

# Evaluation and Comparison of EEG Traces: Latent Structure in Nonstationary Time Series

Mike WEST, Raquel PRADO, and Andrew D KRYSTAL

---

We explore and illustrate the use of time series decomposition methods for evaluating and comparing latent structure in nonstationary electroencephalographic (EEG) traces obtained from depressed patients during brain seizures induced as part of electroconvulsive therapy (ECT). Analysis of the patterns of change over time in the frequency structure of such EEG data provides insight into the neurophysiological mechanisms of action of this effective but poorly understood antidepressant treatment, and allows clinicians to modify ECT treatments to optimize therapeutic benefits while minimizing associated side-effects. Our work has introduced new methods of time:frequency analysis of EEG series that identify the complete pattern of time evolution of frequency structure over the course of a seizure, and usefully assist in these scientific and clinical studies. New methods of decomposition of flexible dynamic models provide time domain decompositions of individual EEG series into collections of latent components in different frequency bands. This allows us to explore ECT seizure characteristics via inferences on the time-varying parameters that characterize these latent components, and to relate differences in such characteristics across seizures to differences in the therapeutic effectiveness and cognitive side-effects of those seizures. This article discusses the scientific context and problems, development of nonstationary time series models and new methods of decomposition to explore time:frequency structure, and aspects of model fitting and analysis. We include applied studies on two data sets from recent clinical ECT studies. One is an initial illustrative analysis of a single EEG trace, the second compares the EEG data recorded during two types of ECT treatment that differ in therapeutic effectiveness and cognitive side-effects. The uses of these models and time series decomposition methods in extracting and contrasting key features of the seizure underlying the EEG signals are highlighted in these studies. Through the use of these models we have quantified, for the first time, decreases in the dominant frequencies of low frequency EEG components during ECT seizures. We have also identified preliminary evidence that such decreases are enhanced under the more effective ECTs at higher electrical dosages, a finding that is consistent with prior reports and the hypothesis that more effective forms of ECT are more potent at eliciting neurophysiological inhibitory processes.

**KEY WORDS:** Bayesian inference; Dynamic latent factors; Dynamic linear models; Electroconvulsive therapy (ECT); Electroencephalographic (EEG) time series; Time:Frequency analysis; Time series decomposition; Time-varying autoregressions

---

## 1. INTRODUCTION

Electroconvulsive therapy (ECT) in human subjects involves the induction of a series of brain seizures for therapeutic purposes and is the most effective treatment known for major depression (Weiner and Krystal 1994). The mechanisms of action of ECT, however, remain poorly understood, and a high level of interest exists in studies of electroencephalographic (EEG) recordings for their potential to improve the understanding of ECT and to impact on clinical practice by guiding ECT treatment designs (Krystal and Weiner 1994; Krystal et al 1993, 1995, 1998; Krystal, Weiner et al 1996; Krystal 1998; Nobler et al 1993). EEG traces are time series of electrical potential fluctuations recorded at various scalp locations, and reflect the physiological behavior of the underlying brain cells (Weiner et al 1991). Studies including those referenced above have identified a number of features of seizure

EEGs that differ as a function of the treatment therapeutic antidepressant potency and side-effects. Hence the attributes of these signals have promise for improving our understanding of the neurophysiological changes in brain function that are necessary to bring about a therapeutic antidepressant response. The clinical potential lies in the prospects for predicting the therapeutic response and side-effects associated with an ECT treatment, as such outcomes are presently impossible to determine until after a delay of several treatments (Krystal and Weiner 1994; Krystal 1998).

ECT seizure EEG series are typically globally nonstationary, exhibiting marked patterns of changes over time in frequency structure and amplitudes. Studies to date point to the importance of low-frequency (less than 5Hz) EEG activity during ECT treatments, and to amplitude suppression of EEG activity immediately after a seizure. It is thought that these aspects of seizure EEGs relate to changes in the relative degree of inhibitory input to cortical neurons elicited by the seizures, and so they serve as indicators of the degree of therapeutic antidepressant effect associated with the treatment (Post et al 1986; Sackeim et al 1991). Decreases in EEG frequency content over the course of a seizure are consistent with increases in inhibitory tone (Staton et al 1981; Weiner et al 1991), and so it is important to isolate and quantify the time-variation in frequency structure. Previous analyses have typically been limited to qualitative ratings by observers or the applica-

---

Mike West is the Arts and Sciences Professor of Statistics and Decision Sciences, and Director of the Institute of Statistics and Decision Sciences, at Duke University. The corresponding address is ISDS, Duke University, Durham, NC 27708-0251, USA, and the web site address is <http://www.stat.duke.edu>. Raquel Prado is Assistant Professor, Departamento de Cómputo Científico y Estadística, Universidad Simón Bolívar, Caracas, Venezuela. Andrew Krystal is Assistant Professor in the Department of Psychiatry and Behavioral Sciences and Associate Director of the Quantitative EEG Laboratory at Duke University Medical Center, Durham, NC 27710. This research was performed while the second author was a PhD candidate in ISDS at Duke University, and was partially supported by the National Science Foundation under grants DMS 9704432 and 9707914. The authors acknowledge the constructive comments of the Editor and an Associate Editor on an earlier version of this paper.

---

© 1999 American Statistical Association  
Journal of the American Statistical Association  
December 1999, Vol 94, No. 448, Applications and Case Studies

tion of standard stationary models to small segments of the series. Our work aims to extend prior approaches via appropriate global modeling of EEG traces that can effectively elucidate the complete pattern of time-evolution of frequency content over seizures. In addition to enhancing basic scientific understanding, improved analysis has potential to aid in assessing how differences in the patterns of EEG frequency content over the course of ECT seizures are related to treatment efficacy and side-effects. Our methods are explored here in two studies. In the first, we study a single channel of seizure EEG data to illustrate both methodology and the scientific insights generated from our approach. A second study explores differences between two EEG series generated under two types of ECT and for which available post-experimental evidence indicated differences in therapeutic anti-depressant potency and side-effects (Krystal et al 1995; Sackeim et al 1991, 1993). This represents the first attempt to identify differences in seizure EEG data as a function of treatment clinical parameters, and is part of a preliminary effort to use such models in improving the understanding of mechanisms of action of ECT and to impact on clinical practice.

We analyze EEG series using nonstationary, time-varying parameter autoregressive (TVAR) models, and use new theory of dynamic model decompositions to explore the time:frequency structure of latent processes underlying EEG signals. Such models are empirically justifiable in this applied context, and generate natural approaches to decomposition of observed time series into collections of latent components that may be interpreted in physical terms. Analysis of various kinds of EEG data using time-varying autoregressions has a history, notably in the work of W. Gersch and coauthors since the early 1980s (e.g., Gersch 1987; see references in Kitagawa and Gersch 1996). These and other authors have demonstrated the flexibility of high-order TVAR models in representing changing stochastic structure in stationary and nonstationary series in various applied fields, and particularly in modeling series with marked and stochastically time-varying periodicities such as are encountered in EEG studies. These authors have typically focused on questions of feedback and time lags in multi-channel EEG recordings via exploration of time variations in frequency structure as evidenced in patterns of change in the “instantaneous” spectral densities implied by the TVAR models. A TVAR model has the form of a standard autoregression but with autoregressive parameters that change through time. At any instant, the “current” values of the parameters may be plugged into the formula for the theoretical spectral density of a standard stationary autoregression to deliver what is referred to as the instantaneous spectrum of the nonstationary model. These ideas, introduced in Kitagawa (1983) and Kitagawa and Gersch (1985), relate closely to nonparametric methods of “evolutionary” spectral analysis using spectral estimates based on moving data “windows” (originating from Priestley 1965), though TVAR methods have proven much more widely useful in adapting to both smooth and more

abrupt changes through time in spectral characteristics (see discussion in Kitagawa and Gersch 1996, section 11.3).

Our methods may be viewed as natural extension of the TVAR literature that both complements and extends methodology in practically important directions. Our focus is on time domain decompositions of observed EEG traces, and the specific class of TVAR models provides a very general and flexible framework for this. Our methods explore time-varying aspects of the latent structure of the series across the continuum of frequencies, and constitute a general approach to time:frequency analysis of nonstationary processes that usefully extends the utility of TVAR models as they have been used to date. In particular, the identification of latent processes at unrestricted and often intersecting frequency ranges represents a first modeling approach to isolating EEG components that reflect the physical superposition of activity of multiple nonstationary neural generators (Geva et al 1995). Our methods isolate and estimate latent subseries that may be interpreted as representing brain wave activity in various nominal frequency bands, and whose time-varying parameters characterize the effects of the seizure on activity in these bands. Comparing such characteristics across EEG channels and between seizures under different ECT protocols provides insights into both basic scientific and clinical issues.

Time series decomposition methodology is based on extensions of existing time series decomposition theory (West 1997a) to time-varying parameter models. The new theoretical results underlying the methodology arise in general classes of dynamic linear models (DLMs) of which TVAR models form a very important subclass. In Section 2 we introduce TVAR models and describe the decomposition structure and theory. Section 3 discusses further aspects of model specification, fitting and posterior analysis. Section 4 discusses application in the initial, illustrative study of an EEG series generated in a recent study of antidepressant electroconvulsive therapies. Section 5 presents the treatment comparison study using data from 2-channel EEG recordings made during three types of ECT seizures which differ in their associated therapeutic potency. The use of model decompositions in isolating key differences in structure between treatments is explored here, indicating the potential for TVAR methods to assist in comparative treatment analysis via the evaluation and comparison of EEG traces. Section 6 provides summary comments.

