

Model-Based Compressive Sensing for Signal Ensembles

Marco F. Duarte

Program in Applied and Computational Mathematics
Princeton University
Princeton, NJ 08544
Email: mduarte@princeton.edu

Volkan Cevher, Richard G. Baraniuk

Department of Electrical and Computer Engineering
Rice University
Houston, TX 77005
Email: {volkan,richb}@rice.edu

Abstract—Compressive sensing (CS) is an alternative to Shannon/Nyquist sampling for acquiring sparse or compressible signals. Instead of taking N periodic samples, we measure $M \ll N$ inner products with random vectors and then recover the signal via a sparsity-seeking optimization or greedy algorithm. A new framework for CS based on unions of subspaces can improve signal recovery by including dependencies between values and locations of the signal’s significant coefficients. In this paper, we extend this framework to the acquisition of signal ensembles under a common sparse supports model. The new framework provides recovery algorithms with theoretical performance guarantees. Additionally, the framework scales naturally to large sensor networks: the number of measurements needed for each signal does not increase as the network becomes larger. Furthermore, the complexity of the recovery algorithm is only linear in the size of the network. We provide experimental results using synthetic and real-world signals that confirm these benefits.

I. INTRODUCTION

Compressive sensing (CS) is a new approach to simultaneous sensing and compression that enables a potentially large reduction in the sampling and computation costs at a sensor for signals having a sparse representation in some basis. CS builds on the work of Candès, Romberg, and Tao [1] and Donoho [2], who showed that a signal having a sparse representation in one basis can be recovered from a small set of projections onto a second measurement basis that is *incoherent* with the first.¹ *Random projections* play a central role as a *universal* measurement basis in the sense that they are incoherent with any fixed basis with high probability. The CS measurement process is nonadaptive, and the recovery process is nonlinear, for which a variety of algorithms have been proposed [1–6].

While CS has relied mostly on a simplistic sparse

¹Roughly speaking, *incoherence* means that no element of one basis has a sparse representation in terms of the other basis.

or compressible signal model, there exists a parallel framework for more general *structured sparsity models* that favor certain configurations for the magnitudes and indices of the significant coefficients of the signal. It is then possible to design recovery algorithms that exploit the knowledge of this structure [7–11]. By reducing the degrees of freedom of a sparse or compressible signal, structured sparsity models provide two immediate benefits to CS. First, they enable a reduction in the number of measurements M required to stably recover a signal. Second, during signal recovery, they enable us to better differentiate true signal information from recovery artifacts, which leads to a more robust recovery.

CS is particularly apt for distributed sensor networks [12–14], where multiple sparse or compressible signals generated by a single physical process are recorded simultaneously. It is possible to improve over standard CS by leveraging the structure present among the sensed signals. As an example, the *common sparse supports* model for signal ensembles assumes that the signals share the location of their nonzero coefficients. In this case, the use of specially tailored signal recovery algorithms provides better performance than standard CS, as it is possible to recover strictly sparse signals from even fewer measurements [15]. Furthermore, this improvement is achieved without requiring collaboration between the sensors during the measurement process; this provides lower communication requirements that are crucial in sensor network applications.

In this paper, we introduce a CS recovery algorithm for signal ensembles with common sparse supports that leverages the model-based CS theory of [8] while conserving the reduced communication requirements of [15]. The model-based CS theory provides us with provable guarantees on the recovery performance of this new algorithm. Additionally, the algorithm reduces the dependence of the recovery’s computational com-

plexity on the number of sensors to be only linear. Our theoretical results also show that the number of measurements needed per sensors does not increase as the sensor network gets larger. Interestingly, we observe experimentally that the number of CS measurements per sensor required for recovery actually decreases as the number of sensors that communicate with each other increases; this provides better scalability for applications that use a large number of sensors, such as surveillance and monitoring.

This paper is organized as follows. A review of the CS theory in Section II lays out the foundational concepts. Section III develops the common sparse supports model as a union of subspaces and introduces a model-based recovery algorithm with robustness guarantees. Section IV reports on a series of numerical experiments that demonstrates the improved performance of the algorithm. We conclude with a discussion in Section V.

