OPTIMAL SPARSE REPRESENTATIONS IN GENERAL OVERCOMPLETE BASES

Dmitry M. Malioutov, Müjdat Çetin, and Alan S. Willsky

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
77 Massachusetts Ave., Cambridge, MA 02139, USA

ABSTRACT

We consider the problem of enforcing a sparsity prior in underdetermined linear problems, which is also known as sparse signal representation in overcomplete bases. The problem is combinatorial in nature, and a direct approach is computationally intractable even for moderate data sizes. A number of approximations have been considered in the literature, including stepwise regression, matching pursuit and its variants, and recently, basis pursuit (ℓ_1) and also ℓ_p-norm relaxations with $p < 1$. Although the exact notion of sparsity (expressed by an ℓ_0-norm) is replaced by ℓ_1 and ℓ_p norms in the latter two, it can be shown that under some conditions these relaxations solve the original problem exactly. The seminal paper of Donoho and Huo establishes this fact for ℓ_1 (basis pursuit) for a special case where the linear operator is composed of an orthogonal pair. In this paper, we extend their results to a general underdetermined linear operator. Furthermore, we derive conditions for the equivalence of ℓ_0 and ℓ_p problems, and extend the results to the problem of enforcing sparsity with respect to a transformation (which includes total variation priors as a special case). Finally, we describe an interesting result relating the sign patterns of solutions to the question of ℓ_1-ℓ_0 equivalence.

1. INTRODUCTION

The topic of enforcing a sparsity prior in underdetermined linear problems has many important applications including feature selection, signal restoration, super-resolution source localization, and subset selection in linear regression, among many others. Mathematically, the basic version of the problem can be described as follows. Given a signal $\mathbf{y} \in \mathbb{C}^M$, and a matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$ with $M < N$, we would like to recover an unknown sparse signal $\mathbf{x} \in \mathbb{C}^N$, which satisfies $\mathbf{y} = \mathbf{A}\mathbf{x}$. The linear system has an infinite number of solutions. The min-ℓ_2-norm solution is not sparse in general. Instead, we would like to choose the sparsest solution, the one with the smallest number of non-zero elements. We denote the number of nonzero elements of a signal $\mathbf{x}$ by $\|\mathbf{x}\|_0^0$. The sparse representation problem (the ℓ_0 problem) is:

$$\min \ \|\mathbf{x}\|_0^0 \ \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x} \quad (\ell_0 \text{ problem}) \tag{1}$$

This problem is combinatorial in nature, and the solution requires searching through all subsets of indices of $\mathbf{x}$. This is not tractable even for moderate values of M and N. It has been proposed to consider ℓ_1-norm (Basis Pursuit) [1] and ℓ_p-norm approximations (with $p < 1$) which can be solved using continuous optimization methods:

$$\min \ \|\mathbf{x}\|_1 \ \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x} \qquad (\ell_1 \text{ problem}) \tag{2}$$

$$\min \ \|\mathbf{x}\|_p^p \ \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x}, \ p < 1 \quad (\ell_p \text{ problem}) \tag{3}$$

We consider the two cases ($p = 1$ and $p < 1$) separately, since the problem is convex only when $p = 1$. The ℓ_1 problem can be solved efficiently by linear programming for real data and by second order cone (SOC) programming for complex data [2]. The ℓ_p problem is nonconvex, but a locally optimal solution can be found by local optimization; for some applications the resulting local minima are observed to give excellent approximations [2].

According to empirical observations, ℓ_1 and ℓ_p problems lead to sparse solutions. However, there are more convincing reasons to use them instead of the exact sparsity approach in (1). The pioneering work of Donoho and Huo [3] establishes that the solution of the ℓ_1 problem is the same as the solution of the ℓ_0 problem, as long as the underlying $\mathbf{x}$ is *sparse enough* with respect to the matrix $\mathbf{A}$ (these notions are to be defined shortly). They consider the special case where $\mathbf{A}$ is composed of two orthogonal bases.

In this work, we extend their results to the case of a general overcomplete basis[1]. Furthermore, we derive conditions for the equivalence of global solutions to the ℓ_0 and ℓ_p problems, which get less restrictive as p decreases to 0. In addition, we prove a result which sheds light on the sign patterns of exact ℓ_1 solutions to the ℓ_0 problem, which is most interesting when the sufficient conditions for equivalence are not satisfied. Finally, we extend the results to the problem of signal representations which are sparse after a linear transformation $\mathbf{D} \in \mathbb{C}^{J \times N}$. The problem has the form $\min \|\mathbf{D}\mathbf{x}\|_1$ subject to $\mathbf{y} = \mathbf{A}\mathbf{x}$. A special case of this problem, where $\mathbf{D}$ is a representation of a discrete gradient operator, goes under the name *total variation*, and has important applications in image processing. The significance of the results on equivalence of the ℓ_1, ℓ_p and ℓ_0 problems under specified conditions lies in being able to replace combinatorial optimization by the much easier tasks of linear programming or continuous ℓ_p optimization.

