

A distributed spatio-temporal EEG/MEG inverse solver

Wanmei Ou ^{a,*}, Matti S. Hämäläinen ^b, Polina Golland ^a

^a Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

^b Department of Radiology, Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA

ARTICLE INFO

Article history:

Received 21 December 2007

Revised 14 May 2008

Accepted 30 May 2008

Available online 14 June 2008

Keywords:

EEG

MEG

Inverse solver

ℓ_1 -norm

Temporal basis functions

Second-order cone programming

ABSTRACT

We propose a novel ℓ_1/ℓ_2 -norm inverse solver for estimating the sources of EEG/MEG signals. Based on the standard ℓ_1 -norm inverse solvers, this sparse distributed inverse solver integrates the ℓ_1 -norm spatial model with a temporal model of the source signals in order to avoid unstable activation patterns and “spiky” reconstructed signals often produced by the currently used sparse solvers. The joint spatio-temporal model leads to a cost function with an ℓ_1/ℓ_2 -norm regularizer whose minimization can be reduced to a convex second-order cone programming (SOCP) problem and efficiently solved using the interior-point method. The efficient computation of the SOCP problem allows us to implement permutation tests for estimating statistical significance of the inverse solution. Validation with simulated and human MEG data shows that the proposed solver yields source time course estimates qualitatively similar to those obtained through dipole fitting, but without the need to specify the number of dipole sources in advance. Furthermore, the ℓ_1/ℓ_2 -norm solver achieves fewer false positives and a better representation of the source locations than the conventional ℓ_2 minimum-norm estimates.

© 2008 Elsevier Inc. All rights reserved.

Introduction

Electroencephalography (EEG) and magnetoencephalography (MEG) are widely used for functional brain mapping. With appropriate source estimation algorithms one can locate the activated regions, as well as estimate their dynamics. The non-invasive nature of EEG and MEG makes these methods particularly suitable for neuroscience research and clinical practice, such as surgical planning for epilepsies (Knake et al., 2006).

Localizing activated regions from EEG/MEG data involves solving an electromagnetic inverse problem. Unfortunately, even with perfect knowledge of the electric and magnetic fields outside of the source region, this problem does not have a unique solution because there are currents which are either electrically or magnetically silent, or both. Moreover, solutions might not depend continuously on the data without regularization, which means small errors in measurements might cause errors of arbitrary size in the estimated sources. These two characteristics make the inverse problem ill-posed in the sense of Hadamard (1902). This paper introduces an integrated spatio-temporal regularizer to overcome the instabilities of standard sparse inverse solvers.

There are two main types of inverse solvers for EEG/MEG source estimation: discrete parametric solvers, also known as dipole fitting, and distributed inverse solvers. The standard dipole fitting algorithms estimate the location, orientation, and amplitudes of a fixed number of

current dipoles (Wood, 1982; Scherg and Von Cramon, 1985; Mosher et al., 1992; Uutela et al., 1998). In contrast, distributed solvers discretize the source space into locations on the cortical surface or in the brain volume without explicitly controlling the number of current dipoles. The desired solution is computed by minimizing a cost function that depends on all sources in the source space, such as an overall minimum power or minimum current (Hämäläinen and Ilmoniemi, 1984; Dale and Sereno, 1993; Wang et al., 1993; Uutela et al., 1999).

Dipole fitting usually provides robust estimates for activation signals, but localization is challenging when several sources are active because the associated cost function depends non-linearly on the dipole locations. Additionally, the quality of the results degrades when the assumed number of dipoles differs from the true number (see, e.g. Wood, 1982; Hari and Forss, 1999). Although it is possible to obtain an initial guess for the number of dipoles through singular-value decomposition (SVD) of the recordings (Huang et al., 1998), this method is sensitive to the user-defined thresholds and is problematic in the presence of correlated source signals. Furthermore, it has been argued that a set of current dipoles may not be a good model for activations with relatively large spatial extents (Jerbi et al., 2004).

Not restricted to a fixed number of dipoles, the distributed solvers estimate the amplitude of all possible source locations. The widely used minimum-norm estimate (MNE) (Hämäläinen and Ilmoniemi, 1984; Dale and Sereno, 1993; Wang et al., 1993) recovers a source distribution with minimum overall energy (or minimum ℓ_2 -norm) that produces data consistent with the measurements. Although the ℓ_2 -norm method leads to an efficient linear inverse operator, the MNE solutions are often too diffuse. In particular, MNE is not appropriate for localization of early sensory activations and focal epilepsy, which have

* Corresponding author.

E-mail addresses: wanmei@mit.edu (W. Ou), msh@nmr.mgh.harvard.edu (M.S. Hämäläinen), polina@csail.mit.edu (P. Golland).

been shown to be focal in intracranial experiments (Barth et al., 1982; Allison et al., 1989). To overcome this property, the Focal Underdetermined System Solver (FOCUSS) (Gorodnitsky and Rao, 1997) augments the MNE solver with a recursive weighting scheme. FOCUSS has been shown to be equivalent to a p -norm solver where $p \leq 1$ (Rao and Kreutz-Delgado, 1999). Other regularizers based on a norm penalty can provide bias towards sparsity. Among them, the minimum current estimate (MCE) (minimum ℓ_1 -norm) is the most popular (Uutela et al., 1999).

One of the drawbacks of the conventional ℓ_1 -norm inverse solvers, as well as other focal solvers such as FOCUSS, is their sensitivity to noise. Similar to other distributed solvers, the conventional ℓ_1 -norm solvers are typically applied to each time sample in the data separately. The solvers' sensitivity to noise causes the estimated activations to "jump" among neighboring spatial locations from one time instant to another. Equivalently, the time course at a particular location can show substantially "spiky" discontinuities when viewed over time. To avoid this problem one commonly averages the time courses across adjacent sites, at the expense of spatial resolution.

Two alternative approaches utilize temporal constraints to improve reconstruction accuracy: a direct application of the temporal constraint as a regularizer in the cost function and a use of temporal basis functions. In Baillet and Garnero (1997), Brooks et al. (1997), Schmitt et al. (2001), Galka et al. (2004), Zhang et al. (2005), and Camilo et al. (2007) a regularizer is explicitly incorporated in the cost function to model the smoothness of the current source distributions between consecutive time instants. For example Baillet and Garnero (1997) encourages small residuals in the least squares estimates of the current sources between the current time point and the previous one. The studies of Galka et al. (2004), Zhang et al. (2005), and Camilo et al. (2007) propose a state-space model with smooth state transitions, i.e., the current source distributions between consecutive time instants. The temporal regularization terms in Brooks et al. (1997) and Schmitt et al. (2001) are expressed as the ℓ_2 -norm of the output of the current sources passed through a pre-designed low-pass filter in the time domain. While conceptually, these methods address the problem of sensitivity to noise, their implementation requires substantial amount of computation, except for a limited number of low-pass filters. The work of Zhang et al. (2005) provides a comprehensive comparison among the above regularization methods. Taking a significantly different approach to reducing the sensitivity to noise, the vector-based spatio-temporal minimum ℓ_1 -norm solver (VESTAL) projects the sample-wise ℓ_1 -norm estimates into the signal subspace defined by a set of temporal basis functions (Huang et al., 2006). Models based on temporal basis functions have also been proposed for other types of inverse solvers. For instance, Geva (1998) constructed a basis set using wavelets and computed inverse solutions for each basis function separately using dipole fitting. Trujillo-Barreto et al. (2007) explored the use of wavelets as a temporal model in the context of distributed solutions.

Similar to MCE and VESTAL, we employ the ℓ_1 -norm regularizer to encourage spatial sparsity. We reduce MCE's sensitivity to noise by incorporating our knowledge of the temporal characteristics of the source signals. Specifically, we assume that the source signals are linear combinations of *multiple* temporal basis functions, and apply the distributed inverse solver to the coefficients of all basis functions simultaneously. We utilize the conventional definition of amplitude, the ℓ_2 -norm, to summarize the activation strength at each location. Since the ℓ_2 -norm does not encourage sparsity, many coefficients for an active location are usually non-zero in the inverse solution.

This integrated spatio-temporal regularizer is at the core of our ℓ_1/ℓ_2 -norm inverse solver. The ℓ_1/ℓ_2 -norm regularizer was suggested in farfield narrowband sensor array applications (Malioutov et al., 2005) to model the diffuse temporal structure of the source signals.

Although we focus on the EEG/MEG application, the proposed framework is also applicable to computed tomography reconstruction, with modifications to the spatial model so as to encourage piece-wise constant solutions (i.e., ℓ_1 -norm on spatial derivatives).

To summarize, the proposed solver imposes ℓ_1 -norm regularization in space and ℓ_2 -norm regularization in the temporal domain. The resulting inverse problem can be formulated as a second-order cone programming (SOCP) problem and solved efficiently using the interior-point method (Alizadeh and Goldfarb, 2001). In contrast to VESTAL, which uses the spatial and the temporal models separately in a two-step estimation procedure, our solver unifies the two models into a single regularizer in order to avoid error propagation from the first estimation step to the second one. Experimental comparisons in the Simulation studies section reveal that the joint spatio-temporal model implicitly increases the signal-to-noise ratio (SNR) and achieves a more accurate reconstruction.

There are various ways to obtain the temporal basis functions to represent the source signals, including Fourier and wavelet decompositions. However, a compact representation of the signals, i.e., a small number of basis functions, can significantly reduce computation requirements. In this work, we generate the basis set through the singular-value decomposition of the sensor data, which often closely reflects the temporal structure of the source signals. We examine the effects of the basis selection on the resulting reconstruction by varying the singular-value cutoff and the noise amplitude in simulation experiments.

Existing ℓ_1 -norm related solvers also lack a consensus in handling the free-orientation source reconstruction. In the conventional MCE (Uutela et al., 1999) and its cortically-constrained variant (Lin et al., 2006), the orientations of the sources are determined prior to invoking the ℓ_1 -norm minimizer. Uutela et al. estimated the orientations from an initial MNE solution, while Lin et al. utilized both the MNE solution and anatomical information. The method proposed in Matsuura and Okabe (1999) alternates between computing the inverse solution and estimating the source orientation, but it suffers from convergence issues and requires intensive computations. VESTAL (Huang et al., 2006) applies the ℓ_1 -norm to each source component via a bias-reduction scheme in the free-orientation case. Since the ℓ_2 -norm is invariant with respect to rotations of the local coordinate system at each source, it is straightforward to extend our method to include free orientations, as we demonstrate in this paper.

We also construct a permutation test for the ℓ_1/ℓ_2 -norm inverse solution. Here, we follow the framework proposed in Pantazis et al. (2005) for constructing the null hypothesis distribution. In their work, permutations can be applied either before or after the reconstruction due to linearity of the MNE inverse operator. In contrast, permutations must be performed prior to the ℓ_1/ℓ_2 -norm inverse operation. Although a large number of samples are required to yield an accurate estimate of the null distribution, we can still apply the permutation testing since the interior-point method for SOCP is quite efficient.

The remainder of this paper is organized as follows. The Methods section describes the ℓ_1/ℓ_2 -norm inverse solver. We then briefly address implementation issues. The Results section presents experimental results using simulated and human MEG data, followed by a discussion and conclusions.

Methods

In this section, we first provide background and define our notation, and then we describe the spatio-temporal model and formulate the ℓ_1/ℓ_2 -norm inverse solver as an SOCP problem for a fixed-orientation source model. The Temporal basis selection section explains the temporal basis construction scheme. We then present an extension of the proposed solver to free-orientation cases and a

permutation test for accessing the significance of the resulting inverse solutions.

Background and notation

Under the *quasi-static* approximation of Maxwell's equations, the observed EEG/MEG signals $\mathbf{y}(t)$ at time t are linear functions of the current sources $\mathbf{s}(t)$:

$$\mathbf{y}(t) = \mathbf{A}\mathbf{s}(t) + \mathbf{e}(t) \tag{1}$$

where \mathbf{A} is $M \times N$ lead-field matrix. $\mathbf{e}(t) \sim \mathcal{N}(\mathbf{0}, \Sigma)$ is the measurement noise; the noise covariance Σ can be estimated from pre-stimulus data. $\mathbf{s}(t)$, $N \times 1$, and $\mathbf{y}(t)$, $M \times 1$, are vectors in the source space and the signal space, respectively. The number of sources $N (\sim 10^3 - 10^4)$ is much larger than the number of measurements $M (\sim 10^2)$, leading to an infinite number of solutions satisfying Eq. (1) even for $\mathbf{e}(t) = \mathbf{0}$. Without loss of generality, we apply spatial whitening based on the estimated noise covariance Σ to both the data and the lead-field matrix, leading $\mathbf{e}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to in the derivations.

