Compressive sensing: how to sample data from what you know!

Laurent Jacques (ICTEAM/ELEN, UCL)

THE 35TH ANNUAL CONFERENCE OF THE EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS

EUROGRAPHICS 2014

Strasbourg, France

CONFERENCE 7-11 APRIL 2014 STRASBOURG PALAIS DES CONGRÈS

First Part:

- Sparsity, low-rankness and relatives:
  "From information to structures"

- Compress while you sample:
  "From structures to scrambled sensing"

- and Reconstruct!
  "From scrambled sensing to information"

First Part:

- Sparsity, low-rankness and relatives:
  "From information to structures"

- Compress while you sample:
  "From structures to scrambled sensing"

- and Reconstruct!
  "From scrambled sensing to information"

Informative signals are composed of structures ...

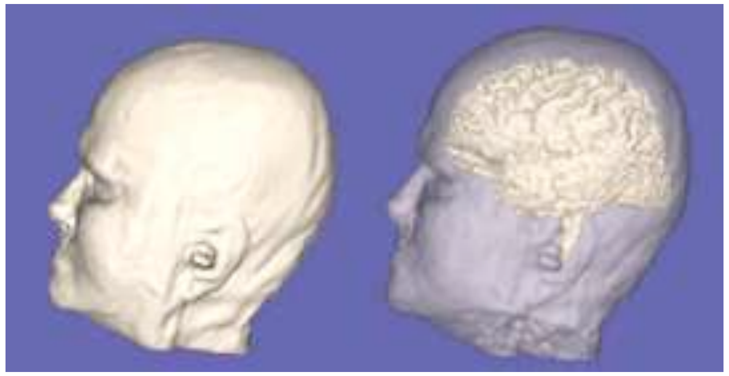

3-D data

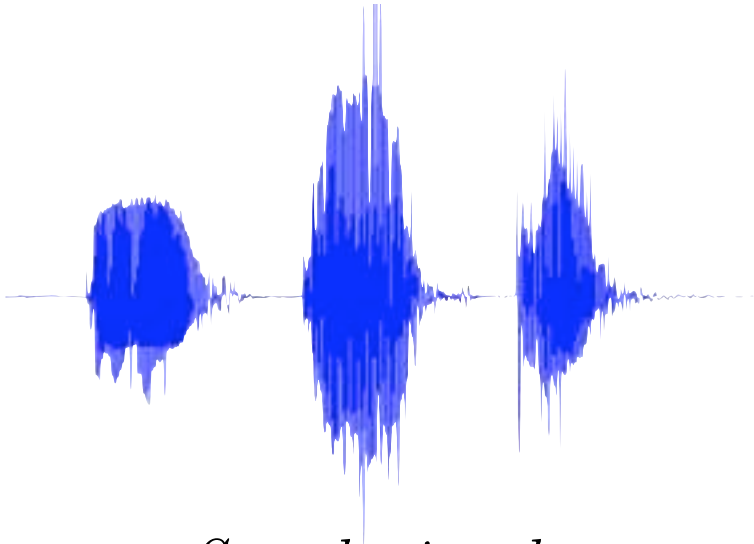

Speech signal

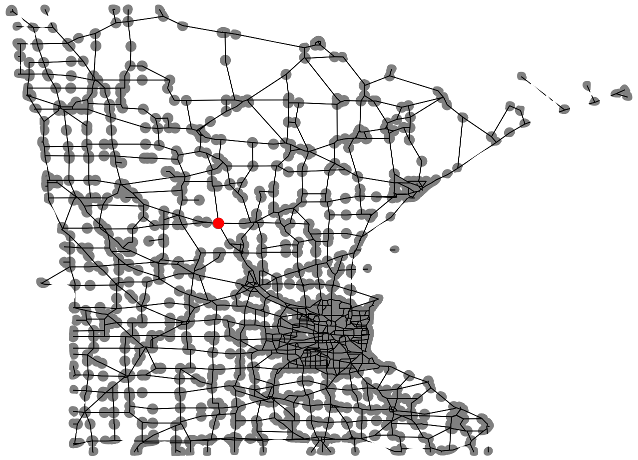

Data on Graph

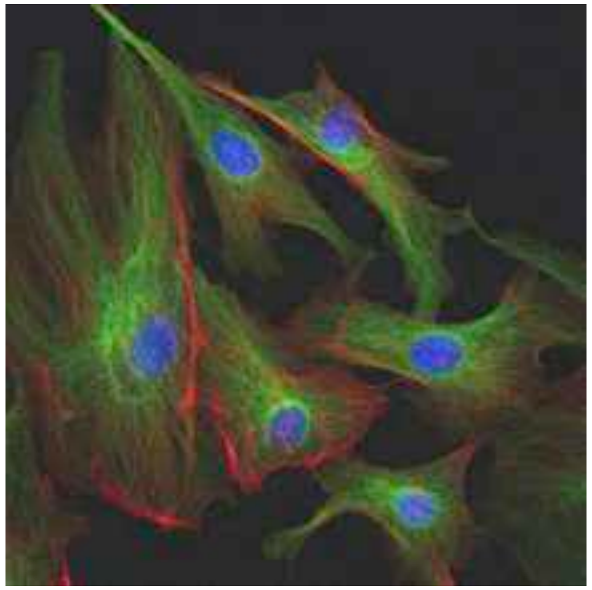

Biology

Video

Astronomy

2-D Example: Using Wavelets!

Representing this image ...

2-D Example: Using Wavelets!

...with those "wavelets"

Representing this image ...

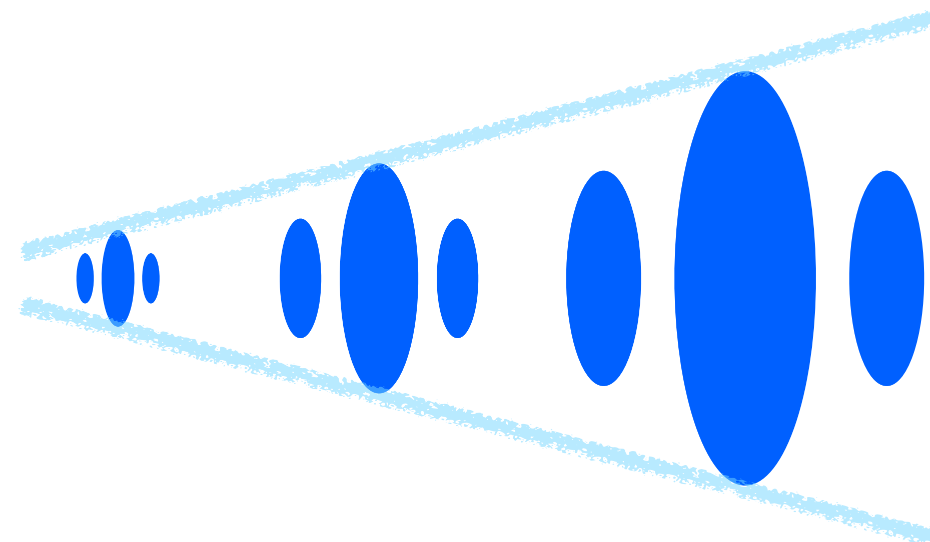

different sizes, scales

e.g.,

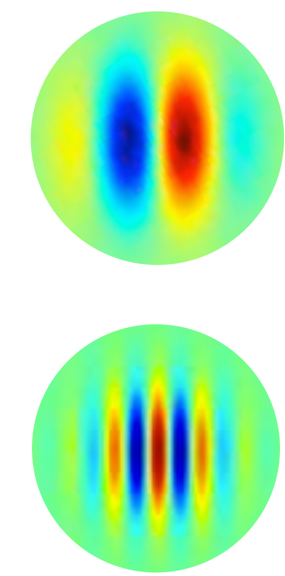

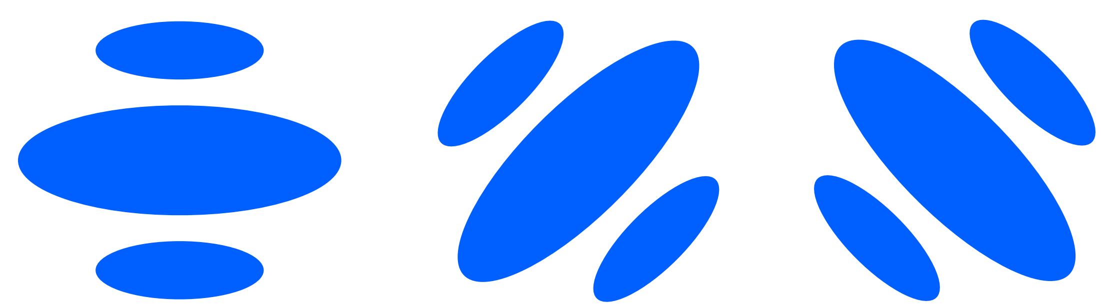

different orientations

2-D Example: Using Wavelets!

...with those "wavelets"

Representing this image ...

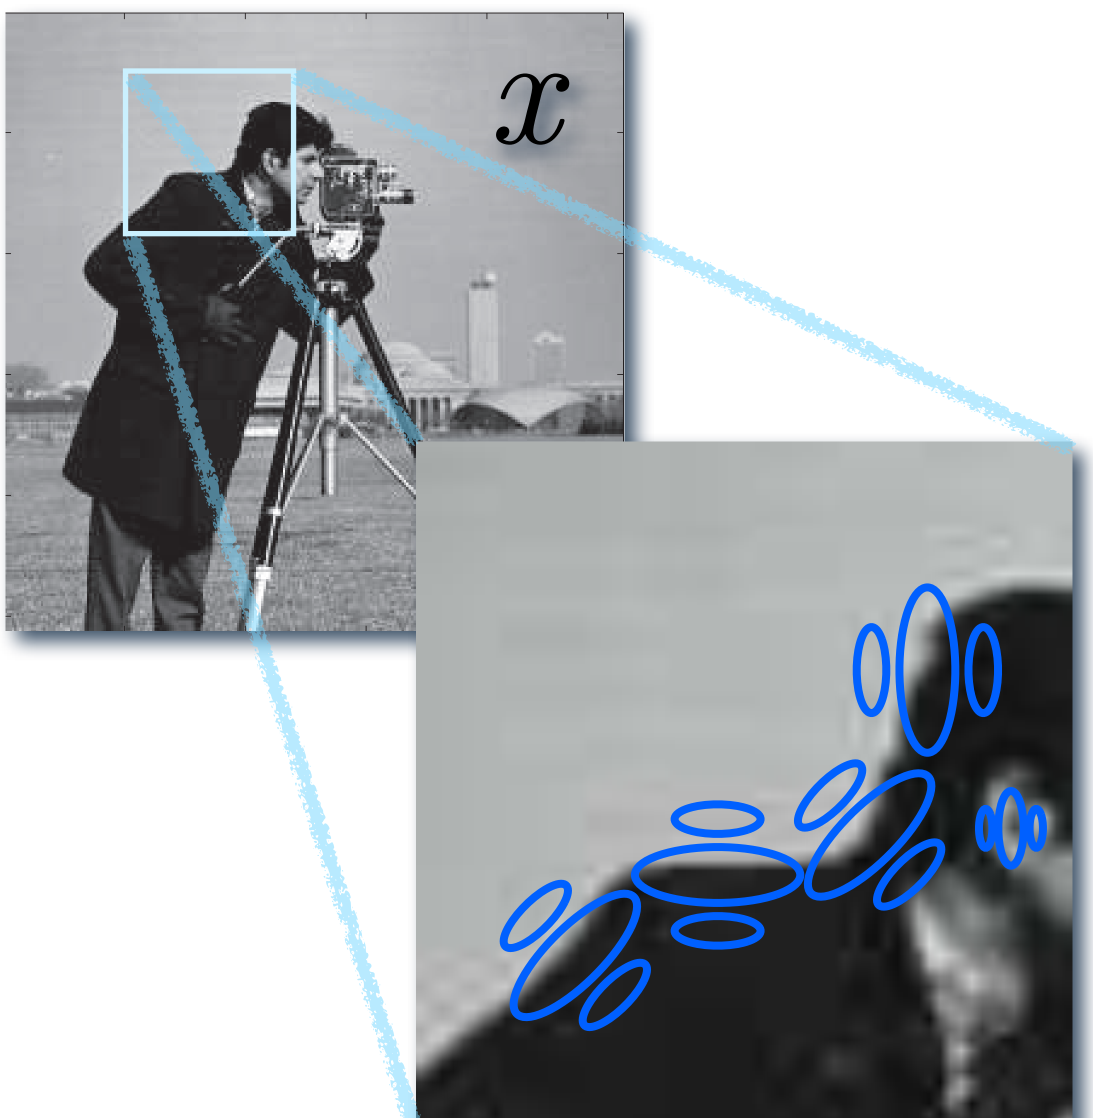

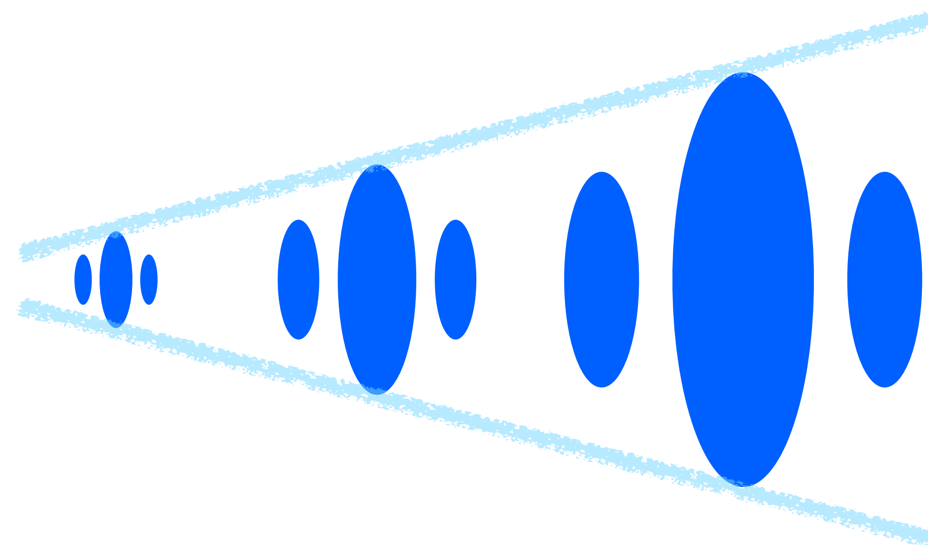

different sizes, scales

e.g.,

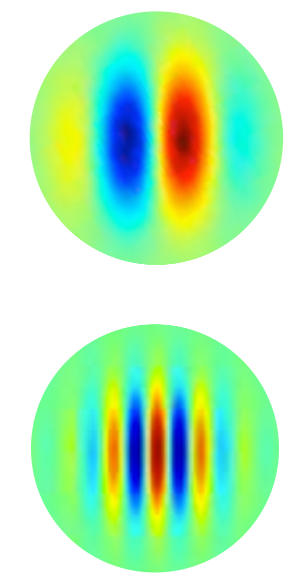

different orientations

2-D Example: Using Wavelets!

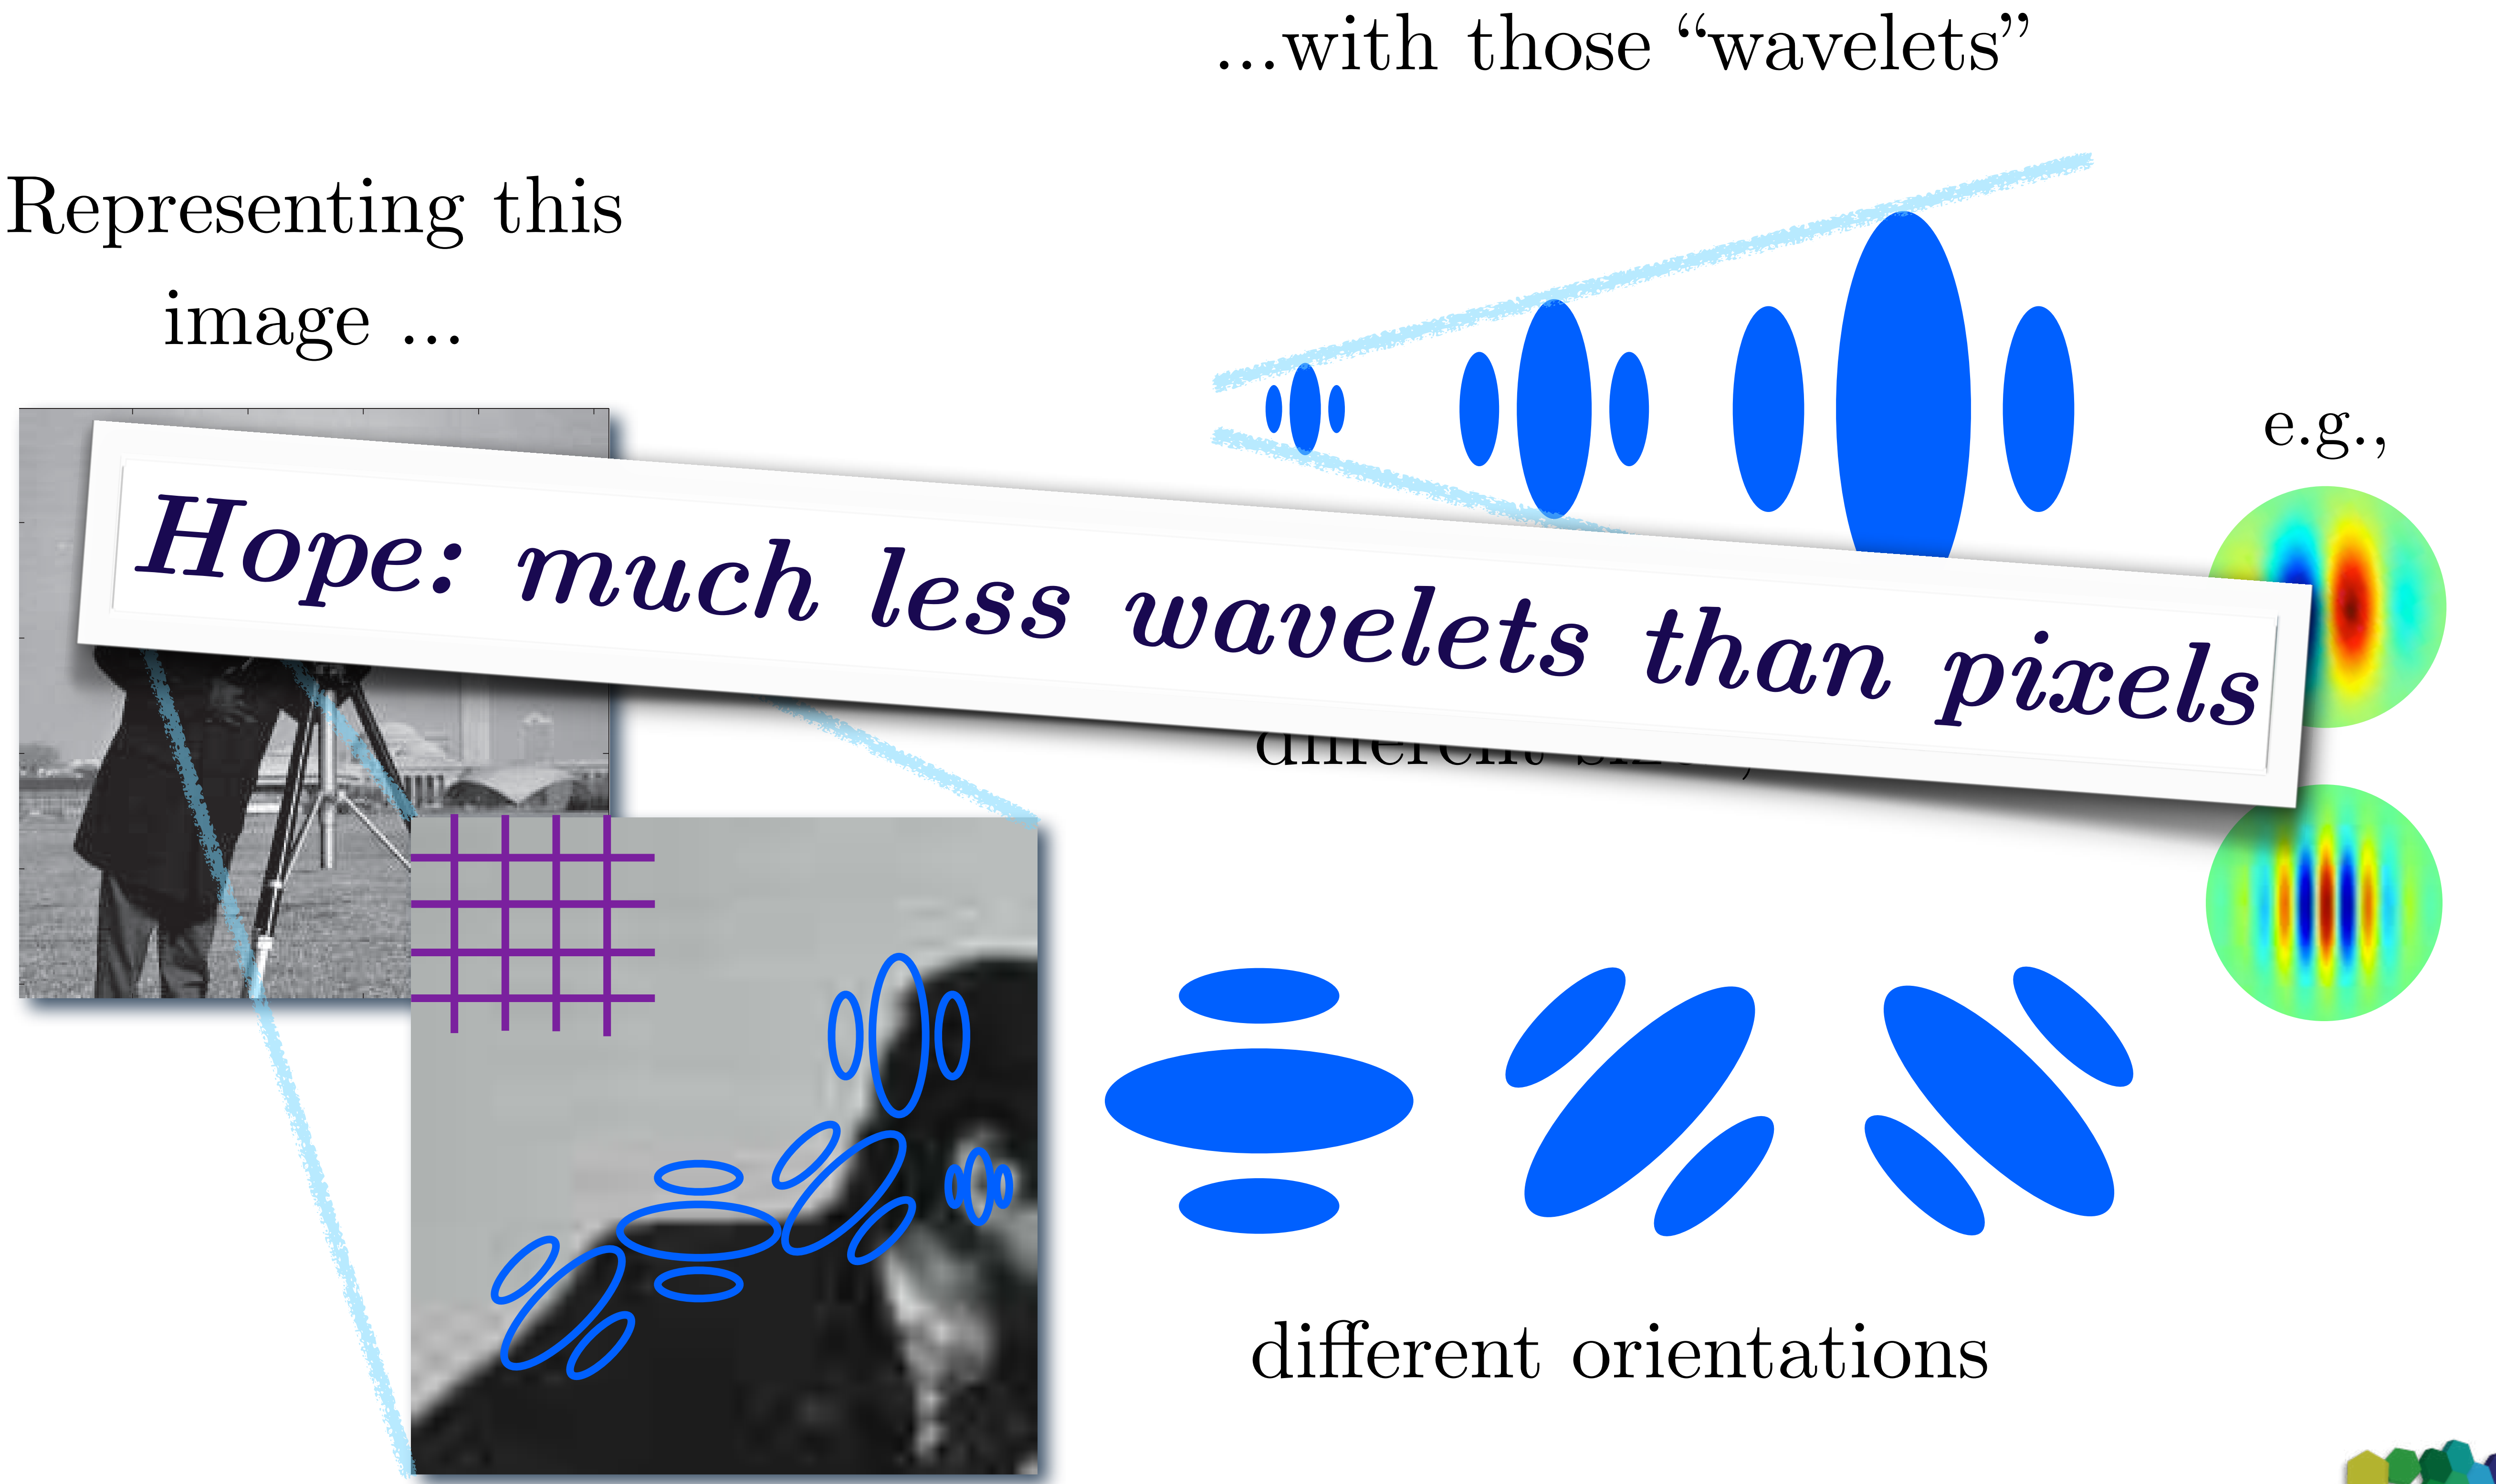

2-D Example: Using Wavelets!

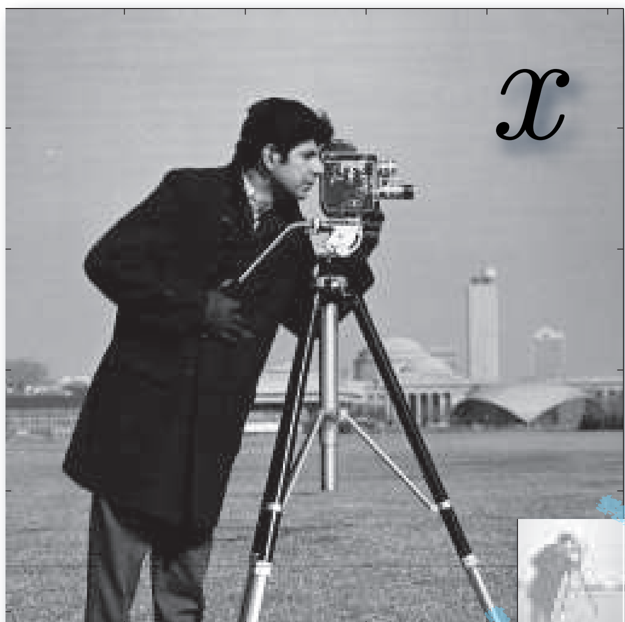

Set of coefficients in Ψ
(gray=0)

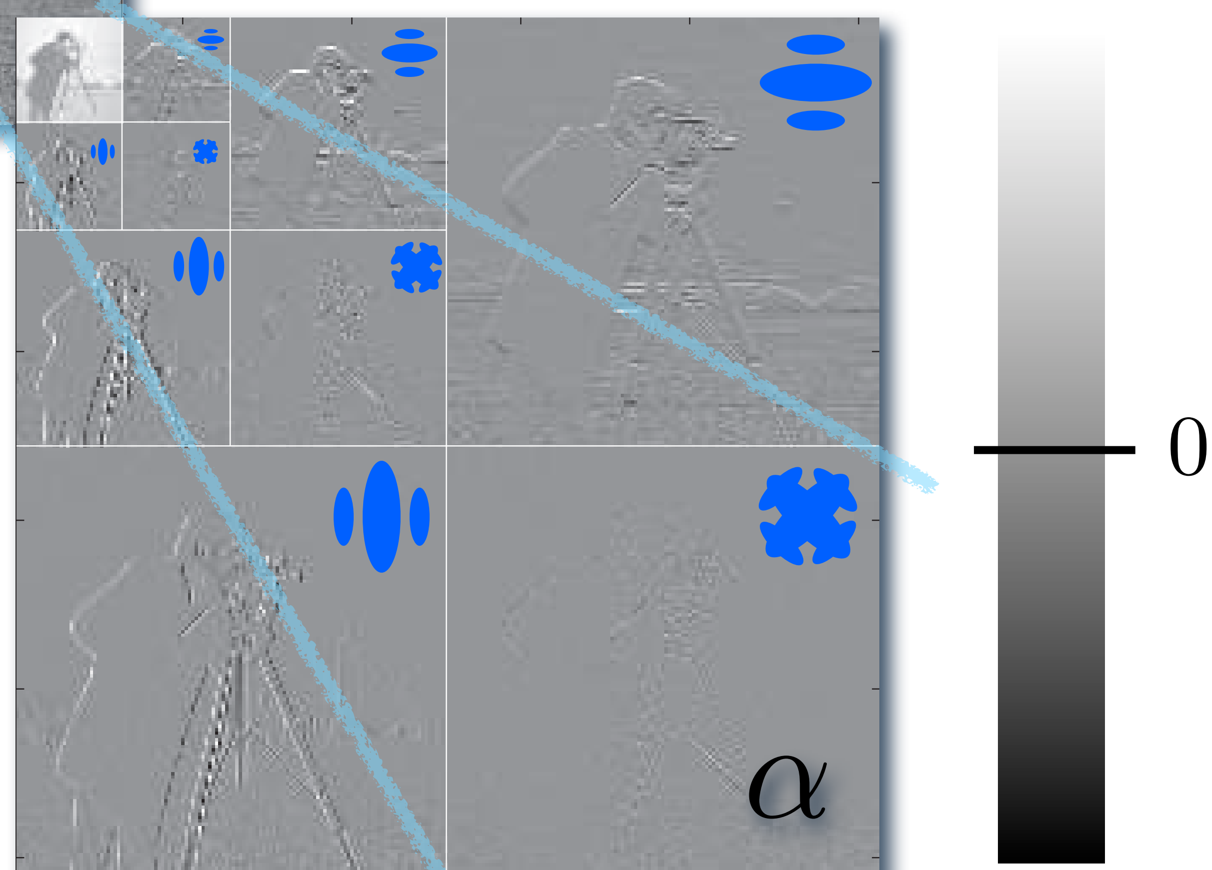

Wavelet
basis Ψ

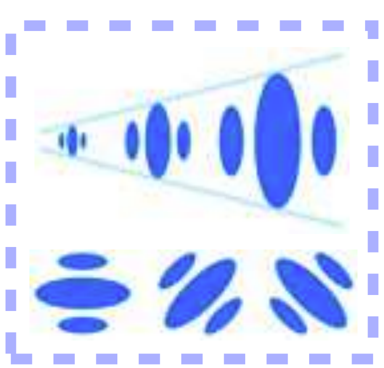

$$\alpha \text{ s.t. } x = \sum_i \alpha_i \Psi_i$$

2-D Example: Using Wavelets!

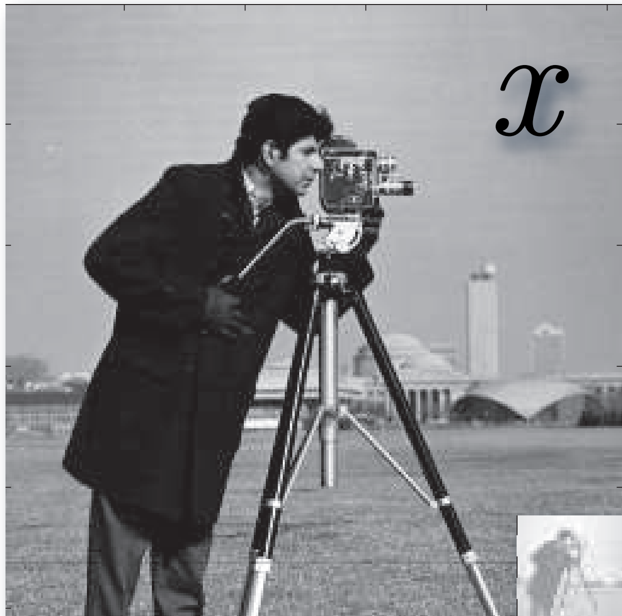

Set of coefficients in Ψ
(gray=0)

Wavelet
basis Ψ

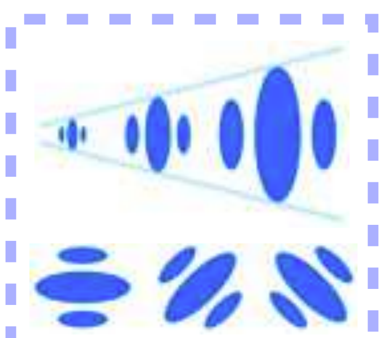

α s.t. $x = \sum_i \alpha_i \Psi_i$

Thresholding =
Keep K first values

α_K

2-D Example: Using Wavelets!

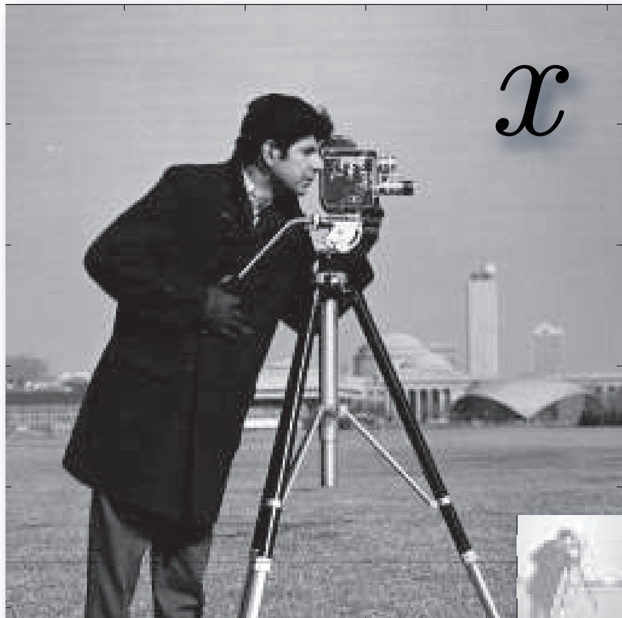

Set of coefficients in Ψ
(gray=0)

Wavelet
basis Ψ

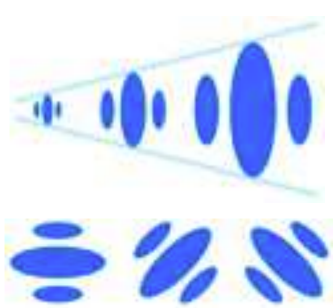

α s.t. $x = \sum_i \alpha_i \Psi_i$

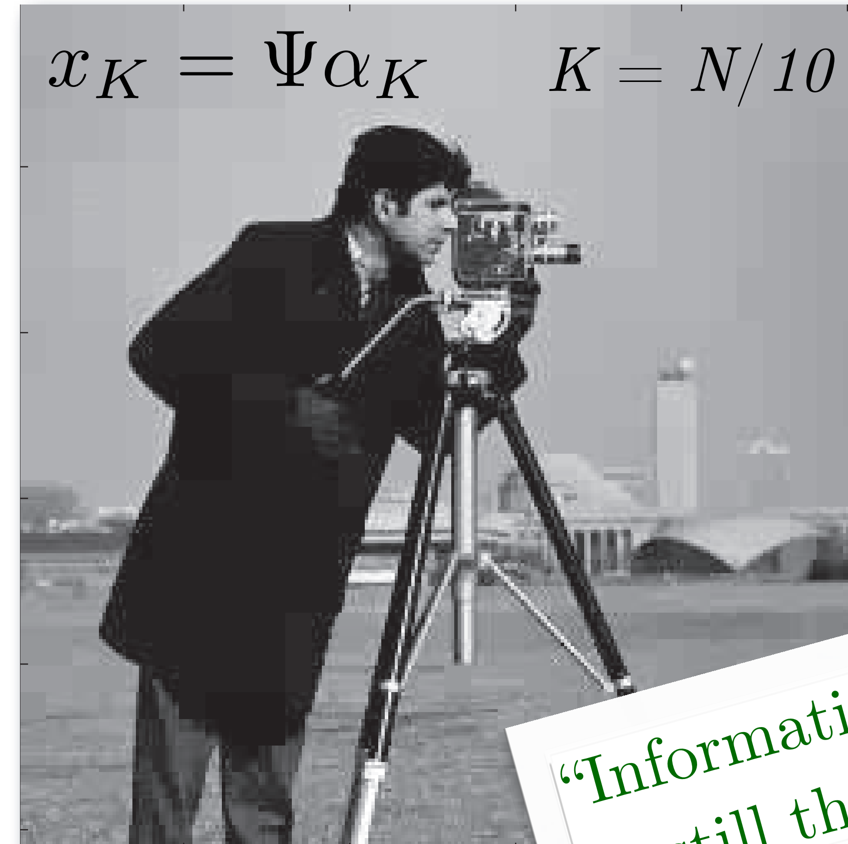

"Information" is still there!

α_K

Thresholding =
Keep K first values

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x = \sum_{j=1}^{D} \alpha_j \Psi_j = \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N \times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

matrix repr.

