



HAL
open science

Regularized Bayesian compressed sensing in ultrasound imaging

Nicolas Dobigeon, Adrian Basarab, Denis Kouamé, Jean-Yves Tournet

► **To cite this version:**

Nicolas Dobigeon, Adrian Basarab, Denis Kouamé, Jean-Yves Tournet. Regularized Bayesian compressed sensing in ultrasound imaging. 20th European Signal and Image Processing Conference (EUSIPCO 2012), European Association for Signal, Speech, and Image Processing (EURASIP), Aug 2012, Bucharest, Romania. pp.2600–2604. hal-03146722

HAL Id: hal-03146722

<https://hal.science/hal-03146722>

Submitted on 22 Feb 2021

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

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

REGULARIZED BAYESIAN COMPRESSED SENSING IN ULTRASOUND IMAGING

Nicolas Dobigeon⁽¹⁾, Adrian Basarab⁽²⁾, Denis Kouamé⁽²⁾ and Jean-Yves Tournet⁽¹⁾

⁽¹⁾ University of Toulouse, IRIT/INP-ENSEEIH, 31071 Toulouse Cedex 7, France

⁽²⁾ University of Toulouse, IRIT/UPS, 31062 Toulouse Cedex 9, France

ABSTRACT

Compressed sensing has recently shown much interest for ultrasound imaging. In particular, exploiting the sparsity of ultrasound images in the frequency domain, a specific random sampling of ultrasound images can be used advantageously for designing efficient Bayesian image reconstruction methods. We showed in a previous work that assigning independent Bernoulli Gaussian priors to the ultrasound image in the frequency domain provides Bayesian reconstruction errors similar to a classical ℓ_1 minimization technique. However, the advantage of Bayesian methods is to estimate the sparsity level of the image by using a hierarchical algorithm. This paper goes a step further by exploiting the spatial correlations between the image pixels in the frequency domain. A new Bayesian model based on a correlated Bernoulli Gaussian model is proposed for that purpose. The parameters of this model can be estimated by sampling the corresponding posterior distribution using an MCMC method. The resulting algorithm provides very low reconstruction errors even when reducing significantly the number of measurements via random sampling.

Index Terms— Ultrasound imaging, compressed sensing, Bayesian inference, Markov random field.

1. INTRODUCTION

Ultrasound (US) imaging is one of the most popular medical imaging techniques and represents the gold standard in many crucial diagnostic exams such as obstetrics and cardiology. The main advantages of US imaging are its relatively low cost, its innocuity for the patient, its ease of use and real time nature. However, the real-time property is sometimes limited by the acquisition time or by the high amount of acquired data, especially in 3D ultrasound imaging. Even in 2D applications, a higher frame rate could be beneficial, i.e., for cardiac US monitoring. For this reason, a few research groups have recently started to evaluate the feasibility of US acquisitions using the compressive sampling (CS) framework [1, 2]. In particular, Friboulet *et al.* have presented in [3] a method for randomly sub-sampling the US raw data (signals before beamforming and classically used in US imaging for obtaining the radiofrequency (RF) lines). The idea

investigated in [3] is missing data reconstruction using the assumption of sparsity in the wave-atom domain and a gradient based optimization technique. Another attempt of using the CS framework in US imaging has been presented in [4]. More precisely, an acquisition scheme and a numerical reconstruction have been introduced starting from the analog domain and based on the hypothesis that US axial profiles of the images (the RF lines) can be modeled by pulse streams.

We have recently studied a new CS technique for US images based on: i) specific sampling schemes adapted to US acquisition and respecting the constraints of incoherence and ii) numerical reconstruction techniques for RF images [5–7]. Concerning the sampling schemes adopted for RF US images, two sampling masks have been investigated. The first mask, which guarantees maximal incoherence with the sparsifying basis (the Fourier domain in our case), is a uniform random pattern in the two spatial directions. The second mask, less incoherent with the Fourier domain but far more adapted to US imaging, considers that randomly chosen columns of the RF images (in other words randomly chosen RF lines) are not sampled at all. With the second sampling pattern, the number of US emitted pulses necessary for the acquisition of one image can be reduced. As a consequence the acquisition time is accelerated. For this reason, this paper concentrates on this second decimation pattern.

