

PAPER • OPEN ACCESS

A hierarchical Bayesian perspective on majorization-minimization for non-convex sparse regression: application to M/EEG source imaging

To cite this article: Yousra Bekhti *et al* 2018 *Inverse Problems* **34** 085010

View the [article online](#) for updates and enhancements.

Related content

- [Mixed-norm estimates for the M/EEG inverse problem using accelerated gradient methods](#)
Alexandre Gramfort, Matthieu Kowalski and Matti Hämäläinen
- [Fast Markov chain Monte Carlo sampling for sparse Bayesian inference in high-dimensional inverse problems using L1-type priors](#)
Felix Lucka
- [Sequential Monte Carlo samplers for semi-linear inverse problems and application to magnetoencephalography](#)
Sara Sommariva and Alberto Sorrentino



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

A hierarchical Bayesian perspective on majorization-minimization for non-convex sparse regression: application to M/EEG source imaging

Yusra Bekhti¹, Felix Lucka^{2,3}, Joseph Salmon¹
and Alexandre Gramfort^{1,4}

¹ LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France

² Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, Netherlands

³ Centre for Medical Image Computing, University College London, London, WC1E 6BT, United Kingdom

⁴ INRIA, Parietal team, Université Paris-Saclay, Saclay, France

E-mail: yusra.bekhti@telecom-paristech.fr

Received 8 December 2017, revised 4 May 2018

Accepted for publication 1 June 2018

Published 28 June 2018



CrossMark

Abstract

Majorization-minimization (MM) is a standard iterative optimization technique which consists in minimizing a sequence of convex surrogate functionals. MM approaches have been particularly successful to tackle inverse problems and statistical machine learning problems where the regularization term is a sparsity-promoting concave function. However, due to non-convexity, the solution found by MM depends on its initialization. Uniform initialization is the most natural and often employed strategy as it boils down to penalizing all coefficients equally in the first MM iteration. Yet, this arbitrary choice can lead to unsatisfactory results in severely under-determined inverse problems such as source imaging with magneto- and electro-encephalography (M/EEG). The framework of hierarchical Bayesian modeling (HBM) is an alternative approach to encode sparsity. This work shows that for certain hierarchical models, a simple alternating scheme to compute fully Bayesian maximum *a posteriori* (MAP) estimates leads to the exact same sequence of updates as a standard MM strategy (see the adaptive lasso). With this parallel outlined, we show how to improve upon these MM techniques by probing the multimodal posterior density using Markov Chain



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Monte-Carlo (MCMC) techniques. Firstly, we show that these samples can provide well-informed initializations that help MM schemes to reach better local minima. Secondly, we demonstrate how it can reveal the different modes of the posterior distribution in order to explore and quantify the inherent uncertainty and ambiguity of such ill-posed inference procedure. In the context of M/EEG, each mode corresponds to a plausible configuration of neural sources, which is crucial for data interpretation, especially in clinical contexts. Results on both simulations and real datasets show how the number or the type of sensors affect the uncertainties on the estimates.

Keywords: majorization-minimization, hierarchical Bayesian modeling, sparse models, M/EEG inverse problems, uncertainty quantification, MCMC sampling

(Some figures may appear in colour only in the online journal)

1. Introduction

Over the last two decades, sparsity has emerged as a key concept to solve inverse problems such as tomographic image reconstruction, deconvolution or inpainting, but also to regularize high dimensional regression problems in the field of machine learning. There are mainly two routes to introduce sparsity in such problems.

The first route, embraced by the optimization community and frequentist statisticians, is to promote sparsity using convex optimization theory. This line of work has led to now mature theoretical guarantees (Foucart and Rauhut 2013) when using regularization functions based on ℓ_1 -norm and other convex variants (Tibshirani 1996). In particular, it has been popularized in the signal processing community under the name of compressed sensing (Candès and Wakin 2008) when combined with incoherent measurements.

There are however some limitations of sparsity-promoting convex penalties based on the ℓ_1 -norm. All the features (also called regressors, atoms or sources depending on the terminology of the community) involved in the solution form what is called the support of the solution. Convex penalties can fail to identify the correct support in the presence of highly noisy data, but also in low noise setups if the forward operator (referred to as design matrix in statistics) is poorly conditioned. Convex regularizations also lead to a systematic underestimation bias in the amplitude of the coefficients (Osher *et al* 2006, Chartrand 2007, Candès *et al* 2008, Saab *et al* 2008, Chzhen *et al* 2017).

To address these limitations of ℓ_1 -type models, re-weighted schemes have been proposed (Candès *et al* 2008, Gasso *et al* 2009, Rakotomamonjy 2011, Zhang and Rao 2011, Strohmeier *et al* 2016), of which the adaptive lasso (Zou 2006) is the most commonly used in the statistics community: starting from the Lasso estimator, which amounts to regressing with a standard ℓ_1 -norm as a regularizer (this estimator is sometimes referred to as basis pursuit denoising (BPDN) (Chen *et al* 1998) in signal processing), the adaptive lasso solves a sequence of weighted Lasso problems, where at each iteration the weights are chosen such that the strongest coefficients are less and less penalized. From the optimization point of view, such an iterative scheme can be derived from so-called majorization-minimization (MM) strategies (Lange *et al* 2000, Schifano *et al* 2010). The idea behind MM is to minimize the objective function by successively minimizing upper bounds that are easier to optimize. Many well-known optimization approaches can be interpreted as instances of MM, e.g. simple gradient descent

or proximal algorithms (Combettes and Pesquet 2011), expectation-maximization (EM) (Dempster *et al* 1977), and difference-of-convex (DC) programming techniques (Horst and Thoai 1999). More recently, re-weighted ℓ_1 -norm schemes based on MM principle have been particularly popular to handle concave, hence non-convex regularizations such as $\ell_{0.5}$ -quasi-norms or logarithmic functions. As such, these schemes are prone to converging to a local minimum determined by the initial, uniformly weighted ℓ_1 -norm solution (i.e. the Lasso estimator) that constitutes the first iterate.

The common way to formulate hierarchical Bayesian modeling (HBM) is to consider the variance parameters of Gaussian prior models as additional random variables which have to be estimated from the data as well. Their prior distributions are referred to as hyper-priors. Plausible solutions to the regression problem, that both fit data and the *a priori* assumption of sparsity, are explicitly characterized as multiple distinct modes of the posterior distribution. This characterization is the Bayesian analogue to local minima in variational regression approaches when working with non-convex functionals. To infer a point estimate of the parameters of interest from the *a posteriori* distribution different strategies exist. For instance variational Bayesian approaches (Jordan *et al* 1999, MacKay 2003, Sato *et al* 2004, Friston *et al* 2008, Shervashidze and Bach 2015), sparse Bayesian learning (SBL) approaches (also referred to as type-I or type-II maximum likelihood estimates) (Tipping 2001, Wipf and Rao 2004, Wipf and Nagarajan 2009, Zhang and Rao 2011) and fully-Bayesian strategies (Calvetti *et al* 2009, Lucka *et al* 2012) are possible. In this work, we focus on the latter one for a non-standard type of HBM examined in Lucka (2014) that combines a non-Gaussian prior with an ℓ_1 -type energy function with a specific Gamma hyper-prior. Interestingly, for this HBM, a simple alternating scheme to compute full maximum *a posteriori* (MAP) estimates leads to exactly the same sequence of problems solved by MM applied to $\ell_{1/2}$ -type regularizations. In other words, the adaptive lasso estimator (Zou 2006) commonly used in machine learning is tightly related to this HBM model. With this observation made, it is natural to revisit and improve these MM schemes by leveraging the ability of the Bayesian framework to explore the modes of the posterior distribution by Markov chain Monte-Carlo (MCMC) schemes (Kaipio and Somersalo 2005, Robert and Casella 2005). This can not only mitigate the aforementioned initialization-dependence of MM, but more importantly, it can offer insights into the structure and importance of potentially multiple plausible sparse solutions. Yet, the benefit comes at the cost of additional computational efforts.

Magnetoencephalography and electroencephalography (M/EEG) are technologies that allow to measure the electromagnetic fields produced by active neurons in a non-invasive way. Localization of foci of neural activations from M/EEG recordings is a high impact problem both for cognitive neuroscience and clinical neuroscience, with applications in pathologies such as epilepsy, sleep research or neurodegenerative disorders. Despite the linearity of the forward problem, this inverse problem is particularly challenging as the forward operator is both under-determined and strongly ill-conditioned. As such, both non-convex optimization strategies with re-weighted schemes (Strohmeier *et al* 2016) and hierarchical Bayesian approaches have been proposed (Sato *et al* 2004, Calvetti *et al* 2009, Sorrentino *et al* 2009, Wipf and Nagarajan 2009, Lucka *et al* 2012) for M/EEG source localization. For this reason, it is an ideal application for our examinations.

The manuscript is organized as follows: first, we present a unified perspective on both routes to sparsity, i.e. re-weighted ℓ_1 MM schemes and specific HBMs. We show that a particular optimization-based inference strategy recovers the MM algorithm. We then describe a HBM inference strategy based upon an MCMC sampling and show on simulated and experimental M/EEG datasets how these stochastic MCMC-based techniques can not only help to improve upon deterministic approaches, but also help to reveal multiple plausible solutions to

the inverse problem. This analysis leads to an uncertainty quantification (UQ) of the support recovery of non-convex sparse regression problems that provides very useful complementary information, in particular for very ill-conditioned and under-determined applications like M/EEG source localization.

