

Time for dithering! Quantized random embeddings with RIP random matrices

Chunlei Xu, Vincent Schellekens, Valerio Cambareri, Laurent Jacques



ICTEAM/ELEN, University of Louvain (UCL), Belgium



Outline

- ▶ (Brief introduction to CS)
- ▶ Quantizing Compressive Sensing theory
- ▶ Properties of dithered, quantized random projections
- ▶ Classification in a quantized world
- ▶ Quantized Random Sketching of Datasets

Brief introduction to CS

Compressed Sensing...

... in a nutshell:

Generalize Dirac/Nyquist sampling:

1°) ask *few* (**linear**) *questions*

about your *informative* signal

2°) and recover it *differently* (**non-linearly**)”

Compressed Sensing...

... *in a nutshell*:

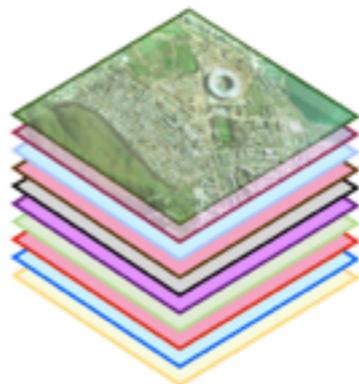
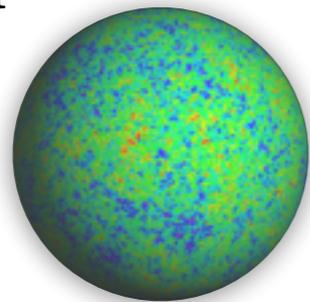
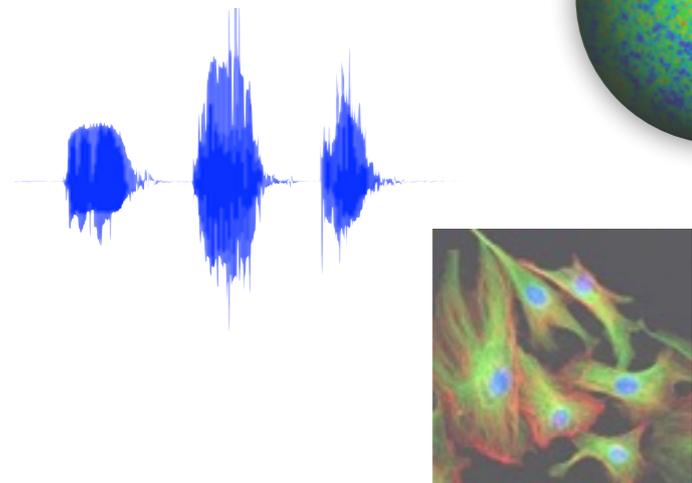
Generalize Dirac/Nyquist sampling:

1° ask *few* (**linear**) *questions*

about your *informative* signal

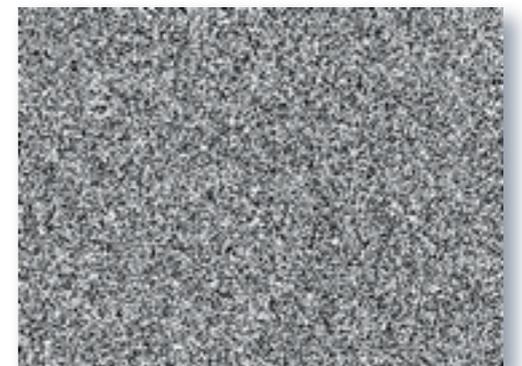
2° and recover it *differently* (**non-linearly**)”

Signals composed
of structures

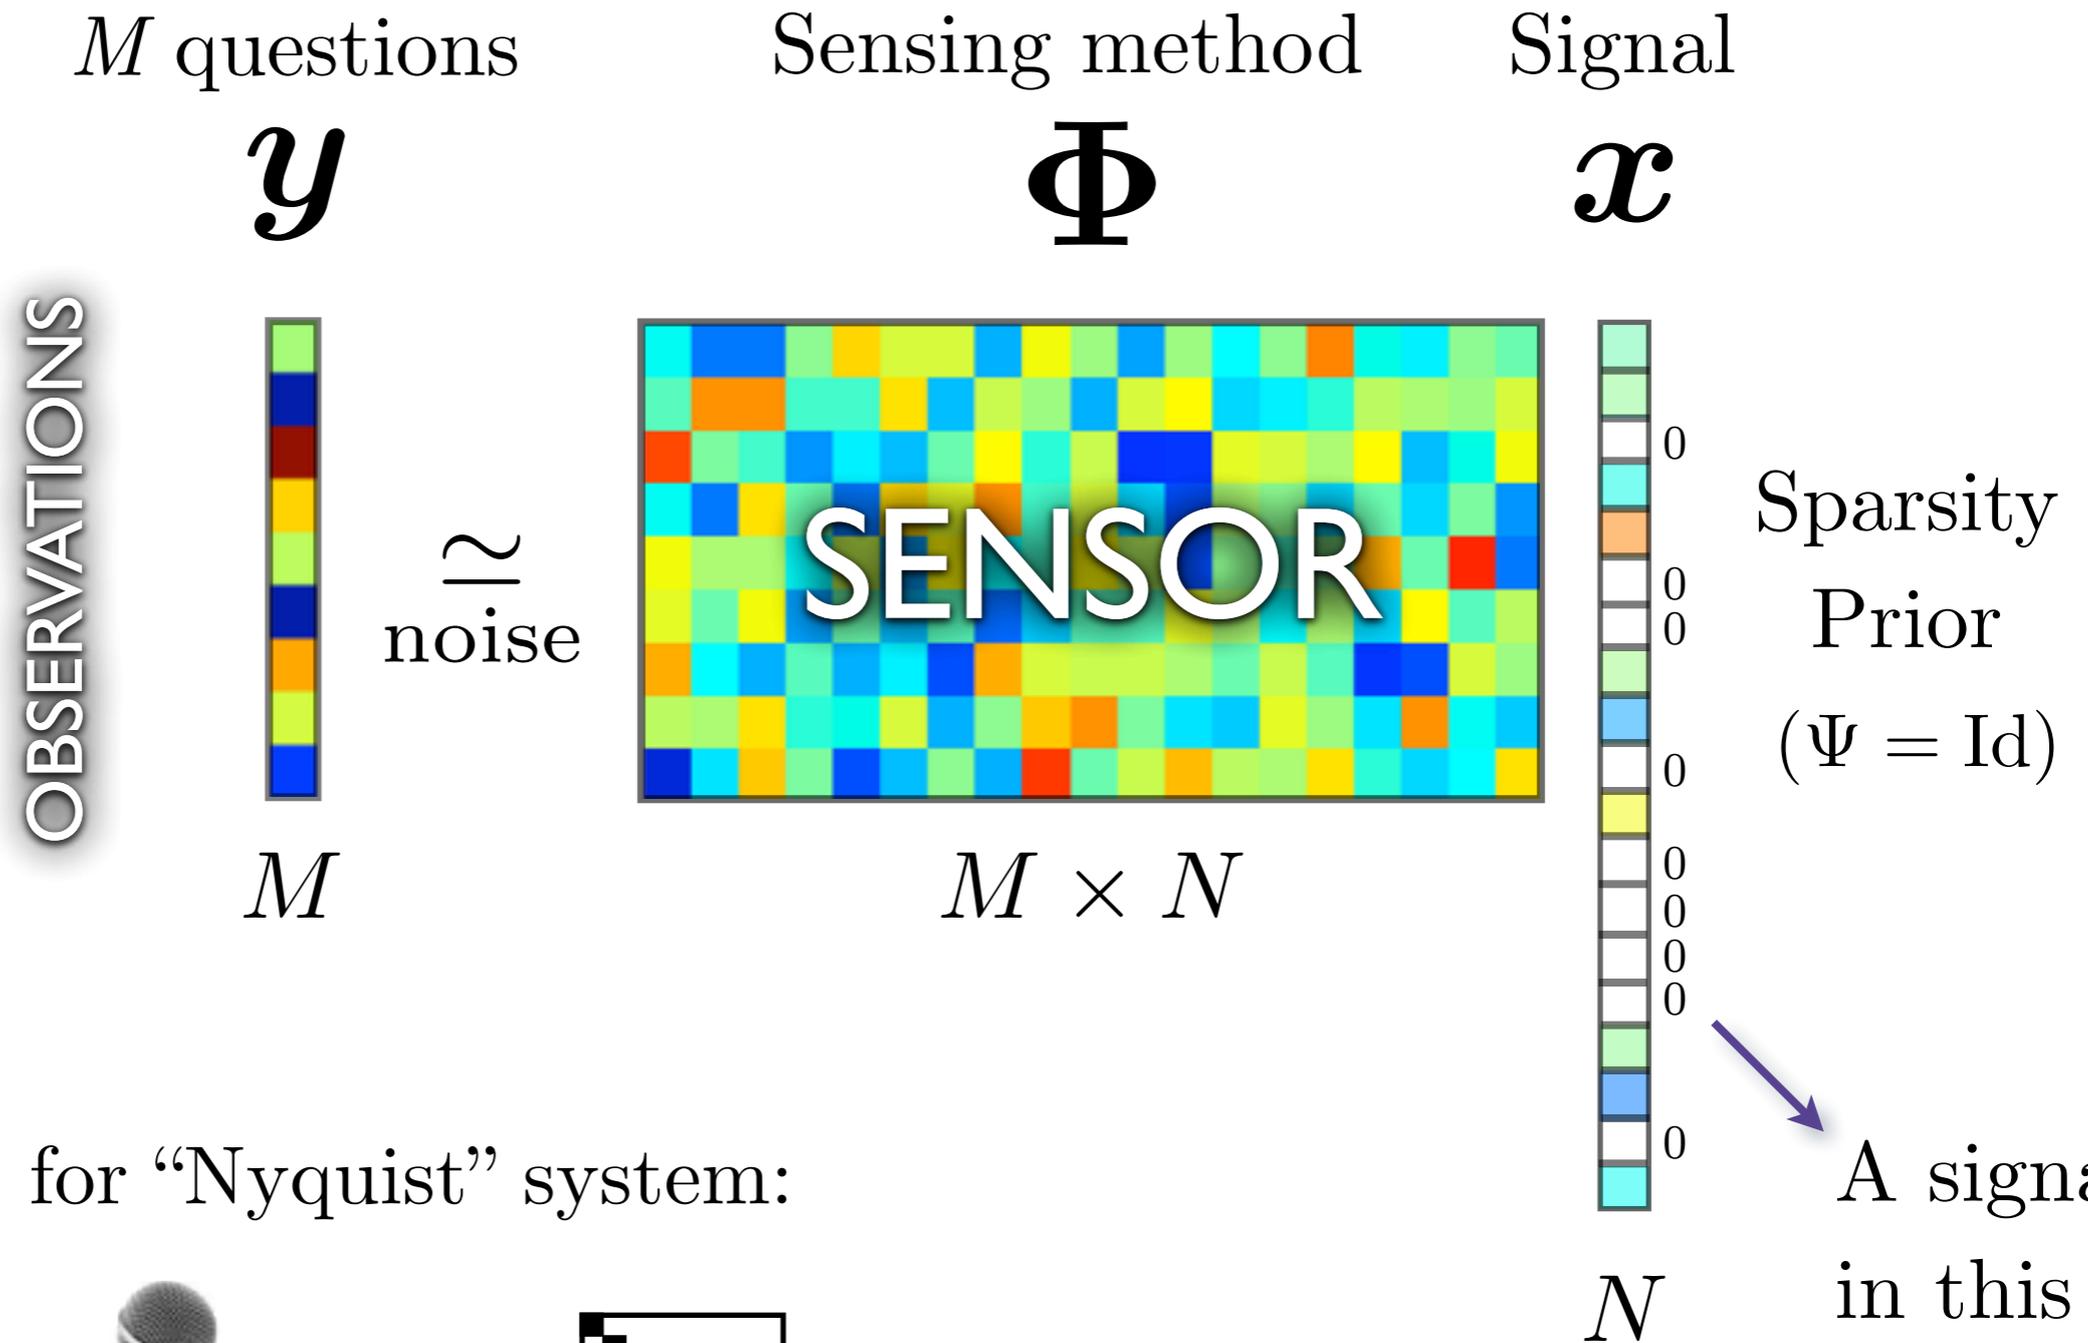


\exists a prior!

\neq



Compressed Sensing...



Reminder: for “Nyquist” system:



A signal
in this
discrete
world

Compressed Sensing...

M questions

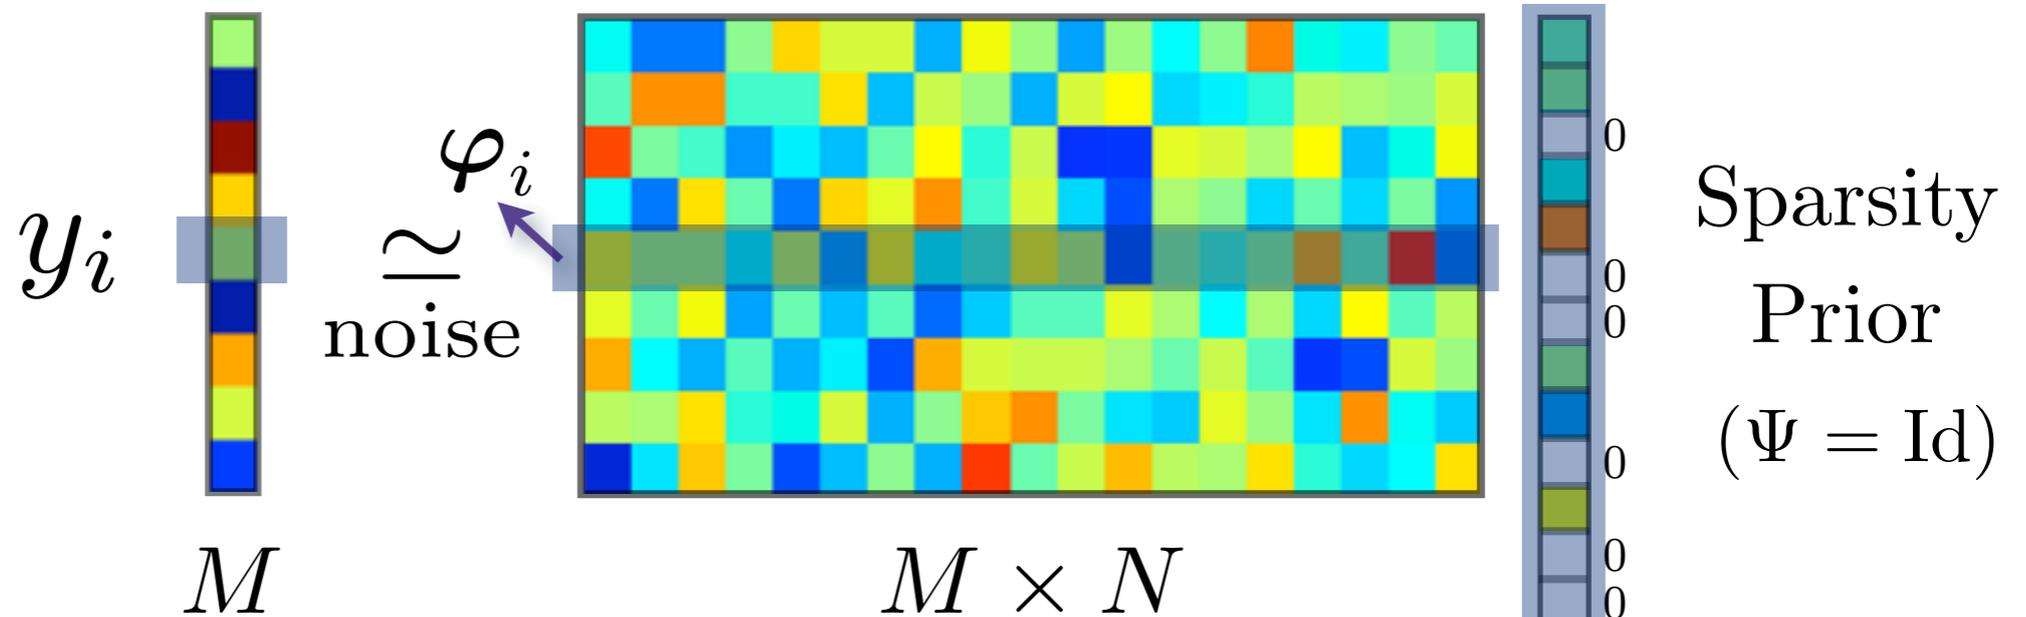
Sensing method

Signal

y

Φ

x



Generalized Linear Sensing!

$$y_i \simeq \langle \varphi_i, x \rangle = \varphi_i^T x$$

$$1 \leq i \leq M$$

e.g., to be realized
optically/analogically

A signal
in this
discrete
world

Compressed Sensing...

M questions

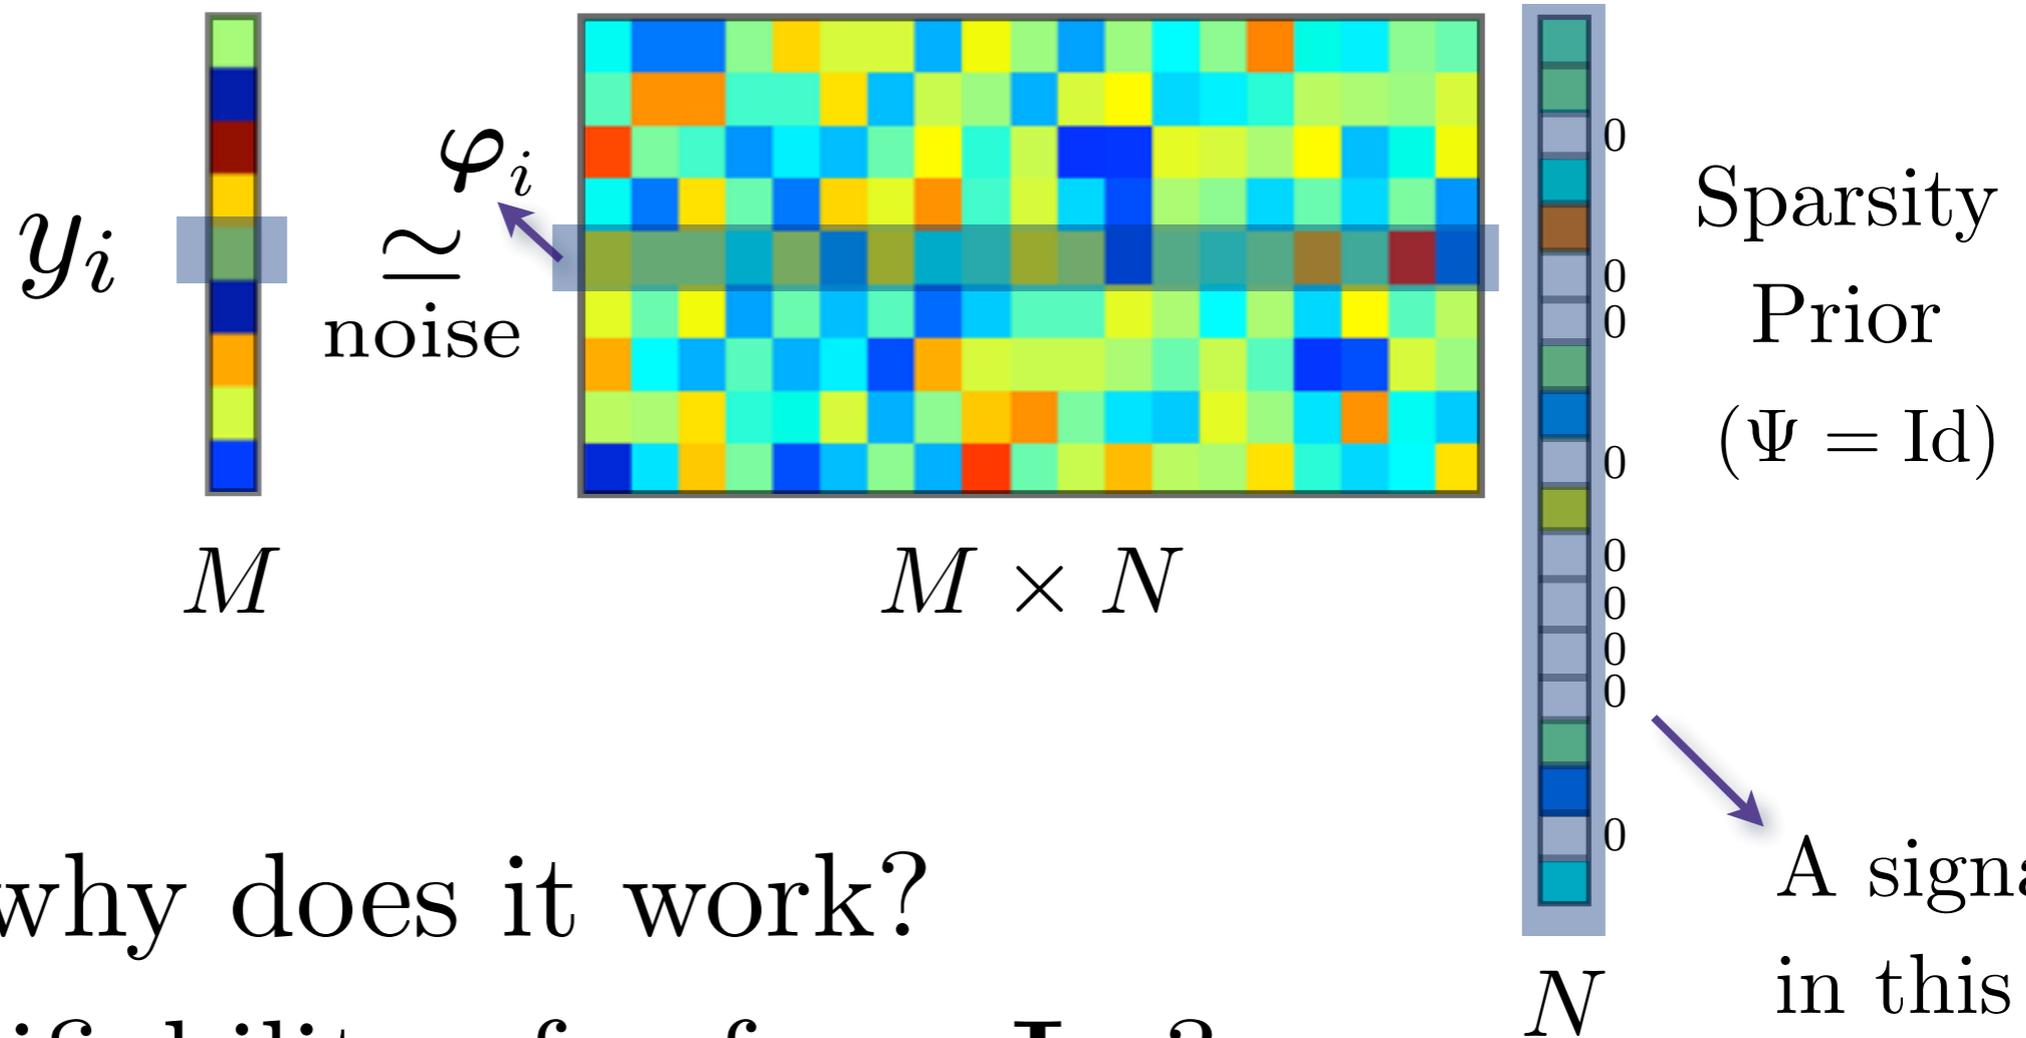
y

Sensing method

Φ

Signal

x



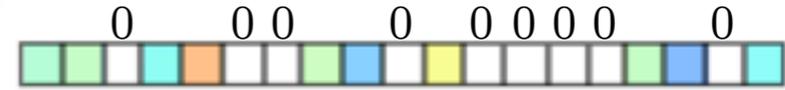
But why does it work?

Identifiability of x from Φx ?

\Rightarrow use the prior!

A signal
in this
discrete
world

Compressed Sensing...

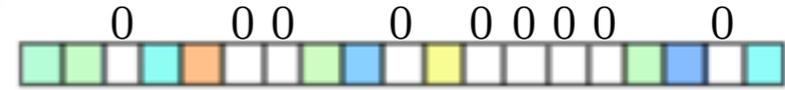


Two K -sparse signals $\mathbf{x}, \mathbf{x}' \in \Sigma_K := \{\mathbf{u} : \|\mathbf{u}\|_0 := |\text{supp } \mathbf{u}| \leq K\}$

For many random constructions of Φ (e.g., Gaussian, Bernoulli, structured) and “ $M \gtrsim K \log(N/K)$ ”, with high probability,

Geometry of $\Phi(\Sigma_K)$
 \approx Geometry of Σ_K

Compressed Sensing...



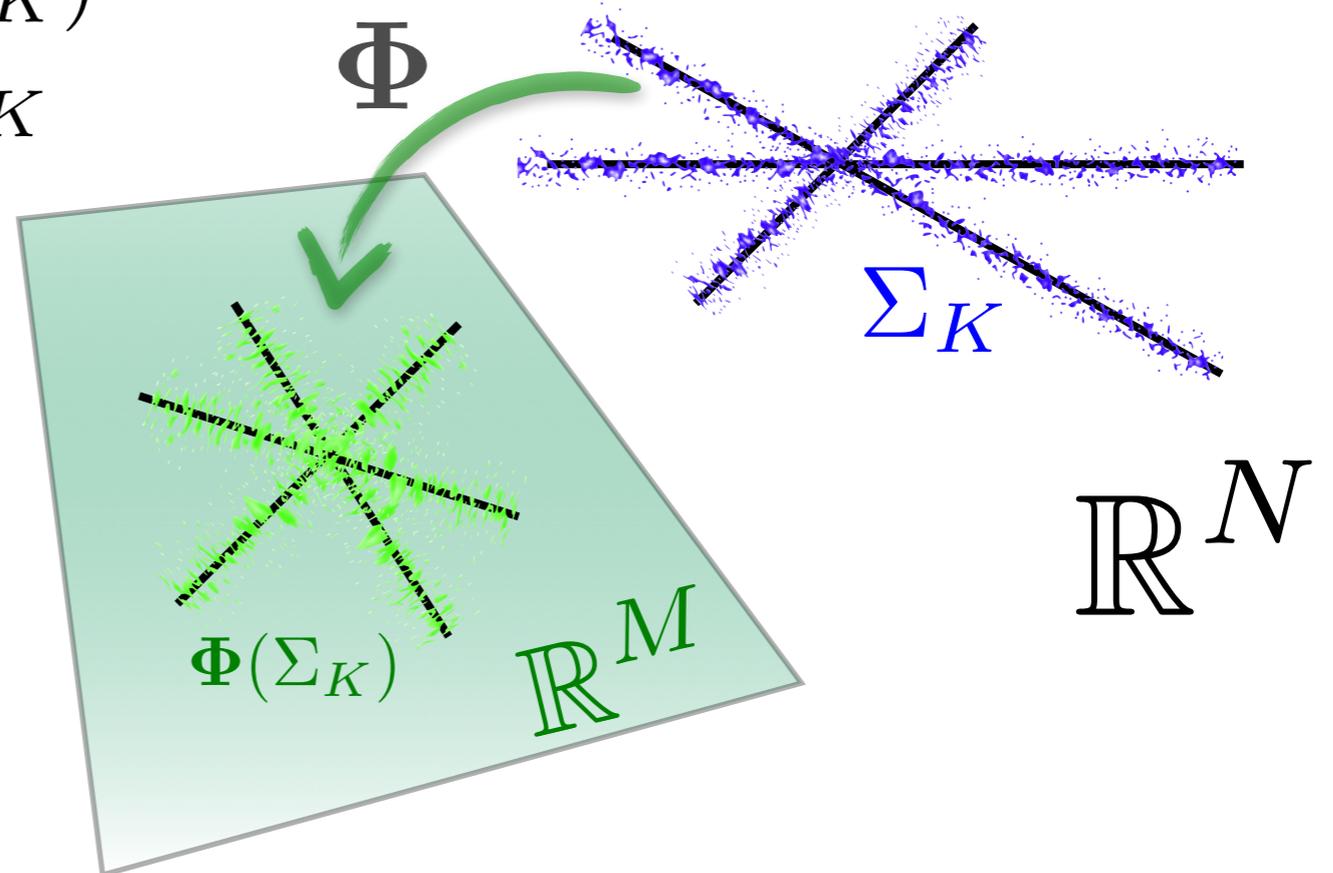
Two K -sparse signals $\mathbf{x}, \mathbf{x}' \in \Sigma_K := \{\mathbf{u} : \|\mathbf{u}\|_0 := |\text{supp } \mathbf{u}| \leq K\}$

For many random constructions of Φ (e.g., Gaussian, Bernoulli, structured) and “ $M \gtrsim K \log(N/K)$ ”, with high probability,

Geometry of $\Phi(\Sigma_K)$
 \approx Geometry of Σ_K

$$\Phi \mathbf{x} \approx \Phi \mathbf{x}' \iff \mathbf{x} \approx \mathbf{x}'$$

observations true signals

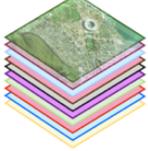


For all $\mathbf{x}, \mathbf{x}' \in \Sigma_K$ and $0 < \rho < 1$,

Restricted Isometry Property

$$(1 - \rho) \|\mathbf{x} - \mathbf{x}'\|^2 \leq \frac{1}{M} \|\Phi \mathbf{x} - \Phi \mathbf{x}'\|^2 \leq (1 + \rho) \|\mathbf{x} - \mathbf{x}'\|^2$$

Compressed Sensing...

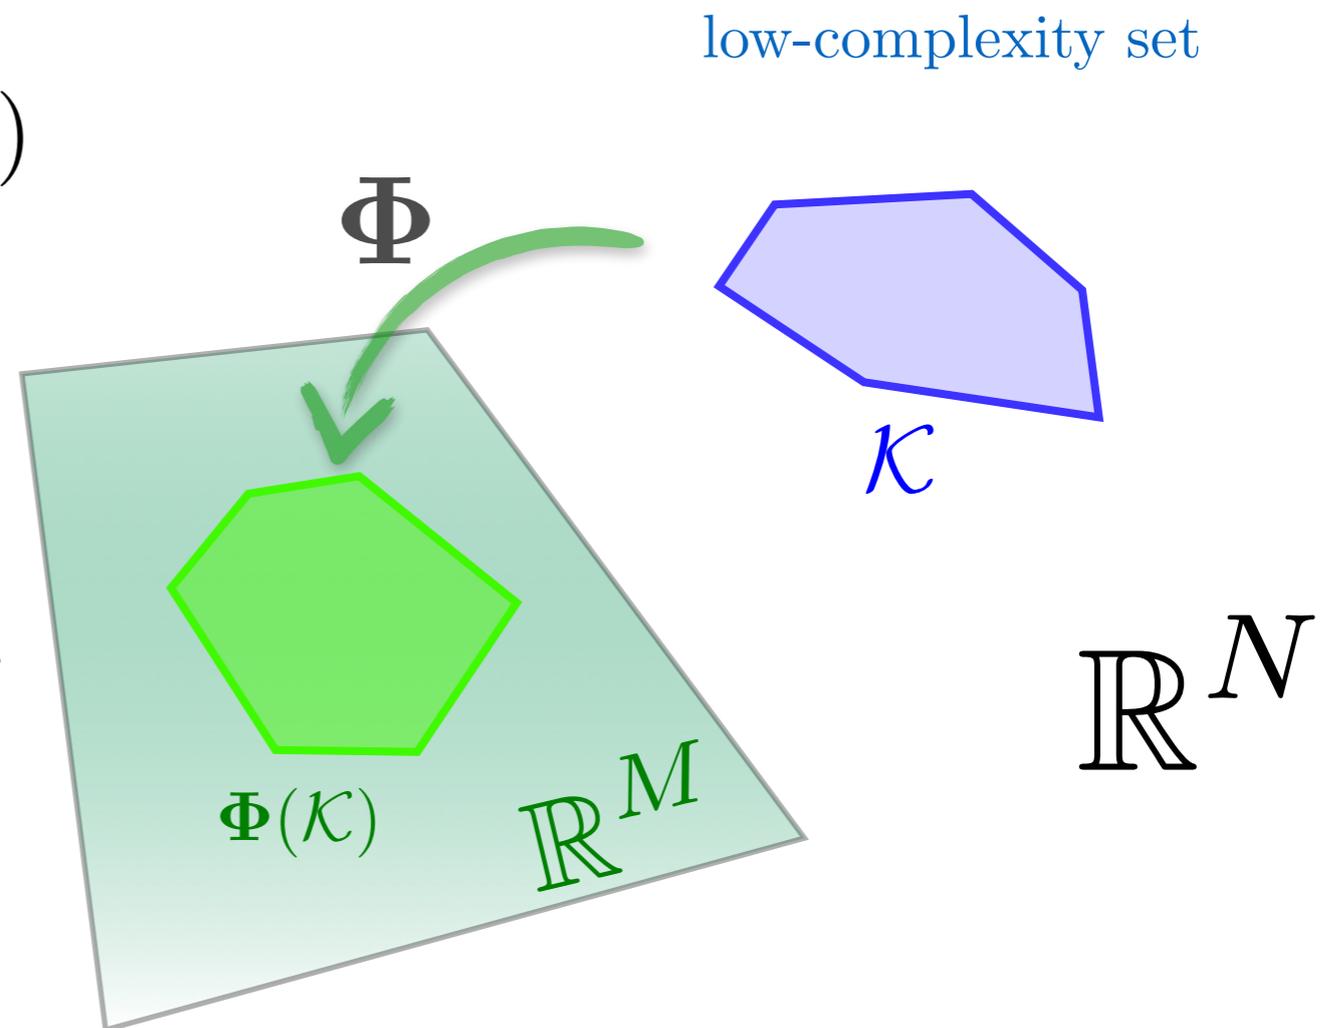
Two low-complexity signals $\mathbf{x}, \mathbf{x}' \in \mathcal{K}$ (e.g., low-rank data )

For many random constructions of Φ (e.g., Gaussian, Bernoulli, structured) and “ $M \gtrsim C_{\mathcal{K}}$ ”, with high probability,

Geometry of $\Phi(\mathcal{K})$
 \approx Geometry of \mathcal{K}

$$\Phi \mathbf{x} \approx \Phi \mathbf{x}' \Leftrightarrow \mathbf{x} \approx \mathbf{x}'$$

And generalization to many other low-complexity sets!



For all $\mathbf{x}, \mathbf{x}' \in \mathcal{K}$ and $0 < \rho < 1$,

Restricted Isometry Property

$$(1 - \rho) \|\mathbf{x} - \mathbf{x}'\|^2 \leq \frac{1}{M} \|\Phi \mathbf{x} - \Phi \mathbf{x}'\|^2 \leq (1 + \rho) \|\mathbf{x} - \mathbf{x}'\|^2$$

Compressed Sensing...

Use non-linear reconstruction methods: e.g.,

Basis Pursuit DeNoise [Chen, Donoho, Saunders, 98]

$$\hat{\mathbf{x}} \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \|\mathbf{y} - \Phi \mathbf{u}\| \leq \epsilon$$

If $\frac{1}{\sqrt{M}} \Phi$ respects the Restricted Isometry Property (RIP)

Then, if $\rho < \sqrt{2} - 1$ [Candès, 09]

(with $f \lesssim g \equiv \exists c > 0 : f \leq cg$)

Robustness: *vs* sparse deviation + noise.

$$\|\mathbf{x} - \hat{\mathbf{x}}\| \lesssim \underbrace{\frac{1}{\sqrt{K}} \|\mathbf{x} - \mathbf{x}_K\|_1}_{e_0(K): \text{ error of the model}} + \underbrace{\frac{\epsilon}{\sqrt{M}}}_{\text{noise}}$$

hidden constant

$e_0(K)$: error of the model

noise

The Power of Random Projections

At the heart of CS: *random projections!*

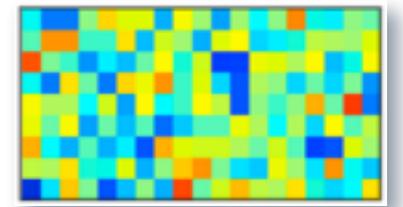
as realized by random sensing matrices

For instance:

- ▶ **random sub-Gaussian ensembles** (e.g., Gaussian, Bernoulli)

e.g., Gaussian: $\Phi \in \mathbb{R}^{M \times N}$, with $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$

or $\Phi_{ij} \sim_{\text{iid}} \pm 1$ (eq. prob), \dots

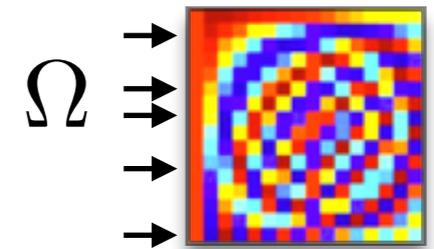


or *structured* sensing matrices (less memory, fast computations):

- ▶ **random Fourier/Hadamard ensembles** (e.g., for CT, MRI);

e.g., $\Phi = \mathbf{F}_\Omega$, with $\mathbf{F} \in \mathbb{C}^{n \times n}$

and random $\Omega \subset \{1, \dots, n\}$, $|\Omega| = m$

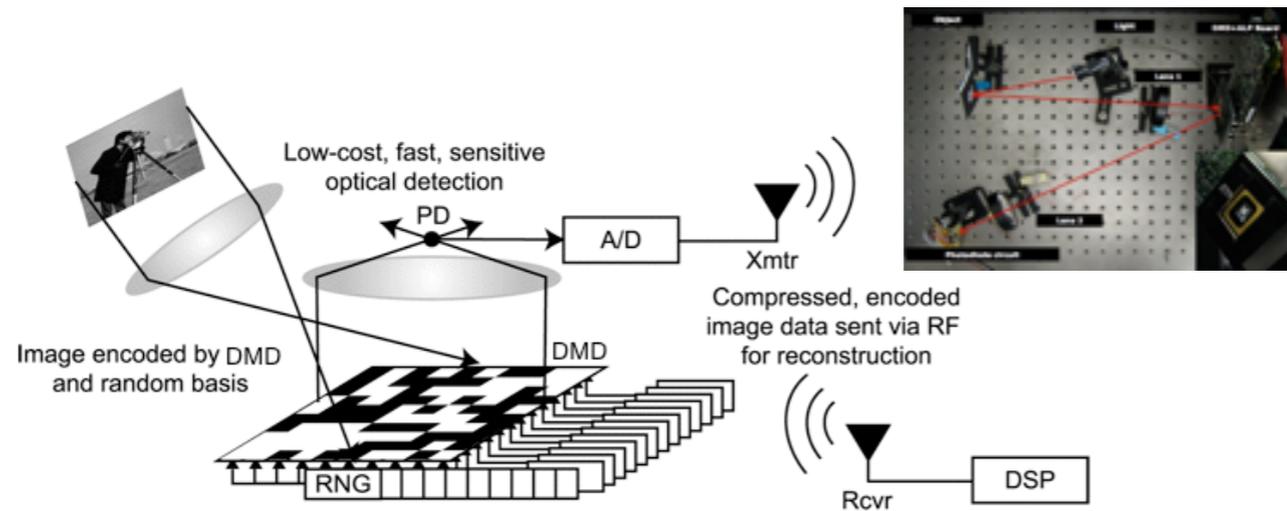


- ▶ **random convolutions, spread-spectrum** (e.g., for imaging), \dots

(see, e.g., [Foucart, Rauhut, 2013])

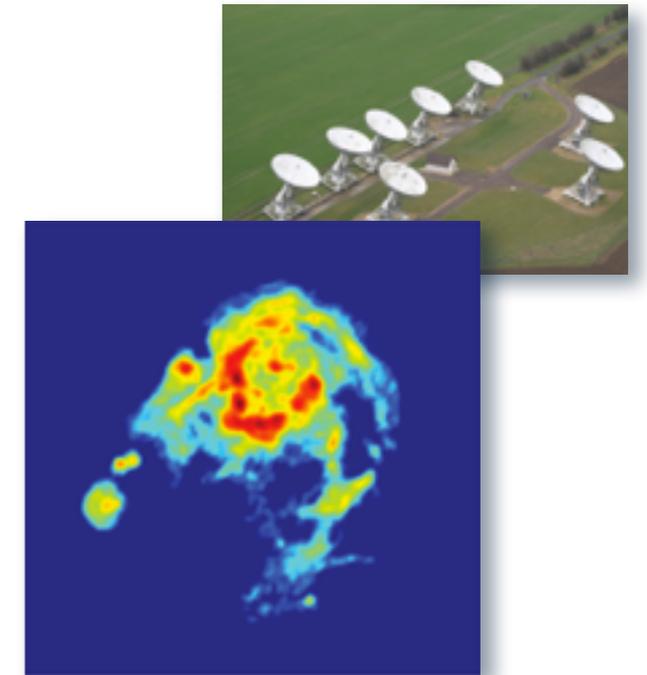
CS Applications

Rice Single Pixel Camera



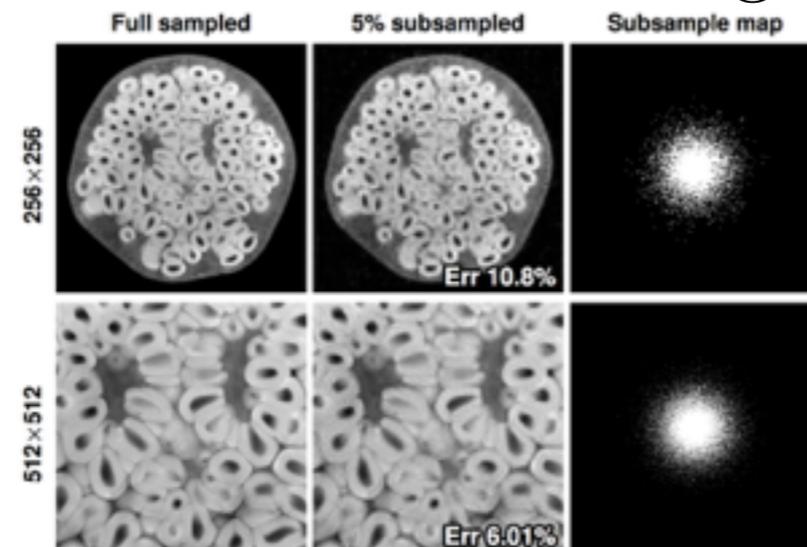
[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 08]

Radio-Astronomy

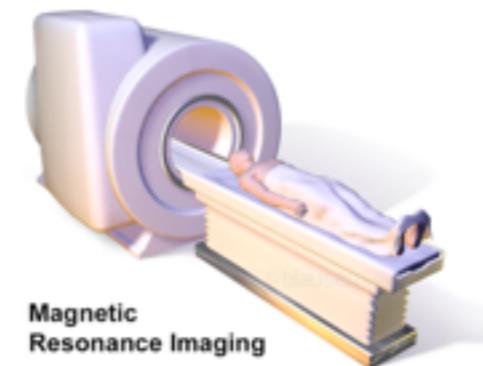


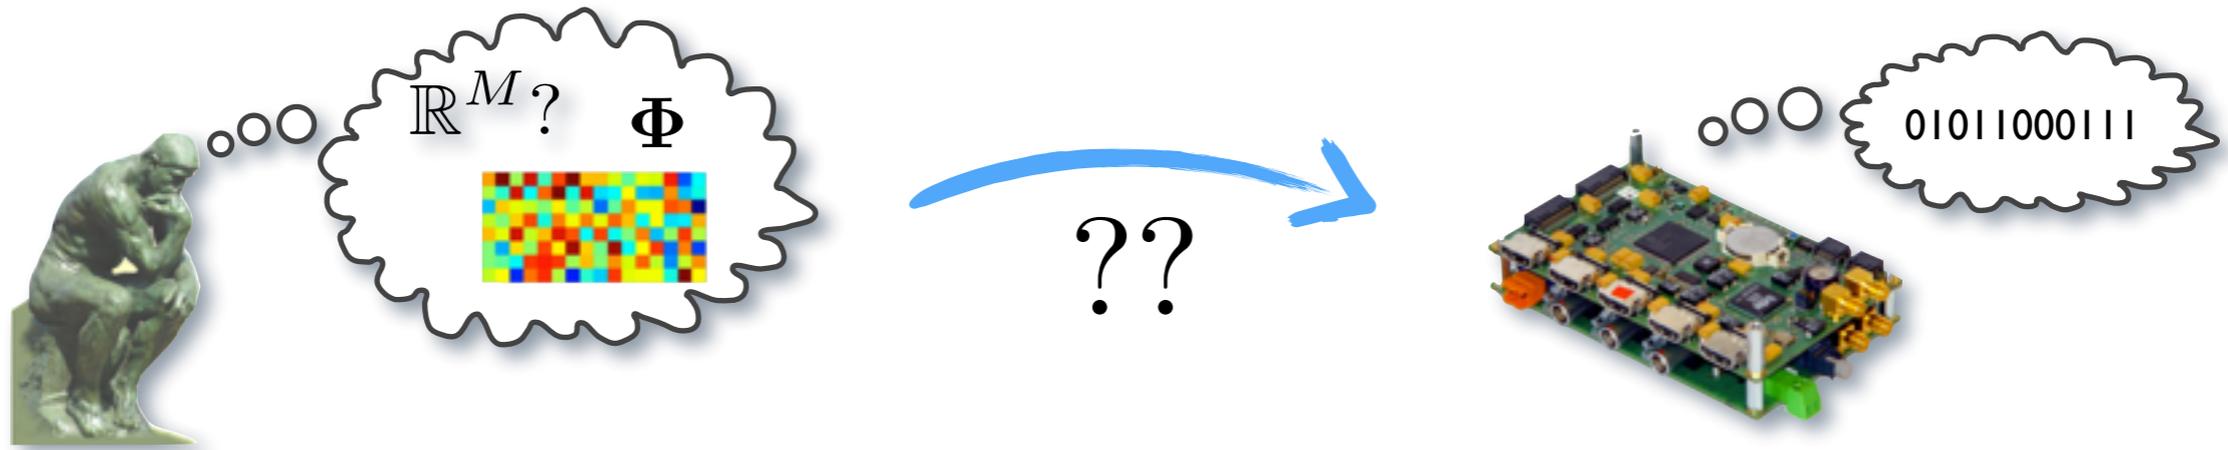
[Carrillo, McEwen, Wiaux, 12]

MRI & Fourier Imaging



[Roman, Adcock, Hansen, 14 + Siemens]



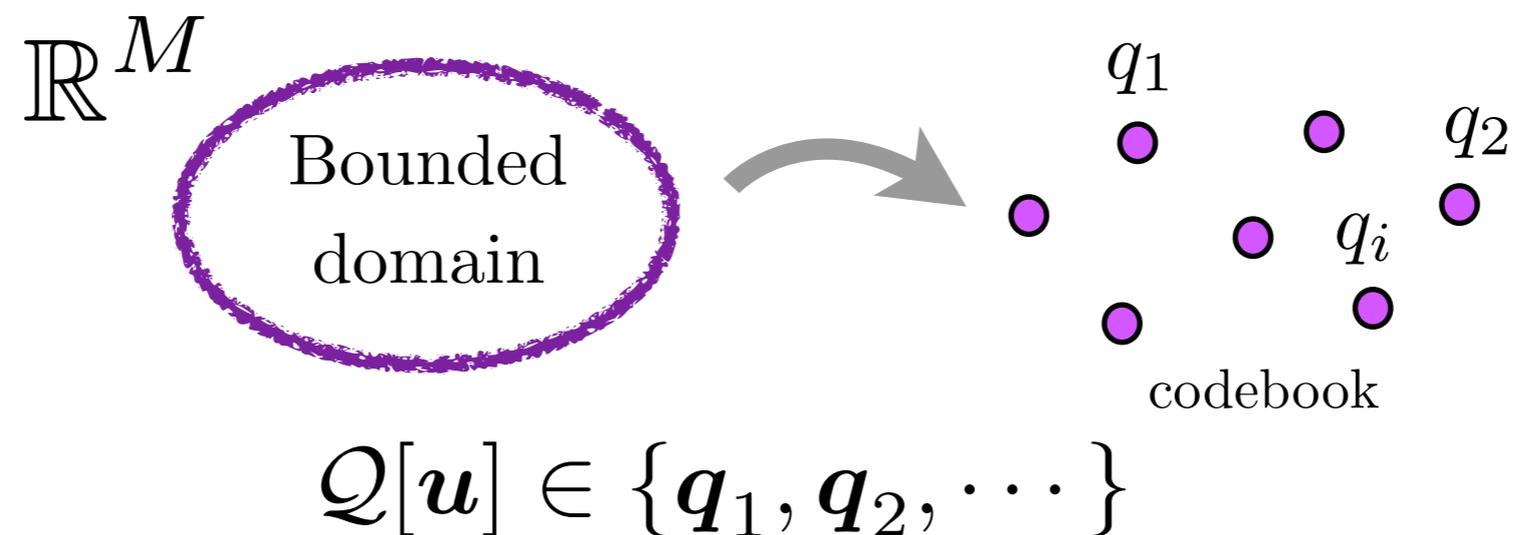


Quantizing Compressive Sensing theory

What is quantization?

