Steining the Steiner formula

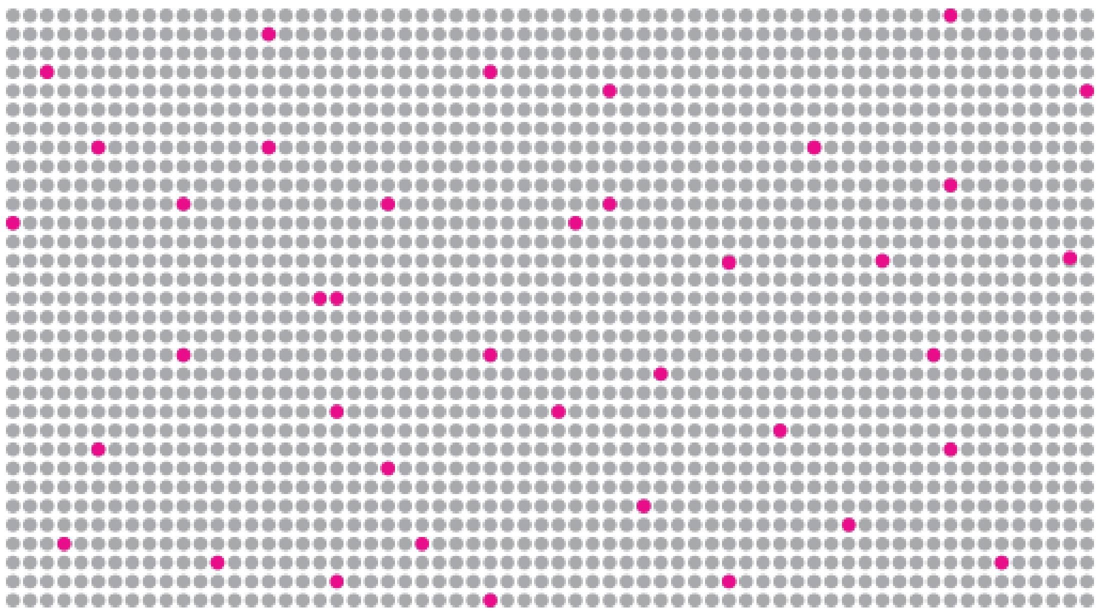

Ivan Nourdin (based on a joint work with L. Goldstein and G. Peccati)

UNIVERSITÉ DU LUXEMBOURG

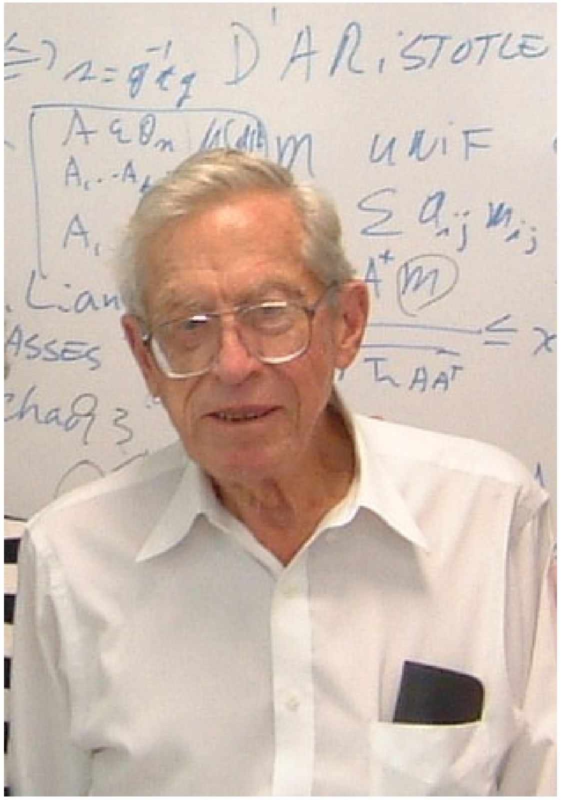

Stein's method: If F is a square integrable real-valued random variable such that $E[F] = 0$ and if $N \sim N(0,1)$, then

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

where the supremum runs over all C^1 functions $\phi : \mathbb{R} \to \mathbb{R}$ with $\|\phi'\|_\infty \leq 2$.

