3. MSRI Graduate Summer School Notes
Jeff Blanchard, 2018
Foundations of Compressed Sensing

We finally begin our exploration of compressed sensing. The compressed sensing problem asks us to recover a sparse vector $x \in \mathbb{R}^n$ from the measurements $y \in \mathbb{R}^m$ whose entries are obtained by linear (inner product) measurements. This provides us with an ill-posed problem, an underdetermined system of linear equations. Let $A \in \mathbb{R}^{m\times n}$ be a matrix whose rows represent the linear measurements. Then $y = Ax$ is a model of this system. We begin by formalizing some terminology used throughout the remainder of the course.

Definition 3.1 Support. The *support* of a vector (signal) $x \in \mathbb{R}^n$ is the index set of the nonzero entries in the vector:

$$\text{supp}(x) = \{i : x_i \neq 0\}.$$

To count the number of nonzeros in a vector, we introduce a *norm-like* counting operator which we incorrectly call the ℓ_0 norm.

Definition 3.2 ℓ_0 Norm. The function $\|\cdot\|_0 : \mathbb{R}^n \to \mathbb{N}$ counts the number of nonzero entries in a vector. Let $x \in \mathbb{R}^n$, then

$$\|x\|_0 = |\text{supp}(x)| = \#\{i : x_i \neq 0\}.$$

One way to measure the sparsity of a vector is via the number of nonzero entries in the vector. We refer to a vector as k sparse when it has no more than k nonzero entries.

Definition 3.3. A vector $x \in \mathbb{R}^n$ is k-sparse if $\|x\|_0 \leq k$. The set of all k-sparse vectors in $\mathbb{R}^n$ is denoted $\chi_n(k)$,

$$\chi_n(k) = \{x \in \mathbb{R}^n : \|x\|_0 \leq k\}.$$

Throughout our analysis with sparse vectors, it will be useful to introduce a restriction of a vector to a subset of the entries. This definition is awkward as it will have two distinct meanings which will be used interchangeably but should be clear from context.

Definition 3.4 Restrictions. Let $S \subseteq \{1, \ldots, n\}$ be an index set.

1. The *restriction* of $x \in \mathbb{R}^n$ to the index set S is denoted x_S and is defined by
$$(x_S)_i = \begin{cases} x_i & \text{if } i \in S \\ 0 & \text{if } i \notin S \end{cases}.$$
2. Also, x_S will be used to denote the vector with $|S|$ entries, $x_S \in \mathbb{R}^{|S|}$, where
$$(x_S)_i = x_{S_i} \text{ for } i = 1, 2, \ldots, |S|.$$
3. If $A \in \mathbb{R}^{m\times n}$, the *restriction* of the matrix A to the index set S is the $m \times |S|$ matrix A_S whose columns are the columns of A indexed by S, namely
$$A_S = \left[A_{S_1}|A_{S_2}|\cdots|A_{S_{|S|}}\right].$$

The restrictions are useful in describing algorithms which select support sets in each iteration. The following exercise is a simple formality to help you get comfortable with restrictions.

Exercise 3.1. Suppose $x \in \mathbb{R}^n$ is a k-sparse vector, i.e. $x \in \chi_n(k)$, and let $S = \text{supp}(x)$. Prove that under the form 1. from Def. 3.4, $x = x_S$. Now prove that under forms 2. and 3. we have $Ax = A_S x_S$.

3.1. **Encoder and Decoders.** The point of compressed sensing is to simultaneously measure and compress the information in a sparse vector. In the traditional sampling paradigm we explored in Chapter 2, we sample a signal at n locations and compress the signal via a change of basis to an orthonormal basis. If the signal x is discrete, i.e. $x \in \mathbb{R}^n$, the sampling consists of taking n point measurements. In terms of inner products, this boils down to taking n inner products of the vector x with the standard basis vectors e_i, $i = 1, \ldots, n$. The matrix whose rows consist of the transposes of the standard basis vectors is clearly the $n \times n$ identity matrix. So our measurement scheme when taking the full n measurements is simply $y = I_{n\times n}x$. In this case, reconstruction of x is automatic since $y = x$. Therefore, we can always reconstruct a vector x from n measurements. In fact, we are not restricted to the identity matrix. If $A \in \mathbb{R}^{n\times n}$ is any full rank matrix,

then if $y = Ax$ we can reconstruct x via $x = A^{-1}y$ (or some more efficient algorithm). In Chap. 2 we did just this. Two examples of A are the Fourier matrix and a wavelet matrix.

