Inferring Brain Signals Synchronicity from a Sample of EEG Readings

Qian Li[1*]
Damla Şentürk[1,2],
Catherine A. Sugar[1,2,3],
Shafali Jeste[3],
Charlotte DiStefano[3],
Joel Frohlich[3],
and
Donatello Telesca[1,*].
[1] Department of Biostatistics, University of California, Los Angeles
[2] Department of Statistics, University of California, Los Angeles
[3] Department of Psychiatry and Biobehavioral Sciences,
University of California, Los Angeles
[*] *email: dtelesca@ucla.edu*

September 29, 2016

*This work was supported by the grant R01 GM111378-01A1 (DS, DT, CS) from the National Institute of General Medical Sciences.

Abstract

Inferring patterns of synchronous brain activity from a heterogeneous sample of electroencephalograms (EEG) is scientifically and methodologically challenging. While it is intuitively and statistically appealing to rely on readings from more than one individual in order to highlight recurrent patterns of brain activation, pooling information across subjects presents non-trivial methodological problems. We discuss some of the scientific issues associated with the understanding of synchronized neuronal activity and propose a methodological framework for statistical inference from a sample of EEG readings. Our work builds on classical contributions in time-series, clustering and functional data analysis, in an effort to reframe a challenging inferential problem in the context of familiar analytical techniques. Some attention is paid to computational issues, with a proposal based on the combination of machine learning and Bayesian techniques.

Keywords: Consensus Clustering, EEG, Hierarchical Mixture Models, Spectral Clustering.

1 Introduction

Functional neuroimaging technologies, including MRI, PET, MEG, and EEG, aim to measure different aspects of brain function as they relate to specific mental processes. This article focuses on the analysis of Electroencephalography (EEG) data in the context of neuropsychology studies. EEG is a well-established noninvasive method for measuring spontaneous and event-related electrical activity across brain regions. The technology captures voltage fluctuation as signals, which reflect the distributed neuronal activities being projected on a cortical patch on which an EEG sensor is placed (Teplan 2002). The general aim of an EEG study is often the identification of neural function and cognitive states. Diverse biomedical applications include epilepsy, sleep disorders, multiple sclerosis, brain tumors, lesions, schizophrenia, and mood disorders (Teplan 2002).

Typical analyses in EEG studies focus primarily on inferring group differences in regions of interest. Such differences are assessed both in the frequency domain, by means of an amplified Signal-to-Noise Ratio (SNR) (Laufs et al. 2003), and, in the case of studies involving external stimuli, in the time domain, by means of averaging and smoothing over repeated applications of the stimuli (Hasenstab et al. 2015).

Beyond differential activation of brain regions, mounting evidence is building a case for the deeper understanding of neural interactions (Di Martino et al. 2014, Craddock et al. 2013). In this setting, magnetic resonance imaging has become an established workhorse for the mapping and annotation of the human connectome at the macro-scale. The key to the success of MRI technologies as a preferred measurement tool in functional connectivity studies lies in their ability to produce measurements at high spatial resolution. This ability comes, however, at the cost of low time resolution, and perhaps most importantly, at the cost of severe hardware limitations, which make MRI studies hard to design in a logistically

and financially feasible fashion.

On the other end, EEG is thought to provide reliable measurements of neuronal activity only for the brain cortical regions, with low spatial resolution and often low SNR. However, compared to other imaging techniques, EEG has the advantage of relying on less bulky hardware and is associated with robust and extremely non-invasive imaging protocols, making the technology readily available for implementation and adaptation to a variety of scientific investigations.

In a seminal contribution, Euan et al. (2015) suggested exploiting EEG's excellent temporal resolution by defining the concept of spectral synchronicity. In particular, a pair of EEG signals are considered spectrally synchronized if they are both dominated by similar frequency oscillations. This idea formalizes the concept of coordinated neuronal activity and reflects recent empirical evidence, which suggests that differential patterns of coordinated neuronal activity may be associated with a range of neuropsychiatric and neurological processes, including memory formation (Fell and Axamcher 2011) and mental disorders (Broyd et al. 2008).

From a statistical perspective, multi-subject studies of functional connectivity still pose substantial methodological challenges. Ideally, statistical inference should provide tools for the understanding of typical functional connectivity patterns, as well as quantification of familiar concepts like sample and population variability, and dependence on clinical phenotypes via regression. Even though some progress in the direction of population level inference has recently been made in the context of fMRI data (Narayan and Allen 2015; Shou et al. 2014), typical analyses are still reliant on untenable assumptions of time-independence. The literature is, in fact, substantially silent on the subject of population level connectivity inference using EEG data. In this work, we aim to solve this problem and introduce a simple and interpretable technique for the analysis of brain synchronicity

from a sample of EEG readings.

Inferring patterns of synchronous brain activity from a heterogeneous sample of EEG measurements is indeed scientifically and methodologically challenging. While it is intuitively and statistically appealing to rely on readings from more than one individual in order to highlight recurrent patterns of brain activation, pooling information across subjects presents non-trivial technical problems. In particular, brain signals variation between subjects is expected to be highly volatile, with little to no experimental control for inter-subject synchronization. This observation defines a natural expectation for EEG response patterns to be highly specific to the sensor-by-subject experimental unit, rendering classical assumptions for statistical models of pooled signals unlikely to be satisfied in naïve applications.

Our approach is based on the definition of cortical maps, identifying areas of synchronous neuronal activity specific to individual subjects and experimental epochs, intended as time intervals. Synchronized cortical regions are estimated via a mixture model of eigen-Laplacian vectors, obtained from appropriately constructed dissimilarity matrices. As the experiment evolves in time, subject and time-specific cerebral maps form a longitudinal ensemble. In this context, we posit that pooled information, within and between subjects, is amenable to statistical analysis via a hierarchical model involving mixture probabilities (Lock and Dunson 2013), which we call Multilevel Integrative Clustering (MIC). Our framework supports both the definition of coordinated neuronal activity via a mixture approach, and the formulation of probability statements describing inter-subject and intra-subject variability via the familiar toolset of hierarchical modeling.

Our manuscript is organized as follows. In Section 2 we describe a general framework for integrative clustering at the epoch, subject and population levels. In Section 3 we assess the operative characteristics of our proposed approach through experiments on engineered

data. In Section 4 we apply the proposed framework to the analysis of a resting-state EEG study on typically developing (TD) children and children diagnosed with Autism Spectrum Disorder (ASD). We conclude with a critical discussion and potential extensions in Section 5.

2 Multilevel Integrative Clustering (MIC)

In the following discussion we proceed to characterize coordinated neuronal activity via time-varying pairwise distances between the time series associated with a set of EEG sensors or electrodes. Our approach builds on Euan et al. (2015), who define synchronicity in relation to pairwise similarities between the power spectral densities of electrode-level signals. In §2.1, we describe a data meta-processing step aimed at obtaining stable time-varying estimates of the EEG spectral profiles. In §2.2, instead of directly operating on spectral densities, we model a set of related d-dimensional eigen-Laplacians via a multilevel model for clustering areas of synchronous neuronal activation. Inferential and computational details are discussed in §2.3 and §2.4.

2.1 From EEG Signals to eigen-Laplacian Matrices

Let $\{Z_\tau,\ \tau = 0, \pm 1, \pm 2, \ldots\}$ be a zero mean, weakly stationary time series, with auto-covariance $\mathcal{C}_Z(h) = E(Z_\tau, Z_{\tau+h})$, $(h = 0, \pm 1, \pm 2, \ldots)$. The second order properties of the series may be described by the spectral density function $\phi_Z(\omega)$ of Z_τ as in Brillinger (1981), so that:

$$\phi_Z(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \mathcal{C}_Z(\tau) \exp(-i\tau\omega), \quad \omega \in [0, \pi].$$

Intuitively, $\phi_Z(\omega)$ may be interpreted as the variance contributed to the entire series by oscillations in a narrow frequency band around $\omega \in [0, \pi]$. The spectral analysis of neural signals is an important workhorse in EEG studies, as frequency bands are thought to be associated with specific cognitive, perceptive and cellular phenomena (Teplan 2002).

