



Sparse estimation using Bayesian hierarchical prior modeling for real and complex linear models

Niels Lovmand Pedersen^{a,*}, Carles Navarro Manchón^a, Mihai-Alin Badiu^a,
Dmitriy Shutin^b, Bernard Henri Fleury^a

^a Department of Electronic Systems, Aalborg University, Niels Jernes Vej 12, DK-9220 Aalborg, Denmark

^b Institute of Communications and Navigation, German Aerospace Center, Oberpfaffenhofen, D-82234 Wessling, Germany

ARTICLE INFO

Article history:

Received 8 October 2014

Received in revised form

2 March 2015

Accepted 12 March 2015

Available online 20 March 2015

Keywords:

Sparse Bayesian learning

Sparse signal representations

Underdetermined linear systems

Hierarchical Bayesian modeling

Sparsity-inducing priors

ABSTRACT

In sparse Bayesian learning (SBL), Gaussian scale mixtures (GSMs) have been used to model sparsity-inducing priors that realize a class of concave penalty functions for the regression task in real-valued signal models. Motivated by the relative scarcity of formal tools for SBL in complex-valued models, this paper proposes a GSM model – the Bessel K model – that induces concave penalty functions for the estimation of complex sparse signals. The properties of the Bessel K model are analyzed when it is applied to Type I and Type II estimation. This analysis reveals that, by tuning the parameters of the mixing pdf different penalty functions are invoked depending on the estimation type used, the value of the noise variance, and whether real or complex signals are estimated. Using the Bessel K model, we derive sparse estimators based on a modification of the expectation–maximization algorithm formulated for Type II estimation. The estimators include as special instances the algorithms proposed by Tipping and Faul [1] and Babacan et al. [2]. Numerical results show the superiority of the proposed estimators over these state-of-the-art algorithms in terms of convergence speed, sparseness, reconstruction error, and robustness in low and medium signal-to-noise ratio regimes.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Compressive sensing and sparse signal representation have attracted the interest of an increasing number of researchers over the recent years [3–6]. This is motivated by the widespread applicability that such techniques have found in a large variety of engineering disciplines. Generally speaking, these disciplines consider the following signal model:

$$\mathbf{y} = \Phi \mathbf{w} + \mathbf{n}. \quad (1)$$

In this expression, \mathbf{y} is an $M \times 1$ vector of measurement samples, $\Phi = [\phi_1, \dots, \phi_N]$ is an $M \times N$ dictionary matrix with $N > M$. The additive term \mathbf{n} is an $M \times 1$ perturbation vector, which is assumed to be Gaussian distributed with zero-mean and covariance $\lambda^{-1} \mathbf{I}$, where $\lambda > 0$ denotes the noise precision and \mathbf{I} is the identity matrix. The objective is to accurately estimate the $N \times 1$ unknown weight vector $\mathbf{w} = [w_1, \dots, w_N]^T$, which is assumed K -sparse in the canonical basis.

We coin the signal model (1) as either real, when Φ , \mathbf{w} , and \mathbf{n} are all real, or as complex, when Φ , \mathbf{w} , and \mathbf{n} are all complex.¹ Historically, real signal models have dominated

* Corresponding author. Tel.: +45 99408616.

E-mail addresses: nlp@es.aau.dk (N.L. Pedersen), cnm@es.aau.dk (C. Navarro Manchón), in_mba@es.aau.dk (M.-A. Badiu), dmitriy.shutin@dlr.de (D. Shutin), bfl@es.aau.dk (B.H. Fleury).

¹ Obviously, one could also consider a mixed model where, e.g., Φ and \mathbf{n} are complex but \mathbf{w} is real. In this paper we focus on the two most relevant cases of real and complex signal models as defined earlier.

the research in sparse signal representation and compressive sensing. However, applications seeking sparse estimation for complex signal models are not uncommon. An example is the estimation of multipath wireless channels [6–9]. The extension of sparse representation from real signal models to complex models is not always straightforward, as we will discuss in this paper.

Many convex [10,11], greedy [12,13], and Bayesian methods have been proposed in the literature in recent years to devise sparse estimators. In this paper, we focus on Bayesian inference methods commonly referred to as sparse Bayesian learning (SBL) [14,15]. In SBL, we design priors for \mathbf{w} that induce sparse representations of $\Phi\mathbf{w}$. Instead of working directly with the prior probability density function (pdf) $p(\mathbf{w})$, SBL typically uses a two-layer hierarchical prior model that involves a conditional prior pdf $p(\mathbf{w}|\gamma)$ and a hyperprior pdf $p(\gamma)$. The goal is to select these pdfs in such a way that we can construct computationally tractable iterative algorithms that estimate both the hyperparameter vector γ and the weight vector \mathbf{w} with the latter estimate being sparse. A widely used two-layer prior model is the model where the entries of \mathbf{w} are independent and identically distributed according to a Gaussian scale mixture (GSM) [16–20]. Specifically, w_i is modeled as $w_i = \sqrt{\gamma_i}u_i$, where u_i is a standard Gaussian random variable and γ_i is a non-negative random scaling factor, also known as the mixing variable.² The pdf $p(\gamma_i)$ of the latter variable is called the mixing pdf of the GSM. Based on a careful selection of $p(\gamma_i)$, an inference algorithm is then constructed. The sparsity-inducing property of the resulting estimator does not only depend on $p(\gamma_i)$ but also on the type of inference method used, as discussed next.

In SBL two widespread inference approaches, referred to as Type I and Type II estimation following [21], have been used. In Type I estimation, the maximum-a-posteriori (MAP) estimate of \mathbf{w} is computed from the observation \mathbf{y} :

$$\begin{aligned}\hat{\mathbf{w}}_I(\mathbf{y}) &= \underset{\mathbf{w}}{\operatorname{argmax}} p(\mathbf{w}|\mathbf{y}) \\ &= \underset{\mathbf{w}}{\operatorname{argmax}} \log \int p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)p(\gamma) d\gamma.\end{aligned}\quad (2)$$

Equivalently, the Type I estimator $\hat{\mathbf{w}}_I$ is obtained as the minimizer of the Type I cost function

$$\mathcal{L}_I(\mathbf{w}) \triangleq \rho \|\mathbf{y} - \Phi\mathbf{w}\|_2^2 + \lambda^{-1} q_I(\mathbf{w}). \quad (3)$$

In the above expression, $\|\cdot\|_p$, $p \geq 1$, is the ℓ_p -norm and the parameter ρ takes values $\rho = 1/2$ when the signal model (1) is real and $\rho = 1$ when it is complex. The pdf $p(\gamma)$ is designed such that the penalization term $q_I(\mathbf{w}) \propto e^{-\log p(\mathbf{w})}$ with $p(\mathbf{w}) = \int p(\mathbf{w}|\gamma)p(\gamma) d\gamma$ enforces a sparse estimate of the weight vector \mathbf{w} .³

In Type II estimation [22,14,15], the MAP estimate of γ is computed from the observation \mathbf{y} :

$$\hat{\gamma}_{II}(\mathbf{y}) = \underset{\gamma}{\operatorname{argmax}} p(\gamma|\mathbf{y})$$

$$= \underset{\gamma}{\operatorname{argmax}} \log \int p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)p(\gamma) d\mathbf{w}. \quad (4)$$

Thus, the estimator $\hat{\gamma}_{II}$ is the minimizer of

$$\mathcal{L}_{II}(\gamma) \triangleq \rho \mathbf{y}^H \mathbf{C}^{-1} \mathbf{y} + \rho \log |\mathbf{C}| - \log p(\gamma) \quad (5)$$

with $\mathbf{C} \triangleq \lambda^{-1} \mathbf{I} + \Phi \Gamma \Phi^H$ and $\Gamma = \operatorname{diag}(\gamma)$. The Type II estimator of \mathbf{w} follows as

$$\hat{\mathbf{w}}_{II}(\mathbf{y}) = \langle \mathbf{w} \rangle_{p(\mathbf{w}|\mathbf{y}; \hat{\gamma}_{II}(\mathbf{y}))} = \left(\Phi^H \Phi + \lambda^{-1} \hat{\Gamma}_{II}^{-1} \right)^{-1} \Phi^H \mathbf{y}, \quad (6)$$

where $\hat{\Gamma}_{II} = \operatorname{diag}(\hat{\gamma}_{II}(\mathbf{y}))$ and $\langle \cdot \rangle_{p(\mathbf{x})}$ denotes expectation over the pdf $p(\mathbf{x})$. The impact of $p(\gamma)$ on the estimator $\hat{\mathbf{w}}_{II}$ is not straightforward. This complicates the task of selecting $p(\gamma)$ inducing a sparse estimate of \mathbf{w} . In [21], the relationship between Type I and Type II estimation has been identified. This result makes it possible to compare the two estimation methods. Invoking [21, Theorem 2], $\hat{\mathbf{w}}_{II}(\mathbf{y})$ is equivalently the minimizer of the Type II cost function

$$\mathcal{L}_{II}(\mathbf{w}) \triangleq \rho \|\mathbf{y} - \Phi\mathbf{w}\|_2^2 + \lambda^{-1} q_{II}(\mathbf{w}) \quad (7)$$

with penalty

$$q_{II}(\mathbf{w}) = \min_{\gamma} \{ \rho \mathbf{w}^H \Gamma^{-1} \mathbf{w} + \rho \log |\mathbf{C}| - \log p(\gamma) \}. \quad (8)$$

Specifically, $\hat{\mathbf{w}}_{II}(\mathbf{y})$ in (6) equals the global minimizer of $\mathcal{L}_{II}(\mathbf{w})$ if and only if $\hat{\gamma}_{II}(\mathbf{y})$ equals the global minimizer of $\mathcal{L}_{II}(\gamma)$. Likewise, $\hat{\mathbf{w}}_*(\mathbf{y}) = \langle \mathbf{w} \rangle_{p(\mathbf{w}|\mathbf{y}; \hat{\gamma}_*(\mathbf{y}))}$ is a local minimizer of $\mathcal{L}_{II}(\mathbf{w})$ if and only if $\hat{\gamma}_*(\mathbf{y})$ is a local minimizer of $\mathcal{L}_{II}(\gamma)$.

The MAP estimates in (2) and (4) cannot usually be computed in closed-form and one must resort to iterative inference methods to approximate them. One method is the Relevance Vector Machine (RVM) [14,15]. In RVM the mixing pdf $p(\gamma_i)$ is equal to an improper constant prior. An instance of the expectation-maximization (EM) algorithm is then formulated to approximate the Type II estimator. Another method, devised for real signal models in [23], uses the EM algorithm to approximate two popular Type I estimators with respectively ℓ_1 -norm and log-sum constrained penalization. These penalization terms arise from selecting the mixing pdf equal to an exponential pdf and the non-informative Jeffreys prior, respectively. In the former case, the marginal prior pdf $p(\mathbf{w})$ is the product of Laplace pdfs and $\mathcal{L}_I(\mathbf{w})$ equals the cost function of Least Absolute Shrinkage and Selection Operator (LASSO) [10] or Basis Pursuit Denoising [11].⁴

The sparse estimators in [14,15,23] inherit the limitation of the instances of the EM algorithm that they embed: high computational complexity and slow convergence [1]. To circumvent this shortcoming, a fast inference framework is proposed in [1] for RVM and later applied to derive the Fast Laplace algorithm [2]. The latter algorithm is derived based on an augmented probabilistic model obtained by adding a third layer to the real GSM model of the Laplace pdf; the third layer introduces a hyper-hyperprior for the rate parameter of the exponential pdf, which coincides with the regularization parameter of the ℓ_1 penalization induced by the Laplace prior.

² In this configuration, γ_i can be seen as the variance of w_i .

³ Here $x \propto e^y$ denotes $\exp(x) = \exp(\nu) \exp(y)$, and thus $x = \nu + y$, for some arbitrary constant ν . We will also make use of $x \propto y$, which denotes $x = \nu y$ for some positive constant ν .

⁴ Let us point out that the hierarchical representation resulting in the ℓ_1 -norm presented in [23] is only valid for real-valued variables. In this paper, we extend this representation to cover complex-valued variables as well.

However, as Fast Laplace is based on Type II estimation, it cannot be seen as the adaptive Bayesian version of the ℓ_1 re-weighted LASSO algorithm [24]. The Bayesian version of this latter estimator is proposed in [25,26].

Even though the fast algorithms in [1] and [2] converge faster than their EM counterparts, they still suffer from slow convergence, especially in low and moderate signal-to-noise ratio (SNR) regimes as we will demonstrate in this paper. Furthermore, in these regimes, the algorithms significantly overestimate the number of non-zero weights. We will show that this behavior is, in fact, a consequence of the prior models used to derive the algorithms.

