

SPARSE BAYESIAN REGULARIZATION USING BERNOULLI-LAPLACIAN PRIORS

Lotfi Chaari, Jean-Yves Tournet and Hadj Batatia

University of Toulouse, IRIT - INP-ENSEEIH, France
firstname.lastname@enseeiht.fr

ABSTRACT

Sparse regularization has been receiving an increasing interest in the literature. Two main difficulties are encountered when performing sparse regularization. The first one is how to fix the parameters involved in the regularization algorithm. The second one is to optimize the inherent cost function that is generally non differentiable, and may also be non-convex if one uses for instance an ℓ_0 penalization. In this paper, we handle these two problems jointly and propose a novel algorithm for sparse Bayesian regularization. An interesting property of this algorithm is the possibility of estimating the regularization parameters from the data. Simulation performed with 1D and 2D restoration problems show the very promising potential of the proposed approach. An application to the reconstruction of electroencephalographic signals is finally investigated.

Index Terms— Sparse Bayesian restoration, MCMC methods, parameter estimation, $\ell_0 + \ell_1$ regularization

1. INTRODUCTION

Sparse signal and image restoration has been of increasing interest during the last decades. It has found several fields of applications such as remote sensing [1] and medical image reconstruction [2], especially after the emergence of the compressed sensing theory [3]. Data volumes are continuously increasing, e.g., in recent applications where imaging systems may deliver multidimensional signals up to 4D (3D + time for instance) [?]. Accounting for sparsity properties while reconstructing such signals is therefore of great interest. Since observation systems are generally ill-posed, regularization is usually requested to improve quality of reconstructed signals. Regularization consists of constraining the search space through some prior information that we inject in the model to stabilize the inverse problem. Such prior information generally involves additional parameters that have to be tuned. Fixing these parameters is actually an open issue, since they deeply impact the target solution quality. One can fix them by cross-validation or by using some Bayesian methods such

as [4]. However, the estimation of these parameters still relies on an external algorithm, which makes the regularization problem not fully automatic. After estimating these parameters, the cost function associated with the restoration problem has to be optimized. In the recent literature, variational approaches have been widely used to solve the corresponding optimization problem. They generally rely on some iterative optimization algorithm since most of the edge-preserving penalizations yield to non-differentiable cost functions whose extrema have no closed-form expression. In the recent image processing literature [5, ?, 6], proximal algorithms have been notably investigated such as forward-backward (FB) [6] and parallel proximal algorithms (PPXA) [7]. However, these algorithms can only handle convex cost functions such as those involving ℓ_1 or $\ell_1 + \ell_2$ penalizations [8]. In other words, regularization problems involving an ℓ_0 pseudo-norm penalization cannot be solved using these algorithms.

In this paper, we propose a novel approach to handle this ℓ_0 regularization problem in a Bayesian framework. The main advantage of our Bayesian restoration method is that the regularization parameters can be estimated from the data, allowing the sparsity level of the target signal/image to be determined. The proposed method is fully automatic and does not need any user interaction. Indeed, Bayesian method have been widely promoted during the last decades due to their flexibility in handling complicated models, especially when model hyperparameters are not easy to set. For instance, such investigation is clearly stated in the biomedical [9] and hyperspectral imaging [10] fields.

In recent literature, different priors have been used for sparse regularization such as Bernoulli-Gaussian [11] or Bernoulli-exponential priors [12]. The Bernoulli-Laplacian prior investigated in this work provides sparser solutions when compared to the Bernoulli-Gaussian one. This increased sparsity is due to the Laplacian term, while accounting for both positive and negative signal values in contrast to the Bernoulli-exponential prior.

The paper is organized as follows. In Section 2, the proposed model and inference scheme are introduced. Experimental validations are presented in Section 3. Conclusions and future work are reported in Section 4.

The authors would like to thank Yoann Altmann for helping them in the Matlab implementation, as well as Prof. Alexandre Gramfort for his help in EEG experiments.