Before introducing models and decomposition theory, some general comments on data are in order. Depending on experimental context and goals, a single EEG data sets comprises between 2 and 19 univariate time series recorded simultaneously at differing scalp locations on the patient under study, each location referred to as an EEG “channel”. Our univariate analyses select one or more channels for study. The raw data are 256Hz digitized records of voltage fluctuations following amplification and band-pass filtering, and we typically study sub-sampled series with between 40 and 50 equally spaced observations per second. Seizure records of between 1 and 3 minutes are typical, so that after sub-sampling we are dealing with time series of a few thousand equally spaced observations.

## 2. TIME SERIES MODELS AND DECOMPOSITION THEORY

A univariate time series  $x_t$ , ( $t = 1, 2, \dots$ ), follows a time-varying autoregressive model of order  $p$ , or TVAR( $p$ ), if

$$x_t = \sum_{j=1}^p \phi_{t,j} x_{t-j} + \epsilon_t \quad (1)$$

where  $\phi_t = (\phi_{t,1}, \dots, \phi_{t,p})'$  is the instantaneous autoregressive parameter vector at time  $t$  and  $\epsilon_t$  is a zero-mean innovation. Typically, and here, the innovations are assumed uncorrelated and normal,  $\epsilon_t \sim N(0, \sigma_t^2)$  with possibly time-varying variances  $\sigma_t^2$ . The model has the form of a standard autoregression at each time, but the autoregressive parameters and innovations variance may change through time. The AR parameters at each time  $t$  are not constrained to the usual stationary region, though often may lie in such a region. In such cases, we may refer to the series as “locally stationary” but the changes in parameters represent global nonstationarities. The model is very general, permitting both slow and more rapid or abrupt changes in parameters, and as a result provides a very flexible framework for modeling patterns of change in the stochastic structure and observed nonstationarities in data. The model is completed by specifying evolution model components for the time-varying parameters  $\phi_t$  and the innovations variance  $\sigma_t^2$ , and this is discussed in the following section together with aspects of model fitting and posterior computation. We first discuss the theoretical basis of our decomposition methodology and its implications in the class of TVAR models defined by equation (1) irrespective of the specific forms of evolution of  $\phi_t$  and  $\sigma_t^2$ .

The theory of time series decompositions in West (1997a) has important generalizations that are relevant to TVAR models and other dynamic models with time-varying parameters (including various classes of nonlinear models). The theory is accessed by casting the TVAR model (1) in dynamic linear model (DLM) form, as

$$x_t = \mathbf{F}' \mathbf{x}_t, \quad \mathbf{x}_t = \mathbf{G}_t \mathbf{x}_{t-1} + \boldsymbol{\omega}_t \quad (2)$$

where  $\mathbf{F} = (1, 0, \dots, 0)'$ ,  $\mathbf{x}_t = (x_t, x_{t-1}, \dots, x_{t-p+1})'$ ,  $\boldsymbol{\omega}_t = \epsilon_t \mathbf{F}$  and

$$\mathbf{G}_t \equiv \mathbf{G}(\phi_t) = \begin{pmatrix} \phi_{t,1} & \phi_{t,2} & \phi_{t,3} & \cdots & \phi_{t,p-1} & \phi_{t,p} \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & \cdots & \cdots & 1 & 0 \end{pmatrix}$$

for each  $t$ . This is one of several possible DLM forms of the TVAR model, and a natural extension of the DLM representation of standard AR models; the latter are obtained when  $\phi_t = \phi$  and  $\mathbf{G}_t = \mathbf{G}(\phi)$  are constant, and their state space forms originated with Akaike (1974). The time-varying versions in equation (2) are special cases of the broader class of DLMs, also known as a state space

models that date back at least to Kalman (1960). Our decomposition theory applies at that level of generality.

The central decomposition result arises simply from standard theory of model structure and the notions of similar models in linear stochastic systems (West and Harrison 1997, chapter 5). Suppose that, at each time  $t$ , the eigenvalues of  $\mathbf{G}_t$  are distinct, as is generally true for TVAR models of practical interest (related decomposition theory exists for important DLMs that have multiple common eigenvalues, but is not developed here; see West and Harrison 1997.) The  $p$  distinct eigenvalues will usually contain complex elements; suppose that, at time  $t$ , there are  $c$  pairs of complex conjugate eigenvalues and  $r = p - 2c$  real and distinct eigenvalues. Now, the numbers of real and complex eigenvalues may differ at different times  $t$ , but to begin here we assume that they are fixed in number so  $c$  and  $r$  are not subscripted by  $t$ . Later we revisit the possibility, and practical implications, of time-varying  $r$  and  $c$ . For now, the structure of the decomposition result is most clearly revealed assuming fixed numbers. Denote the complex eigenvalues by  $r_{t,j} \exp(\pm i\omega_{t,j})$  for  $j = 1, \dots, c$ , and the real eigenvalues by  $r_{t,j}$  for  $j = 2c + 1, \dots, p$ . Then  $\mathbf{G}_t = \mathbf{E}_t \mathbf{A}_t \mathbf{E}_t^{-1}$  where  $\mathbf{A}_t$  is the diagonal matrix of eigenvalues, in the arbitrary order specified, and  $\mathbf{E}_t$  is the  $d \times d$  matrix whose columns are the corresponding, complex-valued, normalized eigenvectors.

Reparametrize the model as follows. For each  $t$ , define the known matrix  $\mathbf{H}_t = \text{diag}(\mathbf{E}_t' \mathbf{F}) \mathbf{E}_t^{-1}$  and linearly transform  $\mathbf{x}_t$  to  $\boldsymbol{\gamma}_t = \mathbf{H}_t \mathbf{x}_t$ . Then equation (2) becomes

$$x_t = \mathbf{1}' \boldsymbol{\gamma}_t, \quad \boldsymbol{\gamma}_t = \mathbf{A}_t \mathbf{K}_t \boldsymbol{\gamma}_{t-1} + \boldsymbol{\delta}_t \quad (3)$$

where  $\mathbf{1} = (1, 1, \dots, 1)'$ ,  $\boldsymbol{\delta}_t = \mathbf{H}_t \boldsymbol{\omega}_t$  is a zero-mean normal innovation with a structured and singular variance matrix, and  $\mathbf{K}_t = \mathbf{H}_t \mathbf{H}_{t-1}^{-1}$ . With  $\boldsymbol{\gamma}_t = (\gamma_{t,1}, \dots, \gamma_{t,d})'$  this means that  $x_t$  is the sum of the individual  $\gamma_{t,j}$  processes. By construction, the final  $r$  elements of  $\boldsymbol{\gamma}_t$  are real, corresponding to the real eigenvalues  $r_{t,j}$  at each time  $t$ . Rename these real-valued processes  $y_{t,j}$ . The initial  $2c$  elements of  $\boldsymbol{\gamma}_t$  occur as complex conjugate pairs. Within each pair  $j = 1, \dots, c$ , the sum  $z_{t,j} = \gamma_{t,2j-1} + \gamma_{t,2j}$  is real and is simply evaluated as twice the common real part. Each real process  $z_{t,j}$  corresponds to the eigenvalue pairs  $r_{t,j} \exp(\pm i\omega_{t,j})$  at time  $t$ . The basic decomposition result may now be expressed as

$$x_t = \sum_{j=1}^c z_{t,j} + \sum_{j=2c+1}^d y_{t,j} \quad (4)$$

where the values of each of the real latent processes  $z_{t,j}$  and  $y_{t,j}$  may be directly evaluated based on the foregoing theory. Given known, estimated or simulated values of  $\mathbf{F}$ ,  $\mathbf{G}_t$  and  $\mathbf{x}_t$  for each  $t$ , we need to compute the eigen-decomposition and the transformations detailed above to evaluate all latent components. Exploring the forms of the  $z_{t,j}$  and  $y_{t,j}$  series over time is central to our time-frequency decomposition method. The framework is completed by identifying the structure of the component processes, and exploring the implications for interpretation of the decomposition.

