Lecture: Introduction to Compressed Sensing
Sparse Recovery Guarantees

http://bicmr.pku.edu.cn/~wenzw/bigdata2021.html

Acknowledgement: this slides is based on Prof. Emmanuel Candes' and Prof. Wotao Yin's lecture notes

Underdetermined systems of linear equations

- $x \in \mathbb{R}^n, A \in \mathbb{R}^{m \times n}, b \in \mathbb{R}^m$

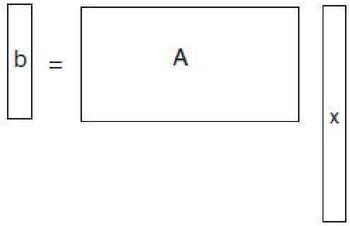

When fewer equations than unknowns

- Fundamental theorem of algebra says that we cannot find x
- In general, this is absolutely correct

Special structure

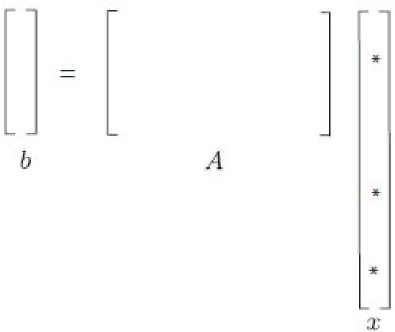

If unknown is assumed to be

- sparse
- low-rank

then one can *often* find solutions to these problems by convex optimization

Compressive Sensing

`http://bicmr.pku.edu.cn/~wenzw/courses/sparse_l1_example.m`

Find the sparest solution

- Given n=256, m=128.
- A = randn(m,n); u = sprandn(n, 1, 0.1); b = A*u;

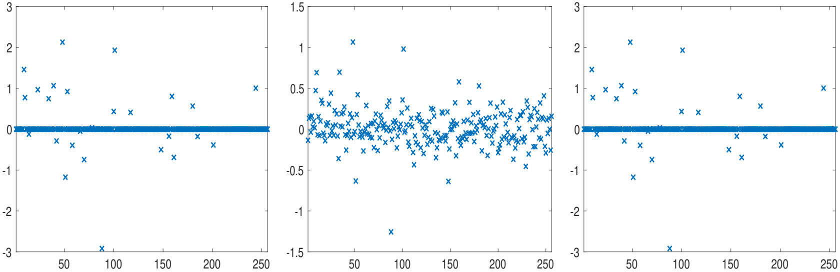

$$\begin{cases} \min_x \|x\|_0 \\ \text{s.t. } Ax = b \end{cases}$$

(a) ℓ_0-minimization

$$\begin{cases} \min_x \|x\|_2 \\ \text{s.t. } Ax = b \end{cases}$$

(b) ℓ_2-minimization

$$\begin{cases} \min_x \|x\|_1 \\ \text{s.t. } Ax = b \end{cases}$$

(c) ℓ_1-minimization

Linear programming formulation

ℓ_0 minimization

$$\begin{array}{ll} \min & \|x\|_0 \\ s.t. & Ax = b \end{array}$$

Combinatorially hard

ℓ_1 minimization

$$\begin{array}{ll} \min & \|x\|_1 \\ s.t. & Ax = b \end{array}$$

Linear program

$$\begin{array}{ll} \text{minimize} & \sum_i |x_i| \\ \text{subject to} & Ax = b \end{array}$$

is equivalent to

$$\begin{array}{ll} \text{minimize} & \sum_i t_i \\ \text{subject to} & Ax = b \\ & -t_i \leq x_i \leq t_i \end{array}$$

with variables $x, t \in \mathbb{R}^n$

$x^\star$ is a solution $\iff$ $(x^\star, t^\star = |x^\star|)$ is a solution

Compressed sensing

- Name coined by David Donoho
- Has become a label for sparse signal recovery
- But really one instance of underdetermined problems

- Informs analysis of underdetermined problems
- Changes viewpoint about underdetermined problems
- Starting point of a general burst of activity in
  - information theory
  - signal processing
  - statistics
  - some areas of computer science
  - ...
- Inspired new areas of research, e. g. low-rank matrix recovery

A contemporary paradox

- Massive data acquisition
- Most of the data is redundant and can be thrown away
- Seems enormously wasteful

Raw: 15MB

JPEG: 150KB

Sparsity in signal processing

1 megapixel image

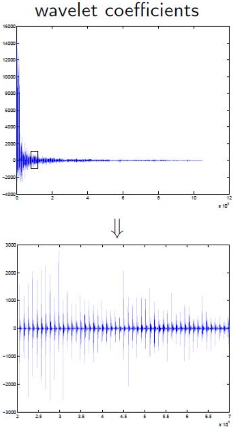

zoom in

Implication: can discard small coefficients without much perceptual loss

Sparsity and wavelet "compression"

Take a mega-pixel image

- Compute $1,000,000$ wavelet coefficients
- Set to zero all but the $25,000$ largest coefficients
- Invert the wavelet transform

