Exploratory Modelling of Multiple Non-Stationary Time Series: Latent Process Structure and Decompositions

Raquel Prado and Mike West*

Duke University
United States

ABSTRACT We describe and illustrate Bayesian approaches to modelling and analysis of multiple non-stationary time series. This begins with univariate models for collections of related time series assumedly driven by underlying but unobservable processes, referred to as dynamic latent factor processes. We focus on models in which the factor processes, and hence the observed time series, are modelled by time-varying autoregressions capable of flexibly representing ranges of observed non-stationary characteristics. We highlight concepts and new methods of time series decomposition to infer characteristics of latent components in time series, and relate univariate decomposition analyses to underlying multivariate dynamic factor structure. Our motivating application is in analysis of multiple EEG traces from an ongoing EEG study at Duke. In this study, individuals undergoing ECT therapy generate multiple EEG traces at various scalp locations, and physiological interest lies in identifying dependencies and dissimilarities across series. In addition to the multivariate and non-stationary aspects of the series, this area provides illustration of the new results about decomposition of time series into latent, physically interpretable components; this is illustrated in data analysis of one EEG data set. The paper also discusses current and future research directions.

Key words and phrases: Dynamic latent factor model, Dynamic linear model, Non-stationary time series, Time series decomposition

---

*This research was supported in part by the National Science Foundation under grant DMS-9311071. The EEG data and context arose from discussions with Dr Andrew Krystal, of Duke University Medical Center, with whom continued interactions have been most valuable. Address for correspondence: Institute of Statistics and Decision Sciences, Duke University, Durham, NC 27708–0251 U.S.A. (http://www.stat.duke.edu)

1 Introduction

A wide variety of scientific and socio-economic problems involve study of multiple time series driven by underlying, latent processes characterising the system. In such problems, interest lies in identifying the time evolution and structure of the underlying system, as well as in characterising the univariate series and how they respond to variations in the underlying system. The notion of "exchangeable time series" (e.g., West and Harrison 1997, Section 16.4) has proven useful in some such problems, as have dynamic hierarchical models (e.g., Gamerman and Migon 1993). In other areas, a more direct approach to modelling and inference about the underlying but latent system process is of interest, leading to the notion of *dynamic factor models* that are the focus of this paper. We describe how such models are structured, and elaborate on exploratory univariate time series analyses that provide insight and inferences on underlying latent processes and cross-sectional structure. Our development involves specific, non-stationary time series models for latent processes, based on time-varying, autoregressive component dynamic models. Practical issues of model fitting and computation are briefly noted and illustrated.

One motivating problem arises from a collaboration with Duke psychiatrists studying issues of clinical design and efficacy of various brain seizure treatments, and also concerned with questions of brain function in various neurological conditions. Electroconvulsive therapy (ECT) is a major tool in brain seizure treatment and in fundamental brain research (Weiner and Krystal 1993), and EEG monitoring is the primary method of observation on brain activity during ECT (as in other contexts). ECT induced seizures are monitored by scalp electrodes measuring resulting EEG waveforms at various scalp locations, simultaneously, throughout the seizure episode. Very long EEG time series, of the order of several to a few tens of thousands of recordings, are available for individuals under varying treatment conditions, each with multiple series recorded. The data analysis summarised here comes from one experimentally induced seizure on a single individual under one treatment. We have 18 parallel series from 18 separate EEG "channels" – measured via electrodes at 18 scalp locations, and each generating a record of over 26000 voltage levels at a sampling rate of 256 observations per second. Some data from channel 4 appears in Figure 1 (see also the top time series plot in Figure 2); the general features are typical of the 18 channels, though levels and amplitudes of fluctuations over time do appear quite variable across channel. Of the 26000 plus observations on channel 4, we removed about 2000 at the start and end of the recording interval, and then subsampled the remainder every 6 observations to produce a series of 3600 observations; four sections of 500 consecutive observations from the start, end and two mid-sections of this series appear in Figure 1. Essentially indistinguishable displays are generated from the original data. The EEG time series exhibit a major pattern of quasi-cyclicity with