Although we do not discuss it in this paper, the application of these ideas typically requires handling noise, which can be simply accomplished for Gaussian noise through the use of the following cost function: $\arg\min \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_p^p$, where λ is a parameter balancing the two terms. This problem is discussed in [2] and [1]. The analysis in this paper, as well as simulations in [2] provide a strong motivation for the use of ℓ_1 and ℓ_p penalties in the noisy case as well.

This work was supported by the Army Research Office under Grant DAAD19-00-1-0466, and the Office of Naval Research under Grant N00014-00-1-0089.

[1] During the preparation of this manuscript we learned that very recent work related to this extension has been done independently in [4], [5], and [6].

2. ℓ_0 UNIQUENESS CONDITIONS

Before relating the ℓ_0, ℓ_1 and ℓ_p problems, we first address the question of uniqueness of solutions of the ℓ_0 problem, to make sure that ℓ_0 solutions are useful. To that end, we introduce a measure of independence of sets of columns of $\mathbf{A}$, index of unambiguity $K(\mathbf{A})$, which leads to a necessary and sufficient condition for the uniqueness of solutions.

Let $\mathbf{A} \in \mathbb{C}^{M \times N}$ with columns $\mathbf{a}_i$, $\mathbf{A} = [\mathbf{a}_1, ..., \mathbf{a}_N]$. The index of ambiguity $K(\mathbf{A})$ of $\mathbf{A}$ is the largest integer such that any set of $K(\mathbf{A})$ columns of $\mathbf{A}$ is linearly independent. Hence, either $K(\mathbf{A}) = N$, and no additional columns exist, or there exists a set of $K(\mathbf{A}) + 1$ columns which is linearly dependent. Using this definition, we can characterize the uniqueness of solutions to the ℓ_0 problem through the following theorem:

Theorem 1 (Uniqueness of solutions to the ℓ_0 cost function) *Let $\mathbf{A} \in \mathbb{C}^{M \times N}$ with $N > M$. Also, suppose that for some $\mathbf{x}^*$, $\mathbf{y} = \mathbf{A}\mathbf{x}^*$, and $\|\mathbf{x}^*\|_0^0 = L$. Then (1) has a unique solution $\hat{\mathbf{x}} = \mathbf{x}^*$ for all such $\mathbf{y}$ and $\mathbf{x}^*$ if and only if $L < (K(\mathbf{A}) + 1)/2$.*

We refer the reader to [2] for the simple proof, which is based on the observation that there is a vector in the nullspace of $\mathbf{A}$ which has $K(\mathbf{A}) + 1$ non-zero entries. This condition is necessary and sufficient, but due to its discontinuous nature it is hard to use it to relate the ℓ_0, ℓ_1, and ℓ_p problems. To move on, we introduce a different measure of independence of columns of $\mathbf{A}$, maximum absolute pairwise dot-product, $M(\mathbf{A})$, which extends the measure introduced in [3] to general overcomplete bases:

$$M(\mathbf{A}) = \max_{i \neq j} |\mathbf{a}_i' \mathbf{a}_j|, \text{ where } \|\mathbf{a}_k\|_2 = 1, \forall k. \tag{4}$$

$M(\mathbf{A})$ measures how spread-out the columns of $\mathbf{A}$ are. Due to Schwartz inequality, $0 \leq M(\mathbf{A}) \leq 1$, and $M(\mathbf{A}) = 0$ if and only if $\mathbf{A}$ has orthogonal columns. Small values of $M(\mathbf{A})$ mean that the set of columns is almost orthogonal, whereas values close to unity mean that there are at least two columns separated by a very small angle. Although $M(\mathbf{A})$ takes into account only the relations between pairs of columns of $\mathbf{A}$, it has a strong tie with linear dependence structure of larger sets of columns, and in particular with $K(\mathbf{A})$, as described by the following theorem:

Theorem 2 (Relation of $K(\mathbf{A})$ to $M(\mathbf{A})$) *Let $\mathbf{A} \in \mathbb{C}^{M \times N}$ with $\|\mathbf{a}_i\|_2 = 1$, $\forall i$ (thus $\mathbf{A}$ does not have $\mathbf{0}$ as a column), then*

$$M(\mathbf{A}) \geq \frac{1}{K(\mathbf{A})} \tag{5}$$

The proof of the theorem (see [2]) involves the following fact from geometry of polytopes[2]:

