Multivariate Stochastic Volatility Modeling of Neural Data

Tung D. Phan, Jessica A. Wachter, Michael J. Kahana

University of Pennsylvania

Abstract

Because multivariate autoregressive models have failed to adequately account for the complexity of neural signals, researchers have predominantly relied on non-parametric methods when studying the relations between brain and behavior. Using a database of medial temporal lobe (MTL) recordings from 96 neurosurgical patients, we show that time series models with volatility described by a multivariate stochastic latent variable process and lagged interactions between signals in different brain regions provide new insights into the dynamics of brain function. We estimate both the parameters describing the latent variable processes and the directional correlations in volatility between brain regions using Bayesian sampling techniques. The implied volatility inferred from our process positively correlates with high-frequency spectral activity, a signal that correlates with neuronal activity and is widely used to study brain function. We show that volatility features derived from our model can reliably decode good vs. poor memory states, and that this classifier performs as well as those using spectral features. Using the multivariate stochastic volatility model, we uncovered hippocampal-perirhinal bidirectional connections in the MTL regions that are associated with successful memory encoding.

Multivariate Stochastic Volatility Modeling of Neural Data

Recent advances in neuroscience have enabled researchers to measure brain function with both high spatial and temporal resolution, leading to significant advances in our ability to relate complex behaviors to underlying neural signals. Because neural activities give rise to electrical potentials, much of our knowledge concerning the neural correlates of cognition derive from the analyses of multi-electrode recordings, which yield a multivariate time series of voltage recorded at varying brain locations (denoted here as $\mathbf{y}_t$). Such signals may be measured non-invasively, using scalp electroencephalography (EEG) methods, or invasively, using subdural grids or intraparynchymal depth electrodes in human neurosurgical patients. In recent years, intracranially recorded (iEEG) signals have yielded detailed information on correlations between time-series measures and a wide range of behaviors including perception, attention, learning, memory, language, problem solving and decision making (Jacobs & Kahana, 2010).

Whereas other fields that grapple with complex multivariate time series have made effective use of parametric models such as economics and engineering Kim et al. (1998); Blanchard & Simon (2001); West (1996), neuroscientists abandoned early parametric approaches (e.g., linear autoregressive models) in favor of non-parametric spectral decomposition methods, as a means of uncovering features of neural activity that may correlate with behavior. A strength of these non-parametric methods is that they have enabled researchers to link fluctuations in iEEG signals to low-frequency neural oscillations observed during certain behavioral or cognitive states, such as slow-wave sleep (Landolt et al., 1996; Chauvette et al., 2011; Nir et al., 2011), eye closure (Klimesch, 1999; Goldman et al., 2002; Laufs et al., 2003; Barry et al., 2007) or spatial exploration (Kahana et al., 2001; Raghavachari et al., 2001; Caplan et al., 2003; Ekstrom et al., 2005; Byrne et al., 2007). High-frequency neural activity, which has also been linked to a variety of cognitive and behavioral states (Maloney et al., 1997; Herrmann et al., 2004; Canolty et al., 2006), is less clearly oscillatory, and may reflect asynchronous stochastic volatility of the underlying

EEG signal (Burke et al., 2015).

Although spectral analysis methods have been used extensively in the neuroscience literature, they assume that there is unique information in each of a discrete set of frequency bands. The number of bands and frequency ranges used in these methods have been the subject of considerable controversy. Indeed, Manning et al. Manning et al. (2009) have shown that broadband power often correlates more strongly with neuronal activity than does power at any narrow band. Also, non-parametric methods implicitly assume that the measured activity is observed independently during each observational epoch, and at each frequency, an assumption which is easily rejected in the data, which show strong temporal autocorrelation as well as correlations among frequency bands Von Stein & Sarnthein (2000); Jensen & Colgin (2007); Axmacher et al. (2010). Moreover, non-parametric methods applied to EEG signals are typically done in a univariate fashion that neglects the spatial correlational structure. By simultaneously modeling the spatial and temporal structure in the data, parametric models confer greater statistical power so long as they are not poorly specified.

Parametric methods have been applied to various types of multivariate neural data including EEG Hesse et al. (2003); Dhamala et al. (2008); Bastos et al. (2015), magnetoencephalography (MEG) David et al. (2006), functional magnetic resonance imaging (FMRI) Roebroeck et al. (2005); Goebel et al. (2003); David et al. (2008), and local field potentials (LFP) Brovelli et al. (2004). These methods typically involve fitting vector autoregressive (VAR) models to multivariate neural data that are assumed to be stationary in a specific time interval of interest. The regression coefficient matrix derived from the VAR models can be used to study the flow of information between neuronal regions in the context of Granger causality (G-causality). Neuroscientists have used Gaussian VAR models to study the *effective connectivity* (directed influence) between activated brain areas during cognitive and visuomotor tasks Zhou et al. (2009); Deshpande et al. (2009); Graham et al. (2009); Roebroeck et al. (2005). Although VAR models and

G-causality methods have been argued to provide useful insights into the functional organization of the brain, their validity relies upon the assumptions of linearity, stationarity, and homoskedasticity (constant variance over time) of the neural data. When one of these assumptions is violated, the conclusions drawn from a G-causality analysis will be inconsistent and misleading Stokes & Purdon (2017). One of the most common violations is the assumption of homoskedasticity (Wong et al., 2006). Therefore, in the present work we adopt a stochastic volatility approach in which volatility is assumed to follow a stochastic process. Such models have been extremely useful in the analyses of financial market data which, like neural data, exhibits high kurtosis Heston (1993); Bates (1996); Barndorff-Nielsen (2002).

We propose a multivariate stochastic volatility (MSV) model with the aim of estimating the time-varying volatility of multivariate neural data and its spatial correlational structure. The MSV model assumes that the volatility series of iEEG signals follows a latent variable vector autoregressive process and it allows for the lagged signals of different brain regions to influence each other by specifying a full persistent matrix (typically assumed to be diagonal) in the VAR process for volatility. We employed a Bayesian approach to estimate the latent volatility series and the parameters of the MSV model using the forward filtering backward sampling and Metropolis Hastings algorithms.

We validated the MSV model in a unique dataset comprising depth-electrode recordings from 96 neurosurgical patients. These patients volunteered to participate in a verbal recall memory task while they were undergoing clinical monitoring to localize the epileptogenic foci responsible for seizure onset. Our analyses focused on the subset of electrodes (n =718) implanted in the medial temporal lobe (MTL) regions, including hippocampus, parahippocampal cortex, entorhinal cortex and perirhinal cortex. We chose to focus on these regions given their prominent role in the encoding and retrieval of episodic memories Davachi et al. (2003); Kirwan & Stark (2004); Kreiman et al. (2000); Sederberg et al. (2007).

We show that the MSV model with interactions between regions provides a substantially superior fit to MTL recordings than univariate stochastic volatility (SV) models. The implied volatility in these models positively correlates with non-parameter estimates of spectral power, especially in the gamma frequency band. We demonstrate the utility of our method for decoding cognitive states by using a logistic regression classifier trained on the implied volatility data across MTL electrodes to predict which studied items will be subsequently recalled. We find that the MSV-derived features outperform spectral features in decoding cognitive states, supporting the value of this model-based time-series analysis approach to the study of human cognition. Furthermore, using the MSV model to construct a directional MTL connectivity network, we find that significant bidirectional connectivity between the perirhinal cortex and the hippocampus predicts successful memory encoding.