2. BAYESIAN RESTORATION METHOD

2.1. Problem formulation

Let $\mathbf{x} \in \mathbb{R}^M$ be our target signal, which is measured by $\mathbf{y} \in \mathbb{R}^P$ through a linear observation operator \mathcal{H} . Accounting for the additive acquisition noise, the observation model we are interested in can be written as

$$\mathbf{y} = \mathcal{H}\mathbf{x} + \mathbf{n}. \quad (1)$$

Without loss of generality, we only focus here on multiplicative linear operators and additive noise. When \mathcal{H} is ill-conditioned, the above inverse problem may be ill-posed. We propose here to adopt a sparse regularization strategy for estimating the unknown signal/image \mathbf{x} via a Bayesian framework. In the following section, the hierarchical Bayesian model used for regularization is detailed.

2.2. Hierarchical Bayesian model

2.2.1. Likelihood

Under the assumption of additive Gaussian noise of variance σ_n^2 , the likelihood can be expressed as follows:

$$f(\mathbf{y}|\mathbf{x}, \sigma_n^2) = \left(\frac{1}{2\pi\sigma_n^2}\right)^{P/2} \exp\left(-\frac{\|\mathbf{y} - \mathcal{H}\mathbf{x}\|^2}{2\sigma_n^2}\right) \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm.

2.2.2. Priors

In our model, the unknown parameter vector to be estimated is denoted by $\boldsymbol{\theta} = \{\mathbf{x}, \sigma_n^2\}$. In what follows, we introduce the prior distributions to be used for these two parameters.

Prior for \mathbf{x}

In order to promote the sparsity of the target signal, we adopt here a Bernoulli-Laplace prior for every x_i ($i = 1, \dots, M$), given by:

$$f(x_i|\omega, \lambda) = (1 - \omega)\delta(x_i) + \frac{\omega}{2\lambda} \exp\left(-\frac{|x_i|}{\lambda}\right) \quad (3)$$

where $\delta(\cdot)$ is the Dirac delta function, $\lambda > 0$ is the parameter of the Laplace distribution, and w is a weight belonging to $[0, 1]$. This prior is similar to the one used in [12]. However, since we are considering signals with both positive and negative coefficients, the exponential distribution is replaced here by the Laplace one. Assuming the coefficients x_i a priori independent, the prior distribution for \mathbf{x} writes:

$$f(\mathbf{x}|\omega, \lambda) = \prod_{i=1}^M f(x_i|\omega, \lambda). \quad (4)$$

Prior for σ_n^2

To guarantee the positivity of σ_n^2 and keep this prior non-informative, we use here a Jeffrey's prior defined as:

$$f(\sigma_n^2) \propto \frac{1}{\sigma_n^2} \mathbb{1}_{\mathbb{R}^+}(\sigma_n^2) \quad (5)$$

where $\mathbb{1}_{\mathbb{R}^+}$ is the indicator function on \mathbb{R}^+ , i.e., $\mathbb{1}_{\mathbb{R}^+}(\xi) = 1$ if $\xi \in \mathbb{R}^+$ and 0 otherwise. Motivations for using this kind of prior for the noise variance can be found in standard textbooks on Bayesian inference such as [13].

2.2.3. Hyperparameter priors

Hyperprior for ω

For simplicity reasons, and to use non-informative priors, we use here a uniform distribution on the simplex $[0, 1]$ for ω , i.e., $\omega \sim \mathcal{U}_{[0,1]}$.

Hyperprior for λ

Since λ is real positive, a conjugate inverse-gamma (IG) distribution has been used as a hyper-prior:

$$f(\lambda|\alpha, \beta) = \mathcal{IG}(\lambda|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{-\alpha-1} \exp\left(-\frac{\beta}{\lambda}\right) \quad (6)$$

where $\Gamma(\cdot)$ is the gamma function, and α and β are hyperparameters to be fixed (in our experiments these hyperparameters have been set to $\alpha = \beta = 10^{-3}$).

