

Greedy Signal Recovery Review

Deanna Needell*, Joel Tropp, Roman Vershynin

Abstract—The two major approaches to sparse recovery are L_1 -minimization and greedy methods. Recently, Needell and Vershynin developed Regularized Orthogonal Matching Pursuit (ROMP) that has bridged the gap between these two approaches. ROMP is the first stable greedy algorithm providing uniform guarantees.

Even more recently, Needell and Tropp developed the stable greedy algorithm Compressive Sampling Matching Pursuit (CoSaMP). CoSaMP provides uniform guarantees and improves upon the stability bounds and RIC requirements of ROMP. CoSaMP offers rigorous bounds on computational cost and storage. In many cases, the running time is just $O(N \log N)$, where N is the ambient dimension of the signal. This review summarizes these major advances.

I. INTRODUCTION

Sparse signals are those that contain much less information than their ambient dimension suggests. The conventional signal compression scheme acquires the entire signal and then compresses it. This methodology has been questioned for decades, and new approaches in compressed sensing have been developed to overcome this seemingly wasteful approach.

Suppose \mathbf{x} is a signal in \mathbb{R}^N , and define the ℓ_0 quasi-norm¹

$$\|\mathbf{x}\|_0 = |\text{supp}(\mathbf{x})| = |\{j : x_j \neq 0\}|.$$

When $\|\mathbf{x}\|_0 \leq s$, we say that the signal \mathbf{x} is s -sparse. In practice, signals are not exactly sparse, but are rather close to sparse vectors. For example, compressible signals are those whose coefficients decay rapidly when sorted by magnitude. We say that a signal \mathbf{x} is p -compressible with magnitude R if the sorted components of the signal decay at the rate

$$|x|_{(i)} \leq R \cdot i^{-1/p} \quad \text{for } i = 1, 2, 3, \dots \quad (\text{I.1})$$

The sparse recovery problem is the reconstruction of such signals from a set of nonadaptive linear measurements. The measurements are of the form $\Phi \mathbf{x}$ where Φ is some $m \times N$ measurement matrix. Although in general this recovery is NP-Hard, work in compressed sensing has shown that for certain kinds of measurement matrices, recovery is possible when the number of measurements m is nearly linear in the sparsity s ,

$$m = s \log^{O(1)}(N). \quad (\text{I.2})$$

The survey [2] contains a discussion of these results.

DN is with the Dept. of Mathematics, University of California, Davis, One Shields Ave., Davis CA 95616, USA e-mail: dneedell@math.ucdavis.edu.

RV is with the Dept. of Mathematics, University of Michigan, 530 Church St., Ann Arbor MI 48105, USA email: romanv@umich.edu.

JAT is with Applied and Computational Mathematics, MC 217-50, California Inst. Technology, Pasadena, CA 91125, USA email: jtropp@acm.caltech.edu

¹We consider real numbers here for simplicity, but similar results hold for the complex case.

The two major algorithmic approaches to sparse recovery are based on L_1 -minimization and on greedy methods (Matching Pursuits). In this review we briefly describe these methods, as well as two new iterative methods that provide the advantages of both approaches. The first method, ROMP, is the first stable algorithm to provide uniform guarantees. The second method, CoSaMP, improves upon the results of ROMP, and provides rigorous computational bounds.

II. MAJOR APPROACHES

A. L_1 -minimization

To recover the sparse signal \mathbf{x} from its measurements $\Phi \mathbf{x}$, one needs to find the the solution to the highly non-convex problem

$$\min \|\mathbf{z}\|_0 \quad \text{subject to} \quad \Phi \mathbf{z} = \Phi \mathbf{x}. \quad (\text{L}_0)$$

Donoho and his associates [6] suggested that for some measurement matrices Φ , the generally NP-Hard problem (L_0) should be equivalent to its convex relaxation:

$$\min \|\mathbf{z}\|_1 \quad \text{subject to} \quad \Phi \mathbf{z} = \Phi \mathbf{x}, \quad (\text{L}_1)$$

where $\|\mathbf{z}\|_1 = \sum_i |z_i|$ denotes the ℓ_1 -norm. The convex problem (L_1) can be solved using methods of linear programming.

Clearly if the measurement matrix Φ is one-to-one on all $2s$ -sparse vectors, then the s -sparse signal \mathbf{x} will be recovered by solving (L_0). Candès and Tao [4] proved that if Φ satisfies a stronger condition then recovery is possible by solving the convex problem (L_1).

