ALGORITHMIC LINEAR DIMENSION REDUCTION IN THE ℓ_1 NORM FOR SPARSE VECTORS

A. C. GILBERT, M. J. STRAUSS, J. A. TROPP, AND R. VERSHYNIN

Abstract. It has recently been observed that sparse and compressible signals can be sketched using very few nonadaptive linear measurements in comparison with the length of the signal. This sketch can be viewed as an embedding of an entire class of compressible signals into a low-dimensional space. To date, most algorithms for approximating or reconstructing the signal from the sketch, such as the linear programming approach proposed by Donoho and his collaborators, along with Candès and Tao, require time polynomial in the signal length.

This paper develops a new method for linearly embedding compressible signals of length d into a space of dimension $O(m \log^2(d))$. We present a reconstruction algorithm whose run time, $O(m \log^2(m) \log^2(d))$, is *sublinear* in the length d of the signal. The reconstruction error is within a logarithmic factor (in m) of the optimal m-term approximation error in ℓ_1. In particular, m-sparse signals are recovered perfectly and compressible signals are recovered with polylogarithmic distortion.

These sketching techniques and the corresponding reconstruction algorithms are an application of an algorithmic dimension reduction in the ℓ_1 norm. In particular, vectors of support m in dimension d can be linearly embedded into $O(m \log^2 d)$ dimensions with polylogarithmic distortion. We can reconstruct a vector from its low-dimensional sketch in time $O(m \log^2(m) \log^2(d))$. Furthermore, this reconstruction is stable under small perturbations.

Date: April 11, 2006.

Key words and phrases. Approximation, embedding, dimension reduction, sketching, sparse approximation, sublinear algorithms.

Gilbert and Tropp are with the Department of Mathematics, The University of Michigan at Ann Arbor, 2074 East Hall, 530 Church St., Ann Arbor, MI 48109-1043. E-mail: `{annacg|jtropp}@umich.edu`. ACG has been supported in part by NSF DMS 0354600. JAT has been supported by NSF DMS 0503299.

Strauss is with the Department of Mathematics and the Department of Electrical Engineering and Computer Science, The University of Michigan at Ann Arbor. E-mail: `martinjs@umich.edu`. MJS has been supported in part by NSF DMS 0354600.

Vershynin is with the Department of Mathematics, The University of California at Davis, Davis, CA 95616. E-mail: `vershynin@math.ucdavis.edu`. RV is an Alfred P. Sloan Research Fellow. He was also partially supported by the NSF grant DMS 0401032.

1. Introduction

A metric space (X, d_X) is said to embed into a metric space (Y, d_Y) with distortion D if there are positive numbers A, B such that $B/A \le D$ and a map $\mathbf{\Phi} : X \to Y$ such that

$$A\, d_X(x, y) \le d_Y(\mathbf{\Phi}(x), \mathbf{\Phi}(y)) \le B\, d_X(x, y) \quad \text{for all } x \in X,\ y \in Y. \tag{1.1}$$

A fundamental problem is to understand when a finite metric space, which is isometrically embedded in some normed space X, admits a dimension reduction; i.e., when it can be embedded in an appropriate normed space Y of low dimension. Dimension reduction techniques enjoy a wide variety of algorithmic applications, including data stream computations [CM03, GGI+02a] and approximate searching for nearest neighbors [IN05] (to cite just a few). The dimension reduction result of Johnson and Lindenstrauss [JL84] is a fundamental one. It states that any set of N points in ℓ_2 can be embedded in ℓ_2^n with distortion $(1+\epsilon)$ and where the dimension $n = O(\log(N)/\epsilon^2)$.

A similar problem in the ℓ_1 space had been a long standing open problem; it was solved in the negative in [BC03] (see another example in [NL04]). There exists a set of N points in ℓ_1 such that any embedding of it into ℓ_1^n with distortion D requires $n = N^{\Omega(1/D^2)}$ dimensions. Thus, a dimension reduction in ℓ_1 norm with constant distortion is not possible. However, it is well known how to do such a dimension reduction with a logarithmic distortion. One first embeds any N-point metric space into ℓ_2 with distortion $O(\log N)$ using Bourgain's theorem [Bou85], then does dimension reduction in ℓ_2 using Johnson-Lindenstrauss result [JL84], and finally embeds ℓ_2^n into ℓ_1^{2n} with constant distortion using Kashin's theorem ([Kas77], see Corollary 2.4 in [Pis89]). For linear embeddings $\mathbf{\Phi}$, even distortions of polylogarithmic order are not achievable. Indeed, Charikar and Sahai [CS02] give an example for which any linear embedding into ℓ_1^n incurs a distortion $\Omega(\sqrt{N/n})$.

Two fundamental questions arise from the previous discussion.

(1) What are spaces for which a dimension reduction in the ℓ_1 norm is possible with constant distortion?
(2) What are spaces for which a linear dimension reduction in the ℓ_1 norm is possible with constant or polylogarithmic distortion?

One important space which addresses question (2) positively consists of all vectors of small support. Charikar and Sahai [CS02] prove that the space of vectors of support m in dimension d can be linearly embedded into ℓ_1^n with distortion $1+\epsilon$ with respect to the ℓ_1 norm, where $n = O((m/\epsilon)^2 \log d)$ (Lemma 1 in [CS02]).

The main result of our paper in an *algorithmic* linear dimension reduction for the space of vectors of small support. The algorithm runs in sublinear time and is stable.

Theorem 1. *Let Y be a set of points in $\mathbb{R}^d$ endowed with the ℓ_1 norm. Assume that each point has non-zero coordinates in at most m dimensions. Then these points can be linearly embedded into ℓ_1 with distortion $O(\log^2(d)\log^3(m))$, using only $O(m\log^2 d)$ dimensions. Moreover, we can reconstruct a point from its low-dimensional sketch in time $O(m\log^2(m)\log^2(d))$.*

This dimension reduction reduces the quadratic order of m in [CS02] to a linear order. Our embedding does, however, incur a distortion of polylogarithmic order. In return for this polylogarithmic distortion, we gain an *algorithmic linear dimension reduction*—there exists a sublinear time algorithm that can reconstruct every vector of small support from its low-dimensional sketch.

The space of vectors of support m in dimension d is a natural and important space as it models closely the space of compressible signals. A *compressible signal* is a long signal that can be represented with an amount of information that is small relative to the length of the signal. Many classes of d-dimensional signals are compressible, e.g.,

- The m-sparse class $B_0(m)$ consists of signals with at most m nonzero entries.
- For $0 < p < 1$, the weak ℓ_p class $B_{\text{weak-}p}(r)$ contains each signal f whose entries, sorted by decaying magnitude, satisfy $|f|_{(i)} \le r\, i^{-1/p}$.

These types of signals are pervasive in applications. Natural images are highly compressible, as are audio and speech signals. Image, music, and speech compression algorithms and coders are vital pieces of software in many technologies, from desktop computers to MP3 players. Many types of automatically generated signals are also highly redundant. For example, the distribution of bytes per source IP address in a network trace is compressible—just a few source IP addresses send the majority of the traffic.

One important algorithmic application of our dimension reduction is the reconstruction of compressible signals. This paper describes a method for constructing a random linear operator $\mathbf{\Phi}$ that maps each signal f of length d to a sketch of size $O(m \log^2 d)$. Given this sketch, we exhibit an algorithm called Chaining Pursuit that constructs an $O(m)$-term approximation of the signal with an error that is within a logarithmic factor (in m) of the optimal m-term approximation error. A compressible signal is well-approximated by an m-sparse signal so the output of Chaining Pursuit is a good approximation to the original signal, in addition to being a compressed represention of the original signal. Moreover, this measurement operator succeeds simultaneously for all signals with high probability. The major innovation of this result is to combine sublinear reconstruction time with stable and robust linear dimension reduction of all compressible signals.

Theorem 2. *With probability at least $(1 - d^{-3})$, the random measurement operator $\mathbf{\Phi}$ has the following property. Suppose that f is a d-dimensional signal whose best m-term approximation with respect to ℓ_1 norm is f_m. Given the $O(m \log^2(d))$ sketch $\boldsymbol{V} = \mathbf{\Phi} f$ and the number m, the Chaining Pursuit algorithm produces a signal $\widehat{f}$ with at most $O(m)$ nonzero entries. This signal estimate satisfies*

$$\|f - \widehat{f}\|_1 \leq (1 + C \log m)\big\|f - f_m\big\|_1. \tag{1.2}$$

In particular, if $f_m = f$, then also $\widehat{f} = f$. The time cost of the algorithm is $O(m \log^2(m) \log^2(d))$.