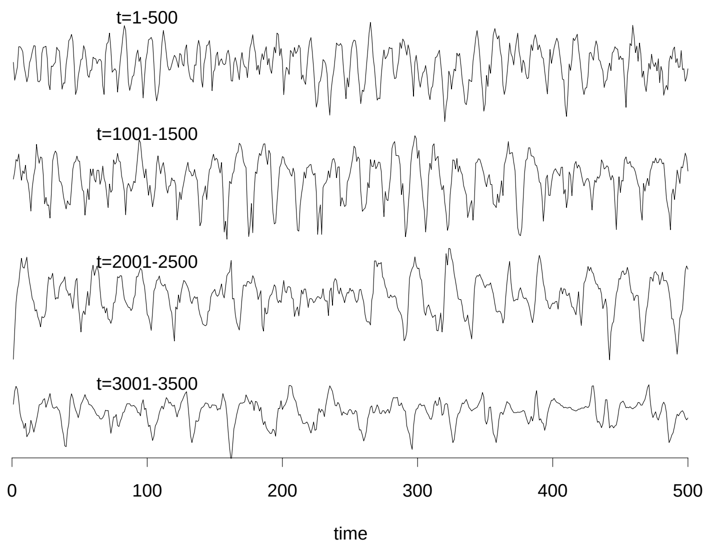

FIGURE 1. Sections of 500 consecutive observations from the recordings of EEG channel 4.

time-varying frequency characteristics, with superimposed high-frequency distortions. The appearance is quite typical of the "slow-wave" activity of epileptiform discharges, with increased amplitudes and wavelengths relative to the common alpha waves, and others, apparent in "normal" EEG traces (Dyro 1989). The frequency of the apparently dominant waveform decays through the course of the seizure, with concurrent time-variation in amplitude and decay characteristics; this represents the initial build up to peak intensity of the seizure, with correspondingly rapid oscillations in EEG levels, followed by gradual dissipation and eventual decline of the fluctuations as the seizure dies out. Modelling these patterns will bear in mind the objective of characterising the seizure through parameters that measure and reflect the level of maturity of the seizure, capturing and quantifying its rate of onset and eventual decay; inferring the beginning, duration and end of actual seizure activity is of critical interest and has proven to be a challenging problem (Weiner and Krystal 1993). Comparing such parameters across channels is of interest in exposing possible differences due to scalp placement of electrodes, and in inferring structure in the driving seizure signal through commonalities across channels. Comparison of inferences about such characterising parameters from EEG records on repeated seizures under differing treatment conditions, and across different individuals, is of further interest, though this is beyond our scope here.

It is apparent from Figure 1 that there is rough stability of the EEG waveform in short periods of perhaps one or two hundred observations, on this time scale. This holds up across channels. Exploratory analysis suggests that, within such short sections, the data appear consistent with a roughly constant AR model of order 10-15, but that the parameters differ as time evolves through the seizure episode. This indicates the global applicability of a time-varying AR model (West and Harrison 1997, Section 9.5). We develop such models below, and explore issues of time series decomposition and multivariate, dynamic factor model structure for the collection of EEG channel series.

2 Non-stationary process models

2.1 Factor models and time-varying autoregressions

Suppose a system under study is characterised by a latent process x_t that drives m parallel time series via

$$y_{i,t} = \beta_i x_t + \nu_{i,t} \tag{1}$$

for $i = 1, \dots, m$ and all t. The β_i are regression parameters, or factor weights, and the $\nu_{i,t}$ assumedly independent noise sequences, here taken as zero-mean normal, $N(\nu_{i,t}|0, s_i)$ for some variances s_i.

Suppose further that the latent x_t series is a time-varying parameter autoregression (TV-AR, as in West and Harrison 1997, Section 9.5). Such models are of particular interest for long series to adapt to time-varying patterns of dependence and non-stationarities; they embody the notion of "local stationarity" but "global" non-stationarity. Specifically, assume that

$$x_t = \sum_{h=1}^{p} \phi_{h,t} x_{t-h} + \epsilon_t \quad \text{and} \quad \phi_{h,t} = \phi_{h,t-1} + \omega_{h,t}, \tag{2}$$

