

Using ICA and realistic BOLD models to obtain joint EEG/fMRI solutions to the problem of source localization

Ted Brookings^a, Stephanie Ortigue^b, Scott Grafton^{b,*}, Jean Carlson^c

^a Department of Biology, Brandeis University, Waltham, MA 02454, USA

^b Department of Psychology, University of California, Santa Barbara, 93106, USA

^c Department of Physics, University of California, Santa Barbara, CA 93106, USA

ARTICLE INFO

Article history:

Received 5 March 2008

Revised 17 August 2008

Accepted 31 August 2008

Available online 23 September 2008

ABSTRACT

We develop two techniques to solve for the spatio-temporal neural activity patterns using Electroencephalogram (EEG) and Functional Magnetic Resonance Imaging (fMRI) data. EEG-only source localization is an inherently underconstrained problem, whereas fMRI by itself suffers from poor temporal resolution. Combining the two modalities transforms source localization into an overconstrained problem, and produces a solution with the high temporal resolution of EEG and the high spatial resolution of fMRI. Our first method uses fMRI to regularize the EEG solution, while our second method uses Independent Components Analysis (ICA) and realistic models of Blood Oxygen-Level Dependent (BOLD) signal to relate the EEG and fMRI data. The second method allows us to treat the fMRI and EEG data on equal footing by fitting simultaneously a solution to both data types. Both techniques avoid the need for ad hoc assumptions about the distribution of neural activity, although ultimately the second method provides more accurate inverse solutions.

© 2008 Elsevier Inc. All rights reserved.

Introduction

Electroencephalogram (EEG) and Functional Magnetic Resonance Imaging (fMRI) are two commonly used modalities for investigating human brain states in cognitive neuroscience experiments. Both are noninvasive, but in other respects they are complimentary. EEG measures voltage changes in electrodes placed on the scalp, whose number ranges commonly from 32 to 256. EEG has millisecond time sensitivity, but spatial information must be inferred through an inversion process, and has at most as many independent spatial measurements as there are electrodes (there may be fewer due to correlations between nearby electrodes) (de Peralta-Menendez and Gonzalez-Andino, 1998; Michel et al., 2004). fMRI measures changes in blood oxygen level (Ogawa et al., 1990; Jens Frahm, 1992) (called the BOLD signal) throughout the brain. It produces a 3D image with a spatial resolution of roughly a few millimeters, but temporal resolution is on the order of a few seconds. Furthermore the BOLD signal is a complicated convolution of brain activity because the blood oxygen level takes several seconds to rise and even longer to fall in response to an impulse of activity. Thus EEG provides an excellent measure of temporal dynamics but a poor measure of spatial locations, and fMRI provides an excellent measure of spatial locations but a poor measure of temporal dynamics.

In this paper we develop two novel methods for source localization using both EEG and fMRI data. By combining the two modalities, the high temporal resolution of EEG can be augmented with the high spatial resolution of fMRI. Existing literature has established the potential gains from combining EEG and Positron Emission Tomography (Heinze et al., 1994), as well as EEG and fMRI (Whittingstall et al., 2007; Gerloff et al., 1996; Dale and Halgren, 2001). However, in past studies the difficulties inherent in combining such dissimilar modalities have led to a reduced scope of the analysis: inclusion of data from only one EEG lead (Calhoun et al., 2006), or avoiding the EEG inverse by using ICA to analyze EEG and fMRI data simultaneously, and thereby obtaining ICA sources that have related EEG and fMRI signals (Moosmann et al., 2008). Other techniques use fMRI to constrain the location of likely sources (Liu et al., 1998, 2006), fitting for the location of dipoles seeded within active fMRI regions (Stancák et al., 2005), or employ an adaptive Wiener filter (Liu and He, 2008) that is updated by EEG and fMRI data.

Here we present two techniques for working with full EEG and fMRI data sets and solving to obtain neural activities throughout the cortex at high spatial and temporal resolution. Our first method uses standard techniques to invert EEG data, but employs fMRI data to constrain the free components of the solution. Throughout this paper, we refer to this technique as our “fMRI regularized inverse”, and we describe it in the fMRI regularized inverse Section. The second method uses model reduction algorithms (Principle Component Analysis, or PCA; and Independent Component Analysis, or ICA) to decrease the size of the inverse problem, and a detailed model of the BOLD signal (discussed in the Model-reduced joint inverse Section) to relate EEG

* Corresponding author.

E-mail addresses: brookings@brandeis.edu (T. Brookings), ortigue@psych.ucsb.edu (S. Ortigue), grafton@psych.ucsb.edu (S. Grafton), carlson@physics.ucsb.edu (J. Carlson).

and fMRI data. This enables us to simultaneously fit the EEG and fMRI data. We refer to this method as our “model-reduced joint inverse”, and it is described in the Model-reduced joint inverse Section. The model-reduced joint inverse has the additional advantage of treating the EEG and fMRI data on equal footing, instead of using the fMRI merely as a constraint. In the Testing the algorithms Section, we evaluate and contrast the effectiveness of these techniques on synthetically generated data to demonstrate the potential effectiveness of using this methodology to analyze data recorded from human subjects.

EEG source localization

We begin with a brief, general description of EEG source localization (Michel et al., 2004). This background provides the starting point for our first method of combining EEG and fMRI (fMRI regularized inverse Section) which begins with EEG, but makes use of the fact that the basic problem of EEG source localization is underconstrained. When EEG is considered alone, different methods employ different techniques for *regularization*—selecting a particular solution out of an infinite family of possibilities, which would otherwise share an equal quality of fit. In the fMRI regularized inverse Section, we introduce a specific method of regularization which uses complementary fMRI data. Our solution optimizes agreement with the fMRI data without altering the quality of fit to EEG data.

To perform source localization and thus obtain spatial information from EEG data, one must model the effect that currents within the brain produce on voltages measured at the scalp. Sources in different regions of the brain contribute to electrode voltages by summing linearly. Thus, given a set of source currents, the measured EEG signal can be calculated with a linear Green's function matrix (or *lead field*).

This matrix is typically calculated by modeling the head as a series of concentric regions with different conductivities, and solving Maxwell's equations for dipole current sources. Often the head model is quite simple, although it is possible to incorporate 3D images from a structural MRI (as in Fig. 1) of the subject's head. More complicated head models may produce a more accurate lead field matrix, but the resulting matrices are used the same way. Since the EEG sources are dipoles, they possess a magnitude and an orientation. Thus typically the lead field maps vector sources to scalar EEG measurements. However, it is possible to calculate a scalar lead field (Grave de Peralta Menendez et al., 2004) and invert to obtain cortical current density, or to reduce a vector lead field to a scalar lead field by using anatomy to fix the orientation of the dipoles (Dale and Sereno, 1993; Liu et al., 1998), and invert to obtain the dipole magnitude.

Once the lead field L has been calculated, if the neural sources \vec{S} are known, then the expected EEG signals \vec{E} can be computed simply through the linear *forward model*

$$\vec{E} = L\vec{S} \quad (1)$$

Source localization is then a matter of inverting the linear equation relating the unknown source currents to the known EEG signal via the Green's function. Therefore for some generalized inverse L^- , we have

$$\vec{S} = L^- \vec{E} \quad (2)$$

This system is highly underdetermined. If the brain is divided up into regions similar in size to fMRI resolution, there are tens of thousands of voxels, but there are only of order one hundred EEG signals to constrain the solution. With no additional constraints, it is possible to add to any solution a set of sources in the null space of the Green's function and obtain another valid solution. Although it is possible to decrease the extra degrees of freedom by appealing to physiology (Phillips et al., 2002) or careful treatment of the electric field, as in ELECTRA (de Peralta-Menendez et al., 2000) and LAURA (de




Fig. 1. Lead field from structural MRI. A structural MRI (top) captures a high resolution 3D image of the brain. This image can be used to identify regions of grey matter within the cortex, which are likely locations for EEG sources (bottom, lateral projection).