Coming back to the original motivation of this work, though complex GSM models have been proposed in the literature [27,28], they have not been extensively applied within the framework of SBL. An example illustrating this fact is the hierarchical modeling of the ℓ_1 -norm in Type I estimation. While this penalty results from selecting the exponential mixing pdf for the entries in γ in real GSM models, said pdf will not induce the ℓ_1 -norm penalty for complex models. Yet to the best of our knowledge, the complex GSM model realizing the ℓ_1 -norm penalty has not been established in the literature. Moreover, it is not evident what sparsity-inducing property the complex GSM model exhibits when applied in Type II estimation. Motivated by the relative scarcity of formal tools for sparse learning in complex models and inspired by the recent analysis of sparse Bayesian algorithms in [21], we propose and investigate an SBL approach that applies to both real and complex signal models.

Starting in Section 2, we first present a GSM model for both real and complex sparse signal representation where the mixing pdf $p(\gamma_i)$ is selected to be a gamma pdf. When \mathbf{w} is real, the marginal prior pdf $p(\mathbf{w})$ equals the product of Bessel K pdfs [17–19].⁵ We extend the Bessel K model to cover complex weights and model for this extension several penalty functions previously introduced for inferring real sparse weights. One important example is the hierarchical prior modeling inducing the ℓ_1 -norm penalty for complex weights. We then analyze the Type I and Type II estimators derived from the Bessel K model. We show that a sparsity-inducing prior for Type I estimation does not necessarily have this property for Type II estimation and, interestingly, a sparsity-inducing prior for real weights is not necessarily sparsity-inducing for complex weights. In the particular case, where the dictionary matrix Φ is orthonormal, we demonstrate, using the EM algorithm, that Type I and Type II estimators derived using the Bessel K model are generalizations of the soft-thresholding rule with degree of sparseness depending on the selection of the shape parameter of the gamma pdf $p(\gamma_i)$. Additionally, we show that this model has a strong connection to the Bayesian formalism of the group LASSO [26,29]. Note that the Bessel K model has been previously introduced for sparse signal representation [30,31]. However, these works were restricted to the inference of real weights and did not consider the relationship between Type I and Type II estimation.

In Section 3, we propose greedy, low-complexity estimators using the Bessel K model. The estimators are based on a modification of the EM algorithm for Type II estimation. As the Bessel K model encompasses the prior models used in [1] and [2], the iterative algorithms derived in these publications can be seen as instances of our estimators for particular settings of the associated parameters of the gamma mixing pdf.

Section 4 provides numerical results obtained via Monte Carlo simulations that reveal the superior performance of the proposed estimators in terms of convergence speed of the algorithms, sparseness, and mean-squared error (MSE) of the estimates. Furthermore, and of great importance in many engineering areas, the estimators show a significant robustness in low and moderate SNR regimes; a property not exhibited by the traditional SBL estimators, such as [1] and [2], and other state-of-the-art non-Bayesian sparse estimators. This result opens for a potential application of our estimators in systems operating in these SNR regimes, e.g., receivers in wireless communications [7,8]. Furthermore, the proposed estimators can inherently incorporate the estimation of the noise variance. In the literature, this parameter is often learned from a training procedure or tuned for optimality. Since the algorithms in [1] and [2] only differ from ours in the choice of parameters of the mixing pdf, we can safely conclude that the observed performance benefits are a direct consequence of our proposed prior model.

Finally, we conclude the paper in Section 5.

2. The Bessel K model for real and complex signal representation

In this section we present the Bessel K model for SBL. We first state the probabilistic model of the signal model (1). Based on this probabilistic model we analyze the Type I and Type II cost functions. We then show how to obtain various estimators with different sparsity-inducing properties by appropriately setting the parameters of the Bessel K model.

2.1. Probabilistic model

We begin with the specification of the probabilistic model for (1) augmented with the two-layer prior model for \mathbf{w} :

$$p(\mathbf{y}, \mathbf{w}, \gamma) = p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)p(\gamma). \quad (9)$$

From (1), $p(\mathbf{y}|\mathbf{w}) = \mathcal{N}(\mathbf{y}; \Phi\mathbf{w}, \lambda^{-1}\mathbf{I})$ if the signal model is real and $p(\mathbf{y}|\mathbf{w}) = \mathcal{CN}(\mathbf{y}; \Phi\mathbf{w}, \lambda^{-1}\mathbf{I})$ if the model is complex.⁶

The sparsity constraints on \mathbf{w} are determined by the joint prior pdf $p(\mathbf{w}|\gamma)p(\gamma)$. Motivated by previous works on GSM modeling and SBL [14,15,23] we select the conditional prior pdf $p(\mathbf{w}|\gamma)$ to factorize in a product of zero-mean Gaussian pdfs: $p(\mathbf{w}|\gamma) = \prod_i p(w_i|\gamma_i)$, where

$$p(w_i|\gamma_i) = \left(\frac{\rho}{\pi\gamma_i}\right)^\rho \exp\left(-\rho\frac{|w_i|^2}{\gamma_i}\right). \quad (10)$$

⁵ The Bessel K pdf is in turn a special case of even a larger class of generalized hyperbolic distributions [17], obtained when the mixing pdf is a Generalized Inverse Gaussian pdf.

⁶ $\mathcal{N}(\cdot; \mathbf{a}, \mathbf{B})$ and $\mathcal{CN}(\cdot; \mathbf{a}, \mathbf{B})$ denote, respectively, a multivariate real and a multivariate complex Gaussian pdf with mean vector \mathbf{a} and covariance matrix \mathbf{B} . We shall also make use of the gamma pdf $\text{Ga}(\cdot; a, b) = \frac{b^a}{\Gamma(a)} x^{a-1} \exp(-bx)$ with shape parameter a and rate parameter b .

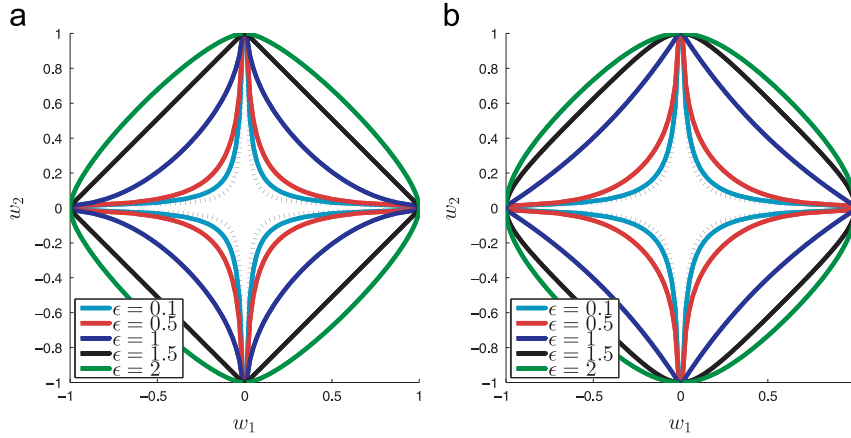


Fig. 1. One contour line of the restriction to \mathbb{R}^2 of (a) $q_I(w_1, w_2; \epsilon, \eta = 1)$ and (b) $q_{II}(w_1, w_2; \epsilon, \eta = 1)$ for selected values of ϵ . In (b) Φ is orthonormal and $\lambda^{-1} = 1/4$. The gray dashed lines depict the contour lines corresponding to the setting $\epsilon = \eta = 0$, i.e., the mixing density equals the Jeffreys prior. (a) Type I (b) Type II.

In the above expression, $\rho = 1/2$ when \mathbf{w} is real and $\rho = 1$ when \mathbf{w} is complex. We choose the mixing pdf $p(\gamma)$ to be a product of identical gamma pdfs, i.e., $p(\gamma) = \prod_i p(\gamma_i; \epsilon, \eta)$ with $p(\gamma_i; \epsilon, \eta) \triangleq \text{Ga}(\gamma_i; \epsilon, \eta)$. The prior pdf for \mathbf{w} is then given by $p(\mathbf{w}; \epsilon, \eta) = \int p(\mathbf{w}|\gamma)p(\gamma; \epsilon, \eta) d\gamma = \prod_i p(w_i; \epsilon, \eta)$ with

$$p(w_i; \epsilon, \eta) = \frac{2(\rho\eta)^{\frac{(\epsilon+\rho)}{2}}}{\pi^\rho \Gamma(\epsilon)} |w_i|^{\epsilon-\rho} K_{\epsilon-\rho}(2\sqrt{\rho\eta}|w_i|). \quad (11)$$

In this expression, $K_\nu(\cdot)$ is the modified Bessel function of the second kind and order $\nu \in \mathbb{R}$. In case \mathbf{w} is real ($\rho = 1/2$), we obtain the GSM model of the Bessel K pdf [17,18]. We will keep the same terminology when \mathbf{w} is complex ($\rho = 1$).⁷ The Bessel K pdf (11) represents a family of prior pdfs for \mathbf{w} parametrized by ϵ and η . By selecting different values for ϵ and η , we realize various penalty functions for Type I and Type II estimation as shown in the following.

2.2. Type I cost function

The Type I cost function $\mathcal{L}_I(\mathbf{w})$ induced by the Bessel K model is given by (3) with penalty $q_I(\mathbf{w}) = \sum_i q_I(w_i; \epsilon, \eta)$ where $q_I(w_i; \epsilon, \eta) \triangleq -\log(|w_i|^{\epsilon-\rho} K_{\epsilon-\rho}(2\sqrt{\rho\eta}|w_i|))$. (12)

Special cases of Type I penalties resulting from the Bessel K pdf have already been considered in the literature for sparse regression when the weights are real [30,31]. We review them together with introducing the corresponding extension to complex weights.

2.2.1. The ℓ_1 -norm penalty

This penalty is of particular importance in sparse signal representation as the convex relaxation of the ℓ_0 -norm.⁸

When \mathbf{w} is real, it is well-known that the Laplace prior induces the ℓ_1 -norm penalty. The Bessel K pdf (11) encompasses the Laplace pdf as a special case with the

selection $\epsilon = 1$ and $\rho = 1/2$.⁹

$$p(w_i; \epsilon = 1, \eta) = \sqrt{\frac{\eta}{2}} \exp(-\sqrt{2\eta}|w_i|), \quad w_i \in \mathbb{R}. \quad (13)$$

The Laplace pdf for real weights is thereby the pdf of a GSM with an exponential mixing pdf [16].

The extension of (13) to \mathbf{w} complex is not straightforward. One approach is to treat the real and imaginary parts of each w_i independently with both parts modeled according to the real GSM representation of the Laplace pdf. Doing so using (13) we obtain $p(w_i) = \frac{\eta}{2} \exp(-\sqrt{2\eta}(|\text{Re}\{w_i\}| + |\text{Im}\{w_i\}|))$. Obviously this approach does not lead to the ℓ_1 -norm penalty for Type I estimation.¹⁰ The complex GSM model with a gamma mixing pdf with shape parameter $\epsilon = 3/2$ does induce this penalty. Indeed, with this setting, (11) becomes, for $\rho = 1$,

$$p(w_i; \epsilon = 3/2, \eta) = \frac{2\eta}{\pi} \exp(-2\sqrt{\eta}|w_i|), \quad w_i \in \mathbb{C}. \quad (14)$$

Throughout the paper, we refer to the pdf in (14) as the Laplace pdf for complex weights.

In summary, the Bessel K model induces the ℓ_1 -norm penalty $q_I(\mathbf{w}) = 2\sqrt{\rho\eta} \sum_i |w_i|$ with the selection $\epsilon = \rho + 1/2$. The introduced GSM model of the Laplace pdf for both real and complex variables is strongly connected with the group LASSO and its Bayesian interpretation [26,29], where sparsity is enforced simultaneously over groups of k variables. In the Bayesian interpretation of the group LASSO a gamma pdf with shape parameter $(k+1)/2$ is employed to model the prior for each of the variables in a group. This choice of shape parameter is consistent with the choice of ϵ in the GSM model of the Laplace prior: in the real case a group consists of $k=1$ variable and, thus, $(k+1)/2 = 1$, whereas in the complex case, a group consists of the real and imaginary parts of a complex variable, hence, $k=2$ and $(k+1)/2 = 3/2$.

⁷ To the authors' best knowledge, the GSM model of the Bessel K pdf has only been presented for real variables.

⁸ The ℓ_0 -norm of the vector \mathbf{x} is the number of non-zero entries in \mathbf{x} . Note that by abuse of terminology $\|\cdot\|_0$ is coined a norm even though it does not fulfill all properties of a norm.

⁹ Here, we make use of the identity $K_{1/2}(z) = \sqrt{\frac{\pi}{2z}} \exp(-z)$ [32].

¹⁰ The ℓ_1 -norm for the complex vector \mathbf{x} is defined as $\|\mathbf{x}\|_1 = \sum_i |x_i| = \sum_i \sqrt{\text{Re}^2\{x_i\} + \text{Im}^2\{x_i\}}$ [33,34].

