

General Strategy for Hierarchical Decomposition of Multivariate Time Series: Implications for Temporal Lobe Seizures

M. A. REPUCCI, N. D. SCHIFF, and J. D. VICTOR

Department of Neurology and Neuroscience, Weill Graduate School of Medical Sciences of Cornell University, 1300 York Avenue, New York, NY

(Received 19 October 2000; accepted 4 October 2001)

Abstract—We describe a novel method for the analysis of multivariate time series that exploits the dynamic relationships among the multiple signals. The approach resolves the multivariate time series into hierarchically dependent underlying sources, each driven by noise input and influencing subordinate sources in the hierarchy. Implementation of this hierarchical decomposition (HD) combines principal components analysis (PCA), autoregressive modeling, and a novel search strategy among orthogonal rotations. For model systems conforming to this hierarchical structure, HD accurately extracts the underlying sources, whereas PCA or independent components analysis does not. The interdependencies of cortical, subcortical, and brainstem networks suggest application of HD to multivariate measures of brain activity. We show first that HD indeed resolves temporal lobe ictal electrocorticographic data into nearly hierarchical form. A previous analysis of these data identified characteristic nonlinearities in the PCA-derived temporal components that resembled those seen in absence (petit mal) seizure electroencephalographic traces. However, the components containing these characteristic nonlinearities accounted for only a small fraction of the power. Analysis of these data with HD reveals furthermore that components containing characteristic nonlinearities, though small, can be at the origin of the hierarchy. This finding supports the link between temporal lobe and absence epilepsy. © 2001 Biomedical Engineering Society. [DOI: 10.1114/1.1424914]

Keywords—Nonlinear dynamics, Autoregression, Epilepsy, Principal components, Independent components, Spatiotemporal, Electroencephalography, Electrocorticography.

INTRODUCTION

Multivariate time series are abundant in biomedical systems, from spatiotemporal signals such as the electroencephalogram (EEG), to temporal patterns of gene expression observed with GeneChip²¹ technologies. One common approach for the analysis of multivariate data involves decomposition into several independent time series (e.g., by principal components analysis⁴), followed

by analysis of the individual time series (e.g., autoregressive modeling, spectral analysis, and nonlinear techniques²⁰). Typically, the goal of the analysis is to extract the underlying biological sources of the multivariate temporal data, in order to characterize each source more fully and to describe the system more precisely. Unfortunately, the criterion of independence of sources enforced by standard decomposition techniques may not be appropriate for real biological systems, in which the underlying sources may well have dynamic interrelationships. With this in mind, we developed a novel approach to multivariate time series analysis—hierarchical decomposition (HD)—that exploits the causal (i.e., forward in time) dynamic content of multivariate temporal data.

In contrast to HD, standard decomposition techniques including principal components analysis (PCA), alone or modified by varimax rotations (VR),⁴ independent components analysis (ICA),² and non-negative matrix factorization (NMF),¹⁰ do not exploit causality. Briefly, PCA seeks sources that are uncorrelated (orthogonal). ICA seeks sources that are independent in an information-theoretic sense. VR seeks a transformation that maximizes the sum of the variances in each extracted source. NMF attempts to find sources whose weights are non-negative (i.e., only additive, not subtractive, combinations of sources are allowed). These assumptions are appropriate for image analysis or for separation of independent time series (e.g., the “cocktail party problem,” separating one voice from the cacophony of other conversations and background noise), but may be unsuitable for multivariate systems originating from dynamically interrelated sources. In particular, these methods assume that the order of data points is irrelevant, and thus would produce equivalent results for the original, time-reversed, or randomly shuffled data.

As described below, the HD method takes advantage of the signal dynamics, and thus may achieve a more useful resolution of the sources of a multivariate time series. HD was designed with physiologically derived

Address correspondence to Michael A. Repucci, Department of Neurology and Neuroscience, Weill Graduate School of Medical Science of Cornell University, 1300 York Avenue, New York, NY 10021.

multivariate time series in mind. Such time series often result from dynamically interrelated generators, in which some of the generators have a driving role. For example, the EEG exhibits synchronizations, desynchronizations, and rhythmic oscillations resulting from cortical, subcortical, and brainstem network activity: these are characteristics anticipated in a dynamically interrelated system.¹⁸ To exploit these observations, we assume that the generators are organized hierarchically. Namely, we assume that there is an autonomous generator whose state depends only on noise input and feedback from itself (i.e., an autoregressive process). Each successive generator is influenced by an independent noise input, its own state, and the state of more dominant generators in the hierarchy, such that the last generator is influenced by noise input and the state of *all* generators. The HD method, as presented below, attempts to account for the temporal dynamics in a multivariate data set by resolving it into a model of this form.

We show that if a system does conform to this hierarchical structure, and the multivariate autoregressive model includes sufficient terms, then the HD resolution is unique and will recover the hierarchically related sources. Moreover, the HD approach will not be misleading for an independent system; if the underlying sources are truly independent, the HD algorithm will uncover them as well. On the other hand, there is no guarantee that all multivariate time series are reducible to hierarchical form—one such example is a system with reciprocally related generators. For such systems, failure of this approach to identify a hierarchical resolution is positive evidence for the presence of a more complex network structure.

In brief, the HD approach consists of building a multivariate linear autoregressive (MLAR) model⁶ from the PCA-derived components of the original data set. This is followed by seeking a coordinate transformation that rotates the MLAR model to make it as consistent as possible with a hierarchical interrelationship among components. We demonstrate that for simulated multivariate time series with truly hierarchical structure, HD accurately extracts the underlying generators of the system. We then use the HD method to analyze ictal electrocorticographic (ECoG) records, the application that inspired the HD approach. Analysis of these data reveals that significant hierarchical structure is present in all records; this fact is not obvious by PCA or ICA decomposition alone. Subsequent analysis of these HD resolved components adds further insight to the dynamic nature of an ictal discharge.

In previous work,¹⁵ we identified certain characteristic dynamic nonlinearities in EEG traces obtained during *absence* seizures (also known as *petit mal* seizures) and in the PCA-derived temporal components from temporal lobe ictal ECoG data. The similarities, most notably the

nonlinear interactions at lags around 90 and 150 ms, suggested a common underlying mechanism. These characteristic dynamics, however, were observed only in components that accounted for a small fraction of the variance of the original data (third, fourth, or fifth ranking components, when ordered by the amount of variance explained), thus leaving their biological significance uncertain. By contrast, the HD method demonstrates that these nonlinearities are present in the more autonomous components (i.e., the components in the system that drive the remaining components). This indicates that the size of the components does not necessarily correlate with their physiological importance, and thus strengthens the evidence that temporal lobe and absence seizures share a common mechanism. Here, the ability of HD to resolve the underlying generators on the basis of dynamical interrelationships, rather than by independence or power, sheds light on a prior puzzle.

METHODS

Multivariate ECoG data were previously collected as briefly described here (for further details, see Schiff *et al.*¹⁵). Epilepsy patients with medication-resistant temporal lobe seizures were studied with subdural electrode grids and strips in the course of presurgical evaluation for temporal lobectomy. ECoG signals were recorded by a 64-channel Telefactor Beehive (West Conshohocken, PA) telemetry system (sampling rate of 200 Hz, low-pass filter at 0.3 Hz and high-pass filter at 70 Hz), and concurrent video images of each patient were obtained. Thirteen to 16 channels of clear artifact-free recordings of ictal events (occurring during wakefulness) were resampled at 100 Hz (2:1); the mean was subtracted and detrending applied. Typically, 5 s of data recorded from within an ongoing ictal event were analyzed as described in the text.

All data analyses herein were carried out via MATLAB functions and scripts (version 5.3, release 11, Pentium II PC running Windows 2000), which are archived, along with a sample data set, at <http://www.med.cornell.edu/research/hds>. We consider a multivariate data set consisting of N evenly spaced observations in time, on each of M channels ($M \leq N$), represented as a matrix X ($M \times N$). In our application to ECoG data $N=512$ (approximately 5 s of data sampled at 100 Hz) and typically $M=14$ (epicortical recording sites).

Principal Components Analysis (PCA)

The original data matrix X is processed by PCA (see file “pca.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>), yielding a second matrix Y ($M \times N$), consisting of linear combinations of P orthogonal principal components,

$$Y = C^T W T. \tag{1}$$

C is a $P \times M$ matrix whose rows are the spatial weights, T is a $P \times N$ matrix whose rows are the temporal weights, and W is a $P \times P$ diagonal matrix of eigenvalues whose elements, squared, are the amount of variance explained by each principal component. For any number of components P [$P \leq \min(M, N)$], PCA guarantees that the unexplained variance between X and Y ,

$$R_{\text{PCA}} = \text{tr}[(X - Y)(X - Y)^T], \tag{2}$$

is minimized (i.e., Y is the “best fit” to X).

This decomposition of X into orthogonal (i.e., independent or instantaneously uncorrelated) components C and T , however, is not unique. Replacing the temporal components T by GT (for any nonsingular linear transformation G) and, accordingly, the spatial components C by $W^{-1}(G^{-1})^T W C$, leaves the matrix Y unchanged: $[(W^{-1}(G^{-1})^T W C)^T W [GT]] = C^T W G^{-1} W^{-1} W G T = C^T W T = Y$, where we have used basic properties of the transpose, and the fact that W is diagonal. The specific, but arbitrary, solution identified by PCA consists of matrices C and T whose rows are orthonormal, and a matrix W whose diagonal elements are non-negative and ordered from largest to smallest. Resolving the nonuniqueness of the decomposition obtained with Eq. (1) is the focus of the remainder of the HD analysis (i.e., finding a unique, nonarbitrary transformation G).