Multivariate Stochastic Volatility Models

Volatility, the instantaneous variance of a time series, has been used extensively in financial applications to forecast aspects of future returns, price derivatives, and study recessions, inflation and monetary policies Engle et al. (2001); Cogley & Sargent (2005); Blanchard & Simon (2001). There is by now a large literature on stochastic volatility models and methods for estimating these models either by closed-form solutions Heston (1993); Heston & Nandi (2000) or by simulation Harvey & Shephard (1996); Kim et al. (1998); Omori et al. (2007). In this study, we utilize a multivariate stochastic volatility model with the aim of using volatility to study cognitive processes, in particular, ones that are related to memory. One innovation that is made in this paper is that we allow for volatility series of different regions in the brain to influence each other, thus, enabling us to study the directionality of the correlational structure of the considered brain regions. Specifically, we allow for the volatility series of different regions to Granger-cause each

other. Let $\mathbf{y}_t = (y_{1t}, \cdots, y_{Nt})'$ be a multivariate iEEG time-series at N electrodes[1]. We model y_{jt}, $1 \leq j \leq N$, as follows:

$$y_{jt} = \exp(\frac{x_{jt}}{2})\epsilon^y_{jt}, \tag{1}$$

and

$$x_{jt} - \mu_j = \sum_{k=1}^{N} \beta_{jk}(x_{kt-1} - \mu_k) + \epsilon^x_{jt}, \tag{2}$$

where the error terms follow multivariate normal ($\mathcal{MVN}$) distributions: $\epsilon^y_t = (\epsilon^y_{1t}, \cdots, \epsilon^y_{Nt}) \sim \mathcal{MVN}(0, I_N)$, $\epsilon^x_t = (\epsilon^x_{1t}, \cdots, \epsilon^x_{Nt}) \sim \mathcal{MVN}(0, \Sigma)$. I_N denotes the identity matrix of dimension N, and $\Sigma = diag(\sigma^2_1, \cdots, \sigma^2_N)$ is assumed to be diagonal. That is, y_{jt} is a time series whose conditional log-variance (log-volatility), x_{jt}, follows an AR(1) process that depends on its past value and the past values of other electrodes. The series $\{y_{1t}\}_{t=1}^T, \cdots, \{y_{Nt}\}_{t=1}^T$ are assumed to be conditionally independent given their log-volatility series $\{x_{1t}\}_{t=1}^T, \cdots, \{x_{Nt}\}_{t=1}^T$. The coefficient, b_{jk}, models how the past value of channel k affects the current value of channel j. We can rewrite Eqn. 2 in a matrix form,

$$\mathbf{x}_t - \mu = B(\mathbf{x}_{t-1} - \mu) + \epsilon^x_t \tag{3}$$

where $\mathbf{x}_t = (x_{1t}, \cdots, x_{Nt}), \mu = (\mu_1, \cdots, \mu_N)$ and $B(j,k) = \beta_{jk}$. The vector error terms ϵ^y_t and ϵ^x_t are assumed to be independent. The parameters in the system above are assumed to be unknown and need to be estimated.

Following a Bayesian perspective, we assume that the parameters are not completely unknown, but they follow some prior distributions. Then, using the prior distributions and the information provided by the data, we can make inferences about the parameters from their posterior distributions.

[1]We detrended the raw time-series using autoregressive models of order p, where p was selected based on the Akaike information criterion (AIC).

Priors and Estimation Method

We specify prior distributions for the set of parameters $\theta = (\mu, B, \sigma^2)$ of the MSV model. The mean vector μ follows a multivariate normal distribution $\mu \sim \mathcal{MVN}(b_\mu, B_\mu)$. Each entry of the persistence matrix $B_{ij} \in (-1, 1)$ is assumed to follow a beta distribution, $(B_{ij} + 1)/2 \sim Beta(a_0, b_0)$ Kim et al. (1998). For volatility of volatility, we utilize a gamma prior, $\sigma_j \sim \Gamma(1/2, 1/2B_\sigma)$ Kastner & Frühwirth-Schnatter (2014), which is equivalent to $\pm\sqrt{\sigma_j^2} \sim N(0, B_\sigma)$. We estimated the latent volatility processes and the parameters of the MSV model using a Metropolis-within-Gibbs sampler Kim et al. (1998); Omori et al. (2007); Kastner & Frühwirth-Schnatter (2014) (See S.I. for derivation and for discussion of parameter identification).

Results

We analyzed the behavioral and electrophysiological data of 96 subjects implanted with subdural and depth electrodes during a free recall memory task. Subjects learned at least 25 lists of 12 unrelated words presented on a screen, with each list followed by a short arithmetic distractor task. Then, subjects recalled as many words from the previously studied list as possible in any order (Fig. 1). For all of our analyses, we only considered the encoding periods of the task. In addition, we focused our analyses on the MTL regions that have been implicated in episodic memory encoding Squire & Zola-Morgan (1991); E. Solomon et al. (2017a); Long & Kahana (2015). To assess a particular effect across subjects, we utilized the maximum a posteriori (MAP) estimate by taking the posterior mean of the variable of interest (whether it be the volatility time series $\mathbf{x}_t$ or the regression coefficient matrix B) Stephan et al. (2010).

[Figure 1 about here]

Model Comparison

To establish the validity of the MSV model, we compared its performance to that of univariate stochastic volatility models (equivalent to setting all the off-diagonal entries of the matrix B in Eqn. 2 to 0) in fitting iEEG data. We applied the MSV model to the multivariate neural data combined across encoding periods (regardless of whether the word items were later recalled) and SV models to datasets of individual electrodes. We utilized the deviance information criterion (DIC) Spiegelhalter et al. (2002); Gelman et al. (2014) considered to be a Bayesian analogue of the Akaike information criterion (AIC) to evaluate the performance of the models. The DIC consists of two components: the negative log-likelihood, $\bar{D} = \mathbb{E}_{\theta,\mathbf{x}|y}[-2\log f(y \mid \theta, \mathbf{x})]$, which measures the goodness-of-fit of the model and the effective number of parameters,
$p_D = \bar{D} - D(\bar{\theta}, \bar{\mathbf{x}}) = \mathbb{E}_{\theta,\mathbf{x}|y}[-2\log f(y \mid \theta, \mathbf{x})] + 2\log f(y \mid \bar{\theta}, \bar{\mathbf{x}})$, which measures the complexity of the model. Where $\bar{\theta}$ and $\bar{\mathbf{x}}$ denote the posterior means of the latent volatility series and the parameters of the MSV model. The DIC balances the trade-off between model fit and model complexity. Models with smaller DICs are preferred. To account for the varying amount of data each subject had, we averaged the DIC by the number of events and electrodes. We found the MSV model to have a consistently lower DIC value than the SV model with a mean difference of 23 ($\pm$ 5.9 SEM). This indicates that the MSV model is more than 150 times as probable as the SV models Kass & Raftery (1995), suggesting that the MSV model is a more appropriate model for iEEG data.

