


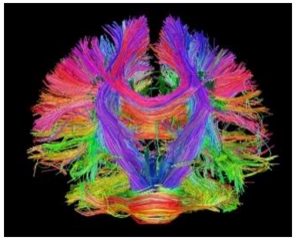


CS-570
Statistical Signal Processing

Lecture 8: Compressed Sensing

Spring Semester 2019

Grigorios Tsagkatakis

Today's Objectives

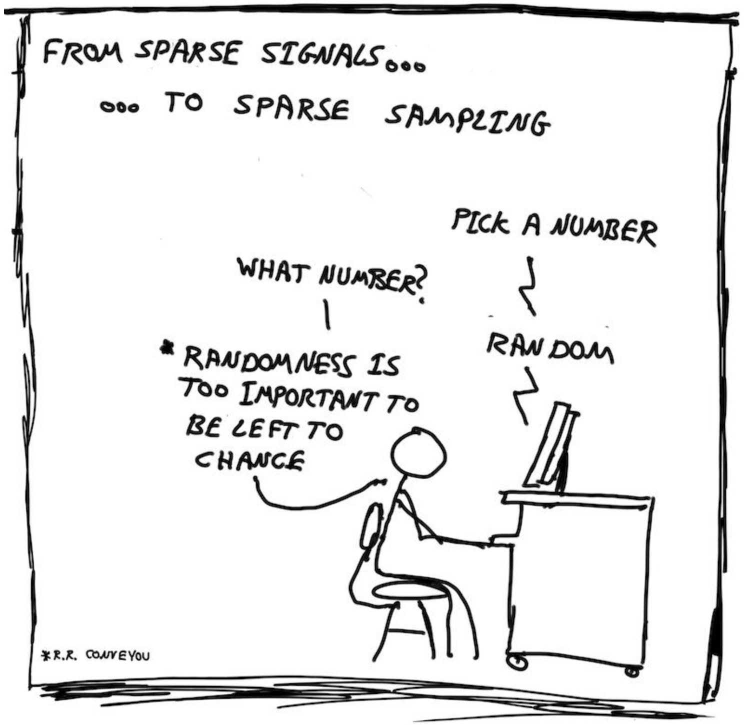

Topics:

- Compressed Sensing

Disclaimer: Material used:

Richard Baraniuk Talks, Rice University

Candès, Emmanuel J., and Michael B. Wakin. "An introduction to compressive sampling" *IEEE signal processing magazine* 25.2 (2008)

challenge 1
data too expensive

Case in Point: MR Imaging

- Measurements **very expensive**
- \$1-3 million per machine
- 30 minutes per scan

Case in Point: IR Imaging

Thermal Imaging Infrared Camera Thermography w/45° Lens

Home > Thermal Imaging Cameras >

Thermal Imaging Infrared Camera Thermography w/45° Lens

Item Id: 55902-2302

Review this item

0

Like — Be the first of your friends to like this.

Your Price: $25,950.00

Availability: Back Order

Add to Wish List

Quantity 1 **Add to Cart**

TruReward$ Points

25950 TruReward$ points will be placed in your account when you buy this item.

High Definition Images

320x240 **320 x 240 pixel resolution**

challenge 2
too much data

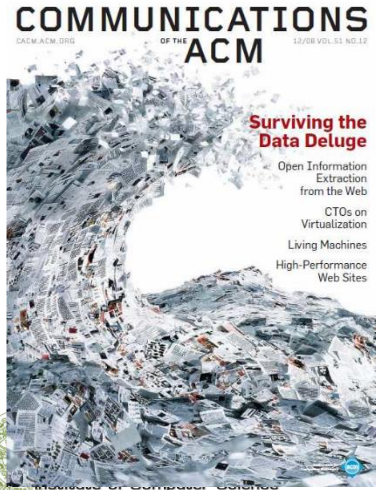

Example

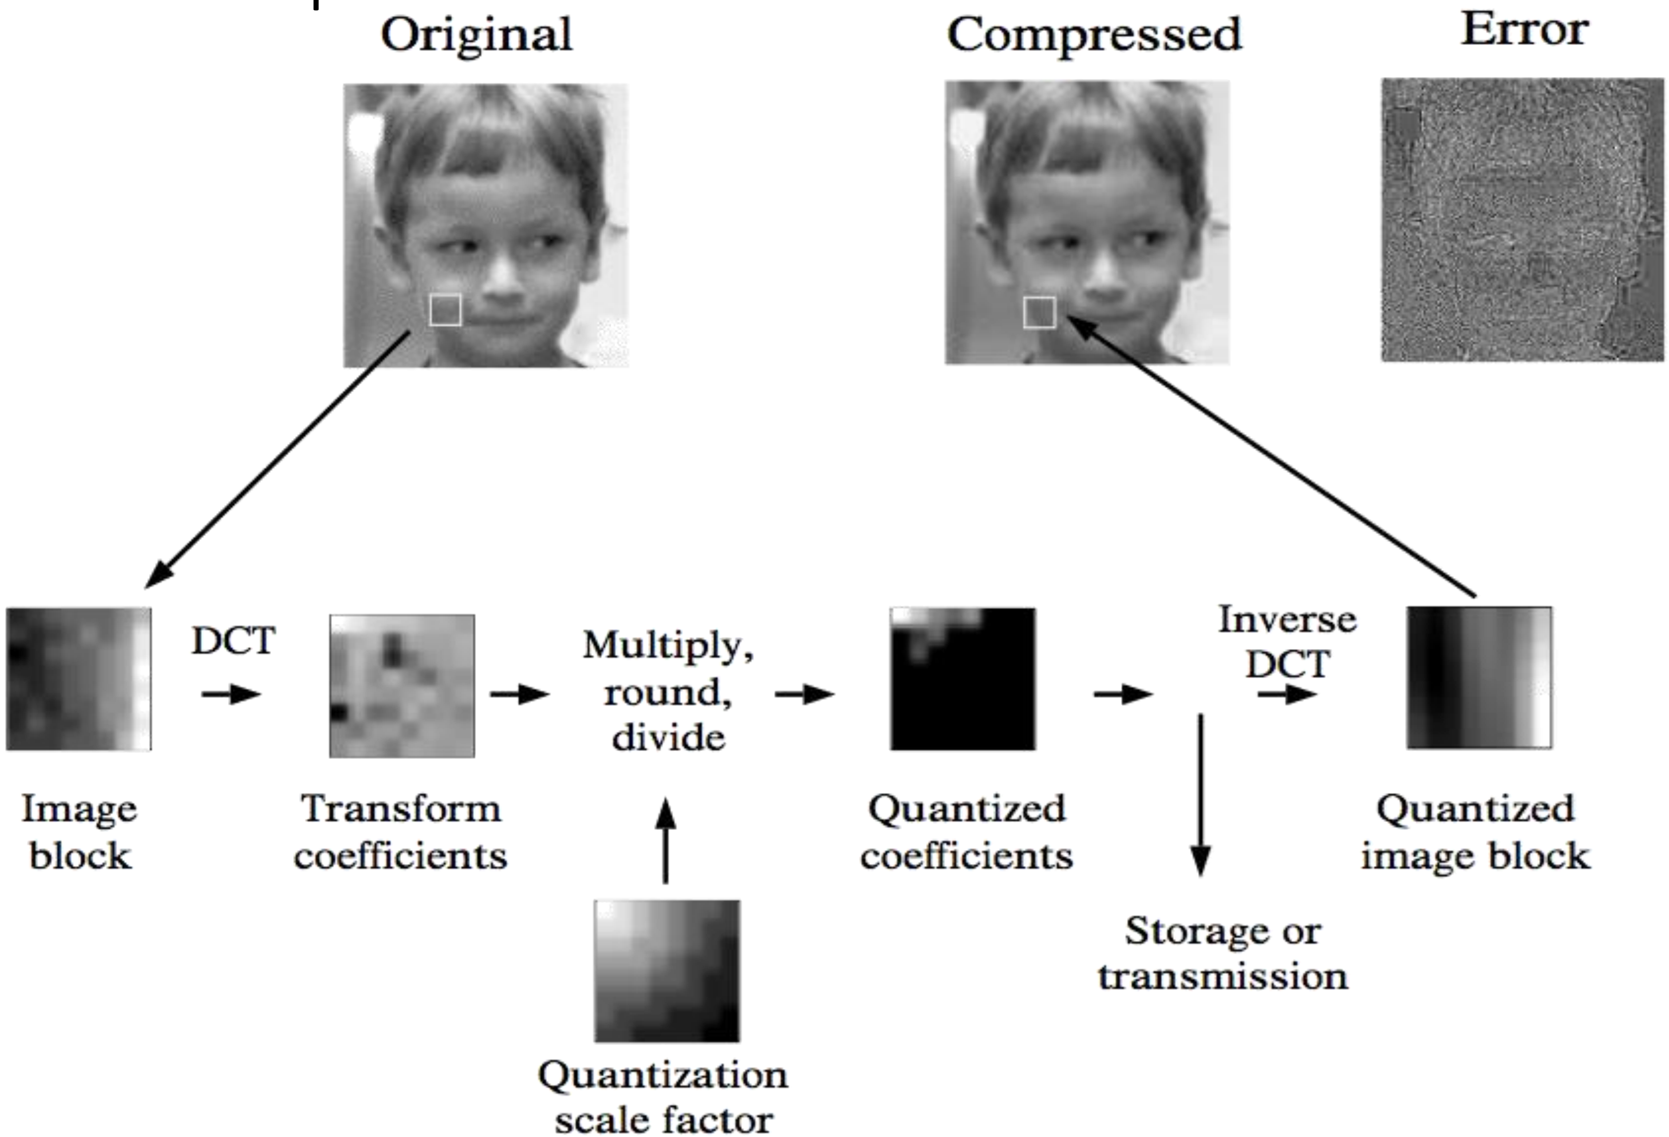
Original
Compressed
Error
DCT
Multiply, round, divide
Inverse DCT
Image block
Transform coefficients
Quantized coefficients
Quantized image block
Storage or transmission
Quantization scale factor

Compressability

N
pixels

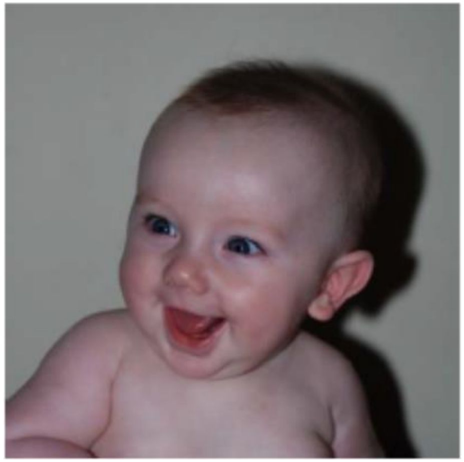

$K \ll N$
large
wavelet
coefficients

(blue = 0)

N
wideband
signal
samples

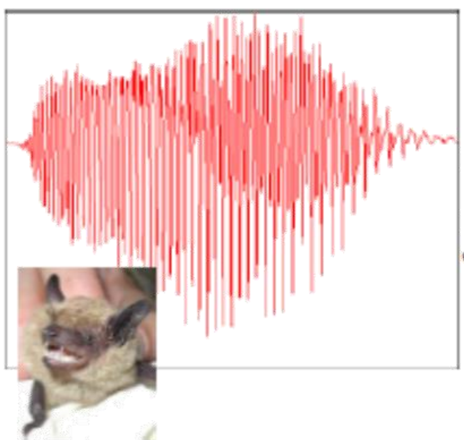

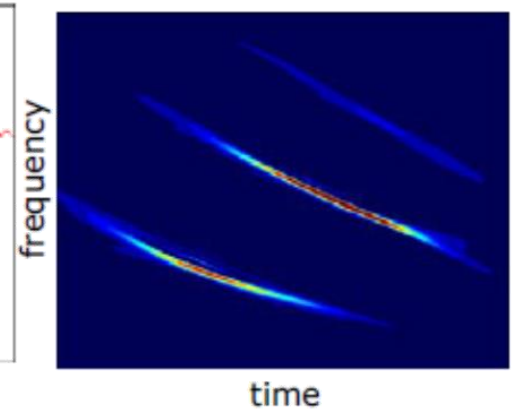

$K \ll N$
large
Gabor (TF)
coefficients

Approximation

1 megapixel image

25k term approx

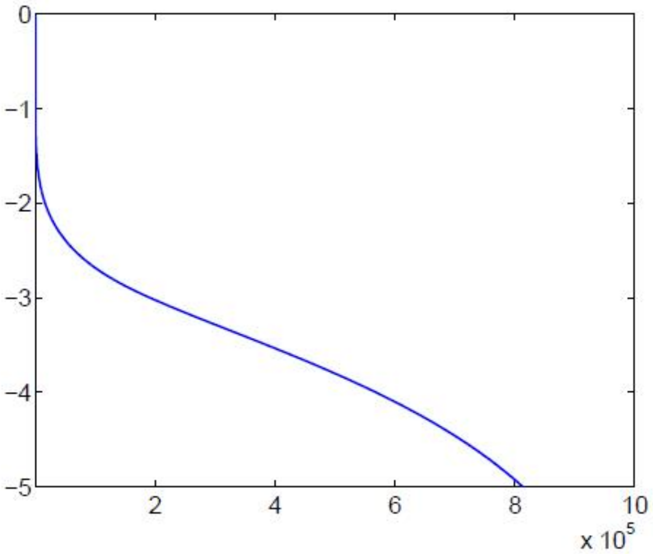

B-term approx error

- Within 2 digits (in MSE) with ~ 2.5% of coeffs

Sensing by Sampling

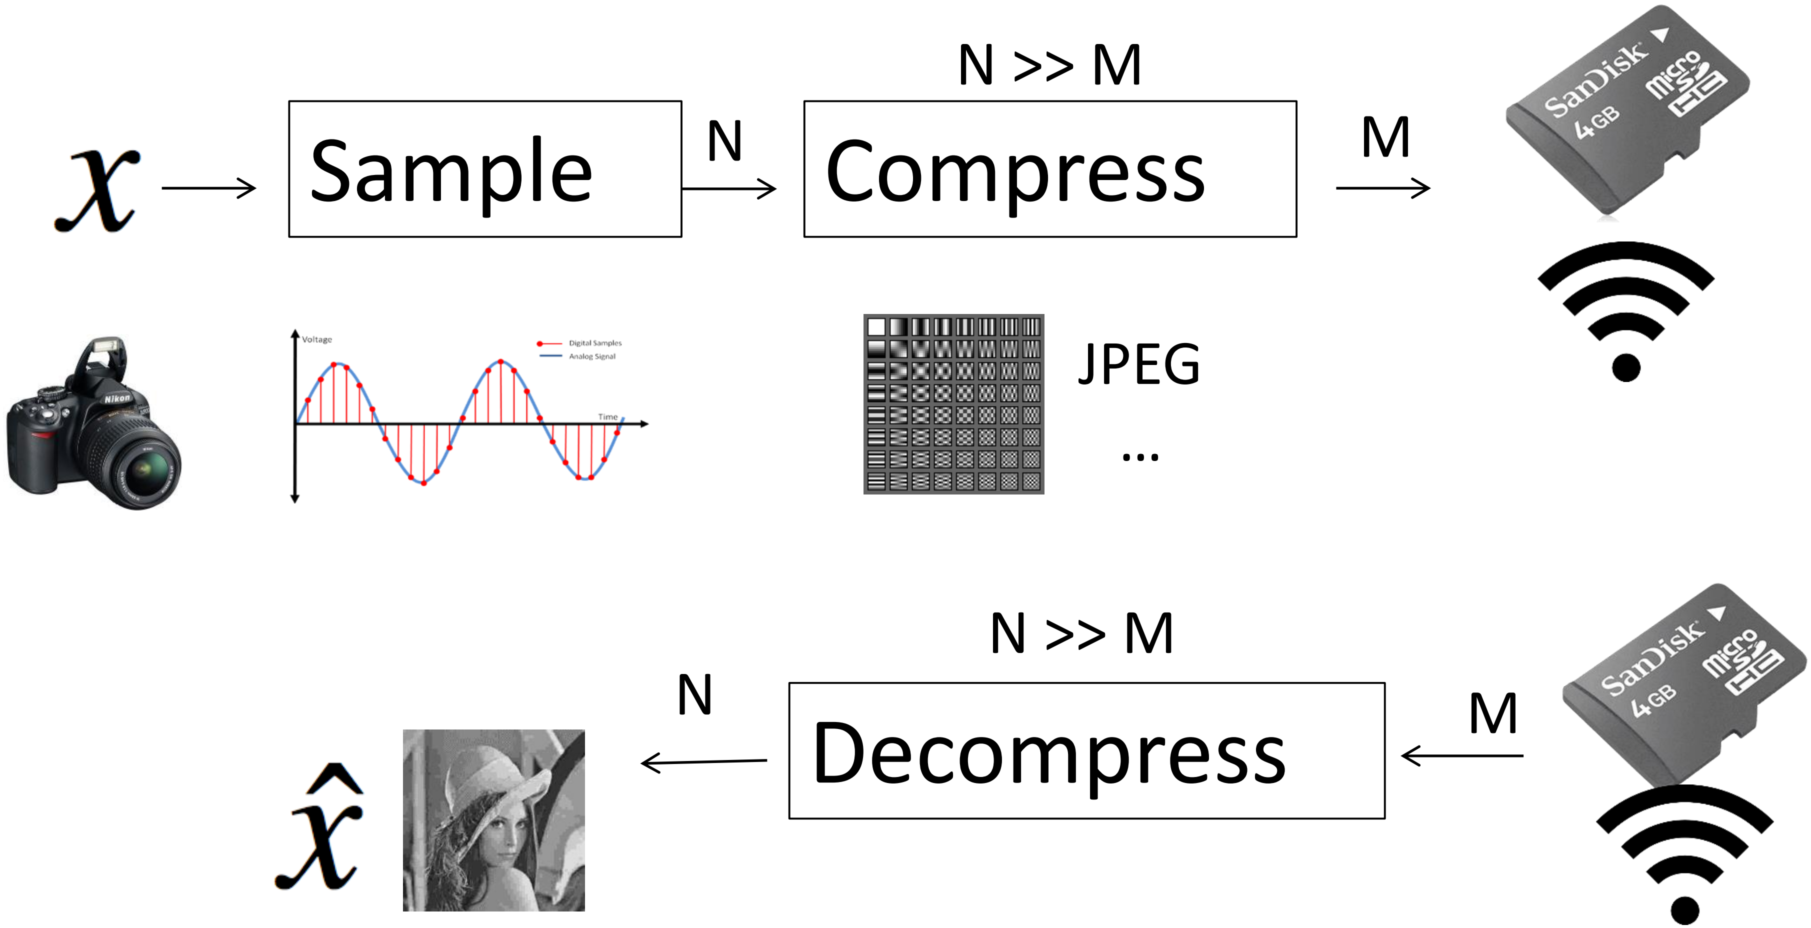

Compressed Sensing

- Directly acquire "***compressed***" data
- Replace samples by more general "measurements"

$$K \approx \underline{M} \ll N$$

$x \rightarrow$ **compressive sensing** $\xrightarrow[y]{M}$ transmit/store

receive $\xrightarrow{M}$ reconstruct $\xrightarrow{N}$ $\widehat{x}$

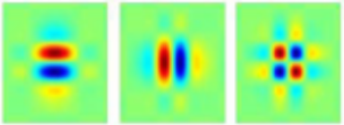

2004—2019

INFORMATION THEORY, VOL. 52, NO. 4, APRIL 2006 1289

Compressed Sensing

David L. Donoho, *Member, IEEE*

Abstract—Suppose x is an unknown vector in $\mathbf{R}^m$ (a digital image or signal); we plan to measure n general linear functionals of x and then reconstruct. If x is known to be compressible by transform coding with a known transform, and we reconstruct via the nonlinear procedure defined here, the number of measurements n can be dramatically smaller than the size m. Thus, certain natural classes of images with m pixels need only $n = O(m^{1/4}\log^{5/2}(m))$ nonadaptive nonpixel samples for faithful recovery, as opposed to the usual m pixel samples.

