

COMPRESSIVE SENSING OF A SUPERPOSITION OF PULSES

Chinmay Hegde and Richard G. Baraniuk

ECE Department, Rice University

ABSTRACT

Compressive Sensing (CS) has emerged as a potentially viable technique for the efficient acquisition of high-resolution signals and images that have a sparse representation in a fixed basis. The number of linear measurements M required for robust polynomial time recovery of S -sparse signals of length N can be shown to be proportional to $S \log N$. However, in many real-life imaging applications, the original S -sparse image may be blurred by an unknown point spread function defined over a domain Ω ; this multiplies the apparent sparsity of the image, as well as the corresponding acquisition cost, by a factor of $|\Omega|$. In this paper, we propose a new CS recovery algorithm for such images that can be modeled as a sparse superposition of pulses. Our method can be used to infer both the shape of the two-dimensional pulse and the locations and amplitudes of the pulses. Our main theoretical result shows that our reconstruction method requires merely $M = O(S + |\Omega|)$ linear measurements, so that M is sublinear in the overall image sparsity $S|\Omega|$. Experiments with real world data demonstrate that our method provides considerable gains over standard state-of-the-art compressive sensing techniques in terms of numbers of measurements required for stable recovery.

Index Terms— Compressive sensing, blind deconvolution, sparse approximation

1. INTRODUCTION

Compressive Sensing (CS) [1, 2] is an alternate framework to the traditional Shannon/Nyquist framework of digital signal and image acquisition. CS can be viewed as a scheme for simultaneous sensing and compression; instead of being proportional to the Fourier bandwidth, the rate of data acquisition need only be proportional to the *sparsity* of the signal, i.e., the number of nonzero coefficients of a signal representation in some basis. Thus, if the sparsity of the target signal/image is known to be much smaller than its ambient dimension, then CS can potentially lead to considerable savings in data acquisition and transmission costs.

Nevertheless, in many real-world sensing applications, the assumption of exact sparsity is an oversimplification. For example, CS has been proposed as an effective method for astronomical imaging [3]. A perfectly sharp high-resolution

image of the night sky would consist of a sparse field of points (corresponding to the locations of the stars). Instead, the image acquired by an ordinary telescope comprises a sparse *superposition of pulses*, owing to atmospheric effects and non-idealities in the imaging apparatus. This translates to an apparent *increase in sparsity* of the underlying image and consequently has an adverse effect on the performance of any state-of-the-art CS imaging system. Similar blurring effects are observed in various other applications, such as radio interferometry, synthetic aperture radar (SAR), and sonar.

In this paper, we develop and analyze a CS framework for the compressive acquisition of sparse images contaminated by a small amount of unknown blur. We motivate a deterministic model for the image classes of interest; we derive a bound on the number of linear measurements required to encode the essential information contained in this model; and finally, we develop an iterative algorithm that estimates both the two-dimensional (2D) locations and amplitudes of the sparse nonzeros of the image, as well as the 2D-blurring function, using far fewer measurements than conventional CS methods. Our approach can be linked to various concepts in the literature, including dictionary learning [4] and blind deconvolution [5]. Our work is an extension of our recent theory and methods dealing with the compressive sensing of streams of 1D pulses [6].

The paper is organized as follows. Section 2 provides a brief review of compressive sensing, as well as deconvolution techniques used in imaging. Section 3 introduces our proposed signal model and describes our main theoretical result and algorithm. Section 4 illustrates the advantages of our method with example reconstructions of a real world image. Section 5 concludes with discussion and extensions.

2. PRELIMINARIES

A signal $x \in \mathbb{R}^N$ is termed as K -sparse in the orthonormal basis Ψ if the corresponding basis expansion $\alpha = \Psi^T x$ contains no more than K nonzero elements. In the sequel, unless otherwise noted, the sparsity basis Ψ is assumed to be the identity matrix for \mathbb{R}^N . Denote the set of all K -sparse signals in \mathbb{R}^N as Σ_K . In terms of geometry, Σ_K can be identified as the union of $\binom{N}{K}$, K -dimensional subspaces of \mathbb{R}^N , with each subspace being the linear span of exactly K canonical unit vectors of \mathbb{R}^N .

