2585-20

Joint ICTP-TWAS School on Coherent State Transforms, Time-Frequency and Time-Scale Analysis, Applications

2 - 20 June 2014

Quantizing compressed sensing: From high resolution to 1-bit quantization scheme

L. Jacques

UCL, Louvain-La-Neuve, Belgium

Quantizing compressed sensing: From high resolution to 1-bit quantization scheme

Laurent Jacques, UCL, Belgium

Coherent state transforms, time-frequency and time-scale analysis, applications

The Abdus Salam International Centre for Theoretical Physics

Compressive Sampling

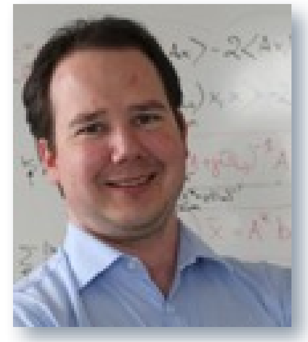

H. Rauhut's tutorial

Compressed Sensing

Highly compressed recap of what is ...

Compressive Sensing

Compressed Sampling

Generally, sampling is ...

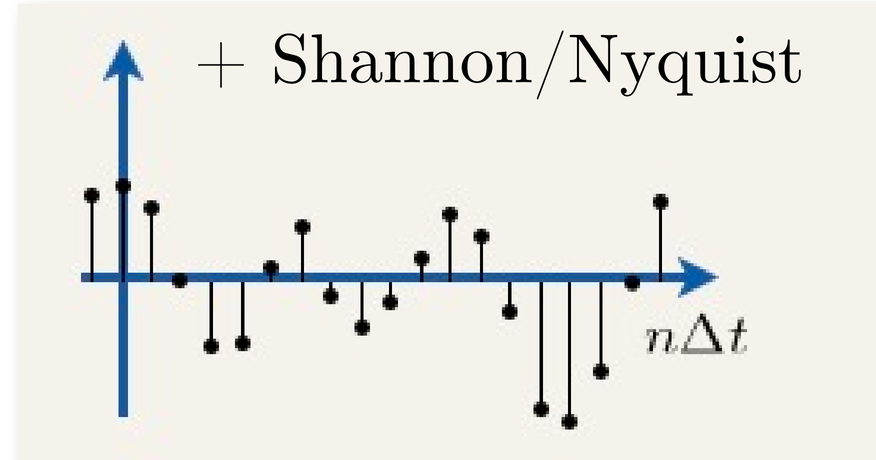

Human readable signal!

Generally, sampling is ...

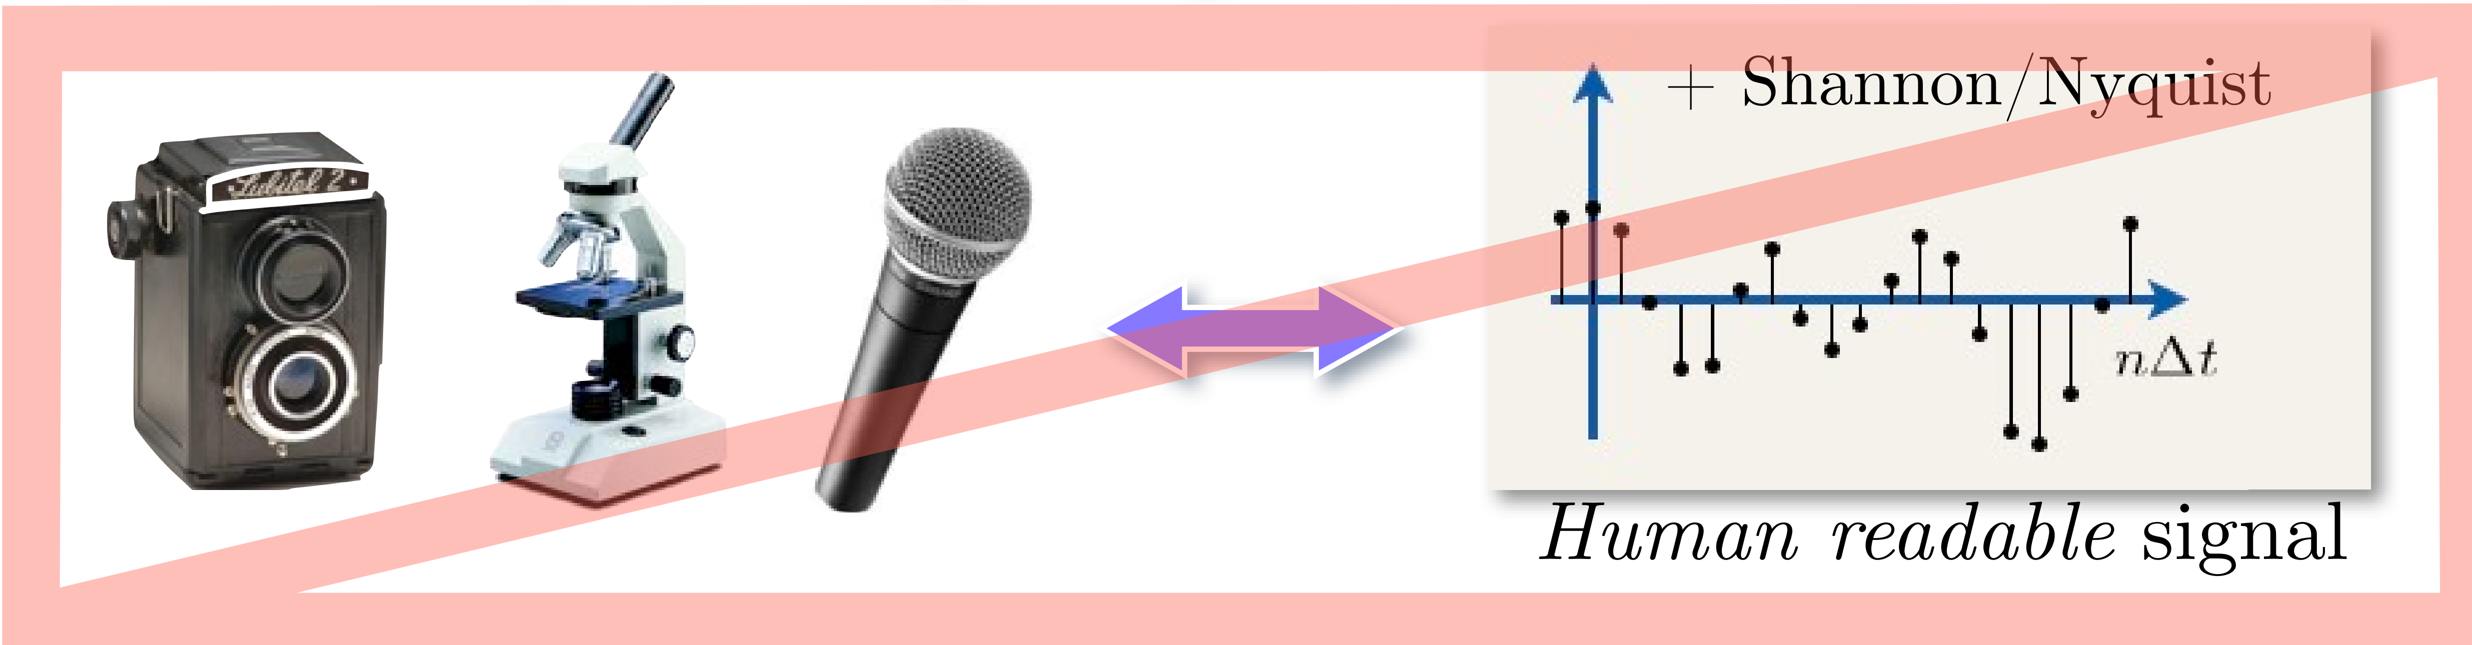

New ways to sample signals ....

"*Computer readable*" sensing + prior information

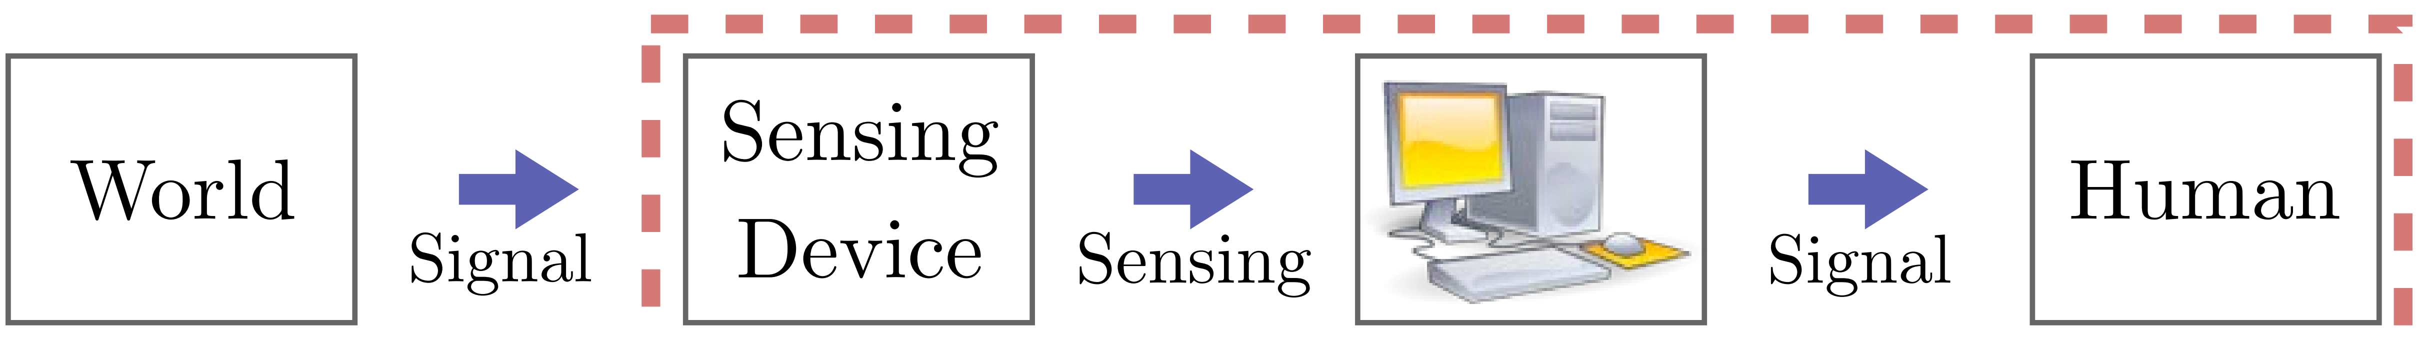

Optimized setup: sampling rate $\propto$ information

Generally, sampling is ...

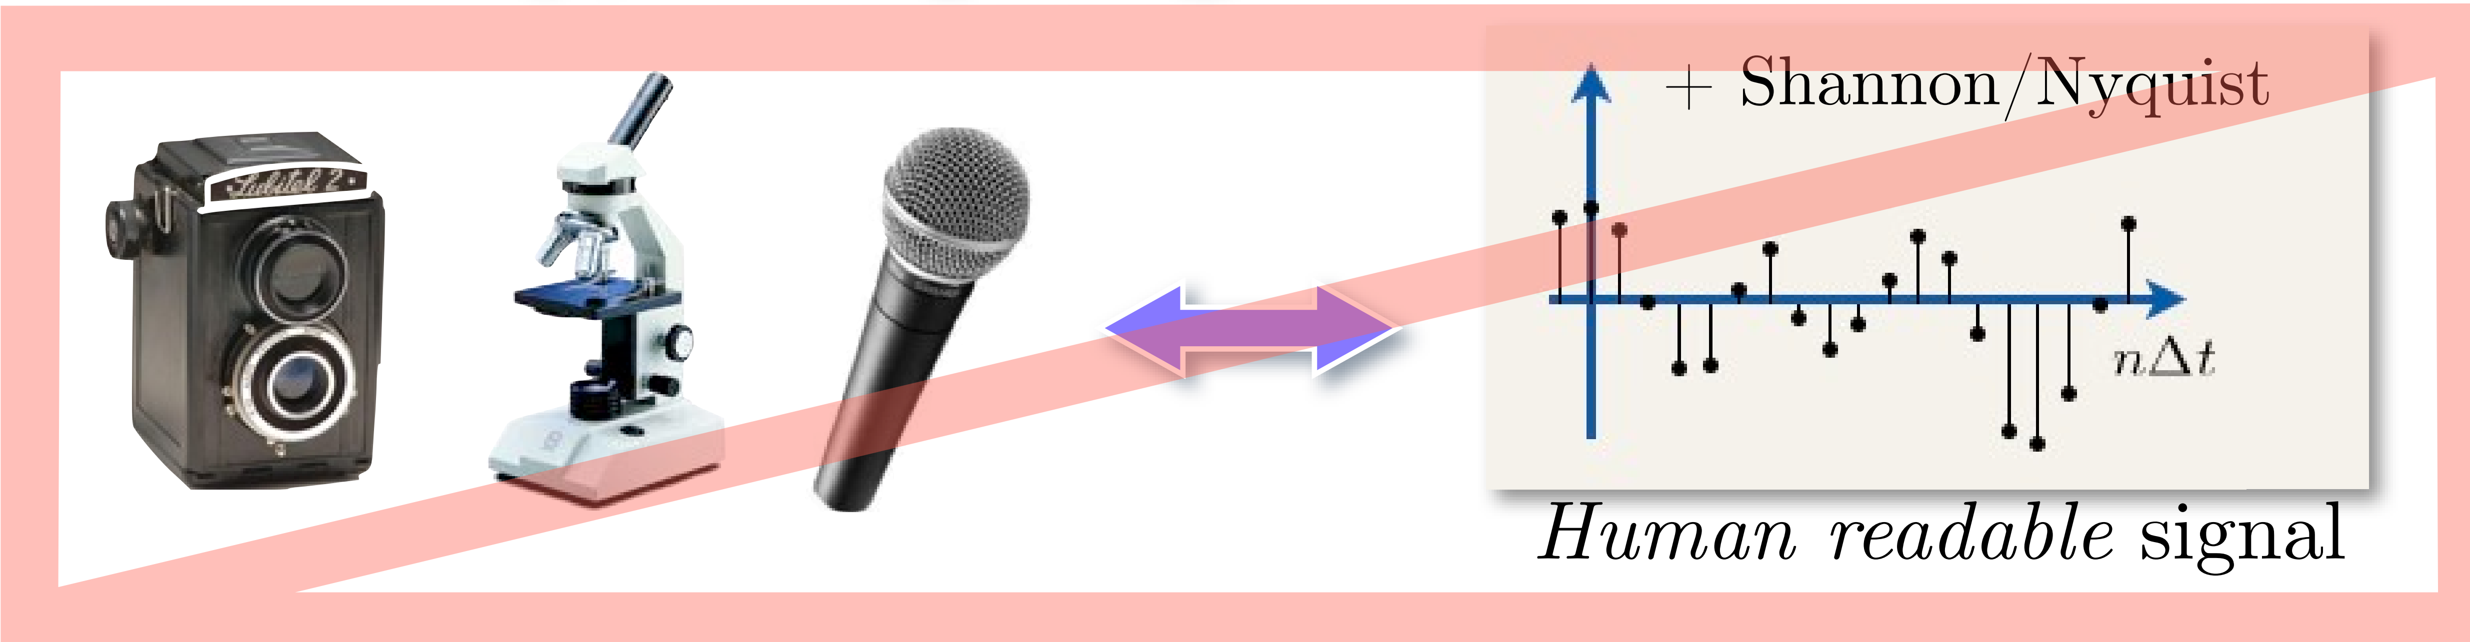

New ways to sample signals ....

structures, sparsity, low-rank, ...

"*Computer readable*" sensing + prior information

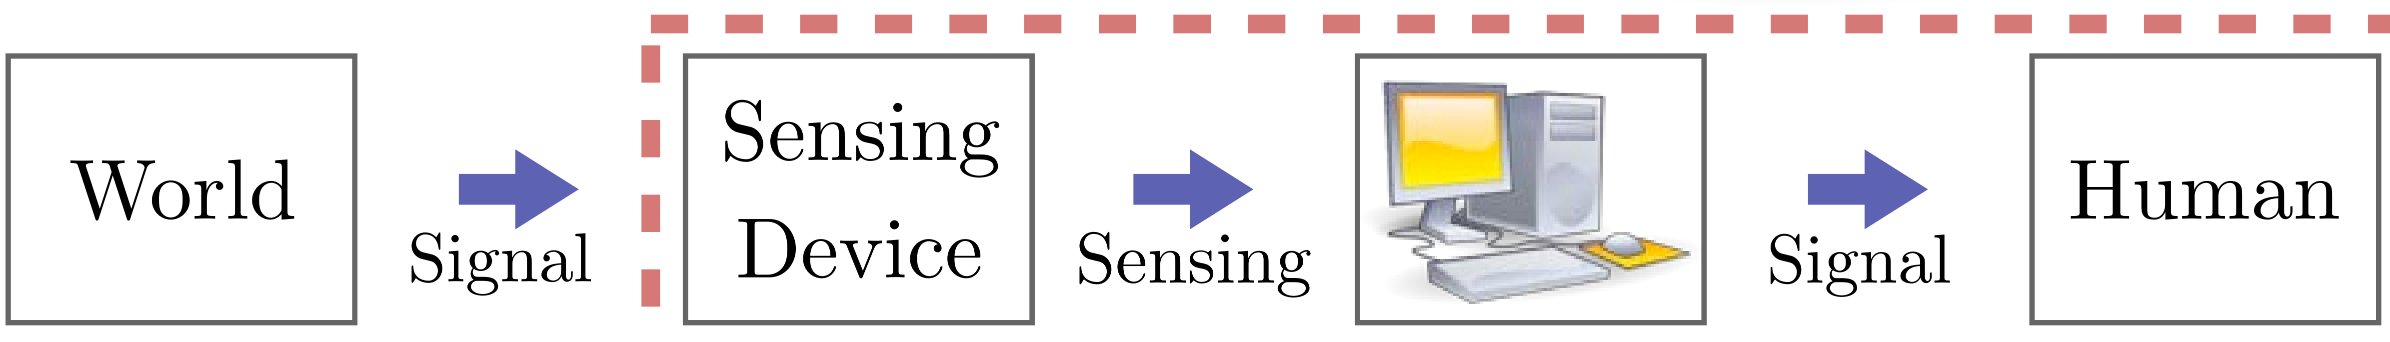

Optimized setup: sampling rate $\propto$ information

Compressed Sensing

... *in a nutshell:*

"Forget" Dirac, forget Nyquist,
ask *few* (**linear**) *questions*
about your informative (**sparse**) signal,
and recover it *differently* (**non-linearly**)"

Compressed Sensing

Assumption: **the probability that our world is totally discrete is very high ...**

Compressed Sensing

Then ...

M questions — Sensing method — Signal

$\boldsymbol{y}$ $\quad\boldsymbol{\Phi}\quad$ $\boldsymbol{x}$

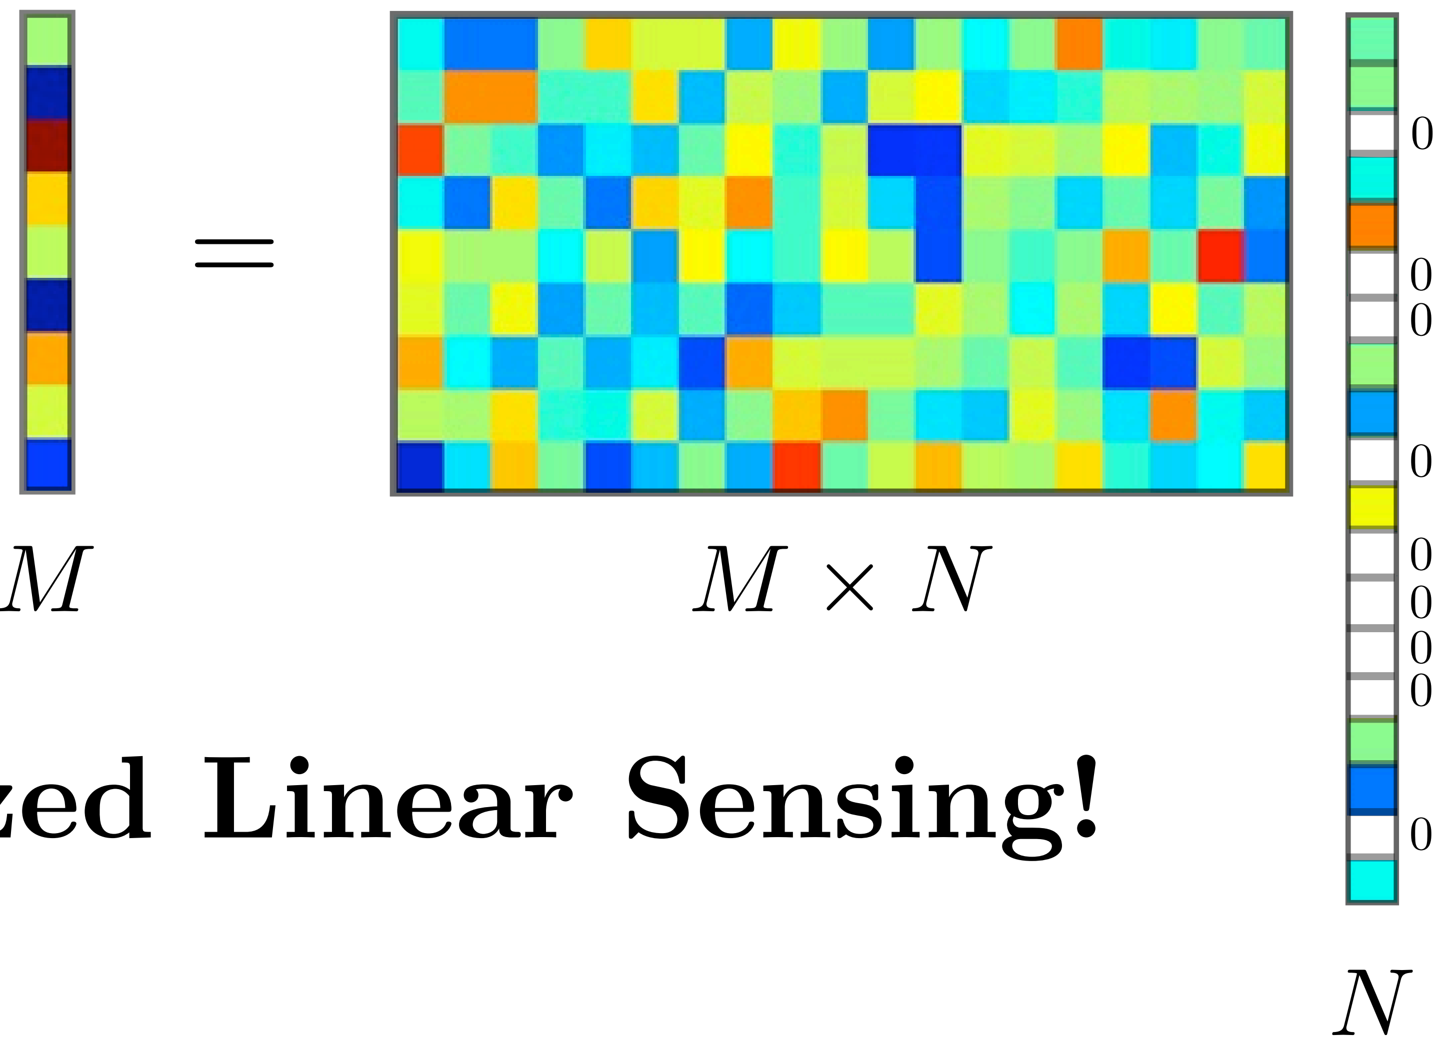

Sparsity Prior ($\Psi = \mathrm{Id}$)

A signal in this discrete world

Generalized Linear Sensing!

Compressed Sensing

Then ...

M questions — Sensing method — Signal

$\boldsymbol{y}$ — $\boldsymbol{\Phi}$ — $\boldsymbol{x}$

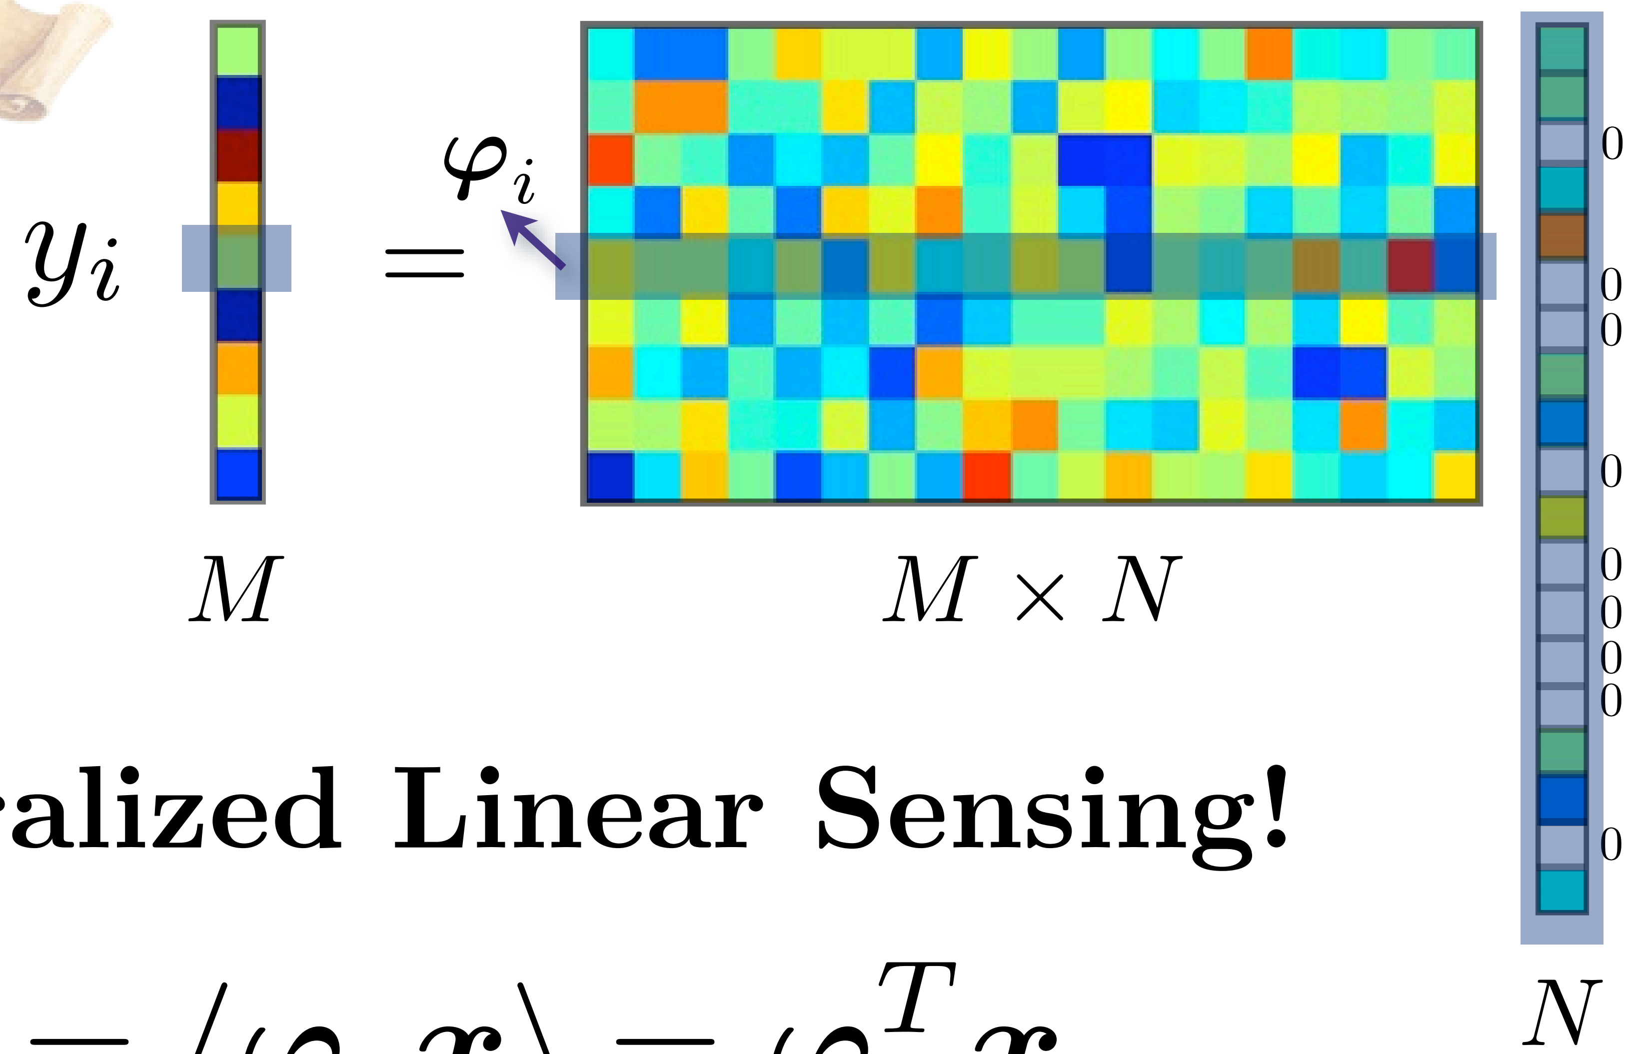

Sparsity Prior ($\Psi = \mathrm{Id}$)

A signal in this discrete world

Generalized Linear Sensing!

$$y_i = \langle \boldsymbol{\varphi}, \boldsymbol{x} \rangle = \boldsymbol{\varphi}^T \boldsymbol{x}$$

$$1 \leq i \leq M$$

Compressed Sensing

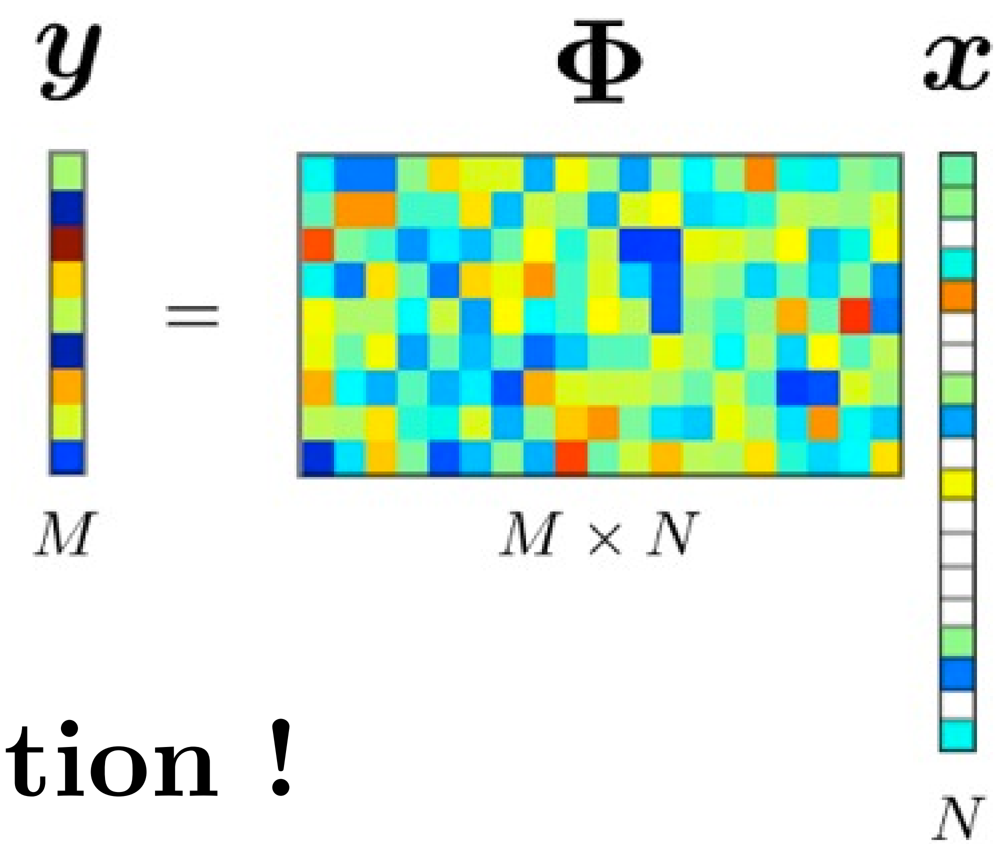

+ Non-linear reconstruction !

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

$$\boldsymbol{x}^* = \underset{\boldsymbol{u} \in \mathbb{R}^N}{\arg\min} \ \|\boldsymbol{u}\|_0 \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{u}$$

$$\|\boldsymbol{u}\|_0 = \#\{j : u_j \neq 0\}$$

Compressed Sensing

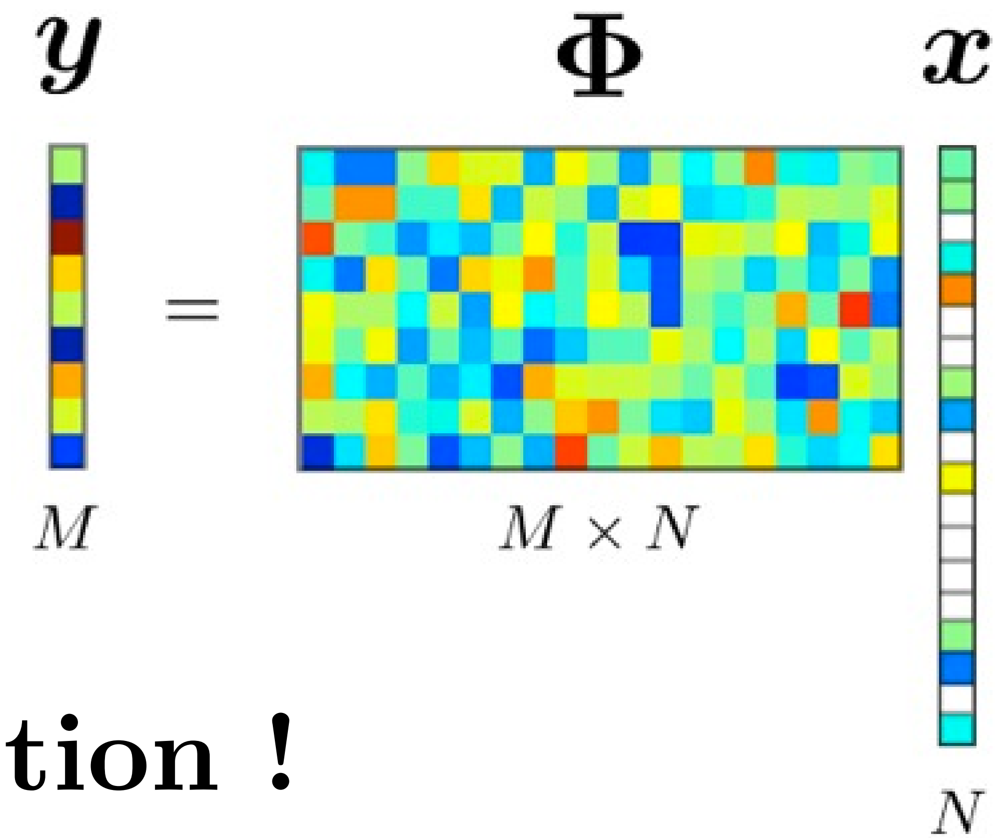

$+$ **Non-linear reconstruction !**

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned"
then:

(relax.)

$$\boldsymbol{x}^* = \underset{\boldsymbol{u} \in \mathbb{R}^N}{\arg\min} \ \|\boldsymbol{u}\|_{\not{0}\ 1} \ \text{s.t.}\ \boldsymbol{y} = \boldsymbol{\Phi u}$$

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Compressed Sensing

Simplifying assumption

$\exists\, \delta \in (0,1)$ Restricted Isometry Property

$$\sqrt{1-\delta}\,\|\boldsymbol{v}\|_2 \ \leqslant\ \|\boldsymbol{\Phi v}\|_2 \ \leqslant\ \sqrt{1+\delta}\,\|\boldsymbol{v}\|_2$$

for all $2K$ sparse signals $\boldsymbol{v}$.

any subset of $2K$ columns is an *isometry*

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u}\in\mathbb{R}^N} \|\boldsymbol{u}\|_{\cancel{0}\,1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

if $\delta < \sqrt{2}-1$ [Candes 08]

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$ (Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Compressed Sensing

Simplifying assumption

$\exists\, \delta \in (0,1)$ Restricted Isometry Property

$$\sqrt{1-\delta}\,\|\boldsymbol{v}\|_2 \leqslant \|\boldsymbol{\Phi v}\|_2 \leqslant \sqrt{1+\delta}\,\|\boldsymbol{v}\|_2$$

for all $2K$ sparse signals $\boldsymbol{v}$.

any subset of $2K$ columns is an *isometry*

Examples:

+ Gaussian
+ Bernoulli
+ Random Fourier
+ ....

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u}\in\mathbb{R}^N} \|\boldsymbol{u}\|_{\cancel{0}\;1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

if $\delta < \sqrt{2}-1$ [Candes 08]

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Compressed Sensing

Simplifying assumption

$\exists\ \delta \in (0,1)$ Restricted Isometry Property

$$\sqrt{1-\delta}\,\|\boldsymbol{v}\|_2 \leqslant \|\boldsymbol{\Phi v}\|_2 \leqslant \sqrt{1+\delta}\,\|\boldsymbol{v}\|_2$$

for all $2K$ sparse signals $\boldsymbol{v}$.

any subset of $2K$ columns is an *isometry*

Examples:

+ Gaussian
+ Bernoulli
+ Random Fourier
+ ....

$$M = O(K \log N/K) \ll N$$

$\boldsymbol{\Phi} \in \mathbb{R}^{M\times N}$, $\Phi_{ij} \sim_{\rm iid} \mathcal{N}(0,1)$

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u}\in\mathbb{R}^N} \|\boldsymbol{u}\|_{\not 0\ 1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

if $\delta < \sqrt{2} - 1$ [Candes 08]

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Compressed Sensing

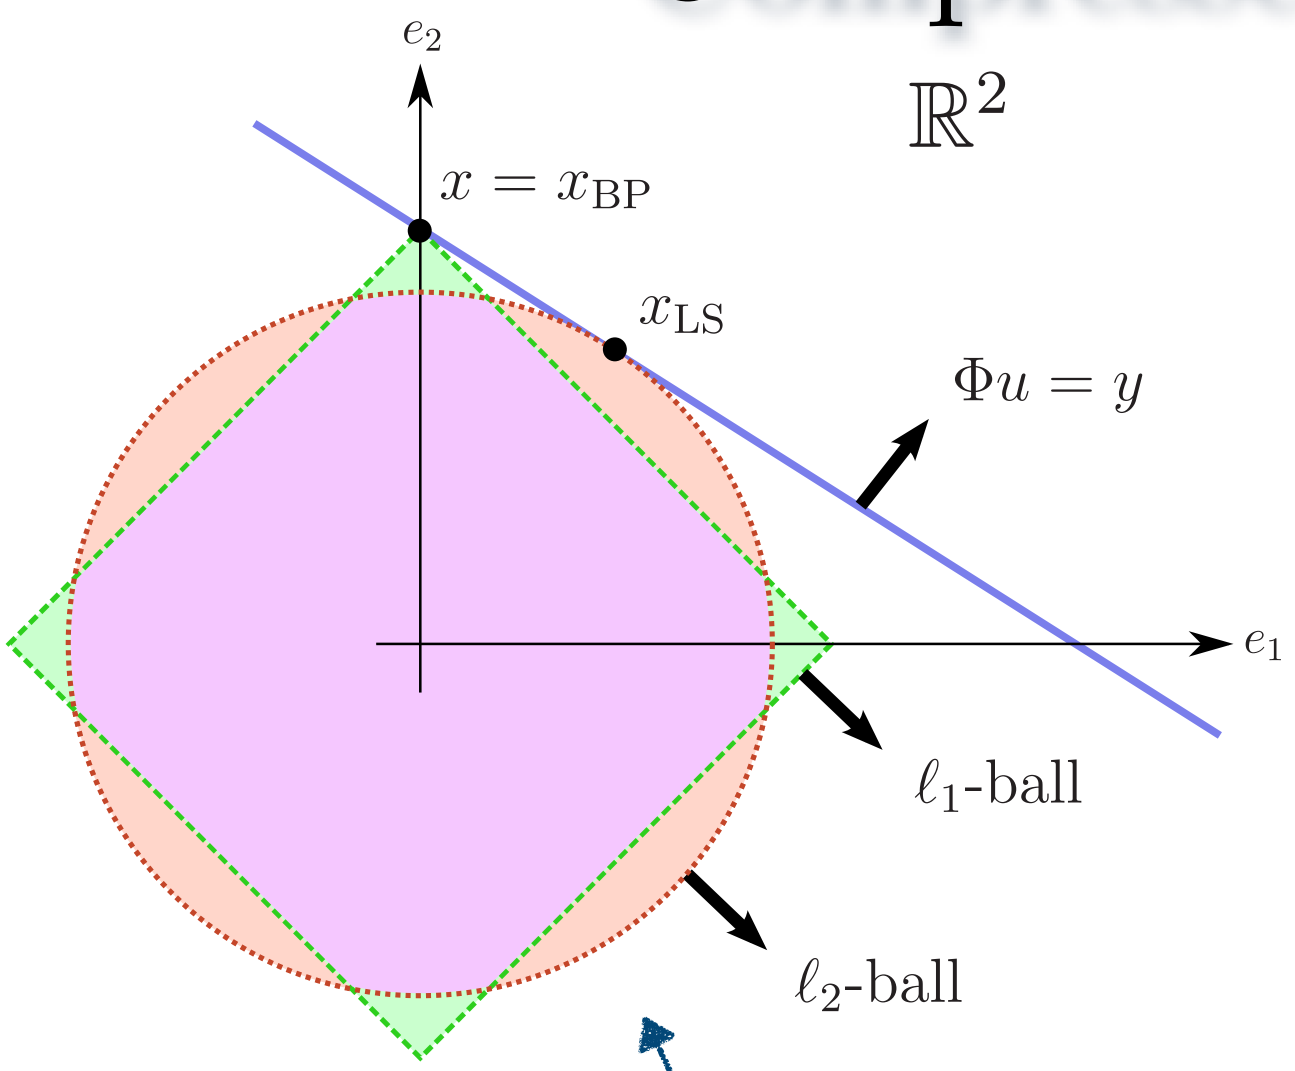

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_{0 \to 1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Solvers:

Linear Programming,
Interior Point Method,
Proximal Methods,
... **Tons** of toolboxes ...

Compressed Sensing

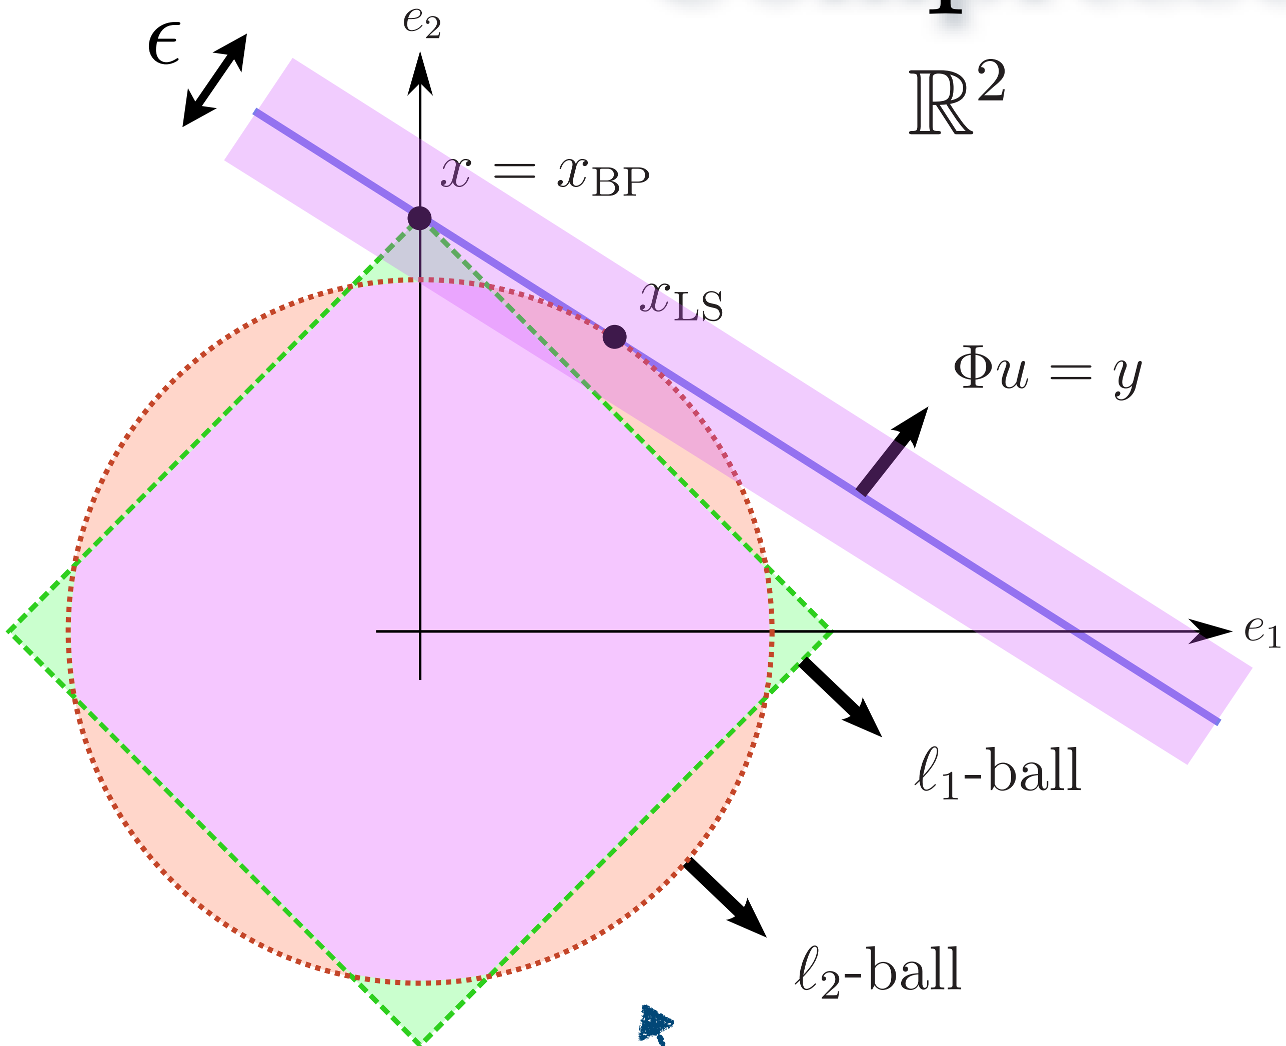

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned"

then: $\|\boldsymbol{y} - \boldsymbol{\Phi u}\| \leqslant \epsilon$

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_{\cancel{2}\,1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Solvers:

Linear Programming,
Interior Point Method,
Proximal Methods,
... **Tons** of toolboxes ...

Compressed Sensing

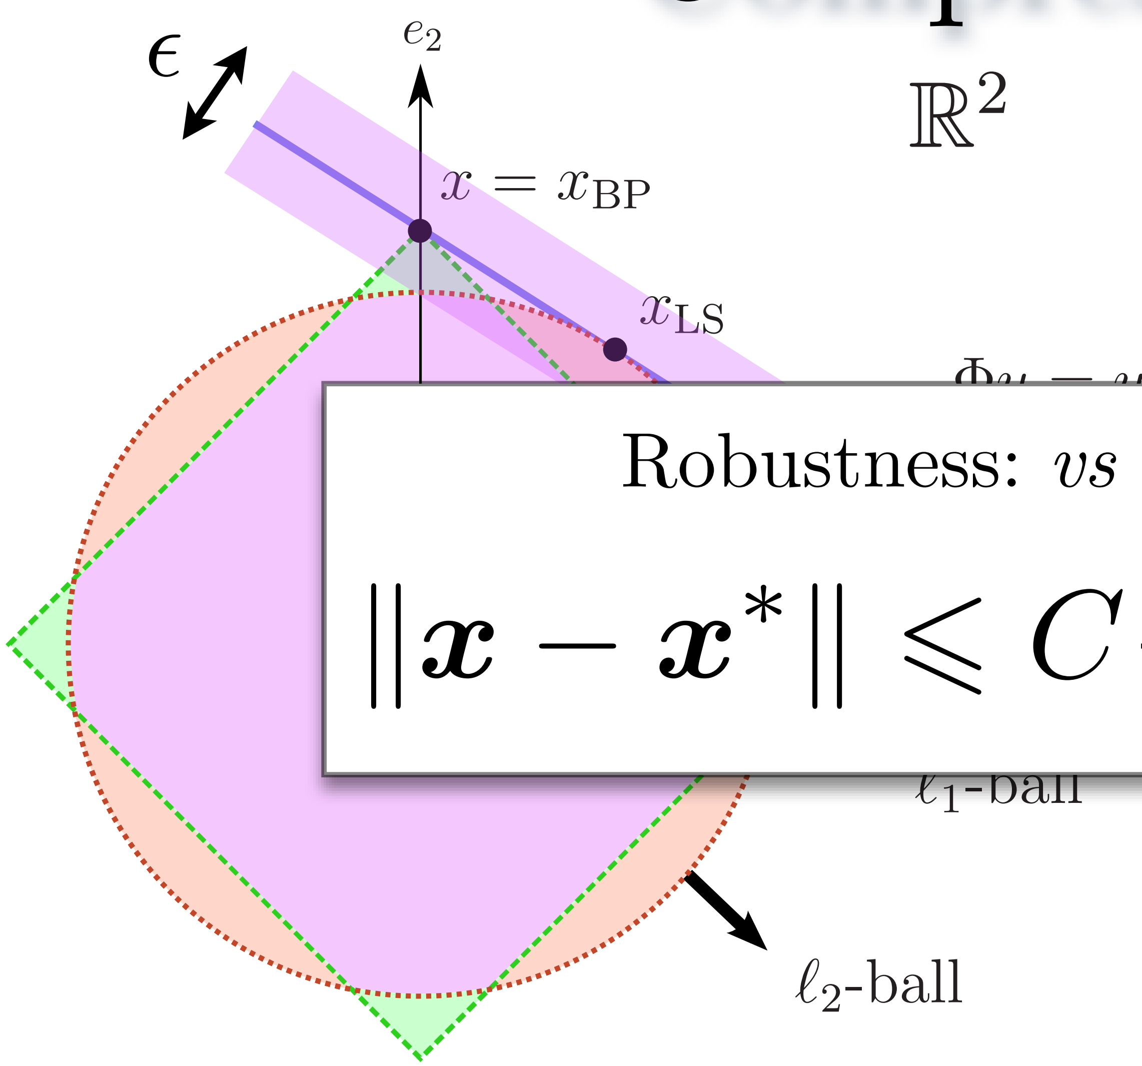

Robustness: *vs* sparse deviation + noise.

$$\|\boldsymbol{x} - \boldsymbol{x}^*\| \leqslant C \frac{1}{\sqrt{K}} \|\boldsymbol{x} - \boldsymbol{x}_K\|_1 + D\epsilon$$

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

$$\|\boldsymbol{y} - \boldsymbol{\Phi u}\| \leqslant \epsilon$$

(relax.)

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_{\cancel{2}\,1} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi u}$$

$\|\boldsymbol{u}\|_1 = \sum_j |u_j|$

(Basis Pursuit) [Chen, Donoho, Saunders, 1998]

Solvers:

Linear Programming,
Interior Point Method,
Proximal Methods,
... **Tons** of toolboxes ...

Part 1

When quantization meets compressed sensing

Outline:

1. Context
2. Former QCS methods and performance limits
3. Consistent Reconstructions
4. Sigma-Delta quantization in CS
5. To saturate or not? And how much?

1. Context

What is quantization?

- Generality:

  Intuitively: "*Quantization maps a bounded continuous domain to a set of finite elements (or codebook)*"

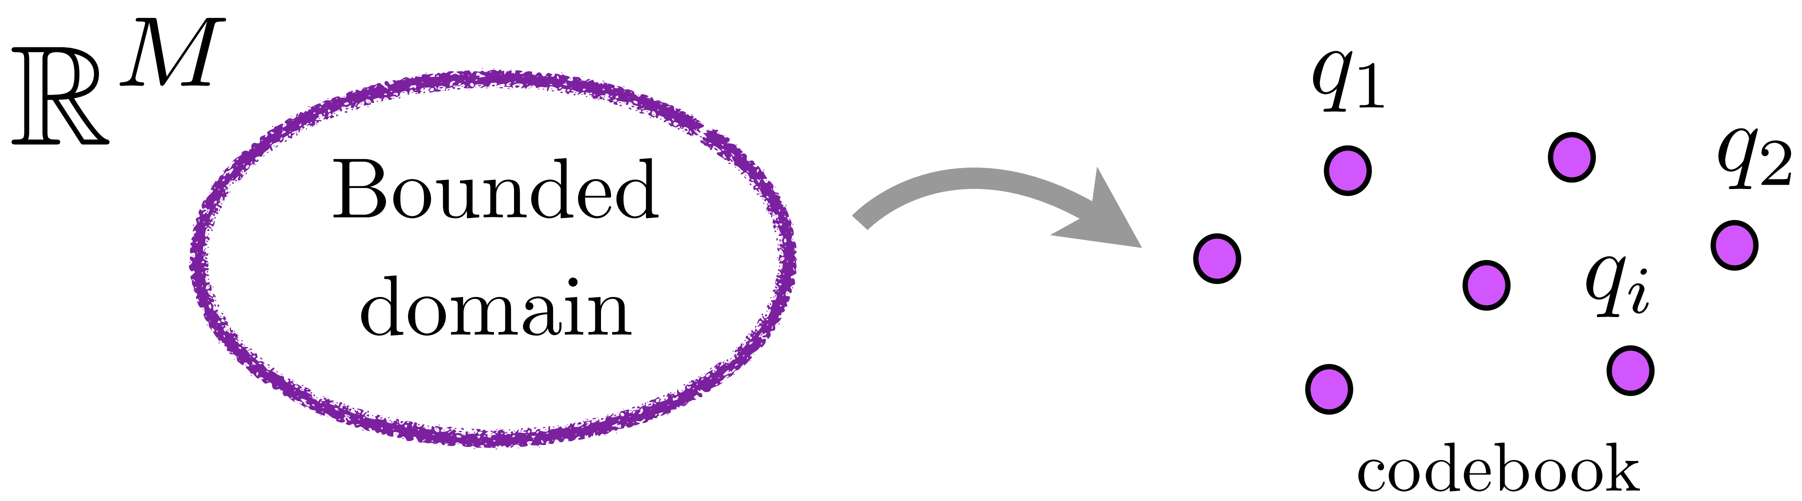

$$\mathcal{Q}[x] \in \{q_1, q_2, \cdots\}$$

- Oldest example: rounding off $\lfloor x \rfloor, \lceil x \rceil, \ldots \quad \mathbb{R} \to \mathbb{Z}$

Example 1: scalar quantization

- In $\mathbb{R}^M$, on each component of M- dimensional vectors:

$$\Omega = \{q_i \in \mathbb{R} : 1 \leqslant i \leqslant 2^B\}, \qquad \text{(levels)} \quad \square$$

$$\mathcal{T} = \{t_i \in \overline{\mathbb{R}} : 1 \leqslant i \leqslant 2^B + 1, t_i \leqslant t_{i+1}\} \quad \text{(thresholds)} \quad \bullet$$

$$\forall \lambda \in \mathbb{R}, \qquad \mathcal{Q}[\lambda] = q_i \;\Leftrightarrow\; \lambda \in \mathcal{R}_i \triangleq [t_i, t_{i+1}), \;\; \text{1-D quantization cell}$$

$$\forall u \in \mathbb{R}^M, \quad (\mathcal{Q}[u])_j = \mathcal{Q}[u_j]$$

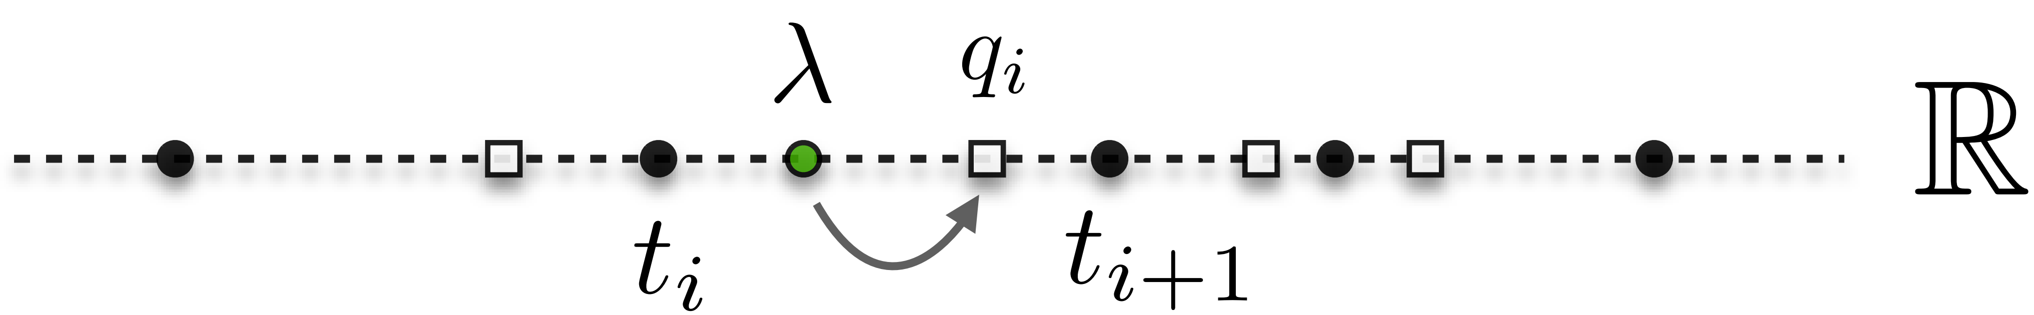

<u>other names:</u>

Pulse Code Modulation - PCM
Memoryless Scalar Quantization - MSQ

Example 1: scalar quantization

- In $\mathbb{R}^M$, on each component of M- dimensional vectors:

$$\Omega = \{q_i \in \mathbb{R} : 1 \leqslant i \leqslant 2^B\}, \qquad \text{(levels)} \quad \square$$

$$\mathcal{T} = \{t_i \in \overline{\mathbb{R}} : 1 \leqslant i \leqslant 2^B + 1, t_i \leqslant t_{i+1}\} \quad \text{(thresholds)} \quad \bullet$$

$$\forall \lambda \in \mathbb{R}, \qquad \mathcal{Q}[\lambda] = q_i \;\Leftrightarrow\; \lambda \in \mathcal{R}_i \triangleq [t_i, t_{i+1}), \; \text{1-D quantization cell}$$

$$\forall u \in \mathbb{R}^M, \quad (\mathcal{Q}[u])_j = \mathcal{Q}[u_j]$$

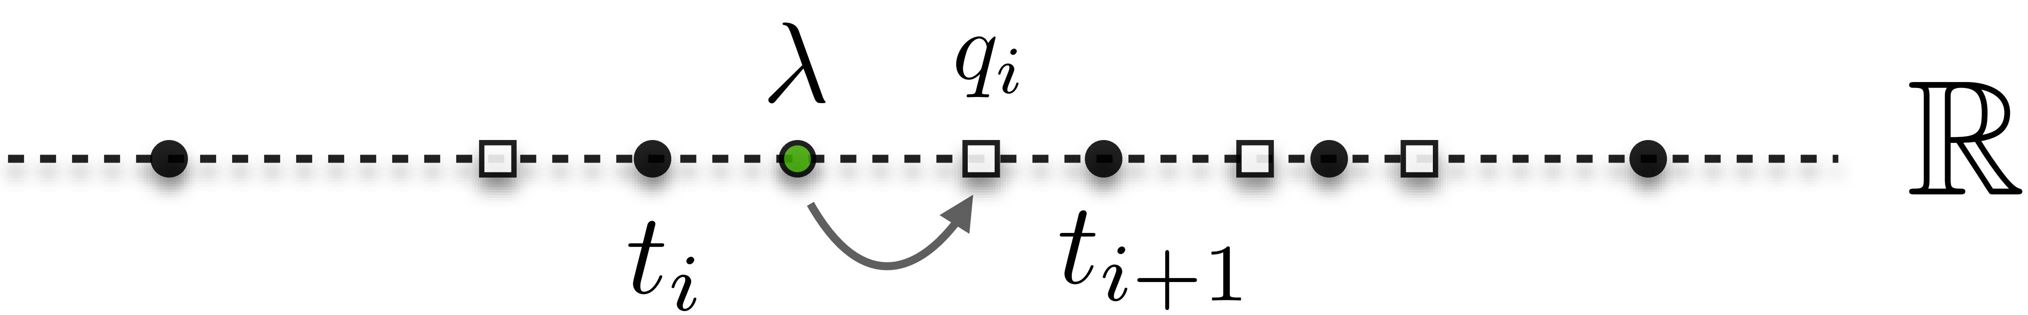

- Globally:

$$\mathcal{Q}[\boldsymbol{z}] = \boldsymbol{q} \in \Omega^M \;\Leftrightarrow\; \boldsymbol{z} \in$$

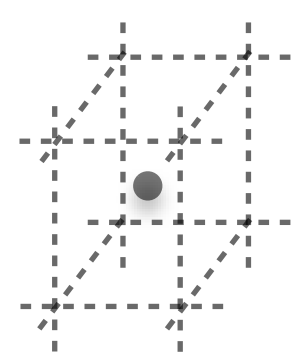

M - D quantization cell

$$\mathcal{R}_{i_1} \times \mathcal{R}_{i_2} \times \cdots \times \mathcal{R}_{i_M}$$

$$:= \mathcal{Q}^{-1}[\boldsymbol{q}]$$

Example 1: scalar quantization

- Regular uniform

$$q_k = (k + 1/2)\alpha$$
$$t_k = k\alpha$$

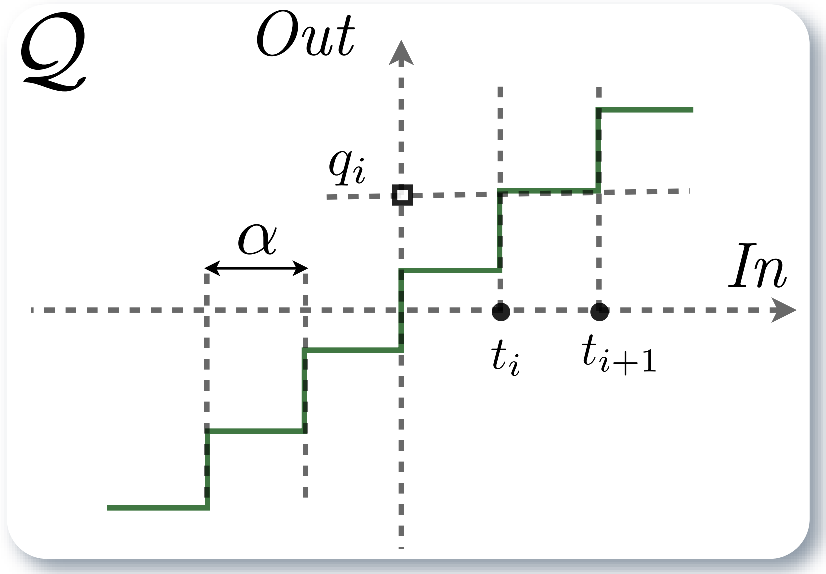

Example 1: scalar quantization

- Regular uniform

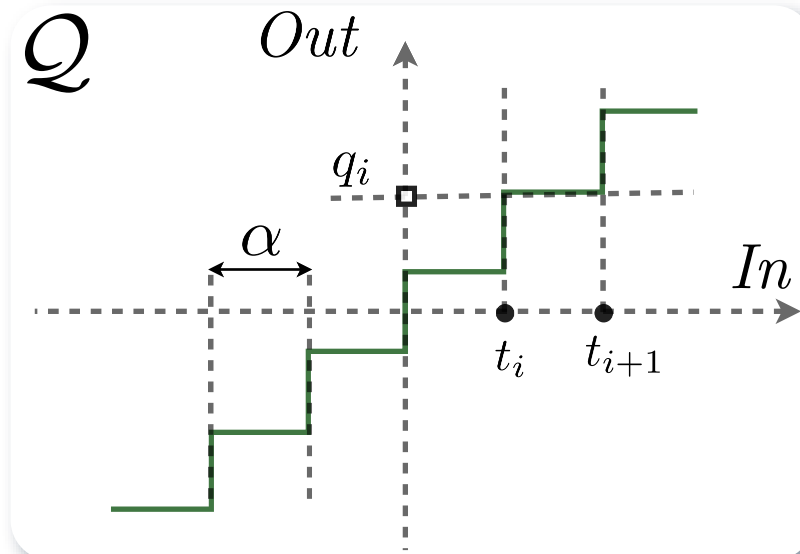

$$q_k = (k + 1/2)\alpha$$
$$t_k = k\alpha$$

- Regular non-uniform

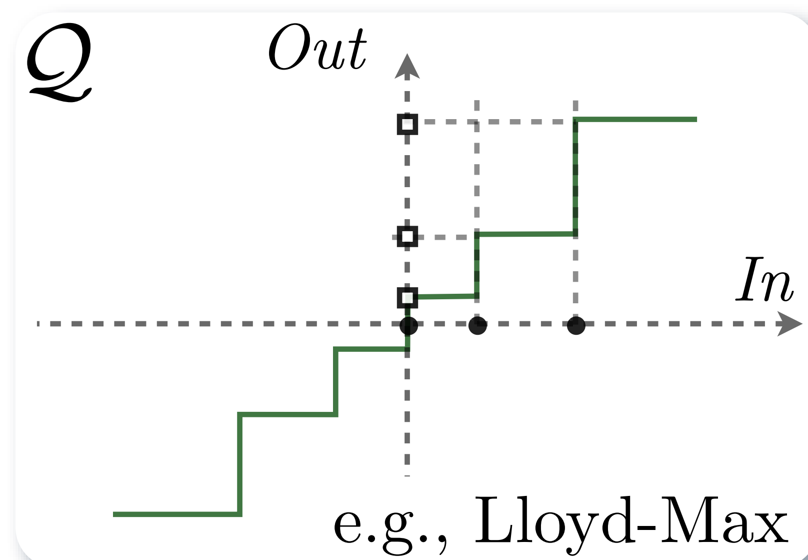

Ω and $\mathcal{T}$ optimized

e.g., wrt an input distribution Z
find minimum distortion, *i.e.*,

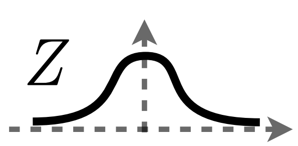

$$\underset{\mathcal{T},\Omega}{\text{argmin}}\ \mathbb{E}_Z \|Z - \mathcal{Q}[Z]\|^2$$

