

COSAMP: ITERATIVE SIGNAL RECOVERY FROM INCOMPLETE AND INACCURATE SAMPLES

D. NEEDELL AND J. A. TROPP

ABSTRACT. Compressive sampling offers a new paradigm for acquiring signals that are compressible with respect to an orthonormal basis. The major algorithmic challenge in compressive sampling is to approximate a compressible signal from noisy samples. This paper describes a new iterative recovery algorithm called CoSaMP that delivers the same guarantees as the best optimization-based approaches. Moreover, this algorithm offers rigorous bounds on computational cost and storage. It is likely to be extremely efficient for practical problems because it requires only matrix–vector multiplies with the sampling matrix. For compressible signals, the running time is just $O(N \log^2 N)$, where N is the length of the signal.

1. INTRODUCTION

In applications, most signals of interest contain scant information relative to their ambient dimension, but the classical approach to signal acquisition ignores this fact. We usually collect a complete representation of the target signal and process this representation to sieve out the actionable information. Then we discard the rest. Contemplating this ugly inefficiency, one might ask if it is possible instead to acquire *compressive samples*. In other words, is there some type of measurement that automatically winnows out the information from a signal? Incredibly, the answer is sometimes *yes*.

Compressive sampling refers to the idea that, for certain types of signals, a small number of nonadaptive samples carries sufficient information to approximate the signal well. Research in this area has two major components:

Sampling: How many samples are necessary to reconstruct signals to a specified precision? What type of samples? How can these sampling schemes be implemented in practice?

Reconstruction: Given the compressive samples, what algorithms can efficiently construct a signal approximation?

The literature already contains a well-developed theory of sampling, which we summarize below. Although algorithmic work has been progressing, the state of knowledge is less than complete. We assert that a practical signal reconstruction algorithm should have all of the following properties.

- It should accept samples from a variety of sampling schemes.
- It should succeed using a minimal number of samples.
- It should be robust when samples are contaminated with noise.
- It should provide optimal error guarantees for every target signal.
- It should offer provably efficient resource usage.

Date: 16 March 2008. Revised 16 April 2008 and 11 July 2008. Initially presented at Information Theory and Applications, 31 January 2008, San Diego.

2000 *Mathematics Subject Classification.* 41A46, 68Q25, 68W20, 90C27.

Key words and phrases. Algorithms, approximation, basis pursuit, compressed sensing, orthogonal matching pursuit, restricted isometry property, signal recovery, sparse approximation, uncertainty principle.

DN is with the Mathematics Dept., Univ. California at Davis, 1 Shields Ave., Davis, CA 95616. E-mail: dneedell@math.ucdavis.edu. JAT is with Applied and Computational Mathematics, MC 217-50, California Inst. Technology, Pasadena, CA 91125. E-mail: jtropp@acm.caltech.edu.

This paper presents and analyzes a novel signal reconstruction algorithm that achieves these desiderata. The algorithm is called CoSaMP, from the acoustic *compressive sampling matching pursuit*. As the name suggests, the new method is ultimately based on orthogonal matching pursuit (OMP) [39], but it incorporates several other ideas from the literature to accelerate the algorithm and to provide strong guarantees that OMP cannot. Before we describe the algorithm, let us deliver an introduction to the theory of compressive sampling.

1.1. Rudiments of Compressive Sampling. To enhance intuition, we focus on sparse and compressible signals. For vectors in \mathbb{C}^N , define the ℓ_0 “quasinorm”

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

We say that a signal \mathbf{x} is *s-sparse* when $\|\mathbf{x}\|_0 \leq s$. Sparse signals are an idealization that we do not encounter in applications, but real signals are quite often *compressible*, which means that their entries decay rapidly when sorted by magnitude. As a result, compressible signals are well approximated by sparse signals. We can also talk about signals that are compressible with respect to other orthonormal bases, such as a Fourier or wavelet basis. In this case, the sequence of coefficients in the orthogonal expansion decays quickly. It represents no loss of generality to focus on signals that are compressible with respect to the standard basis, and we do so without regret. For a more precise definition of compressibility, turn to Section 2.5.

In the theory of compressive sampling, a *sample* is a linear functional applied to a signal. The process of collecting multiple samples is best viewed as the action of a *sampling matrix* Φ on the target signal. If we take m samples, or *measurements*, of a signal in \mathbb{C}^N , then the sampling matrix Φ has dimensions $m \times N$. A natural question now arises: *How many measurements are necessary to acquire s-sparse signals?*

The minimum number of measurements $m \geq 2s$ on account of the following simple argument. The sampling matrix must not map two different *s-sparse* signals to the same set of samples. Therefore, each collection of $2s$ columns from the sampling matrix must be nonsingular. Although certain Vandermonde matrices arising from Reed–Solomon codes satisfy this property, these matrices are not really suitable for signal acquisition because they contain column submatrices that are very badly conditioned. As a result, some sparse signals are mapped to very similar sets of samples, and it is unstable to invert the sampling process numerically.

Instead, Candès and Tao proposed the stronger condition that the *geometry* of sparse signals should be preserved under the action of the sampling matrix [6]. To quantify this idea, they defined the *rth restricted isometry constant* of a matrix Φ as the least number δ_r for which

$$(1 - \delta_r) \|\mathbf{x}\|_2^2 \leq \|\Phi\mathbf{x}\|_2^2 \leq (1 + \delta_r) \|\mathbf{x}\|_2^2 \quad \text{whenever } \|\mathbf{x}\|_0 \leq r. \quad (1.1)$$

We have written $\|\cdot\|_2$ for the ℓ_2 vector norm. When $\delta_r < 1$, these inequalities imply that each collection of r columns from Φ is nonsingular, which is the minimum requirement for acquiring $(r/2)$ -sparse signals. When $\delta_r \ll 1$, the sampling operator very nearly maintains the ℓ_2 distance between each pair of $(r/2)$ -sparse signals. In consequence, it is possible to invert the sampling process stably.

To acquire *s-sparse* signals, one therefore hopes to achieve a small restricted isometry constant δ_{2s} with as few samples as possible. A striking fact is that many types of random matrices have excellent restricted isometry behavior. For example, we can often obtain $\delta_{2s} \leq 0.1$ with

$$m = O(s \log^\alpha N)$$

measurements, where α is a small integer. Unfortunately, no deterministic sampling matrix is known to satisfy a comparable bound. Even worse, it is computationally difficult to check the inequalities (1.1), so it may never be possible to exhibit an explicit example of a good sampling matrix.

As a result, it is important to understand how random sampling matrices behave. The two quintessential examples are Gaussian matrices and partial Fourier matrices.

Gaussian matrices: If the entries of $\sqrt{m}\Phi$ are independent and identically distributed standard normal variables then

$$m \geq \frac{Cr \log(N/r)}{\varepsilon^2} \implies \delta_r \leq \varepsilon \quad (1.2)$$

except with probability e^{-cm} . See [6] for details.

Partial Fourier matrices: If $\sqrt{m}\Phi$ is a uniformly random set of m rows drawn from the $N \times N$ unitary discrete Fourier transform (DFT), then

$$m \geq \frac{Cr \log^5 N \cdot \log(\varepsilon^{-1})}{\varepsilon^2} \implies \delta_r \leq \varepsilon \quad (1.3)$$

except with probability N^{-1} . See [35] for the proof.

Here and elsewhere, we follow the analyst's convention that upright letters (c, C, \dots) refer to positive, universal constants that may change from appearance to appearance.

The Gaussian matrix is important because it has optimal restricted isometry behavior. Indeed, for any $m \times N$ matrix,

$$\delta_r \leq 0.1 \implies m \geq Cr \log(N/r),$$

on account of profound geometric results of Kashin [25] and Garnaev–Gluskin [16]. Even though partial Fourier matrices may require additional samples to achieve a small restricted isometry constant, they are more interesting for the following reasons [2].

- There are technologies that acquire random Fourier measurements at unit cost per sample.
- The sampling matrix can be applied to a vector in time $O(N \log N)$.
- The sampling matrix requires only $O(m \log N)$ storage.

Other types of sampling matrices, such as the *random demodulator* [38], enjoy similar qualities. These traits are essential for the translation of compressive sampling from theory into practice.

1.2. Signal Recovery Algorithms. The major algorithmic challenge in compressive sampling is to approximate a signal given a vector of noisy samples. The literature describes a huge number of approaches to solving this problem. They fall into three rough categories:

Greedy pursuits: These methods build up an approximation one step at a time by making locally optimal choices at each step. Examples include OMP [39], stagewise OMP (StOMP) [14], and regularized OMP (ROMP) [31, 30].

Convex relaxation: These techniques solve a convex program whose minimizer is known to approximate the target signal. Many algorithms have been proposed to complete the optimization, including interior-point methods [2, 26], projected gradient methods [15], and iterative thresholding [11].

Combinatorial algorithms: These methods acquire highly structured samples of the signal that support rapid reconstruction via group testing. This class includes Fourier sampling [20, 22], chaining pursuit [17], and HHS pursuit [18], as well as some algorithms of Cormode–Muthukrishnan [9] and Iwen [24].

At present, each type of algorithm has its native shortcomings. Many of the combinatorial algorithms are extremely fast—sublinear in the length of the target signal—but they require a large number of somewhat unusual samples that may not be easy to acquire. At the other extreme, convex relaxation algorithms succeed with a very small number of measurements, but they tend to be computationally burdensome. Greedy pursuits—in particular, the ROMP algorithm—are intermediate in their running time and sampling efficiency.

CoSaMP, the algorithm described in this paper, is at heart a greedy pursuit. It also incorporates ideas from the combinatorial algorithms to guarantee speed and to provide rigorous error bounds [18]. The analysis is inspired by the work on ROMP [31, 30] and the work of Candès–Romberg–Tao [3] on convex relaxation methods. In particular, we establish the following result.

Theorem A (CoSaMP). *Suppose that Φ is an $m \times N$ sampling matrix with restricted isometry constant $\delta_{2s} \leq c$. 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*

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

where $\mathbf{x}_{s/2}$ is a best $(s/2)$ -sparse approximation to \mathbf{x} . 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 $O(N)$.

Let us expand on the statement of this result. First, recall that many types of random sampling matrices satisfy the restricted isometry hypothesis when the number of samples $m = O(s \log^\alpha N)$. Therefore, the theorem applies to a wide class of sampling schemes when the number of samples is proportional to the target sparsity and logarithmic in the ambient dimension of the signal space.

The algorithm produces an s -sparse approximation whose ℓ_2 error is comparable with the scaled ℓ_1 error of the best $(s/2)$ -sparse approximation to the signal. Of course, the algorithm cannot resolve the uncertainty due to the additive noise, so we also pay for the energy in the noise. This type of error bound is structurally optimal, as we discuss in Section 2.5. Some disparity in the sparsity levels (here, s versus $s/2$) seems to be necessary when the recovery algorithm is computationally efficient [34].

We can interpret the error guarantee as follows. In the absence of noise, the algorithm can recover an s -sparse signal to arbitrarily high precision. Performance degrades gracefully as the energy in the noise increases. Performance also degrades gracefully for compressible signals. The theorem is ultimately vacuous for signals that cannot be approximated by sparse signals, but compressive sampling is not an appropriate technique for this class.