EEG time-series signals are usually collected in relation to a geodesic net of p electrodes. Upon collection, raw signals are segmented into 1024ms time intervals for EEG preprocessing, which typically includes bandpass filtering, electrode and segments rejection, and artifacts inspection. Similar pipelines are common for EEG analysis, which can improve the SNR for spectral analysis (Bigdely-Shamlo et al. 2015).

Let $i = 1, \ldots, n$ index n study subjects, $j = 1, \ldots, p$ index p EEG electrodes, and $s_{\ell i}$, $\ell_i = 1, \ldots, q_i$, index q_i-1024ms segments retained after data quality control. The filtered EEG data can be seen as an ensemble of time-series segments $Y_{ij}(s_{\ell_i})$, each composed of a number of measurements reflective of analog-to-digital sampling rates, typically 256/512Hz.

We are interested in the time-dynamics of neuronal synchronicity through a notion of time-varying spectral density via local stationarity (Florian and Pfurtscheller 1995; Rosen et al. 2012). In our formulation we fully acknowledge common pre-processing practices, which sees qualifying EEG segments being concatenated and re-referenced without time labelling. This practice typically leads to latent gaps in the post-processed series, providing a non-standard inferential framework for time-varying spectral estimation.

In order to obtain time-varying stable estimates of electrode-specific spectra, we operate on a combined set of γ adjacent segments $(s_{\ell_i}, \ldots, s_{(\ell_i+\gamma)})$, which we define as epochs. Furthermore, adjacent epochs smooth over the original time domain by overlapping over a $\delta \in (0, 1)$ fraction of segments. For each subject i, electrode j and epoch $t \in \{1, 2, \ldots, T_i\}$, we obtain estimates $\widehat{\phi}_{ij}(\omega, t)$ of the epoch-specific spectral density by averaging segment specific spectral density estimates obtained as in Ombao et al. (2001). The details of this

procedure are reported in a supplementary document. Our approach stems from the idea introduced by Hasenstab et al. (2015) in the context of time-domain analyses. The use of overlapped sliding windows in the estimation of a time-dependent power spectral density mediates between the need for stable estimates and the potential for non-stationarity over the entire duration of the study. A study of inferential robustness to smoothing choices is reported in §3.

Following the approach by Euan et al. (2015), desynchronicity is measured by *total variation distance* (TVD) between a pair of spectral densities estimated at each epoch, so that, for subject i, desynchronicity between electrode j and electrode k at epoch t is defined as:

$$d_{it}(\widehat{\phi}_{ij}, \widehat{\phi}_{ik}) = 1 - \int \min\{\widehat{\phi}_{ij}(\omega, t),\, \widehat{\phi}_{ik}(\omega, t)\}\, \mathrm{d}\omega.$$

For each subject and epoch, these pairwise distances produce a $p \times p$ dissimilarity matrix $D_i(t) = \left[d_{it}(\widehat{\phi}_{ij}, \widehat{\phi}_{ik})\right]$, summarizing information on differential synchronicity between the p electrodes from different cortical regions.

Before clustering, each matrix is represented in the eigen-space spanned by the largest d eigenvectors of the graph-Laplacian associated with an affinity matrix $A_i(t) = 1 - D_i(t)$. More precisely, we take a *graph cuts* view of clustering and construct a normalized graph-Laplacian $\mathcal{G}_i(t) = diag\left[A_i(t)\mathbf{1}_p\right]^{1/2} A_i(t)\, diag\left[A_i(t)\mathbf{1}_p\right]^{1/2}$, representing a weighted undirected graph between EEG electrodes. In this setting, we follow (Ng et al. 2001) and summarize the information in $\mathcal{G}_i(t)$ with its largest d eigenvectors $X_i(t) \in \mathbb{R}^{p \times d}$.

This strategy is intuitively motivated by the analysis of the isolated connected components "*ideal case*", in which $A_{jk}(t) > 0$ iff components j and k belong to the same cluster, and $A_{jk}(t) = 0$ otherwise. In this simplified setting, considering K clusters, the first K columns of $X_i(t)$ have non-zero elements corresponding to connected components in $A_i(t)$. Row-wise, $X_i(t)$ is piece-wise constant, suggesting K-means as a simple clustering rule to

recover the connected components.

In reality, we work under the assumption that $\mathcal{G}_i(t)$ is a perturbation of the "*ideal case*" and in §2.2 we exploit this intuition to develop model-based clustering of electrodes at the epoch, subject and population level. Crucially, we avoid using a mixture model of spectral densities; instead model-based clustering of EEG signals over potentially non-convex manifolds is achieved using simpler location/scale-mixture models involving vectors in $\mathbb{R}^d$.

2.2 Hierarchical Mixture Priors and Multilevel Inference

Let $X_{ij}(t) \in \mathbb{R}^d$, be a d-dimensional eigen-Laplacian vector associated with the EEG signal for subject i, $(i = 1, 2, \ldots, n)$; electrode j, $(j = 1, 2, \ldots, p)$; at epoch $t = 1, 2, \ldots, T_i$. In practice, we observe subject-specific epochs t_{im_i}, $(m_i = 1, 2, \ldots, T_i)$. However, without loss of generality and for ease of notation, we maintain the lighter epoch indexing t throughout the manuscript.

Within subject, at epoch t, we conceptualize synchronous patterns of cortical activity, by clustering electrodes according to the following mixture model. Denoting with $f\{\cdot \mid \cdot\}$ a generic density with respect the the Lebesgue measure on $\mathcal{B}\left(\mathbb{R}^d\right)$, we assume that each eigen-Laplacian vector $X_{ij}(t)$ is sampled from a K-components mixture distribution, indexed by parameters $\boldsymbol{\theta}_{ik}(t)$ and mixture probabilities $p_{ijk}(t) \in [0, 1]$, such that:

$$X_{ij}(t) \sim \sum_{k=1}^{K} p_{ijk}(t) f\{X_{ij}(t) \mid \boldsymbol{\theta}_{ik}(t)\}, \quad \sum_{k=1}^{K} p_{ijk}(t) = 1. \tag{1}$$

We find it convenient to re-express this sampling model with the equivalent hierarchical

representation, mixing over cluster labels $L_{ij}(t) \in \{1, 2, \ldots, K\}$, s.t.:

$$\begin{aligned} X_{ij}(t) \mid L_{ij}(t) = k &\sim f\{X_{ij}(t) \mid \boldsymbol{\theta}_{ik}(t)\}, \\ \Pr\{L_{ij}(t) = k\} &= p_{ijk}(t). \end{aligned} \tag{2}$$

In this setting, echoing the clustering "*ideal case*" discussed in the previous section, we exploit the connection between K-means and Gaussian mixtures and represent the sampling density in (1) as a K-component location/scale mixture of Gaussian distributions. Specifically, let $\mu_{ik}(t) \in \mathbb{R}^d$ be a d-dimensional mean vector, and $\sigma^2_{ik}(t) > 0$ be a variance parameter. We assume:

$$f\{X_{ij}(t) \mid \boldsymbol{\theta}_{ik}(t)\} = N\{\boldsymbol{\mu}_{ik}(t), \sigma^2_{ik}(t) I_d\}. \tag{3}$$

Given the sampling model in (2), our proposed approach for the integration of information at the subject and population levels follows a conceptually simple strategy, building directly on the setting of multilevel modeling (Gelman and Hill 2007). Crucially, we maintain that mixture means and variances are independent across subjects and epochs, but posit that cluster configurations, conceptualizing synchronicity of brain regions, are likely to adhere to patterns of similarity within and between subjects.

We make this idea precise by specifying a hierarchical prior for the mixture probabilities, $p_{ijk}(t)$. This is achieved by defining conditionally exchangeable mixture configurations, where epoch-level clusters $L_i(t)$ are obtained, *a priori*, as a stochastic perturbation of a time stable subject-level clustering, indexed by C_i. Similarly, subject level configurations, C_i, are obtained as a stochastic perturbation of a population-level cluster, indexed by S.