- $\exists$ Non-regular (P. Boufounos)

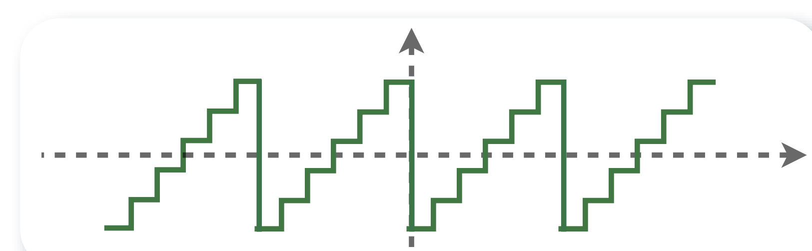

Example 2: vector quantization

(**caveat**: not really covered in this tutorial, ... except ΣΔ, see later)

Quantization = codebook $\boldsymbol{\Omega}$ + quantization cells $\boldsymbol{\mathcal{R}} = \{\mathcal{R}_i \subset \mathbb{R}^M\}$

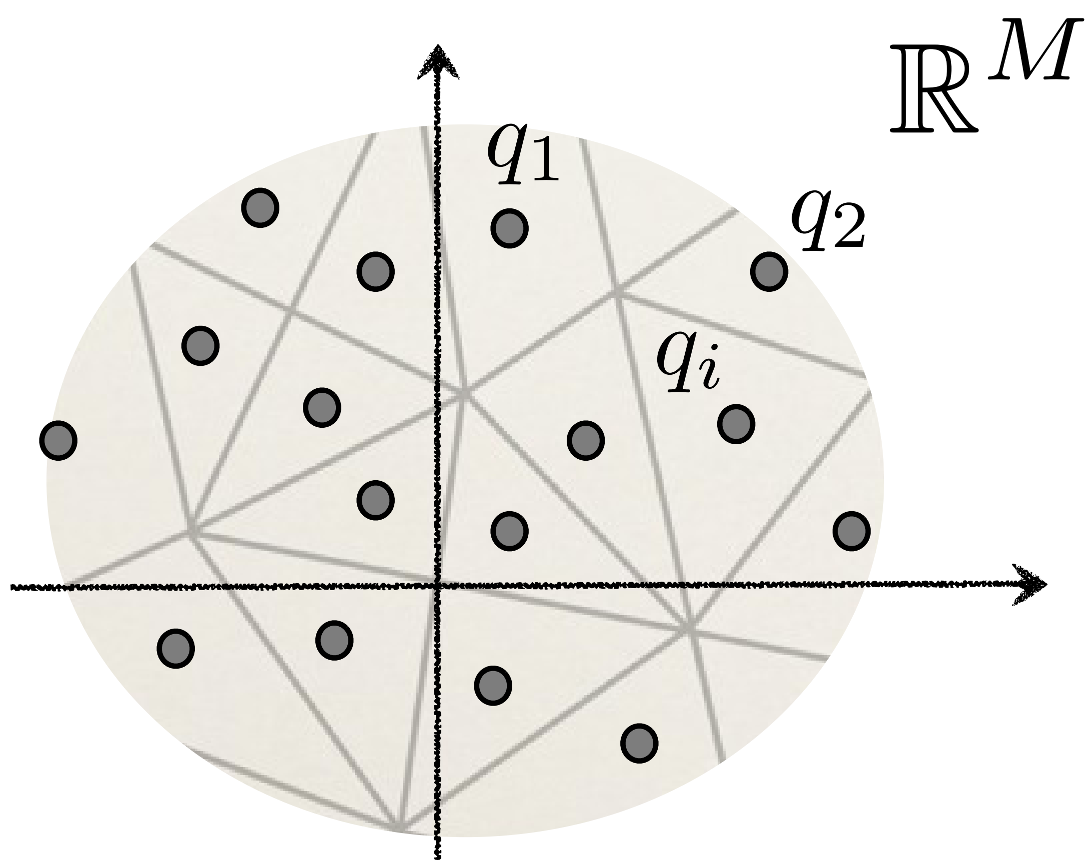

(non-separable quantization)

e.g., $\underset{\boldsymbol{\Omega},\boldsymbol{\mathcal{R}}}{\operatorname{argmin}}\ \mathbb{E}_Z \|\boldsymbol{Z} - \mathcal{Q}[\boldsymbol{Z}]\|^2$

Classical Sampling and Quantization

For acquisition:

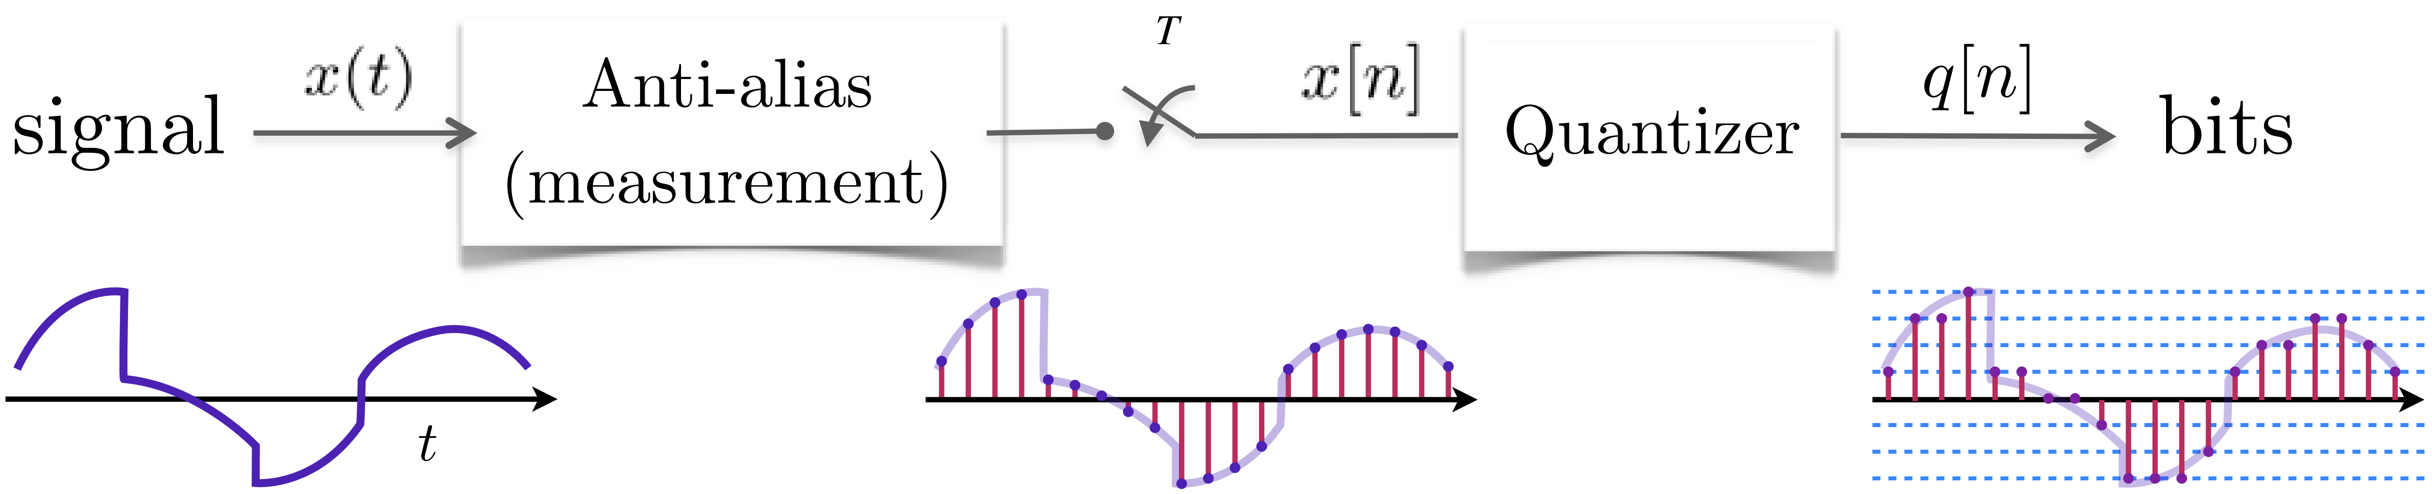

Classical Sampling and Quantization

For acquisition:

signal $\xrightarrow{x(t)}$ Anti-alias (measurement) — T — $\xrightarrow{x[n]}$ Quantizer $\xrightarrow{q[n]}$ bits

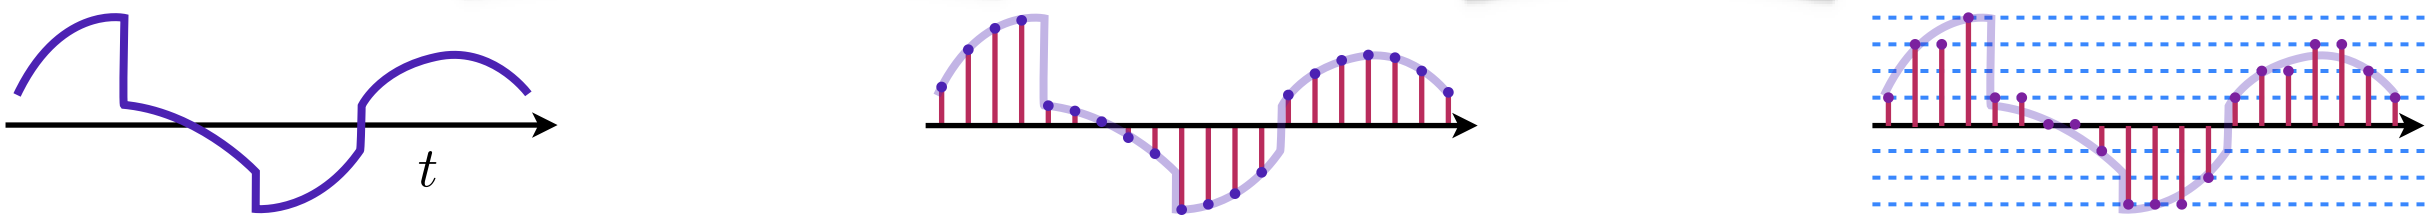

For reconstruction:

$\xrightarrow{q[n]}$ Low-pass Filter (linear reconstruction) $\xrightarrow{\widehat{x}(t)}$

Classical Sampling and Quantization

For acquisition:

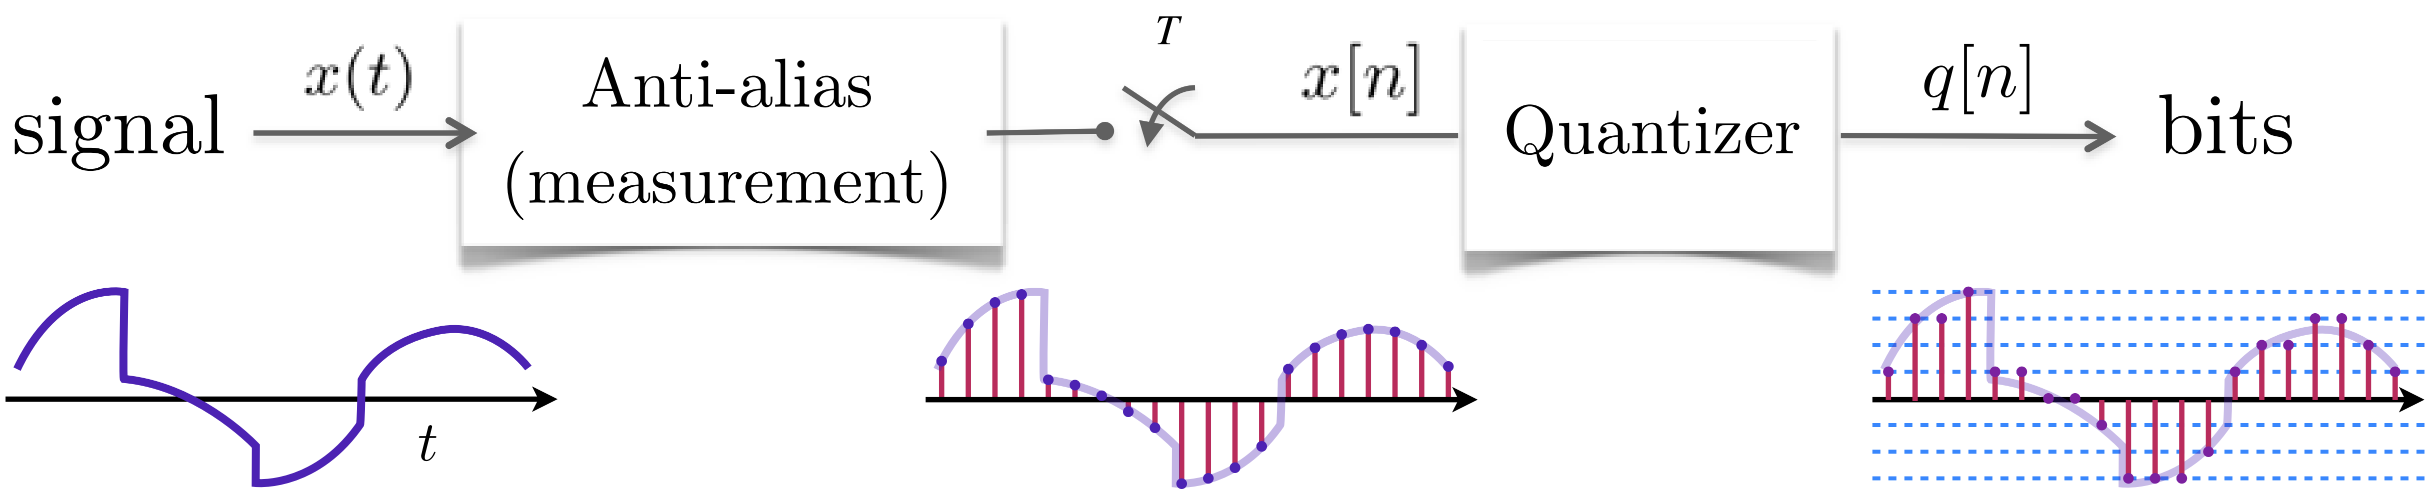

For reconstruction:

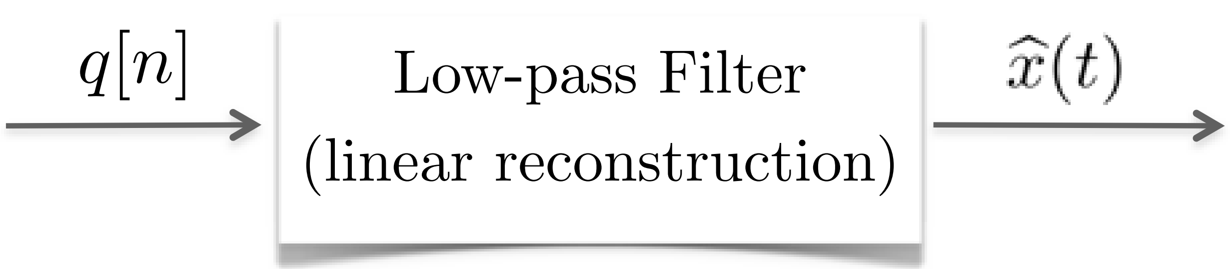

Sampling: discretization in time ⇨ Lossless at the Nyquist rate

Quantization: discretization in amplitude ⇨ Always lossy

Need both for digital data acquisition

Compressive Sampling and Quantization

Compressed sensing theory says:

"Linearly sample a signal

at a rate function of

its intrinsic dimensionality"

Information theory and sensor designer say:

"Okay, but I need to

quantize/digitize my measurements!"

(e.g., in ADC)

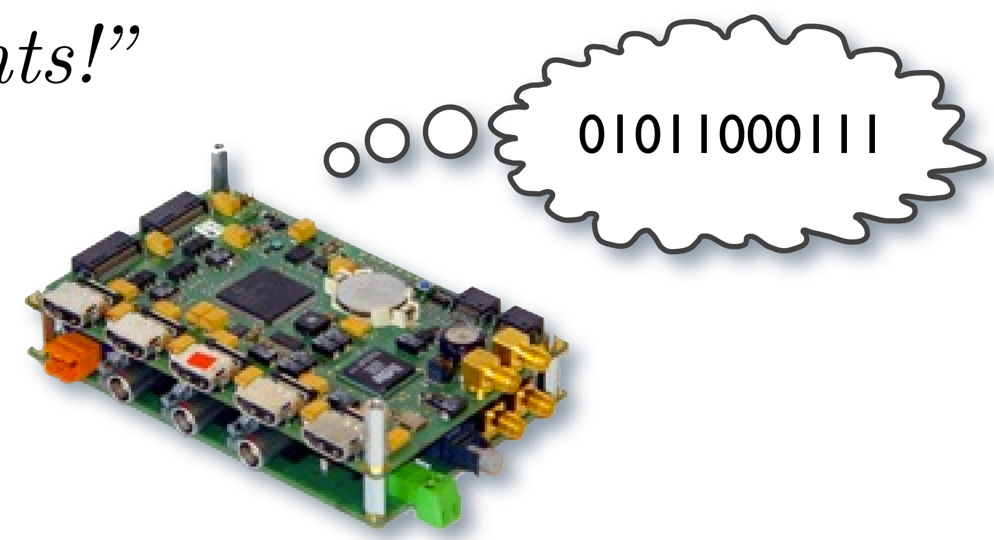

The Quantized CS Problem (QCS)

Natural questions:

- How to integrate quantization in CS?
- What do we loose?
- Are they some theoretical limitations? (related to information theory? geometry?)
- How to minimize quantization effects in the reconstruction?

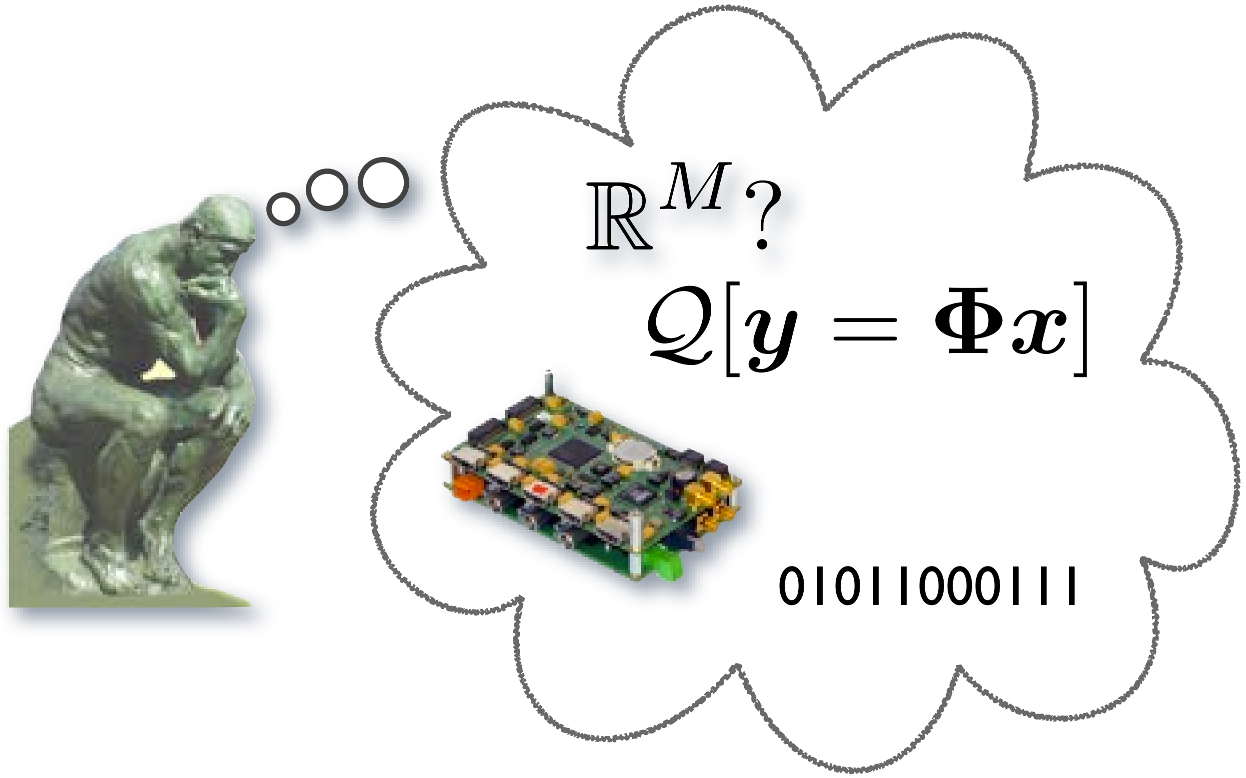

QCS: a system view

With no additional noise:

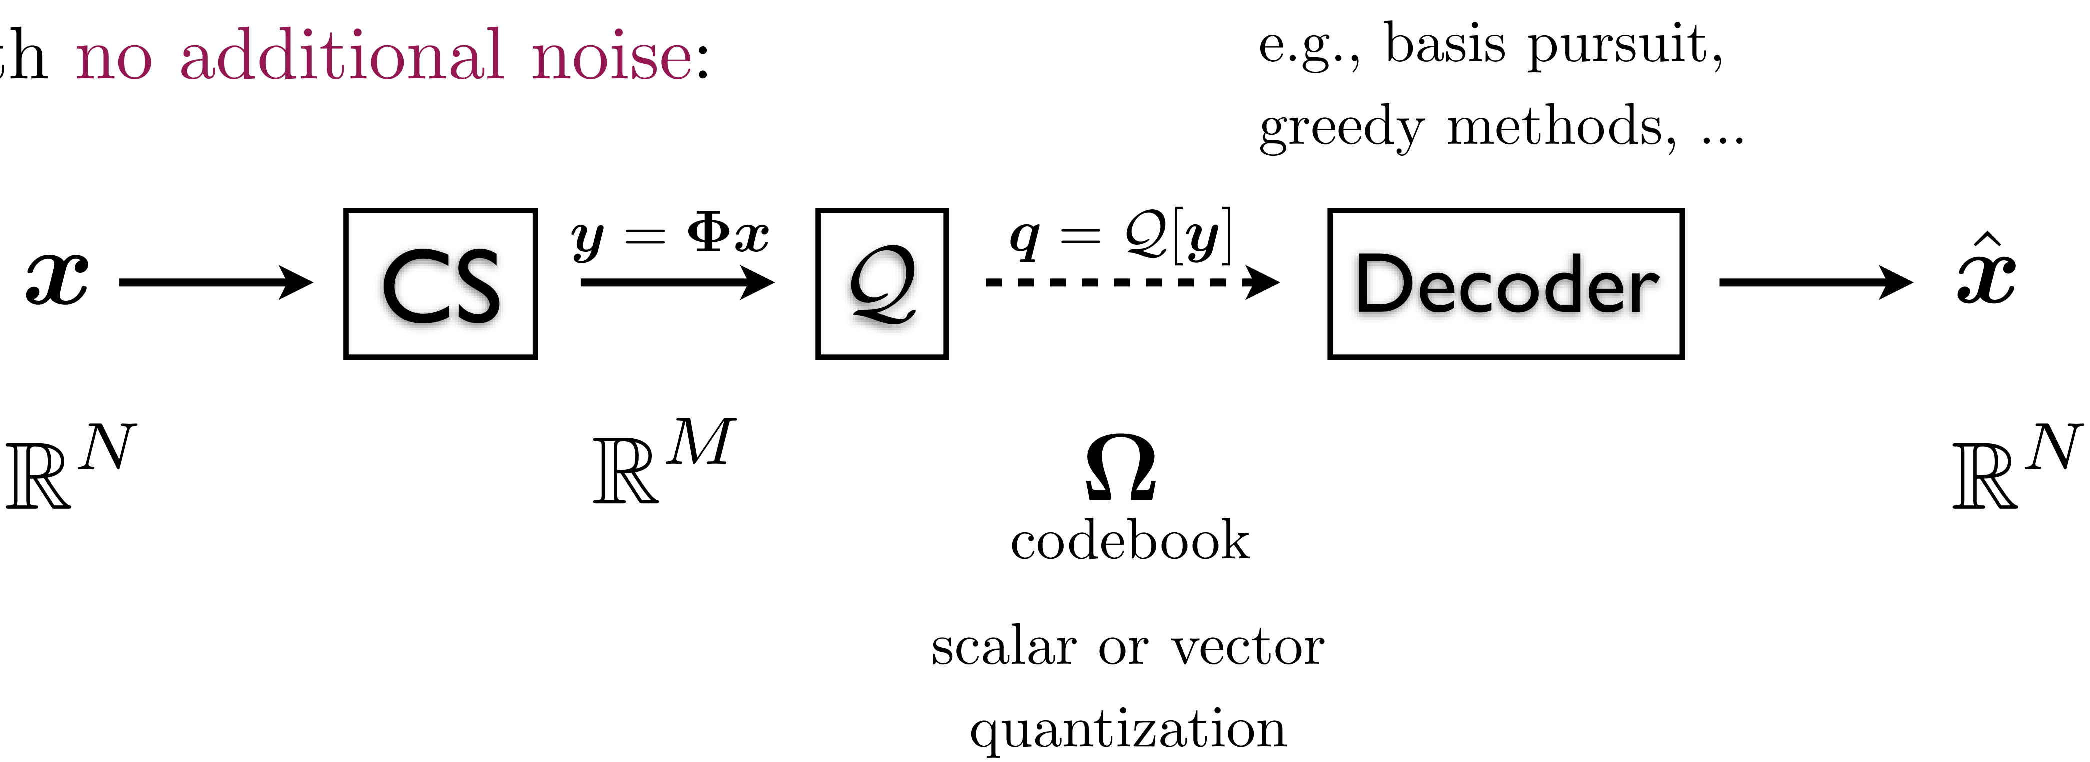

QCS: a system view

With no additional noise:

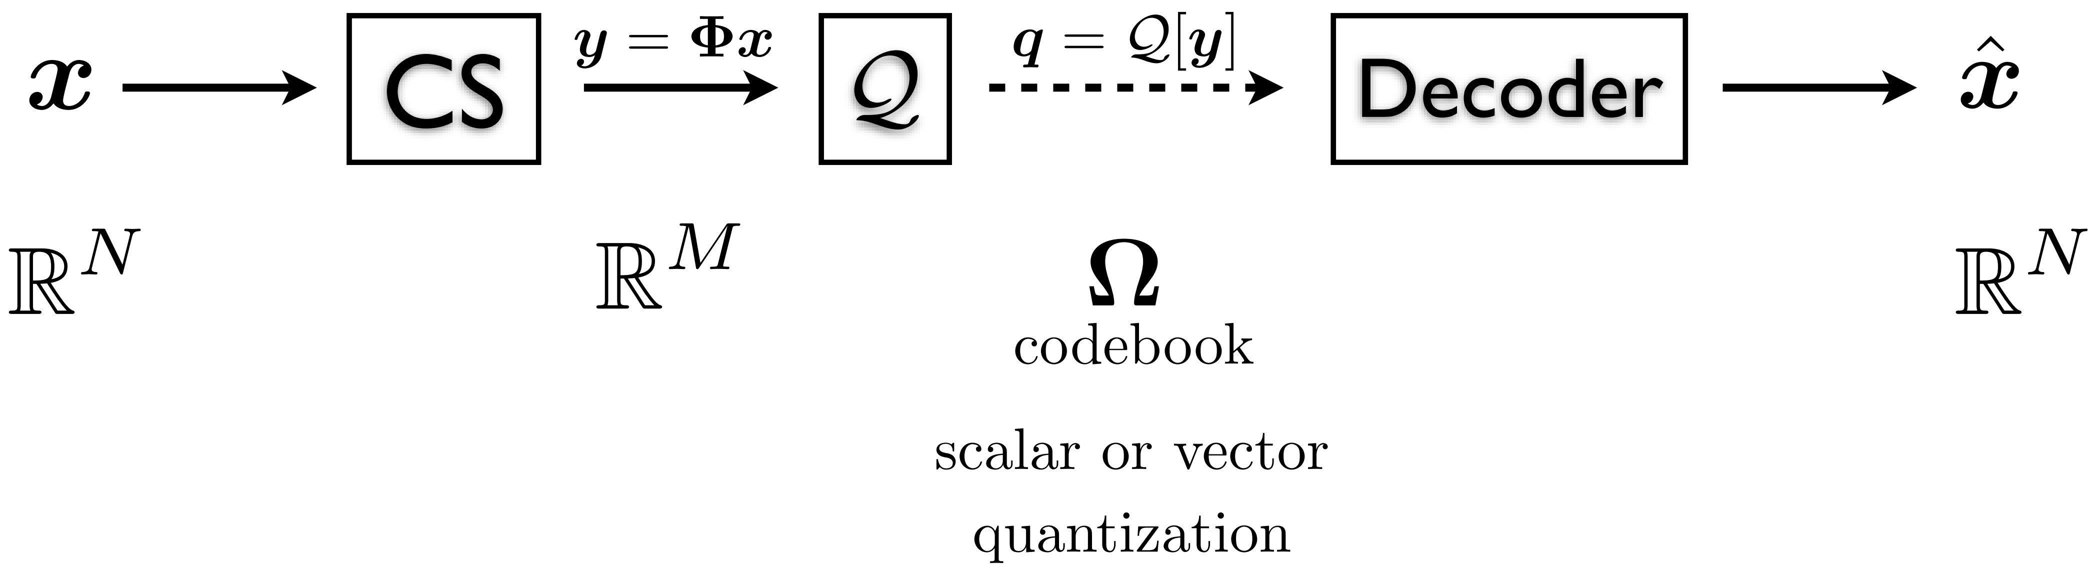

Finite codebook $\Rightarrow \hat{\boldsymbol{x}} \neq \boldsymbol{x}$

(i.e., impossibility to encode continuous domain in a finite number of elements)

QCS: a system view

With no additional noise:

e.g., basis pursuit,
greedy methods, ...

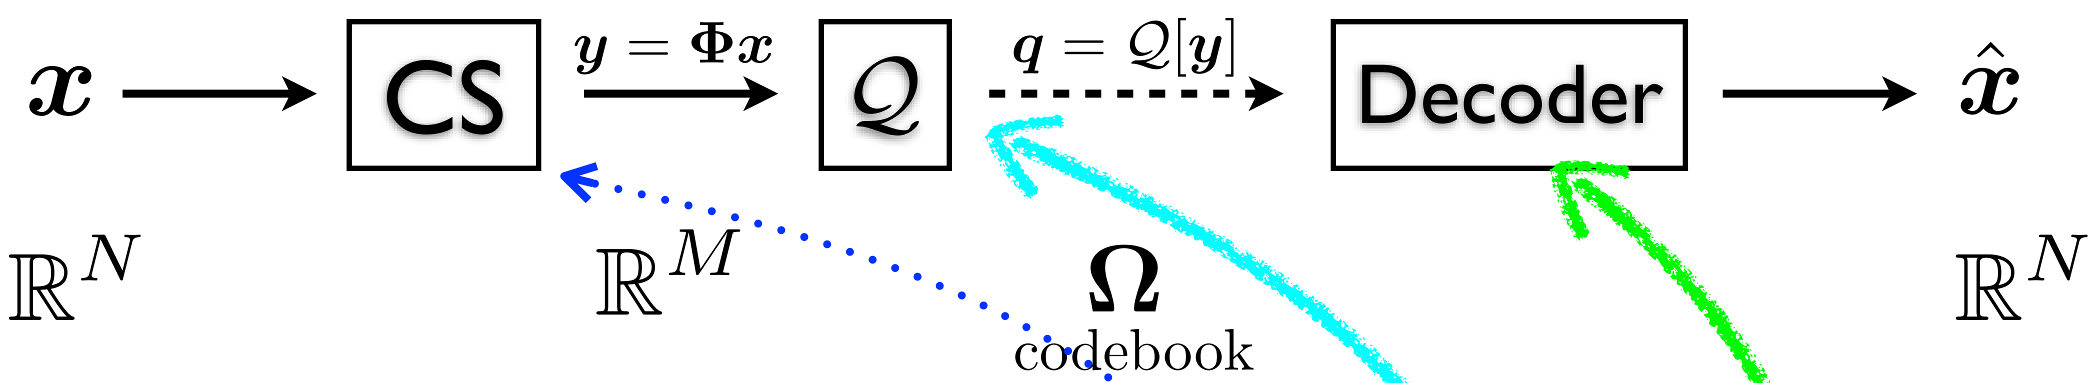

Finite codebook $\Rightarrow \hat{x} \neq x$

Objective: Minimize $\|\hat{x} - x\|$

given a certain number of:

bits, measurements, or bits/meas.

How?

Where to act?

Change CS, Q or decoder?

Some of them? all?

2. Former QCS methods and performance limits

Scalar quantization in CS

Turning measurements into bits $\rightarrow$ scalar quantization

$$q_i = \mathcal{Q}[(\boldsymbol{\Phi x})_i] = \mathcal{Q}[\langle \boldsymbol{\phi}_i, \boldsymbol{x} \rangle] \in \Omega \subset \mathbb{R}$$

$$\boldsymbol{q} \;=\; \mathcal{Q}\big[\boldsymbol{\Phi x}\big] \in \boldsymbol{\Omega} = \Omega^M,$$

Important points:

- Definition of $\boldsymbol{\Phi}$ independent of M (e.g., $\Phi_{ij} \sim_{\rm iid} \mathcal{N}(0,1)$)
  $\rightarrow$ preserves measurement dynamic!
- B bits per measurement
- Total bit budget: $R = BM$
- No further encoding (e.g., entropic)

Former solution (Candès, Tao, ...)

- Quantization is like a noise

$$\boldsymbol{q} \;=\; \mathcal{Q}\big[\boldsymbol{\Phi x}\big] = \boldsymbol{\Phi x} + \boldsymbol{n}$$

quantization distortion

Former solution (Candès, Tao, ...)

- Quantization is like a noise

$$\boldsymbol{q} \;=\; \mathcal{Q}\big[\boldsymbol{\Phi x}\big] = \boldsymbol{\Phi x} + \boldsymbol{n}$$

and CS is robust (e.g., with *basis pursuit denoise*)

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{u}\in\mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{\Phi u} - \boldsymbol{q}\| \leqslant \epsilon \quad \text{(BPDN)}$$

$\ell_2 - \ell_1$ instance optimality:

If $\|\boldsymbol{n}\| \leqslant \epsilon$ and $\frac{1}{\sqrt{M}}\boldsymbol{\Phi}$ is RIP$(\delta, 2K)$ with $\delta \leqslant \sqrt{2}-1$, then

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \;\leqslant\; C\frac{\epsilon}{\sqrt{M}} + D\, e_0(K),$$

for some $C, D > 0$ and $e_0(K) = \|\boldsymbol{x} - \boldsymbol{x}_K\|_1/\sqrt{K}$.

Former solution (Candès, Tao, ...)

- Quantization is like a noise

$$\boldsymbol{q} \;=\; \mathcal{Q}\big[\boldsymbol{\Phi x}\big] = \boldsymbol{\Phi x} + \boldsymbol{n}$$

and CS is robust (e.g., with *basis pursuit denoise*)

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{u}\in\mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{\Phi u} - \boldsymbol{q}\| \leqslant \epsilon \quad \text{(BPDN)}$$

$\ell_2 - \ell_1$ instance optimality:

If $\|\boldsymbol{n}\| \leqslant \epsilon$ and $\frac{1}{\sqrt{M}}\boldsymbol{\Phi}$ is RIP$(\delta, 2K)$ with $\delta \leqslant \sqrt{2} - 1$, then

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \;\leqslant\; C\frac{\epsilon}{\sqrt{M}} + D\, e_0(K),$$

for some $C, D > 0$ and $e_0(K) = \|\boldsymbol{x} - \boldsymbol{x}_K\|_1/\sqrt{K}$.

How to find it?

Former solution (Candès, Tao, ...)

1. For uniform quantization, by construction:

ϵ?

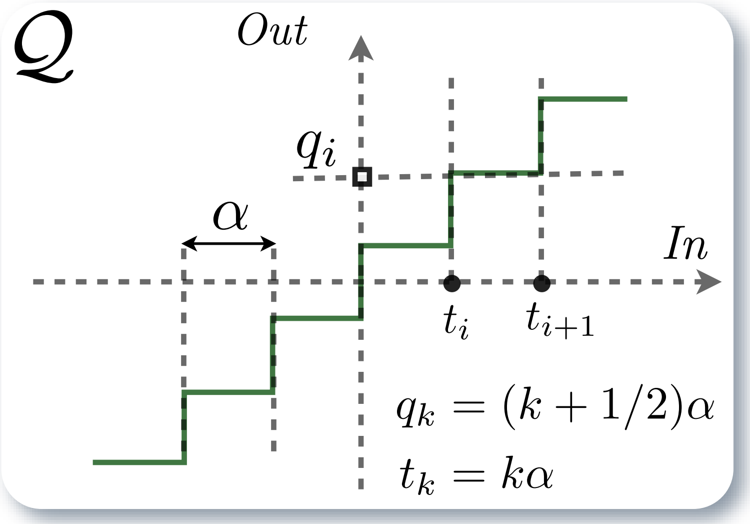

$$\begin{aligned} n_i &= \mathcal{Q}[(\boldsymbol{\Phi x})_i] - (\boldsymbol{\Phi x})_i \\ &\in q_{k_i} - \mathcal{R}_{k_i} = [-\alpha/2, \alpha/2] \\ &\Rightarrow \|\boldsymbol{n}\|_\infty \leq \alpha/2 \end{aligned}$$

$$\Rightarrow \|\boldsymbol{n}\|^2 \leqslant M\|\boldsymbol{n}\|_\infty^2 \leqslant M\alpha^2/4$$

and plug this upper bound in BPDN

Former solution (Candès, Tao, ...)

1. For uniform quantization, by construction:

ϵ?

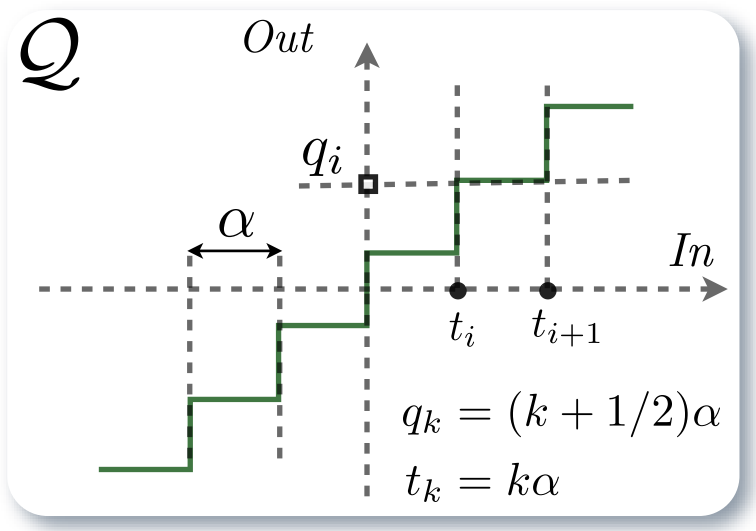

$$n_i = \mathcal{Q}[(\boldsymbol{\Phi x})_i] - (\boldsymbol{\Phi x})_i$$
$$\in q_{k_i} - \mathcal{R}_{k_i} = [-\alpha/2, \alpha/2]$$
$$\Rightarrow \|\boldsymbol{n}\|_\infty \leq \alpha/2$$

$$\Rightarrow \|\boldsymbol{n}\|^2 \leqslant M\|\boldsymbol{n}\|_\infty^2 \leqslant M\alpha^2/4$$

and plug this upper bound in BPDN

can be improved!

Former solution (Candès, Tao, ...)

2. For uniform quantization, uniform model!

ϵ?

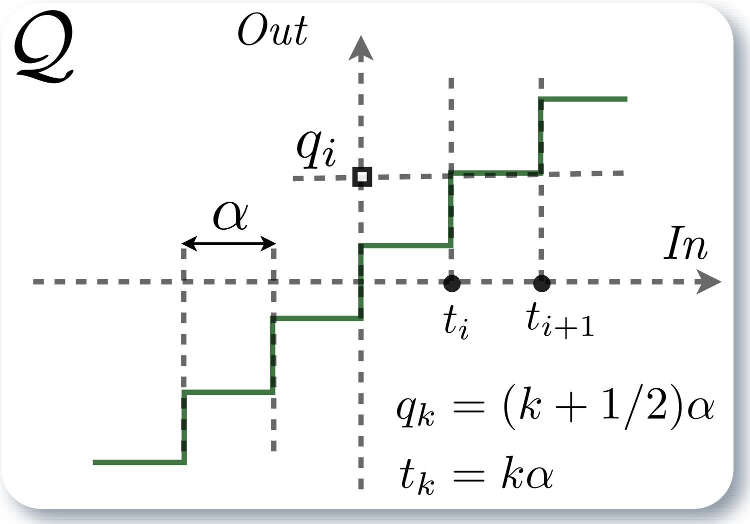

$$n_i = \mathcal{Q}[(\boldsymbol{\Phi x})_i] - (\boldsymbol{\Phi x})_i$$
$$\in q_{k_i} - \mathcal{R}_{k_i} = [-\alpha/2, \alpha/2]$$
$$\sim_{\text{iid}} \text{Uniform}([-\alpha/2, \alpha/2])$$

(HRA - high resolution assumption)

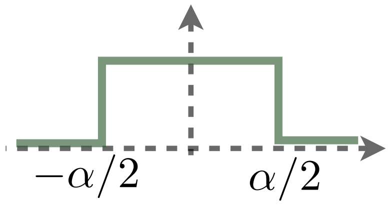

Former solution (Candès, Tao, ...)

2. For uniform quantization, uniform model!

ϵ?

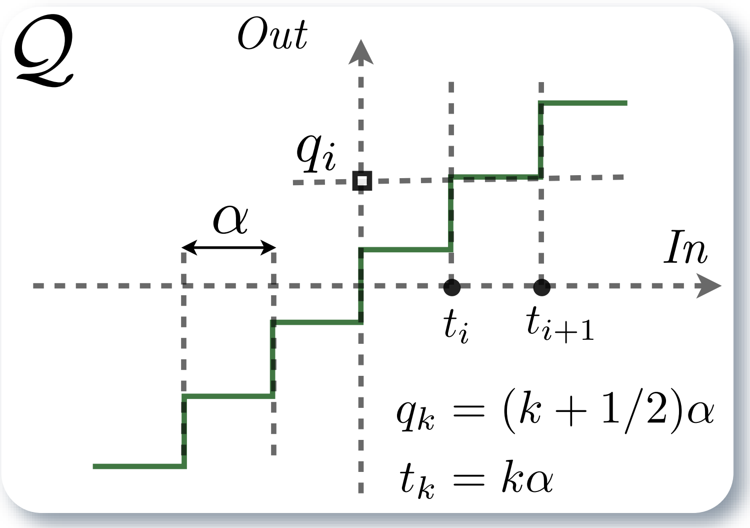

$$n_i = \mathcal{Q}[(\boldsymbol{\Phi x})_i] - (\boldsymbol{\Phi x})_i$$
$$\in q_{k_i} - \mathcal{R}_{k_i} = [-\alpha/2, \alpha/2]$$
$$\sim_{\rm iid} \mathrm{Uniform}([-\alpha/2, \alpha/2])$$

(HRA - high resolution assumption)

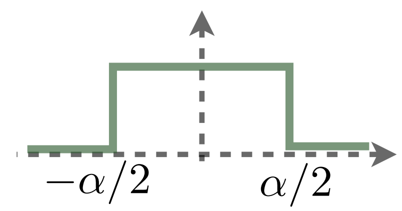

$$\Rightarrow \mathbb{E}|n_i|^2 = \alpha^2/12$$

$$\Rightarrow \|\boldsymbol{n}\|^2 \leqslant \mathbb{E}\|\boldsymbol{n}\|^2 + \kappa\sqrt{\mathrm{Var}\|\boldsymbol{n}\|^2}$$ (Chernoff-Hoeffding, bounded RVs)

$$\leqslant M\frac{\alpha^2}{12} + \kappa\sqrt{M}\frac{\alpha^2}{6\sqrt{5}} = \epsilon_2^2 \simeq M\frac{\alpha^2}{12}$$

with $\Pr > 1 - e^{-2\kappa^2}$

and plug this upper bound in BPDN

Former solution (Candès, Tao, ...)

- Therefore, from BPDN $\ell_2 - \ell_1$ instance optimality:

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \lesssim C\,\alpha + D\,e_0(K),$$

for $C, D > 0$

(for BPDN with ϵ_2, under prev. cond.)

Former solution (Candès, Tao, ...)

- Therefore, from BPDN $\ell_2 - \ell_1$ instance optimality:

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \lesssim C\,\alpha + D\,e_0(K),$$

for $C,D > 0$

(for BPDN with ϵ_2, under prev. cond.)

- Assuming :
  - bounded dynamics: $\|\boldsymbol{\Phi x}\|_\infty = \max_j |(\boldsymbol{\Phi x})_i| \leqslant \rho$ (e.g., by discarding saturation) (see later)
  - B bits per measurements $\Rightarrow \alpha \simeq \frac{2\rho}{2^B}$

$$\Rightarrow \text{BPDN RMSE} \lesssim C'\,2^{-B} + D\,e_0(K)$$

for $C',D > 0$

as soon as RIP holds: $M = O(K \log N/K)$

- Equivalently: BPDN RMSE $\simeq O(2^{-R/M}) + e_0(K)$

  for a rate $R = BM$ bits (total "bid budget" for all meas.)

RMSE Lower bound?

- Let a fixed K-sparse $\boldsymbol{x} \in \mathbb{R}^N$

RMSE Lower bound?

- Let a fixed K-sparse $\boldsymbol{x} \in \mathbb{R}^N$
- Oracle: you know $T = \operatorname{supp} \boldsymbol{x}$

RMSE Lower bound?

- Let a fixed K-sparse $\boldsymbol{x} \in \mathbb{R}^N$
- Oracle: you know $T = \operatorname{supp} \boldsymbol{x}$
- Noisy measurements (random noise):

  Given $\boldsymbol{\Phi} \in \mathbb{R}^{M \times N}$ with $\Phi_{ij} \sim_{\text{iid}} N(0,1)$

  $$\boldsymbol{y} = \boldsymbol{\Phi}_T \boldsymbol{x} + \boldsymbol{n}, \quad \text{with} \quad \mathbb{E}\, \boldsymbol{n}\boldsymbol{n}^T = \sigma^2 \mathbf{Id}_{M \times M}$$

- Assume: $\frac{1}{\sqrt{M}}\boldsymbol{\Phi}$ is RIP(K, δ_K) and RIP$(1, \delta_1)$
- Compute LS solution: $\hat{\boldsymbol{x}}_T = \boldsymbol{\Phi}_T^{\dagger} \boldsymbol{y} = \underbrace{(\boldsymbol{\Phi}_T^* \boldsymbol{\Phi}_T)^{-1} \boldsymbol{\Phi}_T^*}_{\text{pseudo-inverse}} \boldsymbol{y}$

  $\hat{\boldsymbol{x}}_{T^c} = 0$

- Then: $\text{MSE} = \mathbb{E}_{\boldsymbol{n}} \|\boldsymbol{x} - \hat{\boldsymbol{x}}\|^2 \geqslant r^{-1}\, \sigma^2 \left(\frac{1-\delta_1}{1+\delta_K}\right)$

  (as for BPDN) & $\text{MSE} \leqslant \frac{1}{1-\delta_K}\sigma^2$ from [Needell, Tropp, 08]

  for oversampling factor $r = M/K$

- for QCS: $\Rightarrow \text{RMSE} = \Omega(r^{-1/2}\, 2^{-B})$ & $\text{RMSE} = O(2^{-B})$

3. Consistent Reconstructions

Consistent reconstructions in CS?

- Problem in previous case: if $\hat{\boldsymbol{x}}$ solution of BPDN,

  - no Quantization Consistency (QC): $\mathcal{Q}[\boldsymbol{\Phi}\hat{\boldsymbol{x}}] \neq Q[\boldsymbol{\Phi}\boldsymbol{x}]$

$$\|\boldsymbol{\Phi}\hat{\boldsymbol{x}} - \mathcal{Q}[\boldsymbol{\Phi}\boldsymbol{x}]\| \leqslant \epsilon_2 \quad \not\Rightarrow \mathcal{Q}[\boldsymbol{\Phi}\hat{\boldsymbol{x}}] = Q[\boldsymbol{\Phi}\boldsymbol{x}]$$

(from BPDN constraint)

$\Longrightarrow$ sensing information is fully not exploited!

  - ℓ_2 constraint $\approx$ Gaussian distribution (MAP - cond. log. lik.)

But why looking for consistency?

First,

Proposition (Goyal, Vetterli, Thao, 98) *If T is known (with $|T| = K$), the best decoder* Dec() *provides a $\hat{\boldsymbol{x}} = \mathrm{Dec}(\boldsymbol{y}, \boldsymbol{\Phi})$ such that:*

$$\mathrm{RMSE} = (\mathbb{E}\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|^2)^{1/2} \gtrsim r^{-1}\alpha,$$

where $\mathbb{E}$ is wrt a probability measure on $\boldsymbol{x}_T$ in a bounded set $\mathcal{S} \subset \mathbb{R}^K$.

This bound is achieved, at least, for $\boldsymbol{\Phi}_T = \mathrm{DFT} \in \mathbb{R}^{M\times K}$, when Dec() *is* **consistent.**

V. K Goyal, M. Vetterli, N. T. Thao, "Quantized Overcomplete Expansions in R^{N}: Analysis, Synthesis, and Algorithms" , IEEE Tran. IT, 44(1), 1998

But why looking for consistency?

Second,

If $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ is a (random) frame in $\mathbb{R}^N$ $(M \geqslant N)$,

Then, for $\mathcal{Q}(\boldsymbol{y}) = \boldsymbol{y} + \boldsymbol{n}$ with $n_i \sim \mathcal{U}([-\frac{1}{2}\alpha, \frac{1}{2}\alpha])$, and $\hat{\boldsymbol{x}}$ consistent,

$$(\mathbb{E}_{\mathbf{\Phi},\boldsymbol{n}} \|\boldsymbol{x} - \hat{\boldsymbol{x}}\|^2)^{1/2} \lesssim r^{-1}\alpha,$$

[Powell, Whitehouse, 2013]

(unit norm frame)

and

$$\|\boldsymbol{x} - \hat{\boldsymbol{x}}\| \lesssim r^{-1}\alpha \cdot O(\log M, \log N, \log \eta),$$

[LJ 2014]

(Gaussian frame)

with $\Pr \geqslant 1 - \eta$.

or $\frac{K}{M}\alpha \cdot O(\log K, \log M, \log N, \log \eta)$ in K sparse case

In quest of consistency...

$\ell_2 \to \ell_\infty$

- Modify BPDN [W. Dai, O. Milenkovic, 09]

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \mathcal{Q}[\boldsymbol{\Phi u}] = \boldsymbol{q}$$

$\Leftrightarrow \boldsymbol{\Phi u} \in \mathcal{Q}^{-1}[\boldsymbol{q}]$

convex set in $\mathbb{R}^M$

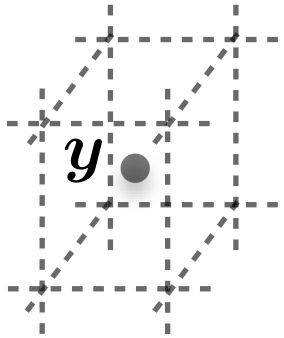

$\Leftrightarrow \|\boldsymbol{\Phi u} - \boldsymbol{q}\|_\infty \leq \alpha/2$

(if uniform quant.)

$\exists$ numerical methods

+ modified greedy algo:
"*subspace pursuit*"

In quest of consistency... $\ell_2 \to \ell_\infty$

- Modify BPDN [W. Dai, O. Milenkovic, 09]

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \mathcal{Q}[\boldsymbol{\Phi u}] = \boldsymbol{q}$$

Simulations: $M = 128, N = 256, K = 6, 1000$ trials $\Rightarrow \lambda \simeq 20$

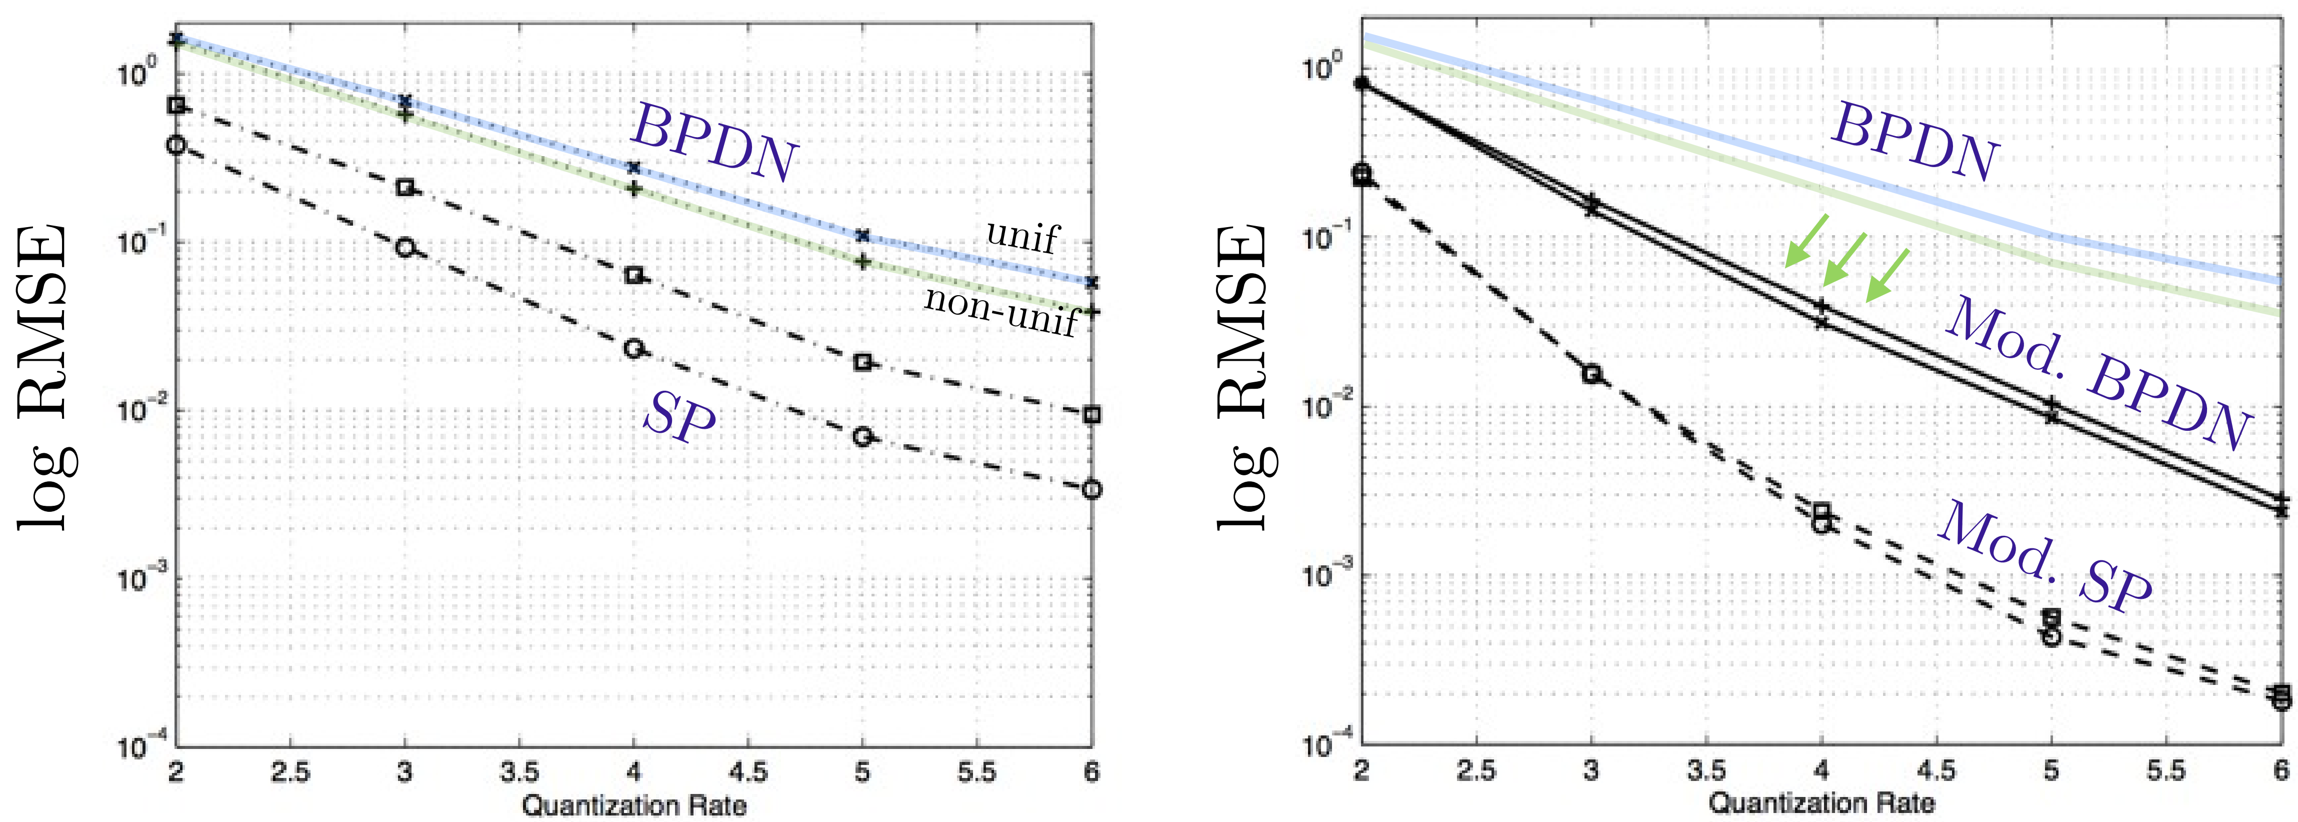

W. Dai, H. V. Pham, and O. Milenkovic, "Quantized Compressive Sensing", preprint, 2009

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

$\ell_2 \to \ell_p \ (p \geq 2)$

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

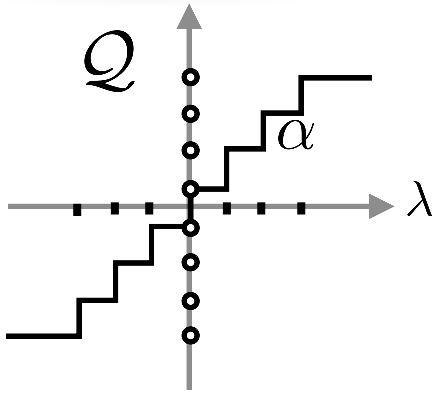

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$

- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

$\ell_2 \to \ell_p \ (p \geq 2)$

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

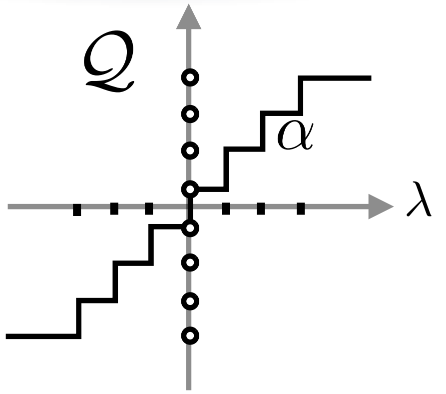

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$
- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

How to find it? again, uniform model:

Dequantizing CS?

$\ell_2 \to \ell_p \; (p \geq 2)$

[LJ, Hammond, Fadili, 2009, 2011]

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

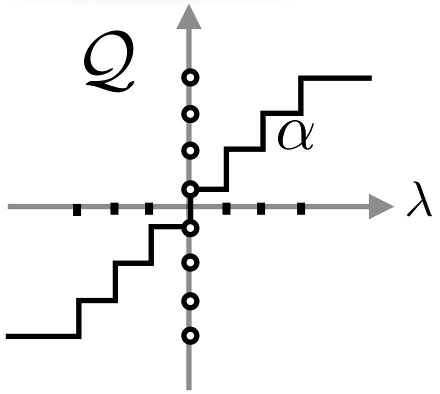

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$
- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

How to find it? again, uniform model:

$$\begin{aligned} n_i &= \mathcal{Q}[(\boldsymbol{\Phi x})_i] - (\boldsymbol{\Phi x})_i \\ &\in q_{k_i} - \mathcal{R}_{k_i} = [-\alpha/2, \alpha/2] \\ &\sim_{\text{iid}} \text{Uniform}([-\alpha/2, \alpha/2]) \end{aligned} \quad \Longrightarrow$$

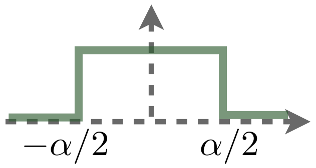

Estimating p^{th} moment:

$$\epsilon_p(\alpha) = \frac{\alpha}{2(p+1)^{1/p}} \left(M + \kappa(p+1)\sqrt{M}\right)^{1/p}$$

works with $\Pr \geq 1 - e^{-2\kappa^2}$

Note: $\epsilon_p(\alpha) \xrightarrow[p \to \infty]{} \frac{\alpha}{2} = \text{QC}\,!$

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

$\ell_2 \rightarrow \ell_p \; (p \geq 2)$

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

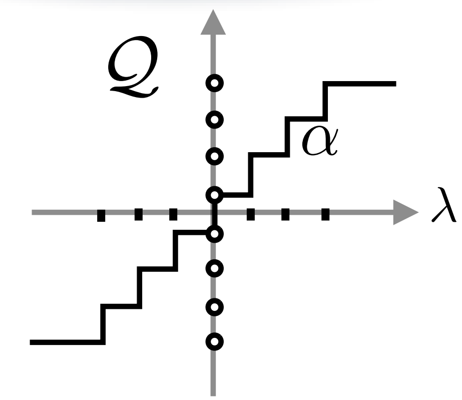

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$

- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \operatorname*{arg\,min}_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

BPDQ Stability ?

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

$\ell_2 \to \ell_p \ (p \geq 2)$

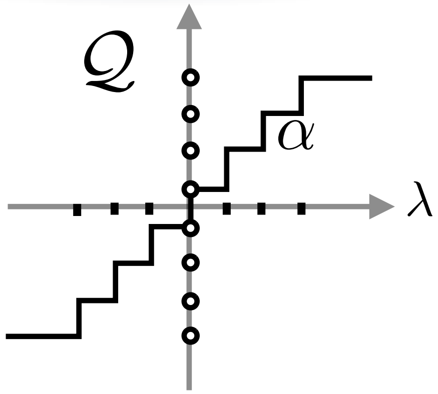

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$

- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

BPDQ Stability ?

Ok, if $\boldsymbol{\Phi}$ is RIP_p of order K, *i.e.*,

$\exists\ \mu_p > 0,\ \delta \in (0,1)$,

$$\sqrt{1-\delta}\,\|\boldsymbol{v}\|_2 \leqslant \tfrac{1}{\mu_p}\,\|\boldsymbol{\Phi v}\|_p \leqslant \sqrt{1+\delta}\,\|\boldsymbol{v}\|_2,$$

for all K sparse signals $\boldsymbol{v}$.

Dequantizing CS?

