

11-13-2015

A note on practical approximate projection schemes in signal space methods

Xiaoyi Gu

University of California, Los Angeles

Deanna Needell

Claremont McKenna College

Shenyinying Tu

University of California, Los Angeles

Recommended Citation

X. Gu, D. Needell, and S. Tu. "On practical approximate projection schemes in signal space methods." *SIAM Undergraduate Research Online*, 2015.

This Article 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.

A note on practical approximate projection schemes in signal space methods

Xiaoyi Gu, Deanna Needell and Shenying Tu *

November 13, 2015

Abstract

Compressive sensing (CS) is a new technology which allows the acquisition of signals directly in compressed form, using far fewer measurements than traditional theory dictates. Recently, many so-called *signal space* methods have been developed to extend this body of work to signals sparse in arbitrary dictionaries rather than orthonormal bases. In doing so, CS can be utilized in a much broader array of practical settings. Often, such approaches often rely on the ability to optimally project a signal onto a small number of dictionary atoms. Such optimal, or even approximate, projections have been difficult to derive theoretically. Nonetheless, it has been observed experimentally that conventional CS approaches can be used for such projections, and still provide accurate signal recovery. In this letter, we summarize the empirical evidence and clearly demonstrate for what signal types certain CS methods may be used as approximate projections. In addition, we provide theoretical guarantees for such methods for certain sparse signal structures. Our theoretical results match those observed in experimental studies, and we thus establish both experimentally and theoretically that these CS methods can be used in this context.

1 Introduction

Compressive sensing (CS) addresses some drawbacks of traditional signal acquisition by suggesting that a signal may be acquired directly in compressed form, and often with far less samples than previously thought. The key assumption behind this methodology is that signals of interest are *sparse* or nearly sparse. Typically, this notion can be quantified by assuming that the signal $\mathbf{x} \in \mathbb{C}^n$ can be represented as $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$ where $\boldsymbol{\alpha}$ is sparse ($\|\boldsymbol{\alpha}\|_0 := |\text{supp}(\boldsymbol{\alpha})| = s \ll n$) and \mathbf{D} is an orthonormal basis called the sparsifying basis. For example, if \mathbf{D} is a wavelet basis, most natural images are considered sparse in that basis. However, in many practical applications, the sparsifying basis is not orthonormal but instead a redundant, highly overcomplete frame. In this setting, most classical CS approaches break down, and new technologies are needed to address this practical issue. Recently, both optimization based and greedy methods have begun to be developed; however, many of these approaches rely on the existence of approximately optimal projections to project signals onto sparse subspaces of the frame. Empirical evidence has suggested that classical CS methods can be used for such projections, but these results have lacked theoretical support.

The mathematical formulation of the CS problem can be described as follows. Denote the signal of interest by $\mathbf{x} \in \mathbb{C}^n$. The measurement vector $\mathbf{y} \in \mathbb{C}^m$ is obtained using the measurement matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ (where $m \ll n$) and written $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, where $\mathbf{e} \in \mathbb{C}^m$ is additive noise. The sparsity condition on the signal \mathbf{x} is that $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$, for some coefficient vector $\boldsymbol{\alpha} \in \mathbb{C}^d$ with $\|\boldsymbol{\alpha}\|_0 \leq k \ll n$, yielding a k -sparse representation of the signal \mathbf{x} with respect to the dictionary $\mathbf{D} \in \mathbb{C}^{n \times d}$ ($n \leq d$).

*This work was supported by NSF grant DMS-1045536, NSF CAREER grant #1348721, and the Alfred P. Sloan Fellowship.

In the simplest classical CS setting, the sparsifying basis \mathbf{D} is just the identity matrix. In this setting, it is clear that as long as the measurement matrix \mathbf{A} is one-to-one on k -sparse vectors, then one can recover a k -sparse vector \mathbf{x} from its noiseless measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$ by searching for the sparsest vector among all those that match the measurements \mathbf{y} . This ℓ_0 -minimization problem of course is NP-hard in general [Mut05], but motivates the seminal work of Candès, Romberg and Tao [CT05, CRT06] which shows that a sparse vector \mathbf{x} can also be recovered using the relaxed ℓ_1 -minimization method, under a slightly stricter assumption on the matrix \mathbf{A} . Candès and Tao [CT05] introduce the Restricted Isometry Property (RIP) which asks that \mathbf{A} satisfy the following for a constant $\delta_k \in (0, 1)$:

$$(1 - \delta_k)\|\mathbf{x}\|_2^2 \leq \|\mathbf{A}\mathbf{x}\|_2^2 \leq (1 + \delta_k)\|\mathbf{x}\|_2^2 \quad \text{for all } \mathbf{x} \text{ with } \|\mathbf{x}\|_0 \leq k. \quad (1.0.1)$$

Fortunately, it is now well-known that many classes of randomly constructed measurement matrices satisfy this condition. For example, if the entries of \mathbf{A} are drawn from a subgaussian distribution and $m \geq Ck/\log(n)$, then \mathbf{A} satisfies the RIP with high probability [CRVT05, RV08]. Analogous results hold for subsampled Discrete Fourier Transform (DFT) matrices and others with a fast-multiply [RV08, RRT12]. Under the assumption of the RIP, the ℓ_1 -minimization problem is guaranteed to robustly reconstruct a k -sparse signal \mathbf{x} from its noisy measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$. Greedy methods such as OMP [TG07, Zha11], ROMP [NV07], CoSaMP [NT09], and IHT [BD09], which aim to identify the signal support iteratively, also provide robust reconstruction guarantees in this setting. We refer the reader to the above references for details about each of the algorithmic guarantees.

2 Signal Space Methods

The above mentioned CS methods provide rigorous reconstruction guarantees when the signal of interest \mathbf{x} is sparse in an orthonormal basis. Indeed, if we write $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$ for some sparse coefficient vector $\boldsymbol{\alpha}$, these approaches aim at recovering the representation $\boldsymbol{\alpha}$. When \mathbf{D} is orthonormal, this of course translates robustly to the recovery of the actual signal \mathbf{x} . However, a vast array of signals in practice are compressible not in an orthonormal basis but in some highly overcomplete dictionary such as an over-sampled DFT, Gabor frame, or many of the redundant dictionaries used in signal and image processing. Due to the effect of redundancy, these classical approaches fail both theoretically and empirically when \mathbf{D} is no longer orthonormal.