II. BACKGROUND

A. Compressive Sensing

Given a basis Ψ , we can represent every signal $x \in \mathbb{R}^N$ in terms of the coefficient vector α as $x = \Psi\alpha$. In this section we will assume without loss of generality that the signal x is sparse or compressible in the canonical domain so that the sparsity basis Ψ is the identity and $\alpha = x$. A signal x is K -sparse if only $K \ll N$ entries of x are nonzero. We call the set of indices corresponding to the nonzero entries the *support* of x and denote it by $\text{supp}(x)$. The set Σ_K of all K -sparse signals is the union of the $\binom{N}{K}$, K -dimensional subspaces aligned with the coordinate axes in \mathbb{R}^N . While many natural and manmade signals are not strictly sparse, the absolute magnitude of their coefficients decay quickly and therefore can be approximated as such; we say that such signals are *compressible*.

Compressive sensing (CS) integrates the signal acquisition and compression steps into a single process [1–3]. In CS we do not acquire x directly but rather acquire $M < N$ linear measurements $y = \Phi x$ using an $M \times N$ measurement matrix Φ . We then recover x by exploiting its sparsity or compressibility. Our goal is to push M as close as possible to K in order to perform as much signal “compression” during acquisition as possible. In order to recover a good estimate of x from the M compressive measurements, the measurement matrix Φ should satisfy the *restricted isometry property* (RIP):

Definition 1: [1] An $M \times N$ matrix Φ has the K -

RIP with constant δ_K if, for all $x \in \Sigma_K$,

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

In words, the K -RIP ensures that all submatrices of Φ of size $M \times K$ are close to an isometry, and therefore distance (and information) preserving. Practical recovery algorithms typically require that Φ have a slightly stronger $2K$ -RIP or higher-order RIP in order to preserve differences of K -sparse vectors (which are $2K$ -sparse in general) and other higher-order structures [1, 5, 16].

While checking whether a measurement matrix Φ satisfies the K -RIP for a fixed value of δ_K is an NP-Complete problem in general [1], random matrices whose entries are i.i.d. subgaussian random variables² work with high probability, provided $M = \mathcal{O}(K \log(N/K))$. These random matrices also have a so-called *universality* property in that, for any choice of orthonormal basis matrix Ψ , $\Phi\Psi$ has the K -RIP with high probability. This is useful when the signal is sparse in a basis other than the identity. A random Φ corresponds to an intriguing data acquisition protocol in which each measurement y_j is a randomly weighted linear combination of the entries of x .

A number of different CS signal recovery algorithms, both from optimization and greedy approaches [1–5], offer provably stable signal recovery with performance close to optimal K -term sparse approximation. For a matrix Φ that holds the $2K$ -RIP and noisy measurements $y = \Phi x + n$, the recovered signal \hat{x} has the guarantee

$$\|x - \hat{x}\|_2 \leq C_1 \|x - x_K\|_2 + \frac{C_2}{\sqrt{K}} \|x - x_K\|_1 + C_3 \|n\|_2,$$

where x_K denotes the K -sparse approximation to x , and C_1 and C_2 denote constants. This result has many implications. Under noiseless measurements, K -sparse signals are recovered perfectly; for compressible signals, the recovery has distortion close to that of sparse approximation; and when noise is present, the effect on recovery distortion is bounded.

B. Distributed Compressive Sensing

Distributed compressive sensing (DCS) [15] is an extension of the CS acquisition framework to correlated signal ensembles. Let x_1, \dots, x_J be a set of sparse or

²A random variable X is called subgaussian if there exists $c > 0$ such that $\mathbb{E}(e^{Xt}) \leq e^{c^2 t^2 / 2}$ for all $t \in \mathbb{R}$. Examples include the Gaussian and Bernoulli random variables, as well as any bounded random variable.

compressible signals that are acquired simultaneously by a group of sensors. Since the different signals correspond to different recordings of the same event, we can expect significant structure to be present between their significant values and coefficients. A relevant example is the *common sparse supports model*, which assumes that all signals share the locations of their significant coefficients. A practical situation that follows this model is where multiple sensors acquire the same Fourier-sparse signal but with phase shifts and attenuations caused by signal propagation.

