COMPRESSED SENSING IMAGE RECONSTRUCTION VIA RECURSIVE SPATIALLY ADAPTIVE FILTERING

Karen Egiazarian, Alessandro Foi, and Vladimir Katkovnik

Institute of Signal Processing, Tampere University of Technology
P.O. Box 553, 33101 Tampere, Finland
firstname.lastname@tut.fi

ABSTRACT

We introduce a new approach to image reconstruction from highly incomplete data. The available data are assumed to be a small collection of spectral coefficients of an arbitrary linear transform. This reconstruction problem is the subject of intensive study in the recent field of "compressed sensing" (also known as "compressive sampling"). Our approach is based on a quite specific recursive filtering procedure. At every iteration the algorithm is excited by injection of random noise in the unobserved portion of the spectrum and a spatially adaptive image denoising filter, working in the image domain, is exploited to attenuate the noise and reveal new features and details out of the incomplete and degraded observations. This recursive algorithm can be interpreted as a special type of the Robbins-Monro stochastic approximation procedure with regularization enabled by a spatially adaptive filter. Overall, we replace the conventional parametric modeling used in CS by a nonparametric one.

We illustrate the effectiveness of the proposed approach for two important inverse problems from computerized tomography: Radon inversion from sparse projections and limited-angle tomography. In particular we show that the algorithm allows to achieve *exact reconstruction* of synthetic phantom data even from a very small number projections. The accuracy of our reconstruction is in line with the best results in the compressed sensing field.

Index Terms— compressed sensing, sparsity, inverse problems, Radon transform, limited-angle tomography.

1. INTRODUCTION

During the last three years, compressed sensing (CS) has received growing attention, mainly motivated by the positive theoretical and experimental results shown in [1], [2], [5], [7], [11], [12], [13]. The basic settings of signal reconstruction under conditions of CS are as follows. An unknown signal of interest is observed (sensed) through a limited number linear functionals. These observations can be considered as an incomplete portion of the spectrum of the signal with respect to a given linear transform $\mathcal{T}$. Thus, conventional linear reconstruction/synthesis (e.g., inverse transform) cannot in general reconstruct the signal. For example, when $\mathcal{T}$ is the Fourier transform, CS considers the case where the available spectrum is much smaller than what is required according to the Nyquist-Shannon sampling theory. It is generally assumed that the signal can be represented sparsely with respect to a different relevant basis (e.g., wavelets) or that, alternatively, it belongs to a specific class of functions (e.g., piecewise constant functions). In the publications cited above, it is shown that under such assumptions, stable reconstruction of the unknown signal is possible and that in some cases the reconstruction can be exact. These techniques typically rely on convex optimization with a penalty expressed by the ℓ_0 or ℓ_1 norm [15] which is exploited to enable the assumed sparsity [6]. It results in parametric modeling of the solution and in problems that are then solved by mathematical programming algorithms.

In this work, we propose to replace the traditional parametric modeling used in CS by a nonparametric one. The nonparametric modeling is implemented by the use of spatially adaptive filters. Overall, this represents a new and alternative paradigm for compressed sensing signal reconstruction. The logic behind of our approach is as follows. The regularization imposed by the ℓ_0 or ℓ_1 norms (or by more general criteria) is essentially only a tool for design of some nonlinear filtering. Let us replace this implicit regularization by explicit filtering, exploiting spatially adaptive filters sensitive to image features and details. If these filters are properly designed we have reasonable hopes to achieve better results than it can be achieved by the formal approach based of formulation of imaging as the variational problem with imposed global constraints. In imaging, the regularizations with *global* sparsity penalties (such as ℓ_p norms in some domain) often results in *inefficient* filtering. It is known (e.g., [14]) that a *higher* quality can be achieved when the regularization criteria are *local* and *adaptive*. This is demonstrated in particular in the context of image denoising, where the performance of advanced spatially adaptive (both local and non-local) methods significantly overcomes that of the traditional approaches (e.g., [4] and references therein).

