# Compressed sensing based compression of SAR raw data

Gabriel Rilling, Mike Davies and Bernard Mulgrew
{G.Rilling, Mike.Davies, B.Mulgrew}@ed.ac.uk
Institute for Digital Communications (IDCOM) & Joint Research Institute for Signal and Image Processing
School of Engineering, University of Edinburgh
Alexander Graham Bell Building, The King's Buildings, Mayfield Road, Edinburgh EH9 3JL, United Kingdom

***Abstract*—(DRAFT) This work addresses the possible use of compressed sensing to compress SAR raw data. Due to their noise-like features, SAR images are difficult to acquire with compressed sensing techniques. However, some parts of the images are usually compressible and we investigate two techniques exploiting that compressibility to allow a compressive acquisition of the whole image. Such parts are usually associated with man-made structures and can therefore be absent from images containing only natural targets. The two new techniques proposed here result in a siginificant enhancement of the image quality compared to classical compressed sensing. Moreover, compared to classical sampling and quantisation of the SAR raw data, they allow a significant reduction of bitrate with a limited increase of the distortion.**

## I. Introduction

Synthetic aperture radar (SAR) is an active ground imaging system based on coherent processing of multiple radar echoes acquired along the path of a moving platform (aircraft or satellite). Due to the low computational resources of the acquisition platforms and the steadily increasing resolution of SAR systems, the data cannot generally be processed on board and must be stored or transmitted to the ground where the image formation process is performed. The amount of image data produced is now constrained by on board storage capabilities and transmission links.

To address this problem, many techniques have been proposed to compress the raw SAR data [1], [2], [3]. However, SAR systems in practice mostly use the simplest methods because of their low computational requirements. In this context, an appealing idea is to apply results of the rapidly developping field of *compressed sensing* introduced in [4], [5]. Unlike traditional compression/decompression methods, compressed sensing allows very simple non-adaptive compression schemes at the expense of a significantly increased complexity for the decompression. The key idea is to exploit redundancy in the data modelled as *sparsity* in an appropriate dictionary.

In the context of SAR, sparsity has been mostly used for denoising [6] and superresolution [6], [7], [8] with excellent performance. Compressed sensing has also been proposed in [9], [10], [11] with encourageing simulation results. However the only known realistic example provided in [11] has shown the application of compressed sensing to SAR to be particularly challenging in practice. In this contribution, we build on the results of [11] and investigate two simple methods aimed at improving the quality of the recontructed images.

The paper is organized as follows. In Section II we propose a brief overview of SAR data processing and of the properties of the resulting images. The compressed sensing based SAR acquisition and decoding framework is presented in Section III together with the two proposed enhancements. The performance of the methods is then assessed in Section IV.

## II. SAR data and image properties

### A. SAR image formation - nature of SAR raw data

In the "spotlight" mode, SAR data is acquired from a moving platform by emitting at close intervals a bandpass microwave radar signal in direction of a specific area or *scene* and sampling the signal backscattered by the ground objects.

When the size of the scene is small compared to its distance to the radar platform, the curvature of the wavefront of the radar signal over the scene can be neglected. This approximation, illustrated in Fig. 1, results in a simple interpretation of SAR data in the Fourier domain (2D Fourier transform of the scene) and is the base of a SAR processing technique referred to as "polar format algorithm". In this approximation, each received signal only contains information averaged over the scene in a direction orthogonal to the direction of propagation of the emitted radar signal. In the Fourier domain, each received signal thus contains information included in a radial line orthogonal to the averaging direction, i.e. in the direction of propagation. Further analysis shows that the actual information is in fact included in a segment, whose radial position is related to the band of the radar signal by a factor $2/c$, where $c$ is the speed of light. Thus, the whole SAR data approximately correspond to a *polar grid in the Fourier domain*, centered away from the origin at a distance corresponding to twice the mean wavelength of the emitted radar signal. A natural consequence of this bandpass property is that SAR images are complex-valued. In order to form the SAR image, the polar grid data are interpolated to a rectangular grid from which the image is computed by means of an inverse DFT.

In practice, the angular range of the polar grid (corresponding to the angle between the dashed gray lines in Fig. 1 (b)) is often very small, in which case the polar grid can already be well approximated by a rectangular grid. For this reason, we will simply assume in this preliminary work that the raw data correspond to the 2D Fourier transform of the SAR image. The effect of the mapping from the polar grid to the rectangular one is expected to be small and is under current investigation as well the impact of more accurate SAR system models.

### B. Statistical properties

Two properties of SAR images will be important to understand the methods proposed in this paper and their results. The first one is the noise-like characteristics of SAR images. The second one is their often very high dynamic range caused by a few very bright objects.

*1) Noise-like properties:* As the output of a coherent imaging system, SAR images are extremely noisy. This feature comprises two aspects. First, the magnitude of the SAR image contains what is referred to as speckle noise which can be efficiently modelled as multiplicative exponential white noise. Second, the phase of the image can also be modelled as uniform white noise in $[0, 2\pi]$.