Now suppose $S = \text{supp}(x)$ and $\|x\|_0 = k < n$. If we are fortunate enough to know S prior to measuring x, say because we are omniscient or because we have some oracle telling us the support, then we can completely observe the information in x via only k measurements.

Exercise 3.2. Suppose $S = \text{supp}(x)$ and $\|x\|_0 = k < n$, and suppose you had a method like a magic eight ball that would tell you the set S. Construct a $k \times n$ matrix A which will capture all the information in x and provide a reconstruction routine. (Don't over think this; it is as simple as it seems.)

Therefore, we know that we must take at least k measurements. If $x \in \chi_n(k)$ and $B \in \mathbb{R}^{l\times k}$, then $y = Bx_S \in \mathbb{R}^l$ is an underdetermined system of equations with no sparsity in the observed vector. This has infinitely many solutions as we recall from linear algebra. So, if we must take at least k measurements, and we never require more than n measurements of $x \in \chi_n(k)$, the goal of compressed sensing is to identify a measurement process with m measurements so that $k < m < n$. Ideally, we want to take only k measurements but lacking the support identification oracle forces us to take more than k measurements. Therefore, we seek m as close to k as possible; in other words, we want to take the fewest number of measurements possible so that we capture all the information in x.

Capturing the information is only one part of the process. It would be useless to capture the information in x if we were unable to the reconstruct x. A nice piece of terminology that we will adopt is the notion of an encoder-decoder pair.

Definition 3.5. Let $x \in \chi_n(k)$. An *encoder* is a matrix $A \in \mathbb{R}^{m\times n}$ so that the measurements $y \in \mathbb{R}^m$, $y = Ax$, capture the information content of x; we say A encodes x into y.

A *decoder* of (y, A) is an optimization method or an algorithm which when given the measurements y and the encoder A the method/algorithm is able to recover (reconstruct, decode) the original k-sparse vector x.

The underdetermined system $y = Ax$ has an infinite number of solutions. In Section 1.3, we saw that one way to select a particular solution is to solve the least squares problem. However, the least squares solution is usually dense, meaning that most of it entries (and very likely all of its entries) will be nonzero. If we want to simultaneously compress the signal, we want to find a sparse solution, not a dense solution. Therefore, in compressed sensing we seek the sparsest solution from the infinitely many solutions. In other words, we seek a vector x with the fewest number of nonzero entries which could have generated the measurements y.

Suppose we have a suitable encoder $A \in \mathbb{R}^{m\times n}$. Our first decoder is the naive decoder which simply searches for the sparsest solution to $y = Ax$.

Definition 3.6 ℓ_0 Decoder. Suppose $y = Ax \in \mathbb{R}^m$. A vector which solves to the optimization problem

$$\min_{z\in\mathbb{R}^n} \|z\|_0 \quad \text{subject to } y = Az = Ax$$

is the sparest solution vector to the underdetermined system of equations $y = Ax$.

Usually we want to identify the solution vector rather than the solution (value) to an optimization problem. It is standard in the literature to call the solution to an optimization problem the *argument*. In other words, if $x^\star$ is the vector which solves the optimization problem in Def. 3.6, we write

$$x^\star = \text{argmin}_{z\in\mathbb{R}^n}\|z\|_0 \quad \text{subject to } y = Az = Ax.$$

Exercise 3.3. Describe a the method for searching for the sparsest solution via the ℓ_0 decoder and count the number of linear systems that must be solved to complete this search.