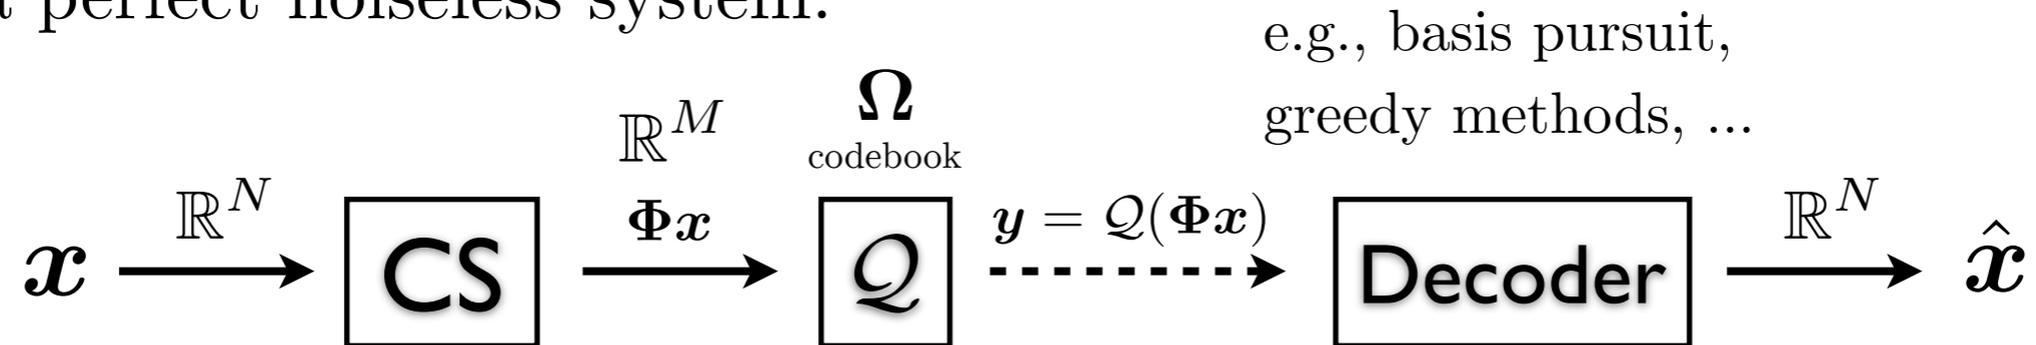
- ▶ Generality:

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



Quantizing Compressed Sensing?

In a perfect noiseless system:



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

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

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

given a certain number of:

bits, measurements, or bits/meas.

Best tradeoff? Large M and low bits/meas. *or* low M and high resolution?

Examples of quantization

- ▶ Simple example: rounding/flooring*

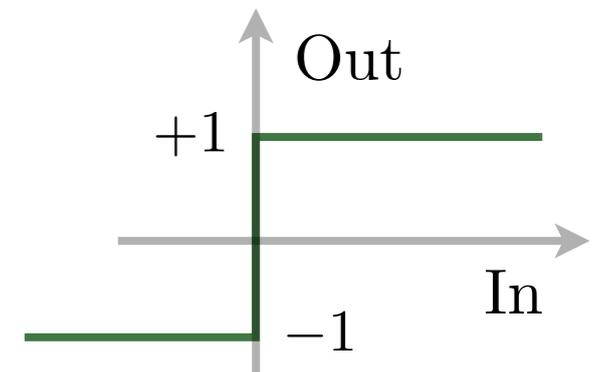
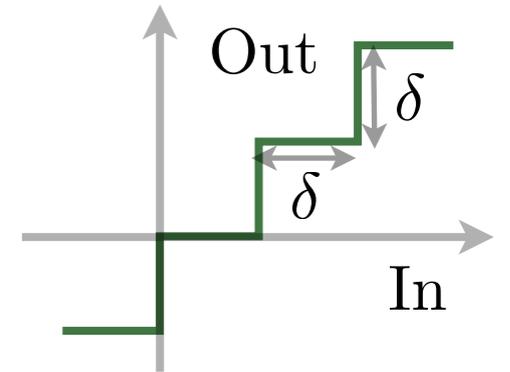
$$Q[\lambda] = \delta \lfloor \frac{\lambda}{\delta} \rfloor \in \delta\mathbb{Z}$$

for some resolution $\delta > 0$ and $Q(\mathbf{u}) = (Q(u_1), Q(u_2), \dots)$.

- ▶ Even simpler: 1-bit quantizer

$$Q[\lambda] = \text{sign } \lambda \in \pm 1$$

(with lost of the global measurement amplitude)

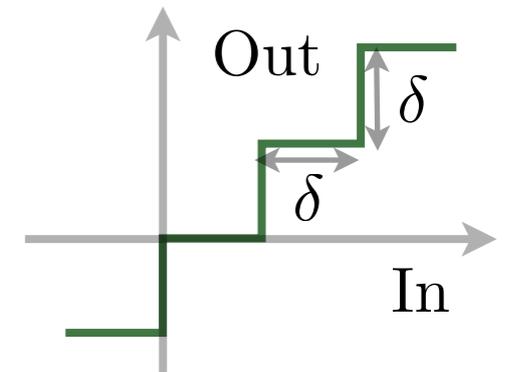


*: Also known as a special case of Pulse Code Modulation - PCM, or Memoryless Scalar Quantization - MSQ

Examples of quantization

- ▶ Simple example: rounding/flooring

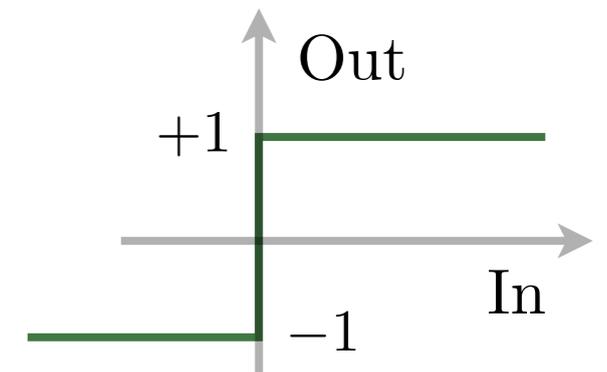
$$Q[\lambda] = \delta \lfloor \frac{\lambda}{\delta} \rfloor \in \delta\mathbb{Z}$$



for some resolution $\delta > 0$ and $Q(\mathbf{u}) = (Q(u_1), Q(u_2), \dots)$.

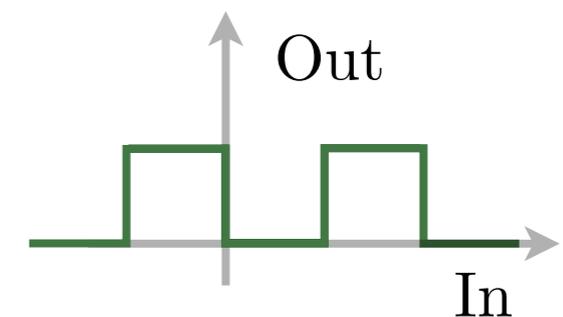
- ▶ Even simpler: 1-bit quantizer

$$Q[\lambda] = \text{sign } \lambda \in \pm 1$$



- ▶ Non-regular: square wave (or LSB)

$$Q[\lambda] := \delta (\lfloor \frac{\lambda}{\delta} \rfloor \bmod 2)$$



- ▶ Other examples (not covered here): Non-uniform scalar quantizer, vector quantizer, $\Sigma\Delta$ quantizer/noise shaping, ...

(see the works of, e.g., [\[Gunturk, Lammers, Powell, Saab, Yilmaz, Goyal\]](#))

First attempts [Candès, Tao, 04]

- ▶ Quantization is like a noise! (e.g., for $Q[\lambda] = \delta \lfloor \frac{\lambda}{\delta} \rfloor \in \delta\mathbb{Z}$)

$$\mathbf{y} = Q(\Phi\mathbf{x}) = \Phi\mathbf{x} + \mathbf{n}, \quad \text{with } \mathbf{n} = Q(\Phi\mathbf{x}) - \Phi\mathbf{x}.$$
$$\text{and } \|\mathbf{n}\|^2 = O(m\delta^2)$$

- ▶ **Problem:**

e.g., for BPDN, with

$$\hat{\mathbf{x}} \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \|\mathbf{y} - \Phi\mathbf{u}\| \leq \epsilon$$

$\Rightarrow \|\mathbf{x} - \hat{\mathbf{x}}\| \lesssim \frac{\epsilon}{\sqrt{m}} = O(\delta)$ does not decay if m increases!

counterintuitive?

First attempts [Candès, Tao, 04]

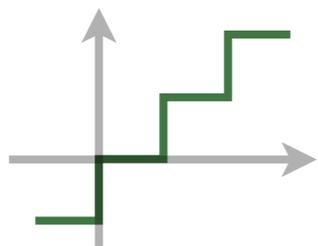
► Causes of the problem:

① Solution is **not consistent** (*i.e.*, we lost something):

$$\|\Phi \hat{x} - y\| \text{ small} \not\Rightarrow Q(\Phi \hat{x}) = y$$

② Quantization is discontinuous (it does not “**dither**”)

$$\exists x : \|x - \hat{x}\| > C \quad (\text{e.g., if } \Phi_{ij} \in \{\pm 1\})!$$



[Plan, Vershynin, 13] [LJ, 17]

① & ② = two main ingredients in this talk

The power of consistency

1

- ▶ Former insight: [Thao, Vetterli, 94]

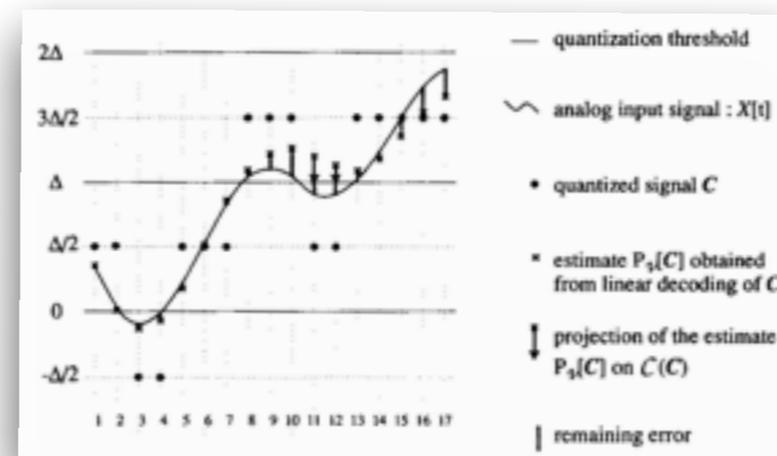
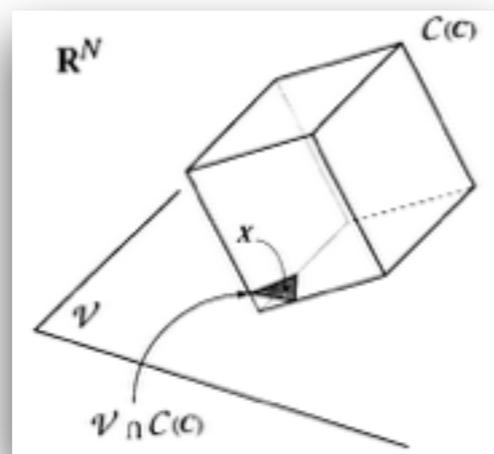
Quantization of oversampled, band-limited signals

$$\|\mathbf{x} - \hat{\mathbf{x}}\| = O(\delta/R)$$

with R the oversampling rate (i.e., M in CS),

and $\hat{\mathbf{x}}$ an estimate that is both band-limited (same prior)

and consistent (non-linear reconstruction needed)



N. Thao and M. Vetterli, “[Deterministic analysis of oversampled A/D conversion and decoding improvement based on consistent estimates](#),” Signal Processing, IEEE Transactions on Signal Processing, vol. 42, no. 3, pp. 519–531, 1994.

The power of dithering

- ▶ Inject a pre-quantization, uniform “noise”:

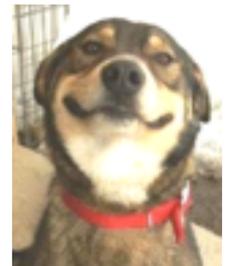
Given a resolution $\delta > 0$,

a scalar quantizer $Q(\cdot) := \delta \lfloor \cdot / \delta \rfloor$,

a “well-behaved” Φ (see later),

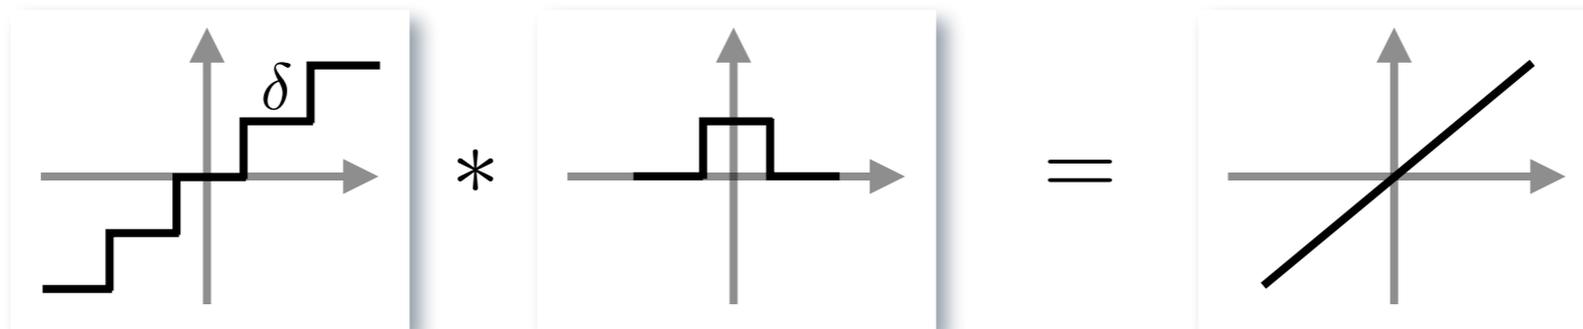
and a dithering $\xi \in \mathbb{R}^m$ with $\xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta])$

The good boy!



$$A(\mathbf{x}) := Q(\Phi \mathbf{x} + \xi)$$

- ▶ Motivation? $\mathbb{E}_{\xi} Q(\mathbf{u} + \xi) = \mathbf{u}$
 $\Rightarrow A(\mathbf{x}) \approx \Phi \mathbf{x}$ if M large



The power of dithering

- ▶ Inject a pre-quantization, uniform “noise”:

Given a resolution $\delta > 0$,

a scalar quantizer $Q(\cdot) := \delta \lfloor \cdot / \delta \rfloor$,

a “well-behaved” Φ (see later),

and a dithering $\xi \in \mathbb{R}^m$ with $\xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta])$

$$A(\mathbf{x}) := Q(\Phi \mathbf{x} + \xi)$$

- ▶ Motivation? $\mathbb{E}_{\xi} Q(\mathbf{u} + \xi) = \mathbf{u}$
 $\Rightarrow A(\mathbf{x}) \approx \Phi \mathbf{x}$ if M large
- ▶ Interesting quantized (discrete) random embedding of low-complexity signals!

Properties of dithered, quantized random projections

Parenthesis:

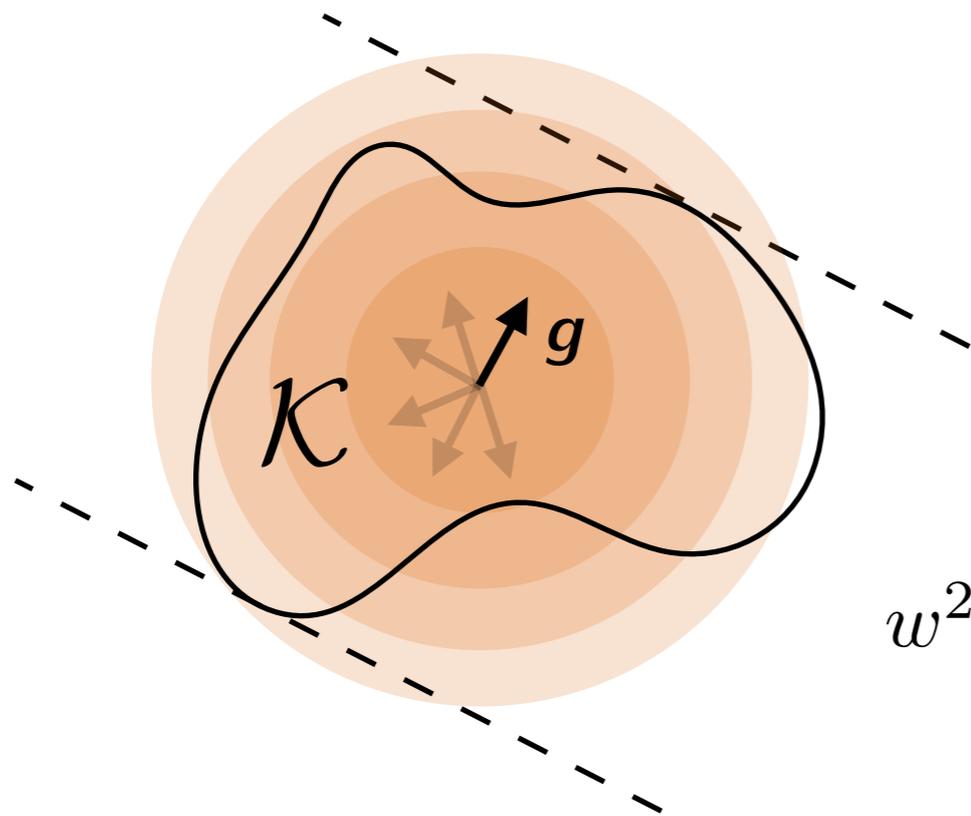
How to measure the dimension
of a low-complexity set?

Parenthesis:

Gaussian (Mean) Width (GW):

Let $\mathcal{K} \subset \mathbb{R}^n$, $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$,

$$w(\mathcal{K}) = \mathbb{E}_{\mathbf{g}} \sup_{\mathbf{x} \in \mathcal{K}} |\langle \mathbf{x}, \mathbf{g} \rangle|$$



Examples:

$$w^2(\mathcal{K}) \lesssim \log |\mathcal{K}|$$

$$w^2(\mathbb{B}^n) \lesssim n$$

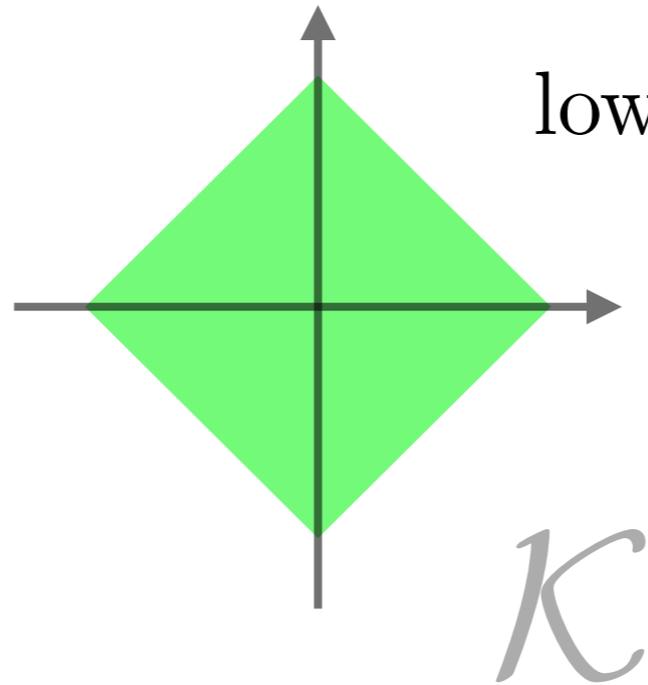
$$w^2(\Sigma_k^n \cap \mathbb{B}^n) \lesssim k \log(n/k)$$

$$w^2(\mathcal{M}_r \cap \mathbb{B}_F^{n \times n}) \lesssim rn$$

$$w^2(\cup_{i=1}^T \mathcal{K}_i) \lesssim \log T + \max_i w^2(\mathcal{K}_i)$$

⋮

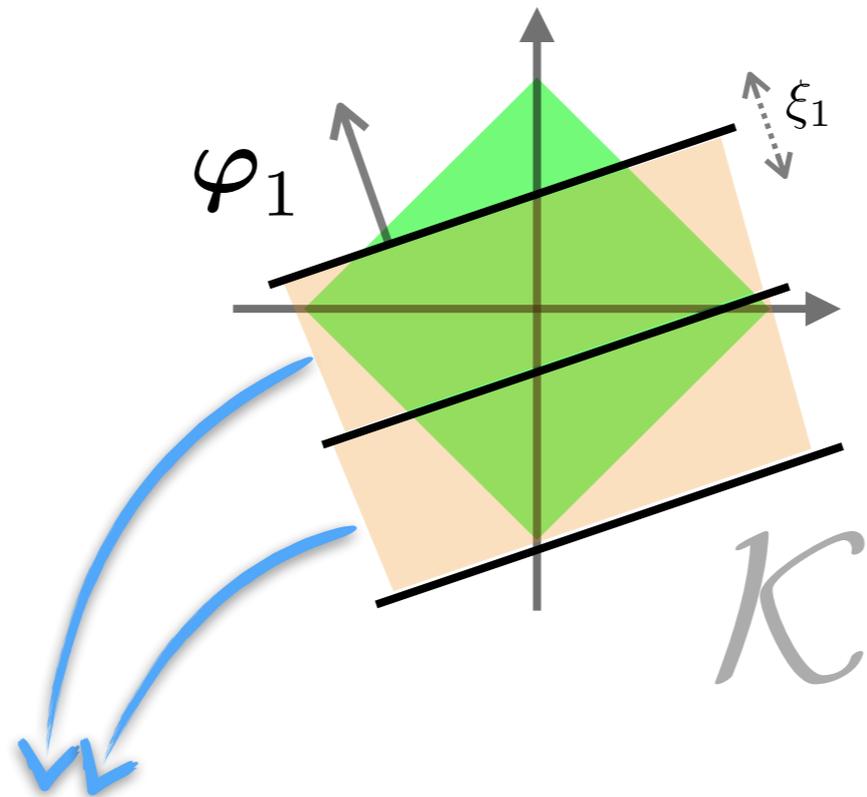
A Control of the “consistency width”



low complexity set \mathcal{K}

(e.g., sparse signals,
low-rank matrix,
compressible signals, ...)

A Control of the “consistency width”

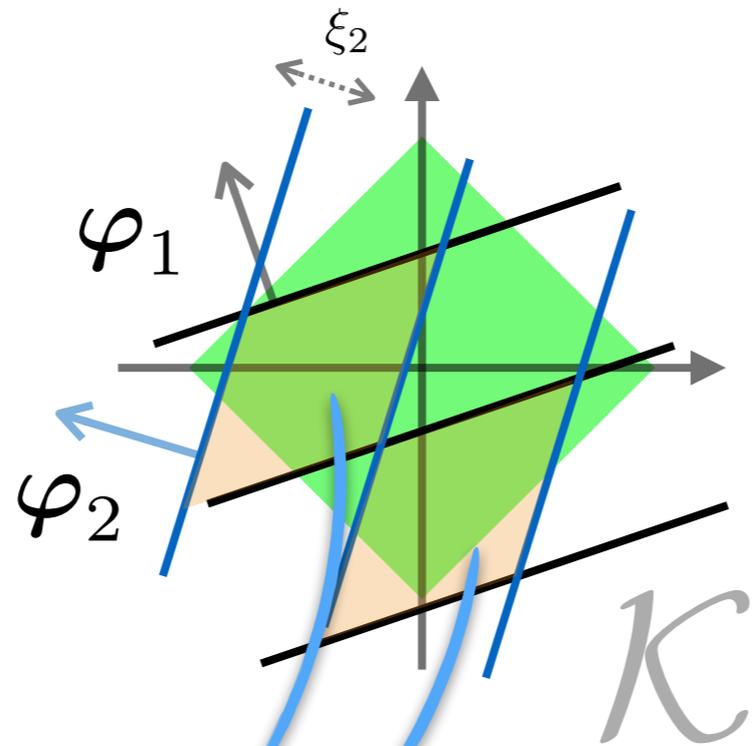


$$\Phi = \begin{pmatrix} \varphi_1^T \\ \vdots \\ \varphi_M^T \end{pmatrix}$$

Signals u s.t.

$$\underbrace{Q(\varphi_1^T u + \xi_1)}_{\delta[(\varphi_1^T u + \xi_1)/\delta]} = \text{cst.}$$

A Control of the “consistency width”

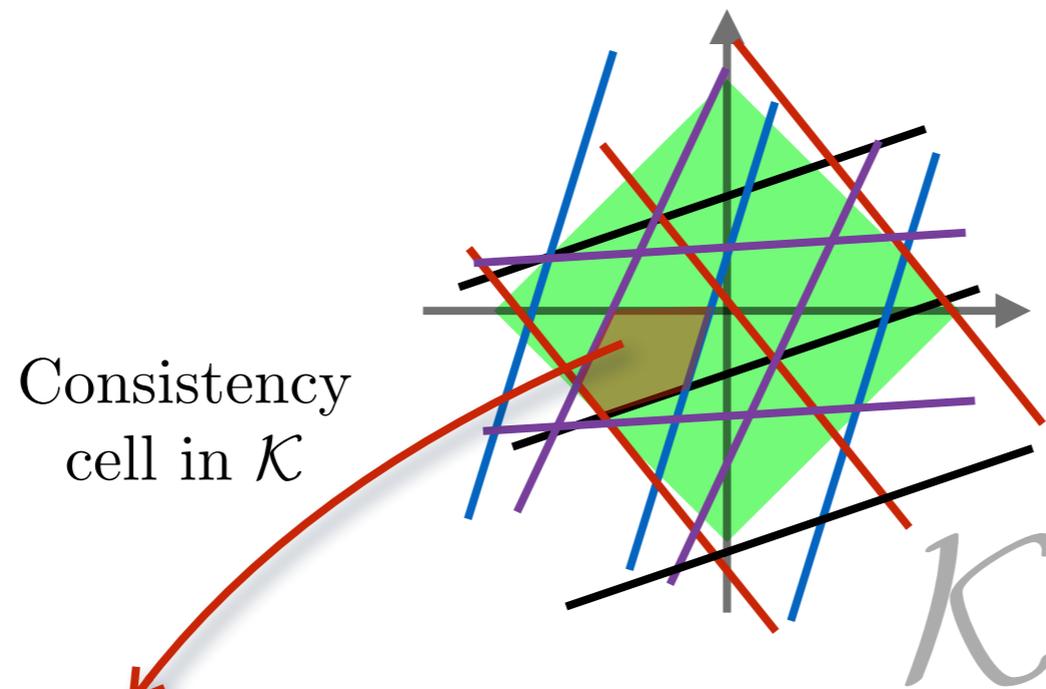


$$\Phi = \begin{pmatrix} \varphi_1^T \\ \vdots \\ \varphi_M^T \end{pmatrix}$$

Signals u s.t.

$$\left. \begin{aligned} Q(\varphi_1^T u + \xi_1) &= \text{cst.} \\ Q(\varphi_2^T u + \xi_2) &= \text{cst.} \end{aligned} \right\}$$

A Control of the “consistency width”



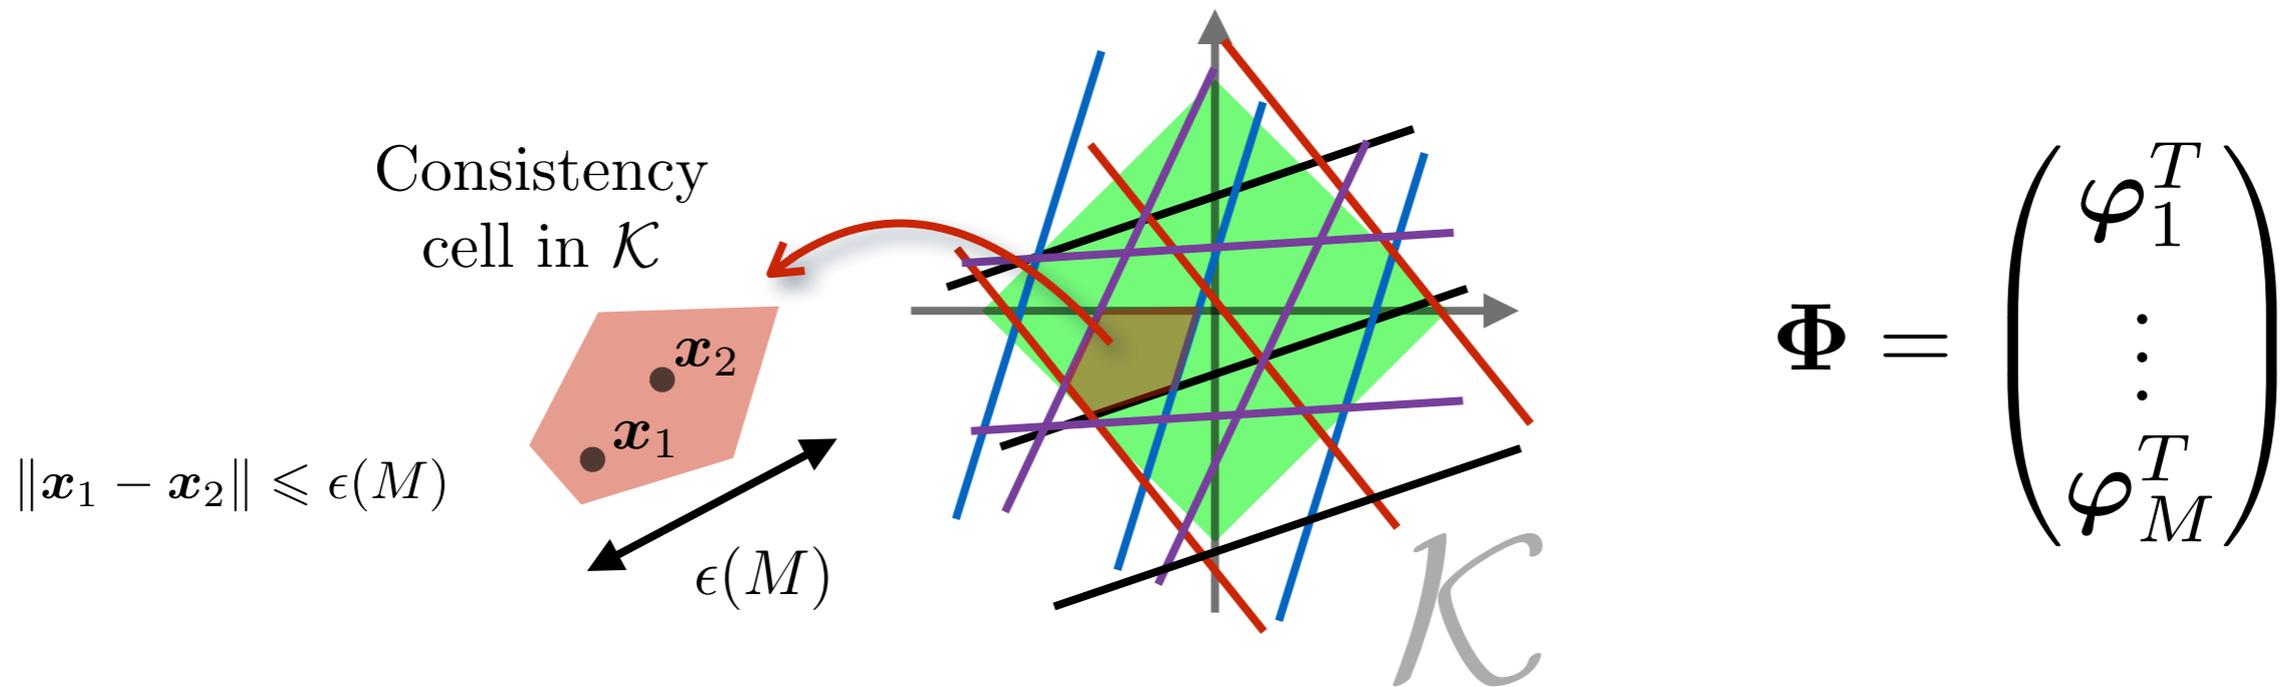
$$\Phi = \begin{pmatrix} \varphi_1^T \\ \vdots \\ \varphi_M^T \end{pmatrix}$$

Signals \mathbf{u} s.t.

$$A(\mathbf{u}) := \mathcal{Q}(\Phi \mathbf{u} + \boldsymbol{\xi}) = \mathbf{y}$$

for some $\mathbf{y} \in \delta \mathbb{Z}^M$

A Control of the “consistency width”

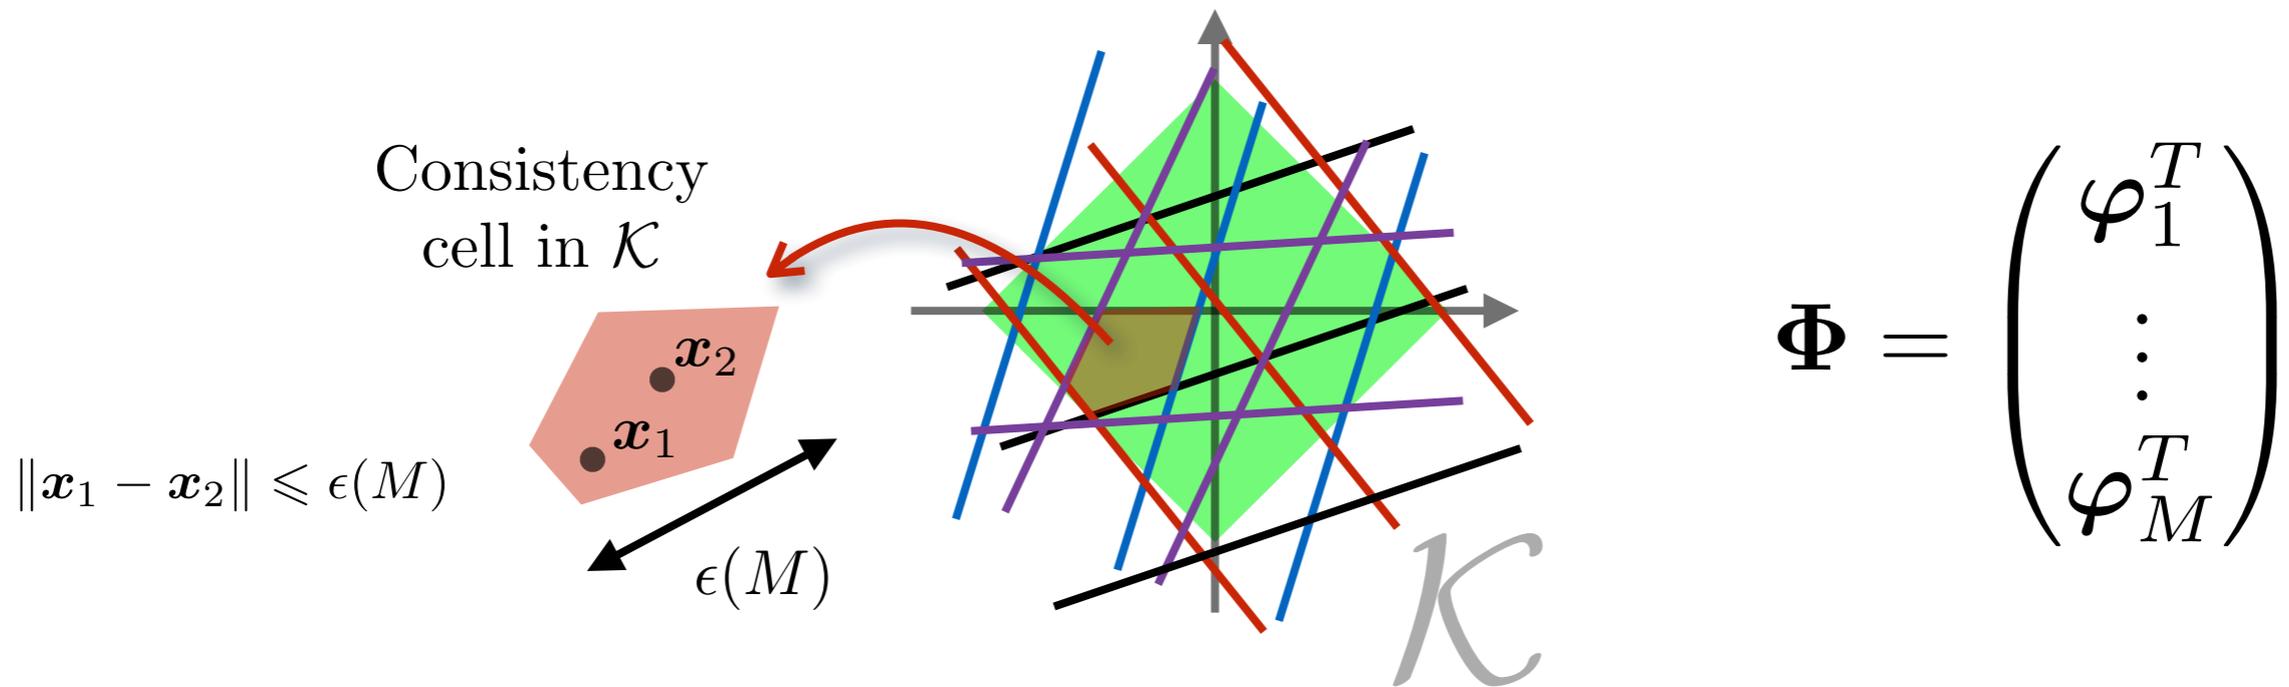


Consistency width:

$$\epsilon(M; \mathcal{K}, \delta) = \max_{\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{K}} \|\mathbf{x}_1 - \mathbf{x}_2\| \text{ s.t. } \mathbf{A}(\mathbf{x}_1) = \mathbf{A}(\mathbf{x}_2)$$

Size $\epsilon(M)$ should decay for large M !

A Control of the “consistency width”



For Φ a random Gaussian matrix, with high probability,

[LJ, 16], [LJ, 17]

$$\epsilon(M) \leq C_{\mathcal{K},\delta} M^{-1/q}.$$

For \mathcal{K} a structured set (e.g., sparse signals, low-rank matrices)

$$q = 1, \quad C_{\mathcal{K},\delta} = (1 + \delta)w(\mathcal{K} \cap \mathbb{B}^n)^2 \times \log \text{ factors},$$

otherwise, for a bounded set,

$$q = 4, \quad C_{\mathcal{K},\delta} = \left(\frac{1+\delta}{\delta^{1/2}}\right) w(\mathcal{K})^{1/2}.$$

Open problem:
Extension to RIP matrices?