Remark 3. *(i) The number $O(m)$ spikes in the output signal $\widehat{f}$ can be improved to $(1+\epsilon)m$ for any fixed $\epsilon \in (0, 1)$ with small modifications of the algorithm; the current version gives $1.105m$ spikes. (ii) The factor $\log m$ is intrinsic to this approach. However, the proof gives a stronger statement—the approximation in the weak-1 norm* without *that factor: $\|f - \widehat{f}\|_{\text{weak}-1} \leq C\|f - f_m\|_1$. (iii) Our argument shows that the reconstruction $\widehat{f}$ is not only stable with respect to noise in the signal, as Equation* (1.2) *shows, but also with respect to inaccuracy in the measurements. Indeed, a stronger inequality holds. For every $\boldsymbol{V}$ (not necessarily the sketch of f), we have*

$$\|f_m - \widehat{f}\|_1 \leq C \log m \Big(\|f - f_m\|_1 + \|\mathbf{\Phi} f - \boldsymbol{V}\|_1\Big).$$

1.1. **Related Work.** The problem of sketching and reconstructing m-sparse and compressible signals has several precedents in the TCS literature, especially the paper [CM03] on detecting heavy hitters in nonnegative data streams and the works [GGI$^+$02b, GMS05] on Fourier sampling. More recent papers from TCS include [CM05, CRTV05]. Sparked by the papers [Don04] and [CT04], the computational harmonic analysis and geometric functional analysis communities have produced an enormous amount of work, including [CRT04, Don05, DT05, CT05, RV05, TG05, MPTJ05].

Most of the previous work has focused on a reconstruction algorithm that involves linear programming (as first investigated and promoted by Donoho and his collaborators) or second-order cone programming [Don04, CT04, CRTV05]. Computation times have not been reported, but they are expected to be cubic in the length d of the signal. This cost is high, since we are seeking an approximation that involves $O(m)$ terms. The paper [TG05] describes another algorithm with running time of order $O(m^2 d \log d)$, which can be reduced to $O(md \log d)$ in certain circumstances. None of these approaches are comparable with the sublinear algorithms described here.

There are a few sublinear algorithms available in the literature. The Fourier sampling paper [GMS05] can be viewed as a small space, sublinear algorithm for signal reconstruction. Its primary shortcoming is that the measurements are not uniformly good for the entire signal class. The recent work [CM05] proposes some other sublinear algorithms for reconstructing compressible signals. Few of these algorithms offer a uniform guarantee. The ones that do require more measurements—$O(m^2 \log d)$ or worse—which means that they are not sketching the signal as efficiently as possible. We present a table in the Appendix which summarizes the state of the art.

1.2. **Organization.** In Section 2, we provide an overview of determining a sketch of the signal f. We give an explicit construction of the linear map $\mathbf{\Phi}$. In Section 3, we detail the reconstruction algorithm, Chaining Pursuit, and in Section 4 we give an analysis of the algorithm, proving our main result. In Section 5 we use our algorithmic analysis to derive a dimension reduction in the ℓ_1 norm for sparse vectors.

2. Sketching the Signal

This section describes a linear process for determining a sketch $\boldsymbol{V}$ of a signal f. Linearity is essential for processing additive updates to the signal. Not only is this property important for applications, but it arises during the iterative algorithm for reconstructing the signal from the sketch. Linearity also makes the computation of the sketch straightforward, which may be important for modern applications that involve novel measurement technologies.

2.1. **Overview of sketching process.** We use the following notation, similar to [CM05].

Definition 4. *Let P be a p-by-d matrix and Q a q-by-d matrix, with rows $\{P_i : 0 \le i < p\}$ and $\{Q_j : 0 \le j < q\}$, respectively. The* row tensor product $S = P \otimes_r Q$ *of P and Q is a pq-by-d matrix whose rows are $\{P_i Q_j : 0 \le i < p,\ 0 \le j < q\}$, where $P_i Q_j$ denotes the componentwise product of two vectors of length d.*

Formally, the measurement operator $\mathbf{\Phi}$ is a row tensor product $\mathbf{\Phi} = \boldsymbol{B} \otimes_r \boldsymbol{A}$. Here, $\boldsymbol{A}$ is a $O(m \log d) \times d$ matrix called the *isolation matrix* and $\boldsymbol{B}$ is a $O(\log d) \times d$ matrix called the *bit test matrix*. The measurement operator applied to a signal f produces a sketch $\boldsymbol{V} = \mathbf{\Phi} f$, which we can regard as a matrix with dimensions $O(m \log d) \times O(\log d)$. Each row of $\boldsymbol{V}$ as a matrix contains the result of the bit tests applied to a restriction of the signal f by a row of A. We will refer to each row of the data matrix as a *measurement* of the signal.

2.2. **The isolation matrix.** The isolation matrix $\boldsymbol{A}$ is a 0–1 matrix with dimensions $O(m \log d) \times d$ and a hierarchical structure. Let a be a sufficiently large constant, to be discussed in the next two sections. The Chaining Pursuit algorithm makes $K = 1 + \log_a m$ passes over the signal, and it requires a different set of measurements for each pass. The measurements for the kth pass are contained in the $O(mk/2^k) \times d$ submatrix $\boldsymbol{A}^{(k)}$. During the kth pass, the algorithm performs $T_k = O(k \log d)$ trials. Each trial t is associated with a further submatrix $\boldsymbol{A}_t^{(k)}$, which has dimensions $O(m/2^k) \times d$.

Each trial submatrix $\boldsymbol{A}_t^{(k)}$ encodes a random partition of the d signal positions into $O(m/2^k)$ subsets. That is, each signal position is assigned uniformly at random to one of $O(m/2^k)$ subsets. So the matrix contains a 1 in the (i, j) position if the jth component of the signal is assigned to subset i. Therefore, the submatrix $\boldsymbol{A}_t^{(k)}$ is a 0–1 matrix with exactly one entry in each column equal to one. The trial submatrix can also be viewed as a random linear hash function from a space of d keys onto a set of $O(m/2^k)$ buckets.

2.3. **The bit test matrix.** Formally, $\boldsymbol{B} = [\,\mathbf{e} \mid \boldsymbol{B}_0\,]^{\mathrm{T}}$ is a 0–1 matrix. Here, the symbol $\mathbf{e}$ represents the d-dimensional vector of ones. The matrix $\boldsymbol{B}_0$ has dimensions $d \times \log_2\lceil d \rceil$. Its ith column has a one in each row whose index has a zero in the ith digit of its binary expansion. (Here we use the convention that the rows of the matrix are indexed beginning with zero for the bottom row.) Therefore, the componentwise product of the ith row of B with f yields a copy of the signal f with the components that have bit i equal to one selected and the others zeroed out.

2.4. **Storage costs.** The bit test matrix requires no storage. The total storage for the isolation matrix is $O(d \log d)$.

2.5. **Encoding time.** The time cost for measuring a signal is $O(\log^2(m) \log^2(d))$ *per nonzero component*. This claim follows by observing that a single column of $\boldsymbol{A}$ contains $O(\log^2(m) \log(d))$ nonzero entries, and we must apply $\boldsymbol{A}$ to each of $O(\log d)$ copies of the signal. Note that this argument assumes random access to the columns of the isolation matrix. We will use this encoding time calculation when we determine the time costs of the Chaining Pursuit algorithm.

3. Signal Approximation with Chaining Pursuit

Suppose that the original signal f is well approximated by a signal with m nonzero entries (spikes). The goal of the Chaining Pursuit algorithm is to use a sketch of the signal to obtain a signal approximation with no more than $O(m)$ spikes. In each pass, the algorithm recovers a constant fraction of the remaining spikes. Then it sketches the recovered spikes and updates the data matrix to reflect the residual signal. After $O(\log m)$ passes, the residual has no significant entries remaining.

The reason for the name "Chaining Pursuit" is that this process decomposes the signal into pieces with supports of geometrically decreasing sizes. It resembles an approach in analysis and probability, also called chaining, that is used to control the size of a function by decomposing it into pieces with geometrically decreasing sizes. A famous example of chaining in probability is to establish bounds on the expected supremum of an empirical process [Tal05]. For an example of chaining in TCS, see [IN05].

3.1. **Overview of Algorithm.** The structure of the Chaining algorithm is similar to other sublinear approximation methods described in the literature [GGI$^+$02b]. First, the algorithm identifies spike locations and estimates the spike magnitudes. Then it encodes these spikes and subtracts them from the sketch to obtain an implicit sketch of the residual signal. These steps are repeated until the number of spikes is reduced to zero. The number a that appears in the statement of the algorithm is a sufficiently large constant that will be discussed further in Section 4.

Algorithm: Chaining Pursuit
Inputs: Number m of spikes, the sketch $\boldsymbol{V}$, the isolation matrix $\boldsymbol{A}$
Output: A list of $O(m)$ spike locations and values
For each pass $k = 0, 1, \dots, \log_a m$:
For each trial $t = 1, 2, \dots, O(k \log d)$:
For each measurement $n = 1, \dots, O(m/2^k)$
Use bit tests to identify the spike position
Use a bit test to estimate the spike magnitude
Retain m/a^k distinct spikes with values largest in magnitude
Retain spike positions that appear in more than 9/10 of trials
Estimate final spike sizes using medians
Encode the spikes using the measurement operator
Subtract the encoded spikes from the sketch

3.2. **Implementation.** Most of the steps in this algorithm are straightforward to implement using standard abstract data structures. The only point that requires comment is the application of bit tests to identify spike positions and values.

