# 5
# Sparse Recovery

In this chapter, we will witness the power of *sparsity* for the first time. Let's get a sense of what it's good for. Consider the problem of solving an underdetermined linear system $Ax = b$. If we are given $A$ and $b$, there's no chance to recover $x$ uniquely, right? Well, not if we know that $x$ is sparse. In that case, there are natural conditions on $A$ where we actually will be able to recover $x$ even though the number of rows of $A$ is comparable to the sparsity of $x$ rather than its dimension. Here we will cover the theory of sparse recovery. And in case you're curious, it's an area that not only has some theoretical gems, but also has had major practical impact.

## 5.1 Introduction

In signal processing (particularly imaging), we are often faced with the task of recovering some unknown signal given linear measurements of it. Let's fix our notation. Throughout this chapter, we will be interested in solving a linear system $Ax = b$ where $A$ is an $m \times n$ matrix and $x$ and $b$ are $n$ and $m$ dimensional vectors, respectively. In our setup, both $A$ and $b$ are known. You can think of $A$ as representing the input-output functionality of some measurement device we are using.

Now, if $m < n$, then we cannot hope to recover $x$ uniquely. At best we could find some solution $y$ that satisfies $Ay = b$ and we would have the promise that $x = y + z$, where $z$ belongs to the kernel of $A$. This tells us that if we want to recover an $n$-dimensional signal, we need at least $n$ linear measurements. This is quite natural. Sometimes you will hear this referred to as the Shannon-Nyquist rate, although I find that a rather opaque way to describe what is going on. The amazing idea that will save us is that if $x$ is sparse — i.e., $b$ is a linear combination of only a few columns of $A$ — then we really will be able to get



https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

away with many fewer linear measurements and still be able to reconstruct $x$ exactly.

What I want to do in this section is explain why you actually should not be surprised by it. If you ignore algorithms (which we won't do later on), it's actually quite simple. It turns out that assuming that $x$ is sparse isn't enough by itself. We will always have to make some structural assumption about $A$ as well. Let's consider the following notion:

**Definition 5.1.1** *The Kruskal rank of a set of vectors $\{A_i\}_i$ is the maximum $r$ such that all subsets of at most $r$ vectors are linearly independent.*

If you are given a collection of $n$ vectors in $n$ dimensions, they can all be linearly independent, in which case their Kruskal rank is $n$. But if you have $n$ vectors in $m$ dimensions — like when we take the columns of our sensing matrix $A$ – and $m$ is smaller than $n$, the vectors can't be all linearly independent, but they can still have Kruskal rank $m$. In fact, this is the common case:

**Claim 5.1.2** *If $A_1, A_2, \ldots, A_n$ are chosen uniformly at random from $\mathbb{S}^{m-1}$, then almost surely their Kruskal rank is $m$.*

Now let's prove our first main result about sparse recovery. Let $\|x\|_0$ be the number of nonzero entries of $x$. We will be interested in the following highly non-convex optimization problem:

$$(P_0) \qquad \min \|w\|_0 \text{ s.t. } Aw = b$$

Let's show that if we could solve $(P_0)$, we could find $x$ from much fewer than $n$ linear measurements:

**Lemma 5.1.3** *Let $A$ be an $m \times n$ matrix whose columns have Kruskal rank at least $r$. Let $x$ be an $r/2$-sparse vector and let $Ax = b$. Then the unique optimal solution to $(P_0)$ is $x$.*

*Proof:* We know that $x$ is a solution to $Ax = b$ that has objective value $\|x\|_0 = r/2$. Now suppose there were any other solution $y$ that satisfies $Ay = b$.

Consider the difference between these solutions, i.e., $z = x - y$. We know that $z$ is in the kernel of $A$. However, $\|z\|_0 \geq r + 1$, because by assumption every set of at most $r$ columns of $A$ is linearly independent. Finally, we have

$$\|y\|_0 \geq \|z\|_0 - \|x\|_0 \geq r/2 + 1$$

which implies that $y$ has larger objective value than $x$. This completes the proof. ■

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

So if we choose the columns of our sensing matrix to be random $m$ dimensional vectors, then from just $m$ linear measurements we can, in principle, recover any $m/2$-sparse vector uniquely. But there is a huge catch. Solving $(P_0)$ — i.e., finding the sparsest solution to a system of linear equations — is *NP*-hard. In fact, this is a simple and important reduction that is worth seeing. Following Khachiyan [97], let's start from the subset sum problem, which is a standard *NP*-hard problem:

**Problem 1** *Given distinct values $\alpha_1, \ldots, \alpha_n \in \mathbb{R}$, does there exist a set $I \subseteq [n]$ so that $|I| = m$ and $\sum_{i \in I} \alpha_i = 0$?*

We will embed an instance of this problem into the problem of finding the sparsest nonzero vector in a given subspace. We will make use of the following mapping, which is called the *weird moment curve*:

$$\Gamma'(\alpha_i) = [1, \alpha_i, \alpha_i^2, \ldots, \alpha_i^{m-2}, \alpha_i^m]$$

The difference between this and the standard moment curve is in the last term, where we have $\alpha_i^m$ instead of $\alpha_i^{m-1}$.

**Lemma 5.1.4** *A set $I$ with $|I| = m$ has $\sum_{i \in I} \alpha_i = 0$ if and only if the vectors $\{\Gamma'(\alpha_i)\}_{i \in I}$ are linearly dependent.*

*Proof:* Consider the determinant of the matrix whose columns are $\{\Gamma'(\alpha_i)\}_{i \in I}$. Then the proof is based on the following observations:

(a) The determinant is a polynomial in the variables $\alpha_i$ with total degree $\binom{m}{2} + 1$, which can be seen by writing the determinant in terms of its Laplace expansion (see, e.g., [88]).
(b) Moreover, the determinant is divisible by $\prod_{i<j} \alpha_i - \alpha_j$, since the determinant is zero if any $\alpha_i = \alpha_j$.

