HAL
open science

Time-Frequency Mixed-Norm Estimates: Sparse M/EEG imaging with non-stationary source activations

Alexandre Gramfort, Daniel Strohmeier, Jens Haueisen, Matti Hämäläinen, Matthieu Kowalski

To cite this version:

Alexandre Gramfort, Daniel Strohmeier, Jens Haueisen, Matti Hämäläinen, Matthieu Kowalski. Time-Frequency Mixed-Norm Estimates: Sparse M/EEG imaging with non-stationary source activations. NeuroImage, 2013, 70, pp.410-422. 10.1016/j.neuroimage.2012.12.051 . hal-00773276

HAL Id: hal-00773276

https://inria.hal.science/hal-00773276v1

Submitted on 12 Jan 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Time-Frequency Mixed-Norm Estimates: Sparse M/EEG imaging with non-stationary source activations

A. Gramfort[1,2,3,4,*], D. Strohmeier[5], J. Haueisen[5,6], M. Hämäläinen[4], M. Kowalski[7]

Abstract

Magnetoencephalography (MEG) and electroencephalography (EEG) allow functional brain imaging with high temporal resolution. While solving the inverse problem independently at every time point can give an image of the active brain at every millisecond, such a procedure does not capitalize on the temporal dynamics of the signal. Linear inverse methods (Minimum-norm, dSPM, sLORETA, beamformers) typically assume that the signal is stationary: regularization parameter and data covariance are independent of time and the time varying signal-to-noise ratio (SNR). Other recently proposed non-linear inverse solvers promoting focal activations estimate the sources in both space and time while also assuming stationary sources during a time interval. However such an hypothesis only holds for short time intervals. To overcome this limitation, we propose time-frequency mixed-norm estimates (TF-MxNE), which use time-frequency analysis to regularize the ill-posed inverse problem. This method makes use of structured sparse priors

*Institut Mines-Telecom, Telecom ParisTech, CNRS LTCI, 37-39 Rue Dareau, 75014 Paris, France

Email address: alexandre.gramfort@telecom-paristech.fr (A. Gramfort)

[1]Institut Mines-Telecom, Telecom ParisTech, CNRS LTCI, Paris, France

[2]INRIA, Parietal team, Saclay, France

[3]NeuroSpin, CEA Saclay, Bat. 145, 91191 Gif-sur-Yvette Cedex, France

[4]Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, and Harvard Medical School, Charlestown MA, USA

[5]Institute of Biomedical Engineering and Informatics, Ilmenau University of Technology, Ilmenau, Germany

[6]Biomagnetic Center, Dept. of Neurology, University Hospital Jena, Jena, Germany

[7]Laboratoire des Signaux et Systèmes (L2S), Supelec-CNRS-Univ Paris-Sud, Plateau de Moulon, 91192 Gif-sur-Yvette Cedex, France

defined in the time-frequency domain, offering more accurate estimates by capturing the non-stationary and transient nature of brain signals. State-of-the-art convex optimization procedures based on proximal operators are employed, allowing the derivation of a fast estimation algorithm. The accuracy of the TF-MxNE is compared to recently proposed inverse solvers with help of simulations and by analyzing publicly available MEG datasets.

Keywords: Inverse problem, Magnetoencephalography (MEG), Electroencephalography (EEG), sparse structured priors, convex optimization, time-frequency, algorithms

---

1. Introduction

Distributed source models in magnetoencephalography and electroencephalography (collectively M/EEG) use thousands of current dipoles that are used as candidate sources to explain the M/EEG measurements. Those dipoles can be located on a dense three-dimensional grid within the brain volume, typically every 5 mm, or over a surface of the segmented cortical mantle [7], both of which can be automatically segmented from high-resolution anatomical Magnetic-Resonance Images (MRIs). Following Maxwell's equations, each dipole adds its contribution linearly to the measured signal. Note that this linearity of the forward problem is not a modeling assumption but a fact based on the fundamental physics of the problem.

The task in the inverse problem is to map the M/EEG measurements to the brain, i.e., to estimate the distribution of dipolar currents that can explain the measured data. Inverse methods that estimate distributed sources are commonly referred to as *imaging methods*. This is motivated by the fact that the current estimate explains the data and can be visualized as an image, at least at a given point in time. The orientations of the dipoles can be either considered to be known, e.g., by aligning them with the estimated cortical surface normals [7], in which case only the dipole amplitudes need to be estimated. Alternatively, the orientations can be considered as unknown in which case both amplitudes and orientations need to be estimated at each spatial location.

One of the challenges for distributed inverse methods is that the number of dipoles by far exceeds the number of M/EEG sensors: the problem is ill-posed. Therefore, constraints using *a priori* knowledge based on the characteristics of the actual source distributions are necessary. Common

priors are based on the Frobenius norm and lead to a family of methods generally referred to as mininum norm estimators (MNE) [45, 19]. Minimum norm estimates can be converted into statistical parameter maps, which take into account the noise level, leading to *noise-normalized* methods such as dSPM [6] or sLORETA [35]. While these methods have some benefits like simple implementation and a good robustness to noise, they do not take into account the natural assumption that only a few brain regions are typically active during a cognitive task. Interestingly, this latter assumption is what justifies a parametric method known as "dipole fitting" [37] routinely used in clinical practice. In order to promote such focal or *sparse* solutions within the distributed source model framework, one uses sparsity-inducing priors such as a ℓ_p norm with $p \leq 1$ [30, 14]. However, with such priors it is challenging to obtain consistent estimates of the source orientations [42] as well as temporally coherent source estimates [34].

In order to promote spatio-temporally coherent focal estimates, several publications have proposed to constrain the active sources to remain the same over the time interval of interest [34, 11, 46, 15]. The implicit assumption is then that the sources are *stationary*. While this conjecture is reasonable for short time intervals, it is not a good model for realistic sources configurations where multiple transient sources activate sequentially during the analysis period, or simultaneously, before returning to baseline at different time instants.

When working with time series with transient and non-stationary effects, relevant signal processing tools are short time Fourier transforms (STFT) and wavelet decompositions. Contrary to a simple Fast Fourier Transform (FFT), they provide information localized in time and frequency (or scale). In particular, time-frequency decompositions, e.g., Morlet wavelet transforms, are routinely used in MEG and EEG analysis to study transient oscillatory signals. Such decompositions have been employed to analyze both sensor-level data and source estimates, but no attempt has been made to use their output in constructing a regularizer for the inverse problem.

In this contribution, we address the problem of localizing non-stationary focal sources from M/EEG data using appropriate sparsity inducing norms. Extending the work from [15] in which we coined the term Mixed-Norm Estimates (MxNE), we propose to use mixed-norms defined in terms of the time-frequency decompositions of the sources. We call this approach the Time-Frequency Mixed-Norms Estimates (TF-MxNE). The benefit is that the estimates can be obtained over longer time intervals while making stan-

dard preprocessing such as filtering or time-frequency analysis on the sensors optional. The inverse problem is formulated as in [15] as a convex optimization problem whose solutions are computed with an efficient solver based on proximal iterations.

We start with a detailed presentation of the problem and the algorithm. Next, we compare the characteristics and performance of various priors with help of realistic simulated data. Finally, we analyze publicly available MEG datasets (auditory and visual stimulations) demonstrating the benefit of TF-MxNE in terms of source localization and estimation of the time courses of the sources.

A preliminary version of this work was presented at the international conference on Information Processing in Medical Imaging (IPMI) [17]. In this paper we improve the solver to support loose orientation constraints, depth compensation as well as a debiasing step to better estimate source amplitudes. We also analyze new experimental data.

Notation: We indicate vectors with bold letters, $\mathbf{a} \in \mathbb{R}^N$ (resp. $\mathbb{C}^N$) and matrices with capital bold letters, $\mathbf{A} \in \mathbb{R}^{N \times N}$ (resp. $\mathbb{C}^{N \times N}$). $\mathbf{a}[i]$ stands for the ith entry in the vector, while $\mathbf{A}[i, \cdot]$ and $\mathbf{A}[\cdot, i]$ denote the ith row and ith column of a matrix, respectively. We denote $\|\mathbf{A}\|_{\text{Fro}}$ the Frobenius norm, $\|\mathbf{A}\|_{\text{Fro}}^2 = \sum_{i,j=1}^N |\mathbf{A}[i,j]|^2$, $\|\mathbf{A}\|_1 = \sum_{i,j=1}^N |\mathbf{A}[i,j]|$ the ℓ_1 norm, and $\|\mathbf{A}\|_{21} = \sum_{i=1}^N \sqrt{\sum_{j=1}^N |\mathbf{A}[i,j]|^2}$ the ℓ_{21} mixed norm. $\mathbf{A}^T$ and $\mathbf{A}^{\mathcal{H}}$ stand for the matrix transpose and a Hermitian transpose, respectively.

2. General model and method

After a short introduction to Gabor time-frequency dictionaries for M/EEG signals, we present the details of our TF-MxNE inverse problem approach. We then detail the proposed optimization strategy, which uses proximal iterations.

2.1. Gabor dictionaries

Here we briefly present some important properties of Gabor dictionaries, see [8] for more details. Given a signal observed over a time interval, its conventional Fourier transform estimates the frequency content but loses the time information. To analyze the evolution of the spectrum with time and hence the non-stationarity of the signal, Gabor introduced windowed Fourier atoms which correspond to a short-time Fourier transform (STFT) with a

Gaussian window. In practice, for numerical computation, a challenge is to properly discretize the continuous STFT. The discrete STFT with a Gaussian window is also known as the discrete Gabor Transform [12].

The setting we are considering is the finite-dimensional one. Let $\mathbf{g} \in \mathbb{R}^T$ be a "mother" analysis window. Let $f_0 \in \mathbb{N}$ and $k_0 \in \mathbb{N}$ be the frequency and the time sampling rate in the time-frequency plane generated by the STFT, respectively. The family of the translations and modulations of the mother window generates a family of Gabor atoms $(\boldsymbol{\phi}_{mf})_{mf}$ forming the dictionary $\boldsymbol{\Phi} \in \mathbb{C}^{T\times K}$, where K denotes the number of atoms. The atoms can be written as

$$\boldsymbol{\phi}_{mf}[n] = \mathbf{g}[n-mk_0]e^{\frac{i2\pi f_0 fn}{T}}, \quad m \in \{0,\dots,\frac{T}{k_0}-1\}, f \in \{0,\dots,\frac{T}{f_0}-1\} \ . \quad (1)$$

If the product f_0k_0 is small enough, *i.e.,* the time-frequency plane is sufficiently sampled, the family $(\boldsymbol{\phi}_{mf})_{mf}$ is a frame of $\mathbb{R}^T$, *i.e.,* one can recover any signal $\mathbf{x} \in \mathbb{R}^T$ from its Gabor coefficients $(\langle \mathbf{x}, \boldsymbol{\phi}_{mf}\rangle) = \boldsymbol{\Phi}^{\mathcal{H}}\mathbf{x}$. More precisely, there exists two constants $A, B > 0$ such that [1]:

$$A\|\mathbf{x}\|_2^2 \leq \sum_{m,f} |\langle \mathbf{x}, \boldsymbol{\phi}_{mf}\rangle|^2 \leq B\|\mathbf{x}\|_2^2 \ . \quad (2)$$

When $A = B$, the frame is *tight*. When the vectors $\boldsymbol{\phi}_{mf}$ are normalized the frame is an orthogonal basis if and only if $A = B = 1$. The Balian-Low theorem says that it is impossible to construct a Gabor frame which is a basis. Consequently, a Gabor transform is redundant or overcomplete and there exists an infinitely number of ways to reconstruct $\mathbf{x}$ from a given family of Gabor atoms. In the following, the considered $\boldsymbol{\Phi}$ dictionaries are tight frames.

