. Since the MVAR model parameters are adapted to each window, this procedure is called Adaptive MVAR (AMVAR) modeling.
4. Assess the variability of the spectral estimators by bootstrapping.

This article is organized as a principled progression of steps that establish the validity of the AMVAR modeling approach to spectral analysis of cortical ERPs. In Sect. 2 we present mathematical formulations of the AMVAR model and related estimation procedures. We discuss the steps leading from autoregressive models to spectral quantities. In Sect. 3 we consider a simple three-variable autoregressive process for which all spectral quantities are known exactly. We demonstrate that one can obtain the correct spectral quantities by fitting an MVAR model to multiple short data segment realizations generated by the process. In Sect. 4 we examine the application of MVAR models to ERPs recorded from macaque monkeys. Here we discuss in detail the four-step AMVAR procedure mentioned above. In Sect. 5 we demonstrate the usefulness of the AMVAR approach for the study of cortical interactions during cognitive processing. Section 6 summarizes and reiterates the main points of the article.

2 Formulation and estimation of MVAR models

Extensive treatment of MVAR time series models can be found in the literature (Whittle 1963; Gersch 1970; Pfurtscheller and Haring 1972; Franaszczuk et al. 1985; Lutkepohl 1993; Schack and Krause 1995). Here we summarize the main elements needed for our work with emphasis on how to incorporate multiple realizations of data into the estimation procedure.

Let $\mathbf{X}(t) = [X(1, t), X(2, t), \dots, X(p, t)]^T$ be a p -dimensional random process. Here T denotes matrix transposition. In ERP studies, p represents the total number of ERP channels. Assume that the process $\mathbf{X}(t)$ is stationary and can be described by the following m th-order autoregressive equation:

$$\mathbf{X}(t) + \mathbf{A}(1)\mathbf{X}(t-1) + \dots + \mathbf{A}(m)\mathbf{X}(t-m) = \mathbf{E}(t) \quad (1)$$

where $\mathbf{A}(i)$ are $p \times p$ coefficient matrices and $\mathbf{E}(t) = [E(1, t), E(2, t), \dots, E(p, t)]^T$ is a zero mean uncorrelated noise vector with covariance matrix Σ .

To estimate $\mathbf{A}(i)$ and Σ , we multiply (1) from the right by $\mathbf{X}^T(t-k)$, where $k = 1, 2, \dots, m$. Taking expectation we obtain the Yule-Walker equations

$$\mathbf{R}(-k) + \mathbf{A}(1)\mathbf{R}(-k+1) + \dots + \mathbf{A}(m)\mathbf{R}(-k+m) = \mathbf{0} \quad (2)$$

where $\mathbf{R}(n) = \langle \mathbf{X}(t)\mathbf{X}^T(t+n) \rangle$ is $\mathbf{X}(t)$'s covariance matrix of lag n . In deriving these equations we have used the fact that $\langle \mathbf{E}(t)\mathbf{X}^T(t-k) \rangle = \mathbf{0}$ as a result of $\mathbf{E}(t)$ being an uncorrelated process.

For a single realization of the \mathbf{X} process, $\{\mathbf{x}(i)\}_{i=1}^N$, we compute the covariance matrix in (2) according to

$$\tilde{\mathbf{R}}(n) = \frac{1}{N-n} \sum_{i=1}^{N-n} \mathbf{x}(i)\mathbf{x}^T(i+n) \quad (3)$$

If multiple realizations of the same process are available, then we compute the above quantity for each realization and average across all the realizations to obtain the final estimate of the covariance matrix. (Note that for a single short trial of data one uses the divisor N for evaluating covariance to reduce inconsistency. Due to the availability of multiple trials we have used the divisor $(N-n)$ in the above definition (3) to achieve an unbiased estimate.) It is quite clear that, for a single realization, if N is small, one will not get good estimates of $\mathbf{R}(n)$ and hence will not be able to obtain a good model. This problem can be overcome if a large number of realizations of the same process are available. In this case the data length can be as short as the model order m plus 1.

Equation (1) contains a total of mp^2 unknown model coefficients. In (2) there are exactly the same number of simultaneous linear equations. One can simply solve these equations to obtain the model coefficients. The Levinson, Wiggins, Robinson (LWR) algorithm (Morf et al. 1978; Haykin and Kesler 1983) is a more robust solution procedure, based on the ideas of maximum entropy. This algorithm was implemented in this study. The noise covariance matrix Σ can be obtained as part of the LWR algorithm. Otherwise one may obtain Σ through

$$\Sigma = \mathbf{R}(0) + \sum_{i=1}^m \mathbf{A}(i)\mathbf{R}(i) \quad (4)$$

Here we note that $\mathbf{R}^T(k) = \mathbf{R}(-k)$.

The above estimation procedure can be carried out for any model order m . The correct m is usually determined by minimizing the Akaike Information Criterion (AIC) (Akaike 1974) defined as

$$\text{AIC}(m) = 2 \log[\det(\Sigma)] + 2p^2m/N_{\text{total}} \quad (5)$$

where N_{total} is the total number of data points from all the trials. Plotted as a function of m the proper model order corresponds to the minimum of this function.