Hence we can write the determinant as

$$\Big(\prod_{\substack{i<j \\ i,j \in I}} (\alpha_i - \alpha_j)\Big)\Big(\sum_{i \in I} \alpha_i\Big).$$

We have assumed that the $\alpha_i$'s are distinct, and consequently the determinant is zero if and only if the sum of $\alpha_i = 0$. ■

We can now prove a double whammy. Not only is solving $(P_0)$ *NP*-hard, but so is computing the Kruskal rank:

**Theorem 5.1.5** *Both computing the Kruskal rank and finding the sparsest solution to a system of linear equations are NP-hard.*

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

*Proof:* First let's prove that computing the Kruskal rank is *NP*-hard. Consider the vectors $\{\Gamma'(\alpha_i)\}_i$. It follows from Lemma 5.1.4 that if there is a set $I$ with $|I| = m$ that satisfies $\sum_{i \in I} \alpha_i = 0$, then the Kruskal rank of $\{\Gamma'(\alpha_i)\}_i$ is at most $m-1$, and otherwise is exactly $m$. Since subset sum is *NP*-hard, so too is deciding whether the Kruskal rank is $m$ or at most $m-1$.

Now let's move on to showing that finding the sparsest solution to a linear system is *NP*-hard. We will use a one-to-many reduction. For each $j$, consider the following optimization problem:

$$(P_j) \qquad \min \|w\|_0 \text{ s.t. } \left[\Gamma'(\alpha_1), \ldots, \Gamma'(\alpha_{j-1}), \Gamma'(\alpha_{j+1}), \ldots, \Gamma'(\alpha_n)\right] w = \Gamma'(\alpha_j)$$

It is easy to see that the Kruskal rank of $\{\Gamma'(\alpha_i)\}_i$ is at most $m-1$ if and only if there is some $j$ so that $(P_j)$ has a solution whose objective value is at most $m-2$. Thus $(P_0)$ is also *NP*-hard. ■

In the rest of this chapter, we will focus on algorithms. We will give simple greedy methods as well as ones based on convex programming relaxations. These algorithms will work under more stringent assumptions on the sensing matrix $A$ than just that its columns have large Kruskal rank. Nevertheless, all of the assumptions we make will still be met by a randomly chosen $A$, as well as many others. The algorithms we give will even come with stronger guarantees that are meaningful in the presence of noise.

## 5.2 Incoherence and Uncertainty Principles

In 1965, Logan [107] discovered a striking phenomenon. If you take a band-limited signal and corrupt it at a sparse set of locations, it is possible to uniquely recover the original signal. This turns out to be a sparse recovery problem in disguise. Let's formalize this:

**Example 1** *The spikes-and-sines matrix $A$ is an $n \times 2n$ matrix*

$$A = [I, D]$$

*where $I$ is the identity matrix and $D$ is the discrete Fourier transform matrix, i.e.,*

$$D_{a,b} = \frac{\omega^{(a-1)(b-1)}}{\sqrt{n}}$$

*and $\omega = e^{2\pi i/n}$ is the $n^{\text{th}}$ root of unity.*

Let $x$ be a sparse $2n$-dimensional vector. The nonzeros in the first $n$ coordinates represent the locations of the corruptions. The nonzeros in the last

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

$n$ coordinates represent the frequencies present in the original signal. Thus we know $A$ and $b$ and are promised that there is a solution $x$ to $Ax = b$ where $x$ is sparse. It took a number of years until the work of Donoho and Stark [64], who realized that this phenomenon wasn't limited to just the spike-and-sines matrix. It's actually a quite general phenomenon. The key is the notion of incoherence:

**Definition 5.2.1** *The columns of $A \in \mathbb{R}^{n \times m}$ are $\mu$-incoherent if for all $i \neq j$*

$$|\langle A_i, A_j \rangle| \leq \mu \|A_i\| \cdot \|A_j\|.$$

Throughout this section, we will focus on just the case when the columns of $A$ are unit vectors. Hence a matrix is $\mu$-incoherent if for all $i \neq j$, $|\langle A_i, A_j \rangle| \leq \mu$. However, all the results we derive here can be extended to general $A$ when the columns are not necessarily unit vectors. As we did for the Kruskal rank, let's show that random vectors are incoherent:

**Claim 5.2.2** *If $A_1, A_2, \ldots, A_m$ are chosen uniformly at random from $\mathbb{S}^{n-1}$, then with high probability they will be $\mu$-incoherent for*

$$\mu = O\left(\sqrt{\frac{\log m}{n}}\right).$$

You can also check that the spike-and-sines matrix is $\mu$-incoherent with $\mu = 1/\sqrt{n}$. In that way, the results we derive here will contain Logan's phenomenon as a special case. Anyway, let's now show that if $A$ is incoherent and if $x$ is sparse enough, then it will be the uniquely sparsest solution to $Ax = b$.

