# A method for decomposing multivariate time series into a causal hierarchy within specific frequency bands

Jonathan D. Drover and Nicholas D. Schiff

July 12, 2018

**Abstract**

We propose a method - Frequency extracted hierarchical decomposition (FEHD) - for studying multivariate time series that identifies linear combinations of its components that possess a causally hierarchical structure - the method orders the components so that those at the "top" of the hierarchy drive those below. The method shares many of the features of the "hierarchical decomposition" method of Repucci, et al. (2001) but makes a crucial advance - the proposed method is capable of determining this causal hierarchy over arbitrarily specified frequency bands. Additionally, a novel minimization strategy is used to generate the decomposition resulting in an increase in stability, reliability, and an improvement in the sensitivity to model parameters. We demonstrate the utility of the method by applying it to both artificial time series constructed to have specific causal graphs, and to the EEG of healthy volunteers and patient subjects who are recovering from a severe brain injury.

## 1 Introduction

The statistical examination of multivariate time series (which we will refer to as vector time series - VTS and the individual variable time series will be referred to as *components*) is ubiquitous. Examples include stock prices, disease propagation, ecological populations, cellular recordings, chemical reactions, and many more. Common methods used to decompose these time series into (possibly) more useable forms include principal component analysis (orthogonal projections chosen to maximize variance explained) and independent component analysis (ICA) (projections chosen to minimize Gaussian-ness). In this paper, we present a method that decomposes a VTS into components that form an ordered hierarchy. The ordering of the hierarchy is one where the components at the end (we will refer to the end of the hierarchy as the "bottom") are the least causal, and those at the beginning ("top") Granger cause those below.

The method that we propose shares many features with the hierarchical decomposition (HD) method described by Repucci et al. in [22]. Specifically, both methods are based on the manipulation of a vector autoregressive model (VAR) by recombining components using

orthogonal rotation matrices. The hierarchy is determined by minimizing the amount of Granger causality in one direction. HD takes all activity at all frequencies, zero to Nyquist, into account. In other words, HD attempts to locate a global hierarchy, taking into account all of the existing processes. It is in this regard that the method we propose provides a major advance. Our method determines components with the desired hierarchical relationship within user-defined frequency bands. We call this method the Frequency Extracted Hierarchical Decomposition (FEHD).

For HD, the decomposition is accomplished by rotating the lag matrices of the VAR model so that the resulting similar matrices are as lower triangular as possible (using the sum of squares of the upper triangular elements as the quantity to be minimized). If it is possible to simultaneously triangularize the lag matrices with a single rotation matrix, then the rotated data will have the property that components near the bottom do not Granger cause activity in those above. The method we propose shares this motivation, but attempts to locate the hierarchy without explicitly trying to triangularize lag matrices, something that is generally not possible.

At the core of FEHD is the machinery of Granger causality. However, the method, along with HD, is distinct in that the goal is not to simply determine the causal relationships between time series. Rather, the goal of FEHD is to recombine the components of a VTS into one where there is (as close as possible to) a hierarchical relationship between the components.

Our primary motivation for the development of this method is the study of the electroencephalogram (EEG). The use of FEHD for the analysis of EEG is built around the assumption that the recorded EEG is the concerted output of multiple generators, interacting by superposition. We assume generators posses their own intrinsic dynamics, and that the interactions between generators can be captured by linearly combining them. In choosing a dynamic model, a multivariate autoregressive model with independent noise terms is a natural choice. The purpose of FEHD is to extract, from recorded EEG, a set of generators which interact in a hierarchical way.

In the brain, there are many different processes occurring simultaneously. In terms of EEG, these processes can often be identified with features in the power spectral density (PSD). It is unlikely that each of these,and other, processes with content at different frequency bands obey the same causal hierarchy. In fact, several recent studies demonstrate evidence for causal interactions within frequency specific activity across distributed cortical networks. In the monkey visual cortex, rhythms in the gamma band have been shown to originate in V1, and then to influence coherent gamma oscillations in V4, while feedback from V4 to V1 is not via the gamma band [6]. Additional feed forward activity has been shown to occur in the theta band (4 Hz.), while alpha and beta activity has been shown to propagate in the feedback direction [4, 5, 15, 29]. In this context, a clear shortcoming of HD is that it is incapable of isolating these individual networks. The ability to work within specific frequency bands is crucial to determining the dynamics for these processes. The advance that FEHD provides over both HD and the methods used in the studies listed above is that rather than simply determine the causality that one component exerts on another, the

method seeks to find generators - groups of channels that, collectively, drive other groups of channels within a given frequency band. Unlike other analyses of the EEG, FEHD does not treat channels as sources.

It is further well-established that the content of the EEG, specifically the various rhythms that can be identified by the power spectrum, correlates with underlying brain function [8]. For example, a consistently organized spatiotemporal pattern characterizes the wakeful EEG with a predominance of high frequency rhythms present over the frontal regions and a dominant alpha rhythm of 8-12Hz appearing over the occipital parietal EEG channels [9]. Also, modulation of specific rhythms is seen in different sensorimotor contexts; e.g. a suppression of high frequency 25Hz beta rhythms with movement [11], or suppression of alpha rhythms with shifts of attention [31]. In the context of recovery of consciousness after severe brain injuries, we have suggested that a sequence of changes in the shape of the EEG spectrum can be expected to emerge over time during the recovery process [10, 23]; preliminary studies in patients recovering consciousness after post-cardiac arrest coma support this model [21]. In specific contexts, low frequency oscillations may indicate the presence of a cortical reserve of healthy, but under excited populations of neurons [10] or evidence of thalamic deafferentation (see Section 3.2.1). Specific patterns of rhythms may arise under different pathological conditions, for example, driven by the intrinsic properties of cortical pyramidal cells [20], or originating from thalamic cells that are slowly bursting due to low threshold calcium channel mechanisms [19]. These rhythms have been shown to have distinct dynamical properties. Below we show how the unique features of FEHD analysis allow us to explore the distinctions in resting EEG dynamics between healthy subjects and patients with severe brain injuries. Though the focus of this paper is on EEG, the method that we present can be used on any VTS that approximates a Gaussian time series. Most notably is ECoG - a modality that is similar to EEG, but provides better-localized responses because the electrodes are placed directly on the cortical surface.

The goal of this paper is to describe and demonstrate the FEHD method. The structure of the paper is as follows - a brief description of linear auto-regression for multivariate time series, with an emphasis on EEG, followed by a description of the FEHD algorithm. Using synthetic time series, we demonstrate how the method resolves networks possessing different graphical structures, including a VTS that is designed so that the Granger causal hierarchy at one frequency is in the opposite direction as the Granger causal hierarchy at a second frequency. We also demonstrate the method applied to the EEG of four healthy subjects, all of whom have repeated recordings over six month time periods. We show that the method is stable, in that the results are consistent both across subjects and across time. We also show that the method is not overly sensitive to model parameters. Finally, we provide examples of how the method can test hypotheses regarding the dynamics of an EEG recorded from patient-subjects recovering from brain injury.

## 2 Methods

FEHD begins with an autoregressive model of data, $\vec{x}(t)$,

$$\vec{x}(t) = \sum_{k=1}^{L} A_k \vec{x}(t - k\Delta t) + \vec{\omega} \tag{1}$$

where

$$< \vec{\omega}(t), \vec{\omega}(t-\tau) >= \begin{cases} \Sigma & \tau = 0 \\ 0 & \tau \neq 0 \end{cases} \tag{2}$$

and where $\Sigma$ is the identity matrix and $< \cdot, \cdot >$ indicates an average over $t$. We seek generators such that the residuals do not covary. Thus, we require that $\vec{\omega}(t)$ are independent (orthogonal) white noise processes. Having orthogonal residuals allows us to calculate the Granger causality between groups of components of the VTS. We will describe a method to convert an arbitrary AR model to one with independent white noise residuals.

The FEHD method relies heavily on the machinery of Granger causality. A time series, $y(t)$, is said to *Granger cause* a second time series, $x(t)$, if knowledge about past values of $y$ allows a better prediction of $x$. Suppose we have two auto-regressive models. One of these models, $M_1$, makes predictions about $x$ using past values of $x$. The other, model $M_2$, also makes predictions about $x$, but uses the past values of both $x$ and $y$. If the residual time series corresponding to $x$ for the model $M_2$ has less variance than the analogous residual time series for model $M_1$ then the addition of $y$ into the model allowed for a better fit of $x$. We say that $y$ Granger causes $x$. In terms of the AR model:

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \sum_{k=1}^{L} A_k \begin{pmatrix} x(t-\Delta t) \\ y(t-\Delta t) \end{pmatrix} + \vec{\omega}(t),$$

if the (1,2) element of any $A_k$ is non-zero, then past values of $y(t)$ contribute to the prediction of $x(t)$ and so we say that $y(t)$ Granger causes $x(t)$.

Granger causality can also be determined in the frequency domain, over any band of frequencies (between 0 and Nyquist). Given an AR model (7), one can obtain the Fourier transform of the model time series. A calculation that needs to be performed repeatedly in FEHD is to calculate the frequency response as

$$H(f) = \left( I - \sum_{k=1}^{L} A_k e^{-i2\pi f k \Delta t} \right)^{-1}. \tag{3}$$

where $A_k$ is the $k$th lag matrix of the AR model. This is a matrix whose elements model (as a function of frequency) the transfer functions between the individual noise inputs and the individual outputs. From this, we can compute the spectral density as

$$S(f) = H(f)\Sigma H(f)^* \tag{4}$$

where the superscript $*$ indicates the conjugate-transpose, and $\Sigma$ is the covariance of the residuals, the identity matrix.

The crux of FEHD lies in determining the amount of Granger causality that a single component applies to the remaining components. To determine this, we organize the data into two parts - the first part is the first $M-1$ components, and the second part is the remaining $M$th component, i.e.

$$\vec{z}(t) = \begin{bmatrix} \vec{z_1}(t) \\ z_2(t) \end{bmatrix}$$

where $\vec{z_1} \in \mathrm{R}^{M-1}$ and $z_2 \in \mathrm{R}$. With this division, we can write the response function (3) as blocks [7, 13],

$$\begin{pmatrix} \vec{z_1}(f) \\ z_2(f) \end{pmatrix} = \begin{pmatrix} H_{z_1 z_1} & H_{z_1 z_2} \\ H_{z_2 z_1} & H_{z_2 z_2} \end{pmatrix} \begin{pmatrix} \vec{\omega}_{z_1}(f) \\ \omega_{z_2}(f) \end{pmatrix}$$

From this we can write the spectral density matrix of $\vec{z}$ as

$$S_{z_1 z_1}(f) = \vec{z_1}(f)\vec{z_1}(f)^* = H_{z_1 z_1}\Sigma_{z_1}H^*_{z_1 z_1} + H_{z_1 z_2}\Sigma_{z_2}H^*_{z_1 z_2}.$$

This is the sum of an internal part and an external part. The Granger causality at frequency $f$ - the part of $\vec{z_1}(f)$ that is explained by $z_2(f)$ - is

$$F_{z_2 \to z_1}(f) = \ln\left(\frac{|S_{z_1 z_1}(f)|}{|H_{z_1 z_1}\Sigma_{z_1}H^*_{z_1 z_1}|}\right) \tag{5}$$

where $|\cdot|$ indicates a determinant [7]. To determine the Granger Causality over a range of frequencies $[f_1, f_2]$, we define

$$G(\vec{z}(t)) = \int_{f_1}^{f_2} F_{z_2 \to \vec{z_1}}(f)df. \tag{6}$$

The goal of FEHD is identification of the rotation matrix $Q_0$ such that $G$ is minimized over all rotation matrices, i.e.

$$G(Q_0\vec{x}(t)) \leq G(Q\vec{x}(t)) \ \forall Q \text{ st. } Q^*Q = I.$$

In the following section, we describe the construction of the AR model so that it satisfies the requirements (2). Then we describe how the FEHD method utilizes (6) to determine a hierarchical VTS in a specified frequency band.

## 2.1 Linear Vector Auto-Regression

In this section we describe our process for obtaining an AR model of the form (1), (2) from the data. The linear AR model is a way to model the dependence of the current state of a VTS on its past values. Given a VTS, $\vec{x}(t_{-L}), ..., \vec{x}(t_0), \vec{x}(t_1), ..., \vec{x}(t_N) \in \mathbb{R}^M$, the AR model is

$$\vec{x}_i = \sum_{k=1}^{L} B_k\vec{x}_{i-k} + \vec{\omega_{B}}_i \tag{7}$$

where $\vec{x}_i = \vec{x}(t_i)$, and $B_k \in \mathbb{R}^{M\times M}$. The residuals, $\vec{\omega_{Bi}} = \vec{\omega}_B(t_i)$, consist of $M$ white noise processes with covariance matrix $\Sigma_B$. It is of note that $\Sigma_B$ need not be diagonal, since the residuals can be white noise sequences that covary. We will transform the model so that the residuals are orthogonal later.

We compute the lag matrices of the model, $B_k$'s, using least squares. Since clean EEG data is approximately Gaussian [26, 27, 22] the least squares method is equivalent to maximum likelihood estimation, to solving the Yule-Walker equations [28], and easier to implement.

There is debate on how to choose the proper order of an AR model, i.e, determining the number of lags to include in the model. FEHD relies on the spectral quantities that one can derive from an AR model (see previous section), and so it is reasonable to require that the model be capable of producing a spectral matrix that is representative of the data. Commonly used criterion for determining the ideal number of lags to use in an AR model include the Bayes Information Criterion and the Akaike Information Criterion (AIC) [2, 1]. In this paper, we use 40 lags for all analyses. We choose 40 because it exceeds the recommendation of the AIC for each data set tested, and the resulting AR models yield estimates of the power spectra that closely match those obtained using conventional means.

