Khoury College of Computer Sciences | Fall 2023
Northeastern University | 20 & 24 October 2023
CS 7870: Seminar in Theoretical Computer Science | Scribe: Dustin Lin

Outline:

- Compressed sensing
- Discrete sparse recovery
- Proof of compressed sensing
- Generating a random A: Johnson-Lindenstrauss implies RIP

These scribe notes cover the general topic of compressed sensing, presented over two lectures. The material is partly based on the lecture notes of Christopher Musco [Mus18], Tim Roughgarden and Gregory Valiant [RV17], Sepehr Assadi [Asa20], and the paper relating the Johnson-Lindenstrauss transform with compressed sensing [BDDW06]. The original references for compressed sensing are [CRT06, Don06].

1 Compressed sensing

The **compressed sensing** problem is the following: consider some unknown vector $z \in \mathbb{R}^n$ that we are trying to find. We can take m linear measurements on z to help us figure out what z is. We can think of these m linear measurements together as a matrix $A \in \mathbb{R}^{m\times n}$. And therefore our problem is the following:

Given some $z \in \mathbb{R}^n$, design an $A \in \mathbb{R}^{m\times n}$ such that $Ax = b$, and we can recover z via b.

Ideally we want $m \ll n$, but in the general case it is requires to have n equations to solve for n unknowns. To get out of this conundrum, we need to assume some structure of z.

Definition 1. A vector z is *k-sparse* if z has non-zero values in at most k entries.

1.1 Discrete Sparse Recovery

As an example, we begin with a discrete setting where $z \in \mathbb{F}_2^n$, and take our linear-measurements as a series of dot-products. How can we go about designing our measurements to recover z? For intuition, let us look at the case where $k = 1$ and $n = 8$.

Given a 1-sparse vector $z \in \mathbb{F}_2^8$ our task boils down to finding which index the singular 1 entry is in. We can design our A to do a binary-like search to iteratively deduce where the entry is. For this example, we proceed as follows:

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix} \in \mathbb{F}_2^{3\times 8}$$

We can look at the entries of $Az = b$ to see (via the first measurement) if the single 1 is in the first or second half of z. If the first entry of $b = 0$ then we know it is in the first set of four coordinates, and similarly if $b = 1$, we know it is in the second set of four coordinates. Continuing the binary search, using the combined information of all the entries of b, we can recover the index of the single 1, and hence z. Note that in the general $k = 1$ case, we can take $A \in \mathbb{F}_2^{\log n \times n}$, we obtain that $m = \log n$ measurements suffice.

In the general k case, the problem is essentially trying to distinguish between the $\binom{n}{k}$ different vectors z can be. With m different measurements, we can at most hope to distinguish between 2^m different values (one guess per measurement). Therefore we want to ensure:

$$2^m \geq \binom{n}{k} \geq \left(\frac{n}{k}\right)^k$$
$$\implies m \geq k \log \frac{n}{k}$$

1.1.1 Discrete sparse recovery via the probabilistic method

The candidate construction for the measurements itself is the following: simply let each one be a random vector over $\mathbb{F}_2^n$. We now analyze its performance.

Suppose x_1 and x_2 are 2 different k-sparse length n vectors. Let's look at the probability that some random measurement (denoted as a particular row A_i) differs on the two:

Claim 1.

$$\Pr[A_i \cdot x_1 \neq A_i \cdot x_2] = 1/2.$$

Proof. Since $x_1 \neq x_2$ there exists an index j such that $x_{1j} \neq x_{2j}$. Let A_i', x_1', and x_2' denote the $(n-1)$-dimensional vectors identical to A_i, x_1, and x_2, respectively, except for coordinate j. We thus have

$$\begin{aligned} A_i \cdot x_1 &= A_i' \cdot x_1 + A_{ij} x_{1j} \\ A_i \cdot x_2 &= A_i' \cdot x_2 + A_{ij} x_{2j} \end{aligned}$$