The running time bound indicates that each matrix–vector multiplication reduces the error by a constant factor (if we amortize over the entire execution). That is, the algorithm has linear convergence¹. We find that the total runtime is roughly proportional to the *reconstruction signal-to-noise ratio*

$$\text{R-SNR} = 10 \log_{10} \left(\frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - \mathbf{a}\|_2} \right) \text{ dB.}$$

For compressible signals, one can show that $\text{R-SNR} = O(\log s)$. The runtime is also proportional to the cost of a matrix–vector multiply. For sampling matrices with a fast multiply, the algorithm is accelerated substantially. In particular, for the partial Fourier matrix, a matrix–vector multiply requires time $O(N \log N)$. It follows that the total runtime is $O(N \log N \cdot \text{R-SNR})$. For most signals of interest, this cost is nearly linear in the signal length!

The algorithm requires the sparsity level s as part of its input. There are several strategies for choosing s . The simplest method is to deduce the sparsity level s from the number m of measurements. When the signal length N is large in comparison with the other parameters, phase transition analysis [13] suggests that we can recover most sparse signals when $m \approx 2s \log N$. Therefore, the estimate $s \approx m/2 \log N$ is often reasonable. For particular applications, it may be preferable to undertake an empirical study of the relationship between the parameters.

A second approach is to run CoSaMP using a range of sparsity levels and to select the best approximation obtained. To limit the time cost, one may vary s along a geometric progression, say $s = 1, 2, 4, 8, \dots, m$, which increases the runtime by a factor no worse than $O(\log m)$. To determine

¹Mathematicians sometimes refer to linear convergence as “exponential convergence.”

the quality of the output \mathbf{a} , we may compute the error $\|\Phi\mathbf{a} - \mathbf{u}\|_2$ in approximating the samples, which reflects the actual approximation error. See Appendix A for more details.

1.3. Notation. Let us instate several pieces of notation that are carried throughout the paper. For $p \in [1, \infty]$, we write $\|\cdot\|_p$ for the usual ℓ_p vector norm. We reserve the symbol $\|\cdot\|$ for the spectral norm, i.e., the natural norm on linear maps from ℓ_2 to ℓ_2 .

Suppose that \mathbf{x} is a signal in \mathbb{C}^N and r is a positive integer. We write \mathbf{x}_r for the signal in \mathbb{C}^N that is formed by restricting \mathbf{x} to its r largest-magnitude components. Ties are broken lexicographically. This signal is a best r -sparse approximation to \mathbf{x} with respect to any ℓ_p norm. Suppose now that T is a subset of $\{1, 2, \dots, N\}$. We define the restriction of the signal to the set T as

$$\mathbf{x}|_T = \begin{cases} x_i, & i \in T \\ 0, & \text{otherwise.} \end{cases}$$

We occasionally abuse notation and treat $\mathbf{x}|_T$ as an element of \mathbb{C}^T , the linear space of complex-valued functions on the set T . We write Φ_T for the the column submatrix of Φ whose columns are listed in the set T . Finally, we define the pseudoinverse of a tall, full-rank matrix \mathbf{A} by the formula $\mathbf{A}^\dagger = (\mathbf{A}^*\mathbf{A})^{-1}\mathbf{A}^*$.

1.4. Organization. The rest of the paper has the following structure. In Section 2 we introduce the CoSaMP algorithm, we state the major theorems in more detail, and we discuss implementation and resource requirements. Section 3 describes some consequences of the restricted isometry property that pervade our analysis. The central theorem is established for sparse signals in Sections 4 and 5. We extend this result to general signals in Section 6. Finally, Section 7 places the algorithm in the context of previous work. The appendix presents variations on the algorithm.

2. THE CoSaMP ALGORITHM

This section gives an overview of the algorithm, along with explicit pseudocode. It presents the major theorems on the performance of the algorithm. Then it covers details of implementation and bounds on resource requirements.

2.1. Intuition. The most difficult part of signal reconstruction is to identify the locations of the largest components in the target signal. CoSaMP uses an approach inspired by the restricted isometry property. Suppose that the sampling matrix Φ has restricted isometry constant $\delta_s \ll 1$. For an s -sparse signal \mathbf{x} , the vector $\mathbf{y} = \Phi^*\Phi\mathbf{x}$ can serve as a proxy for the signal because the energy in each set of s components of \mathbf{y} approximates the energy in the corresponding s components of \mathbf{x} . In particular, the largest s entries of the proxy \mathbf{y} point toward the largest s entries of the signal \mathbf{x} . Since the samples have the form $\mathbf{u} = \Phi\mathbf{x}$, we can obtain the proxy just by applying the matrix Φ^* to the samples.

The algorithm invokes this idea iteratively to approximate the target signal. At each iteration, the current approximation induces a residual, the part of the target signal that has not been approximated. As the algorithm progresses, the samples are updated so that they reflect the current residual. These samples are used to construct a proxy for the residual, which permits us to identify the large components in the residual. This step yields a tentative support for the next approximation. We use the samples to estimate the approximation on this support set using least squares. This process is repeated until we have found the recoverable energy in the signal.

2.2. **Overview.** As input, the CoSaMP algorithm requires four pieces of information:

- Access to the sampling operator via matrix–vector multiplication.
- A vector of (noisy) samples of the unknown signal.
- The sparsity of the approximation to be produced.
- A halting criterion.

The algorithm is initialized with a trivial signal approximation, which means that the initial residual equals the unknown target signal. During each iteration, CoSaMP performs five major steps:

- (1) **Identification.** The algorithm forms a proxy of the residual from the current samples and locates the largest components of the proxy.
- (2) **Support Merger.** The set of newly identified components is united with the set of components that appear in the current approximation.
- (3) **Estimation.** The algorithm solves a least-squares problem to approximate the target signal on the merged set of components.
- (4) **Pruning.** The algorithm produces a new approximation by retaining only the largest entries in this least-squares signal approximation.
- (5) **Sample Update.** Finally, the samples are updated so that they reflect the residual, the part of the signal that has not been approximated.

These steps are repeated until the halting criterion is triggered. In the body of this work, we concentrate on methods that use a fixed number of iterations. Appendix A discusses some other simple stopping rules that may also be useful in practice.

Pseudocode for CoSaMP appears as Algorithm 2.1. This code describes the version of the algorithm that we analyze in this paper. Nevertheless, there are several adjustable parameters that may improve performance: the number of components selected in the identification step and the number of components retained in the pruning step. Another variation appears in Appendix A.

Algorithm 2.1: CoSaMP Recovery Algorithm

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	
<hr/>	
$\mathbf{a}^0 \leftarrow \mathbf{0}$	{ Trivial initial approximation }
$\mathbf{v} \leftarrow \mathbf{u}$	{ Current samples = input samples }
$k \leftarrow 0$	
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 <i>true</i>	

2.3. Performance Guarantees. This section describes our theoretical analysis of the behavior of CoSaMP. The next section covers the resource requirements of the algorithm. Our results depend on a set of hypotheses that has become common in the compressive sampling literature. Let us frame the standing assumptions:

CoSaMP HYPOTHESES

- The sparsity level s is fixed.
- The $m \times N$ sampling operator Φ has restricted isometry constant $\delta_{4s} \leq 0.1$.
- The signal $\mathbf{x} \in \mathbb{C}^N$ is arbitrary, except where noted.
- The noise vector $\mathbf{e} \in \mathbb{C}^m$ is arbitrary.
- The vector of samples $\mathbf{u} = \Phi\mathbf{x} + \mathbf{e}$.

We also define the *unrecoverable energy* ν in the signal. This quantity measures the baseline error in our approximation that occurs because of noise in the samples or because the signal is not sparse.

$$\nu = \|\mathbf{x} - \mathbf{x}_s\|_2 + \frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 + \|\mathbf{e}\|_2. \quad (2.1)$$

We postpone a more detailed discussion of the unrecoverable energy until Section 2.5.

Our key result is that CoSaMP makes significant progress during each iteration where the approximation error is large relative to the unrecoverable energy.

Theorem 2.1 (Iteration Invariant). *For each iteration $k \geq 0$, the signal approximation \mathbf{a}^k is s -sparse and*

$$\|\mathbf{x} - \mathbf{a}^{k+1}\|_2 \leq 0.5\|\mathbf{x} - \mathbf{a}^k\|_2 + 10\nu.$$

In particular,

$$\|\mathbf{x} - \mathbf{a}^k\|_2 \leq 2^{-k} \|\mathbf{x}\|_2 + 20\nu.$$

The proof of Theorem 2.1 will occupy us for most of this paper. In Section 4, we establish an analog for sparse signals. The version for general signals appears as a corollary in Section 6.

Theorem 2.1 has some immediate consequences for the quality of reconstruction with respect to standard signal metrics. In this setting, a sensible definition of the *signal-to-noise ratio* (SNR) is

$$\text{SNR} = 10 \log_{10} \left(\frac{\|\mathbf{x}\|_2}{\nu} \right).$$

The *reconstruction SNR* is defined as

$$\text{R-SNR} = 10 \log_{10} \left(\frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - \mathbf{a}\|_2} \right).$$

Both quantities are measured in decibels. Theorem 2.1 implies that, after k iterations, the reconstruction SNR satisfies

$$\text{R-SNR} \gtrsim \min\{3k, \text{SNR} - 13\} - 3.$$

In words, each iteration increases the reconstruction SNR by about 3 decibels until the error nears the noise floor. To reduce the error to its minimal value, the number of iterations is proportional to the SNR.

Let us consider a slightly different scenario. Suppose that the signal \mathbf{x} is s -sparse, so the unrecoverable energy is $\nu = \|\mathbf{e}\|_2$. Define the *dynamic range*

$$\Delta = 10 \log_{10} \left(\frac{\max |x_i|}{\min |x_i|} \right) \quad \text{where } i \text{ ranges over } \text{supp}(\mathbf{x}).$$

Assume moreover that the minimum nonzero component of the signal is at least 40ν . Using the fact that $\|\mathbf{x}\|_2 \leq \sqrt{s} \|\mathbf{x}\|_\infty$, it is easy to check that $\|\mathbf{x} - \mathbf{a}\|_2 \leq \min |x_i|$ as soon as the number k of

iterations satisfies

$$k \gtrsim 3.3\Delta + \log_2 \sqrt{s} + 1.$$

It follows that the support of the approximation \mathbf{a} must contain every entry in the support of the signal \mathbf{x} .

This discussion suggests that the number of iterations might be substantial if we require a high reconstruction SNR or if the signal has a very wide dynamic range. This initial impression is not entirely accurate. We have established that, when the algorithm performs arithmetic to high enough precision, then a fixed number of iterations suffices to reduce the approximation error to the same order as the unrecoverable energy. See [29] for details about these results.

Remark 2.2. *In the hypotheses, a bound on the restricted isometry constant δ_{2s} also suffices. Indeed, Corollary 3.4 of the sequel implies that $\delta_{4s} \leq 0.1$ holds whenever $\delta_{2s} \leq 0.025$.*

Remark 2.3. *The expression (2.1) for the unrecoverable energy can be simplified using Lemma 7 from [18], which states that, for every signal $\mathbf{y} \in \mathbb{C}^N$ and every positive integer t , we have*

$$\|\mathbf{y} - \mathbf{y}_t\|_2 \leq \frac{1}{2\sqrt{t}} \|\mathbf{y}\|_1.$$

Choosing $\mathbf{y} = \mathbf{x} - \mathbf{x}_{s/2}$ and $t = s/2$, we reach

$$\nu \leq \frac{1.71}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_{s/2}\|_1 + \|\mathbf{e}\|_2. \quad (2.2)$$

In words, the unrecoverable energy is controlled by the scaled ℓ_1 norm of the signal tail.