More specifically, suppose x has a sparse representation in some orthonormal basis (e.g., wavelet, Fourier) or tight frame (e.g., curvelet, Gabor)—so the coefficients belong to an ℓ_p ball

the important information about the signals/images—in effect, not acquiring that part of the data that would eventually just be

23176 citations



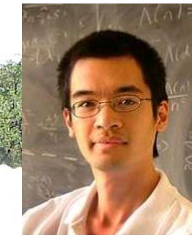

IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. 52, NO. 2, FEBRUARY 2006 489

Robust Uncertainty Principles: Exact Signal Reconstruction From Highly Incomplete Frequency Information

Emmanuel J. Candès, Justin Romberg, *Member, IEEE*, and Terence Tao

Abstract—This paper considers the model problem of reconstructing an object from incomplete frequency samples. Consider a discrete-time signal $f \in \mathbf{C}^N$ and a randomly chosen set of frequencies Ω. Is it possible to reconstruct f from the partial knowledge of its Fourier coefficients on the set Ω?

A typical result of this paper is as follows. Suppose that f is a superposition of $|T|$ spikes $f(t) = \sum_{\tau \in T} f(\tau)\delta(t-\tau)$ obeying

$$|T| \le C_M \cdot (\log N)^{-1} \cdot |\Omega|$$

I. Introduction

In many applications of practical interest, we often wish to reconstruct an object (a discrete signal, a discrete image, etc.) from incomplete Fourier samples. In a discrete setting, we may pose the problem as follows; let $\hat{f}$ be the Fourier transform of a discrete object $f(t)$, $t = (t_1, \ldots, t_d) \in \mathbb{Z}_N^d := \{0, 1, \ldots, N-1\}^d$

14278 citations

dsp.rice.edu/cs archive >1500 papers

nuit-blanche.blogspot.com > 1 posting/sec

Traditional sampling

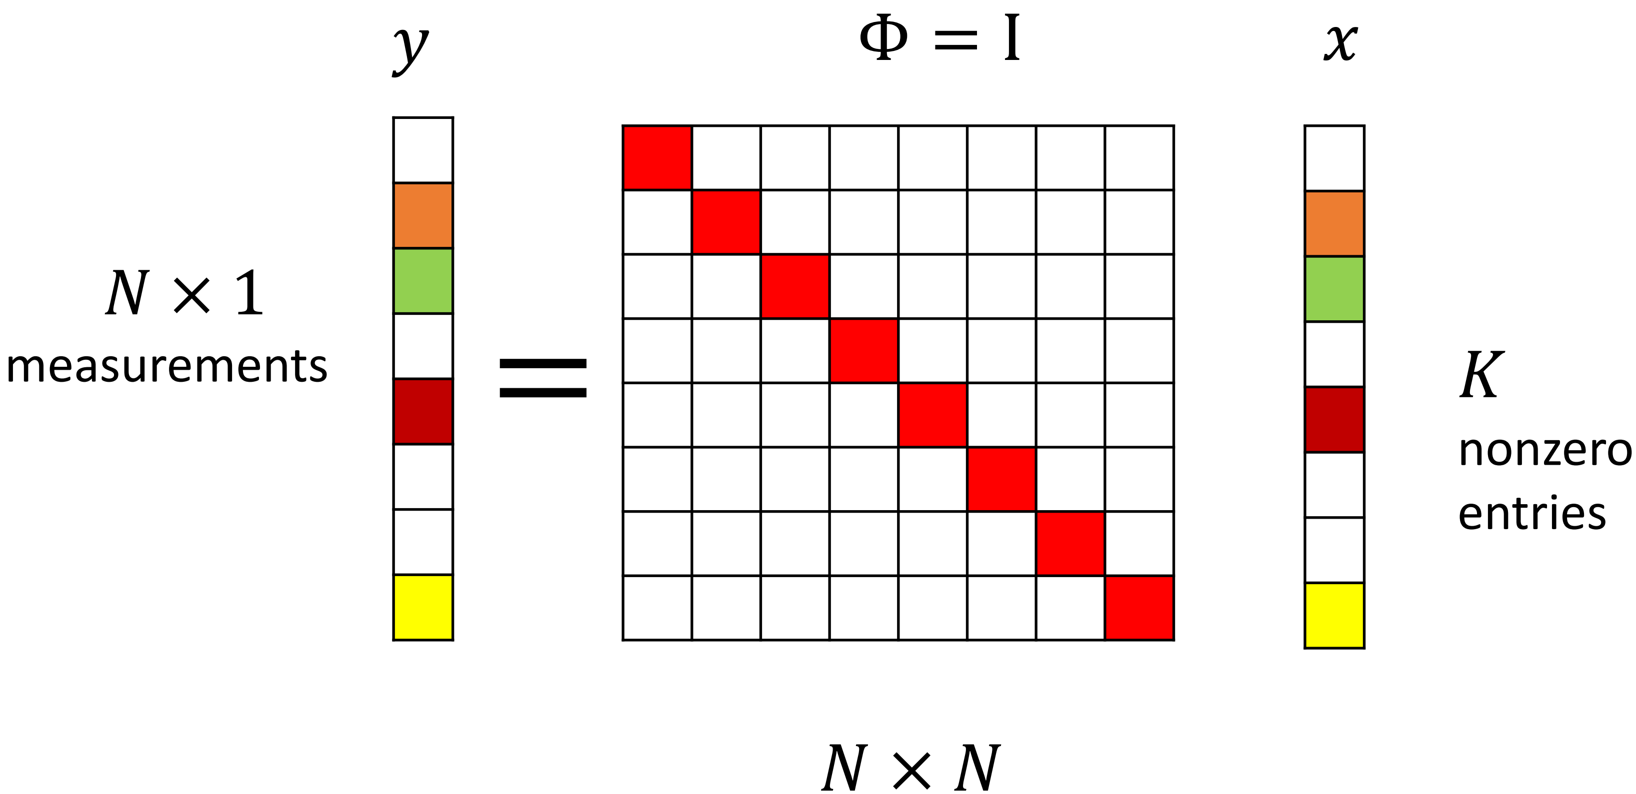

Compressed sampling

$M \times 1$
measurements

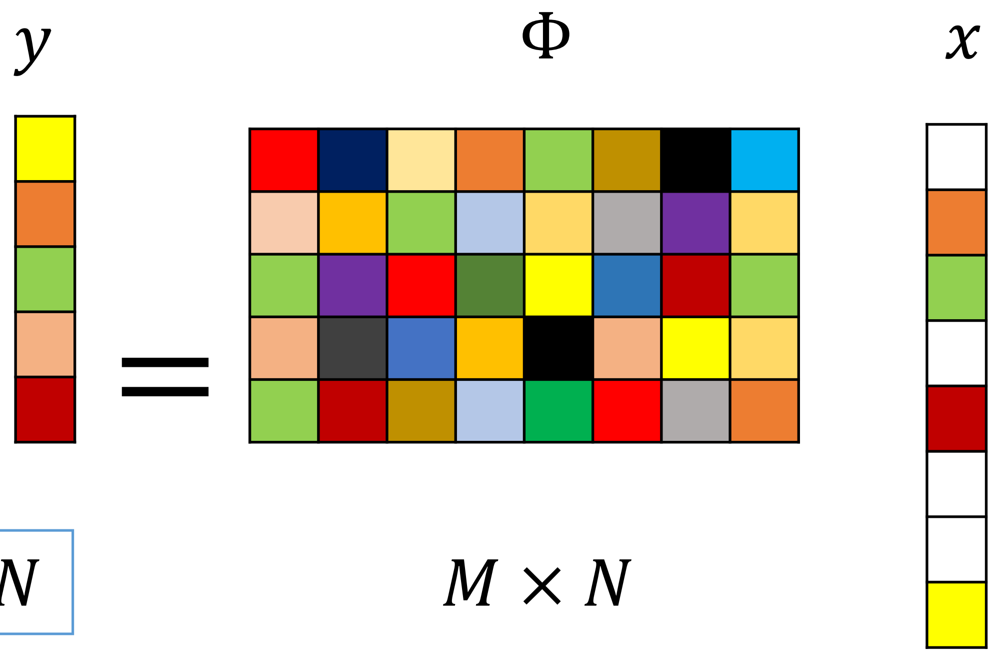

K
nonzero
entries

$K < M \ll N$

Projection **not full rank…**

… and so **loses information** in general

- Ex: Infinitely many x’s map to the same y

Compressed sampling

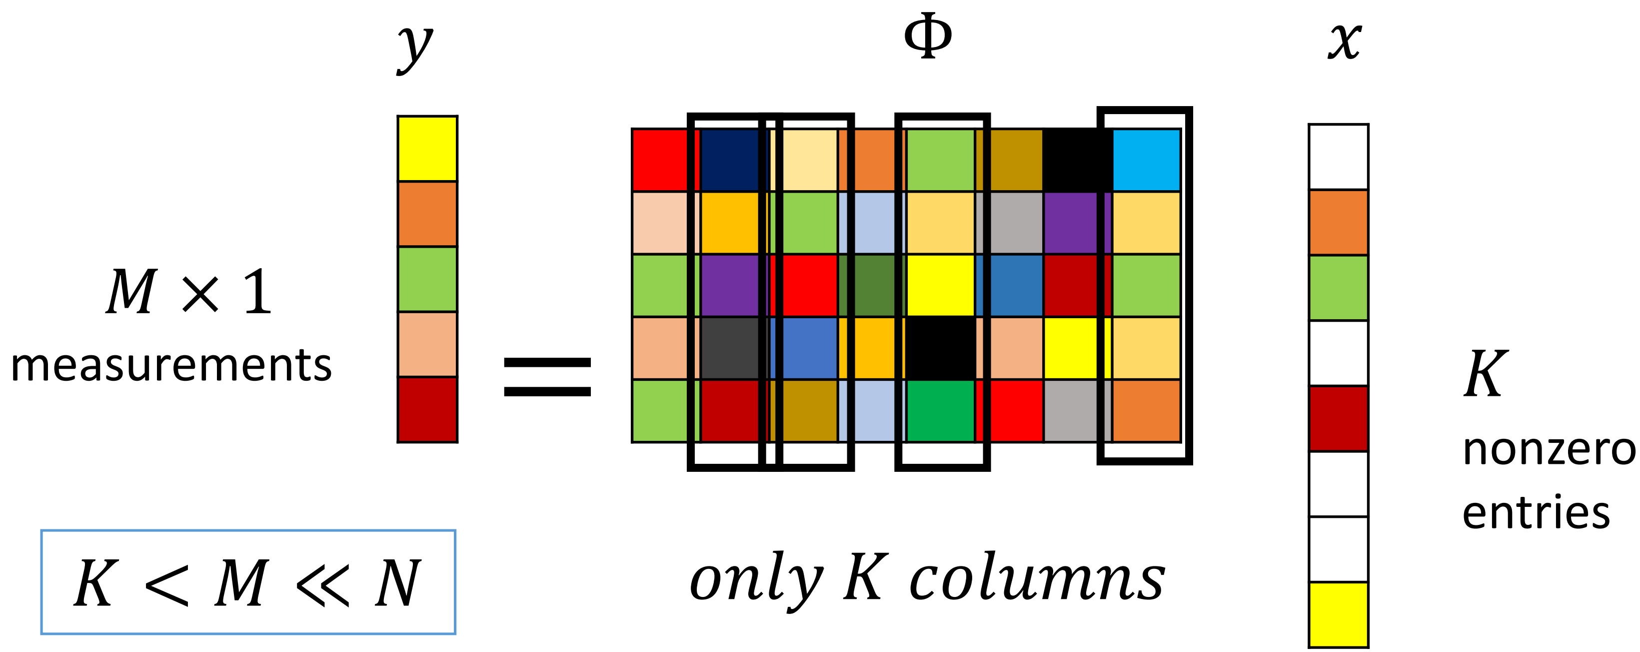

- But we are only interested in ***sparse*** vectors
- Φ is effectively MxK
- Φ should be designed such that each MxK submatrix is full rank
- To preserve distances require al Mx2K submatrices are full rank

NP hard

Stable Embedding

- An information preserving projection Φ preserves the **geometry** of the set of sparse signals

