

A Mathematical Introduction to Compressed Sensing

Stefania Petra
Mathematical Imaging Group
Heidelberg University

Lecture WT 2019/20

Contact

Stefania Petra, MIG - lecture

Room A4 209

Email petra@math.uni-heidelberg.de

Jan Plier, MIG - tutorial

Room A4 212

Email jan.plier@iwr.uni-heidelberg.de

Organizational Matters

Time/Location

Lecture SR 6 - Mathematikon INF 205

Wed. 9:00 - 11:00

Tutorial SR 2 - Mathematikon INF 205

Thu. 16:00 - 18:00

Evaluation Oral exam ← 50% programming exercises

Creditpoints 6CP + 2CP optional programming project

Media forms blackboard / lecture notes / slides

Previous knowledge Linear algebra, analysis I + II
(+ basic tools from probability theory,
convex analysis & optimization)

Web

Check your **email** and

`http://ipa.iwr.uni-heidelberg.de/dokuwiki/doku.php?id=teaching`

for announcements and updates of

- exercise sheets
- handouts
- lecture notes

Please also sign up using **MÜSLI**

`https://muesli.mathi.uni-heidelberg.de`

Literature

S. Foucart, H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Birkhäuser, 2013

S. Boyd, L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004

M. Ledoux, *The Concentration of Measure Phenomenon* American Mathematical Society, 2005


R. Schneider, W. Weil, *Stochastic and Integral Geometry*, Springer, 2008

J.-L. Starck, F. Mutagh, J.M. Fadili, *Sparse Image and Signal Processing*, Cambridge University Press, 2010

Literature

main reference

S. Foucart, H. Rauhut
Birkhäuser, 2013



covers: 2000 - 2012

Content

Theory

sparse reconstruction via l_0/l_1 -minimization;
basic properties: coherence, nullspace property, restricted isometry property; random sensors; phase transitions;
basic tools from convex analysis, probabilities and integral geometry


Algorithms

orthogonal matching pursuit; thresholding based methods;
primal-dual methods

Applications


sparse approximation; image processing (tomographic inversion, deblurring, etc.); low-rank completion

Overview



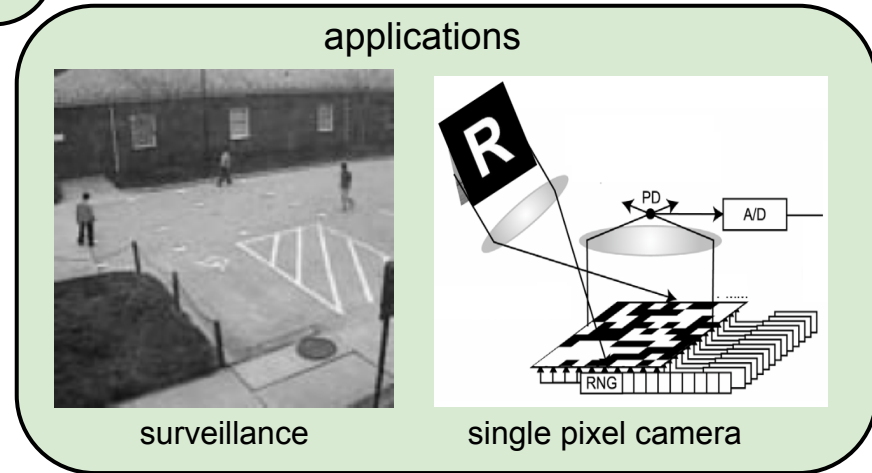


enable

RIP fails

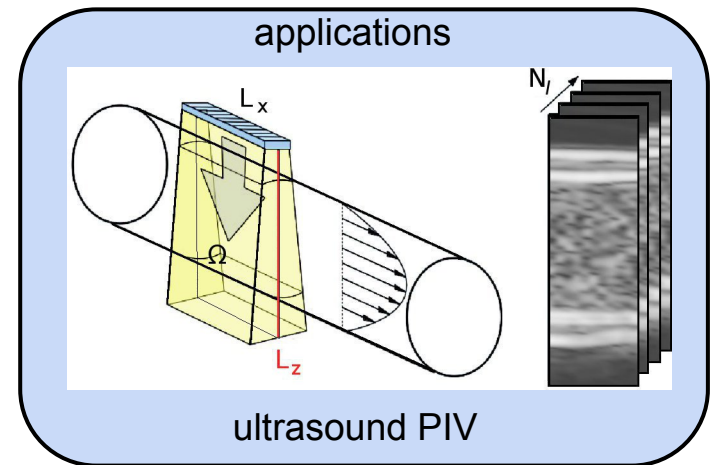


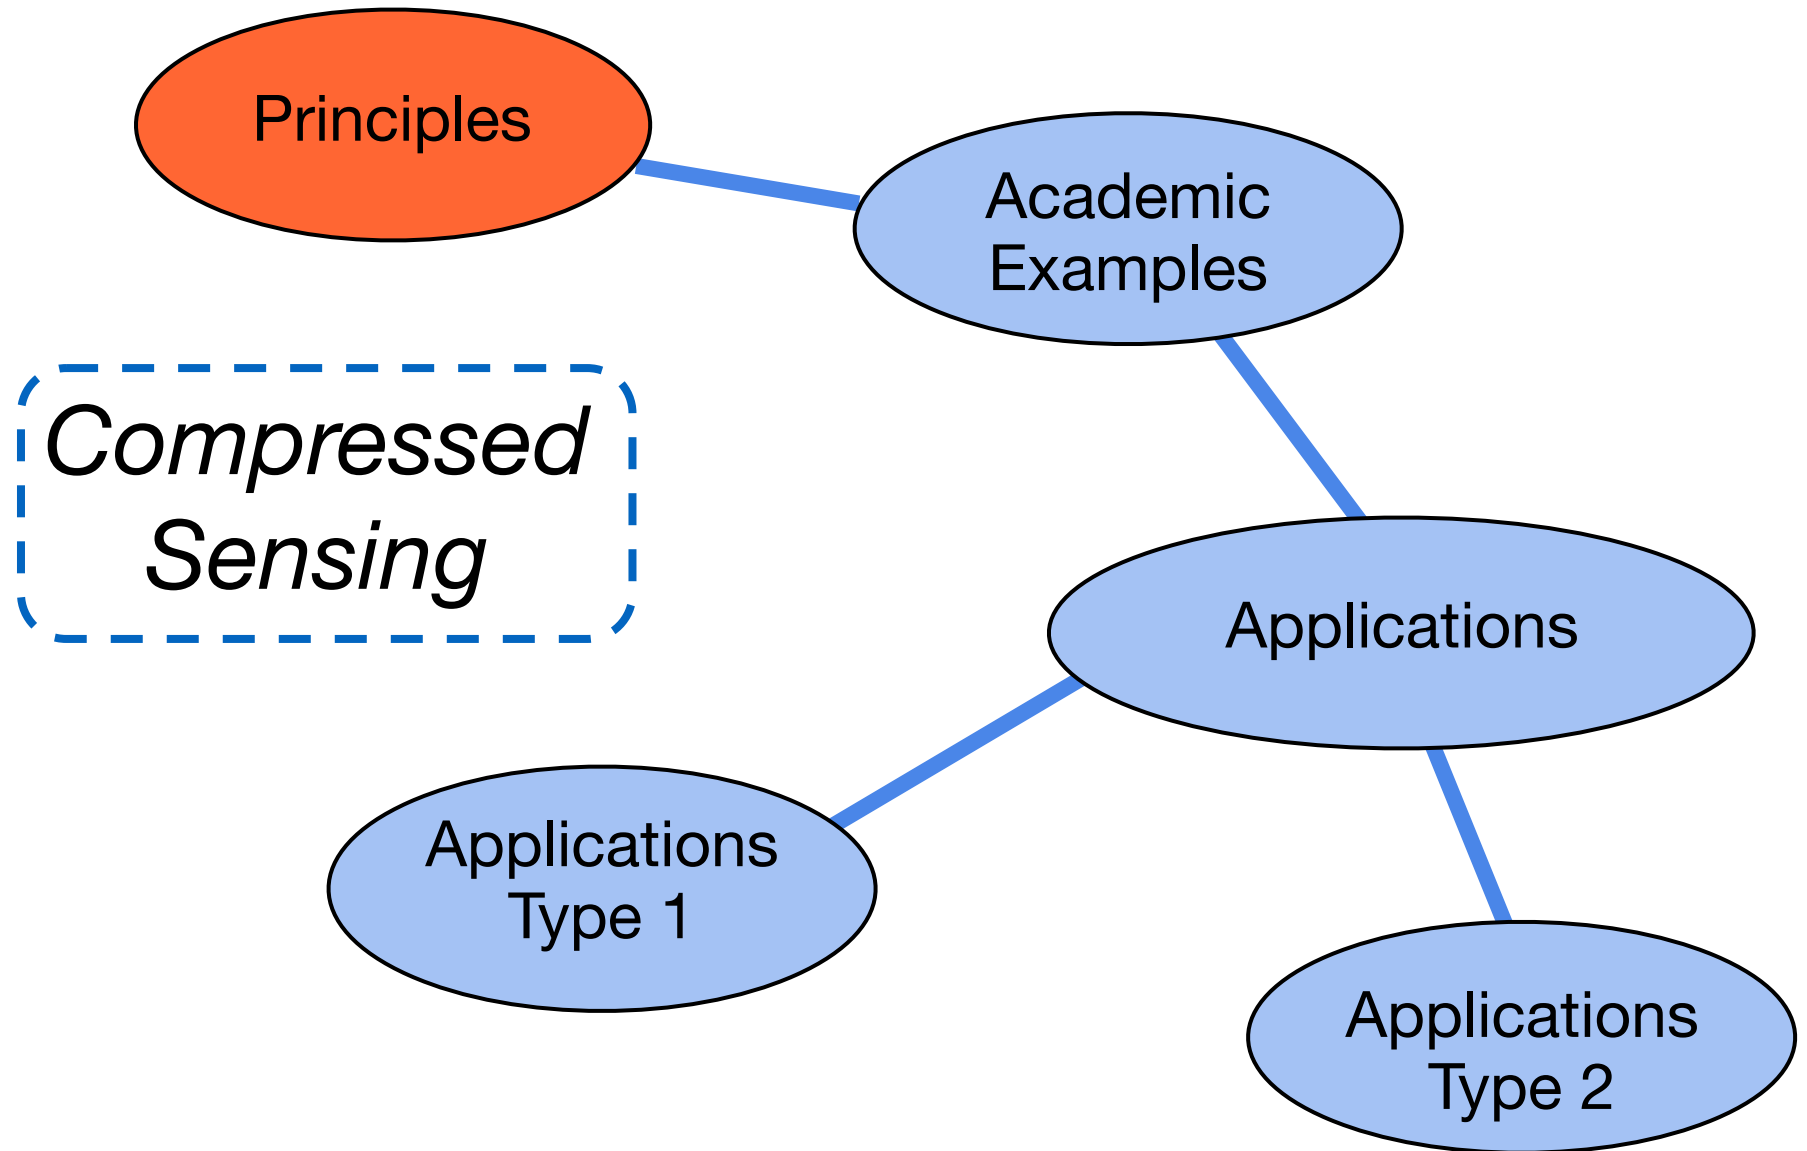
applies

applies



RIP fails





Classical Sampling vs Compressed Sensing



Claude
Shannon



Emmanuel
Candés



Harry
Nyquist



Terence
Tao




David
Donoho

Classical Sampling




Shannon-Nyquist theorem



finite support in
the Fourier domain


correct sampling rate
correct (continuous) recovery

incorrect sampling rate
severe artefacts



... but “big data”