2.4. Implementation and Resource Requirements. CoSaMP was designed to be a practical method for signal recovery. An efficient implementation of the algorithm requires some ideas from numerical linear algebra, as well as some basic techniques from the theory of algorithms. This section discusses the key issues and develops an analysis of the running time for the two most common scenarios.

We focus on the least-squares problem in the estimation step because it is the major obstacle to a fast implementation of the algorithm. The algorithm guarantees that the matrix Φ_T never has more than $3s$ columns, so our assumption $\delta_{4s} \leq 0.1$ implies that the matrix Φ_T is extremely well conditioned. As a result, we can apply the pseudoinverse $\Phi_T^\dagger = (\Phi_T^* \Phi_T)^{-1} \Phi_T^*$ very quickly using an iterative method, such as Richardson's iteration [1, Sec. 7.2.3] or conjugate gradient [1, Sec. 7.4]. These techniques have the additional advantage that they only interact with the matrix Φ_T through its action on vectors. It follows that the algorithm performs better when the sampling matrix has a fast matrix–vector multiply.

Section 5 contains an analysis of the performance of iterative least-squares algorithms in the context of CoSaMP. In summary, if we initialize the least-squares method with the current approximation \mathbf{a}^{k-1} , then the cost of solving the least-squares problem is $O(\mathcal{L})$, where \mathcal{L} bounds the cost of a matrix–vector multiply with Φ_T or Φ_T^* . This implementation ensures that Theorem 2.1 holds at each iteration.

Let us stress that direct methods for least squares are likely to be extremely inefficient in this setting. The first reason is that each least-squares problem may contain substantially different sets of columns from Φ . As a result, it becomes necessary to perform a completely new QR or SVD factorization during each iteration at a cost of $O(s^2 m)$. The second problem is that computing these factorizations typically requires direct access to the columns of the matrix, which is problematic when the matrix is accessed through its action on vectors. Third, direct methods have storage costs $O(sm)$, which may be deadly for large-scale problems.

The remaining steps of the algorithm involve standard techniques. Let us estimate the operation counts.

Proxy: Forming the proxy is dominated by the cost of the matrix–vector multiply $\Phi^* \mathbf{v}$.

Identification: We can locate the largest $2s$ entries of a vector in time $O(N)$ using the approach in [8, Ch. 9]. In practice, it may be faster to sort the entries of the signal in decreasing order of magnitude at cost $O(N \log N)$ and then select the first $2s$ of them. The latter procedure can be accomplished with quicksort, mergesort, or heapsort [8, Sec. II]. To implement the algorithm to the letter, the sorting method needs to be stable because we stipulate that ties are broken lexicographically. This point is not important in practice.

Support Merger: We can merge two sets of size $O(s)$ in expected time $O(s)$ using randomized hashing methods [8, Ch. 11]. One can also sort both sets first and use the elementary merge procedure [8, p. 29] for a total cost $O(s \log s)$.

LS estimation: We use Richardson’s iteration or conjugate gradient to compute $\Phi_T^\dagger \mathbf{u}$. Initializing the least-squares algorithm requires a matrix–vector multiply with Φ_T^* . Each iteration of the least-squares method requires one matrix–vector multiply each with Φ_T and Φ_T^* . Since Φ_T is a submatrix of Φ , the matrix–vector multiplies can also be obtained from multiplication with the full matrix. We prove in Section 5 that a constant number of least-squares iterations suffices for Theorem 2.1 to hold.

Pruning: This step is similar to identification. Pruning can be implemented in time $O(s)$, but it may be preferable to sort the components of the vector by magnitude and then select the first s at a cost of $O(s \log s)$.

Sample Update: This step is dominated by the cost of the multiplication of Φ with the s -sparse vector \mathbf{a}^k .

Table 1 summarizes this discussion in two particular cases. The first column shows what happens when the sampling matrix Φ is applied to vectors in the standard way, but we have random access to submatrices. The second column shows what happens when the sampling matrix Φ and its adjoint Φ^* both have a fast multiply with cost \mathcal{L} , where we assume that $\mathcal{L} \geq N$. A typical value is $\mathcal{L} = O(N \log N)$. In particular, a partial Fourier matrix satisfies this bound.

TABLE 1. Operation count for CoSaMP. Big-O notation is omitted for legibility. The dimensions of the sampling matrix Φ are $m \times N$; the sparsity level is s . The number \mathcal{L} bounds the cost of a matrix–vector multiply with Φ or Φ^* .

Step	Standard multiply	Fast multiply
Form proxy	mN	\mathcal{L}
Identification	N	N
Support merger	s	s
LS estimation	sm	\mathcal{L}
Pruning	s	s
Sample update	sm	\mathcal{L}
Total per iteration	$O(mN)$	$O(\mathcal{L})$

Finally, we note that the storage requirements of the algorithm are also favorable. Aside from the storage required by the sampling matrix, the algorithm constructs only one vector of length N , the signal proxy. The sample vectors \mathbf{u} and \mathbf{v} have length m , so they require $O(m)$ storage. The signal approximations can be stored using sparse data structures, so they require at most $O(s \log N)$ storage. Similarly, the index sets that appear require only $O(s \log N)$ storage. The total storage is at worst $O(N)$.

The following result summarizes this discussion.

Theorem 2.4 (Resource Requirements). *Each iteration of CoSaMP requires $O(\mathcal{L})$ time, where \mathcal{L} bounds the cost of a multiplication with the matrix Φ or Φ^* . The algorithm uses storage $O(N)$.*

2.5. The Unrecoverable Energy. Since the unrecoverable energy ν plays a central role in our analysis of CoSaMP, it merits some additional discussion. In particular, it is informative to examine the unrecoverable energy in a compressible signal. Let p be a number in the interval $(0, 1)$. We say that \mathbf{x} is p -compressible with magnitude R if the sorted components of the signal decay at the rate

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

When $p = 1$, this definition implies that $\|\mathbf{x}\|_1 \leq R \cdot (1 + \log N)$. Therefore, the unit ball of 1-compressible signals is similar to the ℓ_1 unit ball. When $p \approx 0$, this definition implies that p -compressible signals are very nearly sparse. In general, p -compressible signals are well approximated by sparse signals:

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}_s\|_1 &\leq C_p \cdot R \cdot s^{1-1/p} \\ \|\mathbf{x} - \mathbf{x}_s\|_2 &\leq D_p \cdot R \cdot s^{1/2-1/p} \end{aligned}$$

where $C_p = (1/p - 1)^{-1}$ and $D_p = (2/p - 1)^{-1/2}$. These results follow by writing each norm as a sum and approximating the sum with an integral. We see that the unrecoverable energy (2.1) in a p -compressible signal is bounded as

$$\nu \leq 2C_p \cdot R \cdot s^{1/2-1/p} + \|\mathbf{e}\|_2. \quad (2.3)$$

When p is small, the first term in the unrecoverable energy decays rapidly as the sparsity level s increases. For the class of p -compressible signals, the bound (2.3) on the unrecoverable energy is sharp, modulo the exact values of the constants.

With these inequalities, we can see that CoSaMP recovers compressible signals efficiently. Let us calculate the number of iterations required to reduce the approximation error from $\|\mathbf{x}\|_2$ to the optimal level (2.3). For compressible signals, the energy $\|\mathbf{x}\|_2 \leq 2R$, so

$$\text{R-SNR} \leq \log \left(\frac{2R}{2C_p \cdot R \cdot s^{1/2-1/p}} \right) = \log(1/p - 1) + (1/p - 1/2) \log s.$$

Therefore, the number of iterations required to recover a generic p -compressible signal is $O(\log s)$, where the constant in the big- O notation depends on p .

The term ‘‘unrecoverable energy’’ is justified by several facts. First, we must pay for the ℓ_2 error contaminating the samples. To check this point, define $S = \text{supp}(\mathbf{x}_s)$. The matrix Φ_S is nearly an isometry from ℓ_2^S to ℓ_2^m , so an error in the large components of the signal induces an error of equivalent size in the samples. Clearly, we can never resolve this uncertainty.

The term $s^{-1/2} \|\mathbf{x} - \mathbf{x}_s\|_1$ is also required on account of classical results about the Gel’fand widths of the ℓ_1^N ball in ℓ_2^N , due to Kashin [25] and Garnaev–Gluskin [16]. In the language of compressive sampling, their work has the following interpretation. Let Φ be a fixed $m \times N$ sampling matrix. Suppose that, for every signal $\mathbf{x} \in \mathbb{C}^N$, there is an algorithm that uses the samples $\mathbf{u} = \Phi \mathbf{x}$ to construct an approximation \mathbf{a} that achieves the error bound

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq \frac{C}{\sqrt{s}} \|\mathbf{x}\|_1.$$

Then the number m of measurements must satisfy $m \geq cs \log(N/s)$.

3. RESTRICTED ISOMETRY CONSEQUENCES

When the sampling matrix satisfies the restricted isometry inequalities (1.1), it has several other properties that we require repeatedly in the proof that the CoSaMP algorithm is correct. Our first observation is a simple translation of (1.1) into other terms.

Proposition 3.1. *Suppose Φ has restricted isometry constant δ_r . Let T be a set of r indices or fewer. Then*

$$\begin{aligned} \|\Phi_T^* \mathbf{u}\|_2 &\leq \sqrt{1 + \delta_r} \|\mathbf{u}\|_2 \\ \|\Phi_T^\dagger \mathbf{u}\|_2 &\leq \frac{1}{\sqrt{1 - \delta_r}} \|\mathbf{u}\|_2 \\ \|\Phi_T^* \Phi_T \mathbf{x}\|_2 &\stackrel{\leq}{\geq} (1 \pm \delta_r) \|\mathbf{x}\|_2 \\ \|(\Phi_T^* \Phi_T)^{-1} \mathbf{x}\|_2 &\stackrel{\leq}{\geq} \frac{1}{1 \pm \delta_r} \|\mathbf{x}\|_2. \end{aligned}$$

where the last two statements contain an upper and lower bound, depending on the sign chosen.

Proof. The restricted isometry inequalities (1.1) imply that the singular values of Φ_T lie between $\sqrt{1 - \delta_r}$ and $\sqrt{1 + \delta_r}$. The bounds follow from standard relationships between the singular values of Φ_T and the singular values of basic functions of Φ_T . \square

A second consequence is that disjoint sets of columns from the sampling matrix span nearly orthogonal subspaces. The following result quantifies this observation.

Proposition 3.2 (Approximate Orthogonality). *Suppose Φ has restricted isometry constant δ_r . Let S and T be disjoint sets of indices whose combined cardinality does not exceed r . Then*

$$\|\Phi_S^* \Phi_T\| \leq \delta_r.$$

Proof. Abbreviate $R = S \cup T$, and observe that $\Phi_S^* \Phi_T$ is a submatrix of $\Phi_R^* \Phi_R - \mathbf{I}$. The spectral norm of a submatrix never exceeds the norm of the entire matrix. We discern that

$$\|\Phi_S^* \Phi_T\| \leq \|\Phi_R^* \Phi_R - \mathbf{I}\| \leq \max\{(1 + \delta_r) - 1, 1 - (1 - \delta_r)\} = \delta_r$$

because the eigenvalues of $\Phi_R^* \Phi_R$ lie between $1 - \delta_r$ and $1 + \delta_r$. \square

This result will be applied through the following corollary.