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x = \sum_{j=1}^{D} \alpha_j \Psi_j = \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x \;=\; \sum_{j=1}^{D} \alpha_j \Psi_j \;=\; \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

f

$\Psi =$ Dictionary of structures

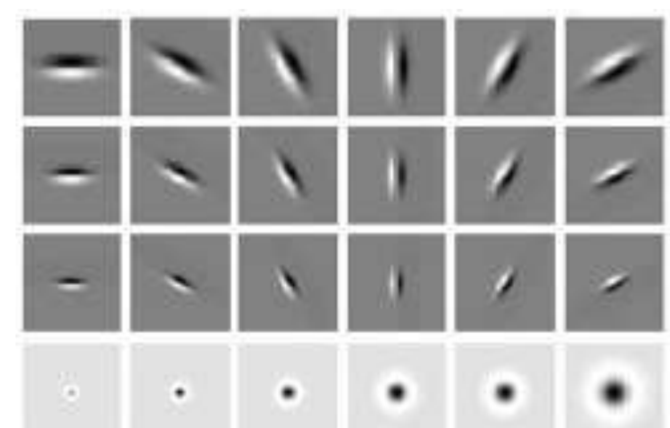

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x = \sum_{j=1}^{D} \alpha_j \Psi_j = \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

f

10 atoms

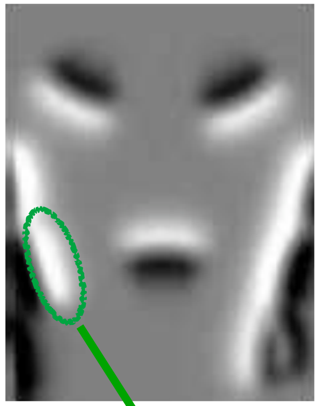

atom

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x \;=\; \sum_{j=1}^{D} \alpha_j \Psi_j \;=\; \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

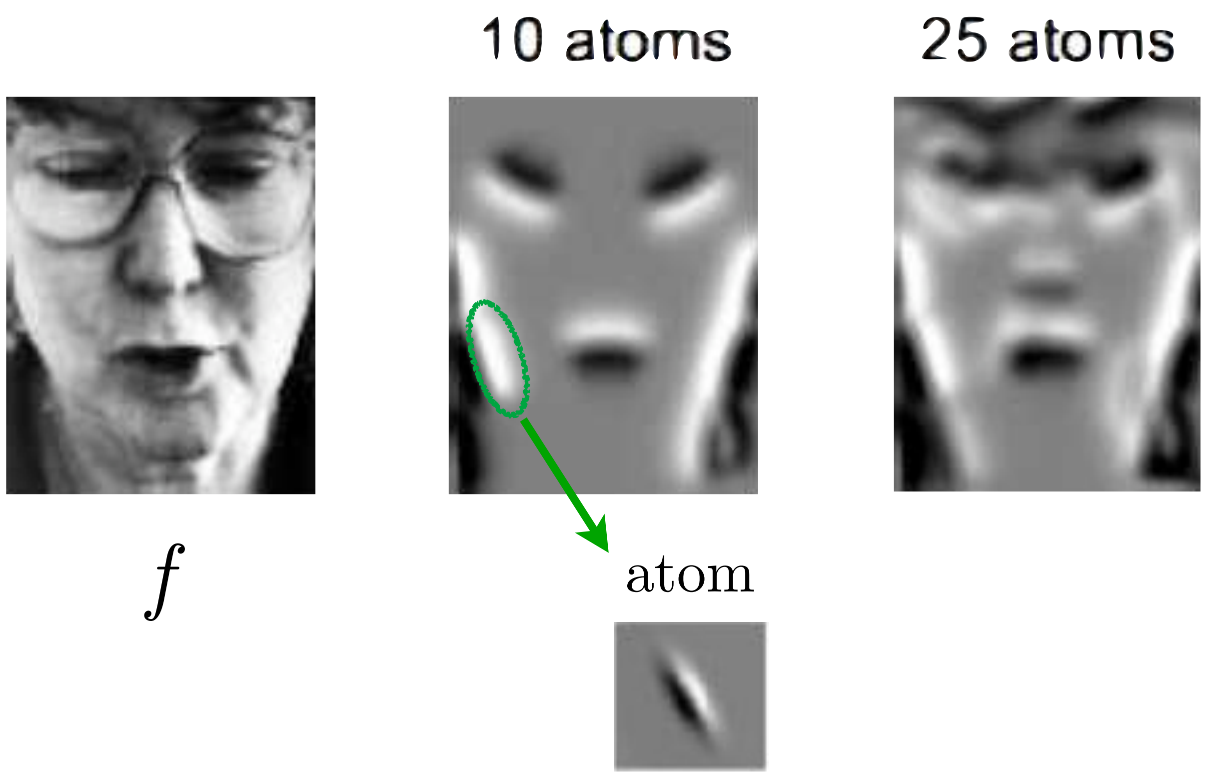

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a "sparsity basis" Ψ with few non-zero elements α :

$$x \;=\; \sum_{j=1}^{D} \alpha_j \Psi_j \;=\; \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

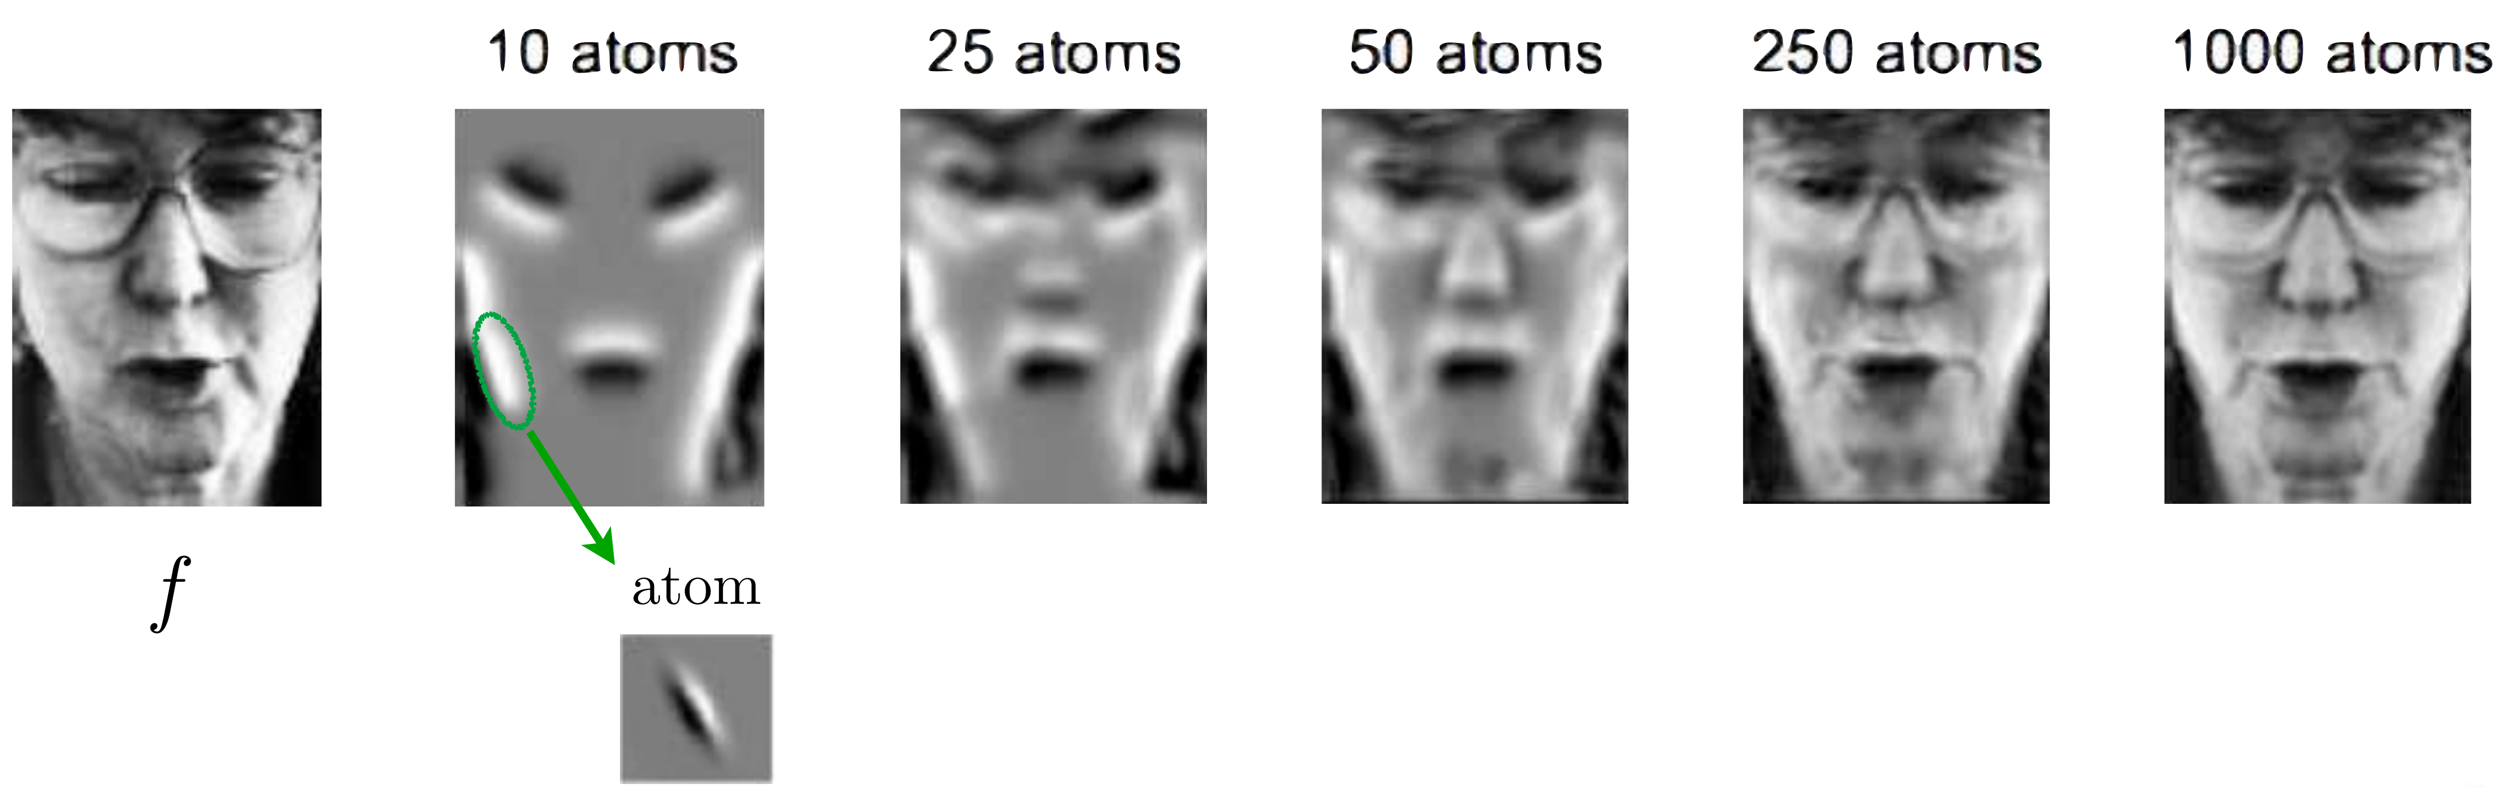

Generalization: Sparsity principle

- Hypothesis: an image (or any signal) can be decomposed in a “sparsity basis” Ψ with few non-zero elements α :

$$x \;=\; \sum_{j=1}^{D} \alpha_j \Psi_j \;=\; \Psi\alpha \qquad \begin{array}{l} \Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N\times D} \\ \alpha = (\alpha_1, \cdots, \alpha_D)^T \in \mathbb{R}^D \end{array}$$

- Ψ can be a ONB (e.g. Fourier, wavelets) or a dictionary (atoms)

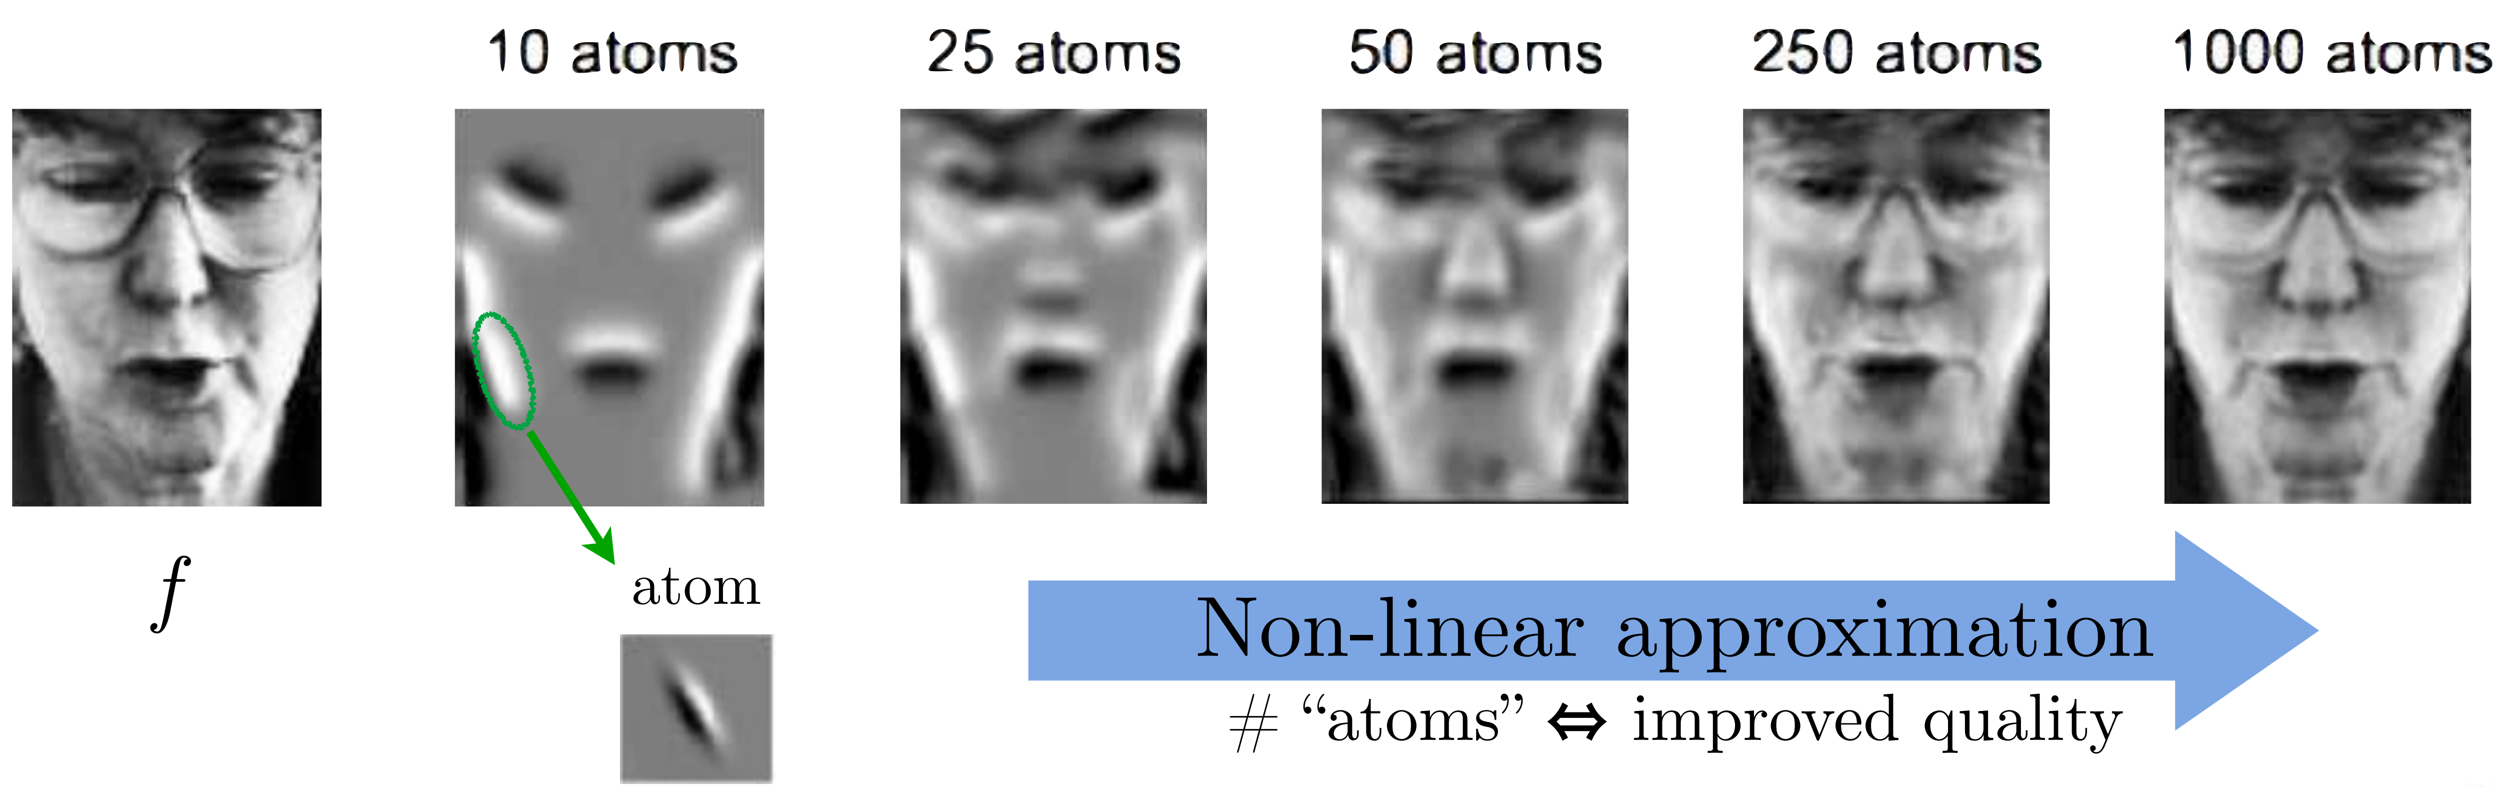

In summary: if “information” ...

... in a signal $x \in \mathbb{R}^N$ (e.g. $N =$ pixel number, voxels, graph nodes, ...)

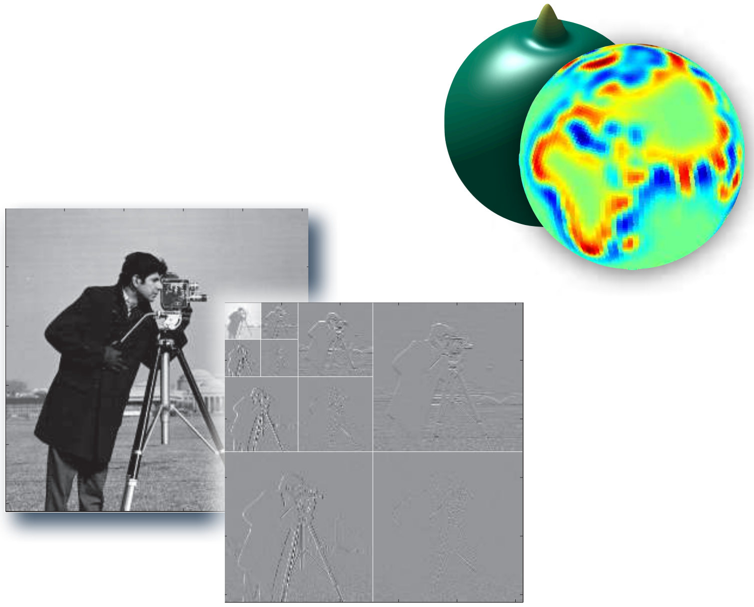

In summary: if "information" ...

... in a signal $x \in \mathbb{R}^N$ (e.g. $N =$ pixel number, voxels, graph nodes, ...)

there exists a "sparsity" basis (e.g. wavelets, Fourier, ...)

$$\Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N \times D}$$

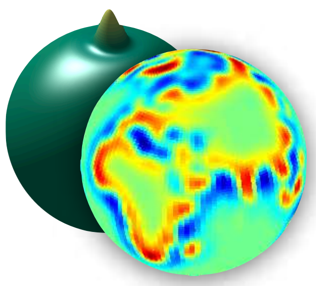

where x has a linear representation

$$x = \sum_{j=1}^{D} \alpha_j \Psi_j = \Psi \alpha$$

and

$$\|\alpha\|_0 := \#\{i : \alpha_i \neq 0\} \ll N \qquad \|\alpha - \alpha_K\| \ll \|\alpha\|$$

$\alpha =$ 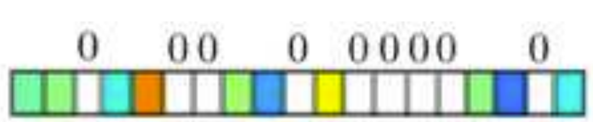

or

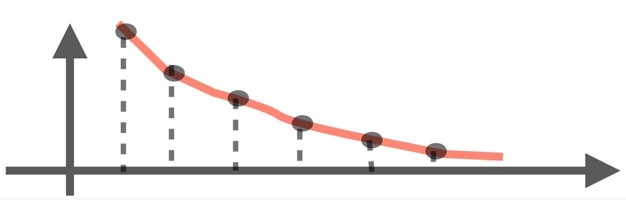

Counterexample: Noise!

In summary: if "information" ...

... in a signal $x \in \mathbb{R}^N$ (e.g. $N =$ pixel number, voxels, graph nodes, ...)

there exists a "sparsity" basis (e.g. wavelets, Fourier, ...)

$$\Psi = (\Psi_1, \cdots, \Psi_D) \in \mathbb{R}^{N \times D}$$

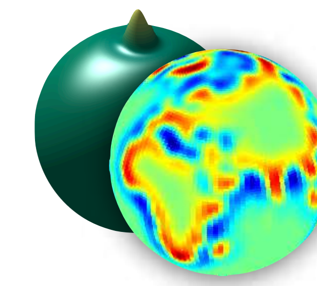

where x has a linear representation

$$x = \sum_{j=1}^{D} \alpha_j \Psi_j = \Psi\alpha$$

and

$\|\alpha\|_0 := \#\{i : \alpha_i \neq 0\} \ll N$

$\alpha =$ 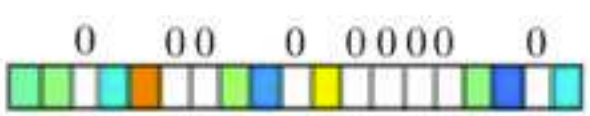

or

$\|\alpha - \alpha_K\| \ll \|\alpha\|$

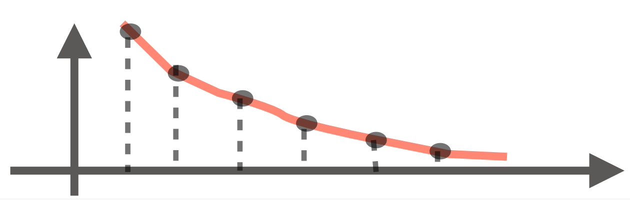

or

Small ℓ_1-norm

$\|\alpha\|_1 = \sum_i |\alpha_i|$

Counterexample: Noise!

Convex! (see after)

Other “informative” models

Other “informative” models

1. Structured sparsity for high-dimensional data

Other "informative" models

1. Structured sparsity for high-dimensional data

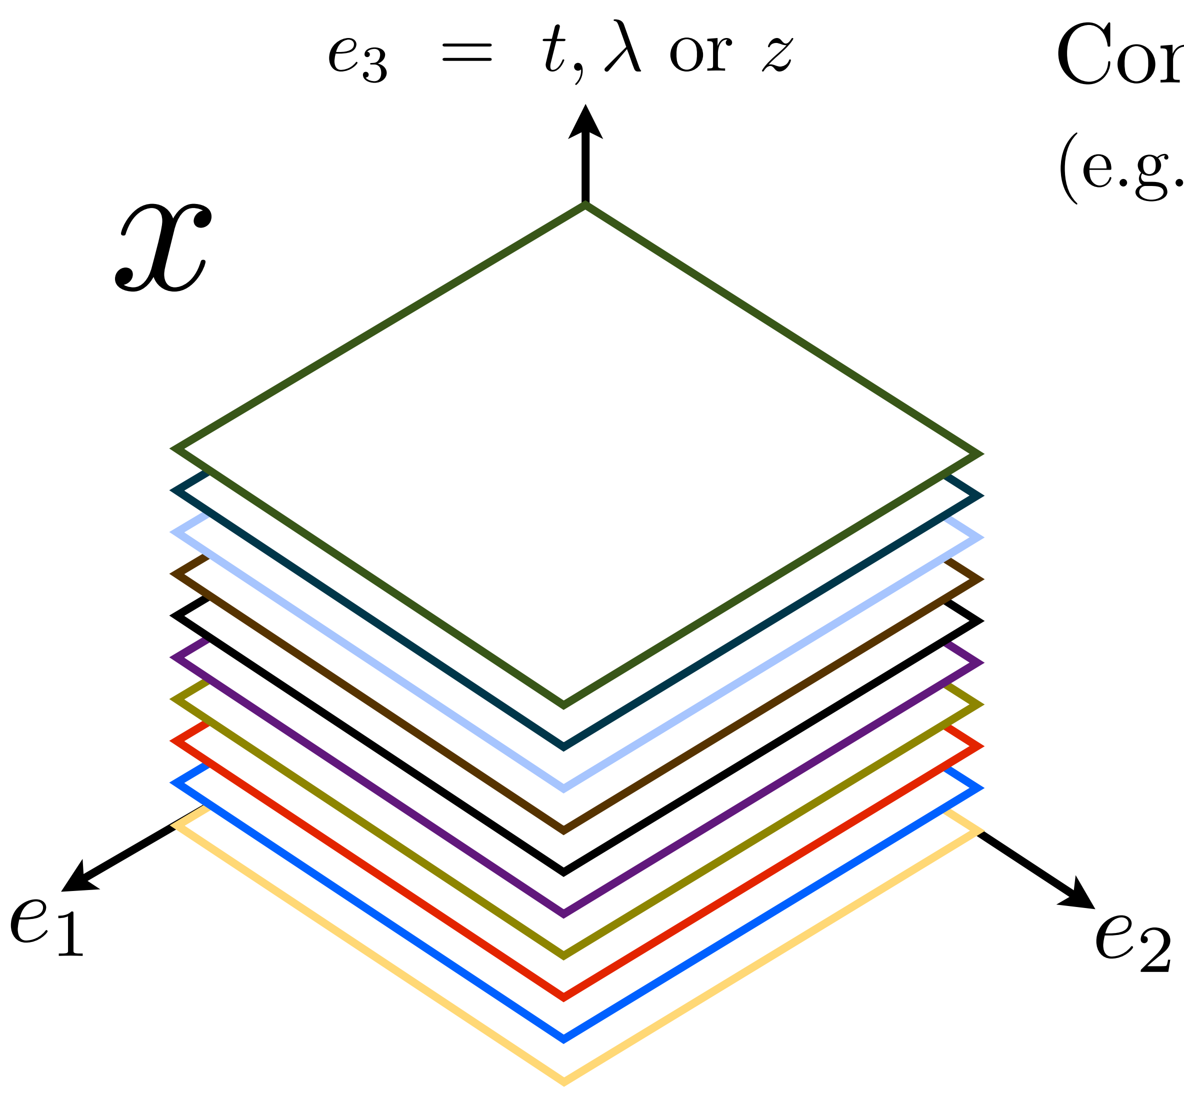

Consider a data volume (3-D or more)
(e.g. video, hyperspectral data, medical data)

Possible models:

- 3-D sparsity basis (see before)
  (sometimes costly, sometimes $\nexists$)
- or *structured sparsity*
  idea: consecutive "slices" vary "slowly"

Other "informative" models

1. Structured sparsity for high-dimensional data

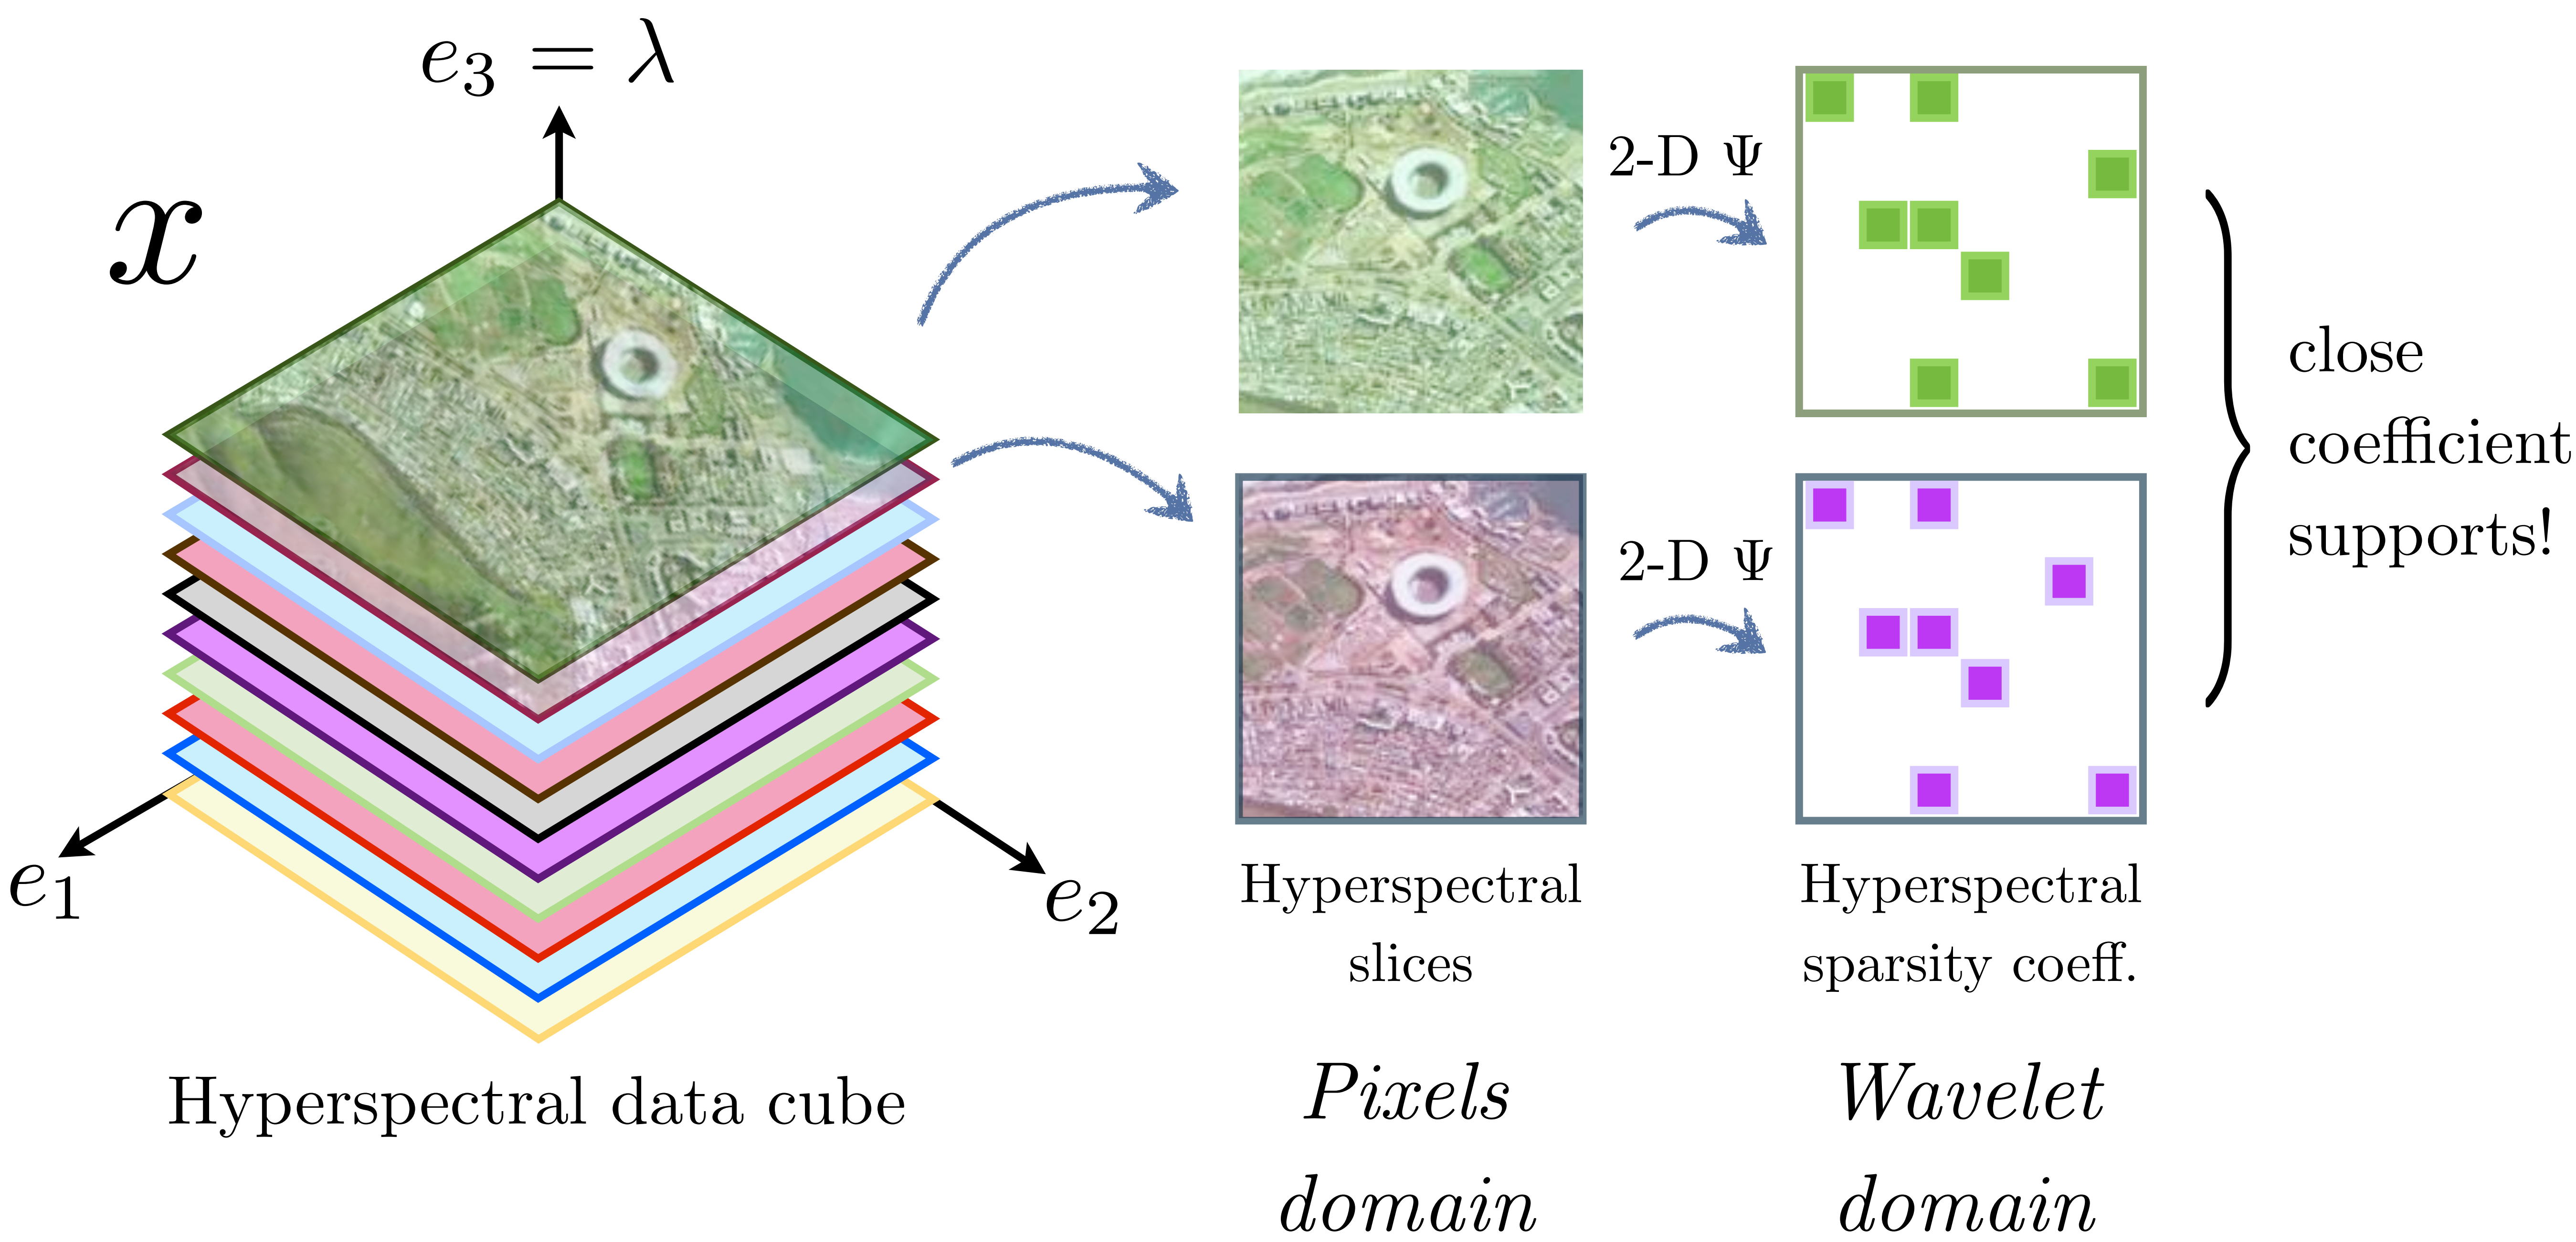

Other "informative" models

1. Structured sparsity for high-dimensional data

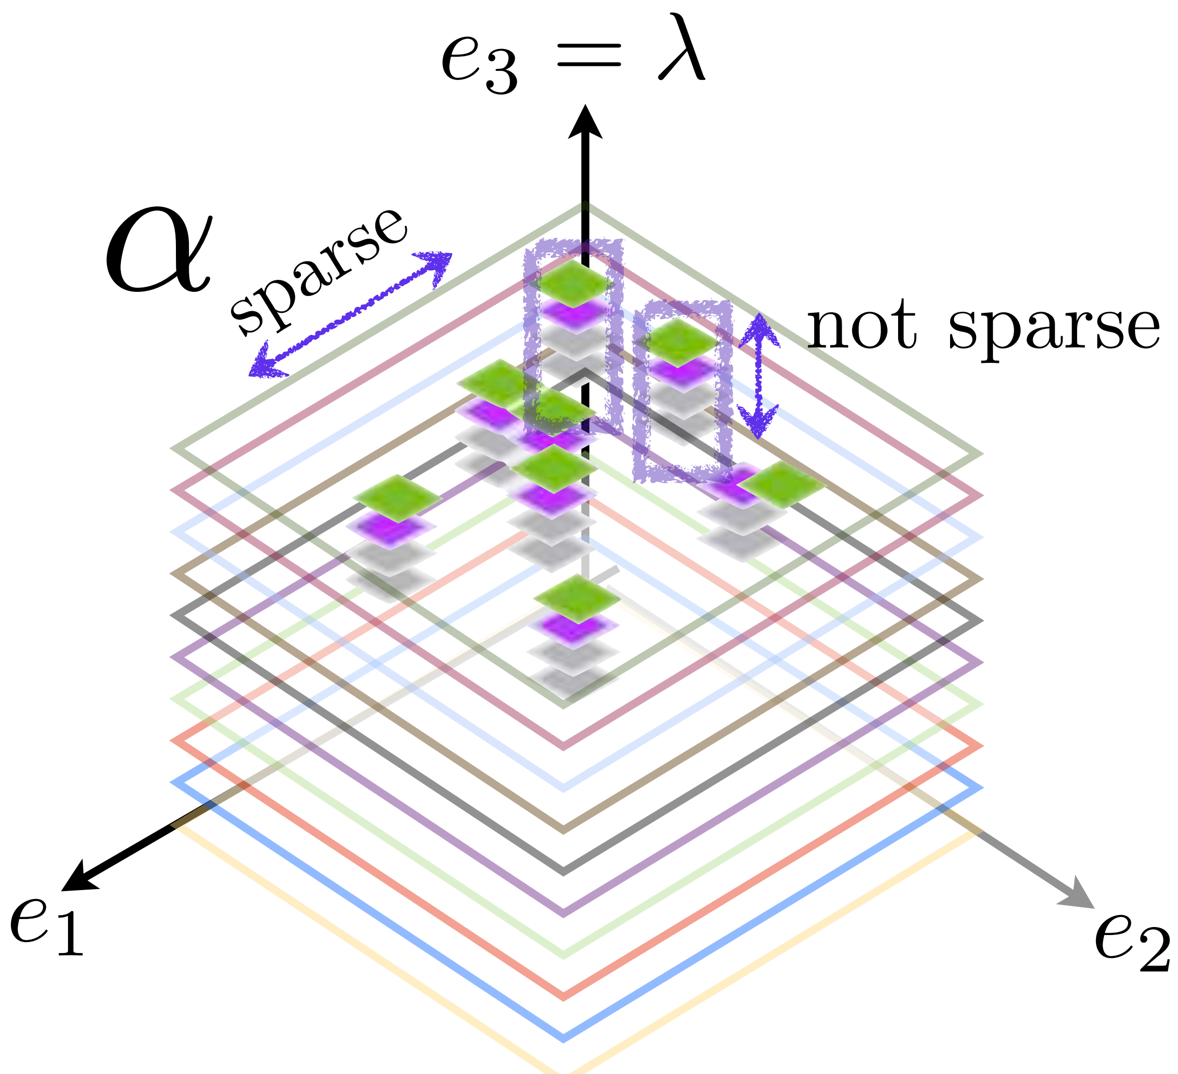

Hyperspectral sparsity coeff.