**Lemma 5.2.3** *Let $A$ be an $n \times m$ matrix that is $\mu$-incoherent and whose columns are unit norm. If $Ax = b$ and $\|x\|_0 < \frac{1}{2\mu}$, then $x$ is the uniquely sparsest solution to the linear system.*

*Proof:* Suppose for the sake of contradiction that we have another solution $y$ that satisfies $Ay = b$ and $\|y\|_0 < \frac{1}{2\mu}$. Then we can look at the difference between these solutions, i.e., $z = x - y$, which satisfies $\|z\|_0 < \frac{1}{\mu}$ and consider the expression

$$z^T A^T A z = 0.$$

If we let $S$ denote the support of $z$ — i.e., the locations where it is nonzero — we have that $A^T A$ restricted to the rows and columns in $S$ is singular. Let this matrix be $B$. Then $B$ has ones along the diagonal, and the entries off the diagonal are bounded by $\mu$ in absolute value. But by Gershgorin's disk theorem, we know that all the eigenvalues of $B$ are contained in a disk in the

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

complex plane centered at one with radius $\mu|S| < 1$. Thus $B$ is nonsingular and we have a contradiction. ■

Actually, we can prove a stronger uniqueness result when $A$ is the union of two orthonormal bases, as is the case for the spikes-and-sines matrix. Let's first prove the following result, which we will mysteriously call an uncertainty principle:

**Lemma 5.2.4** *Let $A = [U, V]$ be an $n \times 2n$ matrix that is $\mu$-incoherent where $U$ and $V$ are $n \times n$ orthogonal matrices. If $b = U\alpha = V\beta$, then $\|\alpha\|_0 + \|\beta\|_0 \geq \frac{2}{\mu}$.*

*Proof:* Since $U$ and $V$ are orthonormal, we have that $\|b\|_2 = \|\alpha\|_2 = \|\beta\|_2$. We can rewrite $b$ as either $U\alpha$ or $V\beta$, and hence $\|b\|_2^2 = |\beta^T(V^T U)\alpha|$. Because $A$ is incoherent, we can conclude that each entry of $V^T U$ has absolute value at most $\mu(A)$, and so $|\beta^T(V^T U)\alpha| \leq \mu(A)\|\alpha\|_1\|\beta\|_1$. Using Cauchy-Schwarz, it follows that $\|\alpha\|_1 \leq \sqrt{\|\alpha\|_0}\|\alpha\|_2$ and thus

$$\|b\|_2^2 \leq \mu(A)\sqrt{\|\alpha\|_0\|\beta\|_0}\|\alpha\|_2\|\beta\|_2.$$

Rearranging, we have $\frac{1}{\mu(A)} \leq \sqrt{\|\alpha\|_0\|\beta\|_0}$. Finally, applying the AM-GM inequality, we get $\frac{2}{\mu} \leq \|\alpha\|_0 + \|\beta\|_0$ and this completes the proof. ■

This proof was short and simple. Perhaps the only confusing part is why we called it an uncertainty principle. Let's give an application of Lemma 5.2.4 to clarify this point. If we set $A$ to be the spikes-and-sines matrix, we get that any non-zero signal must have at least $\sqrt{n}$ nonzeros in the standard basis or in the Fourier basis. What this means is that no signal can be sparse in both the time and frequency domains simultaneously! It's worth taking a step back. If we had just proven this result, you would have naturally associated it with the Heisenberg uncertainty principle. But it turns out that what's really driving it is just the incoherence of the time and frequency bases for our signal, and it applies equally well to many other pairs of bases.

Let's use our uncertainty principle to prove an even stronger uniqueness result:

**Claim 5.2.5** *Let $A = [U, V]$ be an $n \times 2n$ matrix that is $\mu$-incoherent where $U$ and $V$ are $n \times n$ orthogonal matrices. If $Ax = b$ and $\|x\|_0 < \frac{1}{\mu}$, then $x$ is the uniquely sparsest solution to the linear system.*

*Proof:* Consider any alternative solution $A\widetilde{x} = b$. Set $y = x - \widetilde{x}$, in which case $y \in \ker(A)$. Write $y$ as $y = [\alpha_y, \beta_y]^T$, and since $Ay = 0$, we have that $U\alpha_y = -V\beta_y$. We can now apply the uncertainty principle and conclude that $\|y\|_0 = \|\alpha_y\|_0 + \|\beta_y\|_0 \geq \frac{2}{\mu}$. It is easy to see that $\|\widetilde{x}\|_0 \geq \|y\|_0 - \|x\|_0 > \frac{1}{\mu}$, and so $\widetilde{x}$ has strictly more nonzeros than $x$ does, and this completes the proof. ■

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

We can connect incoherence back to our original discussion about Kruskal rank. It turns out that having a matrix whose columns are incoherent is just one easy-to-check way to certify a lower bound on the Kruskal rank. The proof of the following claim is essentially the same as the proof of Lemma 5.2.3. We leave it as an exercise for the reader.

**Claim 5.2.6** *If $A$ is $\mu$-incoherent, then the Kruskal rank of the columns of $A$ is at least $1/\mu$.*

In the next section, we will give a simple greedy algorithm for solving sparse recovery problems on incoherent matrices. The way the algorithm will certify that it is making progress and finding the right nonzero locations of $x$ as it goes along will revolve around the same ideas that underly the uniqueness results we just proved.

## 5.3 Pursuit Algorithms

There is an important class of algorithms for sparse recovery problems called pursuit algorithms. These algorithms are greedy and iterative. They work with incoherent matrices and look for the column in $A$ that explains as much of the observed vector $b$ as possible. They subtract off a multiple of that column and continue on the remainder. The first such algorithm was introduced in an influential paper of Mallat and Zhang [111] and was called matching pursuit. In this section, we will analyze a variant of it called *orthogonal matching pursuit*. What's particularly convenient about the latter is that the algorithm will maintain the invariant that the remainder is orthogonal to all the columns of $A$ we have selected so far. This is more expensive in each step, but is easier to analyze and understand the intuition behind.