*Example: The standard AR( $p$ ) model*

The AR( $p$ ) example, drawn from West (1997a), is useful at this point. In this case,  $\phi_t = \phi$  and hence  $\mathbf{G}_t = \mathbf{G}(\phi)$  is constant over time. The eigenvalues of  $\mathbf{G}$  are the reciprocals of the roots of the usual autoregressive characteristic equation, obviously also constant, i.e.,  $r_{t,j} = r_j$  for  $j = 1, \dots, p$  and, for the complex pairs of eigenvalues,  $\omega_{t,j} = \omega_j$  for  $j = 1, \dots, c$ . It easily follows that  $\mathbf{K}_t = \mathbf{I}$  and the transformed state evolution in equation (3) has a diagonalized form with  $\mathbf{A}_t \equiv \mathbf{A}$ , the matrix of constant eigenvalues. The final  $r$  elements of the transformed state are  $y_{t,j} = \gamma_{t,2c+j}$  which follow standard AR(1) processes with AR parameters  $r_j$ ,

$$y_{t,j} = r_j y_{t-1,j} + \nu_{t,2c+j} \quad (5)$$

for  $j = 1, \dots, r$ . The initial  $2c$  elements of  $\gamma_t$  are complex but once paired deliver the real processes  $z_{t,j} = \gamma_{t,2j-1} + \gamma_{t,2j}$ . Standard theory implies that

$$z_{t,j} = 2r_j \cos(\omega_j) z_{t-1,j} - r_j^2 z_{t-2,j} + \psi_{t,j} \quad (6)$$

where  $\psi_{t,j}$  is a real, zero-mean AR(1) process. Hence  $z_{t,j}$  follows an ARMA(2,1) model whose AR(2) component is quasi-periodic of *characteristic frequency*  $\omega_j$ . Sample paths of the  $z_{t,j}$  are quasi-periodic forms of time-varying amplitudes and phases. Commonly,  $|r_j| < 1$  in each case indicating stationary components, though the theory allows nonstationary components; practical models have  $|r_j| \leq 1$ . Note the close connections with alternative component models in West (1995, 1996a), having  $r_j = 1$ . The decomposition result here applies to give the underlying, latent AR(1) and quasi-cyclical ARMA(2,1) components of  $x_t$  over time  $t$ . Note that the processes are driven by dependent innovations as the  $\nu_{t,j}$  and  $\psi_{t,j}$  terms are all multiples of the original AR( $p$ ) innovation  $\epsilon_t$ . Note further that, in this special model, the decomposition is essentially that delivered by the partial fractions decomposition of an AR( $p$ ) process in the stationary case (Box and Jenkins 1976, section 3.2.1). The approach via DLMS provides a direct method of evaluating each of the latent processes in the decomposition, and also applies in nonstationary cases. Some examples of practical implications in spectral analysis and related time series studies, with model extensions, appear in West (1996b, 1997b), and Huerta and West (1998a,b).

Returning to the general TVAR model, the matrix  $\mathbf{K}_t$  will not be the identity, as is key to the special structure in the constant AR model above. However,  $\mathbf{K}_t$  will be close to the identity when the eigenvectors of  $\mathbf{G}_t$  and  $\mathbf{G}_{t-1}$  are similar, i.e., in cases when  $\phi_t$  changes only slowly. Notice that, given the eigenvectors,  $\mathbf{K}_t$  does not depend at all on the corresponding eigenvalues. Hence consecutive  $\mathbf{G}_t$  matrices with differing eigenvalues but the same eigenvectors will lead to  $\mathbf{K}_t = \mathbf{I}$  and the decomposition is *exactly* as in the standard AR example, but now with time-varying eigenvalues. In such a case, we have the decomposition (4) but

now equations (5) and (6) are generalized to time-varying versions

$$y_{t,j} = r_{t,j} y_{t-1,j} + \nu_{t,2c+j} \quad (7)$$

for  $j = 1, \dots, r$ , and

$$z_{t,j} = 2r_{t,j} \cos(\omega_{t,j}) z_{t-1,j} - r_{t,j}^2 z_{t-2,j} + \psi_{t,j} \quad (8)$$

for  $j = 1, \dots, c$ . Here  $y_{t,j}$  is now a time-varying autoregression, or TVAR(1) process, with AR parameters  $r_{t,j}$  at time  $t$ . Similarly,  $z_{t,j}$  follows a time-varying autoregressive, moving-average model of order (2,1), denoted by TVARMA(2,1). These are quasi-periodic processes with *characteristic frequencies*  $\omega_{t,j}$  changing slowly through time, together with slowly varying moduli  $r_{t,j}$ . Hence  $x_t$  is decomposed as the sum of latent TVAR(1) and TVARMA(2,1) processes, with the latter representing quasi-cyclical component structure at distinct characteristic frequencies  $\omega_{t,j}$  that change slowly through time. In cases of local stationarity with  $|r_{t,j}| < 1$ , the corresponding instantaneous spectral densities are peaked around  $\omega_{t,j}$  and the sharpness of the peaks are increasing functions of the  $r_{t,j}$ . Then the spectrum of the signal process  $x_t$  is time-varying, given at each instant  $t$  as the product of the instantaneous spectra of the  $y_{t,j}$  and of the AR part of the  $z_{t,j}$  processes.

In problems of applied interest with slowly time-varying  $\mathbf{G}_t$ , the  $\mathbf{K}_t$  matrices are very close identity matrices. In studies of a range of EEG series, typical  $\mathbf{K}_t$  matrices differ from the identity only to the order of  $10^{-5}$ , or less, element by element. As a result, the above representations and interpretations of the latent processes are then almost exact, and the above discussion stands. In other potential applications of TVAR models in which changes between consecutive  $\mathbf{G}_t$  matrices may be more marked, some of the  $\mathbf{K}_t$  matrices may depart more substantially from the identity. The result of this is that, between times  $t-1$  and  $t$ , the latent processes  $z_{t,j}$  and  $y_{t,j}$  are “mixed” via  $\mathbf{K}_t$  in (3). Hence the stochastic structure is not quite as represented by the TVAR(1) and TVARMA(2,1) components above, since the latent processes interact in this way. Let us stress that the decomposition theory applies generally, whatever levels of variation exist in the TVAR parameter. Implementation of the theory produces the latent processes and their time-varying parameters whatever the  $\mathbf{K}_t$  may be. It is the specific *interpretation* of the components that is a little less clear in cases of marked parameter changes – the resulting stochastic structure of the latent processes is not quite so neat as described above. Again, unless there are really marked changes in the eigenvectors of  $\mathbf{G}_t$ , the implied mixing of components is relatively benign and may be ignored, and the component processes viewed as having time series structure almost wholly dominated by the TVAR(1) and TVARMA(2,1) elements discussed. Hence, in practical models of our EEG series, we identify the parameters  $r_{t,j}, \omega_{t,j}$  as those of processes whose key features are TVAR(1) and TVARMA(2,1).

Before moving to aspects of model fitting and analysis, we note that the TVAR model characterizes all DLMS

with time-varying  $\mathbf{G}_t$  matrices of distinct eigenvalues in the same way that standard AR models characterize the class of *similar* constant DLMS having the same set of distinct eigenvalues (West and Harrison 1997, chapters 5 and 6 – two DLMS with constant but different system matrices are *similar DLMS* if the two system matrices are mathematically similar, i.e., if they have the same eigenvalues.) Thus, in a real applied sense, these models essentially encompass the class of practically interesting time-varying state space models with distinct eigenvalues.

### 3. MODEL SPECIFICATION AND POSTERIOR ANALYSIS

Completion of the TVAR model specification in equation (1) requires structuring of the time evolutions of  $\phi_t$  and  $\sigma_t^2$ , among other things. We adopt traditional random walk models for these evolutions, and so provide for estimation and adaptation to changing structure through time without anticipating specific directions of changes (West and Harrison 1997 chapter 3, section 9.6 and 10.8). Specifically, model (1) is completed by adding hierarchical components

$$\phi_t = \phi_{t-1} + \xi_t, \quad \xi_t \sim N(0, \mathbf{W}_t), \quad (9)$$

and

$$\sigma_t^2 = \sigma_{t-1}^2(\delta/\eta_t), \quad \eta_t \sim Be(a_t, b_t) \quad (10)$$

and where the series of stochastic elements  $\xi_t$  and  $\eta_t$  are independent and mutually independent, and also independent of the innovations series  $\epsilon_t$ . In equation (9), the TVAR parameters  $\phi_t$  evolve according to a standard first-order polynomial model, and the extent of potential parameter variations is controlled by the evolution variance matrices  $\mathbf{W}_t$ . These are defined via the standard discount method, as detailed in the Appendix. A single discount factor  $\beta$  in (0,1) leads to values of each  $\mathbf{W}_t$  such that low values of  $\beta$  imply high potential variability, and high values, in the range 0.9 – 0.999, are typically relevant in practice. Similarly, the changes in time of  $\sigma_t^2$  are modeled via the multiplicative random walk of equation (10), also standard in DLM analysis with time-varying variance components. The multiplicative innovations  $\eta_t$  are independent Beta variates, with Beta parameters  $(a_t, b_t)$  defined at each time  $t$  by a further discount factor  $\delta$  completely analogous to  $\beta$ . Again, this is fully detailed in the Appendix. In application, we treat the hyperparameters  $p$ ,  $\beta$  and  $\delta$  as tuning parameters to be specified, and report analyses based on values fixed at joint posterior modes with respect to uniform priors.