If $A_i' \cdot x_1 = A_i' \cdot x_2$, then $A_i \cdot x_1$ differs from $A_i \cdot x_2$ if $A_{ij} = 1$, which happens with probability 1/2. If $A_i' \cdot x_1 \neq A_i' \cdot x_2$, then $A_i \cdot x_1$ differs from $A_i \cdot x_2$ if $A_{ij} = 1$, which happens with probability 1/2. The desired claim follows. □

Therefore, with m measurements, each drawn independently and uniformly at random, the probability that the results of the measurements are all equal for two different vectors is $1/2^m$. A union bound over all possible pairs of k-sparse vectors yields the probability that any pair is not distinguished is at most $\binom{n}{k}^2 \cdot$

$1/2^m$. For some error rate ε, we can bound the number of random measurements m we want as follows.

$$\begin{aligned}\binom{n}{k}^2 \frac{1}{2^m} \leq \left(\frac{en}{k}\right)^{2k} \frac{1}{2^m} < \varepsilon \\ \implies m > \log \frac{\left(\frac{en}{k}\right)^{2k}}{\varepsilon} \\ = 2k \log \frac{en}{k} + \log \frac{1}{\varepsilon} \\ = \mathcal{O}(k \log \frac{n}{k}),\end{aligned}$$

where we use the inequality $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$.

Note that this is a bound only on the number of measurements, which is separate from the process of actually trying to recover z. The recovery process involves solving the system of linear equations obtained from the measurements.

1.2 Generalizing to the continuous case

A problem that arises when trying to extend the approach of discrete sparse recovery to continuous setting is that one cannot apply a union bound over all possible pairs of k-sparse vectors, since there are an infinite number of them. To solve this issue, compressed sensing researchers introduced the concept of the restricted isometry property (RIP).

Definition 2. A matrix A is (k, ε)-RIP if for all k-sparse x:

$$(1 - \varepsilon)||x||_2 \leq ||Ax||_2 \leq (1 + \varepsilon)||x||_2$$

We establish the following theorem.

Theorem 1. If A is $(2k, \varepsilon)$-RIP with $\varepsilon < 1$ then z is the unique solution to:

$$\begin{aligned}&\min \|x\|_0 \\ &s.t. \\ &Ax = b \\ &\text{where } b = Az\end{aligned}$$

(Recall that $\|x\|_p = \left(\sum_i x_i^p\right)^{1/p}$ for $p > 0$. And $\|x\|_0$ is the number of nonnegative coordinates in x.)

Proof. Suppose $x \neq z$ was the solution to the minimizing program. This implies that $Ax = b$ and $Az = b$. We can write $z = x + \Delta$, with $A\Delta = 0$. Since the original z is assumed to be k-sparse, this implies x is also k-sparse, and that Δ is $2k$-sparse.

Using the fact that A is $(2k, \varepsilon)$-RIP, we can say $(1 - \varepsilon)\|\Delta\|_2 \leq \|A\Delta\|_2$, which implies that $\|A\Delta\|_2 > 0$ (since $\|\Delta\|_2 > 0$, a contradiction to $A\Delta = 0$). □

While this is a great result showing that a $(2k, \varepsilon)$-RIP A can help us recover z, there is an issue with the minimization problem. We do not know how to minimize ℓ_0-norm efficiently. To circumvent this, we define a different program that tries to minimize the ℓ_1-norm (which is efficiently solvable via a linear program).

Theorem 2. If A is $(3k, \varepsilon)$-RIP with ε sufficiently small then z is the unique solution to the following program.

$$\begin{aligned} &\min \|x\|_1 \\ &s.t. \\ &Ax = b \\ &\text{where } b = Az \end{aligned}$$

Similar to the previous proof (Theorem 1), we argue about the uniqueness of z by supposing some optimal solution $x \neq z$. It follows that we know $Ax = Az = b$, $||x||_1 \leq ||z||_1$, and can write $z = x + \Delta$, knowing that $A\Delta = 0$. We know that z is k-sparse, but unfortunately we can't say the same thing about x anymore, since all we know is the ℓ_1-norm of x is at most that of z, which does not bound the sparsity of x.