Once we have calculated an appropriate AR model, the next step is to linearly recombine the $M$ components so that the requirements on the covariance of residuals (2) are satisfied. This step is similar in purpose and function to the linear transformation of the model described in [7], though we implement it as in [22] - we choose a matrix $A_0$ such that

$$\left\langle A_0\omega_B(t), \omega_B(t-\tau)^T A_0^T \right\rangle_t = \begin{cases} \Sigma & \tau = 0 \\ 0 & \tau \neq 0 \end{cases}, \tag{8}$$

where $\Sigma$ is a diagonal $M \times M$ matrix - specifically, we choose $A_0$ such that $\Sigma$ is the identity matrix. This will allow us to obtain a model where the driving noise term has the same variance for each of the components. The matrix, $A_0$, can always be obtained from the singular value decomposition of the covariance matrix $\Sigma_B$. We let $\vec{y}(t_i) = A_0\vec{x}(t_i)$ and $\omega(t_i) = A_o\omega_B(t_i)$ for all $t_i$, and $A_k = A_0 B_k A_0^{-1}$ for all $k = 1..L$. The model

$$\vec{y}_i = \sum_{k=1}^{L} A_k \vec{y}_{i-k} + \vec{\omega}_i \tag{9}$$

will therefore have the desired properties - a sufficient number of lags to capture the interactions of the data across components and across time, and independent noisy inputs ($\vec{\omega}$) with identical variance.

## 2.2 FEHD algorithm

In this section, we describe the Frequency Extracted Hierarchical Decomposition (FEHD). The first step, taken to reduce the dimensionality of the problem and hence the computational workload, is to compute the principal components and use the first $M$ in the analysis. For

all analyses in this paper, we use 10 principal components. Across subjects, this number accounts for over 85% of the total variance. This does not preclude the possibility that a causal generator is left out of the data, but it does reduce the probability that this will happen (See Discussion). The principal components are scaled to unit variance.

A step by step description of the method is as follows (a schematic is drawn as Figure 1).

1. Begin with an AR model (7) for the $M$-component set of principal components, $\vec{X}_t$. Transform the data $\vec{Y}_t = A_0\vec{X}_t$, where $A_0$ satisfies (8) so that the condition (2) is satisfied.

2. Divide the data into two groups - the first $M-1$ components go into the first group (called $\vec{X}_1(t)$), and the last ($M$th) component is the second group (called $X_2(t)$).

3. Find the rotation matrix, $Q_0$, such that

$$G(Q_0\vec{X}(t)) \tag{10}$$

   is minimal over the set of all rotation matrices of size $M \times M$, where the function $G$ is defined as (6), the integral of causalities over the desired frequency band . One can minimize this however one chooses. For the analysis, we employ gradient descent with many initial conditions.

4. Once the optimal rotation matrix $Q_0$ is found and (10) is minimized, the component in the $M$th position will be the least causal component possible by rotating the VTS. Once this component is identified, it is marked as the least causal and removed. The resulting data set has dimension $M-1$. With this resultant, repeat the process beginning with step 1, which will identify the least causal component of this data set. Continue to repeat until only one component remains.

At the end of the procedure, the inferred hierarchy begins with the last remaining component (was never the least causal), followed by the last removed component, then the second to last component removed, and so on. In this way, a hierarchy is determined for which those components at the bottom (the first to be removed) apply minimal Granger causality to those above, resulting in a multivariate time series where the causal influence of the components at the bottom is minimal.

To perform the minimization of step 4, we use gradient descent. If there are $M$ remaining components, there are $M-1$ rotations to consider. The net rotation can be written

$$Q_{net} = Q_1 Q_2 ... Q_{M-1}$$

where each $Q_i$ is a rotation matrix that combines the $i$th and $M$th component. Each of the $Q_i$'s are the identity matrix, except

$$Q_{i(i,i)} = Q_{i(M,M)} = \cos(\theta_i)$$

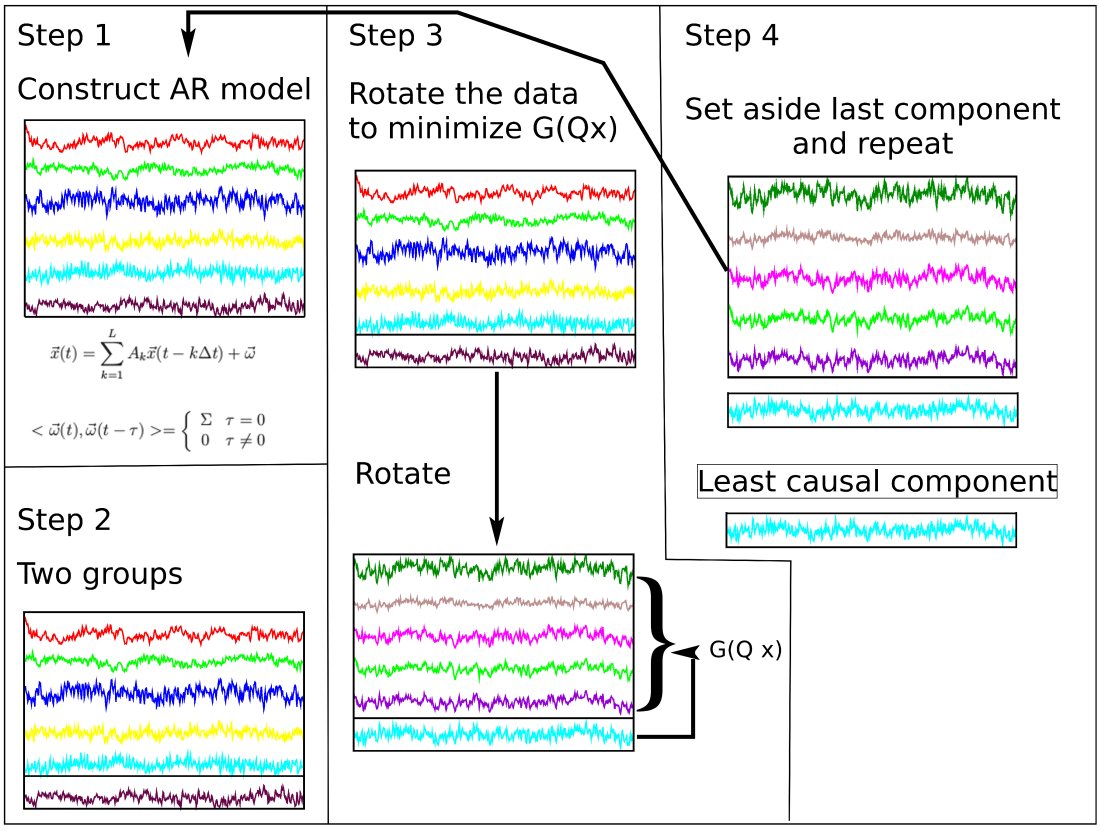


Figure 1: Schematic of the FEHD method, as an iterative loop. In step 1, and AR model is created. On the first pass, the model will be generated using the principal components of the data (Step 0, not shown). On subsequent passes, the AR model is computed using the time series determined in step 4. In step 2, the data is divided into two groups, the first $M-1$ components in one group and the last component in the second group. In step 3 the rotation matrix is found such that the Granger causality from the single component group onto the $M-1$ component group is minimized, over the desired band of frequencies. The function $G$ is defined as (10). In step 4, the least causal component is set aside, and these steps are repeated on the remaining components, until there is only one left.

$$Q_{i(M,i)} = -Q_{i(i,M)} = \sin(\theta_i).$$

We seek values of $\theta_i$ that minimize expression (10). Since gradient descent is susceptible to getting trapped at local minima, we repeat the process with 2500 random initial conditions. We iterate the gradient descent 50 times for each initial condition and take the minimum value found as the minimum.

On output from the algorithm we obtain a transformation matrix, $C$, such that $C\vec{X}(t)$ is the hierarchical time series. For each of the $k = 1..M-1$ minimizations, there will be an optimal rotation matrix $Q_{0k}$, as well as the matrix used to ortho-normalize the residuals of the AR model, $A_{0k}$ (8). After the first time around the FEHD loop, set

$$C_1 = Q_{01}A_{01}.$$

Since we remove the least causal component after identifying it, subsequent iterations of $Q_0$ and $A_0$ will shrink one dimension each step, i.e. $Q_{02}$ will be $M-1 \times M-1$, $Q_{03}$ will be $M-2 \times M-2$, etc. After each removal ($j = 2..M-1$), we set

$$C_j = \begin{bmatrix} Q_{0j}A_{0j} & 0 \\ 0 & I_{j-1} \end{bmatrix} C_{j-1}.$$

where $I_l$ is the $l \times l$ identity matrix. When the algorithm is finished - $M-1$ components have been identified and removed - we let $C = C_{M-1}$ and we have

$$\vec{H}(t) = C_{M-1}\vec{X}(t), \tag{11}$$

the hierarchical time series we desire.

We note that the method cannot be applied to signals that consist of pure sinusoids. The reason is that its starting point is a stationary VAR model. Stability of such models requires that their poles are inside the unit circle. This means that pure sinusoidal inputs cannot be fit, and Step 1 cannot be carried out. This makes sense, since with pure sinusoidal signals, it should not be possible to make inferences about causal influence.

In the Results section, to demonstrate the effectiveness of the method and analyze the output, we will use "pairwise Granger causality", the Granger causality between two components (this differs from the method itself, which is a multivariate method where the cuasality of one component on a group of others is minimized [16]). We compute this as in [7]. Specifically, let $\vec{H}(t)$ be the hierarchical VTS and denote

$$\vec{z}_{i,j}(t) = \begin{bmatrix} h_i(t) \\ h_j(t) \end{bmatrix}$$

where $h_x(t)$ is the $x$th component of $\vec{H}$. We determine the Granger causality of the $j$th component on the $i$th component as

$$\int_{f_1}^{f_2} G(\vec{z}_{i,j})df \tag{12}$$

where $G$ is defined the same as in equation (6), with $\vec{z}_1 = h_i$ and $z_2 = h_j$.

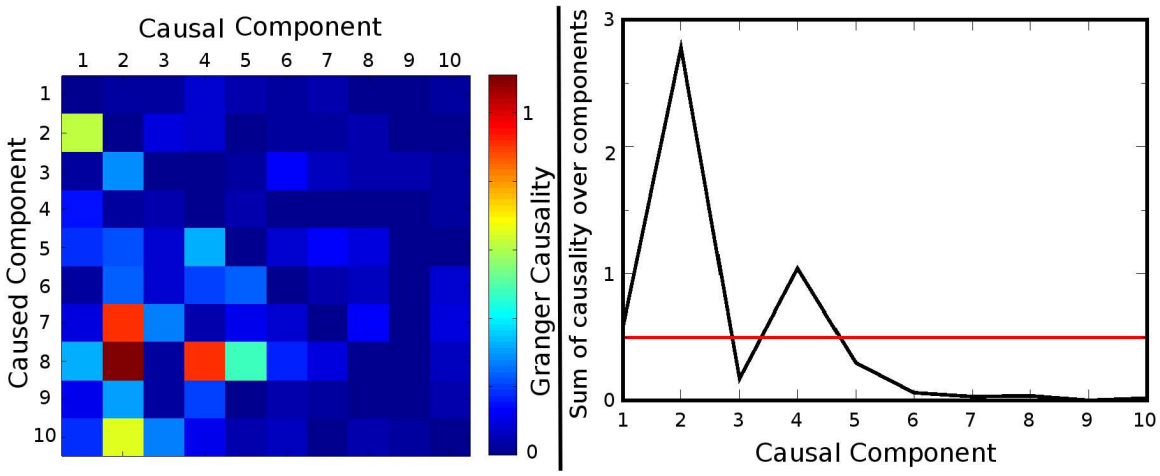


Figure 2: Criterion for a causal generator. The left panel contains the heat-map determined by the pairwise Granger causality for each pair of components. Each block represents, in color scale, the amount of Granger causality exerted from one component (designated by column) onto another (designated by row). Diagonal entries are set to zero. The downstream/upstream ratio is the ratio of the sum of squares of values below the diagonal to the the SOS above the diagonal. In the right panel, we have plotted the column sums (of squares). For each component, this quantity serves as an index of causality. The red line is the criterion chosen for a component to be considered a causal generator. In this example, we would include components 1,2, and 4 as generators.

## 2.3 Analysis of EEG

In this section we describe how we interpret and exposit the results of the FEHD method on electroencephalogram data. Once the FEHD components have been determined, we construct the pairwise Granger causality heat-map. This is an array that contains the magnitude of the Granger causality over the given frequency range for all pairs of components, in both directions. The columns correspond to the causal component and the rows correspond to the caused component. The values are plotted on a color scale where blue is zero and red is the highest value on the grid. If, for example, you wanted to determine how much Granger causality from component $N$ to component $M$, you would look at the color in the $N$th column and $M$th row (see Figure 2).

The pairwise Granger causality heat-map is used to compute the downstream/upstream ratio (D/U ratio) - the ratio of the sum of squares of the lower triangular portion of the heat-map and the sum of squares of the upper triangular portion of the heat-map. This index is identical in purpose to the ratio calculated in [22], used to determine how successful the method was in finding a multivariate time series such that most of the causal intractions are feed-forward. There is more feed-forward connectivity than feedback whenever the D/U-ratio is greater than one. A large D/U-ratio indicates that the method has determined a transformation of the data such that the resulting time series is largely feed-forward, as

desired. To show that a given D/U-ratio is significant, and not probable for a randomly connected network, we create surrogate data sets from the hierarchical components by applying the Fourier transform, randomizing the phase at each frequency, and inverting the transform. We obtain a multivariate time series where the auto-spectra of the hierarchical time series are preserved, but the phase relationships between components are randomized. Since Granger causality depends on these phase relationships, the drive between these components is randomized. We compute the pairwise Granger causality matrix for 250 surrogate time series to obtain a distribution of D/U ratios, with the mean of the logarithm of the D/U ratios approximately zero. The level of significance of the D/U ratio for a hierarchical time series is determined against its corresponding distribution.

We also use this heat-map to determine which FEHD components are causal generators and which are not. To do this, we compute the sum of squares of each column of the heat-map. Each column sum is a rough measure of causality of each corresponding component. We plot these column sums in the right panel of figure 2. From the plot, we choose a cutoff (red line) where those components that exceed the line are considered generators. We do not use a statistically rigorous rationale for the choice of this cutoff. For our purposes here, the difference between generator or not is not a sensitive one, and so we choose the cutoff by eye. Figures 11 and 12 are constructed by displaying the line spectra and topoplot for the components deemed causal generators, omitting those generators that appear to be artifactual.