Spatio-temporal model

The quasi-static assumption allows us to conduct inverse estimation for each time instant independently. However, this often results in highly variable source time courses. The large variability is particularly prominent in the focal solvers, such as the MCE, due to their non-linear nature. To mitigate this problem, we utilize the knowledge of the temporal properties of the source signals to further constrain the solution. To this end, we express the data model in Eq. (1) for all time instants as:

$$\mathbf{Y} = \mathbf{A}\mathbf{S} + \mathbf{E} \tag{2}$$

where $\mathbf{Y} = [\mathbf{y}(1), \mathbf{y}(2), \dots, \mathbf{y}(T)]$ is an $M \times T$ matrix that contains EEG/MEG measurements for all T temporal samples, and \mathbf{S} is an $N \times T$ matrix that represents the source signals. Here, we assume that noise \mathbf{E} is independent in time, i.e., $E[\mathbf{E}^T \mathbf{E}] = \mathbf{I}$. Time-dependent noise models as those suggested in Huizenga et al. (2002) and Bijma et al. (2005) can be incorporated into the estimation procedure as well; this is a topic of future work.

The underlying sources of EEG/MEG measurements, closely related to the postsynaptic potentials (Hämäläinen et al., 1993), are relatively smooth with occasional deflections. For example, a typical response from the contralateral primary somatosensory area has relatively strong deflections immediately after the stimulus (20–40 ms) followed by a smoother time course (Weerd and Kap, 1981). Hence, the activation signals are neither sparse nor diffuse in time. Direct temporal regularization using the ℓ_1 -norm or the ℓ_2 -norm is therefore not appropriate. To model the time-varying frequency content of the signals, we assume that the source signals are linear combinations of multiple orthonormal temporal basis functions, $\mathbf{V} = [\vec{v}_1, \vec{v}_2, \dots, \vec{v}_K]$, that collectively capture the temporal properties of the source signals. $\vec{v}_k, T \times 1$, denotes the k^{th} basis function. In the Temporal basis selection section, we will discuss how to obtain the basis appropriate for the reconstruction. We assume that the basis functions are orthonormal; if they are not, minor modifications of the remaining derivations are needed, as we addressed in the Discussion section.

Projecting both the sensor recordings and the source signals onto the basis functions, the new variables $\tilde{\mathbf{Y}} = \mathbf{Y}\mathbf{V}$ and $\tilde{\mathbf{S}} = \mathbf{S}\mathbf{V}$ are the corresponding projection coefficients. $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{S}}$ are of size $M \times K$ and $N \times K$, respectively. The (n, k) element of $\tilde{\mathbf{S}}$, \tilde{s}_{nk} , indicates the k^{th} coefficient for the source signal at location n . We can rewrite the original data model in Eq. (2) in the transformed domain:

$$\tilde{\mathbf{Y}} = \tilde{\mathbf{A}}\tilde{\mathbf{S}} + \tilde{\mathbf{E}} \tag{3}$$

where $\tilde{\mathbf{E}} = \mathbf{E}\mathbf{V}$. We use $\tilde{\mathbf{e}}_k$ to denote the k^{th} column of $\tilde{\mathbf{E}}$. The temporal independence assumption on \mathbf{E} and orthonormality of \mathbf{V} imply that $\tilde{\mathbf{e}}_k$ and $\tilde{\mathbf{e}}_{k'}$ are independent for $k \neq k'$ and that $\tilde{\mathbf{e}}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Eq. (3) is still under-determined, containing MK equations with NK variables.

To compute inverse solutions for all K basis functions simultaneously, we extend the existing regularizers to use the signal magnitude in the subspace spanned by \mathbf{V} , $\sqrt{\sum_{k=1}^K \tilde{s}_{nk}^2}$, as an indicator of the activation status at location n . In other words, we apply ℓ_2 -norm regularization to the K coefficients for each source location. Because we choose to work with orthonormal basis functions, the ℓ_2 -norm of the reconstructed source signal in the temporal domain is equal to the ℓ_2 -norm in the transformed domain. However, we find it more intuitive to present the model in the transformed domain.

In addition, we assume that the sources exhibit a spatially sparse pattern. This assumption represents the relatively compact source regions typically activated in the sensory areas. To obtain a focal inverse solution, we should ideally employ the ℓ_0 -norm as the spatial regularizer. However, the ℓ_0 -norm regularization leads to an NP-hard optimization problem. In practice, under some regularity conditions (Donoho and Elad, 2003), the ℓ_1 -norm regularizer leads to solutions identical to those produced by the ℓ_0 -norm regularizer. Even when the solution obtained through the ℓ_1 -norm regularization is different from the one produced by the ℓ_0 -norm regularization, it is still more sparse than that obtained with the ℓ_2 -norm regularizer.

With the ℓ_1 -norm regularizer in the spatial domain and the ℓ_2 -norm regularizer in the temporal domain, we incorporate the integrated spatio-temporal ℓ_1/ℓ_2 -norm regularizer

$$\|\tilde{\mathbf{S}}\|_{\ell_1}^2 = \sum_{n=1}^N \sqrt{\sum_{k=1}^K \tilde{s}_{nk}^2} \tag{4}$$

into the estimation problem:

$$\tilde{\mathbf{S}}^* = \arg \min_{\tilde{\mathbf{S}}} \|\tilde{\mathbf{Y}} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}\|_F^2 + \lambda \|\tilde{\mathbf{S}}\|_{\ell_1}^2 \tag{5}$$

$$= \arg \min_{\tilde{\mathbf{S}}} \sum_{k=1}^K \|\tilde{\mathbf{y}}_k - \tilde{\mathbf{A}}\tilde{\mathbf{s}}_k\|_{\ell_2}^2 + \lambda \|\tilde{\mathbf{S}}\|_{\ell_1}^2 \tag{6}$$

where $\tilde{\mathbf{s}}_k$ and $\tilde{\mathbf{y}}_k$ are the k^{th} column vectors in $\tilde{\mathbf{S}}$ and $\tilde{\mathbf{Y}}$. $\|\cdot\|_F$ and $\|\cdot\|_{\ell_2}$ (i.e., $\|\mathbf{x}\|_{\ell_2} = \sqrt{\mathbf{x}^T \mathbf{x}}$) denote the Frobenius norm of a matrix and the standard ℓ_2 -norm of a vector, respectively. λ controls the regularization strength. We will discuss how to select this parameter in the Sensitivity to regularization strength section. After we obtain the optimal coefficients $\tilde{\mathbf{S}}^*$, the reconstructed source signals are linear combinations of the temporal basis functions:

$$\mathbf{S}^* = \tilde{\mathbf{S}}^* \mathbf{V}^T. \tag{7}$$

In this paper, we formulate the inverse problem as a regularized optimization. It also has an equivalent Bayesian interpretation. The first term in Eq. (5) can be considered as the negative log likelihood under white Gaussian noise. The second term corresponds to the negative log prior of the source signals, which in our case is Laplacian in space and Gaussian in time.

From the ℓ_1/ℓ_2 -norm regularizer to second-order cone programming (SOCP)

We cannot directly apply gradient based methods to the optimization problem specified by Eq. (6) since the ℓ_1/ℓ_2 -norm penalty term is not differentiable at zero. However, Eq. (6) can be reduced to the SOCP problem by converting the original unconstrained optimization problem to a constrained one:

$$\langle \tilde{\mathbf{S}}^*, \mathbf{q}^*, \mathbf{z}^*, \mathbf{w}^*, \mathbf{r}^* \rangle = \arg \min_{\langle \tilde{\mathbf{S}}, \mathbf{q}, \mathbf{z}, \mathbf{w}, \mathbf{r} \rangle} (q + \lambda z) \tag{8}$$

$$\text{s.t. } \|\tilde{\mathbf{y}}_k - \mathbf{A}\tilde{\mathbf{s}}_k\|_{\mathcal{L}_2}^2 \leq w_k \quad \forall k = 1, \dots, K \quad (9)$$

$$\sum_{k=1}^K w_k \leq q \quad (10)$$

$$\sqrt{\sum_{k=1}^K \tilde{s}_{nk}^2} \leq r_n \quad \forall n = 1, \dots, N \quad (11)$$

$$\sum_{n=1}^N r_n \leq z \quad (12)$$

New variables, $q, z, \{w_k\}_{k=1}^K$ and $\{r_n\}_{n=1}^N$, are introduced in the conversion procedure. w_k is an upper bound on the discrepancy between the measurements and the signals predicted by the estimated sources in the projection onto $\vec{\mathbf{v}}_k$. q is an upper bound on all w_k 's. r_n is an upper bound on the activation strength for location n . z is an upper bound on the \mathcal{L}_1 -norm of the activation strength of all N locations. At the minimum, the inequality constraints in Eqs. (9–12) are satisfied with equality; otherwise, the objective function can be further reduced.

Mathematically, a second-order cone of dimension D is defined as

$$\mathcal{Q}_D := \left\{ \mathbf{x} = (x_0, \bar{\mathbf{x}}) \in \mathcal{R}^D : x_0 \geq \|\bar{\mathbf{x}}\|_{\mathcal{L}_2} \right\} \quad (13)$$

where x_0 and $\bar{\mathbf{x}}$ denote the first element and the remaining elements of \mathbf{x} , respectively. We can see that Eq. (11) matches with the second-order cone definition. As shown in Alizadeh and Goldfarb (2001), a wide range of constrained formulations, including the quadratic constraint in Eq. (9), can be reduced to the canonical form of a second-order cone. For completeness, we provide the corresponding derivations in Appendix A.

An SOCP problem can be expressed in the canonical form that contains a linear objective function and the feasible set specified as an intersection of an affine linear manifold and the Cartesian product of second-order cones. Since the second-order cone defines a convex set, the feasible set of SOCP is convex. Therefore, SOCP is a convex optimization problem and its local minimum is the global minimum. In fact, for one-dimensional and two-dimensional cones, the second-order cone constraint in Eq. (13) can be reduced to linear constraints. As a result, the corresponding SOCP problem is reduced to a linear programming problem. It is also straightforward to show that the quadratically constrained quadratic programs are a subset of the SOCP problems. Furthermore, the SOCP problem is a special case of a semi-definite program. Therefore, SOCP can be solved efficiently using the primal-dual interior-point method (Alizadeh and Goldfarb, 2001), where Newton's method is employed to reduce the duality gap. Appendix B reviews the primal-dual interior-point method in application to SOCP.

Temporal basis selection

The formulation of our inverse solver is independent of the selected basis \mathbf{V} , but a compact representation of the signals can significantly reduce computation. We estimate the basis using the singular-value decomposition (SVD) of the measurements, which is often able to compactly capture the time-varying frequency content and significant differences in source signals from different regions. Another advantage of using data-adaptive temporal basis functions is that it avoids the difficulty of setting a set of basis functions to accommodate highly variable source signals due to different experimental tasks and subject-to-subject variations. Since according to Eq. (2) the sensor signals are linear combinations of the source signals, the temporal pattern of the source signals is present in the sensor signals as well. In fact, the standard dipole fitting procedure (Moshier et al.,

1992) also performs fitting of the K largest SVD components of the measurements that “adequately” describe the data.

The singular-value decomposition of data \mathbf{Y} is expressed as

$$\mathbf{Y} = \mathbf{U}_Y \mathbf{\Lambda}_Y \mathbf{V}_Y^T \quad (14)$$

each column in $\mathbf{U}_Y = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_M]$ denotes the electromagnetic field pattern; each column in $\mathbf{V}_Y = [\vec{\mathbf{v}}_1, \vec{\mathbf{v}}_2, \dots, \vec{\mathbf{v}}_M]$ denotes the temporal pattern. $\mathbf{\Lambda}_Y$ is a diagonal matrix of the singular values in a descending order. We assume more temporal samples than EEG/MEG sensors, which is usually true in practice due to fast sampling rates.