A CW for consistent reconstruction

Let us take $\mathcal{K} = \mathcal{C}_s = \{\mathbf{u} \in \mathbb{R}^N : \|\mathbf{u}\|_1 \leq \sqrt{s}, \|\mathbf{u}\| \leq 1\}$ & $\mathbf{x}_0 \in \Sigma_s \cap \mathbb{B}^N$

$$\mathbf{x}^* \in \underset{\mathbf{u} \in \mathbb{R}^N}{\operatorname{argmin}} \|\mathbf{u}\|_1 \quad \text{s.t.} \quad \|\Phi \mathbf{u} - \Phi \mathbf{x}_0\| \leq \epsilon, \quad (\text{BPDN})$$

$$\underbrace{\mathbf{A}(\mathbf{u}) = \mathbf{A}(\mathbf{x}_0)}_{\text{convex}}, \quad \mathbf{u} \in \mathbb{B}^N. \quad (\text{CoBP})$$

A CW for consistent reconstruction

Let us take $\mathcal{K} = \mathcal{C}_s = \{\mathbf{u} \in \mathbb{R}^N : \|\mathbf{u}\|_1 \leq \sqrt{s}, \|\mathbf{u}\| \leq 1\}$ & $\mathbf{x}_0 \in \Sigma_s \cap \mathbb{B}^N$

$$\mathbf{x}^* \in \underset{\mathbf{u} \in \mathbb{R}^N}{\operatorname{argmin}} \|\mathbf{u}\|_1 \text{ s.t. } \mathbf{A}(\mathbf{u}) = \mathbf{A}(\mathbf{x}_0), \mathbf{u} \in \mathbb{B}^N. \text{ (CoBP)}$$

$$\|\mathbf{x}^*\| \leq 1 \text{ and } \|\mathbf{x}^*\|_1 \leq \|\mathbf{x}_0\|_1 \leq \sqrt{s}$$

$$\Rightarrow \mathbf{x}^* \in \mathcal{C}_s$$

Moreover $\mathbf{A}(\mathbf{x}^*) = \mathbf{A}(\mathbf{x}_0)$ (consistency)

$$\Rightarrow \|\mathbf{x}^* - \mathbf{x}_0\| \leq \epsilon(M; \mathcal{C}_s, \delta)$$

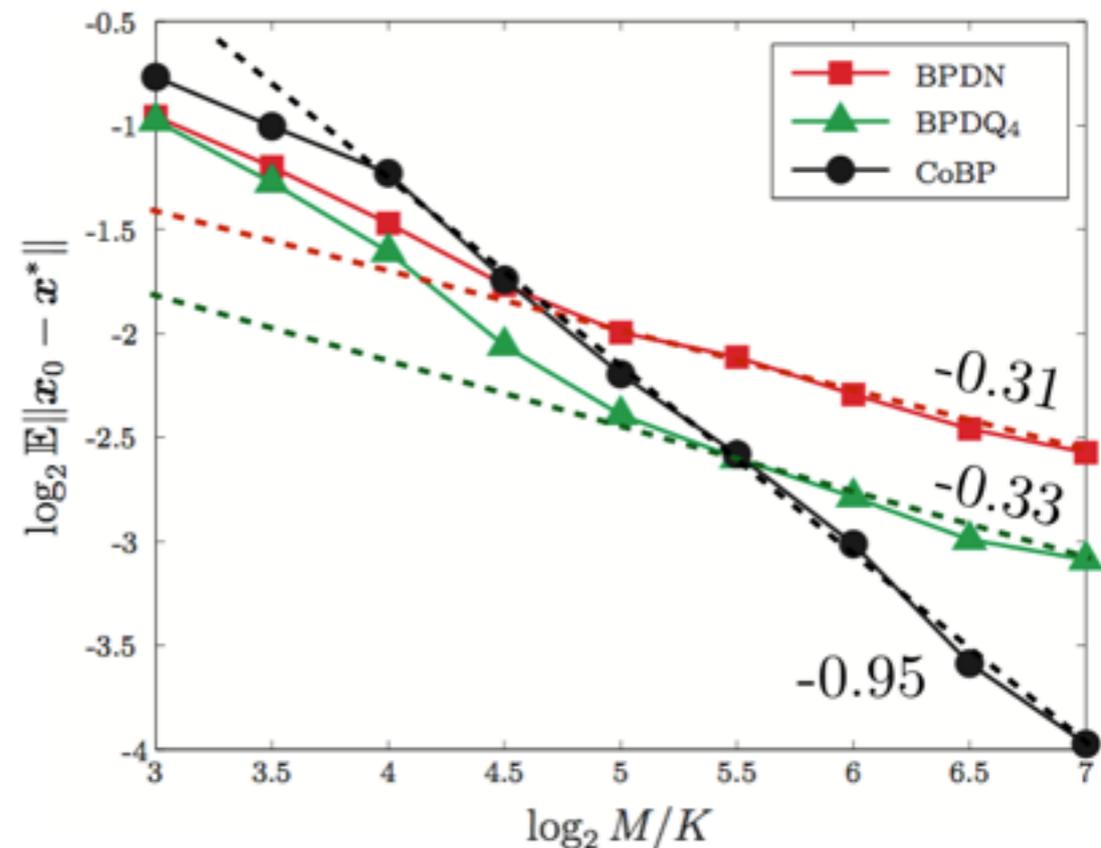
A CW for consistent reconstruction

Let us take $\mathcal{K} = \mathcal{C}_s = \{\mathbf{u} \in \mathbb{R}^N : \|\mathbf{u}\|_1 \leq \sqrt{s}, \|\mathbf{u}\| \leq 1\}$ & $\mathbf{x}_0 \in \Sigma_s \cap \mathbb{B}^N$

$$\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \mathbf{A}(\mathbf{u}) = \mathbf{A}(\mathbf{x}_0), \mathbf{u} \in \mathbb{B}^N. \text{ (CoBP)}$$

K -sparse signals with
random Gaussian sensing

$N = 2048, K = 16,$
 $B = 3, M/K \in [8, 128]$
20 trials per points



A CW for consistent reconstruction

Let us take $\mathcal{K} = \mathcal{C}_s = \{\mathbf{u} \in \mathbb{R}^N : \|\mathbf{u}\|_1 \leq \sqrt{s}, \|\mathbf{u}\| \leq 1\}$ & $\mathbf{x}_0 \in \Sigma_s \cap \mathbb{B}^N$

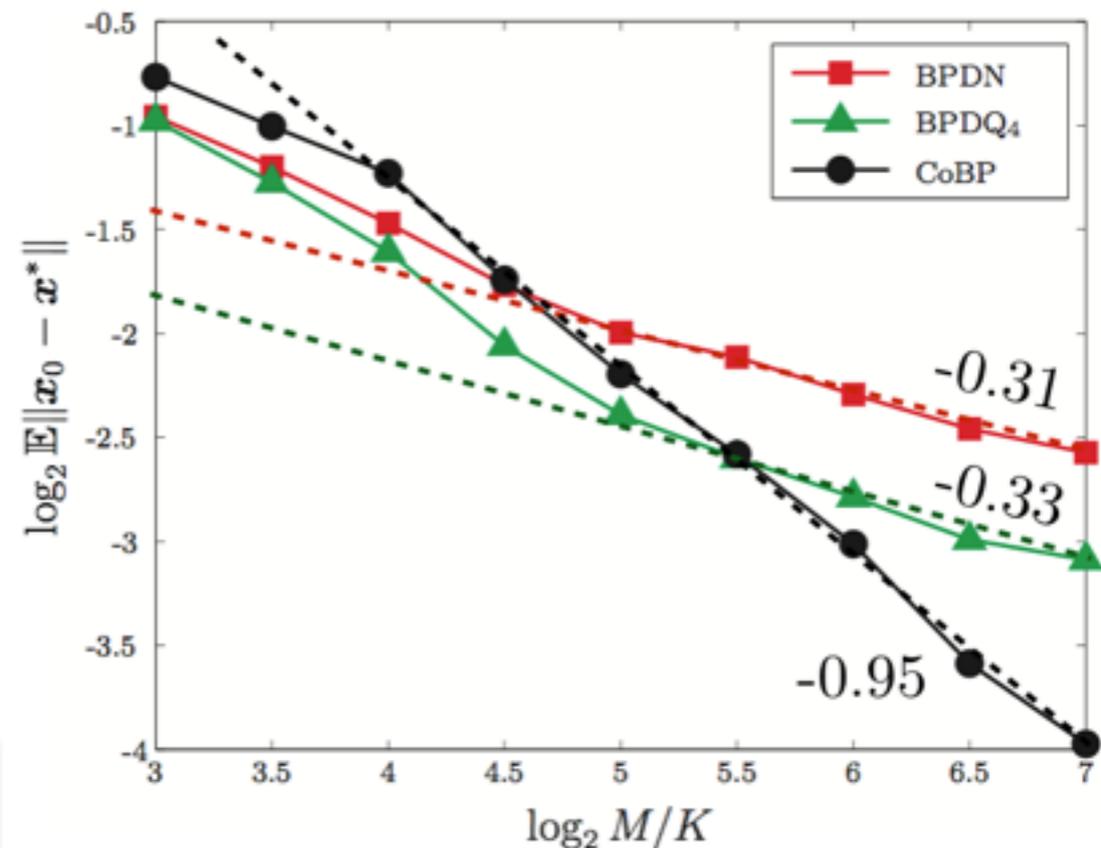
$$\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \mathbf{A}(\mathbf{u}) = \mathbf{A}(\mathbf{x}_0), \mathbf{u} \in \mathbb{B}^N. \text{ (CoBP)}$$

* $\rightarrow \|\mathbf{u}\|_{\#}$

K -sparse signals with
random Gaussian sensing

$N = 2048, K = 16,$
 $B = 3, M/K \in [8, 128]$
20 trials per points

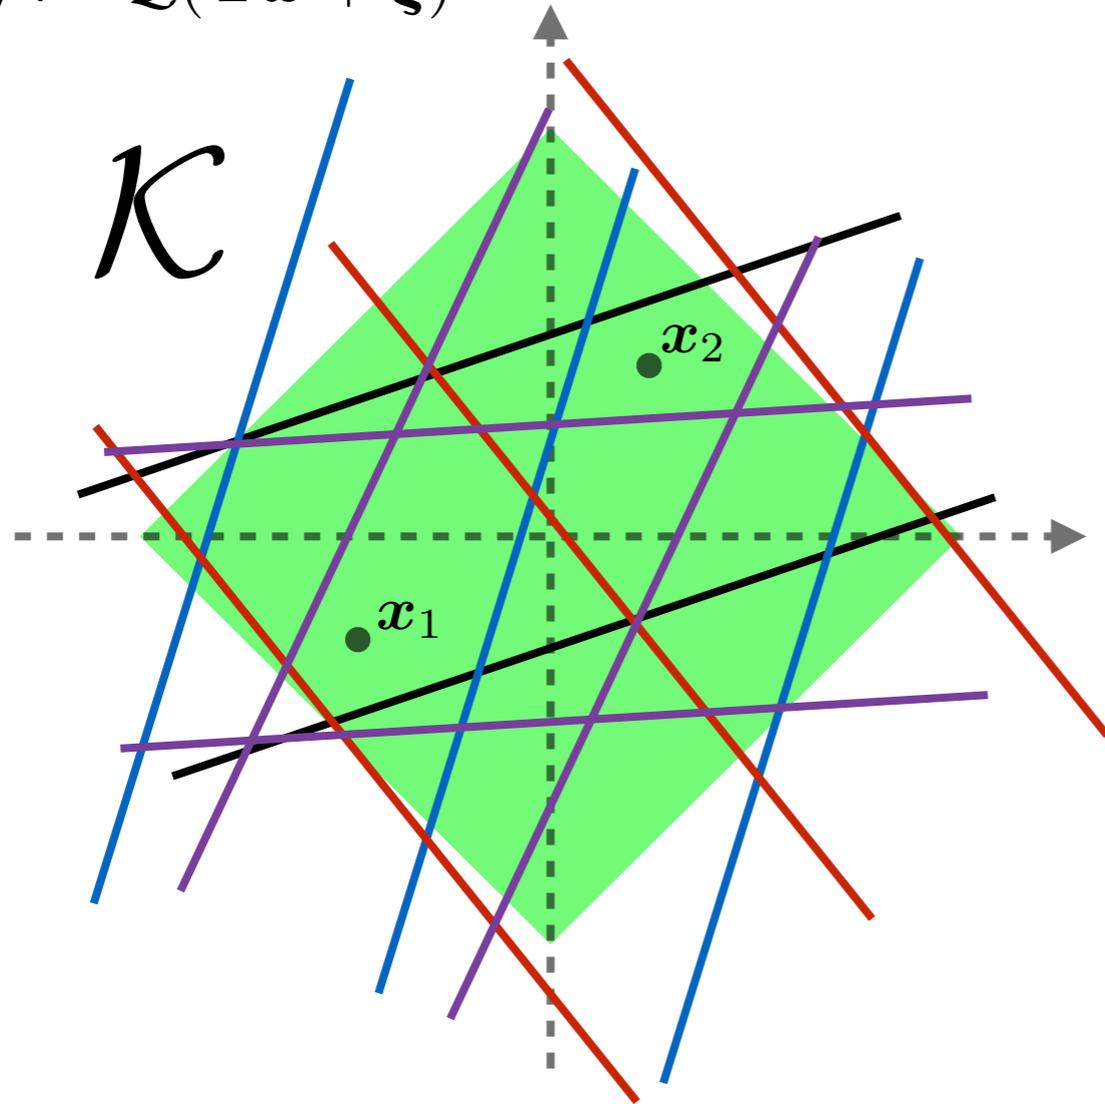
* Can be extended to other gauge/
atomic norms (e.g., nuclear norm)



[Moshtaghpour, LJ, Cambareri, Degraux, De Vleeschouwer, 17]

Ⓑ Quantizing the RIP (approximate consistency)

$$A(\mathbf{x}) := \mathcal{Q}(\Phi \mathbf{x} + \xi)$$



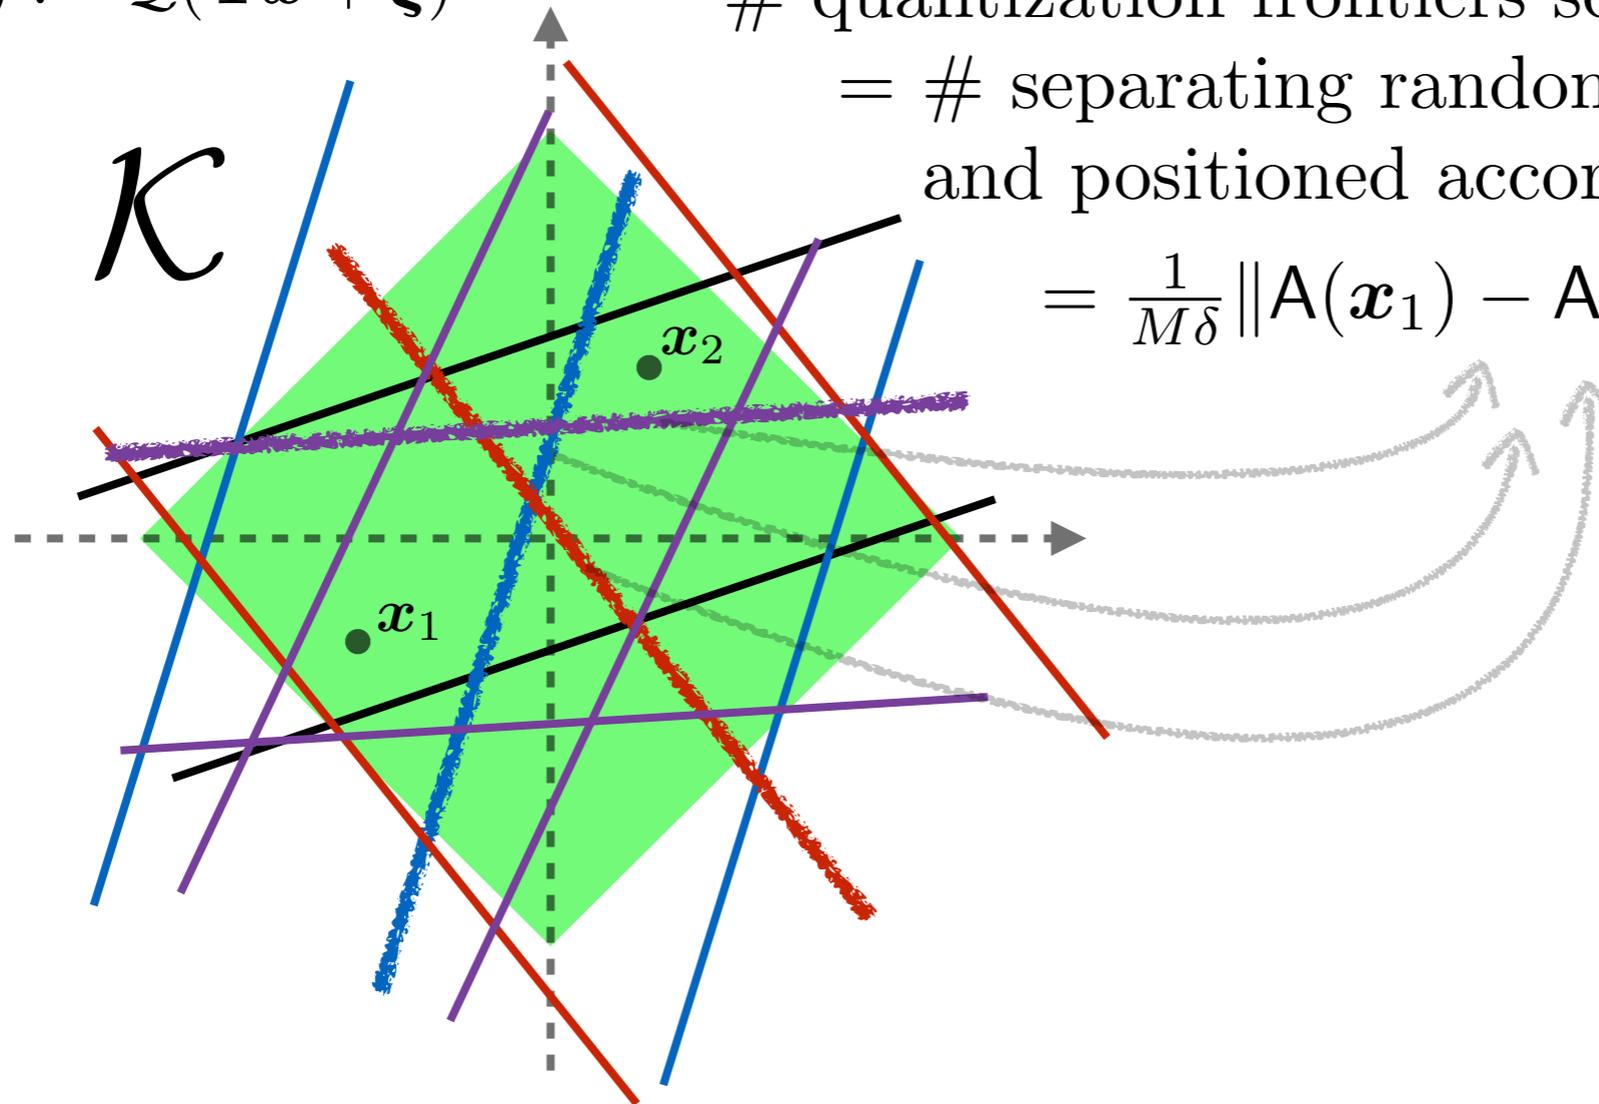
B Quantizing the RIP (approximate consistency)

$$A(\mathbf{x}) := \mathcal{Q}(\Phi \mathbf{x} + \xi)$$

quantization frontiers separating \mathbf{x}_1 and \mathbf{x}_2
= # separating random hyperplanes oriented
and positioned according to (Φ, ξ)

$$= \frac{1}{M\delta} \|A(\mathbf{x}_1) - A(\mathbf{x}_2)\|_1 \approx \|\mathbf{x}_1 - \mathbf{x}_2\|$$

??



Hope: dithering sufficiently smoothen
discontinuities to allow for RIP matrices.

B Quantizing the RIP (approximate consistency)

Let $\mathcal{K} \subset \mathbb{R}^N$ be a structured set (*e.g.*, sparse signals, low-rank matrices).

Let Φ be a (ℓ_1, ℓ_2) -RIP($\epsilon, \mathcal{K} - \mathcal{K}$) matrix, *i.e.*,

$$(1 - \epsilon)\|\mathbf{x}\|^2 \leq \frac{c_\Phi}{m} \|\Phi \mathbf{x}\|_1^2 \leq (1 + \epsilon)\|\mathbf{x}\|^2, \forall \mathbf{x} \in \mathcal{K} - \mathcal{K},$$

(*e.g.*, Gaussian random matrix, circulant Gaussian random matrix for $\mathcal{K} = \Sigma_k$)

[Dirksen, Jung, Rauhut, 17]

Provided that $M \gtrsim \epsilon^{-2} C_{\mathcal{K}} \log(1 + \frac{1}{\delta\epsilon})$, (with $C_{\mathcal{K}} > 0$ an upper bound on $w(\mathcal{K})^2$)
with probability exceeding $1 - C \exp(-\epsilon^2 m)$,

$$(1 - \epsilon)\|\mathbf{x}_1 - \mathbf{x}_2\| - c'\epsilon\delta \leq \frac{1}{m} \|\mathbf{A}(\mathbf{x}_1) - \mathbf{A}(\mathbf{x}_2)\|_1 \leq (1 + \epsilon)\|\mathbf{x}_1 - \mathbf{x}_2\| + c'\epsilon\delta,$$

for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{K} \cap \mathbb{B}^N$.

(\exists other variants with ℓ_2/ℓ_2 and standard RIP)

[LJ, Cambareri, 17]

B Quantizing the RIP (approximate consistency)

Let $\mathcal{K} \subset \mathbb{R}^N$ be a structured set (*e.g.*, sparse signals, low-rank matrices).

Let Φ be a (ℓ_1, ℓ_2) -RIP($\epsilon, \mathcal{K} - \mathcal{K}$) matrix, *i.e.*,

$$(1 - \epsilon)\|\mathbf{x}\|^2 \leq \frac{c_\Phi}{m} \|\Phi \mathbf{x}\|_1^2 \leq (1 + \epsilon)\|\mathbf{x}\|^2, \forall \mathbf{x} \in \mathcal{K} - \mathcal{K},$$

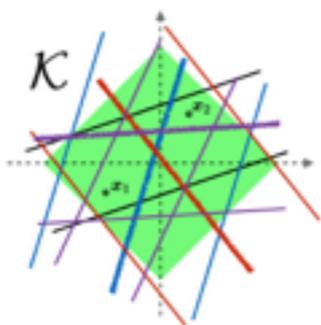
(*e.g.*, Gaussian random matrix, circulant Gaussian random matrix for $\mathcal{K} = \Sigma_k$)

[Dirksen, Jung, Rauhut, 17]

Provided that $M \gtrsim \epsilon^{-2} C_{\mathcal{K}} \log(1 + \frac{1}{\delta\epsilon})$, (with $C_{\mathcal{K}} > 0$ an upper bound on $w(\mathcal{K})^2$)
with probability exceeding $1 - C \exp(-\epsilon^2 m)$,

$$(1 - \epsilon)\|\mathbf{x}_1 - \mathbf{x}_2\| - c'\epsilon\delta \leq \frac{1}{m} \|A(\mathbf{x}_1) - A(\mathbf{x}_2)\|_1 \leq (1 + \epsilon)\|\mathbf{x}_1 - \mathbf{x}_2\| + c'\epsilon\delta,$$

for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{K} \cap \mathbb{B}^N$.



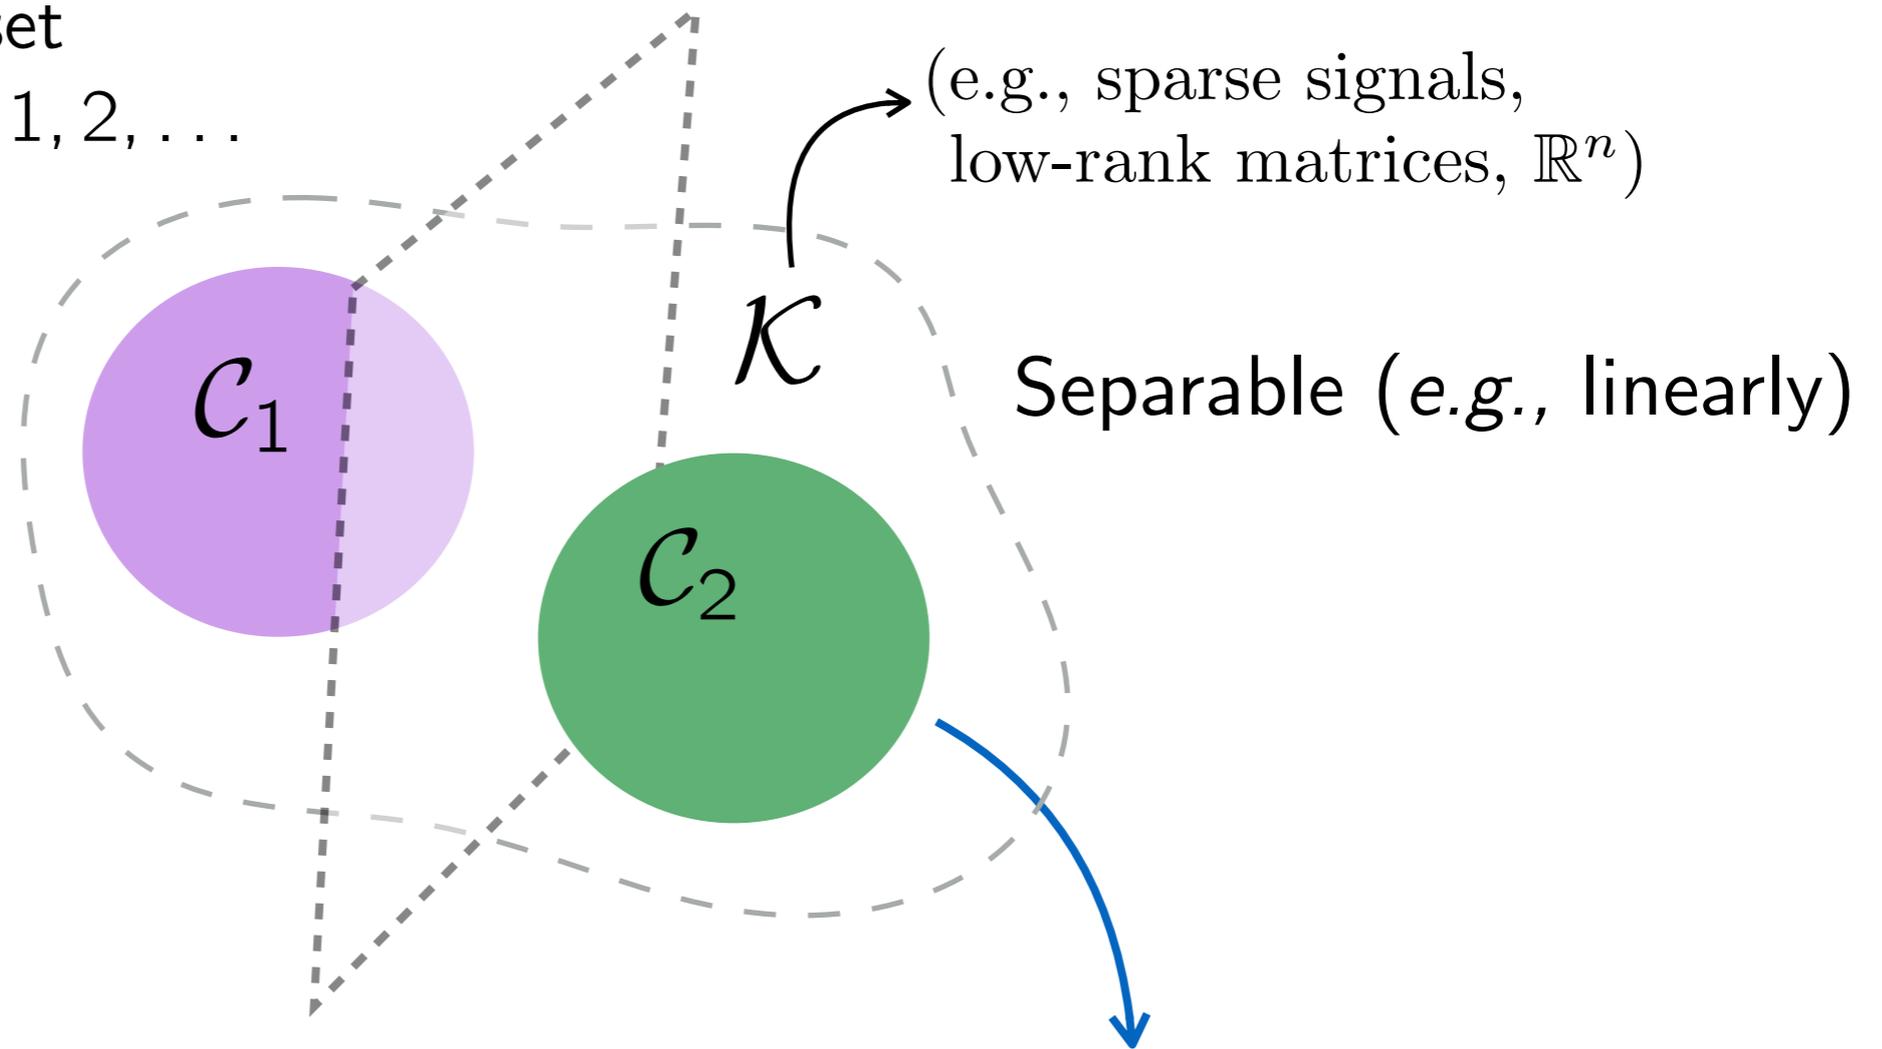
In other words, we can potentially classify signals from their quantized observations!

[LJ, Cambareri, 17]

Classification in a quantized world
“The Rare Eclipse Problem on Tiles”

The Big Picture (an easy classification problem)

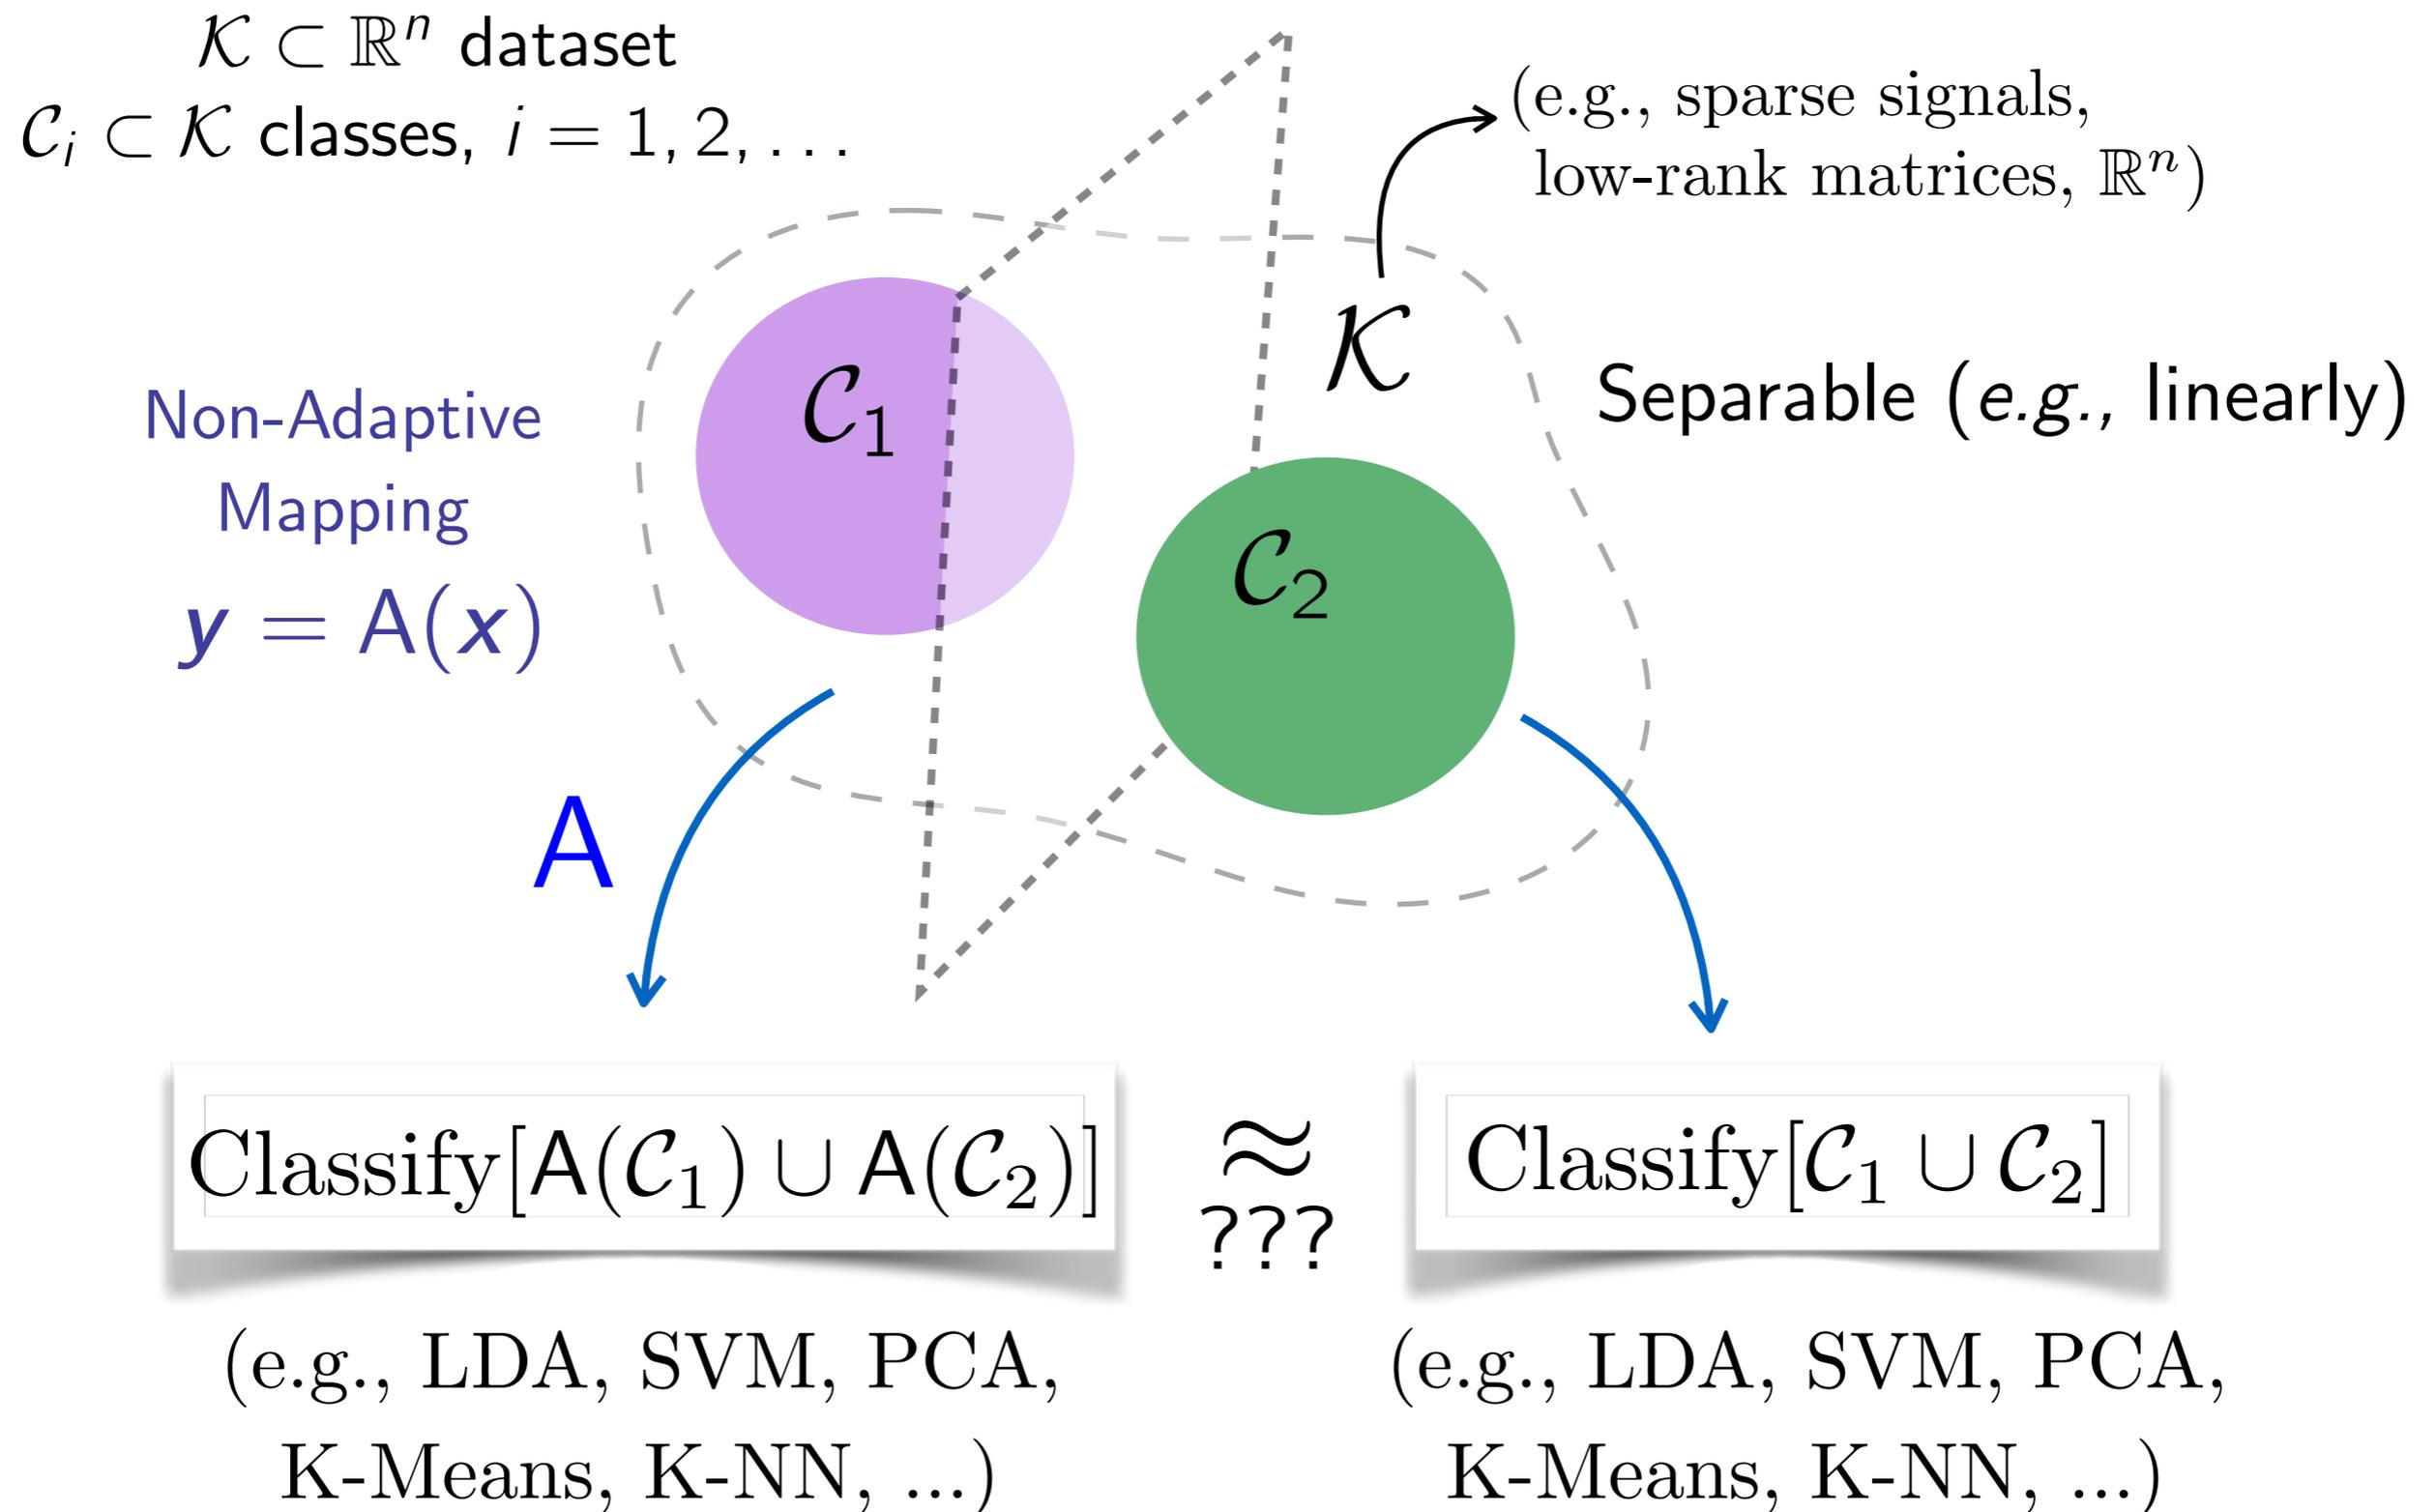
$\mathcal{K} \subset \mathbb{R}^n$ dataset
 $\mathcal{C}_i \subset \mathcal{K}$ classes, $i = 1, 2, \dots$



Classify[$\mathcal{C}_1 \cup \mathcal{C}_2$]

(e.g., LDA, SVM, PCA,
K-Means, K-NN, ...)

The Big Picture (an easy classification problem)



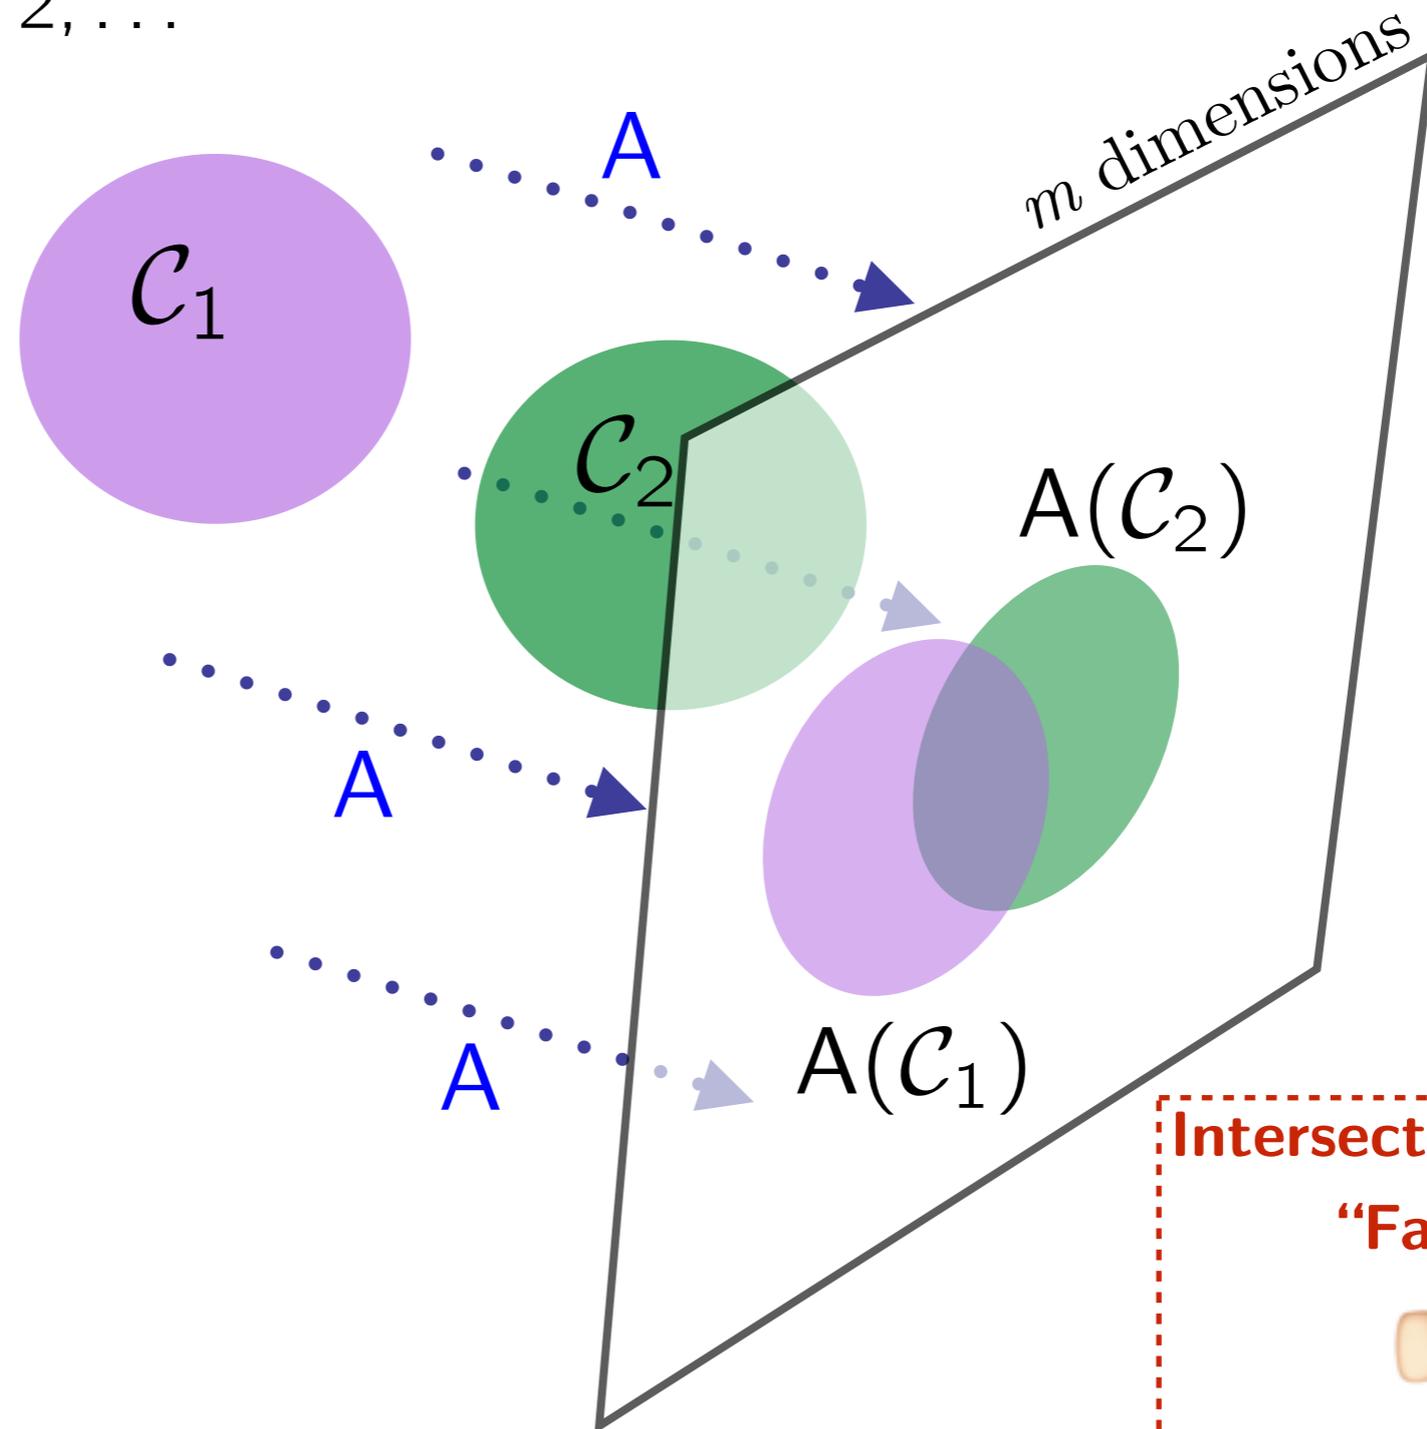
The Big Picture

$\mathcal{K} \subset \mathbb{R}^n$ dataset

$\mathcal{C}_i \subset \mathcal{K}$ classes, $i = 1, 2, \dots$

Non-Adaptive
Mapping

$$y = A(x)$$



Intersecting/Eclipse
"Failure"