Our approach to CS signal reconstruction is realized by a recursive algorithm based on spatially adaptive image denoising. At every iteration the algorithm is excited by injection of random noise in the unobserved portion of the spectrum. The denoising filter working in the image domain attenuates the noise and reveals new features and details out of the incomplete and degraded observations. Roughly speaking, we seek for the solution (reconstructed signal) by stochastic approximations whose search direction is driven by the denoising filter.

We demonstrate the viability of our proposed approach and illustrate its effectiveness for two important inverse problems from computerized tomography: Radon inversion from sparse projections and limited-angle tomography. The former problem has been used by many authors as a benchmark for testing CS reconstruction algorithms. In particular we show that our algorithm allows to achieve *exact reconstruction* of synthetic phantom data even from a very limited number projections. An example of image reconstruction from low-frequency data is also given.

This work was partially supported by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence program 2006 - 2011).

Recursive system

$$\begin{cases} \hat{y}_2^{(0)} = 0, \quad \texttt{(initialization)} & k = 0, \\ \hat{y}_2^{(k)} = \hat{y}_2^{(k-1)} - \gamma_k \left[\hat{y}_2^{(k-1)} - (1-S) .\!* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}\left(y_1 + \hat{y}_2^{(k-1)} \right)\right)\right) + (1-S) .\!* \eta_k \right], & k \geq 1. \end{cases} \tag{1}$$

2. ALGORITHM

2.1. Observation model and notation

Let θ and $y = \mathcal{T}\{\theta\}$ be, respectively, the image intensity and its $2D$ transform. If all elements of the spectrum y are given then the signal can be recovered by inverting the transform, $\theta = \mathcal{T}^{-1}\{y\}$. However, in CS problems only a small portion of the spectrum is available, which makes the reconstruction of θ an ill-posed problem.

Introduce a sampling operator as the characteristic function $S = \chi_\Omega$ (with values 0 or 1) of the available portion Ω of the spectrum. Thus, the pointwise products $S .\!* y$ and $(1-S) .\!* y$ produce a decomposition of the spectrum in two complementary parts $y_1 = S .\!* y,\ y_2 = (1-S) .\!* y,$ with the equation

$$y = y_1 + y_2 = S .\!* y + (1-S) .\!* y.$$

Here, y_1 and y_2 are the observed (known) and the unobserved (unknown) part of y, respectively. The goal is reconstruct θ (or equivalently y_2) from the available data y_1.

2.2. Recursive system

Given an estimate $\hat{y}_2^{(k)}$ of y_2, we define the estimate $\hat{y}^{(k)}$ of y as

$$\hat{y}^{(k)} = y_1 + \hat{y}_2^{(k)},$$

where the superscripts denote the corresponding iteration. With this notation, our algorithm is defined by the recursive system (1) displayed at the top of this page. The flowchart of the system is shown in Figure 1. The system is initialized by setting $\hat{y}_2^{(0)} = 0$. Then, each iteration ($k \geq 1$) comprises of the following steps:

- *Image-domain estimate filtering*. We filter the reconstructed image estimate $\mathcal{T}^{-1}\{\hat{y}^{(k-1)}\}$ with a spatially adaptive filter Φ in the image domain

$$\Phi\left(\mathcal{T}^{-1}\left(y_1 + \hat{y}_2^{(k-1)}\right)\right) = \Phi\left(\mathcal{T}^{-1}\left(\hat{y}^{(k-1)}\right)\right). \tag{2}$$

  The filter Φ specifically takes into consideration the pointwise smoothness of the image estimate. In this way the elements of the image appeared from the spectrum $\hat{y}_2^{(k-1)}$ are recreated from the given y_1. According to the discussion in the introduction, Φ should be a filter with good adaptivity properties which is appropriate for the considered class of signals. In this work we consider a spatially adaptive denoising filter such as [4].

- *Noise addition (excitation)*. Some pseudo-random noise η_k is introduced in the unobserved portion of the spectrum. The additive noise $(1-S) .\!* \eta_k$ works as a random generator of the missing components in the spectrum. During the subsequent iterations, these components are attenuated or enhanced by the action of the filter Φ, depending to what extent they agree with the image features enabled by the observed spectrum y_1.