These two features have the same origin which is that each pixel of the image corresponds to an area whose dimensions are large compared to the wavelength of the radar signal (typically at least $30\mathrm{cm} \times 30\mathrm{cm}$ for an average radar wavelength of $3\mathrm{cm}$) containing multiple objects. These objects generally have different reflectivities which can be modelled as complex-valued: the magnitude corresponds to the intensity of the reflected radiation while the phase

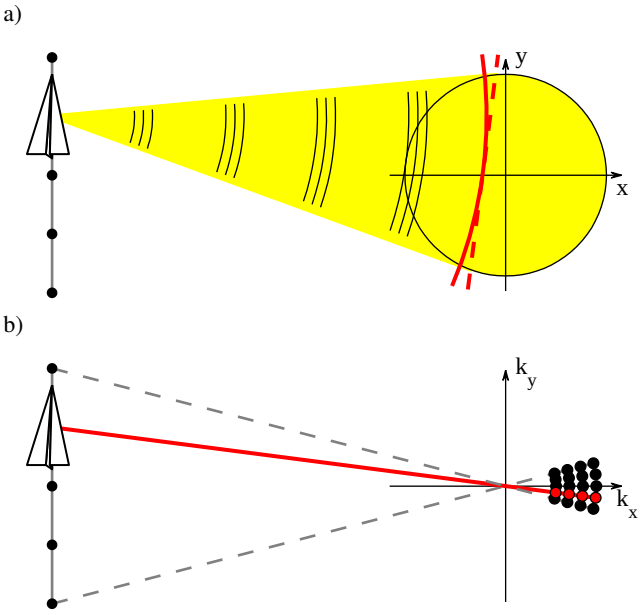


Fig. 1. Geometry of SAR data acquisition. a) in the spatial domain. b) in the Fourier domain. In a) the curved wavefront (red line) is approximated by a straight line (dashed red line).

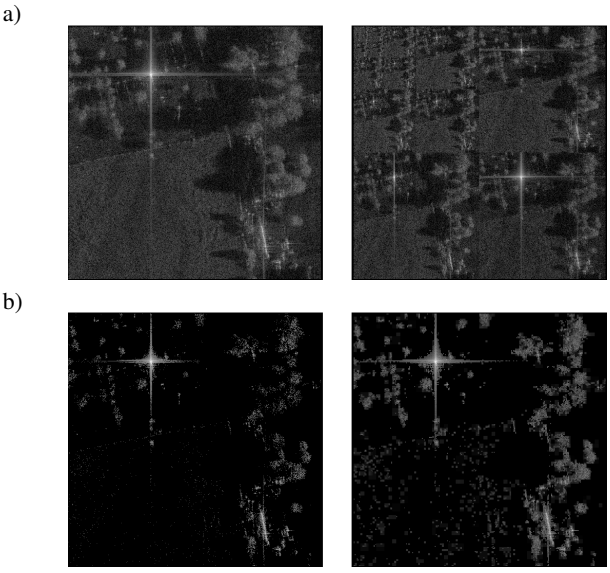


Fig. 2. Compressibility of a SAR image in the spatial and wavelet domains. a) Image and its Haar wavelet transform. b) Compressed images obtained by only keeping the 5% largest coefficients in both domains. All images are in log scale with 70dB dynamics.

corresponds to a phase shift. More importantly, the reflectivities of these objects can typically be modelled as uncorrelated. For this reason, the radiations reflected by them interfere incoherently when added up to obtain the reflectivity of a pixel of the SAR image. The result is that the magnitude of a pixel of a SAR image can take any value in the interval between zero (destructive interference) and the sum of the magnitudes of the subpixel objects (constructive interference). This is the usual explanation for the speckle noise. Since the number of subpixel objects is typically very large, the distribution of the complex-valued reflectivity of a pixel is Gaussian with independent real and imaginary parts both following the same Gaussian distribution. As a consequence the magnitude of a pixel is exponentially distributed and its phase is uniformly distributed in $[0, 2\pi]$. Moreover, the way subpixel objects interfere within a pixel is also independent from the way they interfere for a neighboring pixel. Hence both the multiplicative speckle noise and the phase of the image can be modelled as stationary and white.

Thus, a SAR image $f \in \mathbb{C}^{N\times N}$ can be efficiently modelled at a pixel $(k, l)$ as

$$f_{kl} = \tilde{f}_{kl} s_{kl} e^{i\varphi_{kl}},$$

where $\tilde{f}$, $s$ and $\varphi$ are real-valued images and stand respectively for the SAR image without noise, the speckle noise and the phase. Since both the speckle noise and the phase in this model can be modelled as stationary and white, the global multiplicative noise $se^{i\varphi}$ is a zero-mean stationary white noise. As a consequence, the SAR image $f$ can be modelled as a *zero-mean nonstationary white noise*.