Notation. We consider here the following linear model, known in machine learning as multi-task regression and in signal processing as the multiple measurement vector (MMV) model (Cotter *et al* 2005):

$$\mathbf{M} = \mathbf{G}\mathbf{X} + \mathbf{E}, \quad (1)$$

where $\mathbf{M} \in \mathbb{R}^{m \times t}$. In machine learning m corresponds to the number of samples and t to the number of tasks, and in our application m corresponds to the number of sensors and t to the number of measurements over time. The matrix $\mathbf{G} \in \mathbb{R}^{m \times q}$ is the design matrix, a known instantaneous mixing matrix also referred to as the *forward*, *gain* or *system* matrix. It relates the unknown coefficients $\mathbf{X} \in \mathbb{R}^{q \times t}$, which correspond to amplitudes of neural sources, to the measurements \mathbf{M} . In our application $q = dn$ where n is the number of source locations and d is the number of oriented sources per location ($d = 1$ or $d = 3$ corresponding to estimating a scalar field or a 3D vector field of sources). The matrix \mathbf{E} models the measurement noise, which is assumed to be an *additive, white Gaussian noise* (AWGN). This is a reasonable assumption after performing a proper spatial whitening of the data using an estimate of the noise covariance (Engemann and Gramfort 2015). By $\mathbf{X}_{(i,j)}$ we refer to the entry in the i th row and j th column in \mathbf{X} , while $\mathbf{X}_{(i,:)} \in \mathbb{R}^t$ and $\mathbf{X}_{(:,j)} \in \mathbb{R}^q$ refer to the complete i th row and j th column, respectively. In addition, we denote by $\mathbf{X}_{[i]}$, $i = 1, \dots, n$ the $d \times t$ sub-matrix of \mathbf{X} corresponding to the i th *group*: for $d = 1$, this coincides with $\mathbf{X}_{(i,:)}$, while for $d = 3$, $\mathbf{X}_{[i]} = [\mathbf{X}_{((i-1)d+1,:)}^\top, \mathbf{X}_{((i-1)d+2,:)}^\top, \mathbf{X}_{((i-1)d+3,:)}^\top]^\top$. Note that thereby, the group size is dt . Negative indices are used to exclude certain entries, rows, columns or groups, i.e. $\mathbf{X}_{(:, -j)} \in \mathbb{R}^{q \times (t-1)}$ refers to the matrix obtained by deleting the j th column of \mathbf{X} . Furthermore, let I_m denote the identity matrix of size $m \in \mathbb{N}$. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, the Frobenius norm is given by $\|\mathbf{A}\|_F^2 = \sum_{i,j} \mathbf{A}_{(i,j)}^2$.

2. Methods

We start this section by recalling how MM works when addressing variational formulations with concave, non-convex, regularization. It is followed by an introduction to hierarchical Bayesian models with Gamma hyper-priors. Then, we explain how these seemingly different approaches can lead to the exact same optimization algorithm. From this, we detail how different Bayesian inference strategies using MCMC sampling can more precisely explore the landscape of the posterior distribution of the HBM model, as well as provide multiple possible solutions to the sparse regression problem.

2.1. MM

MM strategies consist in replacing a difficult optimization problem with a series of easier ones that are obtained by upper bounding the objective function, often by a convex majorant. In the context of inverse problems or high-dimensional statistics using sparsity constraints, MM has been successfully applied to address non-convex regularization terms. An example is the

regression model with $\ell_{2,p}$ -quasi-norms regularization when $0 < p < 1$. The desired estimate $\hat{\mathbf{X}}$ is defined as one of potentially multiple minimizers of

$$\hat{\mathbf{X}} \in \arg \min_{\mathbf{X} \in \mathbb{R}^{q \times r}} \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_{\text{F}}^2 + \lambda \sum_{i=1}^n \|\mathbf{X}_{[i]}\|_{\text{F}}^p, \quad (2)$$

where $\lambda > 0$ is the regularization parameter balancing the data fit and the penalty term. One possible MM approach to solve equation (2) with $p = 1/2$ would consist in minimizing a sequence of non-smooth convex surrogate functions where the non-convex regularization is replaced by a weighted $\ell_{2,1}$ norm (Strohmeier *et al* 2016). In each iteration, the weights are derived from the current estimate of \mathbf{X} .

Due to the concavity of the non-decreasing function $\mathbf{X} \mapsto \sqrt{\|\mathbf{X}\|_{\text{F}}}$, it is upper bounded by its tangent and a first order Taylor expansion at the current estimate $\mathbf{X}_{[i]}$ provides an upper bound that can be used to construct the non-smooth convex surrogate problem. By solving this sequence of surrogate problems, the value of the non-convex objective function is guaranteed to decrease. However, due to the non-convexity, only convergence towards a local minimum can be guaranteed.

For the problem in equation (2) with $p = 1/2$, the k th iteration of the MM scheme reads:

$$\hat{\mathbf{X}}^{(k)} \in \arg \min_{\mathbf{X} \in \mathbb{R}^{q \times r}} \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_{\text{F}}^2 + \lambda \sum_{i=1}^n \frac{\|\mathbf{X}_{[i]}\|_{\text{F}}}{w_i^{(k-1)}}, \quad w_i^{(k-1)} = 2\sqrt{\|\mathbf{X}_{[i]}^{(k-1)}\|_{\text{F}}}. \quad (3)$$

As each weight $w_i^{(k)}$ is a non-decreasing function of $\|\mathbf{X}_{[i]}^{(k)}\|_{\text{F}}$, sources with high amplitudes in one iteration will be less penalized in the next iteration. Strong sources are more and more promoted to explain the data \mathbf{M} . Sources for which $\|\mathbf{X}_{[i]}^{(k)}\|_{\text{F}} = 0$ at a certain iteration k are effectively pruned from the model for all following iterations. Using MM therefore leads to a solution that explains the data with fewer active locations i compared to a standard $\ell_{2,1}$ norm regularized solution. Note that a default initialization consists in setting $w_i^{(0)} = 1, \forall i \in [n]$.

To exploit existing fast solvers for the $\ell_{2,1}$ regularized problems (Ndiaye *et al* 2015, Strohmeier *et al* 2016), we reformulate the weighted subproblem and apply the weights by scaling the matrix \mathbf{G} with a diagonal matrix $\mathbf{W}^{(k)} \in \mathbb{R}^{dn \times dn}$ given by:

$$\mathbf{W}^{(k)} = \text{diag}(\mathbf{w}^{(k)} \otimes \mathbf{1}_d), \quad (4)$$

where $\mathbf{w}^{(k)} \in \mathbb{R}^n$, $\mathbf{1}_d \in \mathbb{R}^d$ is a vector of ones and \otimes is the Kronecker product. Defining $\tilde{\mathbf{G}}^{(k)} = \mathbf{G}\mathbf{W}^{(k)}$, the reformulated problem reads:

$$\tilde{\mathbf{X}}^{(k)} \in \arg \min_{\mathbf{X} \in \mathbb{R}^{q \times r}} \frac{1}{2} \|\mathbf{M} - \tilde{\mathbf{G}}^{(k)}\mathbf{X}\|_{\text{F}}^2 + \lambda \sum_{i=1}^n \|\mathbf{X}_{[i]}\|_{\text{F}}. \quad (5)$$

After convergence, we reapply the scaling to $\tilde{\mathbf{X}}$ to obtain $\hat{\mathbf{X}}$:

$$\hat{\mathbf{X}}^{(k)} = \mathbf{W}^{(k-1)}\tilde{\mathbf{X}}^{(k)}. \quad (6)$$

The reformulation through equations (5) and (6) avoids any division by zero when $\mathbf{X}^{(k-1)} = 0$. The above procedure, which matches the strategy of the adaptive lasso estimator (Zou 2006), is expressed as pseudo-code in algorithm 1. More technical details can be found in (Strohmeier *et al* 2016, algorithm 3).

Algorithm 1. $\ell_{2,p}$ MM ALGORITHM WITH $p = 1/2$ (ADAPTIVE LASSO).

input : $\mathbf{M}, \mathbf{G}, \lambda > 0, \mathbf{W}^{(0)} \geq 0, \epsilon > 0, \tau > 0$ and K
for $k = 1$ to K **do**
 $\tilde{\mathbf{G}}^{(k)} = \mathbf{G}\mathbf{W}^{(k-1)}$
 Get $\tilde{\mathbf{X}}^{(k)}$ solving equation (5) at ϵ -precision (e.g. by block coordinate descent).
 Update $\hat{\mathbf{X}}^{(k)} = \mathbf{W}^{(k-1)}\tilde{\mathbf{X}}^{(k)}$
 Update $\mathbf{W}^{(k)} = \text{diag}(\mathbf{w}^{(k)} \otimes \mathbf{1}_d)$ where $\mathbf{w}_{[i]}^{(k)} = 2\sqrt{\|\hat{\mathbf{X}}_{[i]}^{(k)}\|_F}$, $\forall i \in [n]$
 if $\|\hat{\mathbf{X}}^{(k)} - \hat{\mathbf{X}}^{(k-1)}\|_\infty \leq \tau$ **then**
 Break
return $\hat{\mathbf{X}}^{(k)}$

