



Performance evaluation of the Champagne source reconstruction algorithm on simulated and real M/EEG data

Julia P. Owen^{a,b}, David P. Wipf^a, Hagai T. Attias^c, Kensuke Sekihara^d, Srikantan S. Nagarajan^{a,b,*}

^a Biomagnetic Imaging Laboratory, Dept. Radiology and Biomedical Imaging, UCSF San Francisco, CA, USA

^b Joint Graduate Group in Bioengineering UCSF/UC Berkeley, San Francisco, CA, USA

^c Golden Metallic, Inc. San Francisco, CA, USA

^d Dept. of Systems Design and Engineering, Tokyo Metropolitan University, Tokyo, Japan

ARTICLE INFO

Article history:

Received 12 April 2011

Revised 1 December 2011

Accepted 14 December 2011

Available online 23 December 2011

Keywords:

Inverse problem

Magnetoencephalography

Electroencephalography

Bayesian

ABSTRACT

In this paper, we present an extensive performance evaluation of a novel source localization algorithm, Champagne. It is derived in an empirical Bayesian framework that yields sparse solutions to the inverse problem. It is robust to correlated sources and learns the statistics of non-stimulus-evoked activity to suppress the effect of noise and interfering brain activity. We tested Champagne on both simulated and real M/EEG data. The source locations used for the simulated data were chosen to test the performance on challenging source configurations. In simulations, we found that Champagne outperforms the benchmark algorithms in terms of both the accuracy of the source localizations and the correct estimation of source time courses. We also demonstrate that Champagne is more robust to correlated brain activity present in real MEG data and is able to resolve many distinct and functionally relevant brain areas with real MEG and EEG data.

© 2011 Elsevier Inc. All rights reserved.

Introduction

Magnetoencephalography (MEG) and electroencephalography (EEG) non-invasively detect brain activity by measuring minute magnetic fields or electric potentials with an array of sensors positioned around the head. Videos of brain activity can then be reconstructed from M/EEG recordings. In order to transform sensor recordings into videos, both a forward and inverse models must be solved. The forward model relates the underlying electrical/magnetic activity in the brain to expected sensor measurements. The calculation of the forward model requires specification of a brain source model, a volume conductor model, and a model for the measurement system (potentials or gradients of a magnetic field). The field or potential observed by the sensors are, in part, generated by large ensembles of neurons firing synchronously within the brain. The synchronous activation of group of neurons can be modeled as a current dipole, also called a *source*. The magnetic (and electric) field of a current dipole is fully defined by Maxwell's equations and the forward model is calculated for a specific location, orientation, volume conductor type (geometry and conductivity), head model (spherical shell or realistic-head model), and type of measurement (potential or gradient of magnetic field). Once these parameters are fixed, we can obtain a mapping from every candidate source location or voxel to every sensor, called the *lead field*. Solving the inverse problem in M/EEG

involves relating the recorded sensor data to brain activations. This inverse problem is inherently ill-posed as the number of voxels (typically 3000 to 5000) is much greater than the number of sensors (typically 275 sensors for MEG systems and 64, 128, or 256 electrodes for EEG systems). Finding a solution to the inverse problem is also complicated by the presence of correlated sources, sensor noise, artifacts and biological interference (from sources within and outside the brain). Obtaining accurate source locations and time course estimates is central to the validity of M/EEG as a functional brain imaging method.

Currently, a wide variety of source localization algorithms exists for estimating source activity. These algorithms can be grouped into three types: dipole fitting, spatial scanning, and tomographic. Parametric dipole fitting is a common technique, where a small number of point dipole sources are assumed to generate the MEG data, and the problem reduces to determining parameters of the dipoles, such as the location, orientation, and amplitude, usually by non-linear optimization or search (Moshier et al., 1992; Uutela et al., 1998). Spatial scanning techniques estimate the time course at every candidate location while suppressing the interference from activity at the other candidate source locations. Some examples of scanning techniques are minimum-variance adaptive beamforming (MVAB) and other variants of beamformers (Gross et al., 2001; Sekihara and Nagarajan, 2008; Vrba and Robinson, 2001). The third class of algorithms are tomographic techniques, which model the activity at all candidate source locations simultaneously. Tomographic methods include minimum-norm estimation (MNE) (Hämäläinen and Ilmoniemi, 1994), dynamic statistical parametric

* Corresponding author at: Biomagnetic Imaging Laboratory, Dept. Radiology and Biomedical Imaging, UCSF San Francisco, CA, USA.

E-mail address: sri@ucsf.edu (S.S. Nagarajan).

mapping (dSPM) (Dale et al., 2000), and standardized low resolution brain electromagnetic tomography (sLORETA) (Pascual-Marqui, 2002). Some tomographic methods promote sparseness in the solution (Gorodnitsky and Rao, 1997; Uutela et al., 1999), where the majority of the candidate locations do not have significant activity. Empirical evidence shows that a sparse source model can improve the accuracy of the localization in a noisy environment.

Most of the source reconstruction algorithms from the three classes can be framed in a Bayesian schema. This perspective is useful because at a high level, the prior distribution implicitly or explicitly imposed can be used to differentiate and compare the various source localization methods. Algorithms such as MNE, MCE, sLORETA, and MVAB (and other beamformers) assume a known, fixed prior. Alternatively, the parameters of the prior distribution (hyperparameters) can be learned from the data, referred to as *empirical Bayes*. These ideas are extensively discussed in Wipf et al. (2009). Often empirical Bayesian algorithms are more robust in that they include a full statistical model (Friston et al., 2008; Kiebel et al., 2008; Mattout et al., 2006; Nummenmaa et al., 2007; Phillips et al., 2005; Sahani and Nagarajan, 2004; Zumer et al., 2007).

Recently, we have developed *Champagne*, a novel source reconstruction algorithm that is derived in an empirical Bayesian schema and incorporates deep theoretical ideas about sparse-source recovery from noisy, constrained measurements. Champagne improves upon existing methods of source reconstruction in terms of reconstruction accuracy, robustness, and computational efficiency (Wipf et al., 2009, 2010). Experiments with preliminary simulated and real data, presented in Wipf et al. (2009, 2010), show that compared to other commonly-used source localization algorithms, Champagne is more robust to correlated sources and noisy data. However, the full extent of the capabilities of Champagne have not been tested with MEG and EEG data.

In this paper, we present results from both simulated and real MEG data that more rigorously characterize the performance of Champagne and better delineate both its strengths and limits. Ultimately an exhaustive performance analysis of a new source localization algorithm serves two distinct purposes; the first being to motivate the use of the algorithm by showing head-to-head performance with commonly-used methods. The second is to fully describe the conditions under which an algorithm is able to provide accurate localizations and of equal importance, when its performance breaks down. The conditions under which an algorithm fails motivates further development of source localization algorithms to advance beyond these shortcomings.

For the simulated data, we have selected a number of challenging source configurations: a large number of distinct sources, deep sources, and clusters of sources. In addition to investigating these challenging source configurations with simulated data, we also present an analysis of the effect of lead field modeling errors. We present the performance of Champagne on EEG data and various EEG montages. We use real data sets to test the performance of Champagne: auditory-evoked field (AEF), audio-visual task, and face-processing data, both with MEG and with EEG recordings. We compare the results obtained by Champagne to those from three benchmark algorithms, MVAB, sLORETA/dSPM, and MCE, on simulated and real data. Taken together, the results on simulated and real data give a rather complete picture of the source localization capabilities of this novel algorithm, Champagne.

Theory

The Champagne algorithm

This section briefly describes the Champagne algorithm and provides the necessary update rules to implement the method.

Champagne relies on segmenting the data into pre- and post-stimulus periods, learning the statistics of the background activity from the pre-stimulus period, and then applying the statistics of the background activity to the post-stimulus data to uncover the stimulus-evoked activity.

We model post-stimulus sensor data (B_{post}) as:

$$B_{post} = \sum_{r=1}^{d_s} L_r S_r + \mathcal{E}, \quad (1)$$

where $B_{post} \in \mathbb{R}^{d_b \times d_t}$, where d_b equals the number of sensors and d_t is the number of time points at which measurements are made in the post-stimulus period. $L_r \in \mathbb{R}^{d_b \times d_c}$ is the lead field matrix in d_c orientations for the r -th voxel. Each unknown source $S_r \in \mathbb{R}^{d_c \times d_t}$ is a d_c -dimensional neural current-dipole source at d_t time points, projecting from the r -th voxel. There are d_s voxels under consideration. $\mathcal{E} \in \mathbb{R}^{d_b \times d_t}$ is a *noise plus interference* factor that is learned from the pre-stimulus period using Stimulus-Evoked Factor Analysis, SEFA, (Nagarajan et al., 2007), which assumes that the columns (corresponding to sensors) are drawn independently from $\mathcal{N}(0, \Sigma_{\mathcal{E}})$. Thus the noise is uncorrelated over time. Learning the statistics of \mathcal{E} is the first step to performing source localization with the Champagne algorithm.

The second step to the source localization process is to estimate hyperparameters Γ that govern the statistical model of the post-stimulus data. We can fully define the probability distribution of the data conditioned on the sources:

$$p(B_{post} | S) \propto \exp\left(-\frac{1}{2} \left\| B_{post} - \sum_{r=1}^{d_s} L_r S_r \right\|_{\Sigma_{\mathcal{E}}^{-1}}^2\right), \quad (2)$$

where $\Sigma_{\mathcal{E}}^{-1}$ is learned using SEFA and in general form, $\|X\|_W$ denotes the weighted matrix norm $\sqrt{\text{trace}[XWX^T]}$.

We assume the following for the source prior on S :

$$p(S | \Gamma) \propto \exp\left(-\frac{1}{2} \text{trace}\left[\sum_{r=1}^{d_s} S_r^T \Gamma_r^{-1} S_r\right]\right). \quad (3)$$

This is equivalent to applying independently, at each time point, a zero-mean Gaussian distribution with covariance Γ_r to each source S_r . We define Γ to be the $d_s d_c \times d_s d_c$ block-diagonal matrix formed by ordering each Γ_r along the diagonal of an otherwise zero-valued matrix. If the lead field has only one orientation (scalar/orientation-constrained lead field), Γ reduces to a diagonal matrix. Since Γ is unknown, we can find a suitable approximation $\hat{\Gamma} \approx \Gamma$ by integrating out the sources S of the joint distribution $p(S, B_{post} | \Gamma) \propto p(B_{post} | S)p(S | \Gamma)$ and then minimizing the cost function:

$$\mathcal{L}(\Gamma) \triangleq -2 \log p(B_{post} | \Gamma) \equiv \text{trace}[C_b \Sigma_b^{-1}] + \log |\Sigma_b|, \quad (4)$$

where $C_b = d_t^{-1} B_{post} B_{post}^T$ is the empirical covariance and Σ_b is the data model covariance, $\Sigma_b = \Sigma_{\mathcal{E}} + L \Gamma L^T$.

Although automatic-relevance determination (ARD) (Mackay, 1992; Wipf and Nagarajan, 2008) is often cited as the origin for sparsity in this class of models (Friston et al., 2008; Henson et al., 2010), a simpler explanation can be seen by inspecting (4). The objective function consists of two terms, one proportional to Σ_b^{-1} and the other to Σ_b . These terms balance each other, and the function is best minimized by setting the majority of the hyperparameters Γ to zero (or close to zero). As such, only a small fraction of the hyperparameters (Γ) will have values significantly different than zero.

Minimizing the cost function (4) with respect to Γ can be done in a variety of ways, including gradient descent or the expectation-maximization algorithm, but these and other generic methods are exceedingly slow when d_s is large. Instead, we utilize an alternative

optimization procedure that uses convex bounding techniques (Boyd and Vandenberghe, 2004). This method expands upon ideas from Jaakkola (2000); Wipf et al. (2007), handles arbitrary/unknown dipole-source orientations, and converges quickly because it provides a better upper bound (Wipf et al., 2011). This optimization procedure yields a modified cost function:

$$\mathcal{L}(\Gamma, X, Z) = \left\| \bar{B} - \sum_{r=1}^{d_s} L_r X_r \right\|_{\Sigma_\varepsilon^{-1}}^2 + \sum_{r=1}^{d_s} \left[\|X_r\|_{\Gamma_r^{-1}}^2 + \text{trace}(Z_r^T \Gamma_r) \right] - h^*(Z), \quad (5)$$

where by construction $\mathcal{L}(\Gamma) = \min_X \min_Z \mathcal{L}(\Gamma, X, Z)$, the matrix $\bar{B} \in \mathbb{R}^{d_b \times \text{rank}(B_{\text{post}})}$, $\bar{B}\bar{B}^T = C_b$ and X and Z are auxiliary variables.

Minimizing this modified cost function yields three update rules, fully derived in Wipf et al. (2010):

$$X_r^{\text{new}} \rightarrow \Gamma_r L_r^T \Sigma_b^{-1} \bar{B} \quad (6)$$

$$Z_r^{\text{new}} \rightarrow \nabla_{\Gamma_r} \log |\Sigma_b| = L_r^T \Sigma_b^{-1} L_r \quad (7)$$

$$\Gamma_r^{\text{new}} \rightarrow Z_r^{-1/2} \left(Z_r^{1/2} X_r X_r^T Z_r^{1/2} \right)^{1/2} Z_r^{-1/2}, \quad (8)$$

where Γ_r comprise the blocks of the block-diagonal matrix of hyperparameters Γ .

In summary, the Champagne algorithm estimates Γ by iterating between (6), (7), and (8), and with each pass we are theoretically guaranteed to reduce (or leave unchanged) $\mathcal{L}(\Gamma)$.

Each Γ_r was initialized with an identity matrix plus/minus a small random number, $O(1e-5)$, different for each element of each Γ_r . We found that this was the most robust initialization, as opposed to initializing with the source power of another algorithm, such as MVAB. When using a vector lead field, as opposed to a scalar/orientation-constrained lead field, a $d_c \times d_c$ covariance matrix is learned for each source. This covariance can be thought of as describing a noisy or unfixed source orientation.

The source time courses can be calculated using the posterior source distribution $p(S|B_{\text{post}}, \Gamma) \propto p(B_{\text{post}}|S)p(S|\Gamma)$, which is Gaussian. To estimate the source time courses, we choose the source posterior mean, i.e. the mean of $p(S|B_{\text{post}}, \Gamma)$, given by:

$$\hat{s}_r(t) = \Gamma_r L_r^T \left(\Sigma_\varepsilon + L L^T \right)^{-1} B_{\text{post}}(t), \quad (9)$$

where $\hat{s}_r(t) \in \mathbb{R}^{d_c \times 1}$ (a short vector across lead field components).

