

7-16-2013

Signal Space CoSaMP for Sparse Recovery with Redundant Dictionaries

Mark A. Davenport

Georgia Institute of Technology - Main Campus

Deanna Needell

Claremont McKenna College

Michael B. Wakin

Colorado School of Mines

Recommended Citation

Davenport, M. A., Needell, D., Wakin, M. B., "Signal Space CoSaMP for Sparse Recovery with Redundant Dictionaries", IEEE Transactions on Information Theory, vol. 59, iss. 10, 6820 - 6829, 2013. doi: 10.1109/TIT.2013.2273491

This Article - postprint is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

Signal Space CoSaMP for Sparse Recovery with Redundant Dictionaries

Mark A. Davenport^g, Deanna Needell^c, and Michael B. Wakin^{m*}

^gSchool of Electrical and Computer Engineering, Georgia Institute of Technology,
777 Atlantic Dr. NW, Atlanta, GA 30332, USA. Email: mdav@gatech.edu

^cDepartment of Mathematics and Computer Science, Claremont McKenna College,
850 Columbia Ave., Claremont, CA 91711, USA. Email: dneedell@cmc.edu

^mDepartment of Electrical Engineering and Computer Science, Colorado School of Mines,
1500 Illinois St., Golden, CO 80401, USA. Email: mwakin@mines.edu

August 2012 (Revised February 2013, June 2013)

Abstract

Compressive sensing (CS) has recently emerged as a powerful framework for acquiring sparse signals. The bulk of the CS literature has focused on the case where the acquired signal has a sparse or compressible representation in an orthonormal basis. In practice, however, there are many signals that cannot be sparsely represented or approximated using an orthonormal basis, but that *do* have sparse representations in a redundant dictionary. Standard results in CS can sometimes be extended to handle this case provided that the dictionary is sufficiently incoherent or well-conditioned, but these approaches fail to address the case of a truly redundant or overcomplete dictionary. In this paper we describe a variant of the iterative recovery algorithm CoSaMP for this more challenging setting. We utilize the \mathbf{D} -RIP, a condition on the sensing matrix analogous to the well-known restricted isometry property. In contrast to prior work, the method and analysis are “signal-focused”; that is, they are oriented around recovering the *signal* rather than its dictionary *coefficients*. Under the assumption that we have a near-optimal scheme for projecting vectors in signal space onto the model family of candidate sparse signals, we provide provable recovery guarantees. Developing a practical algorithm that can provably compute the required near-optimal projections remains a significant open problem, but we include simulation results using various heuristics that empirically exhibit superior performance to traditional recovery algorithms.

1 Introduction

1.1 Overview

Compressive sensing (CS) is a powerful new framework for signal acquisition, offering the promise that we can acquire a vector $\mathbf{x} \in \mathbb{C}^n$ via only $m \ll n$ linear measurements provided that \mathbf{x} is *sparse* or *compressible*.¹ Specifically, CS considers the problem where we obtain measurements of

^{*}This work was partially supported by NSF grants DMS-1004718 and CCF-0830320, NSF CAREER grant CCF-1149225, and AFOSR grant FA9550-09-1-0465.

¹When we say that a vector \mathbf{z} is k -sparse, we mean that $\|\mathbf{z}\|_0 \stackrel{\text{def}}{=} |\text{supp}(\mathbf{z})| \leq k \ll n$. A compressible vector is one that is well-approximated as being sparse. We discuss compressibility in greater detail in Section 2.2.

the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, where \mathbf{A} is an $m \times n$ sensing matrix and \mathbf{e} is a noise vector. If \mathbf{x} is sparse or compressible and \mathbf{A} satisfies certain conditions, then CS provides a mechanism to recover the signal \mathbf{x} from the measurement vector \mathbf{y} efficiently and robustly.

Typically, however, signals of practical interest are not themselves sparse, but rather have a sparse expansion in some dictionary \mathbf{D} . By this we mean that there exists a sparse coefficient vector $\boldsymbol{\alpha}$ such that the signal \mathbf{x} can be expressed as $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$. One could then ask the simple question: How can we account for this signal model in CS? In some cases, there is a natural way to extend the standard CS formulation—since we can write the measurements as $\mathbf{y} = \mathbf{A}\mathbf{D}\boldsymbol{\alpha} + \mathbf{e}$ we can use standard CS techniques to first obtain an estimate $\hat{\boldsymbol{\alpha}}$ of the sparse coefficient vector. We can then synthesize an estimate $\hat{\mathbf{x}} = \mathbf{D}\hat{\boldsymbol{\alpha}}$ of the original signal. Unfortunately, this is a rather restrictive way to proceed for two main reasons: (i) the application of standard CS results to this problem will require that the matrix given by the product $\mathbf{A}\mathbf{D}$ satisfy certain properties that will not be satisfied for many interesting choices of \mathbf{D} , as discussed further below, and (ii) we are not really interested in recovering $\boldsymbol{\alpha}$ *per se*, but rather in obtaining an accurate estimate of \mathbf{x} . If the dictionary \mathbf{D} is poorly conditioned, the signal space recovery error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2$ could be significantly smaller or larger than the coefficient space recovery error $\|\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}\|_2$. It may be possible to recover \mathbf{x} in situations where recovering $\boldsymbol{\alpha}$ is impossible, and even if we could apply standard CS results to ensure that our estimate of $\boldsymbol{\alpha}$ is accurate, this would not necessarily translate into a recovery guarantee for \mathbf{x} .

In this paper we will consider an alternative approach to this problem and develop an algorithm for which we can provide guarantees on the recovery of \mathbf{x} while making no direct assumptions concerning our choice of \mathbf{D} . Before we describe our approach, however, it will be illuminating to see precisely what goes wrong in an attempt to extend the standard CS formulation. Towards this end, let us return to the case where \mathbf{x} is itself sparse (when $\mathbf{D} = \mathbf{I}$). In this setting, there are many possible algorithms that have been proposed for recovering an estimate of \mathbf{x} from measurements of the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, including ℓ_1 -minimization approaches [5, 12] and greedy/iterative methods such as iterative hard thresholding (IHT) [4], orthogonal matching pursuit (OMP) [18, 22, 28] and compressive sampling matching pursuit (CoSaMP) [21]. For any of these algorithms, it can be shown (see [4, 5, 10, 21]) that \mathbf{x} can be accurately recovered from the measurements \mathbf{y} if the matrix \mathbf{A} satisfies a condition introduced in [8] known as the restricted isometry property (RIP) with a sufficiently small constant δ_k (the precise requirement on δ_k varies from method to method). We say that \mathbf{A} satisfies the RIP of order k with constant $\delta_k \in (0, 1)$ if

$$\sqrt{1 - \delta_k} \leq \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2} \leq \sqrt{1 + \delta_k} \quad (1)$$

holds for all \mathbf{x} satisfying $\|\mathbf{x}\|_0 \leq k$. Importantly, if \mathbf{A} is generated randomly with independent and identically distributed (i.i.d.) entries drawn from a suitable distribution, and with a number of rows roughly proportional to the sparsity level k , then with high probability \mathbf{A} will satisfy (1) [2, 19, 24].

We now return to the case where \mathbf{x} is sparse with respect to a dictionary \mathbf{D} . If \mathbf{D} is unitary (i.e., if the dictionary is an orthonormal basis), then the same arguments used to establish (1) can be adapted to show that for a fixed \mathbf{D} , if we choose a random \mathbf{A} , then with high probability $\mathbf{A}\mathbf{D}$ will satisfy the RIP, i.e.,

$$\sqrt{1 - \delta_k} \leq \frac{\|\mathbf{A}\mathbf{D}\boldsymbol{\alpha}\|_2}{\|\boldsymbol{\alpha}\|_2} \leq \sqrt{1 + \delta_k} \quad (2)$$

will hold for all $\boldsymbol{\alpha}$ satisfying $\|\boldsymbol{\alpha}\|_0 \leq k$. Thus, standard CS algorithms can be used to accurately recover $\boldsymbol{\alpha}$, and because \mathbf{D} is unitary, the signal space recovery error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2$ will exactly equal the coefficient space recovery error $\|\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}\|_2$.

