A Bayesian Spatiotemporal Model for Multi-Subject EEG Source Dynamics and Effective Connectivity

Wesley Thompson[1,2*], Tim Mullen[1,3], Julie Onton[1], Scott Makeig[1]

[1] Swartz Center for Computational Neuroscience/INC, [2] Dept. of Psychiatry, [3] Dept. of Cognitive Science, University of California, San Diego

* wes.stat@gmail.com

Abstract

EEG has millisecond temporal resolution, necessary for analysis of transient cortical dynamics. However, the poor spatial resolution of scalp EEG combined with the confounding effects of volume conduction and non-brain artifacts complicates interpretation of neural dynamics when examined at the level of scalp electrodes. Accurate localization of sources of EEG activity is a difficult, ill-posed problem. One approach is to apply Independent Component Analysis (ICA) to scalp EEG recordings to obtain time courses and scalp maps of maximally-independent sources of EEG activity with projections resembling single or dual symmetric equivalent dipoles. These sources can then be localized using appropriate forward and inverse models, while adaptive vector autoregressive models may be fit to the source time series to model transient information flow. Applied to different subjects, this typically results in varying numbers and locations of source dipoles across subjects which complicates efforts to obtain robust group-level statistics. Here we develop a Bayesian spatiotemporal model for multi-subject source-localized EEG which provides group inferences on the spatial locations and causal relationships among localized sources. Each subject's localized sources are modeled as arising from a mixture distribution of spatial coordinates and time-varying multivariate granger causality. Model inference is obtained via a Markov Chain Monte Carlo algorithm. This approach can be generalized to other non-ICA approaches for separation and localization of dipolar sources, such as beamforming. The utility of this method is initially demonstrated by application to a large multi-subject EEG dataset, where we examine network dynamics underlying error commission in an ERN-producing task.

Theory

MULTI-SUBJECT INFERENCE

Group-level inferences of multi-subject source-localized (dipolar) independent components (ICs) can be problematic. Two or more subjects performing the same task may end up with differing numbers of retained IC sources and different source locations. Thus, unlike scalp channel recordings or region-of-interest (ROI) source analysis there is an inherent uncertainty in matching IC sources across subjects, and therefore in obtaining reliable group inferences regarding functional connectivity between these sources. While various disjoint clustering methods can be used to identify similar sources across subjects, these methods often suffer from poor statistical properties as the number of missing variables increases. It is thus preferable to employ a method which propagates uncertainty regarding source identification to inferences regarding group-level effective connectivity (EC) estimates.

MIXTURE MODEL

Let the M_i x T matrix of IC time series for the i^{th} subject be denoted by $\mathbf{Y}_i$ and let the estimated M_i corresponding dipole spatial locations be denoted by $\mathbf{S}_i$. Let M be the total number of group-level sources (e.g. clusters or ROIs) under consideration for all N subjects (currently M is chosen by a preliminary *a priori* decision, though in the future we will determine this automatically within the inferential framework as in [1]).

The data are modeled as coming from a mixture distribution [2]. To implement this mixture model, for each subject we augment the observed data $\{\mathbf{Y}_i, \mathbf{S}_i\}$ with an M_i x M matrix of latent indicators $\mathbf{Z}_i$. The j^{th} row of $\mathbf{Z}_i$ consists of zeros with exactly one entry equal to one in column k: this indicates that the j^{th} source for subject i corresponds to the k^{th} cluster.

Conditional on the $\mathbf{Z}_i$ we assume

$$\Pr(\mathbf{S}_i \mid \mathbf{Z}_i) = \Pi_{j=1:Mi}\Pi_{k=1:M}\left[N(\mathbf{S}_{ij} \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right]^{zijk} \qquad (1)$$

In addition to the spatial information, we want to incorporate information regarding the dynamics of the source-localized times series $\mathbf{Y}_i$ into the mixture model. Suppose we summarize the EC information contained in the source time series $\mathbf{Y}_i$ via time-varying EC estimates $\mathbf{F}_i(t)$. We include this information in the mixture model via

$$\Pr(\mathbf{F}_i \mid \mathbf{Z}_i) = \Pi_{j=1:Mi}\Pi_{k=1:M}\left[\Pr(\mathbf{F}_i \mid \boldsymbol{\beta}_k) \right]^{zijk} \qquad (2)$$

where $\boldsymbol{\beta}_k$ are parameters which determine the distribution of $\mathbf{F}_i$ conditional on $\mathbf{Z}_i$.