Definition II.1 (Restricted Isometry Condition): A measurement matrix Φ satisfies the *Restricted Isometry Condition* (RIC) with parameters (n, δ) for $\delta \in (0, 1)$ if we have

$$(1-\delta)\|\mathbf{v}\|_2 \leq \|\Phi \mathbf{v}\|_2 \leq (1+\delta)\|\mathbf{v}\|_2 \quad \text{for all } n\text{-sparse vectors.}$$

When δ is small, the restricted isometry condition says that every set of n columns of Φ is approximately an orthonormal system. It has been shown (see [13] and [7]) that random Gaussian, Bernoulli and partial Fourier matrices satisfy the restricted isometry condition with number of measurements as in (I.2). In the more practical case when \mathbf{x} is not exactly sparse and corrupted with noise, we consider the mathematical program:

$$\min \|\mathbf{y}\|_1 \quad \text{subject to} \quad \|\Phi \mathbf{y} - \mathbf{u}\|_2 \leq \varepsilon. \quad (\text{II.1})$$

This program approximately recovers the signal \mathbf{x} even in the presence of noise.

Theorem II.2 (Recovery under RIC [3]): Assume that the measurement matrix Φ satisfies the Restricted Isometry Condition with parameters $(3s, 0.2)$. Let Φ be a measurement matrix and let $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$ be a noisy measurement vector

where $\mathbf{x} \in \mathbb{R}^N$ is an arbitrary signal and $\|e\|_2 \leq \varepsilon$. Then the program (II.1) produces an approximation $\mathbf{x}^\#$ that satisfies:

$$\|\mathbf{x} - \mathbf{x}^\#\|_2 \leq C \left[\frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 + \varepsilon \right],$$

where \mathbf{x}_s denotes the s -sparse vector consisting of the s largest coefficients in magnitude of \mathbf{x} .

In [5], Candès sharpened this theorem to work under the restricted isometry condition with parameters $(2s, \sqrt{2} - 1)$. This theorem also demonstrates that in the noiseless case, the L_1 approach provides exact reconstruction. This was proved initially in [4].

B. Orthogonal Matching Pursuit (OMP)

An alternative approach to sparse recovery is via greedy algorithms. These methods find the support of the signal \mathbf{x} iteratively, and reconstruct the signal using the pseudoinverse.

Orthogonal Matching Pursuit (OMP) is such an algorithm, analyzed by Gilbert and Tropp in [14]. Since we expect the columns of the measurement matrix Φ to be approximately orthonormal, $\Phi^* \Phi \mathbf{x}$ is locally a good approximation to \mathbf{x} . OMP uses this idea to compute the support of a s -sparse signal \mathbf{x} . First, the residual \mathbf{r} is set to the measurement vector \mathbf{u} . At each iteration, the observation vector is set, $\mathbf{y} = \Phi^* \mathbf{r}$, and the coordinate of its largest coefficient in magnitude is added to the index set I . Then by solving a least squares problem, the residual is updated to remove this coordinate's contribution,

$$\mathbf{y} = \arg \min_{\mathbf{z} \in \mathbb{R}^I} \|\mathbf{u} - \Phi \mathbf{z}\|_2; \quad \mathbf{r} = \mathbf{u} - \Phi \mathbf{y}.$$

Repeating this s times yields an index set of s coordinates corresponding to the support of the signal \mathbf{x} . Tropp and Gilbert [14] showed that OMP recovers a sparse signal with high probability.

Theorem II.3 (OMP Recovery [14]): Let Φ be a $m \times N$ subgaussian matrix, and fix a s -sparse signal $\mathbf{x} \in \mathbb{R}^N$. Then OMP recovers (the support of) \mathbf{x} from the measurements $\mathbf{u} = \Phi \mathbf{x}$ correctly with high probability, provided the number of measurements is $m \sim s \log N$.

C. Advantages and challenges of both approaches

The L_1 -minimization method provides *uniform guarantees* for sparse recovery. Once the measurement matrix Φ satisfies the restricted isometry condition, this method works correctly for all sparse signals \mathbf{x} . The method is also *stable*, so it works for non-sparse signals such as those which are compressible, as well as noisy signals. However, the method is based on linear programming, and there is no strongly polynomial time algorithm in linear programming yet.

OMP on the other hand, is quite *fast* both provably and empirically. The speed of OMP is a great advantage, but it lacks the strong guarantees that L_1 provides. Indeed, OMP works correctly for a *fixed* signal and measurement matrix with high probability, and so it must fail for some sparse signals and matrices [12]. It is also unknown whether OMP succeeds for compressible signals or on noisy measurements.