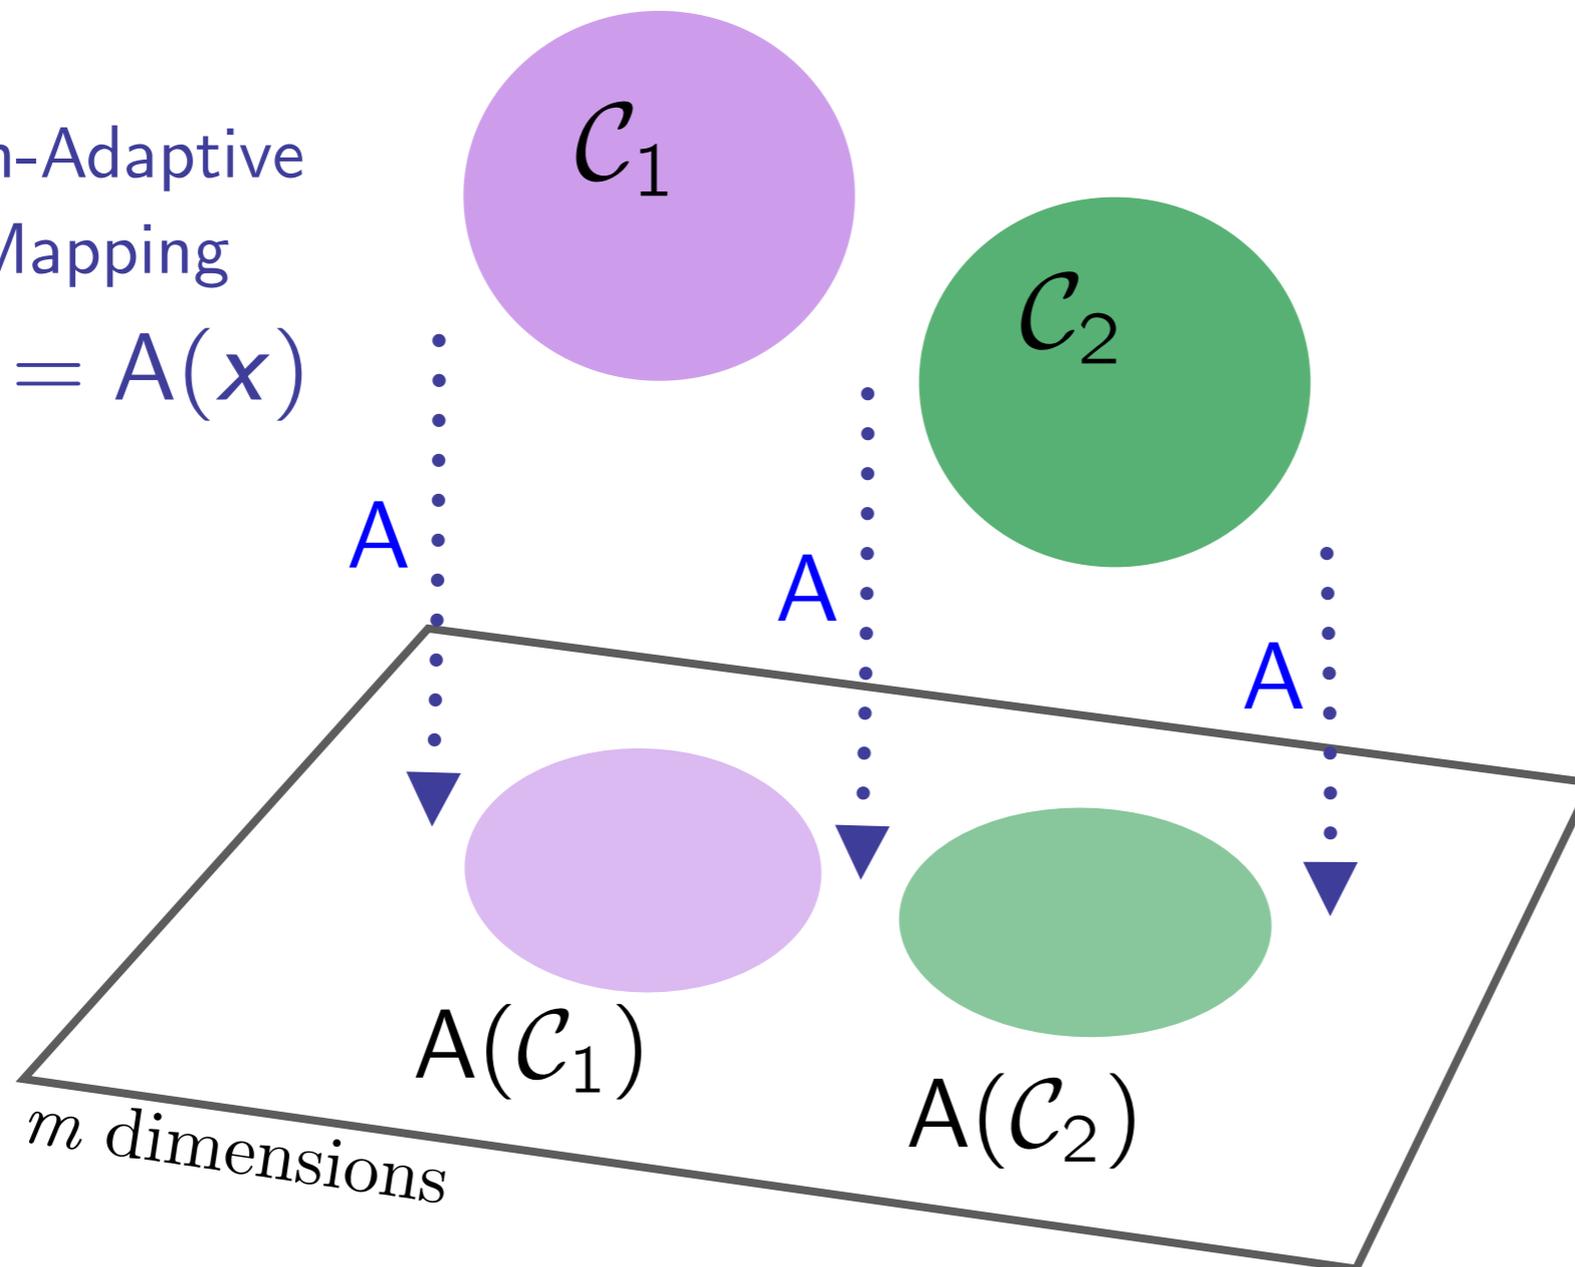
The Big Picture

$\mathcal{K} \subset \mathbb{R}^n$ dataset

$\mathcal{C}_i \subset \mathcal{K}$ classes, $i = 1, 2, \dots$

Non-Adaptive
Mapping

$$y = A(x)$$



Separable, "Success"



The Rare Eclipse Problem (Linear case)

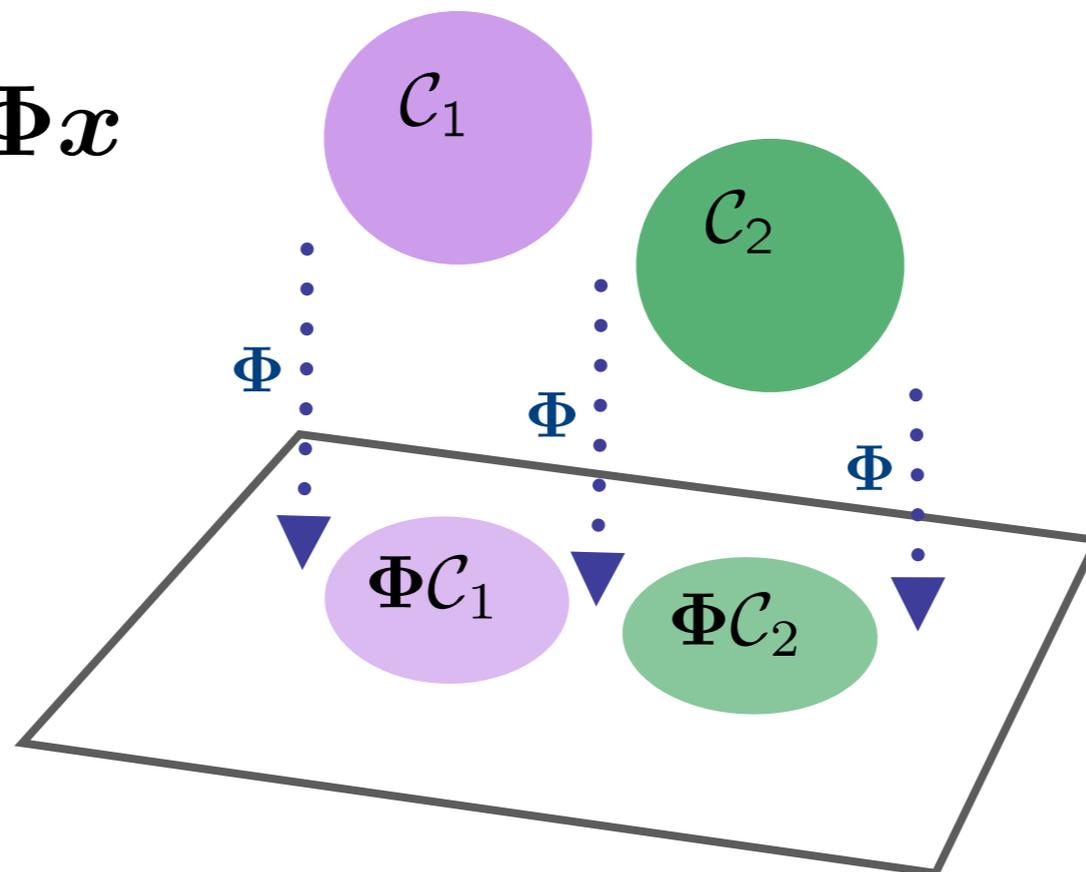


Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$. Given $\eta \in (0, 1)$, find the smallest m so that

$$p_0 := \mathbb{P}_{\Phi}[\Phi\mathcal{C}_1 \cap \Phi\mathcal{C}_2 = \emptyset] \geq 1 - \eta.$$

$$A(\mathbf{x}) = \Phi\mathbf{x}$$



The Rare Eclipse Problem (Linear case)

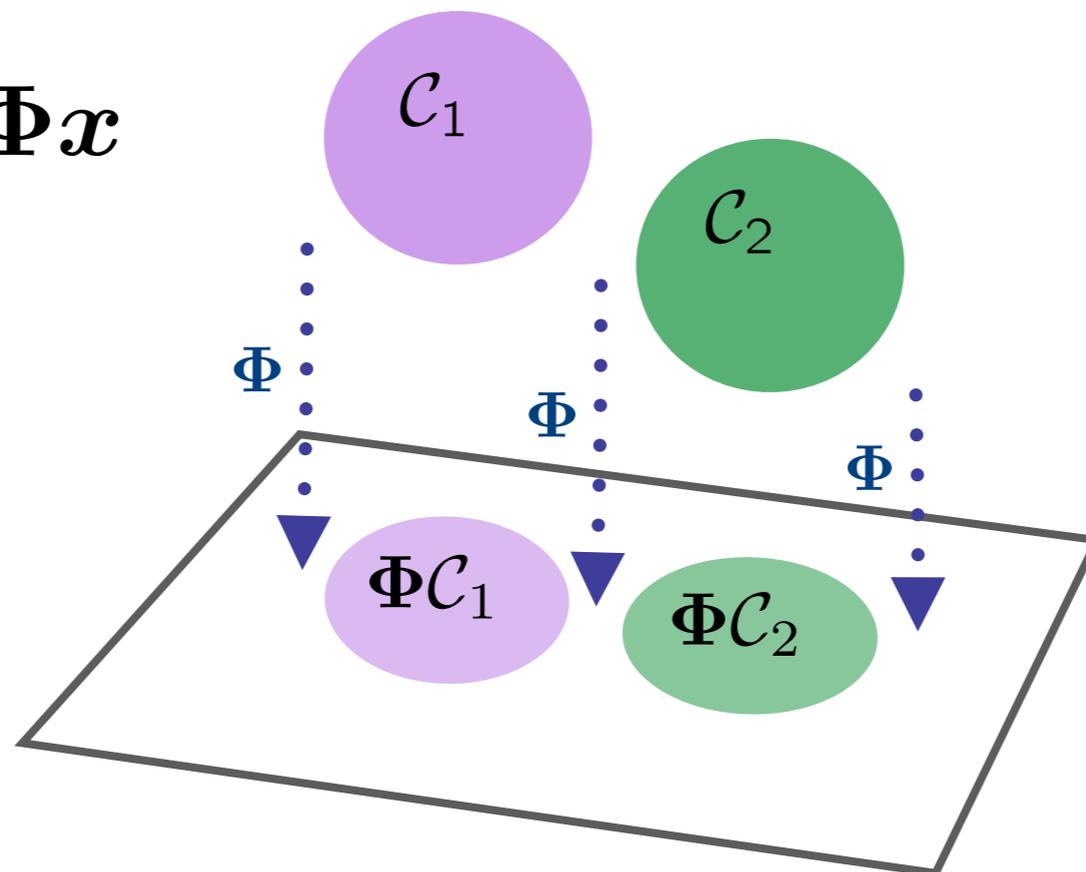


Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$.
Given $\eta \in (0, 1)$, find the smallest m so that

$$p_0 = \mathbb{P}_{\Phi}[\forall \mathbf{x}_1 \in \mathcal{C}_1, \forall \mathbf{x}_2 \in \mathcal{C}_2, \Phi(\mathbf{x}_1 - \mathbf{x}_2) \neq 0] \geq 1 - \eta.$$

$$A(\mathbf{x}) = \Phi \mathbf{x}$$



The Rare Eclipse Problem (Linear case)



Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$.
Given $\eta \in (0, 1)$, find the smallest m so that, with $\mathcal{C}^\ominus = \mathcal{C}_1 - \mathcal{C}_2$,

$$p_0 = \mathbb{P}_\Phi[\mathcal{C}^\ominus \cap \ker \Phi = \emptyset] = \mathbb{P}_\Phi[\mathcal{S} \cap \ker \Phi = \emptyset] \geq 1 - \eta.$$

$$\mathcal{S} = (\mathbb{R}_+ \mathcal{C}^\ominus) \cap \mathbb{S}_2^{n-1}$$

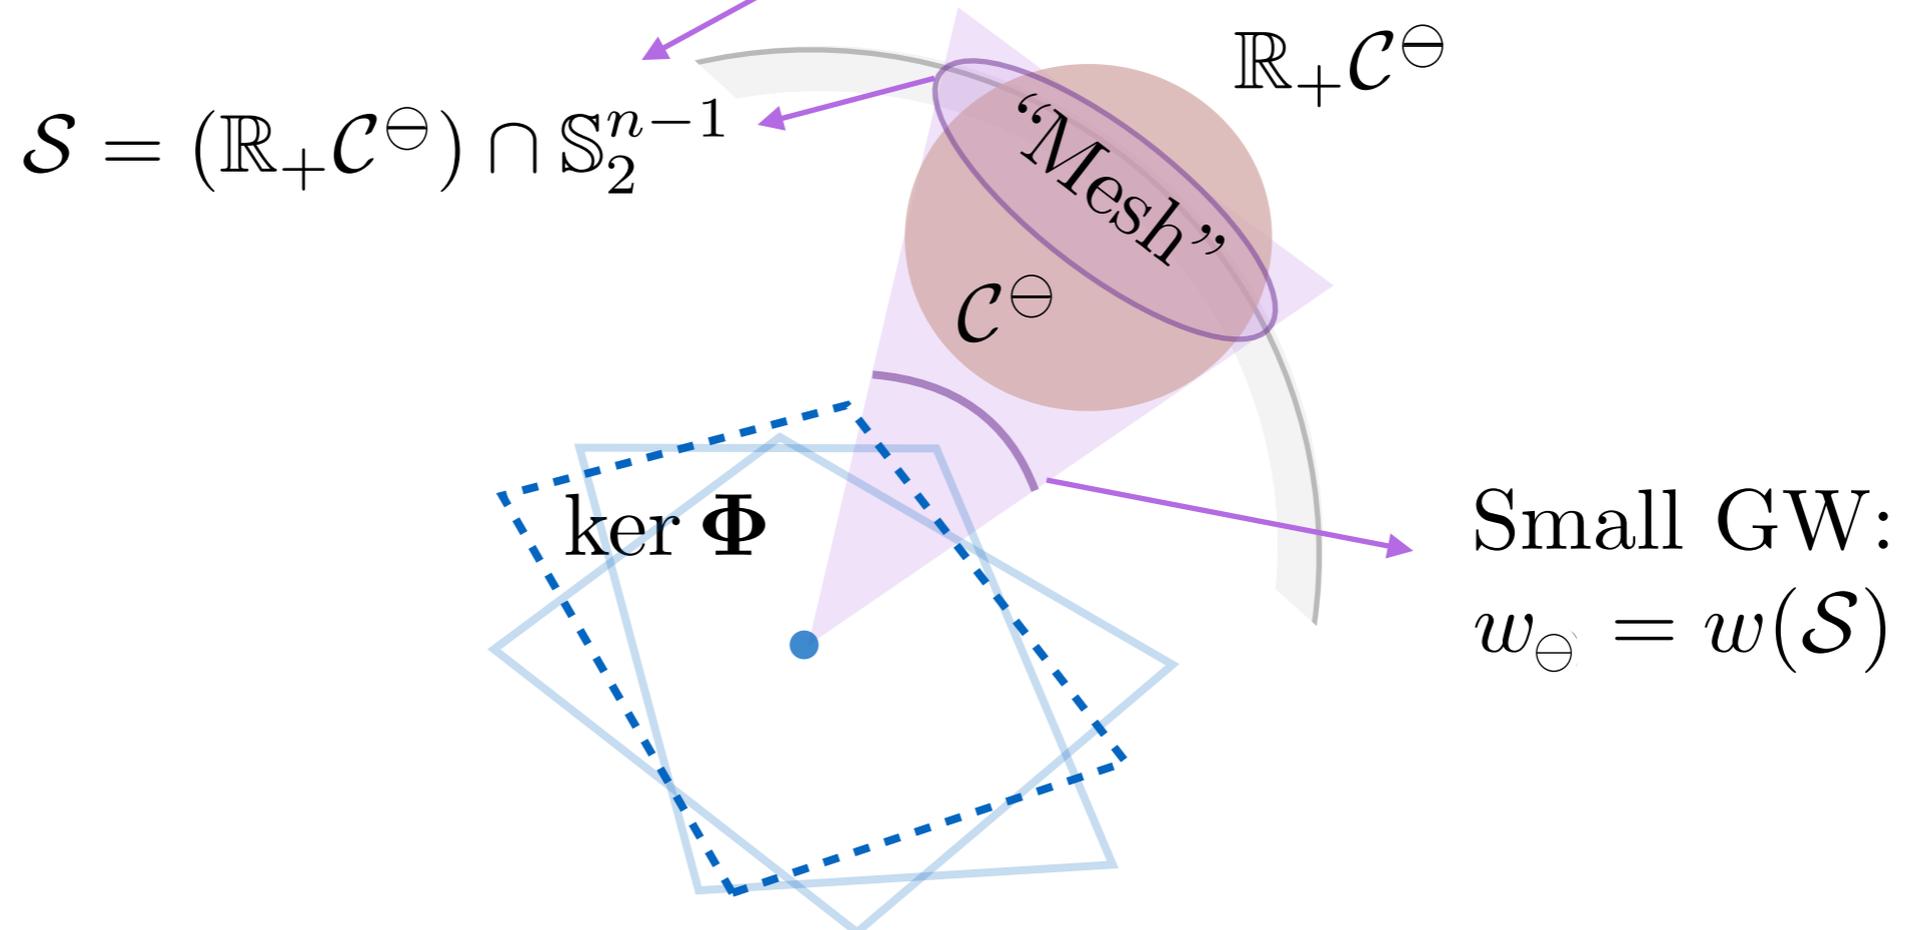
The Rare Eclipse Problem (Linear case)



Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$. Given $\eta \in (0, 1)$, find the smallest m so that, with $\mathcal{C}^\ominus = \mathcal{C}_1 - \mathcal{C}_2$,

$$p_0 = \mathbb{P}_\Phi[\mathcal{C}^\ominus \cap \ker \Phi = \emptyset] = \mathbb{P}_\Phi[\mathcal{S} \cap \ker \Phi = \emptyset] \geq 1 - \eta.$$



The Rare Eclipse Problem (Linear case)



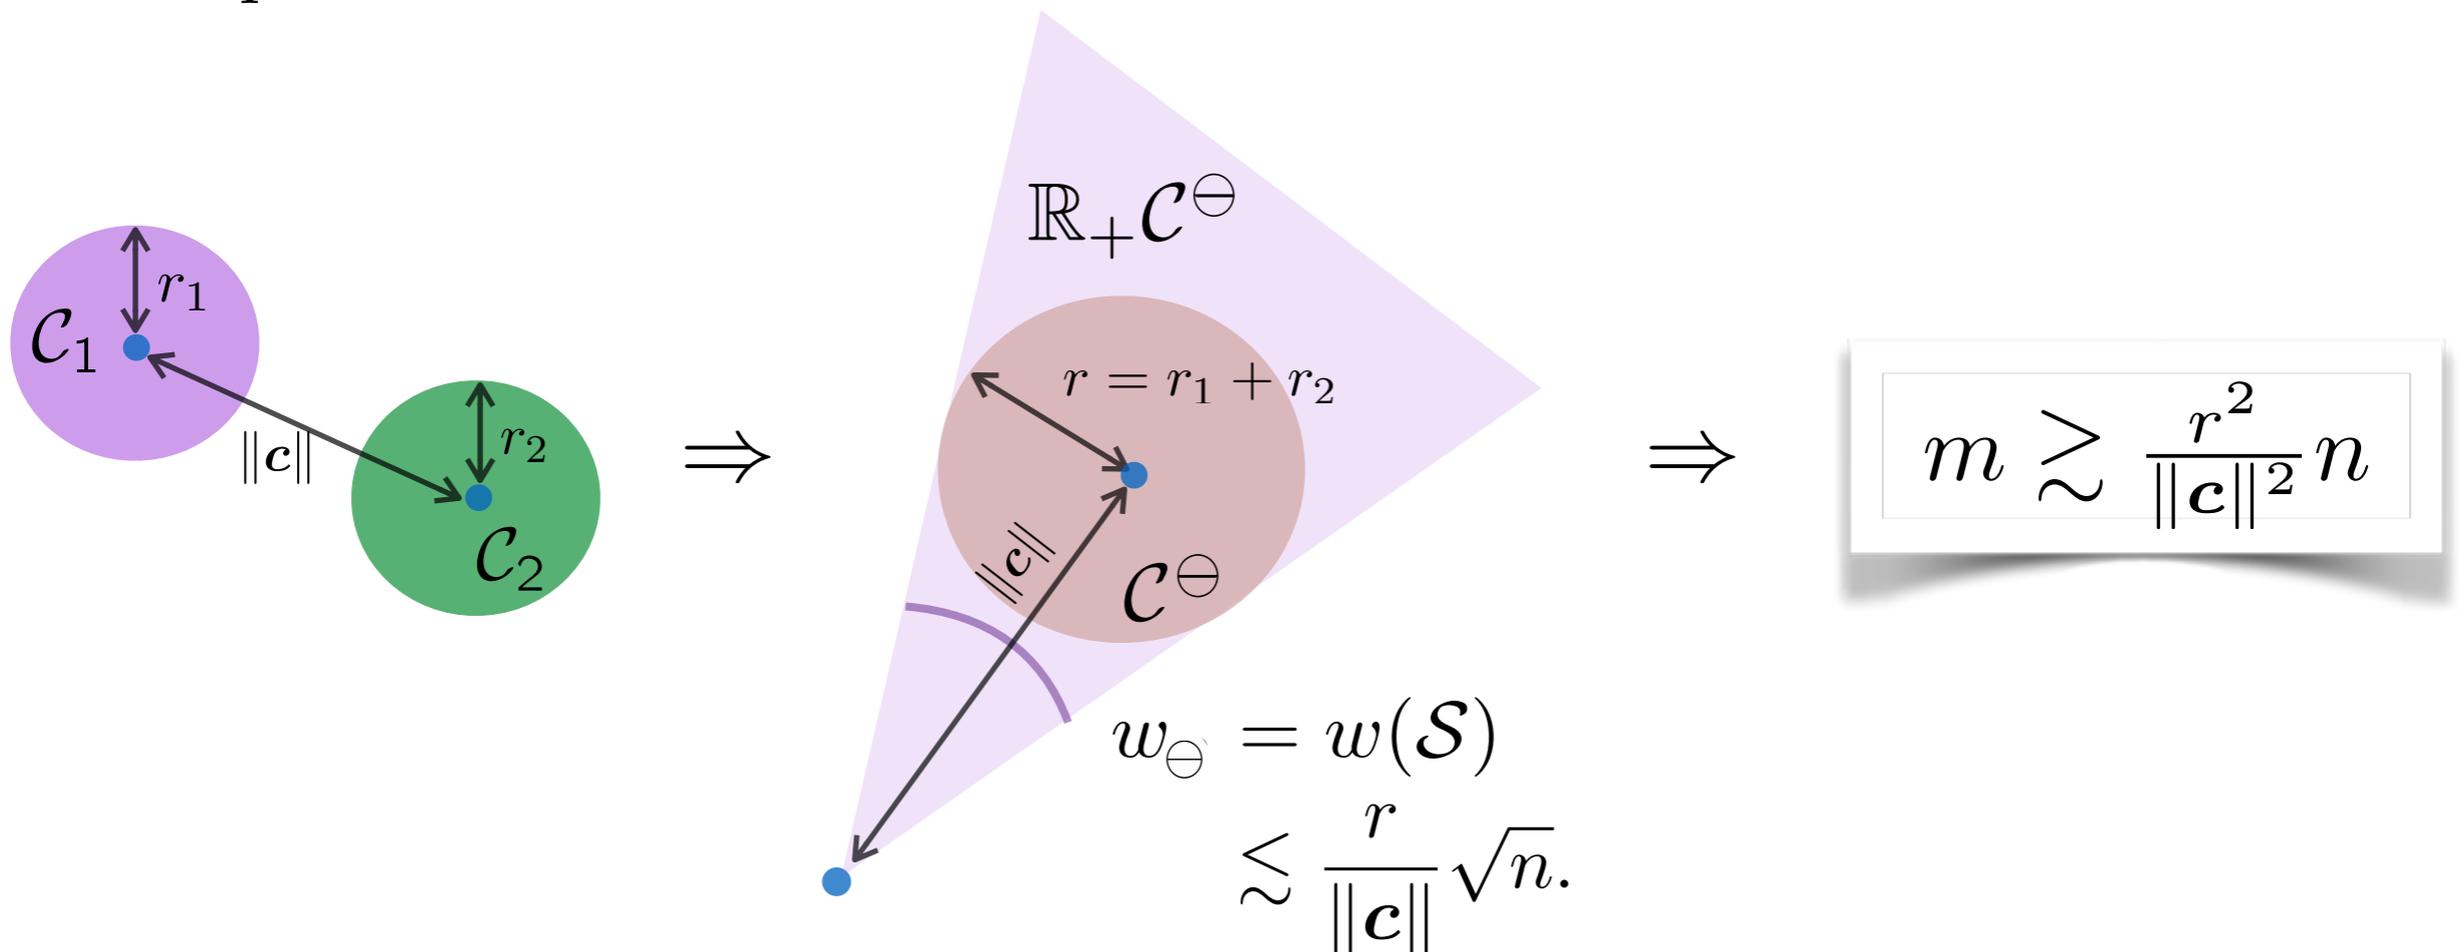
BMR '14: “Gordon’s escape through a mesh” theorem

Proposition (Corollary 3.1 in BMR '14).

(*& really tight* [Amelunxen et al, 13])

Given $\eta \in (0, 1)$, if $m > (w_{\ominus} + \sqrt{2 \log \frac{1}{\eta}})^2 + 1$ then $p_0 \geq 1 - \eta$.

Example:



The Rare Eclipse Problem (alternate proof)

Restricted Isometry Property: (ℓ_1, ℓ_2) -RIP(\mathcal{K}, ϵ)

$$\forall \mathbf{x} \in \mathcal{K}, (1 - \epsilon) \|\mathbf{x}\| \leq \|\Phi \mathbf{x}\|_1 \leq (1 + \epsilon) \|\mathbf{x}\|$$

[Schechtman, 06] [Plan, Vershynin, 14]

If $\mathcal{S} \subset \mathbb{S}^{n-1}$ and $m \gtrsim \epsilon^{-2} w^2(\mathcal{S})$, then, w.h.p.*,

$$(1 - \epsilon) \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{u}\|_1 \leq (1 + \epsilon)$$

for $\Phi \in \mathbb{R}^{m \times n}$ and $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$.

(note: extendable to subsets of \mathbb{B}^n ,
e.g., compressible signals, [Xu, LJ, 18])

*: i.e., $\mathbb{P} \geq 1 - C \exp(-c\epsilon^2 m)$.

The Rare Eclipse Problem (alternate proof)

Restricted Isometry Property: (ℓ_1, ℓ_2) -RIP(\mathcal{K}, ϵ)

$$\forall \mathbf{x} \in \mathcal{K}, (1 - \epsilon) \|\mathbf{x}\| \leq \|\Phi \mathbf{x}\|_1 \leq (1 + \epsilon) \|\mathbf{x}\|$$

[Schechtman, 06] [Plan, Vershynin, 14]

If $\mathcal{S} \subset \mathbb{S}^{n-1}$ and $m \gtrsim \epsilon^{-2} w^2(\mathcal{S})$, then, w.h.p*,

$$(1 - \epsilon) \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{u}\|_1 \leq (1 + \epsilon)$$

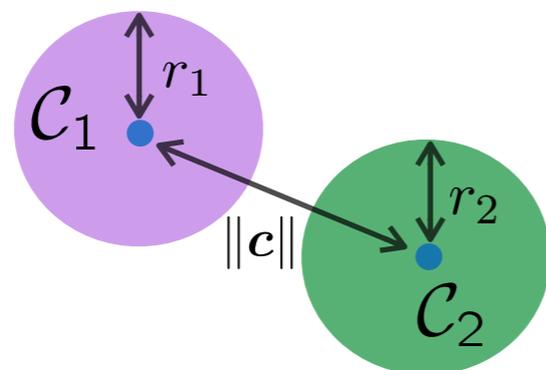
for $\Phi \in \mathbb{R}^{m \times n}$ and $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$.

Therefore:

For $\mathcal{S} = (\mathbb{R}^+ \mathcal{C}^\ominus) \cap \mathbb{S}^{n-1}$, if $m \gtrsim \epsilon^{-2} w_\ominus^2$, w.h.p*, (RIP $_\ominus$)

$$(1 - \epsilon) \|\mathbf{x}_1 - \mathbf{x}_2\| \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{x}_1 - \Phi \mathbf{x}_2\|_1 \leq (1 + \epsilon) \|\mathbf{x}_1 - \mathbf{x}_2\|,$$

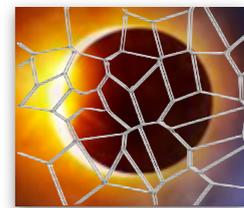
for all $\mathbf{x}_1 \in \mathcal{C}_1$ and all $\mathbf{x}_2 \in \mathcal{C}_2$.



This result also explains REP,
e.g., with $\epsilon = 1/2$, (but less sharply)

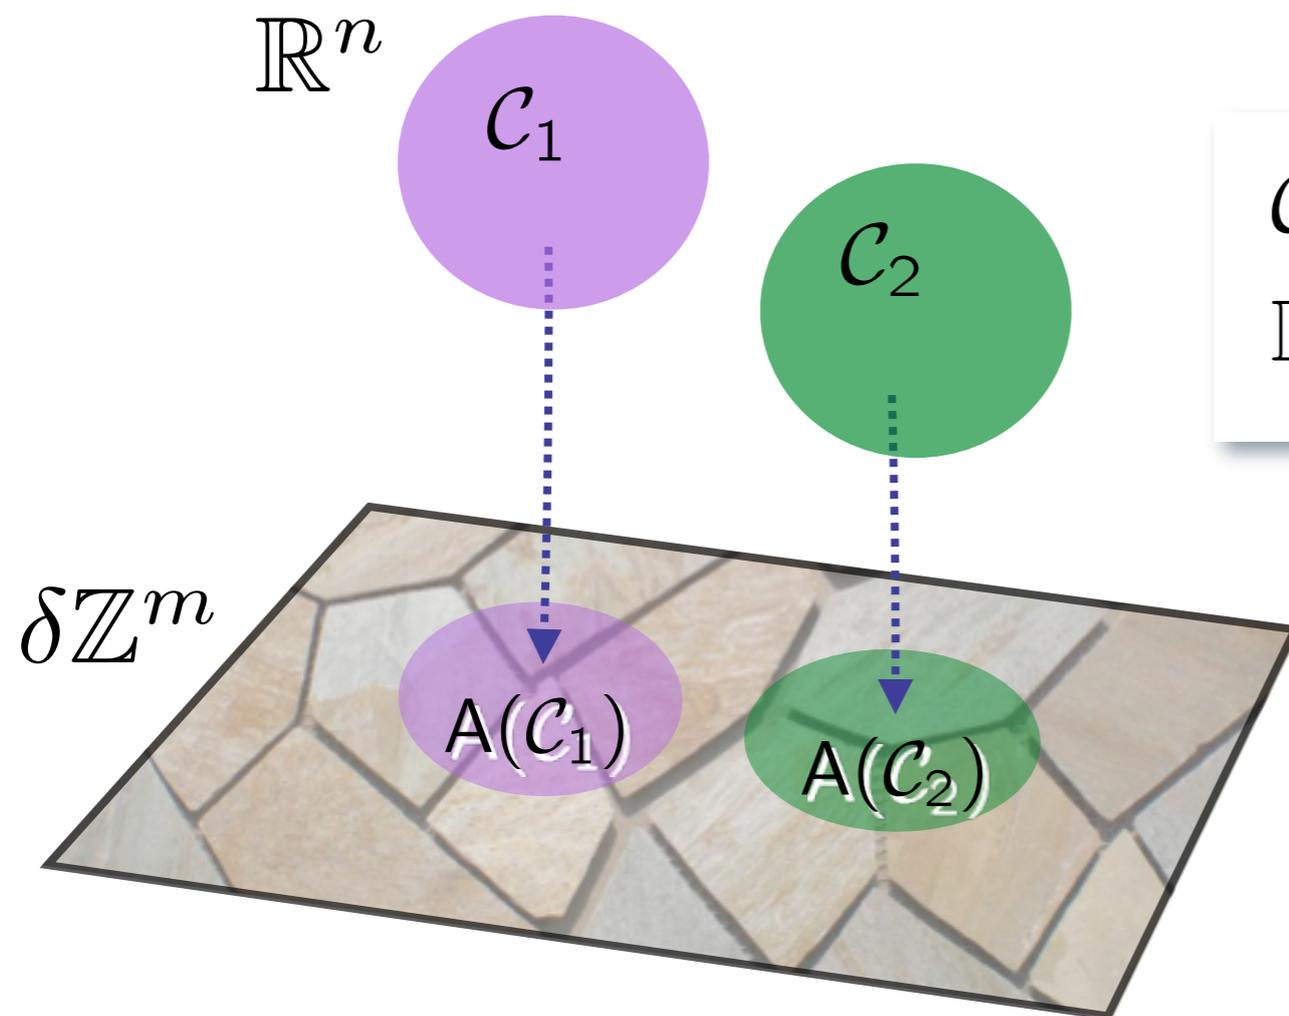
*: i.e., $\mathbb{P} \geq 1 - C \exp(-c\epsilon^2 m)$.

The Rare Eclipse Problem “on Tiles”



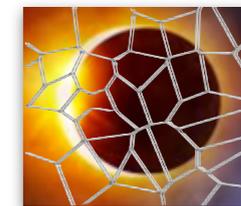
$$A(x) := Q(\Phi x + \xi)$$

with Φ Gaussian random matrix,
 $Q(\lambda) = \delta \lfloor \frac{\lambda}{\delta} \rfloor$, $\xi_i \sim \mathcal{U}([0, \delta])$.



$\mathcal{C}_1, \mathcal{C}_2, m$ and δ such that
 $\mathbb{P}[A(\mathcal{C}_1) \cap A(\mathcal{C}_2) = \emptyset] \geq 1 - \eta$?

The Rare Eclipse Problem “on Tiles”



Given $\sigma := \min_{\mathbf{z} \in \mathcal{C}^\ominus} \|\mathbf{z}\|$ and $w_\cap = w((\mathbb{R}_+ \mathcal{C}^\ominus) \cap \mathbb{S}^{n-1})$.

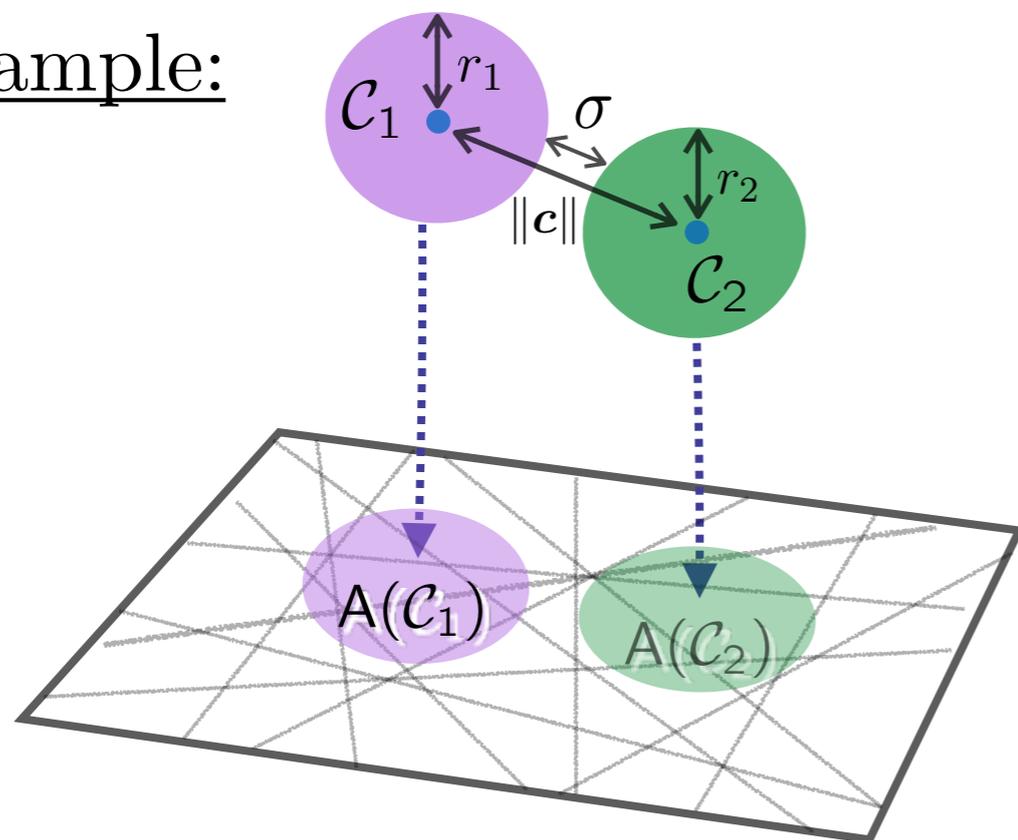
Provided

$$m \gtrsim \left(\underbrace{w_\ominus^2}_{\text{linear}} + \underbrace{n \frac{\delta^2}{\sigma^2}}_{\text{quantiz.}} \right) \left(1 + \underbrace{\log \left(1 + \frac{rm}{\delta n} \right)}_{\text{proof artifact?}} \right) + \underbrace{w_\ominus^{-2} \log \frac{1}{\eta}}_{\text{linear}},$$

we have

$$\mathbb{P}[A(\mathcal{C}_1) \cap A(\mathcal{C}_2) = \emptyset] \geq 1 - \eta.$$

Example:



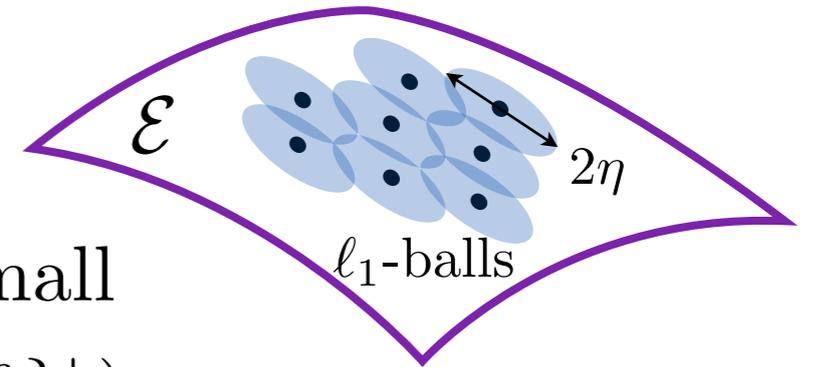
$$\Rightarrow m \gtrsim \left(\frac{r^2}{\|\mathbf{c}\|^2} + \frac{\delta^2}{(\|\mathbf{c}\| - r)^2} \right) n$$

Note: $\delta > \sigma$ is allowed (dithering effect!)
 Note bis: $m > n$ not specially bad ($\delta \mathbb{Z}^m$).

Proof sketch:

A. Embedding $\mathcal{E} \subset \mathbb{R}^m$ into $\delta\mathbb{Z}^m$ with

$$A'(\mathbf{y}) := \underset{\uparrow \delta}{\mathcal{Q}}(\mathbf{y} + \underset{\uparrow \delta}{\boldsymbol{\xi}})$$



If \mathcal{E} has small ℓ_1 -Kolmogorov entropy, *i.e.*, small

$$\mathcal{H}_1(\mathcal{E}, \eta) = \log(\min_{\mathcal{G}} |\{\mathcal{G} : \mathcal{G} \subset \mathcal{E} \subset \mathcal{S} + \eta\mathbb{B}_1^n\}|)$$

Given $\epsilon > 0$, if $m \gtrsim \epsilon^{-2} \mathcal{H}_1(\mathcal{E}, \frac{m\delta\epsilon^2}{1+\epsilon})$, then, w.h.p*, for all $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{E}$ and some $c > 0$,

(P1)

$$\frac{1}{m} \|\mathbf{y}_1 - \mathbf{y}_2\|_1 - c\delta\epsilon \leq \frac{1}{m} \|A'(\mathbf{y}_1) - A'(\mathbf{y}_2)\|_1 \leq \frac{1}{m} \|\mathbf{y}_1 - \mathbf{y}_2\|_1 + c\delta\epsilon.$$

B. For $\mathcal{E} = \Phi\mathcal{K}$ and Φ an (ℓ_1, ℓ_2) -RIP($\mathcal{K} - \mathcal{K}, \epsilon' < 1$),

$$\mathcal{H}_1(\mathcal{E}, 2m\eta) \leq \mathcal{H}_2(\mathcal{K}, \eta) \quad (\text{P2})$$

with \mathcal{H}_2 is bounded for sets (cones, convex spaces, ...)

*: [LJ, Cambareri, '16] [Cambareri, Xu, LJ, '17]

Proof sketch:

C. We apply A and B for $\mathcal{E} = \Phi\mathcal{C}_\cup$ so that (P1) holds if

$$m \gtrsim \epsilon^{-2} n \log\left(1 + \frac{2r(1+\epsilon)^2}{\delta\epsilon^2}\right)$$

D. We rely on (RIP_\ominus) , as in the linear case, i.e.,

For $\mathcal{S} = (\mathbb{R}^+\mathcal{C}^\ominus) \cap \mathbb{S}^{n-1}$, if $m \gtrsim \epsilon_0^{-2} w_\ominus^2$, w.h.p*,

$$(1 - \epsilon_0) \|\mathbf{x}_1 - \mathbf{x}_2\| \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi\mathbf{x}_1 - \Phi\mathbf{x}_2\|_1 \leq (1 + \epsilon_0) \|\mathbf{x}_1 - \mathbf{x}_2\|,$$

for all $\mathbf{x}_1 \in \mathcal{C}_1$ and all $\mathbf{x}_2 \in \mathcal{C}_2$.

E. For $\mathbf{x}_1 \in \mathcal{C}_1$, $\mathbf{x}_2 \in \mathcal{C}_2$,

$$\begin{aligned} \frac{1}{m} \|A'(\Phi\mathbf{x}_1) - A'(\Phi\mathbf{x}_2)\|_1 &\geq \frac{1}{m} \|\Phi(\mathbf{x}_1 - \mathbf{x}_2)\|_1 - c\delta\epsilon \\ &\geq \left(\frac{2}{\pi}\right)^{1/2} (1 - \epsilon_0) \|\mathbf{x}_1 - \mathbf{x}_2\| - c\delta\epsilon \\ &\geq \left(\frac{2}{\pi}\right)^{1/2} (1 - \epsilon_0) \sigma - c\delta\epsilon. \end{aligned}$$

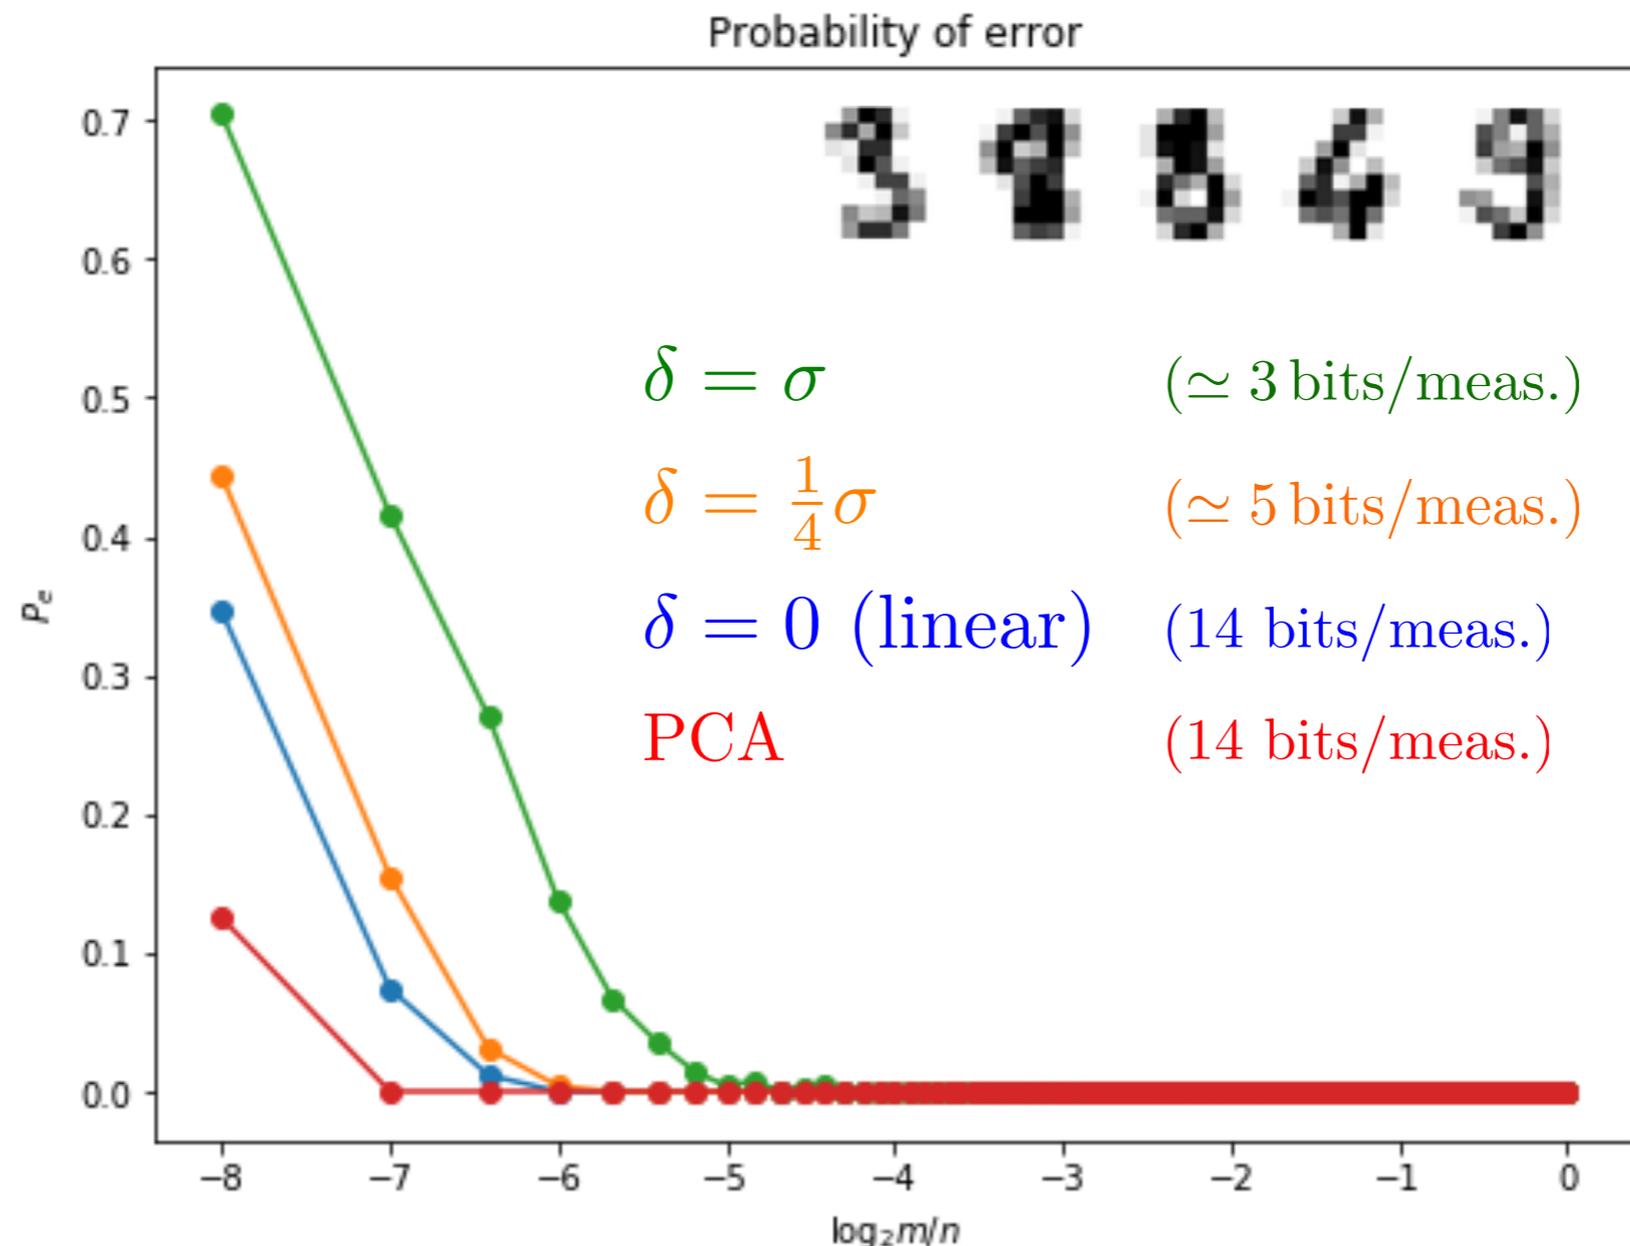
Fixing $\epsilon = \frac{\sqrt{n}}{w_\cap} \epsilon_0$ (to get similar conditions), and forcing $\left(\frac{2}{\pi}\right)^{1/2} (1 - \epsilon_0) \sigma - c\delta\epsilon > c\delta\epsilon > 0$ (for separability) gives the result (after solving for ϵ_0). □

Simulations: Digit dataset (from scikit learn)

10 handwritten digits, 8x8 pixels ($n=64$), samples/class ≈ 12 .

