Medical ultrasound image reconstruction using compressive sampling and lp norm minimization

Adrian Basarab, Alin Achim, Denis Kouamé

Medical ultrasound image reconstruction using compressive sampling and l_p norm minimization

Adrian Basarab[1], Alin Achim[2], Denis Kouamé[1]

[1]Université de Toulouse, IRIT, CNRS UMR 5505, Université Paul Sabatier, Toulouse, France
[2]University of Bristol, Department of Electrical & Electronic Engineering, Bristol, UK

Abstract

In the last four years, a few research groups worked on the feasibility of compressive sampling (CS) in ultrasound medical imaging and several attempts of applying the CS theory may be found in the recent literature. In particular, it was shown that using l_p-norm minimization with p different from 1 provides interesting RF signal reconstruction results. In this paper, we propose to further improve this technique by processing the reconstruction in the Fourier domain. In addition, α-stable distributions are used to model the Fourier transforms of the RF lines. The parameter p used in the optimization process is related to the parameter α obtained by modelling the data (in the Fourier domain) as an α-stable distribution. The results obtained on experimental US images show significant reconstruction improvement compared to the previously published approach where the reconstruction was performed in the spatial domain.

1 Introduction

The low cost, non-ionizing characteristics, ease of use and real-time nature of ultrasound (US) imaging make it one of the most commonly used medical imaging modalities in a number of clinical applications. The real-time property is however limited by the volume of data, especially in 3D applications. Specific 2D applications such as heart monitoring could also benefit from higher acquisition frame rates. Compressive sensing could prove to be a powerful solution to enhance US images frame rate.

Compressive sampling (CS) [1] is a recent theory providing theoretical guarantees of "perfect" signal or image reconstruction from relatively small amount of measurements (below the well known Shannon-Nyquist's criterion), based on two key conditions: i) the signal or image must be sparse in a known basis and ii) the measurements should be incoherent with this basis. If these conditions are respected, the reconstruction is usually done using greedy methods such as orthogonal matching pursuit [2], or l_1 (or more generally l_p) norm minimization algorithms.

In this context, in the last four years, a few research groups worked on the feasibility of compressive sampling in US imaging and several attempts of applying the CS theory may be found in the recent literature. For a complete overview the reader may refer to [3]. Based on various acquisition schemes and sparsity bases, the existing applications of CS in US imaging generally use basis pursuit l_1-norm minimization algorithms. Moreover, it was shown in [4] that using l_p-norms (with p different from 1) provides better RF signal reconstruction. The value of p was automatically chosen by relating it to the α-stable statistics of the RF signals.

In this paper, we show that the results provided in [4] may be further improved by doing the reconstruction in the Fourier domain. In the present work, the value of parameter p was related to the α-stable statistics in the Fourier domain of the RF echoes (I/Q signals). Our reconstruction process relies on the iteratively reweighted least squares approach (IRLS) [5].

2 Basics of compressive sampling theory

Compressive sampling proposes theoretical guarantees and practical implementation schemes for sparse signal reconstruction from a relatively small number of linear measurements. Consider a vector $\boldsymbol{x}$ of N elements, S-sparse in a given basis Ω. That is, the vector $\Omega\boldsymbol{x} \in \mathbb{R}^{\boldsymbol{N\times 1}}$ has only $S < N$ elements different from 0. The main idea behind CS is to reconstruct $\boldsymbol{x}$, or more precisely $\Omega\boldsymbol{x}$, from $M \ll N$ linear measurements regrouped in a vector $\boldsymbol{y} \in \mathbb{R}^{M\times 1}$. Commonly, $\boldsymbol{y}$ is taken as the projection of $\boldsymbol{x}$ on M random Gaussian vectors forming a matrix $\Phi \in \mathbb{R}^{M\times N}$. That is, $\boldsymbol{y} = \Phi\boldsymbol{x}$. The reconstruction is done by solving the following optimization problem:

$$\hat{\boldsymbol{x}} = \min_{\mathrm{x}} \|\Omega \boldsymbol{x}\|_1 \quad \text{subject to} \quad \Phi \boldsymbol{x} = \boldsymbol{y}$$