Throughout this section, let $A$ be an $n \times m$ matrix that is $\mu$-incoherent. Let $x$ be $k$-sparse with $k < 1/(2\mu)$ and let $Ax = b$. Finally, we will use $T$ to denote the support of $x$ — i.e., the locations of the nonzeros in $x$. Now let's formally define orthogonal matching pursuit:

**Orthogonal Matching Pursuit**

Input: matrix $A \in \mathbb{R}^{n \times m}$, vector $b \in \mathbb{R}^n$, desired number of nonzero entries $k \in \mathbb{N}$
Output: solution $x$ with at most $k$ nonzero entries

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

Initialize: $x^0 = 0$, $r^0 = Ax^0 - b$, $S = \emptyset$
For $\ell = 1, 2, \ldots, k$
  Choose column $j$ that maximizes $\frac{|\langle A_j, r^{\ell-1}\rangle|}{\|A_j\|_2^2}$.
Add $j$ to $S$.
  Set $r^\ell = \text{proj}_{U^\perp}(b)$, where $U = \text{span}(A_S)$.
  If $r^\ell = 0$, break.
End
Solve for $x_S$: $A_S x_S = b$. Set $x_{\bar{S}} = 0$.

Our anlaysis will focus on establishing the following two invariants:

(a) Each index $j$ the algorithm selects is in $T$.
(b) Each index $j$ gets chosen at most once.

These two properties immediately imply that orthogonal matching pursuit recovers the true solution $x$, because the residual error $r^\ell$ will be nonzero until $S = T$, and moreover, the linear system $A_T x_T = b$ has a unique solution (which we know from the previous section).

Property **(b)** is straightforward, because once $j \in S$ at every subsequent step in the algorithm, we will have that $r^\ell \perp U$, where $U = \text{span}(A_S)$, so $\langle r^\ell, A_j \rangle = 0$ if $j \in S$. Our main goal is to establish property **(a)**, which we will prove inductively. The main lemma is:

**Lemma 5.3.1** *If $S \subseteq T$ at the start of a stage, then orthogonal matching pursuit selects $j \in T$.*

We will first prove a helper lemma:

**Lemma 5.3.2** *If $r^{\ell-1}$ is supported in $T$ at the start of a stage, then orthogonal matching pursuit selects $j \in T$.*

*Proof:* Let $r^{\ell-1} = \sum_{i \in T} y_i A_i$. Then we can reorder the columns of $A$ so that the first $k'$ columns correspond to the $k'$ nonzero entries of $y$, in decreasing order of magnitude:

$$\underbrace{|y_1| \geq |y_2| \geq \cdots \geq |y_{k'}| > 0}_{\text{corresponds to first } k' \text{ columns of } A}, \quad |y_{k'+1}| = 0, |y_{k'+2}| = 0, \ldots, |y_m| = 0$$

where $k' \leq k$. Hence $\text{supp}(y) = \{1, 2, \ldots, k'\} \subseteq T$. Then, to ensure that we pick $j \in T$, a sufficient condition is that

$$|\langle A_1, r^{\ell-1}\rangle| > |\langle A_i, r^{\ell-1}\rangle| \qquad \text{for all } i \geq k' + 1. \tag{5.1}$$

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

We can lower-bound the left-hand side of (5.1):

$$|\langle r^{\ell-1}, A_1\rangle| = \left|\left\langle \sum_{\ell=1}^{k'} y_\ell A_\ell, A_1 \right\rangle\right| \geq |y_1| - \sum_{\ell=2}^{k'} |y_\ell||\langle A_\ell, A_1\rangle|$$
$$\geq |y_1| - |y_1|(k'-1)\mu \geq |y_1|(1 - k'\mu + \mu),$$

which, under the assumption that $k' \leq k < 1/(2\mu)$, is strictly lower-bounded by $|y_1|(1/2 + \mu)$.

We can then upper-bound the right-hand side of (5.1):

$$|\langle r^{\ell-1}, A_i\rangle| = \left|\left\langle \sum_{\ell=1}^{k'} y_\ell A_\ell, A_i \right\rangle\right| \leq |y_1| \sum_{\ell=1}^{k'} |\langle A_\ell, A_i\rangle| \leq |y_1| k'\mu,$$

which, under the assumption that $k' \leq k < 1/(2\mu)$, is strictly upper-bounded by $|y_1|/2$. Since $|y_1|(1/2+\mu) > |y_1|/2$, we conclude that condition (5.1) holds, guaranteeing that the algorithm selects $j \in T$, and this completes the proof. ■

Now we can prove Lemma 5.3.1:

*Proof:* Suppose that $S \subseteq T$ at the start of a stage. Then the residual $r^{\ell-1}$ is supported in $T$, because we can write it as

$$r^{\ell-1} = b - \sum_{i \in S} z_i A_i, \text{ where } z = \arg\min \|b - A_S z_S\|_2.$$

Applying the above lemma, we conclude that the algorithm selects $j \in T$. ■

This establishes property **(a)** inductively and completes the proof of correctness for orthogonal matching pursuit, which we summarize below:

**Theorem 5.3.3** *Let $A$ be an $n \times m$ matrix that is $\mu$-incoherent and whose columns are unit norm. If $Ax = b$ and $\|x\|_0 < \frac{1}{2\mu}$, then the output of orthogonal matching pursuit is exactly $x$.*

Note that this algorithm works up to exactly the threshold where we established uniqueness. However, in the case where $A = [U, V]$ and $U$ and $V$ are orthogonal, we proved a uniqueness result that is better by a constant factor. There is no known algorithm that matches the best known uniqueness bound there, although there are better algorithms than the one above (see, e.g., [67]).

It is also worth mentioning how other pursuit algorithms differ. For example, in matching pursuit we do not recompute the coefficients $x_i$ for $i \in S$ at the end of each stage. We just keep whatever they are set to and hope that they do not need to be adjusted much when we add a new index $j$ to $S$.

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

This is what makes matching pursuit faster in practice; however, the analysis is more cumbersome because we need to keep track of how the error (due to not projecting $b$ on the orthogonal complement of the columns we've chosen so far) accumulates.