To address this issue, a recent surge of work has been focused on so-called *signal space* methods. Two models to capture sparsity in an overcomplete frame have been studied; the synthesis sparsity model asserts that $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$ for some sparse representation $\boldsymbol{\alpha}$, as discussed above. The analysis sparsity or *cosparsity* model instead captures the sparsity in the analysis coefficients $\mathbf{D}^*\mathbf{x}$ [NDEG13, GNE⁺14, Fou]. We utilize the synthesis sparsity model in this work, and refer the reader to aforementioned references for a description of some alternatives.

In [DW12, DNW12], Davenport et. al. introduce and analyze a greedy method called Signal Space CoSaMP (SSCoSaMP). SSCoSaMP is an adaptation of the CoSaMP method designed to recover the signal \mathbf{x} when \mathbf{D} is an overcomplete frame. Rather than assuming the classical RIP, they instead utilize a generalization of this property called the D -RIP, introduced in [CENR10]. In contrast to other CS algorithms, SSCoSaMP employs a “signal-focused” approach, aimed at determining the signal \mathbf{x} directly rather than its (non-unique) coefficient vector $\boldsymbol{\alpha}$. The method is described by the pseudo-code given in Algorithm 1. Here and throughout, we denote the range of a matrix \mathbf{D} by $\mathcal{R}(\mathbf{D})$.

The identification step in SSCoSaMP requires finding the best k -sparse representation of a vector \mathbf{z} in the dictionary \mathbf{D} :

$$\Omega_{\text{opt}} := \underset{\Lambda:|\Lambda|=k}{\operatorname{argmin}} \|\mathbf{z} - \mathcal{P}_{\Lambda}\mathbf{z}\|_2, \quad (2.0.2)$$

where here and throughout \mathcal{P}_Λ denotes the projection onto the span of the columns of \mathbf{D} indexed by Λ . This problem is itself reminiscent of the classical CS problem; one wishes to recover a sparse representation from an underdetermined linear system. Thus, such an endeavor in general is itself an NP-hard problem. Therefore, we relax this problem and instead ask for a *near-optimal* approximation [DNW12, GN15], writing $\mathcal{S}_\mathbf{D}(\mathbf{w}, k)$ to denote the k -sparse approximation to \mathbf{w} in \mathbf{D} . Because of its similarity to the classical CS problem, one might think to use a classical CS approach to solve this sub-problem. Of course, for typical redundant frames, the RIP will not be in force, and classical results would not suggest accurate recovery. Nonetheless, SSCoSaMP is surprisingly able to accurately recover signals even when \mathbf{D} is highly overcomplete, when using a classical CS algorithm like OMP, CoSaMP or ℓ_1 -minimization for the near-optimal projection $\mathcal{S}_\mathbf{D}$.

Even more interesting is that the behavior of SSCoSaMP varies significantly depending on both the approximate projection used, and the structure of the sparse signal \mathbf{x} . This was originally noted in the experiments of [DNW12], which we demonstrate again here in Figure 1 as in [GGK⁺14]. Here and throughout this paper, the algorithm in parentheses following ‘‘SSCoSaMP’’ is the algorithm used for the approximate projection $\mathcal{S}_\mathbf{D}$ in the identify and prune steps. As is shown, when the support of the signal \mathbf{x} is clustered together, the CoSaMP method as an approximate projection yields accurate recovery whereas ℓ_1 -minimization and OMP do not perform well at all. On the other hand, when the signal support has adequate separation, the exact opposite behavior is seen. This observation led to a recent work which empirically investigated this behavior, cataloging how various approximate projection methods behave for various signal types [GGK⁺14].

Algorithm 1 Signal-Space CoSaMP (SSCoSaMP)

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: $\tilde{\mathbf{w}} = \arg\min_{\mathbf{z}} \|\mathbf{y} - \mathbf{A}\mathbf{z}\|_2$ s.t. $\mathbf{z} \in \mathcal{R}(\mathbf{D}_T)$
 Prune: $\Gamma = \mathcal{S}_\mathbf{D}(\tilde{\mathbf{w}}, k)$
 $\mathbf{x}_{\ell+1} = \mathcal{P}_\Gamma \tilde{\mathbf{w}}$
 $\mathbf{r} = \mathbf{y} - \mathbf{A}\mathbf{x}_{\ell+1}$
 $\ell = \ell + 1$
end while
Output: $\hat{\mathbf{x}} = \mathbf{x}_\ell$

2.1 Contribution

As observed in Figure 1 and in more detail in [GGK⁺14], the accuracy of SSCoSaMP depends on both the approximate projection and the structure of the signal. In addition, none of the approximate projections simultaneously work well for both signal types (however, see [GE13, GGK⁺14] for hybrid algorithms that are accurate for wider variety of signal structures). The discrepancy between the theoretical guarantees for classical CS approaches and the empirical behavior of such methods as approximate projections motivates the work of this note. Here, we provide an analysis of two of the classical approaches which provably show that for certain signal structures, these methods can be used as accurate approximate projections and thus can be utilized in SSCoSaMP for efficient signal recovery. We believe our analysis

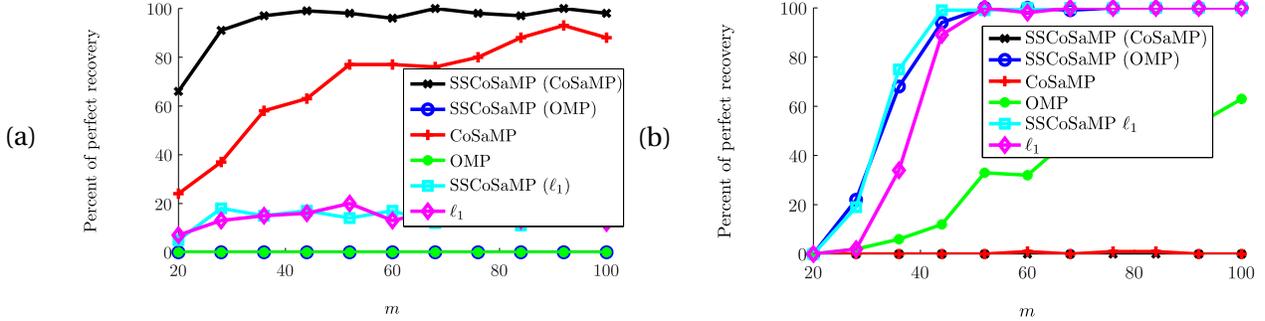


Figure 1: From [GGK⁺14]. Percent of perfect recovery (we define perfect recovery as signal to noise ratio $20 \log(\|\mathbf{x}\|_2 / \|\mathbf{x} - \hat{\mathbf{x}}\|_2)$ greater than 100db) out of 100 trials for different SSCoSaMP variants when the nonzero entries in α are clustered together (left) and when the nonzero entries in α are well-separated (right). Here, $k = 8$, $n = 256$, $d = 1024$, the dictionary $\mathbf{D} \in \mathbb{C}^{n \times d}$ is a $4 \times$ overcomplete DFT and $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a Gaussian matrix.