We use PCA as a first step in the analysis in order to reduce the dimension of the system by removing instrumental noise from the signal (without removing biological noise¹²). For real data, the P temporal components are selected such that each component accounts for a fraction of variance greater than or equal to M^{-1} , except where noted. Typically this allows us to reduce a 13- to 16-channel system to 3–5 components that together account for greater than 80% of the variance. Selecting a small number of components reduces the computational burden for the following analyses.

Multivariate Linear Autoregression (MLAR)

We next create a MLAR model of the temporal components T in a form equivalent to that proposed by Gersh and Yonemoto (see file “mlar.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>).⁶ The P temporal components (i.e., each element $T_{p,n}$ — n th sample of the p th component) are modeled by an L -term autoregression:

$$T_{p,n} = R_{p,n} + \mu_p + \sum_{q=1}^P \sum_{l=1}^L A_{q,p,l} T_{q,n-l}, \tag{3}$$

where R is a $P \times N$ matrix of residual values, μ_p is the mean value of the p th component, and A is a $P \times P \times L$ three-dimensional array of model coefficients. It is helpful to view the model coefficients A as a set of L planes, each of which is a $P \times P$ matrix. The l th plane of A describes the estimated influence of the signals at a time $n-l$, on the signal values at time n . Namely, $A_{q,p,l}$ is the influence of the q th component at the l th time lag on the current value of the p th component.

The model coefficients A are determined by minimizing the sum of squared residual values (the innovations, $R_{p,n}$) in the model

$$R_{\text{MLAR}} = \sum_{p=1}^P \sum_{n=1}^N R_{p,n}^2, \tag{4}$$

via the Yule–Walker equations.²² We choose the maximum lag L (the order of the MLAR model) by the Akaike criterion (AIC),¹ a statistical justification based on whether the amount of reduction in residual variance is sufficient to justify the inclusion of additional linear model terms. In the examples below this typically yields a value of L from 2 to 4 lags.

Decorrelation of the MLAR Innovations

The first step in resolving the nonuniqueness of the PCA decomposition relies upon the $P \times N$ matrix of innovations $R_{p,n}$. These quantities are viewed as random terms that drive the P channels of the MLAR model. Thus, if T represents distinct underlying generators, then the corresponding driving terms in Eq. (3) will be uncorrelated. For this reason, we require that the transformed temporal components GT decorrelate $R_{p,n}$. Additionally, it is analytically convenient to assume that each channel of innovations has equal variance (i.e., no particular source is noisier than another). This amounts to an arbitrary choice for the overall size of each generator, but makes no further assumptions about the dynamical structure of the system.

Under these assumptions, we seek transformations K that orthonormalize the innovations $R_{p,n}$ —that is, $(KR)(KR)^T = KRR^T K^T = I$, where I is the identity matrix. One such matrix K can be obtained by dividing the rows of the eigenvectors of RR^T by the square root of the corresponding eigenvalues. (The rows of the matrix K are necessarily orthogonal, since they are scalar multiples of the eigenvectors of the positive symmetric matrix RR^T .) The transformation K applied to Eq. (3) yields new temporal components $T' = KT$, new autoregressive coefficients $A'_l{}^T = K A_l^T K^{-1}$ (for $l = 1 \cdots L$, where A_l^T denotes the transpose of A_l), and decorrelated, normalized innovations $R' = KR$. Nevertheless, this does not fully resolve the nonuniqueness problem, since premultiplying K

by any orthogonal transformation Q (i.e., Q such that $QQ^T = Q^TQ = I$) yields an alternative matrix $G = QK$, for which GR is likewise decorrelated and normalized: $(GR)(GR)^T = (QKR)(QKR)^T = QKRR^TK^TQ^T = QIQ^T = I$.

Novel Search Among Orthogonal Rotations

At this point, we have reduced the nonuniqueness of the PCA decomposition to an arbitrary scale factor for each channel—the normalization performed above will suffice—and any orthogonal transformation Q . We now search for a rotation of the temporal components QT' that is consistent with the hierarchical structure proposed in the introduction. Recall that the autoregressive coefficients A' from the decorrelated MLAR model specify the dynamic interrelationships among the P components at up to L lags. Accordingly, the HD algorithm seeks a $P \times P$ transformation Q that simultaneously transforms all L matrices A'_l into upper-triangular form (see file “rotation.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>). We require that Q be orthogonal, thereby guaranteeing that the residuals R_{MLAR} are unchanged and that the orthonormality of the innovations R' is preserved. The search is further restricted to the set of rotations [i.e., $\det(Q) = 1$] to avoid arbitrary changes in the sign of an odd number of the components.

Q is obtained iteratively via a procedure motivated by Jacobi matrix diagonalization.^{11,13} We begin with an arbitrary rotation Q_0 (see below), and successively apply a sequence of individual plane rotations J , each about a pair of axes $[u, v]$. The sequence of plane rotations consists of multiple cycles through each of the possible choices of axis pairs. That is,

$$Q = J_{u_d, v_d}(\theta_{fd}) \cdots J_{u_1, v_1}(\theta_{d+1}) J_{u_d, v_d}(\theta_d) \cdots J_{u_1, v_1}(\theta_1) Q_0, \tag{5}$$

where $d = P(P - 1)/2$ is the number of unique axis pairs (i.e., for a 3×3 matrix $d = 3$, for a 4×4 matrix $d = 6$, and so on), f is the number of cycles through all axis pairs, and each $J_{u,v}(\theta)$ is a rotation by an angle θ about axes $[u, v]$. For example,

$$J_{4,2}(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & 0 & \sin(\theta) \\ 0 & 0 & 1 & 0 \\ 0 & -\sin(\theta) & 0 & \cos(\theta) \end{bmatrix}. \tag{6}$$

At the n th stage in the iteration, $\theta_1, \dots, \theta_{n-1}$ are held fixed, and θ_n is determined by minimizing the sum of

squared elements below the diagonal (the residuals, R_{HD}) in all L matrices A'_l (see the Appendix):

$$R_{HD} = \sum_{l=1}^L \sum_{p=2}^P \sum_{q=1}^{p-1} [(QA'_l Q^T)_{p,q}]^2. \tag{7}$$

The algorithm terminates when a full cycle of d iterations passes without further reduction of R_{HD} . In other words, termination occurs when the L matrices A'_l are concurrently transformed to be as close as possible to upper-triangular form. If the generators of the original system can be cast in a hierarchical form, this transformation returns a new set of upper-triangular matrices, A_{HD} ; for other systems, this procedure yields matrices A_{HD} that are as nearly upper triangular as possible, in the sense that R_{HD} has reached the minimum attainable value.

This process of minimization, by successive application of rotations $J_{u,v}(\theta)$, in a manner analogous to the Jacobi matrix diagonalization procedure, falls into the class of “direction set” methods.¹³ Direction set methods entail sequential one-dimensional minimizations—here, determination of θ_{fd} about the axis pair $[u_d, v_d]$ —of a multidimensional—here, d -dimensional—space. The sequence of axis pairs is chosen to cycle through all possible pairs of axes (i.e., for a 4×4 matrix, we use $d = 6$ and the cycle $[2, 1], [3, 1], [4, 1], [3, 2], [4, 2], [4, 3], [2, 1], \dots$). Numerical experimentation indicates that there is no particular advantage to any sequence, provided that each axis pair is examined once per cycle. (In general, however, direction set methods tend to be more efficient when the gradient of the function is used to guide the sequence.¹³) As with any multidimensional minimization, there is a risk of being trapped in a local minimum. While simulated annealing methods could be used to ensure that the global minimum is found, we have found that it is sufficient simply to run several minimizations in parallel, each preceded by a different initializing rotation Q_0 chosen from a collection of well-spaced rotations (see the Appendix and file “rotmesh.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>).

The rotation Q that attains the global minimum for R_{HD} (see file “globalmin.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>) is subsequently used to transform the temporal components T' , thereby identifying components $T_{HD} = QT'$, such that each component is, as nearly as possible, influenced only by itself and the more dominant components in the hierarchy.

Nonlinear Analysis of HD Resolved Components

In our application of HD to ictal ECoG, we are interested in the nonlinear features of the underlying time series identified by the HD procedure. We therefore analyzed the time series $T_{\text{HD}:p}$ (for $p=1\cdots P$) for nonlinear characteristics via nonlinear autoregressive analysis (NLAR), and compared them to the related NLAR analysis of the principal components T_p (for $p=1\cdots P$). The NLAR method is discussed in detail elsewhere.^{16,17}

Briefly, a 20-term linear autoregressive model for each time series is augmented by a single nonlinear term that models a joint influence of the signal at two prior times (times $n-j$ and $n-k$) on the signal at time n . The coefficients B and C of this nonlinear model,

$$T_{\text{HD}:p,n} = R_{\text{NLAR}:p,n} - B_{p,0} - \sum_{l=1}^{20} B_{p,l} T_{\text{HD}:p,n-l} - C_{j,k} T_{\text{HD}:p,n-j} T_{\text{HD}:p,n-k},$$

$$(j=1\cdots 20, k=1\cdots j), \quad (8)$$

are then determined separately for each of the 210 possible choices of a single nonlinear term (i.e., each of the 210 unique assignments of j and k to pairs in the set $\{1, \dots, 20\}$) by minimizing R_{NLAR} . This yields C , a 20×20 symmetric matrix of nonlinear autoregressive coefficients (where each entry in C is resolved by separate models).

The addition of each nonlinear term $C_{j,k}$ to the autoregressive model reduces the model variance V by an amount $\Delta V_{j,k}$ in comparison with the variance associated with the best 20-term linear AR model. Via an extension of the AIC,¹⁹ we use $N\Delta V/V$ to assess whether this reduction in variance is significant. We also examine $N\Delta V_{j,k}/V$ as a function of j and k ; this ‘‘NLAR fingerprint’’ summarizes the nonlinear dynamics of each time series over the lags considered, up to second order.