Compressed Sensing



encoding **below** the Nyquist rate: $m \ll n$

$$b = Ax \quad x \in \mathbb{R}^n, b \in \mathbb{R}^m$$



observation

measurement matrix ("sensor")

Compressed Sensing



encoding **below** the Nyquist rate: $m \ll n$

$$b = Ax \quad x \in \mathbb{R}^n, b \in \mathbb{R}^m$$



observation

measurement matrix (“sensor”)

nonlinear decoding by **convex** programming, e.g.

$$\Delta(b) = \operatorname{argmin}_{x: Ax=b} \|x\|_1$$

performance of en-/decoding pair: $\|x - \Delta(Ax)\|$

Compressed Sensing: Basic Requirements

When does it (provably) work?!

small support in *some transformed domain*

(1) Signals are **sparse** ...

$$\mathcal{X}_k := \{x \in \mathbb{R}^n : \|x\|_0 = |\text{supp}(x)| \leq k\}$$

... or have a **sparse representation** $x = Dz$
(basis, dictionary; henceforth for simplicity: $D = I$)

Compressed Sensing: Basic Requirements

When does it (provably) work?!

small support in *some*
transformed domain

(1) Signals are **sparse** ...

$$\mathcal{X}_k := \{x \in \mathbb{R}^n : \|x\|_0 = |\text{supp}(x)| \leq k\}$$

... or have a **sparse representation** $x = Dz$
(basis, dictionary; henceforth for simplicity: $D = I$)

(2) Sensor matrix A is an isometry on \mathcal{X}_k

$$\exists \delta_k \in (0, 1), \quad \forall x \in \mathcal{X}_k$$

$$(1 - \delta_k) \|x\|_{\mathbb{R}^n}^2 \leq \|Ax\|_{\mathbb{R}^m}^2 \leq (1 + \delta_k) \|x\|_{\mathbb{R}^n}^2$$

“**restricted isometry property (RIP)**”

Some Basic Questions

- Can we trust our model to return an intended sparse signal?
- Does the model have a unique solution? (otherwise, different algorithms may return different answers)
- Is the solution exactly equal to the original sparse signal?
- If not (due to noise), is the solution a faithful approximation of it?
- How much effort is needed to numerically solve the model?

Compressed Sensing: Stable Recovery

Candés, Romberg, Tao 2006

$$\delta_{2k} < \sqrt{2} - 1 \quad (\text{RIP})$$

$$b = Ax^* + \xi, \quad \|\xi\|_2 \leq \varepsilon \quad (\text{bounded noise})$$

$$\Delta(b) = \operatorname{argmin}_{x: \|Ax-b\|_2 \leq \varepsilon} \|x\|_1$$

Stable recovery guarantee

decoding

$$\|x^* - \Delta(b)\|_2 \leq C \left(\frac{\|x^* - x_k^*\|_1}{\sqrt{k}} + \varepsilon \right)$$



$\varepsilon \rightarrow 0$ and k -sparsity \Rightarrow perfect reconstruction

How to Read Recovery Guarantees


Some basic aspects that distinguish different types of guarantees:

- Recoverability (exact) vs stability (inexact)
- General A or special A ?
- Universal (all sparse vectors) or instance (certain sparse vector(s))? Uniform/non uniform recovery guarantees?
- General optimality? or specific to model/algorithm?
- Required property of A : spark, NSP, coherence, RIP, dual certificate?
- If randomness is involved, what is its role?

Sparsity in Coordinate Basis



Sparsity in Orthonormal Basis




$$H_{j,k}^e(x) = 2^j H^e(2^j x - k), \quad e \in V = \{\{0, 1\}, \{1, 0\}, \{1, 1\}\} \text{ and} \\ j \geq 0, \quad k \in \mathbb{Z}^2 \cap 2^j [0, 1)^2$$

Sparsity vs Compressibility





1 megapixel image


wavelet coeffs



(sorted)



zoom in



(log₁₀ sorted)


Sparsity vs Compressibility



1 megapixel image



25k term approx



B -term approx error

- Sparse structure in Wavelet domain: few large coeffs, many small coeffs
- Basis for JPEG2000 image compression standard
- Do not confuse image compression with CS

Good Sensors?

- preserve structure and information in sparse/compressible signals
- models with high probability
- the number of samples should be minimal
- each measurement carries the same amount of information

$$\begin{aligned} b_1 &= \left\langle \begin{array}{c} \text{Image of a woman} \\ \text{Random noise} \end{array} \right\rangle \\ b_2 &= \left\langle \begin{array}{c} \text{Image of a woman} \\ \text{Random noise} \end{array} \right\rangle \\ b_3 &= \left\langle \begin{array}{c} \text{Image of a woman} \\ \text{Random noise} \end{array} \right\rangle \\ &\vdots \\ b^m &= \left\langle \begin{array}{c} \text{Image of a woman} \\ \text{Random noise} \end{array} \right\rangle \end{aligned}$$

Surprise

- measurements do not match image structure at all
- measurements look like random noise
- same measurements can be used for **any** compressible signal class (universal)

Random Sensors: Sufficient RIP Conditions

Mathematical ways to design good sensor matrices A enjoying RIP include


- Gaussian i.i.d. entries $A_{ij} \sim \mathcal{N}(0, \frac{1}{m})$, Bernoulli, ...
- random partial Fourier (DFT) matrices

Random Sensors: Sufficient RIP Conditions



Mathematical ways to design good sensor matrices A enjoying RIP include

- Gaussian i.i.d. entries $A_{ij} \sim \mathcal{N}(0, \frac{1}{m})$, Bernoulli, ...
- random partial Fourier (DFT) matrices

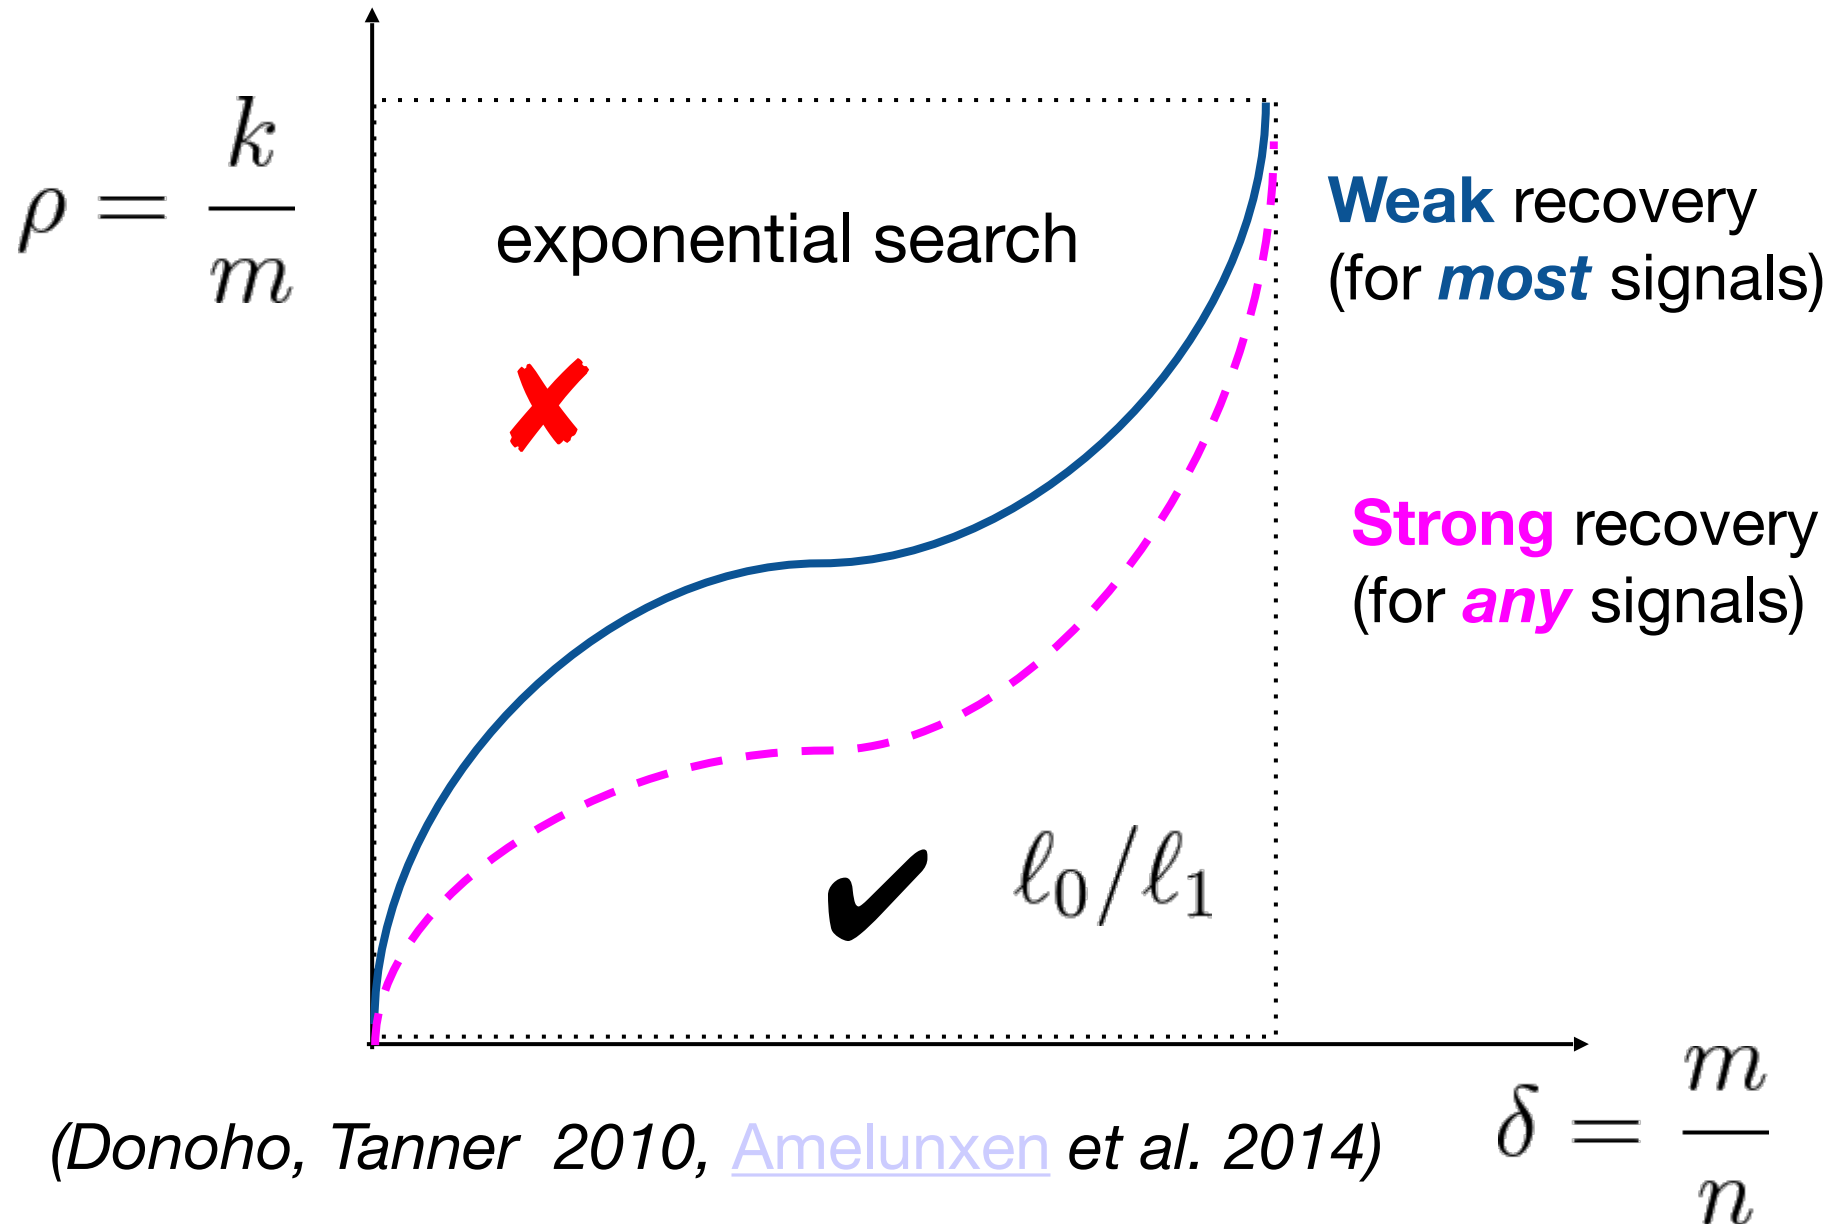
Then $\exists c_1 = c_1(\delta), c_2 = c_2(\delta)$ such that