Corollary 3.3. *Suppose Φ has restricted isometry constant δ_r . Let T be a set of indices, and let \mathbf{x} be a vector. Provided that $r \geq |T \cup \text{supp}(\mathbf{x})|$,*

$$\|\Phi_T^* \Phi \cdot \mathbf{x}|_{T^c}\|_2 \leq \delta_r \|\mathbf{x}|_{T^c}\|_2.$$

Proof. Define $S = \text{supp}(\mathbf{x}) \setminus T$, so we have $\mathbf{x}|_S = \mathbf{x}|_{T^c}$. Thus,

$$\|\Phi_T^* \Phi \cdot \mathbf{x}|_{T^c}\|_2 = \|\Phi_T^* \Phi \cdot \mathbf{x}|_S\|_2 \leq \|\Phi_T^* \Phi_S\| \|\mathbf{x}|_S\|_2 \leq \delta_r \|\mathbf{x}|_{T^c}\|_2,$$

owing to Proposition 3.2. \square

As a second corollary, we show that δ_{2r} gives weak control over the higher restricted isometry constants.

Corollary 3.4. *Let c and r be positive integers. Then $\delta_{cr} \leq c \cdot \delta_{2r}$.*

Proof. The result is clearly true for $c = 1, 2$, so we assume $c \geq 3$. Let S be an arbitrary index set of size cr , and let $\mathbf{M} = \Phi_S^* \Phi_S - \mathbf{I}$. It suffices to check that $\|\mathbf{M}\| \leq c \cdot \delta_{2r}$. To that end, we break the matrix \mathbf{M} into $r \times r$ blocks, which we denote \mathbf{M}_{ij} . A block version of Gershgorin's theorem states that $\|\mathbf{M}\|$ satisfies at least one of the inequalities

$$\|\mathbf{M}\| - \|\mathbf{M}_{ii}\| \leq \sum_{j \neq i} \|\mathbf{M}_{ij}\| \quad \text{where } i = 1, 2, \dots, c.$$

The derivation is entirely analogous with the usual proof of Gershgorin's theorem, so we omit the details. For each diagonal block, we have $\|\mathbf{M}_{ii}\| \leq \delta_r$ because of the restricted isometry inequalities (1.1). For each off-diagonal block, we have $\|\mathbf{M}_{ij}\| \leq \delta_{2r}$ because of Proposition 3.2. Substitute these bounds into the block Gershgorin theorem and rearrange to complete the proof. \square

Finally, we present a result that measures how much the sampling matrix inflates nonsparse vectors. This bound permits us to establish the major results for sparse signals and then transfer the conclusions to the general case.

Proposition 3.5 (Energy Bound). *Suppose that Φ verifies the upper inequality of (1.1), viz.*

$$\|\Phi \mathbf{x}\|_2 \leq \sqrt{1 + \delta_r} \|\mathbf{x}\|_2 \quad \text{when} \quad \|\mathbf{x}\|_0 \leq r.$$

Then, for every signal \mathbf{x} ,

$$\|\Phi \mathbf{x}\|_2 \leq \sqrt{1 + \delta_r} \left[\|\mathbf{x}\|_2 + \frac{1}{\sqrt{r}} \|\mathbf{x}\|_1 \right].$$

Proof. First, observe that the hypothesis of the proposition can be regarded as a statement about the operator norm of Φ as a map between two Banach spaces. For a set $I \subset \{1, 2, \dots, N\}$, write B_2^I for the unit ball in $\ell_2(I)$. Define the convex body

$$S = \text{conv} \left\{ \bigcup_{|I| \leq r} B_2^I \right\} \subset \mathbb{C}^N,$$

and notice that, by hypothesis, the operator norm

$$\|\Phi\|_{S \rightarrow 2} = \max_{\mathbf{x} \in S} \|\Phi \mathbf{x}\|_2 \leq \sqrt{1 + \delta_r}.$$

Define a second convex body

$$K = \left\{ \mathbf{x} : \|\mathbf{x}\|_2 + \frac{1}{\sqrt{r}} \|\mathbf{x}\|_1 \leq 1 \right\} \subset \mathbb{C}^N,$$

and consider the operator norm

$$\|\Phi\|_{K \rightarrow 2} = \max_{\mathbf{x} \in K} \|\Phi \mathbf{x}\|_2.$$

The content of the proposition is the claim that

$$\|\Phi\|_{K \rightarrow 2} \leq \|\Phi\|_{S \rightarrow 2}.$$

To establish this point, it suffices to check that $K \subset S$.

Choose a vector $\mathbf{x} \in K$. We partition the support of \mathbf{x} into sets of size r . Let I_0 index the r largest-magnitude components of \mathbf{x} , breaking ties lexicographically. Let I_1 index the next largest r components, and so forth. Note that the final block I_J may have fewer than r components. We may assume that $\mathbf{x}|_{I_j}$ is nonzero for each j .

This partition induces a decomposition

$$\mathbf{x} = \mathbf{x}|_{I_0} + \sum_{j=0}^J \mathbf{x}|_{I_j} = \lambda_0 \mathbf{y}_0 + \sum_{j=0}^J \lambda_j \mathbf{y}_j$$

where

$$\lambda_j = \|\mathbf{x}|_{I_j}\|_2 \quad \text{and} \quad \mathbf{y}_j = \lambda_j^{-1} \mathbf{x}|_{I_j}.$$

By construction, each vector \mathbf{y}_j belongs to S because it is r -sparse and has unit ℓ_2 norm. We will prove that $\sum_j \lambda_j \leq 1$, which implies that \mathbf{x} can be written as a convex combination of vectors from the set S . As a consequence, $\mathbf{x} \in S$. It emerges that $K \subset S$.

Fix j in the range $\{1, 2, \dots, J\}$. It follows that I_j contains at most r elements and I_{j-1} contains exactly r elements. Therefore,

$$\lambda_j = \|\mathbf{x}|_{I_j}\|_2 \leq \sqrt{r} \|\mathbf{x}|_{I_j}\|_\infty \leq \sqrt{r} \cdot \frac{1}{r} \|\mathbf{x}|_{I_{j-1}}\|_1$$

where the last inequality holds because the magnitude of \mathbf{x} on the set I_{j-1} dominates its largest entry in I_j . Summing these relations, we obtain

$$\sum_{j=1}^J \lambda_j \leq \frac{1}{\sqrt{r}} \sum_{j=1}^J \|\mathbf{x}|_{I_{j-1}}\|_1 = \frac{1}{\sqrt{r}} \|\mathbf{x}\|_1.$$

It is clear that $\lambda_0 = \|\mathbf{x}|_{I_0}\|_2 \leq \|\mathbf{x}\|_2$. We may conclude that

$$\sum_{j=0}^J \lambda_j \leq \|\mathbf{x}\|_2 + \frac{1}{\sqrt{r}} \|\mathbf{x}\|_1 \leq 1$$

because $\mathbf{x} \in K$. □

4. THE ITERATION INVARIANT: SPARSE CASE

We now commence the proof of Theorem 2.1. For the moment, let us assume that the signal is actually sparse. Section 6 removes this assumption.

The result states that each iteration of the algorithm reduces the approximation error by a constant factor, while adding a small multiple of the noise. As a consequence, when the approximation error is large in comparison with the noise, the algorithm makes substantial progress in identifying the unknown signal.

Theorem 4.1 (Iteration Invariant: Sparse Case). *Assume that \mathbf{x} is s -sparse. For each $k \geq 0$, the signal approximation \mathbf{a}^k is s -sparse, and*

$$\|\mathbf{x} - \mathbf{a}^{k+1}\|_2 \leq 0.5 \|\mathbf{x} - \mathbf{a}^k\|_2 + 7.5 \|\mathbf{e}\|_2.$$

In particular,

$$\|\mathbf{x} - \mathbf{a}^k\|_2 \leq 2^{-k} \|\mathbf{x}\|_2 + 15 \|\mathbf{e}\|_2.$$

The argument proceeds in a sequence of short lemmas, each corresponding to one step in the algorithm. Throughout this section, we retain the assumption that \mathbf{x} is s -sparse.

4.1. Approximations, Residuals, etc. Fix an iteration $k \geq 1$. We write $\mathbf{a} = \mathbf{a}^{k-1}$ for the signal approximation at the beginning of the iteration. Define the residual $\mathbf{r} = \mathbf{x} - \mathbf{a}$, which we interpret as the part of the signal we have not yet recovered. Since the approximation \mathbf{a} is always s -sparse, the residual \mathbf{r} must be $2s$ -sparse. Notice that the vector \mathbf{v} of updated samples can be viewed as noisy samples of the residual:

$$\mathbf{v} \stackrel{\text{def}}{=} \mathbf{u} - \Phi \mathbf{a} = \Phi(\mathbf{x} - \mathbf{a}) + \mathbf{e} = \Phi \mathbf{r} + \mathbf{e}.$$

4.2. Identification. The identification phase produces a set of components where the residual signal still has a lot of energy.

Lemma 4.2 (Identification). *The set $\Omega = \text{supp}(\mathbf{y}_{2s})$, where $\mathbf{y} = \Phi^* \mathbf{v}$ is the signal proxy, contains at most $2s$ indices, and*

$$\|\mathbf{r}|_{\Omega^c}\|_2 \leq 0.2223 \|\mathbf{r}\|_2 + 2.34 \|\mathbf{e}\|_2.$$

Proof. The identification phase forms a proxy $\mathbf{y} = \Phi^* \mathbf{v}$ for the residual signal. The algorithm then selects a set Ω of $2s$ components from \mathbf{y} that have the largest magnitudes. The goal of the proof is to show that the energy in the residual on the set Ω^c is small in comparison with the total energy in the residual.

Define the set $R = \text{supp}(\mathbf{r})$. Since R contains at most $2s$ elements, our choice of Ω ensures that $\|\mathbf{y}|_R\|_2 \leq \|\mathbf{y}|_\Omega\|_2$. By squaring this inequality and canceling the terms in $R \cap \Omega$, we discover that

$$\|\mathbf{y}|_{R \setminus \Omega}\|_2 \leq \|\mathbf{y}|_{\Omega \setminus R}\|_2.$$

Since the coordinate subsets here contain few elements, we can use the restricted isometry constants to provide bounds on both sides.

First, observe that the set $\Omega \setminus R$ contains at most $2s$ elements. Therefore, we may apply Proposition 3.1 and Corollary 3.3 to obtain

$$\begin{aligned} \|\mathbf{y}|_{\Omega \setminus R}\|_2 &= \|\Phi_{\Omega \setminus R}^*(\Phi \mathbf{r} + \mathbf{e})\|_2 \\ &\leq \|\Phi_{\Omega \setminus R}^* \Phi \mathbf{r}\|_2 + \|\Phi_{\Omega \setminus R}^* \mathbf{e}\|_2 \\ &\leq \delta_{4s} \|\mathbf{r}\|_2 + \sqrt{1 + \delta_{2s}} \|\mathbf{e}\|_2. \end{aligned}$$

Likewise, the set $R \setminus \Omega$ contains $2s$ elements or fewer, so Proposition 3.1 and Corollary 3.3 yield

$$\begin{aligned} \|\mathbf{y}|_{R \setminus \Omega}\|_2 &= \|\Phi_{R \setminus \Omega}^*(\Phi \mathbf{r} + \mathbf{e})\|_2 \\ &\geq \|\Phi_{R \setminus \Omega}^* \Phi \cdot \mathbf{r}|_{R \setminus \Omega}\|_2 - \|\Phi_{R \setminus \Omega}^* \Phi \cdot \mathbf{r}|_\Omega\|_2 - \|\Phi_{R \setminus \Omega}^* \mathbf{e}\|_2 \\ &\geq (1 - \delta_{2s}) \|\mathbf{r}|_{R \setminus \Omega}\|_2 - \delta_{2s} \|\mathbf{r}\|_2 - \sqrt{1 + \delta_{2s}} \|\mathbf{e}\|_2. \end{aligned}$$

Since the residual is supported on R , we can rewrite $\mathbf{r}|_{R \setminus \Omega} = \mathbf{r}|_{\Omega^c}$. Finally, combine the last three inequalities and rearrange to obtain

$$\|\mathbf{r}|_{\Omega^c}\|_2 \leq \frac{(\delta_{2s} + \delta_{4s}) \|\mathbf{r}\|_2 + 2\sqrt{1 + \delta_{2s}} \|\mathbf{e}\|_2}{1 - \delta_{2s}}.$$

Invoke the numerical hypothesis that $\delta_{2s} \leq \delta_{4s} \leq 0.1$ to complete the argument. \square

4.3. Support Merger. The next step of the algorithm merges the support of the current signal approximation \mathbf{a} with the newly identified set of components. The following result shows that components of the signal \mathbf{x} outside this set have very little energy.