1 megapixel image

25k term approximation

This principle underlies modern lossy coders

Comparison

Sparse representation = good compression

Why? Because there are fewer things to send/store

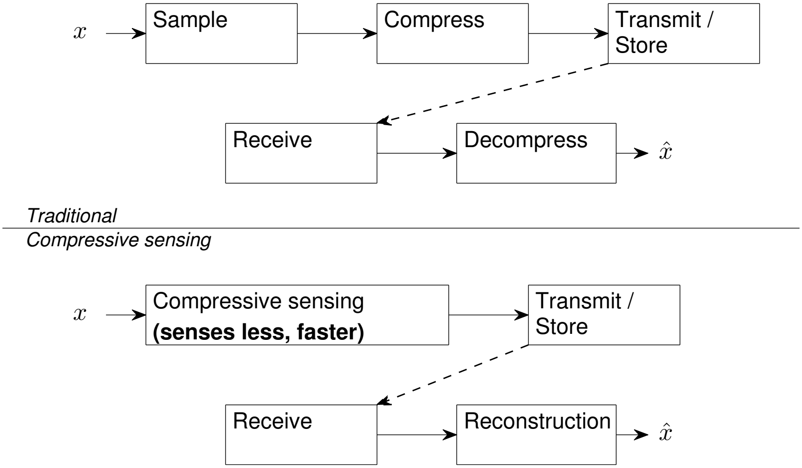

Fundamental Question

The basic question of sparse optimization is:

Can I trust my model to return an intended sparse quantity?

That is

- does my model have a unique solution? (otherwise, different algorithms may return different answers)
- is the solution exactly equal to the original sparse quantity?
- if not (due to noise), is the solution a faithful approximate of it?
- how much effort is needed to numerically solve the model?

How to read guarantees

Some basic aspects that distinguish different types of guarantees:

- Recoverability (exact) vs stability (inexact)
- General A or special A?
- Universal (all sparse vectors) or instance (certain sparse vector(s))?
- General optimality? or specific to model/algorithm?
- Required property of A: spark, RIP, coherence, NSP, dual certificate?
- If randomness is involved, what is its role?
- Condition/bound is tight or not? Absolute or in order of magnitude?

Restricted isometries: C. and Tao (04)

Definition (Restricted isometry constants)

For each $k = 1, 2, \ldots, \delta_k$ is the smallest scalar such that

$$(1 - \delta_k)\|x\|_2^2 \le \|Ax\|_2^2 \le (1 + \delta_k)\|x\|_2^2$$

for all k-sparse x

- Note slight change of normalization
- When δ_k is not too large, condition says that all $m \times k$ submatrices are well conditioned (sparse subsets of columns are not too far from orthonormal)

Interlude: when does sparse recovery make sense?

- x is s-sparse: $\|x\|_0 \leq s$
- **can we recover x from $Ax = b$?**

Perhaps possible if sparse vectors lie away from null space of A

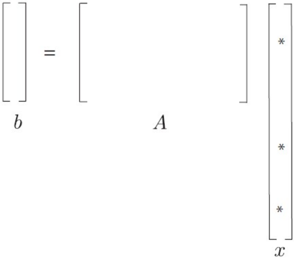

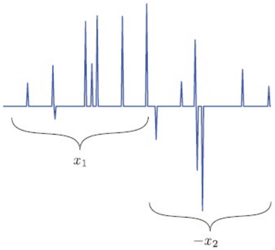

Yes if any $2s$ columns of A are linearly independent

If x_1, x_2 s-sparse such that $Ax_1 = Ax_2 = b$

$A(x_1 - x_2) = 0 \Rightarrow x_1 - x_2 = 0 \Leftrightarrow x_1 = x_2$

Interlude: when does sparse recovery make sense?

- x is s-sparse: $\|x\|_0 \leq s$
- **can we recover x from $Ax = b$?**

Perhaps possible if sparse vectors lie away from null space of A

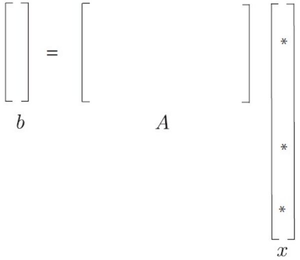

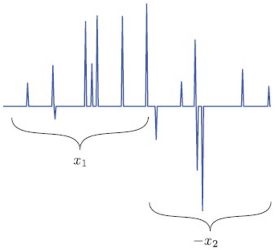

In general, **No** if A has $2s$ linearly dependent columns

$h \neq 0$ is $2s$-sparse with $Ah = 0$

$h = x_1 - x_2 \quad x_1, x_2$ both s-sparse

$Ah = 0 \Leftrightarrow Ax_1 = Ax_2$ and $x_1 \neq x_2$

Equivalent view of restricted isometry property

δ_{2k} is the smallest scalar such that

$$(1-\delta_{2k})\|x_1-x_2\|_2^2 \leq \|Ax_1-Ax_2\|_2^2 \leq (1+\delta_{2k})\|x_1-x_2\|_2^2$$

for all k-sparse vectors x_1, x_2.

The positive lower bounds is that which really matters

- If lower bound does not hold, then we may have x_1 and x_2 both sparse and with disjoint supports, obeying

$$Ax_1 = Ax_2$$

- Lower bound guarantees that distinct sparse signals cannot be mapped too closely (analogy with codes)

With a picture

For all k-sparse x_1 and x_2

$$1 - \delta_{2k} \leq \frac{\|Ax_1 - Ax_2\|_2^2}{\|x_1 - x_2\|_2^2} \leq 1 + \delta_{2k}$$

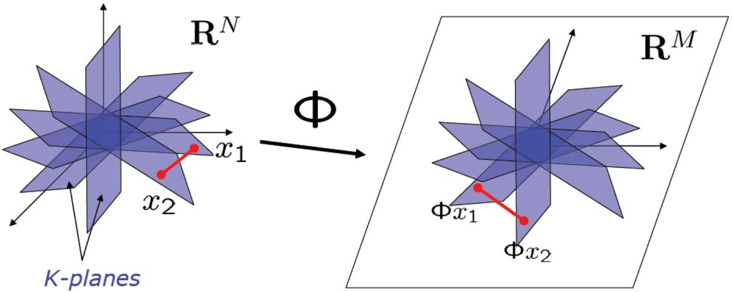

Picture from M. Wakin (Φ is our A)