Definitions of Some Useful Quantities

In assessing the degree of hierarchical structure present in any set of MLAR coefficients, it is useful to define four quantities as follows. For any such set of MLAR model coefficients F (where F may be A , A' , or A_{HD} as below), the sum of all squared elements is

$$R_{\text{TOTAL}} = \sum_{l=1}^L \sum_{p=1}^P \sum_{q=1}^P F_{p,q,l}^2. \quad (9)$$

The size of the decoupled portion of the MLAR model (i.e., the extent to which the components’ dynamics are independent) is summarized by the sum of squared elements on the diagonal in all L matrices F_l :

$$R_{\text{DIAG}} = \sum_{l=1}^L \sum_{p=1}^P F_{p,p,l}^2. \quad (10)$$

The size of the hierarchical dynamic relationships in a MLAR model is summarized by the sum of squared elements above the diagonal in all L matrices F_l :

$$R_{\text{UPPER}} = \sum_{l=1}^L \sum_{p=1}^{P-1} \sum_{q=p+1}^P F_{p,q,l}^2. \quad (11)$$

The size of the dynamic relationships that are ‘‘antihierarchical’’ (i.e., neither hierarchical nor independent) is equal to the sum of squared elements below the diagonal in all L matrices F_l :

$$R_{\text{LOWER}} = \sum_{l=1}^L \sum_{p=2}^P \sum_{q=1}^{p-1} F_{p,q,l}^2. \quad (12)$$

Thus, $R_{\text{TOTAL}} = R_{\text{DIAG}} + R_{\text{UPPER}} + R_{\text{LOWER}}$.

These four quantities will be used to compare the degree of independent and hierarchical drive in the various MLAR models at each stage of analysis described above. R_{TOTAL} before and after the transformation identified by HD must be unchanged, since this transformation is orthogonal. However, the stage of noise decorrelation may change R_{TOTAL} because of the scale factor implicit in the transformation K . For this reason, we will typically focus on the quantities R_{DIAG} , R_{UPPER} , and R_{LOWER} , as fractions of R_{TOTAL} , and refer to these fractions as the independent, hierarchical, and antihierarchical drives, respectively.

We draw the readers attention to the fact that the arbitrary order chosen for the principal and noise-decorrelated components will define one set of values for R_{UPPER} and R_{LOWER} for each model, and that any reordering of those components potentially yields different sets of values. Thus, the $P!$ permutations of those components may change the apparent percent of hierarchical drive in the system. To distinguish the degree to which any such permutation will lead to a more or less hierarchical model versus the degree to which the HD method uncovers the inherent hierarchical structure in the system, we shall use the maximum value of R_{UPPER} —where R_{UPPER} is considered for all permutations of the components—for all models of the principal and noise-decorrelated components. This necessarily results in the minimum possible value for R_{LOWER} , and guarantees that $R_{\text{UPPER}} \geq R_{\text{LOWER}}$.

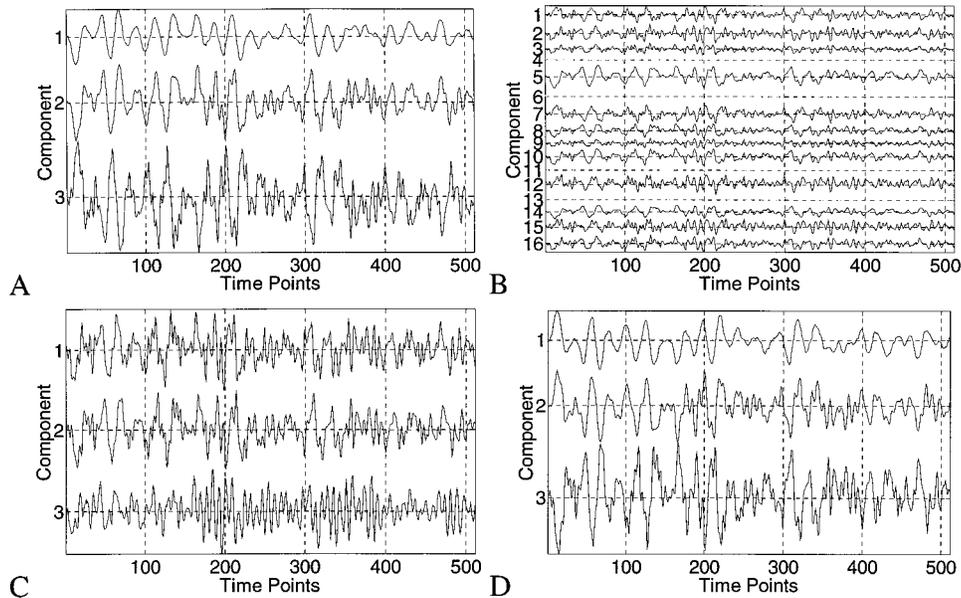


FIGURE 1. HD analysis of a three-generator hierarchical system constructed from a two-lag MLAR model. (A) Time series of the three hierarchical generators, showing one, two, and three dominant frequencies, respectively. (B) Sixteen time series composed of random linear combinations of the three time series in (A). (C) Time series of the first three principal components obtained from (B). (D) Time series of the three resolved hierarchical components obtained by rotation of the principal components. The ability of the HD method to extract hierarchical sources is evidenced by the virtually identical time series (disregarding signal polarity) in (A) and (D), as compared to the scrambled and dissimilar principal components in (C).

RESULTS

HD Analysis of Simulated Multivariate Time Series

To demonstrate proof of principle of the HD method we present several analyses of simulated multivariate systems. We first consider systems composed of three generators driven by Gaussian noise and linked in a hierarchical relationship. For the first system, the autoregressive model has nonzero coefficients up to two lags in the past (i.e., $L=2$); this situation permits unambiguous identification of the underlying generators (this data set and the code to reproduce Fig. 1 is archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>). We then limit the autoregressive model to a single lag (i.e., $L=1$) and show that this leads to degeneracy in the possible minimizing rotations Q , and consequent lack of uniqueness in the hierarchical decomposition. Finally, we examine a system of three *independent* generators, with nonzero autoregressive coefficients up to $L=2$; this example demonstrates that HD is also capable of extracting dynamically independent components.

To create the first simulated system of three hierarchically related time series, we fix a two-lag hierarchical MLAR model, and drive each generator with independent Gaussian white-noise input, generating three time series of length $N=512$. Coefficients above the diagonal in the MLAR matrices—which correspond to hierarchical influences between channels—were randomly chosen

from a Gaussian distribution with a mean of zero and unit variance. (It can be shown that for a hierarchical system off-diagonal coefficients do not affect stability. In our simulations, stability was ensured by choosing the MLAR coefficients on the diagonal to correspond to polynomials of order L whose roots have absolute value less than unity.⁹) Figure 1(A) shows the three resulting time series (the simulated generators). The influence of each generator on the subordinate generators is evident: the first generator exhibits a single dominant frequency, while the second and third generators have two and three dominant frequencies, respectively. (Analysis of the power spectra confirms this.) Estimation of the MLAR coefficients for the simulated generators via the Yule–Walker procedure will approximate (but not equal) the values used to construct the system. That is, the estimated coefficients yield $R_{\text{LOWER}}=0.0003$, a good approximation to the intended $R_{\text{LOWER}}=0$. By comparison, $R_{\text{DIAG}}=8.0316$ and $R_{\text{UPPER}}=4.7208$, indicating $\sim 63\%$ independent drive (i.e., $R_{\text{DIAG}}/R_{\text{TOTAL}}\approx 0.63$) and $\sim 37\%$ hierarchical drive (i.e., $R_{\text{UPPER}}/R_{\text{TOTAL}}\approx 0.37$) in the overall signal.

To simulate the kind of mixing of sources that might occur in EEG recordings, we take 16 random linear combinations of the three generators [see Fig. 1(B) and the file “lincomb.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>]. PCA decomposition of these 16 time series—the first step in the HD analysis—yields components that are remarkably dis-

similar to the original generators [see Fig. 1(C)]. Each principal component contains a mixture of frequencies, and the MLAR model of them gives no hint of the hierarchical structure ($R_{\text{TOTAL}}=22.9593$, $\sim 64\%$ independent and $\sim 24\%$ hierarchical drive). Because these signals are Gaussian, decorrelation by PCA guarantees information-theoretic independence at each time point, and is thus equivalent to ICA. In other words, examination of the higher-order moments *at zero lag* would not help to reveal the hierarchical structure. By contrast, the explicit use of signal dynamics allows HD to resolve hierarchically related components: the resulting MLAR model is characterized by $\sim 64\%$ independent and $\sim 36\%$ hierarchical drive ($R_{\text{TOTAL}}=12.5284$), close to the values obtained from the original generators. Strikingly, the resolved components [see Fig. 1(D)] are virtually identical to the original generators, other than signal polarity. (Amplitudes of the extracted components also match the original generators because we chose equal-variance noises as the driving terms.)

For a three-dimensional system such as this, we can visualize the search for the rotation Q (see the Methods section) by creating a representation of the space of possible three-dimensional rotations. To construct this representation of the rotation space, we recognize that any three-dimensional rotation is specified by an angle of rotation (θ , from 0 to π rad) around a unique axis in space. We can thus represent a three-dimensional rotation as a point r in a solid sphere, where the direction from the center of the sphere to r indicates the axis of the rotation, and the length of this vector represents θ . Figure 2 uses this representation of the rotation space to examine the global behavior of the residual values R_{HD} for the example outlined in Fig. 1. Notice that there is a local minimum in the neighborhood of the global minimum that provides a potential trap for the HD minimization algorithm.