As mentioned before, we further assume that activation signals only lie in the subspace spanned by $\mathbf{V} = [\vec{\mathbf{v}}_1, \vec{\mathbf{v}}_2, \dots, \vec{\mathbf{v}}_K]$, but not in the subspace spanned by $\mathbf{V}^\perp = [\vec{\mathbf{v}}_{K+1}, \vec{\mathbf{v}}_{K+2}, \dots, \vec{\mathbf{v}}_M]$. Performing reconstruction in the signal subspace helps to stabilize the reconstructed source signals since they are constructed as linear combinations of relatively smooth temporal basis functions. In our experiments, the number of basis functions K is fixed in each reconstruction. We examine the performance of the proposed solver with varying K in the Sensitivity to noise and basis selection section and discuss alternative approaches to basis function selection in the Discussion section.

$\mathcal{L}_1/\mathcal{L}_2$ -norm reconstruction for the free-orientation source model

The free-orientation source model has been used both with volumetric source spaces covering the entire brain and with source locations restricted to the cortex only (Dale and Sereno, 1993; Pascual-Marqui et al., 1994). The results of the direct application of the \mathcal{L}_1 -norm regularizer to the three dipole moment coordinates depend on the parametrization of the local coordinates.

To extend our solver to free-orientation sources, we model the current dipole moment magnitude as the \mathcal{L}_2 -norm of the current dipole moments along the three coordinates. This model agrees with the conventional definition of magnitude. The resulting inverse problem is independent of the local coordinate system since the \mathcal{L}_2 -norm is invariant to rotations of orthogonal coordinates. In other words, our method models the spatially sparse activation pattern, but does not enforce sparsity on individual components of the dipole moments. This idea is analogous to the sensor array application (Malioutov et al., 2005) where signals are complex numbers; it was also independently developed and thoroughly evaluated by Ding and He for EEG source localization (Ding and He, 2008).

Extending our formulation in Eq. (4), $\tilde{\mathbf{s}}_{nk}$ is replaced by a three-dimensional vector, denoting the current dipole moments in the three coordinates $\tilde{s}_{x_{nk}}, \tilde{s}_{y_{nk}}$, and $\tilde{s}_{z_{nk}}$. The optimization problem in Eqs. (8–12) remains the same except that Eq. (11) is replaced by a constraint on the three coordinates:

$$\sqrt{\sum_{k=1}^K (\tilde{s}_{x_{nk}}^2 + \tilde{s}_{y_{nk}}^2 + \tilde{s}_{z_{nk}}^2)} \leq r_n \quad \forall n = 1 : N \quad (15)$$

In the original problem, each cone specified in Eq. (11) lies in a $K+1$ -dimensional subspace; in the free-orientation case, the corresponding cone is extended to a $3K+1$ -dimensional subspace. Since the feasible region is an intersection of hyper-cones and hyper-planes, the new formulation is still consistent with the SOCP structure.

Statistical significance testing

The non-linear nature of the \mathcal{L}_1 -norm related inverse operators, including the $\mathcal{L}_1/\mathcal{L}_2$ -norm inverse solver, presents a challenge in obtaining a sufficient statistic for hypothesis testing. Since there is no closed-form solution for the $\mathcal{L}_1/\mathcal{L}_2$ -norm solver, we employ a permutation test. We construct the null distribution by permuting

equal-length pre-stimulus and post-stimulus single-trial recordings. Under the null hypothesis defined as the absence of activation, the pre-stimulus and the post-stimulus recordings are equivalent. As described in Pantazis et al. (2005), in each permutation, we randomly select trials; for each selected trial, we swap its pre-stimulus and post-stimulus labels. Then we apply the inverse solver to the average data. This procedure preserves the noise covariance structure.

All the results presented in this paper are based on permutations. In this work, we control the false discovery rate (FDR) (Genovese et al., 2002; Efron and Tibshirani, 2002), over an amplitude-normalized source space. We first convert the source estimates into p -values, the ratio of permutations whose corresponding amplitude exceeds the original estimate, for each vertex and for each time instant separately. We then compute the FDR threshold (Genovese et al., 2002). We choose to use the p -values instead of the estimated amplitudes since source strength varies among activation regions and varies over time. For instance, the contralateral primary somatosensory (cSI) response is usually substantially larger than that in the ipsilateral secondary somatosensory (iSII) region. In addition, the N20 deflection is often weaker than the later deflections from the cSI area.

We cannot directly compare the activation maps created based on the permutation method with the corresponding statistics for MNE, the dynamic statistical parametric map (dSPM) (Dale and Sereno, 1993), because dSPM typically exhibits higher statistical power due to the quite restrictive Gaussian distribution assumption. On the other hand, the permutation method can capture activations for which the Gaussian assumption is not valid. Since dSPM is one of the most popular estimates in the EEG/MEG inverse community, we visually compare our results with dSPM side-by-side in the experimental section.

Implementation

Source space and lead-field matrix

For the computation of the lead-field matrix, we need a specification of the conductivity structure of the head, i.e., the forward model and the source space. In the forward computations for MEG, we employ the single-compartment boundary-element model (Hämäläinen and Sarvas, 1989; Oostendorp and Van Oosterom, 1989). For the source space, we restrict the locations of the sources to the cortical surface, which, in this work, is extracted using Freesurfer (Dale et al., 1999; Fischl et al., 1999). Due to the organization of the cortex, we can further constrain the source orientation to be perpendicular to the cortical surface. Independent of the choice of source space resolution, the orientation at each vertex is computed from the original triangulation of the cortical surface with a 0.65-mm grid spacing. Similar to other inverse solvers with orientation constraint, the sparse spacing of the source space may result in localization error, which could be avoided by denser sampling (Lin et al., 2006). Moreover, it is straightforward to alleviate this effect in our method by applying the free-orientation model presented in the last section.

In practise, the lead-field matrix \mathbf{A} is often ill-conditioned. That means some of its M singular values are close to zero. It is common to improve the conditional number of \mathbf{A} by employing the truncated SVD regularization. We use $\mathbf{A}^{(m)}$, a rank- m approximation of \mathbf{A} (Kaipio and Somersalo, 2004). In our experience, the inverse solutions obtained using \mathbf{A} and $\mathbf{A}^{(m)}$ are almost identical, which reflects the robustness of our solver. Working with $\mathbf{A}^{(m)}$ further reduces the number of variables in the optimization problem by reducing M to m and significantly accelerates computations. Therefore, all the results reported in this article are based on $\mathbf{A}^{(m)}$ with $m=100$. On the other hand, to obtain realistic simulated data, the forward calculations of the simulated signals are based on the full matrix \mathbf{A} .

Pre-processing for temporal basis function construction

Due to different types of sensors, gradiometers and magnetometers in MEG and electrodes in EEG, the measurements have different units and different ranges of recordings. To construct a set of temporal basis functions, we must first whiten the measurements in the sensor space according to the estimated noise covariance matrix. Without this whitening procedure, some subsets of the sensor recordings, such as the magnetometers, would have been ignored in the construction of basis functions. In addition, we need to exclude eventual stimulus artifacts when computing the SVD of \mathbf{Y} ; otherwise, most of the basis functions in \mathbf{V}_Y would mainly explain the artifacts. For example, in our analysis of median-nerve experiments, measurements from the first 5 ms after the stimulus onset are excluded in the basis function construction.

Multi-resolution approach

In this work, the estimates of the source locations are confined to a mesh. In order to reduce computational complexity, we employ a multi-resolution scheme. We first perform source estimation on a coarse mesh, then we adaptively refine the mesh around the activation regions. In other words, the forward model at a high resolution level includes all the vertices at one level below and the newly introduced vertices around the detected regions.

The ℓ_1 -norm regularization often produces focal estimates, which is more appropriate to model activations in the sensory regions. However, for a spatially extended source, the corresponding source estimates may appear as several activated vertices in the extended patch if the estimation is conducted on a much finer mesh (Uutela et al., 1999). In most of our experiments, we used a double resolution scheme, 20- and 10-mm spacing between vertices. Our reconstruction results for the median-nerve experiments show that the activations in the primary sensory cortex can be accurately represented using this multi-resolution scheme. For auditory experiments, where the sources are slightly more diffuse, our solver detects several adjacent vertices in the auditory areas.

Computation requirements

Compared with the MCE, which solves T linear programming problems with N variables each, our solver performs a single SOCP optimization over NK variables. As described in Eq. (8) and Appendix A, we increase the number of variables to approximately $(N+M)K$ in order to convert the ℓ_1/ℓ_2 -norm solver into the SOCP formulation. In this work, we use the Self-Dual-Minimization software package (SeDuMi) (Strum, 1999) that implements the primal-dual interior-point method with logarithm barrier functions, to solve for the SOCP problem. The primal-dual interior-point method, employed in SeDuMi, has run time $\Theta(((N+M)K)^3)$ per iteration. It converges within thirty iterations in most of our experiments. For $M=100$ (with the truncated SVD-regularized lead-field matrix), $K=3$, and $N\approx 500$, our current implementation takes about 10 s with a standard PC (2.8 GHz CPU and 8 GB RAM) to compute the inverse solution. When $N\approx 2000$, it takes about 100 s. Combining with a two-level multi-resolution scheme, with $N\approx 500$ for the first level and $N\approx 700$ for the second level, it takes 25 s. When K increases to 6, the multi-resolution scheme takes about 100 s.

Results

Due to a lack of ground truth in human EEG/MEG experiments, we first study the behavior of the method and its sensitivity to parameter settings and to noise using simulated data. We then compare the method to standard inverse solvers using human MEG data from somatosensory and auditory studies.

Simulation studies

To simulate MEG measurements, we created active vertices A, B, and C (Fig. 1, top) on the cortical sheet at source spacing of 20 mm, with current source orientation along the normal to the cortical surface. In all experiments in this paper, we scaled the reconstructed surfaces to 10^5 -mm² surface area per hemisphere. Vertex A is located in the lateral frontal region, Vertex B is located at the pre-central gyrus, and Vertex C is located at the Sylvian fissure. The time courses of these three vertices are shown in Figs. 2a–c (black solid curves). We chose the signals to have similar temporal characteristics to those of the auditory evoked responses, but with temporal translation and scaling. The source signals of vertices B and C show activation during overlapping time intervals, which makes the inverse problem difficult.

For the forward calculations, we employed the sensor configuration of the 306-channel Neuromag VectorView MEG system (204 gradiometers and 102 magnetometers) used in our experimental studies. The location of the array with respect to the head and the noise covariance matrix were obtained from human MEG experiments. A single-compartment homogeneous forward model was employed. With Gaussian noise added, the resulting signals have an SNR=3 dB, where the SNR is defined as $10\log_{10}\frac{\|AS\|_F^2}{MT\sigma^2}$, where σ^2 is the noise variance. The resolution of the source space is relatively coarse; nevertheless, this example serves as a good illustration for the method.

In the inverse estimation, we fixed the orientation of the estimated currents to be perpendicular to the cortical mesh. Fig. 1b depicts the inverse solutions at three time frames obtained from the ℓ_1/ℓ_2 -norm solver using three basis functions and $\lambda=10^9$. The parameter values were selected based on our validation experiments presented later in

this section. Curves marked with ‘o’ in Figs. 2a–c correspond to the source signals estimated by the method. The resulting spatial maps and source time courses match well with the ground truth.

ℓ_1/ℓ_2 -norm vs. ℓ_1 -per-coefficient

To further explore the behavior of the ℓ_1/ℓ_2 -norm regularizer, we compared its reconstructions with those obtained by applying the ℓ_1 -norm regularizer to the coefficients of each basis function separately. We will refer to this solver as ℓ_1 -per-coefficient. This comparison reveals the effect of the ℓ_2 -norm regularization for all coefficients. For each basis function, ℓ_1 -per-coefficient computes the least ℓ_1 -norm solution for the coefficients independent of other basis functions. It also achieves stable reconstruction due to the use of temporal basis functions. We applied ℓ_1 -per-coefficient, also employing the interior-point method implemented in Strum (1999) to the data described above. Fig. 1c depicts the reconstruction results, and the corresponding time courses of the three active vertices are presented in Figs. 2a–c (marked with ‘+’). Both ℓ_1/ℓ_2 -norm and ℓ_1 -per-coefficient detect activations in the three vertices (Fig. 1), but the ℓ_1/ℓ_2 -norm solution contains fewer false positive activations than that obtained through the ℓ_1 -per-coefficient method. Moreover, the reconstructed time courses of the ℓ_1/ℓ_2 -norm solution match the ground truth time courses slightly better than those of the ℓ_1 -per-coefficient solution.