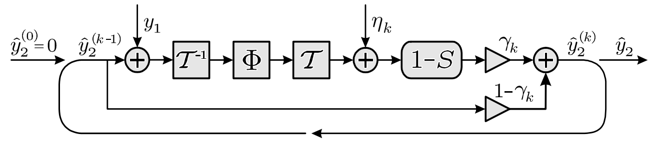

Fig. 1. Flowchart of the recursive system (1).

- *The update* of the estimate $\hat{y}_2^{(k)}$ uses a difference between the estimate $\hat{y}_2^{(k-1)}$ and the prediction of this spectrum obtained after filtering and noise addition. Note that both terms are projected onto the unknown portion of the spectrum by multiplication against $(1-S)$. The factor γ_k scales the prediction difference and controls the rate of evolution (step size) of the algorithm.

2.3. Comments

Stochastic approximation. The recursive algorithm (1) can be treated as the Robbins-Monro stochastic approximation procedure (see, e.g., [9]). If the step size parameter γ_k satisfies the standard conditions

$$\gamma_k > 0, \qquad \sum \gamma_k = \infty, \qquad \sum \gamma_k^2 < \infty \tag{3}$$

and some assumptions on the operator $(1-S) .\!* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}(\cdot)\right)\right)$ hold, as $k \to \infty$ the estimates $\hat{y}_2^{(k)}$ from the recursive system (1) converge in mean squared sense to a solution $\hat{y}_2$ of the equation

$$E\left\{\hat{y}_2 - (1-S) .\!* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right)\right) + (1-S) .\!* \eta_k\right\} = 0,$$

i.e.

$$\hat{y}_2 = (1-S) .\!* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right)\right). \tag{4}$$

If there is no smoothing in the filter Φ, the equation (4) becomes the identity $\hat{y}_2 = (1-S) .\!* \mathcal{T}\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right) = (1-S) .\!* (y_1 + \hat{y}_2) = \hat{y}_2$. Thus, any $\hat{y}_2$ satisfies to the equation (4), there is no image reconstruction and the algorithm does not work. Therefore, in order for the solution $\hat{y}_2$ to be non-trivial, the adaptive smoothing in (4) should be strong enough.

Its strength can be controlled by evaluating the smoothing effects translated in y_1. Indeed, if $\hat{y}_2$ is given, the equation $\hat{y}_1 = S .\!* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right)\right)$ defines the corresponding smoothed spectrum $\hat{y}_1$. The difference between y_1 and $\hat{y}_1$ can be used as an objective measure of the correspondence between the filtered data and the given data at hand.

Excitation noise. The additive noise η_k used in the procedure (1) does not influence the equation (4). There are two arguments in favor of excitation of the algorithm by the random noise. First of all it improves the performance of the algorithm. It accelerates the transition process of the recursive procedure bringing it fast in the area of solution where the random walks steadies. The amplitude of these random walks decreases together with γ_k. It is well known (e.g., [8]) that the random search applied in optimization problems results in random walks well concentrated in areas of global extremum. Thus, the random search imposed by random excitation of the search trajectory can be useful for separation of local and global extrema. In a similar way, if the equation (4) has more than one solution, the randomness can help to find a "strong" solution with better quality of imaging or lower values of some hypothetical criterion function where the gradient (or quasi-gradient) can be defined as the vector corresponding to $\hat{y}_2 - (1 - S) .* \mathcal{T}\left(\Phi\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right)\right)$. Further, by varying the variance of the additive noise η_k one can control the level of smoothing in the recursive procedure and hence the rate of evolution of the algorithm. Thus, in practice, the assumptions (3) can be relaxed and a fixed γ_k can be used provided that var $\{\eta_k\} \underset{k\to\infty}{\rightarrow} 0$.

Stopping rule. The algorithm can be stopped when the estimates $\hat{y}_2^{(k)}$ approach numerical convergence or after a specified number of iterations. We do not excite the algorithm with noise at the final iteration k_{final} and set $\eta_{k_{\text{final}}} \equiv 0$.