Training/Test sets = 50%/50%. $\sigma = \min_{i,j:i \neq j} \min_{\mathbf{u} \in \mathcal{C}_i, \mathbf{v} \in \mathcal{C}_j} \|\mathbf{u} - \mathbf{v}\|$

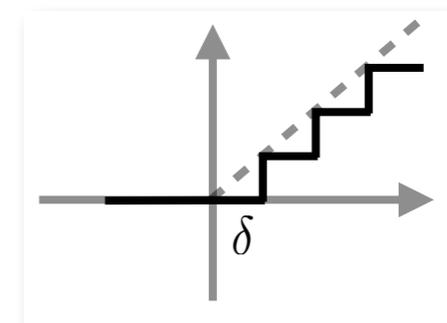
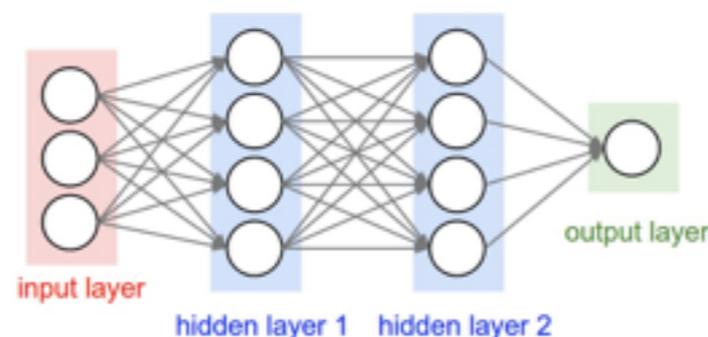
Classification: 5-NN Classifier.



Try some code out here: github.com/VC86/MLSPbox

Take-away messages

- ▶ From CS to QCS (for scalar quantizers)
- ▶ Importance of consistency and dithering
- ▶ Reconstruction still possible in QCS with decaying error as m increases
- ▶ Learning/Classification possible in QCS domain
- ▶ Open problems
- ▶ CW for (other/all?) RIP matrices?
- ▶ Quantizing non-linear embedding (clipping, ReLU)?
- ▶ ...



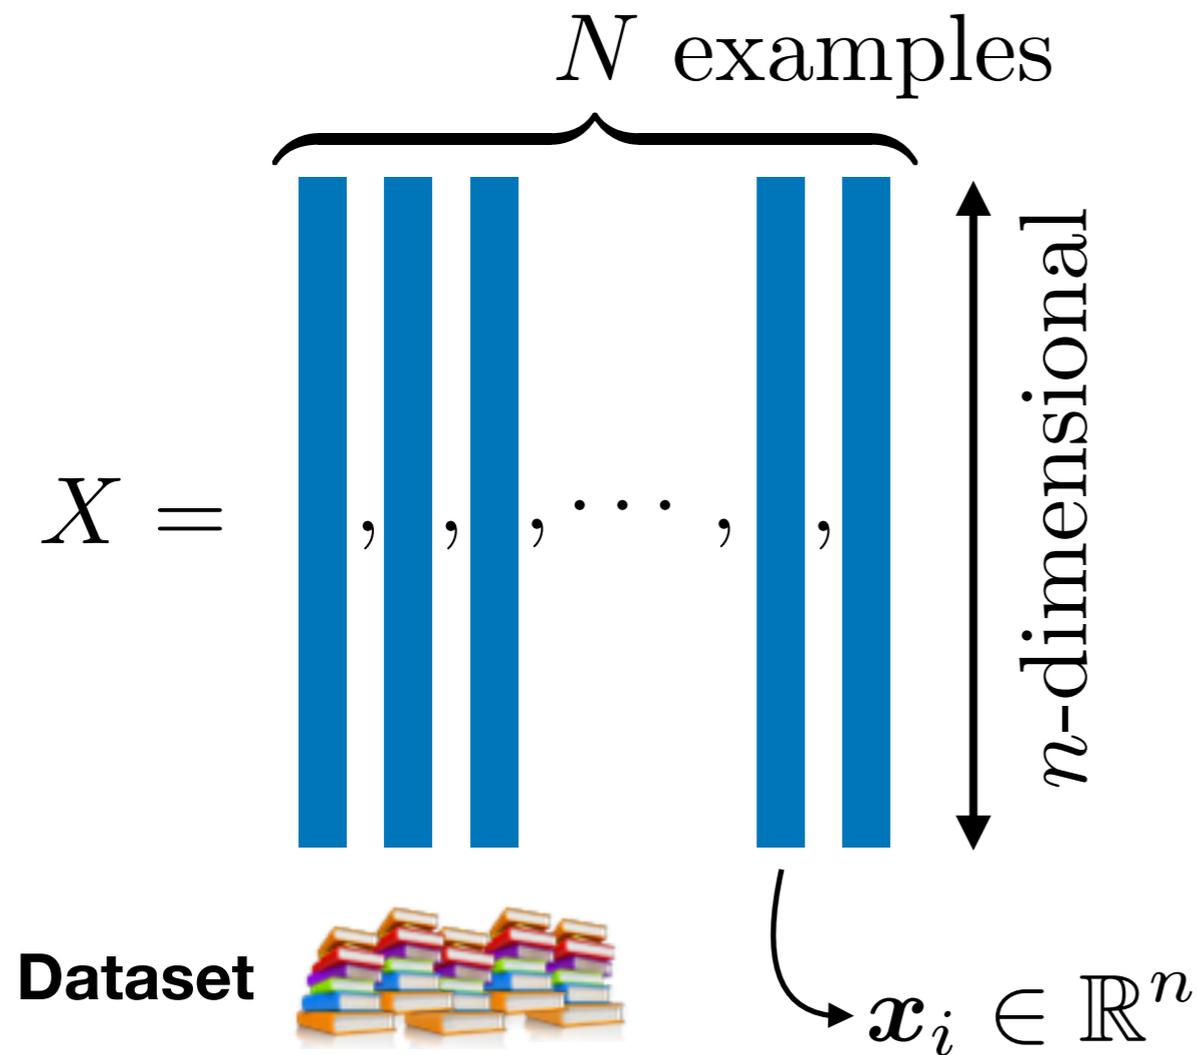
Thank you for your attention!

- ▶ N. Thao and M. Vetterli, **“Deterministic analysis of oversampled A/D conversion and decoding improvement based on consistent estimates,”** Signal Processing, IEEE Transactions on Signal Processing, vol. 42, no. 3, pp. 519–531, 1994.
- ▶ E. J. Candes and T. Tao, **“Decoding by linear programming,”** IEEE transactions on information theory, vol. 51, no. 12, pp. 4203–4215, 2005.
- ▶ Y. Plan and R. Vershynin, **“Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach,”** IEEE Transactions on Information Theory, vol. 59, no. 1, pp. 482–494, 2013.
- ▶ LJ, **“Error Decay of (almost) Consistent Signal Estimations from Quantized Gaussian Random Projections”**, IEEE Transactions in Information Theory, Vol. 62, no.8, p. 4696-4709, 2016
- ▶ LJ and V. Cambareri, **“Time for dithering: fast and quantized random embeddings via the restricted isometry property,”** Information and Inference: A Journal of the IMA, p. iax004, 2017.
- ▶ LJ, **“Small width, low distortions: quantized random embeddings of low-complexity sets,”** IEEE Transactions on information theory, vol. 63, no. 9, pp. 5477–5495, 2015.
- ▶ V. Chandrasekaran, B. Recht, P.A. Parrilo, and A. S. Willsky, **“The convex geometry of linear inverse problems,”** Foundations of Computational mathematics, vol. 12, no. 6, pp. 805–849, 2012.
- ▶ Foundations of Computational mathematics, vol. 12, no. 6, pp. 805–849, 2012.
- ▶ R. Gribonval, G. Blanchard, N. Keriven, Y. Traonmilin. **“Compressive Statistical Learning with Random Feature Moments”**, arXiv:1706.07180
- ▶ P. T. Boufounos, **“Universal rate-efficient scalar quantization,”** IEEE Transactions on Information Theory, vol.58, no. 3, pp. 1861–1872, 2012.
- ▶ G. Schechtman, **“Two observations regarding embedding subsets of Euclidean spaces in normed spaces”**, Adv. Math. 200 (2006), 125–135.

Backup slides

Quantized Random Sketching of Datasets

Compressing a dataset?



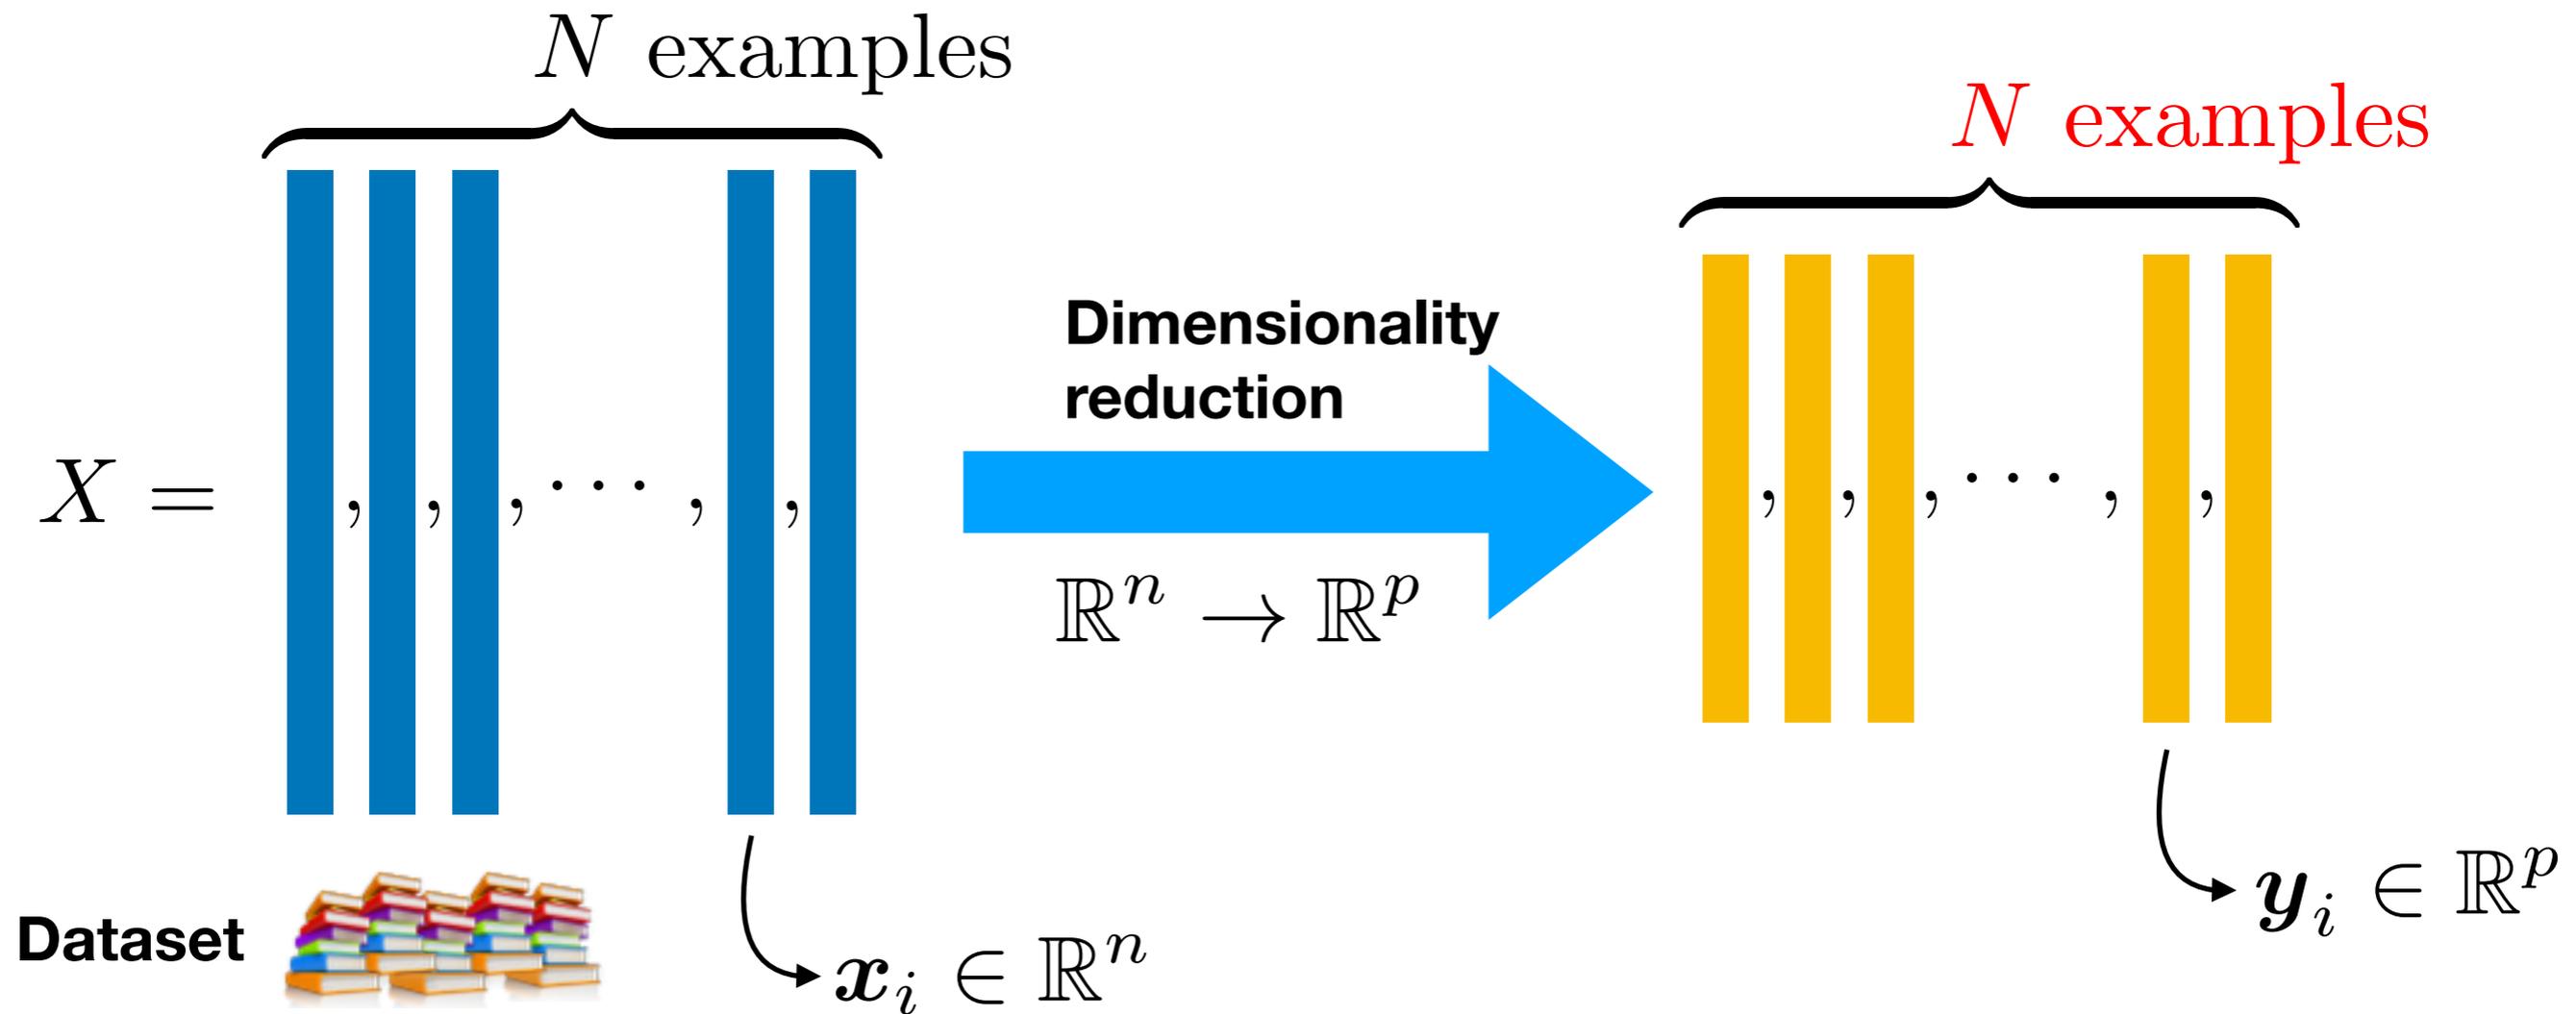
Large datasets means:

- ▶ Large memory required
- ▶ Slow learning algorithm
(e.g., K-Means, (K)SVM, ...)

Complexity:
 $O(nNK)$ per iteration

BUT extracted “knowledge” is “simple”
Do we really need all this data?

Compressing a dataset?



Classical CS methods

[Davenport et al. '07-'10], [Haupt et al. '06]

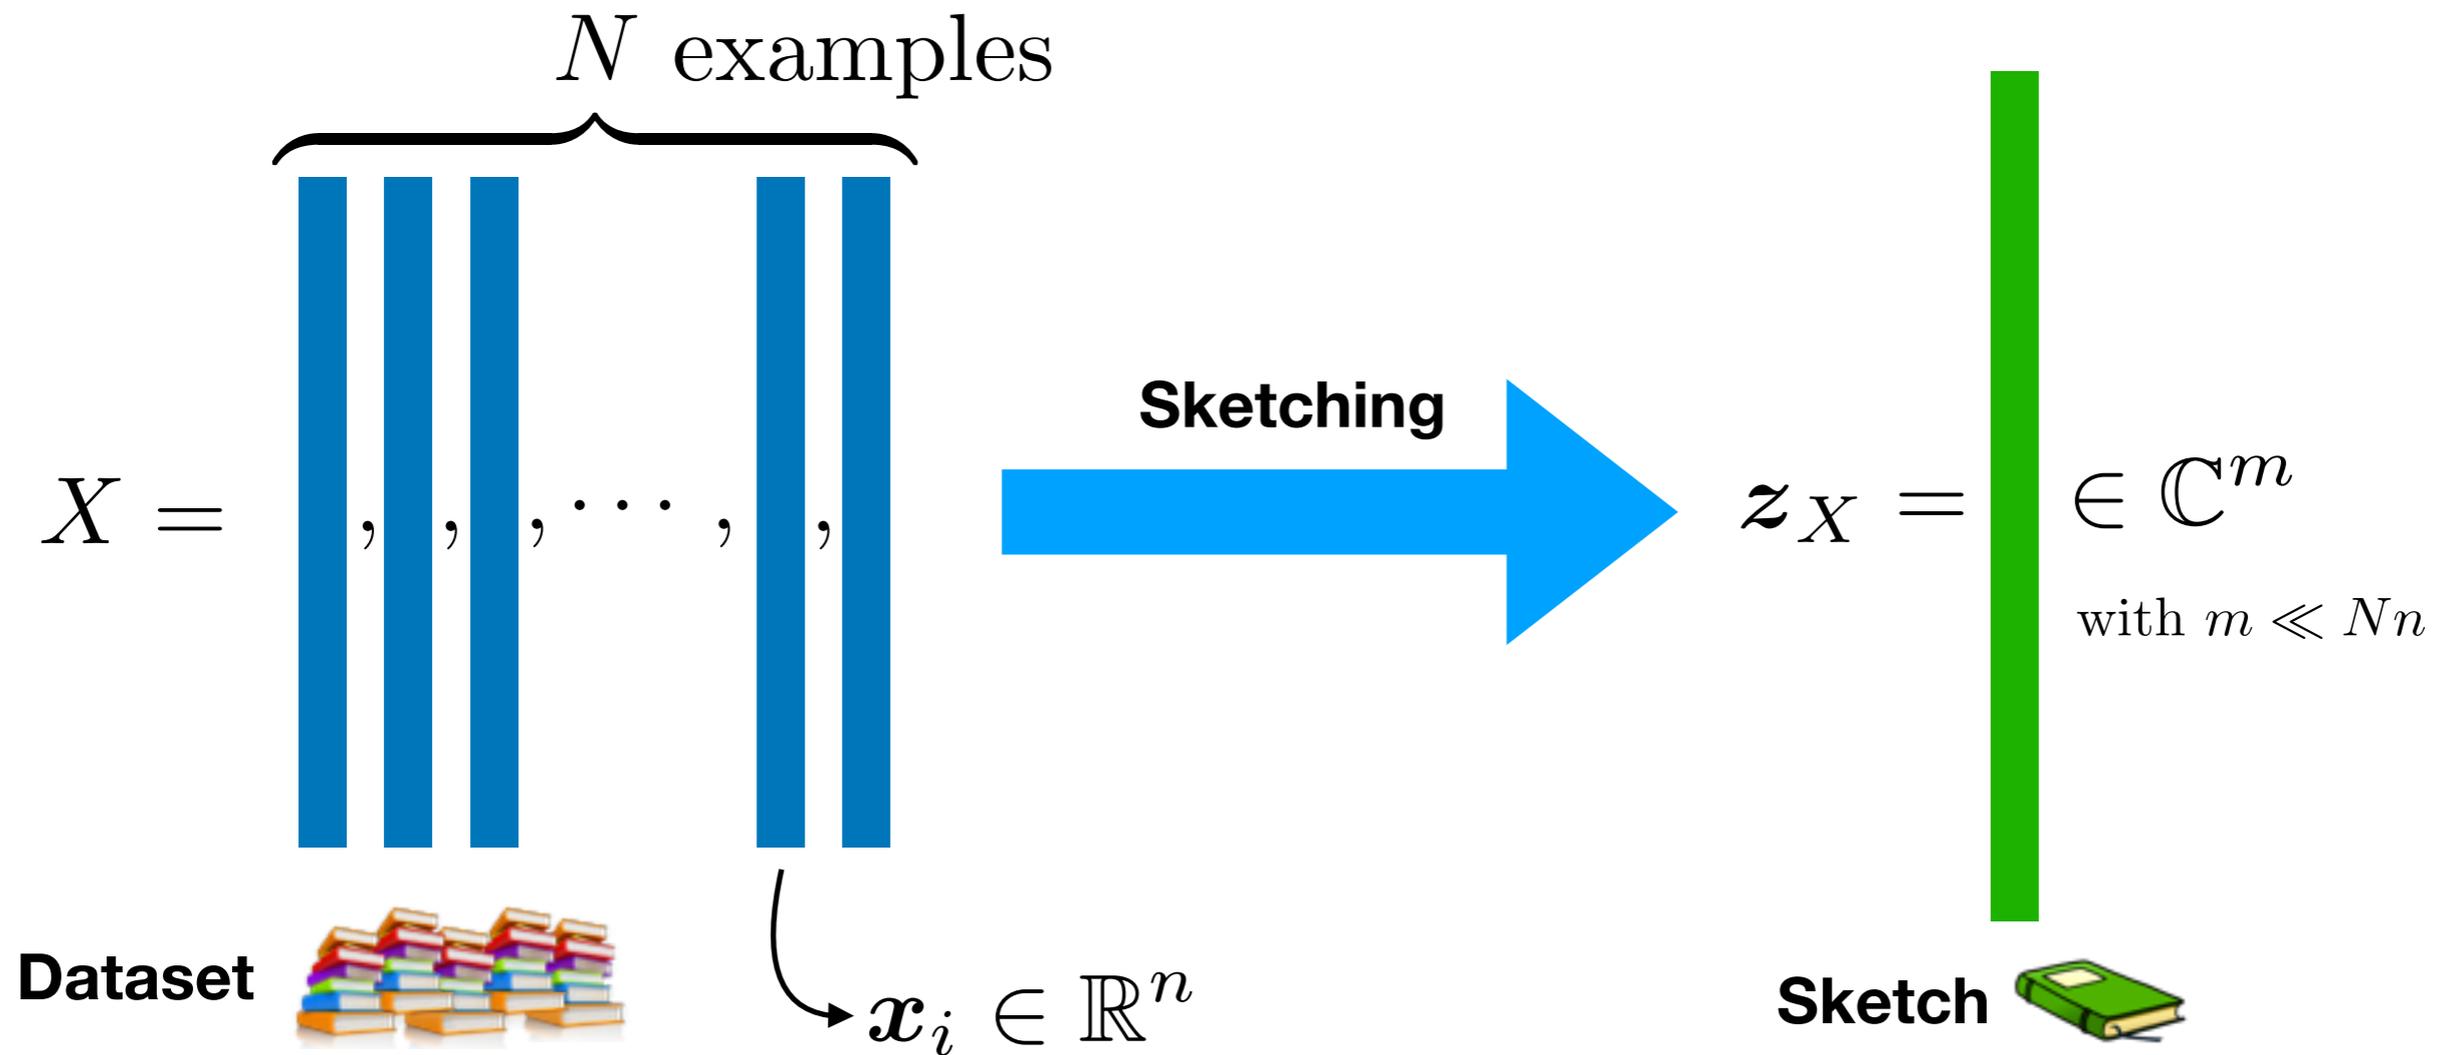
[Reboredo et al. '13-'16],

[Bandeira, Mixon, Recht '14]

- ▶ Compressed representation ✓
- ▶ Preserves relevant information ✓
- ▶ **Constant number of examples** ✗

N can be VERY large (“big data”)!

Compressing a dataset?



- ▶ Compressed representation ✓
- ▶ Preserves relevant information ✓
- ▶ Dataset summary = single vector ✓

[Keriven, Bourrier, Gribonval, Pérez, 16]

[Gribonval, Blanchard, Keriven, Traonmilin, 17]

Compressing a dataset?

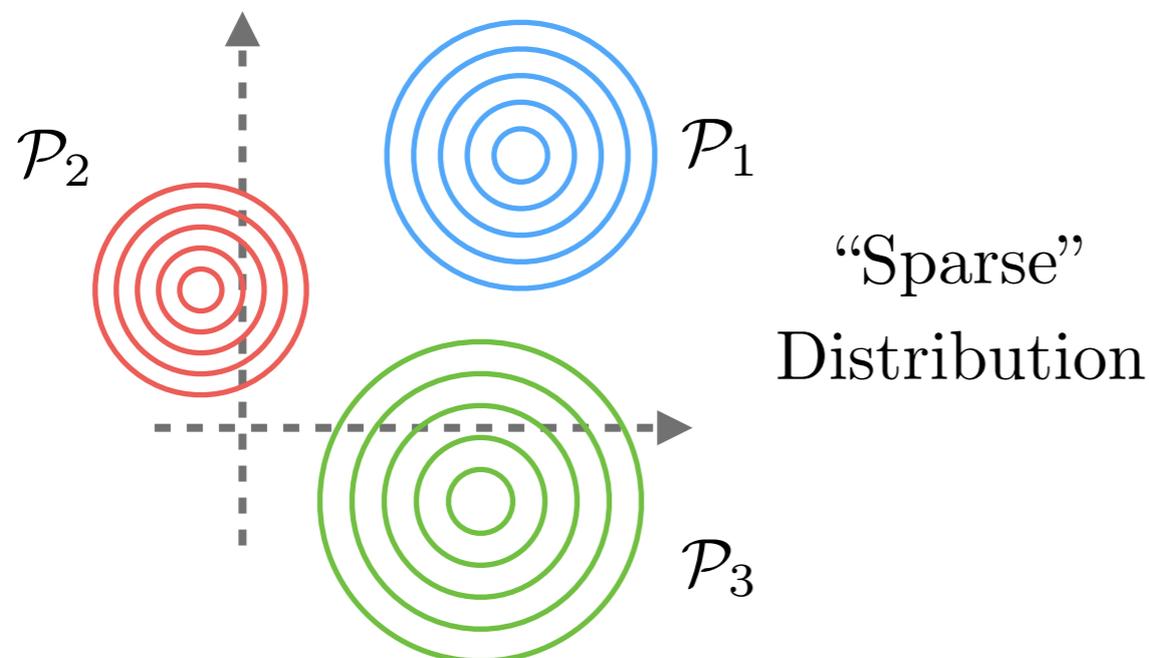
Dataset



$$\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset \mathbb{R}^N$$

Prior information:

$$\mathbf{x}_i \sim_{\text{iid}} \mathcal{P} = \sum_{k=1}^K \alpha_k \mathcal{P}_k$$



(e.g., Gaussian mixture)

[Keriven, Bourrier, Gribonval, Pérez, 16]

[Gribonval, Blanchard, Keriven, Traonmilin, 17]

Compressing a dataset?

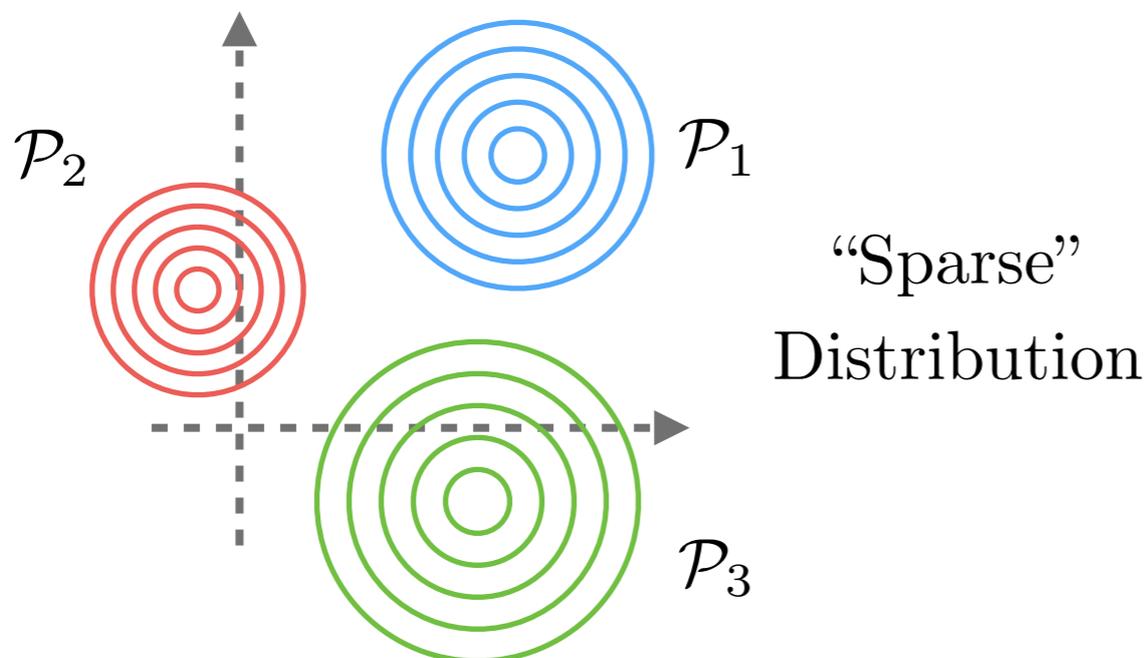
Dataset



$$\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset \mathbb{R}^N$$

Prior information:

$$\mathbf{x}_i \sim_{\text{iid}} \mathcal{P} = \sum_{k=1}^K \alpha_k \mathcal{P}_k$$



(e.g., Gaussian mixture)

[Keriven, Bourrier, Gribonval, Pérez, 16]

[Gribonval, Blanchard, Keriven, Traonmilin, 17]

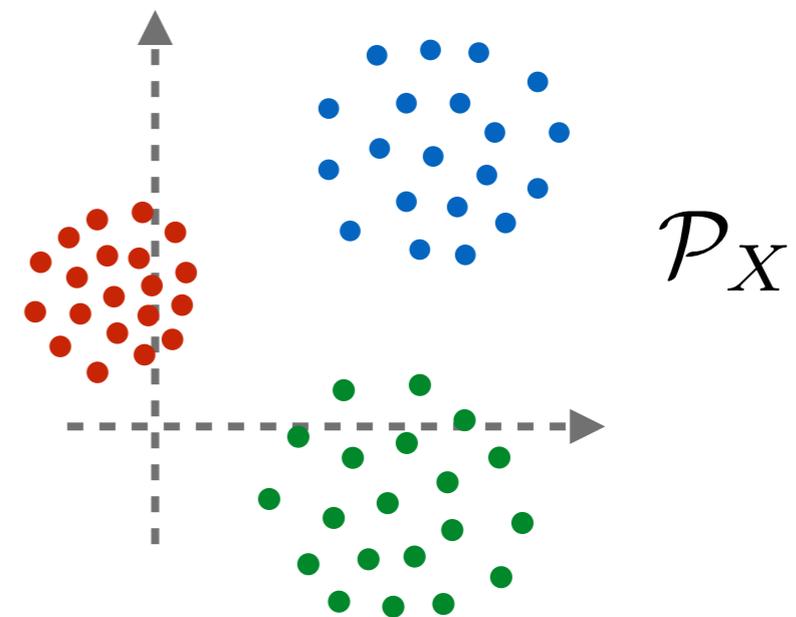
Sketch



$z_X \leftarrow$ information about \mathcal{P}

Observable? Not \mathcal{P} but

$$\mathcal{P}_X := \frac{1}{N} \sum_{j=1}^N \delta_{\mathbf{x}_j}$$



Objective: $z_X = \mathcal{A}(\mathcal{P}_X) \approx \mathcal{A}(\mathcal{P})$

We need to sketch distributions!

Sketch of a distribution

Linear sketch on a distribution \mathcal{P} :

Given m “frequencies” $\{\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m\}$,

$$\mathcal{A}(\mathcal{P}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} \left[e^{-i\boldsymbol{\omega}_j^T \mathbf{x}} \right]_{j=1}^m \in \mathbb{C}^m,$$

with $\mathbb{E}_{\mathbf{x} \sim \mathcal{P}}(e^{-i\boldsymbol{\omega}^T \mathbf{x}}) = \int_{\mathbb{R}^n} e^{-i\boldsymbol{\omega}^T \mathbf{x}} \mathcal{P}(\mathbf{x}) d^n \mathbf{x}$ for $\boldsymbol{\omega} \in \mathbb{R}^n$.

$\Rightarrow \mathcal{A}$ = sampling the *characteristic function* of \mathcal{P}
(*i.e.*, its Fourier transform) over m frequencies

$\Rightarrow \mathcal{A}$ = generalized moments of \mathcal{P}

Reminiscent of CS
with partial Fourier sensing!

Related to
Random Fourier Features
[Rahimi, Recht, 07]

Sketch of a distribution

Sketch of \mathbf{X} ? (the only observable!)

$$\mathcal{A}(\mathcal{P}_X) := \frac{1}{N} \sum_{i=1}^N \left[e^{-i\omega_j^T \mathbf{x}_i} \right]_{j=1}^m \in \mathbb{C}^m,$$

Properties:

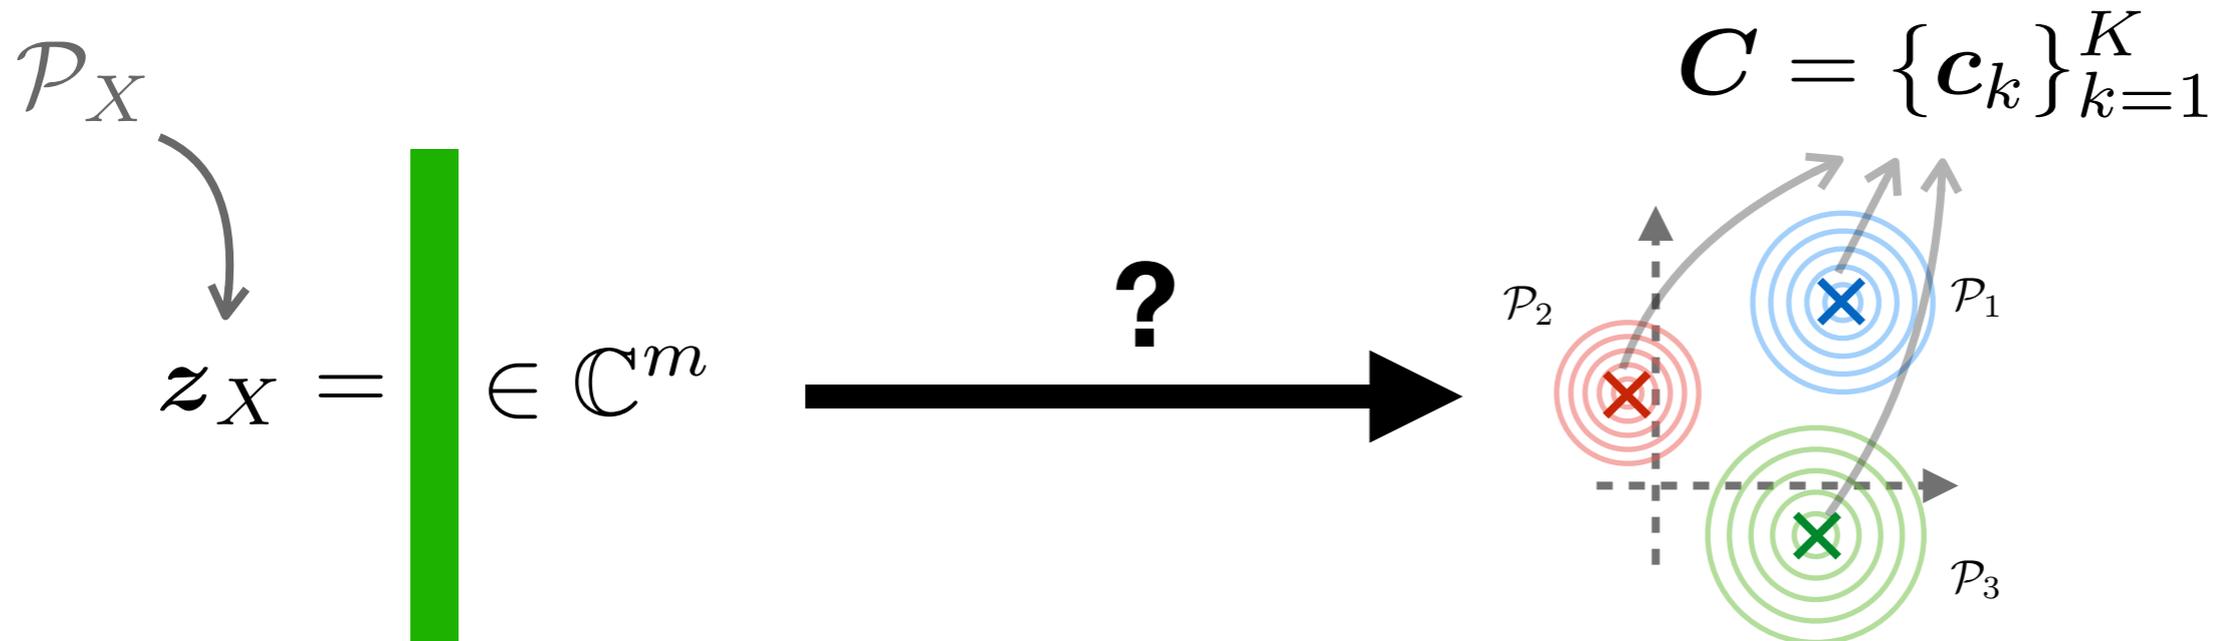
- ▶ *oblivious* of each involved signal (as expected)
- ▶ easily updatable (from the sum, or *pooling*)
- ▶ we “see” \mathcal{P}_X only over certain frequencies!

$$\mathcal{A} \left(\begin{array}{c} \text{red dots} \\ \text{blue dots} \\ \text{green dots} \\ \mathcal{P}_X \end{array} \right) \approx \mathcal{A} \left(\begin{array}{c} \text{red circles } \mathcal{P}_2 \\ \text{blue circles } \mathcal{P}_1 \\ \text{green circles } \mathcal{P}_3 \end{array} \right)$$

Provided we take low-frequencies!

Select $\{\omega_j\}_{j=1}^m$ conveniently!