Benchmark source localization algorithms

We chose to test the Champagne algorithm against four representative source localization algorithms from the literature: an adaptive spatial filtering method, MVAB (Sekihara and Nagarajan, 2008), two non-adaptive spatial filtering methods, sLORETA and dSPM, and a version of MCE. The sLORETA and dSPM algorithms are variants of MNE. Details about the implementation of each of the algorithms can be found in Appendix A. Our experience has shown that sLORETA and dSPM, although differing in their theoretical properties, perform similarly with simulated and experimental data. Thus, we present their results together and denote their results with SL/dSPM.

Methods

Quantifying performance

In order to quantify performance on simulated data, we used two features: localization accuracy and time course estimation accuracy. First, we determined whether sources were correctly localized and

then we determined if the source time courses were accurately reconstructed for those source locations. There are a variety of methods to determine detection accuracy from signal detection theory. Typically these methods take into account the occurrence of both hits and false positives and a ROC (receiver-operator characteristic) curve can be obtained for pairs of numbers of hits and false positives. In our simulations, there is more than one simultaneously active brain area (or source), so the use of the free-response ROC (FROC) curve is applicable as it allows for multiple hits and false positives in a single image (Darvas et al., 2004). The A' metric (Snodgrass and Corwin, 1988) estimates the area under the FROC curve for one hit rate (H_R) and false-positive rate (F_R) pair, or in our case, for each simulation. (If the area under the FROC curve is large, then the hit rate is high compared to the false positive rate.)

To determine the accuracy of the time courses, we calculated the correlation coefficient between the seed and estimated source time courses for each hit. Then, we averaged the correlation coefficients for all the hits for each simulation run; this average correlation coefficient is denoted as \bar{R} . Detailed equations for H_R , A' , and \bar{R} are provided in Appendix B.

We developed a metric that captures both the accuracy of the location and the time courses of the algorithms, which we call the *Aggregate Performance* (AP). It combines the A' , \bar{R} , and H_R in the following equation:

$$AP = \frac{1}{2} \left(A' + H_R \bar{R} \right) \quad (10)$$

We use the H_R as a weight for \bar{R} since we compute only the correlation coefficient for the sources that are correctly localized. AP ranges from 0 to 1. For this paper, we use an AP value of 0.75 as the cutoff for a successful localization.

MEG simulations

The simulated data in this paper was generated by simulating dipole sources, with either fixed or variable orientation. We seeded the voxel locations with damped sinusoidal time courses and then projected the voxel activity to the sensors with the lead field. The brain volume was segmented into 8 mm voxels and a two-orientation ($d_c = 2$) forward lead field was calculated using a single spherical-shell model (Sarvas, 1987) implemented in NUTMEG (Dalal et al., 2004) unless where otherwise noted. The parcellation at this resolution yields approximately 4500 voxels (4578 voxels for the simulations in this paper). The time course was partitioned into pre- and post-stimulus periods. The pre-stimulus period (270 samples) contained only noise and interfering brain activity. The post-stimulus period (450 samples), the activity of interest, or the stimulus-evoked activity, was superimposed on the noise and interference present in the pre-stimulus period. The noise plus interference activity consisted of the resting-state sensor recordings collected from a human subject presumed to have only spontaneous activity and sensor noise. Each source location was seeded with a distinct time course of activity and the sources were only present for half of the post-stimulus period (225 samples). The voxel activity was projected to the sensors through the lead field and the noise was added to achieve a desired signal to noise ratio. The simulated data had 275 sensor recordings.

The specific difficult configurations we tested are:

1. Discriminating two dipoles – We examined the minimum distance at which two sources can be resolved. The spacing between voxels on our grid was 8 mm, thus we tested the localization accuracy when the distance between the two sources (inter-dipole distance) was 16, 24, 32, or 40 mm. The locations of the two sources were chosen randomly with the constraint that the minimum distance from the center of the head was 35 mm (since deeper

- sources are harder to localize). (The maximum distance from the center of the head in the volume of interest (VOI) is 65 mm.)
2. Detecting multiple sources – In order to examine the ability of Champagne to localize many individual sources, we performed extensive simulations with randomly seeded sources. We used a volumetric (two-component) lead field computed in NUTMEG (as described above). We randomly seeded 3 to 30 sources throughout the brain. The locations for the sources were chosen so that there was some minimum distance between sources (at least 10 mm) and a minimum distance from the center of the head (at least 35 mm).
 3. Rotating versus fixed dipole model – Altering the correlation between the lead field components, α_{intra} , changes the degree that the orientation of a source rotates (see Appendix for more explanation). We wanted to investigate the effect of the α_{intra} on the ability of Champagne to localize 10, 15, and 20 sources; we chose intra-dipole correlations of 0.25, 0.75, and 1. If the orientation was truly fixed, $\alpha_{intra} = 1$, the most probable orientation for the sources would be normal to the cortical surface (the pyramidal cells in the cortex that give rise to the majority of the MEG signal are mostly orientated perpendicular to the surface). In addition to testing the vector lead field with varying intra-dipole correlations (as described above), we also tested the scalar lead field when the intra-dipole correlation was 0.25 and 1. In all five cases a vector lead field was used for the forward model and we changed the lead field used for solving the inverse problem. Instead of using a scalar lead field from another software package, which would have a different voxel grid than the vector lead field from NUTMEG, we transformed the vector lead field into a scalar lead field by assuming an optimal orientation for every voxel (Sekihara and Nagarajan, 2008). The lead field used for the inverse is referred to as the *inverse model*.
 4. Effect of cortical surface orientation errors – We extended the previous experiment to investigate the effect of orientation errors on the results obtained with the scalar lead field. When using a scalar lead field, an assumption is being made about the orientation of the sources in the brain. To test the effect of errors between the true and assumed orientation, we randomly rotated the orientation of each voxel after simulating the data and before performing source localization with Champagne. The maximum perturbation to the orientation ranged from 0 to $\frac{\pi}{4}$. The orientation of each voxel was not rotated the same amount, rather every voxel was rotated by a randomly generated angle between zero and the maximum perturbation angle. We chose to run these simulations on 10, 15 and 20 sources since this is the regime for which there is a discrepancy in performance for the scalar and vector lead field. The data were generated to have an SNIR of 10 dB.
 5. Deep sources – To assess the ability of Champagne to localize deep sources, which are typically hard to localize with MEG, we constructed three conditions to compare performance with: 1) no deep sources (or all shallow sources), 2) half deep sources and half shallow sources, 3) all deep sources. A deep source was defined as less than 35 mm from the center of the head and a shallow source was defined as above, at least 35 mm from the center of the head. These conditions were designed to test sensitivity to deep sources when there are only deep sources and when there is a combination of deep and shallow sources. The configurations with only shallow sources were included to give a basis for comparison. The total number of sources was 2, 4, 6 and 10, where each total number of sources had three conditions associated with it. For example, for 4 total sources, there was one condition where there were 4 shallow sources, one condition where there were 2 shallow and 2 deep sources, and one condition where there were 4 deep sources. We chose the maximum number of sources to be 10 because that is the largest number of sources

that Champagne was able to localize in the Detecting multiple sources experiment.

6. Clusters – Given the sparsity of Champagne, we tested its ability to localize distributed activity by simulating clusters of sources. We seeded 5, 10, or 15 clusters each with 10 sources in each cluster. These clusters correspond to 50, 100, and 150 voxels having non-zero activity. The placement of the cluster center was random and the clusters consisted of sources seeded in the 9 nearest neighboring voxels. The source time courses within each cluster had an inter-dipole correlation coefficient of 0.8 and an intra-dipole correlation coefficient of 0.25. The multiple clusters were correlated with a correlation coefficient of 0.5. We made the correlations within the clusters higher than between clusters because nearby voxels are more plausibly correlated than voxels at a distance. For the clusters, we are both interested in whether the cluster is localized and whether the extent of the cluster is accurately reconstructed. To assess the localization of the clusters, we use the A' metric. The A' metric is calculated for the clusters by testing if there is a local peak within the known extent of the cluster. To assess the accuracy of the extent of the clusters, we calculate the fraction of the seeded voxels with power in or above the 80th percentile of all the voxels. We call this fraction the *Cluster Extent Score (CES)*.

We could adjust both the signal-to-noise-plus-interference ratio (SNIR), the correlations between the different voxel time courses, inter-dipole correlation (α_{intra}), and the correlation between the orientations of the dipolar sources, intra-dipole correlation (α_{inter}). Details of these parameters are described in Appendix C. We provide two examples of the simulated sensor data at 0 dB and 10 dB in Fig. 1. The red line indicates the start of the post-stimulus period (effectively the stimulus onset) at 0 ms. In this example, there are 5 sources seeded throughout the brain. If not indicated otherwise, each of the experiments above were conducted with the following settings: SNIR = 0, 10 dB, $\alpha_{inter} = 0.5$, and $\alpha_{intra} = 0.25$.

The results we obtained using simulated data are presented in two forms. First, we show the plots of mean A' , \bar{R} , mean H_R , and/or mean AP, our Aggregate Performance metric. For each configuration, the results are averaged over 50 simulations and we have plotted these averaged results with standard error bars.

For some of the experiments, we also show examples of the localization results from single simulations, which complement our aggregate results. We compute the source power at every voxel and project the activity to the surface of a rendered MNI-template brain. These plots contain a projection of the true source power, called *Ground Truth*.

EEG simulations

We tested Champagne on simulated EEG data using a scalar, cortically-constrained lead field computed in SPM (<http://www.fil.ion.ucl.ac.uk/spm>) by selecting the coarse resolution. This resulted in approximately 5000 voxels at approximately 5 mm spacing. The EEG data was simulated in the same way as the MEG data, described above, only there was no α_{intra} and we repeated the Detecting multiple sources experiment with this lead field. In addition to running on the maximum number of sensors (128), we also investigated the effect of subsampling the number of sensors on the ability to localize sources. We subsampled the 128 sensor lead field to 64, 32, and 16 sensors, while preserving coverage, because many EEG researchers use as few as 10 electrodes for a standard 10–20 montage or 16 sensors for clinical EEG systems.

Real data

All MEG data was acquired in the Biomagnetic Imaging Laboratory at UCSF with a 275-channel CTF Omega 2000 whole-head MEG

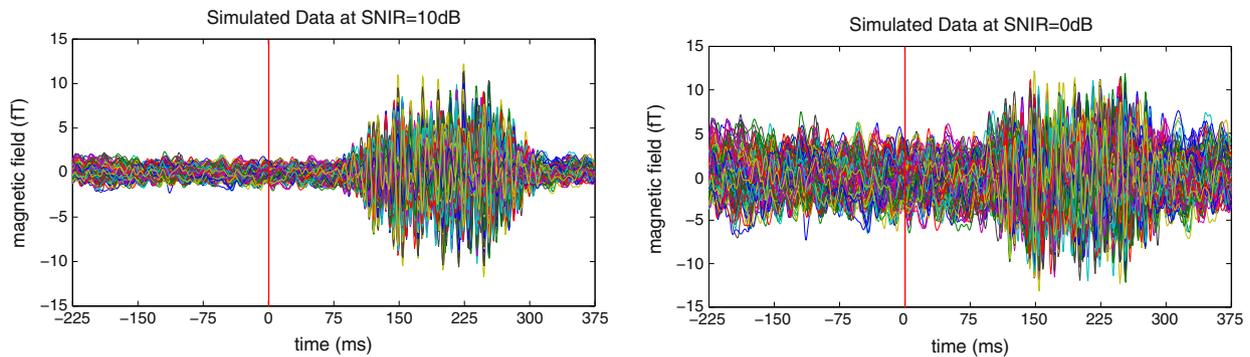


Fig. 1. Simulated data at 10 dB and 0 dB. The red line indicates the start of the post-stimulus period (or the stimulus onset) at 0 ms. There are 5 sources seeded in this example and the noise is real resting-state data.

system from VSM MedTech (Coquitlam, BC, Canada) with a 1200 Hz sampling rate. The lead field for each subject was calculated in NUT-MEG (Dalal et al., 2004) using a single-sphere head model (two-orientation lead field) and an 8 mm voxel grid. The data was digitally filtered from 1 to 160 Hz to remove artifacts and the DC offset was removed.

We ran Champagne and the benchmark algorithms on three real MEG data sets:

1. Auditory evoked field (AEF) – The neural responses of seven subjects to an AEF stimulus were localized. The AEF response was elicited with single 600 ms duration tones (1 kHz) presented binaurally. The data were averaged across 120 trials (after the trials were time-aligned to the stimulus). The pre-stimulus window was selected to be -100 ms to 5 ms and the post-stimulus time window was selected to be 5 ms to 250 ms, where 0 ms is the onset of the tone.
2. Audio-visual – We analyzed a data set designed to examine the integration of auditory and visual information. We presented single 35 ms duration tones (1 kHz) simultaneous with a visual stimulus. The visual stimulus consisted of a white cross at the center of a black monitor screen. The data were averaged across 100 trials (after the trials were time-aligned to the stimulus). The pre-stimulus window was selected to be -100 ms to 5 ms and the post-stimulus time window was selected to be 5 ms to 450 ms, where 0 ms is the onset of the simultaneous auditory and visual stimulation.
3. Face processing – A MEG data set from a subject for which faces and scrambled faces were presented in a random order with an interstimulus interval of 1 s was used as the third data set. For the face stimulus, the pre-stimulus window was -200 ms to 5 ms and the post-stimulus time window was from 5 ms to 450 ms, where 0 ms is the appearance of the visual stimulus. In addition to running Champagne on the face processing data, we also constructed a contrast dataset to localize the brain areas more active when the faces were presented than when scrambled faces were presented. The pre-stimulus period of this data set consisted of the trial-averaged post-stimulus period (5 ms to 450 ms) from the scrambled-face data and the post-stimulus period consisted of the trial-averaged post-stimulus period (5 ms to 450 ms) from the face data. There were 40 trials of both the face and scrambled data; the data (for each condition) were averaged (after the trials were time-aligned to the stimulus).

For the real MEG data plots, we show coronal sections through the voxel with the maximum power. If the coronal section does not adequately demonstrate the localization, we included the sagittal slice. The crosshairs on the MRI images show the location of the voxel whose time course is plotted.