The EEG signals were recorded using the Natus XLTEK FS128 (San Carlos, CA) EEG data acquisition system (sampling rate = 250 Hz., impedance $\leq$ 5 kOhms). EEG used in analysis was recorded from participants (4 healthy controls and 2 brain-injured patients) in the awake, eyes-opened resting state. Artifact-free portions of the recordings were selected via visual inspection, partitioned into 3-second segments, and montaged using a 19-channel bipolar montage. The data were then detrended (first order) segment by segment.

# 3 Results

In this section we apply the FEHD method. First we apply to artificial time series that are constructed specifically to test the capacity of the method to determine generators and hierarchical networks for different graphs. We include an example that demonstrates how FEHD can correctly determine hierarchical networks that are specific to a frequency band. Then, we apply the method to EEG data recorded from four healthy controls and two subjects who have suffered a severe brain injury.

## 3.1 Application of the method to synthetic time series.

We begin by demonstrating the FEHD method on synthetic time series. These time series are generated by filtering white noise through an AR model designed to manifest a specific graph. We use these synthetic time series to accomplish the following:

1. Demonstrate that the method recovers a purely hierarchical time series from a random combination of generators, over the entire bandwidth. For this we use the identical time series used to validate the original method in [22].

2. Determine how independent components (components that are neither driven by nor drive other components, in terms of Granger causality) are treated by the method. In particular, we observe the non-uniqueness of hierarchical time series when there are multiple non-causal components.

3. Characterize the behavior of the method on a synthetic time series that does not have hierarchical structure. For this, we filter noise through an AR model designed to correspond to a cyclical graph.

4. Demonstrate that the method is effective at isolating frequency band-specific hierarchical time series.

5. Provide an example that will illustrate and demonstrate the process of choosing the number of principal components to use in the analysis, as well as show that the method is resilient to shared instantaneous noise.

The synthetic time series that we use correspond to the following graphs:

1. A purely hierarchical causal chain,

2. a pair of independent generators that both influence a third component,

3. a pair of independent networks,

4. a cyclic network,

5. a network where the hierarchy at one frequency is opposite the hierarchy at another frequency.

6. A network with a single causal generator driving multiple caused nodes. We execute the method for two cases, one where there is no shared noise between components, and one where there is substantial shared noise.

Each of the non-frequency dependent time series (1-4) are constructed the same way. We construct an auto-regressive model with lag matrices that correspond to the graph in question, i.e. if component 1 Granger causes component 2 then at least one lag matrix will have a non-zero value for its $(2, 1)$ entry. For each example we use two lag matrices. We use the AR model consisting of these lag matrices to generate a synthetic time series. To do this we choose two values, the initial conditions $\vec{x}_{-1}$ and $\vec{x}_0$, and simulate a vector time series of independent white noise, $\vec{\omega}$. We iterate the AR model (9) to obtain a synthetic time series, $\vec{x}(t)$. Next, we create sixteen random combinations of the components of $\vec{x}$ (as an example, see the upper right panel of figure 3), and compute the principal components of these. We keep the first three (example time series 1,2,4) or four (example time series

3) (lower left panel of Figure 3) principal components. The method is carried out on these principal components, and the hierarchical components are computed (lower right panel of figure 3).

It is important to note that because the synthetic time series have rank precisely 3 (or 4), choosing the number of principal components is trivial. For experimental data, choosing the number of principal components is not trivial.

The first three examples that we provide are hierarchical networks that can be modeled by an AR model with lower triangular lag matrices. The first of these is a *purely hierarchical network*. In such a network, each component Granger causes each of the components beneath it in the hierarchy (the downstream graph is dense), and no component Granger causes any component above. The generating AR model consists of lag matrices that are lower triangular, and no component will be independent (neither Granger caused nor Granger causing).

The synthetic time series used in [22] possesses a purely hierarchical structure. We take advantage of this and use the same synthetic data to test FEHD over the entire bandwidth. Since there is a purely hierarchical structure, HD performs very well, and provides a ground truth for comparison to the results of FEHD. The results in figure 3 show that FEHD successfully recovers the original generators from the principal components of the random combinations.

If the entire bandwidth is considered, and the network can be rotated into a purely hierarchical form, the methods HD and FEHD will perform identically, since they are seeking the same thing. The proof is based on the convergence relation given in [7]:

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} F_{Y\to X}(\lambda)d\lambda \leq \hat{F}_{Y\to X},$$

where $\hat{F}_{Y\to X}$ is the Granger causality measure in the time domain and is trivially zero for when $X$ is above $Y$ and all lag matrices are lower triangular. So, if all upper triangular entries of all of the lag matrices are zero (the HD solution) then the causality in the frequency domain will also be zero.

The hierarchical decomposition for a purely hierarchical network is unique. If the network is not purely hierarchical and there are multiple components that are non-causal, then the time series determined by FEHD will not be unique. To demonstrate this, we consider a pair of synthetic time series designed to have multiple non-causal components.

We consider how the method performs on a multivariate time series containing two components, independent of one another, mutually driving a third component (Figure 4). It can be seen from this figure that the method does not reproduce the original generators - the top two components determined by FEHD are not the same as the first two original generators. However, the method does return a hierarchical time series. FEHD successfully identifies the non-causal component (the third original component is the same as the third hierarchical component). This non-causal component is the only component that is Granger caused so when it is removed from the analysis the remaining components are both non-causal. Hence, any linear combination of these components will also be non-causal. At each step FEHD

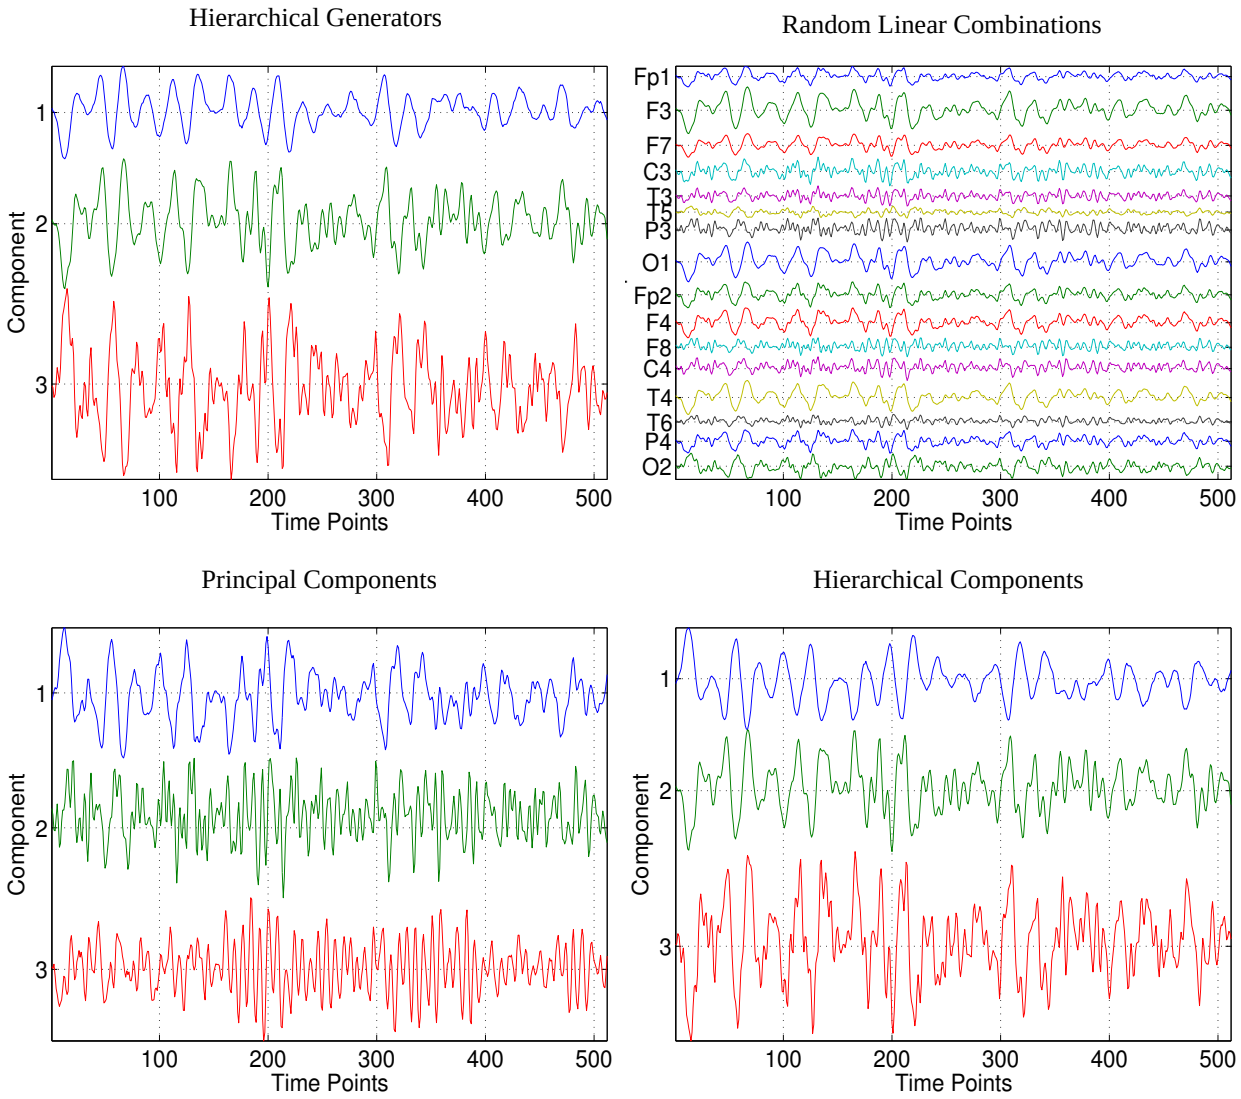


Figure 3: Application of FEHD on hierarchically generated time series. The upper left panel shows the hierarchical generators. In the upper right panel are random linear combinations of the generators. The lower panels are the first three principal components (left) and the Hierarchical components (right).

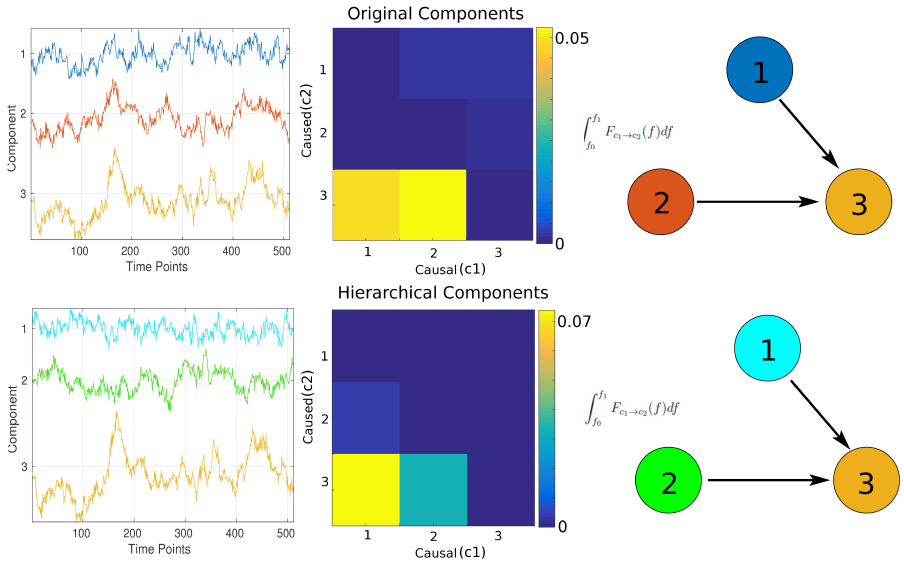


Figure 4: FEHD on network with two sources driving a third component. The left panel shows the time series of the original components (top), and the hierarchical components (bottom) as determined by FEHD. The middle column shows, as a color scale, the causal influence components exert on one another for both the original and hierarchical components. The third column is a cartoon of the graph for the original and hierarchical components. The dots on the right are color coded and numbered to match the time series on the right and the nodes in the graph on the right.

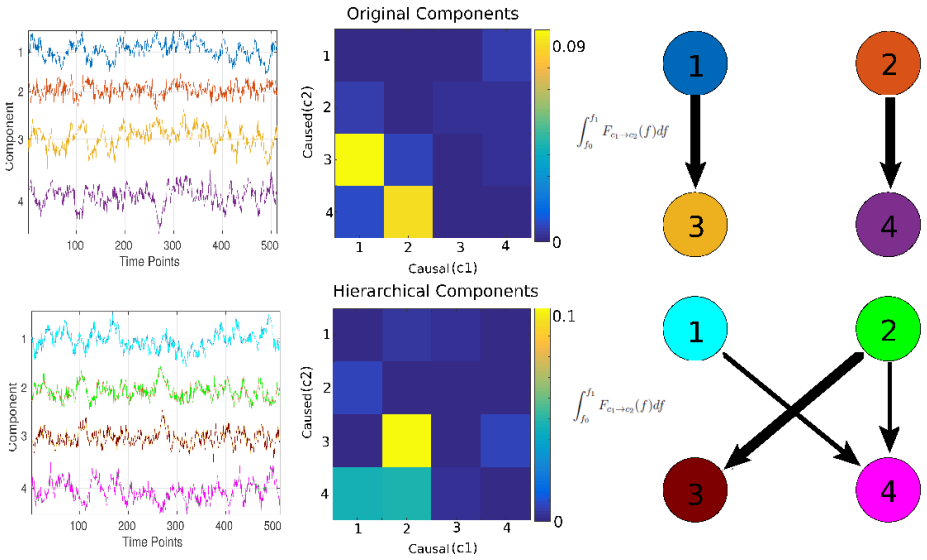


Figure 5: FEHD on network with two independent sub-networks. The layout is the same as in Figure 4, but with line thicknesses reflecting the causality from one component to another (Thicker is greater).

searches for the least causal linear combination. Any linear combination of non-causal components is non-causal, and so any linear combination will yield this minimum (zero). Thus, the determined hierarchical time series is not unique. The specific linear combination that the method arrives at depends on the noise in the coefficients of the model. Even though FEHD did not reproduce the original components, it did produce a valid hierarchical VTS, as intended.