Recall that a measurement is a row of the sketch matrix, which consists of $\log_2\lceil d \rceil + 1$ numbers:

$$\left[\begin{array}{cccc|c} b(1) & b(2) & \dots & b(\log_2\lceil d \rceil) & c \end{array}\right].$$

The number c arises from the final column of the bit test matrix. We obtain an (estimated) spike location from these numbers as follows. If $|b(i)| \geq |c - b(i)|$, then the ith bit of the location is zero. Otherwise, the ith bit of the location is one. To estimate the value of the spike from the measurements, we use the least significant bit as follows. If $|b(1)| > |c - b(1)|$, then the estimated size of the spike is $b(1)$. Otherwise, the estimated size is $c - b(1)$. Any of the other bits would work just as well, or one may take a median over all the bits.

Recall that each measurement arises by applying the bit test matrix to a copy of the signal restricted to a subset of its components. In Section 4, we show that the estimated location and value are accurate if the subset contains a single large component of the signal and the other components have smaller ℓ_1 norm.

We encode the recovered spikes by accessing the columns of the isolation matrix corresponding to the locations of these spikes and then performing a sparse matrix–vector multiplication. Note that this step requires random access to the isolation matrix.

3.3. **Storage costs.** The primary storage cost derives from the isolation matrix $\boldsymbol{A}$. Otherwise, the algorithm requires only $O(m \log d)$ working space.

3.4. **Time costs.** During pass k, the primary cost of the algorithm occurs when we encode the recovered spikes. The number of recovered spikes is at most $O(m/a^k)$, so the cost of encoding these spikes is $O(ma^{-k} \log^2(m) \log^2(d))$. The cost of updating the sketch is the same. Summing over all passes, we obtain $O(m \log^2(m) \log^(d))$ total running time.

4. Analysis of Chaining Pursuit

This section contains a detailed analysis of the Chaining Pursuit algorithm, which yields the following theorem. Fix an isolation matrix $\boldsymbol{A}$ which satisfies the conclusions of Theorem 6 of the sequel.

Theorem 5 (Chaining Pursuit)**.** *Suppose that f is a d-dimensional signal whose best m-term approximation with respect to ℓ_1 norm is f_m. Given the sketch $\boldsymbol{V} = \boldsymbol{\Phi} f$ and the number m, Chaining Pursuit produces a signal $\widehat{f}$ with at most $O(m)$ nonzero entries. This signal estimate satisfies*

$$\|f - \widehat{f}\|_1 \leq (1 + C \log m) \|f - f_m\|_1.$$

In particular, if $f_m = f$, then also $\widehat{f} = f$.

4.1. **Challenges of the analysis.** Chaining Pursuit is an iterative algorithm. The major difficulty of its analysis is to control the approximation error from blowing up in a geometric progression from pass to pass. In pass $k = 0$, the algorithm is working with measurements of the original signal f. This signal can be decomposed as $f = f_m + w$, where f_m is the best m-term approximation of f (*spikes*) and w is the remainder of the signal, called *external noise*. If $w = 0$, the analysis becomes quite simple. Indeed, in that case we exactly recover a constant fraction of spikes in each pass; so we will exactly recover the signal f in $O(\log m)$ passes.

In presence of the external noise $w \neq 0$, we can still recover a constant fraction of spikes in the first pass, however with error whose ℓ_1 norm is proportional to the ℓ_1 norm of the noise w. This error forms the "internal noise", which will add to the external noise in the next round. So, *the total noise doubles at every round*. After the $\log_a m$ rounds (needed to recover all spikes), the error of recovery will become polynomial in m. This is clearly unacceptable: Theorem 5 claims the error to be logarithmic in m.

This calls for a more delicate analysis of the error. Instead of adding the internal noise as a whole to the original noise, we will show that the internal noise spreads out over the subsets of the random partitions. So, most of the measurements will contain a small fraction of the internal noise, which will yield a small error of recovery in the current round. The major difficulty is to prove that this spreading phenomenon is *uniform* – one isolation matrix spreads the internal noise for all signals f at once, with high probability. This is a quite delicate problem. Indeed, in the last passes a constant number of spikes remain in the signal, and we have to find them correctly. So, the spreading phenomenon must hold for all but a constant number of measurements. Allowing so few exceptional measurements would naturally involve a very weak probability of such phenomenon to hold. On the other hand, in the last passes the internal noise is very big (having accumulated in all previous passes). Yet we need the spreading phenomenon to be uniform in all possible choices of the internal noise. It may seem that the weak probability estimates would not be sufficient to control a big internal noise in the last passes.

We will resolve this difficulty by doing a "surgery" on the internal noise, decomposing it in pieces corresonding to the previous passes, proving corresponding uniform probability estimates for each of these pieces, and uniting them in the end. This yields the Isolation Theorem 6, which summarizes the properties of the isolation matrix.

4.2. **Deterministic Part.** The proof of Theorem 5 is by induction on the pass k. We will normalize the signal so that $\|w\|_1 = 1/(400000a)$. We will actually prove a result stronger than Theorem 5. We will show that, in pass k, the signal has the form

$$f^{(k)} = s_k + w + \sum_{j=0}^{k-1} \nu_j \tag{4.1}$$

where s_k contains at most m_k spikes. The vector ν_j is the *internal noise* from pass j, which consists of $3m_j$ or fewer nonzero components with magnitudes at most $2m^{-j}$. When we have done with all passes, that is when $k = 1 + \log_a m$, we will have no more spikes in the signal ($m_k = 0$ thus $s_k = 0$). This at once implies Theorem 5.

The proof of (4.1) will only use the two *deterministic* properties of the isolation matrix. They are summarized in the following theorem, whose proof we will sketch with a series of propositions.

Theorem 6 (Properties of Isolation Matrix)**.** *The isolation matrix has the following quality with probability at least* $(1 - d^{-3})$*. If a signal has the form* (4.1) *for some pass* k*, then at least 99/100 of the trial submatrices have these two properties:*

1. *All but* $\frac{1}{100} m_{k+1}$ *spikes appear alone in a measurements, isolated from the other spikes.*
2. *Except for at most* $\frac{1}{100} m_{k+1}$ *of the measurements, the internal and external noise assigned to that measurement have* ℓ_1 *norm at most* $\frac{1}{1000} m_k^{-1}$*.*

The first proposition considers the performance of the algorithm in one of the 99/100 non-exceptional trials under an artificial assumption that will be removed in the second proposition

Proposition 7 (One Trial, No Inaccuracies)**.** *Suppose that the trial is not exceptional. Assume that each measurement contains at most one spike and that the external noise in each measurement is no greater than* $\varepsilon = \frac{1}{1000} m_k^{-1}$*. Then the trial constructs a list of at most* m_k *spikes.*

1. *If* $\left|f^{(k)}(i)\right| > 2\varepsilon$ *then the list contains a spike with position* i *and estimated value* $f^{(k)}(i) \pm \varepsilon$*.*
2. *If the list contains a spike with position* i *and* $\left|f^{(k)}(i)\right| \le 4\varepsilon$*, then the estimated value of the spike is no more than* 5ε *in magnitude.*

We call list items that satisfy these estimates accurate*.*

Proof. To prove this proposition, we outline a series of lemmas. We begin with a simple observation about the performance of the bit-tests.

Lemma 8. *Assume that a measurement contains a position* i *of value* p *and that the* ℓ_1 *norm of the other positions in the measurement is at most* ϵ*. Then*

1. *The estimated value* p_{est} *is bounded by the total measurement; that is,* $|p_{\text{est}}| \le |p| + \epsilon$*.*

(2) *If $|p| > 2\epsilon$, then the estimated position is i (i.e., the bit-test locates the position correctly) and the estimated value is within ϵ from p, $|p_{\text{est}} - p| \le \epsilon$.*

Proof. Follows immediately from the definitions of the bit-tests and the estimation procedures. □

Let us set $\epsilon = \frac{1}{1000} m_k^{-1}$ and observe that $\frac{1}{m_{k-1}} < \frac{1}{1000 m_k} = \frac{1}{\epsilon}$. Theorem 6 establishes two types of measurements which most satisfy. We refer to these two types as *good measurements* and define them precisely.

Definition 9. *A* good measurement *satisfies one of the two criteria:*

(1) *The measurement is empty; that is, it contains positions with values $\left|f^{(k)}(i)\right| \le \epsilon$ and the total ℓ_1 norm of the positions in the measurement is less than 1.5ϵ.*
(2) *The measurement contains one spike at position i with $\left|f^{(k)}(i)\right| > \epsilon$ and the ℓ_1 norm of all other positions in this measurement is less than 0.5ϵ.*

The next lemma shows that for good measurements, the bit-tests return reasonably accurate estimates.

Lemma 10. *Assume that a measurement is a good one. If the measurement is empty, then the estimated value of $f^{(k)}$ in that measurement is no more than 1.5ϵ. If the measurement contains one spike, then the estimated position is the position of the spike and its estimated value is within 0.5ϵ of the true value of the spike.*

Proof. Follows from the definitions of good measurements and Lemma 8. □

The next lemma follows from the previous argument and demonstrates that if the bit-tests identify a spike position, they do so reasonably accurately and precisely.