further the understanding of the behavior of such methods, which is interesting in its own right as well as in the context of signal space methods.

3 Analytical Justification

Recall that (2.0.2) is NP-hard because it requires examining all k combinations of the columns of \mathbf{D} . A concession one makes is to instead look for a near-optimal projection, as used in Algorithm 1. It has been shown [GN15] that as long as the near optimal projection is good enough, i.e.,

$$\|\mathcal{P}_{\mathcal{S}_D(x,k)}\mathbf{x} - \mathbf{x}\|_2 \leq C\|\mathcal{P}_\Lambda\mathbf{x} - \mathbf{x}\|_2 \quad \text{and} \quad \|\mathcal{P}_{\mathcal{S}_D(x,k)}\mathbf{x}\|_2 \geq c\|\mathcal{P}_\Lambda\mathbf{x}\|_2 \quad (3.0.1)$$

for all \mathbf{x} and suitable constants c and C (and where \mathcal{P}_Λ denotes the optimal projection), then SSCoSaMP provides accurate recovery of the signal. Although there do exist such projections for well behaved dictionaries \mathbf{D} , for truly redundant dictionaries such projections are not known to exist. Nonetheless, empirical studies [DNW12] using traditional compressed sensing algorithms for these projections showed that in certain cases SSCoSaMP using these projections still yields accurate recovery. In this section we build on state of the art results in compressed sensing to provide theoretical justification for this behavior; motivated by Fig. 1, we examine the case when SSCoSaMP(ℓ_1) and SSCoSaMP(OMP) are applied to well-separated signals. We assume here that the dictionary \mathbf{D} is an overcomplete discrete Fourier matrix (DFT); results for similar dictionaries follow analogously.

3.1 SSCoSaMP(OMP)

Numerical experiments conducted on SSCoSaMP(OMP) suggest that when the sparse vector α in $\mathbf{x} = \mathbf{D}\alpha$ is well-separated, SSCoSaMP(OMP) gives accurate recovery of the signal \mathbf{x} . We thus seek theoretical backing for the use of OMP as an approximate projection within SSCoSaMP in this scenario. First, suppose the signal \mathbf{x} has a k -sparse representation in the $n \times d$ matrix Φ ,

$$\mathbf{x} = \sum_{i \in \Omega_{\text{opt}}} \alpha_i \boldsymbol{\varphi}_i \quad (3.1.1)$$

where Ω_{opt} is an index set of size k , and $\boldsymbol{\varphi}_i$ and $\boldsymbol{\alpha}_i$ denote the i th column of $\boldsymbol{\Phi}$ and element of $\boldsymbol{\alpha}$, respectively. Then we can extract the matrix $\boldsymbol{\Phi}_{\Omega_{\text{opt}}}$ from the dictionary $\boldsymbol{\Phi}$ whose columns are listed in Ω_{opt} ,

$$\boldsymbol{\Phi}_{\Omega_{\text{opt}}} = [\boldsymbol{\varphi}_{i_1} \ \boldsymbol{\varphi}_{i_2} \ \cdots \ \boldsymbol{\varphi}_{i_k}]_{i_1, \dots, i_k \in \Omega_{\text{opt}}}. \quad (3.1.2)$$

Then the signal can be expressed as

$$\mathbf{x} = \boldsymbol{\Phi}_{\Omega_{\text{opt}}} \boldsymbol{\alpha}_{\Omega_{\text{opt}}}. \quad (3.1.3)$$

We will utilize the following established result of Cai and Wang [CW11] that improves upon the seminal work of Gilbert and Tropp [TG07]. They prove a sufficient condition to accurately recover the signal from contaminated samples when $\boldsymbol{\Phi}$ is sufficiently incoherent. Also see these works for a detailed description of OMP.

Proposition 3.1. [CW11, Prop. 1] *Let Ω_{opt} denote the support set of the signal $\boldsymbol{\alpha}$, set $M := \max_{i \in \Omega_{\text{opt}}^c} \|\boldsymbol{\Phi}_{\Omega_{\text{opt}}}^\dagger \boldsymbol{\phi}_i\|_1$, let \mathbf{r}_i be the residual vector of OMP in the i th iteration, and let λ_{\min} denote the minimal eigenvalue of the matrix $\boldsymbol{\Phi}_{\Omega_{\text{opt}}}^* \boldsymbol{\Phi}_{\Omega_{\text{opt}}}$. Suppose $\|\mathbf{e}\|_2 \leq \varepsilon$ and $M < 1$. Then the OMP algorithm (with the stopping rule $\|\mathbf{r}_i\|_2 \leq \varepsilon$) recovers Ω_{opt} exactly if all the nonzero coefficients $\boldsymbol{\alpha}_i$ satisfy*

$$|\boldsymbol{\alpha}_i| \geq \frac{2\varepsilon}{(1-M)\lambda_{\min}} \quad (3.1.4)$$

In our setting, $\mathbf{D} = \boldsymbol{\Phi}$ is not incoherent, but we want to show that when the support of $\boldsymbol{\alpha}$ is well-separated, then $\max_{i \in \Omega_{\text{opt}}^c} \|\boldsymbol{\Phi}_{\Omega_{\text{opt}}}^\dagger \boldsymbol{\phi}_i\|_1 \leq 1$. From now on, we take \mathbf{D} to be an $n \times d$ ($n \leq d$) overcomplete DFT matrix (having unit-norm columns), but the technique can be extended to other overcomplete dictionaries. Recall that if \mathbf{D} is an $n \times d$ overcomplete DFT dictionary, then the entries of \mathbf{D} obey (with our normalization):

$$\mathbf{D}_{jk} = \frac{\omega^{jk}}{\sqrt{n}}_{j=0,1,\dots,n-1,k=0,1,\dots,d-1}, \quad \text{where } \omega = e^{-\frac{2\pi i}{d}}. \quad (3.1.5)$$

We will first define a separation condition that will be needed to guarantee accurate recovery.