where $\boldsymbol{\phi}_t = (\phi_{1,t}, \dots, \phi_{p,t})'$ is the time-varying AR parameter vector at time t and $\boldsymbol{\omega}_t = (\omega_{1,t}, \dots, \omega_{p,t})'$ is the stochastic change in the parameter vector at time t. Further assume that the error terms are independent normal, with distributions $N(\epsilon_t|0, v_t)$ and $N(\boldsymbol{\omega}_t|0, \mathbf{W}_t)$ for some sequence of AR innovation variances v_t, and variance-covariance matrices $\mathbf{W}_t$ controlling variation in $\boldsymbol{\phi}_t$. We refer to (2) as a TV-AR(p) model. This can be written as a dynamic linear model (DLM) in various ways (West and Harrison 1997, Sections 9.4 and 9.5). Variants on this model might include other forms of time-variation for $\boldsymbol{\phi}_t$, beyond this basic random walk, non-normal innovations and evolution errors for outliers and more abrupt changes in structure, and additional DLM components for time-varying trends (West 1995, 1996, 1997a,b,c).

Introduce the "instantaneous" characteristic AR polynomial $\phi_t(u) = 1 - \phi_{1,t}u - \cdots - \phi_{p,t}u^p$ for each t, and the standard back-shift operator B. Then $\phi_t(B)x_t = \epsilon_t$ and (1) implies

$$\phi_t(B)y_{i,t} = \beta_i\epsilon_t + \phi_t(B)\nu_{i,t}. \tag{3}$$

This is a time-varying parameter ARMA(p,p) model, or TV-ARMA(p,p). The TV-AR filter $\phi_t(\cdot)$ is common across series i, just the underlying structure of the latent x_t process. Hence fitting individual TV-ARMA models, or similar, to the $y_{i,t}$ should yield inferences about the AR component that are similar across series, and which may be used to identify structure in the underlying x_t process. This is a key to exploratory modelling to isolate factor structure. A practical strategy is to fit high-order TV-AR models to each of the $y_{i,t}$ series and use these to isolate the dominant TV-AR components, attributing the residual high-frequency components to the MA noise structure. This is a valuable and easily implemented strategy whose success and utility is described below in the EEG context. To be specific, suppose that $p = 4$ and we fit, say, a TV-AR(12) model to $y_{i,t}$, with AR polynomial $\phi_{i,t}(B)$ of order $p = 12$. Then we expect that $\phi_{i,t}(B) \approx \phi_t(B)\theta_{i,t}(B)$ where $\phi_t(\cdot)$ is close to the TV-AR(4) polynomial of the x_t process, and $\theta_{i,t}(B)$ represents a TV-AR(8) component that, when inverted, approximates the TV-MA(4) structure of $\beta_i\epsilon_t + \phi_t(B)\nu_{i,t}$. Our practical experience with a variety of data sets verifies the utility of this approach.

2.2 Decompositions of factor and observation processes

A most useful time series decomposition result is relevant to inference about latent structure and the partitioning of the series x_t into key components. We use this in data analysis below, though space precludes presentation of full details here. The results used generalise West (1997c) to the time-varying case and will be fully reported elsewhere; see also West and Harrison (1997, Section 9.5).

Note that $\phi_t(u) = \prod_{j=1}^{p}(1 - \alpha_{j,t}u)$ where the $\alpha_{j,t}$ are (real or complex) reciprocals of the characteristic roots at time t. Suppose, as is usual, that the roots are distinct, occurring as c pairs of complex conjugates and $r = p - 2c$ real and distinct values. Write the complex roots as $r_{j,t}\exp(\pm i\omega_{j,t})$ for $j = 1, \ldots, c$, noting that the real-valued, non-zero elements $\omega_{j,t}$ correspond to the "instantaneous" frequencies of quasi-cyclical component behaviour in the series at time t, varying over time. Write the real roots as $r_{j,t}$ for $j = 2c + 1, \ldots, p$. Based on the development of similar models in DLM theory (West and Harrison 1997, Section 5.4.4), our decomposition result is that x_t may be written as the sum of $c + r$ real processes