The three selected basis functions are illustrated in Fig. 2d, and the reconstructed coefficients from these two algorithms are shown in Figs. 2(g,h). The projection coefficients of the simulated signals are also presented for comparison purpose (Fig. 2f). Although the spatial pattern of the projection coefficients is similar for the two methods, the spatial pattern for all three coefficients is more sparse in the ℓ_1/ℓ_2 -norm solution (Fig. 2g) than in the ℓ_1 -per-coefficient solution (Fig. 2h).




Fig. 1. Activation maps at different time frames. (a) Ground truth activation maps at peak response time for three sources. (b) The spatial maps estimated using the ℓ_1/ℓ_2 -norm solver. (c) The spatial maps estimated using the ℓ_1 -per-coefficient solver. The color codes in (a) do not indicate current directions. Hot/cold colors in (b, c) correspond to outward/inward current flow. The most active areas in the solutions are highlighted, and their estimated time courses are shown in Figs. 2a–c.

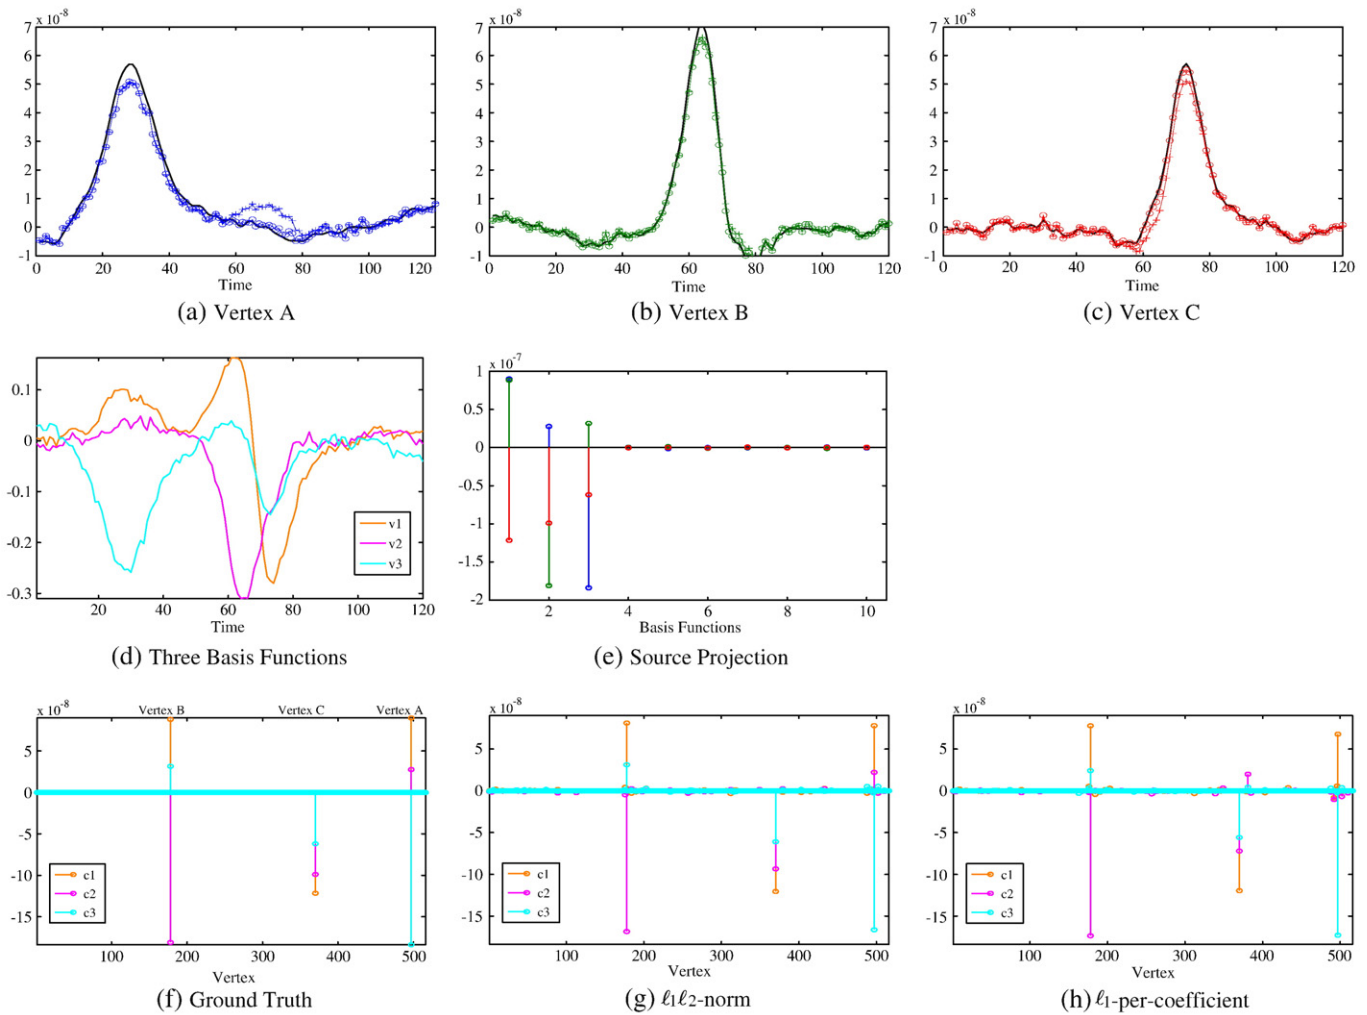


Fig. 2. Reconstructed source signals, the three basis functions, and the reconstructed coefficients \tilde{S} . The top row illustrates the simulated (black solid curves) and the reconstructed source time courses (ℓ_2 -norm marked as 'o' and ℓ_1 -per-coefficient marked as '+') for the three active vertices; the middle row presents the selected basis functions (d) and the projection coefficients of the simulated source time courses onto the top ten basis functions (e). (f) Projection coefficients of the simulated data onto the three basis functions, corresponding to the top three singular components. (g) Reconstructed coefficients from $\ell_1\ell_2$ -norm. (h) Reconstructed coefficients from ℓ_1 -per-coefficient.

Since ℓ_1 -per-coefficient models the coefficients of each basis function separately, vertex with large coefficient for one basis function may have zero coefficient for another basis function. On the other hand, $\ell_1\ell_2$ -norm considers all coefficients jointly in sparsity determination. This method is particularly helpful for basis functions which have a smaller SNR, such as \vec{v}_2 and \vec{v}_3 . That is illustrated by a more sparse distribution of the coefficients in Fig. 2g than that in Fig. 2h, a missing \vec{v}_2 component for Vertex A (Fig. 2h), and a false detection for a vertex close to Vertex B (Fig. 2h). The ℓ_2 -norm regularizer essentially helps

bundle basis functions \vec{v}_2 and \vec{v}_3 with those that are aligned well with the signal subspace, such as \vec{v}_1 , to jointly determine an activation map. Therefore, we can see that sparsity defined by all coefficients is more suitable for the current basis construction method in conjunction with complex neural signals.

Fig. 2e presents the projection coefficients of the simulated source signals onto the temporal basis functions. The coefficients that correspond to basis functions $\{\vec{v}_k : k \geq 4\}$ are close to zero. We only displayed the coefficients corresponding to the first ten basis




Fig. 3. Reconstructed source signals from MNE, MCE, and VESTAL. Each panel illustrates the simulated (black solid curves, same as Figs. 2a–c) and the reconstructed source time courses: MNE (solid), MCE (dashed), and VESTAL (marked as 'Δ').

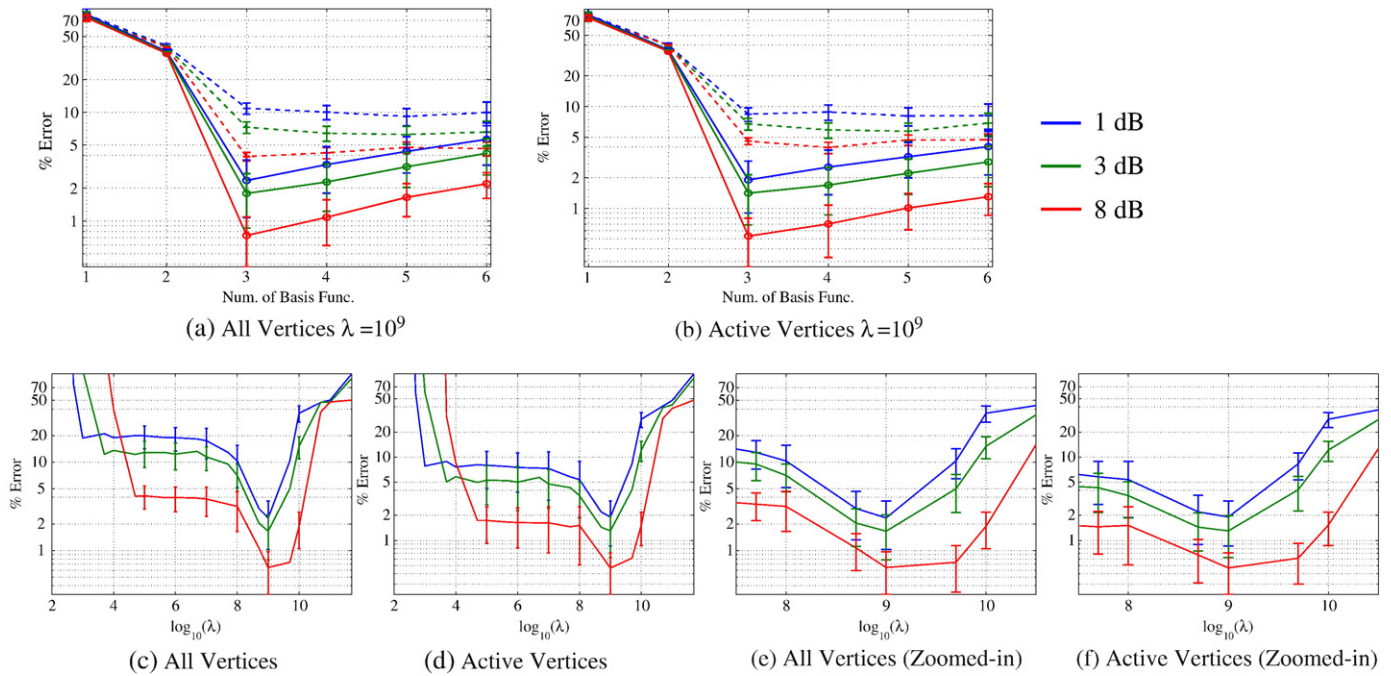


Fig. 4. Relative MSE. Subpanels (a, b) present the percentage relative MSE of the reconstruction results using the ℓ_1/ℓ_2 -norm (solid curves) and the ℓ_1 -per-coefficient (dashed curves) for all vertices and the three active vertices under three different SNR settings. Subpanels (c, d) present the relative MSE of the reconstruction results from ℓ_1/ℓ_2 -norm for all vertices and for the three active vertices under three SNR settings as the regularization strength, λ , varies between 10^2 and 10^{12} . Subpanels e, f are the corresponding zoomed-in versions. Note that the error bars close to the bottom of the figures appear large due to the logarithmic scale.

functions. For this simple example, Fig. 2e verifies that the selected basis well approximates the signal subspace of the simulated signals.