The canonical reconstruction of $\mathbf{x}$ from its Gabor coefficients requires a canonical dual window, denoted by $\tilde{\mathbf{g}}$. Following (1) to define $(\tilde{\boldsymbol{\phi}}_{mf})_{mf}$ we have:

$$\mathbf{x} = \sum_{m,f} \langle \mathbf{x}, \boldsymbol{\phi}_{mf}\rangle \tilde{\boldsymbol{\phi}}_{mf} = \sum_{m,f} \langle \mathbf{x}, \tilde{\boldsymbol{\phi}}_{mf}\rangle \boldsymbol{\phi}_{mf} = \boldsymbol{\Phi}^{\mathcal{H}}\mathbf{x}\tilde{\boldsymbol{\Phi}} = \tilde{\boldsymbol{\Phi}}^{\mathcal{H}}\mathbf{x}\boldsymbol{\Phi} \ ,$$

where $\tilde{\boldsymbol{\Phi}}$ is the Gabor dictionary formed with the dual windows. When the frame is tight, then we have $\tilde{\mathbf{g}} = \mathbf{g}$, and more particularly we have $\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathcal{H}} = \|\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathcal{H}}\|\mathbf{Id}$[8]. The representation being redundant, for any $\mathbf{x} \in \mathbb{R}^T$

[8] We can however say nothing about $\boldsymbol{\Phi}^{\mathcal{H}}\boldsymbol{\Phi}$ in general.

one can find a set of coefficients z_{mf} such that $\mathbf{x} = \sum_{m,f} z_{mf} \boldsymbol{\phi}_{mf}$, while the z_{mf} verify some suitable properties dictated by the application. For example, it is particularly interesting for M/EEG to find a sparse representation of the signal. Indeed, a scalogram, sometimes simply called TF transform of the data in the MEG literature, generally exhibits a few peaks localized in the time-frequency domain. In other words, an M/EEG signal can be expressed as a linear combinations of a few oscillatory atoms. In order to demonstrate this, Fig. 1 shows the STFT of a single planar gradiometer channel MEG signal from a somatosensory experiment, the same STFT restricted to the 50 largest coefficients (approximately only 10% of the coefficients), and the signal reconstructed with only these coefficients compared to the original signal. We observe that the true signal can be well approximated by only a few coefficients, i.e., a few Gabor atoms. In the presence of white Gaussian noise, restricting the time-frequency representation of a signal to the largest coefficients denoises the data. This stems from the fact, that Gaussian white noise in not sparse in the time-frequency domain, but rather spreads energy uniformly over all time-frequency coefficients [40]. Thresholding or shrinking the coefficients therefore reduces noise and smoothes the data. This is further explained in the context of wavelet transforms in [9].

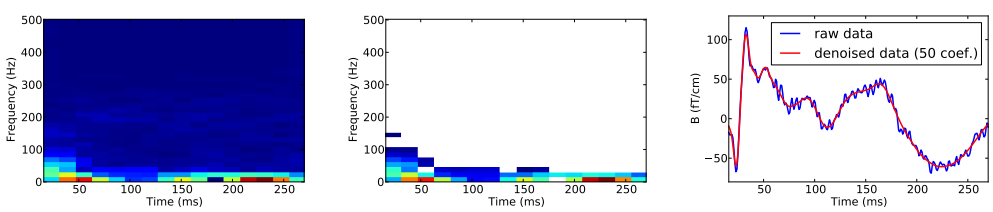

(a) STFT (b) STFT (50 coef.) (c) MEG data

Figure 1: a) Short-time Fourier transform (STFT) of a single channel MEG signal sampled at 1000 Hz showing the sparse nature of the transformation (window size 64 time points and time shift $k_0 = 16$ samples). b) STFT restricted to the 50 largest coefficients c) Data and data reconstructed using only the 50 largest coefficients.

In practice, the Gabor coefficients are computed using the Fast Fourier Transform (FFT) and not by a multiplication by a $\boldsymbol{\Phi}$ matrix as suggested above. Such operations can be efficiently implemented as in the LTFAT toolbox[9] [38]. Another practical concern to keep in mind is the tradeoff

[9]http://ltfat.sourceforge.net/

between the size of the window $\mathbf{g}$ and the time shift k_0. A long window will have a good frequency resolution and a limited time resolution. The time resolution can be improved with a small time shift leading however to a larger computational cost, both in time and memory. Finally, as any computation done with an FFT, the STFT implementations assume circular boundary conditions for the signal. To take this into account and avoid edge artifacts, the signal has to be windowed, e.g., using a Hann window.

2.2. The inverse problem with time-frequency dictionaries

The linearity of Maxwell's equations implies that the signals measured by M/EEG sensors are linear combinations of the electromagnetic fields produced by all current sources. The linear forward operator, called *gain matrix*, predicts the M/EEG measurements due to a configuration of sources based on a given volume conductor model [32]. Given such a linear forward operator $\mathbf{G} \in \mathbb{R}^{N \times P}$, where N is the number of sensors and P the number of sources, the measurements $\mathbf{M} \in \mathbb{R}^{N \times T}$ (T number of time instants) are related to the source amplitudes $\mathbf{X} \in \mathbb{R}^{P \times T}$ by $\mathbf{M} = \mathbf{G}\mathbf{X}$.

The computation of the gain matrix $\mathbf{G}$, e.g., with a Boundary Element Method (BEM) [24, 16], requires modeling of the electromagnetic properties of the head [19] such as the specification of the tissue conductivities. The matrix is then numerically computed. In the inverse problem one computes a best estimate of the neural currents, $\mathbf{X}^\star$, based on the measurements $\mathbf{M}$. However, since $P \gg N$, the problem is ill-posed and priors need to be imposed on $\mathbf{X}$. Historically, the sources amplitudes were computed time instant by time instant using priors based on ℓ_p norms. The ℓ_2 (Frobenius) norm leads to MNE, LORETA, dSPM, or sLORETA while several alternative solvers based on ℓ_p norms with $p \leq 1$ have also been proposed to promote sparse solutions [30, 14]. However, since such solvers work on an instant by instant basis they do not model the oscillatory nature of electromagnetic brain signals. Note that even if the ℓ_2 norm based methods work time instant by time instant, the estimates reflect the temporal characteristics of the data, since they are obtained by linear combinations of sensor data. This, however, implies that the parameters of the inverse solver are independent of time, which corresponds to assuming that the SNR is independent of time. Although MNE type approaches have been used with success, the assumption of constant SNR is clearly wrong since the signal amplitudes vary in time while the noise stays constant, or may be even smaller during an evoked response. The noise is usually estimated from baseline periods such

as prestimulus intervals or periods when the brain is not yet responding to the stimulus.

Beyond single instant solvers, various sparsity-promoting approaches have been proposed [34, 11, 46]. Although they manage to capture the time courses of the activations more accuratly than the instantaneous sparse solvers, they implicitly assume that the active sources are the same over the entire time interval of interest. This also implies that if a source is detected as active at one time point, its activation will be non-zero during the entire time interval of interest. To go beyond this approach, we propose a solver which promotes on the one hand that the source configuration is spatially sparse, and on the other hand that the time course of each active dipole is a linear combination of a limited number of Gabor atoms, as suggested by Fig. 1. Since a Gabor oscillatory atom is localized in time, sources can be marked as active only during a short time period. The model reads:

$$\mathbf{M} = \mathbf{G}\mathbf{X} + \mathbf{E} = \mathbf{G}\mathbf{Z}\mathbf{\Phi}^{\mathcal{H}} + \mathbf{E} \ , \tag{3}$$

where $\mathbf{\Phi}^{\mathcal{H}} \in \mathbb{C}^{K \times T}$ is a dictionary of K Gabor atoms, $\mathbf{Z} \in \mathbb{C}^{P \times K}$ are the coefficients of the decomposition, and $\mathbf{E}$ is additive white noise, $\mathbf{E} \sim \mathcal{N}(0, \lambda \mathbf{I})$. Given a prior on $\mathbf{Z}$, $\mathcal{P}(\mathbf{Z}) \sim \exp(-\Omega(\mathbf{Z}))$, the maximum a posteriori estimate (MAP) is obtained by solving:

$$\mathbf{Z}^{\star} = \arg\min_{\mathbf{Z}} \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{Z}\mathbf{\Phi}^{\mathcal{H}}\|_{\text{Fro}}^2 + \lambda \Omega(\mathbf{Z}) \ , \ \ \lambda > 0 \ . \tag{4}$$

If we consider $\Omega(\mathbf{Z}) = \|\mathbf{Z}\|_1$, (4) corresponds to a LASSO problem [39], a.k.a. Basis Pursuit Denoising (BPDN) [4], where features (or regressors) are spatio-temporal atoms. Similarly to the original formulation of MCE, *i.e.,* ℓ_1 regularization without applying $\mathbf{\Phi}$, such a prior is likely to suffer from inconsistencies over time [34]. Indeed such a norm does not impose a structure for the non-zero coefficients: they are likely to be scattered all over $\mathbf{Z}^{\star}$ (see Fig. 2). Therefore, simple ℓ_1 priors do not guarantee that only a few sources are active during the time window of interest. To promote this, one needs to employ mixed-norms such as the ℓ_{21} norm [34, 15]. By doing so, the estimates have a sparse row structure (see Fig. 2). However the ℓ_{21} prior on $\mathbf{Z}$ does not produce denoised time series as it does not promote source estimates that are formed by a sum of a few Gabor atoms. In order to recover the sparse row structure, while simultaneously promoting sparsity of the decompositions, we propose to use a composite prior formed by the