The EEG data (128-channel ActiveTwo system) was downloaded from the SPM website (<http://www.fil.ion.ucl.ac.uk/spm/data/mmfaces>) and the lead field was calculated in SPM8 using the coarse resolution. The EEG data paradigm involved randomized presentation of at least 86 faces and 86 scrambled faces, although we did not use the scrambled face data. We averaged the time-aligned trials to the presentation of the face and created an averaged data set. The pre-stimulus window was selected to be -200 ms to 5 ms and the post-stimulus time window was selected to be 5 ms to 250 ms. For this real data set we found that using the three-component (vector) lead field in SPM was more robust than the orientation-constrained lead field (we selected the coarse tessellation for our grid resolution). The power was plotted on a 3-D brain and the time courses for the peak voxels are plotted (the arrows point from a particular voxel to its time course).

Computational considerations

The Champagne algorithm typically converged in 75 to 100 iterations. The lead field we used for the majority of the simulations had approximately 5500 voxels, which results in the estimation of approximately 16,500 hyperparameters, and took approximately 10 min to run on a lead field of that size (and trial-averaged data) on an Intel(R) Core(TM)2 Quad CPU @ 3.00 GHz, with 8 GB of memory. The number of hyperparameters Champagne estimates is dependent on the number of lead field orientations. For a scalar (or fixed-orientation) lead field, d_s hyperparameters are estimated. For a 2 or 3 component lead field, $3 \times d_s$ or $6 \times d_s$ unknown hyperparameters are estimated, respectively. (Each precision matrix Γ_r is a symmetric matrix.)

Results

MEG simulations

Discriminating two dipoles

The first column of Fig. 2 shows the hit rate (H_R) plotted against inter-dipolar distance at 10 dB and 0 dB. The H_R results show that at 10 dB Champagne is able to localize two sources at any inter-dipole distance and at 0 dB Champagne is able to distinguish two sources if there is 32 mm (or 3 voxels width) between the sources. The benchmark algorithms do not perform as well. A' , in the second column, echoes the H_R plots across SNIR and algorithm.

Champagne is able to reconstruct the time courses at all inter-dipole distances at 10 dB and at 0 dB, Champagne is successful when there is at least 32 mm between the sources, while the other algorithms have more difficulty reconstructing the time courses at both SNIR levels and all inter-dipole distances.

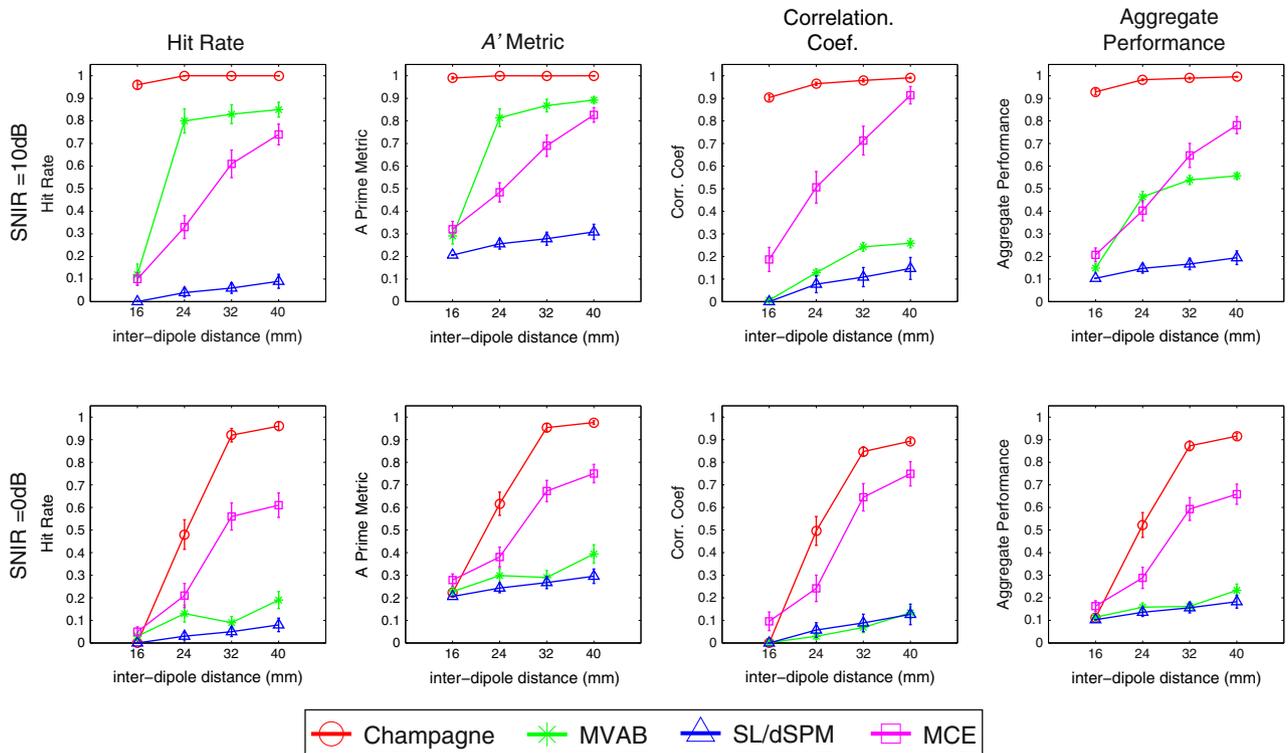


Fig. 2. Simulations with two dipoles seeded with inter-dipole distances of: 16 mm, 24 mm, 32 mm, or 40 mm. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. Components: (a) A' metric, (b) hit rate (H_R) and (c) average time-course correlation coefficient (\bar{R}) used to compute (d) Aggregate Performance metric (AP). The following equation is used: $AP = \frac{1}{2}(A' + H_R\bar{R})$. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

The three values presented in the first three columns of Fig. 2, H_R , A' , and \bar{R} , are combined to compute Aggregate Performance (AP) shown in the fourth column. For Champagne, the AP results look very similar to the plots in the previous three columns since Champagne tended to perform well across metrics (or perform poorly as is the case at 0 dB and the smaller inter-dipole distances). Likewise, MCE shows similar performance across the first three metrics (or columns) and yields similar trends in the AP metric. MCE performs most similarly to Champagne at large inter-dipole distances and performs like MVAB at smaller inter-dipole distances at 10 dB. At 0 dB, however, Champagne and MCE have similar trends across the inter-dipole distances, but Champagne performs better than MCE at the lower SNIR level. The AP results for MVAB at 10 dB demonstrate that MVAB was penalized for having poor time course reconstruction in spite of having good localization results. SL/dSPM does poorly across all of the first three metrics at both noise levels and thus, does poorly in terms of the AP metric.

Detecting multiple sources

In Fig. 3, we plot number of sources versus AP at SNIR levels of 10 dB and 0 dB (inset plot). At both SNIR levels, Champagne outperforms the other source localization algorithms. Champagne is able to accurately reconstruct up to 10 sources at 10 dB and up to 5 sources at 0 dB. For both SNIR levels and across number of sources, the three other source localization algorithms perform at almost the same level.

We also show single simulation examples at 10 dB and 0 dB for both 5 (Fig. 4) and 10 sources (Fig. 5). The 5 source examples demonstrate that Champagne is able to recover 5 sources at both SNIR levels. Both MCE and MVAB are able to recover most of the 5 sources, but also have some false positives and blur around the sources. At 0 dB, MVAB and MCE do not successfully localize the

sources, there is a peak at the center of the head (which gets projected to the surface), and one successful source for MVAB. SL/dSPM did not localize any of the 5 sources at either SNIR level. These results show a peak at the center of the head and while this peak does extend out to where the sources were seeded, there are no distinct peaks at the seeded locations. The 10 source examples demonstrate that Champagne is able to recover 9 out of 10 sources

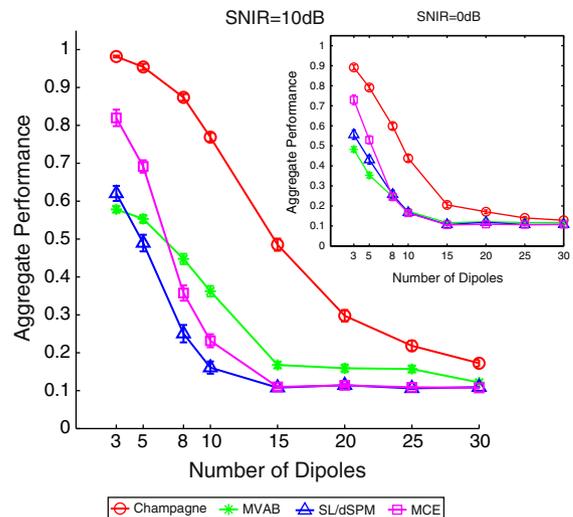


Fig. 3. Results from simulations to test localization on multiple dipoles (3 to 30) with a volumetric lead field (2 components). AP is plotted against number of dipoles for SNIR = 10 dB and SNIR = 0 dB (inset plot). The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

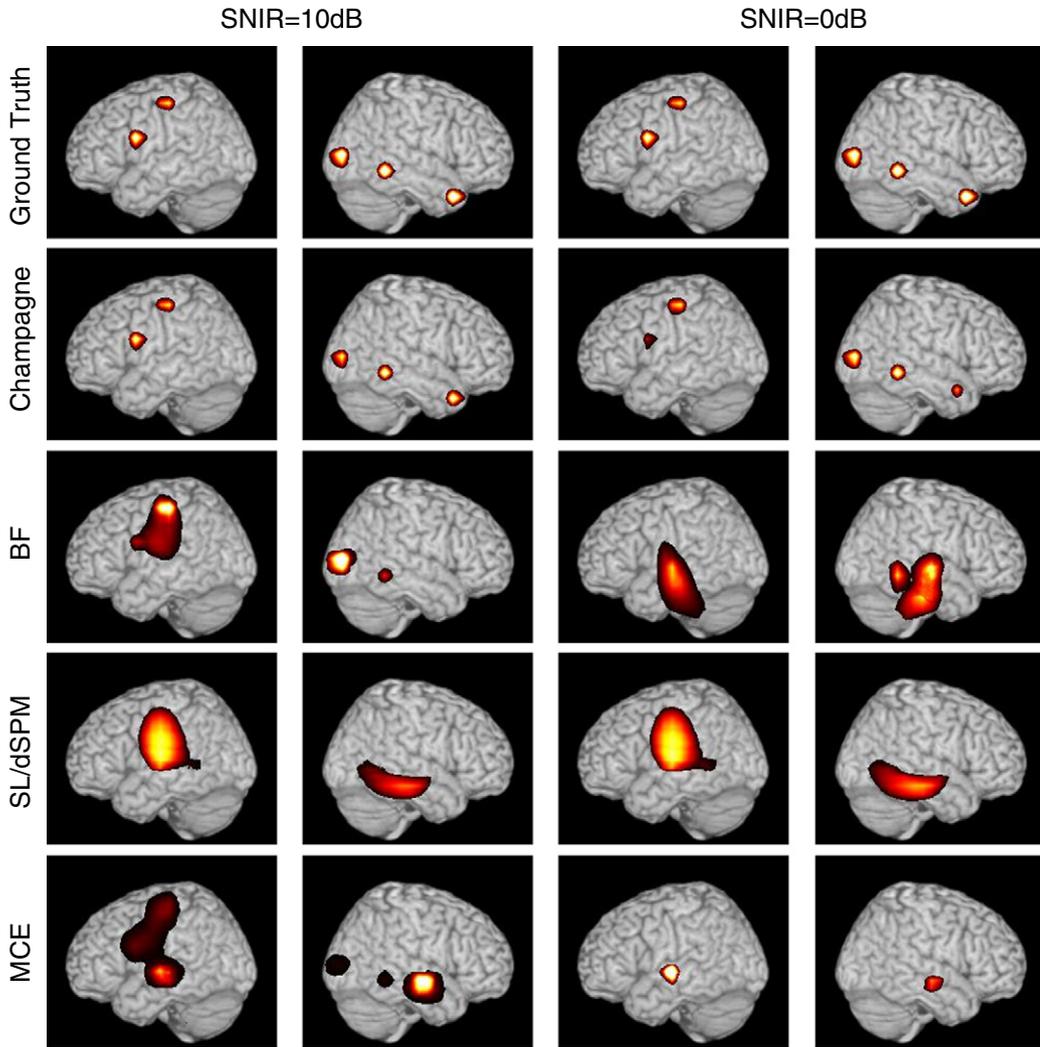


Fig. 4. A single example of the localization results with the vector lead field for 5 dipoles at SNIR = 10 dB (right columns) and SNIR = 0 dB (left columns). The ground truth location of sources are shown for comparison. We project the source power to the surface of a template brain.

at a SNIR of 10 dB. At 0 dB, Champagne is able to recover over half the sources, but also has some false positives, including a peak at the center of the head. Both MVAB and MCE are able to localize some of the 10 sources at 10 dB, but at 0 dB these algorithms are unable to resolve more than one source. As in the 5 source example, SL/dSPM is not able to recover any of the 10 sources at 10 dB or 0 dB; there is only a peak at the center of the head.

Rotating versus fixed dipole model

In Fig. 6, AP is plotted against five conditions: 1) $\alpha_{intra}=1$ and vector inverse model, 2) $\alpha_{intra}=0.75$ and vector inverse model, 3) $\alpha_{intra}=0.25$ and vector inverse model, 4) $\alpha_{intra}=1$ and scalar inverse model, and 5) $\alpha_{intra}=0.25$ and scalar inverse model for 10, 15, and 20 sources. The performance using a vector inverse model (the first three conditions) demonstrates that the presence of intra-dipole correlations makes the localization problem more difficult and that the weaker the correlation, the easier it is to localize multiple sources. The results using a scalar inverse model show that if the sources' orientations are truly fixed, it is advantageous to use a scalar lead field for the inverse problem. The performance on 15 and 20 sources is better with the scalar lead field than the performance obtained using the vector lead field if the orientation is fixed. On the other hand, if the orientation is rotating and a scalar inverse model is used for the reconstruction the performance is

drastically reduced as compared to the experiments where a vector inverse model is used.

Effect of cortical surface orientation errors

In Fig. 7, AP is plotted against the maximum orientation error for 10, 15, and 20 sources. The trend in performance across the perturbations is consistent for 10, 15 and 20 sources. The performance drops by approximately 10% immediately with the smallest maximum perturbation, $\frac{\pi}{64}$. The performance stays the same for maximum perturbations of $\frac{\pi}{32}$ and $\frac{\pi}{16}$, and then starts falling off more rapidly for maximum perturbations of $\frac{\pi}{8}$ and $\frac{\pi}{4}$.

We conducted an additional experiment to assess the effect of lead-field errors that may occur due to errors in coregistration. We used a vector lead field to generate data at 10 dB and then we solved the inverse problem with Champagne with a lead field shifted diagonally the distance of about half the diagonal of a voxel (approximately 5 mm). We tested Champagne with 3 to 30 sources and found that the AP metric was reduced, at worst, 10% for each number of sources (figure not shown).