Comparison with MNE, MCE, and VESTAL

We also compared the proposed method with the standard MNE, MCE, and VESTAL (Fig. 3). The estimates from the standard MNE are smaller than the simulated signals, and it is caused by the diffuse property of the ℓ_2 -norm regularization. The estimated time courses from MCE exhibit substantially “spiky” discontinuities due to the solver’s sensitivity to noise. Projecting MCE results to a set of basis functions, VESTAL removes the discontinuities; however, the amplitude of the estimated time courses is smaller than the true activation signals since the two-step procedure cannot fully compensate for the errors in the original MCE solutions. Therefore, in the rest of the stimulated experiments, we focus on the performance of ℓ_1/ℓ_2 -norm and ℓ_1 -per-coefficient.

Sensitivity to noise and basis selection

To examine the sensitivity of the methods to noise and basis selection, we computed inverse solutions for 100 independently generated data sets for each noise setting (varying from SNR 1 dB to 8 dB) and basis selection cutoff (K varying from 1 to 6). The relative mean squares error (MSE)¹ for the three active vertices and for all vertices of the ℓ_1 -per-coefficient and the ℓ_1/ℓ_2 -norm inverse solutions are shown in Figs. 4(a, b).

The ℓ_1/ℓ_2 -norm outperforms ℓ_1 -per-coefficient under all SNR settings and basis selection cutoffs we examined. The improvement of the relative MSE varies from 4% to 10%, with larger improvement for noisier data. The large improvement in the low SNR cases again demonstrates the importance of the ℓ_2 -norm regularization on the coefficients of the representation. The standard deviation of the reconstructions estimated over the 100 simulated data sets is similar

for the two solvers. It varies between 0.3% and 2.5% for $K \geq 2$ and all three selected SNR settings. For $K=1$, the standard deviation is between 5% and 10% due to more variability in representing the signals using a single temporal basis function.

In general, both solvers achieve the best performance for $K=3$ basis functions. If the chosen number of basis functions is too high, some basis functions represent noise, resulting in slight degradation of the result quality as reflected by the gentle slope on the logarithmic scale. On the other hand, the ℓ_1 -per-coefficient’s performance is not affected by including too many basis functions, because its estimated sources from the noisy basis functions are usually small. Including too few basis functions leads to a significant loss of signals; both solvers fail to recover the missing signals.

Sensitivity to regularization strength

We also investigated the methods’ sensitivity to the value of the regularization parameter λ . Large λ corresponds to a high penalty on the strength of the current sources, in terms of the ℓ_1/ℓ_2 -norm; small λ emphasizes the data fidelity term. Due to whitening, the first term in Eq. (5) is on the order of MK , where M is on the order of 10^2 . For an activated vertex in our experiments, \tilde{s}_{nk} is on the order of 10^{-8} . Hence, $|\tilde{S}|_1^2$ is approximately $10^{-7} K$. Therefore, $\lambda=10^9$ roughly balances between the data fidelity and the regularization terms in Eq. (5). In the experiments using human MEG data, we set $\lambda=10^9$. Since the values in the data fidelity and regularization terms are both linearly proportional to K , the regularization strength should be independent of the number of basis functions participating in the inverse calculation. That means the sensitivity of ℓ_1 -per-coefficient and ℓ_1/ℓ_2 -norm to λ should be the same. Hence, we only present the relative MSE obtained using ℓ_1/ℓ_2 -norm for all vertices and for the three active vertices for various values of λ (Figs. 4c–f). As we can see, λ around 10^9 provides accurate reconstruction results. The regularization shows no effect for $\lambda < 10^3$; when $\lambda > 10^{10}$, the data fidelity term is effectively ignored in the optimization process. For $\lambda=10^9$, the standard deviation of the MSE estimated from the 100 data sets is less than 1%.

¹ We define the relative MSE as $\frac{\|reconstructed\ signals - ground\ truth\ signals\|_2^2}{\|ground\ truth\ signals\|_2^2}$.

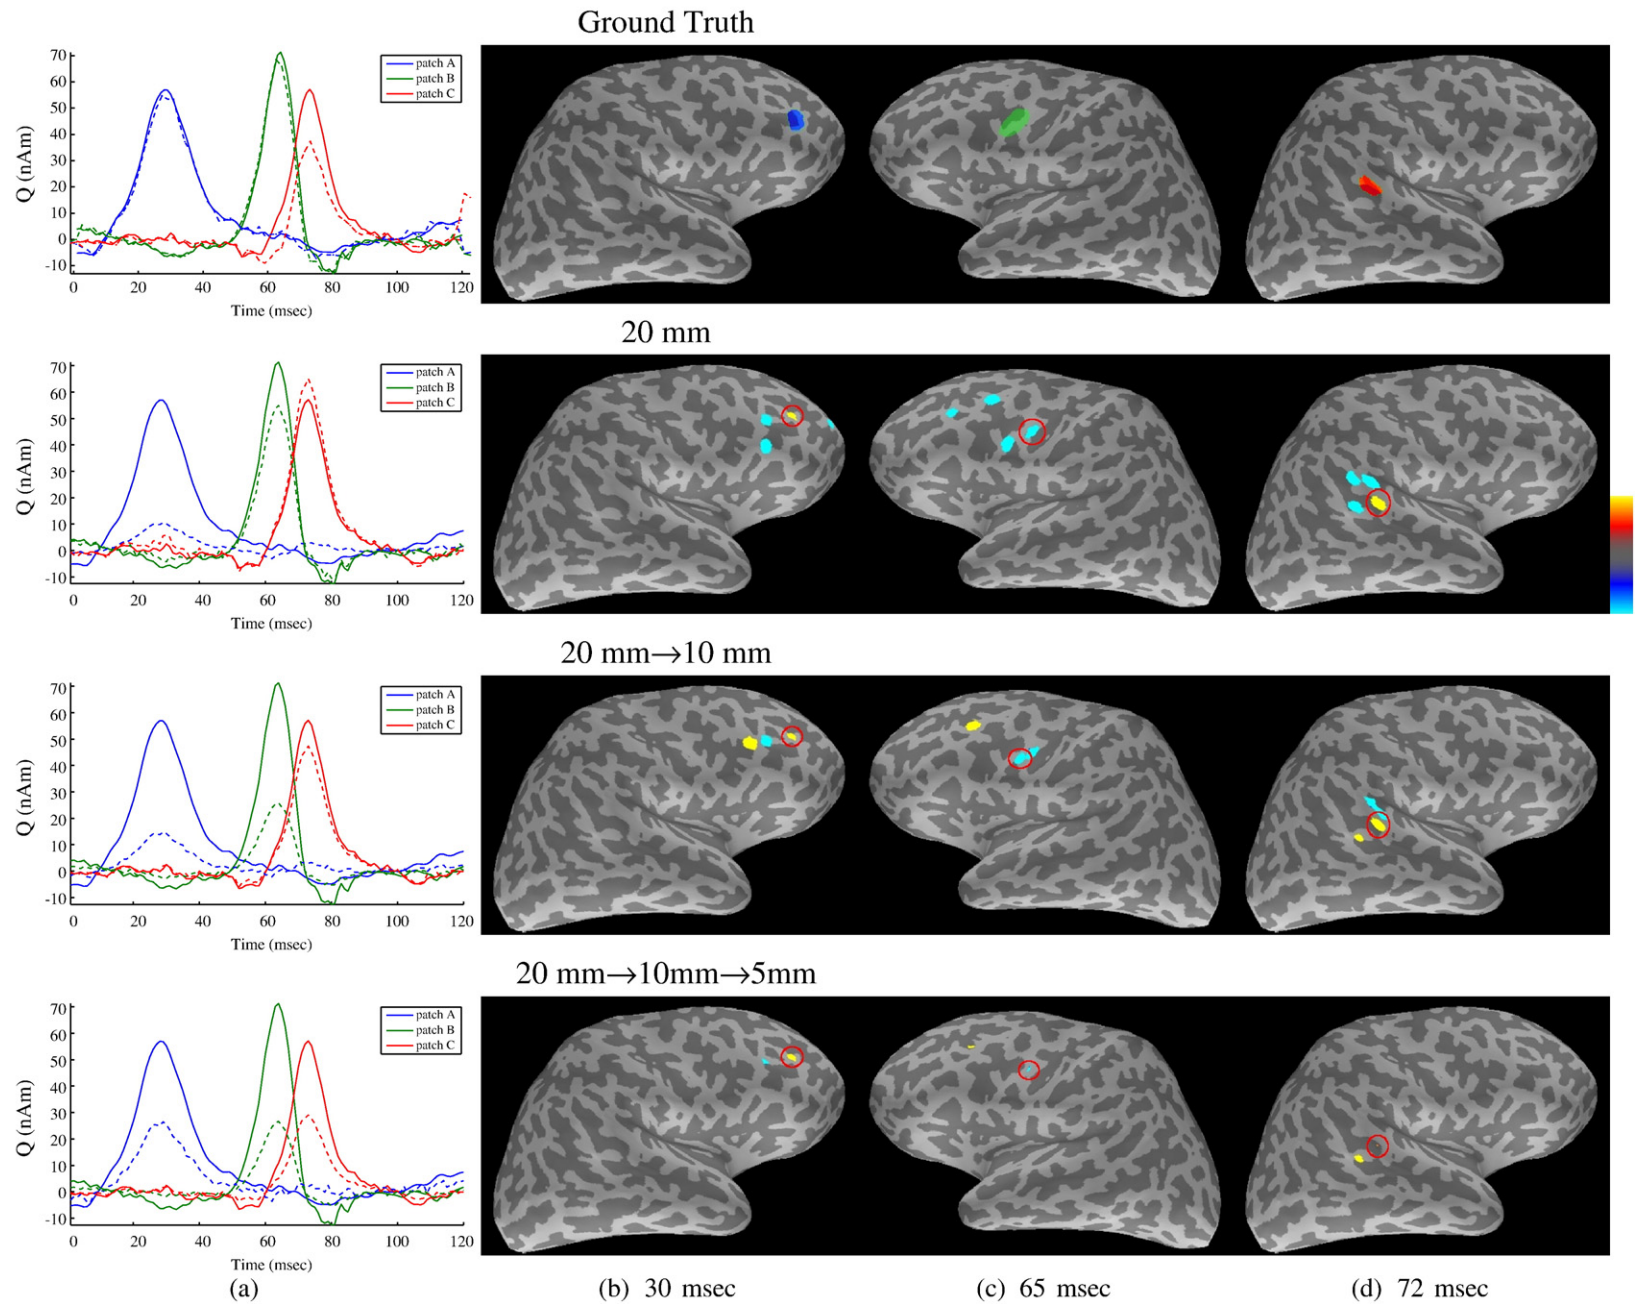


Fig. 5. Reconstructions obtained using the ℓ_1/ℓ_2 -norm solver with different multi-resolution schemes. The top row presents the three simulated activation patches (color is used to label patches and does not indicate current directions). The detected results (row two to four) are shown in hot or cold colors corresponding to current flowing outward and inward, respectively. The time courses of the highlighted areas are shown in column (a). Solid curves in (a) are the simulated time courses. The dashed curves in row one are the dipole fitting results. The dashed curves in row two to four are the reconstructed time courses from ℓ_1/ℓ_2 -norm, which are scaled by a factor of four for illustration purposes.




Fig. 6. Six selected basis functions. The basis functions were obtained from SVD of the MEG measurements between 6 ms and 200 ms after stimulus onset.

Different spatial resolutions

To further examine the ℓ_1/ℓ_2 -norm inverse solutions at different spatial resolutions, we extended the simulated sources at vertices A, B, and C, described in the last experiment, to three patches at 0.65 mm tessellation resolution. The patches have spatial extent of approximately 15-mm in diameter (200 to 500 vertices at a 0.65 mm resolution), indicated by the colored patches in the first row of Fig. 5.

The ground truth source signals are identical to those employed before, shown as solid curves in Fig. 5a. To generate the sensor signals, we added Gaussian noise with covariance matrix estimated from the pre-stimulus recordings of a human MEG data set, with a resulting SNR=3 dB. To avoid an “inverse crime,” the inverse solutions were calculated at the resolution lower than the resolution used in the simulations, including a single-level mesh at 20 mm, a two-level multi-resolution scheme at 20 and 10 mm, and a three-level multi-resolution scheme at 20, 10, and 5 mm. We set $\lambda = 10^9$ and $K = 3$ in this experiment.