$\ell_2 \rightarrow \ell_p \ (p \geq 2)$

[LJ, Hammond, Fadili, 2009, 2011]

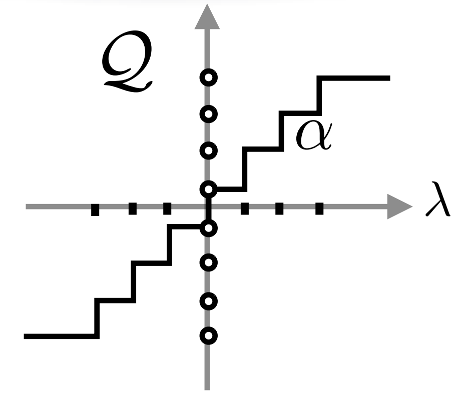

- Distortion model:

$$\boldsymbol{q} = \mathcal{Q}[\boldsymbol{\Phi x}] = \boldsymbol{\Phi x} + \boldsymbol{n}, \quad n_i \sim U(-\tfrac{\alpha}{2}, \tfrac{\alpha}{2})$$

- Observation: $\|\boldsymbol{\Phi x} - \boldsymbol{q}\|_\infty \leq \alpha/2$

- Reconstruction: Generalizing BPDN with BPDQ

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u} \in \mathbb{R}^N} \|\boldsymbol{u}\|_1 \text{ s.t. } \|\boldsymbol{q} - \boldsymbol{\Phi u}\|_p \leq \epsilon_p$$

Towards $p = \infty$

Related to GGD MAP

Gain over BPDN (for tight $\epsilon_p(\alpha, M)$)

$$\Rightarrow \|\boldsymbol{x} - \hat{\boldsymbol{x}}\| = O(\epsilon_p/\mu_p)$$

$$\Rightarrow \|\boldsymbol{x} - \hat{\boldsymbol{x}}\| = O(\alpha/\sqrt{p+1})$$

But no free lunch: for $\boldsymbol{\Phi}$ Gaussian

$$M = O\big((K \log N/K)^{p/2} \big)$$

$\Longrightarrow$ Another reading: limited range of valid p for a given M (and K)!

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

Histograms of $\alpha^{-1}(\boldsymbol{q} - \boldsymbol{\Phi}\hat{\boldsymbol{x}})_i$

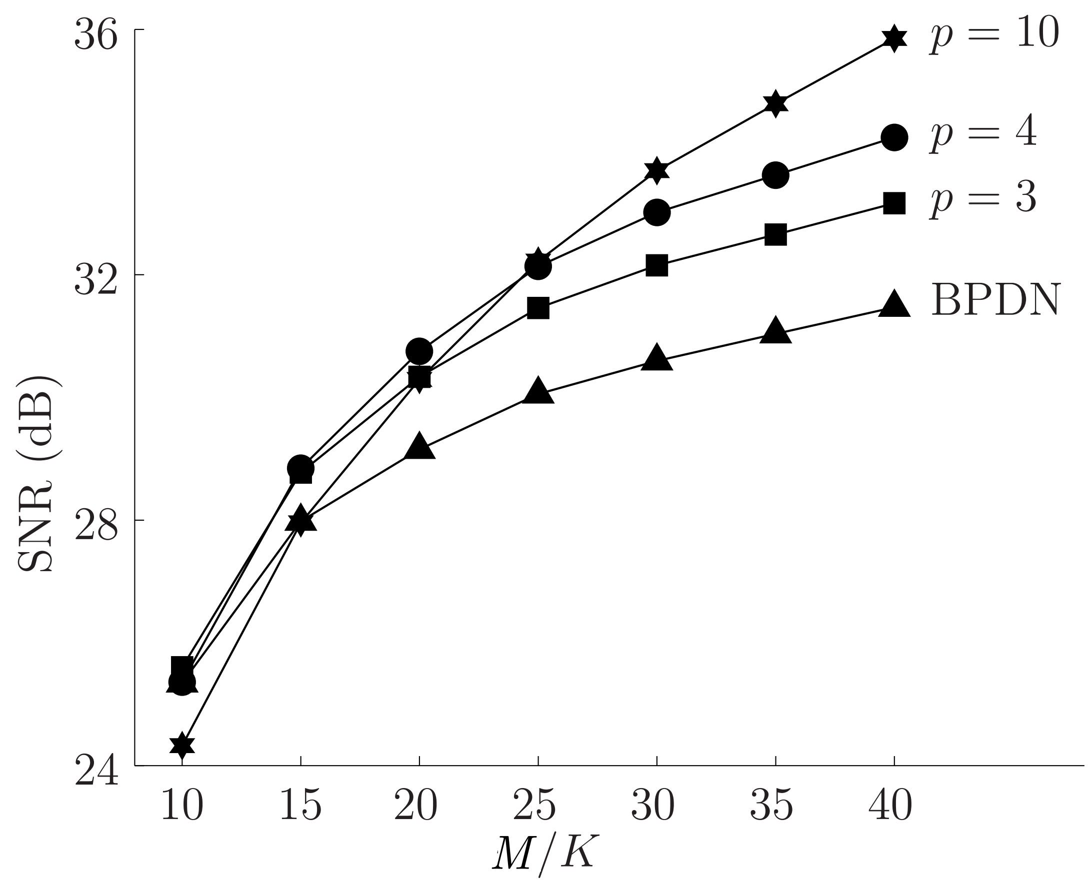

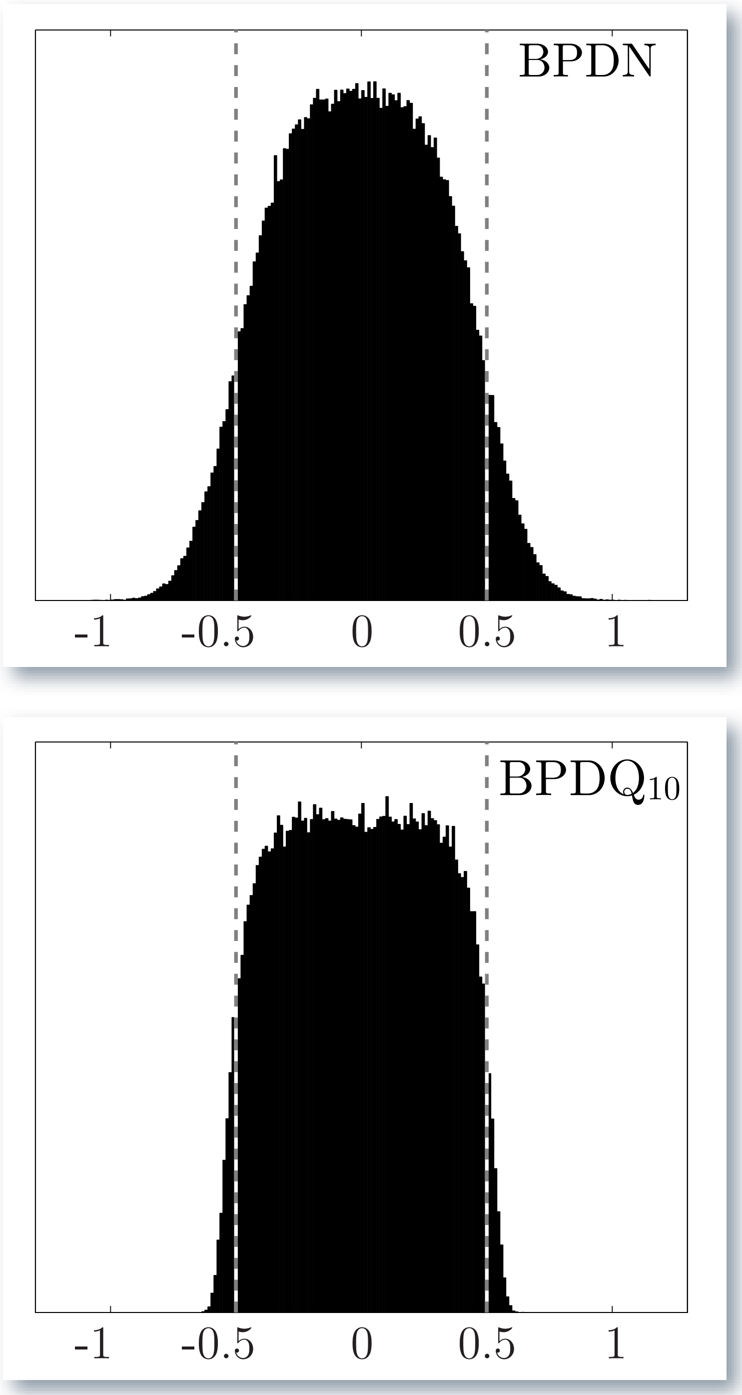

- N=1024, K=16, Gaussian $\boldsymbol{\Phi}$
- 500 K-sparse (canonical basis)
- Non-zero components follow $\mathcal{N}(0,1)$
- Quantiz. bin width $\alpha = \|\boldsymbol{\Phi x}\|_\infty/40$

LJ, D. Hammond, J. Fadili "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine." Information Theory, IEEE Transactions on, 57(1), 559-571.

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

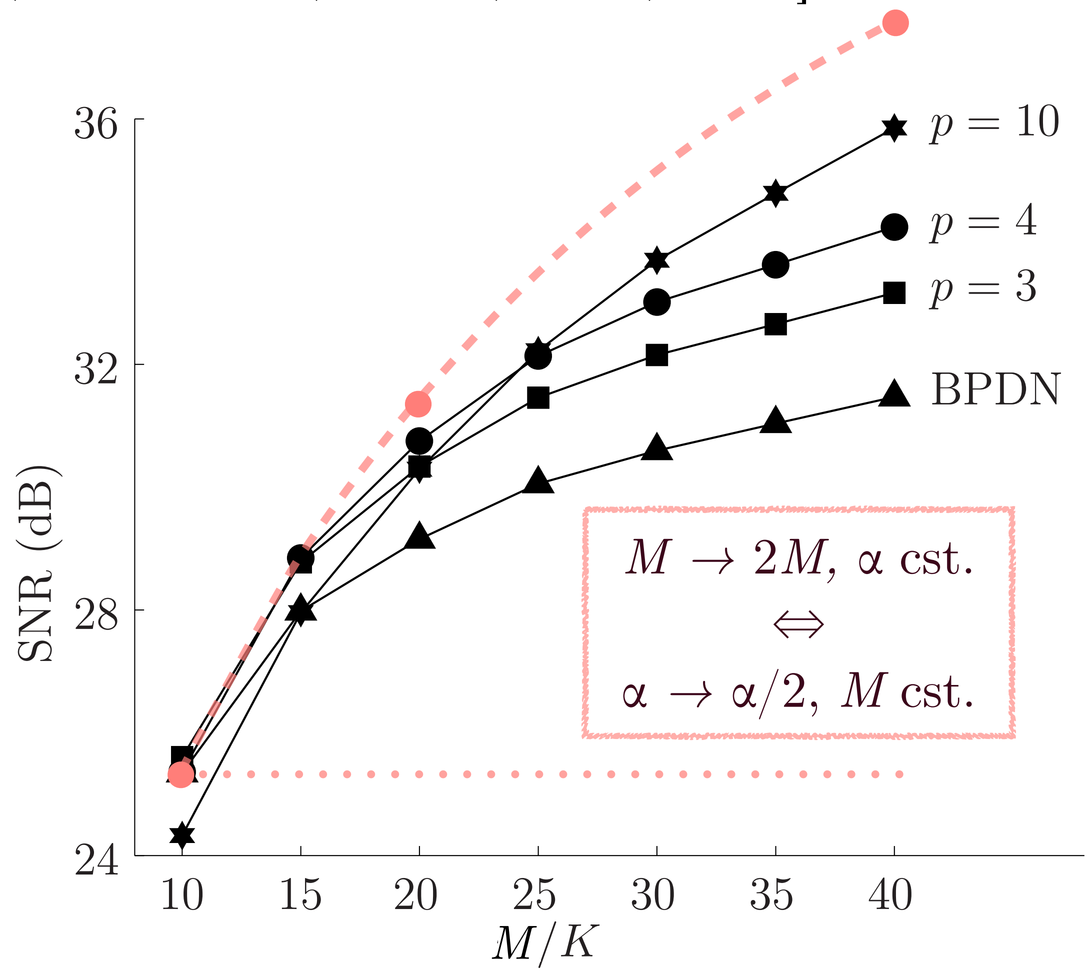

* N=1024, K=16, Gaussian $\boldsymbol{\Phi}$
* 500 K-sparse (canonical basis)
* Non-zero components follow $\mathcal{N}(0,1)$
* Quantiz. bin width $\alpha = \|\boldsymbol{\Phi x}\|_\infty/40$

Histograms of
$\alpha^{-1}(\boldsymbol{q} - \boldsymbol{\Phi}\hat{\boldsymbol{x}})_i$

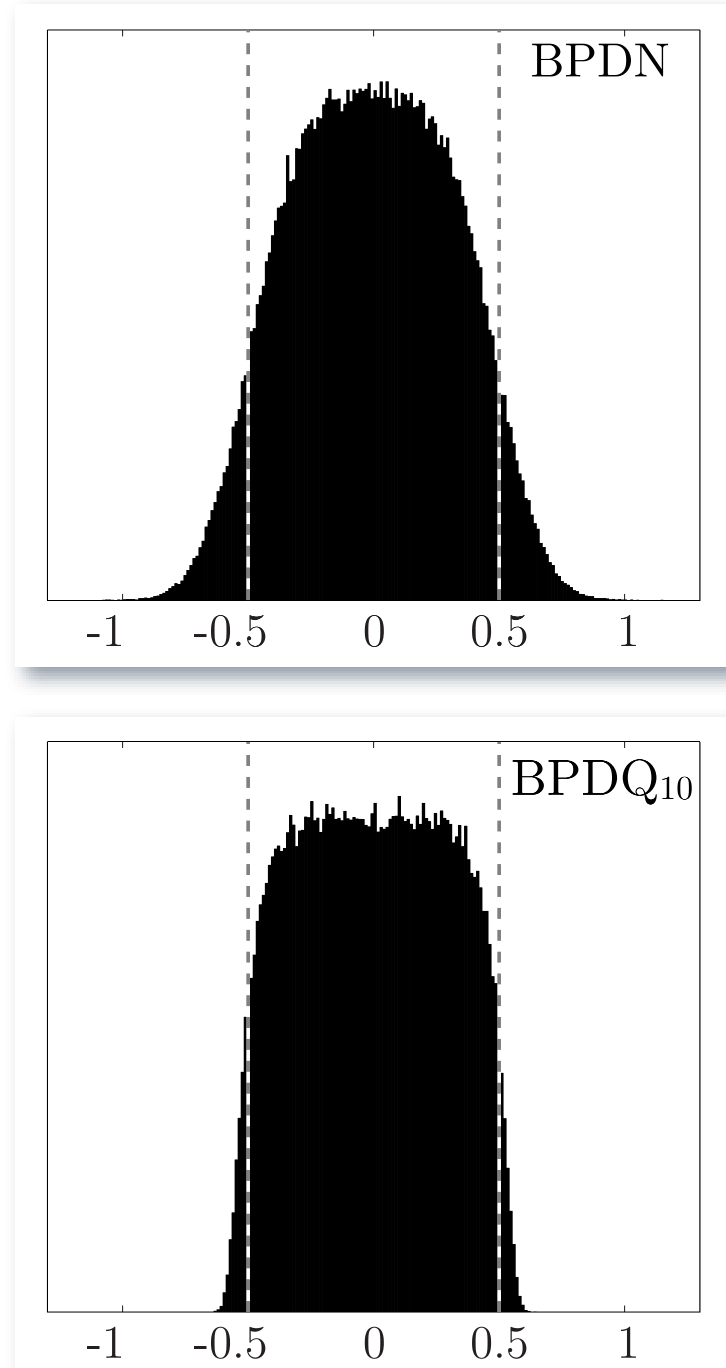

LJ, D. Hammond, J. Fadili "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine." Information Theory, IEEE Transactions on, 57(1), 559-571.

Dequantizing CS?

[LJ, Hammond, Fadili, 2009, 2011]

A bit outside the theory...

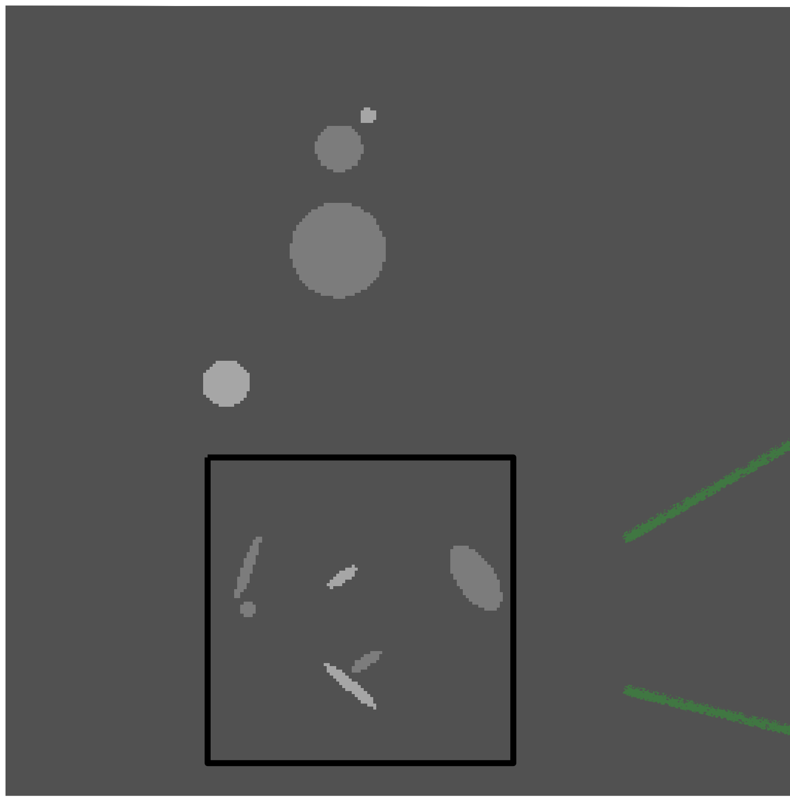

BPDN-TV

SNR: 8.96 dB

BPDQ_{10} -TV

SNR: 12.03 dB

- ✶ Synthetic Angiogram [Michael Lustig 07, SPARCO],
- ✶ $\boldsymbol{\Phi}$: **Random Fourier Ensemble**
- ✶ $N/M = 8$
- ✶ Decoder: $\Delta_{TV,p}(y, \epsilon_p)$
- ✶ Quantiz. bin width = 50 (i.e. 12 bins)

LJ, D. Hammond, J. Fadili "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine." Information Theory, IEEE Transactions on, 57(1), 559-571.

4. Sigma-Delta quantization in CS

Context:

- **Former attempts:** (see prev. slides)

  CS + uniform scalar quantization (or pulse code modulation - PCM)

For K-sparse signals: $\|\mathcal{Q}_\alpha[\boldsymbol{\Phi x}] - \boldsymbol{\Phi x}\|_2 \leqslant c\sqrt{M}\alpha \;\Rightarrow\; \|\boldsymbol{x}^* - \boldsymbol{x}\| \leqslant C\alpha$ (with RIP)

and for high λ, $\|\mathcal{Q}_\alpha[\boldsymbol{\Phi x}] - \boldsymbol{\Phi x}\|_p \leqslant cM^{1/p}\alpha \;\Rightarrow\; \|\boldsymbol{x}^* - \boldsymbol{x}\| \leqslant C\alpha/\sqrt{p+1}$ (with RIP$_p$)

- **No (real) improvement if M increases!**
- **Can we do better?**

Context:

- **Former attempts:** (see prev. slides)

  CS + uniform scalar quantization (or pulse code modulation - PCM)

For K-sparse signals: $\|\mathcal{Q}_\alpha[\boldsymbol{\Phi x}] - \boldsymbol{\Phi x}\|_2 \leqslant c\sqrt{M}\alpha \;\Rightarrow\; \|\boldsymbol{x}^* - \boldsymbol{x}\| \leqslant C\alpha$ (with RIP)

and for high λ, $\|\mathcal{Q}_\alpha[\boldsymbol{\Phi x}] - \boldsymbol{\Phi x}\|_p \leqslant cM^{1/p}\alpha \;\Rightarrow\; \|\boldsymbol{x}^* - \boldsymbol{x}\| \leqslant C\alpha/\sqrt{p+1}$ (with RIP$_p$)

- **No (real) improvement if M increases!**
- **Can we do better?**

  Can we have $\|\boldsymbol{x}^* - \boldsymbol{x}\| \leqslant O(r^{-s}\alpha)$ for some $s > 0$?

- **Staying with PCM,** $s \leqslant 1$ (Goyal-Vetterli-Thao lower bound)
- **Solution: replacing PCM by $\Sigma\Delta$ quantization!**

  [S. Güntürk, A. Powell, R. Saab, Ö. Yılmaz]

$\Sigma\Delta$ quantization (reminder)

- PCM: Signal sensing + unif. quantization (step α)

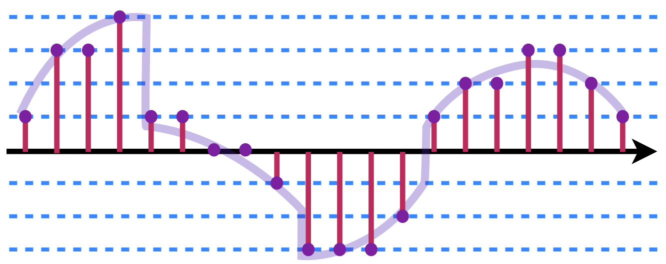

$\boldsymbol{x} \in \mathbb{R}^K \longrightarrow \boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} \in \mathbb{R}^M$

$\boldsymbol{q} = \mathcal{Q}_{\rm PCM}[\boldsymbol{y}]$ with

$$q_k = \mathcal{Q}_{\rm PCM}[y_k] := \operatorname*{argmin}_{u \in \alpha\mathbb{Z}} |y_k - u|, \quad 1 \leqslant k \leqslant M$$

Let $\boldsymbol{A}^{\#}$, a left inverse of $\boldsymbol{A}$, *i.e.*, $\boldsymbol{A}^{\#}\boldsymbol{A} = \mathbf{Id}$.

Then, $\hat{\boldsymbol{x}} := \boldsymbol{A}^{\#}\boldsymbol{q} \Rightarrow \|\boldsymbol{x} - \hat{\boldsymbol{x}}\| = \|\boldsymbol{A}^{\#}\underbrace{(\boldsymbol{y} - \boldsymbol{q})}_{\text{quant. noise}}\|$

→ Goal: $\displaystyle \operatorname*{minimize}_{\boldsymbol{A}^{\#}\boldsymbol{A} = \mathbf{Id}.} \|\boldsymbol{A}^{\#}(\boldsymbol{y} - \boldsymbol{q})\|!$

→ Taking (Moore-Penrose) pseudo-inverse: $\boldsymbol{A}^{\#} = \boldsymbol{A}^{\dagger} = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*$
(or canonical dual of the frame $\boldsymbol{A}$)

- In CS, this could be used if signal support was known (see before)

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta \equiv$ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta \equiv$ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- 1st order $\Sigma\Delta$: (in 1-D) Quantizing the sequence $\{y_j : j \geqslant 0\}$

Use of state variables $\{\rho_j\}$ (1-step memory):

find q_j: $q_j = \mathcal{Q}^{(1)}_{\Sigma\Delta}[y_j] := \mathrm{argmin}_{u \in \alpha\mathbb{Z}} |\rho_{j-1} + y_j - u|$ $= \mathcal{Q}_{\mathrm{PCM}}[\rho_{j-1} + y_j]$

find ρ_j: $(\Delta\rho)_j = \rho_j - \rho_{j-1} = y_j - q_j$ (difference eq.)

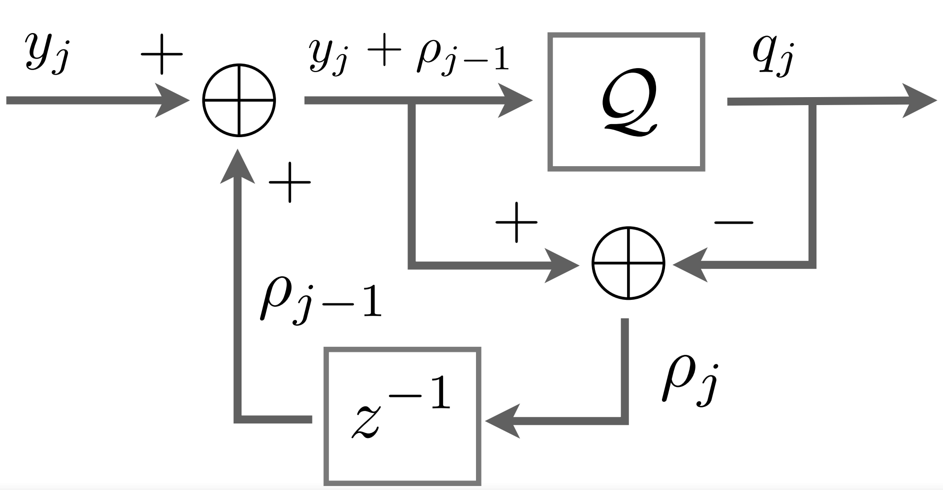

with: $|\rho_j| \leqslant \alpha$

$|y_j - q_j| \leqslant 2\alpha$

bigger than α but still $O(\alpha)$

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta \equiv$ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- s^{th} order $\Sigma\Delta$: (in 1-D) Quantizing the sequence $\{y_j : j \geqslant 0\}$

  Use of state variables $\{\rho_j\}$ (s-step memory):

Remark:
PCM is
0^{th} order $\Sigma\Delta$

find q_j: $\quad q_j = \mathcal{Q}^{(s)}_{\Sigma\Delta}[y_j] := \text{argmin}_{u\,\in\,\alpha\mathbb{Z}} \big|\sum_{i=1}^{s} (-1)^{i-1} \binom{s}{i} \rho_{j-n} + y_j - u\big|$

find ρ_j: $\quad (\Delta^s \rho)_j = y_j - q_j \quad$ (s^{th} order difference eq.)

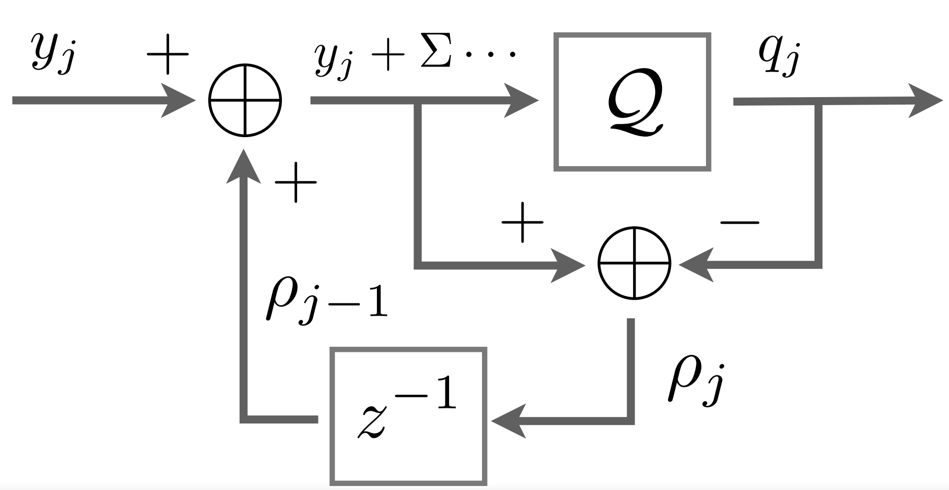

with: $|\rho_j| \leqslant \alpha$

$|y_j - q_j| \leqslant 2^{s-1}\alpha$

bigger than α but still $O(\alpha)$

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta$ ≡ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- s^{th} order $\Sigma\Delta$:

Most important fact: $(\Delta^s \rho)_j = y_j - q_j \;\Leftrightarrow\; \boldsymbol{D}^s \boldsymbol{\rho} = \boldsymbol{y} - \boldsymbol{q}$

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta \equiv$ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- s^{th} order $\Sigma\Delta$:

Most important fact: $(\Delta^s \rho)_j = y_j - q_j \;\Leftrightarrow\; \boldsymbol{D}^s \boldsymbol{\rho} = \boldsymbol{y} - \boldsymbol{q}$

$$\hat{\boldsymbol{x}} := \boldsymbol{A}^{\#}\boldsymbol{q} \Rightarrow \|\boldsymbol{x} - \hat{\boldsymbol{x}}\| = \|\boldsymbol{A}^{\#}\boldsymbol{D}^s(\boldsymbol{y} - \boldsymbol{q})\|$$

$\Sigma\Delta$ quantization (reminder)

- $\Sigma\Delta \equiv$ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- s^{th} order $\Sigma\Delta$:

Most important fact: $(\Delta^s \rho)_j = y_j - q_j \;\Leftrightarrow\; \boldsymbol{D}^s \boldsymbol{\rho} = \boldsymbol{y} - \boldsymbol{q}$

$$\hat{\boldsymbol{x}} := \boldsymbol{A}^{\#}\boldsymbol{q} \Rightarrow \|\boldsymbol{x} - \hat{\boldsymbol{x}}\| = \|\boldsymbol{A}^{\#}\boldsymbol{D}^s(\boldsymbol{y} - \boldsymbol{q})\|$$

minimize $\|\boldsymbol{A}^{\#}\boldsymbol{D}^s(\boldsymbol{y} - \boldsymbol{q})\|$!

~~Pseudo-inverse~~

~~$\boldsymbol{A}^{\dagger} = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*$~~

Sobolev duals

$\boldsymbol{A}_{\text{sob},s} = (\boldsymbol{D}^{-s}\boldsymbol{A})^{\dagger}\boldsymbol{D}^{-s}$

ΣΔ quantization (reminder)

- ΣΔ ≡ noise shaping! Enjoy of:
  - freedom to pick $\boldsymbol{q} \in \alpha\mathbb{Z}^M$
  - freedom to take another left inverse $\boldsymbol{A}^{\#}$
- $s^{\rm th}$ order ΣΔ:

Most important fact: $(\Delta^s \rho)_j = y_j - q_j \;\Leftrightarrow\; \boldsymbol{D}^s \boldsymbol{\rho} = \boldsymbol{y} - \boldsymbol{q}$

$$\hat{\boldsymbol{x}} = \boldsymbol{A}_{\mathrm{sob},s}\boldsymbol{q} \qquad \boldsymbol{A}_{\mathrm{sob},s} = (\boldsymbol{D}^{-s}\boldsymbol{A})^{\dagger}\boldsymbol{D}^{-s}$$

Proposition *Let* $\boldsymbol{A} \in \mathbb{R}^{M\times K}$ *with* $A_{ij} \sim_{\rm iid} \mathcal{N}(0,1)$.
For any $\kappa \in (0,1)$, *if* $r := M/K \geqslant c(\log M)^{1/(1-\kappa)}$, *then with* $Pr > 1 - e^{-c' M/r^{\kappa}}$,

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \;\leqslant\; C_s\, r^{-\kappa(s-\frac{1}{2})}\alpha,$$

for some $c, c', C_s > 0$.

proof: show that
$\sigma_{\min}(\boldsymbol{D}^{-s}\boldsymbol{A}) > C'_s\, r^{\kappa(s-\frac{1}{2})}\sqrt{M}$

ΣΔ quantization in CS

$$\boldsymbol{x} \in \Sigma_K \subset \mathbb{R}^N \;\rightarrow\; \boldsymbol{y} = \boldsymbol{\Phi x} \in \mathbb{R}^M \;\rightarrow\; \boldsymbol{q} = \mathcal{Q}^{(s)}_{\Sigma\Delta}[\boldsymbol{y}]$$

$$\|\boldsymbol{y} - \boldsymbol{q}\| \leqslant 2^{s-1}\,\alpha\,\sqrt{M}$$

<u>Two-steps procedure:</u>

<u>remark</u>: Recent dev. don't require these!

1. find the support T of $\boldsymbol{x}$: coarse approx. with BPDN
2. compute $\hat{\boldsymbol{x}} := (\boldsymbol{\Phi}_T)_{\mathrm{sob},s}\boldsymbol{q} = (\boldsymbol{D}^{-s}\boldsymbol{\Phi}_T)^{\dagger}\boldsymbol{D}^{-s}\boldsymbol{q}$

Proposition *Let* $\boldsymbol{\Phi} \in \mathbb{R}^{M\times K}$ *with* $\Phi_{ij} \sim_{\mathrm{iid}} \mathcal{N}(0,1)$. *Suppose* $\kappa \in (0,1)$ *and* $r := M/K \geqslant c(\log M)^{1/(1-\kappa)}$ *for* $c > 0$. *Then,* $\exists\, c', C, C_s > 0$ *such that, with* $Pr > 1 - e^{-c' M/r^{\kappa}}$, *for all* $\boldsymbol{x} \in \Sigma_K$ *s.t.* $\min_{i\in \mathrm{supp}\,\boldsymbol{x}} |x_i| \geqslant C\alpha$,

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \;\leqslant\; C_s\, r^{-\kappa(s-\frac{1}{2})}\alpha.$$

<u>proof</u>: Union bound on any K-column subset of $\boldsymbol{\Phi}$ + proba having good support.

ΣΔ quantization in CS (Simulations)

$M \in \{100, 200, \cdots, 1000\}$, $K = 10$ and 1000 trials ($x_i \in \{0, \pm 1/\sqrt{K}\}, \|\boldsymbol{x}\| \simeq 1, \alpha = 10^{-2}$)

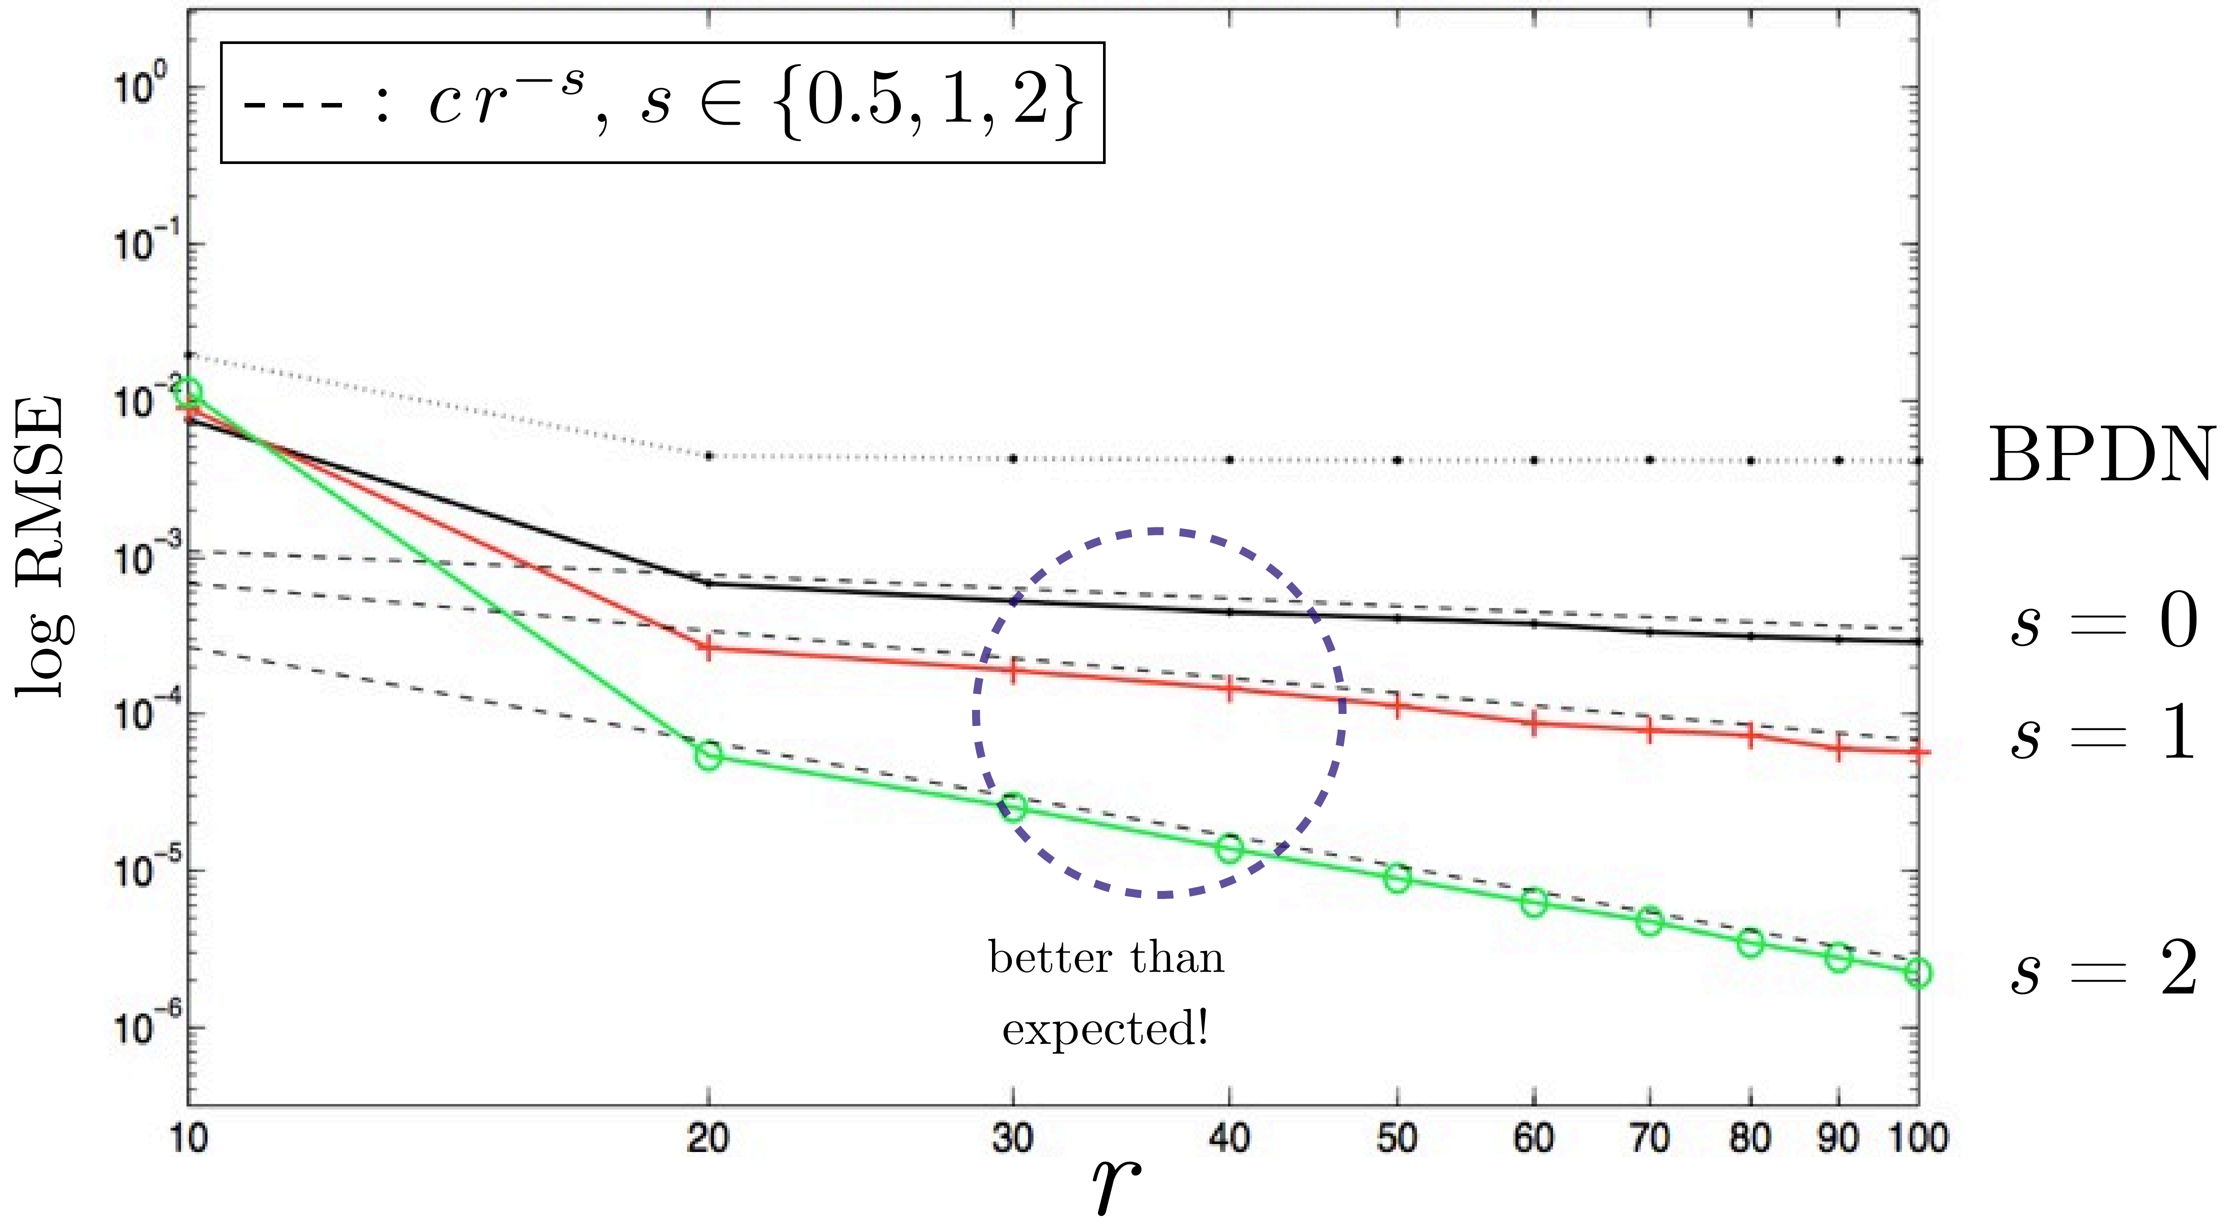

Güntürk, C. S., Lammers, M., Powell, A. M., Saab, R., & Yılmaz, Ö. (2013). **Sobolev duals for random frames and ΣΔ quantization of compressed sensing measurements**. Foundations of Computational Mathematics, 13(1), 1-36.

5. To saturate or not?
And how much?

Saturation phenomenon:

Uniform quantization:

- α quantization interval
- error per measurement bounded:

$$|\lambda - \mathcal{Q}_\alpha[\lambda]| \leqslant \alpha/2$$

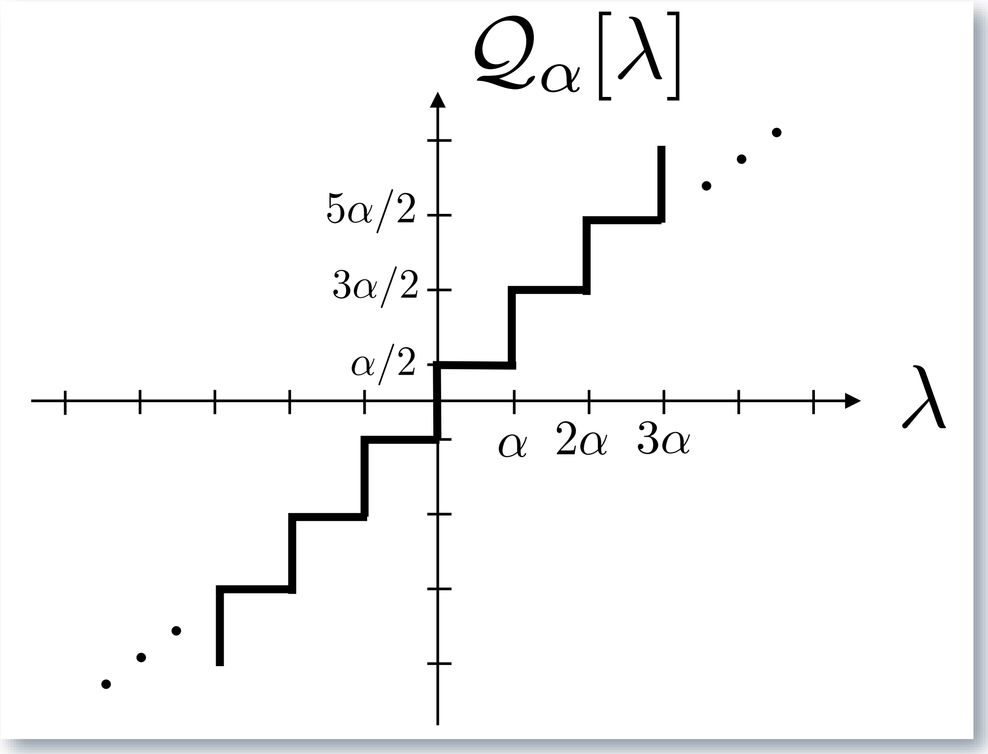

Saturation phenomenon:

Uniform quantization:

- α quantization interval
- error per measurement bounded:

$$|\lambda - \mathcal{Q}_\alpha[\lambda]| \leqslant \alpha/2$$

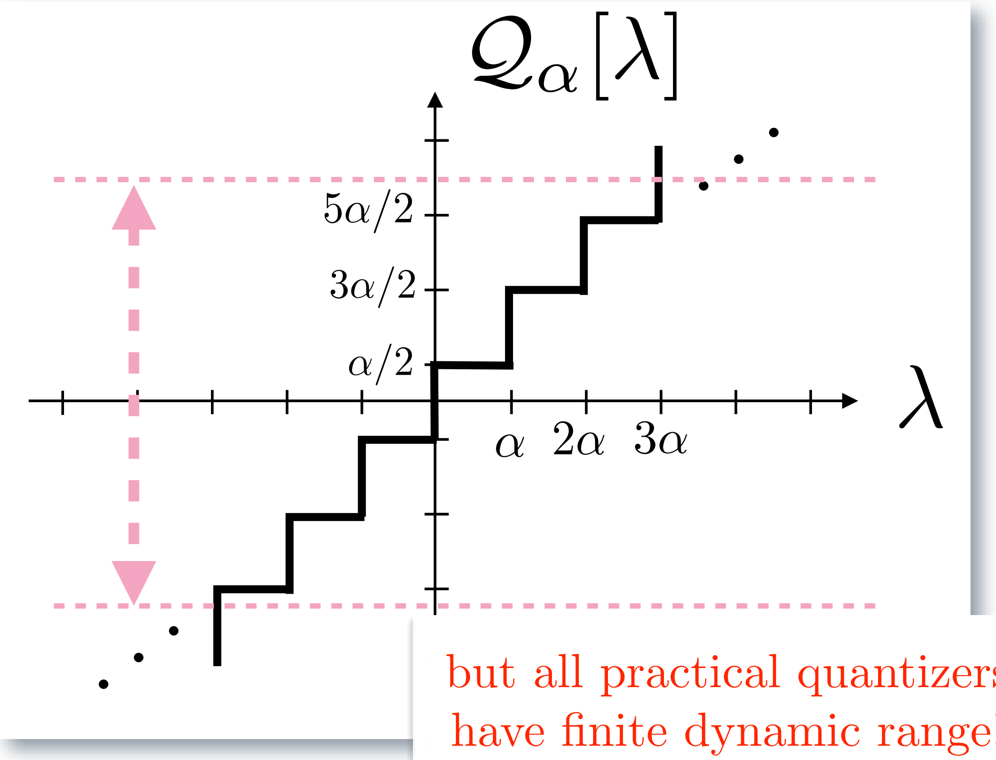

but all practical quantizers have finite dynamic range!

Saturation phenomenon:

Uniform quantization:

- α quantization interval
- error per measurement bounded:

$$|\lambda - \mathcal{Q}_\alpha[\lambda]| \leqslant \alpha/2$$

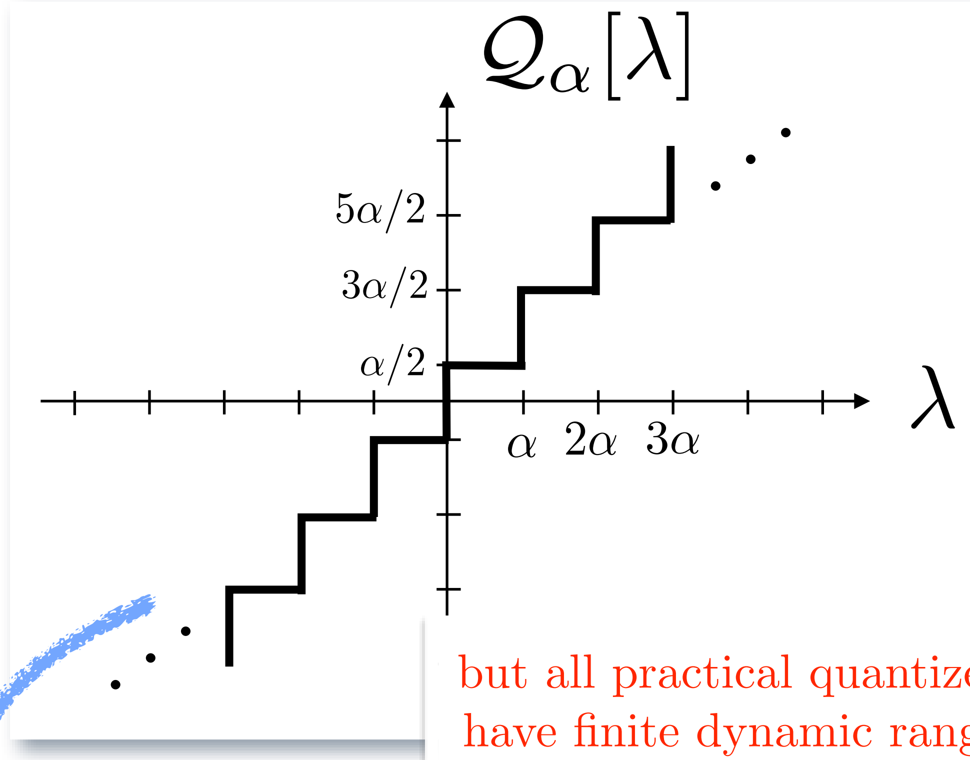

but all practical quantizers have finite dynamic range!

Finite Dynamic Range Quantization:

- G "saturation level"
- B bit rate (bits per measurement)
- quantization interval is $\alpha = 2^{-B+1}G$
- measurements above G saturate
- saturation error is *unbounded*

CS guarantees are for bounded errors only!

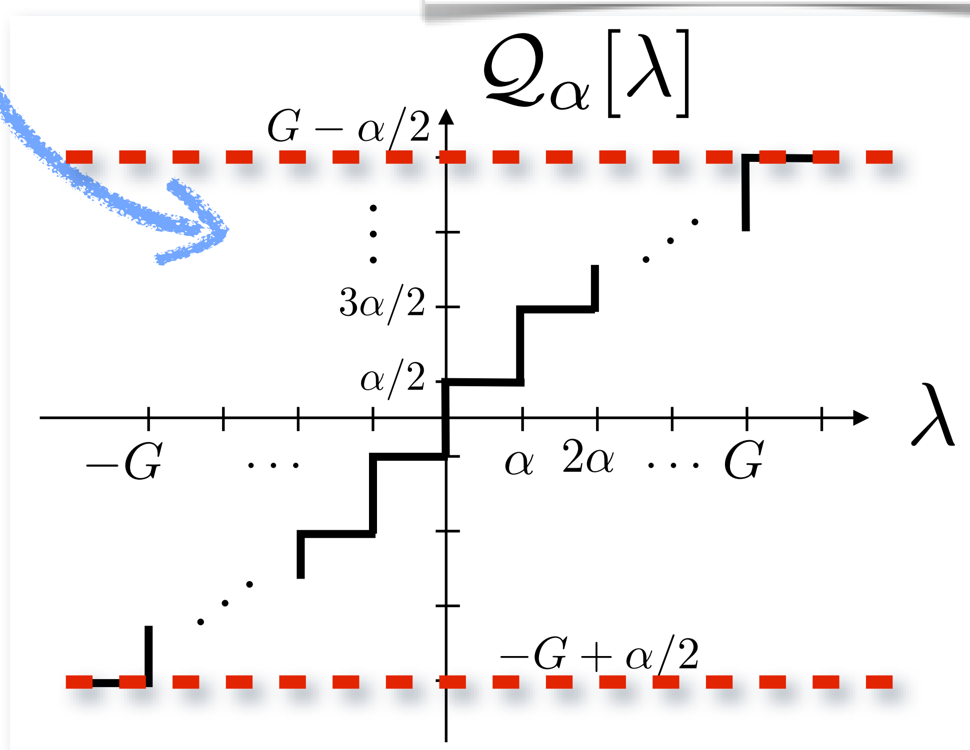

Democracy in Action

[Laska, Boufounos, Davenport, Baraniuk 12]

(i) Saturation Rejection:

Simply discard saturated measurements and corresponding rows of Φ

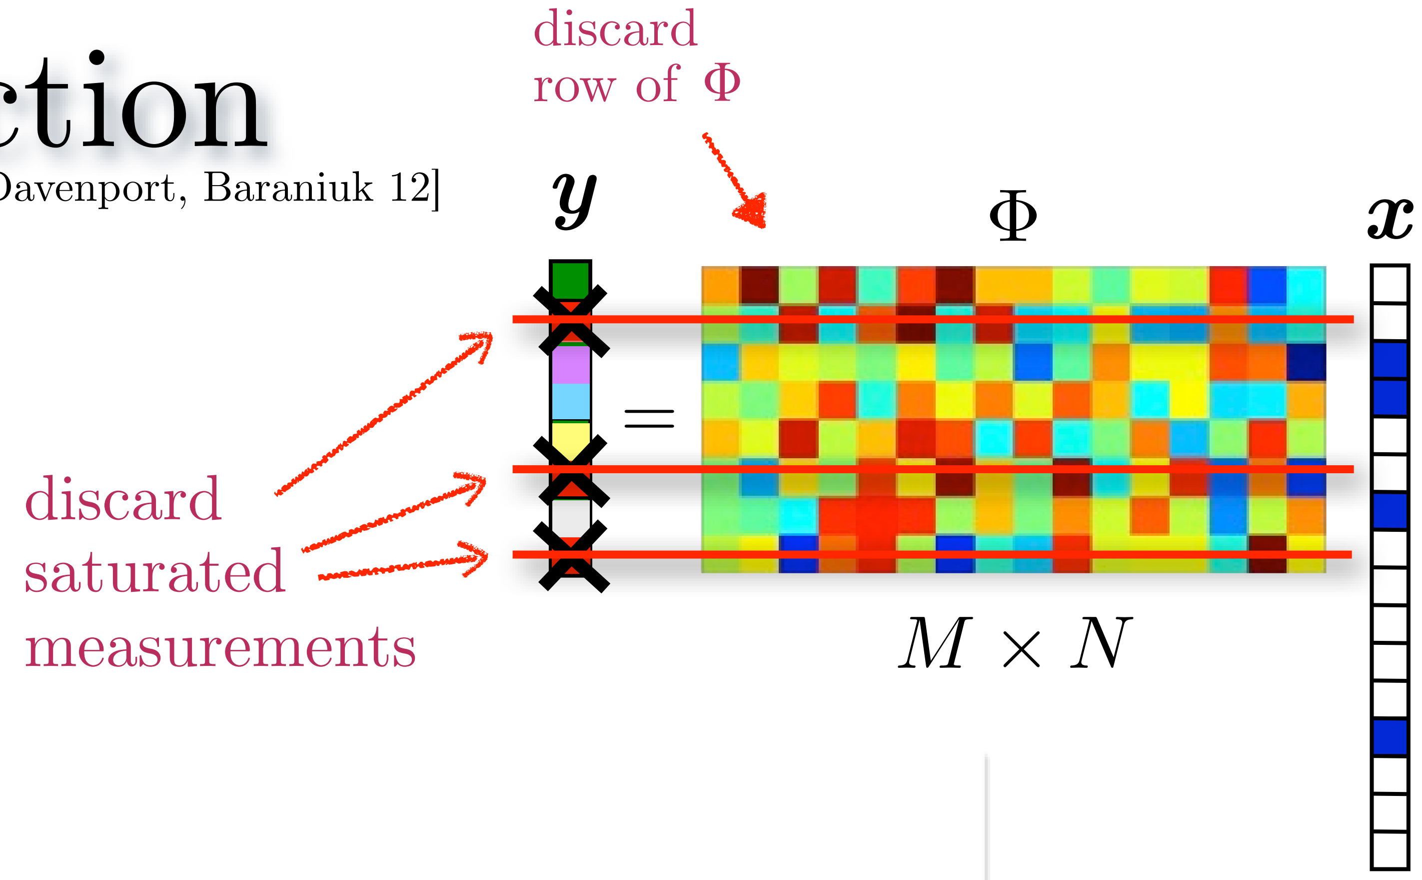

"democratic measurements"

each measurement has roughly same amount of information

→ RIP holds on row subsets of Φ

Democracy in Action

[Laska, Boufounos, Davenport, Baraniuk 12]

(i) Saturation Rejection:

Simply discard saturated measurements and corresponding rows of Φ

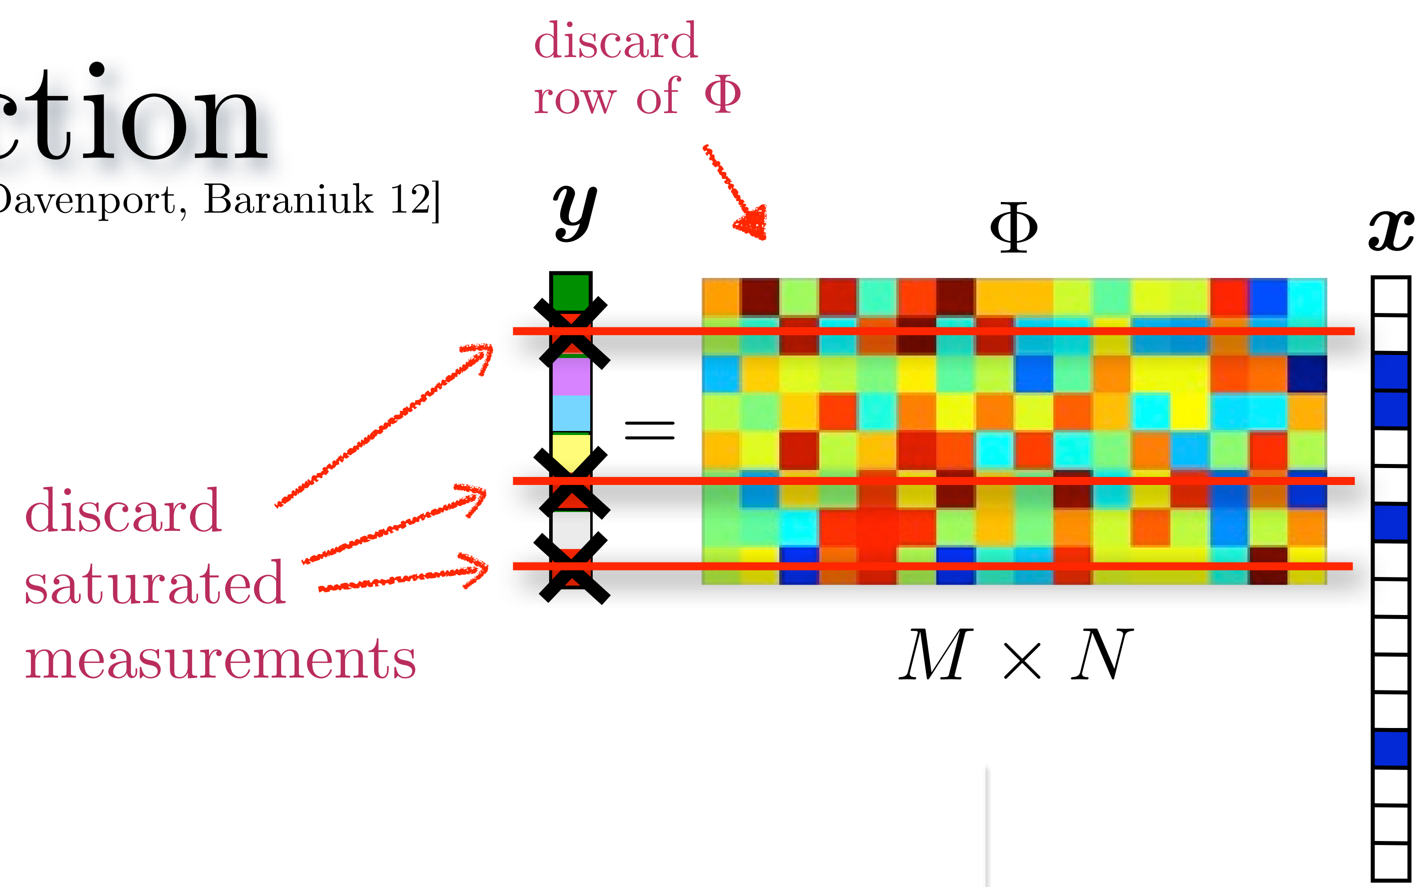

"democratic measurements"

each measurement has roughly same amount of information

RIP holds on row subsets of Φ

(ii) Saturation Consistency:

Include saturated measurements as inequality constraint

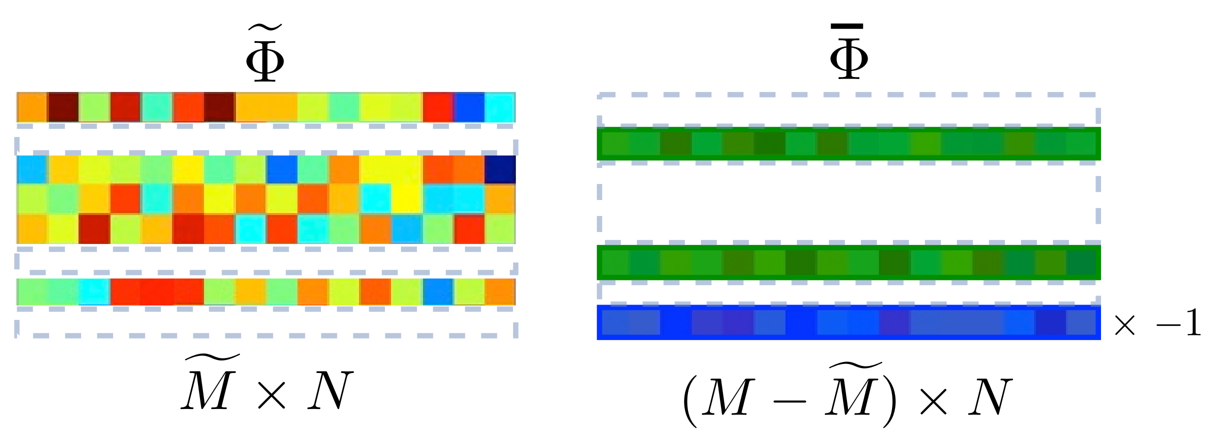

$$\hat{\mathbf{x}} = \mathrm{argmin}_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \|\widetilde{\Phi}\mathbf{x} - \widetilde{\mathbf{y}}\|_2 < \epsilon$$