Deep sources

The results on deep sources are shown in Fig. 8. We have plotted AP against total number of sources (2, 4, 6, or 10) for the three source configurations described above (rows of Fig. 8), at 0 and 10 dB

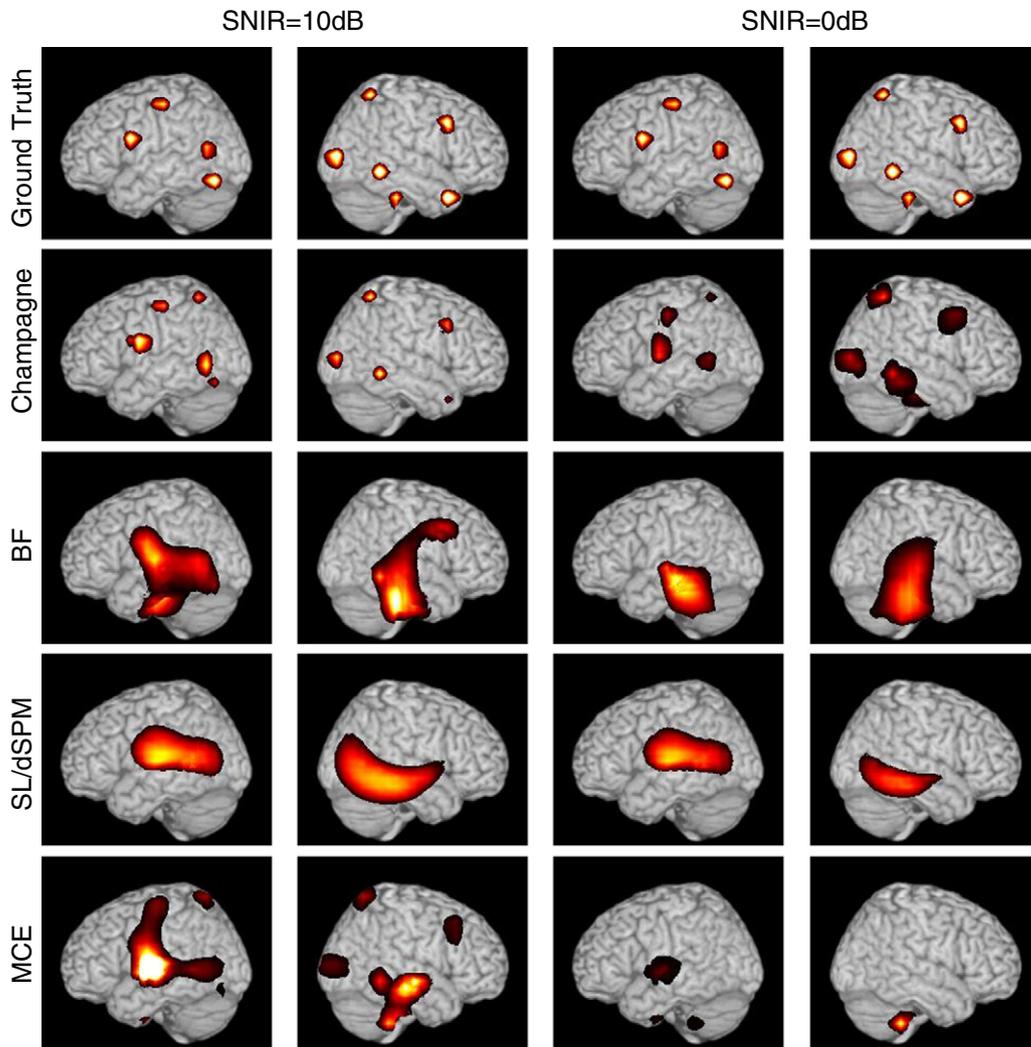


Fig. 5. A single example of the localization results with the vector lead field for 10 dipoles at SNIR = 10 dB (right columns) and SNIR = 0 dB (left columns). The ground truth location of sources are shown for comparison. We project the source power to the surface of a template brain.

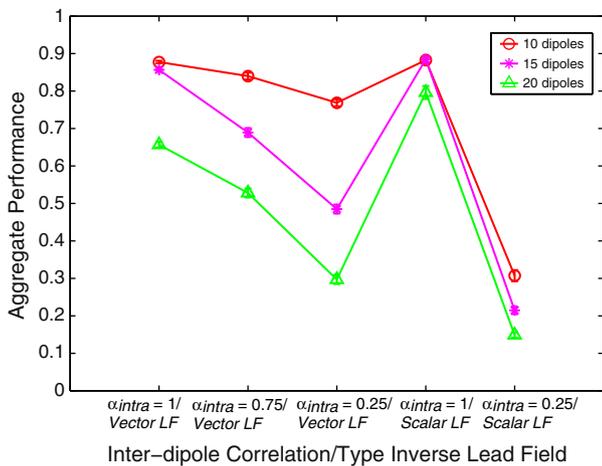


Fig. 6. Exploration of lead field errors with 10, 15, and 20 dipoles (depicted by color) by using one forward model and a different inverse model. AP is plotted against five conditions: 1) $\alpha_{intra}=1$ and vector inverse model, 2) $\alpha_{intra}=0.75$ and vector inverse model, 3) $\alpha_{intra}=0.25$ and vector inverse model, 4) $\alpha_{intra}=1$ and scalar inverse model, and 5) $\alpha_{intra}=0.25$ and scalar inverse model. The inter-dipole correlation is 0.5 for all conditions. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

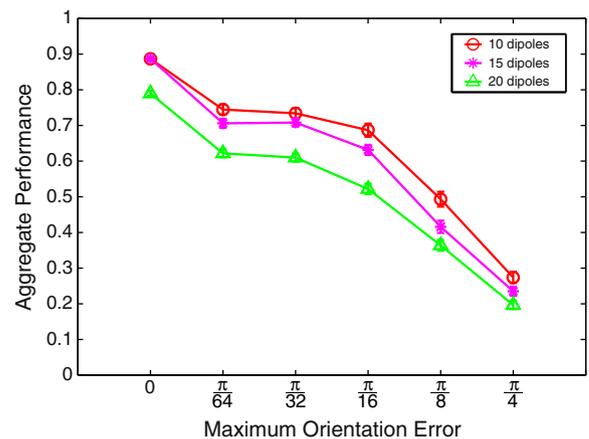


Fig. 7. Localization results for 10, 15, and 20 dipoles (depicted by color) with orientation errors introduced to the scalar lead field. AP is plotted against maximum perturbation angle. The orientation vector of each voxel is perturbed by a random angle from 0 to the maximum angle (shown on the x-axis). The inter-dipole correlation is 0.5 (and the intra-dipole correlation is 0.25 for the vector lead field). The results are averaged over 50 simulations at each data point and the error bars show the standard error.

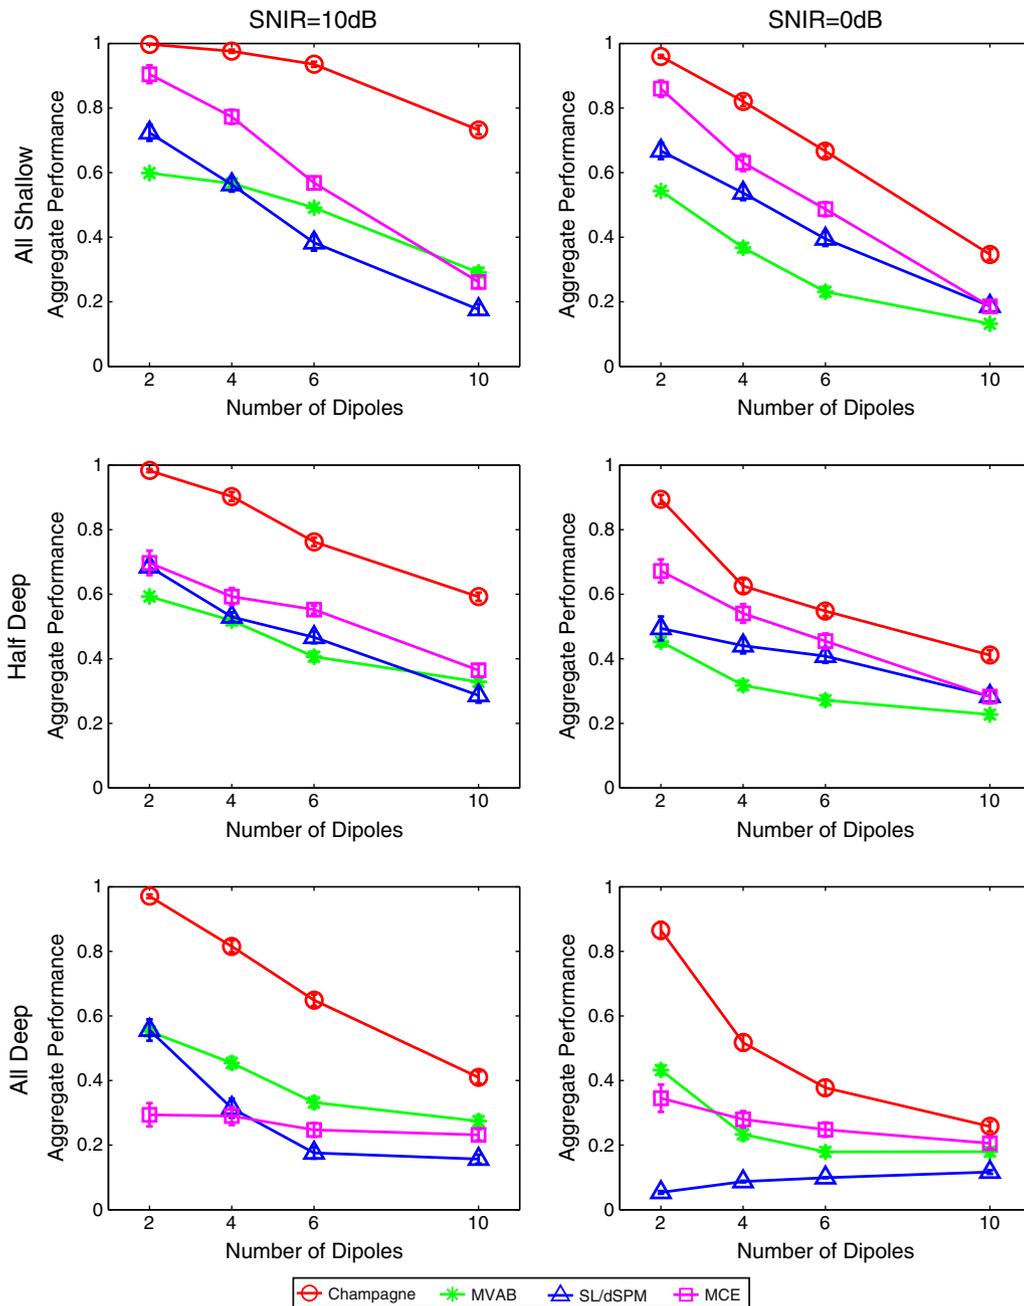


Fig. 8. Localization results in the presence of deep sources at SNIR = 10 dB (left column) and SNIR = 0 dB (right column). AP is plotted against total number of sources in each panel. There were three conditions (corresponding to the rows of the figure): 1) no deep sources (or all shallow sources), 2) half deep sources and half shallow sources, 3) all deep sources. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

(columns of Fig. 8). Across all the source configurations, noise levels, and numbers of sources, Champagne outperforms the other source localization algorithms. As we saw in the *Detecting Multiple Dipoles* experiment, Champagne is able to reconstruct 10 sources at 10 dB and 4 to 6 sources at 0 dB. MCE consistently does better than MVAB and SL/dSPM for the conditions where there are all shallow sources or half shallow (and half deep) sources. The presence of deep sources (whether they constitute half or all the sources) degrades performance for all of the methods. This decrease in performance is especially marked for MCE at both noise levels and SL/dSPM at 0 dB.

Clusters

The results from the experiment with clusters are shown in Fig. 9. In the first row, we show A' metric results for 10 dB and 0 dB. In the

second row, we show the CES results for 10 dB and 0 dB. For 5 clusters, Champagne and MVAB perform equally for both metrics at both SNIR levels, but at 10 and 15 clusters Champagne performs better than MVAB and the other algorithms both in terms of localizing the clusters and reconstructing their extent. MCE does a fairly good job at localizing the clusters at 10 dB as seen in the A' plot, but is not able to reconstruct the extent of the clusters as seen in the CES plot. At 0 dB, MCE and SL/dSPM perform similarly in terms of A' and CES and in general, SL/dSPM is not successful in localizing or reconstructing the extent of the clusters. A single-simulation experiment with 10 clusters is presented in Fig. 10. Champagne is able to localize all 10 clusters in this example. MVAB and MCE are able to localize a fair number of clusters (7 and 5 respectively) and SL/dSPM is only able to localize one cluster (the cluster in the occipital lobe).