In other words, we search for the sparsest representation of $\boldsymbol{x}$ (in the sense of the l_1 norm) in the Ω domain that respects the measurements. Depending on the choice of Ω and Φ (namely the incoherence between the two), several theoretical results may be found in the literature giving the smallest number of measurements required for a perfect reconstruction with exceedingly high probability. The smallest number of measurements required is in the most classical cases a product between S, log(N) and a constant depending on the incoherence between Ω and Φ.

3 Proposed methodology

3.1 α-stable statistics

The statistics of ultrasound images are an ongoing research subject with various applications such as tissue characterization or image segmentation. Recently, a few research teams showed that the α-stable distribution is well suited to model the RF signals [4, 6, 7]. We recall the characteristic function of a symmetric α-stable random variable as:

$$\varphi(\omega) = exp(j\delta\omega - \gamma|\omega|^{\alpha}), \tag{1}$$

where α is the characteristic exponent taking values in the interval (0,2], $\delta \in (-\infty, \infty)$ is the location parameter and γ (strictly positive) is the dispersion of the distribution.
The Fourier transform of the RF lines can also be modeled by an α-stable distribution. In our work, in order to determine the value of α and the parameter p used in the l_p-norm reconstruction process (related to α), we use the same methodology as the one proposed in [4].

3.2 Fourier transforms reconstruction via l_p norm minimisation

We denote hereafter by $\boldsymbol{X} \in \mathbb{R}^{\boldsymbol{N}\times\boldsymbol{J}}$ an US RF image formed by J RF signals of length N, denoted by $\boldsymbol{x_1}$, $\boldsymbol{x_2}$,…,$\boldsymbol{x_J}$. Moreover, we denote by $\boldsymbol{\theta_j} \in \mathbb{C}^{\boldsymbol{N}\times\boldsymbol{1}}$ the 1D Fourier transforms of $\boldsymbol{x_j}$, for j running from 1 to J.

$$\boldsymbol{\theta_j} = \boldsymbol{\mathcal{F}}\boldsymbol{x_j}, \quad j \in \{1,2,\dots,J\}, \tag{2}$$

where $\boldsymbol{\mathcal{F}} \in \mathbb{C}^{\boldsymbol{N}\times\boldsymbol{N}}$ is the 1D Fourier matrix.
In this paper we propose to process the reconstruction using l_p-norm minimization and the α-stable fitting in the Fourier domain. As we compare our results to those obtained using [4], where the reconstruction and the fitting are done in the spatial domain, we define hereafter the two problems. We denote by P_1 the model used in [4] and by P_2 the proposed one. The sampling schemes corresponding to the two problems are:

$$\begin{aligned} P_1&: \quad \boldsymbol{y_j} = \boldsymbol{\Phi_j}\boldsymbol{x_j} \;\; with \;\; j \in \{1,2,\dots,J\} \\ P_2&: \quad \boldsymbol{m_j} = \boldsymbol{\Psi_j}\boldsymbol{\theta_j} = \boldsymbol{\Psi_j}\boldsymbol{\mathcal{F}}\boldsymbol{x_j} \;\; with\, j \in \{1,2,\dots,J\} \end{aligned} \tag{3}$$