*2) Dynamic range:* In most cases, the radiation emitted by the radar antenna is scattered when hitting the ground and only a very small proportion of the energy is reflected in the direction of the antenna. This typically happens in natural areas without any man-made objects. However, in some cases, a much larger proportion of the energy is reflected towards the antenna. A well-known cause for this phenomenon is the presence of corner shapes which are very common in man-made buildings or vehicles. As a consequence, SAR images containing man-made objects typically have very bright pixels localised on those objects while the background of the image is much darker. Moreover, these bright pixels are usually highly localised: for a building, only a few arretes and corners appear as very bright in the SAR image. In practice, the brightest pixels of a SAR image can typically be $10^3$ times larger than the background pixels.

### C. Compressibility

Due to their noise-like properties, complex-valued SAR images are very difficult to compress efficiently. Considering the previous model $f = \tilde{f}se^{i\varphi}$, the image without noise $\tilde{f}$ part typically has the same good compression properties as most natural images. However the multiplicative noise $se^{i\varphi}$ endows the whole image with a high entropy, thus drastically reducing its compressibility in any dictionary. For this reason, typical sparsifying transforms used in image processing such as wavelet transforms do not result in good sparse approximations for SAR images.

Considering for example a Haar wavelet decomposition (see Fig. 2), we observe that the wavelet coefficients are not concentrated in the coarser scales as for usual compressible natural images. The wavelet decomposition instead looks similar to the decomposition of a white noise image where all wavelet scales are similarly populated[1]. As a consequence, if only the 5% largest wavelet coefficients are kept, most of the details of the original image are lost. Comparatively, keeping the 5% largest pixels of the image results in a greater loss of detail but not as drastically as for usual natural images. In both cases the remaining parts correspond mostly to the brightest objects of the image, which are compressible because they are highly localised.

## III. Compressed sensing based coding of SAR data

### A. Compressed sensing basics

According to compressed sensing theory (see e.g. [12] for a tutorial), a discrete signal or image expressed as a vector $f \in \mathbb{C}^n$ can be exactly reconstructed with a reduced number of samples compared to the Nyquist rate provided that it is sparse in some basis: $f = \Psi x$

[1]Using more sophisticated wavelets may increase slightly the concentration of the large coefficients in the coarser scales but the improvement in compressibility is rather small.

where $\mathbf{\Psi} \in \mathbb{C}^{n\times n}$ is a matrix whose columns are the basis vectors, and $x \in \mathbb{C}^n$ is a vector with a small number of non zero components $k \ll n$.

In a compressed sensing framework, the signal/image is acquired through linear projections: $y = \mathbf{\Phi} f$, where $y \in \mathbb{C}^m$ is the measurements vector and $\mathbf{\Phi} \in \mathbb{C}^{m\times n}$ is referred to as the measurement matrix. Considering the $k$-sparse representation $x$, this results in the measurement equation

$$y = \mathbf{\Phi}\mathbf{\Psi} x.$$

In order to recover a $k$-sparse vector $x$, the number of measurements $m$ must be at least greater than $k$ but can be significantly smaller than the signal/image dimension $k < m \ll n$. While this cannot be achieved with any combination of measurement matrix and basis, it has been shown that several classes of random measurement matrices allow it for any basis with high probability.

Given the measurements $y$, the reconstruction of the sparse vector $x$ can be achieved by searching for the sparsest vector $\tilde{x}$ consistent with the measurements. This is usually referred to as the $\ell_0$ optimization problem " $x = \operatorname{argmin}_{\tilde{x}} \|\tilde{x}\|_0$ subject to $y = \mathbf{\Phi}\mathbf{\Psi}\tilde{x}$ ", where the $\ell_0$ pseudo-norm $\|\cdot\|_0$ corresponds to the number of non zero elements. As it is well known, this is a combinatorial problem which cannot be solved directly in practice. The two most common approaches are therefore to replace it with an $\ell_p$ optimization problem with $0 < p \leq 1$ or to use a greedy algorithm such as Orthogonal Matching Pursuit.

In this simplified overview of the compressed sensing theory, we have only focused on the noiseless acquisition of an exact sparse signal. To be applicable in practice, the theory has also been adapted to the noisy and non exactly sparse cases but the principles remain essentially the same.

### B. Compressed sensing for SAR

In order to define a compressed sensing based acquisition scheme for SAR images, three elements must be specified: a basis where the data are assumed sparse (or close to sparse), a measurement operator and a recontruction algorithm.