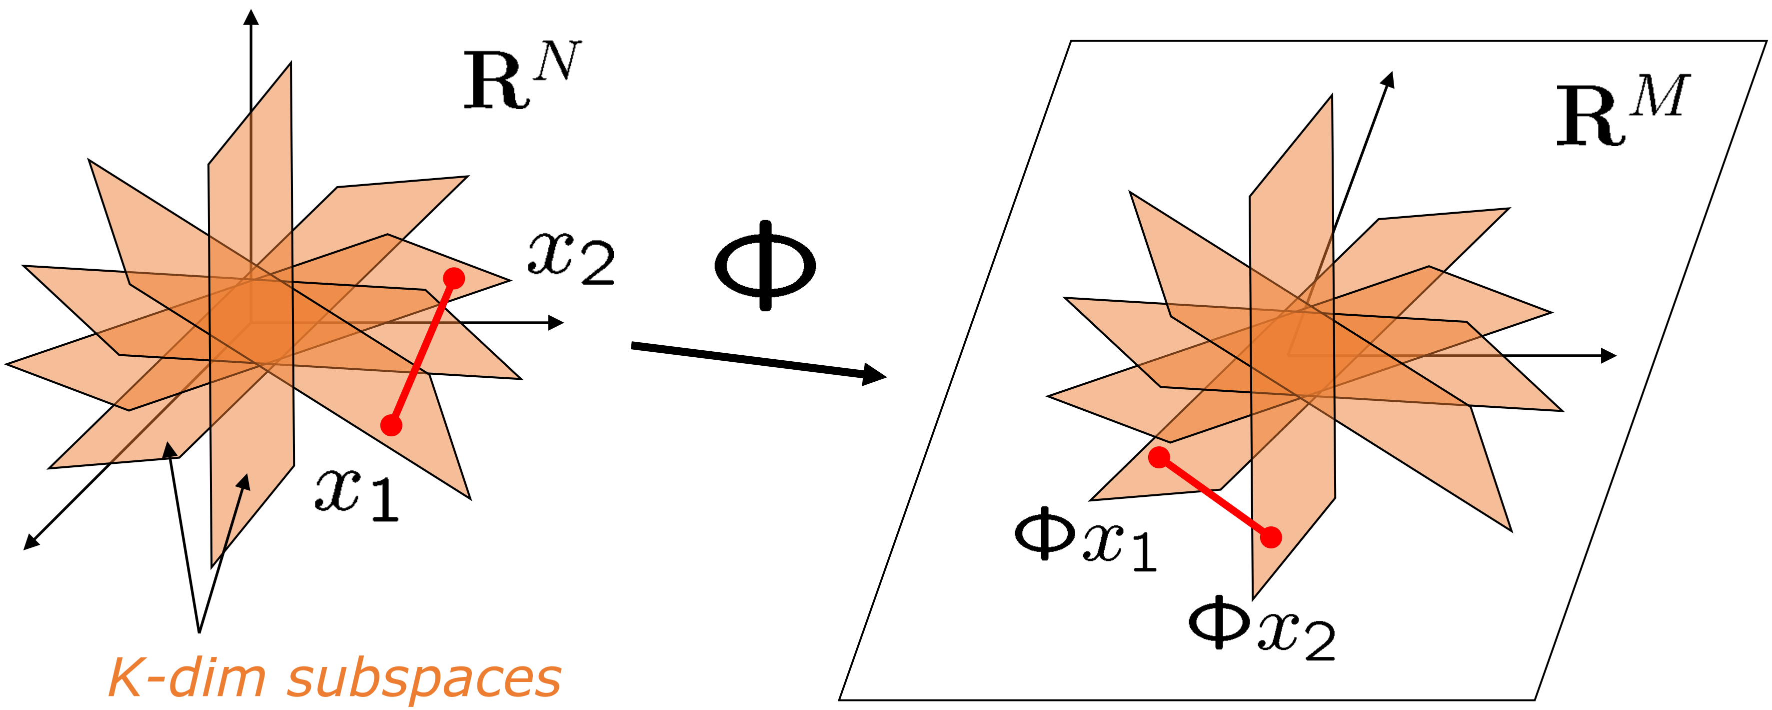

- SE ensures that

$$\|x_1 - x_2\|_2 \approx \|\Phi x_1 - \Phi x_2\|_2$$

Random Embedding is Stable

- Measurements y = **random linear combinations** of the entries of x

- **No information loss** for sparse vectors x whp

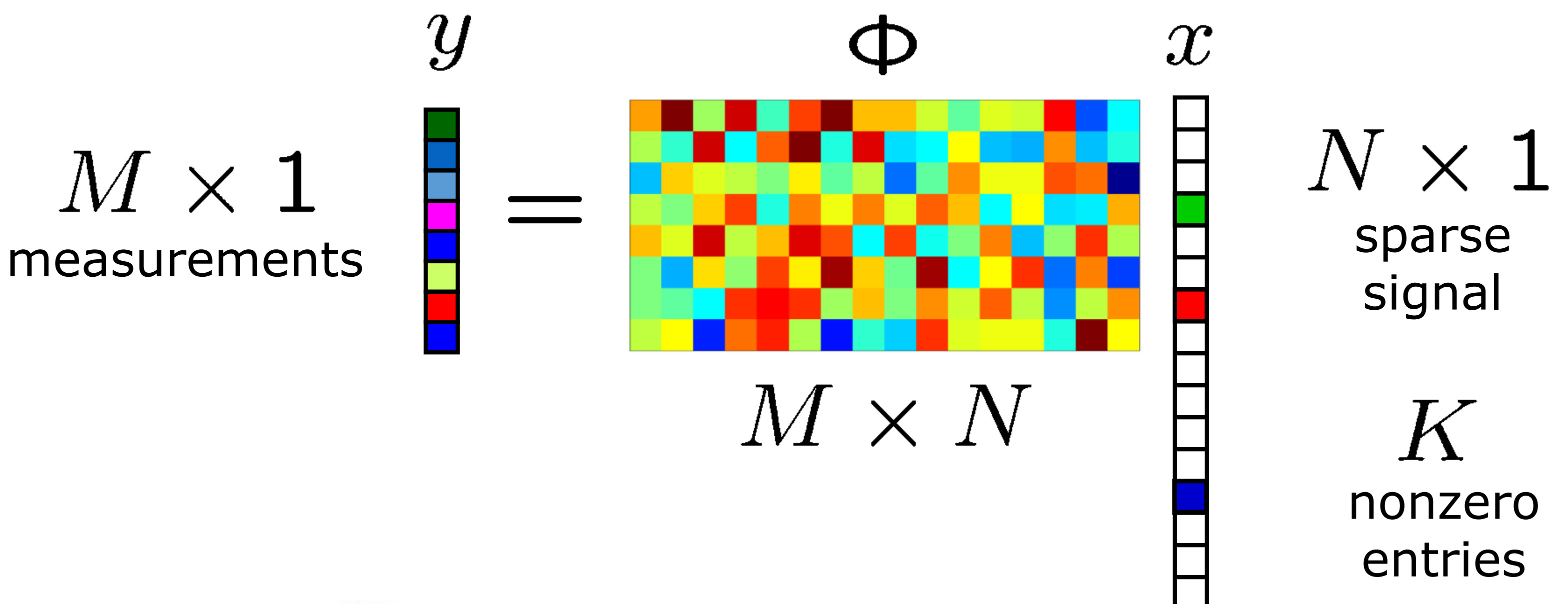

Preservation of distances

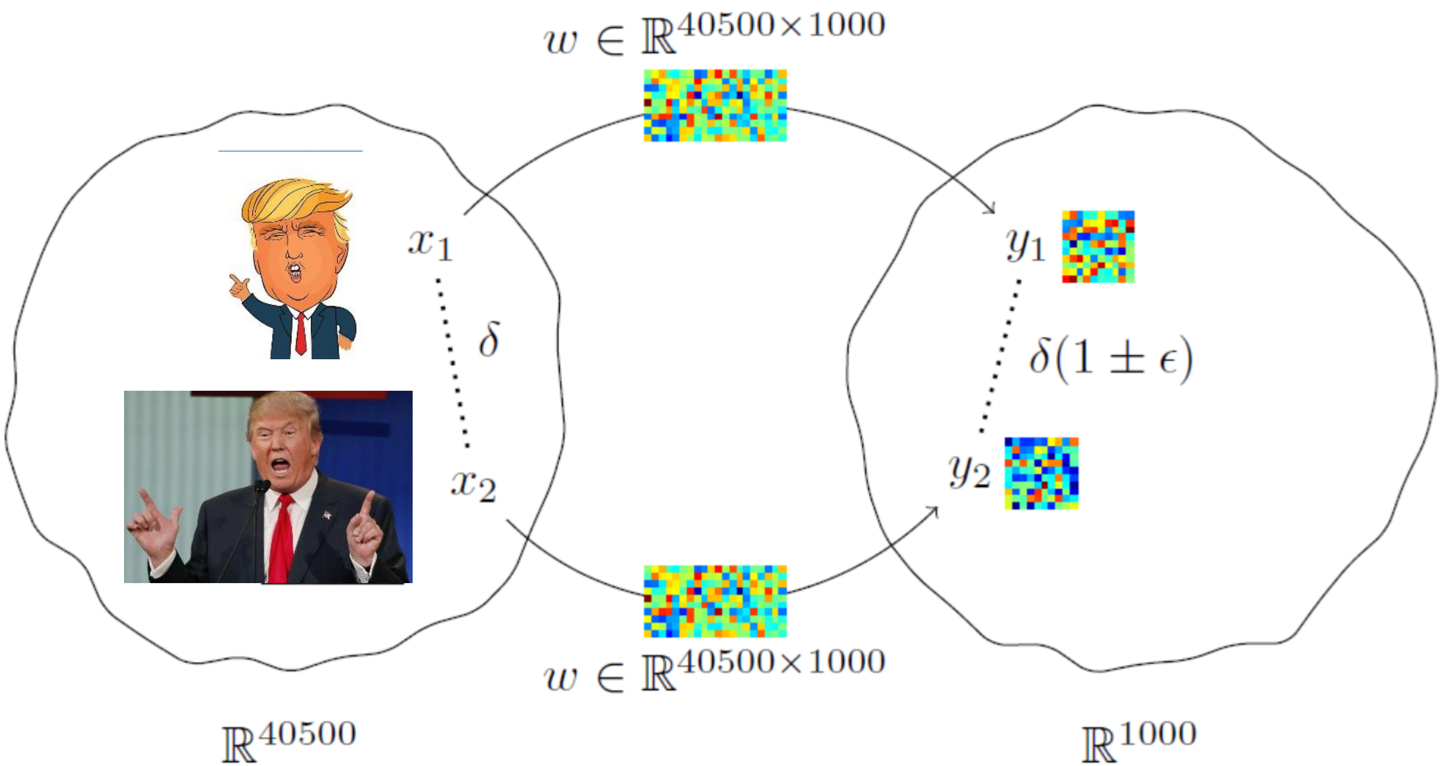

$$(1-\epsilon)\|x_1 - x_2\|^2 \leq \|y_1 - y_2\|^2 \leq (1+\epsilon)\|x_1 - x_2\|^2$$

This result is formalized in the *Johnson-Lindenstrauss Lemma*

The Johnson-Lindenstrauss Lemma

For any $0 < \epsilon < 1/2$ and any integer $m > 4$, let $k = \frac{20 \log m}{\epsilon^2}$. Then, for any set V of m points in $\mathbb{R}^N$ $\exists$ $f : \mathbb{R}^N \to \mathbb{R}^k$ s.t. $\forall$ $\boldsymbol{u}, \boldsymbol{v} \in V$:

$$(1-\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^2 \leq \|f(\boldsymbol{u}) - f(\boldsymbol{v})\|^2 \leq (1+\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^2.$$

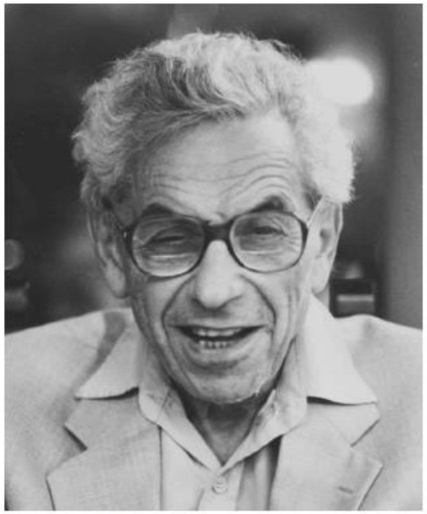

Paul Erdös
1913-1996

Joram Lindenstrauss
1936-2012

William B. Johnson
1944-

Random Projections

Lemma (Frankl Meahara (1987))

Let $\mathbb{D}$ denote an i.i.d. Gaussian distribution for each entry of Φ. Then, $\mathbb{D}$ exhibits the JL property.

Lemma (Achlioptas (2003))

Let $\mathbb{D}$ denote an i.i.d. ± 1 distribution for each entry of Φ. Then, $\mathbb{D}$ exhibits the JL property.

Random subgaussian matrix Φ has the **RIP** w.h.p. if

$$M = O(K + \log \binom{N}{K}) = O(K \log(N/K))$$

$y_1 = \Big\langle$  , 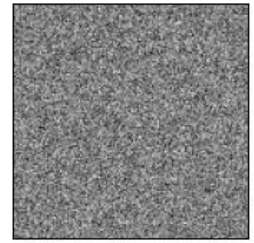 $\Big\rangle$

$y_2 = \Big\langle$  , 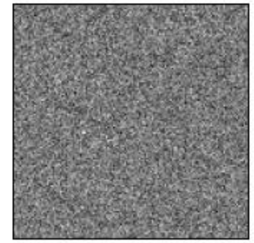 $\Big\rangle$

$y_3 = \Big\langle$  , 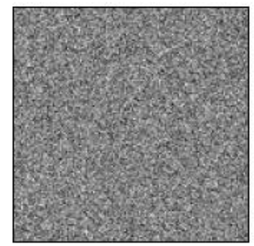 $\Big\rangle$

$\vdots$

$y_M = \Big\langle$  , 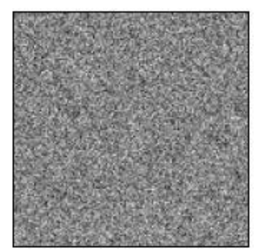 $\Big\rangle$

Representation vs. Measurements

- Image structure: local, coherent

  Good basis functions:

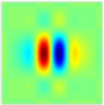
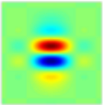

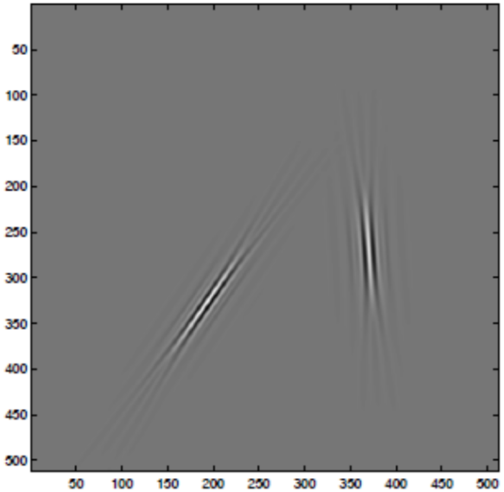

- Measurements: global, incoherent

  Good test functions:

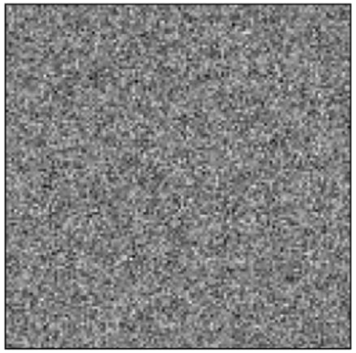
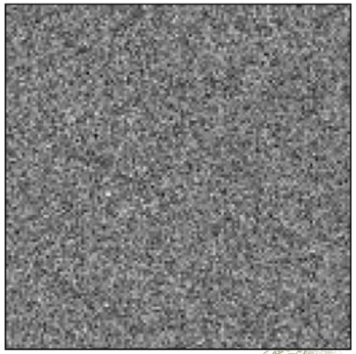
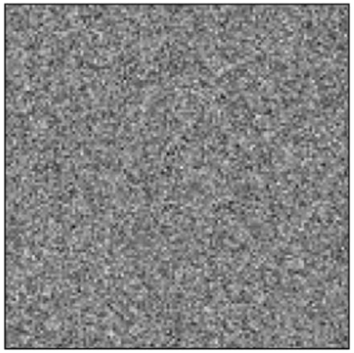

Sparse coding

General framework $\min_{\mathbf{x}} \|\mathbf{x}\|_0 \ \text{s.t.}\ \|\mathbf{y} - \mathbf{\Phi x}\|_2 \leq \varepsilon$

- Challenge: finding $|x|_0$ is NP hard

y Φ x

$M \times 1$
measurements

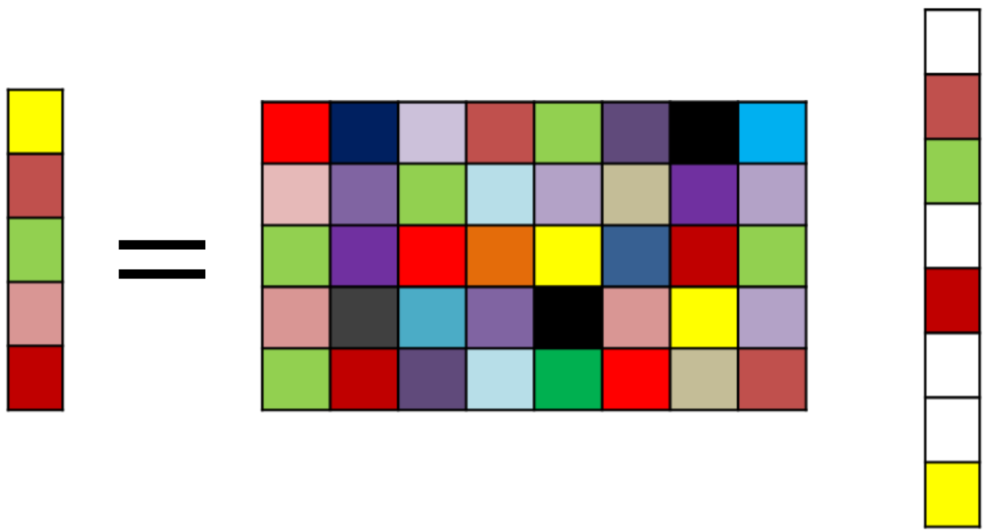

$N \times 1$
sparse signal

K
nonzero
entries

$K < M \ll N$

$M \times N$

Signal Structure

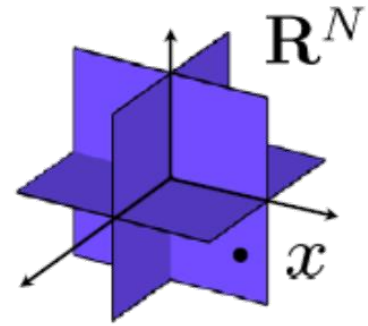

- **Sparse** signal: only K out of N coordinates nonzero
  - model: union of K-dimensional subspaces

- **Compressible** signal: sorted coordinates decay rapidly to zero
  - model: ℓ_p ball: $\|x\|_p^p = \sum_i |x_i|^p \leq 1, \; p \leq 1$

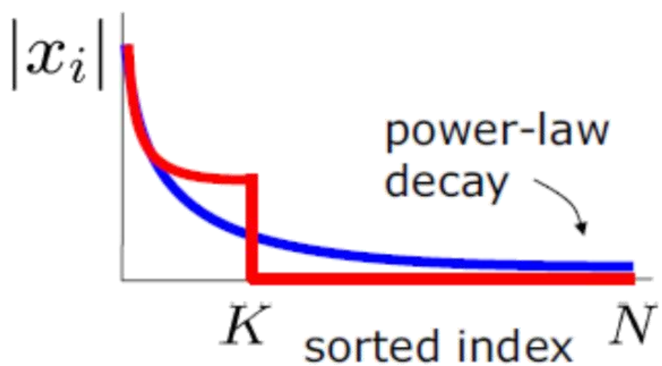

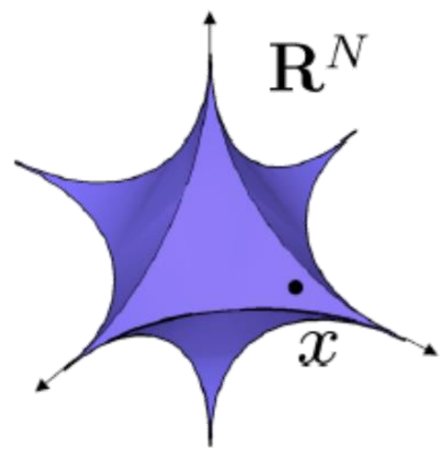

Signal Recovery

- Recovery: given $y = \Phi x$
  (ill-posed inverse problem) find x (sparse)

- ℓ_2 **fast, wrong** $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$

x

$$\widehat{x} = (\Phi^T \Phi)^{-1} \Phi^T y$$

pseudoinverse

Why l2 does not work

for signals sparse in the
space/time domain

$$\widehat{x} = \arg \min_{y=\Phi x'} \|x'\|_2$$

least squares,
minimum ℓ_2 solution
is almost **never sparse**

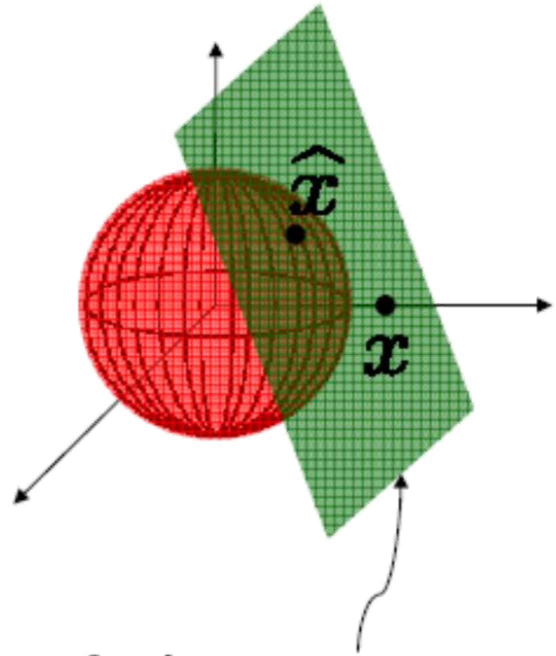

$$\{x' : \ y = \Phi x'\}$$

null space of Φ
translated to x
(random angle)

Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (ill-posed inverse problem) find x

- ℓ_2 fast, wrong $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$

- ℓ_0 **correct:**
  only M=2K
  measurements
  required to
  stably reconstruct
  K-sparse signal

  slow: NP-complete
  algorithm

  $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_0$

  number of nonzero entries

Signal Recovery

- Recovery: given $y = \Phi x$
  (ill-posed inverse problem) find x (sparse)
- ℓ_2 fast, wrong $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$
- ℓ_0 correct, slow $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_0$
- ℓ_1 **correct, efficient**
  mild oversampling $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_1$
  [Candes, Romberg, Tao; Donoho] ***linear program***

number of measurements required

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

Why l1 works

for signals sparse in the
space/time domain

$$\widehat{x} = \arg\min_{y=\Phi x'} \|x'\|_1$$

minimum ℓ_1 solution
= sparsest solution
(with high probability) if

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

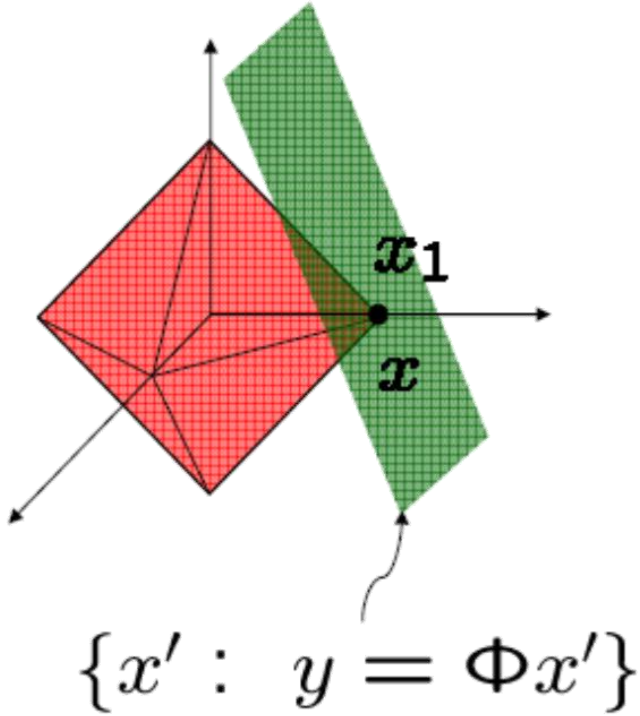

How many measurements?