Peralta-Menendez and Gonzalez-Andino, 2002), ultimately the EEG data cannot produce a unique solution by itself.

To produce a unique solution, the problem must be regularized. This is typically accomplished through a simplifying hypothesis: the solution is of minimum norm (Hämäläinen and Ilmoniemi, 1994), or constrained by a prescribed tradeoff between quality of fit and smoothness as in LORETA (Pascual-Marqui et al., 1994), or by imposing covariance constraints on the solution as in Beamformer (Sekihara et al., 2001) techniques, etc. Alternatively, Bayesian methods (Friston et al., 2008; Kiebel et al., 2008; Karl Friston, 2006; Phillips et al., 2005) incorporate hypotheses through the introduction of explicit *priors*. These priors are used as a starting probability distribution for source activities, and this distribution is updated by EEG data in a manner consistent with Bayesian statistics. However, none of these hypotheses are guaranteed to be correct. Thus, aside from convenience, there is little reason to suppose that they are more “correct” than the infinitely many other solutions that equally agree with data (including those produced by different regularization schemes). More problematically, these regularization techniques actually exclude reasonable source distributions from possible solutions when those source distributions do not conform to the regularizing hypothesis. Because of this, for *any* given EEG-only regularization technique, there are distributions of sources that can *never* be found as a solution, no matter what the experiment is, or what EEG data is produced (de Peralta-Menendez and Gonzalez-Andino, 1998, 2000). If the true source distribution is one of those unobservable distributions, the solution may bear little resemblance to the true sources.

fMRI regularized inverse

One way to improve upon regularization schemes is to use an independent source of data (in our case, fMRI) to choose between the infinitely many solutions allowed by the EEG data. Our regularization strategy is to start with a minimum norm inverse of EEG data, and then alter it in ways that improve agreement with fMRI data, without altering the quality of the fit to EEG data. We accomplish this by adding vectors in the null space of the lead field matrix (Ahlfors and Simpson, 2004). We still arrive at a unique solution, but we do so with more principled assumptions, and it is not *a priori* guaranteed that some source distributions are unobservable.

To construct our fMRI regularized inverse, we begin by considering solutions of the form

$$\vec{S} = L^- \vec{E} + N_L \vec{\alpha}, \tag{3}$$

where L^- is the Moore–Penrose generalized inverse (Menke, 1984), and N_L is a matrix whose columns are vectors in the null space of L . With this form, the unknown vector $\vec{\alpha}$ controls the projection of the solution into the null space. Thus, our solution is the minimum norm solution, plus some unknown vector in the null space of L . Fig. 2

schematically illustrates the sizes of the matrices used by our two inversion methods.

Here we briefly justify this choice for the solution form. We choose our sources to minimize disagreement between the actual EEG data and our modeled EEG data. We obtain the modeled EEG data E_{solution} with the lead field matrix and the forward model (Eq. (1)). Thus, we define the solution error as

$$\sigma_{\vec{E}} = \|\vec{E}_{\text{solution}} - \vec{E}\| = \|L\vec{S} - \vec{E}\|. \tag{4}$$

Here $\|M\|$ denotes the sum of the squares of all the elements of matrix M . Substituting the form of our solution from Eq. (1), we obtain

$$\sigma_{\vec{E}} = \|L(L^- \vec{E} + N_L \vec{\alpha}) - \vec{E}\|. \tag{5}$$

Then since N_L is a matrix of columns in the null space of L , $LN_L = 0$, and therefore

$$\sigma_{\vec{E}} = \|LL^- \vec{E} - \vec{E}\|. \tag{6}$$

Thus the error is independent of the value of $\vec{\alpha}$ – adding the null vector to the solution does not affect the quality of the EEG fit. Regularization is then a process whereby we choose a specific $\vec{\alpha}$ (for

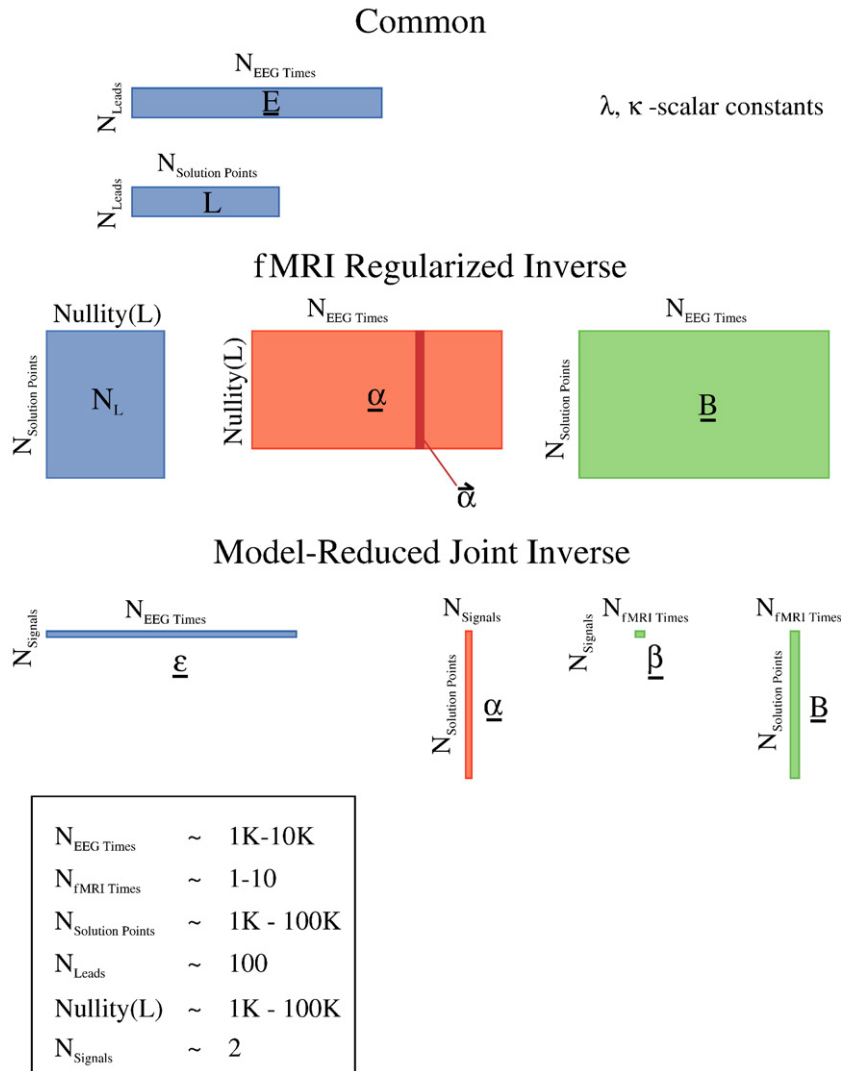


Fig. 2. Matrix sizes by inversion method. Here we depict (heuristically, not to scale) the sizes and shapes of the matrices used by the different inversion methods. “Common” denotes matrices used in both the fMRI regularized inverse and the model-reduced joint inverse. Note that the **B** matrix is much larger in the fMRI regularized inverse because it is interpolated. One consequence of the different strategies is that the Model-Reduced Joint Inverse solves for a much smaller (in terms of matrix size) solution. The vector $\vec{\alpha}$ is not used in the calculation of the fMRI regularized inverse (but is referred to in the fMRI regularized inverse Section). It corresponds to a single column of the matrix **α**.

each time step) and arrive at a unique solution. Our goal is to choose an $\vec{\alpha}$ that causes our solution to be maximally consistent with accompanying fMRI data. Since the fMRI data will vary with different experiments, the constraints placed on the solution will vary as well. Thus, this technique dramatically decreases the number of solutions that *in principle* cannot be found, and thereby avoids eliminating from consideration potentially valid solutions.