Measurement error term (quantization)

$$\text{and} \quad \bar{\Phi}\mathbf{x} \geq G \cdot \mathbf{1}$$

Saturation consistency constraint

Experimental Results

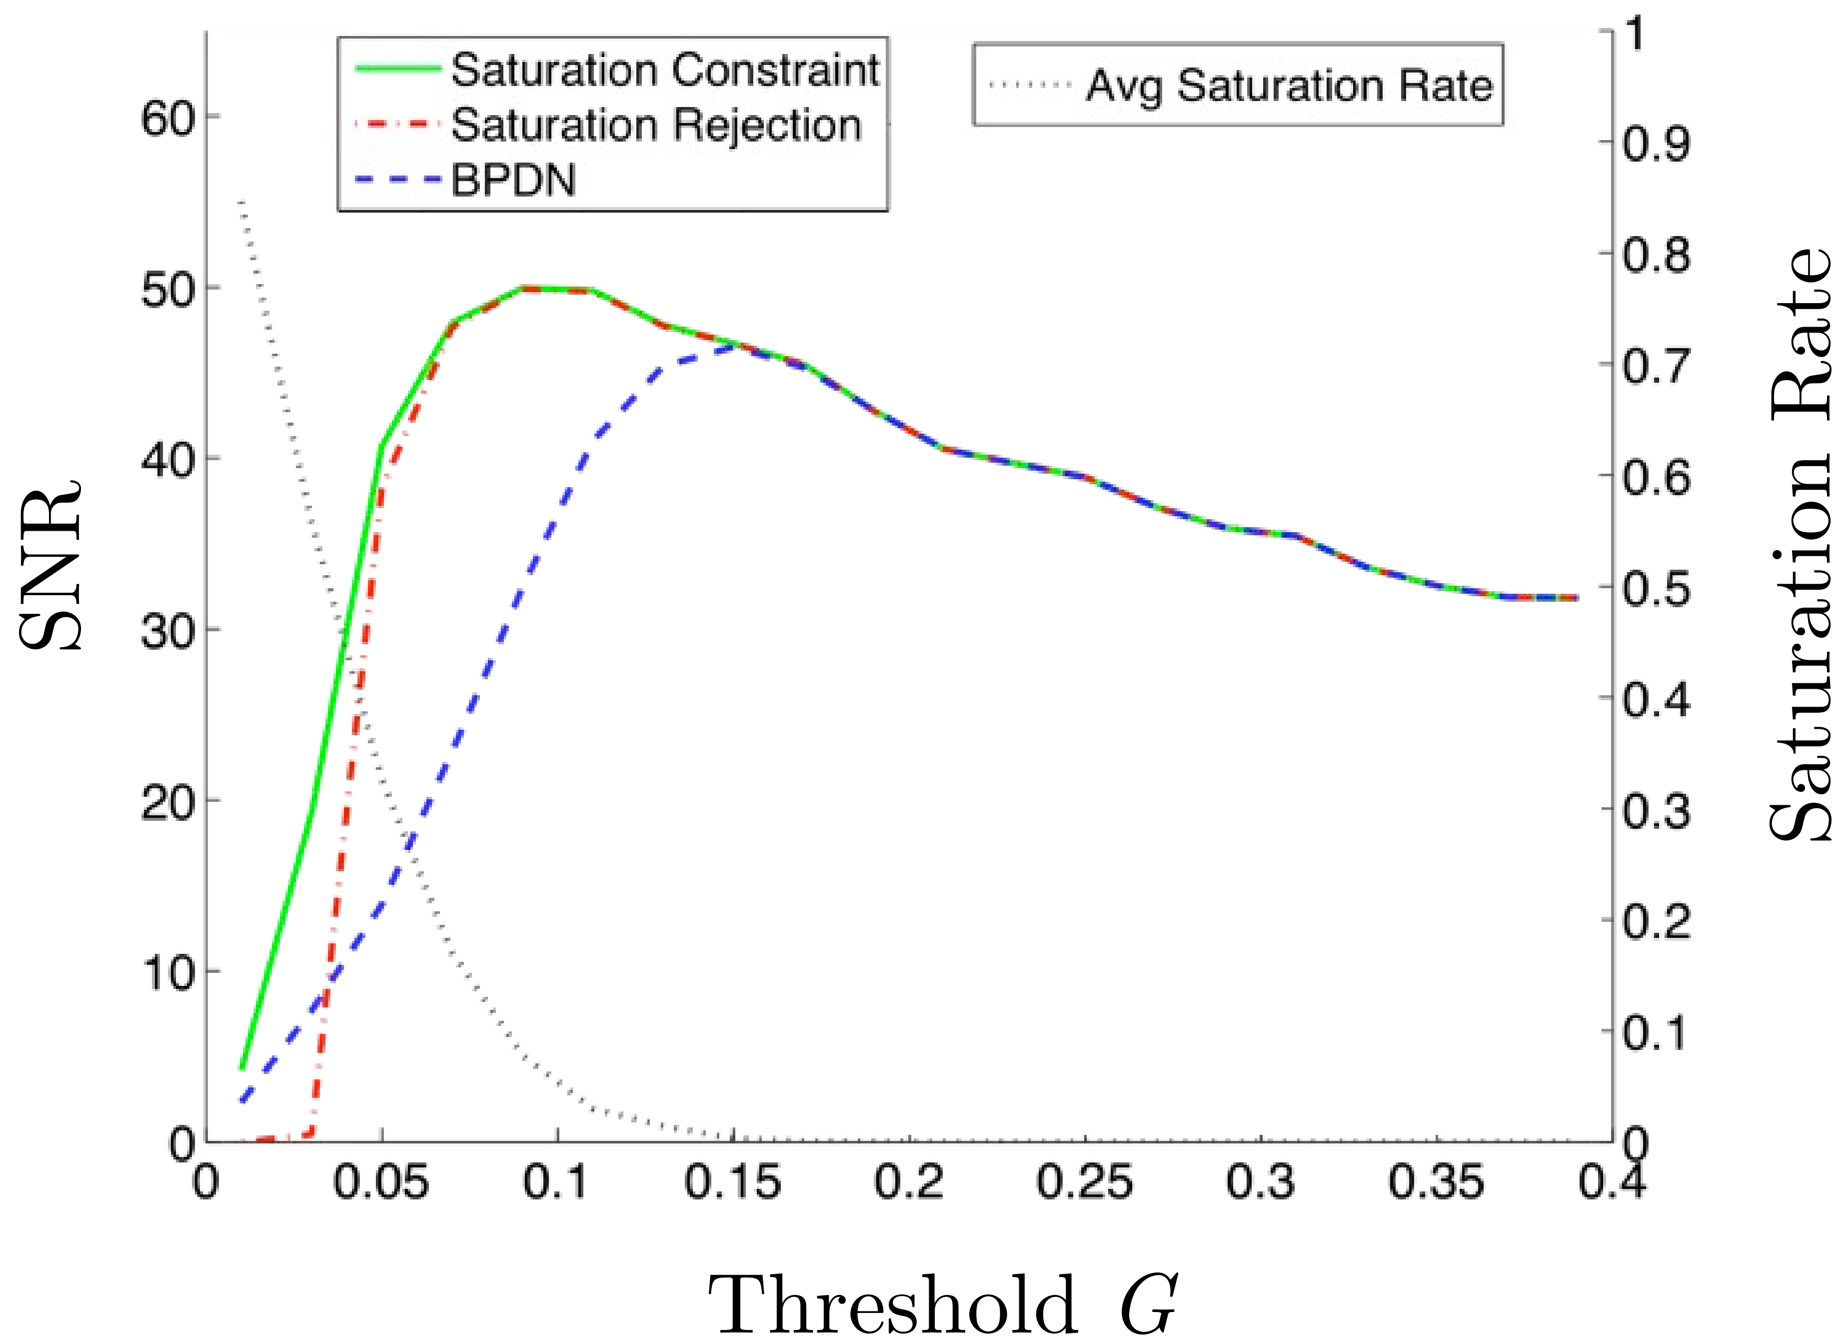

J.N. Laska, P.T. Boufounos, M.A. Davenport, R.G.Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing”. *Applied and Computational Harmonic Analysis*, *31*(3), 429-443. (2011)

Experimental Results

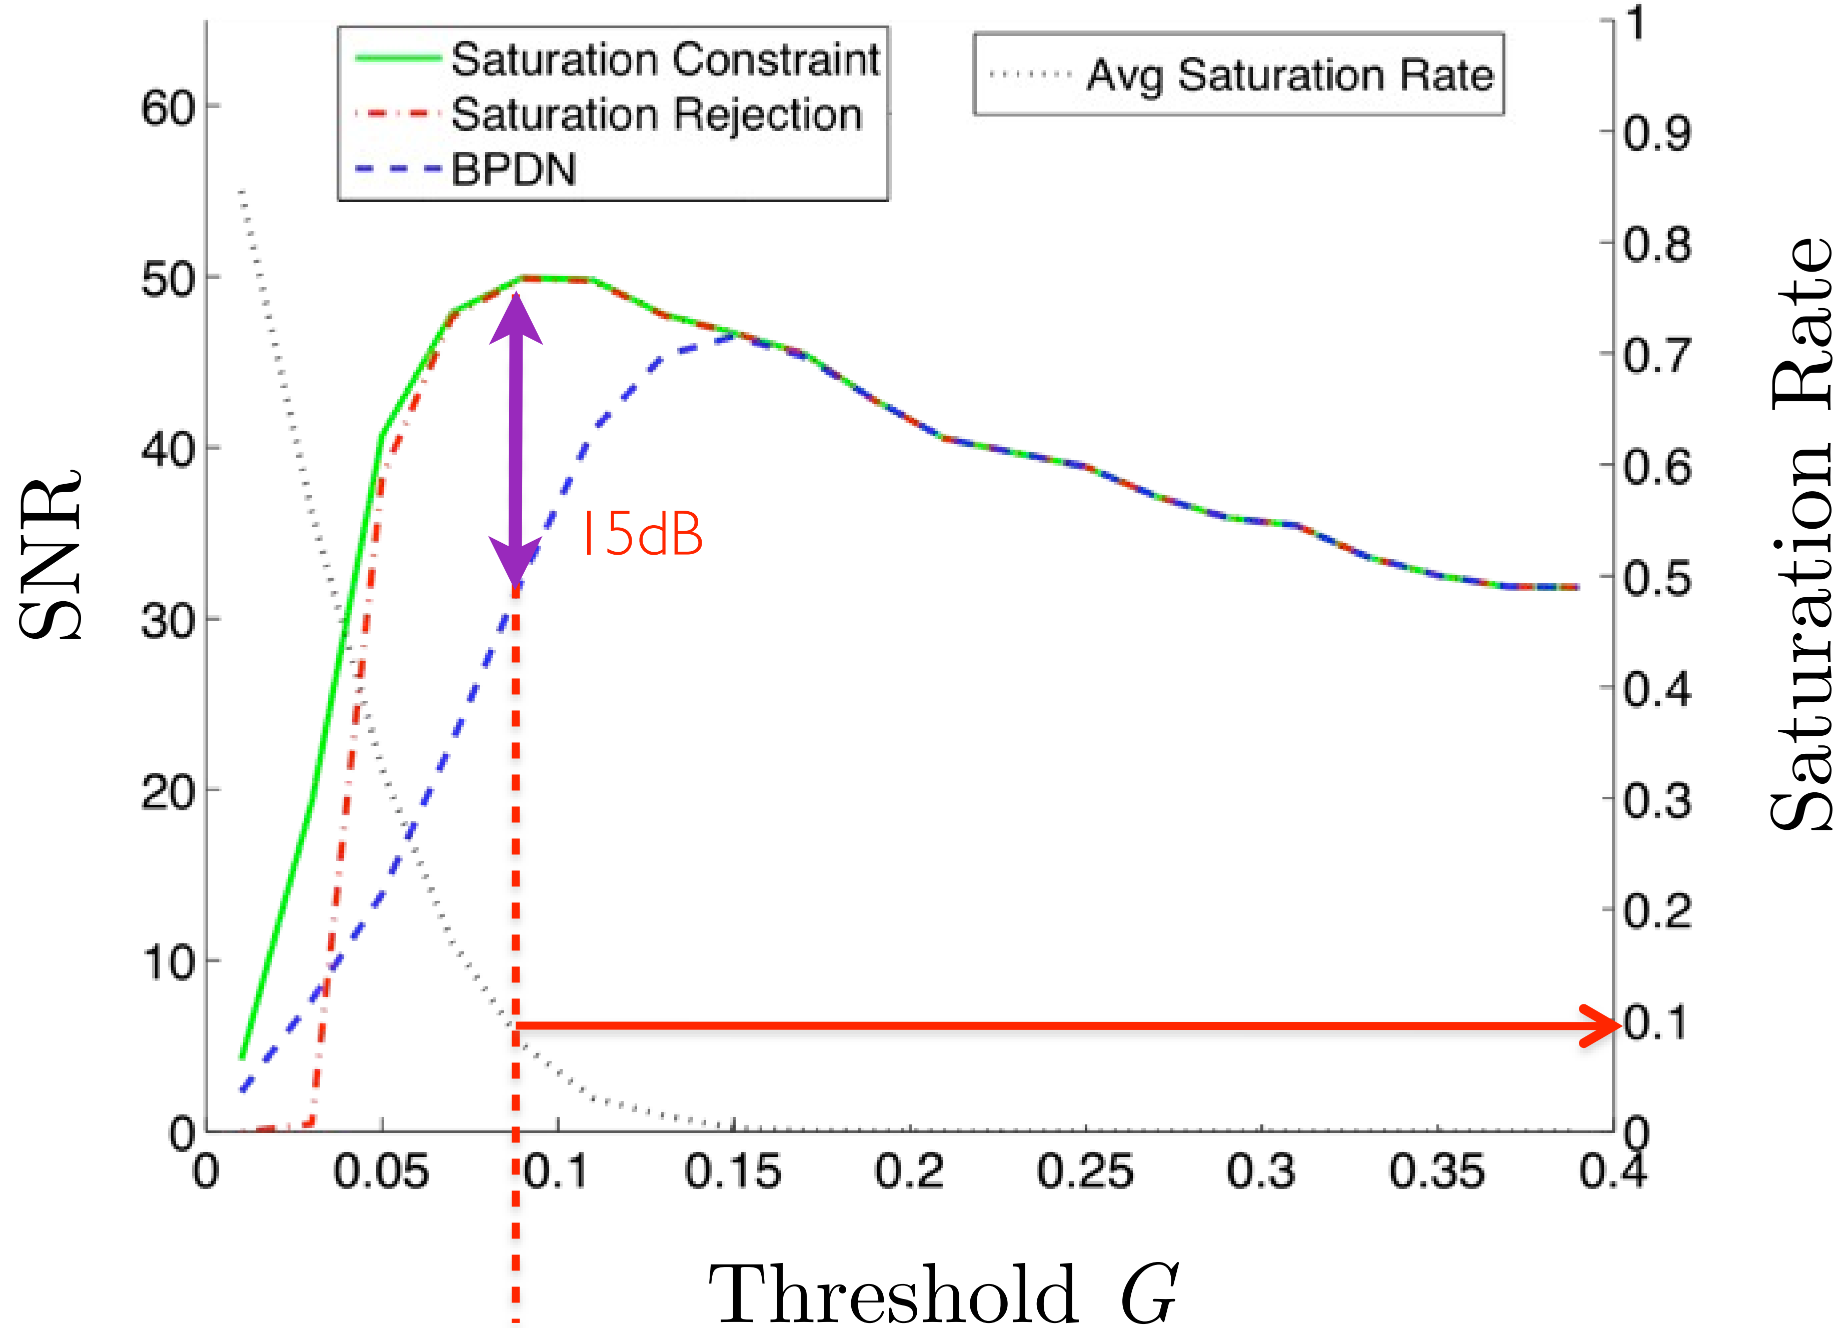

Note: optimal performance requires 10% saturation

J.N. Laska, P.T. Boufounos, M.A. Davenport, R.G.Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing". *Applied and Computational Harmonic Analysis*, *31*(3), 429-443. (2011)

Experimental Results

The “saturation gap”

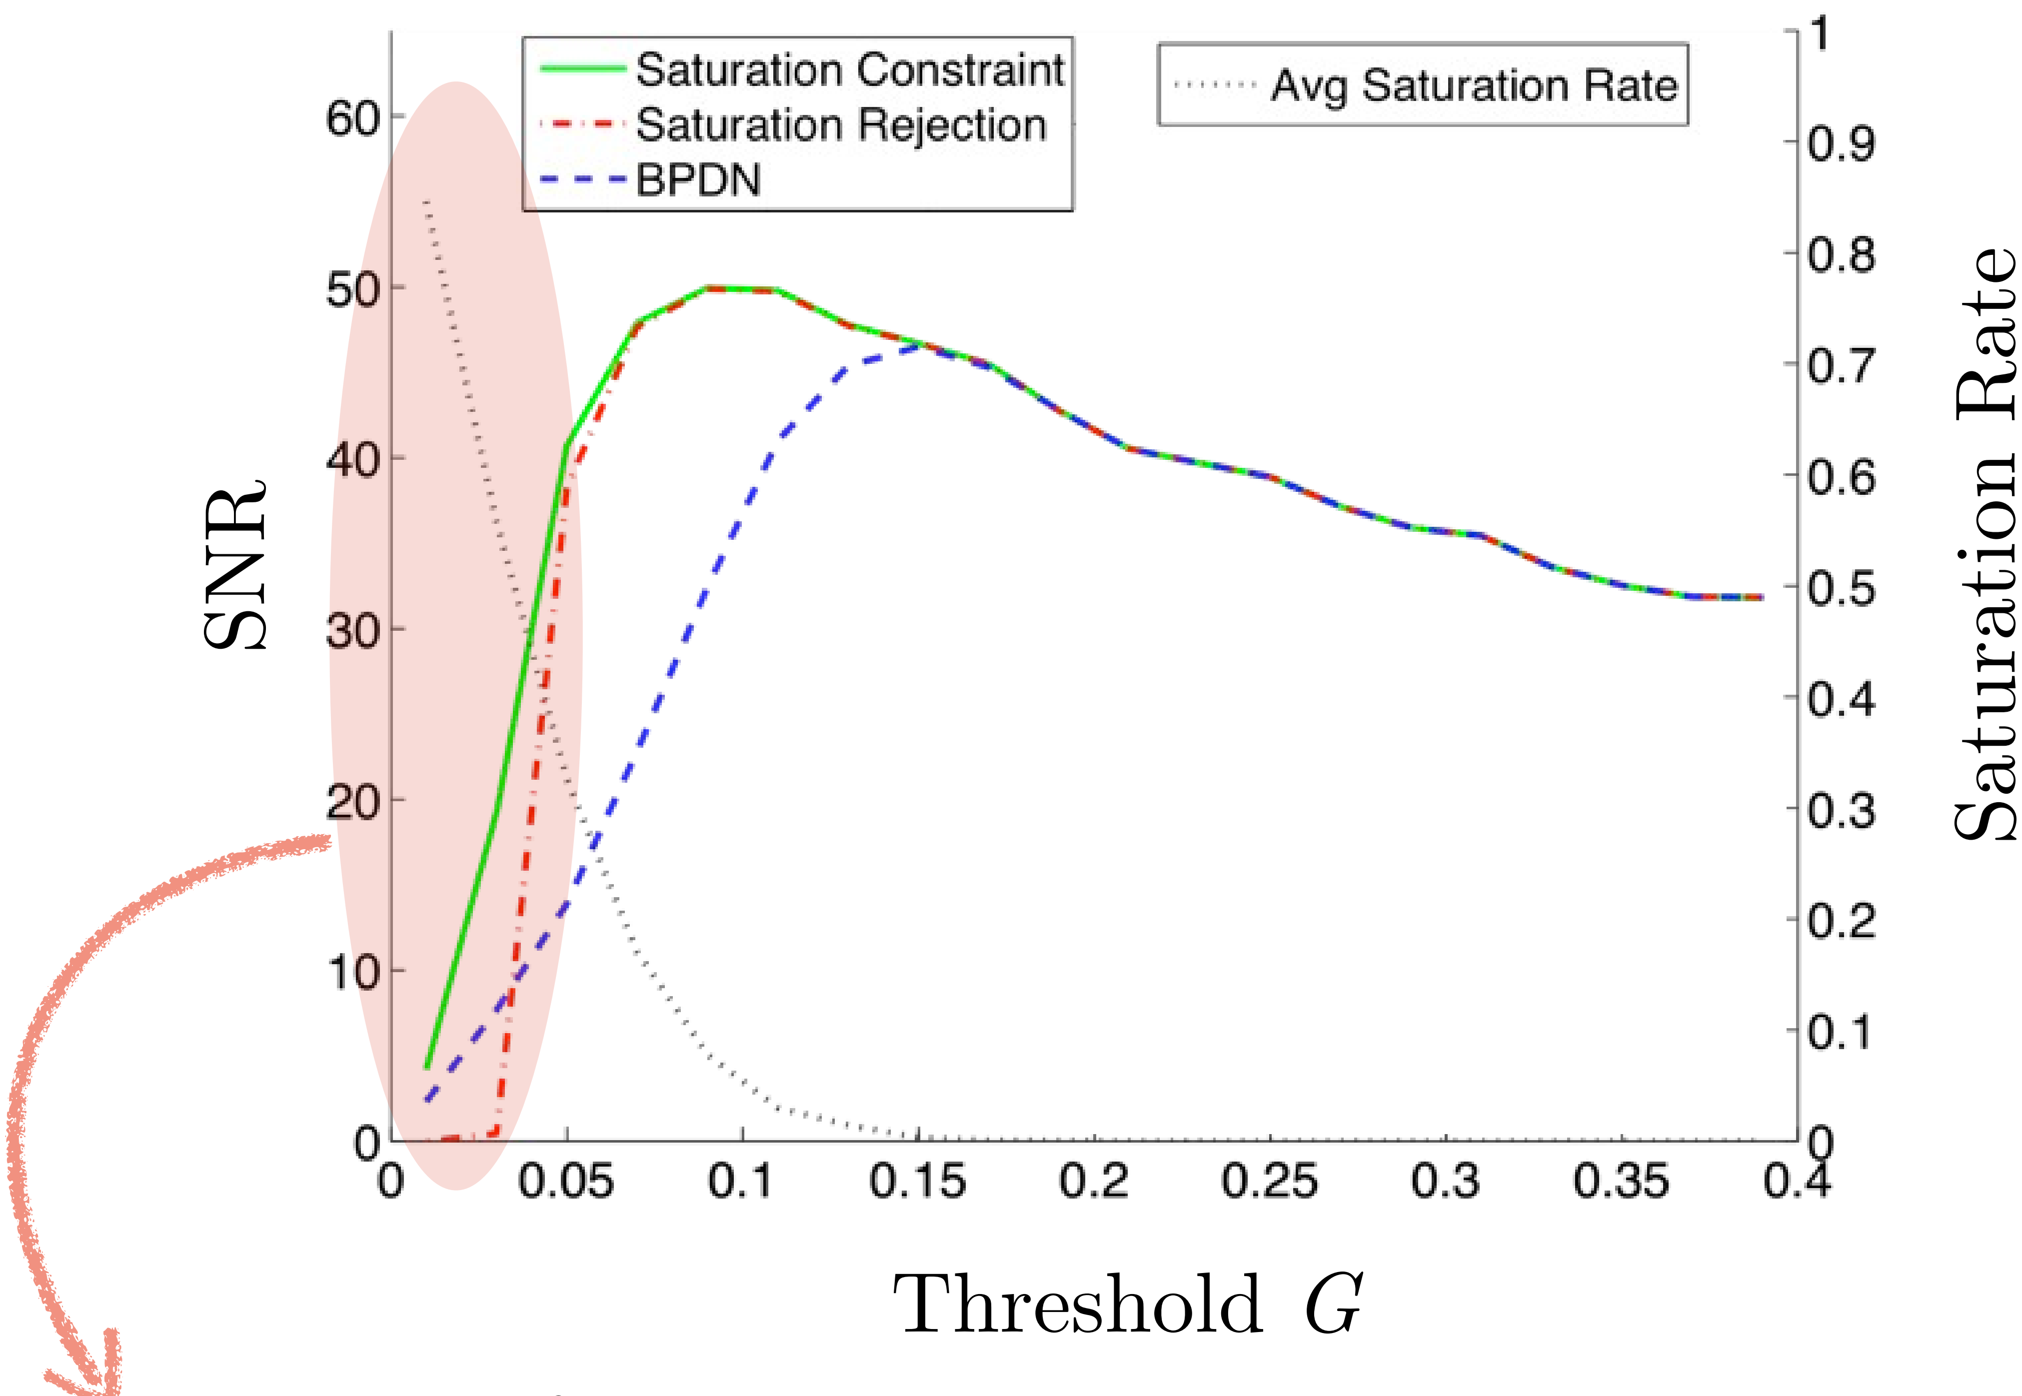

- Majority of measurements saturate
- Recovery fails

J.N. Laska, P.T. Boufounos, M.A. Davenport, R.G.Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing”. *Applied and Computational Harmonic Analysis*, *31*(3), 429-443. (2011)

Experimental Results

The "saturation gap"

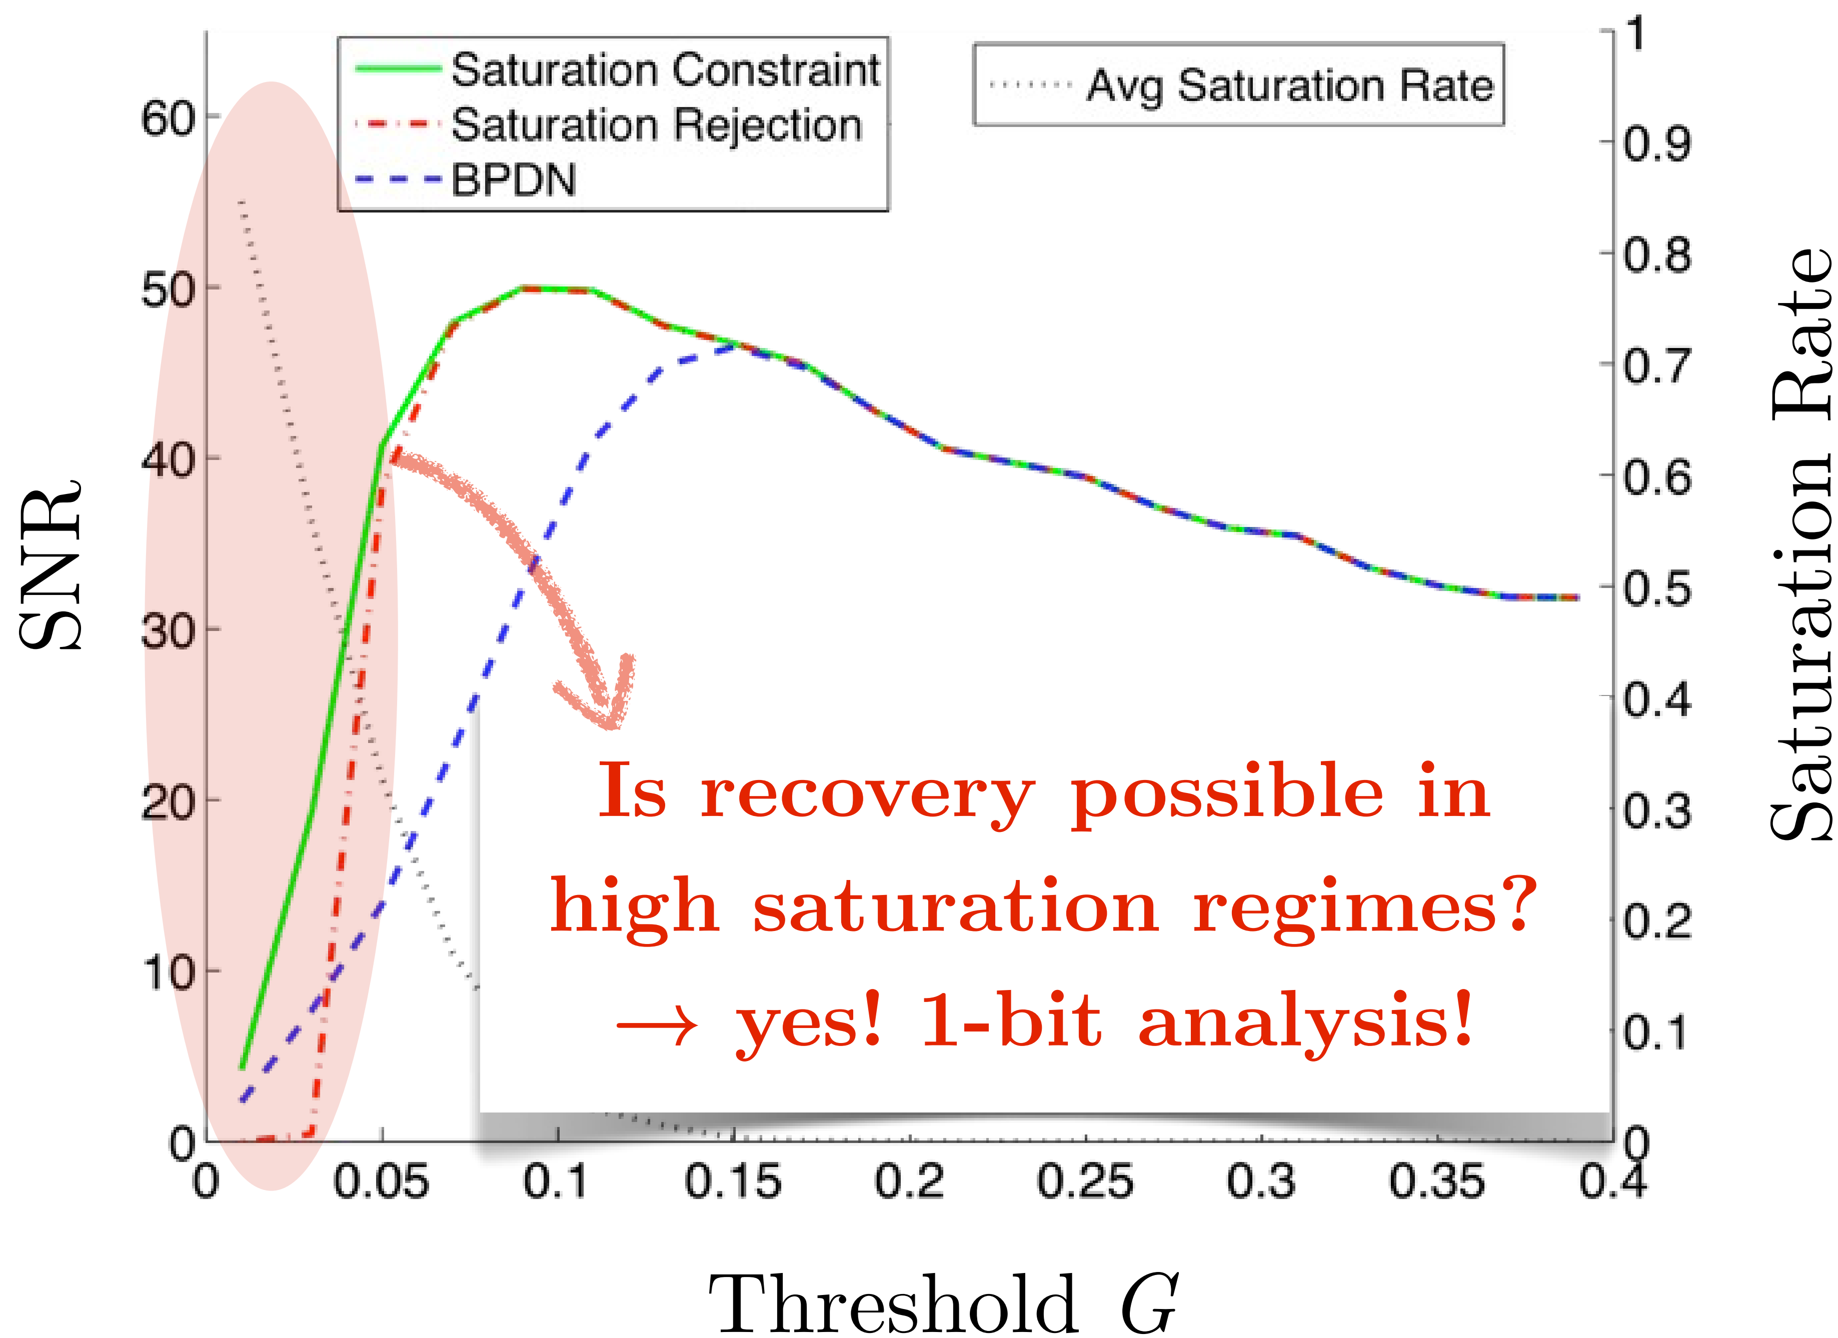

- Majority of measurements saturate
- Recovery fails

J.N. Laska, P.T. Boufounos, M.A. Davenport, R.G.Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing". *Applied and Computational Harmonic Analysis*, *31*(3), 429-443. (2011)

Further Reading

- V. K Goyal, M. Vetterli, N. T. Thao, “Quantized Overcomplete Expansions in RN: Analysis, Synthesis, and Algorithms”, *IEEE Trans. Info. Theory*, 44(1), 1998
- P. T. Boufounos and R. G. Baraniuk, “Quantization of sparse representations,” *Rice University ECE Department Technical Report 0701*. Summary appears in *Proc. Data Compression Conference (DCC)*, Snowbird, UT, March 27-29, 2007
- W. Dai, H. V. Pham, and O. Milenkovic, “Quantized Compressive Sensing”, preprint, 2009
- L. Jacques, D. Hammond, J. Fadili “Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine.” *IEEE Transactions on Information Theory*, 57(1), 559-571, 2011
- J.N. Laska, P.T. Boufounos, M.A. Davenport, R.G.Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing”. *Applied and Computational Harmonic Analysis*, 31(3), 429-443, 2011
- L. Jacques, D. Hammond, J. Fadili, “Stabilizing Nonuniformly Quantized Compressed Sensing with Scalar Companders”, arXiv:1206.6003, 2012
- Güntürk, C. S., Lammers, M., Powell, A. M., Saab, R., & Yılmaz, Ö. “Sobolev duals for random frames and ΣΔ quantization of compressed sensing measurements”. *Foundations of Computational Mathematics*, 13(1), 1-36, 2013
- A. M. Powell, J.T. Whitehouse, “Error bounds for consistent reconstruction: random polytopes and coverage processes”, arXiv:1405.7094, 2013
- L Jacques, “Error Decay of (almost) Consistent Signal Estimations from Quantized Random Gaussian Projections”, arXiv:1406.0022, 2014
- P. T. Boufounos, L. Jacques, F. Krahmer, R. Saab, “Quantization and Compressive Sensing”, arXiv:1405.1194

Part 2

Extreme quantization: 1-bit compressed sensing

Outline:

1. Context
2. Theoretical performance limits
3. Stable embeddings: angles are preserved
4. Generalized Embeddings
5. 1-bit CS Reconstructions?
6. Playing with thresholds in 1-bit CS

1. Context

Central question: 1-bit sampling?

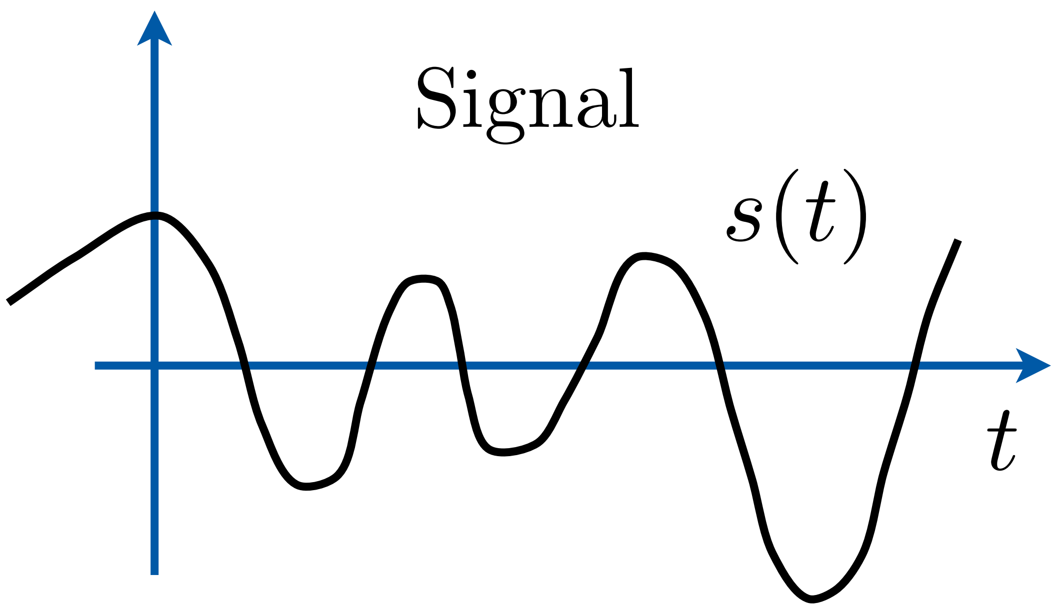

Central question: 1-bit sampling?

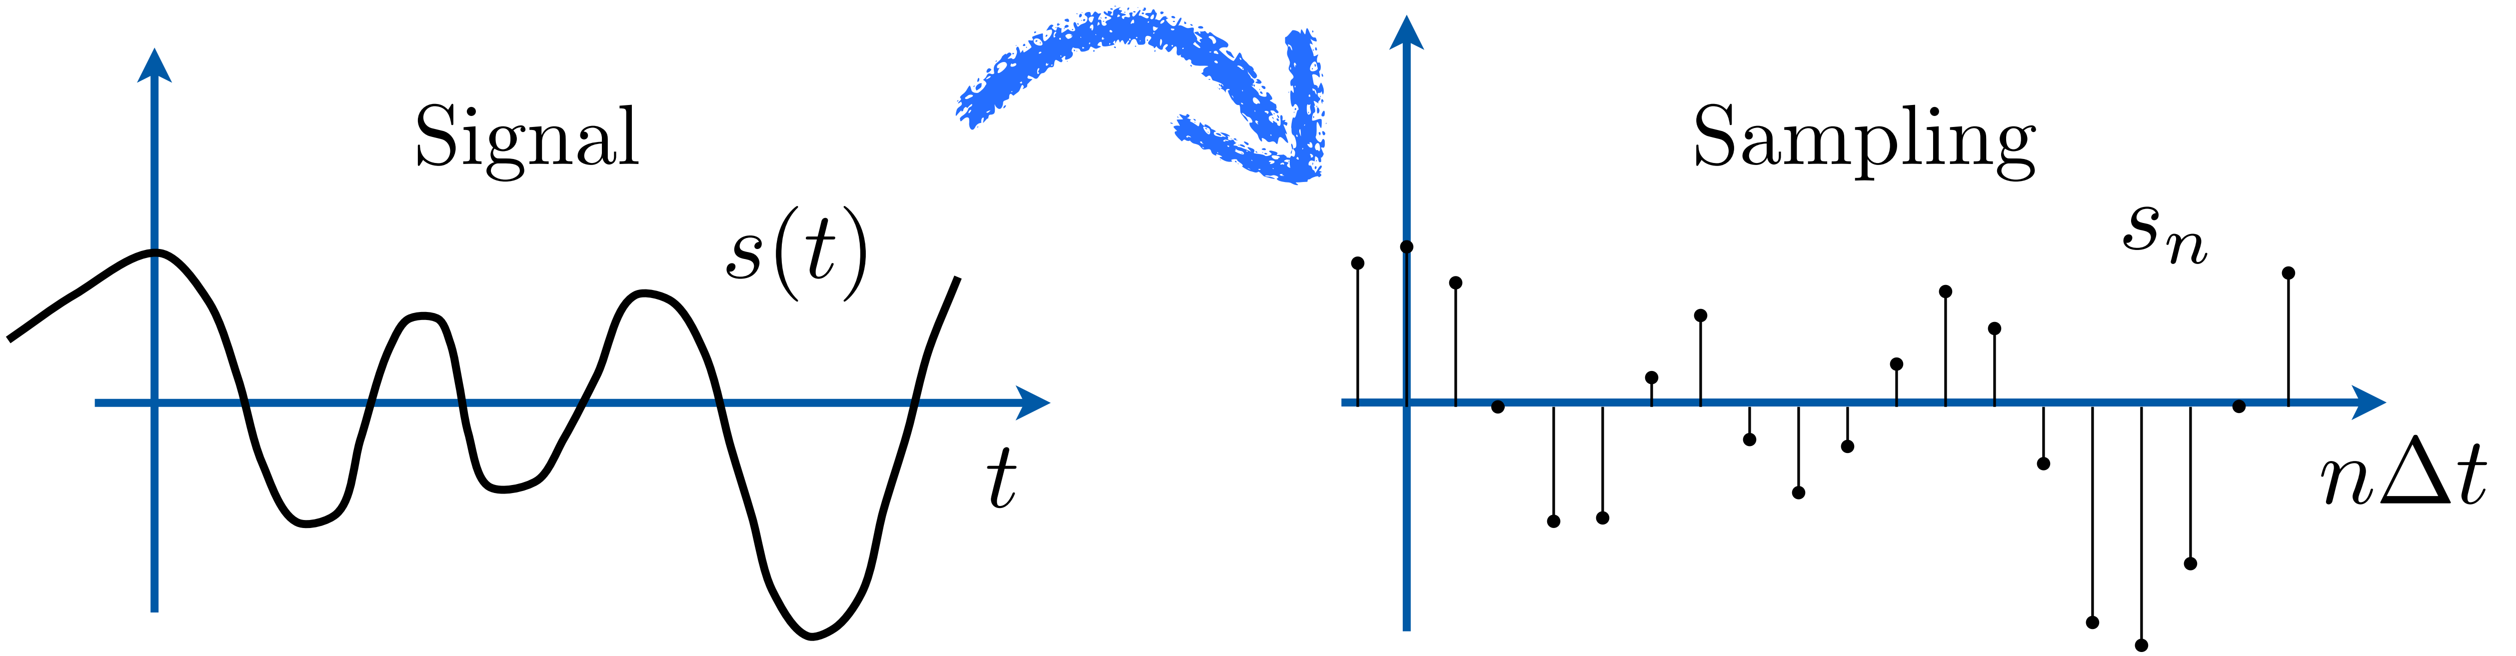

Central question: 1-bit sampling?

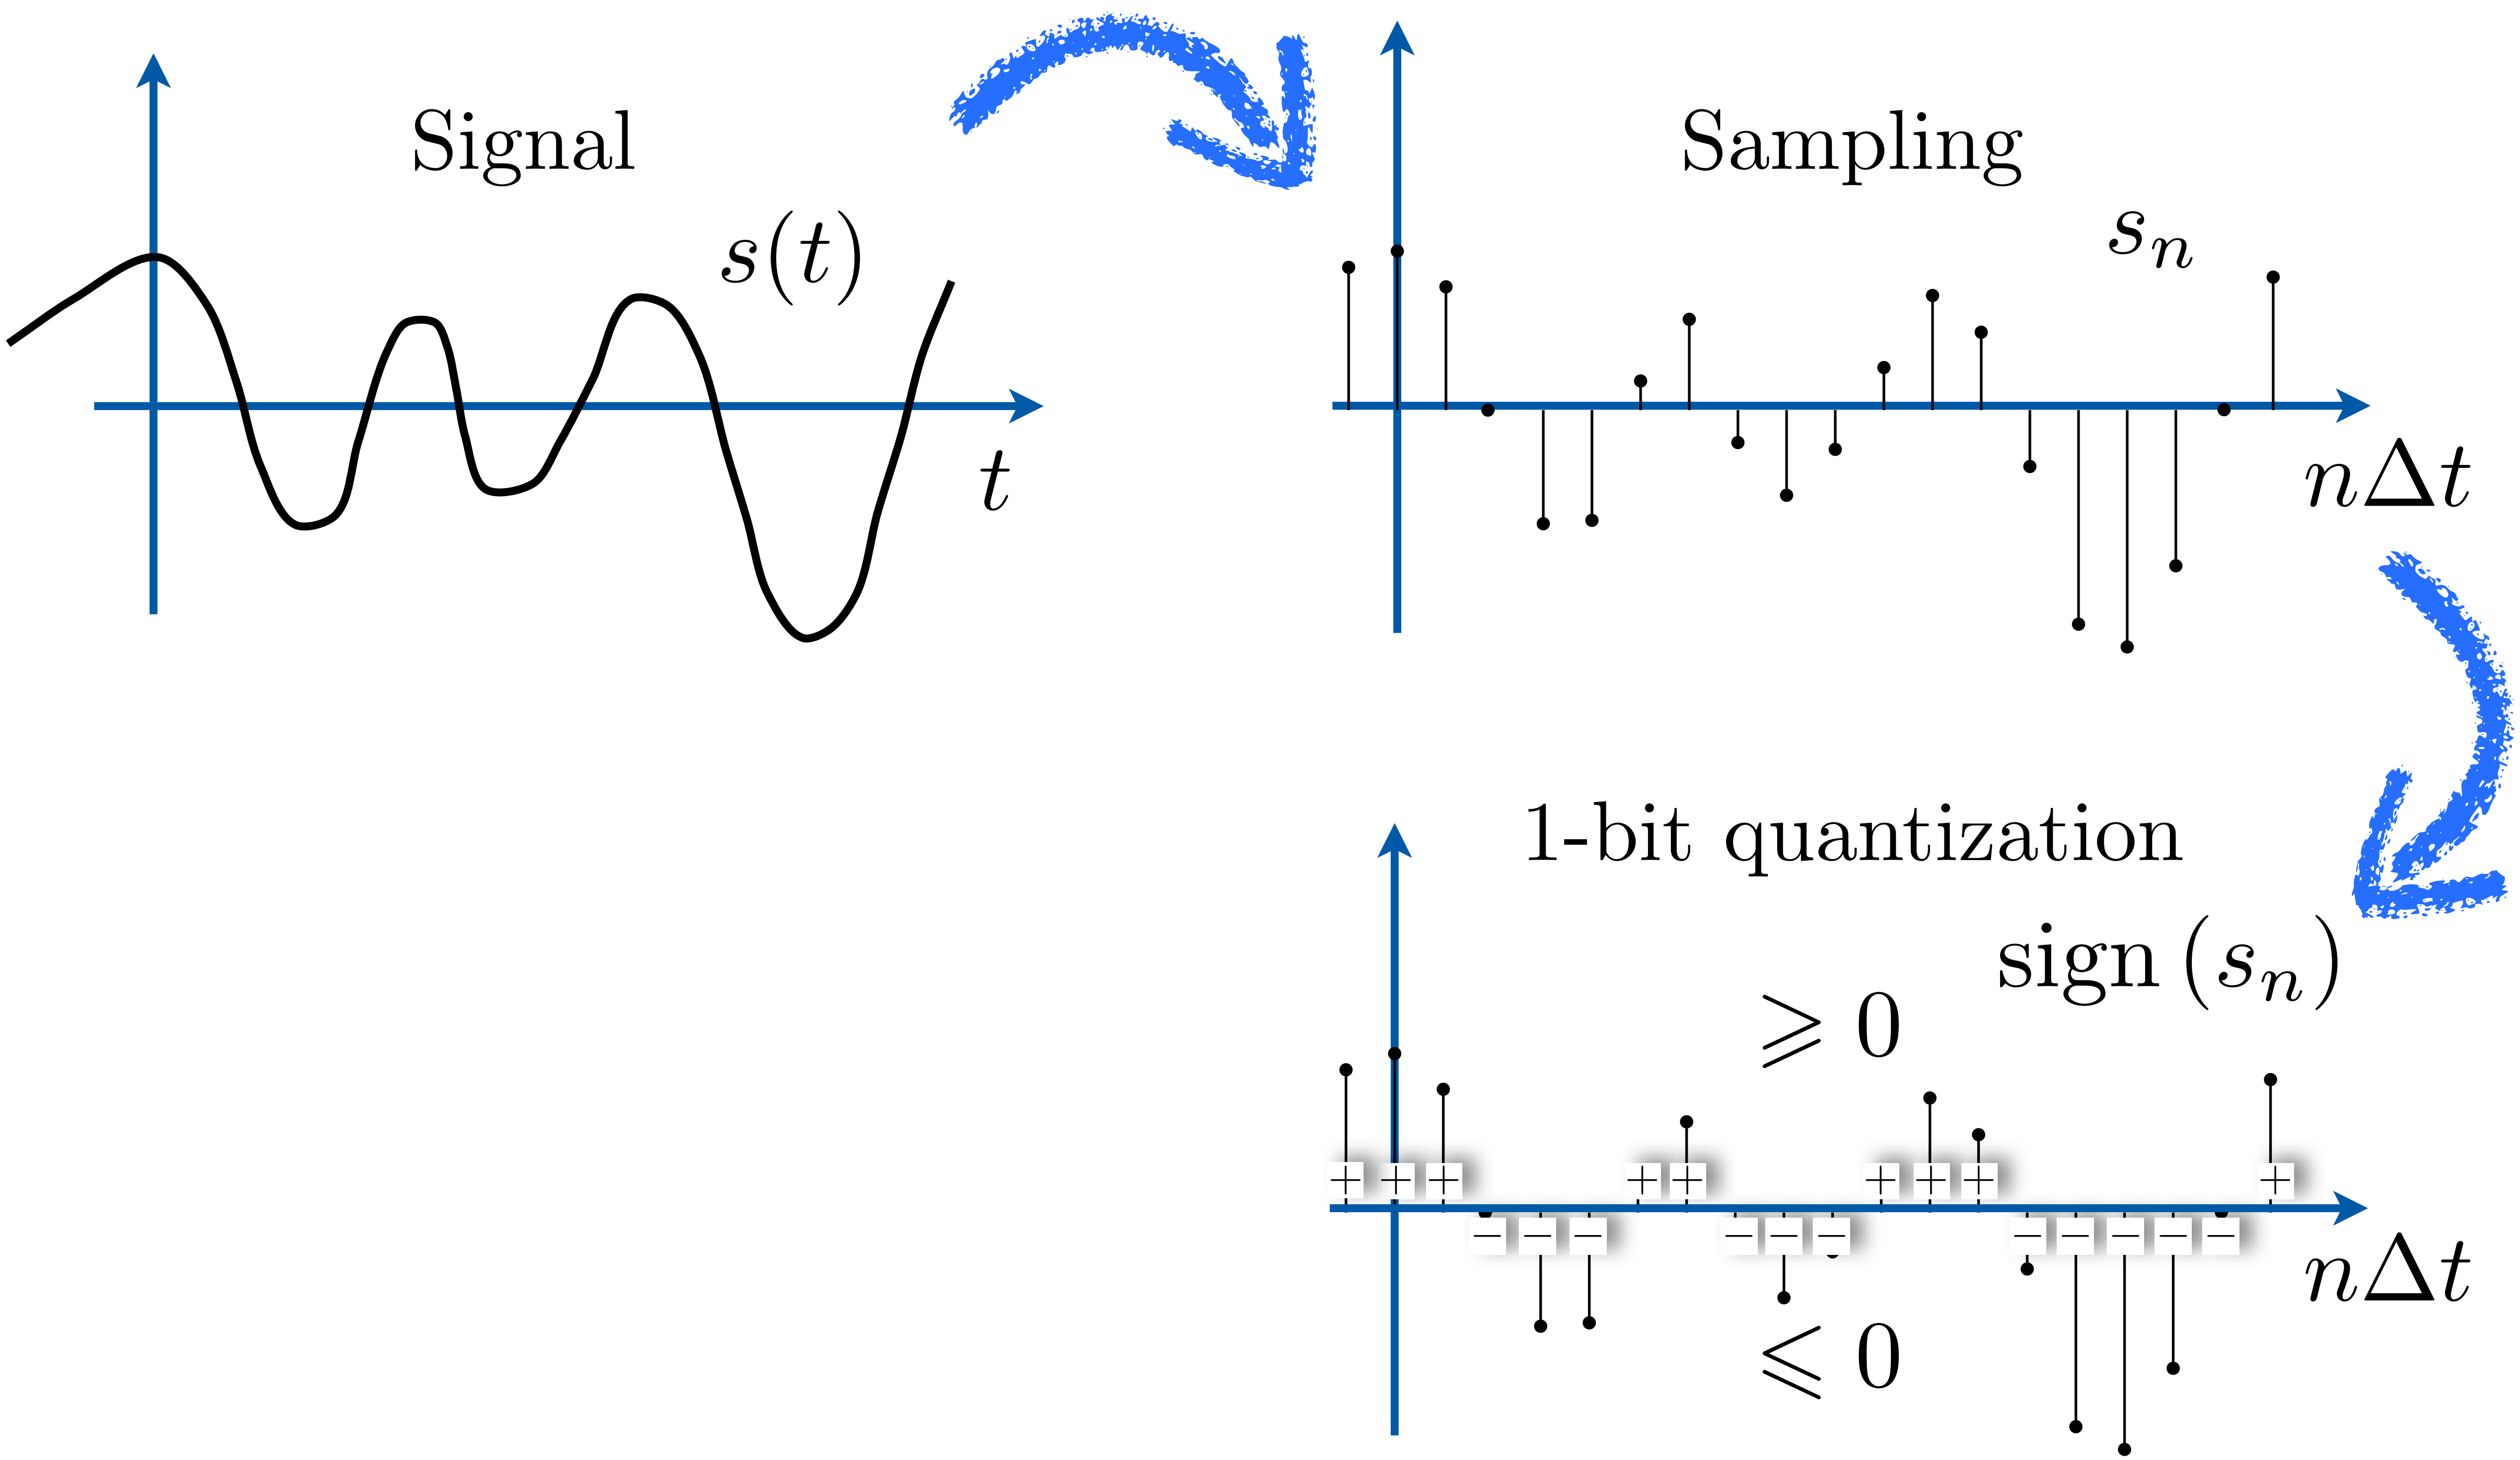

Central question: 1-bit sampling?

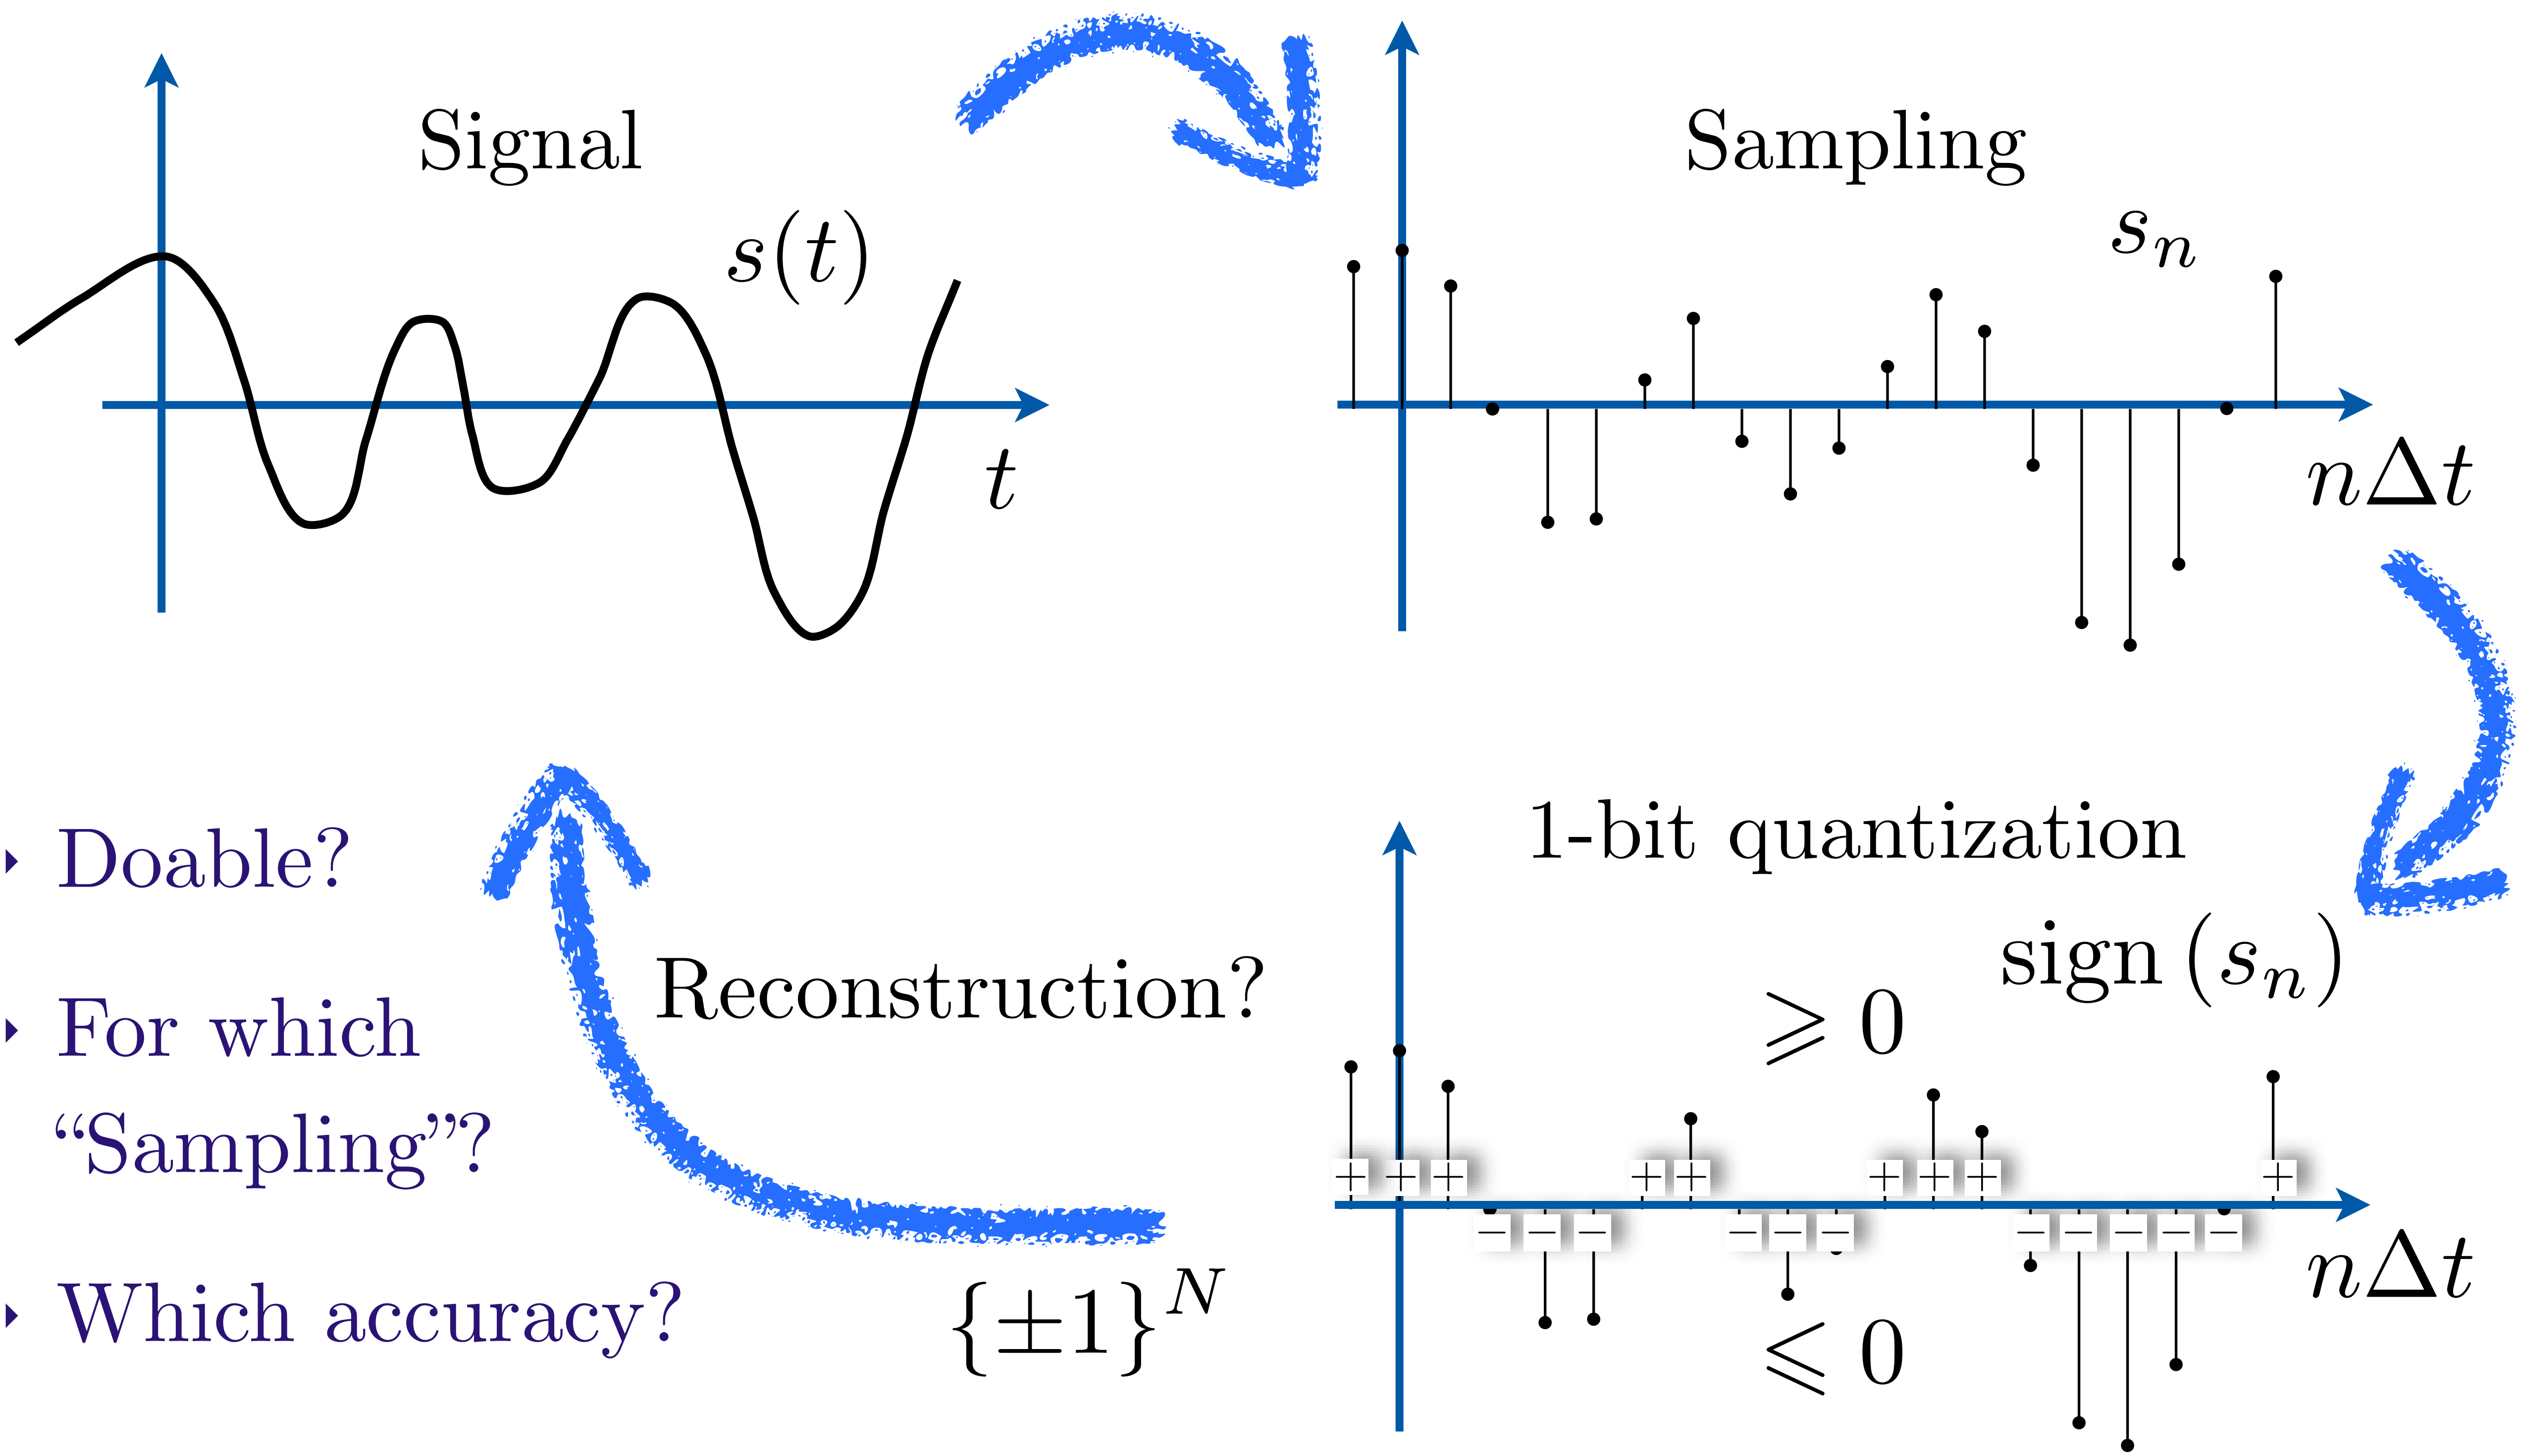

- Doable?
- For which "Sampling"?
- Which accuracy?

Why 1-bit? Very Fast Quantizers!