Φ satisfies Restricted Isometry Property (RIP)

For all x that are K sparse

$$(1-\delta_K)||x||_2^2 \le ||\Phi x||_2^2 \le (1+\delta_K)||x||_2^2$$

When Φ MxN satisfies RIP of order 2K with δ<√(2)-1,

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

- Random (sub-) Gaussian (iid Gaussian, Bernoulli) satisfy RIP

Standard CS Recovery

- **Iterative Thresholding**

Given $y = \Phi x$, recover a sparse x

initialize: $\widehat{x}_0 = 0$, $r = y$, $i = 0$

iteration:

- $i \leftarrow i + 1$
- $b \leftarrow \widehat{x}_{i-1} + \Phi^T r$ **update signal estimate**
- $\widehat{x}_i \leftarrow \text{thresh}(b, K)$ **prune signal estimate** (best *K*-term approx)
- $r \leftarrow y - \Phi \widehat{x}_i$ **update residual**

return: $\widehat{x} \leftarrow \widehat{x}_i$

Adapted from "**Model-based Compressive Sensing**", by Volkan Cevher

Recovery of non-sparse signals

Measurement matrix

$$y = \Phi X$$

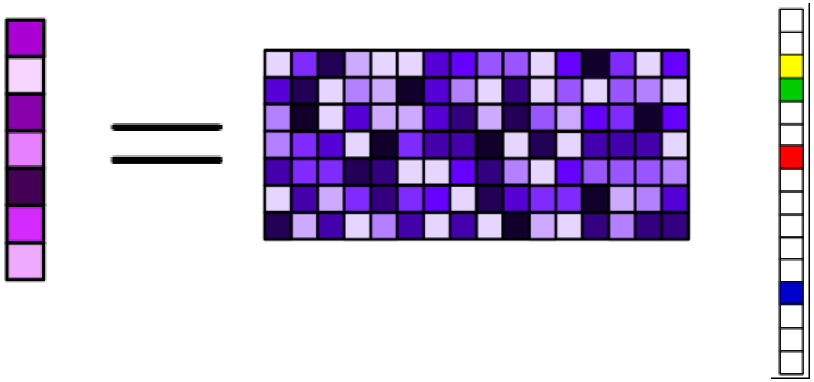

Recovery of non-sparse signals

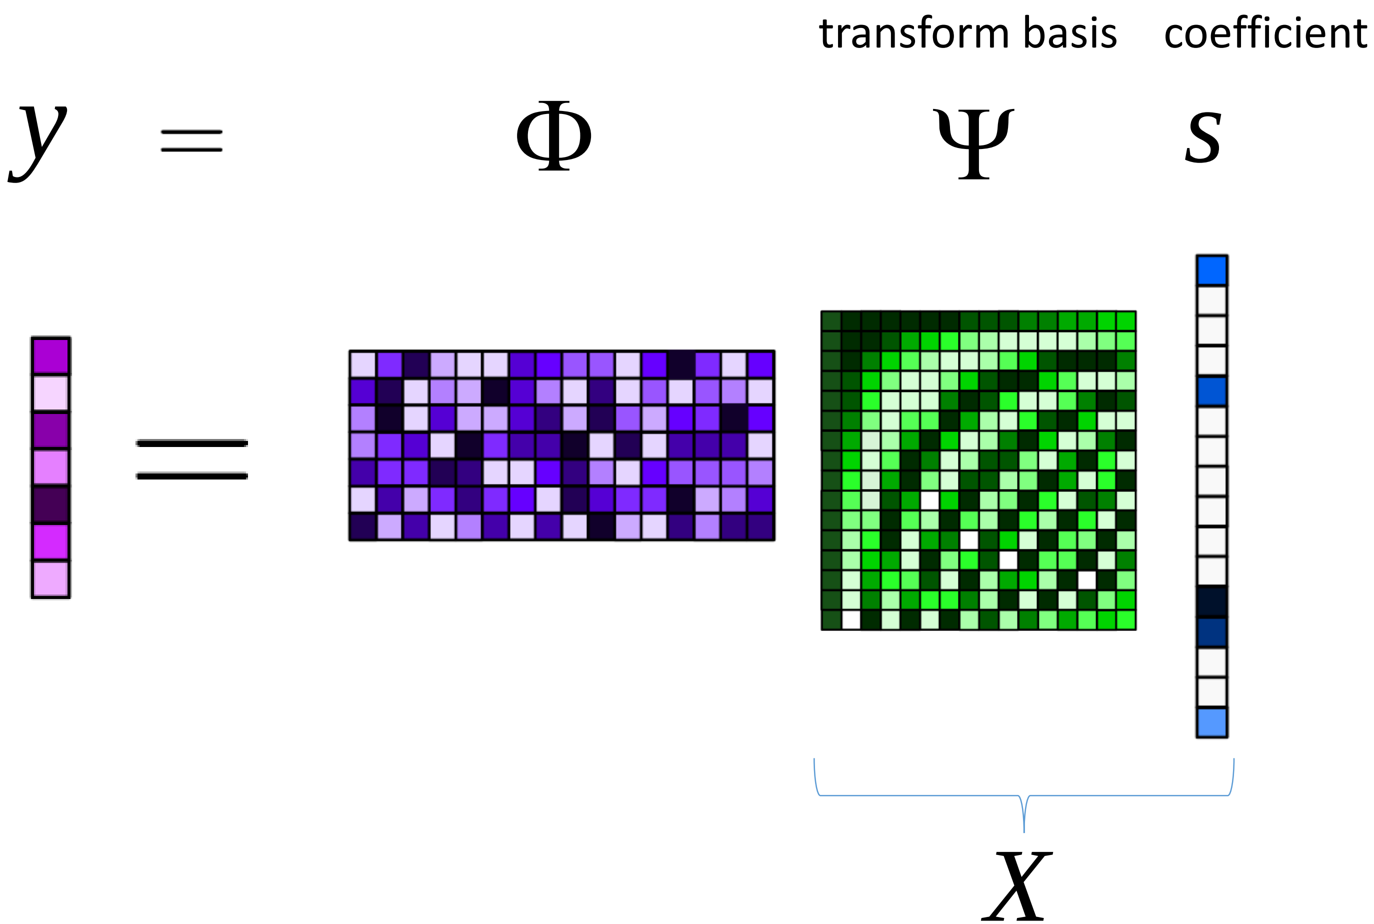

Recovery of non-sparse signals

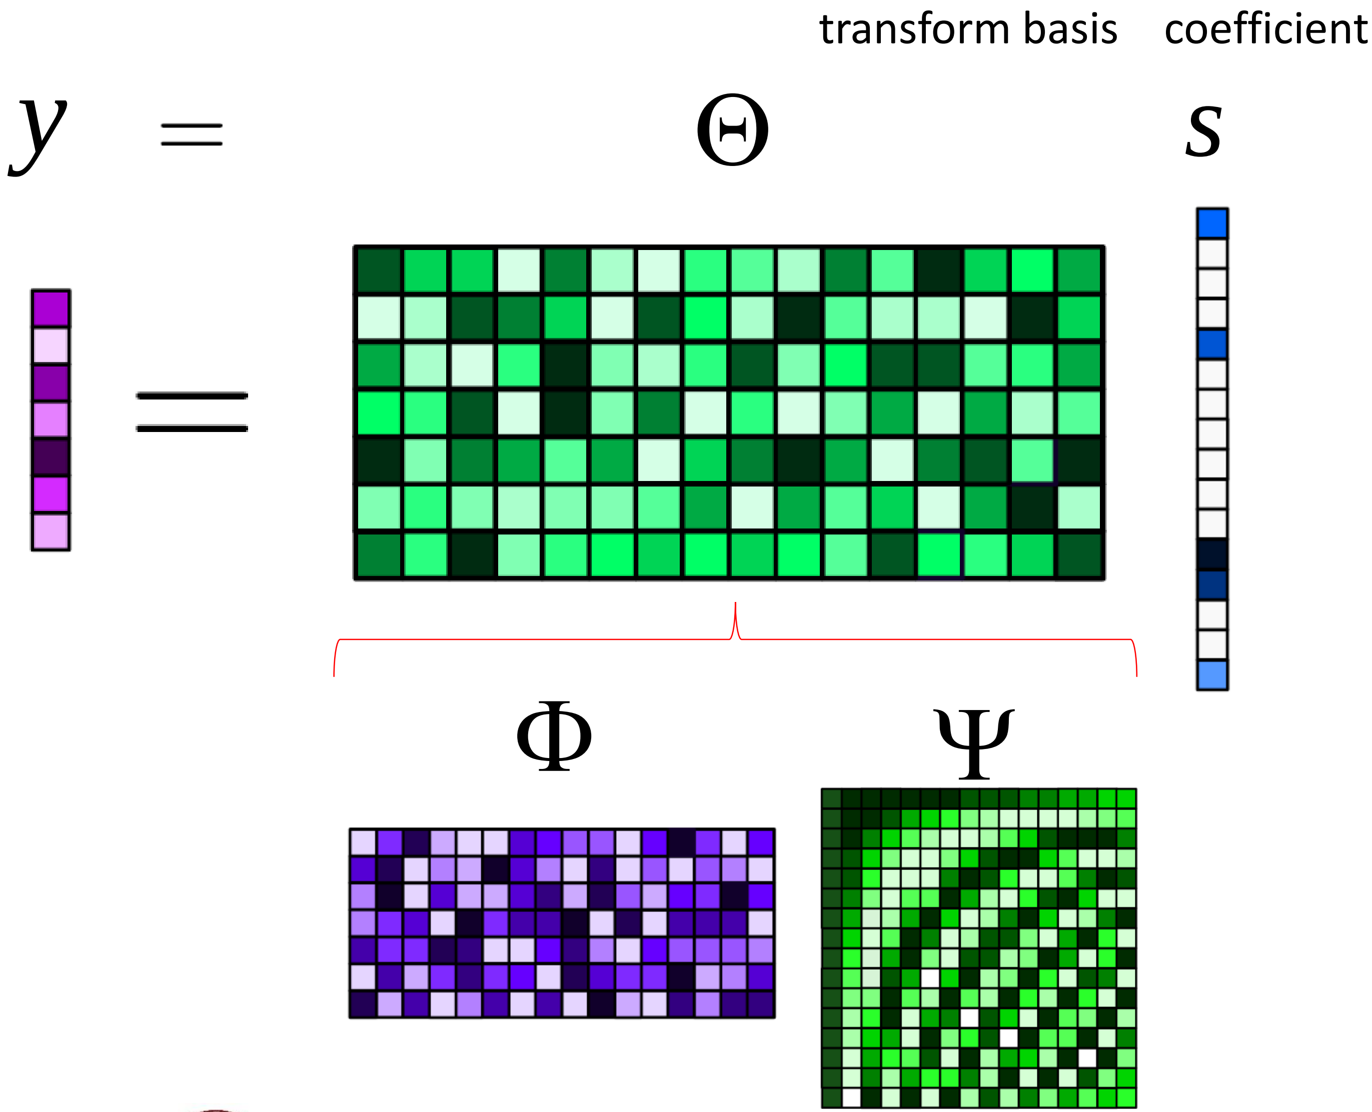

Reconstruction

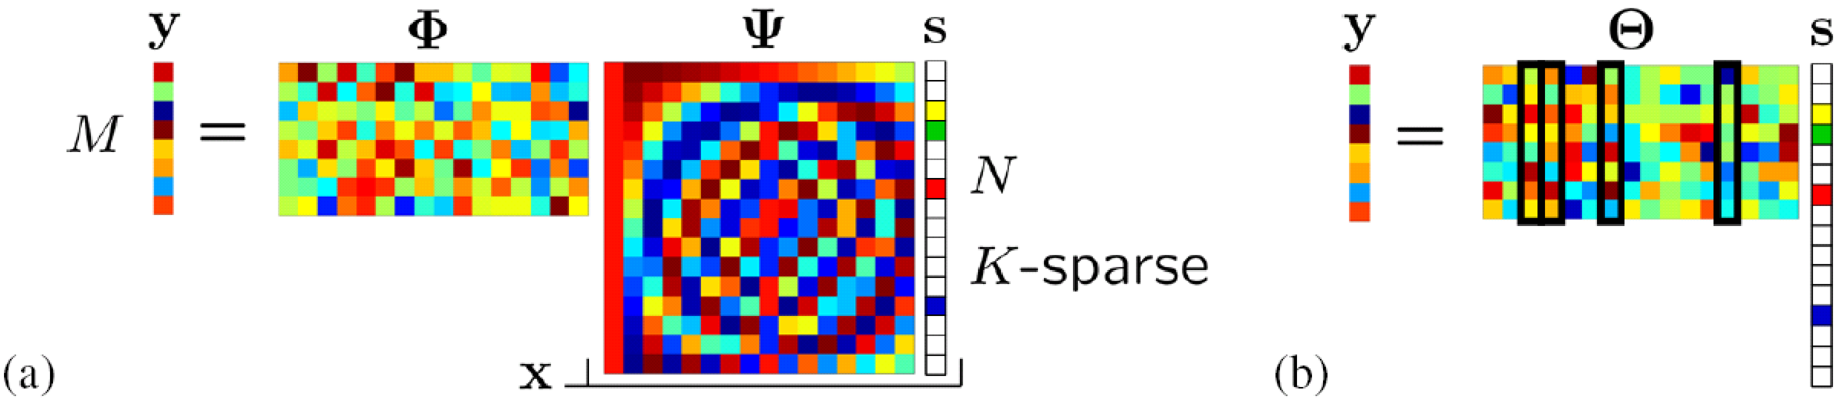

Definition:

$\Theta = \Phi\Psi$ satisfies the Restricted Isometry Property (RIP) of order K iff there exists some constant $\delta_K < 1$ such that $(1-\delta_K)\,||\boldsymbol{\alpha_K}||_2^2 \leq ||\Theta\boldsymbol{\alpha_K}||_2^2 \leq (1+\delta_K)\,||\boldsymbol{\alpha_K}||_2^2$ for all K-sparse vectors $\boldsymbol{\alpha_K}$.

Recovery from CS measurements

Theorem:
If the RIP is satisfied with $\delta_{2K} < \dfrac{2}{3+\sqrt{7/4}} \simeq 0.4627$, then for all signals,

the BP minimization problem $\min_{\bar{\boldsymbol{\alpha}} \in \mathbb{C}^N} ||\bar{\boldsymbol{\alpha}}||_1$ subject to $||\boldsymbol{y} - \Theta\bar{\boldsymbol{\alpha}}||_2 \leq \epsilon$

provides accurate and stable reconstruction $\boldsymbol{\alpha}^*$: $||\boldsymbol{\alpha} - \boldsymbol{\alpha}^*||_2 \leq c\epsilon + \dfrac{d}{\sqrt{K}}||\boldsymbol{\alpha} - \boldsymbol{\alpha_K}||_1$,

for ϵ such that $||\boldsymbol{n}||_2 \leq \epsilon$, $\boldsymbol{\alpha_K}$ the best K-sparse approximation of the signal in the l1 sense, and c and d are constants depending only on δ_{2K} .

Single-Pixel CS Camera [Baraniuk and Kelly, et al.]