Any attempt at combining information from EEG and fMRI invariably confronts the problem that EEG and fMRI data are measuring fundamentally different quantities. fMRI measures blood oxygen level, which is affected by blood volume, flow rate, and of course the oxygen content of the blood. The various factors contributing to the BOLD signal are directly or indirectly related to the metabolism of brain cells, and thus neural activity. On the other hand, the currents responsible for EEG are ionic currents produced when neuronal activity causes alterations in the flow of ions into and out of neurons. Since this activity varies spatially, currents are produced in the brain fluid, roughly synchronous with and proportional to activity. Additionally, current is fundamentally a vector quantity, although as mentioned previously, it is possible to describe the EEG sources as scalars by fixing the direction of the currents (e.g. by anatomical constraints), or by working with current source density, or local field potential (voltage as measured by an electrode inserted into the brain).

Several studies by different researchers have shown that BOLD signal and neural activity are correlated – in fact, if the frequency of the neural activity is fixed, the BOLD signal is roughly linearly proportional to the neural activity (Heeger et al., 2000; Heeger and Ress, 2002; Shmuel et al., 2006; Logothetis et al., 2001). BOLD signal has also been shown to be proportional to local field potential (Shmuel et al., 2006; Logothetis et al., 2001). Since the current sources should be proportional to neural activity (and local field potential) it is reasonable to expect the sources of the EEG signal to be proportional to the BOLD signal. This is the basis of our regularization scheme: modes from the null space of the lead field are added to ensure that the solution is proportional to the BOLD signal \vec{B} .

To regularize, we minimize the disagreement between the solution and linearly-scaled fMRI at each time. We introduce the unknown constants λ (which describes the proportionality between fMRI signal and neural activity) and κ (which is an offset). While we do not expect introducing λ and κ to literally “convert” BOLD signal into EEG signal or even to fully capture the relationship between the two measures, doing so nonetheless provides a concrete relationship between them that we can use to obtain a unique solution to our minimization problem:

$$\sigma_{\vec{B}} = \min_{\{\lambda, \kappa, \vec{\alpha}\}} \left\| |L^- \vec{E} + N_L \vec{\alpha}| - \lambda \vec{B} - \kappa \right\|, \quad (7)$$

with $|M|$ used to denote the matrix whose elements are the absolute values of the elements of M . The parameters λ and κ are unknown *a priori* because they vary from machine to machine (and even between trials on a single machine). Thus λ , κ , and $\vec{\alpha}$ are determined by our minimization procedure. Since we assume the proportionality constants are not changing over the course of a brief period of time, we obtain a single set of values for λ and κ that characterize the relationship between BOLD and neural activity for the entire experiment.

An additional issue is that the BOLD signal is temporally very different from the EEG signal. We will discuss this in detail in the Testing and algorithms Section, but it can be thought of (as a simple approximation) as following neural activity, but delayed by roughly 4 s. Also, fMRI data is sampled at a much lower rate than EEG data (a typical fMRI machine may record a new image every 2 s, whereas EEG is often sampled every 2 ms), and therefore most EEG data will not have accompanying BOLD signal at *precisely* the same time.

Typically analysis involves creating a poststimulus time histogram (PSTH) collecting all images taken after a trial of a certain type: for example, a subject may look at pictures of spiders or puppies, in which

case all “spider” images would be grouped together and categorized by the relative time since the most recent spider picture was displayed (EEG data is averaged through a similar process). If there is a variable inter-trial interval that is not synchronous with the image acquisition, then the BOLD signal is effectively sampled several times per second. Because the actual signal is changing much more slowly than this (over the course of a few seconds), it can be interpolated or shifted in time with a Discrete Fourier Transform or a Fourier Series without missing features. Thus we shift the transformed signal to account for the temporal delay. We group vectors together so that for each time t , $\vec{\alpha}(t)$ for that time is a column in a matrix α . Similarly, \vec{B} is replaced by B and \vec{E} is replaced by E . Then Eq. (7) becomes a single equation to simultaneously solve for λ , κ , and the matrix α (i.e. we solve for every vector $\vec{\alpha}(t)$ simultaneously):

$$\sigma_B = \min_{\{\lambda, \kappa, \alpha\}} \left\| |L^- E + N_L \alpha| - \lambda B - \kappa \right\|. \quad (8)$$

Eq. (8) is not a linear equation. The EEG source currents are inherently signed – large positive and large negative currents both indicate large neuronal activity, but they contribute oppositely to measured EEG signal. Thus it is the absolute value of the EEG currents that is proportional to the BOLD signal. However, this nonlinearity is mild enough that Eq. (8) can be solved iteratively through the following program:

```

 $\underline{S}_0 = L^- E$ 
Repeat (increasing n)
   $\min_{\{\lambda_n, \kappa_n\}} \left\| |S_{n-1}| - \lambda_n B - \kappa_n \right\|$ 
  define  $F_n$  s.t.  $|F_n| = |\lambda_n B + \kappa_n|$  and
     $\text{sign}(F_n) = (S_{n-1})$ 
   $\alpha_n = \arg \min_{\alpha} \left\| |S_0 + N_L \alpha - F_n| \right\|$ 
   $\underline{S}_n = S_0 + N_L \alpha_n$ 
Stop when  $\left\| \underline{S}_n - \underline{S}_{n-1} \right\| < \text{Tolerance}$ 

```

Once the difference between consecutive iterations of this procedure falls below the prescribed tolerance, $\underline{S} = \underline{S}_n$ is a solution to the nonlinear minimization problem. Each of the minimization steps is linear, and reasonably fast. The first, minimizing over λ and κ , is merely a best-fit line (albeit one with many data points). The second is solved to obtain α_n as:

$$\alpha_n = N_L^- \left(F_n - \underline{S}_0 \right), \quad (9)$$

and N_L^- need only be calculated once in advance. We continue the iteration until the change in \underline{S}_n from one iteration to the next is below our specified tolerance. The result of this algorithm, \underline{S} is the fMRI regularized inverse solution \underline{S} that we sought. We present the results from tests of the fMRI regularized inverse in the Testing the algorithms Section.

Model-reduced joint inverse

To better incorporate fMRI data into source localization, we next incorporate a model that provides detail beyond merely noting that the BOLD signal is proportional to neural activity. As mentioned previously, the BOLD signal response is delayed (by roughly 4 s for a brief impulse of neural activity). Our fMRI regularized inverse takes advantage of the correlation between BOLD signal and neural activity, but using the correlation alone is not ideal because the BOLD signal does not simply mirror lagged neural activity. The signal is also temporally broadened, and therefore simply shifting BOLD to eliminate the time lag does not result in a signal that matches neural activity. By incorporating a detailed model of BOLD into an inversion algorithm, we obtain closer agreement between our model and data, and thus a more accurate inverse.

Detailed models of the BOLD response exist – we use a variant of the Balloon/Windkessel model (Buxton et al., 1998). The model we use (Robinson et al., 2006) is a set of coupled nonlinear differential equations that start with neural activity as an input, and model physical changes in the blood vessels, changes in the blood oxygen content, and the measured response of these changes in an fMRI machine. The result is shown in Fig. 3. The differential equations are nonlinear, but over a range of input activities the amplitude of the BOLD response varies approximately linearly with the amplitude of the neural activity. This *must* be the case, because we know that BOLD is proportional to neural activity.

Our fMRI regularized inverse cannot make use of the realistic model of BOLD response because it does not have an estimate of how the neural activity varies with time until it produces its solution. In order to incorporate a modeled BOLD response into our inversion procedure, we take an alternate approach that starts by identifying candidate neural activity signals, and then fits them to data. Principle Component Analysis (PCA) and Independent Component Analysis (ICA) are ideally suited to this job. PCA extracts orthogonal signals that explain the majority of variance in the data (while reducing noise) and ICA forms linear combinations of those orthogonal vectors to produce a basis of signals that are maximally independent of each other. PCA associates a *singular value* with each signal, and more prominent signals have larger singular values. Fig. 4 shows a typical example of this process applied to synthetic EEG data and real experimental data. In the inset, the two largest singular values stand out above the tail of smaller values that correspond to noise signals. This indicates that the optimal analysis would result from selecting two PCA/ICA signals (although it is often advisable to redo the analysis later with one or more additional signals, to see how it affects the final fit).