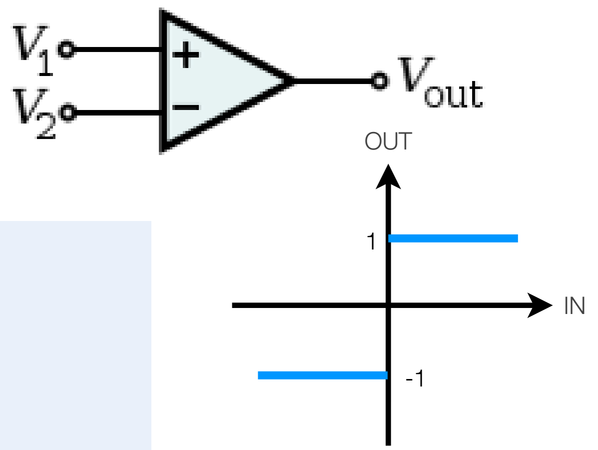

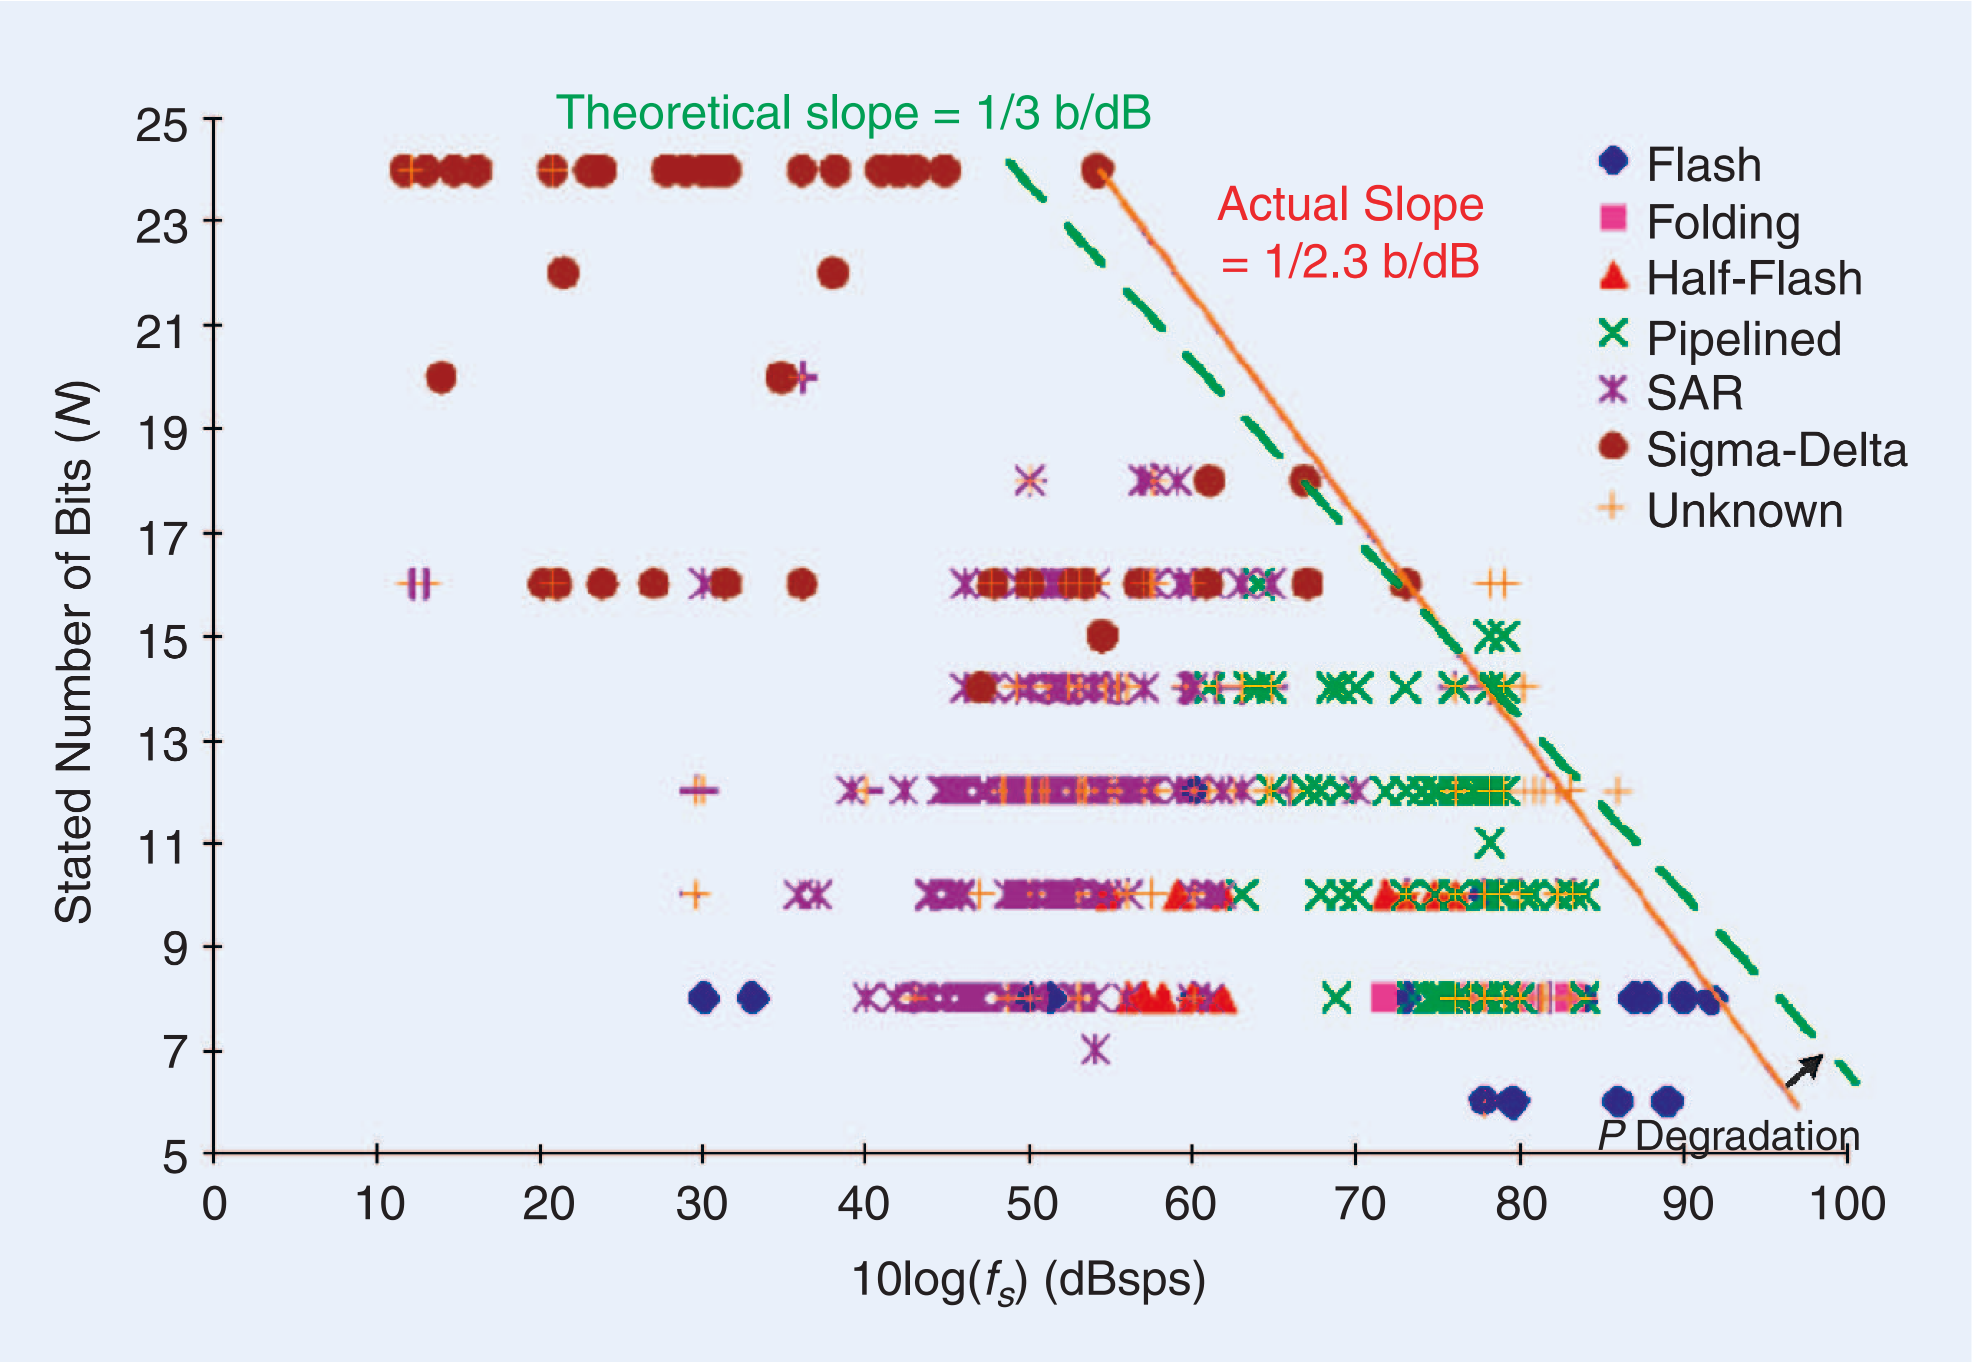

[FIG1] Stated number of bits versus sampling rate.

[From "Analog-to-digital converters" B. Le, T.W. Rondeau, J.H. Reed, and C.W.Bostian, IEEE Sig. Proc. Magazine, Nov 2005]

Why 1-bit? Very Fast Quantizers!

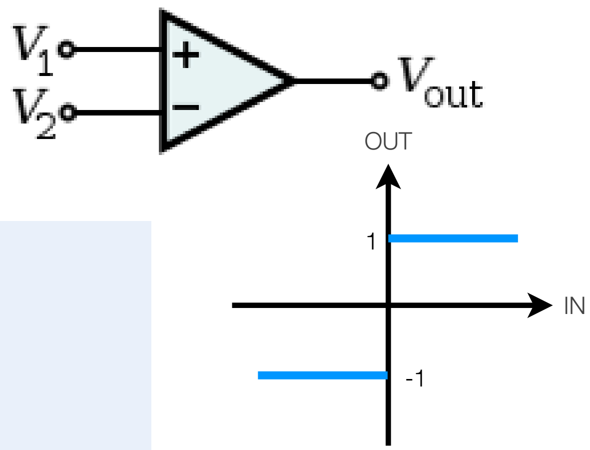

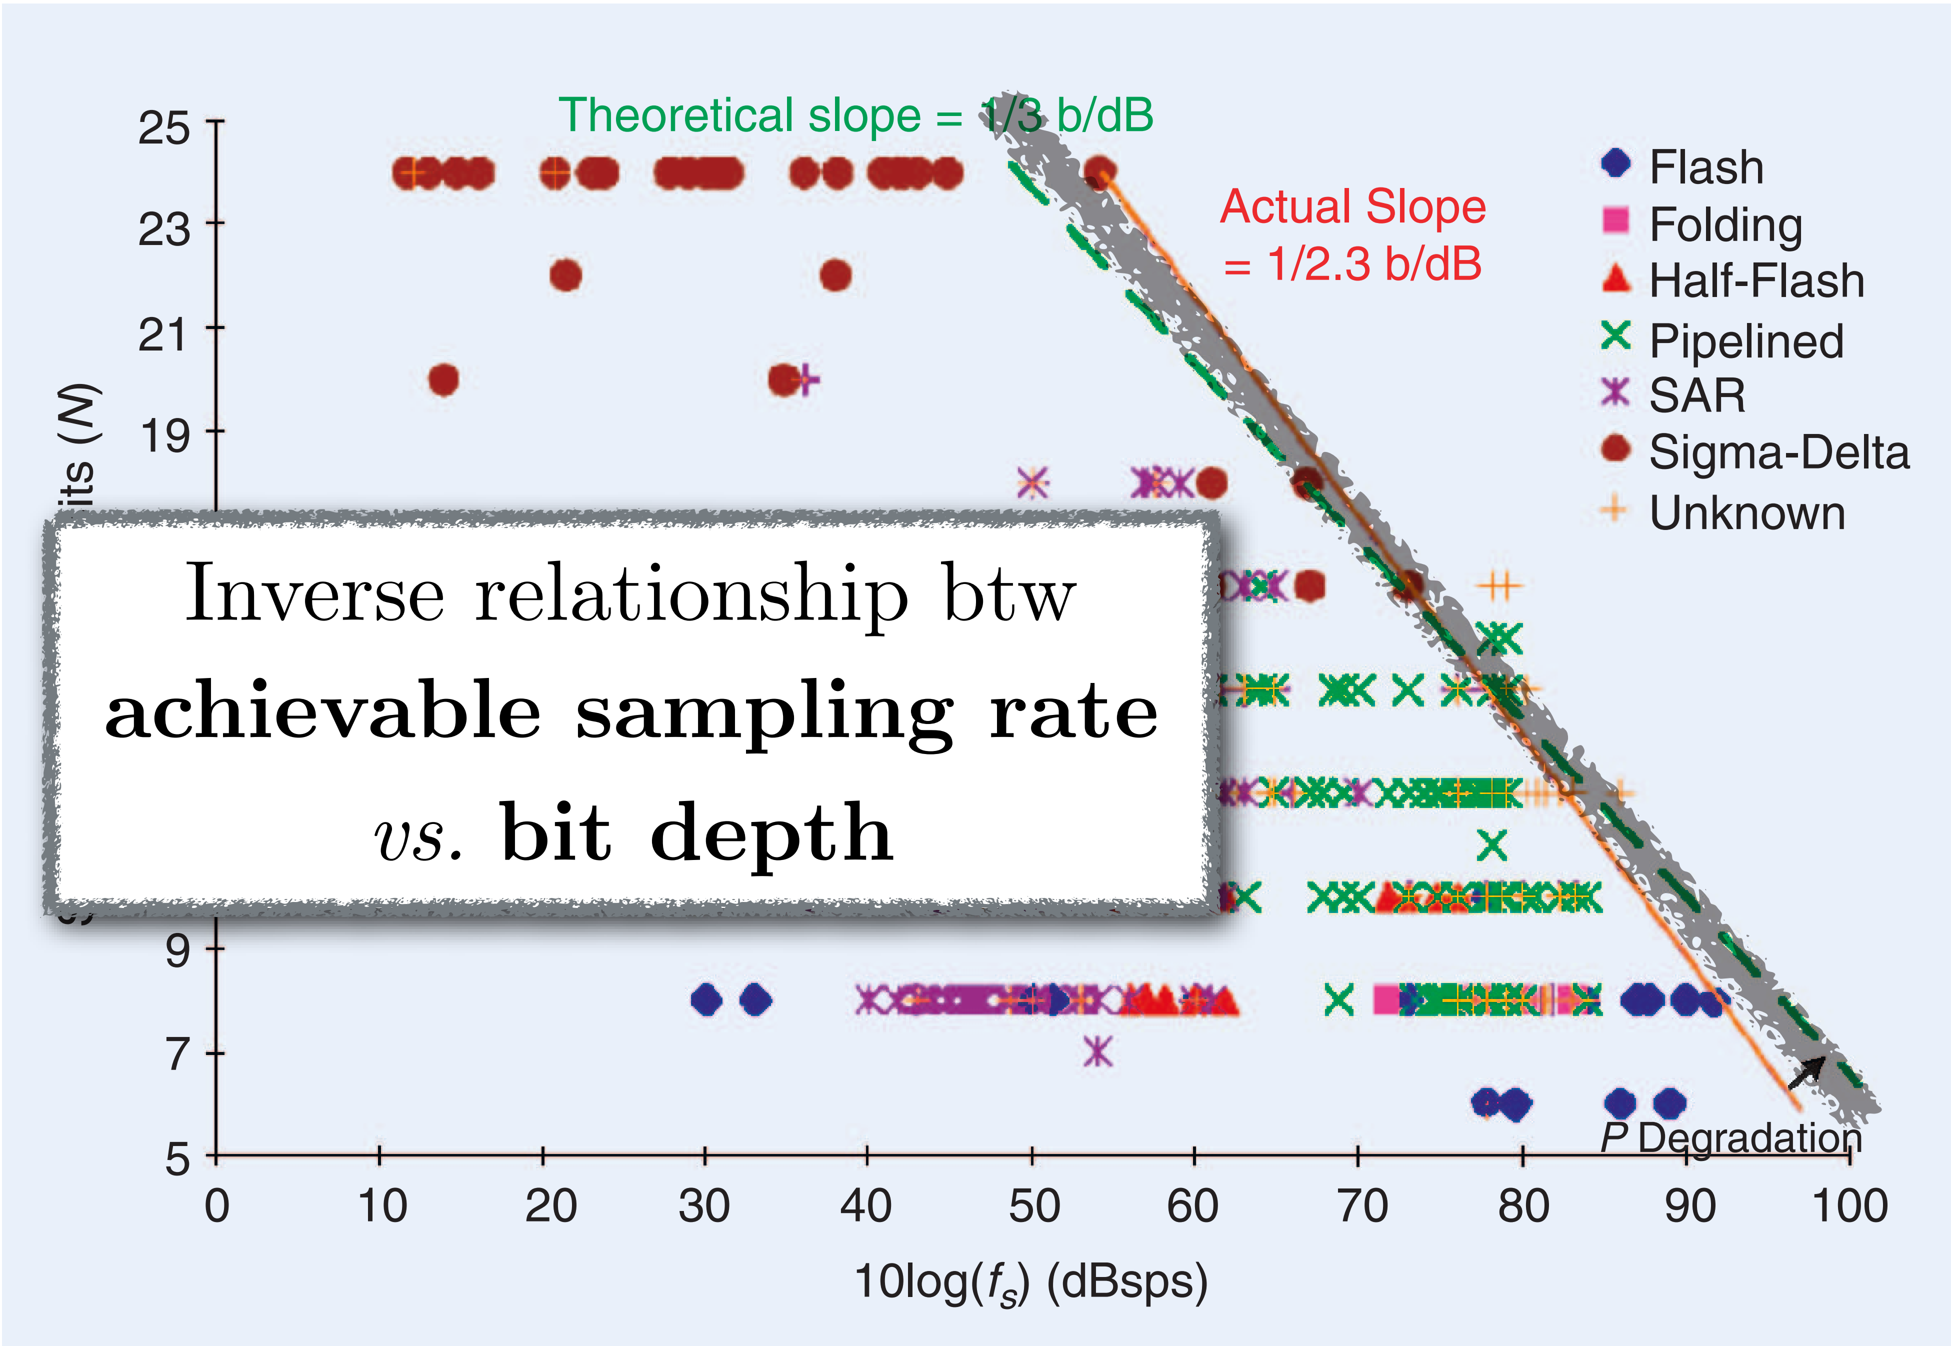

Inverse relationship btw **achievable sampling rate** *vs.* **bit depth**

[FIG1] Stated number of bits versus sampling rate.

[From "Analog-to-digital converters" B. Le, T.W. Rondeau, J.H. Reed, and C.W.Bostian, IEEE Sig. Proc. Magazine, Nov 2005]

Compressed Sensing

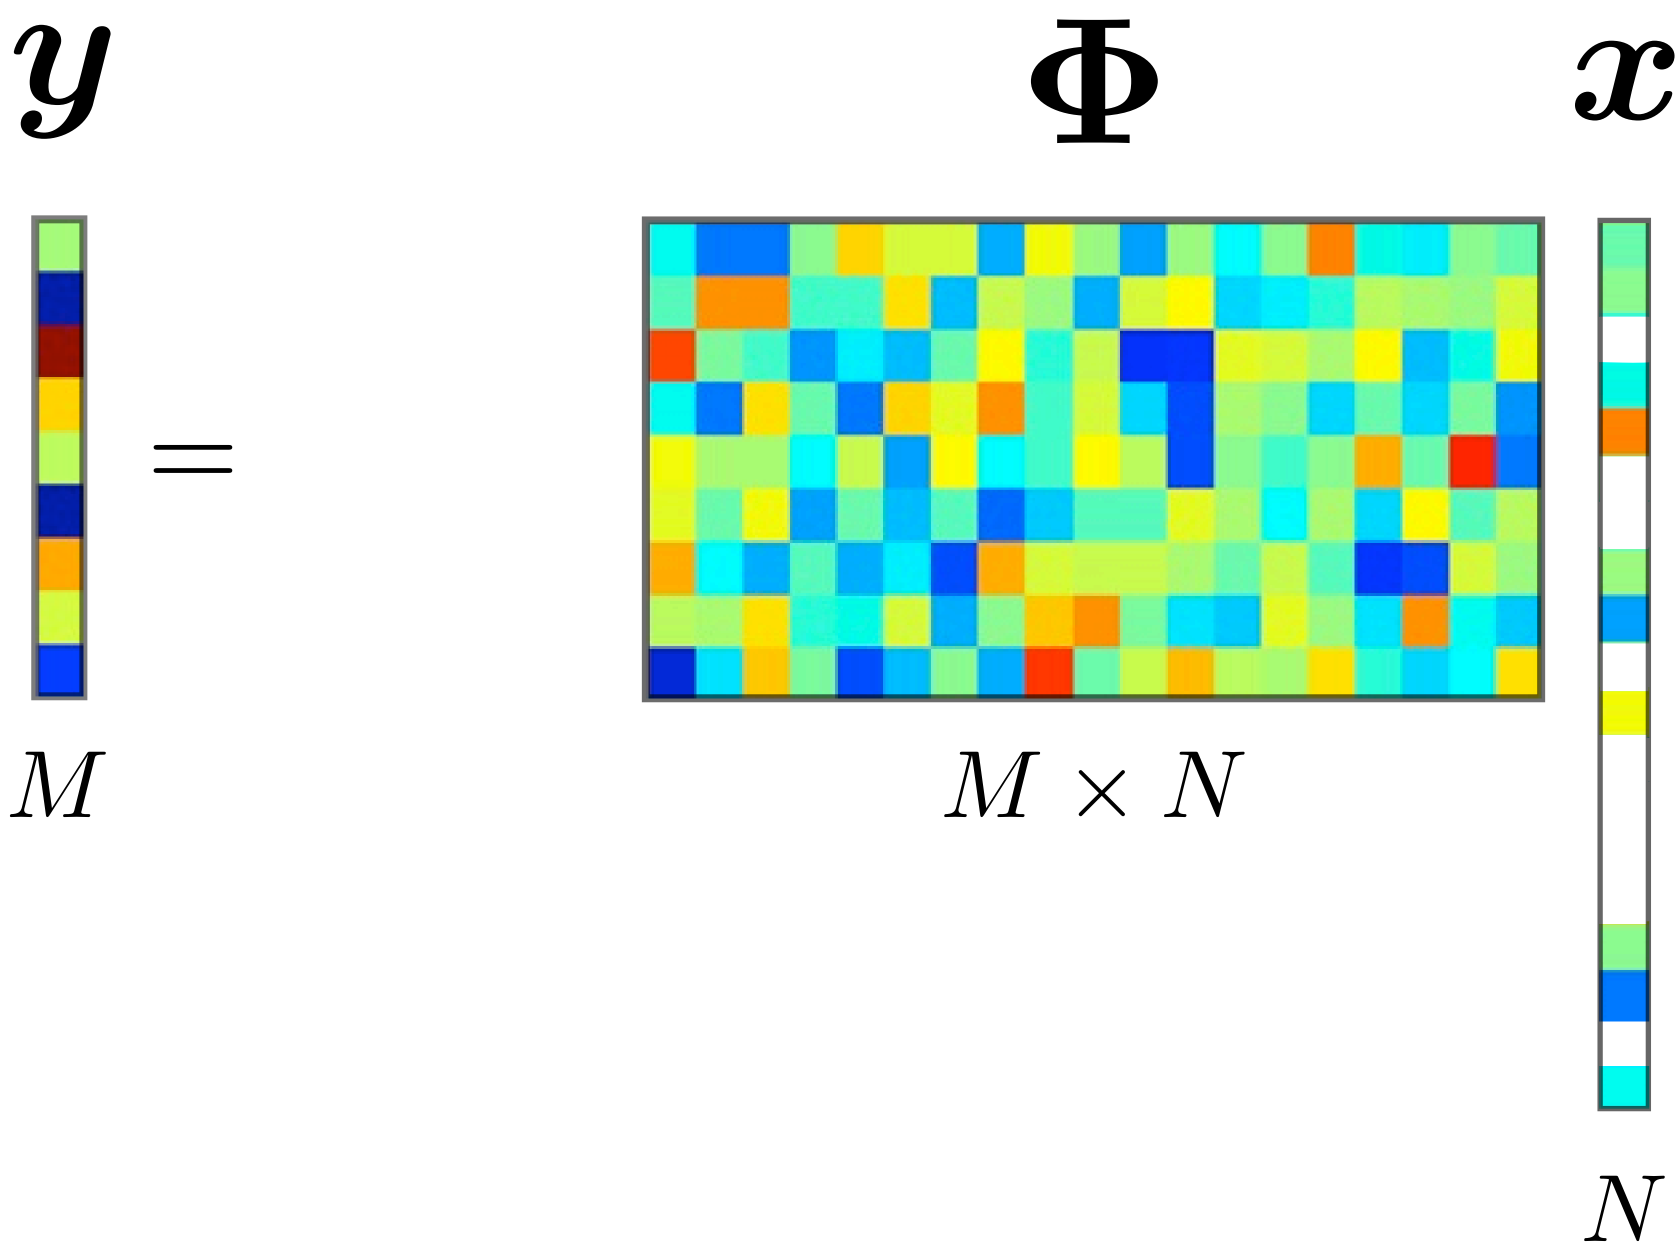

1-bit Compressed Sensing

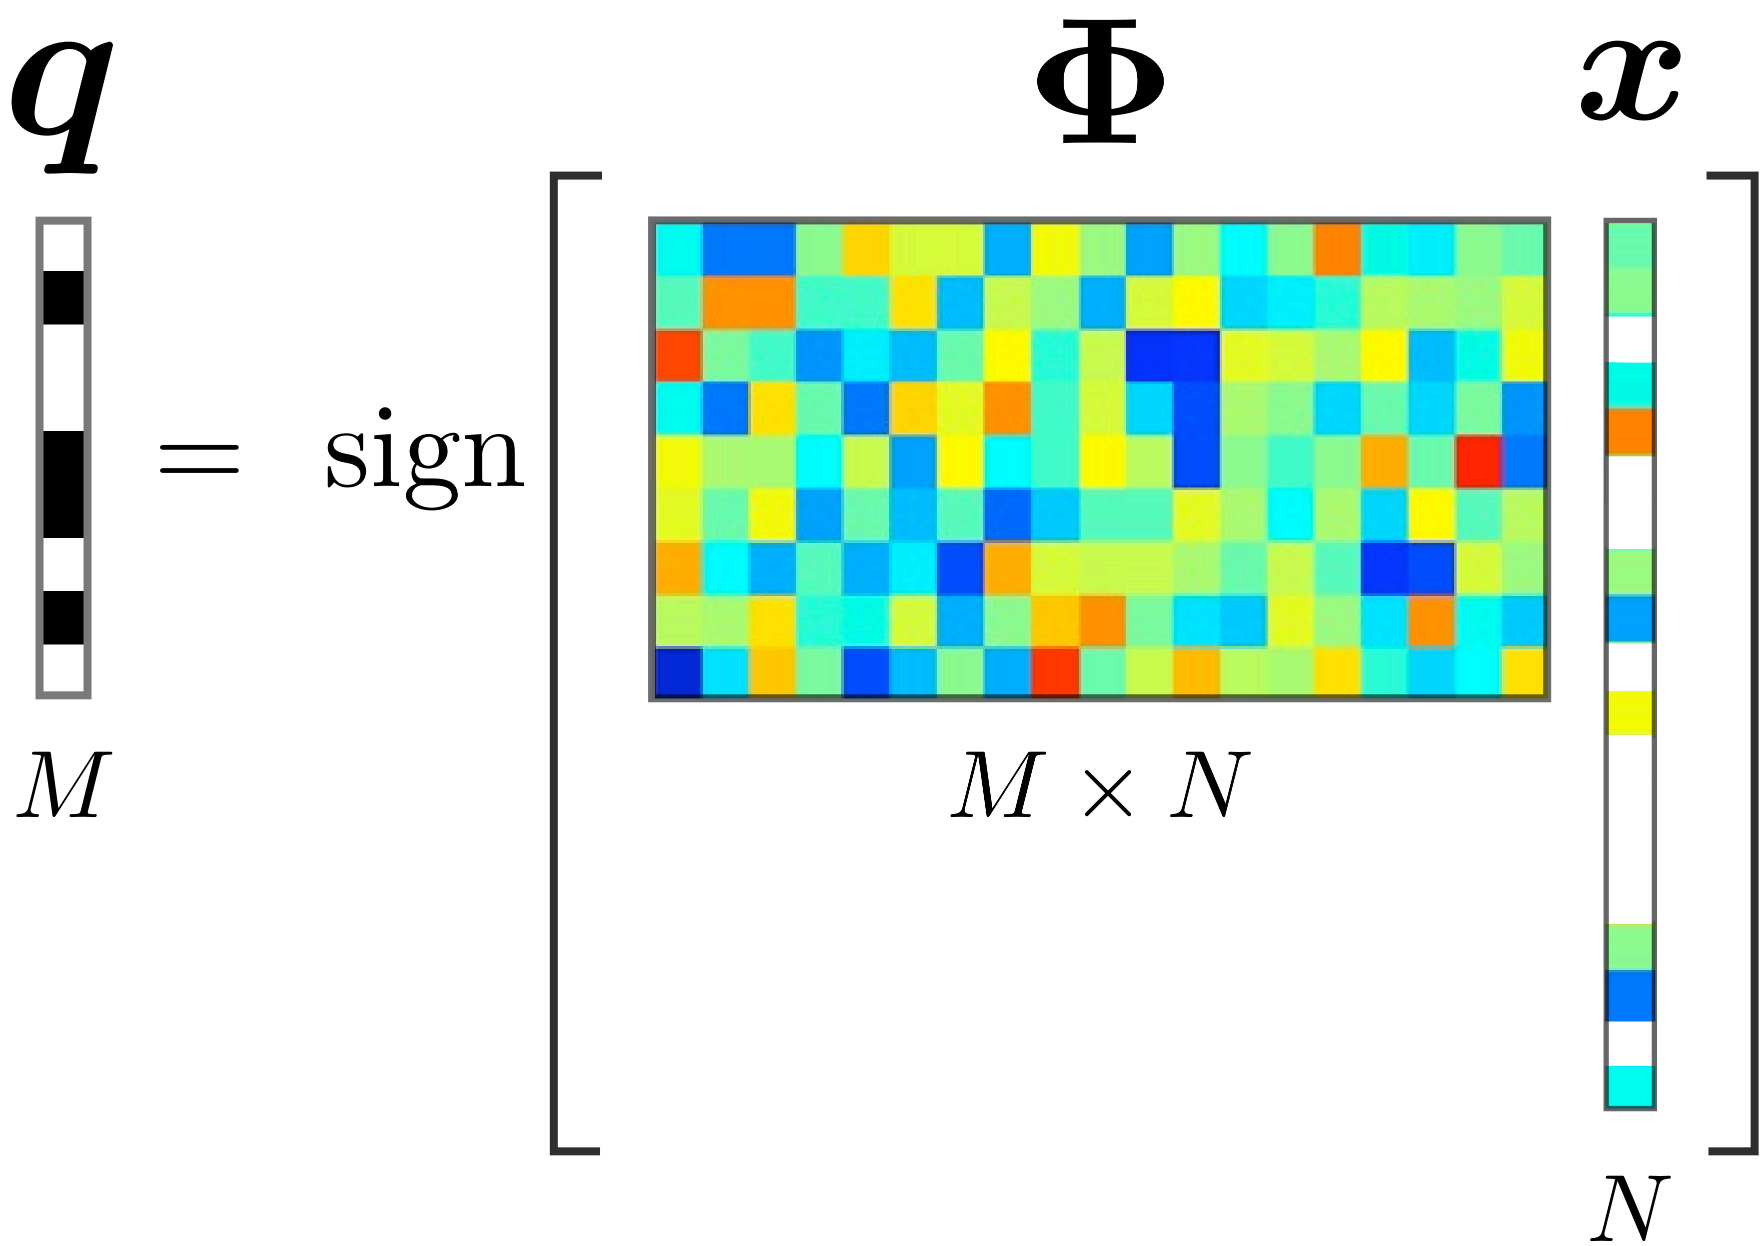

with: $\operatorname{sign} t = \begin{cases} 1 & \text{if } t > 0 \\ -1 & \text{if } t \leqslant 0 \end{cases}$ component-wise

1-bit Compressed Sensing

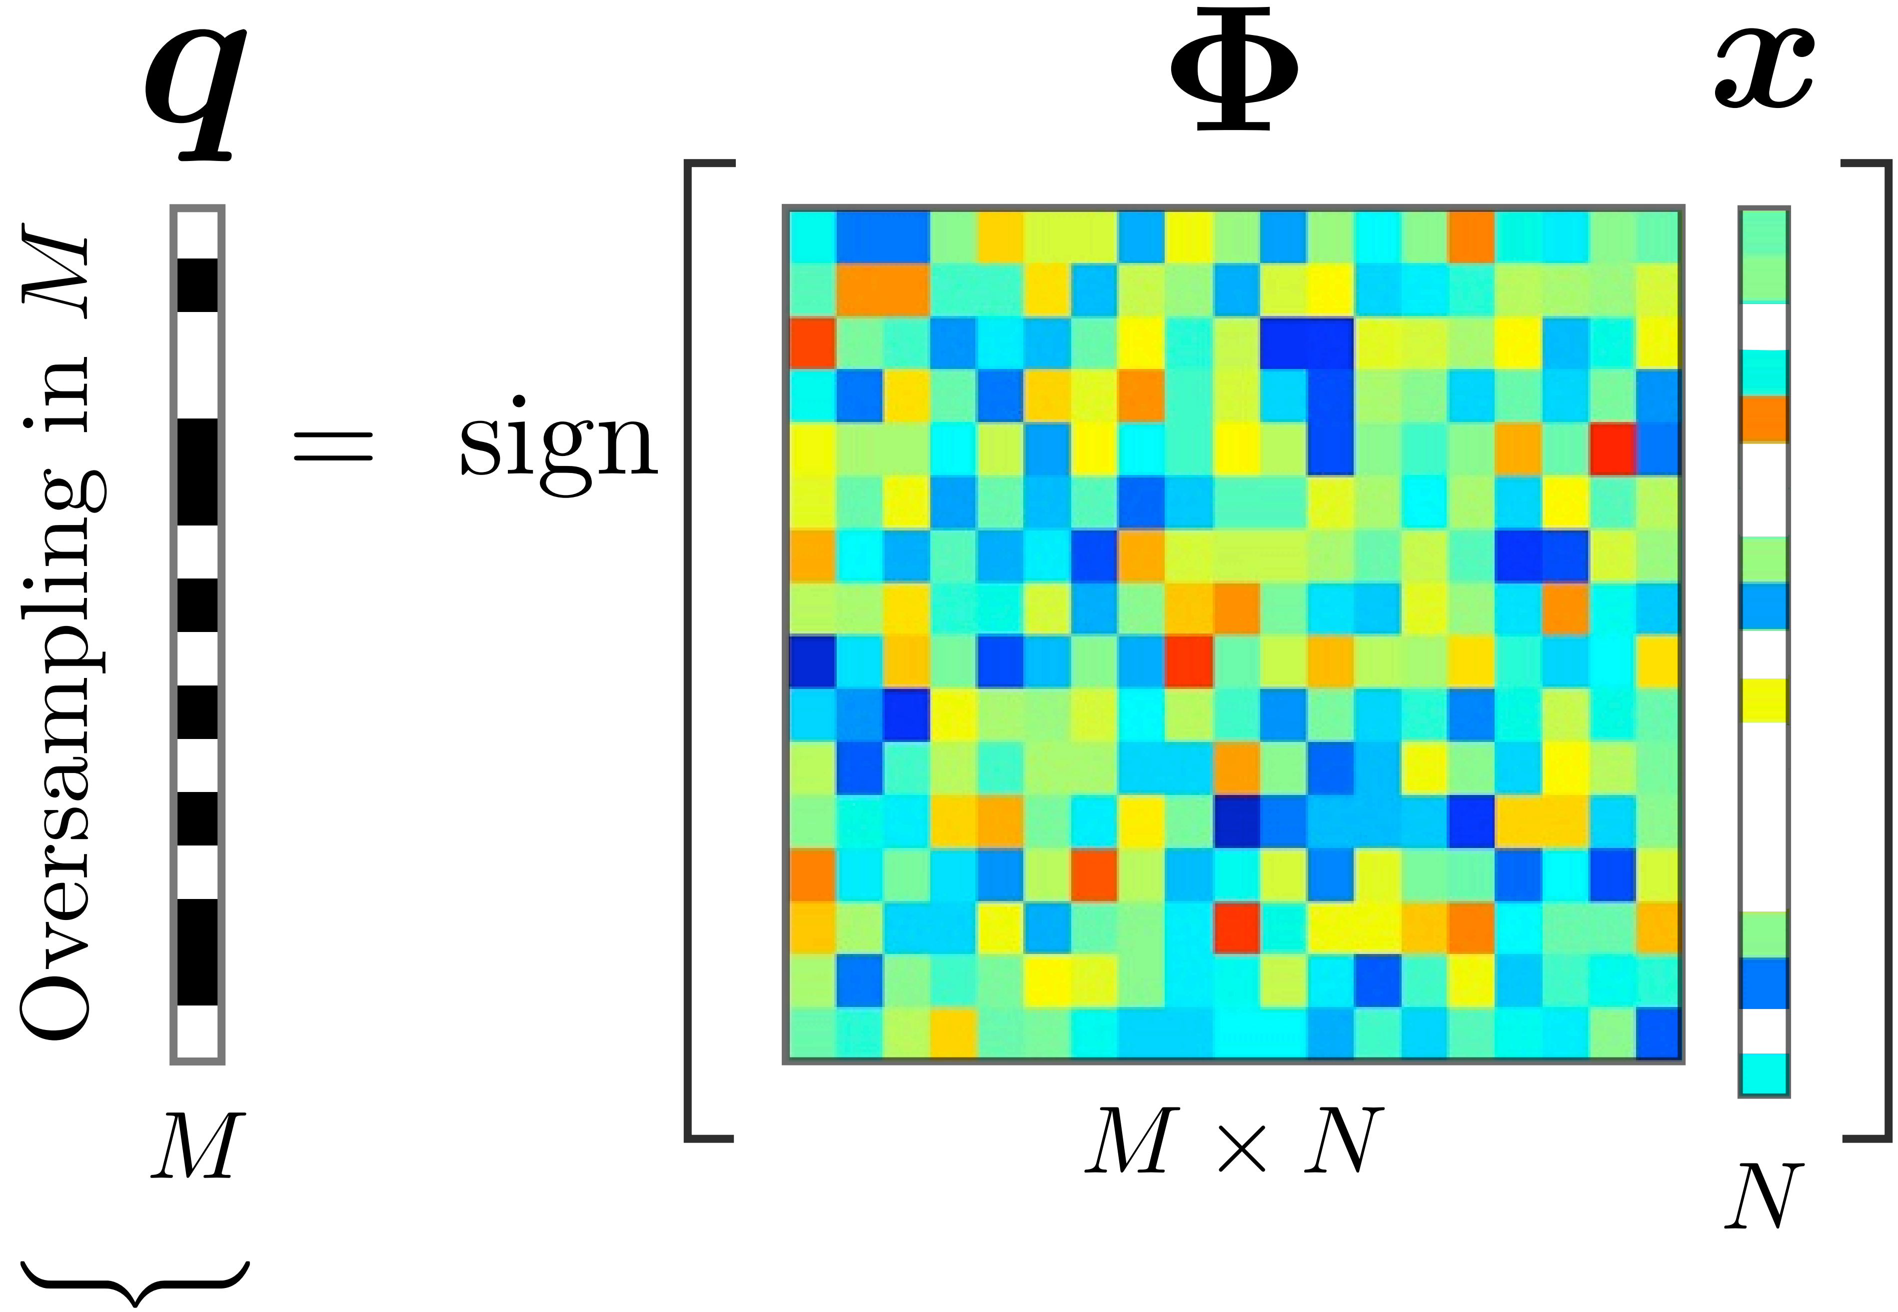

M-bits! But, which information inside $\boldsymbol{q}$?

1-bit ~~Compressed~~ Computational Sensing

bits matter!

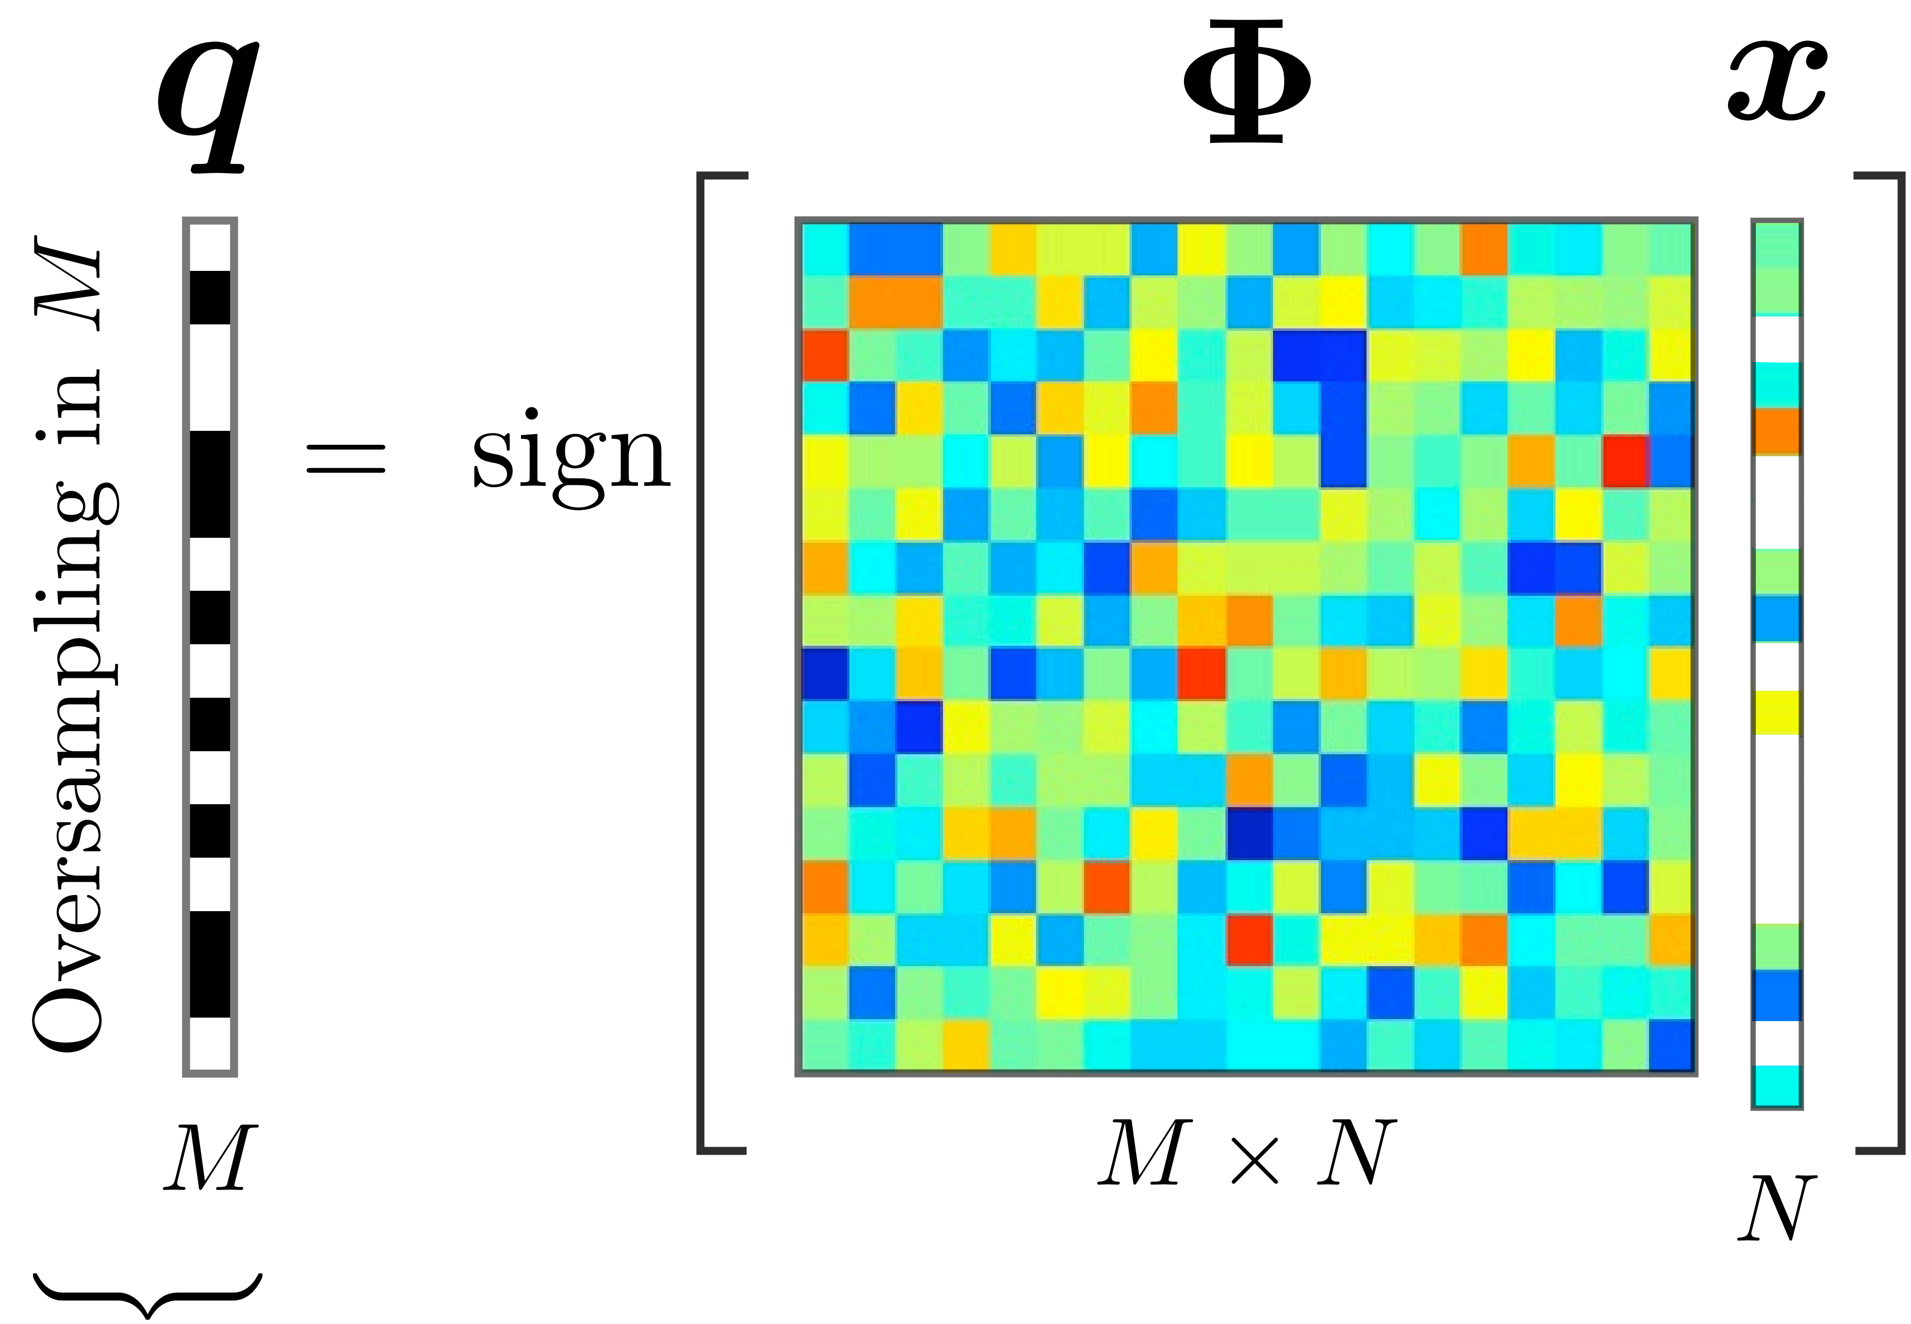

M-bits! But, which information inside $\boldsymbol{q}$?

1-bit Computational ~~Compressed~~ Sensing

bits matter!

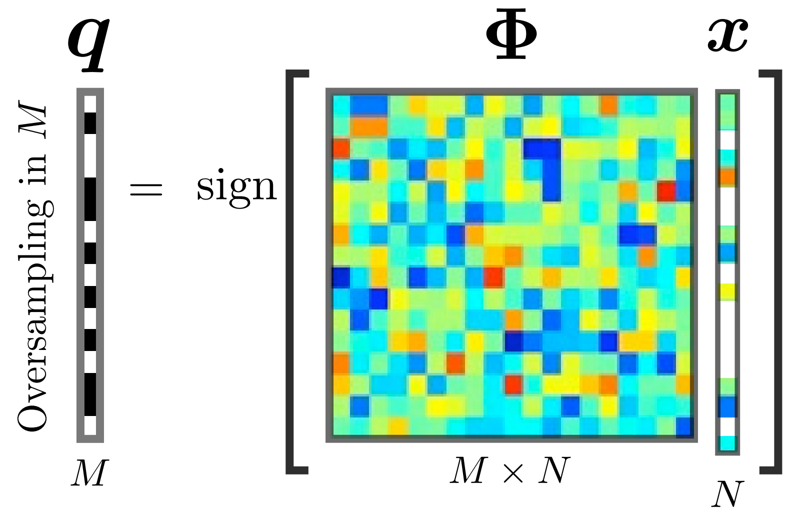

Warning 1: signal amplitude is lost!

$$\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi}(\lambda \boldsymbol{x})) = \mathrm{sign}\,(\boldsymbol{\Phi x}), \quad \forall \lambda > 0$$

$\Rightarrow$ Amplitude is arbitrarily fixed

Examples : $\|\boldsymbol{x}\| = 1$ or $\|\boldsymbol{\Phi x}\|_1 = 1$

1-bit ~~Compressed~~ Computational Sensing

bits matter!

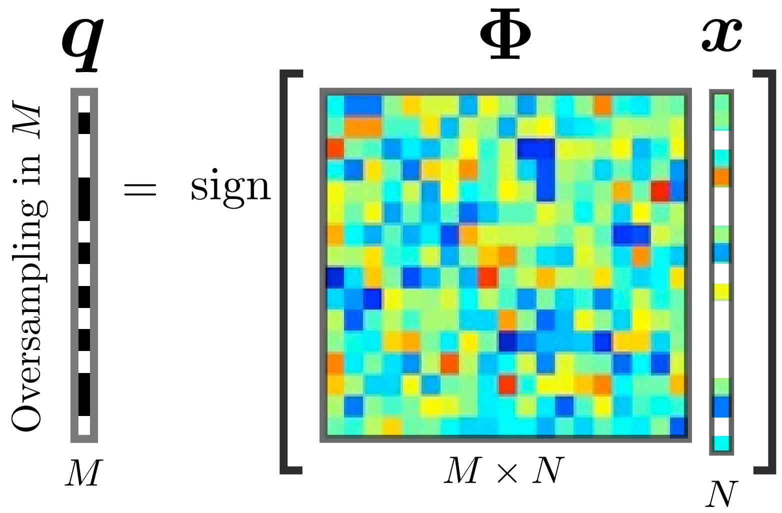

[Plan, Vershynin, 11]

Warning 2: ∃ forbidden sensing!

Let $\boldsymbol{x}_\lambda := (1, \lambda, 0, \cdots, 0)^T \in \mathbb{R}^N$

and $\boldsymbol{\Phi} \in \{\pm 1\}^{M \times N}$ (*e.g.*, Bernoulli).

We have $\|\boldsymbol{x}_0 - \boldsymbol{x}_\lambda\| = \lambda$

but $\boldsymbol{q} = \operatorname{sign}(\boldsymbol{\Phi}\boldsymbol{x}_0) = \operatorname{sign}(\boldsymbol{\Phi}\boldsymbol{x}_\lambda), \ \forall |\lambda| < 1$

$\Rightarrow$ No hope to distinguish them by increasing M!

2. Theoretical performance limits

Lower bound: cell intersection viewpoint

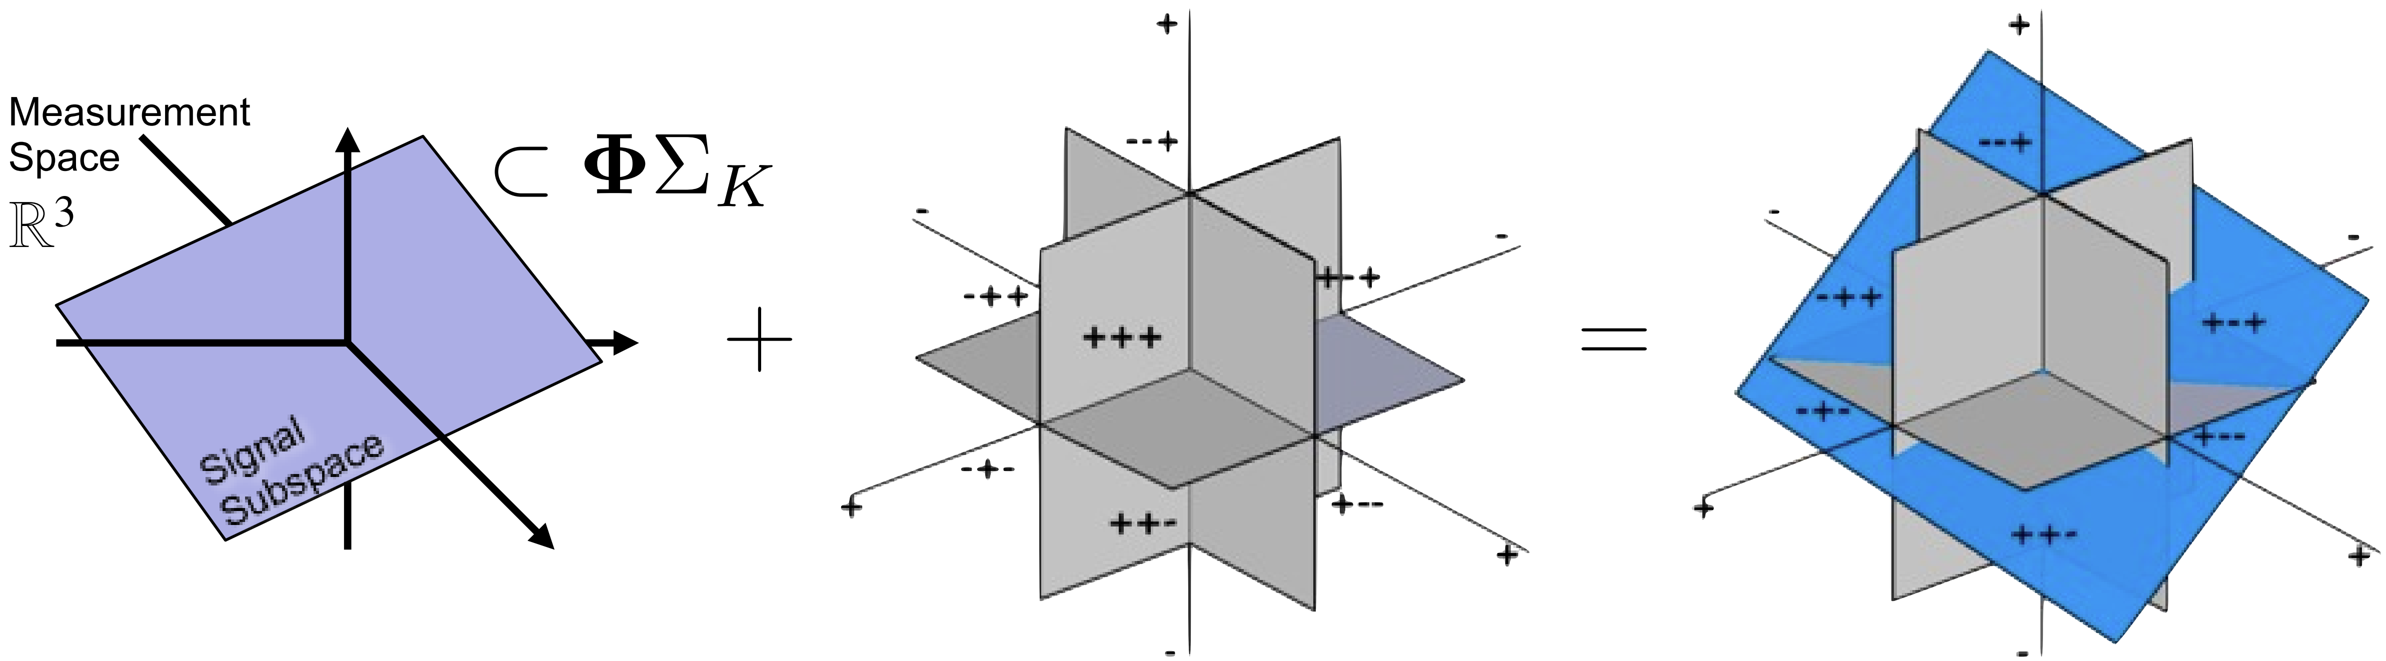

Not all quantization cells intersected!

no more than $C = 2^K \binom{N}{K} \binom{M}{K}$

Lower bound: cell intersection viewpoint

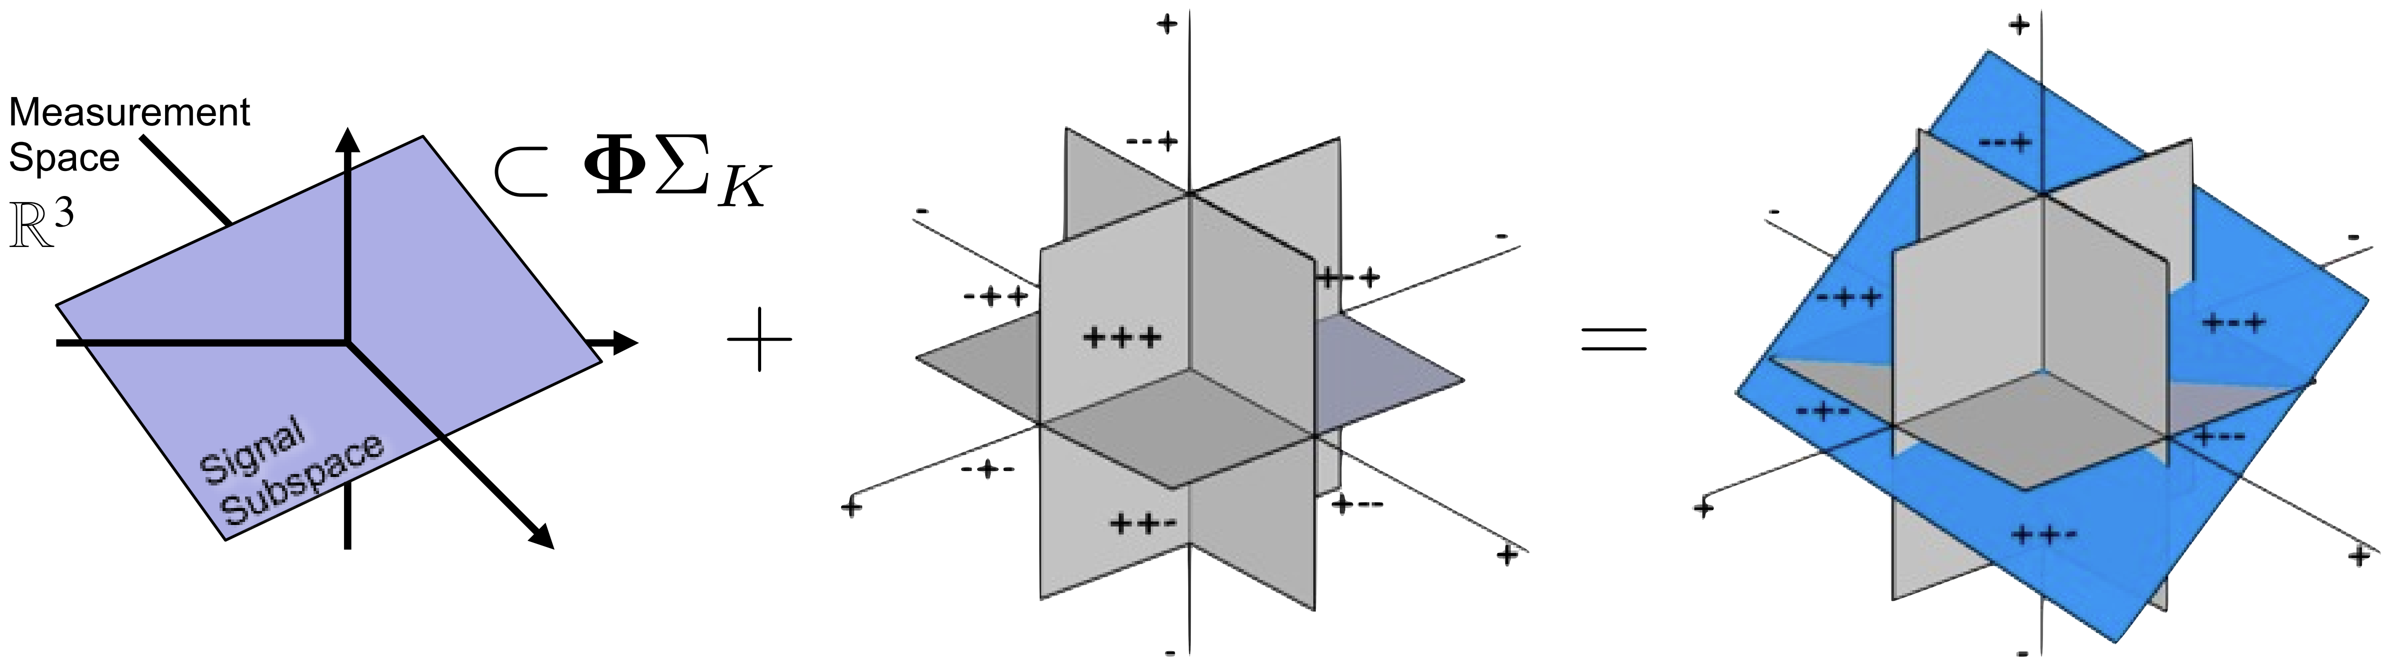

Not all quantization cells intersected!

no more than $C = 2^K \binom{N}{K}\binom{M}{K}$

Most efficient ϵ-covering of $S^{N-1} \cap \Sigma_K$ with ϵ-caps

$\Rightarrow$ lower bound on C by "vol$(S^{N-1} \cap \Sigma_K)$/vol(ϵ-cap)"

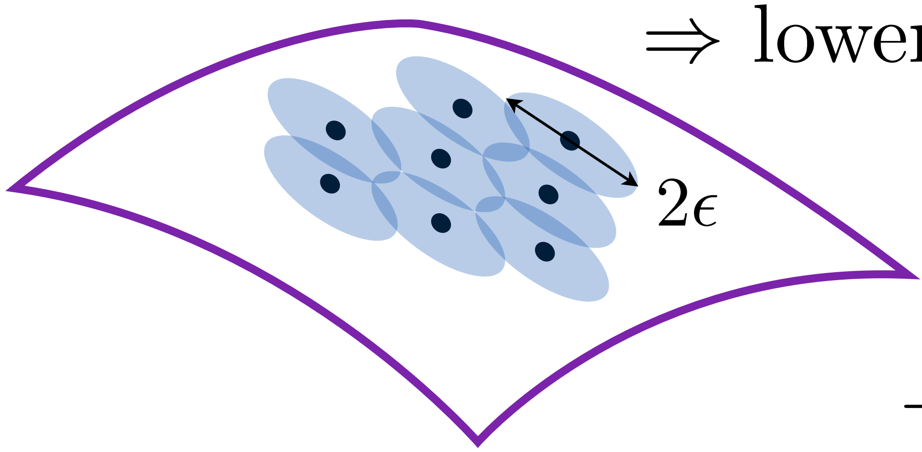

$$\Rightarrow \epsilon = \Omega(K/M)$$

$\rightarrow$ Lower bound on any 1-bit reconstruction error

Reaching this bound ?

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

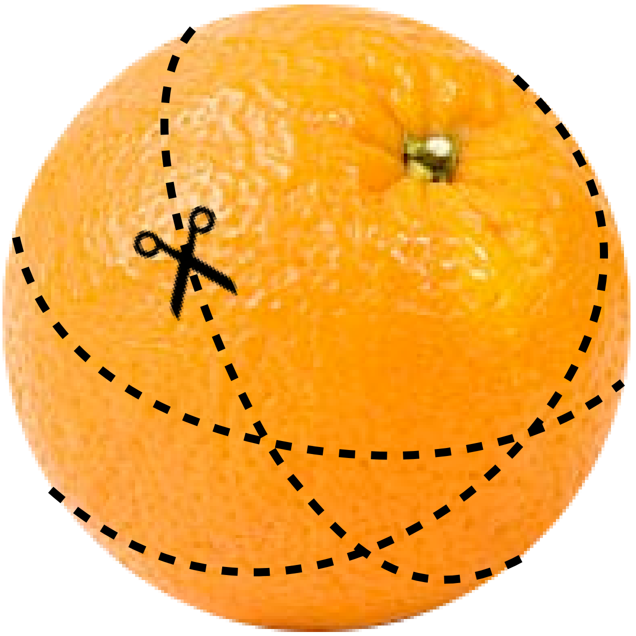

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

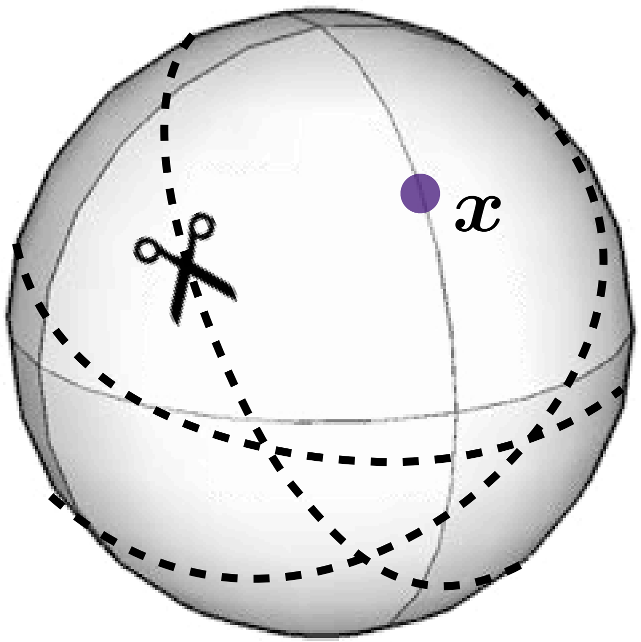

Reaching this bound ?