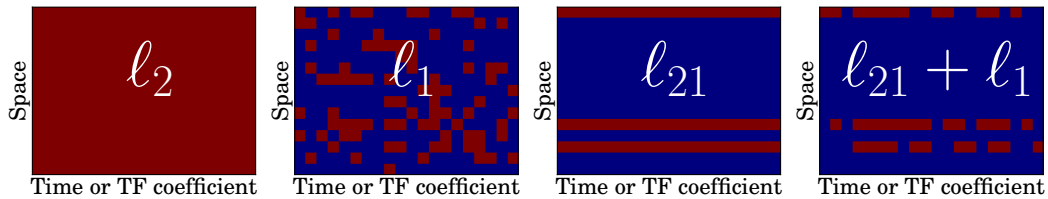

Figure 2: Sparsity patterns promoted by the different priors: ℓ_2 all non-zero, ℓ_1 scattered and unstructured non-zero, ℓ_{21} block row structure, and $\ell_{21} + \ell_1$ block row structure with intra-row sparsity. Red color indicates non-zero coefficients.

sum of ℓ_{21} and ℓ_1 norms. The prior then reads:

$$\lambda\Omega(\mathbf{Z}) = \lambda_{\text{space}}\|\mathbf{Z}\|_{21} + \lambda_{\text{time}}\|\mathbf{Z}\|_1 \ , \ \ \lambda_{\text{space}} > 0, \lambda_{\text{time}} > 0 \ . \tag{5}$$

A large regularization parameter λ_{space} will lead to a spatially very sparse solution, while a large regularization parameter λ_{time} will promote sources with smooth times series. This is due to the uniform spectrum of the noise (see Section 2.1) and the fact that a large λ_{time} will promote source activations made up of few TF atoms, each of which has a smooth waveform.

2.3. Optimization strategy

The optimization strategy, which we propose for minimizing the cost function in (4), is based on the Fast Iterative Shrinkage Thresholding Algorithm (FISTA) [2], a first-order schemes that handles the minimization of any cost function $\mathcal{F}$ that can be written as a sum of two terms: a smooth convex term f_1 with Lipschitz gradient and a convex term f_2, potentially non-differentiable: $\mathcal{F}(\mathbf{Z}) = f_1(\mathbf{Z}) + f_2(\mathbf{Z})$. In order to apply FISTA, we need to be able to compute the so-called proximity operator associated with f_2, *i.e.*, the proximity operator associated with the composite $\ell_{21} + \ell_1$ prior [17].

Definition 1 (Proximity operator). *Let $\varphi : \mathbb{R}^M \to \mathbb{R}$ be a proper convex function. The proximity operator associated to φ, denoted by $\text{prox}_\varphi : \mathbb{R}^M \to \mathbb{R}^M$ reads:*

$$\text{prox}_\varphi(\mathbf{Z}) = \underset{\mathbf{V}\in\mathbb{R}^M}{\arg\min} \frac{1}{2}\|\mathbf{Z} - \mathbf{V}\|_2^2 + \varphi(\mathbf{V}) \ .$$

While the proximity operators of mixed-norms relevant for M/EEG can be found in [15], in the case of the composite prior in (5), the proximity operator is given by the following lemma.

Lemma 1 (Proximity operator for $\ell_{21} + \ell_1$). *Let $\mathbf{Y} \in \mathbb{C}^{P\times K}$ be indexed by a double index (p, k). $\mathbf{Z} = \mathrm{prox}_{\lambda\|.\|_1+\mu\|.\|_{21})}(\mathbf{Y}) \in \mathbb{C}^{P\times K}$ is given for each coordinates (p, k) by*

$$Z[p,k] = \frac{Y[p,k]}{|Y[p,k]|}\left(|Y[p,k]| - \lambda\right)_+ \left(1 - \frac{\mu}{\sqrt{\sum_k (|Y[p,k]| - \lambda)_+^2}}\right)_+ .$$

where for $x \in \mathbb{R}$, $(x)_+ = \max(x, 0)$, and by convention $\frac{0}{0} = 0$.

This result is a corollary of the proximity operator derived for hierarchical group penalties recently proposed in [23]. The penalty described here can indeed be seen as a 2-level hierarchical structure, and the resulting proximity operator reduces to successively applying the ℓ_1 and ℓ_{21} proximity operator. Both of these proximity operators are discussed in detail in [15].

The pseudo code is provided in Algorithm 1. The Lipschitz constant $\mathcal{L}$ of the gradient of the smooth term in (4) is given by the square of the spectral norm of the linear operator $\mathbf{Z} \to \mathbf{G}\mathbf{Z}\mathbf{\Phi}^{\mathcal{H}}$. We estimate it with the power iteration method.

Algorithm 1 FISTA with TF dictionaries to minimize 4

Input: Measurements $\mathbf{M}$, gain matrix $\mathbf{G}$, regularization parameter $\lambda > 0$ and I the number of iterations.

Output: $\mathbf{Z}^\star$

Auxiliary variables : $\mathbf{Y}$ and $\mathbf{Z}_o \in \mathbb{R}^{P\times K}$, and τ and $\tau_o \in \mathbb{R}$.
Estimate the Lipschitz constant $\mathcal{L}$ with the power iteration method.
$\mathbf{Y} = \mathbf{Z}^\star = \mathbf{Z}$, $\tau = 1$, $0 < \mu < \mathcal{L}^{-1}$
for $i = 1$ to I **do**
  $\mathbf{Z}_o = \mathbf{Z}^\star$
  $\mathbf{Z}^\star = \mathrm{prox}_{\mu\lambda\Omega}\left(Y + \mu\mathbf{G}^T(\mathbf{M} - \mathbf{G}\mathbf{Y}\mathbf{\Phi}^{\mathcal{H}})\mathbf{\Phi}\right)$
  $\tau_o = \tau$
  $\tau = \frac{1+\sqrt{1+4\tau^2}}{2}$
  $\mathbf{Y} = \mathbf{Z}^\star + \frac{\tau_o - 1}{\tau}(\mathbf{Z}^\star - \mathbf{Z}_o)$
end for

3. Specific modeling for M/EEG inverse problem

The M/EEG literature has shown that general solvers of the statistics literature need to be adapted to the specificities of the M/EEG inverse problem.

Crucial steps in the computation of the source estimates are noise whitening, depth compensation, handling of source orientations, and amplitude bias correction.

3.1. Spatial whitening

The model in (3) assumes that the additive noise is Gaussian white with $\mathbf{E} \sim \mathcal{N}(0, \lambda \mathbf{I})$. This strong modeling assumption is made realistic by a whitening step that relies on estimating the *noise covariance matrix*. For this purpose, *baseline* data is employed, which is recorded while the subject is at rest e.g. during pre-stimulus periods. If only MEG is recorded, the noise covariance can be estimated from data recorded without subject, often called *empty room* data. This approach provides good estimates of the measurement noise level. Although the noise level depends on the signal frequency, one usually uses a single frequency-unspecific noise covariance matrix. An alternative approach for frequency-dependent spatial whitening is presented in [36].

The whitening step is particularly fundamental when different sensor types are used: EEG and MEG with gradiometers and magnetometers record signals with different units of measure and with different noise levels. The whitening step makes data recorded by different sensors comparable and adapted for joint estimation.

3.2. Source models with unconstrained orientations

When the source orientations given by the normals of the cortical mesh cannot be trusted, it can be interesting to relax this constraint by placing three orthogonal sources at each spatial location. When all three orientations are allowed to explain the data equivalently the model is called *free orientation*. Moreover, it can be of interest to have intermediate models using *loose orientation* constraints [27]. However for such loose and free orientation models, the TF composite prior needs to be adapted. Let each source be indexed by a spatial location i and an orientation $o \in \{1, 2, 3\}$. Let $o = 1$ correspond to the orientation normal to the cortex, and $o = 2$ and $o = 3$ the two tangential orientations. We call $0 < \rho \leq 1$, the parameter controlling how loose is the orientation constraint. The ℓ_1 and ℓ_{21} norms read:

$$\|\mathbf{Z}\|_1 = \sum_{i,k} \sqrt{|\mathbf{Z}[(i, o = 1), k]|^2 + \frac{1}{\rho}|\mathbf{Z}[(i, o = 2), k]|^2 + \frac{1}{\rho}|\mathbf{Z}[(i, o = 3), k]|^2}$$

$$\|\mathbf{Z}\|_{21} = \sum_i \sqrt{\sum_k |\mathbf{Z}[(i, o=1), k]|^2 + \frac{1}{\rho}|\mathbf{Z}[(i, o=2), k]|^2 + \frac{1}{\rho}|\mathbf{Z}[(i, o=3), k]|^2} \ ,$$

where k indexes the TF coefficients. When $\rho = 1$ the orientation is free and it amounts to grouping the orientations in a common ℓ_2 norm such as in [34, 20]. Such priors are a principled way of supporting loose orientation constraints in the context of non-ℓ_2 priors.

Observe here that $\|\mathbf{Z}\|_1$ is not an ℓ_1 norm *per se*. Indeed, it is a ℓ_{21} norm, but we have chosen to keep the same notations as in the constrained orientation case for the sake of readability.

In practice, using free orientation models means that at a given location, the current dipoles selected to explain the data can have an orientation that varies in time similarly to the *rotating dipole model* employed in dipole fitting.

3.3. Depth compensation

The principal contribution to M/EEG data comes from superficial cortical gray matter: deep sources are attenuated due to their larger distance from the sensors. While it is common in statistics to scale the columns of the gain matrix such that $\|\mathbf{G}_i\|_2 = 1$, practice with M/EEG data shows that it is often not a good idea. The rationale in statistics is to avoid favoring regressors, here sources, just due to the amplitude of the corresponding column in the gain matrix. When doing this for M/EEG, it tends to favor too much very deep sources which are less likely to be visible with M/EEG. For this reason, a common practice with MNE type approaches is to use a softer depth bias compensation. Given a parameter $0 \leq \gamma \leq 1$, the three columns ($\mathbf{G}[\cdot, (i, o=1)]$, $\mathbf{G}[\cdot, (i, o=2)]$, $\mathbf{G}[\cdot, (i, o=3)]$) of $\mathbf{G}$ for the three orientations at the same location are normalized by $\sqrt{(\|\mathbf{G}[\cdot, (i, o=1)]\|_2^2 + \|\mathbf{G}[\cdot, (i, o=2)]\|_2^2 + \|\mathbf{G}[\cdot, (i, o=3)]\|_2^2)^\gamma}$. If $\gamma = 0$ it corresponds to no depth bias compensation and $\gamma = 1$ leads to full scaling which may lead to spurious deep sources appearing in the results.