The procedure to compute our model-reduced joint inverse begins by using PCA followed by ICA on EEG data to produce candidate EEG signals $\underline{\varepsilon}$ which are organized so that each signal makes up one row of the matrix, and each column corresponds to an EEG sample time. Next we use $\underline{\varepsilon}$ and the BOLD model to produce candidate BOLD signals $\underline{\beta}$ (with a similar matrix structure, only far fewer rows because fMRI is sampled less frequently – see Fig. 2). $\underline{\beta}$ therefore incorporates distortions of the signals in $\underline{\varepsilon}$ that include time-shifting, broadening, and other effects.

We simultaneously fit linear combinations of these signals to EEG and fMRI data by solving the minimization problem:

$$\min_{\{\alpha_{\lambda,\kappa}\}} \sigma = \|(\underline{E} - \underline{L}\alpha)\underline{\xi}_E\| + \|(\lambda\underline{B} + \kappa - \alpha\underline{\beta})\underline{\xi}_B\|, \quad (10)$$

where $\underline{\xi}_E$ and $\underline{\xi}_B$ are matrices whose entries are weights for the error in the EEG and fMRI components of the fit, corresponding to the




Fig. 3. BOLD response. Here we show the modeled BOLD response after an impulse of activity at $t=0$.

reciprocal of the measurement uncertainty for EEG data and BOLD data respectively. They can be chosen to weight some elements of the data set more heavily (e.g. force the fMRI discrepancy to be less than the EEG), in accordance to prior knowledge that the experimenter may possess. We set both ξ_E and ξ_B equal to one when inverting our synthetic data. It should be noted that we have assumed that BOLD data exists for each solution point. In this paper we work with synthetic data and can make this so by fiat. However, with real data the anatomical MRI used to obtain the lead field and solution points has a higher resolution than the fMRI that measures the BOLD signal. Thus with real data, one must either down-sample the number of solution points, or interpolate the BOLD signal (e.g. by assigning each solution point the value of BOLD at the nearest fMRI voxel).

Although superficially the minimization problem appears to be complicated, it is linear and can be solved efficiently (the computational cost varies approximately as the cube of the number of solution points, and a modern personal computer can invert for ≈ 1000 solution points in a fraction of a second). Once α has been determined, the solution is

$$\underline{S} = \alpha \underline{\varepsilon}. \quad (11)$$

We derive the solution to this minimization problem in Appendix C.

Aside from a generally more accurate treatment of BOLD, this technique has several specific advantages.

- Linearity guarantees a fast, exact solution.
- Since the number of temporal signals is far less than the number of sampled times, the size of the inverse problem is dramatically reduced.
- BOLD does not have to be interpolated or shifted in time.
- Some types of missing fMRI sources (Liu et al., 2006) are treated appropriately. Sources with brief durations or frequencies that evoke minimal bold response will (assuming the signals are found in the EEG data by PCA) result in negligible modeled BOLD response in $\underline{\beta}$ and therefore EEG data will control the fit.
- This procedure allows the EEG and fMRI data to be inverted simultaneously, and each fMRI data point and each EEG data point are treated procedurally the same in the inversion. Thus we treat fMRI and EEG on more equal footing, instead of using fMRI merely as a constraint on an EEG solution.

The basic technique is also very flexible and can be altered to include additional temporal signals that the experimenter may expect to be present. For example, the experimenter could insert a cardiac signal as a column of α to provide an alternative to preprocessing data. Additionally, it is possible to group solution points together that the experimenter may expect to have correlated activity (either due to neurophysiological, functional, or anatomical priors, or from PCA/ICA used on fMRI data). We have opted to weight each BOLD data point the same as each EEG data point, but this might not be desirable, and a better solution may be obtainable by adjusting the relative weight assigned to the fMRI data (Liu et al., 1998). This can be accomplished by multiplying the BOLD norm in the minimization problem by an experimenter-chosen constant.

Testing the algorithms

We tested each source localization technique by generating a data set with known activity patterns and then comparing the actual activities to the solutions produced by the inversion technique. Inversion of multiple sources was a critical test, as many existing EEG inversion techniques (such as those based on ad hoc regularization schemes) have particular difficulty with multiple sources. We tested the techniques against sources that were distributed over moderate-sized regions of the brain (roughly 10% of its width). We chose distributed sources (instead of sources with a single active point) because many regularization schemes have