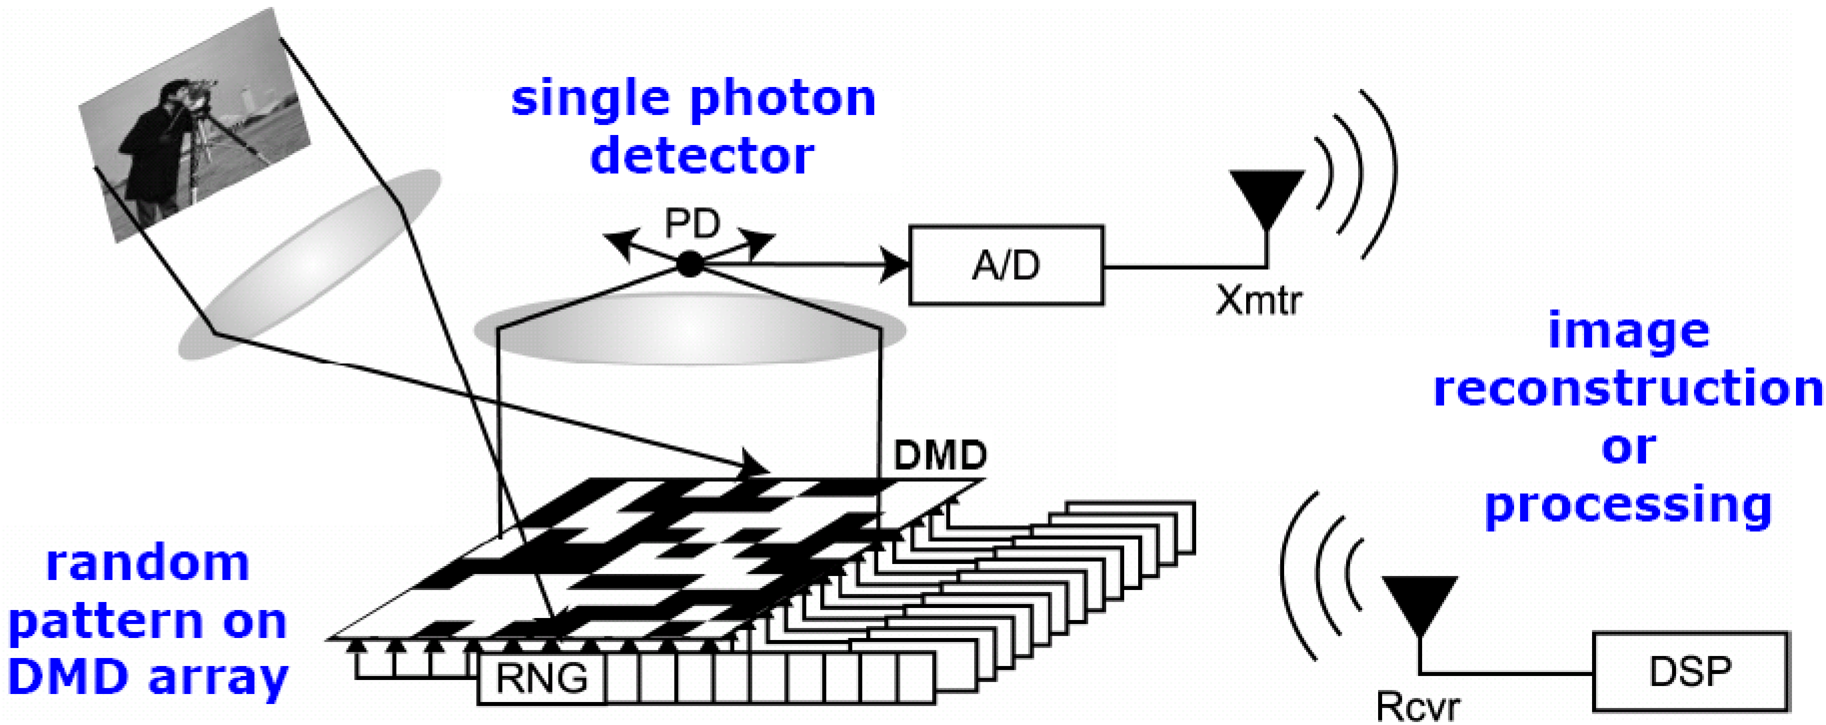

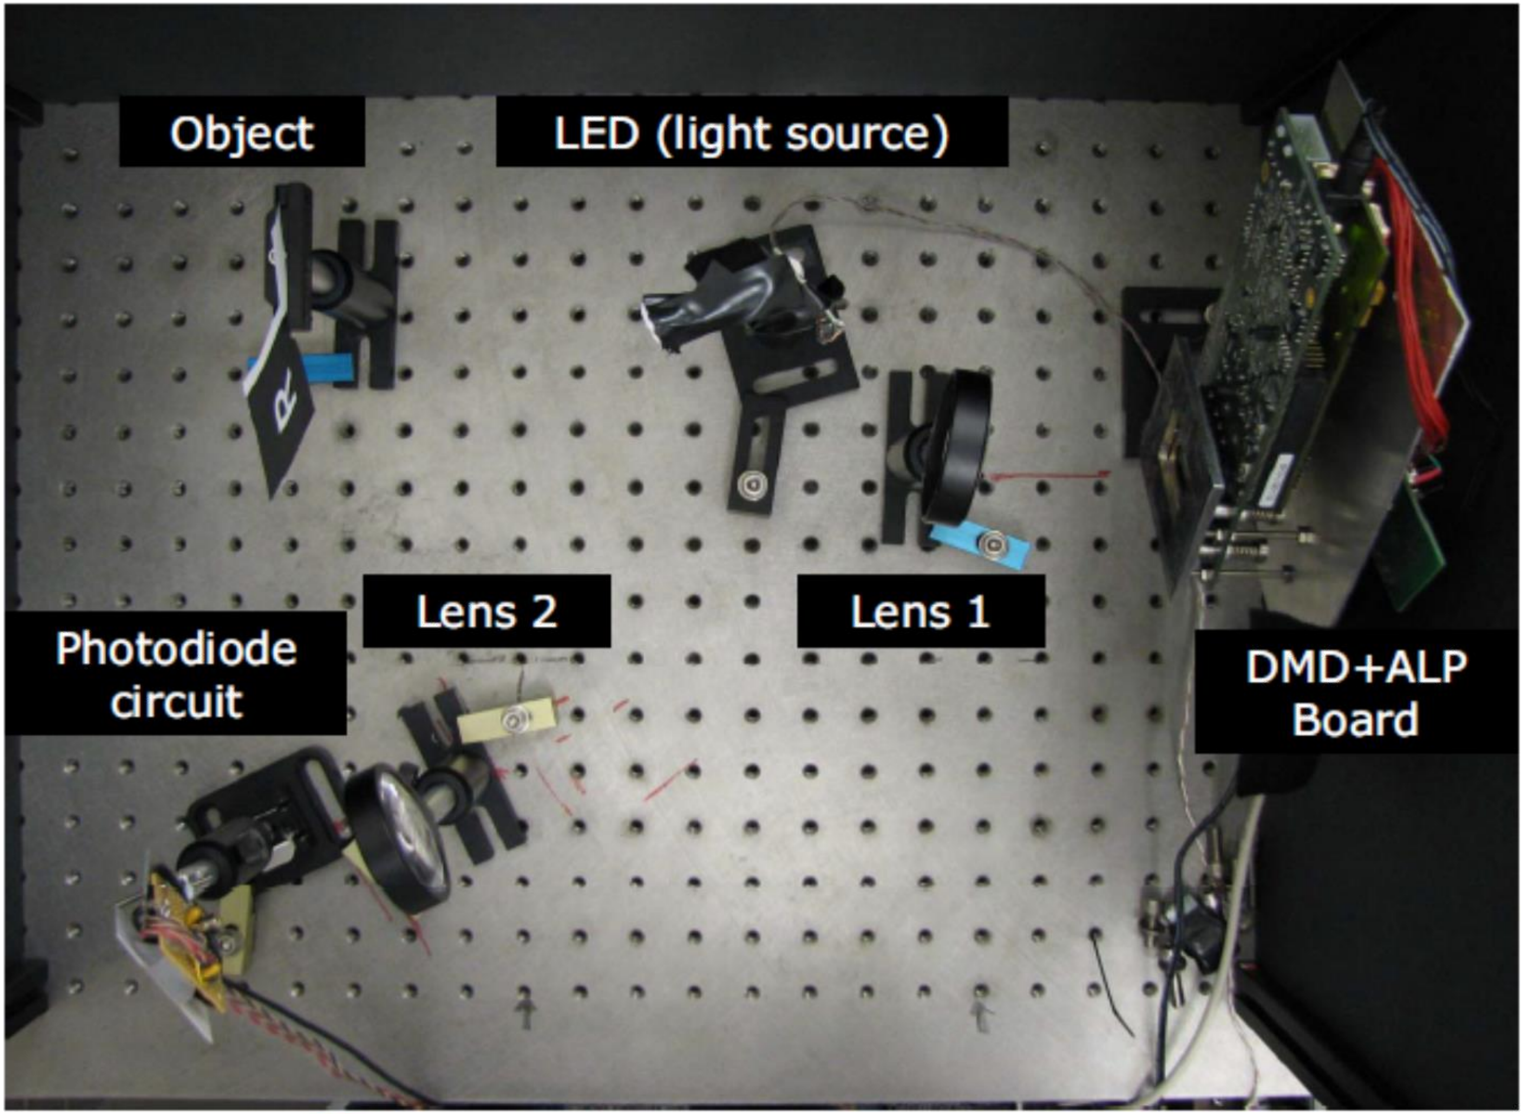
Object
LED (light source)
Lens 2
Lens 1
Photodiode circuit
DMD+ALP Board

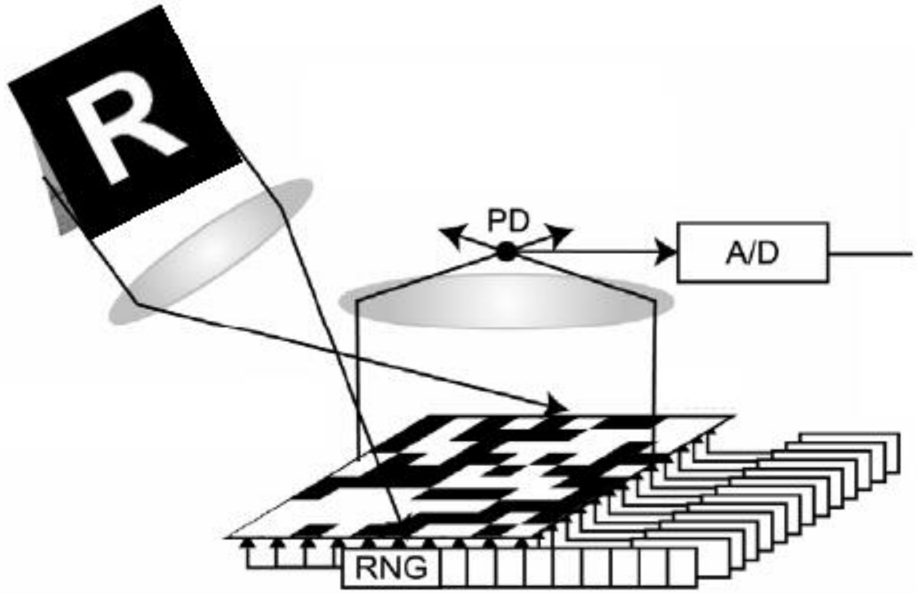

target
65536 pixels

11000 measurements
(16%)

1300 measurements
(2%)

4096
pixels

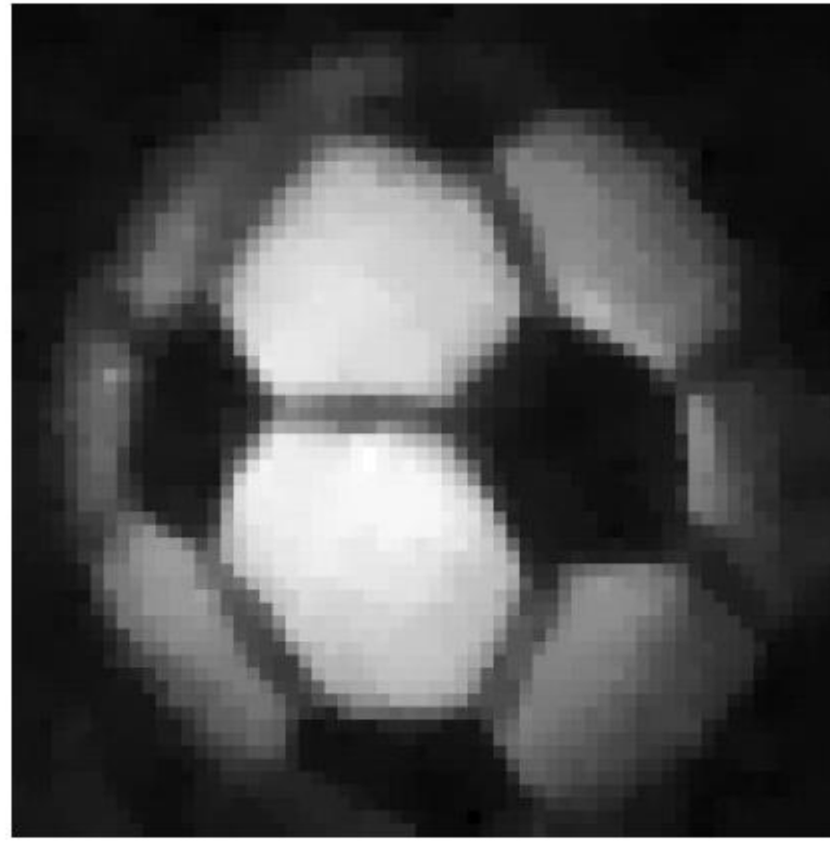

500
random measurements

Background subtraction

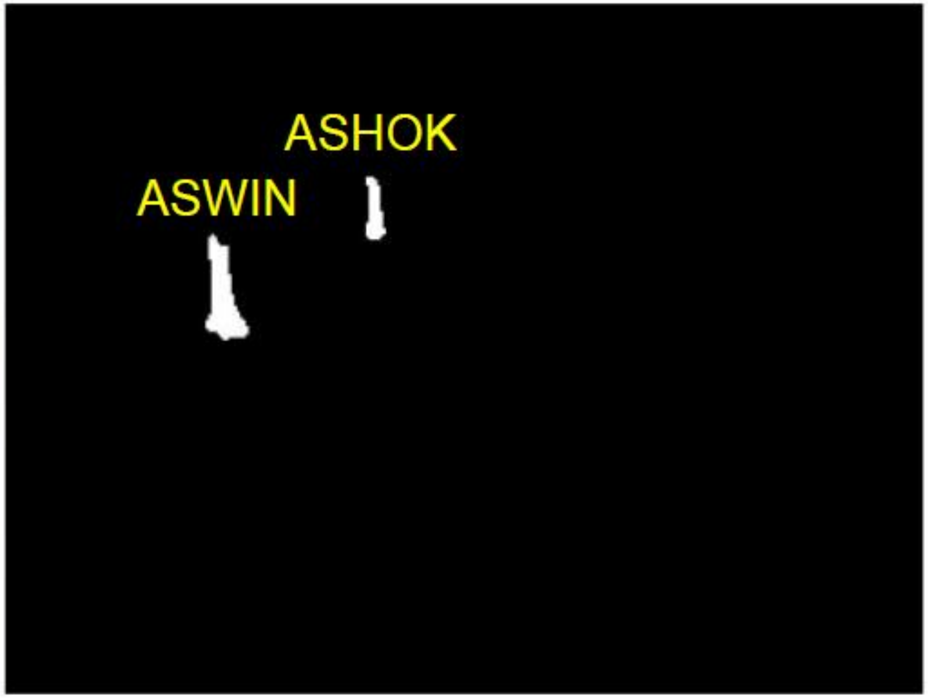

Range Imaging

Objective: Generate 2D depth map

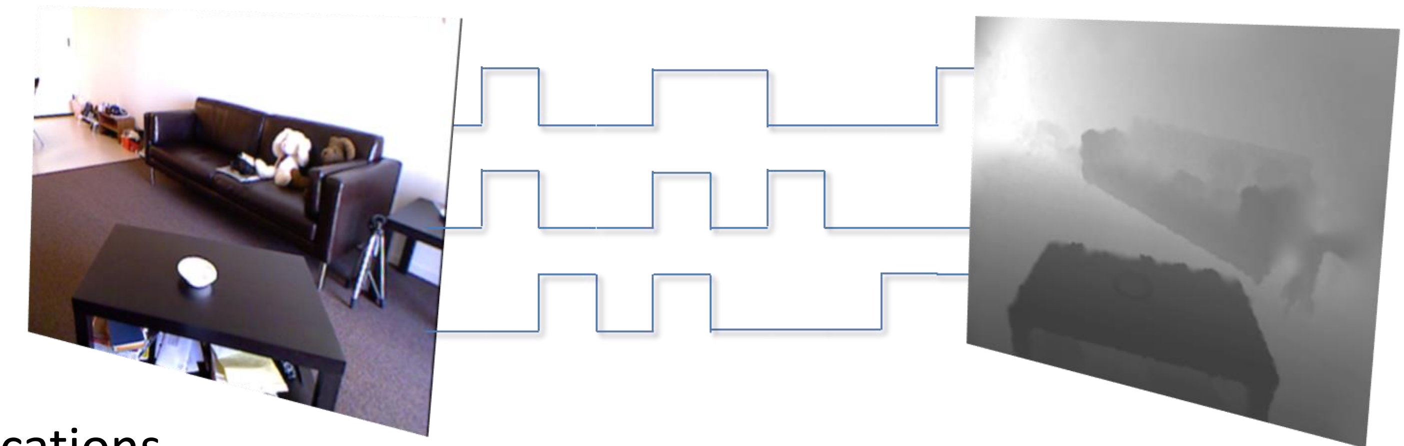

Applications

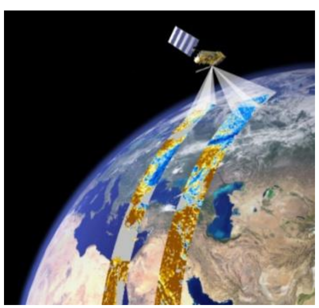

Remote Sensing

Human-Computer Interface

Consumer photography

ToF Cameras

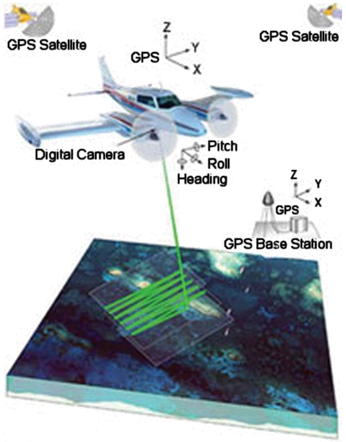

- Constant speed of light/sound/RF etc
- LIght Detection And Ranging (LIDAR)
- Characteristics
  - ➢Direct measurement
  - ➢Limited influence on the environment
  - ➢High accuracy time measurements
  - ➢**Large number of frame necessary**

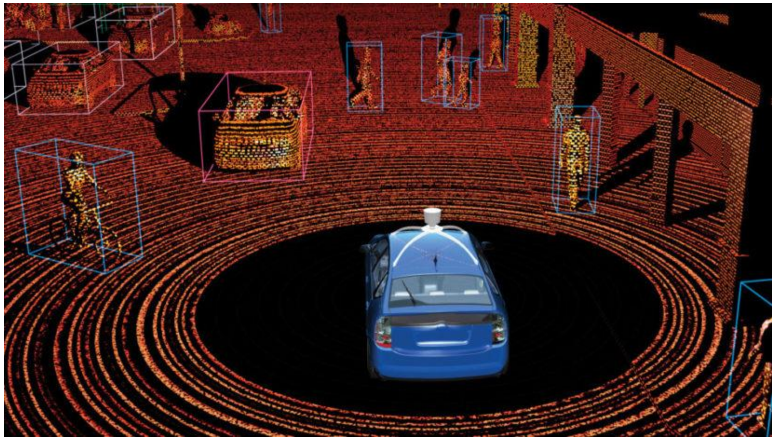

Time-of-Flight principle

- Range Gated Techniques
- Characteristics
  - Direct measurement
  - Limited influence on the environment
  - High accuracy time measurements
  - **Large number of frame necessary**

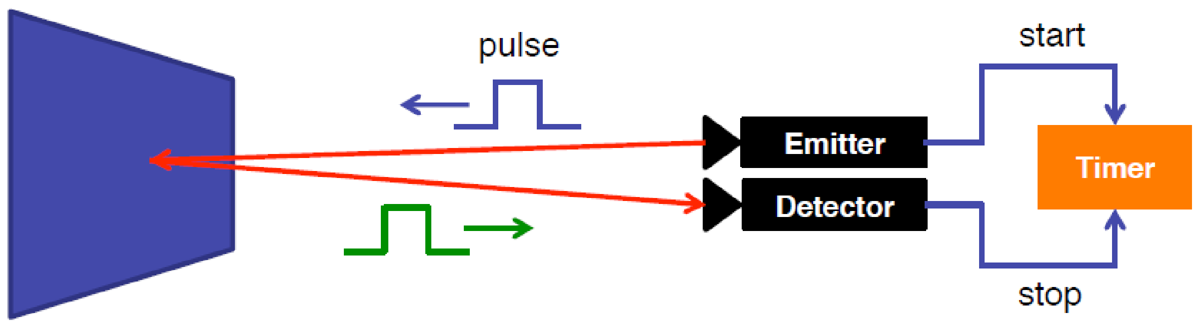

Gated Range Imaging

Process	Techniques
➢Emit light	• Time Slicing - Classical
➢Encode reflection	• Gate Coding - Deterministic
➢Extract depth	• Compressed Sensing