In the following example, we obtain $\mathbf{F}_i(t)$ by computing the graph-normalized Direct Directed Transfer Function (dDTF, [3]) – a frequency-domain measure of multivariate Granger-causal relationships – for each pair of IC sources. We obtain time-varying dDTF estimates using a sliding-window vector autoregressive (VAR) model with a 500 ms window length and 30 ms step size producing 80 time points. The dDTF is integrated over the theta band (3-7 Hz) and modeled as a smooth function of time via a penalized *B*-spline; the $\boldsymbol{\beta}_k$ are the fixed effects group level of the coefficients $\mathbf{b}_i$ of the spline basis functions.

$$\Pr(\mathbf{b}_i \mid \mathbf{Z}_i) = N(\mathbf{b}_i \mid \boldsymbol{\beta}_k, \sigma^2_{bk}\mathrm{I}) \qquad (3)$$

BAYESIAN INFERENCE

We place Dirichlet($\boldsymbol{\alpha}_i$) prior distributions on the allowable patterns of the latent indicator matrices $\mathbf{Z}_i$. Along with the augmented likelihood derived from Eq. (1)-(3), we complete the Bayesian specification of the model by placing Inverse Wishart (IW) prior distributions on the $\boldsymbol{\Sigma}_k$, Inverse Gamma distributions on the σ^2_{bk}, and diffuse normal distributions on the $\boldsymbol{\mu}_k$ and $\boldsymbol{\beta}_k$.

Model inference proceeds via a Markov Chain Monte Carlo (MCMC) algorithm. Full conditional posterior distributions are standard. In particular, the allowable patterns of the indicator matrices $\mathbf{Z}_i$ are multinomial. Allowable patterns have exactly one nonzero element in each row and at most one non-zero element in each column. Since the number of allowable patterns is too large to sample from directly (in general $M_i ! / (M\text{-}M_i)!$), we sample a subset of the allowable patterns at each iteration of the MCMC algorithm as follows:

(i) For subject i, randomly sample two distinct row indices j_1 and j_2 between 1 and M_i.
(ii) Keeping all other indices fixed, compute the conditional posterior over all allowable patterns permuting the column indices k_1 and k_2 for which the j_1 and j_2 rows are nonzero.
(iii) Sample $\mathbf{Z}_i$ from the conditional distribution keeping other rows fixed.

Sampling of all other parameters conditional on the latent data $\mathbf{Z}_i$ is straightforward.

Data

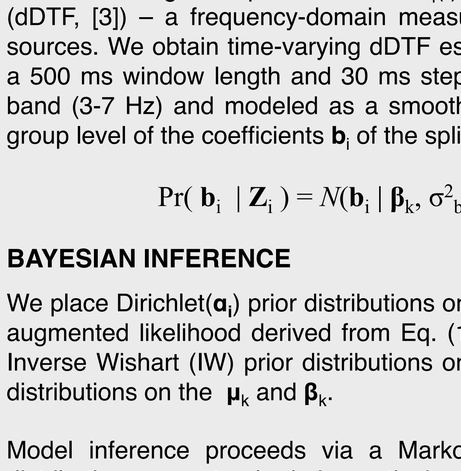

128-channel (256 Hz) EEG data were collected from 24 subjects performing a visual letter two-back task with auditory feedback [4]. Trials were segregated based on response type (Incorrect vs. Correct). Following zero-phase FIR high-pass filtering (1 Hz), response-locked datasets were subjected to Infomax Independent Component Analysis (ICA). ICA is effective at separating source components that are maximally instantaneously independent, which can be further analyzed for transient dependencies [5]. A single (or dual symmetric) equivalent dipole model was then fit to each independent component (IC) using EEGLAB's DIPFIT2 function. We rejected ICs corresponding to artifacts such as eye blinks and muscle activity, and those with a poor dipole fit (> 15% r.v., or lying outside brain volume).