Relation to Spectral Power

We next analyzed the relation between volatility and spectral power (see Materials & Methods) over a wide range of frequencies, from 3 to 180 Hz in 1 Hz steps). For each subject, we computed the correlation between volatility and spectral power for each encoding event and then averaged these correlations across all events. Since spectral powers of close frequencies are highly correlated, we utilized a Gaussian process model Rasmussen

(2004) to estimate the correlation between volatility and spectral power as a function of frequency. Fig. 2 indicates that the correlation between volatility and spectral power is significantly positive across the spectrum and increasing in frequency. This illustrates the broadband nature of the volatility measure, but also suggests that volatility may more closely relate to previous neuroscientific findings observed for high-frequency as compared with low-frequency activity. Having established that the MSV model outperforms the more traditional SV approach, and having shown that the implied volatility of the series reliably correlates with high frequency neural activity, we next asked whether we can use the model-derived time series of volatility to predict subjects' behavior in a memory task.

[Figure 2 about here]

Classification of Subsequent Memory Recall

Extensive previous work on the electrophysiological correlates of memory encoding has shown that spectral power, in both the low frequency (4-8 Hz) theta band and at frequencies about 40 Hz (so called gamma activity), reliably predicts which studied words will be subsequently recalled or recognized Sederberg et al. (2003). Here we ask whether the implied volatility derived from the MSV model during word encoding can also reliably predict subsequent recall. To benchmark our MSV findings we conducted parallel analyses of wavelet-derived spectral power at frequencies ranging between 3-180 Hz. To aggregate across MTL electrodes within each subject we applied an L2-penalized logistic regression classifier using features extracted during the encoding period to predict subsequent memory performance. To estimate the generalization of the classifier, we utilized a nested cross-validation procedure in which we trained the model on $N-1$ sessions using the optimal penalty parameter selected via another inner cross-validation procedure on the same training data. We then tested the classifier on a hold-out session collected on a different day. We computed the receiver operating characteristic (ROC) curve, relating true and false positives, as a function of the criterion used to assign regression output to

response labels. We then use the AUC metric (area under the ROC curve) to characterize model performance. We find that MSV-model implied volatility during item encoding reliably predicts subsequent recall, yielding an average AUC of 0.53 (95% CI, from 0.51 to 0.55). AUCs reliably exceeded chance levels in 73 percent of subjects (33 out of 45 subjects who contributed at least 3 sessions of data). Fig. 3 compares these findings against results obtained using wavelet-derived power. Here we see that implied volatility does as well as, or better than, spectral measures at nearly all frequencies. In order to capture the correlation between spectral powers (thus their corresponding classifiers' performances), we fitted a Gaussian regression model to test the functional form of ΔAUC. We find that the ΔAUC function is significantly different from the 0 function ($\chi^2_{11} = 42$, $P < 10^{-5}$) Benavoli & Mangili (2015), which indicates that on average volatility performs significantly better than spectral power in predicting subsequent memory recall.

[Figure 3 about here]

Directional Connectivity Analysis

Having established that volatility is predictive of subsequent memory recall, we now seek to identify directional connections between MTL subregions that are related to successful memory encoding. To investigate the intra-MTL directional connectivity patterns that correlate with successful memory encoding, we utilize a subsequent memory effect (SME) paradigm in which we compare the MTL directional connectivity patterns (regression coefficient matrix B) associated with recalled (R) word items to those associated with non-recalled (NR) items. The SME paradigm has been widely used in the memory literature to study neural correlates (typically spectral power in a specific frequency band) that predict successful memory formation Sederberg et al. (2003); Long et al. (2014).

The intra-MTL connectivity SME was constructed using the following procedure. First, we separated the word items into recalled and non-recalled items offline. Using the MSV model, we constructed an intra-MTL connectivity network for each memory outcome.

We compared the distribution of the elements of these matrices across subjects. For the analysis, we considered four subregions of the MTL: hippocampus (Hipp), entorhinal cortex (EC), perirhinal cortex (PRC), and parahippocampal cortex (PHC). We then computed the contrast between the two intra-MTL networks corresponding to recalled and non-recalled items for each ordered pair of subregions excluding the ones with fewer than 10 subjects contributing to the analysis. To compute the directional connectivity from region I to region J, we took the average of the "influences" that electrodes in region I have on electrodes in region J: $\Delta_{I\to J} = \frac{1}{|I||J|}\sum_{i\in I, j\in J}(B^R_{ij} - B^{NR}_{ij})$, where $|I|$ denotes the number of electrodes in region I. Finally, we averaged the contrast for each ordered pair of MTL subregions across sessions. Fig. 4 illustrates the intra-MTL connectivity SME for the left and right hemispheres. Directed connections between the left hippocampus and the left PRC reliably decreases (false-discovery-rate-corrected) during successful memory encoding ($\Delta_{Hipp\to PRC} = -0.04, t_{47} = -3.49$, adj. $P < 0.01$ and $\Delta_{PRC\to Hipp} = -0.06, t_{47} = -2.66$, adj. $P < 0.05$). The difference between the directional connections between these two regions is not significant ($t_{47} = 0.53$, $P = 0.60$). The decreases in the bi-directional connections between the hippocampus and the perirhinal cortex are consistent with the findings in E. Solomon et al. (2017a) which suggested that gamma networks desynchronize during encoding. We did not, however, find any other significant directional corrections among the remaining regions (Fig. 4, Tables 2 and 3).

[Figure 4 about here]

Discussion

The ability to record electrophysiological signals from large numbers of brain recording sites has created a wealth of data on the neural basis of behavior and a pressing need for statistical methods suited to the properties of multivariate, neural, time-series data. Because neural data strongly violate homoskedasticity assumptions underlying standard approaches, such as Granger causality Stokes & Purdon (2017), researchers have

generally eschewed these model-based approaches and embraced non-parameter data analytic procedures. The multivariate stochastic volatility framework that we propose allows for non-stationarity while taking full advantage of the Granger-causality approach. This framework allows us to explicitly model the time-varying variance of neural signals. Similar stochastic volatility models have been used extensively in the financial economics literature to characterize a wide range of phenomena.

Applying MSV models to recordings from indwelling electrodes in 96 neurosurgical patients allowed us to model neural connectivity among medial-temporal lobe (MTL) subregions, and show how changes in connectivity predicted successful memory encoding. We further demonstrated how the implied volatility extracted from these models correlated with spectral power over a wide range of frequencies. We found that volatility is significantly positively correlated with spectral power, and this correlation increases with frequency.

We further verified that neural volatility estimated during word encoding reliably predicts subsequent recall. By comparing a penalized logistic-regression classifier on volatility features to one trained on spectral power, we found that volatility performs at least as well as spectral power at any frequency in predicting subsequent recall. Using Gaussian process regression we were able to confirm that our MSV model-derived volatility features significantly outperformed the spectral features in decoding memory processes in the human brain.