Definition 3.2 (Well-separated). *We say that Ω_{opt} is well-separated when $\mathbf{D}_{\Omega_{\text{opt}}}^* \mathbf{D}_{\Omega_{\text{opt}}}$ is strictly diagonally dominant. Recall that we may define for a $k \times k$ matrix \mathbf{J} ,*

$$\Delta_i(\mathbf{J}) := J_{ii} - \sum_{j \neq i} |J_{ij}|, \quad (3.1.6)$$

and then say \mathbf{J} is strictly diagonally dominant if $\Delta_i(\mathbf{J}) > 0$ for $1 \leq i \leq k$. We refer to a signal $\boldsymbol{\alpha}$ as well-separated when its support set Ω_{opt} satisfies this property.

Note that once the minimum separation¹ between any two non-zero elements of $\boldsymbol{\alpha}$, denoted h_{\min} , is large enough, Ω_{opt} is well-separated. Indeed, this is easily due to the fact that the gram matrix of the dictionary \mathbf{D} has quickly decaying off-diagonal terms (see Fig. 2). We can quantify this notion by the following definition.

¹Note that we measure distance in the DFT dictionary cyclically, so that column d and column 1 are 1 column apart.

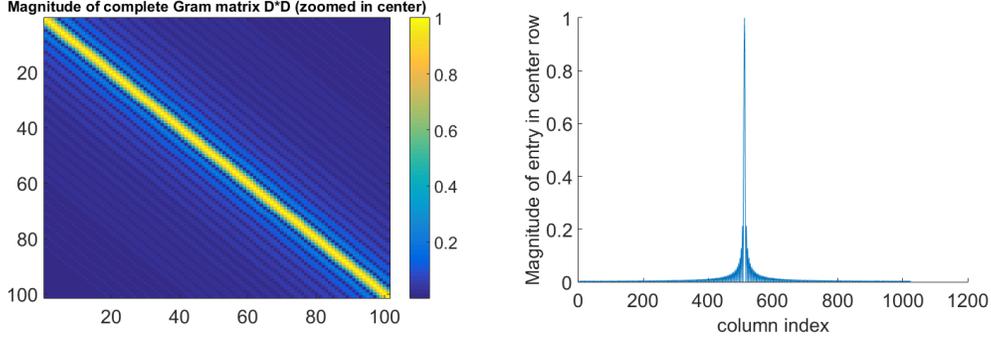


Figure 2: Left: Magnitudes at center of gram matrix $\mathbf{D}^* \mathbf{D}$. Right: Magnitudes in center row.

Definition 3.3 (Dominance factor). For an $n \times d$ DFT dictionary \mathbf{D} , let $\mathcal{S}_{k, h_{\min}}$ denote the set of all support sets Ω of size k that have a minimum separation of at least h_{\min} :

$$\mathcal{S}_{k, h_{\min}} = \left\{ \Omega : |\Omega| = k, \max_{p, q \in \Omega, p > q} \min(p - q, d - p + q) \geq h_{\min} \right\}.$$

We define the dominance factor $\eta_{h_{\min}, k}$ as the worst possible sum of off diagonals $\sum_{q \neq p} |\mathbf{G}_{pq}|$ for $\mathbf{G} = \mathbf{D}_{\Omega}^* \mathbf{D}_{\Omega}$ when Ω has k entries and minimum separation h_{\min} . That is,

$$\eta_{h_{\min}, k} := \max_{\Omega \in \mathcal{S}_{k, h_{\min}}} \max_{p \in \Omega} \sum_{q \in \Omega, q \neq p} |(\mathbf{D}^* \mathbf{D})_{pq}|. \quad (3.1.7)$$

We can also use this quantity to bound the correlations between atoms in a support set and an outside atom. We may define

$$\eta'_{h_{\min}, k} := \max_{\Omega \in \mathcal{S}_{k, h_{\min}}} \max_{p \notin \Omega} \sum_{q \in \Omega} |(\mathbf{D}^* \mathbf{D})_{pq}|. \quad (3.1.8)$$

With this notation, one observes that for a fixed sparsity level k , if $\eta_{h_{\min}, k} < (\mathbf{D}^* \mathbf{D})_{qq}$ (where with our normalization $(\mathbf{D}^* \mathbf{D})_{qq} = 1$) then any support set Ω_{opt} of size k and minimum separation h_{\min} is guaranteed to satisfy the well-separated condition. In this sense, utilizing these quantities will yield “worst-case” bounds, where we fix h_{\min} and provide bounds for the worst possible support set. Our bounds will thus be highly pessimistic but can be presented in a more simple form, depending only on h_{\min} . We also remark that $\eta_{h_{\min}, k} \leq \eta'_{h_{\min}, k}$ but that $\eta_{h_{\min}, k} \approx \eta'_{h_{\min}, k}$ as the redundancy in the dictionary grows large. Nonetheless, we treat these two quantities separately since for some dictionaries they may be significantly different.

Figure 3 shows numerically how separated the support of the signal needs to be to satisfy the well-separated condition². This again confirms that as long as the minimum separation h_{\min} is large enough, the well-separated property holds. Of course note that this computation verifies the minimum separation required for the *worst case* support with that separation, whereas there may be many other support sets with smaller minimum separation which still guarantee the well-separated property. Nonetheless, we will continue to utilize this quantity to bound the performance of OMP so that we may state the results only in terms of the separation h_{\min} .

We also plot a (crude) upper bound³ on $\eta'_{h_{\min}, k}$ in Figure 3. We will see later that we also want this quantity to be smaller than one.

Finally, we will utilize the following lemma, which bounds the restricted isometry constant when the signal has a well-separated support Ω_{opt} .