Results

Below: All retained dipolar sources from all subjects color-coded by cluster membership (k-means clustering).
Right: Frames from a Causal BrainMovie3D showing group inferences on source locations and effective connectivity (integrated over 3-7 Hz theta band) as obtained from Bayesian mixture model. K-means cluster centroids and spatial dispersion used as initial values and priors for MCMC algorithm. A cluster is retained if more than 33% of subjects have greater than 50% probability of cluster membership. We retain only connectivity that deviates significantly from [-750 -500 ms] baseline (p<0.01).

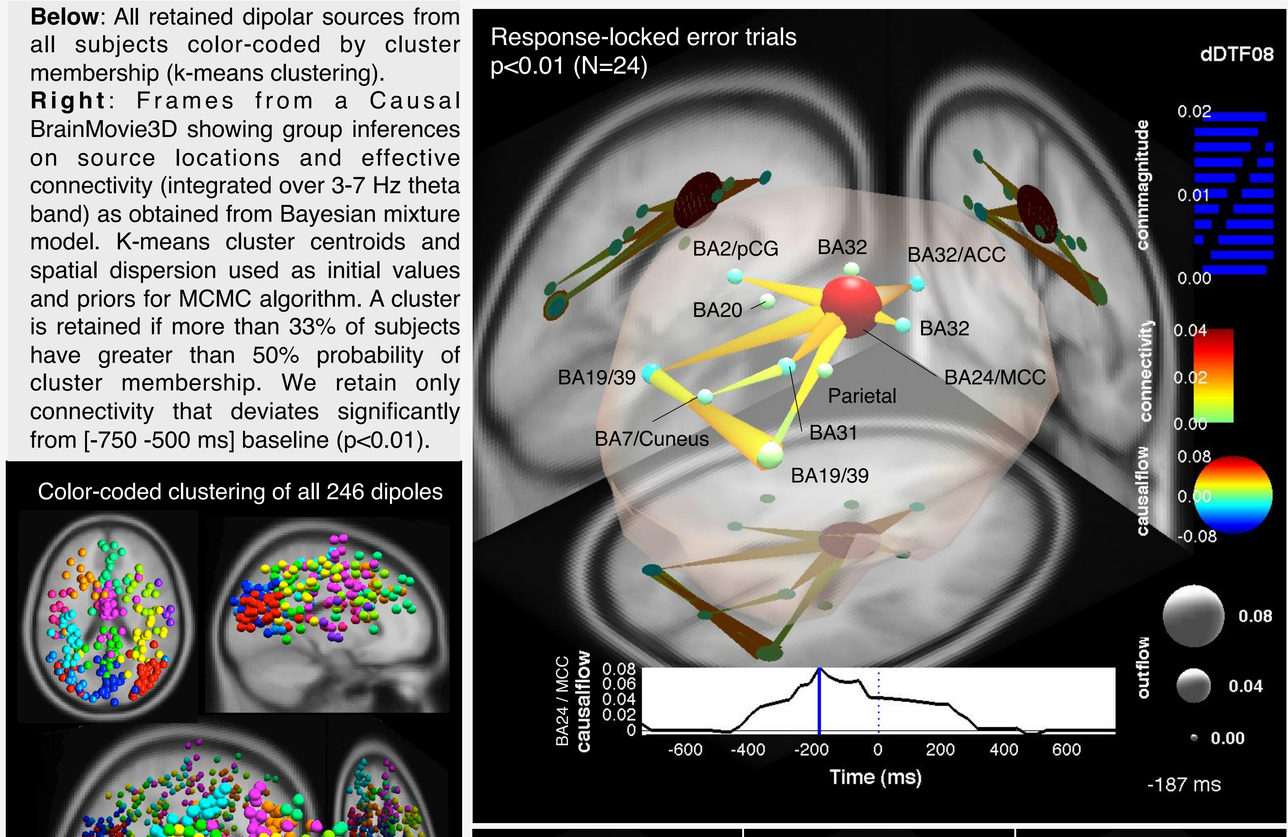

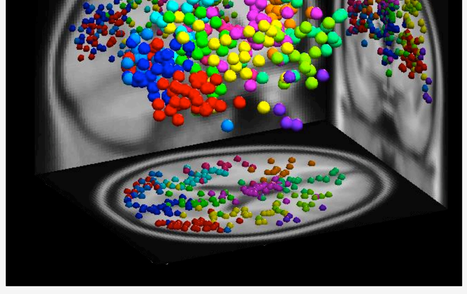