A key strength of the MSV approach is its ability to identify directed interactions between brain regions without assuming stationarity. We thus used this approach to determine the directional connections between MTL sub-regions that correlate with successful memory encoding. Using the regression coefficient matrix of the multivariate volatility process, we found that periods of decreased connectivity between the hippocampus and the perirhinal cortex generally predicted successful memory encoding. Prior studies provide some precedent for this result by showing that brain regions often

exhibit desynchronization of high-frequency activity during memory encoding E. Solomon et al. (2017b); Burke et al. (2015).

To model the complex dynamics of neural time-series data signals, MSV models require a large amount of data to accurately estimate the latent volatility process and model parameters. Although such large datasets are not always available, high-frequency neural recordings provide an excellent testbed for application and further development of this class of models. By presenting a first application of this approach to neural time-series data, we believe researchers will be able to further extend these models to broader classes of neural recordings, and exploit their statistical power to substantially increase our understanding of how behavior emerges from the complex interplay of neural activity across many brain regions.

Materials and Methods

Participants. Ninety six patients with drug-resistant epilepsy undergoing intracranial electroencephalographic monitoring were recruited in this study. Data were collected as part of a study of the effects of electrical stimulation on memory-related brain function at multiple medical centers. Surgery and iEEG monitoring were performed at the following centers: Thomas Jefferson University Hospital (Philadelphia, PA), Mayo Clinic (Rochester, MN), Hospital of the University of Pennsylvania (Philadelphia, PA), Emory University Hospital (Atlanta, GA), University of Texas Southwestern Medical Center (Dallas, TX), Dartmouth-Hitchcock Medical Center (Lebanon, NH), Columbia University Medical Center (New York, NY) and the National Institutes of Health (Bethesda, MD). The research protocol was approved by the Institutional Review Board at each hospital and informed consent was obtained from each participant. Electrophysiological signals were collected from electrodes implanted subdurally on the cortical surface and within brain parenchyma. The neurosurgeons at each clinical site determined the placement of electrodes to best localize epileptogenic regions. Across the clinical sites, the following models of depth and

grid electrodes (electrode diameter in pararentheses) were used: PMT Depthalon (0.86 mm); Adtech Spencer RD (0.86mm); Adtech Spencer SD (1.12mm); Adtech Behnke-Fried (1.28mm); Adtech subdural and grids (2.3mm). The dataset can be requested at `http://memory.psych.upenn.edu/RAM_Public_Data`.

Free-recall task. Each subject participated in a delayed free-recall task in which they were instructed to study a list of words for later recall test. The task is comprised of three parts: encoding, delay, and retrieval. During encoding, the subjects were presented with a list of 12 words that were randomly selected from a pool of nouns (`http://memory.psych.upenn.edu/WordPools`). Each word presentation lasts for 1600 ms followed by a blank inter-stimulus interval (ISI) of 800 to 1200 ms. To mitigate the recency effect (recalling last items best) and the primacy effect (recalling first items better than the middle items), subjects were asked to perform a math distraction task immediately after the presentation of the last word. The math problems were of the form A+B+C = ?, where A,B,C were randomly selected digits. The delay math task lasted for 20 seconds, after which subjects were asked to recall as many words as possible from the recent list of words, in any order during the 30-second recall period. Subjects performed up to 25 lists per session of recording (300 words). Multiple sessions were recorded over the course of the patient's hospital stay.

Electrophysiological Recordings and Data Processing. iEEG signals were recorded from subdural and depth electrodes at various sampling rates (500, 1000, or 1600 Hz) based on the the amplifier and the preference of the clinical team using one of the following EEG systems: DeltaMed XlTek (Natus), Grass Telefactor, and Nihon-Kohden. We applied a 5 Hz band-stop fourth order Butterworth filter centered on 60 Hz to attenuate signal from electrical noise. We re-referenced the data using the common average of all electrodes in the MTL to eliminate potentially confounding large-scale artifacts and noise. We used Morlet wavelet transform (wave number = 5) to compute power as a function of time for our iEEG signals. The frequencies were sample linearly from 3 to 180 Hz with 1 Hz

increments. For each electrode and frequency, spectral power was log-transformed and then averaged over the encoding period. Within a session of recording, the spectral power was z-scored using the distribution of power features across events. To extract volatility feature, we applied the MSV model to the dataset constructed from concatenating encoding events within a session. The implied volatility was then averaged over the encoding period.

Anatomical Localization. The MTL electrodes were anatomically localized using the following procedure. Hippocampal subfields and MTL cortices were automatically labeled in a pre-implant 2 mm thick T2-weighted MRI using the Automatic segmentation of hippocampal subfields (ASHS) multi-atlas segmentation method Yushkevich et al. (2015). A post-implant was co-registered with the MRI using Advanced Normalization Tools Avants et al. (2008). MTL depth electrodes that were visible in the CT were then localized by a pair of neuroradiologists with expertise in MTL anatomy.

Statistical Analyses. To assess an effect across subjects, we applied classical statistical tests on the maximum a posteriori (MAP) estimate of the parameter of interest . This approach has been used in many Bayesian applications to FMRI studies Stephan et al. (2010) to test an effect across subjects. For analyses concerning frequencies, we applied Gaussian regression models Rasmussen (2004) to take the correlations among frequencies into account. We used the Matern (5/2) kernel function for all analyses that used Gaussian regression models. p-values were FDR-corrected at $\alpha = 0.05$ significance level when multiple tests were conducted.

Supplementary Materials

MCMC Algorithm for MSV Models

In this section, we derive an MCMC algorithm, which consists of Metropolis-Hastings steps within a Gibbs sampler, for estimating the latent volatility process and its parameters. As before, let $\mathbf{x}_t$ denote the latent multivariate volatility time-series and $\theta = (\mu, B, \sigma^2)$ the parameters in the MSV model. Following Kim et al. (1998); Omori et al.

(2007), we log-transform Eqn. 1

$$y^*_{jt} = x_{jt} + \log((\epsilon^y_{jt})^2), \tag{4}$$

where $y^*_{jt} = \log(y^2_{jt} + c)$ with a fixed offset constant $c = 10^{-4}$ to avoid values equal to 0. Eqn. 4 is linear but non-Gaussian. To ameliorate the non-Gaussianity problem, we approximate the log-transformed error term $\log((\epsilon^y_{jt})^2) \sim \log(\chi^2_1)$ by a mixture of 10 normal distributions as in Omori et al. (2007):

$$\log(\chi^2_1) \sim \sum_{k=1}^{10} p_k N(m_k, v^2_k).$$

The values of p_k, m_k and v_k are tabulated in Omori et al. (2007). As a result, we introduce a latent mixture component indicator variable, r_{jt}, for channel j at time t such that $\log((\epsilon^y)^2) \mid (r_{jt} = k) \sim N(m_k, v^2_k)$. The indicator variable is also estimated in the MCMC sampler. Given the mixture indicator $\mathbf{r}_t$ and the vector parameter θ, the latent volatility series $\mathbf{x}_t$ can be sampled using a forward-filtering and backward-sampling (FFBS) procedure West (1996). The mixture indicator $\mathbf{r}_t$ can be sampled from a multinomial distribution

$$P(r_{jt} = k \mid \mathbf{x}_t, \theta) \propto P(r_{jt} = k) \frac{1}{v_k} \exp\Big\{ - \frac{(\tilde{y}_{jt} - x_{jt} - m_k)^2}{2v^2_k} \Big\}. \tag{5}$$

Finally, the vector parameter θ can be sampled using an ancillarity-sufficiency interweaving strategy (ASIS) Yu & Meng (2011); Kastner & Frühwirth-Schnatter (2014) which involves sampling the vector parameter θ given the unstandardized volatility series x_{jt} via a Metropolis-Hasting step (non-centered step) and then sampling θ again given the standardized volatility series $\tilde{x}_{jt} = \dfrac{x_{jt} - \mu_j}{\sigma_j}$ (centered step). Yu & Meng (2011) argued that by alternating between the non-centered and centered steps, we obtain a more efficient MCMC sampler that has a better mixing rate and converges faster. In addition, Kastner & Frühwirth-Schnatter (2014) showed that the ASIS can accurately sample latent volatility time-series that have low persistences, which is often the case for iEEG signals.