The second example consists of three hierarchically related generators, but one in which the predefined MLAR model stipulates only single-lag influences ($R_{\text{TOTAL}}=11.8557$, $\sim 7\%$ independent and $\sim 93\%$ hierarchical drive). Figure 3 shows the time series ($N=512$) for this example in the same format as Fig. 1. Again, PCA is unable to resolve the hierarchical generators ($R_{\text{TOTAL}}=1.6037$, $\sim 73\%$ independent and $\sim 23\%$ hierarchical drive). By contrast, the HD algorithm uncovers the hierarchical components ($R_{\text{TOTAL}}=11.8533$, $\sim 8\%$ independent and $\sim 92\%$ hierarchical), as is evidenced by the similarity between Figs. 3(A) and 3(D). However, the hierarchical decomposition is not unique: the global behavior of the residuals, as shown in Fig. 4, demonstrates the existence of multiple global minima.

It can be shown that this nonuniqueness is expected for single-lag hierarchical MLAR models. For a generic P -component system, there are $P!$ rotations Q (not lim-

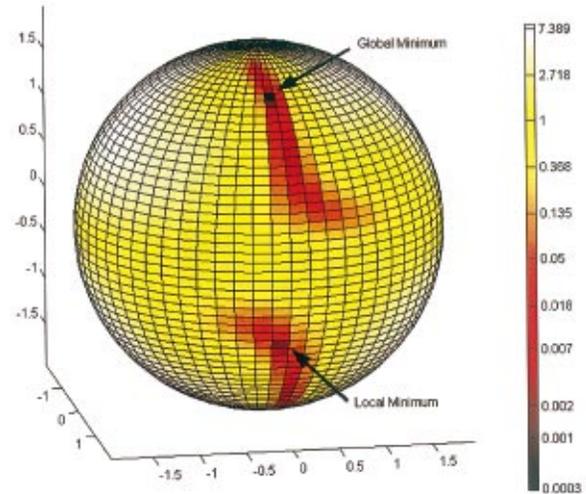


FIGURE 2. Graphical representation of a portion of the rotation space for the HD analysis of the system displayed in Fig. 1. Each rotation Q is specified by a point r in a solid sphere, where the direction from the center of the sphere to r indicates the axis of the rotation, and the length of this vector represents θ . The residual value R_{HD} for each rotation is represented in color (see scale to the right). In this portion of space (angle of rotation $\theta \approx 1.9478$ rad), we can see both a global minimum and a local minimum, which provides a potential trap for the HD algorithm.

ited to permuting the generators or changing their sign) that minimize R_{HD} . These rotations correspond to the $P!$ possible orderings of eigenvalues in the matrix A_{HD} . (To see this, we observe that transforming A' into upper-triangular form by a rotation amounts to choosing an appropriate orthonormal basis to form the columns of Q . When hierarchical generators exist, application of the Gram–Schmidt procedure to the eigenvectors of A' provides such an orthonormal basis. This is because triangular form for a matrix implies a full set of real eigenvalues and eigenvectors.¹¹ Each of the $P!$ orderings of the P eigenvectors yields a different basis via the Gram–Schmidt procedure, and hence, a distinct rotation Q and a distinct set of L triangular matrices $A_{\text{HD}:l} = QA'_lQ^T$.) This degeneracy does not occur with $L \geq 2$, since rotations Q that make QA'_1Q^T triangular do not necessarily make QA'_2Q^T triangular.

An extension of this analysis allows us to state the conditions under which a set of matrices A_1, A_2, \dots, A_L can be simultaneously transformed to upper-triangular form by an orthogonal matrix Q . First, each matrix must have a full set of real eigenvalues, so that each matrix can be individually transformed into upper-triangular form.¹¹ Second, it must be possible to order the eigenvectors for each matrix in such a way that (a) the first eigenvector of each matrix is identical (other than a scale factor), and (b) for each $k \leq Q$, the first k eigenvectors of

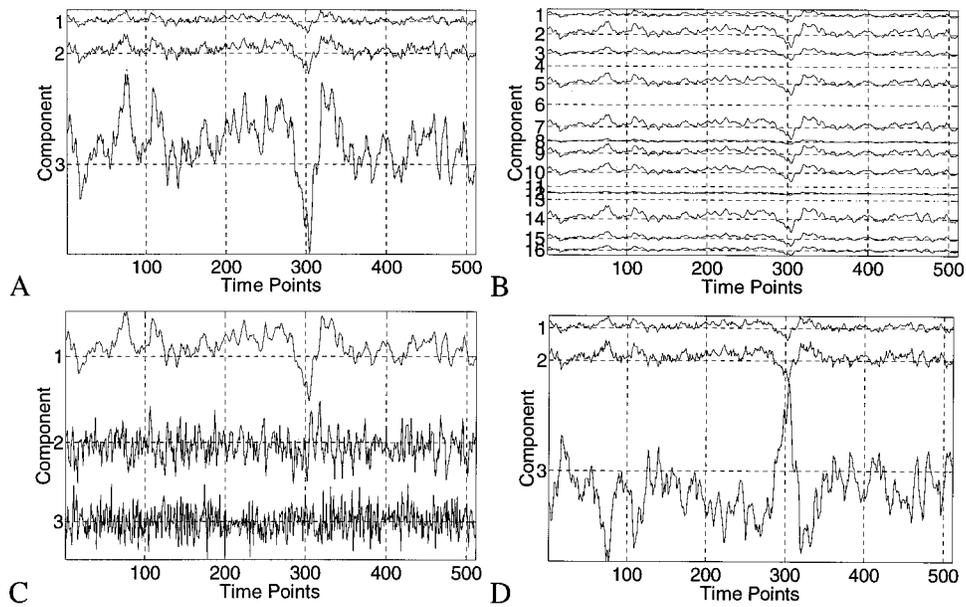


FIGURE 3. HD analysis of a three-generator hierarchical system constructed from a one-lag MLAR model. (A) Time series of the hierarchical generators. (B) Sixteen time series composed of random linear combinations of the three time series in (A). (C) Time series of the first three principal components obtained from (B). (D) Time series of the three resolved hierarchical components obtained by rotation of the principal components. Although HD extracts hierarchical sources, the hierarchical decomposition for the single-lag case is not unique (see Fig. 4).

each matrix span the same k -dimensional subspace. These conditions fix the order of the eigenvectors, provided that the number of lags $L > 1$.

The final simulated example consists of three independent generators, each driven by independent Gaussian noise. To simulate independence, the two-lag MLAR model has no off-diagonal terms. MLAR coefficients estimated from the time series ($N=512$) generated by this model [see Fig. 5(A)] indeed indicate essentially 100%

independent dynamics: $R_{\text{DIAG}}=7.5851$ and $R_{\text{TOTAL}}=7.5875$. When PCA is applied to a set of 16 random linear combinations of these generators [see Fig. 5(B)], the extracted components see [Fig. 5(C)] depend in part on the sizes of the contributions of the individual generators to each channel, and thus do not fully reveal the independent dynamics ($R_{\text{TOTAL}}=7.5888$, $\sim 94\%$ independent, $\sim 3\%$ hierarchical, and $\sim 3\%$ antihierarchical drive).

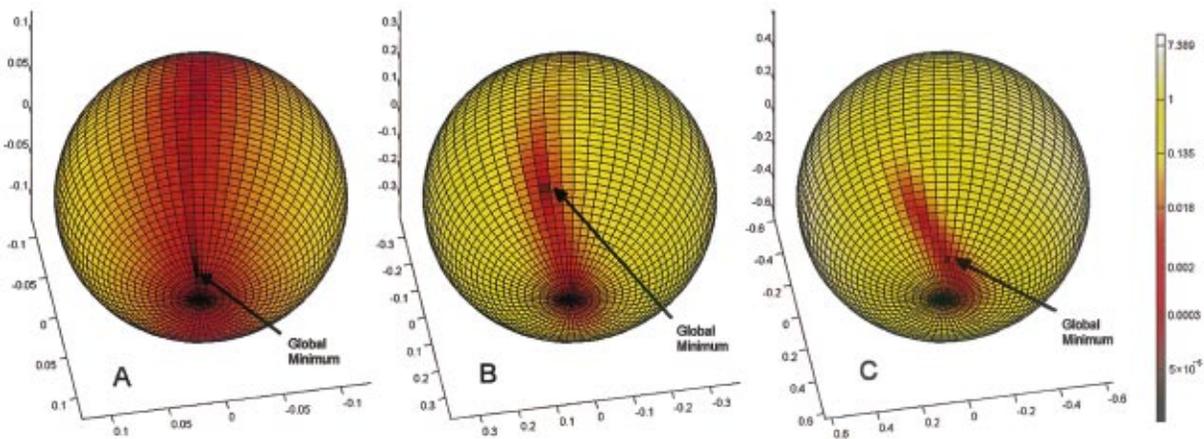


FIGURE 4. Graphical representation of three portions of the rotation space for the HD analysis of the system displayed in Fig. 3. The residual value R_{HD} for each rotation Q is represented as in Fig. 2. The three panels are chosen to contain three of the global minima for R_{HD} ($\theta \approx 0.1257$ rad in panel A, $\theta \approx 0.3770$ rad in panel B, and $\theta \approx 0.6283$ rad in panel C). This illustrates the degeneracy in the rotation Q for single-lag models. (Small differences in the color of each minimum result from the rough, finite sampling required by the graphical presentation.)

