

Lecture 8

Prof. Nick Harvey

University of British Columbia

1 Compressed Sensing

Let $x \in \mathbb{R}^n$ be some unknown vector. We can learn information about x through the following “measurement” operation: we can choose a vector $a \in \mathbb{R}^n$ and we are told the value of $a^T x$. We would like to minimize the number of measurements needed to learn the vector x . Note that performing m measurements is equivalent to choosing an $m \times n$ matrix A then asking for the vector of measurements Ax .

Without any further assumptions, n operations are clearly necessary and sufficient. This is because (assuming that the query vectors a are linearly independent) after i operations, x could be any vector in the $n - i$ dimensional subspace of points consistent with the i measurements.

Now suppose that we know some additional properties of x . We say that a vector is **k -sparse** if it has only k non-zero components. Suppose that x is *almost k -sparse*, meaning that $x = y + z$ where y is k -sparse and z is a small “noise” vector. Can we perform $m < n$ queries and compute a good approximation to x ? Intuitively, we want to learn y exactly and we don’t care about learning z .

We present a solution which involves two interesting ideas. The first idea is to perform *random measurements*, much like Johnson-Lindenstrauss measures the projection of vectors on a random subspace. We will set $m = \Theta(k \log n)$, let A be a matrix of independent Gaussians, then ask for the vector of measurements Ax .

The second idea is to use “small L_1 -norm” as a proxy for sparsity. After obtaining the measurements Ax , we will find the vector x^* with smallest L_1 -norm that is consistent with those measurements. In other words, we solve the optimization problem

$$\min\{ \|x^*\|_1 : Ax^* = Ax \}, \quad (1)$$

where $\|x^*\|_1 = \sum_{i=1}^n |x_i^*|$ is the L_1 -norm. This optimization problem can be reformulated as a linear program, although the details of that transformation are not crucial for us.

Our main theorem is:

Theorem 1 Define $\text{Err}_1^k(x) = \min_{x'} \|x' - x\|_1$, where the minimum is taken over all k -sparse vectors. With probability at least $1 - 2/n$ over the choice of A ,

$$\|x^* - x\|_2 = O(\text{Err}_1^k(x)/\sqrt{k}), \quad (2)$$

To interpret the error guarantee in (2), note that if $x = y + z$ as above, then setting $x' = y$ we get $\|x^* - x\|_2 = O(\|z\|_1/\sqrt{k})$. So the approximation error of x^* only depends on the magnitude of the noise.

In fact, we will prove a stronger result than Theorem 1 states. We will prove that with probability at least $1 - 2/n$, we obtain a matrix A such that (2) holds for every vector $x \in \mathbb{R}^n$.

To prove such a result, we need to define a condition on the matrix A that will ensure that (2) holds for all x . We say that A satisfies the **(k, δ) -Restricted Isometry Property (RIP)** if, for any k -sparse

vector x we have

$$(1 - \delta)\|x\|_2 \leq \|Ax\|_2 \leq (1 + \delta)\|x\|_2. \quad (3)$$

This condition is equivalent to requiring that, for every submatrix of A consisting of k columns, all singular values lie in the interval $[1 - \delta, 1 + \delta]$. (An $m \times k$ matrix with for which all singular values are exactly equal to 1 is an orthogonal matrix if $k = m$ and if $k \leq m$ it is called an isometry.)

We will prove two theorems, which directly imply our main theorem.

Theorem 2 *If $m = \Theta(k \log n)$ and A is an $m \times n$ matrix whose entries are independent normal random variables with distribution $N(0, 1/k)$, then A is $(k, 1/3)$ -RIP with probability at least $1 - 2/n$.*

Theorem 3 *If A is $(25k, 1/3)$ -RIP then (2) holds.*

The second theorem does not involve any new concepts or any randomization; it just performs some clever algebraic manipulations. So we will not discuss it in class, but we include the proof in Section 3 for completeness.

2 Proof of Theorem 2

First of all, let us derive a very crude bound on $\|Ax\|$. If the entries of A all have absolute value at most n then $\|Ax\| \leq n^2\|x\|$ for all x . By standard bounds on the Gaussian tail, the probability that this fails to hold is less than $\sqrt{k} \int_n^\infty e^{-kt^2/2} dt < 1/n$. (Alternatively, we could construct A with binomial random variables instead of Gaussians, in which case our crude bound certainly holds.)

To prove the theorem we must show that (3) holds for every k -sparse vector x . The argument proceeds by a union bound over all sets $T \subset \{1, \dots, n\}$ of indices with $|T| = k$, then considering all vectors x whose non-zeros all lie in the coordinates in T .

So restricting to the coordinates in T , our desired statement is: if B is an $m \times k$ matrix of independent Gaussians, then

$$\frac{2}{3}\|y\| \leq \|By\| \leq \frac{4}{3}\|y\| \quad \forall y \in \mathbb{R}^k. \quad (4)$$

This is almost identical to the statement of the Johnson-Lindenstrauss lemma, except previously we only wanted those inequalities to hold for a fixed vector, or a small collection of vectors. Now we want it to hold for *all* vectors.