In (3), $\boldsymbol{\Phi_j}$ and $\boldsymbol{\Psi_j}$ stand for Gaussian random matrices of size $M \times N$, with $M<<N$ the number of measurements. The measurements are denoted by $\boldsymbol{y_j} \in \mathbb{R}^{\boldsymbol{M}\times\boldsymbol{1}}$ in the spatial domain and by $\boldsymbol{m_j} \in \mathbb{C}^{\boldsymbol{M}\times\boldsymbol{1}}$ in the Fourier domain. The purpose of problems P_1 and P_2 is to respectively recover the RF signals $\boldsymbol{x_j}$ and the Fourier transforms $\boldsymbol{\theta_j}$ from the corresponding measurements.
In [4], it was shown that the problem P_1 is more accurately solved using l_p norm minimization than classical basis pursuit with minimal l_1 prior. Thus, the optimization problem was formulated as follows:

$$P_1: \quad \hat{\boldsymbol{x}}_{\boldsymbol{j}} = \min \left\|\boldsymbol{x_j}\right\|_p \quad \text{subject to} \quad \left\|\boldsymbol{\Phi_j}\boldsymbol{x_j} - \boldsymbol{y_j}\right\|_2^2 < \lambda\,, \quad \text{for } j \in \{1,2,\dots,J\} \tag{4}$$

The problem P_2 addressed here is also solved using the minimization of the l_p norm constrained with respect to the measurements. Thus, the problem P_2 consists in finding the vectors $\boldsymbol{\theta_j}$ with the minimum l_p norm by solving:

$$P_2: \quad \widehat{\boldsymbol{\theta}}_{\boldsymbol{j}} = \min \|\boldsymbol{\theta}_{\boldsymbol{j}}\|_p \quad \text{subject to} \quad \|\boldsymbol{\Psi}_{\boldsymbol{j}}\boldsymbol{\theta}_{\boldsymbol{j}} - \boldsymbol{m}_{\boldsymbol{j}}\|_2^2 < \lambda\,, \quad \text{for } j \in \{1,2,\dots,J\}, \tag{5}$$

where λ is a hyper-parameter accounting for the compromise between the data attachment and l_p-norm minimization. For this, we used the modified IRLS algorithm proposed in [5]. The value of p, following the considerations in [4], was set at $\alpha - 0.01$, where α was obtained by fitting an α-stable distribution to the data. Finally the RF lines are obtained by inverting the estimated Fourier transforms:

$$\widehat{\boldsymbol{x}}_{\boldsymbol{j}} = \boldsymbol{\mathcal{F}}^{-1}\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}, \quad j \in \{1,2,\dots,J\}, \tag{6}$$

3.3 l_p norm minimization via IRLS algorithm

In this paper, the optimization problems P_1 and P_2 given in (4) and (5) are solved using the iteratively reweighted least squares approach (IRLS) [5]. The main idea of IRLS algorithm is to replace the l_p norms in (4) and (5) by weighted l_2 norms. Thus P_2 in (5) is replaced in the IRLS optimization process by the following P_2^{IRLS} problem:

$$P_2^{IRLS}: \quad \widehat{\boldsymbol{\theta}}_{\boldsymbol{j}} = \min \sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{w_i}\theta_{\mathrm{ji}}^2 \quad \text{subject to} \quad \|\boldsymbol{\Psi}_{\boldsymbol{j}}\boldsymbol{\theta}_{\boldsymbol{j}} - \boldsymbol{m}_{\boldsymbol{j}}\|_2^2 < \lambda\,, \quad \text{for } j \in \{1,2,\dots,J\}, \tag{5}$$

Where θ_{ji} stands for the i[th] elements of vector $\boldsymbol{\theta}_{\boldsymbol{j}}$. The weights $\mathrm{w_i}$ and the estimated vector $\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}$ are iteratively updated as shown by the pseudo-code given below. Note that in our work the real and imaginary parts of the Fourier transforms $\boldsymbol{\theta}_{\boldsymbol{j}}$ are estimated separately. Thus, in the pseudo code, all the vectors contain only real values.