Image estimates. An image estimate $\hat{\theta}^{(k)}$ can be obtained at the kth iteration as $\mathcal{T}^{-1}\left(\hat{y}^{(k)}\right) = \mathcal{T}^{-1}\left(y_1 + \hat{y}_2^{(k)}\right)$, although in practice $\Phi\left(\mathcal{T}^{-1}(y_1 + \hat{y}_2)\right)$ or $\mathcal{T}^{-1}\left(y_1 + \hat{y}_2^{(k)} + \gamma_k (1 - S) .* \eta_{k-1}\right)$ are better estimates because of the absence of the excitation noise. All these estimates converge to $\mathcal{T}^{-1}(y_1 + \hat{y}_2)$ as $k \to \infty$.

3. EXPERIMENTS

The following experiments are carried out using a simplified form of the recursion (1), where $\gamma_k \equiv 1$ and η_k is independent Gaussian noise with exponentially decreasing variance var $\{\eta_k\} = \alpha^{-k-\beta}$. For the filter Φ, we use the block-matching and 3D filtering algorithm (BM3D) [3, 4]. The separable 3D Haar wavelet decomposition is adopted as the transform utilized internally by the BM3D algorithm.

We begin with illustrative inverse problems of compressed sensing for computerized tomography. As in [1], we simulate the Radon projections by "approximately" radial lines in the rectangular FFT domain. Note that the initial image estimate $\hat{\theta}^{(0)} = \mathcal{T}^{-1}\left(y_1 + \hat{y}_2^{(0)}\right) = \mathcal{T}^{-1}(y_1)$ always coincides with the conventional back-projection estimate (minimum ℓ_2-norm estimate).

3.1. Radon inversion from sparse projections

First, we reproduce exactly the same experimental setup from [1], where 22 radial lines are sampled from the FFT spectrum of the Shepp-Logan phantom (size 256×256 pixels), as shown in Figure 2(left). Further, we reduce the number of available Radon projections from 22 to 11 (see Figure 2(center)). The initial back-projection estimates are shown in Figure 3. As the recursive algorithm progresses, the reconstruction error improves steadily until numerical convergence, as it can be seen from the plots in Figure 4. For both cases the reconstruction is *exact*, in the sense that the final reconstruction error (PSNR ≃270dB) is comparable with the numerical precision of this particular implementation of the algorithm (double precision floating-point). We remark however that in practice such a high accuracy is never needed: already at a PSNR of about 60dB the image estimates can hardly be distinguished from the original.

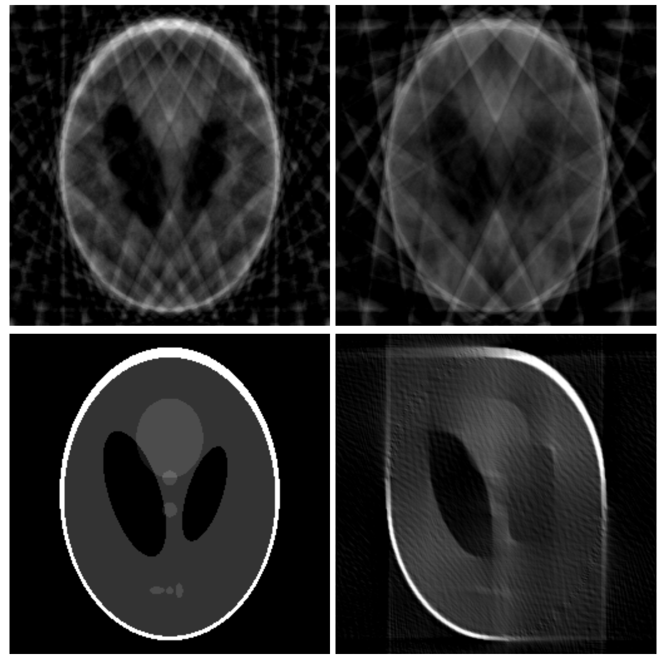

Fig. 3. Clockwise from top-left: back-projection estimates for 22 radial lines, 11 radial lines, 61 radial lines with limited-angle (90 degrees), and original phantom θ (unknown and shown here only as a reference). For all three experiments, the estimates obtained after convergence of the algorithm coincide with the original image.