Lemma 4.3 (Support Merger). *Let Ω be a set of at most $2s$ indices. The set $T = \Omega \cup \text{supp}(\mathbf{a})$ contains at most $3s$ indices, and*

$$\|\mathbf{x}|_{T^c}\|_2 \leq \|\mathbf{r}|_{\Omega^c}\|_2.$$

Proof. Since $\text{supp}(\mathbf{a}) \subset T$, we find that

$$\|\mathbf{x}|_{T^c}\|_2 = \|(\mathbf{x} - \mathbf{a})|_{T^c}\|_2 = \|\mathbf{r}|_{T^c}\|_2 \leq \|\mathbf{r}|_{\Omega^c}\|_2,$$

where the inequality follows from the containment $T^c \subset \Omega^c$. \square

4.4. Estimation. The estimation step of the algorithm solves a least-squares problem to obtain values for the coefficients in the set T . We need a bound on the error of this approximation.

Lemma 4.4 (Estimation). *Let T be a set of at most $3s$ indices, and define the least-squares signal estimate \mathbf{b} by the formulae*

$$\mathbf{b}|_T = \Phi_T^\dagger \mathbf{u} \quad \text{and} \quad \mathbf{b}|_{T^c} = \mathbf{0},$$

where $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$. Then

$$\|\mathbf{x} - \mathbf{b}\|_2 \leq 1.112 \|\mathbf{x}|_{T^c}\|_2 + 1.06 \|\mathbf{e}\|_2.$$

This result assumes that we solve the least-squares problem in infinite precision. In practice, the right-hand side of the bound contains an extra term owing to the error from the iterative least-squares solver. In Section 5, we study how many iterations of the least-squares solver are required to make the least-squares error negligible in the present argument.

Proof. Note first that

$$\|\mathbf{x} - \mathbf{b}\|_2 \leq \|\mathbf{x}|_{T^c}\|_2 + \|\mathbf{x}|_T - \mathbf{b}|_T\|_2.$$

Using the expression $\mathbf{u} = \Phi\mathbf{x} + \mathbf{e}$ and the fact $\Phi_T^\dagger\Phi_T = \mathbf{I}_T$, we calculate that

$$\begin{aligned} \|\mathbf{x}|_T - \mathbf{b}|_T\|_2 &= \|\mathbf{x}|_T - \Phi_T^\dagger(\Phi \cdot \mathbf{x}|_T + \Phi \cdot \mathbf{x}|_{T^c} + \mathbf{e})\|_2 \\ &= \|\Phi_T^\dagger(\Phi \cdot \mathbf{x}|_{T^c} + \mathbf{e})\|_2 \\ &\leq \|(\Phi_T^*\Phi_T)^{-1}\Phi_T^*\Phi \cdot \mathbf{x}|_{T^c}\|_2 + \|\Phi_T^\dagger\mathbf{e}\|_2. \end{aligned}$$

The cardinality of T is at most $3s$, and \mathbf{x} is s -sparse, so Proposition 3.1 and Corollary 3.3 imply that

$$\begin{aligned} \|\mathbf{x}|_T - \mathbf{b}|_T\|_2 &\leq \frac{1}{1 - \delta_{3s}} \|\Phi_T^*\Phi \cdot \mathbf{x}|_{T^c}\|_2 + \frac{1}{\sqrt{1 - \delta_{3s}}} \|\mathbf{e}\|_2 \\ &\leq \frac{\delta_{4s}}{1 - \delta_{3s}} \|\mathbf{x}|_{T^c}\|_2 + \frac{\|\mathbf{e}\|_2}{\sqrt{1 - \delta_{3s}}}. \end{aligned}$$

Combine the bounds to reach

$$\|\mathbf{x} - \mathbf{b}\|_2 \leq \left[1 + \frac{\delta_{4s}}{1 - \delta_{3s}}\right] \|\mathbf{x}|_{T^c}\|_2 + \frac{\|\mathbf{e}\|_2}{\sqrt{1 - \delta_{3s}}}.$$

Finally, invoke the hypothesis that $\delta_{3s} \leq \delta_{4s} \leq 0.1$. \square

4.5. Pruning. The final step of each iteration is to prune the intermediate approximation to its largest s terms. The following lemma provides a bound on the error in the pruned approximation.

Lemma 4.5 (Pruning). *The pruned approximation \mathbf{b}_s satisfies*

$$\|\mathbf{x} - \mathbf{b}_s\|_2 \leq 2\|\mathbf{x} - \mathbf{b}\|_2.$$

Proof. The intuition is that \mathbf{b}_s is close to \mathbf{b} , which is close to \mathbf{x} . Rigorously,

$$\|\mathbf{x} - \mathbf{b}_s\|_2 \leq \|\mathbf{x} - \mathbf{b}\|_2 + \|\mathbf{b} - \mathbf{b}_s\|_2 \leq 2\|\mathbf{x} - \mathbf{b}\|_2.$$

The second inequality holds because \mathbf{b}_s is the best s -sparse approximation to \mathbf{b} . In particular, the s -sparse vector \mathbf{x} is a worse approximation. \square

4.6. Proof of Theorem 4.1. We now complete the proof of the iteration invariant for sparse signals, Theorem 4.1. At the end of an iteration, the algorithm forms a new approximation $\mathbf{a}^k = \mathbf{b}_s$, which is evidently s -sparse. Applying the lemmas we have established, we easily bound the error:

$$\begin{aligned} \|\mathbf{x} - \mathbf{a}^k\|_2 &= \|\mathbf{x} - \mathbf{b}_s\|_2 \\ &\leq 2\|\mathbf{x} - \mathbf{b}\|_2 && \text{Pruning (Lemma 4.5)} \\ &\leq 2 \cdot (1.112\|\mathbf{x}|_{T^c}\|_2 + 1.06\|\mathbf{e}\|_2) && \text{Estimation (Lemma 4.4)} \\ &\leq 2.224\|\mathbf{r}|_{\Omega^c}\|_2 + 2.12\|\mathbf{e}\|_2 && \text{Support merger (Lemma 4.3)} \\ &\leq 2.224 \cdot (0.2223\|\mathbf{r}\|_2 + 2.34\|\mathbf{e}\|_2) + 2.12\|\mathbf{e}\|_2 && \text{Identification (Lemma 4.2)} \\ &< 0.5\|\mathbf{r}\|_2 + 7.5\|\mathbf{e}\|_2 \\ &= 0.5\|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 7.5\|\mathbf{e}\|_2. \end{aligned}$$

To obtain the second bound in Theorem 4.1, simply solve the error recursion and note that

$$(1 + 0.5 + 0.25 + \dots) \cdot 7.5\|\mathbf{e}\|_2 \leq 15\|\mathbf{e}\|_2.$$

This point completes the argument.

5. ANALYSIS OF ITERATIVE LEAST-SQUARES

To develop an efficient implementation of CoSaMP, it is critical to use an iterative method when we solve the least-squares problem in the estimation step. The two natural choices are Richardson's iteration and conjugate gradient. The efficacy of these methods rests on the assumption that the sampling operator has small restricted isometry constants. Indeed, since the set T constructed in the support merger step contains at most $3s$ components, the hypothesis $\delta_{4s} \leq 0.1$ ensures that the condition number

$$\kappa(\Phi_T^* \Phi_T) = \frac{\lambda_{\max}(\Phi_T^* \Phi_T)}{\lambda_{\min}(\Phi_T^* \Phi_T)} \leq \frac{1 + \delta_{3s}}{1 - \delta_{3s}} < 1.223.$$

This condition number is closely connected with the performance of Richardson's iteration and conjugate gradient. In this section, we show that Theorem 4.1 holds if we perform a constant number of iterations of either least-squares algorithm.

5.1. Richardson's Iteration. For completeness, let us explain how Richardson's iteration can be applied to solve the least-squares problems that arise in CoSaMP. Suppose we wish to compute $\mathbf{A}^\dagger \mathbf{u}$ where \mathbf{A} is a tall, full-rank matrix. Recalling the definition of the pseudoinverse, we realize that this amounts to solving a linear system of the form

$$(\mathbf{A}^* \mathbf{A}) \mathbf{b} = \mathbf{A}^* \mathbf{u}.$$

This problem can be approached by *splitting* the Gram matrix:

$$\mathbf{A}^* \mathbf{A} = \mathbf{I} + \mathbf{M}$$

where $\mathbf{M} = \mathbf{A}^* \mathbf{A} - \mathbf{I}$. Given an initial iterate \mathbf{z}^0 , Richardson's method produces the subsequent iterates via the formula

$$\mathbf{z}^{\ell+1} = \mathbf{A}^* \mathbf{u} - \mathbf{M} \mathbf{z}^\ell.$$

Evidently, this iteration requires only matrix-vector multiplies with \mathbf{A} and \mathbf{A}^* . It is worth noting that Richardson's method can be accelerated [1, Sec. 7.2.5], but we omit the details.

It is quite easy to analyze Richardson's iteration [1, Sec. 7.2.1]. Observe that

$$\|\mathbf{z}^{\ell+1} - \mathbf{A}^\dagger \mathbf{u}\|_2 = \|\mathbf{M}(\mathbf{z}^\ell - \mathbf{A}^\dagger \mathbf{u})\|_2 \leq \|\mathbf{M}\| \|\mathbf{z}^\ell - \mathbf{A}^\dagger \mathbf{u}\|_2.$$

This recursion delivers

$$\|\mathbf{z}^\ell - \mathbf{A}^\dagger \mathbf{u}\|_2 \leq \|\mathbf{M}\|^\ell \|\mathbf{z}^0 - \mathbf{A}^\dagger \mathbf{u}\|_2 \quad \text{for } \ell = 0, 1, 2, \dots$$

In words, the iteration converges linearly.

In our setting, $\mathbf{A} = \Phi_T$ where T is a set of at most $3s$ indices. Therefore, the restricted isometry inequalities (1.1) imply that

$$\|\mathbf{M}\| = \|\Phi_T^* \Phi_T - \mathbf{I}\| \leq \delta_{3s}.$$

We have assumed that $\delta_{3s} \leq \delta_{4s} \leq 0.1$, which means that the iteration converges quite fast. Once again, the restricted isometry behavior of the sampling matrix plays an essential role in the performance of the CoSaMP algorithm.

Conjugate gradient provides even better guarantees for solving the least-squares problem, but it is somewhat more complicated to describe and rather more difficult to analyze. We refer the reader to [1, Sec. 7.4] for more information. The following lemma summarizes the behavior of both Richardson's iteration and conjugate gradient in our setting.

