NeuroImage 126 (2016) 81–95

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Spectrally resolved fast transient brain states in electrophysiological data

Diego Vidaurre [a,*], Andrew J. Quinn [a], Adam P. Baker [a], David Dupret [b], Alvaro Tejero-Cantero [b,c], Mark W. Woolrich [a]

[a] Oxford Centre for Human Brain Activity, Department of Psychiatry, University of Oxford, UK
[b] MRC Brain Network Dynamics Unit, Department of Pharmacology, University of Oxford, UK
[c] Computational Neuroscience, Department Biologie II Ludwig Maximilian, University Munich, Germany

ARTICLE INFO

Article history:
Received 27 July 2015
Accepted 19 November 2015
Available online 26 November 2015

Keywords:
multivariate autoregressive model
MEG
Transient connectivity
Bayesian modelling
Spectral estimation
Multitaper
Coherence
Partial directed coherence
Sign ambiguity

ABSTRACT

The brain is capable of producing coordinated fast changing neural dynamics across multiple brain regions in order to adapt to rapidly changing environments. However, it is non-trivial to identify multiregion dynamics at fast sub-second time-scales in electrophysiological data. We propose a method that, with no knowledge of any task timings, can simultaneously identify and describe fast transient multiregion dynamics in terms of their temporal, spectral and spatial properties. The approach models brain activity using a discrete set of sequential states, with each state distinguished by its own multiregion spectral properties. This can identify potentially very short-lived visits to a brain state, at the same time as inferring the state's properties, by pooling over many repeated visits to that state. We show how this can be used to compute state-specific measures such as power spectra and coherence. We demonstrate that this can be used to identify short-lived transient brain states with distinct power and functional connectivity (e.g., coherence) properties in an MEG data set collected during a volitional motor task.

Introduction

The brain is able to coordinate neural oscillations across multiple brain areas in both rest and task (Buzsaki and Draguhn, 2004), (Mantini et al., 2007), (Fries, 2005), (Schnitzler and Gross, 2005). However, the manner in which these neural interactions arise in the brain is not fully understood. Typically, in electrophysiological data these oscillatory interactions are characterised using their multiregion spectral properties, e.g., the power content or the extent of phase locking (e.g., coherence) over different cortical regions (Lachaux et al., 1999). However, since the brain must be able to rapidly reorganise neural oscillations in response to the environment, there is a need to be able to identify how these multiregion spectral properties vary over time at potentially very fast (sub-second) time-scales.

Many existing methods for investigating time-varying patterns of spectral properties or functional connectivity use sliding time windows (Wendling et al., 2009), (Allen et al., 2014). Sliding window approaches pre-specify the temporal resolution of the changing patterns, and make inefficient use of the data when the same patterns occur recurrently at other points of time. The exception to this is when data can be pooled over epochs of a repeated task; however, this necessitates an assumption of stationarity over trials. These approaches also require a choice of the width of the time-window. Short windows can lead to noisy estimations, whereas long ones can miss the quickest changes.

In this paper, we provide a unified framework for characterising oscillatory dynamics in terms of their time-varying spatial and multiregion spectral properties without the knowledge of any task timings. The primary contribution of the method is that it operates simultaneously on the frequency, time and space dimensions, thus allowing for a unique description of transient spectral properties including power spectra and connectivity measures such as coherence. Importantly, it can identify when multiregion spectral patterns repeat at different points in time, and thereby pool over them to provide a better estimation of those patterns.

Although it is broadly applicable to any electrophysiological data modality, we focus here on magnetoencephalography (MEG), of particular interest for research on human connectivity for its fine-grain temporal resolution, wide-brain coverage and non-invasive nature. To this end, we also devise a way to deal with the sign ambiguity inherent to source reconstruction in MEG, which can jeopardise multisession/subject analyses if left unaddressed.

The method combines two well-known models: the multivariate autoregressive (MAR) (Penny and Roberts, 2002) model and the Hidden Markov model (HMM) (Juang and Rabiner, 1985). The MAR model characterises the behaviour of time series by linear historical interactions

* Corresponding author.
E-mail addresses: diego.vidaurre@ohba.ox.ac.uk (D. Vidaurre), andrew.quinn@psych.ox.ac.uk (A.J. Quinn), adam.baker@ohba.ox.ac.uk (A.P. Baker), david.dupret@pharm.ox.ac.uk (D. Dupret), tejero@bio.lmu.de (A. Tejero-Cantero), mark.woolrich@ohba.ox.ac.uk (M.W. Woolrich).

http://dx.doi.org/10.1016/j.neuroimage.2015.11.047

between the observed time series from different brain regions. MARs are able to characterise the frequency structure of the data, and by making the model multivariate, are able to capture interactions (e.g., coherence) between multiple brain regions. The HMM is a mathematical formalism that describes a time series as a sequence of states, where each state has its own model of the observed data (i.e., the observation model). Here, the observation model we use corresponds to a MAR model, and, hence, each state is related to a different set of multiregion autoregression coefficients describing the neural oscillations. In what follows, we will refer to the HMM with MAR observation model as the HMM–MAR.

Although the spectral contents of the states can be obtained directly from the (parametric) MAR model, we propose a non-parametric method based on the multitaper (Thomson, 1982) to obtain the states' spectral information given the state time courses. The motivations of the non-parametric approach are threefold. Firstly, the multitaper is known to provide a reliable estimation, often superior to the parametric approaches. Secondly, the MAR order, which is not needed for the non-parametric estimation, strongly affects the estimation of the spectral information. Finally, MAR orders that produce sensible state discrimination for the HMM–MAR do not necessarily match the MAR orders that are optimum for spectral estimation. We will show below that, even when the state visits are short (around 100 ms or less), the proposed statewise multitaper can provide reliable estimations of the entire range of frequencies of interest, including the low frequencies.