A standard procedure to examine whether an MVAR time series model is suited for a given data set is to check whether the residual noise is white. Here the residual

noise is obtained by computing the difference between model predicted values and the actually measured values. In Sect. 4 we describe an additional test that allows us to assess directly how much statistical structure in the data is reflected by the model.

Once an MVAR model is adequately estimated it becomes the basis for subsequent spectral analysis.¹ In the spectral domain (1) can be written as

$$\mathbf{X}(f) = \mathbf{H}(f)\mathbf{E}(f) \quad (6)$$

where

$$\mathbf{H}(f) = \left(\sum_{j=0}^m \mathbf{A}(j)e^{-ij2\pi f} \right)^{-1} \quad (7)$$

is the transfer function. Note that we define $\mathbf{A}(0)$ to be the identity matrix.

From (6), after proper ensemble averaging, we obtain the spectral matrix

$$\mathbf{S}(f) = \mathbf{H}(f)\mathbf{\Sigma}\mathbf{H}^*(f) \quad (8)$$

from which all the commonly used multivariate spectral quantities such as power, ordinary coherence, partial coherence, and multiple coherence can be derived (Jenkins and Watts 1968). For the purpose of this article, we will only consider the ordinary coherence, which we hereafter refer to simply as coherence. The coherence between two given components $X(i, t)$ and $X(j, t)$ is defined as

$$C_{ij}(f) = \frac{|S_{ij}|^2}{S_{ii}S_{jj}} \quad (9)$$

where S_{ij} is the (i, j) th element of the spectral matrix. (We comment that sometimes the quantity $C_{ij}(f)$ is given the name of squared coherency and $\sqrt{C_{ij}(f)}$ is called the coherence.) The value of coherence is normalized between 0 and 1, and it measures the degree of linear dependence between $X(i, t)$ and $X(j, t)$. If it is near 1 at any frequency f , then the two processes are maximally interdependent at that frequency. On the other hand, a value near 0 indicates independence of the two processes at frequency f .

¹ Given the goal here of achieving short-window spectral analysis, a common question concerns the resolution of spectral quantities at low frequencies. This question in fact highlights the difference between parametric and nonparametric spectral estimations. For classical (nonparametric) spectral analysis, resolution is indeed a problem if the window length is not long enough to cover a substantial segment of the cycle. In AMVAR analysis, however, we first establish that the repeated data observations can be accurately estimated by an underlying linear stochastic process that is approximately stationary within the analysis window. Since this process is stationary, it can be computed for any length of time, no matter how long. The spectral quantities derived by AMVAR analysis are the same as those that would be computed from this stationary process if it were allowed to continue for a very long time. Since that stationary process may contain components of any frequency, we can accurately determine those components, even at low frequencies. This holds equally well for all the components in the full MVAR model

3 Analysis of a simple three-channel autoregressive model

In this section we consider a three-channel autoregressive process:

$$\begin{aligned} x(t) &= \xi(t) \\ y(t) &= x(t-1) + \eta(t) \\ z(t) &= \lambda z(t-1) + x(t-1) + \epsilon(t) \end{aligned} \quad (10)$$

where $|\lambda| < 1$ is a parameter, and $\xi(t)$, $\eta(t)$, $\epsilon(t)$ are three independent white noise processes with zero means and variances σ_1^2 , σ_2^2 , σ_3^2 , respectively. The pairwise coherence is easily calculated to be

$$\begin{aligned} C_{12}(f) &= \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \\ C_{13}(f) &= \frac{\sigma_1}{\sigma_1^2 + \sigma_3^2} \\ C_{23}(f) &= \frac{\sigma_1^4}{(\sigma_1^2 + \sigma_2^2)(\sigma_1^2 + \sigma_3^2)} \end{aligned} \quad (11)$$

which are independent of the frequency.

Setting $\lambda = 0.5$, $\sigma_1 = 1$, $\sigma_2 = 0.2$, and $\sigma_3 = 0.3$ we simulate (10) to generate 100 realizations of 10-point data sets. Assuming no knowledge of (10) we fit an MVAR model to the data according to the steps specified in the previous section. The derived coherence shown in Fig. 1A is in excellent agreement with the analytical results, $C_{12}(f) = 0.96$ (top curve), $C_{13}(f) = 0.92$ (middle curve), $C_{23}(f) = 0.88$ (lower curve), computed from (11). This example demonstrates the principle that stochastic processes of the autoregressive type can be fully recovered from multiple trials of short data segments. It is worth noting that, to examine the sensitivity of the fitting process to the model order specification, we used a third-order MVAR to fit the simulated data. The excellent result demonstrates the good tolerance provided by models whose order is slightly different from that of the generating process. Equally excellent results are obtained when a first-order MVAR is used.