Compressive K-Means



Sketch matching: finding centroids & weights from

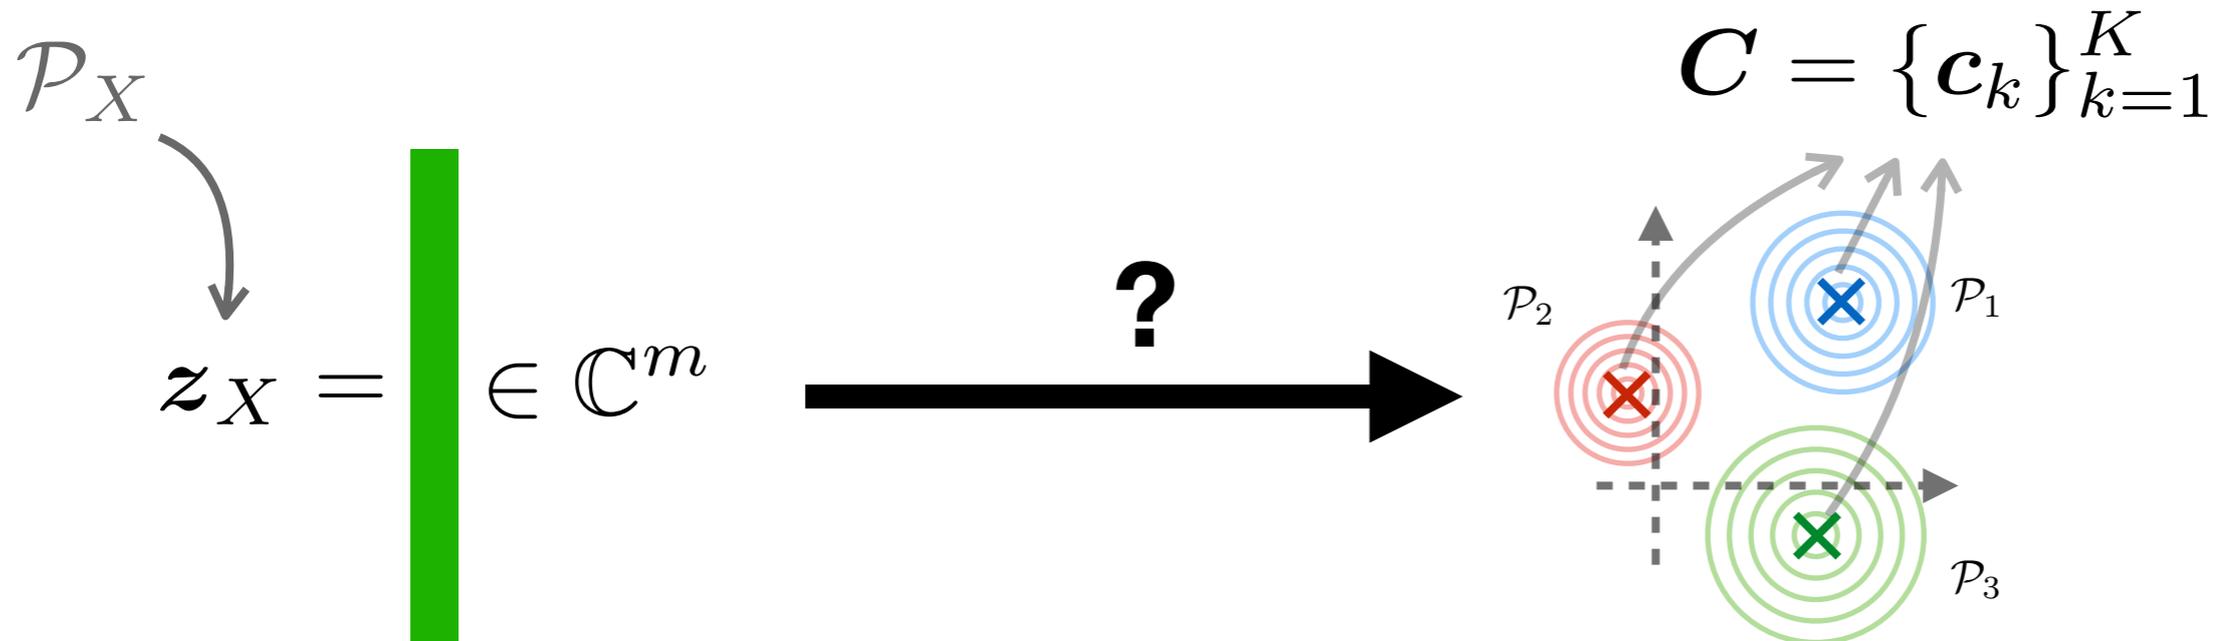
$$\min_{C, \alpha} \left\| z_X - \mathcal{A} \left(\sum_{k=1}^K \alpha_k \delta_{c_k} \right) \right\|^2$$

Sparse mixture
of K Diracs

with $\omega_j \sim_{\text{iid}} \Lambda$, $1 \leq j \leq m$

and Λ a pdf promoting low-frequencies

Compressive K-Means



Sketch matching: finding centroids & weights from

$$\min_{C, \alpha} \left\| z_X - \mathcal{A} \left(\sum_{k=1}^K \alpha_k \delta_{c_k} \right) \right\|^2$$

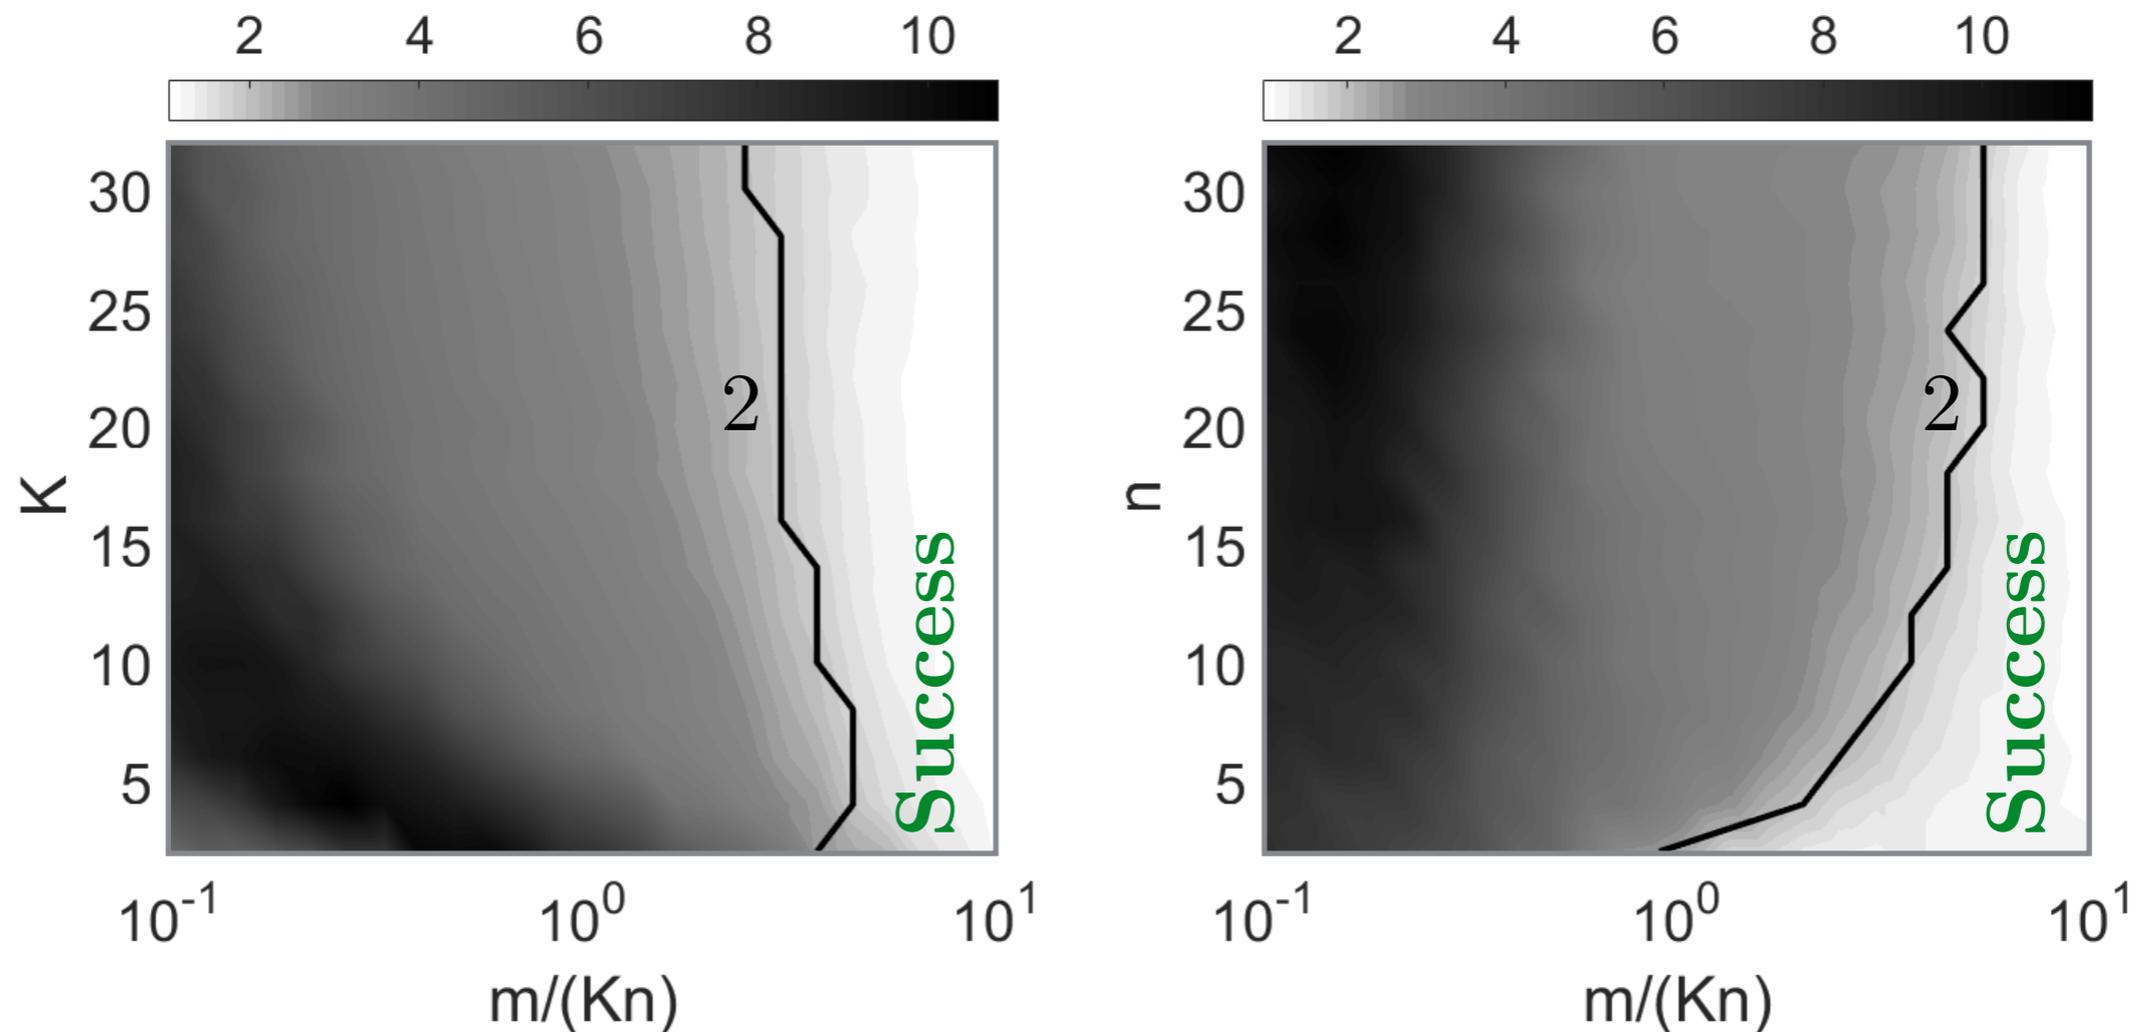
non-convex, infinite dimensional optimization.

At low frequencies,

$$\mathcal{A} \left(\begin{array}{c} \text{[red dots]} \quad \text{[blue dots]} \\ \text{[green dots]} \end{array} \mathcal{P}_X \right) \approx \mathcal{A} \left(\begin{array}{c} \text{[red circles]} \quad \text{[blue circles]} \\ \text{[green circles]} \end{array} \mathcal{P}_2, \mathcal{P}_1, \mathcal{P}_3 \right) \approx \mathcal{A} \left(\begin{array}{c} \text{[red x]} \quad \text{[blue x]} \\ \text{[green x]} \end{array} \right)$$

Compressive K-Means

$$SSE(\mathbf{X}, \mathbf{C}_{\text{KM}}) / SSE(\mathbf{X}, \mathbf{C}_{\text{CKM}})$$

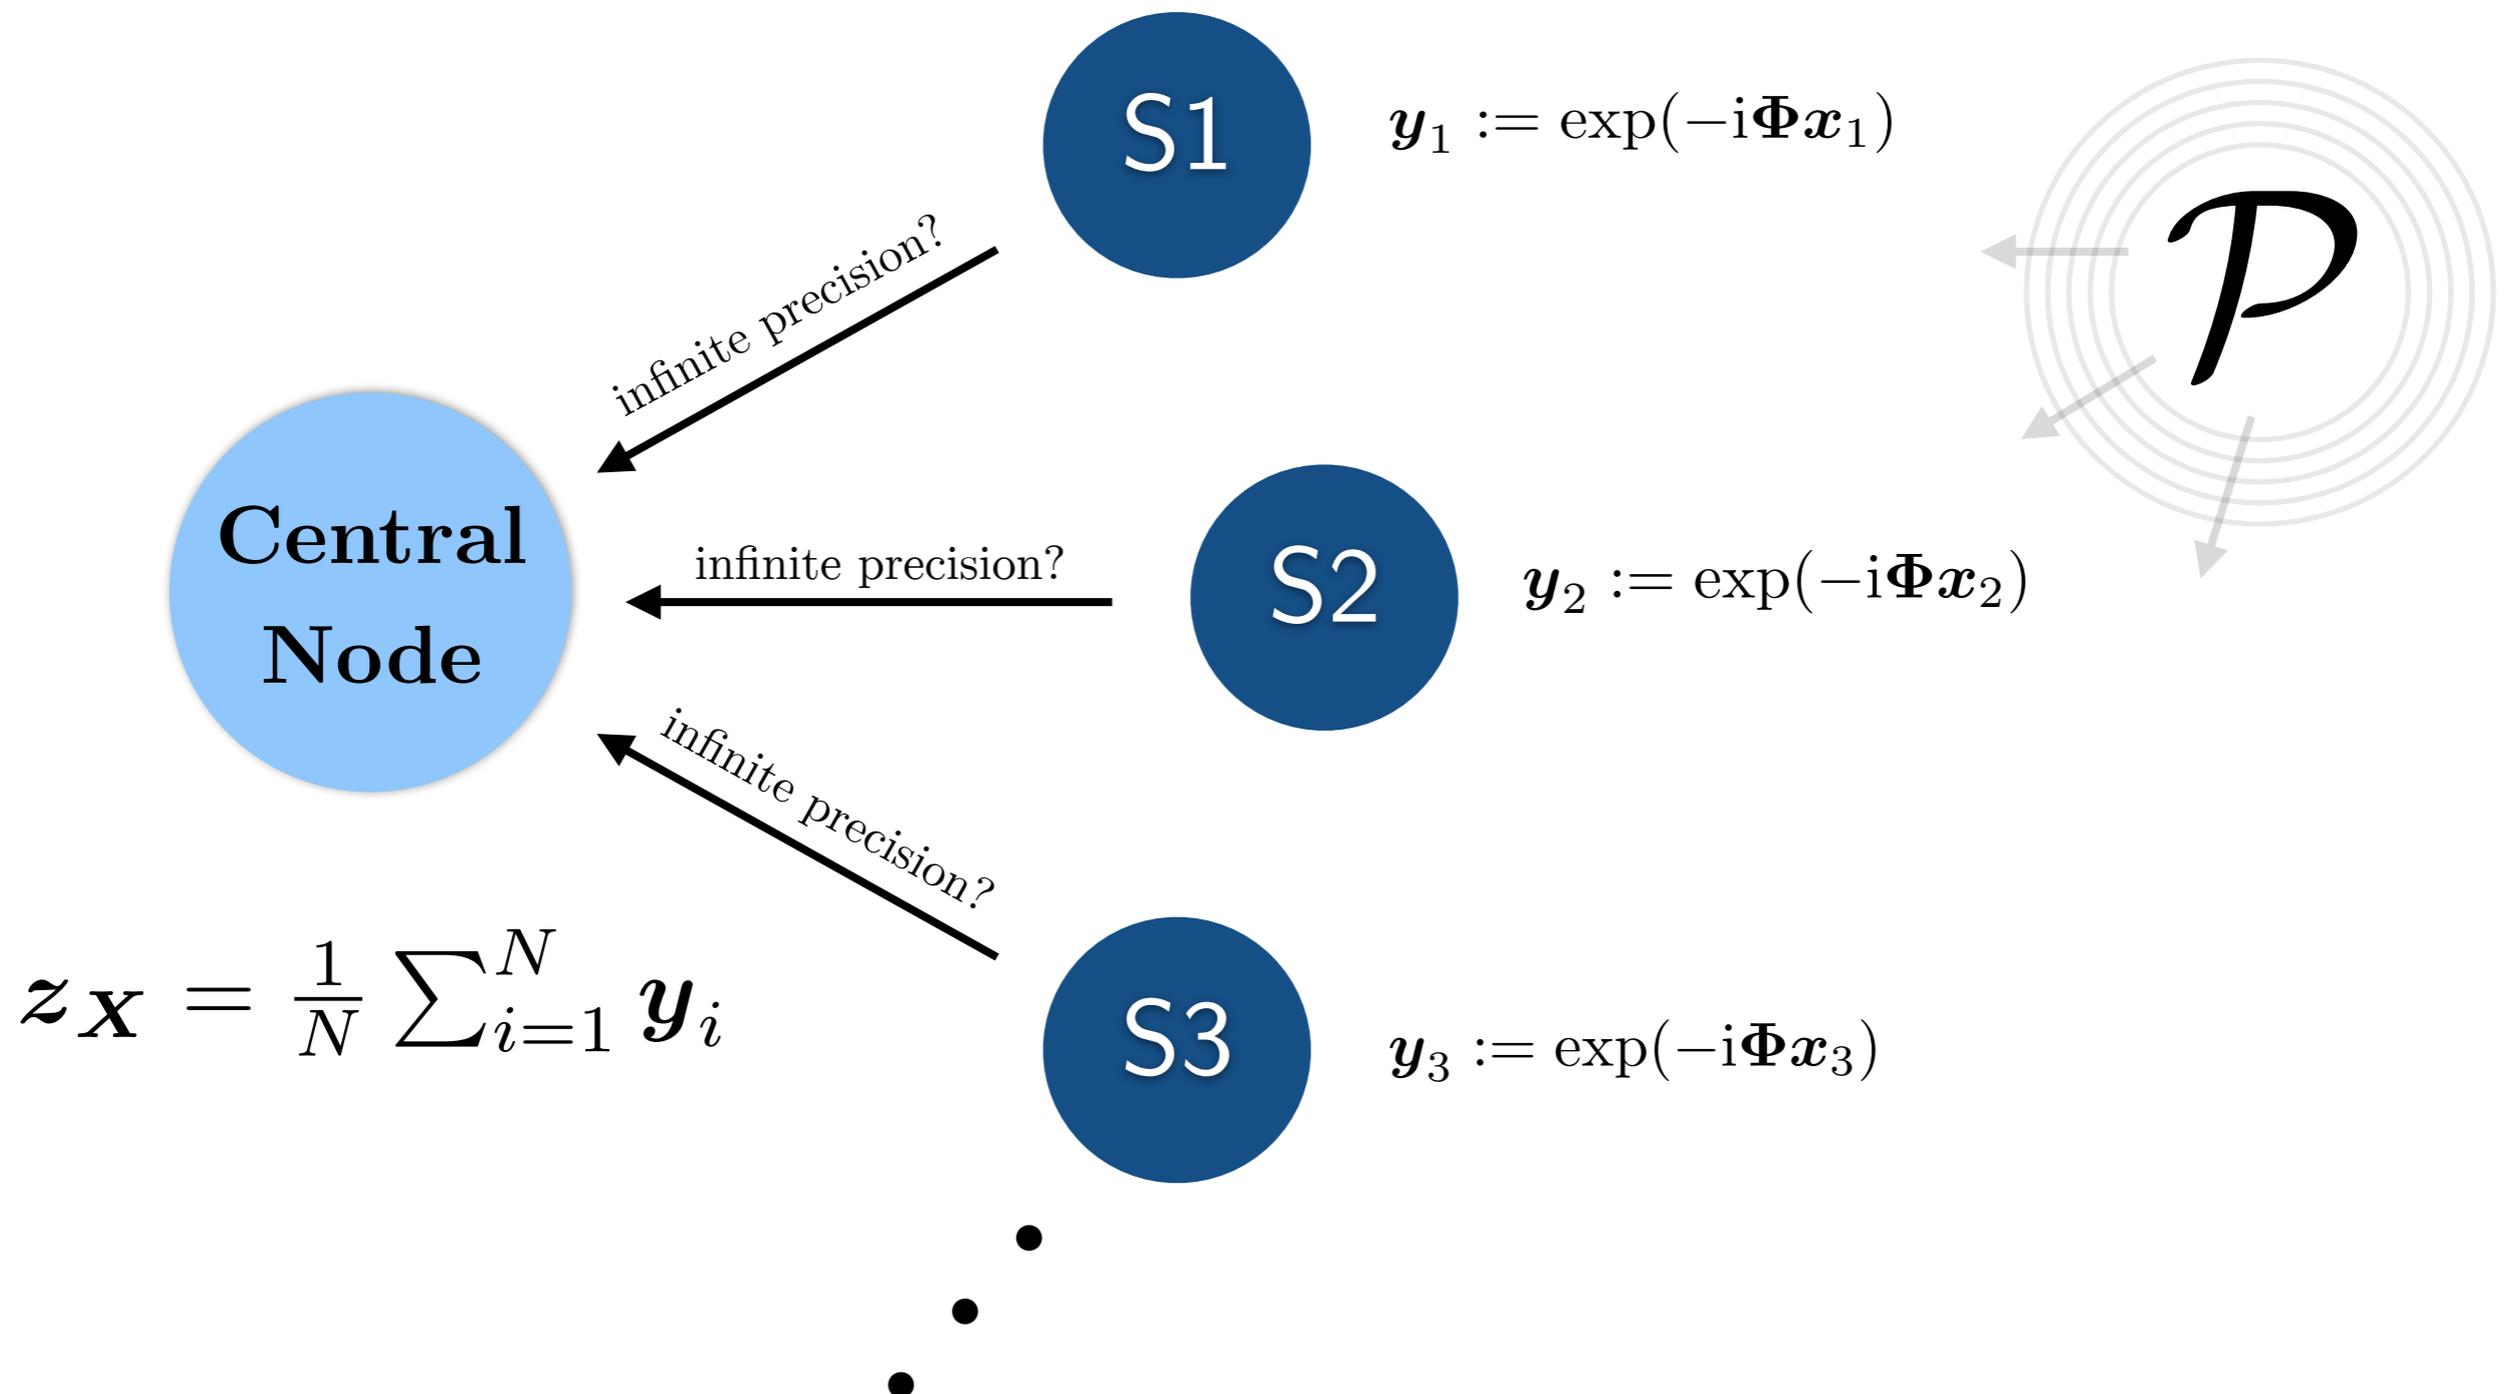


$$SSE(\mathbf{X}, \mathbf{C}) := \sum_{i=1}^N \min_k \|\mathbf{x}_i - \mathbf{c}_k\|^2$$

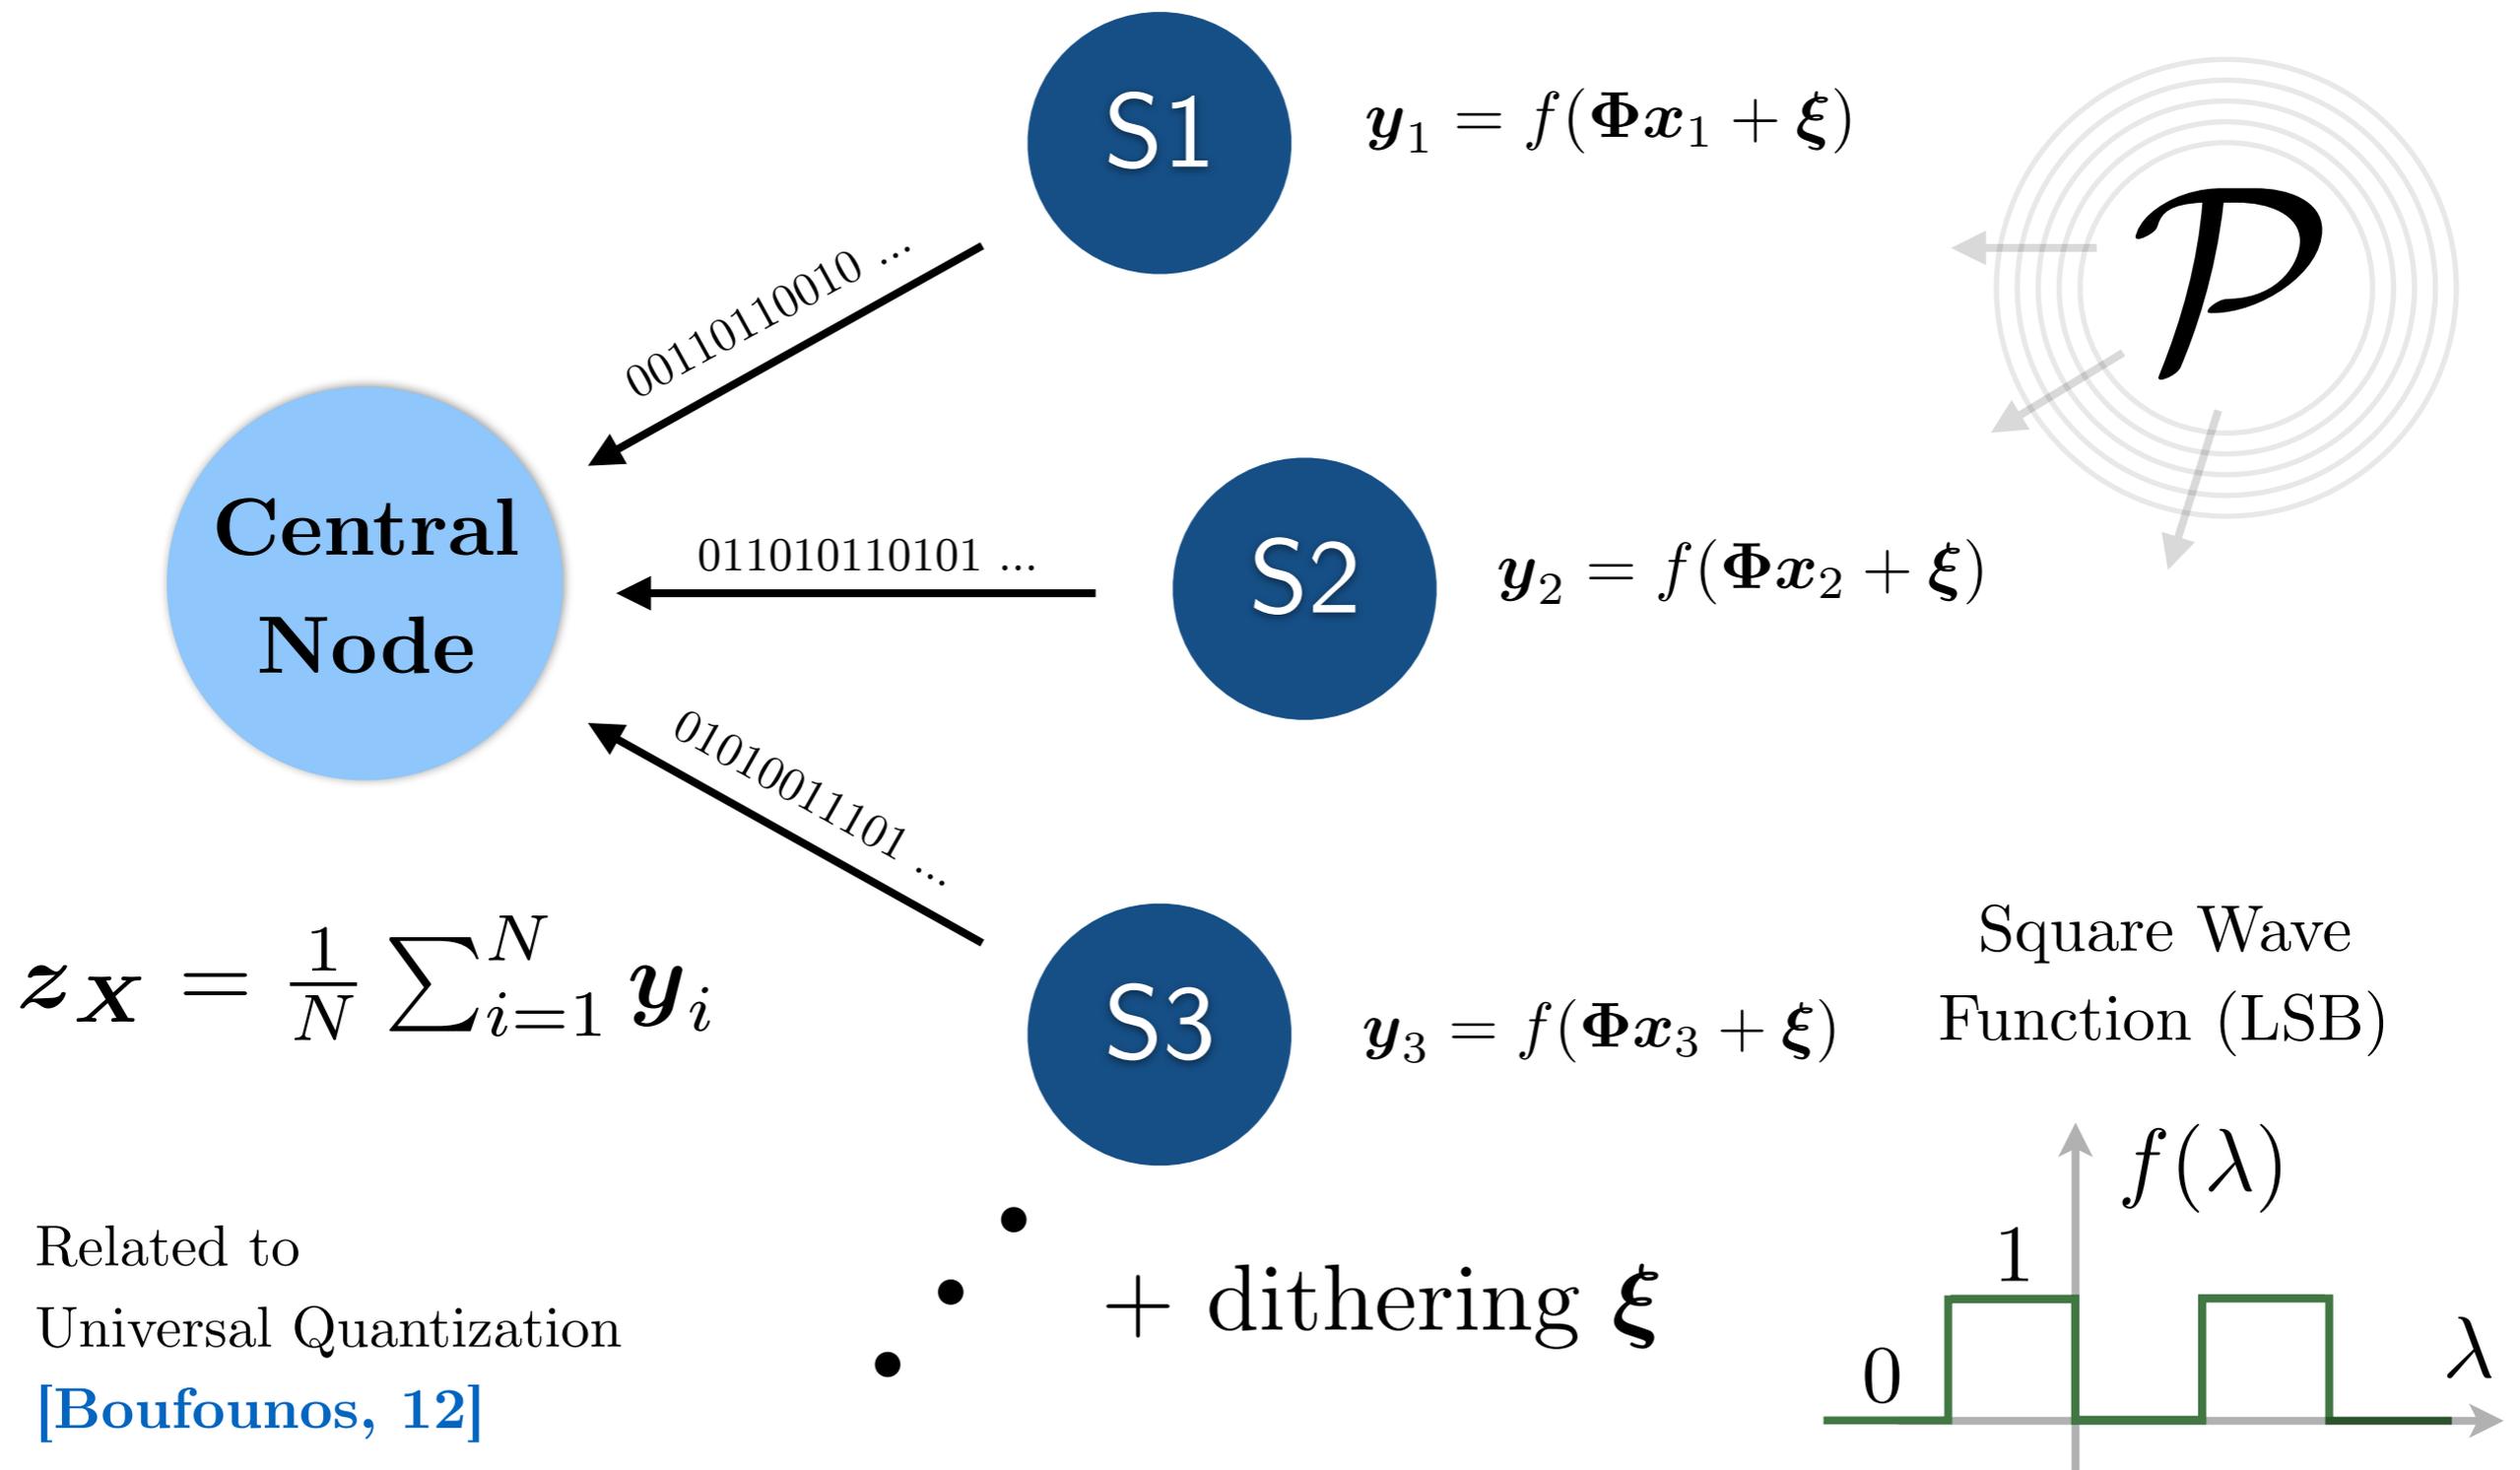
For large N (10^7), CKM up to 150 times faster than KM!

[Keriven, Tremblay, Traonmilin, Gribonval, 17]

Quantizing Sketches?

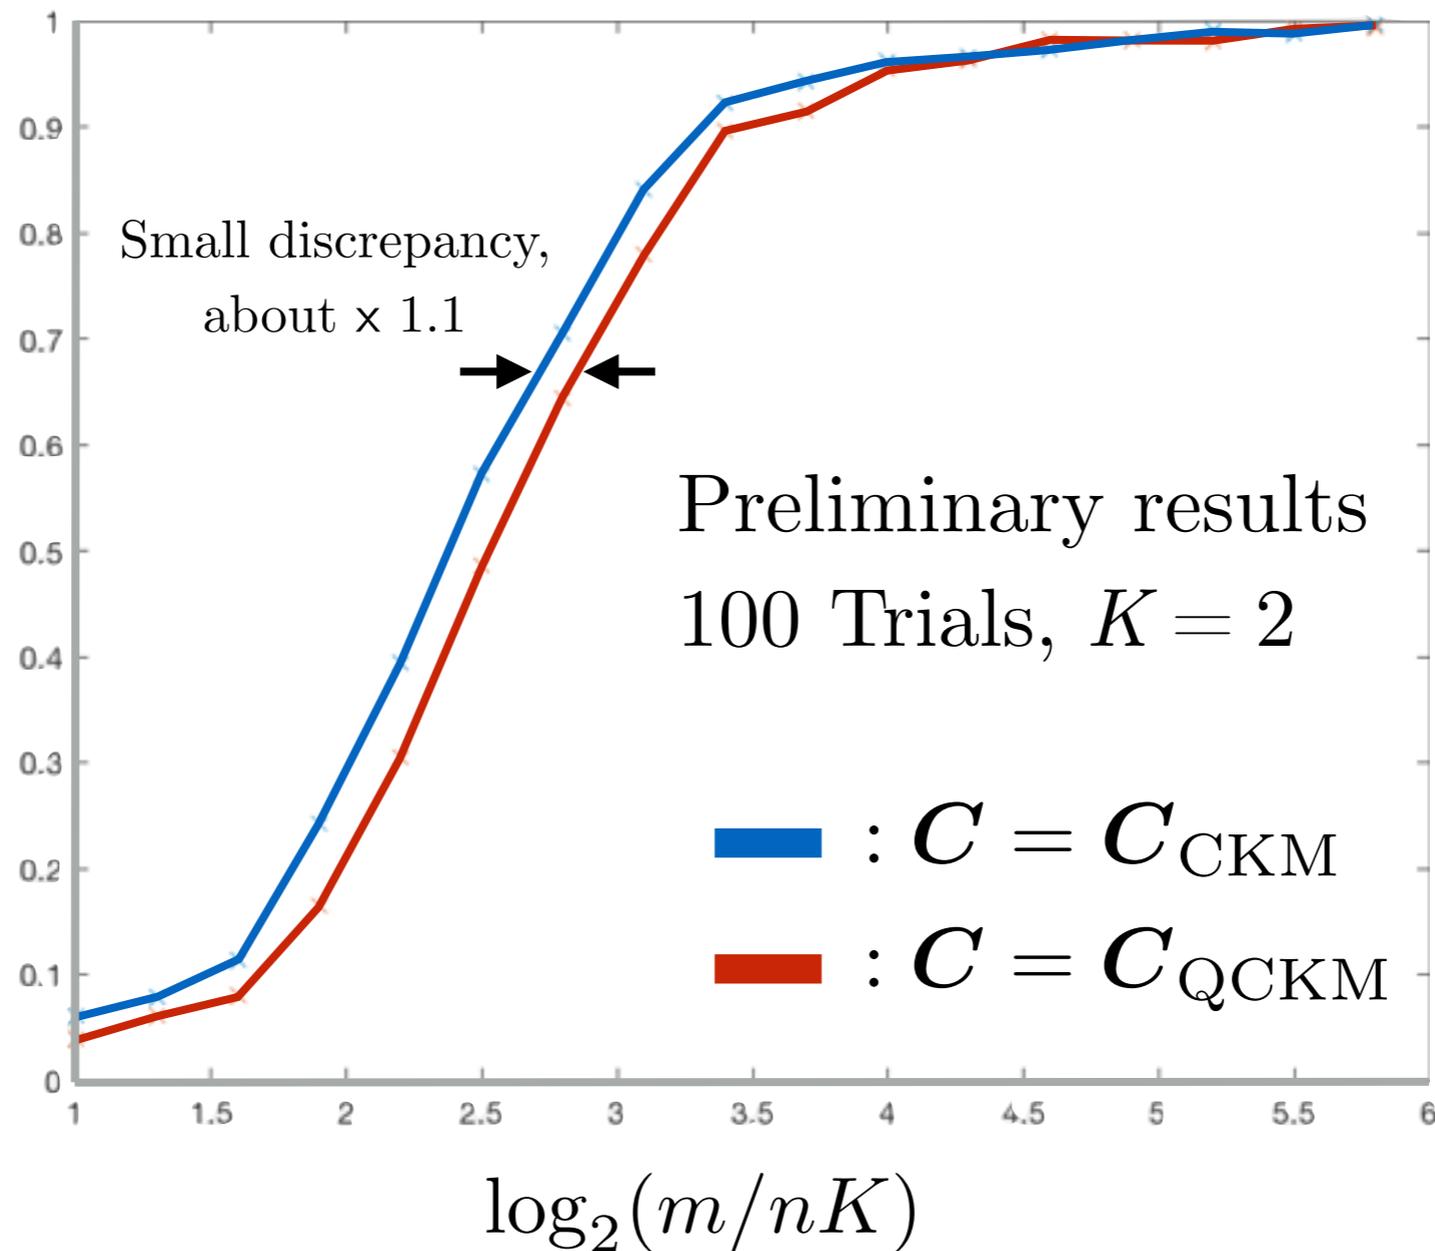


Quantizing Sketches?



Quantized Compressive K-Means

$$\Pr[\text{SSE}(\mathbf{X}, \mathbf{C}) < 1.2 \text{SSE}(\mathbf{X}, \mathbf{C}_{\text{KM}})]$$



Can be proved theoretically, thanks to the dithering
(ongoing work)

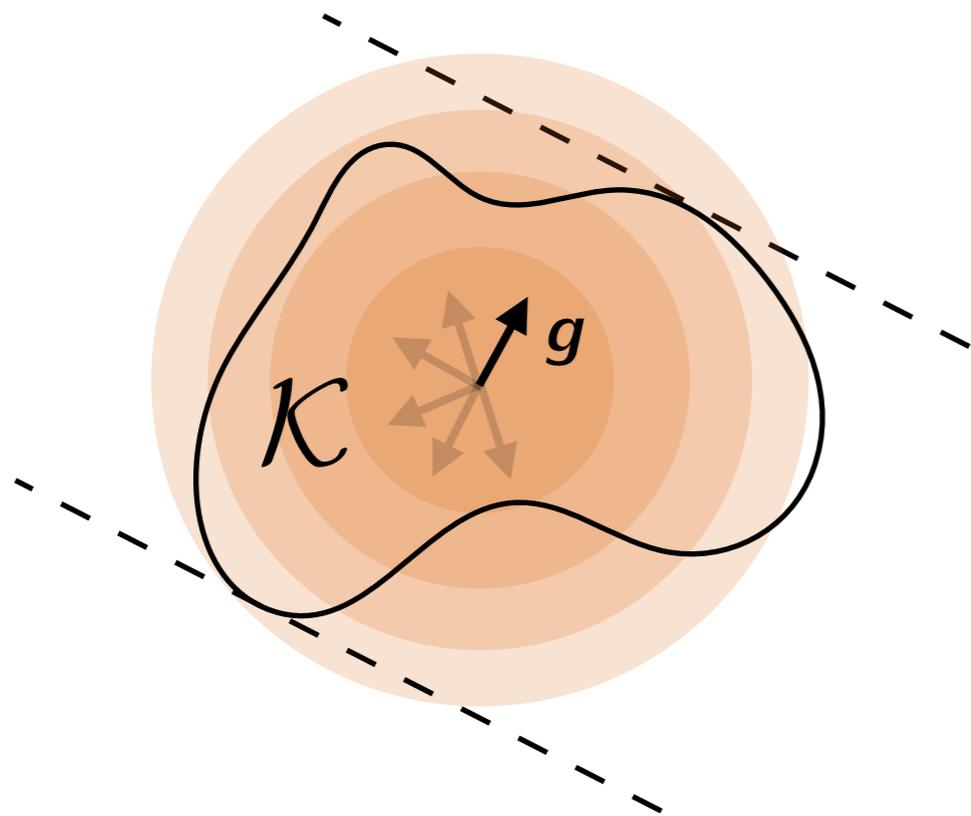
The Linear Case

Bandeira, Mixon, Recht '14 [BMR '14]

Parenthesis: Useful Tool

Gaussian Width (GW):

Let $\mathcal{K} \subset \mathbb{R}^n$

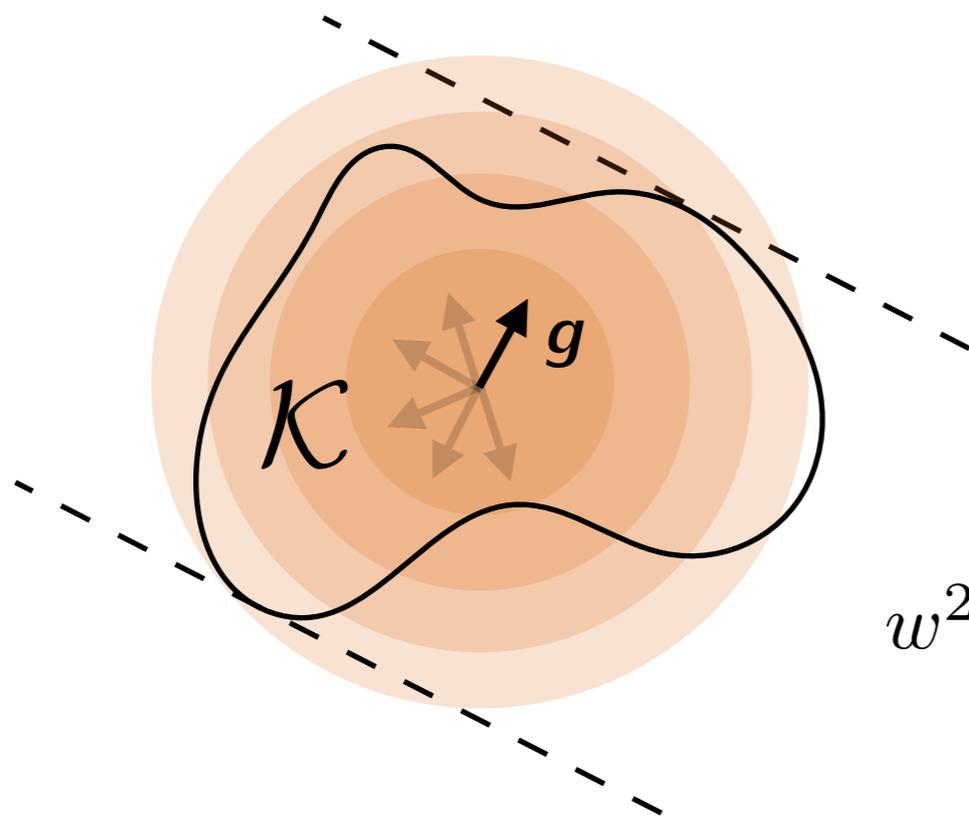


Parenthesis: Useful Tool

Gaussian Width (GW):

Let $\mathcal{K} \subset \mathbb{R}^n$, $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$,

$$w(\mathcal{K}) = \mathbb{E}_{\mathbf{g}} \sup_{\mathbf{x} \in \mathcal{K}} |\langle \mathbf{x}, \mathbf{g} \rangle|$$