We first show how the model works on synthetic data, for which the ground-truth spectra are known. We then use the proposed model to characterise the neural dynamics in the primary motor cortex (M1) during a self-paced button press MEG experiment. We demonstrate that the proposed approach is able to identify HMM states that are task dependent despite training the HMM with no knowledge of the task timings, and that it can produce sensible state-specific estimates of the power spectral density (PSD), coherence and partial directed coherence (PDC) (Sameshima and Baccala, 1999) that are significantly different over states.

The method

We now describe the HMM–MAR, its Bayesian hierarchy and some aspects of model selection and inference. We also provide details about the non-parametric spectral estimation, and about two issues that are central to source space MEG data analysis: sign ambiguity and signal leakage. Fig. 1 illustrates the proposed workflow schematically; each step is described below.

Definition of the states and their Markov dynamics

In this section, we describe the observation model and the state transitions. As mentioned above, the observation model corresponds to a MAR model, and the state transitions follow the (first-order) Markovian assumption.

We first introduce some notation. Let $\boldsymbol{y}_t \in \mathbb{R}^N$ be the multichannel source signal and $x_t \in \{1, \ldots, K\}$ the hidden state variable, with $t = 1, \ldots, T$. Let $\mathcal{A}$ be the set of lags considered by the MAR model. We now present the MAR model leaving $\mathcal{A}$ unspecified, and will get into specifics about the choice of $\mathcal{A}$ in due course. Assuming Gaussian noise and centred data, our observation model is

$$\boldsymbol{y'}_t | x_t = k \sim \mathcal{N}\left(\sum_{l \in \mathcal{A}} \mathbf{y'}_{t-l} \mathbf{W}_l^{(k)}, \boldsymbol{\Sigma}^{(k)}\right), \quad (1)$$

where $\boldsymbol{W}_l^{(k)}$ are $N \times N$ dimensional matrices representing the k-th state autoregression coefficient matrices for lag l and the variance is given by some random noise distribution. We denote $\boldsymbol{W}^{(k)} = [\boldsymbol{W}_1^{(k)}; \ldots; \boldsymbol{W}_P^{(k)}]$. We shall also refer to the expectation of $\mathrm{P}(x_t = k | \boldsymbol{Y})$ as γ_{tk}, and $\boldsymbol{\gamma}_t = (\gamma_{t1}, \ldots, \gamma_{tK})$.

The noise covariance matrix $\boldsymbol{\Sigma}^{(k)}$ can be chosen to be diagonal or a full matrix. In the former case, we assume the zero-lag correlations to be zero. In the latter case, the noise is correlated across channels, which implies that the estimation of the autoregression coefficients has to be done for all channels at the same time (see Appendix B). Another decision to be made is whether we set the noise distribution to be equal for all states, so that $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}^{(k)}$, for all k.

For the hidden state variables, we use Markov dynamics, meaning that the probability $P(x_t = k)$ is conditionally independent of the history of the state variable given x_{t-1}. Hence, we have

$$\mathrm{P}(x_t = k_1 | x_{t-1} = k_2) = \Theta_{k_1 k_2}, \quad \mathrm{P}(x_1 = k) = \eta_k, \quad (2)$$

where $\Theta_{k_1 k_2}$ and η_k are model paramerers that need to be inferred. The model is graphically represented in Fig. 2.

Model complexity and model selection

In this section, we discuss the parametrisation of the MAR model and how to control its complexity. This is crucial, because, if the MAR models are too complex, the inference process (as a consequence of the Bayesian principle of parsimony) will tend to drop most of the states of the model by letting a few (or even one) dominant states to control the entire time series. Albeit good in terms of the tradeoff between predictability and parsimony, this hinders the discovery of quasi-stationary connectivity networks.

Firstly, driven by objective Bayesian principles, we use appropriate automatic relevance determination (ARD) priors on the autoregression coefficients. These ARD priors are Gaussian, and are imposed at two levels: for each lag (regularising on the time–frequency dimension) and for each pair of sources (regularising on the spatial dimension).

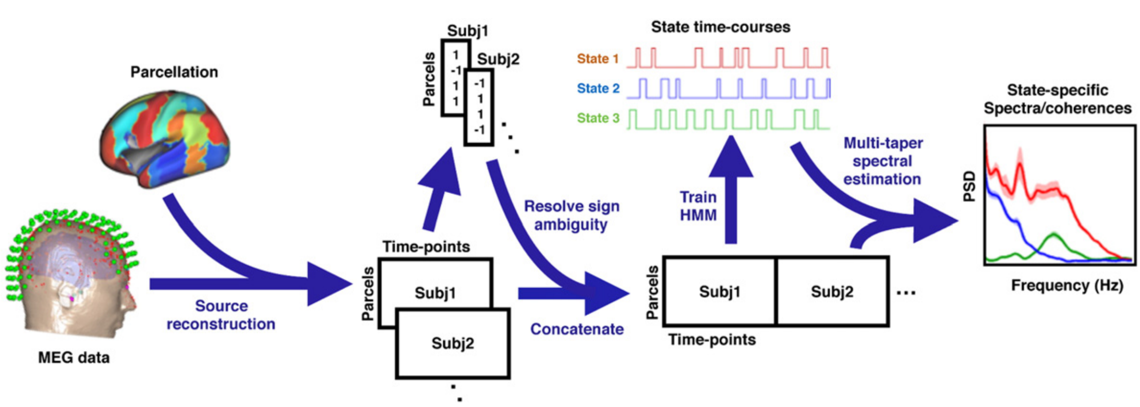

Fig. 1. Workflow of the proposed method.