Characterization of ℓ_1 solutions

Underdetermined system: $A \in \mathbb{R}^{m \times n}, m < n$

$$\min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{s.t. } Ax = b$$

x is solution iff

$$\|x+h\|_1 \geq \|x\|_1 \quad \forall h \in \mathbb{R}^n \ \text{ s.t. } Ah = 0$$

Notations: x supported on $T = \{i : x_i \neq 0\}$

$$\begin{aligned} \|x+h\|_1 &= \sum_{i \in T} |x_i + h_i| + \sum_{i \in T^c} |h_i| \\ &\geq \sum_{i \in T} |x_i| + \sum_{i \in T} \text{sgn}(x_i) h_i + \sum_{i \in T^c} |h_i| \end{aligned}$$

because $|x_i + h_i| \geq |x_i| + \text{sgn}(x_i) h_i$

Necessay and sufficient condition for ℓ_1 recovery

For all $h \in \text{null}(A)$

$$\sum_{i \in T} \text{sgn}(x_i) h_i \leq \sum_{i \in T^c} |h_i|$$

Why is this necessary? If there is $h \in \text{null}(A)$ with

$$\sum_{i \in T} \text{sgn}(x_i) h_i > \sum_{i \in T^c} |h_i|$$

then

$$\|x - h\|_1 < \|x\|_1.$$

Proof: There exists a small enough t such that

$$|x_i - th_i| = \begin{cases} x_i - th_i = x_i - t\text{sgn}(x_i)h_i & \text{if } x_i > 0 \\ -(x_i - th_i) = -x_i - t\text{sgn}(x_i)h_i & \text{if } x_i < 0 \\ t|h_i| & \text{otherwise} \end{cases}$$

Then

$$\|x - th\|_1 = \|x\|_1 - t \sum_{i \in T} \text{sgn}(x_i) h_i + t \sum_{i \in T^c} |h_i| < \|x\|_1$$

Characterization via KKT conditions

$$\min_{x\in\mathbb{R}^n} f(x) \quad \text{s.t.} \quad Ax = b$$

- f convex and differentiable Lagrangian
- $\mathcal{L}(x,\lambda) = f(x) + \langle \lambda, b - Ax \rangle$

$Ax = 0$ if and only if x is orthogonal to each of the row vectors of A.

KKT condition

x is solution iff x is feasible and $\exists \lambda \in \mathbb{R}^m$ s.t.

$$\nabla_x \mathcal{L}(x,\lambda) = 0 = \nabla f(x) - A^\top \lambda$$

Geometric interpretation: $\nabla f(x) \perp \text{null}(A)$.
When f is not differentiable, condition becomes: x feasible and $\exists \lambda \in \mathbb{R}^m$ s.t.

$$A^\top \lambda \text{ is a subgradient of } f \text{ at } x$$

Subgradient

Definition

u is a subgradient of convex f at x_0 if for all x

$$f(x) \geq f(x_0) + u \cdot (x - x_0)$$

if f is differentiable at x_0, the only subgradient is $\nabla f(x_0)$

Subgradients of $f(t) = |t|, t \in \mathbb{R}$

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

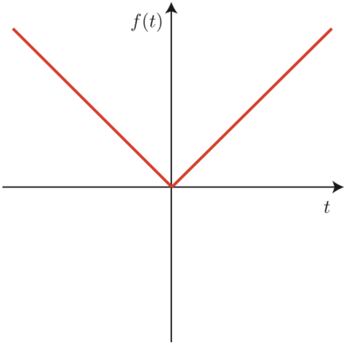

Subgradients of $f(x) = \|x\|_1, x \in \mathbb{R}^n$:
$u \in \partial \|x\|_1$ (u is a subgradient) iff

$$\begin{cases} u_i = \text{sgn}(x_i) & x_i \neq 0 \\ |u_i| \leq 1 & x_i = 0 \end{cases}$$

Optimality conditions II

Given $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.

$$(P) \qquad \min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{s.t.} \quad Ax = b$$

The dual problem is

$$\max_y \quad y^\top b, \qquad \text{s.t. } \|A^\top y\|_\infty \leq 1$$

x optimal solution iff x is feasible and there exists $u = A^\top \lambda (u \perp \text{null}(A))$ with

$$\begin{cases} u_i = \text{sgn}(x_i) & x_i \neq 0 \ (i \in T) \\ |u_i| \leq 1 & x_i = 0 \ (i \in T^c) \end{cases}$$

If in addition

- $|u_i| < 1$ when $x_i = 0$
- A_T has full column rank (implies by RIP)

Then x is the unique solution. We will call such a u or λ a dual certificate.