Concerning the missing data reconstruction, a classical gradient based technique was initially investigated [5–7]. However, a new Bayesian method has been recently proposed in [8]. This Bayesian method is based on a Bernoulli Gaussian prior for the US image Fourier transform (FT). Despite the reconstruction errors equivalent to those obtained by variational techniques (for the 2D random pattern sampling), the Bayesian reconstruction is interesting since it does not require any hyperparameter adjustment and allows the sparsity level in the Fourier domain to be estimated. However, the results presented in [8] focused only on the completely random decimation pattern which is not always realistic in practical applications.

In this paper, we propose to improve the Bayesian model used in [8] by exploiting the spatial correlations between non-zero pixels in the Fourier domain. In the literature, a similar idea has recently been explored for reconstructing cluster structured sparse signals [9]. In US imaging, the characteris-

tics of the RF images (e.g., the RF lines are modulated signals) and the US acquisition (e.g., the properties of the point spread function may be partially known) allow some assumptions on the image FT to be made [10]. More precisely, we can assume that the FT of an RF image is supported by two symmetrical compact non-zero regions, located around the central frequency of the US probe. This assumption is taken into account in the Bayesian model proposed in this paper.

The paper is organized as follows: Section 2 introduces the Bayesian model considered for US image reconstruction and its corresponding posterior distribution. Section 3 studies a Gibbs sampler that can be used to generate samples distributed according to this posterior. Simulation results are presented in Section 4. Results on a realistic simulated US image show that the proposed method outperforms the Bayesian method of [8] and the ℓ_1 minimization technique (based on a conjugate gradient algorithm) recently studied in [5–7]. Conclusions are finally reported in Section 5.

2. BAYESIAN MODEL

CS consists of estimating a sparse signal \mathbf{x} from a noisy observation vector \mathbf{y} , from projections on a complex $M \times N$ random matrix \mathbf{T} (with $M \ll N$)

$$\mathbf{y} = \mathbf{T}\mathbf{x} + \boldsymbol{\eta} \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_M)^T \in \mathbb{C}^M$ is the measurement vector, $\mathbf{x} = (x_1, \dots, x_N)^T \in \mathbb{C}^N$ is the unknown parameter vector and $\boldsymbol{\eta} = (\eta_1, \dots, \eta_M)^T \in \mathbb{C}^M$ is an additive complex noise including instrumentation errors and errors due to the signal being only approximately sparse [8].

In our case, \mathbf{y} represents the measured spatial samples which consist of randomly skipping RF lines in the image lateral direction [6], \mathbf{x} is the FT of the US image to be reconstructed and $\mathbf{T} = \boldsymbol{\Psi}\mathbf{F}^{-1}$ is the sampling matrix, with $\boldsymbol{\Psi}$ indicating the random positions of the measured spatial samples and \mathbf{F}^{-1} standing for the inverse Fourier matrix. The problem addressed in this work consists of estimating the sparse signal \mathbf{x} from the random projections \mathbf{y} . This sparse reconstruction task is formulated as an estimation problem solved within a Bayesian framework. The likelihood function and the parameter priors involved in the Bayesian model are introduced in what follows.

2.1. Likelihood function

Assuming the additive noise in (1) is Gaussian complex with zero mean and unknown variance σ^2 , the likelihood is

$$p(\mathbf{y}|\mathbf{x}, \sigma^2) = (\pi\sigma^2)^{-N} \exp\left(-\frac{1}{\sigma^2} \|\mathbf{y} - \mathbf{T}\mathbf{x}\|_2^2\right)$$

where N is the number of image pixels and $\|x\| = \sqrt{x^T x}$ is the standard ℓ_2 norm.