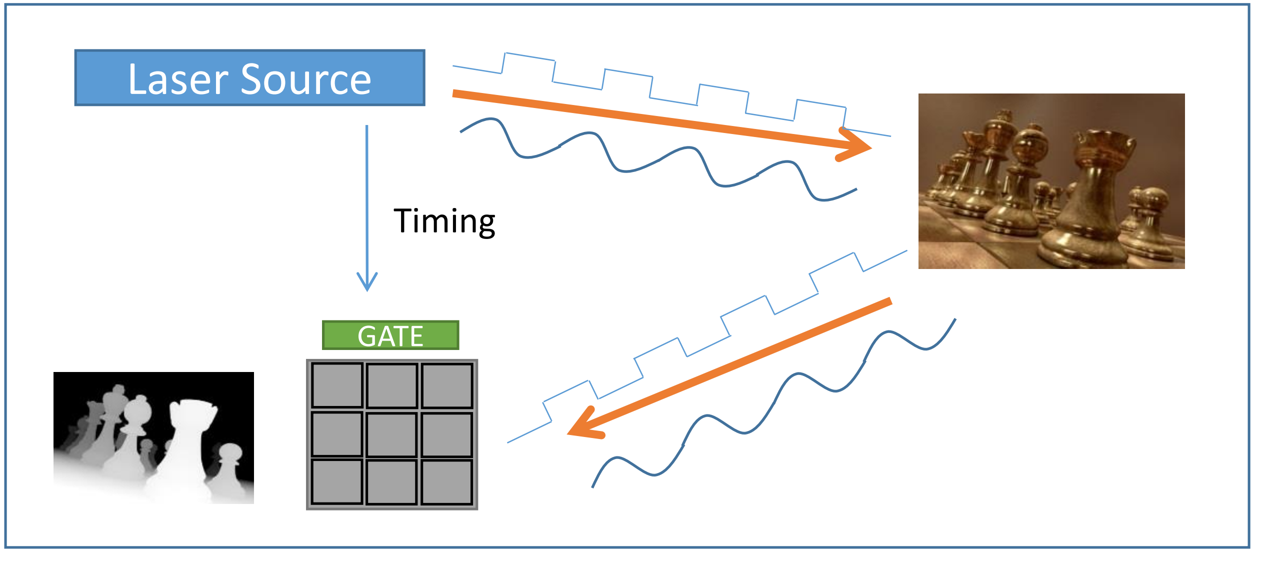

Active Range Imaging

Time Slicing

- Single depth per image
- Baseline approach
- Single object

$$Z = max(m_0, m_1, \ldots, m_k)$$

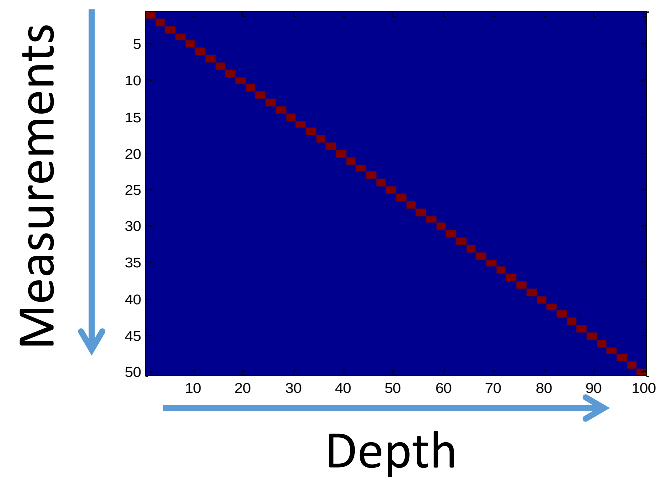

Gate Coding

- Exploit pulse profile
- Enhanced coding
- Single pulse acquisition

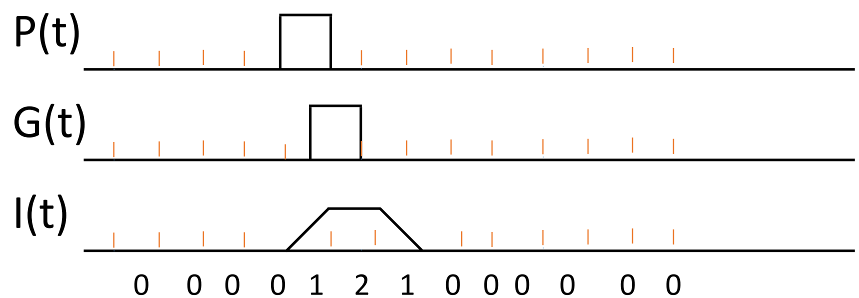

Compressed Gated Range Sensing

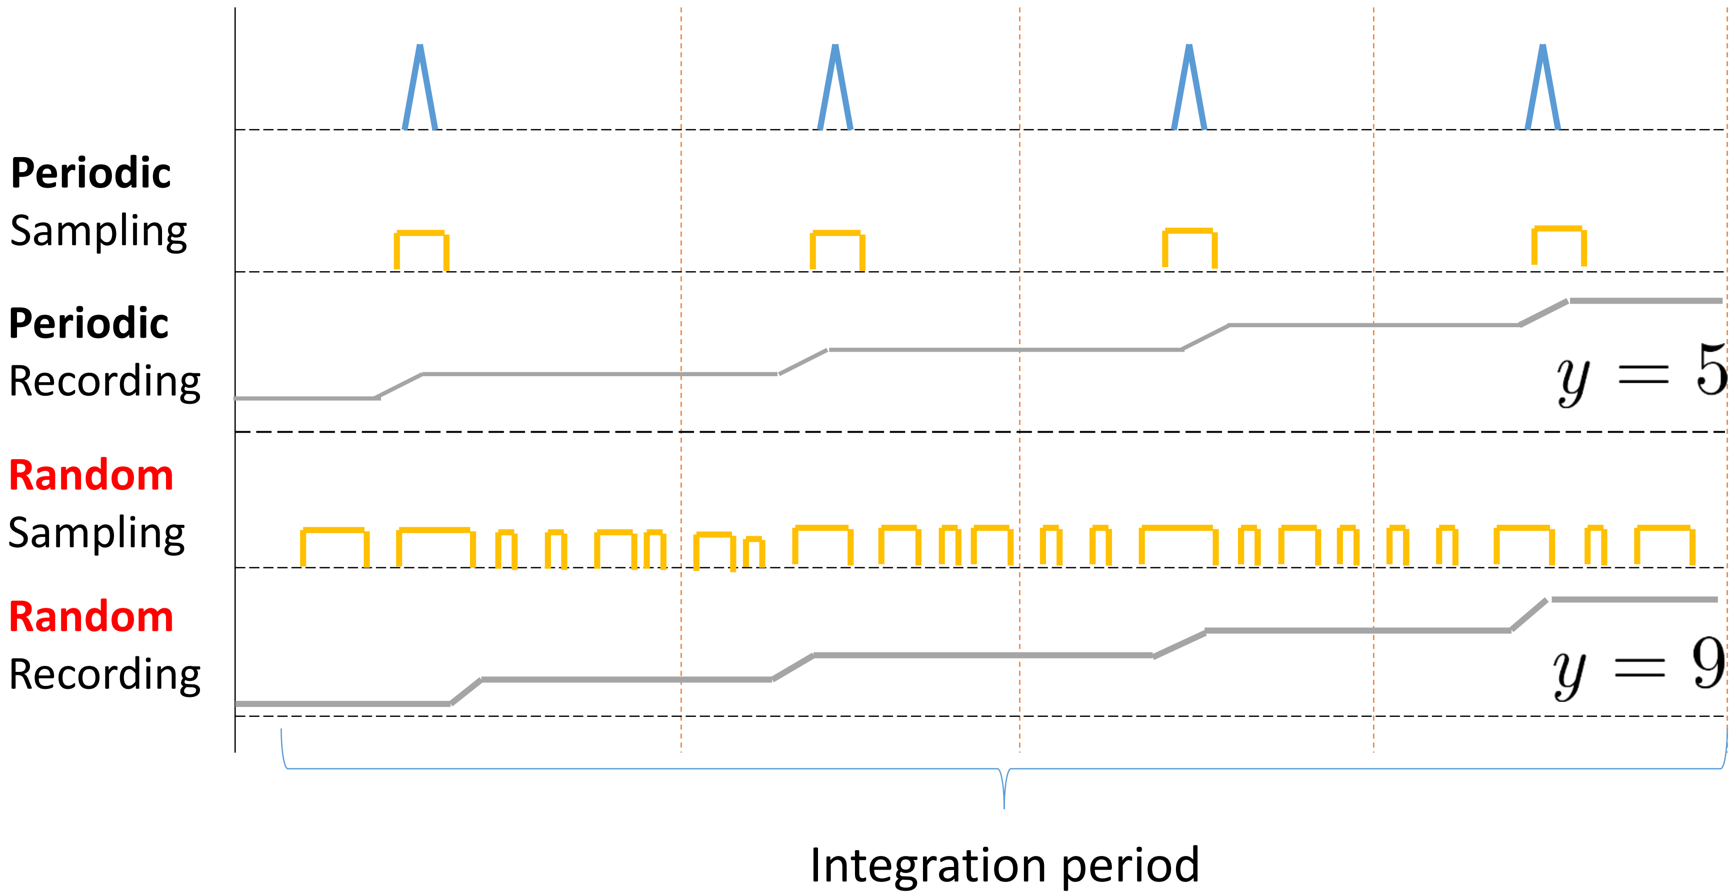

Compressed Gated Range Sensing

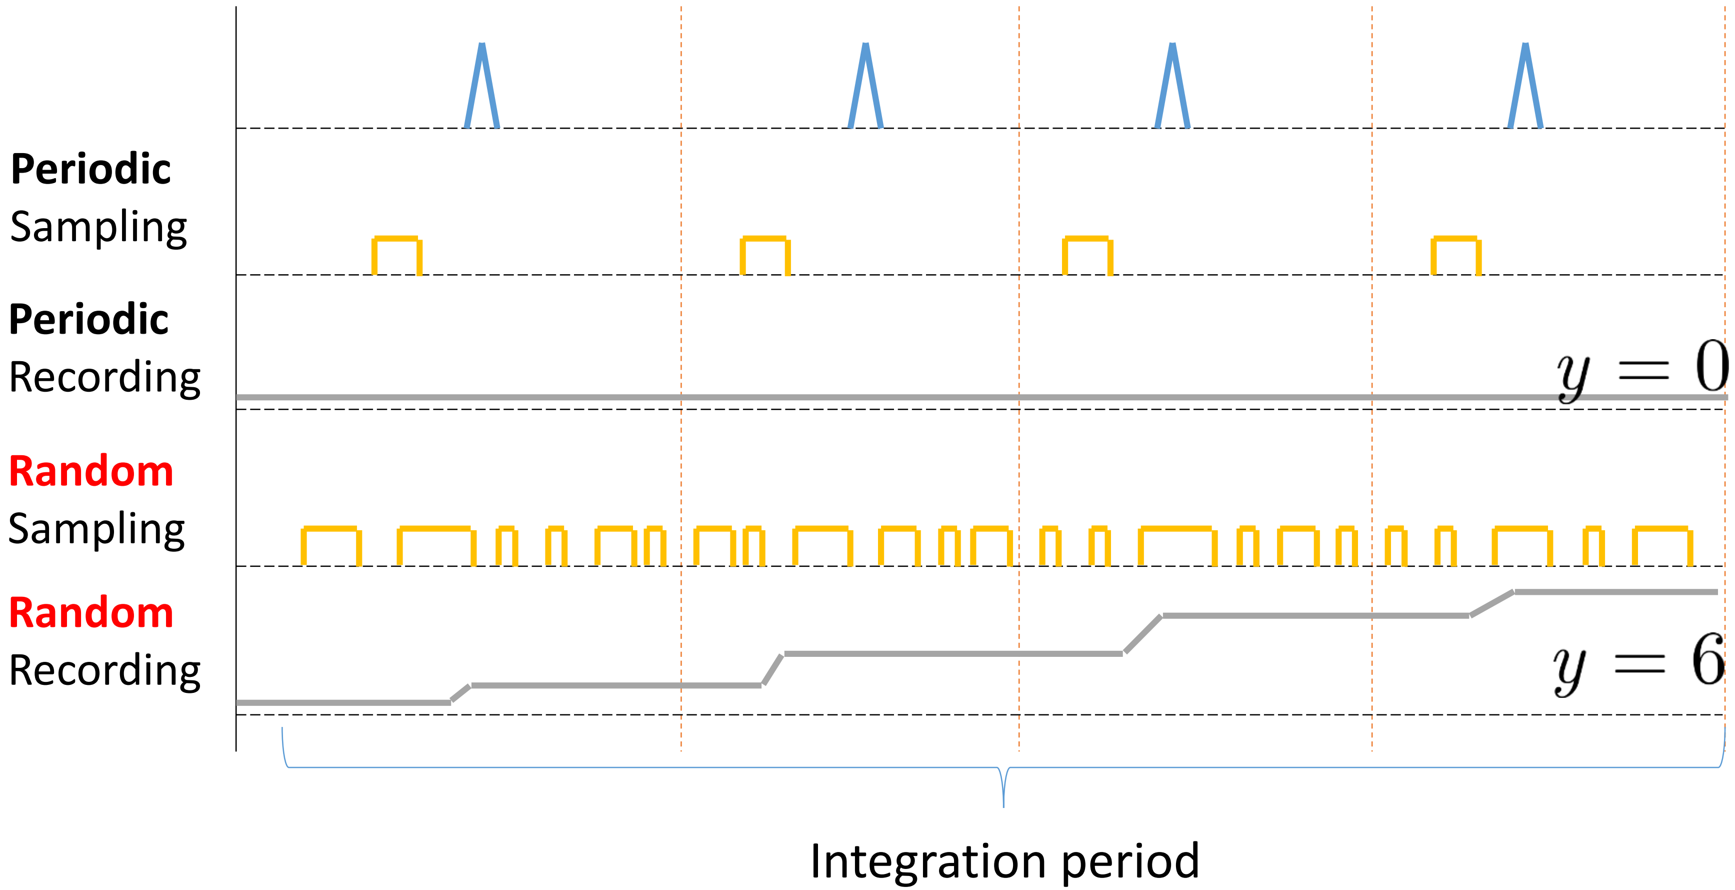

CGRS - overview

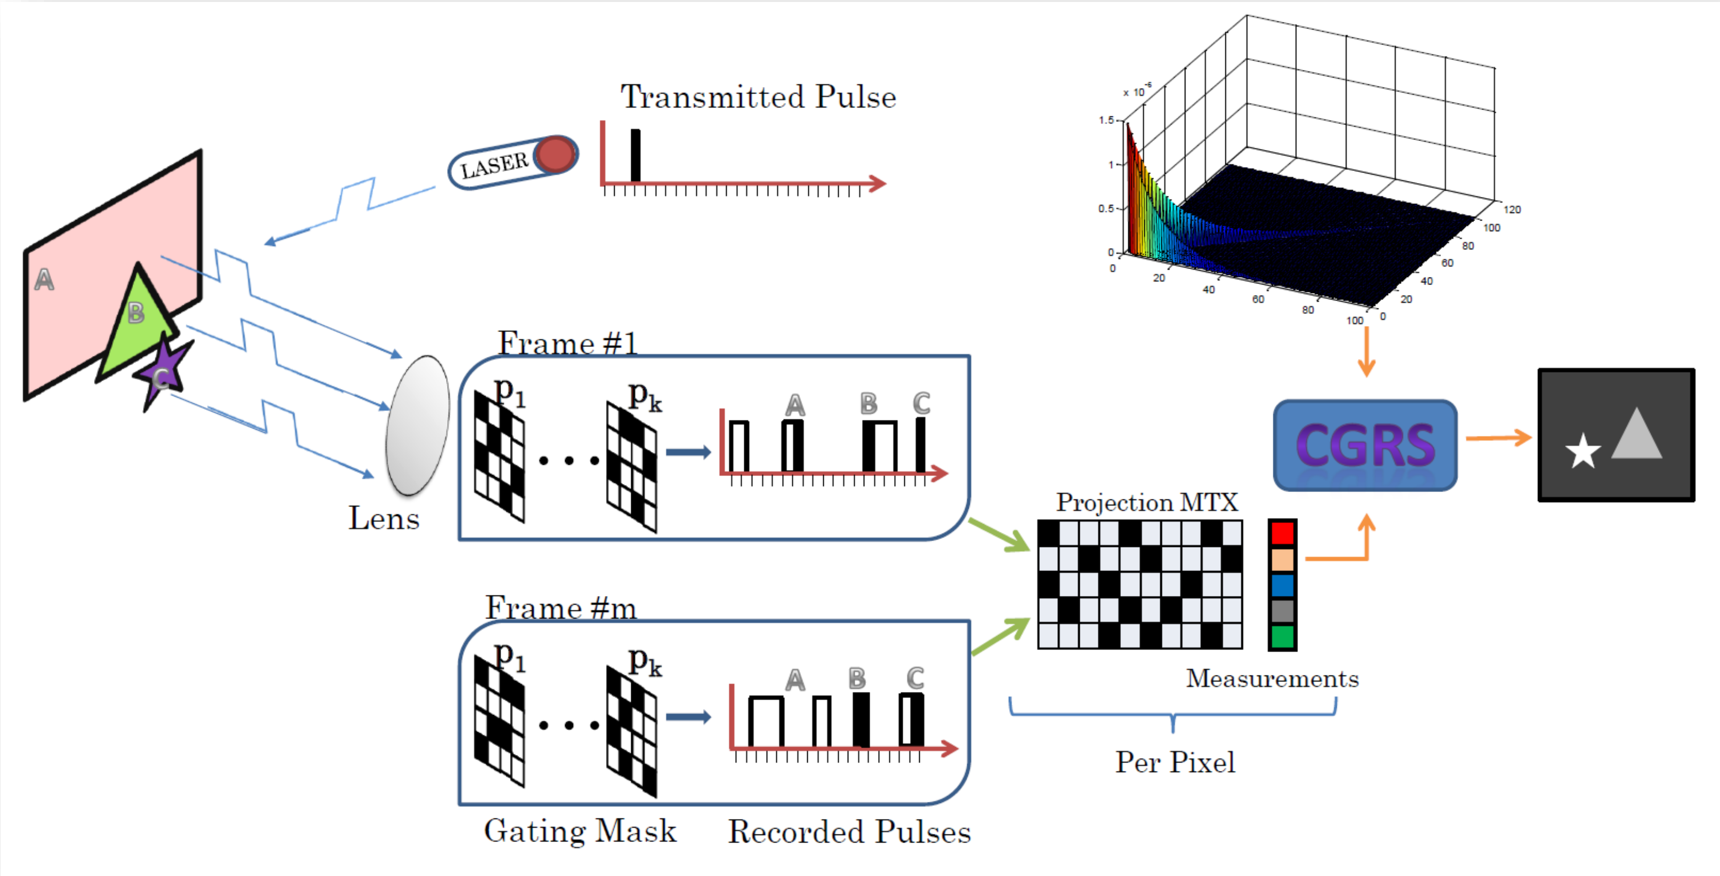

CGRS - details

Sparse (non-negative) depth signal recovery

$\mathbf{G}$ - Random sampling mechanism

$\mathbf{D}$ - Dictionary (backscatter, attenuation, non-ideal gating)

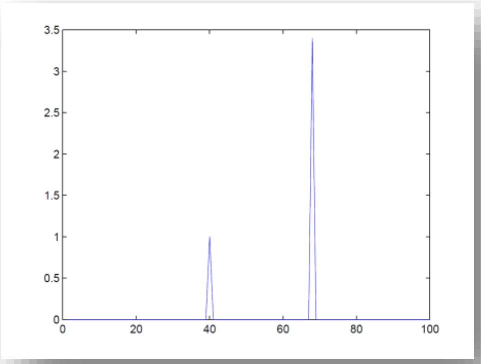

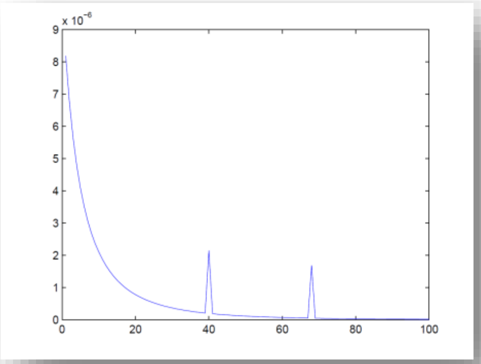

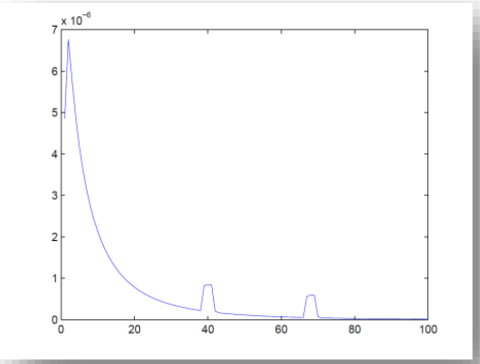

Optimization

$$\begin{aligned} &\min \|\mathbf{s}\|_1 \\ &\text{s.t.} \quad \|\mathbf{y} - \mathbf{GDs}\|_2 \leq \epsilon \\ &\qquad s \geq 0 \end{aligned}$$

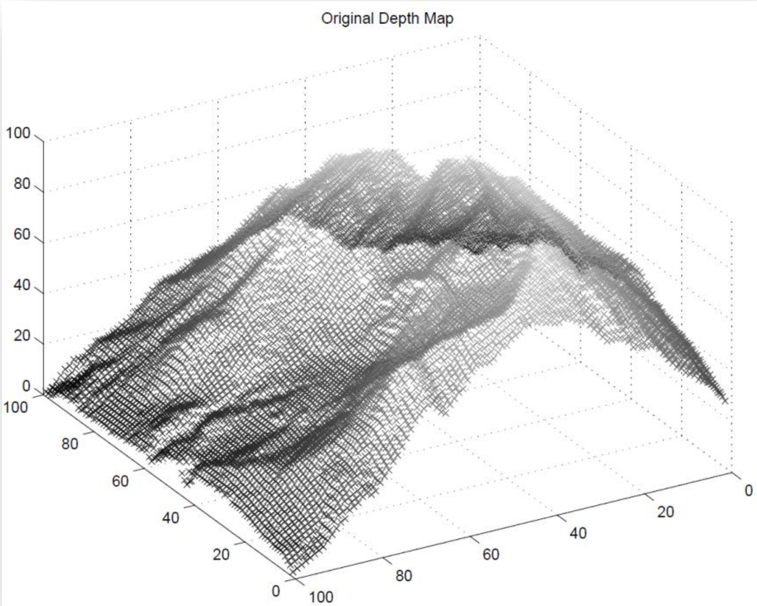
Original Depth Map
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