Let $C_{ij} \in \{1, 2, \ldots, K\}$ be the cluster label for electrode j at the level of subject i. Furthermore, let $\beta_i(t) \in [1/K, 1]$ be an adherence parameter, quantifying conformity

between cluster assignments at epoch t and the subject-level label $C_i = (C_{i1}, \ldots, C_{ip})'$. We assume,

$$\Pr\{L_{ij}(t) = k \mid c_{ij}\} \equiv \nu_c\{k, c_{ij}, \beta_i(t)\} = \begin{cases} \beta_i(t) & \text{if } c_{ij} = k \\ \frac{1-\beta_i(t)}{K-1} & \text{otherwise} \end{cases}, \tag{4}$$

where the probability $\nu_c\{\cdot, \cdot, \cdot\}$ is defined implicitly. This prior defines a probabilistic anchor, relating epoch level patterns of synchronicity at the subject level via simple and interpretable parameters $\beta_i(t)$. The underlying assumption is that epoch-level patterns of synchronicity are allowed to vary dynamically with t, but that variation in cluster configurations is anchored at the subject-level by a consensus pattern C_i.

A similar anchoring strategy is pursued at the population level. Specifically, let $S_j \in \{1, 2, \ldots, K\}$ be a population level cluster label for electrode j, and $\alpha_i \in [1/K, 1]$ be an adherence parameter, quantifying conformity between cluster assignments for subject i and population level labels $S = (S_1, \ldots, S_p)'$. We assume,

$$\Pr(C_{ij} = k \mid s_j) \equiv \nu_s(k, s_j, \alpha_i) = \begin{cases} \alpha_i & \text{if } s_j = k \\ \frac{1-\alpha_i}{K-1} & \text{otherwise} \end{cases}, \tag{5}$$

where probability $\nu_s(\cdot, \cdot, \cdot)$ is defined implicitly. The model is completed by specifying population level prior proportions:

$$\Pr(S_j = k) = \pi_k, \quad (k = 1, 2, \ldots, K).$$

To build intuition about the nature of these priors, we note that, if $\alpha_i = 1$, we expect cluster assignments for subject i to match exactly the population-level labels with probability 1. In contrast, for α_i approaching the value $1/K$, electrode clustering configurations C_i, for subject i, are drawn independently of the population level labels S. Similar considerations

apply to $\beta_i(t)$, as these paratmeters relate subject- and epoch-level cluster configurations.

This modeling strategy is loosely related to the idea of consensus clustering (Nguyen and Caruana 2007), as applied to the integration of multi-source data. Our specific formulation is a direct generalization to multilevel models of the approach taken by Lock and Dunson (2013) to the integration of heterogeneous genomic data.

In our multilevel setting, the conditional posterior distribution for epoch-level cluster labels $L_{ij}(t)$ is easily defined as:

$$\begin{aligned} \Pr\{L_{ij}(t) = k \mid X_{ij}(t), c_{ij}, s_j, \boldsymbol{\theta}_i(t)\} &\propto f\{X_{ij}(t) \mid L_{ij}(t) = k, \boldsymbol{\theta}_i(t)\} \Pr\{L_{ij}(t) = k \mid c_{ij}\} \\ &= f\{X_{ij}(t) \mid \boldsymbol{\theta}_{ik}(t)\} \nu_c\{k, c_{ij}, \beta_i(t)\}. \end{aligned} \tag{6}$$

This form highlights how inference on $L_{ij}(t)$ integrates information from both data $X_i(t)$ at epoch t, and subject-level clustering C_i (assumed stable across epochs), through a weighting scheme proportional to the size of the adherence parameter $\beta_i(t)$.

At the subject-level, conditional posterior probabilities of cluster membership weigh epoch level configurations $L_i(t) = (L_{i1}(t), L_{i2}(t), \ldots, L_{ip}(t))'$ with population level configurations S, through adherences α_i as follows:

$$\begin{aligned} \Pr\{C_{ij} = k \mid \ell_{ij}(1), \ldots, \ell_{ij}(T_i), s_j\} &\propto \Pr\{\ell_{ij}(1), \ldots, \ell_{ij}(T_i) \mid C_{ij} = k\} \Pr\{C_{ij} = k \mid s_j\} \\ &= \textstyle\prod_{t=1}^{T_i} \nu_c\{\ell_{ij}(t), k, \boldsymbol{\beta}_i\} \, \nu_s(k, c_j, \alpha_i). \end{aligned} \tag{7}$$

Finally, at the population level, overall consensus labels S are determined according to the following conditional posterior probability:

$$\Pr(S_j = k \mid c_{1j}, \ldots, c_{nj}, \Pi, \boldsymbol{\alpha}) \propto \pi_k \prod_{i=1}^{n} \nu_s(k, c_{ij}, \alpha_i). \tag{8}$$

2.3 Posterior Inference

We discuss posterior inference for the model in § 2.2 on the basis of MCMC samples from the target distribution. Even though multilevel modeling of cluster labels is a somewhat non-standard approach in a hierarchical setting, conditionally conjugate analysis is indeed possible, resulting in significant simplifications in computation and inference.

Specifically, we consider a standard Dirichlet prior for population-level proportions, so that $\Pi = (\pi_1, \pi_2, \ldots, \pi_K)' \sim \text{Dirichlet}(\eta)$. Epoch-level means and variances, are chosen to be conjugate to the graph Laplacian likelihood in (3). Letting $\boldsymbol{\theta}_{ik}(t_{im}) = (\boldsymbol{\mu}_{ik}(t_{im})', \sigma_k^2(t_{im}))'$, we assume that $\boldsymbol{\theta}_{ik}(t_{im}) \sim N\Gamma^{-1}(\boldsymbol{\mu}_0, \lambda_0, \xi_{01}, \xi_{02})$. Finally, subject-level adherence parameters α_i and epoch-level adherence parameters $\beta_i(t_{im})$ are assigned truncated Beta priors, with left truncation at $1/K$, so that:

$$\alpha_i \sim \text{TBeta}(a_i, b_i, 1/K), \text{ and } \quad \beta_i(t) \sim \text{TBeta}(c_i, d_i, 1/K).$$

A justification for these truncated Beta priors may be obtained by considering the form of the marginal allocation probabilities at subject and epoch level. Given Π, subject-level allocation probabilities are expressed as:

$$p_{ik} = \Pr(C_{ij} = k \mid \pi_k) = \pi_k \alpha_i + (1 - \pi_k)\frac{1 - \alpha_i}{K - 1}.$$

Similarly, at the epoch level, we have:

$$\Pr\{L_{ij}(t) = k \mid \Pi\} = \sum_{c_{ij}} \Pr\{L_{ij}(t) = k \mid c_{ij}\} \Pr(c_{ij} \mid \Pi) = \beta_i(t) p_{ik} + (1 - p_{ik})\frac{1 - \beta_i(t)}{K - 1}.$$

At both levels, an adherence value of $1/K$ corresponds to allocation probabilities, which

are independent of higher-level clustering realizations.

A Gibbs sampler targeting the posterior distribution is easily devised, by iterating through a transition sequence of full conditional posteriors. Specific details about the form of the conditional posterior densities are reported in a supplementary document.

At each level of the model, the posterior probability associated with set of clustering labels, for generality say $p(C \mid \mathbf{Y})$, and the corresponding MCMC samples, summarize our knowledge about potential partitions of cortical regions into synchronously activated areas. Based on the information in this posterior, we may be interested in selecting a representative partition, say C^*. Following Dahl (2006), we avoid using the naïve maximum *a posteriori* (MAP) estimate and instead consider a point estimator based on least squares. More precisely, consider an MCMC sample of M $p-$dimensional label configurations, $\{C^{(r)} : r = 1, 2, \ldots, M\}$. For each sample, we define a $p \times p$ adjacency matrix $\mathcal{A}\left(C^{(r)}\right) = \left[\mathcal{A}\left(C^{(r)}\right)_{ij}\right] = \left[I(C_i^{(r)} = C_j^{(r)})\right]$. Let $\bar{\mathcal{A}}$ be an estimate of the posterior mean $E[A \mid Y]$. The least square estimate C^* is selected from posterior realizations which minimize the following Frobenius norm

$$C^* = \min_{r=1,\ldots,M} || \, \mathcal{A}\left(C^{(r)}\right) - \bar{\mathcal{A}} \, ||_2.$$

Uncertainty about clustering estimates can be obtained from the posterior distribution, locally by quantifying pairwise relative frequencies of synchronization or globally via the distribution of $\mathcal{D} = ||A^{(r)} - \bar{A}||_2$. Examining this quantity facilitates direct comparison between subject and population level clustering results, allowing for low dimensional assessment of cluster quality, population and individual-level variability.