2.2. HBM

In this section, we formulate the inference problem given by equation (1) and the regularization strategy with $\ell_{2,p}$ -quasi-norms from a Bayesian perspective (Kaipio and Somersalo 2005, Lucka 2014): the Bayesian approach incorporates prior beliefs about the model parameters in terms of probability distributions. Under the AWGN assumption the likelihood of the model is given by:

$$p_{\text{like}}(\mathbf{M}|\mathbf{X}) \propto \exp\left(-\frac{1}{2}\|\mathbf{M} - \mathbf{G}\mathbf{X}\|_F^2\right). \quad (7)$$

From equation (2) we can construct the $\ell_{2,p}$ group prior as:

$$p_{\text{prior}}(\mathbf{X}) \propto \exp\left(-\lambda \sum_{i=1}^n \|\mathbf{X}_{[i]}\|_F^p\right) = \prod_{i=1}^n \exp(-\lambda \|\mathbf{X}_{[i]}\|_F^p), \quad (8)$$

which leads to the following posterior probability density using Bayes rule:

$$p_{\text{post}}(\mathbf{X}|\mathbf{M}) \propto \exp\left(-\frac{1}{2}\|\mathbf{M} - \mathbf{G}\mathbf{X}\|_F^2 - \lambda \sum_{i=1}^n \|\mathbf{X}_{[i]}\|_F^p\right). \quad (9)$$

To extend equation (8) to a hierarchical prior model (MacKay 2003), we replace the scalar λ by a vector of hyper-parameters $\gamma \in \mathbb{R}_+^n$ and for any $p \geq 1$ we write the *conditional $\ell_{2,p}$ prior* as:

$$p_{\text{prior}}(\mathbf{X}|\gamma) \propto \exp\left(-\sum_{i=1}^n \left(\frac{\|\mathbf{X}_{[i]}\|_F^p}{\gamma_i} + \frac{d}{p} \log(\gamma_i)\right)\right), \quad (10)$$

where the logarithmic term accounts for the terms of the normalization that depend on γ (Lucka 2014). A common choice for a hyper-prior on each γ_i is given by a *Gamma distribution* (MacKay 2003, Kaipio and Somersalo 2005, Calvetti et al 2009, Lucka et al 2012) with shape and scale parameters α and β :

$$p_{\text{hyper}}(\gamma) \propto \prod_{i=1}^n \gamma_i^{\alpha-1} \exp\left(-\frac{\gamma_i}{\beta}\right) = \exp\left(\sum_{i=1}^n \left(-\frac{\gamma_i}{\beta} + (\alpha-1) \log(\gamma_i)\right)\right). \quad (11)$$

Then, the full posterior over both \mathbf{X} and γ becomes:

$$p_{\text{post}}(\mathbf{X}, \gamma | \mathbf{M}) \propto \exp \left(-\frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_{\mathbb{F}}^2 - \sum_{i=1}^n \left(\frac{\|\mathbf{X}_{[i]}\|_{\mathbb{F}}^p}{\gamma_i} + \frac{\gamma_i}{\beta} - \left(\alpha - 1 - \frac{dt}{p} \right) \log(\gamma_i) \right) \right). \quad (12)$$

The question of how to best derive parameter estimates, in particular how to treat the two different types of parameters \mathbf{X} and γ , distinguishes different HBM-based inference strategies. Variational Bayesian approaches (MacKay 2003, Jordan et al 1999, Sato et al 2004, Friston et al 2008, Shervashidze and Bach 2015) and sparse Bayesian learning (Tipping 2001, Wipf and Rao 2004, Wipf and Nagarajan 2009, Zhang and Rao 2011) approaches rely on approximating or marginalizing the full, joint posterior distribution (12). In contrast, fully-Bayesian strategies (Calvetti et al 2009, Lucka et al 2012) work with it directly. The most popular one is the full maximum-a-posteriori (*full-MAP*) estimate which is defined as

$$(\hat{\mathbf{X}}_{\text{MAP}}, \hat{\gamma}_{\text{MAP}}) \in \underset{(\mathbf{X}, \gamma) \in \mathbb{R}^{q \times r} \times \mathbb{R}_+^n}{\arg \max} \{p_{\text{post}}(\mathbf{X}, \gamma | \mathbf{M})\}. \quad (13)$$

A common strategy to compute it is to minimize the *negative log posterior* $-\log p_{\text{post}}(\mathbf{X}, \gamma | \mathbf{M})$ by alternating minimization over \mathbf{X} and γ (known as *block coordinate descent* in optimization):

$$\mathbf{X}^{(k)} \in \underset{\mathbf{X} \in \mathbb{R}^{q \times r}}{\arg \min} \left\{ \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_{\mathbb{F}}^2 + \sum_{i=1}^n \frac{\|\mathbf{X}_{[i]}\|_{\mathbb{F}}^p}{\gamma_i^{(k-1)}} \right\}, \quad (14)$$

$$\gamma_i^{(k)} \in \underset{\gamma_i \in \mathbb{R}_+}{\arg \min} \left\{ \frac{\|\mathbf{X}_{[i]}^{(k)}\|_{\mathbb{F}}^p}{\gamma_i} + \frac{\gamma_i}{\beta} - \left(\alpha - 1 - \frac{dt}{p} \right) \log(\gamma_i) \right\}, \quad \forall i \in [n]. \quad (15)$$

Other fully-Bayesian estimates are defined as integrals of functions of \mathbf{X} and γ with respect to the posterior distribution, e.g. first or second moment estimates. To compute these high dimensional integrals efficiently, only Markov chain Monte Carlo (MCMC) methods that draw correlated samples from the posterior distribution can be used (Kaipio and Somersalo 2005, Robert and Casella 2005). A commonly used MCMC scheme for HBM is given by *blocked Gibbs sampling* which alternates as:

$$\mathbf{X}^{(k)} \sim p_{\text{post}}(\mathbf{X}, \gamma^{(k-1)} | \mathbf{M}) \propto p_{\text{post}}(\mathbf{X} | \mathbf{M}, \gamma^{(k-1)}), \quad (16)$$

$$\gamma^{(k)} \sim p_{\text{post}}(\mathbf{X}^{(k)}, \gamma | \mathbf{M}) \propto p_{\text{post}}(\gamma | \mathbf{M}, \mathbf{X}^{(k)}). \quad (17)$$

In this study, however, we are not interested in sampling the posterior distribution to compute the integral-based estimators but we rather want to explore the different modes of this multimodal distribution, each of which corresponds to parameters that are both sparse and likely to explain the data.

One can notice similar structures in (14)–(17). In each step, we make use of the conditional structure of the posterior: for γ fixed, we have to solve one qt -dimensional $\ell_{2,p}$ optimization/sampling problem, while for \mathbf{X} fixed, we have to solve n 1-dimensional optimization/sampling problems. We will describe these two steps in more detail in the next two sections.

2.3. HBM optimization

The optimization problem defined in equation (14) reduces to an $\ell_{2,p}$ -norm regularized regression problem that can be solved as described in section 2.1. To solve equation (15), we compute the first order optimality condition for each i :

$$-\frac{\|\mathbf{X}_{[i]}^{(k)}\|_F^p}{\gamma_i^2} + \frac{1}{\beta} - \frac{(\alpha - 1 - \frac{dt}{p})}{\gamma_i} = 0. \quad (18)$$

For $\alpha \geq dt/p + 1$, the problem in equation (15) is convex, and the positive root of equation (18) is given by:

$$\gamma_i = \beta \left(\nu + \sqrt{\nu^2 + \frac{\|\mathbf{X}_{[i]}^{(k)}\|_F^p}{\beta}} \right), \quad \nu := \frac{\alpha - 1 - dt/p}{2}. \quad (19)$$

Note that similar rules to update the noise level were considered in the Bayesian lasso (Park and Casella 2008, Kyung *et al* 2010) and the scaled lasso (see for instance (Städler *et al* 2010, Dalalyan 2012)). A difference though is that the update we perform here is on the penalty term, whereas in the mentioned references, it was rather performed on the data-fitting term.

If we furthermore choose $\alpha = dt/p + 1$, then $\nu = 0$ and most terms disappear; equations (14) and (15) hence read:

$$\mathbf{X}^{(k)} = \arg \min_{\mathbf{X} \in \mathbb{R}^{q \times t}} \left\{ \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_F^2 + \sum_{i=1}^n \frac{\|\mathbf{X}_{[i]}\|_F^p}{\gamma_i^{(k-1)}} \right\}, \quad (20)$$

$$\gamma_i^{(k)} = \sqrt{\beta} \sqrt{\|\mathbf{X}_{[i]}^{(k)}\|_F^p}, \quad \forall i = 1, \dots, n, \quad (21)$$

which can be combined to the fixed point iteration:

$$\mathbf{X}^{(k)} = \arg \min_{\mathbf{X} \in \mathbb{R}^{q \times t}} \left\{ \frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_F^2 + \frac{2}{\sqrt{\beta}} \sum_{i=1}^n \frac{\|\mathbf{X}_{[i]}\|_F^p}{2\sqrt{\|\mathbf{X}_{[i]}^{(k-1)}\|_F^p}} \right\}. \quad (22)$$

If we compare equations (22) with (3) we see that we re-derived the MM algorithm for $p = 1$ as an alternating optimization scheme to compute the *full-MAP* estimate for a specific HBM, namely using a conditional $\ell_{2,1}$ group prior and a Gamma hyper-prior with $\alpha = dt + 1$ and $\beta = 4/\lambda^2$. Using $\mathbf{w}_i^{(0)} := 1$ in the MM scheme corresponds to starting with $\gamma_i^{(0)} := 1/\lambda = \sqrt{\beta}/2$. From previous work (Strohmeier *et al* 2016) we know that due to the non-convexity, a good initialization of the weights $\mathbf{w}_i^{(0)}$ in the MM algorithm is crucial for its performance, but aside from uniform initialization, only heuristic initialization strategies were used, e.g. using the same re-weighting as in the sLORETA method (Pascual-Marqui 2002). In this work, we leverage the re-interpretation of the MM algorithm through the HBM framework to obtain multiple initializations in a systematic fashion, namely as samples drawn from the full posterior. This way, we can not only reach better local minima but more importantly, we can identify and characterize multiple possible sparse solutions. Such plausible solutions to the sparse regression problem in equation (1) are the modes of the posterior distribution (12) with different relative probability masses.

