Computation and Relaxation of Conditions for Equivalence between ℓ^1 and ℓ^0 Minimization

Yoav Sharon, *Student Member, IEEE,* John Wright, *Student Member, IEEE,* and Yi Ma, *Senior Member, IEEE*

***Abstract*—In this paper, we investigate the exact conditions under which the ℓ^1 and ℓ^0 minimizations arising in the context of sparse error correction or sparse signal reconstruction are equivalent. We present a much simplified condition for verifying equivalence, which leads to a provably correct algorithm that computes the exact sparsity of the error or the signal needed to ensure equivalence. In the case when the encoding matrix is imbalanced, we show how an optimal diagonal rescaling matrix can be computed via linear programming, so that the rescaled system enjoys the widest possible equivalence.**

I. INTRODUCTION

Recently, there has been an explosion of research in sparse representation, centering around two mutually complementary mathematical problems:

1) *Sparse Signal Reconstruction:* Finding *sparse* solutions to underdetermined systems of linear equations [1], [2], [3].
2) *Sparse Error Correction:* Finding robust solutions to overdetermined systems of linear equations, subject to arbitrary, but *sparse* errors [4], [5].

The conceptual appeal of sparsity is clear: representations that involve only a few basis elements are far more amenable to human interpretation than arbitrary linear combinations. Sparsity is also the natural goal of data compression: representations involving only a few basis elements are clearly more compact, and may be encoded using fewer bits. Sparsity also arises in the context of error correcting codes, where the error, rather than the signal, is assumed to be sparse [4]. However, the applications of sparse representations range far beyond such traditional problems in signal processing and information theory. For example, the problem of recognizing human faces from images can be cast as a sparse representation problem in which the pattern of sparsity encodes the identity of the subject [6]. Recognition subject to occlusion can be viewed as a special robust decoding problem, in which the sparse errors are concentrated on a fraction of occluded pixels. Further applications exist in image denoising and signal reconstruction [7], as well as in robust state estimation subject to measurement errors.

Much of the promised practical impact of sparse representation is due to a very recently discovered phenomenon, sometimes referred to as ℓ^1-ℓ^0 *equivalence* [3]. Equivalence results roughly state that if the representation to be computed (or the error to be corrected) is *sufficiently sparse*, then the NP-hard problem of finding the sparsest linear representation [8] can be solved efficiently and exactly via linear programming, by minimizing an appropriate ℓ^1-norm. A great deal of effort has been invested into determining, in terms of the properties of the basis or the encoding matrix, how sparse the desired representation must be for equivalence to hold (e.g. [1], [2], [3], [4], [5], amongst others). For example, [9] introduces the so-called *Uniform Uncertainty Principle* for overcomplete bases, which holds if all subsets of bases are approximately orthonormal. Similar hypotheses[1] have proven extremely fruitful in analyzing ℓ^1-ℓ^0 equivalence in random matrix ensembles. One such result states that in overcomplete bases generated at random from a Gaussian distribution, asymptotically, with overwhelming probability, ℓ^1-minimization recovers all sparse representations whose fraction of nonzero elements is less than a fixed constant [4].

Despite these successes, the conditions given in the literature are generally sufficient but not necessary[2], giving very pessimistic indications of the ability of ℓ^1-minimization to recover sparse representations. Simulations (in e.g. [4]) demonstrate a surprisingly large gap between theory and experiment: ℓ^1-minimization significantly outperforms expectations. Positive theoretical results are mostly asymptotic and probabilistic in nature. To the authors' knowledge, no algorithm has been proposed for efficiently verifying equivalence in a given (fixed) linear system of finite size. Yet this question is of the utmost importance in designing engineering systems: knowing exactly how sparse a linear system allows its solutions to be or how many errors the linear system can correct (via ℓ^1-minimization) is essential to analyzing its performance and applicability. Given several proposed linear bases or encoding matrices, knowing their exact ℓ^1-ℓ^0 equivalence properties would allow the engineer to choose the one that maximizes the operating range of ℓ^1-minimization.

Contributions of this paper: The main contribution of this paper is a simple, novel algorithm for determining when ℓ^1-ℓ^0 equivalence holds in a given linear system of equations. We prove the algorithm's correctness and analyze its complexity. In situations where the given encoding matrix is imbalanced, we show how an optimal diagonal rescaling matrix can be computed via linear programming, so that the rescaled system enjoys the widest possible equivalence.

All authors are is with the Coordinated Science Laboratory, Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, IL, 61801 USA e-mail: {ysharon2,jnwright,yima}@uiuc.edu.

[1] Stated, for example, in terms of the *Restricted Isometry Constants* [4].

[2] One noteworthy exception is the necessary and sufficient condition given in [3] in terms of polytope neighborliness, but even there no algorithm is proposed for verifying equivalence.