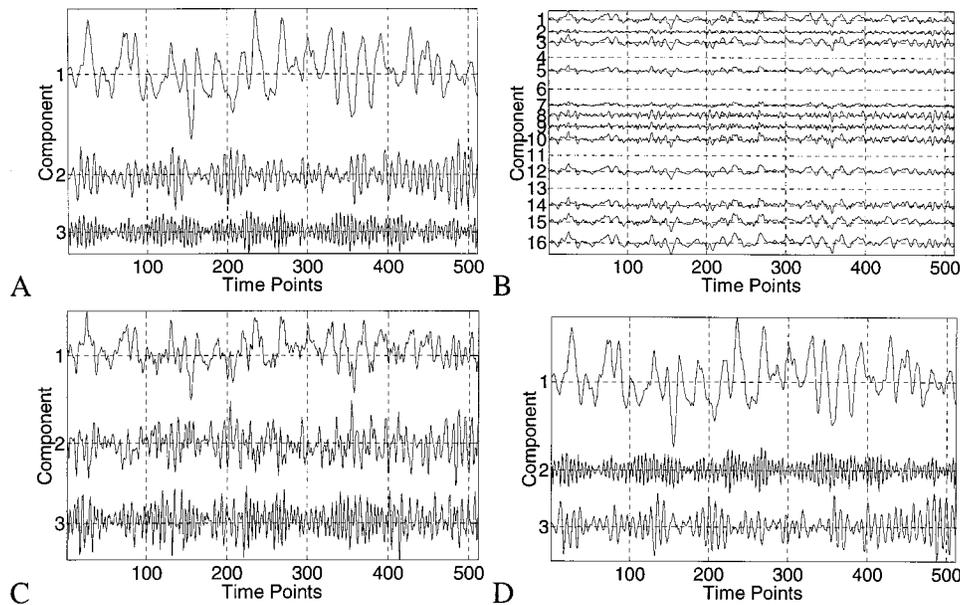


FIGURE 5. HD analysis of a three-generator *independent* system constructed from a two-lag MLAR model. (A) Time series of the independent generators. (B) Sixteen time series composed of random linear combinations of the three time series in (A). (C) Time series of the first three resolved principal components. (D) Time series of the three HD-resolved independent components. Here, we see that the HD method is also capable of extracting dynamically independent sources.

Following application of HD [see Fig. 5(D)], we find approximately 100% independent drive ($R_{\text{TOTAL}} = 7.5862$). However, since there is no hierarchical dependence in this example, the order of the components extracted by HD is arbitrary—here, the second and third generators are interchanged in Fig. 5(D) compared with Fig. 5(A).

HD Analysis of Ictal ECoG Data

Here, we use HD to reanalyze the ictal electrocorticographic records examined by Schiff and co-workers.¹⁵ In the analysis of electroencephalographic data, neither the number nor the time course of the underlying generators is known in advance. Therefore, we select the P components that constitute most of the original signal variance (see the Methods section), and are interested in both the degree of the resolved hierarchical structure and the dynamics of each component. We use the traditional AIC (Ref. 1) to determine the order of the MLAR models. Application of a stricter AIC that doubles the relative weighting of the number of degrees of freedom¹⁹ typically reduces the model order L by 1 but does not otherwise significantly alter the results (data not shown). Since individual seizure generators are commonly thought of as cyclic (i.e., roughly periodic in time), it is important to point out that the hierarchical relationship between generators, which we seek here, is by no means in contradiction to this view. That is, generators that are related to each other in a hierarchical fashion may themselves be cyclic in time.

Strikingly, the components resolved by HD are both proportionally more hierarchically influenced (compare columns labeled “ $R_{\text{UPPER}}/R_{\text{LOWER}}$ ” in Table 1) and more independently driven (compare columns labeled “% R_{DIAG} ” in Table 1) than the components resolved by PCA. For example, the MLAR model coefficients of the principal components— A in Eq. (3)—obtained from the third seizure record of patient 2 are characterized by $\sim 94.3\%$ independent drive (% R_{DIAG}) and $\sim 4.4\%$ hierarchical influence (% R_{UPPER}), yielding a hierarchical to antihierarchical ratio ($R_{\text{UPPER}}/R_{\text{LOWER}}$) of 3.4 (sixth row, first column, Table 1). By contrast, the MLAR model coefficients of the hierarchical components (A_{HD}) reveals that $\sim 96.2\%$ of the drive is independent (% R_{DIAG}), $\sim 3.6\%$ is hierarchical (% R_{UPPER}), and the ratio of hierarchical drive ($R_{\text{UPPER}}/R_{\text{LOWER}}$) is 16.3 (sixth row, third column, Table 1). Some, but not all, of this simplification of model structure is due to the reorganization of the model coefficients by the noise decorrelation procedure, which yields model coefficients A' . For this example, noise decorrelation alone yields $\sim 96.0\%$ independent (% R_{DIAG}) and $\sim 2.5\%$ hierarchical drive (% R_{UPPER}). But, in most instances, noise decorrelation fails to identify a proportionally more hierarchical organization ($R_{\text{UPPER}}/R_{\text{LOWER}}$ is only 1.7), whereas HD extracts components that are both more independent (larger $R_{\text{DIAG}}/R_{\text{TOTAL}}$, all records) and proportionally more hierarchical (greater ratio of R_{UPPER} to R_{LOWER} in 7/9 records), as detailed in Table 1. Moreover, the contribution of each of the L coefficient matrices to the indices

TABLE 1. Summary of the percentage of independent, hierarchical, and antihierarchical drive for the MLAR model coefficients of the principal components decomposition (A), noise-decorrelated PCA decomposition (A'), and hierarchical decomposition (A_{HD}) for all patient records. The percentages are obtained by normalizing the values for the independent (R_{DIAG}), hierarchical (R_{UPPER}), and antihierarchical (R_{LOWER}) drive by their total (R_{TOTAL}) in each instance, here labeled $\%R_{DIAG}$, $\%R_{UPPER}$, and $\%R_{LOWER}$, respectively. The unitless ratio R_{UPPER}/R_{LOWER} reflects the relative proportion of hierarchical influence in the system. Percentages for R_{UPPER} and R_{LOWER} for A and A' reflect the most hierarchical ordering of the respective components, as outlined in the Methods section, such that $R_{UPPER} \geq R_{LOWER}$. All values are rounded to the nearest 0.1 for clarity.

| | Principal components | | | Noise-decorrelated components | | | Hierarchical components | | | | | |
|-----------|----------------------|---------------|---------------|-------------------------------|--------------|---------------|-------------------------|-----------------------|--------------|---------------|---------------|-----------------------|
| | $\%R_{DIAG}$ | $\%R_{UPPER}$ | $\%R_{LOWER}$ | R_{UPPER}/R_{LOWER} | $\%R_{DIAG}$ | $\%R_{UPPER}$ | $\%R_{LOWER}$ | R_{UPPER}/R_{LOWER} | $\%R_{DIAG}$ | $\%R_{UPPER}$ | $\%R_{LOWER}$ | R_{UPPER}/R_{LOWER} |
| Patient 1 | | | | | | | | | | | | |
| Seizure 1 | 89.6 | 8.7 | 1.7 | 5.2 | 97.2 | 1.8 | 1.1 | 1.6 | 98.5 | 1.2 | 0.4 | 3.2 |
| Seizure 2 | 69.8 | 26.2 | 4.0 | 6.5 | 84.4 | 10.6 | 5.0 | 2.1 | 97.4 | 2.3 | 0.3 | 8.3 |
| Seizure 3 | 84.6 | 12.8 | 2.7 | 4.8 | 88.8 | 7.7 | 3.5 | 2.2 | 95.9 | 3.4 | 0.7 | 4.7 |
| Patient 2 | | | | | | | | | | | | |
| Seizure 1 | 82.5 | 11.9 | 5.6 | 2.1 | 84.1 | 10.8 | 5.1 | 2.1 | 88.3 | 9.9 | 1.8 | 5.4 |
| Seizure 2 | 93.2 | 6.3 | 0.6 | 11.3 | 96.7 | 2.4 | 0.9 | 2.6 | 98.1 | 1.8 | 0.1 | 14.7 |
| Seizure 3 | 94.3 | 4.4 | 1.3 | 3.4 | 96.0 | 2.5 | 1.5 | 1.7 | 96.2 | 3.6 | 0.2 | 16.3 |
| Patient 3 | | | | | | | | | | | | |
| Seizure 1 | 93.6 | 4.9 | 1.5 | 3.3 | 94.4 | 4.0 | 1.6 | 2.5 | 95.4 | 4.1 | 0.5 | 8.5 |
| Seizure 2 | 99.7 | 0.2 | 0.1 | 4.0 | 99.8 | 0.2 | 0.0 | 5.6 | 99.8 | 0.2 | 0.0 | 10.4 |
| Patient 4 | | | | | | | | | | | | |
| Seizure 1 | 96.3 | 2.9 | 0.8 | 3.7 | 98.8 | 0.8 | 0.4 | 1.9 | 99.7 | 0.2 | 0.0 | 13.3 |

R_{UPPER} and R_{LOWER} are approximately equal in all cases (data not shown), perhaps as a consequence of the fact that the HD algorithm treats all L matrices equally.

As alluded to in the Introduction, examination of the nonlinear dynamics present in the extracted components provides another means to interpret the results of the HD method. To assess the nonlinear dynamics in an individual time series we use the NLAR fingerprint.¹⁹ This approach identified common nonlinear dynamics between 3/s spike-and-wave seizure traces,^{16,17} and the principal components derived from the partial complex seizure records of patients 1 and 2.¹⁵ Figure 6 compares the NLAR analyses of the principal components (as analyzed in Ref. 15) with that of the hierarchical components. In Fig. 6, the percent of variance of the original signal that each component explains is shown by the height of a bar. (For PCA this necessarily decreases as the component number increases.) The significance of the nonlinear dynamics identified in each component, as summarized by $N\Delta V/V$, is shown by the corresponding point on a solid line [the dashed line indicates significant nonlinearities, $N\Delta V/V=4$ (Ref. 19)]. Whereas the principal components with the largest $N\Delta V/V$ are typically those components that explain only a small fraction of the variance (leaving their significance unclear), HD analysis reveals that the hierarchical components with the greatest nonlinear structure can be the most autonomous components (lowest numbered component). This is seen in seizure 1 of patient 1, and in seizures 1 and 3 of patient 2.