Examples:

$$w^2(\mathcal{K}) \lesssim \log |\mathcal{K}|$$

$$w^2(\mathbb{B}^n) \lesssim n$$

$$w^2(\Sigma_k^n \cap \mathbb{B}^n) \lesssim k \log(n/k)$$

$$w^2(\mathcal{M}_r \cap \mathbb{B}_F^{n \times n}) \lesssim rn$$

$$w^2(\cup_{i=1}^T \mathcal{K}_i) \lesssim \log T + \max_i w^2(\mathcal{K}_i)$$

⋮

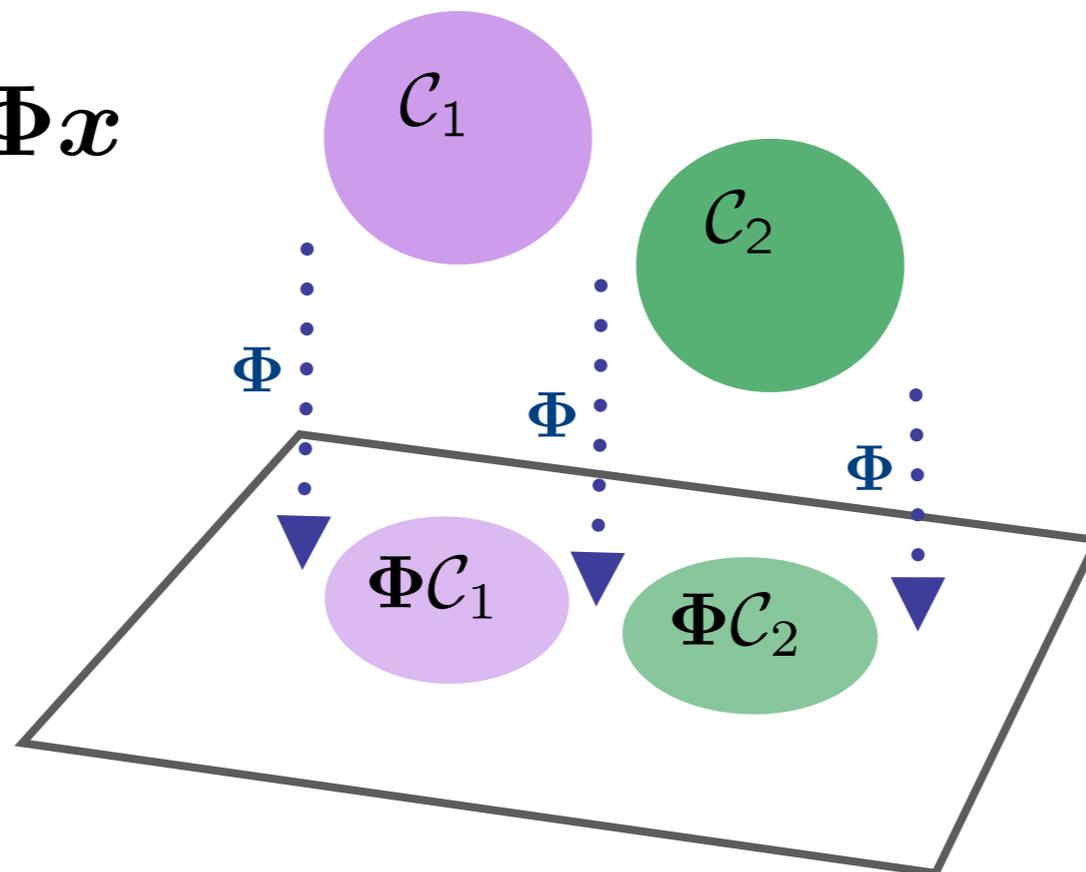
The Rare Eclipse Problem

Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$.
Given $\eta \in (0, 1)$, find the smallest m so that

$$p_0 := \mathbb{P}_{\Phi}[\Phi\mathcal{C}_1 \cap \Phi\mathcal{C}_2 = \emptyset] \geq 1 - \eta.$$

$$A(\mathbf{x}) = \Phi\mathbf{x}$$



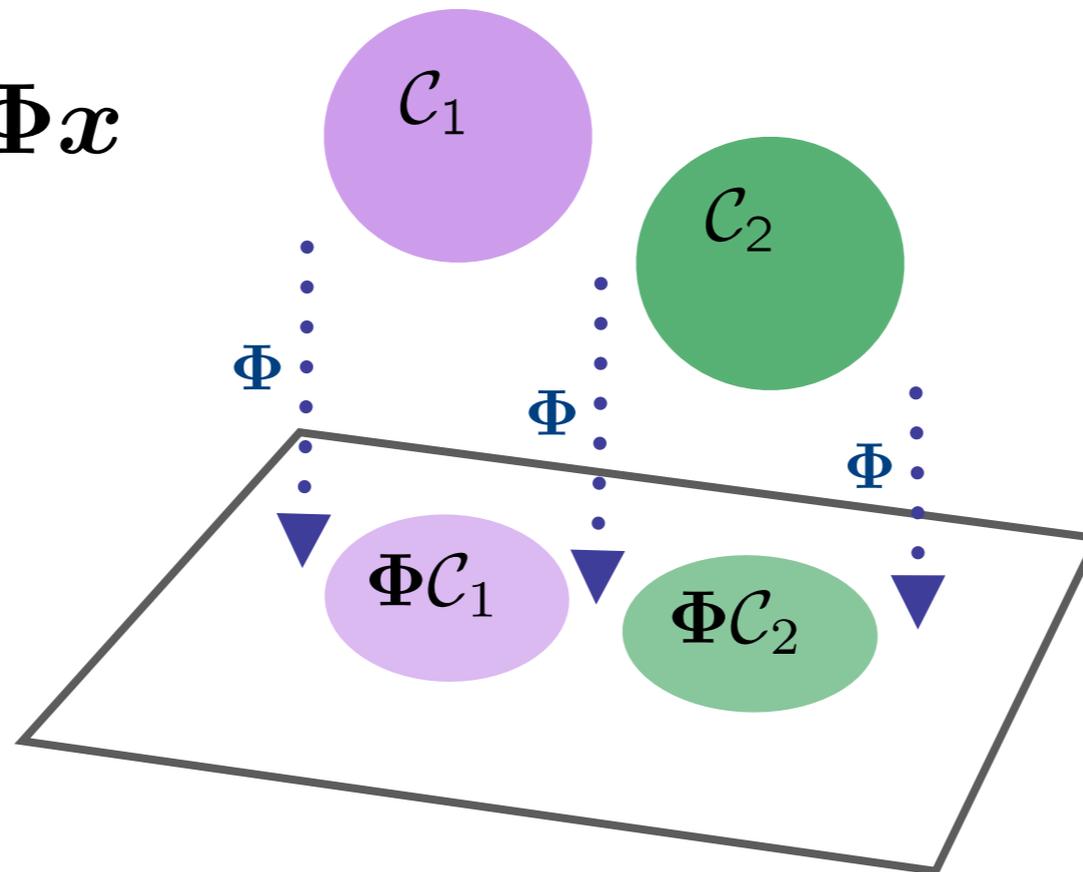
The Rare Eclipse Problem

Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$.
Given $\eta \in (0, 1)$, find the smallest m so that

$$p_0 = \mathbb{P}_{\Phi}[\forall \mathbf{x}_1 \in \mathcal{C}_1, \forall \mathbf{x}_2 \in \mathcal{C}_2, \Phi(\mathbf{x}_1 - \mathbf{x}_2) \neq 0] \geq 1 - \eta.$$

$$A(\mathbf{x}) = \Phi \mathbf{x}$$

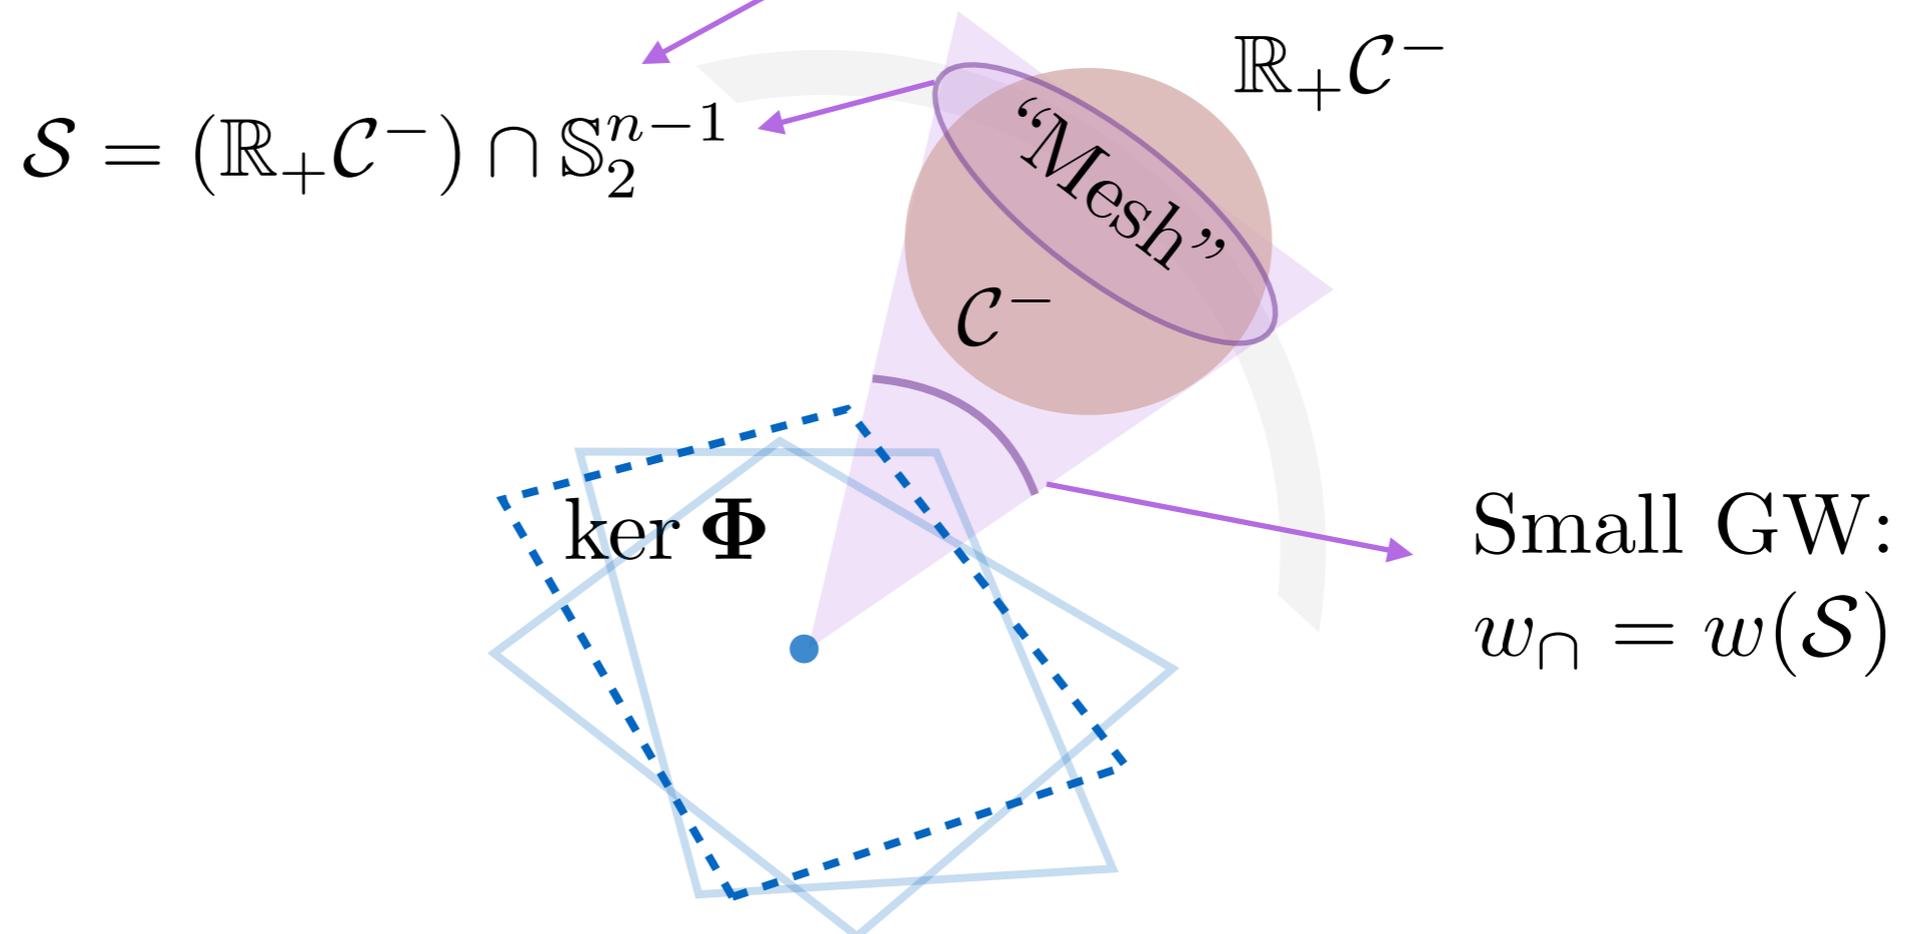


The Rare Eclipse Problem

Problem (Rare Eclipse Problem (Bandeira *et al.* '14)).

Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\Phi \sim \mathcal{N}^{m \times n}(0, 1)$. Given $\eta \in (0, 1)$, find the smallest m so that, with $\mathcal{C}^- = \mathcal{C}_1 - \mathcal{C}_2$,

$$p_0 = \mathbb{P}_{\Phi}[\mathcal{C}^- \cap \ker \Phi = \emptyset] = \mathbb{P}_{\Phi}[\mathcal{S} \cap \ker \Phi = \emptyset] \geq 1 - \eta.$$



The Rare Eclipse Problem

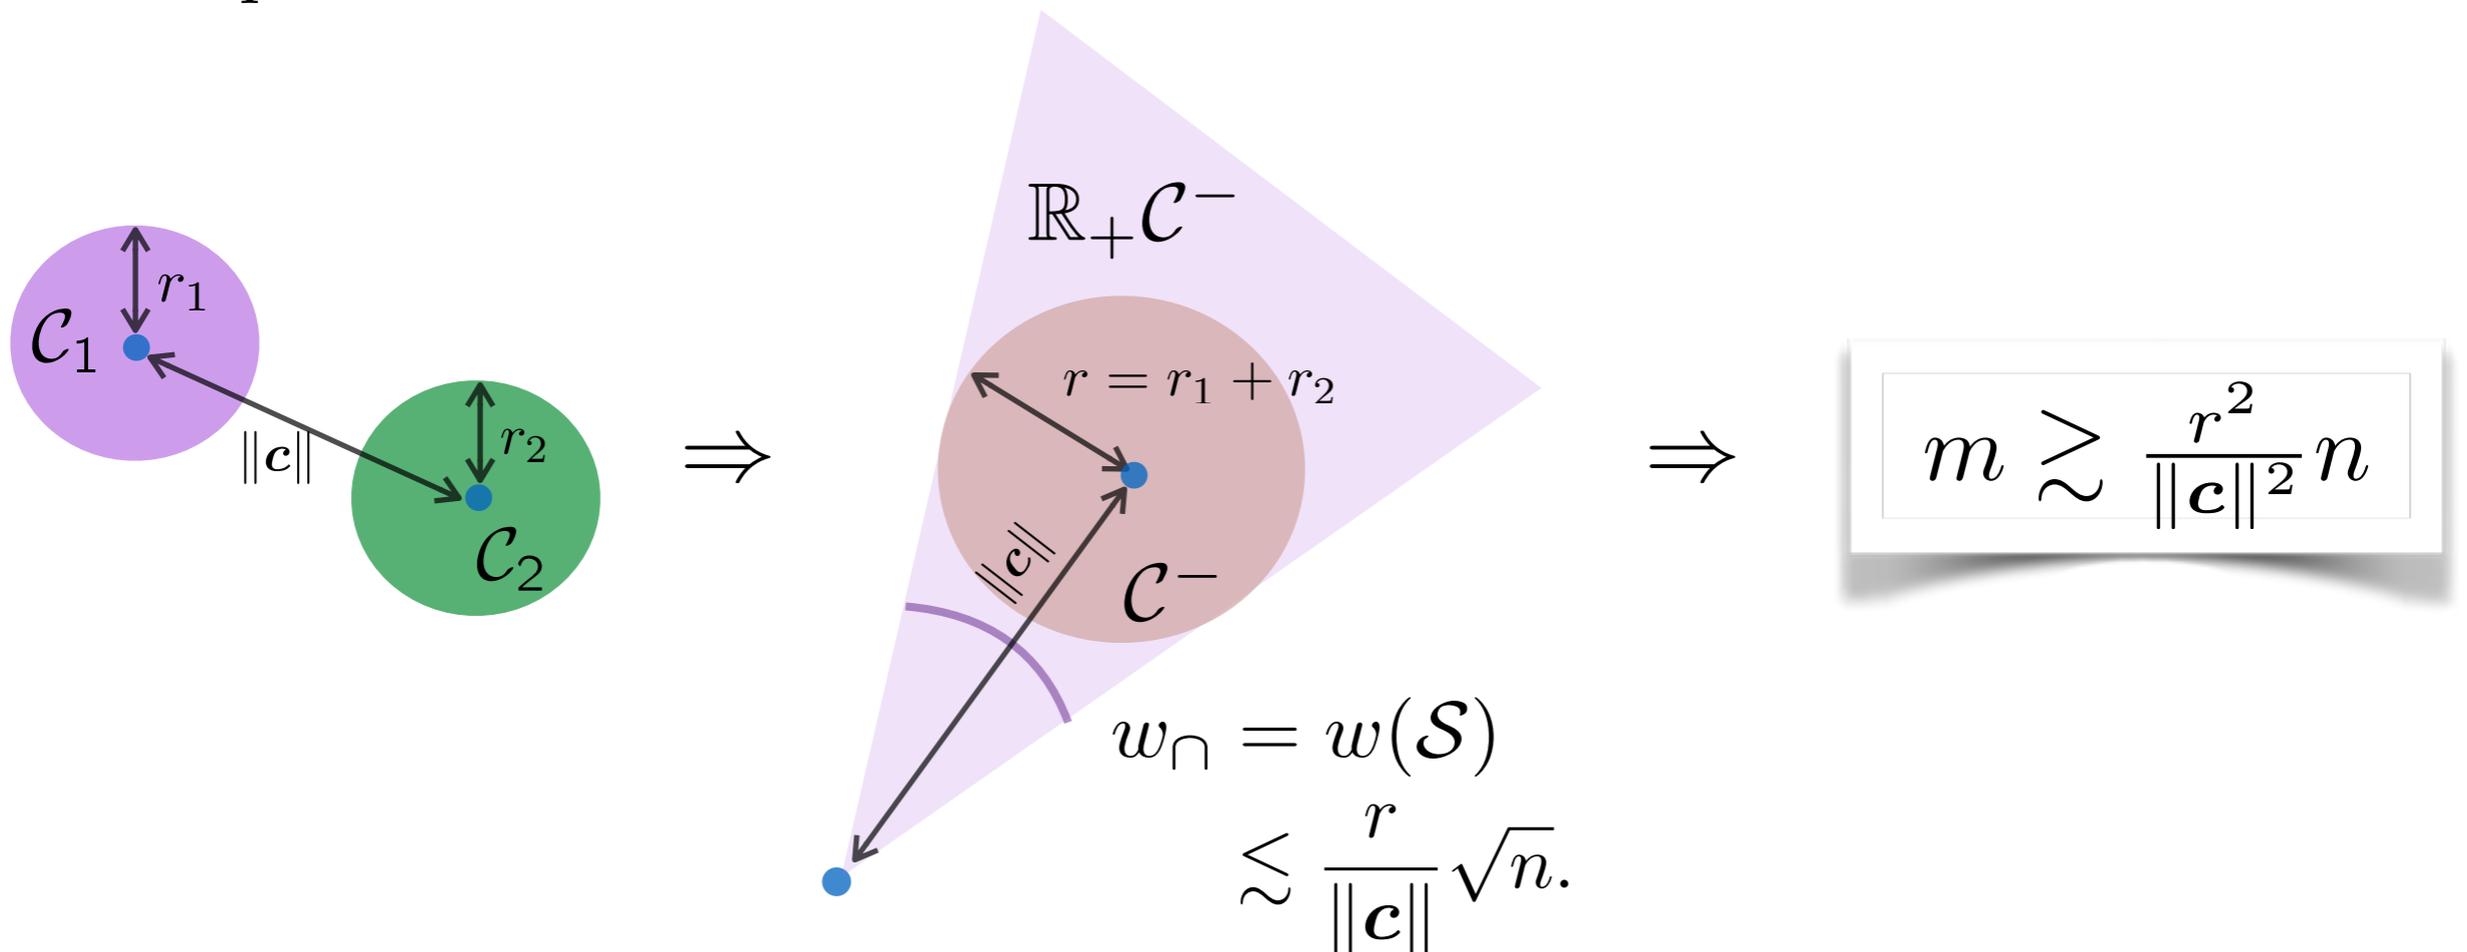
BMR '14: “Gordon’s escape through a mesh” theorem

Proposition (Corollary 3.1 in BMR '14).

(*& really tight* [Amelunxen et al, 13])

Given $\eta \in (0, 1)$, if $m > \left(w_n + \sqrt{2 \log \frac{1}{\eta}}\right)^2 + 1$ then $p_0 \geq 1 - \eta$.

Example:



The Rare Eclipse Problem (alternative)

Restricted Isometry Property: (ℓ_1, ℓ_2) -RIP(\mathcal{K}, ϵ)

$$\forall \mathbf{x} \in \mathcal{K}, (1 - \epsilon) \|\mathbf{x}\| \leq \|\Phi \mathbf{x}\|_1 \leq (1 + \epsilon) \|\mathbf{x}\|$$

[Schechtman, 06] [Plan, Vershynin, 14]

If $\mathcal{S} \subset \mathbb{S}^{n-1}$ and $m \gtrsim \epsilon^{-2} w^2(\mathcal{S})$, then, w.h.p.*,

$$(1 - \epsilon) \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{u}\|_1 \leq (1 + \epsilon)$$

for $\Phi \in \mathbb{R}^{m \times n}$ and $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$.

*: *i.e.*, $\mathbb{P} \geq 1 - C \exp(-c\epsilon^2 m)$.

The Rare Eclipse Problem (alternative)

Restricted Isometry Property: (ℓ_1, ℓ_2) -RIP(\mathcal{K}, ϵ)

$$\forall \mathbf{x} \in \mathcal{K}, (1 - \epsilon) \|\mathbf{x}\| \leq \|\Phi \mathbf{x}\|_1 \leq (1 + \epsilon) \|\mathbf{x}\|$$

[Schechtman, 06] [Plan, Vershynin, 14]

If $\mathcal{S} \subset \mathbb{S}^{n-1}$ and $m \gtrsim \epsilon^{-2} w^2(\mathcal{S})$, then, w.h.p*,

$$(1 - \epsilon) \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{u}\|_1 \leq (1 + \epsilon)$$

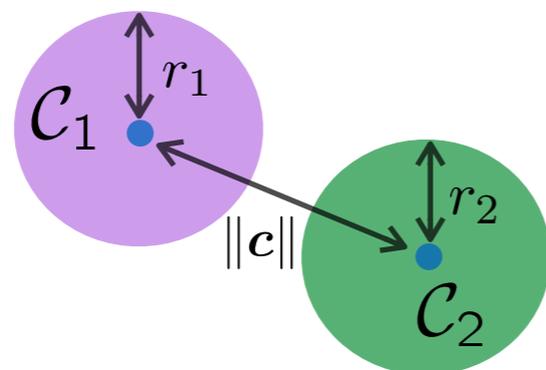
for $\Phi \in \mathbb{R}^{m \times n}$ and $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$.

Therefore:

For $\mathcal{S} = (\mathbb{R}^+ \mathcal{C}^-) \cap \mathbb{S}^{n-1}$, if $m \gtrsim \epsilon^{-2} w_{\cap}^2$, w.h.p*, (P1)

$$(1 - \epsilon) \|\mathbf{x}_1 - \mathbf{x}_2\| \leq \sqrt{\frac{\pi}{2}} \frac{1}{m} \|\Phi \mathbf{x}_1 - \Phi \mathbf{x}_2\|_1 \leq (1 + \epsilon) \|\mathbf{x}_1 - \mathbf{x}_2\|,$$

for all $\mathbf{x}_1 \in \mathcal{C}_1$ and all $\mathbf{x}_2 \in \mathcal{C}_2$.



This result also explains REP!
(but less sharply)

*: *i.e.*, $\mathbb{P} \geq 1 - C \exp(-c\epsilon^2 m)$.

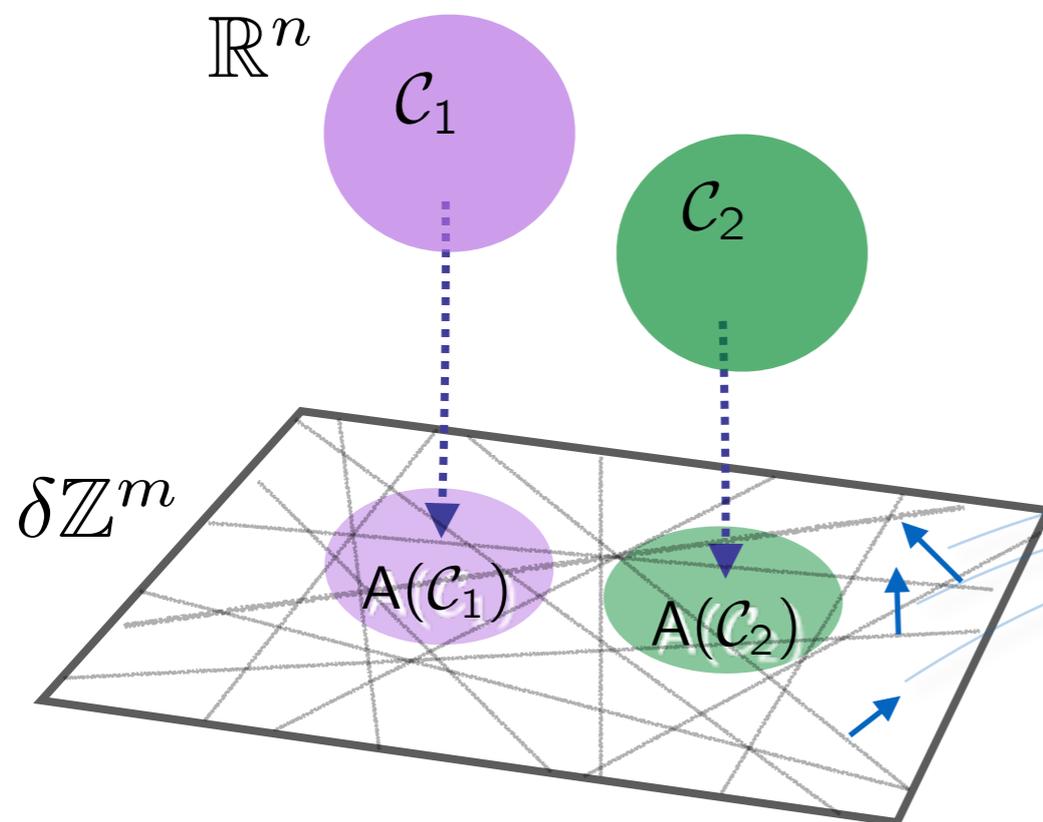
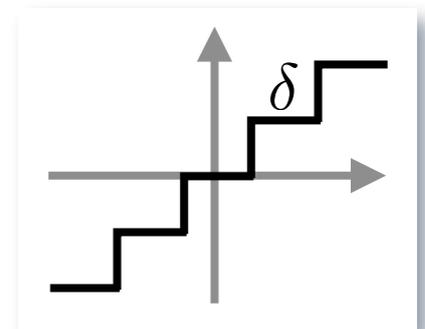
The Quantized Case

This work

Quantized Dithered Random Mapping

$$\left\{ \begin{array}{l} \delta > 0 \\ \mathcal{Q}(\cdot) := \delta \lfloor \cdot / \delta \rfloor \quad (\text{applied componentwise}) \\ \Phi \text{ is } (\ell_1, \ell_2)\text{-RIP}(\mathcal{K}, \epsilon) \\ \text{and a dithering } \xi \in \mathbb{R}^m \text{ with } \xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta]) \end{array} \right. \begin{array}{l} \text{(a resolution),} \\ \text{(scalar quantizer),} \\ \text{(a well-behaved } \Phi\text{),} \\ \text{(your friend)} \end{array}$$

QDRM: $A(\mathbf{x}) := \mathcal{Q}(\Phi \mathbf{x} + \xi)$



$$A_i(\mathbf{x}) = \mathcal{Q}(\varphi_i^\top \mathbf{x} + \xi_i) \in \delta \mathbb{Z}$$

$$\Phi = (\varphi_1, \dots, \varphi_m)^\top$$

Quantized Random Embeddings

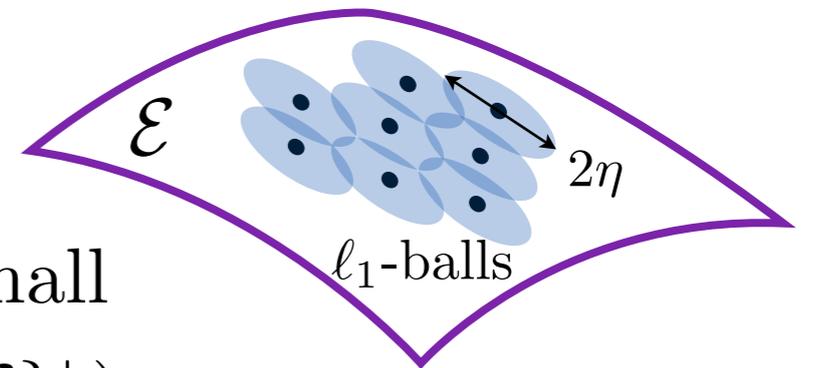
(i.e., only thanks to the dithering!)

Embedding $\mathcal{E} \subset \mathbb{R}^m$ into $\delta\mathbb{Z}^m$ with

$$A'(\mathbf{y}) := \underset{\uparrow \delta}{\mathcal{Q}}(\mathbf{y} + \underset{\uparrow \delta}{\boldsymbol{\xi}})$$

If \mathcal{E} has small ℓ_1 -Kolmogorov entropy, *i.e.*, small

$$\mathcal{H}_1(\mathcal{E}, \eta) = \log(\min_{\mathcal{S}} |\{\mathcal{S} \subset \mathcal{E} : \mathcal{S} + \eta\mathbb{B}_1^n \supset \mathcal{E}\}|)$$



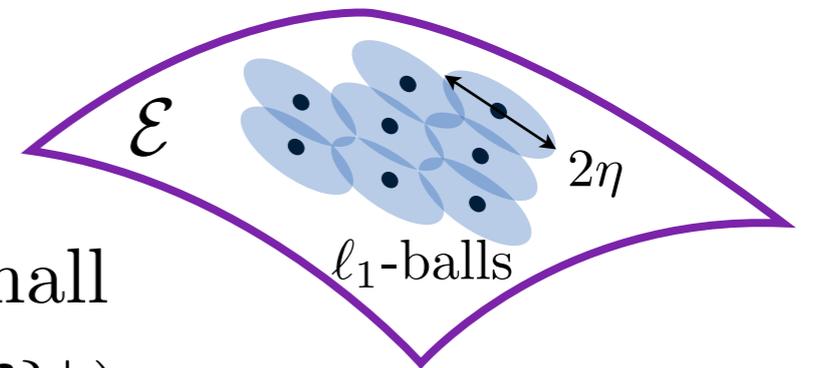
*: [LJ, Cambareri, '16] [Cambareri, Xu, LJ, '17]

Quantized Random Embeddings

(i.e., only thanks to the dithering!)

Embedding $\mathcal{E} \subset \mathbb{R}^m$ into $\delta\mathbb{Z}^m$ with

$$A'(\mathbf{y}) := \underset{\uparrow \delta}{\mathcal{Q}}(\mathbf{y} + \underset{\uparrow \delta}{\boldsymbol{\xi}})$$



If \mathcal{E} has small ℓ_1 -Kolmogorov entropy, *i.e.*, small

$$\mathcal{H}_1(\mathcal{E}, \eta) = \log(\min_{\mathcal{S}} |\{\mathcal{S} \subset \mathcal{E} : \mathcal{S} + \eta\mathbb{B}_1^n \supset \mathcal{E}\}|)$$

Given $\epsilon > 0$, if $m \gtrsim \epsilon^{-2} \mathcal{H}_1(\mathcal{E}, \frac{m\delta\epsilon^2}{1+\epsilon})$, then, w.h.p*, for all $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{E}$ and some $c > 0$,

(P2)

$$\|\mathbf{y}_1 - \mathbf{y}_2\|_1 - c\delta\epsilon \leq \frac{1}{m} \|A'(\mathbf{y}_1) - A'(\mathbf{y}_2)\|_1 \leq \|\mathbf{y}_1 - \mathbf{y}_2\|_1 + c\delta\epsilon.$$

Remarks: \circ For $\mathcal{E} = \Phi\mathcal{K}$ and Φ an (ℓ_1, ℓ_2) -RIP($\mathcal{K} - \mathcal{K}, \epsilon' < 1$),

$$\mathcal{H}_1(\mathcal{E}, 2m\eta) \leq \mathcal{H}_2(\mathcal{K}, \eta) \quad (\text{P3})$$

\circ \mathcal{H}_2 is bounded for sets (cones, convex spaces, ...)

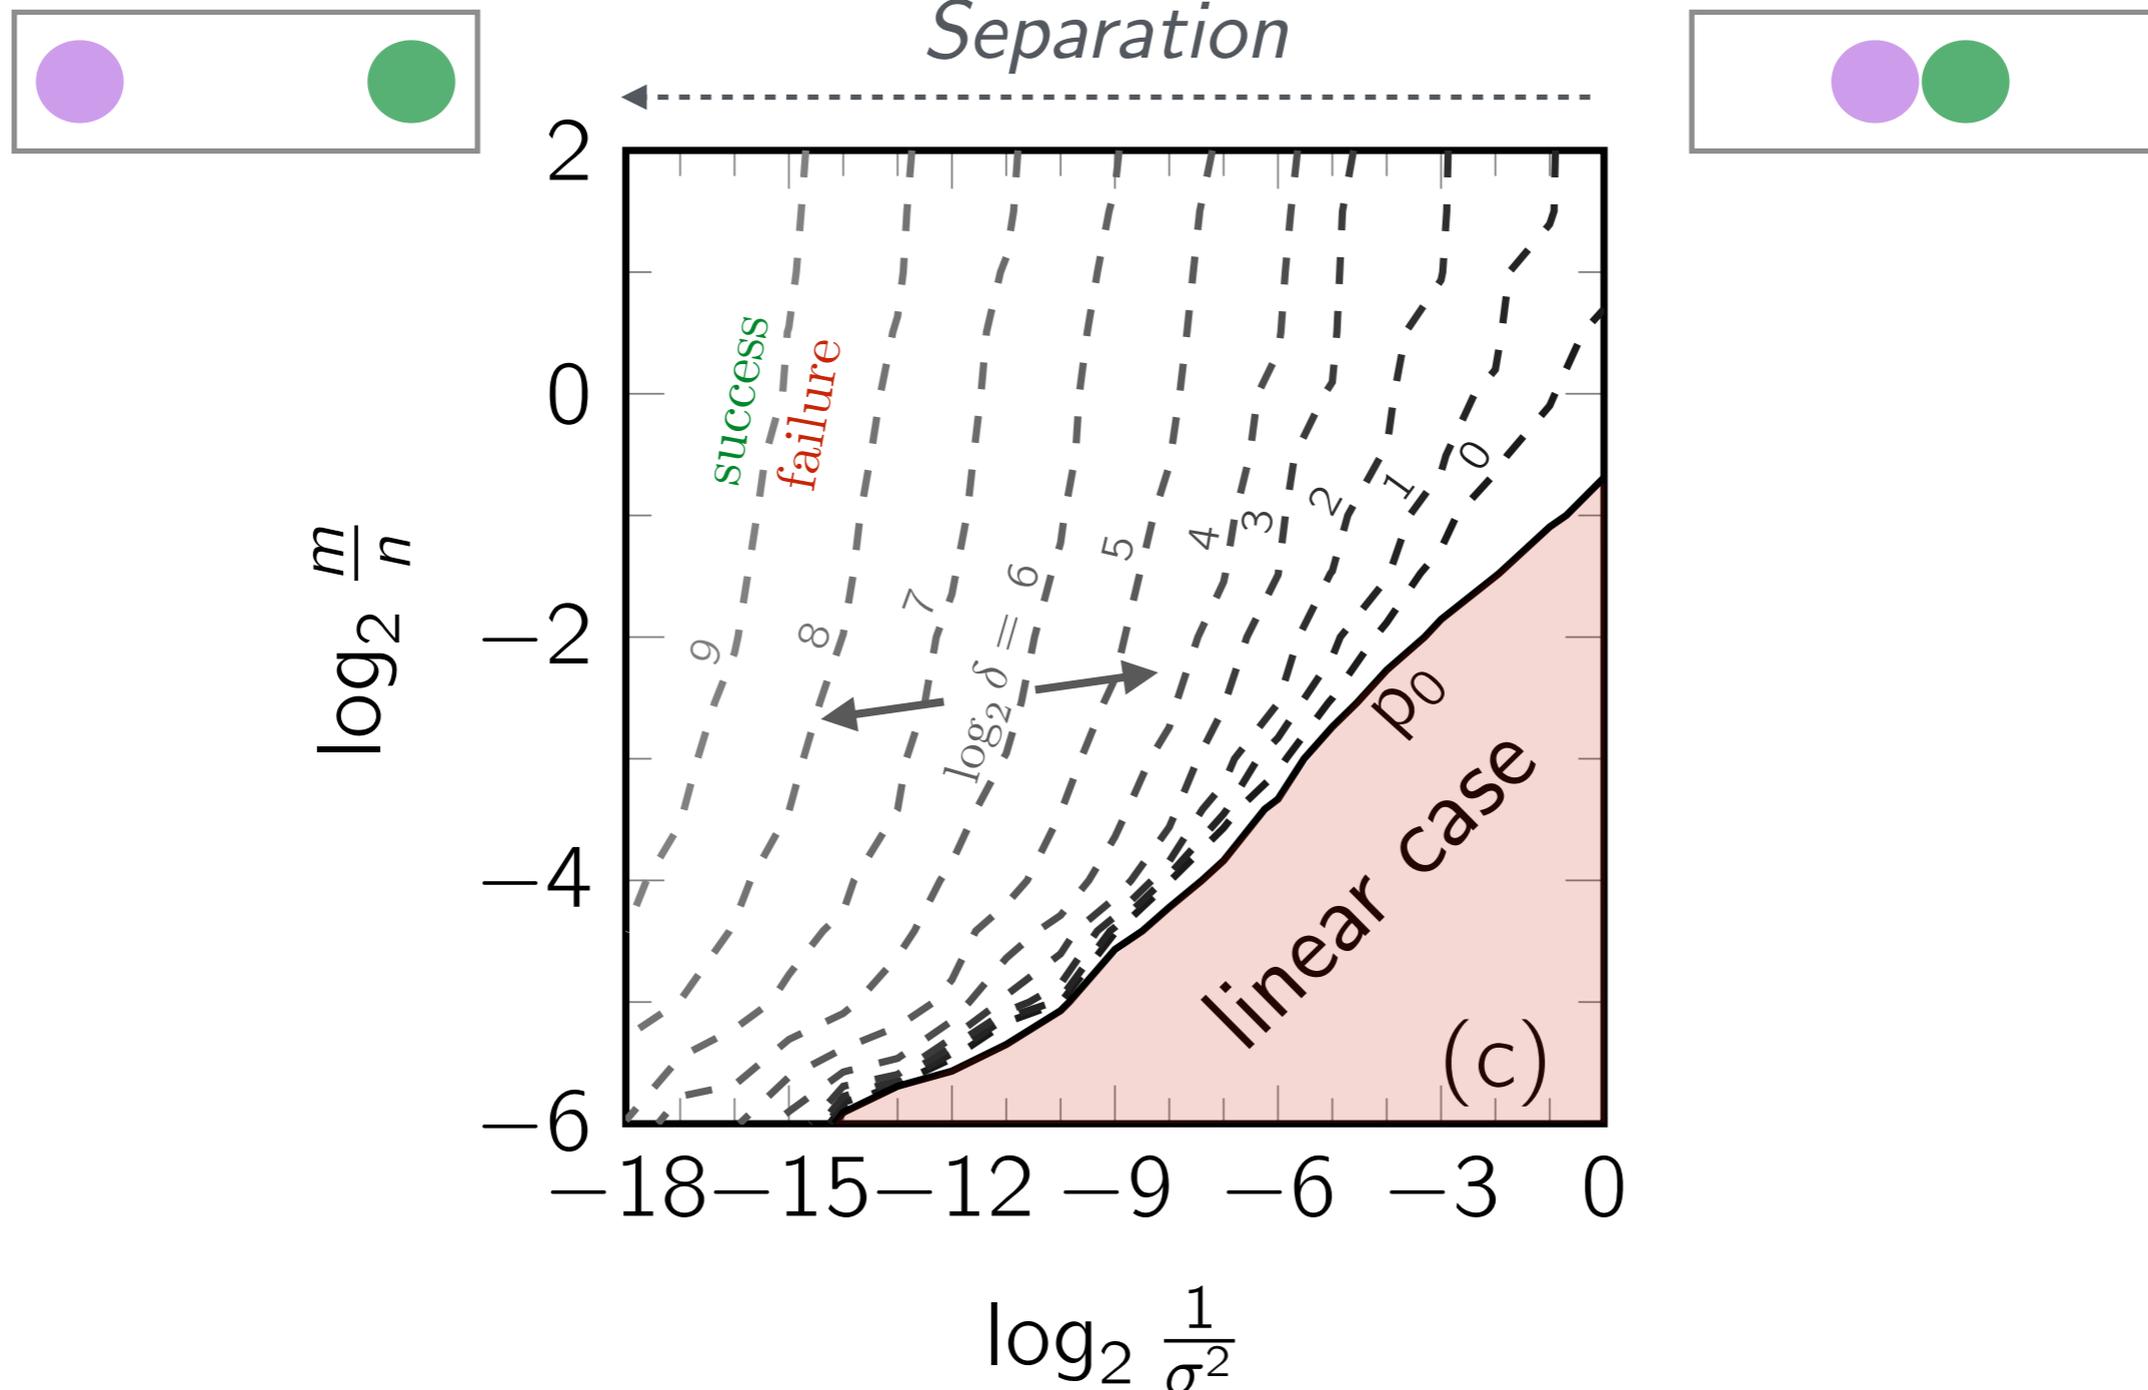
*: [LJ, Cambareri, '16] [Cambareri, Xu, LJ, '17]

Conclusion

- **Take-away messages:**
 - QDRMs: non-adaptive dimensionality reduction, preserve geometry of datasets
 - Extension of BMR '14 to QDRMs, with sample complexity loss in $\frac{\delta^2}{\sigma^2}n$ (quantiz. impact)
- **Future work:**
 - Better bound and testable conditions for empirical tests.
 - Extension of result to $\text{RIP}_{2,2}$ matrices, “fast” random matrices
 - Extension of framework to other non-linear maps (ReLU?)
→ applicability to D/CNN with random weights [Giryes et al. 15]?

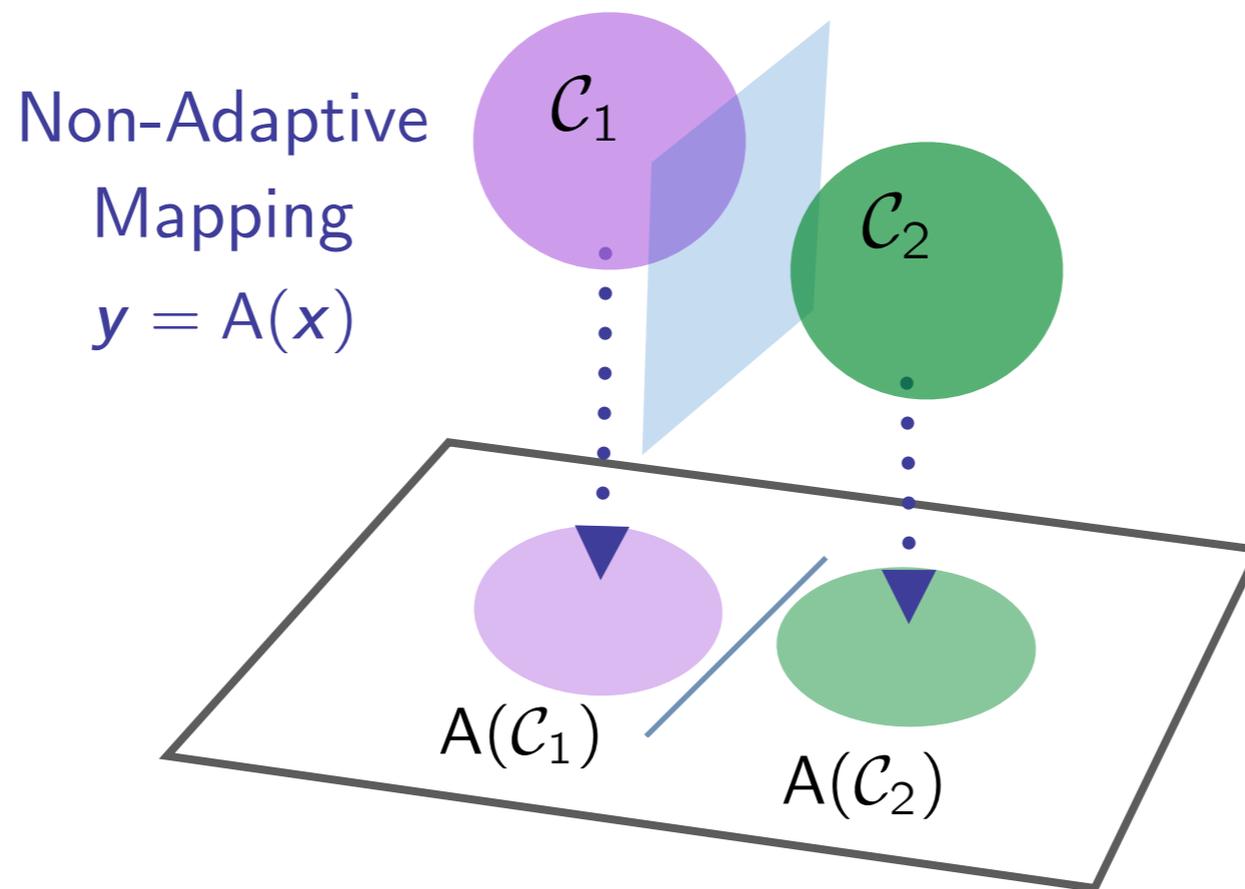
Simulations: Phase transition study

Empirical evaluation, **crude lower bound** (with no dithering),
128 trials, randomly drawn, $n = 64$, fixed $r = 1$, $\delta = \{1, 2, \dots, 512\}$.
Represented: Phase transition level-curves (at 0.9).



The Big Picture (an easy classification problem)

How does the *probability of error* of a (generic) learning task depend on m, A, \mathcal{K} ? What if A is (*mildly*) non-linear?



Related works on “(linear) compressive classification”:

Davenport *et al.* '07-'10, Haupt *et al.* '06,

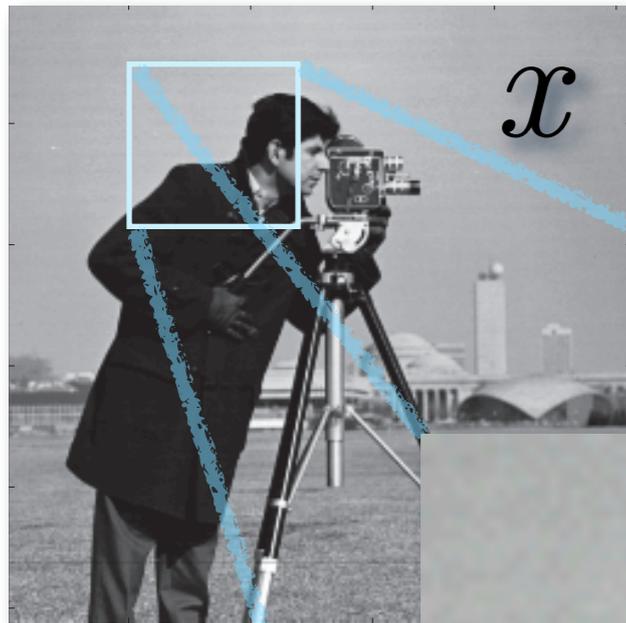
Reboredo *et al.* '13-'16,

[Bandeira, Mixon, Recht '14 \[BMR '14\]](#)

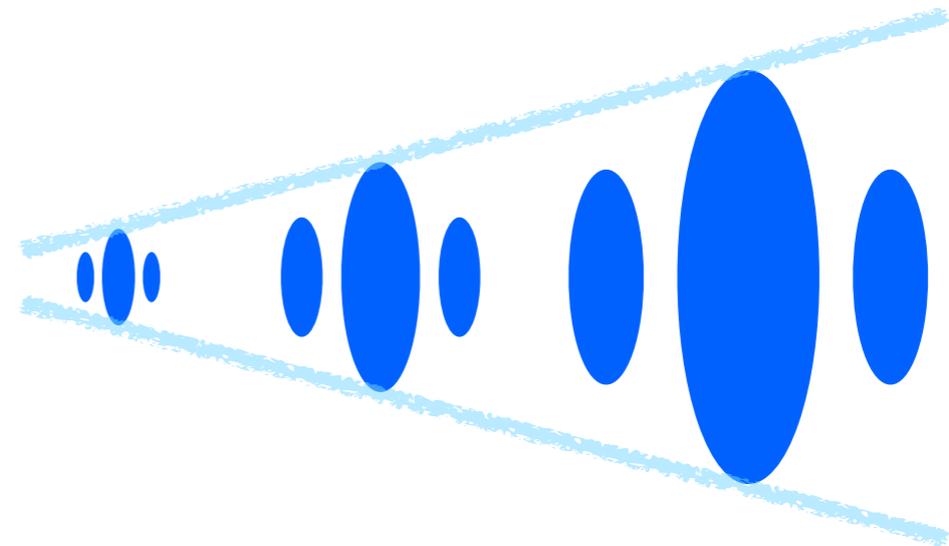
Prior information?