Organization of this paper: The paper is organized as follows. §II formally introduces the sparse error correction and representation problems and discusses their relationship. §III describes our algorithm for verifying ℓ^1-ℓ^0 equivalence, and proves its correctness. §IV presents a novel algorithm that optimally rescales a given linear system to maximize equivalence, and proves its correctness. Concluding remarks and open questions for future research are given in §V. The source codes for all the algorithms in this paper will be available on-line when the paper is accepted for publication.

II. Sparse representation and ℓ^1-ℓ^0 equivalence

A. Two dual ℓ^0 minimization problems

Consider the following ℓ^0-norm[3] minimization problems:

1) *Sparse Error Correction:* Given $\boldsymbol{y} \in \mathbb{R}^m$, $A \in \mathbb{R}^{m\times n}$ $(m > n)$,
$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - A\boldsymbol{x}\|_0 . \tag{1}$$
2) *Sparse Signal Reconstruction:* Given $\boldsymbol{z} \in \mathbb{R}^p$, $B \in \mathbb{R}^{p\times m}$ $(p < m)$,
$$\boldsymbol{w}^* = \arg\min_{\boldsymbol{w}} \|\boldsymbol{w}\|_0 \quad \text{subject to} \quad \boldsymbol{z} = B\boldsymbol{w}. \tag{2}$$

We will assume that these problems have a unique solution. This will be generally be true if there exists $\boldsymbol{x}_0$ such that the error vector, $\boldsymbol{e} = \boldsymbol{y} - A\boldsymbol{x}_0$, is sparse enough, or if there exists $\boldsymbol{w}_0$ sparse enough such that $\boldsymbol{z} = B\boldsymbol{w}_0$. For example, if any set of $2T$ columns of B are linearly independent, then any sparse representation $\boldsymbol{z} = B\boldsymbol{w}_0$ with $\|\boldsymbol{w}_0\| \leq T$ is the unique solution to (2) [1].

The above two problems are mutually *dual*, or complementary, in the sense that one can convert one problem to the other. It has been shown in [4] that the decoding problem (1) can be converted to the sparse representation problem (2). To see how to convert problem (2) to (1), let $n = m - \text{rank}(B)$ and let A be a full-rank $m \times n$ matrix whose columns span the kernel of B: $BA = 0$. Now find any $\boldsymbol{y}$ so that $\boldsymbol{z} = B\boldsymbol{y}$ and define $f(\boldsymbol{x}) \doteq \boldsymbol{y} - A\boldsymbol{x}$. In this notation,
$$\arg\min_{\boldsymbol{w}:\boldsymbol{z}=B\boldsymbol{w}} \|\boldsymbol{w}\| = f\left(\arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - A\boldsymbol{x}\|\right). \tag{3}$$

B. Equivalence between ℓ^0 and ℓ^1 minimization

Problems (1) and (2) are NP-hard in general [8]. Nevertheless, as shown in [1], [4], if the error $\boldsymbol{e}$ or the solution $\boldsymbol{w}$ is sufficiently sparse, then the solutions to (1) and (2) are *the same as* the solutions to:
$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - A\boldsymbol{x}\|_1 , \tag{4}$$
$$\boldsymbol{w}^* = \arg\min_{\boldsymbol{w}} \|\boldsymbol{w}\|_1 \quad \text{subject to} \quad \boldsymbol{z} = B\boldsymbol{w}, \tag{5}$$
respectively. Problems (4) and (5) can be efficiently solved via linear programming. In addition to this dramatic decrease in computational cost, these problems can also be shown to be robust to small measurement errors [2], [5].

In this paper, we are interested in determining *for any given matrix A (or B), exactly how sparse the error $\boldsymbol{e}$ (or the solution $\boldsymbol{w}$) must be in order for the above ℓ^1 minimizations to solve problems* (1) *and* (2). We therefore define the following threshold for sparse error correction:
$$T^*(A) \doteq \max\ T \in \mathbb{Z} \quad \text{such that}$$
$$\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq T \quad \Rightarrow \quad \boldsymbol{x}_0 = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - A\boldsymbol{x}\|_1 . \tag{6}$$
$T^*(A)$ is the largest number of arbitrary errors in the observation $\boldsymbol{y}$ that can be corrected by ℓ^1-minimization. This quantity is closely related to the *Equivalence Breakdown Point* introduced by Donoho [2], [3] for the sparse signal reconstruction problem (2):
$$\text{EBP}(B) \doteq \max\ k \in \mathbb{Z} \quad \text{such that}$$
$$\|\boldsymbol{w}_0\|_0 \leq k \quad \Rightarrow \quad \boldsymbol{w}_0 = \arg\min_{\boldsymbol{w}:B\boldsymbol{w}=B\boldsymbol{w}_0} \|\boldsymbol{w}\|_1. \tag{7}$$
This is the maximum number of arbitrary nonzeros that can be uniquely recovered by ℓ^1-minimization.