Computation and inference for MIC is performed under the R environment. A readily compiled package is available from the corresponding author's GitHub page.

2.4 Number of Clusters and Identifiability

Posterior inference as described in §2.3 presumes a known number of clusters K and a known number of eigen-Laplacian components d. For given d, selection of the number of mixture components, K, may be based on information criteria. In our simulation studies we find that the Bayesian Information Criterion (BIC) (Schwarz et al. 1978) tends to outperform more complicated indices. Our findings are in agreement with Steele and Raftery (2010), who observed that BIC outperforms many other criteria including ICL, DIC, and AIC, especially in the case of Gaussian mixture models.

The choice of d is less trivial, even though, some theoretical results point to the inclusion of the first K eigenvectors as being sufficient in the task of separating K groups, (Ng et al. 2001). Guided by this general principle, we perform a joint search on the dimensionality of the eigen-Laplacian d, and the number of clusters K simultaneously. More precisely, within a specific dimension d, the optimal value of $(K \mid d)$ is determined by the maximal BIC. Starting from low dimensions, usually $K = d = 2$, we allow for up-transitions on dimensionality, when $K^* \mid d > d$. Stopping rules, aiming at achieving stable solutions around the equality of $d^* = K^*$ are determined heuristically. Details are reported in Algorithm 1. Crucially, we avoid complete enumeration over all (d, K) combinations, and propose a search strategy which is linear in the maximum number of clusters. Our empirical studies in §3 show good performance and fast convergence to well behaved solutions.

For given d and K, simulation based procedures, including MCMC, are usually prone to *label switching* (Celeux et al. 2000). In the setting of the model proposed in §2.2 the same phenomenon may occur both within and between data levels. An important aspect of simulation-based inference in multilevel clustering is, therefore, the enforcement of correspondence between component labels of epochs, subjects and population level clustering.

Algorithm 1 (d, K) Selection

```
Set d = 2, K = 2;
current_BIC = BIC(d, K);
while d ≤max_d do
    while BIC(d, K + 1) ≥ current_BIC do
        current_BIC = BIC(d, K + 1);
        K = K + 1;
    if d ≥ K then
        break;
    else
        d = K, K = K − 1;
        current_BIC = BIC(d, K);
return (d, K)
```

Possible remedies include artificial identifiability constraints, relabeling procedures, and label invariant loss functions (Jasra et al. 2005). Within the multilevel setting, we proceed with online class relabeling or alignment. More precisely, we operate within population and subject-level indexes to find permutations of labels that maximize adherence with the population level clustering. Specifically, all newly sampled labels are permuted to insure maximal alignment with the population indexes. If $\mathcal{A}_0$ is an adjacency matrix as defined in §2.3, representing the current state of the population level labels S, and $\mathcal{A}_q$ is an adjacency matrix representing the current state of any other level clustering, optimal alignments are obtained by maximizing $tr\left(\mathcal{A}_0'\mathcal{A}_q\right)$ over $k!$ possible permutations.

3 Monte Carlo Studies

To investigate the operating characteristics of the proposed framework, we simulate EEG signals with the desired oscillation features from a mixture of AR(2) processes. We seek to evaluate: (1) the sensitivity of MIC results to differing sliding window size, γ, at a

fixed overlapping percentage, δ, (2) the accuracy of estimated quantities for varying group adherence, and (3) the performance of the model selection strategy proposed in Algorithm 1.

3.1 Simulation setup for spectrally specified EEGs

We make an effort to tailor the simulation of engineered time series in a way that mimics a sample of EEG readings typically seen in practice. To this end, we note that EEGs are often expected to feature oscillation patterns at different frequency bands: delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-12 Hz), beta (12-30 Hz) and gamma (30-50 Hz). Waveforms that are subdivided into bandwidths are thought to correspond to region-related activities on the cortex, both normally and pathologically. Our strategy, aims to simulate this spectral distinguishability by allowing each spectrum to exhibit concentrated (peak-shaped) energy in at most two frequency bands. Given a family of spectra, EEG time-series are simulated from a linear mixture of second order auto-regressive AR(2) processes. Details about the data generating mechanism are reported in a supplementary document. Furthermore, we represent potential non-stationarity by generating time-series as realizations from a piecewise stationary process, alternating randomly between two spectral configurations: a *main*-state (Fig 1(a)), and an *off*-state shown in Fig 1(b). The *main*-state has a time span $t_{\text{main}} \sim \exp(\lambda)$, with $\lambda = .05$s, followed by the *off*-state which has a time span $t_{\text{off}} \sim N(5, 1)$. Fig 1 (c) depicts this piecewise-stationarity for one electrode from the simulated samples. Cluster labels are generated as follows:

1. At the *population level*, we structure cluster labels S_j, $(j = 1, \ldots, p = 100)$ to partition 100 sensors into 4 balanced clusters.

2. Draw α from a Uniform$(0.5, 1)$ distribution. For $i = 1, \ldots, 9$, and $j = 1, \ldots, 100$; generate *subject level* labels $C_{ij} \in \{1, 2, 3, 4\}$ with probabilities $\Pr(C_{ij} = S_j) = \alpha$ and

$\Pr(C_{ij} \neq S_j) = 1 - \alpha$.

3. Given C_{ij}, generate piecewise stationary processes for 50 seconds, according to the *main*-state / *off*-state mechanism described previously.

Our Monte Carlo study is based on 100 datasets. The number of subjects, electrodes and segments were chosen to mimic the sampling structure in our case study. Note that in this setting, knowledge of the timing of *on*-state, *off*-state would result in perfect agreement of cluster labels within subject. Our simulation is therefore engineered to detect specific sensitivity to alternative metapreprocessing strategies.

3.2 Operating characteristics

In §2.1 we introduced a pre-processing step to smooth over the duration of the EEG recordings in order to obtain time-stable estimates of spectral densities. Because the time-series literature is substantially silent on the estimation of time-varying spectra in series with latent gaps, we set out to assess robustness of our procedure to alternative meta-processing strategies. In particular, we assess sensitivity of window size, $\gamma \in \{4, 6, 8, 10\}$, at a fixed $\delta = 0.5$ fraction of overlap between epochs.

In all 100 realizations, Algorithm 1 successfully selected the correct number of clusters ($K = 4$). Search paths differed by length across simulations. We observe two possibilities in our data: a searching path of length 7 that directly jumps to $d = 4$ after $K = 4$ is selected optimally at $d = 2$; and another path of length 10, which explores $d = 2, 3$ and 4 sequentially before d and K coincide at 4, Fig 2(a). This confirms our intuition that the complexity of the proposed search strategy is $O(p)$, as opposed to a complete enumeration strategy, which would require $max(d) \times p$ runs. We also find that results are relatively

robust in terms of the fit, estimated adherence parameters α, and clustering accuracy. Detailed results are provided in a supplementary document.

We investigate the performance of MIC under varying degrees of subject-specific variability, by examining estimates of adherence between subject and population-level clustering. Fig 2(b) depicts posterior medians $\widehat{\alpha}_i = E(\alpha_i \mid \mathbf{X})$ and their 95% credible intervals, based on the 2.5 and 97.5 percentiles, against the true α's. Posterior estimates are generally close to their true values, and over 99% of the credible intervals cover the true α's. For clarity, in both panels we report findings for $\gamma = 8$, as our findings indicate very low sensitivity to metapreprocessing variations.

Clustering accuracy, defined as the percentage of correctly classified electrodes, is assessed both at the subject and population level, Fig 2(c). Estimated subject-level clusters tend to be recovered accurately ($> 90\%$), regardless of α values. As expected, accuracy in the recovery of population level patterns relies on the magnitude of subject-level adherence to the population, with accuracy approaching 100% as $\alpha \to 1$.