Previously we analyzed the probability of (4) holding for a single vector, then extended that to a small collection of vectors by taking a union bound. Unfortunately we now want (4) to hold for infinitely many vectors, and we don't want to union bound over an infinitely large set!

The solution is to observe that the function $y \mapsto \|By\|/\|y\|$ is a continuous function in y , so if $2/3 \leq \|By\|/\|y\| \leq 4/3$ holds for some y , then it approximately holds for all nearby y . Also, since $\|By\|/\|y\|$ only depends on the direction of y , not its norm, it suffices to consider vectors with $\|y\| = 1$.

So define the sphere $S = \{x \in \mathbb{R}^k : \|x\| = 1\}$. Since we only need to consider points $y \in S$, it seems conceivable that we could find *finitely* many points $P = \{p_1, \dots, p_\ell\}$ such that *every* point in S is “nearby” to some p_i . Such a set P is called an [epsilon-net](#), and they are well-studied objects. The following fact is known.

Fact 4 *There exists a set $P \subset S$ of size $O(1/\epsilon)^k$ that is an ϵ -net of S , meaning that $\min_i \|p_i - x\| \leq \epsilon$ for all $x \in S$.*

We set $\epsilon = \frac{1}{6n^2}$ and let P be an ϵ -net of size $O(1/\epsilon)^k = \exp(O(k \ln n))$. Then we apply the Johnson-Lindenstrauss lemma with error parameter $1/6$ to the points in P . By a union bound, with probability at least $1 - |P| \cdot e^{-\Theta(m)}$ we have

$$\frac{5}{6}\|p\| \leq \|Bp\| \leq \frac{7}{6}\|p\| \quad \forall p \in P.$$

Now consider any point y with $\|y\| = 1$. Because P is an ϵ -net, there is some point $p \in P$ of distance at most ϵ to y . That is, letting $z = y - p$, we have $\|z\| \leq \epsilon$. Now using the crude bound from above, we have $\|Bz\| \leq n^2\epsilon = 1/6$. So, assuming that (4) holds for points in P we get

$$\|By\| \leq \|Bp\| + \|Bz\| \leq \frac{7}{6}\|p\| + \frac{1}{6} = \frac{4}{3}.$$

Similarly,

$$\|By\| \geq \|Bp\| - \|Bz\| \geq \frac{5}{6}\|p\| - \frac{1}{6} = \frac{2}{3}.$$

So far we have only considered a single set of indices T and shown that failure probability is at most $|P| \cdot e^{-\Theta(m)} = \exp(O(k \ln n) - \Theta(m))$. We analyze the total failure probability by a union bound over all choices of T . The number of sets T of size k is $\binom{n}{k} \leq n^k = \exp(k \ln n)$. So the probability that (4) fails to hold for any point in any of our ϵ -nets is at most

$$\binom{n}{k} \cdot |P| \cdot e^{-\Theta(m)} = \exp(O(k \log n) - \Theta(m)),$$

which is less than $1/n$ if $m = ck \ln n$ for a sufficiently large constant c .

3 Appendix: Proof of Theorem 3

Let $h = x^* - x$. Our goal is to bound $\|h\|_2$. To start off, let us order the coordinates such that

- $|x_1|, \dots, |x_k|$ are all at least $|x_{k+1}|, \dots, |x_n|$.
- $|h_{k+1}| \geq \dots \geq |h_n|$.

Now define the sets of indices

$$\begin{aligned} T_0 &= \{1, \dots, k\} \\ T_1 &= \{k+1, \dots, 26k\} \\ T_2 &= \{26k+1, \dots, 51k\}, \\ \dots &= \dots \end{aligned}$$

so $|T_0| = k$ and $|T_i| = 25k$ for $i \geq 1$. Also, define

$$T_{01} = T_0 \cup T_1 \quad \overline{T_0} = \bigcup_{i \geq 1} T_i \quad \overline{T_{01}} = \bigcup_{i \geq 2} T_i.$$

We will use the notation x_S to denote the restriction of x to the coordinates in S . For example x_{T_0} is the vector (x_1, \dots, x_k) . Finally, define

$$\epsilon = \|x_{\overline{T_0}}\|_1 / \sqrt{k} = \text{Err}_1^k(x) / \sqrt{k},$$

since the smallest L_1 error of any k -sparse approximation is obtained letting x' be the k -sparse vector that contains x 's largest k coordinates (in absolute value).

The proof involves three claims.

Claim 5 $\|h_{\overline{T_0}}\|_1 \leq \|h_{T_0}\|_1 + O(\sqrt{k}\epsilon)$.

Claim 6 $\|h_{\overline{T_{01}}}\|_2 \leq \|h_{T_0}\|_2 + O(\epsilon)$.

Claim 7 $\|h_{T_{01}}\|_2 \leq O(\epsilon)$.