This latter record is particularly striking because HD analysis uncovers a hierarchical component with significant nonlinearities ($N\Delta V/V \approx 8$) even though none of the principal components display significant nonlinear characteristics ($N\Delta V/V < 4$). We hypothesize that this reflects a more effective separation of underlying generators by HD, and that the components extracted by PCA reveal only diluted effects of the nonlinearity due to mixing. Additionally, five components were retained for seizure 3 of patient 1 (following Schiff *et al.*)¹⁵—this is an exception to the guidelines mentioned in the Methods section—because the fifth component has an $N\Delta V/V \approx 5.2$. Here, HD analysis with five components exposes significant nonlinear characteristics in the third hierarchical component, which accounts for nearly 50% of the original signal variance. Finally, we note that there are differences in the analyses of individual records within and across patients, both in terms of the relative sizes of the hierarchical components and in the HD components that demonstrate the greatest nonlinearities. We suspect that these differences reflect physiologic differences in the events themselves rather than nonrobustness of the analysis method, since (a) the records themselves appeared artifact free, (b) the results of the analysis are not significantly altered by small changes in the model order,

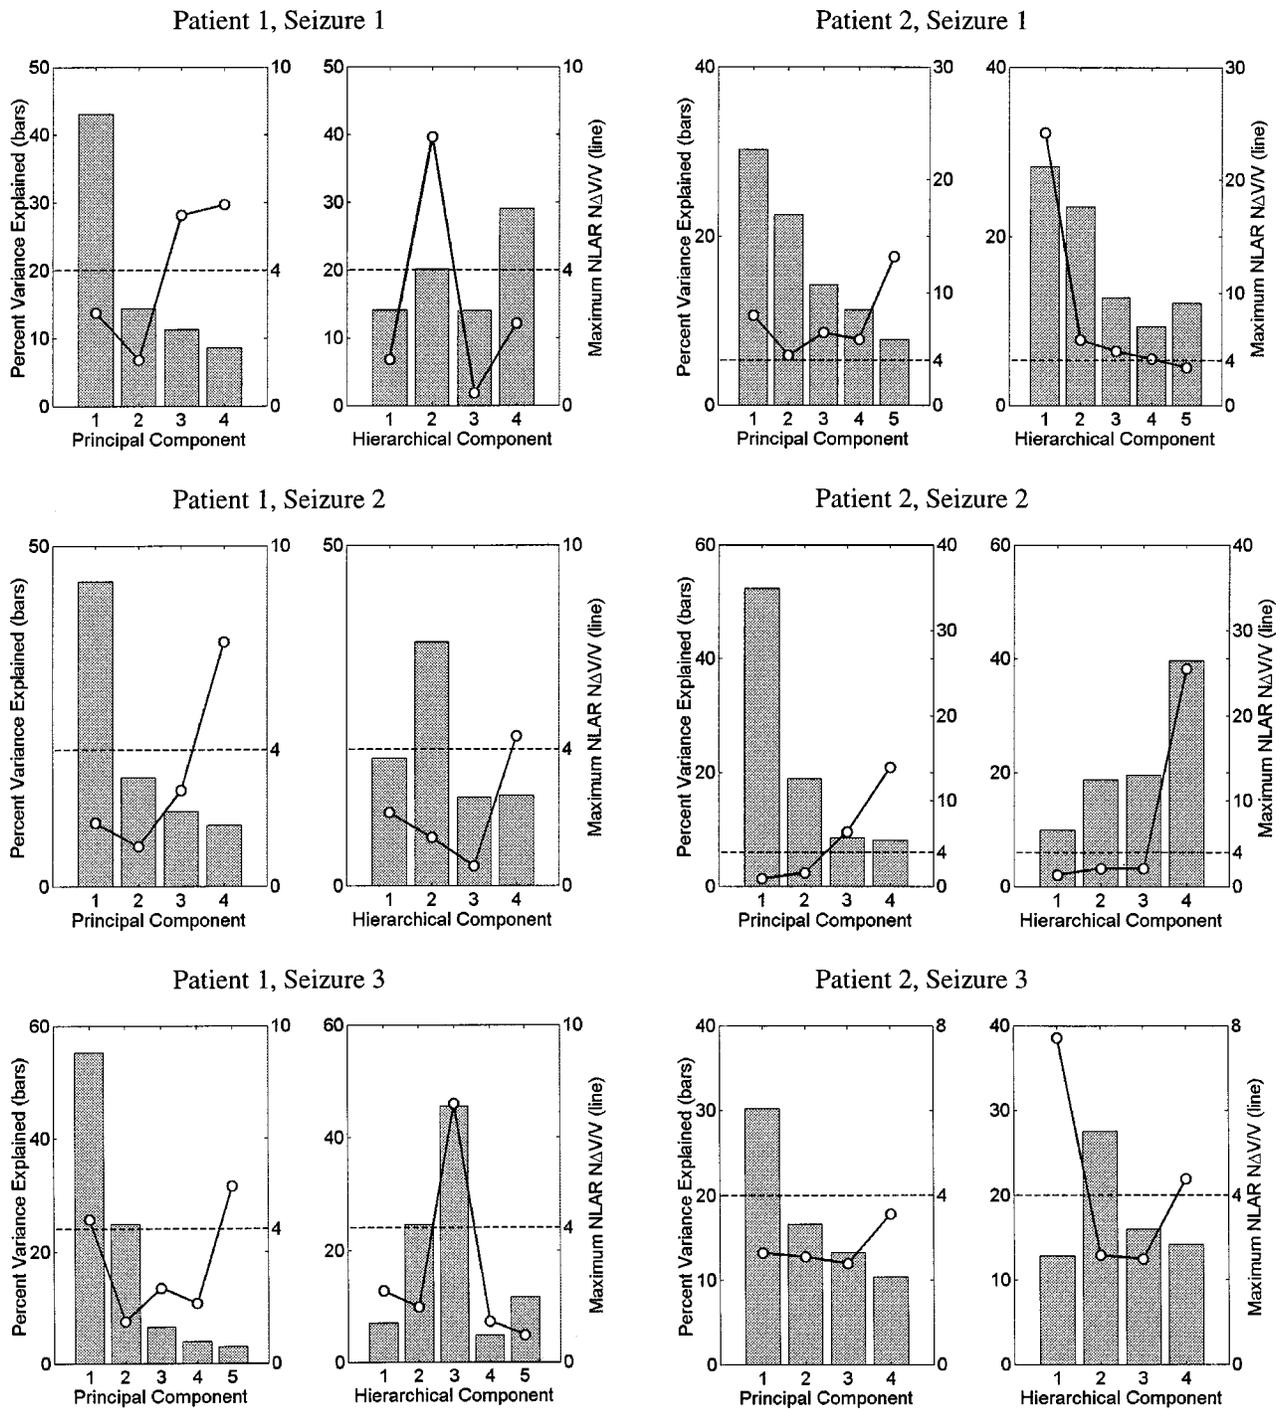


FIGURE 6. Summary of the NLAR analyses of the principal components (see Ref. 15) and hierarchical components derived from each record of patients 1 and 2. The percent variance explained by each component is shown by the height of a bar, while the maximum $N\Delta V/V$ for that component is shown by the corresponding point on a solid line. The dashed line in each graph indicates $N\Delta V/V=4$, the criterion we have used for detection of significant dynamical nonlinearities (see Ref. 19). Often HD analysis reveals that the component with the greatest nonlinear structure is one of the more autonomous components. Patient 2, seizures 1 and 3 are particularly striking examples of this.

and (c) for systems in which the internal structure is known (Figs. 1–5), HD provides an accurate decomposition.

DISCUSSION

Hierarchical Decomposition (HD) Method

In summary, the ability of HD to uncover hierarchically related components results from its fundamental use of the dynamics inherent in any multivariate time series. We use a MLAR model to simultaneously enforce two conditions upon the components. The first assumption is that the noise (i.e., residual values between the observed behavior and the MLAR model) in each component is independent and of equal magnitude. This makes concrete the assumption that the stochastic components of the generators are independent, and it arbitrarily sets a scale for the size of each generator. It is implemented in a single step, and can always be achieved, independent of the dynamics of the system. The second condition seeks a hierarchical relationship among components. It is implemented iteratively, via successive orthogonal rotations of the components to attain a set of MLAR matrices that are as triangular as possible.

Using simulated systems of hierarchically organized generators, we show that HD is capable of extracting the original sources, whereas PCA or ICA cannot (Figs. 1 and 3). The hierarchical minimization procedure is graphically depicted in Figs. 2 and 4, and highlights two important caveats: (a) finding the global minimum for the model (i.e., the most triangular) requires avoiding various local minima, and (b) models with simple dynamics (i.e., where only single-lag influences are significant) are degenerate, having multiple global minima. The latter point underscores the fact that separation of the generators relies intimately on their dynamics. If the MLAR model is trivial (i.e., $L=0$, no significant dynamics in the individual components), then the HD procedure is trivial, too, and the continuum of ambiguities intrinsic to PCA remain. If the MLAR model contains one lag only, then the HD procedure can narrow down possible models of the system to a discrete set of size $P!$ (where P is the number of generators). If the MLAR model contains two or more lags, then the HD decomposition is typically unique, and if the original system has a hierarchical structure, the algorithm will identify it.