Clearly this is far too expensive. In fact this is called a combinatorial search and the ℓ_0 decoder a combinatorial optimization problem. While we focus on this decoder in this chapter, the remainder of the course will focus on alternative, more efficient algorithms for finding the sparsest solution.

3.2. **Spark and Uniqueness.** Suppose x is sparse and that $y = Ax$. If there is a vector z so that $\|x\|_0 > \|z\|_0$ and $y = Az$, then x will not be the sparsest solution to compressed sensing problem. We are operating under the assumption that our vector x is indeed the sparsest solution so that we want to recover x. Now, if there is a vector z with $\|x\|_0 = \|z\|_0$ and $y = Az$, then we still might not get x as the solution for the ℓ_0 decoder; z is also a worthy solution. Therefore, we want to determine the situations where x will be the

unique solution to the ℓ_0 decoder. That is, we want our encoder A to ensure that x is the unique, sparsest vector generating the measurements y.

Our first tool is the notion of *spark* of an encoder introduced by Donoho and Elad. There are several ways to define spark, so we choose one and then prove equivalence of other notions.

Definition 3.7 Spark. The *spark* of a matrix $A \in \mathbb{R}^{m\times n}$ is the number of nonzero entries in the sparsest nonzero vector contained in the null space of A:

$$\text{spark}(A) = \min_{z\in\mathbb{R}^n - \{0\}} \|z\|_0 \quad \text{subject to } Az = 0.$$

Often, spark(A) is defined as the minimum size of a set of linearly dependent columns of A.

Exercise 3.4. Suppose $s = \text{spark}(A)$. Prove that s is minimum size of a set of linearly dependent columns of A.

(*HINT:* Prove this equivalence as follows: show that there exists a set of $s = \text{spark}(A)$ columns of A that is linearly dependent and then show that any set of $t < s$ columns of A must be linearly independent.)

From this exercise we can quickly obtain an upper bound on spark(A).

Exercise 3.5. Prove $spark(A) \leq m + 1$.

If a matrix has *full spark*, this means the matrix achieves this upper bound. This is a highly desirable property. This property is so desirable it has a standard name, *general position*. We will occasionally use general position in the future as an assumption on our matrices. The property is stronger than being full rank in that every $m \times m$ submatrix is full rank, rather than at least one $m \times m$ submatrix being full rank.

Definition 3.8 General Position. A matrix $A \in \mathbb{R}^{m\times n}$ with $m < n$ is in *general position* if every set of m columns of A spans $\mathbb{R}^m$, i.e. $spark(A) = m + 1$.

Ignoring general position for the time being, let's determine a uniform recovery guarantee for the ℓ_0 decoder. A recovery result is *uniform* if it implies that the decoder can recover all k sparse vectors. On the other hand, a *nonuniform* recovery guarantee will state that we can recover a specified k sparse vector. This difference is subtle, but important. Uniform guarantees are very strong; they say the decoder will always succeed and return the sparsest solution no matter what k sparse vector we are trying to find.

To arrive at our uniform recovery guarantee for the ℓ_0 decoder, we prove a lemma in the form of exercises.

Exercise 3.6. Prove that the following statements are equivalent:

(a) Every k sparse vector x is the unique sparsest solution to $Az = Ax$; i.e. if $x, z \in \chi_n(k)$ and $Az = Ax$, then $z = x$.
(b) Every set of $2k$ columns of A is linearly independent.

Let's pause to collect what we have proven.

Theorem 3.2.1. *Let $A \in \mathbb{R}^{m\times n}$. The following statements are equivalent:*

(a) Every k sparse vector x is the unique sparsest solution to $Az = Ax$; i.e. if $x, z \in \chi_n(k)$ and $Az = Ax$, then $z = x$.
(b) Every set of $2k$ columns of A is linearly independent.
(c) $spark(A) > 2k$.
(d) The only $2k$ sparse vector in the null space of A is the zero vector:

$$null(A) \cap \chi_n(2k) = \{0\}.$$

Proof.
- $(a) \Leftrightarrow (b)$: Exer. 3.6.
- $(b) \Leftrightarrow (c)$: Exer. 3.4.
- $(c) \Leftrightarrow (d)$: Def. 3.7.