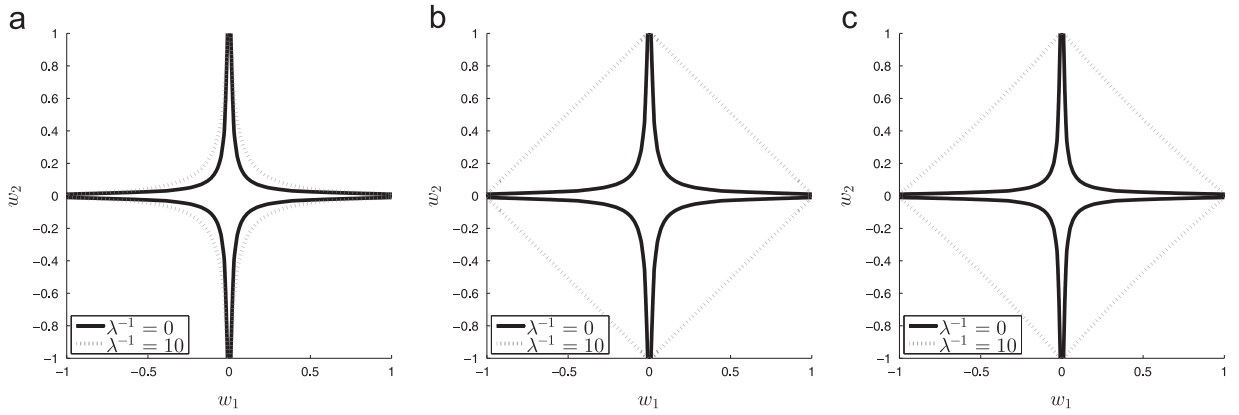


Fig. 2. One contour line of the restriction to \mathbb{R}^2 of (a) $q_{II}(w_1, w_2; \epsilon = 1/2, \eta = 1)$, (b) $q_{II}(w_1, w_2; \epsilon = 1, \eta = 1)$, and (c) $q_{II}(w_1, w_2; \epsilon = 1, \eta = 0)$ with Φ orthonormal and λ^{-1} as a parameter. Note that $q_{II}(w_1, w_2; \epsilon = 1, \eta = 0)$ in (c) coincides with the penalty used in RVM [14,15]. (a) Bessel K, $\epsilon = 1/2$ (b) Bessel K, $\epsilon = 1$ (c) RVM.

2.2.2. The log-sum penalty

The selection $\epsilon = \eta = 0$ in (11) entails the Jeffreys (improper) prior density $p(\gamma_i) \propto \gamma_i^{-1}$ and thereby the improper marginal prior density $p(\mathbf{w}) \propto \prod_i |w_i|^{-2\rho}$. Thus, when the mixing density of the GSM is chosen to be non-informative, the log-sum penalization $q_I(\mathbf{w}) = 2\rho \sum_i \log |w_i|$ is invoked in (3). This penalty has gained much interest in the literature, including [14,15,23,24,35], as it is known to strongly promote sparse estimates.

2.2.3. The Bessel K penalty

The Bessel K pdf can be used with arbitrary values of $\epsilon \geq 0$ controlling its sparsity-inducing property. To illustrate this, Fig. 1(a) depicts one contour line of the restriction¹¹ to \mathbb{R}^2 of $q_I(w_1, w_2; \epsilon, \eta)$ in (12) for selected values of ϵ . As ϵ approaches zero, more probability mass concentrates along the \mathbf{w} -axes. As a consequence, the mode of the resulting posterior pdf $p(\mathbf{w}|\mathbf{y}; \epsilon, \eta)$ is more likely to be close to the axes, thus encouraging a sparse estimate. The behavior of the ℓ_1 -norm penalty that results from the selection $\epsilon = \rho + 1/2 = 3/2$ is also clearly recognized.

2.3. Type II cost function

We invoke Theorem 2 in [21] to obtain the Type II cost function induced by the Bessel K model (see (7) and (8)):

$$\mathcal{L}_{II}(\mathbf{w}) \triangleq \rho \|\mathbf{y} - \Phi \mathbf{w}\|_2^2 + \lambda^{-1} q_{II}(\mathbf{w}) \quad (15)$$

with

$$q_{II}(\mathbf{w}; \epsilon, \eta) = \min_{\gamma} \left\{ \rho \mathbf{w}^H \mathbf{\Gamma}^{-1} \mathbf{w} + \rho \log |\mathbf{C}| + (1 - \epsilon) \sum_i \log \gamma_i + \eta \sum_i \gamma_i \right\}. \quad (16)$$

In contrast to $q_I(\mathbf{w})$, $q_{II}(\mathbf{w})$ is non-separable. This makes an interpretation of $q_{II}(\mathbf{w})$ as done for $q_I(\mathbf{w})$ in Fig. 1(a) rather difficult. However, this interpretation becomes straightforward if Φ is orthonormal, i.e., $\Phi^H \Phi = \mathbf{I}$. In this

case $q_{II}(\mathbf{w})$ is separable, i.e., $q_{II}(\mathbf{w}) = \sum_i q_{II}(w_i)$ with

$$q_{II}(w_i; \epsilon, \eta) = \min_{\gamma_i} \left\{ \rho \frac{|w_i|^2}{\gamma_i} + \rho \log(\lambda^{-1} + \gamma_i) + (1 - \epsilon) \log \gamma_i + \eta \gamma_i \right\}. \quad (17)$$

Fig. 1(b) shows the contours of the restriction to \mathbb{R}^2 of $q_{II}(w_1, w_2; \epsilon, \eta)$ in (17) for different values of ϵ . Again, we observe the same increased concentration of mass around the \mathbf{w} -axes for decreasing values of ϵ . Interestingly, $q_{II}(w_1, w_2; \epsilon = 3/2, \eta)$ is no longer sparsity-inducing as compared to $q_I(w_1, w_2; \epsilon = 3/2, \eta)$. Thus, a sparsity-inducing prior model for Type I estimation is not necessarily sparsity-inducing for Type II estimation. We further investigate this important result in Section 2.4.

Another important property of the Type II penalty is its dependency on the noise variance λ^{-1} . Figs. 2(a) and (b) depict a single contour line of (17) for two values of ϵ and two values of λ^{-1} . Notice that $q_{II}(\mathbf{w}; \epsilon = 1/2, \eta = 1)$ resembles the log-sum penalty even in noisy conditions. For comparison purposes, we show in Fig. 2(c) the Type II penalty computed with the prior model in RVM [14,15] which utilizes a constant prior pdf $p(\gamma_i) \propto 1$ (corresponding to setting $\epsilon = 1$ and $\eta = 0$ in (16)). When $\lambda^{-1} = 0$ the RVM penalty equals the log-sum penalty. However, in noisy conditions, the RVM penalty resembles the ℓ_1 -norm penalty. Note that we cannot simply set λ^{-1} to some small value in order to obtain a strong sparsity-inducing penalty in RVM as λ^{-1} acts as a regularization of $q_{II}(\mathbf{w})$ in (15). Based on this observation, we expect that the Type II estimator derived from the Bessel K model achieves improved sparsity performance as compared to RVM in noisy scenarios. The numerical results conducted in Section 4 demonstrate that this is indeed the case.

2.4. Type I and Type II estimation

Having evaluated the impact of ϵ on $q_I(\mathbf{w})$ and $q_{II}(\mathbf{w})$, we now investigate its effect of this parameter on the corresponding Type I and Type II estimators. We demonstrated that as ϵ decreases, $q_I(\mathbf{w})$ and $q_{II}(\mathbf{w})$ become more and more sparsity-inducing which motivates the selection of a small ϵ for sparse estimation. On the other hand the

¹¹ Let f denote a function defined on a set A . The restriction of f to a subset $B \subset A$ is the function defined on B that coincides with f on this subset.

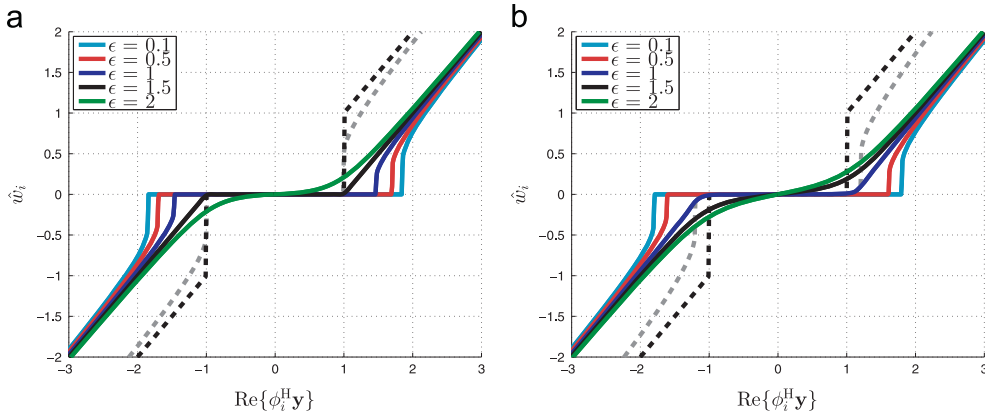


Fig. 3. Restriction to $\text{Im}\{\phi_i^H \mathbf{y}\} = 0$ of the EM-based Type I and Type II estimators of the complex weight w_i when Φ is orthonormal. The gray dashed lines depict the estimator corresponding to the setting $\epsilon = \eta = 0$, i.e., when $p(\gamma_i)$ equals the Jeffreys prior. The black dashed lines represent the hard-threshold rule. All results were generated using $\lambda^{-1} = 1/4$ and η set such that $\lambda^{-1}\sqrt{\eta/\rho} = 1$. (a) Type I (b) Type II.

Bessel K model for Type I and Type II estimation dominates the information contained in the observation \mathbf{y} for decreasing values of ϵ . Specifically, in case of Type I, when $\epsilon \leq \rho$ then $\lim_{w_i \rightarrow 0} q_i(w_i) = -\infty$, hence, the Type I estimator does not exist as $\mathcal{L}_I(\mathbf{w})$ has singularities. Likewise, this is the case for the Type II estimator when $\epsilon < 1$. The unbounded behavior of these penalties naturally questions the practicability of the Bessel K model in SBL. At least one would expect that we should refrain from selecting $\epsilon \leq \rho$ in case of Type I estimation and $\epsilon < 1$ for Type II estimation. Note, however, that utilizing unbounded penalties in SBL is not uncommon. Examples include [30,31] as well as the popular GSM model realizing the log-sum penalty in, e.g., [23]. Furthermore, the sparsity-inducing behavior of the penalty curves in Figs. 1 and 2 provides a strong motivation for using the Bessel K model in SBL. The approach is to formulate approximate inference algorithms, such as EM, for Type I and Type II estimation that overcome the difficulty of the singularities in the objective functions.

2.4.1. Approximate Type I estimation

The EM algorithm approximating the Type I estimator makes use of the complete data $\{\gamma, \mathbf{y}\}$ for \mathbf{w} .¹² The M-step computes an estimate of \mathbf{w} as the maximizer of

$$\langle \log(p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)p(\gamma)) \rangle_{p(\gamma; \hat{\mathbf{w}})}, \quad (18)$$

where $p(\gamma; \hat{\mathbf{w}})$ is computed in the E-step. Notice that as $p(\mathbf{w}|\mathbf{y}, \gamma) \propto p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)$ is proportional to a Gaussian pdf for \mathbf{w} , (18) does not have any singularity in contrast to $\mathcal{L}_I(\mathbf{w})$.

In order to get further insight into the impact of ϵ on the EM algorithm, we follow [23] and let Φ be orthonormal such that the EM update of the estimate of \mathbf{w} decouples into N independent scalar optimization problems. Fig. 3(a) visualizes the EM estimator for different values of ϵ . Clearly, the EM estimator approximates the soft-thresholding rule for large values of $\text{Re}\{\phi_i^H \mathbf{y}\}$ and as ϵ decreases the threshold value increases, thus, encouraging sparsity.

When the Bessel K pdf equals the Laplace pdf (i.e., $\epsilon = \rho + 1/2$), $\hat{\mathbf{w}}_I$ coincides with the soft-thresholding rule, which can be computed in closed form:

$$\hat{w}_{I,i}(\mathbf{y}) = \text{sign}(\phi_i^H \mathbf{y}) \max \left\{ 0, \left| \phi_i^H \mathbf{y} \right| - \lambda^{-1} \sqrt{\frac{\eta}{\rho}} \right\}, \quad i = 1, \dots, N. \quad (19)$$

Here, $\text{sign}(x) = x/|x|$ is the sign function. Notice that the EM estimator with $\epsilon = \rho + 1/2$ approximates (19) as depicted in Fig. 3(a).

