# Just Relax: Convex Programming Methods for Identifying Sparse Signals in Noise

Joel A. Tropp, *Member, IEEE*

***Abstract*— This article studies a difficult and fundamental problem that arises throughout electrical engineering, applied mathematics, and statistics. Suppose that one forms a short linear combination of elementary signals drawn from a large, fixed collection. Given an observation of the linear combination that has been contaminated with additive noise, the goal is to identify which elementary signals participated and to approximate their coefficients. Although many algorithms have been proposed, there is little theory which guarantees that these algorithms can accurately and efficiently solve the problem.**

**This article studies a method called convex relaxation, which attempts to recover the ideal sparse signal by solving a convex program. This approach is powerful because the optimization can be completed in polynomial time with standard scientific software. The paper provides general conditions which ensure that convex relaxation succeeds. As evidence of the broad impact of these results, the article describes how convex relaxation can be used for several concrete signal recovery problems. It also describes applications to channel coding, linear regression, and numerical analysis.**

***Index Terms*— Algorithms, approximation methods, Basis Pursuit, convex program, optimization methods, linear regression, Orthogonal Matching Pursuit, sparse representations**

## I. Introduction

LATELY, there has been a lot of fuss about sparse approximation. This class of problems has two defining characteristics:

1) An input signal is approximated by a linear combination of elementary signals. In many modern applications, the elementary signals are drawn from a large, linearly dependent collection.
2) A preference for "sparse" linear combinations is imposed by penalizing nonzero coefficients. The most common penalty is the number of elementary signals that participate in the approximation.

Sparse approximation problems arise throughout electrical engineering, statistics, and applied mathematics. One of the most common applications is to compress audio [1], images [2], and video [3]. Sparsity criteria also arise in linear regression [4], deconvolution [5], signal modeling [6], preconditioning [7], machine learning [8], de-noising [9], and regularization [10].

Date: 14 February 2004. Revised 6 February 2005 and 21 November 2005.

J. A. Tropp was with the Inst. for Computational Engineering and Sciences (ICES), The University of Texas at Austin, Austin, TX 78712 USA. He is currently with the Dept. of Mathematics, The University of Michigan at Ann Arbor, Ann Arbor, MI 48109 USA. E-mail: jtropp@umich.edu. The research for this article was completed under the support of an NSF Graduate Fellowship.

### A. The Model Problem

In this work, we will concentrate on the model problem of identifying a sparse linear combination of elementary signals that has been contaminated with additive noise. The literature on inverse problems tends to assume that the noise is an arbitrary vector of bounded norm, while the signal processing literature usually models the noise statistically; we will consider both possibilities.

To be precise, suppose we measure a signal of the form

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} + \boldsymbol{\nu} \tag{1}$$

where $\boldsymbol{\Phi}$ is a known matrix with unit-norm columns, $\boldsymbol{c}_{\mathrm{opt}}$ is a sparse coefficient vector (i.e., few components are nonzero), and $\boldsymbol{\nu}$ is an unknown noise vector. Given the signal $\boldsymbol{s}$, our goal is to approximate the coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$. In particular, it is essential that we correctly identify the nonzero components of the coefficient vector because they determine which columns of the matrix $\boldsymbol{\Phi}$ participate in the signal.

Initially, linear algebra seems to preclude a solution—whenever $\boldsymbol{\Phi}$ has a nontrivial null space, we face an ill-posed inverse problem. Even worse, the sparsity of the coefficient vector introduces a combinatorial aspect to the problem. Nevertheless, if the optimal coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$ is sufficiently sparse, it turns out that we can accurately and efficiently approximate $\boldsymbol{c}_{\mathrm{opt}}$ given the noisy observation $\boldsymbol{s}$.

### B. Convex Relaxation

The literature contains many types of algorithms for approaching the model problem, including brute force [4, Sec. 3.7–3.9], nonlinear programming [11], and greedy pursuit [12–14]. In this paper, we concentrate on a powerful method called convex relaxation. Although this technique was introduced over thirty years ago in [15], the theoretical justifications are still shaky. This paper attempts to lay a more solid foundation.

Let us explain the intuition behind convex relaxation methods. Suppose we are given a signal $\boldsymbol{s}$ of the form (1) along with a bound on the $\ell_2$ norm of the noise vector, say $\|\boldsymbol{\nu}\|_2 \leq \varepsilon$. At first, it is tempting to look for the sparsest coefficient vector that generates a signal within distance $\varepsilon$ of the input. This idea can be phrased as a mathematical program:

$$\min_{\boldsymbol{c}} \; \|\boldsymbol{c}\|_0 \quad \text{subject to} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2 \leq \varepsilon \tag{2}$$

where the $\ell_0$ quasi-norm $\|\cdot\|_0$ counts the number of nonzero components in its argument. To solve (2) directly, one must sift

through all possible disbursements of the nonzero components in $\boldsymbol{c}$. This method is intractable because the search space is exponentially large [12, 16].

To surmount this obstacle, one might replace the $\ell_0$ quasi-norm with the $\ell_1$ norm to obtain a convex optimization problem.

$$\min_{\boldsymbol{b}} \ \|\boldsymbol{b}\|_1 \quad \text{subject to} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2 \leq \delta \qquad (\ell_1\text{-ERROR})$$

where the tolerance $\delta$ is related to the error bound $\varepsilon$. Intuitively, the $\ell_1$ norm is the convex function closest to the $\ell_0$ quasi-norm, so this substitution is referred to as *convex relaxation*. One hopes that the solution to the relaxation will yield a good approximation of the ideal coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$. The advantage of the new formulation is that it can be solved in polynomial time with standard scientific software [17].

We have found that it is more natural to study the closely related convex program,

$$\min_{\boldsymbol{b}} \ \tfrac{1}{2}\ \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \ + \ \gamma\ \|\boldsymbol{b}\|_1 \qquad (\ell_1\text{-PENALTY})$$

As before, one can interpret the $\ell_1$ norm as a convex relaxation of the $\ell_0$ quasi-norm. So the parameter $\gamma$ manages a tradeoff between the approximation error and the sparsity of the coefficient vector. This optimization problem can also be solved efficiently, and one expects that the minimizer will approximate the ideal coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$.

Appendix I offers an excursus on the history of convex relaxation methods for identifying sparse signals.

### C. Contributions

The primary contribution of this paper is to justify the intuition that convex relaxation can solve the model problem. Our two major theorems describe the behavior of solutions to ($\ell_1$-PENALTY) and solutions to ($\ell_1$-ERROR). We apply this theory to several concrete instances of the model problem.

Let us summarize our results on the performance of convex relaxation. Suppose that we measure a signal of the form (1) and that we solve one of the convex programs with an appropriate choice of $\gamma$ or $\delta$ to obtain a minimizer $\boldsymbol{b}_\star$. Theorems 8 and 14 demonstrate that

- the vector $\boldsymbol{b}_\star$ is close to the ideal coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$.
- The support of $\boldsymbol{b}_\star$ (i.e., the indices of its nonzero components) is a subset of the support of $\boldsymbol{c}_{\mathrm{opt}}$.
- Moreover, the minimizer is unique.

In words, the solution to the convex relaxation identifies every sufficiently large component of $\boldsymbol{c}_{\mathrm{opt}}$, and it never mistakenly identifies a column of $\boldsymbol{\Phi}$ that did not participate in the signal. As other authors have written, "It seems quite surprising that any result of this kind is possible" [18, Section 6].

The conditions under which the convex relaxations solve the model problem are geometric in nature. Section III-C describes the primary factors influencing their success:

- Small sets of the columns from $\boldsymbol{\Phi}$ should be well-conditioned.
- The noise vector should be weakly correlated with all the columns of $\boldsymbol{\Phi}$.

These properties can be difficult to check in general. This paper relies on a simple approach based on the *coherence parameter* of $\boldsymbol{\Phi}$, which measures the cosine of the minimal angle between a pair of columns. It may be possible to improve these results using techniques from Banach space geometry.

As an application of the theory, we will see that convex relaxation can be used to solve three versions of the model problem. If the coherence parameter is small,

- both convex programs can identify a sufficiently sparse signal corrupted by an arbitrary vector of bounded $\ell_2$ norm (Sections IV-C and V-B).
- The program ($\ell_1$-PENALTY) can identify a sparse signal in additive white Gaussian noise (Section IV-D).
- The program ($\ell_1$-ERROR) can identify a sparse signal in uniform noise of bounded $\ell_2$ norm (Section V-C).

In addition, Section IV-E shows that ($\ell_1$-PENALTY) can be used for subset selection, a sparse approximation problem that arises in statistics. Section V-D demonstrates that ($\ell_1$-ERROR) can solve another sparse approximation problem from numerical analysis.

This paper is among the first to analyze the behavior of convex relaxation when noise is present. Prior theoretical work has focused on a special case of the model problem where the noise vector $\boldsymbol{\nu} = \boldsymbol{0}$. See for example [14, 19–23]. Although these results are beautiful, they simply do not apply to real-world signal processing problems. We expect that our work will have a significant practical impact because so many applications require sparse approximation in the presence of noise. As an example, Dossal and Mallat have used our results to study the performance of convex relaxation for a problem in seismics [24].

After the first draft [25] of this paper had been completed, it came to the author's attention that several other researchers were preparing papers on the performance of convex relaxation in the presence of noise [18, 26, 27]. These manuscripts are significantly different from the present work, and they also deserve attention. Turn to Section VI-B for comparisons and contrasts.

### D. Channel Coding

It may be illuminating to view the model problem in the context of channel coding. The Gaussian channel allows us to send one real number during each transmission, but this number is corrupted by the addition of a zero-mean Gaussian random variable. Shannon's channel coding theorem [28, Chap. 10] shows that the capacity of the Gaussian channel can be achieved (asymptotically) by grouping the transmissions into long blocks and using a random code whose size is exponential in the block length. A major practical problem with this approach is that decoding requires an exponentially large lookup table.

In contrast, we could also construct a large code by forming sparse linear combinations of vectors from a fixed codebook $\boldsymbol{\Phi}$. Both the choice of vectors and the nonzero coefficients carry information. The vector length is analogous with the block length. When we transmit a codeword through the channel, it is contaminated by a zero-mean, white Gaussian

random vector. To identify the codeword, it is necessary to solve the model problem. Therefore, any sparse approximation algorithm—such as convex relaxation—can potentially be used for decoding.

To see that this coding scheme is practical, one must show that the codewords carry almost as much information as the channel permits. One must also demonstrate that the decoding algorithm can reliably recover the noisy codewords.

We believe that channel coding is a novel application of sparse approximation algorithms. In certain settings, these algorithms may provide a very competitive approach. Sections IV-D and V-C present some simple examples that take a first step in this direction. See [29] for some more discussion of these ideas.

### E. Outline

The paper continues with notation and background material in Section II. Section III develops the fundamental lemmata that undergird our major results. The two subsequent sections provide the major theoretical results for ($\ell_1$-PENALTY) and ($\ell_1$-ERROR). Both these sections describe several specific applications in signal recovery, approximation theory, statistics, and numerical analysis. The body of the paper concludes with Section VI, which surveys extensions and related work. The appendices present several additional topics and some supporting analysis.

## II. Background

This section furnishes the mathematical *mise en scène*. We have adopted an abstract point of view that facilitates our treatment of convex relaxation methods. Most of the material here may be found in any book on matrix theory or functional analysis, such as [30–32].

### A. The Dictionary

We will work in the finite-dimensional, complex inner-product space $\mathbb{C}^d$, which will be called the *signal space*[1]. The usual Hermitian inner product for $\mathbb{C}^d$ will be written as $\langle \cdot, \cdot \rangle$, and we will denote the corresponding norm by $\|\cdot\|_2$. A *dictionary* for the signal space is a finite collection $\mathscr{D}$ of unit-norm elementary signals. The elementary signals are called *atoms*, and each is denoted by $\boldsymbol{\varphi}_\omega$, where the parameter $\omega$ is drawn from an index set $\Omega$. The indices may have an interpretation, such as the time–frequency or time–scale localization of an atom, or they may simply be labels without an underlying metaphysics. The whole dictionary structure is written as

$$\mathscr{D} \;=\; \{\boldsymbol{\varphi}_\omega \;:\; \omega \in \Omega\}$$

The letter $N$ will denote the number of atoms in the dictionary. It is evident that $N = |\mathscr{D}| = |\Omega|$, where $|\cdot|$ returns the cardinality of a finite set.

[1]Modifications for real signal spaces should be transparent, but the apotheosis to infinite dimensions may take additional effort.

### B. The Coherence Parameter

A summary parameter of the dictionary is the *coherence*, which is defined as the maximum absolute inner product between two distinct atoms [12, 19].

$$\mu \;\stackrel{\text{def}}{=}\; \max_{\lambda \neq \omega} \left| \langle \boldsymbol{\varphi}_\omega,\; \boldsymbol{\varphi}_\lambda \rangle \right|. \tag{3}$$

When the coherence is small, the atoms look very different from each other, which makes them easy to distinguish. It is common that the coherence satisfies $\mu = O(d^{-1/2})$. We say informally that a dictionary is *incoherent* when we judge that the coherence parameter is small. An incoherent dictionary may contain far more atoms than an orthonormal basis (i.e., $N \gg d$) [14, Sec. II-D].

The literature also contains a generalization of the coherence parameter called the *cumulative coherence function* [14, 21]. For each natural number $m$, this function is defined as

$$\mu_1(m) \;\stackrel{\text{def}}{=}\; \max_{|\Lambda| \leq m} \; \max_{\omega \notin \Lambda} \sum\nolimits_{\lambda \in \Lambda} \left| \langle \boldsymbol{\varphi}_\omega,\; \boldsymbol{\varphi}_\lambda \rangle \right|.$$

This function will often provide better estimates than the coherence parameter. For clarity of exposition, we only present results in terms of the coherence parameter. Parallel results using cumulative coherence are not hard to develop [25].

It is essential to be aware that coherence is not fundamental to sparse approximation. Rather, it offers a quick way to check the hypotheses of our theorems. It is possible to provide much stronger results using more sophisticated techniques.

### C. Coefficient Space

Every linear combination of atoms is parameterized by a list of coefficients. We may collect them into a *coefficient vector*, which formally belongs to the linear space[2] $\mathbb{C}^\Omega$. The canonical basis for this space is given by the vectors whose coordinate projections are identically zero, except for one unit coordinate. The $\omega$-th canonical basis vector will be denoted $\mathbf{e}_\omega$.

Given a coefficient vector $\boldsymbol{c}$, the expressions $c_\omega$ and $\boldsymbol{c}(\omega)$ both represent its $\omega$-th coordinate. We will alternate between them, depending on what is most typographically felicitous.

The *support* of a coefficient vector is the set of indices at which it is nonzero:

$$\operatorname{supp}(\boldsymbol{c}) \;\stackrel{\text{def}}{=}\; \{\omega \in \Omega \;:\; c_\omega \neq 0\}. \tag{4}$$

Suppose that $\Lambda \subset \Omega$. Without notice, we may embed "short" coefficient vectors from $\mathbb{C}^\Lambda$ into $\mathbb{C}^\Omega$ by extending them with zeros. Likewise, we may restrict "long" coefficient vectors from $\mathbb{C}^\Omega$ to their support. Both transubstantiations will be natural in context.

### D. Sparsity and Diversity

The *sparsity* of a coefficient vector is the number of places where it equals zero. The complementary notion, *diversity*, counts the number of places where the coefficient vector does not equal zero. Diversity is calculated with the $\ell_0$ *quasi-norm*:

$$\|\boldsymbol{c}\|_0 \;\stackrel{\text{def}}{=}\; |\operatorname{supp}(\boldsymbol{c})|. \tag{5}$$

[2]In case this notation is unfamiliar, $\mathbb{C}^\Omega$ is the collection of functions from $\Omega$ to $\mathbb{C}$. It is equipped with the usual addition and scalar multiplication to form a linear space.

For any positive number $p$, define

$$\|\boldsymbol{c}\|_p \stackrel{\text{def}}{=} \left[\sum\nolimits_{\omega\in\Omega} |c_\omega|^p\right]^{1/p} \tag{6}$$

with the convention that $\|\boldsymbol{c}\|_\infty \stackrel{\text{def}}{=} \max_{\omega\in\Omega} |c_\omega|$. As one might expect, there is an intimate connection between the definitions (5) and (6). Indeed, $\|\boldsymbol{c}\|_0 = \lim_{p\to 0} \|\boldsymbol{c}\|_p^p$. It is well known that the function (6) is convex if and only if $p \geq 1$, in which case it describes the $\ell_p$ norm.

From this vantage, one can see that the $\ell_1$ norm is the convex function "closest" to the $\ell_0$ quasi-norm (subject to the same normalization on the zero vector and the canonical basis vectors). For a more rigorous approach to this idea, see Proposition 2.1 of [33] and also [22].

### E. Dictionary Matrices

Let us define a matrix $\boldsymbol{\Phi}$, called the *dictionary synthesis matrix*, that maps coefficient vectors to signals. Formally,

$$\boldsymbol{\Phi} : \mathbb{C}^\Omega \longrightarrow \mathbb{C}^d \quad \text{via} \quad \boldsymbol{\Phi} : \boldsymbol{c} \longmapsto \sum\nolimits_{\omega\in\Omega} c_\omega\, \boldsymbol{\varphi}_\omega.$$

The matrix $\boldsymbol{\Phi}$ describes the action of this linear transformation in the canonical bases of the underlying vector spaces. Therefore, the columns of $\boldsymbol{\Phi}$ are the atoms. The conjugate transpose of $\boldsymbol{\Phi}$ is called the *dictionary analysis matrix*, and it maps each signal $\boldsymbol{s}$ to a coefficient vector that lists the inner products between signal and atoms.

$$\boldsymbol{\Phi}^* : \mathbb{C}^d \longrightarrow \mathbb{C}^\Omega \quad \text{via} \quad (\boldsymbol{\Phi}^*\, \boldsymbol{s})(\omega) = \langle \boldsymbol{s},\ \boldsymbol{\varphi}_\omega \rangle\,.$$

The rows of $\boldsymbol{\Phi}^*$ are atoms, conjugate-transposed.

### F. Sub-dictionaries

A *sub-dictionary* is a linearly independent collection of atoms. We will exploit the linear independence repeatedly. If the atoms in a sub-dictionary are indexed by the set $\Lambda$, then we define a synthesis matrix $\boldsymbol{\Phi}_\Lambda : \mathbb{C}^\Lambda \to \mathbb{C}^d$ and an analysis matrix $\boldsymbol{\Phi}_\Lambda^* : \mathbb{C}^d \to \mathbb{C}^\Lambda$. These matrices are entirely analogous with the dictionary synthesis and analysis matrices. We will frequently use the fact that the synthesis matrix $\boldsymbol{\Phi}_\Lambda$ has full column rank.