For the remainder of the paper, we consider only the error correction problem (4). However, the duality between error correction and signal reconstruction implies that if A is a full-rank matrix spanning the kernel of B, then $\text{EBP}(B) = T^*(A)$. Thus, the proposed algorithms can be applied to (2) as well.

III. Algorithm for verifying ℓ^1-ℓ^0 equivalence

In this section, we derive an algorithm for computing $T^*(A)$, and hence precisely verifying ℓ^1-ℓ^0 equivalence. We will need the following definition and propositions:

Definition 1: The "*d-skeleton*" is defined to be the collection of all the d-dimensional faces of the standard ℓ^1-ball $B_1 \doteq \{\boldsymbol{v} \in \mathbb{R}^m : \|\boldsymbol{v}\|_1 \leq 1\}$. In particular the 0-skeleton is all the vertices, the 1-skeleton is all the edges, and so on. We will be denote it as $\text{SK}_d(B_1)$:
$$\text{SK}_d(B_1) \doteq \{\boldsymbol{v} \in \mathbb{R}^m : \|\boldsymbol{v}\|_1 = 1, \|\boldsymbol{v}\|_0 \leq d+1\}.$$

Proposition 2: For every $\boldsymbol{x}_0 \in \mathbb{R}^n$ and $\boldsymbol{y} \in \mathbb{R}^m$ the following implication holds
$$\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq T \quad \Rightarrow \quad \boldsymbol{x}_0 = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - A\boldsymbol{x}\|_1 \tag{8}$$
if and only if
$$\forall \boldsymbol{v} \in \text{SK}_{T-1}(B_1), \quad \forall \boldsymbol{z} \in \mathbb{R}^n \setminus 0, \quad \|\boldsymbol{v} + A\boldsymbol{z}\|_1 > 1. \tag{9}$$

Note that this proposition is true regardless of the assumption that for any $\boldsymbol{y}$ and $\boldsymbol{x}_0$ such that $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq T$, (1) will have a unique solution. Naturally, if the assumption does not hold, then neither do (8) and (9) hold.

Proof: Assume that for some $\boldsymbol{x}_0$ and $\boldsymbol{y}$, with $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq T$, (8) does not hold. Then there exists $\boldsymbol{x} \neq \boldsymbol{x}_0$ such that $\|\boldsymbol{y} - A\boldsymbol{x}\|_1 \leq \|\boldsymbol{y} - A\boldsymbol{x}_0\|_1 \doteq c$. Choose $\boldsymbol{v} = \frac{1}{c}(\boldsymbol{y} - A\boldsymbol{x}_0)$. Since $\|\boldsymbol{v}\|_0 \leq T$ and $\|\boldsymbol{v}\|_1 = 1$ we have $\boldsymbol{v} \in \text{SK}_{T-1}(B_1)$. Choose $\boldsymbol{z} = \frac{1}{c}(\boldsymbol{x}_0 - \boldsymbol{x}) \neq 0$, then $\|\boldsymbol{v} + A\boldsymbol{z}\|_1 = \frac{1}{c}\|\boldsymbol{y} - A\boldsymbol{x}\|_1 \leq 1$. Thus (9) does not hold and this proves the sufficiency. Assume now that (9) does not hold, so there exist $\boldsymbol{v} \in \text{SK}_{T-1}(B_1)$ and $\boldsymbol{z} \in \mathbb{R}^n \setminus 0$ such that $\|\boldsymbol{v} + A\boldsymbol{z}\|_1 \leq 1$. Choose $\boldsymbol{x}_0$ arbitrarily and set $\boldsymbol{y} = A\boldsymbol{x}_0 + \boldsymbol{v}$. Then $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq T$, and there exists $\boldsymbol{x} = \boldsymbol{x}_0 - \boldsymbol{z} \neq \boldsymbol{x}_0$ such that $\|\boldsymbol{y} - A\boldsymbol{x}\|_1 = \|\boldsymbol{v} + A\boldsymbol{z}\|_1 \leq 1 = \|\boldsymbol{y} - A\boldsymbol{x}_0\|_1$. Thus (8) does not hold and this proves the necessity. ■

[3] Here, $\|\boldsymbol{w}\|_0$ counts the number of nonzero entries in $\boldsymbol{w}$.