Parameter Identification

To demonstrate that the Gibbs sampler can accurately estimate the latent volatility process and its associated parameters, we conducted a simulation study in which we generated $N = 5$ time series of length $T = 50,000$ according to Eqn. 1 and 2 with various signal-to-noise ratios (SNR), which is controlled by varying the volatility of the volatility series, to mimic the typical length and the number of electrodes in our iEEG datasets. The SNR is calculated by taking the ratio of the average volatility of volatility across electrodes to the expected standard deviation of the noise term in Eqn. 4. We sampled 10,000 posterior draws and discarded the first 5,000 draws as a burn-in period to allow for convergence to the stationary distribution. Table 1 reports the posterior means of the parameters of the MSV model. Throughout the simulation, we use priors whose means are equal to the true values of the parameters. We observe that the Gibbs sampler can reliably estimate the parameters of the MSV model from datasets with various signal-to-noise ratios. The identification of the parameters in the MSV model comes from the strength of our large iEEG datasets which typically have tens of thousands of data points, an amount of data that rarely exists in financial applications.

Model Fit Plots

We provide a visualization of the latent volatility series. Fig. 5 illustrates the recordings from a hippocampal electrode during encoding of a list of 12 word items from a subject performing a verbal free-recall task at the University of Pennsylvania Hospital. The top panels show the detreneded iEEG series using AR(p) models. The bottom panels show the respective latent volatility series associated with the detrened signals. We observe that the latent volatility series capture the instantaneous variance of the original series.

Output Tables of Statistical Tests

This section reports the statistical tests for directional connection SME that contain at least 10 subjects. Tables 2 and 3 show the results of these tests for the left and right hemispheres.

References

Avants, B. B., Epstein, C. L., Grossman, M., & Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, *12*(1), 26–41.

Axmacher, N., Henseler, M. M., Jensen, O., Weinreich, I., Elger, C. E., & Fell, J. (2010). Cross-frequency coupling supports multi-item working memory in the human hippocampus. *Proceedings of the National Academy of Sciences*, *107*(7), 3228–3233.

Barndorff-Nielsen, O. E. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(2), 253–280.

Barry, R. J., Clarke, A. R., Johnstone, S. J., Magee, C. A., & Rushby, J. A. (2007). Eeg differences between eyes-closed and eyes-open resting conditions. *Clinical Neurophysiology*, *118*(12), 2765–2773.

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., ... Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of financial studies*, *9*(1), 69–107.

Benavoli, A., & Mangili, F. (2015). Gaussian processes for bayesian hypothesis tests on regression functions. In *Artificial intelligence and statistics* (pp. 74–82).

Blanchard, O., & Simon, J. (2001). The long and large decline in us output volatility. *Brookings papers on economic activity*, *2001*(1), 135–164.

Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., & Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: directional influences

revealed by granger causality. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(26), 9849–9854.

Burke, J. F., Ramayya, A. G., & Kahana, M. J. (2015). Human intracranial high-frequency activity during memory processing: neural oscillations or stochastic volatility? *Current Opinion in Neurobiology*, *31*, 104–110.

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological review*, *114*(2), 340.

Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., ... Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, *313*(5793), 1626–1628.

Caplan, J. B., Madsen, J. R., Schulze-Bonhage, A., Aschenbrenner-Scheibe, R., Newman, E. L., & Kahana, M. J. (2003). Human θ oscillations related to sensorimotor integration and spatial learning. *The Journal of Neuroscience*, *23*(11), 4726–4736.

Chauvette, S., Crochet, S., Volgushev, M., & Timofeev, I. (2011). Properties of slow oscillation during slow-wave sleep and anesthesia in cats. *The Journal of Neuroscience*, *31*(42), 14998–15008.

Cogley, T., & Sargent, T. J. (2005). Drifts and volatilities: monetary policies and outcomes in the post wwii us. *Review of Economic dynamics*, *8*(2), 262–302.

Davachi, L., Mitchell, J. P., & Wagner, A. D. (2003). Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proceedings of the National Academy of Sciences*, *100*(4), 2157 – 2162.

David, O., Guillemain, I., Saillet, S., Reyt, S., Deransart, C., Segebarth, C., & Depaulis, A. (2008). Identifying neural drivers with functional mri: an electrophysiological validation. *PLoS biology*, *6*(12), e315.

David, O., Kiebel, S. J., Harrison, L. M., Mattout, J., Kilner, J. M., & Friston, K. J. (2006). Dynamic causal modeling of evoked responses in eeg and meg. *NeuroImage*, *30*(4), 1255–1272.

Deshpande, G., LaConte, S., James, G. A., Peltier, S., & Hu, X. (2009). Multivariate granger causality analysis of fmri data. *Human brain mapping*, *30*(4), 1361–1373.

Dhamala, M., Rangarajan, G., & Ding, M. (2008). Analyzing information flow in brain networks with nonparametric granger causality. *Neuroimage*, *41*(2), 354–362.

Ekstrom, A. D., Caplan, J. B., Ho, E., Shattuck, K., Fried, I., & Kahana, M. J. (2005). Human hippocampal theta activity during virtual navigation. *Hippocampus*, *15*(7), 881–889.

Engle, R. F., Patton, A. J., et al. (2001). What good is a volatility model. *Quantitative finance*, *1*(2), 237–245.

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for bayesian models. *Statistics and Computing*, *24*(6), 997–1016.

Goebel, R., Roebroeck, A., Kim, D.-S., & Formisano, E. (2003). Investigating directed cortical interactions in time-resolved fmri data using vector autoregressive modeling and granger causality mapping. *Magnetic resonance imaging*, *21*(10), 1251–1261.

Goldman, R. I., Stern, J. M., Engel Jr, J., & Cohen, M. S. (2002). Simultaneous eeg and fmri of the alpha rhythm. *Neuroreport*, *13*(18), 2487.

Graham, S., Phua, E., Soon, C. S., Oh, T., Au, C., Shuter, B., ... Yeh, B. (2009). Role of medial cortical, hippocampal and striatal interactions during cognitive set-shifting. *Neuroimage*, *45*(4), 1359–1367.