These claims easily imply Theorem 3, which states that $\|h\|_2 = O(\text{Err}_1^k(x)/\sqrt{k})$. We have

$$\|h\|_2 \leq \|h_{T_{01}}\|_2 + \|h_{\overline{T_{01}}}\|_2 \leq 2 \cdot \|h_{T_{01}}\|_2 + O(\epsilon) \leq O(\epsilon) = O(\text{Err}_1^k(x)/\sqrt{k}),$$

by Claims 6 and 7.

PROOF:(of Claim 5) The main facts we need are that the L_1 -norm satisfies the triangle inequality and it is additive on any partition of the coordinates, i.e., $\|x\|_1 = \|x_S\|_1 + \|x_{\overline{S}}\|_1$ for any set of indices S . Since x^* is an optimal solution to (1), we have

$$\begin{aligned} \|x\|_1 &\geq \|x^*\|_1 \\ &= \|x - h\|_1 \\ &= \|x_{T_0} - h_{T_0}\|_1 + \|x_{\overline{T_0}} - h_{\overline{T_0}}\|_1 \\ &\geq \|x_{T_0}\|_1 - \|h_{T_0}\|_1 - \|x_{\overline{T_0}}\|_1 + \|h_{\overline{T_0}}\|_1, \end{aligned}$$

by the triangle inequality. Rearranging, we get

$$\|h_{\overline{T_0}}\|_1 \leq \|x\|_1 - \|x_{T_0}\|_1 + \|x_{\overline{T_0}}\|_1 + \|h_{T_0}\|_1 = 2\|x_{\overline{T_0}}\|_1 + \|h_{T_0}\|_1,$$

which proves the claim. \square

PROOF:(of Claim 6) By our ordering of the coordinates, the j th coordinate of $h_{\overline{T_0}}$ has absolute value $|(h_{\overline{T_0}})_j| \leq \|h_{\overline{T_0}}\|_1/j$. So

$$\begin{aligned} \|h_{\overline{T_{01}}}\|_2^2 &= \sum_{j=25k+1}^n (h_{\overline{T_0}})_j^2 \\ &\leq \sum_{j=25k+1}^n (h_{\overline{T_0}})_j^2/j^2 \\ &\leq \|h_{\overline{T_0}}\|_1^2 \cdot \sum_{j=25k+1}^n 1/j^2 \\ &\leq \|h_{\overline{T_0}}\|_1^2/k, \end{aligned}$$

where we have bounded the sum by $\int_{25k}^{\infty} j^{-2} dj < 1/k$. Thus

$$\|h_{\overline{T_{01}}}\|_2 \leq \|h_{\overline{T_0}}\|_1/\sqrt{k} \leq \|h_{T_0}\|_1/\sqrt{k} + O(\epsilon) \leq \|h_{T_0}\|_2 + O(\epsilon).$$

Here the second inequality is by Claim 5 and the third is because $\|x\|_1 \leq \sqrt{k}\|x\|_2$ for any k -dimensional vector x . \square

PROOF:(of Claim 7) Our optimization problem ensures that $Ax = Ax^*$, so $Ah = 0$. So,

$$0 = \|Ah\|_2 \geq \|Ah_{T_{01}}\|_2 - \sum_{i \geq 2} \|Ah_{T_i}\|_2 \geq \frac{1}{3}\|h_{T_{01}}\|_2 - \frac{4}{3} \sum_{i \geq 2} \|h_{T_i}\|_2, \quad (5)$$

where the first inequality is by the triangle inequality, and the second is by the RIP property.

Next we derive an upper bound on $\sum_{i \geq 2} \|Ah_{T_i}\|_2$. By our ordering of the coordinates, every coordinate in $h_{T_{j+1}}$ is smaller (in absolute value) than the minimum coordinate in h_{T_j} . This means it is also smaller than the average coordinate, which is $\|h_{T_j}\|_1/25k$. So $\|h_{T_{j+1}}\|_2^2 \leq 25k \cdot (\|h_{T_j}\|_1/25k)^2$, which implies that

$$\|h_{T_{j+1}}\|_2 \leq \|h_{T_j}\|_1/5\sqrt{k} \quad \forall j \geq 2.$$

Summing over all j ,

$$\begin{aligned} \sum_{j \geq 2} \|h_{T_j}\|_2 &\leq \sum_{j \geq 1} \|h_{T_j}\|_1/5\sqrt{k} \\ &= \|h_{\overline{T_0}}\|_1/5\sqrt{k} \\ &\leq \|h_{T_0}\|_1/5\sqrt{k} + O(\epsilon) \\ &\leq \|h_{T_0}\|_2/5 + O(\epsilon) \\ &\leq \|h_{T_{01}}\|_2/5 + O(\epsilon). \end{aligned}$$

Here the second inequality is by Claim 5. Combining this with (5), we get

$$\|h_{T_{01}}\|_2 \leq \frac{12}{3} \left(\|h_{T_{01}}\|_2/5 + O(\epsilon) \right) = \frac{4}{5} \|h_{T_{01}}\|_2 + O(\epsilon),$$

which implies the claim. \square