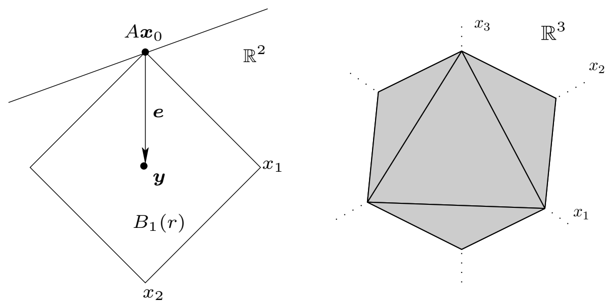

Fig. 1. **Left:** Illustration of the relation between (8) and (9) where $m = 2$, $n = 1$, and $T = 1$. Because $T = 1$ the difference $\boldsymbol{e}$ between $A\boldsymbol{x}$ and $\boldsymbol{y}$ must be parallel to one of the main axes, in this case it is parallel to the vertical axis. The ℓ^1 minimization problem corresponds to finding the ℓ^1 ball of radius r around $\boldsymbol{y}$ that touches the subspace spanned by A. The recovered signal, $\boldsymbol{x}$, is the touching point, and it equals $\boldsymbol{x}_0$ if and only if the subspace spanned by A does not penetrate the ℓ^1 ball. The inequality in (9) verifies that whenever the subspace spanned by A is translated to one of the $(T-1)$-faces of the ℓ^1 ball, penetration does not occur. **Right:** An example where with $m = 3$, there exists such an A that (9) and (8) hold. Now, the one-dimensional span of A is parallel to the reader's line of sight. Here indeed the span of A, translated by any difference vector whose 0-norm equals 1, does not penetrate the ℓ_1 ball. This is illustrated by the fact that all the 6 vertices of the ℓ_1 ball are visible to the reader. From the fact that not all edges are visible, it is concluded that for $T = 2$, (9) does not hold (and thus neither does (8)).

Note that (8) is exactly the condition we want to verify. Figure 1 left illustrates the relation between (8) and (9), and gives the geometric intuition as to why they are equivalent. Figure 1 right shows an example where (8) (and (9)) holds. Although condition (9) is more geometrically intuitive than condition (8), it is still difficult to verify computationally. Without knowing $T^*(A)$, one needs to check starting from $T = 1, 2, \ldots$ until the condition (9) eventually fails. Moreover, condition (9) requires checking that at *every* point on the d-skeleton, the subspace spanned by A does not penetrate the ℓ^1 ball. This leads to the main result of this paper: The condition (9) is in fact equivalent to another condition that is much easier to check.

Proposition 3: Let $A \in \mathbb{R}^{m\times n}$ and $d \in \mathbb{N} \cup 0$ be given and assume the rows of A are in general directions, i.e. any n rows of A are independent[4]. The following holds:

$$\forall \boldsymbol{v} \in \mathrm{SK}_d\left(B_1\right), \quad \forall \boldsymbol{z} \in \mathbb{R}^n \setminus 0 \quad \left\|\boldsymbol{v} + A\boldsymbol{z}\right\|_1 > 1 \tag{10}$$

if and only if for all subsets $I \subset M \doteq \{1, \ldots, m\}$ containing $n-1$ indices, all subsets $J \subset M \setminus I$ containing $T = d+1$ indices, and for some $\boldsymbol{y} \in \mathbb{R}^m$ such that

$$\boldsymbol{y} \in \mathrm{span}\left(A\right) \setminus 0, \quad \forall i \in I \quad y_i = 0, \tag{11}$$

the following holds:

$$\sum_{j\in J} |y_j| < \sum_{j\in M\setminus J} |y_j| \tag{12}$$

Proof: First note that due to the assumption that the rows of A are in general directions, for any $I \subset M$ containing $n-1$ indices, (11) defines $\boldsymbol{y}$ uniquely upto scale ($n-1$ independent equations in n variables). As (12) is invariant to (nonzero) scaling in $\boldsymbol{y}$, it holds for some $\boldsymbol{y}$ satisfying (11) if and only if it holds for all such $\boldsymbol{y}$.

We first prove the sufficiency – the "if" direction. Assume (10) does not hold, then there exists $\boldsymbol{v} \in \mathrm{SK}_d\left(B_1\right)$ and $\boldsymbol{z} \in \mathbb{R}^n \setminus 0$ such that $\left\|\boldsymbol{v} + A\boldsymbol{z}\right\|_1 \leq 1$. Let A_{i*} denote the i-th row of A. Let $P_z \subset M$ be the subset of indices for which $v_i + A_{i*}\boldsymbol{z}$ does not vanish. Let $P_v \subset M$ be the subset of indices for which v_i does not vanish: $|P_v| \leq d+1$. Note that:

$$1 \geq \left\|\boldsymbol{v} + A\boldsymbol{z}\right\|_1 \geq \sum_{i\in P_v} |v_i| - \sum_{i\in P_v} |A_{i*}\boldsymbol{z}| + \sum_{i\in M\setminus P_v} |A_{i*}\boldsymbol{z}|\,.$$

As $\boldsymbol{v} \in \mathrm{SK}_d\left(B_1\right) \Rightarrow \|\boldsymbol{v}\|_1 = \sum_{i\in P_v} |v_i| = 1$ we have

$$\sum_{i\in P_v} |A_{i*}\boldsymbol{z}| \geq \sum_{i\in M\setminus P_v} |A_{i*}\boldsymbol{z}|\,. \tag{13}$$

