


Recovering sparse high dimensional data:

how to do it in the cheapest possible way?



Ivan Nourdin (Université du Luxembourg)

Example 1: group testing.

- Consider a population of d persons (d large), of which a *small* proportion (representing, say, $s \ll d$ persons) is sick.
- One can determine whether a given person is sick or healthy by a blood test. Taking blood is easy, but testing it is costly (in time or in money).
- If one wants to find all the sick persons, the first (bad!) idea is to test each person *individually*. It leads to d blood tests to only find s sick persons.
- But if one thinks more deeply, one can find a much better solution, consisting in only doing about $s \log_2 d$ tests!

Example 2: starry sky.



Example 3: angiography of a hand.




Exemple 4: acquiring-compressing of daily life picture

Representation of a picture in a wavelet basis:




(a)



(b)

- [1st step] Acquiring: 'classically', one measure for each pixel; that is, d measures overall. (Classical size of a picture: $d = 2304 \times 3072 \approx 7.10^6$ pixels.)
- [2nd step] Compressing (jpeg, jpeg2000, ...): one actually keeps only m bits of information (with $m \ll d$) that are enough to recover the picture with a given level of resolution (in general: $m/d \approx 10\%$, even less!)

Here is the kind of pictures which have full spectrum:



Question: Why should we do m measures if, finally, it is to throw most of them away?




In the two papers

- E. Candès, J. Romberg, and T. Tao. “**Robust Uncertainty principles: Exact signal reconstruction from highly incomplete frequency information.**” IEEE Trans. Information Theory, 2006
- D. Donoho. “**Compressed sensing.**” IEEE Trans. Information Theory, 2006

a new method was released, allowing one to significantly reduce the number of measures that are actually needed to recover a *sparse* signal.

It gave raise to a entirely new fields of research, the ‘*Compressed sensing*’ or ‘*Compressive sensing*’.

Since its introduction in the early 2000s, this method has seen a huge buzz (comparable to the buzz around big data)! It has been applied in many different and diverse areas, including medical imaging, astronomy, seismology, ...



Mathematically speaking:

- One considers $x \in \mathbb{R}^d$ (unknown).
- One makes m linear measurements of x : $\langle a_1, x \rangle, \dots, \langle a_m, x \rangle$. That is, one observes $z = Ax \in \mathbb{R}^m$, where $A = \begin{pmatrix} a_1 \\ \vdots \\ a_m \end{pmatrix} \in \mathcal{M}_{m \times d}(\mathbb{R})$.
- How to 'reconstruct' x from Ax ?
- If $m < d$, it is an **undetermined** linear problem. In particular, the set of solutions is infinite.

Crucial assumption: let us suppose that x is s -**sparse**. That is, at most s of its coefficients are nonzero (one writes: $\|x\|_0 \leq s$). How to take advantage of this situation?

- **Proposition**. Assume that $m \geq 2s$ and that any choice of $2s$ rows of A systematically leads to a linearly independent family (such a matrix is easy to build). Then, any s -sparse vector $x \in \mathbb{R}^d$ can be uniquely recovered from $Ax \in \mathbb{R}^m$.

Proof. If x and $x^\#$ are both s -sparse, then their difference $x - x^\#$ is $2s$ -sparse. If, furthermore, they satisfy $Ax^\# = Ax$, then $x^\# = x$ necessarily. ■

- Thus, we are left to solve a minimization problem:


$$(P_0) : \quad \min_{x^\#} \|x^\#\|_0 \quad \text{subject to } Ax^\# = Ax.$$

Unfortunately, to solve (P_0) with a computer takes an exponential time...

- What is easy, in contrast, is to solve


$$(P_2) : \quad \min_{x^\#} \|x^\#\|_2 \quad \text{subject to } Ax^\# = Ax.$$

- Indeed, using the least square method one can check that the solution of (P_2) is explicitly given by $x^\# = {}^tA(A^tA)^{-1}Ax$.




Unfortunately, especially in high dimension, the solution $x^\#$ of (P_2) is likely to be very far away from the expected solution x .

So, despite being easy to implement, this approach is of no help to solve our problem.