Model completion requires specification of an initial prior  $p(\phi_1, \sigma_1^2 | D_0)$  where  $D_0$  denotes initial information and all details of the model structure. Together with model components (9), (10) and chosen values of hyperparameters  $p$ ,  $\beta$  and  $\delta$ , this completes the specification of the prior distribution  $p(\Phi_k, \Sigma_k | D_0)$  where, for any positive integer  $k$ ,  $\Phi_k = \{\phi_1, \dots, \phi_k\}$  and  $\Sigma_k = \{\sigma_1^2, \dots, \sigma_k^2\}$ . The standard framework adopts relatively diffuse and conjugate initial

priors, as discussed in the Appendix. Then, given a series of  $n$  observed data values and the resulting information set  $D_n = \{D_0, y_1, \dots, y_n\}$ , analysis is based on full posterior distributions  $p(\Phi_n, \Sigma_n | D_n)$ . As summarized in the Appendix, standard sequential updating and retrospective filtering/smoothing algorithms for DLMS apply to deliver required components of the posterior. Particular marginal distributions of interest here are the multivariate T posteriors  $p(\phi_t | D_n)$  for  $t = 1, \dots, n$ , and summaries such as the trajectories of parameters  $E(\phi_{t,j} | D_n)$  over time  $t$ . These are useful in exploring the time variation in TVAR parameters. In connection with time series decomposition, focus rests on the eigenstructure of  $\mathbf{G}(\phi_t)$  at each time  $t$ . Given any point estimate of  $\phi_t$ , such as the posterior mean  $E(\phi_t | D_n)$ , we can compute the eigenvectors and eigenvalues of  $\mathbf{G}(\phi_t)$  and hence evaluate the corresponding moduli  $r_{t,j}$  and frequencies  $\omega_{t,j}$  of the root structure of the TVAR model at time  $t$ . To make more formal inferences about these, we simply draw samples from the multivariate T posterior  $p(\phi_t | D_n)$  and perform the eigenstructure analysis for each sampled value; this produces a sample from the joint posterior of the roots, the  $r_{t,j}$  and  $\omega_{t,j}$ , which can be graphically and numerically summarized. The values of the corresponding latent processes  $y_{t,j}$  and  $z_{t,j}$  in the key decomposition (4) can similarly be evaluated based on any specified value of  $\phi_t$ , so that we can deduce posterior estimates of the trajectories of these latent processes through time. Illustrations follow in the next two sections.

There are two complications in moving from the TVAR parameter  $\phi_t$  to the set of eigenvalues/characteristic roots  $r_{t,j} \exp(\pm i\omega_{t,j})$ . First, there is no inherent mathematical identification of the roots, and so no immediate identification of the corresponding latent processes. As frequency and amplitude characteristics of the latent components vary through time, a component that has the lowest frequency at one time may have a higher frequency later, for example. Similar comments apply to moduli and amplitudes of components. So we may experience “switching” effects as we look over the course of a time series. For posterior summary, identification must be enforced through an ordering, and we typically order the eigenvalues/roots at each time  $t$  in terms of their relative frequencies. Then interpretation of the posterior graphical and numerical summaries must bear in mind that components may “switch” from time to time as the data structure, and the model’s response, evolves.

The second, closely related issue is that the number of real/complex pairs of eigenvalues may differ at different times. The analysis described in Section 2 assumed fixed and constant numbers  $r$  and  $c$ , and the decomposition then follows easily, but complications arise as these numbers may change. The key reason for this is that collections of higher frequency components  $z_{t,j}$ , corresponding to complex roots, will often have fairly low moduli  $r_{t,j}$  and be apparent in the model decomposition as representations of high frequency noise; these are important model components but will typically be low in amplitude relative to more dominant, lower frequency and persistent components that have physical interpretations. With very

high frequency ranges, relatively small changes in the  $\phi_t$  parameters can lead to one or more pair of such complex roots “disappearing”, being substituted by real roots with low values and correspondingly low amplitudes of the induced real components in the series. The reverse of this phenomenon, complex roots substituting real roots, is also experienced. This methodological phenomenon needs recognizing although its practical import is negligible, as illustrated in each of the analyses of Sections 4 and 5 below.

#### 4. AN ILLUSTRATIVE EEG ANALYSIS

A first EEG series provides an initial, illustrative analysis to highlight key features of the problem and methodology. The full EEG data series, code named Ictal19 (“ictal” simply means seizure related) comes from a prior study involving 19 channel EEG data recorded in 8 subjects (Krystal, Greenside et al 1996; Zoldi et al 1996). During each seizure, 19 parallel EEG series are recorded in parallel from 19 electrodes located on patient’s scalp. Data recording uses Ag/CL electrodes on each of the 19 channels of the International 10-20 System, and utilizes a linked-ear reference and two additional channels dedicated to detecting eye-movement artifact (Krystal, Greenside et al 1996). The data underwent amplification and filtering with a band-pass of 1.6-70Hz, were then stored on magnetic tape and subsequently digitized off-line at 256Hz with 12 bits accuracy using the PC-based EEG acquisition and analysis system *EEGSYS* (Friends of Medical Science, Inc.). Prior to analysis the data underwent manual artifact rejection by the third author.

Previous analyses of sections of some of the Ictal19 data set have been discussed in Krystal, Greenside et al (1996), Prado and West (1997), and Zoldi et al (1996). Here we comment on a new analysis of just one channel, channel Ictal19-Cz which corresponds to the electrode positioned centrally at the very top of the head, referred to as the vertex in EEG nomenclature. We note that very similar patterns of behavior are evident across all channels; indeed, the issues of similarities and redundancies of multi-channel recordings have been highlighted in Prado and West (1997). This analysis represents the first application of time-varying modeling to ECT seizure EEG data and was intended to determine feasibility and identify its potential for reflecting the frequency content over time in ECT seizure EEG data.

Some sections of the data appear in Figure 1, the full series appearing also in Figure 3. The original recordings were subsampled every sixth observation to deliver 3,600 observations spanning the highest amplitude portion of the seizure episode, about 85 seconds at a rate of 42.67 observations per second. The subsampling is a standard practice adopted mainly for computational efficiency. No significant information is lost in moving from the original sampling rate of 256Hz to 42.67Hz as the key features of the EEG of interest are in moderate to lower frequency bands (4-15Hz). Furthermore, comparisons of analyses of original and subsampled series across various EEG data sets confirm that the resulting inferences on EEG com-

ponent structure are essentially unaffected by the down-sampling. Returning to Figure 1, note that the seizure is underway and reaching maturity at the start of the data window selected, then it continues at high amplitude before beginning to dissipate towards the end of the time series. Figure 1 displays four sections of 500 consecutive observations from near the start, two central sections and near the end of the recording period displayed. This clearly illustrates the intense oscillatory nature of the data and the changes throughout the course of the seizure in amplitude and frequency characteristics. The very rapid oscillations during the mature period of the seizure, decay in apparent frequency and amplitude towards the end of the episode. Evidently, a stable, constant parameter model of essentially any kind is untenable. However, subseries of consecutive observations of shorter lengths – say, 50-150 observations – are adequately represented by autoregressions with quasi-periodic components. Hence global TVAR models are relevant.

We discuss aspects of analysis using a TVAR(12) model with discount factors  $\beta = 0.994$  and  $\delta = 0.99$ . As mentioned earlier and detailed in the Appendix, the values of hyperparameters  $(p, \beta, \delta)$  are chosen with guidance from exploration of the marginal likelihood function computed over grids of values. The picture in this example is quite typical: the likelihood function is extremely flat over the ranges  $8 \leq p \leq 14$ ,  $0.99 \leq \beta \leq 0.999$  and  $0.98 \leq \delta \leq 0.999$ . Maximizing values of  $(p, \beta)$  are essentially constant at  $(11, 0.994)$  across this range of values for  $\delta$ . For the analysis reported, we fix discount factors at the maximizing values and conservatively fix  $p = 12$ , slightly larger than the strict MLE (but with almost the same likelihood value). Results, in terms of inference on component structure over time, are essentially unchanged in models with orders and discount factors in the ranges noted above, as is also typical. Higher order models were fitted obtaining similar results in terms of the decompositions and the trajectories of the estimated characteristic frequencies and amplitudes for the key low frequency components. As model order is increased too far, we tend to over-fit the noise, resulting in more high frequency components – of frequencies higher than 15Hz – with negligible relative contributions in amplitude to the decomposition of the observed series over the time range. Lower order models were also explored; as model order is decreased too far, we fail to capture all aspects of structure of the lowest frequency components at some time intervals, and it becomes difficult to recover scientifically interpretable components from the continuum of frequencies.

With these hyperparameter values, the analysis indicates moderate degrees of variation through time in both  $\phi_t$  and  $\sigma_t^2$ , as the seizure begins, matures and eventually decays. Time-variation in amplitude of the data fluctuations, and in the amplitudes of latent component processes underlying the data, are driven by changes in both  $\phi_t$  and  $\sigma_t^2$ . We note that the component-specific innovations  $\nu_{t,2c+j}$  in equation (7) and  $\psi_{t,j}$  in equation (8) have conditional variances that are increasing functions of  $\sigma_t^2$  but also depend on  $\phi_t$  in complicated ways. To provide a




Figure 1. Sections of EEG voltage levels from EEG data set Ictal19-Cz. From the top down, the graphs displays four sections of 500 observations from the full 3,600, taken at the start, early central, late central and final sections of the full series.

