ℓ_1 minimization

We will now focus on underdetermined systems of equations:

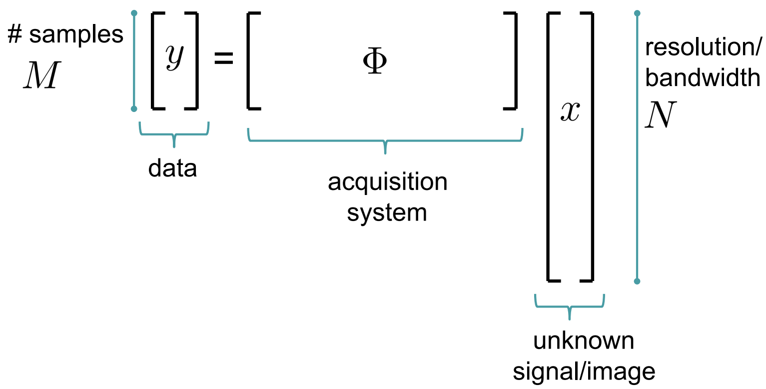

Suppose we observe $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}_0$, and given $\boldsymbol{y}$ we attempt to estimate $\boldsymbol{x}_0$ by applying the pseudo-inverse of $\boldsymbol{\Phi}$ to $\boldsymbol{y}$ — we are implicitly solving the program

$$\min_{\boldsymbol{x}} \ \|\boldsymbol{x}\|_2 \quad \text{subject to} \quad \boldsymbol{\Phi}\boldsymbol{x} = \boldsymbol{y}.$$

Of course, we will recover x_0 exactly only under very special circumstances, namely that $\boldsymbol{x}_0$ is in $\operatorname{Range}(\boldsymbol{\Phi}^{\mathrm{T}})$; if $\boldsymbol{x}_0 = \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\alpha}$ for some $\boldsymbol{\alpha} \in \mathbb{R}^M$, then

$$\begin{aligned}
\boldsymbol{\Phi}^{\mathrm{T}}(\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathrm{T}})^{-1}\boldsymbol{y} &= \boldsymbol{\Phi}^{\mathrm{T}}(\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathrm{T}})^{-1}\boldsymbol{\Phi}\boldsymbol{x}_0 \\
&= \boldsymbol{\Phi}^{\mathrm{T}}(\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathrm{T}})^{-1}\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\alpha} \\
&= \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\alpha} \\
&= \boldsymbol{x}_0.
\end{aligned}$$

So if $\boldsymbol{x}_0$ lies in a particular M-dimensional subspace of $\mathbb{R}^N$, it can be recovered from observations through the $M \times N$ matrix $\boldsymbol{\Phi}$.

If $\boldsymbol{x}_0$ is *sparse* (i.e. has a small number of non-zero terms at unknown locations), we might envision recovery via a different type of optimization program. We find the sparsest vector that explains $\boldsymbol{y}$ by solving

$$\min_{\boldsymbol{x}} \#\{i \,:\, x[i] \neq 0\} \quad \text{subject to} \quad \boldsymbol{\Phi x} = \boldsymbol{y},$$

where $\#\{i \,:\, x[i] \neq 0\}$ is the number of coordinates where $\boldsymbol{x}$ is non-zero.

The optimization program above is incredibly hard to solve directly. Basically, there is no better way to solve it than to see if there is a one sparse vector which matches the measurements $(O(MN))$, then test if $\boldsymbol{y}$ can be written as the superposition of two columns $(O(4MN^2))$, then .... As the cost of testing if $\boldsymbol{y}$ is the superposition of any S columns is $O(MS^2\binom{N}{S})$, this quickly gets out of control.

The ℓ_1 norm has long been used as a heuristic for promoting sparsity. In the past decade, there has been a very large body of work that characterizes its effectiveness quantitatively. Instead of the combinatorial program above, we solve the convex program

$$\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{subject to} \quad \boldsymbol{\Phi x} = \boldsymbol{y}. \tag{1}$$

Our goal today to formalize the conditions under which this recovery procedure is effective.