*1) Sparse representation:* As shown previously, the statistics of SAR images imply that there is no basis or dictionary where the data can be assumed sparse. For this reason, it seems a priori impossible to acquire with a decent quality a whole SAR image in a compressed sensing framework. However, the very bright objects often related to man-made structures or vehicles are typically sparse in the space domain and slightly sparser in a wavelet domain. The image $f \in \mathbb{C}^n$ can thus be decomposed into two components $f = f_s + f_n$, where $f_s$ corresponds to the sparse bright objects and $f_n$ to the remaining non sparse areas. In this decomposition, the sparse component $f_s$ is typically larger than the other component because the bright objects are often several orders of magnitude brighter than the background of the image, thus compensating for their limited spatial support. If the image is represented in an orthonormal wavelet basis $\mathbf{\Psi}$, this property is preserved, leading to a decomposition $f = \mathbf{\Psi} x_s + \mathbf{\Psi} x_n$ where $x_s$ is sparse and larger than $x_n$. Thus, when bright objects are present, the whole SAR image can be assumed close to sparse in a wavelet basis. In the following, we will consider more specifically a *Haar wavelet basis* because more sophisticated wavelet bases appear to result in non significant improvement.

*2) Measurement operator:* As previously mentioned, SAR raw data can be assimilated to samples of the Fourier transform of the SAR image. Among the classes of generic measurement matrices used for compressed sensing, this naturally calls for the *partial Fourier matrix*[4] class where the measurements $y$ correspond to uniformly randomly selected Fourier coefficients of the SAR image. If $\mathcal{F} \in \mathbb{C}^{n\times n}$ is the matrix representing the 2D DFT operator, we define the measurement matrix $\mathbf{\Phi}^{m\times n}$ as a random subset of $m$ lines of $\mathcal{F}$.

*3) Reconstruction algorithm:* Given the measurements $y = \mathbf{\Phi} f$, our aim is to recover the sparse signal $x_s$ such that

$$y = \mathbf{\Phi}\mathbf{\Psi} x_s + \mathbf{\Phi}\mathbf{\Psi} x_n.$$

In this equation, the second term of the right hand side is expected to be much smaller than the first term and can be treated as noise as far as the reconstruction of the sparse component $x_s$ is concerned. It will be the purpose of the next sections to address the reconstruction of $x_n$. Given the above measurement equation, the reconstruction of a sparse approximation $\hat{x}_s$ is obtained by means of a recent greedy algorithm referred to as *"stagewise weak approximate conjugate gradient pursuit"* [13], [14], [15]. This algorithm is well supported by theoretical analysis and its performance is comparable to state of the art algorithms such as CoSaMP [16]. A significant advantage however is that it can be much faster thanks to an approximate estimation of the least squares estimate. In the context of the application to SAR images investigated in this paper, this algorithm allows for a very efficient computation of the sparse approximation, taking typically one minute for a 1.5 million pixels image on a recent computer (using only one CPU core and a non fully optimized Matlab code).

### C. Possible improvements for SAR data

In order to improve the compressed sensing based recovery of whole SAR images, we propose two simple techniques to enhance the quality of the reconstructed image in the non compressible areas which cannot be well described by a sparse approximation. The first one consists in classical compressed sensing acquisition with an additional postprocessing. The second one is inspired by the so-called "Hybrid Compressed Sensing" proposed in [17].

*1) Postprocessing: compressed sensing as an interpolation in the Fourier domain:* The previously described compressed sensing acquisition and reconstruction of a SAR image results in a sparse approximation of the image roughly corresponding to the brightest objects. Assuming that these objects have been perfectly recovered (i.e. $\hat{x}_s = x_s$), the residual $\hat{y}_n = y - \mathbf{\Phi}\mathbf{\Psi} x_s = \mathbf{\Phi} f_n$ carries information corresponding to the non sparse areas of the image. As $\mathbf{\Phi}$ is a partial Fourier matrix, this information corresponds to the knowledge of some of the Fourier coefficients of $f_n$. Without better assumptions on $f_n$, a trivial solution is to choose the estimate $\hat{f}_n = \mathbf{\Phi}^H y_n$, which corresponds to setting the other Fourier coefficients to zero. The result of this choice is that the final estimate of the SAR image $\hat{f}$ is simply the orthogonal projection of the sparse approximation $\hat{f}_s = \mathbf{\Psi}\hat{x}_s$ on the subspace solution to the linear equation $y = \mathbf{\Phi} f$. As a consequence, the distance between the true image $f$ and the estimate $\hat{f}$ is necessarily reduced, which means that this projection always reduces the mean square error of the reconstruction.

From a global point of view, the whole process of computing a sparse approximation and then project on the solution subspace can also be interpreted as an interpolation in the Fourier domain. Indeed, the known Fourier samples $y$ are kept while the unknown Fourier samples are reconstructed using a sparsity hypothesis.

*2) Hybrid compressed sensing:* The underlying idea of the hybrid compressed sensing method proposed in [17] is that the sparse wavelet approximations of natural images typically have full coarser scales while only the finer scales are effectively sparse. The method consists in first separating the image $f = f_a + f_d$ into an approximation component $f_a$ corresponding to the coarser wavelet scales and a detail component $f_d$ corresponding to the finer scales. Then, the approximation is sampled exhaustively (in the wavelet domain) while the detail is sampled and reconstructed using compressed sensing.