We investigate the relationship between subject-level and population-level clustering variance estimates as a function of adherence and meta-processing strategy, Fig 2(d). Our summaries focus on a measure of global variance $\mathcal{D}$, as defined in §2.3. More precisely, denoting the clustering variance by $\mathcal{D}_S$ at the population level, and by $\mathcal{D}_{C_i}$ at the level of subject i, we consider the average difference in clustering variance, defined as:

$$\Delta_{\mathcal{D}} = E(\mathcal{D}_S \mid \mathbf{X}) - \frac{1}{n}\sum_i E(\mathcal{D}_{C_i} \mid \mathbf{X}).$$

As the adherence simulation truth approaches a level of complete agreement ($\alpha \to 1$), the average difference in clustering variance $\Delta_{\mathcal{D}}$ converges to zero, indicating that average subject-level and population-level cluster variances reach similar magnitudes over strongly

adherent clustering patterns. Finally, because variance estimates are likely to be vulnerable to alternative pre-processing strategies, we investigate sensitivity to changes in the smoothing bandwidth, γ. In our simulation studies, clustering variance estimates tend to be robust to reasonable alternative smoothing patterns, Fig 2(d). This feature is likely to be useful in many applications, where it is usually hard to develop meta-processing gold standards.

4 A Case Study on Resting State Brain Activity

Our study originates from an experiment aimed at understanding children's neurocognitive development. The study was carried out in the department of Psychiatry at UCLA and aims to cluster spectrally synchronized EEG signals recorded during *resting-state*. We provide technical background information about the study design and measurement structure in a web-based supplement. Here we investigate neuronal synchronicity in a group of typically developing (TD) children. We contrast group inference for the TD cohort against patterns of synchronicity in a cohort of children diagnosed with Autism Spectrum Disorder (ASD) in § 4.1. To our knowledge this is the first attempt at population level-inference for neuronal synchronicity in the setting of EEG studies.

4.1 MIC Analysis of TD and ASD Children

Autism Spectrum Disorder (ASD) describes a neurodevelopmental condition, characterized by social communication deficits, presence of repetitive behaviors, and/or restricted interest. Clinical presentation is highly variable, with heterogeneity in relation to medical conditions, behavioral challenges, and degree of intellectual impairments (Parr et al. 2011). Such behavioral and neurophysiological heterogeneity poses serious challenges to the study

of the neurophysiological substrate. In this respect, resting-state EEG is a particularly advantageous, and therefore popular, brain imaging choice (Wang et al. 2013).

Here we perform a comparative study between age-matched TD and ASD cohorts, under the framework of Multilevel Integrative Clustering (MIC). The study includes 9 participants (29-60 months of age) from the TD group, and 10 participants (27-99 months of age) from the ASD group. During the experiment, EEG was recorded at 250Hz using 129 channel geodesic nets with Ag/AgCl electrodes. Recordings took place while participants watched videos of bubbles and other non-social images on a computer monitor for 2 to 6 minutes.

Starting with the TD cohort, our analysis follows the scheme detailed in § 2.1 and considers epochs composed of $\gamma = 6$ contiguous 1024ms segments, allowing for a $\delta = 0.5$ overlap between epochs. This choice was based on both substantive and empirical considerations. In particular, we consider a smoothing strategy that guarantees good average adherence. A sensitivity analysis to differential smoothing choices was carried out with respect to both the epoch length and the percent of overlap. For fixed $K = d = 5$ we observe only small changes in the estimated adherence parameters $\hat{\alpha}_i = E(\alpha_i \mid \mathbf{X})$ (Figure 3). Furthermore, the search strategy outlined in Algorithm 1 selects $K = 5$ clusters in all but one of the smoothing configurations, where K is taken to be 6.

For the TD group, Figure 4 illustrates the full depth of multilevel inferences available from the MIC analysis. More precisely, cluster labels, representing patterns of spectrally synchronized neuronal activity, are projected onto a scalp map at the population, subject, and epoch levels. The first row represents the population-level cluster (1.a), together with two subject-level maps in (1b) [worst population adherence] and (1c) [best population adherence]. Variability within-subject is illustrated in the second row, where we compare subject-level inference for a specific subject with subject's epoch estimates. Specifically, (2b) reports the worst adherence epoch and (2c) reports the best. At the population level,

we identify 5 spectrally synchronized areas, corresponding to the following cortical regions: frontal, left parietal and left temporal, top-right parietal and right temporal, right-posterior and right parietal, left posterior and occipital.

An informal comparison between TD and ASD groups is carried out in Figure 5. At the population level, the least square estimates of cortical clusters are remarkably similar between the two cohorts, with the exception of an asymmetrical partition on the posterior and occipital regions, where a small cluster was isolated from the rest and leaning towards the left hemisphere for ASD, but towards the right hemisphere for TD.

Further, we examine local and global sources of cluster variability in both groups. At the electrode level, we report the entropy associated with posterior cluster label probabilities in Figure 5: (1.b) for ASD and (2.b) for TD.

Perfect partitions, e.g. an electrode assigned to cluster k with probability one, yields 0 entropy, whereas uniform assignment probabilities yield entropy equal to 1. We observe that the mid-, right-frontal and mid posterior regions are the most stable regions for both groups. Compared to the ASD group, the TD cohort exhibits more stable regions, for example, in the left-temporal (speech and language related), left-central, as well as some regions in the posterior and occipital areas of the cortex. The high entropy observed on the left-hemisphere among ASD children coincides with the abnormal left-hemispheric asymmetry findings in the literature on individuals with ASD (Stroganova et al. 2007, Burnette et al. 2011).

We gain more insight into the nature of variability of synchronized neuronal patterns by examining global sources of cluster variance at the subject-specific and population levels. More precisely we consider the full posterior distribution of the Frobenius norm of differential affinity, between individual draws and the least square adjacency (§ 2.3). Figures 5, (1.c) for ASD and (2.d) for TD, highlight how in the TD group, subject-level clusters com-

bine to reduce overall variance at the population level. In ASD children we observe higher global heterogeneity, with subject-level variability spreading over a wider range, when compared to the TD cohort. This observation echoes some of our previous findings in EEG studies of implicit-learning in ASD and TD children (Hasenstab et al. 2015; Hasenstab, Sugar, Telesca, Jeste and Senturk 2016; Hasenstab, Telesca, Sugar, DiStefano and Şentürk 2016).

5 Discussion

This paper proposes what to our knowledge is the first comprehensive statistical framework for population level inference of spectrally synchronized brain activity from a heterogeneous sample of EEG readings. A hierarchical model allows for the estimation of population level synchronicity patterns, with full consideration of intra- and inter-subjects variability. Crucially, information is borrowed at the latent level of cluster membership indicators. Dependent mixtures are based on a hierarchical Dirichlet prior, indexed by interpretable and informative parameters, which measure cluster adherence at all levels of the hierarchy.

Our approach melds non-parametric dimension reduction and fully model-based techniques through a graph-partitioning representation of clustering. This two-stage approach is likely to be useful in several experimental settings involving EEG measurements, where different scientific goals and different data meta-processing concerns may require substantial subject-matter input in the definition of similarity between cortical regions.

In our study we operate within the context of spectral synchronicity. It is however important to point out that alternative measures of neuronal affinity, for example partial correlation, coherence, and mutual information, are also amenable to MIC analysis. In this sense, the proposed framework is quite general and can be adapted to handle alterna-

tive neuroimaging data platforms, such as functional Magnetic Resonance Imaging (fMRI). This consideration also applies, with possible minor adjustments, to the integration of multiple imaging modalities. This flexibility traces back to the hierarchical prior, which relates cluster labels rather than cluster-specific parameters (location and scale for example), so that complex data alignment issues are resolved within a higher level of modeling abstraction. Clearly, technical preprocessing pipelines may differ substantially between and within modalities. Therefore, important analytic details should be thoughtfully engineered in practice.

Our simulation results in § 3.2, show that inference is robust to reasonable variants in the meta-processing strategies. In our experiments, simple information criteria like BIC tend do do well in the selection of the number of clusters K, when combined with a search over the number of eigen-Laplacians d. Our model, of course, offers a very simple representation of cluster variability within- and between-subjects. Therefore, modeling refinements are likely needed in applications where one can expect a strong dynamic evolution of synchronicity patterns; such as the setting of stimulus-based EEG studies. Additionally, potentially useful extensions include a formal treatment of group comparison and regression.

References

Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.-M., and Robbins, K. A. (2015), "The PREP pipeline: standardized preprocessing for large-scale EEG analysis," *Frontiers in Neuroinformatics*, 9, 16.
URL: *http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4471356/*

Brillinger, D. R. (1981), *Time series: data analysis and theory* Holden-Day Inc. Oakland, California.