Email: {chinmay,richb}@rice.edu. Web: dsp.rice.edu. This work was supported by the grants NSF CCF-0431150 and CCF-0728867, DARPA/ONR N66001-08-1-2065, ONR N00014-07-1-0936 and N00014-08-1-1112, AFOSR FA9550-07-1-0301, ARO MURIS W911NF-07-1-0185 and W911NF-09-1-0383, and the TI Leadership University Program.

2.1. Compressive Sensing

Suppose instead of collecting all the coefficients of a vector $x \in \mathbb{R}^N$, we merely record M inner products (measurements) of x with $M < N$ pre-selected vectors; this can be represented in terms of a linear transformation $y = \Phi x$, $\Phi \in \mathbb{R}^{M \times N}$. The central tenet of CS is that x can be *exactly* recovered from y , even though Φ is low-rank and has a nontrivial nullspace. In particular, a condition on Φ known as the *restricted isometry property* (RIP) can be defined thus [7]:

Definition 1 An $M \times N$ matrix Φ has the K -RIP with constant δ_K if, for all $x \in \Sigma_K$,

$$(1 - \delta_K)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 + \delta_K)\|x\|_2^2. \quad (1)$$

A matrix Φ with the K -RIP essentially ensures a stable metric embedding of the set of *all* K -sparse signals into a space of dimension M . An important hallmark of CS is that matrices whose elements are chosen as i.i.d. samples from a random subgaussian distribution [8] satisfy the RIP with high probability, provided $M \geq \mathcal{O}(K \log(N/K))$. Thus, M is linear in the sparsity of the signal set K and *only logarithmic* in the ambient dimension N .

CS signal recovery aims to perform stable, feasible inversion of the operator Φ onto its domain Σ_K . A multitude of CS recovery algorithms exist in the literature; recently, iterative support selection algorithms (eg., CoSaMP [9]) have emerged that offer uniform, stable guarantees while remaining computationally efficient. The RIP plays a crucial role in CS recovery; it can be shown that if Φ possesses the $3K$ -RIP, then state-of-the-art recovery algorithms like CoSaMP can recover *any* K -sparse signal x , given measurements $y = \Phi x$.

2.2. Applications of CS in 2D-imaging

The CS approach can be used in a variety of imaging applications. As a consequence of the theory described above, CS techniques for imaging are particularly effective in settings where (i) the cost of acquiring each pixel measurement is high, and (ii) the target image can be sparsely represented by a suitable fixed basis. One example of a CS imaging system is the single-pixel camera [10]. CS has also been proposed as a viable approach in astronomical imaging, since the actual sparsity of a size- N image of the night sky (measured in terms of the canonical basis) is much smaller than N [3].

3. A CS FRAMEWORK FOR 2D PULSES

3.1. Motivation

Due to variable environmental conditions, as well as imperfections in the imaging apparatus (such as a poorly focussed lens), acquired images are often corrupted by a small amount of unknown blur. In applications such as astronomical imaging, the objects in the field of view may be reasonably modeled to be at infinity, so that the blurring function remains

stationary across different parts of the image. Assuming a sparse underlying field of light sources (which are well well-separated in space) and a stationary blurring function, the acquired image appears as a superposition of non-overlapping pulses with identical shapes but varying amplitudes. If the original image consists of S nonzeros and the blurring kernel is defined over a domain Ω , then the overall sparsity of the acquired image jumps to $\mathcal{O}(S|\Omega|)$ (where $|\cdot|$ denotes cardinality of a set). Thus, for a CS imaging system to yield comparable image quality as in the ideal case, the number of measurements increases to $\mathcal{O}(S|\Omega| \log(N/S|\Omega|))$. In many situations, the cost of acquiring these additional measurements may be prohibitively high.