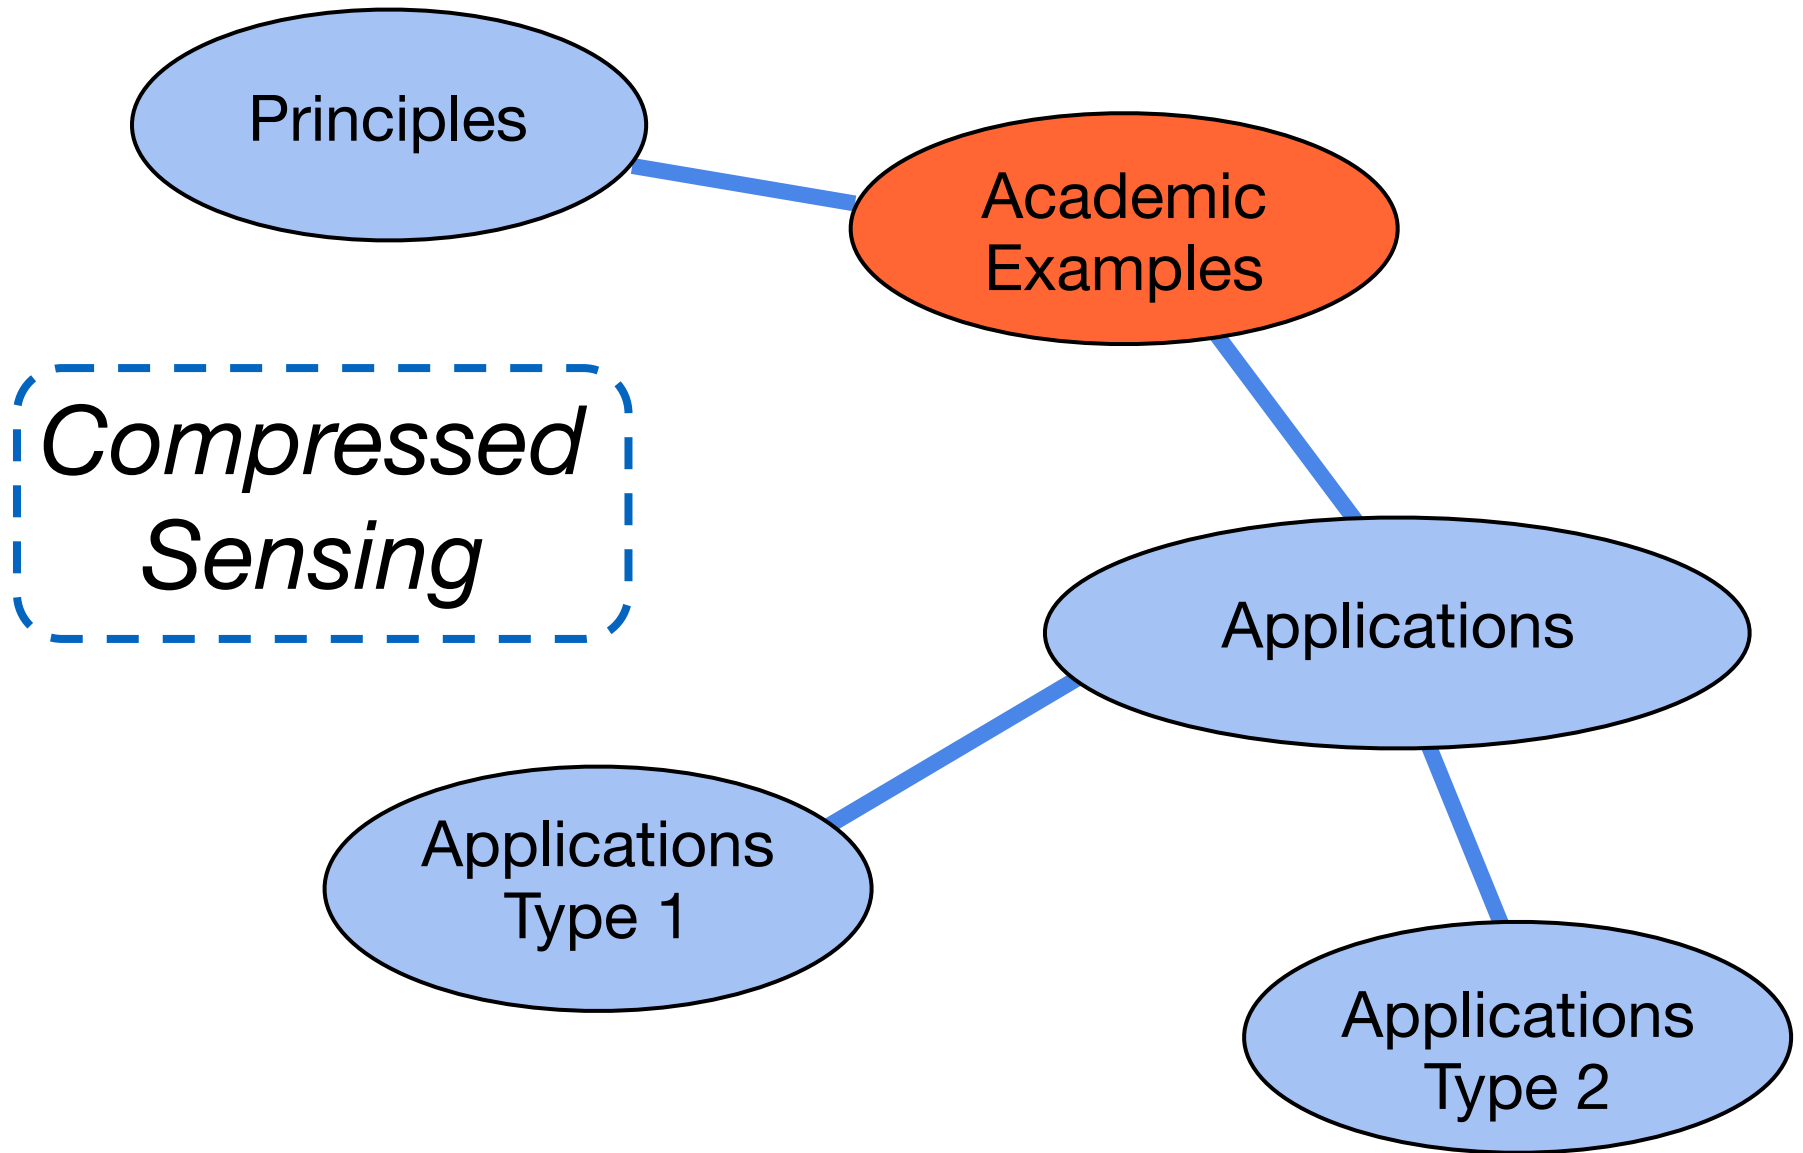
$$\Pr (A \text{ has RIP}_k) \geq 1 - 2e^{-c_1 m}$$