Uniqueness

Notation

- x_T: restriction of x to indices in T
- A_T: submatrix with column indices in T

If supp$(x) \subseteq T$,

$$Ax = A_T x_T.$$

Let $h \in \text{null}(A)$. Since $u \perp \text{null}(A)$, we have

$$\sum_{i \in T} \text{sgn}(x_i) h_i = \sum_{i \in T} u_i h_i = \langle u, h \rangle - \sum_{i \in T^c} u_i h_i = -\sum_{i \in T^c} u_i h_i < \sum_{i \in T^c} |h_i|$$

unless $h_{T^c} \neq 0$. Now if $h_{T^c} = 0$, then since A_T has full column rank,

$$Ah = A_T h_T = 0 \Rightarrow h_T = 0 \Rightarrow h = 0$$

In conclusion, for any $h \in \text{null}(A)$, $\|x + h\|_1 > \|x\|_1$ unless $h \neq 0$

Sufficient conditions

- $T = \text{supp}(x)$ and A_T has full column rank ($A_T^\top A_T$ invertible)
- $\text{sgn}(x_T)$ is the sign sequence of x on T and set

$$\lambda = A_T(A_T^\top A_T)^{-1}\text{sgn}(x_T) \text{ and } u := A^\top \lambda$$

- if $|u_i| \leq 1$ for all $i \in T^c$, then x is solution
- if $|u_i| < 1$ for all $i \in T^c$, then x is the unique solution

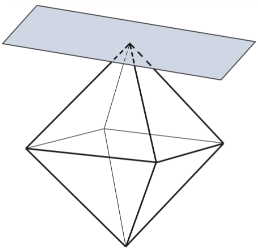

Why?

- $u_i = \text{sgn}(x_i)$ if $i \in T$, since

$$u_T = A_T^\top A_T(A_T^\top A_T)^{-1}\text{sgn}(x_T) = \text{sgn}(x_T)$$

- $u_i = A_i^\top \lambda$ if $i \notin T$.

So u is a valid dual certificate

RIP

- RIP: For each $k = 1, 2, \ldots, \delta_k$ is the smallest scalar such that

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

  for all k-sparse x

- Define the constant $\theta_{S,S'}$ such that :

$$\left\langle A_T c, A_{T'} c' \right\rangle \leq \theta_{S,S'} \|c\| \|c'\|$$

  holds for all disjoint sets T, T' of cardinality $|T| \leq S$ and $|T'| \leq S'$,

- For all S and S', we have

$$\theta_{S,S'} \leq \delta_{S+S'} \leq \delta_{S,S'} + \max\{\delta_S, \delta_{S'}\}$$

Why this dual certificate? Why $|u_i| < 1$ for all $i \in T^c$?

- Let $S \geq 1$ be such that $\delta_S + \theta_{S,S'} + \theta_{S,2S} < 1$. Then there exits a vector λ such that $\lambda^\top A_j = \text{sgn}(x_j)$ for all $j \in T$ and for all $j \notin T$:

$$|u_j| = |\lambda^\top A_j| \leq \frac{\theta_{S,S'}}{(1-\delta_S-\theta_{S,2S})\sqrt{S}}\|\text{sgn}(x)\| \leq \frac{\theta_{S,S'}}{(1-\delta_S-\theta_{S,2S})} < 1$$

- Assume $S \geq 1$ such that $\delta_S + \theta_{S,S'} + \theta_{S,2S} < 1$. Let x be a real vector supported on T such that $|T| \leq S$. Let $b = Ax$. Then x is a unique minimizer to (P).

- Read Lemma 2.1 and Lemma 2.2 in "E. Candes and T. Tao. Decoding by linear programming. IEEE Transactions on Information Theory, 51:4203–4215, 2005".

General setup

- x not necessarily sparse
- observe $b = Ax$
- recover by ℓ_1 minimization

$$\min \|\hat{x}\|_1 \text{ s. t. } A\hat{x} = b$$

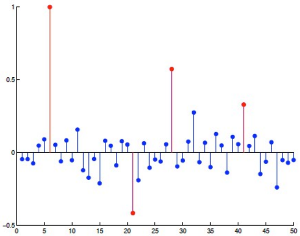

Interested in comparing performance with sparsest approximation x_s:

$$x_s = \arg \min_{\|z\|_0 \le s} \|x - z\|$$

- x_s: s-sparse
- s-largest entries of x are the nonzero entries of x_s

General signal recovery

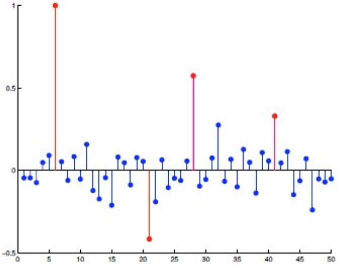

Theorem (Noiseless recovery (C., Romberg and Tao[a]))

If $\delta_{2s} < \sqrt{2} - 1 = 0.414\ldots$, ℓ_1 recovery obeys

$$\begin{aligned} \|\hat{x} - x\|_2 &\lesssim \|x - x_s\|_1/\sqrt{s} \\ \|\hat{x} - x\|_1 &\lesssim \|x - x_s\|_1 \end{aligned}$$

- Deterministic (nothing is random)
- Universal (applies to all x)

- Exact if x is s-sparse
- Otherwise, essentially reconstructs the s largest entries of x
- Powerful if s is close to m

General signal recovery from noisy data

Inaccurate measurements: z error term (stochastic or deterministic)

$$b = Ax + z, \text{ with } \|z\|_2 \le \epsilon$$

Recovery via the LASSO: ℓ_1 minimization with relaxed constraints

$$\min \|\hat{x}\|_1 \text{ s. t. } \|A\hat{x} - b\|_2 \le \epsilon$$

Theorem (C., Romberg and Tao)

Assume $\delta_{2s} < \sqrt{2} - 1$, then

$$\|\hat{x} - x\|_2 \lesssim \frac{\|x - x_s\|_1}{\sqrt{s}} + \epsilon = approx.error + measurement\ error$$

(numerical constants hidden in $\lesssim$ are explicit)

- When $\epsilon = 0$ (no noise), earlier result
- Says when we can solve underdetermined systems of equations accurately

