NeuroImage 87 (2014) 427–443

Contents lists available at ScienceDirect

NeuroImage

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

A Subspace Pursuit-based Iterative Greedy Hierarchical solution to the neuromagnetic inverse problem☆

Behtash Babadi [a,b,*,1], Gabriel Obregon-Henao [a,1], Camilo Lamus [a,b], Matti S. Hämäläinen [c,d], Emery N. Brown [a,b,d,e], Patrick L. Purdon [a,*]

[a] Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital, USA
[b] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, USA
[c] Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, USA
[d] Harvard-MIT Division of Health Sciences and Technology, USA
[e] Institute for Medical Engineering and Science, Massachusetts Institute of Technology, USA

ARTICLE INFO

Article history:
Accepted 5 September 2013
Available online 18 September 2013

Keywords:
MEG
EEG
Source localization
Sparse representations
Compressed sensing
Greedy algorithms
Evoked fields analysis

ABSTRACT

Magnetoencephalography (MEG) is an important non-invasive method for studying activity within the human brain. Source localization methods can be used to estimate spatiotemporal activity from MEG measurements with high temporal resolution, but the spatial resolution of these estimates is poor due to the ill-posed nature of the MEG inverse problem. Recent developments in source localization methodology have emphasized temporal as well as spatial constraints to improve source localization accuracy, but these methods can be computationally intense. Solutions emphasizing spatial sparsity hold tremendous promise, since the underlying neurophysiological processes generating MEG signals are often sparse in nature, whether in the form of focal sources, or distributed sources representing large-scale functional networks. Recent developments in the theory of compressed sensing (CS) provide a rigorous framework to estimate signals with sparse structure. In particular, a class of CS algorithms referred to as greedy pursuit algorithms can provide both high recovery accuracy and low computational complexity. Greedy pursuit algorithms are difficult to apply directly to the MEG inverse problem because of the high-dimensional structure of the MEG source space and the high spatial correlation in MEG measurements. In this paper, we develop a novel greedy pursuit algorithm for sparse MEG source localization that overcomes these fundamental problems. This algorithm, which we refer to as the Subspace Pursuit-based Iterative Greedy Hierarchical (SPIGH) inverse solution, exhibits very low computational complexity while achieving very high localization accuracy. We evaluate the performance of the proposed algorithm using comprehensive simulations, as well as the analysis of human MEG data during spontaneous brain activity and somatosensory stimuli. These studies reveal substantial performance gains provided by the SPIGH algorithm in terms of computational complexity, localization accuracy, and robustness.

Introduction

Magnetoencephalography (MEG) is among the most popular non-invasive modalities of brain imaging, and provides measurements of the collective electromagnetic activity of neuronal populations with high temporal resolution on the order of milliseconds. This technique has been used to study the mechanisms of language, cognition, sensory function, and brain oscillations, as well as the localization of epileptic seizures.

Localizing active regions in the brain from MEG measurements requires solving the high-dimensional ill-posed neuromagnetic inverse problem: given measurements from a limited number of sensors (~10^2 sensors) and a model for the observation process, the goal is to estimate spatiotemporal cortical activity over numerous sources (~10^4 sources). The ill-posedness of this problem is due not only to the limited number of measurements compared to unknowns, but also due to the high spatial dependency between measurements. The ill-posed nature of the inverse problem limits the spatial resolution of MEG. In comparison, other imaging modalities such as functional magnetic resonance imaging (fMRI) or positron emission tomography (PET) have high spatial resolution, but poor temporal resolution.

In the past two decades, various inverse solutions have been proposed for MEG and Electroencephalography (EEG) source localization.

☆ This work has been presented in part at the 2012 IEEE Engineering in Medicine and Biology Conference (Obregon-Henao et al., 2012).

* Corresponding authors at: Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital, USA.
E-mail addresses: behtash@nmr.mgh.harvard.edu (B. Babadi), patrickp@nmr.mgh.harvard.edu (P.L. Purdon).

[1] The first two authors have contributed equally to this work.