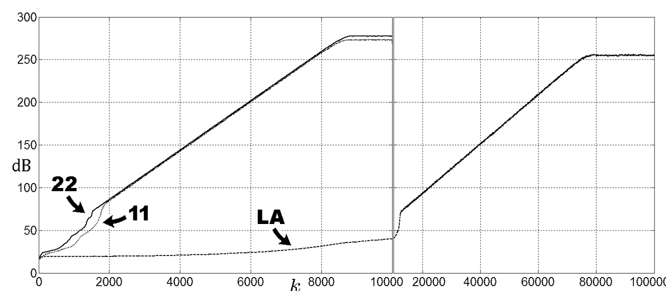

Fig. 4. Progression of the PSNR (dB) of the reconstructed image estimate $\hat{\theta}^{(k)}$ with respect to the iteration count k for the three experiments: 22 and 11 sparse projections ("22" and "11") and limited-angle ("LA").

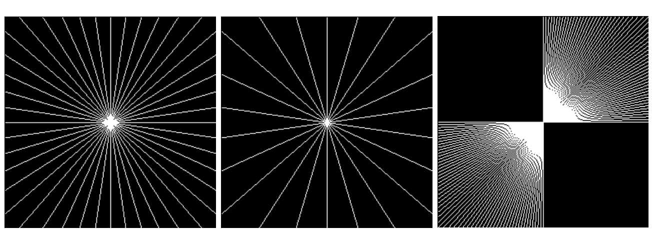

Fig. 2. Sample domain Ω for the FFT spectrum (i.e., the sampling function $S = \chi_\Omega$) for the three experiments: 22 radial lines, 11 radial lines, 90 degrees limited-angle with 61 radial lines.

3.2. Limited-angle tomography

In the two previous experiments, the available Radon projections were uniformly distributed with respect to the projection angle. A more difficult case arises when the angles under which the projections are taken are limited. Similarly to [10], we consider an overall aperture for the projections of 90 degrees. This restriction is essential, since all frequency information is completely missing along half of the orientations, which makes the reconstruction of, e.g., edges across these orientation extremely hard. We complicate the problem further, by taking only a smaller subset of 61 projections (a total of 256 properly-oriented projections would be required to cover a 90 degrees aperture). These sparse, limited-angle projections are illustrated in Figure 2(right). Although the convergence is here much slower than in the previous two experiments, the algorithm eventually achieves exact reconstruction.

In the above three experiments, as soon as the estimate reaches a quality of about 70dB, the recursion enters a phase of improvement at a constant rate (linear in terms of PSNR since var $\{\eta_k\}$ decreases exponentially) which appears to be limited only by the used arithmetic precision.

3.3. Reconstruction from low-frequency data

The proposed recursive procedure can be applied also to more conventional image-processing problems. As a final example, we present the result of reconstruction of a non-synthetic test image, namely *Cameraman* (256×256 pixels), from the low-frequency portion of its Fourier spectrum. In particular, the set Ω is a 128×128 square centered at the DC. In Figure 5 we show the initial estimate $\hat{\theta}^{(0)}$ (by zero-padding in FFT domain, thus minimum ℓ_2-norm) and the reconstructed image obtained after few iterations of the algorithm. Despite the reconstruction is not exact, the salient details of the image are properly restored and there are no significant artifacts (e.g., ringing) thanks to the spatial adaptivity embedded in our procedure.

Further illustrations and animations showing the evolution of the recursive algorithm can be found on our research project's website at http://www.cs.tut.fi/~comsens .

Fig. 5. *Cameraman*: initial estimate $\hat{y}^{(0)}$ (zero-padding) (PSNR=27.32dB), and reconstructed estimate $\hat{y}^{(62)}$ after 62 iterations (PSNR=29.10dB).