Unfortunately, this approach won't do in cases where \mathbf{D} is not unitary and especially in cases where \mathbf{D} is highly redundant/overcomplete. For example, \mathbf{D} might represent the overcomplete Discrete Fourier Transform (DFT), the Undecimated Discrete Wavelet Transform, a redundant Gabor dictionary, or a union of orthonormal bases. The challenges that we must confront when dealing with \mathbf{D} of this form include:

- Redundancy in \mathbf{D} will mean that in general, the representation of a vector \mathbf{x} in the dictionary is not unique—there may exist many possible coefficient vectors $\boldsymbol{\alpha}$ that can be used to synthesize \mathbf{x} .
- Coherence (correlations) between the columns of \mathbf{D} can make it difficult to satisfy (2) with a sufficiently small constant δ_k to apply existing theoretical guarantees. For instance, while the DFT forms an orthonormal basis, a $2\times$ overcomplete DFT is already quite coherent, with adjacent columns satisfying $|\langle d_i, d_{i+1} \rangle| > 2/\pi > 0.63$. Since the coherence provides a bound on δ_k for all $k \geq 2$, this means that \mathbf{D} itself cannot satisfy the RIP with a constant $\delta_{2k} < 0.63$. For the random constructions of \mathbf{A} typically considered in the context of CS, with high probability \mathbf{A} will preserve the conditioning of \mathbf{D} (good or bad) on each subspace of interest. (For such \mathbf{A} , one essentially needs \mathbf{D} itself to satisfy the RIP in order to expect \mathbf{AD} to satisfy the RIP.) Thus, in the case of the $2\times$ overcomplete DFT, we would expect the RIP constant for \mathbf{AD} to be *at least* roughly 0.63—well outside the range for which any of the sparse recovery algorithms described above are known to succeed.
- As noted above, if the dictionary \mathbf{D} is poorly conditioned, the signal space recovery error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2$ could differ substantially from the coefficient space recovery error $\|\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}\|_2$, further complicating any attempt to understand how well we can recover \mathbf{x} by appealing to results concerning the recovery of $\boldsymbol{\alpha}$.

All of these problems essentially stem from the fact that extending standard CS algorithms in an attempt to recover $\boldsymbol{\alpha}$ is a *coefficient-focused* recovery strategy. By trying to go from the measurements \mathbf{y} all the way back to the coefficient vector $\boldsymbol{\alpha}$, one encounters all the problems above due to the lack of orthogonality of the dictionary.

In contrast, in this paper we propose a *signal-focused* recovery strategy for CS. Our algorithm employs the model of sparsity in an arbitrary dictionary \mathbf{D} but directly obtains an estimate of the signal \mathbf{x} , and we provide guarantees on the quality of this estimate in signal space. Our algorithm is a modification of CoSaMP [21], and in cases where \mathbf{D} is unitary, our “Signal-Space CoSaMP” algorithm reduces to standard CoSaMP. However, our analysis requires comparatively weaker assumptions. Our bounds require only that \mathbf{A} satisfy the \mathbf{D} -RIP [6] (which we explain below in Section 1.3)—this is a different and less-restrictive condition to satisfy than requiring \mathbf{AD} to satisfy the RIP. The algorithm does, however, require the existence of a near-optimal scheme for projecting a vector \mathbf{x} onto the set of signals admitting a sparse representation in \mathbf{D} . While the fact that we require only an *approximate* projection is a significant relaxation of the requirements of traditional algorithms like CoSaMP (which require *exact* projections), showing that a practical algorithm can provably compute the required near-optimal projection remains a significant open problem. Nevertheless, as we will see in Section 3, various practical algorithms do lead to empirically favorable performance, suggesting that this challenge might not be insurmountable.

1.2 Related Work

Our work most closely relates to Blumensath's Projected Landweber Algorithm (PLA) [3], an extension of Iterative Hard Thresholding (IHT) [4] that operates in signal space and accounts for a

union-of-subspaces signal model. In several ways, our work is a parallel of this one, except that we extend CoSaMP rather than IHT to operate in signal space. Both works assume that \mathbf{A} satisfies the \mathbf{D} -RIP, and implementing both algorithms requires the ability to compute projections of vectors in the signal space onto the model family. (These requirements are described more thoroughly in Section 1.3 below.) One critical difference, however, is that our analysis allows for near-optimal projections whereas the PLA analysis does not.² Other fine differences are noted below.

Also related are works that employ an assumption of “analysis sparsity,” in which a signal \mathbf{x} is analyzed in a dictionary \mathbf{D} , and recovery from CS measurements is possible if $\mathbf{D}^*\mathbf{x}$ is sparse or compressible. Conventional CS algorithms such as ℓ_1 -minimization [6, 20], IHT [9, 14], and CoSaMP [14] have been adapted to account for analysis sparsity. These works are similar to ours in that they provide recovery guarantees in signal space and do not require \mathbf{AD} to satisfy the RIP. However, the assumption of analysis sparsity is in general different from the “synthesis sparsity” that we assume, where there exists a sparse coefficient vector $\boldsymbol{\alpha}$ such that $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$. For example, in the analysis case, exact sparsity implies that the analysis vector $\mathbf{D}^*\mathbf{x}$ is sparse, whereas in our setting exact sparsity implies the coefficient vector $\boldsymbol{\alpha}$ is. Under both of these assumptions, both the ℓ_1 -analysis method and our Signal Space CoSaMP algorithm provide recovery guarantees proportional to the norm of the noise in the measurements, $\|\mathbf{e}\|_2$ [6]. Without exact sparsity, ℓ_1 -analysis adds an additional factor $\|\mathbf{D}^*\mathbf{x} - (\mathbf{D}^*\mathbf{x})_k\|_1/\sqrt{k}$, where $(\mathbf{D}^*\mathbf{x})_k$ represents the k largest coefficients in magnitude of $(\mathbf{D}^*\mathbf{x})$. In the synthesis sparsity setting, the analogous “tail-term” is less straightforward (see Section 2.2 below for details). In summary, these algorithms are intended for different signal families and potentially different dictionaries. Nevertheless, there are some similarities between our work and analysis CoSaMP (ACoSaMP) [14] as we will see below.

Finally, it is worth mentioning the loose connection between our work and that in “model-based CS” [1]. In the case where \mathbf{D} is unitary, IHT and CoSaMP have been modified to account for structured sparsity models, in which certain sparsity patterns in the coefficient vector $\boldsymbol{\alpha}$ are forbidden. This work is similar to ours in that it involves a projection onto a model set. However, the algorithm (including the projection) operates in coefficient space (not signal space) and employs a different signal model; the requisite model-based RIP is more similar to requiring that \mathbf{AD} satisfy the RIP; and extensions to non-orthogonal dictionaries are not discussed. In fact, our work is in part inspired by our recent efforts [11] to extend the “model-based CS” framework to a non-orthogonal dictionary in which we proposed a similar algorithm to the one considered in this paper.

1.3 Requirements

First, to establish notation, suppose that \mathbf{A} is an $m \times n$ matrix and \mathbf{D} is an arbitrary $n \times d$ matrix. We suppose that we observe measurements of the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} = \mathbf{AD}\boldsymbol{\alpha} + \mathbf{e}$. For an index set $\Lambda \subset \{1, 2, \dots, d\}$ (sometimes referred to as a *support set*), we let \mathbf{D}_Λ denote the $n \times |\Lambda|$ submatrix of \mathbf{D} corresponding to the columns indexed by Λ , and we let $\mathcal{R}(\mathbf{D}_\Lambda)$ denote the column span of \mathbf{D}_Λ . We also use $\mathcal{P}_\Lambda : \mathbb{C}^n \rightarrow \mathbb{C}^n$ to denote the orthogonal projection operator onto $\mathcal{R}(\mathbf{D}_\Lambda)$ and $\mathcal{P}_{\Lambda^\perp} : \mathbb{C}^n \rightarrow \mathbb{C}^n$ to denote the orthogonal projection operator onto the orthogonal complement of $\mathcal{R}(\mathbf{D}_\Lambda)$.³

We will approach our analysis under the assumption that the matrix \mathbf{A} satisfies the \mathbf{D} -RIP [6]. Specifically, we say that \mathbf{A} satisfies the \mathbf{D} -RIP of order k if there exists a constant $\delta_k \in (0, 1)$ such

²Technically, the analysis of the PLA [3] allows for near-optimal projections but only with an additive error term. For sparse models, however, such projections could be made arbitrarily accurate simply by rescaling the signal before projecting. Thus, we consider the PLA to require exact projections for sparse models.

³Note that $\mathcal{P}_{\Lambda^\perp}$ does *not* represent the orthogonal projection operator onto $\mathcal{R}(\mathbf{D}_{\{1,2,\dots,d\}\setminus\Lambda})$.

that

$$\sqrt{1 - \delta_k} \leq \frac{\|\mathbf{A}\mathbf{D}\boldsymbol{\alpha}\|_2}{\|\mathbf{D}\boldsymbol{\alpha}\|_2} \leq \sqrt{1 + \delta_k} \quad (3)$$

holds for all $\boldsymbol{\alpha}$ satisfying $\|\boldsymbol{\alpha}\|_0 \leq k$. We note that this is different from requiring that \mathbf{A} satisfy the RIP—although (1) and (3) appear similar, the RIP requirement demands that this condition holds for vectors \boldsymbol{x} containing few nonzeros, while the \mathbf{D} -RIP requirement demands that this condition holds for vectors \boldsymbol{x} having a sparse representation in the dictionary \mathbf{D} . We also note that, compared to the requirement that $\mathbf{A}\mathbf{D}$ satisfy the RIP (2), it is relatively easy to ensure that \mathbf{A} satisfies the \mathbf{D} -RIP. In particular, we have the following lemma.

Lemma 1.1 (Corollary 3.1 of [11]). *For any choice of \mathbf{D} , if \mathbf{A} is populated with i.i.d. random entries from a Gaussian or subgaussian distribution, then with high probability, \mathbf{A} will satisfy the \mathbf{D} -RIP of order k as long as $m = O(k \log(d/k))$.*

In fact, using the results of [17] one can extend this result to show that given any matrix \mathbf{A} satisfying the traditional RIP, by applying a random sign matrix one obtains a matrix that with high probability will satisfy the \mathbf{D} -RIP.