Steiner formula: if $K \subset \mathbb{R}^d$ is a convex body, then

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

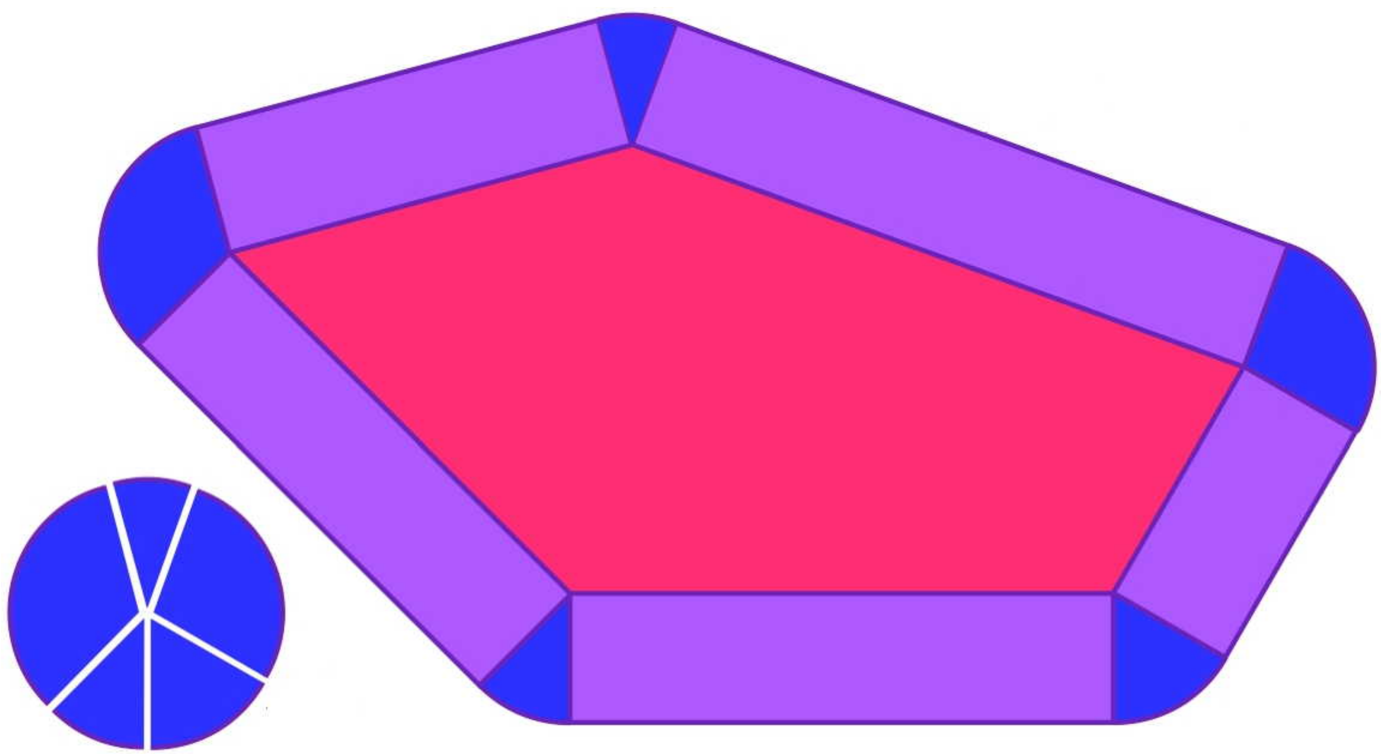

- Let $x_0 \in \mathbb{R}^d$ be unknown, where d is meant to be large ($d = 100$, $d = 1000$, $d = 10^6$, etc.)

Problem: We want to *acquire* x_0 with the *smallest* possible number of *linear* measurements.

- Let $x_0 \in \mathbb{R}^d$ be unknown, where d is meant to be large ($d = 100$, $d = 1000$, $d = 10^6$, etc.)

Problem: We want to *acquire* x_0 with the *smallest* possible number of *linear* measurements.

- One makes m linear measurements of x_0: $\langle a_1, x_0 \rangle, \ldots, \langle a_m, x_0 \rangle$.

That is, one observes $Ax_0 \in \mathbb{R}^m$, where $A = \begin{pmatrix} a_1 \\ \hline \vdots \\ \hline a_m \end{pmatrix} \in \mathbb{R}^{m \times d}$.

Notations that will be used throughout the talk:

- d is the <u>ambient dimension</u>

- m is the <u>number of measurements</u>

- $A \in \mathbb{R}^{m \times d}$ is the <u>measurement matrix</u>

Problem: "I give you A and $y = Ax_0$. From this, are you able to recover x_0?".

Problem: "I give you A and $y = Ax_0$. From this, are you able to recover x_0?".

If $m \geq d$, it is a *determined* system, and it is easy.

Problem: "I give you A and $y = Ax_0$. From this, are you able to recover x_0?".

If $m \geq d$, it is a *determined* system, and it is easy.

If $m < d$, it is an *undetermined* system, so there is no hope to provide a positive answer.