To further illustrate the non-uniqueness of the determined hierarchical time series, we consider a synthetic VTS where there are two independent networks (Figure 5). Each of these networks consist of two components, one that Granger causes the other. Since neither of the "receiving" components exert any Granger causality, any combination of them will be non-causal. This is the same situation as we saw in the previous example, and again FEHD does not reproduce the original generators. However, it does determine a time series that is hierarchical, as desired.

The next example is a cyclic network, the first of our examples that is not hierarchical to begin with. The goal of FEHD is to minimize feedback, and so it is interesting to see what it does on a network that has a strong feedback. Figure 6 shows the results when the method is applied to such a network. The original time series is generated exactly the same as above, except rather than a hierarchical structure, a feedback is included in the AR model. Since all three original components are Granger causal, any linear combination of them will also be Granger causal. FEHD locates the linear combination that exerts the least amount of Granger causality. This minimum will not be zero, and so there will be some feedback. Thus, FEHD is not going to locate a strictly hierarchical VTS, where all Granger causality

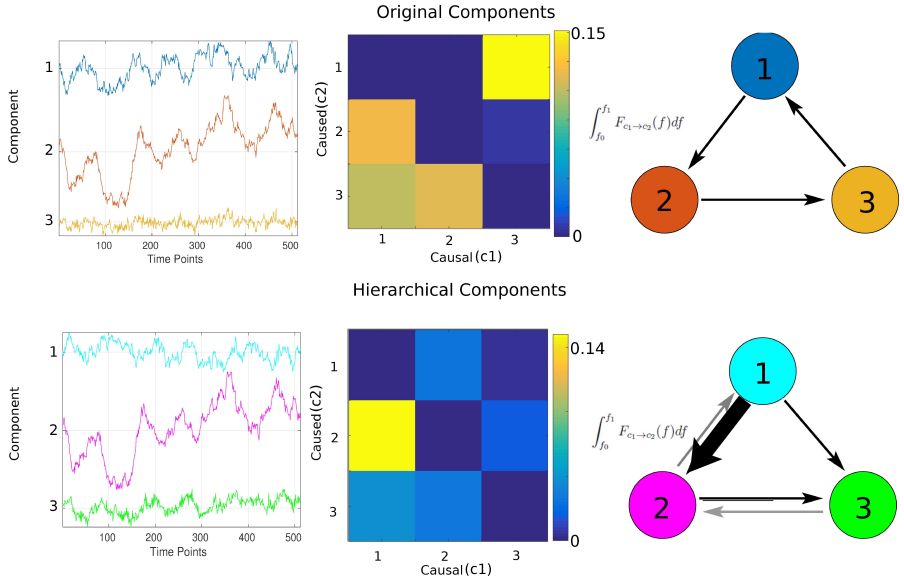


Figure 6: FEHD on a network with feedback. The details are the same as in Figures 4 and 5. In the lower panel, the gray arrows are the feedback connections, which are present but much smaller than the feed forward connection from component 1 to component 2. We note here that although the connection matrix shows a connection from component 1 to component 3 for the original components, because we designed this time series, we know this connection is not there. The appearance of this connection results from the well known inability of Granger Causality to determine conditional causality.

is downward. Rather, it provides as close to hierarchical VTS as possible by minimizing the feedback for each component. One can see from figure 6 that there are small feedbacks remaining after FEHD (lighter blue squares in the upper diagonal).

We note here, for this cyclic network, that the D/U-ratio defined in section 2.3, will be large. Generalizing, consider an N-component VTS organized such that each component causes the next with GC of 1, with the exception of the Nth component which causes the first with a GC of 1. This will have a D/U-ratio of N-1, which is much larger than one would expect from a random network, and thus significant. This example shows that the D/U-ratio is not useful as a classifier of VTS with zero feedback, but rather an indicator of the amount of feed-forward causality relative to the amount of feedback.

The next example we provide illustrates the major advance of FEHD - the ability to determine hierarchical time series within arbitrary frequency bands. To demonstrate this, we construct a multivariate time series such that at frequency $\omega_1 = 0.2$ (where $\Delta t = 1$ so that the Nyquist frequency is 0.5) the first component Granger causes the other two, and at frequency $\omega_2 = 0.4$ the third component Granger causes the first and second components. The auto spectra, and the pairwise Granger causality indices for such a time series are shown in Figure 7.

To create a VTS series that satisfies the above description, we create two AR models. First, we begin by fitting a damped oscillation with frequency 0.2 to a 2 lag AR model. Then we proceed as in the previous examples, creating an AR model that is hierarchical with the damped oscillatory component as the generator. We also do this for a damped oscillator with frequency 0.4. We filter Gaussian white noise through each of these models, and then add the resulting time series (flipping the second so that the causality is upward). We repeat the process of previous examples by randomly combining into 16 components, taking the first three principal components, and applying the FEHD method separately on the frequency ranges $[0.15, 0.25]$ and $[0.35, 0.45]$ using 52 frequency points to approximate the integral (10) (chosen to be consistent with the FFT).

The left panel of Figure 8 contains the original generators. In the center panel, the hierarchical decomposition is shown for the frequency band $[0.15, 0.25]$. The method effectively computes and orders the generators for this frequency band (2 and 3 are switched, but they are independent and so they can appear in either order - see Figure 7). On the right is the FEHD for frequencies $[0.35, 0.45]$. The method successfully orders the components here as well, with generator 3 moving to the top, and 1 moving to the bottom. Here, generator 2 also Granger causes 1, so there is no ambiguity.

The final example we provide is a synthetic time series that corresponds to a type of network that is prevalent in the brain, a single generator driving behavior at multiple (eight) other locations. We do this analysis twice. First, we generate time series so that there is no shared noise when fit to an AR model (residuals are close to orthogonal). When we do the analysis a second time, we generate the time series so that the residuals covary substantially, modeling a shared noisy input. As in the previous examples, we randomly combine the original components to have a multivariate time series with sixteen components.

Figure 9 shows the results for each of these analyses. In each case, the components

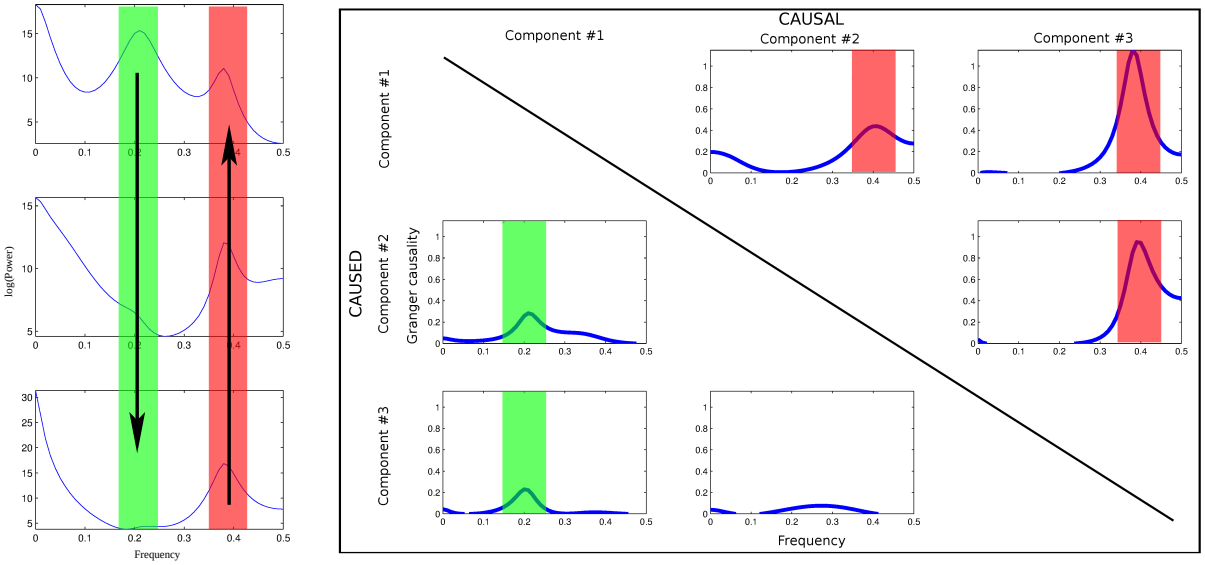


Figure 7: The causality of the example multivariate time series frequency by frequency. The left panel shows the hierarchy: at frequency $\omega_1 = 0.2$ the top component Granger causes the other two, and at frequency $\omega_2 = 0.4$ the third component Granger causes the other two. The right panel shows the pairwise Granger causality at all frequencies. Here, the $i$th column represents the causal index (expression (5)) of the $i$th component at the given frequency upon the component corresponding to the given row.

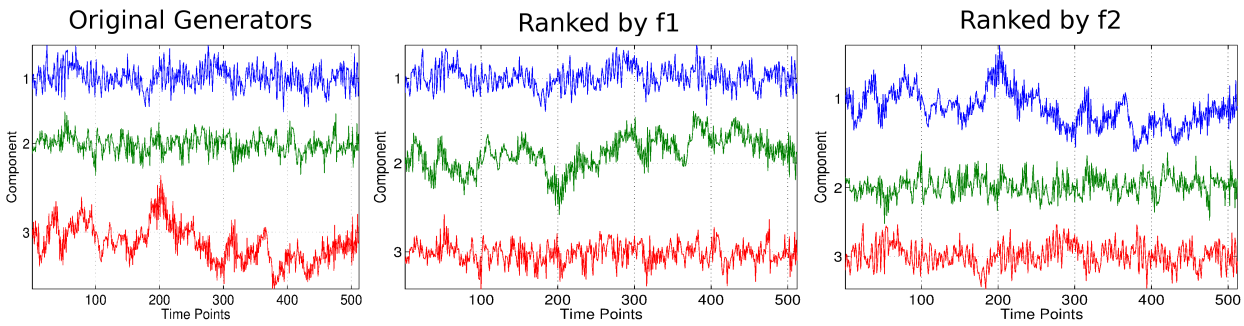


Figure 8: Hierarchical Decomposition by frequency. The left figure is the original generators, chosen so that at $\omega_1 = 0.2$ the first generator causes the other two and at $\omega_2 = 0.4$ , the third generator causes the other two. The middle figure shows the FEHD at $\omega_1$ (integral over $[0.15, 0.25]$) where the first generator on the left is most causal, and the right figure shows the FEHD at $\omega_2$ (integral over $[0.35, 0.45]$) where the third generator on the left is most causal.

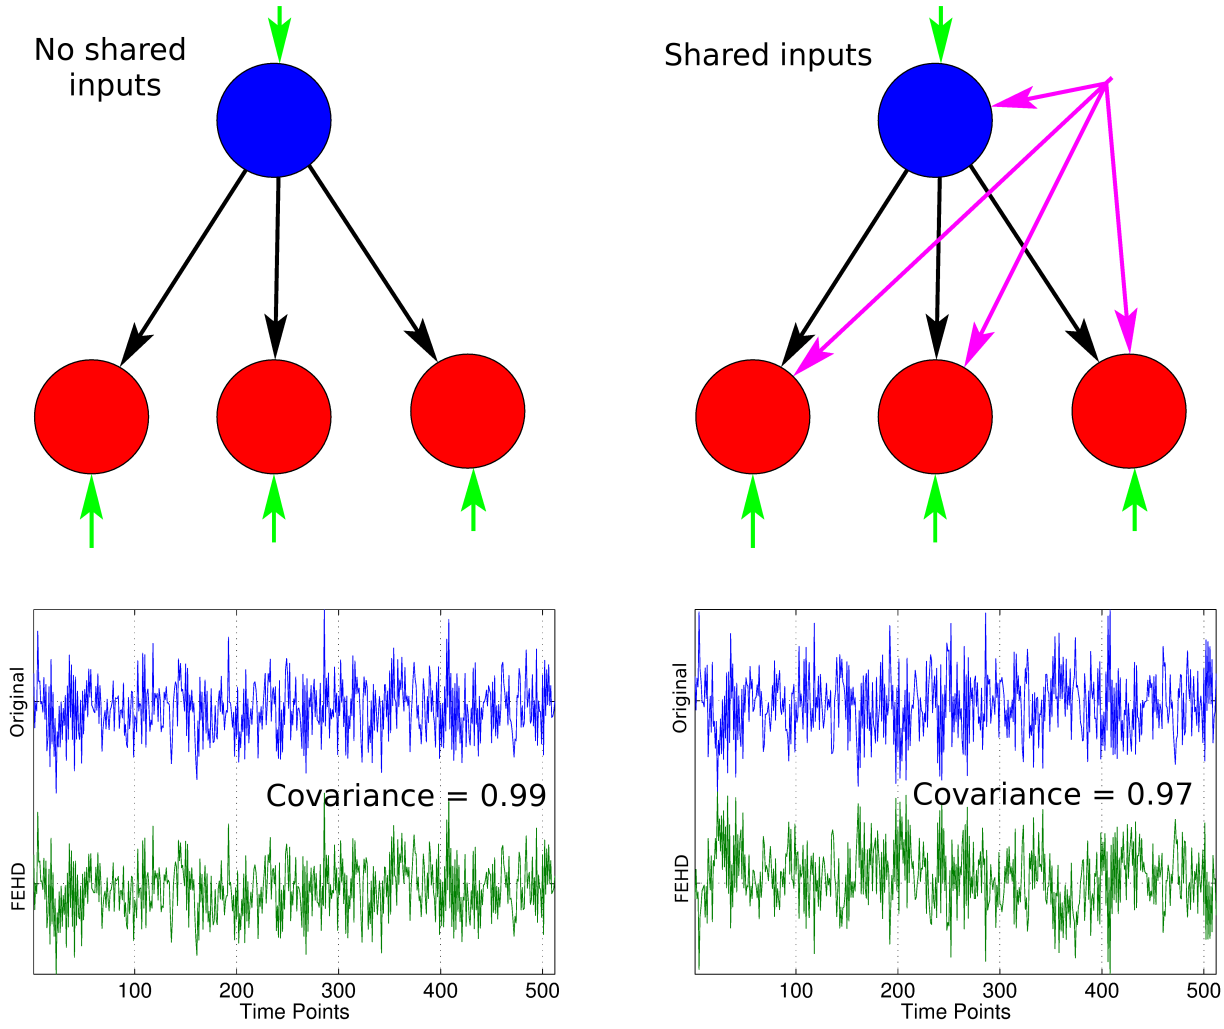


Figure 9: FEHD successfully determines the generator for an all-to-one network, where a single component causes activity in all other components. The upper part of the figure shows cartoons describing the networks. Each has the same structure - a single causal generator that influences behavior in all of the other nodes (black arrows). For the actual analysis done, there is one generator and eight recipients. Each node in each network also has its own, independent input. The difference is that for the network on the right, each node also receives a shared input (magenta arrows). The lower part of the figure shows the time series of the original generator as well as the component determined to be at the top of the hierarchy by FEHD. In each case, the resemblance between the original generator time series, and the time series identified by FEHD is very strong - the covariance is greater than 0.97 (out of 1) for each.