2.4. HBM sampling

As outlined in equations (16) and (17) in section 2.2, we sample the full posterior $p_{\text{post}}(\mathbf{X}, \gamma | \mathbf{M})$ by blocked Gibbs sampling, i.e. we alternate between sampling the conditional distributions $p_{\text{post}}(\mathbf{X} | \mathbf{M}, \gamma^{(k-1)})$ and $p_{\text{post}}(\gamma | \mathbf{M}, \mathbf{X}^{(k)})$. The conditional $p_{\text{post}}(\mathbf{X} | \mathbf{M}, \gamma^{(k-1)})$ is a high dimensional distribution composed of a Gaussian likelihood and an $\ell_{2,p}$ prior, where our main interest here is $p = 1$. It was demonstrated in Lucka (2012) that *single component Gibbs sampling (SC Gibbs)* is an efficient MCMC technique to sample such distributions. For the specific $\ell_{2,p}$ priors used here, *slice sampling* can be used to perform the sub-steps in SC Gibbs sampling, namely the sampling of the one-dimensional single-component conditional densities. The resulting *slice-within-Gibbs* sampler was examined in Lucka (2016). For completeness, the details of the implementation are given in appendix A.

The conditional $p_{\text{post}}(\gamma | \mathbf{M}, \mathbf{X}^{(k)})$ factorizes over groups i :

$$p_{\text{post}}(\gamma_i | \mathbf{M}, \mathbf{X}^{(k)}) \propto \exp \left(-\frac{\|\mathbf{X}_{[i]}^{(k)}\|_F^p}{\gamma_i} - \frac{\gamma_i}{\beta} + (\alpha - 1 - dt/p) \log(\gamma_i) \right). \quad (23)$$

For the case of $\alpha = dt/p + 1$, which is our main interest due to its connection to MM revealed in the previous section, equation (23) reduces to:

$$p_{\text{post}}(\gamma_i | \mathbf{M}, \mathbf{X}^{(k)}) \propto \exp \left(-\frac{\|\mathbf{X}_{[i]}^{(k)}\|_F^p}{\gamma_i} \right) \exp \left(-\frac{\gamma_i}{\beta} \right), \quad (24)$$

which can be sampled with a simple accept-reject algorithm as described in appendix B. The complete procedure is described in algorithm 2. Therein, K_0 refers to the burn-in size, i.e. the initial samples that are discarded, K to the sample size of the blocked Gibbs sampler. We denote K_{SC} and K_{SS} the sample sizes of the SC Gibbs and the slice sampler that carry out the sampling in the sub-steps.

Algorithm 2. BLOCK GIBBS SAMPLING SCHEME.

input : $\mathbf{M}, \mathbf{G}, \mathbf{X}^{(-K_0)}, \gamma^{(-K_0)}, K_0, K, K_{\text{Gibbs}}, K_{\text{SC}}, K_{\text{SS}}, \alpha, \beta$
for $k = -K_0 + 1$ to K **do**
 Set $\mathbf{X}^{(k)} = \mathbf{X}^{(k-1)}$.
 for $k_{\text{SC}} = 1$ to K_{SC} **do**
 Draw a random permutation P of $\{1, \dots, n\}$
 for $l \in P$ **do**
 Sample $\mathbf{X}_{(ij)}^{(k)} \sim p_{\text{post}}(\mathbf{X}_{(ij)} | \mathbf{X}_{-(ij)}^{(k-1)}, \mathbf{M}, \gamma^{(k)})$, $\forall (i,j) \in [l]$ - via K_{SS}
 steps of **Slice Sampling** algorithm 4.
 Sample $\gamma_i^{(k)} \sim p_{\text{post}}(\gamma_i | \mathbf{M}, \mathbf{X}^{(k)})$, $\forall i = 1, \dots, n$ via **Accept-Reject**
 Algorithm 5.
return $\{\mathbf{X}^{(k)}, \gamma^{(k)}\}_{k=1}^K$

2.5. Combining sampling and optimization

Finding the correct support in a sparse under-determined regression problem like (1) is inherently of combinatorial complexity. In the two approaches we examined, this is reflected in the non-convexity of the objective function (2) and the multi-modality of the joint posterior

distribution (12), respectively. Here, we want to investigate whether the link between MM and the HBM framework can be used to quantify the ambiguity and uncertainty posed by sparse support identification. Traditional uncertainty quantification (UQ) measures such as covariance estimates of \mathbf{X} or γ may fail to do so as they cannot capture the multi-modality of the posterior distribution in a satisfactory way. In addition, no sample $\mathbf{X}^{(k)}$ is exactly sparse: as the posterior distribution is a continuous density function, the event $\{\mathbf{X}_{[i]}^{(k)} = 0\}$ has zero probability. This means that the whole support of $\mathbf{X}^{(k)}$ is active with probability 1. Even a thresholded average of the support of $\mathbf{X}^{(k)}$ will only reveal the average probability of a location being part of the support. In our application to M/EEG source analysis, an arguably more interesting statistical output is the set sources forming the network of brain areas active in a given data set. This question needs a more profound spatial analysis of the structure of the most prominent modes of the posterior, and is left open by the above mentioned measures. Here, we propose to tackle it in a different way: algorithm 3 describes a combination of first sampling the posterior and then using each sample to initialize MM to optimize the posterior distribution. This yields a chain of different posterior modes $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$, i.e. approximate solutions to equation (1) that fulfill our *a priori* knowledge of a sparse support. If we assume that the division of $\mathbb{R}^{q \times t}$ into attractors of the MM algorithm roughly overlaps with the division of $\mathbb{R}^{q \times t}$ into modes of the marginalized posterior over \mathbf{X} within the HBM framework, the relative frequency with which these modes occur in $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$ corresponds to their relative posterior mass. While a mathematically more profound and detailed analysis of this heuristic is left for future work, we will illustrate in the following numerical examples how this mode analysis can be used to reveal and quantify some of the ambiguity of sparse under-determined regression problems.

Algorithm 3. COMBINATION OF GIBBS SAMPLER AND MM ALGORITHM.

input : $\mathbf{M}, \mathbf{G}, \lambda, \mathbf{X}^{(-K_0)}, \gamma^{(-K_0)}, K_0, K, K_{SC}, K_{SS}, K_{MM}, \epsilon > 0, \tau > 0$
 Use algorithm 2 with input
 $(\mathbf{M}, \mathbf{G}, \mathbf{X}^{(-K_0)}, \gamma^{(-K_0)}, K_0, K, K_{SC}, K_{SS}, \alpha = dt + 1, \beta = 4/\lambda^2)$ to obtain
 MCMC chain $\{\mathbf{X}^{(k)}, \gamma^{(k)}\}_{k=1}^K$.
for $k = 1$ *to* K **do**
 Set $\mathbf{w}_i^{(0)} = \lambda \gamma_i^{(k)}, \forall i = 1, \dots, n$ and run algorithm 1 with input
 $(\mathbf{M}, \mathbf{G}, \lambda, \mathbf{W}_i^{(0)}, \epsilon, \tau, K_{MM})$ to obtain $\hat{\mathbf{X}}^{(k)}$.
return $\{\hat{\mathbf{X}}^{(k)}, \mathbf{X}^{(k)}, \gamma^{(k)}\}_{k=1}^K$

3. Results

We now examine the benefits of our re-interpretation of the MM algorithm described in section 2.1 as a specific way to compute a full-MAP estimate for a specific HBM as described in sections 2.2 and 2.3. We first illustrate basic properties of the methods in a one dimensional toy problem before we examine a simulated MEG dataset and two experimental M/EEG datasets.

3.1. One dimensional illustrations

We start with a toy problem where $m = 10, q = 20, d = t = 1$ and the true unknown \mathbf{X} is all zero except for $\mathbf{X}_5 = \mathbf{X}_{15} = 1$.




Figure 1. Setup for the first one dimensional illustration. (a) Random Gaussian design covariance matrix \mathbf{C} ; (b) design matrix \mathbf{G} ; (c) $\mathbf{G}^T \mathbf{G}$.