Extra assumption: x_0 is s-sparse, that is, at most s of its entries are nonzero. (But, of course, we don't know which ones!)

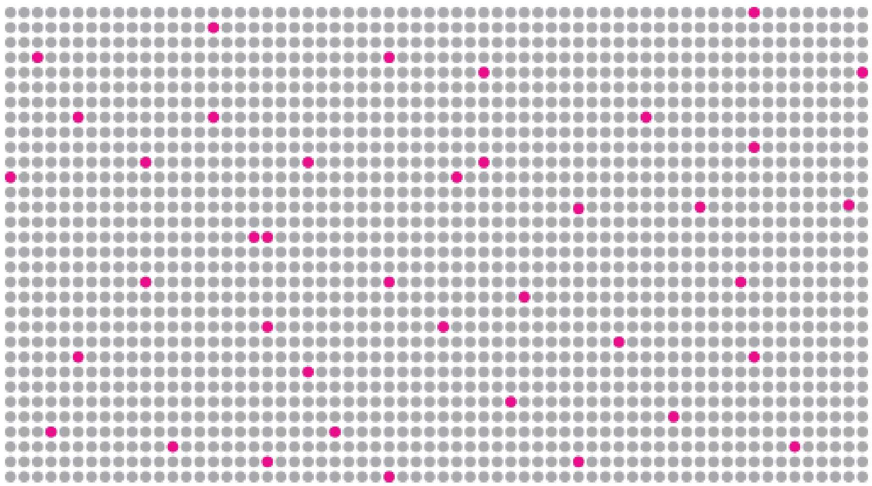

Questions:

- is such an assumption realistic in practice?

- what is the gain of doing such an assumption?

Example 1: pictures taken with smartphones

- Assume x_0 encodes a picture of size $n \times m$, e.g. $n = 2592$, $m = 1944 \Rightarrow d = 5\,038\,848$

- Each entry of x_0 has a value between 0 (black) and 15 (white), depending on the luminosity at the corresponding pixel

- x_0 itself is not sparse. But the vector of differences between adjacent elements of x_0 is $416\,686$-sparse ($\approx 8\%$).

Example 2: Medical Resonance Imagery

Example 3: Seismology

Example 4: High-resolution radar

Example 5: Analog-to-digital converters

....

....

..

Extra assumption: x_0 is s-sparse, that is, at most s of its entries are nonzero. (But, of course, we don't know which ones!)

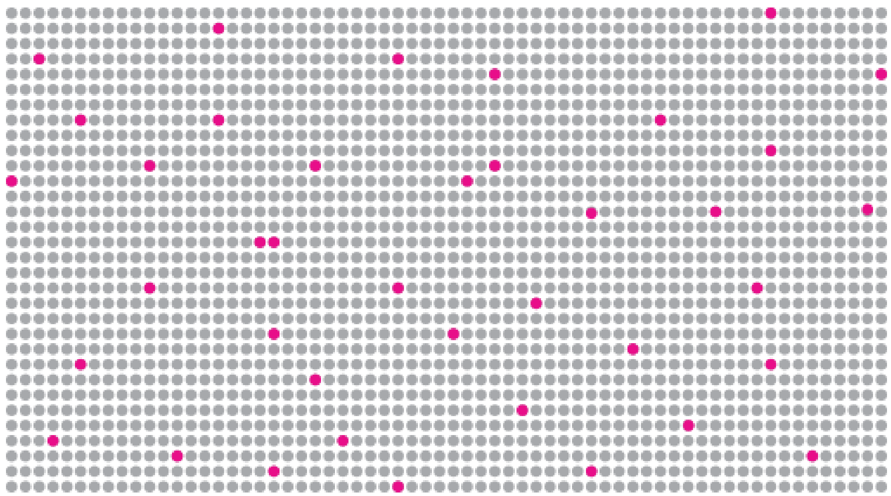

Questions:

- is such an assumption realistic in practice? **YES**

- what is the gain of doing such an assumption?

What is the gain of assuming sparsity?

- 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

To recover $x_0 \in \mathbb{R}^d$ from $A \in \mathbb{R}^{m \times d}$ and $Ax_0 \in \mathbb{R}^m$, let us consider a minimization problem. The following one is, at first glance, the most natural to consider:

$$(P_0): \qquad \min_x \|x\|_0 \quad \text{subject to } Ax = Ax_0,$$

where $\|x\|_0$ is the cardinality of the support of x.

Question: Is x_0 the unique solution to (P_0)?

Answer: Yes, provided that $m \geq 2s$ (and that easy-to-check conditions on A are satisfied)

So, is the problem (already) over?