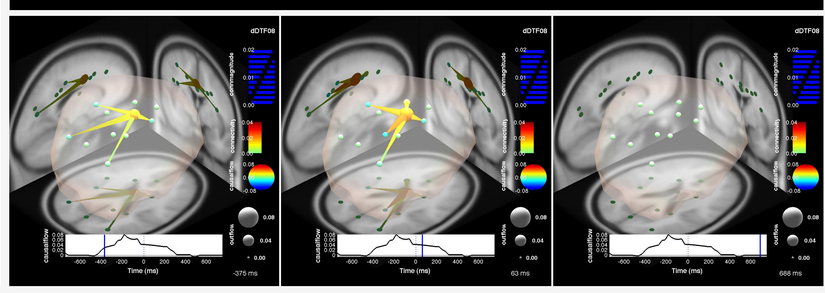

Below: Time-varying theta-band (3-7 Hz) dDTF group-level inferences with 99% confidence intervals. Mean source locations with Talairach coordinates and anatomical designations (Talairach Daemon) are shown on the marginals. Translucent regions indicate time intervals that deviate significantly from the [-750 to -500] ms baseline (p<0.01, uncorrected). Note the significant outflow from a source in the dorsal middle cingulate cortex (BA24/MCC) – likely the rostral cingulate zone (RCZ) – immediately before, during, and following responses made in error. All connectivity analysis and visualizations are produced using the EEGLAB-compatible Source Information Flow Toolbox [7].