Necessary and sufficient conditions for ℓ_1 recovery

Let $\boldsymbol{x}_0$ be a vector supported on a set $\Gamma \subset \{1, 2, \ldots, N\}$, and let $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}_0$. It is clear that $\boldsymbol{x}_0$ is a solution to (1) if and only if

$$\|\boldsymbol{x}_0 + \boldsymbol{h}\|_1 \geq \|\boldsymbol{x}_0\|_1 \qquad \forall \boldsymbol{h} \text{ with } \boldsymbol{\Phi}\boldsymbol{h} = \boldsymbol{0}. \tag{2}$$

It is always true that

$$\begin{aligned} \|\boldsymbol{x}_0 + \boldsymbol{h}\|_1 - \|\boldsymbol{x}_0\|_1 &= \sum_{\gamma \in \Gamma} (|x_0[\gamma] + h[\gamma]| - |x_0[\gamma]|) + \sum_{\gamma \in \Gamma^c} |h[\gamma]| \\ &\geq \sum_{\gamma \in \Gamma} \operatorname{sign}(x_0[\gamma]) h[\gamma] + \sum_{\gamma \in \Gamma^c} |h[\gamma]| \end{aligned}$$

since

$$|a + b| - |a| \geq \operatorname{sign}(a) b \quad \text{for all } a, b \in \mathbb{R}.$$

Thus (2) holds when

$$-\sum_{\gamma \in \Gamma} \operatorname{sign}(x_0[\gamma]) h[\gamma] \leq \sum_{\gamma \in \Gamma^c} |h[\gamma]| \quad \text{for all} \quad \boldsymbol{h} \in \operatorname{Null}(\boldsymbol{\Phi}). \tag{3}$$

This condition is also necessary, since if there is an $\boldsymbol{h} \in \operatorname{Null}(\boldsymbol{\Phi})$ with

$$\sum_{\gamma \in \Gamma^c} |h[\gamma]| < -\sum_{\gamma \in \Gamma} \operatorname{sign}(x_0[\gamma]) h[\gamma],$$

then the same is true for $\epsilon\boldsymbol{h}$ for an $\epsilon > 0$. Since $\epsilon\boldsymbol{h} \in \operatorname{Null}(\boldsymbol{\Phi})$, $\boldsymbol{\Phi}(\boldsymbol{x}_0 + \epsilon\boldsymbol{h}) = \boldsymbol{y}$ and

$$\begin{aligned} \|\boldsymbol{x}_0 + \epsilon\boldsymbol{h}\|_1 &= \sum_{\gamma \in \Gamma} |x_0[\gamma] + \epsilon h[\gamma]| + \sum_{\gamma \in \Gamma^c} |\epsilon h[\gamma]| \\ &< \sum_{\gamma \in \Gamma} |x_0[\gamma] + \epsilon h[\gamma]| - \epsilon \operatorname{sign}(x_0[\gamma]) h[\gamma] \\ &\leq \sum_{\gamma \in \Gamma} |x_0[\gamma]| \quad \text{for some small enough } \epsilon > 0 \\ &= \|\boldsymbol{x}_0\|_1. \end{aligned}$$

So this would imply that there is another vector (namely $\boldsymbol{x}_0 + \epsilon \boldsymbol{h}$) that has smaller ℓ_1 norm and is also feasible.

It is interesting to note that given the observation matrix $\mathbf{\Phi}$, our ability to recover a vector $\boldsymbol{x}_0$ is determined only by

1. the set Γ on which $\boldsymbol{x}_0$ is supported (the locations of its "active elements"), and
2. the signs of the elements on this set.

