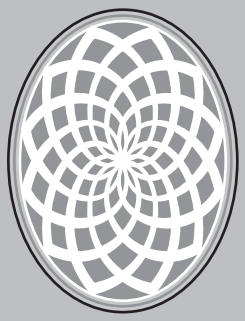

Industrial Mathematics Institute

2004:06

Stable recovery of sparse overcomplete representations in the presence of noise

D.L. Donoho, M. Elad and V. Temlyakov

IMI

Preprint Series

Department of Mathematics
University of South Carolina

Stable Recovery of Sparse Overcomplete Representations in the Presence of Noise

David L. Donoho* Michael Elad† Vladimir Temlyakov‡

Abstract

Overcomplete representations are attracting interest in signal processing theory, particularly due to their potential to generate sparse representations of signals. However, in general, the problem of finding sparse representations must be unstable in the presence of noise. We prove the possibility of stable recovery under a combination of sufficient sparsity and favorable structure of the overcomplete system.

We consider the situation where an ideal underlying signal indeed has a sparse representation but we can observe only a noisy version. We assume that the overcomplete system has a property of *mutual incoherence*, and that the ideal underlying signal has a sufficiently sparse representation, according to a sparsity threshold defined using the mutual coherence of the overcomplete system. Under these conditions, we show that the optimally–sparse approximation to the noisy data, to within the noise level, differs from the optimally–sparse decomposition of the ideal noiseless signal by at most a constant multiple of the noise level.

In general, this optimal-sparsity method requires heavy computational effort (e.g. brute-force combinatorial optimization). However, we show that stable reconstruction is also available by solving a convex quadratic programming problem. In this approach, the sparsity objective is replaced by the ℓ^1 objective; one finds the approximate representation whose coefficients have the minimal ℓ^1 norm while fitting the data to within the noise level. This explains certain successes of Basis Pursuit in signal processing and Lasso in statistical modeling. We also consider greedy stepwise least-squares and show that, when stopped at the point where the size of the residual equals the noise level, this has at least a *local* stability property. This explains certain successes of Matching Pursuit in signal processing and Stepwise Regression in statistical modeling.

These methods can also be applied with an exaggerated noise tolerance. When done properly, the resulting sparse approximation of the noisy data will actually contain only 'correct' nonzeros, i.e. only terms also appearing in the unique sparsest representation of the ideal noiseless sparse signal.

Keywords. Sparse representation, Overcomplete Representation, Incoherent Dictionary, Stability, Mixed ℓ^1-ℓ^2 norm minimization, ℓ^1-norm penalization, Basis Pursuit, Matching Pursuit, Lasso, Greedy Approximation, Stepwise Regression, Kruskal rank, Superresolution.

Acknowledgements. Partial support from NSF DMS 00-77261, 01-40698 (FRG), 02-00187, NSF ANI-008584 (ITR), DARPA ACMP, and ONR-MURI. We would like to thank Matthew Ahlert and Inam Rahman for help in preparing the figures.

*Department of Statistics, Stanford University, Stanford 94305-9025 CA. USA.

†Department of Computer Science, Technion, Haifa 32000, Israel

‡Department of Mathematics, University of South Carolina, Columbia, SC 29208 USA.

1 Introduction

1.1 Overcomplete Representation

Researchers spanning a diverse range of viewpoints have recently advocated the use of *overcomplete* signal representations [23, 26, 1, 5, 4, 30, 34, 32]. Generally speaking, they suppose we have a signal vector $y \in R^n$, and a collection of vectors $\phi_i \in R^n$, $i = 1, ..., m$, with $m > n$ such vectors, so that the collection forms 'more than a basis'; since [23] such collections are usually called *dictionaries*, and their elements are called *atoms*. We want a representation for our signal $y = \sum_i \alpha_i \phi_i$ as a linear combination of atoms in this dictionary.

Such representations differ from the more traditional basis representation because they offer a wider range of generating elements; potentially, this wider range allows more flexibility in signal representation, and more effectiveness at tasks like signal extraction and data compression. Proposals for overcomplete representations have included multiscale Gabor functions [23, 26], systems defined by algebraic codes [30], amalgams of wavelets and sinusoids [3, 4, 14], libraries of windowed cosines with windows of a range of different widths and locations [5, 40], multiscale windowed ridgelets [29], systems generated at random [12], and amalgams of wavelets and linelike generating elements [19].

A number of interesting arguments, both heuristic and theoretical, have been advanced to support the benefits of overcompleteness; in theoretical neuroscience it has been argued that overcomplete representations are probably necessary for use in biological settings in the mammalian visual system [24]; in approximation theory, there are persuasive examples where approximation from overcomplete systems outperforms any known basis [2]; in signal processing, it has been reported that decomposition into separate transforms gives improved compression [1, 9] and improved equalization [7]; and in image processing, it has been shown that one can separate images into disjoint signal types using such decompositions [28, 29, 19].

At the same time, there is an apparent obstacle to overcomplete representations, based on elementary linear algebra. We can think of the atoms in our dictionary as columns in a matrix $\mathbf{\Phi}$, so that $\mathbf{\Phi}$ is n by m and $m > n$. A representation of $y \in R^n$ can be thought of as a vector $\alpha \in R^m$ satisfying $y = \mathbf{\Phi}\alpha$. However, linear algebra tells us that because $m > n$, the problem of representation is underdetermined. Hence, as is widely taught in elementary courses, there is no unique solution to the representation problem, and far more disturbingly, if the data are even slightly inaccurate, some familiar algorithms will be staggeringly unstable. That this can be a real issue was shown by Wohlberg [41] who considered a dictionary of sinusoids with frequencies spaced finer than the usual discrete Fourier frequencies, and documented the extreme ill-posedness that can result.

In this article we consider the impact of sparsity constraints on this situation, and study algorithms which can in certain circumstances generate sparse representations in an overcomplete dictionary. We derive rigorous bounds showing that, when the dictionary $\mathbf{\Phi}$ has a property of *mutual incoherence* (defined below), and when it offers a sufficiently sparse representation for the ideal noiseless signal, the algorithms are locally stable, i.e. under addition of small amounts of noise, the algorithms recover the ideal sparse representation with an error that grows at most proportionally to the noise level. Some of the algorithms are even globally stable, i.e. they recover the ideal noiseless reconstruction with an error at worst proportional to noise level even under the addition of arbitrary amounts of noise. Under sufficient sparsity the constants of proportionality are very reasonable.

In short we show that, although the problem of recovering the underlying overcomplete representation is admittedly very ill-posed in general, when the underlying representation is sparse, and the dictionary is incoherent, the ill-posedness can disappear.

1.2 Sparse Representation

To fix ideas, consider the problem of finding the sparsest representation possible in an overcomplete dictionary $\mathbf{\Phi}$. As a measure of sparsity of a vector α, we take the so-called ℓ^0 norm $\|\alpha\|_0$, which is simply the number of non-zero elements in α. The sparsest representation is then the solution to the optimization problem

$$(P_0): \quad \min_{\alpha} \ \|\alpha\|_0 \ \text{ subject to } \ y = \mathbf{\Phi}\alpha. \tag{1.1}$$

As stated, this problem seems to be a general combinatorial optimization problem, requiring that one enumerate all possible k-element collections of columns of $\mathbf{\Phi}$, for $k = 1, 2, \ldots, m$, looking for the smallest collection permitting representation of the signal. Such an algorithm would cost at least $O(2^m)$ flops to carry out in general, and at least $O(m^k)$ even when a sparse k-element representation existed. We therefore turn to approximations/relaxations of (P_0).

Orthogonal Greedy Algorithm. One heuristic approach builds up k-element approximate representations a step at a time, adding to an existing $(k-1)$-element approximation a new term chosen in a greedy fashion to minimize the ℓ^2 error in the resulting augmented approximation among all possible choices of the single additional term. When stopped after $N \ll m$ stages, one gets a sparse approximate representation. In more detail, the procedure starts from an initial residual $r^{(0)} = y$ and a current decomposition $\hat{y}^0 = 0$; then for $k = 1, \ldots,$ it augments the decomposition $\hat{y}^{(k-1)} \to \hat{y}^{(k)}$ and updates the residual $\hat{r}^{(k-1)} \to \hat{r}^{(k)}$ stepwise, always maintaining $y = \hat{y}^{(k)} + \hat{r}^{(k)}$. In more detail, we suppose that the dictionary has normalized atoms, so that each $\|\phi_i\|_2 = 1$. At the k-th stage, the algorithm selects an atom to be added to the decomposition based on correlation with the current residual

$$i_k = \text{argmax}_{1 \le i \le m} |\langle r^{(k-1)}, \phi_i \rangle|;$$

it builds a decomposition consisting of the atoms selected through that stage

$$\hat{y}^{(k)} = \sum_{l=1}^{k} a_{i_l}^k \phi_{i_l}, \tag{1.2}$$

where the coefficients $(a_{i_l}^k)$ are fitted by least squares to minimize $\|y - \hat{y}^{(k)}\|^2$; and it subtracts this model from y, obtaining a new residual

$$r^{(k)} = y - \hat{y}^k,$$

which can be input to the next stage of the algorithm. At that point, the basic iteration is repeated. The algorithm stops when the residual norm is below some predetermined threshold, representing a tolerated numerical inaccuracy.

In the setting of statistical modelling, greedy stepwise least-squares is called *forward stepwise regression*, and has been widely practiced since the 1960's [8, 18]. When used in the signal processing setting this goes by the name of *Matching Pursuit* (MP) [23]; actually we have described a variant called Orthonormal Matching Pursuit (OMP) [25]. Following [10], we call this the Orthogonal Greedy Algorithm (OGA).

ℓ^1 *Penalization.* A more formal approach convexifies (P_0) by replacing the ℓ^0-norm with an ℓ^1-norm:

$$(P_1): \quad \min_{\alpha} \ \|\alpha\|_1 \ \text{ subject to } \ y = \mathbf{\Phi}\alpha. \tag{1.3}$$

This can be cast as a Linear Programming (LP) problem, for which solutions are available even in large-scale problems, owing to modern interior-point linear programming methods. This approach

to overcomplete signal representation was called Basis-Pursuit (BP) in [4], which observed that it sometimes gave highly sparse solutions to problems known to have such sparse solutions, and showed that it could, in specific cases, outperform the greedy Matching Pursuit approach in generating sparse solutions.

Formal Justification. The key point about both MP and BP is that they are much more practical than direct solution of (P_0). Perhaps surprisingly, these approaches can, in certain conditions, correctly solve (P_0). Thus, practical methods can solve problems that otherwise on the surface seem computationally intractable. Previous work [13, 14, 12, 17] established that both the OGA and the BP approaches can be successful for signals having sparse representations; under appropriate conditions on $\mathbf{\Phi}$ and y, these algorithms produce the globally optimal solution of (P_0). The concept of *mutual coherence* of the dictionary $\mathbf{\Phi}$ plays a major role in these results. It is defined, assuming that the columns of $\mathbf{\Phi}$ are normalized to unit ℓ^2-norm, in terms of the Gram matrix $\mathbf{G} = \mathbf{\Phi}^T\mathbf{\Phi}$. With $G(k,j)$ denoting entries of this matrix, the mutual coherence is

$$M = M(\mathbf{\Phi}) = \max_{1\le k,j\le m,\ k\neq j} |G(k,j)|. \tag{1.4}$$

A dictionary is incoherent is M is small. There are overcomplete systems with $m \approx n^2$ and $M \approx 1/\sqrt{n}$ [30]. The results in [13, 14, 12, 17] showed that, if there exists a representation $y = \mathbf{\Phi}\alpha$ with sparsity $N = \|\alpha\|_0$, and N does not exceed a threshold $(1 + M^{-1})/2$ defined by M alone, then (a) this is the unique sparsest representation, and (b) these algorithms would find it. If, for example, $M = 1/\sqrt{n}$ this result promises, for large n, an ideal form of atomic decomposition even of fairly complex objects. In such cases, provided the object y is made from $< \sqrt{n}/2$ atoms in the dictionary, this sparse decomposition can be uniquely recovered.

1.3 Presence of Noise

In most practical situations it is not sensible to assume that the available data y obey precise equality $y = \mathbf{\Phi}\alpha$ with a sparse representation α. A more plausible scenario assumes *sparse approximate representation*: that there is an ideal noiseless signal x_0 with a sparse representation, $x_0 = \mathbf{\Phi}\alpha_0$ with $\|\alpha_0\|_0$ small, but that we can observe only a noisy version $y = x_0 + z$, where $\|z\|_2 < \epsilon$.

Noise-Aware Variant of (P_0). We can adapt to this noisy setting by modifying (P_0) to include a noise allowance:

$$(P_{0,\delta}): \quad \min_{\alpha} \ \|\alpha\|_0 \ \text{ subject to } \ \|y - \mathbf{\Phi}\alpha\|_2 \le \delta. \tag{1.5}$$

Note that $(P_{0,0}) \equiv (P_0)$. Also, if we apply this with $\delta > \epsilon = \|y - x\|_2$, the problem has a sparse solution; in fact the solution $\hat{\alpha}$ obeys $\|\hat{\alpha}\|_0 \le \|\alpha_0\|_0$, or more formally

$$val(P_{0,\delta}) \le val(P_0), \qquad \delta > \epsilon. \tag{1.6}$$

Noise-Aware Variant of OGA. Just as with (P_0), $(P_{0,\delta})$ demands exorbitant computational efforts in general, and so again we may resort to heuristics and relaxations. On the one hand, OGA can be employed for approximating the solution of (1.5); the stepwise procedure can simply be stopped when the representation error gets below δ.

Noise-Aware Variant of (P_1). On the other hand, we can pursue a strategy of convexification, replacing the ℓ^0-norm in (1.5) by an ℓ^1-norm,

$$(P_{1,\delta}): \quad \min_{\alpha} \ \|\alpha\|_1 \ \text{ subject to } \ \|y - \mathbf{\Phi}\alpha\|_2 \le \delta. \tag{1.7}$$

This can be cast as a convex quadratic program which can be solved by many standard approaches, including Iteratively-Reweighted Least Squares (IRLS) (Appendix B, [20]), interior-point algorithms [4], and active-set methods. It is also closely related to Basis Pursuit Denoising (BPDN) [4] and to the LASSO technique employed in machine learning to avoid over-fitting [31] when combining predictors; both of those proposals amount to solving a corresponding convex optimization in Lagrangian form

$$(P_{1',\lambda}): \quad \min_{\alpha} \quad \|\alpha\|_1 + \quad \|y - \mathbf{\Phi}\alpha\|_2^2/\lambda; \tag{1.8}$$

for appropriate $\lambda = \lambda(y, \delta)$ the solutions of $(P_{1,\delta})$ and $(P_{1',\lambda(y,\delta)})$ are the same.

1.4 Stability Properties

In this paper we develop several results exhibiting stable recovery of sparse representations in the presence of noise. We now briefly sketch their statements.

First, we show that when sufficient sparsity is present, where 'sufficient' is defined relative to the degree of mutual coherence, solving $(P_{0,\epsilon})$ enables stable recovery. We suppose that we have a possibly overcomplete system $\mathbf{\Phi}$ with mutual coherence $M = M(\mathbf{\Phi})$. Suppose that we have a noisy signal y and that the ideal noiseless signal x_0 has a sparse representation α_0 with at most N nonzeros. Then if $N < (M^{-1} + 1)/2$, it follows both that (a) α_0 is the unique sparsest representation of x_0, and (b) the solution $\hat{\alpha}_{0,\epsilon}$ of $(P_{0,\epsilon})$ obeys

$$\|\hat{\alpha}_{0,\epsilon} - \alpha_0\|_2 \le \Lambda_0(M, N) \cdot \epsilon; \quad \forall \epsilon > 0 \tag{1.9}$$

here the stability coefficient $\Lambda_0(M, N)^2 \le 4/(1 - M(2N - 1))$. The proportionality constant Λ_0 (which we also call the *stability coefficient*) can be quite moderate given sufficient sparsity. For example, if the sparsity measure $N \le 1/(4M)$, we get a stability coefficient ≤ 3. In words, provided the underlying object is sparsely represented and the noise level is known, recovery by explicitly imposing sparsity yields an approximation to the ideal sparse decomposition of the noiseless signal in which the error is at worst proportional to the input noise level.

Next, we develop a parallel result for ℓ^1 minimization. Making parallel assumptions, tightened so that the ideal noiseless signal x_0 has a sparse representation α_0 with $N \le (M^{-1} + 1)/4$, we show both that (a) α_0 is the unique sparsest representation of x_0, and (b) the solution $\hat{\alpha}_{1,\epsilon}$ of $(P_{1,\epsilon})$ obeys

$$\|\hat{\alpha}_{1,\epsilon} - \alpha_0\|_2 \le \Lambda_1(M, N) \cdot \epsilon, \quad \forall \epsilon > 0, \tag{1.10}$$

where the stability coefficient $\Lambda_1(M, N)^2 \le 4/(1 - M(4N - 1))$. In words, ℓ^1-based reconstruction in incoherent overcomplete systems has an error which is at worst proportional to the input noise level. The sparsity requirement is twice as stringent for the ℓ^1-based result as for the ℓ^0-based result.