Lemma 11. *Assume that in a single measurement the bit-tests identify and estimate a spike at position i and that the estimated value p_{est} is greater than 1.5ϵ, $|p_{\text{est}}| > 1.5\epsilon$. Then the measurement contains a spike at position i and the true value of p is within 0.5ϵ of p_{est}.*

Strictly speaking, we perform multiple trials at each round k. We obtain, in a single trial, an estimate position and its estimated value, which we call the preliminary estimated value for that position. If, after performing all the trials, we have more than one preliminary estimated value for an estimated position, we simply use the preliminary value with the largest absolute value as the estimate assigned to this position. We then identify the m_k positions with the largest assigned estimates. A simple argument (which we omit here for brevity) demonstrates that the true value p of a spike at position i with $|p| > 2\epsilon$ is assigned an estimate p_{est} within 0.5ϵ of p. Furthermore, if the assigned estimate p_{est} of a spike at position i satisfies $|p_{\text{est}}| > 1.5\epsilon$, then the true value p is within 0.5ϵ of p_{est}. To simplify our arguments in what follows, we simply refer to the estimated values in one trial as the assigned estimate values.

With the above lemmas, we are able to complete the proof of the proposition. Our previous discussion shows that those positions i with $\left|f^{(k)}(i)\right| > \epsilon$ include the positions with estimated values larger than 1.5ϵ; i.e.,

$$\left\{ i \,\middle|\, |f^{(k)}(i)| > 1.5\epsilon \right\} \subset \left\{ i \,\middle|\, |f^{(k)}(i)| > \epsilon \right\}.$$

Our inductive hypothesis assumes that there are at most m_k positions in the right set above, so there are at most m_k positions in the left set as well. Our algorithm (for one trial) identifies all of these positions and reports estimated values that are within 0.5ϵ of the true values. Hence, our list of identified positions includes those i with $\left|f^{(k)}(i)\right| > 2\epsilon$. If a position i is identified and if $\left|f^{(k)}(i)\right| \le 4\epsilon$, the its estimated value is at most 5ϵ in magnitude by the previous lemmas. This proves the proposition. □

The second proposition removes the artificial assumption on the spikes and noise.

Proposition 12 (One Trial)**.** *Suppose that the trial is not an exceptional trial. In this trial, suppose each measurement is a good one, except for at most* $\frac{1}{50}m_{k+1}$*. Then the trial constructs a list of at most* m_k *spikes. All items in the list are accurate, except at most* $\frac{3}{50}m_{k+1}$*.*

Proof. We begin the proof with a lemma that shows the list produced by the algorithm is stable with respect to changes in a few measurements.

Lemma 13. *Assume that we perform one trial of the algorithm with two different signals and that their measurements (in this one trial) are identical except for* b *measurements. Then the estimated signals are identical except in* $2b$ *positions.*

Proof sketch. Prove this for $b = 2$ and then proceed by induction. □

Let us now consider the set of fewer than $\frac{1}{50}m_{k+1}$ measurements and set to zero the signal positions that fall into these measurements. This procedure creates two signals: the original signal and the restricted signal (with zeroed out positions). The restricted signal satisfies the conditions in Proposition 7 so all its measurements are good ones and agree with those of the original signal except for $\frac{1}{50}m_{k+1}$ measurements. The previous lemma guarantees that the estimated signals for the original and restricted signals are identical in all but $\frac{1}{25}m_{k+1}$ positions. By our inductive hypothesis, there are at most $\frac{1}{50}m_{k+1}$ positions of the original signal with value greater than 2ϵ in the exceptional measurements. Let us gather these $\frac{1}{25}m_{k+1}$ and $\frac{1}{50}m_{k+1}$ exceptional positions into one set of $\frac{3}{50}m_{k+1}$ exceptions. It is straightforward to show that the positions not in this exceptional set are good positions and if they are identified, they are identified accurately. □

We combine results from all trials. The algorithm considers positions identified in at least $\frac{9}{10}$ of the total trials T. It then takes the median (over all trials) to estimate the values of these positions.

Lemma 14 (Combining Trials)**.** *The number of list items that are inaccurate in more than 1/10 of the trials is at most* m_{k+1}*. The total number of positions that appear in 9/10 of the trials is at most* $\frac{10}{9}m_k$*.*

Proof sketch. Both parts follow by simple counting arguments. The list items that are inaccurate come from the accumulation of the $\frac{3}{50}m_{k+1}$ exceptional positions in Proposition 12. These become the spikes in the next pass. See the Appendix for a detailed proof. □

Now we are ready to prove the induction step. Recall that after round k, the new signal is the difference between the current signal and its estimate: $f^{(k)} = f^{(k-1)} - \tilde{f}^{(k-1)}$, with the convention that if a signal position is not considered and not changed by the algorithm, its estimated value is zero.

Lemma 15 (Induction Hypothesis)**.** *After pass* k*, there are at most* m_{k+1} *spikes remaining. The contribution* ν_k *to the internal noise contains at most* $3m_k$ *components with values at most* $2/m_k$*.*

Proof sketch. The list items that are inaccurate in more than 1/10 of the become spikes in the next pass. The list items that are identified in 9/10 of the trials consist of at least 8/10 accurate estimates, so medians over all values also yield accurate estimates. Since accurate estimates are either close to the actual spike value or close to zero, the value of the residual signal in these locations is less than $9\varepsilon < m_k^{-1}$. The difference between the spikes in the signal $f^{(k)}$ and the large entries in the update signal yields contains at most $m_{k+1} + m_k + \frac{10}{9}m_k < 3m_k$ terms of size less than $m_k^{-1} + m_{k+1}^{-1} < 2m_k^{-1}$. See the Appendix for a more detailed proof. □

This lemma proves the induction hypothesis.

The final lemma controls the recovery error and completes the proof of Theorem 6

Lemma 16 (Total Spikes and Recovery Error)**.** *Chaining Pursuit recovers at most* $O(m)$ *spikes. The total recovery error is at most* $(1 + C \log m)\,\|w\|_1$*.*

Proof Sketch. After pass $K = \log_a m$, there are no more spikes remaining since $m_k < 1$. At most $\frac{10}{9} m_k$ spikes are recovered in pass k. Since m_k decays geometrically, the total number of spikes is $O(m)$. The error after the last pass is the ℓ_1 norm of the signal $f^{(K+1)}$. This signal consists of the external noise, which has norm $\|w\|_1$, and the internal noise, which satisfies

$$\left\| \sum\nolimits_{j=0}^{K} \nu_j \right\|_1 \leq \sum\nolimits_{j=1}^{K} 6 m_j m_j^{-1} = 6 \log_a m.$$

This value is a constant times $\|w\|_1$. □

4.3. **Probabilistic Part.** Here we prove that the random isolation matrix $\boldsymbol{A}$ indeed has the two properties described in Theorem 6.

Lemma 17 (Isolations)**.** *Fix a round k. With probability at least $1 - \exp\{-4 m_k \log d\}$, the following is true. In pass k, at least 99/100 of the trial submatrices isolate all but $\frac{1}{100}$ of the m_k spikes.*

Proof sketch. By the inductive hypothesis, the signal has at most $n := m_k$ positions of value larger than $1/m_{k-1}$. We put these in $N := \overline{m}_k = C'(a) m / 2^k$ buckets. We can choose the function $C'(a)$ so that in the proceeding lemma we have $C'(a) = C(\frac{1}{100a})$. Let us now apply the following lemma, which states that all except $\frac{1}{100} m_{k+1}$ positions are isolated with probability $1 - \delta$, where $\delta = 2e^{-9 m_k}$.

Lemma 18 (Balls and Bins)**.** *Put n balls randomly and independently into $N > C(M) n$ buckets. Then, with probability $1 - 2e^{-9n}$, all except n/M balls are isolated in their buckets.*

Proof. One complication in the proof comes from the absence of independence among buckets. We would rather let buckets choose balls. However, the contents of different buckets is dependent because the total number of balls is limited. So we will replace the original model with n balls with an *independent* model. We provide a complete analysis of this model in the Appendix. □

Let I be the event that all but $\frac{1}{100} m_k$ of the spikes are isolated. Let us repeat the random assignment above independently T times (for T trials) and let δ_t be independent Bernoulli random variables, $\mathbb{E}\, \delta_t = \delta$. We see by Chernoff's inequality that

$$\begin{aligned} \mathbb{P}\{I \text{ fails in more than } \frac{1}{100} T \text{ trials}\} &\leq \mathbb{P}\Big\{ \sum_{t=1}^{T} \delta_t > \frac{1}{100} T \Big\} \leq (100 e \delta)^{\frac{1}{100} T} \\ &= \exp\big(- m_k \cdot \frac{1}{100} T \big) \leq \exp(-4 m_k \log d). \end{aligned}$$

□

Lemma 19 (External noise)**.** *In pass k, in every trial, the number of measurements where the ℓ_1 norm of the external noise exceeds $\frac{1}{2000} m_k^{-1}$ is at most $\frac{1}{200} m_{k+1}$.*

Proof sketch. This is an easy part of the argument. The $(1,1)$ operator norm of each matrix $\boldsymbol{A}_t^{(k)}$ equals one, so it does not inflate the norm of the external noise. We use Markov's inequality to bound the number of measurements with too much noise. □

Lemma 20 (Internal noise)**.** *Fix a round k. With probability at least $1 - \exp\{-4 m_k \log d\}$, the following is true. In pass k, in at least 99/100 trials, the number of measurements where the ℓ_1 norm of the internal noise exceeds $\frac{1}{2000} m_k^{-1}$ is at most $\frac{1}{200} m_{k+1}$.*

Proof sketch. The internal noise control is a delicate statement, which forms a main technical part of the argument. First we pass from the random partition model to a model where measurements are independent, using conditioning. Each piece ν_j of the internal noise exceeds $\lambda_j m_k$ in a single measurement with probability $\mathrm{e}^{-O(m_j T_k)}$. We apply Chernoff's bound to find the probability that more than $\varepsilon_j m_k T_k$ measurements are noisy during all T_k trials. Then Markov's inequality controls the number of trials where more than $\varepsilon_j m_k$ measurements are contaminated. The numbers λ_j and ε_j are chosen so that the quantitative pigeonhole principle yields a bound on the total number of measurements with internal noise greater than $\frac{1}{2000}m_k^{-1}$. We must refer the reader to the Appendix for details. □

5. Algorithmic Dimension Reduction

We have been able to establish the following dimension reduction theorem for sparse vectors.