While it is possible to perform CS acquisition at each sensor and then perform separate recovery for each individual signal, this naive approach would ignore the structure present between the signals in the ensemble. In DCS, we employ this additional structure through specially tailored signal recovery algorithms. To reduce the communication required between the sensors during the measurement process, each sensor obtains measurements of its own signal $y_j = \Phi_j x_j$ independently; these measurements are then sent to a central location that is interested in recovering the signal. We stack the signal and measurement vectors into single vectors $X = [x_1^T \dots x_j^T]^T$ and $Y = [y_1^T \dots y_j^T]^T$. The distributed measurement structure yields a single measurement matrix Φ of the form

$$\Phi = \begin{bmatrix} \Phi_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \Phi_2 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \Phi_J \end{bmatrix} \quad (2)$$

that provides us with the standard CS measurement equation $Y = \Phi X$. We then can recover the signal ensemble from Y using the matrix Φ with standard CS recovery algorithms. For example, the Simultaneous Orthogonal Matching Pursuit (SOMP) [17] algorithm is a greedy algorithm for recovery of signal ensembles with common sparse supports that exploits the structure of the matrix Φ shown in (2). While SOMP provides good experimental performance [15], the theoretical guarantees obtained for greedy algorithms in [6] do not apply due to the non-dense structure of the matrix Φ .

C. Model-based Compressive Sensing

While many natural and manmade signals and images can be described to first-order as sparse or compressible, the support of their large coefficients often has an underlying inter-dependency structure. A new framework for CS captures such structure using a union-of-subspaces model [8]. Such a model reduces

the degrees of freedom of a sparse/compressible signal by permitting only certain configurations of supports for the large coefficients. We also develop the example union-of-subspaces model of common sparse supports for signal ensembles. As we will show, a model allows us to reduce — in some cases significantly — the number of measurements M required to stably recover a signal.

Recall that a K -sparse signal vector x lives in $\Sigma_K \subset \mathbb{R}^N$, which is a union of $\binom{N}{K}$ subspaces of dimension K . Other than its K -sparsity, there are no further constraints on the support or values of its coefficients. A *union-of-subspaces signal model* (a *signal model* in the sequel for brevity) endows the K -sparse signal x with additional structure that allows certain K -dimensional subspaces in Σ_K and disallows others [9, 10]. More formally, let $x|_\Omega$ represent the entries of x corresponding to the set of indices $\Omega \subseteq \{1, \dots, N\}$, and let Ω^c denote the complement of the set Ω . A signal model \mathcal{M}_K is then defined as the union of m_K canonical K -dimensional subspaces

$$\mathcal{M}_K = \bigcup_{m=1}^{m_K} \mathcal{X}_m, \quad \mathcal{X}_m := \{x : x|_{\Omega_m} \in \mathbb{R}^K, x|_{\Omega_m^c} = 0\};$$

each subspace \mathcal{X}_m contains all signals x with $\text{supp}(x) \in \Omega_m$. Thus, the signal model \mathcal{M}_K is defined by the set of possible supports $\{\Omega_1, \dots, \Omega_{m_K}\}$. Signals from \mathcal{M}_K are called *K -model sparse*. Clearly, $\mathcal{M}_K \subseteq \Sigma_K$ and contains $m_K \leq \binom{N}{K}$ subspaces. In the sequel, we will use an algorithm $\mathbb{M}(x, K)$ that returns the best K -term approximation of the signal x under the model \mathcal{M}_K .