By comparison, OGA obeys a local stability result. Again suppose a possibly overcomplete system with $M = M(\mathbf{\Phi})$, and an ideal noiseless signal x_0 having a representation with at most N atoms. Suppose that the smallest among the N nonzeros in the representation of x_0 has amplitude A. Assume that we know the noise level $\epsilon = \|y - x_0\|_2$ and run the OGA just until the representation error $\le \epsilon$. Call the result of this greedy algorithm $\hat{\alpha}_{OGA,\epsilon}$. Set $\Lambda_{OGA}(M, N) = (1 - M(N - 1))^{-1/2}$. Then if

$$\epsilon \le \epsilon_0 = \epsilon_0(M, N, A) \equiv A/2(1 - M(2N - 1))^{1/2} \tag{1.11}$$

the recovered representation $\hat{\alpha}_{OGA,\epsilon}$ obeys

$$\|\hat{\alpha}_{OGA,\epsilon} - \alpha_0\|_2 \le \Lambda_{OGA}(M, N) \cdot \epsilon. \tag{1.12}$$

This is a local stability result because for large values of $\epsilon = \|y - x_0\|_2$ the condition (1.11) will necessarily fail.

Note the parallel nature of the bounds and the conclusions. Three quite different algorithms all obey stability results in which having N a fraction of M^{-1} is the key assumption.

1.5 Support Properties

A fundamental question about efforts to obtain sparse representation: *do we actually recover the* **correct** *sparsity pattern*? (Our stability results do not address this question, since it is possible for a nonsparse representation to be close to a sparse representation in an ℓ^2 sense.)

The question is fundamental and broadly significant. Throughout science and technology, it is habitual to fit sparse models to noisy data, and then simply *assume* that terms appearing in such fitted models are dependable features.

In this paper, we are able to shed some light on this situation. Our results show that, under appropriate conditions, the empirical representation $\hat{\alpha}$ is not only at least as sparse as the ideal sparse representation but *it only contains atoms also appearing in the ideal sparse representation*. Since that ideal sparse representation is, by our other results, unique and well-defined, these insights endow the empirical support of $\hat{\alpha}$ with a, perhaps surprising, significance.

Our first result is obtained in the course of analyzing OGA; it shows that, under condition (1.11) and $N < M^{-1}/2$, the ideal noiseless representation is unique, and the support of $\hat{\alpha}_{OGA}$ is contained in the support of α_0.

Our second result concerns solution of $(P_{1,\delta})$ with $\delta = C \cdot \epsilon$. Here $C = C(M, N) > 1$, and so we are solving the ℓ^1 minimization problem using an exaggeration of the noise level. It shows, with $M = M(\mathbf{\Phi})$ and $\|\alpha_0\|_0 \leq N$, that the solution $\hat{\alpha}_{1,\delta}$ has its support contained in the support of α_0. Here $C \approx 3\sqrt{N}$ if $MN \leq 1/2$, so ordinarily it requires considerable overstatement of the noise level to achieve this level of conservatism. However, it does provide the very interesting *epistemological* benefit that the atoms appearing in the representation have more than incidental meaning.

1.6 Contents

The next three sections supply the analysis behind the stability bounds just quoted. Then comes a discussion of support properties. Later sections extend our work in several ways:

- *Variants of the Greedy Algorithm.* We give local stability results for the non-orthogonal Pure Greedy algorithm, and the Weak Orthogonal Greedy Algorithm.
- *Numerical Results.* We study the actual stability behavior of the ℓ^1 and OGA on synthetic examples, finding typical behavior far more favorable than our theoretical bounds.
- *Approximation Bounds.* We also study the behavior of the Greedy algorithm and its variants in an approximation setting, where we are trying to approximate y rather than represent x_0, sharpening and extending previously known results of Temlyakov [32, 33], Gilbert at al. [15] and Tropp [36].
- *Superresolution.* We situate our work with respect to the problem of superresolution, in which astronomers, seismologists, spectroscopists and others attempt to 'deblur' sparse spike trains.
- *Geometry.* We develop a geometric viewpoint explaining why stability can sometimes be expected for the ℓ^1 penalization scheme, under conditions of sufficient sparsity.

An Appendix develops ties with the literature on Greedy Approximation, gives alternate proofs of results in Section 2, and proofs of several results stated in the body of the paper.

We have recently learned that in parallel to our efforts, J.A. Tropp has been working independently on some of the same problems [38]. After some recent discussions and a careful study we find that his methods and results have a rather different flavor, ensuring that both works are of interest in studying sparse approximation under noise

2 Stability Using $(P_{0,\epsilon})$

Suppose again the existence of an ideal noiseless signal $x_0 = \mathbf{\Phi}\alpha_0$ and noisy observations $y = x_0 + z$ with $\|y - x_0\|_2 \leq \epsilon$. Consider applying $(P_{0,\delta})$ with $\delta = \epsilon$ to obtain a sparse approximation to y. The following establishes the stability estimate mentioned in the introduction.

Theorem 2.1 *Let the dictionary $\mathbf{\Phi}$ have mutual coherence $M = M(\mathbf{\Phi})$. Suppose the noiseless signal $x_0 = \mathbf{\Phi}\alpha_0$, where α_0 satisfies*

$$\|\alpha_0\|_0 = N < (1/M + 1)/2. \tag{2.1}$$

Then:

(a) α_0 is the unique sparsest such representation of x_0; and

(b) the reconstruction $\hat{\alpha}_{0,\epsilon}$ from applying $(P_{0,\epsilon})$ to the noisy data y approximates α_0:

$$\|\hat{\alpha}_{0,\epsilon} - \alpha_0\|_2^2 \leq \frac{4\epsilon^2}{1 - M(2N-1)}\ , \qquad \forall \epsilon > 0. \tag{2.2}$$

2.1 Uniqueness of the Target

Our proof of Theorem 2.1 goes in two stages. The first stage considers Claim (a), that α_0 is the unique sparsest representation of x. This claim confers special status on α_0 as a distinguished target for reconstruction. This claim actually follows from known results in [13, 14, 12, 17], which we now sketch because they provide a foundation for the second stage of the proof. (An alternate proof for Claim (a) is provided in the Appendix, using a greedy algorithms approach).

Definition 2.2 [12] *Given an $n \times m$ matrix $\mathbf{A}$, let $\sigma = Spark(\mathbf{A})$ be the largest $\sigma \leq n + 1$ such that every $\sigma - 1$ columns from $\mathbf{A}$ are linearly independent.*

After [12] appeared, we learned that Kruskal [21] worked with a related notion in the context of fitting trilinear arrays by simple models; this notion has later been called the Kruskal rank [42, 22]. We stick with the new terminology in this paper, because we will extend it to a broader setting in Section 2.2.

Several properties of *Spark* were given in [12]. First, the mutual coherence M can control *Spark*:

Lemma 2.3 *If the columns of $\mathbf{A}$ have unit ℓ^2 norm, then $Spark(\mathbf{A}) \geq 1/M + 1$, where $M = M(\mathbf{A})$ is the mutual coherence (1.4).*

Second, *Spark* in turn controls the sparsity of solutions of linear equations:

Lemma 2.4 *If $u = \mathbf{A}v_1$ and $u = \mathbf{A}v_2$, then $\|v_1\|_0 + \|v_2\|_0 \geq Spark(\mathbf{A})$.*

Third, *Spark* gives a criterion for uniqueness of sparse solutions to linear equations:

Lemma 2.5 *If $u = \mathbf{A}v_0$ and $\|v_0\|_0 < Spark(\mathbf{A})/2$, then v_0 is the unique sparsest possible solution to the equation $u = \mathbf{A}v$.*

Combining these facts, the sparsity assumption (2.1) made in the statement of Theorem 2.1 implies that $\|\alpha_0\|_0 < \mathrm{Spark}(\mathbf{\Phi})/2$, and hence α_0 is the unique sparsest solution of the equations $x = \mathbf{\Phi}\alpha$. Hence Claim (a) of Theorem 2.1 is proven.

2.2 Stability of the Reconstruction

We now turn to Claim (b) of Theorem 2.1. We know two proofs. One is given in Lemma A.2 of the Appendix. The other, developed here, extends the notion of *Spark*.

Definition 2.6 *Given a matrix $\mathbf{A}$, let $\sigma = Spark_\eta(\mathbf{A})$ be the largest $\sigma \leq n+1$ such that every $\sigma - 1$ columns from $\mathbf{A}$ form a matrix with smallest singular value $> \eta$.*

For $\eta = 0$ we have $Spark_0 = Spark$ since having all singular values > 0 is equivalent to linear independence of the columns. To our knowledge formalization of this concept may be new, however, we are aware that there are other signal processing problems where related notions are useful, compare [39].

We begin our discussion of $Spark_\eta$ by recording the obvious:

Lemma 2.7 *$Spark_\eta(\mathbf{A})$ is monotone decreasing in η, and $\forall\ \eta \geq 0$, $n+1 \geq Spark(\mathbf{A}) \geq Spark_\eta(\mathbf{A}) \geq 1$.*

We now turn to a more fundamental observation:

Lemma 2.8 *If $\mathbf{A}$ has normalized columns and Mutual Coherence $M = M(\mathbf{A})$, $Spark_\eta(\mathbf{A}) \geq (1 - \eta^2)/M + 1$.*

Proof. We need a simple observation about eigenvalues of diagonally-dominant matrices.

Lemma 2.9 *Given an s-by-s symmetric matrix $\mathbf{H}$ with diagonal entries equal to one and off-diagonal entries $\leq M$ in amplitude, the smallest eigenvalue of $\mathbf{H}$ is at least $1 - M(s-1)$.*

Indeed, if v is an eigenvector, suppose the largest absolute value in v is v^*, occurring at index i. Then from $(\mathbf{H} - \lambda I)v = 0$ we get $(1-\lambda)v_i = -\sum_{i \neq j} H_{ij}v_j$; using $|H_{ij}| \leq M$ for $i \neq j$ and $|v_j| \leq v^*$ we have

$$(1-\lambda)v^* \leq M(s-1)v^*.$$

It follows that $\lambda \geq 1 - M(s-1)$. □

With $\mathbf{A}$ obeying the stated assumptions, every leading minor $\mathbf{H}$ in the Gram matrix $\mathbf{G} = \mathbf{A}^H\mathbf{A}$ has ones on the diagonal because the columns of $\mathbf{A}$ are normalized. Also, its off-diagonal entries are bounded by M in absolute value, by definition of $M = M(\mathbf{A})$. Lemma 2.9 gives a bound on the minimum eigenvalue of every such $\mathbf{H}$ of size s by s, which implies that for $s \leq (1-\eta^2)/M + 1$, the minimum eigenvalue exceeds η^2. □

Lemma 2.10 *If $\|\mathbf{A}v\|_2 \leq \eta$ with $\|v\|_2 = 1$, then $\|v\|_0 \geq Spark_\eta(\mathbf{A})$.*

Proof. Suppose $\sigma = Spark_\eta(\mathbf{A})$ and assume that $\|v\|_0 = \sigma - 1$. The non-zero entries in v, pick out a subset J of $\sigma - 1$ columns in $\mathbf{A}$. Let $\mathbf{H}_J$ denote the corresponding leading minor of the Gram matrix and let v_J denote the corresponding vector just containing the nonzero entries in v. Then $\|\mathbf{A}v\|_2^2 = v_J^H \mathbf{H}_J v_J$. As $\eta^2 \geq \|\mathbf{A}v\|_2^2$ we conclude that $\mathbf{H}_J$ has an eigenvalue smaller than η^2. This contradicts the assumption that $\sigma = Spark_\eta(\mathbf{A})$. □

We have the following η-analog of Lemma 2.4.

Lemma 2.11 *If α_1 and α_2 satisfy $\|y - \mathbf{\Phi}\alpha_i\|_2 \leq \epsilon$, $i = 1, 2$, then*

$$\|\alpha_1\|_0 + \|\alpha_2\|_0 \geq Spark_\eta(\mathbf{\Phi}), \quad \eta = (2\epsilon)/\|\alpha_1 - \alpha_2\|_2. \tag{2.3}$$

Proof. The triangle inequality yields $\|\mathbf{\Phi}(\alpha_1 - \alpha_2)\|_2 \leq 2\epsilon$. Put differently, we have $\|\mathbf{\Phi}v\|_2 \leq \eta$, where $v = (\alpha_1 - \alpha_2)/\|\alpha_1 - \alpha_2\|_2$. ¿From Lemma 2.10, we get $\|v\|_0 \geq Spark_\eta(\mathbf{\Phi})$. But

$$\|\alpha_1\|_0 + \|\alpha_2\|_0 \geq \|\alpha_1 - \alpha_2\|_0 = \|v\|_0, \tag{2.4}$$

which establishes (2.3). □

For an η-analog of Lemma 2.5, we obtain, not uniqueness, but localization to a single Euclidean ball.

Theorem 2.12 *Given D and ϵ, set $\eta = 2\epsilon/D$. Suppose there are two approximate representations α_i, $i = 1, 2$ both obeying*

$$\|y - \mathbf{\Phi}\alpha_i\|_2 \leq \epsilon \quad \textit{and} \quad \|\alpha_i\|_0 < \frac{1}{2} Spark_\eta(\mathbf{\Phi}). \tag{2.5}$$

Then $\|\alpha_1 - \alpha_2\|_2 \leq D$.

Proof. Suppose that α_1 satisfies (2.5). Consider any α_2 at least D-separated from α_1. Set $\nu = 2\epsilon/\|\alpha_1 - \alpha_2\|$ and $\eta = 2\epsilon/D$, so $\eta \geq \nu$. Due to Lemmas 2.10 and 2.11 we know that

$$\|\alpha_1\|_0 + \|\alpha_2\|_0 \geq Spark_\nu(\mathbf{\Phi}) \geq Spark_\eta(\mathbf{\Phi}). \tag{2.6}$$

Thus, if $\|\alpha_1\|_0 < \frac{1}{2} Spark_\eta(\mathbf{\Phi})$, then $\|\alpha_2\|_0 > \frac{1}{2} Spark_\eta(\mathbf{\Phi})$. Hence α_2 cannot satisfy (2.5), and thus it must necessarily be in the ball of radius D around α_1. □

We are now in position to prove assertion (b) of Theorem 2.1. The solution to $(P_{0,\epsilon})$ is at least as sparse as the ideal sparse representation α_0, since by (1.6),

$$\|\hat{\alpha}_{0,\epsilon}\|_0 = val(P_{0,\epsilon}) \leq val(P_0) = \|\alpha_0\|_0.$$

For fixed $N = \|\alpha_0\|_0$, and arbitrarily small $\delta > 0$, let $\eta = \eta(M, N, \delta)$ satisfy

$$2N = \frac{1 - \eta^2}{M} + 1 + \delta.$$

By Lemma 2.8, $Spark_\eta(\mathbf{\Phi}) > 2N$. We conclude that

$$\|\hat{\alpha}_{0,\epsilon}\|_0 \leq \|\alpha_0\|_0 < \frac{1}{2} Spark_\eta(\mathbf{\Phi}).$$

At the same time, by definition of the noise level ϵ and the optimization problem $(P_{0,\epsilon})$,

$$\|y - \mathbf{\Phi}\hat{\alpha}_{0,\epsilon}\|_0 \leq \epsilon, \qquad \|y - \mathbf{\Phi}\alpha_0\|_0 \leq \epsilon.$$

Adapt the argument behind Theorem 2.12 to conclude that $\|\hat{\alpha}_{0,\epsilon} - \alpha_0\| \leq D = D(\eta, \epsilon) = 2\epsilon/\eta$. Setting $\eta = \eta(M, N, \delta)$, and substituting $D\eta = 2\epsilon$, we get an inequality like (2.2), only with right-hand side $\Lambda(M, N, \delta)^2\epsilon^2$ with coefficient

$$\Lambda(M, N, \delta)^2 = 4/(1 - M(2N - 1 - \delta)).$$

Letting $\delta \to 0$, we get inequality (2.2), and Theorem 2.1 follows. □

3 Stability Using $(P_{1,\epsilon})$

As in the introduction, we are given a signal $y = \mathbf{\Phi}\alpha_0 + z$, where z is an additive noise, known to satisfy $\|z\|_2 \leq \epsilon$. We apply $(P_{1,\delta})$ to this signal (not necessarily with $\delta = \epsilon$); i.e. we solve (1.7) and obtain a solution $\hat{\alpha}_{1,\delta,\epsilon}$. We study its deviation from the ideal representation α_0.

3.1 Stability Result

Theorem 3.1 *Let the overcomplete system $\mathbf{\Phi}$ have mutual coherence $M(\mathbf{\Phi})$. If some representation of the noiseless signal $x_0 = \mathbf{\Phi}\alpha_0$ satisfies*

$$N = \|\alpha_0\|_0 \leq (1/M + 1)/4, \tag{3.1}$$

then this is the unique sparsest representation of x_0; moreover, the deviation of the $(P_{1,\delta})$ representation from α_0 can be bounded by

$$\|\hat{\alpha}_{1,\delta,\epsilon} - \alpha_0\|_2^2 \leq \frac{(\epsilon + \delta)^2}{1 - M(4N - 1)} \ . \tag{3.2}$$

Proof. First, the assertion that α_0 is the unique sparsest representation follows from Theorem 2.1 and the fact that $\frac{1+M}{4M} < \frac{1+M}{2M}$.