Next, recall that one of the key steps in the traditional CoSaMP algorithm is to project a vector in signal space onto the model family of candidate sparse signals. In the traditional setting (when \mathbf{D} is an orthonormal basis), this step is trivial and can be performed by simple thresholding of the entries of the coefficient vector. Our Signal Space CoSaMP algorithm (described more completely in Section 2) involves replacing thresholding with a more general projection of vectors in the signal space onto the signal model. Specifically, for a given vector $\mathbf{z} \in \mathbb{C}^n$ and a given sparsity level k , define

$$\Lambda_{\text{opt}}(\mathbf{z}, k) := \arg \min_{\Lambda: |\Lambda|=k} \|\mathbf{z} - \mathcal{P}_{\Lambda}\mathbf{z}\|_2.$$

The support $\Lambda_{\text{opt}}(\mathbf{z}, k)$ —if we could compute it—could be used to generate the best k -sparse approximation to \mathbf{z} ; in particular, the nearest neighbor to \mathbf{z} among all signals that can be synthesized using k columns from \mathbf{D} is given by $\mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z}$. Unfortunately, computing $\Lambda_{\text{opt}}(\mathbf{z}, k)$ may be difficult in general. Therefore, we allow for near-optimal projections to be used in our algorithm. For a given vector $\mathbf{z} \in \mathbb{C}^n$ and a given sparsity level k , we assume a method is available for producing an estimate of $\Lambda_{\text{opt}}(\mathbf{z}, k)$, denoted $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ and having cardinality $|\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)| = k$, that satisfies

$$\|\mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z} - \mathcal{P}_{\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)}\mathbf{z}\|_2 \leq \min \left(\epsilon_1 \|\mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z}\|_2, \epsilon_2 \|\mathbf{z} - \mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z}\|_2 \right) \quad (4)$$

for some constants $\epsilon_1, \epsilon_2 \geq 0$. Setting ϵ_1 or ϵ_2 equal to 0 would lead to the requirement that $\mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z} = \mathcal{P}_{\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)}\mathbf{z}$ exactly. Note that our metric for judging the quality of an approximation to $\Lambda_{\text{opt}}(\mathbf{z}, k)$ is entirely in terms of its impact in signal space. It might well be the case that $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ could satisfy (4) while being substantially different (or even disjoint) from $\Lambda_{\text{opt}}(\mathbf{z}, k)$. Thus, while computing $\Lambda_{\text{opt}}(\mathbf{z}, k)$ may be extremely challenging when \mathbf{D} is highly redundant, there is hope that efficiently computing an approximation that satisfies (4) might still be possible. However, determining whether this is the case remains an open problem.

It is important to note that although computing a near-optimal support estimate that satisfies the condition (4) remains a challenging task in general, several important related works have run into the same problem. As we previously mentioned, the existing analysis of the PLA [3] actually requires exact computation of $\Lambda_{\text{opt}}(\mathbf{z}, k)$. The analysis of ACoSaMP [14] allows a near-optimal projection to be used, with a near-optimality criterion that differs slightly from ours. Simulations of ACoSaMP, however, have relied on practical (but not theoretically backed) methods for computing

Algorithm 1 Signal Space CoSaMP

input: \mathbf{A} , \mathbf{D} , \mathbf{y} , k , stopping criterion
initialize: $\mathbf{r} = \mathbf{y}$, $\mathbf{x}^0 = 0$, $\ell = 0$, $\Gamma = \emptyset$
while not converged **do**
 proxy: $\tilde{\mathbf{v}} = \mathbf{A}^* \mathbf{r}$
 identify: $\Omega = \mathcal{S}_{\mathbf{D}}(\tilde{\mathbf{v}}, 2k)$
 merge: $T = \Omega \cup \Gamma$
 update: $\tilde{\mathbf{x}} = \arg \min_{\mathbf{z}} \|\mathbf{y} - \mathbf{A}\mathbf{z}\|_2$ s.t. $\mathbf{z} \in \mathcal{R}(\mathbf{D}_T)$
 $\Gamma = \mathcal{S}_{\mathbf{D}}(\tilde{\mathbf{x}}, k)$
 $\mathbf{x}^{\ell+1} = \mathcal{P}_{\Gamma} \tilde{\mathbf{x}}$
 $\mathbf{r} = \mathbf{y} - \mathbf{A}\mathbf{x}^{\ell+1}$
 $\ell = \ell + 1$
end while
output: $\hat{\mathbf{x}} = \mathbf{x}^{\ell}$

this projection. In Section 3, we present simulation results for Signal Space CoSaMP using practical (but not theoretically backed) methods for computing $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$. We believe that computing provably near-optimal projections is an important topic worthy of further study, as it is really the crux of the problem in all of these settings.

2 Algorithm and Recovery Guarantees

Given noisy compressive measurements of the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, our Signal Space CoSaMP algorithm for recovering an estimate of the signal \mathbf{x} is specified in Algorithm 1.

2.1 A Bound for the Recovery of Sparse Signals

For signals having a sparse representation in the dictionary \mathbf{D} , we have the following guarantee.

Theorem 2.1. *Suppose there exists a k -sparse coefficient vector $\boldsymbol{\alpha}$ such that $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$, and suppose that \mathbf{A} satisfies the \mathbf{D} -RIP of order $4k$. Then the signal estimate $\mathbf{x}^{\ell+1}$ obtained after $\ell+1$ iterations of Signal Space CoSaMP satisfies*

$$\|\mathbf{x} - \mathbf{x}^{\ell+1}\|_2 \leq C_1 \|\mathbf{x} - \mathbf{x}^{\ell}\|_2 + C_2 \|\mathbf{e}\|_2, \quad (5)$$

where C_1 and C_2 are constants that depend on the isometry constant δ_{4k} and on the approximation parameters ϵ_1 and ϵ_2 . In particular,

$$C_1 = ((2 + \epsilon_1)\delta_{4k} + \epsilon_1)(2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \quad \text{and} \quad C_2 = \left(\frac{(2 + \epsilon_2)((2 + \epsilon_1)(1 + \delta_{4k}) + 2)}{\sqrt{1 - \delta_{4k}}} \right).$$

Our proof of Theorem 2.1 appears in Appendix A and is a modification of the original CoSaMP proof [21]. Through various combinations of ϵ_1 , ϵ_2 , and δ_{4k} , it is possible to ensure that $C_1 < 1$ and thus that the accuracy of Signal Space CoSaMP improves at each iteration. Taking $\epsilon_1 = \frac{1}{10}$, $\epsilon_2 = 1$, and $\delta_{4k} = 0.029$ as an example, we obtain $C_1 \leq 0.5$ and $C_2 \leq 12.7$. Applying the relation (5) recursively, we then conclude the following.

Corollary 2.1. *Suppose there exists a k -sparse coefficient vector $\boldsymbol{\alpha}$ such that $\boldsymbol{x} = \boldsymbol{D}\boldsymbol{\alpha}$, and suppose that \boldsymbol{A} satisfies the \boldsymbol{D} -RIP of order $4k$ with $\delta_{4k} = 0.029$. Suppose that Signal Space CoSaMP is implemented using near-optimal projections with approximation parameters $\epsilon_1 = \frac{1}{10}$ and $\epsilon_2 = 1$. Then the signal estimate \boldsymbol{x}^ℓ obtained after ℓ iterations of Signal Space CoSaMP satisfies*

$$\left\| \boldsymbol{x} - \boldsymbol{x}^\ell \right\|_2 \leq 2^{-\ell} \|\boldsymbol{x}\|_2 + 25.4 \|\boldsymbol{e}\|_2. \quad (6)$$

By taking a sufficient number of iterations ℓ , the first term on the right hand side of (6) can be made arbitrarily small, and ultimately the recovery error depends only on the level of noise in the measurements. For a precision parameter η , this shows that at most $O(\log(\|\boldsymbol{x}\|_2/\eta))$ iterations are needed to ensure that

$$\|\boldsymbol{x} - \widehat{\boldsymbol{x}}\|_2 = O(\eta + \|\boldsymbol{e}\|_2) = O(\max\{\eta, \|\boldsymbol{e}\|_2\}).$$

The cost of a single iteration of the method is dominated by the cost of the **identify** and **update** steps, where we must obtain sparse approximations to $\tilde{\boldsymbol{v}}$ and $\tilde{\boldsymbol{x}}$, respectively. We emphasize again that there is no known algorithm for computing the approximation $\mathcal{S}_{\boldsymbol{D}}$ efficiently, and the ultimate computational complexity of the algorithm will depend on this choice. However, in the absence of a better choice, one natural option for estimating $\mathcal{S}_{\boldsymbol{D}}$ is to use a greedy method such as OMP or CoSaMP (see Section 3 for experimental results using these choices). The running time of these greedy methods on an $n \times d$ dictionary \boldsymbol{D} are $O(knd)$ or $O(nd)$, respectively [21]. Therefore, using these methods as approximations in the **identify** and **update** steps yields an overall running time of $O(knd \log(\|\boldsymbol{x}\|_2/\eta))$ or $O(nd \log(\|\boldsymbol{x}\|_2/\eta))$. For sparse signal recovery, these running times are in line with state-of-the-art bounds for traditional CS algorithms such as CoSaMP [21], except that Signal Space CoSaMP can be applied in settings where the dictionary \boldsymbol{D} is not unitary. The error bounds of Corollary 2.1 also match those of classical results, except that here we assume suitable accuracy of $\mathcal{S}_{\boldsymbol{D}}$. Of course, since no efficient near-optimal projection method is known, this presents a weakness in these results, but it is one shared by all comparable results in the existing literature.