$$x_t = \sum_{j=1}^{c} z_{j,t} + \sum_{j=1}^{r} a_{j,t} \tag{4}$$

in which the summands $z_{j,t}$ correspond to the complex roots, and the $a_{j,t}$ to the real roots. The real-valued processes $a_{j,t}$ follow individual, time-varying AR(1) models $a_{j,t} = r_{j,t}a_{j,t-1} + \eta_{j,t}$ for some zero-mean innovations $\eta_{j,t}$, (correlated across component index j). The real-valued $z_{j,t}$ processes follow quasi-periodic, time-varying ARMA$(2,1)$ models, $z_{j,t} = 2r_{j,t}\cos(\omega_{j,t})z_{j,t-1} - r_{j,t}^2 z_{j,t-2} + \gamma_{j,t}$ for additional, zero-mean innovations $\gamma_{j,t}$ that are mutually correlated and also correlated with the $\eta_{j,t}$. Based on a specified $\boldsymbol{\phi}_t$ vector and estimated or known values of x_t, extensions of the eigenanalysis results in West (1997c) apply to allow calculation of the defining quantities $r_{j,t}$ and $\omega_{j,t}$, and of the actual values of the latent component processes $a_{j,t}$ and $z_{j,t}$. The time-varying variances of the innovation sequences $\eta_{j,t}$ and $\gamma_{j,t}$ are similarly computed; these are relevant in assessing relative amplitudes of the latent components.

Under the factor model (1) it follows that the latent processes $a_{j,t}$ and $z_{j,t}$ drive the observed series $y_{i,t}$, so that fitting higher-order TV-AR models to the individual $y_{i,t}$ series should yield estimated latent processes that are similar across series (up to multiplicative constant factors β_i). This constructive result is illustrated in decomposition of EEG signals below. We note, in passing, the connections with the AR component models of West (1995, 1996).

2.3 Multivariate dynamic factor models

More general dynamic factor models extend (1) to include $k > 1$ latent factor processes and possibly time-varying factor weights. A dynamic, $k-$factor model has the following ingredients: label the $k-$factor processes $x_{j,t}$ for $j = 1, \dots, k$; let $\beta_{i,j,t}$ be the (possibly time-varying) regression parameter relating $y_{i,t}$ to $x_{j,t}$; write $\mathbf{y}_t = (y_{1,t}, \dots, y_{m,t})'$ and $\boldsymbol{\beta}_{j,t} = (\beta_{1,j,t}, \dots, \beta_{m,j,t})'$ for each j and t. Then the multivariate time series $\{\mathbf{y}_t\}$ is modelled as

$$\mathbf{y}_t = \sum_{j=1}^{k} \boldsymbol{\beta}_{j,t} x_{j,t} + \boldsymbol{\nu}_t \qquad (5)$$

over $t = 1, \dots, n$, where $\boldsymbol{\nu}_t \sim N(\boldsymbol{\nu}_t|\mathbf{0}, \mathbf{V})$ are independent noise terms. It will usually be appropriate to assume $\mathbf{V} = \text{diag}(V_1, \dots, V_n)$ so that the instantaneous dependencies among the $y_{i,t}$ are due entirely to the latent factor processes $x_{j,t}$. Defining the $m \times k$ matrices $\mathbf{B}_t = [\boldsymbol{\beta}_{1,t}, \cdots, \boldsymbol{\beta}_{k,t}]$ for each t, (5) can be written as

$$\mathbf{y}_t = \mathbf{B}_t \mathbf{x}_t + \boldsymbol{\nu}_t \qquad (6)$$

where $\mathbf{x}_t = (x_{1,t}, \dots, x_{k,t})'$ is the latent factor vector at time t.

A rather general framework models the collection of latent processes $\mathbf{x}_t$ via a DLM. Then (6) results in a highly structured, multivariate DLM for $\mathbf{y}_t$, a fact that has modelling and technical ramifications. One important class of models is that based on lagged latent factors. Suppose that the first factor $x_{1,t} = x_t$ is a TV-AR process as above, and that additional factors are simply lagged values, namely $x_{2,t} = x_{t-1}, \ldots, x_{k,t} = x_{t-k+1}$. In this case, it easily follows that, for each i,

$$\phi_t(B) y_{i,t} = \sum_{j=1}^{k} \beta_{i,j} \epsilon_{t-j+1} + \phi_t(B) \nu_{i,t}, \tag{7}$$

so $y_{i,t}$ is TV-ARMA(p, q) with $q = \max(p, k)$. The comments about fitting higher-order TV-AR models applies here as above.