provided the number of measurements satisfies

$$m \geq c_2 k \log(n/k) \quad (m \propto \text{sparsity !})$$


Known Phase Transitions: Gaussian A





Academic Example

$$\begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$



x_j deviation of the j th coin
from the correct mass

b_i i th weighting

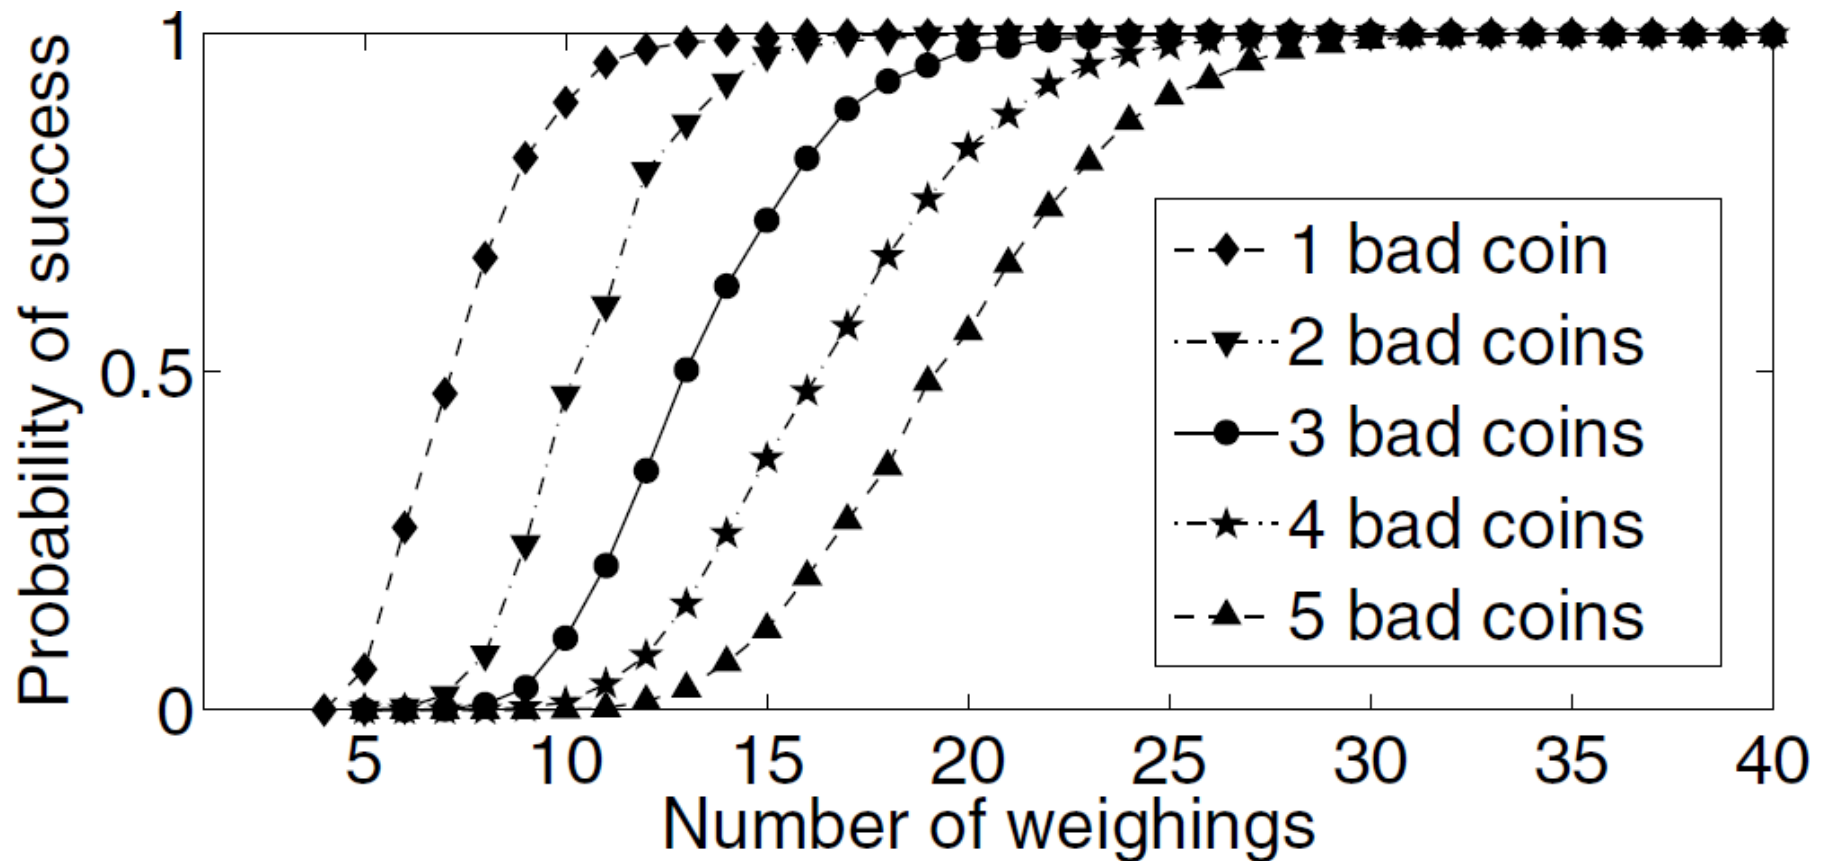
One coin is false. Which one?




Weigh Random Subsets

Example: $n = 100$ coins, m weighings


A *Bernoulli* $m \times n$ underdetermined !




Academic Example



$x \in \mathbb{R}^{100}$, $\|x\|_0 = 10$, $b = Ax$, $A \in \mathbb{R}^{30 \times 100}$ Gaussian