Let $\Lambda$ index a sub-dictionary. The *Gram matrix* of the sub-dictionary is given by $\boldsymbol{\Phi}_\Lambda^*\, \boldsymbol{\Phi}_\Lambda$. This matrix is Hermitian, it has a unit diagonal, and it is invertible. The *pseudo-inverse* of the synthesis matrix is denoted by $\boldsymbol{\Phi}_\Lambda^\dagger$, and it may be calculated using the formula $\boldsymbol{\Phi}_\Lambda^\dagger = (\boldsymbol{\Phi}_\Lambda^*\boldsymbol{\Phi}_\Lambda)^{-1}\, \boldsymbol{\Phi}_\Lambda^*$. The matrix $\boldsymbol{P}_\Lambda$ will denote the orthogonal projector onto the span of the sub-dictionary. This projector may be expressed using the pseudo-inverse: $\boldsymbol{P}_\Lambda = \boldsymbol{\Phi}_\Lambda \boldsymbol{\Phi}_\Lambda^\dagger$.

On occasion, the distinguished index set is $\Lambda_{\text{opt}}$ instead of $\Lambda$. In this case, the synthesis matrix is written as $\boldsymbol{\Phi}_{\text{opt}}$, and the orthogonal projector is denoted $\boldsymbol{P}_{\text{opt}}$.

### G. Signals, Approximations, and Coefficients

Let $\Lambda$ be a sub-dictionary, and let $\boldsymbol{s}$ be a fixed input signal. This signal has a unique $\ell_2$ best approximation $\boldsymbol{a}_\Lambda$ using the atoms in $\Lambda$, which is determined by $\boldsymbol{a}_\Lambda = \boldsymbol{P}_\Lambda\, \boldsymbol{s}$. Note that the approximation $\boldsymbol{a}_\Lambda$ is orthogonal to the residual $(\boldsymbol{s} - \boldsymbol{a}_\Lambda)$. There is a unique coefficient vector $\boldsymbol{c}_\Lambda$ supported on $\Lambda$ that synthesizes the approximation: $\boldsymbol{a}_\Lambda = \boldsymbol{\Phi}_\Lambda\, \boldsymbol{c}_\Lambda$. We may calculate that $\boldsymbol{c}_\Lambda = \boldsymbol{\Phi}_\Lambda^\dagger\, \boldsymbol{s} = \boldsymbol{\Phi}_\Lambda^\dagger\, \boldsymbol{a}_\Lambda$.

### H. Operator Norms

The $(p, q)$ *operator norm* of a matrix $\boldsymbol{A}$ is given by

$$\|\boldsymbol{A}\|_{p,q} \stackrel{\text{def}}{=} \max_{\boldsymbol{x}\neq\boldsymbol{0}} \frac{\|\boldsymbol{A}\,\boldsymbol{x}\|_q}{\|\boldsymbol{x}\|_p}.$$

An immediate consequence is the upper norm bound

$$\|\boldsymbol{A}\,\boldsymbol{x}\|_q \ \leq\ \|\boldsymbol{A}\|_{p,q}\, \|\boldsymbol{x}\|_p\,.$$

Suppose that $1/p + 1/p' = 1$ and $1/q + 1/q' = 1$. Then we have the identity

$$\|\boldsymbol{A}^*\|_{q',p'} \ =\ \|\boldsymbol{A}\|_{p,q}\,. \tag{7}$$

We also have the following lower norm bound, which is proved in [25, Section 3.3].

*Proposition 1:* For every matrix $\boldsymbol{A}$,

$$\min_{\substack{\boldsymbol{x}\in\mathscr{R}(\boldsymbol{A}^*)\\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\|\boldsymbol{A}\,\boldsymbol{x}\|_q}{\|\boldsymbol{x}\|_p} \quad\geq\quad \left\|\boldsymbol{A}^\dagger\right\|_{q,p}^{-1}. \tag{8}$$

If $\boldsymbol{A}$ has full row rank, equality holds in (8). When $\boldsymbol{A}$ is invertible, this result implies

$$\min_{\boldsymbol{x}\neq\boldsymbol{0}} \frac{\|\boldsymbol{A}\,\boldsymbol{x}\|_q}{\|\boldsymbol{x}\|_p} \quad=\quad \left\|\boldsymbol{A}^{-1}\right\|_{q,p}^{-1}.$$

The symbol $\mathscr{R}(\cdot)$ denotes the range (i.e., column span) of a matrix.

## III. Fundamental Lemmata

Fix an input signal $\boldsymbol{s}$ and a positive parameter $\gamma$. Define the convex function

$$L(\boldsymbol{b}) \ =\ \tfrac{1}{2}\ \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \ +\ \gamma\ \|\boldsymbol{b}\|_1\,. \tag{L}$$

In this section, we study the minimizers of this function. These results are important to us because (L) is the objective function of the convex program ($\ell_1$-PENALTY), and it is essentially the Lagrangian function of ($\ell_1$-ERROR). By analyzing the behavior of (L), we can understand the performance of convex relaxation methods.

The major lemma of this section provides a sufficient condition which ensures that the minimizer of (L) is supported on a given index set $\Lambda$. To develop this result, we first characterize the (unique) minimizer of (L) when it is restricted to coefficient vectors supported on $\Lambda$. We use this characterization to obtain a condition under which every perturbation away from the restricted minimizer must increase the value of the objective function. When this condition is in force, the global minimizer of (L) must be supported on $\Lambda$.

### A. Convex Analysis

The proof relies on standard results from convex analysis. As it is usually presented, this subject addresses the properties of real-valued convex functions defined on real vector spaces. Nevertheless, it is possible to transport these results to the complex setting by defining an appropriate real-linear structure on the complex vector space. In this section, therefore, we use the bilinear inner product $\operatorname{Re}\langle \boldsymbol{x},\ \boldsymbol{y}\rangle$ instead of the usual sesquilinear (i.e., Hermitian) inner product. Note that both inner products generate the $\ell_2$ norm.

Suppose that $f$ is a convex function from a real-linear inner-product space $\mathbb{X}$ to $\mathbb{R}$. The *gradient* $\nabla f(\boldsymbol{x})$ of the function $f$ at the point $\boldsymbol{x}$ is defined as the usual (Fréchet) derivative, computed with respect to the real-linear inner product. Even when the gradient does not exist, we may define the *subdifferential* of $f$ at a point $\boldsymbol{x}$:

$$\partial f(\boldsymbol{x}) \stackrel{\text{def}}{=} \{\boldsymbol{g} \in \mathbb{X} : f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \operatorname{Re}\langle \boldsymbol{y}-\boldsymbol{x},\ \boldsymbol{g}\rangle \quad \text{for every } \boldsymbol{y} \in \mathbb{X}\}.$$

The elements of the subdifferential are called *subgradients*. If $f$ possesses a gradient at $\boldsymbol{x}$, the unique subgradient is the gradient. That is,

$$\partial f(\boldsymbol{x}) \ = \ \{\nabla f(\boldsymbol{x})\}.$$

The subdifferential of a sum is the (Minkowski) sum of the subdifferentials. Finally, if $f$ is a closed, proper convex function then $\boldsymbol{x}$ is a global minimizer of $f$ if and only if $\mathbf{0} \in \partial f(\boldsymbol{x})$. The standard reference for this material is [34].

*Remark 2:* The subdifferential of a convex function provides a dual description of the function in terms of its supporting hyperplanes. In consequence, the appearance of the subdifferential in our proof is analogous with the familiar technique of studying the dual of a convex program.

### B. Restricted Minimizers

First, we must characterize the minimizers of the objective function (L). Fuchs developed the following result in the real setting using essentially the same method [23].

*Lemma 3:* Suppose that $\Lambda$ indexes a linearly independent collection of atoms, and let $\boldsymbol{b}_\star$ minimize the objective function (L) over all coefficient vectors supported on $\Lambda$. A necessary and sufficient condition on such a minimizer is that

$$\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star \ = \ \gamma\,(\boldsymbol{\Phi}_\Lambda^*\boldsymbol{\Phi}_\Lambda)^{-1}\,\boldsymbol{g} \tag{9}$$

where the vector $\boldsymbol{g}$ is drawn from $\partial\,\|\boldsymbol{b}_\star\|_1$. Moreover, the minimizer $\boldsymbol{b}_\star$ is unique.

*Proof:* Suppose that $\operatorname{supp}(\boldsymbol{b}) \subset \Lambda$. Then the vectors $(\boldsymbol{s}-\boldsymbol{a}_\Lambda)$ and $(\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}\,\boldsymbol{b})$ are orthogonal to each other. Apply the Pythagorean Theorem to see that minimizing (L) over coefficient vectors supported on $\Lambda$ is equivalent to minimizing the function

$$F(\boldsymbol{b}) \ \stackrel{\text{def}}{=} \ \tfrac{1}{2}\,\|\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}_\Lambda\,\boldsymbol{b}\|_2^2 \ + \ \gamma\,\|\boldsymbol{b}\|_1 \tag{10}$$

over coefficient vectors from $\mathbb{C}^\Lambda$. Recall that $\boldsymbol{\Phi}_\Lambda$ has full column rank. It follows that the quadratic term in (10) is strictly convex, and so the whole function $F$ must also be strictly convex. Therefore, its minimizer is unique.

The function $F$ is convex and unconstrained, so $\mathbf{0} \in \partial F(\boldsymbol{b}_\star)$ is a necessary and sufficient condition for the coefficient vector $\boldsymbol{b}_\star$ to minimize $F$. The gradient of the first term of $F$ at $\boldsymbol{b}_\star$ equals $(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)\,\boldsymbol{b}_\star - \boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{a}_\Lambda$. From the additivity of subdifferentials, it follows that

$$(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)\,\boldsymbol{b}_\star \ - \ \boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{a}_\Lambda \ + \ \gamma\,\boldsymbol{g} \ = \ \mathbf{0}$$

for some vector $\boldsymbol{g}$ drawn from the subdifferential $\partial\,\|\boldsymbol{b}_\star\|_1$. We pre-multiply this relation by $(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}$ to reach

$$\boldsymbol{\Phi}_\Lambda^\dagger\boldsymbol{a}_\Lambda - \boldsymbol{b}_\star \ = \ \gamma\,(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\,\boldsymbol{g}.$$

Apply the fact that $\boldsymbol{c}_\Lambda = \boldsymbol{\Phi}_\Lambda^\dagger\boldsymbol{a}_\Lambda$ to reach the conclusion. ∎

Now we identify the subdifferential of the $\ell_1$ norm. To that end, define the signum function as

$$\operatorname{sgn}(r\,\mathrm{e}^{\mathrm{i}\,\theta}) \ \stackrel{\text{def}}{=} \ \begin{cases} \mathrm{e}^{\mathrm{i}\,\theta} & \text{for } r > 0 \\ 0 & \text{for } r = 0. \end{cases}$$

One may extend the signum function to vectors by applying it to each component.

*Proposition 4:* Let $\boldsymbol{z}$ be a complex vector. The complex vector $\boldsymbol{g}$ lies in the subdifferential $\partial\,\|\boldsymbol{z}\|_1$ if and only if

- $|g_k| \leq 1$ whenever $z_k = 0$, and
- $g_k = \operatorname{sgn} z_k$ whenever $z_k \neq 0$.

Indeed, $\|\boldsymbol{g}\|_\infty = 1$ unless $\boldsymbol{z} = \mathbf{0}$, in which case $\|\boldsymbol{g}\|_\infty \leq 1$.

We omit the proof. At last, we may develop bounds on how much a solution to the restricted problem varies from the desired solution $\boldsymbol{c}_\Lambda$.

*Corollary 5 (Upper Bounds):* Suppose that $\Lambda$ indexes a sub-dictionary, and let $\boldsymbol{b}_\star$ minimize the function (L) over all coefficient vectors supported on $\Lambda$. The following bounds are in force:

$$\|\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star\|_\infty \ \leq \ \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty} \tag{11}$$

$$\|\boldsymbol{\Phi}_\Lambda\,(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star)\|_2 \ \leq \ \gamma\,\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,1}. \tag{12}$$

*Proof:* We begin with the necessary and sufficient condition

$$\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star \ = \ \gamma\,(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\,\boldsymbol{g} \tag{13}$$

where $\boldsymbol{g} \in \partial\,\|\boldsymbol{b}_\star\|_1$. To obtain (11), we take the $\ell_\infty$ norm of (13) and apply the upper norm bound.

$$\begin{aligned}\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_\infty \ &= \ \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\,\boldsymbol{g}\right\|_\infty \\ &\leq \ \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty}\,\|\boldsymbol{g}\|_\infty.\end{aligned}$$

Proposition 4 shows that $\|\boldsymbol{g}\|_\infty \leq 1$, which proves the result.

To develop the second bound (12), we pre-multiply (13) by the matrix $\boldsymbol{\Phi}_\Lambda$ and compute the $\ell_2$ norm:

$$\begin{aligned}\|\boldsymbol{\Phi}_\Lambda\,(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star)\|_2 \ &= \ \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\boldsymbol{g}\right\|_2 \\ &\leq \ \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^\dagger)^*\right\|_{\infty,2}\,\|\boldsymbol{g}\|_\infty.\end{aligned}$$

As before, $\|\boldsymbol{g}\|_\infty \leq 1$. Finally, we apply the identity (7) to switch from the $(\infty, 2)$ norm to the $(2, 1)$ norm. ∎

### C. The Correlation Condition

Now, we develop a condition which ensures that the global minimizer of (L) is supported inside a given index set. This result is the soul of the analysis.

*Lemma 6 (Correlation Condition):* Assume that $\Lambda$ indexes a linearly independent collection of atoms. Let $\boldsymbol{b}_\star$ minimize the function (L) over all coefficient vectors supported on $\Lambda$. Suppose that

$$\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma \left[1 \;-\; \max\nolimits_{\omega \notin \Lambda} \left|\langle \boldsymbol{\Phi}_\Lambda^\dagger \boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\rangle\right|\right]$$

where $\boldsymbol{g} \in \partial \|\boldsymbol{b}_\star\|_1$ is determined by (9). It follows that $\boldsymbol{b}_\star$ is the unique global minimizer of (L). In particular, the condition

$$\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma \left[1 \;-\; \max\nolimits_{\omega \notin \Lambda} \left\|\boldsymbol{\Phi}_\Lambda^\dagger \boldsymbol{\varphi}_\omega\right\|_1\right]$$

guarantees that $\boldsymbol{b}_\star$ is the unique global minimizer of (L).

In this work, we will concentrate on the second sufficient condition because it is easier to work with. Nevertheless, the first condition is significantly more powerful. Note that either sufficient condition becomes worthless when the bracket on its right-hand side is negative.

We typically abbreviate the bracket in the second condition:

$$\mathrm{ERC}(\Lambda) \;\stackrel{\text{def}}{=}\; 1 \;-\; \max\nolimits_{\omega \notin \Lambda} \left\|\boldsymbol{\Phi}_\Lambda^\dagger \boldsymbol{\varphi}_\omega\right\|_1.$$

The notation "ERC" stands for *Exact Recovery Coefficient*, which reflects the fact that $\mathrm{ERC}(\Lambda) > 0$ is a sufficient condition for several different algorithms to recover the optimal representation of an exactly sparse signal [14]. Roughly, $\mathrm{ERC}(\Lambda)$ measures how distinct the atoms in $\Lambda$ are from the remaining atoms. Observe that if the index set satisfies $\mathrm{ERC}(\Lambda) > 0$, then the sufficient condition of the lemma will hold as soon as $\gamma$ becomes large enough. But if $\gamma$ is too large, then $\boldsymbol{b}_\star = \boldsymbol{0}$.

Given a nonzero signal $\boldsymbol{v}$, define the function

$$\mathrm{maxcor}(\boldsymbol{v}) \;\stackrel{\text{def}}{=}\; \frac{\max_{\omega \in \Omega} \;|\langle \boldsymbol{v},\; \boldsymbol{\varphi}_\omega\rangle|}{\|\boldsymbol{v}\|_2},$$

and place the convention that $\mathrm{maxcor}(\boldsymbol{0}) = 0$. This quantity measures the maximum correlation between the signal $\boldsymbol{v}$ and any atom in the dictionary. To interpret the left-hand side of the sufficient conditions in the lemma, observe that

$$\begin{aligned}\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;&=\; \max_{\omega \in \Omega} \;|\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| \\ &=\; \mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\; \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2 \,.\end{aligned}$$

Therefore, the lemma is strongest when the magnitude of the residual and its *maximum correlation* with the dictionary are both small. If the dictionary is not exponentially large, a generic vector is weakly correlated with the dictionary on account of measure concentration phenomena. Since the atoms are normalized, the maximum correlation never exceeds one. This fact yields a (much weaker) result that depends only on the magnitude of the residual.

*Corollary 7:* Let $\Lambda$ index a sub-dictionary, and let $\boldsymbol{s}$ be an input signal. Suppose that the residual vector satisfies

$$\|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2 \quad\leq\quad \gamma\, \mathrm{ERC}(\Lambda).$$

Then any coefficient vector $\boldsymbol{b}_\star$ that minimizes the function (L) must be supported inside $\Lambda$.

One might wonder whether the condition $\mathrm{ERC}(\Lambda) \geq 0$ is really necessary to develop results of this type. The answer is a qualified affirmative. Appendix II offers a partial converse of Lemma 6.

### D. Proof of the Correlation Condition Lemma

We now establish Lemma 6. The argument presented here is different from the original proof in the technical report [25]. It relies on a perturbation technique that can be traced to the independent works [18, 25].

*Proof:* Let $\boldsymbol{b}_\star$ be the unique minimizer of the objective function (L) over coefficient vectors supported on $\Lambda$. In particular, the value of the objective function increases if we change any coordinate of $\boldsymbol{b}_\star$ indexed in $\Lambda$. We will develop a condition which guarantees that the value of the objective function also increases when we change any other component of $\boldsymbol{b}_\star$. Since (L) is convex, these two facts will imply that $\boldsymbol{b}_\star$ is the global minimizer.

Choose an index $\omega$ not listed in $\Lambda$, and let $\delta$ be a nonzero scalar. We must develop a condition which ensures that

$$L(\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;>\; 0$$

where $\mathbf{e}_\omega$ is the $\omega$-th canonical basis vector. To that end, expand the left-hand side of this relation to obtain

$$\begin{aligned}L(\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;&= \\ &\tfrac{1}{2} \left[\|(\boldsymbol{s} - \boldsymbol{\Phi}\, \boldsymbol{b}_\star) - \delta\, \boldsymbol{\varphi}_\omega\|_2^2 \;-\; \|\boldsymbol{s} - \boldsymbol{\Phi}\, \boldsymbol{b}_\star\|_2^2\right] \\ &\qquad + \;\gamma \left[\|\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega\|_1 \;-\; \|\boldsymbol{b}_\star\|_1\right].\end{aligned}$$

Next, simplify the first bracket by expanding the $\ell_2$ norms and canceling like terms. Simplify the second bracket by recognizing that $\|\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega\|_1 = \|\boldsymbol{b}_\star\|_1 + \|\delta\, \mathbf{e}_\omega\|_1$ since the two vectors have disjoint supports. Hence

$$\begin{aligned}L(\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;&= \\ &\tfrac{1}{2}\, \|\delta\, \boldsymbol{\varphi}_\omega\|_2^2 \;-\; \mathrm{Re}\,\langle \boldsymbol{s} - \boldsymbol{\Phi}\, \boldsymbol{b}_\star,\; \delta\, \boldsymbol{\varphi}_\omega\rangle \;+\; \gamma\, |\delta|\,.\end{aligned}$$

Add and subtract $\boldsymbol{\Phi}_\Lambda\, \boldsymbol{c}_\Lambda$ in the left-hand side of the inner product, and use linearity to split the inner product into two pieces. We reach

$$\begin{aligned}L(\boldsymbol{b}_\star + \delta\, \mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;&= \\ &\tfrac{1}{2}\, \|\delta\, \boldsymbol{\varphi}_\omega\|_2^2 \;-\; \mathrm{Re}\,\langle \boldsymbol{s} - \boldsymbol{\Phi}_\Lambda\, \boldsymbol{c}_\Lambda,\; \delta\, \boldsymbol{\varphi}_\omega\rangle \\ &\qquad -\; \mathrm{Re}\,\langle \boldsymbol{\Phi}_\Lambda\, (\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star),\; \delta\, \boldsymbol{\varphi}_\omega\rangle \;+\; \gamma\, |\delta|\,.\end{aligned}$$

We will bound the right-hand side below. To that end, observe that the first term is strictly positive, so we may discard it.