2.3. Bayesian inference scheme

We adopt here a maximum a posteriori (MAP) strategy in order to estimate the model parameter vector $\boldsymbol{\theta}$ based on the likelihood, the priors and hyperpriors introduced hereabove. If we denote by $\Phi = \{\lambda, \omega\}$ the model hyperparameters, the joint posterior distribution of $\{\boldsymbol{\theta}, \Phi\}$ can be expressed as

$$f(\boldsymbol{\theta}, \Phi|\mathbf{y}, \alpha, \beta) \propto f(\mathbf{y}|\boldsymbol{\theta})f(\boldsymbol{\theta}|\Phi)f(\Phi|\alpha, \beta). \quad (7)$$

Akin to [12], we propose here to use a Gibbs algorithm [13] that iteratively samples according to the conditional posteriors $f(\mathbf{x}|\mathbf{y}, \omega, \lambda, \sigma_n^2)$, $f(\sigma_n^2|\mathbf{y}, \mathbf{x})$, $f(\lambda|\mathbf{x}, \alpha, \beta)$ and $f(\omega|\mathbf{x})$.

2.3.1. Sampling according to $f(\sigma_n^2|\mathbf{y}, \mathbf{x})$

Straightforward calculations combining the likelihood and the prior distribution of σ_n^2 lead to the following posterior:

$$\sigma_n^2|\mathbf{x}, \mathbf{y} \sim \mathcal{IG}\left(\sigma_n^2|P/2, \|\mathbf{y} - \mathcal{H}\mathbf{x}\|^2/2\right) \quad (8)$$

which is easy to sample.

2.3.2. Sampling according to $f(\lambda|\mathbf{x}, \alpha, \beta)$

Calculations similar to [12] lead to the following posterior which is also simple to sample:

$$\lambda|\mathbf{x}, \alpha, \beta \sim \mathcal{IG}\left(\lambda|\alpha + \|\mathbf{x}\|_0, \beta + \|\mathbf{x}\|_1\right) \quad (9)$$

where $\|\cdot\|_0$ denotes the ℓ_0 pseudo-norm calculating the number of non-zero coefficients, and $\|\cdot\|_1$ is the ℓ_1 norm defined as $\|\mathbf{x}\|_1 = \sum_{i=1}^M |x_i|$

2.3.3. Sampling according to $f(\omega|\mathbf{x})$

Straightforward calculations show that the posterior of ω is a beta distribution:

$$\omega \sim \mathcal{B}(1 + \|\mathbf{x}\|_0, 1 + M - \|\mathbf{x}\|_0) \quad (10)$$

according to which it is easy to sample.

2.3.4. Sampling according to $f(\mathbf{x}|\mathbf{y}, \omega, \lambda, \sigma_n^2)$

We can easily derive the posterior distribution of each signal element x_i conditionally to the rest of the signal. Straightforward computations lead to the following form of this posterior

$$f(x_i|\mathbf{y}, \mathbf{x}_{-i}, \omega, \lambda) = \omega_{1,i} \delta(x_i) + \omega_{2,i} \mathcal{N}^+(\mu_i^+, \sigma_i^2) + \omega_{3,i} \mathcal{N}^-(\mu_i^-, \sigma_i^2) \quad (11)$$

where \mathcal{N}^+ and \mathcal{N}^- denote the truncated Gaussian distribution on \mathbb{R}^+ and \mathbb{R}^- , respectively. By decomposing \mathbf{x} on the orthonormal basis $B = \{\mathbf{e}_1, \dots, \mathbf{e}_M\}$ such that $\mathbf{x} = \tilde{\mathbf{x}}_{-i} + x_i \mathbf{e}_i$ where $\tilde{\mathbf{x}}_{-i}$ is nothing but \mathbf{x} whose i^{th} element is set to 0, and denoting $\mathbf{v}_i = \mathbf{y} - \mathcal{H}\mathbf{x}_{-i}$ and $\mathbf{h}_i = \mathcal{H}\mathbf{e}_i$, the weights $(\omega_{l,i})_{1 \leq l \leq 3}$ are given by