## 5.4 Prony's Method

There is a widespread misconception that sparse recovery algorithms are a modern invention. Actually, sparse recovery dates back to 1795, to an algorithm called Prony's method. It will give us almost everything we want. We will have an explicit $2k \times n$ sensing matrix $A$ for which we will be able to recover any $k$-sparse signal exactly and with an efficient algorithm. It even has the benefit that we can compute the matrix-vector product $Ax$ in $O(n \log n)$ time using the fast Fourier transform.

The caveat to this method is that it is very unstable, since it involves inverting a Vandermonde matrix, which can be very ill-conditioned. So when you hear about compressed sensing as breaking the Shannon-Nyquist barrier, you should remember that Prony's method already does that. What sets apart the algorithms we will study later on is that they work in the presence of noise. That's the crucial aspect that makes them so practically relevant. Nevertheless, Prony's method is very useful from a theoretical standpoint, and the types of results you can get out of it have a habit of being rediscovered under other names.

### Properties of the Discrete Fourier Transform

Prony's method will make crucial use of various properties of the discrete Fourier transform. Recall that as a matrix, this transformation has entries

$$F_{a,b} = \left(\frac{1}{\sqrt{n}}\right) \exp\left(\frac{i2\pi(a-1)(b-1)}{n}\right).$$

As we did before, we will simplify the notation and write $\omega = e^{i2\pi/n}$ for the $n^{\text{th}}$ root of unity. With this notation, the entry in row $a$, column $b$ is $\omega^{(a-1)(b-1)}$.

The matrix $F$ has a number of important properties, including:

(a) $F$ is orthonormal: $F^H F = FF^H$, where $F^H$ is the Hermitian transpose of $F$.
(b) $F$ diagonalizes the convolution operator.

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

We haven't defined convolution, so let's do that now. Actually, let's do that through its corresponding linear transformation:

**Definition 5.4.1 (Circulant matrix)** *For a vector $c = [c_1, c_2, \ldots, c_n]$, let*

$$M^c = \begin{bmatrix} c_n & c_{n-1} & c_{n-2} & \ldots & c_1 \\ c_1 & c_n & c_{n-1} & \ldots & c_2 \\ \vdots & & & & \vdots \\ c_{n-1} & \ldots & \ldots & \ldots & c_n \end{bmatrix}.$$

Then the matrix-vector product $M^c x$ is the vector we get out of convolving $c$ and $x$, which we will denote by $c * x$. Intuitively, if you think of $c$ and $x$ as representing the probability distribution of discrete random variables, then $c*x$ represents the distribution of the random variable you get by adding the two of them and wrapping around $n$ using modular arithmetic.

As we asserted above, we can diagonalize $M^c$ using $F$. More formally, we have the following fact, which we will use without proof:

**Claim 5.4.2** *$M^c = F^H diag(\widehat{c})F$, where $\widehat{c} = Fc$.*

This tells us that we can think of convolution as coordinate-wise multiplication in the Fourier representation. More precisely:

**Corollary 5.4.3** *Let $z = c * x$; then $\widehat{z} = \widehat{c} \odot \widehat{x}$, where $\odot$ indicates coordinate-wise multiplication.*

*Proof:* We can write $z = M^c x = F^H \mathrm{diag}(\widehat{c})Fx = F^H \mathrm{diag}(\widehat{c})\widehat{x} = F^H(\widehat{c} \odot \widehat{x})$, and this completes the proof. ■

## The helper Polynomial

Prony's method revolves around the following helper polynomial:

**Definition 5.4.4 (helper polynomial)**

$$\begin{aligned} p(z) &= \prod_{b \in \mathrm{supp}(x)} \omega^{-b}(\omega^b - z) \\ &= 1 + \lambda_1 z + \ldots + \lambda_k z^k \end{aligned}$$

**Claim 5.4.5** *If we know $p(z)$, we can find supp$(x)$.*

*Proof:* In fact, an index $b$ is in the support of $x$ if and only if $p(\omega^b) = 0$. So we can evaluate $p$ at powers of $\omega$, and the exponents where $p$ evaluates to a nonzero are exactly the support of $x$. ■

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

The basic idea of Prony's method is to use the first $2k$ values of the discrete Fourier transform to find $p$, and hence the support of $x$. We can then solve a linear system to actually find the values of $x$. Our first goal is to find the helper polynomial. Let

$$v = [1, \lambda_1, \lambda_2, \ldots, \lambda_k, 0, \ldots, 0], \text{ and } \widehat{v} = Fv.$$

It is easy to see that the value of $\widehat{v}$ at index $b+1$ is exactly $p(\omega^b)$.

**Claim 5.4.6** $supp(\widehat{v}) = \overline{supp(x)}$

That is, the zeros of $\widehat{v}$ correspond to the roots of $p$, and hence nonzeros of $x$. Conversely, the nonzeros of $\widehat{v}$ correspond to the zeros of $x$. We conclude that $x \odot \widehat{v} = 0$, and so:

**Corollary 5.4.7** $M^{\widehat{x}} v = 0$

*Proof:* We can apply Claim 5.4.2 to rewrite $x \odot \widehat{v} = 0$ as $\widehat{x} * v = \widehat{0} = 0$, and this implies the corollary. ■

Let us write out this linear system explicitly:

$$M^{\widehat{x}} = \begin{bmatrix} \widehat{x}_n & \widehat{x}_{n-1} & \ldots & \widehat{x}_{n-k} & \ldots & \widehat{x}_1 \\ \widehat{x}_1 & \widehat{x}_n & \ldots & \widehat{x}_{n-k+1} & \ldots & \widehat{x}_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \widehat{x}_{k+1} & \widehat{x}_k & \ldots & \widehat{x}_1 & \ldots & \widehat{x}_{k+2} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \widehat{x}_{2k} & \widehat{x}_{2k-1} & \ldots & \widehat{x}_k & \ldots & \widehat{x}_{2k+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$

Recall, we do not have access to all the entries of this matrix, since we are only given the first $2k$ values of the DFT of $x$. However, consider the $k \times k+1$ submatrix whose top left value is $\widehat{x}_{k+1}$ and whose bottom right value is $\widehat{x}_k$. This matrix only involves the values that we do know!