Example: sparse images in wavelets

Representing this
image ...



\mathcal{X}

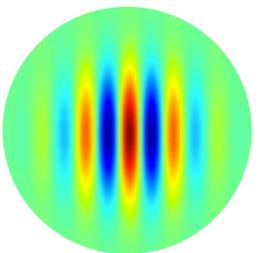
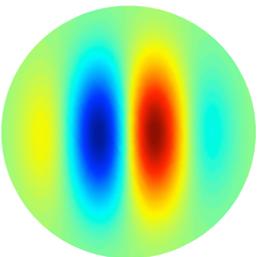


different sizes, scales



different orientations

e.g.,



Compressed Sensing...

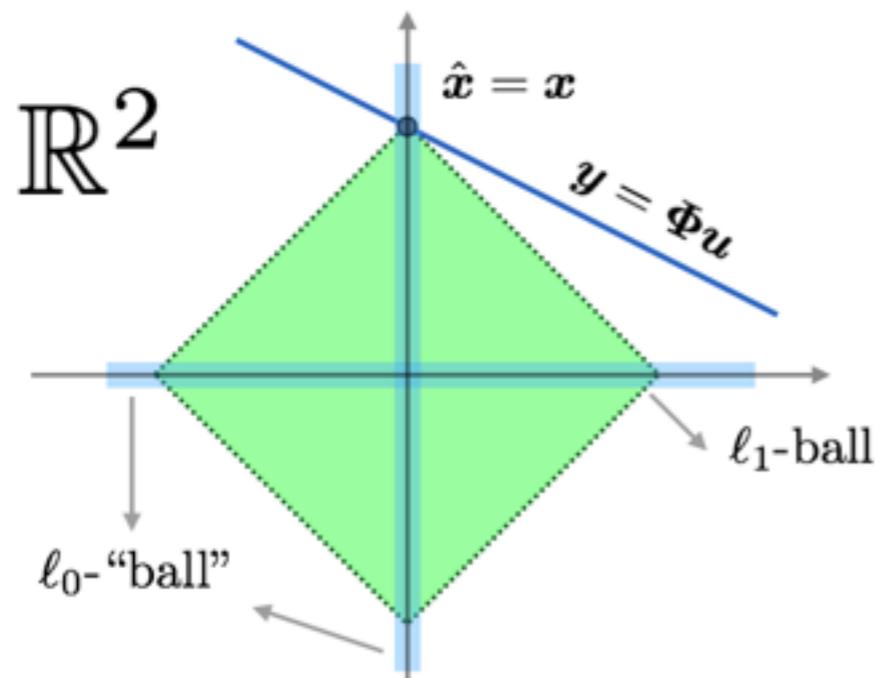
Use non-linear reconstruction methods:

Basis Pursuit DeNoise [Chen, Donoho, Saunders, 98]

$$\hat{\mathbf{x}} \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \|\mathbf{y} - \Phi\mathbf{u}\| \leq \epsilon$$

Sparsity promotion
 $\|\mathbf{u}\|_1 = \sum_j |u_j|$

Level of "noise"
 $\mathbf{y} = \Phi\mathbf{x} + \mathbf{n}, \|\mathbf{n}\| \leq \epsilon$



Many toolboxes:

SPGL1, Numerical Tours, ...

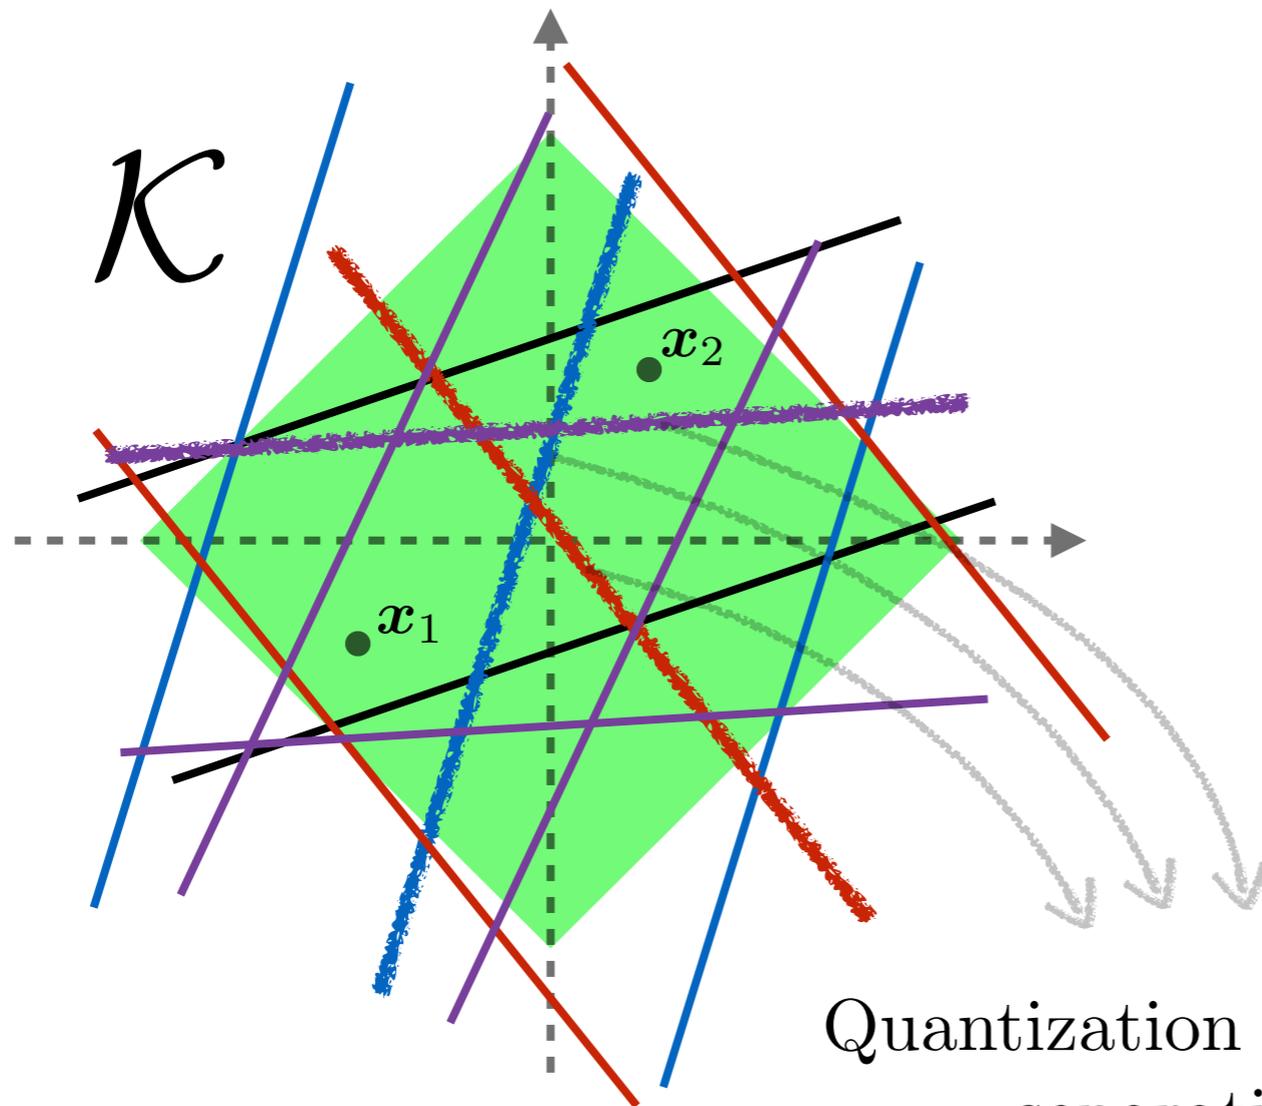
+ Many other algorithms
(ex. greedy algorithms)

ℓ_1 -ball in high dimension
 \approx set of bounded sparse signals

*: with other norms for other low-complexity sets (e.g., nuclear norm)

Properties of $A(\mathbf{x}) := \mathcal{Q}(\Phi\mathbf{x} + \xi)$

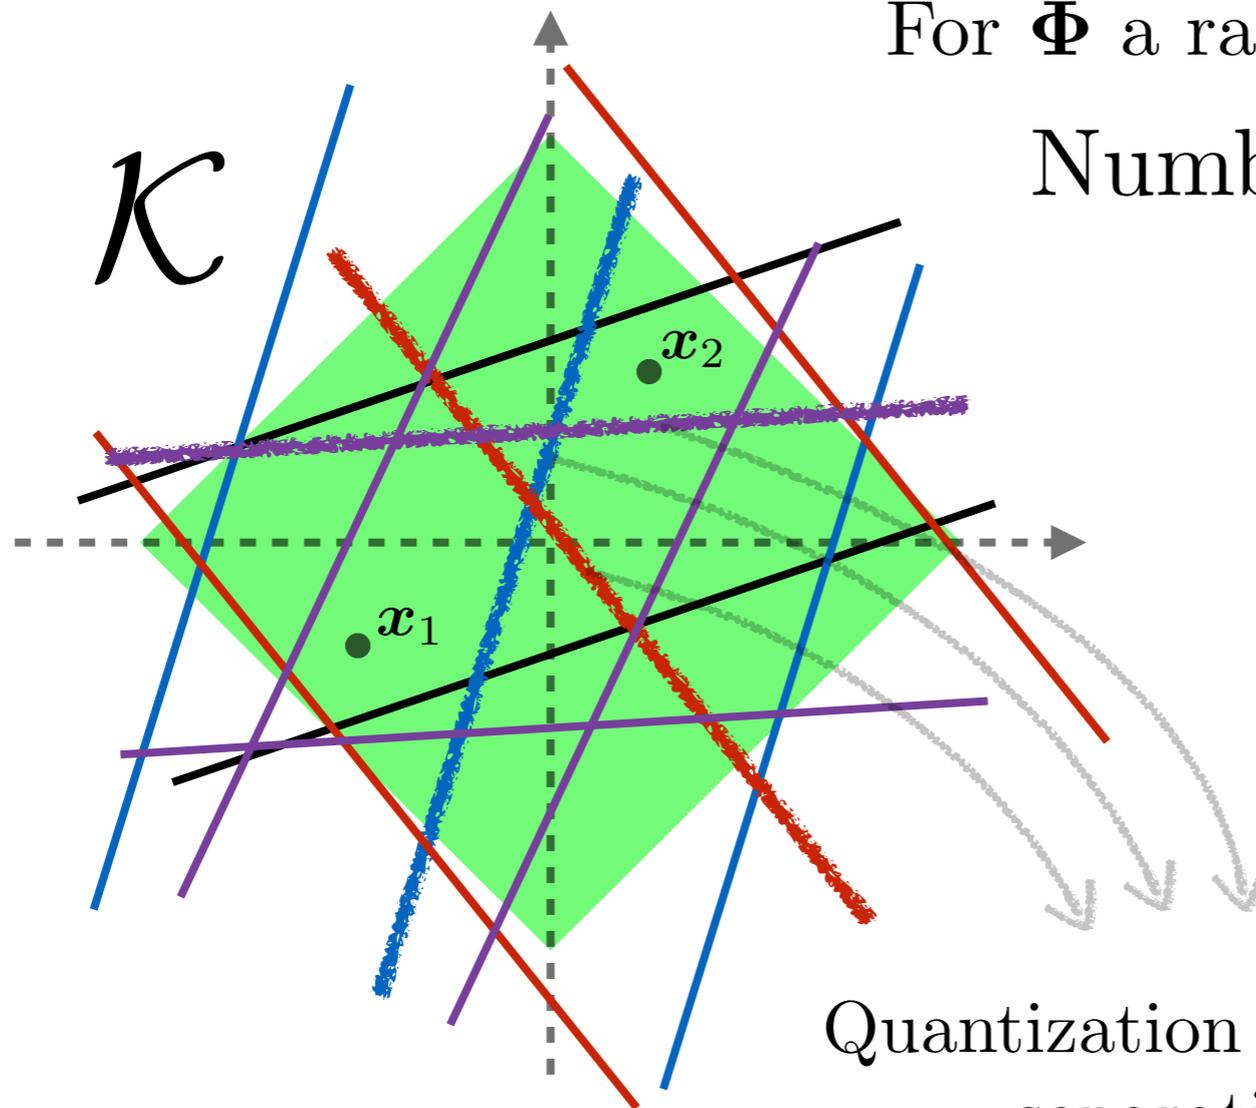
- ▶ 3. Quantizing the RIP (approximate consistency)



Quantization frontiers separating \mathbf{x}_1 and \mathbf{x}_2
= separating random hyperplanes oriented
and positioned according to (Φ, ξ)

Properties of $A(\mathbf{x}) := \mathcal{Q}(\Phi\mathbf{x} + \xi)$

- ▶ 3. Quantizing the RIP (approximate consistency)



For Φ a random Gaussian matrix,

Number of such hyperplanes

$$\approx \|\mathbf{x}_1 - \mathbf{x}_2\|$$

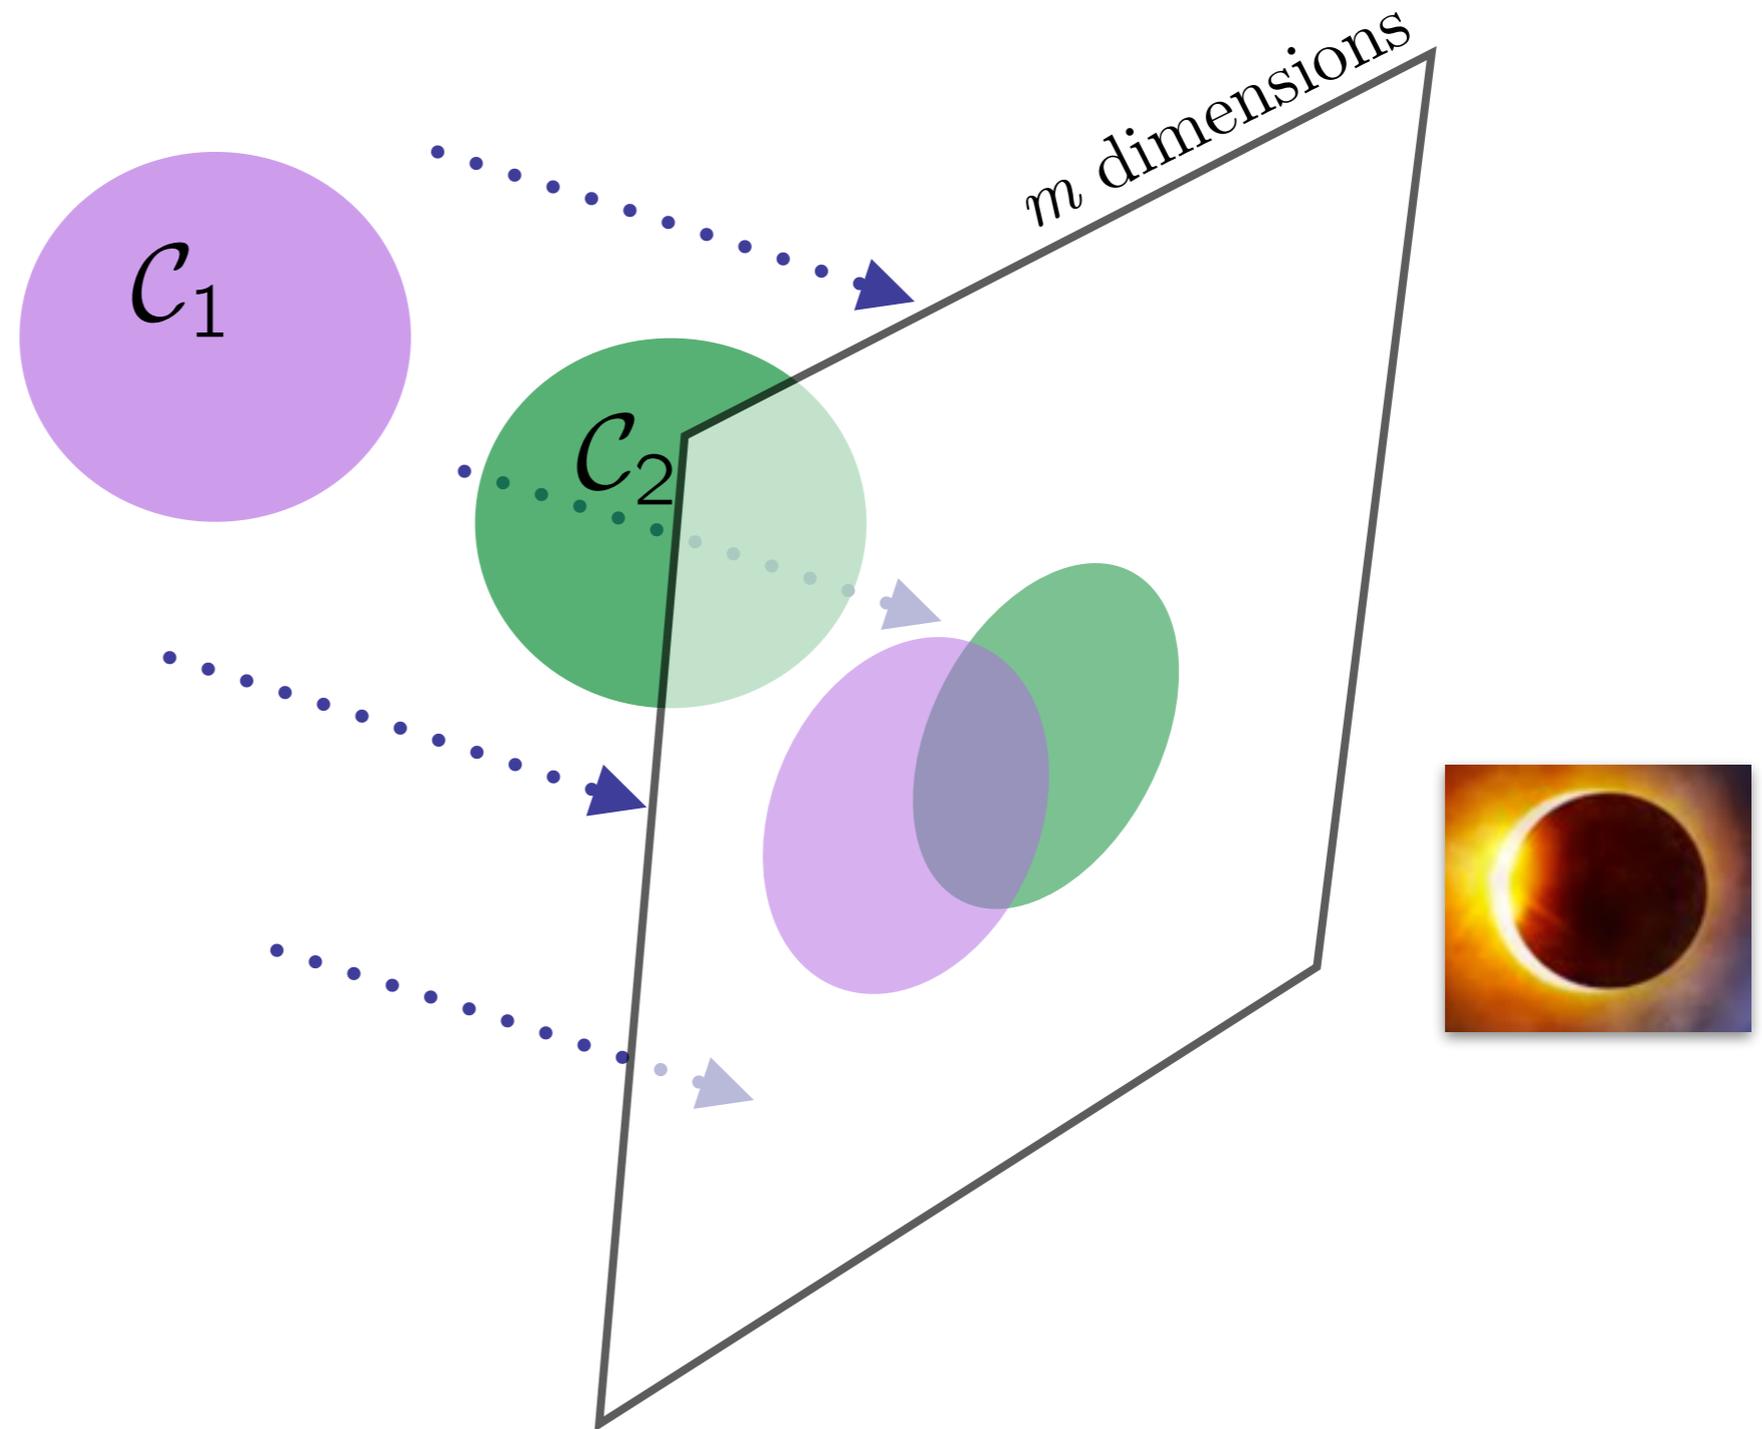
with high probability

provided $M \geq C_{\mathcal{K}} \epsilon^{-2} \log(1/\delta\epsilon^2)$

(e.g., for sparse signals,
low-rank matrices)

Quantization frontiers separating \mathbf{x}_1 and \mathbf{x}_2
= separating random hyperplanes oriented
and positioned according to (Φ, ξ)

The Big Picture (an easy classification problem)



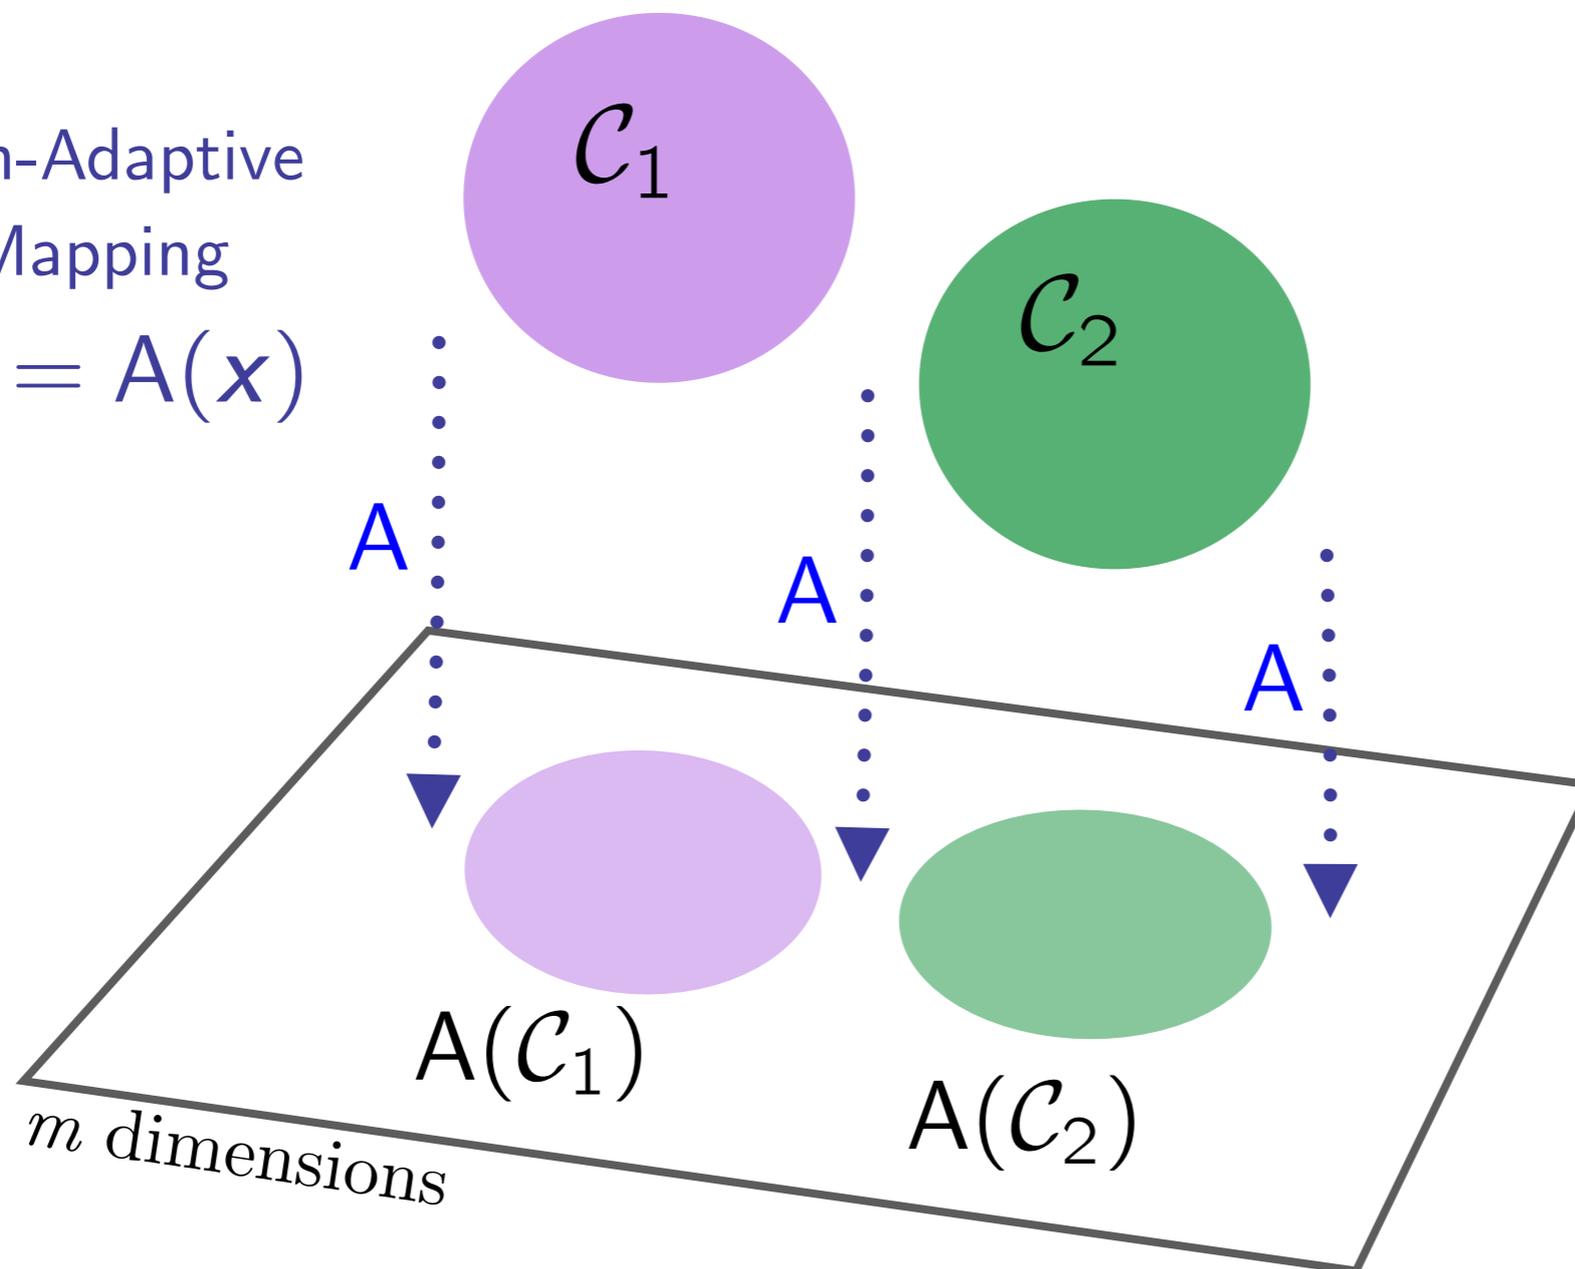
The Big Picture (an easy classification problem)

$\mathcal{K} \subset \mathbb{R}^n$ dataset

$\mathcal{C}_i \subset \mathcal{K}$ classes, $i = 1, 2, \dots$

Non-Adaptive
Mapping

$$y = A(x)$$



Separable, "Success"



First attempts [Candès, Tao, 04]

- ▶ Quantization is like a noise! (e.g., for $Q[\lambda] = \delta \lfloor \frac{\lambda}{\delta} \rfloor \in \delta\mathbb{Z}$)

$$\mathbf{y} = Q(\Phi\mathbf{x}) = \Phi\mathbf{x} + \mathbf{n}, \quad \text{with } \mathbf{n} = Q(\Phi\mathbf{x}) - \Phi\mathbf{x}.$$
$$\text{and } \|\mathbf{n}\|^2 = O(m\delta^2)$$

- ▶ Problem: (e.g., for $\hat{\mathbf{x}} \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1$ s.t. $\|\mathbf{y} - \Phi\mathbf{u}\| \leq \epsilon$, aka BPDN)

$$\|\mathbf{x} - \hat{\mathbf{x}}\| \lesssim \frac{\epsilon}{\sqrt{m}} = O(\delta) \text{ does not decay if } m \text{ increases!}$$

Remark: r -order $\Sigma\Delta$ quantizer + CS achieves anyway

$$\|\mathbf{x} - \hat{\mathbf{x}}\| = O(M^{-r+1/2}),$$

and, in specific cases, exponentially decaying rate.

(see the works of, e.g., [Gunturk, Lammers, Powell, Saab, Yilmaz])

but, more complex implementation (memory)

Properties of $A(\mathbf{x}) := \mathcal{Q}(\Phi\mathbf{x} + \boldsymbol{\xi})$

- ▶ 2. For a non-consistent reconstruction method:

Given $\frac{1}{\sqrt{m}}\Phi$ RIP over Σ_{4K} , and $\xi_i \sim_{\text{iid}} \mathcal{U}([0, \delta])$,
for all $\mathbf{x} \in \Sigma_K$ observed through $\mathbf{y} = A(\mathbf{x}) = \mathcal{Q}(\Phi\mathbf{x} + \boldsymbol{\xi})$,

The Projected Back Projection estimate

$$\hat{\mathbf{x}} = \mathcal{H}_K\left(\frac{1}{m}\Phi^\top \mathbf{y}\right),$$

is such that

$$\|\mathbf{x} - \hat{\mathbf{x}}\| \leq C_K(1 + \delta)M^{-1/2},$$

with very high probability.

Really simple
(almost dumb)
reconstruction

This extends to other low-complexity sets
and it works for any RIP matrix!

Properties of $A(\mathbf{x}) := \mathcal{Q}(\Phi \mathbf{x} + \xi)$

- 2. For a non-consistent reconstruction method:

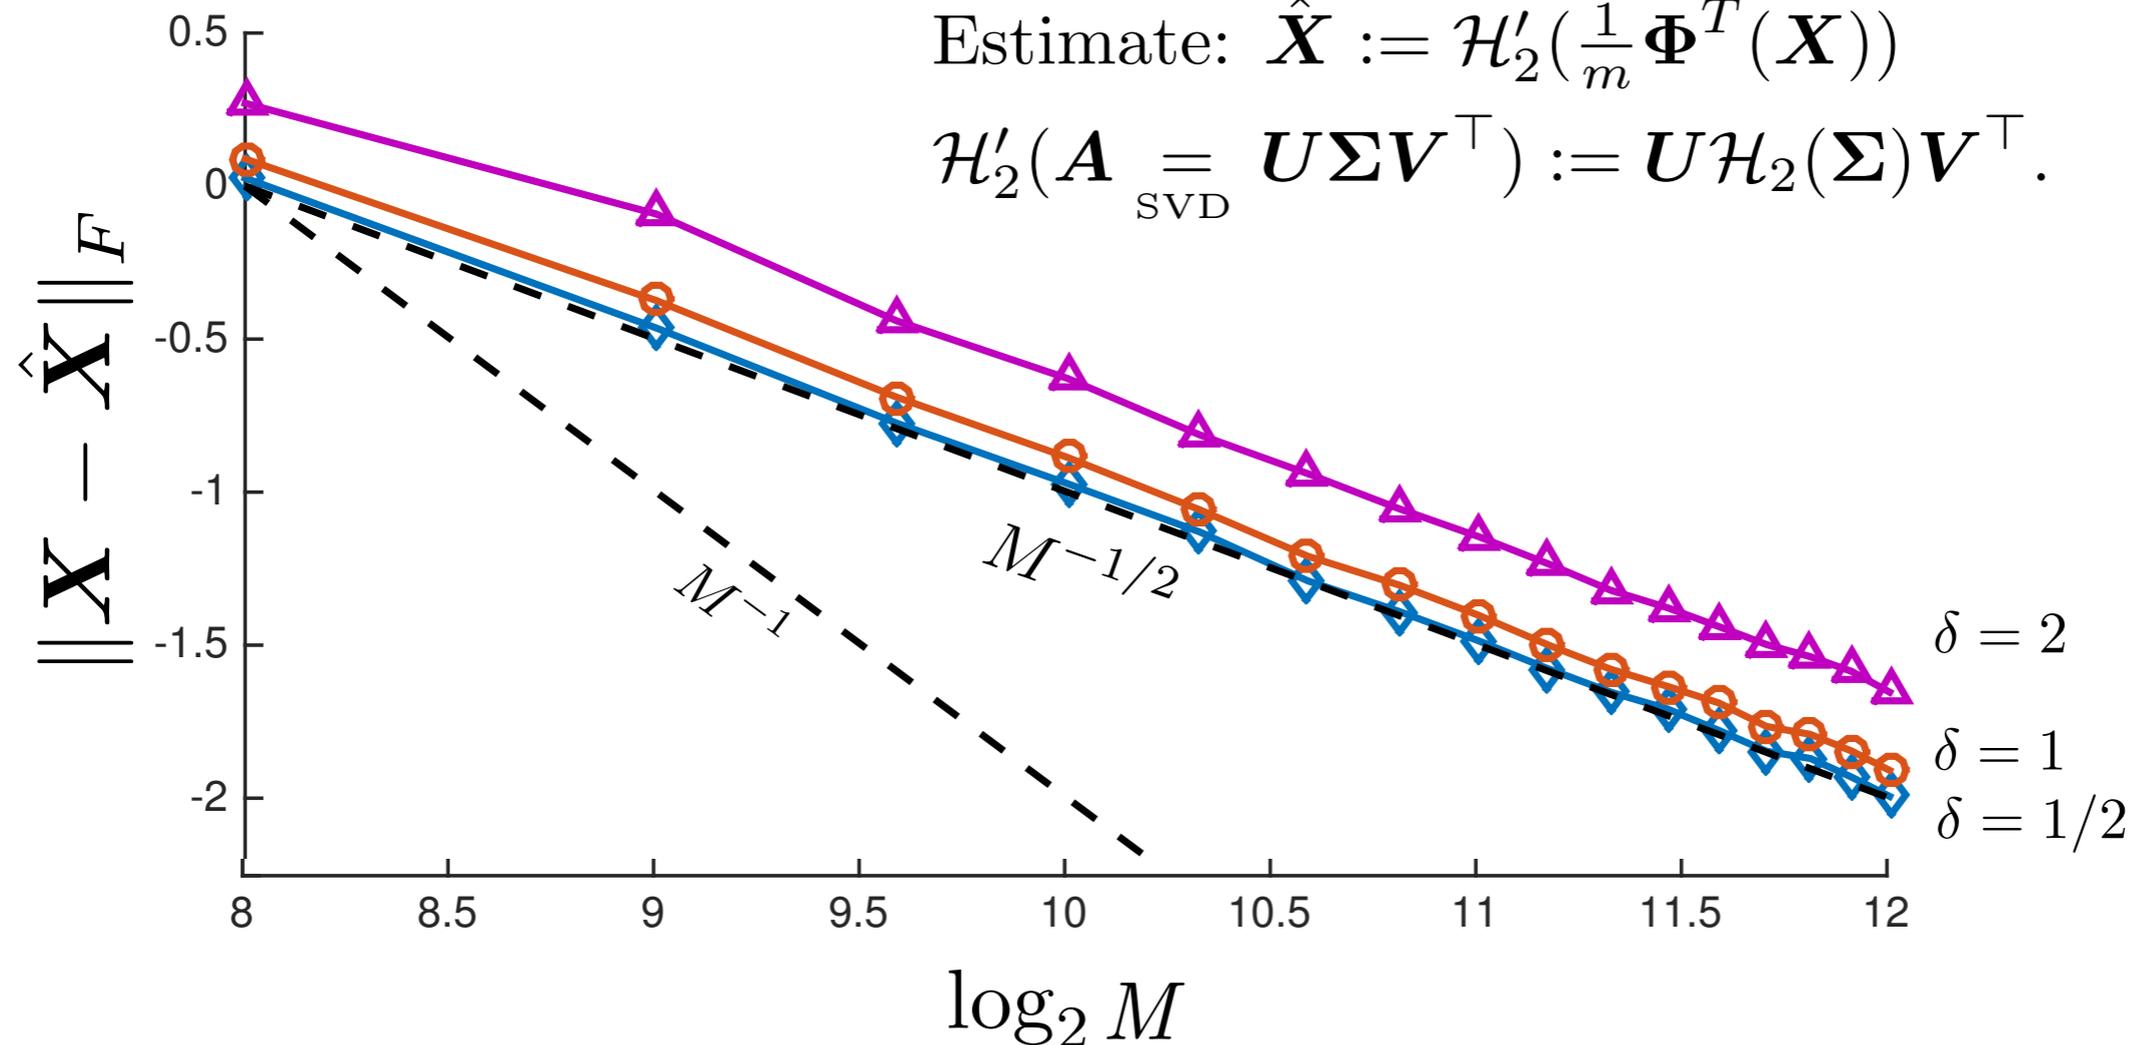
Example:

\mathbf{X} : Rank-2 64×64 -matrix ($n = 64^2$, vectorized)

Φ : Gaussian random matrix

Estimate: $\hat{\mathbf{X}} := \mathcal{H}'_2\left(\frac{1}{m}\Phi^T(\mathbf{X})\right)$

$\mathcal{H}'_2(\mathbf{A} \underset{\text{SVD}}{=} \mathbf{U}\Sigma\mathbf{V}^\top) := \mathbf{U}\mathcal{H}_2(\Sigma)\mathbf{V}^\top.$



Compressed Sensing...

Use non-linear reconstruction methods:

Basis Pursuit DeNoise [Chen, Donoho, Saunders, 98]

$$\hat{\mathbf{x}} \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \|\mathbf{y} - \Phi\mathbf{u}\| \leq \epsilon$$

Level of "noise"
 $\mathbf{y} = \Phi\mathbf{x} + \mathbf{n}, \|\mathbf{n}\| \leq \epsilon$

$\|\mathbf{u}\|_1 = \sum_j |u_j|$
Sparsity promotion

Solved by many toolboxes:

SPGL1, Numerical Tours, CVX ...

+ Many other algorithms
(ex. greedy algorithms)



ℓ_1 -ball in high dimension

Compressive Sampling and Quantization

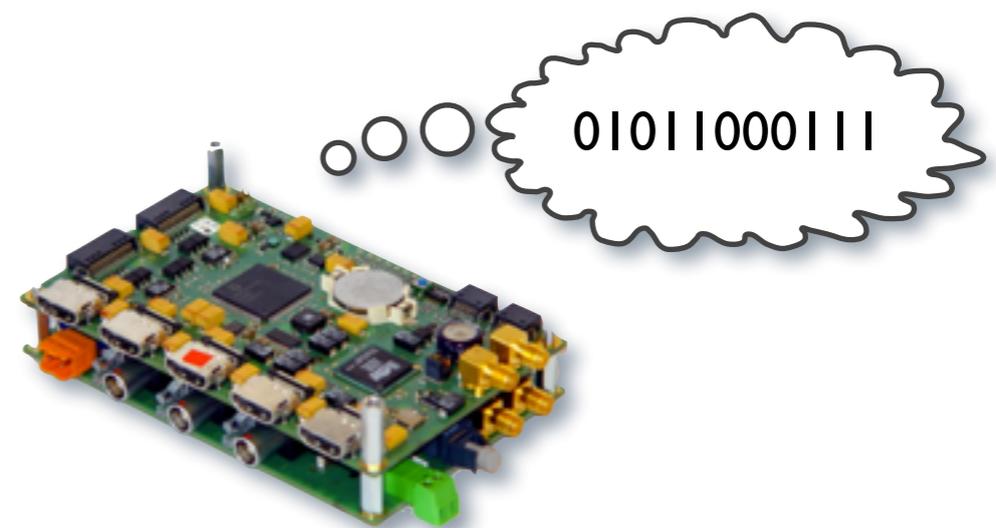
Compressed sensing theorist says:

*“Linearly sample a signal
at a rate function of
its intrinsic dimensionality”*



Information theorist and sensor designer say:

*“Okay, but I need to quantize/digitize my measurements!”
(e.g., in ADC)*



Integration?

QCS theory?

Theoretical Bounds

The Rare Eclipse Problem (Linear case)

Adapted* from [Bandeira, Mixon, Recht 14]

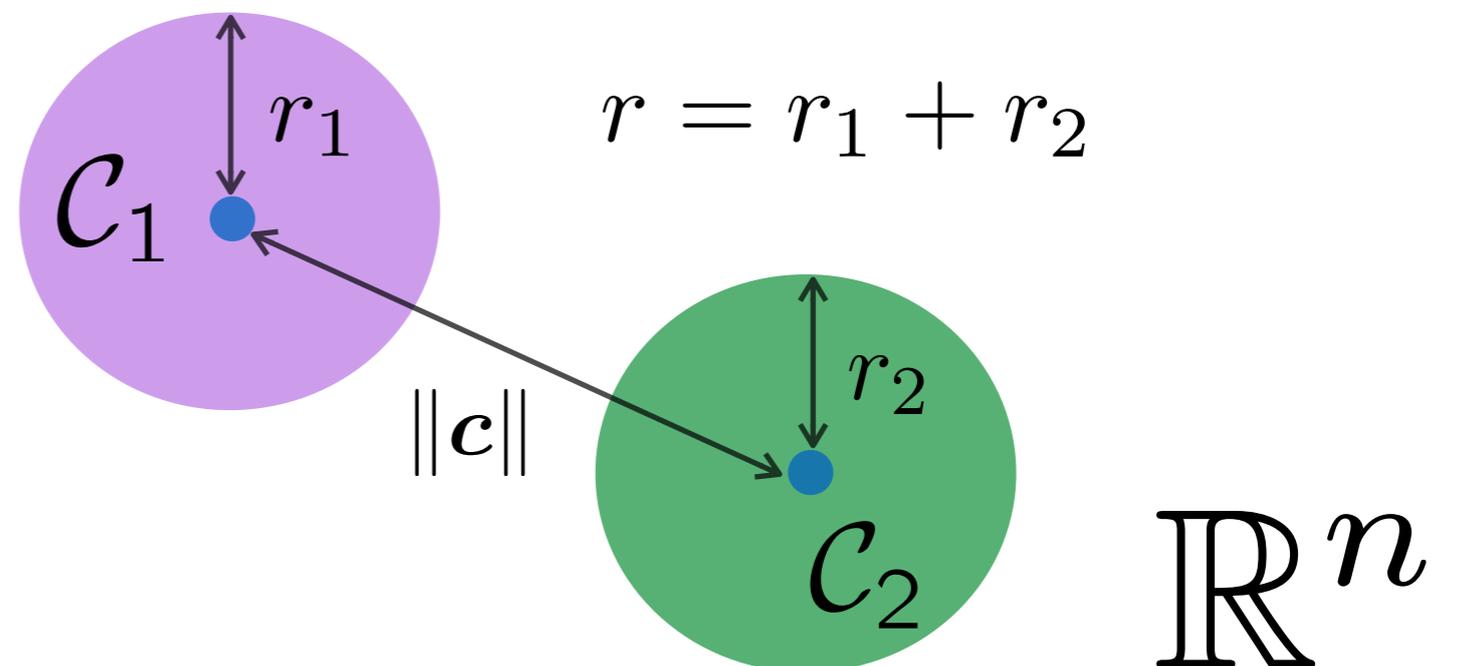
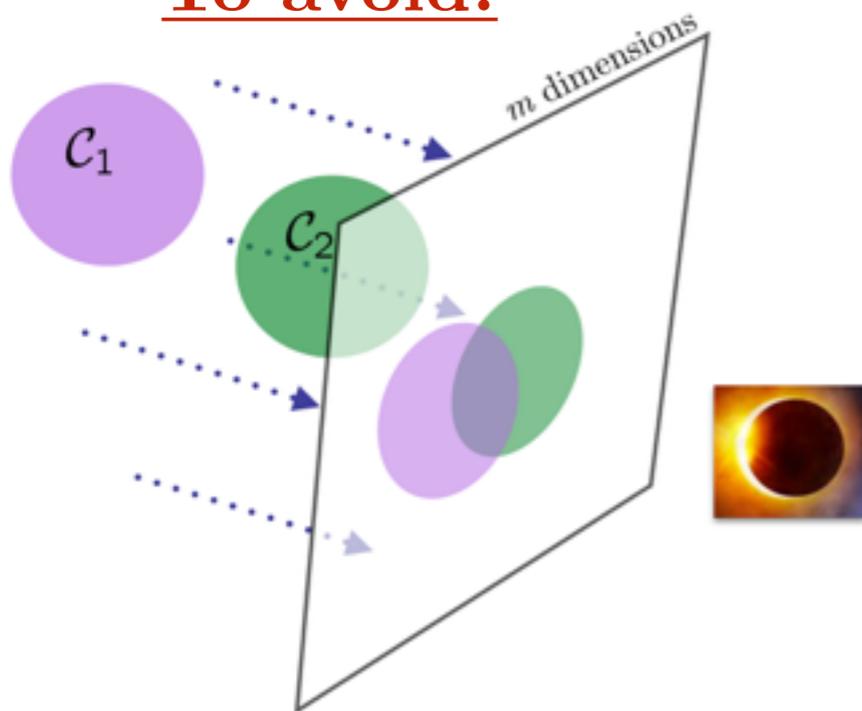
Given $0 < \eta < 1$ and $\Phi \in \mathbb{R}^{m \times n}$ a random Gaussian matrix,
if

$$m \gtrsim \frac{r^2}{\|\mathbf{c}\|^2} n + \log \frac{1}{\eta},$$

then

$$\mathbb{P}[\Phi\mathcal{C}_1 \cap \Phi\mathcal{C}_2 = \emptyset] \geq 1 - \eta.$$

To avoid!



*: what is presented here is a special case.