Example 1. First, \mathbf{G} is a random matrix constructed in the following way: its rows are drawn from a Gaussian distribution with zero mean and a block-diagonal covariance matrix $\mathbf{C} = \text{blkdiag}(\mathbf{C}_1, \mathbf{C}_2)$, where $(\mathbf{C}_1)_{ij} = 0.5^{|i-j|}$, $(\mathbf{C}_2)_{ij} = 0.95^{|i-j|}$, $i, j = 1, \dots, 10$. Then, each column is normalized to have unit ℓ_2 norm. Figure 1 illustrates the set-up. Notice that due to the asymmetry in the design, the correct recovery of the source at location 15 is more difficult due to the stronger correlation between columns 11-20 of \mathbf{G} . We generate \mathbf{M} by adding AWGN scaled by $20\% \|\mathbf{M}\|_\infty$ to $\mathbf{G}\mathbf{X}$. We first run the MM algorithm 1 using a uniform initialization, i.e. $\mathbf{w}_i^{(0)} = 1, \forall i = 1 \in [n]$, with $\lambda = 0.2\lambda_{\max}$ where $\lambda_{\max} = \max_{1 \leq i \leq n} \|(\mathbf{G}^T \mathbf{M})_{[i]}\|_F^2$ is the smallest regularization value for which no source is found as active using an $\ell_{2,1}$ regularization (Ndiaye *et al* 2015, Strohmeier *et al* 2016). It recovers an \mathbf{X} supported at locations 5 and 11, i.e. it is not able to locate the second source correctly. Then, we run algorithm 3 with $K = K_0 = 10\,000$, $K_{\text{SC}} = K_{\text{SS}} = 10$ and the same settings for the MM algorithm as before to obtain chains of posterior samples $\{\mathbf{X}^{(k)}\}_{k=1}^K$, and the corresponding posterior modes $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$. We clustered the modes based on their spatial support which reveals that a total of 16 different modes were found. Figure 2 depicts the spatial support of the modes listed based on the relative frequency with which they were found. It reveals that, indeed, there is a larger uncertainty in the location of the second source (at true location 15) and that in this scenario, the support of the mode which is found most often coincides with that of the true solution. To check that the MCMC sampler described in algorithm 2 is not simply stuck in this mode for a long time, we compute how many steps (with respect to index k) it takes on average before solutions $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$ change. The result is 1.63 steps, which means that the sampler switches between modes very frequently and should be able to explore the posterior sufficiently well. In traditional UQ, the covariance matrix of the posterior samples $\{\mathbf{X}^{(k)}\}_{k=1}^K$ would be used to characterize uncertainty and correlation between activity at different locations. Here, we want to compare it to a matrix whose (i, j) th entry shows the relative frequency with which two locations i and j are simultaneously active in the support of the modes $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$. Such a matrix is another way to visualize the information given by figure 2. Figure 3 shows that the covariance matrix reveals very little information about the true, sparse source locations and the larger ambiguity about the source at location 15 induced by the asymmetric design of \mathbf{G} .

Example 2. While the posterior mode whose support coincided with that of the true solution was also found with the highest relative frequency, it is not clear whether this frequency is a reliable indication of the mode's true relative posterior mass. In general, this question is difficult to examine for high dimensional problems. Nonetheless, we constructed a second ex-




Figure 2. Each row of the matrix indicates the spatial support of a mode found by algorithm 3 in example 1 (the true locations are 5 and 15). The rows are ordered by the relative frequency (in %), by which they are found. The row marked by * corresponds to the mode found by MM with uniform initialization.




Figure 3. Comparison between (a) posterior sample covariance matrix; (b) corresponding correlation matrix; (c) a matrix which shows the frequency with which locations i and j are simultaneously found active in the $\hat{\mathbf{X}}^{(k)}$ in its (i,j) th entry.

ample to at least show that the frequencies are consistent: we now draw the rows of a 10×10 matrix $\tilde{\mathbf{G}}$ from a Gaussian distribution with zero mean and the covariance matrix \mathbf{C}_2 as for the previous example. Then, we set $\mathbf{G} = [\tilde{\mathbf{G}}, \tilde{\mathbf{G}}]$, i.e. the first and last 10 columns of \mathbf{G} are exactly the same. This means that the regression problem (1) and the posterior distribution are invariant with respect to switching the first and last 10 entries. Every mode has a corresponding copy ‘on the other side’, which should be found with the same relative frequency. All other settings are the same as in the previous example, except that we choose $\lambda = 0.5\lambda_{\max}$ larger than before to boost modes which are only supported at a single location. Figure 4 reveals that, indeed, all modes found are supported only at a single location and are found with a similar frequency as their corresponding copy. The average number of steps for the sampler to switch between different modes is now 1.40 steps which is, again, very low. In addition, even the average number of steps it needs to switch between modes supported in locations 1–10 to modes supported in 11–20 is only 2.34 steps.




Figure 4. Each row of the matrix indicates the spatial support of a mode found by algorithm 3 in example 2 (the true locations are 5 and 15). The rows are ordered by the relative frequency (in %), by which they are found and colored in such a way that it is easy to identify modes which are similar when locations 1–10 and 11–20 are switched.

3.2. Simulated MEG data

We generated a realistic simulation based on a free-orientation ($d = 3$) source model with $n = 7498$ cortical locations and $m = 306$ MEG sensors. Two of these locations were selected to be active, one in each hemisphere. One of the sources had a deep ventral location in the inferior occipital gyrus (figure 5(c)), and the second one had a more superficial location in the motor cortex (figure 5(a)). Their corresponding waveforms are shown in figure 5(b). When passed to the solvers, they are cropped between 40 to 180 ms to keep only the two peaks. This leads to $t = 43$ time samples.

First, we re-visit the question whether we are able to find better source estimates using MCMC-derived initializations than with the uniformly initialized MM algorithm 1, also in this high-dimensional example. For this, we first run the MM algorithm 1 using a uniform initialization, i.e. $\mathbf{w}_i^{(0)} = 1, \forall i = 1 \in [n]$, with $\lambda = 0.05\lambda_{\max}$. Then, we run algorithm 3 with $K_0 = 300, K = 900, K_{SC} = K_{SS} = 1$ and the same settings for the MM algorithm as before to obtain a chain of MM-optimized solutions $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^K$. Figure 6(c) shows the histogram of the objective function values reached by these solutions (computed with equation (5)). The vertical black bar shows the value of the objective function of the uniformly initialized MM solver and we can see that some initializations indeed lead to source estimates with a lower objective value. Figures 6(a) and (b) show the locations of the estimated sources resulting from uniform and best MCMC-based initialization. For the artificial source in figure 6(a), both results find the exact location, so they are superposed. For the deeper source in figure 6(b), neither result finds the exact position, but the MCMC-based initialization is closer. This means that the result did not only improve from an optimization point of view, but also judged by the quality criteria of the given application.

Now, we examine how the posterior modes found by algorithm 3 react to changes in the measurement design. To do so, we switch from using all 306 MEG sensors to using only 204 gradiometers or each other gradiometer (102 sensors). By reducing the number of sensors we increase the under-determinedness of the problem, and the intuition is that it should lead to




Figure 5. Simulated MEG dataset. (a) and (c) show superficial and deep source (hidden in the medial view) locations, respectively. (b) gives their corresponding waveforms color-coded by location.




Figure 6. Location of simulated and estimated sources using the uniformly initialized MM solver (denoted as ‘MM’) and best MCMC-based initialization in terms of objective function value. Left: estimation of the artificial source on the left hemisphere. Middle: estimation of the deep source on the right hemisphere. Right: histogram of the objective function value for 900 MCMC initializations (algorithm 3). The uniform initialization used for the MM (black vertical line) is not very bad, meaning that the basic MM is able to recover a good source estimates for some configurations. See figure 9 for a case where the basic MM fails.

more variability among the plausible sparse solutions. The graphical analysis, which is more involved for this high-dimensional scenario, is presented and described in figures 7 and 8. A first observation is that the superficial source in the premotor cortex was correctly identified as part of the support of every local minima when using the full 306 MEG sensors. It was however sometimes mis-localized when reducing the number of sensors (figure 7). A second observation is that the spatial spread of these miss-localizations is smaller for this superficial source than it is for the deep source. This deep source in the ventral cortex is more difficult to find even with all sensors. Indeed, none of the 100 best initialization perfectly localized the deep simulated source. In general, we can clearly see how the ambiguity increases when decreasing the number of sensors, and how the distribution of source networks gets more fuzzy. However, our analysis also provides useful local measures of these phenomena.

3.3. Experimental MEG data

We now repeat our analysis with two experimental open datasets. The first one is a recording of auditory evoked fields (MNE sample dataset (Gramfort *et al* 2013)). The second one contains visual evoked fields (visual condition of MNE sample dataset) for which source

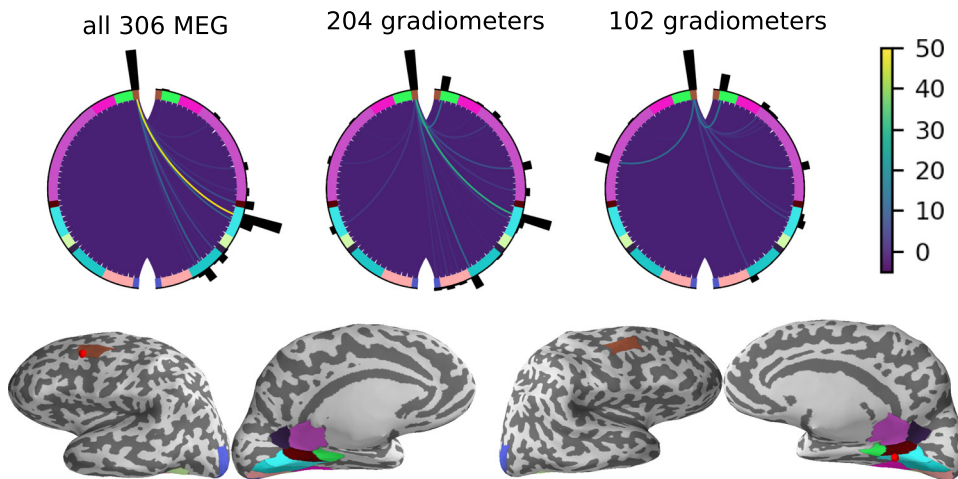


Figure 7. Source network analysis for simulated data: for a clearer presentation, the chain $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^{900}$ of modes was thinned to the 100 ones with the lowest objective function (equation (2)). The first row of sub-figures displays the support of these best modes in the following way: each position in the circle represents a source location that was part of the support of at least one mode for one sensor configuration. The black bar attached to each position corresponds to the relative frequency with which this source location appeared as part of the support. Two positions are connected by a line if they were simultaneously part of the support and the color of this line corresponds to the relative frequency with which this happened. Note that the background of the circle is white, but it is densely covered by purple lines indicating rare connections. The positions are placed left or right, depending on which hemisphere they belong to. For symmetry, for each active source location, its counterpart on the other hemisphere was included in the graphic as well. In addition, the positions are grouped and colored based on a parcellation of the brain into anatomical regions (taken from an atlas). The second row of sub figures shows these regions in the brain and the simulated sources.

localization is a more difficult task due to the proximity between neural sources. The true nature of the underlying source network is also less clear for this second dataset.