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$

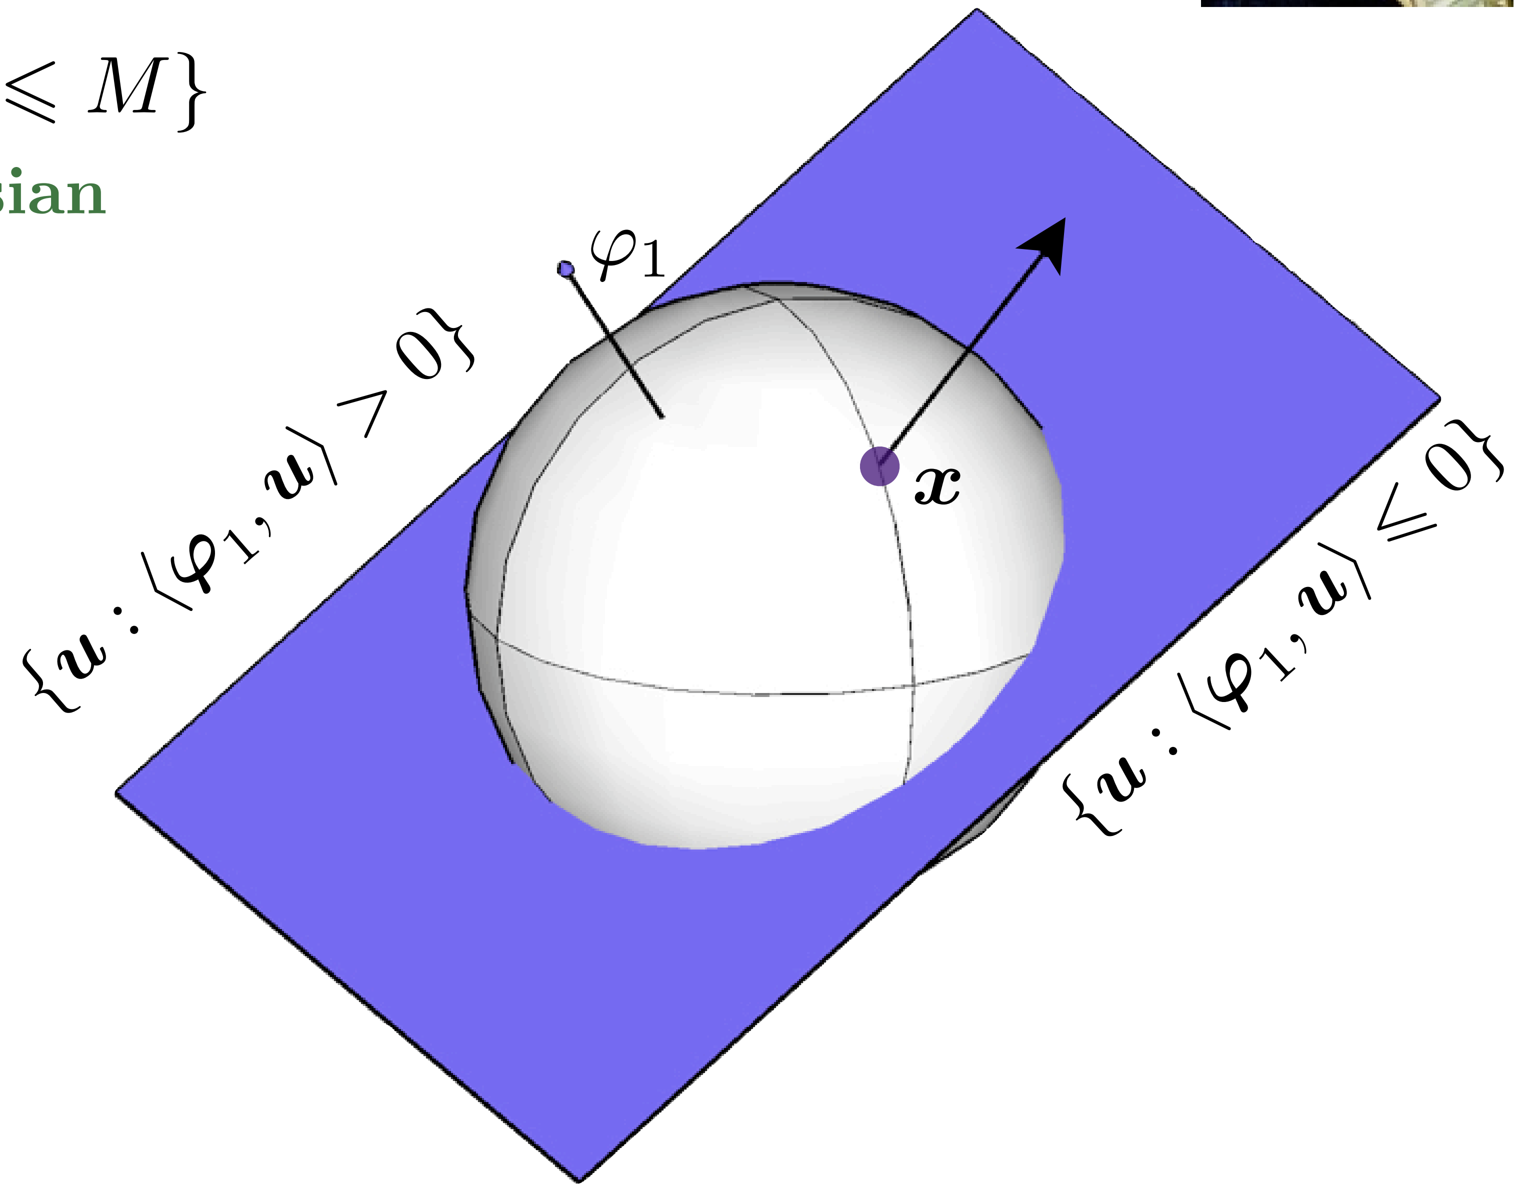

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

$\boldsymbol{x}$ on S^2
M vectors:
$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$
iid Gaussian

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$

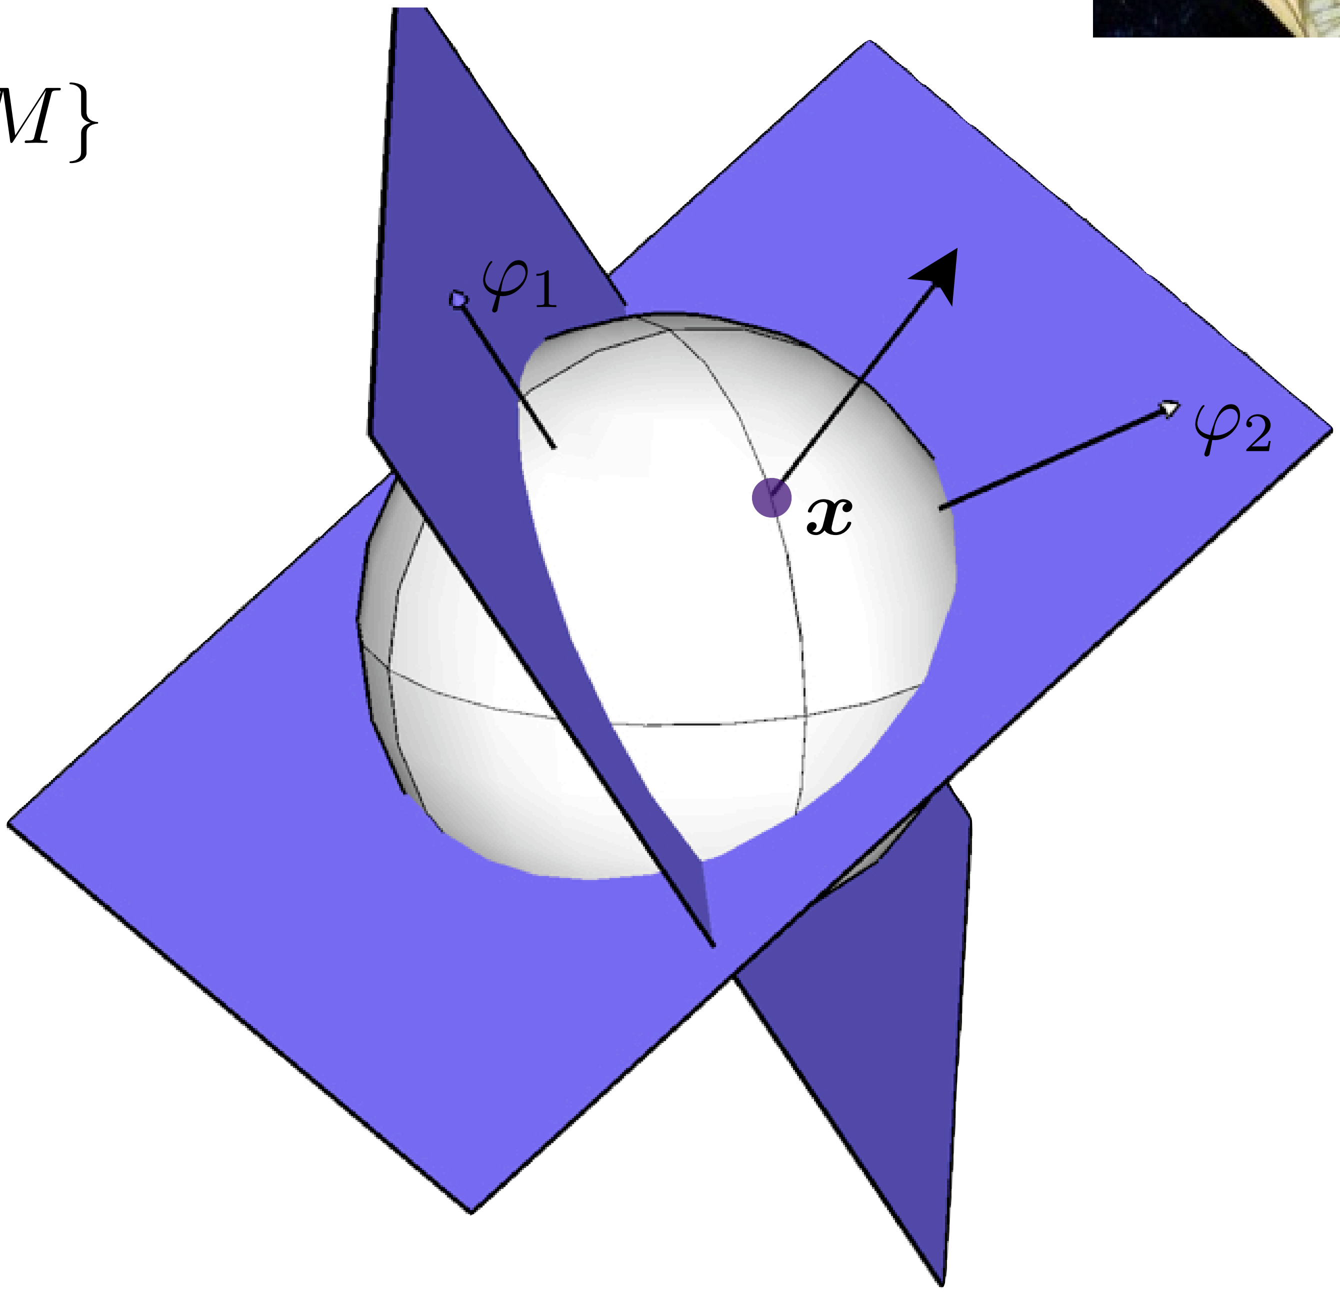

Reaching this bound ?

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_3, \boldsymbol{x} \rangle \leqslant 0$

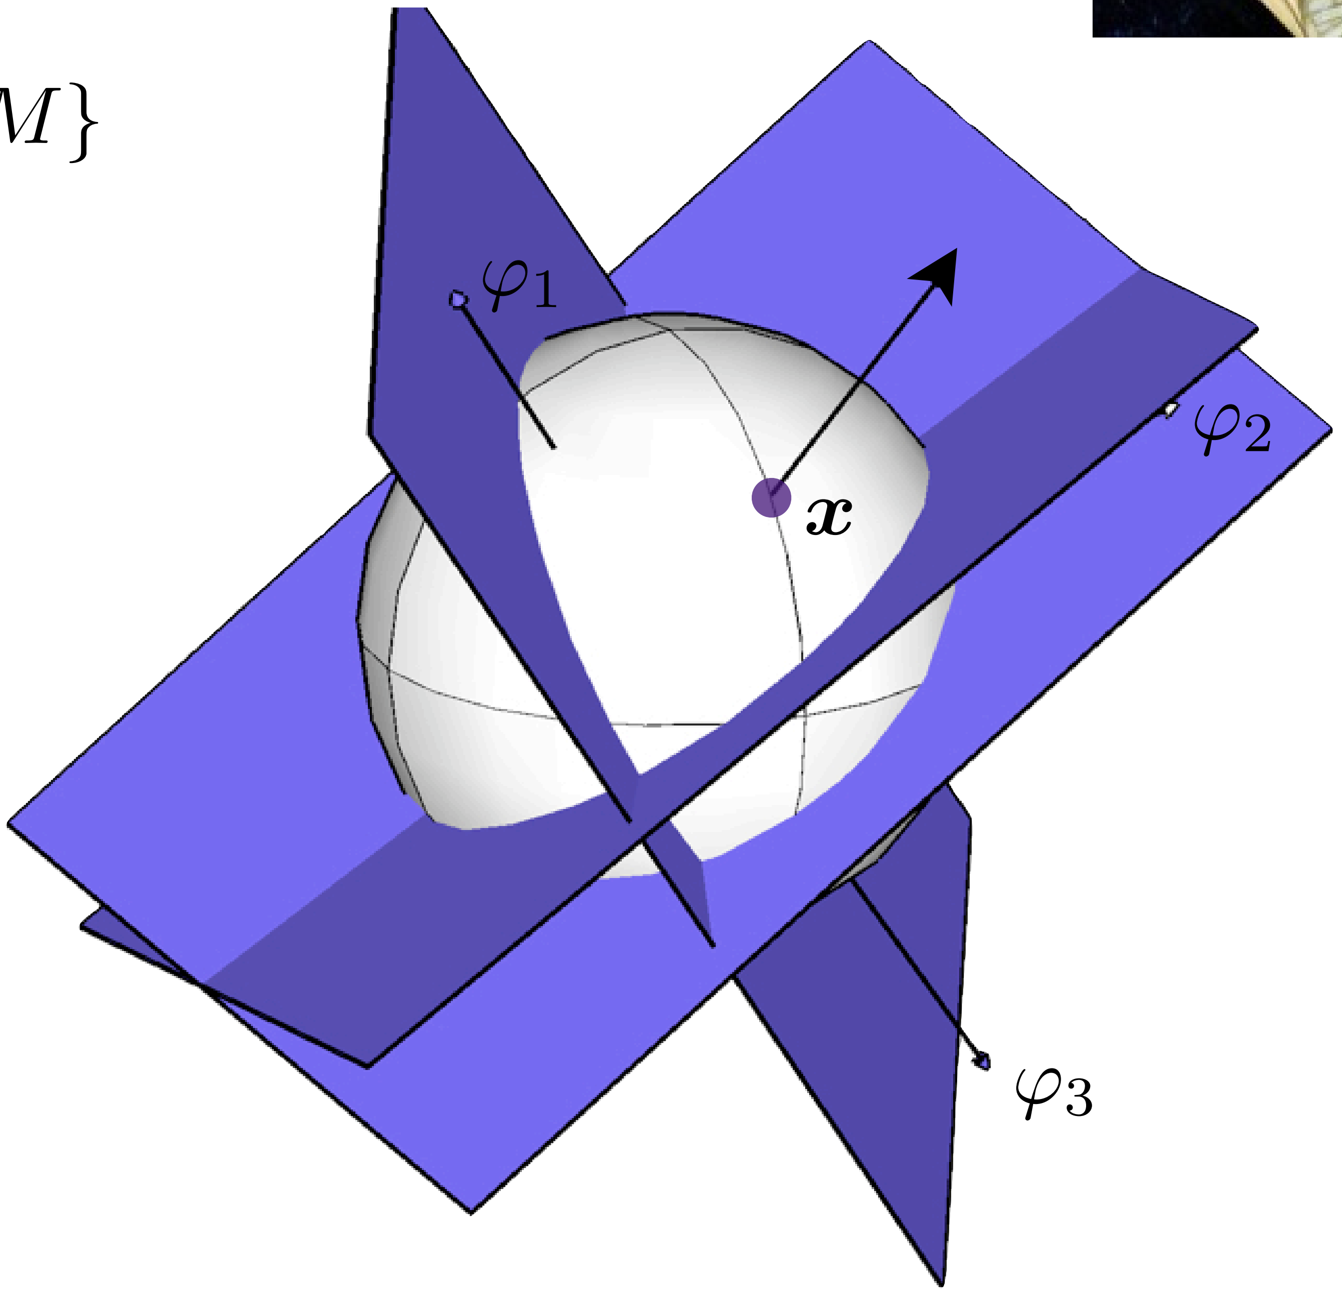

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

$\boldsymbol{x}$ on S^2
M vectors:
$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$
iid Gaussian

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_3, \boldsymbol{x} \rangle \leqslant 0$
$\langle \boldsymbol{\varphi}_4, \boldsymbol{x} \rangle > 0$

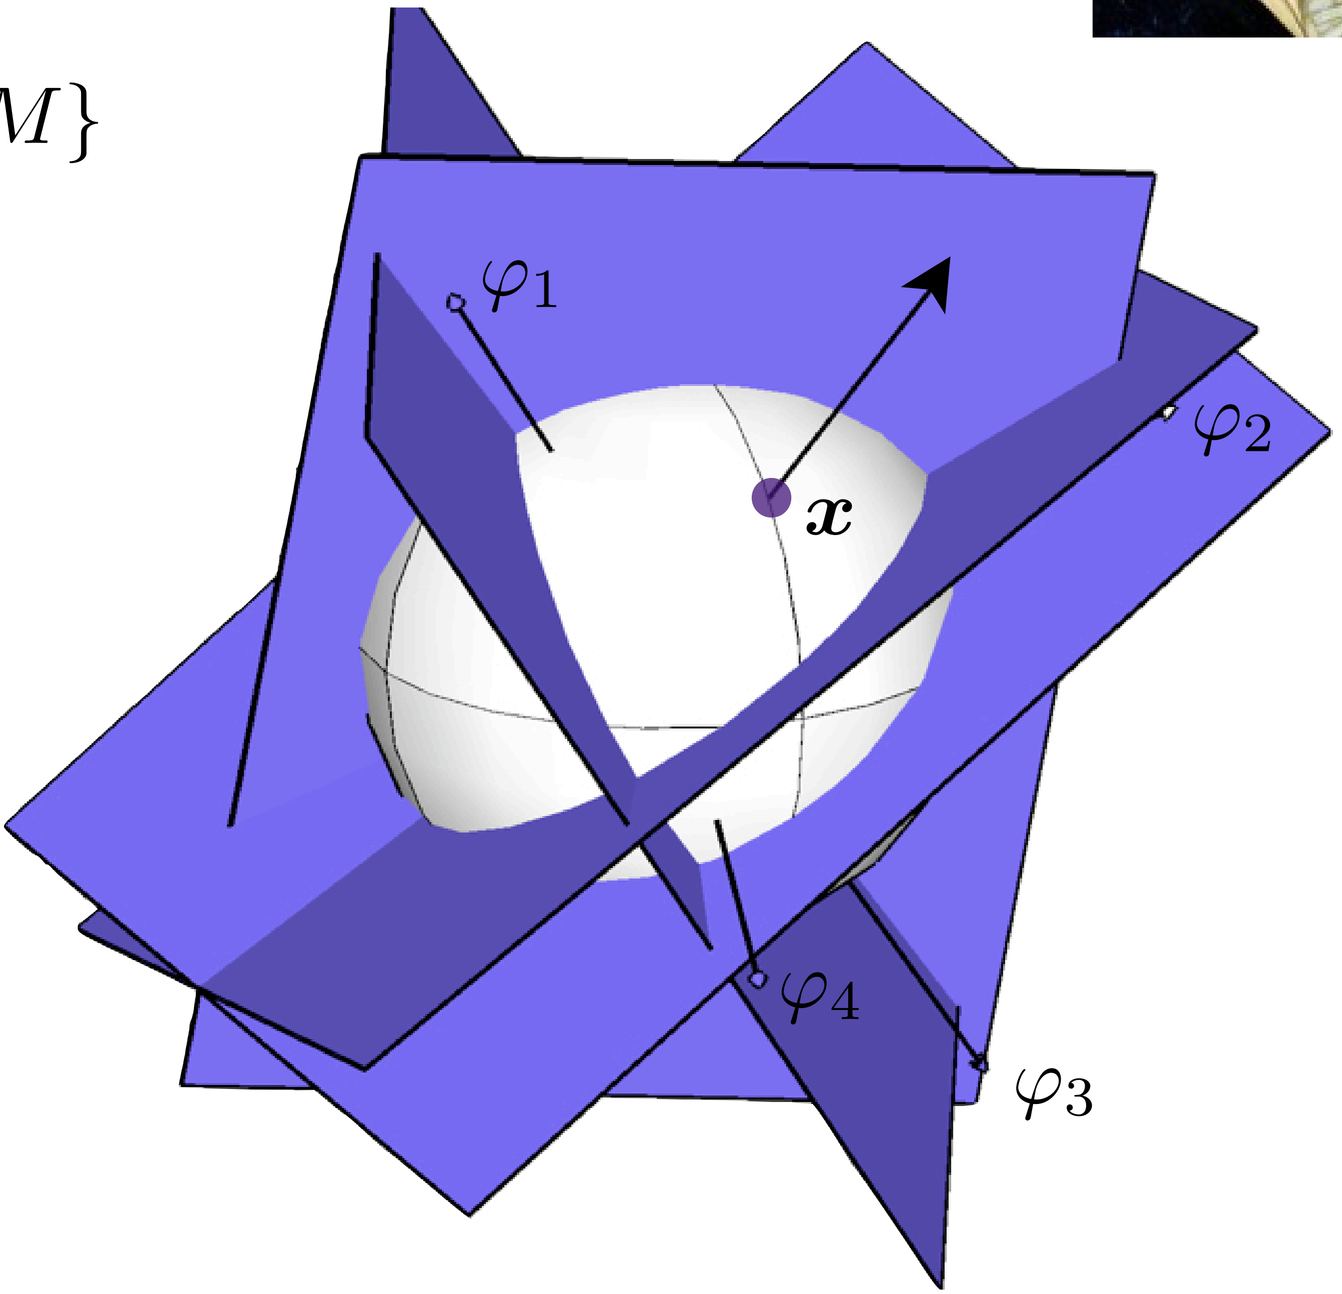

Reaching this bound ?

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_3, \boldsymbol{x} \rangle \leqslant 0$
$\langle \boldsymbol{\varphi}_4, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_5, \boldsymbol{x} \rangle > 0$

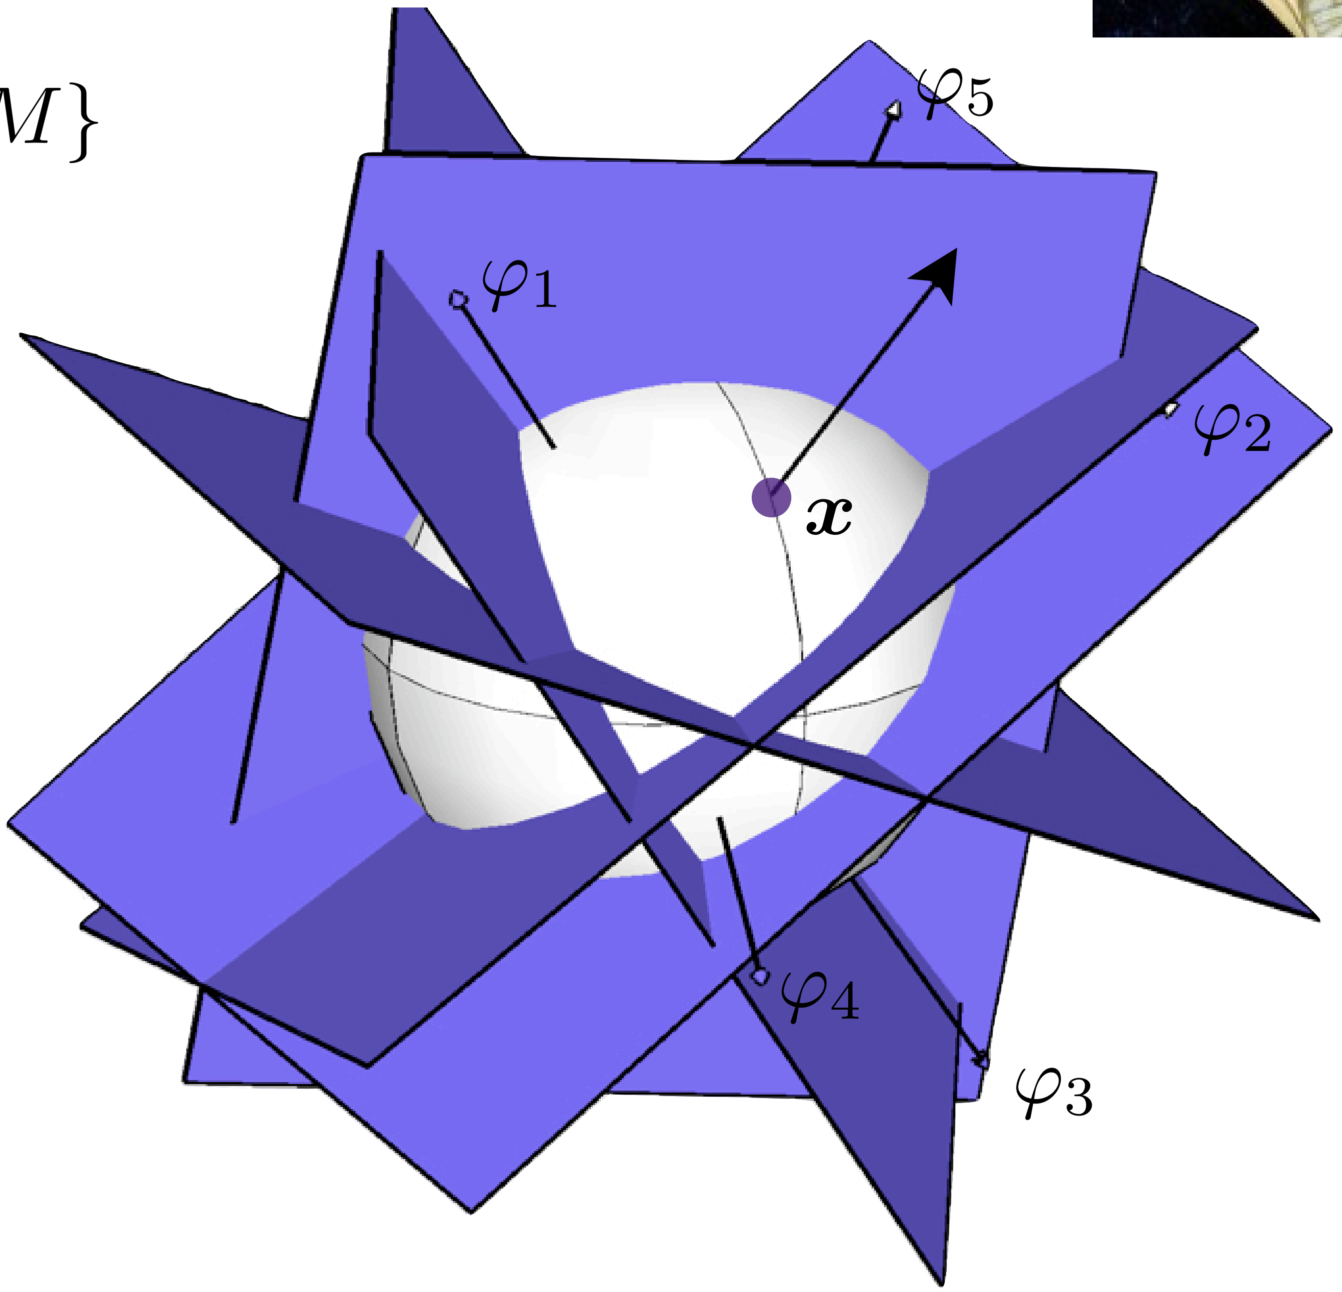

Reaching this bound ?

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

<u>Carl Friedrich Gauss:</u>
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_3, \boldsymbol{x} \rangle \leqslant 0$
$\langle \boldsymbol{\varphi}_4, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_5, \boldsymbol{x} \rangle > 0$
$\vdots$

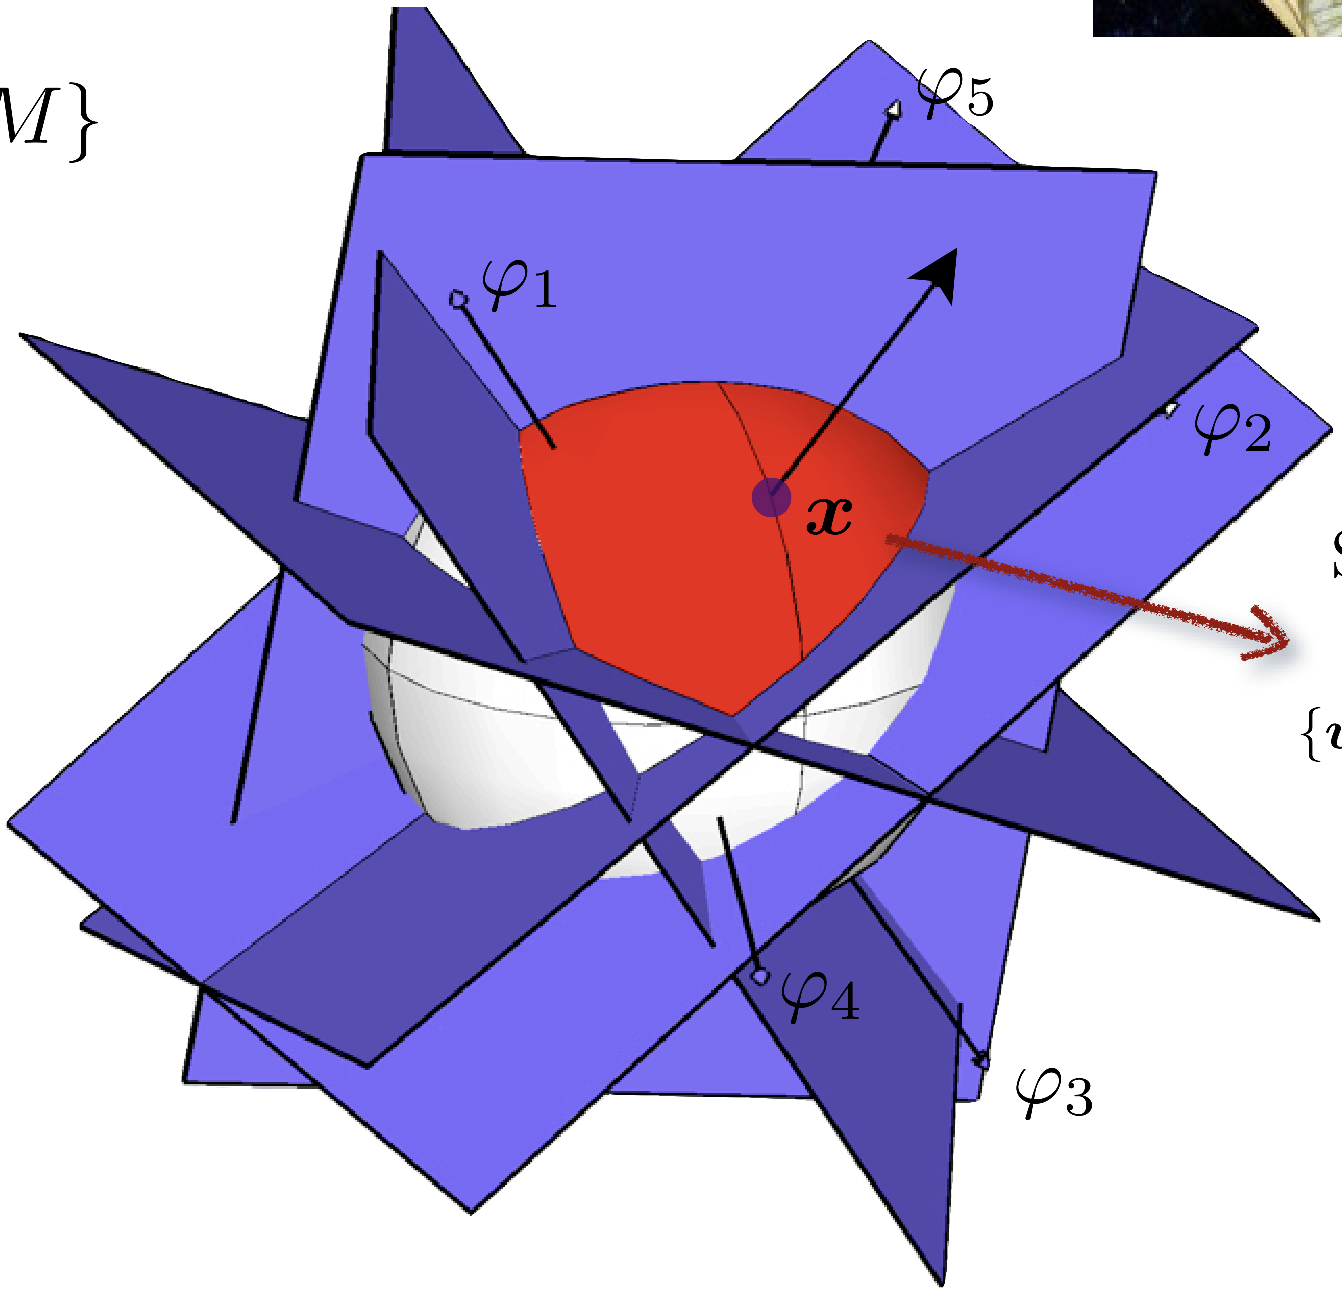

Smaller and smaller when M increases
$\{\boldsymbol{u} : \mathrm{sign}\,(\boldsymbol{\Phi u}) = \mathrm{sign}\,(\boldsymbol{\Phi x})\}$

Reaching this bound ?

$\boldsymbol{x}$ on S^2

M vectors:

$\{\boldsymbol{\varphi}_i : 1 \leqslant i \leqslant M\}$

iid Gaussian

1-bit Measurements

$\langle \boldsymbol{\varphi}_1, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_2, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_3, \boldsymbol{x} \rangle \leqslant 0$
$\langle \boldsymbol{\varphi}_4, \boldsymbol{x} \rangle > 0$
$\langle \boldsymbol{\varphi}_5, \boldsymbol{x} \rangle > 0$
$\vdots$

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

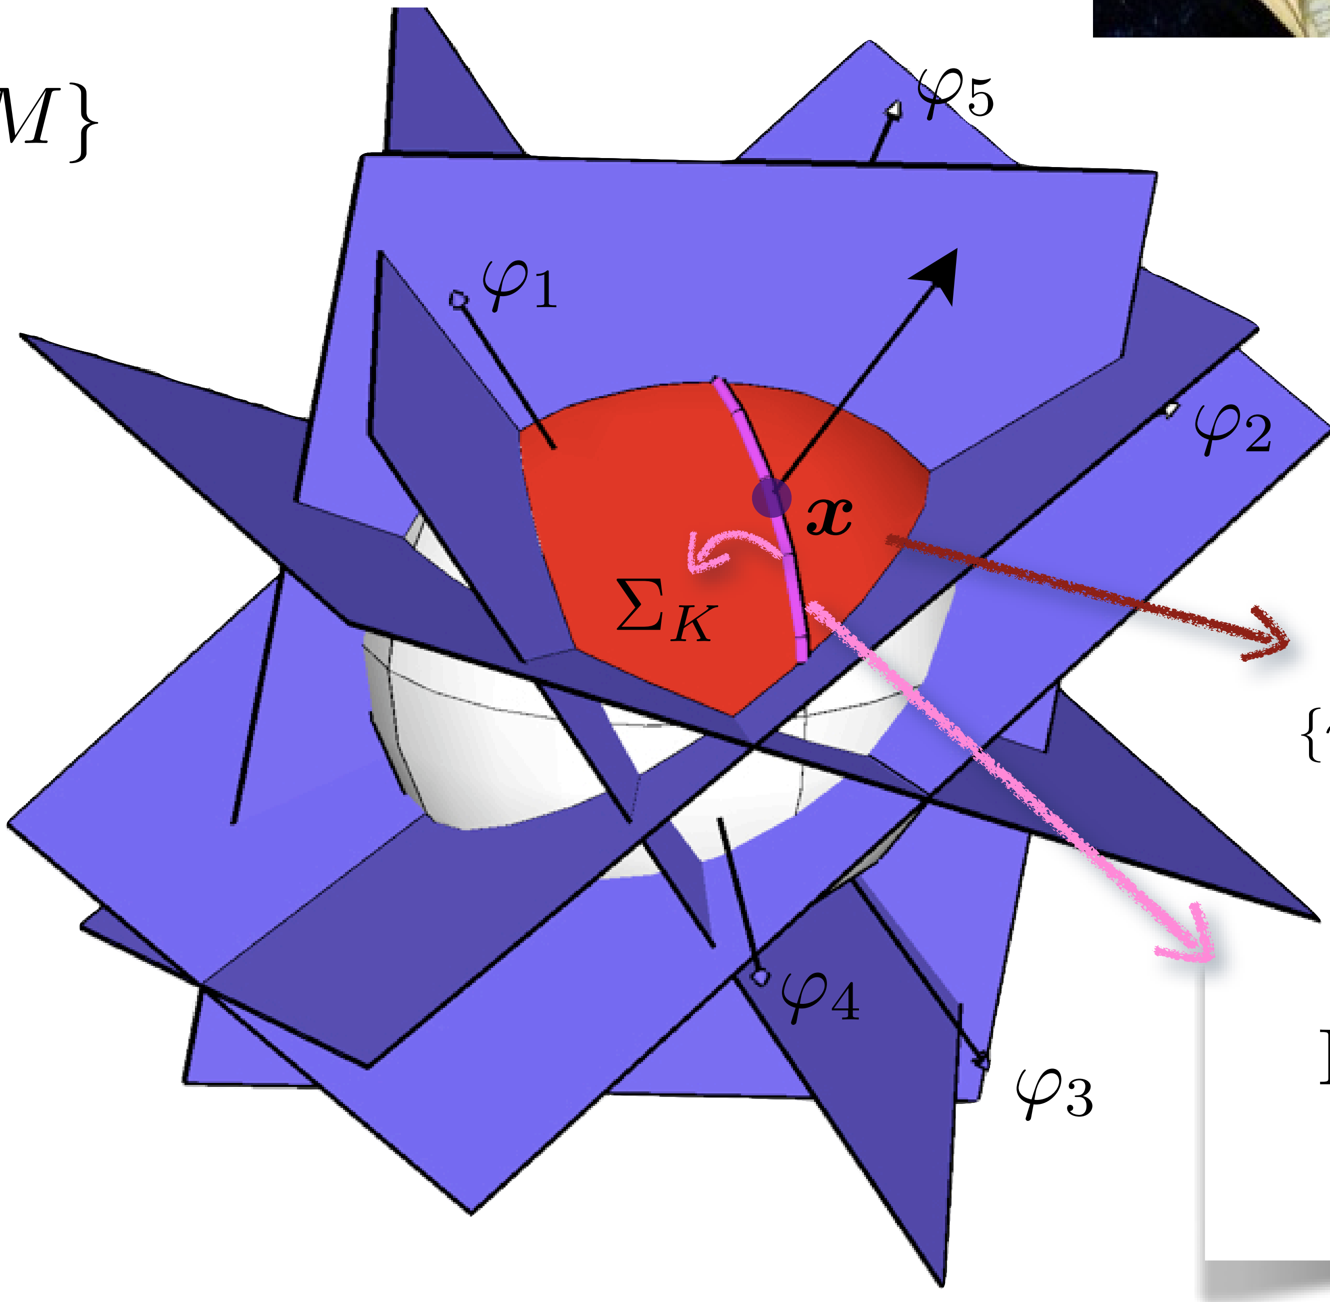

Smaller and smaller when M increases

$\{\boldsymbol{u} : \mathrm{sign}(\boldsymbol{\Phi u}) = \mathrm{sign}(\boldsymbol{\Phi x})\}$

Lower bound on this width?

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

Let $A(\cdot) := \mathrm{sign}\,(\boldsymbol{\Phi}\,\cdot)$ with $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$.

If $M = O(\epsilon^{-1} K \log N)$, then, w.h.p,

for any two unit K-sparse vectors $\boldsymbol{x}$ and $\boldsymbol{s}$,

$$A(\boldsymbol{x}) = A(\boldsymbol{s}) \quad \Rightarrow \quad \|\boldsymbol{x} - \boldsymbol{s}\| \leq \epsilon$$

$$\Leftrightarrow \epsilon = O\big(\tfrac{K}{M} \log \tfrac{MN}{K}\big)$$

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

Let $A(\cdot) := \operatorname{sign}(\boldsymbol{\Phi}\,\cdot)$ with $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$.

If $M = O(\epsilon^{-1} K \log N)$, then, w.h.p,

for any two unit K-sparse vectors $\boldsymbol{x}$ and $\boldsymbol{s}$,

$$A(\boldsymbol{x}) = A(\boldsymbol{s}) \quad \Rightarrow \quad \|\boldsymbol{x} - \boldsymbol{s}\| \leq \epsilon$$

$$\Leftrightarrow \epsilon = O\big(\tfrac{K}{M} \log \tfrac{MN}{K}\big)$$

almost optimal

Reaching this bound ?

Carl Friedrich Gauss:
"1-bit CS? I solved it at breakfast by randomly slicing my orange!"
http://www.gaussfacts.com

Let $A(\cdot) := \mathrm{sign}\,(\boldsymbol{\Phi}\,\cdot)$ with $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$.

If $M = O(\epsilon^{-1} K \log N)$, then, w.h.p,

for any two unit K-sparse vectors $\boldsymbol{x}$ and $\boldsymbol{s}$,

$$A(\boldsymbol{x}) = A(\boldsymbol{s}) \quad \Rightarrow \quad \|\boldsymbol{x} - \boldsymbol{s}\| \leq \epsilon$$

$$\Leftrightarrow \epsilon = O\big(\tfrac{K}{M} \log \tfrac{MN}{K}\big)$$

almost optimal

Note: You can even afford a small error, *i.e.*,

if only b bits are different between $A(\boldsymbol{x})$ and $A(\boldsymbol{s})$ $\Rightarrow \|\boldsymbol{x} - \boldsymbol{s}\| \leqslant \frac{K+b}{K}\,\epsilon$

3. Stable embeddings: angles are preserved

What's known?

- Let's define

$$A(\boldsymbol{u}) := \mathrm{sign}\,(\boldsymbol{\Phi u}) \;\Leftrightarrow\; A_j(\boldsymbol{u}) = \mathrm{sign}\,(\boldsymbol{\varphi}_j \cdot \boldsymbol{u}) \in \{\pm 1\}$$

$j^{\rm th}$ row of $\boldsymbol{\Phi}$

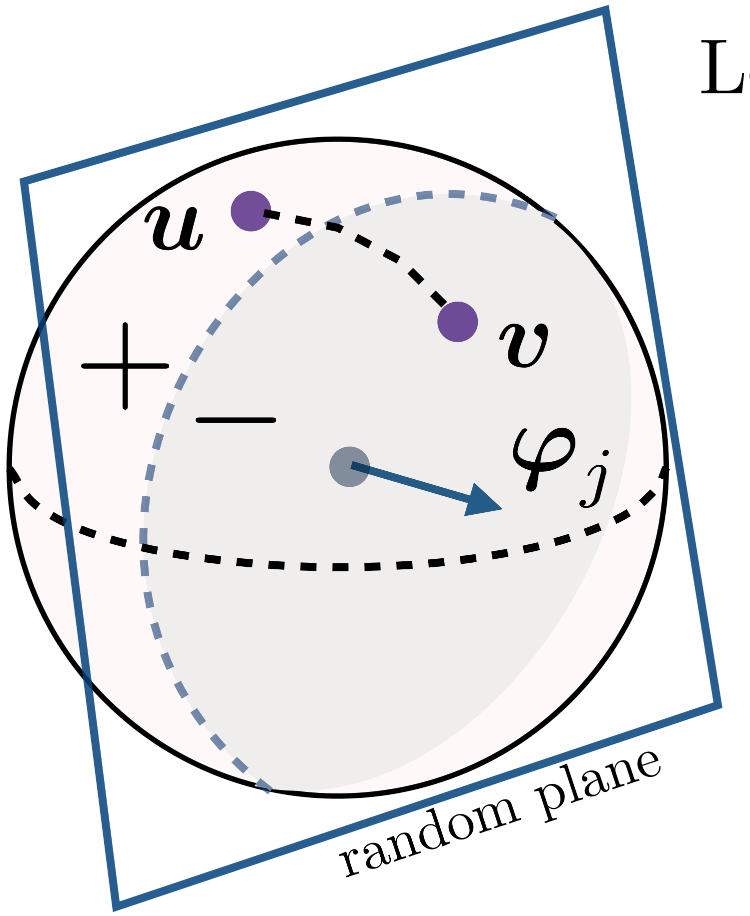

Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{S}^{N-1}$ (wlog)

$$\mathbb{P}[A_j(\boldsymbol{u}) \neq A_j(\boldsymbol{v})] \;=\; ?$$

What's known?

- Let's define

$$A(\boldsymbol{u}) := \operatorname{sign}(\boldsymbol{\Phi}\boldsymbol{u}) \;\Leftrightarrow\; A_j(\boldsymbol{u}) = \operatorname{sign}(\boldsymbol{\varphi}_j \cdot \boldsymbol{u}) \in \{\pm 1\}$$

j^{th} row of $\boldsymbol{\Phi}$

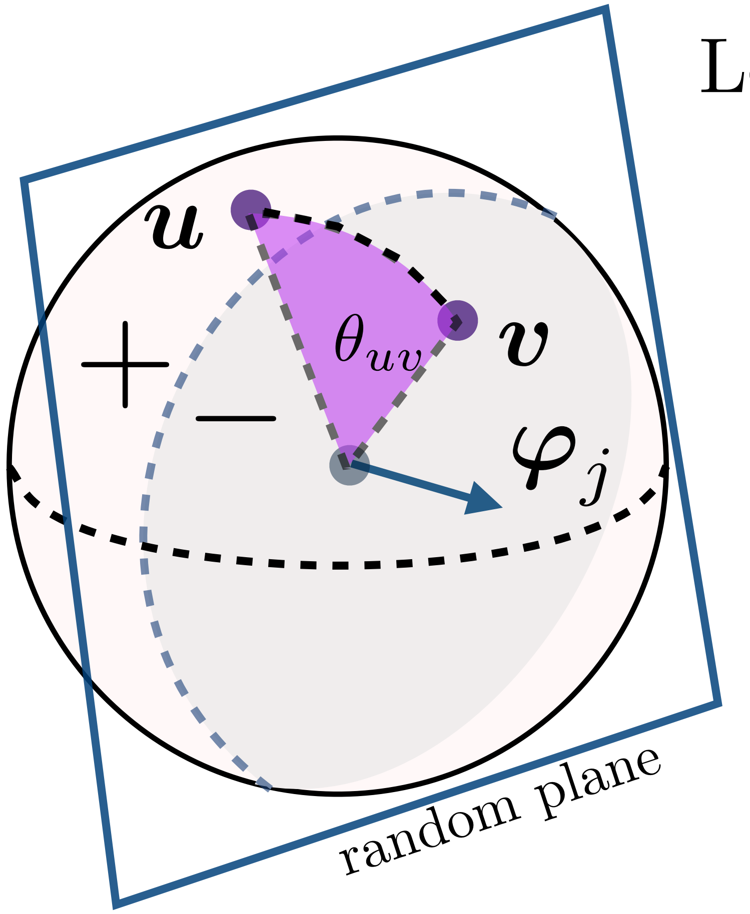

Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{S}^{N-1}$ (wlog)

$$\mathbb{P}[A_j(\boldsymbol{u}) \neq A_j(\boldsymbol{v})] = \frac{1}{\pi}\,\text{angle}(\boldsymbol{u}, \boldsymbol{v})$$

$$= \frac{1}{\pi}\theta_{uv}$$

What's known?

- Let's define

$$A(\boldsymbol{u}) := \mathrm{sign}\,(\boldsymbol{\Phi}\boldsymbol{u}) \;\Leftrightarrow\; A_j(\boldsymbol{u}) = \mathrm{sign}\,(\boldsymbol{\varphi}_j \cdot \boldsymbol{u}) \in \{\pm 1\}$$

j^{th} row of $\boldsymbol{\Phi}$

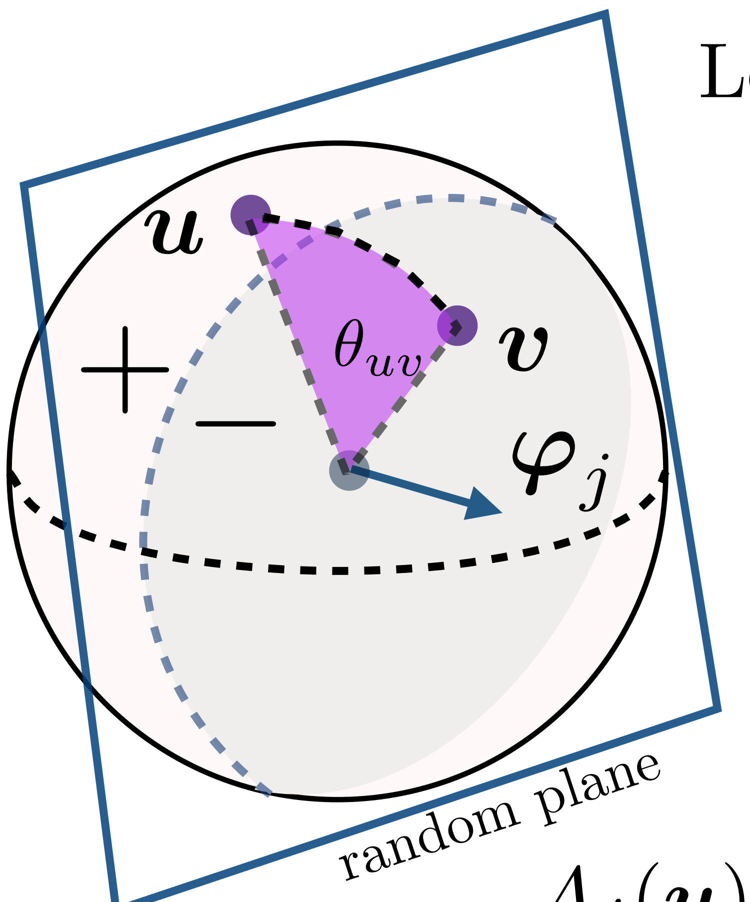

Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{S}^{N-1}$ (wlog)

$$\mathbb{P}[A_j(\boldsymbol{u}) \neq A_j(\boldsymbol{v})] = \tfrac{1}{\pi}\,\mathrm{angle}(\boldsymbol{u}, \boldsymbol{v})$$

$$= \tfrac{1}{\pi}\theta_{uv}$$

$$\Rightarrow \quad X_j = \tfrac{1}{2}\big|\overbrace{A_j(\boldsymbol{u}) - A_j(\boldsymbol{v})}^{A_j(\boldsymbol{u}) \oplus A_j(\boldsymbol{v})\ \ (\mathrm{XOR})}\big| \underset{\mathrm{iid}}{\sim} \mathrm{Bernoulli}(\tfrac{\theta_{uv}}{\pi}) \in \{0, 1\}$$

Starting point: Hamming/Angle Concentration

- Metrics of interest:

$$d_H(\boldsymbol{u}, \boldsymbol{v}) = \tfrac{1}{M} \textstyle\sum_i (u_i \oplus v_i) \quad \text{(norm. Hamming)}$$

$$d_{\text{ang}}(\boldsymbol{x}, \boldsymbol{s}) = \tfrac{1}{\pi} \arccos(\langle \boldsymbol{x}, \boldsymbol{s} \rangle) \quad \text{(norm. angle)}$$

Starting point: Hamming/Angle Concentration

- Metrics of interest:

$$d_H(\boldsymbol{u},\boldsymbol{v}) = \tfrac{1}{M}\textstyle\sum_i (u_i \oplus v_i) \quad \text{(norm. Hamming)}$$
$$d_{\rm ang}(\boldsymbol{x},\boldsymbol{s}) = \tfrac{1}{\pi}\arccos(\langle \boldsymbol{x},\boldsymbol{s}\rangle) \quad \text{(norm. angle)}$$

- Known fact: if $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$ [*e.g.*, Goemans, Williamson 1995]

Let $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, $A(\cdot) = \mathrm{sign}\,(\boldsymbol{\Phi}\ \cdot) \in \{-1,1\}^M$ and $\epsilon > 0$.
For any $\boldsymbol{x},\boldsymbol{s} \in S^{N-1}$, we have

$$\mathbb{P}_{\boldsymbol{\Phi}}\Big[\,\big|\underbrace{d_H\big(A(\boldsymbol{x}),A(\boldsymbol{s})\big)} - d_{\rm ang}(\boldsymbol{x},\boldsymbol{s})\big| \leqslant \epsilon\,\Big] \geqslant 1 - 2\,e^{-2\epsilon^2 M}.$$

$$\tfrac{1}{M}\textstyle\sum_{i=1}^{M} X_i = \tfrac{1}{M}\sum_i A_i(\boldsymbol{x}) \oplus A_i(\boldsymbol{s})$$

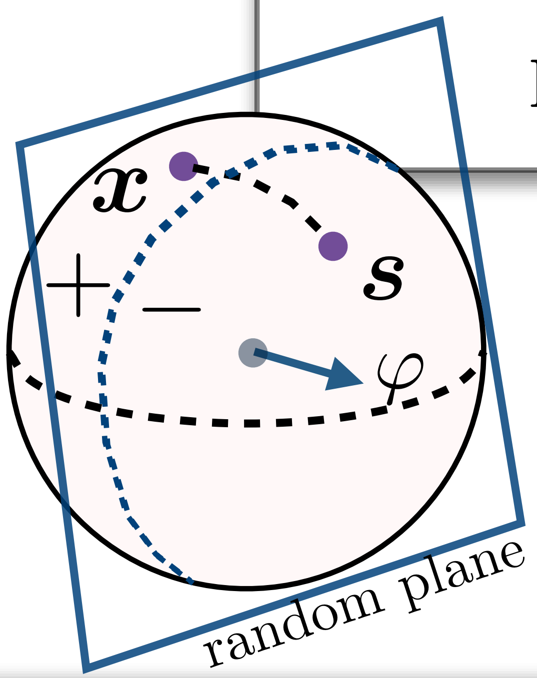

$\Rightarrow$ Thanks to $A(.)$, Hamming distance concentrates around vector angles!

Binary ϵ Stable Embedding (BϵSE)

A mapping $A : \mathbb{R}^N \to \{\pm 1\}^M$ is a **binary ϵ-stable embedding** (BϵSE) of order K for sparse vectors if

$$d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon \leq d_H(A(\boldsymbol{x}), A(\boldsymbol{s})) \leq d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) + \epsilon$$

for all $\boldsymbol{x}, \boldsymbol{s} \in S^{N-1}$ with $\boldsymbol{x} \pm \boldsymbol{s}$ K-sparse.

kind of "binary restricted (quasi) isometry"

Binary ϵ Stable Embedding (BϵSE)

A mapping $A : \mathbb{R}^N \to \{\pm 1\}^M$ is a **binary ϵ-stable embedding** (BϵSE) of order K for sparse vectors if

$$d_{\text{ang}}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon \leq d_H(A(\boldsymbol{x}), A(\boldsymbol{s})) \leq d_{\text{ang}}(\boldsymbol{x}, \boldsymbol{s}) + \epsilon$$

for all $\boldsymbol{x}, \boldsymbol{s} \in S^{N-1}$ with $\boldsymbol{x} \pm \boldsymbol{s}$ K-sparse.

kind of "binary restricted (quasi) isometry"

- *Corollary*: for any algorithm with output $\boldsymbol{x}^*$ jointly K-sparse and consistent (*i.e.*, $A(\boldsymbol{x}^*) = A(\boldsymbol{x})$),

$$d_{\text{ang}}(\boldsymbol{x}, \boldsymbol{x}^*) \leqslant 2\epsilon!$$

- If limited binary noise, d_{ang} still bounded
- If not exactly sparse signals (but almost), d_{ang} still bounded

BϵSE existence? Yes!

Let $\mathbf{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, fix $0 \leqslant \eta \leqslant 1$ and $\epsilon > 0$. If

$$M \;\geqslant\; \tfrac{4}{\epsilon^2}\big(K\,\log(N) + 2K\,\log(\tfrac{50}{\epsilon}) + \log(\tfrac{2}{\eta})\big),$$

then $\mathbf{\Phi}$ is a BϵSE with $\Pr > 1-\eta$.

$$M = O(\epsilon^{-2} K \log N)$$

BϵSE existence? Yes!

Let $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, fix $0 \leqslant \eta \leqslant 1$ and $\epsilon > 0$. If

$$M \;\geqslant\; \tfrac{4}{\epsilon^2}\big(K\,\log(N) + 2K\,\log(\tfrac{50}{\epsilon}) + \log(\tfrac{2}{\eta})\big),$$

then $\boldsymbol{\Phi}$ is a BϵSE with $\Pr > 1-\eta$.

$$M = O(\epsilon^{-2} K \log N)$$

Proof sketch:

1) Generalize

$$\mathbb{P}_{\boldsymbol{\Phi}}\big[\,\big|\,d_H\big(A(\boldsymbol{x}), A(\boldsymbol{s})\big) \;-\; d_{\rm ang}(\boldsymbol{x},\boldsymbol{s})\big| \;\leqslant\; \epsilon\,\big] \;\geqslant\; 1 - 2\,e^{-2\epsilon^2 M}.$$

to

$$\mathbb{P}_{\boldsymbol{\Phi}}\Big[\,\big|\,d_H\big(A(\boldsymbol{u}), A(\boldsymbol{v})\big) \;-\; d_{\rm ang}(\boldsymbol{x},\boldsymbol{s})\big| \;\leqslant\; \epsilon + (\tfrac{\pi}{2}D)^{1/2}\,\delta\,\Big] \;\geqslant\; 1 - 2\,e^{-2\epsilon^2 M}.$$

for $\boldsymbol{u}, \boldsymbol{v}$ in a D-dimensional neighborhood of width δ around $\boldsymbol{x}$ and $\boldsymbol{s}$ resp.

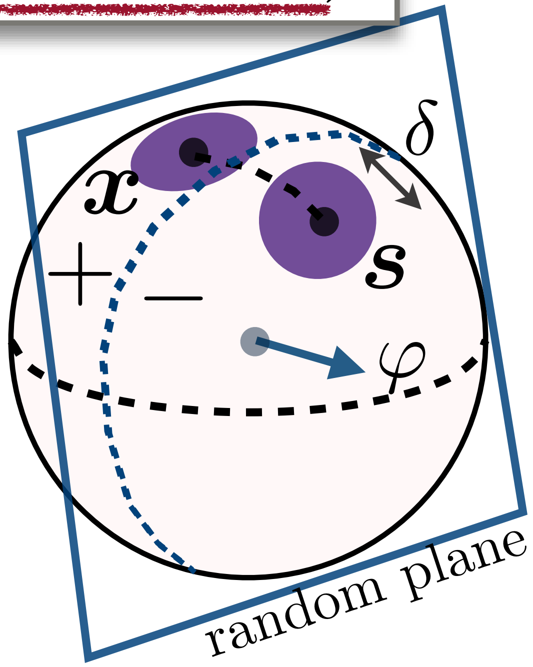

2) Covers the space of "K-sparse signal pairs" in $\mathbb{R}^N$ by

$$O\big(\tbinom{N}{K}\delta^{-2K}\big) = O\big((\tfrac{eN}{K\delta^2})^K\big) \text{ neighborhoods.}$$

3) Apply Point 1 with union bound, and "stir until the proof thickens"

BϵSE existence? Yes!

Let $\mathbf{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, fix $0 \leqslant \eta \leqslant 1$ and $\epsilon > 0$. If

$$M \;\geqslant\; \tfrac{4}{\epsilon^2}\big(K\,\log(N) + 2K\,\log(\tfrac{50}{\epsilon}) + \log(\tfrac{2}{\eta})\big),$$

then $\mathbf{\Phi}$ is a BϵSE with $\Pr > 1-\eta$.

$$M = O(\epsilon^{-2} K \log N)$$

$\Rightarrow$ BϵSE consistency "width":

$$\epsilon \;=\; O\big((\tfrac{K}{M}\log\tfrac{MN}{K})^{1/2}\big)$$

BϵSE existence? Yes!

Let $\mathbf{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, fix $0 \leqslant \eta \leqslant 1$ and $\epsilon > 0$. If

$$M \geqslant \tfrac{4}{\epsilon^2}\big(K\,\log(N) + 2K\,\log(\tfrac{50}{\epsilon}) + \log(\tfrac{2}{\eta})\big),$$

then $\mathbf{\Phi}$ is a BϵSE with $\Pr > 1-\eta$.

$$M = O(\epsilon^{-2} K \log N)$$

$\Rightarrow$

BϵSE consistency "width":

$$\epsilon = O\big((\tfrac{K}{M}\log\tfrac{MN}{K})^{1/2}\big)$$

not as optimal but stronger result!

$d_H \leftrightarrow d_{\rm ang}$

4. Generalized Embeddings

Beyond strict sparsity ... [Plan, Vershynin]

Let $\mathcal{K} \subset S^{N-1}$ (*e.g.*, compressible signals s.t. $\|\boldsymbol{x}\|_2/\|\boldsymbol{x}\|_1 \leqslant \sqrt{K}$)
$\neq \Sigma_K$

What can we say on $d_H(A(\boldsymbol{x}), A(\boldsymbol{s}))$ for $\boldsymbol{x}, \boldsymbol{s} \in \mathcal{K}$?

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452
Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity ... [Plan, Vershynin]

Let $\mathcal{K} \subset S^{N-1}$ ($e.g.$, compressible signals s.t. $\|\boldsymbol{x}\|_2/\|\boldsymbol{x}\|_1 \leqslant \sqrt{K}$)
$\neq \Sigma_K$

What can we say on $d_H(A(\boldsymbol{x}), A(\boldsymbol{s}))$ for $\boldsymbol{x}, \boldsymbol{s} \in \mathcal{K}$?

Uniform tesselation: [Plan, Vershynin, 11]

$$\mathrm{P}\big(\underbrace{\#\text{ random hyperplanes btw }\boldsymbol{x}\text{ and }\boldsymbol{s}}_{d_H(A(\boldsymbol{x}),A(\boldsymbol{s}))} \propto d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s})\big)\ ?$$

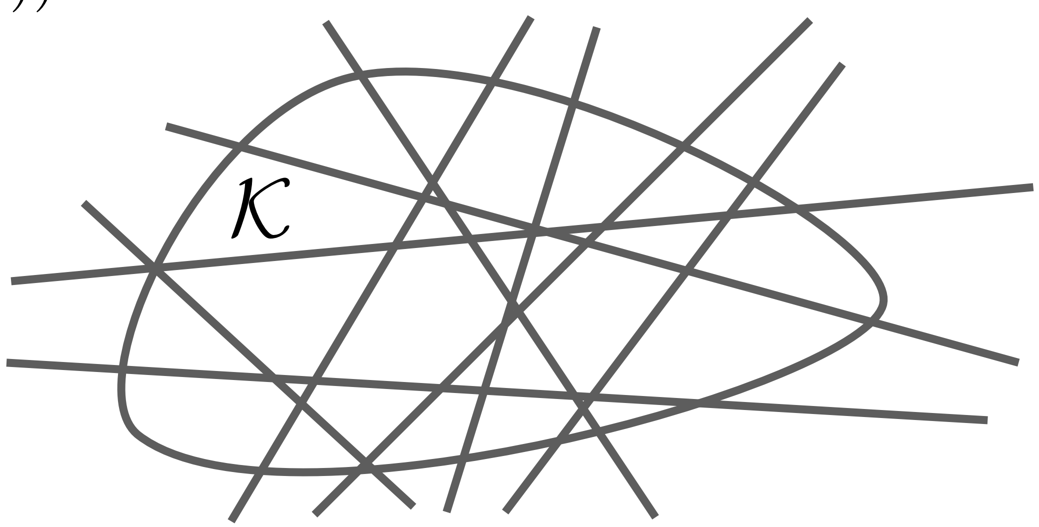

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity ... [Plan, Vershynin]

Measuring the "dimension" of $\mathcal{K}$ → Gaussian mean width:

$$w(\mathcal{K}) := \mathbb{E} \sup_{\boldsymbol{u} \in \mathcal{K} - \mathcal{K}} \langle \boldsymbol{g}, \boldsymbol{u} \rangle, \text{ with } g_k \sim_{\text{iid}} \mathcal{N}(0,1)$$

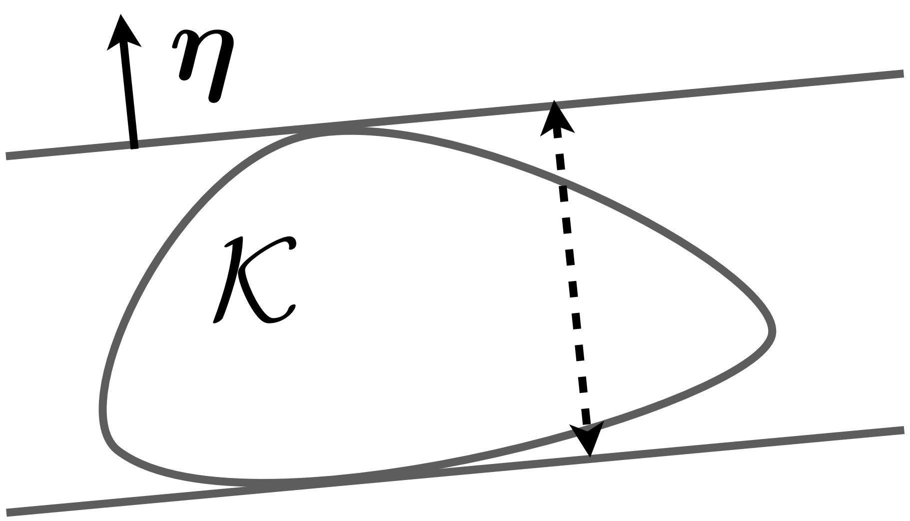

width in direction $\boldsymbol{\eta}$

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity ... [Plan, Vershynin]

Measuring the "dimension" of $\mathcal{K}$ → Gaussian mean width:

$$w(\mathcal{K}) := \mathbb{E} \sup_{\boldsymbol{u}\in\mathcal{K}-\mathcal{K}} \langle \boldsymbol{g}, \boldsymbol{u}\rangle, \text{ with } g_k \sim_{\text{iid}} \mathcal{N}(0,1)$$

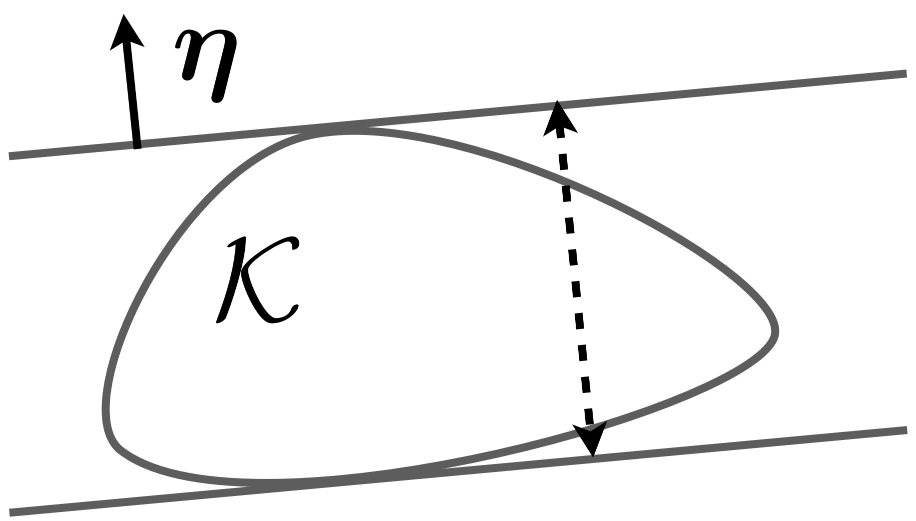

width in direction $\boldsymbol{\eta}$

Examples:

$w^2(\mathcal{S}^{N-1}) \leqslant 4N$

$w^2(\mathcal{K}) \leqslant C \log|\mathcal{K}|$ (for finite sets)

$w^2(\mathcal{K}) \leqslant L$ if subspace with $\dim \mathcal{K} = L$

$w^2(\Sigma_K) \simeq K \log(2N/K)$

⋮

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity ... [Plan, Vershynin]

Proposition *Let* $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$ *and* $\mathcal{K} \subset \mathbb{R}^N$. *Then, for some* $C, c > 0$, *if*

$$M \;\geqslant\; C\epsilon^{-6} w^2(\mathcal{K}),$$

then, with $Pr \geqslant 1 - e^{-c\epsilon^2 M}$, *we have*

$$d_{\rm ang}(\boldsymbol{x},\boldsymbol{s}) - \epsilon \;\leqslant\; d_H(A(\boldsymbol{x}), A(\boldsymbol{s})) \;\leqslant\; d_{\rm ang}(\boldsymbol{x},\boldsymbol{s}) - \epsilon, \quad \forall \boldsymbol{x},\boldsymbol{s} \in \mathcal{K}.$$

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452
Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity … [Plan, Vershynin]

Proposition *Let* $\mathbf{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$ *and* $\mathcal{K} \subset \mathbb{R}^N$. *Then, for some* $C, c > 0$, *if*

$$M \geqslant C\epsilon^{-6} w^2(\mathcal{K}),$$

not as optimal but stronger result!

then, with $Pr \geqslant 1 - e^{-c\epsilon^2 M}$, *we have*

$$d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon \leqslant d_H(A(\boldsymbol{x}), A(\boldsymbol{s})) \leqslant d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon, \quad \forall \boldsymbol{x}, \boldsymbol{s} \in \mathcal{K}.$$

Generalize BϵSE to more general sets.