2.2 A Discussion Concerning the Recovery of Arbitrary Signals

We can extend our analysis to account for signals \boldsymbol{x} that do not exactly have a sparse representation in the dictionary \boldsymbol{D} . For the sake of illustration, we again take $\epsilon_1 = \frac{1}{10}$, $\epsilon_2 = 1$, and $\delta_{4k} = 0.029$, and we show how (6) can be extended.

We again assume measurements of the form $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}$ but allow \boldsymbol{x} to be an arbitrary signal in \mathbb{C}^n . For any vector $\boldsymbol{\alpha}_k \in \mathbb{C}^d$ such that $\|\boldsymbol{\alpha}_k\|_0 \leq k$, we can write

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{A}\boldsymbol{D}\boldsymbol{\alpha}_k + \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k) + \boldsymbol{e}.$$

The term $\tilde{\boldsymbol{e}} := \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k) + \boldsymbol{e}$ can be viewed as noise in the measurements of the k -sparse signal $\boldsymbol{D}\boldsymbol{\alpha}_k$. Then by (6) we have

$$\left\| \boldsymbol{D}\boldsymbol{\alpha}_k - \boldsymbol{x}^\ell \right\|_2 \leq 2^{-\ell} \|\boldsymbol{D}\boldsymbol{\alpha}_k\|_2 + 25.4 \|\tilde{\boldsymbol{e}}\|_2 \leq 2^{-\ell} \|\boldsymbol{D}\boldsymbol{\alpha}_k\|_2 + 25.4 (\|\boldsymbol{e}\|_2 + \|\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k)\|_2), \quad (7)$$

and so using the triangle inequality,

$$\left\| \boldsymbol{x} - \boldsymbol{x}^\ell \right\|_2 \leq 2^{-\ell} \|\boldsymbol{D}\boldsymbol{\alpha}_k\|_2 + 25.4 \|\boldsymbol{e}\|_2 + \|\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k\|_2 + 25.4 \|\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k)\|_2. \quad (8)$$

One can then choose the coefficient vector $\boldsymbol{\alpha}_k$ to minimize the right hand side of (8).

Bounds very similar to this appear in the analysis of both the PLA [3] and ACoSaMP [14]. Such results are somewhat unsatisfying since it is unclear how the term $\|\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}_k)\|_2$ might

behave. Unfortunately, these bounds are difficult to improve upon when \mathbf{D} is not unitary. Under some additional assumptions, however, we can make further modifications to (8). For example, the following proposition allows us to bound $\|\mathbf{A}(\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k)\|_2$.

Proposition 2.1 (Proposition 3.5 of [21]). *Suppose that \mathbf{A} satisfies the upper inequality of the RIP, i.e., that $\|\mathbf{A}\mathbf{x}\|_2 \leq \sqrt{1 + \delta_k} \|\mathbf{x}\|_2$ holds for all $\mathbf{x} \in \mathbb{C}^n$ with $\|\mathbf{x}\|_0 \leq k$. Then, for every signal $\mathbf{z} \in \mathbb{C}^n$,*

$$\|\mathbf{A}\mathbf{z}\|_2 \leq \sqrt{1 + \delta_k} \left[\|\mathbf{z}\|_2 + \frac{1}{\sqrt{k}} \|\mathbf{z}\|_1 \right]. \quad (9)$$

Plugging this result in to (8), we have

$$\begin{aligned} \left\| \mathbf{x} - \mathbf{x}^\ell \right\|_2 &\leq 2^{-\ell} \|\mathbf{D}\boldsymbol{\alpha}_k\|_2 + 25.4 \|\mathbf{e}\|_2 + (25.4\sqrt{1 + \delta_k} + 1) \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k\|_2 \\ &\quad + \frac{25.4\sqrt{1 + \delta_k}}{\sqrt{k}} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k\|_1. \end{aligned} \quad (10)$$

For any $\mathbf{x} \in \mathbb{C}^n$, one could define the model mismatch quantity

$$\text{mismatch}(\mathbf{x}) := \inf_{\boldsymbol{\alpha}_k: \|\boldsymbol{\alpha}_k\|_0 \leq k} \left[\|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k\|_2 + \frac{1}{\sqrt{k}} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k\|_1 \right].$$

We remark that this mismatch quantity is analogous to the tail bounds in the literature for methods which do not allow for redundant dictionaries. In particular, the ℓ_1 -norm term in the classical setting is required on account of geometric results about Gelfand widths [13, 16]. If this quantity is large, then the signal is far from compressible and we are not in a setting for which our method is designed. Plugging this definition into (10), we obtain

$$\left\| \mathbf{x} - \mathbf{x}^\ell \right\|_2 \leq 2^{-\ell} \|\mathbf{D}\boldsymbol{\alpha}_k\|_2 + 25.4 \|\mathbf{e}\|_2 + 26.4\sqrt{1 + \delta_k} \cdot \text{mismatch}(\mathbf{x}). \quad (11)$$

In some sense, one can view $\text{mismatch}(\mathbf{x})$ as the “distance” from \mathbf{x} to the set of signals that are k -sparse in the dictionary \mathbf{D} , except that the actual “distance” being measured is a mixed ℓ_2/ℓ_1 norm. In cases where one expects this distance to be small, (11) guarantees that the recovery error will be small.

We close this discussion by noting that if we make the stronger assumption that \mathbf{AD} actually satisfies the RIP, we can also measure the model mismatch in the coefficient space rather than the signal space. Let $\boldsymbol{\alpha} \in \mathbb{C}^d$ be any vector that satisfies $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$. Then using a natural extension of Proposition 2.1, we conclude that

$$\|\mathbf{A}(\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k)\|_2 = \|\mathbf{AD}(\boldsymbol{\alpha} - \boldsymbol{\alpha}_k)\|_2 \leq \sqrt{1 + \delta_k} \left(\|\boldsymbol{\alpha} - \boldsymbol{\alpha}_k\|_2 + \frac{1}{\sqrt{k}} \|\boldsymbol{\alpha} - \boldsymbol{\alpha}_k\|_1 \right).$$

When $\boldsymbol{\alpha}$ is compressible, the recovery error (8) will be reasonably small.

3 Simulations

As we discussed in Section 1.3, the main difficulty in implementing our algorithm is in computing projections of vectors in the signal space onto the model family of candidate sparse signals. One such projection is required in the **identify** step in Algorithm 1; another such projection is required in the **update** step. Although our theoretical analysis can accommodate near-optimal support estimates

$\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ that satisfy the condition (4), computing even near-optimal supports can be a challenging task for many dictionaries of practical interest. In this section, we present simulation results using practical (but heuristic) methods for attempting to find near-optimal supports $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$. Specifically, we find ourselves in a situation that mirrors the early days of the sparse recovery literature—we would like to identify a sparse vector that well-approximates \mathbf{z} . This is precisely the scenario where recovery algorithms like OMP and ℓ_1 -minimization were first proposed, so despite the lack of a theoretical guarantee, we can still apply these algorithms. Of course, if we are leaving the solid ground of theory and entering the world of heuristics, we can also just consider applying standard algorithms like OMP, CoSaMP, and ℓ_1 -minimization algorithms directly to the CS recovery problem to see how they perform. We will see, however, that the “Signal Space CoSaMP” algorithms resulting from using standard solvers for $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ —even though they are not quite covered by our theory—can nevertheless outperform these classical CS reconstruction techniques.

In all simulations that follow, we set the signal length $n = 256$. We let \mathbf{D} be an $n \times d$ dictionary (two different dictionaries are considered below), and we construct a length- d coefficient vector $\boldsymbol{\alpha}$ with $k = 8$ nonzero entries chosen as i.i.d. Gaussian random variables. We set $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$, construct \mathbf{A} as a random $m \times n$ matrix with i.i.d. Gaussian entries, and collect noiseless measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$. After reconstructing an estimate of \mathbf{x} , we declare this recovery to be perfect if the SNR of the recovered signal estimate is above 100 dB. All of our simulations were performed via a MATLAB software package that we have made available for download at <http://users.ece.gatech.edu/~mdavenport/software>.

3.1 Renormalized Orthogonal Dictionary

As an instructive warm-up, we begin with one almost trivial example where the optimal projection can be computed exactly: we construct \mathbf{D} by taking an orthobasis and renormalizing its columns while maintaining their orthogonality. To compute $\Lambda_{\text{opt}}(\mathbf{z}, k)$ with such a dictionary, one merely computes $\mathbf{D}^*\mathbf{z}$, divides this vector elementwise by the column norms of \mathbf{D} , and sets Λ equal to the positions of the k largest entries.