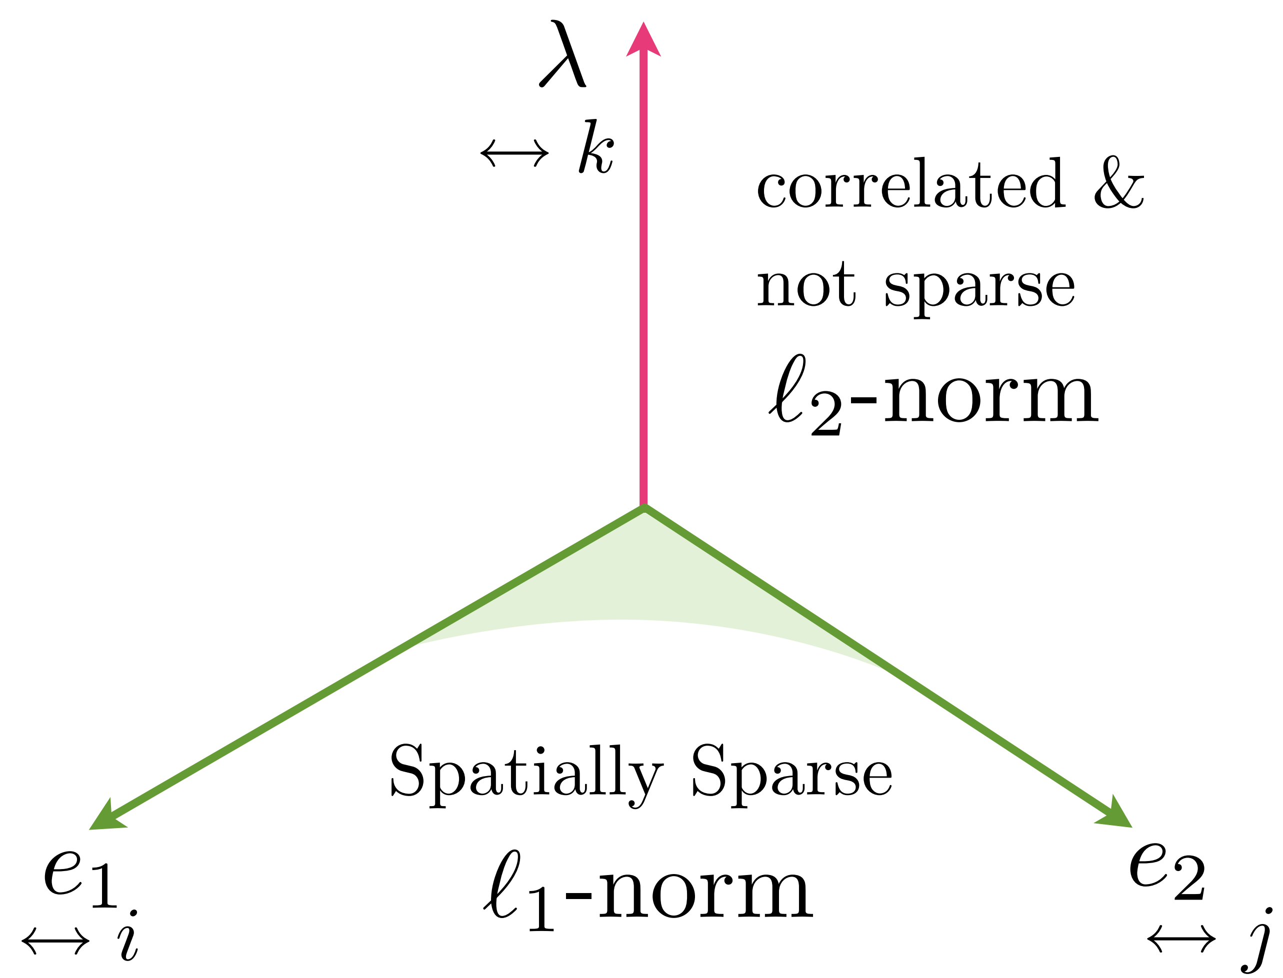

Other "informative" models

1. Structured sparsity for high-dimensional data

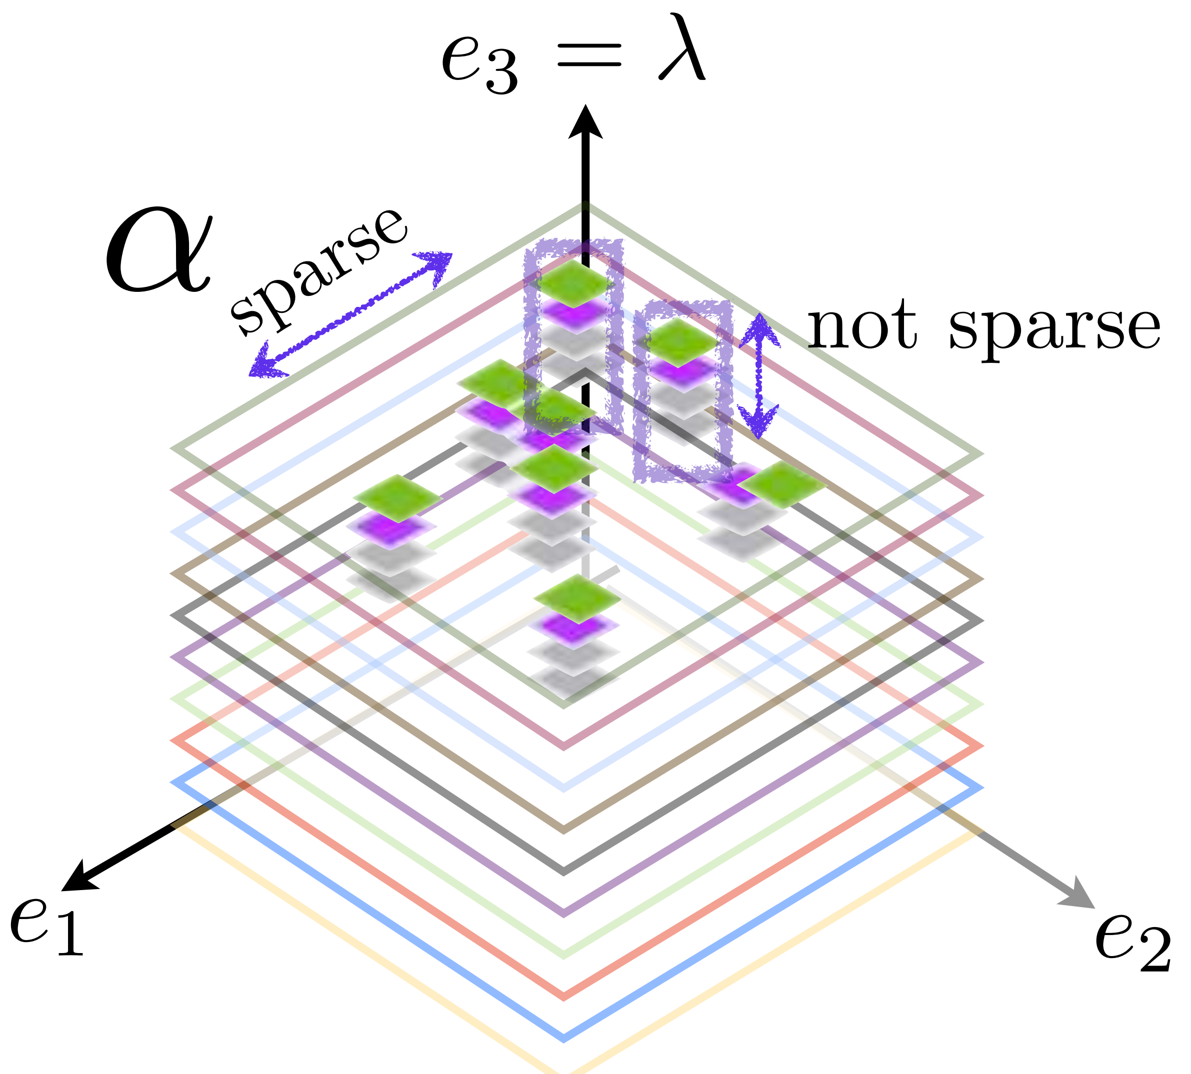

Hyperspectral
sparsity coeff.

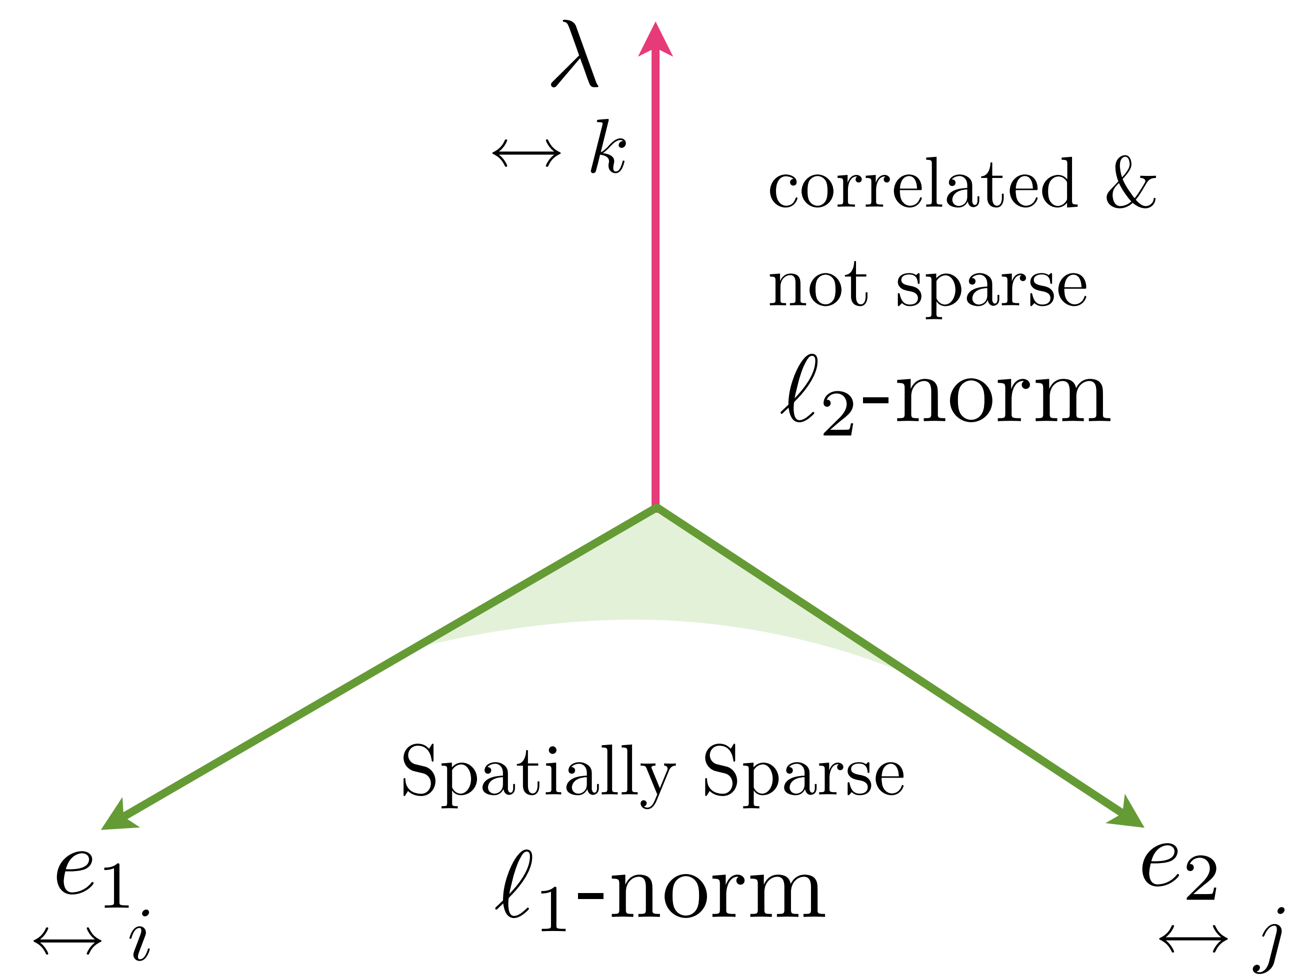

$\Rightarrow$ small $\ell_{2,1}$-norm

$$\|\alpha\|_{2,1} = \textstyle\sum_{ij} \big(\sum_k |\alpha_{i,j,k}|^2\big)^{1/2}$$

Other "informative" models

2. Low-rank models in high dimensions

Other "informative" models

2. Low-rank models in high dimensions

Example:
Video
(PETS DB)

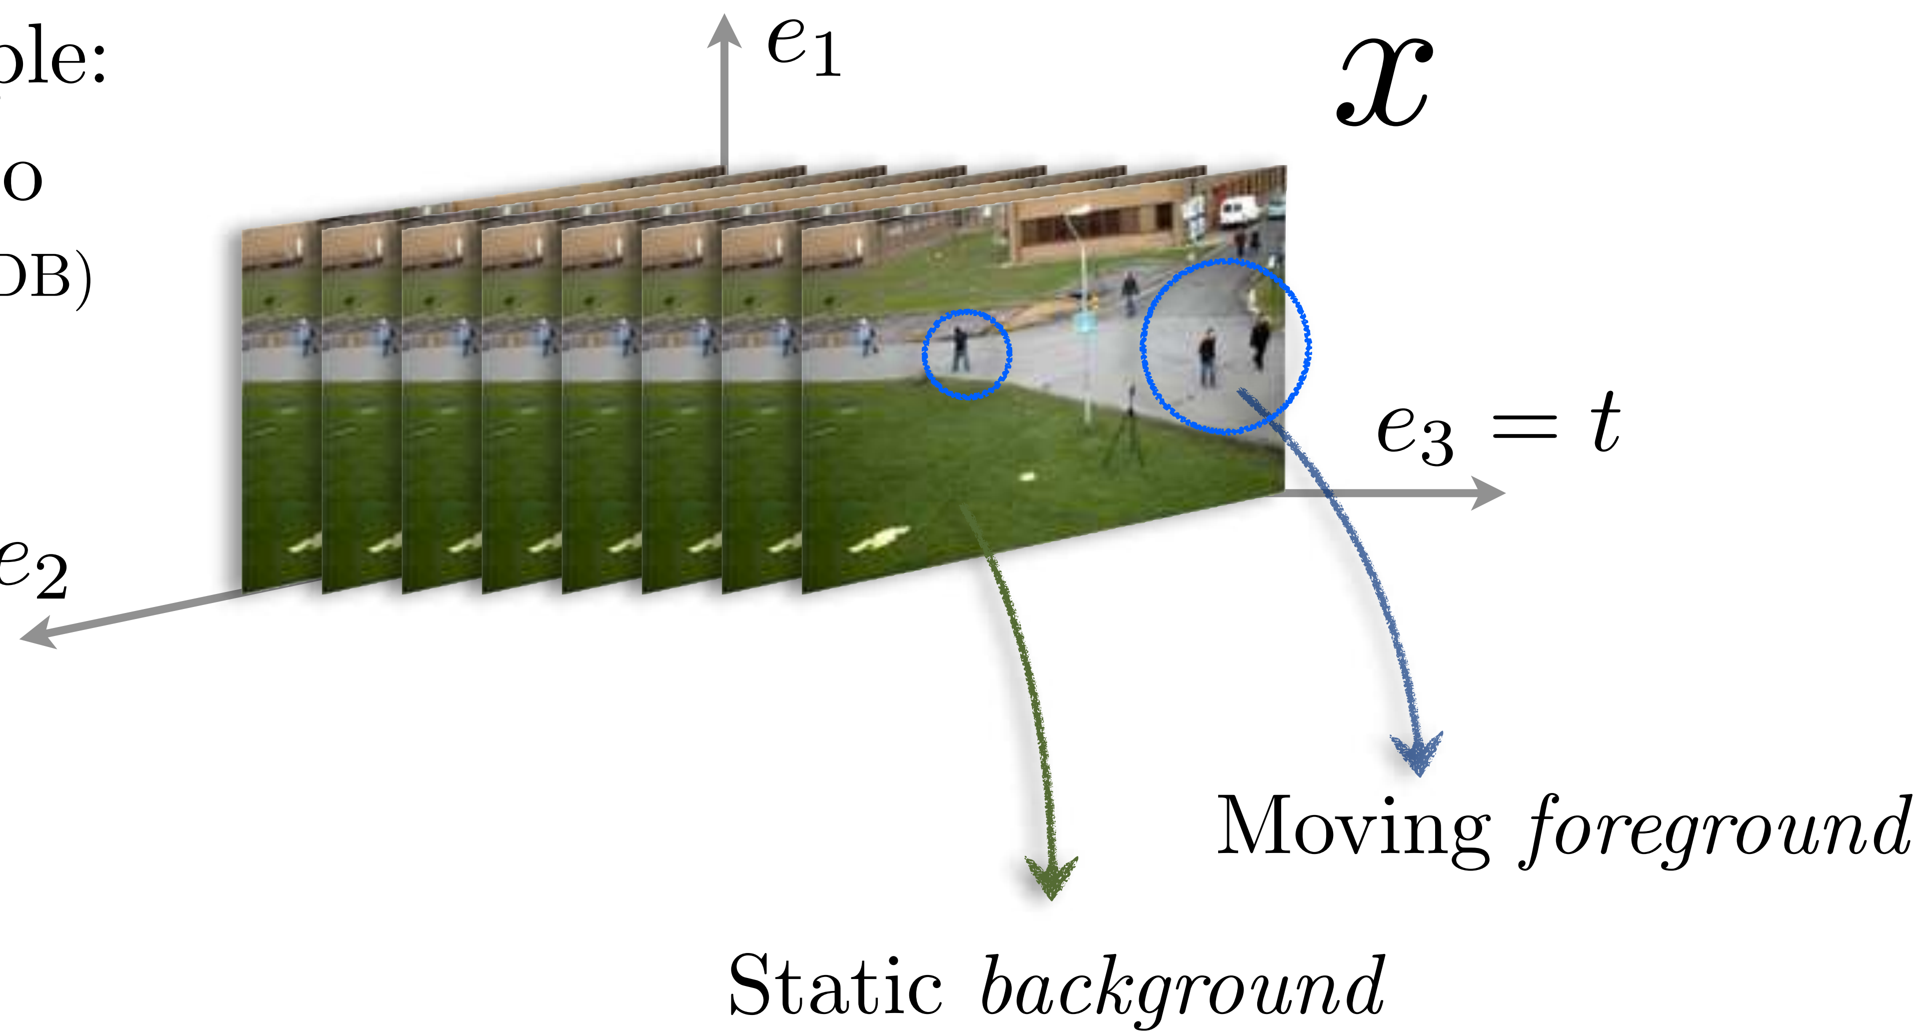

Other "informative" models

2. Low-rank models in high dimensions

Example:

Video

(PETS DB)

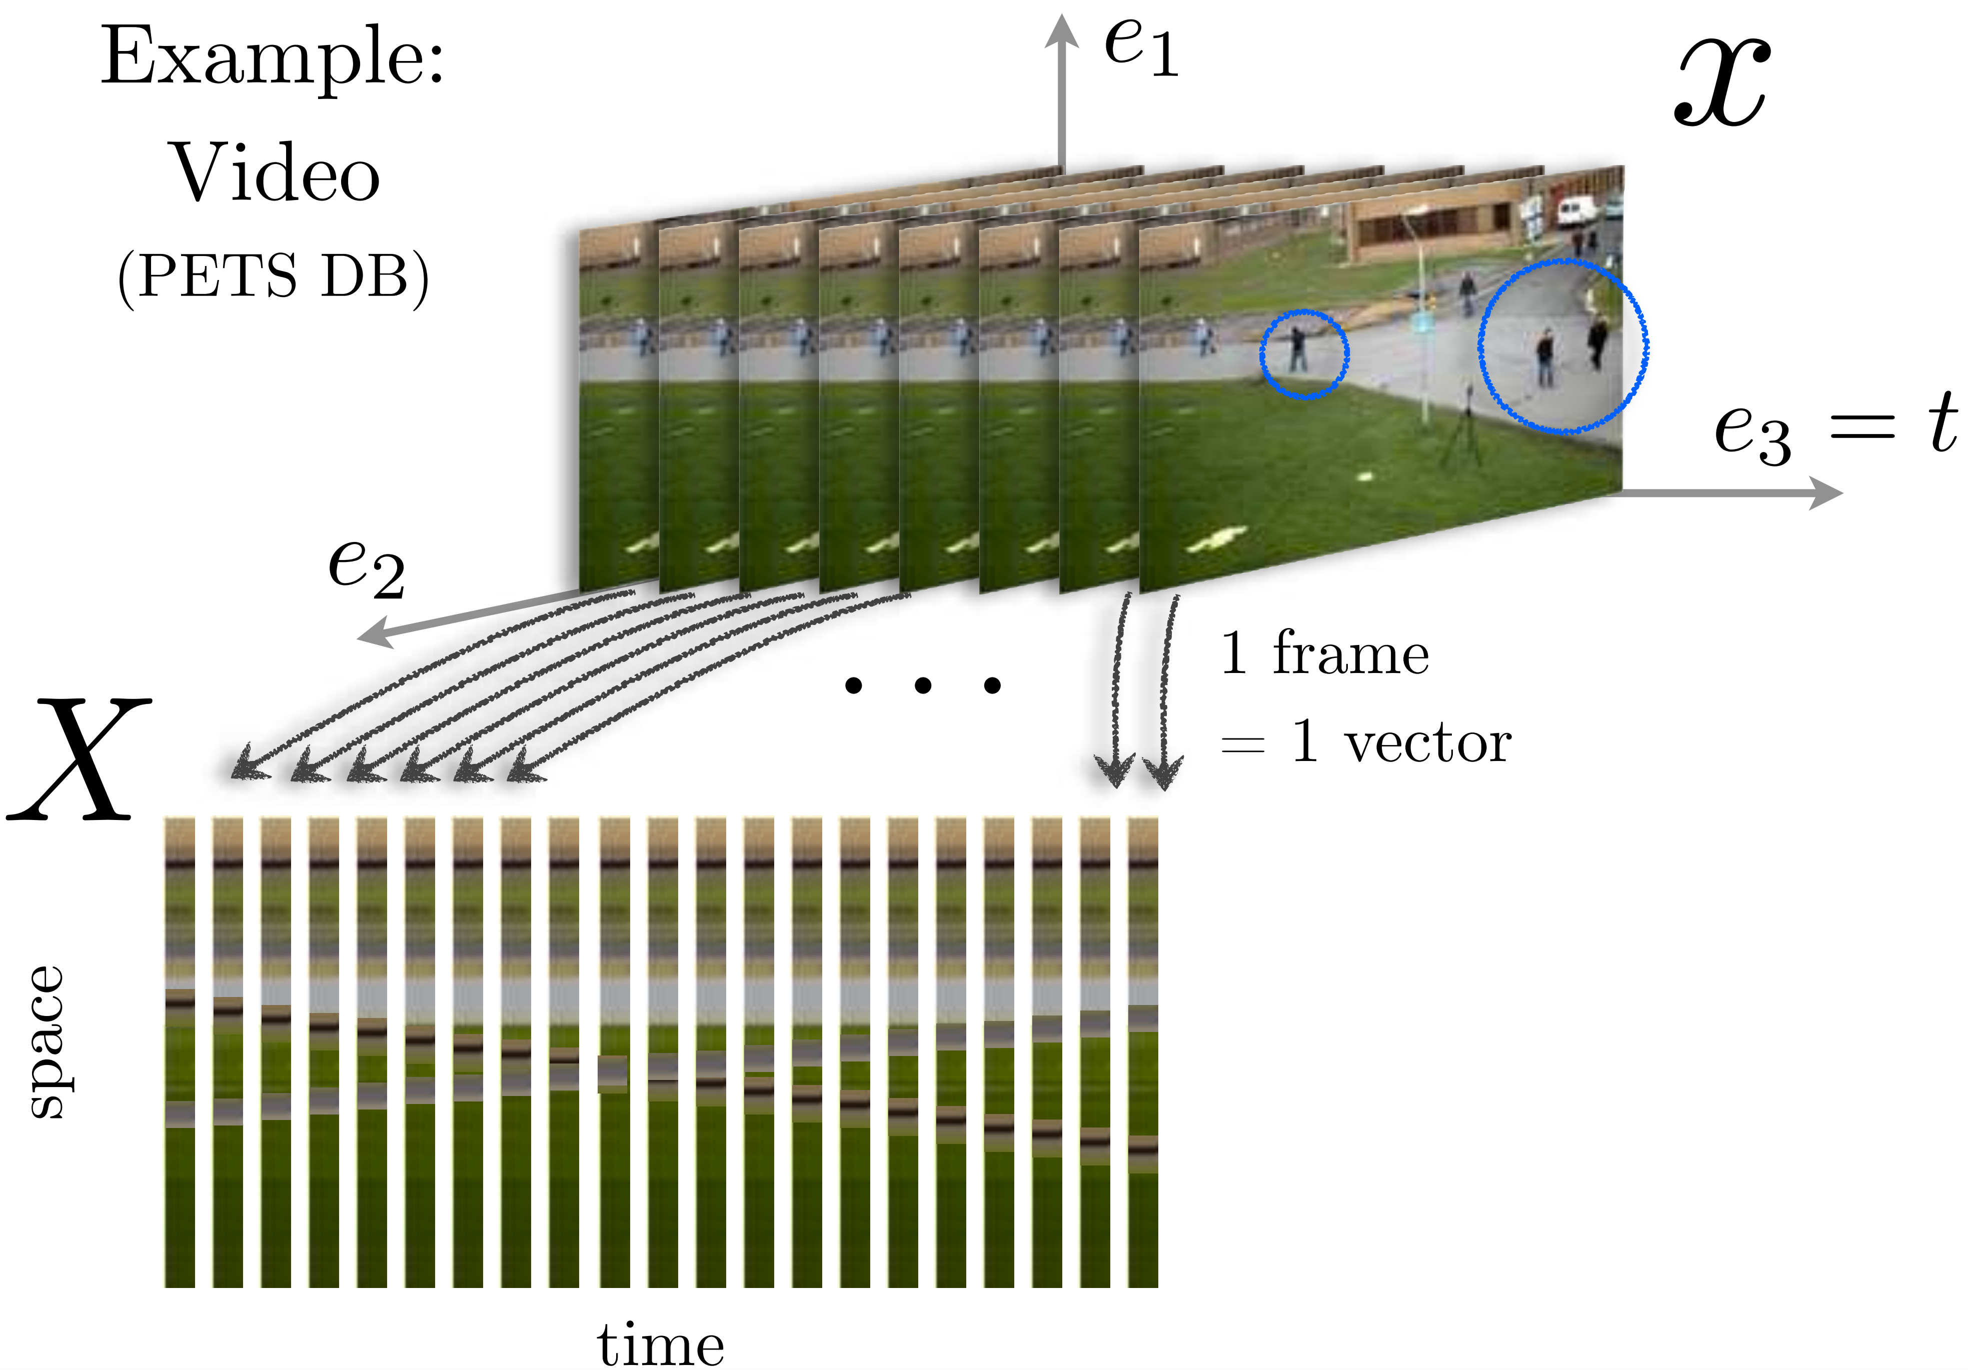

Other "informative" models

2. Low-rank models in high dimensions

Example:
Video
(PETS DB)

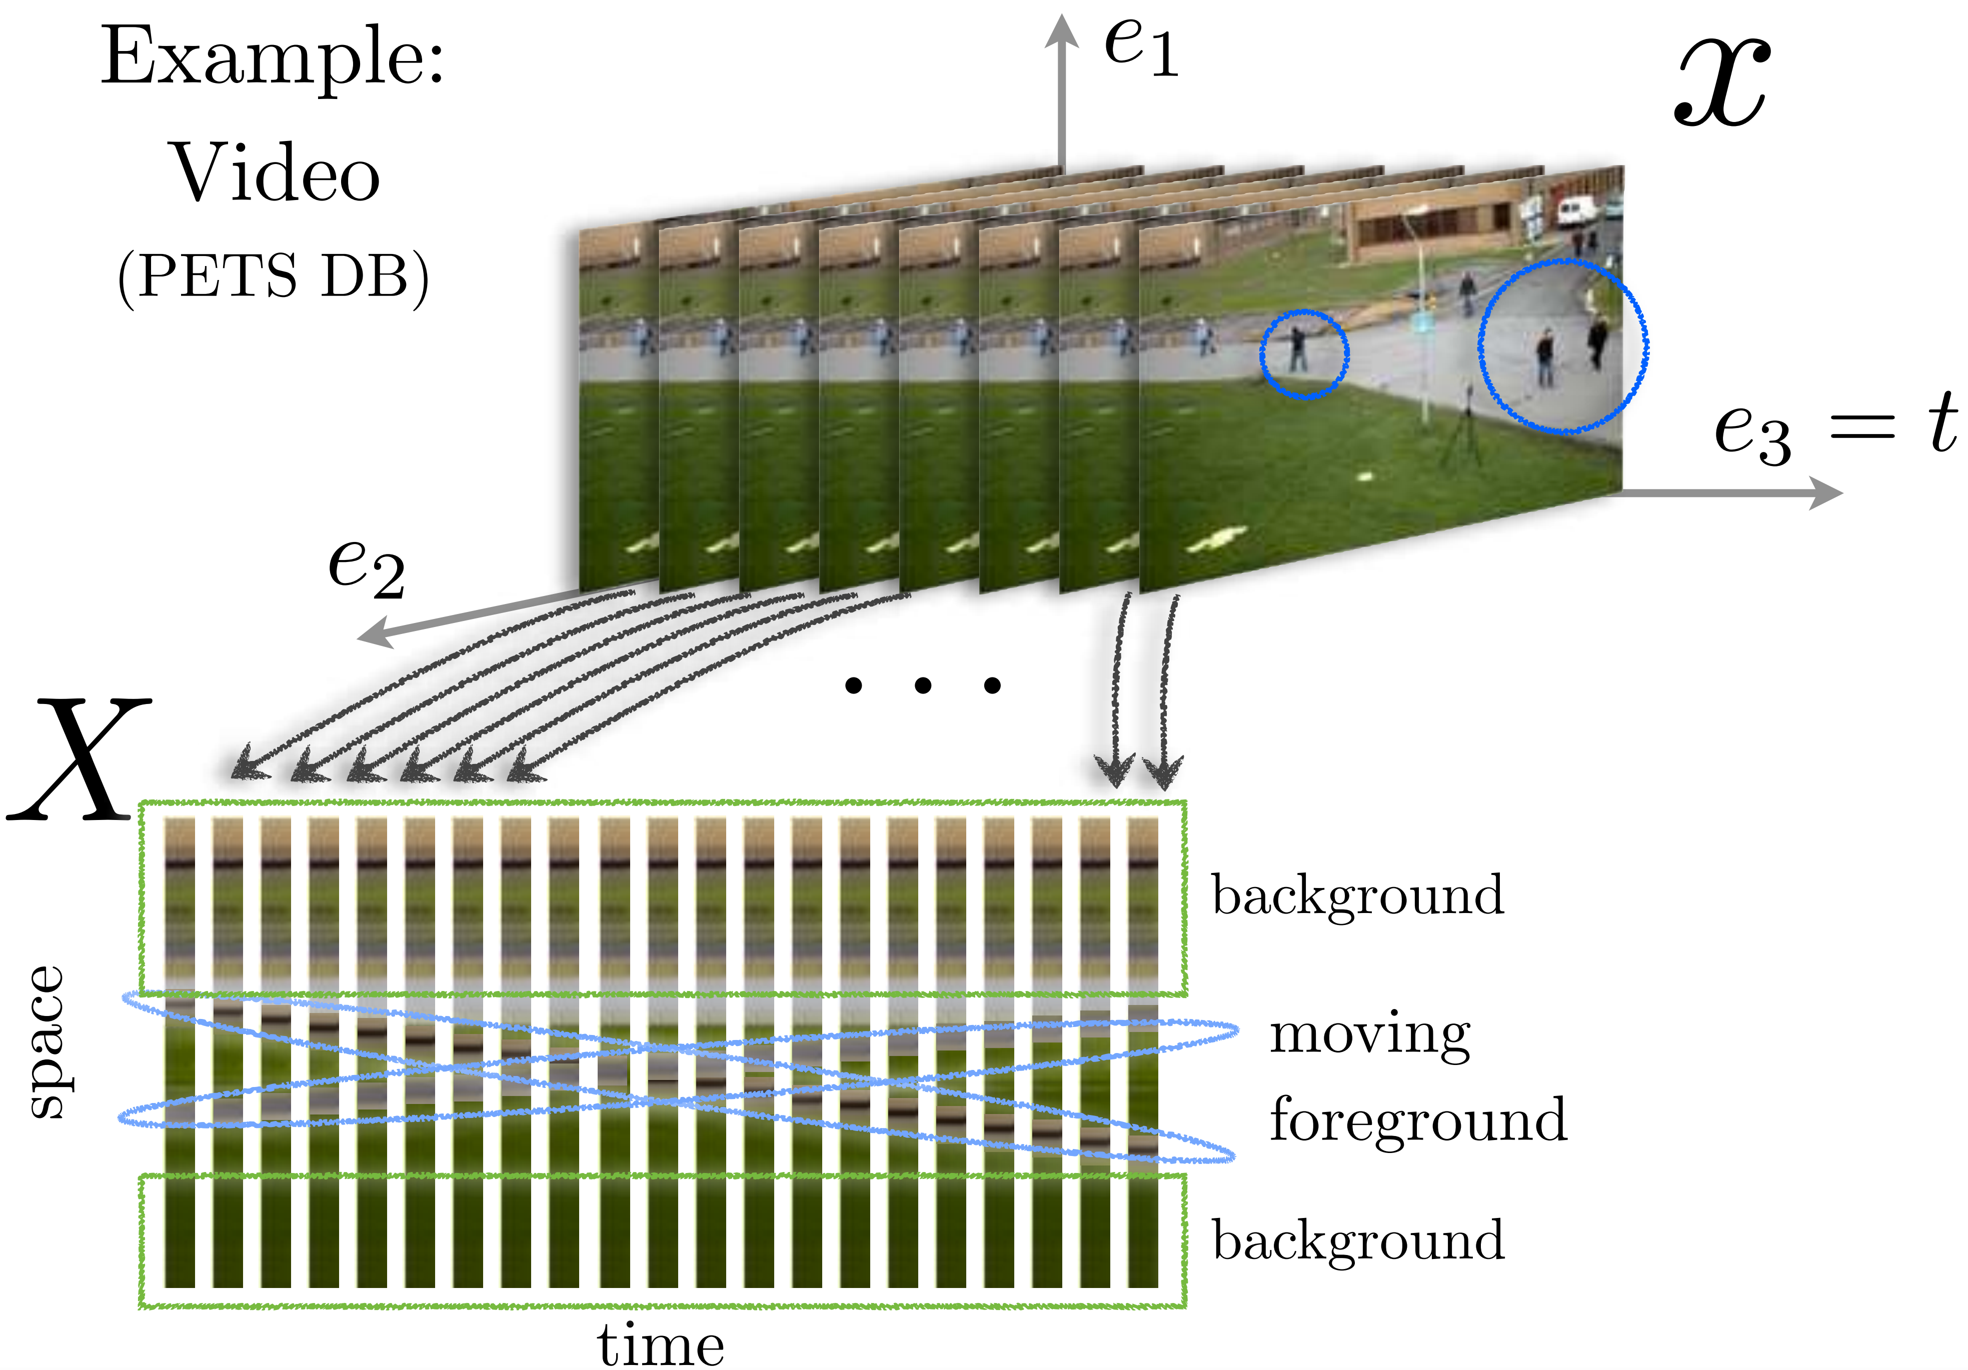

Other "informative" models

2. Low-rank models in high dimensions

Example:

Video

(PETS DB)