- In order to solve (P_0), we have to consider all the possible supports for x_0 and then to solve the corresponding systems.

- For instance, suppose that $d = 1000$ and $s = 10$. We have to solve $\binom{1000}{10} \geq 10^{20}$ linear systems of size 10×10. Each such system can be solved in 10^{-10} seconds. Then, the time required to solve (P_0) is around 10^{10} seconds, i.e., more than...

So, is the problem (already) over?

- In order to solve (P_0), we have to consider all the possible supports for x_0 and then to solve the corresponding systems.

- For instance, suppose that $d = 1000$ and $s = 10$. We have to solve $\binom{1000}{10} \geq 10^{20}$ linear systems of size 10×10. Each such system can be solved in 10^{-10} seconds. Then, the time required to solve (P_0) is around 10^{10} seconds, i.e., more than... 300 years!

- What is easy and quick, contrary to (P_0), is to solve

$$(P_2): \qquad \min_x \|x\|_2 \quad \text{subject to } Ax = Ax_0.$$

- Indeed (least square method): one can check that the solution of (P_2) is explicitely given by $x = A^T(AA^T)^{-1}Ax_0$.

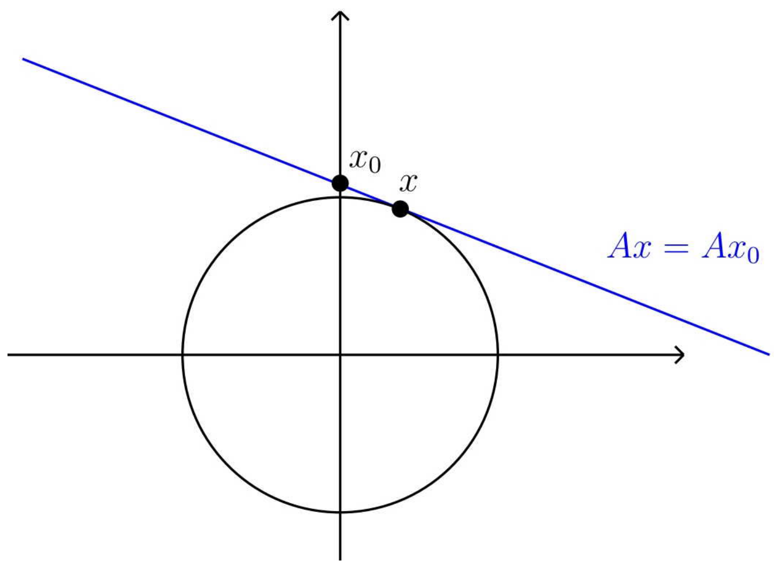

Unfortunately, especially in high dimension, the solution x of (P_2) is likely to be very far away from the expected solution x_0.

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

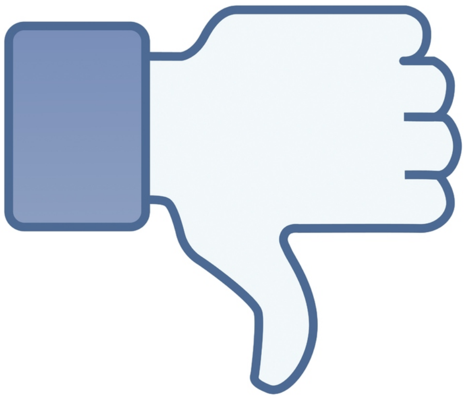

As we will see in a minute, it turns out that a much better idea is to use the ℓ_1 norm, that is, to consider

$$(P_1): \qquad \min_x \|x\|_1 \quad \text{subject to } Ax = Ax_0.$$

Nota: the *simplex algorithm*, which is quick and efficient, allows to solve (P_1) in practice.

Why does ℓ_1-norm sound like a more reasonable candidate?