What we will argue is that Δ is "almost" k-sparse. Let S be the set of indices where z has non-zero entries. For any set I of indices and any vector y, let y_I represent the vector y restricted to the indices in I, and all other index values set to zero.

Lemma 1. $\|\Delta_S\|_1 \geq \|\Delta_{\bar{S}}\|_1$

Proof. Intuitively, it makes sense that the ℓ_1-norm of Δ should be greater on z's non-zero entries since otherwise x's ℓ_1 norm would be high. We can show this formally in the following way. We know $z|_{\bar{S}}$ is the all zero string, and express $\|x\|_1$ in terms of $\|z\|_1$ and $\|\Delta\|_1$ via the triangle inequality.

$$\begin{aligned} \|x\|_1 &\geq \|z\|_1 - \|\Delta_S\|_1 + \|\Delta_{\bar{S}}\|_1 \\ \implies \|\Delta_S\|_1 &\geq \|z\|_1 - \|x\|_1 + \|\Delta_{\bar{S}}\|_1 \\ &\geq \|\Delta_{\bar{S}}\|_1, \end{aligned}$$

where the last inequality holds since $\|x\|_1 \leq \|z\|_1$. ☐

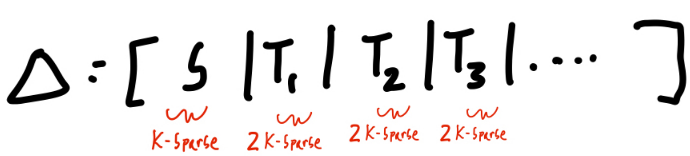

Figure 1: Rearranging indices of Δ

For the second lemma it will be useful to "rearrange" the indices/values of Δ and organize them in chunks. We renumber the indices so that the first k indices are from S and the remaining indices are sorted from largest to smallest value; so $\Delta_{k+1} \geq \Delta_{k+2} \geq \ldots \geq \Delta_n$. Let T_i denote the set of indices $\{k+(i-1)2k+j : 1 \leq j \leq 2k\}$, for $i \geq 1$. So, T_1 consists of the first $2k$ indices after S, T_2 the next $2k$ indices, etc. The following claim at a high level states that the ℓ_2 norm from the S indices is lower-bounded by a constant times the sum of the rest of Δ.

Lemma 2. $\|\Delta_S\|_2 \geq \sqrt{2} \cdot \sum_{j\geq 2} \|\Delta_{T_j}\|_2$

Proof. For k-sparse vectors v we know the following, where the first inequality follows from Cauchy-Schwartz and the second inequality follows from elementary algebra.

$$\sqrt{k}\|v\|_2 \geq \|v\|_1 \geq \|v\|_2 \tag{1}$$

Using this fact and Lemma 1, we derive the following.

$$\begin{aligned} \|\Delta_S\|_2 &\geq \frac{\|\Delta_s\|_1}{\sqrt{k}} && \text{Equation 1} \\ &\geq \frac{1}{\sqrt{k}} \sum_{j\geq 1} \|\Delta_{T_j}\|_1 && \text{Lemma 1} \\ &\geq \frac{\sqrt{2k}}{\sqrt{k}} \sum_{j\geq 2} \|\Delta_{T_j}\|_2 \end{aligned}$$

where we justify the last inequality as follows. It would suffice to show that $||\Delta_{T_j}||_1 \geq \sqrt{2k}||\Delta_{T_{j+1}}||_2$. For some T_j and T_{j+1}, we know from numbering of the indices and the definition of T_j that the value of Δ at each index of T_j is at least as large as that at each index of T_{j+1}. If $\Delta_{T_j} = (a_1, a_2, \ldots, a_{2k})$, then $\Delta_{T_{j+1}} \leq (a_{2k}, \ldots, a_{2k})$. This leads to an upper bound of $\sqrt{2k}a_{2k} \leq \|\Delta_{T_j}\|_1/\sqrt{2k}$ on $\|\Delta_{T_{j+1}}\|_2$, as desired. □

With these two claims, we are ready to prove Theorem 2.