A common preprocessing step in time series analysis is to subtract the temporal mean from the time series. We have found that this step is not advisable for individual realizations that are of short duration. This is because the temporal mean estimated from the small number of samples in a short data segment is not a reliable estimate of the true process mean and can vary greatly from realization to realization. Figure 1B shows the effect of subtracting the temporal mean from each realization of the simulated three-variable data set before model fitting. The resulting coherence spectrum from the fitted model is severely underestimated in the low frequency components. We will argue in the next section that a necessary preprocessing step prior to model fitting is subtraction of the ensemble mean. In the present example, subtraction of the ensemble mean is not required since it is already zero due to the model assumptions. Actual ERP data typically contain a sig-




Fig. 1. **a** Estimated coherence spectra for the pairwise combinations of channels in the three-variable autoregressive model (10). **b** Estimated coherence spectra as in **a**, but with prior removal of the temporal mean from each realization. It is demonstrated that such removal leads to poor estimation of the coherence spectrum of the underlying process

nificant ensemble mean, the averaged ERP, which is commonly studied in electrophysiology.

4 Application of AMVAR to cortical event-related potentials

ERP data from macaque monkeys were used to test the AMVAR approach outlined in Sect. 2. The ERPs were recorded from chronically implanted transcortical bipolar electrodes at 15 distributed sites in multiple cortical areas of one hemisphere as the monkey performed a visuomotor pattern discrimination task (Bressler et al. 1993). The prestimulus stage began when the monkey depressed a hand lever. This was followed 0.5–1.25 s later by the appearance of a visual stimulus (a four-dot pattern) on a computer screen. The monkey made a GO or NO-GO response depending on the stimulus category and the session contingency. The entire trial lasted about

500 ms, during which the ERPs were recorded at a sampling rate of 200 Hz. The specific data set considered here consisted of 888 trials from one monkey with the GO response.

For the purpose of data analysis, each of the 15 electrode recording sites was considered to be a separate data channel. As the first step of preprocessing, we detrended the single-trial ERP time series of each channel, subtracted the temporal mean of the entire trial, and divided by the temporal standard deviation. The result was that the data from each channel and each trial were given equal weight in model estimation.

4.1 Inherent nonstationarity of ERP data

The visuomotor pattern discrimination task studied here did not involve a single uniform cognitive function. Rather, it required that the monkey undergo several distinct cognitive states during the course of the trial: motor maintenance, visual stimulus anticipation, visual feature extraction, visual pattern discrimination, GO/NO-GO decision, and final motor execution in the GO condition. This succession of states was reflected in an inherent nonstationarity of the recorded ERPs. As mentioned above, the use of MVAR analysis is based on the assumption that the underlying stochastic process is stationary. Therefore, MVAR analysis of the ERP data necessitated that we first assess the degree of this nonstationarity.

The simplest nonstationarity measure is the ensemble mean (first-order statistic), calculated by averaging the measured amplitude values at a given channel at each time point across trials. Figure 2 shows an example of the ensemble mean from a site in the striate cortex. The time-varying nature of the signal shows that the underlying process is nonstationary in the mean. It is worth noting that this ensemble mean is the commonly studied averaged ERP that is often considered to be the only signal related to cognitive processing. An important result of the




Fig. 2. Averaged event-related potential (with standard deviation as error bars) from a site in the striate cortex

present study is that this view is not correct. We find that the so-called “noise” component of the ERP, obtained by subtracting the ensemble mean from the single-trial ERP, actually contains rich task-relevant information that cannot be inferred from the ensemble mean.

Figure 2 also plots as error bars at each time point the ensemble standard deviation, which is a second-order statistic. The time-varying nature of this quantity is also apparent.

The nonstationarity embodied in the mean and standard deviation can be easily removed by subtracting the ensemble mean, point-by-point, from each trial and then dividing the result, again point-by-point, by the standard deviation. In fact, these two procedures constitute our second and third preprocessing steps. They are applied separately to the data from each channel. After these two preprocessing steps are applied, the mean of the resulting ensemble of trials is zero, and the standard deviation is one, at every time point. Note that the step of standard deviation normalization is crucial for allowing the dynamical changes in model-derived spectral quantities to be compared at each stage of task processing. Unless otherwise noted, the above three preprocessing steps were applied prior to all other analysis steps described below. We will return to the topic of ensemble mean subtraction in more detail in the next subsection.

A deeper source of nonstationarity that cannot be removed easily is the correlation structure in the data. Correlation nonstationarity refers to the fact that the auto-correlations of single channels, and the cross-correlations between channels, at a fixed lag, vary as a function of time during the task. To demonstrate this correlation nonstationarity, we averaged the zero-lag cross-correlation between two channels from the striate cortex over all trials successively for each time point during the task. The time-varying structure of the resulting curve, shown in Fig. 3, clearly demonstrates the nonstationarity of the cross-correlation statistic.




Fig. 3. The zero-lag cross-correlation between two channels from the striate cortex computed in 1-point windows at every time point during the task

4.2 Short-window approach to dealing with correlation nonstationarity

Our strategy for dealing with correlation nonstationarity is to perform MVAR analysis in short, highly overlapped time windows in which the underlying stochastic processes are considered to be locally stationary. (Using overlapped windows makes the estimated models vary smoothly.) By adaptively estimating the MVAR model parameters in each window, this AMVAR procedure yields finely resolved dynamical information about the cortical processes related to cognitive state. A crucial question is how to choose the window length in an objective way. It is clear that it must be longer than the model order (see Sect. 2). On the other hand, it should not be so long as to lose the temporal dynamics that we wish to resolve.