Then invoke the lower triangle inequality, and use the linearity of the inner products to draw out $|\delta|$.

$$\begin{aligned} L(\boldsymbol{b}_\star + \delta\,\mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;> \\ \gamma\,|\delta| \;-\; |\langle \boldsymbol{\Phi}_\Lambda\,(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star),\; \boldsymbol{\varphi}_\omega\rangle|\;|\delta| \\ -\;|\langle \boldsymbol{s} - \boldsymbol{\Phi}_\Lambda\,\boldsymbol{c}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle|\;|\delta|\,. \end{aligned} \quad (14)$$

It remains to rewrite (14) in a favorable manner.

First, identify $\boldsymbol{a}_\Lambda = \boldsymbol{\Phi}_\Lambda\,\boldsymbol{c}_\Lambda$ in the last term on the right-hand side of (14). Next, let us examine the second term. Lemma 3 characterizes the difference $(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star)$. Introduce this characterization, and identify the pseudo-inverse of $\boldsymbol{\Phi}_\Lambda$ to discover that

$$\begin{aligned} |\langle \boldsymbol{\Phi}_\Lambda\,(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star),\; \boldsymbol{\varphi}_\omega\rangle| &= \gamma\,\big|\big\langle (\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\boldsymbol{g},\; \boldsymbol{\varphi}_\omega\big\rangle\big| \\ &= \gamma\,\big|\big\langle \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\big\rangle\big| \end{aligned} \quad (15)$$

where $\boldsymbol{g} \in \partial\,\|\boldsymbol{b}_\star\|_1$ is determined by (9). Substitute (15) into the bound (14).

$$\begin{aligned} L(\boldsymbol{b}_\star + \delta\,\mathbf{e}_\omega) \;-\; L(\boldsymbol{b}_\star) \;> \\ \Big[\gamma \;-\; \gamma\,\big|\big\langle \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\big\rangle\big| \;-\; \big|\big\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\big\rangle\big|\Big]\;|\delta|\,. \end{aligned} \quad (16)$$

Our final goal is to ensure that the left-hand side of (16) is strictly positive. This situation occurs whenever the bracket is nonnegative. Therefore, we need

$$|\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| \quad\leq\quad \gamma\,\Big[1 \;-\; \big|\big\langle \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\big\rangle\big|\Big]\,.$$

We require this expression to hold for every index $\omega$ that does not belong to $\Lambda$. Optimizing each side over $\omega$ in $\Omega \setminus \Lambda$ yields the stronger condition

$$\begin{aligned} \max\nolimits_{\omega\notin\Lambda} |\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| \;\leq \\ \gamma\,\Big[1 \;-\; \max\nolimits_{\omega\notin\Lambda}\big|\big\langle \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\big\rangle\big|\Big]\,. \end{aligned}$$

Since $(\boldsymbol{s} - \boldsymbol{a}_\Lambda)$ is orthogonal to the atoms listed in $\Lambda$, the left-hand side does not change if we maximize over all $\omega$ from $\Omega$. Therefore,

$$\begin{aligned} \max\nolimits_{\omega\notin\Lambda} |\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| &= \max\nolimits_{\omega\in\Omega} |\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| \\ &= \|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty\,. \end{aligned}$$

We conclude that the relation

$$\|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma\,\Big[1 \;-\; \max\nolimits_{\omega\notin\Lambda}\big|\big\langle \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega,\; \boldsymbol{g}\big\rangle\big|\Big]$$

is a sufficient condition for every perturbation away from $\boldsymbol{b}_\star$ to increase the objective function. Since (L) is convex, it follows that $\boldsymbol{b}_\star$ is the unique global minimizer of (L).

In particular, since $\|\boldsymbol{g}\|_\infty \leq 1$, it is also sufficient that

$$\|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma\,\Big[1 \;-\; \max\nolimits_{\omega\notin\Lambda}\big\|\boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\varphi}_\omega\big\|_1\Big]\,.$$

This completes the argument. ∎

## IV. $\ell_1$ Penalization

Suppose that $\boldsymbol{s}$ is an input signal. In this section, we study applications of the convex program

$$\min_{\boldsymbol{b}\in\mathbb{C}^\Omega} \quad \tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \gamma\,\|\boldsymbol{b}\|_1\,. \qquad (\ell_1\text{-PENALTY})$$

The parameter $\gamma$ controls the tradeoff between the error in approximating the input signal and the sparsity of the approximation.

We begin with a theorem that provides significant information about the solution to ($\ell_1$-PENALTY). Afterward, this theorem is applied to several important problems. As a first example, we show that the convex program can identify the sparsest representation of an exactly sparse signal. Our second example shows that ($\ell_1$-PENALTY) can be used to recover a sparse signal contaminated by an unknown noise vector of bounded norm, which is a well-known inverse problem. We will also see that a statistical model for the noise vector allows us to provide more refined results. In particular, we will discover that ($\ell_1$-PENALTY) is quite effective for recovering sparse signals corrupted with Gaussian noise. The section concludes with an application of the convex program to the subset selection problem from statistics.

The reader should also be aware that convex programs of the form ($\ell_1$-PENALTY) have been proposed for numerous other applications. Geophysicists have long used them for deconvolution [5, 35]. Certain support vector machines, which arise in machine learning, can be reduced to this form [8]. Chen, Donoho, and Saunders have applied ($\ell_1$-PENALTY) to de-noise signals [9], and Fuchs has put it forth for several other signal processing problems, e.g., in [36, 37]. Daubechies, Defrise, and De Mol have suggested a related convex program to regularize ill-posed linear inverse problems [10]. Most intriguing, perhaps, Olshausen and Field have argued that the mammalian visual cortex may solve similar optimization problems to produce sparse representations of images [38].

Even before we begin our analysis, we can make a few immediate remarks about ($\ell_1$-PENALTY). Observe that, as the parameter $\gamma$ tends to zero, the solution to the convex program approaches a point of minimal $\ell_1$ norm in the affine space $\{\boldsymbol{b} : \boldsymbol{\Phi}\,\boldsymbol{b} = \boldsymbol{s}\}$. It can be shown that, except when $\boldsymbol{s}$ belongs to a set of signals with Lebesgue measure zero in $\mathbb{C}^d$, no point of the affine space has fewer than $d$ nonzero components [14, Prop. 4.1]. So the minimizer of ($\ell_1$-PENALTY) is generically nonsparse when $\gamma$ is small. In contrast, as $\gamma$ approaches infinity, the solution tends toward the zero vector.

It is also worth noting that ($\ell_1$-PENALTY) has an analytical solution $\boldsymbol{b}_\star$ whenever $\boldsymbol{\Phi}$ is unitary. One simply computes the orthogonal expansion of the signal $\boldsymbol{s}$ with respect to the columns of $\boldsymbol{\Phi}$ and applies the soft thresholding operator to each coefficient. More precisely, $\boldsymbol{b}_\star(\omega) = T_\gamma(\langle \boldsymbol{s},\; \boldsymbol{\varphi}_\omega\rangle)$ where

$$T_\gamma(r\,\mathrm{e}^{\mathrm{i}\,\theta}) \;\stackrel{\text{def}}{=}\; \begin{cases} (r-\gamma)\,\mathrm{e}^{\mathrm{i}\,\theta}, & \text{if } r > \gamma \\ 0, & \text{otherwise.} \end{cases}$$

This result will be familiar to anyone who has studied the process of shrinking empirical wavelet coefficients to de-noise functions [39, 40].

### A. Performance of ($\ell_1$-PENALTY)

Our major theorem on the behavior of ($\ell_1$-PENALTY) simply collects the lemmata from the last section.

*Theorem 8:* Let $\Lambda$ index a linearly independent collection of atoms for which $\mathrm{ERC}(\Lambda) \geq 0$. Suppose that $\boldsymbol{s}$ is an input signal whose $\ell_2$ best approximation over $\Lambda$ satisfies the correlation condition

$$\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma \, \mathrm{ERC}(\Lambda).$$

Let $\boldsymbol{b}_\star$ solve the convex program ($\ell_1$-PENALTY) with parameter $\gamma$. We may conclude that

- the support of $\boldsymbol{b}_\star$ is contained in $\Lambda$, and
- the distance between $\boldsymbol{b}_\star$ and the optimal coefficient vector $\boldsymbol{c}_\Lambda$ satisfies

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_\infty \;\leq\; \gamma \left\|(\boldsymbol{\Phi}_\Lambda^* \boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty}.$$

- In particular, $\mathrm{supp}(\boldsymbol{b}_\star)$ contains every index $\lambda$ in $\Lambda$ for which

$$|\boldsymbol{c}_\Lambda(\lambda)| \;>\; \gamma \left\|(\boldsymbol{\Phi}_\Lambda^* \boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty}.$$

- Moreover, the minimizer $\boldsymbol{b}_\star$ is unique.

In words, we approximate the input signal over $\Lambda$, and we suppose that the remaining atoms are weakly correlated with the residual. It follows then that the minimizer $\boldsymbol{b}_\star$ of the convex program involves only atoms in $\Lambda$ and that this minimizer is not far from the coefficient vector that synthesizes the best approximation of the signal over $\Lambda$.

To check the hypotheses of this theorem, one must carefully choose the index set $\Lambda$ and leverage information about the relationship between the signal and its approximation over $\Lambda$. Our examples will show how to apply the theorem in several specific cases. At this point, we can also state a simpler version of this theorem that uses the coherence parameter to estimate some of the key quantities. One advantage of this formulation is that the index set $\Lambda$ plays a smaller role.

*Corollary 9:* Suppose that $m\,\mu \leq \frac{1}{2}$, and assume that $\Lambda$ contains no more than $m$ indices. Suppose that $\boldsymbol{s}$ is an input signal whose $\ell_2$ best approximation over $\Lambda$ satisfies the correlation condition

$$\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma \frac{1 \;-\; (2m-1)\,\mu}{1 \;-\; (m-1)\,\mu}.$$

Let $\boldsymbol{b}_\star$ solve ($\ell_1$-PENALTY) with parameter $\gamma$. We may conclude that

- the support of $\boldsymbol{b}_\star$ is contained in $\Lambda$, and
- the distance between $\boldsymbol{b}_\star$ and the optimal coefficient vector $\boldsymbol{c}_\Lambda$ satisfies

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_\infty \;\leq\; \frac{\gamma}{1 \;-\; (m-1)\,\mu}.$$

- In particular, $\mathrm{supp}(\boldsymbol{b}_\star)$ contains every index $\lambda$ in $\Lambda$ for which

$$|\boldsymbol{c}_\Lambda(\lambda)| \;>\; \frac{\gamma}{1 \;-\; (m-1)\,\mu}.$$

- Moreover, the minimizer $\boldsymbol{b}_\star$ is unique.

*Proof:* The corollary follows directly from the theorem when we apply the coherence bounds of Appendix III. ∎

This corollary takes an especially pleasing form if we assume that $m\,\mu \leq \frac{1}{3}$. In that case, the right-hand side of the correlation condition is no smaller than $\frac{1}{2}\,\gamma$. Meanwhile, the coefficient vectors $\boldsymbol{b}_\star$ and $\boldsymbol{c}_\Lambda$ differ in $\ell_\infty$ norm by no more than $\frac{3}{2}\,\gamma$.

Note that, when $\boldsymbol{\Phi}$ is unitary, the coherence parameter $\mu = 0$. Corollary 9 allows us to conclude that the solution to the convex program identifies precisely those atoms whose inner products against the signal exceed $\gamma$ in magnitude. Their coefficients are reduced in magnitude by at most $\gamma$. This description matches the performance of the soft thresholding operator.

### B. Example: Identifying Sparse Signals

As a simple application of the foregoing theorem, we offer a new proof that one can recover an exactly sparse signal by solving the convex program ($\ell_1$-PENALTY). Fuchs has already established this result in the real case [23].

*Corollary 10 (Fuchs):* Assume that $\Lambda$ indexes a linearly independent collection of atoms for which $\mathrm{ERC}(\Lambda) \geq 0$. Choose an arbitrary coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$ supported on $\Lambda$, and fix an input signal $\boldsymbol{s} = \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$. Let $\boldsymbol{b}_\star(\gamma)$ denote the unique minimizer of ($\ell_1$-PENALTY) with parameter $\gamma$. We may conclude that

- There is a positive number $\gamma_0$ for which $\gamma < \gamma_0$ implies that $\mathrm{supp}(\boldsymbol{b}_\star(\gamma)) = \Lambda$.
- In the limit as $\gamma \downarrow 0$, we have $\boldsymbol{b}_\star(\gamma) \to \boldsymbol{c}_{\mathrm{opt}}$.

*Proof:* First, note that the best approximation of the signal over $\Lambda$ satisfies $\boldsymbol{a}_\Lambda = \boldsymbol{s}$ and that the corresponding coefficient vector $\boldsymbol{c}_\Lambda = \boldsymbol{c}_{\mathrm{opt}}$. Therefore, $\|\boldsymbol{\Phi}^* (\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty = 0$. The Correlation Condition is in force for every positive $\gamma$, so Theorem 8 implies that the minimizer $\boldsymbol{b}_\star(\gamma)$ of the program ($\ell_1$-PENALTY) must be supported inside $\Lambda$. Moreover, the distance from $\boldsymbol{b}_\star(\gamma)$ to the optimal coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$ satisfies

$$\|\boldsymbol{b}_\star(\gamma) - \boldsymbol{c}_{\mathrm{opt}}\|_\infty \;\leq\; \gamma \left\|(\boldsymbol{\Phi}_\Lambda^* \boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty}.$$

It is immediate that $\boldsymbol{b}_\star(\gamma) \to \boldsymbol{c}_{\mathrm{opt}}$ as the parameter $\gamma \downarrow 0$. Finally, observe that $\mathrm{supp}(\boldsymbol{b}_\star)$ contains every index in $\Lambda$ provided that

$$\frac{\min_{\lambda \in \Lambda} |\boldsymbol{c}_{\mathrm{opt}}(\lambda)|}{\|(\boldsymbol{\Phi}_\Lambda^* \boldsymbol{\Phi}_\Lambda)^{-1}\|_{\infty,\infty}} \;>\; \gamma.$$

Note that the left-hand side of this inequality furnishes an explicit value for $\gamma_0$. ∎

As a further corollary, we obtain a familiar result for Basis Pursuit that was developed independently in [14, 23], generalizing the work in [19–22]. We will require this corollary in the sequel.

*Corollary 11 (Fuchs, Tropp):* Assume that $\mathrm{ERC}(\Lambda) > 0$. Let $\boldsymbol{c}_{\mathrm{opt}}$ be supported on $\Lambda$, and fix a signal $\boldsymbol{s} = \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$. Then $\boldsymbol{c}_{\mathrm{opt}}$ is the unique solution to

$$\min_{\boldsymbol{b} \in \mathbb{C}^{\Omega}} \|\boldsymbol{b}\|_1 \quad \text{subject to} \quad \boldsymbol{\Phi}\,\boldsymbol{b} = \boldsymbol{s}. \tag{BP}$$

*Proof:* As $\gamma \downarrow 0$, the solutions to ($\ell_1$-PENALTY) approach a point of minimum $\ell_1$ norm in the affine space $\{\boldsymbol{b} : \boldsymbol{\Phi}\,\boldsymbol{b} = \boldsymbol{s}\}$. Since $\boldsymbol{c}_{\mathrm{opt}}$ is the limit of these minimizers, it follows that $\boldsymbol{c}_{\mathrm{opt}}$ must also solve (BP). The uniqueness claim seems to require a separate argument and the assumption that $\mathrm{ERC}(\Lambda)$ is strictly positive. See [23, 41] for two different proofs. ∎

### C. Example: Identifying Sparse Signals in Noise

This subsection shows how ($\ell_1$-PENALTY) can be applied to the inverse problem of recovering a sparse signal contaminated with an arbitrary noise vector of bounded $\ell_2$ norm.

Let us begin with a model for our ideal signals. We will use this model for examples throughout the paper. Fix a $\mu$-coherent dictionary $\boldsymbol{\Phi}$ containing $N$ atoms in a real signal space of dimension $d$, and place the coherence bound $m\,\mu \le \frac{1}{3}$.

Each ideal signal is a linear combination of $m$ atoms with coefficients of $\pm 1$. To formalize things, select $\Lambda$ to be any nonempty index set containing $m$ atoms. Let $\boldsymbol{c}_{\mathrm{opt}}$ be a coefficient vector supported on $\Lambda$, and let the $m$ nonzero entries of $\boldsymbol{c}_{\mathrm{opt}}$ equal $\pm 1$. Then each ideal signal takes the form $\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$. To correctly recover one of these signals, it is necessary to determine the support set $\Lambda$ of the optimal coefficient vector as well as the signs of the nonzero coefficients.

Observe that the total number of ideal signals is

$$2^m \binom{N}{m}$$

since the choice of atoms and the choice of coefficients both carry information. The coherence estimate in Proposition 26 allows us to establish that the power $\|\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}\|_2^2$ of each ideal signal satisfies

$$\|\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}\|_2^2 \;\le\; \|\boldsymbol{\Phi}_\Lambda\|_{2,2}^2\,\|\boldsymbol{c}_{\mathrm{opt}}\|_2^2 \;\le\; 4\,m/3.$$

Similarly, no ideal signal has power less than $2\,m/3$.

In this example, we will form input signals by contaminating the ideal signals with an arbitrary noise vector $\boldsymbol{\nu}$ of bounded norm, say $\|\boldsymbol{\nu}\|_2 \le \varepsilon$. Therefore, we measure a signal of the form

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\nu}.$$

We would like to apply Corollary 9 to find circumstances in which the minimizer $\boldsymbol{b}_\star$ of ($\ell_1$-PENALTY) identifies the ideal signal. That is, $\operatorname{sgn}\boldsymbol{b}_\star = \operatorname{sgn}\boldsymbol{c}_{\mathrm{opt}}$. Observe that this is a concrete example of the model problem from the Introduction.