There has thus existed a gap between the approaches. The development of Regularized Orthogonal Matching Pursuit

(ROMP) bridges this gap by providing a greedy algorithm with the same advantages as the L_1 method. Compressive Sampling Matching Pursuit (CoSaMP) improves upon these results and provides rigorous runtime guarantees. We now discuss these new algorithms.

III. REGULARIZED OMP

Regularized Orthogonal Matching Pursuit (ROMP) is a new algorithm developed by Needell and Vershynin in [11], [10] for sparse recovery that performs correctly for all measurement matrices Φ satisfying the restricted isometry condition, and for all sparse signals. Again, since the restricted isometry condition guarantees every set of s columns forms approximately an orthonormal system, every s coordinates of the observation vector $\mathbf{y} = \Phi^* \mathbf{u}$ are in a loose sense good estimators of the corresponding s coefficients of \mathbf{x} . This notion suggests to use the s largest coefficients of the observation vector \mathbf{y} rather than only the largest, as in OMP. We also include a regularization step to ensure that each coordinate carries close to an even share of the information. The ROMP algorithm is described in Algorithm III.1.

Algorithm III.1: ROMP Recovery Algorithm [11]

ROMP(Φ, \mathbf{u}, s)

Input: Measurement matrix Φ , measurement vector \mathbf{u} , sparsity level s

Output: Index set $I \subset \{1, \dots, d\}$

- 1) Let the index set $I = \emptyset$ and the residual $\mathbf{r} = \mathbf{u}$. Repeat the following steps s times or until $|I| \geq 2s$:
- 2) Choose a set J of the s biggest coordinates in magnitude of the observation vector $\mathbf{y} = \Phi^* \mathbf{r}$, or all of its nonzero coordinates, whichever set is smaller.
- 3) Among all subsets $J_0 \subset J$ with comparable coordinates:
$$|\mathbf{y}(i)| \leq 2|\mathbf{y}(j)| \quad \text{for all } i, j \in J_0,$$
 choose J_0 with the maximal energy $\|\mathbf{y}|_{J_0}\|_2$.
- 4) Add the set J_0 to the index set: $I \leftarrow I \cup J_0$, and update the residual:

$$\mathbf{w} = \arg \min_{\mathbf{z} \in \mathbb{R}^I} \|\mathbf{u} - \Phi \mathbf{z}\|_2; \quad \mathbf{r} = \mathbf{u} - \Phi \mathbf{w}.$$

For measurement matrices that satisfy the RIC, ROMP stably recovers all s -sparse signals. This is summarized in the following theorem from [10].

Theorem III.1 (Recovery by ROMP [10]): Assume a measurement matrix Φ satisfies the restricted isometry condition with parameters $(8s, \epsilon)$ for $\epsilon = 0.01/\sqrt{\log s}$. Consider an arbitrary vector \mathbf{x} in \mathbb{R}^N . Suppose that the measurement vector $\Phi \mathbf{x}$ becomes corrupted, so we consider $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$ where \mathbf{e} is some error vector. Then ROMP produces a good approximation $\hat{\mathbf{x}}$ to \mathbf{x} :

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_2 \leq C \sqrt{\log s} \left(\|e\|_2 + \frac{\|\mathbf{x} - \mathbf{x}_s\|_1}{\sqrt{s}} \right). \quad (\text{III.1})$$

In the case where the signal \mathbf{x} is exactly sparse without noise, this theorem guarantees exact reconstruction. Note also that in the noisy case, ROMP needs no knowledge about the error vector \mathbf{e} to approximate the signal. In the special case where \mathbf{x} is a compressible signal as in (I.1), the theorem provides the bound

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq R' \frac{\sqrt{\log s}}{s^{p-1/2}} + C\sqrt{\log s}\|\mathbf{e}\|_2. \quad (\text{III.2})$$

ROMP thus provides the first greedy approach with these uniform and stable guarantees. See [11], [10] for runtime analysis and empirical results. The results are optimal up to the logarithmic factors, and the runtime is similar to that of OMP. There is, however, room for improvement which leads us to the CoSaMP algorithm, developed by Needell and Tropp [9]. CoSaMP provides optimal guarantees as well as an important implementation analysis.

IV. COSAMP

CoSaMP is an iterative recovery algorithm that provides the same guarantees as even the best optimization approaches. As in the case of ROMP and the L_1 approach, CoSaMP recovers signals using measurement matrices that satisfy the RIC. Thus as before, the observation vector $\mathbf{y} = \Phi^* \mathbf{u}$ serves as a good proxy for the signal \mathbf{x} . Using the largest coordinates, an approximation to the signal is formed at each iteration. After each new residual is formed, reflecting the missing portion of the signal, the measurements are updated. This is repeated until all the recoverable portion of the signal is found. (See [9] for halting criteria.) The algorithm is described in Algorithm IV.2.