ℓ_0 quasi-norm





ℓ_1 norm




ℓ_2 norm


Academic Example



l_0 quasi-norm




l_1 norm

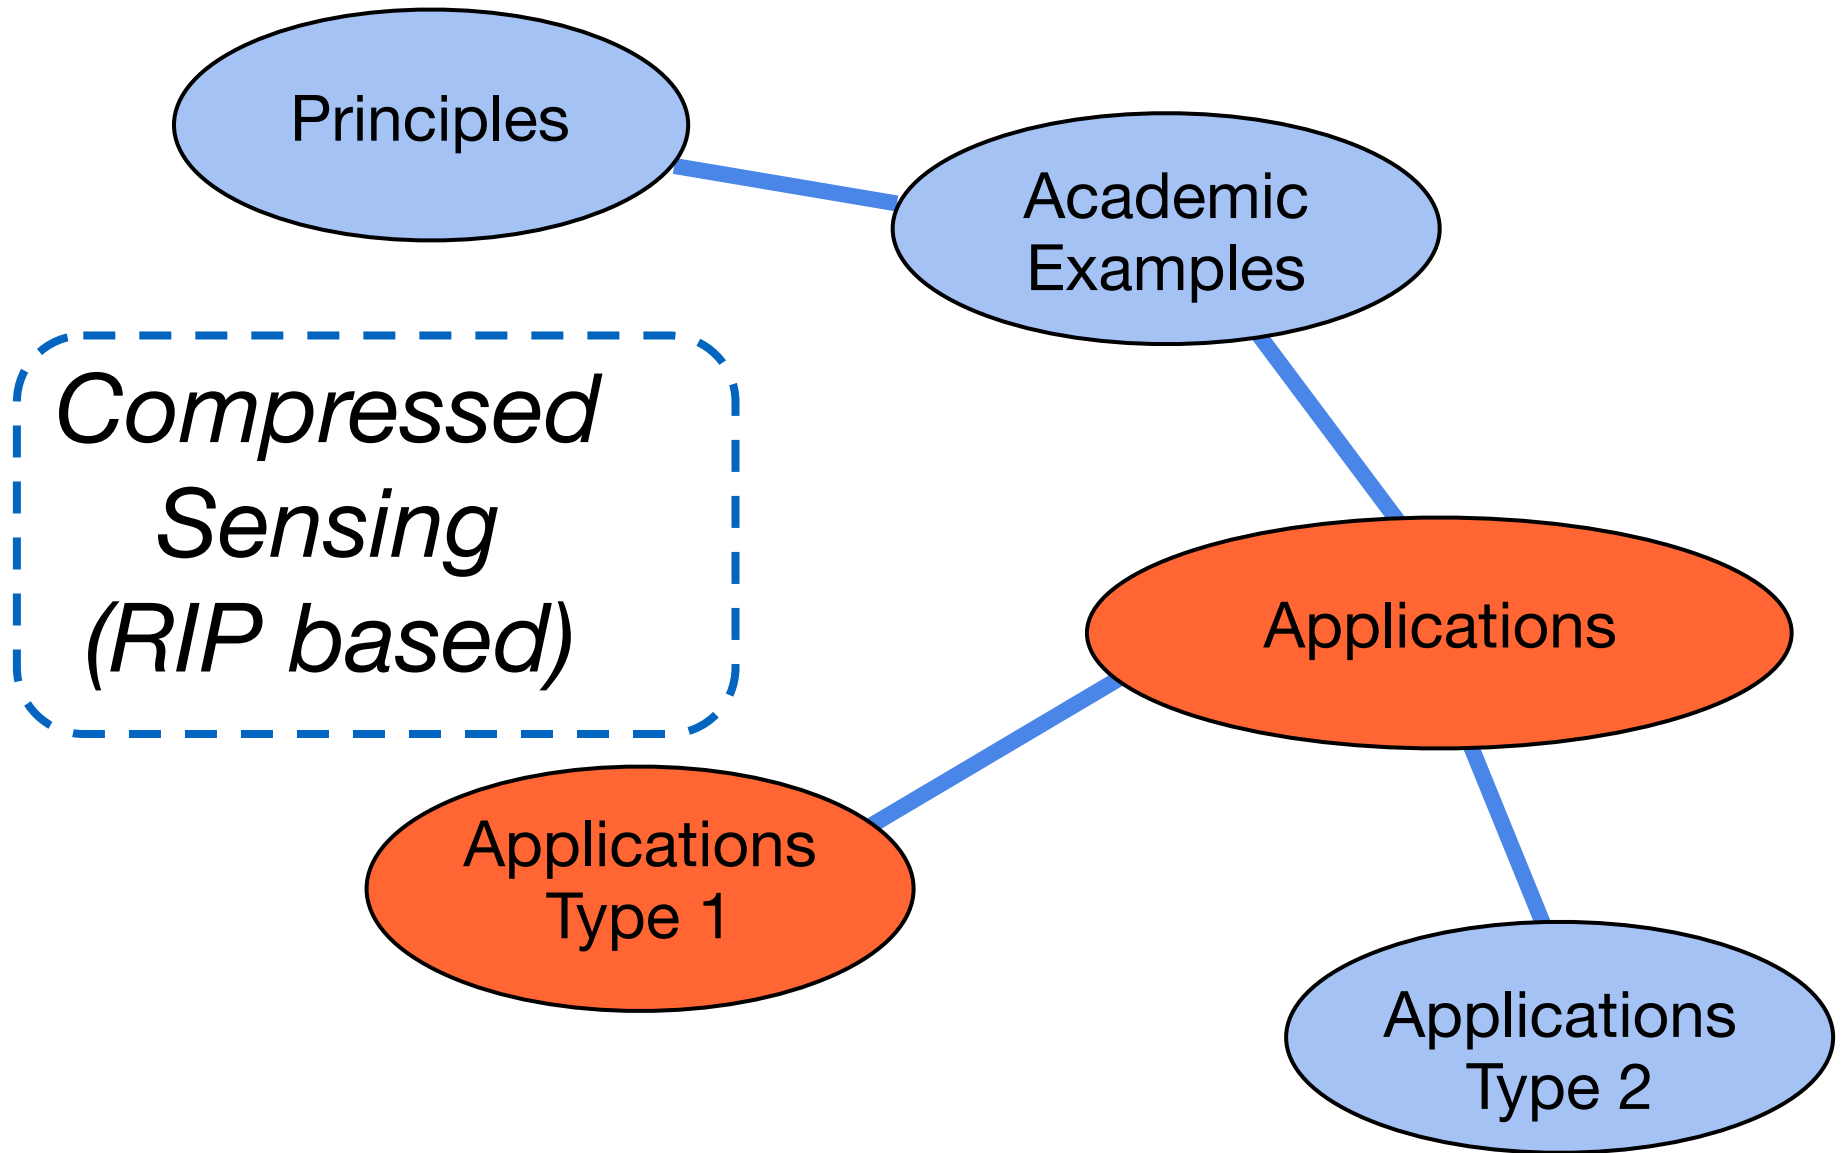


l_2 norm

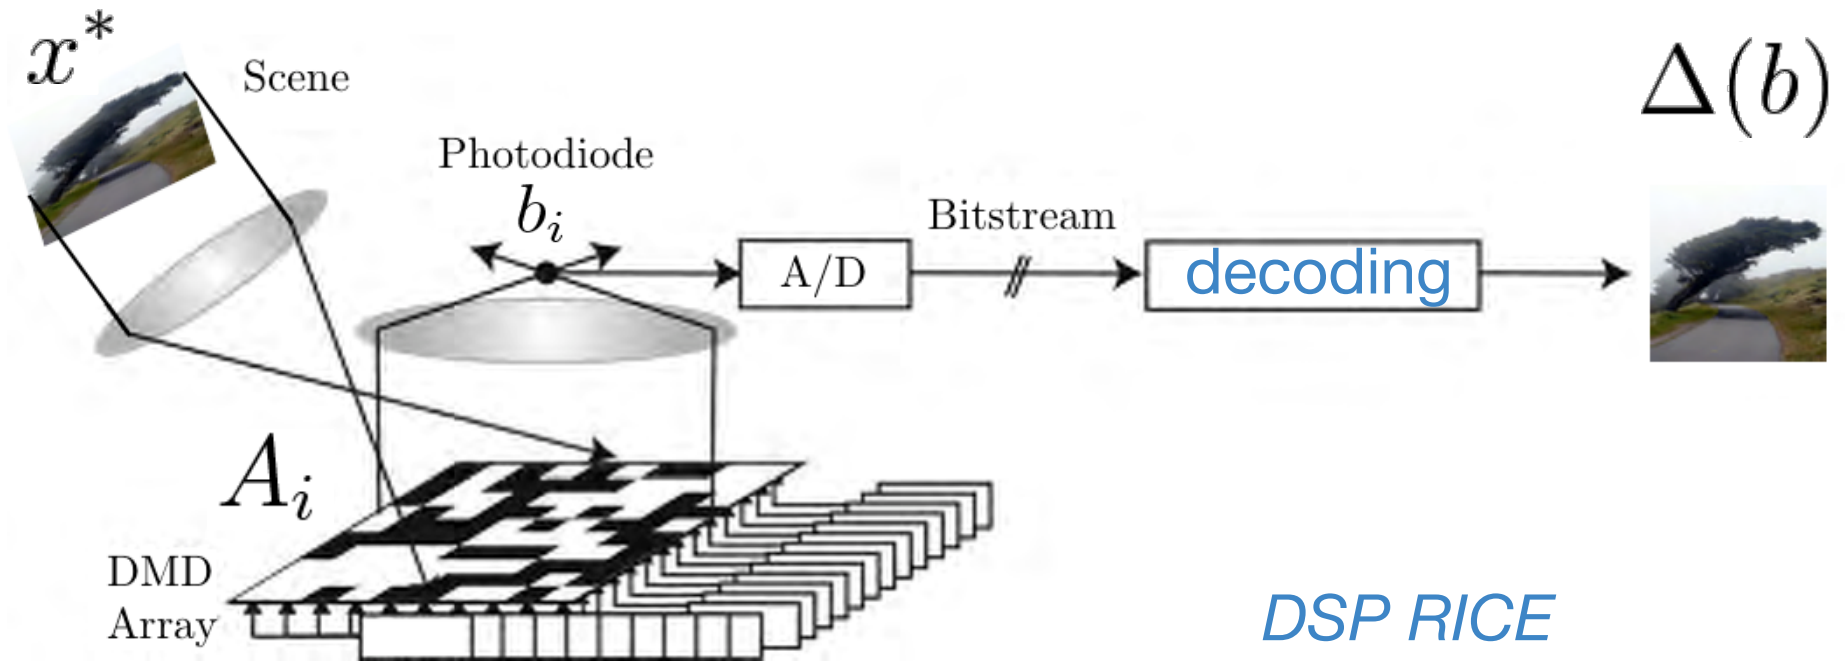
Academic Example




- ℓ_1 can be used as a proxy for ℓ_0
- This is a convex program and can be solved in polynomial time



Single Pixel Camera



$$m = O(n^{1/4} \log^5(n)) \ll n$$



Single Pixel Camera

target
65536 pixels



11000 measurements
(16%)



1300 measurements
(2%)



Rapid Sparse Magnetic Resonance Imaging

Lustig, Donoho, Pauly 2007

(Google Scholar: 3105 citations

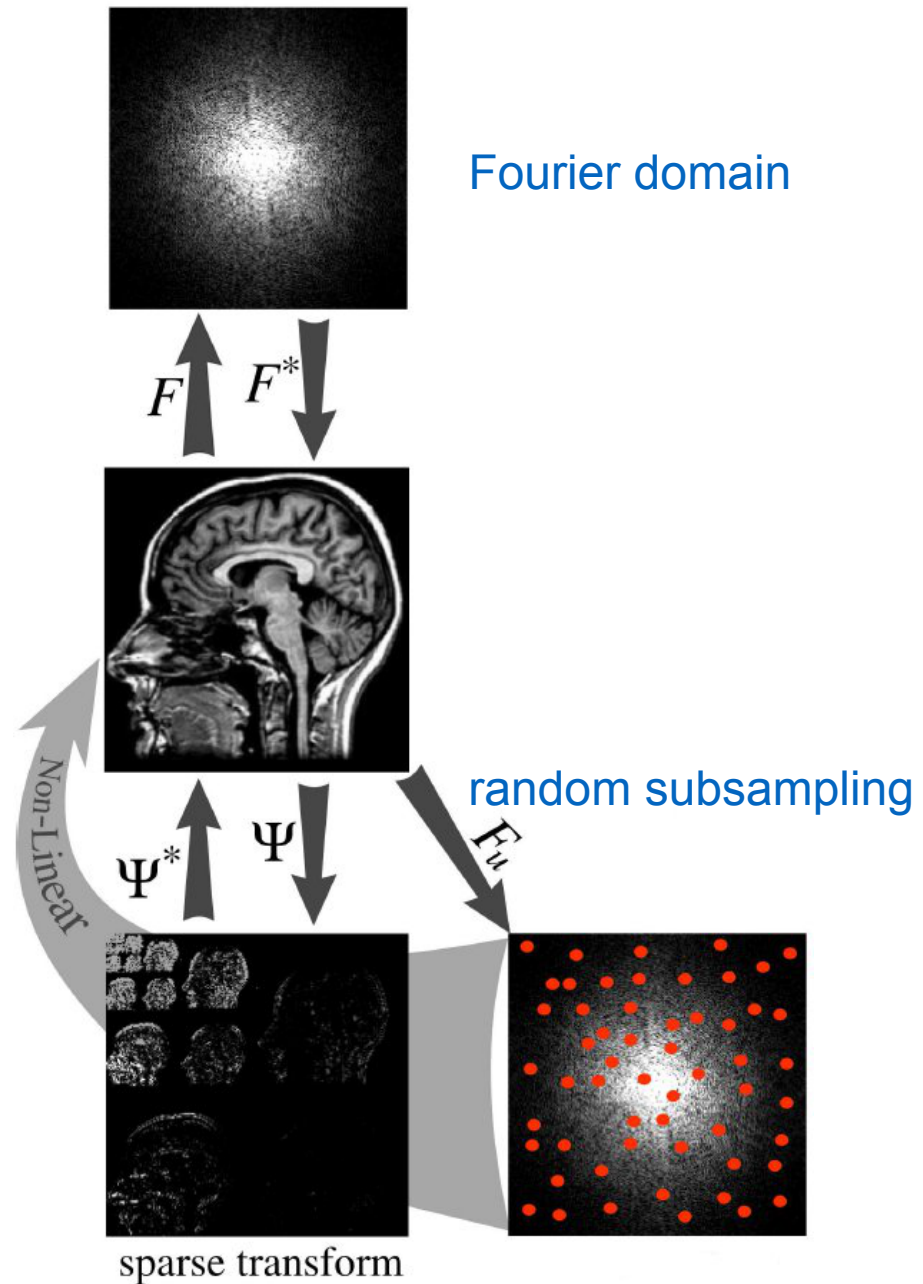
October 2016)

compression &
speed-up factor: 8

nonlinear (convex)
decoding & recovery


signal

Fourier domain



Compressed Video Sensing

© Baraniuk et al.,
NIPS'11



illum. changes

$$F = L + S$$

motion

video = low-rank + sparse signal

Compressed Video Sensing

$$F = L + S$$

video = low-rank + sparse signal

Recovery by *convex* programming

$$\min_{L,S} \|L\|_* + \lambda \|\text{vec}(S)\|_1$$

subject to

$$\|f - \mathcal{A}(L + S)\|_2 \leq \varepsilon$$

low-rank

sparse

sampling

Scientific Imaging: 2007 - 2018

Medical Imaging

- Magnetic Resonance Imaging
- Computerized Tomography

Imaging in other Sciences

- Thermoacoustic Tomography
- Photoacoustic Tomography
- Electrical Impedance Tomography
- Electron Tomography
- Seismic Tomography
- Fluorescence Microscopy
- Radio Interferometry