For each j running from 1 to J

a. Initialization:

Iteration number $k = 0$

Residual tolerance $\varepsilon = 1$

Initial estimate $\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(0)} = \boldsymbol{\Psi}_{\boldsymbol{j}}^{T}\boldsymbol{m}_{\boldsymbol{j}}$

Set p equal to $\alpha - 1$, where α is obtained by α -stable fitting as shown in next section

b. **While** $\varepsilon > 10^{-8}$ and $k < 1000$

i. $k = k + 1$

ii. Find the weights $\mathrm{w_i} = \left(\left(\hat{\theta}_{\mathrm{ji}}\right)^2 + \varepsilon\right)^{\frac{\mathrm{p}}{2}-1}$

iii. Form a diagonal matrix $\boldsymbol{Q}_k$ whose entries are $\frac{1}{\mathrm{w_i}}$

Calculate the current estimate: $\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k)} = \boldsymbol{Q}_k\boldsymbol{\Psi}_{\boldsymbol{j}}^{T}\left(\boldsymbol{\Psi}_{\boldsymbol{j}}\boldsymbol{Q}_k\boldsymbol{\Psi}_{\boldsymbol{j}}^{T}\right)^{-1}\boldsymbol{m}_{\boldsymbol{j}}$

iv. **If** $\left\|\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k)} - \widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k-1)}\right\|_2 < \frac{\sqrt{\varepsilon}}{100}$ **then** $\varepsilon = \frac{\varepsilon}{10}$

v. $\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k-1)} = \widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k)}$

Output: for each j, $\widehat{\boldsymbol{\theta}}_{\boldsymbol{j}}^{(k_e)}$ with k_e the index of the final iteration for each j.

3 Results

3.1 α-stable modelling

Before discussing the reconstruction accuracy, we show some results related to fitting the data with α-stable distribution. As this is directly related to the choice of p, it largely influences the quality of the reconstructions. For this, we considered a true ultrasound image acquired on a phantom using an Ultrasonix RP scanner, with a central frequency of 6.6 MHz. The RF lines were sampled at 20 MHz. We show in Figure 1 the result of fitting an α-stable distribution on the histogram of one RF signal extracted from this image and on the histogram of its 1D Fourier transform (real part).

Among all the parameters of the α-stable distribution, we are particularly interested in α, as it will determine the value of p considered in the reconstruction process. For the results shown in Figure 1, we found α equal to 1.21 when modelling the RF in space (or time) domain, and 0.74 for its Fourier transform. The values were similar for all the RF lines of the US image. This result confirms that the RF line is sparser in the Fourier domain than in the spatial domain, and motivates our choice of doing the reconstruction in the Fourier domain [8, 9, 10].

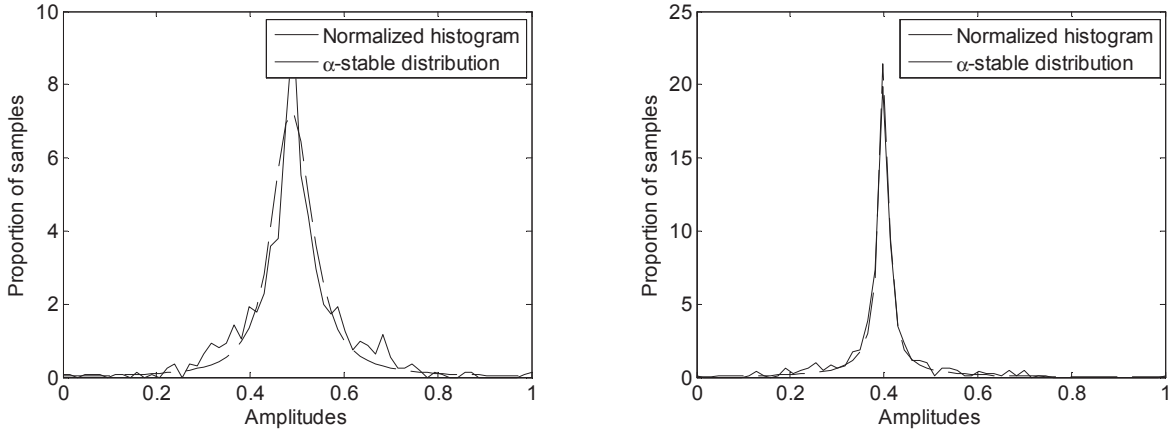

(a) (b)

Figure 1. Example of α-stable modelling of one RF line (a) and of the real part of its 1D Fourier transform (b).