If we know that the signal x being acquired is K -model sparse, then we can relax the RIP constraint on the CS measurement matrix Φ and still achieve stable recovery from the compressive measurements $y = \Phi x$. A *model-based RIP* requires that (1) holds *only* for signals $x \in \mathcal{M}_K$ [9, 10]; we denote this new property as \mathcal{M}_K -RIP to specify the dependence on the chosen signal model, and change the model-based RIP constant from δ_K to $\delta_{\mathcal{M}_K}$ for clarity. Blumensath and Davies [9] have quantified the number of measurements M necessary for a subgaussian CS matrix to have the \mathcal{M}_K -RIP with constant $\delta_{\mathcal{M}_K}$ and with probability $1 - e^{-t}$ to be

$$M \geq \frac{2}{c\delta_{\mathcal{M}_K}^2} \left(\ln(2m_K) + K \ln \frac{12}{\delta_{\mathcal{M}_K}} + t \right). \quad (3)$$

This bound can be used to recover the conventional CS result by substituting $m_K = \binom{N}{K} \approx (Ne/K)^K$.

To take practical advantage of signal models in CS, we need to integrate them into a standard CS recovery

algorithm. For most greedy algorithms, the key modification is simple: we merely replace the best K -term approximation step (usually applied using thresholding) with a best K -term model-based approximation. Since at each iteration we need to search only over the m_K subspaces of \mathcal{M}_K rather than $\binom{N}{K}$ subspaces of Σ_K , fewer measurements will be required for the same degree of robust signal recovery. Or, alternatively, using the same number of measurements, more accurate recovery can be achieved.

We choose to modify the CoSaMP algorithm [5] for two reasons. First, it has robust recovery guarantees based on the K -RIP that are on par with the best convex optimization-based approaches. Second, it has a simple iterative, greedy structure based on a best BK -term approximation (with B a small integer) that is easily modified to incorporate a best BK -term model-based approximation.

III. COMMON SPARSE SUPPORTS MODEL

In this section, we formulate the common sparse supports model for signal ensembles as a union of subspaces. This enables us to leverage the model-based CS framework of [8] to obtain recovery algorithms with provable guarantees.

A. Formulation as union of subspaces

Consider once again an ensemble of length- N signals $\{x_1, \dots, x_J\}$ that are captured simultaneously. It is possible to stack the signals as rows of a matrix $\tilde{X} \in \mathbb{R}^{N \times J}$ or into a single vector $X \in \mathbb{R}^{NJ}$, where we assign indices to the entries of X that identify both the signal $j \in \{1, \dots, J\}$ and entry $n \in \{1, \dots, N\}$ that is observed. Under the common sparse supports model, all signals share the locations of their nonzero coefficients. We formalize the model using a union-of-subspaces formulation:

Definition 2: Define the set of K -sparse signals with common supports as

$$\mathcal{S}_K = \{X = [x_1^T \dots x_N^T]^T \in \mathbb{R}^{JN} \text{ s.t. } x_j(n) = 0 \text{ for } n \notin \Omega, \Omega \subseteq \{1, \dots, N\}, |\Omega| = K\}.$$

In a slight abuse of notation, we label this model \mathcal{S}_K even though the dimensionality of the subspaces contained is JK .

To formulate a model-based recovery algorithm, we must obtain model-based approximations using \mathcal{S}_K . For this purpose, we define the mixed norm of a matrix:

Definition 3: The (p, q) mixed norm of the matrix $\tilde{X} = [x_1 \ x_2 \ \dots \ x_N]$ is defined as

$$\|\tilde{X}\|_{(p,q)} = \left(\sum_{n=1}^N \|x_n\|_p^q \right)^{1/q}.$$

When $q = 0$, $\|\tilde{X}\|_{(p,0)}$ simply counts the number of nonzero columns in X . We immediately find that $\|\tilde{X}\|_{(p,p)} = \|X\|_p$.

Intuitively, we pose the algorithm $\mathbb{S}(X, K)$ to obtain the best approximation of the signal X under the model \mathcal{S}_K as follows:

$$\tilde{X}_K^{\mathbb{S}} = \arg \min_{X' \in \mathbb{R}^{J \times N}} \|\tilde{X} - X'\|_{(2,2)} \text{ s.t. } \|X'\|_{(2,0)} \leq K,$$

with X_K denoting the row-concatenated vector corresponding to the matrix \tilde{X}_K . It is easy to show that to obtain the approximation, it suffices to perform column-wise hard thresholding: let ρ be the K^{th} largest ℓ_2 -norm among the columns of \tilde{X} . Our approximation algorithm is $\mathbb{S}(\tilde{X}, K) = \tilde{X}_K^{\mathbb{S}} = [\mathcal{T}(x_1) \ \dots \ \mathcal{T}(x_N)]$, where