Lemma 1 (Optimality of the simplex for line packing in $\mathbb{C}^K$) *Let $\tilde{\mathbf{A}} \in \mathbb{C}^{K \times (K+1)}$. Then $M(\tilde{\mathbf{A}}) \geq \frac{1}{K}$. The equality is achieved for the regular simplex (allowing rotations and reflections of vertices around the origin).*

Now using Theorems 1 and 2 we are led to a different condition for the uniqueness of solutions to the ℓ_0 problem:

[2]There are close connections between the problem of characterizing $M(\mathbf{A})$, the problem of spherical code design, and the problem of ray and line packing on Euclidean spheres, where we borrowed this result from. We thank R. Blume-Kohout and P. Shor for a helpful discussion on the subject, and the proof of Lemma 1.

Theorem 3 (Uniqueness of solutions for ℓ_0 through $M(\mathbf{A})$) *Let $\mathbf{A} \in \mathbb{C}^{M \times N}$ with $N > M$. Also, suppose that for some $\mathbf{x}^*$, $\mathbf{y} = \mathbf{A}\mathbf{x}^*$, and $\|\mathbf{x}^*\|_0^0 = L$. Then (1) has a unique solution $\hat{\mathbf{x}} = \mathbf{x}^*$ if $L < \frac{1/M(\mathbf{A})+1}{2}$.*

This condition is sufficient only, and the ℓ_0 problem may have unique solutions even when it is not satisfied. The reason for introducing a looser condition is that we next use $M(\mathbf{A})$ to relate the ℓ_0 and ℓ_1 problems, and the same sufficient condition leads to the equivalence of the two problems.

3. EQUIVALENCE CONDITION FOR ℓ_0 AND ℓ_1 PROBLEMS FOR GENERAL OVERCOMPLETE BASES

Donoho and Huo [3], and later Elad and Bruckstein [7] give a sufficient condition for the equivalence of solutions to the ℓ_0 and ℓ_1 problems for the case when $\mathbf{A}$ is composed of two orthogonal bases. We extend the condition to the general overcomplete basis case using the measure $M(\mathbf{A})$ introduced in (4):

Theorem 4 (Equivalence of ℓ_0 and ℓ_1 problems) *Suppose that the ℓ_0 problem (1) has a unique solution $\hat{\mathbf{x}}$ with $\|\hat{\mathbf{x}}\|_0^0 = L$. If $L < \frac{1+1/(M(\mathbf{A}))}{2}$, then $\hat{\mathbf{x}}$ is also the solution of the ℓ_1 problem in (2).*

Proof. In the beginning stages, the structure of the proof follows that of [3] and [7], generalizing some of the notions for a general overcomplete basis. One novel aspect of the proof is the derivation of Lemma 2. Suppose that $\hat{\mathbf{x}}$ is the optimal solution to (1). To satisfy $\mathbf{y} = \mathbf{A}\tilde{\mathbf{x}}$, any other candidate $\tilde{\mathbf{x}}$ must have the form $\tilde{\mathbf{x}} = \hat{\mathbf{x}} + \boldsymbol{\delta}$, where $\boldsymbol{\delta} \in Null(\mathbf{A})$, the nullspace of $\mathbf{A}$. In order for $\hat{\mathbf{x}}$ to be the optimal ℓ_1 solution as well, we need:

$$\|\hat{\mathbf{x}} + \boldsymbol{\delta}\|_1 > \|\hat{\mathbf{x}}\|_1 \text{ for any } \boldsymbol{\delta} \in Null(\mathbf{A}),\ \boldsymbol{\delta} \neq \mathbf{0} \tag{6}$$

Let $\mathcal{I}_x$ denote the set of indices where the optimal ℓ_0 solution $\hat{\mathbf{x}}$ has non-zero values (the support of $\hat{\mathbf{x}}$). Also its complement, the set of zero-valued indices of $\hat{\mathbf{x}}$, is denoted by $\mathcal{I}_x^C$. We can divide the ℓ_1 norm into the components on and off the support of $\hat{\mathbf{x}}$ and then use the triangle inequality to manipulate (6) as follows:

$$\|\hat{\mathbf{x}} + \boldsymbol{\delta}\|_1 - \|\hat{\mathbf{x}}\|_1 = \left(\sum_{i \in \mathcal{I}_x} |\hat{x}_i + \delta_i| + \sum_{i \in \mathcal{I}_x^C} |\delta_i| \right) - \sum_{i \in \mathcal{I}_x} |\hat{x}_i| \tag{7}$$
$$\geq \sum_{i \in \mathcal{I}_x^C} |\delta_i| - \sum_{i \in \mathcal{I}_x} |\delta_i| = \|\boldsymbol{\delta}\|_1 - 2 \sum_{i \in \mathcal{I}_x} |\delta_i| > 0$$