²To obtain a crude upper bound on $\eta_{h_{\min}, k}$ we simply assume that Ω_{opt} contains equally separated non-zero elements, a distance h_{\min} apart and sum the coherence between those atoms and the worst off-support column. For that reason, when

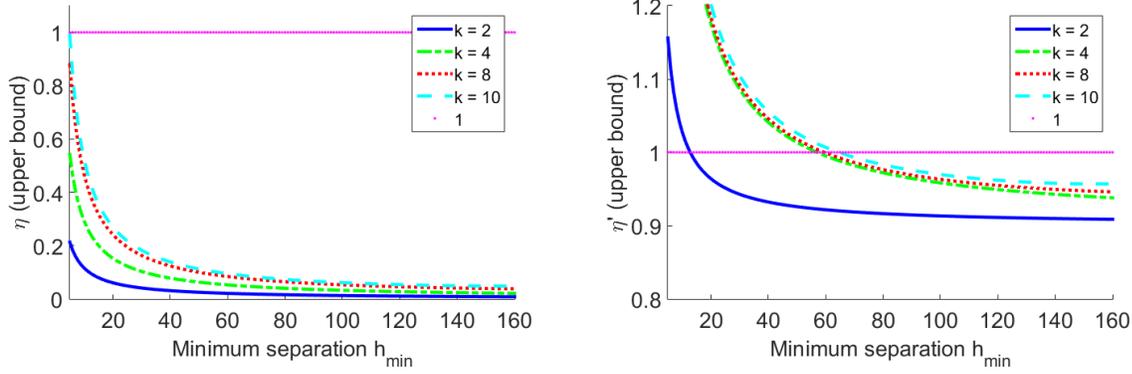


Figure 3: Dictionary \mathbf{D} is 256×1024 overcomplete DFT. Left: upper bound on the dominance factor $\eta_{h_{\min},k}$ for support size k and minimum separation h_{\min} . Horizontal line shows value of diagonal entry. When curve falls below horizontal line the well-separated property is guaranteed to hold. Right: upper bound on $\eta'_{h_{\min},k}$ as a function of h_{\min} .

Lemma 3.4. \mathbf{D} satisfies the following variant of the Restricted Isometry Property (1.0.1) with parameters (k, δ_k) :

$$(1 - \delta_k) \|\alpha\|_2^2 \leq \|\mathbf{D}\alpha\|_2^2 \leq (1 + \delta_k) \|\alpha\|_2^2 \quad (3.1.9)$$

for any well-separated k -sparse vector α . In particular, if the minimum separation distance between the nonzeros of α is denoted by h_{\min} , then

$$\delta_k \leq \eta_{h_{\min},k}, \quad (3.1.10)$$

where $\eta_{h_{\min},k}$ is defined in (3.1.7).

Proof. Fix a signal α with well-separated support $\Omega = \Omega_{\text{opt}}$, set $\mathbf{G} := \mathbf{D}_{\Omega}^* \mathbf{D}_{\Omega}$ and let g_{pq} be the entries of \mathbf{G} . According to the Gershgorin circle theorem [Ger31], for every singular value λ_i of \mathbf{G} , we have

$$1 - R_p \leq \lambda_i \leq 1 + R_p \quad (3.1.11)$$

for some $p \in \Omega$, where $R_p = \sum_{q \neq p} |g_{pq}|$ is the radius of the Gershgorin disc. By definition,

$$\|\mathbf{G} - \mathbf{I}\mathbf{d}\|_2 = \max_i (\lambda_i - 1) \leq \max_{p \in \Omega} R_p = \max_{p \in \Omega} \sum_{q \in \Omega, q \neq p} |g_{pq}|. \quad (3.1.12)$$

Since we know the minimum distance (with respect to the full dictionary \mathbf{D}) between any two columns in $\mathbf{D}_{\Omega_{\text{opt}}}$ is h_{\min} , then we have by the definition (3.1.7) that

$$\|\mathbf{G} - \mathbf{I}\mathbf{d}\|_2 = \max_{p \in \Omega} \sum_{q \in \Omega, q \neq p} |g_{pq}| \leq \eta_{h_{\min},k}. \quad (3.1.13)$$

Lastly, the following observation proves the claim:

$$\|\mathbf{G} - \mathbf{I}\mathbf{d}\|_2 = \sup_{\|\alpha\|_2=1} |\langle (\mathbf{D}_{\Omega_{\text{opt}}}^* \mathbf{D}_{\Omega_{\text{opt}}} - \mathbf{I}\mathbf{d}) \alpha, \alpha \rangle| = \sup_{\|\alpha\|_2=1} \left| \|\mathbf{D}\alpha\|_2^2 - \|\alpha\|_2^2 \right|. \quad (3.1.14)$$

□

this bound is less than one, we guarantee that any support set with separation h_{\min} is well-separated.

³To bound $\eta'_{h_{\min},k}$, we first upper bound the coherence between two columns separated by a distance h by $f(h) := 1/n \cdot \text{csc}(h\pi/d)$. Then one observes that because $f(h)$ is convex, the maximum in (3.1.8) is attained when column p is a neighbor of an element in Ω_{opt} . Hence $\eta'_{h_{\min},k}$ is bounded by the sum of the coherence between columns in Ω_{opt} with respect to such a column p . The bound can then be written formally as $f(1) + \sum_{j=1}^r f(jh+1) + f(jh-1)$ where $r = \lfloor (k+1)/2 \rfloor$.

3.1.1 Guarantees for SSCoSAMP(OMP)

We consider here the case when the measurements are corrupted with noise, allowing the noiseless case to follow as a special instance when the noise has norm zero. If our signal is contaminated by noise, then (3.1.1) becomes

$$\mathbf{x} = \sum_{i \in \Omega_{opt}} \boldsymbol{\alpha}_i \boldsymbol{\phi}_i + \mathbf{e}, \quad \|\mathbf{e}\|_2 \leq \varepsilon. \quad (3.1.15)$$

We are now prepared to state our main result, which guarantees accurate recovery of well-separated signals under modest amounts of noise.