$$\omega_{l,i} = \frac{u_{l,i}}{\sum_{l=1}^3 u_{l,i}} \quad (12)$$

where

$$\begin{aligned} u_{1,i} &= 1 - \omega \\ u_{2,i} &= \frac{\omega}{2\lambda} \exp\left(\frac{\mu_{i+}^2}{2\sigma_i^2}\right) \sqrt{2\pi\sigma_i^2} C(\mu_{i+}, \sigma_i^2) \\ u_{3,i} &= \frac{\omega}{2\lambda} \exp\left(\frac{\mu_{i-}^2}{2\sigma_i^2}\right) \sqrt{2\pi\sigma_i^2} C(-\mu_{i-}, \sigma_i^2) \end{aligned} \quad (13)$$

and

$$\begin{aligned} \sigma_i^2 &= \frac{\sigma_n^2}{\|\mathbf{h}_i\|^2} \\ \mu_{i+} &= \sigma_i^2 \left(\frac{\mathbf{h}_i^T \mathbf{v}_i}{\sigma_n^2} - \frac{1}{\lambda} \right), \\ \mu_{i-} &= \sigma_i^2 \left(\frac{\mathbf{h}_i^T \mathbf{v}_i}{\sigma_n^2} + \frac{1}{\lambda} \right) \\ C(\mu, \sigma^2) &= \sqrt{\frac{\sigma^2 \pi}{2}} \left(1 + \operatorname{erf}\left(\frac{\mu}{\sqrt{2\sigma^2}}\right) \right). \end{aligned} \quad (14)$$

The main steps of the proposed sampling algorithm are summarized in Algorithm 1.

Algorithm 1 Gibbs sampler.

```

Initialize with some  $\mathbf{x}^{(0)}$ 
repeat
  Sample  $\sigma_n^2$  according to Eq. (8).
  Sample  $\lambda$  according to Eq. (9).
  Sample  $\omega$  according to Eq. (10).
  for  $i = 1$  to  $M$  do
    Sample  $x_i$  according to Eq. (11).
  end for
until convergence

```

After convergence, the proposed algorithm ends up with sampled sequences that will be used to compute the minimum mean square error (MMSE) estimator of the unknown parameter vector, allowing us to compute the estimated signal $\hat{\mathbf{x}}$, in addition to $\hat{\sigma}_n^2$, $\hat{\lambda}$ and $\hat{\omega}$.

3. EXPERIMENTAL VALIDATION

In order to validate the proposed method for sparse signal and image restoration, three experiments have been conducted. The first two experiments correspond to 1D and 2D signal restorations. The third experiment handles an electroencephalography (EEG) reconstruction problem. Since we are in a simulation context, results are evaluated in terms of signal to noise ratio (SNR) given by $20 \log_{10} \frac{\|\mathbf{x}^0\|}{\|\mathbf{x}^0 - \hat{\mathbf{x}}\|}$, where \mathbf{x}^0 and $\hat{\mathbf{x}}$ are the reference and estimated signals, respectively.

3.1. 1D signal restoration

In this experiment, a 1D sparse signal \mathbf{x} of size 100 is recovered from its distorted version \mathbf{y} observed according to the model in Eq. (1), where the observation operator \mathcal{H} is the second order difference operator. The observation \mathbf{y} has been simulated by adding a Gaussian noise \mathbf{n} of variance $\sigma_n^2 = 1$. The regularization scheme detailed in Section 2 is used with the same sparsity promoting priors and hyperparameter setting. Fig. 1 shows the original signal (Reference), the restored signal obtained with the proposed algorithm (MCMC) and the

restored signal obtained with a PPXA algorithm (referred to as ℓ_1) whose parameters have been initialized by the output of the proposed algorithm.