3.4. Source weighting: fMRI priors?

Mixed-norm regularizations [15] can be written with spatially dependent scalar weights. It can be used to promote some sources by reducing their regularization. For example, given a weight vector $\mathbf{w} \in \mathbb{R}_+^P$, one weight per physical location, the TF-MxNE prior can be modified as:

$$\lambda \Omega_{\mathbf{w}}(\mathbf{Z}) = \sum_{i=1}^{P} \mathbf{w}[i] \left(\lambda_{\text{space}} \|\mathbf{Z}[i, \cdot]\|_2 + \lambda_{\text{time}} \|\mathbf{Z}[i, \cdot]\|_1\right) \ ,$$

where $\mathbf{Z}[i, \cdot]$ stands for the ith row of $\mathbf{Z}$. If $\mathbf{w}[i]$ is small the regularization for the source at location i will be small and the source i is likely to be selected to explain the data. Assuming that additional location information of the sources is available, e.g., from fMRI, information is known about the sources, such as fMRI localizations, it would be possible to inject this knowledge in the prior in order to have fMRI informed sparse estimates. Note that sparsity promoting priors do not lead to source estimates where every dipole in the source space has a non-zero activation. It means that, although some regions are promoted by the weights, they may not contain any estimated source. It indeed may happen that MEG misses sources, for example if they are radially oriented. In this sense the proposed weighted scheme does not act as a strong prior on the MEG source localization.

For computational reasons, one can also exploit fast solvers such as dSPM or sLORETA to derive scalar weights that can help reduce the number of candidate sources. Typically, one can threshold dSPM/sLORETA estimates and restrict the TF-MxNE solver to a small portion of the cortex, further improving the computational efficiency of the optimization algorithm. It corresponds to setting $\mathbf{w}[i]$ to infinity (or very large) if the ith spatial location yields very low dSPM values at all points in time.

3.5. Amplitude bias compensation

Methods based on ℓ_1 priors, such as TF-MxNE, are known to impose an amplitude bias on the solution. This is due to the general bias-variance trade-off in statistical estimation. With ℓ_1 based priors, the high sparsity of the solution comes at the price of a strong amplitude bias. Given the waveforms for the selected sources it is possible to post-process them and correct the amplitude bias leading to meaningful amplitudes of the source activations. See [18] for an example of amplitude bias correction in the context of fMRI decoding.

A first natural approach to correct for the amplitude bias is to compute the least squares solution restricted to the active set of sources provided by the TF-MxNE solution. It amounts to computing a dipole fit with a known set of dipoles, which is no longer an ill-posed problem. However, this procedure affects the source time courses, and the signal smoothness promoted by the TF-MxNE is lost. Hence, rather than re-estimating the source time courses using least squares, we correct the amplitude bias by scaling the TF-MxNE results. For this purpose, we introduce a diagonal scaling matrix $\mathbf{D}$, whose diagonal elements are scaling factors for all sources in

the active set. These scaling factors are constrained to be above 1 to actually remove the bias, and are constant over time. Furthermore, in the case of free orientation, they are identical for all orientations at a given location in order to preserve the source characteristics and orientations estimated using TF-MxNE. The bias corrected source estimate $\tilde{\mathbf{X}}$ is computed using $\mathbf{D}$ as $\tilde{\mathbf{X}} = \mathbf{DX} = \mathbf{DZ\Phi}^{\mathcal{H}}$. We estimated the scaling matrix $\mathbf{D}$ based on the following convex optimization problem:

$$\mathbf{D}^{\star} = \arg\min_{\mathbf{D}} \|\mathbf{M} - \mathbf{GDX}\|_{\text{Fro}}^2 \text{ s.t. } \begin{cases} \mathbf{D_{ij}} \geq 1, & i = j \\ \mathbf{D_{ij}} = 0, & i \neq j \end{cases}.$$

The optimization problem can also be solved efficiently with FISTA after writing the constraint on $\mathbf{D}$ as an indicator function over a convex set $\mathcal{C} = \{\mathbf{D} \text{ s.t. } \mathbf{D_{ii}} \geq 1, \text{ and } \mathbf{D_{ij}} = 0, \text{ if } i \neq j\}$:

$$f_2(\mathbf{D}) = \iota_{\mathcal{C}}(\mathbf{D}) = \begin{cases} 0 & \text{if } \mathbf{D} \in \mathcal{C} \\ \infty & \text{otherwise} \end{cases}.$$

4. Practical details

This section presents the details in the efficient implementation of Algorithm 1. We also discuss the choice of the hyperparameters (regularization parameters).

4.1. Implementation

Algorithm 1 requires to compute Gabor transforms at each iteration which can be computationally demanding. However, due to the ℓ_{21} sparsity inducing prior, only a few rows of $\mathbf{Z}$ have non-zero coefficients. The Gabor transform is therefore computed for only a limited number of rows, equivalently a small number of active sources. This makes the computation of $\mathbf{Y\Phi}^{\mathcal{H}}$ (cf. Algorithm 1 line 6) much faster.

Also when a tight frame is used, the ℓ_{21} norm of a signal does not change when $\mathbf{\Phi}$ is applied. This means that the ℓ_{21} proximity operator can be applied to temporal data to discard some sources from the active set without computing the STFT. This comes from the fact that if $\text{prox}_{\|\ \|_{21}}(\mathbf{x}) = \mathbf{0}$ for a time series $\mathbf{x}$ then $\text{prox}_{\|\ \|_1 + \|\ \|_{21}}(\mathbf{\Phi x}) = \mathbf{0}$.

Since the proposed optimization problem is convex, the solution does not depend on initial conditions. Hence, in order to further reduce the computation time, it is beneficial to initialize the TF-MxNE solver with the ℓ_{21}

MxNE solution obtained with the same spatial regularization, since MxNE can be computed efficiently using active set strategies [15]. Note again that the ℓ_{21} MxNE solution is used as an initialization and not for restricting the source space.

4.2. Selection of the regularization parameters

Model selection in the present case amounts to setting the regularization parameters λ_{space} and λ_{time}, as well as the parameter of the Gabor transform, namely the time resolution with k_0 and the frequency resolution, function of the window length T. The parameter k_0 and T will depend on the length of a time interval during which signals can be considered stationary. A too high sampling of the time-frequency plane will also lead to high computational costs. The regularization parameters have an effect on the spatial sparsity, the number of active dipoles, and the temporal smoothness of the source time series. Different strategies exist to set such model parameters (cross-validation, discrepancy principle etc.).

In the case of ℓ_{21} priors, one can prove that there exists a value of $\lambda_{space}^{\max}$ for λ_{space} such that if $\lambda_{\text{space}} \geq \lambda_{space}^{\max}$, then $\mathbf{Z}^*$ is filled with zeros, i.e., no source is active. This provides a convenient way to specify the regularization parameter as the ratio of λ_{space} and $\lambda_{space}^{\max}$, between 0 and 1. In the next section if λ_{space} is given as a percentage it corresponds to this ratio, rescaled to percents. For convenience, the parameter λ_{time} can then be also scaled by $\lambda_{\text{space}}^{\max}$. The benefit of the reparametrization of the regularization parameter is that they become much less sensitive to the dataset. Assuming $\mathbf{\Phi}$ is a tight frame then $\|\mathbf{X}\|_{21} = \|\mathbf{X\Phi}\|_{21} = \|Z\|_{21}$, then one can show based on the optimality conditions for the ℓ_{21} mixed-norm [15] that:

$$\lambda_{space}^{\max} = \max_i \|(\mathbf{G^T M})[i, \cdot]\|_2$$

5. Results

In the following, we first evaluate the accuracy of our solver with simulations. We then apply our solver to two MEG/EEG datasets.

5.1. Simulation study

In order to have a reproducible and reasonably fast comparison of different priors, we generated a small simulation dataset with 20 EEG electrodes and 200 sources. Four of these sources were randomly selected to be active. The

ECD waveforms (Fig. 4(a)) represent 1 high and 3 low frequency components. The time course of the oscillatory high frequency component is modeled by a Gabor atom, whereas the time courses of the low frequency components were obtained from a somatosensory evoked potential study [22] by fitting manually ECDs to the P15, N20 and P23 components. To make the comparison of the priors independent of the forward model and the sources spatial configuration, the linear forward operator was a random matrix, whose columns were normalized to 1. The scaling of the columns is not mandatory here but simplifies the parameter setting by making it independent of the number of simulated sensors. White Gaussian noise or realistic 1/f noise, simulated with an auto-regressive (AR) process of order five calibrated on real MEG data, was added to the signals to achieve a desired signal-to-noise ratio (SNR). Following the notation of (3), we define SNR as $20\log_{10}(\|M\|_{\text{Fro}}/\|E\|_{\text{Fro}})$.

Figure 3 presents the RMSE on the estimation for different solvers as a function of λ_{space} (RMSE $= \|\mathbf{X}_{\text{sim}} - \mathbf{X}^{\star}_{\Omega}\|^2_{\text{Fro}}$) for two different values of λ_{time}, with and without correcting for the amplitude bias. High and low noise conditions were investigated, with both white and AR noise. The parameter λ_{space} was chosen on a logarithmic grid from 10^{-2} to $10^{2.5}$. The Gabor dictionary is tight, constructed with a 64 sample window $\mathbf{g}$ with $k_0 = 4$ samples time shift.

Many observations can be made from these results. First, the spatio-temporal solvers, namely MxNE and TF-MxNE, outperform instantaneous MNE and MCE. Second, we observe that a small value of λ_{time} in TF-MxNE leads to results similar to MxNE which is fully expected since we use a tight frame that does not change the ℓ_2 norm of a signal. Finally, we observe that correcting for the amplitude bias clearly improves the results in all conditions and particularly with high values of λ_{space}. This suggests to run the solver with a high value of λ_{space} (high percentage of $\lambda^{\max}_{space}$) which also leads in practice to a faster convergence of the solver. Note also that, as expected, violation of the modeling assumption about the noise, namely Gaussian white noise, produces a degradation in the performance of the solver in the AR condition. Figure 4 shows the reconstructions for the best λ_{space} according to Figure 3 for the ℓ_1, ℓ_{21} MxNE and the TF-MxNE (with and without bias correction). It can be observed, that the TF-MxNE method with the composite TF prior is able to reconstruct the smooth time course of the simulated sources contrary to ℓ_1 and ℓ_{21} priors.