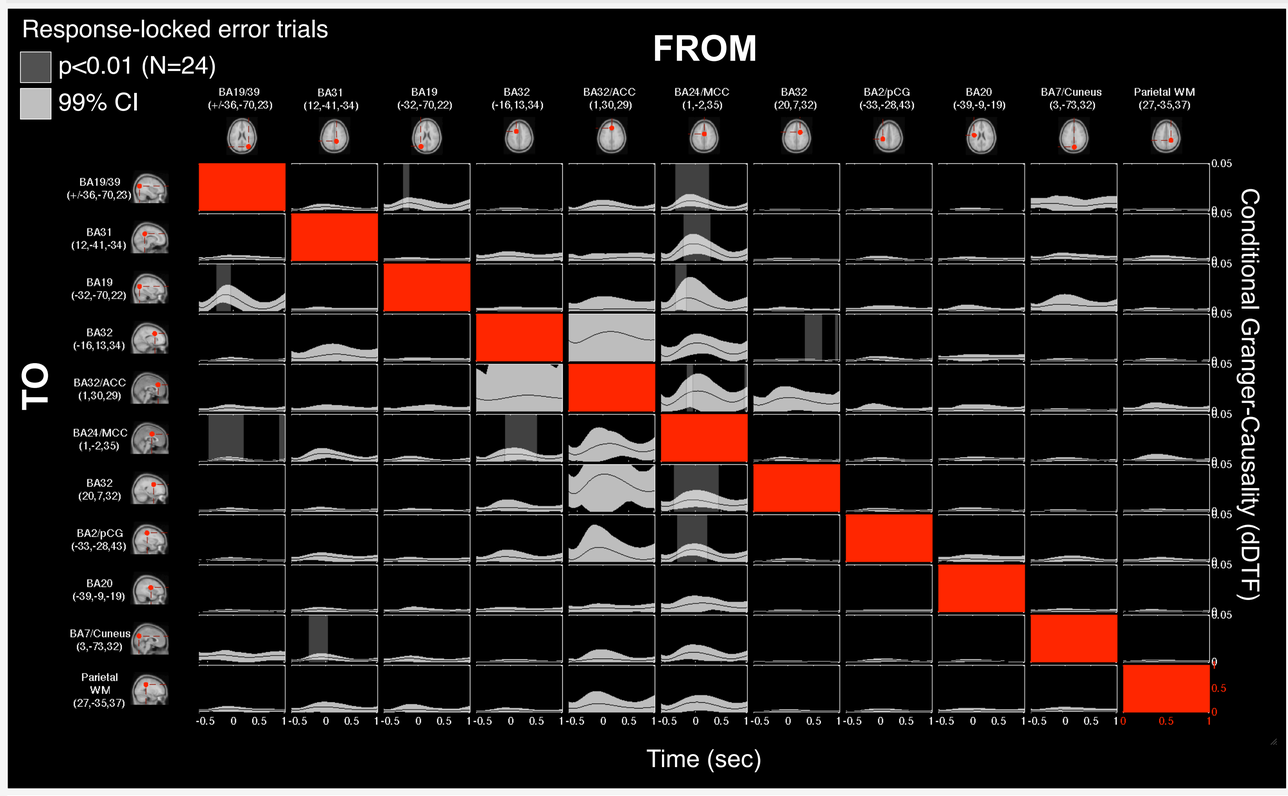

Conclusions and Future Work

We have demonstrated a preliminary application of a novel Bayesian spatiotemporal model for obtaining group-level inferences and confidence intervals on expected dipolar source locations and dynamics (e.g. connectivity). In this application we demonstrated the emergence of statistically significant causal relationships between dorsal MCC and several cortical and cingulate structures during error commission. This is commensurate with theoretical and experimental evidence for a significant causal role of MCC in error processing [6]. We realize this model represents a first step which can be further improved upon. We are currently working on expanding this to the 2D time-frequency plane using a tensor product of 1D splines (allowing different degrees of smoothing across time or frequency). We also plan to use a Dirichlet hyperprior to automatically select the optimal number of clusters as in [1]. The method also can be adapted to gracefully handle outliers, which should help improve the confidence interval estimates from those shown above. Finally, while in this example prior distribution parameters were determined using mean and covariance information from an initial k-means clustering step, it is straightforward to incorporate biologically-plausible priors for source locations and dispersion, which can be determined via existing numerical simulation data as well as task-specific prior expectations. The method can also be extended naturally to modeling statistical interactions between multiple experimental conditions via hierarchical modeling, which is a current avenue of research for us.

Once fully developed, we expect this approach will have a significant impact on the ability to flexibly obtain robust group-level inferences and statistics on the spatiotemporal dynamics and/or interactions of point-process (dipolar) sources. The approach may also have utility when used with distributed source localization algorithms and we are currently exploring the use of source spatial distributions obtained from Sparse Bayesian Learning.

References

[1] Ishwaran, H. and James, L., (2002) *JCGS* [2] Fruhwirth-Schnatter, S., (2006) Finite Mixture and Markov Switching Models. Springer, NY, NY. [3] Korziniewska, et al, (2008) *Human Brain Mapping* [4] Onton, J. and Makeig, S (2007) *Society for Neuroscience* [5] Makeig, S., et al (2002) *Science* [6] Shackman et al (2011) *Nature Reviews* [7] Mullen, T., et al (2010) www.sccn.ucsd.edu/wiki/SIFT

Support: San Diego Fellowship (TM), Glushko Fellowhsip (TM), ONR, NIMH, NINDS, and The Swartz Foundation (Old Field, NY)