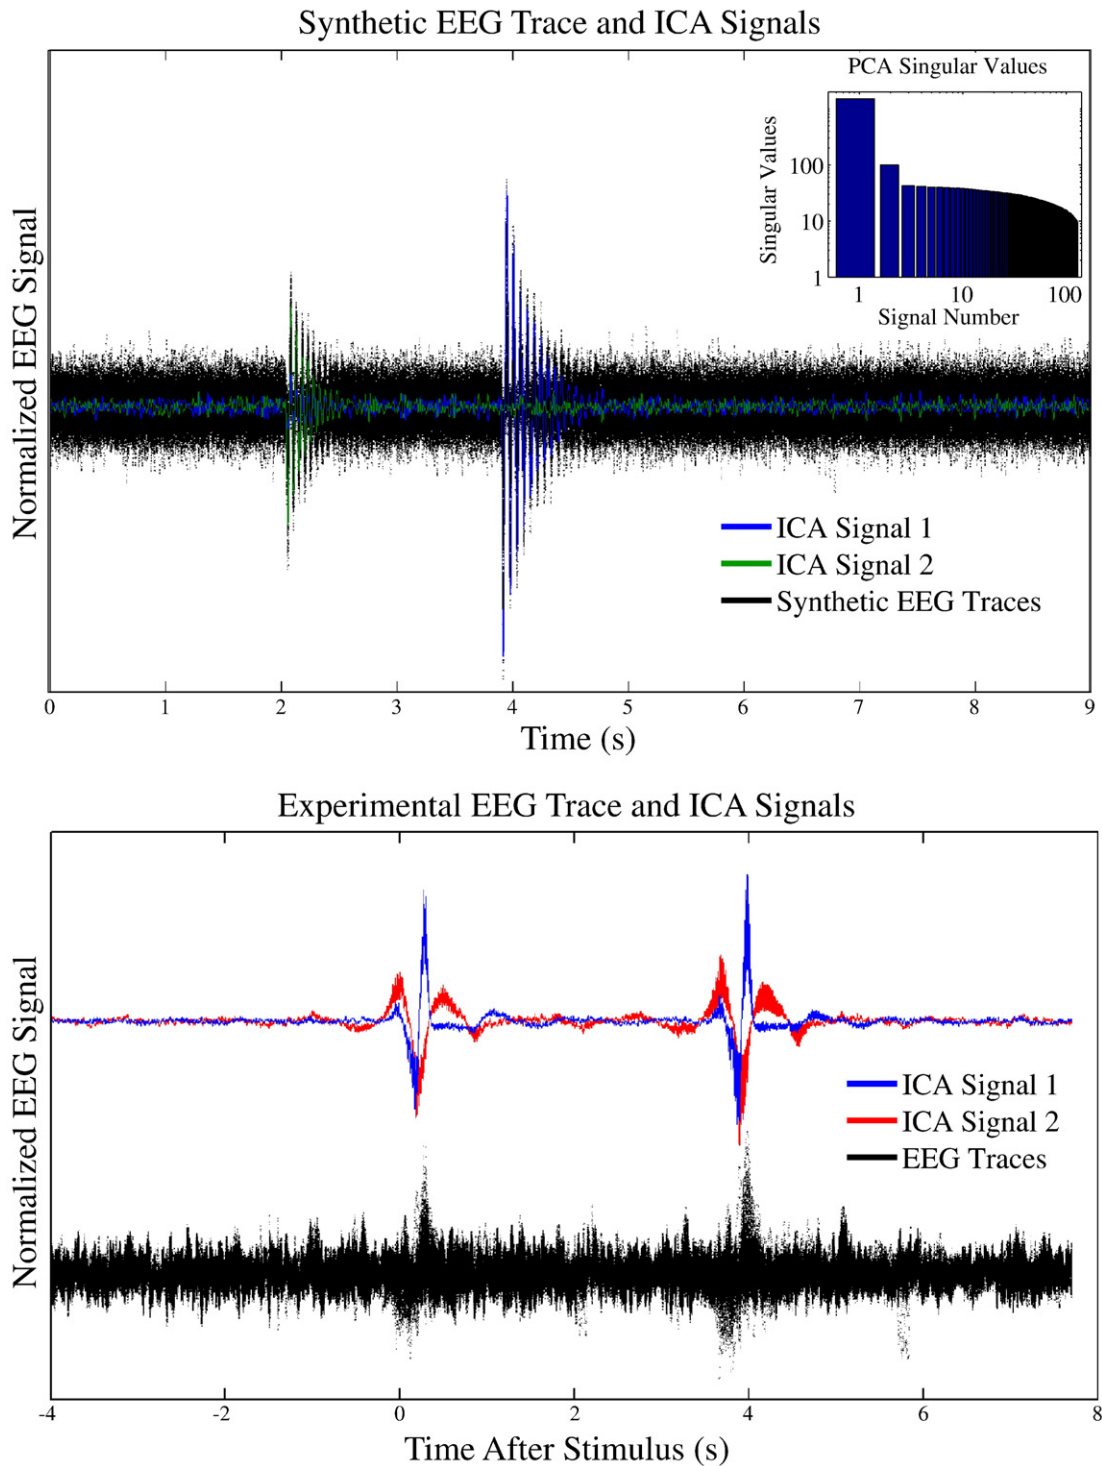


Fig. 4. ICA basis signals. Here we show an example of how ICA extracts the critical features of EEG measurements. The top image shows synthetic EEG data from all 128 leads in black (from the 25% noise condition), with two extracted ICA signals in blue and green. Inset is a bar plot of the magnitude of corresponding singular values: the first two clearly stand out above the tail of noise, indicating that there are two large signals. The bottom image depicts data from a real data set. Human volunteers watched a series of movies with different pairs of everyday objects. Either a human would be present in the scene and look at one object, or one object would be highlighted. Here we show EEG data averaged over the condition where during two sequential movies the same object is looked at by a human, but from one movie to the next the object has switched sides (e.g. from left to right). The first two EEG signals determined by ICA are plotted in blue and red, respectively. The 128 original EEG signals are plotted in black.

difficulty with distributed sources. Of course, the presence of multiple, distributed sources is also to be expected in any experiment involving human subjects.

To produce the synthetic data, we began by calculating a lead field and an associated set of solution points from an anatomical MRI image, using the SMAC (Spinelli et al., 2000) toolkit. We chose the

lead field to cover one thousand solution points and reduced it from a vector lead field to a scalar by fixing the orientation of source dipoles. We then randomly selected one or more of those solution points to have significant activity. For the majority of our trials, activity consisted of two randomly chosen locations, with a cluster of nearby sites all having the same activity pattern. To produce a

realistic time course (Robinson et al., 2006) for such a signal we used the form

$$z(t) = ae^{-\beta t} \sin(\beta \delta t). \quad (12)$$

Parameter values were randomly generated, using the 2006 Robinson et al. paper as a guide for parameter ranges. We used the

“extremely large” range of a (0–80) to ensure that the inversion techniques were tested in the more problematic nonlinear regime of BOLD (although note that the large amplitude did not provide a benefit to the signal-to-noise ratio, because we scaled synthetic noise by the maximum signal amplitude). β ranged from 1.7 to 10.0 (we increased the minimum β to guarantee that the signals fit comfortably within our trial length) and δ ranged from 1.0 to 25.0.

Any sites in an active cluster were deemed “active sites” and assigned an activity $z(t)$, and all other points had an activity $z(t)$ equal to zero. We grouped the activities for all solution points together into the vector $\vec{z}(t)$. Synthetic EEG data was then generated by multiplying the sources $\vec{z}(t)$ by the lead field matrix L . The BOLD signal was generated via the differential equations, acting directly on $z(t)$ for each solution point. The BOLD signal was sampled so as to simulate a series of five trials with fMRI scan times not synchronized to the trial design, collected in a post stimulus time histogram. That is, BOLD was sampled on average five times every 2 s, but with irregular spacing. Gaussian noise (uncorrelated in space and time) was added to all the values for these synthetically produced EEG and BOLD signals. As stated before, the noise level was scaled so that the standard deviation was a pre-chosen fraction of the maximal signal value. For example, if the noise fraction was 0.1, then the Gaussian noise added to the EEG signal would have a standard deviation of 0.1 times the maximum EEG signal amplitude (maximum across all channels and the whole trial), and the Gaussian noise added to the fMRI signal would have a standard deviation of 0.1 times the maximal fMRI signal.

After noise was added, the signals were then preprocessed (fMRI data was whitened, and the EEG signal was passed through a notch filter to remove signal components above 50 Hz). After performing the inversion procedure, the solution sources were compared to the true sources, z . The most critical comparisons were true source location versus solution source location, and true signal time courses versus solution time courses.

For our model reduction procedure, PCA and ICA were performed by the FastICA algorithm (Hyvärinen, 1999), using a “skew” non-linearity. We found that FastICA’s “symmetric approach” produces better quality solutions, but was more likely to fail to converge than the “deflation approach”; thus our procedure made up to two attempts with the symmetric approach, and if they both failed, used the deflation option as a fall-back. Over the course of thousands of test trials with synthetic data sets, this strategy was sufficient to produce a solution in each case. Since it was impossible to manually oversee the generation of so many solutions, invariably some low quality ICA results were obtained. Refinements to the ICA algorithm (or finding a more robust alternative) could further improve the quality of the model-reduced joint inverse.

As expected, a simple minimum norm inverse produces results of limited usefulness. While some high activity areas overlap the true sources, there are high activity areas found far away as well (Fig. 5), and also there are true sources that are not found. The time courses of correctly identified points are of decreased amplitude in low noise




Fig. 5. Typical location accuracy of source inversion. A lateral view of solutions found by different inverse algorithms applied to the same synthetic data set: the minimum norm inverse (top), fMRI regularized inverse (middle), and model-reduced joint inverse (bottom). The different true sources are circles colored magenta and yellow, with X's colored green and red marking the locations found by the inversion procedure (green corresponds to magenta, and red to yellow). As was typical in our trials with low noise data, the model-reduced joint inverse correctly identifies the location of each solution point, while the other two techniques have false positives and false negatives. The fMRI regularized inverse performed marginally better in this example, which was also typical in our low noise tests.

data (Fig. 6) and are lost altogether in high noise data (Fig. 7). Regularizing the inversion with fMRI produces only slightly better results. The fMRI regularized inverse solution contains false positives and negatives, and still finds a greatly decreased signal amplitude. By contrast the model-reduced joint inverse method correctly identifies the true locations of the sources, and the time courses it finds are closer to the correct amplitude. When the synthetic data has very little noise added to it, the model-reduced joint inverse has slightly less noise in its solution than do the other two techniques – this is because the model-reduced joint inverse produces a solution that is a linear combination of the (presumably smooth) signals found by ICA. However, the model-reduced joint inverse also responds more robustly to increased noise, producing useful solutions even after the minimum norm and fMRI regularized inverses produce solutions dominated by noise.