Consider

$$\begin{bmatrix} \widehat{x}_k & \widehat{x}_{k-1} & \ldots & \widehat{x}_1 \\ \vdots & & & \\ \widehat{x}_{2k-1} & \widehat{x}_{2k-1} & \ldots & \widehat{x}_k \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_k \end{bmatrix} = - \begin{bmatrix} \widehat{x}_{k+1} \\ \vdots \\ \vdots \\ \widehat{x}_{2k} \end{bmatrix}$$

It turns out that this linear system is full rank, so $\lambda$ is the unique solution to the linear system (the proof is left to the reader[1]). The entries in $\lambda$ are the

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

coefficients of $p$, so once we solve for $\lambda$, we can evaluate the helper polynomial on $\omega^b$ to find the support of $x$. All that remains is to find the values of $x$. Indeed, let $M$ be the restriction of $F$ to the columns in $S$ and its first $2k$ rows. $M$ is a Vandermonde matrix, so again $Mx_S = \widehat{x}_{1,2,\ldots 2k}$ has a unique solution, and we can solve this linear system to find the nonzero values of $x$.

The guarantees of Prony's method are summarized in the following theorem:

**Theorem 5.4.8** *Let $A$ be the $2k \times n$ matrix obtained from taking the first $2k$ rows of the discrete Fourier transform matrix $F$. Then for any $k$-sparse signal $x$, Prony's method recovers $x$ exactly from $Ax$.*

In case you're curious, this is yet another topic in sparse recovery that we can relate back to Kruskal rank. It is easy to show that the columns of $A$ have Kruskal rank equal to $2k$. In fact, this is true regardless of which $2k$ rows of $F$ we choose. Moreover, it turns out that there are settings where Prony's method and related methods can be shown to work in the presence of noise, but only under some separation conditions on the nonzero locations in $x$. See Moitra [113] for further details.

## 5.5 Compressed Sensing

In this section we will introduce a powerful new assumption about our sensing matrix $A$, called the restricted isometry property. You can think of it as a robust analogue of the Kruskal rank, where not only do we want every set of (say) $2k$ columns of $A$ to be linearly independent, we also want them to be well-conditioned. We will show that a simple convex programming relaxation is amazingly effective. With a good choice for $A$, we will be able to recover a $k$-sparse signal from $O(k \log(n/k))$ linear measurements. The algorithm runs in polynomial time and, moreover, it is robust to noise in the sense that even if $x$ is not $k$-sparse, we will still be able to approximately recover its $k$ largest coordinates. This is a much stronger type of guarantee. After all, natural signals aren't $k$-sparse. But being able to recover their $k$ largest coordinates is often good enough.

Now let's define the restricted isometry property:

**Definition 5.5.1** *A matrix $A$ satisfies the $(k, \delta)$-restricted isometry property if for all $k$-sparse vectors $x$ we have*

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2.$$

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

As with the other assumptions we have considered, the restricted isometry property holds on randomly chosen sensing matrices with high probability:

**Lemma 5.5.2** *Let $A$ be an $m \times n$ matrix where each entry is an independent standard Gaussian random variable. Provided that $m \geq 10k \log n/k$, then with high probability $A$ satisfies the $(k, 1/3)$-restricted isometry property.*

Next let's work toward formalizing what we mean by approximately recovering the $k$ largest coordinates of $x$. Our goal will be formulated in terms of the following function:

**Definition 5.5.3** $\gamma_k(x) = \min_{w \text{ s.t. } \|w\|_0 \leq k} \|x - w\|_1$

To put this in more plain terms, $\gamma_k(x)$ is the sum of the absolute values of all but the $k$ largest magnitude entries of $x$. And if $x$ really is $k$-sparse, then $\gamma_k(x) = 0$.

Our goal is to find a $w$ that approximates $x$ almost as well as any $k$-sparse vector does. More formally, we want to find a $w$ that satisfies $\|x - w\|_1 \leq C\gamma_k(x)$, and we want to do so using as few linear measurements as possible. This learning goal already subsumes our other exact recovery results from previous sections, because when $x$ is $k$-sparse, then, as we discussed, $\gamma_k(x)$ is zero, so we have no choice but to recover $w = x$.

In this section, our approach will be based on a convex programming relaxation. Instead of trying to solve the *NP*-hard optimization problem $(P_0)$, we will consider the now famous $\ell_1$-relaxation:

$$(P_1) \qquad \min \|w\|_1 \text{ s.t. } Aw = b$$

Let's first state some of the well-known results about using $(P_1)$ for sparse recovery:

**Theorem 5.5.4** *[43] If $\delta_{2k} + \delta_{3k} < 1$, then if $\|x\|_0 \leq k$, we have $w = x$.*

**Theorem 5.5.5** *[42] If $\delta_{3k} + 3\delta_{4k} < 2$, then*

$$\|x - w\|_2 \leq \frac{C}{\sqrt{k}}\gamma_k(x).$$

The guarantees above are a bit different (and often stronger) than the others, because the bound is in terms of the $\ell_2$ norm of the error $x - w$.