Figure 9 shows the equivalent to figure 6 for both datasets. Again, we see that lower objective function value can be obtained using MCMC-based initializations. The auditory sample dataset is commonly assumed to be generated by two bilateral focal sources around the auditory cortices in the superior temporal gyrus of the temporal lobe. Due to the superficial nature of these sources and their large distance, estimation of their position is regarded as a relatively simple task. Indeed, the histogram shows that using MCMC-based initializations does not help a lot to reduce the objective function compared to a uniformly initialized MM solution. However, in the case of the visual dataset, where several closed-by sources are active, the difference is quite drastic. The majority of the MCMC-based initializations lead to lower values of the objective function. Looking at the source distribution plots on the brain for both datasets, one can also observe more complex source configurations for the visual data.

Next, we repeat the graphical source network analysis from figure 7 for the two datasets. Figure 10 shows the results for the auditory dataset and three sensor configurations: all 364 EEG + MEG sensors, all 306 MEG sensors or each other sensor resulting in 182 EEG + MEG sensors. One can see how adding EEG to MEG sensors reduces the ambiguity of the regression problem. The plots show less but more prominent modes, i.e. the posterior mass is concentrated on fewer stable source configurations. We also see that the locations




Figure 8. The support of $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^{900}$ was extracted to build an uncertainty map by smoothing. The relative frequencies with which each source location was part of the support was computed and plotted on the brain surface together with the two simulated sources (green dots). Each column corresponds to the results for each of the three sensor setups examined. Less the number of sensors and/or more the source is deep, more uncertainty in the brain map. Note that the deep source is not in the support of the $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^{900}$, it seems to be recovered only due to smoothing.

of the most prominent modes shift. This is consistent with results of other studies on EEG-MEG combination (Molins *et al* 2008, Lucka 2014, Aydin *et al* 2014) as EEG is sensitive to some sources that MEG is almost blind to, e.g. sources with a strong radial component. If we subsample the EEG + MEG sensors by only using every other location, the ambiguity and spatial spread of the recovered support increases. One can see that there is more activity in the dark green label, which corresponds to a brain area commonly not associated with auditory responses. The connections between source locations show that none of the modes found really stands out, i.e. is found much more often compared to the others. Most of the connections do not occur more than 200 times within the 900 samples, so they are part of the purple background of low frequency connections in the plots.

Figure 11 shows the same results for the more complex visual dataset. Compared to the auditory dataset, we see that even with all sensors, the ambiguity of the regression problem seems to be a lot higher compared to the auditory dataset: we see that the posterior mass is distributed among many more source configurations. For the other two sensor configurations, we see similar effects as in the auditory data set. Nevertheless, it can be noticed that the large majority of identified sources with all MCMC initializations are on the right hemisphere. This is consistent with the known functional organization of the visual cortex. Indeed, in this

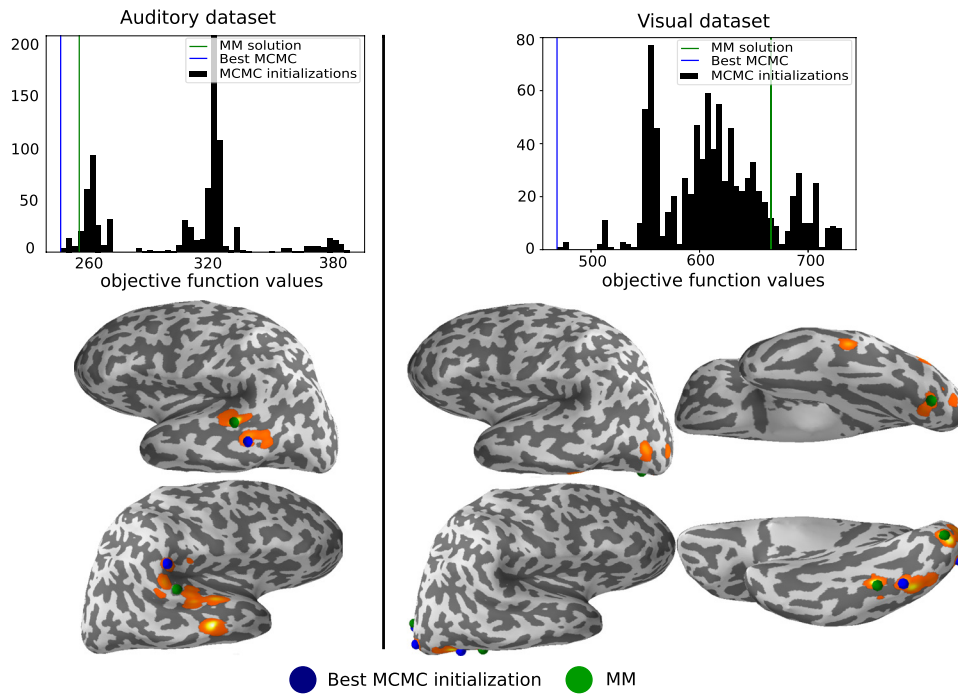


Figure 9. Histogram of the objective function value of $\{\hat{\mathbf{X}}^{(k)}\}_{k=1}^{900}$ for auditory and visual datasets (306 MEG sensors). The histogram for visual dataset shows more uniform one in the MM solution. Under each histogram, these source configurations are shown on the left and right hemisphere.




Figure 10. Source network analysis for auditory data. The figures are constructed in the same way as described in figure 7 except that all 900 mode samples are displayed.




Figure 11. Source network analysis for visual data. The figures are constructed in the same way as described in figure 7 except that all 900 mode samples are displayed.

experimental condition the subject was presented with checker board flashes on the left visual hemifield which is known to primarily project onto the right hemisphere of the cortex.

4. Discussion

Scientific literatures relying either on frequentist or on Bayesian statistical inference often coexist in many fields ranging from machine learning, inverse problems, signal processing to computational biology. In this work, we started from an under-determined, ill-conditioned MMV/multi-task regression problem and examined two seemingly unrelated approaches—MM as an optimization technique for tackling non-convex optimization problems arising in frequentist regression, and HBM as a Bayesian prior modeling framework. We showed that one obtains the same algorithms, and therefore the same solutions, when considering some specific choices of models, parameters and inference strategies. In particular the parallel was established between the $\ell_{2,1/2}$ -norm regularized regression by MM and the full-MAP estimation for $\ell_{2,1}$ hierarchical priors with specific Gamma hyper-priors. We further showed that this conceptual parallel can be exploited to improve the MM solution by providing well-informed algorithmic initializations.

For this, we first constructed a multi-layered Gibbs sampler for the joint posterior density of our HBM. Each sample is then used to initialize the MM step done with a state-of-the-art convex solver using block coordinate descent techniques and acceleration strategies based on active sets. The sampler used has also an efficient sub-sampler for $\ell_{2,1}$ priors at its core. Despite the multi-modality of the posterior, the MCMC scheme is able to jump rapidly between the different attractors of the MM scheme. Indeed, using each sample as an initialization to the MM computation, one ends up in many different local minima/modes of the posterior (see figures 9 and 11). Therefore, this procedure allows us to reveal and explore different plausible source configurations in more details.

Based on this observation, we showcased how one can use the chain of local minima found by MCMC-initialized MM to analyze the variability of the different sparse solutions and how this yields different information compared to traditional and generic Bayesian uncertainty quantification techniques that use for example covariance estimates or credible sets

derived from posterior samples (Szabó *et al* 2015). It is also different from methods developed specifically for parametric M/EEG source localization based on dipole fitting (Fuchs *et al* 2004, Darvas *et al* 2005). These latter approaches cannot easily be transferred to sparse, non-parametric approaches. On the other hand, our approach can easily be extended to include space-time-frequency structured sparsity constraints that can model more sophisticated source configurations than examined here (Castao-Candamil *et al* 2015). Using our developed techniques on simulations and actual data, one could observe that uncertainty in M/EEG is location specific and also source configuration specific. This is of course well-known by experts in this field, but here we provide a computational approach to visualize it and quantify it. This is an important incentive to develop such automated, data-dependent methods to quantify uncertainties in the context of M/EEG source imaging. In more conventional imaging methods such as computer tomography (CT) or magnetic resonance imaging (MRI), the signal originates from weak tissue interaction with strong external fields and the forward operator \mathbf{G} depends almost exclusively on the physical properties of the scanner. In this situation, uncertainty is usually distributed in a smooth, well-known way over the image domain. Artifacts as well as real anatomical features are also easy to distinguish for a trained radiologist. The situation for M/EEG is very different. The weak signals originate from endogenous activity, and they are very dependent on dataset specific factors such as source orientation, location and attenuation which all depend on the geometry of the head of the analyzed subject. That is also why the forward matrix \mathbf{G} needs to be constructed for each individual patient, after fixing the electrical properties of the head issues, which if wrong, increases the uncertainties.