So, for $\hat{\mathbf{x}}$ to be ℓ_1-optimal, it suffices to have $\frac{\sum_{i \in \mathcal{I}_x} |\delta_i|}{\|\boldsymbol{\delta}\|_1} < \frac{1}{2}$. This is a good start, but hard to test numerically. To move further, similarly to [7], we consider a family of problems indexed by i:

$$\min\ \|\boldsymbol{\delta}\|_1 \text{ subject to } \boldsymbol{\delta} \in Null(\mathbf{A}) \text{ and } \delta_i = 1 \tag{8}$$

Suppose the minimum value is Q_i, when index i is fixed. Define $Q(\mathbf{A}) = \min_i\ Q_i$. Next, we use Lemma 2, which states that $Q(\mathbf{A}) \geq (1 + \frac{1}{M(\mathbf{A})})$. Consider the condition that we are trying to prove, (7):

$$\frac{\sum_{i \in \mathcal{I}_x} |\delta_i|}{\|\boldsymbol{\delta}\|_1} = \sum_{i \in \mathcal{I}_x} \frac{|\delta_i|}{\|\boldsymbol{\delta}\|_1} \leq \sum_{i \in \mathcal{I}_x} \frac{1}{Q_i} \tag{9}$$
$$\leq \sum_{i \in \mathcal{I}_x} \frac{1}{Q(\mathbf{A})} = \|\hat{\mathbf{x}}\|_0^0 \frac{1}{Q(\mathbf{A})} \leq \|\hat{\mathbf{x}}\|_0^0 (1 + \frac{1}{M(\mathbf{A})})^{-1}$$

We need to have $\frac{\sum_{i \in \mathcal{I}_x} |\delta_i|}{\|\boldsymbol{\delta}\|_1} < \frac{1}{2}$. In order for that to happen, it is sufficient that $\|\hat{\mathbf{x}}\|_0^0 < \frac{1}{2}(1 + \frac{1}{M(\mathbf{A})})$. This proves the theorem. ■ In the proof we used Lemma 2. Its derivation appears in [2].

Lemma 2 (Bound on $Q(\mathbf{A})$ for a general overcomplete $\mathbf{A}$) *$Q(\mathbf{A})$, the minimum of the optimum values of the problems in (8) over all i, can be bounded by $Q(\mathbf{A}) \geq (1 + \frac{1}{M(\mathbf{A})})$.*

4. EQUIVALENCE CONDITION FOR ℓ_0 AND ℓ_P

The approximation of the ℓ_0 norm by an ℓ_p quasi-norm with $p < 1$ is more accurate than by the ℓ_1 norm. We show that under certain conditions the global minimum of the ℓ_p problem also achieves the minimum of the ℓ_0 problem, and moreover, as $p \to 0$, the sufficient condition for the equivalence of the two problems approaches the necessary and sufficient condition for the uniqueness of solutions to the ℓ_0 problem, stated in Theorem 1.

To relate the ℓ_0 and ℓ_p problems, we consider a new measure of $\mathbf{A}$ based on the order statistics of vectors in $Null(\mathbf{A})$. For any $\boldsymbol{\delta} \in Null(\mathbf{A})$ define $\tilde{\boldsymbol{\delta}}$ to be a permutation of $\boldsymbol{\delta}$ in which the absolute values of the coordinates δ_i are sorted in decreasing order. Thus, $\tilde{\delta}_1 = \max_i |\delta_i|$, $\tilde{\delta}_N = \min_i |\delta_i|$. We define $S(\mathbf{A})$ as:

$$S(\mathbf{A}) = \min\ \tilde{\delta}_{K(\mathbf{A})+1} \text{ over all } \boldsymbol{\delta} \in Null(\mathbf{A}), \text{ with } \|\boldsymbol{\delta}\|_\infty = 1 \tag{10}$$

Note that we consider the $(K(\mathbf{A}) + 1)$-st ranking element, hence $S(\mathbf{A})$ is strictly greater than zero (otherwise there would exist a set of $K(\mathbf{A})$ linearly dependent columns of $\mathbf{A}$). Also $S(\mathbf{A}) \leq 1$, since $\|\boldsymbol{\delta}\|_\infty = \tilde{\delta}_1 = 1$. Using the new measure of $\mathbf{A}$, we are led to the following theorem:

Theorem 5 (Equivalence of ℓ_0 and ℓ_p with $p \leq 1$) *Suppose that the ℓ_0 problem (1) has a unique solution $\hat{\mathbf{x}}$ with $\|\hat{\mathbf{x}}\|_0^0 = L$. If $L < \frac{S(\mathbf{A})^p (K(\mathbf{A})+1)}{1+S(\mathbf{A})^p}$, then $\hat{\mathbf{x}}$ is also the solution of the ℓ_p problem in (3).*

Proof outline. In order for $\hat{\mathbf{x}}$ to be the unique solution of (3), it must be true that $\|\hat{\mathbf{x}}\|_p^p < \|\hat{\mathbf{x}} + \boldsymbol{\delta}\|_p^p$ for any $\boldsymbol{\delta} \in Null(\mathbf{A})$, $\boldsymbol{\delta} \neq \mathbf{0}$. The triangle inequality holds for $\| \bullet \|_p^p$ with $p < 1$:

$$\|\hat{\mathbf{x}}\|_p^p - \|\boldsymbol{\delta}\|_p^p \leq \|\hat{\mathbf{x}} + \boldsymbol{\delta}\|_p^p \leq \|\hat{\mathbf{x}}\|_p^p + \|\boldsymbol{\delta}\|_p^p \tag{11}$$

We again split $\mathbf{x}$ and $\boldsymbol{\delta}$ according to indices on and off the support of $\mathbf{x}$, and similarly to the proof for ℓ_1, using (11), we obtain:

$$\sum_{i \in \mathcal{I}_x^C} |\delta_i|^p - \sum_{i \in \mathcal{I}_x} |\delta_i|^p > 0 \tag{12}$$

Finally, considering the worst case scenario for $\mathbf{x}$ and $\boldsymbol{\delta}$ (see [2]), we have $\sum_{i \in \mathcal{I}_x^C} |\delta_i|^p - \sum_{i \in \mathcal{I}_x} |\delta_i|^p \geq S(\mathbf{A})^p (K(\mathbf{A}) + 1 - L) - L$, which leads to the desired result in Theorem 5.

Now let us analyze the new condition. Since $S(\mathbf{A})$ is below unity except for degenerate cases, for $p < 1$ the bound is less restrictive than for $p = 1$. As p goes to zero, the condition approaches $(K(\mathbf{A}) + 1)/2$, as long as $S(\mathbf{A})$ is non-zero. But this is the necessary and sufficient condition for uniqueness of solutions to the ℓ_0 problem which we have derived in Theorem 1! Thus the ℓ_p problem with small p is a very good approximation of the ℓ_0 problem, which is quite intuitive.

5. SIGN PATTERNS OF EXACT SOLUTIONS

Now we describe an interesting observation that we have made regarding the optimal solutions to the ℓ_1 problem, (2). We consider the case where the sufficient condition in Theorem 4 is not met. Hence, some ℓ_0-optimal solutions are also ℓ_1-optimal, while others are not. We would like to characterize these two sets of $\mathbf{x}$'s. We show that all ℓ_1-optimal solutions $\hat{\mathbf{x}}$ which are also ℓ_0-optimal can be determined by considering a finite number of test cases, according to the supports and sign patterns of $\hat{\mathbf{x}}$. We define the sign pattern of $\mathbf{x}$ to be a vector $\mathbf{s} \in \mathbb{R}^L$, where $\|\mathbf{x}\|_0^0 = L$, containing the signs of the elements of $\mathbf{x}$ on the support of $\mathbf{x}$. Using this definition, we get the following theorem:

Theorem 6 (Sign patterns of solutions) *Let $\mathbf{A} \in \mathbb{R}^{M \times N}$, $M < N$, and $K(\mathbf{A}) = M$. Suppose that $\|\hat{\mathbf{x}}\|_0^0 < \frac{M+1}{2}$, and $\hat{\mathbf{x}}$ is the optimal solution to both the ℓ_0 and ℓ_1 problems, for a given $\mathbf{y}$, i.e.*

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{y}=\mathbf{A}\mathbf{x}} \|\mathbf{x}\|_0^0, \textit{ and } \hat{\mathbf{x}} = \arg \min_{\mathbf{y}=\mathbf{A}\mathbf{x}} \|\mathbf{x}\|_1 \tag{13}$$