Lemma 5.1 (Error Bound for LS). *Richardson's iteration produces a sequence $\{\mathbf{z}^\ell\}$ of iterates that satisfy*

$$\|\mathbf{z}^\ell - \Phi_T^\dagger \mathbf{u}\|_2 \leq 0.1^\ell \cdot \|\mathbf{z}^0 - \Phi_T^\dagger \mathbf{u}\|_2 \quad \text{for } \ell = 0, 1, 2, \dots$$

Conjugate gradient produces a sequence of iterates that satisfy

$$\|\mathbf{z}^\ell - \Phi_T^\dagger \mathbf{u}\|_2 \leq 2 \cdot \rho^\ell \cdot \|\mathbf{z}^0 - \Phi_T^\dagger \mathbf{u}\|_2 \quad \text{for } \ell = 0, 1, 2, \dots$$

where

$$\rho = \frac{\sqrt{\kappa(\Phi_T^* \Phi_T)} - 1}{\sqrt{\kappa(\Phi_T^* \Phi_T)} + 1} \leq 0.072.$$

5.2. Initialization. Iterative least-squares algorithms must be seeded with an initial iterate, and their performance depends heavily on a wise selection thereof. CoSaMP offers a natural choice for the initializer: the current signal approximation. As the algorithm progresses, the current signal approximation provides an increasingly good starting point for solving the least-squares problem. The following shows that the error in the initial iterate is controlled by the current approximation error.

Lemma 5.2 (Initial Iterate for LS). *Let \mathbf{x} be an s -sparse signal with noisy samples $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$. Let \mathbf{a}^{k-1} be the signal approximation at the end of the $(k-1)$ th iteration, and let T be the set of components identified by the support merger. Then*

$$\|\mathbf{a}^{k-1} - \Phi_T^\dagger \mathbf{u}\|_2 \leq 2.112 \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 1.06 \|\mathbf{e}\|_2$$

Proof. By construction of T , the approximation \mathbf{a}^{k-1} is supported inside T , so

$$\|\mathbf{x}|_{T^c}\|_2 = \|(\mathbf{x} - \mathbf{a}^{k-1})|_{T^c}\|_2 \leq \|\mathbf{x} - \mathbf{a}^{k-1}\|_2.$$

Using Lemma 4.4, we may calculate how far \mathbf{a}^{k-1} lies from the solution to the least-squares problem.

$$\begin{aligned} \|\mathbf{a}^{k-1} - \Phi_T^\dagger \mathbf{u}\|_2 &\leq \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + \|\mathbf{x} - \Phi_T^\dagger \mathbf{u}\|_2 \\ &\leq \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 1.112 \|\mathbf{x}|_{T^c}\|_2 + 1.06 \|\mathbf{e}\|_2 \\ &\leq 2.112 \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 1.06 \|\mathbf{e}\|_2. \end{aligned}$$

□

5.3. Iteration Count. We need to determine how many iterations of the least-squares algorithm are required to ensure that the approximation produced is sufficiently good to support the performance of CoSaMP.

Corollary 5.3 (Estimation by Iterative LS). *Suppose that we initialize the LS algorithm with $\mathbf{z}^0 = \mathbf{a}^{k-1}$. After at most three iterations, both Richardson's iteration and conjugate gradient produce a signal estimate \mathbf{b} that satisfies*

$$\|\mathbf{x} - \mathbf{b}\|_2 \leq 1.112 \|\mathbf{x}|_{T^c}\|_2 + 0.0022 \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 1.062 \|\mathbf{e}\|_2.$$

Proof. Combine Lemma 5.1 and Lemma 5.2 to see that three iterations of Richardson's method yield

$$\|\mathbf{z}^3 - \Phi_T^\dagger \mathbf{u}\|_2 \leq 0.002112 \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 0.00106 \|\mathbf{e}\|_2.$$

The bound for conjugate gradient is slightly better. Let $\mathbf{b}|_T = \mathbf{z}^3$. According to the estimation result, Lemma 4.4, we have

$$\|\mathbf{x} - \Phi_T^\dagger \mathbf{u}\|_2 \leq 1.112 \|\mathbf{x}|_{T^c}\|_2 + 1.06 \|\mathbf{e}\|_2.$$

An application of the triangle inequality completes the argument. □

5.4. CoSaMP with Iterative least-squares. Finally, we need to check that the sparse iteration invariant, Theorem 4.1 still holds when we use an iterative least-squares algorithm.

Theorem 5.4 (Sparse Iteration Invariant with Iterative LS). *Suppose that we use Richardson's iteration or conjugate gradient for the estimation step, initializing the LS algorithm with the current approximation \mathbf{a}^{k-1} and performing three LS iterations. Then Theorem 4.1 still holds.*

Proof. We repeat the calculation in Section 4.6 using Corollary 5.3 instead of the simple estimation lemma. To that end, recall that the residual $\mathbf{r} = \mathbf{x} - \mathbf{a}^{k-1}$. Then

$$\begin{aligned} \|\mathbf{x} - \mathbf{a}^k\|_2 &\leq 2 \|\mathbf{x} - \mathbf{b}\|_2 \\ &\leq 2 \cdot (1.112 \|\mathbf{x}|_{T^c}\|_2 + 0.0022 \|\mathbf{r}\|_2 + 1.062 \|\mathbf{e}\|_2) \\ &\leq 2.224 \|\mathbf{r}|_{\Omega^c}\|_2 + 0.0044 \|\mathbf{r}\|_2 + 2.124 \|\mathbf{e}\|_2 \\ &\leq 2.224 \cdot (0.2223 \|\mathbf{r}\|_2 + 2.34 \|\mathbf{e}\|_2) + 0.0044 \|\mathbf{r}\|_2 + 2.124 \|\mathbf{e}\|_2 \\ &< 0.5 \|\mathbf{r}\|_2 + 7.5 \|\mathbf{e}\|_2 \\ &= 0.5 \|\mathbf{x} - \mathbf{a}^{k-1}\|_2 + 7.5 \|\mathbf{e}\|_2. \end{aligned}$$

This bound is precisely what is required for the theorem to hold. \square

6. EXTENSION TO GENERAL SIGNALS

In this section, we finally complete the proof of the main result for CoSaMP, Theorem 2.1. The remaining challenge is to remove the hypothesis that the target signal is sparse, which we framed in Theorems 4.1 and 5.4. Although this difficulty might seem large, the solution is simple and elegant. It turns out that we can view the noisy samples of a general signal as samples of a sparse signal contaminated with a different noise vector that implicitly reflects the tail of the original signal.

Lemma 6.1 (Reduction to Sparse Case). *Let \mathbf{x} be an arbitrary vector in \mathbb{C}^N . The sample vector $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$ can also be expressed as $\mathbf{u} = \Phi \mathbf{x}_s + \tilde{\mathbf{e}}$ where*

$$\|\tilde{\mathbf{e}}\|_2 \leq 1.05 \left[\|\mathbf{x} - \mathbf{x}_s\|_2 + \frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 \right] + \|\mathbf{e}\|_2.$$

Proof. Decompose $\mathbf{x} = \mathbf{x}_s + (\mathbf{x} - \mathbf{x}_s)$ to obtain $\mathbf{u} = \Phi \mathbf{x}_s + \tilde{\mathbf{e}}$ where $\tilde{\mathbf{e}} = \Phi(\mathbf{x} - \mathbf{x}_s) + \mathbf{e}$. To compute the size of the error term, we simply apply the triangle inequality and Proposition 3.5:

$$\|\tilde{\mathbf{e}}\|_2 \leq \sqrt{1 + \delta_s} \left[\|\mathbf{x} - \mathbf{x}_s\|_2 + \frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 \right] + \|\mathbf{e}\|_2.$$

Finally, invoke the fact that $\delta_s \leq \delta_{4s} \leq 0.1$ to obtain $\sqrt{1 + \delta_s} \leq 1.05$. \square

This lemma is just the tool we require to complete the proof of Theorem 2.1.

Proof of Theorem 2.1. Let \mathbf{x} be a general signal, and use Lemma 6.1 to write the noisy vector of samples $\mathbf{u} = \Phi \mathbf{x}_s + \tilde{\mathbf{e}}$. Apply the sparse iteration invariant, Theorem 4.1, or the analog for iterative least-squares, Theorem 5.4. We obtain

$$\|\mathbf{x}_s - \mathbf{a}^{k+1}\|_2 \leq 0.5 \|\mathbf{x}_s - \mathbf{a}^k\|_2 + 7.5 \|\tilde{\mathbf{e}}\|_2.$$

Invoke the lower and upper triangle inequalities to obtain

$$\|\mathbf{x} - \mathbf{a}^{k+1}\|_2 \leq 0.5 \|\mathbf{x} - \mathbf{a}^k\|_2 + 7.5 \|\tilde{\mathbf{e}}\|_2 + 1.5 \|\mathbf{x} - \mathbf{x}_s\|_2.$$

Finally, recall the estimate for $\|\tilde{\mathbf{e}}\|_2$ from Lemma 6.1, and simplify to reach

$$\begin{aligned} \|\mathbf{x} - \mathbf{a}^{k+1}\|_2 &\leq 0.5 \|\mathbf{x} - \mathbf{a}^k\|_2 + 9.375 \|\mathbf{x} - \mathbf{x}_s\|_2 + \frac{7.875}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 + 7.5 \|\mathbf{e}\|_2 \\ &< 0.5 \|\mathbf{x} - \mathbf{a}^k\|_2 + 10\nu. \end{aligned}$$

where ν is the unrecoverable energy (2.1). \square

6.1. Proof of Theorem A. We have now collected all the material we need to establish the main result. Fix a precision parameter η . After at most $O(\log(\|\mathbf{x}\|_2/\eta))$ iterations, CoSaMP produces an s -sparse approximation \mathbf{a} that satisfies

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq C \cdot (\eta + \nu)$$

in consequence of Theorem 2.1. Apply inequality (2.2) to bound the unrecoverable energy ν in terms of the ℓ_1 norm. We see that the approximation error satisfies

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

According to Theorem 2.4, each iteration of CoSaMP is completed in time $O(\mathcal{L})$, where \mathcal{L} bounds the cost of a matrix–vector multiplication with Φ or Φ^* . The total runtime, therefore, is $O(\mathcal{L} \log(\|\mathbf{x}\|_2/\eta))$. The total storage is $O(N)$.

Finally, in the statement of the theorem, we replace δ_{4s} with δ_{2s} by means of Corollary 3.4, which states that $\delta_{cr} \leq c \cdot \delta_{2r}$ for any positive integers c and r .

7. DISCUSSION AND RELATED WORK

CoSaMP draws on both algorithmic ideas and analytic techniques that have appeared before. This section describes the other major signal recovery algorithms, and it compares them with CoSaMP. It also attempts to trace the key ideas in the algorithm back to their sources.

7.1. Algorithms for Compressive Sampling. We begin with a short discussion of the major algorithmic approaches to signal recovery from compressive samples. We focus on provably correct methods, although we acknowledge that some *ad hoc* techniques provide excellent empirical results.