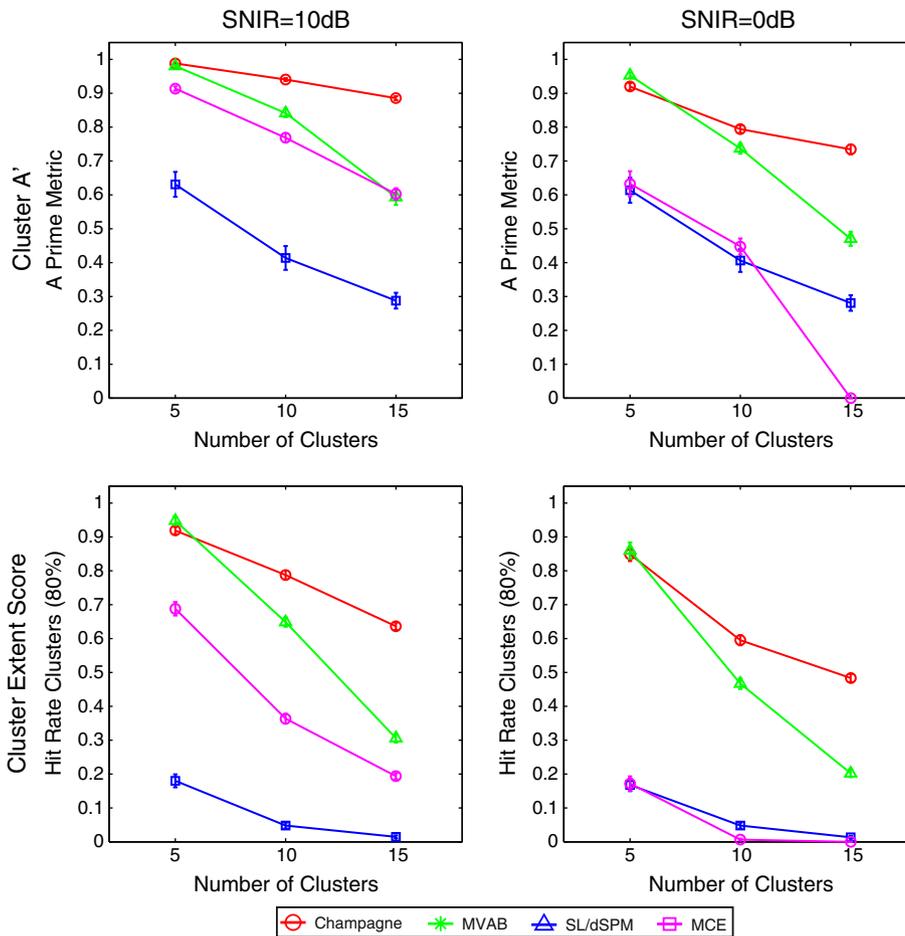


Fig. 9. Performance on 5, 10 and 15 clusters of dipoles (each cluster has 10 contiguous active voxels) using a vector lead field. Cluster A' is plotted against number of clusters in the top row, showing the localization accuracy and Cluster Extent Score (CES) is plotted against number of clusters in the bottom row, showing the algorithms' ability to correctly model the extent of the clusters. The inter-dipole correlation within the cluster is 0.8, the inter-dipole correlation between clusters is 0.5, and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

EEG simulations

In Fig. 11 we present the EEG simulation results. The left column shows the results at 10 dB and the results at 0 dB are shown in the right column. The results obtained with the full 128-sensor lead field are shown in the top row, and the number of sensors in the lead field decrease in the subsequent rows. At almost every SNIR level and number of sources and sensors, Champagne outperforms the other algorithms. In the one-source case, only MCE is able to localize that source better than Champagne with 128 and 64 sensors at both 10 dB and 0 dB. The performance of MCE after one source drops off quickly in these plots. In general, the performance of Champagne is the same for 128 and 64 sensors, we do not see a degradation in performance until 32 sensors. This is true for the other localization algorithms as well.

Performance on real data

Auditory evoked field

The localization results from the AEF data from seven subjects are shown in Fig. 12. In each subplot, we show the power at each voxel in a 50–75 ms window around the M100 peak. Champagne is able to localize bilateral auditory activity for all seven subjects (shown in the first column of Fig. 12). The activity is in Heschel's gyrus, which is the location of primary auditory cortex. SL/dSPM is able to localize bilateral auditory activity in five of the seven subjects. In these five subjects, the activations are diffuse and in most cases biased to one side.

MVAB is only able to localize bilateral activity in one subject (Subject 5) and in the other six cases it localizes the activity to the center of the head. MCE is only able to localize the auditory activity on one side in most subjects. In Subjects 4 and 6, MCE is able to localize bilateral activity, but the activity is more lateral than Heschel's gyrus. The MCE algorithm favors voxels on the edge of the voxel grid, and often does not accurately localize cortical areas.

Audio-visual task

The results Champagne yields for the audio-visual task are presented in Fig. 13. In the first and second rows, we show the brain activations associated with the auditory stimulus. Champagne is able to localize bilateral auditory activity in Heschel's gyrus in the window around the M100 peak, shown in Figs. 13(a) and (c). The time courses for the left and right auditory sources are shown in Figs. 13(b) and (d), along with the window used around the M100 peak. The two auditory sources had the maximum power in the window around the M100 peak. Second, we present the early visual response in Figs. 13(e) and (f). Champagne is able to localize a source in medial, occipital gyrus with a peak around 150 ms. We plot the power in the window around this peak in Fig. 13(e) and the time course of the source marked with the crosshairs in Fig. 13(f). Using a later time window, shown in Fig. 13(g), we can localize a later visual response with a time course (Fig. 13(h)) that has power extending past 200 ms.

The benchmark algorithms were run on this data set, the results can be found in the supplementary materials (Figs. S1 to S3). In brief, MVAB and SL/dSPM are able to localize some of the auditory

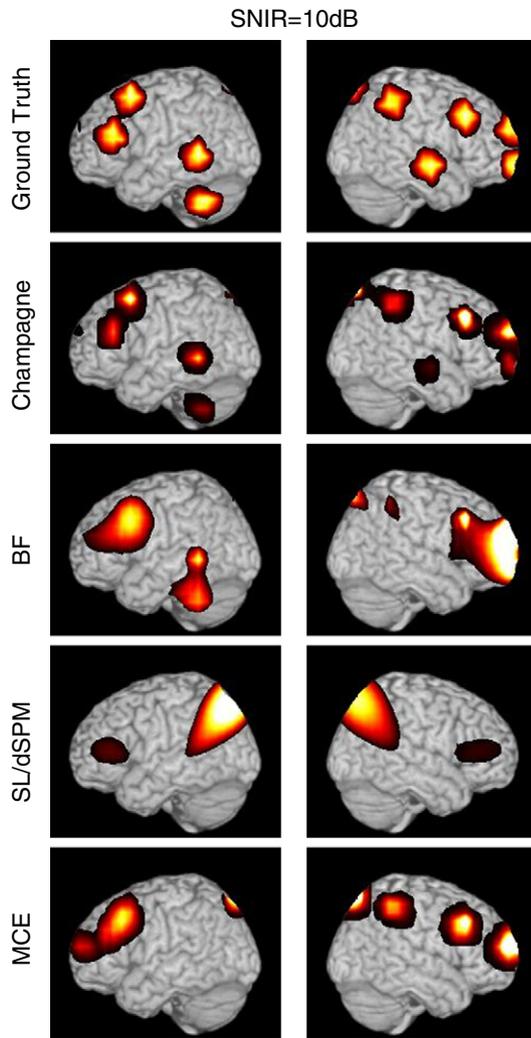


Fig. 10. A single example of the localization results for 10 clusters (each with 10 dipoles) at SNIR = 10 dB with the vector lead field. The ground truth location of the clusters are shown for comparison. We project the source power to the surface of a template brain.

and visual areas found with Champagne, but do not provide focal localizations and the time courses are not as distinct for the auditory and visual sources. MCE provides results similar to Champagne, though it favors the voxels on the periphery of the volume.

Face-processing task: MEG

The face processing results are shown in Figs. 14 and 15. First, we examine the brain areas involved in the processing of the faces. Fig. 14(a) shows an early visual response to the presentation of the face visual stimulus in medial occipital cortex with time course in Fig. 14(b). Fig. 14(c) shows a later visual response more lateral to the early response with the time course in Fig. 14(d). In Figs. 14 (e) and (g) we show a bilateral activation in the fusiform gyrus with time courses shown in Figs. 14 (f) and (h) that show peaks around 170 ms. Second, in the contrast condition, we can see a bilateral activation in the fusiform gyrus Figs. 15 (a) and (c) with a peak around 170 ms (time courses shown in Figs. 15 (b) and (d)). The activations in the fusiform gyrus are the maximum in the window shown in Figs. 15 (b) and (d), around 170 ms. The results concur with those from the literature (Kanwisher et al., 1997). An area in fusiform gyrus (on the ventral surface of the occipital lobe), called the fusiform face area (FFA), has been shown to be an area preferentially activated

by the presentation of faces versus other visual stimuli, such as scrambled faces.

The benchmark algorithms were also applied to this face processing data set; the results can be found in the supplementary materials (Figs. S4 to S6). In brief, SL/dSPM and MCE are able to localize a visual activation and bilateral FFA. MVAB localizes a visual area, but the peak in the time course does not correspond to early visual activity.

Face-processing task: EEG

In Fig. 16, we present the results from using Champagne on the face-processing EEG data set. In Fig. 16(a), we see that Champagne is able to localize early visual areas that have a peak around 100 ms. In Fig. 16 (b), the ventral surface of the brain is shown. There are a few activations in and around fusiform gyrus. These activations are larger in extent and have peaks around 170 ms corresponding to the N170 seen in the sensor data. These results are consistent with those obtained in (Henson et al., 2010) using the same EEG data set.

The benchmark algorithms were applied to the EEG data set; these results can be found in the supplementary materials (Fig. S7). In summary, the benchmark algorithms are able to localize visual and in the case of MCE, fusiform activations. However, the time courses for all the peak locations have a similar shape and cannot be used to differentiate the activations.

Discussion and conclusion

In this paper, we present the results from a rigorous battery of performance tests of Champagne, our novel source-localization algorithm. We have compared the performance of Champagne to commonly-used source localization algorithms, MVAB, SL/dSPM, and MCE. The tests use simulated and real data; the simulated data explores performance with difficult source configurations and the real data demonstrates the ability of Champagne localize real brain activity in the face of source correlations and real noise, interference, and signal-to-noise ratios.

The experiments with simulated data exemplify that Champagne provides robust localization and time course estimation with complex source configurations and noisy, correlated sensor data. Champagne is able to distinguish two neighboring, correlated sources, which could have real implications when analyzing real M/EEG data with brain areas activated in close proximity to one another, such as auditory and secondary somatosensory cortex or visual cortical areas V1/V2/V3/V4.

We also found that Champagne is able to successfully localize up to 10 sources at high SNR with vector fields, and up to 20 sources with scalar lead fields with accurate orientation estimates. However, even slight inaccuracies in orientation can cause performance to degrade with scalar lead fields below the performance of vector lead-fields.

Interestingly, Champagne is able to provide accurate localization of deep sources even when there is a mix of deep and shallow sources and only deep sources. Champagne produces the largest improvement when localizing only deep sources at both noise levels as compared to the other algorithms. Localizing the individual clusters and capturing the cluster extent is also done most accurately by Champagne at both SNIR levels. MVAB does nearly as well as Champagne; the localization of the clusters is most likely aided by the increase in the signal at the cluster locations and accurately modeling the extent of the sources is likely aided by the inherent blur in beamformer solution. MCE does well at localization, but does not accurately model the extent, due to its sparsity profile. Thus, it is notable that Champagne is able to both get the location and the extent of the clusters despite its sparsity profile. Champagne is sparse, but not so sparse that it cannot model a total of 50 to 150 active voxels at one time. (We also tried clusters of smaller and larger source extents, for which the results were similar.)

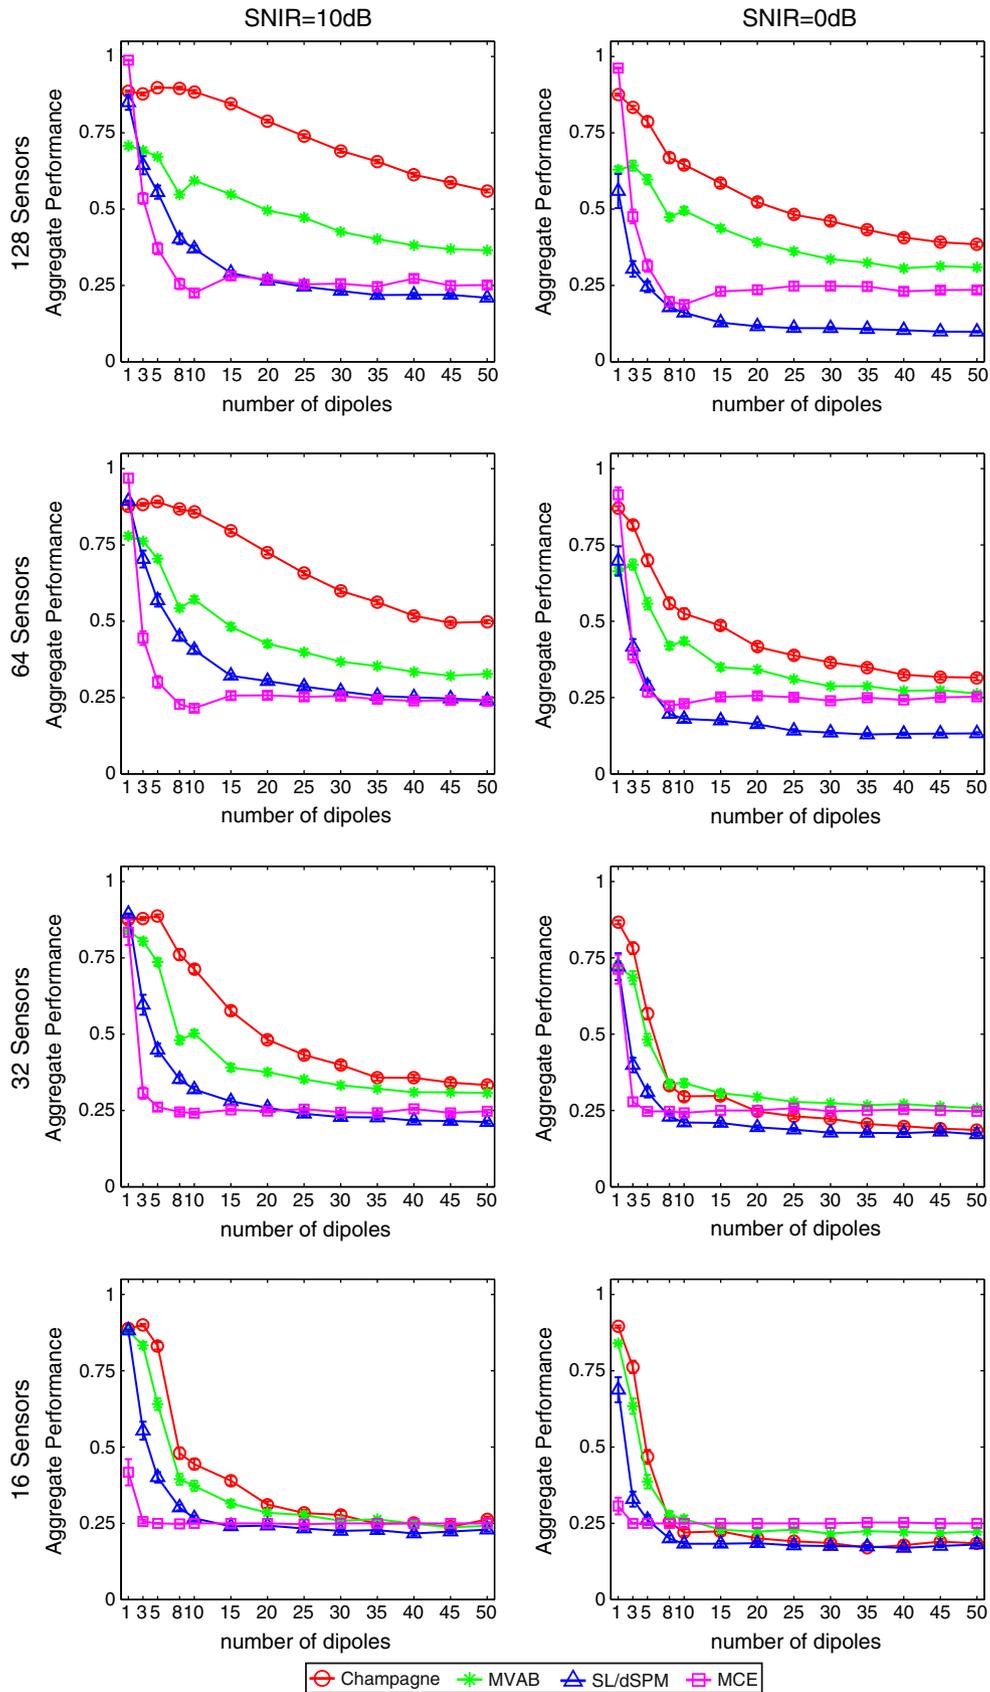


Fig. 11. Results for EEG simulations with a scalar EEG lead field, AP is plotted against number of dipoles (1 to 50). Also shown is the effect of downsampling the number sensors on performance (shown in each row). The inter-dipole correlation between clusters is 0.5. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