In the case of SAR images, the assumption that the wavelet coefficients corresponding to the coarser scales are typically larger than the finer scales is not valid. However, exhaustively sampling an approximation of the image can still be useful insofar as most parts of the image are difficult to recover using compressive sampling. Thus, using a hybrid compressed sensing scheme guarantees at least a low resolution everywhere while the brightest objects can be acquired with a better resolution using compressed sensing on the detail part. This might make sense for surveillance applications if one is e.g. interested in detecting vehicles while keeping a coarse monitoring of the area.

In order to reduce the computational load on the sensing platform, the hybrid compressed sensing method can be further adapted to avoid the computation of a partial wavelet transform. Indeed, the above approximation $f_a$ is a specific low-pass filtered version of the image $f$ but other low-pass filters may be as good for our purpose. Since the SAR raw data are assimilated to Fourier samples, the simplest low-pass filter to implement is the perfect low-pass filter corresponding to a rectangular subset of the 2D Fourier transform of the image. Thus we define the approximation as a specific rectangular subset of the 2D Fourier samples while the detail corresponds to the remaining Fourier coefficients.

The corresponding measurement matrix for the approximation is referred to as $\mathbf{\Phi}_a$ and results in the measurements $y_a = \mathbf{\Phi}_a$. Since the approximation is exhaustively sampled it can be reconstructed as $f_a = \mathbf{\Phi}_a^H y_a$.

The measurement matrix $\mathbf{\Phi}_d$ for the detail component is not a random partial Fourier matrix anymore: it contains all the Fourier coefficients corresponding to the approximation (whose values are zero for the detail) and random coefficients among the remaining Fourier coefficients. It results in measurements $y_d$ which are used to reconstruct a sparse approximation $\hat{f}_d$ of the detail component using the compressed sensing procedure defined above with $\mathbf{\Phi}_d$ instead of $\mathbf{\Phi}$.

Given $f_a$ and $\hat{f}_d$, the estimate of the SAR image $\hat{f}$ is defined as

$$\hat{f} = \left(1 - \mathbf{\Phi}_a^H \mathbf{\Phi}_a\right) \hat{f}_d + f_a.$$

This definition means that $\hat{f}$ is equal to $\hat{f}_d$ except for the Fourier coefficients corresponding to the approximation which are replaced by their value from $f_a$.

In addition, the postprocessing proposed in the previous paragraph can also be applied to $\hat{f}_d$ before it is combined with the approximation to form the final estimate. The performance of both cases, with and without postprocessing, will be assessed in the simulation section.

### D. Quantisation

In order to provide a complete coding scheme, the next step is to quantise the compressive measurements. The traditional way of coding SAR raw data is to quantise the samples with a Block Adaptive Quantiser (BAQ) [1]. Better quantisers have been proposed in the litterature [2], [3] but the BAQ remains popular because of its simplicity. The latter consists of two steps. First, the raw data are divided into small blocks, which are normalized by their standard deviation. The real and imaginary parts of the data within each block are then quantised independently with scalar normalised Gaussian quantisers.

In our simulations, we use a slightly modified BAQ that seem to perform slightly better in our case. The first modification is to normalize the data by the maximum modulus value of each block instead of the standard deviation. The second one is to use a vector BAQ [2] instead of a standard BAQ, whereby scalar quantisers are replaced by a vector quantiser trained with the statistics of the block normalised data, which are almost Gaussian provided the blocks are not too small.

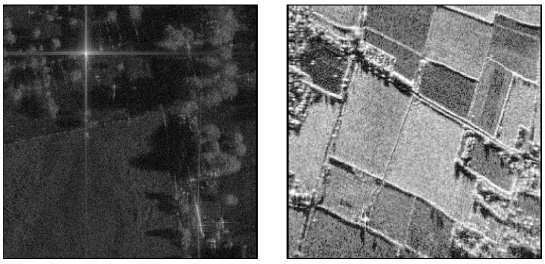

Fig. 3. Test SAR images. The image on the left contains very bright objects and is represented in log scale with 70dB dynamics. The image on the right does not contain any very bright objects and is represented in log scale with 30dB dynamics.

## IV. Simulations

In order to assess the performance of the proposed modifications of the usual compressed sensing framework, we consider two SAR images with and without very bright objects (see Fig. 3). Typical results of the proposed methods are shown in Fig. 4.

In case of moderate subsampling (like 50% missing samples), we observe that standard compressed sensing (second row of Fig. 4) allows the recovery of the main structures of the images, whether they contain very bright objects or not. However the flat areas where the statistics of the images are similar to stationary white noise are badly recovered and are only sparsely filled with small blocks. This behaviour is consistent with the fact that sparse approximations in the wavelet domain badly describe such flat areas. The addition of the proposed postprocessing (third row of Fig. 4) allows the suppression of the blocky artifacts and thus enhances the visual quality of the image. Moreover, this also significantly reduces the mean square error.