We will show that there exists $\tilde{\boldsymbol{z}} = \boldsymbol{z} + \boldsymbol{x}$ such that at least $n-1$ of $A_{i*}\tilde{\boldsymbol{z}}$, $i \in M \setminus P_v$ vanish and

$$\sum_{i\in P_v} |A_{i*}\tilde{\boldsymbol{z}}| \geq \sum_{i\in P_v} |A_{i*}\boldsymbol{z}|\,, \tag{14}$$

$$\sum_{i\in M\setminus P_v} |A_{i*}\tilde{\boldsymbol{z}}| \leq \sum_{i\in M\setminus P_v} |A_{i*}\boldsymbol{z}|\,. \tag{15}$$

If such $\tilde{\boldsymbol{z}}$ exists, then (12) will not hold for $I = \{i \in M \setminus P_v\, |A_{i*}\tilde{\boldsymbol{z}} = 0\}$, $\boldsymbol{y} = A\tilde{\boldsymbol{z}}$ and $J = P_v$.

If $|M \setminus (P_z \cup P_v)| \geq n-1$, then we can take $\tilde{\boldsymbol{z}} = \boldsymbol{z}$ and we are done. Otherwise define

$$w_i = \mathrm{sign}\left(A_{i*}\boldsymbol{z}\right)$$

and consider the set of $|M \setminus (P_z \cup P_v)| + 1$ equations in $\boldsymbol{x}$:

$$\sum_{i\in P_v} w_i A_{i*}\boldsymbol{x} = 0 \tag{16}$$

$$A_{i*}\boldsymbol{x} = 0 \quad \forall i \in M \setminus \left(P_z \cup P_v\right). \tag{17}$$

Since the number of equations is less than n, there is at least a one dimensional subspace $\{\alpha\boldsymbol{x} : \alpha \in \mathbb{R}\}$ of solutions. For any solution $\boldsymbol{x}$, (14) holds with $\tilde{\boldsymbol{z}} = \boldsymbol{z} + \boldsymbol{x}$, since:

$$\sum_{i\in P_v} |A_{i*}\left(\boldsymbol{z} + \boldsymbol{x}\right)| \geq \sum_{i\in P_v} w_i A_{i*}\left(\boldsymbol{z} + \boldsymbol{x}\right) = \sum_{i\in P_v} |A_{i*}\boldsymbol{z}| + 0.$$

Now choose α small enough so that

$$-w_i\alpha A_{i*}\boldsymbol{x} \leq w_i A_{i*}\boldsymbol{z} \quad \forall i \in P_z \setminus P_v \tag{18}$$

(note that $w_i A_{i*}\boldsymbol{z} > 0\ \forall i \in P_z \setminus P_v$). We then have

$$\sum_{i\in M\setminus P_v} |A_{i*}\left(\boldsymbol{z} + \alpha\boldsymbol{x}\right)| = \sum_{i\in P_z\setminus P_v} w_i A_{i*}\boldsymbol{z} + \alpha \sum_{i\in P_z\setminus P_v} w_i A_{i*}\boldsymbol{x}$$

where we used (17) and the definition of P_z to restrict the summation to be over $P_z \setminus P_v$ instead of over $M \setminus P_v$. By taking the sign of α to be opposite to that of $\sum_{i\in P_z\setminus P_v} w_i A_{i*}\boldsymbol{x}$, then as long as (18) holds (15) also holds. If we now choose the magnitude of α such that

$$|\alpha| = \min_{i:\mathrm{sign}(\alpha A_{i*}\boldsymbol{x})\neq\mathrm{sign}(A_{i*}\boldsymbol{z})} \frac{|A_{i*}\boldsymbol{z}|}{|A_{i*}\boldsymbol{x}|}, \tag{19}$$

[4] Our approach can still be applied if the rows of A are not in general direction, yet it will make the algorithm more complicated.

then (18) still holds, but we also have that for some $j \in P_z \backslash P_v$, $A_{j*}\tilde{\boldsymbol{z}} = A_{j*}(\boldsymbol{z} + \alpha \boldsymbol{x}) = 0$. If now at least $n-1$ of $A_{i*}\tilde{\boldsymbol{z}}$, $i \in M \setminus P_v$ vanish then there exists $I \subset M \setminus P_v$ such that $A_{i*}\tilde{\boldsymbol{z}} = 0 \ \forall i \in I$ and with $\boldsymbol{y} = A\tilde{\boldsymbol{z}}$ and $J = P_v$ (12) does not hold. If less than $n-1$ of $A_{i*}\tilde{\boldsymbol{z}}$, $i \in M \setminus P_v$ vanish, then we can replace $\boldsymbol{z}$ with $\tilde{\boldsymbol{z}}$, redefine P_z, and repeat the process. Since at each iteration $P_z \setminus P_v$ decreases by at least 1, we are guaranteed to eventually find $\tilde{\boldsymbol{z}}$ for which $n-1$ of $A_{i*}\tilde{\boldsymbol{z}}$, $i \in M \setminus P_v$ do vanish. This proves sufficiency.