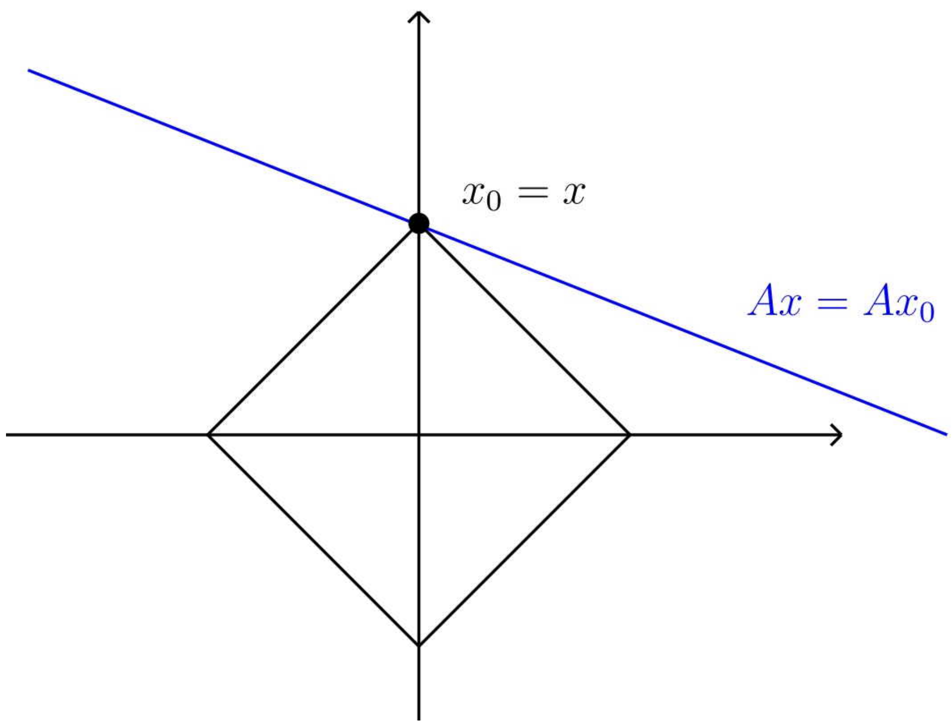

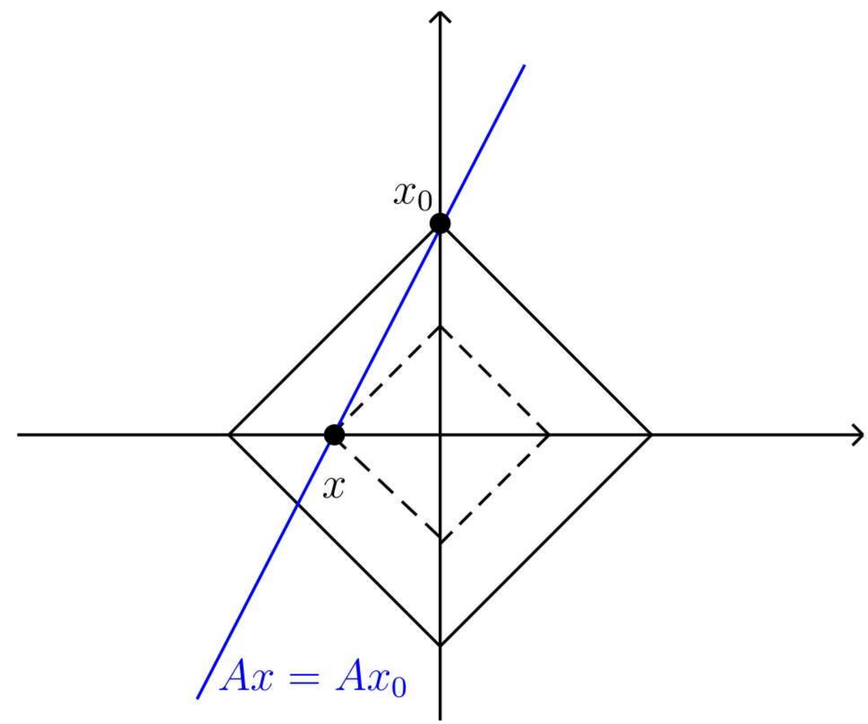

A famous and representative result in the theory of *compressed sensing* is the following theorem (or "how to solve a deterministic problem by introducing randomness")

Theorem (à la Candès, Romberg and Tao). Assume that the number of measurements m satisfies $m \geq 2\beta s \log d + s$, where $\beta > 1$ is fixed. Assume that $A \in \mathbb{R}^{m\times d}$ is *Gaussian*. Finally, let x_0 be an s-sparse vector of $\mathbb{R}^d$. Then, with probability at least

$$1 - \frac{2}{d^{f(\beta,s)}},$$

one has that x_0 is the unique minimizer to the program

$$(P_1): \qquad \min_x \|x\|_1 \quad \text{subject to } Ax = Ax_0.$$

The 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).

A classical experiment (following Donoho and Tanner)

- Fix a large d, say $d = 100$.

- Consider a pair $(s, m) \in \{1, \ldots, d\}^2$ (the values for s and m will then vary).

- Pick a s-sparse vector $x_0 \in \mathbb{R}^d$ at random.

- Compute Ax_0 with $A \in \mathbb{R}^{m \times d}$ a random Gaussian matrix. Apply the simplex algorithm. If you (don't) get x_0, then consider it is a success (failure).