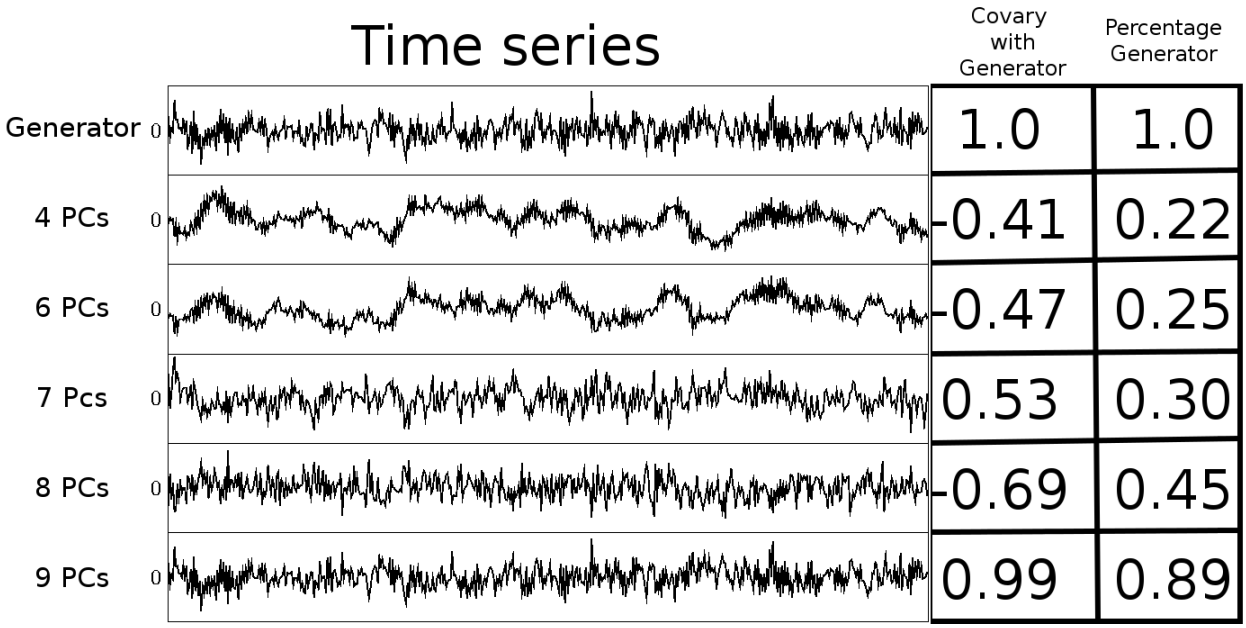


Figure 10: The performance of FEHD when the generator exhibits low variance. In this figure, we show how FEHD resolves the network described in figure REF using a range of principal components. The columns are as follows: The number of PCs used in the analysis, the time series of the first FEHD component (the first row is the true generator), the plotted time series, the covariance with the true generator, and the percentage of the identified first hierarchical component made up of the true generator.

determined by FEHD to be at the top of the hierarchy are a very close match to the original generator. The rest of the FEHD components are not unique, since they are just random combinations of non-causal components, and so we do not show them.

The above time series was constructed to have another property - the time series corresponding to the node that drives all of the other nodes has relatively small variance. When determining principal components, this results in most of this generator being contained in smaller principal components. In figure 10, we show the results for a range of principal components used (4-9). We measure the closeness of the identified generator to the true generators using covariance, as well as the relative contribution of the true generator to the identified one, in terms of weight. Since the generator is not largely present in the first few principal components, FEHD will not be able to locate it until sufficiently small principal components are used. Figure 10 shows that as more principal components are added the isolation of the true generator becomes increasingly precise.

In addition to the capacity to derive the hierarchical decomposition for user-specified frequency bands, the method provides another, more subtle, improvement over HD. If a multivariate time series has one or more independent components (neither Granger caused nor causal), then the decomposition obtained by HD is not unique. That is, an independent

component can appear anywhere in the hierarchy, since it would not contribute anything to the upper triangular part of any of the lag matrices, regardless of where it was. To contrast, FEHD would place all such components at the bottom of the hierarchy, since they exert no causal influence. If there are multiple such components, neither method is unique, however FEHD will place the components at the bottom of the hierarchy whereas HD would place them randomly throughout.

## 3.2 Application of the method to the EEG

The primary advance that we make with FEHD, over HD, is that FEHD is able to determine hierarchical time series over any band of frequencies. To demonstrate and justify the method for use with EEG, we perform the FEHD analysis on the EEG recorded from four healthy subjects, each of whom have been recorded at two time points (figures 11 and 12). Additionally, we apply FEHD to the EEG of two subjects who have suffered a severe brain injury.

We use the data from these six subjects to demonstrate the following:

1. Functionality - simply put, we want to demonstrate that the method does what we claim it does. We show that FEHD successfully forms components of a highly hierarchical time series, where components at the bottom exert a minimal amount of Granger causality on those above. Moreover, the method is capable of determining distinct causal networks for different frequency ranges. The causal structure of a hierarchical network uncovered in one frequency band will not generally apply to another frequency band.

2. Stability - Small perturbations to the data must result in proportionately small changes in the results. We show that the method is able to extract similar generators for EEG recorded at time points 6 months apart. This test-retest reliability is strong evidence of a stable method.

3. Biological usefulness, specifically for EEG - We aim to show that the method has the capacity to provide information about the brain using EEG data. We do this by viewing the results in light of some "ground truths" of resting state activity in the healthy brain. We show how the method facilitates a hypothesis regarding the mechanism of dysfunction in the global dynamics of an injured brain. We also provide a case where removing the pre-image of pathological FEHD components from the data allows isolation of a previously hidden (and indicative of healthy thalamocortical loops) rhythm.

We compute the downstream/upstream Granger causality ratios to determine how hierarchical the components determined by FEHD are. In terms of the pairwise Granger causality heat-maps, we divide the sum of squares of the lower triangular entries by the sum of squares of the upper triangular entries. A large ratio indicates that the method successfully ordered the components so that components apply the minimal amount of Granger causality to those

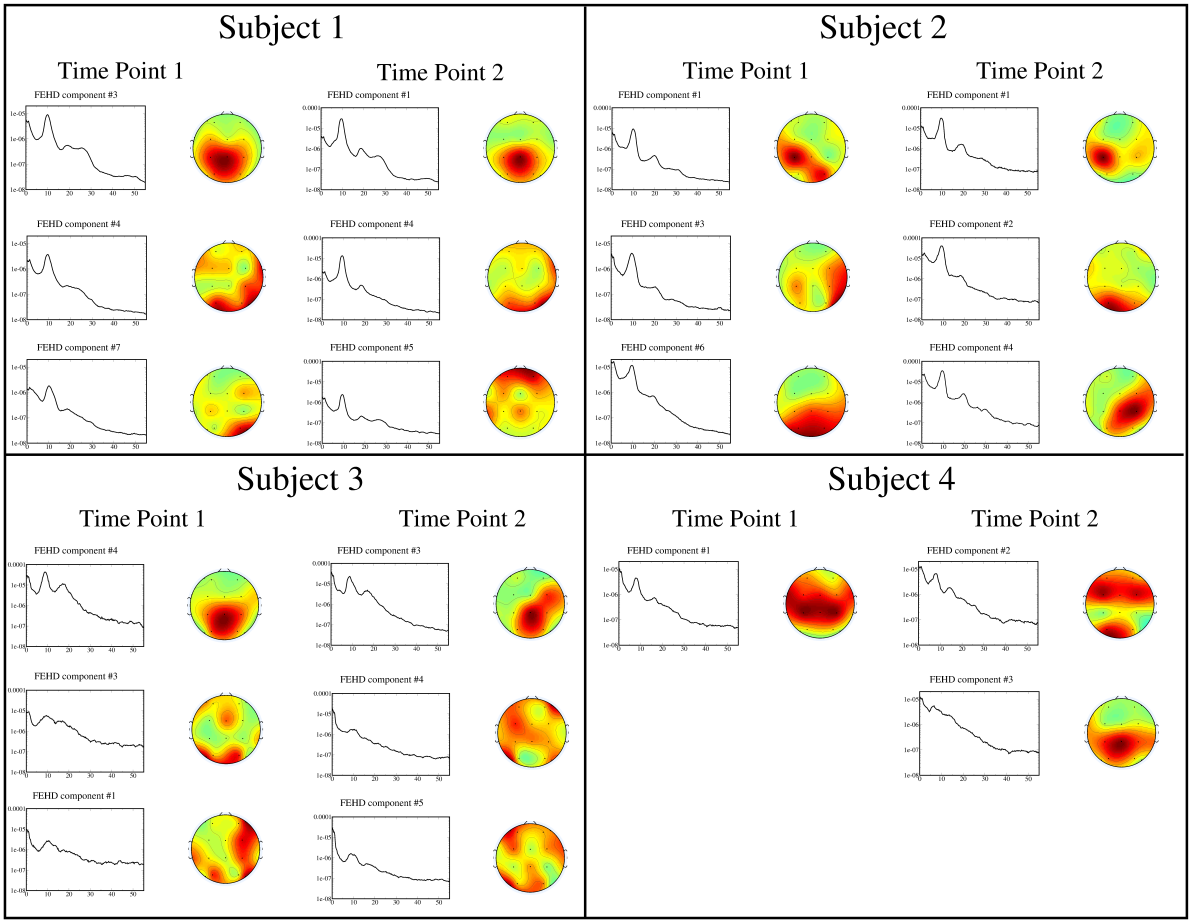


Figure 11: FEHD analysis over the $\alpha$ band (8-12 Hz) of the EEG of four healthy subjects. Each of the four large panels represents one subject, tested at two time points six months apart. Plotted are the line spectra and topographic map of the location on the scalp of the generators determined by FEHD, as described in the methods section.The generators are shown in the order determined by FEHD, and so the hierarchy is from top to bottom, with the exception of the first time point of subject 3. In this case, the generators do not exert a large amount of drive on one another, and so clearly identifying similarities between the analysis at the two time points took precedence over the exact output of FEHD.

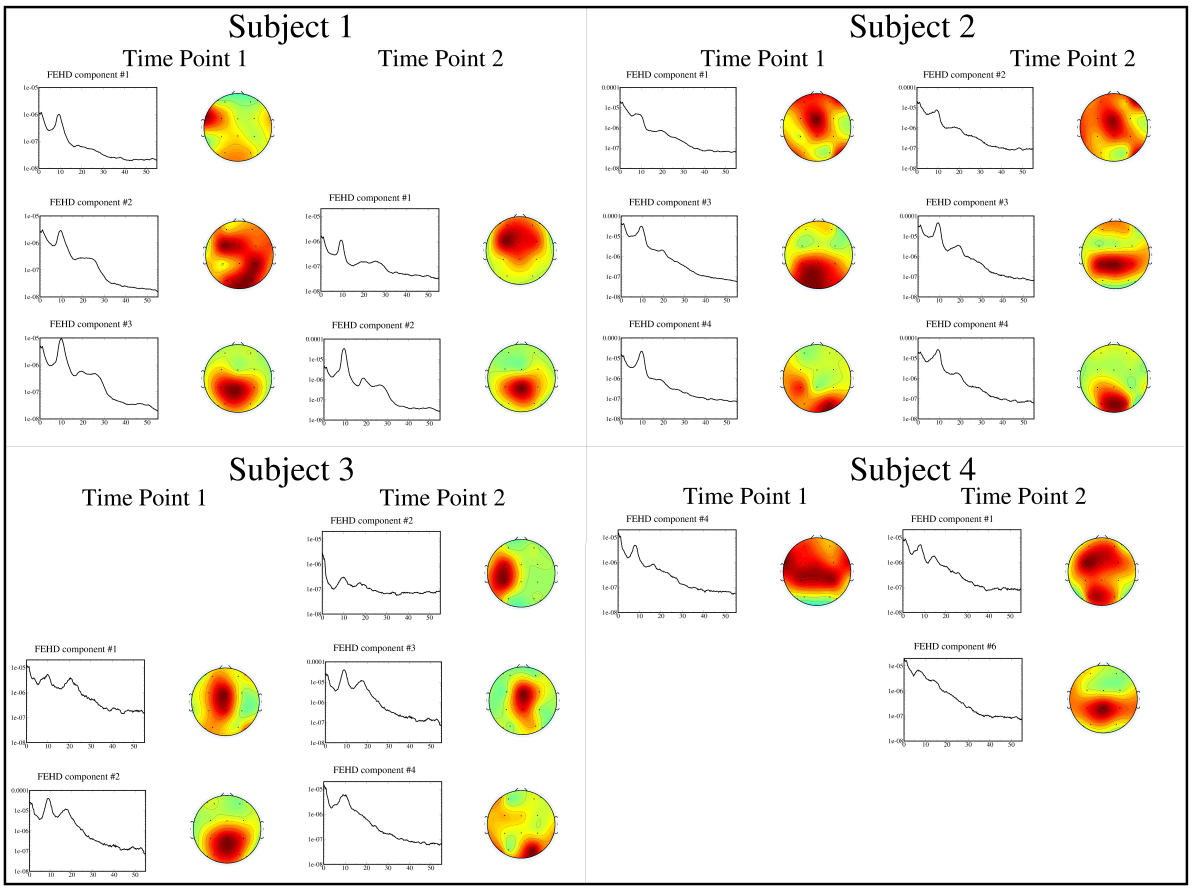


Figure 12: FEHD analysis over the $\beta$ band (the specific range varies) of the EEG of four healthy subjects. Details are identical to that of figure 11. Frequency ranges for subjects : Subject 1, 20-30 Hz; Subject 2, 18-22 Hz; Subject 3, 20-22 Hz; Subject 4, 14-18 Hz.