Theorem 3.5. *Suppose that $\|\mathbf{e}\|_2 < \varepsilon$, and*

$$B(h_{\min}) := \frac{\eta'_{h_{\min},k}}{1 - \eta_{h_{\min},k}} < 1, \quad (3.1.16)$$

where both terms are defined in (3.1.7) and (3.1.8). Then OMP (with the stopping rule $\|\mathbf{r}_i\|_2 \leq \varepsilon$) will exactly recover Ω_{opt} if the minimum magnitude of the nonzero elements of the signal $\boldsymbol{\alpha}$ satisfies

$$\min_{i \in \Omega_{opt}} |\boldsymbol{\alpha}_i| \geq \frac{2\varepsilon}{1 - \eta_{h_{\min},k} - \eta'_{h_{\min},k}}. \quad (3.1.17)$$

Proof. The proof of this result follows in the same fashion as the proof of [CW11, Prop. 1]. Denote the set of all selected indices in the first t iterations by Ω_t , and by $\Omega_t^* := \Omega_{opt} \setminus \Omega_t$ the remaining coefficients of the support yet to be identified. We maintain the notation that the dictionary \mathbf{D} has columns $\boldsymbol{\phi}_i$ for $i = 1, 2, \dots, d$.

We proceed by induction on t , and suppose that $\mathbf{D}_{\Omega_t} \subset \mathbf{D}_{\Omega_{opt}}$ after t iterations. Define $\mathcal{P}_t := \mathbf{D}_{\Omega_t} \mathbf{D}_{\Omega_t}^\dagger$. Then the residual after t steps is

$$\mathbf{r}_t = (\mathbf{I} - \mathcal{P}_t) \mathbf{y} = (\mathbf{I} - \mathcal{P}_t) \mathbf{D} \boldsymbol{\alpha} + (\mathbf{I} - \mathcal{P}_t) \mathbf{e} =: \mathbf{s}_t + \mathbf{n}_t \quad (3.1.18)$$

where $\mathbf{s}_t := (\mathbf{I} - \mathcal{P}_t) \mathbf{D} \boldsymbol{\alpha}$ and $\mathbf{n}_t := (\mathbf{I} - \mathcal{P}_t) \mathbf{e}$. Let

$$M_t = \max_{i \in \Omega_t^*} |\boldsymbol{\phi}_i^* \mathbf{s}_t|, \quad M'_t = \max_{i \in \Omega_{opt}^c} |\boldsymbol{\phi}_i^* \mathbf{s}_t|, \quad N_t = \max_i |\boldsymbol{\phi}_i^* \mathbf{n}_t|. \quad (3.1.19)$$

In order for OMP to select a correct index at the next iteration, we need $\max_{i \in \Omega_t^*} |\boldsymbol{\phi}_i^* \mathbf{r}_t| > \max_{i \in \Omega_{opt}^c} |\boldsymbol{\phi}_i^* \mathbf{r}_t|$. It suffices to require that $M_t - M'_t > 2N_t$ since that would imply

$$\max_{i \in \Omega_t^*} |\boldsymbol{\phi}_i^* \mathbf{r}_t| \geq M_t - N_t > M'_t + N_t \geq \max_{i \in \Omega_{opt}^c} |\boldsymbol{\phi}_i^* \mathbf{r}_t|. \quad (3.1.20)$$

We utilize the following lemma.

Lemma 3.6. [Tro04, Proof of Theorem 4.2] *In our notation, $B(h_{\min})M_t > M'_t$ for all t .*

By this lemma, $M_t - M'_t > (1 - B(h_{\min}))M_t$. Hence it suffices to require $M_t > \frac{2}{1 - B(h_{\min})} N_t$ for OMP to select a correct index. Then we have by definition of Ω_t^* ,

$$M_t = \|\mathbf{D}_{\Omega_t^*}^* \mathbf{s}_t\|_\infty = \|\mathbf{D}_{\Omega_t^*}^* (\mathbf{I} - \mathcal{P}_t) \mathbf{D} \boldsymbol{\alpha}\|_\infty = \|\mathbf{D}_{\Omega_t^*}^* (\mathbf{I} - \mathcal{P}_t) \mathbf{D}_{\Omega_t^*} \boldsymbol{\alpha}_{\Omega_t^*}\|_\infty. \quad (3.1.21)$$

The following lemma allows us to relate the above expression to $\mathbf{D}_{\Omega_{opt}}^* \mathbf{D}_{\Omega_{opt}}$.

Lemma 3.7. [CW11, Lemma 5] *In our notation, $\lambda_{\min}(\mathbf{D}_{\Omega_{opt}}^* \mathbf{D}_{\Omega_{opt}}) \leq \lambda_{\min}(\mathbf{D}_{\Omega_t^*}^* (\mathbf{I} - \mathcal{P}_t) \mathbf{D}_{\Omega_t^*})$.*

Thus we can bound,

$$\|\mathbf{D}_{\Omega_t^*}^*(\mathbf{I} - \mathcal{P}_t)\mathbf{D}_{\Omega_t^*}\boldsymbol{\alpha}_{\Omega_t^*}\|_2 \geq \lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})\|\boldsymbol{\alpha}_{\Omega_t^*}\|_2. \quad (3.1.22)$$

Combining this with the fact that $|\Omega_t^*| = k - t$, we have

$$M_t \geq \frac{1}{\sqrt{k-t}}\|\mathbf{D}_{\Omega_t^*}^*\mathbf{s}_t\|_2 = \frac{1}{\sqrt{k-t}}\|\mathbf{D}_{\Omega_t^*}^*(\mathbf{I} - \mathcal{P}_t)\mathbf{D}_{\Omega_t^*}\boldsymbol{\alpha}_{\Omega_t^*}\|_2 \geq \frac{\lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})}{\sqrt{k-t}}\|\boldsymbol{\alpha}_{\Omega_t^*}\|_2. \quad (3.1.23)$$

Therefore, for the sufficient condition $M_t - M'_t > 2N_t$ to hold, it is enough that

$$\|\boldsymbol{\alpha}_{\Omega_t^*}\|_2 > \frac{2\sqrt{k-t}N_t}{(1 - B(h_{\min}))\lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})}. \quad (3.1.24)$$

Since $\|\mathbf{e}\|_2 \leq \varepsilon$,

$$\|\mathbf{n}_t\|_2 = \|(\mathbf{I} - \mathcal{P}_t)\mathbf{e}\|_2 \leq \|\mathbf{e}\|_2 \leq \varepsilon. \quad (3.1.25)$$