For the necessity – the "only if" direction, assume there exists $I \subset M$, $|I| = n-1$, $J \subset M \setminus I$, $|J| = d+1$, and some $\boldsymbol{y}$ for which (11) holds but (12) does not. Set $c = \sum_{j \in J} |y_j|$ and $\boldsymbol{v}$ such that $\forall i \in J$, $v_i = \frac{1}{c} y_i$ and $\forall i \in M \setminus J$, $v_i = 0$. Let $\boldsymbol{x}$ such that $\boldsymbol{y} = A\boldsymbol{x}$, and set $\boldsymbol{z} = -\frac{1}{c}\boldsymbol{x}$. We now have that $\boldsymbol{v} \in \mathrm{SK}_d(B_1)$ and $\boldsymbol{z} \in \mathbb{R}^n \setminus 0$. Furthermore,

$$\|\boldsymbol{v} + A\boldsymbol{z}\|_1 = \frac{1}{c} \sum_{j \in M \setminus J} |A_{j*}\boldsymbol{x}| = \frac{\sum_{j \in M \setminus J} |y_j|}{\sum_{j \in J} |y_j|} \leq 1,$$

where we used the assumption that (12) does not hold for the last inequality. Thus (10) does not hold and this proves the necessity. ∎

The following algorithm exploits Proposition 3 to compute the maximum T, $T^*(A)$, for which (8) holds:

Algorithm 1 Computing $T^*(A)$

Input: $A \in \mathbb{R}^{m \times n}$.

1: Set $T \leftarrow m$ and let $I_1, \ldots, I_N$, $N = \binom{m}{n-1}$, be all the subsets of $M \doteq \{1, \ldots, m\}$ containing $n-1$ indices.
2: **for** $k = 1 : N$ **do**
3: Find a nontrivial solution $\boldsymbol{x} \in \mathbb{R}^n$ to
$$A_{i*}\boldsymbol{x} = 0 \qquad \forall i \in I_k.$$
4: Set $\boldsymbol{y} \leftarrow A\boldsymbol{x}$ and reorder the elements of $\boldsymbol{y}$ such that
$$|y_{r_1}| \geq |y_{r_2}| \geq \cdots \geq |y_{r_m}|.$$
5: Find the largest integer, s, such that
$$\sum_{i=1}^{s} |y_{r_i}| < \sum_{i=s+1}^{m} |y_{r_i}|. \tag{20}$$
6: Set $T \leftarrow \min\{T, s\}$.
7: **end for**

Output: T.

Theorem 4: Algorithm 1 returns the maximum $T^*(A)$ for which (8) holds.

Proof: Note that due to the sorting, (20) holds for some integer s if and only if (12) holds for all subsets $J \in M \setminus I_k$ containing s elements. Proposition 2, Proposition 3, and the fact that the algorithm returns the minimum integer s for which (20) holds over all subsets $I \subset M$, $|I| = n-1$, and all subsets $J \subset M \setminus I$, $|J| = s$ proves the theorem. ∎

The required computation time is polynomial for any fixed m, but exponential[5] in general. This complexity is typical of combinatorial problems involving polytopes, see [10]. Algorithm 1 is therefore most suitable for small-to-moderate size problems. For example, if we fix $m = 3n$, our straightforward Matlab implementation can process matrices upto size 27×9 on a standard PC. If we fix $m = 6n$, the algorithm succeeds upto size 36×6. When the computational complexity exceeds available resources, a good upper bound for $T(A)$ can still be obtained by randomly sampling supports I_j. Alternatively, if all sparsity patterns are considered equally likely, the algorithm can be easily modified to instead compute, for each T, the probability that ℓ^1-minimization will correct T errors.

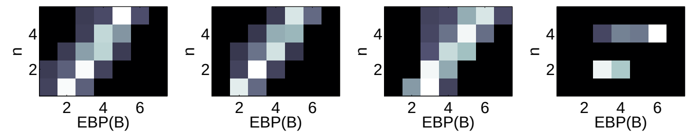

(a) Gaussian (b) Sphere (c) DCT (d) Hadamard

Fig. 2. Histogram of EBP in various $5n \times 6n$ random dictionaries. Vertical axis: n. Horizontal axis: histogram of $\mathrm{EBP}(B)$.

Example 5 (Sparse recovery with random dictionaries): We apply Algorithm 1 to test the capability of ℓ^1-minimization to recover sparse solutions in dictionaries sampled from several random matrix ensembles. The ensembles considered are Gaussian (entries iid normal), Partial DCT (uniformly chosen rows of the DCT matrix), Partial Hadamard (uniformly chosen rows of a Hadamard matrix) [11], and Uniform Sphere (columns sampled uniformly from the sphere). In each case, we generate a $5n \times 6n$ overcomplete dictionary B and calculate $\mathrm{EBP}(B) = T^*(B^\perp)$. We repeat $1{,}000$ times for each $n = 1, \ldots, 5$. Figure 2 plots the histograms for each of these ensembles. Notice that when the Hadamard matrix exists[6], it exhibits better sparse recovery capability than the others. Interestingly, the Partial DCT matrix performs competitively with Gaussian and Spherical, despite asymptotics that are worse by a log factor [12].