Theorem 21. *Let X be the union of all m-sparse signals in $\mathbb{R}^d$ and endow $\mathbb{R}^d$ with the ℓ_1 norm. The linear map $\mathbf{\Phi} : \mathbb{R}^d \to \mathbb{R}^n$ in Theorem 2 satisfies*

$$A\|f - g\|_1 \le \|\mathbf{\Phi}(f) - \mathbf{\Phi}(g)\|_1 \le B\|f - g\|_1$$

for all f and g in X, where $1/A = C\log^2(m)\log^2(d)$ and $B = C\log(m)$ and $n = O(m\log^2 d)$.

Proof. The upper bound is equivalent to saying that the $\ell_1 \to \ell_1$ operator norm satisfies $\|\mathbf{\Phi}\|_{1\to1} \le B$. This norm is attained at an extreme point of the unit ball of ℓ_1^d, which thus have support 1. Then the upper bound follows at once from the definition of $\mathbf{\Phi}$. That is, any $0-1$ vector of support 1 vector gets mapped by $\mathbf{\Phi}$ to a $0-1$ vector of support bounded by the total number of bit-tests in all trials and passes, which is $\sum_{k=0}^{\log_a m} O(k\log d)\log_2 d \le B$.

The lower bound follows from Theorem 2. Let f_m and g_m be any d-dimensional signals f and g with support m. Let $f = f_m$ and $V = \mathbf{\Phi}g_m$. Then the reconstruction $\widehat{f}$ from V will be exact: $\widehat{f} = g_m$. Then Remark 3 implies $\|f_m - g_m\|_1 \le \mathrm{C}\log(m)\ \|\mathbf{\Phi}f_m - \mathbf{\Phi}g_m\|_1$, which completes the proof. □

We are interested not only in the distortion and dimension reduction properties of our embedding but also in the stability and robustness properties of the embedding. Our previous analysis guarantees that $\mathbf{\Phi}^{-1}\mathbf{\Phi}$ is the identity on X and that the inverse can be computed in sublinear time since Chaining Pursuit perfectly recovers m-sparse signals. Our previous analysis also shows that our dimension reduction is stable and robust. In other words, our embedding and the reconstruction algorithm can tolerate errors η in the data $x \in X$, as well as errors ν in the measurements:

Theorem 22. *The linear map $\mathbf{\Phi} : \mathbb{R}^d \to \mathbb{R}^n$ in Theorem 2 and the reconstruction map $\mathbf{\Psi} : \mathbb{R}^n \to \mathbb{R}^d$ given by the Chaining Pursuit algorithm satisfy the following for every $\eta \in \mathbb{R}^d$ and every $\nu \in \mathbb{R}^n$ and for all m-sparse signals x in $\mathbb{R}^d$:*

$$\|x - \mathbf{\Psi}(\mathbf{\Phi}(x+\eta)+\nu)\|_1 \le (1 + \mathrm{C}\log m)(\|\eta\|_1 + \|\nu\|_1).$$

Proof. This is just a reformulation of our observations in Remark 3 with $x = f_m$, $\eta = f - f_m$, $\nu = \mathbf{\Phi}f - V$. □

References

[BC03] B. Brinkman and M. Charikar. On the impossibility of dimension reduction in ℓ_1. In *Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science (2003)*, 2003.

[Bou85] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel J. Math.*, 52:46–52, 1985.

[CDD06] A. Cohen, W. Dahmen, and R. DeVore. Remarks on compressed sensing. Working draft, 2006.

[CM03] G. Cormode and S. Muthukrishnan. What's hot and what's not: Tracking most frequent items dynamically. In *Proc. ACM Principles of Database Systems*, pages 296–306, 2003.

[CM05] G. Cormode and S. Muthukrishnan. Towards an algorithmic theory of compressed sensing. Technical report, DIMACS, July 2005.

[CM06] G. Cormode and S. Muthukrishnan. Combinatorial algorithms for compressed sensing. In *Proc. 40th IEEE Conference on Information Sciences and Systems*, Princeton, Mar. 2006.

[CRT04] E. Candès, J. Romberg, and T. Tao. Exact signal reconstruction from highly incomplete frequency information. Submitted for publication, June 2004.

[CRTV05] E. J. Candès, M. Rudelson, T. Tao, and R. Vershynin. Error correction via linear programming. In *Proc. FOCS 2005*, Pittsburgh, Oct. 2005.

[CS02] M. Charikar and A. Sahai. Dimension reduction in the ℓ_1 norm. In *Proceedings of the 43rd Annual IEEE Conference on Foundations of Computer Science (2002)*. IEEE Press, 2002.

[CT04] E. J. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? Submitted for publication, revised April 2005, Oct. 2004.

[CT05] E. J. Candès and T. Tao. Decoding by linear programming. Available from `arXiv:math.MG/0502327`, Feb. 2005.

[Don04] D. L. Donoho. Compressed sensing. Unpublished manuscript, Oct. 2004.

[Don05] D. L. Donoho. Neighborly polytopes and sparse solution of underdetermined linear equations. Dept. of Statistics TR 2005-4, Stanford Univ., 2005.

[Don06] D. L. Donoho. Neighborly polytopes and sparse solution of underdetermined linear equations. *IEEE Trans. Inform. Theory*, 2006. To appear.

[DT05] D. L. Donoho and J. Tanner. Sparse nonnegative solution of underdetermined linear equations by linear programming. Dept. of Statistics TR 2005-6, Stanford Univ., Apr. 2005.

[DT06] D. L. Donoho and J. Tanner. Thresholds for the recovery of sparse solutions via l1 minimization. In *Proc. 40th IEEE Conference on Information Sciences and Systems*, Princeton, Mar. 2006.

[GGI$^+$02a] A. C. Gilbert, S. Guha, P. Indyk, Y. Kotidis, S. Muthukrishnan, and M. J. Strauss. Fast, small-space algorithms for approximate histogram maintenance. In *ACM Symposium on Theoretical Computer Science*, 2002.

[GGI$^+$02b] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. J. Strauss. Near-optimal sparse Fourier representations via sampling. In *ACM Symposium on Theoretical Computer Science*, 2002.

[GMS05] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss. Improved time bounds for near-optimal sparse Fourier representation via sampling. In *Proc. SPIE Wavelets XI*, San Diego, 2005.

[GSTV06] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin. Algorithmic dimension reduction in the ℓ_1 norm. 2006.

[IN05] P. Indyk and A. Naor. Nearest neighbor preserving embeddings. Submitted, 2005.

[JL84] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mapping into hilbert space. *Contemporary Mathematics*, (26):189–206, 1984.

[Kas77] B. Kashin. Sections of some finite dimensional sets and classes of smooth functions. *Izv. Acad. Nauk SSSR*, 41:334–351, 1977.

[MPTJ05] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subgaussian processes. *Comptes Rendus Acad. Sci.*, 340:885–888, 2005.

[NL04] A. Naor and J. R. Lee. Embedding the diamond graph in l_p and dimension reduction in l_1. *Geometric and Functional Analysis*, 14(4):745–747, 2004.

[Pis89] G. Pisier. *The volume of convex bodies and Banach space geometry.* Cambridge University Press, 1989.

[RV05] M. Rudelson and R. Veshynin. Geometric approach to error correcting codes and reconstruction of signals. Available from `arXiv:math.MG/0502299`, Feb. 2005.

[RV06] M. Rudelson and R. Veshynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Proc. 40th IEEE Conference on Information Sciences and Systems*, Mar. 2006.

[Tal05] M. Talagrand. *The Generic Chaining.* Springer, Berlin, 2005.

[TG05] J. A. Tropp and A. C. Gilbert. Signal recovery from partial information via Orthogonal Matching Pursuit. Submitted to *IEEE Trans. Inform. Theory*, April 2005.