Second, the stability bound can be posed as the solution to an optimization problem of the form:

$$\max_{\alpha_0,\ z} \ \|\hat{\alpha} - \alpha_0\|_2^2 \ \text{ subject to } \left\{ \begin{array}{c} \hat{\alpha} = \arg\min_\alpha \ \|\alpha\|_1 \ \text{ subject to } \ \|\mathbf{\Phi}\alpha - y\|_2 \leq \delta \\ y = \mathbf{\Phi}\alpha_0 + z, \quad \|z\|_2 < \epsilon, \quad \|\alpha_0\|_0 \leq N. \end{array} \right\}. \tag{3.3}$$

Put in words, we consider all representation vectors α_0 of bounded support, and all possible realizations of bounded noise, and we ask for the largest error between the ideal sparse decomposition and its reconstruction from noisy data. Defining $v = \alpha - \alpha_0$, and similarly $w = \hat{\alpha} - \alpha_0$, we can rewrite the above problem as

$$\max_{\alpha_0,\ z} \ \|w\|_2^2 \ \text{ subject to } \left\{ \begin{array}{cc} w = \arg\min_v & \|\alpha_0 + v\|_1 \ \text{ subject to } \ \|\mathbf{\Phi}v - z\|_2 \leq \delta \\ & \|z\|_2 \leq \epsilon, \quad \|\alpha_0\|_0 \leq N. \end{array} \right\}. \tag{3.4}$$

We develop an upper bound on $val(3.4)$ in a sequence of relaxations, each one expanding the feasible set and increasing the maximal value. To begin, note that if w is the minimizer of $\|\alpha_0 + v\|_1$ under these

constraints, then relaxing the constraints to all w satisfying $\|\alpha_0 + w\|_1 \leq \|\alpha_0\|$ expands the feasible set. Thus, we consider

$$\left\{ w \;\middle|\; \begin{array}{c} \|\alpha_0 + w\|_1 \leq \|\alpha_0\|_1 \;\&\; \|\mathbf{\Phi} w - z\|_2 \leq \delta \\ \|z\|_2 \leq \epsilon, \quad \|\alpha_0\|_0 \leq N \end{array} \right\} . \tag{3.5}$$

We now expand this set by exploiting the relation

$$\|\alpha_0 + w\|_1 - \|\alpha_0\| \geq \|w\|_1 - 2\sum_{k\in\mathcal{S}} |w(k)|,$$

where $\mathcal{S}$ is the support of the nonzeros in α_0 with complement $\mathcal{S}^c$, and we used $|a+b|-|a| \geq |a|-|b|-|a| = -|b|$. Therefore, we get a further increase in value by replacing the feasible set in (3.5) with

$$\left\{ w \;\middle|\; \begin{array}{c} \|w\|_1 \leq 2\sum_{k\in\mathcal{S}} |w(k)| \;, \quad \|\mathbf{\Phi} w - z\|_2 \leq \delta \\ \|z\|_2 \leq \epsilon, \quad \#\mathcal{S} \leq N \end{array} \right\} . \tag{3.6}$$

Writing this out yields a new optimization problem with still larger value:

$$\max_{w,\,\mathcal{S},\,z} \; \|w\|_2^2 \;\text{ subject to }\; \left\{ \begin{array}{c} \|w\|_1 \leq 2\sum_{k\in\mathcal{S}} |w(k)| \;, \quad \|\mathbf{\Phi} w - z\|_2 \leq \delta \\ \|z\|_2 \leq \epsilon, \quad \#\mathcal{S} \leq N \end{array} \right\}. \tag{3.7}$$

We next simplify our analysis by eliminating the noise vector z, using

$$\{w \mid \|\mathbf{\Phi} w - z\|_2 \leq \delta \;\;\&\;\; \|z\|_2 < \epsilon\} \subseteq \{w \mid \|\mathbf{\Phi} w\|_2 \leq \epsilon + \delta\}\,. \tag{3.8}$$

Expanding the feasible set of (3.7) using this observation gives

$$\max_{\mathcal{S},\,w} \; \|w\|_2^2 \;\text{ subject to }\; \left\{ \begin{array}{c} \|w\|_1 < 2\sum_{k\in\mathcal{S}} |w(k)| \;, \quad \|\mathbf{\Phi} w\|_2 \leq \Delta \\ \#\mathcal{S} \leq N \end{array} \right\}. \tag{3.9}$$

where we introduced $\Delta = \epsilon + \delta$.

The constraint $\|\mathbf{\Phi} w\|_2 \leq \Delta$ is not posed in terms of the absolute values in the vector w, complicating the analysis; we now relax this constraint using incoherence of $\mathbf{\Phi}$. Again the Gram matrix is $\mathbf{G} = \mathbf{\Phi}^T\mathbf{\Phi}$, and the mutual coherence is the maximal off-diagonal amplitude: $M = \max_{k\neq j} \;\; |G(k,j)|$. For a vector w, let $|w|$ be vector containing absolute values from w; similarly for matrices. Also, let $\mathbf{1}$ be the the m-by-m matrix of all ones, The constraint

$$\|\mathbf{\Phi} w\|_2 = w^T\mathbf{G} w \leq \Delta^2$$

can be relaxed:

$$\begin{array}{rcl} \Delta^2 \geq w^T\mathbf{G} w & = & \|w\|_2^2 + w^T\left(\mathbf{G} - \mathbf{I}\right) w \\ & \geq & \|w\|_2^2 - |w|^T\left|\mathbf{G} - \mathbf{I}\right||w| \\ & \geq & \|w\|_2^2 - M|w|^T\left|\mathbf{1} - \mathbf{I}\right||w| = (1+M)\|w\|_2^2 - M\|w\|_1^2. \end{array} \tag{3.10}$$

Using this, $val(3.9)$ is upper-bounded by the value

$$\max_{\mathcal{S},\,w} \; \|w\|_2^2 \;\text{ subject to }\; \left\{ \begin{array}{c} \|w\|_1 < 2\sum_{k\in\mathcal{S}} |w(k)| \\ (1+M)\|w\|_2^2 - M\|w\|_1^2 \leq \Delta^2 \\ \#\mathcal{S} \leq N \end{array} \right\}. \tag{3.11}$$

This problem is invariant under permutations of the entries in w which preserve membership in $\mathcal{S}$ and $\mathcal{S}^c$. It is also invariant under relabelling of coordinates. So assume that all non-zeros in α_0 are concentrated in the initial slots of the vector, i.e. that $\mathcal{S} = \{1, \ldots, N\}$.

Putting $w = (w_0,\ w_1)$ where w_0 gives the first N entries in w, and w_1 the remaining $m - N$ entries of w, we obviously have

$$\|w\|_2^2 = \|w_0\|_2^2 + \|w_1\|_2^2 \quad , \quad \|w\|_1 = \|w_0\|_1 + \|w_1\|_1. \tag{3.12}$$

The ℓ^1 norm on R^k dominates the ℓ^2 norm and is dominated by $\sqrt{k}$ times the ℓ^2 norm. Thus

$$\|w_0\|_1 \geq \|w_0\|_2 \geq \frac{\|w_0\|_1}{\sqrt{N}} \quad , \quad \|w_1\|_1 \geq \|w_1\|_2 \geq \frac{\|w_1\|_1}{\sqrt{m-N}}. \tag{3.13}$$

We define

$$A = \|w_0\|_1 \quad B = \|w_1\|_1 \quad c_0 = \left(\frac{\|w_0\|_2}{\|w_0\|_1}\right)^2 \quad c_1 = \left(\frac{\|w_1\|_2}{\|w_1\|_1}\right)^2. \tag{3.14}$$

Returning to the problem given in (3.11), and using our notations, we obtain a further reduction, from an optimization problem on R^m to an optimization problem on $(A, B, c_0, c_1) \in R^4$:

$$\begin{aligned} &\max \quad c_0 A^2 + c_1 B^2 \\ &\text{subject to} \quad \left\{ \begin{array}{c} A > B \\ (1+M)(c_0 A^2 + c_1 B^2) - M(A+B)^2 \leq \Delta^2 \\ A, B \geq 0 \ , \ \frac{1}{N} \leq c_0 \leq 1 \ , \ 0 < c_1 \leq 1 \end{array} \right\}. \end{aligned} \tag{3.15}$$

We further define $B = \rho A$, where $0 \leq \rho < 1$ and rewrite (3.15) as

$$\begin{aligned} &\max \quad (c_0 + \rho^2 c_1) A^2 \\ &\text{subject to} \quad \left\{ \begin{array}{c} (1+M)(c_0 + \rho^2 c_1) A^2 - M(1+\rho)^2 A^2 \leq \Delta^2 \\ A \geq 0 \ , \ \frac{1}{N} \leq c_0 \leq 1 \ , \ 0 < c_1 \leq 1 \ , \ 0 \leq \rho < 1 \end{array} \right\}. \end{aligned} \tag{3.16}$$

Define $\mu = (1+\rho)^2/(c_0 + \rho^2 c_1)$. Then $1 \leq \mu \leq 4N$ over the region (3.16). Setting $V = A^2(c_0 + \rho^2 c_1)$, the first constraint defining that region takes the form

$$(1+M)V - M\mu V \leq \Delta^2. \tag{3.17}$$

We now focus attention on the subset of the feasible region for (3.16) where

$$(1+M) - M\mu > 0. \tag{3.18}$$

Hence,

$$V \leq \frac{\Delta^2}{1 - M(\mu - 1)} \leq \frac{\Delta^2}{1 - M(4N-1)}, \tag{3.19}$$

as stated by (3.2) with the choice $\mu = 4N$.

The requirement (3.18) puts a restriction on N and M, being free parameters of the problem. Using $\mu = 4N$ leads to the sparsity requirement in (3.1), since $(1+M) - 4NM > 0$. □

3.2 Interpretation of the Stability Result

Theorem 3.1 prompts several remarks.

- Setting $\epsilon = \delta = 0$ puts us in the noiseless case $(P_{1,0})$. In that setting, Theorem 3.1 tells us that if $N < (1 + M^{-1})/4$, there will be zero error in finding the unique sparsest representation – i.e. solving the l^1 optimization problem $(P_{1,0})$ solves the l^0 problem $(P_{0,0})$. As the l^1 problem is convex and the l^0 problem combinatorial in general, this is by itself significant. The same general phenomenon described has been observed before in [13, 14, 12, 17]. The sharpest results, in [12, 17], established that this phenomenon occurs for any sparsity N smaller than $(1+M^{-1})/2$, which means that the new result is slack by a factor of 2 in the $\epsilon = 0$ case. Perhaps a tighter inequality could be achieved with more care.

- If the signal is not noisy (i.e. $\epsilon = 0$) but nevertheless $(P_{1,\delta})$ is employed with $\delta > 0$, an approximate solution is assured, with a bound on the deviation of the approximate representation from the ideal noiseless representation. So in 'needlessly' going from $(P_{1,0})$ to $(P_{1,\delta})$ we tolerate errors in the decomposition, but the errors are controlled.

- Perhaps more surprisingly, if the signal is noisy – $\epsilon > 0$ – and we set $\delta = 0$ as if there were no noise, a stability result is still obtained! This is initially confusing as it questions the need for adapting to noise by introducing the whole concept of $(P_{1,\delta})$ in the first place. However, note that our bound on the representation error is based on worst-case analysis, and will typically provide a generous overestimate on the actual error, as we will see in Section 5 below.

4 Local Stability of the Greedy Algorithm

Observe that both $(P_{0,\epsilon})$ and $(P_{1,\epsilon})$ refer to global optimization problems, while the orthogonal greedy algorithm (OGA) described in the introduction is based on greedy stagewise approximation. Paralleling this distinction, the stability result we now develop for OGA is a local one, valid only for sufficiently small $\epsilon < \epsilon^*(\alpha_0)$.

For ease of exposition we shall hereafter assume that the order of the columns ϕ_1, $\phi_2, \ldots$ in the overcomplete system matrix $\mathbf{\Phi}$ has been chosen so that in the ideal noiseless signal $y = x_0 + z = \mathbf{\Phi}\alpha_0 + z$, the first N entries in α_0 are the non-zero entries, and that these are ordered:

$$|\alpha_0(1)| \geq |\alpha_0(2)| \geq \quad \cdots \quad \geq |\alpha_0(N)|. \tag{4.1}$$

Theorem 4.1 *Suppose the ideal noiseless signal x_0 has a representation $x_0 = \mathbf{\Phi}\alpha_0$ satisfying*

$$N = \|\alpha_0\|_0 \leq \frac{1+M}{2M} - \frac{1}{M} \cdot \frac{\epsilon}{|\alpha_0(N)|}. \tag{4.2}$$

Then α_0 is the unique sparsest representation of x_0. Denote by $\hat{\alpha}_{OGA,\epsilon}$ the result of greedy stepwise least-squares fitting which stops as soon as the representation error $\leq \epsilon$. Then

(a) *$\hat{\alpha}_{OGA,\epsilon}$ has the correct sparsity pattern:*

$$supp(\hat{\alpha}_{OGA,\epsilon}) = supp(\alpha_0); \tag{4.3}$$

(b) *$\hat{\alpha}_{OGA,\epsilon}$ approximates the ideal noiseless representation:*

$$\|\hat{\alpha}_{g,\epsilon} - \alpha_0\|_2^2 \leq \frac{\epsilon^2}{1 - M(N-1)} \ . \tag{4.4}$$

We break the analysis in two stages, considering claims (4.3) and (4.4) in turn.

4.1 Getting The 'Correct' Sparsity Pattern

Lemma 4.2 *Suppose that we have a signal y satisfying $y = x_0 + z$ where x_0 admits sparse synthesis $x_0 = \mathbf{\Phi}\alpha_0$ using at most N atoms, where*

$$N < \frac{1+M}{2M} - \frac{1}{M} \cdot \frac{\epsilon}{\|\alpha\|_\infty}, \tag{4.5}$$

and where $\|z\| < \epsilon$. Then the first step of the greedy algorithm selects an atom from among the $\leq N$ nonzeros in α_0.

Proof. The greedy algorithm operates by projecting y onto each atom ϕ_k in turn, selecting an atom where the projection magnitude is largest. The Lemma will follow from

$$\max_{1 \leq k \leq N} |\langle y, \phi_k \rangle| > \max_{k > N} |\langle y, \phi_k \rangle|. \tag{4.6}$$

We now develop a lower bound on the left side and an upper bound on the right side which guarantees this. Assuming that the largest amplitude entry in α_0 occurs in slot 1, the left-hand side of (4.6) is lower-bounded by

$$\begin{aligned} |\langle y, \phi_1 \rangle| &= |\langle x + z, \phi_1 \rangle| \\ &\geq |\langle \sum_1^N \alpha(j)\phi_j, \phi_1 \rangle| - |\langle z, \phi_1 \rangle| \\ &\geq |\alpha_0(1)| - \sum_{j=2}^N |\alpha_0(j)| \cdot |\langle \phi_j, \phi_1 \rangle| - \epsilon \\ &\geq |\alpha_0(1)| - |\alpha_0(1)| \cdot (N-1)M - \epsilon. \end{aligned} \tag{4.7}$$

We used: $\|\phi_j\|_2^2 = 1$ for all j; $|\langle \phi_j, \phi_1 \rangle| \leq M$ for $j \neq 1$; $\|z\|_2 \leq \epsilon$; and the ordering of the $|\alpha(k)|$. The right-hand side of (4.6) can be upper-bounded by the same approach, leading to, for $k > N$

$$\begin{aligned} |\langle y, \phi_k \rangle| &= |\langle x + z, \phi_k \rangle| \\ &\leq \sum_{j=1}^N |\alpha_0(j)| \cdot |\langle \phi_j, \phi_k \rangle| + |\langle z, \phi_k \rangle| \\ &\leq |\alpha_0(1)| \cdot NM + \epsilon. \end{aligned} \tag{4.8}$$

Imposing (4.5) and using the two bounds (4.7)-(4.8), we see that

$$|\alpha_0(1)| - |\alpha_0(1)| \cdot (N-1)M - \epsilon > |\alpha_0(1)| \cdot NM + \epsilon. \tag{4.9}$$

Relation (4.6) follows; the greedy algorithm therefore chooses at Stage 1 one of the nonzeros in the ideal representation of the noiseless signal. □

To continue to later stages, we need

Lemma 4.3 *Let $x_0 = \sum_{i=1}^{N} \alpha_0(i)\phi_i$, and $y = x_0 + z$ with $\|z\| \le \epsilon$. Let $\mathcal{S}_k$ be a set of k indices in $\{1, \ldots, m\}$. Let a_k be a vector of m coefficients with k nonzeros located at the indices in $\mathcal{S}_k$. Define a new signal y_k by subtracting k atoms with nonzero coefficients in a_k:*

$$y_k = y - \sum_{i \in \mathcal{S}_k} a_k(i)\phi_i.$$

Similarly, define

$$x_k = x_0 - \sum_{i \in \mathcal{S}_k} a_k(i)\phi_i.$$

Then

- *If $\mathcal{S}_k \subset \{1, \ldots, N\}$ where $N < (1 + M^{-1})/2$, x_k has a unique sparsest representation $x_k = \mathbf{\Phi}\alpha_k$ made of at most N atoms; these are all atoms originally appearing in the representation of x_0.*
- *The new signal y_k can be viewed as a superposition of x_k and noise z_k, with noise level $\epsilon_k = \|z_k\| \le \epsilon$.*

Proof. Define the vector

$$\alpha_k(i) = \begin{cases} \alpha_0(i) - a_k(i) & i \in \mathcal{S}_k \\ \alpha_0(i) & i \notin \mathcal{S}_k \end{cases}$$

Then clearly $x_k = \mathbf{\Phi}\alpha_k$. Also, $supp(\alpha_k) \subseteq supp(\alpha)$. Since then

$$\|\alpha_k\|_0 \le \|\alpha_0\|_0 \le N < (1 + M^{-1})/2,$$

we conclude that α_k is the unique sparsest representation of x_k. Moreover,

$$\begin{aligned} \epsilon_k &= \|y_k - x_k\| \\ &= \| \left(y - \sum a_k(i)\phi_i\right) - \left(x_0 - \sum a_k(i)\phi_i\right) \|_2 \\ &= \|y - x_0\|_2 = \|z\|_2 \le \epsilon. \end{aligned}$$

Hence we have established the two claims. □

The impact of the preceding two Lemmas is that selection of a term, followed by the formation of the residual signal, leads us to a situation like before, where the ideal noiseless signal has no more atoms than before, and the noise level is the same.