Proof of noisy recovery result

Let $h = \hat{x} - x$. Since $\hat{x}$ and x are feasible, we obtain

$$\|Ah\|_2 \leq \|A\hat{x} - b\|_2 + \|b - Ax\|_2 \leq 2\epsilon$$

The RIP gives

$$|\langle Ah_T, Ah\rangle| \leq \|Ah_T\|_2 \|Ah\|_2 \leq 2\epsilon\sqrt{1+\delta_{2s}}\|h_T\|_2.$$

Hence,

$$\begin{aligned}\|h\|_2 &\leq C_0 \frac{\|x - x_s\|_1}{\sqrt{s}} + C_1 \frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2} \quad \text{lemma 4} \\ &\leq C_0 \frac{\|x - x_s\|_1}{\sqrt{s}} + C_1 2\epsilon\sqrt{1+\delta_{2s}}\end{aligned}$$

Preliminaries: Lemma 1

Let $\Sigma_k = \{x \in \mathbb{R}^n \mid x$ has k nonzero components$\}$

1. If $u \in \Sigma_k$, then $\|u\|_1/\sqrt{k} \le \|u\|_2 \le \sqrt{k}\|u\|_\infty$.
   Proof: $\|u\|_1 = |\langle u, \text{sgn}(u)\rangle| \le \|u\|_2\|\text{sgn}(u)\|_2$.

2. Let u, v be orthogonal vectors. Then $\|u\|_2 + \|v\|_2 \le \sqrt{2}\|u+v\|_2$.
   Proof: Apply the first statement with $w = (\|u\|_2, \|v\|_2)^\top$

3. Let A satisfies RIP of order $2k$. then for any $x, x' \in \Sigma_k$ with disjoint supports

$$|<Ax, Ax'>| \le \delta_{s+s'}\|x\|_2\|x'\|_2$$

Proof: Suppose x and x' are unit vectors as above. Then $\|x+x'\|_2^2 = 2$, $\|x-x'\|_2^2 = 2$ due to the disjoint supports. The RIP gives

$$2(1-\delta_{s+s'}) \le \|Ax \pm Ax'\|_2^2 \le 2(1+\delta_{s+s'})$$

Parallelogram identity

$$|<Ax, Ax'>| = \frac{1}{4}\left|\|Ax+Ax'\|_2^2 - \|Ax-Ax'\|_2^2\right| \le \delta_{s+s'}$$

1. Let T_0 be any subset $\{1, 2, \dots, n\}$ such that $|T_0| \leq s$. For any $u \in \mathbb{R}^n$, define T_1 as the index set corresponding to the s entries of $u_{T_0^c}$ with largest magnitude, T_2 as indices of the next s largest coefficients, and so on. Then

$$\sum_{j \geq 2} \|u_{T_j}\|_2 \leq \frac{\|u_{T_0^c}\|_1}{\sqrt{s}}$$

Proof: We begin by observing that for $j \geq 2$,

$$\|u_{T_j}\|_\infty \leq \frac{\|u_{T_{j-1}}\|_1}{s}$$

since the T_j sort u to have decreasing magnitude. Using Lemma 1.1, we have

$$\sum_{j \geq 2} \|u_{T_j}\|_2 \leq \sqrt{s} \sum_{j \geq 2} \|T_j\|_\infty \leq \sum_{j \geq 1} \frac{\|u_{T_j}\|_1}{\sqrt{s}} = \frac{\|u_{T_0^c}\|_1}{\sqrt{s}}$$

Preliminaries: Lemma 3

- Let A satisfies the RIP with order $2s$. Let T_0 be any subset $\{1, 2, \ldots, n\}$ such that $|T_0| \leq s$ and $h \in \mathbb{R}^n$ be given. Define T_1 as the index set corresponding to the s entries of $h_{T_0^c}$ with largest magnitude, and set $T = T_0 \cup T_1$.Then

$$\|h_T\|_2 \leq \alpha \frac{\|h_{T_0^c}\|_1}{\sqrt{s}} + \beta \frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2}$$

where $\alpha = \frac{\sqrt{2}\delta_{2s}}{1-\delta_{2s}}$ and $\beta = \frac{1}{1-\delta_{2s}}$

Proof: Since $h_T \in \Sigma_{2s}$, the RIP gives

$$(1 - \delta_{2s})\|h_T\|_2^2 \leq \|Ah_T\|_2^2.$$

Continue: Proof Lemma 3

Define T_j as Lemma 2. Since $Ah_T = Ah - \sum_{j\geq 2} Ah_{T_j}$, we have

$$(1-\delta_{2s})\|h_T\|_2^2 \leq \|Ah_T\|_2^2 =< Ah_T, Ah > - < Ah_T, \sum_{j\geq 2} Ah_{T_j} >$$

Lemma 1.3 gives

$$| < Ah_{T_i}, Ah_{T_j} > | \leq \delta_{2s}\|Ah_T\|_2\|Ah\|_2$$

Note that $\|h_{T_0}\|_2 + \|h_{T_1}\|_2 \leq \sqrt{2}\|h_T\|_2$, we have

$$\begin{aligned}
& | < Ah_T, \sum_{j\geq 2} Ah_{T_j} > | = | \sum_{j\geq 2} < Ah_{T_0}, Ah_{T_j} > + \sum_{j\geq 2} < Ah_{T_1}, Ah_{T_j} > | \\
\leq \quad & \delta_{2s}(\|h_{T_0}\|_2 + \|h_{T_1}\|_2) \sum_{j\geq 2} \|h_{T_j}\|_2 \leq \sqrt{2}\delta_{2s}\|h_T\|_2 \sum_{j\geq 2} \|h_{T_j}\|_2 \\
\leq \quad & \sqrt{2}\delta_{2s}\|h_T\|_2 \frac{\|u_{T_0^c}\|_1}{\sqrt{s}}
\end{aligned}$$