6. Appendix

6.1. **Related Work.** Table 1 compares the major algorithmic contributions. Some additional comments on this table may help clarify the situation. A class-optimal error bound means that the algorithm outputs an approximation whose error is within a constant factor of the worst error for any signal in the class. An instance-optimal error bound means that, for any signal, the algorithm outputs an approximation whose error is within a constant factor of the optimal m-term error for that particular signal. A mixed-norm instance optimal error bound simply means that we measure the errors in the instance optimal bound with two different norms. In particular, an ℓ_2–ℓ_1 mixed-norm bound gives

$$\|f - \widehat{f}\|_2 \le \frac{C}{\sqrt{m}} \|f - f_m\|_1.$$

6.2. **Proofs: Probabilistic analysis of Isolation Theorem.** **Lemma 18** *Put n balls randomly and independently into $N > C(M)n$ buckets. Then, with probability $1 - 2e^{-9n}$, all except n/M balls are isolated in their buckets.*

Proof. One complication in the proof comes from the absence of independence among buckets. We would rather let buckets choose balls. However, the contents of different buckets is dependent because the total number of balls is limited. So we will replace the original model with n balls with an *independent* model. The independent model will be easier to handle by the standard large deviation technique; the independent model reduces the original model from it by conditioning on the number of balls.

The independent model is the following assignment. We divide each bucket into n sub-buckets, and let δ_{ki} be independent $0,1$ valued random variables with expectation $\mathbb{E}\,\delta_{ki} = 1/N$, for all buckets $k = 1, \dots, N$ and sub-buckets $i = 1, \dots, n$. The independent random variables $X_k = \sum_{i=1}^n \delta_{ki}$ will be called the number of balls in bucket k in the independent model. If we condition on the total number of such "balls", we obtain the distribution of the numbers of true balls X'_k in bucket k in the original model:

$$(X'_1, \dots, X'_N) \equiv \Big(X_1, \dots, X_N \;\Big|\; \sum_{k=1}^N X_k = n\Big).$$

To prove the Lemma, we have to show that the number of non-isolated balls is small. The number of non-isolated balls in bucket k is $Y'_k = X'_k \cdot 1_{\{X'_k > 1\}}$ so the conclusion of the Lemma is that

$$\mathbb{P}\Big\{\sum_{k=1}^N Y'_k > n/M\Big\} \le 2e^{-9n}. \tag{6.1}$$

We will now transfer this problem to the independent model. First, without loss of generality we can change the n balls in the lemma and in the original model to $0.9n$ balls. We do not change the independent model, so the number of non-isolated balls in the independent model is

Approach, References	Signal Class	Uniform	Error bd.	# Measurements	Storage	Decode time
OMP + Gauss [TG05]	m-sparse	No	No error	$m \log d$	$md \log d$	$m^2 d \log d$
Group testing [CM06]	m-sparse	No	No error	$m \log^2 d$	$\log d$	$m \log^2 d$
ℓ_1 min. + Gauss [Don06, DT05, DT06]	m-sparse	Yes	No error	No closed form	$\Omega(md)$	$\Omega(md)$ LP
Group testing [CM06]	m-sparse	Yes	No error	$m^2 \log^2 d$	$m \log(d/m)$	$m^2 \log^2 d$
Group testing [CM06]	weak ℓ_p	Yes	ℓ_2 class optimal	$m^{\frac{3-p}{1-p}} \log^2 d$	$m^{\frac{2-p}{1-p}} \log d$	$m^{\frac{4-2p}{1-p}} \log^3 d$
ℓ_1 min. + Gauss [CT04, CDD06]	Arbitrary	Yes	Mixed ℓ_2–ℓ_1 instance	$m \log(d/m)$	$d \log(d/m)$	$\Omega(md)$ LP
ℓ_1 min. + Fourier [CT04, RV06, CDD06]	Arbitrary	Yes	Mixed ℓ_2–ℓ_1 instance	$m \log^4 d$	$m \log^5 d$	$d \log d$ (empirical)
Chaining Pursuit [GSTV06]	Arbitrary	Yes	Mixed ℓ_1–weak ℓ_1 instance	$m \log^2 d$	$d \log^2 d$	$m \log^2 d$
Fourier sampling [GGI$^+$02b, GMS05]	Arbitrary	No	ℓ_2 instance optimal	$m \operatorname{polylog} d$	$m \operatorname{polylog} d$	$m \operatorname{polylog} d$
Group testing [CM06]	Arbitrary	No	ℓ_2 instance optimal	$m \log^{5/2} d$	$\log^2 d$	$m \log^{5/2} d$

Notes: $\Omega(md)$ LP denotes the time to solve a linear program with $\Omega(md)$ variables. We suppress big-O notation for legibility.

Table 1. Comparison of algorithmic results for compressed sensing

$Y_k = X_k \cdot 1_{\{X_k>1\}}$. We have to bound

$$\begin{aligned}
\mathbb{P}\Big\{\sum_{k=1}^N Y'_k > n/M\Big\} &= \mathbb{P}\Big\{\sum_{k=1}^N Y_k > n/M \Big| \sum_{k=1}^N X_k = 0.9n\Big\} \\
&\le \mathbb{P}\Big\{\sum_{k=1}^N Y_k > n/M \Big| \sum_{k=1}^N X_k \ge 0.9n\Big\} \\
&\le \frac{\mathbb{P}\Big\{\sum_{k=1}^N Y_k > n/M\Big\}}{\mathbb{P}\Big\{\sum_{k=1}^N X_k \ge 0.9n\Big\}}.
\end{aligned}$$

By Prokhorov-Bennett inequality,

$$\mathbb{P}\Big\{\sum_{k=1}^N X_k \ge 0.9n\Big\} = \mathbb{P}\Big\{\sum_{k=1}^N \sum_{i=1}^n \delta_{ki} \ge 0.9n\Big\} \ge 1/2.$$

Therefore, proving Equation (6.1) in the original model reduces to proving that

$$\mathbb{P}\Big\{\sum_{k=1}^N Y_k > n/M\Big\} \le e^{-9n} \tag{6.2}$$

in the independent model.

A standard way to prove deviation inequalities such as Equation (6.2) is through the moment generating function. By Markov's inequality and independence of the variables Y_k, we have

$$\begin{aligned}
\mathbb{P}\Big\{\sum_{k=1}^N Y_k > n/M\Big\} &= \mathbb{P}\Big\{e^{10M\sum_{k=1}^N Y_k} > e^{10n}\Big\} \le e^{-10n} \cdot \mathbb{E}\Big[e^{10M\sum_{k=1}^N Y_k}\Big] \\
&= e^{-10n} \cdot \Big(\mathbb{E}\Big[e^{10MY_1}\Big]\Big)^N
\end{aligned}$$

To complete the proof of Equation (6.2) it remains to show that for $Y = (\sum_{i=1}^n \delta_i) \cdot 1_{\{\sum_{i=1}^n \delta_i > 1\}}$, its moment generating function satisfies

$$\Big(\mathbb{E}\Big[e^{10MY}\Big]\Big)^N \le e^n \tag{6.3}$$

where δ_i are $0,1$ valued independent random variables with $\mathbb{E}\,\delta_i = 1/N$. To estimate the moment generating function $\mathbb{E}[e^{MY}]$ in Equation (6.3), it suffices to know the tail probability $\mathbb{P}\{Y > t\}$ for large t. For large t, we estimate this tail probability by removing the restriction onto non-isolated bucket in the definition of Y and applying Chernoff's inequality for independent random variables. The tail probability is, however, much smaller if we do restrict onto non-isolated buckets. We take this into account for small t by computing the expectation of Y (which is straightforward).

Let us start with the first moment of Y. We claim that

$$\mathbb{E}[Y] \le C\Big(\frac{n}{N}\Big)^2. \tag{6.4}$$

Compare this with the average number of balls without conditioning on being non-isolated, $\mathbb{E}[X_k] = \frac{n}{N}$. Indeed, by the linearity of expectation,

$$\begin{aligned}
\mathbb{E}[Y] &= n \cdot \mathbb{E}\left[\delta_1 \cdot 1_{\{\sum_{i=1}^n \delta_i > 1\}}\right] \\
&= n \cdot \mathbb{P}\left\{\delta_1 = 1 \text{ and there exists } i \in \{2,\dots,n\} :\ \delta_i = 1\right\} \\
&= \mathbb{P}\left\{\delta_1 = 1\right\} \cdot \left(1 - \mathbb{P}\left\{\forall i \in \{2,\dots,n\} : \delta_i = 0\right\}\right) \\
&= \frac{n}{N} \cdot \left(1 - \left(1 - \frac{1}{N}\right)^{n-1}\right) \le \frac{n}{N} \cdot \left(1 - e^{-n/N}\right) \le C\left(\frac{n}{N}\right)^2.
\end{aligned}$$

Next, by the Chernoff inequality, for $s > 2$, we have

$$\mathbb{P}\{Y > s\} \le \mathbb{P}\{\sum_{i=1}^n \delta_i > s\} \le (s\frac{N}{n})^{-s}. \tag{6.5}$$

Now we are ready to bound the moment generating function. Let $K = 10M$ and change variables $t = e^{Ks}$, so that

$$\begin{aligned}
\mathbb{E}\left[e^{KY}\right] &= \int_0^\infty \mathbb{P}\left\{e^{KY} > t\right\} dt = 1 + \int_1^\infty \mathbb{P}\left\{e^{KY} > t\right\} dt \\
&= 1 + K\int_0^\infty \mathbb{P}\left\{Y > s\right\} e^{Ks}\, ds.
\end{aligned}$$