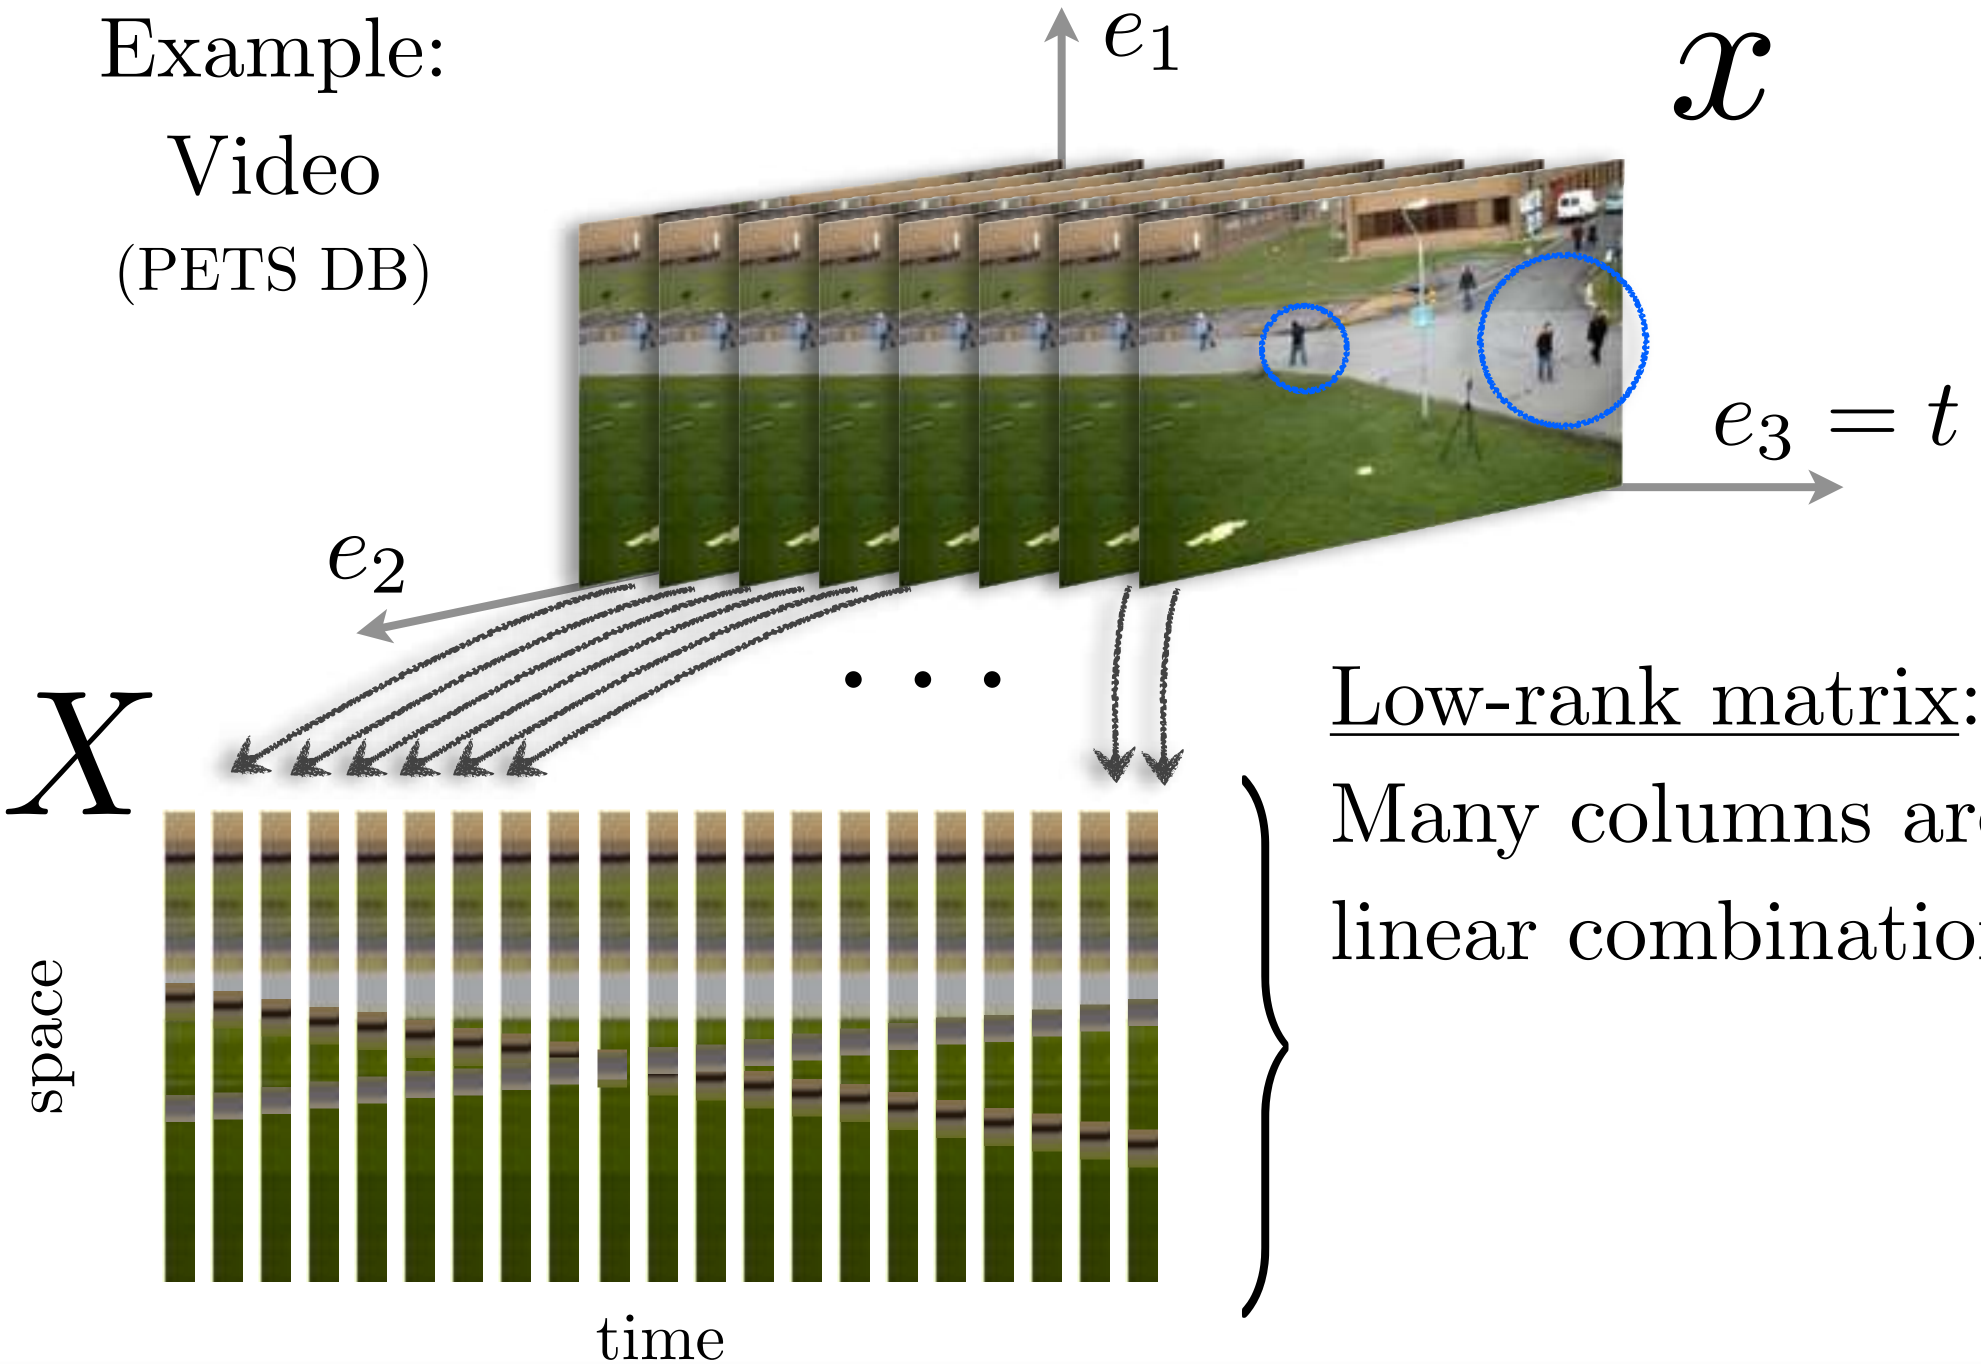

Low-rank matrix:

Many columns are (almost) linear combination of others

Other "informative" models

2. Low-rank models in high dimensions

Example:
Video
(PETS DB)

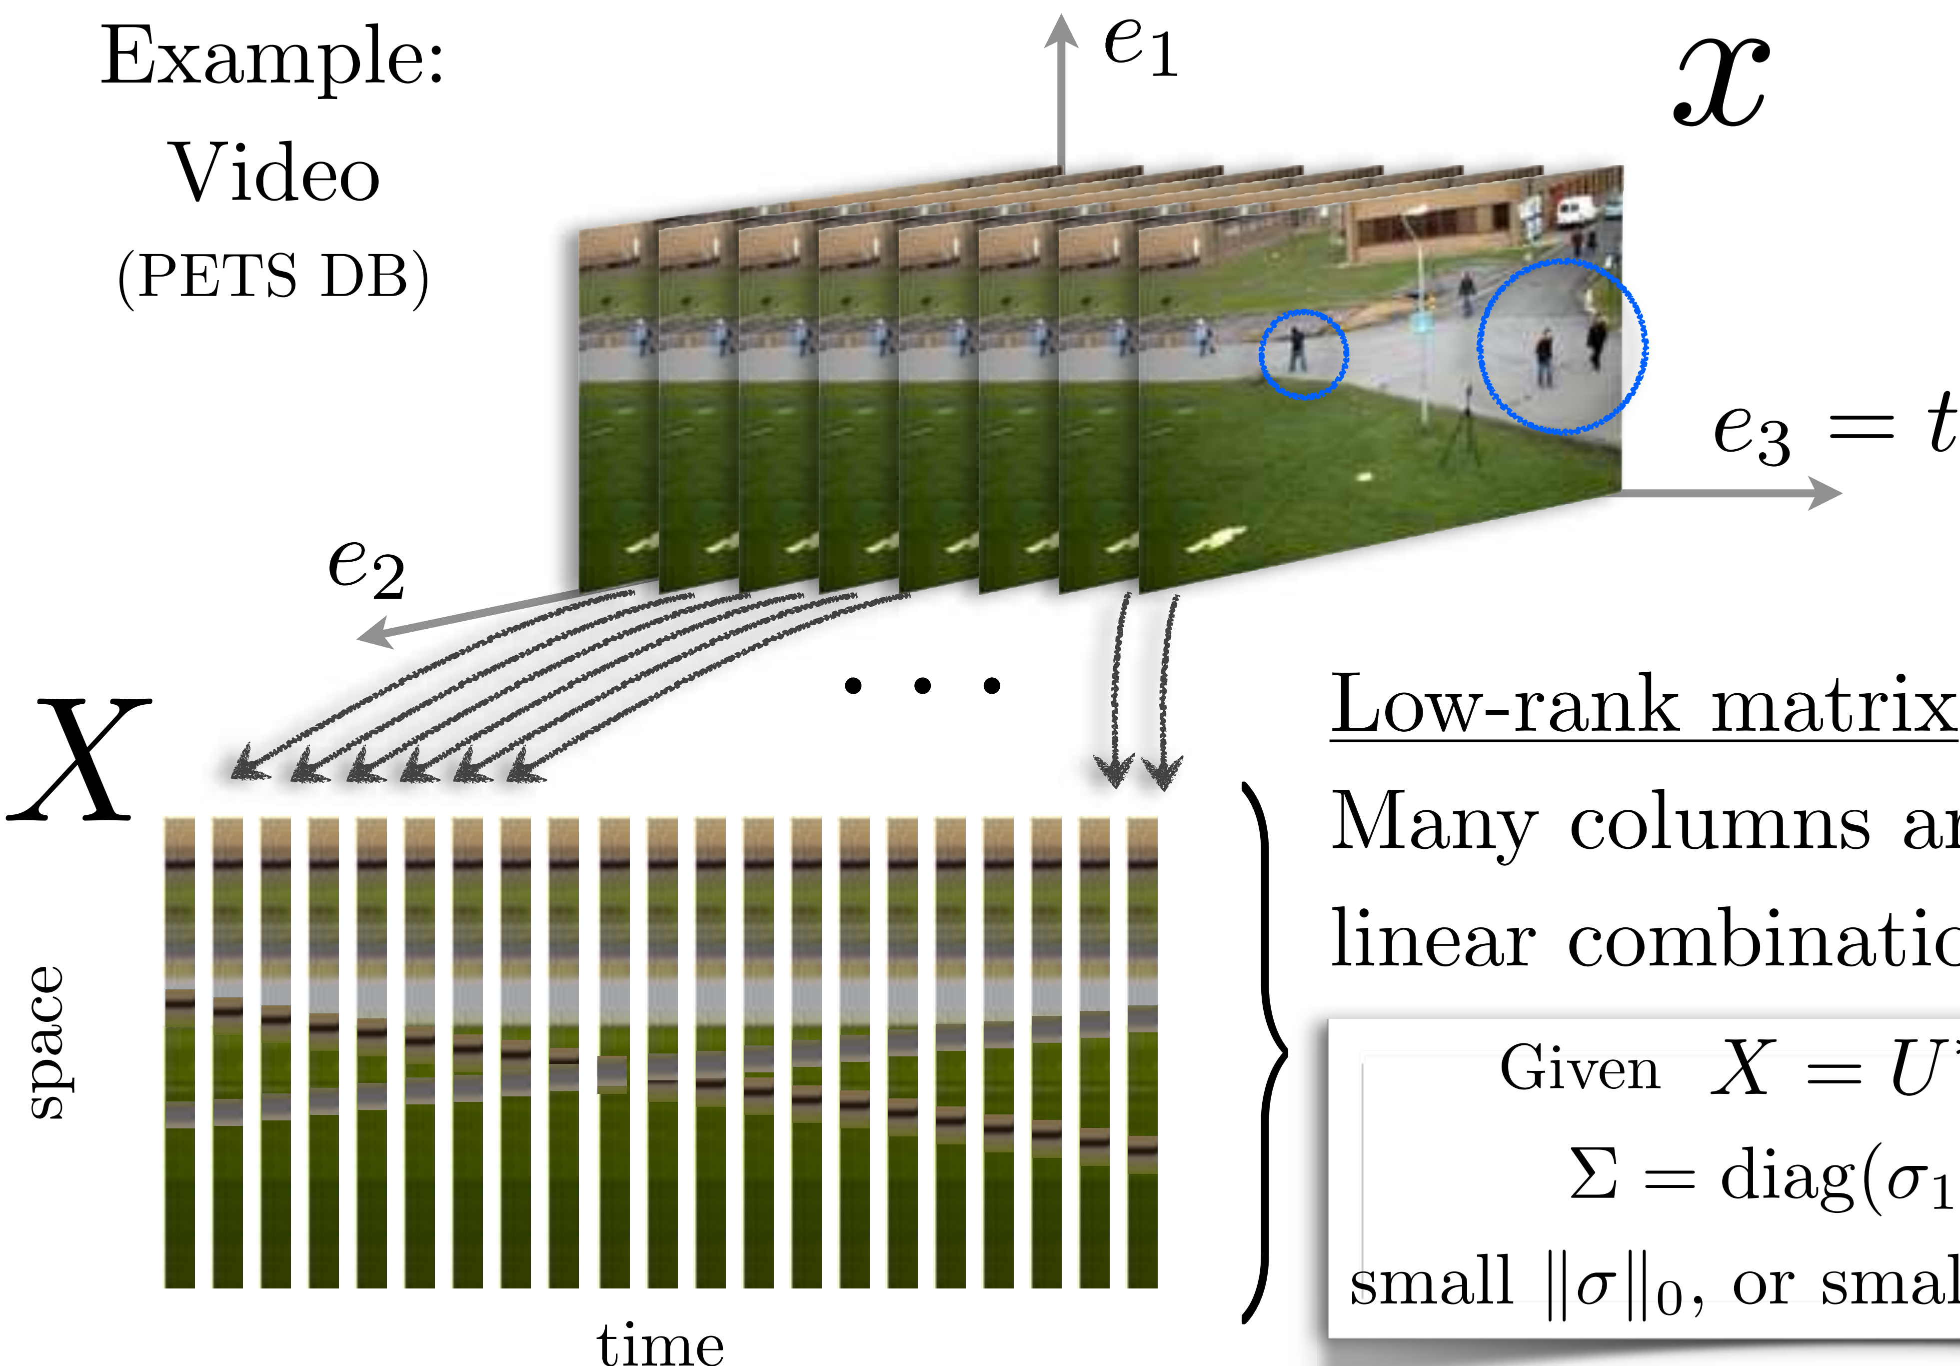

Low-rank matrix:
Many columns are (almost) linear combination of others

Given $X = U^* \Sigma V$ (SVD)

$\Sigma = \mathrm{diag}(\sigma_1, \cdots, \sigma_N)$

small $\|\sigma\|_0$, or small $\|\sigma\|_1 = \|X\|_*$

General Sparsity Applications

1. Data Compression/Transmission (by definition);
2. Data restoration :
   - Denoising,
   - Debluring,
   - Inpainting, ...

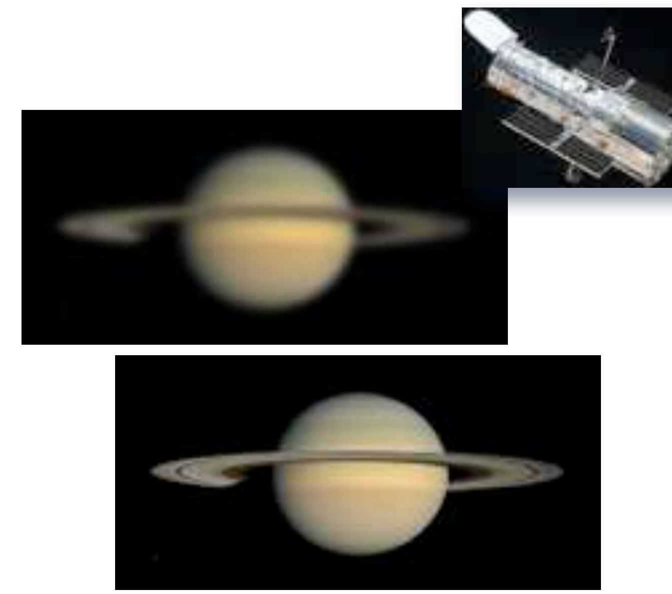

3. Simplified model and interpretation (*e.g.*, in ML)

General Sparsity Applications

1. Data Compression/Transmission (by definition);
2. Data restoration :
   - Denoising,
   - Debluring,
   - Inpainting, ...

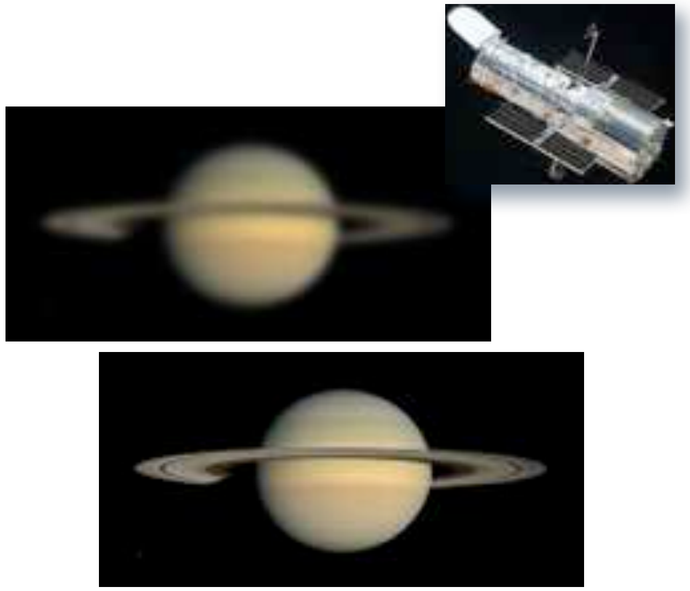

3. Simplified model and interpretation (*e.g.*, in ML)

More generally,
For regularizing (stabilizing) inverse problems
+ Impact on data sampling philosophy ! (see after)

e.g., in Ivo's talk

First Part:

- Sparsity, low-rankness and relatives:
  "From information to structures"

- Compress while you sample:
  "From structures to scrambled sensing"

- and Reconstruct!
  "From scrambled sensing to information"

Sampling with Sparsity

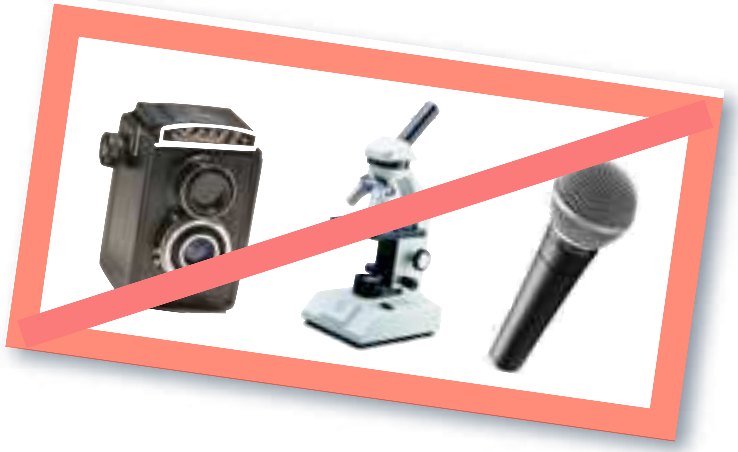

- Paradigm shift:

  "Computer readable" sensing

  + prior information (structures)

Sampling with Sparsity

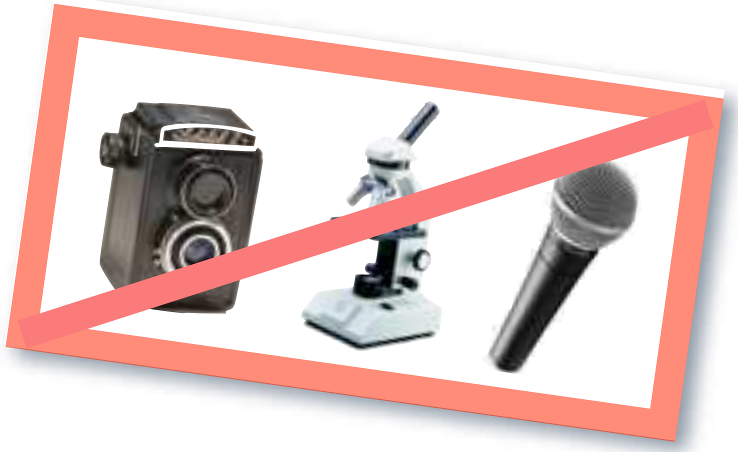

- Paradigm shift:
  "Computer readable" sensing
  + prior information (structures)

World → Signal → Sensing Device → Sensing →  → Signal → Human

Optimized setup: sampling rate $\propto$ information




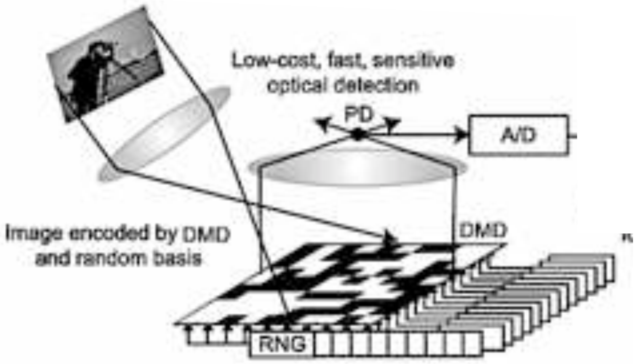

- Examples:
  Radio-Interferometry, Compressed Sensing, MRI, Deflectometry, Seismology, ...

Sampling with Sparsity

but ... non-linear reconstruction schemes!

<u>Regularized inverse problems:</u>

Reconstruct $x \in \mathbb{R}^N$ from $y = \mathrm{Sensing}(x) \in \mathbb{R}^M$

given a sparse model on x.

<u>Examples</u>: Tomography,
frequency/partial observations, ...

Sampling with Sparsity

but ... non-linear reconstruction schemes!

Regularized inverse problems:

Reconstruct $x \in \mathbb{R}^N$ from $y = \text{Sensing}(x) \in \mathbb{R}^M$ given a sparse model on x.

Examples: Tomography, frequency/partial observations, ...

$$x^* = \operatorname*{argmin}_{u \in \mathbb{R}^N} \mathcal{S}(u) \text{ s.t. } \text{Sensing}(u) \approx \text{Sensing}(x)$$

Sparsity metric:

e.g., small $\mathcal{S}(\alpha) = \|\alpha\|_1$ if $u = \Psi\alpha$,
small Total Variation $\mathcal{S}(u) = \|\nabla u\|$

Noise: Gaussian, Poisson, ...

Compressed Sensing

Compressed Sensing

CS in a nutshell:

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

Compressed Sensing

M questions — Sensing method — Signal

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

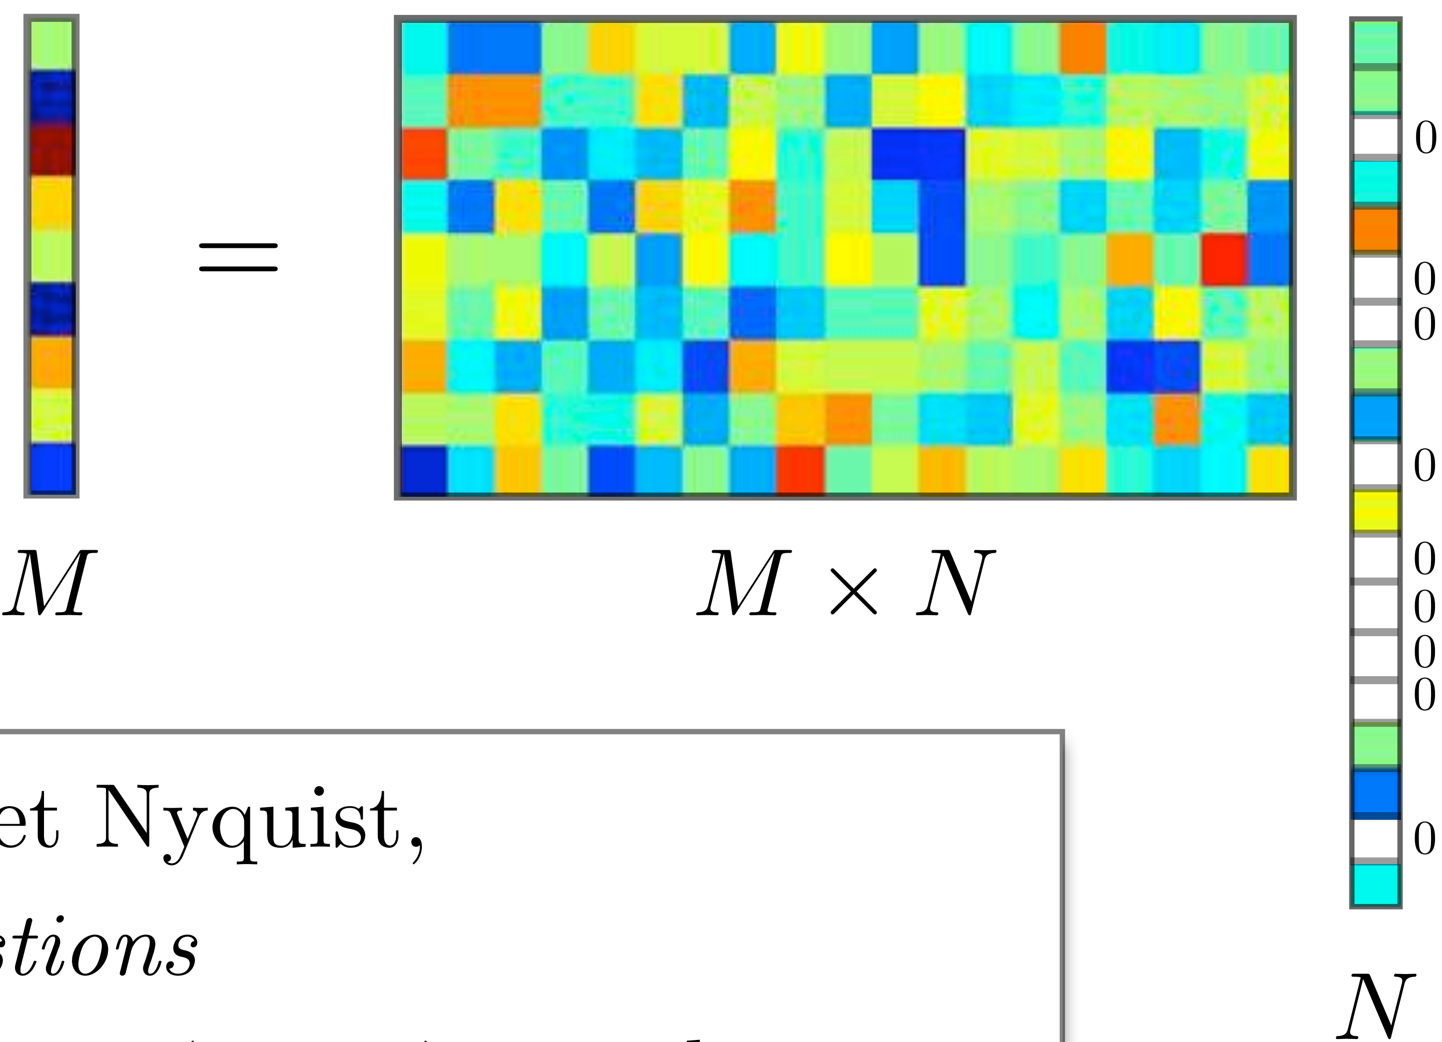

Sparsity Prior ($\Psi = \mathrm{Id}$)

CS in a nutshell:

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

Compressed Sensing

Low complexity signal

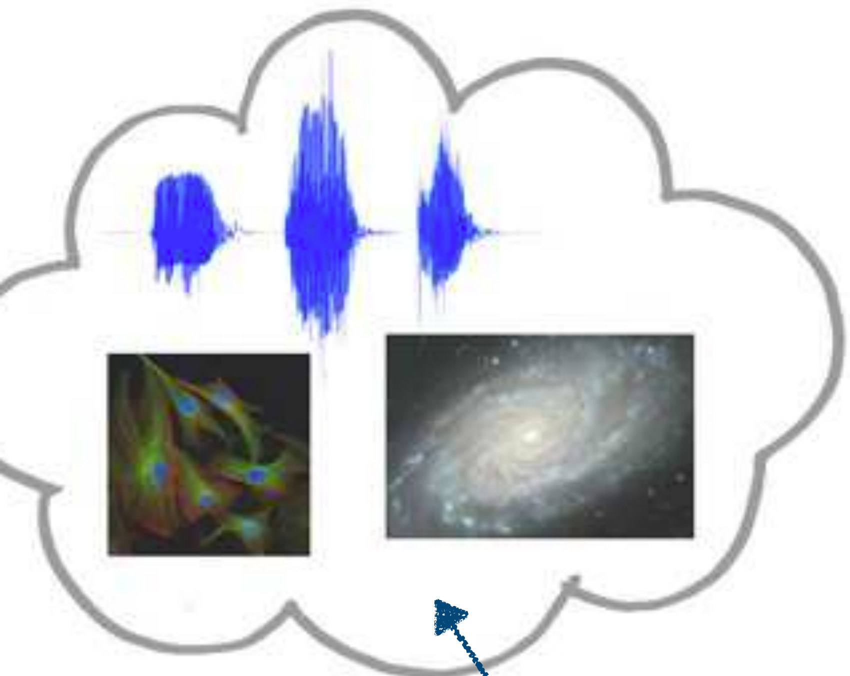

$\boldsymbol{y}$ $\qquad$ $\boldsymbol{\Phi}$ $\qquad$ $\boldsymbol{x}$

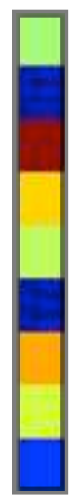

$=$

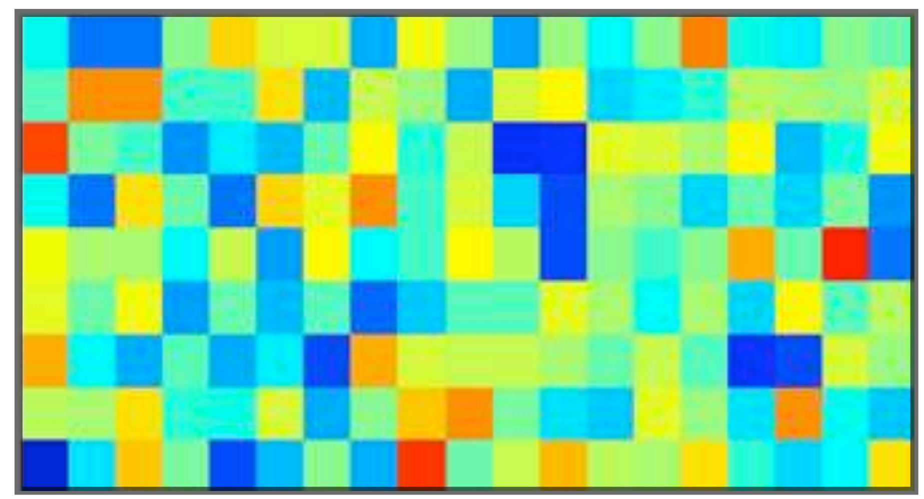

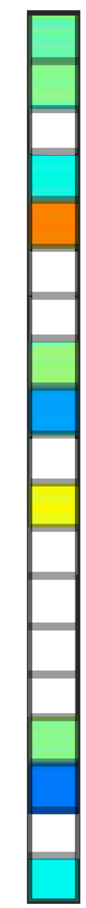

M $\qquad$ $M \times N$ $\qquad$ N

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

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

$\|\boldsymbol{u}\|_0 = \#\{j : u_j \neq 0\}$ **Non-linear reconstruction**

Compressed Sensing

$$\underset{M}{\boldsymbol{y}} = \underset{M \times N}{\boldsymbol{\Phi}} \; \underset{N}{\boldsymbol{x}}$$

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

(relax.)

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

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

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

Compressed Sensing

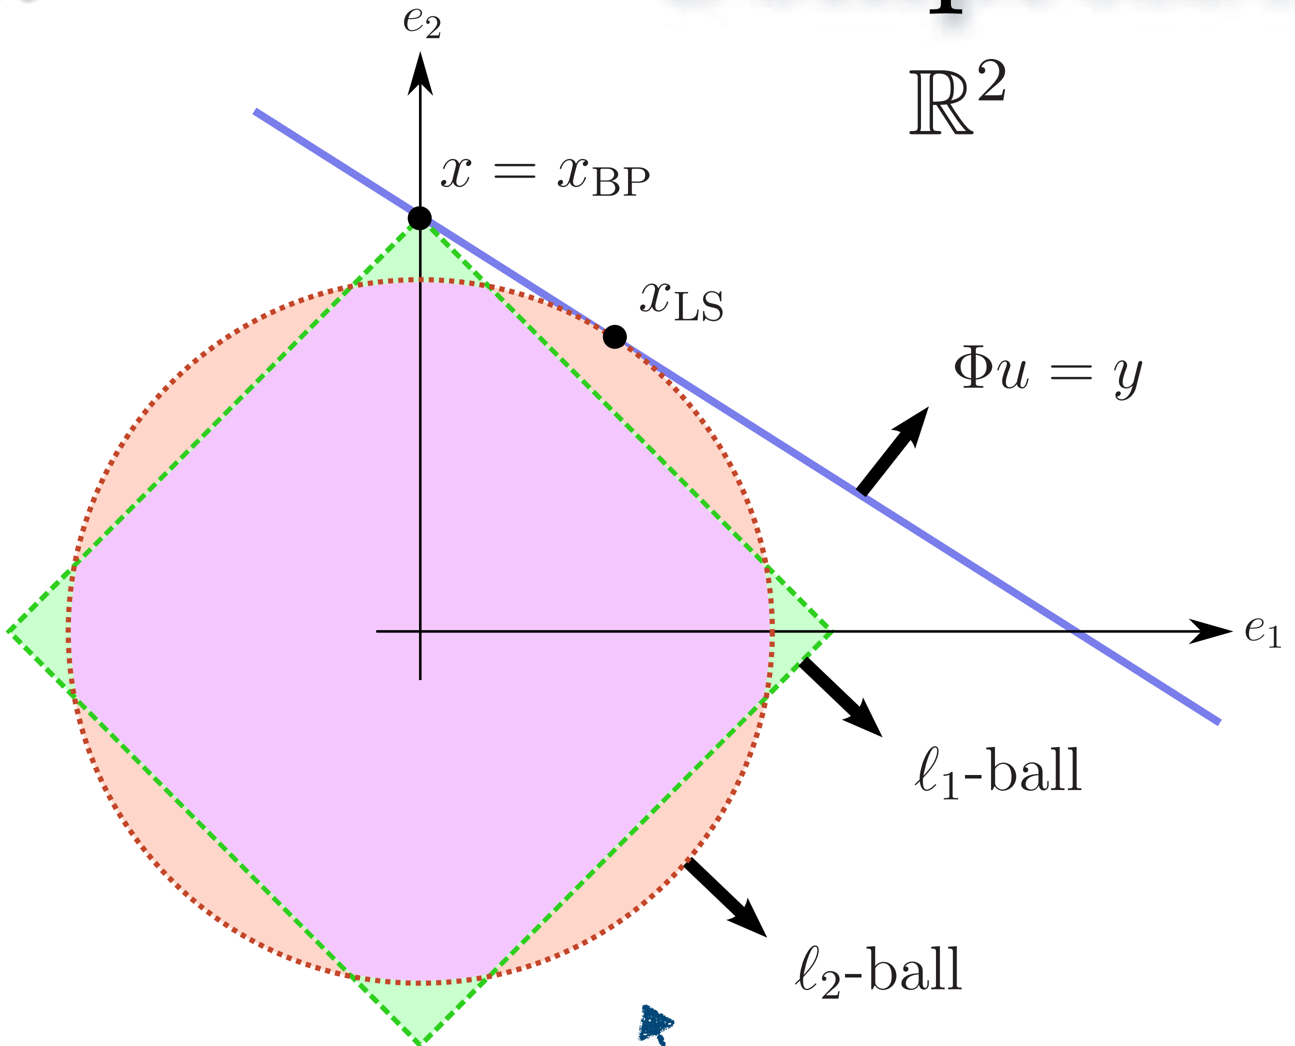

If $\boldsymbol{x}$ is K-sparse and if $\boldsymbol{\Phi}$ well "conditioned" then:

(relax.)

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

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

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

Solvers:

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

Compressed Sensing

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

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

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

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

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

(relax.)

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

if $\delta < 1/\sqrt{2}$ [Cai & Zhang, 14]

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

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

Compressed Sensing

$\exists\ \delta \in (0,1)$ <u>Restricted Isometry Property</u>

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

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

any subset of $2K$ columns

is an *isometry*

<u>Examples:</u>

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

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

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

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

then:

(relax.)

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

if $\delta < 1/\sqrt{2}$ [Cai & Zhang, 14]

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

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

Compressed Sensing

$\|\boldsymbol{y} - \boldsymbol{\Phi u}\| \leqslant \epsilon$ (if noise)

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

If $\quad \boxed{\begin{array}{l} \underline{M = O(K \ln N/K)} \ll N \\ \boldsymbol{\Phi} \in \mathbb{R}^{M \times N}, \ \Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0,1) \end{array}}$

Robustness: *vs* sparse deviation + noise.

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

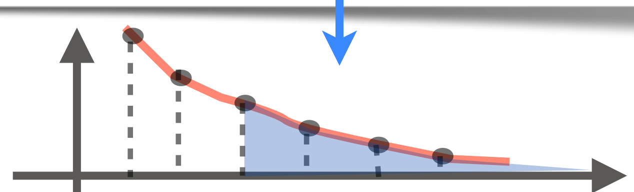

… in summary, CS is

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

$\boldsymbol{x}$

or $\boldsymbol{x}^* \simeq \boldsymbol{x}$

if noise, quantization,
non-linearities

HEAVY RECONSTRUCTION

M questions — Sensing method — Signal

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

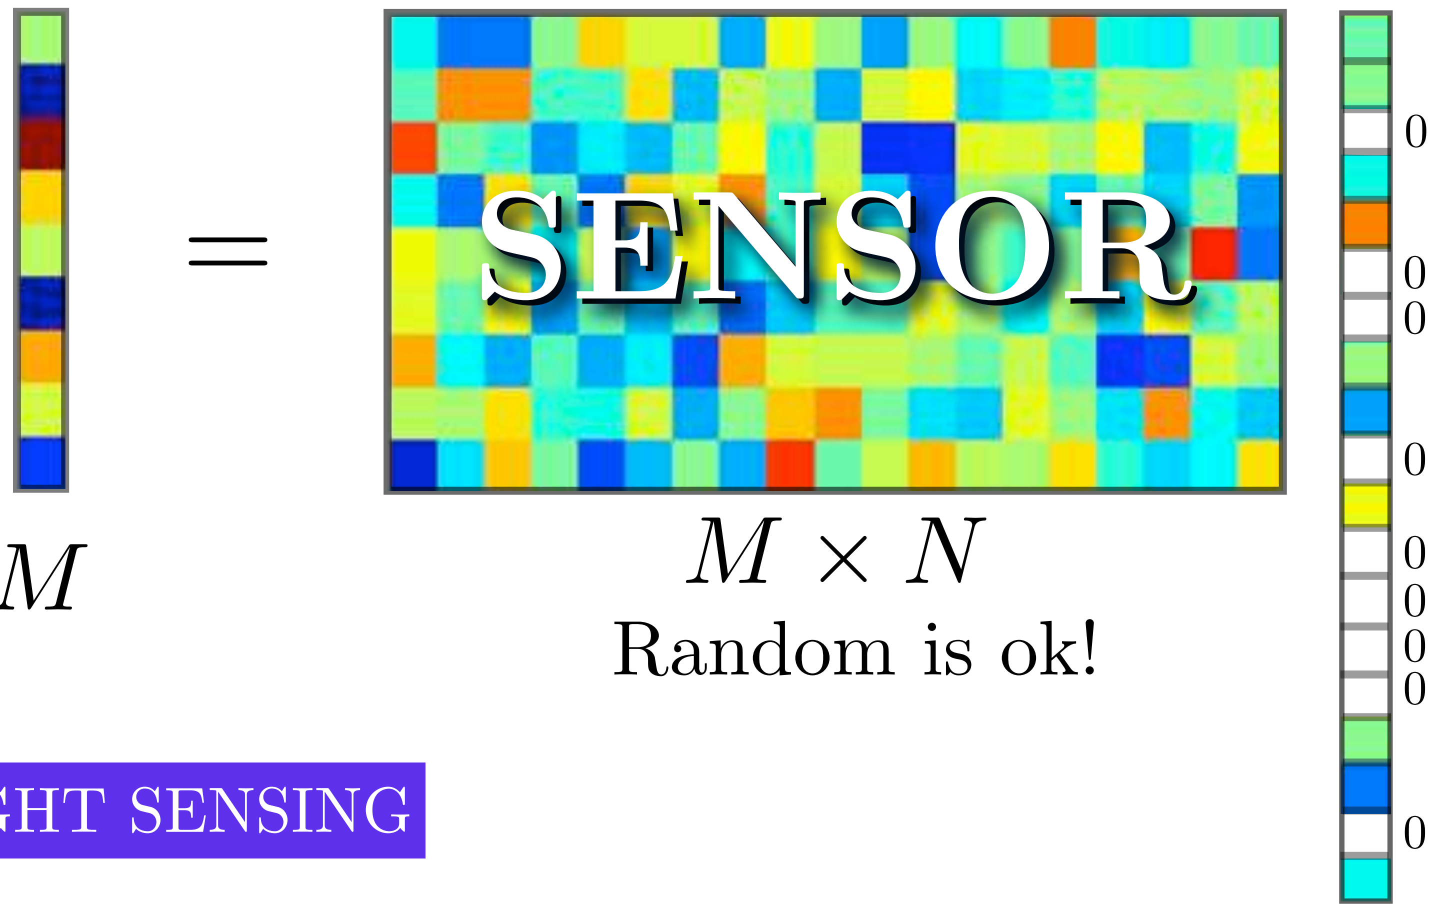

M

$M \times N$

Random is ok!