**First Time Point**

| | | Subject 1 FEHD | | Subject 2 FEHD | | Subject 3 FEHD | | Subject 4 FEHD | |
|---|---|---|---|---|---|---|---|---|---|
| | | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ |
| PGC | $\alpha$ | 7.0 | 1.94 | 28.7 | 7.17 | 27.8 | 3.81 | 31.2 | 14.7 |
| | $\beta$ | 1.19 | 5.10 | 3.37 | 14.17 | 3.54 | 28.6 | 26.1 | 81.9 |

**Second Time Point**

| | | Subject 1 FEHD | | Subject 2 FEHD | | Subject 3 FEHD | | Subject 4 FEHD | |
|---|---|---|---|---|---|---|---|---|---|
| | | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ | $\alpha$ | $\beta$ |
| PGC | $\alpha$ | 28.1 | 1.53 | 36.9 | 8.53 | 26.4 | 3.34 | 130 | 45.0 |
| | $\beta$ | 5.29 | 3.51 | 1.09 | 22.5 | 3.19 | 12.51 | 45.7 | 75.1 |

Table 1: Table containing the downstream/upstream ratios of the pairwise Granger causality for each of the four healthy subjects. The columns correspond to the FEHD analysis done, by subject and frequency range ($\alpha$ or $\beta$). The rows correspond to the frequency range over which the pairwise Granger causality (PGC) is computed. The yellow cells contain the ratios for when the FEHD analysis and pairwise Granger causality are computed over the same frequency band. The uncolored cells contain the ratios for when the FEHD analysis and pairwise Granger causality are computed over different frequency ranges. All of the ratios in yellow cells are significant ($p < 0.01$).

above them in the hierarchy. The results of this, for four healthy subjects, are shown in Table 1 (yellow cells). All of the D/U ratios in these cells are significant ($p < 0.01$). In all cases, the method was able to determine a hierarchy where there is much more Granger causality in one direction (downstream) than the other (upstream). Moreover, the method determines different causal networks for different frequency ranges. We compute the hierarchical components for one frequency range, and check the pairwise Granger causality over another. For example, we compute the FEHD components over the $\alpha$-range, but check the pairwise Granger causality over the $\beta$-range. The results are shown in Table 1 (uncolored cells). For each subject in each visit, the D/U ratio is higher when the pairwise Granger causality is computed over the same frequency range for which the FEHD components were determined. This result shows that the method is capable of determining hierarchical networks in different frequency bands, and that the networks are unique to the frequency range analyzed.

Figures 11 and 12 demonstrate a robust test-retest reliability of the method. For data recorded at time points months apart, FEHD extracts a set of generators that exhibit remarkable similarities within subject 1-3, both in terms of spectral content and location on the scalp. Specifically, in the $\alpha$-range (Figure 11), the first two generators determined for Subject 1, the first for Subject 2, and the first two for Subject 3 are very similar across time points. This reliability can also be inferred for analysis in the $\beta$-range (see Figure 12). The generators for Subject 1 possess the same frontal location on the head, though for time

| | Subject 1 | | Subject 2 | | Subject 3 | | Subject 4 | |
|---|---|---|---|---|---|---|---|---|
| | $\alpha_1 D_2$ | $\alpha_2 D_1$ | $\alpha_1 D_2$ | $\alpha_2 D_1$ | $\alpha_1 D_2$ | $\alpha_2 D_1$ | $\alpha_1 D_2$ | $\alpha_2 D_1$ |
| D/U ratio $\alpha$ | 6.39 | 7.19 | 3.26 | 15.91 | 3.32 | 38.0 | 3.97 | 4.96 |

Table 2: Table of downstream/upstream ratios when $\alpha$-range transformations are applied across visits. For each subject, $\alpha_i$, $i = 1, 2$ represents the $\alpha$-range FEHD transformation matrix from time point $i$, and $D_j$ is the data from time point $j$. All of the D/U-ratios listed are significant, with $p < 0.01$ for columns with green headers and $p < 0.05$ for columns with blue headers.

point 1 the method found an additional generator and did not remove the contribution from occipital channels. The first generators determined for Subject 2 are remarkably similar. The analysis of Subject 3 is similar to that of subject 1, in that an additional generator was determined by the analysis for time point 2. For each time point, a frontal-central generator was determined. For both the $\alpha$-range and the $\beta$-range, the analysis of Subject 4 produced the worst results, as the consistency between time points is not as clear as in the other subjects. Moreover, the method extracts the same generators for each frequency band. We suspect that these differences are in the data since the method determines the most hierarchical networks for this subject, in terms of downstream/upstream ratios, of the four subjects presented.

By visual inspection, we have seen similarities in the generators. We also want to know how similar the determined networks are between time points. To do this, we once again rely on the downstream/upstream ratio of Granger causality. In what we refer to as cross-transforming, we apply the transformations obtained by FEHD from the data at one of the time points to the data recorded at the other time point. For example, we apply FEHD to the data recorded at time point 1, obtaining a matrix specific to the $\alpha$ range. We then apply this linear transformation to the data recorded at time point 2. We will refer to this component set as $\alpha_1 D_2$. Then, we carry out the pairwise Granger causality analysis on this set of components, for both the $\alpha$ and $\beta$ ranges. If one transformation results in a significantly hierarchical network when applied to either data set, then we infer that the generators determined by the method are similarly causal across time points.

The results of cross-transforming over the $\alpha$ frequency range are shown in Table 2. From this table, it can be seen that all of the cross transformed multivariate time series are hierarchical ( D/U ratio values greater than 1). Moreover, cross-transformed time series are significant (Blue headers indicate $p < 0.05$ significance and green header indicate $p < 0.01$ significance, see section 2.3 for significance testing).

We repeat the cross-transformation analysis for the $\beta$-range. (Table 3). As above, each of the cross transformed time series are hierarchical. The values are often significant as well, as denoted by the blue and green headers.

In terms of biological inference, the generators for the $\alpha$-range are mostly occipital. On the other hand, the generators for higher frequency activity ($\beta$) are largely frontal. For healthy brains, the occipital nature of the $\alpha$-rhythm and the frontal origins of $\beta$ activity are

| | Subject 1 | | Subject 2 | | Subject 3 | | Subject 4 | |
|---|---|---|---|---|---|---|---|---|
| | $\beta_1 D_2$ | $\beta_2 D_1$ | $\beta_1 D_2$ | $\beta_2 D_1$ | $\beta_1 D_2$ | $\beta_2 D_1$ | $\beta_1 D_2$ | $\beta_2 D_1$ |
| D/U ratio $\beta$ | 1.68 | 1.97 | 3.79 | 5.40 | 2.61 | 2.47 | 5.73 | 3.16 |

Table 3: Table of downstream/upstream ratios when $\beta$-range transformations are applied across visits. The color of the header indicates significance: green corresponds to $p < 0.01$, blue to $p < 0.05$ and red to $p < 0.1$.

well known. Because the method is able to capture this fundamental property of healthy brains, we have reason to believe that this method is capable of lending insight regarding the global dynamics of the injured and recovering brain. In the next two subsections, we apply these methods to the EEG recorded from two patient subjects recovering from severe brain injuries.

### 3.2.1 FEHD as a tool to study recovery from severe brain injury

We provide two examples that illustrate how FEHD can allow for hypothesis development and testing in human subjects recovering from severe brain injuries. This area of research is one likely to benefit from the availability of FEHD analysis since the EEG in the severely brain injured subject typically combines a mix of signals from recovering brain networks and those of pathological origin. The latter arise from transient functional deafferentation or, in many cases, permanent structural injuries [24, 23].

Because recovery of consciousness is linked to measures of global brain dynamics, the ability to assess the presence of specific causal dynamic processes is important. These processes can either suggest restoration of function or the isolation of pathological rhythms. Here we provide two relevant examples. The first example is the evaluation of an abnormal dynamical element detectable across the scalp EEG of an injured subject. This element arises coincident with recovery of brain function (along with behavioral improvement). We show how FEHD analysis supports a biological interpretation of the origin of this EEG feature. In a second example we illustrate, using data obtained from a different patient subject, how the use of FEHD to isolate pathological dynamics can allow for further evaluation of recovery of brain function. Removal of the pathological component(s), identified by FEHD, and reanalysis of what remains allow us to uncover and isolate a normal, dominant, occipital $\alpha$-rhythm originating in a more preserved cerebral hemisphere. Both examples demonstrate the utility of the features of FEHD analysis over those of the full-bandwidth HD method.

#### Isolation of pathological dynamics: thalamocortical dysrhythmia

In this example, we evaluate findings from a patient subject who spontaneously recovered from the minimally conscious state to a level of full, interactive spoken communication with persistent disorientation (see [17] for topical review) over a 10 year time period. In this subject, an initial EEG assessment showed a minor feature in the power spectrum

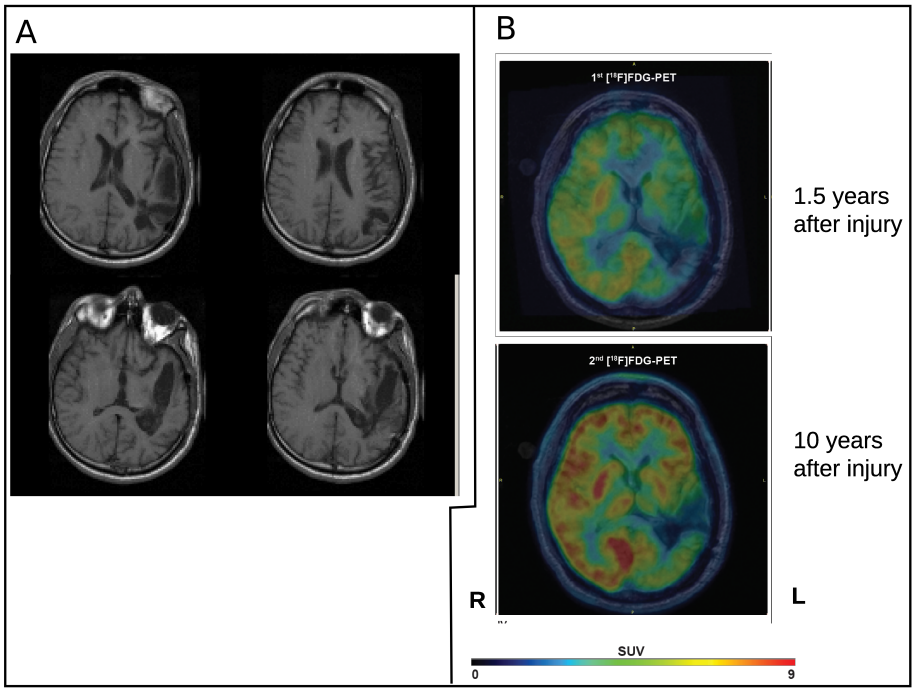


Figure 13: Structural MRI and PET for patient subject IN388R. Panel A shows four representative magnetic resonance images revealing an asymmetric injury of the left hemisphere. Panel B shows changes in regional cerebral metabolic rates measured by fluorodeoxyglucose positron emission tomography (18FDG-PET). A global increase in cerebral metabolic rate is seen between the first measurement 1.5 years post-injury and the second time point at 10 years post-injury (coincident with the EEG measurements shown in Figure 14).

around 4-8 Hz that, after 10 years, became a prominent feature on the majority of recorded channels. This emerged along with a new, very prominent broad peak in the power spectrum in the range of 25 to 35 Hz (see Figure 14). The emergence of the combined 4-8Hz and 25-35Hz features of the resting EEG correlated with broad increases in fluorodeoxyglucose-positron emission tomography (18FDG-PET) measured cerebral metabolism across both hemispheres with a marked asymmetry of near normal metabolic rates across the right hemisphere and relatively lower resting metabolism in the left hemisphere (Figure 13). Both the EEG and 18FDG-PET changes coincided with the patient's recovery of spoken language and interactive communication between 1.5 and 10 years after severe brain injury.

The combination of these features in the power spectrum (4-8 Hz for the low frequency peak, and 25-35 Hz for the high frequency peak) are consistent with a physiological phenomenon reported in the EEG of several neurological disorders known as thalamocortical dysrhythmia, TCD [14, 19]. TCD is described as an EEG or magnetoencephalography (MEG) phenomena arising from a deafferented thalamus that produces low threshold T type calcium channel bursting during wakeful periods as result of relative hyperpolarization

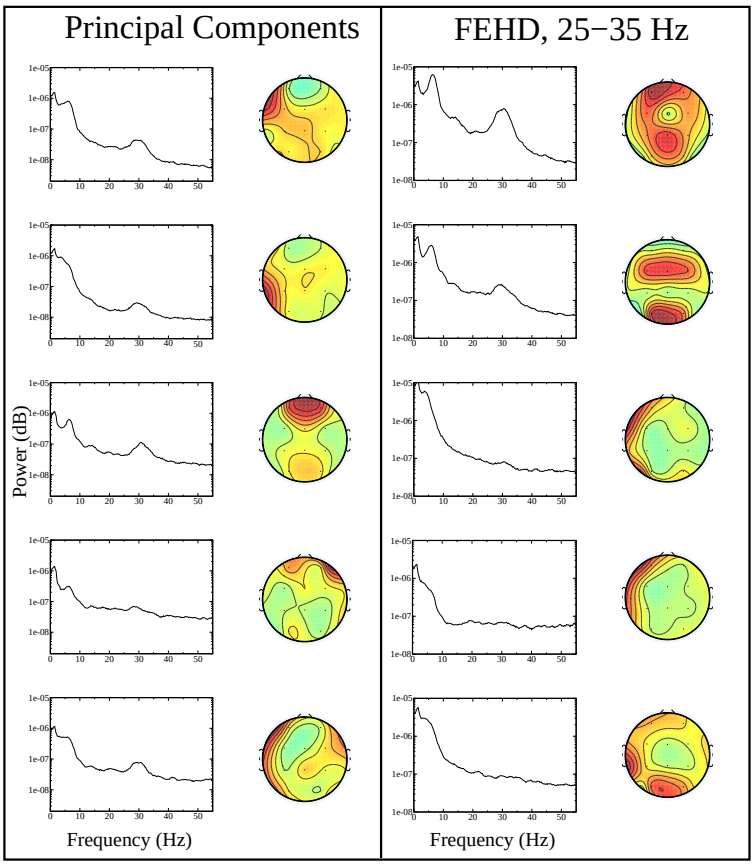


Figure 14: Principal components and FEHD components of EEG recorded from patient subject IN388R at 10 years into recovery from a severe brain injury. In the left column are the first five principal components, line spectra and the spatial weights. In the right column are the first five hierarchical components as determined by FEHD over the frequency range $25 - 35$Hz. The hierarchy is from top to bottom.

because of loss of excitatory inputs [14, 19]. The theta frequency (4-8Hz) burst of the thalamic neurons drive cortical cells and produce slow oscillations in the EEG. The arriving bursts are theorized to create a lateral inhibition effect within the surrounding cortex, resulting in the generation of the linked higher frequency rhythm (this is known as the *edge effect* [19, 18]).