□

Notice that the conditions in Lem. 3.2.1 are statements about the encoder A capturing the complete information in any k sparse vector in a unique fashion. This allows us to reconstruct k sparse vectors. The following theorem is a simple Corollary of Thm. 3.2.1. We'll call it a theorem anyway!

Theorem 3.2.2. *Suppose $A \in \mathbb{R}^{m\times n}$ and $x \in \chi_n(k)$ and let $y = Ax$. If*

$$k < \frac{spark(A)}{2}$$

then x is the unique, sparsest solution to the underdetermined system $y = Ax$.

Moreover, the ℓ_0 decoder will recover x from (y, A).

Exercise 3.7. Prove Thm. 3.2.2.

3.3. **Coherence.** A common method for analyzing decoders is the *coherence* of the encoder matrix A. The coherence of a matrix is the maximum magnitude of the inner products between any two distinct columns. This provides a quantitative description of how similar the columns of a matrix would be. In order to avoid the obvious problems imposed by simply scaling a single column of the matrix, we assume that each column in normalized, $\|A_j\|_2 = 1$ for all $j = 1, \dots, n$. In fact, we will assume this property of the encoder matrix A for the rest of the course.

- For every encoder matrix $A \in \mathbb{R}^{m\times n}$, we henceforth assume that $\|A_j\|_2 = 1$ for each column $j = 1, \dots, n$.
- We use the notation x^*y for the inner product, even when we assume the entries of x and y are real. (This will help you read other compressed sensing literature.)

Definition 3.9 Coherence. The *coherence* of a matrix $A \in \mathbb{R}^{m\times n}$ is the maximum magnitude of the inner product between any two distinct columns of A, namely

$$\mu(A) = \max_{i\neq j} |A_i^* A_j| \,. \tag{3}$$

As mentioned above, the coherence measures the worst case similarity between two columns in A. Equivalently, it is a measure of how distinct any two columns in A might be, thereby providing a measure of the redundancy in the n columns. A matrix with a low coherence will offer a better chance at recovering a sparse vector since the linear combination of the rows will be easier to identify. In fact, the case of an orthogonal matrix demonstrates this idea well. If $Q \in \mathbb{R}^{n\times n}$ is orthogonal, then identifying the nonzero entries in a sparse vector x from the measurements $y = Ax$ is straightforward.

Exercise 3.8. Prove that if $Q \in \mathbb{R}^{n\times n}$ is orthogonal (or unitary), then $\mu(Q) = 0$.

Of course, it will be impossible for a matrix $A \in \mathbb{R}^{m\times n}$ to have $\mu(A) = 0$ since there are more columns than rows. While we won't prove the following fact, the lower bound on the coherence of a matrix with m rows and n columns ($m < n$) is now well known.

Lemma 3.1. *For $A \in \mathbb{R}^{m\times n}$,*

$$\mu(A) \geq \sqrt{\frac{n-m}{m(n-1)}}.$$

When m and n are large, the right hand side of the inequality is approximately $1/\sqrt{m}$.

While the coherence is a more subtle indicator of independence of columns than spark, the two are closely related.

Theorem 3.3.1 (Spark and Coherence)**.** *For any matrix $A \in \mathbb{R}^{m\times n}$ with $\|A_j\|_2 = 1$,*

$$spark(A) \geq \min\{m + 1, \mu(A)^{-1} + 1\}. \tag{4}$$

This relationship between spark and coherence allow us to develop another uniqueness result which is a uniform recovery guarantee for the ℓ_0 decoder.