First, let us determine when the Correlation Condition holds for $\Lambda = \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$. According to the Pythagorean Theorem,

$$\begin{aligned}\|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2 \;&=\; \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}\|_2^2 \;-\; \|\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} - \boldsymbol{a}_\Lambda\|_2^2 \\ &=\; \|\boldsymbol{\nu}\|_2^2 \;-\; \|\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} - \boldsymbol{a}_\Lambda\|_2^2\,.\end{aligned}$$

Therefore, $\|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2 \le \varepsilon$. Each atom has unit $\ell_2$ norm, so

$$\|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;=\; \max_{\omega\in\Omega} |\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle| \;\le\; \varepsilon.$$

Referring to the paragraph after Corollary 9, we see that the correlation condition is in force if we select $\gamma = 2\,\varepsilon$. Invoking Corollary 9, we discover that the minimizer $\boldsymbol{b}_\star$ of ($\ell_1$-PENALTY) with parameter $\gamma$ is supported inside $\Lambda$ and also that

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_\infty \;\le\; 3\,\gamma/2 \;=\; 3\,\varepsilon. \tag{17}$$

Meanwhile, we may calculate that

$$\begin{aligned}\|\boldsymbol{c}_\Lambda - \boldsymbol{c}_{\mathrm{opt}}\|_\infty \;&=\; \left\|\boldsymbol{\Phi}_\Lambda^\dagger\,(\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}})\right\|_\infty \\ &=\; \left\|\boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\nu}\right\|_\infty \\ &\le\; \left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,\infty}\,\|\boldsymbol{\nu}\|_2\,.\end{aligned}$$

The coherence bound in Proposition 26 delivers

$$\|\boldsymbol{c}_\Lambda - \boldsymbol{c}_{\mathrm{opt}}\|_\infty \;\le\; \varepsilon\,\sqrt{3/2}. \tag{18}$$

In consequence of the triangle inequality and the bounds (17) and (18),

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_\infty \;\le\; \left(3 + \sqrt{3/2}\right)\,\varepsilon.$$

It follows that $\operatorname{sgn}\boldsymbol{b}_\star = \operatorname{sgn}\boldsymbol{c}_{\mathrm{opt}}$ provided that

$$\varepsilon \;<\; \frac{1}{3 + \sqrt{3/2}} \;\approx\; 0.2367.$$

In conclusion, the convex program ($\ell_1$-PENALTY) can always recover the ideal signals provided that the $\ell_2$ norm of the noise is less than $(3 + \sqrt{3/2})^{-1}$. At this noise level, the signal-to-noise ratio (SNR) is no greater than

$$\mathrm{SNR} \;\le\; \frac{4\,m/3}{(3 + \sqrt{3/2})^{-2}} \;=\; 2\,m\,(1 + \sqrt{6})^2.$$

Similarly, the SNR is no smaller than $m\,(1 + \sqrt{6})^2$. Note that the SNR grows linearly with the number of atoms in the signal.

Let's consider some specific numbers. Suppose that we are working in a signal space with dimension $d = 2^8$, that the dictionary contains $N = 2^{15}$ atoms, that the coherence $\mu = 2^{-4}$, and that the sparsity level $m = 5$. We can recover any of the ideal signals provided that the noise level $\varepsilon < 0.2367$. For the supremal value of $\varepsilon$, the SNR lies between 17.74 dB and 20.76 dB. The total number of ideal signals is just over $2^{68}$, so each one encodes 68 bits of information.

On the other hand, if we have a statistical model for the noise, we can obtain good performance even when the noise level is substantially higher. In the next subsection, we will describe an example of this stunning phenomenon.

### D. Example: The Gaussian Channel

A striking application of $\ell_1$ penalization is to recover a sparse signal contaminated with Gaussian noise. The fundamental reason this method succeeds is that, with high probability, the noise is weakly correlated with every atom in the dictionary (provided that the number of atoms $N$ is sub-exponential in the dimension $d$). Note that Fuchs has developed qualitative results for this type of problem in

his conference publication [26]; see Section VI-B for some additional details on his work.

Let us continue with the ideal signals described in the last subsection. This time, the ideal signals are corrupted by adding a zero-mean Gaussian random vector $\boldsymbol{\nu}$ with covariance matrix $\sigma^2 \mathbf{I}$. We measure the signal

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\, \boldsymbol{c}_{\text{opt}} \;+\; \boldsymbol{\nu},$$

and we wish to identify $\boldsymbol{c}_{\text{opt}}$. In other words, a codeword is sent through a Gaussian channel, and we must decode the transmission. We will approach this problem by solving the convex program ($\ell_1$-PENALTY) with an appropriate choice of $\gamma$ to obtain a minimizer $\boldsymbol{b}_\star$.

One may establish the following facts about this signal model and the performance of convex relaxation.

- For each ideal signal, the SNR satisfies
$$\frac{2\,m}{3\,d}\,\sigma^{-2} \;\le\; \text{SNR} \;\le\; \frac{4\,m}{3\,d}\,\sigma^{-2}.$$
- The capacity of a Gaussian channel [28, Chap. 10] at this SNR is no greater than
$$\tfrac{1}{2}\,\log_2\left[1 + \frac{4\,m}{3\,d}\sigma^{-2}\right].$$
- For our ideal signals, the number of bits per transmission equals
$$\frac{1}{d}\,\left[m + \log_2 \binom{N}{m}\right].$$
- The probability that convex relaxation correctly identifies the ideal signal exceeds
$$\left[1 \;-\; \exp\left\{-\tfrac{1}{8}\,\gamma^2\,\sigma^{-2}\right\}\right]^{N-m} \left[1 \;-\; \exp\left\{-\tfrac{1}{3}\,(1-\tfrac{3}{2}\,\gamma)^2\sigma^{-2}\right\}\right]^{m}.$$
It follows that the failure probability decays exponentially as the noise power drops to zero.

To establish the first item, we compare the power of the ideal signals against the power of the noise, which is $\sigma^2\,d$. To determine the likelihood of success, we must find the probability (over the noise) that the Correlation Condition is in force so that we can invoke Corollary 9. Then we must ensure that the noise does not corrupt the coefficients enough to obscure their signs. The difficult calculations are consigned to Appendix IV-A.

We return to the same example from the last subsection. Suppose that we are working in a signal space with dimension $d = 2^8$, that the dictionary contains $N = 2^{15}$ atoms, that the coherence $\mu = 2^{-4}$, and that the sparsity level is $m = 5$. Assume that the noise level $\sigma = 0.05$, which is about as large as we can reliably handle. In this case, a good choice for the parameter $\gamma$ is 0.53. With these selections, the SNR is between 7.17 dB and 10.18 dB, and the channel capacity does not exceed 1.757 bits per transmission. Meanwhile, we are sending about 0.267 bits per transmission, and the probability of perfect recovery exceeds 95% over the noise. Although there is now a small failure probability, let us emphasize that the SNR in this example is over 10 dB lower than in the last example. This improvement is possible because we have accounted for the direction of the noise.

Let's be honest. This example still does not inspire much confidence in our coding scheme: the theoretical rate we have established is nowhere near the actual capacity of the channel. The shortcoming, however, is not intrinsic to the coding scheme. To obtain rates close to capacity, it is necessary to send linear combinations whose length $m$ is on the order of the signal dimension $d$. Experiments indicate that convex relaxation can indeed recover linear combinations of this length. But the analysis to support these empirical results requires tools much more sophisticated than the coherence parameter. We hope to pursue these ideas in a later work.

### *E. Example: Subset Selection*

Statisticians often wish to predict the value of one random variable using a linear combination of other random variables. At the same time, they must negotiate a compromise between the number of variables involved and the mean squared prediction error to avoid overfitting. The problem of determining the correct variables is called *subset selection*, and it was probably the first type of sparse approximation to be studied in depth. As Miller laments, statisticians have made limited theoretical progress due to numerous complications that arise in the stochastic setting [4, Prefaces].

We will consider a deterministic version of subset selection that manages a simple tradeoff between the squared approximation error and the number of atoms that participate. Let $\boldsymbol{s}$ be an arbitrary input signal. Suppose $\tau$ is a threshold that quantifies how much improvement in the approximation error is necessary before we admit an additional term into the approximation. We may state the formal problem

$$\min_{\boldsymbol{c}\,\in\,\mathbb{C}^\Omega} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2^2 \;+\; \tau^2\,\|\boldsymbol{c}\|_0\,. \qquad \text{(SUBSET)}$$

Were the support of $\boldsymbol{c}$ fixed, then (SUBSET) would be a least-squares problem. Selecting the optimal support, however, is a combinatorial nightmare. In fact, if the dictionary is unrestricted, it must be NP-hard to solve (SUBSET) in consequence of results from [12, 16].

The statistics literature contains dozens of algorithmic approaches to subset selection, which [4] describes in detail. A method that has recently become popular is *the lasso*, which replaces the difficult subset selection problem with a convex relaxation of the form ($\ell_1$-PENALTY) in hope that the solutions will be related [42]. Our example provides a rigorous justification that this approach can succeed. If we have some basic information about the solution to (SUBSET), then we may approximate this solution using ($\ell_1$-PENALTY) with an appropriate choice of $\gamma$.

We will invoke Theorem 8 to show that the solution to the convex relaxation has the desired properties. To do so, we require a theorem about the behavior of solutions to the subset selection problem. The proof appears in Appendix V.

*Theorem 12:* Fix an input signal $\boldsymbol{s}$, and choose a threshold $\tau$. Suppose that the coefficient vector $\boldsymbol{c}_{\text{opt}}$ solves the subset selection problem, and set $\boldsymbol{a}_{\text{opt}} = \boldsymbol{\Phi}\,\boldsymbol{c}_{\text{opt}}$.

- For $\lambda \in \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$, we have $|\boldsymbol{c}_{\mathrm{opt}}(\lambda)| \geq \tau$.
- For $\omega \notin \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$, we have $|\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\ \boldsymbol{\varphi}_\omega \rangle| \leq \tau$.

In consequence of this theorem, any solution $\boldsymbol{c}_{\mathrm{opt}}$ to the subset selection problem satisfies the Correlation Condition with $\Lambda_{\mathrm{opt}} = \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$, provided that $\gamma$ is chosen so that

$$\tau \ \leq \ \gamma \operatorname{ERC}(\Lambda_{\mathrm{opt}}).$$

Applying Theorem 8 yields the following result.

*Corollary 13 (Relaxed Subset Selection):* Fix an input signal $\boldsymbol{s}$. Suppose that

- the vector $\boldsymbol{c}_{\mathrm{opt}}$ solves (SUBSET) with threshold $\tau$;
- the set $\Lambda_{\mathrm{opt}} = \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$ satisfies $\operatorname{ERC}(\Lambda_{\mathrm{opt}}) > 0$; and
- the coefficient vector $\boldsymbol{b}_\star$ solves ($\ell_1$-PENALTY) with parameter $\gamma \geq \tau \,/\, \operatorname{ERC}(\Lambda_{\mathrm{opt}})$.

Then it follows that

- the relaxation never selects a nonoptimal atom since $\operatorname{supp}(\boldsymbol{b}_\star) \subset \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$;
- the solution to the relaxation is nearly optimal since

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_\infty \ \leq \ \frac{\left\|(\boldsymbol{\Phi}_{\mathrm{opt}}^* \boldsymbol{\Phi}_{\mathrm{opt}})^{-1}\right\|_{\infty,\infty}}{\operatorname{ERC}(\Lambda_{\mathrm{opt}})} \ \tau.$$

- In particular, $\operatorname{supp}(\boldsymbol{b}_\star)$ contains every index $\lambda$ for which

$$|\boldsymbol{c}_{\mathrm{opt}}(\lambda)| \ > \ \frac{\left\|(\boldsymbol{\Phi}_{\mathrm{opt}}^* \boldsymbol{\Phi}_{\mathrm{opt}})^{-1}\right\|_{\infty,\infty}}{\operatorname{ERC}(\Lambda_{\mathrm{opt}})} \ \tau.$$

- Moreover, $\boldsymbol{b}_\star$ is the unique solution to ($\ell_1$-PENALTY).

In words, if any solution to the subset selection problem satisfies a condition on the Exact Recovery Coefficient, then the solution to the convex relaxation ($\ell_1$-PENALTY) for an appropriately chosen parameter $\gamma$ will identify every significant atom in the solution to (SUBSET) and it will never involve any atom that does not belong in that optimal solution.

It is true that, in the present form, the hypotheses of this corollary may be difficult to verify. Invoking Corollary 9 instead of Theorem 8 would yield a more practical result involving the coherence parameter. Nevertheless, this result would still involve a strong assumption on the sparsity of an optimal solution to the subset selection problem. It is not clear that one could verify this hypothesis in practice, so it may be better to view these results as philosophical support for the practice of convex relaxation.

For an incoherent dictionary, one could develop a converse result of the following shape.

> Suppose that the solution to the convex relaxation ($\ell_1$-PENALTY) is sufficiently sparse, has large enough coefficients, and yields a good approximation of the input signal. Then the solution to the relaxation must approximate an optimal solution to (SUBSET).

As this paper was being completed, it came to the author's attention that Gribonval et al. have developed some results of this type [43]. Their theory should complement the present work nicely.

## V. Error-Constrained $\ell_1$ Minimization

Suppose that $\boldsymbol{s}$ is an input signal. In this section, we study applications of the convex program.

$$\min_{\boldsymbol{b} \in \mathbb{C}^\Omega} \ \|\boldsymbol{b}\|_1 \quad \text{subject to} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2 \ \leq \ \delta. \quad (\ell_1\text{-ERROR})$$

Minimizing the $\ell_1$ norm of the coefficients promotes sparsity, while the parameter $\delta$ controls how much approximation error we are willing to tolerate.

We will begin with a theorem that yields significant information about the solution to ($\ell_1$-ERROR). Then, we will tour several different applications of this optimization problem. As a first example, we will see that it can recover a sparse signal that has been contaminated with an arbitrary noise vector of bounded $\ell_2$ norm. Afterward, we show that a statistical model for the bounded noise allows us to sharpen our analysis significantly. Third, we describe an application to a sparse approximation problem that arises in numerical analysis.

The literature does not contain many papers that apply the convex program ($\ell_1$-ERROR). Indeed, the other optimization problem ($\ell_1$-PENALTY) is probably more useful. Nevertheless, there is one notable study [18] of the theoretical performance of ($\ell_1$-ERROR). This manuscript will be discussed in more detail in Section VI-B.

Let us conclude this introduction by mentioning some of the basic properties of ($\ell_1$-ERROR). As the parameter $\delta$ approaches zero, the solutions will approach a point of minimal $\ell_1$ norm in the affine space $\{\boldsymbol{b} : \boldsymbol{\Phi}\,\boldsymbol{b} = \boldsymbol{s}\}$. On the other hand, as soon as $\delta$ exceeds $\|\boldsymbol{s}\|_2$, the unique solution to ($\ell_1$-ERROR) is the zero vector.

### A. Performance of ($\ell_1$-ERROR)

The following theorem describes the behavior of a minimizer of ($\ell_1$-ERROR). In particular, it provides conditions under which the support of the minimizer is contained in a specific index set $\Lambda$. We reserve the proof until Section V-E.

*Theorem 14:* Let $\Lambda$ index a linearly independent collection of atoms for which $\operatorname{ERC}(\Lambda) > 0$, and fix an input signal $\boldsymbol{s}$. Select an error tolerance $\delta$ no smaller than

$$\left[1 \ + \ \left(\frac{\operatorname{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\,\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,1}}{\operatorname{ERC}(\Lambda)}\right)^2\right]^{1/2} \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2 \,.$$

Let $\boldsymbol{b}_\star$ solve the convex program ($\ell_1$-ERROR) with tolerance $\delta$. We may conclude that

- the support of $\boldsymbol{b}_\star$ is contained in $\Lambda$,
- the distance between $\boldsymbol{b}_\star$ and the optimal coefficient vector $\boldsymbol{c}_\Lambda$ satisfies

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_2 \ \leq \ \delta\,\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}.$$

- In particular, $\boldsymbol{b}_\star$ contains every index $\lambda$ from $\Lambda$ for which

$$|\boldsymbol{c}_\Lambda(\lambda)| \ > \ \delta\,\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}.$$

- Moreover, the minimizer $\boldsymbol{b}_\star$ is unique.

In words, if the parameter $\delta$ is chosen somewhat larger than the error in the optimal approximation $\boldsymbol{a}_\Lambda$, then the solution

to ($\ell_1$-ERROR) identifies every significant atom in $\Lambda$ and it never picks an incorrect atom. Note that the manuscript [18] refers to the first conclusion as a *support result* and the second conclusion as a *stability result*.

To invoke the theorem, it is necessary to choose the parameter $\delta$ carefully. Our examples will show how this can be accomplished in several specific cases. Using the coherence parameter, it is possible to state a somewhat simpler result.

*Corollary 15:* Suppose that $m\,\mu \leq \frac{1}{2}$, and assume that $\Lambda$ lists no more than $m$ atoms. Suppose that $\boldsymbol{s}$ is an input signal, and choose the error tolerance $\delta$ no smaller than

$$\left[1 \;+\; \frac{m\,(1 \;-\; (m-1)\,\mu)}{(1 \;-\; (2m-1)\,\mu)^2}\,\operatorname{maxcor}(\boldsymbol{s}-\boldsymbol{a}_\Lambda)^2\right]^{1/2} \|\boldsymbol{s}-\boldsymbol{a}_\Lambda\|_2\,.$$

Let $\boldsymbol{b}_\star$ solve ($\ell_1$-ERROR) with tolerance $\delta$. We may conclude that $\operatorname{supp}(\boldsymbol{b}_\star) \subset \Lambda$. Furthermore, we have $\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_2 \leq \delta \left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}$.

*Proof:* The corollary follows directly from the theorem when we apply the coherence bounds from Appendix III. ∎

This corollary takes a most satisfying form under the assumption that $m\,\mu \leq \frac{1}{3}$. In that case, if we choose the tolerance

$$\delta \;\geq\; \sqrt{1+6\,m\,\operatorname{maxcor}(\boldsymbol{s}-\boldsymbol{a}_\Lambda)^2}\;\|\boldsymbol{s}-\boldsymbol{a}_\Lambda\|_2\,,$$

then we may conclude that $\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_2 \leq \delta\,\sqrt{3/2}$.

### B. Example: Identifying Sparse Signals in Noise (Redux)

In Section IV-C, we showed that ($\ell_1$-PENALTY) can be used to solve the inverse problem of recovering sparse signals corrupted with arbitrary noise of bounded magnitude. This example will demonstrate that ($\ell_1$-ERROR) can be applied to the same problem.

We proceed with the same ideal signals described in Section IV-C. For reference, we quickly repeat the particulars. Fix a $\mu$-coherent dictionary $\boldsymbol{\Phi}$ containing $N$ atoms in a $d$-dimensional real signal space, and assume that $m\,\mu \leq \frac{1}{3}$. Let $\Lambda$ index $m$ atoms, and let $\boldsymbol{c}_{\mathrm{opt}}$ be a coefficient vector supported on $\Lambda$ whose $m$ nonzero entries equal $\pm 1$. Each ideal signal takes the form $\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$.