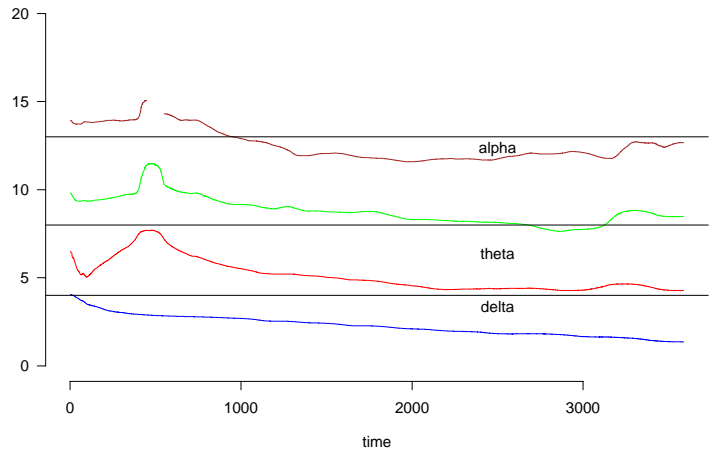


Figure 2. Trajectories of estimated characteristic frequencies of the four latent components of lowest frequencies in series Ictal19-Cz. The horizontal broken lines roughly demark the so-called delta, theta and alpha ranges of frequencies of neural rhythms (Dyro 1989).




Figure 3. Data and estimated latent components in the decomposition of Ictal19-Cz. From the bottom up, the graph displays the time series followed by the first four estimated components in order of increasing characteristic frequencies, each being a quasi-periodic TVARMA(2,1) sub-process. Components and data are plotted on the same vertical scale for direct comparison of relative amplitudes.




Figure 4. A central section of 500 observations of series Ictal19-Cz together with the corresponding sections of the first six underlying latent processes.

formal assessment of relative amplitudes of the latent component processes, we simply take the formulae for the theoretical marginal variances of each of the AR and ARMA models of (5) and (6), and evaluate these variances at each time point based on the posterior means of  $\phi_t$  and  $\sigma_t^2$ . The resulting instantaneous variances are used as measures of the relative amplitudes of the components.

Across the time course of the series the model eigenstructure exhibits a sustained set of at least three, usually four or more complex components with moduli and arguments that vary slowly throughout the seizure. Also, the

higher moduli correspond to the low frequency components in this analysis.

Figure 2 displays estimated values of the arguments of the first four eigenvalues in increasing order; these are evaluated at the posterior means  $E(\phi_t|D_n)$  at each time  $t \leq n$ . These arguments  $\omega_{t,j}$  are the instantaneous characteristic frequencies of the corresponding latent processes in equation (8). Each has an apparently decreasing form over time, consistent with the greater intensity of the seizure initially that forces higher frequency oscillations in the early and mature parts of the seizure, followed by gradual dissipation and decreases in characteristic frequencies. It must be remembered that these  $\omega_{t,j}$  represent peak frequencies

of the corresponding component processes which do have broad-band characteristics; these trajectories summarize movement over time in the peak of the corresponding component evolutionary spectra. The lowest frequency component lies in the “delta” frequency range which is the characteristic range of slow-waves manifested in the EEG during some types of seizures including the middle and late phases of ECT seizures (Niedermeyer 1993; Staton et al 1981; Weiner et al 1991; Weiner and Krystal 1993).

Figures 3 and 4 display estimates of the corresponding latent components  $z_{t,j}$  over time. These components are each plotted with the vertical scale set as the range of the data, so that their relative amplitudes are easily assessed. They are also ordered by  $\omega_{t,j}$ ; component  $j = 1$  is the delta/slow-wave. Notice that this delta wave dominates in amplitude and appears as a smoothed version of the data series; it can be viewed as the “seizure wave” for this episode. The second component in the “theta” frequency range is much lower in amplitude but represents a significant component process. Subsidiary components at higher frequencies, the “alpha” band, appear from the decomposition, and are also persistent, though are much lower in amplitude than the lower frequency components. This decomposition is an explicit representation of expected phenomena: the seizure EEG is dominated by delta frequency slow-waves, which decrease in frequency as the seizure progresses (Staton et al 1981; Weiner et al 1991; Weiner and Krystal 1993). Since these slow-waves are thought to reflect the degree of inhibitory input to cortical neurons during seizures, these findings are consistent with a relative increasing dominance of inhibitory tone elicited during ECT seizures (Pedley and Traub 1990; Sackeim et al 1991). The elicitation of inhibition is of particular interest in the light of preliminary evidence that this inhibitory, anti-convulsant response is central to the anti-depressant efficacy of ECT (Krystal et al 1995, 1998; Krystal and Weiner 1996; Nobler et al 1993; Post et al 1986; Sackeim et al 1991). The decrease in slow-wave frequency, evident in Figure 2, represents the first time that this phenomenon has been quantified for ECT seizure EEG data. Several intriguing features are evident. First, the lowest frequency component manifests a consistent decrease in frequency content over the entire seizure which is initially nonlinear and then appears to be linear and gradual thereafter. The higher frequency components are different but all appear to have an initial increase in frequency followed by a consistent gradual decline in frequency. These findings suggest that there may be at least two different physiologic phenomena contributing to the frequency content manifest in the seizure EEG data. Second, this analysis suggests that it is feasible to quantify the degree of decline in frequency of the latent components over the course of the seizure and therefore the pattern of decline for these components may be studied further for their capacity to differentiate ECT seizures as a function of clinical treatment; this aspect of our work is illustrated in the study of Section 5.

Finally, we comment again on the statistical issues of identification in isolating latent components. The above

discussion is based on the eigenstructure of the estimated state matrix  $\mathbf{G}(\phi_t)$ , evaluated at  $E(\phi_t|D_n)$ , with the ordered in terms of increasing arguments, or frequencies, at each  $t$ . Throughout most of the time series the state matrix has a full six pairs of complex conjugate eigenvalues, corresponding to six latent  $z_{t,j}$  processes as in equation (8). For short periods however, one or more of these roots at very high frequencies vanishes, being substituted by a pair of distinct real roots. This is evident just once in the analysis as displayed here, at around  $t = 450 - 500$  where, as illustrated in Figure 2, the frequency trajectory of the higher frequency component exhibits a “break.” In this brief time interval, the corresponding  $\mathbf{G}(\phi_t)$  matrix “loses” its fourth pair of complex conjugate eigenvalues, switching back to (at least) four again after a few time periods. Evidently, however, the event has little impact on our interpretation and understanding of the component structure.

## 5. ECT TREATMENT COMPARISONS

Much of our scientific interest is focused on questions of the identification of EEG differences between ECT seizures that differ in therapeutic antidepressant potency and treatment side-effects. In this connection, this section concerns comparative discussion of EEG series obtained from two seizures elicited in the same subject by forms of ECT that differ in their associated therapeutic potency and side-effects. The data arises from a larger experiment in which subjects were randomized to three different types of ECT treatments and provides a test-bed for a preliminary determination of the extent to which TVAR models can aid in identifying patterns of evolution in frequency content that are related to the effectiveness and side-effects of ECT treatments. In this analysis we compared right unilateral (UL) ECT treatments administered at two different electrical stimulus intensities. The intensities compared were Low (barely above the seizure threshold) and Mod (moderately supra-threshold). Several studies suggest that the latter form of treatment is significantly more therapeutically effective but also associated with more side-effects (Krystal et al 1995; Sackeim et al 1991). Based on prior preliminary work and the analysis of Section 4, we expected that there would be evidence of smaller slow-wave amplitude, and less evidence of suppression of EEG activity after the end of the seizure in Low ECT compared with Mod ECT. The data in this study were recorded with Ag/Cl electrodes from left prefrontal to left mastoid and right prefrontal to right mastoid derivations. The EEG data were digitized as in Section 4, and we select one channel from each seizure for analysis.

Figure 5 displays several sections of the EEG series from an individual labeled S26 under the moderate stimulus level, referred to as S26.Mod. The format is exactly as in Figure 1, and it is immediately clear that the waveforms are radically different, evidencing a more “spiked” seizure wave with more detailed high frequency structure. The full seizure record for S26.Mod also appears in Figure 6. The data window begins just after the onset of seizure, and again the data graphed are at a rate of 42.67 observations

per second so that the time frame of the 3,200 observations graphed is 75 seconds.




Figure 5. Sections of EEG voltage levels from EEG series S26.Mod. From the top down, the graph displays four sections of 500 observations from the full 3,200, taken at the start, early central, late central and final sections of the full series.

The second seizure of this same individual under the lower treatment level produces the EEG trace denoted by S26.Low and graphed in Figure 7. We compare aspects of the two seizures through examination of inferences on the component structure over time for each of the two series. Some aspects of analysis are illustrated in Figures 6-11 inclusive. These are based on TVAR(20) models for both series, and discount factors  $\beta = 0.994$  and  $\delta = 0.95$ , with values again selected after exploring the marginal likelihood functions for hyperparameters and then compromising on common values close to the MLEs for the two individual analyses. Extensive reanalysis on discrete grids of hyperparameters around those selected confirm that the marginal likelihood functions are, in each of the two cases, very flat around the selected values which are close to the MLEs, and inferences reported and graphed do not vary meaningfully under such sensitivity analyses. Following model fitting, the decomposition analysis is applied as in the previous section, and resulting estimates of several latent components in each of S26.Mod and S26.Low are graphed over time in Figures 6 and 7. It turns out that each series appears to be composed of two distinct delta waves and two distinct theta waves, together with a collection of higher frequency alpha waves and residual components. Compared to the analysis of the previous section, a higher order model is supported and needed here due to the more complicated structure in terms of the high frequency components and “spiky” oscillations in the series. High frequency components have a major relevant contribution in amplitude to the series in both treatments, as illustrated in Figures 6 and 7 where the components are all graphed on the same scale. This explains the need for a

higher order model to account, not only for slow frequency components that are key components in the decomposition of the observed series, but also for high frequency components that represent rapid neural oscillations, neural and experimental noise.