In these simulations, we know exactly how many generators (i.e., three) contribute to each 16-channel data set we create. Moreover, the first three principal components necessarily account for 100% of the variance of the system. This highlights two potentially important limitations of the HD method. In real-world data, by selecting only the most important principal components (see the Methods section), we make the explicit compromise of not

accounting for all of the variance of the signal ($\sim 20\%$ in the current application to EEG). Although in principle HD analysis could be performed on the raw data, this compromise is often necessary because identification of a minimum in the extremely high-dimensional space of all possible rotations is not likely to be robust. Additionally, one might wonder how HD would behave if performed on fewer components than the number of generators known to be present. Here, the results of HD analysis on numerical simulations provide reassurance (data not shown), since main features of the decomposition are preserved. For example, three hierarchical sources extracted by HD from a system built with four generators look more similar to the four original generators than do any of the principal components, in so far as they seem to be less arbitrarily mixed.

For many previous approaches to the analysis of multivariate time series (including methods based exclusively upon PCA and related methods), the underlying rationale is that the observed signals represent linear mixtures of independent sources (or generators). In many systems, such as the brain, the observed time series represent mixtures of sources that are not strictly independent, but rather dynamically influence one another, often in stereotyped patterns.¹⁸ This is the rationale for the present approach, in which we separate multivariate time series by assuming instead that the sources influence each other in a hierarchical fashion. More generally, the method can be adapted to seek interrelationships of other canonical types (e.g., cyclic) by appropriate modification of Eq. (7). Thus, the HD method may be considered to be a special case of a large set of canonical decomposition procedures.

The use of causality (i.e., the influence of prior observations on the current observation) and its importance in the HD algorithm deserves emphasis. The MLAR technique used by HD is implicitly dependent upon causal relations among the observations in a multivariate time series data set. That is, the MLAR model description of the data relies upon the order of the observations, as it describes the effect of prior observations upon current observations. This description would be obscured by time-reversing or shuffling the data. HD uses these causal relationships to guide the search for a rotation Q (see the Methods section) that yields a hierarchical interrelationship among components. The hierarchy of components delineates a specific type of causal relationship in a data set, namely, that the current value for each component is affected by prior values of itself and more dominant components in the hierarchy.

This approach is in direct contrast to other causal techniques for multivariate data analysis, including Granger causality,⁷ generalized autoregressive conditional heteroskedasticity (GARCH),³ adaptive multivariate autoregressive modeling (AMVAR),⁵ and directed

coherence.⁸ These methods assume that the original time series represents the unmixed sources, and then look for and characterize the degree and type of causal relations among the data. In contrast, HD assumes a general form for the causal interaction, and uses this assumption to direct the decomposition into sources. (Granger⁷ notes that multivariate data sets can have a “spurious causal-chain appearance,” and examples of this can be readily constructed by linear mixing of independent time series. However, as pointed out above, for multivariate data sets whose MLAR structure includes more than one lag, reduction to hierarchical form is not guaranteed, and if such a reduction is possible it is typically unambiguous.) From this point of view, HD appears to be unique.

Application of HD to Ictal ECoG

The degree of hierarchical and independent structure in all components derived from ictal ECoG records of patients 1–4 is summarized in Table 1. Examination of Table 1 indicates that the components resolved by HD are both proportionally more hierarchical (i.e., see the columns labeled “ R_{UPPER}/R_{LOWER} ” under each type of component) and more independently driven (i.e., see the columns labeled “% R_{DIAG} ” under each type of component) than the components resolved by PCA. Also evident from Table 1 is the fact that reorganization of the components by noise decorrelation is partially responsible for this simplification of model structure. However, noise decorrelation alone typically results in a more even distribution of off-diagonal coefficients (i.e., values for R_{UPPER}/R_{LOWER} that are closer to 1). In addition, notice that the proportion of hierarchical drive in these systems revealed by HD is sometimes quite large, particularly for seizures 2 and 3 of patient 2 (R_{UPPER}/R_{LOWER} equals 14.7 and 16.3, respectively). Nevertheless, there is also variability, both from patient to patient and within patients, in the degree of hierarchical relationship between the underlying seizure generators.

Previously, we showed that PCA decomposition of temporal lobe seizure records uncovers components whose nonlinear dynamics, as characterized by NLAR fingerprints, resemble those of the EEG during absence seizures.¹⁵ Because the principal components exhibiting significant nonlinear characteristics account for relatively small fractions of the original signal variance, the biological significance of the nonlinear signal was unclear. In contrast, HD analysis of these same records uncovers components containing the same nonlinearities, but they can appear relatively early in the hierarchy (see Fig. 6). The position of these nonlinearities in the hierarchy appears to be independent of the amount of original signal variance that these components explain. Furthermore, HD analysis may clarify the biological significance of these nonlinearities: they possibly reflect deep generators

that contribute little to the surface EEG, but affect a large portion of the entire neuronal system. More generally, the HD analyses demonstrate that although the nonlinearities are not always a major contribution to the overall signal variance, they nevertheless may exert influence over all or most of the system.

In addition, the NLAR fingerprints of the hierarchical components support the connection between temporal lobe and absence seizures, and strengthen the suggestion that both seizure types may result from similar underlying neuronal mechanisms.^{14,15} Significant nonlinearities in these fingerprints appear in the vicinity of 90 and 150 ms,¹⁴ as is seen in the NLAR fingerprints of absence seizure traces,¹⁷ and as reported earlier for the NLAR fingerprints of the principal components.¹⁵ Notice, however, that the nonlinearities tend to be more statistically significant in the hierarchical components than in the principal components. Our findings may help explain why even focal temporal lobe seizures produce global alterations in awareness and behavior: the importance of a neuronal circuit in driving global brain activity (i.e., its position in the hierarchy) may be disproportionate to the fraction of the variance of the surface activity that it explains, especially for generators that are deep. The results of HD modeling are thus consistent with the notion that temporal lobe and absence seizures reflect alteration of selective circuit mechanisms that play a key role in forebrain integration.¹⁵ The application of the HD technique to ECoG records of temporal lobe epilepsy suggests that there is an interpretive benefit gained by extracting hierarchical, rather than independent, components.

ACKNOWLEDGMENTS

The authors thank The Lederman Family Foundation and the Biomedical Engineering Program of Cornell University for partial support of this research. M.A.R. is supported by EY7138, N.D.S. is supported by NS02172, and J.D.V. is supported by EY9314 and EY7977.

APPENDIX

This Appendix provides additional details for two aspects of the algorithm we have presented above (see the Methods section). The description is keyed to MATLAB code, archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>, but does not fully describe the code itself (see embedded documentation within the code for further details). Where the code differs slightly from the calculations presented in this exposition such differences are noted. We make no claim that this code is optimized but rather hope that it will provide a useful starting point for other researchers.

Determination of the Rotation Angle (θ_n)

This calculation is carried out by the file “rotation.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>.

For any square matrix A , transformation via a plane rotation $J_{u,v}(\theta)$ [see Eqs. (6) and (7)] produces a new square matrix $J_{u,v}(\theta)AJ_{u,v}^T(\theta)=A'$. This transformation only alters the elements in both the rows and columns u and v in matrix A —whose elements are denoted $a_{i,j}$. We are interested in minimizing the sum of squared elements below the diagonal in A' —whose elements are denoted $a'_{i,j}$. Thus we require equations for those elements $a'_{i,j}$ below the diagonal (i.e., such that $i>j$) that are altered by the rotation. [For historical reasons, the code minimizes the sum of squared elements *above* the diagonal of A'^T , and each rotation induces the transformation $J_{u,v}^T(\theta)A^TJ_{u,v}(\theta)=A'^T$. Because $J_{u,v}(\theta)$ is antisymmetric this is fundamentally equivalent to performing the rotation as described in Eq. (7).]

For convenience we can divide the analysis of matrix entries below the diagonal of A' (above the diagonal of A'^T) into five groups of elements (depending on the axes values u and v , some of these groups may be empty). In what follows, $s \in \{u,v\}$, $p = \min(u,v)$, $r = \max(u,v)$, and $n < p < q < r < t$: (1) $a'_{s,n}$, (2) $a'_{t,s}$, (3) $a'_{r,q}$, (4) $a'_{q,p}$, and (5) $a'_{r,p}$. Elementary matrix algebra used in conjunction with trigonometric identities shows that groups 1 and 2 will never affect the residual value R_{HD} [see Eq. (7)]. For group 1 elements

$$a'_{u,n}{}^2 + a'_{v,n}{}^2 = [a_{u,n} \cos(\theta) + a_{v,n} \sin(\theta)]^2 + [a_{v,n} \cos(\theta) - a_{u,n} \sin(\theta)]^2 = a_{u,n}^2 + a_{v,n}^2,$$

and similarly for group 2 elements,

$$a'_{t,u}{}^2 + a'_{t,v}{}^2 = [a_{t,u} \cos(\theta) + a_{t,v} \sin(\theta)]^2 + [a_{t,v} \cos(\theta) - a_{t,u} \sin(\theta)]^2 = a_{t,u}^2 + a_{t,v}^2.$$

In contrast, elements in groups 3 and 4, when present, typically affect the residual value R_{HD} , as will the group 5 element, which is always present.