Typically, applied contexts involve far fewer factor processes than observation series, i.e., k much smaller than m. In the EEG context, for example, we surmise that a small number of underlying factor processes should explain the observed variability across the $m = 18$ channels; two primary, latent though real processes are the invoked seizure waveform and the underlying normal brain activity, respectively. We may need $k > 2$ factor processes if, for example, lagged values of these two major processes are evident in the multi-channel series.

Note that, in the case of constant parameter AR models for the latent $x_{j,t}$ processes, and with constant factor weights $\beta_{i,j,t} = \beta_{i,j}$ for all t, we are in a context close to that of the foundational theory in Peña and Box (1987), which provides useful background material.

3 Some analyses of the EEG data

Following earlier comments, we have undertaken various exploratory analyses of the 18 EEG channels individually using TV-AR(12) models. Inference is based on couching the model in DLM form and then applying standard results, as in West and Harrison (1997, Section 9.5) for example. In particular, we model time-variation in the AR parameters using discount factors, and select model order and discount factors to maximise marginal likelihood functions. Then posterior distributions for the individual TV-AR parameter vector for each series at each time point are computed using standard updating and filtering recursions (West and Harrison 1997, chapter 4); results summarised below are based on posterior means of all TV-AR coefficients. Across the 18 channels, $p = 12$ and a common discount factor of 0.994 are chosen and used in the analyses summarised here. Full details will be reported elsewhere; we note that fitting higher-order models does not materially impact the broad conclusions of interest here.

Across all 18 series, we find that posterior distributions for the TV-AR(12) parameters support moderate degrees of variation through time as the seizure begins, matures and eventually decays. Across series $i = 1, \ldots, 18$ and all time t, TV-AR polynomials exhibit and maintain two key pairs of complex conjugate roots, with moduli and arguments that vary slowly throughout the seizure. The frequencies of these two roots, though time-varying, correspond closely to the known and expected ranges of the seizure slow-wave and the basic alpha rhythm of background brain activity; as a result we identify these four roots as common. In each series at each time point, we now factor the TV-AR polynomial $\phi_{i,t}(B)$ as the product of a TV-AR(4) polynomial with these four roots, and a resulting TV-AR(8) component $\theta_{i,t}(B)$. Inverting each $\theta_{i,t}(B)$ provides an inferred, infinite-order MA representation, and we evaluate the estimated MA coefficients as functions of those of $\theta_{i,t}(B)$. This confirms that, across all channels i and over all time t, only the first three or four coefficients are non-negligible, providing support for the representation (3) and hence for the single factor structure (1) with an underlying TV-AR(4) process (2).

We illustrate the model for a single channel, number $i = 4$. Based on posterior mean estimates of the TV-AR(12) coefficients for this series, we apply the decomposition result of Section 2.2 to isolate estimated latent components. The two dominant pairs of complex-conjugate roots correspond to two quasi-cyclical TV-ARMA$(2, 1)$ components and it happens that these two components dominate the rest in amplitude. We note a basic identification problem: any AR parameter vector generates characteristic roots, and resulting series components, that have no inherent ordering. We surmount this by summarising inferences based on ordering components by estimated amplitude. Various interesting technical and conceptual issues arise due to this inherent identifiability; for example, as amplitude, wavelength and moduli characteristics are time-varying, distinct components may switch positions under any specific ordering. Such issues will be thoroughly explored in a forthcoming article.

Figure 2 displays the first four estimated components of the channel 4 series. Figure 3 is a similar display over a short subseries in the central section of the seizure, more clearly displaying structure in the data and components. The remaining roots contribute components of lower amplitudes than those displayed. Evidently, the first component is dominant, representing the periodic forcing of aggregate transmission of cortical neural networks that drive the slow-wave EEG signal. The second component apparently contributes to the signal, though with a much lower amplitude than the primary wave. Analyses of each of the 18 EEG channels lead to very similar decompositions; channel 4 here is quite typical.

Inferences about the "instantaneous" moduli and frequencies of each of the two dominant components displayed follow directly from inferences on the TV-AR parameter vector. At the posterior mean of the AR vector, at each time point, we obtain estimates of moduli and frequency directly by computing roots of the characteristic polynomial. This is done for each of the 18 series and illustrated in Figures 4 and 5. Figure 4 displays estimates of instantaneous frequency of the dominant seizure waveform over time throughout the seizure episode and for each of the 18 channels.