Harvey, A. C., & Shephard, N. (1996). Estimation of an asymmetric stochastic volatility model for asset returns. *Journal of Business & Economic Statistics*, *14*(4), 429–434.

Herrmann, C. S., Munk, M. H., & Engel, A. K. (2004). Cognitive functions of gamma-band activity: memory match and utilization. *Trends in cognitive sciences*, *8*(8), 347–355.

Hesse, W., Möller, E., Arnold, M., & Schack, B. (2003). The use of time-variant eeg granger causality for inspecting directed interdependencies of neural assemblies. *Journal of neuroscience methods*, *124*(1), 27–44.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, *6*(2), 327–343.

Heston, S. L., & Nandi, S. (2000). A closed-form garch option valuation model. *The Review of Financial Studies*, *13*(3), 585–625.

Jacobs, J., & Kahana, M. J. (2010). Direct brain recordings fuel advances in cognitive electrophysiology. *Trends in cognitive sciences*, *14*(4), 162–171.

Jensen, O., & Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in cognitive sciences*, *11*(7), 267–269.

Kahana, M. J., Seelig, D., & Madsen, J. R. (2001). Theta returns. *Current Opinion in Neurobiology*, *11*(6), 739–744.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, *90*(430), 773–795.

Kastner, G., & Frühwirth-Schnatter, S. (2014). Ancillarity-sufficiency interweaving strategy (asis) for boosting mcmc estimation of stochastic volatility models. *Computational Statistics & Data Analysis*, *76*, 408–423.

Kim, S., Shephard, N., & Chib, S. (1998, July). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *The Review of Economic Studies*, *65*(3), 361–393.

Kirwan, C. B., & Stark, C. E. (2004). Medial temporal lobe activation during encoding and retrieval of novel face-name pairs. *Hippocampus*, *14*(7), 919–930.

Klimesch, W. (1999). Eeg alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Research Reviews*, *29*(2), 169–195.

Kreiman, G., Koch, C., & Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nature Neuroscience*, *3*, 946–953.

Landolt, H.-P., Dijk, D.-J., Achermann, P., & Borbély, A. A. (1996). Effect of age on the sleep eeg: slow-wave activity and spindle frequency activity in young and middle-aged men. *Brain Research*, *738*(2), 205–212.

Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., & Krakow, K. (2003). Eeg-correlated fmri of human alpha activity. *Neuroimage*, *19*(4), 1463–1476.

Long, N. M., Burke, J. F., & Kahana, M. J. (2014). Subsequent memory effect in intracranial and scalp eeg. *Neuroimage*, *84*, 488–494.

Long, N. M., & Kahana, M. J. (2015). Successful memory formation is driven by contextual encoding in the core memory network. *NeuroImage*, *119*, 332-337.

Maloney, K. J., Cape, E., Gotman, J., & Jones, B. (1997). High-frequency γ electroencephalogram activity in association with sleep-wake states and spontaneous behaviors in the rat. *Neuroscience*, *76*(2), 541–555.

Manning, J. R., Jacobs, J., Fried, I., & Kahana, M. J. (2009). Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *Journal of Neuroscience*, *29*(43), 13613–13620.

Nir, Y., Staba, R., Andrillon, T., Vyazovskiy, V. V., Cirelli, C., Fried, I., & Tononi, G. (2011). Regional slow waves and spindles in human sleep. *Neuron*, *70*(1), 153-169.

Omori, Y., Chib, S., Shephard, N., & Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, *140*(2), 425–449.

Raghavachari, S., Kahana, M. J., Rizzuto, D. S., Caplan, J. B., Kirschen, M. P., Bourgeois, B., ... Lisman, J. E. (2001). Gating of human theta oscillations by a working memory task. *The Journal of Neuroscience*, *21*(9), 3175–3183.

Rasmussen, C. E. (2004). Gaussian processes in machine learning. In *Advanced lectures on machine learning* (pp. 63–71). Springer.

Roebroeck, A., Formisano, E., & Goebel, R. (2005). Mapping directed influence over the brain using granger causality and fmri. *Neuroimage*, *25*(1), 230–242.

Sederberg, P. B., Kahana, M. J., Howard, M. W., Donner, E. J., & Madsen, J. R. (2003). Theta and gamma oscillations during encoding predict subsequent recall. *Journal of Neuroscience*, *23*(34), 10809–10814.

Sederberg, P. B., Schulze-Bonhage, A., Madsen, J. R., Bromfield, E. B., McCarthy, D. C., Brandt, A., ... Kahana, M. J. (2007). Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cerebral Cortex*, *17*(5), 1190–1196.

Solomon, E., Kragel, J., Sperling, M., Sharan, A., Worrell, G., Kucewicz, M., ... Kahana, M. (2017a). Widespread theta synchrony and high-frequency desynchronization

underlies enhanced cognition. *Nature Communications*, *8*(1), 1704. doi: 10.1038/s41467-017-01763-2

Solomon, E., Kragel, J., Sperling, M. R., Sharan, A., Worrell, G., Kucewicz, M., ... others (2017b). Widespread theta synchrony and high-frequency desynchronization underlies enhanced cognition. *Nature communications*, *8*(1), 1704.

Solomon, E. A., Stein, J. M., Das, S., Gorniak, R., Sperling, M. R., Worrell, G., ... Kahana, M. J. (2018). Functional wiring of the human medial temporal lobe. *bioRxiv*. Retrieved from `https://www.biorxiv.org/content/early/2018/01/31/257899`

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, *253*(5026), 1380.

Stephan, K., Penny, W., Moran, R., den Ouden, H., Daunizeau, J., & Friston, K. (2010). Ten simple rules for dynamic causal modeling. *NeuroImage*, *49*(4), 3099.

Stokes, P. A., & Purdon, P. L. (2017). A study of problems encountered in granger causality analysis from a neuroscience perspective. *Proceedings of the National Academy of Sciences*, *114*(34), E7063–E7072.

Von Stein, A., & Sarnthein, J. (2000). Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International journal of psychophysiology*, *38*(3), 301–313.

West, M. (1996). *Bayesian forecasting*. Wiley Online Library.

Wong, K. F. K., Galka, A., Yamashita, O., & Ozaki, T. (2006). Modelling non-stationary variance in eeg time series by state space garch model. *Computers in biology and medicine*, *36*(12), 1327–1335.

Yu, Y., & Meng, X.-L. (2011). To center or not to center: That is not the questionâĂŤan ancillarity–sufficiency interweaving strategy (asis) for boosting mcmc efficiency. *Journal of Computational and Graphical Statistics*, *20*(3), 531–570.

Yushkevich, P. A., Pluta, J. B., Wang, H., Xie, L., Ding, S.-L., Gertje, E. C., ... Wolk, D. A. (2015). Automated volumetry and regional thickness analysis of hippocampal subfields and medial temporal cortical structures in mild cognitive impairment. *Human brain mapping*, *36*(1), 258–287.

Zhou, Z., Ding, M., Chen, Y., Wright, P., Lu, Z., & Liu, Y. (2009). Detecting directional influence in fmri connectivity analysis using pca based granger causality. *Brain research*, *1289*, 22–29.

Table 1

Simulation Study.