For strong subsampling (like 90% missing samples), standard compressed sensing only allows the reconstruction of the very bright objects when such objects are present. In this case postprocessing does not result in any noticeable visual improvement even though the mean square error is still reduced. When no bright objects are present, the compressed sensing reconstructed image looks like a random collection of wavelet blocks with no apparent correlation with the original image.

In the hybrid case, the approximation image is combined with the sparse reconstruction of the detail image. For moderate subsampling, this results in an image with a high resolution on the main structures which benefit from the sparse reconstruction, and a reduced resolution otherwise. Compared to standard compressed sensing with postprocessing, the image looks much closer to the original image and is especially significantly better contrasted. However, the hybrid method is slightly worse in terms of mean square error. This can be compensated by applying the proposed postprocessing in the hybrid case too but the contrast of the image is then reduced without any siginificant improvement of the visual image quality.

Quantitatively, the performance of the proposed methods can be represented as rate distortion curves showing the normalised mean square error between the complex-valued original and recontructed images as a function of the number of bits/pixel (see Fig. 5). For a given number of bits/pixel, multiple combinations of subsampling ratios and numbers of bits per sample are generally possible. The displayed rate distortion curves correspond to the ones leading to the best performance.

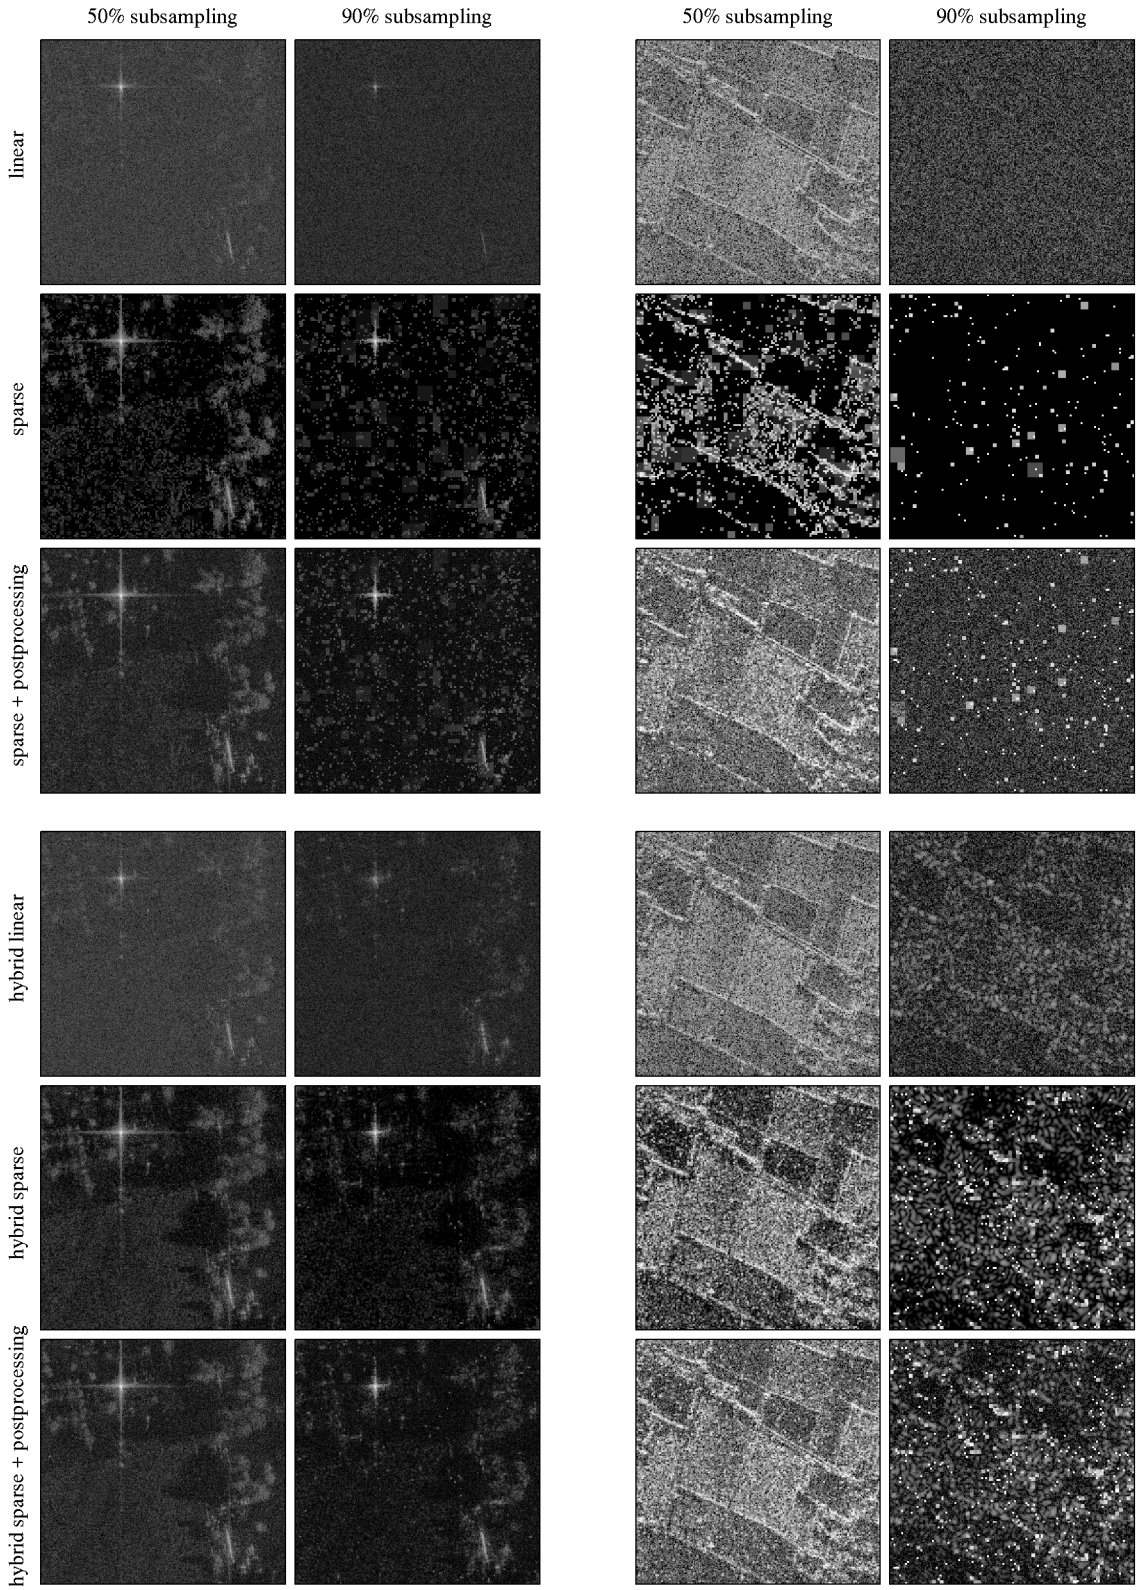


Fig. 4. Examples of reconstructed images using the proposed methods. The four columns correspond to the two test images with different subsampling ratios. The top three rows correspond to subsampling using random Fourier coefficients. The first row shows the result of a linear (or minimum energy) reconstruction where the unknown Fourier samples are assumed equal to zero. The second row shows the sparse approximation obtained using standard compressed sensing. The third row shows the result after postprocessing. The bottom three rows correspond to hybrid subsampling where half of the samples are taken from a square area in the Fourier domain while the other samples