2.4.2. Approximate Type II estimation

The EM algorithm approximating Type II estimation is devised using $\{\mathbf{w}, \mathbf{y}\}$ as the complete data for γ .¹³ The M-step computes an estimate of γ as the maximizer of

$$\langle \log(p(\mathbf{y}|\mathbf{w})p(\mathbf{w}|\gamma)p(\gamma)) \rangle_{p(\mathbf{w}; \hat{\gamma})}, \quad (20)$$

with $p(\mathbf{w}|\hat{\gamma})$ computed in the E-step. As $p(\gamma|\mathbf{w}) \propto p(\mathbf{w}|\gamma)p(\gamma)$ is a Generalized Inverse Gaussian (GIG) pdf for γ , (20) does not exhibit any singularity as opposed to $\mathcal{L}_{II}(\gamma)$.

In Fig. 3(b), we show the EM estimate of w_i for different settings of ϵ . Similar to Type I, the Type II estimate approaches the soft-thresholding rule as $\text{Re}\{\phi_i^H \mathbf{y}\}$ becomes larger and as ϵ decreases a sparser estimate is obtained. However, when $\epsilon = 3/2$, i.e., utilizing the Laplace GSM model for the complex weights, the Type I estimator coincides with the soft-threshold rule while the Type II estimator does not have this threshold-like behavior and is not sparse. This was already indicated by the behavior of $q_{II}(\mathbf{w}; \epsilon = 3/2, \eta)$ in Fig. 1(b).

From Fig. 3 we conclude that the EM-based Type I estimator is a sparse estimator for $\epsilon \leq \rho + 1/2$, whereas the EM-based Type II estimator only exhibits this property for $\epsilon \leq 1$. In Fig. 4, we illustrate this important difference in the behavior of these estimators for real and complex signal representation when utilizing the GSM model of the Laplace prior: the EM-based Type I estimator achieves a sparse solution for both real and complex weights, whereas for the EM-based Type II estimator this is only the case for real weights.

¹² This EM algorithm is derived in Appendix A.

¹³ This EM algorithm is derived in Section 3.1.




Fig. 4. EM-based Type I and Type II estimates using the Laplace GSM model for (a)–(b) real and (c)–(d) complex weights. For these simulations, $\Phi \in \mathbb{C}^{50 \times 128}$ with its entries drawn independently according to $\phi_{mi} \sim \text{CN}(0, 1/M)$. The $K=12$ non-zero entries in \mathbf{w} are of the form $w_k = \exp(j\theta_k)$ with $\theta_k, k = 1, \dots, K$, drawn independently according to a uniform distribution on $[0, 2\pi)$. The SNR is fixed at 60 dB. (a) $\rho = 1/2, \epsilon = 1$ (b) $\rho = 1/2, \epsilon = 1$ (c) $\rho = 1, \epsilon = 3/2$ (d) $\rho = 1, \epsilon = 3/2$.

3. Sparse Bayesian inference

In this section we derive a Bayesian inference scheme that relies on the Bessel K model presented in Section 2. First, we obtain an EM algorithm that approximates the Type II estimator of the weight vector \mathbf{w} in (6). Inspired by [1] and [36] we then derive a fast algorithm based on a modification of the EM algorithm. We show that this algorithm actually encompasses the fast algorithms in [1] and [2] as special instances.

Naturally, the approach presented here can also be applied to derive algorithms approximating the Type I estimator. However, numerical investigations not reported here indicate that these algorithms often fail to produce sparse estimates of \mathbf{w} when small values of the parameter ϵ are selected. Hence, we restrict the discussion in this section to algorithms approximating the Type II estimator.

3.1. Sparse Bayesian inference using EM

We adapt the EM algorithm approximating the Type II estimator previously used for SBL [14,15,37,2] to the Bessel K model. As the value of λ is in general unknown and has a significant impact on the sparsity-inducing property on $q_H(\mathbf{w})$ (see Section 2), we include the estimation of this parameter in the inference framework. We seek the MAP estimate of $\{\gamma, \lambda\}$, i.e., the maximizer of

$$\mathcal{L}(\gamma, \lambda) = \log p(\mathbf{y}, \gamma, \lambda) = \log(p(\mathbf{y}|\gamma, \lambda)p(\gamma)p(\lambda)). \quad (21)$$

We use the EM algorithm to approximate the MAP estimator. We specify $\{\mathbf{w}, \mathbf{y}\}$ to be the complete data for $\{\gamma, \lambda\}$. With this choice the E-step of the EM algorithm computes the conditional expectation

$$\langle \log p(\mathbf{y}, \mathbf{w}, \gamma, \lambda) \rangle_{p(\mathbf{w}|\mathbf{y}, \gamma^{[t]}, \lambda^{[t]})} \quad (22)$$

with $p(\mathbf{w}|\mathbf{y}, \gamma^{[t]}, \lambda^{[t]}) = \text{N}(\mathbf{w}; \boldsymbol{\mu}^{[t]}, \boldsymbol{\Sigma}^{[t]})$ or $p(\mathbf{w}|\mathbf{y}, \gamma^{[t]}, \lambda^{[t]}) = \text{CN}(\mathbf{w}; \boldsymbol{\mu}^{[t]}, \boldsymbol{\Sigma}^{[t]})$ depending on whether the underlying signal model is real or complex. Here, $(\cdot)^{[t]}$ denotes the estimate of the parameter given as an argument at iteration t . In either case, the parameters of the conditional pdf of \mathbf{w} read

$$\boldsymbol{\Sigma}^{[t]} = \left(\lambda^{[t]} \Phi^H \Phi + (\Gamma^{[t]})^{-1} \right)^{-1}, \quad (23)$$

$$\boldsymbol{\mu}^{[t]} = \lambda^{[t]} \boldsymbol{\Sigma}^{[t]} \Phi^H \mathbf{y}. \quad (24)$$

The M-step of the EM algorithm updates the estimate of $\{\gamma, \lambda\}$ as the maximizer of (22):

$$\gamma_i^{[t+1]} = \frac{\epsilon - \rho - 1 + \sqrt{(\epsilon - \rho - 1)^2 + 4\rho\eta\langle |w_i|^2 \rangle^{[t]}}}{2\eta}, \quad i = 1, \dots, N, \quad (25)$$

$$\lambda^{[t+1]} = \frac{M}{\|\mathbf{y} - \Phi \boldsymbol{\mu}^{[t]}\|_2^2 + \text{tr}(\Phi^H \Phi \boldsymbol{\Sigma}^{[t]})}. \quad (26)$$

Here, $\langle |w_i|^2 \rangle^{[t]}$ is the i th diagonal element of $\boldsymbol{\Sigma}^{[t]} + \boldsymbol{\mu}^{[t]}(\boldsymbol{\mu}^{[t]})^H$ computed in the E-step and $\text{tr}(\cdot)$ is the trace operator.

3.2. Modified update of $\gamma_i^{[t+1]}$

One of the major drawbacks of the EM algorithm approximating the Type II estimator is its slow convergence, as observed in [1].¹⁴ In this section, we discuss a modification of the EM algorithm that improves the convergence speed. The proposed algorithm is inspired by [1] and [36]. To this end, we focus on the update of a single estimate of γ_i and express this update as a (non-linear recurrent) function of the previous update. Then, we analyze the fixed points of this function for different settings of the hyperparameters ϵ and η and formulate a new update rule for the estimate of γ_i at iteration $t+1$ based on these fixed points. From this point on, we restrict our analysis to the Bessel K model with $\epsilon \leq 1$ since, as discussed in Section 2, the setting $\epsilon > 1$ does not yield a sparse Type II estimator.

To begin, we consider the update in (25) for a single parameter γ_i while considering the estimates $\gamma_k^{[t]}, k \neq i$, and $\lambda^{[t]}$ as fixed quantities. In Appendix B.1, we show that the dependency of $\langle |w_i|^2 \rangle^{[t]}$ on $\gamma_i^{[t]}$ is expressed as

$$\langle |w_i|^2 \rangle^{[t]} = \frac{(\gamma_i^{[t]})^2 (s_i^{[t]} + |q_i^{[t]}|^2) + \gamma_i^{[t]} (s_i^{[t]})^2}{(\gamma_i^{[t]} + s_i^{[t]})^2} \quad (27)$$

with $s_i^{[t]} \triangleq \mathbf{e}_i^T \boldsymbol{\Sigma}^{[t]} \mathbf{e}_i$, $q_i^{[t]} \triangleq \lambda^{[t]} \mathbf{e}_i^T \boldsymbol{\Sigma}^{[t]} \Phi^H \mathbf{y}$, $\boldsymbol{\Sigma}_{-i}^{[t]} \triangleq (\lambda^{[t]} \Phi^H \Phi + \sum_{k \neq i} (\gamma_k^{[t]})^{-1} \mathbf{e}_k \mathbf{e}_k^T)^{-1}$ and \mathbf{e}_i denoting an $N \times 1$ vector of all zeros but 1 at the i th position. By inserting (27) into (25), we obtain an update expression of the form

$$\gamma_i^{\text{new}} = \phi_i^{[t]}(\gamma_i^{\text{old}}) \quad (28)$$

where the function $\phi_i^{[t]}$ is parametrized by $\epsilon, \eta, s_i^{[t]}$, and $q_i^{[t]}$. Next, we explore the hypothetical behavior of the estimates of γ_i that we would obtain by recursively applying $\phi_i^{[t]}$ ad infinitum. We do so by analyzing the existence of fixed points

¹⁴ The selected mixing pdf also has a significant impact on the convergence speed as shown in Section 4.

of the function $\varphi_i^{[t]}$. A fixed point $\tilde{\gamma}_i$ of $\varphi_i^{[t]}$ must fulfill

$$\tilde{\gamma}_i = \varphi_i^{[t]}(\tilde{\gamma}_i) = \frac{\epsilon - \rho - 1 + \sqrt{(\epsilon - \rho - 1)^2 + 4\rho\eta \frac{\tilde{\gamma}_i^2(s_i + |q_i|^2) + \tilde{\gamma}_i s_i^2}{(\tilde{\gamma}_i + s_i)^2}}}{2\eta} \quad (29)$$

where, for notational simplicity, we have dropped the iteration index for s_i and q_i . By inspection of (29), it is clear that $\tilde{\gamma}_i = 0$ is always a fixed point of $\varphi_i^{[t]}$ when $\epsilon \leq 1$. We look for other positive fixed points by solving (29). These fixed points are solutions of the fourth order equation

$$0 = \gamma_i(\eta\gamma_i^3 + \gamma_i^2[2\eta s_i - (\epsilon - \rho - 1)] + \gamma_i[\eta s_i^2 - 2(\epsilon - \rho - 1)s_i - \rho(s_i + |q_i|^2)] - (\epsilon - 1)s_i^2). \quad (30)$$

Hence, if any strictly positive fixed point $\tilde{\gamma}_i$ of $\varphi_i^{[t]}$ exists, it must be a solution of the cubic equation

$$0 = \eta\gamma_i^3 + \gamma_i^2[2\eta s_i - (\epsilon - \rho - 1)] + \gamma_i[\eta s_i^2 - 2(\epsilon - \rho - 1)s_i - \rho(s_i + |q_i|^2)] - (\epsilon - 1)s_i^2. \quad (31)$$

As we show in Appendix B.2, the positive solutions of (31) correspond, in fact, to the stationary points of (21) when all variables except γ_i are kept fixed at their current estimates, i.e., of

$$\ell_i^{[t]}(\gamma_i) \propto \log(p(\mathbf{y}|\gamma_i, \gamma_{-i}^{[t]}, \lambda^{[t]})p(\gamma_i)). \quad (32)$$

Based on the above analysis, we formulate a new update rule for γ_i at iteration $t+1$. Given the values of all estimates at iteration t , we calculate the fixed points of the corresponding function $\varphi_i^{[t]}$ by solving (30). Then

- if no strictly-positive fixed points of $\varphi_i^{[t]}$ exist, we set $\gamma_i^{[t+1]} = 0$, which, remember, is also a fixed point of $\varphi_i^{[t]}$;
- if strictly-positive fixed points of $\varphi_i^{[t]}$ exist, we select the fixed point $\tilde{\gamma}_i$ which yields the largest value $\ell_i^{[t]}(\tilde{\gamma}_i)$ among all strictly positive fixed points. We then set $\gamma_i^{[t+1]} = \tilde{\gamma}_i$.

Note that the above selection criterion for $\gamma_i^{[t+1]}$ is a heuristic choice. In fact, we have no guarantee that, by iteratively applying the recurrent function $\varphi_i^{[t]}$, convergence to the selected fixed point will occur. This depends on the initialization $\gamma_i^{[t]}$. Moreover, when $\epsilon < 1$, selecting a strictly-positive fixed point instead of zero does not guarantee that the objective function (21) is increased, as (32) diverges to infinity when γ_i tends to zero.¹⁵ With this selection, however, we hope to obtain an improved convergence speed at the expense of sacrificing the monotonicity property of the EM algorithm. The numerical results obtained with this heuristic choice, shown in Section 4, confirm the effectiveness of the approach.