			Truth							MSV						
Dataset	SNR	Channel	α	b_1	b_2	b_3	b_4	b_5	σ	α	b_1	b_2	b_3	b_4	b_5	σ
1	0.16	1	3.39	0.85	0.20	0.00	0.00	0.00	0.17	3.39	0.86	0.19	0.00	-0.04	0.00	0.17
		2	3.60	0.00	0.88	-0.10	0.00	0.00	0.19	3.60	0.00	0.87	-0.10	-0.00	-0.01	0.19
		3	3.55	0.00	0.00	0.87	0.30	0.00	0.19	3.53	-0.00	-0.00	0.85	0.29	0.01	0.20
		4	3.51	0.00	0.00	0.00	0.71	0.00	0.12	3.51	0.01	-0.00	0.01	0.69	-0.03	0.14
		5	3.38	0.00	0.00	0.00	0.00	0.80	0.15	3.38	0.00	-0.01	-0.01	0.06	0.79	0.15
2	0.27	1	3.60	0.95	0.20	0.00	0.00	0.00	0.21	3.93	0.95	0.19	-0.00	0.00	0.00	0.24
		2	3.74	0.00	0.90	-0.10	0.00	0.00	0.25	3.82	-0.00	0.90	-0.10	0.00	-0.00	0.25
		3	3.89	0.00	0.00	0.93	0.30	0.00	0.29	3.81	-0.00	0.00	0.93	0.30	0.00	0.29
		4	3.05	0.00	0.00	0.00	0.91	0.00	0.28	3.04	-0.00	0.01	0.01	0.91	0.00	0.28
		5	3.96	0.00	0.00	0.00	0.00	0.95	0.22	3.97	-0.00	0.00	0.00	-0.00	0.94	0.22
3	0.42	1	3.25	0.52	0.20	0.00	0.00	0.00	0.42	3.25	0.51	0.20	0.01	-0.02	0.01	0.43
		2	3.64	0.00	0.57	-0.10	0.00	0.00	0.46	3.63	0.01	0.55	-0.14	0.02	0.02	0.46
		3	3.18	0.00	0.00	0.58	0.30	0.00	0.31	3.18	0.03	0.01	0.56	0.32	0.03	0.31
		4	3.44	0.00	0.00	0.00	0.65	0.00	0.41	3.45	0.00	-0.00	-0.00	0.64	0.00	0.41
		5	3.26	0.00	0.00	0.00	0.00	0.60	0.36	3.26	-0.02	-0.00	0.00	-0.01	0.60	0.36

We generated 3 datasets with different signal-to-noise ratios. The observed multivariate time-series $\mathbf{y}_t$ was simulated according to the data-generating process specified by the MSV model with pre-specified parameters (truth). We then applied the MSV model to the simulated series $\mathbf{y}_t$ to recover the parameters of the MSV model. In this simulation study, the non-zero off-diagonal entries of the matrix B were fixed across datasets. The diagonal elements of B were generated from a uniform distribution on $[0.7, 0.9]$, $[0.9, 1.0]$, and $[0.5, 0.7]$ respectively. The volatilities of volatility of the electrodes were generated from a uniform distribution on $[0.1, 0.2]$, $[0.2, 0.3]$, and $[0.3, 0.5]$ respectively.

Table 2

intra-MTL directional connectivity in the left hemisphere

region I → region J	mean $\Delta_{I\to J}$	se	t	N	p	adj. p
Hipp → PRC	-0.044	0.013	-3.494	48	0.001	0.009**
PRC → Hipp	-0.060	0.022	-2.667	48	0.011	0.045*
Hipp → EC	0.010	0.032	0.312	14	0.768	0.953
EC → Hipp	-0.077	0.067	-1.158	14	0.285	0.569
PHC → PRC	-0.006	0.040	-0.146	16	0.889	0.953
PRC → PHC	-0.037	0.028	-1.348	16	0.212	0.564
PRC → EC	0.001	0.022	0.061	21	0.953	0.953
EC → PRC	0.005	0.030	0.168	21	0.872	0.953

Table 3

intra-MTL directional connectivity in the right hemisphere

region I → region J	mean $\Delta_{I\to J}$	se	t	N	p	adj. p
Hipp → PRC	-0.010	0.020	-0.471	40	0.645	0.838
PRC → Hipp	-0.016	0.029	-0.575	40	0.574	0.838
Hipp → EC	0.011	0.030	0.361	14	0.733	0.838
EC → Hipp	0.054	0.047	1.144	14	0.290	0.838
Hipp → PHC	-0.027	0.032	-0.837	15	0.432	0.838
PHC → Hipp	0.044	0.035	1.232	15	0.254	0.838
PRC → EC	0.020	0.052	0.378	14	0.722	0.838
EC → PRC	-0.002	0.081	-0.020	14	0.985	0.985

$*$: P$<$ 0.05, $**$: P$<$ 0.01, adjusted p-values were calculated using the Benjamini-Hochberg procedure.

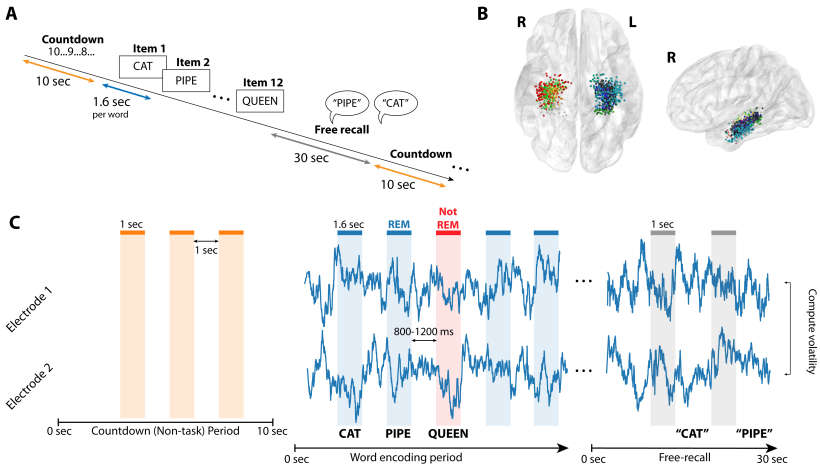

Figure 1. Figure reproduced from E. A. Solomon et al. (2018). Task Design and Analysis. (A) Subjects performed a verbal free-recall which consists of three phases: (1) countdown (orange), (2) word encoding (blue), and (3) free recall (gray). (B) 96 Participants were implanted with depth electrodes in the medial temporal lobe (MTL) with localized subregions: CA1, CA3, dentate gyrus (DG), subiculum (Sub), perirhinal cortex (PRC), entorhinal cortex (EC), or parahippocampal cortex (PHC). (C) For each subject, we applied the MSV model to the iEEG time-series recorded during the word encoding period to estimate the latent volatility series and the parameters of the MSV model.