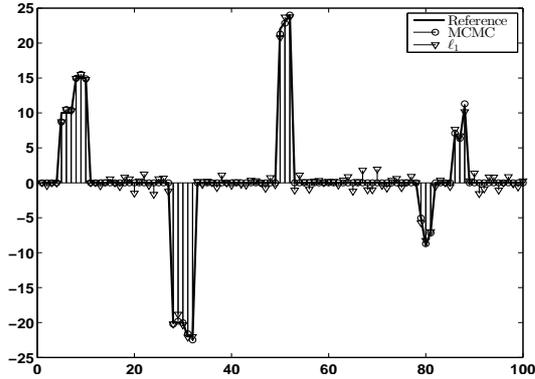


Fig. 1. Original and restored signals using the proposed method and ℓ_1 regularization.

The proposed method provides the smallest estimation error since the corresponding markers are always closer to the ground truth signal. Quantitatively speaking, with an initial SNR of 5.71 dB (observed signal), we achieve an SNR = 31.33 dB with our method, while the ℓ_1 regularization can only reach SNR = 20.05 dB. In addition to the automatic estimation of the regularization parameter, the proposed method ensures a better sparsity of the target signal due to the Bernoulli-Laplace prior, which is equivalent to an $\ell_0 + \ell_1$ regularization.

Regarding parameter estimation, Fig. 2 illustrates the posterior distributions of parameters σ_n^2 , λ and ω , as well as their estimated values. We can easily notice that the estimation of σ_n^2 is very accurate since the estimated value (0.98) is very close to the reference one which is indicated by a vertical line in Fig. 2[top-right]. Moreover, the restored signal estimated with our method shows higher sparsity level since $\|\hat{\mathbf{x}}_{\text{MCMC}}\|_0 = 20$ and $\|\hat{\mathbf{x}}_{\ell_1}\|_0 = 100$.

Finally, it is interesting to note that these results have been obtained after 600 iterations including a burn-in period of 200 iterations. The requested computation time using a Matlab implementation on a 64-bit 2.00GHz i7-3667U architecture was about 22 seconds.

3.2. 2D image restoration

In this experiment, a 2D sparse image \mathbf{x} of size 26×26 is recovered from its distorted version \mathbf{y} observed according to the model in Eq. (1) (with the same observation operator and noise level as in the first experiment). Fig. 3 illustrates the original, observed, and reconstructed images using our method and an ℓ_1 regularization.

The same conclusions apply to this example. The reported SNR values corroborate the clear advantage of the autocalibrated $\ell_0 + \ell_1$ Bayesian regularization made possible by our

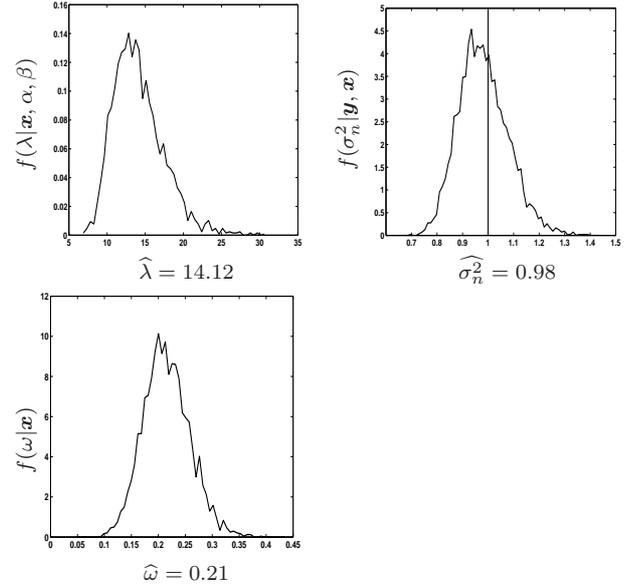


Fig. 2. Estimated posterior distributions of parameters σ_n^2 , λ and ω .