Since the fMRI regularized inverse did not represent a dramatic improvement over the minimum norm but the model-reduced joint inverse did, this raises the question of whether the fMRI information is really aiding the inverse, or if the improvement is solely due to noise-reduction from ICA. To address this issue, we can drop the fMRI term

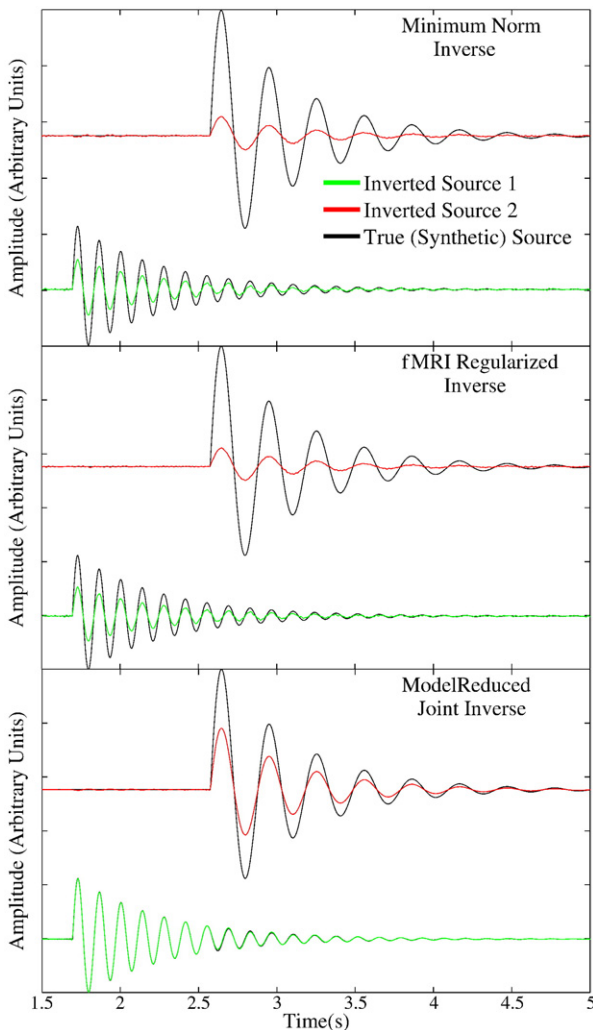


Fig. 6. Typical time course accuracy of source inversion. In this test with low noise data, the signals recovered from minimum norm inverse (top) and fMRI regularized inverse (middle) are significantly reduced in amplitude compared to signals from the model-reduced joint inverse (bottom). The true synthetic signals are in black with the inverted signals in red and green (different colors indicate that the algorithm found two distinct sources). In the bottom signal for the model-reduced joint inverse, the inverted signal overlays the true signal.

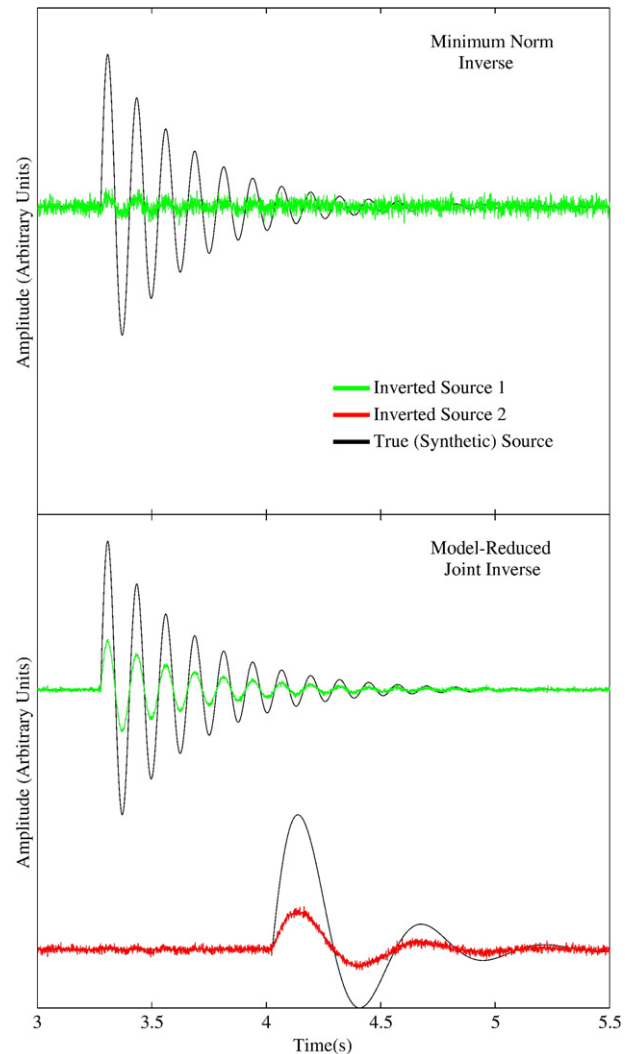


Fig. 7. Time courses in a high noise test. Here are the time courses found by the minimum norm inverse (top) and the model-reduced joint inverse (bottom) in a test with increased noise. The minimum norm solution finds only one source, which is dominated by noise. The model-reduced joint inverse has produced a good solution (albeit one with a decreased amplitude). The fMRI regularized inverse (not shown here) also produces a solution that is dominated by noise.

from the minimization problem for the model-reduced joint inverse, leaving the simpler problem,

$$\min_{\{\alpha_{\lambda, \kappa}\}} \sigma = \|(E - L\alpha_{\lambda, \kappa})\sigma_E\|. \quad (13)$$

Without the fMRI information, the problem is once again underdetermined, and must be regularized. For purposes of comparison, we regularize with a minimum norm. Specifically,

$$\alpha = L^+ E \tilde{\epsilon}. \quad (14)$$

As before, the solution is

$$\underline{S} = \alpha \tilde{\epsilon}. \quad (15)$$

We refer to this solution as the model-reduced minimum norm inverse. The model-reduced minimum norm responds robustly to increased noise, however performance at low noise levels is comparable to the ordinary minimum norm inverse. The model-reduced joint inverse provides a superior solution at all the noise levels we tested, especially when comparing solution location. This suggests that the

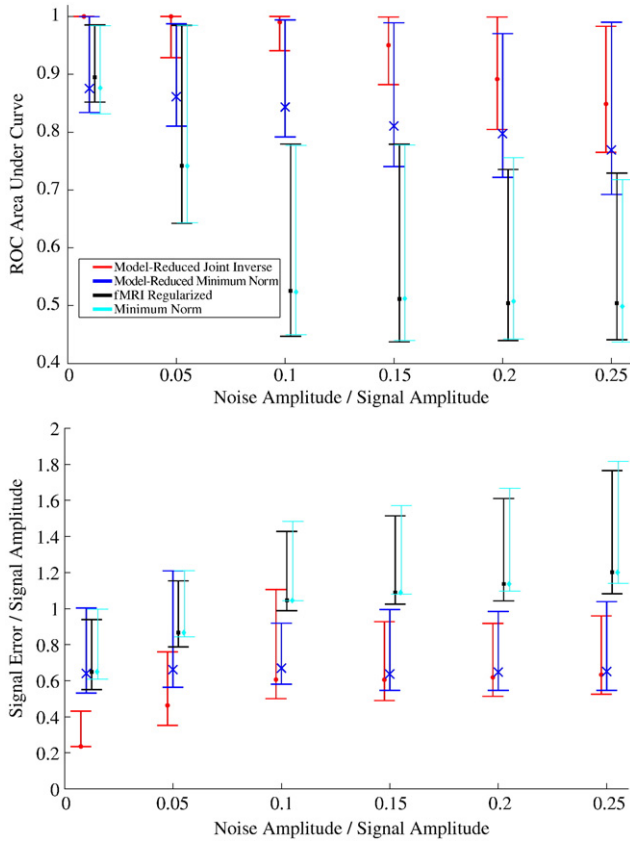


Fig. 8. Accuracy at progressively higher noise levels. Cumulative results for 1000 trials at each of six noise levels. Gaussian noise (uncorrelated in space and time) was added to the synthetic EEG and fMRI data, with the noise amplitude scaling as a fraction of the maximal EEG and fMRI signal strength respectively. The top graph shows area under the ROC curve (for identifying source locations), and the bottom graph shows signal error normalized to signal magnitude (taken in the frequency domain). The results are color-coded by inversion algorithm, with minimum norm inverse as cyan, fMRI regularized inverse as black, model-reduced minimum norm as blue, and model-reduced joint inverse as red. The bars indicate a 90% confidence interval. Note that an ROC area of 0.5 indicates that a test has provided no information, while an ROC area of 1.0 indicates perfect discrimination.