In this example, the ideal signals are contaminated by an unknown vector $\boldsymbol{\nu}$ with $\ell_2$ norm no greater than $\varepsilon$ to obtain an input signal

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\nu}.$$

We wish to apply Corollary 15 to determine when the minimizer $\boldsymbol{b}_\star$ of ($\ell_1$-ERROR) can identify the ideal signals. That is, $\operatorname{sgn}\boldsymbol{b}_\star = \operatorname{sgn}\boldsymbol{c}_{\mathrm{opt}}$.

First, we must determine what tolerance $\delta$ to use for the convex program. Since there is no information on the direction of the residual, we must use the bound $\operatorname{maxcor}(\boldsymbol{s}-\boldsymbol{a}_\Lambda) \leq 1$. As in Section IV-C, the norm of the residual satisfies $\|\boldsymbol{s}-\boldsymbol{a}_\Lambda\|_2 \leq \varepsilon$. Therefore, we may choose $\delta = \varepsilon\,\sqrt{1+6\,m}$.

Invoke Corollary 15 to determine that the support of $\boldsymbol{b}_\star$ is contained in $\Lambda$ and that

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_\Lambda\|_2 \;\leq\; \delta\,\sqrt{3/2}.$$

A modification of the argument in Section IV-C yields

$$\|\boldsymbol{c}_\Lambda - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; \left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}\,\|\boldsymbol{s}-\boldsymbol{a}_\Lambda\|_2\,.$$

Proposition 26 provides an upper bound for this operator norm, from which we conclude that

$$\|\boldsymbol{c}_\Lambda - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; \varepsilon\,\sqrt{3/2}.$$

The triangle inequality furnishes

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; (\varepsilon+\delta)\,\sqrt{3/2}.$$

Recalling the value of $\delta$, we have

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; \varepsilon\,(1 \;+\; \sqrt{1+6\,m})\,\sqrt{3/2}.$$

Therefore, $\operatorname{sgn}\boldsymbol{b}_\star = \operatorname{sgn}\boldsymbol{c}_{\mathrm{opt}}$ provided that

$$\varepsilon \;<\; \frac{\sqrt{2/3}}{1+\sqrt{1+6\,m}}.$$

Note that this upper bound decreases as $m$ increases.

Let us consider the same specific example as before. Suppose that we are working in a real signal space of dimension $d = 2^8$, that the dictionary contains $2^{15}$ atoms, that the coherence level $\mu = 2^{-4}$, and that the level of sparsity $m = 5$. Then we may select $\delta = \varepsilon\,\sqrt{31} \approx 5.568\,\varepsilon$. To recover the sign of each coefficient, we need $\varepsilon < 0.1243$. At this noise level, the SNR is between 23.34 dB and 26.35 dB. Comparing this example with Section IV-C, we conclude that ($\ell_1$-PENALTY) provides a more robust method for recovering our ideal signals.

The manuscript [18] of Donoho et al. contains some results related to Corollary 15. That article splits Corollary 15 into two pieces. Their Theorem 6.1 guarantees that the minimizer of ($\ell_1$-ERROR) has the desired support properties, while their Theorem 3.1 ensures that the minimizer approximates $\boldsymbol{c}_{\mathrm{opt}}$, the optimal coefficient vector. We will use our running example to illustrate a difference between their theory and ours. Their support result also requires that $\delta \geq \varepsilon\,\sqrt{31}$ to ensure that the support of the minimizer $\boldsymbol{b}_\star$ is contained in the ideal support. Their stability result requires that $m\,\mu \leq \frac{1}{4}$, so it does not allow us to reach any conclusions about the distance between $\boldsymbol{b}_\star$ and $\boldsymbol{c}_{\mathrm{opt}}$.

### C. Example: The Uniform Channel

A compelling application of Theorem 14 is to recover a sparse signal that is polluted by uniform noise with bounded $\ell_2$ norm. As before, the fundamental reason that convex relaxation succeeds is that, with high probability, the noise is weakly correlated with every atom in the dictionary. This subsection describes a simple example that parallels the case of Gaussian noise.

We retain the same model for our ideal signals. This time, we add a random vector $\boldsymbol{\nu}$ that is uniformly distributed on the $\ell_2$ ball of radius $\varepsilon$. The measured signals look like

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\nu}.$$

In words, we transmit a codeword through a (somewhat unusual) communication channel. Of course, we will attempt to recover the ideal signal by solving ($\ell_1$-ERROR) with an appropriate tolerance $\delta$ to obtain a minimizer $\boldsymbol{b}_\star$. The minimizer $\boldsymbol{b}_\star$ correctly identifies the optimal coefficients if and only if $\operatorname{sgn} \boldsymbol{b}_\star = \operatorname{sgn} \boldsymbol{c}_{\mathrm{opt}}$.

Let $t$ be an auxiliary parameter, and suppose that we choose the tolerance

$$\delta \;=\; \left[1 + 6\,m\,t^2\right]^{1/2}\,\varepsilon.$$

We may establish the following facts about the minimizer $\boldsymbol{b}_\star$ of ($\ell_1$-ERROR).

- The probability (over the noise) that we may invoke Corollary 15 exceeds

$$1 \;-\; 2\,(N-m)\,\exp\left\{-\tfrac{1}{2}\,t^2\,(d-m)\right\}.$$

- In this event, we conclude that $\operatorname{supp}(\boldsymbol{b}_\star) \subset \Lambda$ and that
- the distance between the coefficient vectors satisfies

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\le\; (\varepsilon + \delta)\,\sqrt{3/2}.$$

In particular, $\operatorname{sgn} \boldsymbol{b}_\star = \operatorname{sgn} \boldsymbol{c}_{\mathrm{opt}}$ whenever

$$(\varepsilon + \delta)\,\sqrt{3/2} \;<\; 1.$$

To calculate the success probability, we must study the distribution of the correlation between the noise and the nonoptimal atoms. We relegate the difficult calculations to Appendix IV-B. The other two items were established in the last subsection.

Let's do the numbers. Suppose that we are working in a real signal space of dimension $d = 2^8$, that the dictionary contains $N = 2^{15}$ atoms, that the coherence level $\mu = 2^{-4}$, and that the level of sparsity $m = 5$. To invoke Corollary 15 with probability greater than 95% over the noise, it suffices to choose $t^2 = 0.1123$. This corresponds to the selection $\delta = 2.090\,\varepsilon$. In the event that the corollary applies, we recover the index set and the signs of the coefficients provided that $\varepsilon < 0.2642$. At this noise level, the SNR is between 16.79 dB and 19.80 dB.

Once again, by taking the direction of the noise into account, we have been able to improve substantially on the more naïve approach described in the last subsection. In contrast, none of the results in [18] account for the direction of the noise.

### D. Example: Error-Constrained Sparse Approximation

In numerical analysis, a common problem is to approximate or interpolate a complicated function using a short linear combination of more elementary functions. The approximation must not commit too great an error. At the same time, one pays for each additional term in the linear combination whenever the approximation is evaluated. Therefore, one may wish to maximize the sparsity of the approximation subject to an error constraint [16].

Suppose that $\boldsymbol{s}$ is an arbitrary input signal, and fix an error level $\varepsilon$. The sparse approximation problem we have described may be stated as

$$\min_{\boldsymbol{c} \in \mathbb{C}^\Omega} \|\boldsymbol{c}\|_0 \quad \text{subject to} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2 \;\le\; \varepsilon. \tag{19}$$

Observe that this mathematical program will generally have many solutions that use the same number of atoms but yield different approximation errors. The solutions will always become sparser as the error tolerance increases. Instead, we consider the more convoluted mathematical program

$$\min_{\boldsymbol{c} \in \mathbb{C}^\Omega} \|\boldsymbol{c}\|_0 \;+\; \tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2\,\varepsilon^{-1} \quad \text{subject to} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2 \;\le\; \varepsilon. \tag{ERROR}$$

Any minimizer of (ERROR) also solves (19), but it produces the smallest approximation error possible at that level of sparsity.

We will attempt to produce a solution to the error-constrained sparse approximation problem by solving the convex program ($\ell_1$-ERROR) for a value of $\delta$ somewhat larger than $\varepsilon$. Our major result shows that, under appropriate conditions, this procedure approximates the solution of (ERROR).

*Corollary 16 (Relaxed Sparse Approximation):* Fix an input signal $\boldsymbol{s}$. Suppose that

- the vector $\boldsymbol{c}_{\mathrm{opt}}$ solves (ERROR) with tolerance $\varepsilon$;
- the set $\Lambda_{\mathrm{opt}} = \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$ satisfies the condition $\operatorname{ERC}(\Lambda_{\mathrm{opt}}) > 0$; and
- the vector $\boldsymbol{b}_\star$ solves the convex relaxation ($\ell_1$-ERROR) with threshold $\delta$ such that

$$\delta \;\ge\; \left[1 \;+\; \left(\frac{\left\|\boldsymbol{\Phi}_{\mathrm{opt}}^\dagger\right\|_{2,1}}{\operatorname{ERC}(\Lambda_{\mathrm{opt}})}\right)^2\right]^{1/2} \varepsilon.$$

Then it follows that

- the relaxation never selects a nonoptimal atom since $\operatorname{supp}(\boldsymbol{b}_\star) \subset \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$;
- yet $\boldsymbol{b}_\star$ is no sparser than a solution to (ERROR) with tolerance $\delta$.
- The solution of the relaxation is nearly optimal since

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\le\; \delta\,\left\|\boldsymbol{\Phi}_{\mathrm{opt}}^\dagger\right\|_{2,2};$$

- in particular, $\operatorname{supp}(\boldsymbol{b}_\star)$ contains every index $\lambda$ for which

$$|\boldsymbol{c}_{\mathrm{opt}}(\lambda)| \;>\; \delta\,\left\|\boldsymbol{\Phi}_{\mathrm{opt}}^\dagger\right\|_{2,2}.$$

- Moreover, $\boldsymbol{b}_\star$ is the unique solution of ($\ell_1$-ERROR).

*Proof:* Note that (ERROR) delivers a maximally sparse vector in the set $\{\boldsymbol{c} : \|\boldsymbol{s} - \boldsymbol{\Phi}\boldsymbol{c}\|_2 \le \delta\}$. Since $\boldsymbol{b}_\star$ is also in this set, it certainly cannot be any sparser. The other claims follow directly by applying Theorem 14 to $\boldsymbol{s}$ with the index set $\Lambda_{\mathrm{opt}} = \operatorname{supp}(\boldsymbol{c}_{\mathrm{opt}})$. We also use the facts that $\|\boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}}\|_2 \le \varepsilon$ and that the maximum correlation function never exceeds one. Finally, ∎

It may be somewhat difficult to verify the hypotheses of this corollary. In consequence, it might be better to view the result as a philosophical justification for the practice of convex relaxation.

In the present setting, there is a sort of converse to Corollary 16 that follows from the work in [18]. This result allows us

to verify that a solution to the convex relaxation approximates the solution to (ERROR).

*Proposition 17:* Assume that the dictionary has coherence $\mu$, and let $\boldsymbol{s}$ be an input signal. Suppose that $\boldsymbol{b}_\star$ is a coefficient vector whose support contains $m$ indices, where $m\,\mu < \frac{1}{2}$, and suppose that

$$\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2 \;\leq\; \varepsilon.$$

Then we may conclude that the optimal solution $\boldsymbol{c}_{\mathrm{opt}}$ to (ERROR) with tolerance $\varepsilon$ satisfies

$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; \frac{2\,\varepsilon}{\sqrt{1 - 2\,m\,\mu}}.$$

*Proof:* Theorem 2.12 of [18] can be rephrased as follows. Suppose that

$$\|\boldsymbol{s} - \boldsymbol{\Phi}\boldsymbol{b}_\star\|_2 \leq \varepsilon \quad\text{and}\quad \|\boldsymbol{b}_\star\|_0 < \tfrac{1}{2}\,\mathrm{spark}_\eta(\boldsymbol{\Phi});$$
$$\|\boldsymbol{s} - \boldsymbol{\Phi}\boldsymbol{c}_{\mathrm{opt}}\|_2 \leq \varepsilon \quad\text{and}\quad \|\boldsymbol{c}_{\mathrm{opt}}\|_0 < \tfrac{1}{2}\,\mathrm{spark}_\eta(\boldsymbol{\Phi}).$$

Then $\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \leq 2\varepsilon/\eta$. We may apply this result in the present context by making several observations.

If $\boldsymbol{b}_\star$ has sparsity level $m$, then the optimal solution $\boldsymbol{c}_{\mathrm{opt}}$ to (ERROR) with tolerance $\varepsilon$ certainly has a sparsity level no greater than $m$. Set $\eta = \sqrt{1 - 2\,m\,\mu}$, and apply Lemma 2.8 of [18] to see that $\mathrm{spark}_\eta(\boldsymbol{\Phi}) \geq 2\,m + 1$. ∎

More results of this type would be extremely valuable, because they make it possible to verify that a given coefficient vector approximates the solution to (ERROR). For the most recent developments, we refer the reader to the report [43].

### *E. Proof of Theorem 14*

Now we demonstrate that Theorem 14 holds. This argument takes some work because we must check the Karush–Kuhn–Tucker necessary and sufficient conditions for a solution to the convex program ($\ell_1$-ERROR).

*Proof:* Suppose that $\Lambda$ indexes a linearly independent collection of atoms and that $\mathrm{ERC}(\Lambda) > 0$. Fix an input signal $\boldsymbol{s}$, and choose a tolerance $\delta$ so that

$$\delta^2 \;\geq\; \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2 \;+\; \frac{\|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty^2\;\big\|\boldsymbol{\Phi}_\Lambda^\dagger\big\|_{2,1}^2}{\mathrm{ERC}(\Lambda)^2} \qquad (20)$$

This is the same bound as in the statement of the theorem, but it has been rewritten in a more convenient fashion. Without loss of generality, take $\|\boldsymbol{s}\|_2 > \delta$. Otherwise, the unique solution to ($\ell_1$-ERROR) is the zero vector.

Consider the restricted minimization problem

$$\min_{\{\boldsymbol{b}\,:\,\mathrm{supp}(\boldsymbol{b})\subset\Lambda\}} \|\boldsymbol{b}\|_1 \quad\text{subject to}\quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2 \;\leq\; \delta. \qquad (21)$$

Since the objective function and the constraint set are both convex, the Karush–Kuhn–Tucker conditions are necessary and sufficient [34, Chapter 28]. They guarantee that the (unique) minimizer $\boldsymbol{b}_\star$ of the program (21) satisfies

$$\boldsymbol{b}_\star \;\in\; \operatorname*{arg\,min}_{\{\boldsymbol{b}\,:\,\mathrm{supp}(\boldsymbol{b})\subset\Lambda\}} \zeta_\star\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \|\boldsymbol{b}\|_1 \qquad (22)$$

$$\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2 \;=\; \delta \qquad (23)$$

where $\zeta_\star$ is a strictly positive Lagrange multiplier. If we define $\gamma_\star = 1/(2\,\zeta_\star)$, then we may rewrite (22) as

$$\boldsymbol{b}_\star \;\in\; \operatorname*{arg\,min}_{\{\boldsymbol{b}\,:\,\mathrm{supp}(\boldsymbol{b})\subset\Lambda\}} \tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \gamma_\star\,\|\boldsymbol{b}\|_1. \qquad (24)$$

We will show that $\boldsymbol{b}_\star$ is also a global minimizer of the convex program ($\ell_1$-ERROR).

Since $\boldsymbol{a}_\Lambda$ is the best $\ell_2$ approximation of $\boldsymbol{s}$ using the atoms in $\Lambda$ and the signal $\boldsymbol{\Phi}\,\boldsymbol{b}_\star$ is also a linear combination of atoms in $\Lambda$, it follows that the vectors $(\boldsymbol{s} - \boldsymbol{a}_\Lambda)$ and $(\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}\,\boldsymbol{b}_\star)$ are orthogonal to each other. Applying the Pythagorean Theorem to (23), we obtain

$$\|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2 \;+\; \|\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2^2 \;=\; \delta^2$$

whence

$$\|\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2 \;=\; \left[\delta^2 \;-\; \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2\right]^{1/2}. \qquad (25)$$

The coefficient vector $\boldsymbol{b}_\star$ also satisfies the hypotheses of Corollary 5, and therefore we obtain

$$\|\boldsymbol{a}_\Lambda - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2 \;\leq\; \gamma_\star\,\big\|\boldsymbol{\Phi}_\Lambda^\dagger\big\|_{2,1}.$$

Substituting (25) into this relation, we obtain a lower bound on $\gamma_\star$:

$$\left[\delta^2 - \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2\right]^{1/2}\,\big\|\boldsymbol{\Phi}_\Lambda^\dagger\big\|_{2,1}^{-1} \;\leq\; \gamma_\star.$$

Next, we introduce (20) into this inequality and make extensive (but obvious) simplifications to reach

$$\|\boldsymbol{\Phi}^*\,(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma_\star\,\mathrm{ERC}(\Lambda).$$

In view of this fact, the Correlation Condition Lemma allows us to conclude that $\boldsymbol{b}_\star$ is the (unique) global minimizer of the convex program

$$\min_{\boldsymbol{b}\,\in\,\mathbb{C}^\Omega} \;\tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \gamma_\star\,\|\boldsymbol{b}\|_1.$$

Dividing through by $\gamma_\star$, we discover that

$$\boldsymbol{b}_\star \;\in\; \operatorname*{arg\,min}_{\boldsymbol{b}\in\mathbb{C}^\Omega} \;\zeta_\star\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \|\boldsymbol{b}\|_1.$$

Since we also have (23) and $\zeta_\star > 0$, it follows that $\boldsymbol{b}_\star$ satisfies the KKT sufficient conditions for a minimizer of

$$\min_{\boldsymbol{b}\,\in\,\mathbb{C}^\Omega} \;\|\boldsymbol{b}\|_1 \quad\text{subject to}\quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2 \leq \delta. \qquad (\ell_1\text{-ERROR})$$

Therefore, we have identified a global minimizer of ($\ell_1$-ERROR).

Now we must demonstrate that the coefficient vector $\boldsymbol{b}_\star$ provides the *unique* minimizer of the convex program. This requires some work because we have not shown that every minimizer must be supported on $\Lambda$.

Suppose that $\boldsymbol{b}_{\mathrm{alt}}$ is another coefficient vector that solves ($\ell_1$-ERROR). First, we argue that $\boldsymbol{\Phi}\,\boldsymbol{b}_{\mathrm{alt}} = \boldsymbol{\Phi}\,\boldsymbol{b}_\star$ by assuming the contrary. Since $\|\boldsymbol{s}\|_2 > \delta$, the error constraint in ($\ell_1$-ERROR) is binding at every solution. In particular, we must have

$$\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}_{\mathrm{alt}}\|_2 \;=\; \delta \;=\; \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}_\star\|_2.$$

Since $\ell_2$ balls are strictly convex,

$$\big\|\boldsymbol{s} - \tfrac{1}{2}\,\boldsymbol{\Phi}\,(\boldsymbol{b}_{\mathrm{alt}} + \boldsymbol{b}_\star)\big\|_2 \;<\; \delta.$$

It follows that the coefficient vector $\frac{1}{2}(\boldsymbol{b}_{\mathrm{alt}} + \boldsymbol{b}_\star)$ cannot solve ($\ell_1$-ERROR). Yet the solutions to a convex program must form a convex set, which is a contradiction.