We have to find another idea!




$$p = \infty$$




$$p = 2$$



$$p = 1$$



$$0 < p < 1$$



$$p = 0$$



Unit balls in dimension 3 for the norm $\|\cdot\|_p$



Use the ℓ_1 norm, that is, consider

$$(P_1) : \quad \min_{x^\#} \|x^\#\|_1 \quad \text{subject to } Ax^\# = Ax.$$


(In practice, one can solve (P_1) by using the simplex algorithm.)



$$d = 2, \quad m = 1, \quad A = (a_1 \ a_2) \text{ with } a_1, a_2 \neq 0, \quad x = (0, x_2), \quad x^\# = (x_1^\#, x_2^\#)$$

$$Ax^\# = Ax \iff a_1 x_1^\# + a_2 x_2^\# = a_2 x_2$$

In high dimension, specialists (Milman in particular) utilize the following picture for a mental representation of ℓ_1 unit ball in \mathbb{R}^d :



A famous and representative result in the theory of *compressed sensing* is the following (or “how to solve a deterministic problem by introducing randomness”)

Theorem (à la Candès, Romberg and Tao). Consider an integer $m \geq 2\beta s \log d + s$ where $\beta > 1$ is fixed. Assume that $A \in \mathcal{M}_{m \times d}(\mathbb{R})$ is *Gaussian*, more precisely that its entries are independent $N(0, 1/m)$ random variables. Finally, let x be an s -sparse vector of \mathbb{R}^d . Then, with probability at least

$$1 - \frac{2}{df(\beta, s)},$$

one has that x is the unique minimizer to the program


$$(P_1) : \quad \min_{x^\#} \|x^\#\|_1 \quad \text{subject to } Ax^\# = Ax.$$

Fonction f is given by $f(\beta, s) = \left[\sqrt{\frac{\beta}{2s} + \beta} - \sqrt{\frac{\beta}{2s}} \right]^2$. It is increasing in s (for fixed β) and in β (for fixed s).

An experiment (following Amelunxen, Lotz, McCoy and Tropp)

- Fix a large d , say $d = 100$.
- Consider a pair $(s, m) \in \{1, \dots, d\}^2$ (the values for s and m will then vary).
- Pick a s -sparse random vector $x \in \mathbb{R}^d$ at random.
- Compute $z = Ax$ with $A \in \mathcal{M}_{m,d}(\mathbb{R})$ a random Gaussian matrix. Apply the simplex algorithm. If you (don't) get x with an error $\leq 10^{-5}$, then consider it is a success (failure).
- For each possible value of s and m , apply this experiment 50 times, and color the point of coordinates (s, m) with the rule:

50 successes $\rightarrow \circ$... 25 successes $\rightarrow \bullet$... no success $\rightarrow \bullet$



One observes a **phase transition**. The equation for the boundary is very close to $m = 2s \log(d/s) + 2s$ and, as such, agrees Candès, Romberg and Tao's result.

- In order to study this phenomenon, we will analyze it in a more general framework. It will require different mathematical tools, mainly coming from three distinct areas:

* probability theory

* geometry of convex cones

* concentration of measure


- D. Amelunxen, M. Lotz, M.B. McCoy, and J.A. Tropp. **"Living on the edge: phase transitions in convex programs with random data."** *Inform. Inference*, to appear.

- M.B. McCoy and J.A. Tropp. **"From Steiner formulas for cones to concentration of intrinsic volumes"**. *Discrete Comput. Geom.*, 2014.



Let $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function and let us consider the minimization problem:

$$(P) : \quad \min_x f(x) \quad \text{subject to } Ax = Ax_0.$$

Definition. The descent cone of f at x_0 is $\mathcal{D}(f, x_0) = \{y \in \mathbb{R}^d : \exists \tau > 0 \text{ s.t. } f(x_0 + \tau y) \leq f(x_0)\}$.



Fact 1. One has that x_0 is the unique solution to (P) if and only if $\mathcal{D}(f, x_0) \cap \text{null}(A) = \{0\}$.




Fact 2. Since A is Gaussian, its law is invariant by any orthogonal transformation. As a result, $\text{null}(A)(= \ker A)$ is distributed as $Q(\mathbb{R}^{d-m} \times \{0\})$ where Q is chosen at random in $O(d)$.

Consequence. After translation by $-x_0$, our problem reduces to the following *stochastic geometry* problem.


A closed convex cone C and a subspace $L_{d-m} \subset \mathbb{R}^d$ of dimension $d - m$ being given, compute the probability that

$$C \cap QL_{d-m} \neq \{0\},$$

where Q is chosen at random in $O(d)$.




Towards the Crofton formula (polyhedral case)



$$v_0(C) = \frac{1}{2} - \frac{\alpha}{2\pi}, \quad v_1(C) = \frac{1}{2}, \quad v_2(C) = \frac{\alpha}{2\pi}$$

Let $\pi_C : \mathbb{R}^d \rightarrow C$ be the projection onto the polyhedral cone C . Set

$v_k(C) = P\{\pi_C(g) \text{ belongs to } k\text{-dimensional face of } C\}$,
where g denotes a standard Gaussian vector of \mathbb{R}^d .