The TF composite prior was then challenged on a realistic EEG configuration with a 4-shell spherical head model (radii 0.90, 0.92, 0.96 and 1) and

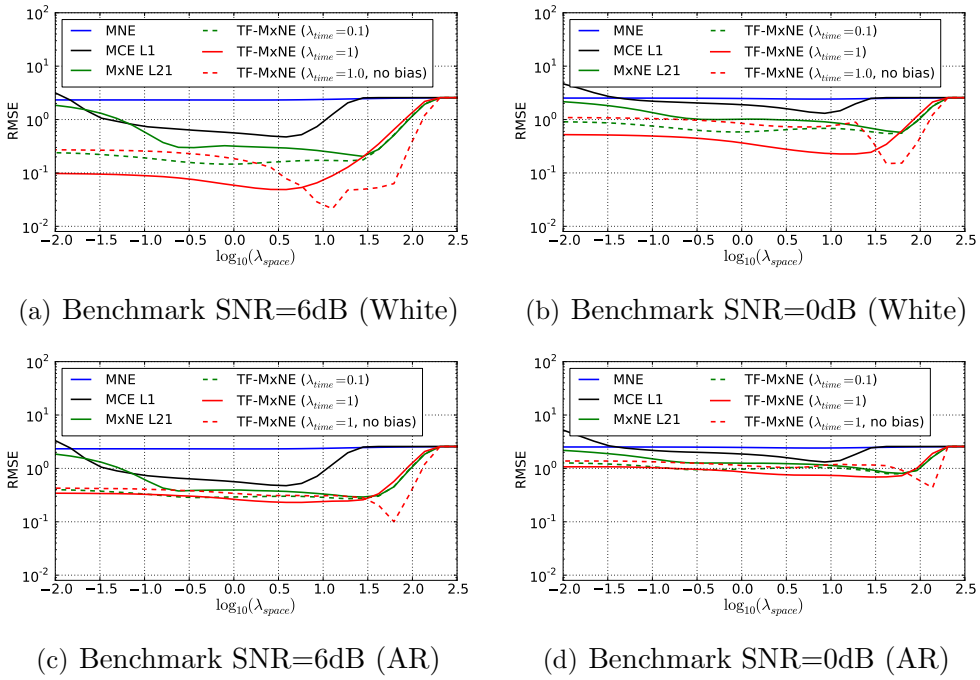

(a) Benchmark SNR=6dB (White)

(b) Benchmark SNR=0dB (White)

(c) Benchmark SNR=6dB (AR)

(d) Benchmark SNR=0dB (AR)

Figure 3: Comparison of RMSE in the source space as a function of λ_{space} (SNR=6bB and SNR=0bB with Gaussian white noise and AR noise). TF priors improve the reconstruction and the best accuracy is obtained with the TF $\ell_{21}+\ell_1$ prior for both noise conditions. The bias correction improves the performance of the reconstruction in both cases, particularly for high values of λ_{space}.

60 electrodes placed according to the international 10-5 electrode system. The source waveforms were the same as before. The source space in Fig. 5 consisted of 152 sources in a regular grid (0.2 spacing) inside the innermost sphere. Source orientations were randomly selected. Figure 5 shows the head model and reconstructions obtained with ℓ_{21} MxNE and the TF-MxNE. Even if the performance drops down due to the limited spatial resolution of EEG, the TF composite prior gives the best RMSE and is able to reconstruct and separate the high frequency component.

5.2. Experimental results with MEG/EEG data

We also applied the TF-MxNE solver to publicly available data: the auditory and visual conditions in the data shipped as sample data with the MNE software (http://martinos.org/mne/). In this MEG/EEG experiment, checkerboard patterns were presented into the left and right visual field,

interspersed by tones to the left or right ear. The interval between the stimuli was 750 ms. Occasionally a smiley face was presented at the center of the visual field. The subject was asked to press a key with the right index finger as soon as possible after the appearance of the face. We will report results obtained on the responses evoked by the auditory stimuli presented to the left ear and visual stimuli shown in the left hemifield.

Data were collected in a magnetically shielded room using the whole-head Elekta Neuromag Vector View 306 MEG system (Neuromag Elekta LTD, Helsinki) equipped with 102 triple-sensor elements (two orthogonal planar gradiometers and one magnetometer per location). EEG data from a 60-channel electrode cap was acquired simultaneously with the MEG.

The signals were recorded with a bandpass of 0.1 - 172 Hz, digitized at 600 samples/s and averaged offline triggered by the stimulus onset. All epochs containing EOG signals higher than 150 μV peak-to-peak amplitude were discarded from the averages. Peak-to-peak rejection parameters for EEG was set to 80 μV, for magnetometers 4000 fT and for gradiometers 2000 fT/cm. This resulted in respectively 55 and 67 averaged epochs for the left auditory and left visual conditions. Evoked data were baseline corrected using 200 ms of pre-stimulus data. The same data segment was used to estimate the noise covariance matrix for spatial whitening. The M/EEG recordings contained two bad channels (1 EEG and 1 MEG), leading to a total of 364 combined M/EEG channels. Signal-space projection (SSP) correction was applied to MEG magnetometers data to suppress environmental noise and biological artifacts [41].

The anatomical MRI data were collected with a Siemens Trio 3T scanner with a T1-weighted sagittal MPRAGE protocol, which were employed for cortical surface reconstruction using FreeSurfer. Two multi-echo 3D Flash acquisitions were also performed to extract the inner skull surface for the 3-layers boundary-element model used for forward model computation. The dipolar sources were sampled over the cortical mantle with 5 mm average distance between sources leading to a total of 7498 dipoles.

5.2.1. Auditory data

The sources were estimated assuming loose orientation constraint ($\rho = 0.2$) and depth bias compensation with $\gamma = 0.9$. The Gabor dictionary was tight, constructed with a 16 samples ($\simeq 27$ ms) long window $\mathbf{g}$ with $k_0 = 4$ samples time shift and $f_0 = 1$ sample frequency shift. A scalar weighting as described in Section 3.4, was performed with a dSPM solution obtained with

the same depth weighting and loose orientation. The λ_{time} parameter was set to 1%. Results are presented in Figure 6.

In order to compare the estimated sources of the TF-MxNE, we computed the solution using an ℓ_{21} MxNE with the same spatial regularization. It amounts to setting λ_{time} in the TF-MxNE to zero and hence not promoting any temporal regularization. This can be observed in the Figure 6 where both MxNE and TF-MxNE lead to bilateral auditory sources given an $\lambda_{\text{space}} = 50\%$. However the TF-MxNE, leads to smooth time courses and true zero activations during baseline. Note that any sparse solver that would only promote spatial sparsity without modeling the dynamics of the signal would fail to estimate true null activations during baseline.

Further looking at the time series of the sources, it can be observed that the contralateral auditory cortex activates before the ipsilateral. This is consistent with the literature. The peak to peak latency difference between right and left cortices for the main activity, known as the N100, can be clearly quantified on the smooth waveforms provided by the TF-MxNE. It is for this subject equal to 16 ms. These experimental results show how the proposed method can provide both fine spatial localization and temporal smoothness, which e.g. allows to fully exploit the temporal resolution of M/EEG for chronometry.

In order to illustrate, what fraction of the data has actually been explained by the bilateral auditory sources obtained with our choice of hyperparameters λ_{space} and λ_{time}, we present in Figure 6-a the original data restricted to the gradiometers as well as in Figure 6-b the data predicted by the sources. Note that all MEG and EEG sensors were used for the localization. However, we show only the gradiometer signals due to visualization purposes. We observe that with the regularization parameters used, the evoked components predicted by the sources are mostly before 150 ms, including the P50 and N100.

For comparison with alternative solvers, we show in Figure 7 the TF-MxNE source locations superposed to dSPM reconstructions and LCMV beamformer [44] outputs 100 ms after stimulation. Since LCMV does not exploit the cortical orientation information, the map peaks are located at the gyri, below the auditory cortex in the temporal lobe. The LCMV solution was here computed using only the left temporal channels in order to avoid running a beamformer on correlated sources, here coming from bilateral auditory cortices.

5.2.2. Visual data

We used the same model parameters (loose orientation, depth bias, scalar weighting, λ_{time} and Gabor dictionary) as for the auditory condition. The spatial regularization was however changed to $\lambda_{\text{space}} = 30\%$ of $\lambda_{\text{space}}^{\max}$. Results are presented in Figure 8.

Figure 8(a) presents the raw evoked response, restricted to the gradiometers. Sources reconstructions lead to three dipoles. According the automatic parcellation of the cortex provided by FreeSurfer, two sources are localized in the early visual cortex V1, while a third one is positioned on the dorsal part of the secondary visual cortex (V2). One can then quantify the temporal latencies. Between the largest peaks for the early V1 sources the latency is zero which suggests that it could be the same source not properly modeled with a single cortically constraint dipole. Between the V1 sources and the later V2 source the latency is equal to 7 ms, which is very reasonable according to the literature [5]. Here again one can observe the source time series are truely set to zero during baseline.

Using the implementation details provided in Section 4.1, the computation on each dataset presented in this section takes about 30 seconds on a standard laptop computer.

For comparison with alternative solvers, we show in Figure 9 the TF-MxNE source locations superposed to dSPM reconstructions and LCMV beamformer [44] outputs at the time instants corresponding to the two peaks in the evoked response. For LCMV, the data covariance was computed from 40 ms to 150 ms. One can observe that for the early peak all three methods agree to locate the main source in V1 along the calcarine fissure, although dSPM and LCMV lead to more smeared maps. For the second peak, it is pleasing to see how well LCMV and TF-MxNE agree on the site of the blue source, located in the dorsal part of V2 according to the FreeSurfer parcellation.

6. Discussion

While time-frequency analysis is commonly used in the context of M/EEG both in the sensor and source space, it has rarely been used to better model the source dynamics in the context of the inverse problem. Some earlier contributions, such as [10, 13, 26, 28], apply a two-step approach. First TF atoms are estimated from sensor data, typically with greedy approaches like Matching Pursuit. Subsequently, the inverse problem is solved on the selected

components using parametric [13], scanning [26] or distributed methods [10, 28]. Such methods suffer from several limitations. They implicitly assume that the source waveforms correspond to single TF atoms, while real brain signals are better modeled by a combination of atoms as proposed in this contribution. In addition, estimation errors made in the first step have a direct impact on the accuracy of the sources estimates. This is a particularly critical issue since the first step does not take into account the biophysics of the problem, *i.e.*, the solution of the forward problem.

Certainly motivated by the ability of dipole fitting methods to explain M/EEG data, spatial sparsity of source configurations has also been a recurrent assumption to improve the performance of the M/EEG distributed source models. However, variational formulations based on mixed-norms [34, 15], or Bayesian methods with sparsity inducing mechanisms [11, 46] that have been proposed so far in the literature do not model the transient oscillatory nature of brain signals. These approaches make the strong assumption that the source time courses are stationary. For example, results are strongly dependent on the time interval considered. Also the solutions obtained by these solvers are invariant with respect to the permutation of the columns of $\mathbf{M}$, *i.e.*, the temporal sequence of the data is immaterial. This contribution does not suffer from this limitation by explicitly modeling the temporal dynamics of the signals.