As a reminder we know the following (assuming towards a contradiction that there is some $x \neq z$ that is an optimal solution):

- $Ax = b$
- $Az = b$
- $||x||_1 \leq ||z||_1$
- $z = x + \Delta$
- $A\Delta = 0$

We want to contradict the fact that $A\Delta = 0$. Using our "rearranged" representation of Δ we can say:

$$\|A\Delta\|_2 \geq \|A\Delta_{S\cup T_1}\|_2 - \sum_{j\geq 2} \|A\Delta_{T_j}\|_2$$

Notice that $\Delta_{S \cup T_1}$ is a $3k$-sparse vector (as with each of the T_j's) which means we can use the $(3k, \varepsilon)$-RIP of A to derive

$$\|A\Delta\|_2 \geq (1-\varepsilon)\|\Delta_{S \cup T_1}\|_2 - (1+\varepsilon)\|\Delta_{T_j}\|_2$$

By Lemma 2, we can replace the Δ_{T_j} terms and derive

$$\begin{aligned} \|A\Delta\|_2 &\geq (1-\varepsilon)\|\Delta_S\| - \frac{(1+\varepsilon)}{\sqrt{2}}\|\Delta_S\|_2 \\ &\geq + \left((1-\varepsilon) - \frac{(1+\varepsilon)}{\sqrt{2}}\right)\|\Delta_S\|, \end{aligned}$$

which can be made positive by choosing $\varepsilon < 0.17$. This arrives at our contradiction that $A\Delta = 0$.

We close by observing that the mathematical program in Theorem 2 can be implemented as a linear program since any term of the form $|y|$ can be replaced by u with the additional constraints $u \geq y$ and $u \geq -y$. Thus, Theorem 2 gives us an efficient way of recovering z, assuming we have a $(3k, \varepsilon)$-RIP matrix A.

1.3 Constructing A: J-L implies RIP

Now we know that if we have an RIP matrix, then we can do compressed sensing. One missing piece of the puzzle is to figure out how to construct RIP matrices. It turns out that we can actually just generate a random matrix! Recall the Johnson-Lindenstrauss Lemma [JL].

J-L Lemma: For any vector $v \in \mathbb{R}^n$, if A is a random $m \times n$ matrix in $\mathbb{F}_2$ (each entry is drawn from the normal distribution $N(0,1)$) (scaled by $1/\sqrt{m}$). Then, with probability $1 - e^{-c\varepsilon^2 m}$ for a suitable constant $c > 0$, we have

$$(1-\varepsilon)\|v\|_2 \leq \|Av\|_2 \leq (1+\varepsilon)\|v\|_2.$$

The last theorem we will show is that J-L implies RIP (which implies compressed sensing). What is noteworthy is that the J-L lemma is a probabilistic statement (a random matrix satisfies a certain condition with high probability) while RIP is not. A key proof strategy is to reason about the infinite set of k-sparse vectors (needed for RIP) by coming up with a finite representation.

More specifically, consider fixing the k indices of a k-sparse vector. Without loss of generality, suppose the first k indices are non-zero, and the rest are 0. Define the set Q for $i \in \mathbb{Z}$ and $\delta > 0$ as follows:

$$Q = \left\{ x = \left(i_1 \frac{\delta}{\sqrt{k}}, i_2 \frac{\delta}{\sqrt{k}}, \ldots i_k \frac{\delta}{\sqrt{k}}, 0, \ldots, 0 \right) : |x| \leq 1 \right\}.$$

We can bound the size of Q by looking at the number of possible values the i's can take to keep the norm ≤ 1. For each i, there are a maximum of $\frac{2\sqrt{k}}{\delta}$ values (the factor two coming from positive and negative values). Therefore, considering all the k indices we get $|Q| \leq (2\sqrt{k}/\delta)^k$.

With Q defined, we ask: how close is a k-sparse vector v with the assumed form to Q (where we define the distance to Q as the distance to the closest point x_v in Q). We observe that

$$\|v - x_v\|_2 \leq \sqrt{k(\delta/\sqrt{k})^2} = \delta.$$

What we will now show is that if we apply the J-L transform on Q, by choosing m (the number of measurements) sufficiently large, the set of measurements A will work "correctly" on Q, which will then imply it will work for all k-sparse points with nonzero entries in the first k indices.