Theorem 3.3.2. *Suppose $A \in \mathbb{R}^{m\times n}$ and $x \in \chi_n(k)$ and let $y = Ax$. If*

$$k < \frac{\mu(A)^{-1} + 1}{2}$$

then x is the unique, sparsest solution to the underdetermined system $y = Ax$.

Moreover, the ℓ_0 decoder will recover x from (y, A).

Exercise 3.9. Prove Thm. 3.3.2. (For now, assume $\mu(A)^{-1} < m$ and use Thms. 3.2.2 and 3.3.1. A complete proof uses Exercise 3.14 directly.)

3.4. **Coherence Exercises.** Let's compute coherence a few times. We know the bound that $\mu(A) \gtrsim 1/\sqrt{m}$. There are special encoders that achieve this lower bound.

Exercise 3.10. Let F_m be the $m \times m$ Fourier matrix and I_m be the $m \times m$ identity matrix. Concatenating these two bases into a single redundant representation system,

$$A = [F_m \mid I_m] .$$

Compute $\mu(A)$.

Exercise 3.11. Write a Matlab script that will generate many random matrices and compute the coherence; be sure to normalize the columns. Then do some sort of visualization and/or descriptive statistics about the coherence of random matrices from various distributions (Normal - `randn(m,n)`; Uniform - `rand(m,n)`; Bernoulli - `sign(randn(m,n))/sqrt(m)`;)

To prove Thm. 3.3.1, we rely on a classical result about eigenvalues. The Gershgorin Disc Theorem states the every eigenvalue of a square matrix lies in a union of discs defined by the rows of the matrix A.

Theorem (Gershgorin Disc Theorem)**.** *Let $A \in \mathbb{R}^{n\times n}$. Every eigenvalue of A lies in the union of the closed discs $\overline{D(A_{ii}, r_i)}$ where $r_i = \sum_{j\neq i} |A_{ij}|$.*

This is a powerful theorem in that it provides a simple method for bounding the spectrum of a matrix. You can prove this theorem via two a process outlined at the end of the exercises. First, let's employ the Gershgorin Disc Theorem to prove Thm. 3.3.1. The following sequence of exercises build on each other and yield the proof of Thm. 3.3.1 which was the key ingredient in the proof of our uniform recovery result for the ℓ_0 decoder, Thm. 3.3.2.

Exercise 3.12. If $A \in \mathbb{R}^{m\times n}$ with $\|A_j\|_2 = 1$ for $j = 1, \ldots, n$, prove that the matrix $B = A_S^* A_S$ for any index set S has diagonal entries equal to 1 and the absolute values of all off diagonal entries are bounded above by $\mu(A)$.

Exercise 3.13. If $A \in \mathbb{R}^{m\times n}$ with $\|A_j\|_2 = 1$ for $j = 1, \ldots, n$ and $S \subset \{1, \ldots, n\}$ is an index set with $|S| = k$, prove that every singular value of A_S lies in the closed disc $\overline{D(1, (k-1)\mu(A))}$.

Exercise 3.14. If $A \in \mathbb{R}^{m\times n}$ with $\|A_j\|_2 = 1$ for $j = 1, \ldots, n$, prove that if

$$k < \frac{1}{\mu(A)} + 1,$$

then every set of k columns of A is linearly independent.

Exercise 3.15. Use the results of this sequence of exercises to prove

$$\operatorname{spark}(A) \geq \min\{m + 1, \mu(A)^{-1} + 1\}.$$

The following exercises provide a proof of the Gershgorin Disc Theorem. It's fun and you should do it if you have the time.

Exercise 3.16. Let λ be an eigenvalue of $A \in \mathbb{R}^{n\times n}$. Prove that there exists a diagonal entry A_{ll} such that

$$|A_{ll} - \lambda| \leq \sum_{j\neq l} |A_{lj}| .$$

(Hint: Let v be an eigenvector associated with λ such that $\|v\|_\infty = 1$ and let $1 \leq l \leq n$ be the index so that $|v_l| = 1$.)

Exercise 3.17. Use the previous exercise to prove the Gershgorin Disc Theorem.