The significant levels of high frequency activity here contrast markedly with the structure of the series studied in Section 4. This may be due in part to individual differences between the subjects studied, but is also likely a reflection of the fact that this data set was recorded pre-frontally; that in Section 4 was recorded over the vertex where the slow wave activity dominates the signal to a greater extent (Krystal, Greenside et al 1996; Zoldi et al 1996). Also in contrast to Section 4, note that the graphs in Figures 6-11 display data well beyond the seizure end points, into the post-ictal period. This later data has been found to be important in comparing differences between ECT treatments as a function of efficacy and side-effects (Krystal et al 1993, 1995, 1996, 1998; Nobler et al 1993). The end of the seizure in Figure 6 is evident by the decrease in amplitude, which occurs at around  $t = 1800$ ; that in Figure 7 occurs at around  $t = 1000$ . To more explicitly compare the frequency content of Low and Mod ECT seizures, Figure 9 graphs estimated trajectories of the instantaneous frequencies of three key slow wave components - two in the delta frequency range and one in the slightly higher theta frequency band - and Figures 10 and 11 compare their estimated amplitudes. These model based inferences confirm our hypothesis that the magnitudes of the slow-waves and their evolution in frequency across the seizure are greater in the Mod than in the Low ECT seizure.

Figure 9 indicates that the characteristic frequencies of the key slow waves decrease to a greater extent over the course of the seizure in S26.Mod than in S26.Low. It is important to note that the two seizures end at different points; the S26.Low seizure course runs from about  $t = 0$  to  $t = 1000$ , while that for S26.Mod is about  $t = 0$  to  $t = 1800$ . The characteristic frequencies of slow waves of S26.Low are stable over time, whereas those for S26.Mod exhibit a general decay during the key later part of the seizure. The relative stability in these characteristic frequencies for S26.Low is also notable in Figure 8, where during the seizure the lower frequency trajectories remain essentially flat. These findings indicate that, indeed, the ECT seizure associated with the higher electrical dosage, S26.Mod, exhibits decreases in the frequencies of slow wave activity during the seizure, suggesting that, in contrast to the less potent S26.Low, the seizure elicits an increasing degree of inhibition of cortical neurons over the seizure. As higher stimulus seizures are generally found to be more therapeutically potent, this lends support to the hypothesis that more effective treatments elicit a greater amount of inhibitory activity (Post et al 1986; Sackeim et al 1991). This analysis represents the first time that this behavior has been quantified and suggests that TVAR models may be useful for improving the utility of ECT seizure EEG data through reflecting the change in frequency content over the course of these seizures.




Figure 6. Data and some of the estimated latent components in the decomposition of EEG series S26.Mod. From the bottom up, the graph displays the time series followed by estimated components in order of increasing characteristic frequencies. The low frequency delta and theta components are individual quasi-periodic TVARMA(2,1) processes, and the alpha component is the sum of two such processes in the higher frequency alpha band. The final “fast” component is the sum of all remaining, very high frequency components representing both rapid neural oscillations, neural and experimental noise.




Figure 7. Data and some of the estimated latent components in the decomposition of EEG series S26.Low, in a format exactly as for S26.Mod in Figure 6.




Figure 8. Trajectories of estimated characteristic frequencies of all latent components in EEG series S26.Low. Note the identification-based “switching” behavior apparent at very high frequencies, especially after the seizure ends, and the relative stability of the estimated characteristic frequencies of the dominant low frequency components.

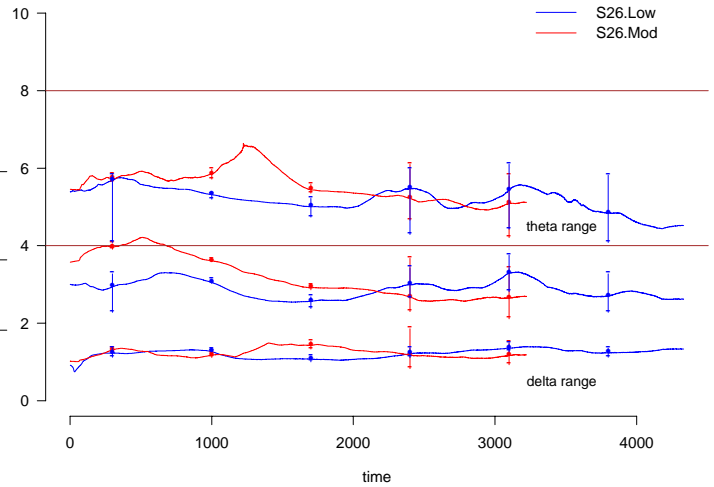


Figure 9. Trajectories of estimated characteristic frequencies of key low frequency components in decompositions of EEG series S26.Mod and S26.Low, now with approximate 95% posterior intervals for instantaneous frequencies at selected time points. The intervals are computed by direct simulation of the posterior distributions for TVAR parameters at the selected time points, the sampled values leading by direct calculation to corresponding samples from the posterior distributions for the instantaneous frequencies. Note that the seizure for S26.Low ends at approximately  $t=1000$ , while S26.Mod ends at about  $t=1800$ .

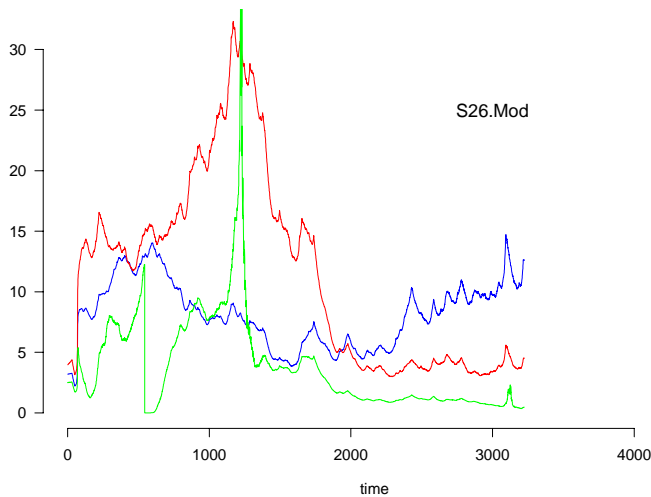


Figure 10. Trajectories of estimated amplitudes of key low frequency components in decomposition of EEG series S26.Mod

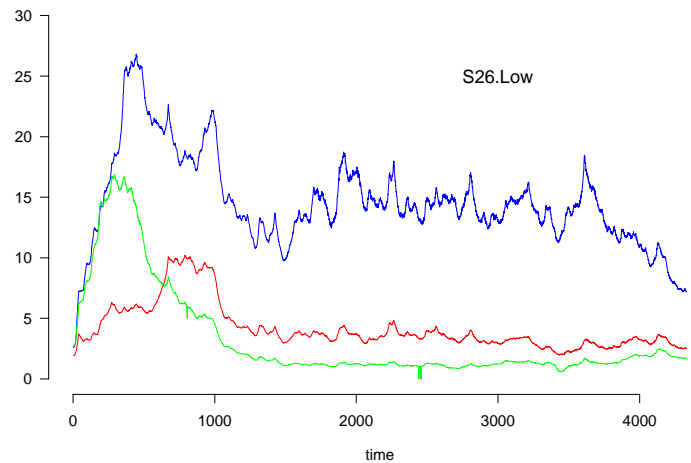


Figure 11. Trajectories of amplitudes of key low frequency components in decomposition of EEG series S26.Low.

Evidence of greater inhibition elicited by the higher dosage S26.Mod can also be seen in Figures 10 and 11. These depict the estimated amplitudes of the low frequency components of the two seizures. The magnitude of the slow-wave component rises more rapidly and to a much greater extent for S26.Mod than for S26.Low during the seizure. This finding is also consistent with a greater degree of inhibition elicited by the more effective, higher electrical dosage, S26.Mod (Pedley and Traub 1990; Sackeim et al 1991). Also consistent with this is evidence of greater immediate post-ictal suppression of EEG activity with this form of treatment. In Figures 6 and 7 it is apparent that the amplitude of seizure activity decreases more dramatically after the end of S26.Mod (after roughly  $t = 1800$ ) than after the end of the S26.Low seizure (after roughly  $t = 1000$ ). This same phenomenon is also reflected particularly nicely in Figures 10 and 11 where the differences in amplitudes between the seizure and post-seizure periods in low frequency components is much greater in Figure 10 (note the dramatic decrease in amplitude at about  $t = 1800$ ) than in Figure 11. This phenomenon is also suggestive of greater inhibition in the higher dosage S26.Mod seizure, consistent with previous studies of seizures that are higher above the seizure threshold and are more therapeutically effective (Krystal et al 1993, 1995, 1996, 1998; Nobler et al 1993).