IV. Computing Scaling Parameters for Higher Robustness

We motivate our next problem through an example:

Example 6: Consider the case of $m = 3$, $n = 1$, $A = [1, 0.1, 0.1]^T$. If we translate span(A) to the vertex $\boldsymbol{v} = [1, 0, 0]^T$ of the unit ℓ^1-ball, B_1, it will penetrate B_1. Thus $T^*(A) = 0$, and ℓ^1 minimization does not correct even a single arbitrary error. However, we can rotate span(A) so it will not penetrate B_1, for example by multiplying A on the right with $D \doteq \mathrm{diag}(1, 10, 10)$. Assume we are given $\boldsymbol{y}$ and $\boldsymbol{x}_0$ are such that $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \leq 1$. Because D is diagonal, $\|D\boldsymbol{y} - DA\boldsymbol{x}_0\|_0 \leq 1$. By construction, the condition (9) holds for DA with $T = 1$. Thus, by Proposition 2, we can recover $\boldsymbol{x}_0$ as:

$$\boldsymbol{x}_0 = \arg\min_{\boldsymbol{x}} \|D\boldsymbol{y} - DA\boldsymbol{x}_0\|_1. \tag{21}$$

[5] If $m, n \to \infty$ in a fixed ratio $n/m = \rho$, the time complexity is $\asymp 2^{mH(\rho)}$, where $H(\cdot)$ is the binary entropy function.

[6] Since Hadamard matrices do not exist for arbitrary n, this ensemble is only tested in the 10×12 and 20×24 cases.

Example 6 shows that rescaling by an appropriate diagonal[7] matrix can increase T^*, allowing ℓ^1 minimization (4) to correct more errors. The following two results show how to systematically compute such a D, if one exists.

Proposition 7: Let $A \in \mathbb{R}^{m \times n}$ be full-rank, $I \subset \{1, \ldots, m\}$, $|I| = n-1$ and $D \in \mathbb{R}^{m \times m}$ a strictly positive diagonal matrix ($D_{ii} > 0$). For every $\boldsymbol{y} \in \mathbb{R}^m$, $\boldsymbol{y} \in \text{span}(A)$ and every $\boldsymbol{y}' \in \mathbb{R}^m$, $\boldsymbol{y}' \in \text{span}(DA)$, such that $\forall i \in I$, $y_i = y'_i = 0$, there exists $\alpha \in \mathbb{R}$ such that

$$y'_j = \alpha D_{jj} y_j, \qquad \forall j \in M \doteq \{1, \ldots, m\}\,. \tag{22}$$

Proof: Let $\boldsymbol{x}$ and $\boldsymbol{x}'$ such that $\boldsymbol{y} = A\boldsymbol{x}$ and $\boldsymbol{y}' = DA\boldsymbol{x}'$. From $\forall i \in I$, $y_i = y'_i = 0$ and the strict positivity of D we have that $\forall i \in I$ $A_{i*}\boldsymbol{x} = A_{i*}\boldsymbol{x}' = 0$. Since $|I| = n-1$ and the rank of A is n, there exists α such that $\alpha A\boldsymbol{x} = A\boldsymbol{x}'$. Multiplying both sides by D gives (22). ■

Theorem 8: Let $A \in \mathbb{R}^{m \times n}$ be full-rank. For every $\boldsymbol{x}_0 \in \mathbb{R}^n$ and $\boldsymbol{y} \in \mathbb{R}^m$ such that $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 \le T$ one can recover $\boldsymbol{x}_0$ using the following ℓ^1-minimization:

$$\boldsymbol{x}_0 = \arg\min_{\boldsymbol{x}} \|D\boldsymbol{y} - DA\boldsymbol{x}\|_1 \tag{23}$$

where D is a strictly positive diagonal matrix *if and only if* for all subsets $I \subset M$ containing $n-1$ elements, all subsets $J \subset M \setminus I$ containing $T = d+1$ elements, and for some $\boldsymbol{y} \in \mathbb{R}^m$ such that

$$\boldsymbol{y} \in \text{span}\,(A)\,, \quad \forall i \in I \quad y_i = 0, \tag{24}$$

the following holds:

$$\sum_{j \in J} D_{jj}\,|y_j| < \sum_{j \in M \setminus J} D_{jj}\,|y_j| \tag{25}$$

Proof: Modify Propositions 2 and 3 by substituting DA in place of A. Using the fact that D is a diagonal matrix $\|D\boldsymbol{y} - DA\boldsymbol{x}\|_0 \le T$ from the modified (8) can be replaced with $\|\boldsymbol{y} - A\boldsymbol{x}\|_0 \le T$. Using Proposition 7 the modified (11) and (12) can be replaced with (24) and (25). This completes the proof. ■

Since (25) is linear with respect to the elements of D, we can write (25) for all subsets $I \subset M$ containing $n-1$ indices, all subsets $J \subset M \setminus I$ containing $T = d+1$ elements, and some $\boldsymbol{y}$ for which (24) holds, and then solve for a feasible D^* using linear programming.