For the sake of demonstration, we set \mathbf{D} equal to the $n \times n$ identity matrix, but we then rescale its first $n/2$ diagonal entries to equal 100 instead of 1. We construct sparse coefficient vectors $\boldsymbol{\alpha}$ as described above, with supports chosen uniformly at random. As a function of the number of measurements m , we plot in Figure 1 the percent of trials in which Signal Space CoSaMP recovers \mathbf{x} exactly. We see that with roughly 50 or more measurements, the recovery is perfect in virtually all trials. In contrast, this figure also shows the performance of traditional CoSaMP using the combined dictionary $\mathbf{A}\mathbf{D}$ to first recover $\boldsymbol{\alpha}$. Because of the non-normalized columns in \mathbf{D} , CoSaMP almost never recovers the correct signal; in fact its support estimates almost always are contained in the set $\{1, 2, \dots, n/2\}$. Of course, if presented with this problem in practice one would naturally want to modify traditional CoSaMP to account for the various column norms in \mathbf{D} ; the point here is merely that our algorithm gives a principled way to make this (trivial) modification.

3.2 Overcomplete DFT Dictionary

As a second example, we set $d = 4n$ and let \mathbf{D} be a $4 \times$ overcomplete DFT dictionary. In this dictionary, neighboring columns are highly coherent, while distant columns are not. We consider two scenarios: one in which the $k = 8$ nonzero entries of $\boldsymbol{\alpha}$ are randomly positioned but well-separated (with a minimum spacing of 8 zeros in between any pair of nonzero entries), and one in which the $k = 8$ nonzero entries all cluster together in a single, randomly-positioned block. Because




Figure 1: Performance in recovering signals having a $k = 8$ sparse representation in a dictionary \mathbf{D} with orthogonal, but not normalized, columns. The plot shows, for various numbers of measurements m , the percent of trials in which each algorithm recovers the signal exactly. Signal Space CoSaMP (in which we can compute optimal projections) outperforms an unmodified CoSaMP algorithm.

of the nature of the columns in \mathbf{D} , we see that many recovery algorithms perform differently in these two scenarios.

3.2.1 Well-separated coefficients

Figure 2(a) plots the performance of six different recovery algorithms for the scenario where the nonzero entries of α are well-separated. Two of these algorithms are the traditional OMP and CoSaMP algorithms from CS, each using the combined dictionary \mathbf{AD} to first recover α . We actually see that OMP performs substantially better than CoSaMP in this scenario, apparently because it can select one coefficient at a time and is less affected by the coherence of \mathbf{D} . It is somewhat remarkable that OMP succeeds at all, given that \mathbf{AD} will not satisfy the RIP and we are not aware of any existing theory that would guarantee the performance of OMP in this scenario.

We also show in Figure 2(a) two variants of Signal Space CoSaMP: one in which OMP is used for computing $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ (labeled “SSCoSaMP (OMP)”), and one in which CoSaMP is used for computing $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ (labeled “SSCoSaMP (CoSaMP)”). That is, these algorithms actually use OMP or CoSaMP as an inner loop inside of Signal Space CoSaMP to find a sparse solution to the equation $\mathbf{z} = \mathbf{D}\alpha$. In this scenario, we see that the performance of SSSCoSaMP (OMP) is substantially better than OMP, while the performance of SSSCoSaMP (CoSaMP) is poor. We believe that this happens for the same reason that traditional OMP outperforms traditional CoSaMP. In general, we have found that when OMP performs well, SSSCoSaMP (OMP) may perform even better, and when CoSaMP performs poorly, SSSCoSaMP (CoSaMP) may still perform poorly.

Figure 2(a) also shows the performance of two algorithms that involve convex optimization for sparse regularization. One, labeled “ ℓ_1 ,” uses an ℓ_1 -minimization approach [27] to find a sparse coefficient vector α' subject to the constraint that $\mathbf{y} = \mathbf{AD}\alpha'$. This algorithm actually outperforms traditional OMP in this scenario. The other, labeled “SSCoSaMP (ℓ_1),” is a variant of Signal Space CoSaMP in which ℓ_1 -minimization is used for computing $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$.⁴ Specifically, to compute $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$, we search for the vector α' having the smallest ℓ_1 norm subject to the constraint that $\mathbf{z} = \mathbf{D}\alpha'$, and we then choose the support that contains the k largest entries of this vector. Remarkably, this algorithm performs best of all. We believe that this is likely due to the fact that,

⁴We are not unaware of the irony of using ℓ_1 -minimization inside of a greedy algorithm.

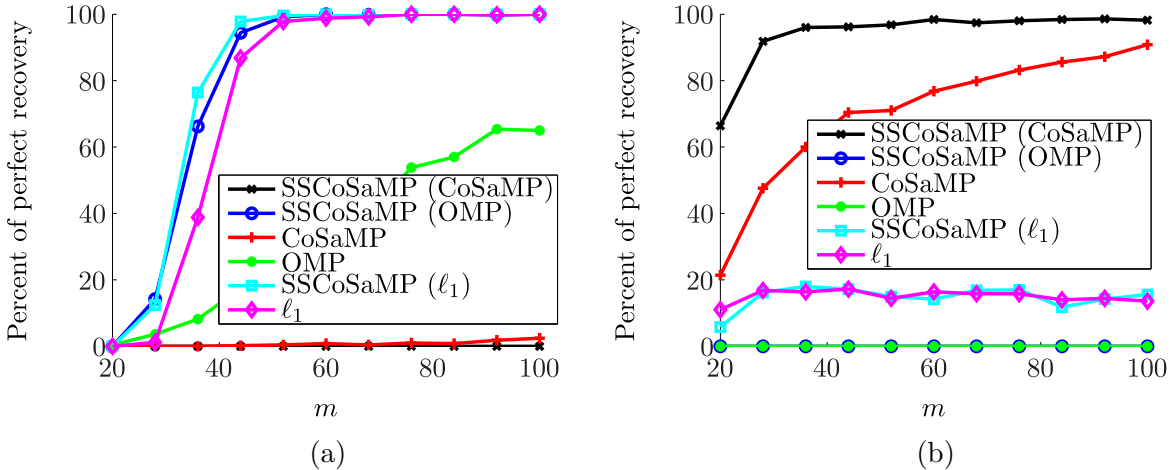


Figure 2: Performance in recovering signals having a $k = 8$ sparse representation in a $4 \times$ overcomplete DFT dictionary. Two scenarios are shown: (a) one in which the $k = 8$ nonzero entries of α are randomly positioned but well-separated, and (b) one in which the $k = 8$ nonzero entries all cluster together in a single, randomly-positioned block. Algorithms involving OMP and ℓ_1 -minimization perform well in the former scenario; algorithms involving CoSaMP perform well in the latter. In general, the Signal Space CoSaMP variants outperform the corresponding traditional CS algorithm.

for the overcomplete DFT dictionary, ℓ_1 -minimization is known to be capable of finding $\Lambda_{\text{opt}}(\mathbf{z}, k)$ exactly when $\mathbf{z} = \mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)} \mathbf{z}$ and the entries of $\Lambda_{\text{opt}}(\mathbf{z}, k)$ are sufficiently well-separated [7]. While we do not guarantee that this condition will be met within every iteration of Signal Space CoSaMP, the fact that the original coefficient vector α has well-separated coefficients seems to be intimately related to the success of ℓ_1 and SSCoSaMP (ℓ_1) here.

We note that all of the above algorithms involve a step where a least-squares problem must be solved on an estimated support set.⁵ This might seem somewhat contrary to our signal-focused approach, since solving this least-squares problem essentially involves recovering a set of coefficients α and then computing $\tilde{\mathbf{x}} = \mathbf{D}_T \alpha$. However, recall that at this point in the algorithm we are solving an over-determined system, so there is no significant difference between solving for α versus $\tilde{\mathbf{x}}$. The problem with coefficient-focused strategies and analysis is that they rely on identifying the subset T containing the “correct” subset of coefficients, whereas for us, T could be very different from the “correct” subset, and the vector α has no particular significance—all that matters is the vector $\tilde{\mathbf{x}}$ that this step ultimately synthesizes. Nevertheless, it is worth noting that when a dictionary \mathbf{D} is highly coherent, it can be numerically challenging to solve this problem, as the resulting submatrix can be very poorly conditioned. Following our discussion in [11], we employ *Tikhonov regularization* [15, 23, 25, 26] and solve a norm-constrained least-squares problem to improve its conditioning. For this we must provide our algorithms with an upper bound on the norm of the sparse coefficient vector. In the simulations above, we have selected this bound to be $10 \times$ the true norm of the original α . The selection of this bound does not have a substantial impact on the performance of OMP or SSCoSaMP (OMP), but we have noticed that CoSaMP and SSCoSaMP (CoSaMP) perform somewhat better when the true norm $\|\alpha\|_2$ is provided as an oracle.

⁵We use this for debiasing after running ℓ_1 .

3.2.2 Clustered coefficients

Figure 2(b) plots the performance of the same six recovery algorithms for the scenario where the nonzero entries of α are clustered into a single block. Although one could of course employ a block-sparse recovery algorithm in this scenario, our intent is more to study the impact that neighboring active atoms have on the algorithms above.