2.2. Parameter priors

The prior distribution chosen for the image pixels \mathbf{x} should exploit the Gaussian properties of ultrasound images [11] and the image sparsity in the frequency domain. In [8], each pixel x_i of the image has been assigned a prior defined as a mixture of a centered complex normal distribution $\mathcal{CN}(0, \sigma_x^2)$ and a mass at the origin

$$p(x_i|w, \sigma_x^2) = (1-w)\delta(|x_i|) + \left(\frac{w}{\pi\sigma_x^2}\right) \exp\left(-\frac{|x_i|^2}{\sigma_x^2}\right) \quad (2)$$

where the hyperparameter w is the prior probability of having a non-zero pixel in the image. A similar distribution has already been proposed in [12] for sparse real signal reconstruction. In [8] the image pixels of the 2D FT of the US image have been assumed to be *a priori* independent. However, ultrasound signals are bandlimited and are commonly assumed to be narrowband around the central frequency of the probe [10]. As a consequence, we propose to enforce the non-zero pixels of the 2D US image FT to be aggregated in homogeneous regions. The strategy advocated to introduce correlations between non-zero pixels is described below.

As noticed in [13], the prior distribution in (2) can be rewritten

$$p(x_i|w, \sigma_x^2) = \sum_{\epsilon \in \{0,1\}} p(x_i|q_i = \epsilon, \sigma_x^2) P[q_i = \epsilon|w] \quad (3)$$

where q_i is a hidden binary variable indicating if x_i is active or not

$$q_i = \begin{cases} 1, & \text{if } x_i \neq 0; \\ 0, & \text{otherwise.} \end{cases}$$

Assigning independent Bernoulli distributions $\mathcal{B}e(w)$ as priors for the q_i 's, i.e.,

$$\begin{aligned} P[q_i = 1|w] &= w \\ P[q_i = 0|w] &= 1 - w \end{aligned} \quad (4)$$

and the conditional prior for x_i

$$p(x_i|q_i, \sigma_x^2) = (1 - q_i)\delta(|x_i|) + \left(\frac{q_i}{\pi\sigma_x^2}\right) \exp\left(-\frac{|x_i|^2}{\sigma_x^2}\right).$$

exactly leads to (2). Forcing the non-zero pixels to be grouped implies that the indicator variables q_i are correlated such that the prior probabilities (4) are no longer appropriate. Conversely, we propose to model the indicator vector $\mathbf{q} = [q_1, \dots, q_n]$ as a Markov random field (MRF), which reduces here to the Ising model

$$p(\mathbf{q}|\beta) \propto \exp\left[\beta \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \kappa(q_i - q_j)\right] \quad (5)$$

where the granularity coefficient β tunes the degree of homogeneity of the regions, $\kappa(\cdot)$ is the Kronecker function and \mathcal{V}_i denotes the index set of the neighboring pixels of x_i .

As in [8], the Bayesian model is complemented by a classical Jeffreys prior for the noise variance

$$f(\sigma^2) \propto \frac{1}{\sigma^2}$$

and a conjugate inverse gamma distribution with parameters α_0 and α_1 for the non-zero pixel variance

$$\sigma_x^2 \sim \mathcal{IG}(\alpha_0, \alpha_1).$$

The hyperparameters α_0 and α_1 are fixed to constant values ensuring the resulting inverse gamma distribution is non-informative.

The image to be reconstructed is then estimated from the joint posterior distribution $p(\mathbf{x}, \mathbf{q}, \sigma^2, \sigma_x^2 | \mathbf{y})$, computed from the following hierarchical Bayesian structure

$$p(\mathbf{x}, \mathbf{q}, \sigma^2, \sigma_x^2 | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}, \sigma^2) p(\sigma^2) \times p(\mathbf{x} | \mathbf{q}, \sigma_x^2) p(\mathbf{q} | \beta) p(\sigma_x^2). \quad (6)$$

We propose to use a Markov chain Monte Carlo (MCMC) method (known as the Gibbs sampler) for generating vectors that are asymptotically distributed according to the joint posterior distribution (6). These generated vectors will be denoted as $\{\mathbf{x}^{(t)}, \mathbf{q}^{(t)}, \sigma^{2(t)}, \sigma_x^{2(t)}\}_{t=1, \dots, N_{MC}}$. The Bayesian estimators such as the maximum *a posteriori* (MAP) or the minimum mean square error (MMSE) estimators relative to the posterior (6) can then be approximated by using these generated samples (see [13] for more details).

