NeuroImage 144 (2017) 142–152

Contents lists available at ScienceDirect

NeuroImage

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

Bayesian EEG source localization using a structured sparsity prior

Facundo Costa [*,a], Hadj Batatia [a], Thomas Oberlin [a], Carlos D'Giano [b], Jean-Yves Tourneret [a]

[a] University of Toulouse, INP/ENSEEIHT - IRIT, 2 rue Charles Camichel, BP 7122, 31071 Toulouse Cedex 7, France
[b] FLENI Hospital, Buenos Aires, Argentina

ARTICLE INFO

Article history:
Received 19 February 2016
Accepted 30 August 2016
Available online 15 September 2016

Keywords:
EEG
MCMC
Inverse problem
Source localization
Structured-sparsity
Hierarchical Bayesian model
ℓ_{20} norm regularization
Medical imaging

ABSTRACT

This paper deals with EEG source localization. The aim is to perform spatially coherent focal localization and recover temporal EEG waveforms, which can be useful in certain clinical applications. A new hierarchical Bayesian model is proposed with a multivariate Bernoulli Laplacian structured sparsity prior for brain activity. This distribution approximates a mixed ℓ_{20} pseudo norm regularization in a Bayesian framework. A partially collapsed Gibbs sampler is proposed to draw samples asymptotically distributed according to the posterior of the proposed Bayesian model. The generated samples are used to estimate the brain activity and the model hyperparameters jointly in an unsupervised framework. Two different kinds of Metropolis–Hastings moves are introduced to accelerate the convergence of the Gibbs sampler. The first move is based on multiple dipole shifts within each MCMC chain, whereas the second exploits proposals associated with different MCMC chains. Experiments with focal synthetic data shows that the proposed algorithm is more robust and has a higher recovery rate than the weighted ℓ_{21} mixed norm regularization. Using real data, the proposed algorithm finds sources that are spatially coherent with state of the art methods, namely a multiple sparse prior approach and the Champagne algorithm. In addition, the method estimates waveforms showing peaks at meaningful timestamps. This information can be valuable for activity spread characterization.

1. Introduction

EEG source localization problem has attracted considerable attention in the literature resulting in a wide range of methods developed in the last years. These can be classified into two groups: (i) the dipole-fitting models that represent the brain activity as a small number of dipoles with unknown positions; and (ii) the distributed-source models that represent the brain activity as a large number of dipoles in fixed positions. Dipole-fitting models (Sommariva and Sorrentino, 2014; da Silva and Van Rotterdam, 1998) try to estimate the amplitudes, orientations and positions of a few dipoles that explain the measured data. Unfortunately, the corresponding estimators are very sensitive to the initial guess of the number of dipoles and their initial locations (Grech et al., 2008). On the other hand, the distributed-source methods model the brain activity using a large number of dipoles with fixed positions and try to estimate their amplitudes (Grech et al., 2008) by solving an ill-posed inverse problem. One of the most simple ways to solve this inverse problem is to use an ℓ_2 norm regularization as the minimum norm estimator (Pascual-Marqui, 1999) or its variants Loreta (Pascual-Marqui et al., 1994) and sLoreta (Pascual-Marqui et al., 2002). However, these methods usually overestimate the active area size (Grech et al., 2008).

Sparsity constraints can remedy the overestimation issue when dealing with applications with discretely localized activity such as certain kinds of epilepsy (Berg et al., 2010). In distributed activity applications, promoting sparsity should provide spatially coherent localization even though it is unable to estimate the activity extension. To apply sparsity, ideally an ℓ_0 pseudo norm regularization (Candes, 2008) should be used. Unfortunately, this procedure is intractable in an optimization framework. As a consequence, the ℓ_0 pseudo norm is usually approximated by the ℓ_1 norm via convex relaxation (Uutela et al., 1999), even if the two regularizations do not always provide the same solution (Candes, 2008). In a previously reported work, we proposed to combine them in a Bayesian framework (Costa et al., 2015), using the ℓ_0 pseudo norm to locate the non-zero positions and the ℓ_1 norm to estimate their amplitudes. However the methods studied in Candes (2008), Uutela et al. (1999), and Costa et al. (2015) consider each time sample independently leading in some cases to unrealistic solutions (Gramfort et al., 2012).