- For each possible value of s and m, repeat this experiment 10 times, and color the point of coordinates (s, m) with the rule:

10 successes $\to$ ● ... 5 successes $\to$ ● ... no success $\to$ ●

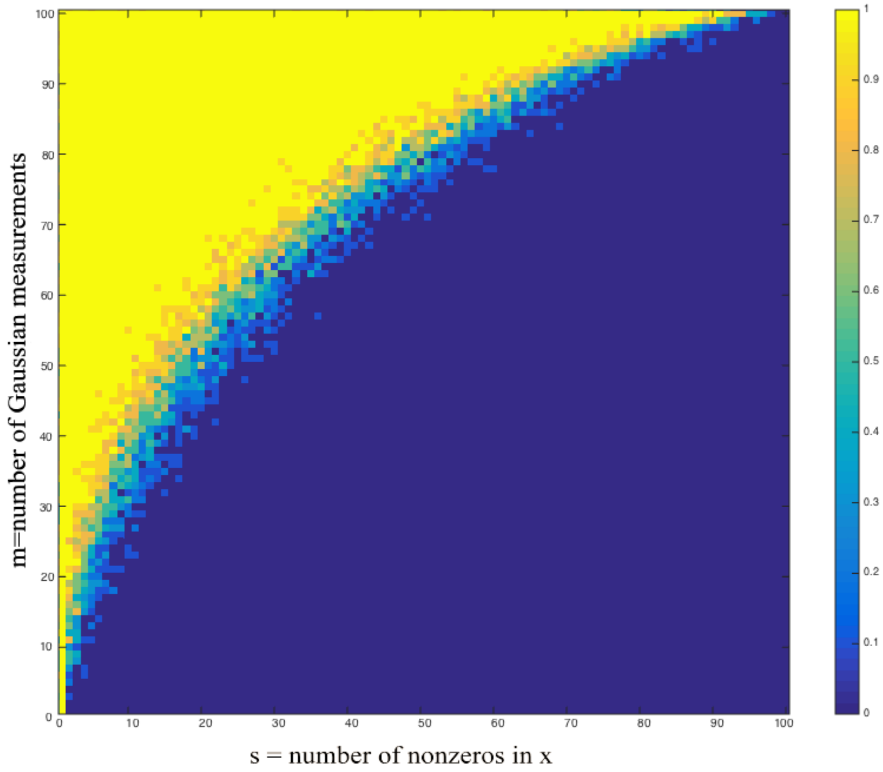

One observes a strong **phase transition**. The equation for the boundary is very close to $m = s\log(d/s) + s$.

In order to understand this threshold phenomenon, we will analyze it in a more general framework, following the two references:

- 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 $x_0 \in \mathbb{R}^d$, let $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ be a convex function, let $A \in \mathbb{R}^{m \times d}$ be a Gaussian matrix, and let us consider the minimization problem:

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

Goal: We want to compute $P(x_0$ is the unique solution of (P)$)$.

Definition. The descent cone of f at x_0 is the cone generated by the perturbations of f at x_0 that do not increase f:

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

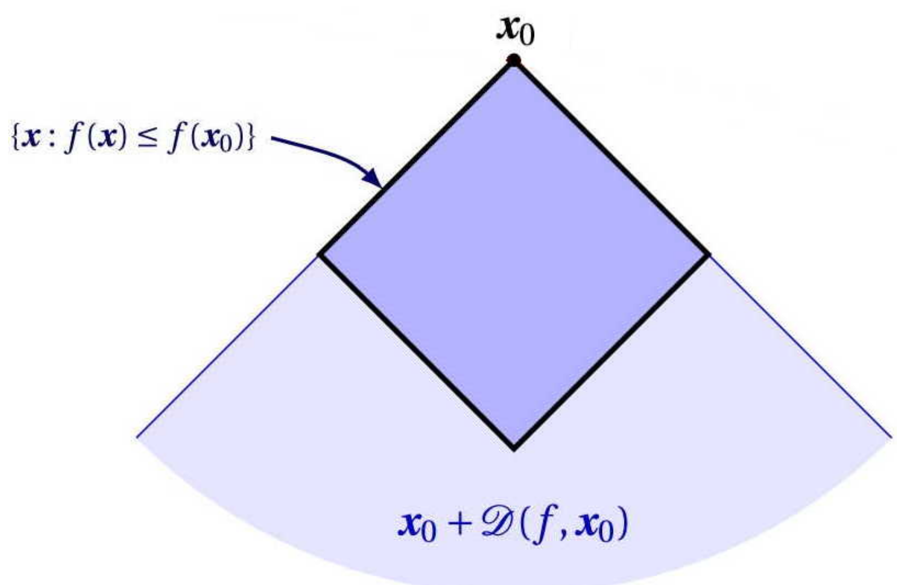

Fact 1. One has that x_0 is the unique solution to (P) if and only if $\mathcal{D}(f, x_0) \cap \mathsf{null}(A) = \{0\}$. $(\mathsf{null} A = \mathsf{ker} A)$