3. GIBBS SAMPLER

To generate samples asymptotically distributed according to (6), we propose to use a Gibbs sampler that iteratively samples according to the conditional distributions associated with the joint posterior $p(\mathbf{x}, \mathbf{q}, \sigma^2, \sigma_x^2 | \mathbf{y})$. The successive steps of the Gibbs sampler are summarized below.

3.1. Sampling according to $p(x_i | \mathbf{x}_{\setminus i}, q_i, \sigma^2, \sigma_x^2, \mathbf{y})$

The conditional distribution of each pixel x_i given $\mathbf{x}_{\setminus i}$ ($\mathbf{x}_{\setminus i}$ denotes the vector \mathbf{x} whose i th element has been removed) is governed by the associated indicator variable q_i

$$\begin{aligned} x_i | q_i = 0 & \sim \delta(|x_i|) \\ x_i | q_i = 1, \mathbf{x}_{\setminus i}, \sigma^2, \sigma_x^2 & \sim \mathcal{CN}(\mu_i, \eta_i^2) \end{aligned} \quad (7)$$

with

$$\begin{cases} \eta_i^2 & = \left(\frac{1}{\sigma_x^2} + \frac{\|\mathbf{t}_i\|^2}{\sigma^2} \right)^{-1} \\ \mu_i & = \mathbf{t}_i^H \mathbf{e}_i \frac{\eta_i^2}{\sigma^2} \\ \mathbf{e}_i & = \mathbf{y} - \sum_{j \neq i} x_j \mathbf{t}_j \end{cases}$$

where \mathbf{t}_i is the i th column of \mathbf{T} .

3.2. Sampling according to $f(q_i | \mathbf{q}_{\setminus i}, x_i, \sigma^2, \sigma_x^2, \mathbf{y})$

The conditional probabilities of the indicator variables q_i are

$$\begin{aligned} P[q_i = 1 | \mathbf{q}_{\setminus i}, x_i, \sigma_x^2] & \propto \\ w_i^{(1)} & \triangleq \frac{1}{\pi \sigma_x^2} \exp \left[\sum_{j \in \mathcal{V}_i} \beta \kappa (1 - q_j) - \frac{|x_i|^2}{\sigma_x^2} \right] \end{aligned} \quad (8)$$

and

$$P[q_i = 0 | \mathbf{q}_{\setminus i}] \propto w_i^{(0)} \triangleq \exp \left[\sum_{j \in \mathcal{V}_i} \beta \kappa (1 - q_j) \right]. \quad (9)$$

Consequently, the indicator q_i is distributed according to the Bernoulli distribution $\mathcal{B}e(\tilde{w}_i^{(1)})$ with $\tilde{w}_i^{(1)} = \frac{w_i^{(1)}}{w_i^{(1)} + w_i^{(0)}}$.

3.3. Sampling according to $f(\sigma^2 | \mathbf{x}, \mathbf{q}, \sigma_x^2, \mathbf{y})$

The conditional distribution of the noise variance is shown to be the following inverse gamma distribution

$$\sigma^2 | \mathbf{x}, \mathbf{y} \sim \mathcal{IG}(M, \|\mathbf{y} - \mathbf{T}\mathbf{x}\|^2). \quad (10)$$

3.4. Sampling according to $f(\sigma_x^2 | \mathbf{x}, \sigma^2, \mathbf{q}, \mathbf{y})$

The conditional distribution of the hyperparameter σ_x^2 (a priori variance for non zero pixels) is

$$\sigma_x^2 | \mathbf{x} \sim \mathcal{IG}(n_1 + \alpha_0, \|\mathbf{x}\|^2 + \alpha_1). \quad (11)$$

4. SIMULATION RESULTS

This section studies the reconstruction results of a simulated ultrasound image. The US image, representing an homogeneous medium crossed by a vessel, has been simulated using Field II simulation program [14]. The simulation parameters were: central frequency = 3MHz, axial sampling frequency = 20MHz, number of simulated RF lines = 256, number of scatterers = 10,000. Note that for visualization reasons, we do not show the RF images, but the corresponding so called B-mode images, obtained after axial demodulation of each RF line and log-compression.