To improve over MNE type methods, by removing the assumption of a constant SNR over time and to correlate the time instants together, one should mention recently proposed state-space models based on Kalman filters and smoothers [29, 25]. Such methods, although promising, are still computationally demanding and could suffer from more practical issues like rather long time series to estimate the latent parameters reliably and the necessity to work with fixed orientations.

In [40], an inverse solver that models the transient and non-stationary responses in M/EEG is proposed. A probabilistic model with wavelet shrinkage is employed to promote spatially smooth time courses. The estimation however relies on model approximations with no guarantee on the solution obtained. The most related work to ours, beyond the field of M/EEG, is probably [31] where sparsity is also promoted on the TF decompositions. The related optimization problem, is however solved with a truncated Newton method which can only be applied to differentiable problems. The non-differentiability of the cost function is tackled by using smooth approximation in the minimization. Moreover, Newton methods are known to be fast in the

neighborhood of the solution, but little is known about the global convergence rate. In [33], it is proved that a suitable Newton technique has the same rate of convergence as the accelerated first order schemes like the one we are employing in TF-MxNE. In this contribution, we do not address the problem of learning spatial basis functions as in [43] as doing so would make the cost function non-convex, which would deteriorate the speed of convergence and would also make the solver dependent on the initialization. However, using a pre-defined dictionary of spatial basis functions in line with [3, 21] and multiplying the gain matrix with this dictionary, our prior could be used to estimate spatially extended sources with temporally smooth waveforms. This would, however, be significantly more computationally expensive.

In this work, we demonstrated how physiologically motivated priors for brain activations can be accounted for in a mathematically principled framework in M/EEG source analysis. Using a composite prior, the sparsity of spatial patterns, the temporal smoothness, and the non-stationarity of the source signals were well recovered. Thanks to the structure of the cost function considered, mainly its convexity, an efficient optimization strategy was proposed. The problem being convex, the solver is not affected by improper initialization and cannot be trapped in local minima. Simulations indicated benefits of the approach over alternative solvers, while results with well understood MEG data confirm the accuracy of the reconstruction with real signals. Both results show that our solver is a promising new approach for mining M/EEG data.

Acknowledgements. This work was supported by National Center for Research Resources P41 RR014075-11, the National Institute of Biomedical Imaging and Bioengineering grants 5R01EB009048 and R01 EB006385, the German Research Foundation (Ha 2899/8-2), and the French ANR ViMAGINEANR-08-BLAN-0250-02.

References

[1] Balian, R., 1981. Un principe d'incertitude en théorie du signal ou en mécanique quantique. Compte Rendu de l'Académie des Sciences, Paris, Série 2 292.

[2] Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences 2 (1), 183–202.

[3] Bolstad, A., Veen, B. V., Nowak, R., Jul 2009. Space-time event sparse penalization for magneto-/electroencephalography. NeuroImage 46 (4), 1066–81.

[4] Chen, S., Donoho, D., Saunders, M., 1998. Atomic decomposition by basis pursuit. SIAM Journal on Scientific Computing 20 (1), 33–61.

[5] Cottereau, B., Lorenceau, J., Gramfort, A., Clerc, M., Thirion, B., Baillet, S., 2011. Phase delays within visual cortex shape the response to steady-state visual stimulation. NeuroImage 54 (3), 1919 – 1929.

[6] Dale, A., Liu, A., Fischl, B., Buckner, R., 2000. Dynamic statistical parametric neurotechnique mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. Neuron 26, 55–67.

[7] Dale, A., Sereno, M., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. J. Cogn. Neurosci. 5 (2), 162–176.

[8] Daubechies, I., 1992. Ten lectures on Wavelets. SIAM-CBMS Conferences Series.

[9] Donoho, D., may 1995. De-noising by soft-thresholding. Information Theory, IEEE Transactions on 41 (3), 613 –627.

[10] Durka, P. J., Matysiak, A., Montes, E. M., Valdés-Sosa, P., Blinowska, K. J., 2005. Multichannel matching pursuit and EEG inverse solutions. Journal of Neuroscience Methods 148 (1), 49 – 59.

[11] Friston, K., Harrison, L., Daunizeau, J., Kiebel, S., Phillips, C., Trujillo-Barreto, N., Henson, R., Flandin, G., Mattout, J., Feb 2008. Multiple sparse priors for the M/EEG inverse problem. Neuroimage 39 (3), 1104–20.

[12] Gabor, D., 1946. Theory of communication. J. IEEE 93, 429 – 457.

[13] Geva, A. B., 1996. Spatio-temporal matching pursuit (SToMP) for multiple source estimation of evoked potentials. In: Electrical and Electronics Eng. pp. 113 –116.

[14] Gorodnitsky, I., George, J., Rao, B., Jan 1995. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. Electroencephalography and Clinical Neurophysiology 95 (4), 231–251.

[15] Gramfort, A., Kowalski, M., Hämäläinen, M., Mar 2012. Mixed-norm estimates for the M/EEG inverse problem using accelerated gradient methods. Physics in Medicine and Biology 57 (7), 1937–1961.

[16] Gramfort, A., Papadopoulo, T., Olivi, E., Clerc, M., 2010. OpenMEEG: opensource software for quasistatic bioelectromagnetics. BioMed Eng OnLine 9 (1), 45.

[17] Gramfort, A., Strohmeier, D., Haueisen, J., Hämäläinen, M., Kowalski, M., 2011. Functional brain imaging with m/eeg using structured sparsity in time-frequency dictionaries. In: Székely, G., Hahn, H. (Eds.), Information Processing in Medical Imaging. Vol. 6801 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 600–611.

[18] Grosenick, L., Klingenberg, B., Knutson, B., Taylor, J. E., 2011. A family of interpretable multivariate models for regression and classification of whole-brain fMRI data. pre-print.

[19] Hämäläinen, M., Ilmoniemi, R., Jan 1994. Interpreting magnetic fields of the brain: minimum norm estimates. Med Biol Eng Comput 32 (1), 35–42.

[20] Haufe, S., Nikulin, V. V., Ziehe, A., Müller, K.-R., Nolte, G., Aug 2008. Combining sparsity and rotational invariance in EEG/MEG source reconstruction. NeuroImage 42 (2), 726–38.

[21] Haufe, S., Tomioka, R., Dickhaus, T., Sannelli, C., Blankertz, B., Nolte, G., Müller, K.-R., 2011. Large-scale EEG/MEG source localization with spatial flexibility. NeuroImage 54 (2), 851–859.

[22] Jaros, U., Hilgenfeld, B., Lau, S., Curio, G., Haueisen, J., 2008. Nonlinear interactions of high-frequency oscillations in the human somatosensory system. Clin Neurophysiol 119 (11), 2647–57.

[23] Jenatton, R., Mairal, J., Obozinski, G., Bach, F., Jul. 2011. Proximal methods for hierarchical sparse coding. The Journal of Machine Learning Research 12, 2297–2334.

[24] Kybic, J., Clerc, M., Abboud, T., Faugeras, O., Keriven, R., Papadopoulo, T., 2005. A common formalism for the integral formulations of the forward EEG problem. IEEE Transactions on Medical Imaging 24 (1), 12–28.

[25] Lamus, C., Hmlinen, M. S., Temereanca, S., Brown, E. N., Purdon, P. L., 2012. A spatiotemporal dynamic distributed solution to the meg inverse problem. NeuroImage 63 (2), 894 – 909.

[26] Lelic, D., Gratkowski, M., Valeriani, M., Arendt-Nielsen, L., Drewes, A. M., 2009. Inverse modeling on decomposed electroencephalographic data: A way forward? Journal of Clinical Neurophysiology 26 (4), 227–235.

[27] Lin, F., Belliveau, J., Dale, A., Hämäläinen, M., 2006. Distributed current estimates using cortical orientation constraints. Hum Brain Mapp 27, 1–13.
URL `http://doi.wiley.com/10.1002/hbm.20155`

[28] Lina, J., Chowdhury, R., Lemay, E., Kobayashi, E., Grova, C., 2012. Wavelet-based localization of oscillatory sources from magnetoencephalography data. Biomedical Engineering, IEEE Transactions on PP (99), 1.

[29] Long, C. J., Purdon, P. L., Temereanca, S., Desai, N. U., Hämäläinen, M. S., Brown, E. N., 2011. State-space solutions to the dynamic magnetoencephalography inverse problem using high performance computing. Annals of Applied Statistics 5 (2B), 1207–1228.

[30] Matsuura, K., Okabe, Y., June 1995. Selective minimum-norm solution of the biomagnetic inverse problem. IEEE Trans Biomed Eng 42 (6), 608–615.

[31] Model, D., Zibulevsky, M., 2006. Signal reconstruction in sensor arrays using sparse representations. Signal Processing 86 (3), 624 – 638.

[32] Mosher, J., Leahy, R., Lewis, P., 1999. EEG and MEG: Forward solutions for inverse methods. IEEE Transactions on Biomedical Engineering 46 (3), 245–259.

[33] Nesterov, Y., Polyak, B., 2006. Cubic regularization of newton's method and its global performance. Mathematical Programming 108 (1), 177–205.

[34] Ou, W., Hämaläinen, M., Golland, P., Feb 2009. A distributed spatio-temporal EEG/MEG inverse solver. NeuroImage 44 (3), 932–946.

[35] Pascual-Marqui, R., 2002. Standardized low resolution brain electromagnetic tomography (sLORETA): technical details. Methods Find. Exp. Clin. Pharmacology 24 (D), 5–12.

[36] Ramírez, R., Kopell, B., Butson, C., Hiner, B., Baillet, S., May 2011. Spectral signal space projection algorithm for frequency domain MEG and EEG denoising, whitening, and source imaging. Neuroimage 56 (1), 78 – 92.

[37] Scherg, M., Von Cramon, D., Jan. 1985. Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. Electroencephalogr Clin Neurophysiol 62 (1), 32–44.

[38] Soendergard, P., Torrésani, B., Balazs, P., 2009. The linear time frequency toolbox. Tech. rep., Technical University of Denmark.

[39] Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Serie B 58 (1), 267–288.

[40] Trujillo-Barreto, N. J., Aubert-Vázquez, E., Penny, W. D., 2008. Bayesian M/EEG source reconstruction with spatio-temporal priors. Neuroimage 39 (1), 318–35.

[41] Uusitalo, M., Ilmoniemi, R., 1997. Signal-space projection method for separating meg or eeg into components. Medical and Biological Engineering and Computing 35 (2), 135–140.

[42] Uutela, K., Hämäläinen, M., Somersalo, E., 1999. Visualization of magnetoencephalographic data using minimum current estimates. Neuroimage 10, 173–180.