Then, if $\tilde{\mathbf{x}}$ has the same support and sign pattern as $\hat{\mathbf{x}}$, then:

$$\tilde{\mathbf{x}} = \arg \min_{\tilde{\mathbf{y}}=\mathbf{A}\mathbf{x}} \|\mathbf{x}\|_0^0 \textit{ and } \tilde{\mathbf{x}} = \arg \min_{\tilde{\mathbf{y}}=\mathbf{A}\mathbf{x}} \|\mathbf{x}\|_1, \textit{ where } \tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}}$$

Proof outline. The support of $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ is the same, so $\|\tilde{\mathbf{x}}\|_0^0 = \|\hat{\mathbf{x}}\|_0^0 < \frac{M+1}{2}$. By Theorem 1, since $\tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}}$, and $\tilde{\mathbf{x}}$ is sparse, $\tilde{\mathbf{x}}$ is the optimal solution to its ℓ_0 problem. Also, since $\hat{\mathbf{x}}$ is the ℓ_1-optimal solution, then $\|\hat{\mathbf{x}}\|_1 < \|\hat{\mathbf{x}} + \mathbf{n}\|_1, \forall \mathbf{n} \in Null(\mathbf{A}), \mathbf{n} \neq \mathbf{0}$. It remains to show that since $\tilde{\mathbf{x}}$ has the same sign pattern and support as $\hat{\mathbf{x}}$ then $\|\tilde{\mathbf{x}}\|_1 < \|\tilde{\mathbf{x}} + \mathbf{n}\|_1, \forall \mathbf{n} \in Null(\mathbf{A}), \mathbf{n} \neq \mathbf{0}$. Suppose this is false, and for $\tilde{\mathbf{x}}$, its ℓ_1 problem has a better solution:

$$\exists \tilde{\mathbf{n}} \in Null(\mathbf{A}), \tilde{\mathbf{n}} \neq \mathbf{0}, \text{ such that } \|\tilde{\mathbf{x}} + \tilde{\mathbf{n}}\|_1 \leq \|\tilde{\mathbf{x}}\|_1 \tag{14}$$

We introduce some notation for indexing: let "$\mathcal{I}_{on}$" be the set of indices on the support of $\tilde{\mathbf{x}}$, and $\mathcal{I}_{off} = \mathcal{I}_{on}^C$, indices off the support of $\tilde{\mathbf{x}}$. Let "$\mathcal{I}_{ss}$" be the indices on the support of $\tilde{\mathbf{x}}$ where $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{n}}$ have the same sign, and "$\mathcal{I}_{ds}$" be the indices on the support of $\tilde{\mathbf{x}}$ where $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{n}}$ have different signs. After some algebra, splitting the vectors according to these indices, we have:

$$\sum_{i \in \mathcal{I}_{ds}} |\tilde{n}_i| \geq \sum_{i \in \mathcal{I}_{off}\, \cup\, \mathcal{I}_{ss}} |\tilde{n}_i| \tag{15}$$

Using this inequality we now find a contradiction to the ℓ_1-optimality of $\hat{\mathbf{x}}$. Since $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ have the same support and the same sign pattern, then $\mathcal{I}_{on}, \mathcal{I}_{off}, \mathcal{I}_{ss}$, and $\mathcal{I}_{ds}$ have the same meaning for $\hat{\mathbf{x}}$ as for $\tilde{\mathbf{x}}$. That is to say, $\mathcal{I}_{on}$ is the support of $\tilde{\mathbf{x}}$ and also of $\hat{\mathbf{x}}$. Signs of $\tilde{\mathbf{n}}$ and $\tilde{\mathbf{x}}$ are the same on $\mathcal{I}_{ss}$, and so are the signs of $\tilde{\mathbf{n}}$ and $\hat{\mathbf{x}}$.

Let $\check{\mathbf{n}} = \alpha\tilde{\mathbf{n}}$, where $\alpha > 0$ is selected such that $|\check{n}_i| < |\hat{x}_i|$ for all $i \in \mathcal{I}_{ds}$. This is possible since $|\hat{x}_i| > 0$ for all $i \in \mathcal{I}_{ds}$. Rewriting ℓ_1-optimality of $\hat{\mathbf{x}}$ in terms of our indices, it follows that $\forall \mathbf{n} \in Null(\mathbf{A}), \mathbf{n} \neq \mathbf{0}$ we must have $\sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i| - \sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i + n_i| < \sum_{i \in \mathcal{I}_{off}\, \cup\, \mathcal{I}_{ss}} |n_i|$. However, for our particular $\check{\mathbf{n}} \in Null(\mathbf{A})$:

$$\sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i| - \sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i + \check{n}_i| = \sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i| - \sum_{i \in \mathcal{I}_{ds}} |\hat{x}_i| + \sum_{i \in \mathcal{I}_{ds}} |\check{n}_i| = \sum_{i \in \mathcal{I}_{ds}} |\check{n}_i| \geq \sum_{i \in \mathcal{I}_{off}\, \cup\, \mathcal{I}_{ss}} |\check{n}_i| \tag{16}$$