Example 9 (Improving robustness of Gaussian codebooks): We randomly sample 250 Gaussian code books of size 15×3, A, and compute $T^*(A)$ before and after scaling by the optimal D^* calculated using Theorem 8. The median $T^*(A)$ is 2, while the median $T^*(D^*A)$ is 3. In 95% of cases, $T^*(D^*A) > T^*(A)$, and in 10% of cases $T^*(D^*A) = 4$, which is also an upper bound for T for this matrix size [3]. Thus, the vast majority of Gaussian codebooks are suboptimal, and many cannot even be rescaled to an optimal codebook by any diagonal matrix D!

[7]Note that if D is not necessarily diagonal, then even if (9) holds for DA with some T, it does not imply that we can use (21) to solve for (1) with the same T since we can find $\boldsymbol{y}$ and $\boldsymbol{x}_0$ such that $\|\boldsymbol{y} - A\boldsymbol{x}_0\|_0 = T$ but $\|D\boldsymbol{y} - DA\boldsymbol{x}_0\|_0 > T$. While this, however, does not necessarily rule out the existence of a non diagonal matrix that can increase the number of cases in which we can use (21) to solve for (1), in this current paper we decided to consider only diagonal matrices.

V. Discussions

The conditions (and the associated algorithms) given in this paper allow exact analysis of the ℓ^1-ℓ^0 equivalence properties of small to moderate sized linear systems, allowing the engineer to guarantee optimal robustness or sparse recovery via ℓ^1-minimization. Small matrix results as in Examples 5 and 9 provide a useful counterpart to the asymptotics in the literature, allowing a more informed selection of measurement ensembles for compressed sensing or robust error correction codes. Exact polynomial-time algorithms for the problems considered here seem unlikely, but their NP-hardness is open. Real problems demand good bounds on $T^*(A)$ for large A (e.g. $5,000 \times 800$ in [6]). An important future question is therefore to analyze how many randomly sampled $\boldsymbol{y}$ in Algorithm 1 are needed to accurately estimate T^* with high probability.

Acknowledgment

The authors would like to thank professor Harm Derksen of the Mathematics Department of the University of Michigan, Dr. Mike Wakin of the Mathematics Department of the California Institute of Technology, and professor Robert Fossum of the Mathematics Department of the University of Illinois, for stimulating discussions on the problems studied in this paper.

References

[1] D. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ^1-norm near solution approximates the sparsest solution," *preprint, http://www-stat.stanford.edu/ donoho/Reports/*, 2004.
[2] ——, "For most large underdetermined systems of linear equations the minimal ℓ^1-norm solution is also the sparsest solution," *Comm. Pure and Applied Math.*, vol. 59, no. 6, pp. 797–829, 2006.
[3] ——, "Neighborly polytopes and sparse solution of underdetermined linear equations," *IEEE. Trans. Info. Theory*, 2006.
[4] E. Candes, M. Rudelson, T. Tao, and R. Vershynin, "Error correction via linear programming," in *IEEE Symposium on FOCS*, 2005, pp. 295–308.
[5] E. Candes and P. A. Randall, "Highly robust error correction by convex programming," *preprint, http://arxiv.org/abs/cs.IT/0612124*, 2006.
[6] J. Wright, A. Yang, A. Ganesh, and Y. Ma, "Recognition via sparse representation: Robustness, occlusion and feature selection," *CSL Technical Report UILU-ENG-07-2205, DC-228, University of Illinois at Urbana Champaign*, 2007.
[7] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, pp. 489–509, 2004.
[8] E. Amaldi and V. Kann, "On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems," *Theoretical Computer Science*, vol. 209, pp. 237–260, 1998.
[9] E. Candes and T. Tao, "Near-optimal signal recovery from random projections: universal encoding strategies," *IEEE Trans. Inform. Theory*, vol. 52, pp. 5406–5425, 2004.
[10] V. Kaibel and M. Pfetsch, "Some algorithmic problems in polytope theory," in *Algebra, Geometry and Software Systems*, M. Joswig and N. Takayama, Eds., 2003.
[11] E. W. Weisstein, "Hadamard matrix," *MathWorld–Wolfram, http://mathworld.wolfram.com/HadamardMatrix.html*.
[12] E. Candes, "Compressive sampling," in *Proc. International Congress of Mathematicians*, 2006.