For groups 3 and 4, we use the trigonometric identities $\cos^2(\theta) = 0.5 + 0.5 \cos(2\theta)$, $\sin^2(\theta) = 0.5 - 0.5 \cos(2\theta)$, and $2 \sin(\theta)\cos(\theta) = \sin(2\theta)$, and find:

$$\begin{aligned} a'_{r,q}{}^2 + a'_{q,p}{}^2 &= [a_{r,q} \cos(\theta) - a_{p,q} \sin(\theta)]^2 \\ &\quad + [a_{q,p} \cos(\theta) + a_{q,r} \sin(\theta)]^2 \\ &= a_{r,q}^2 \cos^2(\theta) - 2a_{r,q}a_{p,q} \cos(\theta)\sin(\theta) \\ &\quad + a_{p,q}^2 \sin^2(\theta) + a_{q,p}^2 \cos^2(\theta) \end{aligned}$$

$$\begin{aligned} &+ 2a_{q,p}a_{p,r} \cos(\theta)\sin(\theta) \\ &+ a_{q,r}^2 \sin^2(\theta) \\ &= (a_{r,q}^2 + a_{q,p}^2) \cos^2(\theta) \\ &\quad + 2(a_{q,p}a_{q,r} - a_{r,q}a_{p,q}) \cos(\theta)\sin(\theta) \\ &\quad + (a_{p,q}^2 + a_{q,r}^2) \sin^2(\theta) \\ &= (a_{q,p}a_{p,r} - a_{r,q}a_{p,q}) \sin(2\theta) \\ &\quad + 0.5(a_{r,q}^2 + a_{q,p}^2 - a_{p,q}^2 - a_{q,r}^2) \cos(2\theta) \\ &\quad + 0.5(a_{r,q}^2 + a_{q,p}^2 + a_{p,q}^2 + a_{q,r}^2). \end{aligned}$$

This equation may be further simplified by combining the sine and cosine terms into a single trigonometric function. (The addition of sine and cosine waves of the same frequency yields another sinusoidal wave at the same frequency, whose amplitude is the square root of the sum of the squared amplitudes of the original waves, and whose phase is the arctangent of the ratio of the amplitudes of the original waves.) Accordingly, we make the substitutions $E \sin(\phi) = (a_{q,p}a_{q,r} - a_{r,q}a_{p,q})$ and $E \cos(\phi) = 0.5(a_{r,q}^2 + a_{q,p}^2 - a_{p,q}^2 - a_{q,r}^2)$ which, used in conjunction with the trigonometric identity $\sin(\alpha)\sin(\beta) + \cos(\alpha)\cos(\beta) = \cos(\alpha - \beta)$, leads to

$$a'_{r,q}{}^2 + a'_{q,p}{}^2 = E \cos(2\theta - \phi) + 0.5(a_{r,q}^2 + a_{q,p}^2 + a_{p,q}^2 + a_{q,r}^2). \quad (\text{A1})$$

Additionally, the equation for the group 5 element simplifies in an analogous manner, using the substitutions $V \sin(\psi) = 0.5(a_{r,r} - a_{p,p})$ and $V \cos(\psi) = 0.5(a_{r,p} + a_{p,r})$, and the trigonometric identities above:

$$\begin{aligned} a'_{r,q} &= [a_{r,p} \cos(\theta) + a_{r,r} \sin(\theta)] \cos(\theta) \\ &\quad - [a_{p,p} \cos(\theta) + a_{p,r} \sin(\theta)] \sin(\theta) \\ &= a_{r,p} \cos^2(\theta) \\ &\quad + (a_{r,r} - a_{p,p}) \cos(\theta)\sin(\theta) \\ &\quad - a_{p,r} \sin^2(\theta) \\ &= 0.5(a_{r,r} - a_{p,p}) \sin(2\theta) \\ &\quad + 0.5(a_{r,p} + a_{p,r}) \cos(2\theta) \\ &\quad + 0.5(a_{r,p} - a_{p,r}) \\ &= V \cos(2\theta - \psi) + 0.5(a_{r,p} - a_{p,r}). \end{aligned}$$

Finally, we calculate the square of the equation for the group 5 element and are left with

$$\begin{aligned} a'_{r,p}{}^2 &= 0.5V^2 \cos(4\theta - 2\psi) + V(a_{r,p} - a_{p,r}) \cos(2\theta - \psi) \\ &\quad + 0.25(a_{r,p} - a_{p,r})^2 + 0.5V^2. \end{aligned} \quad (\text{A2})$$

Thus, we must minimize—as a function of θ —an expression that consists of the above contribution from Eq. (A2) for the group 5 element, and a contribution from Eq. (A1) for each pair of elements in groups 3 and 4 that are present (elements in a pair are always *both* absent or *both* present). Elements in groups 1 and 2 may be neglected, since they contribute a constant that does not depend on θ , as may the constant terms present in Eqs. (A1) and (A2).

In our code, matrix entries in A that contribute to the trigonometric functions governing group 5 (see lines 81–89) and, when necessary, groups 3 and 4 (see lines 91–104) are collected and combined into matrix constants, including E and V , ϕ , and ψ . The trigonometric equation to be minimized is assembled from Eq. (A2) for each group 5 element (see lines 109–112) and, as needed, from Eq. (A1) for pairs of elements in groups 3 and 4. When elements in groups 3 and 4 are present, the function to be minimized may be quite large (see lines 125–127), whereas the minimization is simpler when groups 3 and 4 are empty (see lines 128–131). Finally, we use a customized bisection routine to find the value for θ that minimizes the complete trigonometric function (see subfunction “trigmin,” lines 240–266).

Creating Well-Spaced Initializing Rotations (Q_0)

The initializing rotations Q_0 are determined as follows (see file “rotmesh.m” archived in “HD_ABME_code.zip” at <http://www.med.cornell.edu/research/hds>). For each possible choice of a pair of axes $[u, v]$ we build plane rotations at two rotation angles $\theta = \pm 2\pi/3$ (see lines 56–74). Thus, for a system consisting of P components there are $d = P(P-1)/2$ possible pairs of axes, and $2d$ possible initializing rotations Q_0 . We randomly order all possible initializing rotations Q_0 (see lines 76–78) and use one for each parallel minimization performed [see Eq. (5) and associated text].

REFERENCES

- ¹Akaike, H. A new look at statistical model identification. *IEEE Trans. Autom. Control* 19:716–723, 1974.
- ²Bell, A. J., and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* 7:1129–1159, 1995.
- ³Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J. Econometr.* 31:307–327, 1986.
- ⁴Borg, I., and P. Groenen. *Moder Multidimensional Scaling: Theory and Applications*. New York: Springer, 1997, pp. 132–134 and 399–402.
- ⁵Ding, M., S. L. Bressler, W. Yang, and H. Liang. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: Data processing, model validation, and variability assessment. *Biol. Cybern.* 83:35–45, 2000.
- ⁶Gersh, W. and J. Yonemoto. Parametric time series models for multivariate EEG analysis. *Comput. Biomed. Res.* 10:113–125, 1977.
- ⁷Granger, C. W. J. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37:424–438, 1969.
- ⁸Kaminski, M. J., and K. J. Blinowska. A new method of the description of the information flow in the brain structures. *Biol. Cybern.* 65:203–210, 1991.
- ⁹Kay, S. M. *Modern Spectral Estimation: Theory and Application*. Englewood Cliffs, NJ: Prentice-Hall, 1988, pp. 141 and 142.
- ¹⁰Lee, D. D., and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature (London)* 401:788–791, 1999.
- ¹¹Mirsky, L. *An Introduction to Linear Algebra*. New York: Dover, 1990, pp. 236–247 and 306–311.
- ¹²Mitra, P. P., and B. Pesaran. Analysis of dynamic brain imaging data. *Biophys. J.* 76:691–708, 1999.
- ¹³Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. New York: Cambridge University Press, 1992, pp. 412–413 and 463–466.
- ¹⁴Repucci, M. A., N. D. Schiff, and J. D. Victor. Implications for the dynamics of temporal lobe seizures derived from multivariate autoregressive analysis of electroencephalograms. *Proceedings of the Society for Neuroscience 29th Annual Meeting, Miami, FL (242.9)*, 1999.
- ¹⁵Schiff, N. D., D. R. Labar, and J. D. Victor. Common dynamics in temporal lobe seizures and absence seizures. *Neuroscience (Oxford)* 91:417–428, 1999.
- ¹⁶Schiff, N. D., J. D. Victor, and A. Canel. Nonlinear autoregressive analysis of the 3/s ictal electroencephalogram: Implications for underlying dynamics. *Biol. Cybern.* 72:527–532, 1995.
- ¹⁷Schiff, N. D., J. D. Victor, A. Canel, and D. R. Labar. Characteristic nonlinearities of the 3/s ictal electroencephalogram identified by nonlinear autoregressive analysis. *Biol. Cybern.* 72:519–526, 1995.
- ¹⁸Steriade, M., P. Gloor, R. R. Llinás, F. H. Lopes da Silva, and M.-M. Mesulam. Report of IFCN committee on basic mechanisms: Basic mechanisms of cerebral rhythmic activities. *Electroencephalogr. Clin. Neurophysiol.* 76:481–508, 1990.
- ¹⁹Victor, J. D., and A. Canel. A relation between the Akaike criterion and reliability of parameter estimates, with application to nonlinear autoregressive modeling of the ictal EEG. *Ann. Biomed. Eng.* 20:167–180, 1992.
- ²⁰Wei, W. W. S. *Time Series Analysis: Univariate and Multivariate Methods*. New York: Addison-Wesley, 1990, 478 pp.
- ²¹Yoshioka, K., F. Matsuda, K. Takakura, Y. Noda, K. Imakawa, and S. Sakai. Determination of genes involved in the process of implantation: Application of GeneChip to scan 6500 genes. *Biochem. Biophys. Res. Commun.* 272:531–538, 2000.
- ²²Yule, G. U. On a method of investigating periodicities in disturbed series with special reference to Wolfer’s sunspot numbers. *Philos. Trans. R. Soc. London, Ser. A* 226:267–298, 1927.