In this scenario, between the traditional greedy algorithms, CoSaMP now outperforms OMP, apparently because it is designed to select multiple indices at each step and will not be as affected by the coherence of neighboring active columns in \mathbf{D} . We also see that the performance of SSCoSaMP (CoSaMP) is somewhat better than CoSaMP, while the performance of SSCoSaMP (OMP) is poor. We believe that this happens for the same reason that traditional CoSaMP outperforms traditional OMP. In general, we have found that when CoSaMP performs well, SSCoSaMP (CoSaMP) may perform even better, and when OMP performs poorly, SSCoSaMP (OMP) may still perform poorly.

In terms of our condition for perfect recovery (estimating \mathbf{x} to within an SNR of 100 dB or more), neither of the algorithms that involve ℓ_1 -minimization perform well in this scenario. However, we do note that both ℓ_1 and SSCoSaMP (ℓ_1) do frequently recover an estimate of \mathbf{x} with an SNR of 50 dB or more, though still not quite as frequently as SSCoSaMP (CoSaMP) does.

In these simulations, we again use Tikhonov regularization with a norm bound that is $10\times$ the true norm of the original α . However, we have not found that changing this norm bound has a significant impact on CoSaMP or SSCoSaMP (CoSaMP) in this scenario. We also note that in this scenario we found it beneficial to run CoSaMP and the three Signal Space CoSaMP methods for a few more iterations than in the case of well-separated coefficients; convergence to exactly the correct support can be slow in this case where multiple neighboring atoms in a coherent dictionary are active.

3.2.3 Additional investigations

We close with some final remarks concerning additional investigations. First, our simulations above have tested three heuristic methods for attempting to find near-optimal supports $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$, and we have evaluated the performance of these methods based on the ultimate success or failure of SSCoSaMP in recovering the signal. In some problems of much smaller dimension (where we could use an exhaustive search to find $\Lambda_{\text{opt}}(\mathbf{z}, k)$) we monitored the performance of CoSaMP, OMP, and ℓ_1 -minimization for computing $\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)$ in terms of the effective ϵ_1 and ϵ_2 values they attained according to the metric in (4). For scenarios where the nonzero entries of α were well-separated, we observed typical ϵ_1 and ϵ_2 values for OMP and ℓ_1 -minimization on the order of 1 or less. For CoSaMP, these values were larger by one or two orders of magnitude, as might be expected based on the signal recovery results presented in Section 3.2.1.⁶ For scenarios where the nonzero entries of α were clustered, the ϵ_2 values for OMP and ℓ_1 -minimization increased by about one order of magnitude, but the ϵ_1 and ϵ_2 values for CoSaMP did not change significantly. The primary reason for this, despite the superior signal recovery performance of SSCoSaMP (CoSaMP) in Section 3.2.2, appears to be that even when the nonzero entries of α are clustered, the support of the optimal approximation of $\tilde{\mathbf{v}}$ in the **identify** step will not necessarily be clustered, and so CoSaMP will struggle to accurately identify this support.

Second, we remark that our simulations in Sections 3.2.1 and 3.2.2 have tested two extremes: one scenario in which the nonzero entries of α were well-separated, and one scenario in which the nonzero entries clustered together in a single block. Among the heuristic methods that we

⁶The occasional exception in some of these simulations occurred when it happened that $\|\mathbf{z} - \mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z}\|_2 \approx 0$ but $\|\mathcal{P}_{\Lambda_{\text{opt}}(\mathbf{z}, k)}\mathbf{z} - \mathcal{P}_{\mathcal{S}_{\mathbf{D}}(\mathbf{z}, k)}\mathbf{z}\|_2$ was not correspondingly small, and so the effective ϵ_2 value was large or infinite.

have used for attempting to find near-optimal supports $\mathcal{S}_{\mathcal{D}}(\mathbf{z}, k)$, the question of which method performs best has been shown to depend on the sparsity pattern of $\boldsymbol{\alpha}$. Although we do not present detailed results here, we have also tested these same algorithms using a hybrid sparsity model for $\boldsymbol{\alpha}$ in which half of the nonzero entries are well-separated while the other half are clustered. As one might expect based on the discussions above, all three of the SSCoSaMP methods struggle in this scenario (as do the three standard CS methods). This is yet another reminder that more work is needed to understand what techniques are appropriate for approximating $\mathcal{S}_{\mathcal{D}}(\mathbf{z}, k)$ and how to optimize these techniques depending on what is known about the signal's sparsity pattern.

A Proof of Theorem 2.1

The proof of Theorem 2.1 requires four main lemmas, which are listed below and proved in Sections A.1–A.4. In the lemmas below, $\mathbf{v} = \mathbf{x} - \mathbf{x}^\ell$ denotes the recovery error in signal space after ℓ iterations.

Lemma A.1 (Identify). $\|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2 \leq ((2 + \epsilon_1)\delta_{4k} + \epsilon_1) \|\mathbf{v}\|_2 + (2 + \epsilon_1)\sqrt{1 + \delta_{4k}} \|\mathbf{e}\|_2$.

Lemma A.2 (Merge). $\|\mathcal{P}_{T^\perp} \mathbf{x}\|_2 \leq \|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2$.

Lemma A.3 (Update). $\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \|\mathcal{P}_{T^\perp} \mathbf{x}\|_2 + \frac{2}{\sqrt{1 - \delta_{4k}}} \|\mathbf{e}\|_2$.

Lemma A.4 (Estimate). $\|\mathbf{x} - \mathbf{x}^{\ell+1}\|_2 \leq (2 + \epsilon_2) \|\mathbf{x} - \tilde{\mathbf{x}}\|_2$.

Combining all four statements above, we have

$$\begin{aligned}
\|\mathbf{x} - \mathbf{x}^{\ell+1}\|_2 &\leq (2 + \epsilon_2) \|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \\
&\leq (2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \|\mathcal{P}_{T^\perp} \mathbf{x}\|_2 + \frac{4 + 2\epsilon_2}{\sqrt{1 - \delta_{4k}}} \|\mathbf{e}\|_2 \\
&\leq (2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2 + \frac{4 + 2\epsilon_2}{\sqrt{1 - \delta_{4k}}} \|\mathbf{e}\|_2 \\
&\leq (2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} ((2 + \epsilon_1)\delta_{4k} + \epsilon_1) \|\mathbf{v}\|_2 \\
&\quad + (2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} (2 + \epsilon_1) \sqrt{1 + \delta_{4k}} \|\mathbf{e}\|_2 + \frac{4 + 2\epsilon_2}{\sqrt{1 - \delta_{4k}}} \|\mathbf{e}\|_2 \\
&= ((2 + \epsilon_1)\delta_{4k} + \epsilon_1)(2 + \epsilon_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \|\mathbf{x} - \mathbf{x}^\ell\|_2 + \left(\frac{(2 + \epsilon_2)((2 + \epsilon_1)(1 + \delta_{4k}) + 2)}{\sqrt{1 - \delta_{4k}}} \right) \|\mathbf{e}\|_2.
\end{aligned}$$

This completes the proof of Theorem 2.1.

A.1 Proof of Lemma A.1

In order to prove the four main lemmas, we require two supplemental lemmas, the first of which is a direct consequence of the \mathbf{D} -RIP.

Lemma A.5 (Consequence of \mathbf{D} -RIP). *For any index set B and any vector $\mathbf{z} \in \mathbb{C}^n$,*

$$\|\mathcal{P}_B \mathbf{A}^* \mathbf{A} \mathcal{P}_B \mathbf{z} - \mathcal{P}_B \mathbf{z}\|_2 \leq \delta_{|B|} \|\mathbf{z}\|_2.$$

Proof. We have

$$\begin{aligned}
\delta_{|B|} &\geq \sup_{\mathbf{x}: |B|\text{-sparse in } \mathbf{D}} \frac{|\|\mathbf{A}\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2|}{\|\mathbf{x}\|_2^2} \\
&\geq \sup_{\mathbf{x}} \frac{|\|\mathbf{A}\mathcal{P}_B\mathbf{x}\|_2^2 - \|\mathcal{P}_B\mathbf{x}\|_2^2|}{\|\mathcal{P}_B\mathbf{x}\|_2^2} \\
&\geq \sup_{\mathbf{x}} \frac{|\|\mathbf{A}\mathcal{P}_B\mathbf{x}\|_2^2 - \|\mathcal{P}_B\mathbf{x}\|_2^2|}{\|\mathbf{x}\|_2^2} \\
&= \sup_{\|\mathbf{x}\|_2=1} |\|\mathbf{A}\mathcal{P}_B\mathbf{x}\|_2^2 - \|\mathcal{P}_B\mathbf{x}\|_2^2| \\
&= \sup_{\|\mathbf{x}\|_2=1} |\langle \mathcal{P}_B^* \mathbf{A}^* \mathbf{A} \mathcal{P}_B \mathbf{x} - \mathcal{P}_B^* \mathcal{P}_B \mathbf{x}, \mathbf{x} \rangle| \\
&= \sup_{\|\mathbf{x}\|_2=1} |\langle \mathcal{P}_B \mathbf{A}^* \mathbf{A} \mathcal{P}_B \mathbf{x} - \mathcal{P}_B \mathbf{x}, \mathbf{x} \rangle| \\
&= \|\mathcal{P}_B \mathbf{A}^* \mathbf{A} \mathcal{P}_B - \mathcal{P}_B\|_2,
\end{aligned}$$

where the third line follows because $\|\mathcal{P}_B\mathbf{x}\|_2 \leq \|\mathbf{x}\|_2$ for all \mathbf{x} , and the last line follows from the fact that $\mathcal{P}_B \mathbf{A}^* \mathbf{A} \mathcal{P}_B - \mathcal{P}_B$ is self-adjoint. \square

We'll also utilize an elementary fact about orthogonal projections.