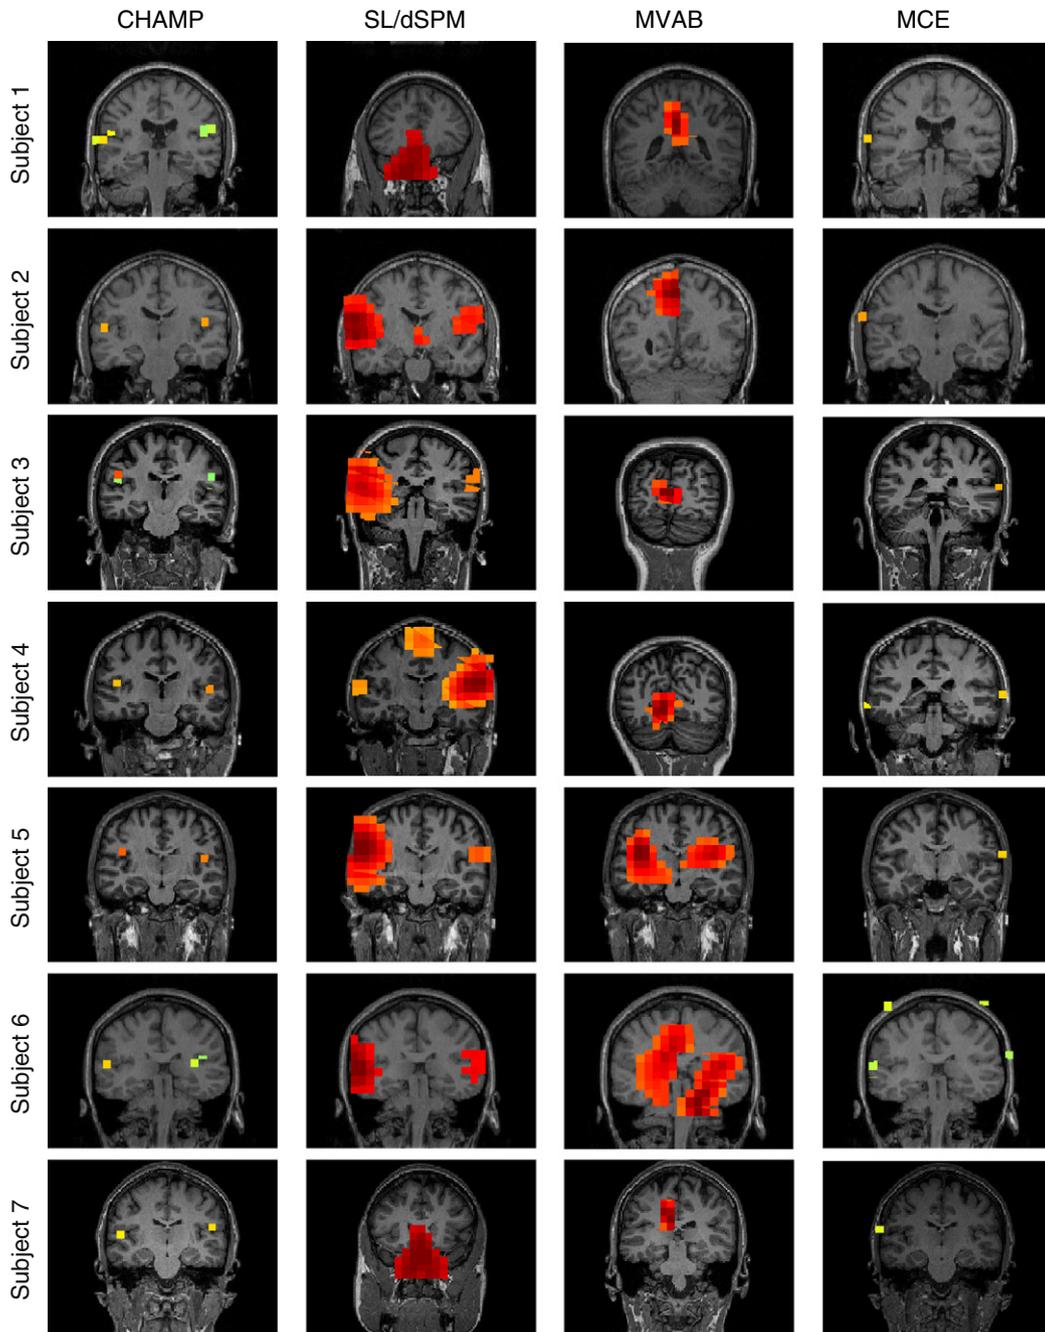


Fig. 12. Auditory evoked field (AEF) results for 7 subjects. The results from using Champagne are shown in the right-most column and the results from using the benchmark algorithms are shown the other three columns.

This paper is the first evaluation of the performance of Champagne on EEG data. We found that Champagne performed similarly on the EEG data with 128 sensors to the MEG results with 275 sensors. The results were not drastically affected by reducing the number of sensors by half, i.e. going down to only 64 sensors. With even fewer sensors, 32 or 16, performance is still superior to benchmarks.

Using simulated data to assess the performance of a source localization algorithm is necessary as one must know the true locations and time courses of the sources in order to definitively recognize a successful localization. But, while simulated data can be generated in a realistic fashion, it is still artificial data. Also, the way in which data is simulated can be biased towards a particular algorithm. The experiments in this paper use simulated data that has been generated in a manner that best replicates key aspects of real M/EEG data. We have also tried to generate data that does not favor Champagne by

making the following choices. We have chosen realistic signal-to-noise ratios and number of time points to best model true experimental data. We use real brain noise instead of Gaussian noise; the noise models in both SEFA and Champagne assume Gaussian distributions. The source time courses are damped sinusoids, while the Champagne model assumes Gaussian time courses. Also, the inter- and intradipole correlations make the source localization problem more difficult, but more closely model the complexity of brain activity.

When simulations are drawn from the generative model for Champagne, i.e. when the source time courses are Gaussian, the SNIR level is 100 dB, and the inter-dipoles correlation is zero, we have found that Champagne is able to localize on the order of 100 sources.

There are aspects of brain activity that we have not explored in this work, such as sources that move over the course of the post-

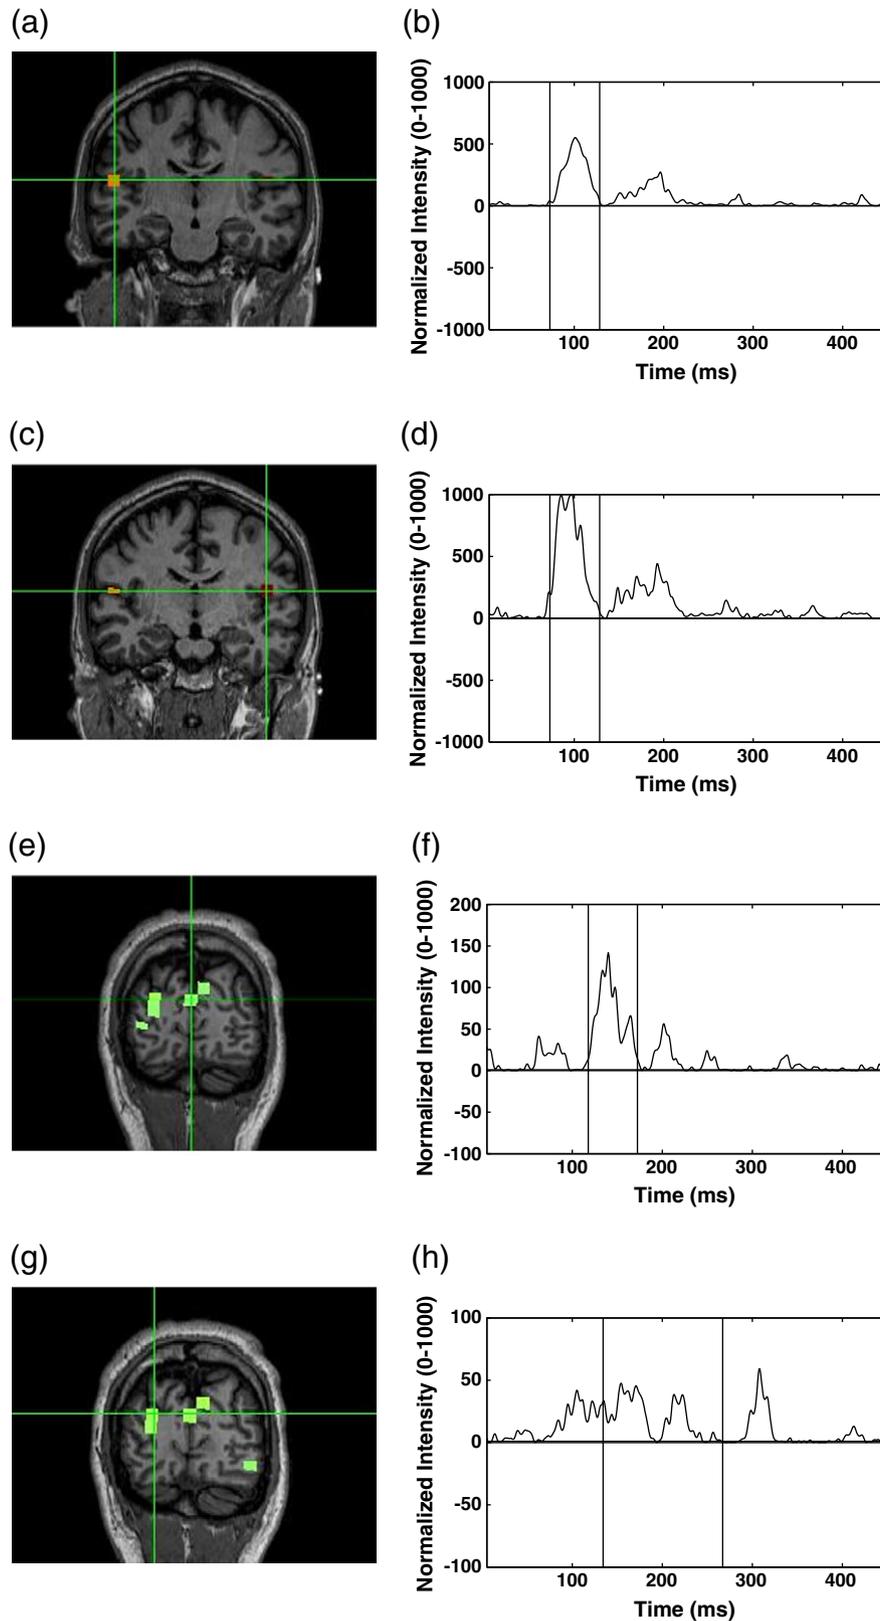


Fig. 13. Audio-visual data localization results from Champagne. Champagne localizes a bilateral auditory response at 100 ms after the simultaneous presentation of tones and a visual stimulus. Champagne localized bilateral auditory activity, (a) and (c), with time courses shown in (b) and (d). Champagne localizes an early visual response at 150 ms after the simultaneous presentation of tones and a visual stimulus. The time course in (f) corresponds to the location indicated by the crosshairs in the coronal sections (e). Champagne localizes a later visual response later than 150 ms after simultaneous presentation of tones and a visual stimulus. The time course in (h) corresponds to the location indicated by the crosshairs in the coronal sections (g).