ICA procedure helps to smooth out errors in the signal time course, but the fMRI adds crucial information about source location. Fig. 8 shows plots of the receiver operating characteristic (ROC) analysis of source location, as well as signal error at increasing values of noise.

Conclusions

The fMRI regularized inverse solution provides results that are only a slight improvement over a simple minimum norm solution. Any regularized solution involves selection of a single point near or within the (very large) null space of the lead field matrix. The main advantage of our fMRI regularized inverse is that it arrives at a unique solution in a data-driven manner, and therefore represents an improvement largely in conceptual framework rather than solution quality. We suspect that any solution that takes the form of an EEG-only solution plus a null vector will tend to be biased in the same ways that the minimum norm solution is. Specifically, which solution points are active will be due to influences from EEG data as well as distortions introduced because the resolution matrix $R=L^T L$ is far from the identity. Put another way, the EEG-only part of the solution is distorted because the lead field matrix can only be inverted approximately, and in a manner that biases the result towards certain solution locations. Thus starting with the EEG-only solution and then adjusting it through some regularization process necessarily entails a

poor starting point. The result of regularized EEG inverse solutions is therefore likely to include false positives and false negatives.

By contrast, the model-reduced joint inverse algorithm outperforms minimum norm consistently, and over every range of test parameters we tried. The model-reduced inverse has several advantages, including a general reduction in noise, inherent uniqueness of the solution, and more symmetric treatment of EEG and fMRI data via simultaneous fitting. It is perhaps unsurprising that the more detailed treatment of BOLD results in better determination of source location, but using ICA to extract source signals results in superior characterization of signal time courses as well.

One potential source for further improvement in our algorithm is in the ICA process itself. ICA is an iterative technique, that occasionally fails to converge, and occasionally converges to a sub-optimal solution. Our procedure (skew nonlinearity, two tries with symmetric approach followed by deflation) was able to overcome this difficulty and produce unsupervised solutions to thousands of synthetic trials. However, further improvements are likely possible through additional adjustment of the ICA parameters, human supervision of the process (e.g. experimenter judgment of whether the ICA procedure has been successful), or replacement of ICA with another algorithm that may be superior for this particular purpose.

The principle remaining difficulty with combined EEG-fMRI techniques is accurate registration of the EEG solution points to the underlying structural MRI image – a problem we avoided by dealing with purely synthetic data. However, this subject is an active area of research (Spinelli et al., 2000; Ermer et al., 2001) with solutions already in existence. We expect that the advantages of the model-reduced joint inverse as applied to synthetic data will carry over in application to human data.

Acknowledgments

This work was supported by the David and Lucile Packard Foundation, NSF grant number DMR-0606092, and the Institute for Collaborative Biotechnologies through grants DAAD19-03-D-0004 and W911NF-07-1-0072.

Appendix A. The minimization problem for the model-reduced joint inverse

In this Appendix we show the steps necessary to solve the minimization problem presented in the Model-reduced joint inverse Section. Although the problem and derivation are novel, the numerical algorithms required to obtain the solution are all widely available. The minimization problem is:

$$\min_{\{\alpha, \lambda, \kappa\}} \sigma = \|(E - L\alpha)\epsilon\| + \|(\lambda B + \kappa - \alpha)\beta\|, \quad (\text{A.1})$$

Henceforth we drop the underline notation – all the vectors have been converted to matrices. To begin, we first compute the variation of σ with respect to α , and set it to zero:

$$L^T E \sigma_{\epsilon}^2 \epsilon - L^T L \alpha \epsilon \sigma_{\epsilon}^2 \epsilon^T + (\lambda B + \kappa) \sigma_{\beta}^2 \beta^T - \alpha \beta \sigma_{\beta}^2 \beta^T = 0. \quad (\text{A.2})$$

Right-multiplying this equation by $(\epsilon \sigma_{\epsilon}^2 \epsilon^T)^{-}$ puts this equation into the form of a Sylvester equation,

$$L^T L \alpha + \alpha \beta \sigma_{\beta}^2 \beta^T (\epsilon \sigma_{\epsilon}^2 \epsilon^T)^{-} + (\lambda B + \kappa) \sigma_{\beta}^2 \beta^T (\epsilon \sigma_{\epsilon}^2 \epsilon^T)^{-} - L^T E \sigma_{\epsilon}^2 \epsilon (\epsilon \sigma_{\epsilon}^2 \epsilon^T)^{-} = 0 \quad (\text{A.3})$$

Sylvester equations (equations of the form $A\alpha + \alpha B + C = 0$) are solvable numerically, and have a nice linear property. For matrices $C1$, $C2$ and scalars u , v :

$$\alpha(A, B, uC1 + vC2) = u\alpha(A, B, C1) + v\alpha(A, B, C2). \quad (\text{A.4})$$

Thus we can write

$$\alpha = \alpha_0 + \lambda\alpha_\lambda + \kappa\alpha_\kappa, \quad (\text{A.5})$$

where α_0 , α_λ and α_κ are the solutions to the Sylvester equation considering only the constant parts, the parts with coefficient λ , and the parts with coefficient κ respectively.

Now we substitute this form into Eq. (A.1) to obtain

$$\min_{\{\lambda, \kappa\}} \sigma = \|(E - L(\alpha_0 + \lambda\alpha_\lambda + \kappa\alpha_\kappa)\varepsilon)\sigma_E\| + \|(\lambda B + \kappa C - (\alpha_0 + \lambda\alpha_\lambda + \kappa\alpha_\kappa)\beta)\sigma_B\|, \quad (\text{A.6})$$

In this seemingly complicated expression, we are only interested in solving for λ and κ . Since λ and κ are linear inside the norms, this is just a best-fit plane. The data points are organized into two separate matrices, but that is irrelevant. We simply treat every entry in the matrices as a separate data point, and thereby reduce the problem to one of the form

$$\min_{a,b} \|\|a\vec{x} + b\vec{y} + \vec{z}\|. \quad (\text{A.7})$$

After fitting for λ and κ , we can now substitute their values into Eq. (A.5) and thereby obtain the solution,