Next, observe that $\boldsymbol{b}_{\mathrm{alt}}$ and $\boldsymbol{b}_\star$ share the same $\ell_1$ norm because they both solve ($\ell_1$-ERROR). Under our hypothesis that $\mathrm{ERC}(\Lambda) > 0$, Corollary 11 states that $\boldsymbol{b}_\star$ is the *unique* solution to the program

$$\min_{\boldsymbol{b} \in \mathbb{C}^\Omega} \|\boldsymbol{b}\|_1 \quad \text{subject to} \quad \boldsymbol{\Phi}\,\boldsymbol{b} = \boldsymbol{\Phi}\,\boldsymbol{b}_\star.$$

Thus $\boldsymbol{b}_{\mathrm{alt}} = \boldsymbol{b}_\star$. We conclude that $\boldsymbol{b}_\star$ is the unique minimizer of the convex relaxation.

Finally, let us estimate how far $\boldsymbol{b}_\star$ deviates from $\boldsymbol{c}_\Lambda$. We begin with equation (25), which can be written

$$\|\boldsymbol{\Phi}_\Lambda(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star)\|_2 \;=\; \left[\delta^2 \;-\; \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2\right]^{1/2}.$$

The right-hand side clearly does not exceed $\delta$, while the left-hand side may be bounded below as

$$\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}^{-1} \|\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star\|_2 \quad\leq\quad \|\boldsymbol{\Phi}_\Lambda(\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star)\|_2.$$

Combine the two bounds and rearrange to complete the argument. ∎

## VI. Discussion

This final section attempts to situate the work of this paper in context. To that end, it describes some results for another sparse approximation algorithm, Orthogonal Matching Pursuit, that are parallel with our results for convex relaxation. Then we discuss some recent manuscripts that contain closely related results for convex relaxation. Afterward, some extensions and new directions for research are presented. We conclude with a vision for the future of sparse approximation.

### A. Related Results for OMP

Another basic technique for sparse approximation is a greedy algorithm called Orthogonal Matching Pursuit (OMP) [12, 44]. At each iteration, this method selects the atom most strongly correlated with the residual part of the signal. Then it forms the best approximation of the signal using the atoms already chosen, and it repeats the process until a stopping criterion is met.

The author feels that it is instructive to compare the theoretical performance of convex relaxation against OMP, so this subsection offers two theorems about the behavior of the greedy algorithm. We extract the following result from Theorem 5.3 of [45].

*Theorem 18 (Tropp–Gilbert–Strauss):* Suppose that $\Lambda$ contains no more than $m$ indices, where $m\,\mu < \frac{1}{2}$. Suppose that $\boldsymbol{s}$ is an input signal whose best approximation over $\Lambda$ satisfies the correlation condition

$$\|\boldsymbol{\Phi}^*(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\|_\infty \;\leq\; \gamma\,\frac{1 \;-\; 2\,m\,\mu}{1 \;-\; m\,\mu}.$$

Apply Orthogonal Matching Pursuit to the input signal, and halt the algorithm at the end of iteration $t$ if the maximum correlation between the computed residual $\boldsymbol{r}_t$ and the dictionary satisfies

$$\|\boldsymbol{\Phi}^*\,\boldsymbol{r}_t\|_\infty \;\leq\; \gamma.$$

We may conclude that

- the algorithm has selected $t$ indices from $\Lambda$, and
- it has chosen every index $\lambda$ from $\Lambda$ for which

$$|\boldsymbol{c}_\Lambda(\lambda)| \;>\; \frac{\gamma}{1 \;-\; 2\,m\,\mu}.$$

- Moreover, the absolute error in the computed approximation $\boldsymbol{a}_t$ satisfies

$$\|\boldsymbol{s} - \boldsymbol{a}_t\|_2^2 \;\leq\; \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2^2 \;+\; \gamma^2\,\frac{m - t}{1 \;-\; (m - t)\,\mu}.$$

Note that this theorem relies on the coherence parameter, and no available version describes the behavior of the algorithm in terms of more fundamental dictionary attributes. Compare this result with Corollary 9 for ($\ell_1$-PENALTY).

There is also a result for OMP that parallels Corollary 15 for ($\ell_1$-ERROR). We quote Theorem 5.9 from [33].

*Theorem 19 (Tropp):* Let $\Lambda$ index a linearly independent collection of atoms for which $\mathrm{ERC}(\Lambda) > 0$. Suppose that $\boldsymbol{s}$ is an input signal, and let $\delta$ be a number no smaller than

$$\left[1 \;+\; \left(\frac{\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,1}}{\mathrm{ERC}(\Lambda)}\right)^2\right]^{1/2} \|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2.$$

Suppose that we apply Orthogonal Matching Pursuit to the input signal, and we halt the algorithm at the end of iteration $t$ if the norm of the computed residual $\boldsymbol{r}_t$ satisfies

$$\|\boldsymbol{r}_t\|_2 \;\leq\; \delta.$$

Then we may conclude that the algorithm has chosen $t$ indices from $\Lambda$, and (obviously) the error in the calculated approximation $\boldsymbol{a}_t$ satisfies $\|\boldsymbol{s} - \boldsymbol{a}_t\|_2^2 \;\leq\; \delta$.

It is remarkable how closely results for the greedy algorithm correspond with our theory for convex relaxation. The manuscript [18] contains some related results for OMP.

### B. Comparison with Related Work

As the first draft [25] of the present work was being released, several other researchers independently developed some closely related results [18, 26, 27]. This subsection provides a short guide to these other manuscripts.

In [26], J.-J. Fuchs has studied the qualitative performance of the $\ell_1$ penalty method for the problem of recovering a sparse signal contaminated with zero-mean, white Gaussian noise. Fuchs shows that the minimizer of ($\ell_1$-PENALTY) correctly identifies all of the atoms that participate in the sparse signal, provided that the signal-to-noise ratio is high enough and the parameter $\gamma$ is suitably chosen. He suggests that quantifying the analysis would be difficult. The example in Section IV-D of this paper shows that his pessimism is unwarranted.

Later, Fuchs extended his approach to study the behavior of the $\ell_1$ penalty method for recovering a sparse signal corrupted by an unknown noise vector of bounded $\ell_2$ norm [26]. Theorem 4 of his paper provides a quantitative result in terms of the coherence parameter. The theorem states that solving ($\ell_1$-PENALTY) correctly identifies all of the atoms that participate in the sparse signal, provided that none of the ideal coefficients are too small and that the parameter $\gamma$ is chosen correctly. It is possible to obtain the same result as an application of our Theorem 8.

The third manuscript [18] provides a sprawling analysis of the performance of (19), ($\ell_1$-ERROR), and OMP for the problem of recovering a sparse signal $\mathbf{\Phi}\, \boldsymbol{c}_{\mathrm{opt}}$ polluted by an unknown noise vector whose $\ell_2$ norm does not exceed $\varepsilon$. It is difficult to summarize their work in such a small space, but we will touch on the highlights.

- Their Theorem 2.1 shows that any solution $\boldsymbol{c}_\star$ to (19) with tolerance $\varepsilon$ satisfies
$$\|\boldsymbol{c}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; K_0\, \varepsilon,$$
where the constant $K_0$ depends on the coherence of the dictionary and the sparsity of the ideal signal.
- Theorem 3.1 shows that any solution $\boldsymbol{b}_\star$ to ($\ell_1$-ERROR) with tolerance $\delta$ satisfies
$$\|\boldsymbol{b}_\star - \boldsymbol{c}_{\mathrm{opt}}\|_2 \;\leq\; K_1\, (\varepsilon + \delta)$$
where the constant $K_1$ also depends on the coherence and the sparsity levels.
- Theorem 6.1 is analogous with Corollary 15 of this paper. Their result, however, is much weaker because it does not take into account the direction of the noise, and it estimates everything in terms of the coherence parameter.

The most significant contributions of [18], therefore, are probably their stability results, Theorems 2.1 and 3.1, which are not paralleled elsewhere.

We hope this section indicates how these three manuscripts illuminate different facets of convex relaxation. Along with the present work, they form a strong foundation for the theoretical performance of convex relaxation methods.

### C. Extensions

There are many different ways to generalize the work in this paper. This subsection surveys some of the most promising directions.

*1) The Correlation Condition:* In this article, we have not used the full power of Lemma 6, the Correlation Condition. In particular, we never exploited the subgradient that appears in the first sufficient condition. It is possible to strengthen our results significantly by taking this subgradient into account. Due to the breadth of the present work, we have chosen not to burden the reader with these results.

*2) Beyond Coherence:* Another shortcoming of the present approach is the heavy reliance on the coherence parameter to make estimates. Several recent manuscripts [46–50] have exploited methods from Banach space geometry to study the average-case performance of convex relaxation, and they have obtained results that far outstrip simple coherence bounds. None of these works, however, contains theory about recovering the correct support of a synthetic sparse signal polluted with noise. The author considers this to be one of the most important open questions in sparse approximation.

*3) Other Error Norms:* One can imagine situations where the $\ell_2$ norm is not the most appropriate way to measure the error in approximating the input signal. Indeed, it may be more effective to use the convex program

$$\min_{\boldsymbol{b} \in \mathbb{C}^\Omega} \; \|\boldsymbol{s} - \mathbf{\Phi}\, \boldsymbol{b}\|_p \;+\; \gamma\, \|\boldsymbol{b}\|_1 \tag{26}$$

where $p \in [1, \infty]$. Intuitively, large $p$ will force the approximation to be close to the signal in every component, while small $p$ will make the approximation more tolerant to outliers. Meanwhile, the $\ell_1$ penalty on the coefficients will promote sparsity of the coefficients. The case $p = \infty$ has recently been studied in [51]. See also [52, 53].

To analyze the behavior of (26), we might follow the same path as in Section III. For a given index set $\Lambda$, we would characterize the solution to the convex program when the feasible set is restricted to vectors supported on $\Lambda$. Afterward, we would perturb this solution in a single component outside $\Lambda$ to develop a condition which ensures that the restricted minimizer is indeed the global minimizer. The author has made some progress with this approach for the case $p = \infty$, but no detailed results are presently available.

*4) Bregman Divergences:* Suppose that $f$ is a differentiable, convex function. Associated with $f$ is a directed distance measure called a *Bregman divergence*. The divergence of $\boldsymbol{x}$ from $\boldsymbol{y}$ is defined as

$$D_f(\boldsymbol{x}; \boldsymbol{y}) \;\stackrel{\mathrm{def}}{=}\; f(\boldsymbol{x}) - f(\boldsymbol{y}) - \operatorname{Re} \langle \boldsymbol{x} - \boldsymbol{y},\; \nabla f(\boldsymbol{y}) \rangle\,.$$

The most common examples of Bregman divergences are the squared $\ell_2$ distance and the (generalized) information divergence. Like these prototypes, a Bregman divergence has many attractive properties with respect to best approximation [54]. Moreover, there is a one-to-one correspondence between Bregman divergences and exponential families of probability distributions [55].

In the model problem, suppose that the noise vector $\boldsymbol{\nu}$ is drawn from the exponential family connected with the Bregman divergence $D_f$. To solve the problem, it is natural to apply the convex program

$$\min_{\boldsymbol{b} \in \mathbb{C}^\Omega} \; D_f(\mathbf{\Phi}\, \boldsymbol{b}; \boldsymbol{s}) \;+\; \gamma\, \|\boldsymbol{b}\|_1\,. \tag{27}$$

Remarkably, the results in this paper can be adapted for (27) with very little effort. This coup is possible because Bregman divergences exhibit behavior perfectly analogous with the squared $\ell_2$ norm.

*5) The Elastic Net:* The statistics literature contains several variants of ($\ell_1$-PENALTY). In particular, Zuo and Hastie have studied the following convex relaxation method for subset selection [56].

$$\min_{\boldsymbol{b} \in \mathbb{C}^\Omega} \; \tfrac{1}{2}\, \|\boldsymbol{s} - \mathbf{\Phi}\, \boldsymbol{b}\|_2^2 \;+\; \gamma\, \|\boldsymbol{b}\|_1 \;+\; \beta\, \|\boldsymbol{b}\|_2^2\,. \tag{28}$$

They call this optimization problem *the elastic net*. Roughly speaking, it promotes sparsity without permitting the coefficients to become too large. It would be valuable to study the solutions to (28) using our methods.

*6) Simultaneous Sparse Approximation:* Suppose that we had several different observations of a sparse signal contaminated with additive noise. One imagines that we could use the extra observations to produce a better estimate for the ideal signal. This problem is called *simultaneous sparse approximation*. The papers [45, 57–60] discuss algorithms for approaching this challenging problem.

### D. Conclusions

Even though convex relaxation has been applied for more than thirty years, the present results are unprecedented in the published literature because they apply to the important case where a sparse signal is contaminated by noise. Indeed, we have seen conclusively that convex relaxation can be used to solve the model problem in a variety of concrete situations. We have also shown that convex relaxation offers a viable approach to the subset selection problem and the error-constrained sparse approximation problem. These examples were based on detailed general theory that describes the behavior of minimizers to the optimization problems ($\ell_1$-PENALTY) and ($\ell_1$-ERROR). This theory also demonstrates that the efficacy of convex relaxation is related intimately to the geometric properties of the dictionary. We hope that this report will have a significant impact on the practice of convex relaxation because it proves how these methods will behave in realistic problem settings.

Nevertheless, this discussion section makes it clear that there is an enormous amount of work that remains to be done. In particular, we still do not understand the average-case behavior of convex relaxation in the presence of noise. But as Vergil reminds us, "*Tantae molis erat Romanam condere gentem.*" Such a burden it was to establish the Roman race.

## ACKNOWLEDGMENTS

I wish to thank A. C. Gilbert and M. J. Strauss for their patience and encouragement. The detailed and insightful remarks of the anonymous referees prompted me to see my results *de novo*. I believe that the current work is immeasurably better than the original submission on account of their attention and care. Thanks also to the Austinites, who have told me for years that I should just relax.

## APPENDIX I
## A BRIEF HISTORY OF RELAXATION

The ascendance of convex relaxation for sparse approximation was propelled by two theoretical–technological developments of the last half century. First, the philosophy and methodology of robust statistics—which derive from work of von Neumann, Tukey, and Huber—show that $\ell_1$ loss criteria can be applied to defend statistical estimators against outlying data points. Robust estimators qualitatively prefer a few large errors and many tiny errors to the armada of moderate deviations introduced by mean-squared-error criteria. Second, the elevation during the 1950s of linear programming to the level of *technology* and the interior-point revolution of the 1980s have made it both tractable and commonplace to solve the large-scale optimization problems that arise from convex relaxation.

It appears that a 1973 paper of Claerbout and Muir is the crucible in which these reagents were first combined for the express goal of yielding a sparse representation [15]. They write,

> In deconvolving any observed seismic trace, it is rather disappointing to discover that there is a nonzero spike at every point in time regardless of the data sampling rate. One might hope to find spikes only where real geologic discontinuities take place. Perhaps the $L_1$ norm can be utilized to give a [sparse] output trace....

This idea was subsequently developed in the geophysics literature [5, 61, 62]. Santosa and Symes, in 1986, proposed the convex relaxation ($\ell_1$-PENALTY) as a method for recovering sparse spike trains, and they proved that the method succeeds under moderate restrictions [35].

Around 1990, the work on $\ell_1$ criteria in signal processing recycled to the statistics community. Donoho and Johnstone wrote a pathbreaking paper which proved that one could determine a nearly optimal minimax estimate of a smooth function contaminated with noise by solving ($\ell_1$-PENALTY) with the dictionary $\mathbf{\Phi}$ an appropriate wavelet basis and the parameter $\gamma$ related to the variance of the noise. Slightly later, Tibshirani proposed this convex program, which he calls *the lasso*, as a method for solving the subset selection problems in linear regression [42]. From here, it is only a short step to Basis Pursuit and Basis Pursuit de-noising [9].

This history could not be complete without mention of parallel developments in the theoretical computer sciences. It has long been known that some combinatorial problems are intimately bound up with continuous convex programming problems. In particular, the problem of determining the maximum value that an affine function attains at some vertex of a polytope can be solved using a linear program [63]. A major theme in modern computer science is that many other combinatorial problems can be solved approximately by means of convex relaxation. For example, a celebrated paper of Goemans and Williamson proves that a certain convex program can be used to produce a graph cut whose weight exceeds 87% of the maximum cut [64]. The present work draws deeply on the fundamental idea that a combinatorial problem and its convex relaxation often have closely related solutions.

## APPENDIX II
## IS THE ERC NECESSARY?

Corollary 10 shows that if $\mathrm{ERC}(\Lambda) \geq 0$, then any superposition of atoms from $\Lambda$ can be recovered using ($\ell_1$-PENALTY) for a sufficiently small penalty $\gamma$. The following theorem demonstrates that this type of result cannot hold if $\mathrm{ERC}(\Lambda) < 0$. It follows from results of Feuer and Nemirovsky [65] that there are dictionaries in which some sparse signals cannot be recovered by means of ($\ell_1$-PENALTY).

*Theorem 20:* Suppose that $\mathrm{ERC}(\Lambda) < 0$. Then we may construct an input signal that has an exact representation using

the atoms in $\Lambda$ and yet the minimizer of the function (L) is *not* supported on $\Lambda$ when $\gamma$ is small.

*Proof:* Since $\mathrm{ERC}(\Lambda) < 0$, there must exist an atom $\boldsymbol{\varphi}_\omega$ for which $\left\|\boldsymbol{\Phi}_\Lambda^\dagger \boldsymbol{\varphi}_\omega\right\|_1 > 1$ even though $\omega \notin \Lambda$. Perversely, we select the input signal to be $\boldsymbol{s} = \boldsymbol{P}_\Lambda \boldsymbol{\varphi}_\omega$. We can synthesize $\boldsymbol{s}$ exactly with the atoms in $\Lambda$ by means of the coefficient vector $\boldsymbol{c}_\Lambda = \boldsymbol{\Phi}_\Lambda^\dagger \boldsymbol{\varphi}_\omega$.

According to Corollary 5, the minimizer $\boldsymbol{b}_\star$ of the function

$$L(\boldsymbol{b}) \;=\; \tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{b}\|_2^2 \;+\; \gamma\,\|\boldsymbol{b}\|_1$$

over all coefficient vectors supported on $\Lambda$ must satisfy $\|\boldsymbol{c}_\Lambda - \boldsymbol{b}_\star\|_\infty \;\le\; \gamma\,\left\|(\boldsymbol{\Phi}_\Lambda^*\boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty}$. Since $\|\boldsymbol{c}_\Lambda\|_1 > 1$ by construction, we may choose $\gamma$ small enough that the bound $\|\boldsymbol{b}_\star\|_1 > 1$ is also in force. Define the corresponding approximation $\boldsymbol{a}_\star = \boldsymbol{\Phi}\,\boldsymbol{b}_\star$.