In particular, to

$$\mathcal{C}_K = \{\boldsymbol{u} \in \mathbb{R}^N : \|\boldsymbol{u}\|_2/\|\boldsymbol{u}\|_1 \leqslant \sqrt{K}\} \supset \Sigma_K$$

with $w^2(\mathcal{C}_K) \leqslant cK \log N/K$.

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

Beyond strict sparsity ... [Plan, Vershynin]

Proposition *Let* $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$ *and* $\mathcal{K} \subset \mathbb{R}^N$. *Then, for some* $C, c > 0$, *if*

$$M \geqslant C\epsilon^{-6} w^2(\mathcal{K}),$$

not as optimal but stronger result!

then, with $Pr \geqslant 1 - e^{-c\epsilon^2 M}$, *we have*

$$d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon \leqslant d_H(A(\boldsymbol{x}), A(\boldsymbol{s})) \leqslant d_{\rm ang}(\boldsymbol{x}, \boldsymbol{s}) - \epsilon, \quad \forall \boldsymbol{x}, \boldsymbol{s} \in \mathcal{K}.$$

Generalize BϵSE to more general sets.

In particular, to

$$\mathcal{C}_K = \{\boldsymbol{u} \in \mathbb{R}^N : \|\boldsymbol{u}\|_2/\|\boldsymbol{u}\|_1 \leqslant \sqrt{K}\} \supset \Sigma_K$$

with $w^2(\mathcal{C}_K) \leqslant cK \log N/K$.

$\Rightarrow$ Extension to "1-bit Matrix Completion" possible!

i.e., $w^2(r\text{-rank } N_1 \times N_2 \text{ matrix}) \leqslant c\,r(N_1 + N_2)$!

Y. Plan, R. Vershynin, "Dimension reduction by random hyperplane tessellations", 2011, arXiv:1111.4452

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.

5. 1-bit CS Reconstructions?

Dumbest 1-bit reconstruction

Fact: If $M = O(\epsilon^{-2} K \log N/K)$ (for $\boldsymbol{x} \in \Sigma_K$ fixed, $\forall\, \boldsymbol{s} \in \Sigma_K$)
or, if $M = O(\epsilon^{-6} K \log N/K)$ ($\forall\, \boldsymbol{x}, \boldsymbol{s} \in \Sigma_K$), then, w.h.p,

$$\left|\tfrac{\sqrt{\pi}/2}{M}\langle \mathrm{sign}\,(\boldsymbol{\Phi x}), \boldsymbol{\Phi s}\rangle - \langle \boldsymbol{x}, \boldsymbol{s}\rangle\right| \leq \epsilon$$ [Plan, Vershynin, 12]

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.
LJ, K. Degraux, C. De Vleeschouwer, "Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing", SAMPTA2013

Dumbest 1-bit reconstruction

<u>Fact</u>: If $M = O(\epsilon^{-2} K \log N/K)$ (for $\boldsymbol{x} \in \Sigma_K$ fixed, $\forall\, \boldsymbol{s} \in \Sigma_K$)
or, if $M = O(\epsilon^{-6} K \log N/K)$ ($\forall\, \boldsymbol{x}, \boldsymbol{s} \in \Sigma_K$), then, w.h.p,

$$\left|\tfrac{\sqrt{\pi}/2}{M}\langle \mathrm{sign}\,(\boldsymbol{\Phi x}), \boldsymbol{\Phi s}\rangle - \langle \boldsymbol{x}, \boldsymbol{s}\rangle\right| \leq \epsilon \qquad \text{[Plan, Vershynin, 12]}$$

- Implication? [LJ, Degraux, De Vleeschouwer, 13]

Let $\boldsymbol{x} \in \Sigma_K \cap S^{N-1}$ and $\boldsymbol{q} = \mathrm{sign}(\boldsymbol{\Phi x})$.
Compute

$$\hat{\boldsymbol{x}} = \tfrac{\pi}{2M}\, \mathcal{H}_K(\boldsymbol{\Phi}^* \boldsymbol{q})$$

Then, if previous property holds,

$$\|\boldsymbol{x} - \hat{\boldsymbol{x}}\| \leq 2\epsilon.$$

Non-uniform case ($\boldsymbol{x}$ given):
$\Rightarrow \epsilon = O\big((\tfrac{K}{M} \log \tfrac{MN}{K})^{1/2}\big)$
Uniform case:
$\Rightarrow \epsilon = O\big((\tfrac{K}{M} \log \tfrac{MN}{K})^{1/6}\big)$

Y. Plan, R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach", IEEE TIT 2012, arXiv:1202.1212.
LJ, K. Degraux, C. De Vleeschouwer, "Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing", SAMPTA2013

Initial approach

- Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x}) =: A(\boldsymbol{x})$
- Initially: [Boufounos, Baraniuk 2008]

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u}} \|\boldsymbol{u}\|_1 \quad \text{s.t.} \quad \mathrm{diag}(\boldsymbol{q})\,\boldsymbol{\Phi u} > 0 \quad \text{and} \quad \|\boldsymbol{u}\|_2 = 1$$

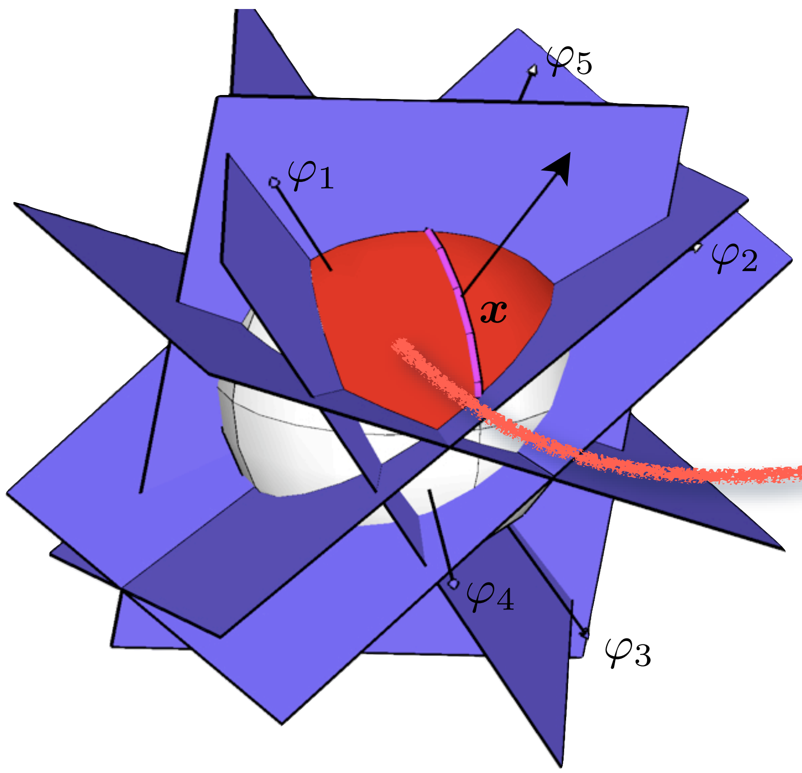

Non-convex! 2 numerical choices :

1. relax + projection on $S^{N\text{-}1}$
2. *"trust region methods"*
   → *Restricted-Step Shrinkage (RSS)*

Consistency constraint:

$$\{\boldsymbol{u} \in \mathbb{R}^N \cap S^{N-1} : \boldsymbol{q} = A(\boldsymbol{u})\}$$

$$\Leftrightarrow \{\boldsymbol{u} \in \mathbb{R}^N \cap S^{N-1} : \mathrm{diag}(\boldsymbol{q})\boldsymbol{\Phi u} > 0\}$$

$$\ni \boldsymbol{x}$$

Initial approach

- Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x}) =: A(\boldsymbol{x})$
- Initially: [Boufounos, Baraniuk 2008]

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u}} \|\boldsymbol{u}\|_1 \quad \text{s.t.} \quad \mathrm{diag}(\boldsymbol{q})\,\boldsymbol{\Phi u} > 0 \quad \text{and} \quad \|\boldsymbol{u}\|_2 = 1$$

(*e.g.*, take the 1st choice)

(relaxed) $\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u}} \|\boldsymbol{u}\|_1 + \lambda \|(\mathrm{diag}(\boldsymbol{q})\,\boldsymbol{\Phi u})_-\|^2 \quad \text{s.t.} \quad \|\boldsymbol{u}\|_2 = 1$

$\rightarrow$ Solved by projected gradient descent

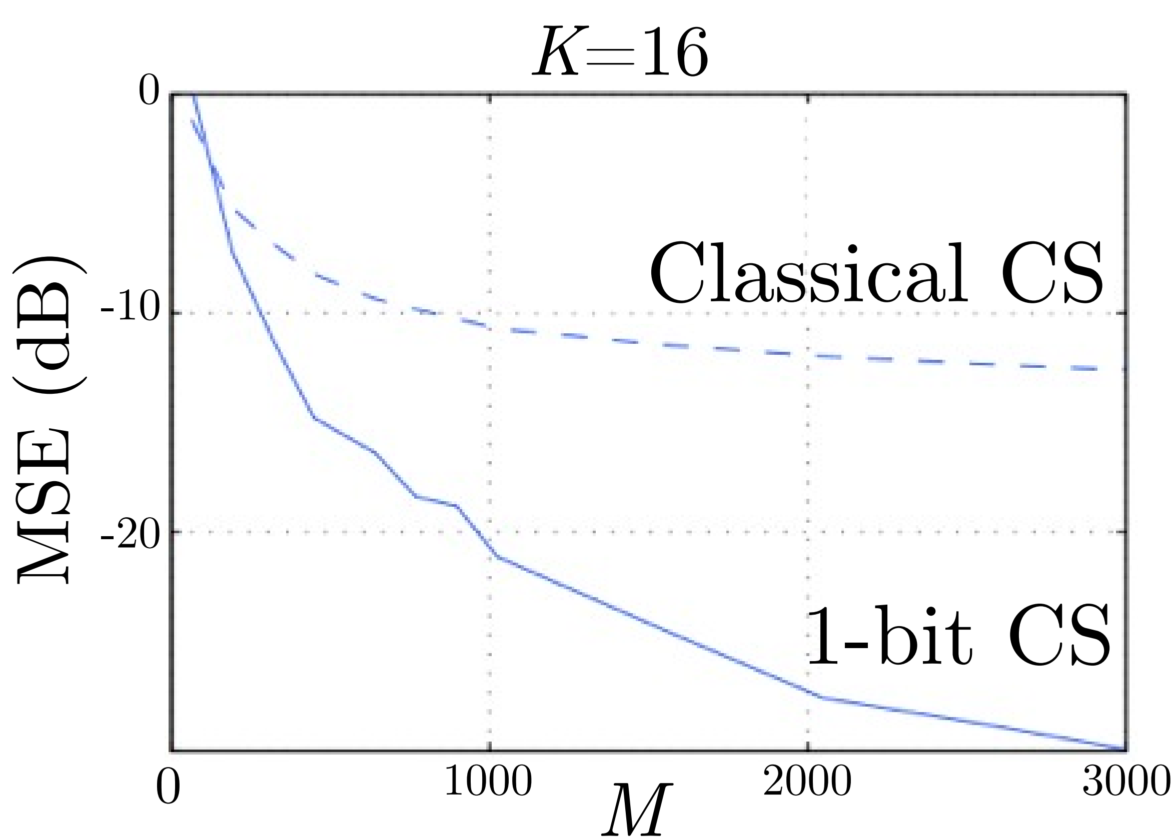

gain brought by (almost)consistency

Initial approach

- Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x}) =: A(\boldsymbol{x})$
- Initially: [Boufounos, Baraniuk 2008]

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u}} \|\boldsymbol{u}\|_1 \quad \text{s.t.} \quad \mathrm{diag}(\boldsymbol{q})\,\boldsymbol{\Phi u} > 0 \quad \text{and} \quad \|\boldsymbol{u}\|_2 = 1$$

($e.g.$, take the 1st choice)

(relaxed) $$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{u}} \|\boldsymbol{u}\|_1 + \lambda\|(\mathrm{diag}(\boldsymbol{q})\,\boldsymbol{\Phi u})_-\|^2 \quad \text{s.t.} \quad \|\boldsymbol{u}\|_2 = 1$$

$\rightarrow$ Solved by projected gradient descent

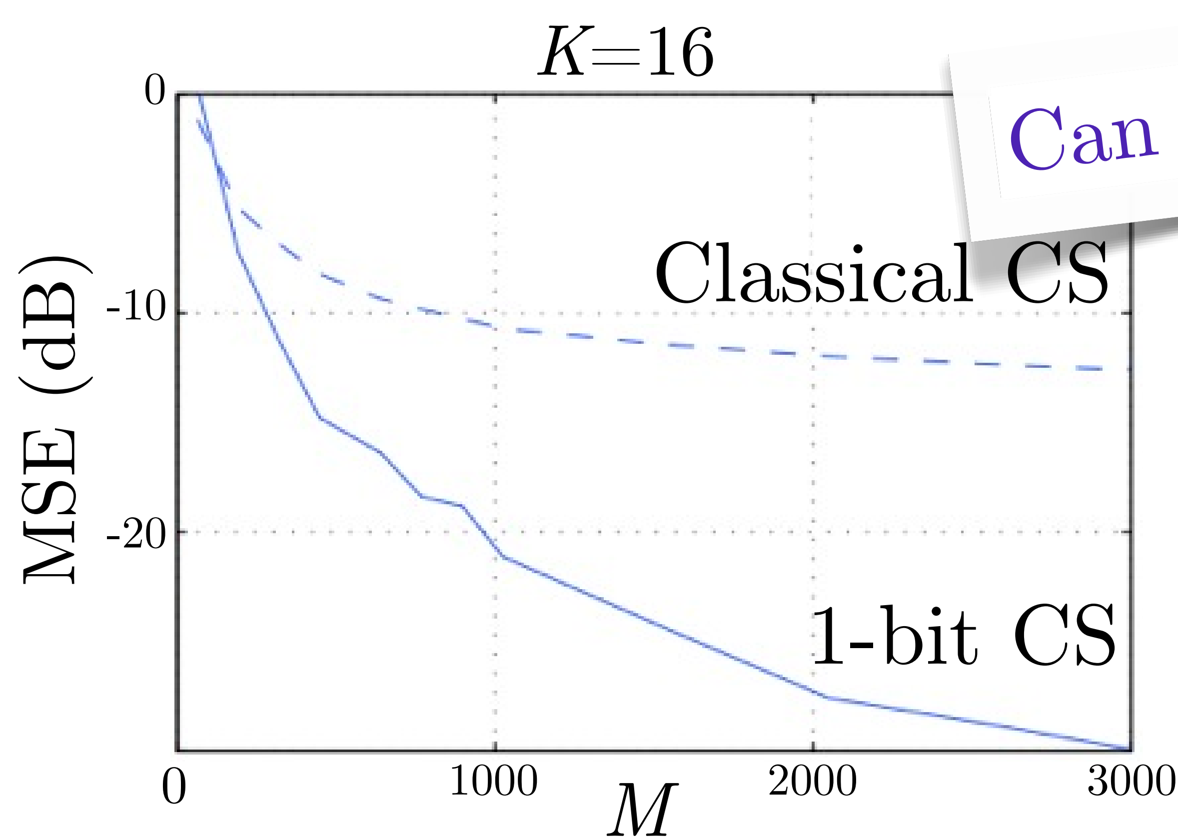

Can we do better?

gain brought by (almost)consistency

Other methods:

- ✓ Matching Sign Pursuit [Boufounos]
- Restricted-Step Shrinkage (RSS) [Laska, We, Yin, Baraniuk]
- ✓ Binary Iterative Hard Thresholding [Jacques, Laska, Boufounos, Baraniuk]
- ✓ Convex Optimization [Plan, Vershynin]
- ...

Matching Sign Pursuit (MSP)

- Iterative greedy algorithm, similar to CoSaMP [Needell, Tropp, 08]
- Maintains running signal estimate and its support T.
- MSP iteration:
  - Identify sign violations $\longrightarrow \boldsymbol{r} = (\mathrm{diag}(\boldsymbol{y})\,\boldsymbol{\Phi}\widehat{\boldsymbol{x}})_-$
  - Compute proxy $\longrightarrow \boldsymbol{p} = \boldsymbol{\Phi}^T\boldsymbol{r}$
  - Identify support $\longrightarrow \Omega = \mathrm{supp}\,\boldsymbol{p}|_{2K} \cup T$
  - Consistent Reconstruction over support estimate:

$$\boldsymbol{b}|_\Omega = \arg\min_{\boldsymbol{u}\in\mathbb{R}^N} \|(\mathrm{diag}(\boldsymbol{y})\boldsymbol{\Phi}\boldsymbol{u})_-\|_2^2 \text{ s.t } \|\boldsymbol{u}\|_2 = 1 \text{ and } \boldsymbol{u}|_{T^c} = 0$$

  - Truncate, normalize, and update estimate: $\widehat{\boldsymbol{x}} \leftarrow \boldsymbol{b}|_K \,/\, \|\boldsymbol{b}|_K\|_2$

Matching Sign Pursuit (MSP)

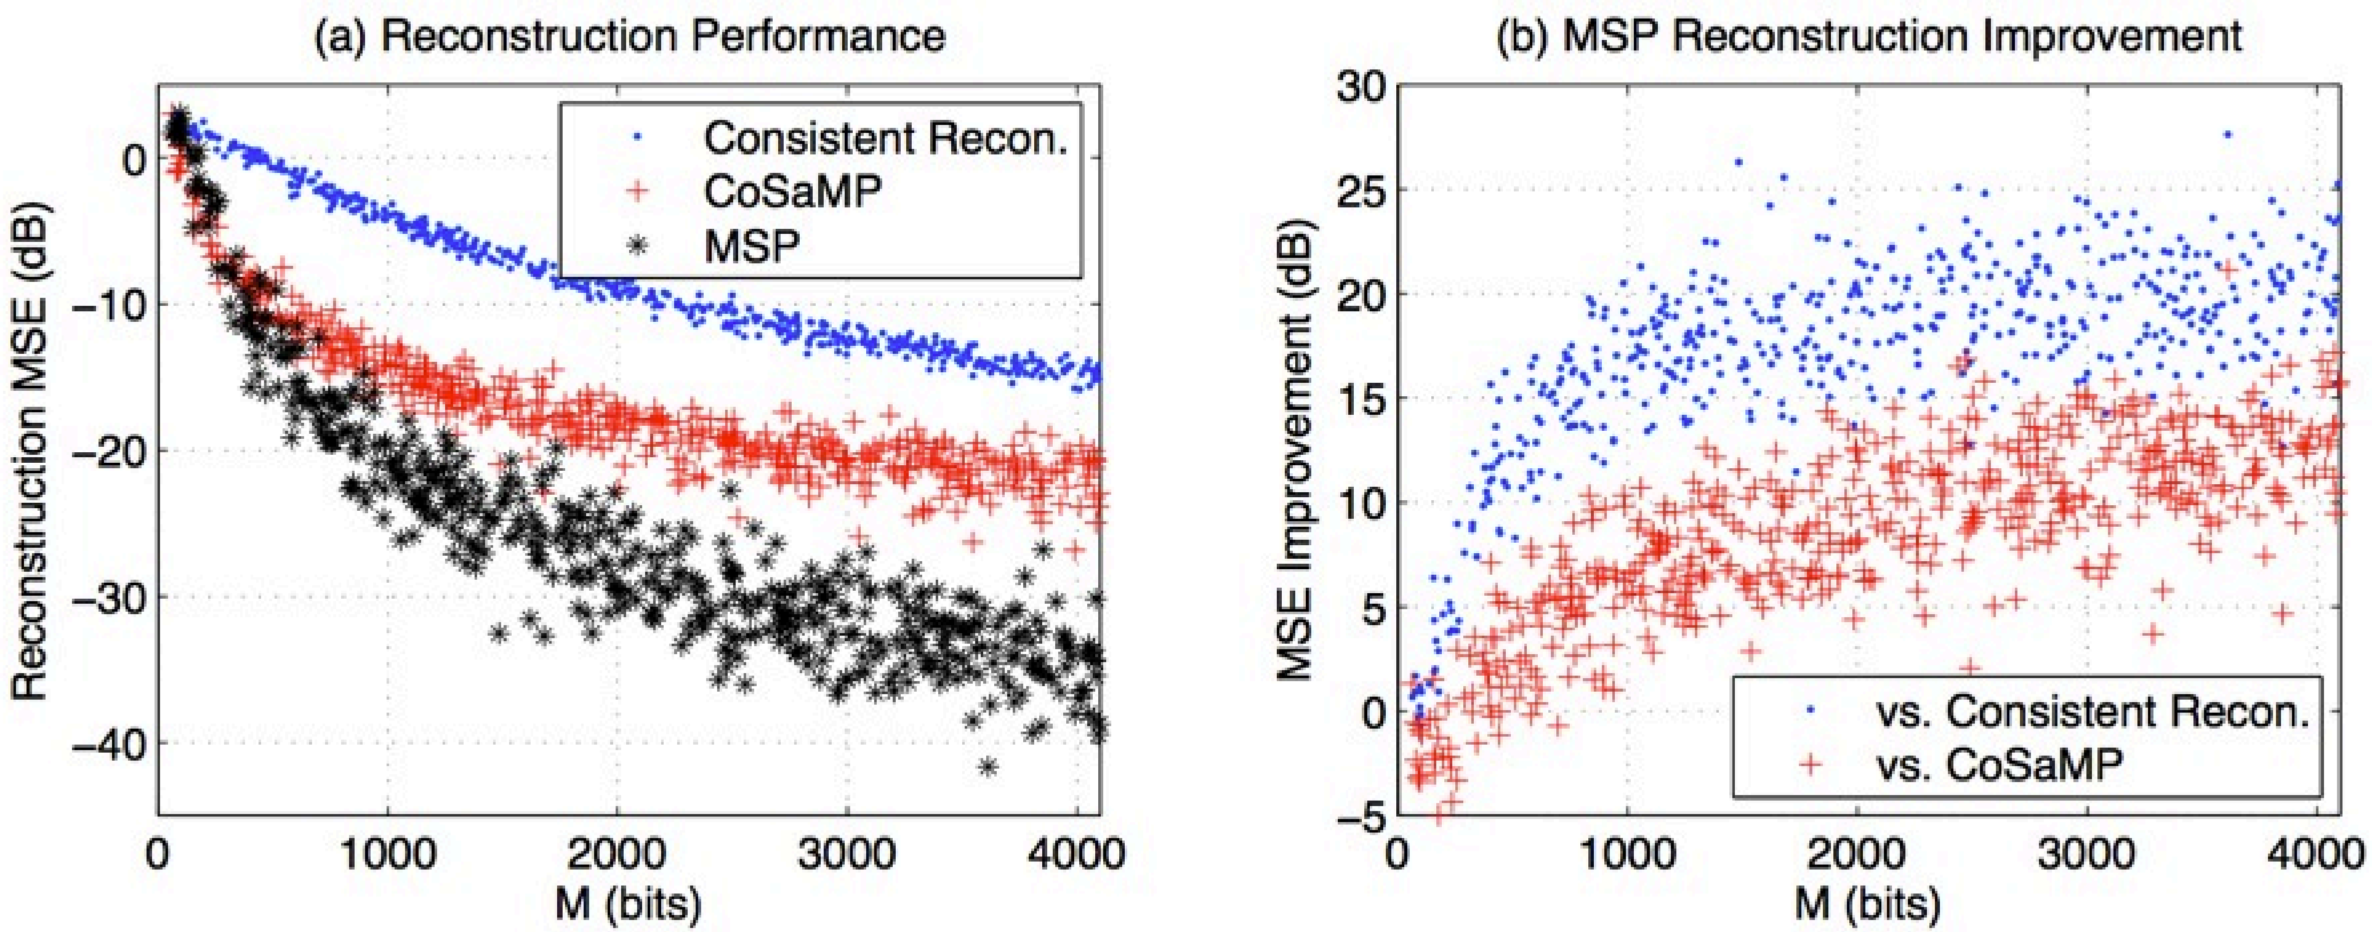

Boufounos, P. T. (2009, November). "Greedy sparse signal reconstruction from sign measurements".
In Signals, Systems and Computers, 2009 Conference Record of the Forty-Third Asilomar Conference on (pp. 1305-1309). IEEE.

Binary Iterative Hard Thresholding

Given $\boldsymbol{q} = A(\boldsymbol{x})$ and K, set $l = 0$, $\boldsymbol{x}^0 = 0$:

$$\boldsymbol{a}^{l+1} = \boldsymbol{x}^l + \tfrac{\tau}{2}\boldsymbol{\Phi}^T\big(\boldsymbol{q} - A(\boldsymbol{x}^l)\big),$$

("gradient" towards consistency)
($\tau > 0$ controls gradient descent)

$$\boldsymbol{x}^{l+1} = \mathcal{H}_K(\boldsymbol{a}^{l+1}), \quad l \leftarrow l+1$$

(proj. K-sparse signal set)

with $\mathcal{H}_K(\boldsymbol{u}) = K$-term hard thresholding

Stop when $d_H(\boldsymbol{q}, A(\boldsymbol{x}^{l+1})) = 0$ or $l =$ max. iter.

minimizes $\mathcal{J}(\boldsymbol{x}') = \|[\,\mathrm{diag}(\boldsymbol{q})(\boldsymbol{\Phi}\boldsymbol{x}')]_-\|_1$ with $(\lambda)_- = (\lambda - |\lambda|)/2$

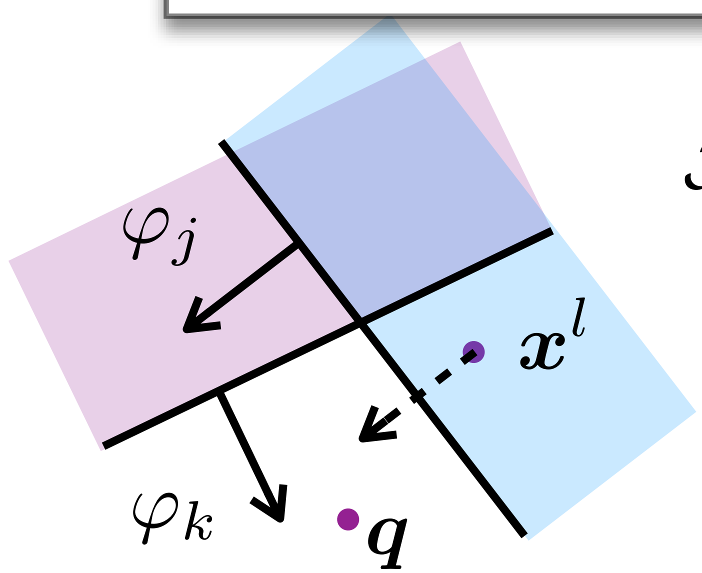

$$\mathcal{J}(\boldsymbol{x}') = \sum_{j=1}^{M} \Big|\big(\overbrace{\mathrm{sign}\,(\langle\boldsymbol{\varphi}_j, \boldsymbol{x}\rangle)}^{q_j}\,\langle\boldsymbol{\varphi}_j, \boldsymbol{x}'\rangle\big)_-\Big|$$

$$q_k - A(\boldsymbol{x}^l)_k = 0$$

$$q_j - A(\boldsymbol{x}^l)_j > 0$$

(connections with ML hinge loss, 1-bit classification)

Binary Iterative Hard Thresholding

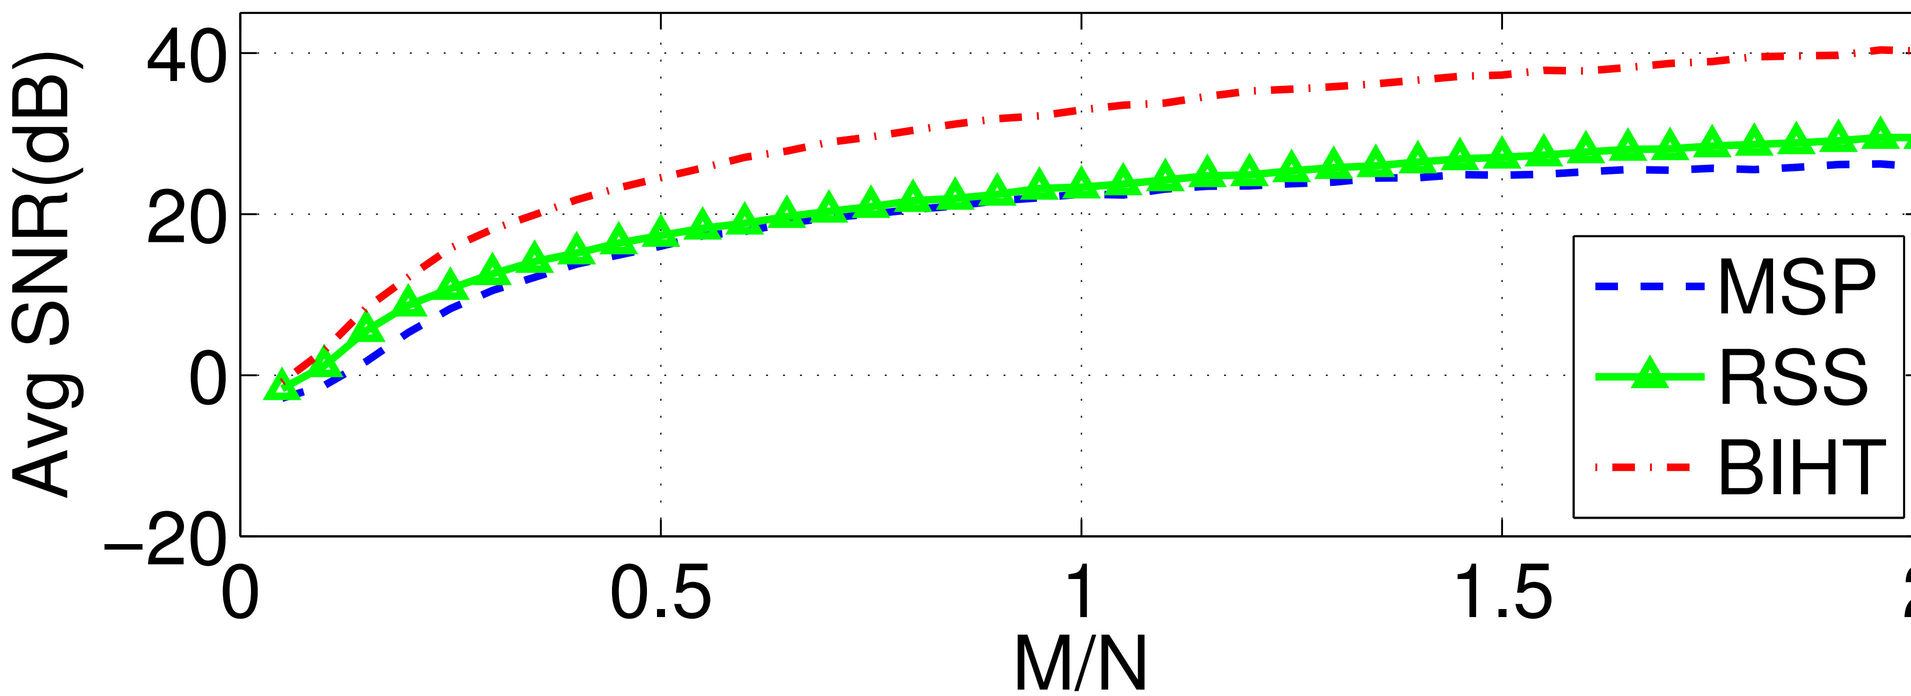

$N = 1000,\ K = 10$
Bernoulli-Gaussian model
normalized signals
1000 trials

Matching Sign pursuit (MSP)
Restricted-Step Shrinkage (RSS)
Binary Iterative Hard Thresholding (BIHT)

Binary Iterative Hard Thresholding

- Testing BϵSE: $d_{\rm ang}(\boldsymbol{x}, \boldsymbol{x}^*) \leqslant d_H(A(\boldsymbol{x}), A(\boldsymbol{x}^*)) + \epsilon(M)$

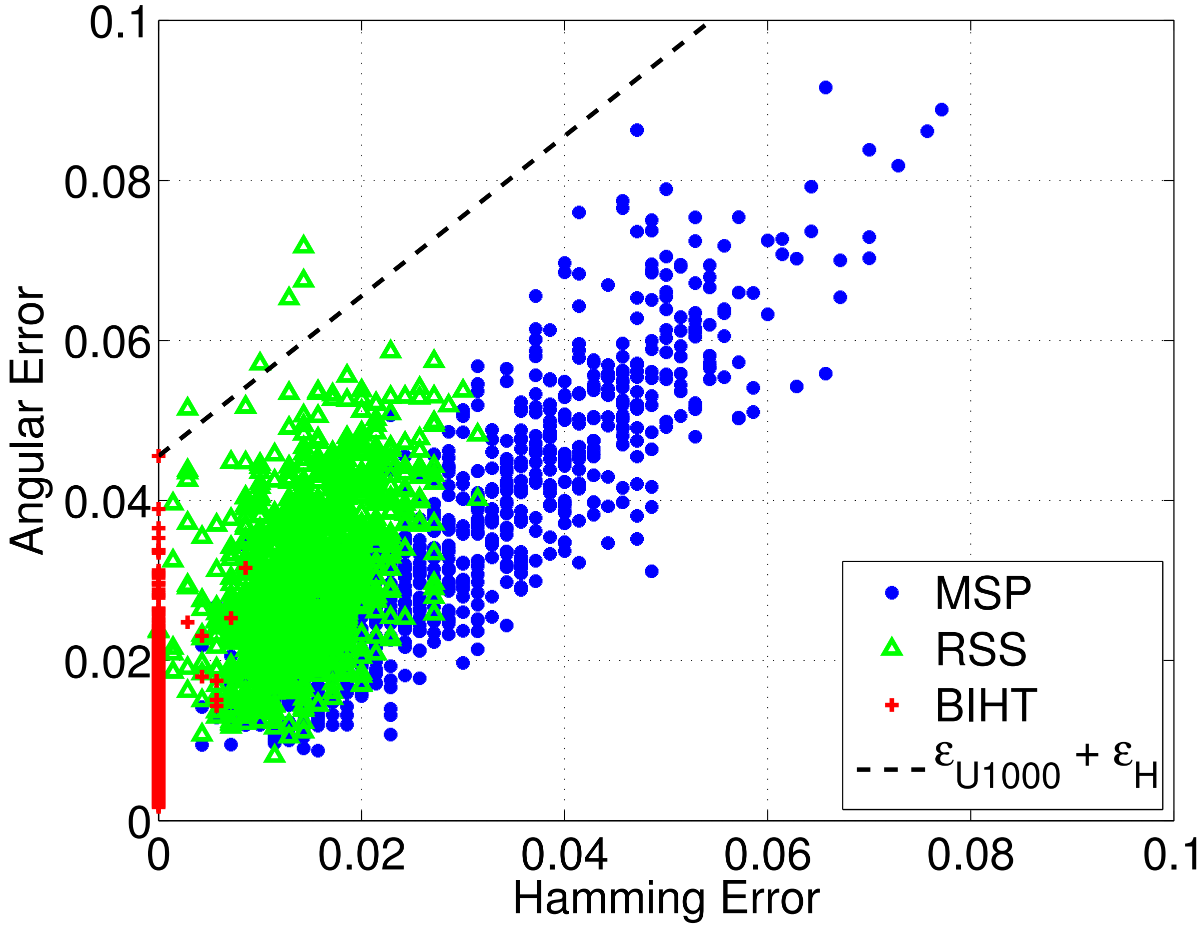

$M/N = 0.7$

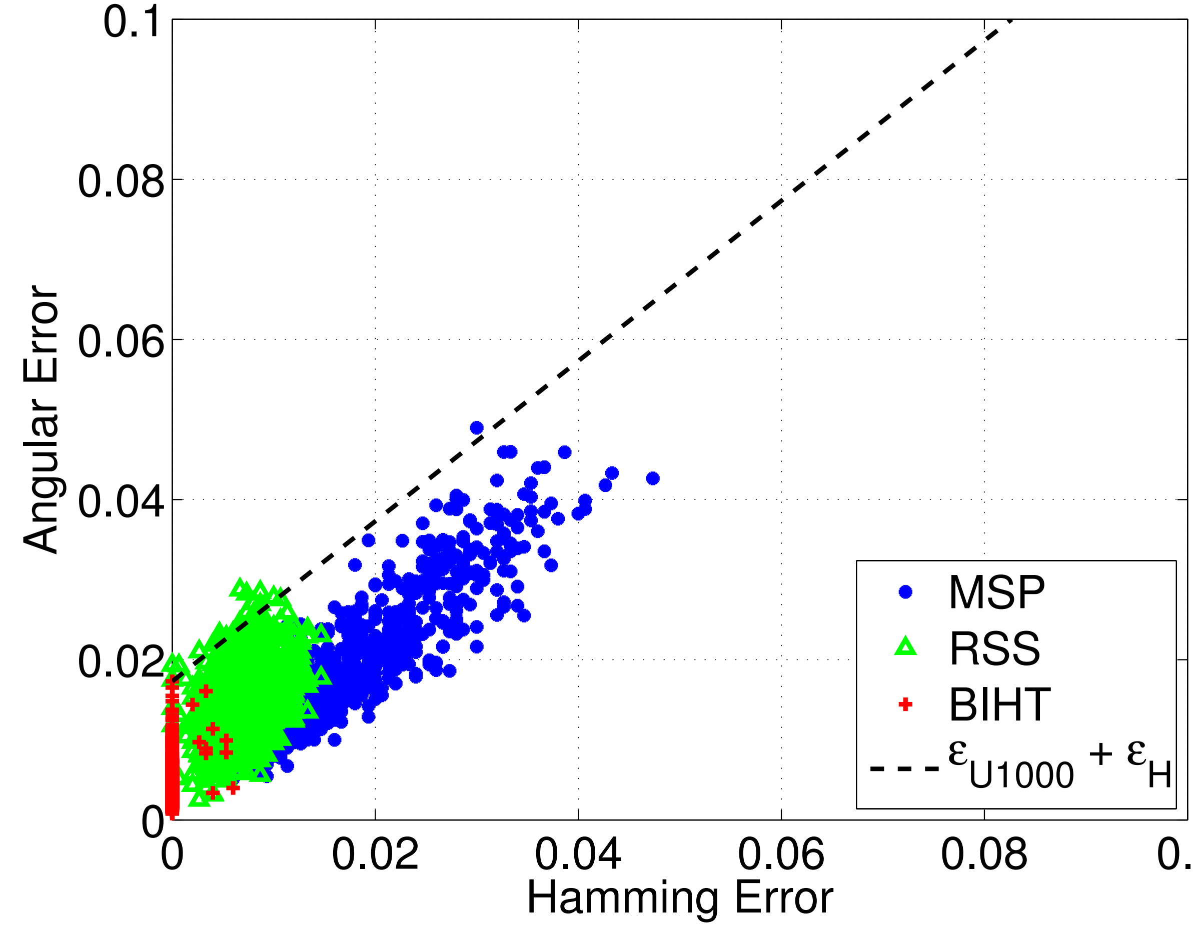

$M/N = 1.5$

Remark: CS vs bits/meas.

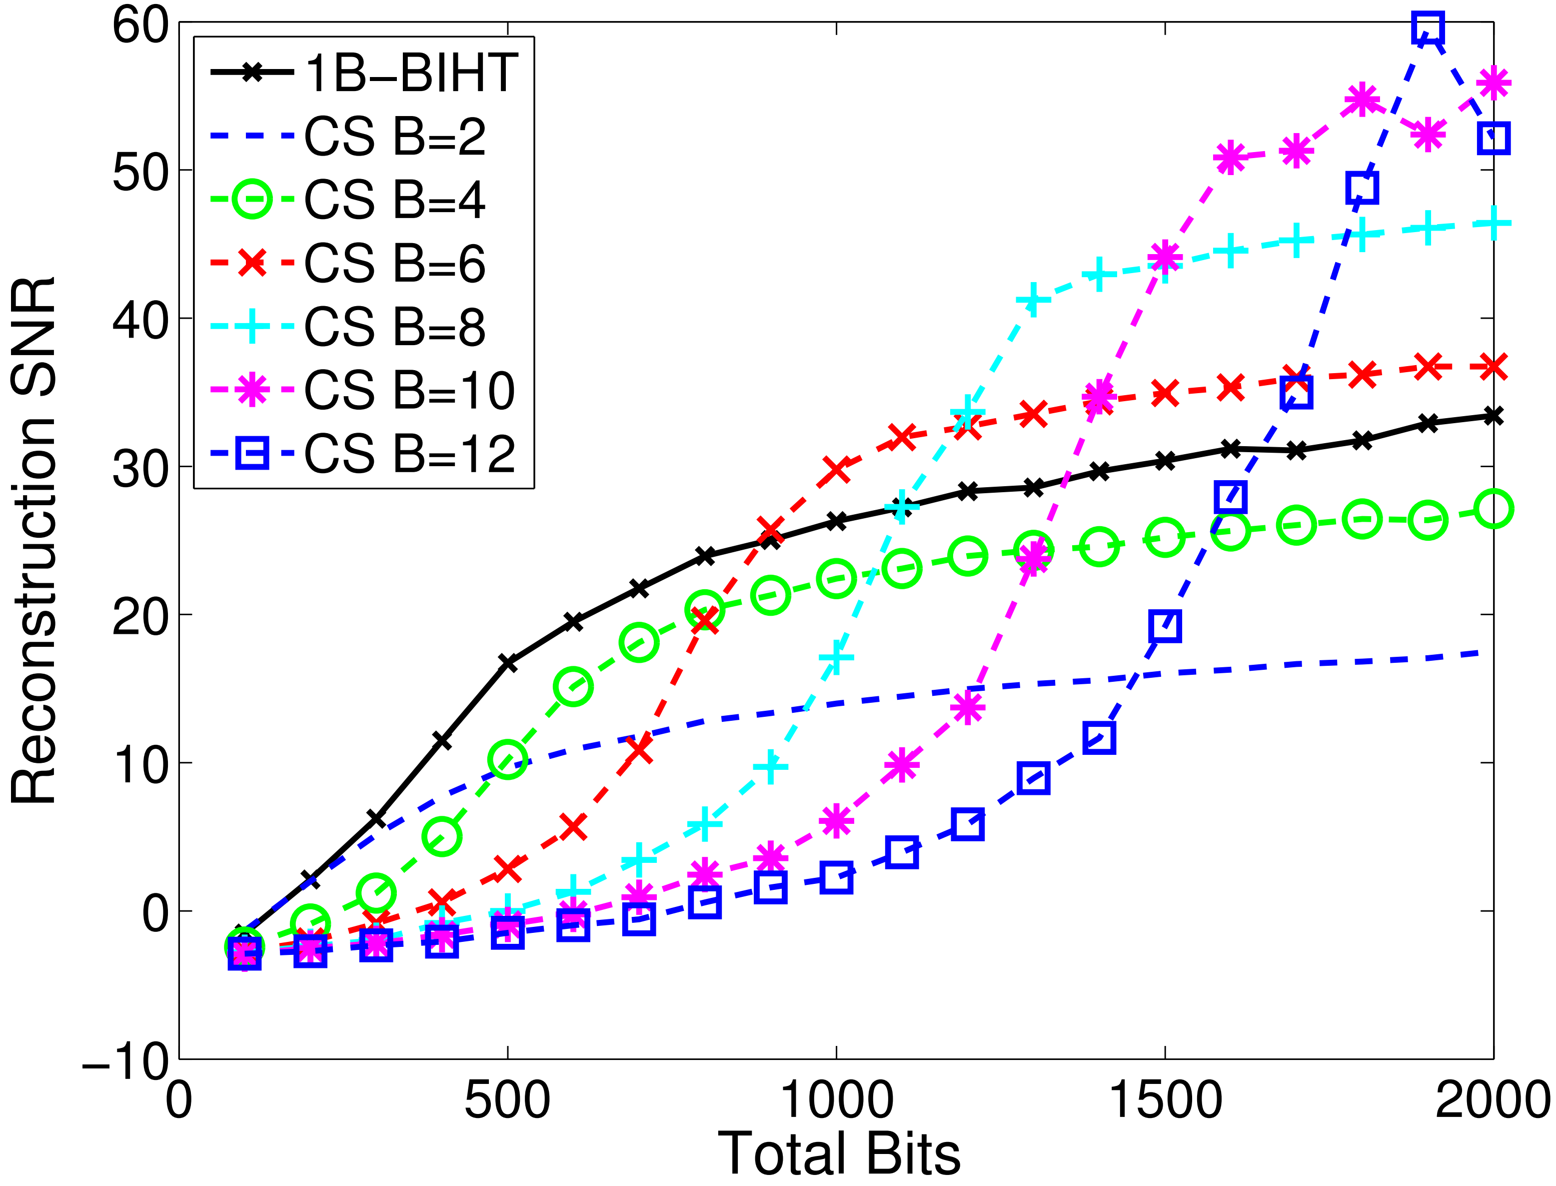

$N = 2000, \; K = 20$

Bernoulli-Gaussian model

normalized signals

B bits/measurement

$B = 1, \ldots, 12$

$M = \text{Total Bits}/B$

1000 trials

Convex Optimization [Plan, Vershynin, 12]

Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x})$ for some signal $\boldsymbol{x} \in \mathcal{K} \subset B_2^N$

e.g., sparse, compressible, low-rank matrix

Compute $$\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{u} \in \mathbb{R}^N} \boldsymbol{q}^T \boldsymbol{\Phi u} \quad \text{s.t.} \quad \boldsymbol{u} \in \mathcal{K}$$

maximize consistency

Convex problem if $\mathcal{K}$ convex!
No ambiguous amplitude definition
($\boldsymbol{u} = 0$ avoided)

S. Bahmani, P.T. Boufounos, B. Raj, "Robust 1-bit Compressive Sensing via Gradient Support Pursuit", arxiv:1304.6626

Convex Optimization [Plan, Vershynin, 12]

Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x})$ for some signal $\boldsymbol{x} \in \mathcal{K} \subset B_2^N$

e.g., sparse, compressible, low-rank matrix

Compute

$$\hat{\boldsymbol{x}} \;=\; \arg\max_{\boldsymbol{u}\in\mathbb{R}^N} \; \boldsymbol{q}^T \boldsymbol{\Phi u} \quad \text{s.t.} \quad \boldsymbol{u} \in \mathcal{K}$$

maximize consistency

Convex problem if $\mathcal{K}$ convex!
No ambiguous amplitude definition
($\boldsymbol{u} = 0$ avoided)

Remark: (PV-L0 problem) [Bahmani, Boufounos, Raj, 13]

$$\hat{\boldsymbol{x}} \;=\; \frac{1}{\|\mathcal{H}_K(\boldsymbol{\Phi}^*\boldsymbol{q})\|}\,\mathcal{H}_K(\boldsymbol{\Phi}^*\boldsymbol{q}) \text{ if } \mathcal{K} = \Sigma_K \text{ !!}$$

S. Bahmani, P.T. Boufounos, B. Raj, "Robust 1-bit Compressive Sensing via Gradient Support Pursuit", arxiv:1304.6626

Convex Optimization [Plan, Vershynin, 12]

Let $\boldsymbol{q} = \mathrm{sign}\,(\boldsymbol{\Phi x})$ for some signal $\boldsymbol{x} \in \mathcal{K} \subset B_2^N$

e.g., sparse, compressible, low-rank matrix

Compute $$\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{u}\in\mathbb{R}^N} \boldsymbol{q}^T \boldsymbol{\Phi u} \quad \text{s.t.} \quad \boldsymbol{u} \in \mathcal{K}$$

maximize consistency

-2 if $\boldsymbol{x}$ is fixed

Proposition *(assuming $\|\boldsymbol{x}\| = 1$) For some $C, c > 0$, if $M \geqslant C\epsilon^{-6} w^2(\mathcal{K})$, then, with $Pr \geqslant 1 - e^{-c\epsilon^2 M}$, we have $\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|^2 \leqslant \sqrt{\frac{\pi}{2}}\,\epsilon$.*

Convex Optimization [Plan, Vershynin, 12]

Let $\boldsymbol{q} = \operatorname{sign}(\boldsymbol{\Phi x})$ for some signal $\boldsymbol{x} \in \mathcal{K} \subset B_2^N$

Compute $\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{u} \in \mathbb{R}^N} \boldsymbol{q}^T \boldsymbol{\Phi u} \quad \text{s.t.} \quad \boldsymbol{u} \in \mathcal{K}$

Proposition *(assuming $\|\boldsymbol{x}\| = 1$) For some $C, c > 0$, if $M \geqslant C\epsilon^{-6} w^2(\mathcal{K})$, then, with $Pr \geqslant 1 - e^{-c\epsilon^2 M}$, we have $\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|^2 \leqslant \sqrt{\frac{\pi}{2}}\,\epsilon$.*

\+ Robust to noise:

noise (bit flip) — noise power

Let $\boldsymbol{q}_{\rm n} = \operatorname{diag}(\boldsymbol{\eta})\,\boldsymbol{q}$ with $\eta_i \in \{\pm 1\}^M$, and assume $d_H(\boldsymbol{q}, \boldsymbol{q}_{\rm n}) \leqslant p$

(under the same conditions)

$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|^2 \leqslant \epsilon\sqrt{\log e/\epsilon} + 11\,p\sqrt{\log e/p}$

Note: if $M = O(\epsilon^{-2}(p - 1/2)^{-2} K \log N/K)$
this term disappears if $\eta_i = \pm 1$ are iid RVs (with $P(\eta_i = 1) = p$)

5. Playing with thresholds in 1-bit CS

Thresholds?

- Given $\boldsymbol{x} \in \mathbb{R}^N$ ($e.g.$, sparse)
  Is there an interest in sensing

$$\mathrm{sign}\left(\langle \boldsymbol{\varphi}, \boldsymbol{x} \rangle - \tau\right)$$

  for some (random) $\boldsymbol{\varphi}$ and $\tau \in \mathbb{R}$?

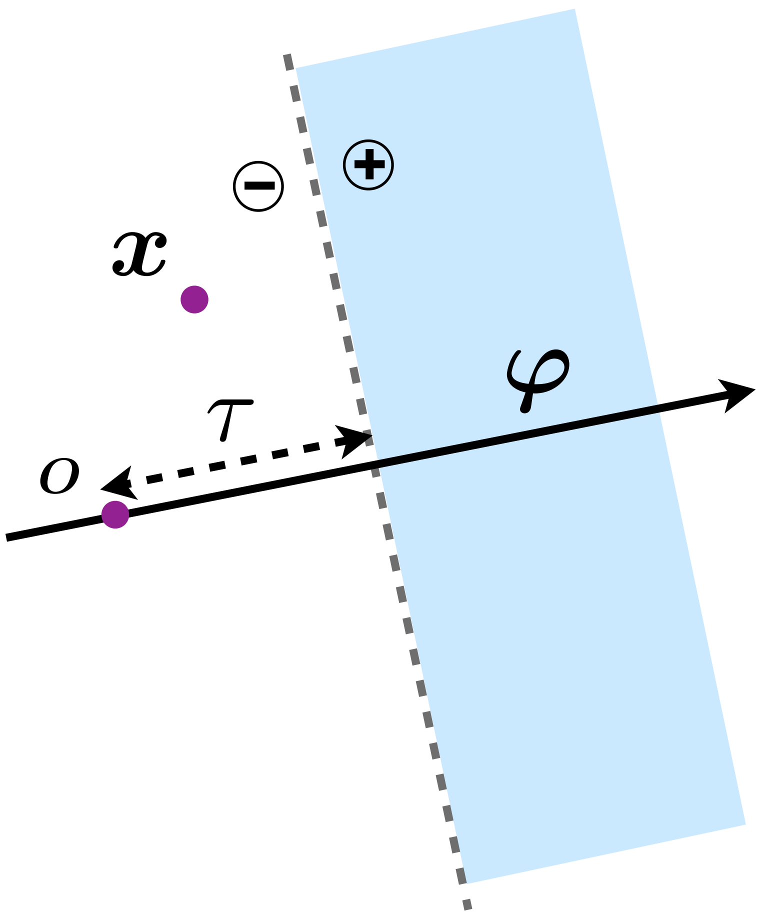

Thresholds?

- Given $\boldsymbol{x} \in \mathbb{R}^N$ (*e.g.*, sparse)
  Is there an interest in sensing

$$\operatorname{sign}(\langle \boldsymbol{\varphi}, \boldsymbol{x} \rangle - \tau)$$

  for some (random) $\boldsymbol{\varphi}$ and $\tau \in \mathbb{R}$?