A rough estimate of the upper limit of the window length may be obtained with the following procedure. First, as in Fig. 3, compute the zero-lag cross-correlation between two channels in a 1-point window at all time points during the task. Second, repeat the first step for progressively longer windows, averaging both within the window and across all the trials. As the window gets longer, more smoothing of the correlation structure occurs, and dynamical variation is lost. This point is illustrated in Fig. 4, which shows the same (1-point window) cross-correlation curve as in Fig. 3, with additional curves using 10-point (50 ms) and 20-point (100 ms) windows. Clearly the 20-point window loses a great deal of the variability in the correlation structure. The 10-point window, on the other hand, is able to track much of the basic pattern of variation. For model estimation, we must also consider the smoothness of the estimated spectral quantities. Our experience indicates that the 10-point (50 ms) window is a good compromise between preserving correlation variability and maintaining smoothness of the estimated spectral quantities. Although one could tailor the window size to fit different




Fig. 4. The zero-lag cross-correlation between the same two channels as in Fig. 3 for 1-, 10-, and 20-point windows

parts of the trial, a 10-point window size is used in the remainder of this study for uniformity in evaluation and display of the results.

4.3 Model order determination and model validation

Thus far our focus has been on data preprocessing and window length determination. We now examine the issue of fitting MVAR models in short time segments. The main objective is to show that MVAR time series models fitted to data sets containing segments of short duration from many trials can indeed capture the statistical properties of the underlying stochastic process.

The first problem in MVAR modeling is to choose the model order, which we did using the AIC. Figure 5 shows the $AIC(m)$ [defined in (5)] as a function of model order for a representative 50-ms window centered 120 ms after the stimulus onset. The shape of the curve was extremely similar for all windows in the trial. The rationale in using the AIC for order determination is to select the model order for which the AIC reaches a minimum. However, the AIC decreases monotonically with increasing order. In fact, this monotonic decrease is typical for our data set and also appears to be a general feature of electroencephalographic (EEG) time series (Jansen 1991). We infer from Fig. 5 that model order 5 is sufficient since there is very little change in the AIC beyond that value. This model order also appears to be quite consistent with other EEG studies (Jansen 1991). To test our choice of model order, we calculated a set of spectral quantities for different model orders around 5 and found that the results are robust against small changes in model order.

An important assumption in MVAR modeling is that the noise process $\mathbf{E}(t)$ in (1) is a white (uncorrelated) process. This requirement can be intuitively understood




Fig. 5. The Akaike Information Criterion (AIC) as a function of model order, computed in a representative 50-ms window centered 120 ms after stimulus onset. The shape of the curve is similar for all other windows

in the sense that if the left-hand side of (1) has captured all the temporal structures in the data, then what remains on the right-hand side (the residual) has no temporal structure. The question is whether the residuals remaining after MVAR modeling of the ERP segments are indeed uncorrelated. We tested this idea by using auto- and cross-correlation measures of the residuals.

Using 15 ERP channels, a 10-point (50-ms) window, and a model order of 5, we obtained five 15-dimensional residual noise vectors. (The five vectors corresponded to the five points in a 10-point window that could be predicted with a 5-point model; each vector contained a residual element for each channel.) We computed auto- and cross-correlations up to lag 3, excluding lag 0, for all pairwise combinations of the 15 channels. The null hypothesis was that the residual noise had no temporal correlation. For this to be true, fewer than 95% of the correlation coefficients were expected to fall within the interval $[-2/\sqrt{5}, 2/\sqrt{5}]$. This allowed about 5% of the coefficients to fall outside the interval by pure chance. We computed the percentage of coefficients lying outside the above interval for each window for model orders from 1 to 6. For all model orders, the percentage of auto- and cross-correlation coefficients that were outside of the interval was below 2%. This proved that the residual noise was indeed white and that the data could be represented as an AR process.

As a last step of model validation, we now consider the question of what portion of the correlation structure in the data is captured by the fitted MVAR models. For this, we devised a statistical consistency test that compared the correlation structure of the real ERP data segments with that of simulated data segments generated by iterating the corresponding equations of the fitted models. For both the real and simulated data sets, containing 15 channels and 888 trials, we computed all auto-correlations and pairwise cross-correlations up to lag 5 in a sequence of overlapping 50-ms windows. In each window, there are 1245 such correlations, treated as a 1245-dimensional vector. The statistical consistency between the correlation structure of the real ERP data and that of the simulated data can be measured by the following percent consistency measure:

$$PC(t) = \left(1 - \frac{|\mathbf{R}_s(t) - \mathbf{R}_r(t)|}{|\mathbf{R}_r(t)|} \right) \times 100 \quad (12)$$

where \mathbf{R}_r denotes the correlation vector of the real data and \mathbf{R}_s was the correlation vector of the simulated data. The smaller the distance between the two vectors, the larger was the consistency. In Fig. 6, we see that, as the window was moved across the duration of the task, this measure was reliably around 90%, indicating a high degree of consistency.