We wish to repeatedly apply these Lemmas. Starting with $\alpha = \alpha_0$, we will get an α_1 and an i_1; we then hope to apply the observations again, getting α_2 and i_2, etc. If we are allowed to continue invoking these Lemmas for N steps, we produce in this way series $\alpha_0, \ldots, \alpha_N$, and i_1,...,i_N. Naturally, the sets $\mathcal{S}_k = \{i_1, ..., i_k\}$ are nested.

Note, however, that a series of conditions must be satisfied for the repeated use of the first Lemma. At the k-th iteration we need the following analog of (4.5):

$$N \le \frac{1 + M}{2M} - \frac{1}{M} \cdot \frac{\epsilon}{\|\alpha_{k-1}\|_\infty}. \tag{4.10}$$

This will follow from our original assumption (4.2) and

Lemma 4.4 *Let α_k differ from α_0 in at most k places, and let α_0 be ordered as in (4.1). Then*

$$\|\alpha_k\|_\infty \geq |\alpha_0(k+1)|.$$

Proof. Indeed, the largest element in α_k is at least as big as the largest untouched element in α_0, which is at least as big as the k-th from largest element in α_0. □

Lemma 4.4 and the ordering assumption on the coefficients in α_0 show that $\|\alpha_k\|_\infty \geq |\alpha_0(N)|$ for $1 \leq k \leq N$, and so the sequence of conditions (4.10) is implied by the final one at $i = N$, which is equivalent to (4.2). Hence, assumption (4.2) allows us to repeatedly apply Lemmas 4.2-4.3, and conclude that atoms selected at stages $1 \leq k \leq N$ obey $1 \leq i_k \leq N$: only correct atoms are selected.

In fact we can say much more. The coefficient sequence a_k generated at stage k solves the least-squares problem

$$\min_a \|y - \sum_{i \in \mathcal{S}_k} a(i)\phi_i\|_2. \tag{4.11}$$

This ensures that the signal y_k is actually orthogonal to each atom selected at stages 1,...k. Hence, OGA is forced to select from among the atoms in $\{1, ..., N\}$ always one of the previously unselected ones. It therefore by stage N selects all N atoms in $\{1, ..., N\}$. Now by assumption, the residual at that stage has ℓ^2 norm $\leq \epsilon$. Hence the stopping criterion must be reached by stage N. At the same time, by inspecting (4.9) we see that at stages $k < N$, each selected term $|\alpha_0(i_k)| > \epsilon$. This implies that the stopping criterion cannot be met before stage N, as $\|y_k\|_2^2 - \|y_{k-1}\|_2^2 \geq |\alpha_0(i_k)|^2$. Thus we have proved:

Lemma 4.5 *OGA stops after precisely N steps.*

This, in turn, proves Claim (a) of the Theorem, (4.3) □

Note that the argument assumes the noise level ϵ is known, to enable the stopping rule in the algorithm. This parallels the assumption $\delta = \epsilon$ we made in Theorems 2.1 and 3.1.

The general idea – that the support properties of α_0 and $\hat{\alpha}_{OGA,\epsilon}$ are the same, seems worthy of its own study. In the appendix below, we call this the *Trapping* property, and develop it further.

4.2 Stability Result

Now we turn to claim (b) of Theorem 4.1. We may partition $\mathbf{\Phi} = [\mathbf{\Psi}\, \mathbf{\Omega}]$, where $\mathbf{\Psi}$ denotes the first N columns in $\mathbf{\Phi}$ and $\mathbf{\Omega}$ the remainder. The OGA solves (4.11) with $\mathcal{S}_N = \{1, \ldots, N\}$, or

$$\hat{\alpha}_{OGA,\epsilon} = \arg\min_\alpha \ \|\mathbf{\Psi}\alpha - y\|_2^2 = \mathbf{\Psi}^+ y. \tag{4.12}$$

Recall that the signal has a representation $y = \mathbf{\Phi}\alpha_0 + z$ with $\|z\|_2 \leq \epsilon$. Thus, using the LS formula above we have

$$\hat{\alpha}_{OGA,\epsilon} = \mathbf{\Psi}^+ y = \mathbf{\Psi}^+ \left([\mathbf{\Psi},\ \mathbf{\Omega}]\alpha_0 + z\right), \tag{4.13}$$

where $\mathbf{\Psi}^+$ denotes the (spectral) generalized inverse of $\mathbf{\Psi}$. We may partition $\alpha_0 = [\beta\, 0]$ where β contains the first N entries of α, and similarly $\hat{\alpha}_{OGA,\epsilon} = [\hat{\beta}\, 0]$. We obtain

$$\hat{\beta} = \beta + \mathbf{\Psi}^+ z. \tag{4.14}$$

The vector $\mathbf{\Psi}^+ z$ represents reconstruction error, and we have the error bound

$$\|\hat{\alpha}_{OGA,\epsilon} - \alpha_0\|_2^2 \leq \|\mathbf{\Psi}^+ z\|_2^2 \leq \|\mathbf{\Psi}^+\|_2^2 \cdot \|z\|_2^2 \leq \frac{\epsilon^2}{\sigma_{min}^2\{\mathbf{\Psi}\}}. \tag{4.15}$$

We can lower-bound the minimal singular value of $\mathbf{\Psi}$ using Lemma 2.8.

$$\sigma_{min}^2\{\mathbf{\Psi}\} \geq \eta^2 \quad \text{such that } \text{Spark}_\eta(\mathbf{\Phi}) \geq \#\text{Columns}(\mathbf{\Psi}). \tag{4.16}$$

This gives $\sigma_{min}^2\{\mathbf{\Psi}\} \geq 1 - M(N-1)$. Hence

$$\|\hat{\alpha}_{OGA,\epsilon} - \alpha_0\|_2^2 \leq \frac{\epsilon^2}{\sigma_{min}^2\{\mathbf{\Psi}\}} \leq \frac{\epsilon^2}{1 - M(N-1)}, \tag{4.17}$$

as claimed. □

4.3 Extensions

Our analysis of the support properties of OGA used rather general properties of the greedy algorithm which are satisfied by other algorithms as well. We briefly consider two variants on OGA and stability results for those variants.

4.3.1 Pure Greedy Algorithm

The Pure Greedy algorithm (PGA) is similar to OGA, only without successive orthogonalization. Like the OGA, it builds up a k-element approximate representation a step at a time, adding to an existing $k-1$-element model a new term chosen in a greedy fashion to minimize the ℓ^2 error in the resulting approximation at that stage. Paralleling the OGA, it defines an initial residual $r^{(0)} = y$ and a current model $\hat{y}^0 = 0$; then for $k = 1, \ldots,$ it augments the model $\hat{y}^{(k-1)} \to \hat{y}^{(k)}$ and updates the residual $\hat{r}^{(k-1)} \to \hat{r}^{(k)}$ stepwise, always maintaining $y = \hat{y}^{(k)} + \hat{r}^{(k)}$. However, the update rule is different than the one for the OGA discussed so far: it is even greedier.

At the k-th stage, PGA selects an atom to be added to the model which offers the highest correlation with the current residual

$$i_k = \text{argmax}_{1 \leq i \leq m} |\langle r^{(k-1)}, \phi_i \rangle|; \tag{4.18}$$

it builds a model consisting of the previous model together with the newly selected atom:

$$\hat{y}^{(k)} = \hat{y}^{(k-1)} + a_{i_k} \phi_{i_k}, \tag{4.19}$$

where the coefficient $a_{i_k} = \langle r^{(k-1)}, \phi_{i_k} \rangle$ is fitted by least squares; and it updates the residual correspondingly:

$$r^{(k)} = r^{(k-1)} - a_{i_k} \phi_{i_k},$$

which can be input to the next stage of the algorithm.

This differs from the orthogonal greedy algorithm OGA discussed so far in that, in the PGA, the model at stage k is obtained without any re-adjustment of the coefficients of terms obtained at earlier stages, while OGA readjusts all coefficients of terms associated with indices $\{i_1, \ldots, i_{k-1}\}$; compare the discussion of (4.19) with that for (1.2). This variant of the greedy algorithm is the same as the original proposal for Matching Pursuit (MP) in [23]; it has also been applied in statistics under the name

stagewise regression [18], and has been used in Projection Pursuit Regression [18], and in approximation theory [32, 33], where the name Pure Greedy Algorithm was used. Further discussion of greedy-type algorithms is available in the survey [35].

We now point out that with the right stopping rule, there is a local stability result for PGA. We prove the following in the Appendix.

Theorem 4.6 *Suppose the ideal noiseless signal x_0 has a representation $x_0 = \mathbf{\Phi}\alpha_0$ satisfying (4.2). Let $M = M(\mathbf{\Phi})$, and set*

$$\eta \equiv \epsilon \cdot \frac{3 - M}{1 - M(2N - 1)}.$$

Denote by $\hat{\alpha}_{PGA,\eta}$ the result of pure greedy fitting which stops immediately when the prospective next term has coefficient not exceeding η. That is, it stops at the smallest $m = m(\eta, y)$ such that $|a_{i_{m+1}}| \le \eta$. Let $\delta = \|r^m\|$.

(a) *$\hat{\alpha}_{PGA,\eta}$ has at least part of the correct sparsity pattern:*

$$supp(\hat{\alpha}_{PGA,\eta}) \subseteq supp(\alpha_0); \tag{4.20}$$

(b) *$\hat{\alpha}_{PGA,\eta}$ approximates the ideal noiseless representation:*

$$\|\hat{\alpha}_{PGA,\eta} - \alpha_0\|_2^2 \le \frac{(\epsilon + \delta)^2}{1 - M(N - 1)} \ . \tag{4.21}$$

4.3.2 Weak Greedy Optimization

The Weak Orthogonal Greedy Algorithm (WOGA) [33] does not presume to be able to solve the maximization problem (4.18) underlying both OGA and PGA, but instead merely delivers an index i_k giving a near-maximum. For a constant $\rho \in (0, 1)$, it guarantees that each i_k obeys merely:

$$|\langle r^{(k-1)}, \phi_{i_k} \rangle| \ge \rho \cdot \max_{1 \le i \le m} |\langle r^{(k-1)}, \phi_i \rangle|. \tag{4.22}$$

Otherwise, the algorithm is the same as OGA. The point is that for certain multiscale dictionaries [23], it can be very much faster to locate an approximate maximizer than a genuine maximizer.

Theorem 4.7 *Let $K = 1 + \rho^{-1}$. Suppose the ideal noiseless signal x_0 has a representation $x_0 = \mathbf{\Phi}\alpha_0$ satisfying*

$$N = \|\alpha_0\|_0 < (1 + M^{-1})/K \tag{4.23}$$

and that the noise level obeys

$$\Lambda_\rho \epsilon \le |\alpha_0(N)|, \tag{4.24}$$

where

$$\Lambda_\rho \equiv \frac{K/\rho}{1 - M(KN - 1)}$$

Then (as usual) α_0 is the unique sparsest representation of x_0. Denote by $\hat{\alpha}_{WOGA,\epsilon}$ the result of weak orthogonal greedy fitting which stops as soon as the representation error $\le \epsilon$. Then

(a) *$\hat{\alpha}_{WOGA,\epsilon}$ has the correct sparsity pattern:*

$$supp(\hat{\alpha}_{WOGA,\epsilon}) \equiv supp(\alpha_0); \tag{4.25}$$

(b) *$\hat{\alpha}_{WOGA,\epsilon}$ approximates the ideal noiseless representation:*

$$\|\hat{\alpha}_{WOGA,\epsilon} - \alpha_0\|_2^2 \leq \frac{\epsilon^2}{1 - M(N-1)} \ . \tag{4.26}$$

This result imposes a requirement for sparsity which is more stringent than before, according as ρ is smaller than 1. The proof is in the Appendix.

5 Excessive Pessimism?

The conditions for stability developed here are unduly restrictive. We used worst-case reasoning exclusively, deriving conditions which must apply to *every* dictionary, *every* sparse representation and *every* bounded noise vector. The bounds we have proven are consequently very loose and do not describe typical behavior; the sparsity conditions we have posed are much too strict. To illustrate this, we conducted numerous experiments to study the stability of various algorithms in concrete cases.

5.1 Experiments with ℓ^1 penalization

We first consider the phenomenon described in Theorem 3.1. We work with a dictionary $\mathbf{\Phi} = [\mathbf{I},\ \mathbf{H}]$ obtained by concatenating two orthonormal bases – the standard and Hadamard bases for signals of length $n = 128$ each, yielding $M = 1/\sqrt{128}$. We used randomly-generated ideal representations α_0 satisfying the conditions of Theorem 3.1; since $(1 + M^{-1})/4 < 3.07$, we use $\|\alpha_0\|_0 = N = 1, 2, 3$. The non-zero entries of α_0 were located in uniform random positions, and the values of those entries were drawn from a normal distribution with zero mean and unit variance. The ideal noiseless signal $x_0 = \mathbf{\Phi}\alpha_0$ was contaminated with zero-mean white Gaussian noise z rescaled to enforce a specified noise level $\epsilon = \|z\|$, obtaining $y = \mathbf{\Phi}\alpha_0 + z$. We numerically solved $(P_{1,\delta})$ for either $\delta = 0$ or $\delta = \epsilon$ and calculated the error measure $\|\alpha_0 - \hat{\alpha}_{1,\delta}\|_2^2$.

Our first experiment considered $\delta = 0$, which leads to $(P_{1,0})$, and which we solved by linear programming. Figure 1 compares the upper bound from Theorem 3.1 with empirical results. The noise level was chosen at several levels $\epsilon \in [0, 1]$, and the number N of nonzeros ranged through $\{1, 2, 3\}$. The three slanted lines in the figure display our theoretical stability bounds, which increase as N grows. The experiments (20 runs per each N, ϵ combination) show much smaller errors than allowed by the bounds. The empirical representation error *does* grow linearly with the noise level. However, the empirical errors show no apparent dependency on N, with all the simulation results exhibiting the same slope, which suggests that our upper bound's dependency on N is a severe overstatement of generic behavior.

Figure 2 refers again to the case $\delta = 0$. With a fixed noise level $\epsilon^2 = 0.5$, we considered a wide range of choices of N. This figure displays the resulting representation error, averaged across 20 experiments at each level of N. As can be seen, stability is obtained far beyond the limit $(1 + M^{-1})/4$ needed for applicability of Theorem 3.1, and again, the error does not seem to depend markedly on the sparsity level in the direct way seen in our upper bound. "Explosion" of the representation error eventually appears as sparsity abates, but only gradually and only well beyond $N = 30$, while our Theorem seems to suggest that we ought to be concerned starting as soon as $N > 3$.

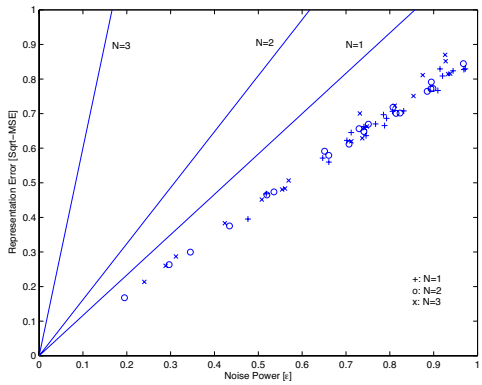

Figure 1: ℓ^1 method, $\delta = 0$: Representation error $\|\alpha_0 - \hat{\alpha}_{1,0}\|_2$ as a function of noise level ϵ.

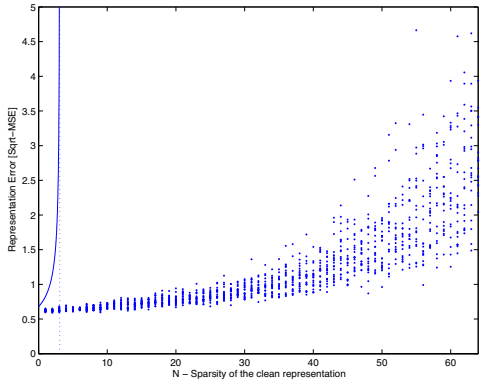

Figure 2: ℓ^1 method, $\delta = 0$: Representation error $\|\alpha_0 - \hat{\alpha}_{1,0}\|_2$ as a function of $N = \|\alpha_0\|_0$, assuming $\epsilon^2 = 0.5$.

We also considered $\delta = \epsilon$ (i.e., the noise power is known and imposed in the recovery process). We solved $(P_{1,\epsilon})$ by Iteratively Reweighted Least-Squares (IRLS). We repeated the experimental setup used for Figure 1; the results appear in Figure 3. Our bounds in this case permit twice the previous error, as now $\delta = \epsilon$. However, as can be seen, the empirical results suggest the contrary – much smaller errors are seen in the recovered representation than were seen before. This is intuitively appealing - it seems clear that $(P_{1,\epsilon})$ should be more noise-cognizant than $(P_{1,0})$.