Next we investigate the solutions of (30) for different selections of ϵ and η . We show that for some particular selections of these parameters, the modified update of $\gamma_i^{[t+1]}$ coincides with the updates in the algorithms presented in [1] and [2]. For brevity, we omit the algorithmic iteration index t throughout the rest of the section.

3.2.1. Fixed points for $0 \leq \epsilon < 1$ and $\eta \geq 0$

We consider an arbitrary value of ϵ in the range $0 \leq \epsilon < 1$. First, as $-(\epsilon - 1)s_i^2 \geq 0$ for $\epsilon < 1$, (31) has at least

one negative solution. If no positive solution exists we set $\hat{\gamma}_i = 0$. If (31) has at least one positive solution, it is easily shown that it has exactly two, denoted by $\gamma_i^{(1)}$ and $\gamma_i^{(2)}$. If $\gamma_i^{(1)} = \gamma_i^{(2)}$ then this point is a saddle point of ℓ_i and therefore we set $\hat{\gamma}_i = 0$. If $\gamma_i^{(2)} > \gamma_i^{(1)}$ then $\hat{\gamma}_i = \gamma_i^{(2)}$ or if $\gamma_i^{(1)} > \gamma_i^{(2)}$ then $\hat{\gamma}_i = \gamma_i^{(1)}$ (the proof is straightforward and is omitted). Thus, we always select the right-most positive solution.

For the special case $\epsilon = \eta = 0$, i.e., when the mixing density coincides with the Jeffreys prior, (31) reduces to a quadratic equation. It is easy to show that, in this case either two positive solutions exist or none exists.

3.2.2. Fixed points for $\epsilon = 1$ and $\eta = 0$

In this case the mixing density coincides with a constant improper prior, which leads to the same GSM model as used in RVM [14,15]. With this setting (31) simplifies to

$$\hat{\gamma}_i = |q_i|^2 - s_i. \quad (33)$$

From (33), a positive solution of (31) exists if and only if $|q_i|^2 > s_i$. If this condition is not satisfied we set $\hat{\gamma}_i = 0$. It is interesting to note that (33) is independent of ρ and thus is the same regardless of whether the signal model (1) is real or complex.

Next, we show the equivalence between (33) and the corresponding update in Fast RVM [1]. In [1], the estimate of γ_i is computed as the maximizer of the marginal log-likelihood function $\ell_i(\gamma_i, \epsilon = 1, \eta = 0)$ in (32). Hence, the estimate of γ_i in [1] equals that in (33), because (33) maximizes $\ell_i(\gamma_i, \epsilon = 1, \eta = 0)$. As the updates of μ , Σ , and $\hat{\lambda}$ are identical to those in Fast RVM, the two algorithms coincide when $\epsilon = 1$ and $\eta = 0$.

3.2.3. Fixed points for $\epsilon = 1$ and $\eta > 0$

In this case the mixing pdf coincides with an exponential pdf, so the GSM model is the same as that used in Fast Laplace [2]. The solution

$$\hat{\gamma}_i = \frac{-(2\eta s_i + \rho) + \sqrt{\rho^2 + 4\rho\eta|q_i|^2}}{2\eta} \quad (34)$$

is positive if and only if $|q_i|^2 - s_i > \eta s_i^2 / \rho$ otherwise we set $\hat{\gamma}_i = 0$. The case $\epsilon = 1$ and $\rho = 1/2$ corresponds to the GSM model of the Laplace prior for real weights. Obviously, (34) can also be used for complex weights, with $\rho = 1$. Yet in this case the marginal prior for \mathbf{w} is no longer Laplacian, as showed in Section 2, but some other sparsity-inducing member of the Bessel K density family. The estimate of γ_i in Fast Laplace [2] is the maximizer of $\ell_i(\gamma_i, \epsilon = 1, \eta)$ and, hence, is identical to the estimate in (34).

3.3. Fast sequential inference scheme

The modified update of $\gamma_i^{[t+1]}$, $i = 1, \dots, N$, described in Section 3.2 can be directly used to speed up the EM algorithm presented in Section 3.1. With this modification, every time an estimate of a given γ_i is set to zero, we remove the corresponding column vector ϕ_i from the dictionary matrix Φ . This effectively reduces the model complexity “on the fly.” However, the first iterations still suffer from a high computational complexity due to the update (23). To avoid this, we follow the approach outlined in [1, Section 4], which consists of

¹⁵ See the discussion in Section 2.4.

starting with an “empty” dictionary Φ and incrementally filling the dictionary by possibly adding one column vector at each iteration of the algorithm. Specifically, at a given iteration of the algorithm, each $\hat{\gamma}_i$, $i = 1, \dots, N$, is computed from (30) and the one, say $\hat{\gamma}_i$, that gives rise to the greatest increase in $\exp(\ell(\cdot))$ between two consecutive algorithmic iterations, is selected. Depending on the value of this $\hat{\gamma}_i$, the corresponding vector ϕ_i is then added, deleted, or kept. The quantities Σ , μ , and $\hat{\lambda}$ are updated using (23), (24), and (26) together with s_i and q_i , $i = 1, \dots, N$. If the estimate of λ is not updated between two consecutive iterations, Σ , μ , s_i , and q_i can be updated efficiently using the update procedures proposed in [1,36].

We refer to the above sequential algorithm as Fast-BesselK.

4. Numerical results

In this section we analyze the performance of the Fast-BesselK algorithm proposed in Section 3.

The purpose is to characterize the impact of the prior model on the performance of the iterative algorithm in terms of MSE, sparseness of $\hat{\mathbf{w}}$, and convergence speed. Section 3 shows that Fast-RVM [1], Fast-Laplace [2], and Fast-BesselK are all instances of the same greedy inference scheme each algorithm resulting from a particular selection of the parameters of the mixing (gamma) pdf. Hence, by comparing the performances of these algorithms, we can draw conclusions on the sparsity-inducing property of their respective prior models.¹⁶

4.1. Simulation scenarios and performance metrics

The performance of the considered sparse algorithms (see Section 4.2) is evaluated by means of Monte Carlo simulations. In order to test the algorithms on a realistic benchmark, we use a random $M \times N$ dictionary matrix Φ , with $M = 100$ and $N = 256$, whose entries are iid zero-mean complex symmetric Gaussian random variables with variance M^{-1} . The weight vector \mathbf{w} has K non-zero entries with associated indices uniformly drawn without repetition from the set $\{1, 2, \dots, N\}$. The set of these indices together with its cardinality K are unknown to the algorithms. The non-zero entries in \mathbf{w} are independent and drawn from a zero-mean circular-symmetric complex Gaussian distribution with unit variance. Other distributions for the entries in \mathbf{w} are considered at the end of this section. All reported performance curves are computed based on a total of 1000 Monte Carlo trials. For each trial, new realizations of the dictionary matrix Φ , the vector \mathbf{w} , and the random perturbation vector \mathbf{n} are drawn independently.¹⁷

All numerical investigations were replicated for an equivalent real-valued signal model. Due to space limitations, we do not include the results of these studies in this

contribution, as most of the conclusions are similar to those drawn from the complex-valued signal model. We will, however, shortly discuss the relation between the performance of the estimators for real and complex models at the end of this section.

The performance is evaluated with respect to the following metrics:

normalized mean-squared error: $\text{NMSE} \triangleq \langle \|\hat{\mathbf{w}} - \mathbf{w}\|_2^2 \rangle / \langle \|\mathbf{w}\|_2^2 \rangle$.

support error rate $\triangleq \#\{i: \hat{w}_i = 0 \text{ and } w_i \neq 0\} \cup \{i: \hat{w}_i \neq 0 \text{ and } w_i = 0\} / N$

We also report the convergence speed, measured in terms of the number of algorithmic iterations used, of the Bayesian inference methods as they share the same computational complexity.

4.2. Inference algorithms considered

The proposed Fast-BesselK algorithm is tested with two settings for ϵ and η :

- Fast-BesselK ($\epsilon = 0$): we set $\epsilon = 0$ and $\eta = 0$ corresponding to the use of the Jeffreys prior as mixing density.¹⁸
- Fast-BesselK ($\epsilon = 0.5$): we set $\epsilon = 0.5$ and $\eta = 1$.

Instead of selecting a particular value of η , we could have included this parameter in the inference framework as done in [2]. Our investigations, however, show that for $\epsilon \ll 1$ the performance of Fast-BesselK becomes largely independent of the choice of η , and we have therefore simply selected $\eta = 1$.¹⁹

The performance of Fast-BesselK is contrasted with the state-of-the-art sparse estimators listed below:

1. Fast-RVM [1,37]: Equivalent to Fast-BesselK with $\epsilon = 1$ and $\eta = 0$ (see Section 3).²⁰
2. Fast-Laplace [2]: Equivalent to Fast-BesselK with $\epsilon = 1$ when including the update for η in [2] (see Section 3).²¹
3. OMP, see, e.g., [12]: OMP terminates when the greedy algorithm has included $K + 10$ column vectors in Φ . We empirically observed that this choice induces a better NMSE performance than when including K columns only.
4. SpaRSA [34]: The sparse reconstruction by separable approximation (SpaRSA) algorithm for solving the LASSO cost function. Following [34], we use the adaptive continuation procedure for the regularization κ of the ℓ_1 -norm penalty in the LASSO cost function. Here SpaRSA repeatedly solves the LASSO cost function with decreasing values for κ until a minimum value of κ is reached. The minimum value of κ is set through training. Specifically, we run 50 Monte Carlo trials for each specific settings of M , N , K , and SNR value. We then

¹⁸ We also considered Fast-BesselK with $\epsilon = 0$ and $\eta = 1$. However, this setting led to similar performance to Fast-BesselK ($\epsilon = 0, \eta = 0$).

¹⁹ If the Fast-BesselK is implemented with a “top-down” approach (starting out with the full dictionary Φ) including individual rate parameters η_i for each w_i , $i = 1, \dots, N$, may be beneficial [7].

²⁰ The software is available on-line at <http://people.ee.duke.edu/lcarin/BCS.html>.

²¹ The software is available on-line at <http://ivpl.eecs.northwestern.edu/>.

¹⁶ Naturally, the practical implementation of the inference schemes also impacts the performance. For the subsequent analysis, Fast-RVM, Fast-Laplace, and Fast-BesselK are all implemented based on the Matlab-code for Fast-RVM located at <http://people.ee.duke.edu/lcarin/BCS.html>.

¹⁷ In this paper we have not included an investigation on a specific application. We refer to the work [8] where such a performance assessment is made.

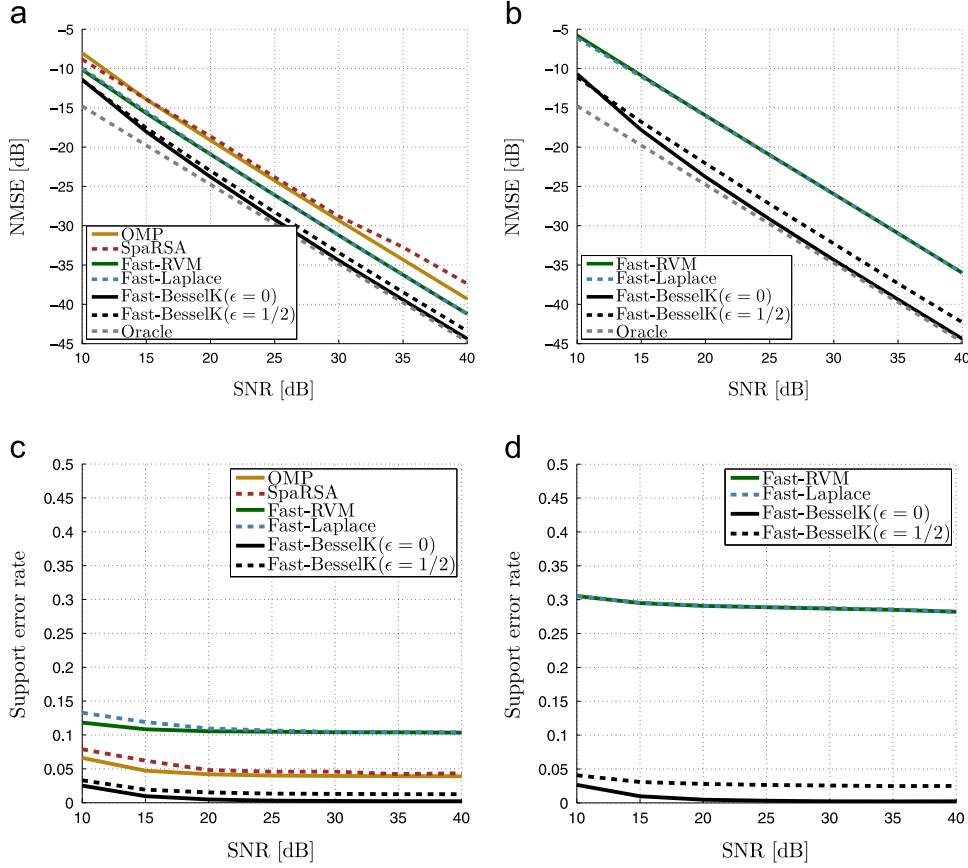


Fig. 5. Performance versus SNR when λ is known ((a), (c)) and λ is unknown and estimated ((b), (d)). Selected system parameter settings: $M=100$, $N=256$, and $K=25$.

choose the value of κ from a set of 50 candidate values in the range $[0.001\|\Phi^H y\|_\infty, 0.1\|\Phi^H y\|_\infty]$ that leads to the smallest error $\|\mathbf{w} - \hat{\mathbf{w}}\|_2^2$.²²

For brevity, we refer to Fast-RVM, Fast-Laplace, and Fast-BesselK as Bayesian algorithms. We initialize these algorithms as outlined in [1, Sec. 4]. They stop when either the stopping criterion $\|\hat{\mu}^{[t+1]} - \hat{\mu}^{[t]}\|_\infty \leq 10^{-8}$ is fulfilled or the number of iterations has reached its max limit set to 1000.