Lemma 3. Suppose $\forall x \in Q$ $(1-\varepsilon)||x||_2 \leq ||Ax||_2 \leq (1+\varepsilon)||x||_2$. Then the above bound also holds for all all k-sparse v of our assumed form.

Proof. For $\|v\|_2 = 1$ we know from above that $\exists x_v$ s.t $||v - x_v||_2 \leq \delta$.

$$\begin{aligned}
||Av||_2 &\leq ||Ax_v||_2 + ||A(v - x_v)||_2 && (2)\\
&\leq (1+\varepsilon)||x_v||_2 + ||A(v - x_v)||_2 && (3)\\
&\leq (1+\varepsilon)||x_v||_2 + (1+\alpha)||v - x_v||_2 && (4)\\
&\leq (1+\varepsilon)\cdot 1 + (1+\alpha)\cdot\delta && (5)
\end{aligned}$$

Equation 2 is from an application of the triangle inequality, equation 3 from using the assumption that $x_v \in Q$. In equation 4 for simplicity we are bounding $||A(v - x_0)||_2 \leq (1+\alpha)||v - x_0||_2$, with no guarantees on α at the moment. Equation 5 comes from simplifying the norms based on what we know about vectors in Q and x_v.

This leads us to an upper bound $||Av||_2 \leq 1 + \varepsilon + \delta + \alpha\delta$. Since we can pick ε to be arbitrarily small, we can say $1 + \varepsilon + \delta + \alpha\delta \leq 1 + \mathcal{O}(\delta)$. This shows that if the J-L lemma holds for points in Q, then a similar bound holds for the v's. □

The final step is to ensure the probabilities for the J-L lemma on the v's are not too large (recall we assumed that $\forall x \in Q$ the J-L lemma holds). The probability that a point doesn't satisfy the J-L lemma is $e^{-c\varepsilon^2 m}$. We need to union bound over $|Q|$ *for each* possible Q. Recall that Q is defined for a *fixed* configuration of a k-sparse vector. There are $\binom{n}{k}$ different k-sparse vectors, each defining its own set Q. We can ensure success with probability $1/2$ by setting m sufficiently large so that it satisfies

$$\binom{n}{k}\left(\frac{2\sqrt{k}}{\delta}\right)^k e^{-c\varepsilon^2 m} \leq \frac{1}{2}.$$

If m satisfies the above condition, then with probability greater than 1/2 (which can be replaced with your favorite constant by selecting m suitably) a random $n \times m$ matrix for A will work for *all* points and *all* possible k.

We can solve for m to get a bound on how many measurements are needed (size of A). We can think of δ and ε as constants for the bound below. We need

$$e^{c\varepsilon^2 m} > \binom{n}{k}\left(\frac{2\sqrt{k}}{\delta}\right)^k \approx \mathcal{O}(kn)^k,$$

which is satisfied by $m = \mathcal{O}(k \ln n)$. We thus have the main compressed sensing result: there exists a set of $m = O(k \ln n)$ linear measurements that can be used to recover any k-sparse signal; furthermore, a randomly chosen set of m measurements will work with high probability.

References

[Asa20] Sepehr Asadi. Compressed sensing. Lecture notes, available from https://people.cs.rutgers.edu/~sa1497/courses/cs514-s20/lec6.pdf, 2020.

[BDDW06] Richard Baraniuk, M Davenport, Ronald DeVore, and M Wakin. The johnson-lindenstrauss lemma meets compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 01 2006.

[CRT06] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489—-509, 2006.

[Don06] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289—-1306, 2006.

[JL] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math.*, 26:189–206. Conference in modern analysis and probability.

[Mus18] Christopher Musco. Compressed sensing. Lecture notes, available from https://www.cs.princeton.edu/courses/archive/fall18/cos521/Lectures/lec21.pdf, 2018.

[RV17] T. Roughgarden and G. Valiant. Compressive sensing. Lecture notes, available from https://timroughgarden.org/s17/l/l17.pdf, 2017.