Fig. 5 shows the inverse solution from each of the three multi-resolution schemes at 30, 65, and 72 ms, corresponding to the peaks of the three simulated source signals. The detected areas are either in blue-cyan or red-yellow corresponding to current flowing inward or outward with respect to the cortex. Each reconstruction was thresholded such that all three areas were detected at their peak times. Good performance is indicated by fewer false positives. A smaller amplitude in the dipole fitted time course for patch C (Fig. 5a row one) indicates that some of the vertices in this patch have current orientations silent with respect to the MEG sensors. Smaller amplitudes in the reconstructed time courses from our solver are expected due to the magnetically silent sources, as well as the distributed nature of the model where some nearby vertices are detected despite the regularization promoting spatial sparsity. We scaled all the reconstructed time courses by a factor of four for illustration purposes.




Fig. 7. Reconstruction maps. Top panel: significance statistics of the ℓ_1/ℓ_2 -norm solver and MNE for the median-nerve experiment. Hot/cold color corresponds to outward/inward current flow. The most active areas in the ℓ_1/ℓ_2 -norm solutions are highlined, and their estimated time courses are shown in Fig. 9. Bottom panel: dipole fitting results with three sources.




Fig. 8. Coronal slices for the detected iSII activations from l_1l_2 -norm and from dipole fitting. Subfigures (c, d) are the corresponding zoomed-in versions.

We observe that 20 mm resolution is too coarse for reconstruction, as reflected by more ambiguity in the source locations. The reconstruction results at the 5 mm resolution are too focal. At this reconstruction resolution, the vertices in a simulated patch are close, and some of them have similar orientation. Mathematically speaking, the signal distributions corresponding to those vertices, indicated by the column vectors in the lead-field matrix, are almost linearly dependent. The l_1 -norm encourages spatial sparsity, and it usually allocates all source current to one of those vertices. Reconstruction at the 10 mm resolution achieves the most accurate results: fewer false positives and a better representation of the spatial extent of the simulated patches. Therefore, we employed the two-level multi-resolution scheme in the experiments using human MEG data.

Human MEG data

Experiments with synthetic data reveal the potential of the l_1l_2 -norm solver to provide accurate and stable solutions when handling focal and correlated sources, even in a noisy environment. Next, we compare the performance of the solver to the MNE and dipole fitting using two human MEG data sets from median-nerve and auditory experiments. Both experiments were acquired using a 306-channel Neuromag VectorView system. The anatomical images were collected with a Siemens Avanto 1.5-T scanner with a T1-weighted sagittal MPRAGE protocol, which were employed for cortical surface reconstruction (Dale et al., 1999; Fischl et al., 1999). A multi-echo 3D Flash acquisition was performed to extract the inner skull surface for the boundary-element model (Hämäläinen and Sarvas, 1989; Oostendorp and Van Oosterom, 1989; Hämäläinen). Informed consent in accordance with the Massachusetts General Hospital ethical committee was obtained from subjects prior to participation.


Median-nerve experiments

We present results for one subject, a 40-year old male, in the study. The median nerve was stimulated at the left wrist according to an event-related protocol, with a random inter-stimulus-interval ranging from 1.5 to 2 s. Data were acquired at sampling rate of 2 kHz; a 200-msec baseline before the stimulus was used to estimate the noise covariance matrix. Approximately 300 trials remained after rejecting trials with eye-movements and other artifacts², from which we computed the average signal used as the input to the inverse solvers. We first applied baseline correction and whitened the data spatially based on the pre-stimulus measurements. For the l_1l_2 -norm solver, we used six basis functions shown in Fig. 6, and they explain 80% variance of the data. SVD was performed on signals between 6 ms and 200 ms after stimulus onset to avoid post-stimulus artifacts.


It has been shown that the median-nerve stimulus activates a complex cortical network (Hari and Forss, 1999). The first activation of

the contralateral primary somatosensory cortex (cSI) peaks around 20 ms and continues over 100 ms; then the secondary somatosensory cortex (SII) activates bilaterally around 70 ms and lasts up to 200 ms. Whether SI and SII form a sequential or parallel architecture is still a topic of active debate (Kass et al., 1979; Rowe et al., 1996). The posterior parietal cortex (PPC), located on the wall of the post-central sulcus, medial and posterior to the SI cortex hand area, activates around 70–110 ms. This area, also known as the parietal association area, most probably functions as an integrator between sensory and motor processing. Although the SI-SII network exhibits robust activation, there is significant variation from subject to subject especially in the time courses of SII activations.

In this experiment, we controlled FDR at 0.05, computed from 5000 permutations as described in Statistical significance test section. We also compared our results with the MNE computed using a standard software package (Hämäläinen). In practice, experts often interpret MNE through its statistics, dSPM, with manually adjusted thresholds. For the purpose of comparison, we selected the threshold for dSPM so that all four regions of interest, cSI, cSII, iSII, and PPC, were



(a) l_1l_2 -norm



(b) dipole fitting

Fig. 9. Reconstructed time courses obtained from the l_1l_2 -norm solver and dipole fitting. The corresponding activation maps are reported in Fig. 7.

² Trials with peak-to-peak amplitude of the EOG signals exceeding 150 μ V, gradiometer signals exceeding 3000 fT/cm, or magnetometer signals exceeding 3.5 pT were rejected. These rejection criteria are the same for the auditory experiment.

included. Fig. 7 presents the activation maps obtained using ℓ_1/ℓ_2 -norm and MNE. At 20 ms, ℓ_1/ℓ_2 -norm pinpoints cSI on the postal wall of the central sulcus. MNE produces a more diffuse solution leading to false positives in the post-central sulcus. The ℓ_1/ℓ_2 -norm clearly demonstrates change of polarity in cSI, reflected by the change of current direction between 20 and 35 ms. The polarities estimated using the ℓ_1/ℓ_2 -norm solver agree with the literature (Wikström et al., 1996): outwards intra-cellular current at 20 ms associated with N20 and inwards intra-cellular current at 35 ms associated with P35. At 75 ms, both MNE and ℓ_1/ℓ_2 -norm capture signals from cSII. ℓ_1/ℓ_2 -norm successfully localizes PPC at the post-central sulcus, but the location of PPC is ambiguous in the MNE results. According to the findings reported in Forss et al. (1994) and Hari and Forss (1999), the signals from iSII is weaker than those from cSII. By controlling FDR at 0.05, the ℓ_1/ℓ_2 -norm solver detects iSII activation at 85 ms, but places it at the superior temporal lobe instead of the inferior parietal lobe. As shown in the volumetric display (Fig. 8), these two regions are very close, making the inverse problem challenging. MNE also presents weak iSII signals; the location is ambiguously spread between the iSII region and the superior temporal lobe.

We estimated the current source dipoles and their corresponding time courses through the standard dipole fitting procedure (Nelder and Mead, 1965; *xfit* software). Dipole fitting was performed using the corresponding channels at 20–38 ms, 75 ms, and 85 ms after the stimulus onset. The source estimates are summarized in one map as illustrated in Fig. 7(bottom), and the corresponding time courses are presented in Fig. 9(b). Dipole fitting did not correctly localize PPC from these data because PPC is very close to cSI. The locations for cSI and cSII identified by our solver match with the dipole fitting results. The correct localization of iSII using dipole fitting required manual intervention in selecting appropriate channels in contrast to the automatic ℓ_1/ℓ_2 -norm solver. The highly folded cortical pattern along the Sylvian fissure presents a significant challenge for most inverse solvers, including both distributed and discrete parametrization approaches (Fig. 8). One way to resolve this problem is to utilize measurements from other modalities, such as fMRI, to further constrain the solution (Liu et al., 1998). We defer such extensions to a future study.

Fig. 9 shows the time courses of the activated regions detected by our solver. The general shape of these time courses agrees with the neuroscience literature (Forss et al., 1994; Hari and Forss, 1999). Our method yields stable time courses that capture the main deflections precisely. The first deflection in cSI occurs at 20 ms. cSI soon changes its polarity and reaches its maximum at 35 ms. Although cSII has stronger signal than iSII, they have similar temporal signature: onset at 60 ms and peak at 82–85 ms. The time courses are quite similar to those estimated through dipole fitting (Fig. 9(b)), except for the cSI activation between 70 and 150 ms. This is most likely because the PPC activation was missed by the dipole fitting and its time course incorporated into the estimate of cSI. Even though the magnitude of the time courses obtained from the two solvers are not directly comparable, this comparison demonstrates the ability of the ℓ_1/ℓ_2 -norm regularization to achieve high-quality reconstructions of source signals. Furthermore, ℓ_1/ℓ_2 -norm does not restrict itself to a fixed number of dipole sources.

Auditory experiments

In the auditory experiments, 500 Hz tone bursts were presented to either the right ear or the left ear of the subject according to an event-related paradigm, with a random inter-stimulus-interval between 1.2 and 1.5 s. Temporal sampling rate of these MEG data was 1.25 KHz. As before, a 200 ms baseline period was used for noise estimation.

After standard pre-processing, described in the last section, we applied ℓ_1/ℓ_2 -norm, MNE, and dipole fitting to the average data. Fig. 10 illustrates one frame of the reconstructed signals, at 110 ms after the stimulus onset. The statistics and the thresholds for ℓ_1/ℓ_2 -norm were computed using the same permutation method as before. Both the ℓ_1/ℓ_2 -norm and the dSPM detected auditory

activations in both hemispheres. Due to close distance between the inferior parietal and the superior temporal regions, all three solvers have false positives in the parietal lobe. The false positives are weaker in the ℓ_1/ℓ_2 -norm solutions than in the MNE solutions. Given the MNE results, it is more ambiguous whether the sources originate from the auditory region or from the SII region. We also examined the polarities of the estimated sources, and they all agree with the findings reported in the literature (Tuomisto et al., 1983).

The corresponding estimated source signals from the ℓ_1/ℓ_2 -norm solver and dipole fitting are depicted in Fig. 11. Both methods detected that the early auditory response occurs around 60 ms and that the contralateral auditory region activates slightly stronger than the ipsilateral region. Compared with the ℓ_1/ℓ_2 -norm solution, the dipole fitting solution captures the temporal details slightly more accurately, as reflected by the 6–8 ms difference between the contralateral and the ipsilateral activations. Nevertheless, the ℓ_1/ℓ_2 -norm solver is more flexible than dipole fitting in capturing the spatial extent of the activation regions.

Discussion

The proposed inverse solver utilizes ℓ_1 -norm regularization to capture spatial sparsity of the activations and ℓ_2 -norm regularization on the projection coefficients in the signal subspace to model the time-varying frequency content in the activation signals. While considering all vertices in the brain as candidate activation sources, our solver can still obtain focal activation maps and capture activation signals with precise deflection signatures. The ℓ_1/ℓ_2 -norm inverse solutions share some similar characteristics with dipole fitting results; however, the number of dipole sources is not required to be known *a priori* for ℓ_1/ℓ_2 -norm. As demonstrated in the simulations, the performance of the ℓ_1/ℓ_2 -norm solver is robust to the chosen number of basis functions. This feature makes the method particularly suitable for neuroscience applications where the number of dipole sources is usually unknown.

We model the activation signals as linear combinations of multiple temporal basis functions. There are various approaches to obtain the basis functions such as the Fourier and wavelet decompositions. If the Fourier decomposition is employed, the selected basis functions must capture the frequency components of the neural signals. If wavelets are used, we need to choose a wavelet family appropriately. If the temporal structure of the source signals at a particular region were known, we would incorporate it as one of the temporal basis functions. In this case, the assumption of the linear combinations of multiple basis functions would need to be modified. Furthermore, we chose to work with orthonormal basis functions. If the basis functions are not orthonormal, the general idea of this paper is still valid, but the ℓ_2 -norm would have to be replaced with the Mahalanobis distance.

Obtaining a compact representation of the signals can significantly reduce the computational requirements. Because of the time-varying frequency content and substantial variability in the signals across activation regions, subjects, and tasks, data-independent basis sets such as Fourier and wavelets may not be the best choice to compactly represent the signals. In this work, we constructed the temporal basis functions through the SVD decomposition of the data. We chose a set of basis functions that correspond to the largest singular values. The cutoff was determined by our knowledge of the source signals and the temporal structure of the singular vectors. Fixing the cutoff may lead to a loss of signals that lie in the orthogonal subspace spanned by \mathbf{V}^\perp . A possible improvement is to alternate between modifying the basis functions and performing reconstruction.