Broyd, S., Demanuele, C., Debener, S., Helps, K., James, C., and Sonuga-Barke, E. (2008), "Default-mode brain dysfunction in mental disorders: a systematic review," *Nature Biobehavioral Review*, 33(3), 279–296.

Burnette, C. P., Henderson, H. A., Inge, A. P., Zahka, N. E., Schwartz, C. B., and Mundy, P. C. (2011), "Anterior EEG asymmetry and the modifier model of autism," *Journal of autism and developmental disorders*, 41(8), 1113–1124.

Celeux, G., Hurn, M., and Robert, C. P. (2000), "Computational and inferential difficulties with mixture posterior distributions," *Journal of the American Statistical Association*, 95(451), 957–970.

Craddock, R. C., Jbadi, S., Yan, C. G., Vogelstein, J. T., Castellanos, F. X., Di Martino, A., Kelly, C., Heberlein, K., Colcombe, S., and Milham, M. P. (2013), "Imaging human connectomes at the macroscale," *Nature Methods*, 10(6), 524–539.

Dahl, D. (2006), "Model-based clustering for expression data via a Dirichlet process mixture model," in *Bayesian inference for gene expression and proteomics*, eds. M. Vannucci, K.-A. Do, and P. Müller Cambridge: Cambridge University Press, pp. 201–215.

Di Martino, A., A, F. D., Satterthwaite, T. D., Castellanos, F. X., Thomason, M. E., Craddock, R. C., Luna, B., Leventhal, B. L., N, Z. X., and Milham, M. P. (2014), "Unraveling the miswired connectome: a developmental perspective," *Neuron*, 17(83), 1335 – 1353.

Euan, C., Ombao, H., and Ortega, J. (2015), "Spectral Synchronicity in Brain Signals," *arXiv preprint arXiv:1507.05018*, .

Fell, J., and Axamcher, N. (2011), "The role of phase synchronization in memory processes," *Nature Reviews Neuroscience*, 12(2), 105–118.

Florian, G., and Pfurtscheller, G. (1995), "Dynamic spectral analysis of event-related EEG data," *Electroencephalography and clinical neurophysiology*, 95(5), 393–396.

Gelman, A., and Hill, J. (2007), *Data analysis using regression and multilevel/hierarchical models* Cambridge University Press , New York.

Hasenstab, K., S. A., Telesca, D., Sugar, C. A., DiStefano, C., and Şentürk, D. (2016), A multi-dimensional functional principal components analysis of EEG data„ Technical report, UCLA-Biostatistics.

Hasenstab, K., Sugar, C., Jeste, S., Telesca, D., McEvoy, K., and Senturk, D. (2015), "Identifying longitudinal trends within EEG experiments," *Biometrics*, 71(1090 - 1100).

Hasenstab, K., Sugar, C., Telesca, D., Jeste, S., and Senturk, D. (2016), "Robust functional clustering of ERP data with application to a study of implicit learning in autism," *Biostatistics*, 3, 484–498.

Jasra, A., Holmes, C., and Stephens, D. (2005), "Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling," *Statistical Science*, pp. 50–67.

Laufs, H., Krakow, K., Sterzer, P., Eger, E., Beyerle, A., Salek-Haddadi, A., and A, K. (2003), "Electroencephalographic signatures of attentional and cognitive default modes in spontaneous brain activity fluctuations at rest," *Proceedings of the national academy of sciences of the United States of America*, 100(19), 11053–11058.

Lock, E. F., and Dunson, D. B. (2013), "Bayesian consensus clustering," *Bioinformatics*, 29(20), 2610–2616.

Narayan, M., and Allen, G. I. (2015), Population Inference for Node Level Differences in Functional Connectivity„ in *IEEE International Workshop on Pattern Recognition in Neuroimaging*.

Ng, A. Y., Jordan, M. I., and Weiss, Y. (2001), On Spectral Clustering: Analysis and an Algorithm„ in *Advances in Neural Information Processing Systems*, MIT Press, pp. 849–856.

Nguyen, N., and Caruana, R. (2007), Consensus clusterings„ in *In Proceedings of the 7th IEEE International Conference on Data Mining*, ed. I. C. Society, pp. 607–612.

Ombao, H. C., Raz, J. A., Strawderman, R. L., and R, V. S. (2001), "A simple generalized cross-validation method of span selection for periodogram smoothing," *Biometrika*, 88(4), 1186–1192.

Parr, J. R., Le Couteur, A., Baird, G., Rutter, M., Pickles, A., Fombonne, E., Bailey, A. J., of Autism Consortium, I. M. G. S. et al. (2011), "Early developmental regression in autism spectrum disorder: Evidence from an international multiplex sample," *Journal of autism and developmental disorders*, 41(3), 332–340.

Rosen, O., Wood, S., and Stoffer, D. S. (2012), "AdaptSPEC: Adaptive spectral estimation for nonstationary time series," *Journal of the American Statistical Association*, 107(500), 1575–1589.

Schwarz, G. et al. (1978), "Estimating the dimension of a model," *The annals of statistics*, 6(2), 461–464.

Shou, H adn Eloyan, A., Nebel, M. B., Mejia, A., Pekar, J. J., Mostofsky, S., Caffo, B.,

Lindquist, M. A., and Crainiceanu, C. (2014), "Shrinkage prediction of seed-voxel brain connectivity using resting state fMRI," *Neuroimage*, 15(102), 938–944.

Steele, R. J., and Raftery, A. E. (2010), "Performance of Bayesian model selection criteria for Gaussian mixture models," *Frontiers of statistical decision making and bayesian analysis*, pp. 113–130.

Stroganova, T. A., Nygren, G., Tsetlin, M. M., Posikera, I. N., Gillberg, C., Elam, M., and Orekhova, E. V. (2007), "Abnormal EEG lateralization in boys with autism," *Clinical Neurophysiology*, 118(8), 1842–1854.

Teplan, M. (2002), "Fundamentals of EEG Measurements," *Measurement Science Review*, 2(2).

Wang, J., Barstein, J., Ethridge, L. E., Mosconi, M. W., Takarae, Y., and Sweeney, J. A. (2013), "Resting state EEG abnormalities in autism spectrum disorders," *Journal of neurodevelopmental disorders*, 5(1), 1.

Welch, P. D. (1967), "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Transactions on audio and electroacoustics*, 15(2), 70–73.

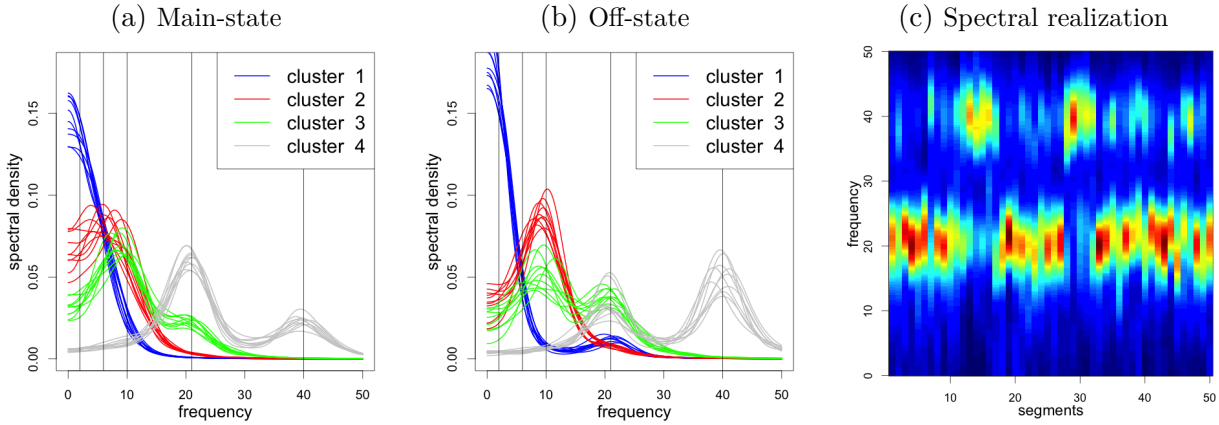

Figure 1: **Simulated spectral configurations**: (a) *main-state* spectral densities. (b) *off-state* spectral densities. (c) Segment-by-segment normalized power spectral densities for a piecewise stationary process simulated from cluster 4.