N

LIGHT SENSING

Ask *few* (linear) *questions*
about your informative (sparse) signal,
and recover it *differently* (non-linearly)"

Candès, Romberg, Tao, 2006
Donoho, 2006
1-Pixel Camera, Wakin et al., 2006

First Part:

- Sparsity, low-rankness and relatives:
  "From information to structures"

- Compress while you sample:
  "From structure to scrambled sensing"

- and Reconstruct!
  "From scrambled sensing to information"

(very broad and active field ... just one slide)

Reconstruct? (just one slide)

- For solving:

e.g., sparsity, low-rank, TV, ...

e.g., L2/L1 distance, robust to Gaussian/Poisson noise, ...

$$\boxed{x^* = \operatorname*{argmin}_{u \in \mathbb{R}^N} \ \mathcal{S}(u) \text{ s.t. } \mathrm{Sensing}(u) \approx \mathrm{Sensing}(x)}$$

many possibilities/solvers ...

Reconstruct? (just one slide)

- For solving:

e.g., sparsity, low-rank, TV, ...

e.g., L2/L1 distance, robust to Gaussian/Poisson noise, ...

$$\boxed{x^* = \operatorname*{argmin}_{u \in \mathbb{R}^N} \ \mathcal{S}(u) \text{ s.t. } \mathrm{Sensing}(u) \approx \mathrm{Sensing}(x)}$$

many possibilities/solvers ...

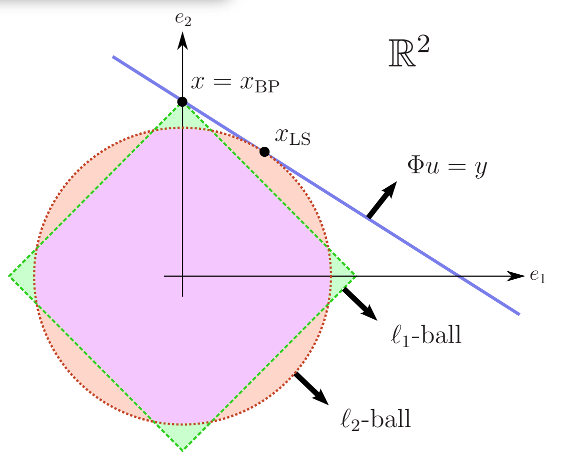

- Convex optimization: tons of toolboxes
  - SPGL1, L1Magic, (F)ISTA, ADMM, ...
  - Proximal algorithms (see also B. Goldluecke's part)
- Iterative (*greedy*) methods:
  - matching pursuit and relatives (OMP)
  - iterative hard thresholding, CoSAMP, SP, smoothed L0, ...
  - Approximate Message Passing Algorithms, Bayesian, ...

Second Part: a few realizations

- Compressive imaging appetizer:
  The Rice single pixel camera
- Other case studies:
  - Radio-interferometry and aperture synthesis
  - Hyperspectral CASSI imaging
  - Highspeed Coded Strobing Imaging

Second Part: a few realizations

- Compressive imaging appetizer:
  The Rice single pixel camera
- Other case studies:
  - Radio-interferometry and aperture synthesis
  - Hyperspectral CASSI imaging
  - Highspeed Coded Strobing Imaging

Rice Single-pixel Camera

[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 2008]

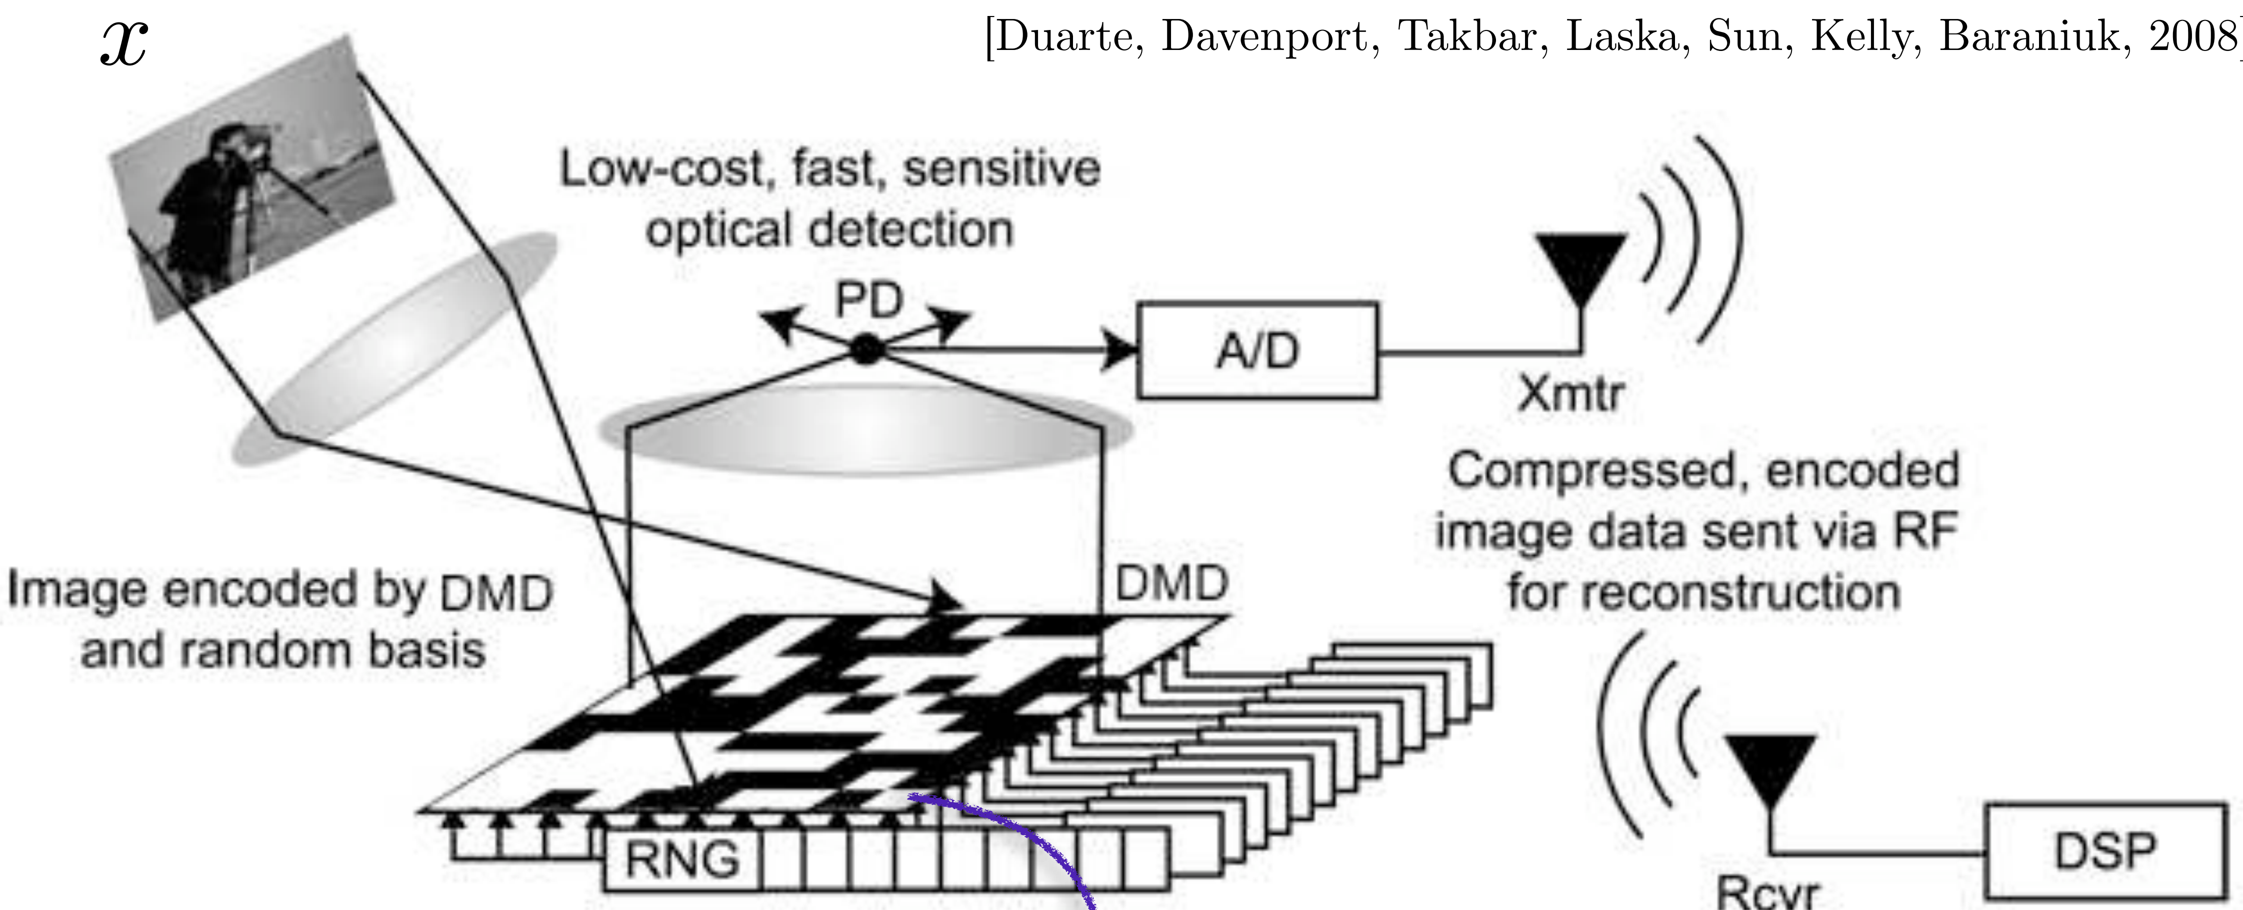

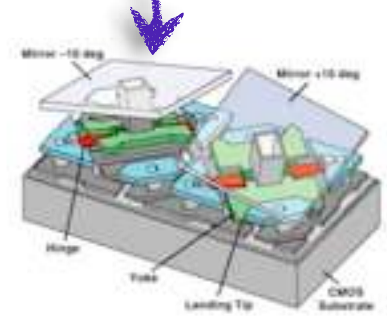

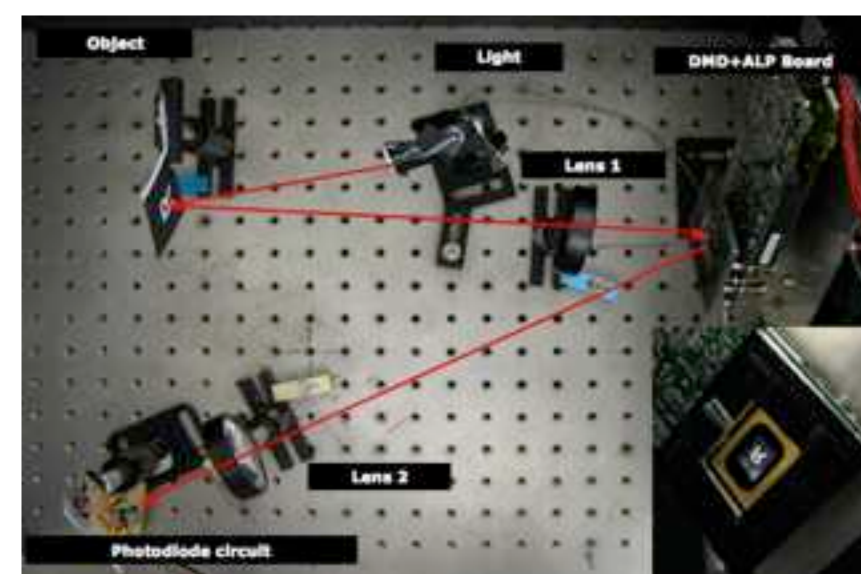

Rice Single-pixel Camera

x

[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 2008]

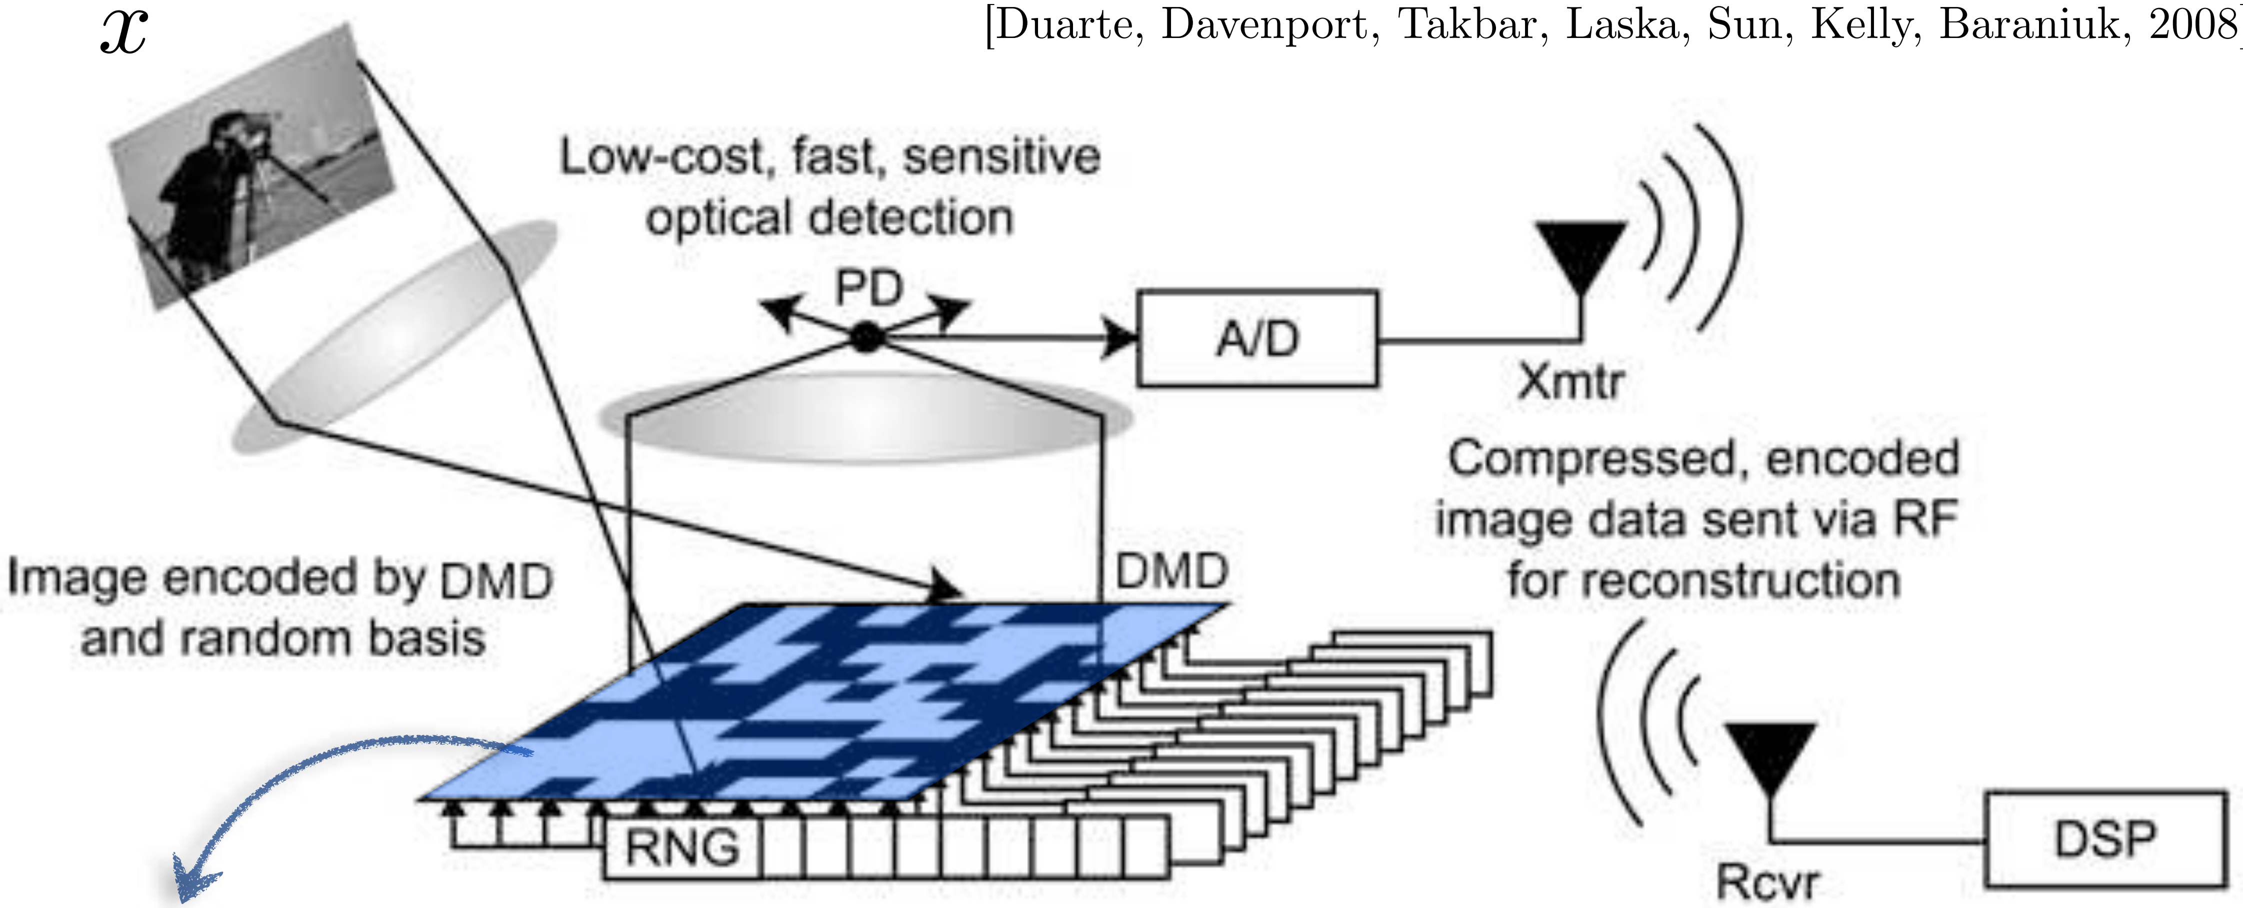

$$\varphi_{ji}\, x_i, \quad \varphi_{ji} \in \{0, 1\}$$

j^{th} random pattern $\varphi_j \in \mathbb{R}^N$

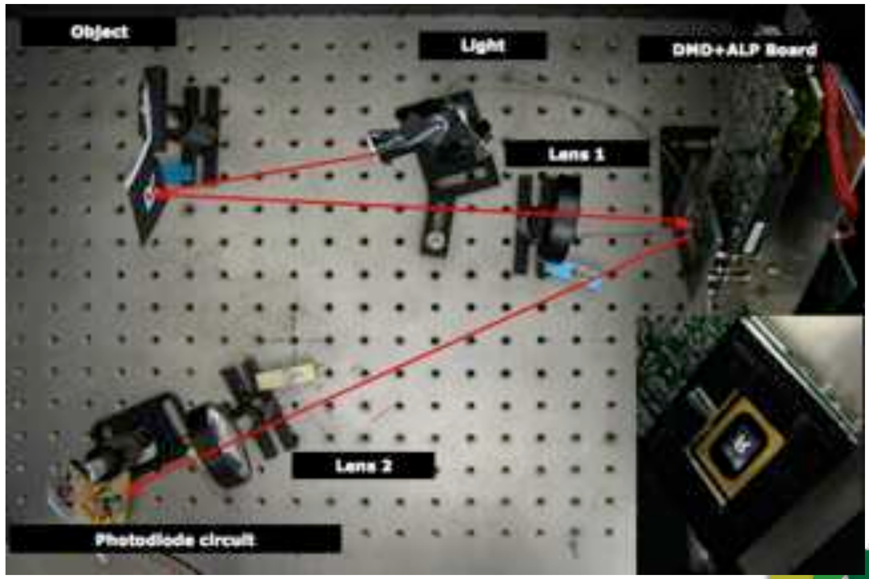

Rice Single-pixel Camera

[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 2008]

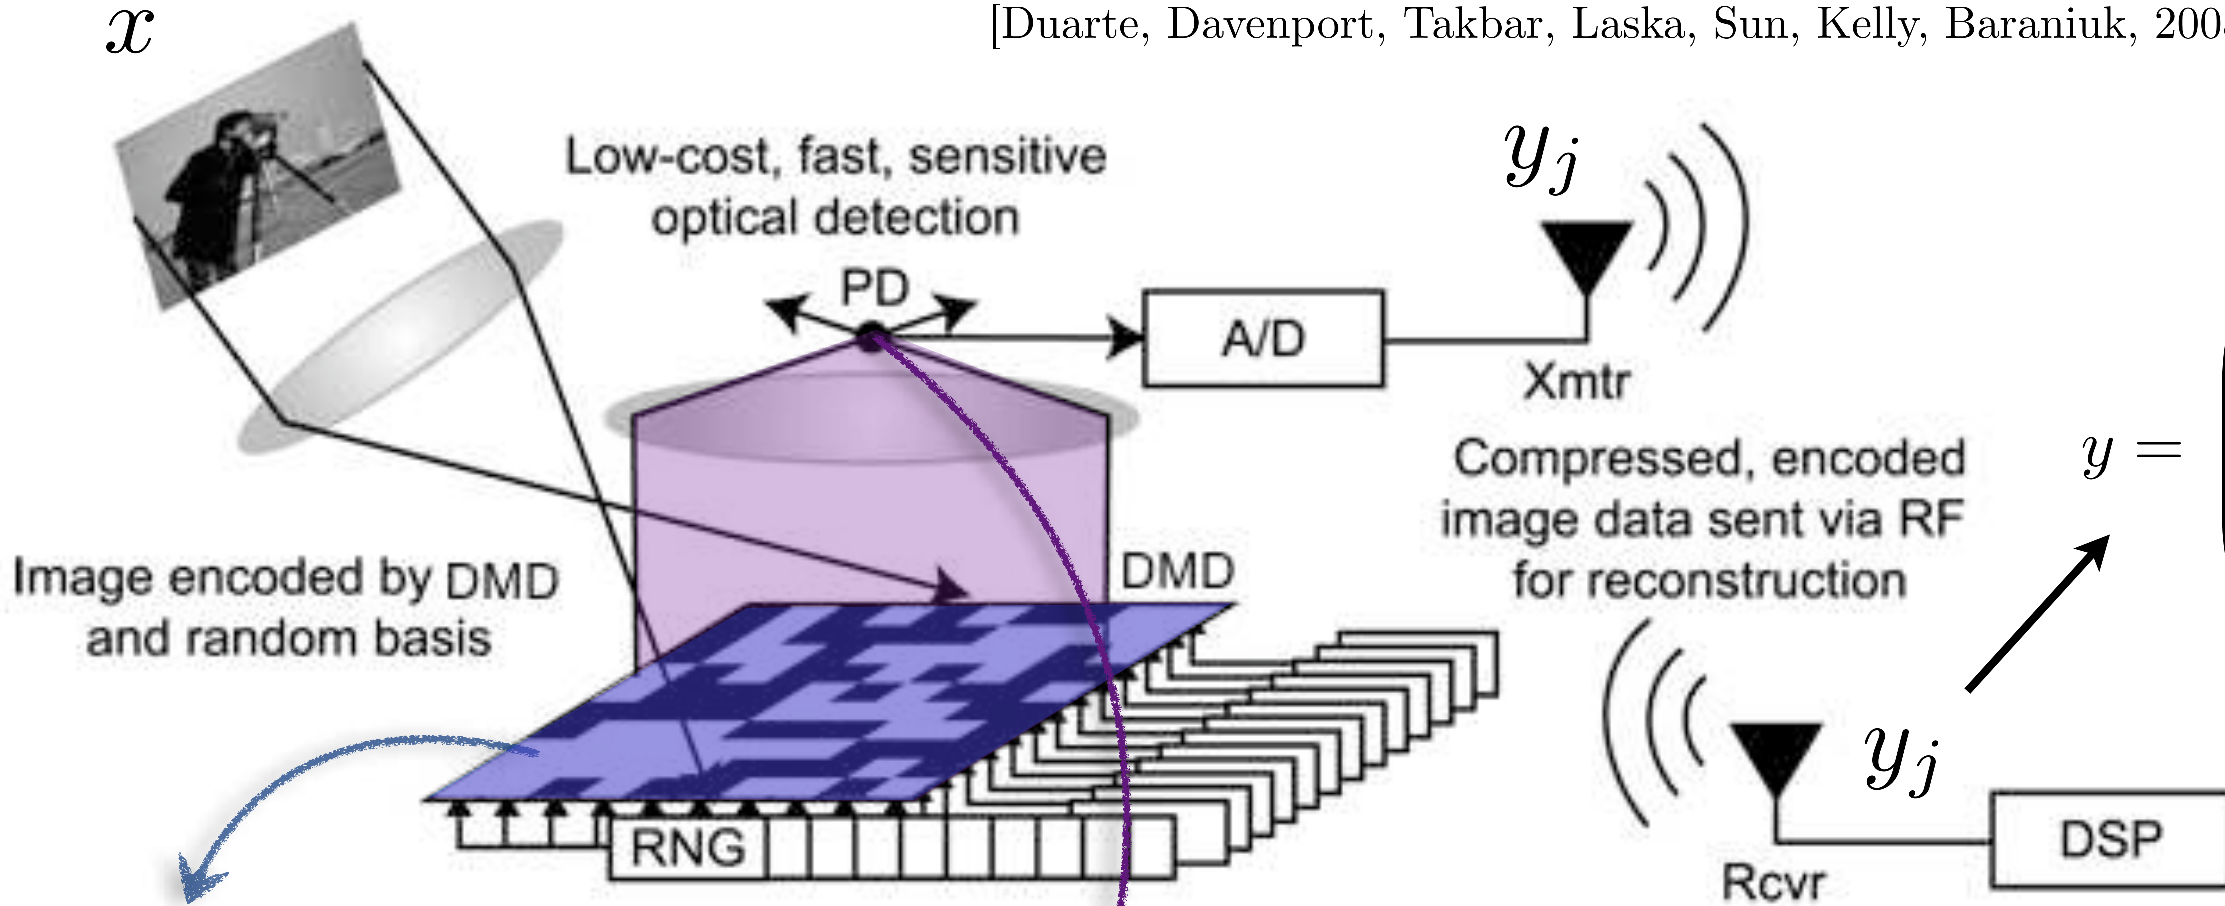

$$y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_M \end{pmatrix}$$

$\varphi_{ji}\, x_i, \quad \varphi_{ji} \in \{0, 1\}$

j^{th} random pattern $\varphi_j \in \mathbb{R}^N$

Optical sum

$$y_j = \sum_i \varphi_{ji}\, x_i$$

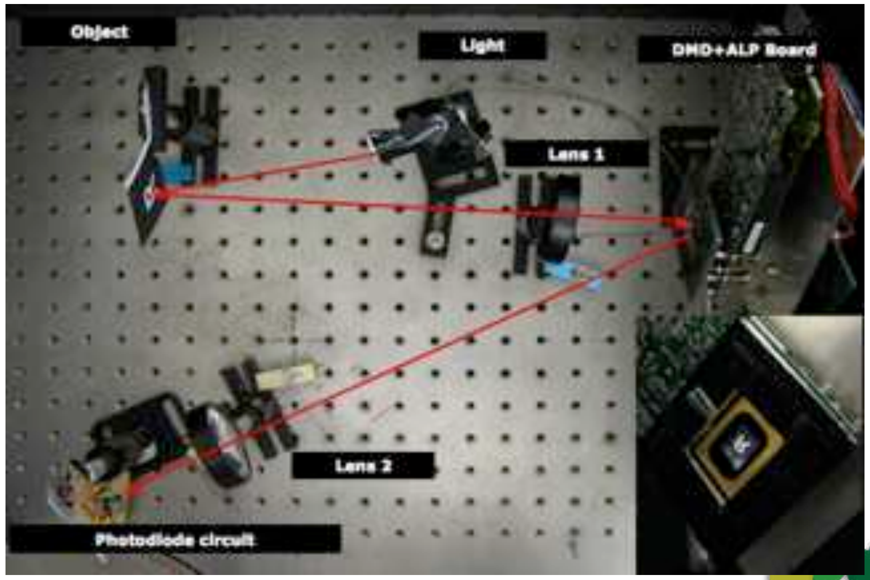

Rice Single-pixel Camera

[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 2008]

x

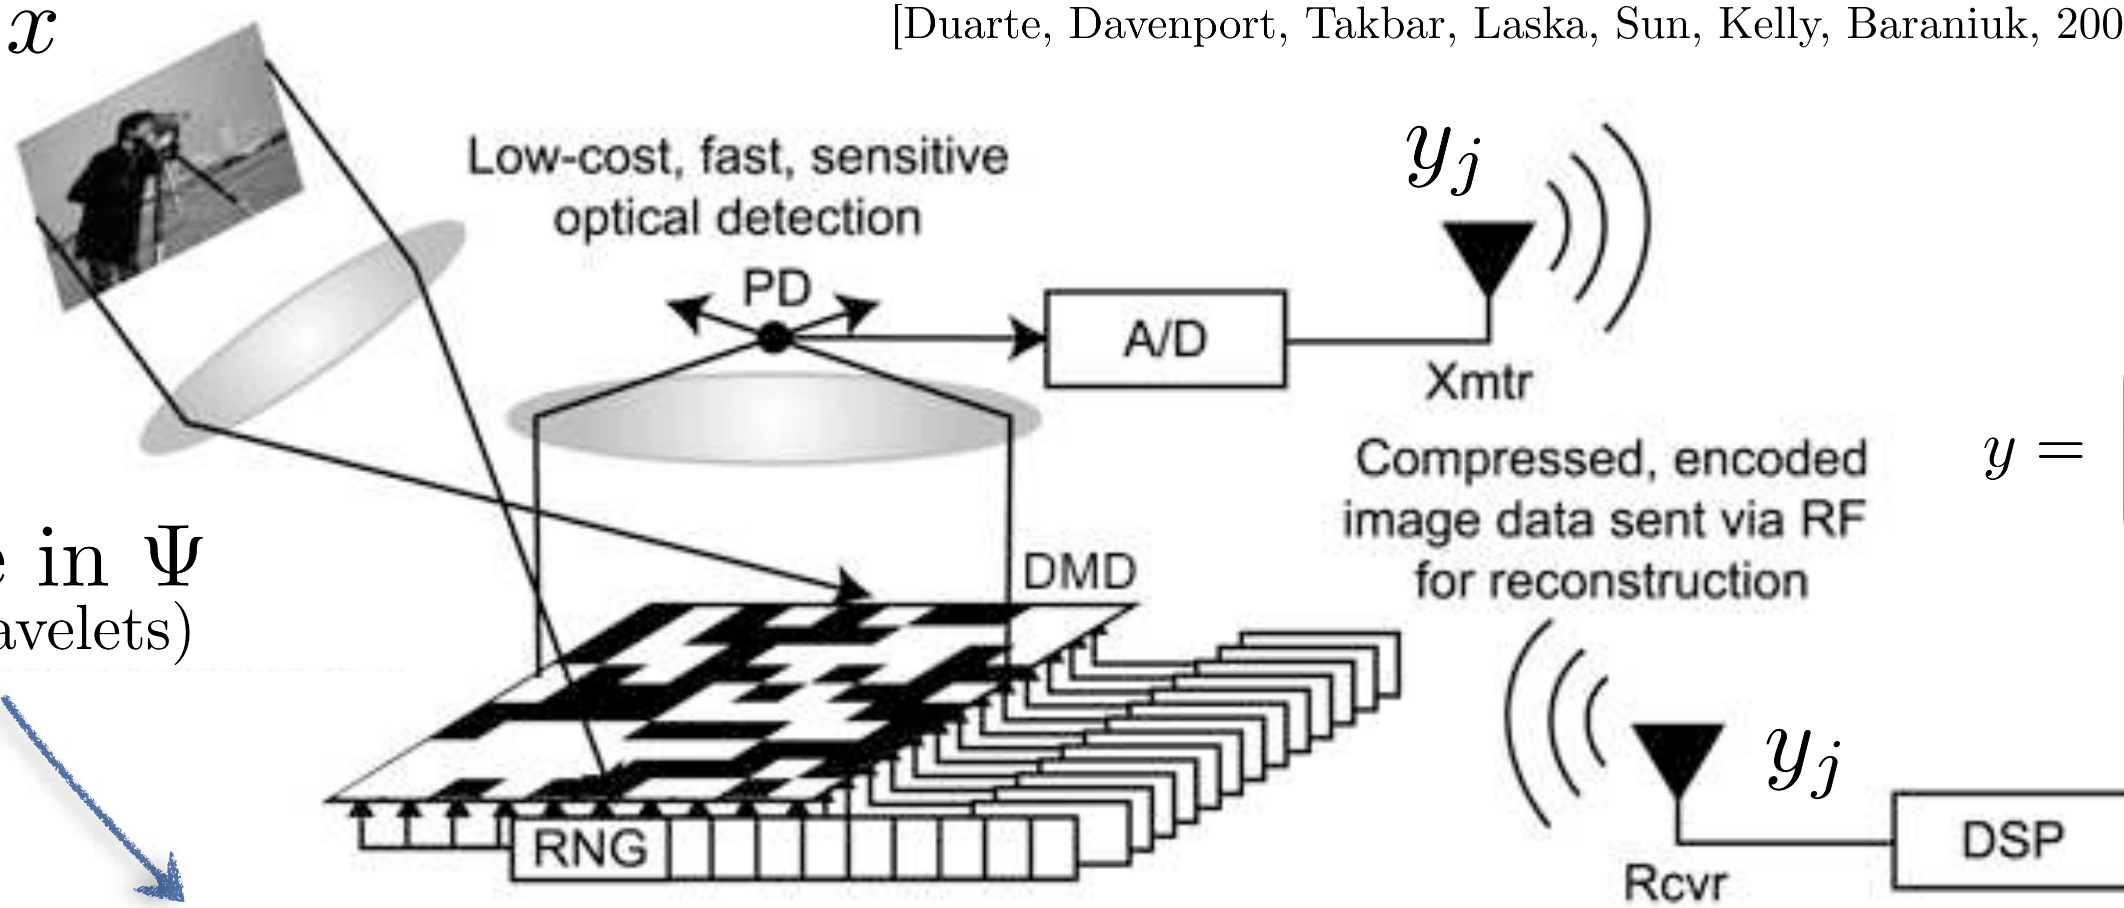

$$y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_M \end{pmatrix}$$

sparse in Ψ
(e.g., wavelets)

$$x^* = \operatorname{argmin} \|\alpha\|_1 \text{ s.t. } \Phi\Psi\alpha = y$$

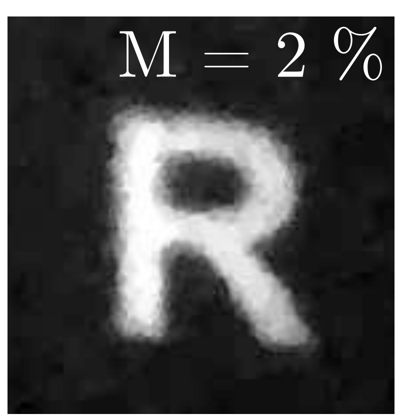

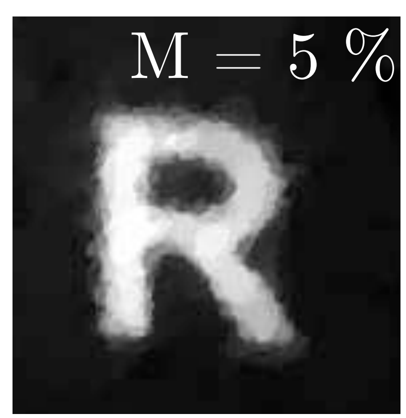

Rice Single-pixel Camera

[Duarte, Davenport, Takbar, Laska, Sun, Kelly, Baraniuk, 2008]

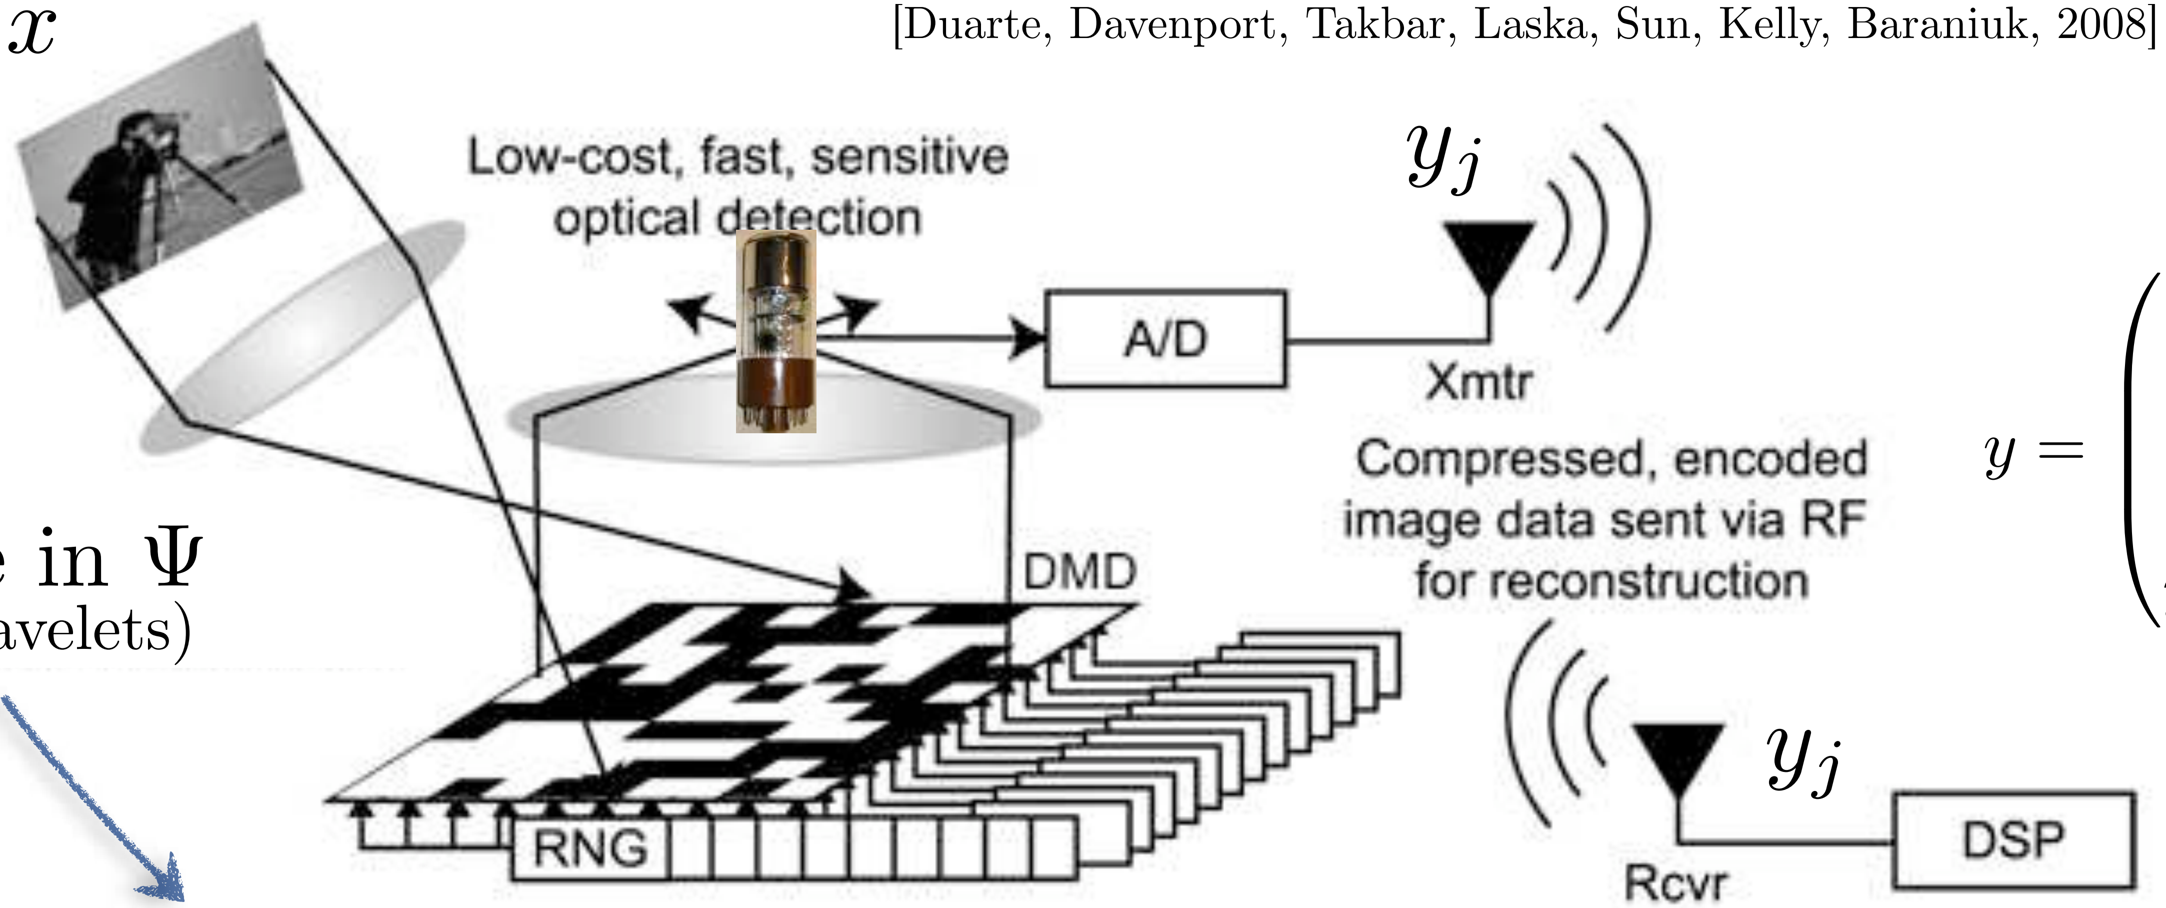

$$y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_M \end{pmatrix}$$

sparse in Ψ
(e.g., wavelets)

$$x^* = \operatorname{argmin} \|\alpha\|_1 \text{ s.t. } \Phi\Psi\alpha = y$$

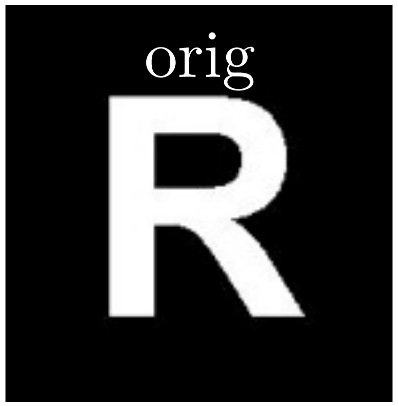

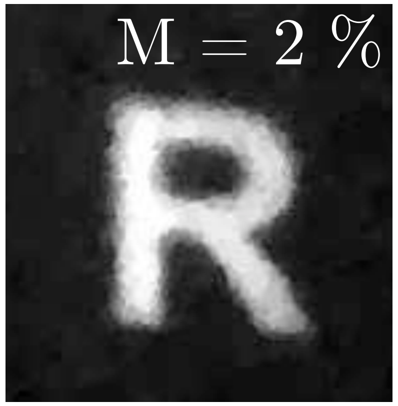

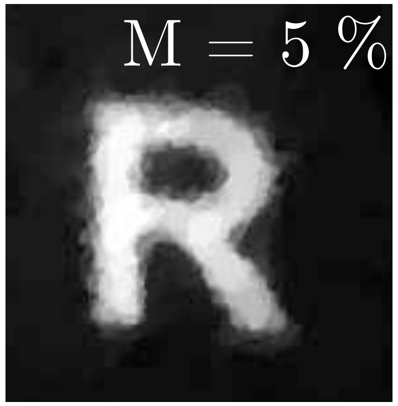

Proof of concept **but** PD is unique!