Accurate estimation of the spatial extent of the sources is one of the main challenges for any source modeling approach. Compared with dipole fitting, the ℓ_1/ℓ_2 -norm solver demonstrated a better ability to capture the auditory activations as shown in the Auditory experiments section. To further improve the ℓ_1/ℓ_2 -norm solver

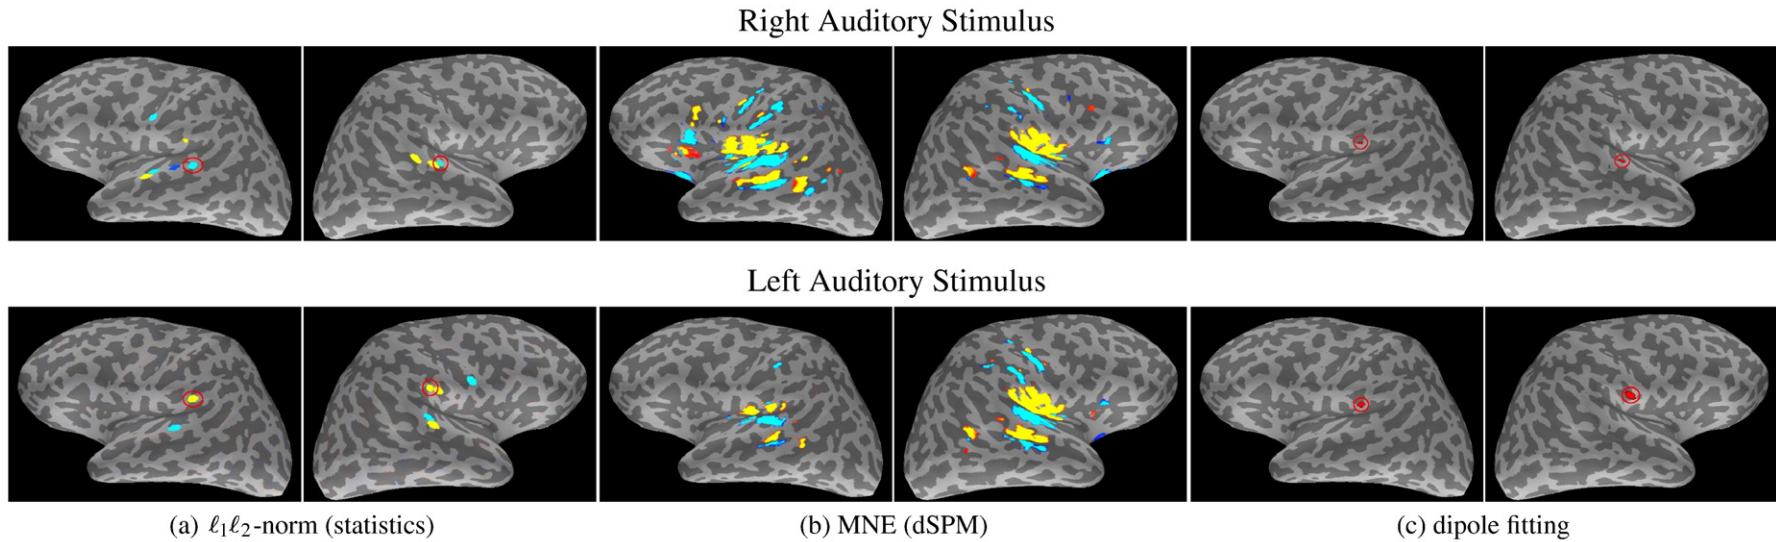


Fig. 10. Significance statistics of the l_1l_2 -norm solver and MNE, and dipole fitting at 110 ms after right (top) and left (bottom) auditory stimulus onset. Hot/cold color corresponds to outward/inward current flow. The most active areas are highlighted, and their estimated time courses are shown in Fig. 11.

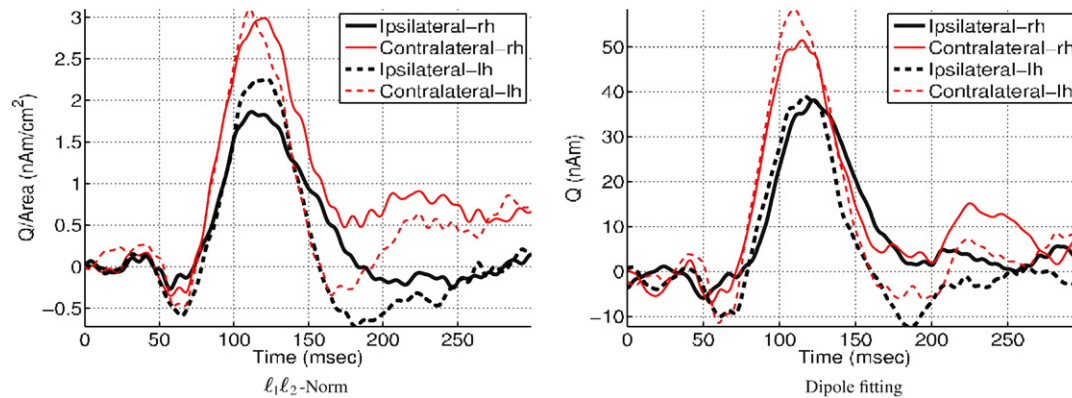


Fig. 11. Reconstructed time courses from l_1l_2 -norm and dipole fitting.

performance for extended sources, we can apply the ℓ_1 -norm regularizer to model the difference among neighboring vertices, rather than the vertices directly. Thus, the new model would favor a piece-wise constant activation pattern. This idea combined with an MNE estimator is at the core of the LORETA reconstruction method (Pascual-Marqui et al., 1994). The work of Auranen et al. (2007) proposed a Bayesian approach in which the measurements and the hyper-prior determine the spatial extent of the activations through estimating the joint posterior distribution of the inverse solution and the exponent in the regularizer. Alternatively, the activation pattern could be expressed using a set of spatial basis functions (Limpiti et al., 2006) or the current multipolar expansions (Cottareau et al., 2007).

Conclusions

The proposed inverse solver takes advantage of the relatively smooth nature of the underlying EEG/MEG source signals through performing inverse operation for all temporal samples simultaneously. To overcome the overly diffuse inverse solutions, the ℓ_1/ℓ_2 -norm captures spatial sparsity through ℓ_1 -norm regularization. It also applies an ℓ_2 -norm regularizer to the projection coefficients of the temporal basis functions spanning the signal subspace. Performing reconstruction in the signal subspace while jointly considering the coefficients for all selected basis functions leads to stable estimates with a smaller number of false positives as confirmed by our experiments using simulated and human MEG data. The ℓ_1/ℓ_2 -norm solver is formulated as an SOCP problem. Its fast optimization enables us to perform a statistical significance testing for the ℓ_1/ℓ_2 -norm inverse solutions via a permutation method. The ℓ_1/ℓ_2 -norm can be straightforwardly applied with and without orientation constraints. Its flexible formulation should also allow incorporation of fMRI information as a constraint.

Acknowledgments

We thank Dr. Dmitry Malioutov and Dr. Xin Huang for the interesting discussion on SOCP. We also thank Biswajit Bose, Hesheng Liu, and Thomas Yeo for helpful comments on this work, and Dr. Van Leemput for proof reading this draft. This work was supported in part by NIH NIMH U54-EB005149, NIH NCRN mBIRN U24-RR021382, NIH NINDS R01-NS051826, NIH BRP R01-E006385, NIH NCRN NAC P41-RR13218 and NIH NCRN CFNT P41-RR14075 grants, by the NSF CAREER Award 0642971, and by U.S. DOE Award Number DE-FG02-99ER62764 to the MIND Institute. Wanmei Ou is partially supported by an NSF graduate fellowship.

Appendix A

A quadratic inequality constraint can be revised into the canonical form of a second-order cone. We start with the standard form of a quadratic inequality constraint:

$$\|\mathbf{y}-\mathbf{A}\mathbf{s}\|_{\ell_2}^2 \leq w \quad (16)$$

With straightforward derivations including expanding the ℓ_2 -norm and completing squares, we can show that Eq. (16) is equivalent to

$$\mathbf{s}^T \mathbf{A}^T \mathbf{A} \mathbf{s} + \left(\frac{1-2\mathbf{y}^T \mathbf{A} \mathbf{s} + \mathbf{y}^T \mathbf{y} - w}{2} \right)^2 \leq \left(\frac{1+2\mathbf{y}^T \mathbf{A} \mathbf{s} - \mathbf{y}^T \mathbf{y} + w}{2} \right)^2 \quad (17)$$

Eq. (16) also implies $1+2\mathbf{y}^T \mathbf{A} \mathbf{s} - \mathbf{y}^T \mathbf{y} + w \geq 0$. Therefore, setting $x_0 = \frac{1+2\mathbf{y}^T \mathbf{A} \mathbf{s} - \mathbf{y}^T \mathbf{y} + w}{2}$ and $\bar{\mathbf{x}} = \begin{bmatrix} \mathbf{A} \mathbf{s} \\ \frac{1-2\mathbf{y}^T \mathbf{A} \mathbf{s} + \mathbf{y}^T \mathbf{y} - w}{2} \end{bmatrix}$, we arrive at the equivalent second-order cone constraint $x_0 \geq \|\bar{\mathbf{x}}\|_{\ell_2}$. As we can see that the conversion introduces a set of new variables \mathbf{x} , of size $M+2$.

Appendix B

Second-order cone programming (SOCP) problems are defined by: (1) a linear objective function, (2) a feasible set that is an intersection of an affine linear manifold with the Cartesian product of second-order cones. Since the linear objective function and the feasible set are convex, SOCP problems can be solved by convex optimization techniques. The canonical primal form of SOCP is as follows:

$$\min \mathbf{c}^T \mathbf{x} \quad (18)$$

$$\text{s.t. } \mathbf{A}\mathbf{x} = \mathbf{b} \quad (19)$$

$$\mathbf{x} \in \mathcal{Q} \quad (20)$$

where $\mathcal{Q} = \{\mathbf{x} : \mathbf{x}_0 \geq \|\bar{\mathbf{x}}\|_{\ell_2}\}$. \mathcal{Q} is also referred to as the Lorentz cone; it is self-dual. The dual cone \mathcal{Q}^C is defined as

$$\mathcal{Q}^C := \{\mathbf{y} : \forall \mathbf{x} \in \mathcal{Q}, \mathbf{y}^T \mathbf{x} \geq 0\} \quad (21)$$

It is straightforward to prove that the self-dual property, $\mathcal{Q} = \mathcal{Q}^C$, using proof by contradiction.

Similar to linear programming, SOCP problems can be solved using the interior-point method with the logarithmic barrier function for the constraints. Even though the primal or dual interior-point methods developed for linear programming can be directly extended to SOCP, as described in Nesterov and Nemirovski (1994), the primal-dual interior-point method is preferred due to its numerical robustness.

The dual form of Eqs. (18–20) is defined as follows:

$$\max \mathbf{b}^T \mathbf{y} \quad (22)$$

$$\text{s.t. } \mathbf{A}^T \mathbf{y} + \mathbf{z} = \mathbf{c} \quad (23)$$

$$\mathbf{z} \in \mathcal{Q}^C = \mathcal{Q} \quad (24)$$

The general procedure of the primal-dual interior-point algorithm combines the primal and dual feasibility and the complementarity conditions and yields

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad \mathbf{x} \in \mathcal{Q} \quad (25)$$

$$\mathbf{A}^T \mathbf{y} + \mathbf{z} = \mathbf{c} \quad \mathbf{z} \in \mathcal{Q} \quad (26)$$

$$\mathbf{x}^T \mathbf{z} = 0 \quad (27)$$

The above system of linear equations is almost identical to the corresponding one for linear programming, except for the extra conic constraints in the primal and dual feasibility equations in Eq. (25) and Eq. (26). In fact, one can combine the two conic constraints and the complementary slackness condition in Eq. (27), and reduce them to a more suitable form (Eq. (30)) for numerical solvers. The revised system of linear equations becomes

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad (28)$$

$$\mathbf{A}^T \mathbf{y} + \mathbf{z} = \mathbf{c} \quad (29)$$

$$\mathbf{x} \circ \mathbf{z} = 0 \quad (30)$$

$$\text{where } \mathbf{x} \circ \mathbf{z} := \begin{bmatrix} \mathbf{x}^T \mathbf{z} \\ x_0 \bar{\mathbf{z}} + z_0 \bar{\mathbf{x}} \end{bmatrix} = \begin{bmatrix} \mathbf{x}^T \mathbf{z} \\ x_0 z_1 + z_0 x_1 \\ \vdots \\ x_0 z_n + z_0 x_n \end{bmatrix}. \quad (31)$$

We refer readers to Alizadeh and Goldfarb (2001) for detailed derivations. The primal-dual interior-point method solves this system of linear equations (Eqs. (28–30)) using Newton's method. The

optimization begins with a relaxed version of the complementary condition (Eq. (30)), and slowly strengthens this condition as iterations proceed. Iterations stop once the residual is less than a pre-selected threshold.