Preliminaries: Lemma 4

- Let A satisfies the RIP with order $2s$ with $\delta_{2s} < \sqrt{2} - 1$. Let $x, \hat{x}$ be given and define $h = \hat{x} - x$. Let T_0 denote the index set corresponding to the s entries of x with largest magnitude. Define T_1 be the index set corresponding to the s entries of $h_{T_0^c}$. Set $T = T_0 \cup T_1$. If $\|\hat{x}\|_1 \leq \|x\|_1$. Then

$$\|h\|_2 \leq C_0 \frac{\|x - x_s\|_1}{\sqrt{s}} + C_1 \frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2}$$

where $C_0 = 2\frac{1-(1-\sqrt{2})\delta_{2s}}{1-(1+\sqrt{2})\delta_{2s}}$ and $C_1 = \frac{2}{1-(1+\sqrt{2})\delta_{2s}}$

Proof: Note that $h = h_T + h_{T^c}$, then $\|h\|_2 \leq \|h_T\|_2 + \|h_{T^c}\|_2$. Let T_j be defined similarly as Lemma 2, then we have

$$\|h_{T^c}\|_2 = \|\sum_{j\geq 2} h_{T_j}\|_2 \leq \sum_{j\geq 2} \|h_{T_j}\|_2 \leq \frac{\|h_{T_0^c}\|_1}{\sqrt{s}}$$

Continue: Proof Lemma 4

Since $\|\hat{x}\|_1 \le \|x\|_1$, we obtain

$$\|x\|_1 \ge \|x_{T_0} + h_{T_0}\|_1 + \|x_{T_0^c} + h_{T_0^c}\|_1 \ge \|x_{T_0}\|_1 - \|h_{T_0}\|_1 + \|h_{T_0^c}\|_1 - \|x_{T_0^c}\|_1.$$

Rearranging and again applying the triangle inequality

$$\|h_{T_0^c}\|_1 \le \|x\|_1 - \|x_{T_0}\|_1 + \|h_{T_0}\|_1 + \|x_{T_0^c}\|_1 \le \|x - x_{T_0}\|_1 + \|h_{T_0}\|_1 + \|x_{T_0^c}\|_1.$$

Hence, we have $\|h_{T_0^c}\|_1 \le \|h_{T_0}\|_1 + 2\|x - x_s\|_1$. Therefore,

$$\|h_{T^c}\|_2 \le \frac{\|h_{T_0}\|_1 + 2\|x - x_s\|_1}{\sqrt{s}} \le \|h_{T_0}\|_2 + \frac{2\|x - x_s\|_1}{\sqrt{s}}.$$

Since $\|h_{T_0}\|_2 \le \|h_T\|_2$, we have

$$\|h\|_2 \le 2\|h_T\|_2 + \frac{2\|x - x_s\|_1}{\sqrt{s}}$$

Continue: Proof Lemma 4

Lemma 3 gives

$$\begin{aligned}\|h_T\|_2 &\leq \alpha\frac{\|h_{T_0^c}\|_1}{\sqrt{s}} + \beta\frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2} \\ &\leq \alpha\frac{\|h_{T_0}\|_1 + 2\|x - x_s\|_1}{\sqrt{s}} + \beta\frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2} \\ &\leq \alpha\|h_{T_0}\|_2 + 2\alpha\frac{\|x - x_s\|_1}{\sqrt{s}} + \beta\frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2}\end{aligned}$$

Using $\|h_{T_0}\|_2 \leq \|h_T\|_2$ gives

$$(1-\alpha)\|h_T\|_2 \leq 2\alpha\frac{\|x - x_s\|_1}{\sqrt{s}} + \beta\frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2}.$$

Dividing by $1-\alpha$ gives

$$\|h\|_2 \leq \left(\frac{4\alpha}{1-\alpha} + 2\right)\frac{\|x - x_s\|_1}{\sqrt{s}} + \frac{2\beta}{1-\alpha}\frac{|\langle Ah_T, Ah\rangle|}{\|h_T\|_2}$$

Spark

First questions for finding the sparsest solution to $Ax = b$

- Can sparsest solution be unique? Under what conditions?
- Given a sparse x, how to verify whether it is actually the sparsest one?

Definition (Donoho and Elad 2003)

The spark of a given matrix A is the smallest number of columns from A that are linearly dependent, written as spark(A).

rank(A) is the largest number of columns from A that are linearly independent. In general, spark(A) $\neq$ rank(A) + 1; except for many randomly generated matrices.

Rank is easy to compute, but spark needs a combinatorial search.

Spark

Theorem (Gorodnitsky and Rao 1997)

If $Ax = b$ has a solution x obeying $\|x\|_0 < spark(A)/2$, then x is the sparsest solution.

- **Proof idea**: if there is a solution y to $Ax = b$ and $x - y \neq 0$, then $A(x - y) = 0$ and thus

$$\|x\|_0 + \|y\|_0 \geq \|x - y\|_0 \geq spark(A),$$

  or $\|y\|_0 \geq spark(A) - \|x\|_0 > spark(A)/2 > \|x\|_0$

- The result does not mean this x can be efficiently found numerically.
- For many random matrices $A \in \mathbb{R}^{m \times n}$, the result means that if an algorithm returns x satisfying $\|x\|_0 < (m + 1)/2$, then x is optimal with probability 1.
- What to do when spark(A) is difficult to obtain?