Specific imaging application:

high energy photons with
single costly photomultiplier

Second Part:

- Compressive imaging appetizer:
  The Rice single pixel camera
- Other case studies:
  - Radio-interferometry and aperture synthesis
  - Hyperspectral CASSI imaging
  - Highspeed Coded Strobing Imaging

Aperture Synthesis in Radio-Astronomy

Observation of the sky in radio frequency

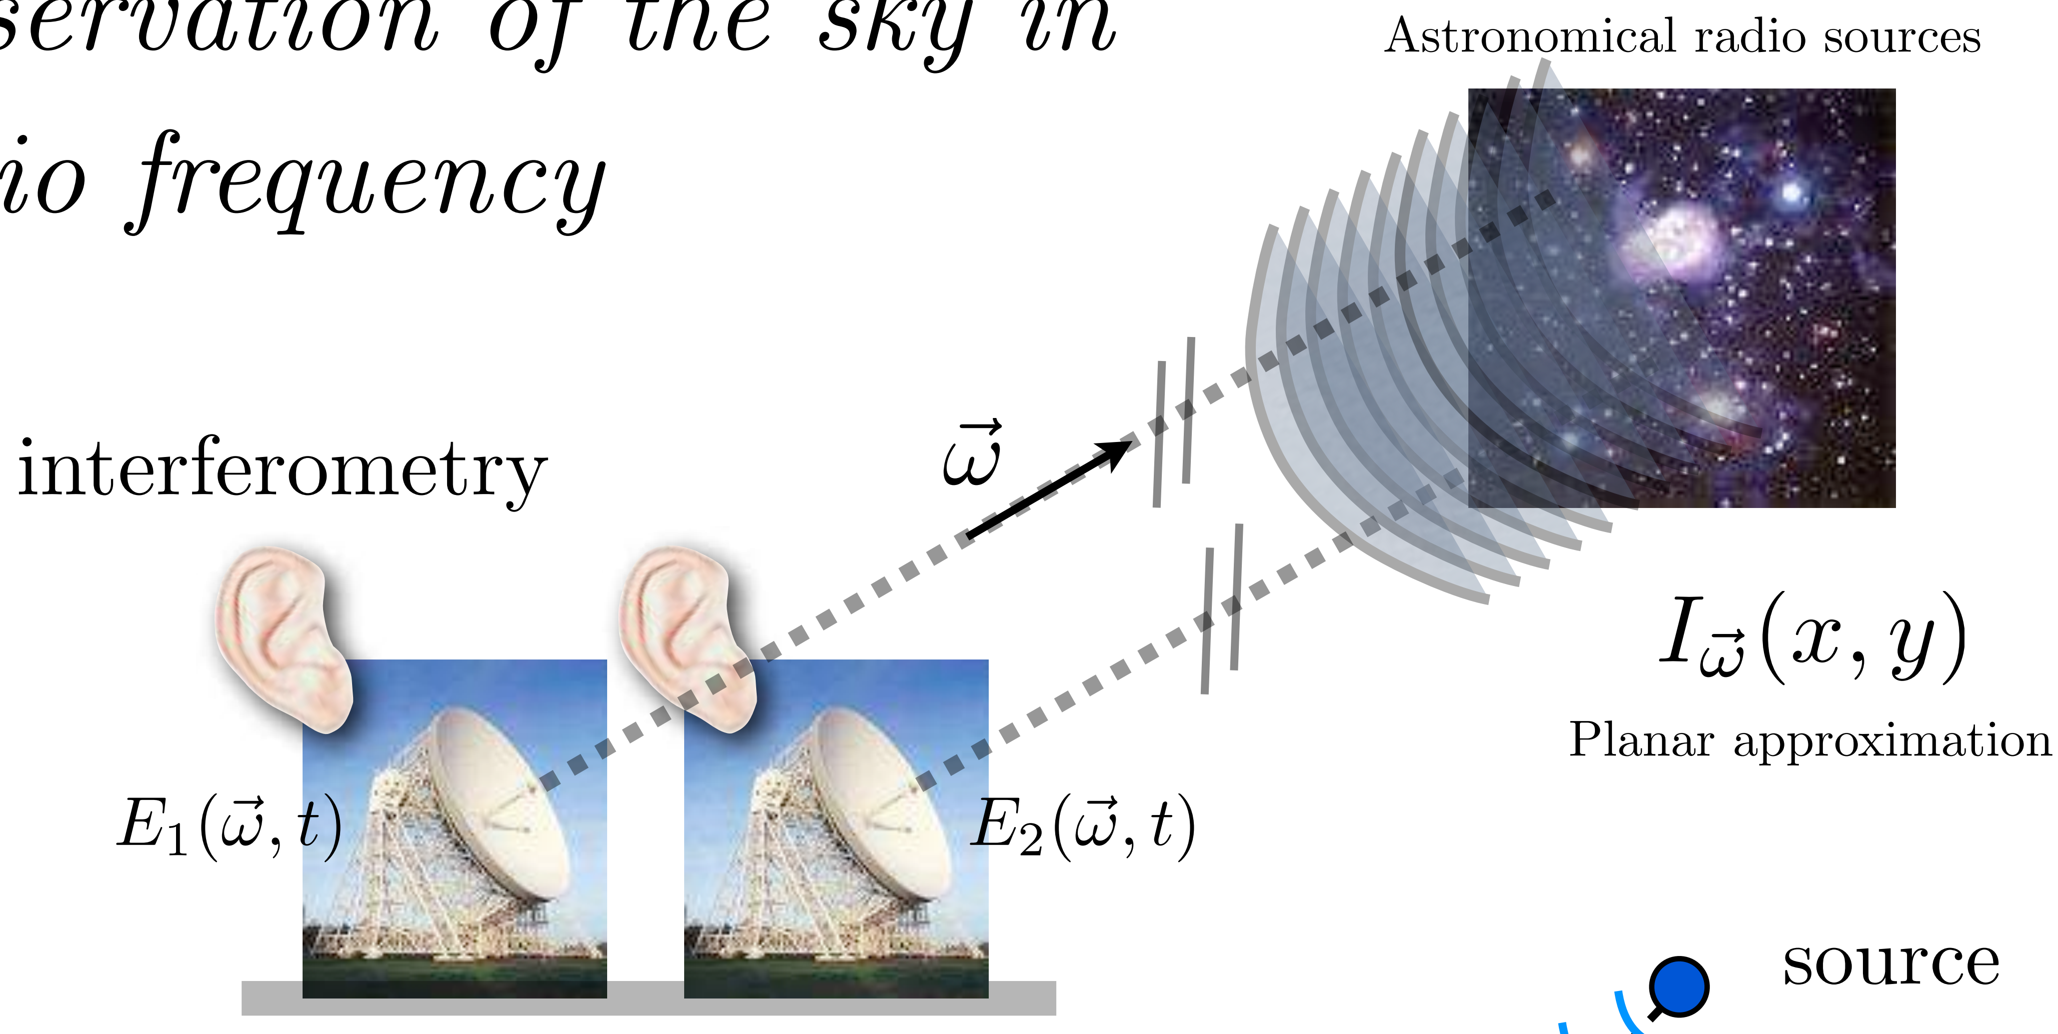

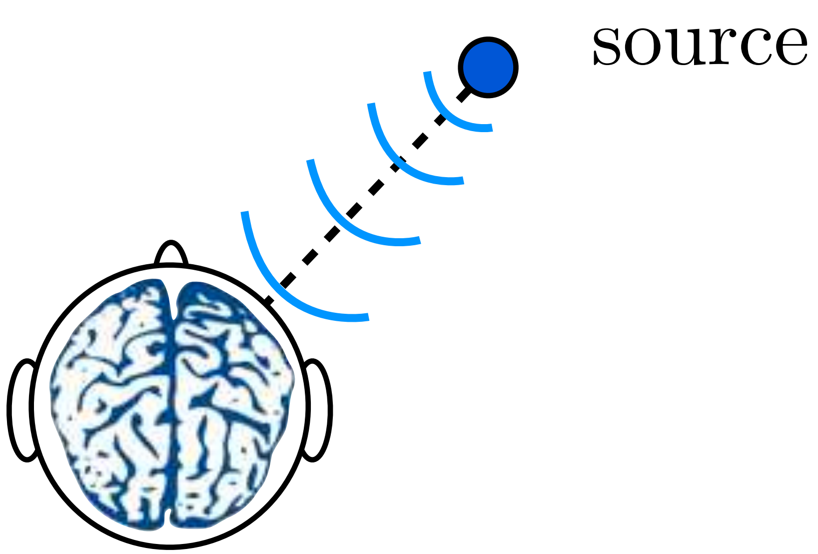

Aperture Synthesis in Radio-Astronomy

Observation of the sky in radio frequency

interferometry

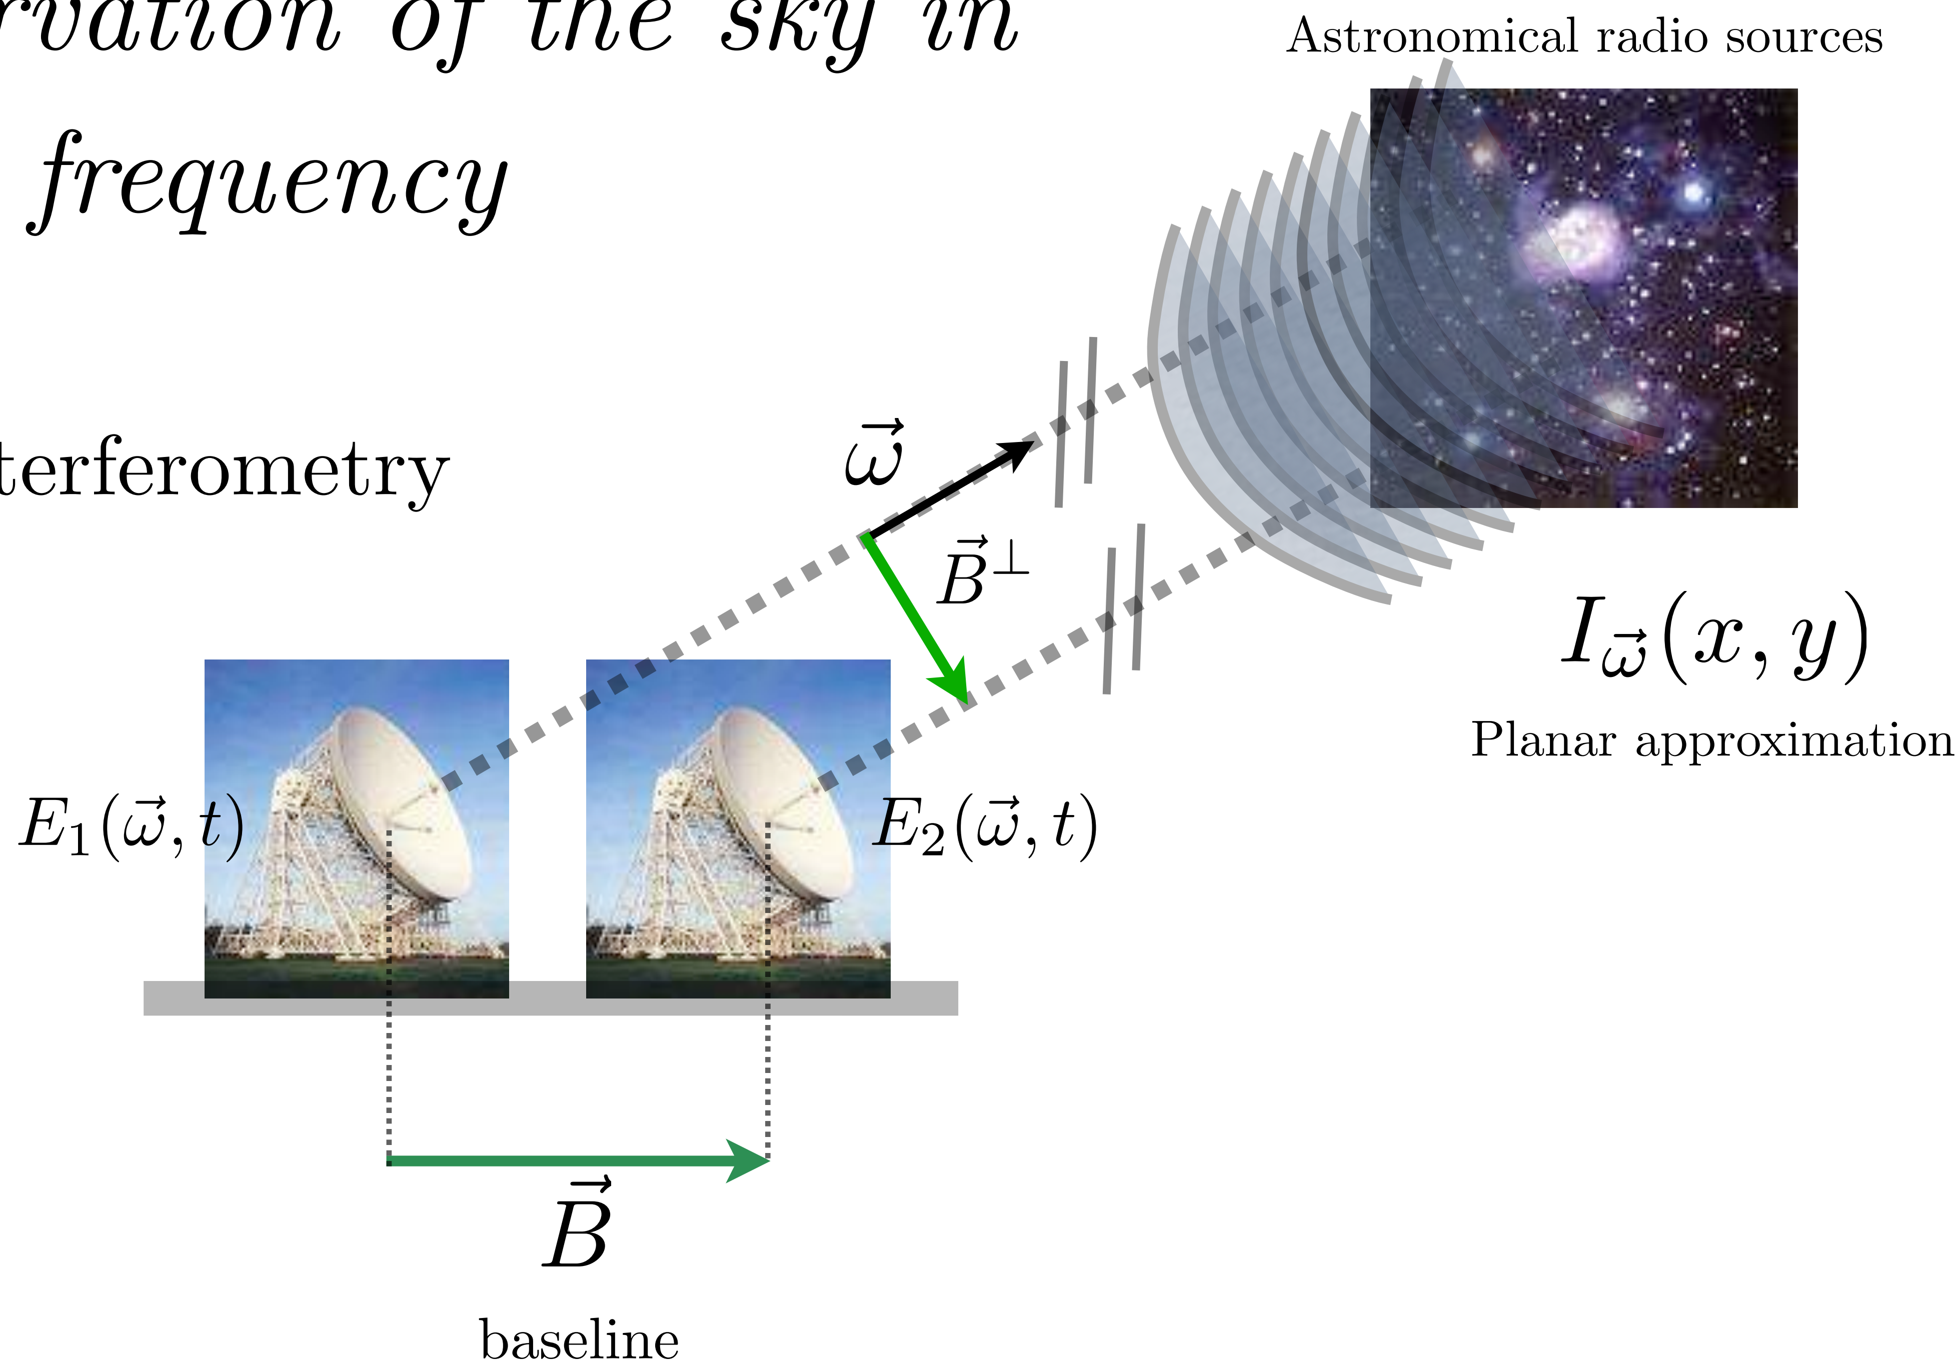

Aperture Synthesis in Radio-Astronomy

Observation of the sky in radio frequency

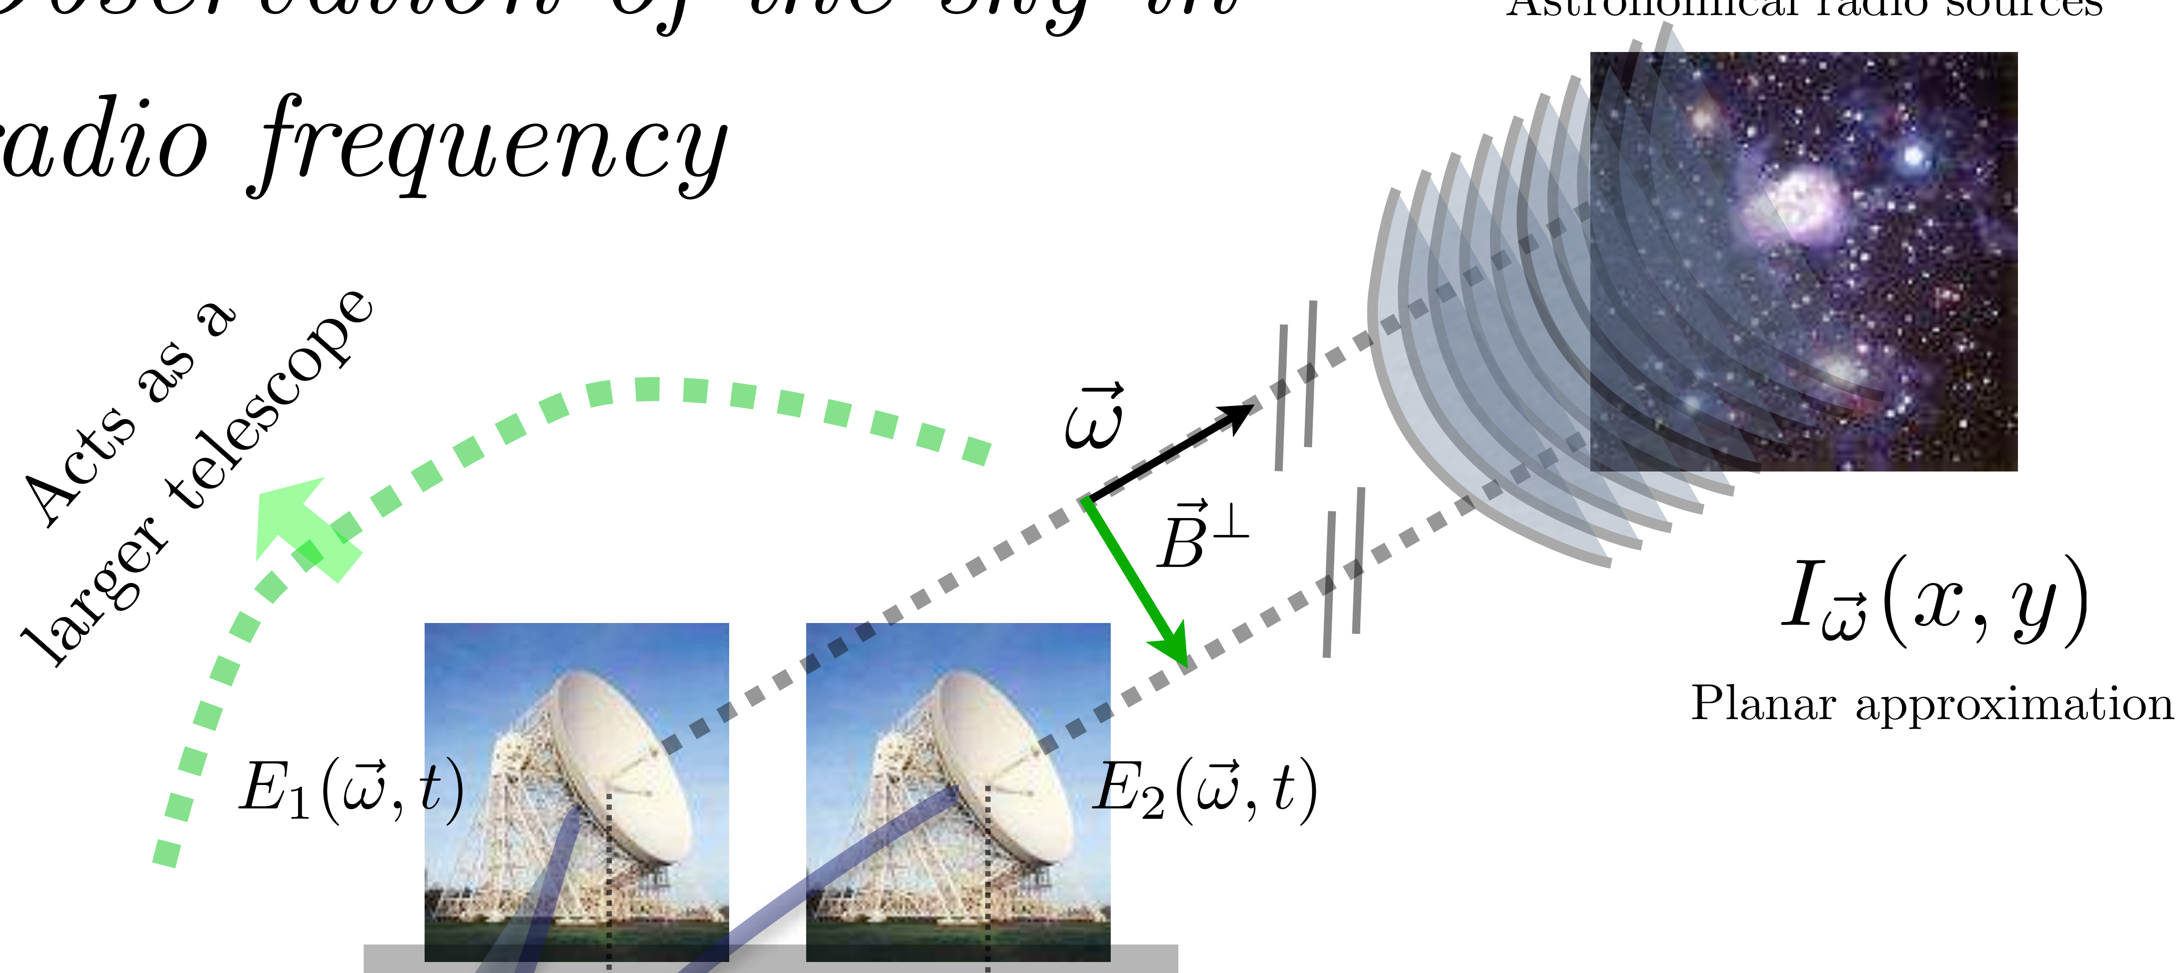

Time correlation :

$$\boxed{\hat{I}_{\vec{\omega}}(\vec{B}^{\perp})}$$

$\vec{B}$

baseline

Van Cittert Zernike Theorem :

$$\hat{I}_{\vec{\omega}}(\vec{B}^{\perp}) = \langle E_1 E_2^* \rangle_{\Delta t}$$

= Link between correlation and 2-D Fourier transform of I !

E pur si muove! "And yet it moves!"

- using N telescopes, $\binom{N}{2}$ possible Fourier observations
- and baselines undergo Earth rotation !

E pur si muove! "And yet it moves!"

- using N telescopes, $\binom{N}{2}$ possible Fourier observations
- and baselines undergo Earth rotation !
- Example :

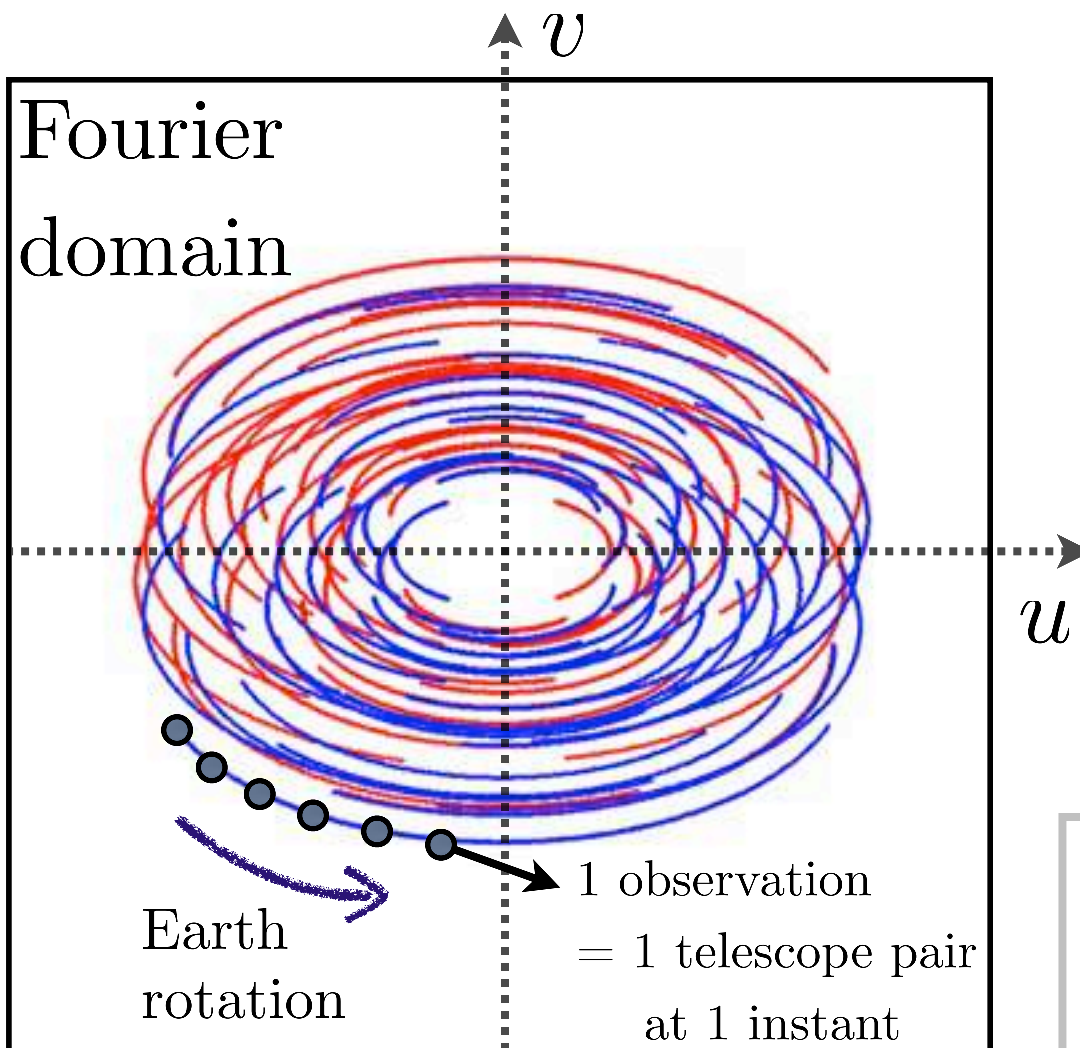

Each telescope pair = 1 elliptic path

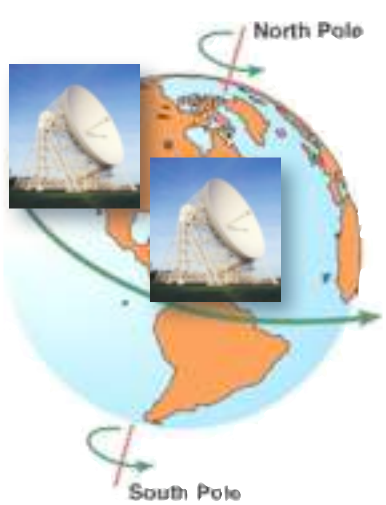

Arcminute Microkelvin Imager
AMI

18 - 120m

4 - 20m

Other configurations :
- Very Large Array VLA
- Square Kilometer Array SKA

E pur si muove! "And yet it moves!"

- using N telescopes, $\binom{N}{2}$ possible Fourier observations
- and baselines undergo Earth rotation !
- Example :

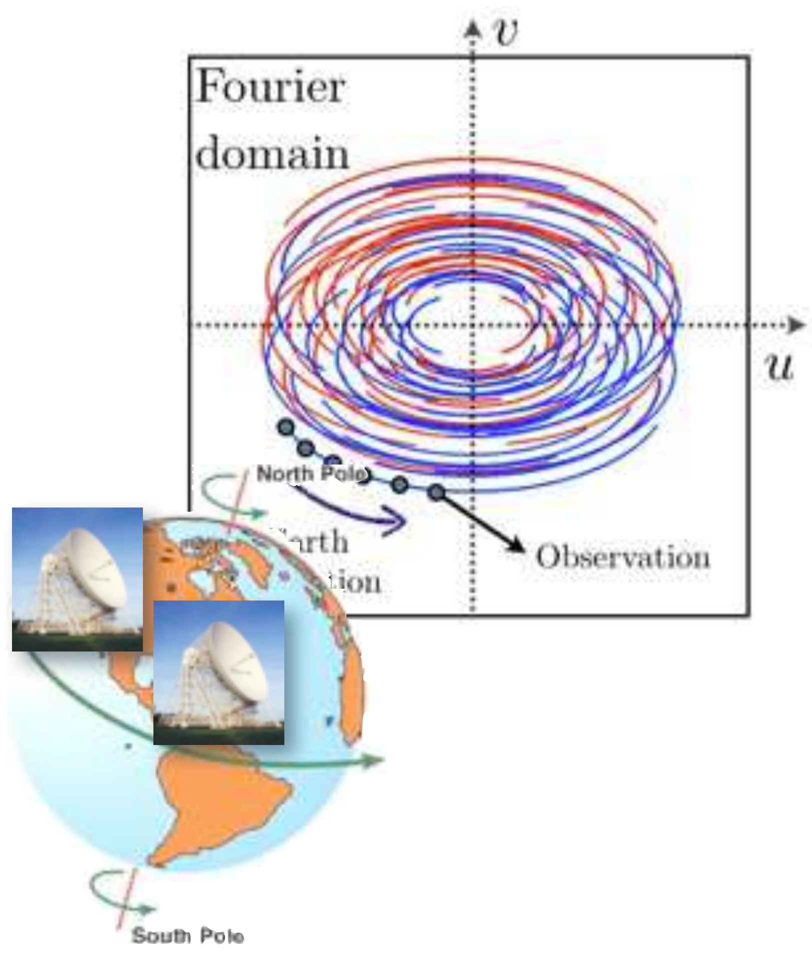

Partial Fourier Sampling Problem!

$$y = SFx + n$$

Selection of a few frequencies

2-D Fourier Transform

E pur si muove! "And yet it moves!"

- using N telescopes, $\binom{N}{2}$ possible Fourier observations
- and baselines undergo Earth rotation !
- Example :

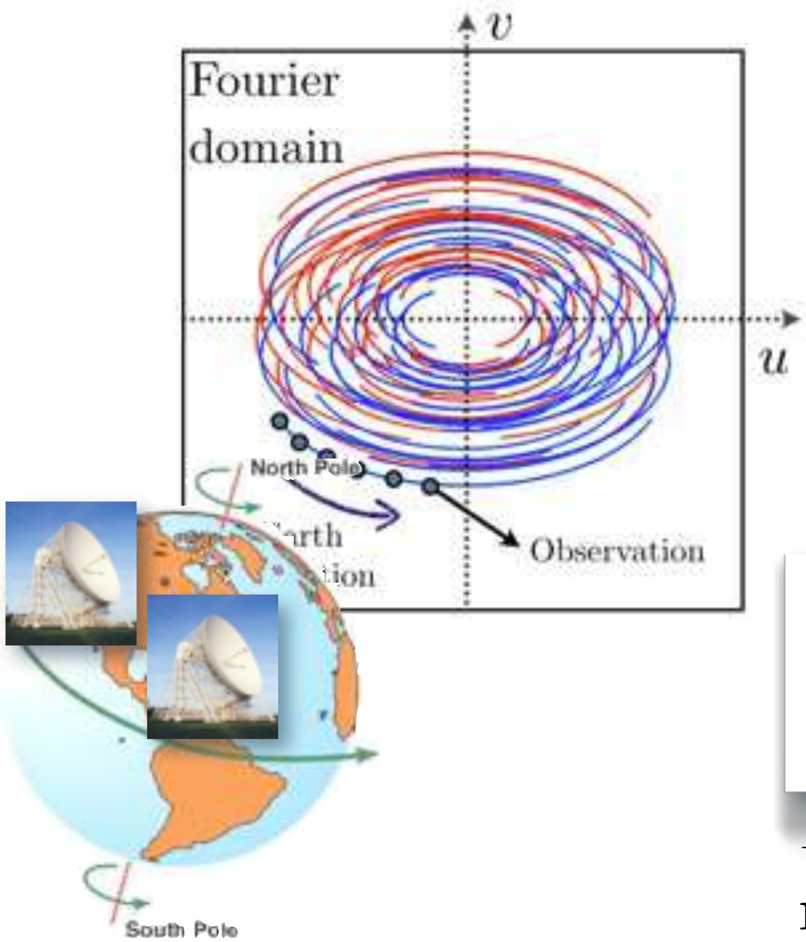

Partial Fourier Sampling Problem!

$$y = SFx + n$$

Selection of a few frequencies

2-D Fourier Transform

Image prior: SPARSITY!