Algorithm IV.2: CoSaMP Recovery Algorithm [9]

COSAMP(Φ , \mathbf{u} , s)
Input: Sampling matrix Φ , noisy sample vector \mathbf{u} , sparsity level s
Output: An s -sparse approximation \mathbf{a} of the target signal

$\mathbf{a}^0 \leftarrow \mathbf{0}$, $\mathbf{v} \leftarrow \mathbf{u}$, $k \leftarrow 0$ { Trivial initial approximation }

repeat

$k \leftarrow k + 1$

$\mathbf{y} \leftarrow \Phi^* \mathbf{v}$ { Form signal proxy }

$\Omega \leftarrow \text{supp}(\mathbf{y}_{2s})$ { Identify large components }

$T \leftarrow \Omega \cup \text{supp}(\mathbf{a}^{k-1})$ { Merge supports }

$\mathbf{b}|_T \leftarrow \Phi_T^\dagger \mathbf{u}$ { Signal estimation by least-squares }

$\mathbf{b}|_{T^c} \leftarrow \mathbf{0}$

$\mathbf{a}^k \leftarrow \mathbf{b}_s$ { Prune to obtain next approximation }

$\mathbf{v} \leftarrow \mathbf{u} - \Phi \mathbf{a}^k$ { Update current samples }

until halting criterion *true*

For a recovery algorithm to be used efficiently in practice, the least squares step must be analyzed carefully. By the restricted isometry condition, the matrix Φ_T in the estimation step is very well-conditioned. This suggests the use of an

iterative method such as Richardson's iteration [1, Sec. 7.4] to apply the pseudoinverse $\Phi_T^\dagger = (\Phi_T^* \Phi_T)^{-1} \Phi_T^*$. This method is analyzed in the context of CoSaMP and shown to provide an efficient means of acquiring the estimation [9]. The following theorem from [9] summarizes the fundamentally optimal recovery guarantees and rigorous computational costs of CoSaMP.

Theorem IV.1 (Recovery by CoSaMP [9]): Suppose that Φ is an $m \times N$ measurement matrix satisfying the restricted isometry condition with parameters $(2s, c)$. Let $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$ be a vector of samples of an arbitrary signal, contaminated with arbitrary noise. For a given precision parameter η , the algorithm CoSaMP produces an s -sparse approximation \mathbf{a} that satisfies

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq C \cdot \max \left\{ \eta, \frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_{s/2}\|_1 + \|\mathbf{e}\|_2 \right\}$$

The running time is $O(\mathcal{L} \cdot \log(\|\mathbf{x}\|_2 / \eta))$, where \mathcal{L} bounds the cost of a matrix–vector multiply with Φ or Φ^* . Working storage is $O(N)$.

In [8] it is shown that only a fixed number of iterations is required to reduce the error to an optimal amount. The report [8] also discusses variations on the CoSaMP algorithm that may improve results in practice.

REFERENCES

- [1] Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.
- [2] E. Candès. Compressive sampling. In *Proc. International Congress of Mathematics*, volume 3, pages 1433–1452, Madrid, Spain, 2006.
- [3] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [4] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51:4203–4215, 2005.
- [5] E. J. Candès. The restricted isometry property and its implications for compressed sensing. *C. R. Math. Acad. Sci. Paris, Serie I*, 346:589–592, 2008.
- [6] D. L. Donoho and P. B. Stark. Uncertainty principles and signal recovery. *SIAM J. Appl. Math.*, 49(3):906–931, June 1989.
- [7] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Uniform uncertainty principle for Bernoulli and subgaussian ensembles. To appear, *Constr. Approx.*, 2009.
- [8] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. ACM Technical Report 2008-01, California Institute of Technology, Pasadena, July 2008.
- [9] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from noisy samples. *Appl. Comput. Harmon. Anal.*, 2008. DOI: 10.1016/j.acha.2008.07.002.
- [10] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. Submitted for publication, October 2007.
- [11] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. DOI: 10.1007/s10208-008-9031-3, 2007.
- [12] H. Rauhut. On the impossibility of uniform sparse reconstruction using greedy methods. To appear, *Sampl. Theory Signal Image Process.*, 2008.
- [13] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Comm. Pure Appl. Math.*, 61:1025–1045, 2008.
- [14] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Info. Theory*, 53(12):4655–4666, 2007.