General Recovery - Spark

- Rank is easy to compute, but spark needs a combinatorial search.
- However, for matrix with entries in general positions, spark(A) = rank(A)+1.
- For example, if matrix $A \in \mathbb{R}^{m\times n}$ $(m < n)$ has entries $A_{ij} \sim \mathcal{N}(0,1)$, then $rank(A) = m = spark(A) - 1$ with probability 1.
- In general, any full rank matrix $A \in \mathbb{R}^{m\times n}$ $(m < n)$, any $m+1$ columns of A is linearly dependent, so

$$spark(A) \leq m + 1 = rank(A) + 1$$

Coherence

Definition (Mallat and Zhang 1993)

The (mutual) coherence of a given matrix A is the largest absolute normalized inner product between different columns from A. Suppose $A = [a_1, a_2, \ldots, a_n]$. The mutual coherence of A is given by

$$\mu(A) = \max_{k,j,k\neq j} \frac{|a_k^\top a_j|}{\|a_k\|_2 \cdot \|a_j\|_2}$$

- It characterizes the dependence between columns of A
- For unitary matrices, $\mu(A) = 0$
- For recovery problems, we desire a small $\mu(A)$ as it is similar to unitary matrices.
- For $A = [\Psi, \Phi]$ where Φ and Ψ are $n \times n$ unitary, it holds $n^{-1/2} \leq \mu(A) \leq 1$. Note $\mu(A) = n^{-1/2}$ is achieved with [I, Fourier], [I, Hadamard]. ($|a_k^\top a_j| = 1, \|a_j\| = \sqrt{n}$)
- if $A \in \mathbb{R}^{m\times n}$ where n > m, then $\mu(A) \geq m^{-1/2}$

Coherence

Theorem (Donoho and Elad 2003)

$$spark(A) \geq 1 + \mu^{-1}(A)$$

Proof Sketch

- $\bar{A} \leftarrow A$ with columns normalized to unit 2-norm
- $p \leftarrow spark(A)$
- $B \leftarrow$ a $p \times p$ minor of $\bar{A}^\top \bar{A}$
- $|B_{ii}| = 1$ and $\sum_{j \neq i} |B_{ij}| \leq (p-1)\mu(A)$
- Suppose $p < 1 + \mu^{-1}(A) \Rightarrow 1 > (p-1)\mu(A) \Rightarrow |B_{ii}| > \sum_{j \neq i} |B_{ij}|, \forall i$
- Then $B \succ 0$ (Gershgorin circle theorem) $\Rightarrow spark(A) > p$. Contradiction.

Coherence-base guarantee

Corollary

If $Ax = b$ has a solution x obeying $\|x\|_0 < (1 + \mu^{-1}(A))/2$, then x is the unique sparsest solution.

Compare with the previous

Theorem (Gorodnitsky and Rao 1997)

If $Ax = b$ has a solution x obeying $\|x\|_0 < spark(A)/2$, then x is the sparsest solution.

- For $A \in \mathbb{R}^{m \times n}$ where m < n, $(1 + \mu^{-1}(A))$ is at most $1 + \sqrt{m}$ but spark can be $1 + m$. spark is more useful.
- Assume $Ax = b$ has a solution with $\|x\|_0 = k < spark(A)/2$. It will be the unique ℓ_0 minimizer. Will it be the ℓ_1 minimizer as well? Not necessarily. However, $\|x\|_0 < (1 + \mu^{-1}(A))/2$ is a sufficient condition.

Coherence-based $\ell_0 = \ell_1$

Theorem (Donoho and Elad 2003, Gribonval and Nielsen 2003)

If A has normalized columns and $Ax = b$ has a solution x satisfying $\|x\|_0 \leq (1 + \mu^{-1}(A))/2$, then x is the unique minimizer with respect to both ℓ_0 and ℓ_1.

Proof Sketch

- Previously we know x is the unique ℓ_0 minimizer; let $S := supp(x)$
- Suppose y is the ℓ_1 minimizer but not x; we study $h := y - x$
- h must satisfy $Ah = 0$ and $\|h\|_1 < 2\|h_S\|_1$ since $0 > \|y\|_1 - \|x\|_1 = \sum_{i \in S^c} |y_i| + \sum_{i \in S}(|y_i| - |x_i|) \geq \|h_{S^c}\|_1 - \sum_{i \in S} |y_i - x_i| = \|h_{S^c}\|_1 - \|h_S\|_1$
- $A^\top Ah = 0 \Rightarrow |h_j| \leq (1 + \mu(A))^{-1}\mu(A)\|h\|_1$, $\forall j$. (Expand $A^\top A$ and use $\|h\|_1 = \sum_{k \neq j} |h_k| + |h_j|$)
- the last two points together contradict the assumption

Result bottom line: allow $\|x\|_0$ up to $O(\sqrt{m})$ for exact recovery

The null space of A

- Definition: $\|x\|_p = (\sum_i |x_i|^p)^{1/p}$
- Lemma: Let $0 < p \leq 1$. If $\|(y-x)_{S^c}\|_p > \|(y-x)_S\|_p$, then $\|x\|_p < \|y\|_p$.
  Proof: Let $h = y - x$.
  $\|y\|_p^p = \|x+h\|_p^p = \|x_S + h_S\|_p^p + \|h_{S^c}\|_p^p =$
  $\|x\|_p^p + (\|h_{S^c}\|_p^p - \|h_S\|_p^p) + (\|x_S + h_S\|_p^p - \|x_S\|_p^p + \|h_S\|_p^p))$
  The last term is nonnegative for $0 < p \leq 1$. Hence, a sufficient condition is $\|h_{S^c}\|_p^p > \|h_S\|_p^p$.
- If the condition holds for $0 < p \leq 1$, it also holds for $q \in (0, p]$.
- **Definition** (null space property $NSP(k, \gamma)$). Every nonzero $h \in \mathcal{N}(A)$ satisfies $\|h_S\|_1 < \gamma \|h_{S^c}\|_1$ for all index sets S with $|S| \leq k$.