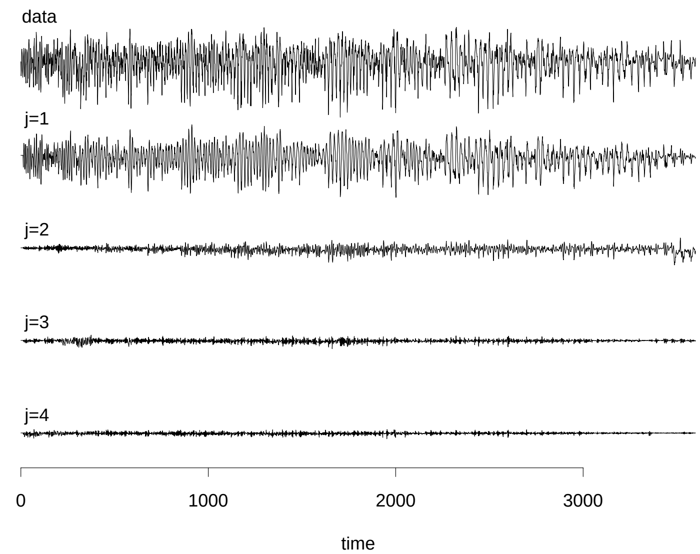

FIGURE 2. First four components of the decomposition of EEG channel 4 signal, evaluated based on posterior mean estimates of the time-varying AR parameter vector.

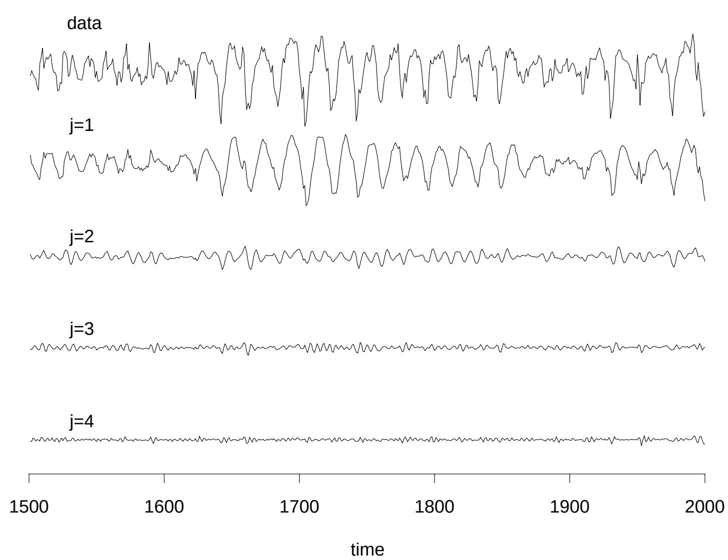

FIGURE 3. First four components of the decomposition of EEG channel 4 signal, as in Figure 2 but restricted to a central section of 500 time points.

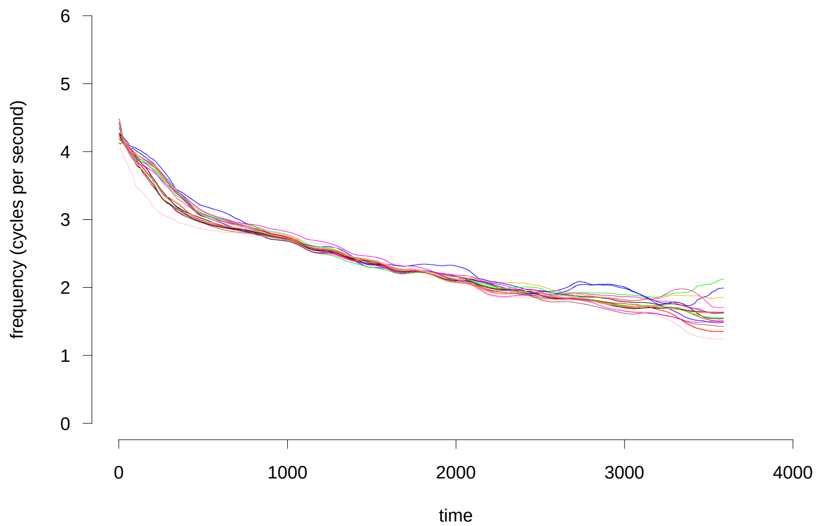

FIGURE 4. Estimated time trajectories of frequencies of the dominant quasi-cyclical components of each of the 18 EEG channels.