To improve source localization, it is possible to make use of the temporal structure of the data. This can be done by considering sparse Bayesian learning using multiple measurement vectors (Zhang and Rao, 2011) or by using the STOUT (Castaño-Candamil et al., 2015) and dMAP-EM (Lamus et al., 2012) methods that apply

* Corresponding author at: INP - ENSEEIHT Toulouse 2, rue Charles Camichel, B.P. 7122, 31071 Toulouse Cedex 7, France. Tel.: +33 05 34 32 22 52.
E-mail address: facundo.costa@enseeiht.fr (F. Costa).

http://dx.doi.org/10.1016/j.neuroimage.2016.08.064

physiological considerations to the source representation. It is also possible to model the time evolution of the dipole activity and estimate it using Kalman filtering (Galka et al., 2004; Long et al., 2011), particle filters (Somersalo et al., 2003; Sorrentino et al., 2013; Chen and Godsill, 2013) or by encouraging spatio-temporal structures by promoting structured sparsity (Huang and Zhang, 2010).

Structured sparsity has been shown to improve results in several applications including audio restoration (Kowalski et al., 2013), image analysis (Yu et al., 2012) and machine learning (Huang et al., 2011). Structured sparsity has also been applied to M/EEG source localization by Gramfort et al. by using the ℓ_{21} mixed norm (Gramfort et al., 2012). This approach promotes sparsity among different dipoles (via the ℓ_1 portion of the norm) and groups all the time samples of the same dipole together, forcing them to be either jointly active or inactive (with the ℓ_2 norm portion). This work was reconsidered by the same authors yielding the iterative reweighted mixed norm estimator (Strohmeier et al., 2014) and the time–frequency mixed-norm estimator (Gramfort et al., 2013). However, all these methods require the manual tuning of the regularization parameters.

Several Bayesian methods have also been used to solve the inverse problem (Friston et al., 2008; Stahlhut et al., 2013; Wipf et al., 2010; Lucka et al., 2012). Friston et al. (2008) developed the multiple sparse priors (MSP) approach, in which they segment the brain into different pre-defined regions and promote all the dipoles in each region to be active or inactive jointly. In contrast, Wipf et al. developed the Champagne algorithm to promote activity to be concentrated on a sparse set of dipoles (Wipf et al., 2010). Lucka et al. (2012) studied a hierarchical Bayesian model (HBM) offering significant improvements over established methods such as MNE and sLoreta.

Similar to Wipf et al., this paper develops a new method encouraging sparse activity considering each dipole separately (Friston et al., 2008). The proposed method uses a multivariate Bernoulli Laplace prior (approximating the weighted ℓ_{20} mixed norm) for the dipole amplitudes without assuming any additional prior information such as the amount or position of the active dipoles. Since the parameters of the proposed model cannot be computed with closed-form expressions, we investigate a Markov chain Monte Carlo sampling technique to draw samples that are asymptotically distributed according to the posterior of the proposed model. Then the brain activity, the model parameters and hyperparameters are jointly estimated in an unsupervised framework. In order to avoid the sampler to becoming stuck around local maxima, specific Metropolis–Hastings moves are introduced. These moves significantly accelerate the convergence speed of the proposed sampler. From the medical point of view, the proposed approach aims at providing the localization of the main sources of the brain activity to help making decisions when selecting candidate patients for recessive surgery, in the case of discretely localized epilepsy (Berg et al., 2010). In addition, considering several time samples simultaneously allows us to estimate the temporal waveforms of the activity. Estimating these waveforms can be useful in some clinical applications, such as the estimation of the spread patterns of the activity in epilepsy (Quintero-Rincón et al., 2016).

The paper is organized as follows: Section 2 presents the proposed Bayesian model. Section 3 introduces the partially collapsed Gibbs sampler used to generate samples distributed according to the posterior of this model and the Metropolis–Hastings moves that are used to accelerate the convergence of the sampler. Experimental results conducted for both synthetic and real data are presented in Section 4. Conclusions are finally reported in Section 5.