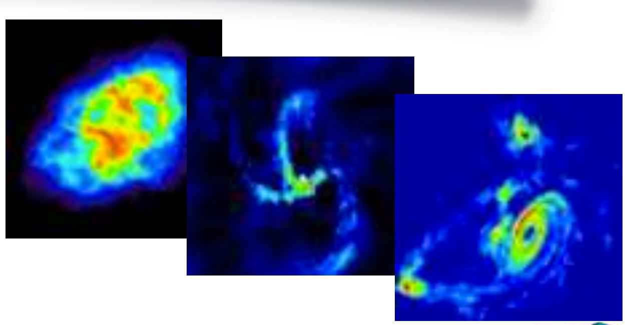

+ Additional improvements:
multiple sparsity basis, reweighted L1, ...

Reconstruction results

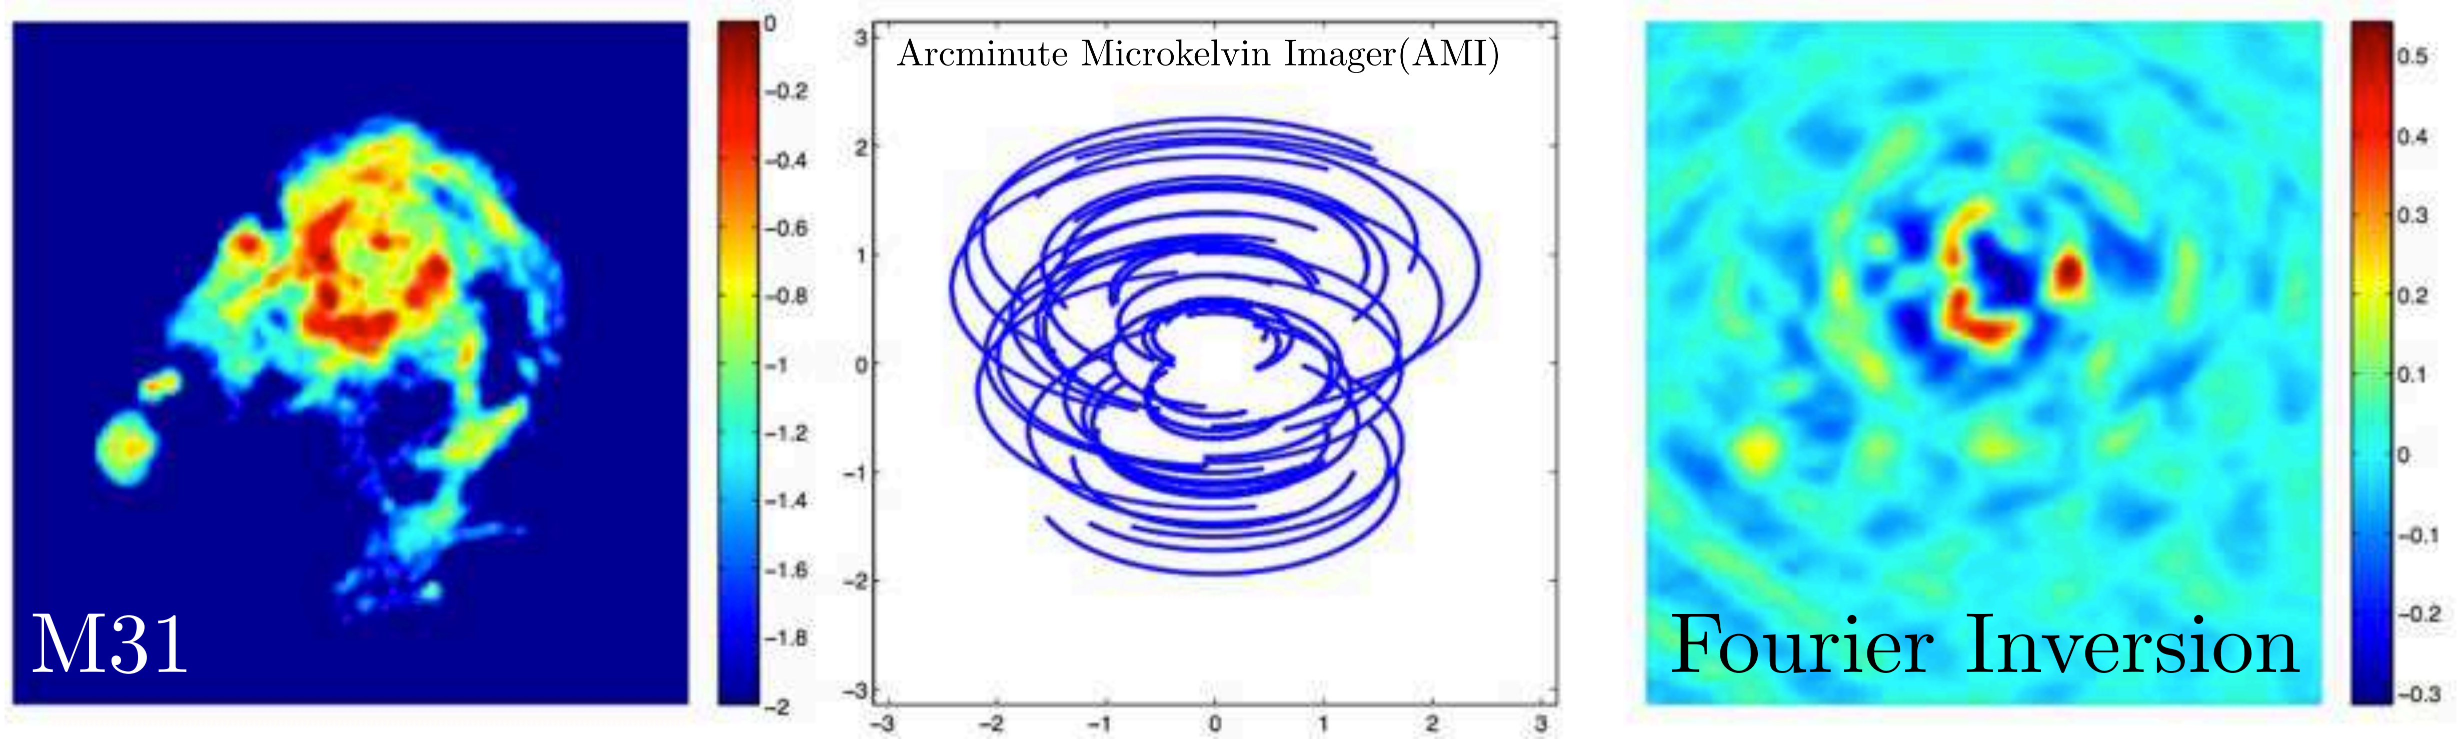

R. Carrillo, J. McEwen, Y. Wiaux, "PURIFY: a new approach to radio-interferometric imaging", Accepted MNRAS, 2014

Reconstruction results

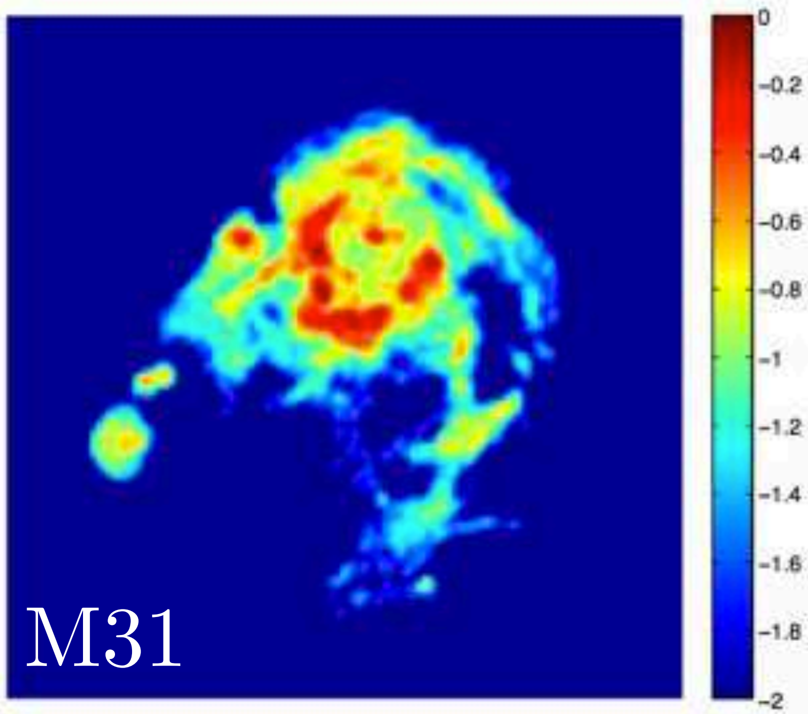

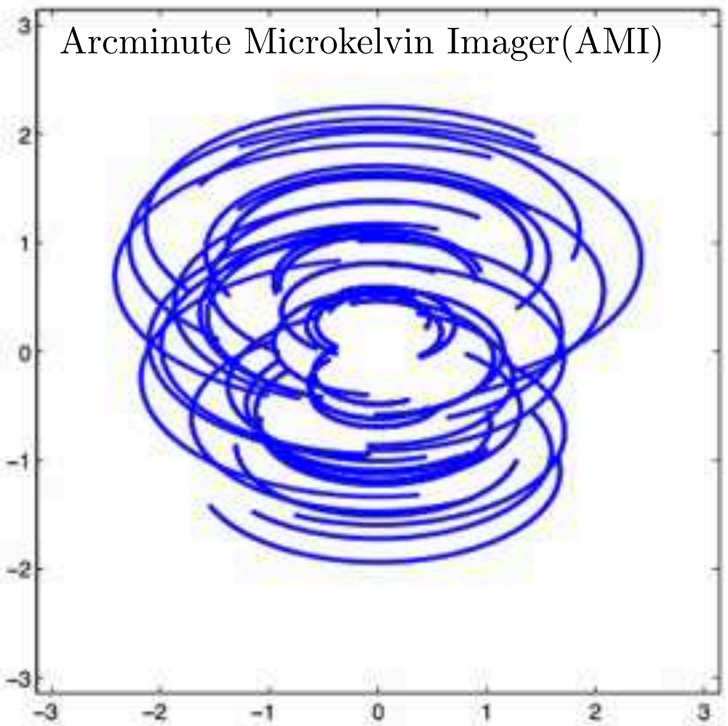

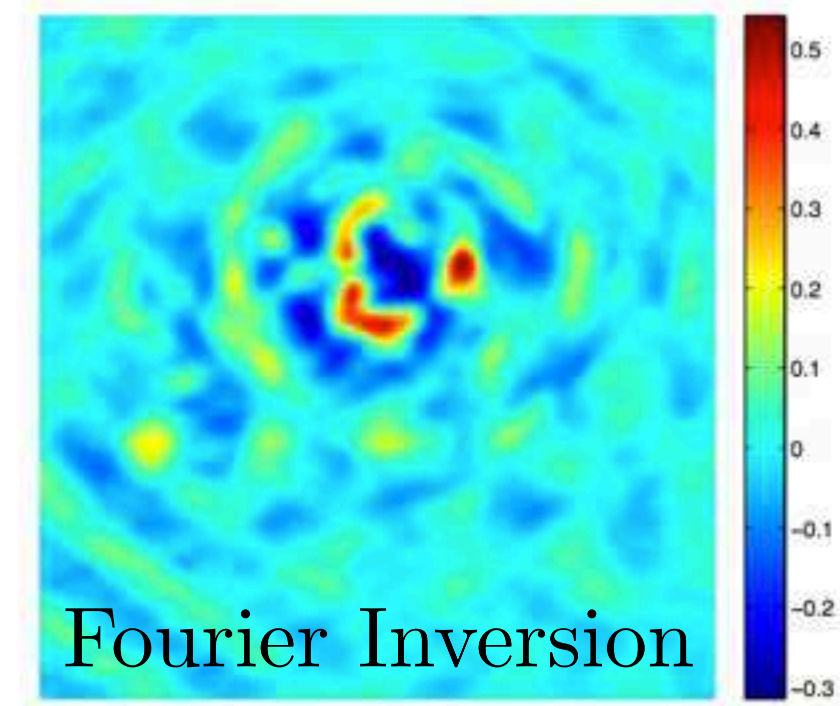

Sparsity
Averaging
Reweighted
Analysis
(SARA)

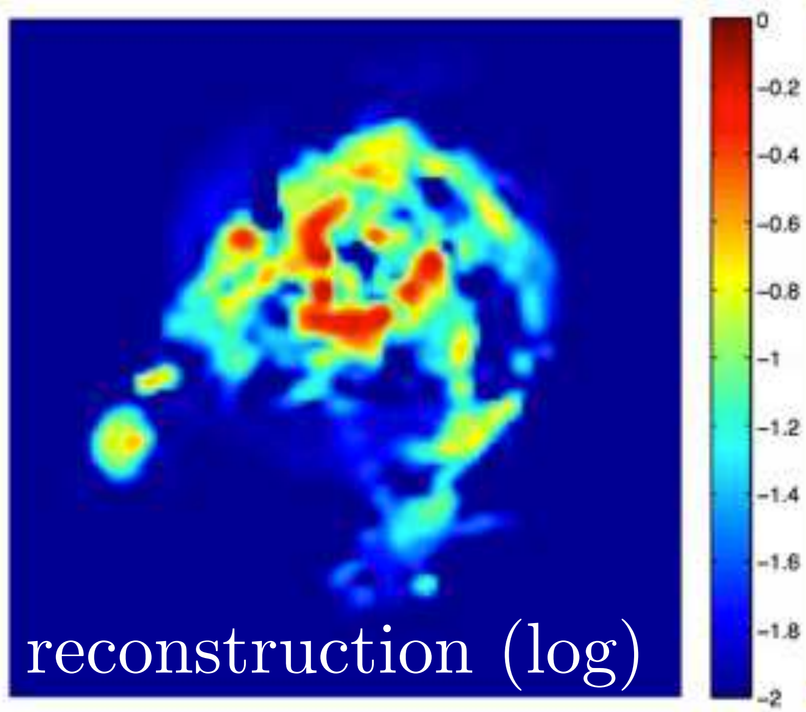

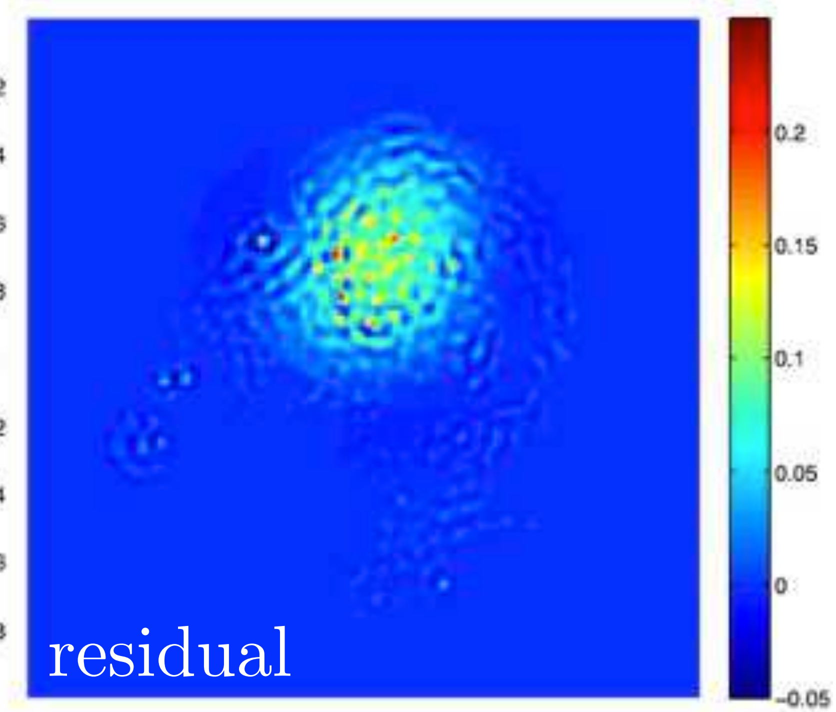

R. Carrillo, J. McEwen, Y. Wiaux, "PURIFY: a new approach to radio-interferometric imaging", Accepted MNRAS, 2014

Second Part:

- Compressive imaging appetizer:
  The Rice single pixel camera
- Other case studies:
  - Radio-interferometry and aperture synthesis
  - Hyperspectral CASSI imaging
  - Highspeed Coded Strobing Imaging

Hyperspectral imaging

- Fusion of spectrometry and imaging

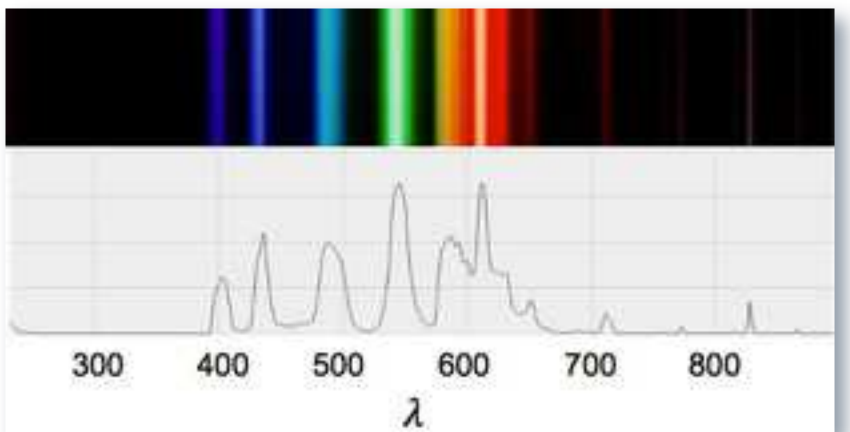

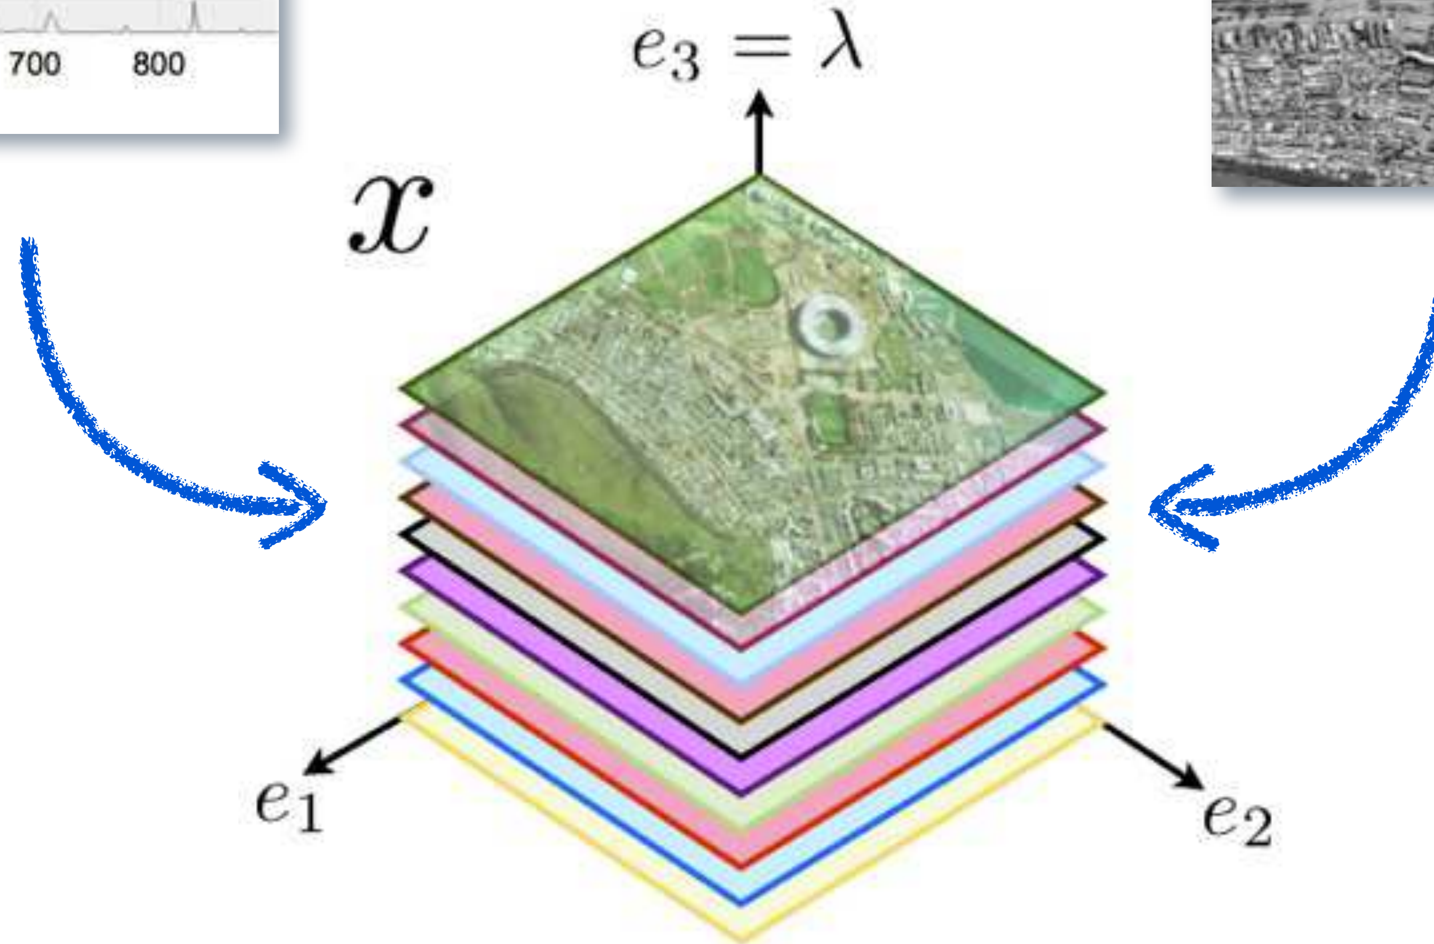

Hyperspectral imaging

- Fusion of spectrometry and imaging
- Applications:
  - material classification/segmentation

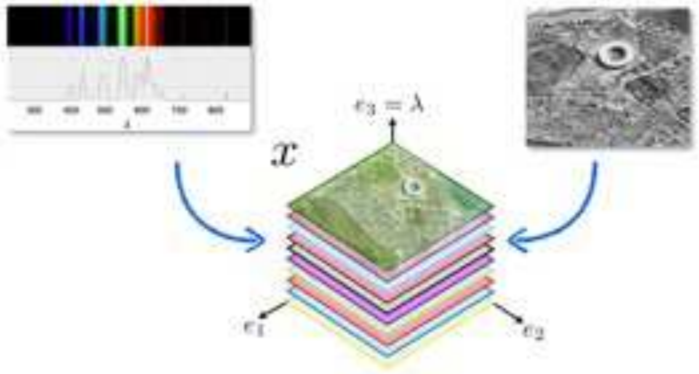

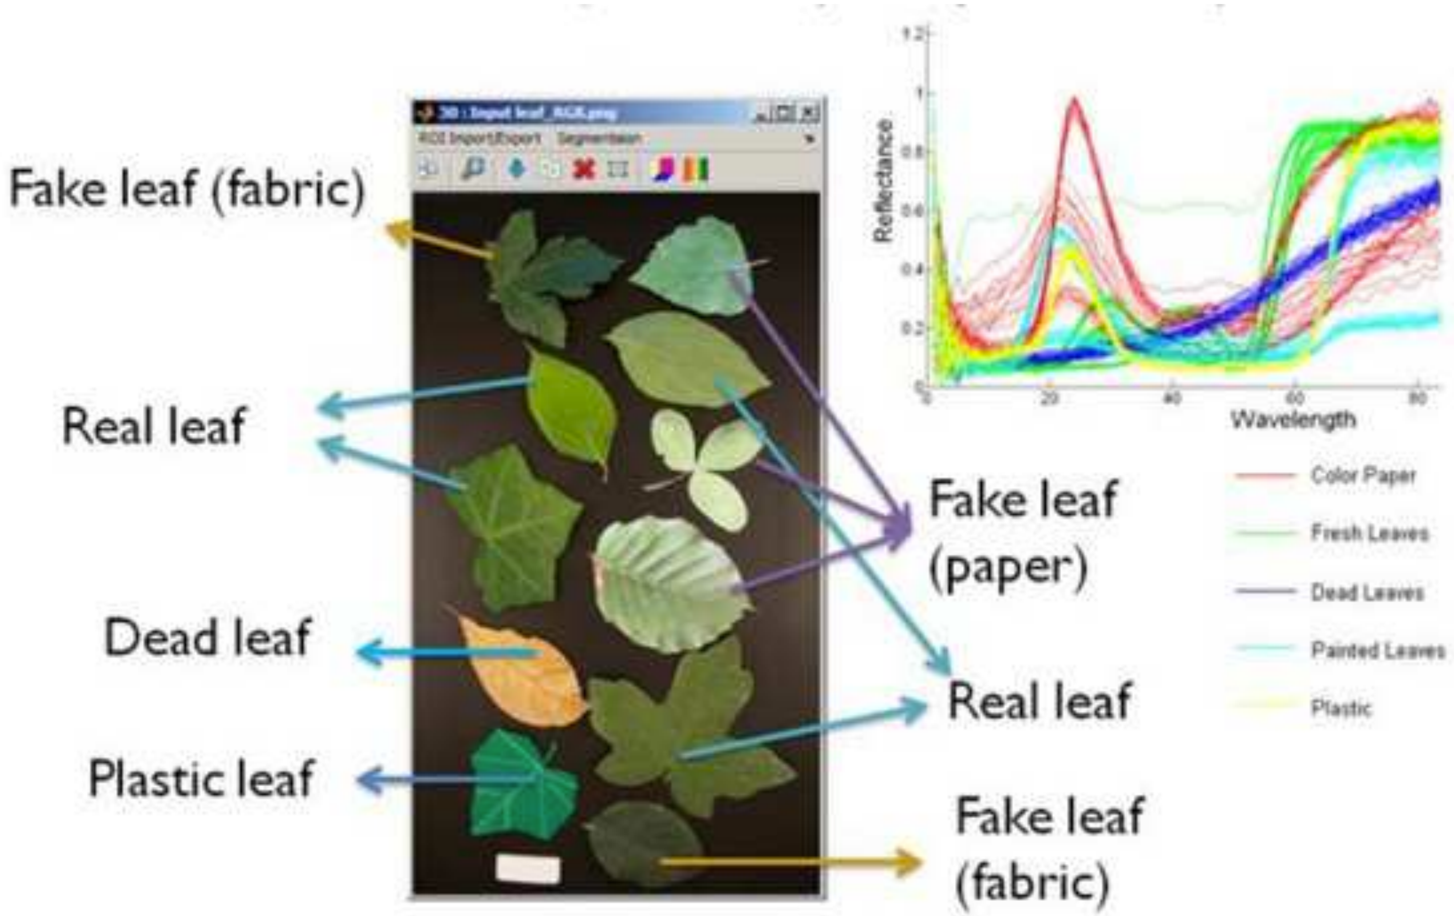

Hyperspectral imaging

- Fusion of spectrometry and imaging
- Applications:
  - material classification/segmentation
  - microscopy/spectroscopy
  - counterfeit detection
  - environmental monitoring
  - skin decease detection
  - ...

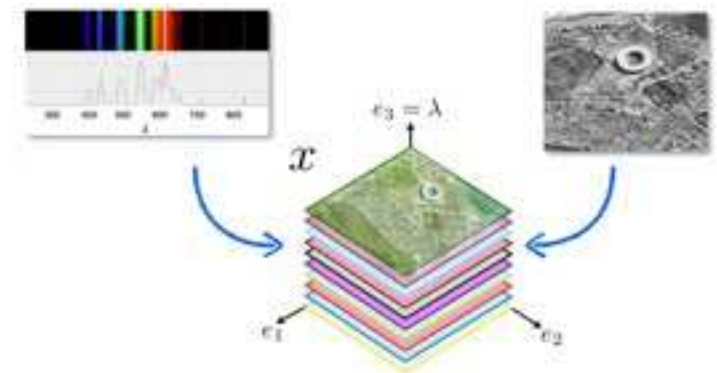

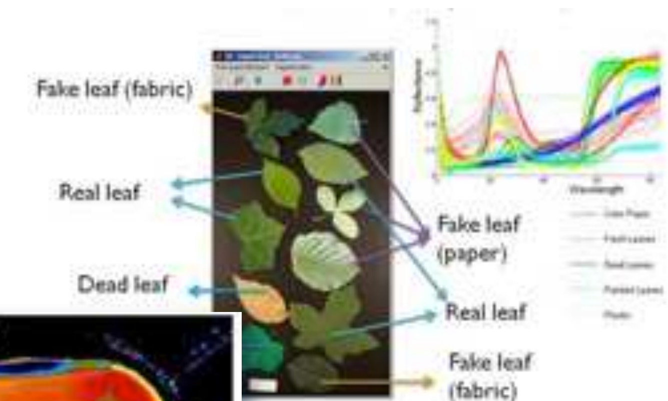

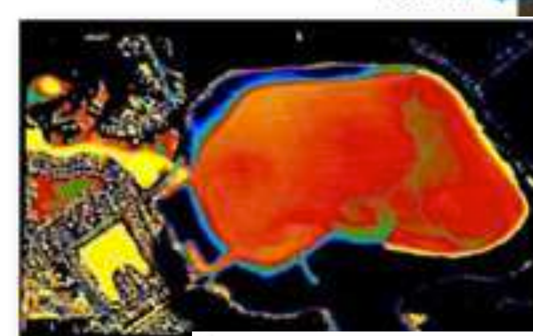

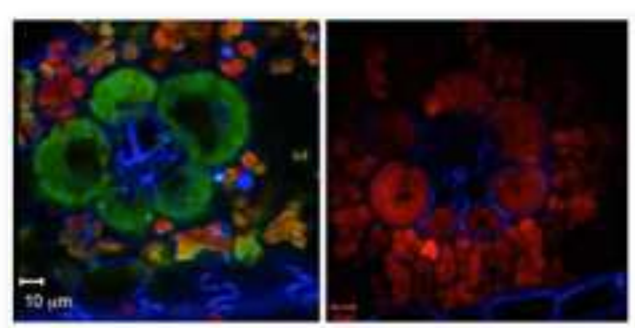

How is it usually done?

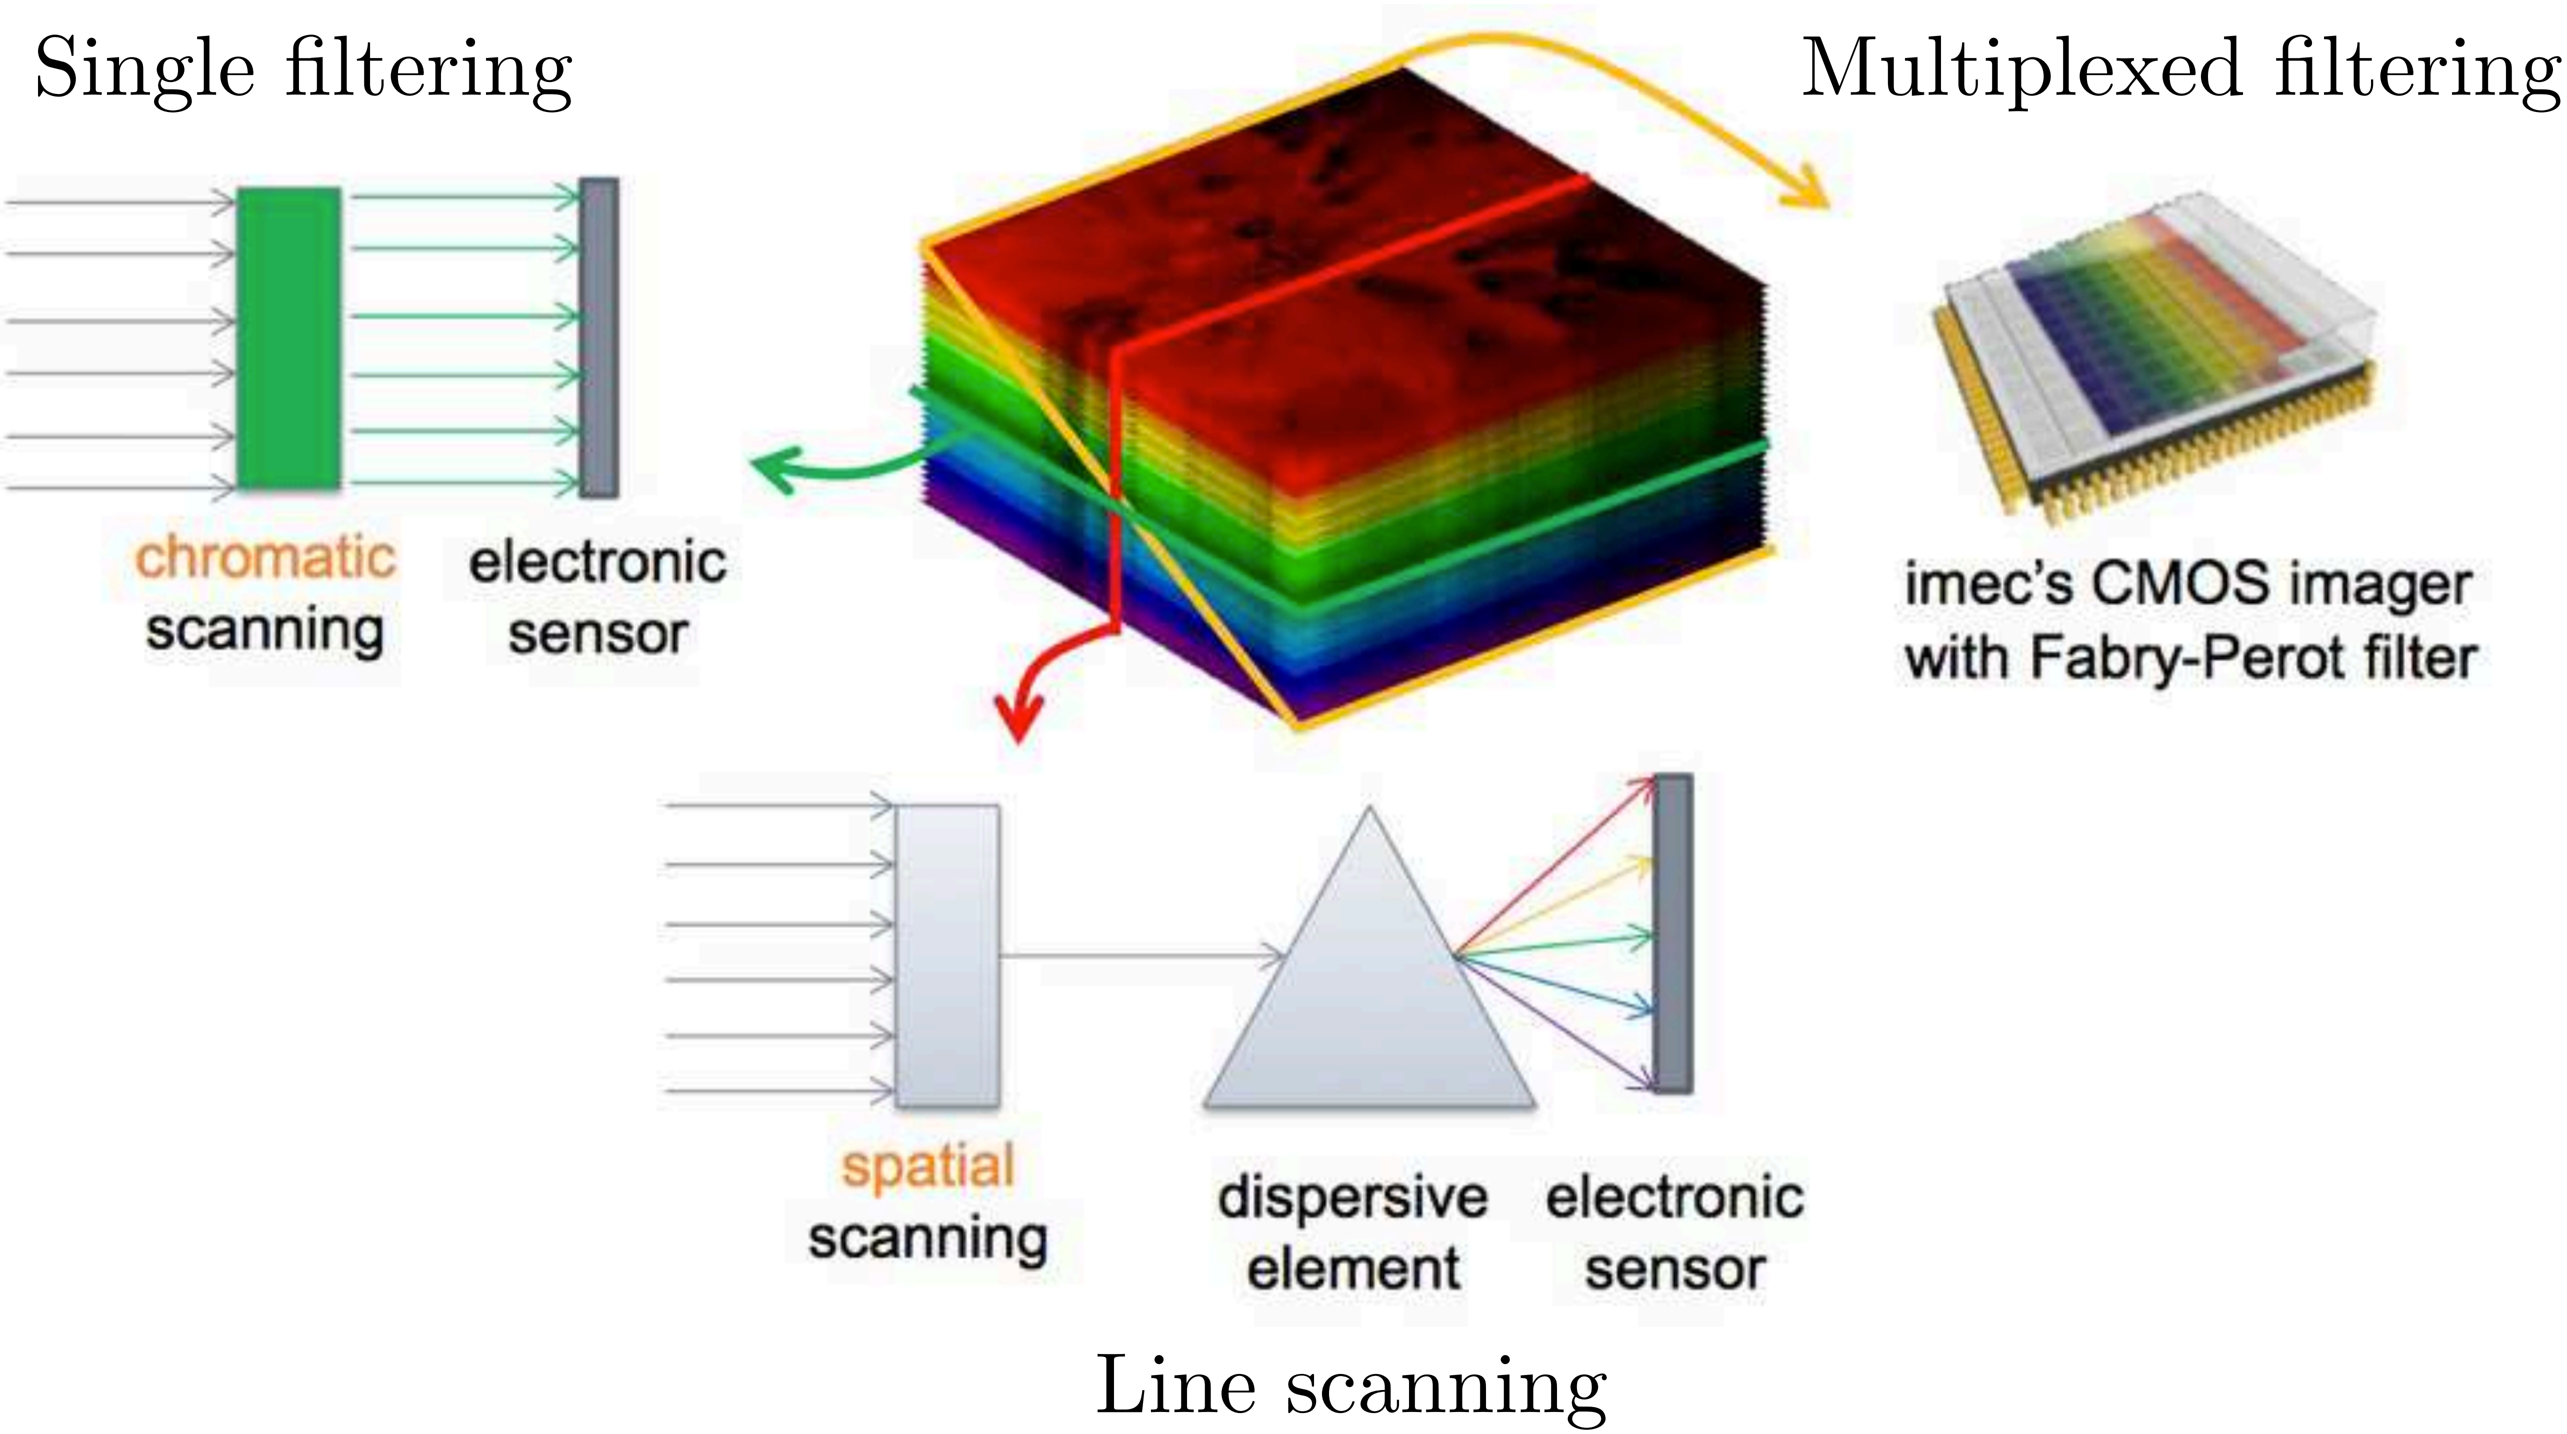

How is it usually done?