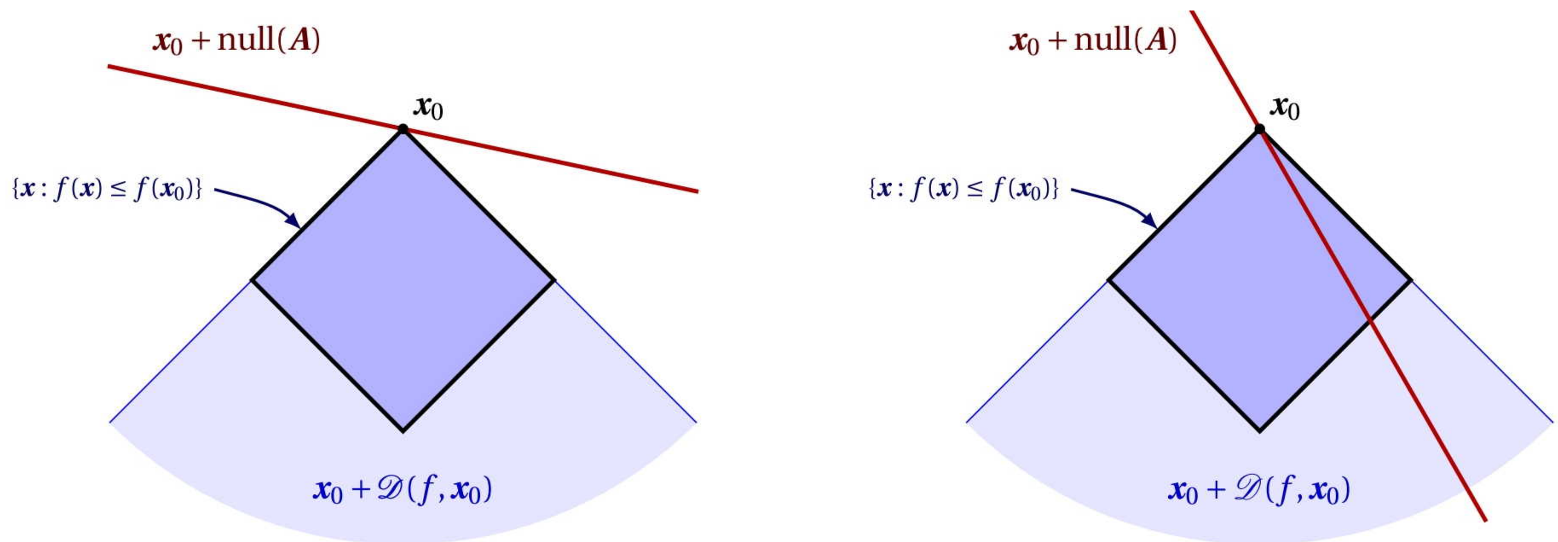

Fact 2. Since A is Gaussian, its law is invariant by any orthogonal transformation. As a result, $\mathsf{null}(A)$ is distributed as QL_{d-m}, where Q is chosen at random in $O(d)$ and $L_{d-m} \subset \mathbb{R}^d$ stands for any (fixed) subspace of dimension $d - m$.

As a consequence of these two facts, we have that

$$\boxed{P(x_0 \text{ is the unique solution of (P)}) = P(C \cap QL_{d-m} = \{0\})}.$$

where C is the descent cone of f at x_0, $L_{d-m} \subset \mathbb{R}^d$ is a (fixed) subspace of dimension $d-m$, and Q is chosen at random in $O(d)$