The initial discovery works on compressive sampling proposed to perform signal recovery by solving a convex optimization problem [2, 12]. Given a sampling matrix Φ and a noisy vector of samples $\mathbf{u} = \Phi\mathbf{x} + \mathbf{e}$ with $\|\mathbf{e}\|_2 \leq \varepsilon$, consider the mathematical program

$$\min \|\mathbf{y}\|_1 \quad \text{subject to} \quad \|\Phi\mathbf{y} - \mathbf{u}\|_2 \leq \varepsilon. \quad (7.1)$$

In words, we look for a signal reconstruction that is consistent with the samples but has minimal ℓ_1 norm. The intuition behind this approach is that minimizing the ℓ_1 norm promotes sparsity, so allows the approximate recovery of compressible signals. Candès, Romberg, and Tao established in [3] that a minimizer \mathbf{a} of (7.1) satisfies

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq C \left[\frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 + \varepsilon \right] \quad (7.2)$$

provided that the sampling matrix Φ has restricted isometry constant $\delta_{4s} \leq 0.2$. In [4], the hypothesis on the restricted isometry constant is sharpened to $\delta_{2s} \leq \sqrt{2} - 1$. The error bound for CoSaMP is equivalent, modulo the exact value of the constants.

The literature describes a huge variety of algorithms for solving the optimization problem (7.1). The most common approaches involve interior-point methods [2, 26], projected gradient methods [15], or iterative thresholding [11]. The interior-point methods are guaranteed to solve the problem to a fixed precision in time $O(m^2 N^{1.5})$, where m is the number of measurements and N is the signal length [32]. Note that the constant in the big-O notation depends on some of the problem data. The other convex relaxation algorithms, while sometimes faster in practice, do not currently offer rigorous guarantees. CoSaMP provides rigorous bounds on the runtime that are much better than the available results for interior-point methods.

Tropp and Gilbert proposed the use of a greedy iterative algorithm called *orthogonal matching pursuit* (OMP) for signal recovery [39]. The algorithm initializes the current sample vector $\mathbf{v} = \mathbf{u}$. In each iteration, it forms the signal proxy $\mathbf{y} = \Phi^* \mathbf{v}$ and identifies a component of the proxy with

largest magnitude. It adds the new component to the set T of previously identified components. Then OMP forms a new signal approximation by solving a least-squares problem: $\mathbf{a} = \Phi_T^\dagger \mathbf{u}$. Finally, it updates the samples $\mathbf{v} = \mathbf{u} - \Phi \mathbf{a}$. These steps are repeated until a halting criterion is satisfied.

Tropp and Gilbert were able to prove a weak result for the performance of OMP [39]. Suppose that \mathbf{x} is a fixed, s -sparse signal, and let $m = Cs \log N$. Draw an $m \times N$ sampling matrix Φ whose entries are independent, zero-mean subgaussian² random variables with equal variances. Given noiseless measurements $\mathbf{u} = \Phi \mathbf{x}$, OMP reconstructs \mathbf{x} after s iterations, except with probability N^{-1} . In this setting, OMP must fail for some sparse signals [33], so it does not provide the same uniform guarantees as convex relaxation. It is unknown whether OMP succeeds for compressible signals or whether it succeeds when the samples are contaminated with noise.

Donoho et al. invented another greedy iterative method called *stagewise OMP*, or StOMP [14]. This algorithm uses the signal proxy to select multiple components at each step, using a rule inspired by ideas from wireless communications. The algorithm is faster than OMP because of the selection rule, and it sometimes provides good performance, although parameter tuning can be difficult. There are no rigorous results available for StOMP.

Very recently, Needell and Vershynin developed and analyzed another greedy approach, called *regularized OMP*, or ROMP [31, 30]. This algorithm is similar to OMP but uses a more sophisticated selection rule. Among the s largest entries of the signal proxy, it identifies the largest subset whose entries differ in magnitude by at most a factor of two. The work on ROMP represents an advance because the authors establish under restricted isometry hypotheses that their algorithm can approximately recover any compressible signal from noisy samples. More precisely, suppose that the sampling matrix Φ has restricted isometry constant $\delta_{8s} \leq 0.01/\sqrt{\log s}$. Given noisy samples $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$, ROMP produces a $2s$ -sparse signal approximation \mathbf{a} that satisfies

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq C\sqrt{\log s} \left[\frac{1}{\sqrt{s}} \|\mathbf{x} - \mathbf{x}_s\|_1 + \|\mathbf{e}\|_2 \right].$$

This result is comparable with the result for convex relaxation, aside from the extra logarithmic factor in the restricted isometry hypothesis and the error bound. The results for CoSaMP show that it does not suffer these parasitic factors, so its performance is essentially optimal.

After we initially presented this work, Dai and Milenkovic developed an algorithm called Subspace Pursuit that is very similar to CoSaMP. They established that their algorithm offers performance guarantees analogous with those for CoSaMP. See [10] for details.

Finally, we note that there is a class of sublinear algorithms for signal reconstruction from compressive samples. A sublinear algorithm uses time and space resources that are asymptotically smaller than the length of the signal. One of the earliest such techniques is the Fourier sampling algorithm of Gilbert et al. [20, 22]. This algorithm uses random (but structured) time samples to recover signals that are compressible with respect to the discrete Fourier basis. Given s polylog(N) samples³, Fourier sampling produces a signal approximation \mathbf{a} that satisfies

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq C \|\mathbf{x} - \mathbf{x}_s\|_2$$

except with probability N^{-1} . The result for Fourier sampling holds for each signal (rather than for all). Later, Gilbert et al. developed two other sublinear algorithms, chaining pursuit [17] and HHS pursuit [18], that offer uniform guarantees for all signals. Chaining pursuit has an error bound

$$\|\mathbf{x} - \mathbf{a}\|_1 \leq C \log N \|\mathbf{x} - \mathbf{x}_s\|_1$$

which is somewhat worse than (7.2). HHS pursuit achieves the error bound (7.2). These methods all require more measurements than the linear and superlinear algorithms (by logarithmic factors),

²A subgaussian random variable Z satisfies $\mathbb{P}\{|Z| > t\} \leq ce^{-ct^2}$ for all $t > 0$.

³The term polylog indicates a function that is dominated by a polynomial in the logarithm of its argument.

	CoSaMP	OMP	ROMP	Convex opt.	Fourier Samp.	HHS pursuit
General samples	RIP	Subgauss.	RIP	RIP	no	no
Opt. # samples	yes	yes	no	yes	no	no
Uniformity	yes	no	yes	yes	no	yes
Stability	yes	?	yes	yes	yes	yes
Running time	$O(mN)$	$O(smN)$	$O(smN)$	$LP(N, m)$	$s \text{ polylog}(N)$	$\text{poly}(s \log N)$

TABLE 2. Comparison of several signal recovery algorithms. The notation s refers to the sparsity level; m refers the number of measurements; N refers to the signal length. See the text for comments on specific designations.

and these measurements must be highly structured. As a result, the sublinear algorithms may not be useful in practice.

The sublinear algorithms are all combinatorial in nature. They use ideas from group testing to identify the support of the signal quickly. There are several other combinatorial signal recovery methods due to Cormode–Muthukrishnan [9] and Iwen [24]. These algorithms have drawbacks similar to the sublinear approaches.

7.2. Relative Performance. Table 2 summarizes the relative behavior of these algorithms in terms of the following criteria.

General samples: Does the algorithm work for a variety of sampling schemes? Or does it require structured samples? The designation “RIP” means that a bound on a restricted isometry constant suffices. “Subgauss.” means that the algorithm succeeds for the class of subgaussian sampling matrices.

Optimal number of samples: Can the algorithm recover s sparse signals from $O(s \log N)$ measurements? Or are its sampling requirements higher (by logarithmic factors)?

Uniformity: Does the algorithm recover all signals given a fixed sampling matrix? Or do the results require a sampling matrix to be drawn at random for each signal?

Stability: Does the algorithm succeed when (a) the signal is compressible but not sparse and (b) when the samples are contaminated with noise? In most cases, stable algorithms have error bounds similar to (7.2). See the discussion above for details.

Running time: What is the worst-case cost of the algorithm to recover a real-valued s -sparse signal to a fixed relative precision, given a sampling matrix with no special structure? The designation $LP(N, m)$ indicates the cost of solving a linear program with N variables and m constraints, which is $O(m^2 N^{1.5})$ for an interior-point method. Note that most of the algorithms can also take advantage of fast matrix–vector multiplies to obtain better running times.

Of the linear and superlinear algorithms, CoSaMP achieves the best performance on all these metrics. Although CoSaMP is slower than the sublinear algorithms, it makes up for this shortcoming by allowing more general sampling matrices and requiring fewer samples. See [23] for a much more detailed comparison of compressive sampling recovery algorithms.

7.3. Key Ideas. We conclude with a historical overview of the ideas that inform the CoSaMP algorithm and its analysis.

The overall greedy iterative structure of CoSaMP has a long history. The idea of approaching sparse approximation problems in this manner dates to the earliest algorithms. In particular, methods for variable selection in regression, such as forward selection and its relatives, all take this form [28]. Temlyakov’s survey [36] describes the historical role of greedy algorithms in nonlinear

approximation. Mallat and Zhang introduced greedy algorithms into the signal processing literature and proposed the name *matching pursuit* [27]. Gilbert, Strauss, and their collaborators showed how to incorporate greedy iterative strategies into fast algorithms for sparse approximation problems, and they established the first rigorous guarantees for greedy methods [19, 20]. Tropp provided a new theoretical analysis of OMP in his work [37]. Subsequently, Tropp and Gilbert proved that OMP was effective for compressive sampling [39].

Unlike the simplest greedy algorithms, CoSaMP identifies many components during each iteration, which allows the algorithm to run faster for many types of signals. It is not entirely clear where this idea first appeared. Several early algorithms of Gilbert et al. incorporate this approach [21, 40], and it is an essential feature of the Fourier sampling algorithm [20, 22]. More recent compressive sampling recovery algorithms also select multiple indices, including chaining pursuit [17], HHS pursuit [18], StOMP [14], and ROMP [31].

CoSaMP uses the restricted isometry properties of the sampling matrix to ensure that the identification step is successful. Candès and Tao isolated the restricted isometry conditions in their work on convex relaxation methods for compressive sampling [5]. The observation that restricted isometries can also be used to ensure the success of greedy methods is relatively new. This idea plays a role in HHS pursuit [18], but it is expressed more completely in the analysis of ROMP [31].

The pruning step of CoSaMP is essential to maintain the sparsity of the approximation, which is what permits us to use restricted isometries in the analysis of the algorithm. It also has significant ramifications for the running time because it impacts the speed of the iterative least-squares algorithms. This technique originally appeared in HHS pursuit [18].

The iteration invariant, Theorem 2.1, states that if the error is large then CoSaMP makes substantial progress. This approach to the overall analysis echoes the analysis of other greedy iterative algorithms, including the Fourier sampling method [20, 22] and HHS pursuit [18].

Finally, mixed-norm error bounds, such as that in Theorem A, have become an important feature of the compressive sampling literature. This idea appears in the work of Candès–Romberg–Tao on convex relaxation [3]; it is used in the analysis of HHS pursuit [18]; it also plays a role in the theoretical treatment of Cohen–Dahmen–DeVore [7].

Acknowledgment. We would like to thank Martin Strauss for many inspiring discussions. He is ultimately responsible for many of the ideas in the algorithm and analysis. We would also like to thank Roman Vershynin for suggestions that drastically simplified the proofs.

APPENDIX A. ALGORITHMIC VARIATIONS

This appendix describes other possible halting criteria and their consequences. It also proposes some other variations on the algorithm.