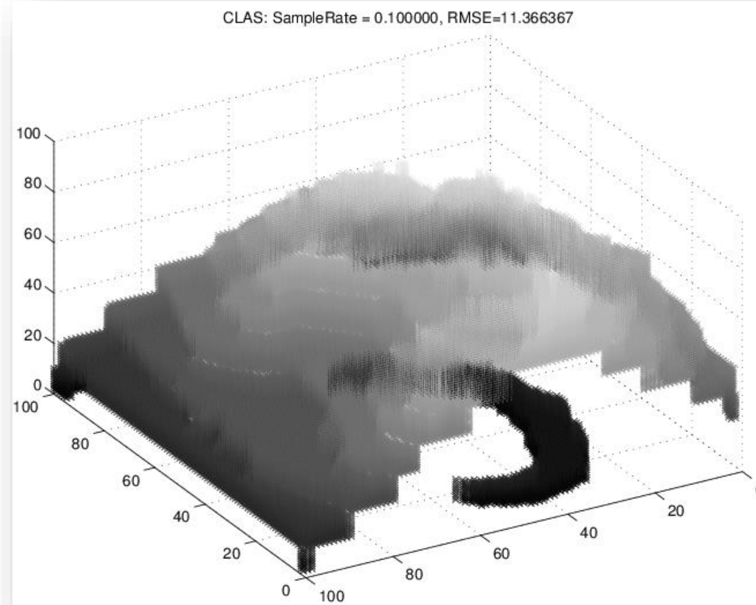
CLAS: SampleRate = 0.100000, RMSE=11.366367
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

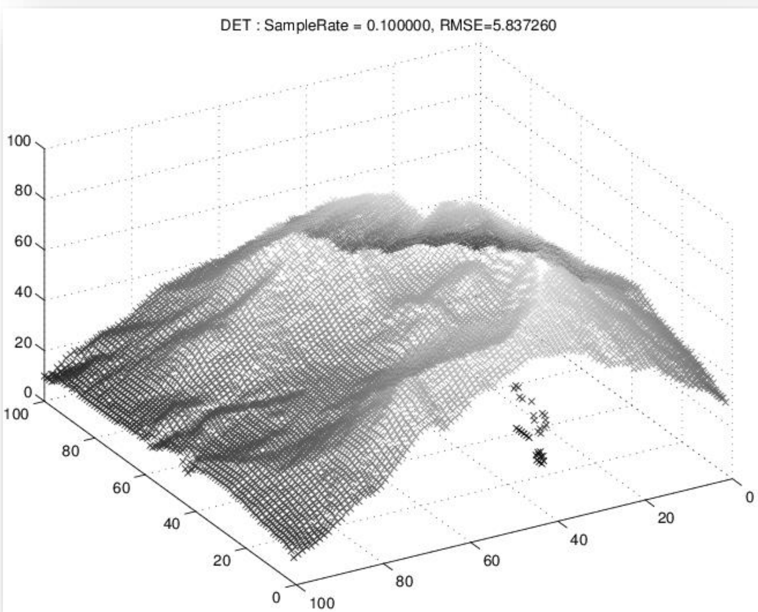
DET : SampleRate = 0.100000, RMSE=5.837260
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

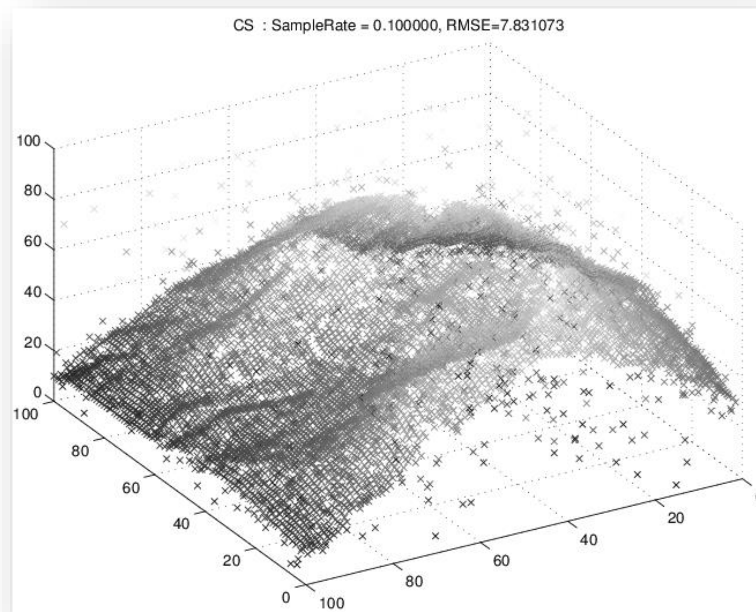
CS : SampleRate = 0.100000, RMSE=7.831073
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

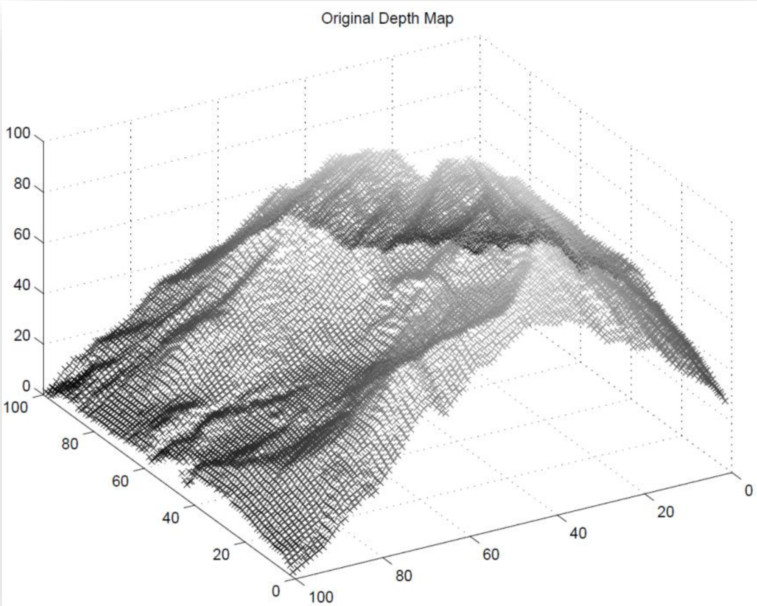
Original Depth Map

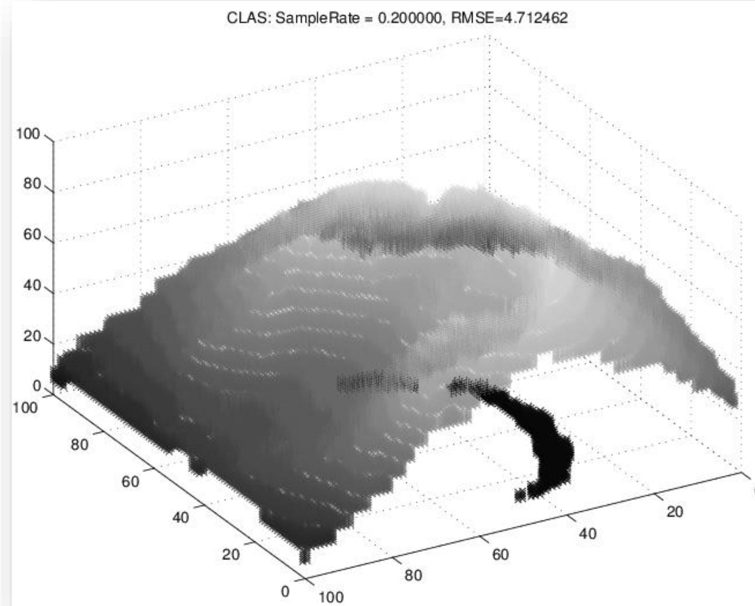
CLAS: SampleRate = 0.200000, RMSE=4.712462

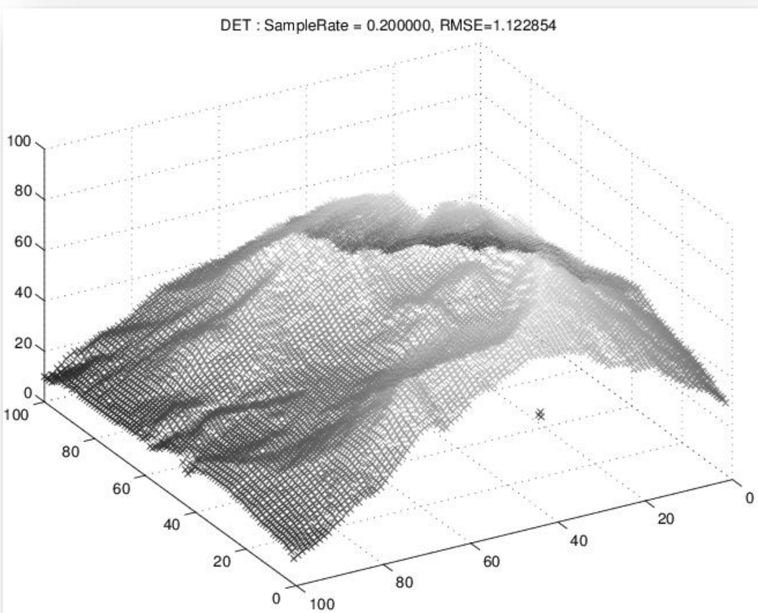
DET : SampleRate = 0.200000, RMSE=1.122854

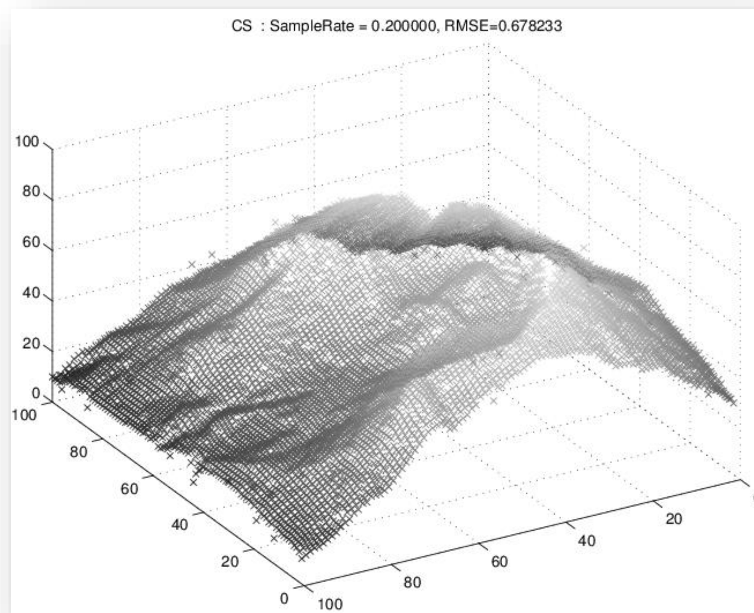
CS : SampleRate = 0.200000, RMSE=0.678233

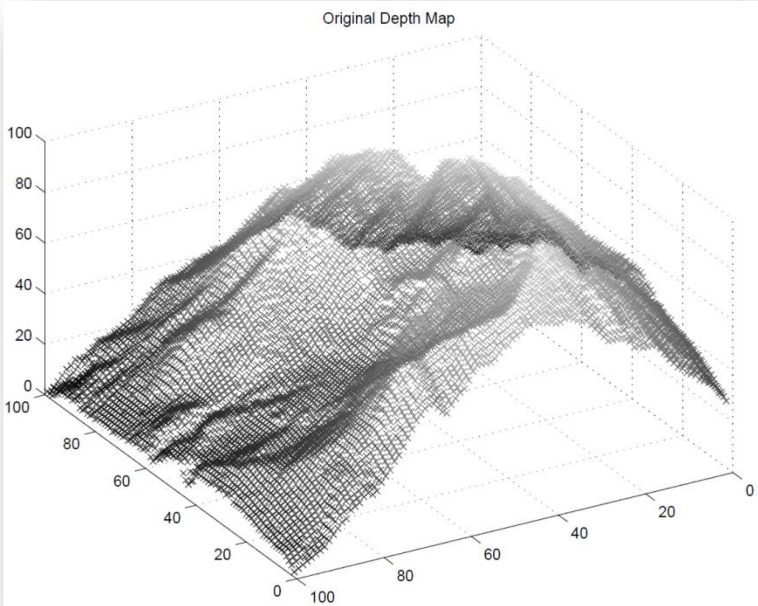
Original Depth Map

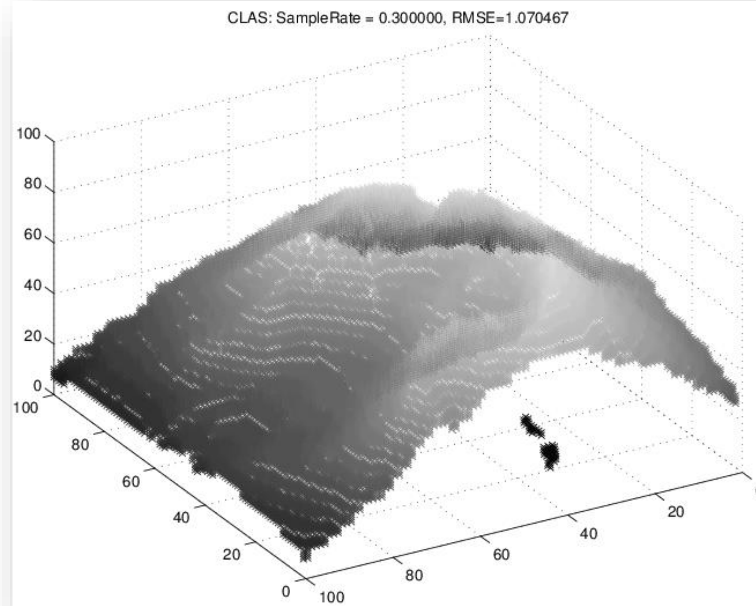
CLAS: SampleRate = 0.300000, RMSE=1.070467

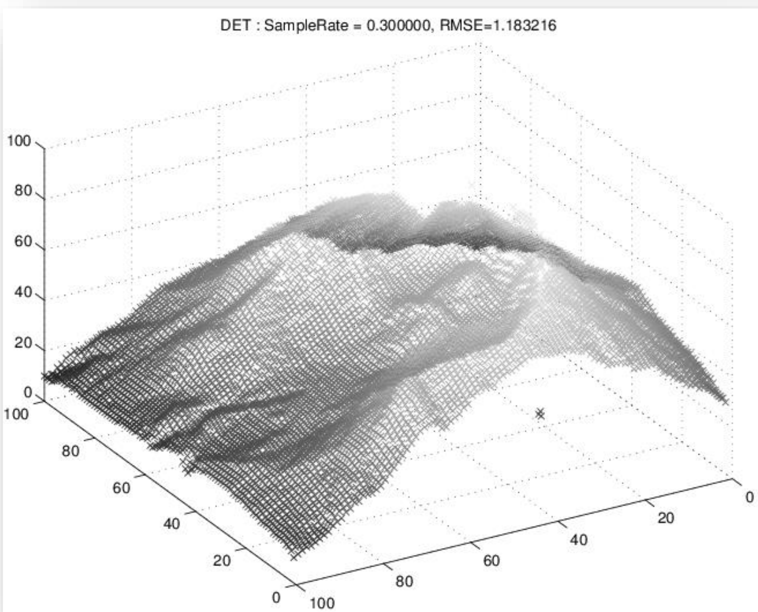
DET : SampleRate = 0.300000, RMSE=1.183216

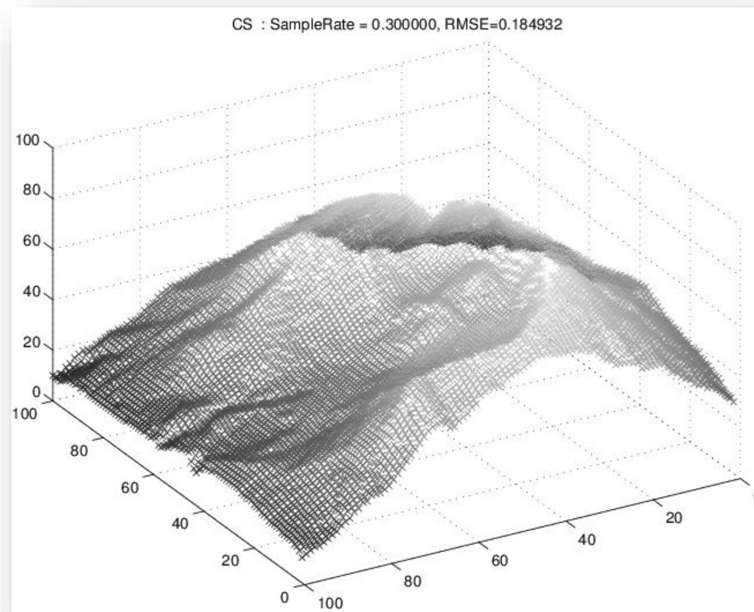
CS : SampleRate = 0.300000, RMSE=0.184932