We split the integral in two parts. We use Equation (6.4) to estimate the integral near zero as

$$\begin{aligned}
\int_0^3 \mathbb{P}\left\{Y > s\right\} e^{Ks}\, ds &\le e^{3K}\int_0^\infty \mathbb{P}\left\{Y > s\right\} ds \\
&= e^{3K}\, \mathbb{E}[Y] \le Ce^{3K}\left(\frac{n}{N}\right)^2,
\end{aligned}$$

and we use Equation (6.5) to estimate the integral near infinity as

$$\begin{aligned}
\int_3^\infty \mathbb{P}\left\{Y > s\right\} e^{Ks}\, ds &\le \int_3^\infty \left(\frac{N}{n}s\right)^{-s} e^{Ks}\, ds \le \int_3^\infty \left(\frac{N}{n}e^{-K}\right)^{-s} ds \\
&= \frac{1}{\ln(\frac{N}{n}e^{-K})}\left(\frac{N}{n}e^{-K}\right)^{-3} \le e^{2K}\left(\frac{n}{N}\right)^2.
\end{aligned}$$

Combining these, we conclude that

$$\mathbb{E}\left[e^{KY}\right] \le 1 + CKe^{3K}\left(\frac{n}{N}\right)^2 \le 1 + \frac{n}{10N}.$$

Hence we obtain Equation (6.3)

$$\left(\mathbb{E}\left[e^{KY}\right]\right)^N \le \left(1 + \frac{n}{10N}\right)^N \le e^{\frac{n}{5N}\cdot N} \le e^n.$$

This proves the lemma. □

Lemma 20 *With probability at least* $1 - \exp\{-4m_k \log d\}$*, the following is true. In pass* k*, in at least 99/100 trials, the number of measurements where the* ℓ_1 *norm of the internal noise exceeds* $\frac{1}{2000}m_k^{-1}$ *is at most* $\frac{1}{200}m_{k+1}$.

Proof. Let us recall the inductive hypothesis in Equation 4.1. This lemma is a statement is about the signal $f^{(k-1)}$ in the positions with values smaller than $1/m_{k-1}$. Lemma 19 gives us a proof for $k = 0$.

Let $k \geq 1$. We can absorb the first summand in the second part of Equation 4.1 into the first term; i.e. we can assume that our signal has the form

$$f^{(k)} = \sum_{j=0}^{k-1} \nu_j \tag{6.6}$$

where ν_j consists of $4m_j$ or fewer nonzero components with magnitudes at most $\frac{3}{m_j}$.

To prove this result, we introduce positive parameters λ_j, ϵ_j, $j = 0, \dots, k-1$, which satisfy

$$\sum_{j=0}^{k-1} \lambda_j \leq \frac{1}{C'} \tag{6.7}$$

and

$$\sum_{j=0}^{k-1} \epsilon_j \leq \frac{1}{C'a} \tag{6.8}$$

where C' is a positive absolute constant to be chosen later. Next, we will prove the following separate claim about the internal noise ν_j.

Claim 23. *Assume that a signal satisfies Equation* (6.6)*. Let* $j \in \{0, \dots, k-1\}$*. Then*

$$\sum_{t=1}^{T} \#\Big(\textit{measurements in trial } t \textit{ s.t. } \|\nu_j\|_1 > \frac{\lambda_j}{m_k}\Big) < \epsilon_j m_k T \tag{6.9}$$

with probability

$$1 - e^{-\gamma m_j T}, \tag{6.10}$$

where γ *is some positive number such that*

$$\gamma T \geq 10 \log d. \tag{6.11}$$

Proof. **Claim implies Lemma 20.** We will first show that this Claim implies Lemma 20. Assume the claim holds. By the definition of T, the exceptional probability is

$$e^{-\gamma m_j T} \leq e^{-10 m_j \log d} \leq \binom{d}{4m_j}^{-4},$$

while the number of choices of round j signal is $\binom{d}{4m_j}$. Hence, with probability $1 - \binom{d}{4m_j}^{-3}$, the inequality in Equation (6.9) holds uniformly for all choices of internal noise ν_j. Summing up these exceptional probabilities for all rounds $j = 0, \dots, k-1$, we conclude that:

> *with probability* $1 - \binom{d}{4m_j}^{-2}$*, the system of measurements is such that the inequality in Equation* (6.9) *holds uniformly for all choices of the signal satisfying Equation* (6.6)*.*

Fix a system of measurements which meets these requirements so that we can dismiss the probability issues. Let us also fix a measurement v and a trial t. We refer to $\nu^t_{k,v}$ as the signal in measurement v and trial t. By Equation (6.7), for a fixed trial and a fixed measurement v, we have the containment of events:

$$\Big\{\|\nu^t_{k,v}\|_1 > \frac{1}{2000 m_k}\Big\} \subseteq \bigcup_{j=0}^{k-1} \Big\{\|\nu^t_{j,v}\|_1 > \frac{\lambda_j}{m_k}\Big\}.$$

Counting the measurements that satisfy each side of this containment, then summing over the trials, we obtain:

$$\begin{aligned}
\sum_{t=1}^{T} \#\Big\{\|\nu_{k,v}^t\|_1 > \frac{1}{2000 m_k}\Big\} &\le \sum_{t=1}^{T}\sum_{j=0}^{k-1} \#\Big\{\|\nu_{j,v}^t\|_1 > \frac{\lambda_j}{m_k}\Big\} \\
&\le \sum_{j=0}^{k-1} \epsilon_j m_k T \quad \text{by (6.9)} \\
&\le \frac{1}{C'a} m_k T \quad \text{by (6.8)} \\
&= \frac{1}{C'} m_{k+1} T.
\end{aligned}$$

By Markov's inequality, this implies (provided C' is chosen large enough) that in at most $\frac{1}{100}T$ trials t, the number of measurements where the ℓ_1 norm of the internal noise exceeds $\frac{1}{2000 m_k}$ is greater than $\frac{1}{200} m_{k+1}$. This implies the conclusion of Lemma 20

Proof of the claim. Now we prove the claim itself. This is a purely probabilistic problem. We call the nonzero positions of ν_j "balls" and we informally refer to the measurements as "buckets". By the definition of ν_j, the 1-norm of ν_j in a measurement (or bucket) v is large, $\|\nu_{j,v}\|_1 > \frac{\lambda_j}{m_k}$ if and only if the measurement contains at least $\frac{1}{3}\lambda_j \frac{m_j}{m_k}$ balls. Then Equation (6.9) is equivalent to

$$\sum_{t=1}^{T} \#(\text{measurement in trial } t \text{ which contain} > \frac{1}{3}\lambda_j \frac{m_j}{m_k} \text{ balls}) \le \epsilon_j m_k T. \tag{6.12}$$

To prove this with required probability (6.10), we will transfer the problem to an independent model—similar to the proof of Lemma 18.

Recall that the *original model* with n balls and T trials, in which we want to prove Equation (6.12), is to put $n = 4m_j$ balls into $N = \overline{m}_k$ buckets (or measurements) independently, and repeat this T times (trials) independently.

We want to replace this by the following *independent model*, where the contents of buckets are independent. There are N buckets in each of T trials. Divide each bucket into $S = nT$ sub-buckets. Let δ_{tli} be independent $0, 1$ valued random variables with expectation $\mathbb{E}\,\delta_{tli} = 1/NT$, for all trials $t = 1, \dots, T$, buckets $l = 1, \dots, N$ and sub-buckets $i = 1, \dots, S$. The independent random variables

$$X_{tl} = \sum_{i=1}^{S} \delta_{tli} \tag{6.13}$$

will be called the number of balls in bucket l, trial t, in the independent model. Note that

$$\mathbb{E}\, X_{tl} = \frac{S}{NT} = \frac{n}{N}.$$

Thus the average total number of balls in buckets in one trial is n. Let E_{tot} be the event that in each of at least $T/2$ trials, the total number of balls in buckets is at least $n/2$. It is then easy to deduce by Chernoff and Prokhorov-Bennett's inequalities that with probability at least $1/2$ (actually $1 - e^{-nT}$), E_{tot} holds; that is, $\mathbb{P}\{E_{\text{tot}}\} \ge 1/2$. We have to bound above the probability of the event

$$E := \{ \text{ Equation (6.12) does not hold}\}$$

in the original model. We reduce it to the independent model as follows:

$$\begin{aligned}\mathbb{P}\{\text{E in independent model}\} &\geq \mathbb{P}\{\text{E in independent model} \mid E_{\text{tot}}\} \cdot \mathbb{P}\{E_{\text{tot}}\} \\ &\geq \frac{1}{2}\, \mathbb{P}\{\text{E in independent model} \mid E_{\text{tot}}\}\end{aligned}$$

The probability can only decrease when we consecutively do the following changes:

(1) remove both occurences of "at least" in E_{tot}, resulting in exactly $T/2$ trials and exactly $n/2$ balls;
(2) restrict the sum in Equation (6.12) to the $T/2$ trials included in the new (exact) version of E_{tot}; and
(3) fix the set of $T/2$ trials in E_{tot}—say, require that these be the first $T/2$ trials.

After doing this, the law becomes the original model with $n/2$ balls and T trials. Hence,

$$\mathbb{P}\{\text{E in independent model}\} \geq \frac{1}{2}\, \mathbb{P}\{\text{E in original model with } n/2 \text{ balls and } T/2 \text{ trials}\}.$$

Therefore, it suffices to prove that Equation (6.12) holds in the independent model, with probability as in (6.10), i.e. with probability $1 - \frac{1}{2}e^{-\gamma m_j T}$, where γ is as in Equation (6.11).

