

Time for dithering! Quantized random embeddings with RIP random matrices

Chunlei Xu, Vincent Schellekens, Amirafshar Moshtaghpour, Valerio Cambareri, Laurent Jacques



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




Outline

- ▶ 1. Brief introduction to
 - ▶ compressive sensing (CS)
 - ▶ & quantized CS (QCS)
- ▶ 2. Quantized dithered random mapping
 - ▶ Dimensionality reduction
 - ▶ Recovering low-complexity vectors in QCS with any RIP matrix
 - ▶ Classification in a quantized world
- ▶ Conclusion

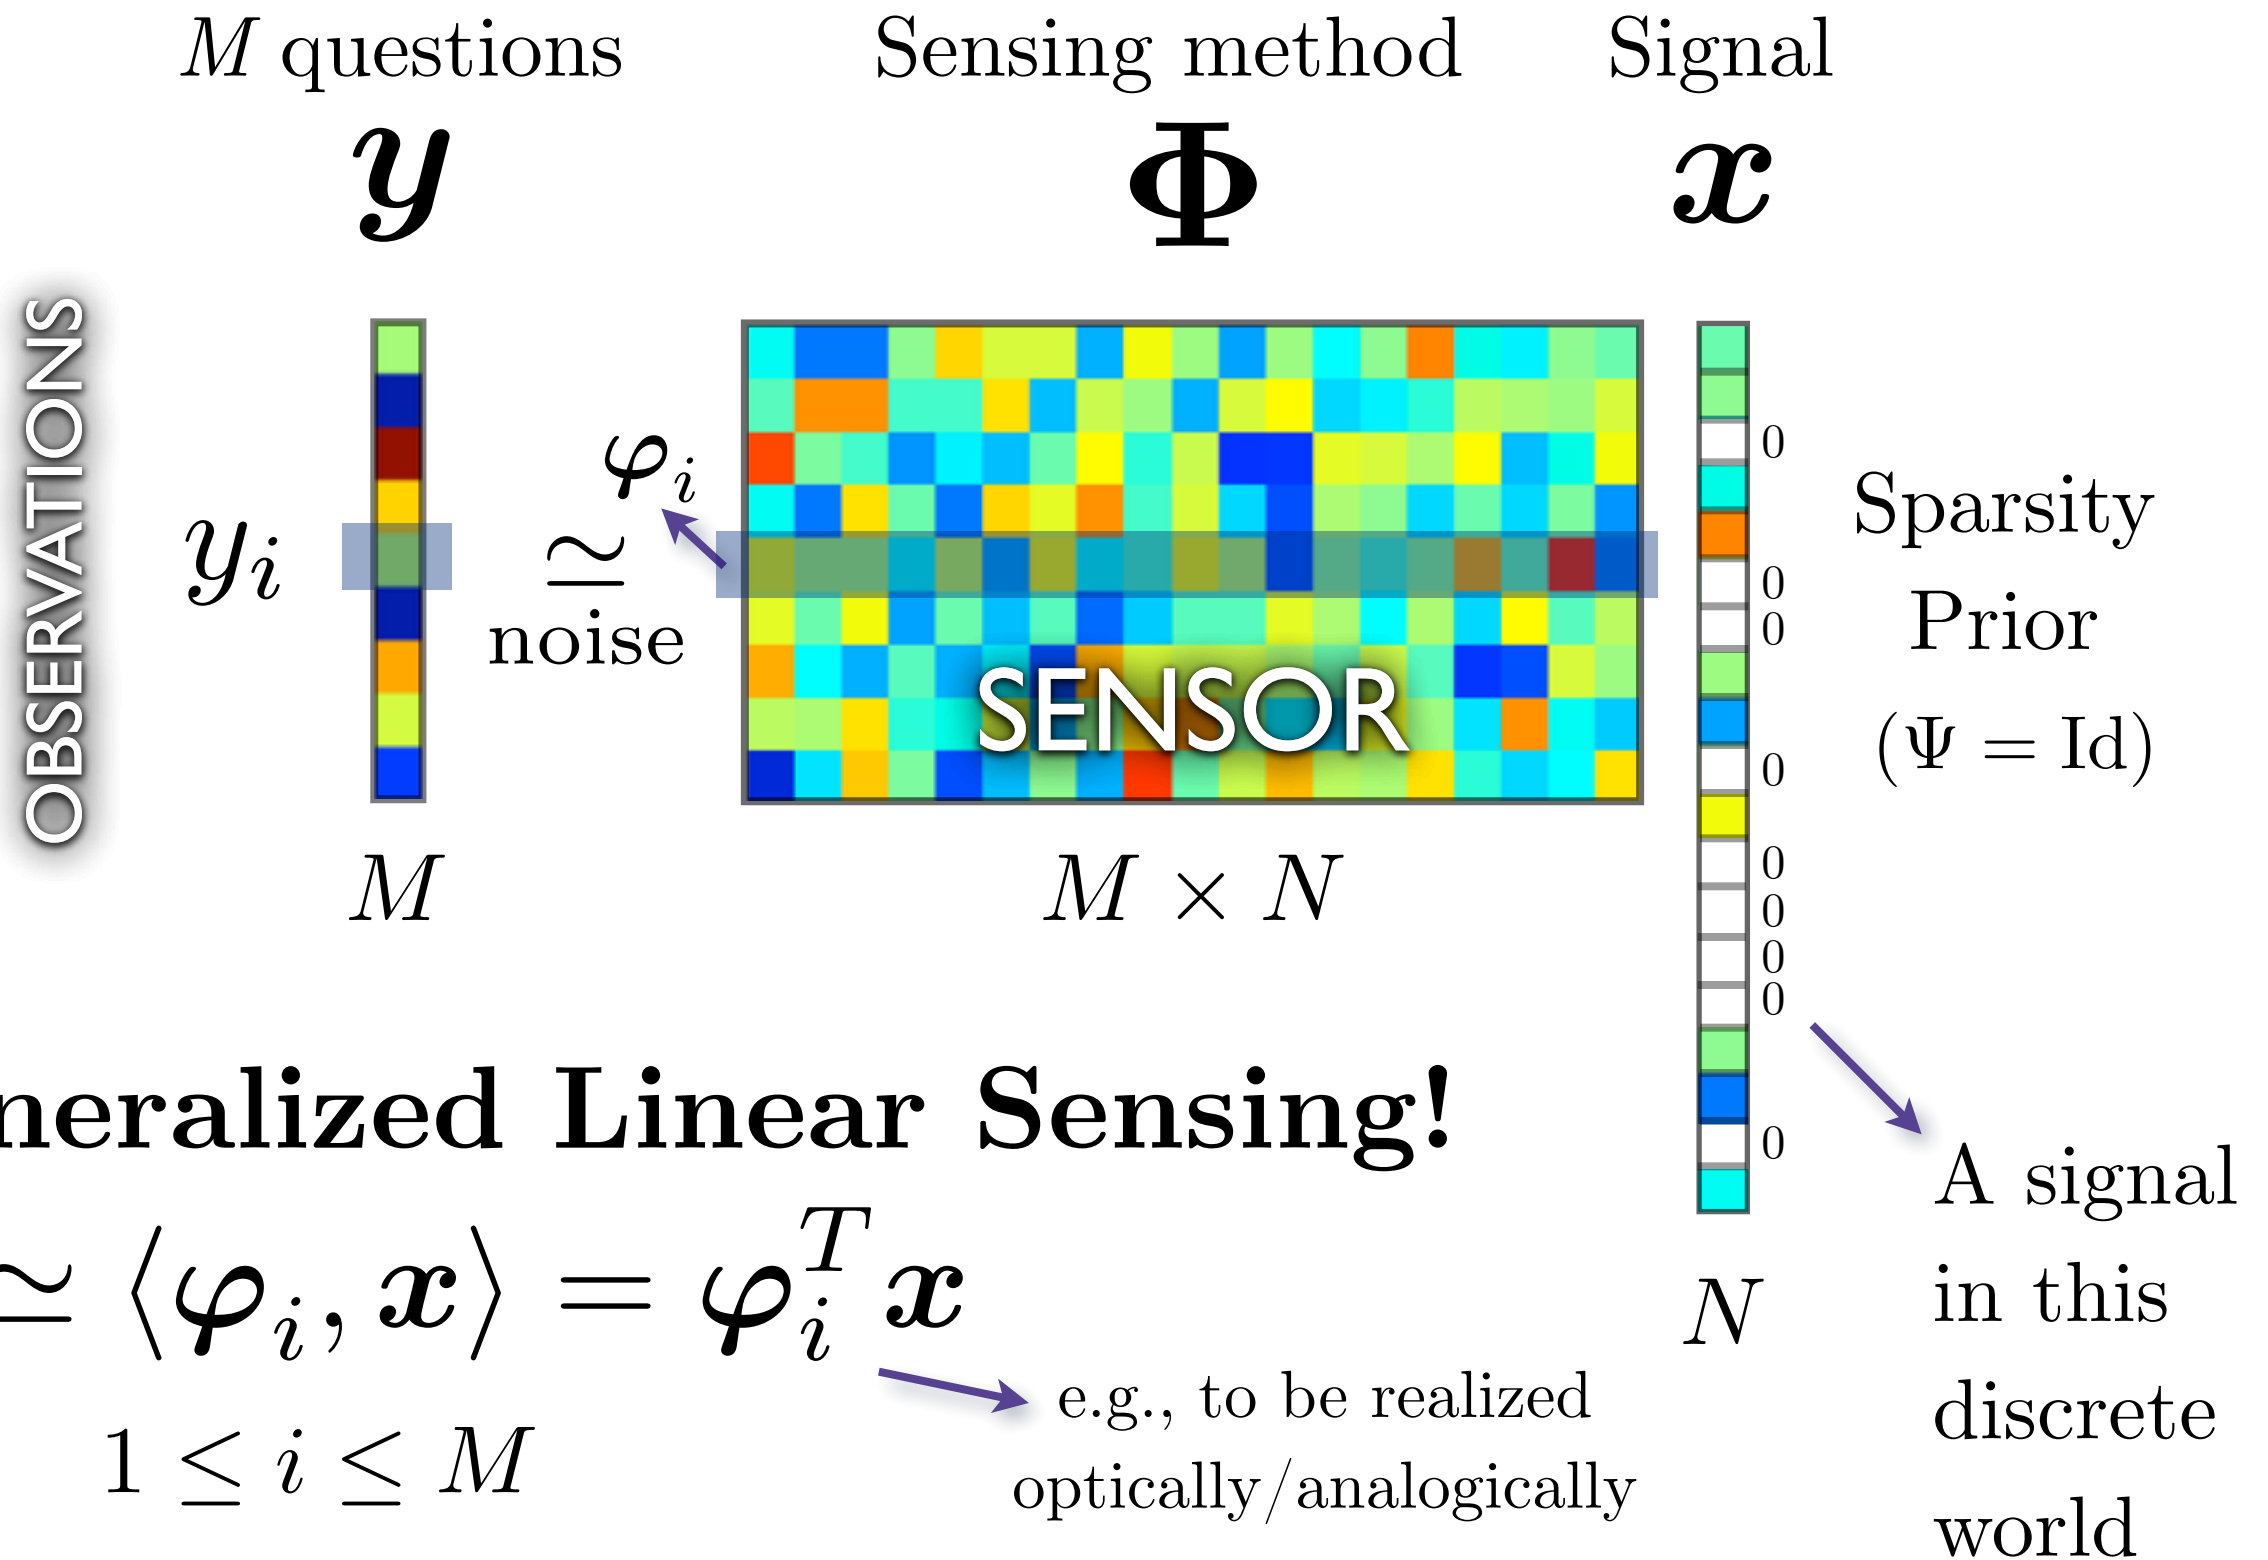
1. Brief introduction to CS & QCS



Compressive sensing...



Compressive sensing...



Identifiability of x from Φx ?

Compressive sensing of sparse signals


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,

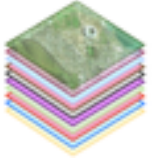
Geometry of $\Phi(\Sigma_K)$
 \approx Geometry of Σ_K

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

observations true signals




Compressive sensing of *l.c.* signals

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,

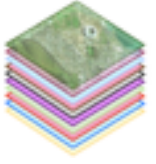
Geometry of $\Phi(\mathcal{K})$
 \approx Geometry of \mathcal{K}

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



For instance: $C_{\mathcal{K}} = w^2(\mathcal{K})$, the Gaussian mean width, *i.e.*, $C_{\mathcal{K}} \lesssim K \log N/K$ for k -sparse vectors, or $C_{\mathcal{K}} \lesssim rn$ for rank- r $n \times n$ matrices.

Compressive sensing of *l.c.* signals

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,

Restricted Isometry Property (RIP)

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

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

Signal reconstruction: ensured with if RIP matrix and with non-linear methods, e.g., Basis Pursuit DeNoise (BPDN), greedy methods (MP, OMP, ...).


Random sensing matrix market?