$$\mathcal{T}(x) = \begin{cases} x & \|x\|_2 \geq \rho, \\ 0 & \|x\|_2 < \rho, \end{cases}$$

for each $1 \leq j \leq J$ and $1 \leq n \leq N$. Alternatively, a recursive approximation algorithm can be obtained by sorting the rows of X by their ℓ_2 norms, and then selecting the rows with largest norm. The complexity of this sorting process is $\mathcal{O}(NJ + N \log N)$.

B. Model-based Recovery

Pseudocode for a modified version of the CoSaMP algorithm tailored to the common sparse supports model is given in Algorithm 1, where it has been specialized for matrices with the structure given in (2). Thanks to this structure, the operations on the matrix Φ can be distributed into the individual matrices Φ_j , thus making the computational complexity of the recovery algorithm only linear in the number of sensors J . Due to its resemblance to SOMP, we dub this new algorithm CoSOMP.

We obtain the following guarantee for signal recovery using CoSOMP, proven in [8].

Theorem 1: Let X be a signal from model \mathcal{S} , and let $Y = \Phi X + n$ be a set of noisy CS measurements. If Φ has the \mathcal{S}_K^4 -RIP with $\delta_{\mathcal{S}_K^4} \leq 0.1$, then the estimate \hat{X} obtained from iteration i of CoSOMP, using the approximation algorithm (4), satisfies

$$\|X - \hat{X}_i\|_2 \leq 2^{-i} \|X\|_2 + 20 \|X - X_K\|_2 + \frac{20}{\sqrt{K}} \|X - X_K\|_1 + 20 \|n\|_2.$$

Algorithm 1 CoSOMP algorithm. $\mathbb{A}(x, K)$ obtains the best K -term approximation of the vector x .

Inputs: CS matrices $\{\Phi\}_{j=1}^J$, measurements $\{y_j\}_{j=1}^J$

Output: K -sparse approximations $\{\hat{x}_j\}_{j=1}^J$ to true signals $\{x_j\}_{j=1}^J$

For each $j = 1, \dots, J$: $\hat{x}_{j,0} = 0$, $r_j = y_j$; $i = 0$ {initialize}

while halting criterion false **do**

1. $i \leftarrow i + 1$

2. $e_j \leftarrow \Phi_j^T r_j$, $j = 1, \dots, J$ {form signal residual estimates}

3. $e = \sum_{j=1}^J (e_j \cdot e_j)$ {merge signal residual estimates in squared ℓ_2 norm}

4. $\Omega \leftarrow \text{supp}(\mathbb{A}(e, 2K))$ {prune merged signal residual estimates}

5. $T \leftarrow \Omega \cup \text{supp}(\hat{x}_{j,i-1})$ {merge supports}

6. $b_j|_T \leftarrow \Phi_j|_T^\dagger y_j$, $b_j|_{T^c} \leftarrow 0$, $j = 1, \dots, J$ {form signal estimates}

7. $b = \sum_{j=1}^J (b_j \cdot b_j)$ {merge signal estimates in squared ℓ_2 norm}

8. $\Lambda \leftarrow \text{supp}(\mathbb{A}(b, K))$ {obtain pruned signal estimate support}

9. $\hat{x}_{j,i}|_\Lambda \leftarrow b_j|_\Lambda$, $\hat{x}_{j,i}|_{\Lambda^c} \leftarrow 0$, $j = 1, \dots, J$ {prune signal estimates}

10. $r_j \leftarrow y_j - \Phi_j \hat{x}_{j,i}$, $j = 1, \dots, J$ {update measurement residuals}

end while

return $\hat{x}_j \leftarrow \hat{x}_{j,i}$, $j = 1, \dots, J$

For DCS matrices Φ that have the structure described in (2), it is easy to show that requiring Φ to have the S_K^4 -RIP with constant δ is equivalent to requiring each matrix Φ_j , $1 \leq j \leq J$ to have the $4K$ -RIP with the same constant. Remarkably, this implies that *the number of measurements required per sensor does not increase as the number of sensors increases*; in fact, the experimental results in the next section show that the number of measurements required will decrease as the network becomes larger.