One has

$$P(\text{blue cone} \cap Q(\text{red line})) = 2 \frac{\text{angle}}{2\pi} = 2v_2(C)$$

Crofton's formula. Provided C is not a subspace, one has


$$\begin{aligned}
 P(C \cap QL_{d-m} \neq \{0\}) &= 2 \sum_{\substack{j=m+1 \\ j-m-1 \text{ even}}}^d v_j(C) \\
 &= 2v_{m+1}(C) + 2v_{m+3}(C) + \dots
 \end{aligned}$$

In the general case (that is when C is possibly not polyhedral), the *intrinsic volumes* $\{v_k(C)\}_{k=0,\dots,d}$ of a closed convex cone $C \subset \mathbb{R}^d$ are defined through

$$\text{Vol}(x \in \mathbb{S}^{d-1} : \text{dist}^2(x, C) \leq \lambda) = \sum_{k=0}^d \beta_{k,d}(\lambda) v_k(C).$$

It is the “conic” version of the celebrated Steiner formula, which is valid for any compact convex set $K \subset \mathbb{R}^d$:

$$\text{Vol}(x \in \mathbb{R}^d : \text{dist}^2(x, K) \leq \lambda) = \sum_{k=0}^d \lambda^{d-k} \text{Vol}(B_{d-k}) \mathcal{V}_k(K).$$



Intrinsic volumes of a closed convex cone are positive and sum to 1. One can therefore consider a random variable V_C defined as

$$P(V_C = k) = v_k(C), \quad k = 0, 1, \dots, d.$$

And by playing a little bit with the Crofton's formula and the definition of V_C , one obtains the 'interlacing property':

$$P(V_C \leq m - 1) \leq P(C \cap QL_{d-m} = \{0\}) \leq P(V_C \leq m).$$

Thus: $\boxed{P(x_0 \text{ is the unique solution of (P)}) \approx P(V_C \leq m)}.$

We are now left to study $P(V_C \leq m)$. To do so, we will rely on a last and crucial ingredient, the *Master Steiner Formula* of McCoy and Tropp (2013).

This formula is particularly useful in our context, as it allows to make a bridge between the abstract random variable V_C and the concrete random variable $\pi_C(\mathbf{g})$, where π_C is the projection onto the cone C and $\mathbf{g} \sim N(0, I_d)$.

It is our gateway towards classical results of concentration of measure (Talagrand, Ledoux, ...). By pushing this idea further, one can mathematically prove the phase transition.

Master Steiner Formula (McCoy, Tropp) Let $\pi_C : \mathbb{R}^d \rightarrow C$ be the projection onto the closed convex cone C . Let $\mathbf{g} \sim N(0, I_d)$. Let $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a measurable function. One has

$$E[\varphi(\|\pi_C(\mathbf{g})\|^2)] = \sum_{k=0}^d E[\varphi(X_k)]P(V_C = k),$$

where X_k is distributed according to the χ^2 law with k degrees of freedom.

Otherwise stated, $\|\pi_C(\mathbf{g})\|^2 \stackrel{(\text{law})}{=} \sum_{i=1}^{V_C} \eta_i^2$, where $\eta_1, \eta_2, \dots \sim N(0, 1)$ are independent and also independent from V_C .

Corollary/Definition: Statistical dimension δ_C of a closed convex cone C is defined as $E[\|\pi_C(\mathbf{g})\|^2] = E[V_C]$.

Another corollary of the Master Steiner Formula is that

$$E[e^{\eta V_C}] = E[e^{\xi \|\pi_C(\mathbf{g})\|^2}], \quad \text{with } \xi = \frac{1}{2}(1 - e^{-2\eta}).$$

An interesting consequence is the following. If C_d is a sequence of closed convex cone of \mathbb{R}^d such that $E(V_{C_d}) = \delta_{C_d} \rightarrow \infty$ and $\liminf \text{Var}(V_{C_d})/\delta_{C_d} > 0$ as $d \rightarrow \infty$, then

$$\frac{V_{C_d} - \delta_{C_d}}{\sqrt{\text{Var}(V_{C_d})}} \rightarrow N(0, 1) \quad \text{iff} \quad \frac{\|\pi_{C_d}(\mathbf{g})\|^2 - \delta_{C_d}}{\sqrt{\text{Var}(\|\pi_{C_d}(\mathbf{g})\|^2)}} \rightarrow N(0, 1).$$

Theoreme (Goldstein, N., Peccati, 2014): For all closed convex cone C , one has, with $N \sim N(0, 1)$,

$$d_{TV} \left(\frac{\|\pi_C(\mathbf{g})\|^2 - \delta_C}{\sqrt{\text{Var}(\|\pi_C(\mathbf{g})\|^2)}}, N \right) \leq \frac{8}{\sqrt{\delta_C}}.$$

Corollary. Let $x_0 \in \mathbb{R}^d$, let $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function and let $C = \{y \in \mathbb{R}^d : \exists \tau > 0 \text{ such that } f(x_0 + \tau y) \leq f(x_0)\}$ be the descent cone of f at x_0 .