Dense & unstructured sensing matrices:

- ▶ 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




Structured sensing matrices (less memory, fast computations):

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

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

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




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

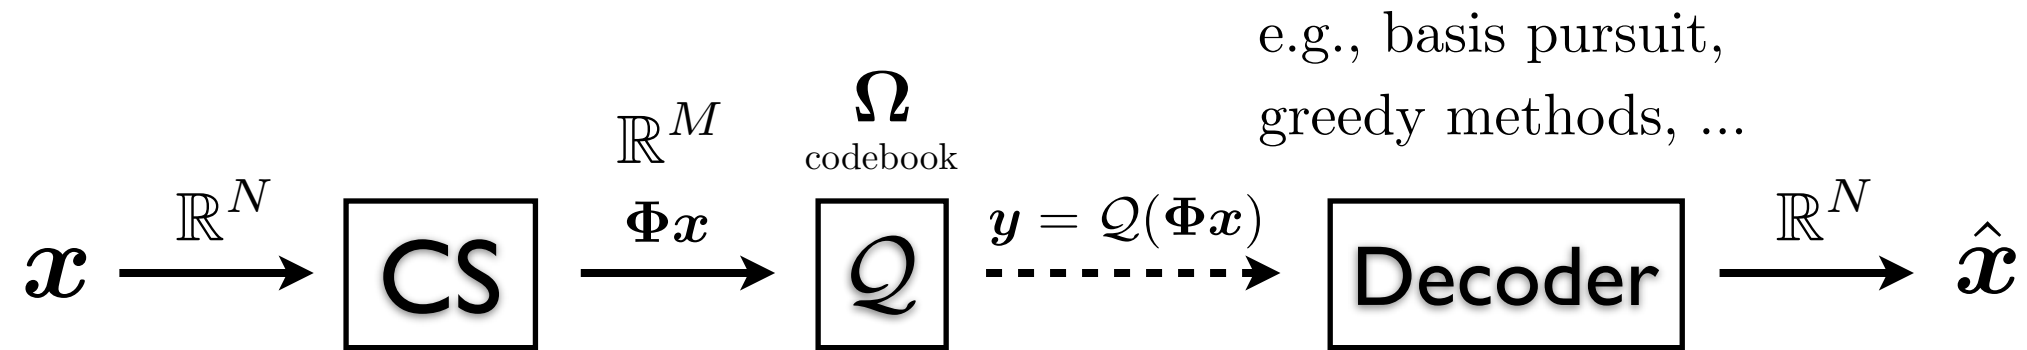
(see, e.g., [Foucart, Rauhut, 2013])

- ▶ or random sensing with natural processes

(e.g., LightOn!)



Quantizing CS?




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 per meas.




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


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



Other examples (not covered here):

Non-regular, e.g., square wave (or LSB)

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



Non-uniform scalar quantizer, vector quantizer, $\Sigma\Delta$ quantizer/noise shaping, ... (see the works of, e.g., [\[Gunturk, Lammers, Powell, Saab, Yilmaz, Goyal\]](#))

QCS, first attempt [Candès, Tao, 04]

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

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

- ▶ **Problem**: e.g., for BPDN,


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

counterintuitive?

QCS, first attempt [Candès, Tao, 04]

Cause of the problem:

Quantization is discontinuous (it does not “dither”)




$$\exists \mathbf{x} \neq \mathbf{x}' : \mathcal{Q}(\Phi \mathbf{x}) = \mathcal{Q}(\Phi \mathbf{x}') \Rightarrow \|\mathbf{x} - \hat{\mathbf{x}}\| > C$$

(e.g., take $\Phi_{ij} \in \{\pm 1\}$ and $\mathcal{Q} = \text{sign}$)!

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

Problem with too sparse signals!

2. Quantized dithered random mapping






The power of dithering (an old trick, revisited)

- Inject a pre-quantization, uniform “noise”:
i.e., a dithering $\xi \in \mathbb{R}^m$ with $\xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta])$ (your friend)

→ The good boy!

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




The power of dithering (an old trick, revisited)


- Inject a pre-quantization, uniform “noise”:
i.e., a dithering $\xi \in \mathbb{R}^m$ with $\xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta])$ (your friend)

→ The good boy!

$$A(x) := Q(\Phi x + \xi) \quad (\text{QDRM})$$



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




The power of dithering (an old trick, revisited)

- ▶ Inject a pre-quantization, uniform “noise”:
i.e., a dithering $\xi \in \mathbb{R}^m$ with $\xi_j \sim_{\text{iid}} \mathcal{U}([0, \delta])$ (your friend)

→ The good boy!


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



- ▶ Motivation? $\mathbb{E}_{\xi} \mathcal{Q}(u + \xi) = u$
 $\Rightarrow A(x) \approx \Phi x$ if M large
- ▶ Possibility to define
quantized dimensionality reduction/embedding!

2.1. Quantized Dimensionality Reduction


Control of the “consistency width”



low complexity set \mathcal{K}

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

Control of the “consistency width”




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

Signals \mathbf{u} s.t.