Three reconstruction methods have been tested: the proposed spatially regularized Bayesian reconstruction, the Bayesian method proposed in [8] and a ℓ_1 minimization using a conjugate gradient optimization technique. Note that the last method was combined with the algorithm in [15], which avoids an empirical choice of the hyperparameter controlling the sparsity term with respect to the data fidelity term.

First, the simulated US image has been cropped to a 32×32 pixel block, as shown in Fig. 1(a). The magnitude of its FT is shown in Fig. 1(b).




Fig. 1. (a) Original 32×32 pixel block extracted from a simulated ultrasound image and (b) Magnitude of its 2D FT.

The image has been decimated using a pattern adapted to US imaging. The decimated image is shown in Fig. 2(a) where only part of the columns contain measured samples, at random positions. The remaining columns (RF lines) are not sampled at all. In our example, the reconstructions have been obtained by using 10% of the original images samples. This means that roughly one RF line out of three (at random lateral positions) has been sampled. For each sampled line, roughly 33% of the samples have been considered. Fig. 2(b) shows the FT magnitude corresponding to the measured samples shown in Fig. 2(a). We can observe that the proposed sampling results in a non localized noise in the Fourier domain.




Fig. 2. (a) Random measured spatial samples using the sampling pattern adapted to ultrasound imaging (only part of the columns are randomly sampled) and (b) Magnitude of the FT of the image in (a).

We show in Fig. 3 the reconstruction results obtained with the three methods, as well as the three reconstructed FTs of the image. We observe in Fig. 3(a) and 3(b) that the Bayesian method presented in [8] totally fails to reconstruct the original Fourier domain because the number of missing samples is too large. The algorithm based on a reweighted ℓ_1 minimization by conjugate gradient and the proposed method provide better reconstruction results depicted in Fig. 3(c)–3(f).

To evaluate quantitatively the differences between the different methods, we have computed the normalized mean square errors (NRMSEs) between the actual and reconstructed US images for different decimation factors. The results displayed in Fig. 4 are in favor of the proposed method




Fig. 3. Reconstructed ultrasound image and the magnitude of the reconstructed Fourier domain using: (a) and (b) Bernoulli Gaussian Bayesian approach presented in [8], (c) and (d) reweighted conjugate gradient optimization, (e) and (f) the proposed Bayesian optimization method.

(note that the results obtained with the method of [8] have not been shown here in order to outline the differences between the two other methods). It shows the evolution of the NRMSE for both the proposed method and ℓ_1 minimization by conjugate gradient, for several percentages of measured samples. Obviously, for both methods, the error is smaller when the number of spatial measured samples increases. However, we observe that our method outperforms the classical one for all the percentages. We also remark that the gap between the two methods decreases with the number of measured samples.

With a crude Matlab implementation of the Bayesian algorithm, out of memory errors occur when large images have to be reconstructed. In order to show a result on the whole 128×128 simulated image (from which the 32×32 block shown below was extracted), a block wise reconstruction has been performed with the proposed method. Note that the ℓ_1 conjugate gradient minimization can be performed without problem on the entire US image. The true image, the measured samples and the two reconstructed images are shown in Fig. 5. Despite its non optimal block wise implementation,




Fig. 4. NRMSEs versus the proportion of measured samples for the reweighted ℓ_1 method (red line) and the proposed Bayesian method (black line).

the proposed Bayesian method outperforms the reweighted ℓ_1 minimization strategy. Note that the dark region in the image, corresponding to the vessel, shows more reconstruction errors. Indeed the vessel region is characterized by an absence of scatterer therefore yielding a weak signal that is more difficult to be reconstructed than other parts of the image.