4. CONCLUSIONS AND FUTURE WORK

The proposed approach is rather different from the other cited techniques. The spatially adaptive filter replaces the global penalties based on ℓ_p-norms. It means that traditional parametric modeling for the unknown signal is replaced by a nonparametric one. We argue that this modeling provides a more flexible and effective form of regularization for approximating the sought solution. The introduction of noise at every iteration and the consequent stochastic nature of our recursion are in sharp contrast with the deterministic iterations used by the standard mathematical programming applied for CS. On the other hand, the reader may find similarities between our approach and other random search techniques such as simulated annealing. We note that this kind of stochastic algorithms have been shown to be very effective alternatives to conventional deterministic approaches also for the solution of a number of different large-scale optimization problems (e.g., traveling salesman). Our results confirm this observation for the case of CS reconstruction. In particular, in the experiments we demonstrated that our procedure can match and also overcome the performance of sophisticated techniques such as [1].

The algorithm and results shown in this paper are mostly aiming at suggesting an alternative and innovative approach to CS problems. A lot of work remains to be done. Sufficiency conditions which can guarantee the asymptotic perfect reconstruction have not yet been established. We believe that these conditions shall be based on the sparsity of the signal with respect to some (overcomplete) basis (e.g., the particular local transforms which are used within the spatially adaptive filter Φ). We are currently working on accelerating convergence of the recursive algorithm, enabling adaptive feedback between the action of the filter Φ and the excitation noise η_k, aiming at the reduction of the overall complexity.

Applications other than those shown in this paper have already been considered and have not been included here because of length limitation. They will be presented on our project's website and in forthcoming publications.

5. REFERENCES

[1] Candès, E., and J. Romberg, "Practical signal recovery from random projections", *IEEE Trans. Signal Process.*, submitted, 2005.

[2] Candès, E., 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, 2006.

[3] Dabov, K., A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising with block-matching and 3D filtering", *Proc. SPIE El. Imaging 2006, Image Process.: Algorithms and Systems V*, 6064A-30, San Jose, California USA, 2006.

[4] Dabov, K., A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering", *IEEE Trans. Image Process.*, 2007 (in press).

[5] Donoho, D.L., "Compressed sensing", *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289-1306, April 2006.

[6] Donoho, D.L., and M. Elad, "Maximal sparsity representation via l_1 minimization", *Proc. Nat. Aca. Sci.*, vol. 100, pp. 2197-2202, 2003.

[7] Gribonval, R., and M. Nielsen, "Sparse representation in unions of bases", *IEEE Trans. Inf. Theory*, vol. 49, no. 12, pp. 3320-3325, 2003.

[8] Katkovnik, V., and O.Yu. Kulchitskii, "Parametric Statistical Estimates for Multiextremum Optimization Problems", *Engineering Cybernetics*, no. 6, 1975.

[9] Kushner, H., G.G. Yin, *Stochastic approximation and recursive algorithms and applications*, 2nd ed., Springer, 2003.

[10] Olson, T., "Stabilized inversion for limited angle tomography", *IEEE Eng. Med. Biol. Mag.*, vol. 14, no. 5, pp. 612-620, 1995.

[11] Tropp, J.A., "Just relax: Convex programming methods for subset selection and sparse approximation", *IEEE Trans. Inf. Theory*, vol. 51, no. 3, pp. 1030-1051, March 2006.

[12] Tropp, J.A., and A.C. Gilbert, "Signal recovery from partial information via Orthogonal Matching Pursuit", *IEEE Trans. Inf. Theory*, submitted, 2005.

[13] Tsaig, Y., and D.L. Donoho, "Extensions of compressed sensing", *Signal Process.*, vol. 86, no. 3, pp. 549-571, March 2006.

[14] Vanzella, W., F.A. Pellegrino, V. Torre, "Self-adaptive regularization", *IEEE Trans. Pattern Anal. Mach. Intelligence*, vol. 26, no. 6, pp. 804-809, 2004.

[15] Wainwright, M.J., "Sharp thresholds for high-dimensional and noisy recovery of sparsity", *Proc. 44th Allerton Conf. Comm., Control Computing*, Monticello, IL, September 2006.