Both the general increase in global cerebral metabolism and the persistent left sided hypometabolism seen in the patient's PET images (Figure 13) support the possible emergence of the combined 4-8Hz and 25-35Hz features in the resting EEG as resulting from this mechanism. Based on this hypothesis we computed the generators for each of these rhythms using FEHD over 4-8 Hz for the low frequency peak, and 25-35 Hz for the high frequency peak (Figure 14) and found two elements of evidence in favor of this hypothesis. The first observation is in comparison to findings of a separation of dominant low and high

frequency features in the power spectrum under FEHD for a typical healthy subject (see Figures 11 and 12). Figure 14 shows that the application of FEHD consolidates the high-frequency rhythm into the two first components, whereas this rhythm was pervasive among the principal components. Moreover, the low frequency theta peak is also largely contained in these first two components. This suggests a dynamical linkage of these two rhythms that was not seen in the analysis of healthy individuals (Figures 11 and 12). The second element of evidence is that the generators underlying these rhythms appear to be spatially located adjacent to the most severely injured regions of the cerebral cortex, corresponding to partially innervated cortical regions with some intact thalamocortical connections and residual cerebral metabolism.

The results of the FEHD analysis, taken together with the patient's marked behavioral improvement - becoming alert and communicative - are supportive of the hypothesis that there is a TCD process emanating from deafferented cortical tissue surrounding the large lesion in the left hemisphere (see Figure 13).

**Removal of pathological dynamical elements and analysis of remaining residual resting EEG signals**

FEHD provides a linear transformation of the data, and one can use this linear operator to selectively remove various features that FEHD analysis identifies with the extraction of hierarchical components. Procedurally, the removal process is the same as when one removes a principal component or an independent component, via ICA, from a multivariate vector series: If $W = C_{M-1}$ is the matrix identified by FEHD (equation (11)) so that

$$Y(t) = WX(t),$$

where $X(t)$ is the original data and $Y(t)$ are the transformed components, we can "remove" components of $Y$ from the original data, $X$, using linear projections. For example, if we want to remove the first FEHD component, we can adjust the original data as

$$\tilde{X}(t) = X(t) - \frac{X(t) \cdot \vec{W}_1}{\|\vec{W}_1\|^2} \vec{W}_1,$$

where $\vec{W}_1$ is the first row of $W$. This ensures that

$$\vec{W}_1 \cdot \tilde{X}(t) = 0,$$

thus removing the first component of $Y$ from the original data $X$. Similarly, removal of multiple components from the original data simply requires determining an orthogonal basis for the relevant rows of $W$ and removal of the projection of the data onto each of these basis vectors, as above, one at a time.

Since FEHD identifies components based on the Granger causality, removal of components can allow the user to remove such isolated dynamical systems that are identified by the analysis. This approach in principle can uncover rhythms that might originally be hidden in

the signal by the increased variance contributed by the large pathological component. Figure 15 demonstrates the validity of this approach using data from a second severely brain-injured patient subject . In this subject, there is a dominant rhythm in the 3-6Hz range that originates from the severely injured right posterior regions of the cortex (Figure 15 C); this 3-6Hz rhythm dominates the global dynamics of the EEG and is most prominent within the first and fourth hierarchical components (Figure 15 A). This same generator appears for the $\alpha$ frequency range when applying FEHD to the raw EEG. Low frequency rhythms in the EEG are common for injured brains, and often pathological behavior. We remove the three FEHD components (3-6 Hz) that have a concentrated peak in this range (components 1,4, and 5 - outlined in red) as described above. We repeat the FEHD analysis in the $\alpha$-range for the residual data set of remaining EEG signals. As shown in Figure 15 B, FEHD of the residual EEG now identifies and isolates a posterior occipital generator containing a very prominent $\alpha$-rhythm. This generator lies within the intact left hemisphere, which demonstrates a normal pattern of cerebral metabolism (Figure 15 C). We propose that the FEHD analysis of the residual EEG uncovers a normal alpha rhythm generated by the left hemisphere that is masked by dominant low frequency activity emanating from the damaged right hemisphere.

Following the example shown in Figure 15, the findings in patient IN388R above (Figure 14) also suggest the possibility that the right hemisphere in this patient which shows a normal metabolic profile (PET) and structural integrity in MRI images might generate a normal alpha rhythm. By removing the image of the data that projects onto these two components and performing the FEHD analysis over the $\alpha$ range, we are able to isolate a 10 Hz rhythm for this subject as well (data not shown). This application of FEHD may have considerable diagnostic value if it allows identification of functional brain activity hidden by pathological processes that add strong signal variance to the resting EEG and mask evidence of recovery. This approach may be particularly important in the evaluation of patients in intensive care settings [24].

# 4 Discussion

The purpose of this manuscript is to provide proof of concept for the FEHD method. To demonstrate the advantages of the method we show analyses of a series of toy examples and then evaluate several different applications to the resting state EEG recorded from both healthy volunteers and patient subjects with severe brain injuries. We show the consistency of the FEHD decomposition of low and high frequency background rhythms in the EEG of four healthy volunteers studied longitudinally. We demonstrate the robustness and stability of FEHD measurements (Figures 11 and 12, and Tables 1 and 2). With the exception of one subject, our results show a consistency in the generators determined at the different time points, both in terms of scalp location and power spectrum.

A more rigorous analysis of the similarity is done by computing the downstream/upstream ratio of pairwise Granger causality.

Additionally, computing the downstream/upstream pairwise ratio of Granger causality provides for a more rigorous analysis. This analysis is intended to determine the robustness

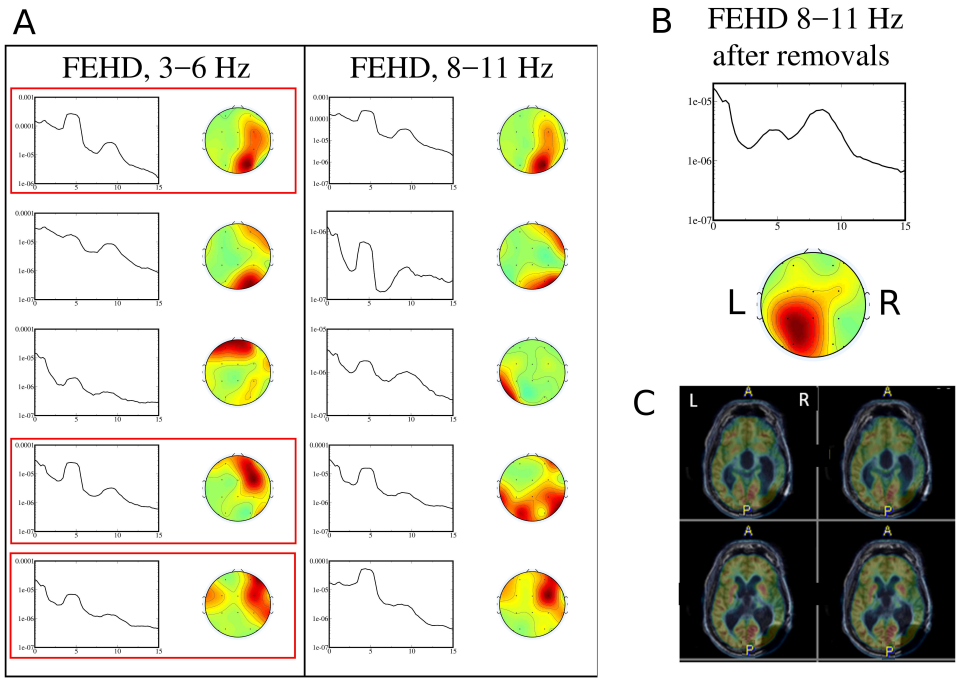


Figure 15: Demonstration of the use of FEHD to project a dominant pathological rhythm from the raw EEG and to examine the dynamics of the remaining residual EEG signal. A. Patient subject IN373M demonstrates a 3-6Hz peak in the power spectrum that is generated over the severely injured right posterior regions of the cortex (Panel C); this 3-6Hz rhythm dominate the global dynamics of the EEG and is most prominent with the first and fourth hierarchical component; the same generator also appears for the $\alpha$-range when applying FEHD to the raw EEG. B. Removal of the first and fourth hierarchical component identified by FEHD for the 3-6 range (the most concentrated peak in this range) from the data (outlined in red boxes) as described in text followed by FEHD analysis of the residual EEG signals on the $\alpha$-range for the residual data set of remaining EEG signals. As shown in Panel B, FEHD of the residual EEG now identifies and isolates a very prominent $\alpha$- rhythm on the left side of the head. This components appears as the second hierarchical component of the FEHD decomposition of the residual EEG and its generators are within the intact left hemisphere which demonstrates a normal pattern of cerebral metabolism. C. Fused MRI/FDG-PET images of patients brain illustrating large right sided region of cortical injury and hypometabolism.

of these determined networks. In every case, the method is successful in determining a hierarchical time series (Table 1). Moreover, for the $\alpha$ range, it is often the case that the VTS made by applying the FEHD transformation derived at one time point to the data from the other time point maintains the hierarchical structure of the non-cross-transformed data set. The results for the $\beta$ range are also positive, in the sense that the cross-transformed networks are more causal in the $\beta$ range than the $\alpha$ range; although the cross-transformed networks have much lower ratios. This was not unexpected for high frequency EEG activity since it is much more variable and includes more types of artifact than $\alpha$ activity. Test-retest EEG differences in high frequency activity were observed in the power spectrum, so we were not surprised that the HD time series were more reproducible for the $\alpha$ range than for the $\beta$-range.

Across subjects, the results are consistent with what is known about global brain dynamics in the frequency bands tested - $\alpha$ generators are largely occipital in location, while $\beta$ activity generally originates from frontal areas. The ability to capture this "ground truth" behavior indicates that this method can aid in the study of recovery from severe brain injury. Furthermore, FEHD can identify distinct dynamical features present in the injured brain which are opposite to those characteristic of healthy human subjects. We further demonstrate that FEHD can be used to investigate these pathological features of the injured brain and generate hypotheses as to the origins of these features. Removal of corresponding components of the FEHD determined time series can uncover hidden evidence of recovery of normal brain function in the injured brain. Each of the observations shown depends on the capacity of FEHD to resolve causal hierarchies within selected band limited frequency ranges.

## 4.1 Relationship between FEHD and other methods

FEHD can be compared with several other methods of causal inference that operate within specific frequency bands. For example, the directed transfer function (DTF), partial directed coherence (PDC), and Granger causality are each methods that provide an index for how much one signal influences another within a frequency band. Also, it is possible to determine transfer entropy on band-passed data to establish cross-frequency interactions but with more generality to include nonlinear interactions not testable utilizing FEHD, DTF, or any other method based on AR models. The critical distinction when directly comparing FEHD, as a tool, with other approaches is that only FEHD can quantify the relationship among hidden generators that are not sampled directly by electrodes, but exist as a superposition across sensors. For example, in section 3.2.1 we demonstrate the extraction of a narrow-band generator with hidden dynamics from the resting EEG; this finding derives from the isolation and removal of a component after the FEHD unmixing process followed by a reanalysis again utilizing FEHD. This sequence of analysis, biological inference and interpretation, and reanalysis of a residual dataset is unique to FEHD among these other related methods.

FEHD is compatible with other mathematical methods utilizing causal inference and can be thought of as a minimization strategy. In this paper, we chose to use Granger causality of a single component onto others as the function to be minimized. One could alternatively choose

to minimize the upstream DTF or PDC component by component, and obtain a hierarchical decomposition. The same approach can be taken using transfer entropy of band passed data, though non-parametric methods likely introduce rather large computation times.

## 4.2 Causality within frequency bands

The major advance of FEHD over HD is the ability to perform the hierarchical decomposition across specific frequency bands. The scalp EEG has many stereotypical rhythms with known behavioral correlates. For example, the posterior alpha rhythm is a kind of idling rhythm which may have more causal influences in behaviors (e.g. role in reaction times) which is reduced when the subject focuses on a task. A major motivation for the development of this method is the identification of a consistent typology of global dynamics in patients recovering consciousness after severe brain injuries utilizing standard HD [10]. Our working hypothesis is that particular pathological features of the EEG in the structurally injured brain will appear during the recovery process and have dynamically isolated frequency content and spatial localization that changes with levels of functional reafferentation across the corticothalamic system [10, 30, 25, 23]. The key advantages of the FEHD method as illustrated here provide the essential tools for further exploration and testing of these ideas.

Other methods have been used to determine the frequency by frequency causality between two time series generated by neural data. For example, in [5], the authors use transfer entropy on band-passed data to demonstrate the causal relationships within the visual cortex both for the spontaneous (baseline condition) and stimulated (movie condition). In [6], the authors use Granger causality over the gamma band to show that propagation in the gamma frequency band is mostly uni-directional between V1 and V4 in the primate visual cortex. Granger causality has also been used to demonstrate the directionality of theta, alpha, and beta rhythms in the primate visual cortex [4, 15, 29]. Each of these studies used the causal inference apparatus (transfer entropy or Granger causality) to establish the causal relationships within pairs of locations. It is in this regard the FEHD (and HD) are different. Rather than establish the causal relationships between locations, FEHD seeks to recombine the components of a multivariate time series so that the resulting VTS possesses a hierarchical structure - or as close to a hierarchical structure as possible. This approach makes FEHD unique. FEHD and HD are the only methods where the components are determined to have the desired causal structure. This allows the user to locate generators.