Despite this apparent increase in the sparsity of the acquired image, it is clear that the number of “independent” degrees of freedom in the image is much smaller than $S|\Omega|$. This is due to the fact that the sparse signal model does not capture the inter-dependencies among the nonzero coefficients. This motivates us to develop a new CS framework for 2D-fields of pulses. As with existing CS approaches, we introduce and analyze three concepts: a high-dimensional model for our signals of interest, a sampling bound (akin to the RIP), and a feasible recovery algorithm.

3.2. Signal model

Without loss of generality, we will assume that the signals of interest are square 2D-images of N pixels. Thus, a typical signal z consists of a sum of S spikes x in a 2D-field circularly convolved with an unknown 2D-blurring kernel h :

$$z = x * h,$$

Owing to the commutative nature of convolution, any such element z can be represented in multiple ways:

$$z = Hx = Hx,$$

where H (respectively, X) is a square circulant matrix with its columns comprising two-dimensional circular shifts of the vector h (respectively, x). In general, for a given z , H and x need not be unique. To avoid possible ambiguities, we make the following two assumptions: (i) The blurring kernel is *minimum phase*. For square 2D-signals, this implies that the nonzero coefficients of the blurring kernel are concentrated in a circular region Ω of diameter d around the *center* of the square. This is in order to remove any possible ambiguities arising from the shift-invariant nature of convolution. (ii) The S nonzero spikes comprising x are separated by at least Δ pixels in discrete 2D-space, where $\Delta > d$. This is in order to ensure that the pulses do not overlap. Hence, we are implicitly assuming a special *structure* among the sparse nonzero coefficients of x ; a similar model for 1D signals has been introduced and studied in [11]. We denote this special 2D-structured sparsity model as \mathcal{M}_S^Δ . A general signal model for a superposition of 2D-pulses can be defined thus:

Definition 2 Let \mathcal{M}_S^Δ be the structured sparsity model in the space of N -dimensional square 2D-images as defined above. Let \mathcal{M}_Ω be the space of minimum phase kernels with nonzeros restricted to the region Ω . Define the set:

$$\mathcal{M}(S, \Omega, \Delta) := \{z \in \mathbb{R}^N : z = x * h, \\ \text{such that } x \in \mathcal{M}_S^\Delta \text{ and } h \in \mathcal{M}_\Omega\}.$$

Then, $\mathcal{M}(S, \Omega, \Delta)$ is called a 2D pulse-field model.

3.3. Sampling bound

We derive an RIP-like sampling bound for signals belonging to the 2D pulse-field model. Observe that any signal $z \in \mathcal{M}(S, \Omega, \Delta)$ consists of at most $S \times |\Omega|$ nonzeros; thus, $\mathcal{M}(S, \Omega, \Delta)$ is a subset of the set of all $(S|\Omega|)$ -sparse signals. On the other hand, only a small fraction of all $(S|\Omega|)$ -sparse signals can be written as a superposition of pulses. Thus, to achieve an RIP-like stable embedding for this reduced set of signals, intuition suggests that we require far fewer than $\mathcal{O}(S|\Omega| \log N)$ linear measurements. The following theorem makes this notion precise.

Theorem 1 Suppose $\mathcal{M}(S, \Omega, \Delta)$ is a 2D pulse-field model as in Definition 2. Then, for any $t > 0$ and

$$M \geq \mathcal{O} \left(\frac{1}{\delta^2} \left((S + |\Omega|) \ln \left(\frac{1}{\delta} \right) + S \log(N/S - \Delta) + t \right) \right)$$

an $M \times N$ i.i.d. subgaussian matrix Φ will satisfy the following property with probability at least $1 - e^{-t}$: for every pair $z_1, z_2 \in \mathcal{M}(S, \Omega, \Delta)$,

$$(1 - \delta) \|z_1 - z_2\|_2^2 \leq \|\Phi z_1 - \Phi z_2\|_2^2 \leq (1 + \delta) \|z_1 - z_2\|_2^2.$$

Proof sketch. A full proof of this theorem is presented in the expanded final version of the manuscript [12]. The proof mechanism involves constructing a net of points in \mathbb{R}^N which can be shown to lie within a squared distance δ of any normalized point z belonging to the model, and applying the Johnson-Lindenstrauss lemma for approximate metric preservation of finite point clouds. \square

Theorem 1 indicates that the number of measurements required for the stable geometric embedding of signals in \mathcal{M} is proportional to $(S + |\Omega|)$; thus, it is *sublinear* in the maximum sparsity $S|\Omega|$ of the images of interest.