We note consistency of the range of frequencies evident here with known and expected ranges of 1-5 cycles per second, gradually decaying as the seizure dies out (Dyro 1989). Figure 5 provides a similar display for estimated moduli of the seizure waveform as a function of time, and for all channels. Similarly, Figure 6 graphs estimated trajectories of amplitudes of the dominant component across channels. Similar figures may be constructed for the second important component, that of the underlying alpha rhythm. The frequency trajectories lie in the 4-9 cycles per second range, consistent with expected ranges of normal brain activity.

The figures indicate that the pattern of time evolution of the dominant seizure waveform is consistent across univariate analyses of the 18 channel series: we see consistency in the decreasing frequency content, in the stable but eventually decaying modulus, and in the form of amplitude trajectories that map the onset, process of maturing and eventual decay of the seizure. Consistency of estimated frequency and amplitude trajectories, and similar consistency of estimated dominant components (not shown), support the notion of an underlying factor model (1) in which the seizure waveform is at least a component of the process x_t. There is similar consistency across channels for the second component corresponding to the background brain activity. Hence we conclude that figures (4-6) give an overall, exploratory description of the nature of the seizure episode, and that an underlying,

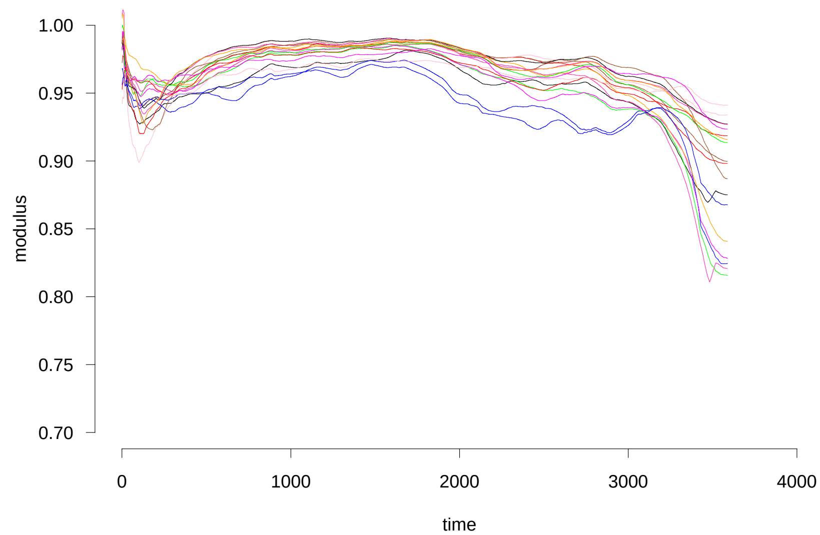

FIGURE 5. Estimated time trajectories of moduli of the dominant quasi-cyclical components of each of the 18 EEG channels.

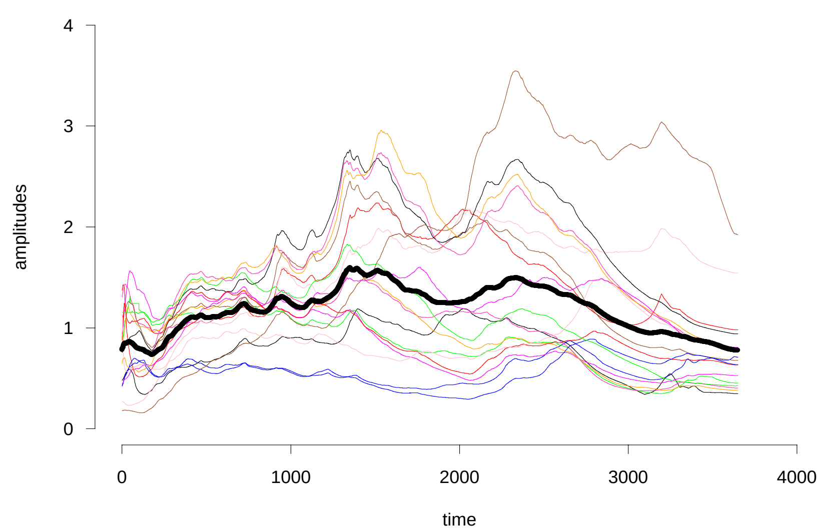

FIGURE 6. Estimated time trajectories of amplitudes of the dominant quasi-cyclical components of each of the 18 EEG channels. The full line through the trajectories represents the average.

latent seizure process contains a seizure signal with frequency, modulus and amplitude characteristics varying as some average of the three sets of trajectories graphed.