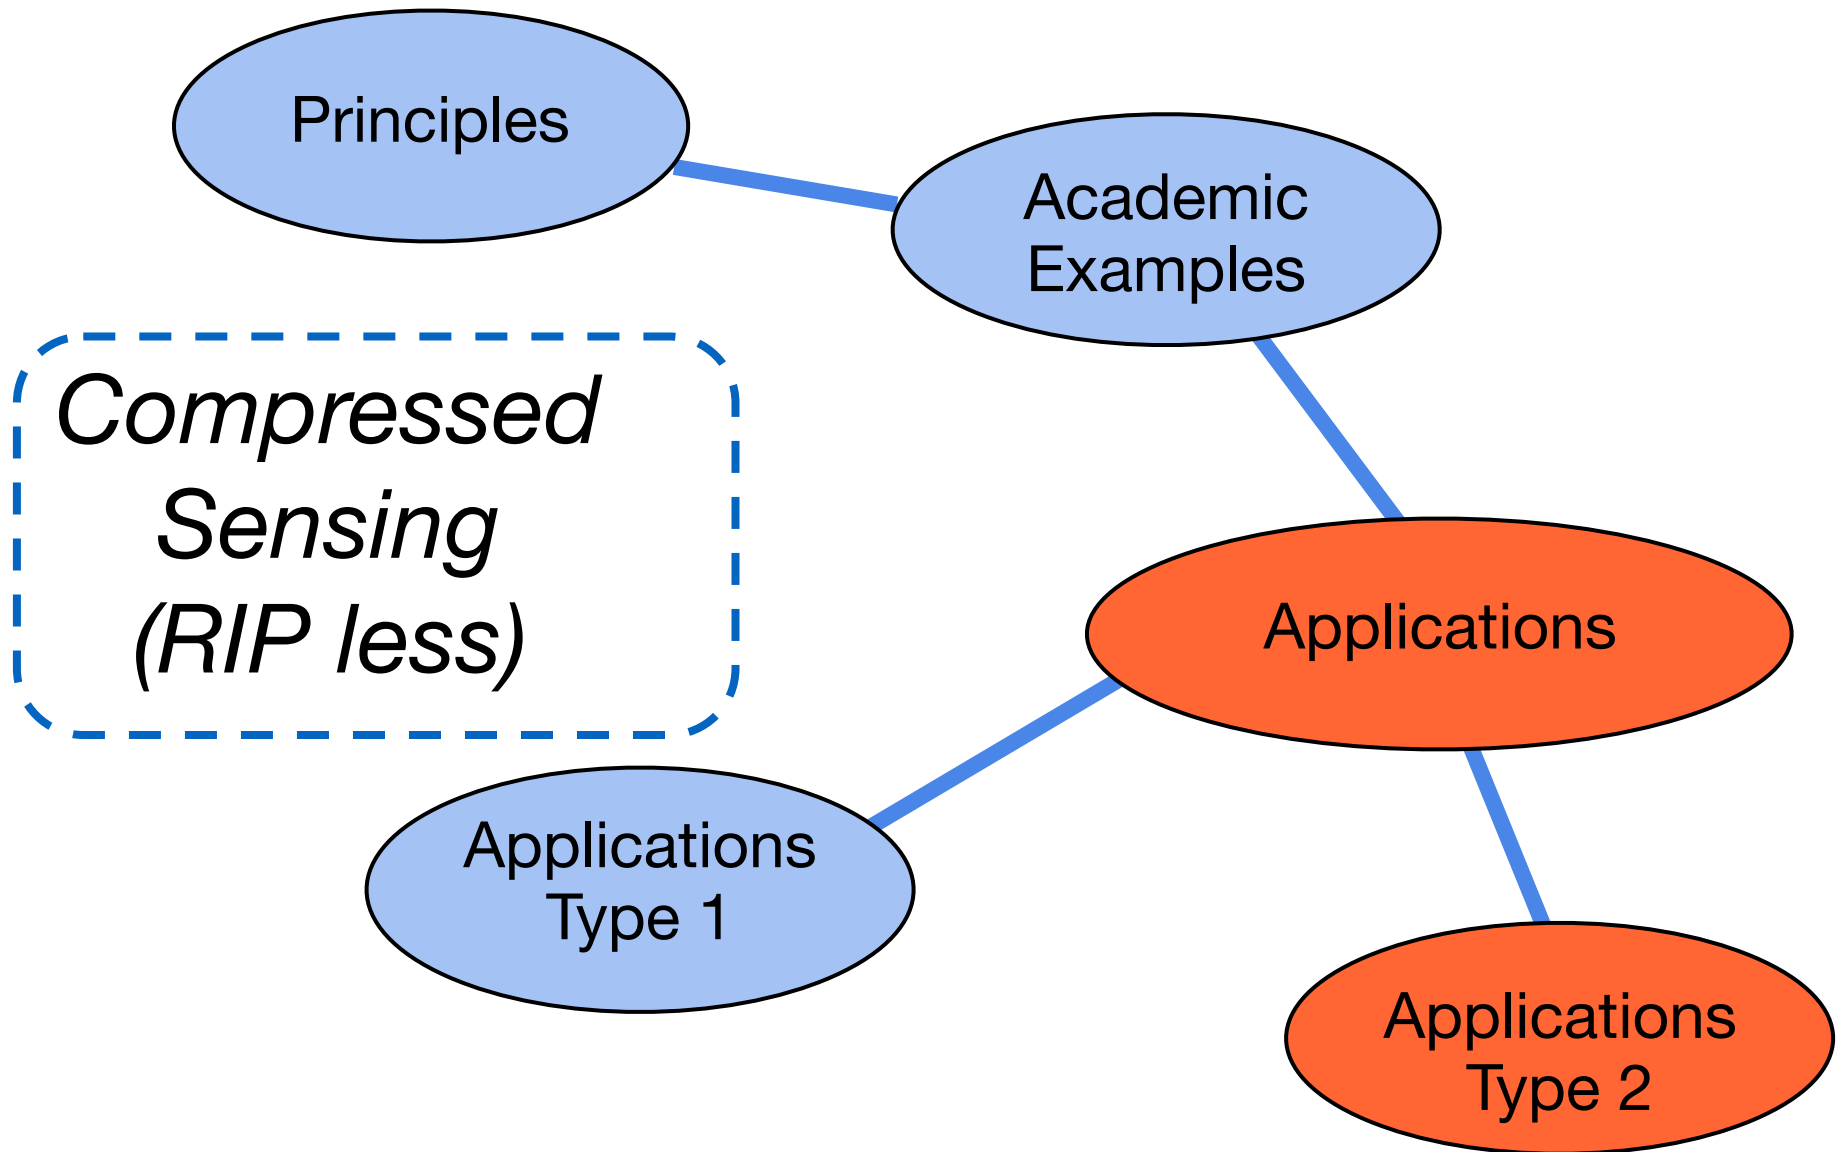
Established RIP-based
CS-theory cannot
explain much better
empirical performance

Significant gap between
mathematical theory and
applied fields

Major recent trend (mathematics)

⇒ dispense with RIP and universal sensing operators

⇒ exploit *structured sparsity*



Example (© A. Hansen, 2014)




Cambridge Advanced Imaging Center (CAIC)
Fluorescence Microscopy: zebra fish cells