3.2 Reconstruction results

The reconstruction experiments were conducted using two ultrasound images. The first one was acquired on a phantom object using an Ultrasonix RP scanner, with a central frequency of 6.6 MHz and a sampling frequency of 20 MHz. The second image represents an *in vivo* healthy thyroid and was acquired with a Siemens Sonoline Elegra scanner using a 7.5 MHz linear probe and a sampling frequency of 50 MHz.

The proposed approach of solving the problem P_2 was compared to the one presented in [4] and called P_1 herein.

Figures 2, 3 and 4 highlight the reconstruction results obtained with the approach in [4] and with the proposed one. All three results show the qualitative and quantitative improvements provided by the new reconstruction scheme. Moreover, the normalized root mean square errors provided in Table 1 confirm the visual impression given by the B-mode images and the RF plots.

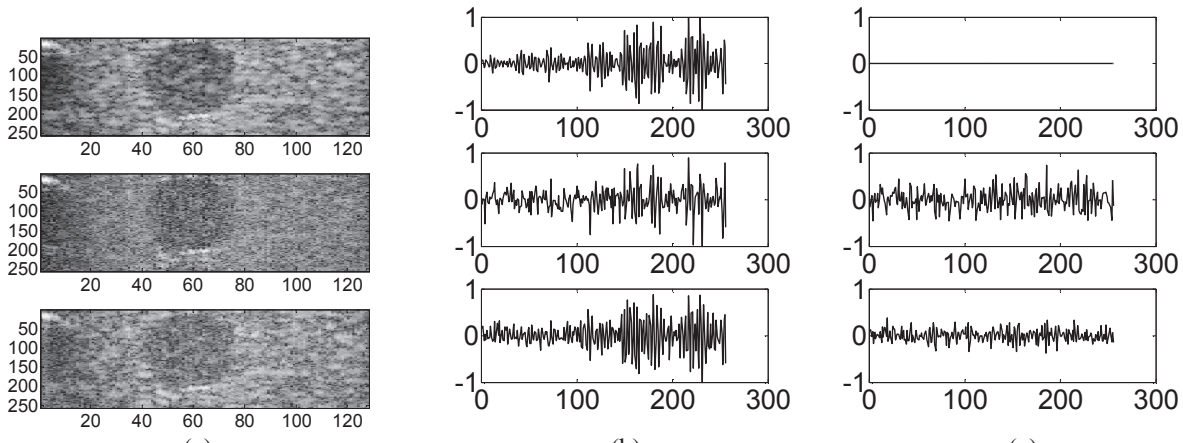

(a) (b) (c)

Figure 2. Reconstruction results for the phantom image, using twice less measurements than the number of initial samples for each RF line. (a) B-mode images corresponding to true and reconstructed RF lines, (b) one RF line, (c) Reconstruction error for one RF line. The first row shows the original data, the second row corresponds to the reconstruction with P_1 and the third row highlights our results (P_2).

	Normalized Root Mean Square Error (NRMSE)		
	Phantom (M=N/2)	*Thyroid (M=N/2)*	*Thyroid (M=N/3)*
Following the scheme P_1	0.68	0.51	0.69
Following the proposed scheme P_2	0.44	0.28	0.54

Table 1. Normalized root mean square error for P_1 and P_2, for the three experiments detailed in the paper. The NRMSEs are calculated using the RF images. N is the number of initial samples and M the number of measurements.