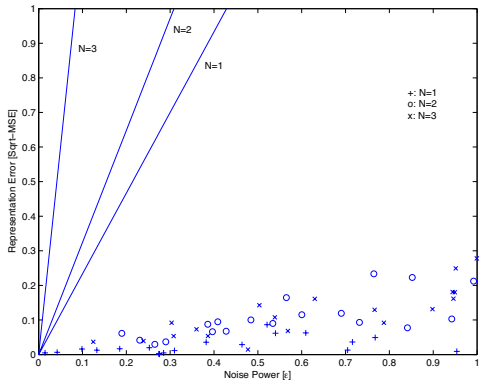

Figure 3: ℓ^1 method, $\delta = \epsilon$: Representation error $\|\alpha_0 - \hat{\alpha}_{1,\epsilon}\|_2$ as a function of the noise level ϵ. Note the substantially smaller error level than in the $\delta = 0$ case.

5.2 Experiments with Greedy Optimization

We now compare the reconstruction errors for OGA versus the bounds in Theorem 4.1. Paralleling the experiment for ℓ^1, we display theoretical bounds and empirical errors in Figure 4. Evidently, OGA behaves stably with results comparable to those obtained by the ℓ^1 penalization. Moreover, the upper bounds provided in Theorem 4.1 are seen to be exactly that – dramatic overestimates of the reconstruction error. Note that in this experiment we have used $N = 1, 2, 3$, totally disregarding the condition as posed in (4.2), tying the allowed sparsity to the coefficients' amplitude. Yet, the results show stability and controlled amount of error.

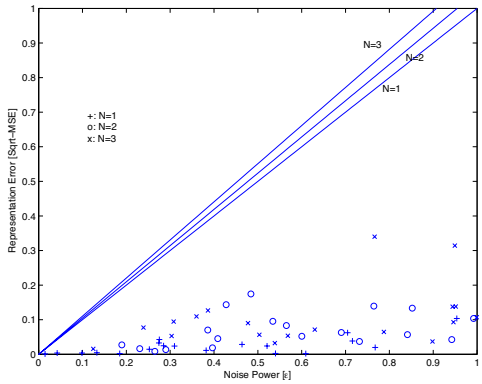

Figure 4: OGA representation error $\|\alpha_0 - \hat{\alpha}_{OGA}\|_2$ as a function of the noise level ϵ.

5.3 Geometric Heuristics

As indicated above, our very general reasoning is to blame for the looseness of our theoretical bounds; by developing bounds valid for a wide range of dictionaries and a wide range of sparsely-represented signals, we are forced to consider the behavior fat the worst-possible combination of dictionary, signal, and noise.

We might get tighter results, by developing tools adapted to each specific $(\mathbf{\Phi}, \alpha_0)$ combination. Unfortunately, the closer we get to case-by-case analysis, the more difficult it becomes to get an intellectually-digestible overview of the situation. At least for ℓ^1 minimization, it seems clear to the authors that, even at values N far greater than those covered in Theorem 3.1 it will generally be true that:

- A sparse vector α_0 generating $x_0 = \mathbf{\Phi}\alpha_0$ will be the unique solution of $(P_{1,0})$; and
- The solution of $(P_{1,\epsilon})$ based on noisy data $y = x_0 + z$ with noise level ϵ will stably recover α_0.

It is less clear to us that we can expect the solution to the ℓ^1 problem to agree with the solution to the ℓ^0 problem with the same degree of generality.

Some insight may be gleaned by considering the geometry of minimal ℓ^1 decomposition; see Figure 5 below. The minimal ℓ^1 decomposition in an overcomplete system is the point in the subspace $A_{x_0} = \{\alpha : x_0 = \mathbf{\Phi}\alpha\}$ having the smallest ℓ^1 norm. Denote this norm by $R_0 = val(P_{1,0})$. Alternatively, if we consider the collection of balls $B_1(R) = \{\alpha : \|\alpha\|_1 \leq R\}$ in R^m, it is the 'first point in A_{x_0}' to 'meet' the family of balls as R grows from 0 to R_0. When this meeting occurs, if it is in a unique point, then the ℓ^1 decomposition is unique. Now note that if α_0 has few nonzeros, then it sits in a low-dimensional face of $B_1(R_0)$. Denote by F_{α_0} the smallest-dimensional face of $B_1(R_0)$ containing α_0 in its interior.

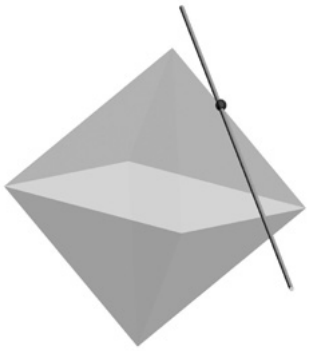

Figure 5: Geometry favorable to unique ℓ^1 decomposition. Intersection of A_{x_0} with $B_1(R_0)$ in a unique point. This point is the unique solution $\hat{\alpha}_{1,0}$

Figure 5 shows clearly a situation where F_{α_0} is *transversal* to A_{x_0} – the two subspaces meet nicely in a single point. More than this: all the faces of $B_1(R_0)$ touching F_{α_0} intersect A_{x_0} tranversally. Now the cleanness of these intersections imply that α_0 is the unique ℓ^1 minimizer in A_{x_0}.

A key observation is that the faces of the ball $B_1(R_0)$ run through a finite list of specific orientations. If we take a *generic* $\mathbf{\Phi}$, there would never be a fortuitous alignment of any subspace A_{x_0} with any of the low-dimensional faces of $B_1(R_0)$; hence transversal intersections should be generic, and we can expect to have unique ℓ^1 minimizers *except* when $\|\alpha_0\|_0$ and $\dim A_{x_0}$ demand non-uniqueness.

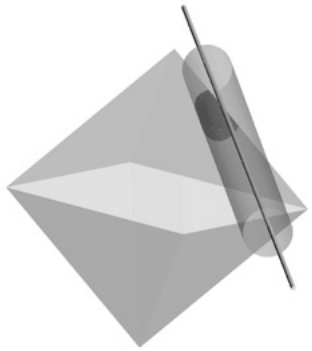

Figure 6: Geometry favorable to stable ℓ^1 decomposition. Intersection of $A_{x_0,2\epsilon}$ with $B_1(R_0)$ in a tubular wedge. $\hat{\alpha}_{1,\epsilon}$ must lie in wedge. Small size of the wedge indicates stability.

What about stability? A geometric explanation of stability for $\hat{\alpha}_{1,\epsilon}$ is illustrated in Figure 6. Because of

$$\|\hat{\alpha}_{1,\epsilon}\|_1 \leq \|\alpha_0\|_1$$

$\hat{\alpha}_{1,\epsilon}$ must belong to the cone C_{1,α_0} with vertex at α_0 consisting of all points α such that for some $t > 0$

$$\|(1-t)\alpha + t\alpha_0\|_1 \leq \|\alpha_0\|_1.$$

On the other hand, because of

$$\|x_0 - \mathbf{\Phi}\hat{\alpha}_{1,\epsilon}\|_2 \leq 2\epsilon,$$

$\hat{\alpha}_{1,\epsilon}$ must belong to the cylinder $A_{x_0,2\epsilon}$ consisting of all vectors α obeying

$$\|x_0 - \mathbf{\Phi}\alpha\|_2 \leq 2\epsilon.$$

In short, for a Cone C_{1,α_0} and a cylinder $A_{x_0,2\epsilon}$, we have

$$\hat{\alpha}_{1,\epsilon} \in C_{1,\alpha_0} \cap A_{x_0,2\epsilon}.$$

Roughly speaking, the size of this intersection is controlled by the angle between A_{x_0} and C_{1,α_0}. That this angle can be positive we know already; because that is the content of the transversality we have already discussed.

There is an analytical framework to quantify the above heuristic notions. There is a stability estimate adapted to a specific $(\mathbf{\Phi}, \alpha_0)$ pair:

$$\|\hat{\alpha}_{1,\epsilon} - \alpha_0\| \leq \Lambda_1(\mathbf{\Phi}, \alpha_0) \cdot 2\epsilon, \qquad \epsilon > 0;$$

where

$$\Lambda_1(\mathbf{\Phi}, \alpha_0) = \sup\{\frac{\|\gamma - \alpha_0\|}{\|\mathbf{\Phi}(\gamma - \alpha_0)\|} : \|\gamma\|_1 \leq \|\alpha_0\|_1\}.$$

Equivalently,

$$\Lambda_1(\mathbf{\Phi}, \alpha_0) = \sup\{\frac{\|v\|}{\|\mathbf{\Phi}v\|} : v \in \dot{C}\},$$

where $\dot{C}$ denotes the tangent cone to C_{1,α_0} at α_0, i.e. the collection of vectors v such that $\alpha_0 + tv \in C_{1,\alpha_0}$ for all sufficiently small $t > 0$.

This last display makes the point that we are trying to optimize a ratio of quadratic forms subject to membership in a cone. This makes us say that Λ_1 is akin to the secant of the angle between A_{x_0} and C_{1,α_0}. Unfortunately, to our knowledge, the problem of finding the angle between a cone and a subspace does not have a convenient computational solution. Hence, although the bound depends intimately on $\mathbf{\Phi}$ and α_0, we know of no way to easily compute this dependence at the moment.

6 Getting the 'Correct' Support with ℓ^1

Our results on local stability of the Greedy Algorithm all use the fact that, under appropriate conditions, the support of the approximate representation for noisy data is a subset of the support of the underlying ideal sparse representation. Now we develop a parallel result for $(P_{1,\delta})$ which is nonlocal, i.e. it it true not merely for small noise level ϵ.

In our formulation, we run $(P_{1,\delta})$ with a specially chosen $\delta \gg \epsilon$.

Theorem 6.1 *Suppose that $y = x_0 + z$ where $x_0 = \mathbf{\Phi}\alpha_0$, $\|\alpha_0\|_0 \le N$ and $\|z\| \le \epsilon$. Let $M = M(\mathbf{\Phi})$ and suppose $\beta \equiv MN < 1/2$. Set*

$$\gamma = \frac{\sqrt{1-\beta}}{1-2\beta}.$$

Solve $(P_{1,\delta})$ with exaggerated noise level $\delta = C \cdot \epsilon$, where $C = C(M,N) = \gamma\sqrt{N}$. Then supp$(\hat{\alpha}_{1,\delta}) \subset$ supp(α_0).

As an example, if $\beta = 1/4$, then $\gamma = \sqrt{3}$; so exaggerating the noise level by a factor $\sqrt{3N}$ leads to getting at least part of the support correctly. While this is rather a severe exaggeration of the noise level, it seems quite surprising that any result of this kind is possible. The $\sqrt{N}$ dependence is intrinsic to the problem. Even in the case where $\mathbf{\Phi}$ is orthogonal, so $M(\mathbf{\Phi}) = 0$, the requirement to have a result of this form is that $\delta = \sqrt{1+N} \cdot \epsilon$.

Proof. Let $\mathcal{S}$ be the support of the ideal noiseless representation α_0, and consider the support-constrained optimization problem $(P_{1,\delta,\mathcal{S}})$ where feasible vectors α must be supported in $\mathcal{S}$. Let α_1 be a solution of this problem. We claim that, in fact, α_1 is actually the solution of the *unconstrained* problem $(P_{1,\delta})$, i.e. $\alpha_1 = \hat{\alpha}_{1,\delta}$.

To do this, we consider perturbations u of α_1 i.e. representations of the form $\alpha_1 + t \cdot u$, for $t > 0$ small. We will show that a perturbation which does not increase the ℓ^1 objective, definitely violates the constraint. Formally

$$\|\alpha_1 + tu\|_1 \le \|\alpha_1\|_1 \quad \text{for small } t > 0, \tag{6.1}$$

implies

$$\|y - \mathbf{\Phi}(\alpha_1 + tu)\|_2 > \delta \quad \text{for small } t > 0. \tag{6.2}$$

By convexity, this local condition implies global optimality. To formalize the local optimality condition, we need two initial observations.

Lemma 6.2 *If $\#\mathcal{S} = N$ where $N < (M^{-1}+1)/2$ then the solution to $(P_{1,\delta,\mathcal{S}})$ is unique.*

Proof. We know that for an appropriate $\lambda > 0$ every constrained solution of $(P_{1,\delta,\mathcal{S}})$ is a minimizer of the form

$$Q_\lambda(\alpha) = \|y - \mathbf{\Phi}\alpha\|_2 + \lambda \sum_{j \in \mathcal{S}} |\alpha(j)|.$$

Consider Q_λ in the vicinity of α_1:

$$Q_\lambda(\alpha_1 + tu) = \|r\|^2 - 2t\langle r, \mathbf{\Phi} u\rangle + t^2\|\Phi u\|^2 + \lambda \sum_{j\in\mathcal{S}_0} |\alpha_1(j) + tu(j)| + t\cdot\lambda \sum_{j\in\mathcal{S}\setminus\mathcal{S}_0} |u(j)|$$

where $r \equiv y - \mathbf{\Phi}\alpha$ and $\mathcal{S}_0 = \mathrm{supp}(\alpha_1) \subset \mathcal{S}$. Then for small $t > 0$:

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) = -2t\langle r, \mathbf{\Phi} u\rangle + t^2\|\Phi u\|^2 + t\cdot\lambda \sum_{j\in\mathcal{S}_0} \sigma(j)u(j) + t\cdot\lambda \sum_{j\in\mathcal{S}\setminus\mathcal{S}_0} |u(j)|,$$

where $\sigma(j) = sign(\alpha_1(j))$ for $j \in \mathcal{S}_0$, and 0 otherwise. In order to be a minimizer of Q_λ, we must have that for all vectors u supported in $\mathcal{S}_0$,

$$0 = -2t\langle r, \mathbf{\Phi} u\rangle + t\cdot\lambda \sum_{j\in\mathcal{S}_0} \sigma(j)u(j),$$

which implies that

$$(\mathbf{\Phi}^T r) = (\lambda/2)\sigma(j), \quad j \in \mathcal{S}_0.$$

It follows that for small enough $t > 0$,

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) = t^2\|\Phi u\|^2 + t\cdot\lambda \sum_{j\in\mathcal{S}\setminus\mathcal{S}_0} |u(j)|.$$

We now invoke Lemma 6.3 below, with $\Omega = \mathcal{S}$ and $\beta = t^{-1}$, to conclude that for every nonzero u supported in $\mathcal{S}$,

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) > 0 \qquad \text{for small } t > 0.$$

It follows from convexity of Q_λ that α_1 is the unique global optimizer of Q_1. □

Lemma 6.3 *Let $\Omega \subset \{1, \ldots m\}$. Let $\mathcal{S}_0 \subset \Omega$. Let u be a nonzero vector supported in Ω, and $\beta > 0$. Then if $M = M(\mathbf{\Phi})$ and $N = \#\mathcal{S}_0 < (M^{-1} + 1)/2$,*

$$0 < \|\Phi u\|^2 + \beta \cdot \sum_{j\in\Omega\setminus\mathcal{S}_0} |u(j)|.$$

Proof. If the support of $u \in R^m$ intersects $\Omega\setminus\mathcal{S}_0$, then $0 < \beta\sum_{j\in\Omega\setminus\mathcal{S}_0} |u(j)|$. On the other hand, if $u \in R^m$ has support $\subset \mathcal{S}_0$ then $\|\Phi u\|_2^2 > 0$ whenever $u \neq 0$. Indeed, the N columns of Φ corresponding to nonzeros in u will be linearly independent. □

We now show that α_1 is the unique global optimum of $(P_{1,\delta})$ – i.e. the original problem, without the support constraint. We introduce the notation $[,]$ for inner product in R^m, $[,]_0$ for inner product restricted to coordinates in $\mathcal{S}_0$, and $[,]_1$ for inner product restricted to coordinates in $\mathcal{S}_0^c$. We introduce the notation $\|\|_{1,0}$ for the ℓ^1 norm restricted to $\mathcal{S}$, etc. We recall the notation r, σ and Q_λ introduced in the proof of Lemma 6.2. Our aim is now to show that for *every*nonzero u without regard to support,

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) > 0 \qquad \text{for small } t > 0. \tag{6.3}$$

This is an expansion of our claim from Lemma 6.2, which applied only to u supported in $\mathcal{S}$. Using arguments similar to those above, we have that for small $t > 0$,

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) = t^2\|\Phi u\|^2 + t \cdot 2\langle r, \mathbf{\Phi} u\rangle_1 + t \cdot \lambda\|u\|_{1,1}.$$

We wish to show that, for some $\eta > 0$,

$$2\langle r, \mathbf{\Phi} u\rangle_1 + \lambda\|u\|_{1,1} \geq \eta \cdot \|u\|_{1,1}. \tag{6.4}$$

It will then follow that

$$Q_\lambda(\alpha_1 + tu) - Q_\lambda(\alpha_1) \geq t^2\|\Phi u\|^2 + t \cdot \eta \cdot \|u\|_{1,1}.$$

applying Lemma 6.3 with $\Omega = \{1, \ldots, m\}$ and $\beta = \eta t^{-1}$. (6.3) follows.