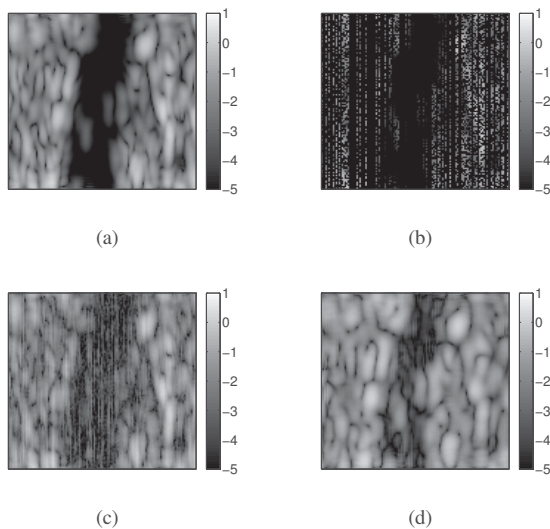


Fig. 5. (a) Original US image, (b) measured samples, (c) image reconstructed using reweighted conjugate gradient optimization and (d) image reconstructed using the proposed Bayesian method.

5. CONCLUSION

This paper presented a new Bayesian ultrasound image reconstruction technique. Contrary to the Bernoulli Gaussian prior considered in [8], the proposed model managed to take advantage of the spatial connectivity between non-zero pixels of the US image Fourier transform. The resulting algorithm

provided better reconstruction errors than the algorithm of the [8] and the ℓ_1 minimization technique of [7], especially for a small amount of measured samples. These very encouraging reconstructions promote the use of a new random sampling scheme adapted to ultrasound imaging, allowing only part of the RF lines to be acquired. Future works include the implementation of the proposed algorithm using graphics processing units (GPU) to decrease the execution times that are still prohibitive for large images.

6. REFERENCES

- [1] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Jan. 2006.
- [2] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, March 2008.
- [3] D. Friboulet, H. Liebgott, and R. Prost, "Compressive sensing for raw RF signals reconstruction in ultrasound," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, oct. 2010, pp. 367–370.
- [4] R. Tur, Y. Eldar, and Z. Friedman, "Innovation rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1827–1842, April 2011.
- [5] C. Quinsac, A. Basarab, J.-M. Girault, and D. Kouame, "Compressed sensing of ultrasound images: Sampling of spatial and frequency domains," in *Proc. IEEE Workshop on Signal Process. Syst. (SIPS)*, Oct. 2010, pp. 231–236.
- [6] C. Quinsac, A. Basarab, J.-M. Gregoire, and D. Kouame, "3D compressed sensing ultrasound imaging," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, San Diego, CA, Oct. 2010.
- [7] C. Quinsac, F. de Vieilleville, A. Basarab, and D. Kouame, "Compressed sensing of ultrasound single-orthant analytical signals," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, San Diego, CA, Oct. 2010.
- [8] C. Quinsac, N. Dobigeon, A. Basarab, J.-Y. Tournet, and D. Kouamé, "Bayesian compressed sensing in ultrasound imaging," in *Proc. IEEE Int. Workshop Comput. Adv. in Multi-Sensor Adaptive Process. (CAM-SAP)*, San Juan, Puerto Rico, Dec. 2011, pp. 101–104.
- [9] L. Yu, H. Sun, J. P. Barbot, and G. Zheng, "Bayesian compressive sensing for cluster structured sparse signals," *Signal Process.*, vol. 92, no. 1, pp. 259–269, 2012.
- [10] T. Hergum, S. Langeland, E. W. Remme, and H. Torp, "Fast ultrasound imaging simulation in k-space," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 56, no. 6, pp. 1159–1167, June 2009.
- [11] R. F. Wagner, M. F. Insana, and D. G. Brown, "Statistical properties of radio-frequency and envelope-detected signals with applications to medical ultrasound," *J. Opt. Soc. Am. A*, vol. 4, no. 5, pp. 910–922, May 1987.
- [12] N. Dobigeon, A. O. Hero, and J.-Y. Tournet, "Hierarchical Bayesian sparse image reconstruction with application to MRFM," *IEEE Trans. Image Process.*, vol. 18, no. 9, pp. 2059–2070, Sept. 2009.
- [13] N. Dobigeon and J.-Y. Tournet, "Bayesian orthogonal component analysis for sparse representation," *IEEE Trans. Signal Process.*, vol. 58, no. 5, pp. 2675–2685, May 2010.
- [14] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 39, no. 2, pp. 262–267, March 1992.
- [15] E. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, no. 5, pp. 877–905, Dec. 2008.