The magnitudes of the entries are not involved at all. Geometrically, the support set coupled with the sign sequence specifies the *facet* of the ℓ_1 ball on which $\boldsymbol{x}_0$ lives:

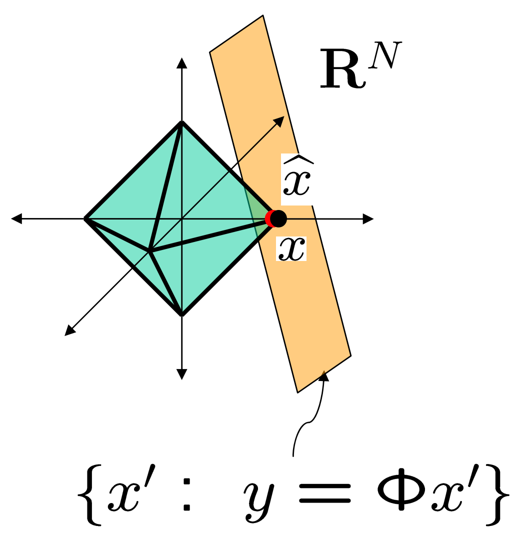

Duality and optimality

We can get a more workable sufficient condition for exact recovery of $\boldsymbol{x}_0$ by looking at the dual of the convex program

$$\min_{\boldsymbol{x}} \ \|\boldsymbol{x}\|_1 \quad \text{subject to} \quad \boldsymbol{\Phi}\boldsymbol{x} = \boldsymbol{y}. \tag{4}$$

Let's start by considering a general optimization program with linear equality constraints:

$$\min_{\boldsymbol{x}} \ f(\boldsymbol{x}) \quad \text{subject to} \quad \boldsymbol{\Phi}\boldsymbol{x} = \boldsymbol{y}. \tag{5}$$

The *Lagrangian* for this problem is

$$L(\boldsymbol{x},\boldsymbol{\nu}) = f(\boldsymbol{x}) + \boldsymbol{\nu}^{\mathrm{T}}(\boldsymbol{\Phi}\boldsymbol{x} - \boldsymbol{y})$$

If f is differentiable, then $\boldsymbol{x}$ is a solution to (5) if it is feasible, $\boldsymbol{\Phi}\boldsymbol{x} = \boldsymbol{y}$, and there exists a $\boldsymbol{\nu} \in \mathbb{R}^M$ such that

$$\nabla_x L(\boldsymbol{x},\boldsymbol{\nu}) = \nabla_x f(\boldsymbol{x}) + \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\nu} = \mathbf{0}.$$

We are interested in the particular functional $f(\boldsymbol{x}) = \|\boldsymbol{x}\|_1$, which is not differentiable but is convex. Fortunately, there is an easy modification to the condition above in this case — $\boldsymbol{x}$ is a solution to (5) if it is feasible and there exists a $\boldsymbol{\nu} \in \mathbb{R}^M$ such that

$$\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\nu} \text{ is a "subgradient" of } f \text{ at } \boldsymbol{x}.$$

A vector $\boldsymbol{u} \in \mathbb{R}^N$ is a subgradient of a function f at $\boldsymbol{x}_0$ if it is normal to a "supporting hyperplane"; that is if

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}_0) + \langle \boldsymbol{x} - \boldsymbol{x}_0, \boldsymbol{u} \rangle.$$

If f is differentiable at $\boldsymbol{x}_0$, then $\nabla f(\boldsymbol{x}_0)$ is the only subgradient.

To make this concept more concrete, consider the (very relevant) one dimensional example $f(t) = |t|$.