1053-8119/$ – see front matter
http://dx.doi.org/10.1016/j.neuroimage.2013.09.008

The earliest proposals appearing in the literature aimed at regularized least squares solutions to the neuromagnetic inverse problem (Gorodnitsky et al., 1995; Hämäläinen and Ilmoniemi, 1994; Pascual-Marqui et al., 1994; Uutela et al., 1998; Van Veen et al., 1997). Later on, inverse solutions in the framework of Bayesian estimation were introduced, with the underlying assumption of temporal independence (Mattout et al., 2006; Nummenmaa et al., 2007; Phillips et al., 2005; Sato et al., 2004; Wipf and Nagarajan, 2009). In order to impose spatio-temporal smoothness on the inverse solution, subsequent algorithms in the Bayesian framework considered the design of spatio-temporal priors (Baillet and Garnero, 1997; Bolstad et al., 2009; Daunizeau and Friston, 2007; Daunizeau et al., 2006; Friston et al., 2008; Greensite, 2003; Limpiti et al., 2009; Trujillo-Barreto et al., 2008; Zumer et al., 2008) or employed linear state-space models (Galka et al., 2004; Lamus et al., 2007; Long et al., 2011; Yamashita et al., 2004). Despite their improved accuracy in source localization, many of these more recent solutions suffer from unwieldy computational complexity.

Sparse solutions to the MEG/EEG source localization problem have received renewed attention in recent years (Durka et al., 2005; Gorodnitsky et al., 1995; Gramfort et al., 2011, 2012; Ou et al., 2009; Tian and Li, 2011; Tian et al., 2012; Valdés-Sosa et al., 2009; Vega-Hernández et al., 2008). Although many of these methods are computationally demanding, they have been shown to enhance the accuracy of source localization. The sparsity constraints underlying these methods express the intuition that out of the $\sim 10^4$ potential sources, only a small number are truly active. In many applications of MEG/EEG source localization, such as sensory or cognitive studies, the underlying cortical domains responsible for the processing are relatively focal and thus sparse. Processes that are spatially-distributed can also be sparse in some basis: for instance, resting-state fMRI dynamics are broadly distributed across the cortex, yet appear to be organized within a small number of specific networks (Damoiseaux et al., 2006). Spatial sparsity is therefore not only a mathematically attractive constraint, but also one that is consistent with the neurophysiological processes underlying EEG and MEG.

The advent of compressed sensing (CS) theory has paved the way to establish a rigorous framework for efficient sampling and estimation of signals with underlying sparse structures (Candès et al., 2006; Donoho, 2006). CS methods have found applications in various disciplines such as communication systems, computational biology, geophysics, and medical imaging (see Bruckstein et al., 2009 for a survey of the CS results). The problem of recovering a sparse unknown signal given limited observations is combinatorial and NP-hard in nature (Bruckstein et al., 2009). Several solutions to this problem have been proposed, which can be categorized into optimization-based methods, greedy pursuit methods, and coding theoretic/Bayesian methods. These solution categories pertain to different regimes of sparseness as well as different ranges of computational complexity. Moreover, they all require certain notions of regularity in the measurement structure which must be satisfied in order to guarantee convergence and sparse recovery.

In this work, we develop an inverse solution to the MEG inverse problem in the context of CS theory that achieves both high localization accuracy and very low computational cost. This algorithm is based on a class of greedy pursuit algorithms known as Subspace Pursuit (Dai and Milenkovic, 2009) or Compressive Sampling Matching Pursuit (CoSaMP) (Needell and Tropp, 2009). These algorithms can achieve high recovery accuracy with low computational complexity, but are difficult to apply directly to the MEG inverse problem because of the high-dimensional structure of the MEG source space and the high spatial correlation in MEG measurements. We introduce novel algorithms that address both of these issues in a principled manner. Through comprehensive simulation studies and analysis of human MEG data, we demonstrate the utility of the proposed method, and its superior performance compared to existing methods.

Methods

The MEG observation model

The MEG data consists of a multidimensional time series recorded using an array of sensors located over the scalp. Let N denote the number of such sensors, and $t = 1,\cdots,T$ denote the discrete time stamps corresponding to the sampling frequency f_s. We denote by $y_{i,t}$ the magnetic measurement of the ith sensor at time t, for $1 \le i \le N$ and $1 \le t \le T$. Let $\mathbf{y}_t := [y_{1,t}, y_{2,t}, \cdots, y_{N,t}]'$ denote the vector of measurements at time t. Finally, we denote the multidimensional observation time series in the interval $[0,T]$ by the $N \times T$ matrix $\mathbf{Y} := [\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_T]$.