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

$$\delta \lfloor (\varphi_1^T \mathbf{u} + \xi_1) / \delta \rfloor$$

Control of the “consistency width”




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

Signals u s.t.

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

Control of the “consistency width”




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

Signals \mathbf{u} s.t.

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

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

Control of the “consistency width”



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

For Φ a random Gaussian matrix, with high probability,

[LJ, 16], [LJ, 17]


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

with $q = 1$ (for, *e.g.*, sparse signals, low-rank matrices), or $q = 4$ for convex sets.

Open problem:
Extension to RIP matrices?


Quantizing the RIP (approximate consistency)

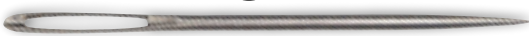
$$A(x) := \mathcal{Q}(\Phi x + \xi)$$




Quantizing the RIP (approximate consistency)

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



Length of  = $\|x_1 - x_2\|$

(thanks to the dithering)
Buffon's needle problem




$$\mathbb{E}(\text{intersections}) \propto \text{length}$$

<http://www.buffon.cnrs.fr>

(In 1733)

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)

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

||| ??



Buffon's needle problem



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

$$\mathbb{E}(\text{intersections}) \propto \text{length}$$

<http://www.buffon.cnrs.fr>
 (In 1733)

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]

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 $M \gtrsim \epsilon^{-2} C_K \log(1 + \frac{1}{\delta\epsilon})$, (with $C_K > 0$ 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 $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{K} \cap \mathbb{B}^N$.

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

[LJ, Cambareri, 17]

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 $M \gtrsim \epsilon^{-2} C_K \log(1 + \frac{1}{\delta\epsilon})$, (with $C_K > 0$ 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 $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{K} \cap \mathbb{B}^N$.

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


Dimensionality reduction!

Classification?


[LJ, Cambareri, 17]

2.2. Recovering low-complexity vectors in QCS with any RIP matrix

The reconstruction problem



Projected Back Projection (PBP)



$$\hat{x} = \mathcal{P}_{\mathcal{K}}\left(\frac{1}{M} \Phi^{\top} A(x)\right)$$

(1st iteration of other potential methods)

PBP Error Analysis

Limited projection property (LPD):

A respect the $\text{LPD}(\mathcal{K}, \Phi, \nu)$ if

$$\frac{1}{M} \left| \langle A(\mathbf{u}), \Phi \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle \right| \leq \nu, \quad \forall \mathbf{u}, \mathbf{v} \in \mathcal{K} \cap \mathbb{B}^N.$$

\equiv How close A is from Φ

If $\frac{1}{\sqrt{M}} \Phi$ is $\text{RIP}(\Sigma_{2K}, \epsilon)$ & A is $\text{LPD}(\Sigma_{2K}, \Phi, \nu)$, then

$$\|\mathbf{x} - \hat{\mathbf{x}}\| \leq 2(\epsilon + \nu).$$

(same result for, e.g., union of low-dimensional spaces, low-rank matrices,
and convex sets with square rooted error)

PBP Error Analysis

Limited projection property (LPD):

A respect the $\text{LPD}(\mathcal{K}, \Phi, \nu)$ if

$$\frac{1}{M} \left| \langle A(\mathbf{u}), \Phi \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle \right| \leq \nu, \quad \forall \mathbf{u}, \mathbf{v} \in \mathcal{K} \cap \mathbb{B}^N.$$

\equiv How close A is from Φ ?

If $\frac{1}{\sqrt{M}} \Phi$ is $\text{RIP}(\Sigma_{2K}, \epsilon)$ & A is $\text{LPD}(\Sigma_{2K}, \Phi, \nu)$, then

$$\|\mathbf{x} - \hat{\mathbf{x}}\| \leq 2(\epsilon + \nu).$$

(same result for, e.g., union of low-dimensional spaces, low-rank matrices, and convex sets with square rooted error)

Question: Which matrices do satisfy the LPD? All RIP ones!

If Φ is $\text{RIP}(\Sigma_{2K}, \epsilon)$ and

$$M \gtrsim \epsilon^{-2} K \log(N/K) \log(1 + \epsilon^{-3}) \log(1/\zeta),$$

then A respects $\text{LPD}(\Sigma_{2k}, \Phi, \epsilon)$ with $\text{Pr} \geq 1 - \zeta$.