As a reference, we also consider the performance of the oracle estimator of \mathbf{w} [38] that “knows” the support of \mathbf{w} , denoted by $\text{supp}(\mathbf{w}) \triangleq \{i: w_i \neq 0\}$. The oracle estimate reads

$$\hat{\mathbf{w}}_o(\mathbf{y}) = \begin{cases} (\Phi_o^H \Phi_o)^{-1} \Phi_o^H \mathbf{y}, & \text{on } \text{supp}(\mathbf{w}) \\ 0, & \text{elsewhere,} \end{cases} \quad (35)$$

where Φ_o is the $M \times K$ dictionary matrix constructed from those columns of Φ that correspond to the non-zero entries in \mathbf{w} .

4.3. Performance comparison

As our analysis in Section 2 shows, the noise precision λ greatly impacts the sparsity property of the Type II penalty. We

therefore investigate the impact of this parameter on the algorithms. First, we assume this quantity to be known to the Bayesian algorithms. Note that SpaRSA and OMP do not estimate λ . In the next step, this parameter is considered unknown and estimated by the Bayesian algorithms.

4.3.1. Performance versus SNR

The goal of this investigation is to evaluate whether the algorithms can achieve sparse and accurate estimates in conditions of low and medium SNR. In these simulations, we set $K=25$. In Figs. 5(a) and (c), λ is known by the Bayesian algorithms. Fig. 5(a) shows that Fast-BesselK($\epsilon=0$) and Fast-BesselK($\epsilon=0.5$) achieve the lowest NMSE among the algorithms across the whole SNR range. Their performance is close to that of the oracle estimator in the high SNR regime, i.e., above 20 dB. These algorithms also achieve the lowest support error rate across the whole SNR range with a value close to zero as shown in Fig. 5(c).

We repeat the investigation for the Bayesian algorithms but this time with the noise precision λ unknown and being estimated alongside \mathbf{w} and γ using (26). The estimate $\hat{\lambda}$ is updated at every third iteration. We observe a significant performance degradation in NMSE and support error rate for Fast-RVM and Fast-Laplace in Figs. 5(b) and (d). The reason is that Fast-RVM and Fast-Laplace heavily overestimate λ , thus, K is overestimated as well (results not

²² The software is available on-line at <http://www.lx.it.pt/mtf/SpaRSA/>.

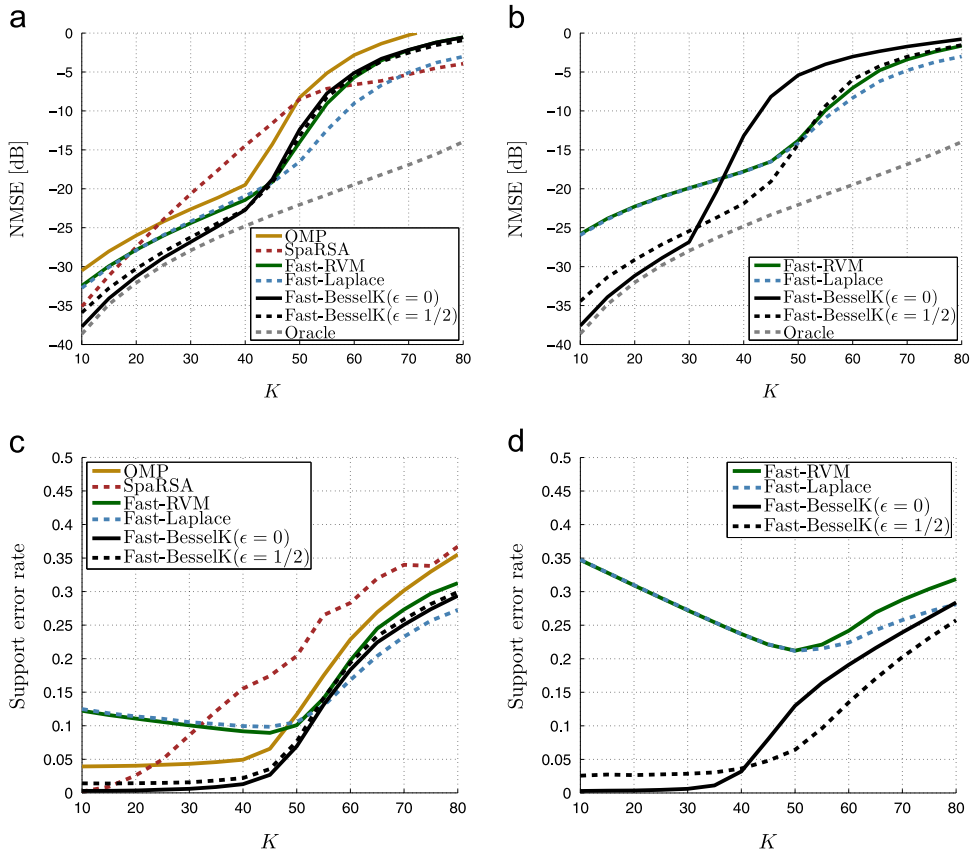


Fig. 6. Performance versus K at 20 dB SNR when λ is known ((a), (c)) and λ is unknown and estimated ((b), (d)). Selected system parameter settings: $M=100$, $N=256$, and $K=25$.

shown).²³ Consequently, the support error rate and NMSE is high. In contrast, the Fast-BesselK algorithms perform essentially the same as when λ is known.

In summary, the results presented in Fig. 5 corroborate the significant impact of the estimation of the noise precision on the performance of the Fast Bayesian algorithms. When λ is known, all algorithms achieve an acceptable performance, both in terms of NMSE and support error rate. However, when λ is unknown and estimated by the algorithms, only Fast-BesselK is able to produce accurate estimates of this parameter, resulting in greatly improved performance as compared to Fast-Laplace and Fast-RVM. This is an important result as, in many applications, the noise precision parameter is not known in advance and, hence, needs to be estimated.

4.3.2. Performance versus K

We fix the SNR at 20 dB and compare the performance of the algorithms versus the number K of non-zero entries in \mathbf{w} . In Figs. 6(a) and (c) we assume λ to be known to the Bayesian algorithms. The NMSE curves in Fig. 6(a) show that when $K \leq 40$ the algorithms achieve an accurate reconstruction of \mathbf{w} . Fast-BesselK($\epsilon=0$) and Fast-BesselK($\epsilon=0.5$) yield the

lowest NMSE which turns out to be close to that of the oracle estimator. In this range, these algorithms exhibit a support error rate close to zero as depicted in Fig. 6(c).

When λ is estimated the NMSE and support error rate performance achieved by Fast-RVM and Fast-Laplace degrade as depicted in Figs. 6(c) and (d). Fast-BesselK($\epsilon=0$) achieves the lowest NMSE but only for $K \leq 30$, as it only accurately estimates λ in this range. Consequently, its support error rate decreases for $K > 30$. In turn, Fast-BesselK($\epsilon=0.5$) achieves similar performance to when λ is known. Hence, the selection of $\epsilon=0.5$ seems to be a good trade-off between achieved sparseness and reconstruction error.

4.3.3. Number of performed algorithmic iterations

We evaluate the convergence speed for the Bayesian algorithms in terms of the number of performed algorithmic iterations. Fig. 7 reports the number of algorithmic iterations until either the stopping criterion is fulfilled or the number of iterations has reached its max limit of 1000 (see Section 4.2) versus SNR and K . The Fast-BesselK algorithms perform significantly less number of iterations across the whole SNR range as compared to Fast-RVM and Fast-Laplace, especially in low to medium SNR as seen in Figs. 7(a) and (b). The same superior performance is observed when K is varied in Figs. 7(c) and (d). Notice that the iteration count of greedy algorithms inherently depends on K . As Fast-RVM and Fast-Laplace tend to heavily overestimate K , they inevitably require

²³ In some cases, the sequence of estimates of λ produced by Fast-RVM and Fast-Laplace did not converge. Due to this, a maximum value of 10^8 was set for $\hat{\lambda}$.

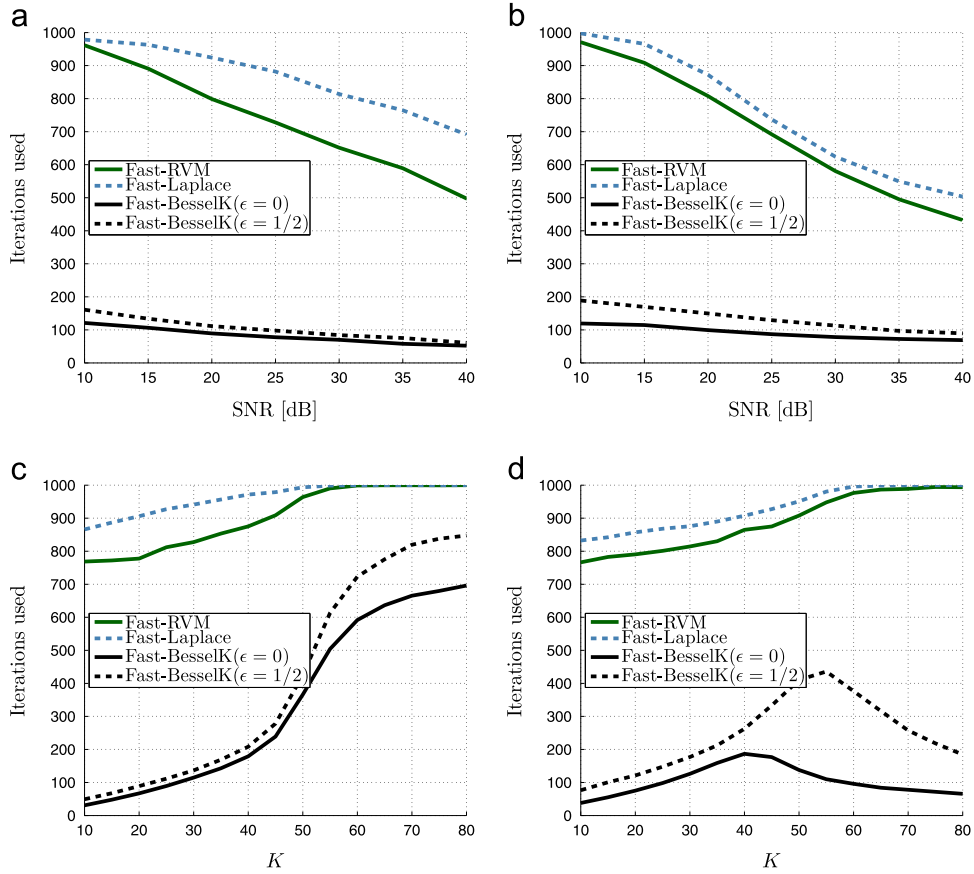


Fig. 7. Number of used iterations versus SNR and K when λ is known ((a), (c)) and λ is unknown and estimated ((b), (d)). Selected system parameter settings: $M=100$, $N=256$. In ((a), (b)) $K=25$ and in ((c), (d)) the SNR is fixed at 20 dB.

a larger number of iterations than algorithms achieving sparser estimates. The Fast-BesselK algorithms exhibit a modest increase of used iterations when $K \leq 40$ as they achieve good reconstruction error in this range, see Fig. 6. When $K \geq 40$, the different performance behavior for Fast-BesselK in Figs. 7(c) and (d) is attributed to the fact that Fast-BesselK significantly underestimates λ in this range. In this case, the penalty $q_H(\mathbf{w})$ has a high impact on the estimate $\hat{\mathbf{w}}$, which leads to a very sparse estimate $\hat{\mathbf{w}}$ and, thus, a low number of algorithmic iterations.