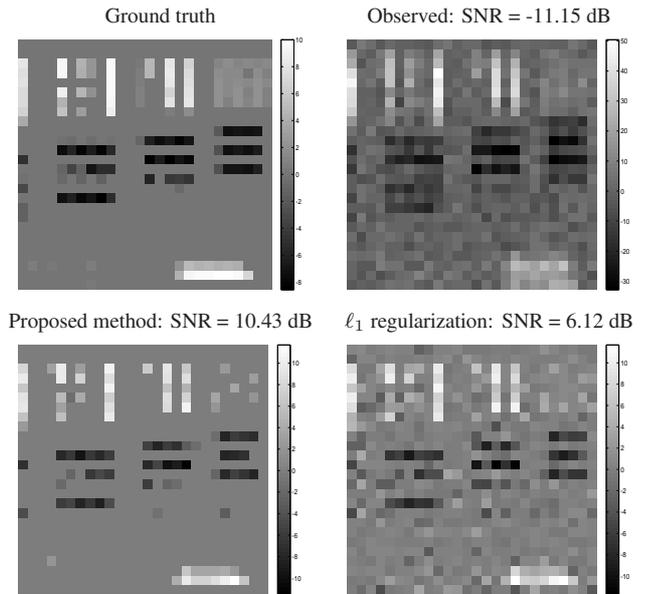


Fig. 3. Original, observed and restored images using the proposed method and ℓ_1 regularization.

algorithm, compared to the variational ℓ_1 regularization (performed using parameters estimated by our algorithm). Sparsity levels of the estimated images show also the improvement obtained with our method since $\|\hat{\mathbf{x}}_{\text{MCMC}}\|_0 = 103$ and $\|\hat{\mathbf{x}}_{\ell_1}\|_0 = 660$. For this experiment, the computational time with the same implementation and number of iterations as in the previous experiment was about 94 seconds.

3.3. EEG reconstruction

The last experiment addresses an EEG reconstruction problem where the observed signals correspond to the activity measured by each electrode during the acquisition time. Using the MNE software¹, we simulated a small EEG dataset using 20 electrodes with 35 sources (35 voxels on the brain surface), where only 4 of them have been chosen to be active. The simulation involved 21 time points, which means that we have to recover a 35×21 image \mathbf{x} from an observation \mathbf{y} of size 20×21 . The linear operator here represents simply the brain model geometry. Fig. 4 displays the ground truth and reconstructed images using our method and an ℓ_1 regularization. Visually speaking, our method gives sparser recovered signal while preserving most of the activated sources. The obtained SNR values show also that the proposed algorithm outperforms the ℓ_1 regularization from a quantitative viewpoint. As regards sparsity levels, the same conclusion as for the first two experiments applies, since we have $\|\hat{\mathbf{x}}_{\text{MCMC}}\|_0 = 51$ and $\|\hat{\mathbf{x}}_{\ell_1}\|_0 = 730$.

Computational time with the same implementation and number of iterations as in the previous two experiments was about 95 seconds.

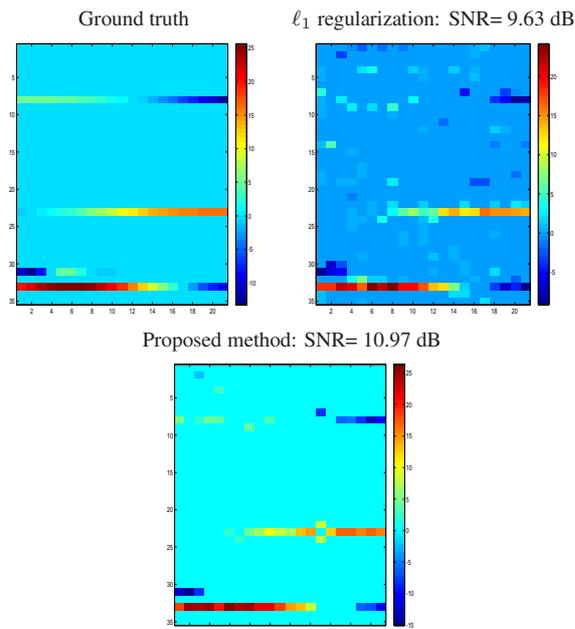


Fig. 4. Original, observed and restored images using the proposed method and ℓ_1 regularization.