[43] Valdés-Sosa, P. A., Vega-Hernández, M., Sánchez-Bornot, J. M., Martínez-Montes, E., Bobes, M. A., Jun 2009. EEG source imaging with spatio-temporal tomographic nonnegative independent component analysis. HBM 30 (6), 1898–910.

[44] Veen, B. V., Drongelen, W. V., Yuchtman, M., Suzuki, A., Jan 1997. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. Biomedical Engineering, IEEE Transactions on 44 (9), 867—880.

[45] Wang, J.-Z., Williamson, S. J., Kaufman, L., July 1992. Magnetic source images determined by a lead-field analysis: the unique minimum-norm least-squares estimation. Biomedical Engineering, IEEE Transactions on 39 (7), 665–675.

[46] Wipf, D., Nagarajan, S., Feb 2009. A unified Bayesian framework for MEG/EEG source imaging. Neuroimage 44 (3), 947–966.

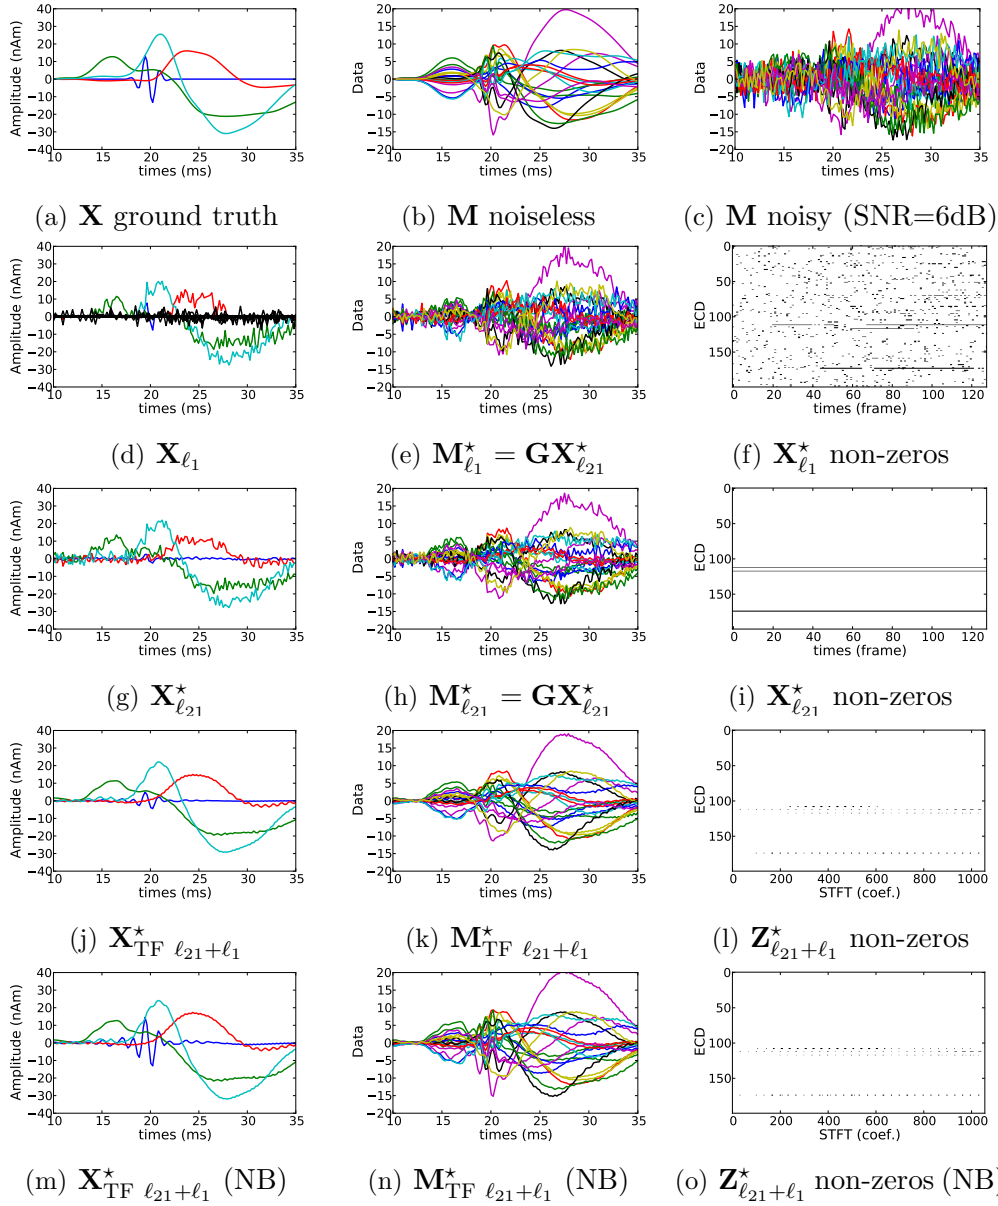

(a) $\mathbf{X}$ ground truth (b) $\mathbf{M}$ noiseless (c) $\mathbf{M}$ noisy (SNR=6dB)

(d) $\mathbf{X}_{\ell_1}$ (e) $\mathbf{M}^\star_{\ell_1} = \mathbf{G}\mathbf{X}^\star_{\ell_{21}}$ (f) $\mathbf{X}^\star_{\ell_1}$ non-zeros

(g) $\mathbf{X}^\star_{\ell_{21}}$ (h) $\mathbf{M}^\star_{\ell_{21}} = \mathbf{G}\mathbf{X}^\star_{\ell_{21}}$ (i) $\mathbf{X}^\star_{\ell_{21}}$ non-zeros

(j) $\mathbf{X}^\star_{\text{TF } \ell_{21}+\ell_1}$ (k) $\mathbf{M}^\star_{\text{TF } \ell_{21}+\ell_1}$ (l) $\mathbf{Z}^\star_{\ell_{21}+\ell_1}$ non-zeros

(m) $\mathbf{X}^\star_{\text{TF } \ell_{21}+\ell_1}$ (NB) (n) $\mathbf{M}^\star_{\text{TF } \ell_{21}+\ell_1}$ (NB) (o) $\mathbf{Z}^\star_{\ell_{21}+\ell_1}$ non-zeros (NB)

Figure 4: Simulations results with $SNR = 6\ dB$. (a) A simulated source activations. (b) Noiseless simulated measurements. (c) Simulated measurements corrupted by noise. (d-e-f) Estimation with ℓ_1 prior. (g-h-i) Estimation with ℓ_{21} prior [34]. (j-k-l) Estimation with composite TF prior. (m-n-o) Estimation with composite TF prior and debiasing. (f-i-l-o) show the sparsity patterns obtained by the 3 different priors as explained in Fig. 2. Result (j) shows how the composite TF prior improves over (d) and (g). (l) presents also a higher level of sparsity compared to (f) and (i).

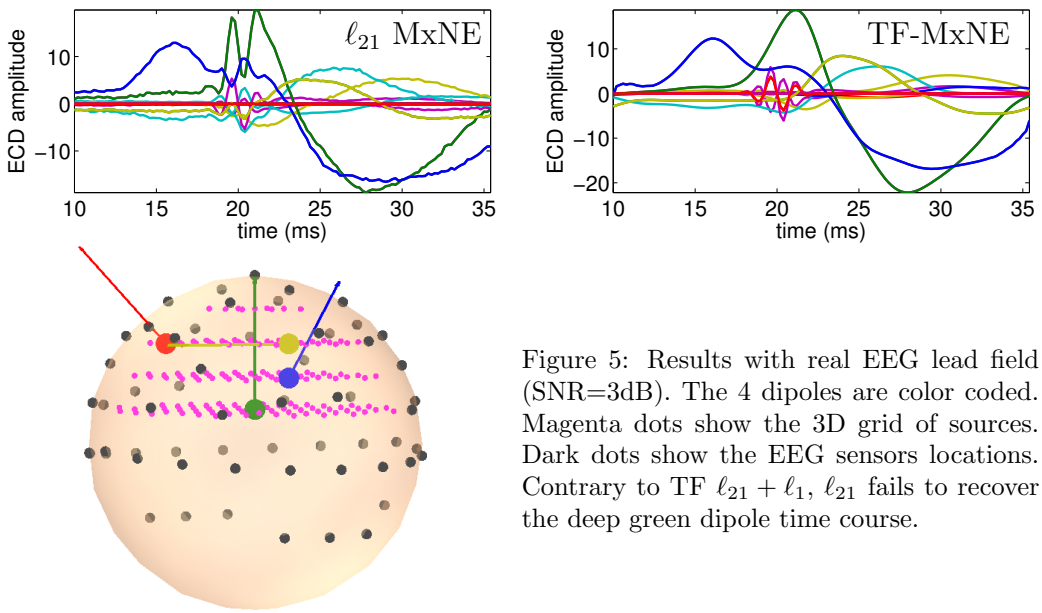

Figure 5: Results with real EEG lead field (SNR=3dB). The 4 dipoles are color coded. Magenta dots show the 3D grid of sources. Dark dots show the EEG sensors locations. Contrary to TF $\ell_{21} + \ell_1$, ℓ_{21} fails to recover the deep green dipole time course.

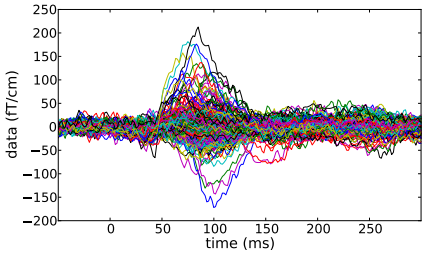

(a) MEG data (Gradiometers only)

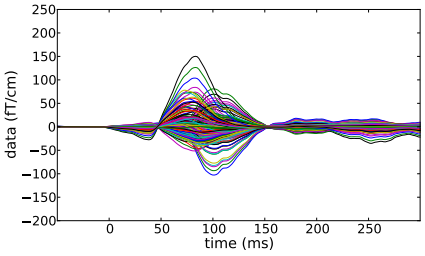

(b) $\mathbf{G}\mathbf{X}^{\star}_{\text{TF-MxNE}}$ (explained data)