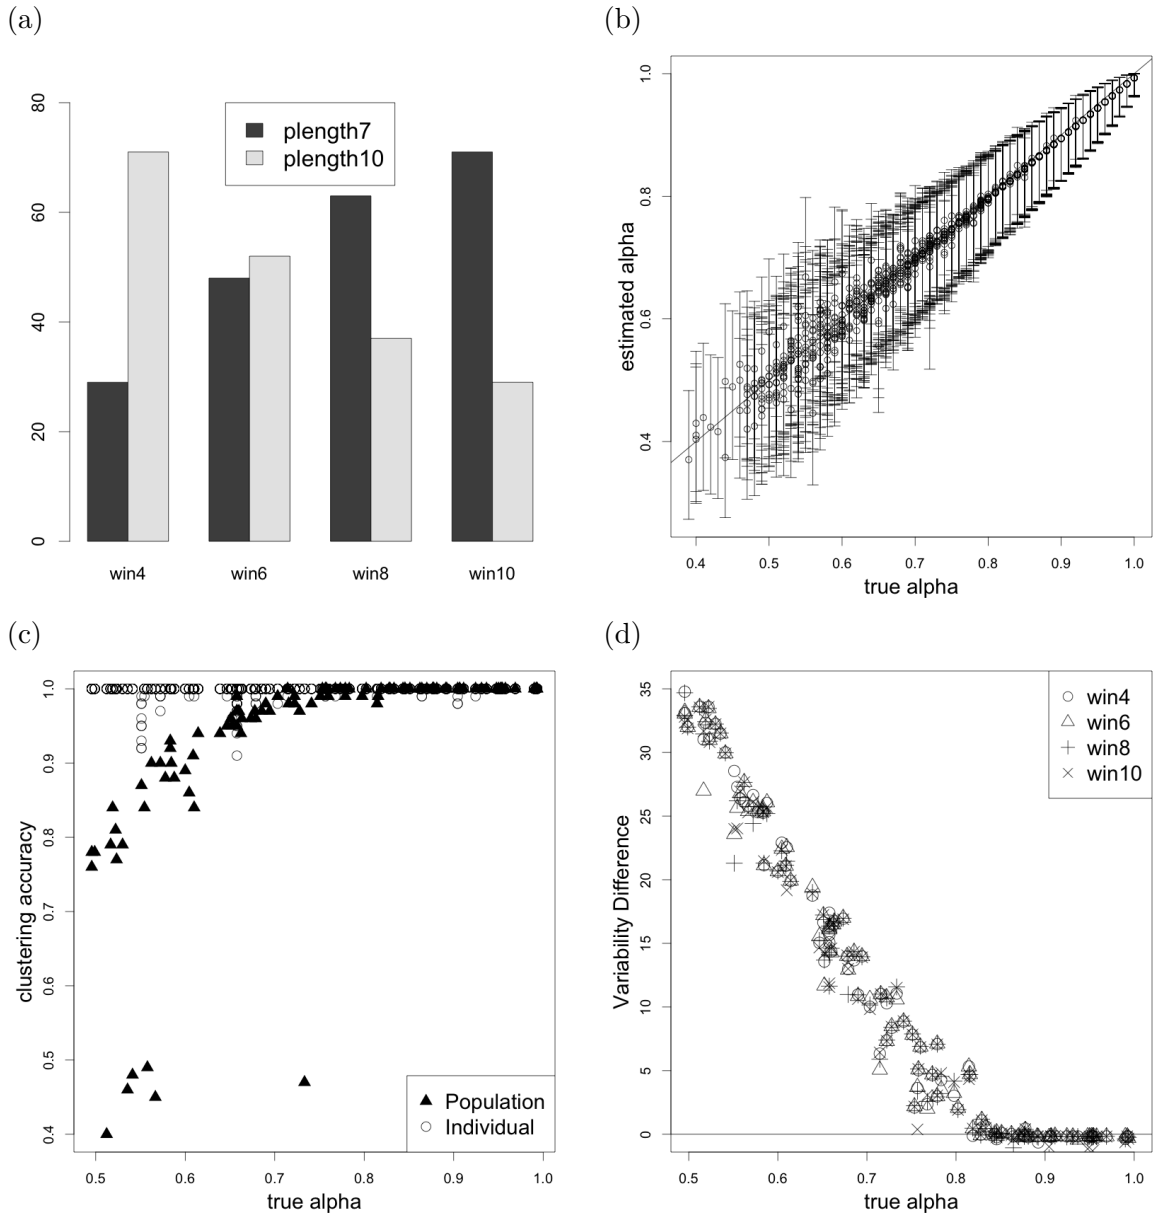

Figure 2: **Simulation results:** (a) Path-length for the search in Algorithm 1 for varying smoothing configurations in γ. (b) Estimated adherence parameters $\widehat{\alpha}$'s and 95% credible intervals against the data generating truth. (c) Clustering accuracy against generating α's at the subject- and population-level. (d) Average difference in clustering variance against true α's.

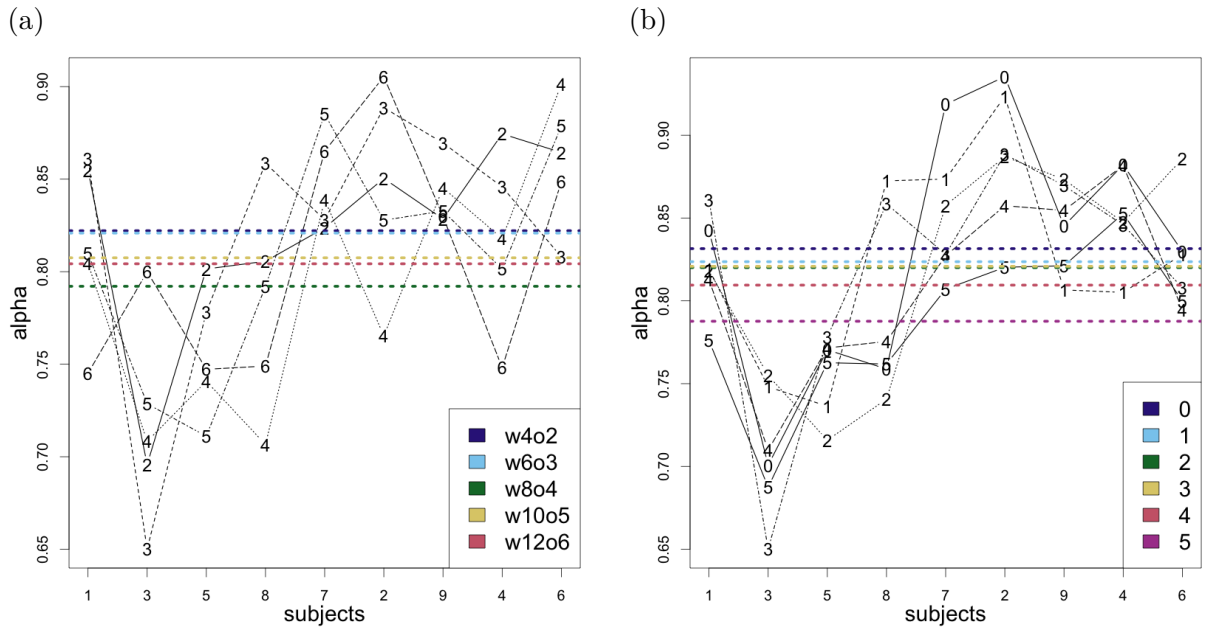

Figure 3: **Sensitivity to smoothing**: (a) Estimated adherence parameters $\hat{\alpha}_i$ and their average for varying smoothing configurations $\gamma \in \{4, 6, 8, 10, 12\}$, $\delta = \gamma/2$. (b) Estimated adherence parameters $\hat{\alpha}_i$ and their average for varying smoothing configurations $\delta \in \{0, 1, \ldots, 5\}$, $\gamma = 6$.