Previous research has developed several algorithms for the recovery of signals with common sparse supports [11, 15, 17, 18]. However, the robustness guarantees for such algorithms either are restricted to exactly sparse signals and noiseless measurements, do not have explicit bounds on the number of necessary measurements, or are asymptotic in nature.

IV. EXPERIMENTS

In this section, we consider three datasets [19] for temperature, humidity, and light readings from a group of 48 nodes deployed at the offices of Intel Research Labs in Berkeley, CA.³ The signals were recorded in an office environment and therefore exhibit periodic behavior caused by the activity levels during day and night. Therefore, we expect the signals to be compressible in

³For the purposes of our experiments, we select signals of length $N = 1024$ and interpolate small amounts of missing data.

TABLE I
PERFORMANCE COMPARISON FOR STANDARD CoSaMP AND CoSOMP RECOVERY ON 49 ENVIRONMENTAL SENSING SIGNALS FROM THE INTEL BERKELEY DATASET.

Dataset	M	CoSaMP	CoSOMP
Light	200	14.07dB	17.87dB
Humidity	80	20.45dB	26.68dB
Temperature	400	19.10dB	26.40dB

the wavelet domain. Since the signals are observations of physical processes, they are smoothly varying in time and space; this causes the sensor readings to be close in value to each other, a situation well captured by the common sparse supports model.

We consider the recovery from CS measurements for these signals. We obtain M CS measurements for each signal using a matrix with random Gaussian distributed entries. We then compare model-based recovery using CoSOMP with standard CoSaMP recovery, where the parameter K is chosen to achieve best performance.⁴ Figure 1 shows the recovery for a representative example: the light intensity signal from sensor 35. The CoSOMP algorithm exploits the common sparse supports structure, recovering salient common features for all signals in the ensemble, and thus obtaining better recovery performance than standard CoSaMP from the same sets of measurements. Table I summarizes similar results for the different datasets.

We also study the performance of these algorithms

⁴We use Daubechies-8 wavelets for compression throughout this section.

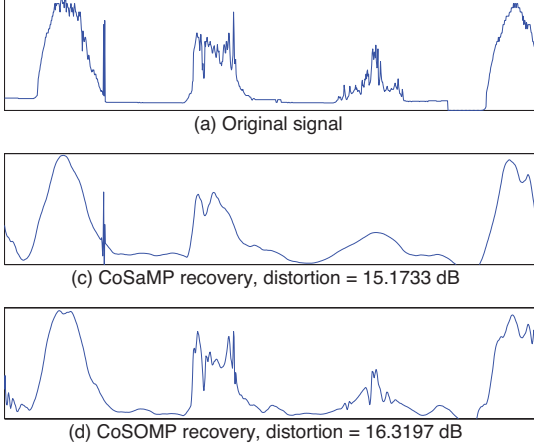


Fig. 1. Recovery of light intensity signal 35 from the Intel Berkeley sensor network using the CoSaMP and CoSOMP algorithm. $N = 1024$, $M = 400$. When the common sparse supports model is used in CoSOMP, the features that are salient in all signals are preserved.

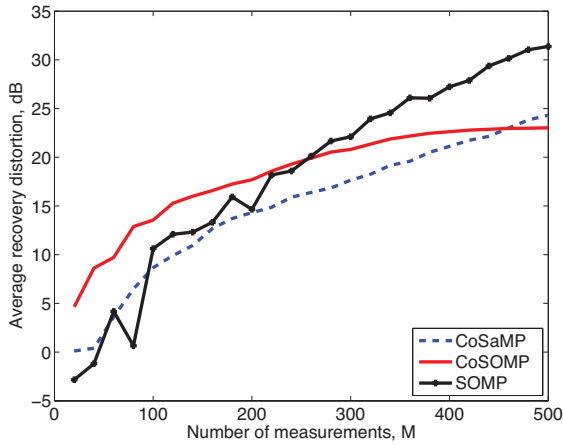


Fig. 2. Performance of CoSaMP, SOMP and CoSOMP on a group of light signals from the Intel Berkeley sensor network as a function of the number of measurements M .