To establish (6.4), we note that

$$\langle r, \mathbf{\Phi} u\rangle = [\mathbf{\Phi}^T r, u]_0 + [\mathbf{\Phi}^T r, u]_1$$

we write $r = r_0 + r_1$, where r_1 is the component of r not in the span of $(\phi_j : j \in \mathcal{S}_0)$, which has norm $\leq \epsilon$ while r_0 is the component of r in the span of the $(\phi_j : j \in \mathcal{S}_0)$,with norm $\leq \delta$. Hence

$$[\mathbf{\Phi}^T r, u]_1 = [\mathbf{\Phi}^T r_0, u]_1 + [\mathbf{\Phi}^T r_1, u]_1$$

$$[\mathbf{\Phi}^T r_1, u]_1 \leq \|u\|_{1,1} \left(\max_{j \in \mathcal{S}_0^c} |(\mathbf{\Phi}^T r_1)_j| \right) = \|u\|_{1,1}\epsilon$$

$$[\mathbf{\Phi}^T r_0, u]_1 = \langle \mathbf{\Phi}_0 v, \mathbf{\Phi}_1 u\rangle = \sum_{i \in \mathcal{S}_0, j \in \mathcal{S}_0^c} G_{ij} v(i) u(j) \leq M\sqrt{N}\|u\|_{1,1}\|v\|_2$$

where $\Phi_0 v = r_0$, and $\mathbf{\Phi}_0$ is the matrix with columns from $\mathcal{S}$ only. We calculate from Lemma A.1 in the Appendix that

$$\|v\|_2(1 - M(N-1))^{1/2} \leq \|r_0\|_2$$

and so

$$[\mathbf{\Phi}^T r_0, u]_1 \leq \delta \cdot |M\sqrt{N}|/(1 - M(N-1))^{1/2}\|u\|_{1,1}$$

and so

$$|[\mathbf{\Phi}^T r, u]_1| \leq \epsilon\|u\|_{1,1}(1 + C|M\sqrt{N}|/(1 - M(N-1))^{1/2}).$$

At the same time, we note that by the optimality of α_1 for $(P_{1,\delta,\mathcal{S}})$, there is a Lagrange multiplier $\lambda > 0$ with

$$[\mathbf{\Phi}^T r, u]_0 = -\lambda[\sigma, u]_0, \qquad \forall u \in R^m.$$

Now by inspection

$$\lambda = \|\mathbf{\Phi}^T r\|_{2,0}/\sqrt{N}.$$

Using $\|r\|_2 = \delta$, and norm equivalence (eg. as in Lemma A.1 or else (3.10)),

$$\|\mathbf{\Phi}^T r\|_{2,0} \geq \delta \cdot (1 - M(N-1))^{1/2}.$$

We conclude that (6.4) holds provided $A > B$, where $A = \delta(1 - M(N-1))^{1/2}/\sqrt{N}$ while $B = \epsilon(1 + C|M\sqrt{N}|/(1 - M(N-1))^{1/2})$. Dividing by ϵ and setting $MN = \beta$ and $C = \gamma\sqrt{N}$ we get $A > B$ when γ is defined as in the statement of the Theorem. (n.b. γ is defined precisely to make $A > B$). (6.4) follows, implying (6.3). □

7 Comparison to Other Work

We now sketch relationships between ideas presented here and other work. The main relationship is to the great deal of ongoing work where image and signal processing tasks are being carried out in overcomplete systems. Much of that work is motivated by the idea that there are sparse representations of the images/signals in the overcomplete representation. What we demonstrate here is that the sparsity heuristic actually can lead rigorously to methods which are stable against the presence of noise. In that sense, those heuristic methods now have an intellectual basis, or at least the beginnings of one. There are two other relations we would like to discuss.

7.1 N-Term Approximation

There has recently been a great deal of interest in the greedy algorithm as a method to generate near-best N-term approximations. Relevant literature includes [32, 33, 15, 36, 35].

The questions asked in this literature concern the quality of N-term approximations $x_N = \mathbf{\Phi}\alpha_N$, where $\|\alpha_N\|_0 = N$, to approximate a general object x. More specifically, let $x_{N,OGA}$ be the N-term approximation to a vector x by the orthogonal greedy algorithm run through N steps, and $x_{N,0}$ be the optimal N-term approximation, obtained by

$$x_{N,0} = \mathrm{argmin}\{\|x - \mathbf{\Phi}\alpha_N\|_2 : \|\alpha_N\|_0 \le N\}.$$

The central question is to compare the approximation errors of the two approaches:

$$\epsilon_{N,0} \equiv \|x - x_{N,0}\| \approx \epsilon_{N,OGA} \equiv \|x - x_{N,OGA}\|.$$

In this direction, the first result was provided by Gilbert et al. [15]

$$\epsilon_{N,OGA} \le 8\sqrt{N}\epsilon_{N,0}, \qquad N < 1/(8M).$$

This was then improved by Tropp to [36]

$$\epsilon_{N,OGA} \le (1 + 6\sqrt{N})^{1/2} \cdot \epsilon_{N,0}, \qquad N < 1/(8M). \tag{7.1}$$

These results show that, at least in its initial stages, the greedy algorithm performs quite well compared to the optimal algorithm.

The results of Section 4 have only indirect connection to this problem. Note that all our estimates in the rest of this paper concern errors such as $\|\alpha_0 - \hat{\alpha}_{OGA}\|$ on the *representation* scale, and the errors of interest in this subsection, such as $\|x_0 - x_{OGA}\|$, are measured on the *reconstruction* scale. Nevertheless, in an incoherent dictionary, the two are connected for small N. In the appendix we prove:

Lemma 7.1 *Let $MN < 1$ and let $\|\alpha_0\|_0, \|\alpha_1\|_0 \le N$. Then*

$$(1 - M(2N-1))^{1/2} \cdot \|\alpha_1 - \alpha_0\|_2 \le \|\mathbf{\Phi}(\alpha_0 - \alpha_1)\|_2 \le (1 + M(2N-1))^{1/2} \cdot \|\alpha_1 - \alpha_0\|_2.$$

In [38] a similar result was obtained for a variant of the OGA discussed here. If, say, $N \le 1/(4M)$, the coefficients in these inequalities are bracketed below by $1/\sqrt{2}$ and above by $\sqrt{3/2}$, respectively. Hence for such N there is a reasonably tight connection between the representation norm and the approximation norm. It follows that the ideas in this paper are very similar to ideas underlying the N-term approximation results cited above. Pursuing this thread leads to results like the following.

Theorem 7.2 *Let $\epsilon_{N,WOGA}$ denote the orthogonal greedy algorithm with weakly-optimal term selection, as discussed above, using parameter ρ. Set $K = 1 + \rho^{-1}$. Then*

$$\epsilon_{N,WOGA} \le \Lambda_\rho \cdot \sqrt{N} \cdot \epsilon_{N,0},$$

where now

$$\Lambda_\rho \equiv 1 + \frac{K \cdot (1 + M(N-1))^{1/2}}{1 - M(KN-1)}.$$

Comparing with Tropp's result (7.1) – in a case where both apply – say $N = 1/3M$ and $\rho = 1$, gives a coefficient $\Lambda_\rho \le 1 + \frac{2\cdot\sqrt{4/3}}{2/3} = 1 + 2/\sqrt{3}$, which is somewhat better than the coefficent $\sqrt{6}$ in (7.1); moreover the new result extends to weak term selection rules. See the Appendix.

Moreover, it is possible to extend the analysis to show that the order $N^{1/2}$ growth in the above analysis can be substantially reduced. In a triangular array setting, where we consider a sequence of n-dimensional approximation problems, suppose that $M = M_n \to 0$. We can exhibit a sequence of improved bounds, far outdistancing the $\sqrt{N}$ formulations of e.g. (7.1).

Theorem 7.3 *Consider a sequence of n-dimensional problems with dictionaries $\mathbf{\Phi}_n$ obeying $M = M(\mathbf{\Phi}_n) \to 0$. For each $\eta \in (0, 1/2)$, there exist constants $c_{i,\eta}$, $i = 1, \ldots, 3$ so that for $N \le (1/M)^{c_{1,\eta}}$,*

$$\epsilon_{c_{2,\eta}N,OGA} \le c_{3,\eta} \cdot N^\eta \cdot \epsilon_{N,0}.$$

See the Appendix.

7.2 Domain of Applicability

An *apparent* application of the results of this paper concerns the problem of resolving a spectrum at a resolution finer than the usual Rayleigh spacing. As a simple model we could consider the complex-valued dictionary with atoms

$$\phi_i(k) \propto \exp\{\sqrt{-1}\frac{2\pi}{\nu n} i(k-1)\}, \qquad k = 1, \ldots, n,\ i = 1, \ldots, \nu n.$$

Here ν, an integer > 1 is the *superresolution* factor, and the implicit constant of proportionality is chosen to enforce the normalization $\|\phi_i\|_2 = 1$. In this overcomplete system, the frequencies are spaced $\frac{2\pi}{\nu n}$ apart, which is ν times as closely as the usual spacing $\frac{2\pi}{n}$ of the Fourier frequencies, hence the term superresolution. If we simply chose $\nu = 1$, we'd have an orthogonal dictionary. If we choose $\nu = 2$, we get an overcomplete system with $m = 2n$. It would be very attractive to be able to solve this problem, getting finer frequency resolution out of a given signal length. However, Wohlberg [41] showed that in general this problem will lead to extreme ill-posedness, even under sparsity constraints.

We remark that, while superresolution is an attractive and important problem, this is *not* the setting we envisioned for applying our results. In the superresolving case, with $\nu = 2$, the dictionary has mutual coherence $M = M_n = \frac{1}{n \sin(\pi/n)}$, which tends to π^{-1} as n increases. This is quite large, and the dictionary is coherent rather than incoherent. It yields the sparsity threshold $(1 + M^{-1})/2 \approx 2.07$ which allows to disentangle at most two atoms, at any n!

The kind of situation we have in mind for applying our results is quite different; we are interested in cases where the mutual coherence is comparable, for large n, to some power $n^{-\beta}$, so that, at least for large n, there is the potential to disentangle fairly complex superpositions of many atoms. Previous

work has given several examples of this type of situation: random dictionaries [12], Grassmannian frames [30], and dictionaries for $3-D$ voxel data made of digital points, lines, and planes [12].

For those interested in supperresolution, we remark that, in our opinion, the analysis in [41] adopts a framework which is unduly pessimistic. The careful theoretical work on superresolution [11] explains that stable superresolution is possible using sparsity; however, the notion of sparsity needs to be *adapted* to the setting. Specifically, it becomes important to define sparsity in terms of 'number of nonzeros per Rayleigh interval' rather than simply 'number of nonzeros'. When this definitional convention is adopted, it is possible to prove that sufficient sparsity again enables superresolution, in accord with a considerable body of empirical work, also cited in [11].

Appendix: Proofs

The Greedy Algorithm

Our earlier description of the greedy algorithm is informal, so we revert here to standard notations for discussion of greedy algorithms in approximation theory [32, 33]. We let H be a Hilbert space containing our signals, and let now g,f, etc. denote elements of this space. In terms of earlier usage, $H = R^n$, $f = y$, etc., but the reasoning is more general.

Let $\mathbf{\Phi}$ be a dictionary for the Hilbert space H; this is a finite, countable or continuous collection of elements $\{\varphi\}$ normalized so that $\|\varphi\| = 1$ and having closed span H. We consider the maximal correlation between f and any member of the dictionary:

$$\chi(f) = \sup_{\phi \in \mathbf{\Phi}} |\langle f, \phi \rangle|,$$

and let $\varphi(f) \in \mathbf{\Phi}$ be an element from $\mathbf{\Phi}$ which maximizes he correlation $|\langle f, \varphi \rangle|$. For simplicity, suppose that a unique maximizer exists; if not, obvious modifications can be supplied. (See also the remark below) We start from the simplest notion of Greedy Algorithm. We define

$$G(f, \mathbf{\Phi}) \equiv \langle f, \varphi(f) \rangle \varphi(f)$$

and

$$R(f, \mathbf{\Phi}) \equiv f - G(f, \mathbf{\Phi}).$$

Pure Greedy Algorithm (PGA) *We define $R_0(f, \mathbf{\Phi}) \equiv f$ and $G_0(f, \mathbf{\Phi}) \equiv 0$. Then, for each $m \geq 1$, we inductively define*

$$G_m(f, \mathbf{\Phi}) \equiv G_{m-1}(f, \mathbf{\Phi}) + G(R_{m-1}(f, \mathbf{\Phi}), \mathbf{\Phi})$$
$$R_m(f, \mathbf{\Phi}) \equiv f - G_m(f, \mathbf{\Phi}) = R(R_{m-1}(f, \mathbf{\Phi}), \mathbf{\Phi}).$$

We next turn to the orthogonal greedy algorithm. If H_0 is a finite-dimensional subspace of H, we let P_{H_0} be the orthogonal projector from H onto H_0. That is, $P_{H_0}(f)$ is the best approximation to f from H_0.

Orthogonal Greedy Algorithm (OGA) *We define $R_0^o(f) \equiv R_0^o(f, \mathbf{\Phi}) \equiv f$ and $G_0^o(f) \equiv G_0^o(f, \mathbf{\Phi}) \equiv 0$. Then for each $k \geq 1$, we inductively define*

$$\begin{aligned} H_k \equiv \quad & H_k(f) \equiv span\{\varphi(R_0^o(f)), \ldots, \varphi(R_{k-1}^o(f))\} \\ G_k^o(f) \equiv \quad & G_k^o(f, \mathbf{\Phi}) \equiv P_{H_k}(f) \\ R_k^o(f) \equiv \quad & R_k^o(f, \mathbf{\Phi}) \equiv f - G_k^o(f). \end{aligned}$$

$$(A-1)$$

Following [33] we define the **Weak Orthogonal Greedy Algorithm (WOGA)**. *Let $\rho \in (0,1]$ be given, and set $f_0^{o,\rho} \equiv f$. Then for each $m \geq 1$ define:*

1. *$\varphi_m^{o,\rho} \in \mathbf{\Phi}$ is any element satisfying*

$$|\langle f_{m-1}^{o,\rho}, \varphi_m^{o,\rho} \rangle| \geq \rho \cdot \sup_{\phi \in \mathbf{\Phi}} |\langle f_{m-1}^{o,\rho}, \phi \rangle|; \tag{A-2}$$

2. $$G_m^{o,\rho}(f) \equiv G_m^{o,\rho}(f, \mathbf{\Phi}) \equiv P_{H_m^\rho}(f), \quad \textit{where} \quad H_m^\rho \equiv \textit{span}(\varphi_1^{o,\rho}, \ldots, \varphi_m^{o,\rho});$$

3. $$f_m^{o,\rho} \equiv f - G_m^{o,\rho}(f, \mathbf{\Phi}).$$

In a finite-dimensional space H the WOGA terminates after $m = \dim H$ steps and $G_m^{o,\rho}(f) = f$.

Remark: All variants of the GA fail to specify a unique sequence of approximants, since ties are possible in the definition of $\varphi(f)$. However, this nonuniqueness can be easily eliminated in the case of a countable dictionary $\mathbf{\Phi}$. Then we can give the elements of the dictionary and ordering $\mathbf{\Phi} = \{\phi_i\}_{i=1}^\infty$, and modify the definition of $\varphi(f)$ to specify the element ϕ_i appearing earliest in the ordering and satisfying

$$|\langle f, \phi_i \rangle| = \chi(f).$$

Comparable modifications of (A-2) can be imposed as well.

Norm Equivalence

Lemma 6.1 can be restated in the following useful form.

Lemma A.1 *Let $\mathbf{\Phi}$ have mutual incoherence M. Then, for any N distinct $\phi_j \in \mathbf{\Phi}$, and for any a_j, $j = 1, \ldots, N$ we have*

$$(\sum_{j=1}^N a_j^2)(1 - M(N-1)) \leq \| \sum_{j=1}^N a_j \phi_j \|_2^2 \leq (\sum_{j=1}^N a_j^2)(1 + M(N-1)).$$

Proof. We have

$$\| \sum_{j=1}^N a_j \phi_j \|_2^2 = \sum_{j=1}^N a_j^2 + \sum_{i \neq j} a_i a_j \langle \phi_i, \phi_j \rangle. \tag{A-3}$$

Next,

$$|\sum_{i \neq j} a_i a_j \langle \phi_i, \phi_j \rangle| \leq M \sum_{i \neq j} |a_i a_j| = M(\sum_{i,j} |a_i a_j| - \sum_{i=1}^N a_i^2)$$

$$= M((\sum_{i=1}^N |a_i|)^2 - \sum_{i=1}^N a_i^2) \leq (\sum_{i=1}^N a_i^2)(N-1)M.$$

Using this inequality to both upper and lower bound $\sum_{i\neq j} a_i a_j \langle \phi_i, \phi_j \rangle$ in the identity (A-3) gives the desired conclusion. □

This has an immediate application, valuable below.

Lemma A.2 *Let M be the mutual incoherence of $\mathbf{\Phi}$. Assume that f has a representation*

$$f = \sum_{i=1}^{N} a_i \psi_i + z, \quad \psi_i \in \mathbf{\Phi}, \quad i = 1, \ldots, N, \quad N < \frac{1}{2M} + \frac{1}{2}$$

with $\|z\| \leq \epsilon$. Let $g = \sum_{j=1}^{N} b_j \varphi_j$, $\varphi_j \in \mathbf{\Phi}$, $j = 1, \ldots, N$ with the property

$$\|f - g\| \leq \delta,$$

we have

$$\sum_{i:\psi_i \neq \varphi_j, j=1,\ldots,N} a_i^2 + \sum_{i:\psi_i = \varphi_j} (a_i - b_j)^2 + \sum_{j:\varphi_j \neq \psi_i, i=1,\ldots,N} b_j^2 \leq \frac{(\epsilon + \delta)^2}{1 - M(2N-1)}. \tag{A-4}$$

Proof. Writing $f = f_0 + z$, where $f_0 = \sum_{i=1}^{N} a_i \psi_i$, we have $\|f_0 - g\| \leq \|f - g\| + \|z\| \leq \delta + \epsilon$. But $f_0 - g$ can be written as a sum of at most $2N$ terms from $\mathbf{\Phi}$, $\sum_{l=1}^{2N} c_\ell \phi_\ell$, $\phi_\ell \in \mathbf{\Phi}$. Applying Lemma A.1 to this sum gives the result. □

We note that this lemma gives an alternate proof of the stability estimate (2.2) in Theorem 2.1. For indeed we let $f \leftrightarrow y$ and $g \leftrightarrow \hat{x} = \mathbf{\Phi}\hat{\alpha}_{0,\epsilon}$, and we apply the lemma with $\epsilon = \delta$.