4.3.4. Performance versus different distributions of the non-zero entries in \mathbf{w}

We investigate the dependency of the performance of the considered algorithms on the underlying prior distribution of the non-zero entries in \mathbf{w} . To this end we repeat the previous numerical studies while considering two additional prior distributions for these entries. The first distribution results from selecting the non-zero entries to be of the form $\exp(j\phi_k)$, $k=1, \dots, K$ with the phases $\{\phi_k\}$ drawn independently and uniformly on the interval $[0, 2\pi)$. The second distribution results from drawing the non-zero entries independently according to a complex Laplace distribution, see (14), with unit variance. We show results only for Fast-RVM, Fast-Laplace, and Fast-BesselK($\epsilon=0.5$), as the performance

gain achieved by Fast-BesselK($\epsilon=0.5$) as compared to OMP and SpARSA is similar to the performance observed in the previous investigations. We conclude from Figs. 8 and 9 that Fast-BesselK($\epsilon=0.5$) still maintains its superior performance. Furthermore, we again observe the important fact that Fast-BesselK($\epsilon=0.5$) achieves similar performance in scenarios with known or unknown noise precision. This is in direct contrast to the other Bayesian methods.

4.3.5. Performance for real signal models

We conclude this section by briefly commenting on the performance achieved by the considered algorithms when they are devised for and applied to real-valued signal models. To distinguish between the algorithms devised based on real signal model from those devised for a complex signal model, in the subsequent discussion we refer to the former (latter) as real (complex) algorithms.

In general, all considered complex algorithms perform better than their real variant. In particular, complex algorithms produce accurate results for less sparse weight vectors than their real counterpart. This is explained by the fact that the former use both real and imaginary parts to prune components in $\hat{\mathbf{w}}$, thus, improving the sparse signal representation.

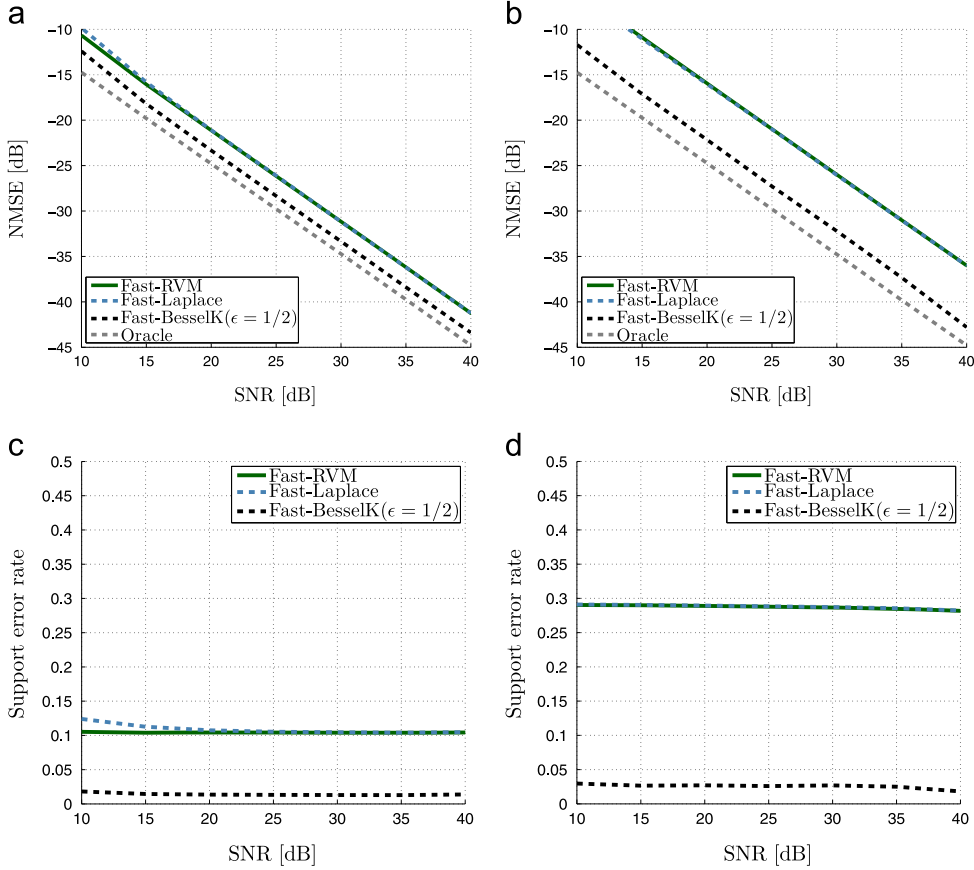


Fig. 8. Performance versus SNR when λ is known ((a), (c)) and λ is unknown and estimated ((b), (d)). The non-zero entries in \mathbf{w} are complex uniformly distributed. Selected system parameter settings: $M=100$, $N=256$, and $K=25$.

The relative performances of the real algorithms compared to each other show the same trends as that observed for their complex variant. As an illustration, real Fast-BesselK($\epsilon=0$) is especially sensitive to high values of K . This is a well-known effect that arises when using the Jeffreys prior as the mixing density. This again emphasizes our conclusion that Fast-BesselK($\epsilon=0.5$) is a good trade-off between sparseness and reconstruction error.

5. Conclusion

In this paper, we proposed a hierarchical prior model for sparse Bayesian learning (SBL) that applies to sparse signal representation in complex and real-valued signal models. Our motivation was on the one hand to overcome the lack of sparsity-inducing prior models for complex signals and on the other hand to propose prior models that induce sparse, accurate, and robust signal representations in conditions of low and medium signal-to-noise ratio (SNR). Both aspects are of particular importance in many engineering applications of sparse signal representation, e.g., in wireless communications.

In the proposed hierarchical prior model, the entries of the parameter vector of interest are modeled as independent complex Gaussian scale mixtures (GSMs) with mixing hyperparameters identically distributed according to a gamma

distribution with shape parameter ϵ and rate parameter η . This model – we termed it the Bessel K model – comprises a family of hierarchical prior probability density functions (pdfs) indexed by these parameters.

We analyzed the properties of Type I and Type II estimators derived from the Bessel K model. Our analysis revealed that the ability of a given element in the density family to induce sparse estimates heavily depends on the inference method used and, interestingly, whether real or complex signals are inferred. In the case of Type I estimation, the Bessel K model invokes, with the right setting of parameters ϵ and η , classical penalties such as the ℓ_1 -norm or the log-sum as special cases. The hierarchical Bayesian formulation of the ℓ_1 -norm penalty in the complex case is especially interesting as, to the authors' knowledge, it has not been proposed before. In the case of Type II estimation, the resulting penalties are also strongly influenced by the variance of the measurement noise, as pointed out by [21]. Nonetheless, we showed that the Bessel K model with $\epsilon < 1$ promotes sparse Type II estimators even when the noise variance is high. In contrast, traditional prior models lose this property in such conditions.

Finally, we derived a greedy algorithm of low complexity based on the modification of the expectation-maximization algorithm formulated for Type II estimation. As the Bessel K model encompasses as special cases previously proposed prior

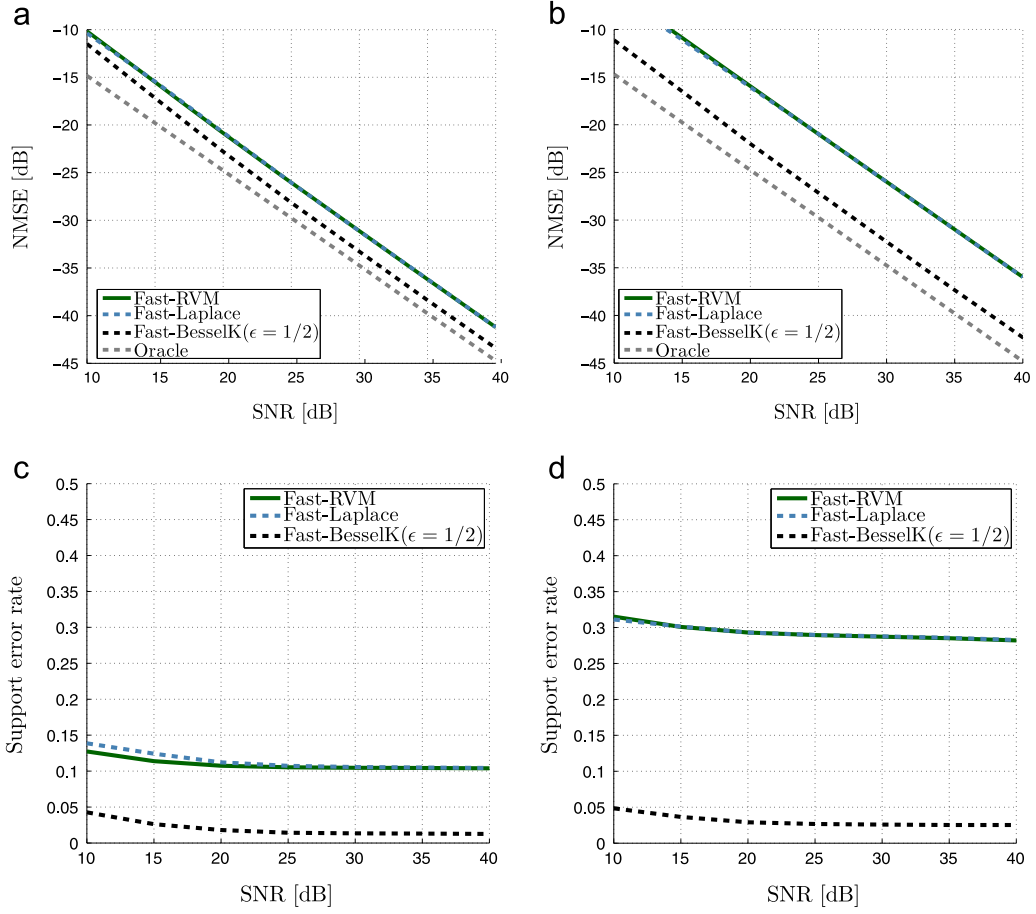


Fig. 9. Performance versus SNR when λ is known ((a), (c)) and λ is unknown and estimated ((b), (d)). The non-zero entries in \mathbf{w} are complex Laplace distributed. Selected system parameter settings: $M=100$, $N=256$, and $K=25$.

models, the algorithm generalizes existing fast SBL methods, allowing us to directly compare the impact of the different prior models on the performance of the resulting estimators.

The numerical results demonstrated that the Bessel K model with $\epsilon < 1$ leads to estimators with superior convergence speed, sparseness, and lower mean-squared estimation error as compared to state-of-the-art sparse Bayesian estimators. We showed a significant robustness compared to the latter estimators in low and moderate SNR regimes. This is in agreement with the superior sparsity-inducing property of the Bessel K model with $\epsilon < 1$ for highly noisy measurements, as shown in Section 2. Furthermore, the results corroborate that the proposed estimators effectively include the estimation of the noise variance, thus avoiding the need for a training procedure for this parameter.

Acknowledgement

This work was supported by the 4GMCT cooperative research project funded by Intel Mobile Communications, Agilent Technologies, Aalborg University and the Danish National Advanced Technology Foundation.

Appendix A. Approximate Type I estimation using EM

Remember that the Type I estimator is the maximizer of

$$\mathcal{L}(\mathbf{w}) = \log(p(\mathbf{y}|\mathbf{w}, \lambda)p(\mathbf{w})). \quad (\text{A.1})$$

We formulate the EM algorithm approximating the Type I estimator by selecting $\{\gamma, \mathbf{y}\}$ to be the complete data for \mathbf{w} . The E-step of the EM algorithm computes the conditional expectation

$$\langle \log p(\mathbf{y}, \mathbf{w}, \gamma) \rangle_{p(\gamma, \hat{\mathbf{w}})} \quad (\text{A.2})$$

with

$$p(\gamma; \hat{\mathbf{w}}) \propto \prod_i \gamma_i^{\epsilon - \rho - 1} \exp(-\gamma_i^{-1} \rho |\hat{\mathbf{w}}_i|^2 - \gamma_i \eta) \quad (\text{A.3})$$

computed in the E-step. The right-hand side expression in (A.3) is recognized as the product of GIG pdfs [39], i.e., $p(\gamma) = \prod_i p(\gamma_i; \nu, a, b_i)$, where, $p(\gamma_i; \nu, a, b_i) = \frac{(a/b_i)^{\frac{\nu}{2}}}{2K_\nu(\sqrt{ab_i})} \gamma_i^{\nu-1} \exp(-\frac{a}{2}\gamma_i - \frac{b_i}{2}\gamma_i^{-1})$ with order $\nu = \epsilon - \rho$ and parameters $a = 2\eta$ and $b_i = 2\rho|\hat{\mathbf{w}}_i|^2$. The moments of the GIG distribution are

given in closed form [39]:

$$\langle \gamma_i^n \rangle = \left(\frac{\rho |\hat{w}_i|^2}{\eta} \right)^{\frac{n}{2}} \frac{K_{\nu+n}(2\sqrt{\rho\eta}|\hat{w}_i|)}{K_{\nu}(2\sqrt{\rho\eta}|\hat{w}_i|)}, \quad n \in \mathbb{R}. \quad (\text{A.4})$$

The M-step of the EM algorithm updates the estimate of \mathbf{w} as the maximizer of (A.2):

$$\hat{\mathbf{w}} = (\Phi^H \Phi + \lambda^{-1} \langle \Gamma^{-1} \rangle)^{-1} \Phi^H \mathbf{y}. \quad (\text{A.5})$$

In case we use the Laplace GSM model ($\nu = \epsilon - \rho = 1/2$), (A.4) with $n = -1$ simplifies to

$$\langle \gamma_i^{-1} \rangle = \frac{\sqrt{\eta/\rho}}{|\hat{w}_i|}, \quad (\text{A.6})$$

where we have invoked the identity $K_{\nu}(\cdot) = K_{-\nu}(\cdot)$ [32].