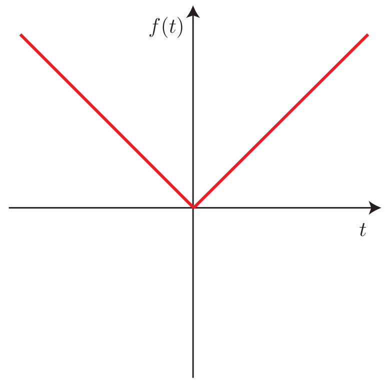

In this case, the set of subgradients is simply the derivative (± 1) for t away from the origin, and all slopes between -1 and $+1$ at the origin:

$$\{\text{subgradients}\} = \begin{cases} \text{sign}(t) & t \neq 0 \\ [-1, 1] & t = 0 \end{cases}.$$

In $\mathbb{R}^N$, for $f(x) = \|\boldsymbol{x}\|_1$, the subgradients of f at $\boldsymbol{x}_0$ are the collection of vectors u such that

$$\begin{aligned} u[\gamma] &= \text{sign}(x_0[\gamma]), \quad & \gamma \in \Gamma \\ |u[\gamma]| &\leq 1, & \gamma \in \Gamma^c, \end{aligned}$$

where Γ is the support set of $\boldsymbol{x}_0$.

Given an $M \times N$ matrix $\boldsymbol{\Phi}$ and a vector $\boldsymbol{x}_0 \in \mathbb{R}^N$ supported on Γ, set $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}_0$. Then $\boldsymbol{x}_0$ is a solution to (4) if there exits a $\boldsymbol{\nu} \in \mathbb{R}^M$ such that

$$\begin{aligned} (\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\nu})[\gamma] &= \operatorname{sign}(x_0[\gamma]), \quad && \gamma \in \Gamma \\ |(\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\nu})[\gamma]| &\leq 1, && \gamma \in \Gamma^c. \end{aligned}$$

In addition, if $\boldsymbol{\Phi}$ is injective on the set of all vectors supported on Γ and we can find a $\boldsymbol{\nu}$ such that the second condition is

$$|(\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\nu})[\gamma]| < 1, \qquad \gamma \in \Gamma^c$$

then $\boldsymbol{x}_0$ is the unique solution.

Choosing a particular dual vector

Our recovery conditions boxed above simply ask that we be able to find one such v (there could be many). Many of the results in the fields of sparse recovery and compressed sensing have narrowed down the condition down by simply testing a prescribed vector. Let

$\boldsymbol{\Phi}_\Gamma =$ the $M \times |\Gamma|$ submatrix containing the columns of $\boldsymbol{\Phi}$ indexed by Γ

and let

$$\boldsymbol{z}_0 \in \mathbb{R}^{|\Gamma|} \text{ contain the signs of } \boldsymbol{x}_0 \text{ on } \Gamma.$$

We set

$$\boldsymbol{u}_0 = \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}_\Gamma(\boldsymbol{\Phi}_\Gamma^{\mathrm{T}}\boldsymbol{\Phi}_\Gamma)^{-1}\boldsymbol{z}_0. \tag{6}$$

Now sufficient conditions for $\boldsymbol{x}_0$ to be the unique minimizer of (4) are

1. $\boldsymbol{\Phi}_\Gamma^{\mathrm{T}}\boldsymbol{\Phi}_\Gamma$ is invertible. If this is true, then the expression above for $\boldsymbol{u}_0$ is well-behaved, and b construction we will have $u_0[\gamma] = \operatorname{sign}(x_0[\gamma])$;
2. If 1) holds, then we need $|u_0[\gamma]| < 1$.

Example: Dictionaries with bounded coherence

Now suppose that $\boldsymbol{\Phi}$ is an $M \times N$ matrix with normalized columns[1],

$$\|\boldsymbol{\Phi}_\gamma\|_2 = 1, \quad \gamma = 1, 2, \ldots, N$$

and *coherence*

$$\mu = \max_{\substack{1\leq\gamma_1,\gamma_2\leq N \\ \gamma_1\neq\gamma_2}} |\langle \boldsymbol{\Phi}_{\gamma_1}, \boldsymbol{\Phi}_{\gamma_2}\rangle|.$$

The quantity μ is essentially a measure of how closely aligned any two columns of $\boldsymbol{\Phi}$ are.