4.4 Ensemble mean subtraction and stability of MVAR models

For the MVAR process defined by (1) to be stable, the roots to the following equation




Fig. 6. Percent consistency as a function of time, demonstrating the similarity between the correlation structure of the real data and that of AMVAR model simulated data during the entire task




Fig. 7. The Stability Index as a function of time without ensemble mean subtraction (*dashed line*) and with it (*solid line*), demonstrating the benefit of such subtraction

$$\det(\lambda^m \mathbf{I} + \lambda^{m-1} \mathbf{A}(1) + \dots + \lambda \mathbf{A}(m-1) + \mathbf{A}(m)) = 0 \quad (13)$$

must satisfy the condition $|\lambda| < 1$. Here by stable we mean that the random process coming from the equation is stationary and does not diverge. If any one of the eigenvalues in (13) has a magnitude greater than 1, the process generated by (1) will blow up, and it cannot be a model of a stationary random process.

Thus, as part of model validation, one should also check whether a fitted model is stable in the sense defined above. Here we discuss this issue in conjunction with the preprocessing step of ensemble mean subtraction prior to model fitting. A common procedure in AR modeling is to subtract the temporal mean within a given window from the single trial time series before fitting. Although this is acceptable in the case of long stationary time series, the results of Sect. 3 above show that this is a bad practice when the window is short. In the short-window case, provided that there exists an ensemble of trials, a much more important procedure is to subtract the ensemble mean from each trial. By doing so, we (1) remove the first-order nonstationarity from the data; and (2) make the ensemble mean equal to zero, which is a requirement for model fitting.² Additionally, we now

² To illustrate further the importance of this step in the analysis of interdependency, we consider the following hypothetical example in which two subjects take part in the same experiment. Record the event-related brain response from the same electrode position on the two subjects. It is reasonable to assume that the averaged event-related responses (ensemble means) from the same electrode position will be similar for the two subjects. If the ensemble means are not removed from these two channels, then we will find significant coherence between the two channels at the frequency of the ensemble mean. This is clearly not a sensible result. For scalp electroencephalogram (or magnetoencephalogram) this factor may not be as significant since the ensemble means are small relative to the raw signals. For the intracranial recordings used here this factor is very important as shown in what follows

show that removal of the ensemble mean is also crucial for achieving stability of the fitted models.

To determine the stability of a fitted MVAR model, we only need to evaluate the largest eigenvalue in (13), denoted λ_1 . We used the quantity $\log |\lambda_1|$ as the Stability Index (SI). Figure 7 shows SI as a function of time during the task. In the case of not removing the ensemble mean (*dashed line*), the values of SI became positive for a period after the onset of stimulus (at 0 ms), rendering the fitted models invalid during this time period. After the removal of the ensemble mean, all the values of SI were negative (*solid line*), indicating that all of the fitted models throughout the task were stable.

4.5 Variability assessment by bootstrapping

Whether using parametric modeling or traditional nonparametric analysis, it is often desirable to determine the statistical significance of spectral quantities to make comparisons between different task conditions or different times during the task. In the case of cognitive processing, these comparisons are frequently crucial for understanding of the underlying neural dynamics. If the entire ensemble of available trials is used to estimate a spectral quantity, only one value of that estimate can be obtained, and there is no indication of the estimation variability. It is just this variability that is needed to determine statistical significance. We note that, besides the variability associated with a finite sample estimation of a putative underlying stochastic process, there are also other sources of variability for real-world data. To assess the effect of all sources of variability, and their effect on the estimated model, we propose to use the bootstrap resampling technique (Efron 1982) in conjunction with AMVAR analysis.

In bootstrap resampling, one randomly picks, with replacement, a set of resample trials, where the size of the resample should be the same as that of the original col-

lection. Replacement means that any given trial's probability of selection is the same every time that a trial is selected, so that the same trial may appear more than once in the same resample, as well as in different resamples. An MVAR model is fitted separately for each resample, and spectral quantities are derived from the model in the usual manner. By this procedure, a number of estimates equal to the number of resamples (which is usually chosen to be 100–200) is obtained for each spectral quantity. The mean and standard deviation are computed for the distribution across resamples of any spectral quantity. We emphasize that before fitting a model to each resample one needs to treat it as a new ensemble and perform the preprocessing steps of ensemble mean removal and standard deviation normalization.

As indicated above the original bootstrap procedure uses a resample size that equals the total number of trials to be analyzed. For many experimental data the trial size can be very large (in the thousands, for example). This presents a computational problem. We investigated the question of whether one can achieve an adequate variability assessment using smaller resample sizes. The example we treat here is the coherence between a pair of channels from prestriate and motor cortical areas, which showed a peak at 12 Hz in a 50-ms window centered at 120 ms poststimulus. Figure 8A displays the results of bootstrap resampling 100 resamples, where the resample size was systematically varied. The mean coherence is stable for resample sizes at or above 150, and the standard deviation is stable for resample sizes at or above 350. Although the standard deviation continues to decrease with increasing resample size, the rate of decrease is very slow. Thus, for practical purposes, we inferred that a resample size of approximately 350 trials was adequate for assessing variability in the present data set.