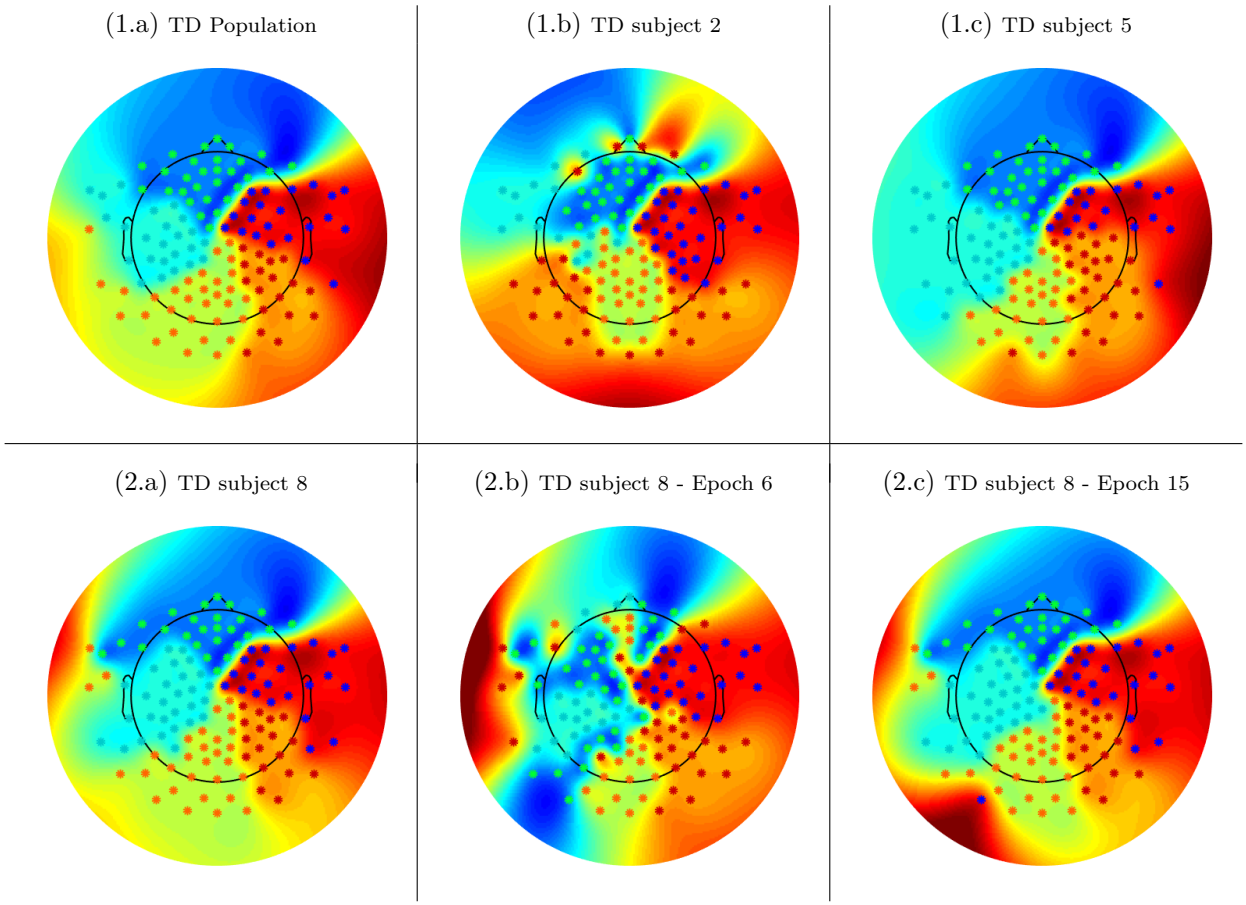

Figure 4: **TD cohort synchronicity**: (1.a) Population-level posterior least square cluster. (1.b) Individual map for subject 2 ($\widehat{\alpha}_i = .73$). (1.c) Individiual map for subject 2 ($\widehat{\alpha}_i = .95$). (2.a). Individual map for subject 8. (2.b) Clustering for subject 8, epoch 6 ($\widehat{\beta}_i(t) = .48$). (2.c) Clustering for subject 8, epoch 15 ($\widehat{\beta}_i(t) = .99$).

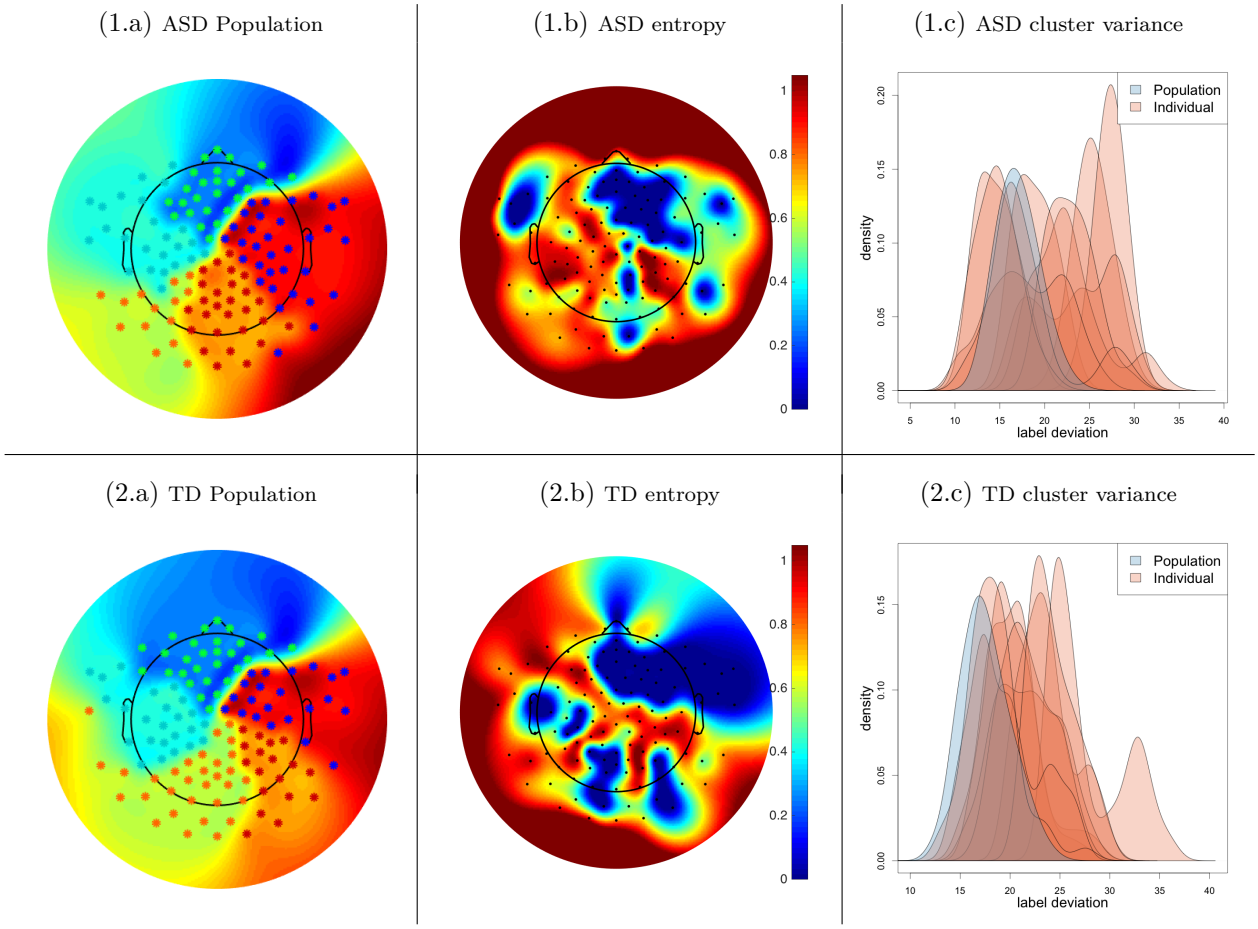

Figure 5: **Group contrasts, ASD (1) vs TD (2)**: (1.a) TD-cohort posterior least square synchronicity. (1.b) TD-cohort normalized posterior entropy. (1.c) TD-cohort global cluster variance at the subject- and population-level. (2.a) ASD-cohort posterior least square synchronicity. (2.b) ASD-cohort normalized posterior entropy. (2.c) ASD-cohort global cluster variance at the subject- and population-level.