3.4. Recovery algorithm

The CS recovery problem can be stated as follows: given measurements of a superposition of pulses:

$$y = \Phi z = \Phi H x = \Phi X h,$$

the goal is to reconstruct the best possible $z \in \mathcal{M}(S, \Omega, \Delta)$ from the measurements y . Standard or structured sparsity methods for CS recovery are unsuitable for this problem,

Algorithm 1 CS recovery of 2D pulse-fields

Inputs: Projection matrix Φ , measurements y , model parameters S, Ω, Δ .

Output: $\mathcal{M}(S, \Omega, \Delta)$ -approximation \hat{z} to true signal z

Initialize: $\hat{x} = 0, \hat{h} = (\mathbf{1}_\Omega^\top, 0, \dots, 0); i = 0$

while halting criterion false **do**

(estimate spike locations and amplitudes)

1. $i \leftarrow i + 1$

2. $\hat{z} \leftarrow \hat{x} * \hat{h}$

3. $\hat{H} = \mathbb{C}(\hat{h}), \Phi_h = \Phi \hat{H}$

4. $e \leftarrow \Phi_h^\top (y - \Phi_h \hat{x})$

5. $\Omega \leftarrow \text{supp}(\mathbb{D}_2(e))$

6. $T \leftarrow \Omega \cup \text{supp}(\hat{x}_{i-1})$

7. $b|_T \leftarrow (\Phi_h)^\dagger_T y, b|_{T^c} = 0$

8. $\hat{x} \leftarrow \mathbb{D}(b)$

(estimate blurring kernel)

9. $\hat{X} = \mathbb{C}(\hat{x}), \Phi_x = \Phi \hat{X}$

10. $\hat{h} \leftarrow \Phi_x^\dagger y$

end while

return $\hat{z} \leftarrow \hat{x} * \hat{h}$

since both x (respectively, X) and h (respectively, H) are unknown and have to be simultaneously inferred. This task is similar to performing *blind deconvolution* [5], which attempts simultaneous inference of the spike locations and kernel coefficients; the key difference is that in our case, we are only given access to the random measurements y and not the Nyquist-rate samples x . We adopt a two-stage iterative approach for signal recovery, akin to the Richardson-Lucy algorithm for blind deconvolution [13]. We fix estimates of the spikes \hat{x}_i and kernel coefficients \hat{h}_i , and update the configuration of the spike locations. This can be shown to be equivalent to performing one iteration of CoSaMP, followed by solving a simple linear program (refer [11] for a detailed description of the 1D-analogue of the procedure.) Once a new candidate for the spike vector \hat{x}_{i+1} has been chosen, we solve for the kernel coefficients \hat{h}_{i+1} using a least-squares procedure. This process is iterated until convergence.

The full algorithm is detailed in pseudocode form in Algorithm 1. Steps 1 to 8 within each iteration are carried out in order to update the estimates of the spike locations and amplitudes given a fixed estimate of the 2D-kernel; Steps 9 and 10 are subsequently used to update the kernel coefficients. The algorithm possesses the same iterative structure as CoSaMP, with three main differences: (i) the kernel estimate \hat{h} can be transformed into its circular convolution operator \hat{H} via the one-to-one mapping \mathbb{C} and is subsequently used to form the overall linear operator Φ_h ; (ii) the sparse approximation steps (5 and 8) are instead replaced by linear programs \mathbb{D}, \mathbb{D}_2 in order to perform best approximations to the \mathcal{M}_S^Δ -structured sparsity model; and, (iii) the kernel is updated in steps 9 and




Fig. 1. (a) Black-and-white image of V838 Monocerotis, a nova-like star, captured by the Hubble Space Telescope on February 8, 2004 (<http://heritage.stsci.edu/gallery/bwgallery/bw0405/about.shtml>). (b) Test image is a zoomed in-version of the region highlighted in green (resolution $N = 64 \times 64 = 4096$.) Reconstruction of test image from $M = 330$ random Gaussian measurements was performed using (c) CoSaMP and (d) Algorithm 1. Our algorithm is robust to variations in the pulse shapes across different 2D pulses.