The simplest "candidate" factor model, therefore, is a single factor model as in (1) in which the latent driving process x_t is TV-AR(4) with complex root structure generating the decomposition into two quasi-periodic TV-ARMA$(2,1)$ subseries; these two subseries represent the seizure slow-wave and background alpha rhythm, as identified and described above. Were this appropriate then, as discussed earlier, the TV-AR(12) models for each of the channels would approximately decompose into these two subseries plus low amplitude, higher frequency components; we have seen that this is the case. However, the data exhibit more complex structure. Under this candidate model, (1) implies a simple linear regression of each of the $y_{i,t}$ on the average series $\bar{y}_t = \sum_{i=1}^{18} y_{i,t}/18$, and with each regression having zero-mean and independent errors. This can be directly assessed, and is found to be wanting; residual time series from such simple linear models evidence residual quasi-periodicities reflecting both seizure and alpha rhythms. This suggests a more elaborate, multi-factor structure in which two separate latent processes drive the data, i.e., a model of the form (5) in which $x_{1,t}$ is a TV-AR(2) "seizure" process and $x_{2,t}$ is a TV-AR(2) "alpha rhythm" process. Even with constant factor weights $\beta_{i,j,t} = \beta_{i,j}$ for all t, each channel i and factors $j = 1, 2$, this model is quite substantially more complex; accepting that time-varying factor weights may be of relevance too adds further complications. More formal approaches to analysis of these more complex but, as illustrated here, practically very relevant multivariate models is under current investigation. This development will involve generalisation of the framework, and resulting Bayesian simulation methods of analysis, in West (1996, 1997c). Nevertheless, the style of exploratory modelling and analysis illustrated here has been of key importance in elucidating the structure of the EEG series, as it is of similar utility in other application areas involving multiple, non-stationary time series.

4 References

[1] Dyro, F.M. (1989) *The EEG Handbook.* Little, Brown and Co., Boston.

[2] Gamerman, D., and Migon H.S. (1993) Dynamic hierarchical models. *Journal of the Royal Statistical Society, Series B,* 55, 629-642.

[3] Krystal, A.D. and Weiner, R.D. (1994) ECT seizure therapeutic adequacy. *Convulsive Therapy*, 10, 153-164.

[4] 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 place-

ment on the ictal electroencephalogram: An intraindividual crossover study. *Biological Psychiatry*, 34, 759-767.

[5] Peña, D., and Box, G.E.P. (1987) Identifying a simplifying structure in time series. *Journal of the American Statistical Association,* **82**, 836-843.

[6] Pole, A., West, M., and Harrison, P.J. (1994) *Applied Bayesian Forecasting and Time Series Analysis.* Chapman-Hall, New York.

[7] Weiner, R.D., and Krystal, A.D. (1993) EEG monitoring of ECT seizures. In *The Clinical Science of Electroconvulsive Therapy,* American Psychiatric Press, Inc., 93-109.

[8] West, M. (1995) Bayesian inference in cyclical component dynamic linear models. *Journal of the American Statistical Association,* **90**, 1301-1312.

[9] West, M. (1996) Some statistical issues in Palæoclimatology (with discussion). In *Bayesian Statistics 5*, (eds: J.O. Berger, J.M. Bernardo, A.P. Dawid and A.F.M. Smith), Oxford University Press.

[10] West, M. (1997a) 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), Kluwer.

[11] 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), IMS Monographs.

[12] West, M. (1997c) Time series decomposition. Under review at *Biometrika.*

[13] West, M., and Harrison, P.J. (1997) *Bayesian Forecasting and Dynamic Models* (2nd Edition). Springer-Verlag, New York.