Trapping Dictionaries

The behavior of the OGA developed in Section 4.1 can be described more vividly as follows: when f is a sparse sum, then greedy term selection only selects from among the terms appearing in the definition of f: it is 'trapped'. This phenomenon holds under broader conditions than just the incoherence of the dictionary. To discuss this clearly, we isolate the phenomenon with a definition.

Definition A.3 *(F-trapping dictionary) Let F be an increasing function on $[1, \Omega)$. The dictionary $\mathbf{\Phi}$ is F-trapping (for OGA) if, for any $N \in [1, \Omega)$ and any f of the form*

$$f = \sum_{i=1}^{N} a_i \varphi_i, \tag{A-5}$$

where φ_i are N distinct elements of $\mathbf{\Phi}$, there exist $\phi_1, \ldots, \phi_m \in \mathbf{\Phi}$ with $m \leq F(N)$ such that for all $k = 0, 1, \ldots$ we have

$$f_k^o \in \mathit{span}(\phi_1, \ldots, \phi_m).$$

Note that when $F(u) = u$ for $1 \leq u \leq 1/2M$, we are simply saying what earlier in the paper would have been formulated as: the OGA preserves support properties. As an example of the potential need for generality beyond the $F(u) = u$ case emphasized here, we note that, from results in [31], it follows that the dictionary $\{\varphi_J = |J|^{-1/2} 1_J, \quad J \subset [0,1]\}$, consisting of normalized indicator functions $1_J(x)$, is F-trapping with $F(u) = 2u + 1$ and $[1, \Omega) = [1, \infty)$. Note that the dictionary is uncountable.

For an F-trapping dictionary, and any f of the form (A-5) with $N \in [1, \Omega)$, the OGA perfectly reconstructs f after $F(N)$ steps:

$$G^o_{F(N)}(f) = f.$$

As already said, incoherent dictionaries are F-trapping with $F(u) = u$. Formally:

Theorem A.4 *Let $M = M(\mathbf{\Phi})$. Then $\mathbf{\Phi}$ is F-trapping with $F(u) = u$ and $\Omega = \frac{1}{2M} + \frac{1}{2}$.*

This can be obtained immediately from the following repackaging of arguments in Section 4.1.

Observation A.5 *Assume $\mathbf{\Phi}$ has mutual incoherence M. Let $N < \frac{1}{2M} + \frac{1}{2}$. For any f of the form*

$$f = \sum_{i=1}^{N} a_i \phi_i,$$

where the ϕ_i are distinct elements of $\mathbf{\Phi}$, we have that $\varphi(f) \in \{\phi_i : 1 \le i \le N\}$.

The F-trapping property can be proved without using incoherence.

Theorem A.6 *Let N be given and suppose that for this N, the dictionary $\mathbf{\Phi}$ has the property: for any distinct $\phi_0, \phi_1, \ldots, \phi_N \in \mathbf{\Phi}$ we have $\sum_{i=1}^{N} |\langle \phi_0, \phi_i \rangle| < 1/2$. Then for any f of the form*

$$f = \sum_{i=1}^{k} a_i \phi_i, \quad k \le N,$$

where ϕ_i are distinct elements of $\mathbf{\Phi}$, we have that $\varphi(f) \in \{\phi_i : 1 \le i \le N\}$.

The idea to replace coherence conditions (i.e. on $M(\mathbf{\Phi})$) by conditions on N-term sums of entries in individual rows of the Gram matrix $(\langle \varphi_i, \varphi_j \rangle)_{i,j=1}^{m}$ originated in [12, 36].

The trapping approach allows a proof of Claim (a) in Theorem 2.1.

Theorem A.7 *Assume $\mathbf{\Phi}$ has mutual incoherence M. Then if f has a representation*

$$f = \sum_{i=1}^{N} a_i \psi_i,$$

where the a_i are nonzero and the ψ_i are distinct elements of $\mathbf{\Phi}$ and $N < \frac{1}{2M} + \frac{1}{2}$ then it is the unique such representation.

Proof. We first remark that $(\psi_i)_{i=1}^{N}$ must be linearly independent; for example using arguments from Section 2 and $N < (1 + M^{-1})/2$. Hence there can't be two distinct representation using that specific collection of N atoms.

Now suppose there were 2 distinct expansions, each using $\le N$ atoms, but different collections of $\le N$ atoms We could then also write $f = \sum_{j=1}^{m} b_j \phi_j$ with nonzero coefficients (b_j), where $m \le N$, and $\{\phi_j\}$ contains at least one atom not belonging to $\{\psi_i\}$. Incoherence allows to apply the trapping principle, which guarantees that $\varphi(f_k^o) \in \{\psi_i\}$ for $k = 0, \ldots, N-1$, while it would also guarantee that $\varphi(f_k^o) \in \{\phi_j\}$ for $k = 0, \ldots, m-1$. Thus we must have that the first m selections lie in *both* $\{\psi_i\}$ and $\{\phi_j\}$. The OGA will never select the same atom twice. Thus the first m selections are distinct, and so

exhaust $\{\phi_j\}$. But $m = \#\{\phi_j\} \le N = \#\{\psi_i\}$ and so we conclude that $\{\phi_j\} \subset \{\psi_i\}$, contradicting the existence of a 'distinct' atoms being used. □

The upper bound for N in Theorem A.7 is sharp. We cannot replace the condition $N < \frac{1}{2M} + \frac{1}{2}$ by $N \le \frac{1}{2M} + \frac{1}{2}$ in the theorem. Here is the corresponding example. We set $\mathbf{\Phi} = \{\phi_i\}_{i=1}^4 \subset \mathbb{R}^3$ with $\phi_1 = 3^{-1/2}(1,1,1)$, $\phi_2 = 3^{-1/2}(1,-1,-1)$, $\phi_3 = 3^{-1/2}(1,1,-1)$, $\phi_4 = 3^{-1/2}(1,-1,1)$. Then $M = 1/3$ and we have nonuniqueness for $S = 2$: $\phi_1 + \phi_2 = \phi_3 + \phi_4$.

Proof of Theorem 4.6

We restate the Theorem so that y is replaced by f, etc. We first show that at initial stages, where PGA extracts terms with coefficient exceeding η, the trapping property holds. Formally, if $\phi_j \in \mathbf{\Phi}$, $j = 1, \ldots, N$ and f is such that

$$\|f - \sum_{i=1}^{N} a_i \phi_i\| \le \epsilon \quad \text{and} \quad |\langle f, \varphi(f)\rangle| > \eta,$$

then $\varphi(f) \in \{\phi_j : 1 \le j \le N\}$. Indeed, denote

$$A \equiv \max_{1 \le i \le N} |a_i| = |a_q|; \quad B \equiv |\langle f, \varphi(f)\rangle|;$$

then

$$A(1 - M(N-1)) - \epsilon \le B \le A(1 + M(N-1)) + \epsilon.$$

Our assumption $B > \eta$ implies that

$$A(1 - M(2N-1)) > 2\epsilon.$$

Therefore,

$$\max_{1 \le j \le N} |\langle f, \phi_j\rangle| \ge |\langle f, \phi_q\rangle| \ge A(1 - M(N-1)) - \epsilon > ANM + \epsilon \ge |\langle f, \phi\rangle|$$

for $\phi \in \mathbf{\Phi}$ different from ϕ_j, $j = 1, \ldots, N$. Hence $\varphi(f)$ belongs to $\{\phi_j : 1 \le j \le N\}$.

Let then $m = m(\eta, f)$ be the stopping stage for the PGA; we may check that $m \le N$. Because of the trapping property, we may write $G_m(f) = \sum_{j=1}^{N} b_j \phi_j$ for coefficients (b_j) of which at most m are nonzero. Invoking Lemma A.2 with f as above and $g = G_m(f)$ we complete the proof. □

Proof of Theorem 4.7

For given $\epsilon > 0$ and $\rho \in (0, 1]$ we denote by $m(\epsilon, \rho)$ the smallest m such that there exists a realization $G_m^{o,\rho}(f)$ with the property

$$\|f - G_{m-1}^{o,\rho}(f)\| > \epsilon \quad \text{and} \quad \|f - G_m^{o,\rho}(f)\| \le \epsilon.$$

We restate Theorem 4.7.

Theorem A.8 *Let $\rho \in (0, 1]$ and $M = M(\mathbf{\Phi})$. Let $N \in \mathbb{N}$ fulfill $\rho - M((1+\rho)N - \rho) > 0$. Suppose that*

$$f = \sum_{j=1}^{S} a_j \phi_j + z, \quad \phi_j \in \mathbf{\Phi}, \quad j = 1, \ldots, S$$

satisfies the conditions

$$S \le N; \quad \epsilon \equiv \|z\|.$$

Suppose also that this is the smallest S for which we can write such a decomposition with $\|z\| \le \epsilon$, and that

$$\epsilon < \frac{\rho - M((1+\rho)N - \rho)}{1+\rho} \cdot \min_j |a_j|. \tag{A-6}$$

Then $m(\epsilon, \rho) = S$,

$$G_S^{o,\rho}(y) = \sum_{j=1}^{S} b_j \varphi_j, \tag{A-7}$$

and

$$\sum_{j=1}^{S} (a_j - b_j)^2 \le \frac{\epsilon^2}{1 - M(N-1)}.$$

Proof. We first show that at initial stages of WOGA the trapping property holds. Put $A \equiv \max_j |a_j| = |a_p|$, and suppose

$$\epsilon < \frac{\rho - M((1+\rho)N - \rho)}{1+\rho} \cdot A. \tag{A-8}$$

Then we have

$$|\langle f, \varphi_p \rangle| \ge A(1 - M(S-1)) - \epsilon$$

and for any ϕ different from $\varphi_1, \ldots, \varphi_S$

$$|\langle f, \phi \rangle| \le AMS + \epsilon.$$

By (A-8) we have

$$AMS + \epsilon < \rho(A(1 - M(S-1) - \epsilon). \tag{A-9}$$

Hence $\varphi_1^{o,\rho}$ of the WOGA coincides with one of the $\varphi_1, \ldots, \varphi_S$, say, $\varphi_1^{o,\rho} = \varphi_{i_1}$ - i.e. we have trapping.

We now argue as in Section 4.1 that the WOGA, in passing from step 1 to step 2 to ... to step S lands us each time back in the same situation, of studying f_1, f_2, ... which have the same structure at each iteration as a superposition $f_k = \sum_{j=1}^{N} b_j^{(k)} \phi_j$, with possibly different coefficients $b^{(k)}$. The condition (A-6) guarantees that at each step until step S, $A = \max_j |b_j^{(k)}|$ obeys (A-8) and hence (A-9) follows; trapping happens at all steps $1, \ldots, S$.

Writing $f = f_0 + z$ with $f_0 \in H = span\{\phi_j, j = 1, \ldots, S\}$, we have $G_S^{o,\rho}(f) \in H$ as well. Hence, $\|f - G_S^{o,\rho}(f)\| = \|(f_0 + z) - P_H(f_0 + z))\| \le \|(I - P_H)(z)\| \le \epsilon$. Therefore, $m(\epsilon, \rho) \le S$. However, by hypothesis, S is minimal for the ability to write f as a superposition of S terms from $\mathbf{\Phi}$ with an error of size $\le \epsilon$. Hence $m(\epsilon, \rho) = S$. We now invoke Lemma A.2 with $g = G_S^{o,\rho}(f)$. □

Proof of Theorem 7.2

We reformulate Theorem 7.2 to simplify the proof.

Theorem A.9 *Let $\rho \in (0, 1]$ and $M = M(\mathbf{\Phi})$. Let $N \in \mathbb{N}$ fulfill $\rho - M((1+\rho)N - \rho) > 0$. Then for any signal*

$$f = \sum_{j=1}^{S} a_j \varphi_j + z, \quad \varphi_j \in \mathbf{\Phi}, \quad j = 1, \ldots, S$$

satisfying the following conditions

$$S \le N; \quad \|z\| \le \epsilon$$

we have

$$\|f - G_S^{o,\rho}(f)\| \le \epsilon S^{1/2}(1 + \frac{(1+\rho)(1+M(S-1))^{1/2}}{\rho - M((1+\rho)N - \rho)}).$$

Proof. We start with a central observation. Denote

$$\delta \equiv \frac{\epsilon(1+\rho)}{\rho - M((1+\rho)N - \rho)}.$$

Let g be a function which can be written in a fashion similar to f, $g = \sum_j b_j \varphi_j + z$ as a superposition of the same elements, with possibly different coefficients. Suppose $\{b_j\}$ are small: $|b_j| \le \delta$, $j = 1, \ldots, S$. By Lemma A.1 we get

$$\|g\| \le \|\sum_{j=1}^{S} b_j \varphi_j\| + \|z\| \le \delta S^{1/2}(1 + M(S-1))^{1/2} + \epsilon \tag{A-10}$$

$$= \epsilon S^{1/2}(1 + \frac{(1+\rho)(1+M(S-1))^{1/2}}{\rho - M((1+\rho)N - \rho)}).$$

We now apply WOGA term selection to $f = \sum_j a_j \varphi_j$. Put $A \equiv \max_j |a_j| = |a_p|$ say. Suppose $A > \delta$. Using the same arguments as in the proof of Theorem 4.7, we find that $\varphi_1^{o,\rho}$ of the WOGA coincides with one of the $\varphi_1, \ldots, \varphi_S$, say, $\varphi_1^{o,\rho} = \varphi_{i_1}$ - i.e. we have trapping.

Denote $\Lambda_1 \equiv \{i_1\}$ and $H_1 \equiv \text{span}\{\varphi_{i_1}\}$. Consider now

$$f_1 \equiv f_1^{o,\rho} = f - G_1^{o,\rho}(f) = f - P_{H_1}(f).$$

It is clear that

$$f_1 = \sum b_j^{(1)} \varphi_j + z_1, \quad \|z_1\| \le \epsilon.$$

In applying WOGA term selection to f_1, we get a new value of $A = A_1 \equiv \max_j |b_j^{(1)}|$; but supposing that $A_1 > \delta$, trapping continues. Hence the next iteration of WOGA will have its term selection among $\varphi_1, \ldots, \varphi_S$.

We now argue (recalling Section 4.1) that in studying f_1, f_2, ..., f_k, as long as $A = A_k \equiv \max_j |b_j^{(k)}| > \delta$ continues to hold in successive steps, the trapping continues, and we repeatedly see the same structure at each iteration: a superposition $f_k = \sum_{j=1}^{N} b_j^{(k)} \phi_j$, with possibly different coefficients $b^{(k)}$.

So suppose that we reach $A_k < \delta$ at stage $k < S$. then the estimate (A-10) already applies at that early stage (with the more favorable value of k in place of S). As the algorithm continues through S steps, the error at stage S will not be larger than the error at stage k, and so the estimate (A-10) applies and gives the desired conclusion.

Suppose we do not reach $A_k < \delta$ at any stage up to and including the S-th stage. Then, with $H = span\{\varphi_1, \ldots, \varphi_S\}$, we have $G_S^{o,\rho}(f) = P_H(f)$, and by the hypothesis $\|z\| \le \epsilon$

$$\|f - G_S^{o,\rho}(f)\| \le \epsilon.$$

which is even better than we claimed. □

Proof of Theorem 7.3

We begin with a simple technical lemma.

Theorem A.10 *Suppose that $\phi_1, \ldots, \phi_N$ are such that $\|\phi_i\| = 1$, $i = 1, \ldots, N$; $|\langle \phi_i, \phi_j \rangle| \le M$, $1 \le i \ne j \le N$. Let $H_N \equiv span(\phi_1, \ldots, \phi_N)$. Then for any f we have*

$$\sum_{i=1}^{N} |\langle f, \phi_i \rangle|^2 \ge (\sum_{i=1}^{N} c_i^2)((1+M)^2 - (2M + M^2)N),$$

where $\{c_i\}$ are coefficients from least-squares projection:

$$P_{H_N}(f) = \sum_{i=1}^{N} c_i \phi_i.$$

Proof. We have $\langle f - P_{H_N}(f), \phi_i \rangle = 0$, $i = 1, \ldots, N$ and therefore

$$|\langle f, \phi_i \rangle| = |\langle P_{H_N}(f), \phi_i \rangle| = |\sum_{j=1}^{N} c_j \langle \phi_j, \phi_i \rangle| \ge |c_i|(1+M) - M \sum_{j=1}^{N} |c_j|.$$

Next, denoting $\sigma \equiv \sum_{j=1}^{N} |c_j|$ we get

$$\sum_{i=1}^{N} |\langle f, \phi_i \rangle|^2 \ge \sum_{i=1}^{N} (|c_i|(1+M) - M\sigma)^2 = (1+M)^2 \sum_{i=1}^{N} c_i^2 - 2(1+M)M\sigma^2 + M^2\sigma^2.$$

Using the inequality $\sigma^2 \le N \sum_{j=1}^{N} c_j^2$ we continue

$$\ge (\sum_{i=1}^{N} c_i^2)(1 - (2M + M^2)(N-1)).$$

□

Theorem 7.3 also depends on this intermediate result.