FEHD provides a different, unique analysis over the methods cited above. Those methods treat channels as sources, and the goal of the analysis is to determine the causal influence between these channels. FEHD, alternatively, seeks linear combinations of channels to form a hierarchy. The goal of FEHD is to determine superpositions of channels (linear combinations) so that these combinations have a hierarchical structure. In other words, rather than knowing how much A causes B, and B causes C and so on, FEHD seeks to determine combinations of A,B,and C to form a hierarchy.

The purpose of this manuscript is to describe and demonstrate FEHD in the context of EEG. However, the method is capable of working on arbitrary (albeit, close to Gaussian) time series such as, for example, local field potential recordings. While local field potential

(LFP) recordings have high fidelity of their origin within the sampled local networks, many such studies sample a small, heterogeneous sets of neuronal populations in complex tissues, particularly studies done within subcortical structures (for example [3] where $> 30$ microwire LFP electrodes were simultaneously sampled across many frontostriatal populations). It is possible that broader participation and independent control of specific frequency bands within each sampled LFP signal may depend on larger networks hidden from the sensors and afford a similar simplification of the dynamics for building biological inference. In this way FEHD offers a powerful and complementary window on causality. Future studies will determine the value of FEHD in such contexts. As noted above, even closely sampled local field potentials have already been shown to carry distinct causal interactions within specific frequency band [3].

## 4.3 FEHD as an extension of HD

FEHD is an extension of the hierarchical decomposition (HD) [22]. FEHD and HD seek a "causal hierarchy", which can also be described as a feed-forward network. Goldman showed that any multivariate time series can be viewed as a feed-forward network (with two-component feedback loops corresponding to complex eigenvalues) when viewed as a linear iterative map [12]. The transformation relies on the Schur factorization, which entails finding a lower triangular matrix that is similar, via an orthogonal matrix, to the original iteration matrix. A linear auto-regressive model is also a linear iterative map, but unlike the iterative maps described by Goldman, the predictions are obtained from multiple time points, or lags. There is no analogue to the Schur factorization for simultaneously triangularizing multiple matrices simultaneously - it is not generally possible (A necessary condition is given in [22] - for $j = 1..M$, where $M$ is the dimension of a lag matrix, there is a $j$ dimensional subspace that contains $j$ eigenvectors of all matrices). Thus, triangularizing the lag matrices of an autoregressive model is not generally possible.

Since simultaneous triangulation is generally not possible, HD adds some flexibility by minimizing the sum of squares of the upper triangular elements. When there are significant feedbacks, this minimization can lead to very erratic results where components determined to be at the bottom of the hierarchy, when checked for pairwise Granger causality, exert a lot of feedback on one or more components above. In using HD to analyze EEG data, we found this unpredictability to cause a lack of stability - the results could vary widely within subjects and conditions, and small parameter changes (such as the addition or subtraction of a lag, or component) proved very disruptive. Rather than attempt to simultaneously triangularize multiple matrices, FEHD simply locates, iteratively, the least causal component that can be obtained by linearly recombining the data. This least causal component is marked as such and removed, and the operation is repeated with the remaining components. This strategy leads to more predictable, stable results, which we demonstrate by showing reproducible results across time and across healthy subjects (Figures 11 and 12).

Because of the strategy we choose, the method is very resilient to the number of lags in the AR model. We use 40 lags for all of the analyses in this manuscript, but we have also verified that the results do not change much as a function of this parameter. The method

is as resilient to the number of principal components used as the data will allow. Principal component analysis treats each time point independently, whereas Granger Causality is largely determined by the statistical relationships between different points in time. This means that a very small principal component can be quite causal. If an additional principal component Granger causes one of the earlier components then it will rise up in the hierarchy, changing the results. The more components you use, the less likely the addition of a new one will introduce a previously orthogonal causal element. However, it cannot be guaranteed that it will not happen (see the last example provided in the synthetic time series section).

It is of note that we locate the minimally causal component at each step, rather than a maximally causal one. Though maximization of causal influence might prove useful for some applications, it is not useful in determining a hierarchy. That a component exerts a lot of Granger causality onto other components does not preclude it from being heavily caused itself.

## 4.4 Drawbacks of the method

At this time, we must use parametric Granger causality, where the kernel estimations of TE are replaced by much more efficient matrix rotations. The use of a linear AR model has drawbacks - primarily the lack of ability to determine non-linear relationships, and the assumption that there are no instantaneous interactions (this assumption is implicit in assuming orthogonal residuals). However, because of the technological restriction to parametric models, FEHD provides the only means of determining global dynamics in arbitrary frequency bands. One cannot, for example, use frequency-band passed data with a AR model, especially over very small frequency windows. First, AR models have only poles, so the zero power outside the pass band is a difficulty. Second, the sharp increase from zero power (outside the band pass) to the power of the band passed signal makes fitting equation (4) to the actual PSD difficult, similar to approximating the Heaviside function with a sigmoid,

$$\frac{1}{1+\sum_i A_i e^{-\lambda_i t}}$$

by changing only the $A_i's$. Doing so requires a lot of lags. This is further complicated when using least squares to compute the AR model. For small enough frequency bands, versions of the data that are shifted differently can become linearly dependant. This causes the normal matrix to become singular, and one cannot use enough lags to adequately model the data.

While not part of the FEHD method itself, the suggested use of the D/U-ratio to determine the success of the method in determining a largely feed-forward VTS is subject to same shortcomings as Granger causality. In particular, the inability of GC to distinguish between direct causality and indirect causality (ie. A causes B causes C vs. A causes B and C) can skew the statistic. One should check for this possibility of by, for example, looking at the grid of pairwise Granger causalities for suspect cascades. If there are three offending components, than one can use other methods to determine the directness or indirectness of these connenctions (for example, see [7]). It is important to note that this issue only applies

to the test statistic. Since FEHD works by peeling away non-causal combinations, it will order the components correctly regardless of whether the connections are direct or indirect.

### 4.5 Summary

In this paper, we describe a modification of the hierarchical decomposition (HD) method which we refer to as Frequency Extracted Hierarchical Decomposition (FEHD). FEHD is capable of determining the hierarchical decomposition of a multivariate time series within arbitrary frequency bands. Moreover, by approaching the minimization process differently, FEHD possesses large advantages over the original implementation of HD in terms of stability, robustness, and sensitivity to modeling parameters.

FEHD is implemented in FORTRAN, and has been tested using the gfortran compiler (of the gcc compiler collection, a product of the GNU project - www.gcc.org), and openmp for parallelization of the particle swarm optimization algorithm (www.openmp.org).

## 5 Acknowledgements

We would like to acknowledge the NIH-NINDS RO1 HD91251, the James S. McDonnell Foundation and the Jerold B. Katz Foundation for their support. We would also like to thank Jonathan Victor for many helpful conversations and suggestions, and Esteban Fridman with assistance with displays of structural and metabolic brain imaging data. Finally, we would like to thank Mary Conte for reading the manuscript and offering an enormous amount of helpful feedback.

## References

[1] Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, AC-19(6), December 1974.

[2] Hirotugu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer, 1998.

[3] Jonathan L. Baker, Jae-Wook Ryou, Xuefeng F. Wei, Christopher R. Butson, Nicholas D. Schiff, and Keith P. Purpura. Robust modulation of arousal regulation, performance, and frontostriatal activity through central thalamic deep brain stimulation in healthy nonhuman primates. *Journal of Neurophysiology*, 116(5):2383–2404, 2016.

[4] Andre Moraes Bastos, Julien Vezoli, Conrado Arturo Bosman, Jan-Mathijs Schoffelen, Robert Oostenveld, Jarrod Robert Dowdall, Peter De Weerd, Henry Kennedy, and Pascal Fries. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85:390–401, January 2015.

[5] Michel Besserve, Bernhard Scholkopf, Nikos K. Logothetis, and Stefano Panzeri. Causal relationaships between frequnecy bands of extracellular signals in visual cortex revealed by an information theoretic analysis. *Journal Computational Neuroscience*, 29:547–566, 2010.

[6] Conrado A. Bosman, Jan-Mathijs Schoffelen, Nicolas Brunet, Robert Oostenveld, Andre M. Bastos, Thilo Womelsdorf, Birthe Rubehn, Thomas Stieglitz, Peter De Weerd, and Pascal Fries. Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron*, 75:875–888, September 2012.

[7] Yonghong Chen, Steven Bressler, and Mingzhou Ding. Frequency decomposition of conditional granger causality and application to multivariate neural field potential data. *Journal of Neuroscience Methods*, 150:228–237, 2006.

[8] Schomer D and Lopes da Silva. *Niedermeyer's Electroencephalography. Basic Principles, clinical applications, and related fields.* Lippincott Williams & Wilkins, 6th edition, 2011.

[9] F. Lopes da Silva. Eeg and meg: relevance to neuroscience. *Neuron*, 80(5):1112–1128, 2013.

[10] J.D. Drover, M.M. Conte, A.M. Goldfine, H.U. Voss, J.D. Victor, and N.D. Schiff. Are low frequency oscillations in the eeg of severely injured brains a marker for functional reserve of cortical neurons? Poster session presented at: Society for Neuroscience, November 15, 2011, Washington DC.

[11] Pfurtschellar G. and Lopes da Silva FH. Event-related eeg/meg synchronization and desynchronization: basic principles. *Clin Neurophysiol*, 110(11):1842–57, Nov 1999.

[12] Mark S Goldman. Memory without feedback in a neural network. *Neuron*, 61:621–634, February 2009.

[13] Geweke J. Measurement of linear dependance and feedback between multiple time series. *Journal American Statistical Association*, 77:304–13, 1982.

[14] D Jeanmonod, M Magnin, and A Morel. Low-threshold calcium spike bursts in the human thalamus. common physiopathology for sensory, motor and limbic positive symptoms. *Brain*, 119:363–75, April 1996.

[15] Nathaniel J. Killian and Elizabeth A. Buffalo. Distinct frequnecies mark the direction of cortical communication. *PNAS*, 111(40):14316–14317, October 2014.

[16] R. Kus, M. Kaminski, and Blinowska KJ. Determination of eeg activity propogation: Pair-wise versus multichannel estimate. *IEEE Trans Biomed Eng*, 51(9):1501–1510, 2004.

[17] S. Laureys and N.D. Schiff. Coma and consciousness: Paradigms (re)framed by neuroimaging. *Neuroimage*, 61:478–491, 2012.

[18] Rodolfo Llinas, Francisco J. Urbano, Elena Leznik, Rey R. Ramizrez, and Hein J.F. van Marle. Rhythmic and dysrhythmic thalamocortical dynamics: Gaba systems and the edge effect. *TRENDS in Neurosciences*, 28(6), June 2005.

[19] Rodolfo R. Llinas, Urs Ribary, Daniel Jeanmonod, Eugene Kronberg, and Partha P. Mitra. Thalamocortical dysrhythmia: A neurological and neuropsychiatric syndrone characterized by magnetoencephalography. *PNAS*, 96(26):15222–15227, Dec 1999.

[20] Silva LR, Amitai Y, and Connors BW. Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science*, 251(4992):432–435, Jan 1991.

[21] Forgacs PB, Frey HP, Velazquez A, Thompson S, Brodie D, Moitra V, Rabani L, Park S, Agarwal S, Falo MC, Schiff ND, and Claassen J. Dynamic regimes of neocortical activity linked to corticothalamic integrity correlate with outcomes in acute anoxic brain injury after cardiac arrest. *Ann Clin Transl Neurol.*, 4(2):119–129, Jan 6 2017.

[22] M.A. Repucci, N.D. Schiff, and J.D. Victor. General strategy for hierarchical decomposition of multivariate time series: Implications for temporal lobe seizures. *Annals of Biomedical Engineering*, 29:1135–1149, 2001.

[23] ND Schiff. *Brain function and responsiveness in disorders of consciousness*, chapter Mesocircuit mechanisms underlying recovery of consciousness following severe brain injuries: models and predictions. Springer Verlag, 2016.

[24] ND Schiff. Uncovering hidden integrative cerebral function in the intensive care unit schiff. *Brain*, 140(9):2259–2262, September 2017.

[25] ND Schiff, T Nauval, and JD Victor. Large-scale brain dynamics in disorders of consciousness. *Current Opinion in Neurobiology*, 25:7–14, 2014.

[26] N.D Schiff, J.D. Victor, and A. Canel. Nonlinear autoregressive analysis of the 3 second ictal EEG: implications for underlying dynamics. *Biological Cybernetics*, 72:527–532, 1995.

[27] N.D. Schiff, J.D. Victor, A. Canel, and D.R. Labar. Characteristic nonlinearities of the 3 second ictal EEG identified by nonlinear autoregressive analysis. *Biological Cybernetics*, 72:519–526, 1995.

[28] J. Songsiri, J. Dahl, and L. Vandenberghe. Graphical models of autoregressive processes. In Y. Eldar and D. Palomar, editors, *Convex Optimization in Signal Processing and Communications*, pages 89–116. Cambridge University Press, 2010.

[29] Timo van Kerkoerle, Matthew W. Self, Bruno Dagnino, Marie-Alice Gariel-Mathis, Jasper Poort, Chris van der Togt, , and Pieter R. Roelfsema. Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *PNAS*, 111(40):14332–14341, October 2014.

[30] Shawniqua T Williams, Mary M Conte, Andrew M Goldfine, Quentin Noirhomme, Olivia Gosseries, Marie Thonnard, Bradley Beattie, Jennifer Hersh, Douglas I Katz, Jonathan D Victor, Steven Laureys, and Nicholas D Schiff. Common resting brain dynamics indicate a possible mechanism underlying zolpidem response in severe brain injury. *eLife*, November 2013.

[31] Michael S. Worden, John J. Foxe, Norman Wang, and Gergory V. Simpson. Anticipatory biasing of visuospatial attention indexed by retinotopically specific $\alpha$-band electroencephalography increases over occipital cortex. *The Journal of Neuroscience*, 20, 2000.