for different numbers of measurements. Figures 2 – 4 plot the probability of exact recovery for the standard CoSaMP and CoSOMP algorithms for the three environmental sensing datasets; we also show the performance of SOMP as a baseline. CoSOMP recovery is superior at low and moderate rates, yet it is surpassed by standard CoSaMP at high rates. This illustrates the applicability of the common sparse supports model, which becomes less valid as the very fine features of each signal (which vary between sensors) are incorporated. While the performance of CoSOMP is similar to that of SOMP, CoSOMP has the added benefit of the proven recovery guarantees.




Fig. 3. Performance of CoSaMP, SOMP and CoSOMP on a group of humidity signals from the Intel Berkeley sensor network as a function of the number of measurements M .

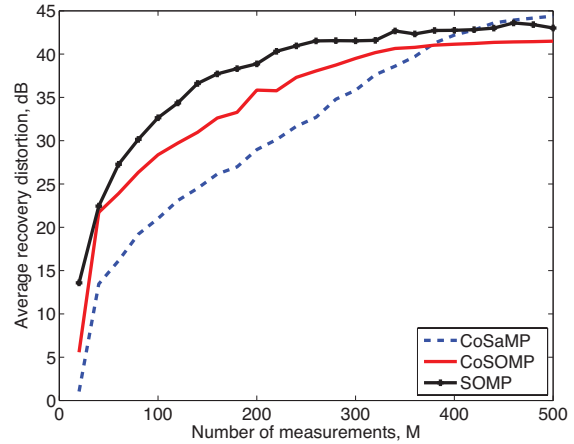


Fig. 4. Performance of CoSaMP, SOMP and CoSOMP on a group of temperature signals from the Intel Berkeley sensor network as a function of the number of measurements M .

Finally, we study the dependence of CoSOMP performance on the number of signals in the ensemble. Figure 5 compares the performance of the standard CoSaMP and CoSOMP algorithms on synthetically generated exactly sparse signals with common sparse supports. Over 100 repetitions, we select the signal supports at random and assign coefficients from a standard Gaussian distribution. We then obtain CS measurements for each signal using matrices with entries following a standard Gaussian distribution. The figure shows that while standard CoSaMP recovery requires more measurements to achieve high probability of successful recovery — as each sensor must succeed independently — the CoSOMP algorithm requires fewer measurements as the number of signals increases, as it is simpler to establish the common sparse support structure. We also see that the

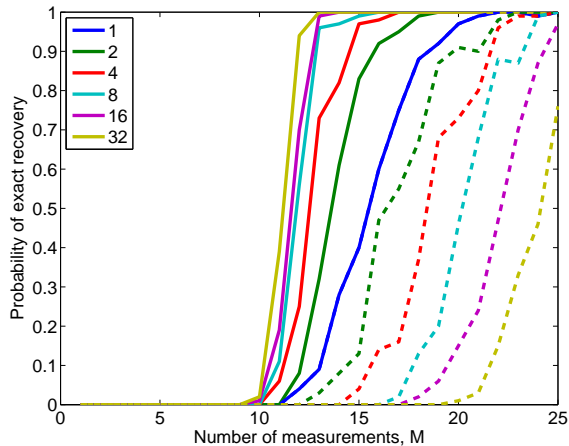


Fig. 5. Performance of CoSaMP (dashed lines) and CoSOMP (solid lines) on a class of signals with common sparse support ($K = 5$) as a function of M for several numbers of sensors J . While more measurements are required with CoSaMP as J increases, CoSOMP requires a decreasing number of measurements, appearing to converge to $M = 2K$ as $J \rightarrow \infty$.

number of measurements necessary for recovery appears to converge to $M = 2K$ as the number of sensors becomes larger; in comparison, for the SOMP algorithm this number of measurements converged to $M = K$ [12]. We believe that this increase in the bound is due to the enlarged support estimate obtained in step 4 of CoSOMP.