We pose the source localization problem over a distributed source model comprising dipoles with fixed locations and possibly variable orientations, representing the cortical activity. Let M be the total number of dipole sources distributed across the cortex, and let $x_{i,t}$ denote the amplitude of the ith dipole at time t. Denoting by $\mathbf{x}_t := [x_{1,t}, x_{2,t}, \cdots, x_{M,t}]'$ the vector of dipole amplitudes at time t, the source space can be fully characterized by the $M \times T$ matrix $\mathbf{X} := [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_T]$.

For a fixed configuration of dipoles, the observation matrix $\mathbf{Y}$ can be related to the source activity matrix $\mathbf{X}$ as follows:

$$\mathbf{Y} = \mathbf{G}\mathbf{X} + \mathbf{V}, \tag{1}$$

where $\mathbf{G} \in \mathbb{R}^{N \times M}$ is the *lead field matrix* and $\mathbf{V} := [\mathrm{v}_1, \mathrm{v}_2, \cdots, \mathrm{v}_T] \in \mathbb{R}^{N \times T}$ is the observation noise matrix. The lead field matrix $\mathbf{G}$ is a mapping from the source space to the sensor space and can be computed using a quasi-static approximation to the Maxwell's equations (Hämäläinen et al., 1993). The observation noise is assumed to be zero-mean Gaussian, with spatial covariance matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$ with no temporal correlation.

The MEG inverse problem corresponds to estimating $\mathbf{X}$ given $\mathbf{Y}$, $\mathbf{G}$ and the statistics of $\mathbf{V}$. The traditional source spaces used for MEG source localization have a dimension of $M \sim 10^3$–10^5, whereas the number of sensors is typically $N \sim 10^2$. Since $M \gg N$, the MEG inverse problem is highly ill-posed and hence requires constructing appropriate spatio-temporal priors in order to estimate $\mathbf{X}$ reliably.

The Minimum Norm Estimate

As outlined in the Introduction, various source localization techniques have been developed since the invention of MEG. Arguably, the Minimum Norm Estimate (MNE) inverse solution is the most widely-used MEG source localization technique (Hämäläinen and Ilmoniemi, 1994). In what follows, we give a brief overview of the MNE and one of its recent variants.

The MNE inputs $\mathbf{Y}$, $\mathbf{G}$, $\mathbf{C}$ and a spatial prior covariance $\mathbf{Q}/\lambda^2 \in \mathbb{R}^{M \times M}$ on $\mathbf{X}$ with λ being a scaling factor, and solves for:

$$\hat{\mathbf{X}}^{\mathrm{MNE}}(\mathbf{Y}, \mathbf{G}, \mathbf{C}, \mathbf{Q}, \lambda) := \arg\min_{\mathbf{X}} \sum_{t=1}^{T} \left\{ \|\mathbf{y}_t - \mathbf{G}\mathbf{x}_t\|^2_{\mathbf{C}^{-1}} + \lambda^2 \|\mathbf{x}_t\|^2_{\mathbf{Q}^{-1}} \right\}. \tag{2}$$

The role of the spatial prior covariance $\mathbf{Q}$ is to weight the penalization of the energy of the estimated sources across the source space (Hämäläinen and Ilmoniemi, 1994; Lamus et al., 2012). Moreover, the scaling factor λ can be viewed as a regularization parameter, controlling the $\mathbf{Q}$-weighted ℓ_2-norm of the estimate. The minimization is separable in time, and the estimate can be expressed in closed form as:

$$\hat{\mathbf{X}}^{\mathrm{MNE}}(\mathbf{Y}, \mathbf{G}, \mathbf{C}, \mathbf{Q}, \lambda) = \mathbf{Q}\mathbf{G}' \left(\mathbf{G}\mathbf{Q}\mathbf{G}' + \lambda^2 \mathbf{C} \right)^{-1} \mathbf{Y}. \tag{3}$$

It is more convenient to compute the MNE estimate in a whitened model in favor of computational stability. The whitened model is obtained by the pre-multiplication of the observation model by $\mathbf{C}^{-1/2}$. Let $\widetilde{\mathbf{Y}} := \mathbf{C}^{-1/2}\mathbf{Y}$ and $\widetilde{\mathbf{G}} := \mathbf{C}^{-1/2}\mathbf{G}$ denote the whitened observation