When considering real data, the source to recover is often poorly understood, especially when it comes to pathological brain activity such as ictal or inter-ictal epileptic activity. In such a situation, providing a single source configuration as a result, together with an ad-hoc uncertainty quantification based on previous studies or acquired expertise, might not be an optimal use of the M/EEG data. Instead, providing multiple hypotheses together, along with a quantification of their uncertainty, can be more useful. Indeed for applications such as pre-surgical epilepsy diagnosis, where M/EEG recordings are one of several diagnostic modalities, each candidate source configuration can provide some evidence for or against a diagnostic hypothesis that could lead to a surgery decision. We therefore believe that future extensions of this work towards a consistent framework for interpreting and quantifying the multitude of potential results of sparse M/EEG source reconstruction approaches can have a significant impact on clinical settings.

Acknowledgments

The authors would like to thank Mathurin Massias for its valuable suggestions and proofreading of the manuscript. This work was supported by the French National Research Agency (ANR-14-NEUC-0002-01), the European Research Council Starting Grant SLAB ERC-YStG-676943 and in parts by the Engineering and Physical Sciences Research Council, UK (EP/K009745/1), the European Union's Horizon 2020 research and innovation programme H2020 ICT 2016-2017 under grant agreement No 732411 (as an initiative of the Photonics Public Private Partnership) and the Netherlands Organisation for Scientific Research (NWO 613.009.106/2383).

Appendix A. Slice-within-Gibbs sampler for parameter update

Within the algorithm 2, to update a group $\mathbf{X}_{[l]}$, we need to sample from the all the one-dimensional, SC densities

$$p_{\text{post}}(\mathbf{X}_{(i,j)} | \mathbf{X}_{-(i,j)}^{(k-1)}, \mathbf{M}, \gamma^{(k)}), \quad (i,j) \in [l], \quad (\text{A.1})$$

where $\mathbf{X}_{-(i,j)}$ refers all the coefficient of the matrix \mathbf{X} except the term (i,j) .

To implement this efficiently, we can precompute several terms and make use of the specific spatio-temporal group structure of the posterior. We first derive the part of the likelihood (7) that depends on a given index pair $(i,j) \in [l]$:

$$\frac{1}{2} \|\mathbf{M} - \mathbf{G}\mathbf{X}\|_{\text{F}}^2 = \sum_{j'} \frac{1}{2} \|\mathbf{M}_{(:,j')} - \mathbf{G}\mathbf{X}_{(:,j')}\|_2^2 \stackrel{j}{\propto} \quad (\text{A.2})$$

$$\frac{1}{2} \|\mathbf{M}_{(:,j)} - \mathbf{G}\mathbf{X}_{(:,j)}\|_2^2 = \frac{1}{2} \|\mathbf{M}_{(:,j)} - (\mathbf{G}_{(:,,-i)}\mathbf{X}_{(-i,j)} + \mathbf{G}_{(:,i)}\mathbf{X}_{(i,j)})\|_2^2 \quad (\text{A.3})$$

$$\stackrel{i}{\propto} \frac{1}{2} \|\mathbf{G}_{(:,i)}\|_2^2 \mathbf{X}_{(i,j)}^2 + \mathbf{G}_{(:,i)}^T (\mathbf{M}_{(:,j)} - \mathbf{G}_{(:,,-i)}\mathbf{X}_{(-i,j)}) \mathbf{X}_{(i,j)} \quad (\text{A.4})$$

$$= \frac{1}{2} \|\mathbf{G}_{(:,i)}\|_2^2 \mathbf{X}_{(i,j)}^2 + \left((\mathbf{G}^T \mathbf{M})_{(i,j)} - (\mathbf{G}_{(:,i)}^T \mathbf{G}) \mathbf{X}_{(:,j)} - \|\mathbf{G}_{(:,i)}\|_2^2 \right) \mathbf{X}_{(i,j)} \quad (\text{A.5})$$

$$:= a z^2 + b z, \quad \text{with } z := \mathbf{X}_{(i,j)}, \quad a := \frac{1}{2} \|\mathbf{G}_{(:,i)}\|_2^2, \quad (\text{A.6})$$

$$b := (\mathbf{G}^T \mathbf{M})_{(i,j)} - (\mathbf{G}_{(:,i)}^T \mathbf{G}) \mathbf{X}_{(:,j)} - \|\mathbf{G}_{(:,i)}\|_2^2. \quad (\text{A.7})$$

Note that $\|\mathbf{G}_{(:,i)}\|_2^2$ and $(\mathbf{G}^T \mathbf{M})$ can be precomputed. The challenging part in the computation of b is to compute $(\mathbf{G}_{(:,i)}^T \mathbf{G})$, as one typically does not want to pre-compute the $q \times q$ matrix $\mathbf{G}^T \mathbf{G}$ and hold it in memory. However, to update all the td components of the l th group (e.g. in the visual evoked fields example, $t = 211$, $d = 3$) one only needs the $d \times q$ matrix $(\mathbf{G}_{(:,l]}^T \mathbf{G})$. Thus, we compute $(\mathbf{G}_{(:,l]}^T \mathbf{G})$ at the start of updating group $\mathbf{X}_{[l]}$ and hold it memory throughout the bloc update. Besides this, the most costly operation to compute b is a dot product of vectors of size q . Next, we derive the part of the prior (10) that depends on $\mathbf{X}_{(i,j)}$, $(i,j) \in [l]$:

$$\begin{aligned} \sum_{l=1}^n \left(\frac{\|\mathbf{X}_{[l]}\|_{\text{F}}^p}{\gamma_l} + \frac{dt}{p} \log(\gamma_l) \right) \stackrel{\mathbf{X}_{(i,j)}, (i,j) \in [l]}{\propto} \gamma_l^{-1} \|\mathbf{X}_{[l]}\|_{\text{F}}^p &= \gamma_l^{-1} \left(\sum_{(i',j') \in [l]} \mathbf{X}_{(i',j')}^2 \right)^{p/2} \\ &= \gamma_l^{-1} \left(\mathbf{X}_{(i,j)}^2 + \sum_{\substack{(i',j') \in [l] \\ (i',j') \neq (i,j)}} \mathbf{X}_{(i',j')}^2 \right)^{p/2} =: c (z^2 + d)^{p/2}, \end{aligned} \quad (\text{A.8})$$

with c and d defined and computed in an obvious way. Taken together, to update $\mathbf{X}_{(i,j)}$, we have to sample from the one-dimensional density:

$$p(z) \propto \exp(-a z^2 - b z) \exp(-c (z^2 + d)^{p/2}) =: p_1(z) p_2(z). \quad (\text{A.9})$$

We take advantage of the fact that (A.9) factorizes in a Gaussian likelihood part $p_1(z)$ and a symmetric, log-concave prior part $p_2(z)$, and use a generalized form of slice sampling (Neal 2003, Robert and Casella 2005) as described in more detail in Lucka (2016) and summarized in algorithm 4. Determining S_2^y in our case is trivial:

$$p_2(z) \geq y \Leftrightarrow c(z^2 + d)^{p/2} \leq -\log(y) \Leftrightarrow |x| \leq \left(\left(\frac{-\log(y)}{c} \right)^{2/p} - d \right)^{1/2}. \quad (\text{A.10})$$

Then, we use a slightly modified, more robust, version of the fast table-based algorithm described in Chopin (2011) to sample from the truncated Gaussian distribution $p_1(z)\mathbb{1}_{S_2^y}(z)$. As initialization for z , we always chose the current value of $\mathbf{X}_{(i,j)}$.

Algorithm 4. SLICE SAMPLING.

input : $p(z) \propto p_1(z)p_2(z)$, z , K_{SS}
for $k = 1$ **to** K_{SS} **do**
 Draw y uniform from $[0, p_2(z)]$ (vertical move).
 Determine $S_2^y := \{z \mid p_2(z) \geq y\}$
 Draw z from $p_1(z)\mathbb{1}_{S_2^y}(z)$ (weighted horizontal move).
return z as a sample of $p(z)$

Appendix B. Accept-reject sampler for hyperparameter update

The conditional density (24) is of the type

$$p(x) \propto \exp\left(-\frac{c}{x}\right) \exp\left(-\frac{x}{\beta}\right), \quad c, \beta \geq 0. \quad (\text{B.1})$$