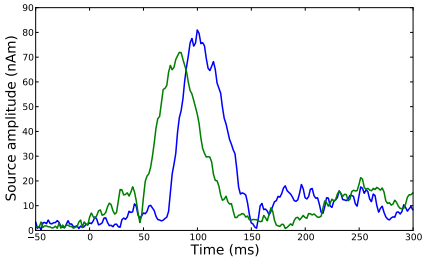

(c) $\mathbf{X}^{\star}_{\text{MxNE}}$

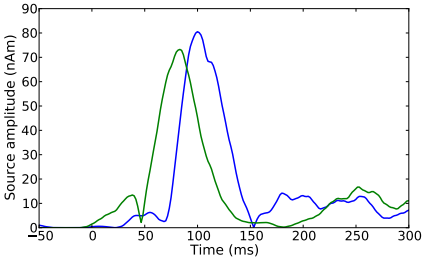

(d) $\mathbf{X}^{\star}_{\text{TF-MxNE}}$

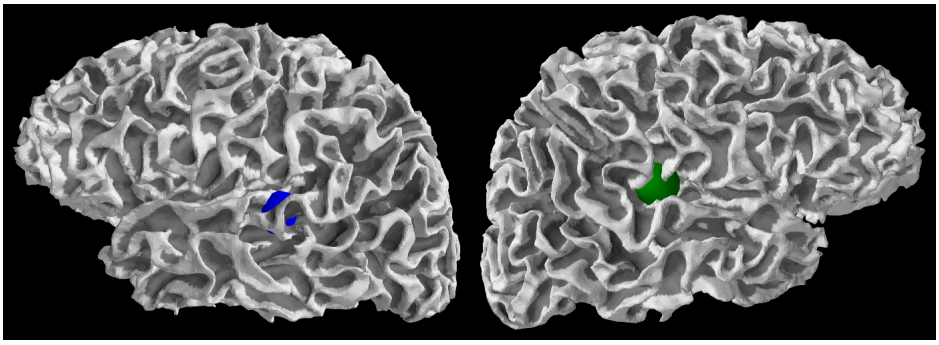

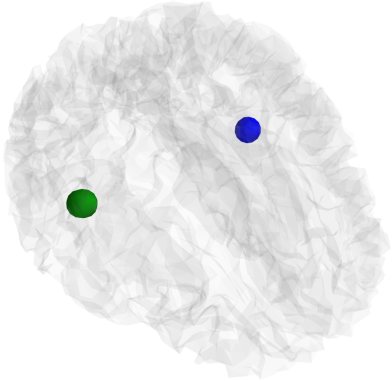

(e) Sources in bilateral auditory cortices

Figure 6: Results obtained with TF-MxNE and MxNE for left-ear auditory stimulation with unfiltered combined MEG/EEG data. Estimation was performed with the loose orientation parameter 0.2, with depth compensation of 0.9 on a set of 7498 cortical locations ($\mathbf{G} \in \mathbb{R}^{364\times 22494}$). The estimation with $\lambda_{\text{space}} = 50\%$ of $\lambda^{\max}_{\text{space}}$ leads to 2 active brain locations at the auditory cortices. TF-MxNE leads to smooth time courses and zeros during baseline.

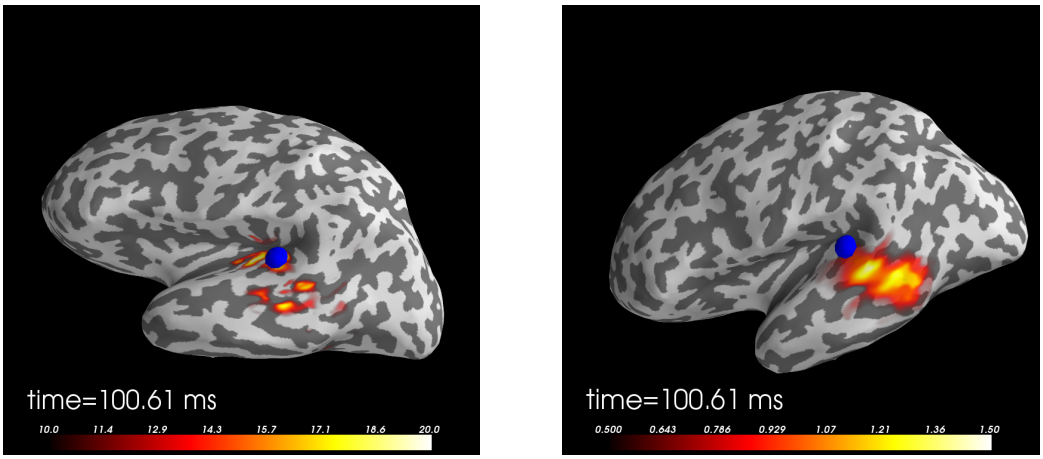

(a) TF-MxNE with dSPM estimates

(b) TF-MxNE with LCMV maps

Figure 7: Comparison of TF-MxNE, dSPM, and LCMV for an auditory stimulation. The images show the TF-MxNE solutions overlayed with dSPM estimates (a) and LCMV output map (b) 100 ms after stimulation. Since LCMV does not exploit the cortical orientation information, the map peaks are located at the gyri, below the auditory cortex in the temporal lobe.

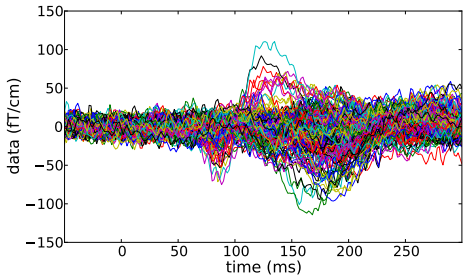

(a) MEG data (Gradiometers only)

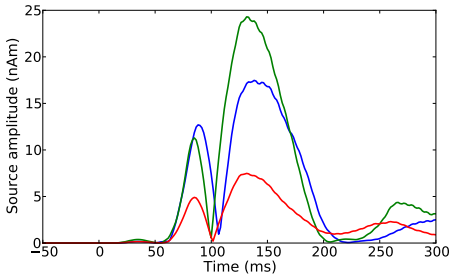

(b) $\mathbf{X}^{\star}_{\text{TF-MxNE}}$

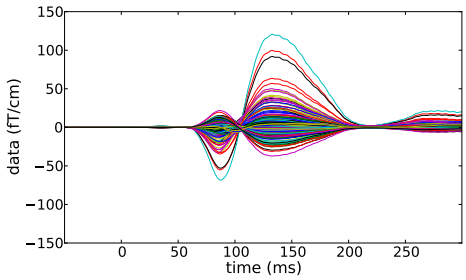

(c) $\mathbf{G}\mathbf{X}^{\star}_{\text{TF-MxNE}}$ (explained data)

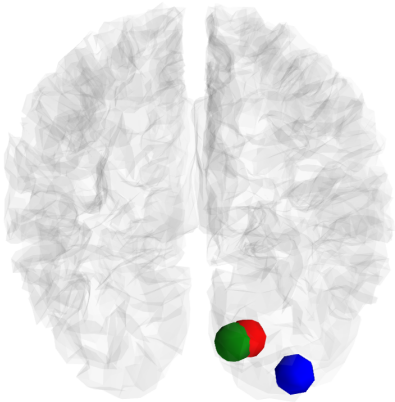

(d) Source locations in the right occipital cortex (visual cortex)

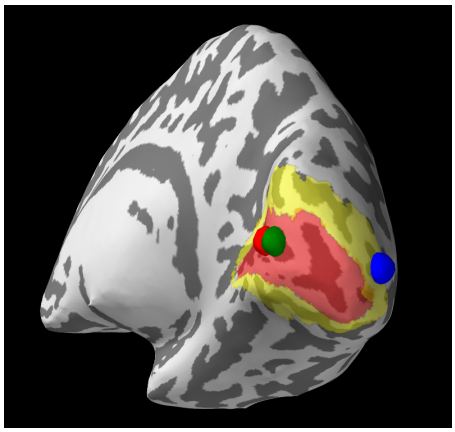

(e) Source locations in right hemisphere with V1 (red) and V2 (yellow) labels.

Figure 8: Results obtained with TF-MxNE for a visual stimulation (checkerboard stimulus in the left visual field) with unfiltered MEG data. Estimation was performed with a loose orientation (parameter 0.2), with a depth compensation of 0.9 on a set of 7498 cortical locations ($\mathbf{G} \in \mathbb{R}^{364\times 22494}$). Estimation with $\lambda_{\text{space}} = 30\%$ of $\lambda^{\text{max}}_{\text{space}}$ leads to 3 active brain locations in the contralateral visual cortex (V1 and V2). TF-MxNE leads to smooth time courses and zeros during baseline.

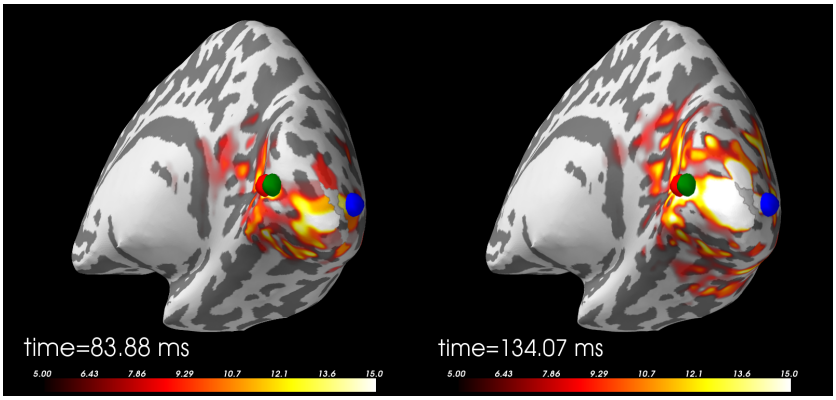

(a) TF-MxNE source locations overlayed with dSPM estimates

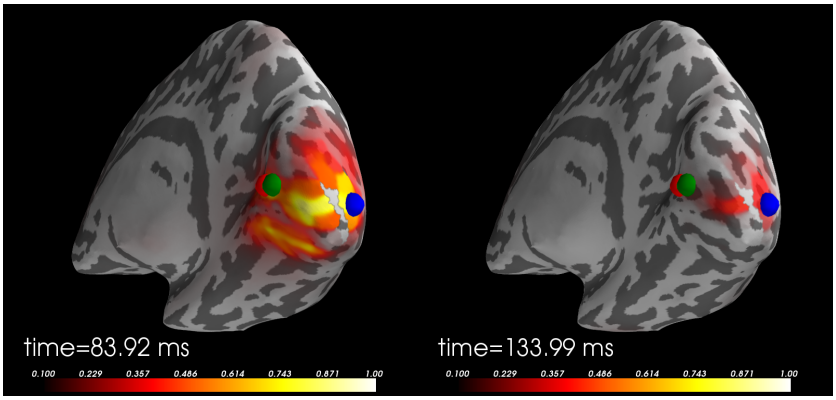

(b) TF-MxNE source locations overlayed with LCMV maps

Figure 9: Comparison of TF-MxNE, dSPM, and LCMV for a visual stimulation (checkerboard stimulus in the left visual field). The images show the TF-MxNE solutions overlayed with dSPM estimates (a) and LCMV output maps (b) for two different time points corresponding to the main peaks in Figure 8(a). For the early peak, there is a very good agreement between all three methods although dSPM and LCMV maps are smeared in space. For the later peak, LCMV and TF-MxNE agree on the location of the blue source.