To decide whether or not this result is a special property of the particular channel pair that we used, we calculated the standard deviations for all pairs at the same frequency and time window. The result, shown in Fig. 8B, is that for all pairs, as for the single pair (Fig. 8A), the standard deviation decreases with increasing resample size. The decrease asymptotes around a resample size of 150.

4.6 Influence of model dimension and hidden variables on coherence estimation

In our multichannel ERP measurements, we were only able to sample a limited number of cortical sites. In other words, many relevant variables were hidden from our probe. A natural question is whether spectral measures derived from MVAR models fitted over all the available channels are robust against addition or deletion of channels. For example, does the coherence spectrum between two channels depend on the inclusion of other channels in the MVAR model from which it is derived? We performed the following test to examine this issue. First, for a 50-ms window centered at 120 ms, we fitted a model using data from all 15 available channels. The three pairwise coherence spectra among

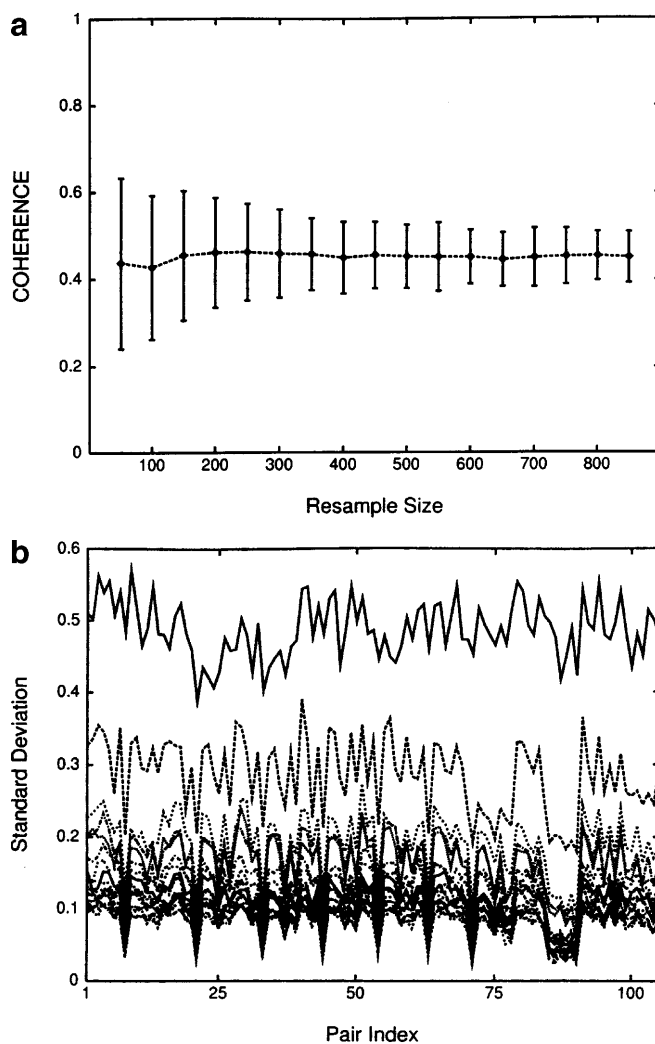


Fig. 8. **a** Convergence of the mean and standard deviation of the peak coherence measure (12 Hz at 120 ms) with increasing resample size for a single pair (prestriate and motor) of channels. **b** Convergence of the standard deviations of the same peak coherence measure for all pairs with increasing resample size. The largest standard deviation values at around 0.5 are for a resample size of 10, and the values progressively drop to around 0.1 with increasing resample sizes up to 400 (in increments of 20)

three channels (one striate and two prestriate) are shown in Fig. 9. Then, we used only data from the three involved channels to fit an MVAR model and derive the corresponding coherence spectra. The results from the 3-channel model, displayed in Fig. 10, show the same spectral features, such as number and location of peaks, as those from the 15-channel model, although there are some differences in the magnitudes of the peaks. This indicates that the coherence between two ERP channels is a robust quantity that is not greatly affected by inclusion or exclusion of other channels.

5 Results of AMVAR analysis of the visuomotor pattern discrimination task

The coherence results from the visuomotor pattern discrimination task illustrate the usefulness of the