raster scan

recovery from 6.25 %


established scheme

2048 × 2048



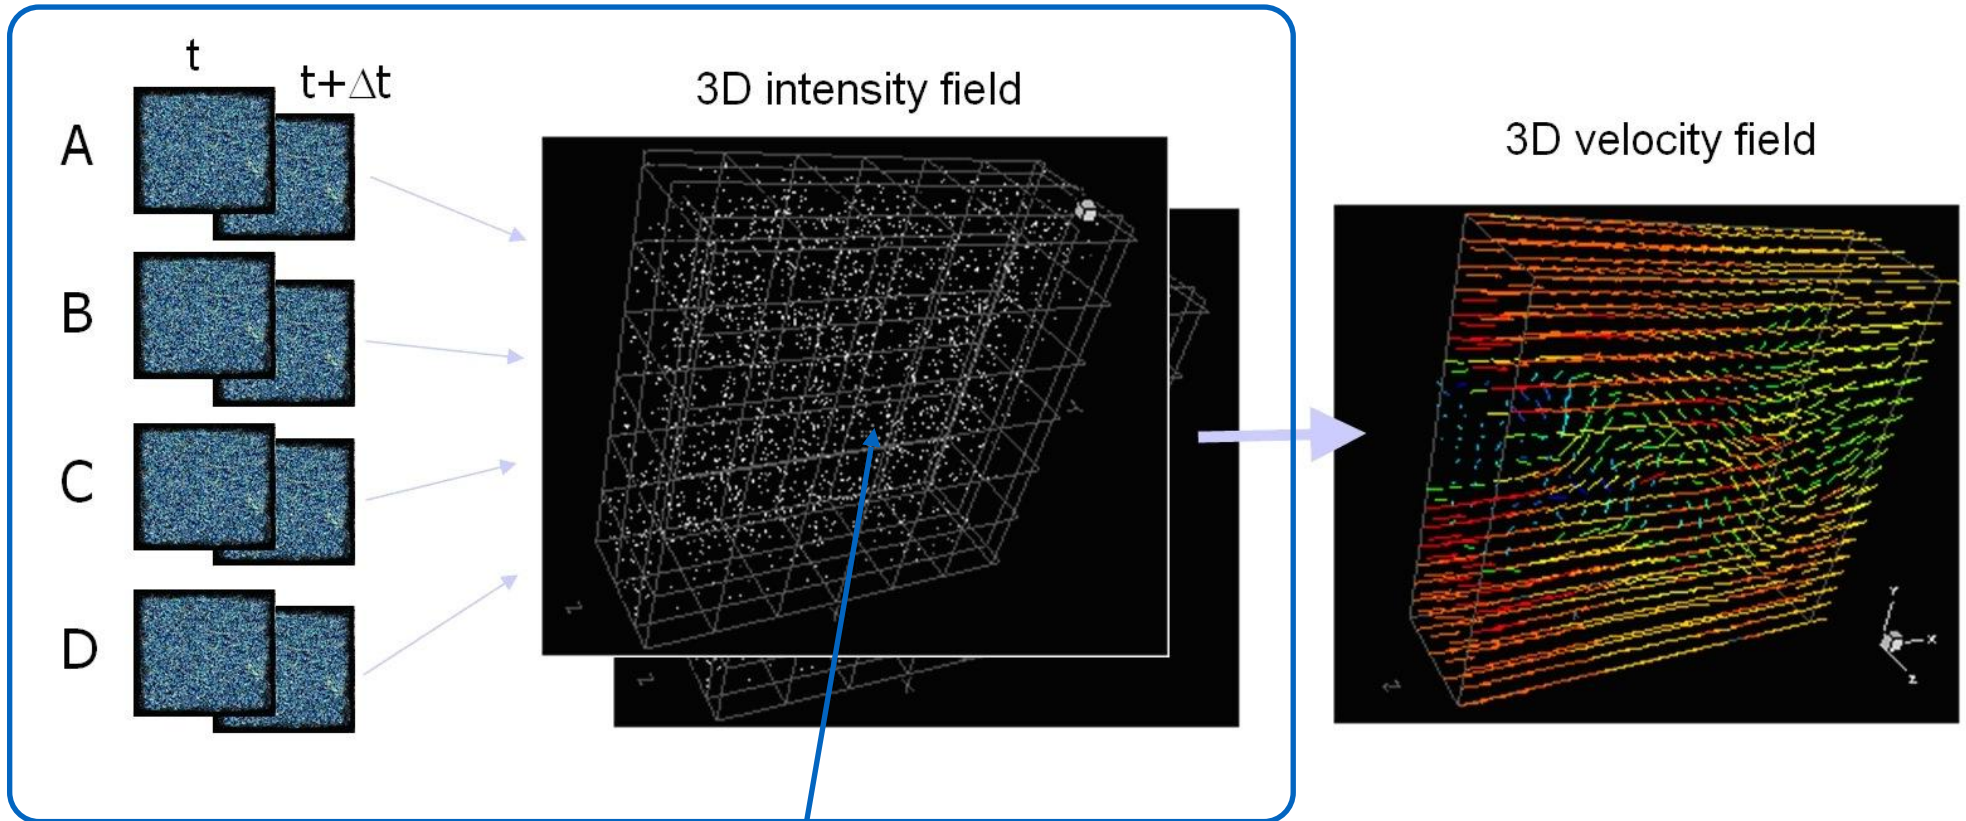
CS and Non Standard Tomography

Experiments in Fluids
DFG SPP 1147



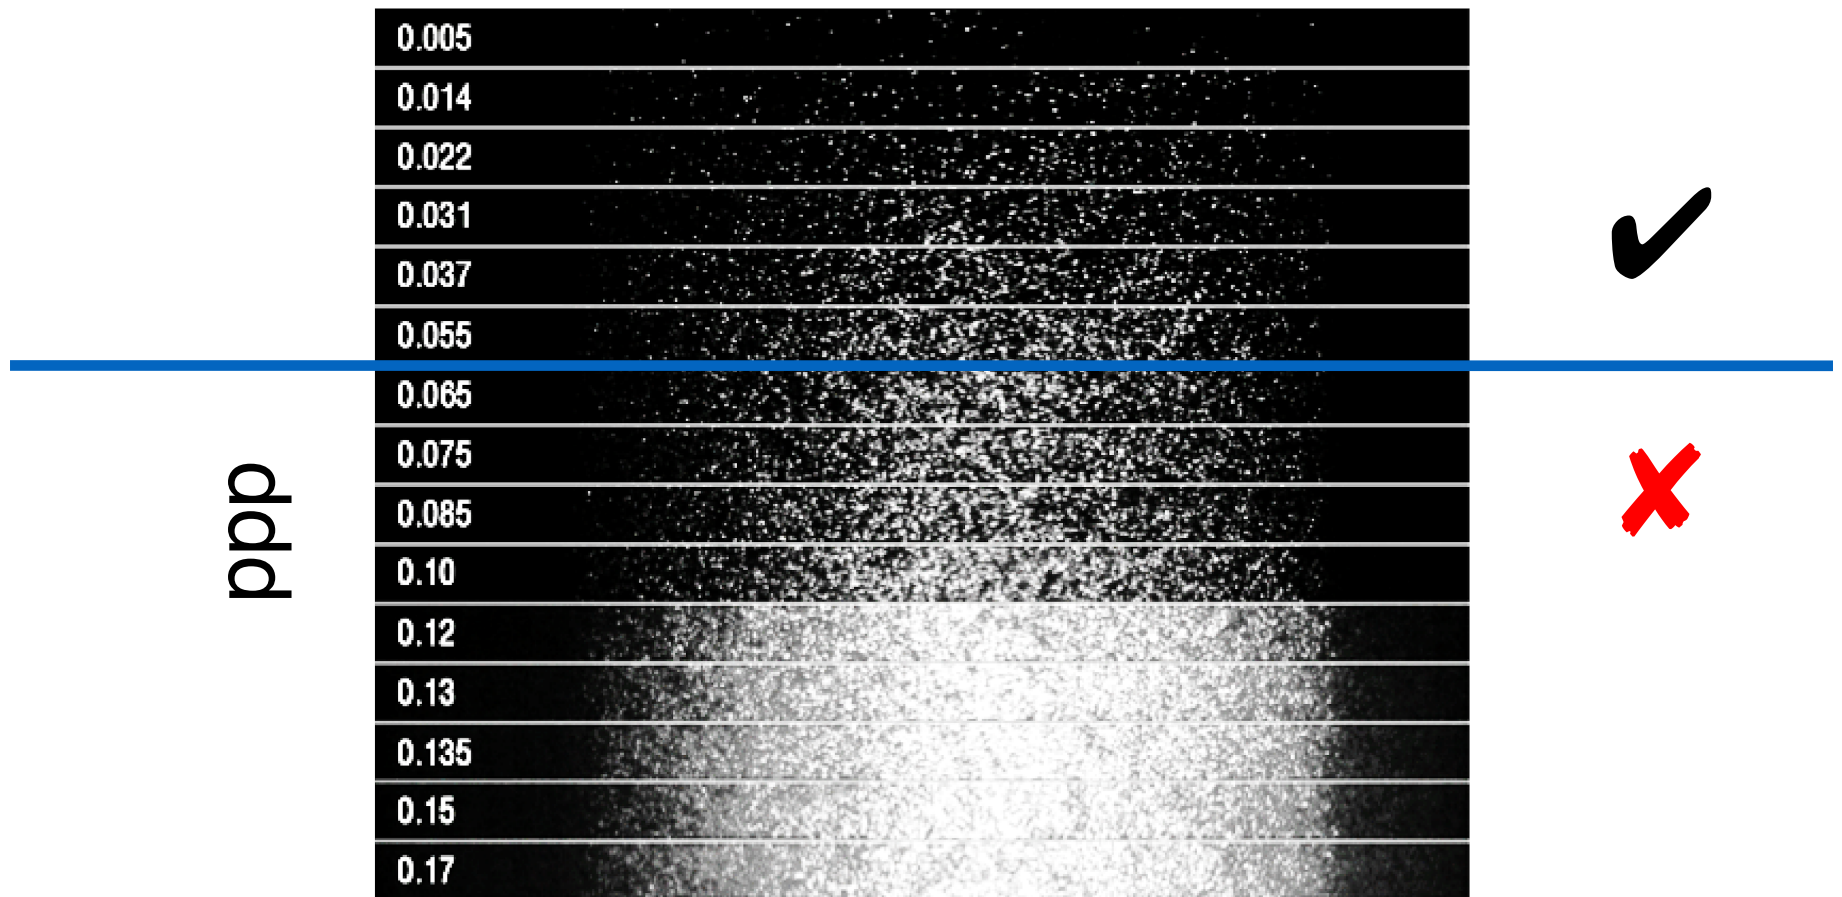
Tomo PIV

Sparsity in Tomo PIV



key parameter: seeding density


Motivation



key parameter: seeding density

Imaging Set-Up

Few measurements



d problem size, resolution

$$n = d^3$$

$$A \in \mathbb{R}_+^{m \times n}$$

$$\boxed{m} = p d^2 < \boxed{n}$$


p projections

k particles

$$Ax = b, \quad x \in \{0, 1\}^n, \quad x \in \mathbb{R}_+^n$$

Poor CS-Sensor

Our sensors have **poor strong** recovery properties.
(P. & Schnörr, PMA
2009)



RIP, neighborly polytopes,
mutual coherence, ...

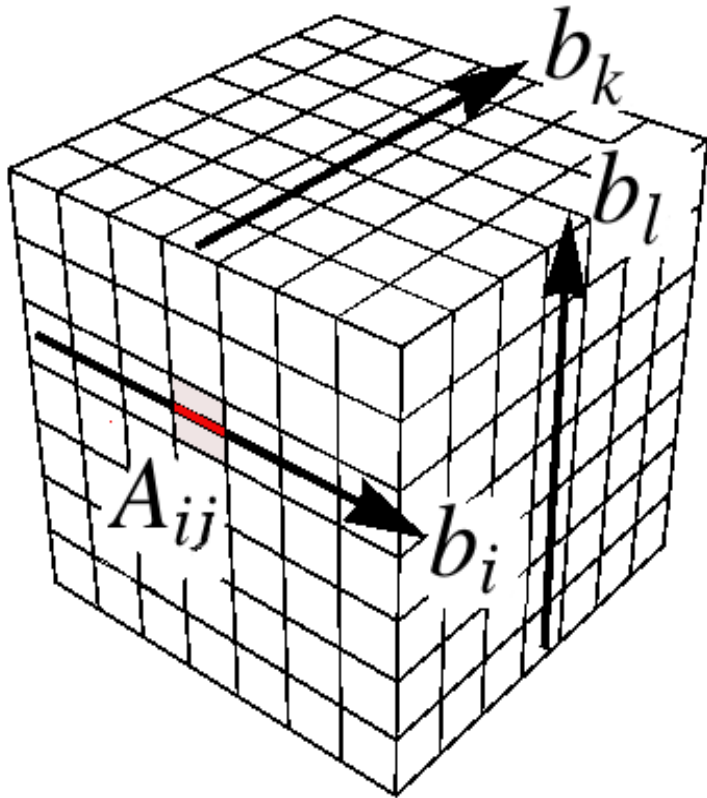
Recovery of *any* vector from

$$\mathcal{X}^n(k)$$

$\Rightarrow k$ **very small**

CS Theory Does Not Apply

Our sensors have **poor strong** recovery properties.
(P. & Schnörr, PMA
2009)




RIP, neighborly polytopes,
mutual coherence, ...

Recovery of *any* vector from



$$\mathcal{X}^n(k)$$

$\Rightarrow k$ **very small**




Practice Matches Theory

of k particles



Non-Destructive Testing (NDT)



key parameter: cosparsity,
sparse

“object complexity”
interfaces

∇x
material

Recovery Performance

(Needell, Ward 2013)

$$\Delta(b) = \operatorname{argmin}_{x: \|Ax-b\|_2 \leq \varepsilon} \|\nabla x\|_1$$

decoding

Recovery Performance

(Needell, Ward 2013)

$$b = Ax^* + \xi, \|\xi\|_2 \leq \varepsilon. \quad (\text{bounded noise})$$

and $AH^\top \leftarrow RIP_{\delta_{5k}}, \delta_{5k} < \frac{1}{3}$ H Haar transform

$$\Delta(b) = \operatorname{argmin}_{x: \|Ax-b\|_2 \leq \varepsilon} \|\nabla x\|_1$$

decoding

Stable recovery guarantee

$$\|x^* - \Delta(b)\|_2 \leq C \left(\frac{\|\nabla x^* - (\nabla x^*)_k\|_1}{\sqrt{k}} + \varepsilon \right)$$




$\varepsilon \rightarrow 0$ and k -sparsity of $\nabla x^* \Rightarrow$ perfect reconstruction

Worst Case vs Practical Recovery

Thm. (Needell, Ward 2013)
our scenario, 8 views



Image gradient ∇x^*
can be at most 6-sparse



But: perfect recovery from
8 views, $\|\nabla x^*\|_0 = 109930$

Cosparsity & Recovery

(*Nam et al, 2013*) $\Lambda := \{i : (\nabla x)_i = 0\}$ cosupport

cosparsity $\ell := |\Lambda| = \dim(\nabla x) - \|\nabla x\|_0$

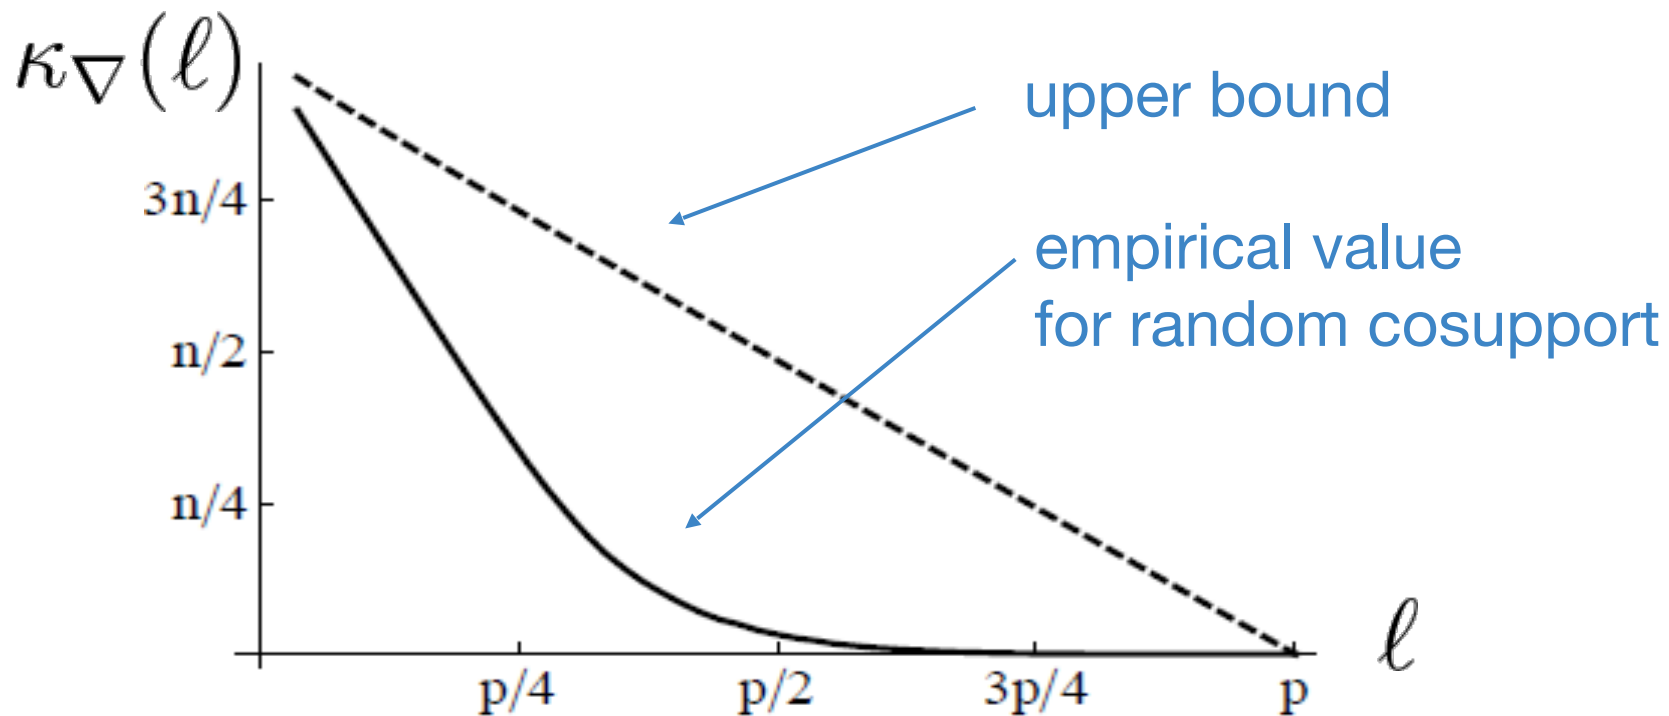
key quantity $\kappa_{\nabla}(\ell) := \max_{|\Lambda|=\ell} \dim \ker \nabla_{\Lambda}$