(extends to union of low-dimensional spaces, low-rank matrices, & convex sets)

Error:
 $O\left(\frac{\sqrt{K}}{\sqrt{M}}\right)$


Error bound of PBP: benefit of dithering

Example: PBP reconstruction over $\mathcal{K} = \Sigma_k^n$ with partial DCT from dithered and non-dithered quantized measurements

$$\begin{cases} \mathbf{x} \in \mathcal{K} \cap \mathbb{B}^n, \quad n = 512, k = 4 \\ \Phi \text{ is a random partial DCT, } \xi_i \sim_{iid} U([0, \delta]) \\ \delta = 0.5 \text{ (diamond)}, \delta = 1 \text{ (circle)}, \delta = 2 \text{ (triangle)} \end{cases}$$



dithered quan. measurements



non-dithered quan. measurements


Remark : thanks to dithering, reconstruction error decays as m increases.

2.3. Classification in a quantized world

(caveat: in this part, $M \rightarrow m$, $N \rightarrow n$)

The Big Picture (an easy classification problem)

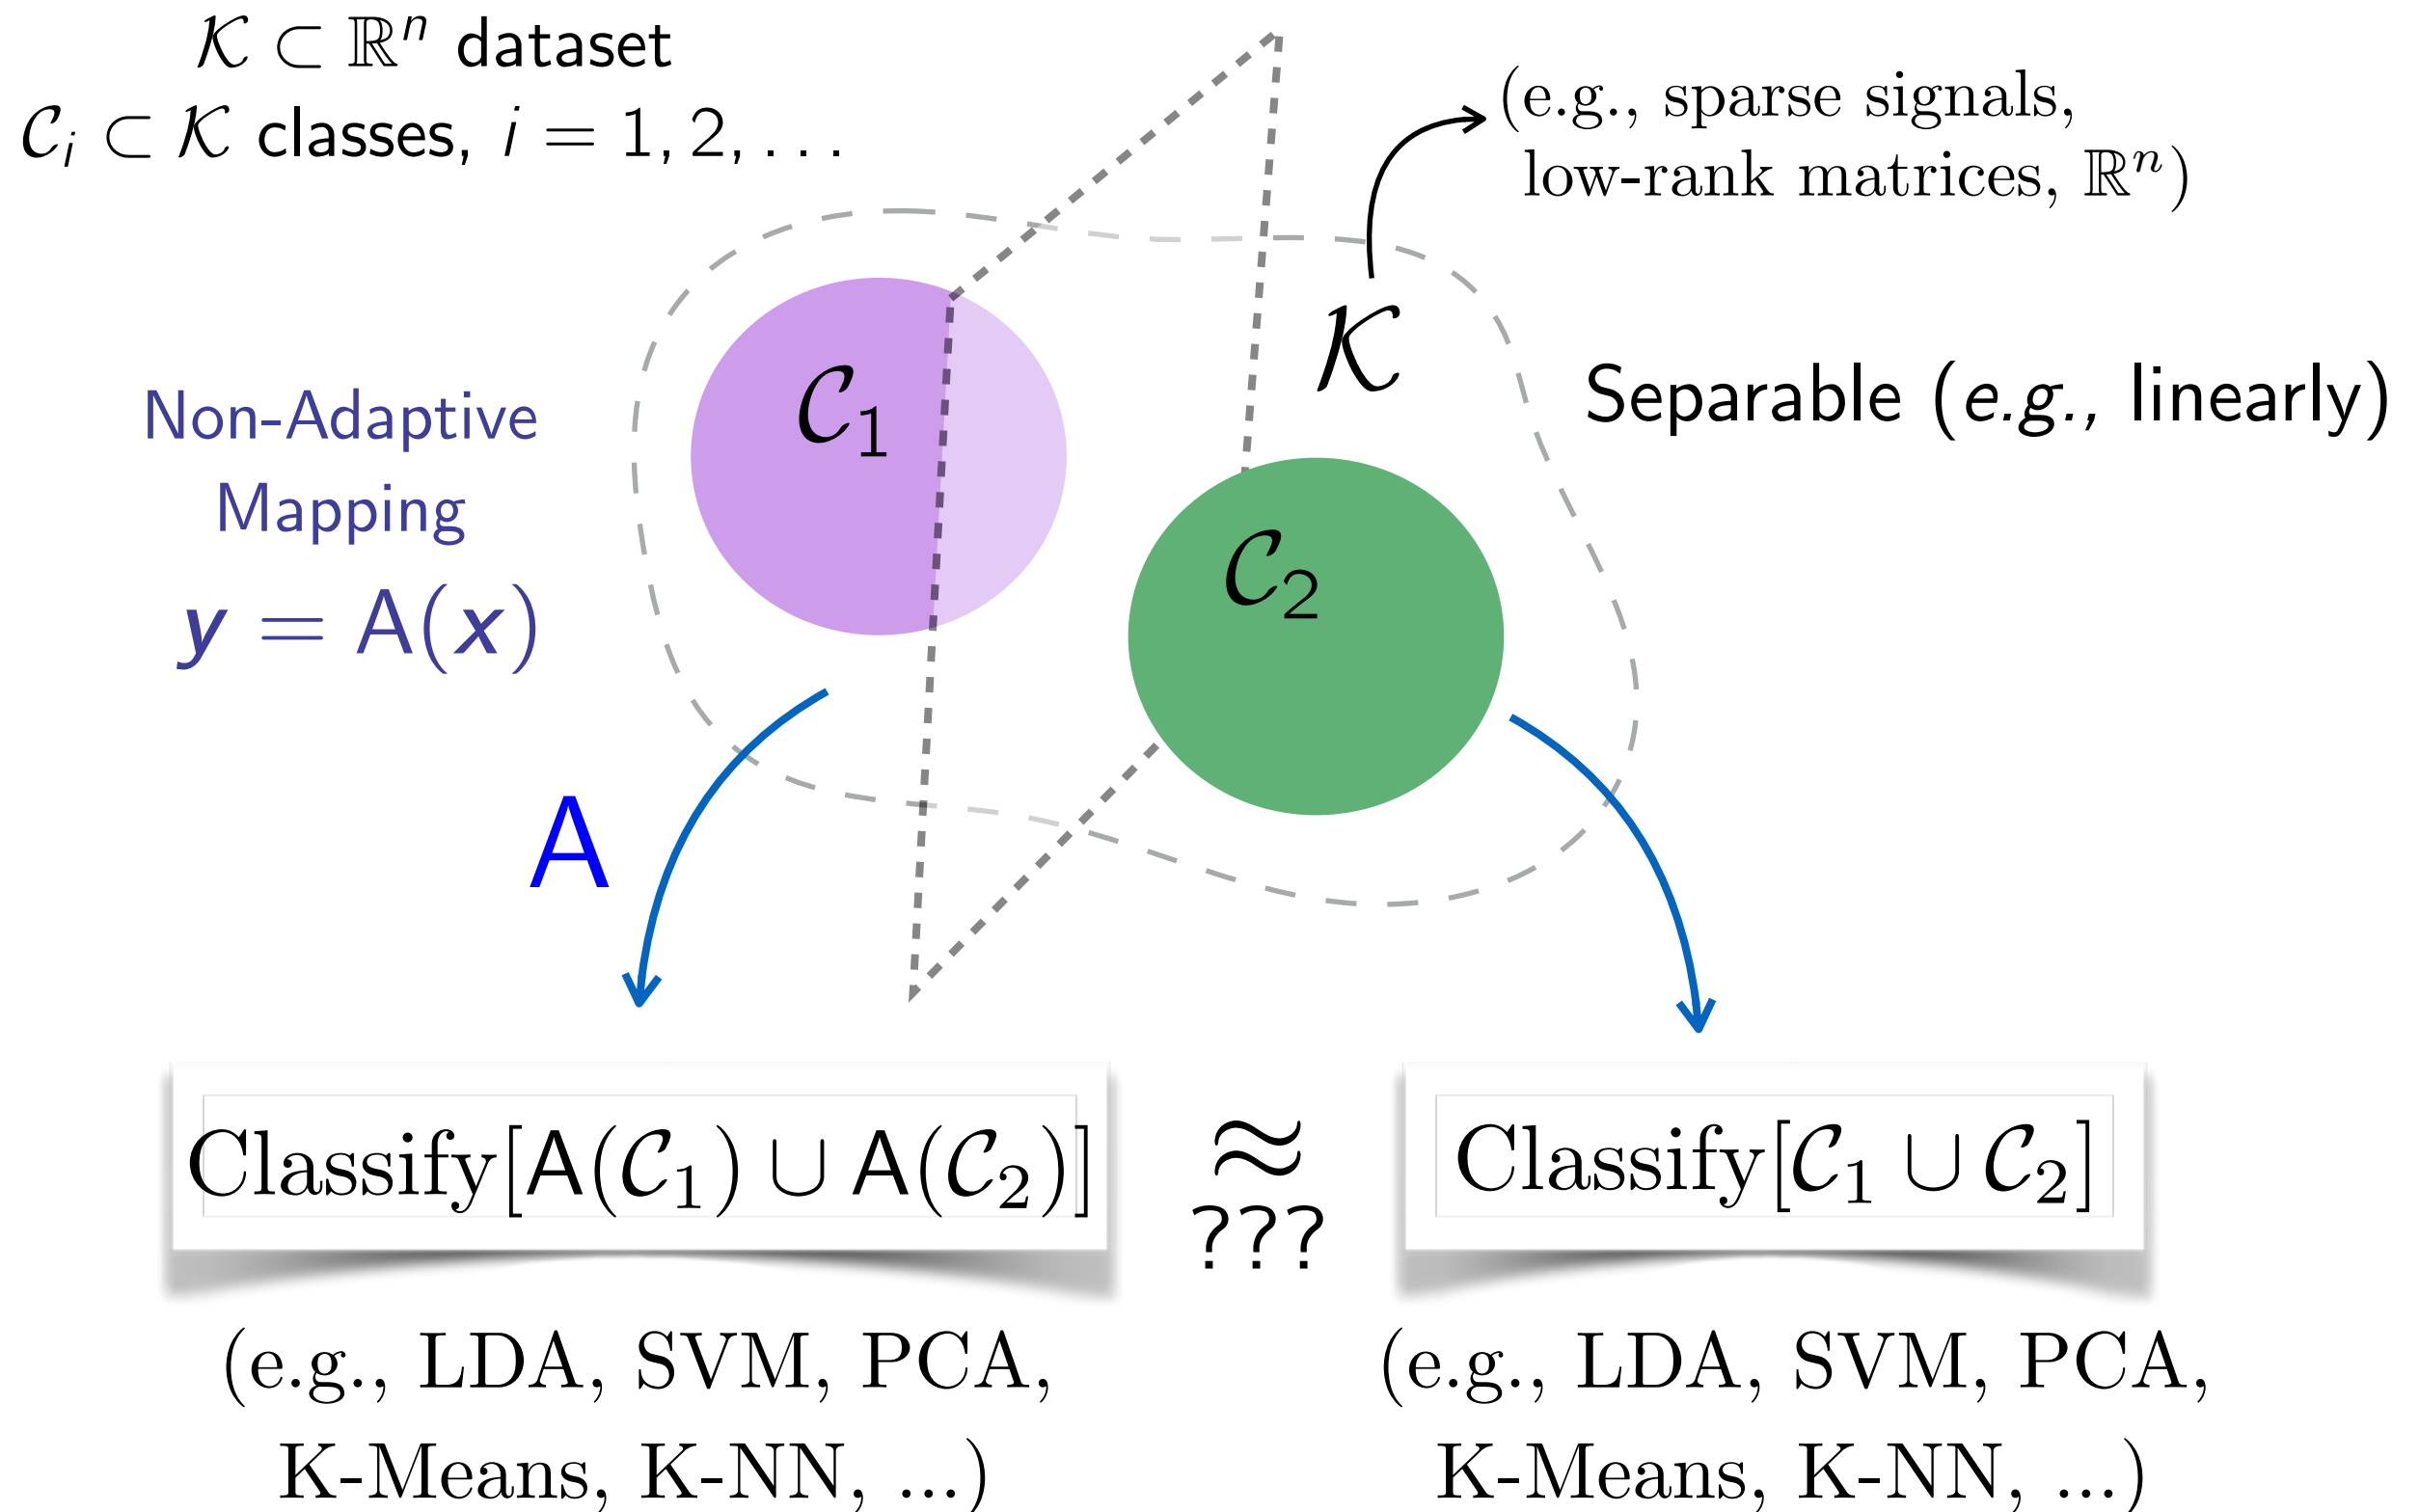
$\mathcal{K} \subset \mathbb{R}^n$ dataset
 $\mathcal{C}_i \subset \mathcal{K}$ classes, $i = 1, 2, \dots$



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

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




The Big Picture

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

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

Non-Adaptive
Mapping

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



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

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

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

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 $m > \left(w_{\ominus} + \sqrt{2 \log \frac{1}{\eta}}\right)^2 + 1$ then $p_0 \geq 1 - \eta$.

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

Example:




The Rare Eclipse Problem “on Tiles”



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

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




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

Idea: use the QRIP, i.e.,

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

w.h.p.

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

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

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 ($\delta \mathbb{Z}^m$).

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

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. Candès 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.
- ▶ A. S. Bandeira, D. G. Mixon, B. Recht, **“Compressive classification and the rare eclipse problem,”** In Compressed Sensing and its Applications 2017 (pp. 197-220). Birkhäuser, Cham.
- ▶ 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.
- ▶ 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.