The null space of A

Theorem (Donoho and Huo 2001, Gribonval and Nielsen 2003)

$\min \ \|x\|_1,$ *s.t.* $Ax = b$ *uniquely recovers all k-sparse vectors* x^o *from measurements* $b = Ax^o$ *if and only if* A *satisfies* $NSP(k, 1)$.

Proof:

- Sufficiency: Pick any k-sparse vector x^o. Let $S := supp(x^o)$. For any non-zero $h \in \mathcal{N}(A)$, we have $A(x^o + h) = Ax^o = b$ and

$$\begin{aligned} \|x^0 + h\|_1 &= \|x_S^o + h_S\|_1 + \|h_{S^c}\|_1 \\ &\geq \|x_S^o\|_1 - \|h_S\|_1 + \|h_{S^c}\|_1 \\ &= \|x_S^o\|_1 - (\|h_S\|_1 - \|h_{S^c}\|_1) \end{aligned}$$

- Necessity. The inequality holds with equality if $\text{sgn}(x_S^o) = -\text{sgn}(h_S)$ and h_S has a sufficiently small scale. Therefore, basis pursuit to uniquely recovers all k-sparse vectors x^o, $NSP(k, 1)$ is also necessary.

The null space of A

- Another sufficient condition (Zhang [2008]) for $\|x\|_1 < \|y\|_1$ is

$$\|x\|_0 < \frac{1}{4}\left(\frac{\|y-x\|_1}{\|y-x\|_2}\right)^2$$

- Proof:

$$\|h_S\|_1 \le \sqrt{|S|}\|h_S\|_2 \le \sqrt{|S|}\|h\|_2 = \sqrt{\|x\|_0}\|h\|_2.$$

Then, the above inequality and the sufficient condition gives $\|y-x\|_1 > 2\|(y-x)_S\|_1$ which is $\|(y-x)_{S^c}\|_1 > \|(y-x)_S\|_1$.

Theorem (Zhang, 2008)

$\min \ \|x\|_1$, *s.t.* $Ax = b$ *recovers* x *uniquely if*

$$\|x\|_0 < \min\left\{\frac{1}{4}\left(\frac{\|h\|_1}{\|h\|_2}\right)^2, \quad h \in \mathcal{N}(A)\backslash\{0\}\right\}$$

The null space of A

- $1 \leq \frac{\|v\|_1}{\|v\|_2} \leq \sqrt{n}, \quad \forall v \in \mathbb{R}^n \setminus \{0\}$
- Garnaev and Gluskin established that for any natural number $p < n$, there exist p-dimensional subspaces $V_p \subset \mathbb{R}^n$ in which

$$\frac{\|v\|_1}{\|v\|_2} \geq \frac{C\sqrt{n-p}}{\sqrt{\log(n/(n-p))}}, \forall v \in V_p \setminus \{0\},$$

- vectors in the null space of A will satisfy, with high probability, the Garnaev and Gluskin inequality for $V_p = \mathrm{Null}(A)$ and $p = n - m$.
- for a random Gaussian matrix A, $\bar{x}$ will uniquely solve ℓ_1-min with high probability whenever

$$\|\bar{x}\|_0 < \frac{C^2}{4} \frac{m}{\log(n/m)}.$$

Formal equivalence

Suppose there is an s-sparse solution to $Ax = b$

- $\delta_{2s} < 1$ solution to combinatorial optimization (min ℓ_0) is unique
- $\delta_{2s} < 0.414$ solution to LP relaxation is unique **and the same**

Comments:

- RIP needs a matrix to be properly scaled
- the tight RIP constant of a given matrix A is difficult to compute
- the result is universal for all s-sparse
- $\exists$ tighter conditions (see next slide)
- all methods (including ℓ_0) require $\delta_{2s} < 1$ for universal recovery; every s-sparse x is unique if $\delta_{2s} < 1$
- the requirement can be satisfied by certain A (e.g., whose entries are i.i.d samples following a subgaussian distribution) and lead to exact recovery for $\|x\|_0 = O(m/log(m/k))$.

More Comments

- (Foucart-Lai) If $\delta_{2s+2} < 1$, then $\exists$ a sufficiently small p so that ℓ_p minimization is guaranteed to recovery any s-sparse x
- (Candes) $\delta_{2s} < \sqrt{2} - 1$ is sufficient
- (Foucart-Lai) $\delta_{2s} < 2(3 - \sqrt{2})/7 \sim 0.4531$ is sufficient
- RIP gives $\kappa(A_S) \leq \sqrt{(1+\delta_s)/(1-\delta_s)}$, $\forall |S| \leq k$. so $\delta_{2s} < 2(3 - \sqrt{2})/7$ gives $\kappa(A_S) \leq 1.7$, $\forall |S| \leq 2m$, very well-conditioned.
- (Mo-Li) $\delta_{2s} < 0.493$ is sufficient
- (Cai-Wang-Xu) $\delta_{2s} < 0.307$ is sufficient
- (Cai-Zhang) $\delta_{2s} < 1/3$ is sufficient and necessary for universal ℓ_1 recovery