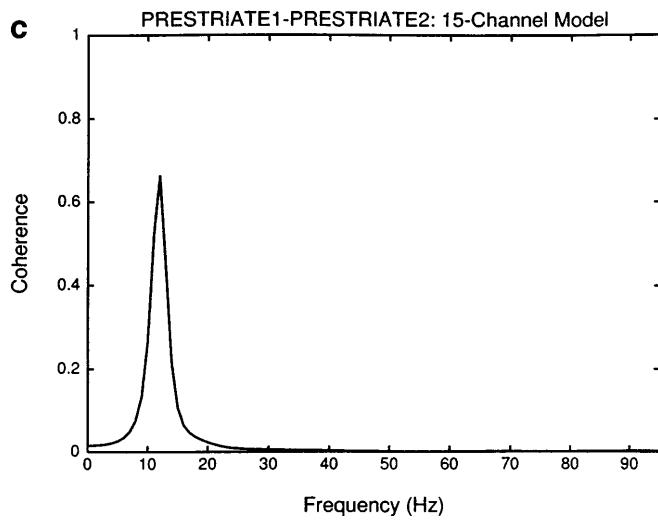

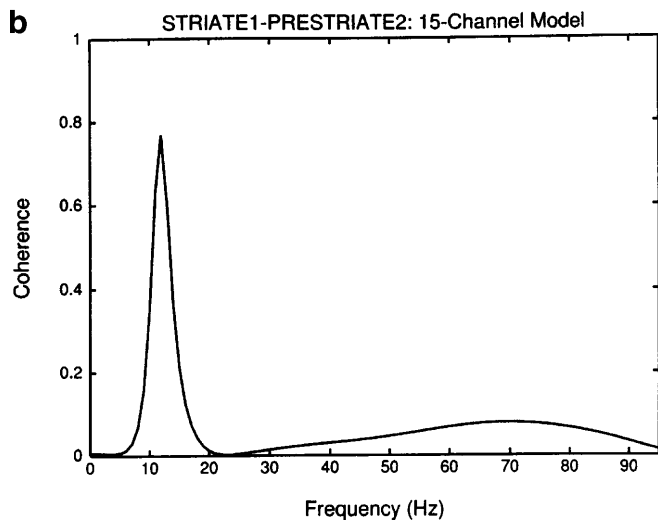




Fig. 9a-c. Pairwise coherence spectra among three channels (one striate and two prestriate) derived from a model based on all 15 channels

Fig. 10a-c. Pairwise coherence spectra among the same three channels as in Fig. 9, derived from a model based only on those three channels

AMVAR approach to spectral analysis of ERPs for understanding the neural substrate of cognitive processing. As mentioned above, the task required the animal to go through a series of cognitive states in the roughly one-half-second task period. We know of no other technique than the short-window AMVAR coherence analysis proposed here for assessing the rapidly shifting patterns of interdependency among cortical sites that




Fig. 11a-c. Patterns of peak intersite coherence in one monkey from the visuomotor pattern discrimination task, demonstrating the evolution of functional interdependency in three successive stages of the task. **a** The pattern of peak coherence in the period prior to arrival of the visual stimulus (0 ms). Two networks are observed: one involving ventral sites at 15 Hz (*dashed lines*), and the other involving dorsal sites at 22 Hz (*solid lines*). **b** The pattern of peak coherence following the onset of visual stimulus processing (120 ms). A large number of site pairs now exhibit significant coherence peaks at 12 Hz. **c** The pattern when the monkey executes a movement of the hand from the lever (280 ms). A new set of channel pairs are interdependent, with significant coherence peaks at 5 Hz

accompany these cognitive state changes. Figure 11 demonstrates the patterns of intersite coherence that develop over successive 50-ms windows at three different stages of task processing. The lines connecting cortical sites represent significant peak coherences. For a given channel pair significance was determined by comparison of its peak coherence value to a threshold level. This threshold was computed for a given window from a baseline distribution of coherence. The baseline distribution consisted of the coherence values at the same frequency over all channel pairs and all windows over the entire task. The threshold was set at two standard deviations above the mean of this distribution.

First, Fig. 11A shows the pattern of peak coherence in the period prior to arrival of the visual stimulus, but after the monkey has started the trial by depressing the hand lever. Here, in a window centered at 0 ms, two distinct networks are observed, one involving ventral sites at 15 Hz (*dashed lines*), and the other involving dorsal sites at 22 Hz (*solid lines*).

In Fig. 11B, we see a dramatic transformation in the pattern of peak coherence following the onset of visual stimulus processing. In this window, centered at 120 ms, the two earlier networks have been replaced by a single network at 12 Hz. Furthermore, a great many more site pairs are significantly involved in this network. This pattern reveals an extensive reorganization of the interdependency relations in the cortex as the cognitive state changes from the prestimulus state of anticipation and focused attention, to the poststimulus state of active processing of the visual feature information. The reorganization occurs in terms of changes in both the frequency of interdependency and the sites that are interdependent.

Finally, in Fig. 11C, the pattern once again undergoes an impressive change at a time, 280 ms, when the monkey actually executes a movement of the hand from the lever. Not only does the peak frequency shift to 5 Hz, but again there is a new set of channel pairs having coherence peaks at this frequency. Although it is beyond the scope of this report to detail all the coherence results from this experiment, we note that this pattern from the GO condition is strikingly different from the pattern from the same time window in the NO-GO condition.

From this cursory examination, we can clearly observe that these three different stages of cognitive processing are starkly resolved by the interdependency relations exhibited in the coherence patterns. Furthermore, these patterns are consistent with the known physiological and anatomical properties of the cortical areas involved. The fact that the patterns changed so completely over such a brief time span further underscores the necessity of using a short-window analysis to uncover these relations. [For details on the dynamic reorganization of large-scale networks seen here, see Bressler et al. (1999, submitted)].