Now we construct a parameterized coefficient vector

$$\boldsymbol{b}(t) \;\stackrel{\mathrm{def}}{=}\; (1-t)\,\boldsymbol{b}_\star \;+\; t\,\mathbf{e}_\omega \qquad \text{for } t \text{ in } [0,1].$$

For positive $t$, it is clear that the support of $\boldsymbol{b}(t)$ is not contained in $\Lambda$. We will prove that $L(\boldsymbol{b}(t)) < L(\boldsymbol{b}_\star)$ for small, positive $t$. Since $\boldsymbol{b}_\star$ minimizes $L$ over all coefficient vectors supported on $\Lambda$, no global minimizer of $L$ can be supported on $\Lambda$.

To proceed, calculate that

$$\begin{aligned} L(\boldsymbol{b}(t)) \;&=\; \tfrac{1}{2}\,\|\boldsymbol{s} - \boldsymbol{a}_\star\|_2^2 \\ &\quad + t\,\operatorname{Re}\langle \boldsymbol{s} - \boldsymbol{a}_\star,\; \boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega\rangle + \tfrac{1}{2}t^2\,\|\boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega\|_2^2 \\ &\qquad + \gamma\,(1-t)\,\|\boldsymbol{b}_\star\|_1 + t\gamma. \end{aligned}$$

Differentiate this expression with respect to $t$ and evaluate the derivative at $t = 0$:

$$\left.\frac{\mathrm{d}L(\boldsymbol{b}(t))}{\mathrm{d}t}\right|_{t=0} \;=\; \operatorname{Re}\langle \boldsymbol{s} - \boldsymbol{a}_\star,\; \boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega\rangle + \gamma\,(1 - \|\boldsymbol{b}_\star\|_1).$$

By construction of $\boldsymbol{b}_\star$, the second term is negative. The first term is nonpositive because

$$\begin{aligned} \langle \boldsymbol{s} - \boldsymbol{a}_\star,\; \boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega\rangle \;&=\; \langle \boldsymbol{P}_\Lambda(\boldsymbol{s} - \boldsymbol{a}_\star),\; \boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega\rangle \\ &=\; \langle \boldsymbol{s} - \boldsymbol{a}_\star,\; \boldsymbol{P}_\Lambda(\boldsymbol{a}_\star - \boldsymbol{\varphi}_\omega)\rangle \\ &=\; \langle \boldsymbol{s} - \boldsymbol{a}_\star,\; \boldsymbol{a}_\star - \boldsymbol{s}\rangle \\ &=\; -\,\|\boldsymbol{s} - \boldsymbol{a}_\star\|_2^2. \end{aligned}$$

Therefore, the derivative is negative, and $L(\boldsymbol{b}(t)) < L(\boldsymbol{b}(0))$ for small, positive $t$. Observe that $\boldsymbol{b}(0) = \boldsymbol{b}_\star$ to complete the argument. ∎

## Appendix III
## Coherence Bounds

In this section, we collect some coherence bounds from Section 3 of the technical report [25]. These results help us to understand the behavior of sub-dictionary synthesis matrices. We begin with a bound for the singular values.

*Proposition 21:* Suppose that $(m-1)\,\mu < 1$, and assume that $|\Lambda| \le m$. Each singular value $\sigma$ of the matrix $\boldsymbol{\Phi}_\Lambda$ satisfies

$$1 \;-\; (m-1)\,\mu \quad\le\quad \sigma^2 \quad\le\quad 1 \;+\; (m-1)\,\mu.$$

The last result yields a useful new estimate.

*Proposition 22:* Suppose that $(m-1)\,\mu < 1$, and assume that $|\Lambda| \;\le\; m$. Then $\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,\infty} \;\le\; [1 - (m-1)\,\mu]^{-1/2}$. Equivalently, the rows $\boldsymbol{\Phi}_\Lambda^\dagger$ have $\ell_2$ norms no greater than $[1-(m-1)\,\mu]^{-1/2}$.

*Proof:* Recall that the operator norm $\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,\infty}$ calculates the maximum $\ell_2$ norm among the rows of $\boldsymbol{\Phi}_\Lambda^\dagger$. It is easy to see that this operator norm is bounded above by $\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2}$, the maximum singular value of $\boldsymbol{\Phi}_\Lambda^\dagger$. But the maximum singular value of $\boldsymbol{\Phi}_\Lambda^\dagger$ is the reciprocal of the minimum singular value of $\boldsymbol{\Phi}_\Lambda$. Proposition 21 provides a lower bound on this singular value, which completes the argument. ∎

We also require another norm estimate on the pseudo-inverse.

*Proposition 23:* Suppose that $|\Lambda| \le m$ and $(m-1)\,\mu < 1$. Then

$$\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,1} \;\le\; \left[\frac{m}{1 \;-\; (m-1)\,\mu}\right]^{1/2}.$$

The next result bounds an operator norm of the Gram matrix.

*Proposition 24:* Suppose that $|\Lambda| \le m$ and $(m-1)\,\mu < 1$. Then

$$\left\|(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty} \quad\le\quad \frac{1}{1 \;-\; (m-1)\,\mu}.$$

We conclude by estimating the Exact Recovery Coefficient.

*Proposition 25:* Suppose that $|\Lambda| \le m$, where $m\,\mu \le \frac{1}{2}$. A lower bound on the Exact Recovery Coefficient is

$$\mathrm{ERC}(\Lambda) \quad\ge\quad \frac{1 \;-\; (2m-1)\,\mu}{1 \;-\; (m-1)\,\mu}.$$

## Appendix IV
## Details of Channel Coding Examples

This appendix contains the gory details behind the channel coding examples from Sections IV-D and V-C. In preparation, let us recall the model for ideal signals. We work with a $\mu$-coherent dictionary $\boldsymbol{\Phi}$ containing $N$ atoms in a $d$-dimensional real signal space. For each set $\Lambda$ that lists $m$ atoms, we consider coefficient vectors $\boldsymbol{c}_{\mathrm{opt}}$ supported on $\Lambda$ whose nonzero entries equal $\pm 1$. Each ideal signal has the form $\boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$.

It will be handy to have some numerical bounds on the quantities that arise. Using the estimates in Appendix III, we obtain the following bounds.

*Proposition 26:* Suppose that $\Lambda$ indexes a sub-dictionary containing $m$ atoms, where $m\mu \le \frac{1}{3}$. Then the following bounds are in force:

- $\left\|\boldsymbol{\Phi}_\Lambda^\dagger\right\|_{2,2} \le \sqrt{3/2}$.
- $\left\|(\boldsymbol{\Phi}_\Lambda^*\,\boldsymbol{\Phi}_\Lambda)^{-1}\right\|_{\infty,\infty} \le 3/2$.
- $\mathrm{ERC}(\Lambda) \ge 1/2$.
- The rows of $\boldsymbol{\Phi}_\Lambda^\dagger$ have $\ell_2$ norms no greater than $\sqrt{3/2}$.

### A. The Gaussian Channel

In this example, we attempt to recover an ideal signal from a measurement of the form

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\nu}$$

where $\boldsymbol{\nu}$ is a zero-mean Gaussian random vector with covariance matrix $\sigma^2\,\mathbf{I}$. We study the performance of ($\ell_1$-PENALTY) for this problem.

First, consider the event that Corollary 9 applies. In other words, it is necessary to compute the probability that the Correlation Condition is in force, i.e.,

$$\|\boldsymbol{\Phi}^*\,(\boldsymbol{s}-\boldsymbol{a}_\Lambda)\|_\infty \;\le\; \gamma\;\mathrm{ERC}(\Lambda)$$

where $\Lambda = \mathrm{supp}(\boldsymbol{c}_{\mathrm{opt}})$. To that end, we will develop a lower bound on the probability

$$\begin{aligned} E_1 \;&\stackrel{\mathrm{def}}{=}\; \mathrm{Prob}\left\{\|\boldsymbol{\Phi}^*\,(\boldsymbol{s}-\boldsymbol{a}_\Lambda)\|_\infty \;\le\; \tfrac{1}{2}\,\gamma\right\} \\ &=\; \mathrm{Prob}\left\{\max\nolimits_{\omega\notin\Lambda}|\langle \boldsymbol{s}-\boldsymbol{a}_\Lambda,\;\boldsymbol{\varphi}_\omega\rangle| \;\le\; \tfrac{1}{2}\,\gamma\right\}. \end{aligned}$$

Observe that $(\boldsymbol{s}-\boldsymbol{a}_\Lambda) \;=\; (\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{s} \;=\; (\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}$. Use adjointness to transfer the projector to the other side of the inner product, and scale $\boldsymbol{\nu}$ by $1/\sigma$ to obtain

$$E_1 \;=\; \mathrm{Prob}\left\{\max\nolimits_{\omega\notin\Lambda}|\langle \boldsymbol{z},\;(\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle| \;\le\; \tfrac{1}{2}\,\gamma/\sigma\right\}$$

where $\boldsymbol{z}$ is a zero-mean Gaussian vector with identity covariance. A direct application of Sidak's Lemma [66, Lemma 2] yields a lower bound on this probability.

$$E_1 \;\ge\; \prod\nolimits_{\omega\notin\Lambda}\mathrm{Prob}\left\{|\langle \boldsymbol{z},\;(\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle| \;\le\; \tfrac{1}{2}\,\gamma/\sigma\right\}.$$

The atoms have unit norm, so the norm of $(\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega$ does not exceed one. Therefore,

$$E_1 \;\ge\; \prod\nolimits_{\omega\notin\Lambda}\mathrm{Prob}\left\{\frac{|\langle \boldsymbol{z},\;(\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle|}{\|(\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\|_2} \;\le\; \tfrac{1}{2}\,\gamma/\sigma\right\}.$$

Suppose that $\boldsymbol{v}$ is an arbitrary unit vector. Integrating the Gaussian kernel over the hyperplane $\boldsymbol{v}^\perp$ yields the bound

$$\begin{aligned} \mathrm{Prob}\left\{|\langle \boldsymbol{z},\;\boldsymbol{v}\rangle| \;\le\; t\right\} \;&=\; \frac{1}{\sqrt{2\pi}}\int_{-t}^{t}\mathrm{e}^{-x^2/2}\,\mathrm{d}x \\ &\ge\; 1-\exp\{-t^2/2\}. \end{aligned} \qquad (29)$$

Apply (29) in our bound for $E_1$, and exploit the fact that the product contains $(N-m)$ identical terms to reach

$$E_1 \;\ge\; \left[1 \;-\; \exp\left\{-\tfrac{1}{8}\,(\gamma/\sigma)^2\right\}\right]^{N-m}.$$

In words, we have developed a lower bound on the probability that the correlation condition holds so that we may apply Corollary 9.

Assume that the conclusions of Corollary 9 are in force. In particular, $\mathrm{supp}(\boldsymbol{b}_\star)\subset\Lambda$ and also $\|\boldsymbol{b}_\star-\boldsymbol{c}_\Lambda\|_\infty \le \tfrac{3}{2}\,\gamma$. The upper triangle inequality implies that

$$\|\boldsymbol{b}_\star-\boldsymbol{c}_{\mathrm{opt}}\|_\infty \;\le\; \tfrac{3}{2}\,\gamma \;+\; \|\boldsymbol{c}_\Lambda-\boldsymbol{c}_{\mathrm{opt}}\|_\infty.$$

Each nonzero component of $\boldsymbol{c}_{\mathrm{opt}}$ equals $\pm1$, so the corresponding component of $\boldsymbol{b}_\star$ has the same sign provided that

$$\|\boldsymbol{c}_\Lambda-\boldsymbol{c}_{\mathrm{opt}}\|_\infty \;<\; 1-\tfrac{3}{2}\,\gamma.$$

We will bound the probability that this event occurs. By definition, $\boldsymbol{c}_\Lambda = \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{s}$, which yields

$$\boldsymbol{c}_\Lambda \;=\; \boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\nu}.$$

It follows that the noise contaminating $\boldsymbol{c}_\Lambda$ depends only on $\boldsymbol{P}_\Lambda\,\boldsymbol{\nu}$, which is independent from $(\boldsymbol{s}-\boldsymbol{a}_\Lambda) = (\mathbf{I}-\boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}$. We calculate that

$$\begin{aligned} E_2 \;&\stackrel{\mathrm{def}}{=}\; \mathrm{Prob}\left\{\|\boldsymbol{c}_{\mathrm{opt}}-\boldsymbol{c}_\Lambda\|_\infty \;<\; 1-\tfrac{3}{2}\,\gamma\right\} \\ &=\; \mathrm{Prob}\left\{\left\|\boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{\nu}\right\|_\infty \;\le\; 1-\tfrac{3}{2}\,\gamma\right\} \\ &\ge\; \mathrm{Prob}\left\{\left\|\boldsymbol{\Phi}_\Lambda^\dagger\,\boldsymbol{z}\right\|_\infty \;\le\; (1-\tfrac{3}{2}\,\gamma)/\sigma\right\} \\ &=\; \mathrm{Prob}\left\{\max_{\lambda\in\Lambda}\left|\langle \boldsymbol{z},\;(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\mathbf{e}_\lambda\rangle\right| \;\le\; (1-\tfrac{3}{2}\,\gamma)/\sigma\right\} \end{aligned}$$

where $\boldsymbol{z}$ is a Gaussian random vector with identity covariance. Applying Sidak's Lemma again, we obtain

$$E_2 \;\ge\; \prod\nolimits_{\lambda\in\Lambda}\mathrm{Prob}\left\{\left|\langle \boldsymbol{z},\;(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\mathbf{e}_\lambda\rangle\right| \;\le\; (1-\tfrac{3}{2}\,\gamma)/\sigma\right\}.$$

The $\ell_2$ norm of each row of $\boldsymbol{\Phi}_\Lambda^\dagger$ is bounded above by $\sqrt{3/2}$. Renormalizing each row and applying (29) gives

$$\begin{aligned} E_2 \;&\ge\; \prod_{\lambda\in\Lambda}\mathrm{Prob}\left\{\frac{\left|\langle \boldsymbol{z},\;(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\mathbf{e}_\lambda\rangle\right|}{\left\|(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\mathbf{e}_\lambda\right\|_2} \;\le\; \frac{(1-\frac{3}{2}\,\gamma)/\sigma}{\left\|(\boldsymbol{\Phi}_\Lambda^\dagger)^*\,\mathbf{e}_\lambda\right\|_2}\right\} \\ &\ge\; \prod\nolimits_{\lambda\in\Lambda}\mathrm{Prob}\left\{|\langle \boldsymbol{z},\;\boldsymbol{v}\rangle| \;\le\; \sqrt{2/3}\,(1-\tfrac{3}{2}\,\gamma)/\sigma\right\} \\ &=\; \left[1 \;-\; \exp\left\{-\tfrac{1}{3}\,\left(1-\tfrac{3}{2}\,\gamma\right)^2\,\sigma^{-2}\right\}\right]^m \end{aligned}$$

where $\boldsymbol{v}$ denotes any fixed unit vector.

The event that we may apply the corollary and the event that the coefficients in $\boldsymbol{c}_\Lambda$ are sufficiently large are independent from each other. We may conclude that the probability of both events occurring is $E_{\mathrm{succ}} = E_1\,E_2$. Ergo,

$$\begin{aligned} E_{\mathrm{succ}} \;\ge\; &\left[1 \;-\; \exp\left\{-\tfrac{1}{8}\,\gamma^2\,\sigma^{-2}\right\}\right]^{N-m} \\ &\left[1 \;-\; \exp\left\{-\tfrac{1}{3}\,\left(1-\tfrac{3}{2}\,\gamma\right)^2\,\sigma^{-2}\right\}\right]^m. \end{aligned}$$

Using the inequality $(1-x)^k \ge 1-kx$, which holds for $x\le1$ and $k\ge1$, we see that the failure probability satisfies

$$\begin{aligned} 1-E_{\mathrm{succ}} \;\le\; &(N-m)\,\exp\left\{-\tfrac{1}{8}\,\gamma^2\,\sigma^{-2}\right\} \\ &+\; m\,\exp\left\{-\tfrac{1}{3}\,\left(1-\tfrac{3}{2}\,\gamma\right)^2\,\sigma^{-2}\right\}. \end{aligned}$$

In particular, the failure probability decays exponentially as the noise power approaches zero. If the noise level is known, it is possible to optimize this expression over $\gamma$ to find the parameter that minimizes the probability of failure.

### B. The Uniform Channel

In this example, we attempt to recover ideal signals from measurements of the form

$$\boldsymbol{s} \;=\; \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}} \;+\; \boldsymbol{\nu}$$

where $\boldsymbol{\nu}$ is uniformly distributed on the $\ell_2$ ball of radius $\varepsilon$. We study the performance of ($\ell_1$-ERROR) for this problem.

Let $\Lambda = \mathrm{supp}(\boldsymbol{c}_{\mathrm{opt}})$. First, observe that the residual vector can be rewritten as

$$\boldsymbol{s} - \boldsymbol{a}_\Lambda \;=\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{s} \;=\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu},$$

so the $\ell_2$ norm of the residual does not exceed $\varepsilon$. According to the remarks after Corollary 15, we should select $\delta$ no smaller than

$$\delta \;\geq\; \left[1 + 6\,m\,\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)^2\right]^{1/2}\,\varepsilon.$$

The rest of the example will prove that the maximum correlation exhibits a concentration phenomenon: it is rather small with high probability. Therefore, we can choose $\delta$ much closer to $\varepsilon$ than a naïve analysis would suggest.