2. Proposed method

EEG source localization is an inverse problem consisting in estimating the brain activity of a patient from EEG measurements taken from M electrodes during T time samples. In a distributed source model, the brain activity is represented by a finite number of dipoles located at fixed positions on the brain cortex. More precisely, we consider N dipoles located on the cortical surface and oriented orthogonally to it (see Hallez et al., 2007 for motivation). The EEG measurement matrix $\boldsymbol{Y} \in \mathbb{R}^{M\times T}$ can be written as

$$\boldsymbol{Y} = \boldsymbol{H}\,\boldsymbol{X} + \boldsymbol{E} \tag{1}$$

where $\boldsymbol{X} \in \mathbb{R}^{N\times T}$ contains the dipole amplitudes, $\boldsymbol{H} \in \mathbb{R}^{M\times N}$ is the lead-field matrix and $\boldsymbol{E}$ is the additive noise.

2.1. Likelihood

It is very classical to assume that the noise samples are independent and identically distributed according to a Gaussian distribution (Grech et al., 2008). Note that when this assumption does not hold it is possible to estimate the noise covariance matrix from measurements that do not contain the signal of interest and use it to whiten the data (Maris, 2003). Denoting as ${\sigma_n}^2$ the noise variance, the independence assumption leads to the likelihood

$$f(\boldsymbol{Y}|\boldsymbol{\theta}) = \prod_{t=1}^{T} \mathcal{N}\left(\boldsymbol{y}^t \middle| \boldsymbol{H}\boldsymbol{x}^t, \sigma_n^2 \mathbb{I}_M\right) \tag{2}$$

where $\mathbb{I}_M$ is the identity matrix of size M and $\boldsymbol{\theta} = \{\boldsymbol{X}, \sigma_n^2\}$ contains the unknown parameters.

2.2. Prior distributions

2.2.1. Brain activity X

To promote structured sparsity of the source activity, we consider the weighted ℓ_{20} mixed pseudo-norm

$$\| \boldsymbol{X} \|_{20} = \#\{ i | \sqrt{v_i}\, \| \boldsymbol{x}_i \|_2 \neq 0 \} \tag{3}$$

where $v_i = \| \boldsymbol{h}^i \|_2$ is a weight introduced to compensate the depth-weighting effect (Grech et al., 2008; Uutela et al., 1999) and $\#\mathcal{S}$ denotes the cardinal of the set $\mathcal{S}$. Since this prior leads to intractable computations, we propose to approximate it by a multivariate Laplace Bernoulli prior for each row of $\boldsymbol{X}$[1]

$$f(\boldsymbol{x}_i|z_i, \lambda) \propto \begin{cases} \delta(\boldsymbol{x}_i) & \text{if } z_i = 0 \\ \exp\left(-\frac{1}{\lambda}\sqrt{v_i}\, \| \boldsymbol{x}_i \|_2\right) & \text{if } z_i = 1 \end{cases} \tag{4}$$

where $\propto$ means "proportional to", λ is the parameter of the exponential distribution and $\boldsymbol{z} \in \{0, 1\}^N$ is a vector indicating if the rows of $\boldsymbol{X}$ are non-zero. To make the analysis easier we introduce the hyperparameter $a = \frac{\sigma_n^2}{\lambda^2}$ leading to

$$f(\boldsymbol{x}_i|z_i, a, \sigma_n^2) \propto \begin{cases} \delta(\boldsymbol{x}_i) & \text{if } z_i = 0 \\ \exp\left(-\sqrt{\frac{v_i a}{\sigma_n^2}}\, \| \boldsymbol{x}_i \|_2\right) & \text{if } z_i = 1. \end{cases} \tag{5}$$

The elements z_i are then assigned a Bernoulli prior with parameter $\omega \in [0, 1]$

$$f(z_i|\omega) = \mathcal{B}(z_i|\omega). \tag{6}$$

Note that the Dirac delta function $\delta(.)$ in the prior of $\boldsymbol{x}_i$ promotes sparsity while the Laplace distribution regulates the amplitudes of the non-zero rows. The parameter ω allows the importance of these two terms to be balanced. In particular, $\omega = 0$ yields $\boldsymbol{X} = 0$ whereas $\omega = 1$ leads to the Bayesian formulation of the group-lasso (Yuan

[1] In this paper, we will denote as $\boldsymbol{m}_i$ the i-th row of the matrix $\boldsymbol{M}$ and as $\boldsymbol{m}^j$ its j-th column.