Note that the first factor is monotonically increasing with limit 0 for $x \searrow 0$ and limit 1 for $x \rightarrow \infty$ while the second factor is proportional to a simple exponential distribution (see figure B1). We can therefore easily construct a dominating density $g(x) \geq p(x)$ to carry out accept-reject sampling (Robert and Casella (2005), section 2.3.2) to generate a sample $z \sim p$: we generate $y \sim g$, $u \sim \mathcal{U}_{[0,1]}$ and accept $z = y$ if $u \leq p(y)/g(y)$ and repeat otherwise. Choosing $g(x) = \exp(-x/\beta)$ would yield a valid sampling density but this choice becomes inefficient with increasing c . Therefore, we split the sampling density into two parts:

$$g(x) = \begin{cases} \hat{p} & \text{if } x < \tilde{x} \\ \exp\left(-\frac{x}{\beta}\right) & \text{otherwise} \end{cases}, \quad (\text{B.2})$$

where $\hat{p} = \max_x p(x)$ is the maximal probability attained at $\hat{x} = \arg \max_x p(x) = \sqrt{\beta c}$ and $\tilde{x} = \beta c / \hat{x} + \hat{x}$ is the solution to $\exp(-x/\beta) = \hat{p}$ (see figure B1). Sampling from (B.2) is then straight-forward using $v, w \sim \mathcal{U}_{[0,1]}$: if one computes

$$G_{x \geq \tilde{x}} = \int_{\tilde{x}}^{\infty} g(x) dx = \beta \exp(-\tilde{x}/\beta), \quad G_{x < \tilde{x}} = \int_0^{\tilde{x}} g(x) dx = \hat{p}\tilde{x}, \quad (\text{B.3})$$

then $v < G_{x \geq \tilde{x}} / (G_{x \geq \tilde{x}} + G_{x < \tilde{x}})$ determines that we are in the tail, $x > \tilde{x}$, where we can use a simple inverse cumulative distribution method to draw a proposal from $g(x)$ using w . If v determines that we are in $x \leq \tilde{x}$, then $x = w\tilde{x}$ is the proposal. For numerical precision, we only compute logarithms of probabilities and use that for $a > 0, b \geq 0$:

$$\log(a + b) = \log a + \log(1 + \exp(b - a)). \quad (\text{B.4})$$




Figure B1. Sketch of the quantities used in the accept-reject sampling algorithm 5.

The whole sampling scheme is shown in algorithm 5. We found the scheme to be efficient enough for all of our computations, i.e. the chosen $g(x)$ is close enough to $p(x)$ to result in an accepted sample after a few trails. If this would become a problem, it would be easy to adaptively improve the dominating density.

Algorithm 5. ACCEPT-REJECT ALGORITHM FOR HYPERPARAMETER UPDATE.

input : $c \geq 0, \beta > 0$
Set $\hat{x} = \sqrt{\beta c}$.
Set $\log \hat{p} = -c/\hat{x} - \hat{x}/\beta$.
Set $\tilde{x} = \beta c/\hat{x} + \hat{x}$.
Set $\log G_{x \geq \tilde{x}} = \log \beta - \tilde{x}/\beta$.
Set $\log G_{x < \tilde{x}} = \log \hat{p} + \log \tilde{x}$.
Set $\log G_{\text{tot}} = \log G_{x \geq \tilde{x}} + \log (1 + \exp(G_{x < \tilde{x}} - G_{x \geq \tilde{x}}))$.
while true do
 Draw u, v, w uniform from $(0, 1)$.
 if $\log v + \log G_{\text{tot}} < \log G_{x \geq \tilde{x}}$ **then**
 Set $W = \log w - \tilde{x}/\beta$.
 Propose $x = -\beta W$:
 if $\beta \log(u) < c/W$ **then**
 return x (acceptance)
 else
 Propose $x = w\tilde{x}$:
 if $\log u + \log \hat{p} < -c/x - x/\beta$ **then**
 return x (acceptance)

ORCID iDs

Yousra Bekhti  <https://orcid.org/0000-0002-6804-8516>

Felix Lucka  <https://orcid.org/0000-0002-8763-5177>

Alexandre Gramfort  <https://orcid.org/0000-0001-9791-4404>

References

- Aydin M et al 2014 *PLoS One* **9** e93154
- Calvetti D, Hakula H, Pursiainen S and Somersalo E 2009 *SIAM J. Imaging Sci.* **2** 879–909
- Candès E J, Wakin M B and Boyd S P 2008 *J. Fourier Anal. Appl.* **14** 877–905
- Candès E and Wakin M 2008 *IEEE Signal Process. Mag.* **25** 21–30
- Castao-Candamil S, Hhne J, Martnez-Vargas J D, An X W, Castellanos-Domnguez G and Haufe S 2015 *NeuroImage* **118** 598–612
- Chartrand R 2007 *IEEE Signal Process. Lett.* **14** 707–10
- Chen S S, Donoho D L and Saunders M A 1998 *SIAM J. Sci. Comput.* **20** 33–61
- Chopin N 2011 *Stat. Comput.* **21** 275–88
- Chzhen E, Hebiri M and Salmon J 2017 arXiv:1707.05232
- Combettes P L and Pesquet J C 2011 Proximal splitting methods in signal processing *Fixed-Point Algorithms for Inverse Problems in Science and Engineering* (New York: Springer) pp 185–212
- Cotter S F, Rao B D, Engan K and Kreutz-Delgado K 2005 *IEEE Trans. Signal Process.* **53** 2477–88
- Dalalyan A S 2012 SOCP based variance free Dantzig selector with application to robust estimation *C. R. Math.* **350** 785–8
- Darvas F, Rautiainen M, Pantazis D, Baillet S, Benali H, Mosher J, Garnero L and Leahy R 2005 *NeuroImage* **25** 355–68
- Dempster A P, Laird N M and Rubin D B 1977 Maximum likelihood from incomplete data via the EM algorithm *J. R. Stat. Soc. B* **39** 1–38
- Engemann D A and Gramfort A 2015 *NeuroImage* **108** 328–42
- Foucart S and Rauhut H 2013 *A Mathematical Introduction to Compressive Sensing* (Basel: Birkhäuser)
- Friston K J, Harrison L, Daunizeau J, Kiebel S J, Phillips C, Trujillo-Barreto N J, Henson R N, Flaxin G and Mattout J 2008 *Neuroimage* **39** 1104–20
- Fuchs M, Wagner M and Kastner J 2004 *Clin. Neurophysiol.* **115** 1442–51
- Gasso G, Rakotomamonjy A and Canu S 2009 *IEEE Trans. Signal Process.* **57** 4686–98
- Gramfort A et al 2013 *Frontiers Neurosci.* **7** 267
- Horst R and Thoai N V 1999 Dc programming: overview *J. Optim. Theory Appl.* **103** 1–43
- Jordan M I, Ghahramani Z, Jaakkola T S and Saul L K 1999 *Mach. Learn.* **37** 183–233
- Kaipio J P and Somersalo E 2006 *Statistical and Computational Inverse Problems (Applied Mathematical Sciences vol 160)* (New York: Springer)
- Kyung M, Gill J, Ghosh M and Casella G 2010 Penalized regression, standard errors, and bayesian lassos *Bayesian Anal.* **5** 369–411
- Lange K, Hunter D R and Yang I 2000 *J. Comput. Graph. Stat.* **9** 1–20
- Lucka F 2012 *Inverse Problems* **28** 125012
- Lucka F 2014 Bayesian inversion in biomedical imaging *PhD Thesis* University of Muenster
- Lucka F 2016 *Inverse Problems* **32** 115019
- Lucka F, Pursiainen S, Burger M and Wolters C 2012 *NeuroImage* **61** 1364–82
- MacKay D J 2003 *Information Theory, Inference and Learning Algorithms* (Cambridge: Cambridge University Press)
- Molins A, Stufflebeam S M, Brown E N and Hämäläinen M 2008 *Neuroimage* **42** 1069–77
- Ndiaye E, Fercoq O, Gramfort A and Salmon J 2015 *NIPS* pp 811–9
- Neal R M 2003 *Ann. Stat.* **31** 705–67
- Osher S, Burger M, Goldfarb D, Xu J and Yin W 2006 *Multiscale Modeling Simul.* **4** 460–89
- Park T and Casella G 2008 *J. Am. Stat. Assoc.* **103** 681–6
- Pascual-Marqui R D 2002 *Methods Findings Exp. Clin. Pharmacol.* **24** 5–12
- Rakotomamonjy A 2011 *Signal Process.* **91** 1505–26

- Robert C P and Casella G 2005 *Monte Carlo Statistical Methods (Springer Texts in Statistics)* (New York: Springer)
- Saab R, Chartrand R and Yilmaz O 2008 *Proc. of the IEEE Int. Conf. on Acoustics, Speech, and Signal Processing* (IEEE) pp 3885–8
- Sato M A, Yoshioka T, Kajihara S, Toyama K, Goda N, Doya K and Kawato M 2004 *NeuroImage* **23** 806–26
- Schifano E D *et al* 2010 *Electron. J. Stat.* **4** 1258–99
- Shervashidze N and Bach F 2015 *IEEE Trans. Signal Process.* **63** 4894–902
- Sorrentino A, Parkkonen L, Pascarella A, Campi C and Piana M 2009 *Hum. Brain Mapp.* **30** 1911–21
- Städler N, Bühlmann P and van de Geer S 2010 *TEST* **19** 209–56
- Strohmeier D, Bekhti Y, Haueisen J and Gramfort A 2016 *IEEE Trans. Med. Imaging* **35** 2218–28
- Szabó B *et al* 2015 *Ann. Stat.* **43** 1391–428
- Tibshirani R 1996 *J. R. Stat. Soc. B* **58** 267–88
- Tipping M E 2001 *J. Mach. Learn. Res.* **1** 211–44
- Wipf D and Nagarajan S 2009 *NeuroImage* **44** 947–66
- Wipf D P and Rao B D 2004 *IEEE Trans. Signal Process.* **52** 2153–64
- Zhang Z and Rao B D 2011 *IEEE J. Sel. Top. Signal Process.* **5** 912–26
- Zou H 2006 *J. Am. Stat. Assoc.* **101** 1418–29