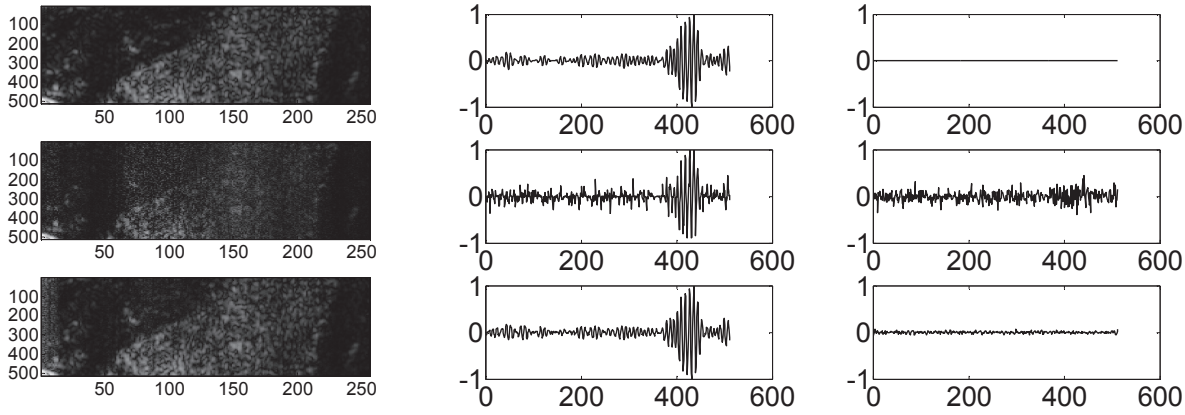

(a) (b) (c)

Figure 3. Reconstruction results for the thyroid image, using twice less measurements than the number of initial samples for each RF line. (a) B-mode images corresponding to true and reconstructed RF lines, (b) one RF line, (c) Reconstruction error for one RF line. The first row shows the original data, the second row corresponds to the reconstruction with P_1 and the third row highlights our results (P_2).

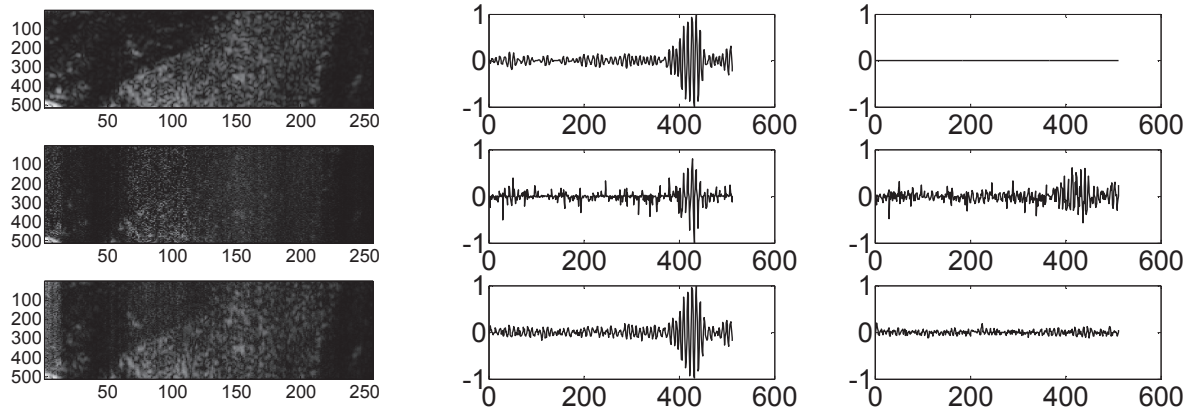

(a) (b) (c)

Figure 4. Reconstruction results for the thyroid image, using three times less measurements than the number of initial samples for each RF line. (a) B-mode images corresponding to true and reconstructed RF lines, (b) one RF line, (c) Reconstruction error for one RF line. The first row shows the original data, the second row corresponds to the reconstruction with P_1 and the third row highlights our results (P_2).

Finally, we show in Figure 5 the reconstruction result obtained with the proposed approach, for an entire RF image containing 256 RF lines. Note that in this case, for memory space requirements imposed by the size of the matrices, the reconstruction was processed blockwise.