Lemma A.6. *For any pair of index sets A, B with $A \subset B$, $\mathcal{P}_A = \mathcal{P}_A \mathcal{P}_B$.*

Now, to make the notation simpler, note that $\tilde{\mathbf{v}} = \mathbf{A}^* \mathbf{A} \mathbf{v} + \mathbf{A}^* \mathbf{e}$ and that $\Omega = \mathcal{S}_{\mathbf{D}}(\tilde{\mathbf{v}}, 2k)$. Let $\Omega^* = \Lambda_{\text{opt}}(\tilde{\mathbf{v}}, 2k)$ denote the optimal support of size $2k$ for approximating $\tilde{\mathbf{v}}$, and set $R = \mathcal{S}_{\mathbf{D}}(\mathbf{v}, 2k)$. Using this notation we have

$$\begin{aligned}
\|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2 &= \|\mathbf{v} - \mathcal{P}_\Omega \mathbf{v}\|_2 \\
&\leq \|\mathbf{v} - \mathcal{P}_\Omega \tilde{\mathbf{v}}\|_2 \\
&= \|(\mathbf{v} - \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}) + (\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_{\Omega^*} \tilde{\mathbf{v}}) + (\mathcal{P}_{\Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Omega \tilde{\mathbf{v}})\|_2 \\
&\leq \|\mathbf{v} - \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}\|_2 + \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_{\Omega^*} \tilde{\mathbf{v}}\|_2 + \|\mathcal{P}_{\Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Omega \tilde{\mathbf{v}}\|_2 \\
&\leq \|\mathbf{v} - \mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v}\|_2 + \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2 + \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_{\Omega^*} \tilde{\mathbf{v}}\|_2 + \|\mathcal{P}_{\Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Omega \tilde{\mathbf{v}}\|_2, \quad (12)
\end{aligned}$$

where the second line follows from the fact that $\mathcal{P}_\Omega \mathbf{v}$ is the nearest neighbor to \mathbf{v} among all vectors in $\mathcal{R}(\mathbf{D}_\Omega)$ and the fourth and fifth lines use the triangle inequality.

Below, we will provide bounds on the first and second terms appearing in (12). To deal with the third term in (12), note that for any Π which is a subset of $R \cup \Omega^*$, we can write

$$\tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}} = (\tilde{\mathbf{v}} - \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}) + (\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}),$$

where $\tilde{\mathbf{v}} - \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}$ is orthogonal to $\mathcal{R}(\mathbf{D}_{R \cup \Omega^*})$, and $\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}$ is contained in $\mathcal{R}(\mathbf{D}_{R \cup \Omega^*})$. Thus we can write

$$\|\tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}\|_2^2 = \|\tilde{\mathbf{v}} - \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}\|_2^2 + \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}\|_2^2.$$

Recall that over all index sets Π with $|\Pi| = 2k$, $\|\tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}\|_2$ is minimized by choosing $\Pi = \Omega^*$. Thus, over all Π which are subsets of $R \cup \Omega^*$ with $|\Pi| = 2k$, $\|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_\Pi \tilde{\mathbf{v}}\|_2^2$ must be minimized

by choosing $\Pi = \Omega^*$. In particular, we have the first inequality below:

$$\begin{aligned}
\|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_{\Omega^*} \tilde{\mathbf{v}}\|_2 &\leq \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_R \tilde{\mathbf{v}}\|_2 \\
&= \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_R \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}\|_2 \\
&\leq \|\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_R(\mathbf{v} + \mathbf{A}^* \mathbf{e})\|_2 \\
&= \|(\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}) + (\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e} - \mathcal{P}_R \mathbf{A}^* \mathbf{e})\|_2 \\
&\leq \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 + \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e} - \mathcal{P}_R \mathbf{A}^* \mathbf{e}\|_2 \\
&\leq \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 + \|(I - \mathcal{P}_R) \mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2 \\
&\leq \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 + \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2. \tag{13}
\end{aligned}$$

The second line above uses Lemma A.6, the third line follows from the fact that $\mathcal{P}_R \mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}$ must be the nearest neighbor to $\mathcal{P}_{R \cup \Omega^*} \tilde{\mathbf{v}}$ among all vectors in $\mathcal{R}(\mathbf{D}_R)$, the fourth line uses the fact that $\mathcal{P}_R \mathbf{v} = \mathbf{v}$ because $R = \mathcal{S}_{\mathbf{D}}(\mathbf{v}, 2k)$ and both \mathbf{x} and \mathbf{x}^ℓ are k -sparse in \mathbf{D} , the fifth line uses the triangle inequality, the sixth line uses Lemma A.6, and the seventh line follows from the fact that $(I - \mathcal{P}_R)$ is an orthogonal projection and hence has norm bounded by 1.

To deal with the fourth term in (12), note that from the definition of Ω^* and from (4), we have the first inequality below:

$$\begin{aligned}
\|\mathcal{P}_{\Omega^*} \tilde{\mathbf{v}} - \mathcal{P}_{\Omega} \tilde{\mathbf{v}}\|_2 &\leq \epsilon_1 \|\mathcal{P}_{\Omega^*} \tilde{\mathbf{v}}\|_2 \\
&= \epsilon_1 \|\mathcal{P}_{\Omega^*} (\mathbf{A}^* \mathbf{A} \mathbf{v} + \mathbf{A}^* \mathbf{e})\|_2 \\
&\leq \epsilon_1 \|\mathcal{P}_{\Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v}\|_2 + \epsilon_1 \|\mathcal{P}_{\Omega^*} \mathbf{A}^* \mathbf{e}\|_2 \\
&= \epsilon_1 \|\mathcal{P}_{\Omega^*} \mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v}\|_2 + \epsilon_1 \|\mathcal{P}_{\Omega^*} \mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2 \\
&\leq \epsilon_1 \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v}\|_2 + \epsilon_1 \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2 \\
&\leq \epsilon_1 \|\mathbf{v}\|_2 + \epsilon_1 \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 + \epsilon_1 \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2, \tag{14}
\end{aligned}$$

The third line above uses the triangle inequality, the fourth line uses Lemma A.6, the fifth line uses the fact \mathcal{P}_{Ω^*} is an orthogonal projection and hence has norm bounded by 1, and the sixth line uses the triangle inequality.

Combining (12), (13), and (14) we see that

$$\|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2 \leq (2 + \epsilon_1) \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 + (2 + \epsilon_1) \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{e}\|_2 + \epsilon_1 \|\mathbf{v}\|_2.$$

Since $\mathbf{v} \in \mathcal{R}(\mathbf{D}_R)$, it follows that $\mathbf{v} \in \mathcal{R}(\mathbf{D}_{R \cup \Omega^*})$, and so

$$\|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathbf{v} - \mathbf{v}\|_2 = \|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^* \mathbf{A} \mathcal{P}_{R \cup \Omega^*} \mathbf{v} - \mathcal{P}_{R \cup \Omega^*} \mathbf{v}\|_2 \leq \delta_{4k} \|\mathbf{v}\|_2,$$

where we have used Lemma A.5 to get the inequality above. In addition, we know that the operator norm of $\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^*$ satisfies

$$\|\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^*\|_2 = \|(\mathcal{P}_{R \cup \Omega^*} \mathbf{A}^*)^*\|_2 = \|\mathbf{A} \mathcal{P}_{R \cup \Omega^*}\|_2 \leq \sqrt{1 + \delta_{4k}},$$

which follows from the \mathbf{D} -RIP. Specifically, for any \mathbf{x} ,

$$\frac{\|\mathbf{A} \mathcal{P}_{R \cup \Omega^*} \mathbf{x}\|_2}{\|\mathbf{x}\|_2} \leq \frac{\|\mathbf{A} \mathcal{P}_{R \cup \Omega^*} \mathbf{x}\|_2}{\|\mathcal{P}_{R \cup \Omega^*} \mathbf{x}\|_2} \leq \sqrt{1 + \delta_{4k}}.$$

Putting all of this together, we have

$$\|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2 \leq ((2 + \epsilon_1) \delta_{4k} + \epsilon_1) \|\mathbf{v}\|_2 + (2 + \epsilon_1) \sqrt{1 + \delta_{4k}} \|\mathbf{e}\|_2.$$