4. CONCLUSION

We proposed in this paper a method for Bayesian sparse regularization using a Bernoulli-Laplace prior. This prior makes the $\ell_0 + \ell_1$ non-convex regularization problem feasible in a

Bayesian framework, in contrast to variational methods which require the convexity of the cost function. In addition, the proposed method estimates the regularization parameters directly from the data. Validation on 1D and 2D restoration problems, as well as in EEG reconstruction, show the efficiency of the proposed method for recovering sparse signals. As a perspective, we will handle the problem of correlation between signal coefficients. Future work will also focus on the validation of our method on real EEG signals, as well as for other real world applications such as image deconvolution.

5. REFERENCES

- [1] C. Chaux, P. L. Combettes, J.-C. Pesquet, and V. Wajs, "Iterative image deconvolution using overcomplete representations," in *European Signal Processing Conference*, Florence, Italy, Sep. 4-8 2006, 5 pp.
- [2] L. Chaari, J.-C. Pesquet, A. Benazza-Benyahia, and P. Ciuciu, "A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging," *Med. Image Anal.*, vol. 15, no. 2, pp. 185–201, Nov. 2011.
- [3] D.L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [4] L. Chaari, J.-C. Pesquet, J.-Y. Tourneret, Ph. Ciuciu, and A. Benazza-Benyahia, "A hierarchical Bayesian model for frame representation," *IEEE Trans. on Signal Process.*, vol. 18, no. 11, pp. 5560–5571, Nov. 2010.
- [5] L. Chaari, J.-C. Pesquet, A. Benazza-Benyahia, and P. Ciuciu, "Minimization of a sparsity promoting criterion for the recovery of complex-valued signals," in *Signal Processing with Adaptive Sparse Structured Representations Workshop (SPARS)*, Saint-Malo, France, April 6-9 2009.
- [6] C. Chaux, P. Combettes, J.-C. Pesquet, and V.R. Wajs, "A variational formulation for frame-based inverse problems," *Inverse Problems*, vol. 23, no. 4, pp. 1495–1518, Aug. 2007.
- [7] P. L. Combettes and J.-C. Pesquet, "A proximal decomposition method for solving convex variational inverse problems," *Inverse Problems*, vol. 24, no. 6, Dec. 2008, 27 p.
- [8] A. Gramfort, D. Strohmeier, J. Haueisen, M.S. Hmlinen, and M. Kowalski, "Time-frequency mixed-norm estimates: Sparse M/EEG imaging with non-stationary source activations," *NeuroImage*, vol. 70, pp. 410–422, 2013.
- [9] L. Chaari, T. Vincent, F. Forbes, M. Dojat, and P. Ciuciu, "Fast joint detection-estimation of evoked brain activity in event-related fMRI using a variational approach," *IEEE Trans. on Medical Imaging*, vol. 32, no. 5, pp. 821–837, May 2013.
- [10] Y. Altmann, N. Dobigeon, S. McLaughlin, and J.-Y. Tourneret, "Non-linear spectral unmixing of hyperspectral images using Gaussian processes," *IEEE Trans. Signal Processing*, vol. 61, no. 10, pp. 2442–2453, May 2013.
- [11] M. Lavielle, "Bayesian deconvolution of Bernoulli-Gaussian processes," *Signal Process.*, vol. 33, no. 1, pp. 67–79, Jul. 2009.
- [12] N. Dobigeon, A. O. Hero, and J.-Y. Tourneret, "Hierarchical Bayesian sparse image reconstruction with application to MRFM," *IEEE Trans. Image Process.*, vol. 18, no. 9, pp. 2059–2070, Sept. 2009.
- [13] C. Robert and G. Casella, *Monte Carlo statistical methods*, Springer, New York, 2004.

¹<http://www.martinos.org/mne/>