Cosparsity & Recovery



(Denitiu et al, 2014)

$$m \geq 2n - (\ell + \sqrt{2\ell + 1} - 1) \quad (2D)$$


$$m \geq 2n - \frac{2}{3} \left(\ell + \sqrt[3]{3\ell^2} + 2\sqrt[3]{\frac{\ell}{3}} - 2 \right) \quad (3D)$$




Recovery for Solid Objects




Compressed Motion Sensing



entire
video



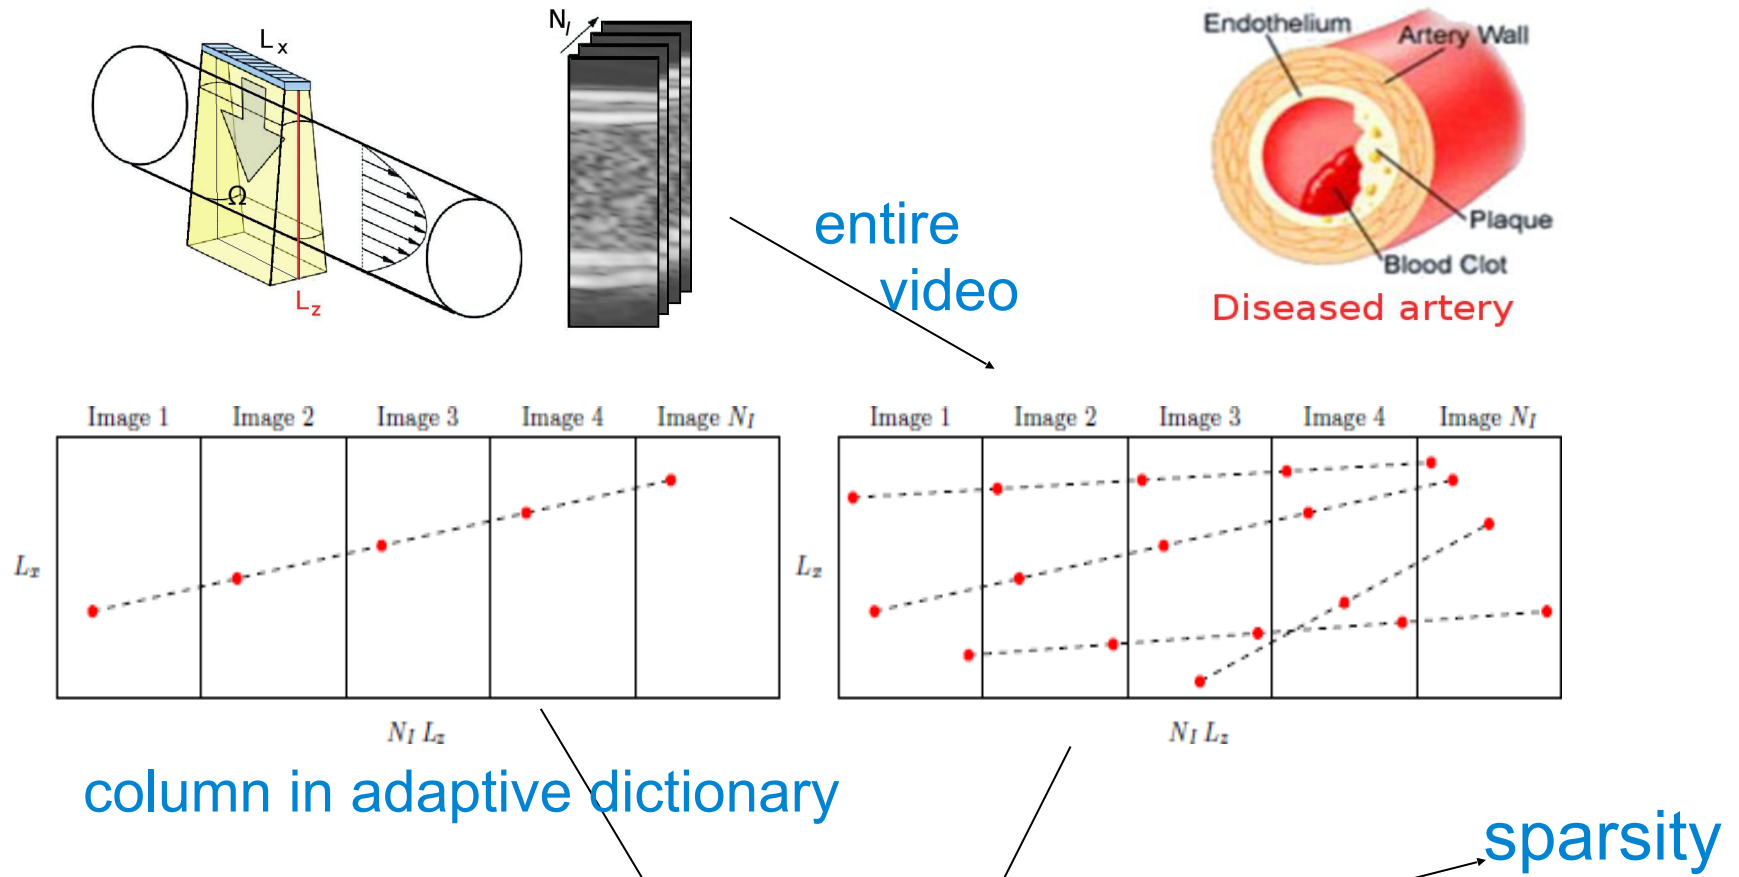
ultrasound imaging
Echo PIV



Few trajectories

(Bodnariuc et al, EMMCVPR 2014)

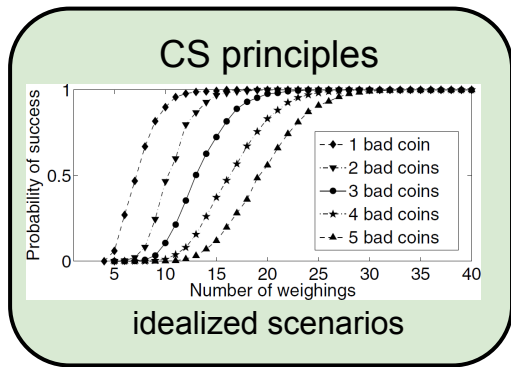
Compressed Motion Sensing



$$\min_v \min_{u \in [0,1]^n} \|D(v)u - F\|_1 + \lambda \|u\|_1$$


(Bodnariuc et al, EMMCVPR 2014)

Overview



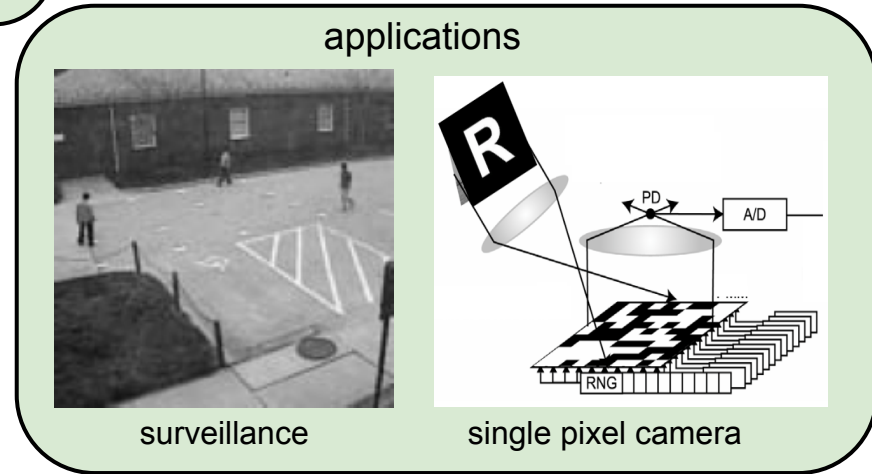


enable

RIP fails




applies

applies



RIP fails



A Mathematical Introduction to Compressed Sensing

Stefania Petra
Mathematical Imaging Group
Heidelberg University

Lecture WT 2018/19

Content

1. Introduction

1.1. Motivation

1.2. Overview

2. Sparse Solutions of Underdetermined Linear Systems

2.1. Vector Spaces

2.2. Sparsity and Compressibility

2.3. Unique Recovery of Sparse Vectors

2.4. NP Hardness of ℓ_0 -Minimization

2.5. ℓ_0 - vs. ℓ_1 -Minimization for Bounded Systems of Linear Equalities and Inequalities

Content

3. Basis Pursuit

3.1. Reformulation of l_1 -Minimization as a Linear Program

3.2. Null Space Property

3.3. Certifying Recovery of Individual Vectors via Dual Certificates

3.4. Alternative Recovery Conditions for l_1 -Minimization

4. Restricted Isometry Property

4.1. Stability and Robustness

4.2. RIP and Measurement Bounds

5. Coherence

5.1. Relation to Spark and RIP

Content

6. Sparse Recovery with Random Matrices

6.1. Basic Set-Up and Definitions

6.2. Some Basic Relations and Facts

6.3. Subgaussian Random Variables and Matrices

6.4. RIP and Uniform Sparse Recovery

6.5. Johnson-Lindenstrauss Embeddings

7. Beyond ℓ_1 -Minimization: Further Low Complexity Models

7.1. Introduction

7.2. Sparse Representations and Norms

7.3. Recovery from Random Measurements

8. Recovery Algorithms

8.1. Greedy Methods for Sparse Recovery

8.2. Basic Algorithms for Non-Smooth Convex Programming

8.3. A Non-Convex Optimization Approach

Content

Appendix A. Notation

Appendix B. Supplementary Definitions and Relations

B.1. Basic Inequalities

B.2. Factorials, Stirling's Approximation

B.3. Gamma Function

B.4. Singular Values, SVD

Appendix C. Concepts and Results from Convex Analysis

C.1. Some Basic Concepts

C.2. Conjugation, Legendre-Fenchel Transform

C.3. Fenchel Duality of Convex Problems

References