Thus for any column $\boldsymbol{\phi}_i$ of \mathbf{D} ,

$$|\boldsymbol{\phi}_i^*\mathbf{n}_t| \leq \|\boldsymbol{\phi}_i\|_2\|\mathbf{n}_t\|_2 \leq \varepsilon, \quad (3.1.26)$$

which implies that $N_t \leq \varepsilon$. Since $\lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}}) \geq 1 - \eta_{h_{\min},k}$ by Lemma 3.4, we have that

$$\|\boldsymbol{\alpha}_{\Omega_t^*}\|_2 \geq \frac{2\sqrt{k-t}\varepsilon}{(1 - B(h_{\min}))(1 - \eta_{h_{\min},k})} = \frac{2\sqrt{k-t}\varepsilon}{(1 - \eta_{h_{\min},k} - \eta'_{h_{\min},k})} \quad (3.1.27)$$

suffices to ensure that a correct index will be selected at this step. Observing that $|\Omega_t^*| = k - t$ yields the desired condition (3.1.17).

It remains to argue that the stopping criterion is sufficient for termination. It suffices to show that at each step OMP doesn't stop early – that with this stopping criterion it runs the full k iterations. Equivalently, we want to show that for $t < k$, $\|\mathbf{r}_t\|_2 > \varepsilon$. We observe that

$$\begin{aligned} \|\mathbf{r}_t\|_2 &= \|(\mathbf{I} - \mathcal{P}_t)\mathbf{D}\boldsymbol{\alpha} + (\mathbf{I} - \mathcal{P}_t)\mathbf{e}\|_2 \\ &\geq \|(\mathbf{I} - \mathcal{P}_t)\mathbf{D}\boldsymbol{\alpha}\|_2 - \|(\mathbf{I} - \mathcal{P}_t)\mathbf{e}\|_2 \\ &\geq \|(\mathbf{I} - \mathcal{P}_t)\mathbf{D}_{\Omega_t^*}\boldsymbol{\alpha}_{\Omega_t^*}\|_2 - \varepsilon. \end{aligned} \quad (3.1.28)$$

Again using Lemma 3.7 along with (3.1.27), we have for $t \leq k$,

$$\begin{aligned} \|(\mathbf{I} - \mathcal{P}_t)\mathbf{D}_{\Omega_t^*}\boldsymbol{\alpha}_{\Omega_t^*}\|_2 &\geq \lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})\|\boldsymbol{\alpha}_{\Omega_t^*}\|_2 \\ &\geq \lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})\frac{2\sqrt{k-t}\varepsilon}{(1 - B(h_{\min}))\lambda_{\min}(\mathbf{D}_{\Omega_{\text{opt}}}^*\mathbf{D}_{\Omega_{\text{opt}}})} \\ &= \frac{2\sqrt{k-t}\varepsilon}{(1 - B(h_{\min}))} > 2\varepsilon. \end{aligned}$$

Thus, until $|\Omega_t^*| = k - t = 0$ (i.e. until we have identified all of the support), as desired we have

$$\|\mathbf{r}_t\|_2 > 2\varepsilon - \varepsilon = \varepsilon. \quad (3.1.29)$$

□

We again numerically display the bounds required by Theorem 3.5. For dimensions $n = 256$, $d = 1024$, noise level $\varepsilon = 10^{-3}$, and again using the same bounds on $\eta_{h_{\min},k}$ and $\eta'_{h_{\min},k}$ as described above, we empirically show several bounds with increasing sparsity levels in Figure 4.

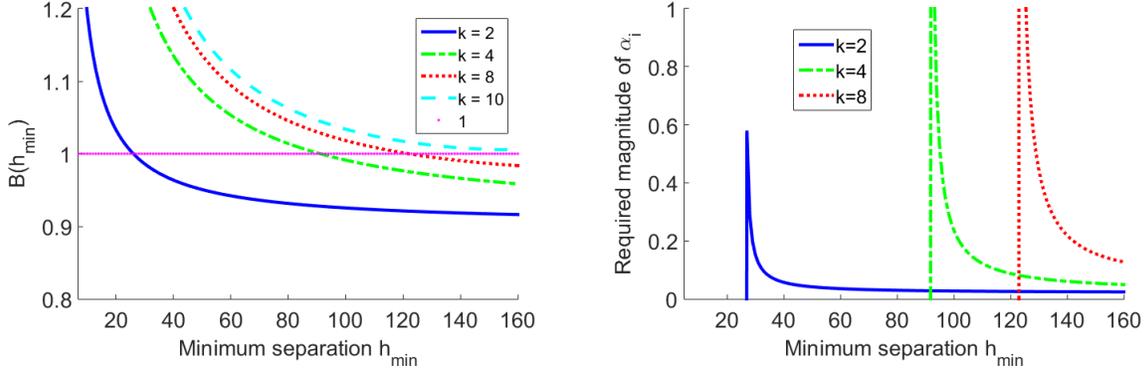


Figure 4: Dictionary is 256×1024 DFT. Left: Plot of upper bound on $B(h_{\min})$ vs minimum separation for various sparsity levels. When the curve falls below 1 (dotted line), Theorem 3.5 guarantees that OMP provides accurate recovery. Right: Lower bound on non-zeros in signal α as in Theorem 3.5 when noise is bounded by 10^{-3} as a function of minimum separation h_{\min} .

3.2 SSCoSaMP(L1)

We conclude with a simple argument showing that the ℓ_1 -minimization method can provide a set of good near-optimal projection components $\mathcal{S}_D(\mathbf{x}, k)$ even if \mathbf{D} is an overcomplete dictionary (partial DFT again for simplicity) under the condition that the signal is well-separated. We will utilize an immediate corollary of [CFG14].

Corollary 3.8. [CFG14] Consider the program

$$\underset{\tilde{\mathbf{x}}}{\operatorname{argmin}} \|\tilde{\mathbf{x}}\|_1 \text{ subject to } \mathbf{F}_n \tilde{\mathbf{x}} = \mathbf{F}_n \mathbf{x}, \quad (3.2.1)$$

where \mathbf{F}_n is the partial Fourier matrix in a low-pass band $[-f_{l_0}, f_{l_0}]$. Then if $T \subset \{0, 1, \dots, N-1\}$ is the support of $\{\mathbf{x}_t\}_{t=0}^{N-1}$ obeying

$$\min_{t, t' \in T: t \neq t'} \frac{1}{N} |t - t'| \geq 2/f_{l_0}, \quad (3.2.2)$$

the solution to (3.2.1) is exact, $\tilde{\mathbf{x}} = \mathbf{x}$.