A.1. Halting Rules. There are three natural approaches to halting the algorithm. The first, which we have discussed in the body of the paper, is to stop after a fixed number of iterations. Another possibility is to use the norm $\|\mathbf{v}\|_2$ of the current samples as evidence about the norm $\|\mathbf{r}\|_2$ of the residual. A third possibility is to use the magnitude $\|\mathbf{y}\|_\infty$ of the entries of the proxy to bound the magnitude $\|\mathbf{r}\|_\infty$ of the entries of the residual.

It suffices to discuss halting criteria for sparse signals because Lemma 6.1 shows that the general case can be viewed in terms of sampling a sparse signal. Let \mathbf{x} be an s -sparse signal, and let \mathbf{a} be an s -sparse approximation. The residual $\mathbf{r} = \mathbf{x} - \mathbf{a}$. We write $\mathbf{v} = \Phi\mathbf{r} + \mathbf{e}$ for the induced noisy samples of the residual and $\mathbf{y} = \Phi^*\mathbf{v}$ for the signal proxy.

The discussion proceeds in two steps. First, we argue that an *a priori* halting criterion will result in a guarantee about the quality of the final signal approximation.

Theorem A.1 (Halting I). *The halting criterion $\|\mathbf{v}\|_2 \leq \varepsilon$ ensures that*

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq 1.06 \cdot (\varepsilon + \|\mathbf{e}\|_2).$$

The halting criterion $\|\mathbf{y}\|_\infty \leq \eta/\sqrt{2s}$ ensures that

$$\|\mathbf{x} - \mathbf{a}\|_\infty \leq 1.12\eta + 1.17\|\mathbf{e}\|_2.$$

Proof. Since \mathbf{r} is $2s$ -sparse, Proposition 3.1 ensures that

$$\sqrt{1 - \delta_{2s}}\|\mathbf{r}\|_2 - \|\mathbf{e}\|_2 \leq \|\mathbf{v}\|_2.$$

If $\|\mathbf{v}\|_2 \leq \varepsilon$, it is immediate that

$$\|\mathbf{r}\|_2 \leq \frac{\varepsilon + \|\mathbf{e}\|_2}{\sqrt{1 - \delta_{2s}}}.$$

The definition $\mathbf{r} = \mathbf{x} - \mathbf{a}$ and the numerical bound $\delta_{2s} \leq \delta_{4s} \leq 0.1$ dispatch the first claim.

Let $R = \text{supp}(\mathbf{r})$, and note that $|R| \leq 2s$. Proposition 3.1 results in

$$(1 - \delta_{2s})\|\mathbf{r}\|_2 - \sqrt{1 + \delta_{2s}}\|\mathbf{e}\|_2 \leq \|\mathbf{y}|_R\|_2.$$

Since

$$\|\mathbf{y}|_R\|_2 \leq \sqrt{2s}\|\mathbf{y}|_R\|_\infty \leq \sqrt{2s}\|\mathbf{y}\|_\infty,$$

we find that the requirement $\|\mathbf{y}\|_\infty \leq \eta/\sqrt{2s}$ ensures that

$$\|\mathbf{r}\|_\infty \leq \frac{\eta + \sqrt{1 + \delta_{2s}}\|\mathbf{e}\|_2}{1 - \delta_{2s}}.$$

The numerical bound $\delta_{2s} \leq 0.1$ completes the proof. \square

Second, we check that each halting criterion is triggered when the residual has the desired property.

Theorem A.2 (Halting II). *The halting criterion $\|\mathbf{v}\|_2 \leq \varepsilon$ is triggered as soon as*

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq 0.95 \cdot (\varepsilon - \|\mathbf{e}\|_2).$$

The halting criterion $\|\mathbf{y}\|_\infty \leq \eta/\sqrt{2s}$ is triggered as soon as

$$\|\mathbf{x} - \mathbf{a}\|_\infty \leq \frac{0.45\eta}{s} - \frac{0.68\|\mathbf{e}\|_2}{\sqrt{s}}.$$

Proof. Proposition 3.1 shows that

$$\|\mathbf{v}\|_2 \leq \sqrt{1 + \delta_{2s}}\|\mathbf{r}\|_2 + \|\mathbf{e}\|_2.$$

Therefore, the condition

$$\|\mathbf{r}\|_2 \leq \frac{\varepsilon - \|\mathbf{e}\|_2}{\sqrt{1 + \delta_{2s}}}$$

ensures that $\|\mathbf{v}\|_2 \leq \varepsilon$. Note that $\delta_{2s} \leq 0.1$ to complete the first part of the argument.

Now let R be the singleton containing the index of a largest-magnitude coefficient of \mathbf{y} . Proposition 3.1 implies that

$$\|\mathbf{y}\|_\infty = \|\mathbf{y}|_R\|_2 \leq \sqrt{1 + \delta_1}\|\mathbf{v}\|_2.$$

By the first part of this theorem, the halting criterion $\|\mathbf{y}\|_\infty \leq \eta/\sqrt{2s}$ is triggered as soon as

$$\|\mathbf{x} - \mathbf{a}\|_2 \leq 0.95 \cdot \left(\frac{\eta}{\sqrt{2s}\sqrt{1 + \delta_1}} - \|\mathbf{e}\|_2 \right).$$

Since $\mathbf{x} - \mathbf{a}$ is $2s$ -sparse, we have the bound $\|\mathbf{x} - \mathbf{a}\|_2 \leq \sqrt{2s}\|\mathbf{x} - \mathbf{a}\|_\infty$. To wrap up, recall that $\delta_1 \leq \delta_{2s} \leq 0.1$. \square

A.2. Other Variations. Here is a version of the algorithm that is, perhaps, simpler than Algorithm 2.1. At each iteration, we approximate the current residual rather than the entire signal. This approach is similar to HHS pursuit [18]. The inner loop changes in the following manner.

Identification: As before, select $\Omega = \text{supp}(\mathbf{y}_{2s})$.

Estimation: Solve a least-squares problem with the *current samples* instead of the original samples to obtain an approximation of the *residual signal*. Formally, $\mathbf{b} = \Phi_{\Omega}^{\dagger} \mathbf{v}$. In this case, one initializes the iterative least-squares algorithm with the zero vector to take advantage of the fact that the residual is becoming small.

Approximation Merger: Add this approximation of the residual to the previous approximation of the signal to obtain a new approximation of the signal: $\mathbf{c} = \mathbf{a}^{k-1} + \mathbf{b}$.

Pruning: Construct the s -sparse signal approximation: $\mathbf{a}^k = \mathbf{c}_s$.

Sample Update: Update the samples as before: $\mathbf{v} = \mathbf{u} - \Phi \mathbf{a}^k$.

By adapting the argument in this paper, we have been able to show that this algorithm also satisfies Theorem A. We believe this version is quite promising for applications. See [29] for other variations.

REFERENCES

- [1] Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.
- [2] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete Fourier information. *IEEE Trans. Info. Theory*, 52(2):489–509, Feb. 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. J. Candès. The restricted isometry property and its implications for compressed sensing. *C. R. Math. Acad. Sci. Paris, Serie I*, 346:589–592, 2008.
- [5] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Info. Theory*, 51(12):4203–4215, Dec. 2005.
- [6] E. J. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Info. Theory*, 52(12):5406–5425, Dec. 2006.
- [7] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. Submitted for publication, 2006.
- [8] T. Cormen, C. Lesierson, L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2nd edition, 2001.
- [9] G. Cormode and S. Muthukrishnan. Combinatorial algorithms for compressed sensing. Technical report, DIMACS, 2005.
- [10] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing: Closing the gap between performance and complexity. Submitted for publication, Mar. 2008.
- [11] I. Daubechies, M. DeFrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57:1413–1457, 2004.
- [12] D. L. Donoho. Compressed sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, Apr. 2006.
- [13] D. L. Donoho and J. Tanner. Counting faces of randomly projected polytopes when the projection radically lowers dimension. Dept. of Statistics Technical Report 2006-11, Stanford Univ., July 2006.
- [14] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck. Sparse solution of underdetermined linear equations by stagewise Orthogonal Matching Pursuit (StOMP). Submitted for publication, 2007.
- [15] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE J. Selected Topics in Signal Processing: Special Issue on Convex Optimization Methods for Signal Processing*, 1(4):586–598, 2007.
- [16] A. Garnaev and E. Gluskin. On widths of the Euclidean ball. *Sov. Math. Dokl.*, 30:200–204, 1984.
- [17] A. Gilbert, M. Strauss, J. Tropp, and R. Vershynin. Algorithmic linear dimension reduction in the ℓ_1 norm for sparse vectors. Submitted for publication, August 2006.
- [18] A. Gilbert, M. Strauss, J. Tropp, and R. Vershynin. One sketch for all: Fast algorithms for compressed sensing. In *Proc. 39th ACM Symp. Theory of Computing*, San Diego, June 2007.
- [19] A. C. Gilbert, S. Guha, P. Indyk, Y. Kotidis, S. Muthukrishnan, and M. J. Strauss. Fast, small-space algorithms for approximate histogram maintenance. In *ACM Symposium on Theoretical Computer Science*, 2002.

- [20] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. J. Strauss. Near-optimal sparse Fourier representations via sampling. In *Proc. of the 2002 ACM Symposium on Theory of Computing STOC*, pages 152–161, 2002.
- [21] A. C. Gilbert, M. Muthukrishnan, and M. J. Strauss. Approximation of functions over redundant dictionaries using coherence. In *Proc. of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, Jan. 2003.
- [22] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss. Improved time bounds for near-optimal sparse Fourier representation via sampling. In *Proceedings of SPIE Wavelets XI*, San Diego, CA, 2005.
- [23] P. Indyk and R. Berinde. Sparse recovery using sparse matrices. CSAIL technical report, Massachusetts Institute of Technology, Jan. 2008.
- [24] M. Iwen. A deterministic sub-linear time sparse Fourier algorithm via non-adaptive compressed sensing methods. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.
- [25] B. Kashin. The widths of certain finite dimensional sets and classes of smooth functions. *Izvestia*, 41:334–351, 1977.
- [26] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. A method for large-scale ℓ_1 -regularized least-squares problems with applications in signal processing and statistics. Submitted for publication, 2007.
- [27] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process.*, 41(12):3397–3415, 1993.
- [28] A. J. Miller. *Subset Selection in Regression*. Chapman and Hall, London, 2nd edition, 2002.
- [29] 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.
- [30] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. Submitted for publication, October 2007.
- [31] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. to appear, *Found. Comput. Math.*, 2007.
- [32] Y. E. Nesterov and A. S. Nemirovski. *Interior Point Polynomial Algorithms in Convex Programming*. SIAM, Philadelphia, 1994.
- [33] H. Rauhut. On the impossibility of uniform sparse reconstruction using greedy methods. To appear, *Sampl. Theory Signal Image Process.*, 2008.
- [34] G. Reeves and M. Gastpar. Sampling bounds for sparse support recovery in the presence of noise. Submitted for publication, Jan. 2008.
- [35] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Proc. 40th Annual Conference on Information Sciences and Systems*, Princeton, Mar. 2006.
- [36] V. Temlyakov. Nonlinear methods of approximation. *Foundations of Comput. Math.*, 3(1):33–107, 2003.
- [37] J. A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Info. Theory*, 50(10):2231–2242, Oct. 2004.
- [38] J. A. Tropp. Beyond Nyquist: Efficient sampling of sparse, bandlimited signals. Presented at SampTA, June 2007.
- [39] 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.
- [40] J. A. Tropp, A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss. Improved sparse approximation over quasi-incoherent dictionaries. In *Proc. 2003 IEEE International Conference on Image Processing*, Barcelona, 2003.