References

- Alizadeh, F., Goldfarb, D., 2001. Second-order cone programming. Technical Report. RRR Report number 51-2001, RUTCOR, Rutgers University, 2001.
- Allison, T., McCarthy, G., Wood, C.C., Darcey, T.M., Spencer, D.D., Williamson, P.D., 1989. Human cortical potentials evoked by stimulation of the median nerve. I. Cytoarchitectonic areas generating short-latency activity. *J. Neurophysiol.* 62, 694–710.
- Auranen, T., Nummenmaa, A., Hämäläinen, M.S., Jääskeläinen, I.P., Lampinen, J., Vehtari, A., Sams, M., 2007. Bayesian analysis of the neuromagnetic inverse problem with P -norm priors. *NeuroImage* 26, 870–884.
- Baillet, S., Garnero, L., 1997. A Bayesian approach to introducing anatomic-functional priors in the EEG/MEG inverse problem. *IEEE Trans. Biomed. Eng.* 44, 374–385.
- Barth, D.S., Sutherling, W., Engel, J., Beatty, J., 1982. Neuromagnetic localization of epileptiform spike activity in the human brain. *Science* 218, 891–894.
- Bijma, F., de Munck, J.C., Heethaar, R.M., 2005. The spatiotemporal MEG covariance matrix modeled as a sum of Kronecker products. *NeuroImage* 27, 402–415.
- Brooks, D.H., Ahmad, G.F., MacLeod, R.S., Maratos, G.M., 1997. Inverse electrocardiography by simultaneous imposition of multiple constraints. *IEEE Trans. Biomed. Eng.* 44, 374–385.
- Camilo, L., Long, C.J., Hämäläinen, M.S., Brown, E.N., Purdon, P.L., 2007. Parameter estimation and dynamic source localization for the magnetoencephalography (MEG) inverse problem. *ISBI*, 1092–1095.
- Cottareau, B., Jerbi, K., Baillet, S., 2007. Multiresolution imaging of MEG cortical sources using an explicit piecewise model. *NeuroImage* 38, 439–451.
- Dale, A.M., Sereno, M., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. *J. Cogn. Neurosci.* 5, 162–176.
- Dale, A.M., Fischl, B., Sereno, M., 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. *NeuroImage* 9, 179–194.
- Ding, L., He, B., 2008. Sparse source imaging in electroencephalography with accurate field modeling. *Hum. Brain Mapp.* 29 (9), 1053–1067.
- Donoho, D.L., Elad, M., 2003. Maximal sparsity representation via L_1 minimization. *Proc. Natl. Acad. Sci. U. S. A.* 100, 2197–2202.
- Efron, B., Tibshirani, R., 2002. Empirical Bayes methods and false discovery rates for microarrays. *Genet. Epidemiol.* 23, 70–86.
- Fischl, B., Sereno, M., Dale, A.M., 1999. Cortical surface-based analysis: II. Inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9, 195–207.
- Fors, N., Hari, R., Salmelin, R., Ahonen, A., Hämäläinen, M., Kajola, M., Knuutila, J., Simola, J., 1994. Activation of the human posterior parietal cortex by median nerve stimulation. *Exp. Brain Res.* 99, 309–315.
- Galka, A., Yamashita, O., Ozaki, T., Biscay, R., Valdes-Sosa, P., 2004. A solution to the dynamical inverse problem of EEG generation using spatiotemporal Kalman filtering. *NeuroImage* 23, 435–453.
- Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* 15, 870–878.
- Geva, A.B., 1998. Bioelectric sources estimation using spatio-temporal matching pursuit. *J. Appl. Signal Processing* 5, 195–208.
- Gorodnitsky, I.F., Rao, B.D., 1997. Sparse signal reconstruction from limited data using FOCUSS: a re-weighted minimum norm algorithm. *IEEE Trans. Signal Processing* 45, 600–616.
- Hadamard, J., 1902. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin* 49–52.
- Hämäläinen, M.S., MNE software user's guide. NMR Center, Mass General Hospital, Harvard University.
- Hämäläinen, M.S., Ilmoniemi, R., 1984. Interpreting measured magnetic fields of the brain: Estimates of current distributions. Technical Report TKK-F-A559.
- Hämäläinen, M.S., Sarvas, J., 1989. Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE Trans. Biomed. Eng.* 36, 165–171.
- Hämäläinen, M.S., Hari, R., Ilmoniemi, R., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalography – theory, instrumentation, and applications to non-invasive studies of the working human brain. *Rev. Mod. Phys.* 65, 413–497.
- Hari, R., Fors, N., 1999. Magnetoencephalography in the study of human somatosensory cortical processing. *Philos. Trans. R. Soc. Lond. B, Biol. Sci.* 354, 1145–1154.
- Huang, M., Aine, C.J., Supek, S., Best, E., Ranken, D., Flynn, E.R., 1998. Multi-start downhill simplex method for spatio-temporal source localization in magnetoencephalography. *Electroencephalogr. Clin. Neurophysiol.* 108, 32–44.
- Huang, M.X., Dale, A.M., Song, T., Halgren, E., Harrington, D.L., Podgomy, I., Canive, J.M., Lewis, S., Lee, R.R., 2006. Vector-based spatial-temporal minimum L_1 -norm solution for MEG. *NeuroImage* 31, 1025–1037.
- Huizenga, H.M., de Munck, J.C., Waldorp, L.J., Grasman, R.P.P.P., 2002. Spatiotemporal EEG/MEG source analysis based on a parametric noise covariance model. *IEEE Trans. Biomed. Eng.* 49, 533–539.
- Kaipio, J., Somersalo, E., 2004. *Statistical and Computational Inverse Problems*. Springer.
- Jerbi, K., Baillet, S., Mosher, J.C., Nolte, G., Gamero, L., Leahy, R.M., 2004. Localization of realistic cortical activity in MEG using current multipoles. *NeuroImage* 22, 779–793.
- Kass, J., Nelson, R., Sur, M., Lin, C.-S., Merzenich, M., 1979. Multiple representations of the body within the primary somatosensory cortex of primates. *Science* 204, 521–523.
- Knake, S., et al., 2006. The value of multichannel MEG and EEG in the presurgical evaluation of 70 epilepsy patients. *Epilepsy Res.* 69, 80–86.
- Lin, F.-H., Belliveau, J.W., Dale, A.M., Hämäläinen, M.S., 2006. Distributed current estimates using cortical orientation constraints. *Hum. Brain Mapp.* 27, 1–13.
- Limpiti, T., Van Veen, B.D., Wakai, R.T., 2006. Cortical patch basis model for spatially extended neural activity. *IEEE Trans. Biomed. Eng.* 53, 1740–1754.
- Liu, A.K., Belliveau, J.W., Dale, A.M., 1998. Spatiotemporal imaging of human brain activity using functional MRI constrained magnetoencephalography data: Monte Carlo simulations. *Proc. Natl. Acad. Sci. U. S. A.* 95, 8945–8950.
- Malioutov, M., Çetin, M., Willsky, A.S., 2005. A sparse signal reconstruction perspective for source localization with sensor arrays. *IEEE Trans. Signal Processing* 53, 3010–3022.
- Matsuura, K. and Okabe, Y., Multiple current-dipole distribution reconstruction reconstructed by modified selective minimum-norm method. *Advances in Biomagnetism Res.: Biomag 96*, Springer, New York, 290–293.
- Mosher, J.C., Lewis, P.S., Leahy, R.M., 1992. Multiple dipole modeling and localization from spatio-temporal MEG data. *IEEE Trans. Biomed. Eng.* 39, 541–557.
- Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. *Computer Journal* 7, 308–313.
- Nesterov, Y., Nemirovski, A., 1994. Interior-point polynomial methods in convex programming. *Studies in Applied Mathematics*. SIAM, Philadelphia, PA.
- Oostendorp, T.F., Van Oosterom, A., 1989. Source parameter estimation in inhomogeneous volume conductors of arbitrary shape. *IEEE Trans. Biomed. Eng.* 36, 382–391.
- Pantazis, D., Nichols, T.E., Baillet, S., Leahy, R.M., 2005. A comparison of random field theory and permutation methods for the statistical analysis of MEG data. *NeuroImage* 25, 383–394.
- Pascual-Marqui, R.D., Michel, C.M., Lehmann, D., 1994. Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. *Int. J. of Psychophysiol.* 18, 49–65.
- Rao, B.D., Kreutz-Delgado, K., 1999. An affine scaling methodology for best basis selection. *IEEE Trans. Signal Processing* 47, 187–200.
- Rowe, M., Turman, A., Murray, G., Zhang, H., 1996. Parallel organization of somatosensory cortical areas I and II for tactile processing. *Clin. Exp. Pharmacol. Physiol.* 23, 931–938.
- Scherg, M., Von Cramon, D., 1985. Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. *Electroencephalogr. Clin. Neurophysiol.* 62, 32–44.
- Schmitt, U., Louis, A.K., Darvas, F., Buchner, H., Fuchs, M., 2001. Numerical aspects of spatio-temporal current density reconstruction from EEG-/MEG-data. *IEEE Trans. Med. Imaging* 20, 314–324.
- Source modeling software (xfit), Elekta-Neuromag Oy, Helsinki, Finland.
- Strum J.F., 1999. SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.* 11 & 12, 625–653.
- Trujillo-Barreto, N., Aubert-Vazquez, E., Penny, W., 2007. Bayesian M/EEG source reconstruction with spatio-temporal priors. *NeuroImage* 39, 318–335.
- Tuomisto, T., Hari, R., Katila, T., Poutanen, T., Varpula, T., 1983. Studies of auditory evoked magnetic and electric responses: modality specificity and modelling. *Nuovo Cimento D* 2, 471–483.
- Uutela, K., Hämäläinen, M.S., Salmelin, R., 1998. Global optimization in the localization of neuromagnetic sources. *IEEE Trans. Biomed. Eng.* 45, 716–723.
- Uutela, K., Hämäläinen, M.S., Somersalo, E., 1999. Visualization of magnetoencephalographic data using minimum current estimates. *NeuroImage* 10, 173–180.
- Wang, J.Z., Williamson, S.J., Kaufman, L., 1993. Magnetic source imaging based on the minimum-norm least-squares inverse. *Brain Topography* 5, 365–371.
- Weerd, J.P.C., Kap, J.I., 1981. Spectro-temporal representations and time-varying spectra of evoked potentials. *Biol. Cybern.* 41, 101–107.
- Wikström, H., Huttunen, J., Korvenoja, A., Virtanen, J., Salonen, O., Aronen, H., Ilmoniemi, R.J., 1996. Effects of interstimulus interval on somatosensory evoked magnetic fields (SEFs): a hypothesis concerning SEF generation at the primary sensorimotor cortex. *Electroencephalogr. Clin. Neurophysiol.* 100, 479–487.
- Wood, C.C., 1982. Application of dipole localization methods to source identification of human evoked potentials. *Ann. N.Y. Acad. Sci.* 388, 139–155.
- Zhang, Y., Ghodrati, A., Brooks, D.H., 2005. An analytical comparison of three spatio-temporal regularization methods for dynamic linear inverse. *Problems in a common statistical framework. Inverse Probl.* 21, 357–382.