In our context, the overcomplete dictionary \mathbf{D} can only collect the lowest $2f_{l_0} + 1$ frequencies. When the dictionary \mathbf{D} has dimension $n \times d$ with $n < d$, $\mathbf{D}\mathbf{x}$ only gives the spectrum of \mathbf{x} in the frequency range $[-\frac{n}{2}, \frac{n}{2}]$. Here then, $\frac{n}{2} = f_{l_0}$. Therefore \mathbf{D} acts as a low-pass filter in our problem and we seek to recover α by only knowing parts of the spectrum of α . We can thus translate the above result to the context of SSCoSaMP projections.

Corollary 3.9. Let \mathbf{D} be an $n \times d$ ($n \leq d$) overcomplete DFT dictionary that gives the Fourier transform of a vector in the frequency domain $[-f_{l_0}, f_{l_0}] := [-n/2, n/2]$. Suppose \mathbf{x} has a k -sparse expansion in \mathbf{D} , i.e. $\mathbf{x} = \mathbf{D}\alpha$ where α has support T with $|T| \leq k$. If the support obeys

$$\min_{t, t' \in T: t \neq t'} \frac{1}{d} |t - t'| \geq 2/f_{l_0} \quad (3.2.3)$$

then the solution to

$$\underset{\hat{\alpha}}{\operatorname{argmin}} \|\hat{\alpha}\|_1 \text{ subject to } \mathbf{D}\hat{\alpha} = \mathbf{x} \quad (3.2.4)$$

is exact ($\hat{\alpha} = \alpha$) and can thus be used as a projection within SS-CoSaMP.

This result provides theoretical backing to the behavior observed by SS-CoSaMP(ℓ_1) in Figure 1.

4 Conclusion

In this work we study the behavior of several classical CS methods as approximate projections within a greedy signal space method. It has been observed that the accuracy of the methods in this setting depends heavily on the signal structure, even though their classical recovery guarantees are independent of signal structure. Here, we analyze the behavior of these approaches when the dictionary may be highly overcomplete and thus does not satisfy typical properties like the RIP or incoherence. We prove that two of the methods, under specific assumptions on the signal structure, can be used as accurate approximate projections. Our analysis thus provides theoretical backing to explain the observed phenomena. It would be useful future work to study additional methods, like the CoSaMP method which we conjecture can serve as an approximate projection when the signal has a clustered support.

References

- [BD09] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. A.*, 27(3):265–274, 2009.
- [CENR10] E. J. Candès, Y. C. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Appl. Comput. Harmon. A.*, 31(1):59–73, 2010.
- [CFG14] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Commun. Pur. Appl. Math.*, 67(6):906–956, 2014.
- [CRT06] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pur. Appl. Math.*, 59(8):1207–1223, 2006.
- [CRVT05] E. J. Candès, M. Rudelson, R. Vershynin, and T. Tao. Error correction via linear programming. *46th Annual Symp. Found. Computer Science*, pages 668–681, 2005.
- [CT05] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE T. Inform. Theory*, 51:4203–4215, 2005.
- [CW11] T. Cai and L. Wang. Orthogonal matching pursuit for sparse signal recovery with noise. *Information Theory, IEEE Transactions*, 57:4680–4688, 2011.
- [DNW12] M. Davenport, D. Needell, and M. B. Wakin. Signal space CoSaMP for sparse recovery with redundant dictionaries. *IEEE T. Inform. Theory*, 59(10):6820, 2012.
- [DW12] M. A. Davenport and M. B. Wakin. Compressive sensing of analog signals using discrete prolate spheroidal sequences. *Appl. Comput. Harmon. A.*, 33(3):438–472, 2012.
- [Fou] S. Foucart. Dictionary-sparse recovery via thresholding-based algorithms. *J. Fourier Anal. Appl.* To appear.

- [GE13] R. Giryes and M. Elad. Omp with highly coherent dictionaries. In *10th Int. Conf. on Sampling Theory Appl.(SAMPTA)*, 2013.
- [Ger31] S. A. Gershgorin. Über die abgrenzung der eigenwerte einer matrix. *Izv. Akad. Nauk. USSR Otd. Fiz.-Mat.*, (6):749–754, 1931.
- [GGK⁺14] C. Garnatz, X. Gu, A. Kingman, J. LaManna, D. Needell, and S. Tu. Practical approximate projection schemes in greedy signal space methods. In *Proc. Allerton Conf. on Communication, Control, and Computing*, 2014.
- [GN15] R. Giryes and D. Needell. Greedy signal space methods for incoherence and beyond. *Applied and Computational Harmonic Analysis*, 39(1):1–20, 2015.
- [GNE⁺14] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. E. Davies. Greedy-like algorithms for the cospase analysis model. *Linear Algebra Appl.*, 441:22–60, 2014.
- [Mut05] S. Muthukrishnan. Data streams: Algorithms and applications. 2005.
- [NDEG13] S. Nam, M. E. Davies, M. Elad, and R. Gribonval. The cospase analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013.
- [NT09] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comput. Harmon. A.*, 26(3):301–321, 2009.
- [NV07] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. *IEEE J. Sel. Top. Signa.*, 4:310–316, 2007.
- [RRT12] H. Rauhut, J. Romberg, and J. A. Tropp. Restricted isometries for partial random circulant matrices. *Appl. Comput. Harmon. A.*, 32(2):242–254, 2012.
- [RV08] M. Rudelson and R. Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Comm. Pure Appl. Math.*, 61:1025–1045, 2008.
- [TG07] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via Orthogonal Matching Pursuit. *IEEE T. Inform. Theory*, 53(12):4655–4666, 2007.
- [Tro04] J. A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE T. Inform. Theory*, 50(10):2231–2242, Oct. 2004.
- [Zha11] T. Zhang. Sparse recovery with orthogonal matching pursuit under RIP. *IEEE T. Inform. Theory*, 57(9):6215–6221, 2011.