This is a contradiction, thus $\tilde{\mathbf{x}}$ is the optimal solution to $\min \|\mathbf{x}\|_1$ subject to $\tilde{\mathbf{y}} = \mathbf{A}\mathbf{x}$. This proves our theorem about sign patterns.

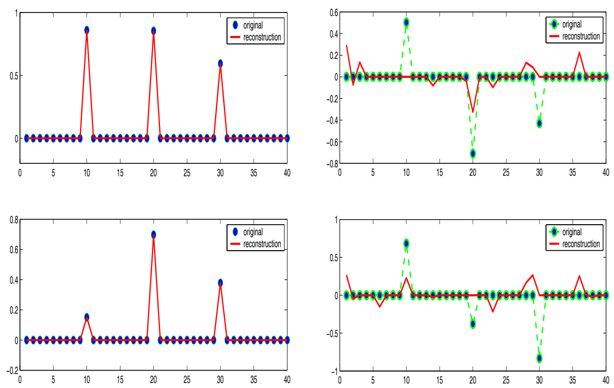

(a) (b)

Fig. 1. Sign patterns of exact solutions. Matrix $\mathbf{A}$ is 10×40. Sparsity profile (support of $\mathbf{x}$) is (10, 20, 30). (a) Sign pattern for exact solutions: $(+, +, +)$. Two signals along with exact reconstructions. (b) Sign pattern for wrong solutions: $(+, -, -)$. Two signals, and the corresponding wrong reconstructions.

An example of sign patterns of solutions appears in Figure 1. In (a), two signals sharing a sign pattern both equal their ℓ_1 reconstructions. In (b), two different signals sharing another sign pattern both yield incorrect reconstructions. In summary, whether or not the ℓ_1 reconstruction will equal the original signal depends on the support and the sign pattern, and not on the signal amplitudes.

6. SPARSITY AFTER A SPECIFIED TRANSFORMATION

Now we consider a more general problem, for a given $\mathbf{D} \in \mathbb{C}^{J \times N}$:

$$\min \|\mathbf{D}\mathbf{x}\|_p^p \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x} \tag{17}$$

where $0 \leq p \leq 1$. This means that $\mathbf{x}$ does not have to be sparse, but we would like the vector $\mathbf{D}\mathbf{x}$ to be sparse. We show how to reduce (17) to the problem of representation with sparsity in the standard basis. As a corollary, this reduction establishes the results of Sections 2, 3 and 4 for the new problem in (17).

For the case where $J = N$, and $\mathbf{D}$ is invertible, this is trivial. Let $\mathbf{z} = \mathbf{D}\mathbf{x}$. Then $\mathbf{x} = \mathbf{D}^{-1}\mathbf{z}$, and the problem (17) can be rewritten as $\min \|\mathbf{z}\|_p^p$ such that $\mathbf{y} = \mathbf{A}\mathbf{D}^{-1}\mathbf{z}$. Let $\tilde{\mathbf{A}} = \mathbf{A}\mathbf{D}^{-1}$, then this is exactly in the form of (2).

The case where $\mathbf{D}$ is rectangular, full-row rank and has a nullspace (e.g. when $\mathbf{D}$ is a pairwise difference operator for Total Variation (TV)) is more interesting. In order to have a unique solution, we must have $Null(\mathbf{A}) \cap Null(\mathbf{D}) = \{\mathbf{0}\}$, which we now assume. Using the same definition of $\mathbf{z}$, we can rewrite the problem as

$$\min \|\mathbf{z}\|_p^p \text{ such that } (\exists \boldsymbol{\delta} \in Null(\mathbf{D}) : \mathbf{y} = \mathbf{A}(\mathbf{D}^\dagger \mathbf{z} + \boldsymbol{\delta})) \tag{18}$$