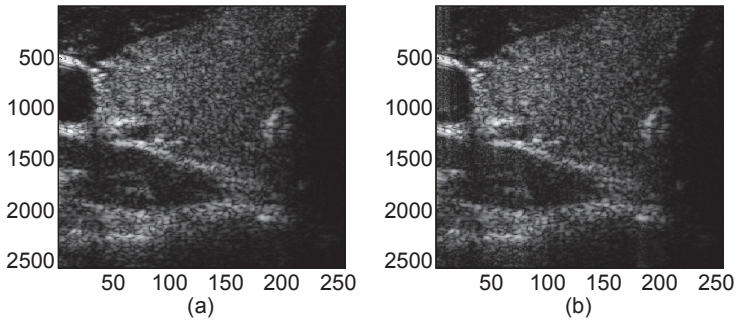

Figure 5. Reconstruction result obtained with the proposed method for the hole thyroid image, using twice less measurements than the number of initial samples for each RF line. (a) B-mode image corresponding to the true RF data, (b) B-mode image calculated with the reconstructed RF data. Both images are plotted with the same dynamic range.

4 Conclusion

In this paper, we proposed a CS-based l_p-norm reconstruction of US images by modelling the Fourier transform of RF signals as α-stable distributions. Moreover, we related the value of p used in the optimization process to α. The results show significant reconstruction improvement compared to the previously published approach where the reconstruction was performed in the spatial domain. More precisely, the normalized root mean square error was shown to be significantly smaller with the proposed approach for two experimental RF images (phantom and in vivo data). The method was evaluated for two amounts of measurements, two and respectively three times lower than the number of samples to be reconstructed.

Acknowledgement

This work was partially supported by ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02 of the University of Toulouse.

References

[1] E.J. Candes, J. Romberg and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” IEEE Transactions on Information Theory, vol. 52, no. 2, pp. 489-509, 2006.
[2] J. Tropp and A. C. Gilbert, “Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit,” IEEE Transactions on Information Theory, vol. 53, no. 12, pp. 4655-4666, 2007.
[3] H. Liebgott, A. Basarab, D. Kouamé, O. Bernard, D. Friboulet, “Compressive sensing in medical ultrasound,” IEEE Ultrasonics Symposium, 6 pages, 2012.
[4] A. Achim, B. Buxton, G. Tzagkarakis, P. Tsakalides, “Compressive sensing for ultrasound RF echoes using α-Stable Distributions,” IEEE EMBC, pp. 4304-4307, 2010.
[5] R. Chartrand, W. Yin, “Iteratively reweighted algorithms for compressive sampling,” IEEE ICASSP, pp. 3869-3872, 2008.
[6] M.A. Kutay, A.P. Petropulu, C.W. Piccoli, "On modeling biomedical ultrasound RF echoes using a power-law shot-noise model," IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, vol.48, no.4, pp. 953-968, 2001.
[7] M. Pereyra and H. Batatia, "Modeling ultrasound echoes in skin tissues using symmetric α-stable processes," IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control, vol. 59, no. 1, pp. 60 - 72, 2012.
[8] C. Quinsac, A. Basarab, D. Kouamé, “Frequency domain compressive sampling for ultrasound imaging,” Advances in Acoustics and Vibration, Special issue on Advances in Acoustic Sensing, Imaging, and Signal Processing, Vol. 2012, 16 pages, 2012.
[9] C. Quinsac, A. Basarab, J-M. Girault, D. Kouamé, “Compressed sensing of ultrasound images: sampling of spatial and frequency domains,” IEEE Workshop on Signal Processing Systems, pp. 231-236, 2010.
[10] N. Dobigeon, A. Basarab, D. Kouamé, J-Y. Tourneret, “Regularized Bayesian compressed sensing in ultrasound imaging,” European Signal and Image Processing Conference, pp. 2600-2604, 2012.