are drawn randomly from the remaining Fourier coefficients. Similarly to the first three rows, the fourth row shows the linear reconstruction; the fifth row shows the combination of the approximation and of the sparse reconstruction of the detail; the sixth row shows the result after postprocessing. All images are based on quantised samples using the block adaptive vector quantiser described in paragraph III-D with 6 bits per sample. All sparse reconstructions contain arbitrarily $\approx 0.1m$ nonzero wavelet coefficients, where $m$ is the number of samples.

The results show that compressed sensing based methods can outperform classical Nyquist rate sampling for low bitrates when the image contains very bright objects. Moreover, both the proposed postprocessing and the hybrid compressed sensing investigated in this paper lead to significant performance improvement over standard compressed sensing. However, the mean square error appears to be slightly deceptive here: the combination of the two proposed methods results in significant error reduction whereas the visual quality of the image may appear reduced because of the loss of contrast.

When no bright objects are present in the image, the performance of standard compressed sensing is always worse than Nyquist rate sampling. The cause of this bad performance is simply that the image is far from being sparse, and therefore the sparsity hypothesis is not useful. In this case, the proposed modifications still result in significant improvement over standard compressed sensing but none of them achieves better results than Nyquist rate sampling.

## REFERENCES


[1] R. Kwok and W.T.K. Johnson. Block adaptive quantization of magellan SAR data. *IEEE Transactions on Geoscience and Remote Sensing*, 27(4):375–383, July 1989.

[2] U. Benz, K. Strodl, and A. Moreira. A comparison of several algorithms for sar raw data compression. *Geoscience and Remote Sensing, IEEE Transactions on*, 33(5):1266–1276, Sep 1995.

[3] Vito Pascazio and Gilda Schirinzi. Sar raw data compression by subband coding. *Geoscience and Remote Sensing, IEEE Transactions on*, 41(5):964–976, May 2003.

[4] E.J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, Feb. 2006.

[5] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, April 2006.

[6] Müjdat Çetin. *Feature-enhanced synthetic aperture radar imaging*. PhD thesis, Boston University, 2001.