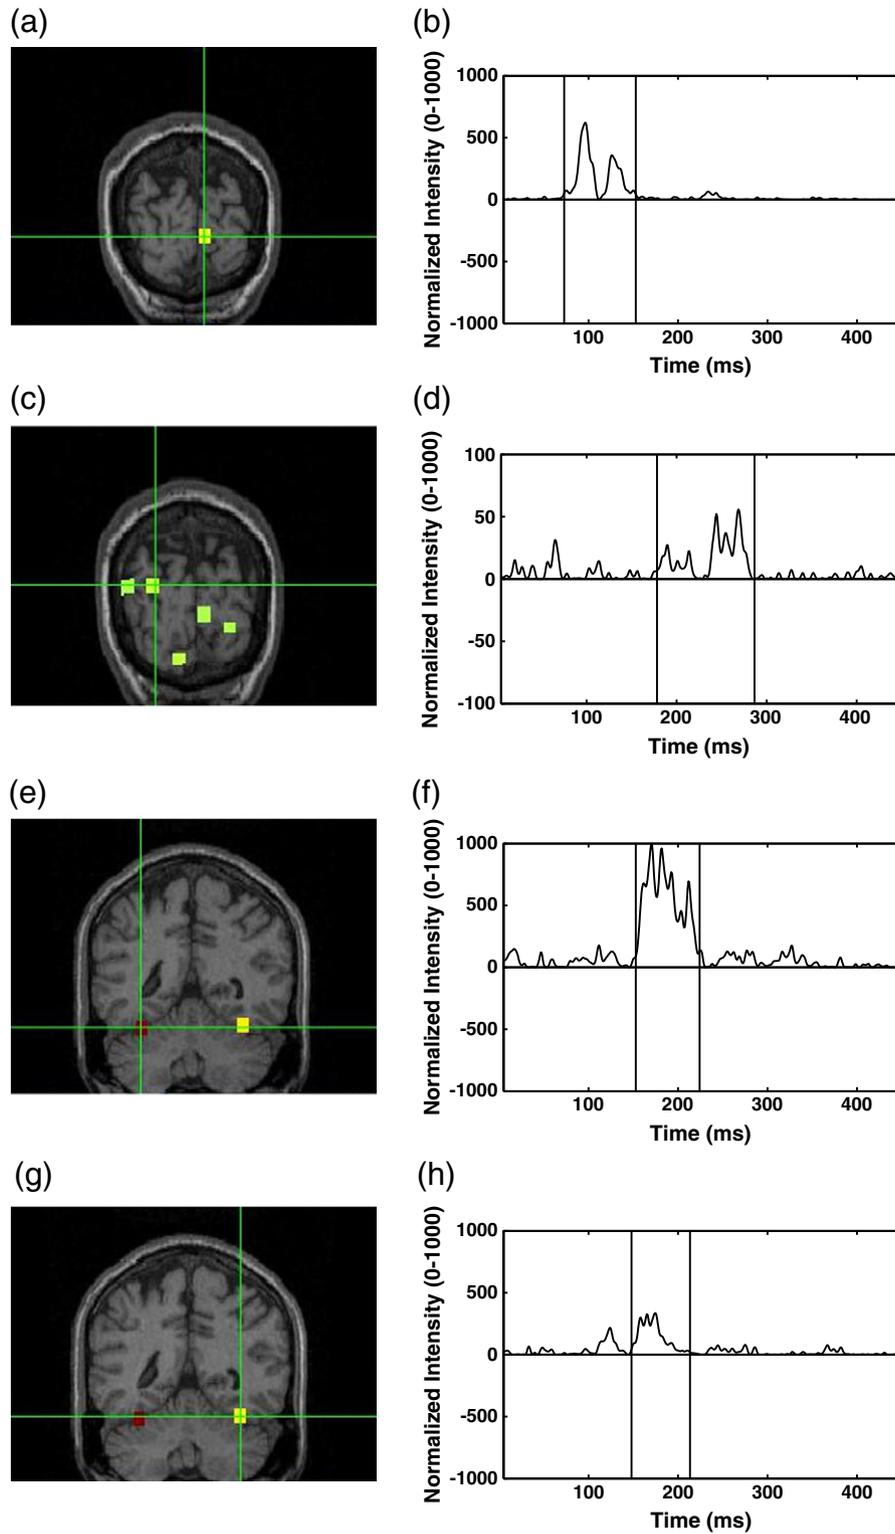


Fig. 14. Results for Champagne run on the face-processing data set. Champagne localizes an early visual response 100 ms after the presentation of a face stimulus, shown in (a). The time course for this source is shown in (b) and corresponds to the location indicated by the crosshairs in the coronal section (a). Champagne localizes a later visual response around 200 ms after the presentation of a face stimulus, seen in (c), with the time course shown in (d). Champagne localizes a bilateral activation in fusiform gyrus that is thought to be in FFA, shown in (e) and (g). The peak for these sources is around 170 ms after the presentation of a face stimulus, time courses shown in (f) and (h).

stimulus period, referred to as dynamic sources. This issue is particularly pertinent when localizing epileptic activity; we plan to investigate the performance of Champagne on dynamic sources in the future.

The experiments with real data highlight the source localization abilities of Champagne. It is harder to evaluate localization accuracy

with real data since the ground truth is not known. For this reason, we have chosen real data sets that have well-established patterns of brain activity; AEF, audio-visual, and face-processing data. For each of the seven AEF data sets presented, Champagne is able to localize bilateral auditory activity in the physiologically accepted region. The other source localization algorithms are not able to provide equally

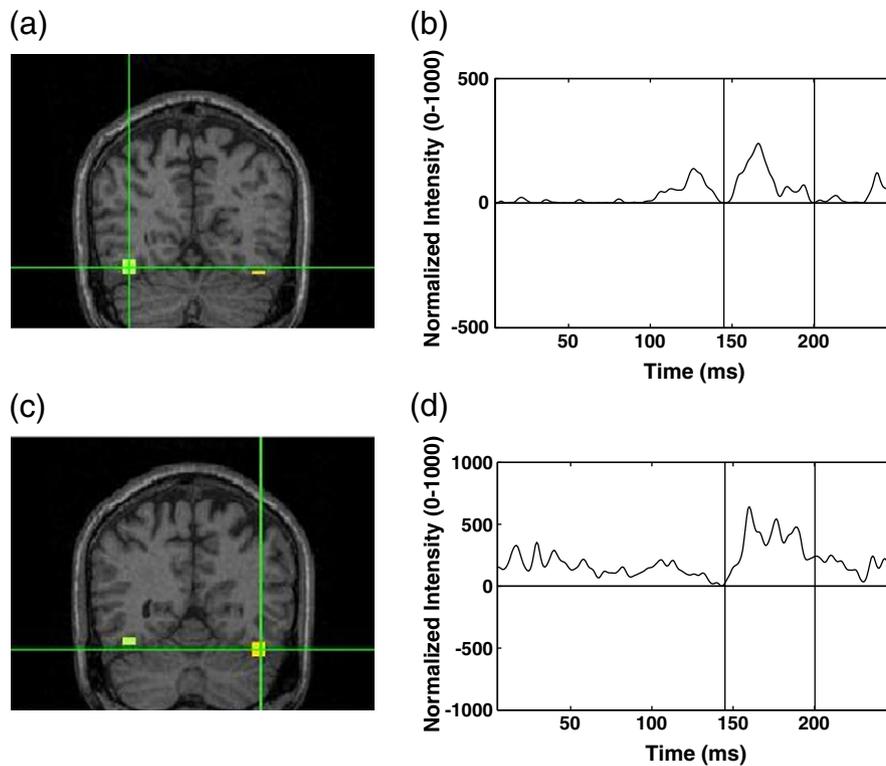


Fig. 15. Results for the face versus scrambled contrast. Champagne localizes a bilateral activation in face-fusiform area (FFA) seen in (a) and (c). The peak is around 170 ms after the presentation of a face stimulus, as seen in the time courses for these sources in (b) and (d).

focal, bilateral activations. Historically, AEF is a difficult data set to obtain accurate activations due to the highly correlated sources. Variants of beamforming have been developed in order to handle the correlated sources, such as coherent suppression beamformers and dual-core beamformers (Dalal et al., 2006; Diwakar et al., 2010). Champagne is able to localize correlated activity without any modifications and it provides a robust solution to this long-standing deficit in source localization. The results from the more complicated data sets, the face processing and the audio-visual task, demonstrate that Champagne is able to localize many distinct, functionally-relevant brain activations. The benchmark algorithms are able to localize some of the sources that Champagne uncovers, but they cannot localize every source nor do they provide focal peaks, as is the case with MVAB and SL/dSPM. A major point of divergence between Champagne and the benchmark algorithms is the estimation of the time courses. In general, Champagne is able to produce time courses that tease apart brain areas, while the other algorithms tend to have very similar time courses at all locations in the brain. The results from the EEG data set shows that Champagne is an effective source localization method for EEG data, localizing activity in both visual areas and in the fusiform gyrus. The benchmark algorithms have better success at localizing the visual activity than the face-specific activations in fusiform gyrus with the EEG data set.

The generative model for Champagne is related to the multiple sparse priors model (MSP) (Friston et al., 2008). Both are covariance component analysis algorithms with some key differences. First, while MSP assumes a set of hundreds of covariance components that represent activity correlated across voxels, Champagne assumes thousands of covariance components representing uncorrelated activity in each voxel. Second, Champagne uses flat hyperpriors for variances and computes MAP estimates for variances with convergent update rules, whereas MSP assumes a Gamma-distribution for hyperpriors for variances, and computes a posterior distribution under a Laplace approximation. Finally, the noise is simultaneously estimated in the MSP algorithm, rather than separately estimated with the pre-

stimulus data, as in Champagne. Given both the similarities and points of divergence, we plan to investigate the comparison of the MSP model to Champagne in future research.

Applying SEFA to the pre- and post-stimulus period can also enable Champagne to be run on “cleaned” post-stimulus data with low-noise covariance. We have found that using SEFA to remove noise can improve the performance of Champagne, particularly on real data. This is an extension of the Champagne algorithm that we plan to explore further. Additionally, Champagne does not make assumptions about the smoothness of sources, nor does it consider uncertainties in the lead field. We are currently investigating the use of spatial- and temporal-smoothness priors, either in the form of basis functions, as in (Zumer et al., 2008), or in the form of autoregressive smoothness priors. Incorporating spatial priors and better noise models, which also model spatiotemporal correlations in the background noise, could potentially improve performance. These extensions of the Champagne method hold promise for improving upon an already robust source localization algorithm.

Although Champagne is embedded in deep theoretical ideas, the resulting algorithm essentially iterates between four simple steps. We believe that the results presented in this paper demonstrate that Champagne makes a significant breakthrough in the reconstruction of brain activity with M/EEG data.

Acknowledgments

We would like to acknowledge the following people who aided in the research presented here. We acknowledge Susanne Honma and Anne Findlay for collecting much of the MEG data in the Biomagnetic Imaging Laboratory. We would also like to thank Carsten Stahlhut for helping with the analyses in SPM and Thomas Jochmann for his help with extending the Champagne algorithm to EEG data. This work was funded in part by an ARCS Graduate Student Fellowship, NIH grants R21 NS076171, and R01 DC004855, DC006435, DC010145, NS067962, NS64060 and by NSF grant BCS-0926196.

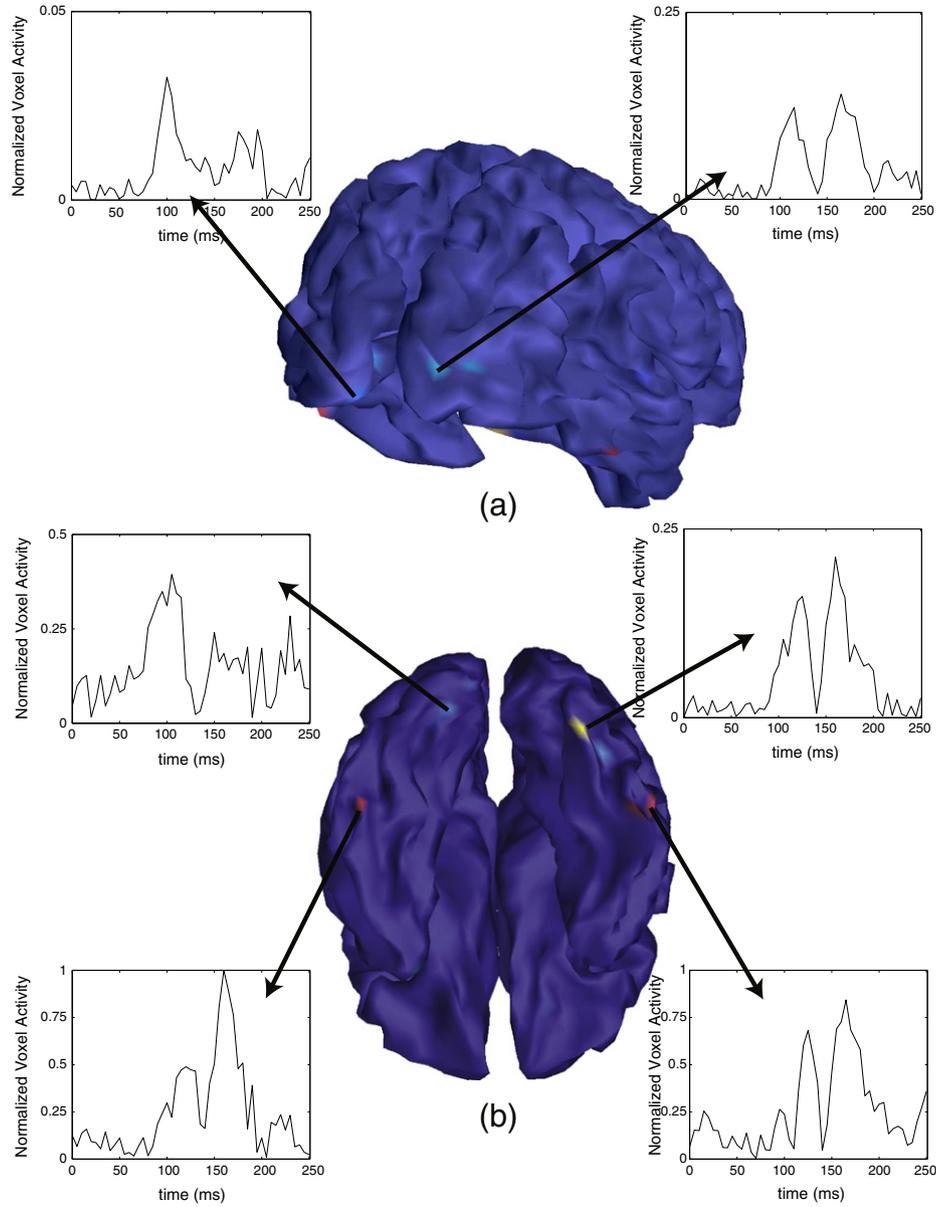


Fig. 16. Results for EEG face processing data from Champagne. (a) Shows two early visual responses in occipital cortex with the time courses. (b) Shows four ventral activations in (or near) face fusiform area (FFA) with time courses showing peaks around 170 ms. (Ventral side of brain shown in (b), with right on the right.)

Appendix A. Benchmark source localization algorithms

For all the algorithms, we localize sources only in the post-stimulus period, referring to as B_{post} .

The estimate for the source-time course ($\hat{s}_r(t)$) at a particular voxel (r) and time point (t) for MVAB is:

$$\hat{s}_r(t) = (L_r^T R^{-1} L_r)^{-1} L_r^T R^{-1} B_{post}(t), \quad (11)$$

where $R = B_{post} B_{post}^T$ is the covariance of the sensor data.

The estimate for the source-time course for sLORETA is:

$$\hat{s}_r(t) = \text{trace} \left((L_r^T G^{-1} L_r)^{-\frac{1}{2}} \right) L_r^T G^{-1} B_{post}(t), \quad (12)$$

where G is called the Gram matrix, $G = LL^T$; it gives an indication of the spatial correlation or overlap between the sensitivity profile of the sensors. Since neighboring sensors have similar sensitivities, this Gram matrix is too close to singular to invert stably. Tikhonov

regularization is needed to invert G , which involves adding a small number to the diagonal entries of the Gram matrix prior to inversion: $G^{-1} = (LL^T + \sigma I)^{-1}$, where σ is a scalar that is chosen empirically, ranging from .001 to 100, and λ is the maximum eigenvalue of the data covariance matrix (R).