## 6. DISCUSSION

The studies presented are indicative of the uses of decomposition methods using TVAR models for EEG series, and of their promise as aids to improving scientific understanding and to assist in clinical ECT practice. The analysis of Section 4 provides a nice illustration of the uses of the approach in extracting and estimating latent EEG components in various key frequency bands, and in exploring the changes over time in frequency structure of

the EEG via this decomposition. For the first time, the changes in characteristics of seizure slow waves have been explicitly inferred, with resulting identification of the decay over time of such dominant low frequency activity over the course of some ECT seizures. The analysis of Section 5 represents a first model based comparison of EEG data obtained from ECT seizures that differ in terms of electrical dosage, therapeutic potency, and side-effects. This analysis was consistent with prior work suggesting a more rapid onset and higher amplitude slow-wave activity with greater post-ictal EEG suppression in the more potent, higher dosage form of treatment. It also provided the first evidence that the more effective, higher dosage, treatment is associated with seizures that appear to decrease in the characteristic frequency of low frequency components to a greater extent. This finding provides further support for the hypothesis that higher dosage ECT may be more potent at eliciting neurophysiological inhibition and that this process may be a central part of the antidepressant mechanism of action of ECT. A determination of whether there is a greater decrease in characteristic frequency of low frequency seizure components in higher dosage, more effective ECT treatments is deserving of further study. The data presented in this article provide preliminary support that such a study may improve the understanding of the mechanism of action of ECT and may make seizure EEG data more useful for allowing practitioners to optimize the dosing of ECT treatments. We anticipate a range of studies of new EEG data arising from clinical investigations geared to that goal, and also wider use of the TVAR modeling and decomposition approach by other investigators.

On models and methodology, we have discussed aspects of more-or-less standard Bayesian analysis with nonstationary dynamic models, the TVAR framework being a key and important subclass of a broader class of dynamic models in which the decomposition theory arises. In related studies (Prado and West 1997; Aguilar et al 1998) we have

touched on multivariate time series issues through comparisons of univariate analyses, as presented here, across several or all of the EEG channels recorded during a single seizure episode. The notion of multivariate modeling to explore issues of data reduction and the spatio-temporal connectivities between EEG channels is exciting in connection with the potential to contribute to further advances in both basic science and in clinical decision making. We anticipate such advances in further collaborative work.

[Received October 1997. Revised December 1998.]

## REFERENCES

- Aguilar, O., Huerta, G., Prado, R., and West, M. (1998), "Bayesian Inference on Latent Structure in Time Series (with discussion)," in *Bayesian Statistics 6*, eds: J.O. Berger, J.M. Bernardo, A.P. Dawid and A.F.M. Smith, Oxford: University Press (to appear).
- Akaike, H. (1974), "Markovian Representation of Stochastic Processes and its Application to the Analysis of Autoregressive Moving Average Processes," *Ann. Inst. Statist. Math.*, **26**, 363-387.
- Box, G.E.P., and Jenkins, G.M. (1976), *Time Series Analysis: Forecasting and Control*, (2nd Edition), San Francisco: Holden-Day.
- Dyro, F.M. (1989), *The EEG Handbook*, Boston: Little, Brown and Co.
- Gersch, W. (1987), "Non-stationary Multichannel Time Series Analysis," in *EEG Handbook, Revised Series*, (Vol. 1), ed. A. Gevins, New York: Academic Press.
- Geva, A.B., Pratt, H., and Zeevi, Y.Y. (1995), "Spatio-temporal Multiple Source Localization by Wavelet-Type Decomposition of Evoked Potentials," *Electroenceph. Clin. Neurophysiol.*, **96**, 278-286.
- Huerta, G., and West, M. (1998a), "Priors and Component Structures in Autoregressive Time Series," *J. Roy. Statist. Soc., Ser. B.* (to appear).
- Huerta, G., and West, M. (1998b), "Bayesian Inference on Periodicities and Component Spectral Structure in Time Series," *J. Time Series Anal.*, (to appear).
- Kalman, R.E. (1960), "A New Approach to Linear Filtering Theory and Prediction Problems," *Trans. ASME J. Basic Engrg.*, Ser. D, **82**, 35-45.
- Kitagawa, G. (1983), "Changing Spectrum Estimation," *J. Sound and Vibration*, **89**, 443-445.
- Kitagawa, G., and Gersch, W. (1985), "A Smoothness Priors Time Varying AR Coefficient Modeling of Nonstationary Time Series," *IEEE Trans. Aut. Control*, **AC-30**, 48-56.
- Kitagawa, G., and Gersch, W. (1996), *Smoothness Priors Analysis of Time Series*, Lecture Notes in Statistics #116, New York: Springer-Verlag.
- Krystal A.D. (1998), "The Clinical Utility of Ictal EEG Seizure Adequacy Models," *Psych. Ann.*, **28**, 30-35.
- Krystal, A.D., Coffey, C.E., Weiner, R.D., and Holsinger, T. (1998), "Changes in Seizure Threshold Over the ECT Course Affect Therapeutic Response and are Detected by Ictal EEG Ratings," *J. Neuropsychiatry Clin. Neurophysiol.*, **10**, 178-186.
- Krystal, A.D., Greenside, H.S., Weiner, R.D., and Gassert, D. (1996), "A Comparison of EEG Signal Dynamics in Waking, after Anesthesia Induction and during Electroconvulsive Therapy Seizures," *Electroenceph. Clin. Neurophysiol.*, **99**, 129-140.
- Krystal, A.D., and Weiner, R.D. (1994), "ECT Seizure Therapeutic Adequacy," *Convulsive Therapy*, **10**, 153-164.
- Krystal, A.D., Weiner, R.D., and Coffey, C.E. (1995), "The Ictal EEG as a Marker of Adequate Stimulus Intensity with Unilateral ECT," *J. Neuropsychiatry Clin. Neurophysiol.*, **7**, 295-303.
- Krystal, A.D., Weiner, R.D., Gassert, D., McCall, W.V., Coffey, C.E., Sibert, T., and Holsinger, T. (1996), "The Relative Ability of Three Ictal EEG Frequency Bands to Differentiate ECT Seizures on the Basis of Electrode Placement, Stimulus Intensity, and Therapeutic Response," *Convulsive Therapy*, **12**, 13-24.
- Krystal, A.D., Weiner, R.D., McCall, W.V., Shelp, F.E., Arias, R., and Smith, P. (1993), "The Effects of ECT Stimulus Dose and Electrode Placement on the Ictal Electroencephalogram: An Intra-individual Crossover Study," *Biological Psychiatry*, **34**, 759-767.
- Niedermeyer, E. (1993), "Epileptic Seizure Disorders," in *Electroencephalography*, (3rd Edition), eds: E. Niedermeyer and F. Lopes da Silva, Baltimore: Williams and Wilkins.
- Nobler, M.S., Sackeim, H.A., and Solomou M. (1993), "EEG Manifestations during ECT: Effects of Electrode Placement and Stimulus Intensity," *Biol. Psychiatry*, **34**, 321-330.
- Pedley, T.A., and Traub, R.D. (1990), "Physiological Basis of the EEG," in *Current Practice of Clinical Electroencephalography*, (2nd Edition), eds: D.D. Daly and T.A. Pedly, New York: Raven Press.
- Post, R.M., Putnam, F., and Uhde, T.W. (1986), "Electroconvulsive Therapy as an Anticonvulsant: Implications for its Mechanism of Action in Affective Illness," *Ann. Nat. Acad. Sci.*, **462**, 376-388.
- Prado, R., and West, M. (1997), "Exploratory Modelling of Multiple Non-stationary Time Series: Latent Process Structure & Decompositions," in *Modelling Longitudinal and Spatially Correlated Data*, ed: T. Gregoire, New York: Springer-Verlag.
- Priestley, M.B. (1965), "Evolutionary Spectra for Nonstationary Processes," *J. Roy. Statist. Soc., Ser. B.*, **27**, 204-229.
- Sackeim, H.A., Devanand, D.P., and Prudic, J. (1991), "Stimulus Intensity, Seizure Threshold, and Seizure Duration: Impact on the Efficacy and Safety of Electroconvulsive Therapy," *Psychiatric Clin. North Am.*, **14**, 803-843.
- Sackeim, H.A., Prudic, J., Devanand D.P., Kiersky J.E., Fitzsimons L., Moody B.J., McElhiney, M.A., Coleman E.A., and Settembrino J.M. (1993), "Effects of Stimulus Intensity and Electrode Placement on the Efficacy and Cognitive Effects of Electroconvulsive Therapy," *New Eng. J. Med.*, **328**, 839-846.
- Staton, R.D., Hass, P.J., and Brumback, R.A. (1981), "Electroencephalographic Recording during Bitemporal and Unilateral Non-dominant Hemisphere (Lancaster Position) Electroconvulsive Therapy," *J. Clin. Psychiatry*, **42**, 264-269.
- Weiner, R.D., Coffey, C.E., and Krystal, A.D. (1991), "The Monitoring and Management of ECT Seizures," *Psychiatric Clin. North Am.*, **14**, 845-869.
- Weiner, R.D., and Krystal, A.D. (1993), "EEG Monitoring of ECT Seizures," in *The Clinical Science of Electroconvulsive Therapy*, ed: C.E. Coffey, Washington, D.C.: American Psychiatric Press, Inc.
- Weiner, R.D., and Krystal, A.D. (1994), "The Present Use of Electroconvulsive Therapy," *Ann. Rev. of Medicine*, **45**, 273-281.
- West, M. (1995), "Bayesian Inference in Cyclical Component Dynamic Linear Models," *Journal of the American Statistical Association*, **90**, 1301-1312.
- West, M. (1996a), "Some Statistical Issues in Palaeoclimatology (with discussion)," in *Bayesian Statistics 5*, eds: J.O. Berger, J.M. Bernardo, A.P. Dawid and A.F.M. Smith, Oxford: University Press.
- West, M. (1996b), "Bayesian Time Series: Models and Computations for the Analysis of Time Series in the Physical Sciences," in *Maximum Entropy and Bayesian Methods 15*, eds: K. Hanson and R. Silver, Dordrecht: Kluwer.
- West, M. (1997a), "Time Series Decomposition," *Biometrika*, **84**, 489-494.
- West, M. (1997b), "Modelling and Robustness Issues in Bayesian Time Series Analysis (with discussion)," in *Bayesian Robustness 2*, eds: J. Berger, F. Ruggeri, and L. Wasserman, Lecture Notes-Monograph Series (Vol. 29), Hayward CA: Institute of Mathematical Statistics.
- West, M., and Harrison, P.J. (1997), *Bayesian Forecasting and Dynamic Models* (2nd Edition), New York: Springer-Verlag.
- Zoldi, S.M., Greenside, H.S., and Krystal, A.D. (1996), "Redundancy and Stationarity of 21-Channel EEG Data Recorded During ECT Seizures," *J. Clin. Neurophysiol.*, **5**, 440.