Let Γ be a fixed subset of $\{1, 2, \ldots, N\}$ of size $|\Gamma| = S$, and let $\boldsymbol{x}_0$ be a vector supported on Γ with sign sequence $\boldsymbol{z}_0$ (so $\boldsymbol{x}_0$ is S-*sparse*). For the first optimality condition note that we can write $\boldsymbol{\Phi}_\Gamma^{\mathrm{T}}\boldsymbol{\Phi}_\Gamma$ as

$$\boldsymbol{\Phi}_\Gamma^{\mathrm{T}}\boldsymbol{\Phi}_\Gamma = \mathbf{I} + \boldsymbol{G}$$

where each entry of the $S \times S$ matrix $\boldsymbol{G}$ is less than or equal to μ. We have

$$\begin{aligned} \|\boldsymbol{G}\| &\leq \|\boldsymbol{G}\|_F \\ &= \sqrt{\sum_{j=1}^{S}\sum_{k=1}^{S} |G[j,k]|^2} \\ &\leq \mu S, \end{aligned}$$

[1]It should be clear that we are using $\boldsymbol{\Phi}_i$ to denote the ith column of $\boldsymbol{\Phi}$ here.

and so we can guarantee $\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\Gamma$ is invertible when

$$S \;\leq\; \frac{1}{\mu}.$$

To check the second condition, we have for $\gamma \in \Gamma^c$

$$\begin{aligned} |u_0[\gamma]| &\leq \|(\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\Gamma)^{-1}\| \, \|\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\gamma\|_2 \, \|\boldsymbol{z}_0\|_2 \\ &\leq \frac{1}{1-\mu S}\|\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\gamma\|_2\sqrt{S}. \end{aligned}$$

Since $\gamma \notin \Gamma$, all the entries of $\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\gamma$ will have magnitude less than μ, and so $\|\mathbf{\Phi}_\Gamma^{\mathrm{T}}\mathbf{\Phi}_\gamma\|_2 \leq \mu\sqrt{S}$, and

$$|u_0[\gamma]| \leq \frac{\mu S}{1-\mu S}.$$

Thus

$$S \;<\; \frac{1}{2\mu}$$

ensures that $\boldsymbol{x}_0$ will be recoverable from observations $\boldsymbol{y} = \mathbf{\Phi}\boldsymbol{x}_0$.

It is interesting to note here that typically, the coherence μ cannot be much smaller than $1/\sqrt{M}$. In fact, if the columns of $\mathbf{\Phi}$ are drawn independently at random on the unit sphere in $\mathbb{R}^M$, then with very high probability $\mu \approx 1/\sqrt{M}$ to within a logarithmic factor. So the result above tells us that "random" underdetermined systems of equations can be solved if the solution is sparse with

$$S \lesssim M^2.$$

Using more subtle arguments than the ones above, we can even argue that we can take $S \sim M$ (to within a logarithmic factor). See, for example, [**?**].

Note on complex vectors: Almost everything we have said so far about ℓ_1 minimization extends in a straightforward manner to complex-valued vectors. First, it is worth mentioning that if $\boldsymbol{x} \in \mathbb{C}^N$, then

$$\|\boldsymbol{x}\|_1 = \sum_{n=1}^{N} |x[n]| = \sum_{n=1}^{N} \sqrt{\operatorname{Re}\{x[n]\}^2 + \operatorname{Im}\{x[n]\}^2}.$$

In this case, the ℓ_1 minimization program can no longer be re-cast as a linear program, but rather is what is called a "sum of norms" program (which is a particular type of "second order cone program"). This type of problem, however, is not too much more difficult to solve from a practical perspective.

The sufficient conditions for recovery are the same, but now we take $\operatorname{sign}(z)$ to be the *phase* of the complex number z. That is, if $z = Ae^{j\theta}$, then $\operatorname{sign}(z) = \theta$.