10 via a simple matrix pseudoinverse. A more detailed description of the mechanism and properties of the algorithm can be found in the expanded version of this paper [12].

4. NUMERICAL EXAMPLES

To demonstrate the utility of our approach, we test our proposed algorithm on a real astronomical image. Our test image is a 64×64 region of a high-resolution image of V838 Monocerotis (a nova-like variable star) captured by the Hubble Space Telescope (highlighted by the green square in Figure 1(a)). Notice the significant variations in the shapes of the 3 large pulses in the test image (Figure 1(b)), as well as smaller, spurious pulses. We measure this image using $M = 330$ random measurements and reconstruct using both the sparse approximation approach (CoSaMP) as well as our proposed Algorithm 1. For our reconstruction methods, we used $S = 3$, $|\Omega| = 120$, $K = 360$ and $\Delta = 20$. As is visually evident from Figure 1, conventional CS does not provide useful results with this reduced set of measurements. In contrast, our approach gives us excellent estimates for the locations of the pulses. Further, our algorithm also provides an circular pulse shape estimate that could be viewed to be equivalent to be a weighted average of the 3 original pulses.

5. DISCUSSION

We have introduced a new CS framework for a superposition of 2D pulses. The key notion is to establish a particular geometric model governing the signal set of interest. This enabled us to quantitatively derive a reduced bound on the number of random measurements required for the stable embedding of this set. Further, we have developed a recovery algorithm that estimates both the spike locations as well as the 2D-profile of the pulses, and numerically demonstrated its benefits over state-of-the-art methods for CS recovery. We have discussed sparse signals and images as represented in the identity basis; our method could be extended in principle

to wavelet- and Fourier-sparse 2D-functions. While our theoretical results are promising, we still do not possess a complete characterization of the convergence properties of our proposed algorithm, as well as its sensitivity to factors such as noise and model-mismatch; we defer these challenging theoretical issues to future research.

6. REFERENCES

- [1] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, September 2006.
- [2] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Info. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [3] J. Bobin, J.-L. Starck, and R. Ottensamer, "Compressed sensing in astronomy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 5, pp. 718–726, 2008.
- [4] M. Yaghoobi, T. Blumensath, and M. E. Davies, "Dictionary learning for sparse approximations using the majorization method," 2009, To appear in *IEEE Trans. Sig. Proc.*
- [5] S. Haykin, "Blind deconvolution," 1994, Prentice Hall.
- [6] C. Hegde and R. G. Baraniuk, "Compressive sensing of streams of pulses," in *Proc. Allerton Conf. on Comm., Control, and Comp.*, 2009.
- [7] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematicians*, Madrid, Spain, 2006, vol. 3, pp. 1433–1452.
- [8] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann, "Uniform uncertainty principle for Bernoulli and subgaussian ensembles," *Constructive Approximation*, Feb. 2008.
- [9] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *ACHA*, June 2008.
- [10] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Mag.*, vol. 25, pp. 83–91, Mar. 2008.
- [11] C. Hegde, M. F. Duarte, and V. Cevher, "Compressive sensing recovery of spike trains using a structured sparsity model," in *SPARS*, Saint Malo, France, April 2009.
- [12] C. Hegde and R. G. Baraniuk, "Sampling and recovery of a sparse sum of pulses," 2010, Preprint. Available at <http://dsp.rice.edu/cs>.
- [13] J.-L. Starck, E. Pantin, and F. Murtagh, "Deconvolution in astronomy: a review," vol. 114, pp. 1051–1069, 2002.