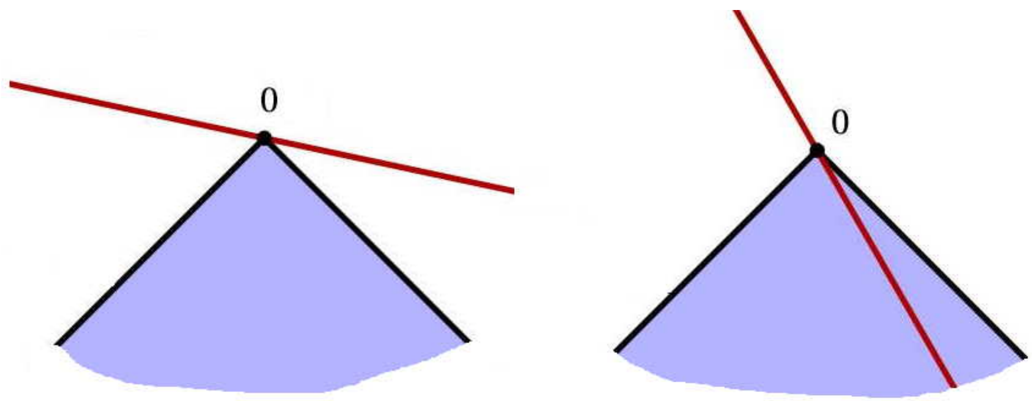

The conic version of the Steiner formula reads as follows. If $C \subset \mathbb{R}^d$ is a closed convex cone, then

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

where $\beta_{k,d}(\lambda)$ is an (explicit) quantity only depending on λ (but not on C).

The quantities $\{v_k(C)\}_{k=0,\dots,d}$ are called the *intrinsic volumes* of the cone C. They are positive and sum to 1.

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

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

By playing a little bit with the Crofton's formula, one can show the 'interlacing property':

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

where the (abstract) random variable V_C defined as

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

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 shall rely on a last and final ingredient, the *Master Steiner Formula*.

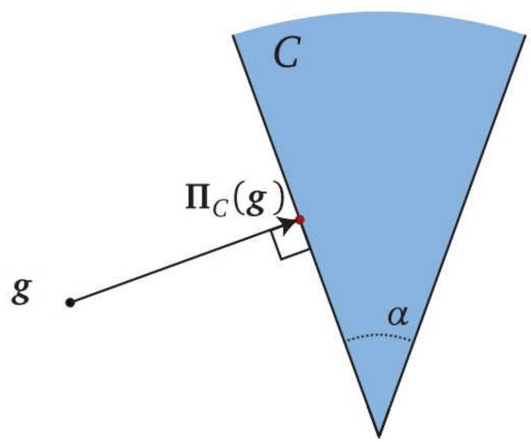

Master Steiner Formula (McCoy, Tropp, 2013) Let $\Pi_C : \mathbb{R}^d \to C$ be the projection onto the closed convex cone C. One has

$$\|\Pi_C(\mathbf{g})\|^2 \overset{\text{(law)}}{=} \chi^2_{V_C}.$$

Corollary 1/Definition: The statistical dimension δ_C of a closed convex cone C is defined as $E[\|\Pi_C(\mathbf{g})\|^2] = E[V_C]$.

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

An important consequence of Corollary 2 is the following.

If C_d is a sequence of closed convex cone of $\mathbb{R}^d$ such that $E(V_{C_d}) = \delta_{C_d} \to \infty$ and $\liminf \mathsf{Var}(V_{C_d})/\delta_{C_d} > 0$ as $d \to \infty$, then

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

Theorem (Goldstein, Nourdin, Peccati). Let $x_0 \in \mathbb{R}^d$, let $f : \mathbb{R}^d \to \mathbb{R}$ 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

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

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

Suppose that $E(V_C) = \delta_C \to \infty$ and that $\liminf \mathrm{Var}(V_C)/\delta_C > 0$ as $d \to \infty$. Then, as $d \to \infty$,

$$\boxed{P(x_0 \text{ is the unique solution of (P)}) \to \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} e^{-\frac{u^2}{2}} du.}$$

Reading the phase transition.

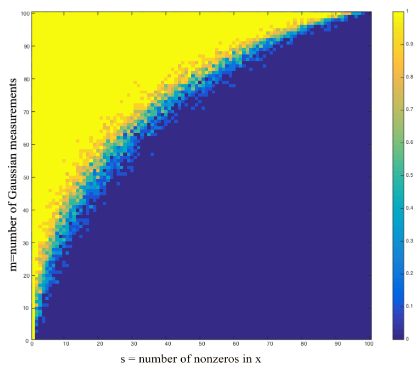

For instance, selecting $m \geq \delta_C + 1.6\sqrt{\mathsf{Var}(V_C)}$, one has

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

In contrast, for $m \leq \delta_C - 1.6\sqrt{\mathsf{Var}(V_C)}$,

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

(*Remark*: A crude bound is $\mathsf{Var}(V_C) \leq 2\delta_C$.)

Theorem (G.-N.-P., 2014): Let C be any closed convex cone and let Π_C denote the projection onto C. One then has, with $N \sim N(0,1)$ and $\mathbf{g} \sim N(0, I_d)$,

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