[7] Ivana Stojanovic, Müjdat Çetin, and William C. Karl. Joint space aspect reconstruction of wide-angle sar exploiting sparsity. In *Algorithms for Synthetic Aperture Radar Imagery XV, Proc. of SPIE*, volume 6970, 2008.

[8] Lee C. Potter, Philip Schnitter, and Justin Ziniel. Sparse reconstruction for RADAR. In *Algorithms for Synthetic Aperture Radar Imagery XV, Proc. of SPIE*, volume 6970, 2008.

[9] Matthew Herman and Thomas Strohmer. Compressed sensing radar. *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 1509–1512, 31 2008-April 4 2008.

[10] Richard Baraniuk and Philippe Steeghs. Compressive radar imaging. *Radar Conference, 2007 IEEE*, pages 128–133, April 2007.

[11] Sujit Bhattacharya, Thomas Blumensath, Bernard Mulgrew, and Mike Davies. Fast encoding of synthetic aperture radar raw data using compressed sensing. In *IEEE Workshop on Statistical Signal Processing, Madison, USA, August 2007*, Madison, USA, August 2007.

[12] E.J. Candes and M.B. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, March 2008.

[13] Mike E. Davies and Thomas Blumensath. Faster & greedier: algorithms for sparse reconstruction of large datasets. In *third IEEE International Symposium on Communications Control, and Signal Processing*, St Julians, Malta, March 2008. (invited paper).

[14] Thomas Blumensath and Mike E. Davies. Stagewise weak gradient pursuits. Part I: Fundamentals and numerical studies. *IEEE Transactions on Signal Processing*, :, . (submitted).

[15] Thomas Blumensath and Mike E. Davies. Stagewise weak gradient pursuits. Part II: Theoretical properties. *IEEE Transactions on Signal Processing*, :, . (submitted).

[16] Deanna Needell and Joel A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, :, 2008.

[17] Yaakov Tsaig and David L. Donoho. Extensions of compressed sensing. *Signal Processing*, 86(3):549–571, March 2006.


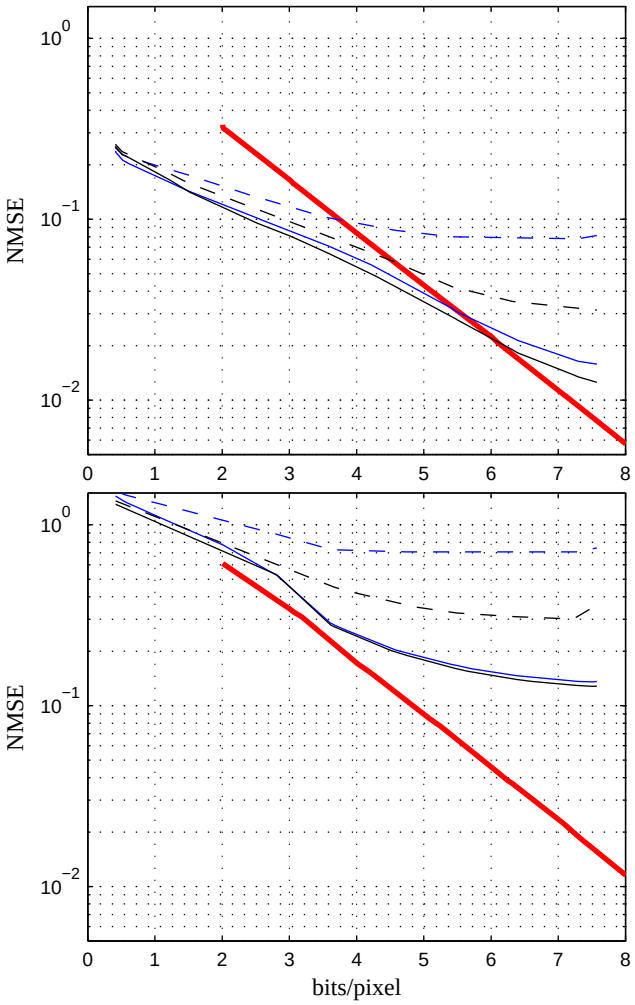


Fig. 5. Optimal rate-distortion curves for the proposed compression methods. The top diagram corresponds to the test image with very bright objects. The bottom diagram corresponds to the test image without bright objects. The thick red line corresponds to exhaustive sampling. The performance of standard compressed sensing is represented as a dashed blue line. The same with post-processing is represented as a solid blue line. The performance of the hybrid method is represented as black lines: dashed line without postprocessing and solid line with postprocessing. All these curves are obtained as convex hulls of the scatter plots corresponding to various subsampling rates (from 10% to 90%), number of nonzero coefficients in the sparse approximations, block sizes, bits/sample and bits/maximum value inside each block. In the case of hybrid sampling, the acquired samples are split equally between the detail and the approximation components.