Theorem A.11 *Let $M = M(\mathbf{\Phi})$. Then for any $S \le 1/(4M)$ we have the following inequalities*

$$\|f_S^o\|^2 \le 2\|f\|(\sigma_S(f) + 6MS\|f\|),$$

$$\|f_S\|^2 \le 2\|f\|(\sigma_S(f) + 7MS\|f\|).$$

Proof. As earlier, denote

$$\chi(f) \equiv \sup_{\phi \in \mathbf{\Phi}} |\langle f, \phi \rangle|. \tag{A-11}$$

For simplicity we assume that the maximizer in (A-11) exists. Then

$$\|f_m\|^2 = \|f_{m-1}\|^2 - \chi(f_{m-1})^2 \quad \text{and} \quad \|f_m^o\|^2 \le \|f_{m-1}^o\|^2 - \chi(f_{m-1}^o)^2.$$

Denote by $\phi_1, \ldots, \phi_S \subset \mathbf{\Phi}$ the elements that have the biggest inner products with f:

$$|\langle f, \phi_1\rangle| \geq |\langle f, \phi_2\rangle| \geq \ldots \geq |\langle f, \phi_S\rangle| \geq \sup_{\phi\in\mathbf{\Phi}, \phi\neq\phi_i, i=1,\ldots,S} |\langle f, \phi\rangle|.$$

We carry out the proof for the OGA and later point out the necessary changes for the PGA. Let $m \leq S$ and

$$f_m^o = f - P_{H_m}(f), \quad H_m = \mathrm{span}(\varphi_1, \ldots, \varphi_m), \quad \varphi_j \in \mathbf{\Phi}.$$

There exists an index $i \in [1, m+1]$ such that $\phi_i \neq \varphi_j$, $j = 1, \ldots, m$. For this i we estimate

$$\langle f_m^o, \phi_i\rangle = \langle f, \phi_i\rangle - \langle P_{H_m}(f), \phi_i\rangle. \tag{A-12}$$

Let

$$P_{H_m}(f) = \sum_{j=1}^{m} c_j \varphi_j.$$

Clearly, $\|P_{H_m}(f)\| \leq \|f\|$. Then by Lemma A.1

$$(\sum_{j=1}^{m} c_j^2)^{1/2} \leq \|f\|(1 - M(m-1))^{-1/2}.$$

We continue

$$|\langle P_{H_m}(f), \phi_i\rangle| \leq M \sum_{j=1}^{m} |c_j| \leq M m^{1/2} (\sum_{j=1}^{m} c_j^2)^{1/2} \leq M S^{1/2} \|f\| (1 - MS)^{-1/2}. \tag{A-13}$$

Thus we get from (A-12) and (A-13) that

$$\chi(f_m^o) \geq |\langle f_m^o, \phi_i\rangle| \geq |\langle f, \phi_i\rangle| - M S^{1/2} \|f\| (1 - MS)^{-1/2}.$$

Therefore

$$(\sum_{v=0}^{S-1} \chi(f_v)^2)^{1/2} \geq (\sum_{i=1}^{S} |\langle f, \phi_i\rangle|^2)^{1/2} - MS\|f\|(1 - MS)^{-1/2}. \tag{A-14}$$

Next, let

$$\sigma_S(f) = \|f - P_{H(S)}(f)\|, \quad P_{H(S)}(f) = \sum_{j=1}^{S} b_j \psi_j.$$

Then

$$\|P_{H(S)}(f)\| \geq \|f\| - \sigma_S(f)$$

and by Lemma A.1

$$\sum_{j=1}^{S} b_j^2 \geq (\|f\| - \sigma_S(f))^2 (1 + MS)^{-1}. \tag{A-15}$$

By Theorem A.10

$$\sum_{j=1}^{S} |\langle f, \psi_j\rangle|^2 \geq (\sum_{j=1}^{S} b_j^2)(1 - 3MS). \tag{A-16}$$

We get from (A-15) and (A-16)

$$\sum_{i=1}^{S}|\langle f,\phi_i\rangle|^2 \geq \sum_{j=1}^{S}|\langle f,\psi_j\rangle|^2 \geq (\|f\|-\sigma_S(f))^2(1+MS)^{-1}(1-3MS).$$

Finally, by (A-14) we get from here

$$(\sum_{v=0}^{S-1}\chi(f_v^o)^2)^{1/2} \geq (\|f\|-\sigma_S(f))(\frac{1-3MS}{1+MS})^{1/2} - MS\|f\|(1-MS)^{-1/2}$$

and

$$\begin{aligned}\|f_S^o\|^2 \leq \|f\|^2 - \sum_{v=0}^{S-1}\chi(f_v)^2 &\leq 2\|f\|(\|f\| - (\sum_{v=0}^{S-1}\chi(f_v)^2)^{1/2}) \\ &\leq 2\|f\|(\sigma_S(f)+6MS\|f\|).\end{aligned}$$

This completes the proof of Theorem A.11 for the OGA. A few changes adapt the proof to the PGA setting. As above we write

$$f_m = f - G_m(f); \quad G_m(f) = \sum_{j=1}^{m} b_j\psi_j, \quad \psi_j \in \mathbf{\Phi}$$

and estimate $|\langle f_m,\phi_i\rangle|$ with $i\in[1,m+1]$ such that $\phi_i \neq \psi_j$, $j=1,\dots,m$. Using instead of $\|P_{H_m}(f)\| \leq \|f\|$ the inequality

$$\|G_m(f)\| \leq \|f\| + \|f_m\| \leq 2\|f\|$$

we obtain the following analog of (A-13)

$$|\langle G_m(f),\phi_i\rangle| \leq 2MS^{1/2}\|f\|(1-MS)^{-1/2}. \tag{A-17}$$

The rest of the proof is the same with (A-13) replaced by (A-17). □

We now show how to combine the inequalities from Theorems A.10 and A.11 to get Theorem 7.3. Write these inequalities in the form ($S \leq N$)

$$\|f_S^o\|^2 \leq C_1 S\sigma_S(f)^2, \tag{A-18}$$

$$\|f_S^o\|^2 \leq 2\|f\|(\sigma_S(f)+C_2MS\|f\|). \tag{A-19}$$

We take m such that $3m \leq N$ and apply (A-19) with $f=f_m^o$, $S=2m$. We get

$$\|f_{3m}^o\|^2 \leq 2\|f_m^o\|(\sigma_{2m}(f_m^o)+C_2M(2m)\|f_m^o\|).$$

Next we use a trivial estimate $\sigma_{2m}(f_m^o)\leq\sigma_m(f)$ and (A-18)

$$\|f_{3m}^o\|^2 \leq 2(C_1m)^{1/2}\sigma_m(f)^2(1+C_3Mm^{3/2}).$$

Therefore for small m ($Mm^{3/2}\leq 1$) we get

$$\|f_{3m}^o\| \leq C_4m^{1/4}\sigma_m(f) \tag{A-20}$$

which is better then (A-18).

We can repeat the above argument with (A-18) replaced by (A-20). Continuing in this way, we obtain Theorem 7.3. □

Appendix B - Iteratively Reweighted Least-Squares

In simulating the Basis Pursuit Denoising, we need to solve the optimization problem

$$(P_{1,\delta}): \quad \min_{\alpha} \ \|\alpha\|_1 \ \text{ subject to } \ \|y - \mathbf{\Phi}\alpha\|_2 \leq \delta. \tag{B-1}$$

It is easier to treat the constraint as an additional penalty and form the alternative problem

$$(\tilde{P}_{1,\lambda}): \quad \min_{\alpha} \ \|\alpha\|_1 + \lambda\|y - \mathbf{\Phi}\alpha\|_2, \tag{B-2}$$

but then we need to search for the proper value of λ so as to satisfy the constraint $\|y - \mathbf{\Phi}\alpha\|_2 \leq \delta$. Thus, we are facing two difficulties: (i) For a given λ solve (B-2); and (ii) Find the value of λ to satisfy the constraint.

The problem posed in (B-2) can be cast in a Quadratic Programming form, and solved as such. However, we chose a different way that was found to be more efficient in the low dimensions we experiment with. The penalty function in (B-2) can be written alternatively as

$$(\tilde{P}_{1,\lambda}): \quad \min_{\alpha} \ \alpha^H \mathbf{W}(\alpha)\alpha + \lambda\|y - \mathbf{\Phi}\alpha\|_2, \tag{B-3}$$

where $\mathbf{W}(\alpha)$ is a diagonal weight matrix with $1/|\alpha_k|$ as its main diagonal entries (assume that for $\alpha_k = 0$ the weight is chosen to be some finite high value in order to avoid infinity). Formed as such, we can use simple Least-Squares to solve $(\tilde{P}_{1,\lambda})$ with $\mathbf{W}(\alpha)$ assumed to be fixed. This solution is obtained as

$$\hat{\alpha}_{opt} = \left(\mathbf{W}(\alpha) + \lambda\mathbf{\Phi}^H\mathbf{\Phi}\right)^{-1}\mathbf{\Phi}^H y. \tag{B-4}$$

The basic idea of the Iterative Reweighted Least-Squares is an initialization by some $\hat{\alpha}_0$, computing the weights as described above, and then update the solution using (B-4). Relaxation of this process could be proposed where the new solution at the n^{th} iteration is computed as

$$\hat{\alpha}_n = \beta \cdot \left(\mathbf{W}(\hat{\alpha}_{n-1}) + \lambda\mathbf{\Phi}^H\mathbf{\Phi}\right)^{-1}\mathbf{\Phi}^H y + (1-\beta)\hat{\alpha}_{n-1}. \tag{B-5}$$

β can be chosen in the range $[0, 1]$ trading stability with speed. Analysis of this algorithm and its convergence properties can be found in [20].

As to the search for λ, our implementation applies a line search process, exploiting the fact that the function we are optimizing with respect to is unimodal.

References

[1] A.P. Berg & W.B. Mikhael, A survey of mixed transform techniques for speech and image coding, In *Proceedings of the 1999 IEEE International Symposium on Circuits and Systems*, Volume 4, pages 106–109, 1999.

[2] Candès, E.J. and Donoho, D.L. (2002) New tight frames of curvelets and the problem of approximating piecewise C^2 images with piecewise C^2 edges. to appear *Comm. Pure Appl. Math.*

[3] S.S. Chen, *Basis Pursuit*, PhD thesis, Stanford University, November 1995.

[4] Chen, S.S., Donoho, D.L. & Saunders, M.A. (2001) Atomic decomposition by basis pursuit, *SIAM Review*, Volume 43, number 1, pages 129–59.

[5] R. Coifman, Y. Meyer, and V. Wickerhauser. Adapted wave form analysis, wavelet-packets and applications. In *ICIAM 91. Proceedings of the Second International Conference on Industrial and Applied Mathematics*, pages 41–50, Philadelphia, PA, 1992. SIAM.

[6] S.F. Cotter, K. Kreutz-Delgado and B.D. Rao (2002) Efficient backward elimination algorithm for sparse signal representation using overcomplete dictionaries, *IEEE Signal Processing Letters*, **9**, 145 -147.

[7] S.F. Cotter and B.D. Rao (2002) Sparse channel estimation via matching pursuit with application to equalization. *IEEE Trans. Communications* **50**, March 2002

[8] Daniel, Cuthbert and Wood, Fred S. [1980]. *Fitting Equations to Data: Computer Analysis of Multifactor Data*, 2nd Edition, John Wiley and Sons.

[9] V.E. DeBrunner, L.X. Chen & H.J. Li, Lapped multiple basis algorithms for still image compression without blocking effect, *IEEE Trans. Image Proc.*, Volume 6, pages 1316–1322, September 1997.

[10] DeVore, R.A. and Temlyakov, V.N. (1996) Some remarks on greedy algorithms, *Adv. Comput. Math.* **12** 213-227.

[11] Donoho, D.L. (1992) Superresolution via sparsity constraints, *SIAM J. Math. Anal.* **23** 1309–1331

[12] Donoho, D.L. & Elad, M. (2002) Optimally sparse representation in general (non-orthogonal) dictionaries via ℓ^1 minimization, *Proc. Nat. Aca. Sci.* **100** 2197-2202.

[13] Donoho, D.L. & Huo, X. (2001) Uncertainty principles and ideal atomic decomposition, *IEEE Trans. on Inf. Theory*, volume 47, number 7, pages 2845–62.

[14] Elad, M. & Bruckstein, A.M. (2002) A generalized uncertainty principle and sparse representation in pairs of $\Re^N$ bases, *IEEE Trans. on Inf. Theory*, Volume 48, pages 2558–2567.

[15] Gilbert, A.C. Muthukrishnan, S. and Strauss, M.J. (2003) Approximation of functions over redundant dictionaries using coherence. 14th Ann. ACM-SIAM Symposium Discrete Algorithms.

[16] I.F. Gorodnitsky and B.D. Rao, Sparse Signal Reconstruction from Limited Data Using FOCUSS: A Re-Weighted Norm Minimization Algorithm, *IEEE Trans. On Signal Processing*, pp. 600-616, Vol. 45, no. 3, March 1997.

[17] Gribonval, R. & Nielsen, M. (2002) Sparse representations in unions of bases, submitted to the *IEEE Trans. on Inf. Theory.*

[18] Hastie, T. Tibshirani, R, and Friedman, J.H. (2001) *Elements of Statistical Learning.* New York: Springer.

[19] Huo, X. (1999) Sparse Image representation Via Combined Transforms, PhD thesis, Stanford.

[20] Karlovitz, L.A. (1970) Construction of nearest points in the ℓ^p, p even and ℓ^1 norms, *Journal of Approximation Theory*, Vol. 3, pp. 123-?127.

[21] Kruskal, J.B. (1977) Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics, *Linear Algebra and Its Applications*, Vol. 18, No. 2, p. 95–138.

[22] Kolda, T.G. (2001) Orthogonal Tensor Decompositions *SIAM Journal on Matrix Analysis and Applications*, Vol. 23, No. 1, pp. 243-255.

[23] Mallat, S. & Zhang, Z. (1993) Matching Pursuit in a time-frequency dictionary, *IEEE Transactions on Signal Processing*, Volume 41, pages 3397–3415.

[24] B. A. Olshausen and D. J. Field. (1997) Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, **37**, 311-325.

[25] Pati, Y.C. & Rezaiifar, R. & Krishnaprasad, P.S. (1993) Orthogonal Matching Pursuit: Recursive Function Approximation with Applications to Wavelet Decomposition, Proceedings of the 27 th Annual Asilomar Conference on Signals, Systems, and Computers.

[26] Qian, S. and Chen, D. (1994) Signal representation using adaptive normalized Gaussian functions. *Signal Processing*, **36** 1-11.

[27] A. Ron & X. Shen, Compactly supported tight affine spline frames in $L_2(R^d)$, *Math. Comp.* Volume 65(216), pages 1513–1530, February 1996.

[28] Starck, J.-L. & Elad, M & Donoho, D.L. (2003) Image Decomposition: Separation of Texture from Piecewise Smooth Content, *SPIE meeting*, August 2003, San-Diego.

[29] Starck, J.-L., Candès, E. & Donoho, D.L. (2002) The curvelet transform for image denoising *IEEE Transactions on Image Processing* 11(6), pp. 131–141.

[30] Strohmer, T. & Heath Jr, R. (2003) Grassmannian frames with applications to coding and communcations, to appear in *Appl.Comp.Harm.Anal.*.

[31] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso, *J. Royal Statist. Soc. B.*, Volume 58, number 1, pages 267-288.

[32] Temlyakov, V.N. (1999) Greedy algorithms and m-term approximation, *J. Approx. Theory* **98** 117-145.

[33] Temlyakov, V.N. (2000) Weak greedy algorithms , *Adv. Comput. Math.* **5** 173-187.

[34] Temlyakov, V.N. (2000) Greedy Algorithms with regard to multivariate systems with special structure. *Constr. Approx.* **16** 399-425.

[35] Temlyakov, V.N. (2003) Nonlinear methods of approximation. *Found. Comput. Math.* **3** 33-107.

[36] Tropp, J.A. (2003) Greed is Good: Algorithmic results for sparse approximation. CAM Tech Report, Univ. Texas. To appear, *IEEE Trans IT*.

[37] Tropp, J.A., Gilbert, A.C., Muthukrishnan, S., & Strauss M.J. (2003) Improved sparse approximation over quasi-incoherent dictionaries. IEEE International Conference on Image Processing, Barcelona, September.

[38] Tropp, J.A. (2004) Just relax: Convex programming methods for subset selection and sparse approximation. ICES Report 04-04, UT-Austin, February.

[39] Venkataramani, R. & Bresler, Y. (2001) Optimal sub-Nyquist nonuniform sampling and reconstruction of multiband signals, *IEEE Trans. Signal Processing*, Vol. 49, No. 10, pp. 2301–2313.

[40] M.V. Wickerhauser, *Adapted Wavelet Analysis from Theory to Software*, Wellesley, 1994.

[41] Wohlberg, B. (2003) Noise sensitivity of sparse signal representations: Reconstruction error bounds for the inverse problem. *IEEE Trans. Signal Processing*, to appear.

[42] Xiangqian, L. & Sidiropoulos, N.D. (2001) Cramer-Rao lower bounds for low-rank decomposition of multidimensional arrays, *IEEE Transactions on Signal Processing*, Vol. 49, No. 9, pp. 2074–2086.