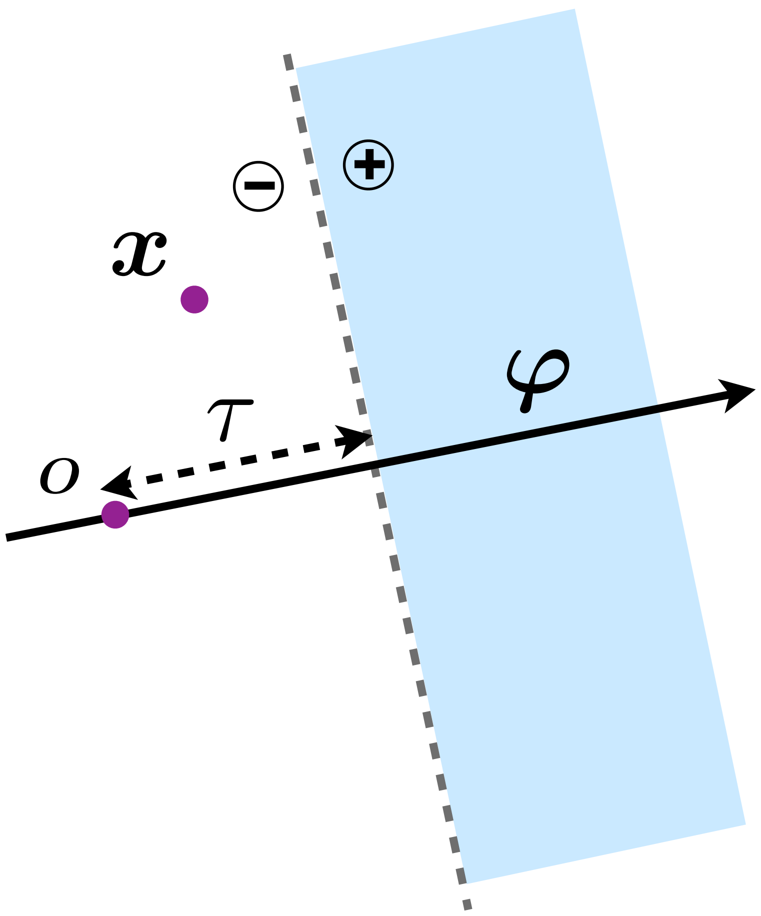

- Two recent applications:
  - adaptive thresholds [Kamilov, Bourquard, Amini, Unser, 12]
  - bridging 1-bit and B-bits QCS [LJ, Degraux, De Vleeschouwer, 13]

1-bit CS with adaptive thresholds

Non-adaptive 1-bit CS $(\tau = 0)$

Reminder

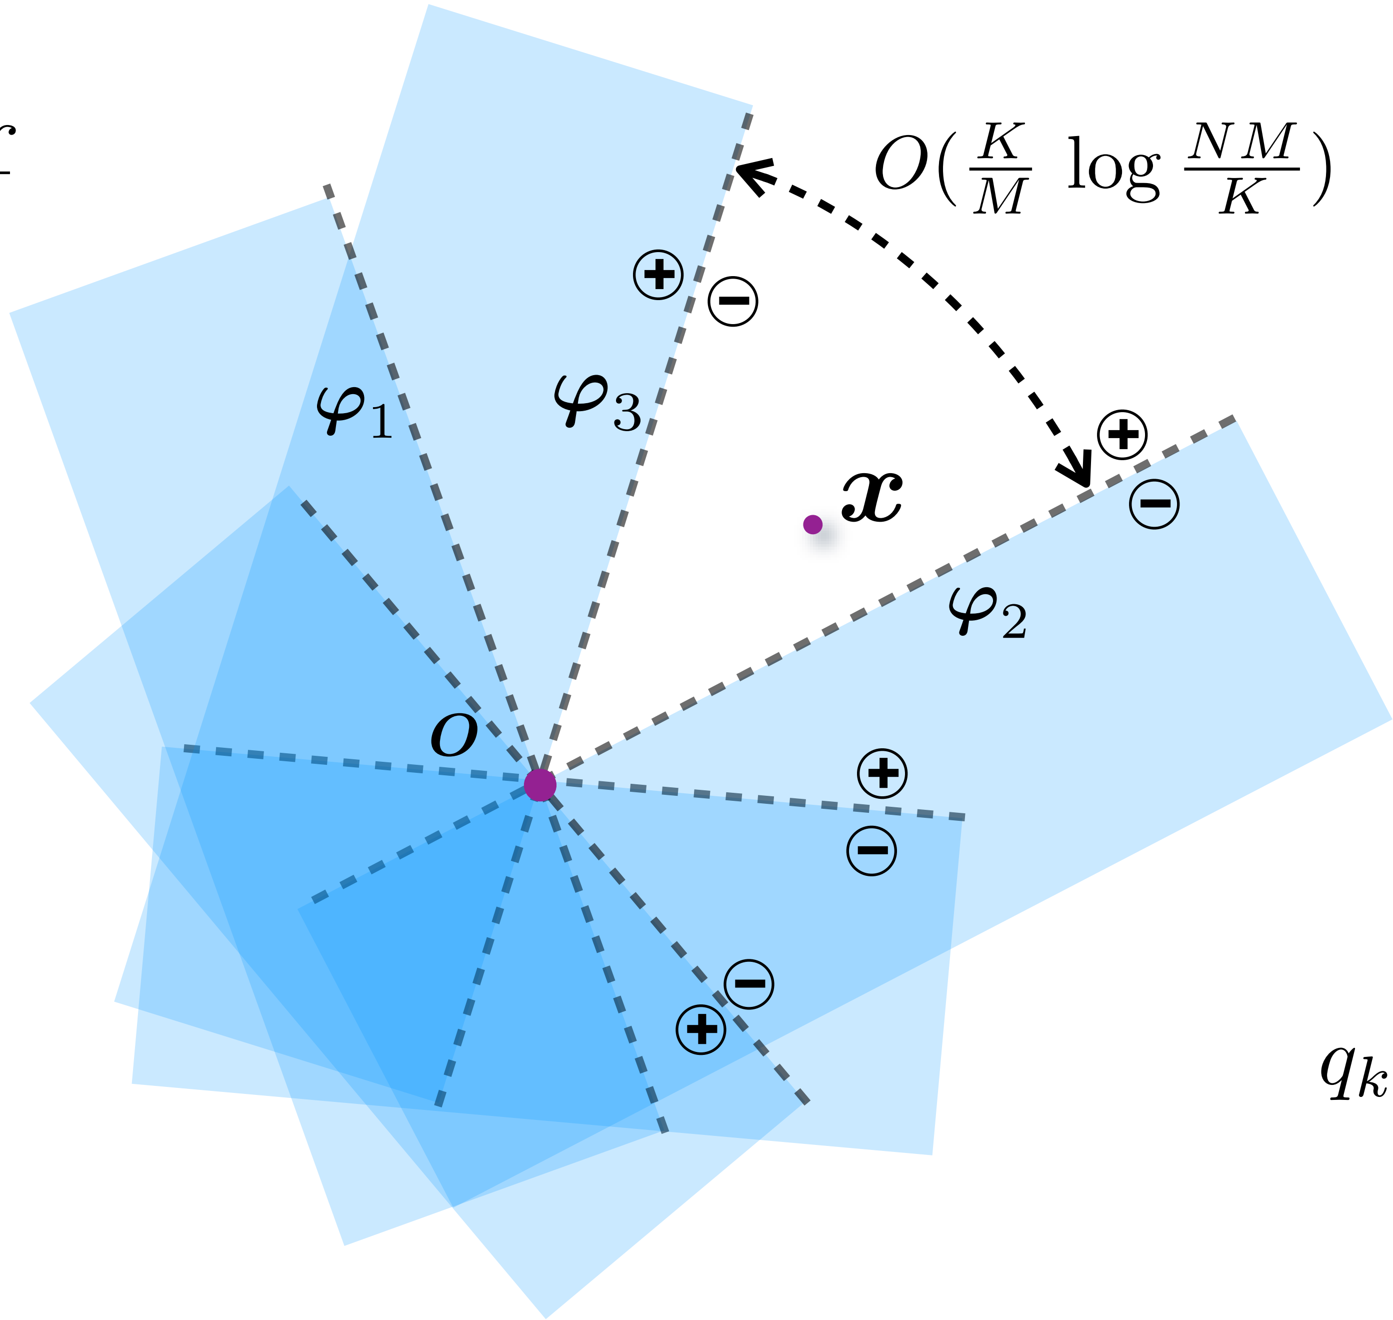

$q_k = \text{sign}\,(\langle \varphi_k, x \rangle)$

1-bit CS with adaptive thresholds

Adaptive 1-bit CS [Kamilov, Bourquard, Amini, Unser, 12]

Given a decoder Rec()

adapted from prev. meas.

$$q_k = \mathrm{sign}\left(\langle \boldsymbol{\varphi}_k, \boldsymbol{x}\rangle - \tau_k\right)$$

$$\begin{cases} \hat{\boldsymbol{x}}_k := \mathrm{Rec}(y_1, \cdots, y_k, \boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_k, \tau_1, \cdots, \tau_k) \\ \tau_{k+1} \text{ s.t. } \langle \boldsymbol{\varphi}_{k+1}, \hat{\boldsymbol{x}}_k\rangle - \tau_{k+1} = 0 \end{cases}$$

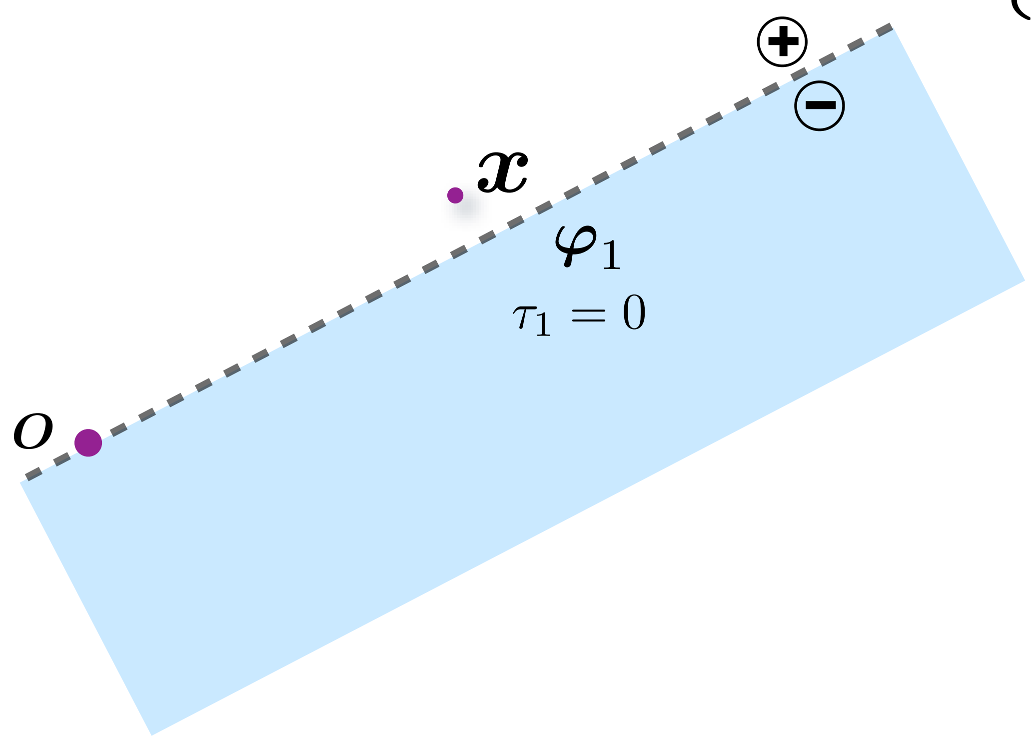

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

Adaptive 1-bit CS [Kamilov, Bourquard, Amini, Unser, 12]

Given a decoder Rec()

adapted from prev. meas.

$$q_k = \mathrm{sign}\left(\langle \boldsymbol{\varphi}_k, \boldsymbol{x} \rangle - \tau_k\right)$$

$$\begin{cases} \hat{\boldsymbol{x}}_k := \mathrm{Rec}(y_1, \cdots, y_k, \boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_k, \tau_1, \cdots, \tau_k) \\ \tau_{k+1} \text{ s.t. } \langle \boldsymbol{\varphi}_{k+1}, \hat{\boldsymbol{x}}_k \rangle - \tau_{k+1} = 0 \end{cases}$$

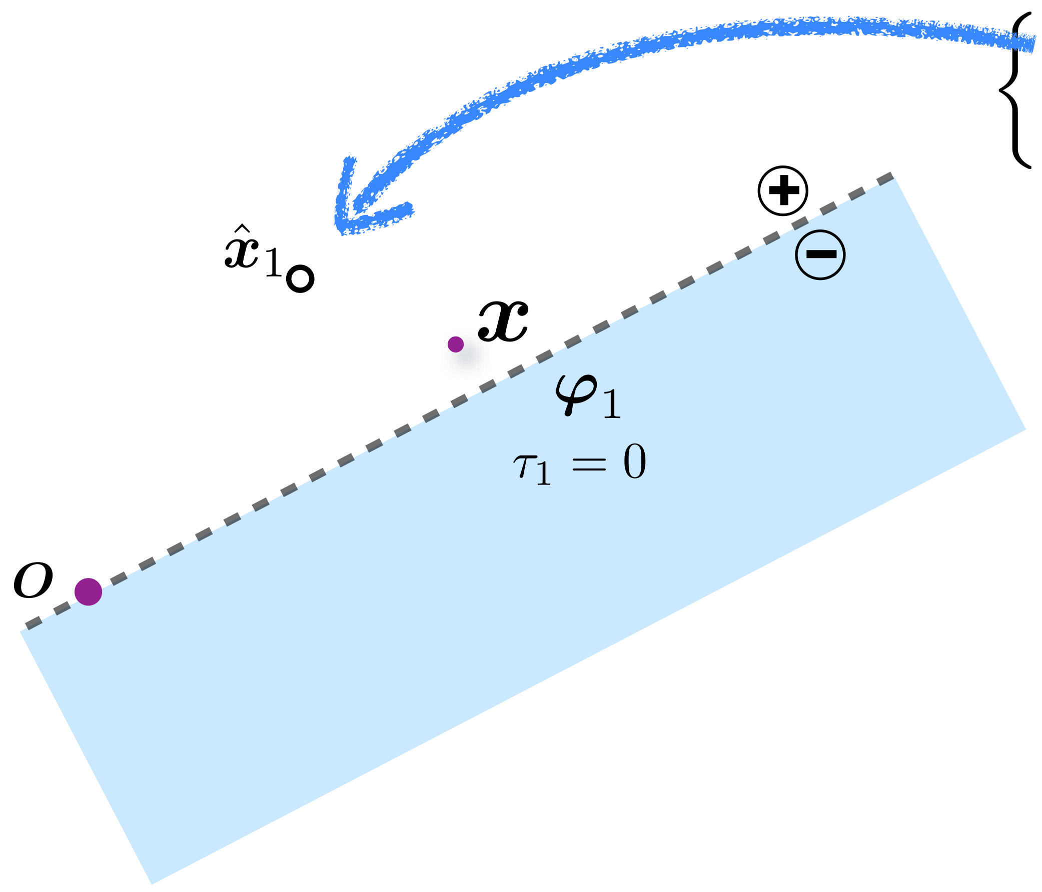

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

Adaptive 1-bit CS [Kamilov, Bourquard, Amini, Unser, 12]

Given a decoder Rec()

adapted from prev. meas.

$$q_k = \mathrm{sign}\left(\langle \boldsymbol{\varphi}_k, \boldsymbol{x}\rangle - \tau_k\right)$$

$$\begin{cases} \hat{\boldsymbol{x}}_k := \mathrm{Rec}(y_1, \cdots, y_k, \boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_k, \tau_1, \cdots, \tau_k) \\ \tau_{k+1} \text{ s.t. } \langle \boldsymbol{\varphi}_{k+1}, \hat{\boldsymbol{x}}_k\rangle - \tau_{k+1} = 0 \end{cases}$$

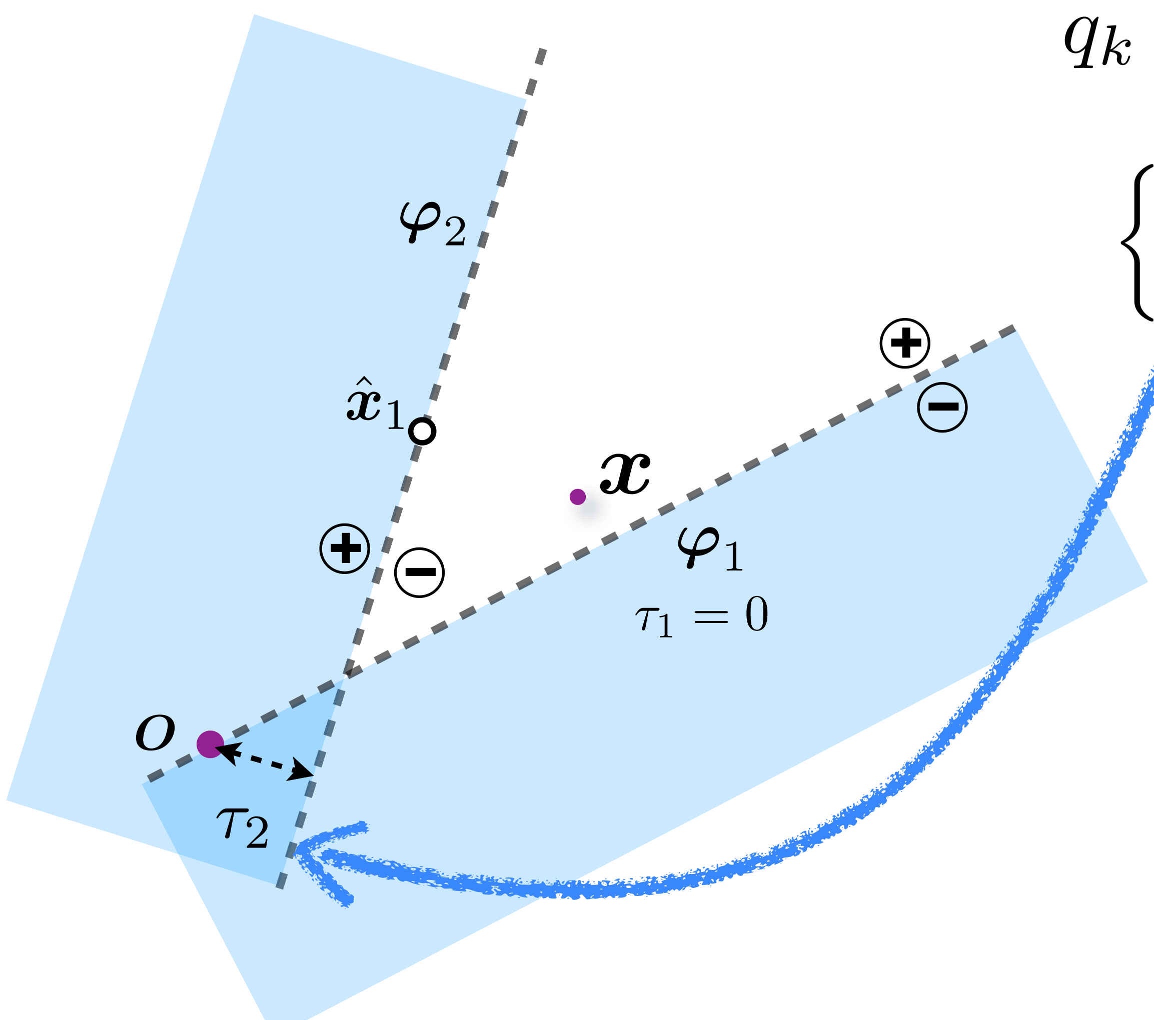

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

Adaptive 1-bit CS [Kamilov, Bourquard, Amini, Unser, 12]

Given a decoder Rec()

adapted from prev. meas.

$$q_k = \mathrm{sign}\left(\langle \boldsymbol{\varphi}_k, \boldsymbol{x}\rangle - \tau_k\right)$$

$$\begin{cases} \hat{\boldsymbol{x}}_k := \mathrm{Rec}(y_1, \cdots, y_k, \boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_k, \tau_1, \cdots, \tau_k) \\ \tau_{k+1} \text{ s.t. } \langle \boldsymbol{\varphi}_{k+1}, \hat{\boldsymbol{x}}_k\rangle - \tau_{k+1} = 0 \end{cases}$$

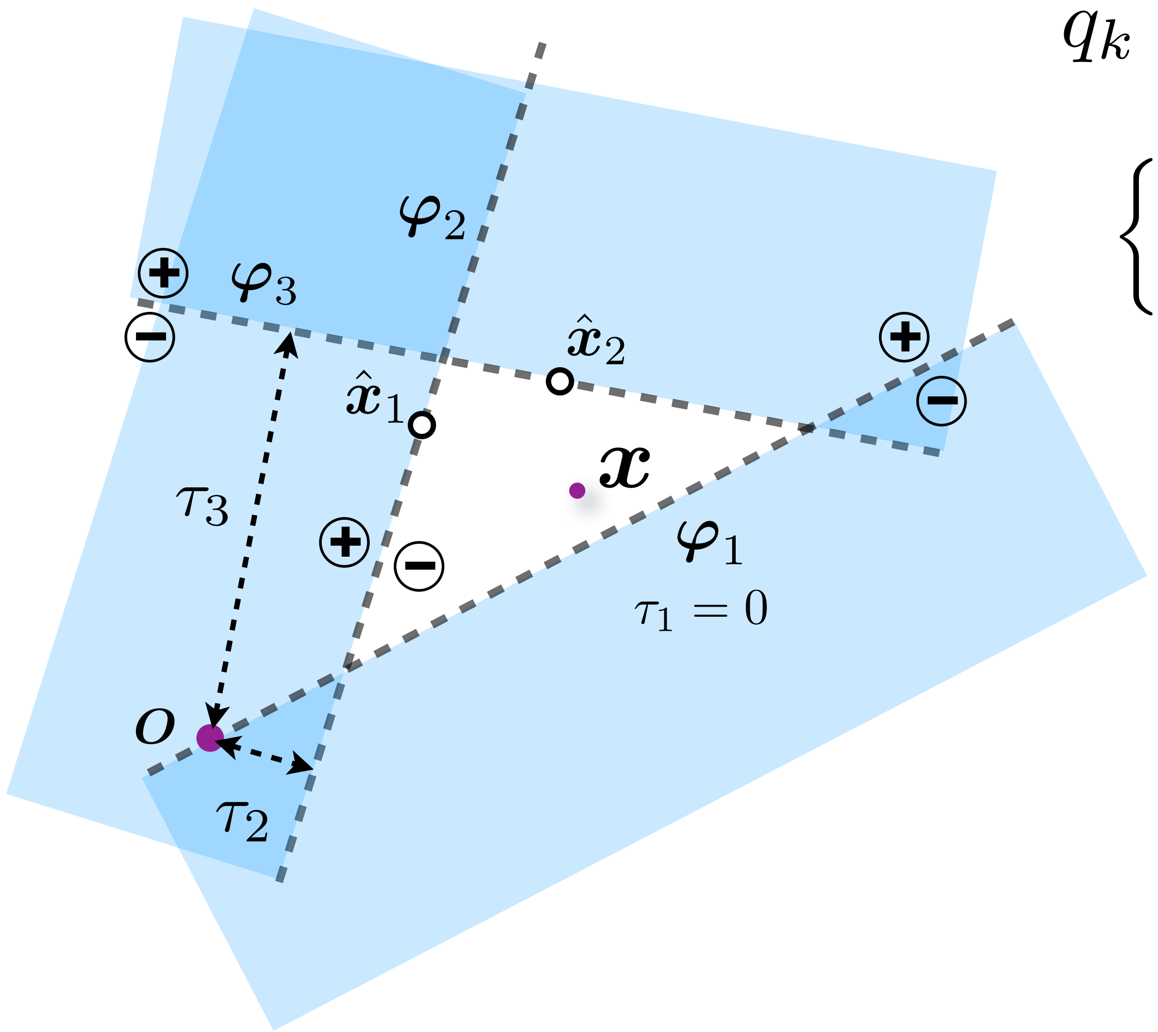

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

Adaptive 1-bit CS [Kamilov, Bourquard, Amini, Unser, 12]

Given a decoder Rec()

adapted from prev. meas.

$$q_k = \mathrm{sign}\left(\langle \boldsymbol{\varphi}_k, \boldsymbol{x}\rangle - \tau_k\right)$$

$$\begin{cases} \hat{\boldsymbol{x}}_k := \mathrm{Rec}(y_1, \cdots, y_k, \boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_k, \tau_1, \cdots, \tau_k) \\ \tau_{k+1} \text{ s.t. } \langle \boldsymbol{\varphi}_{k+1}, \hat{\boldsymbol{x}}_k\rangle - \tau_{k+1} = 0 \end{cases}$$

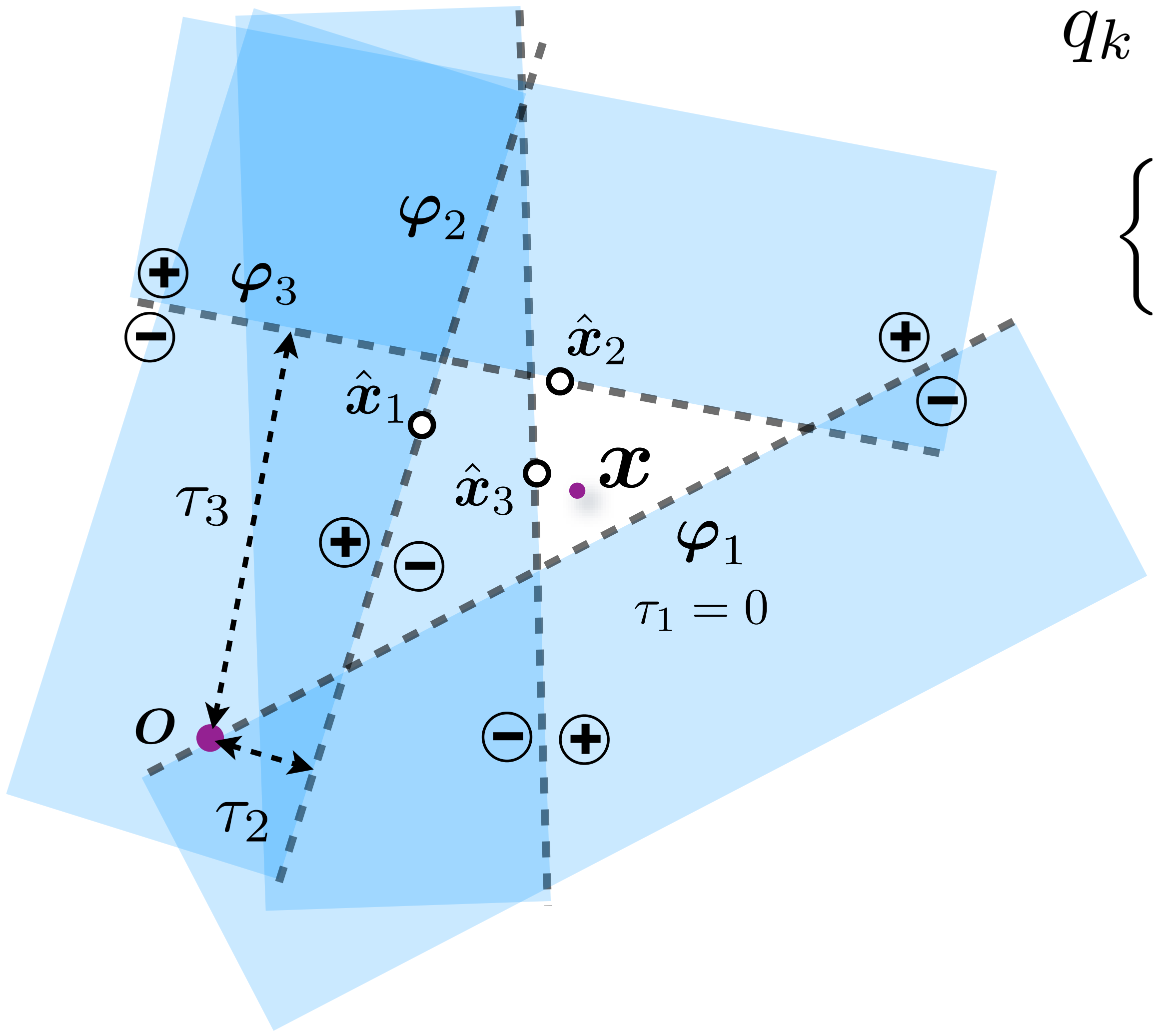

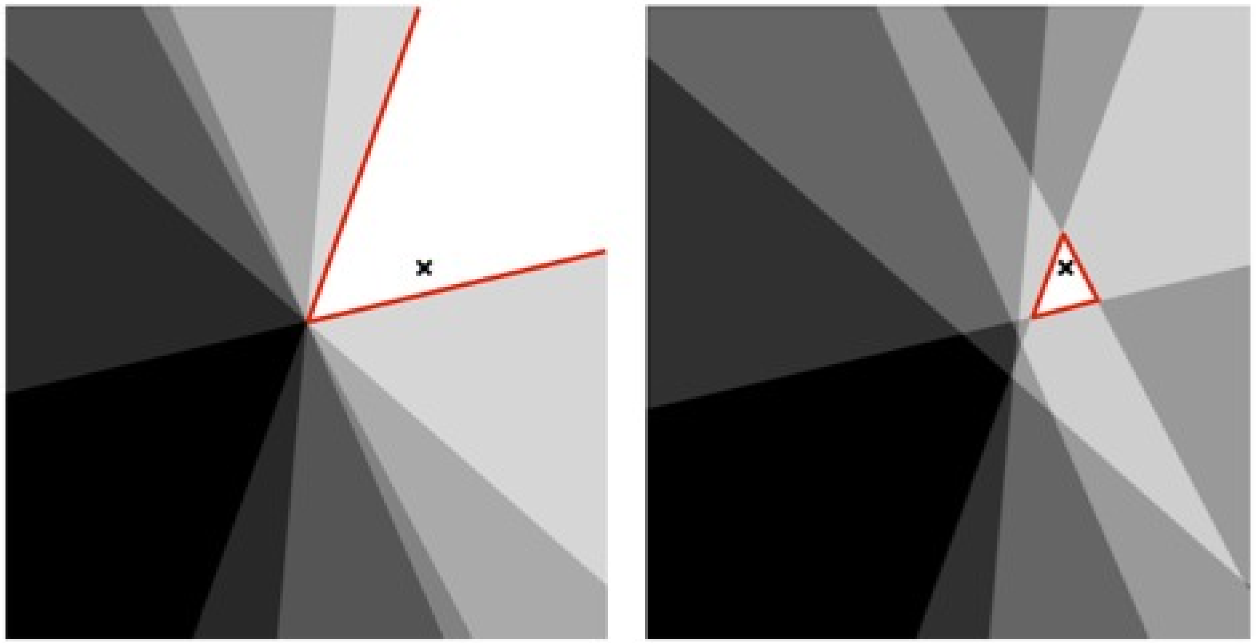

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

System view:

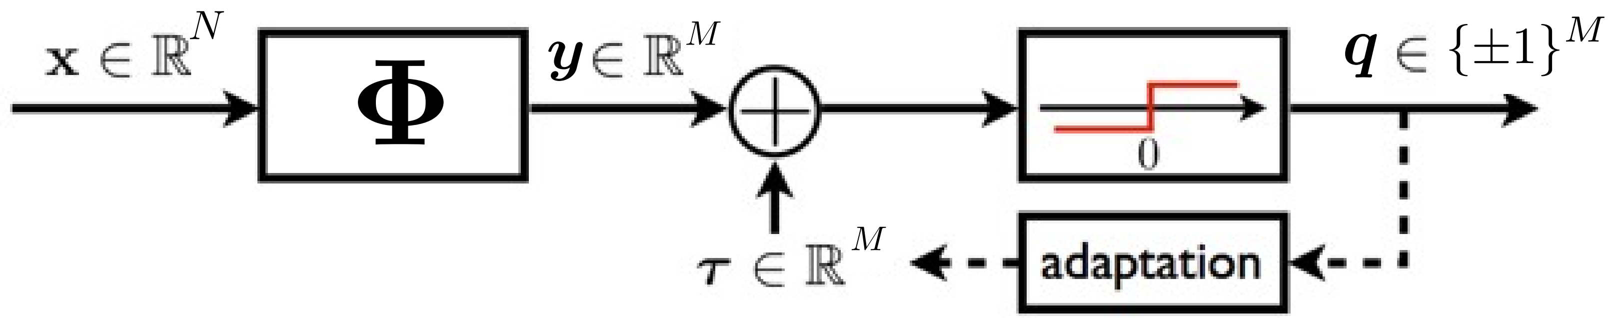

Kind of
$\Sigma\Delta$ loop

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

1-bit CS with adaptive thresholds

System view:

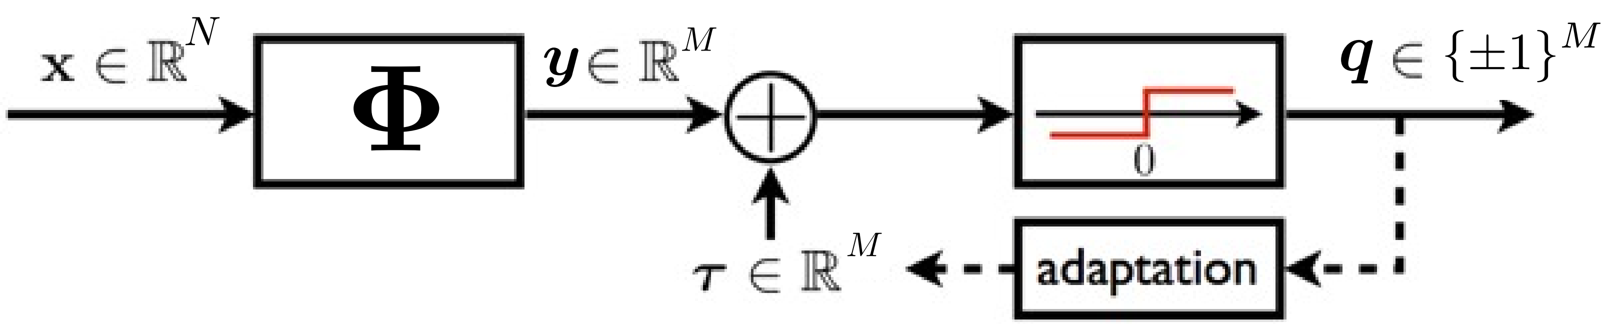

Kind of
$\Sigma\Delta$ loop

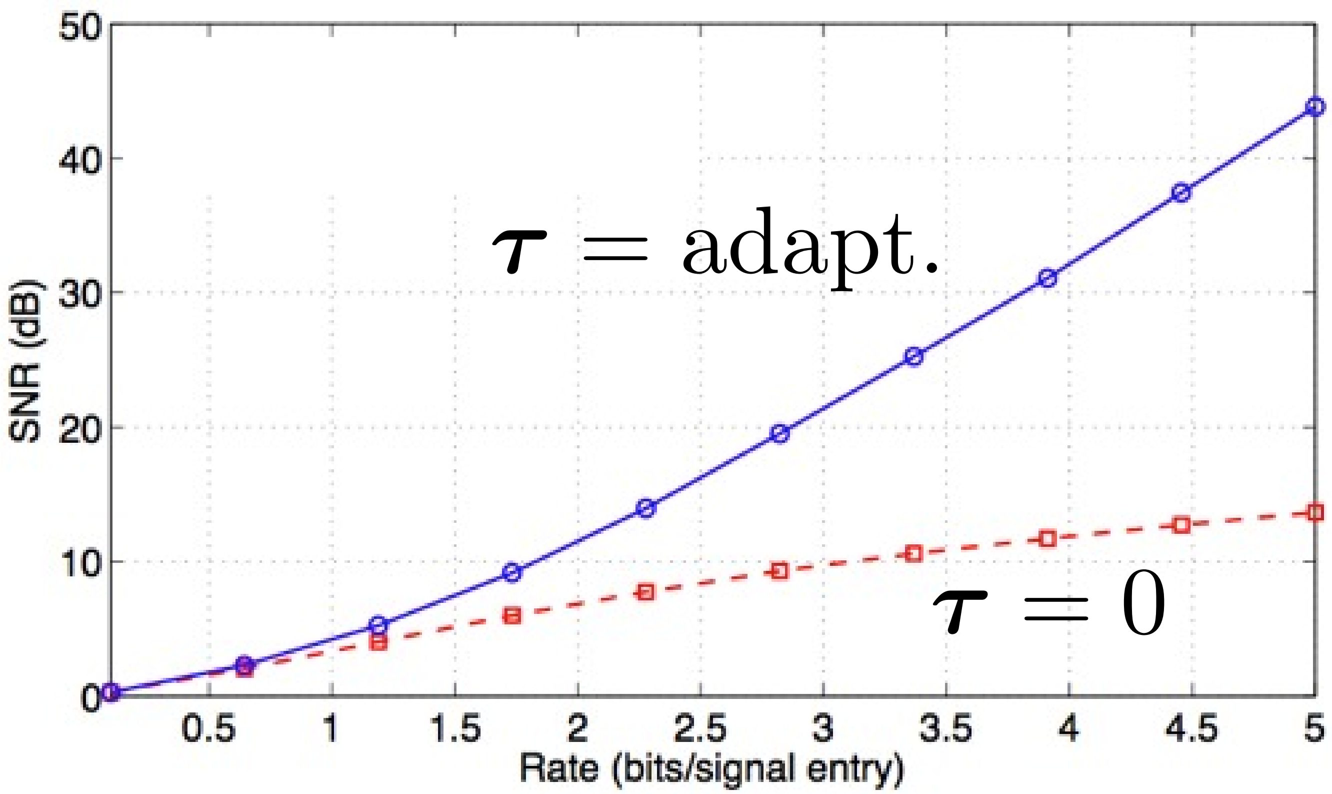

Rec() set to
Generalized
Approximate
Message
Passing

U.S. Kamilov, A. Bourquard, A. Amini, M. Unser,
"One-bit measurements with adaptive thresholds". Signal Processing Letters, IEEE, 19(10), 607-610.

Bridging 1-bit & B-bit CS?

Bridging 1-bit & B-bit CS?

- B-bit quantizer defined with thresholds:

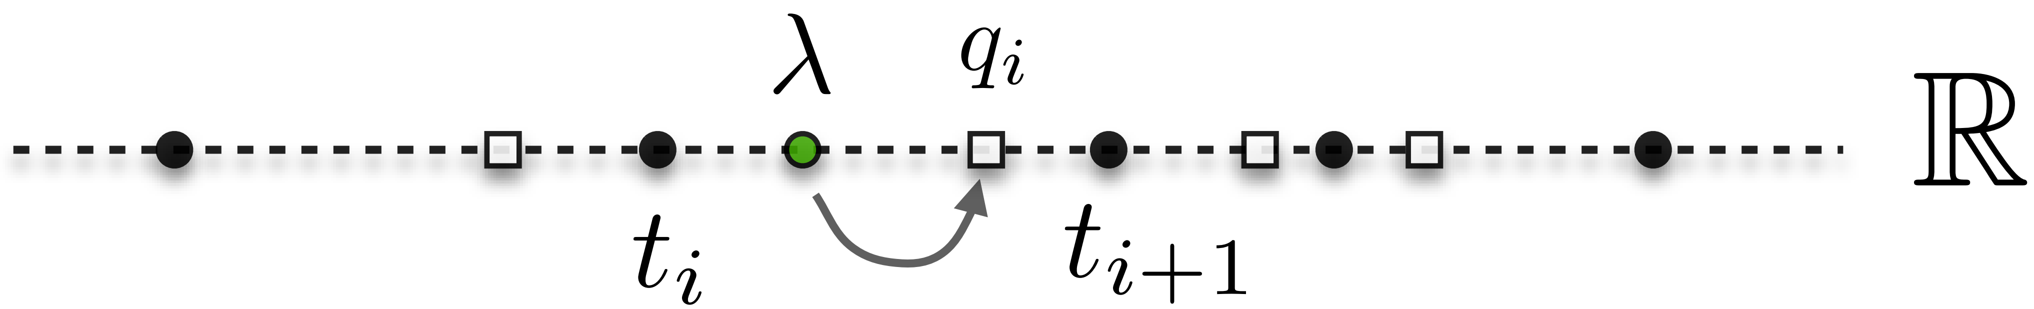

$$\lambda \in \mathcal{R}_i = [t_i, t_{i+1}) \Leftrightarrow \operatorname{sign}(\lambda - t_i) = +1 \ \& \ \operatorname{sign}(\lambda - t_{i+1}) = -1$$

- Can we combine multiple thresholds in 1-bit CS?

Bridging 1-bit & B-bit CS?

Given $\mathcal{T} = \{\tau_j\}$ and $\Omega = \{q_j\}$ ($|\mathcal{T}| = 2^B + 1 = |\Omega| + 1$), let's define

$$J(\nu, \lambda) = \sum_{j=2}^{2^B} w_j \left| \big(\mathrm{sign}\,(\lambda - \tau_j)\,(\nu - \tau_j)\big)_- \right|,$$

with $w_j = q_j - q_{j-1}$.

<u>Illustration:</u> $\lambda \in [\tau_{j-1}, \tau_j),\ \nu \in [\tau_j, \tau_{j+1})$

"delocalized"
BIHT ℓ_1-sided norm

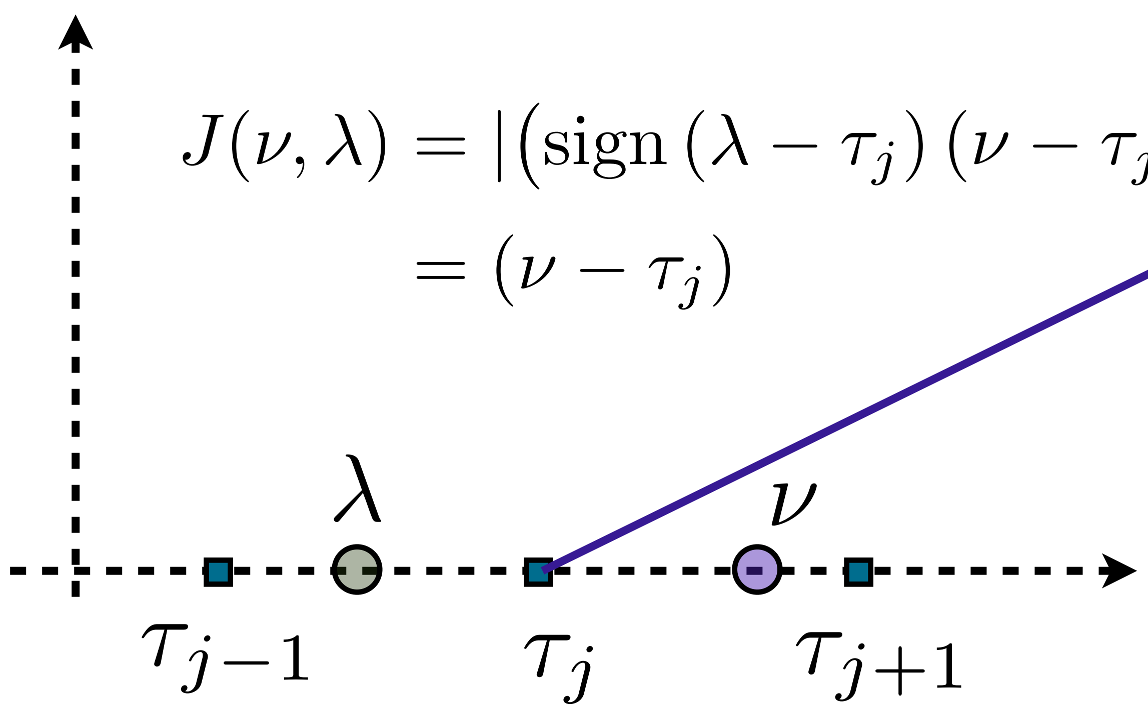

(for $w_j = 1$)

Bridging 1-bit & B-bit CS?

Given $\mathcal{T} = \{\tau_j\}$ and $\Omega = \{q_j\}$ ($|\mathcal{T}| = 2^B + 1 = |\Omega| + 1$), let's define

$$J(\nu, \lambda) = \sum_{j=2}^{2^B} w_j \left| \big(\mathrm{sign}\,(\lambda - \tau_j)\,(\nu - \tau_j)\big)_- \right|,$$

with $w_j = q_j - q_{j-1}$.

<u>Illustration:</u> $\lambda \in [\tau_{j-1}, \tau_j),\ \nu \in [\tau_{j+1}, \tau_{j+2})$

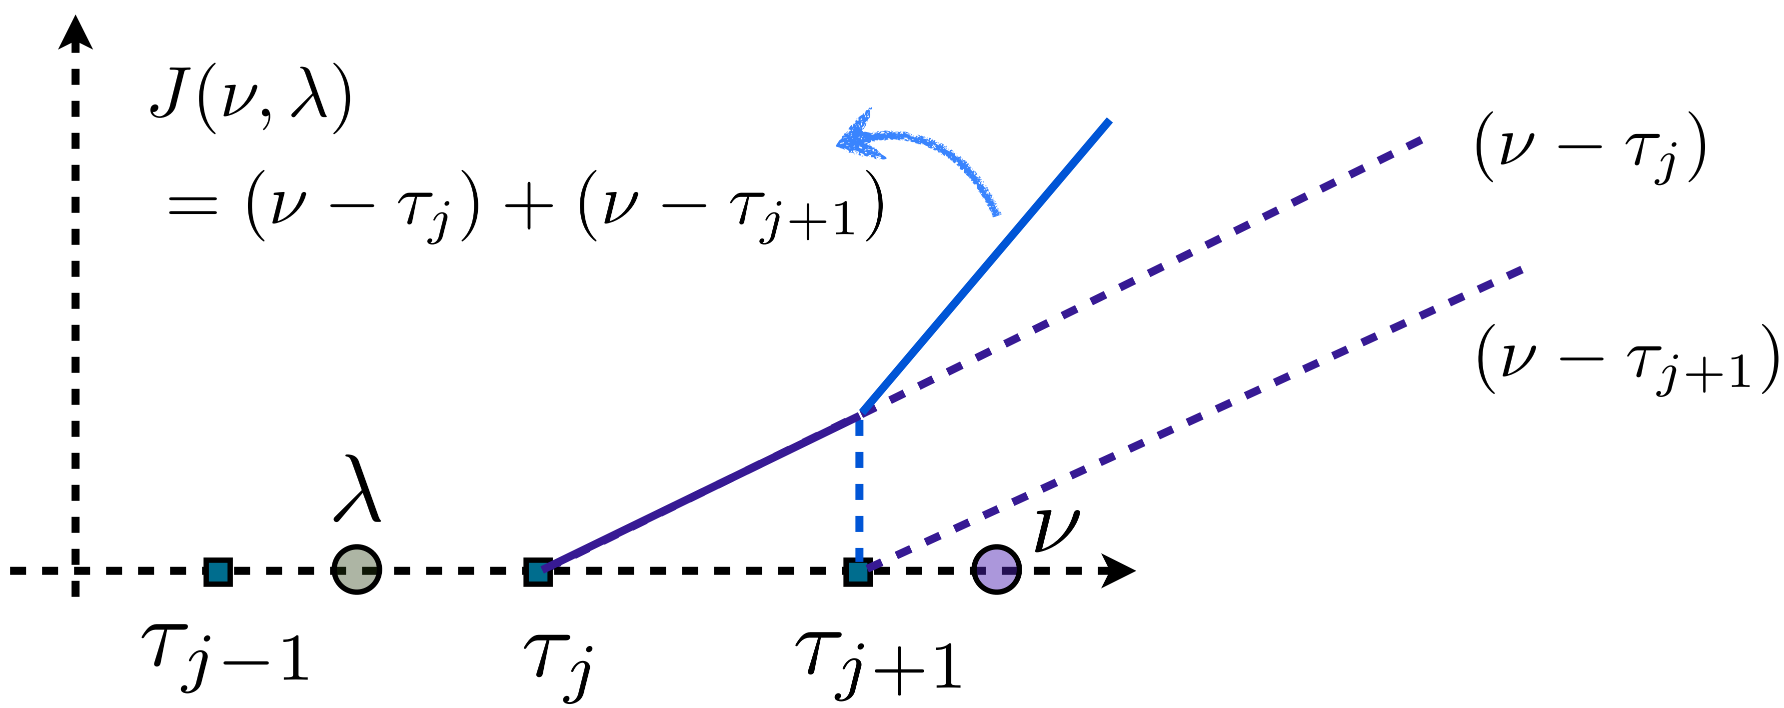

(for $w_j = 1$)

Bridging 1-bit & B-bit CS?

Given $\mathcal{T} = \{\tau_j\}$ and $\Omega = \{q_j\}$ ($|\mathcal{T}| = 2^B + 1 = |\Omega| + 1$), let's define

$$J(\nu, \lambda) = \sum_{j=2}^{2^B} w_j \left|\big(\mathrm{sign}\,(\lambda - \tau_j)\,(\nu - \tau_j)\big)_-\right|,$$

with $w_j = q_j - q_{j-1}$.

Illustration:

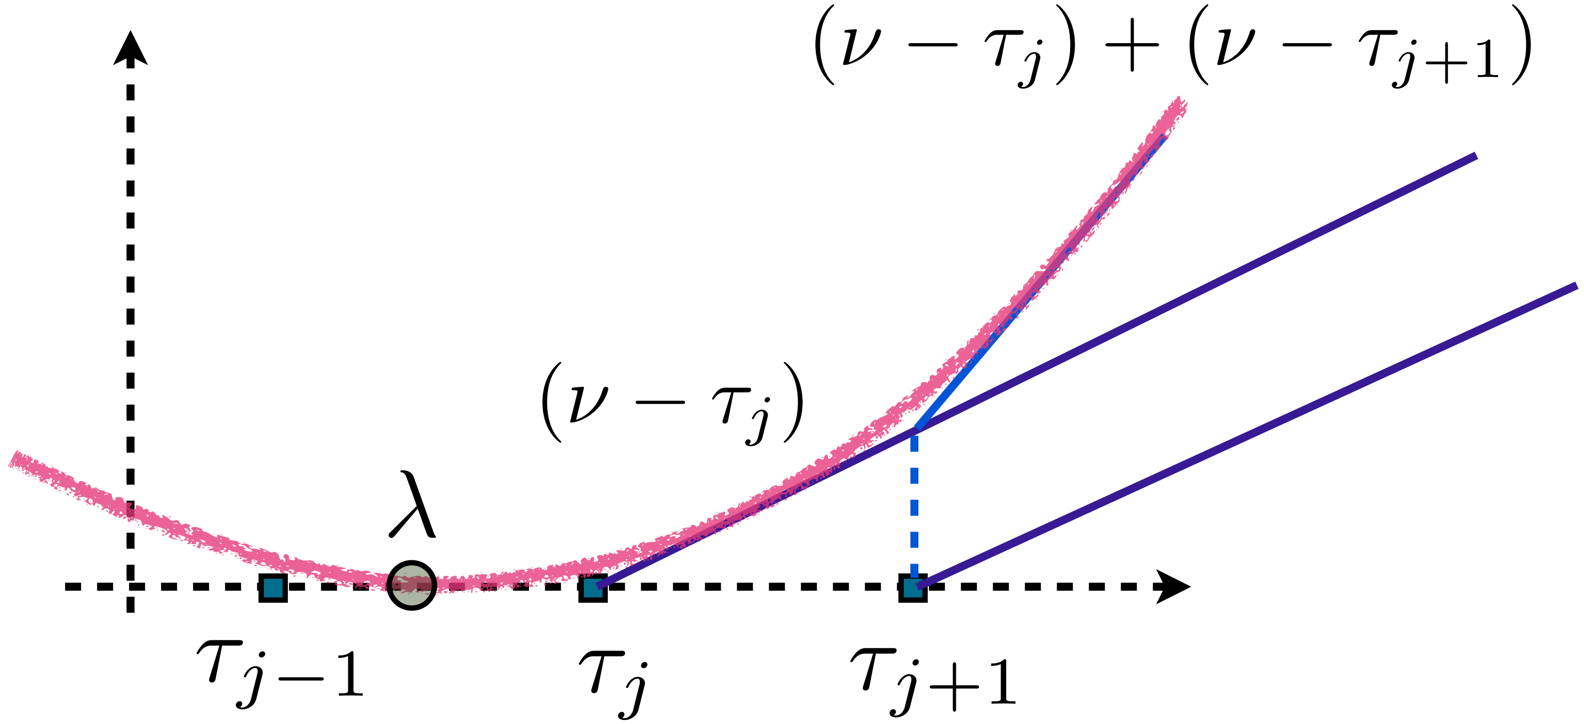

(for $w_j = 1$)

Bridging 1-bit & B-bit CS?

Given $\mathcal{T} = \{\tau_j\}$ and $\Omega = \{q_j\}$ ($|\mathcal{T}| = 2^B + 1 = |\Omega| + 1$), let's define

$$J(\nu, \lambda) = \sum_{j=2}^{2^B} w_j \left| \left(\mathrm{sign}\,(\lambda - \tau_j)\,(\nu - \tau_j) \right)_- \right|,$$

with $w_j = q_j - q_{j-1}$.

<u>Illustration:</u> *more bins*

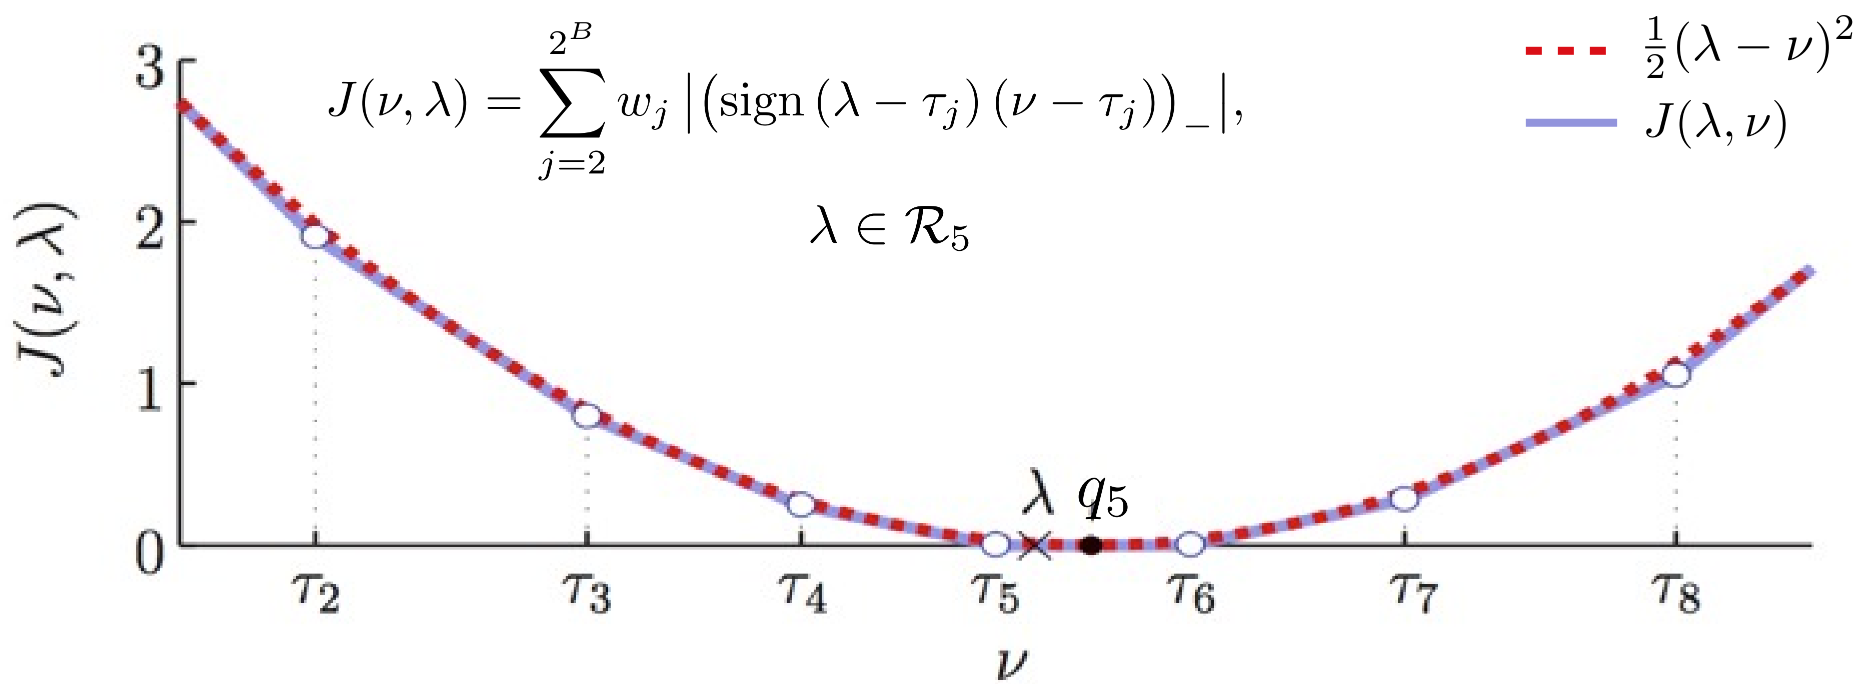

Bridging 1-bit & B-bit CS?

Given $\mathcal{T} = \{\tau_j\}$ and $\Omega = \{q_j\}$ ($|\mathcal{T}| = 2^B + 1 = |\Omega| + 1$), let's define

$$J(\nu, \lambda) = \sum_{j=2}^{2^B} w_j \left| \big(\mathrm{sign}\,(\lambda - \tau_j)\,(\nu - \tau_j)\big)_- \right|,$$

with $w_j = q_j - q_{j-1}$.

For $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^M$: $\mathcal{J}(\boldsymbol{u}, \boldsymbol{v}) := \sum_{k=1}^{M} J(u_k, v_k)$

Remarks:

- J is convex in ν
- For $B = 1$ ($j = 2$ only):

  $\mathcal{J}(\boldsymbol{u}, \boldsymbol{v}) \propto \|(\mathrm{sign}\,(\boldsymbol{v}) \odot \boldsymbol{u})_-\|_1 \to \ell_1$-sided 1-bit energy
- For $B \gg 1$:

  $J(\nu, \lambda) \to \frac{1}{2}(\nu - \lambda)^2$ and $\mathcal{J}(\boldsymbol{u}, \boldsymbol{v}) \to \frac{1}{2}\|\boldsymbol{u} - \boldsymbol{v}\|^2$ (quadratic energy)

Bridging 1-bit & B-bit CS?

- Let's define an *inconsistency* energy:

$$\mathcal{E}_B(\boldsymbol{u}) := \mathcal{J}(\boldsymbol{\Phi u}, \boldsymbol{q}) \text{ with } \boldsymbol{q} = \mathcal{Q}_B[\boldsymbol{\Phi x}] \text{ and } \mathcal{E}\_B(\boldsymbol{x}) = 0$$

- Idea: Minimize it in Σ_K (as for Iterative Hard Thresholding)

[Blumensath, Davies, 08]

$$\min_{\boldsymbol{u} \in \mathbb{R}^N} \mathcal{E}_B(\boldsymbol{u}) \text{ s.t. } \|\boldsymbol{u}\|_0 \leqslant K,$$

T. Blumensath, M.E. Davies, "Iterative thresholding for sparse approximations". *Journal of Fourier Analysis and Applications*, *14*(5-6), 629-654. (2008).

Bridging 1-bit & B-bit CS?

- Let's define an *inconsistency* energy:

$$\mathcal{E}_B(\boldsymbol{u}) := \mathcal{J}(\boldsymbol{\Phi u}, \boldsymbol{q}) \text{ with } \boldsymbol{q} = \mathcal{Q}_B[\boldsymbol{\Phi x}] \text{ and } \mathcal{E}\_B(\boldsymbol{x}) = 0$$

- Idea: Minimize it in Σ_K (as for Iterative Hard Thresholding)

[Blumensath, Davies, 08]

$$\min_{\boldsymbol{u}\in\mathbb{R}^N} \mathcal{E}_B(\boldsymbol{u}) \text{ s.t. } \|\boldsymbol{u}\|_0 \leqslant K,$$

- NP Hard but greedy solution (as for IHT):

$$\boldsymbol{x}^{(n+1)} = \mathcal{H}_K[\boldsymbol{x}^{(n)} - \mu\, \partial\, \mathcal{E}_B(\boldsymbol{x}^{(n)})] \text{ and } \boldsymbol{x}^{(0)} = 0.$$

(sub) gradient

$$\boldsymbol{\Phi}^*(\text{sign}\,(\boldsymbol{\Phi u}) - \text{sign}\,(\boldsymbol{\Phi x})) \xleftarrow{B=1} \partial\, \mathcal{E}_B(\boldsymbol{u}) = \boldsymbol{\Phi}^*(\mathcal{Q}_B(\boldsymbol{\Phi u}) - \boldsymbol{q}) \xrightarrow{B \gg 1} \boldsymbol{\Phi}^*(\boldsymbol{\Phi u} - \boldsymbol{q})$$

BIHT! Quantized IHT (QIHT) IHT!

T. Blumensath, M.E. Davies, "Iterative thresholding for sparse approximations". *Journal of Fourier Analysis and Applications*, *14*(5-6), 629-654. (2008).
LJ, K. Degraux, C. De Vleeschouwer, "Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing", SAMPTA2013

Bridging 1-bit & B-bit CS?

$N = 1024,\ K = 16,\ R = BM \in \{64, 128, \cdots, 1280\}$, 100 trials (+ Lloyd-Max Gauss. Q.)

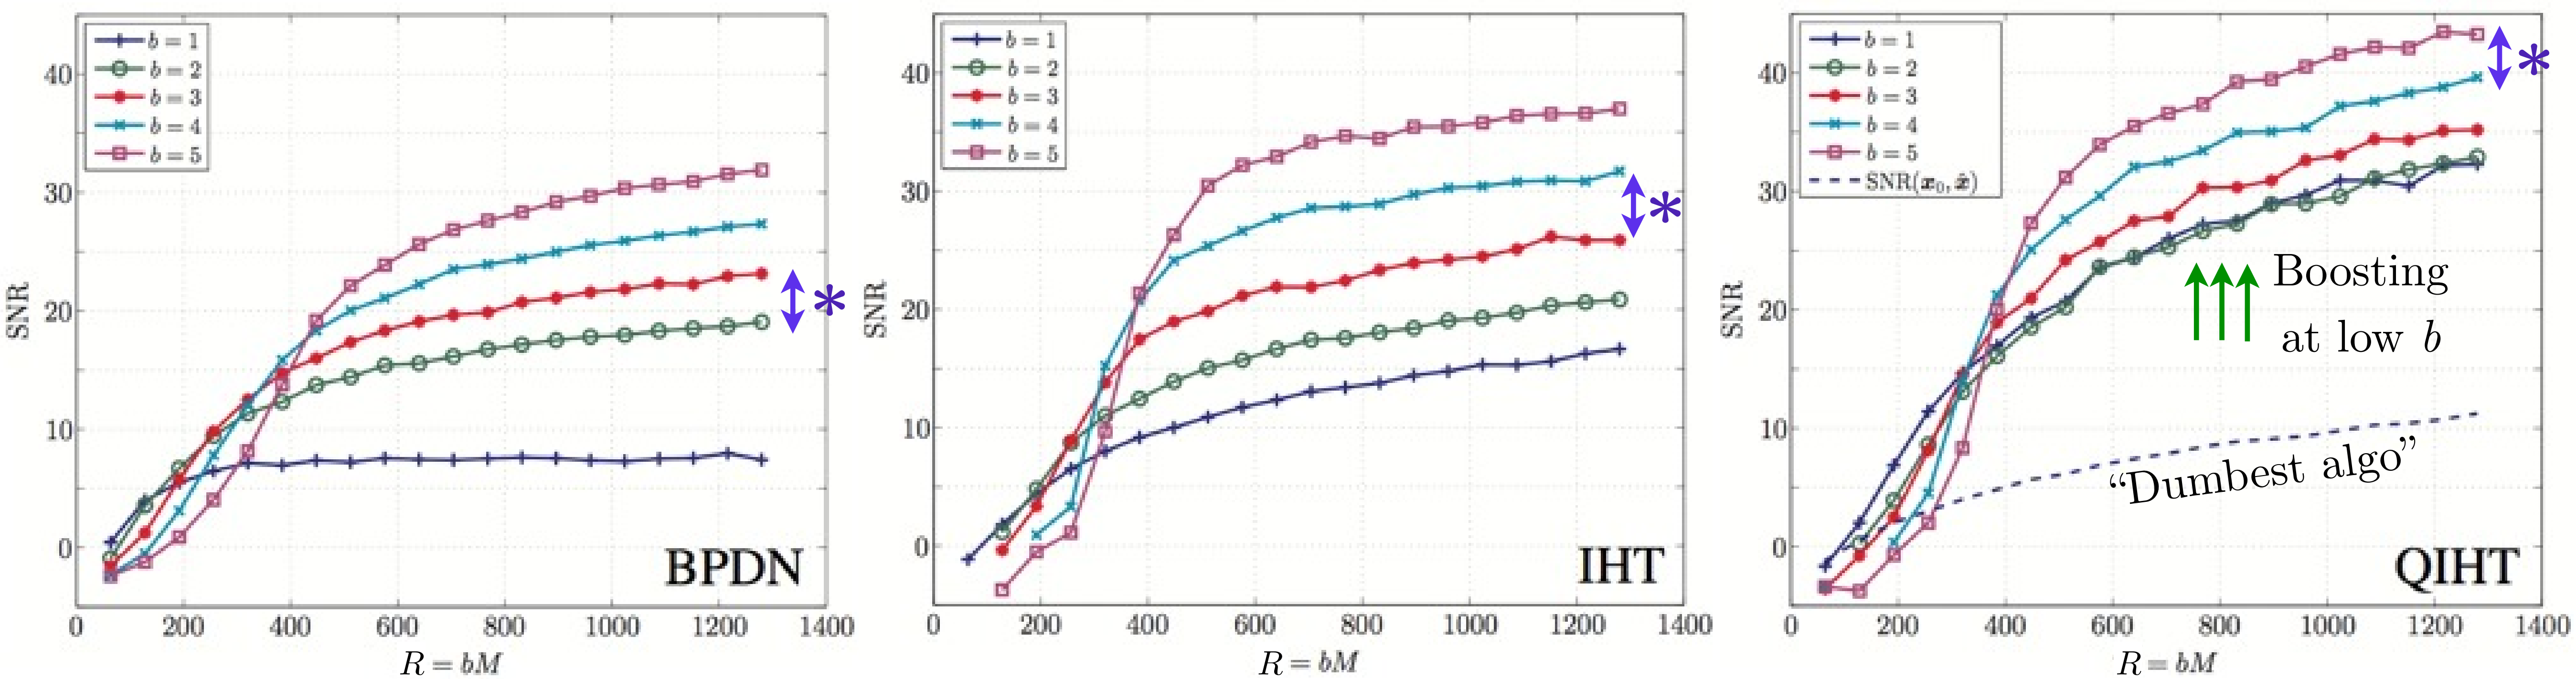

R: total bit budget (BM)

*: almost "6dB per bit" gain

$\mu = \frac{1}{M}(1 - \sqrt{2K/M})$

Adjusted by limit case analysis: BIHT and IHT

Note: entropy could be computed instead of B (*e.g.*, for further efficient coding)

LJ, K. Degraux, C. De Vleeschouwer, "Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing", SAMPTA2013

Bridging 1-bit & B-bit CS?

$$N = 1024,\ K = 16,\ R = BM \in \{64, 128, \cdots, 1280\}, 100 \text{ trials}$$

Interesting transition at $R_0 \simeq 375$

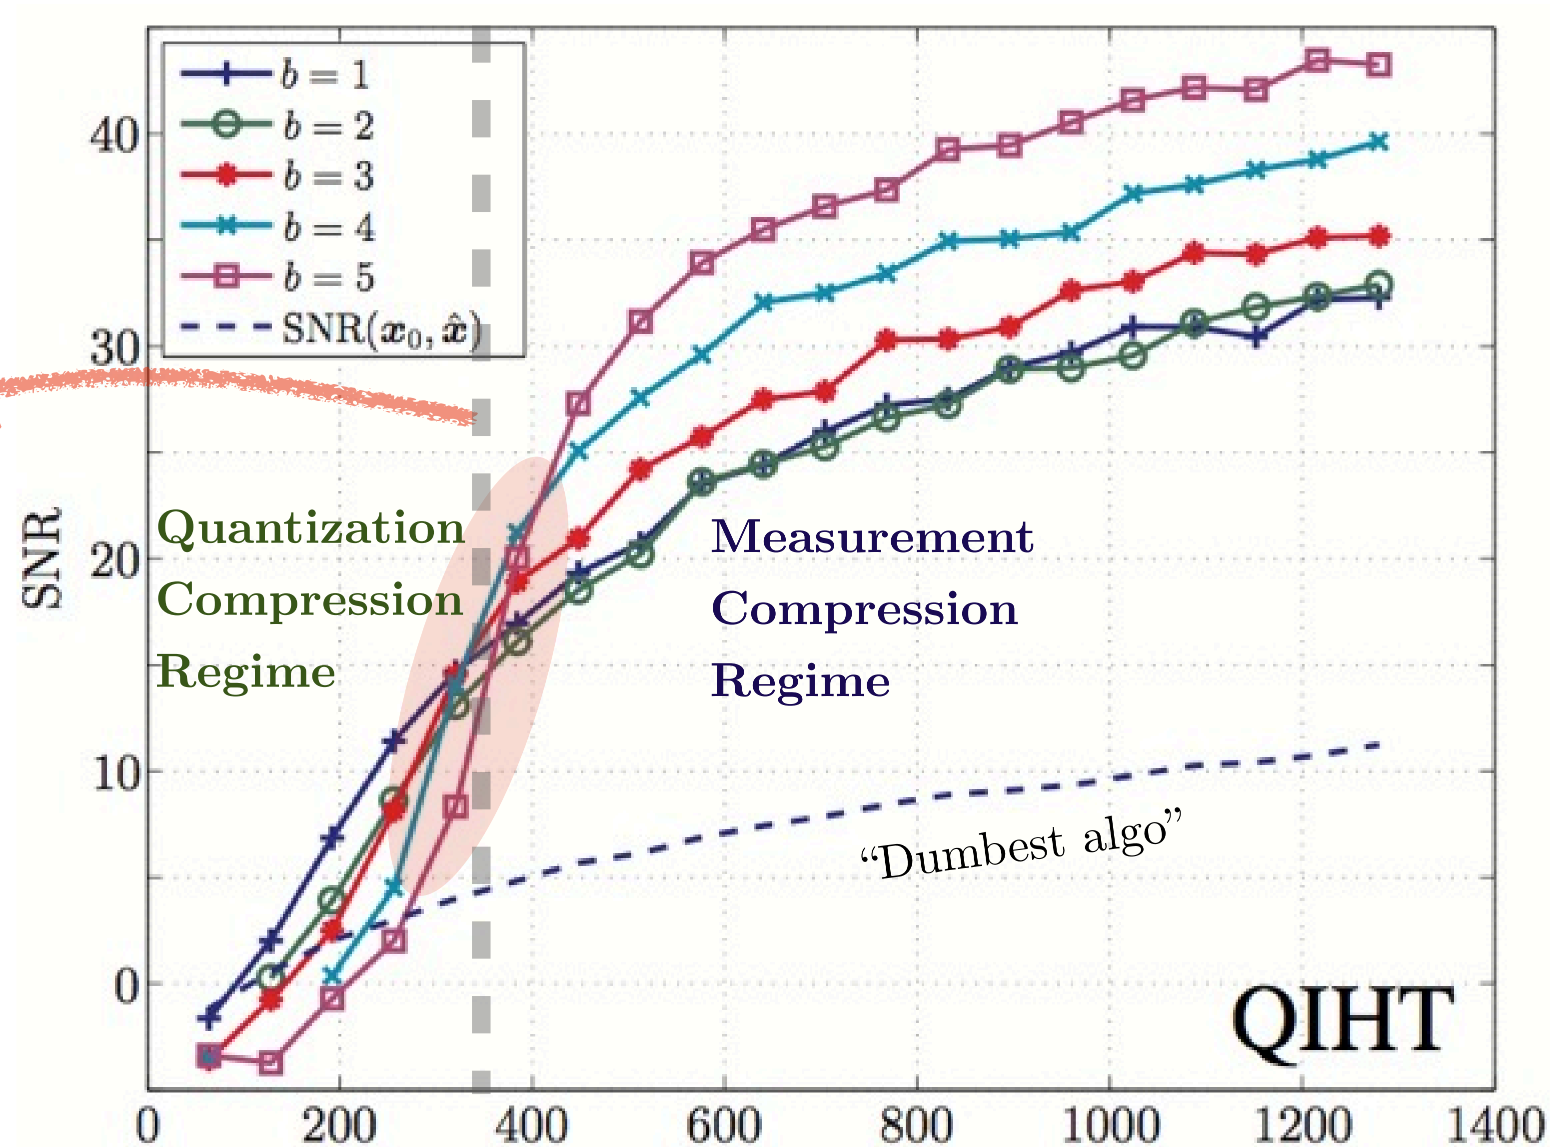

$R\!=\!BM$

"**Regime Change**?" [Laska, Baraniuk, 12]

R_0 could increase with input noise power.

J. N. Laska, R. G. Baraniuk, 'Regime change: Bit-depth versus measurement-rate in compressive sensing", *Signal Processing, IEEE Transactions on*, *60*(7), 3496-3505. (2012)

Further Reading

- T. Blumensath, M.E. Davies, “Iterative thresholding for sparse approximations”. *Journal of Fourier Analysis and Applications*, 14(5-6), pp. 629-654, 2008
- P. T. Boufounos and R. G. Baraniuk, “1-Bit compressive sensing,” *Proc. Conf. Inform. Science and Systems (CISS)*, Princeton, NJ, March 19-21, 2008.
- Boufounos, P. T. (2009, November). “Greedy sparse signal reconstruction from sign measurements”. In *Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers*, 2009
- Y. Plan, R. Vershynin, “Dimension reduction by random hyperplane tessellations”, arXiv:1111.4452, 2011.
- Y. Plan, R. Vershynin, “Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach”, *IEEE Trans. Info. Theory*, arXiv:1202.1212, 2012.
- J. N. Laska, R. G. Baraniuk, ‘Regime change: Bit-depth versus measurement-rate in compressive sensing”, *IEEE Trans. Signal Processing*, 60(7), pp. 3496-3505, 2012.
- U.S. Kamilov, A. Bourquard, A. Amini, M. Unser, “One-bit measurements with adaptive thresholds”. *IEEE Signal Processing Letters*, 19(10), pp. 607-610, 2012
- L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, “Robust 1-Bit Compressive Sensing via Binary Stable Embeddings of Sparse Vectors,” *IEEE Trans. Info. Theory*, 59(4), 2013.
- L. Jacques, K. Degraux, C. De Vleeschouwer, “Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing", SAMPTA 2013, to appear.

Thank you!