Appendix B. Results for Section 3.2

This appendix contains the derivations of some results used in Section 3.2.

B.1. Computation of $\langle |w_i|^2 \rangle$

We follow the approach in [36] to compute $\langle |w_i|^2 \rangle$. We can express $\langle |w_i|^2 \rangle$ as $\langle |w_i|^2 \rangle = \mathbf{e}_i^T (\Sigma + \mu \mu^H) \mathbf{e}_i$ with \mathbf{e}_i being an $N \times 1$ vector of all zeros with 1 at the i th position. First, we consider the dependency of Σ in (23) on a single parameter γ_i . We note that $\Sigma = (\lambda \Phi^H \Phi + \sum_{k \neq i} \gamma_k^{-1} \mathbf{e}_k \mathbf{e}_k^T + \gamma_i^{-1} \mathbf{e}_i \mathbf{e}_i^T)^{-1}$. Making use of the matrix inversion lemma [40] we recast Σ as

$$\Sigma = \Sigma_{-i} - \frac{\Sigma_{-i} \mathbf{e}_i \mathbf{e}_i^T \Sigma_{-i}}{\gamma_i + \mathbf{e}_i^T \Sigma_{-i} \mathbf{e}_i}, \quad (\text{B.1})$$

where $\Sigma_{-i} \triangleq (\lambda \Phi^H \Phi + \sum_{k \neq i} \gamma_k^{-1} \mathbf{e}_k \mathbf{e}_k^T)^{-1}$. After some straightforward algebraic manipulations, $\langle |w_i|^2 \rangle$ can be expressed as

$$\langle |w_i|^2 \rangle = \frac{\gamma_i^2 (s_i + |q_i|^2) + \gamma_i s_i^2}{(\gamma_i + s_i)^2} \quad (\text{B.2})$$

with the definitions $s_i \triangleq \mathbf{e}_i^T \Sigma_{-i} \mathbf{e}_i$ and $q_i \triangleq \lambda \mathbf{e}_i^T \Sigma_{-i} \Phi^H \mathbf{y}$.

B.2. Computation of the stationary points of $\ell(\gamma_i)$

We define

$$\ell_i(\gamma_i) \propto \log(p(\mathbf{y}|\gamma_i, \gamma_{-i}, \lambda)p(\gamma_i)). \quad (\text{B.3})$$

Following the steps in [1] we can write $\ell(\gamma_i)$ as

$$\ell_i(\gamma_i) \triangleq -\rho \log[1 + \gamma_i \tilde{s}_i] + \rho \frac{|\tilde{q}_i|^2}{\gamma_i^{-1} + \tilde{s}_i} + (\epsilon - 1) \log \gamma_i - \eta \gamma_i \quad (\text{B.4})$$

with the definitions $\tilde{s}_i \triangleq \Phi_i^H \mathbf{C}_{-i}^{-1} \Phi_i$, $\tilde{q}_i \triangleq \mathbf{y}^H \mathbf{C}_{-i}^{-1} \Phi_i$, and $\mathbf{C}_{-i} \triangleq \lambda^{-1} \mathbf{I} + \sum_{k \neq i} \gamma_k \Phi_k \Phi_k^H$. Taking the derivative of ℓ with respect to γ_i and equating the result to zero yields

$$0 = \eta \tilde{s}_i^2 \gamma_i^3 + \gamma_i^2 [2\eta \tilde{s}_i - (\epsilon - \rho - 1) \tilde{s}_i^2] + \gamma_i [\eta + \rho(\tilde{s}_i - |\tilde{q}_i|^2) - 2(\epsilon - 1) \tilde{s}_i] - (\epsilon - 1). \quad (\text{B.5})$$

Making use of the matrix inversion lemma for \mathbf{C}_{-i}^{-1} , we show the identities $s_i = \tilde{s}_i^{-1}$ and $|q_i|^2 = |\tilde{q}_i|^2 / \tilde{s}_i^2$ [36]. By substituting these identities into (31), we arrive at the

cubic equation in (B.5). Thus, the positive solutions of (31) are the stationary points of (B.4).

References

- [1] M.E. Tipping, A.C. Faul, Fast marginal likelihood maximisation for sparse Bayesian models, in: Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics, 2003.
- [2] S.D. Babacan, R. Molina, A.K. Katsaggelos, Bayesian compressive sensing using Laplace priors, *IEEE Trans. Image Process.* 19 (1) (2010) 53–63.
- [3] R. Baraniuk, Compressive sensing, *IEEE Signal Process. Mag.* 24 (4) (2007) 118–121.
- [4] E.J. Candes, M.B. Wakin, An introduction to compressive sampling, *IEEE Signal Process. Mag.* 25 (2) (2008) 21–30.
- [5] D.G. Tzikas, A.C. Likas, N.P. Galatsanos, The variational approximation for Bayesian inference, *IEEE Signal Process. Mag.* 25 (6) (2008) 131–146.
- [6] W. Bajwa, J. Haupt, A. Sayeed, R. Nowak, Compressed channel sensing: a new approach to estimating sparse multipath channels, *Proceedings of the IEEE* 98 (6) (2010) 1058–1076.
- [7] N.L. Pedersen, C.N. Manchón, D. Shutin, B.H. Fleury, Application of Bayesian hierarchical prior modeling to sparse channel estimation, in: Proceedings of the IEEE International conference on Communications (ICC), 2012, pp. 3487–3492.
- [8] N.L. Pedersen, C.N. Manchón, B.H. Fleury, A fast iterative Bayesian inference algorithm for sparse channel estimation, in: Proceedings of the IEEE International conference on Communications (ICC), 2013, pp. 4591–4596.
- [9] D. Shutin, B.H. Fleury, Sparse variational Bayesian SAGE algorithm with application to the estimation of multipath wireless channels, *IEEE Trans. Signal Process.* 59 (2011) 3609–3623.
- [10] R. Tibshirani, Regression shrinkage and selection via the LASSO, *J. R. Statist. Soc. B* (1994) 267–288.
- [11] S.S. Chen, D.L. Donoho, M.A. Saunders, Atomic decomposition by basis pursuit, *SIAM J. Sci. Comput.* 20 (1998) 33–61.
- [12] J.A. Tropp, Greed is good: Algorithmic results for sparse approximation, *IEEE Trans. Inform. Theory* 50 (2004) 2231–2242.
- [13] D. Needell, J. Tropp, CoSaMP: Iterative signal recovery from incomplete and inaccurate samples, *Appl. Comput. Harmon. Anal.* 26 (3) (2009) 301–321.
- [14] M. Tipping, Sparse Bayesian learning and the relevance vector machine, *J. Mach. Learning Res.* 1 (2001) 211–244.
- [15] D.P. Wipf, B.D. Rao, Sparse Bayesian learning for basis selection, *IEEE Trans. Signal Process.* 52 (8) (2004) 2153–2164.
- [16] D.F. Andrews, C.L. Mallows, Scale mixtures of normal distributions, *J. R. Statist. Soc. B (Methodol.)* 36 (1974) 99–102.
- [17] O. Barndorff-Nielsen, J. Kent, M. Sorensen, Normal variance-mean mixtures and z distributions, *Int. Statist. Rev.* 50 (1982) 145–159.
- [18] T. Gneiting, Normal scale mixtures and dual probability densities, *J. Statist. Comput. Simul.* 59 (1997) 375–384.
- [19] T. Eltoft, T. Kim, T.-W. Lee, Multivariate scale mixture of Gaussians modeling, *Lect. Notes Comput. Sci.* 3889 (2006) 799–806.
- [20] J.A. Palmer, D.P. Wipf, K. Kreutz-Delgado, B.D. Rao, Variational EM algorithm for non-Gaussian latent variable models, in: Advances in Neural Information Processing Systems, NIPS, 2006.
- [21] D.P. Wipf, B.D. Rao, S. Nagarajan, Latent variable Bayesian models for promoting sparsity, *IEEE Trans. Inform. Theory* 57 (9) (2011) 6236–6255.
- [22] D.J.C. MacKay, Bayesian interpolation, in: *Neural Computation*, vol. 4, 1992, pp. 415–447.
- [23] M. Figueiredo, Adaptive sparseness for supervised learning, *IEEE Trans. Pattern Anal. Mach. Intel.* 25 (9) (2003) 1150–1159.
- [24] E.J. Candes, M.B. Wakin, S. Boyd, Enhancing sparsity by reweighted ℓ_1 minimization, *J. Fourier Anal. Appl.* 14 (5) (2008) 877–905.
- [25] J.E. Griffin, P.J. Brown, Bayesian adaptive lassos with non-convex penalization, Technical Report, Dept. of Statistics, University of Warwick, 2007.
- [26] A. Lee, F. Caron, A. Doucet, C. Holmes, A hierarchical Bayesian framework for constructing sparsity-inducing priors, 2010, [arXiv:1009.1914](https://arxiv.org/abs/1009.1914).
- [27] J.A. Palmer, D.P. Wipf, K. Kreutz-Delgado, B.D. Rao, Probabilistic formulation of independent vector analysis using complex Gaussian scale mixtures, in: Advances in Neural Information Processing Systems, NIPS, 2009.
- [28] Y. Rakwongthai, A. Vo, S. Orintara, Complex Gaussian scale mixtures of complex wavelet coefficients, *IEEE Trans. Signal Process.* 58 (7) (2010) 3545–3556.

- [29] M. Kyung, J. Gill, M. Ghosh, G. Casella, Penalized regression, standard errors, and Bayesian lassos, *Bayesian Anal.* 5 (2) (2010) 369–412.
- [30] F. Caron, A. Doucet, Sparse Bayesian nonparametric regression, in: *Proceedings of the 25th International Conference on Machine Learning*, 2008, pp. 88–95.
- [31] J.E. Griffin, P.J. Brown, Inference with normal-gamma prior distributions in regression problems, *Bayesian Anal.* 5 (1) (2010) 171–188.
- [32] M. Abramowitz, I.A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover, 1972.
- [33] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, D. Gorinevsky, An interior-point method for large-scale ℓ_1 -regularized least squares, *IEEE J. Sel. Topics Signal Process.* 1 (4) (2007) 606–617.
- [34] S.J. Wright, R.D. Nowak, M.A.T. Figueiredo, Sparse reconstruction by separable approximation, *IEEE Trans. Signal Process.* 57 (7) (2009) 2479–2493.
- [35] D.P. Wipf, S. Nagarajan, Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions, *IEEE J. Sel. Topics Signal Process.* 4 (2) (2010) 317–329.
- [36] D. Shutin, T. Buchgraber, S.R. Kulkarni, H.V. Poor, Fast variational sparse Bayesian learning with automatic relevance determination for superimposed signals, *IEEE Trans. Signal Process.* 59 (12) (2011) 6257–6261.
- [37] S. Ji, Y. Xue, L. Carin, Bayesian compressive sensing, *IEEE Trans. Signal Process.* 56 (6) (2008) 2346–2356.
- [38] E.J. Candes, T. Tao, The Dantzig selector: Statistical estimation when p is much larger than n , *Ann. Statist.* 36 (6) (2007) 2313–2351.
- [39] B. Jorgensen, *Statistical Properties of the Generalized Inverse Gaussian Distribution* (Lecture Notes in Statistics 9), Springer-Verlag, New York, 1982.
- [40] K.V. Mardia, J.T. Kent, J.M. Bibby, *Multivariate Analysis. Probability and Mathematical Statistics*, Academic Press, London, 1979.