**Theorem 5.5.6** *[51] If $\delta_{2k} < 1/3$, then*

$$\|x - w\|_1 \leq \frac{2 + 2\delta_{2k}}{1 - 3\delta_{2k}}\gamma_k(x).$$

We won't prove exactly these results. But we will prove something similar following the approach of Kashin and Temlyakov [96], which (to my taste)

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

greatly streamlines these analyses. But before we get to analyzing $(P_1)$, we need to introduce a notion from functional analysis called an almost Euclidean subsection.

## Almost Euclidean Subsections

Informally, an almost Euclidean subsection is a subspace where the $\ell_1$ and $\ell_2$ norms are almost equivalent after rescaling. We will just assert the fact that a random subspace is an almost Euclidean subsection with high probability. Instead, we will spend most of our time establishing various geometric properties about Euclidean subsections that we will use when we return to compressed sensing. The crucial definition is the following:

**Definition 5.5.7** *A subspace $\Gamma \subseteq \mathbb{R}^n$ is a $C$-almost Euclidean subsection if for all $v \in \Gamma$,*

$$\frac{1}{\sqrt{n}}\|v\|_1 \leq \|v\|_2 \leq \frac{C}{\sqrt{n}}\|v\|_1.$$

Actually, the first inequality is trivial. For any vector, it's always true that $\frac{1}{\sqrt{n}}\|v\|_1 \leq \|v\|_2$. The action is all happening in the second inequality. The first time you see them, it's not obvious that such subspaces exist. Indeed, Garnaev and Gluskin [75] proved that there are plenty of almost Euclidean subsections:

**Theorem 5.5.8** *If $\Gamma$ is a subspace chosen uniformly at random with $dim(\Gamma) = n - m$, then for*

$$C = O\left(\sqrt{\frac{n}{m}\log\frac{n}{m}}\right)$$

*we have that $\Gamma$ will be a $C$-almost Euclidean subsection with high probability.*

Let's end with a nice picture to keep in mind. Consider the unit ball for the $\ell_1$ norm. It's sometimes called the cross polytope, and to visualize it you can think of it as the convex hull of the vectors $\{\pm e_i\}_i$ where $e_i$ are the standard basis vectors. Then a subspace $\Gamma$ is almost Euclidean if, when we intersect it and the cross polytope, we get a convex body that is almost spherical.

## Geometric Properties of $\Gamma$

Here we will establish some important geometric properties of $C$-almost Euclidean subsections. Throughout this section, let $S = n/C^2$. First we show that $\Gamma$ cannot contain any sparse, nonzero vectors:

**Claim 5.5.9** *Let $v \in \Gamma$, then either $v = 0$ or $|supp(v)| \geq S$.*

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

*Proof:* From Cauchy-Schwartz and the $C$-almost Euclidean property, we have

$$\|v\|_1 = \sum_{j \in \text{supp}(v)} |v_j| \leq \sqrt{|\text{supp}(v)|} \cdot \|v\|_2 \leq \sqrt{|\text{supp}(v)|} \frac{C}{\sqrt{n}} \|v\|_1.$$

The proof now follows from rearranging terms. ■

It's worth noting that there is a nice analogy with linear error correcting codes, which are also subspaces of large dimension (but over $GF_2$), where we want every nonzero vector to have at least a constant fraction of its coordinates be nonzero. In any case, let's move on to some even stronger properties of almost Euclidean subsections, which have to do with how well the $\ell_1$ norm is spread out. First let's give a useful piece of notation:

**Definition 5.5.10** *For $\Lambda \subseteq [n]$, let $v_\Lambda$ denote the restriction of $v$ to coordinates in $\Lambda$. Similarly, let $v^\Lambda$ denote the restriction of $v$ to $\bar{\Lambda}$.*

With this notation in hand, let's prove the following:

**Claim 5.5.11** *Suppose $v \in \Gamma$ and $\Lambda \subseteq [n]$ and $|\Lambda| < S/16$. Then*

$$\|v_\Lambda\|_1 < \frac{\|v\|_1}{4}.$$

*Proof:* The proof is almost identical to that of Claim 5.5.9. Again using Cauchy-Schwartz and the $C$-almost Euclidean property, we have

$$\|v_\Lambda\|_1 \leq \sqrt{|\Lambda|} \|v_\Lambda\|_2 \leq \sqrt{|\Lambda|} \|v\|_2 \leq \sqrt{|\Lambda|} \frac{C}{\sqrt{n}} \|v\|_1,$$

which, plugging in terms, completes the proof. ■

And now we have all the tools we need to give our first results about $(P_1)$:

**Lemma 5.5.12** *Let $w = x + v$ and $v \in \Gamma$, where $\|x\|_0 \leq S/16$. Then $\|w\|_1 > \|x\|_1$.*

*Proof:* Set $\Lambda = \text{supp}(x)$. Then

$$\|w\|_1 = \|(x+v)_\Lambda\|_1 + \|(x+v)^\Lambda\|_1 = \|(x+v)_\Lambda\|_1 + \|v^\Lambda\|_1.$$

Now we can invoke triangle inequality:

$$\|w\|_1 \geq \|x_\Lambda\|_1 - \|v_\Lambda\|_1 + \|v^\Lambda\|_1 = \|x\|_1 - \|v_\Lambda\|_1 + \|v^\Lambda\|_1 = x_\Lambda\|_1 - 2\|v_\Lambda\|_1 + \|v\|_1$$

However, $\|v\|_1 - 2\|v_\Lambda\|_1 \geq \|v\|_1/2 > 0$ using Claim 5.5.11. This completes the proof. ■

Plugging in the bounds from Theorem 5.5.8, we have shown that we can recover a $k$-sparse vector $x$ of dimension $n$ with

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

$$k \leq S/16 = \Omega(n/C^2) = \Omega\Big(\frac{m}{\log n/m}\Big)$$

from $m$ linear measurements.