6 Summary

The purpose of this article has been to establish that spectral analysis can be carried out for multichannel

ERPs during cognitive tasks by the use of Adaptive MultiVariate AutoRegressive (AMVAR) time series modeling over short overlapping time windows. The main results can be summarized as follows:

1. We pointed out the theoretical foundation of parametric short-window spectral analysis and demonstrated its effectiveness in a simple analytical model in Sect. 3.
2. We proposed a procedure in Sect. 4 for the objective selection of the window length for nonstationary ERP recordings.
3. We identified a set of preprocessing steps in Sect. 4 that are essential for the interpretable application of the model-fitting process: (i) detrend the single-trial ERP time series from each recording site, remove their temporal means, and normalize by the temporal standard deviation; (ii) compute the mean of the ensemble of trials at each site and subtract this mean from single-trial ERPs; and (iii) compute the ensemble standard deviation at each time point for a give site and divide single-trial ERPs by it.
4. We designed a set of tests in Sect. 4 that can be used to examine the validity of the model-fitting process. These tests are: (i) AIC model order selection; (ii) residual whiteness test; (iii) model stability test; and (iv) statistical percent consistency test.
5. We proposed a variability assessment strategy using bootstrapping in Sect. 4. We showed that, as a time-saving measure, one can use fewer numbers of trials in the resample than that of the original trial collection. It is worth emphasizing that for each bootstrap resample one needs to perform steps (ii) and (iii) of preprocessing, since each resample will be treated as a new ensemble for model fitting.
6. Finally, in Sect. 5 we applied the AMVAR spectral analysis technique to ERPs recorded during the GO condition of a visuomotor pattern discrimination task. The results clearly show the remarkable effectiveness of the technique by revealing task-relevant patterns of cortical interdependency during different stages of cognitive processing.

Acknowledgements. We thank Dr. Richard Nakamura for providing us with the experimental data used in this study and Dr. Partha Mitra for stimulating discussions on data analysis. This work was supported by NIMH grant MH-58190-02, NSF grant IBN-9723240, and ONR grant N00014-99-1-0062.

References

- Akaike H (1974) A new look at statistical model identification. *IEEE Trans Autom Control* AC-19: 716–723
- Bendat JS, Piersol AG (1986) *Random data: analysis and measurement procedures*. Wiley, New York
- Bressler SL, Coppola R, Nakamura R (1993) Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature* 366: 153–156
- Bressler SL, Ding M, Yang W (1999) Investigation of cooperative cortical dynamics by multivariate autoregressive modeling of event-related local field potentials. *Neuro Computing* 26–27: 625–631
- Efron B (1982) *The jackknife, the bootstrap, and other resampling plans*. SIAM, Philadelphia
- Florian G, Pfurtscheller G (1995) Dynamic spectral analysis of event-related EEG data. *Electroencephalogr Clin Neurophysiol* 95: 393–396
- Franaszczuk PJ, Blinowska KJ, Kowalczyk M (1985) The application of parametric multichannel spectral estimates in the study of electrical brain activity. *Biol Cybern* 51: 239–247
- Gersch W (1970) Spectral analysis of EEG's by autoregressive decomposition of time series. *Math Biosci* 7: 205–222
- Granger CWJ, Hughes AO (1968) Spectral analysis of short series—A simulation study. *J R Stat Soc Ser A* 130: 83–99
- Haykin S, Kesler S (1983) Prediction-error filtering and maximum-entropy spectral estimation. In: S Haykin (ed) *Nonlinear methods of spectral analysis*. Springer, Berlin Heidelberg New York, pp 9–72
- Jansen BH (1991) Time series analysis by means of linear modeling. In: Weitkunat R (ed) *Digital biosignal processing*. Elsevier, New York, pp 157–180
- Jenkins GM, Watts DG (1968) *Spectral analysis and its applications*. Holden-Day, San Francisco
- Lutkepohl H (1993) *Introduction to multiple time series analysis*, 2nd edn, Springer, Berlin Heidelberg New York
- Morf M, Vieira A, Lee D, Kailath T (1978) Recursive multichannel maximum entropy spectral estimation. *IEEE Trans Geoscience Electronics* 16: 85–94
- Muthuswamy J, Thakor NV (1998) Spectral analysis methods for neurological signals. *J Neurosci Methods* 83: 1–14
- Percival DB, Walden AT (1993) *Spectral analysis for physical applications*. Cambridge University Press, New York
- Pfurtscheller G, Haring G (1972) The use of an EEG autoregressive model for the time saving calculation of spectral power density distributions with a digital computer. *Electroencephalogr Clin Neurophysiol* 33: 113–115
- Schack B, Krause W (1995) Dynamics power and coherence analysis of ultra short-term cognitive processes – a methodical study. *Brain Topogr* 8: 127–136
- Whittle P (1963) On the fitting of multivariate autoregressions and the approximate factorization of a spectral matrix. *Biometrika* 50: 129–134