To begin the calculation, rewrite

$$\begin{aligned}\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda) \;&=\; \max_{\omega \notin \Lambda} \frac{|\langle \boldsymbol{s} - \boldsymbol{a}_\Lambda,\; \boldsymbol{\varphi}_\omega\rangle|}{\|\boldsymbol{s} - \boldsymbol{a}_\Lambda\|_2} \\ &=\; \max_{\omega \notin \Lambda} \frac{|\langle (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu},\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle|}{\|(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}\|_2}.\end{aligned}$$

It is not hard to check that $(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}$ is a spherically symmetric random variable on the orthogonal complement of $\mathscr{R}(\boldsymbol{P}_\Lambda)$. Therefore, the random variable

$$\boldsymbol{u} \;\stackrel{\mathrm{def}}{=}\; \frac{(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}}{\|(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\nu}\|_2}$$

is uniformly distributed on the unit sphere of $\mathscr{R}(\boldsymbol{P}_\Lambda)^\perp$, which is a $(d-m)$-dimensional subspace. In consequence, $\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)$ has the same distribution as the random variable

$$\max_{\omega \notin \Lambda} \,|\langle \boldsymbol{u},\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle|\,.$$

Now, let us bound the probability that the maximum correlation is larger than some number $t$. In symbols, we seek

$$\begin{aligned}E \;&\stackrel{\mathrm{def}}{=}\; \mathrm{Prob}\,\{\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda) \;\geq\; t\} \\ &=\; \mathrm{Prob}\,\{\max\nolimits_{\omega \notin \Lambda} \,|\langle \boldsymbol{u},\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle| \;\geq\; t\} \\ &\leq\; \sum\nolimits_{\omega \notin \Lambda} \mathrm{Prob}\,\{|\langle \boldsymbol{u},\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle| \;\geq\; t\}\,.\end{aligned}$$

Fix an index $\omega \notin \Lambda$, and consider

$$\mathrm{Prob}\,\{|\langle \boldsymbol{u},\; (\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega\rangle| \;\geq\; t\}\,.$$

Since $\boldsymbol{u}$ is spherically symmetric, this probability cannot depend on the direction of the vector $(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega$. Moreover, the probability only increases when we increase the length of this vector. It follows that we may replace $(\mathbf{I} - \boldsymbol{P}_\Lambda)\,\boldsymbol{\varphi}_\omega$ by an arbitrary unit vector $\boldsymbol{v}$ from $\mathscr{R}(\boldsymbol{P}_\Lambda)^\perp$. Therefore,

$$\begin{aligned}E \;&\leq\; \sum\nolimits_{\omega \notin \Lambda} \mathrm{Prob}\,\{|\langle \boldsymbol{u},\; \boldsymbol{v}\rangle| \;\geq\; t\} \\ &=\; (N - m)\;\mathrm{Prob}\,\{|\langle \boldsymbol{u},\; \boldsymbol{v}\rangle| \;\geq\; t\}\,.\end{aligned}$$

The probability can be interpreted geometrically as the fraction of the sphere covered by a pair of spherical caps. Since we are working in a subspace of dimension $(d-m)$, each of the two caps covers no more than $\exp\{-\frac{1}{2}\,t^2\,(d-m)\}$ of the sphere [67, Lemma 2.2]. Therefore,

$$E \;\leq\; 2\,(N - m)\;\exp\{-\tfrac{1}{2}\,t^2\,(d-m)\}.$$

Since we can write $E = \mathrm{Prob}\left\{\mathrm{maxcor}(\boldsymbol{s} - \boldsymbol{a}_\Lambda)^2 \;\geq\; t^2\right\}$, we conclude that the choice

$$\delta \;\geq\; \left[1 + 6\,m\,t^2\right]^{1/2}\,\varepsilon$$

allows us to invoke Theorem 14 with probability greater than $(1 - E)$ over the noise.

## Appendix V
## Solutions to (subset)

In this appendix, we establish Theorem 12, a structural result for solutions to the subset selection problem. Recall that the problem is

$$\min_{\boldsymbol{c} \in \mathbb{C}^\Omega} \quad \|\boldsymbol{s} - \boldsymbol{\Phi}\,\boldsymbol{c}\|_2^2 \;+\; \tau^2\;\|\boldsymbol{c}\|_0\,. \qquad \text{(SUBSET)}$$

Now we restate the theorem and prove it.

*Theorem 27:* Fix an input signal $\boldsymbol{s}$, and choose a threshold $\tau$. Suppose that the coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$ solves the subset selection problem, and set $\boldsymbol{a}_{\mathrm{opt}} = \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$.

- For $\lambda \in \mathrm{supp}(\boldsymbol{c}_{\mathrm{opt}})$, we have $|\boldsymbol{c}_{\mathrm{opt}}(\lambda)| \geq \tau$.
- For $\omega \notin \mathrm{supp}(\boldsymbol{c}_{\mathrm{opt}})$, we have $|\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\; \boldsymbol{\varphi}_\omega\rangle| \leq \tau$.

*Proof:* For a given input signal $\boldsymbol{s}$, suppose that the coefficient vector $\boldsymbol{c}_{\mathrm{opt}}$ is a solution of (SUBSET) with threshold $\tau$, and define $\Lambda_{\mathrm{opt}} = \mathrm{supp}(\boldsymbol{c}_{\mathrm{opt}})$. Put $\boldsymbol{a}_{\mathrm{opt}} = \boldsymbol{\Phi}\,\boldsymbol{c}_{\mathrm{opt}}$, and let $\boldsymbol{P}_{\mathrm{opt}}$ denote the orthogonal projector onto the span of the atoms in $\Lambda_{\mathrm{opt}}$. A quick inspection of the objective function makes it clear that $\boldsymbol{a}_{\mathrm{opt}}$ must be the best $\ell_2$ approximation of $\boldsymbol{s}$ using the atoms in $\Lambda_{\mathrm{opt}}$. Therefore, $\boldsymbol{a}_{\mathrm{opt}} = \boldsymbol{P}_{\mathrm{opt}}\,\boldsymbol{s}$.

We begin with the second conclusion of the theorem. Take any index $\omega$ outside $\Lambda_{\mathrm{opt}}$. The best approximation $\boldsymbol{a}'$ of $\boldsymbol{s}$ using the atoms in $\Lambda_{\mathrm{opt}} \cup \{\omega\}$ is

$$\boldsymbol{a}' \;\stackrel{\mathrm{def}}{=}\; \boldsymbol{a}_{\mathrm{opt}} \;+\; \frac{\langle \boldsymbol{s},\; (\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\rangle}{\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2}\,\frac{(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega}{\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2}.$$

Orthogonal projectors are self-adjoint, so

$$\boldsymbol{a}' \;=\; \boldsymbol{a}_{\mathrm{opt}} \;+\; \frac{\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\; \boldsymbol{\varphi}_\omega\rangle}{\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2^2}\,(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega.$$

Since the two terms are orthogonal, the Pythagorean Theorem furnishes

$$\|\boldsymbol{s} - \boldsymbol{a}'\|_2^2 \;=\; \|\boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}}\|_2^2 \;-\; \frac{|\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\; \boldsymbol{\varphi}_\omega\rangle|^2}{\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2^2}.$$

The second term on the right-hand side measures how much the squared approximation error diminishes if we add $\boldsymbol{\varphi}_\omega$ to the approximation. Notice that the second term must be less than or equal to $\tau^2$, or else we could immediately construct a solution to (SUBSET) that is strictly better than $\boldsymbol{c}_{\mathrm{opt}}$ by using the additional atom. Therefore,

$$|\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\; \boldsymbol{\varphi}_\omega\rangle|^2 \;\leq\; \tau^2\;\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2^2$$

Atoms have unit $\ell_2$ norm, and the projector only attenuates their $\ell_2$ norms. It follows that $\|(\mathbf{I} - \boldsymbol{P}_{\mathrm{opt}})\,\boldsymbol{\varphi}_\omega\|_2^2 \leq 1$, and so $|\langle \boldsymbol{s} - \boldsymbol{a}_{\mathrm{opt}},\; \boldsymbol{\varphi}_\omega\rangle| \leq \tau$.

The argument behind the first conclusion of the theorem is similar. Choose an index $\xi$ from $\Lambda_{\mathrm{opt}}$, and let $\boldsymbol{P}$ denote the

orthogonal projector onto the span of the atoms in $\Lambda_{\text{opt}} \setminus \{\xi\}$. Since $\boldsymbol{a}_{\text{opt}} = \sum_{\Lambda} \boldsymbol{c}_{\text{opt}}(\lambda)\, \boldsymbol{\varphi}_{\lambda}$, the best approximation of $\boldsymbol{s}$ using the reduced set of indices is given by

$$\begin{aligned} \boldsymbol{P}\,\boldsymbol{s} \;&=\; \boldsymbol{P}\,\boldsymbol{a}_{\text{opt}} \;=\; \boldsymbol{a}_{\text{opt}} \;-\; (\mathbf{I} - \boldsymbol{P})\,\boldsymbol{a}_{\text{opt}} \\ &=\; \boldsymbol{a}_{\text{opt}} \;-\; \boldsymbol{c}_{\text{opt}}(\xi)\,(\mathbf{I} - \boldsymbol{P})\,\boldsymbol{\varphi}_{\xi} \end{aligned}$$

Thus, removing the atom $\boldsymbol{\varphi}_{\xi}$ from the approximation would increase the squared error by exactly

$$|\boldsymbol{c}_{\text{opt}}(\xi)|^2 \;\|(\mathbf{I} - \boldsymbol{P})\,\boldsymbol{\varphi}_{\xi}\|_2^2 \,.$$

This quantity must be at least $\tau^2$, or else the reduced set of atoms would afford a strictly better solution to (SUBSET) than the original set. Since $(\mathbf{I} - \boldsymbol{P})$ is an orthogonal projector, $\|(\mathbf{I} - \boldsymbol{P})\,\boldsymbol{\varphi}_{\xi}\|_2^2 \leq 1$. We conclude that $|\boldsymbol{c}_{\text{opt}}(\xi)| \geq \tau$. ∎

One could obviously prove much more about the solutions of the subset selection problem using similar techniques, but these results are too tangential to pursue here.

## REFERENCES


[1] R. Gribonval and E. Bacry, "Harmonic decomposition of audio signals with Matching Pursuit," *IEEE Trans. Signal Processing*, vol. 51, no. 1, pp. 101–111, 2003.

[2] P. Frossard, P. Vandergheynst, R. M. F. i Ventura, and M. Kunt, "A posteriori quantization of progressive Matching Pursuit streams," *IEEE Trans. Signal Processing*, vol. 52, no. 2, pp. 525–535, Feb. 2004.

[3] T. Nguyen and A. Zakhor, "Matching Pursuits based multiple description video coding for lossy environments," in *Proc. 2003 IEEE International Conference on Image Processing*, Barcelona, 2003.

[4] A. J. Miller, *Subset Selection in Regression*, 2nd ed. London: Chapman and Hall, 2002.

[5] H. L. Taylor, S. C. Banks, and J. F. McCoy, "Deconvolution with the $\ell_1$ norm," *Geophysics*, vol. 44, no. 1, pp. 39–52, 1979.

[6] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, pp. 465–471, 1979.

[7] M. Grote and T. Huckle, "Parallel preconditioning with sparse approximate inverses," *SIAM J. Sci. Comput.*, vol. 18, no. 3, pp. 838–853, 1997.

[8] F. Girosi, "An equivalence between sparse approximation and Support Vector Machines," *Neural Comput.*, vol. 10, no. 6, pp. 1455–1480, 1998.

[9] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by Basis Pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1999.

[10] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Comm. Pure Appl. Math.*, vol. 57, pp. 1413–1457, 2004.

[11] B. D. Rao and K. Kreutz-Delgado, "An affine scaling methodology for best basis selection," *IEEE Trans. Signal Processing*, vol. 47, no. 1, pp. 187–200, 1999.

[12] G. Davis, S. Mallat, and M. Avellaneda, "Greedy adaptive approximation," *J. Constr. Approx.*, vol. 13, pp. 57–98, 1997.

[13] V. Temlyakov, "Nonlinear methods of approximation," *Foundations of Comput. Math.*, July 2002.

[14] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[15] J. F. Claerbout and F. Muir, "Robust modeling of erratic data," *Geophysics*, vol. 38, no. 5, pp. 826–844, October 1973.

[16] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, pp. 227–234, 1995.

[17] S. Boyd and L. Vanderberghe, *Convex Optimization*. Cambridge Univ. Press, 2004.

[18] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," February 2004, accepted to *IEEE Trans. Inform. Theory*.

[19] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inform. Theory*, vol. 47, pp. 2845–2862, Nov. 2001.

[20] M. Elad and A. M. Bruckstein, "A generalized uncertainty principle and sparse representation in pairs of bases," *IEEE Trans. Inform. Theory*, vol. 48, no. 9, pp. 2558–2567, 2002.

[21] D. L. Donoho and M. Elad, "Maximal sparsity representation via $\ell_1$ minimization," *Proc. Natl. Acad. Sci.*, vol. 100, pp. 2197–2202, March 2003.

[22] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inform. Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.

[23] J.-J. Fuchs, "On sparse representations in arbitrary redundant bases," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1341–1344, June 2004.

[24] C. Dossal and S. Mallat, "Sparse spike deconvolution with minimum scale," in *Proc. Signal Processing with Adaptive Sparse Structured Representations 2005*. Rennes: INRIA, Nov. 2005, pp. 123–126.

[25] J. A. Tropp, "Just relax: Convex programming methods for subset selection and sparse approximation," The University of Texas at Austin, ICES Report 04-04, 2004.

[26] J.-J. Fuchs, "Recovery of exact sparse representations in the presence of noise," in *Proc. 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Montréal, May 2004.

[27] ——, "Recovery of exact sparse representations in the presence of bounded noise," *IEEE Trans. Inform. Theory*, vol. 51, no. 10, pp. 3601–3608, Oct. 2005.

[28] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley and Sons, 1991.

[29] A. C. Gilbert and J. A. Tropp, "Applications of sparse approximation in communications," in *Proc. 2005 IEEE International Symposium on Information Theory*, Adelaide, Sept. 2005, pp. 1000–1004.

[30] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge Univ. Press, 1985.

[31] E. Kreyszig, *Introductory Functional Analysis with Applications*. John Wiley & Sons, 1989.

[32] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Johns Hopkins Univ. Press, 1996.

[33] J. A. Tropp, "Topics in sparse approximation," Ph.D. dissertation, Computational and Applied Mathematics, The University of Texas at Austin, August 2004.

[34] R. T. Rockafellar, *Convex Analysis*. Princeton Univ. Press, 1970.

[35] F. Santosa and W. W. Symes, "Linear inversion of band-limited reflection seismograms," *SIAM J. Sci. Stat. Comput.*, vol. 7, no. 4, pp. 1307–1330, 1986.

[36] J.-J. Fuchs, "Extension of the Pisarenko Method to sparse linear arrays," *IEEE Trans. Signal Processing*, vol. 45, pp. 2413–2421, Oct. 1997.

[37] ——, "On the application of the global matched filter to DOA estimation with uniform circular arrays," *IEEE Trans. Signal Processing*, vol. 49, no. 4, pp. 702–709, Apr. 2001.

[38] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, pp. 607–609, 1996.

[39] D. L. Donoho and I. M. Johnstone, "Minimax estimation via wavelet shrinkage," Stanford Univ.," Statistics Dept. Technical Report, 1992.

[40] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. London: Academic Press, 1999.

[41] J. A. Tropp, "Recovery of short, complex linear combinations via $\ell_1$ minimization," *IEEE Trans. Inform. Theory*, vol. 51, no. 4, April 2005, to appear.

[42] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal Statist. Soc.*, vol. 58, 1996.

[43] R. Gribonval, R. F. i Venturas, and P. Vandergheynst, "A simple test to check the optimality of a sparse signal approximation," Université de Rennes I, Tech. Rep. 1661, Nov. 2004.

[44] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal Matching Pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Annual Asilomar Conference on Signals, Systems and Computers*, Nov. 1993.

[45] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit," *EURASIP J. Signal Processing*, Nov. 2004, to appear.

[46] E. Candès, J. Romberg, and T. Tao, "Exact signal reconstruction from highly incomplete frequency information," June 2004, submitted for publication.

[47] E. J. Candès and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," Nov. 2004, submitted for publication.

[48] E. J. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" Nov. 2004, submitted for publication.

[49] D. Donoho, "For most large underdetermined systems of linear equations, the mimimal $\ell_1$-norm solution is also the sparsest solution," Sept. 2004, unpublished manuscript.

[50] ——, "For most large underdetermined systems of linear equations, the mimimal $\ell_1$-norm solution approximates the sparsest near-solution," Sept. 2004, unpublished manuscript.
[51] D. L. Donoho and M. Elad, "On the stability of Basis Pursuit in the presence of noise," Nov. 2004, unpublished manuscript.
[52] J.-J. Fuchs, "Some further results on the recovery algorithms," in *Proc. Signal Processing with Adaptive Sparse Structured Representations 2005*. Rennes: INRIA, Nov. 2005, pp. 87–90.
[53] L. Granai and P. Vandergheynst, "Sparse approximation by linear programming using an $l_1$ data-fidelity term," in *Proc. Signal Processing with Adaptive Sparse Structured Representations 2005*. Rennes: INRIA, Nov. 2005, pp. 135–138.
[54] H. H. Bauschke and J. M. Borewein, "Legendre functions and the method of random Bregman projections," *J. Convex Anal.*, vol. 4, no. 1, pp. 27–67, 1997.
[55] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," in *Proc. 4th SIAM International Conference on Data Mining*. SIAM, Apr. 2004.
[56] H. Zuo and T. Hastie, "Regularization and variable selection via the elastic net," Sept. 2004, unpublished manuscript.
[57] D. Leviatan and V. N. Temlyakov, "Simultaneous approximation by greedy algorithms," Univ. of South Carolina at Columbia, IMI Report 2003:02, 2003.
[58] S. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, "Sparse solutions of linear inverse problems with multiple measurement vectors," *IEEE Trans. Signal Processing*, vol. 53, no. 7, pp. 2477–2488, Jul. 2005.
[59] J. Chen and X. Huo, "Theoretical results about finding the sparsest representations of multiple measurement vectors (MMV) in an overcomplete dictionary using $\ell_1$ minimization and greedy algorithms," Oct. 2004, submitted for publication.
[60] J. A. Tropp, "Algorithms for simultaneous sparse approximation. Part II: Convex relaxation," *EURASIP J. Signal Processing*, Nov. 2004, to appear.
[61] S. Levy and P. K. Fullagar, "Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution," *Geophysics*, vol. 46, no. 9, pp. 1235–1243, 1981.
[62] D. W. Oldenburg, T. Scheuer, and S. Levy, "Recovery of the acoustic impedance from reflection seismograms," *Geophysics*, vol. 48, pp. 1318–1337, 1983.
[63] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Dover Press, 1998, corrected republication.
[64] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *J. Assoc. Comput. Mach.*, vol. 42, pp. 1115–1145, 1995.
[65] A. Feuer and A. Nemirovsky, "On sparse representation in pairs of bases," *IEEE Trans. Inform. Theory*, vol. 49, no. 6, pp. 1579–1581, June 2003.
[66] K. Ball, "Convex geometry and functional analysis," in *Handbook of Banach Space Geometry*, W. B. Johnson and J. Lindenstrauss, Eds. Elsevier, 2002, pp. 161–193.
[67] ——, "An elementary introduction to modern convex geometry," in *Flavors of Geometry*, ser. MSRI Publications. Cambridge University Press, 1997, vol. 31, pp. 1–58.

**Joel A. Tropp** (S'03) earned the B.A. degree in Plan II Liberal Arts Honors and the B.S. degree in Mathematics from the University of Texas at Austin (UT-Austin) in May 1999, graduating *magna cum laude*. He continued his graduate studies in Computational and Applied Mathematics (CAM) at UT-Austin where he completed the M.S. degree in May 2001 and the Ph.D. degree in August 2004.

Dr. Tropp is presently a T. H. Hildebrandt Research Assistant Professor in the Mathematics Department of the University of Michigan at Ann Arbor, and he holds an NSF Mathematical Sciences Postdoctoral Research Fellowship.

As an undergraduate, Dr. Tropp was a recipient of the Barry M. Goldwater National Science Scholarship and a semi-finalist for the British Marshall. He won a CAM Graduate Fellowship in 1999 and a National Science Foundation Graduate Fellowship in 2000.