V. CONCLUSIONS

In this paper, we have aimed to demonstrate that there are significant performance gains to be made by exploiting more realistic and richer signal models beyond the simplistic sparse and compressible models that dominate the CS literature. Building on the model-based CS theory of [8], we have provided an algorithm for recovery of signal ensembles under a common sparse support model for which we can provide state-of-the-art guarantees on recovery performance. When distributed sensing is performed, we can reduce the dependence of the recovery computational complexity to be linear in the number of signals present. Furthermore, one can obtain a distributed recovery algorithm when sensors communicate efficiently to each other; we only need to share the proxy estimates for each signal obtained in steps 2 and 6 of CoSOMP.

ACKNOWLEDGMENT

This work was supported by the grants NSF CCF-0431150, CCF-0728867, CNS-0435425, and CNS-0520280, DARPA/ONR N66001-08-1-2065, ONR

N00014-07-1-0936, N00014-08-1-1067, N00014-08-1-1112, and N00014-08-1-1066, AFOSR FA9550-07-1-0301, ARO MURI W911NF-07-1-0185, and the Texas Instruments Leadership University Program.

REFERENCES

- [1] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematicians*, vol. 3, (Madrid, Spain), pp. 1433–1452, 2006.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, pp. 1289–1306, Sept. 2006.
- [3] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Mag.*, vol. 24, no. 4, pp. 118–120, 124, July 2007.
- [4] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Info. Theory*, vol. 52, pp. 4036–4048, Sept. 2006.
- [5] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, June 2008.
- [6] J. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," *IEEE Trans. Info. Theory*, vol. 53, pp. 4655–4666, Dec. 2007.
- [7] V. Cevher, M. F. Duarte, C. Hegde, and R. G. Baraniuk, "Sparse signal recovery using Markov Random Fields," in *Proc. Workshop on Neural Info. Proc. Sys. (NIPS)*, (Vancouver, Canada), Dec. 2008.
- [8] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," 2008. Preprint. Available at <http://dsp.rice.edu/cs>.
- [9] T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Trans. Info. Theory*, Dec. 2008.
- [10] Y. M. Lu and M. N. Do, "Sampling signals from a union of subspaces," *IEEE Signal Processing Mag.*, vol. 25, pp. 41–47, Mar. 2008.
- [11] Y. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Info. Theory*, 2008. To be published.
- [12] M. F. Duarte, M. B. Wakin, D. Baron, and R. G. Baraniuk, "Universal distributed sensing via random projections," in *International Workshop on Inf. Processing in Sensor Networks (IPSN '06)*, (Nashville, TN), pp. 177–185, Apr. 2006.
- [13] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressive wireless sensing," in *Proc. 5th International Workshop on Inf. Processing in Sensor Networks (IPSN '06)*, (Nashville, TN), pp. 134–142, Apr. 2006.
- [14] W. Wang, M. Garofalakis, and K. Ramchandran, "Distributed sparse random projections for refinable approximation," in *International Workshop on Inf. Processing in Sensor Networks (IPSN '06)*, (Cambridge, MA), pp. 331–339, Apr. 2007.
- [15] D. Baron, M. F. Duarte, S. Sarvotham, M. B. Wakin, and R. G. Baraniuk, "Distributed compressive sensing," 2005. Preprint.
- [16] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Compte Rendus de l'Academie des Sciences, Series I*, vol. 346, pp. 589–592, May 2008.
- [17] J. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit," *Signal Processing*, vol. 86, pp. 572–588, Apr. 2006.
- [18] M. Stojnic, F. Parvaresh, and B. Hassibi, "On the reconstruction of block-sparse signals with an optimal number of measurements," *IEEE Trans. Signal Processing*, vol. 47, pp. 3075–3085, Aug. 2009.
- [19] F. Koushanfar, N. Taft, and M. Potkonjak, "Sleeping coordination for comprehensive sensing using isotonic regression and domatic partitions," in *Proc. IEEE INFOCOM*, (Barcelona, Spain), Apr. 2006.