Next we will consider stable recovery. Our main theorem is:

**Theorem 5.5.13** *Let $\Gamma = ker(A)$ be a $C$-almost Euclidean subsection. Let $S = \frac{n}{C^2}$. If $Ax = Aw = b$ and $\|w\|_1 \leq \|x\|_1$, we have*

$$\|x - w\|_1 \leq 4\,\sigma_{\frac{S}{16}}(x)\,.$$

*Proof:* Let $\Lambda \subseteq [n]$ be the set of $S/16$ coordinates maximizing $\|x_\Lambda\|_1$. We want to upper-bound $\|x - w\|_1$. By repeated application of the triangle inequality, $\|w\|_1 = \|w^\Lambda\|_1 + \|w_\Lambda\|_1 \leq \|x\|_1$, and the definition of $\sigma_t(\cdot)$, it follows that

$$\begin{aligned}
\|x - w\|_1 &= \|(x-w)_\Lambda\|_1 + \|(x-w)^\Lambda\|_1 \\
&\leq \|(x-w)_\Lambda\|_1 + \|x^\Lambda\|_1 + \|w^\Lambda\|_1 \\
&\leq \|(x-w)_\Lambda\|_1 + \|x^\Lambda\|_1 + \|x\|_1 - \|w_\Lambda\|_1 \\
&\leq 2\|(x-w)_\Lambda\|_1 + 2\|x^\Lambda\|_1 \\
&\leq 2\|(x-w)_\Lambda\|_1 + 2\,\sigma_{\frac{S}{16}}(x)\,.
\end{aligned}$$

Since $(x-w) \in \Gamma$, we can apply Claim 5.5.11 to conclude that $\|(x-w)_\Lambda\|_1 \leq \frac{1}{4}\|x-w\|_1$. Hence

$$\|x - w\|_1 \leq \frac{1}{2}\|x-w\|_1 + 2\sigma_{\frac{S}{16}}(x)\,.$$

This completes the proof. ■

## Epilogue

Finally, we will end with one of the main open questions in compressed sensing, which is to give a *deterministic* construction of matrices that satisfy the restricted isometry property:

**Question 7 (Open)** *Is there a deterministic algorithm to construct a matrix with the restricted isometry property? Alternatively, is there a deterministic algorithm to construct an almost Euclidean subsection $\Gamma$?*

Avi Wigderson likes to refer to these types of problems as "finding hay in a haystack." We know that a randomly chosen $A$ satisfies the restricted isometry property with high probability. Its kernel is also an almost Euclidean subspace with high probability. But can we remove the randomness? The best known deterministic construction is due to Guruswami, Lee, and Razborov [82]:

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press

**Theorem 5.5.14** *[82] There is a polynomial time deterministic algorithm for constructing an almost Euclidean subspace* $\Gamma$ *with parameter* $C \sim (\log n)^{\log\log\log n}$.

This has got to be too strange a bound to be the best we can do, right?

## 5.6 Exercises

**Problem 5-1:** In this question, we will explore uniqueness conditions for sparse recovery and conditions under which $\ell_1$-minimization provably works.

(a) Let $A\widehat{x} = b$, and suppose $A$ has $n$ columns. Further suppose $2k \leq m$. Prove that for *every* $\widehat{x}$ with $\|\widehat{x}\|_0 \leq k$, $\widehat{x}$ is the uniquely sparsest solution to the linear system if and only if the k-rank of the columns of $A$ is at least $2k$.

(b) Let $U = \text{kernel}(A)$, and $U \subset \mathbb{R}^n$. Suppose that for each nonzero $x \in U$, and for any set $S \subset [n]$ with $|S| \leq k$,

$$\|x_S\|_1 < \frac{1}{2}\|x\|_1$$

where $x_S$ denotes the restriction of $x$ to the coordinates in $S$. Prove that

$$(P1) \qquad \min \|x\|_1 \text{ s.t. } Ax = b$$

recovers $x = \widehat{x}$, provided that $A\widehat{x} = b$ and $\|\widehat{x}\|_0 \leq k$.

(c) **Challenge:** Can you construct a subspace $U \subset \mathbb{R}^n$ of dimension $\Omega(n)$ that has the property that every nonzero $x \in U$ has at least $\Omega(n)$ nonzero coordinates? *Hint:* Use an expander.

**Problem 5-2:** Let $\widehat{x}$ be a $k$-sparse vector in $n$-dimensions. Let $\omega$ be the $n$th root of unity. Suppose we are given $v_\ell = \sum_{j=1}^{n} \widehat{x}_j \omega^{\ell j}$ for $\ell = 0, 1, \ldots, 2k-1$. Let $A, B \in \mathbb{R}^{k\times k}$ be defined so that $A_{i,j} = v_{i+j-2}$ and $B_{i,j} = v_{i+j-1}$.

(a) Express both $A$ and $B$ in the form $A = VD_AV^T$ and $B = VD_BV^T$, where $V$ is a Vandermonde matrix and $D_A, D_B$ are diagonal.

(b) Prove that the solutions to the generalized eigenvalue problem $Ax = \lambda Bx$ can be used to recover the locations of the nonzeros in $\widehat{x}$.

(c) Given the locations of the nonzeros in $\widehat{x}$ and $v_0, v_1, \ldots, v_{k-1}$, give an algorithm to recover the values of the nonzero coefficients in $\widehat{x}$.

This is called the matrix pencil method. If you squint, it looks like Prony's method (Section 5.4) and has similar guarantees. Both are (somewhat) robust to noise if and only if the Vandermonde matrix is well-conditioned, and exactly when that happens is a longer story. See Moitra [113].

https://doi.org/10.1017/9781316882177.006 Published online by Cambridge University Press