$$S = \alpha\varepsilon. \quad (\text{A.8})$$

References

- Ahlfors, S.P., Simpson, G.V., 2004. Geometrical interpretation of fmri-guided meg/eeeg inverse estimates. *NeuroImage* 22 (1), 323–332 (May).
- Buxton, R., Wong, E., Frank, L., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magn. Reson. Med.* 39, 855–864.
- Calhoun, V., Adali, T., Pearlson, G., Kiehl, K., 2006. Neuronal chronometry of target detection: Fusion of hemodynamic and event-related potential data. *NeuroImage* 30, 544–553 (Apr.).
- Dale, A., Sereno, M., 1993. Improved localization of cortical activity by combining eeg and meg with mri cortical surface reconstruction. *J. Cogn. Neurosci.* 5, 162–176.
- Dale, A.M., Halgren, E., 2001. Spatiotemporal mapping of brain activity by integration of multiple imaging modalities. *Curr. Opin. Neurobiol.* 11 (2), 202–208 (Apr.).
- de Peralta-Menendez, R.G., Gonzalez-Andino, S.L., 1998. A critical analysis of linear inverse solutions to the neuroelectromagnetic inverse problem. *IEEE Trans. Biomed. Eng.* 45, 440–448.
- de Peralta-Menendez, R.G., Gonzalez-Andino, S.L., 2000. Discussing the capabilities of laplacian minimization. *Brain Topogr.* 13, 97–104. <http://dx.doi.org/10.1023/A:1026603017734> (Dec.).
- de Peralta-Menendez, R.G., Gonzalez-Andino, S.L., 2002. Comparison of algorithms for the localization of focal sources: Evaluation with simulated data and analysis of experimental data. *Int. J. Bioelectromagn.* 4 (1) <http://www.rgi.tut.fi/ijbem/volume4/number1/menendez/index.htm>.
- de Peralta-Menendez, R.G., Gonzalez-Andino, S., Morand, S., Michel, C., Landis, T., 2000. Imaging the electrical activity of the brain: *Electra. Hum. Brain Mapp.* 9, 1–12.
- Ermer, J.J., Mosher, J.C., Baillet, S., Leahy, R.M., 2001. Rapidly recomputable eeg forward models for realistic head shapes. *Phys. Med. Biol.* 46 (4), 1265–1281. <http://stacks.iop.org/0031-9155/46/1265>.
- Friston, K., Harrison, L., Daunizeau, J., Kiebel, S., Phillips, C., Trujillo-Barreto, N., Henson, R., Flandin, G., Mattout, J., 2008. Multiple sparse priors for the m/eeeg inverse problem. *NeuroImage* 39 (3), 1104–1120 (Feb.).
- Gerloff, C., Grodd, W., Altenmüller, E., Kolb, R., Klose, T.N.U., Voigt, K., Dichgans, J., 1996. Coregistration of eeg and fmri in a simple motor task. *Hum. Brain Mapp.* 4 (3), 199–209.
- Grave de Peralta Menendez, R., Murray, M.M., Michel, C.M., Martuzzi, R., Gonzalez Andino, S.L., 2004. Electrical neuroimaging based on biophysical constraints. *NeuroImage* 21 (2), 527–539 (Feb.).

- Heeger, D.J., Ress, D., 2002. What does fmri tell us about neuronal activity? *Nat. Rev. Neurosci.* 3, 142–151 <http://dx.doi.org/10.1038/nrn730> (Feb.).
- Heeger, D.J., Huk, A.C., Geisler, W.S., Albrecht, D.G., 2000. Spikes versus bold: what does neuroimaging tell us about neuronal activity? *Nat. Neurosci.* 3, 631–633. <http://dx.doi.org/10.1038/76572> (Jul.).
- Heinze, H.J., Mangun, G.R., Burchert, W., Hinrichs, H., Scholz, M., Munte, T.F., Gos, A., Scherg, M., Johannes, S., Hundseshagen, H., Gazzaniga, M.S., Hillyard, S.A., 1994. Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature* 372 (6506), 543–546. <http://dx.doi.org/10.1038/372543a0> (Dec.).
- Hmlinen, M.S., Ilmoniemi, R.J., 1994. Interpreting magnetic fields of the brain: minimum norm estimates. *Med. Biol. Eng. Comput.* 32, 35–42. <http://www.springerlink.com/content/4474345w652h8581>.
- Hyvärinen, A., 1999. Survey on independent component analysis. *Neural Comput. Surv.* 2, 94–128. citeseer.ist.psu.edu/hyv99survey.html.
- Jens Frahm, H.B.K.D.M.W.H., 1992. Dynamic mr imaging of human brain oxygenation during rest and photic stimulation. *J. Magn. Reson. Imaging* 2 (5), 501–505. <http://dx.doi.org/10.1002/jmri.1880020505>.
- Karl Friston, R.H.C.P.J.M., 2006. Bayesian estimation of evoked and induced responses. *Hum. Brain Mapp.* 27 (9), 722–735. <http://dx.doi.org/10.1002/hbm.20214>.
- Kiebel, S.J., Daunizeau, J., Phillips, C., Friston, K.J., 2008. Variational bayesian inversion of the equivalent current dipole model in eeg/meg. *NeuroImage* 39 (2), 728–741 (Jan.).
- 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.* 95 (15), 8945–8950. <http://www.pnas.org/cgi/content/abstract/95/15/8945>.
- Liu, Z., He, B., 2008. fmri-eeg integrated cortical source imaging by use of time-variant spatial constraints. *NeuroImage* 39 (3), 1198–1214 (Feb.).
- Liu, Z., Kecman, F., He, B., 2006. Effects of fmri-eeg mismatches in cortical current density estimation integrating fmri and eeg: a simulation study. *Clin. Neurophysiol.* 117 (7), 1610–1622 (Jul.).
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fmri signal. *Nature* 412, 150–157. <http://dx.doi.org/10.1038/35084005> (Jul.).
- Menke, W., 1984. *Geophysical Data Analysis: Discrete Inverse Theory*. Academic Press, Inc., New York.
- Michel, C.M., Murray, M.M., Lantz, G., Gonzalez-Andino, S., Spinelli, L., Grave de Peralta-Menendez, R., 2004. Eeg source imaging. *Clin. Neurophysiol.* 115 (10), 2195–2222 (Oct.).
- Moosmann, M., Eichele, T., Nordby, H., Hugdahl, K., Calhoun, V.D., 2008. Joint independent component analysis for simultaneous eeg-fmri: Principle and simulation. *Int. J. Psychophysiol.* 67 (3) (Mar.).
- Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. U. S. A.* 87 (24), 9868–9872. <http://www.jstor.org/view/00278424/di993925/99p0076q/0> (dec.).
- 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. Psychophysiol.* 18, 49–65 (Oct.).
- Phillips, C., Rugg, M.D., Friston, K.J., 2002. Anatomically informed basis functions for eeg source localization: combining functional and anatomical constraints. *NeuroImage* 16 (3, Part 1), 678–695 (Jul.).
- Phillips, C., Mattout, J., Rugg, M.D., Maquet, P., Friston, K.J., 2005. An empirical bayesian solution to the source reconstruction problem in eeg. *NeuroImage* 24 (4), 997–1011 (Feb.).
- Robinson, P., Drysdale, P., der Merwe, H.V., Kyriakou, E., Rigozzi, M., Germanoska, B., Rennie, C., 2006. Bold responses to stimuli: dependence on frequency, stimulus form, amplitude, and repetition rate. *NeuroImage* 31, 585–599.
- Sekihara, K., Nagarajan, S.S., Poeppel, D., Marantz, A., Miyashita, Y., 2001. Reconstructing spatio-temporal activities of neural sources using an meg vector beamformer technique. *IEEE Trans. Biomed. Eng.* 48 (7), 760–771 (Jul.).
- Shmuel, A., Augath, M., Oeltermann, A., Logothetis, N.K., 2006. Negative functional mri response correlates with decreases in neuronal activity in monkey visual area v1. *Nat. Neurosci.* 9, 569–577. <http://dx.doi.org/10.1038/nn1675> (Apr.).
- Spinelli, L., Andino, S.G., Lantz, G., Seeck, M., Michel, C.M., 2000. Electromagnetic inverse solutions in anatomically constrained spherical head models. *Brain Topogr.* 13 (2), 115–125. <http://www.springerlink.com/content/j245h18536313151>.
- Stancák, A., Poláček, H., Vrána, J., Rachmanová, R., Hoehstetter, K., Tintra, J., Scherg, M., 2005. Eeg source analysis and fmri reveal two electrical sources in the fronto-parietal operculum during subepidermal finger stimulation. *NeuroImage* 25 (1), 8–20 (Mar.).
- Whittingstall, K., Stroink, G., Schmidt, M., 2007. Evaluating the spatial relationship of event-related potential and functional mri sources in the primary visual cortex. *Hum. Brain Mapp.* 28 (2), 134–142.