Consider the minimization problem

$$(P) : \quad \min_x f(x) \quad \text{subject to } Ax = Ax_0,$$

where $A \in \mathcal{M}_{m \times d}(\mathbb{R})$ is *Gaussian* (all its entries are independent $N(0, 1)$ random variables) and where $m = \lfloor \delta_C + t\sqrt{\text{Var}(V_C)} \rfloor$, $t \in \mathbb{R}$.

Suppose that $E(V_C) = \delta_C \rightarrow \infty$ and that $\liminf \text{Var}(V_C)/\delta_C > 0$ as $d \rightarrow \infty$.

Then, as $d \rightarrow \infty$,

$$P(x_0 \text{ is the unique solution of (P)}) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-\frac{u^2}{2}} du + O\left(\frac{1}{\sqrt{\log \delta_C}}\right).$$

Reading the phase transition.

For instance, selecting $m \geq \delta_C + 1.6\sqrt{2\delta_C}$, one has

$$P(x_0 \text{ is the unique solution of (P)}) \geq 0.95.$$

In contrast, for $m \leq \delta_C - 1.6\sqrt{2\delta_C}$,

$$P(x_0 \text{ is the unique solution of (P)}) \leq 0.05$$

It follows that the phase transition happens on an interval of length $3.2\sqrt{\delta_C}$.

A few ingredients allowing us to prove the theorem

The total variation distance between F and G is defined as

$$d_{TV}(F, G) = \sup_{A \subset \mathbb{R}} |P(F \in A) - P(G \in A)|.$$

Stein's lemma. If $E[F] = 0$ and $E[F^2] = 1$, then

$$d_{TV}(F, N) \leq \sup_{\phi} |E[\phi'(F)] - E[F\phi(F)]|$$

where $N \sim N(0, 1)$ and where the supremum is taken over all the C^1 functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\|\phi'\|_{\infty} \leq 2$.

Malliavin-Stein approach.

Let $H : \mathbb{R}^d \rightarrow \mathbb{R}$ be a C^1 -function, and let $\mathbf{g} \sim N(0, I_d)$ be a standard Gaussian vector of \mathbb{R}^d .

Let $F = H(\mathbf{g})$ and let $m = E[F]$ and $\sigma^2 = \text{Var}(F)$.

For $t \geq 0$, set $\mathbf{g}_t = e^{-t}\mathbf{g} + \sqrt{1 - e^{-2t}}\hat{\mathbf{g}}$, where $\hat{\mathbf{g}}$ stands for an independent copy of \mathbf{g} . Write \hat{E} to indicate the expectation with respect to $\hat{\mathbf{g}}$.

Then, with $N \sim N(m, \sigma^2)$,

$$d_{TV}(F, N) \leq \frac{2}{\sigma^2} \sqrt{\text{Var} \left(\int_0^\infty e^{-t} \langle \nabla H(\mathbf{g}), \hat{E}(\nabla H(\hat{\mathbf{g}}_t)) \rangle dt \right)}.$$

Poincaré inequality. Let $H : \mathbb{R}^d \rightarrow \mathbb{R}$ be of class C^1 and let $\mathbf{g} \sim N(0, I_d)$ be a standard Gaussian vector of \mathbb{R}^d . One has

$$\text{Var}(H(\mathbf{g})) \leq E \|\nabla H(\mathbf{g})\|^2.$$

Theoreme (Goldstein, N., Peccati, 2014): For all closed convex cone C , one has, with $N \sim N(0, 1)$,

$$d_{TV} \left(\frac{\|\pi_C(\mathbf{g})\|^2 - \delta_C}{\sqrt{\text{Var}(\|\pi_C(\mathbf{g})\|^2)}}, N \right) \leq \frac{8}{\sqrt{\delta_C}}.$$

Sketch of the proof: Set $F = H(\mathbf{g}) = \|\pi_C(\mathbf{g})\|^2 - \delta_C$. One has, with $\sigma_C^2 = \text{Var}(F)$,

$$d_{TV}(F, N(0, \sigma_C^2)) \leq \frac{2}{\sigma_C^2} \sqrt{\text{Var} \left(\int_0^\infty e^{-t} \langle \nabla H(\mathbf{g}), \hat{E}(\nabla H(\hat{\mathbf{g}}_t)) \rangle dt \right)}.$$

One observes that

$$\nabla H(\mathbf{x}) = \nabla \|\pi_C(\mathbf{x})\|^2 = 2\pi_C(\mathbf{x})$$

and that $\text{Var}(\varphi(\mathbf{g})) \leq E \|\nabla \varphi(\mathbf{g})\|^2$ (Poincaré). And everything fits particularly well, leading to the desired conclusion!

Thanks for your attention!