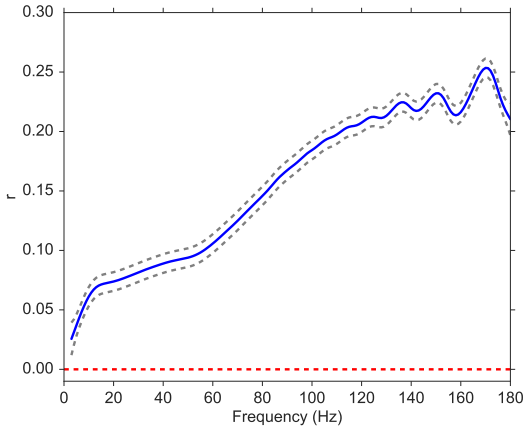

Figure 2. Correlation between volatility and spectral power over a frequency range from 3 to 180 Hz. We fitted a Gaussian process model to estimate the functional form of the correlation function between volatility and spectral power (solid blue line) and its 95% confidence bands (dashed gray lines). The red line shows the null model. We observe a significantly positive correlation between volatility and spectral power, and the correlation increases with frequency.

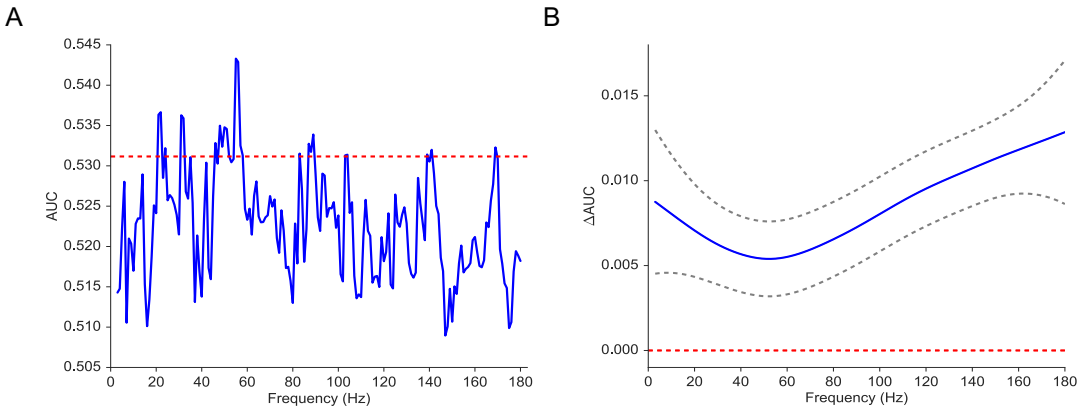

Figure 3. (A) Average AUC of the classifier trained on spectral power across 45 subjects with at least 3 sessions of recording (blue). The red line indicates the average AUC of the classifier trained on volatility. (B) ΔAUC = AUC$_{vol}$ - AUC$_{power}$ as a function of frequency estimated by using a Gaussian regression model (dashed gray lines indicate 95% confidence bands). The red line shows the null model. We observe that the classifier trained on volatility performs at least as well as the one trained on spectral power across the frequency spectrum. We find that functional form of ΔAUC is significantly different from the 0 function ($\chi^2_{11} = 42$, P $< 10^{-5}$) using a Gaussian process model, suggesting that the difference in performance between the volatility classifier and the spectral power classifier is significant.

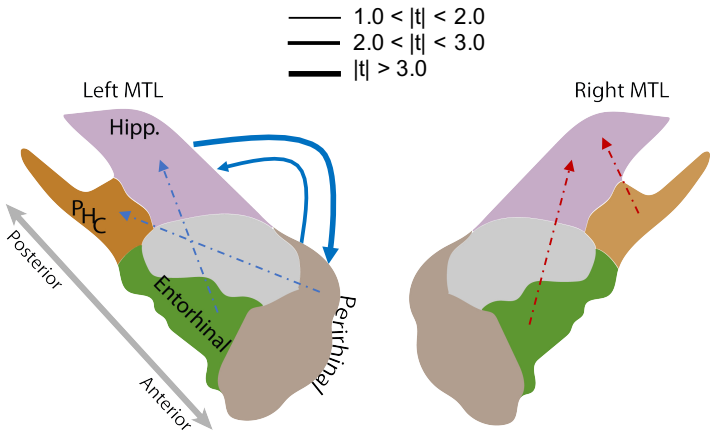

Figure 4. MTL Directional Connectivity Network. The MTL electrodes were divided into four subregions: hippocampus (Hipp.), parrahippocampal cortex (PHC), entorhinal cortex (EC), and perirhinal cortex (PRC). The directional connectivity from region I to region J, $C_{I\to J}$, was calculated by averaging the entries of the sub-matrix of the regression coefficient matrix B, whose rows and columns correspond to region I and J respectively. We computed the contrast bewteen the directional connectivity of recalled and non-recalled events: $\Delta_{I\to J} = C^{R}_{I\to J} - C^{NR}_{I\to J}$ for each subject. Solid lines show significant (FDR-corrected) connections between two regions and dashed lines show trending but insignificant connections. Red indicates positive changes and blue indicates negative changes. The directional connectivity from Hipp. to PRC is significant (adj. $P < 0.01$) and the reverse directional connectivity is also significant (adj. $P < 0.05$).

.

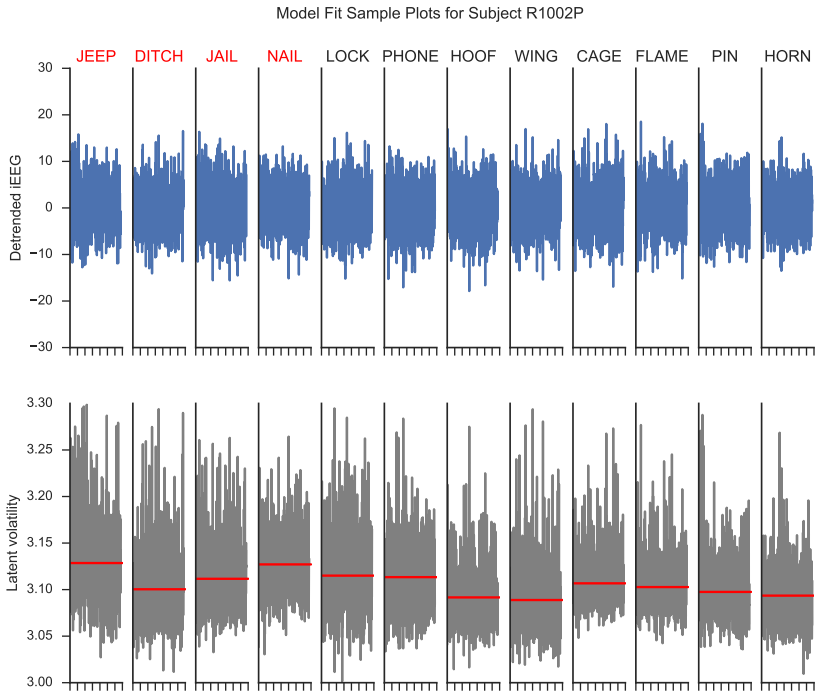

Figure 5. Model fit plots for a hippocampal electrode. The upper panels show the detrended iEEG signals using an AR(p) model for encoding periods of a list of words. The lower panels show the associated estimated latent volatility processes. The red lines indicate the average volatility during the encoding period. Red words are later recalled and blue words are not recalled.