Let us take a basis $\mathbf{N}$ for $Null(\mathbf{D})$, then this is equivalent to $\min \|\mathbf{z}\|_p^p$ subject to $\exists \boldsymbol{\eta}$ such that $\mathbf{y} = \mathbf{A}(\mathbf{D}^\dagger \mathbf{z} + \mathbf{N}\boldsymbol{\eta})$. The set $\{\mathbf{A}\mathbf{N}\boldsymbol{\eta} \mid \boldsymbol{\eta} \in \mathbb{C}^{N-J}\}$ forms a subspace, thus denoting $\boldsymbol{\Phi} = \mathbf{A}\mathbf{N}$, we can split $\mathbf{y}$ into two parts: project $\mathbf{y}$ onto the range space of $\boldsymbol{\Phi}$, $\mathbf{y}^\parallel = \Pi_{\boldsymbol{\Phi}}\mathbf{y}$, and on the orthogonal complement, $\mathbf{y}^\perp = \Pi_{\boldsymbol{\Phi}}^\perp \mathbf{y}$. As long as we can find $\mathbf{z}$ to approximate $\mathbf{y}$ in the subspace orthogonal to the range space of $\boldsymbol{\Phi}$, then we can always find $\boldsymbol{\eta}$ to represent the residual in the range space of $\boldsymbol{\Phi}$. That is to say, we need to solve

$$\hat{\mathbf{z}} = \arg\min \|\mathbf{z}\|_p^p \text{ subject to } \mathbf{y}^\perp = \Pi_{\boldsymbol{\Phi}}^\perp \mathbf{A}\mathbf{D}^\dagger \mathbf{z} \tag{19}$$

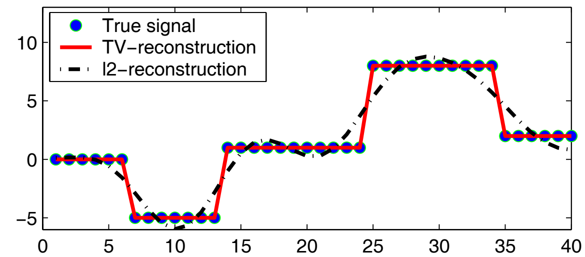

Fig. 2. Total variation reconstruction of a piecewise constant signal. Matrix $\mathbf{A}$ is 10×40, and $\mathbf{D}$ is a 39×40 pairwise difference operator. The original signal, and the ℓ_1-TV reconstruction match exactly. The ℓ_2 reconstruction, however blurs the edges, and does not recover the original signal.

and then we find $\hat{\boldsymbol{\eta}} = \boldsymbol{\Phi}^\dagger(\mathbf{y} - \mathbf{A}\mathbf{D}^\dagger \hat{\mathbf{z}})$. To find $\hat{\mathbf{x}}$ we put the two components back together: $\hat{\mathbf{x}} = \mathbf{D}^\dagger \hat{\mathbf{z}} + \mathbf{N}\hat{\boldsymbol{\eta}}$. Now the theoretical results of preceding sections can be directly applied to the problem in (17). The transformed conditions require checking if $\mathbf{z}$ is sparse enough with respect to $\tilde{\mathbf{A}} = \Pi_{\boldsymbol{\Phi}}^\perp \mathbf{A}\mathbf{D}^\dagger$. An example of exact reconstruction for the ℓ_1-case with $\mathbf{D}$ being a pairwise difference operator appears in Figure 2. The ℓ_1-TV reconstructed signal matches exactly the original signal. For comparison, we plot the ℓ_2 solution, obtained by setting $p = 2$ in (17), which does not favor sparsity. The ℓ_2 solution has blurred edges, and does not match the original signal.

7. CONCLUSION

We have presented theoretical analysis justifying the use of ℓ_1 and ℓ_p approximations to the problem of sparse signal representation in *general* overcomplete bases. If signals are sufficiently sparse, these approximations lead to exact solutions, and if not, we have characterized subsets for which ℓ_1-ℓ_0 equivalence holds. Finally, we have extended the results to sparsity in a transformed domain.

8. REFERENCES

[1] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.

[2] D. M. Malioutov, "A sparse signal reconstruction perspective for source localization with sensor arrays," M.S. thesis, MIT, EECS, July 2003.

[3] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

[4] D. L. Donoho and M. Elad, "Maximal sparsity representation via l1 minimization," *submitted to IEEE Trans. Inf. Theory*, 2002, Stanford Univ. Tech. report: http://www-sccm.stanford.edu/pub/sccm/sccm02-17.pdf.

[5] R. Gribonval and M. Nielsen, "Sparse representation in unions of bases," Nov. 2002, IRISA Publication interne no. 1499.

[6] J.J. Fuchs, "On sparse representations in arbitrary redundant bases," *submitted to IEEE Trans. Inf. Theory*, Dec. 2002.

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