A.2 Proof of Lemma A.2

First note that by the definition of T , $\mathbf{x}^\ell \in \mathcal{R}(\mathbf{D}_T)$, and hence $\mathcal{P}_{T^\perp} \mathbf{x}^\ell = 0$. Thus we can write,

$$\|\mathcal{P}_{T^\perp} \mathbf{x}\|_2 = \left\| \mathcal{P}_{T^\perp} (\mathbf{x} - \mathbf{x}^\ell) \right\|_2 = \|\mathcal{P}_{T^\perp} \mathbf{v}\|_2.$$

Finally, since $\Omega \subseteq T$, we have that

$$\|\mathcal{P}_{T^\perp} \mathbf{v}\|_2 \leq \|\mathcal{P}_{\Omega^\perp} \mathbf{v}\|_2.$$

A.3 Proof of Lemma A.3

To begin, we note that $\mathbf{x} - \tilde{\mathbf{x}}$ has a $4k$ -sparse representation in \mathbf{D} , thus, applying the \mathbf{D} -RIP (of order $4k$) we have

$$\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq \frac{\|\mathbf{A}\mathbf{x} - \mathbf{A}\tilde{\mathbf{x}}\|_2}{\sqrt{1 - \delta_{4k}}}.$$

By construction,

$$\|\mathbf{A}\mathbf{x} - \mathbf{A}\tilde{\mathbf{x}} + \mathbf{e}\|_2 \leq \|\mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{z} + \mathbf{e}\|_2$$

for any $\mathbf{z} \in \mathcal{R}(\mathbf{D}_T)$, in particular for $\mathbf{z} = \mathcal{P}_T \mathbf{x}$. Thus,

$$\begin{aligned} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2 &\leq \frac{\|\mathbf{A}\mathbf{x} - \mathbf{A}\tilde{\mathbf{x}}\|_2}{\sqrt{1 - \delta_{4k}}} \\ &\leq \frac{\|\mathbf{A}\mathbf{x} - \mathbf{A}\tilde{\mathbf{x}} + \mathbf{e}\|_2 + \|\mathbf{e}\|_2}{\sqrt{1 - \delta_{4k}}} \\ &\leq \frac{\|\mathbf{A}\mathbf{x} - \mathbf{A}\mathcal{P}_T \mathbf{x} + \mathbf{e}\|_2 + \|\mathbf{e}\|_2}{\sqrt{1 - \delta_{4k}}} \\ &\leq \frac{\|\mathbf{A}\mathbf{x} - \mathbf{A}\mathcal{P}_T \mathbf{x}\|_2 + 2\|\mathbf{e}\|_2}{\sqrt{1 - \delta_{4k}}} \end{aligned}$$

where the second and fourth lines use the triangle inequality. By applying the \mathbf{D} -RIP we obtain

$$\|\mathbf{A}\mathbf{x} - \mathbf{A}\mathcal{P}_T \mathbf{x}\|_2 \leq \sqrt{1 + \delta_{4k}} \|\mathbf{x} - \mathcal{P}_T \mathbf{x}\|_2 = \sqrt{1 + \delta_{4k}} \|\mathcal{P}_{T^\perp} \mathbf{x}\|_2.$$

Combining all of this,

$$\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq \frac{\sqrt{1 + \delta_{4k}} \|\mathcal{P}_{T^\perp} \mathbf{x}\|_2 + 2\|\mathbf{e}\|_2}{\sqrt{1 - \delta_{4k}}}.$$

A.4 Proof of Lemma A.4

Using the triangle inequality, we have

$$\left\| \mathbf{x} - \mathbf{x}^{\ell+1} \right\|_2 = \left\| \mathbf{x} - \tilde{\mathbf{x}} + \tilde{\mathbf{x}} - \mathbf{x}^{\ell+1} \right\|_2 \leq \|\mathbf{x} - \tilde{\mathbf{x}}\|_2 + \left\| \tilde{\mathbf{x}} - \mathbf{x}^{\ell+1} \right\|_2.$$

Recall that $\Gamma = \mathcal{S}_D(\tilde{\mathbf{x}}, k)$ and $\mathbf{x}^{\ell+1} = \mathcal{P}_\Gamma \tilde{\mathbf{x}}$. Let $\Gamma^* = \Lambda_{\text{opt}}(\tilde{\mathbf{x}}, k)$ denote the optimal support of size k for approximating $\tilde{\mathbf{x}}$. Then we can write

$$\begin{aligned} \left\| \tilde{\mathbf{x}} - \mathbf{x}^{\ell+1} \right\|_2 &\leq \|\tilde{\mathbf{x}} - \mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}}\|_2 + \|\mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}} - \mathcal{P}_\Gamma \tilde{\mathbf{x}}\|_2 \\ &\leq \|\tilde{\mathbf{x}} - \mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}}\|_2 + \epsilon_2 \|\tilde{\mathbf{x}} - \mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}}\|_2, \end{aligned}$$

where the first line follows from the triangle inequality, and the second line uses (4). Combining all of this, we have

$$\begin{aligned} \left\| \mathbf{x} - \mathbf{x}^{\ell+1} \right\|_2 &\leq \left\| \mathbf{x} - \tilde{\mathbf{x}} \right\|_2 + (1 + \epsilon_2) \left\| \tilde{\mathbf{x}} - \mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}} \right\|_2 \\ &\leq \left\| \mathbf{x} - \tilde{\mathbf{x}} \right\|_2 + (1 + \epsilon_2) \left\| \tilde{\mathbf{x}} - \mathbf{x} \right\|_2 \\ &= (2 + \epsilon_2) \left\| \mathbf{x} - \tilde{\mathbf{x}} \right\|_2, \end{aligned}$$

where the second line follows from the fact that $\mathcal{P}_{\Gamma^*} \tilde{\mathbf{x}}$ is the nearest neighbor to $\tilde{\mathbf{x}}$ among all vectors having a k -sparse representation in \mathbf{D} .

References

- [1] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56(4):1982–2001, 2010.
- [2] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Const. Approx.*, 28(3):253–263, 2008.
- [3] T. Blumensath. Sampling and reconstructing signals from a union of linear subspaces. *IEEE Trans. Inform. Theory*, 57(7):4660–4671, 2011.
- [4] T. Blumensath and M. Davies. Iterative hard thresholding for compressive sensing. *Appl. Comput. Harmon. Anal.*, 27(3):265–274, 2009.
- [5] E. Candès. The restricted isometry property and its implications for compressed sensing. *Comptes rendus de l’Académie des Sciences, Série I*, 346(9-10):589–592, 2008.
- [6] E. Candès, Y. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Appl. Comput. Harmon. Anal.*, 31(1):59–73, 2011.
- [7] E. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Arxiv preprint arXiv:1203.5871*, 2012.
- [8] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51(12):4203–4215, 2005.
- [9] V. Cevher. An ALPS view of sparse recovery. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, 2011.
- [10] M. Davenport and M. Wakin. Analysis of orthogonal matching pursuit using the restricted isometry property. *IEEE Trans. Inform. Theory*, 56(9):4395–4401, 2010.
- [11] M. Davenport and M. Wakin. Compressive sensing of analog signals using discrete prolate spheroidal sequences. *Appl. Comput. Harmon. Anal.*, 33(3):438–472, 2012.
- [12] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [13] A. Garnaev and E. Gluskin. On widths of the euclidean ball. *Sov. Math. Dokl.*, 30:200–204, 1984.
- [14] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. Davies. Greedy-like algorithms for the cospase analysis model. *Arxiv preprint arXiv:1207.2456*, 2012.
- [15] P. Hansen. Regularization Tools Version 4.0 for Matlab 7.3. *Numerical Algorithms*, 46:189–194, 2007.
- [16] B. Kashin. The widths of certain finite dimensional sets and classes of smooth functions. *Izvestia*, 41:334–351, 1977.
- [17] F. Kraher and R. Ward. New and improved Johnson Lindenstrauss embeddings via the restricted isometry property. *SIAM J. Math. Anal.*, 43(3):1269–1281, 2011.
- [18] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Processing*, 41(12):3397–3415, 1993.
- [19] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Uniform uncertainty principle for Bernoulli and subgaussian ensembles. *Const. Approx.*, 28(3):277–289, 2008.
- [20] S. Nam, M. Davies, M. Elad, and R. Gribonval. The cospase analysis model and algorithms. *Arxiv preprint arXiv:1106.4987*, 2011.

- [21] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comput. Harmon. Anal.*, 26(3):301–321, 2009.
- [22] Y. Pati, R. Rezaifar, and P. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 1993.
- [23] D. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *J. ACM*, 9:84–97, 1962.
- [24] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Proc. IEEE Conf. Inform. Science and Systems (CISS)*, Princeton, NJ, Mar. 2006.
- [25] A. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Dokl. Akad. Nauk. SSSR*, 151:1035–1038, 1963.
- [26] A. Tikhonov and V. Arsenin. *Solutions of Ill-Posed Problems*. Winston & Sons, Washington, D.C., 1977.
- [27] E. van den Berg and M. Friedlander. Probing the Pareto frontier for basis pursuit solutions. *SIAM J. Sci. Comput.*, 31(2):890–912, 2008.
- [28] T. Zhang. Sparse recovery with orthogonal matching pursuit under RIP. *IEEE Trans. Inform. Theory*, 57(9):6215–6221, 2011.