Similar to sLORETA, the source-time course estimate for dSPM is:

$$\hat{s}_r(t) = \text{trace} \left((L_r^T G^{-2} L_r)^{-\frac{1}{2}} \right) L_r^T G^{-1} B_{post}(t). \quad (13)$$

We used a version of MCE that is specially tailored to handle multiple time-points and unconstrained source orientations (Wipf et al., 2009). This method extends standard MCE by applying a l_2 penalty across time. In this version there is an l_1 -norm over space and an l_2 -norm over time, sometimes called an $l_{1,2}$ -norm in signal processing. The source orientation components are also included within the l_2 penalty. Similar to Champagne, MCE favors sparse/compact source reconstructions.

The MCE cost function is expressed as:

$$\hat{s}(t) = \arg \min_{\hat{s}} \left\| \left\| B - \sum_{r=1}^{d_s} L_r \hat{s}_r \right\|_{\Sigma_\varepsilon^{-1}}^2 + \sum_{r=1}^{d_s} \|\hat{s}_r\|_{\mathcal{F}} \right\|, \quad (14)$$

where $\|\hat{s}_r\|_{\mathcal{F}} = \sqrt{\text{tr}(\hat{s}_r^T \hat{s}_r)}$ and Σ_ε^{-1} is either set to be the inverse noise covariance learned with SEFA from the pre-stimulus period (as in Champagne) or it is set to be $\Sigma_\varepsilon^{-1} = \sigma I$, where σ is a scalar that is empirically selected and I is an identity matrix of size $d_s \times d_s$.

The estimate for the source time course ($\hat{s}_r(t)$) at a particular voxel (r) and time point (t) is obtained by iterating the following equations:

$$\hat{s}_r^{(k+1)}(t) \rightarrow \alpha_r^{(k)} L_r^T \left(\Sigma_\varepsilon + \sum_{r=1}^{d_s} \alpha_r^{(k)} L_r L_r^T \right)^{-1} B_{\text{post}}(t), \quad (15)$$

where $\alpha_r^{(k)} = \sqrt{\frac{1}{d_s} \sum_{t=1}^{d_t} \|\hat{s}(r, t)^{(k)}\|_2^2}$ and k is the iteration number.

Appendix B. Details of performance metrics calculation

Given a power map (the power of the time course for each voxel) for a given simulation, we determined the value of the Aggregate Performance metric (AP) for that simulation, which combines H_R , A' and \bar{R} . In order to calculate the A' metric, we determined H_R and F_R for each run of the simulation in the following way. First, we found all the “local peaks” in the power map of the sources. A local peak is defined as a voxel that is greater than its 20 three-dimensional nearest neighbors and is at least 10% of the maximum activation of the image. (This thresholding at 10% is designed to filter out any spurious peaks or ripples in the image that are much weaker than the maximum peak.)

After all the local peak locations were obtained, we tested whether each local peak is within 10 mm of a true source location. If a particular peak was within 10 mm of a seeded source, that seeded location gets labeled as a hit and if there was not a local peak within 10 mm, that seeded location gets labeled as a false positive. H_R is then calculated by dividing the number of hits by the true number of seeded sources. Determining F_R is more tenuous, as there is not a clear maximum number of possible false positives, as there is with hits. (Note, that a rate is a ratio between the actual amount and the maximum amount.) We empirically determined the maximum number of false positives for each algorithm, for each given experiment. Since the spatially filtering techniques (MVAB, sLORETA, and dSPM) are inherently smoother than the sparse solutions obtained from Champagne and MCE, the maximum number of false positives was determined empirically across all SNIR levels and 50 simulations for each algorithm. Then, the number of false positives was divided by this maximum false positive number in order to calculate F_R for each simulation and SNIR level. Lastly, the A' metric was calculated for each H_R/F_R pair with the following equations:

$$A' = \begin{cases} \frac{1}{2} + \frac{(H_R - F_R)(1 + H_R - F_R)}{4H_R(1 - F_R)} & \text{for } H_R \geq F_R \\ \frac{1}{2} + \frac{(F_R - H_R)(1 + F_R - H_R)}{4F_R(1 - H_R)} & \text{for } F_R > H_R \end{cases} \quad (16)$$

To assess the accuracy of the time course estimates, we used the correlation coefficient between the true and estimated time courses. Note that we only computed the time course correlation if a local peak is deemed a hit. For a particular simulation, we average the correlation coefficients to obtain one number that reflects the time course reconstruction. The correlation coefficient also ranges from 0 to 1, with 1 implying perfect time course estimation. The equation used to obtain the average correlation coefficient (\bar{R}) between the

true source time course $S_i^{\text{true}}(t)$ and the estimated source time course $S_i^{\text{est}}(t)$ for the i^{th} correctly localized source is as follows:

$$\bar{R} = \frac{1}{N} \sum_{i=1}^N \frac{C(S_i^{\text{true}}(t), S_i^{\text{est}}(t))}{\sqrt{C(S_i^{\text{true}}(t), S_i^{\text{true}}(t)) C(S_i^{\text{est}}(t), S_i^{\text{est}}(t))}} \quad (17)$$

where $C(x_i, x_j)$ is the covariance of x_i and x_j and N is the total number of correctly localized sources.

With the values for H_R , A' and \bar{R} , we can calculate AP for the run of the simulations. For most of the simulations, we ran 50 trials of each configuration of the parameters, i.e. source configuration, SNIR, inter/intra-dipole correlations, and then averaged the AP values to obtain one value for each configuration.

Appendix C. Simulated data parameters

For our purposes, SNIR is defined as:

$$\text{SNIR} \triangleq 20 \log \frac{\|LS\|_{\mathcal{F}}}{\|\mathcal{E}\|_{\mathcal{F}}}, \quad (18)$$

where L is the lead field matrix, S are the source time courses, \mathcal{E} is the noise or non-stimulus evoked activity, see (1) and \mathcal{F} denotes the Froebenius norm or the \downarrow_2 -norm.

We simulated data to replicate the signal-to-noise ratios of real stimulus-evoked data (trial-averaged data). We chose 10 dB and 0 dB as the upper and lower bound for typical trial-averaged, real M/EEG data.

The inter-dipole correlations allow us to test the algorithms' performance when disparate areas of the brain have correlated activity. The intra-dipole correlations allow us to model deviations in the orientation of a dipole; a correlation of 1 would imply a fixed orientation and a correlation of 0 would model a dipole that is fully rotating and not stable in orientation. Another interpretation of the intra-dipole correlation is that the brain volume within a single voxel might contain two or more current sources, and these sources might be aligned in different directions. This would be true of voxels that fall on a 90 degree fold of the cortical surface. We impose intra- and inter-dipolar correlations (for a two-orientation lead field) for n dipoles in the following manner.

First we generate the time courses for the two orientations of each dipole: $S_n^1(t)$, the time course of the n_{th} dipole and the first orientation and $S_n^2(t)$, the time course of the n_{th} dipole and the second orientation. Then we correlate the two orientations of each individual dipole. If the intra-dipolar correlation coefficient is α_{intra} , then $S_n^2(t)$ is defined to be:

$$S_n^2(t) \rightarrow \alpha_{\text{intra}} S_n^1(t) + \sqrt{1 - \alpha_{\text{intra}}^2} S_n^2(t) \quad (19)$$

After we set the intra-dipolar correlation, we correlate all the dipoles (n in number) to the first dipole, $S_1(t)$. $S_n(t)$ is a vector formed by concatenating the time courses for each orientation for the n_{th} dipole: $S_n(t) = \begin{bmatrix} S_n^1(t) \\ S_n^2(t) \end{bmatrix}$, for a two-component lead field. If the inter-dipolar correlation coefficient is α_{inter} , then $S_n(t)$ (for $n \neq 1$) is defined to be:

$$S_n(t) \rightarrow \begin{bmatrix} \alpha_{\text{inter}} & 0 \\ 0 & \alpha_{\text{inter}} \end{bmatrix} S_1(t) + \begin{bmatrix} \sqrt{1 - \alpha_{\text{inter}}^2} & 0 \\ 0 & \sqrt{1 - \alpha_{\text{inter}}^2} \end{bmatrix} S_n(t) \quad (20)$$

Appendix D. Supplementary data

Supplementary data to this article can be found online at doi:10.1016/j.neuroimage.2011.12.027.

References

- Boyd, S., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press.
- Dalal, S., Zumer, J., Agrawal, V., Hild, K., Sekihara, K., Nagarajan, S., 2004. NUTMEG: a neuromagnetic source reconstruction toolbox. *Neurol. Clin. Neurophysiol.* 52, 2004–2052.
- Dalal, S., Sekihara, K., Nagarajan, S., 2006. Modified beamformers for coherent source region suppression. *IEEE Trans. Biomed. Eng.* 53, 1357–1363.
- Dale, A., Liu, A., Fischl, B., Buckner, R., Belliveau, J., Lewine, J., Halgren, E., 2000. Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron* 26, 55–67.
- Darvas, F., Pantazis, D., Kucukaltun-Yildirim, E., Leahy, R.M., 2004. Mapping human brain function with MEG and EEG: methods and validation. *NeuroImage* 23 (Suppl. 1), S289–S299.
- Diwakar, M., Huang, M., Srinivasand, R., Harrington, D., Robb, A., Angelesc, A., Muzzattia, L., Pakdamana, R., Songa, T., Theilmanna, R., Leea, R., 2010. Dual-core beamformer for obtaining highly correlated neuronal networks in MEG. *NeuroImage* 54, 253–263.
- 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 MEG/EEG inverse problem. *NeuroImage* 39, 1104–1120.
- Gorodnitsky, I., Rao, B., 1997. Sparse signal reconstruction from limited data using FOCUS: a re-weighted minimum norm algorithm. *IEEE Trans. Signal. Process.* 45, 600–616.
- Gross, J., Kujala, J., Hamalainen, M., Timmermann, L., Schnitzler, A., Salmelin, R., 2001. Dynamic imaging of coherent sources: studying neural interactions in the human brain. *Proc. Natl. Acad. Sci. U. S. A.* 98, 694–699.
- Hämäläinen, M., Ilmoniemi, R., 1994. Interpreting magnetic fields of the brain: minimum norm estimates. *Med. Biol. Eng. Comput.* 32, 35–42.
- Henson, R., Flandin, G., Friston, K., Mattout, J., 2010. A parametric empirical Bayesian framework for fMRI-constrained MEG/EEG source reconstruction. *Hum. Brain Mapp.* 31, 1512–1531.
- Jaakkola, T., 2000. *Advanced Mean Field Methods: Theory and Practice*. MIT Press. Chapter tutorial on variational approximation methods.
- Kanwisher, N., McDermott, J., Chun, M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17, 4302–4431.
- Kiebel, S., Daunizeau, J., Phillips, C., Friston, K., 2008. Variational Bayesian inversion of the equivalent current dipole model in EEG/MEG. *NeuroImage* 39, 728–741.
- Mackay, D., 1992. Bayesian interpolation. *Neural Comput.* 4, 415–447.
- Mattout, J., Phillips, C., Penny, W., Rugg, M., Friston, K., 2006. MEG source localization under multiple constraints: an extended Bayesian framework. *NeuroImage* 30, 753–767.
- Mosher, J., Lewis, P., Leahy, R., 1992. Multiple dipole modeling and localization from spatiotemporal MEG data. *IEEE Trans. Biomed. Eng.* 39, 541–557.
- Nagarajan, S., Attias, H., Hild, K., Sekihara, K., 2007. A probabilistic algorithm for robust interference suppression in bioelectromagnetic sensor data. *Stat. Med.* 26, 3886–3910.
- Nummenmaa, A., Auranen, T., Hämäläinen, M., I. J., Lampinen, J., Sams, M., Vehtari, A., 2007. Hierarchical Bayesian estimates of distributed MEG sources: theoretical aspects and comparison of variational and MCMC methods. *NeuroImage* 35, 669–685.
- Pascual-Marqui, R., 2002. Standardized low-resolution brain electromagnetic tomography (sLORETA): technical details. *Methods Find. Exp. Clin. Pharmacol.* 24 (Suppl. D), 5–12.
- Phillips, C., Mattout, J., Rugg, M., Maquet, P., Friston, K., 2005. An empirical Bayesian solution to the source reconstruction problem in EEG. *NeuroImage* 24, 997–1011.
- Sahani, M., Nagarajan, S., 2004. Reconstructing MEG sources with unknown correlations. *Adv. Neural. Inf. Process. Syst.* 16.
- Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.* 32, 11–22.
- Sekihara, K., Nagarajan, S., 2008. *Adaptive Spatial Filters for Electromagnetic Brain Imaging*, 1st ed. Springer.
- Snodgrass, J., Corwin, J., 1988. Pragmatics of measuring recognition memory: applications to dementia and amnesia. *J. Exp. Psychol. Gen.* 117, 34–50.
- Uutela, K., Hämäläinen, M., Salmelin, R., 1998. Global optimization in the localization of neuromagnetic sources. *IEEE Trans. Biomed. Eng.* 45, 716–723.
- Uutela, K., Hämäläinen, M., Somersalo, E., 1999. Visualization of magnetoencephalographic data using minimum current estimates. *NeuroImage* 10, 173–180.
- Vrba, J., Robinson, S., 2001. Signal processing in magnetoencephalography. *Methods* 25, 249–271.
- Wipf, D., Nagarajan, S., 2008. A new view of automatic relevance determination. *Adv. Neural. Inf. Process. Syst.* 20.
- Wipf, D., Ramirez, R., Palmer, J., Makeig, S., Rao, B., 2007. Analysis of empirical Bayesian methods for neuroelectromagnetic source localization. *Adv. Neural. Inf. Process. Syst.* 19, 1505–1512.
- Wipf, D., Owen, J., Attias, H., Sekihara, K., Nagarajan, S., 2009. Estimating the location and orientation of complex, correlated neural activity using MEG. *Adv. Neural. Inf. Process. Syst.* 21, 1777–1784.
- Wipf, D., Owen, J., Attias, H., Sekihara, K., Nagarajan, S., 2010. Robust Bayesian estimation of the location, orientation, and time course of multiple correlated neural sources using MEG. *NeuroImage* 49, 641–655.
- Wipf, D., Rao, B., Nagarajan, S., 2011. Latent variable Bayesian models for promoting sparsity. *IEEE Transactions on Information Theory* 57, 6236–6255.
- Zumer, J., Attias, H., Sekihara, K., Nagarajan, S., 2007. A probabilistic algorithm integrating source localization and noise suppression for MEG and EEG data. *NeuroImage* 37, 102–115.
- Zumer, J., Attias, H., Sekihara, K., Nagarajan, S., 2008. Probabilistic algorithms for MEG/EEG source reconstructions using temporal basis functions learned from data. *NeuroImage* 41, 924–940.