## APPENDIX: MODEL FITTING AND POSTERIOR COMPUTATION

The equations defining required posterior distributions all arise via direct application of standard DLM theory, as detailed in West and Harrison (1997, especially sections 4.3-4.7 and 10.8). Summary details and formulæ for the analyses reported are given here.

### Sequential Updating

The TVAR model (1) has the form

$$x_t = \mathbf{x}'_{t-1}\boldsymbol{\phi}_t + \epsilon_t$$

for all  $t > 0$ . Coupled with equations (9) and (10) and conditional on a specified set of values of the hyperparameters  $(p, \beta, \delta)$ , this defines a standard dynamic regression DLM with time-varying observation variance  $\sigma_t^2$

In this context, we first sequentially process the  $x_t$  observations to obtain sequentially updated or “on-line” posteriors  $p(\boldsymbol{\phi}_t, \sigma_t^2 | D_t)$  over  $t = 1, \dots, n$ . These have the conjugate normal/inverse gamma forms. The important posterior margins for each  $\boldsymbol{\phi}_t$  and  $\sigma_t^2$  are given by  $(\boldsymbol{\phi}_t | D_t) \sim T_{k_t}(\mathbf{m}_t, \mathbf{C}_t)$ , a multivariate T distribution with  $k_t$  degrees of freedom, location vector  $\mathbf{m}_t$  and scale matrix  $\mathbf{C}_t$ , and  $(\sigma_t^{-2} | D_t) \sim Ga(k_t/2, d_t/2)$ , a gamma distribution with shape parameter  $k_t/2$  and scale parameter  $d_t/2$ . The updating equations for defining parameters are summarized here; these formulæ are used to sequentially compute the parameters for all  $t$ . Note that the usual point estimates of  $\sigma_t^2$ , namely  $s_t = d_t/k_t$ , appear in these formulæ. For  $t = 2, \dots, n$ , we have:

$$\begin{aligned} \mathbf{m}_t &= \mathbf{m}_{t-1} + \mathbf{P}_t e_t, \\ \mathbf{C}_t &= (\mathbf{R}_t - \mathbf{P}_t \mathbf{P}'_t q_t)(s_t/s_{t-1}) \\ k_t &= \delta k_{t-1} + 1, \\ d_t &= \delta d_{t-1} + s_{t-1} e_t^2 / q_t, \end{aligned} \quad (11)$$

where

$$\begin{aligned} e_t &= x_t - f_t, \\ f_t &= \mathbf{x}'_{t-1} \mathbf{m}_{t-1}, \\ q_t &= \mathbf{x}'_{t-1} \mathbf{R}_t \mathbf{x}_{t-1} + s_{t-1}, \\ \mathbf{P}_t &= \mathbf{R}_t \mathbf{x}_{t-1} / q_t, \\ \mathbf{R}_t &= \mathbf{C}_{t-1} + \mathbf{W}_t, \\ \mathbf{W}_t &= \mathbf{C}_{t-1} (1 - \beta) / \beta. \end{aligned} \quad (12)$$

Two comments on discount specifications are in order. First, the specific form of the key variance matrix  $\mathbf{W}_t$  here, and in equation (9), is based on the standard discount factor method using the specified discount factor  $\beta$ . This induces a trivial computation of  $\mathbf{R}_t$  as the last two equations above lead to  $\mathbf{R}_t = \mathbf{C}_{t-1} / \beta$  (see West and Harrison, section 6.3). Second, the variance discount factor  $\delta$  is explicitly involved in the update equations for  $k_t$  and  $d_t$ . In the underlying multiplicative evolution equation (10), the innovation  $\eta_t$  has the Beta distribution  $Be(a_t, b_t)$  where  $a_t = \delta k_{t-1} / 2$  and  $b_t = (1 - \delta) k_{t-1} / 2$  (see West and Harrison 1997, section 10.8).

The analysis requires an initial, conjugate normal/inverse gamma prior  $p(\boldsymbol{\phi}_1, \sigma_1^2 | D_0)$ , which may be taken as a reference prior or a diffuse but informed proper prior. An attractive strategy is available in applications where, as in our EEG context, we have very long time series and the initial period of observations is relatively uninteresting (prior to seizure onset). We can then simply fit a standard constant autoregression to a short initial section of the data (of length of, say, three times the maximum order  $p$  contemplated), producing the standard reference normal/inverse gamma distribution. After inflating the scale parameters of this distribution (by a factor of 10, for example) we then obtain an appropriately located but still relatively very diffuse initial prior for  $(\boldsymbol{\phi}_1, \sigma_1^2)$  for analysis of the series with the initial section discarded.

### Treatment of Hyperparameters

The sequential analysis above provides calculation of the resulting predictive density of the data  $p(x_1, \dots, x_n | D_0) = \prod_{t=1}^n p_t$  where  $p_t = p(x_t | D_{t-1})$  is the value of the observed one-step ahead predictive density at time  $t$ , the ordinate at  $x_t$  of the univariate T density  $T_{\delta k_{t-1}}(f_t, q_t)$ . As the analysis is run at a specified set of values of the three key hyperparameters  $(p, \beta, \delta)$ , the numerical value of  $p(x_1, \dots, x_n | D_0)$  so produced is proportional to the likelihood function for these hyperparameters at that triple, this likelihood being marginalized with respect to the parameters  $\Phi_n$  and  $\Sigma_n$ . Hence, we may run the sequential analysis above repeatedly across a grid of values of (some of) the hyperparameters, to deliver evaluations of the likelihood function across “model space.” A prior for  $(p, \beta, \delta)$  modifies this to a marginal posterior. Treating this triple as a tuning parameter, we may then identify a posterior mode or other point value to condition on for inferences about  $\Phi_n$  and  $\Sigma_n$ . Under a uniform hyperprior, this will select maximum (marginal) likelihood values. Applied work focuses on high values of the two discount factors, and particularly of the variance discount factor  $\delta$ , as very high values already permit marked changes in parameters with rapidly sampled data such as we have in the EEG context. This influences choice of ranges of values of the hyperparameters over which to run the model and compute the likelihood function. Once a triple is selected, it defines a particular TVAR model and the results of sequential analysis at that model provide the basis for retrospective filtering to compute full posterior distributions for the TVAR parameters based on the fixed hyperparameter values, as follows.

### Retrospective Filtering

Following sequential processing of the data, retrospective filtering/smoothing computations produce the margins of the full posterior  $p(\Phi_n, \Sigma_n | D_n)$ . The key margins of interest are, for all  $t \leq n$ , given by  $(\boldsymbol{\phi}_t | D_n) \sim T_{k_{t,n}}(\mathbf{m}_{t,n}, \mathbf{C}_{t,n})$  and  $(\sigma_t^{-2} | D_n) \sim Ga(k_{t,n}/2, d_{t,n}/2)$ . The recursive equations for computing the defining parameters are summarized here. Some of these are naturally written in terms of the point estimates  $s_{t,n} = d_{t,n} / k_{t,n}$

of the variances  $\sigma_t^2$ . Initializing at time  $t = n$  with  $\mathbf{m}_{n,n} = \mathbf{m}_n$ ,  $\mathbf{C}_{n,n} = \mathbf{C}_n$ ,  $k_{n,n} = k_n$ ,  $d_{n,n} = d_n$  and  $s_{n,n} = s_n$ , we proceed sequentially backwards in time over  $t = n - 1, \dots, 1$  to compute:

$$\begin{aligned} \mathbf{m}_{t,n} &= (1 - \beta)\mathbf{m}_t + \beta\mathbf{m}_{t+1,n}, \\ \mathbf{C}_{t,n} &= [(1 - \beta)\mathbf{C}_t + \beta^2\mathbf{C}_{t+1,n}](s_{t,n}/s_t), \\ k_{t,n} &= (1 - \delta)k_t + \delta k_{t+1,n}, \\ 1/s_{t,n} &= (1 - \delta)/s_t + \delta/s_{t+1,n}, \end{aligned} \quad (13)$$

and with  $d_{t,n} = k_{t,n}s_{t,n}$ , for  $t = n - 1, \dots, 1$ .