Multi-Pulse Recovery

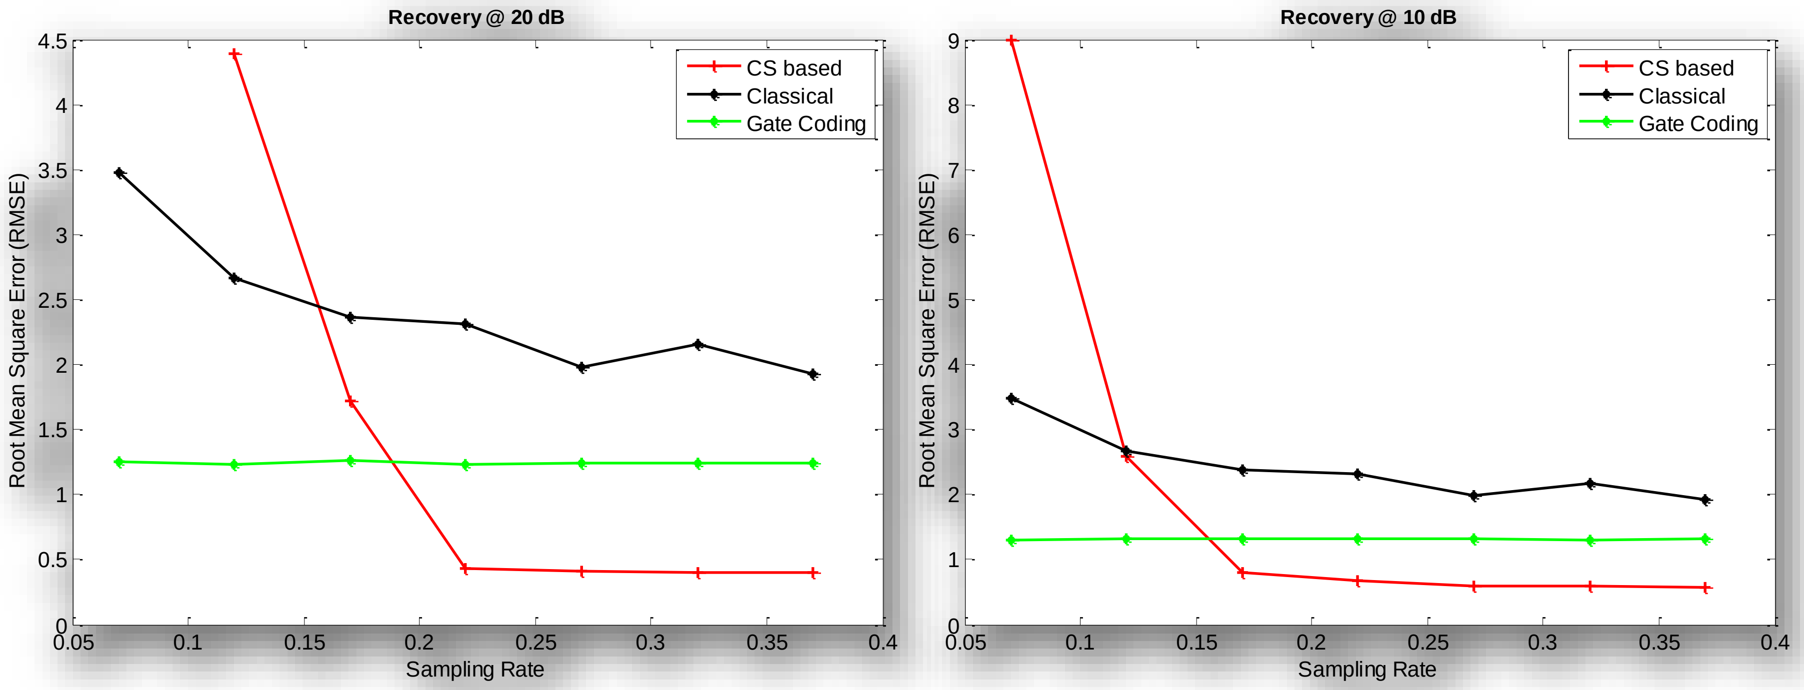

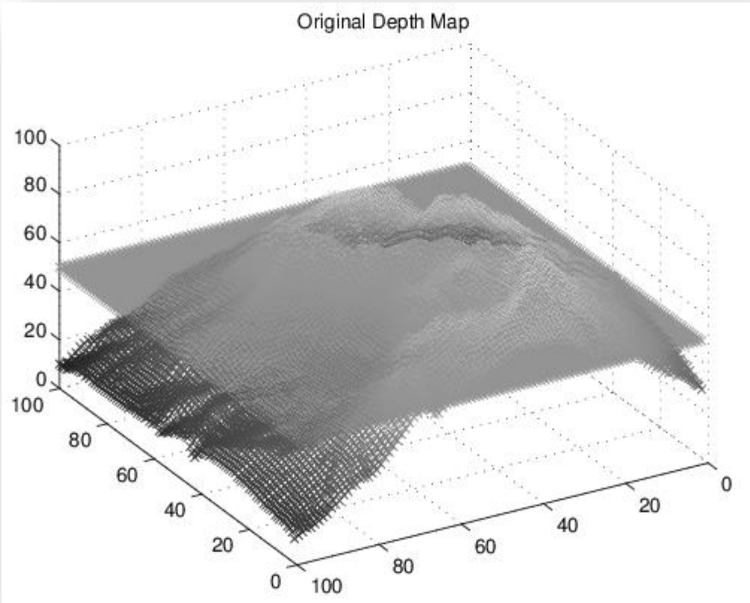
Original Depth Map

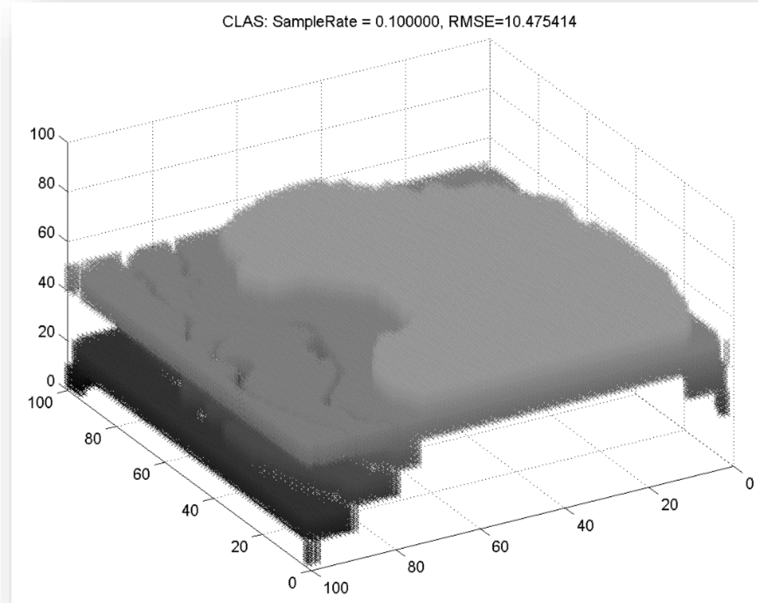
CLAS: SampleRate = 0.100000, RMSE=10.475414

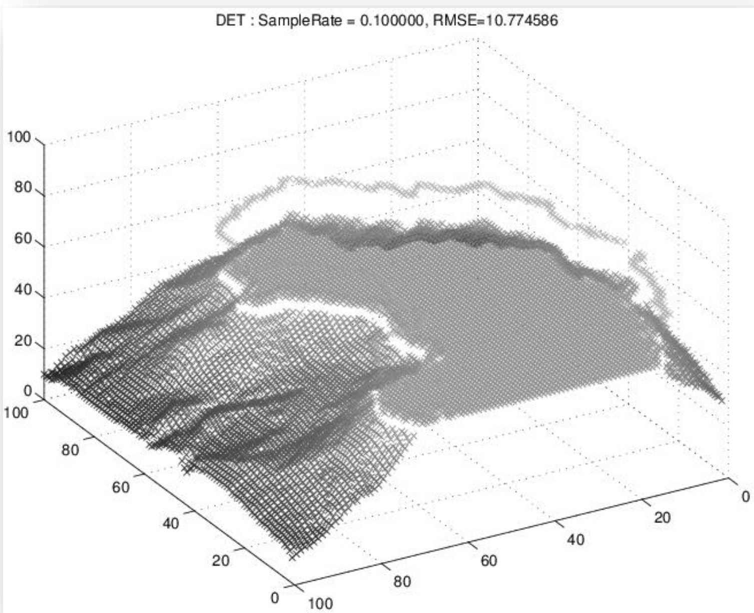
DET : SampleRate = 0.100000, RMSE=10.774586

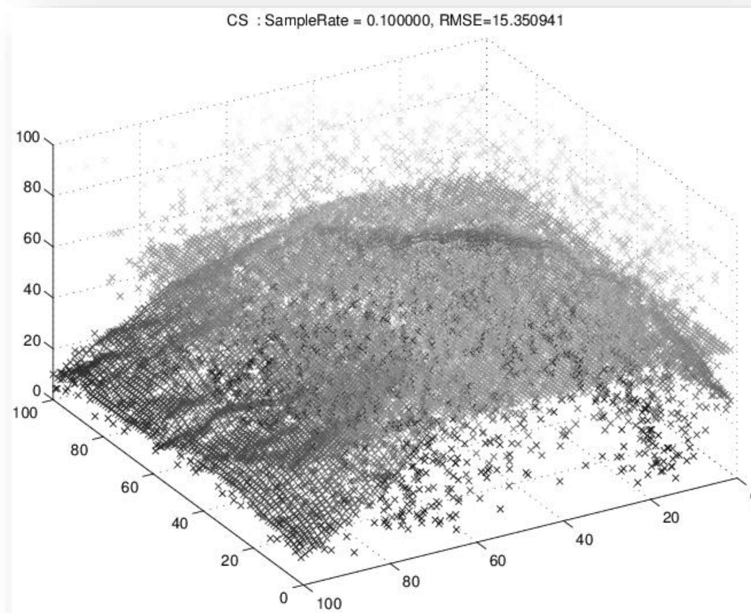
CS : SampleRate = 0.100000, RMSE=15.350941

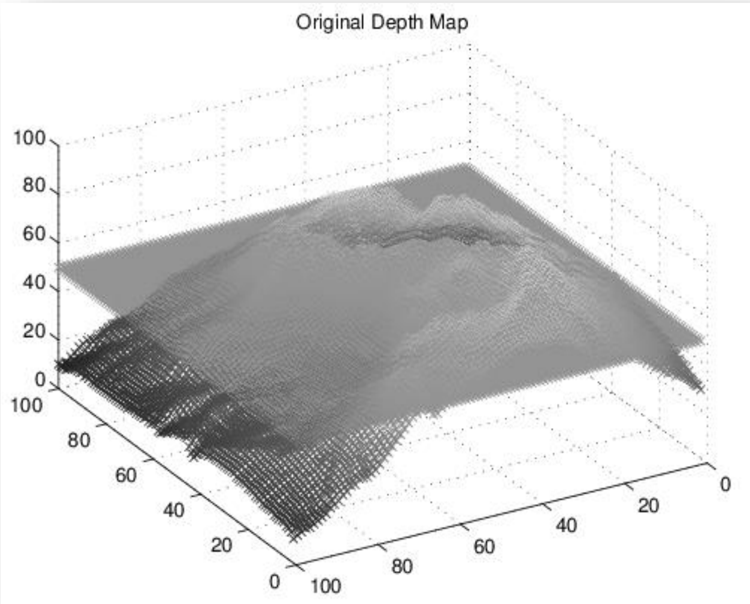
Original Depth Map

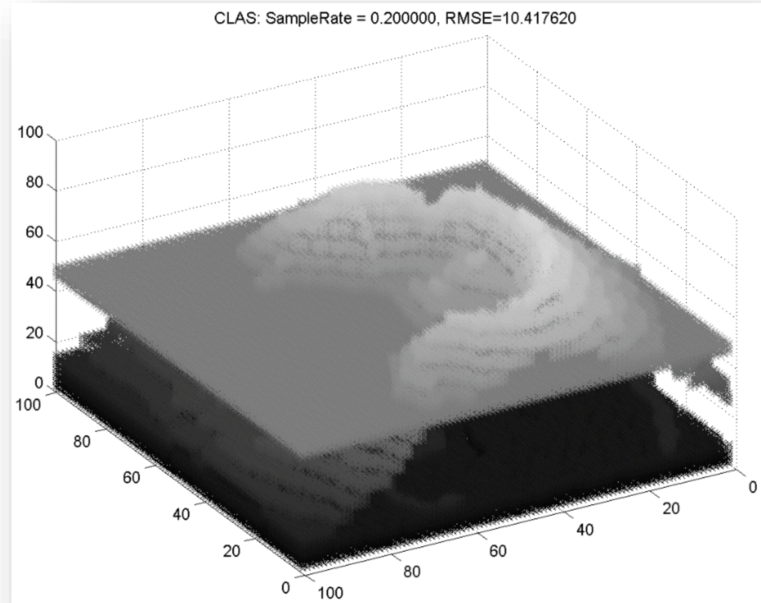
CLAS: SampleRate = 0.200000, RMSE=10.417620

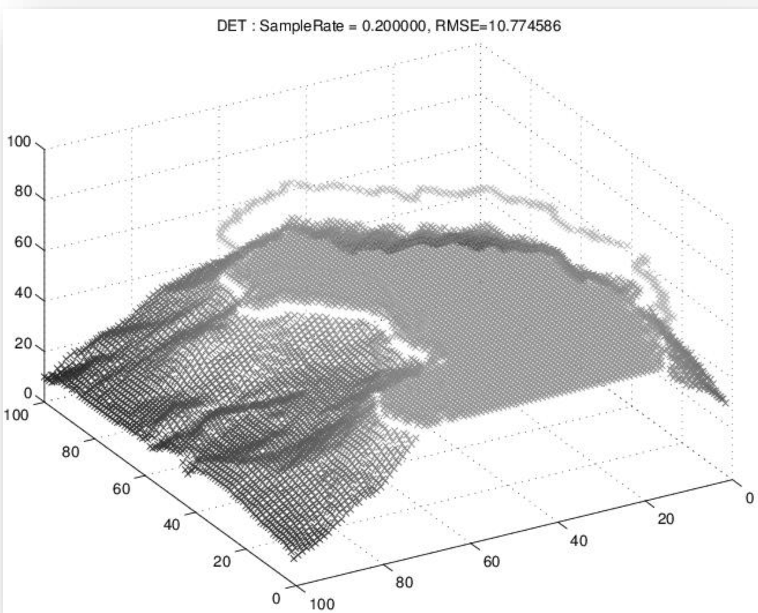
DET : SampleRate = 0.200000, RMSE=10.774586

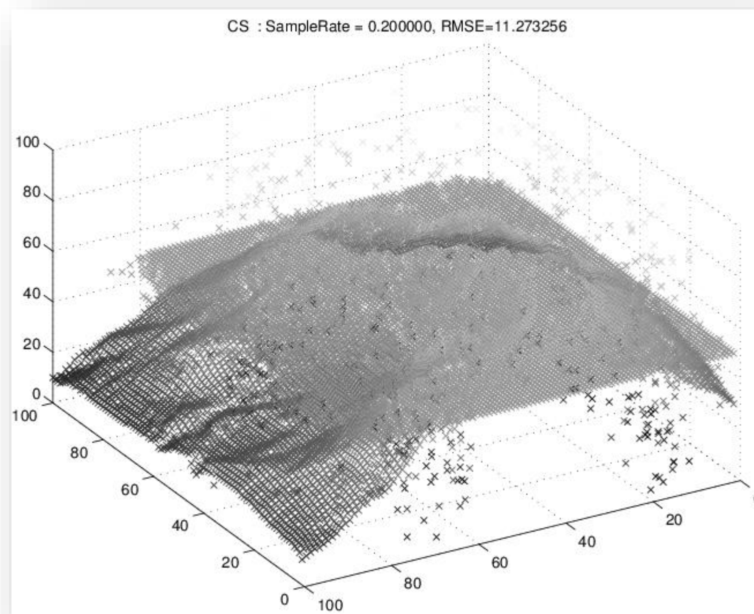
CS : SampleRate = 0.200000, RMSE=11.273256

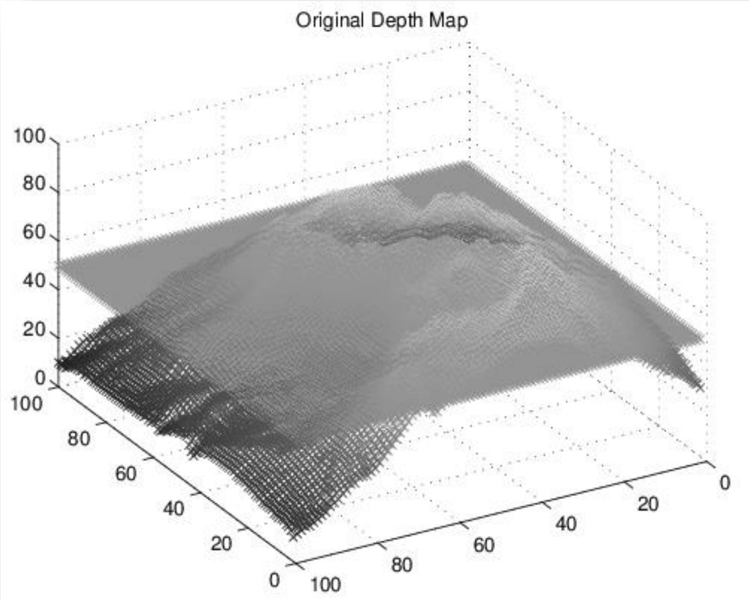
Original Depth Map
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

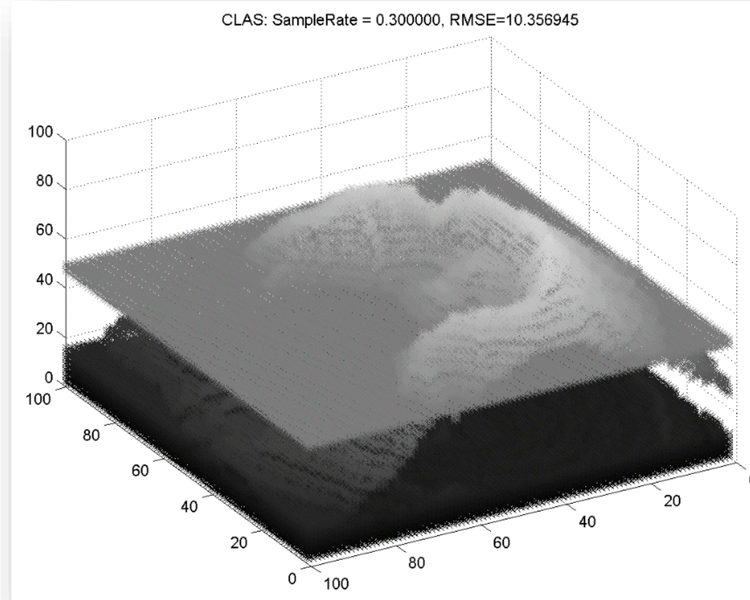
CLAS: SampleRate = 0.300000, RMSE=10.356945
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

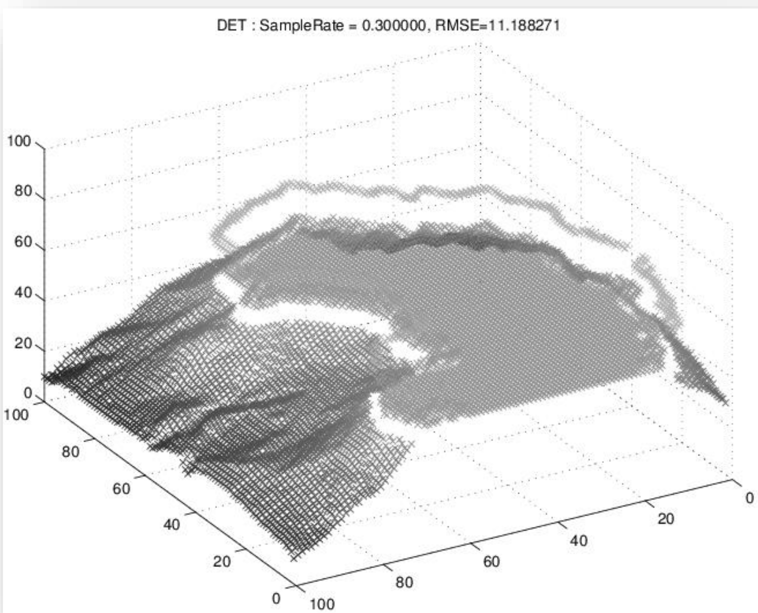
DET : SampleRate = 0.300000, RMSE=11.188271
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0

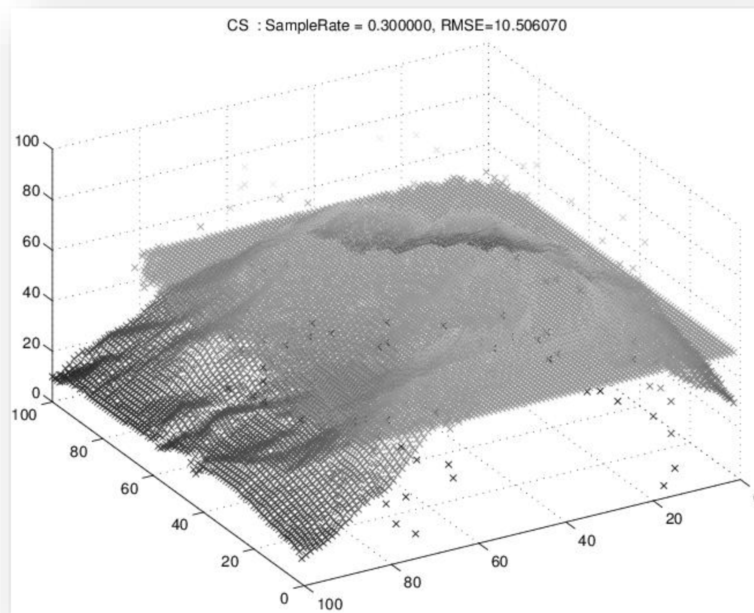
CS : SampleRate = 0.300000, RMSE=10.506070
100
80
60
40
20
0
100
80
60
40
20
0
100
80
60
40
20
0