In order to prove Equation (6.12) with the requisite probability, we first estimate the number of balls X_{tl} in one bucket, see Equation (6.13). It is a sum of $S = nT$ independent Bernoulli random variables with expectations $\frac{1}{NT}$. Then by the Chernoff inequality,

$$\mathbb{P}\Big\{X_{tl} > \frac{1}{3}\lambda_j \frac{m_j}{m_k}\Big\} \leq \Big(\frac{1}{12e}\lambda_j \frac{\overline{m}_k}{m_k}\Big)^{-\frac{1}{3}\lambda_j \frac{m_j}{m_k}}. \tag{6.14}$$

Let us call this probability η. We have to estimate the sum in Equation (6.12) which equals $\sum_{t=1}^{T}\sum_{l=1}^{N}\delta_{tl}$ where

$$\delta_{tl} = 1_{\{X_{tl} > \frac{1}{3}\lambda_j \frac{m_j}{m_k}\}}$$

are independent Bernoulli random variables whose expectations are $\mathbb{E}\,\delta_{tl} \leq \eta$ by Equation (6.14). Then by the Chernoff inequality, the probability that Equation (6.12) fails to hold is

$$\mathbb{P}\Big\{\sum_{t=1}^{T}\sum_{l=1}^{N}\delta_{tl} > \epsilon_j m_k T\Big\} \leq (\beta/e)^{-\epsilon_j m_k T},$$

where

$$\beta = \frac{\epsilon_j m_k}{\eta N} = \epsilon_j \cdot \frac{m_k}{\overline{m}_k} \cdot \frac{1}{\eta} = \epsilon_j \cdot \frac{m_k}{\overline{m}_k} \cdot \Big(\frac{1}{12e}\lambda_j \frac{\overline{m}_k}{m_k}\Big)^{\frac{1}{3}\lambda_j \frac{m_j}{m_k}} \tag{6.15}$$

To complete the proof, we need to show that

$$(\beta/e)^{-\epsilon_j m_k T} \leq \frac{1}{2}e^{-\gamma m_j T},$$

which would follow from

$$(\beta/e)^{\epsilon_j \cdot \frac{m_k}{m_j}} \geq e^{2\gamma}, \tag{6.16}$$

where γ must satisfy Equation (6.11).

Now we specify our choice of λ_j and ϵ_j. Set

$$\begin{aligned}\lambda_j &:= \max\Big\{12e\Big(\frac{m_k}{\overline{m}_k}\Big)^{1/2},\ 24\frac{m_k}{m_j},\ \frac{1}{C'k}\Big\}, \\ \epsilon_j &:= \max\Big\{\frac{m_k}{\overline{m}_k},\ \frac{1}{C'ak}\Big\}.\end{aligned}$$

Then Equations (6.7) and (6.8) are clearly satisfied (recall that $j < k$). Next, the base in Equation (6.15) is estimated as

$$\frac{1}{12e}\lambda_j\frac{\overline{m}_k}{m_k} \geq \left(\frac{\overline{m}_k}{m_k}\right)^{1/2}$$

and the exponent can be estimated using

$$\lambda_j\frac{m_j}{m_k} \geq 24.$$

Also, the first factor of Equation (6.15) is estimated as

$$\epsilon_j \cdot \frac{m_k}{\overline{m}_k} \geq \left(\frac{m_k}{\overline{m}_k}\right)^2.$$

Combining these three estimates, we obtain

$$\beta \geq \left(\frac{m_k}{\overline{m}_k}\right)^2 \cdot \left(\frac{\overline{m}_k}{m_k}\right)^{\frac{1}{6}\lambda_j\frac{m_j}{m_k}} \geq \left(\frac{\overline{m}_k}{m_k}\right)^{\frac{1}{8}\lambda_j\frac{m_j}{m_k}}.$$

Now we can check Equation (6.16).

$$\begin{aligned}
(\beta/e)^{\epsilon_j\cdot\frac{m_k}{m_j}} &\geq \left(\frac{\overline{m}_k}{em_k}\right)^{\frac{1}{8}\lambda_j\epsilon_j} \\
&\geq (a/2)^{\frac{\lambda_j\epsilon_j}{8}\cdot k} \quad \text{by the definition of } m_k,\ \overline{m}_k \\
&\geq \exp\left(\frac{\ln(a/2)}{(C')^2 a}\cdot\frac{k}{k^2}\right) \quad \text{by the definition of } \lambda_j \text{ and } \epsilon_j \\
&= e^{c(a)/k}.
\end{aligned}$$

Therefore Equation (6.16) holds for $\gamma = c(a)/k$, and this choice of γ satisfies the required condition in Equation (6.11), since $T = C(a)(k+1)\log(d)$. □

This completes the proof of the lemma. □

6.3. **Proofs: Deterministic Portion of Isolation Theorem. Lemma 14** *The number of list items that are inaccurate in more than 1/10 of the trials is at most m_{k+1}. The total number of positions that appear in 9/10 of the trials is at most $\frac{10}{9}m_k$.*

Proof. We prove the first part of the lemma with a simple counting argument. Let T denote the total number of trials. We have to bound b where

$$b = \#\{\text{positions bad in} \geq \tfrac{T}{10} \text{ trials}\} \leq \#\{\text{positions bad in} \geq \tfrac{T}{11} \text{ good trials}\}.$$

Let

$$\Sigma = \sum_{j\in \text{ good trials}} \#\{\text{positions bad in trial } j\}.$$

We have $\Sigma \geq \frac{bT}{11}$. Let T' be the number of good trials. Then Proposition 12 tells us that $\Sigma \leq \frac{3}{50}m_{k+1}T' \leq \frac{3\cdot T}{50}m_{k+1}$. Therefore, $\frac{bT}{11} \leq \frac{3T}{50}m_{k+1}$ and hence $b \leq \frac{33}{50}m_{k+1}$.

To prove the second part, recall that the algorithm updates a position if and only if the position is identified in at least 9/10 of the trials. Let ℓ denote the number of such positions. In every trial, m_k positions are identified. Hence,

$$m_kT = \sum_{t=1}^{T}\#\{\text{positions identified in trial } t\} \geq \frac{9}{10}T\ell$$

and thus $\ell \leq \frac{10}{9}m_k$. □

Lemma 15 *After pass k, there are at most m_{k+1} spikes remaining. The contribution ν_k to the internal noise contains at most $3m_k$ components with values at most $2/m_k$.*

Proof. Recall that $4.5\epsilon < m_k^{-1}$. It suffices to prove that for the non-exceptional positions i that satisfy the conclusions of the previous lemma, the value $\left|f^{(k)}(i)\right| \le 4.5\epsilon < \frac{2}{m_k}$. Let us fix such a position i which is good in at least $9/10T$ trials and show that

$$\left|f^{(k-1)}(i) - \tilde{f}^{(k-1)}(i)\right| \le 4.5\epsilon.$$

If $\left|f^{(k-1)}(i)\right| > 2\epsilon$, then the goodness of i in $9/10$ trials implies that i is identified in these trials and hence i is considered by the Algorithm. In each of these $9/10$ trials, the goodness of i also means that the assigned estimated value of i is within 0.5ϵ from its true value $f^{(k-1)}(i)$. Since $\tilde{f}^{(k-1)}(i)$ is the median of the assigned estimated values of i in each trial, it follows that it is also within 0.5ϵ from the true value $f^{(k-1)}(i)$.

Suppose that $\left|f^{(k-1)}(i)\right| \le 2\epsilon$. If i is not considered by the algorithm, then the value of the signal at this position is not changed, so $|f^{(k)}(i)| = |f^{(k-1)}(i)| \le 2\epsilon$. We can assume that i is considered by the algorithm. There are $(9/10)T$ trials in which i is identified and there are $(9/10)T$ trials in which i is good. Hence, there are at least $8/10T$ trials in which i is both good and identified. By the definition of goodness, this means that in each of these trials, the assigned estimated value $|\tilde{f}^{(k-1)}(i)| \le 2.5\epsilon$. Since the estimated value is the median of the assigned estimated values of i in each trial, it follows that $|\tilde{f}^{(k-1)}(i)| \le 2.5\epsilon$. Then

$$|f^{(k-1)}(i) - \tilde{f}^{(k-1)}(i)| \le |f^{(k-1)}(i)| + |\tilde{f}^{(k-1)}(i)| \le 2\epsilon + 2.5\epsilon = 4.5\epsilon.$$

This completes the proof of the first part of the lemma.

The first part of this lemma shows formally that the difference between the spikes in the signal $f^{(k)}$ and the large entries in the update signal (i.e., those with absolute values greater than m_k^{-1}) contains at most m_{k+1} terms. By the inductive step, the same holds for the previous rounds. In addition, Lemma 14 tells us that the algorithm updates at most $1.1m_k$ positions in the signal. By the triangle inequality it follows that the difference contains at most $m_{k+1} + m_k + 1.1m_k \le 3m_k$ terms. The maximal absolute value of the difference signal is

$$\frac{1}{m_k} + \frac{1}{m_{k-1}} \le \frac{2}{m_k}.$$

□