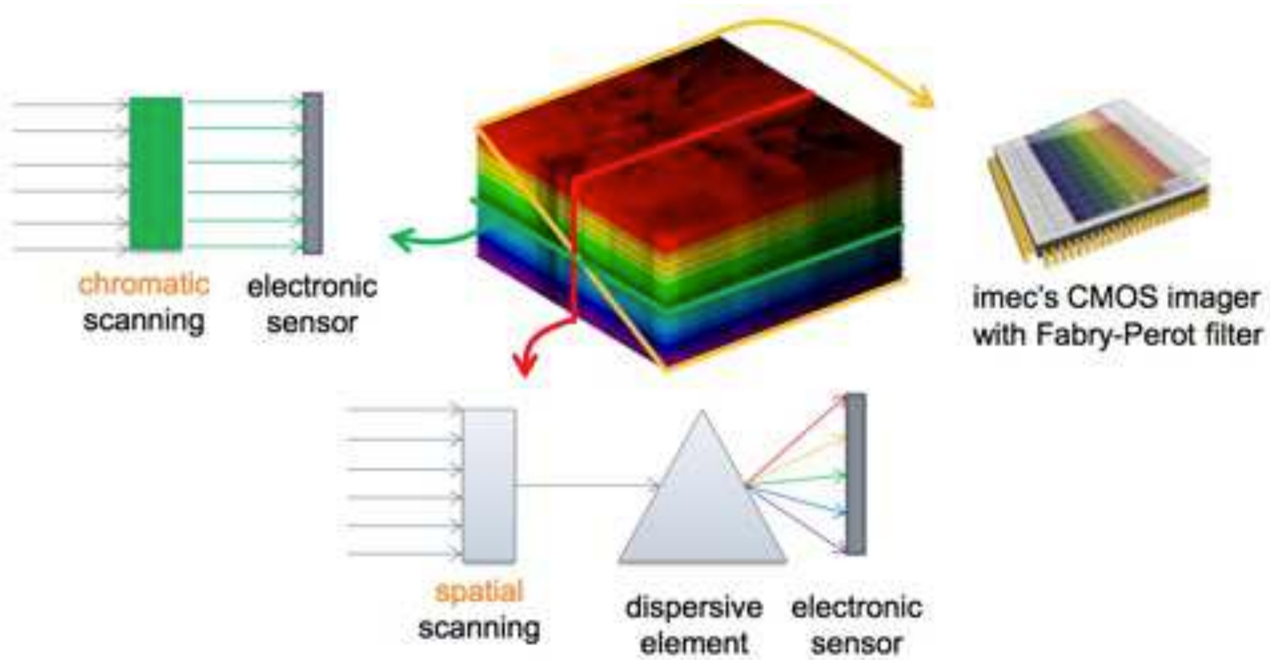

Issues:

- acquisition time is slow
- low spatial/spectral/temporal resolution

  (depending on selected sensing)
- Huge amount of data at sensing
- But "low complexity" (sparse/low-rank) signals

Compressive HS imaging

- high-dimensional data = natural field for CS!
- Coded Aperture Snapshot Spectral Imaging (CASSI)
  - Mixing dispersive element + coded aperture

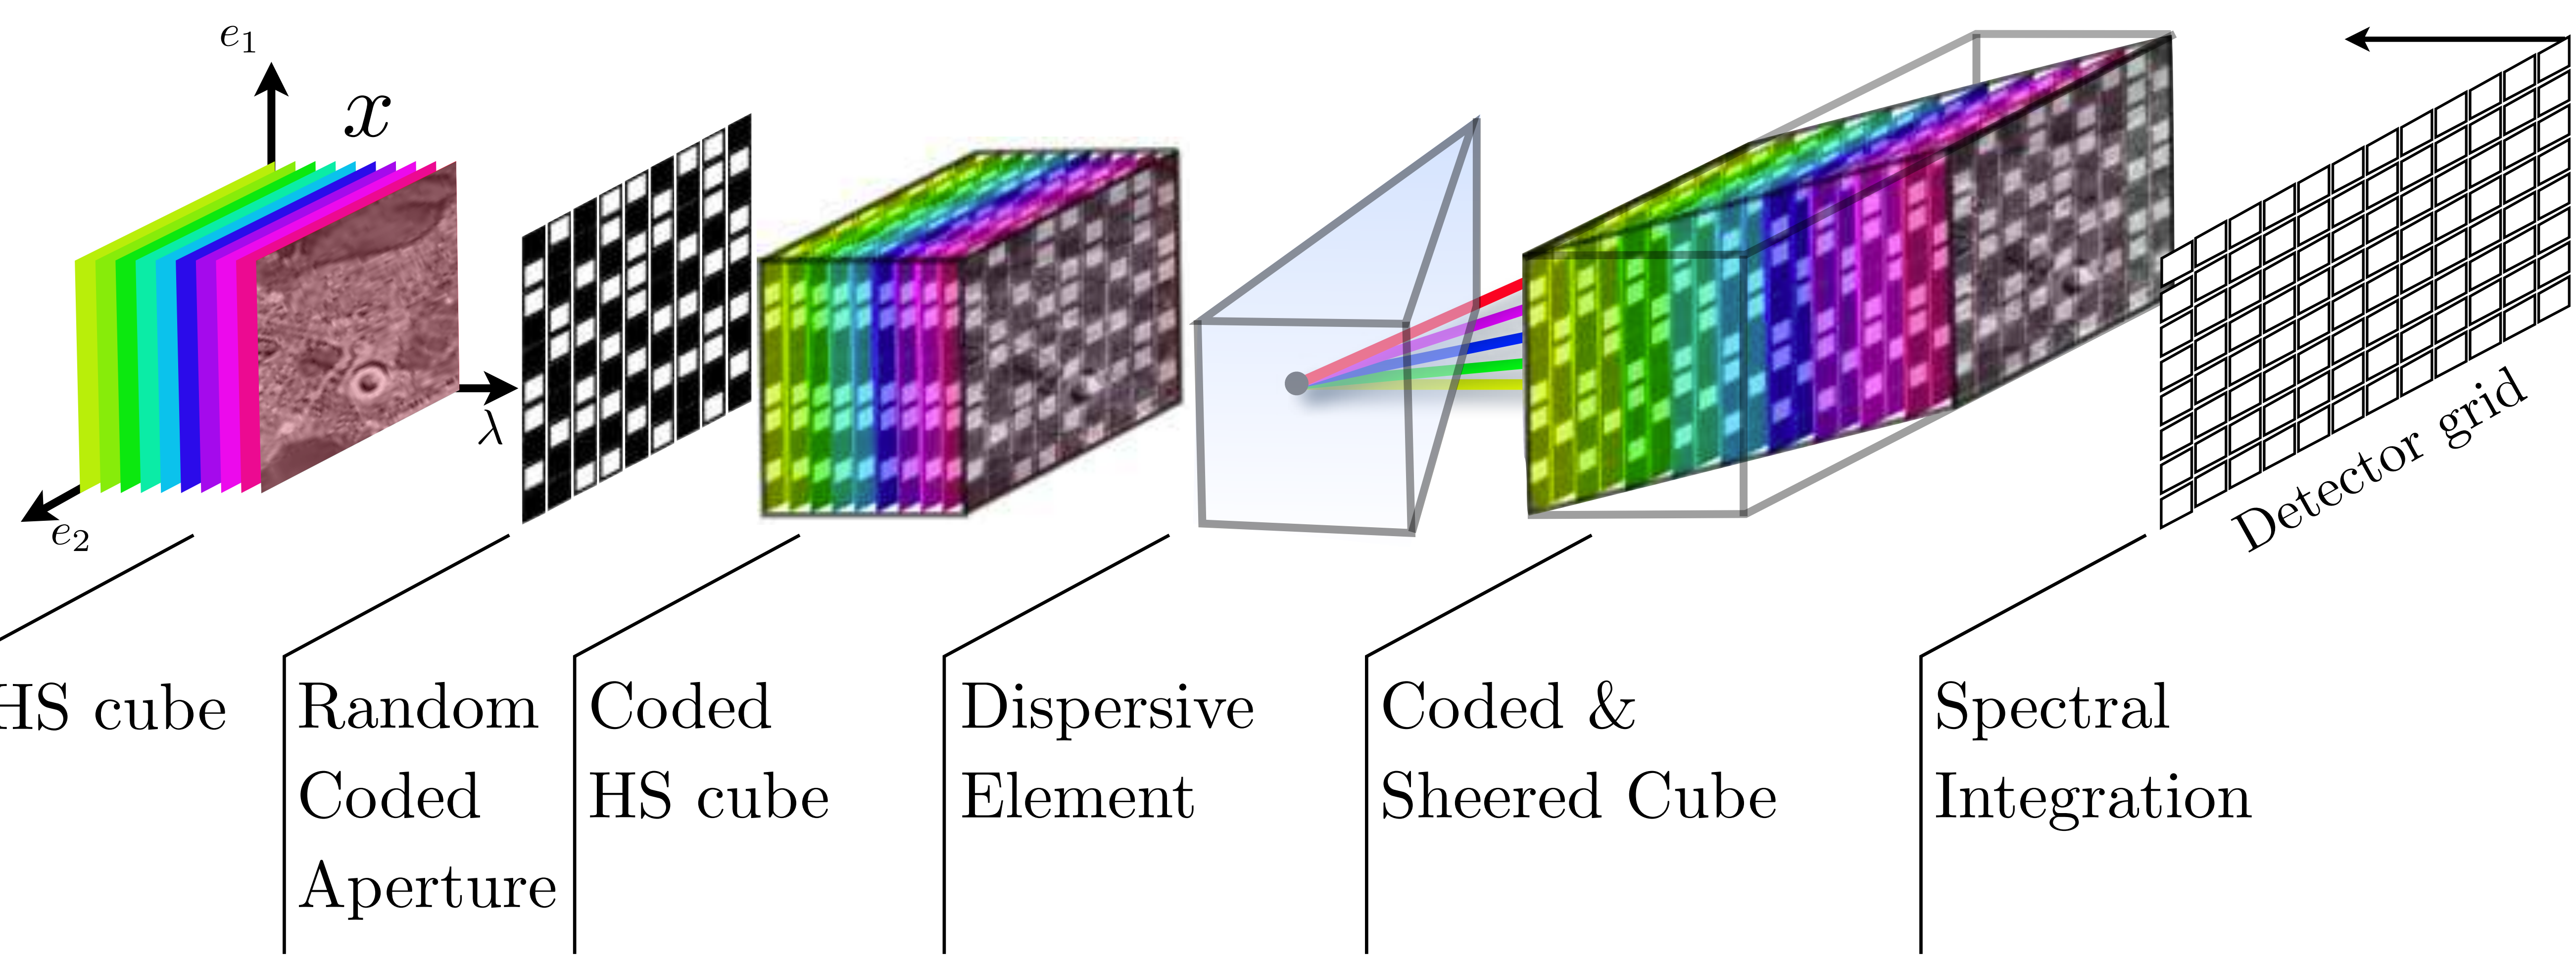

Compressive HS imaging

- high-dimensional data = natural field for CS!
- Coded Aperture Snapshot Spectral Imaging (CASSI)
  - Mixing dispersive element + coded aperture

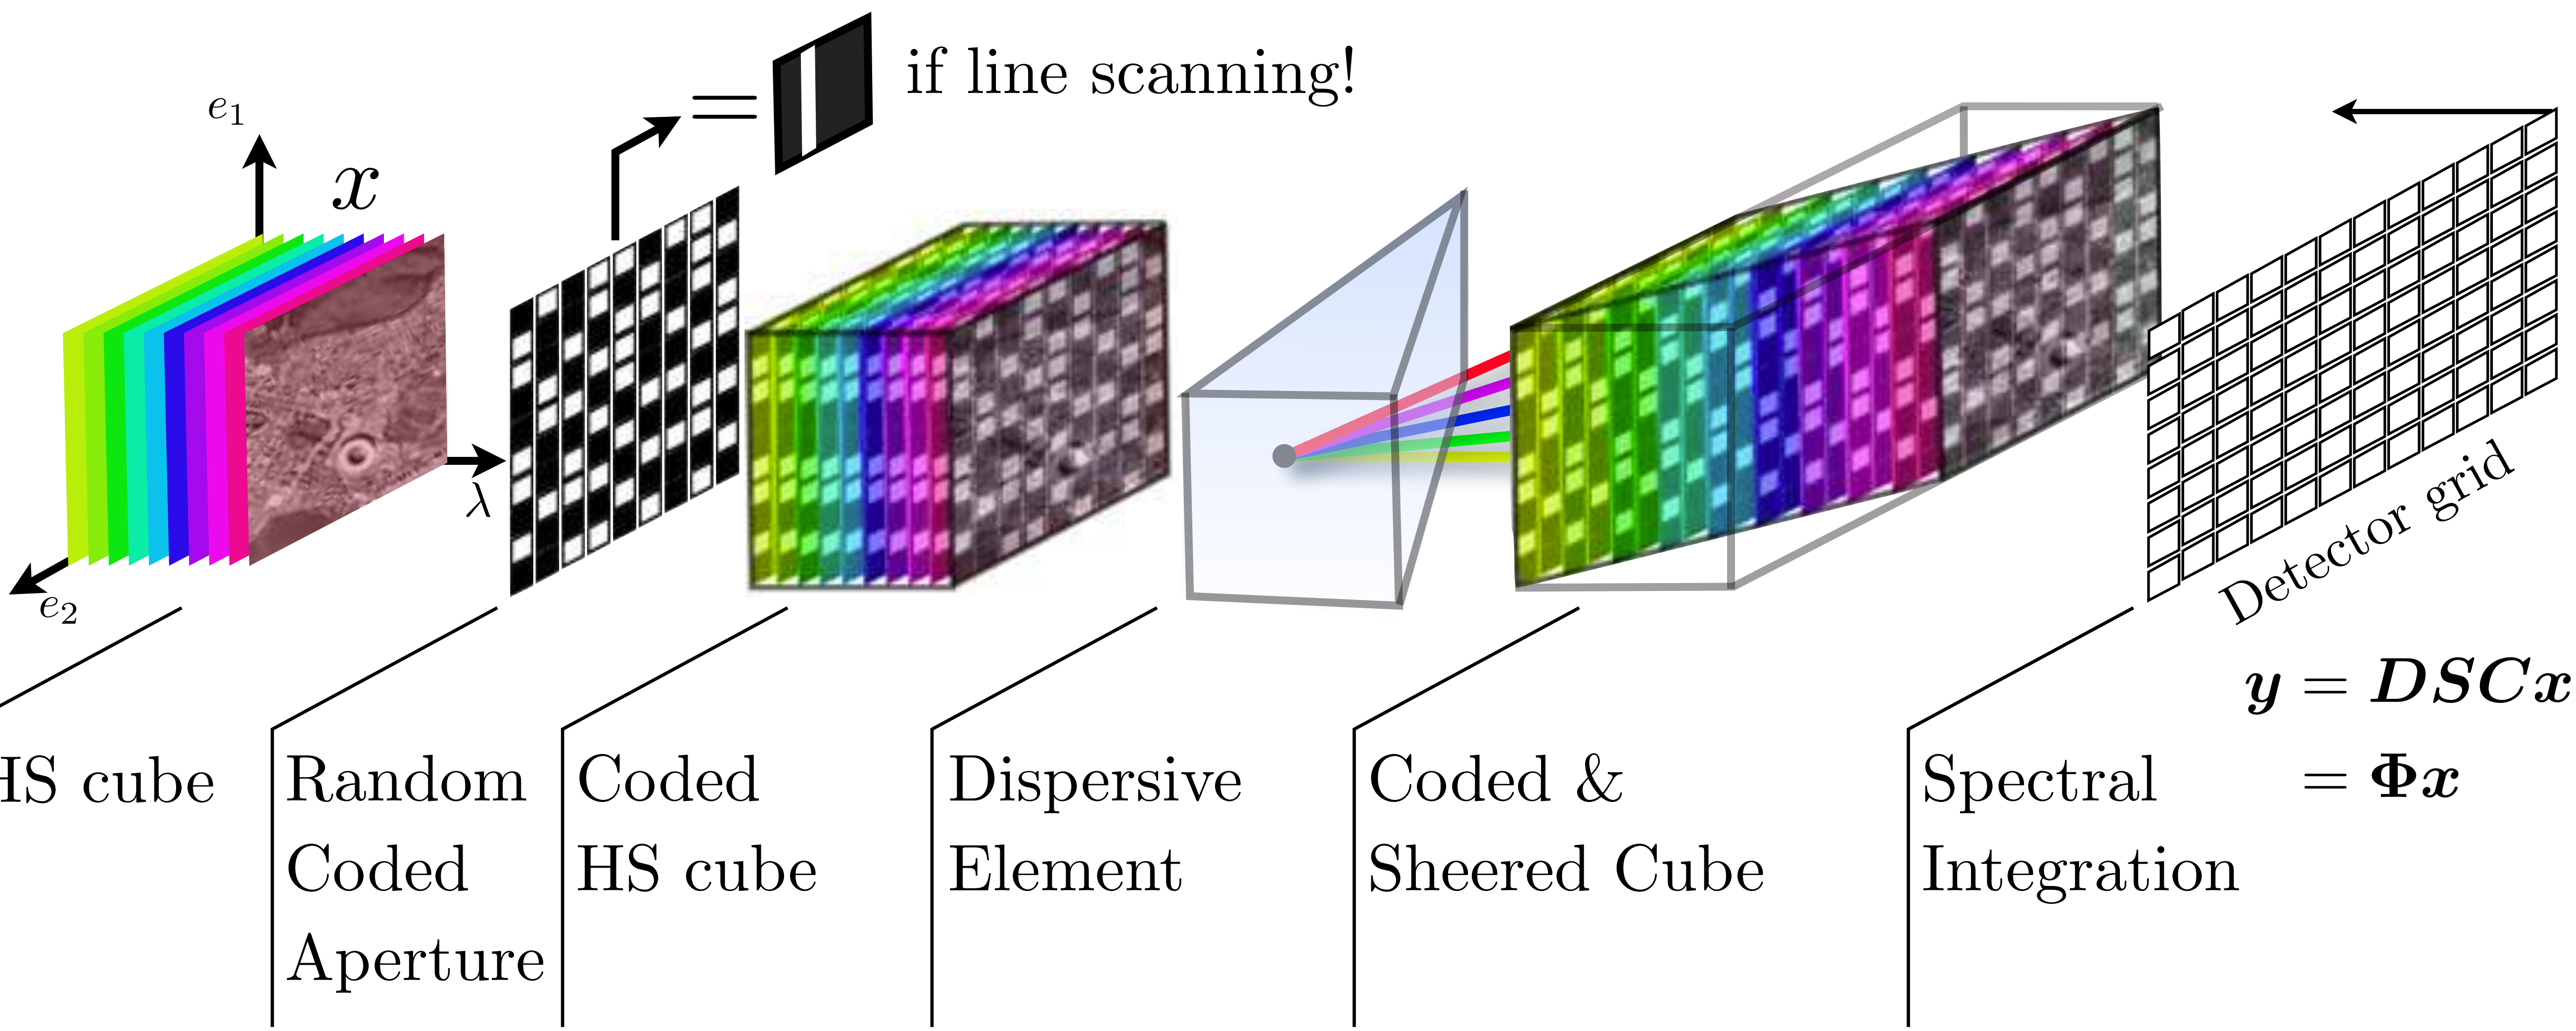

Compressive HS imaging

- high-dimensional data = natural field for CS!
- Coded Aperture Snapshot Spectral Imaging (CASSI)
  - Mixing dispersive element + coded aperture

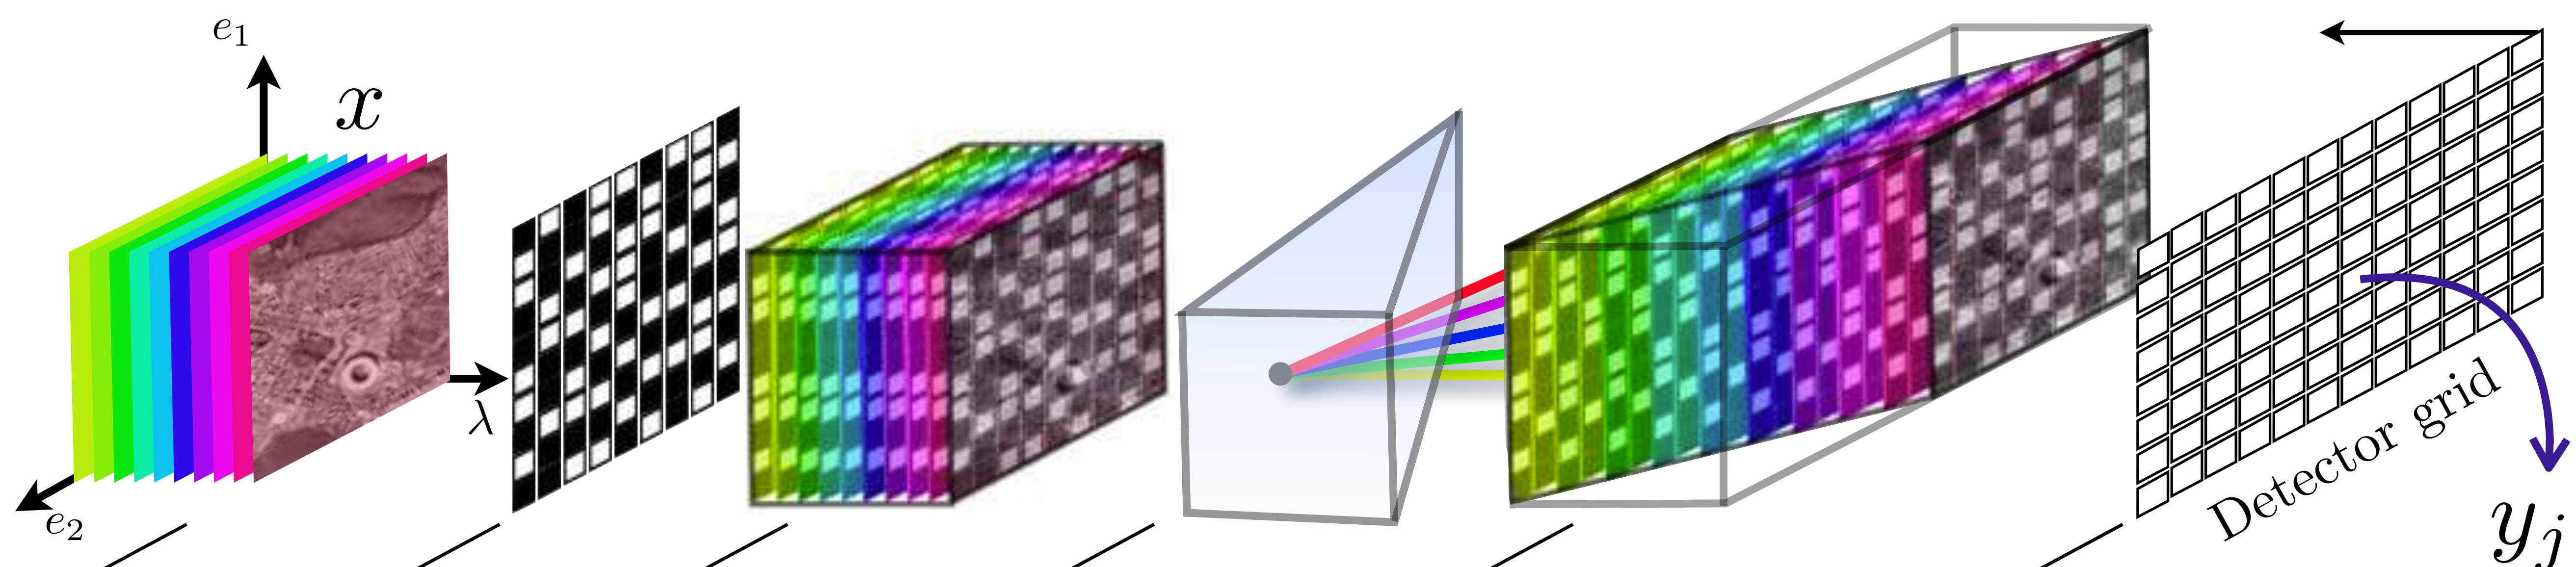

$\boldsymbol{x} \rightarrow \boldsymbol{C} = \mathrm{diag}(\boldsymbol{c}) \rightarrow \boldsymbol{S} \rightarrow \boldsymbol{D} \rightarrow \boldsymbol{y}$

Sensing model: $\boldsymbol{y} = \boldsymbol{DSCx} = \boldsymbol{\Phi x}$

+ with multiple $\boldsymbol{c}$

Compressive HS imaging

- high-dimensional data = natural field for CS!
- Coded Aperture Snapshot Spectral Imaging (CASSI)
  - Mixing dispersive element + coded aperture

Optically:

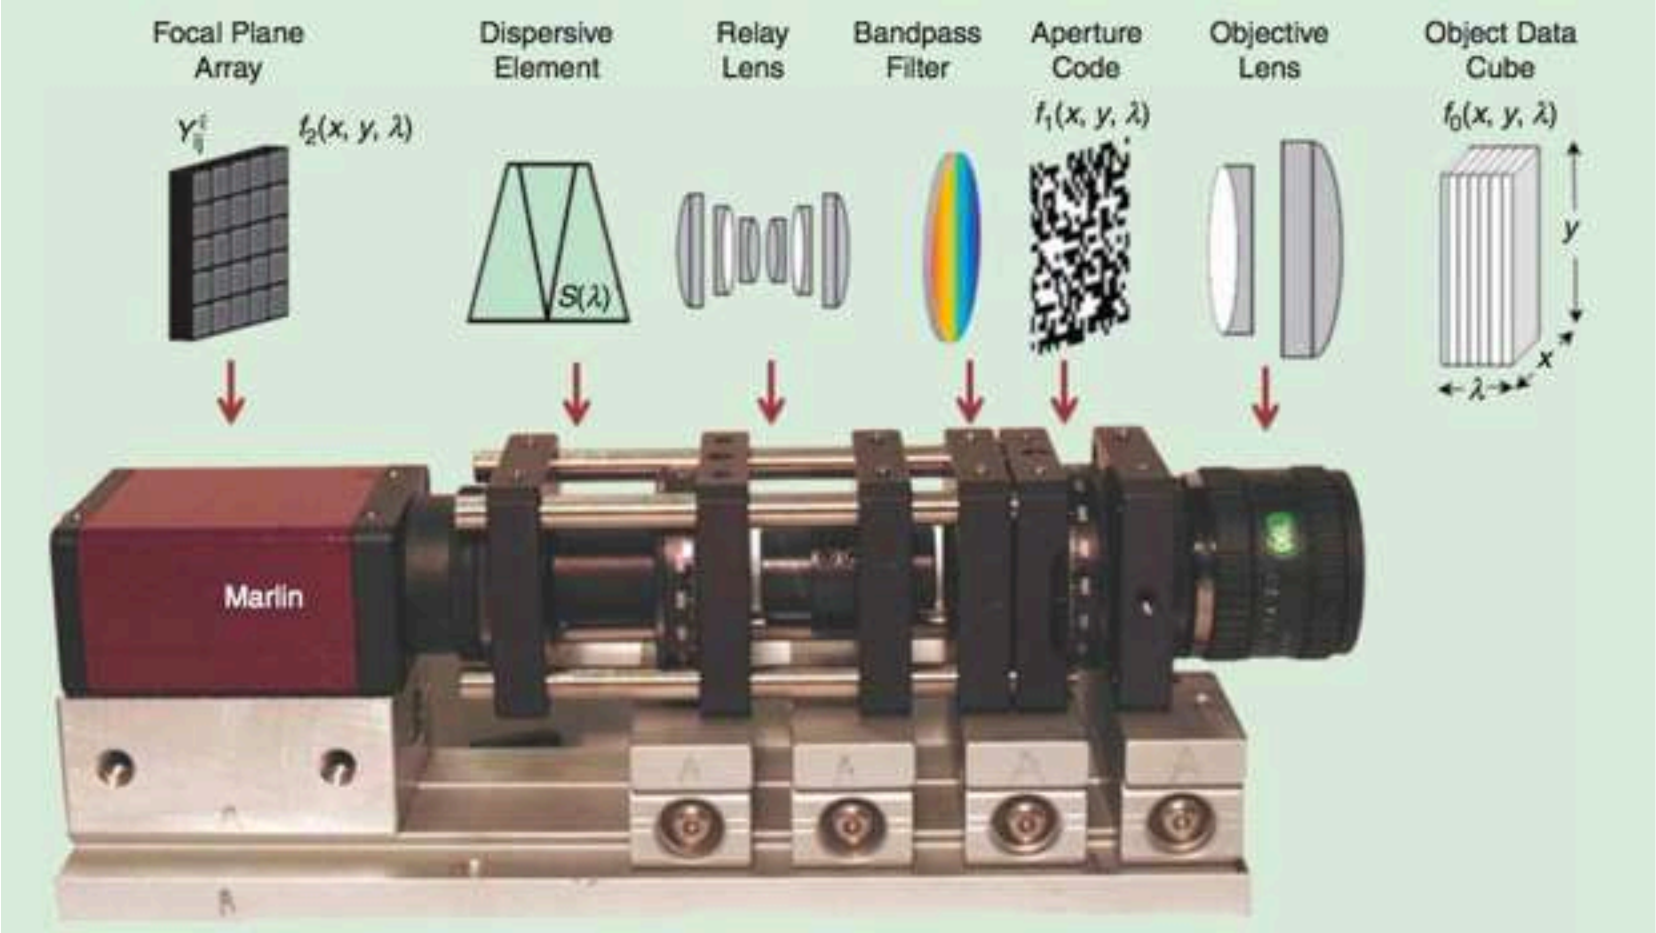

Gonzalo R. Arce, David J. Brady, Lawrence Carin, Henry Arguello, and David S. Kittle, "Compressive Coded Aperture Spectral Imaging", IEEE Sig. Proc, vol. 1, 2014

Compressive HS imaging

- Reconstruction: solving

$$x^* = \Psi^T(\arg\min_{\alpha} \tau\|\alpha\|_1 + \tfrac{1}{2}\,\|y - DSC\Psi\alpha\|^2)$$

$$\Psi = \Psi_1 \otimes \Psi_2$$

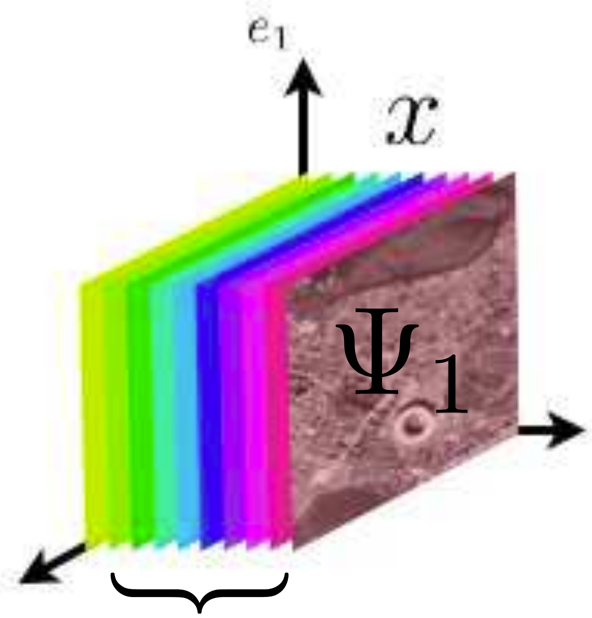

2-D wavelet
Symmlet-8

Ψ_2
DCT

Gonzalo R. Arce, David J. Brady, Lawrence Carin, Henry Arguello, and David S. Kittle, "Compressive Coded Aperture Spectral Imaging", IEEE Sig. Proc, vol. 1, 2014

Compressive HS imaging

- Reconstruction:

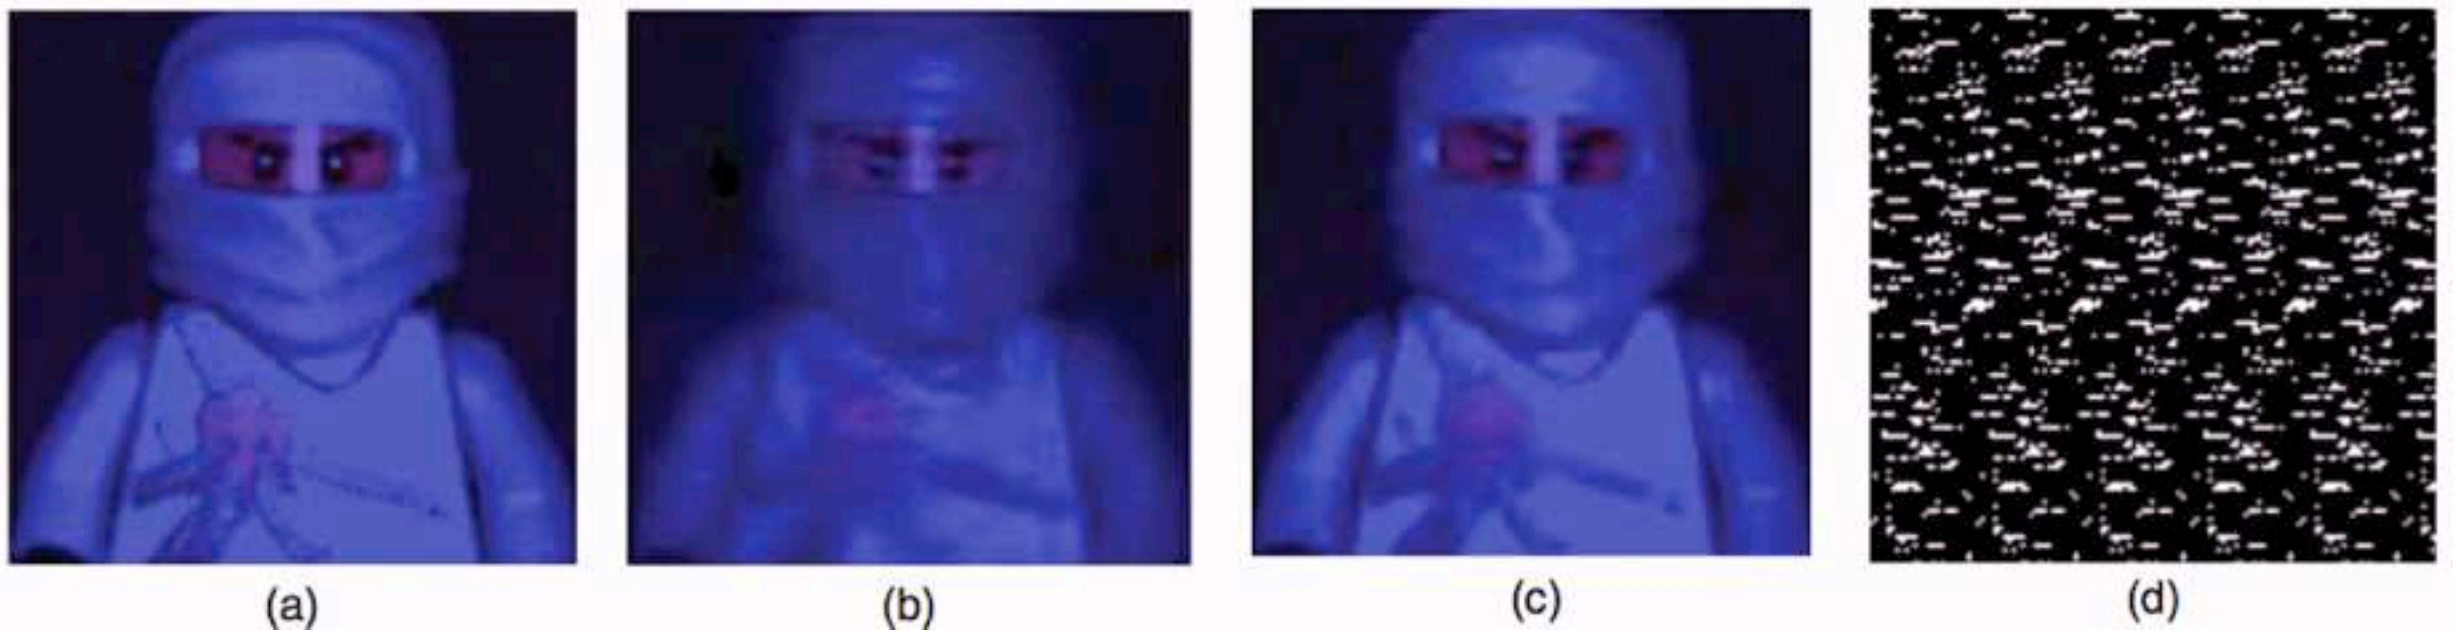

Original (several wav.)	With 12 shots and random $\boldsymbol{c}$	With 12 shots and optimized $\boldsymbol{c}$	example of optimized $\boldsymbol{c}$

Gonzalo R. Arce, David J. Brady, Lawrence Carin, Henry Arguello, and David S. Kittle, "Compressive Coded Aperture Spectral Imaging", IEEE Sig. Proc, vol. 1, 2014

Second Part:

- Compressive imaging appetizer:
  The Rice single pixel camera
- Other case studies:
  - Radio-interferometry and aperture synthesis
  - Hyperspectral CASSI imaging
  - Highspeed Coded Strobing Imaging

High Speed Imaging

- Imaging high speed object
  lead to blurry image
  if low shutter frequency

(source: wikipedia)

- But hardware limitation in # fps (e.g., O(20fps))
- <u>Solution</u>: "Highspeed Coded Strobing Imaging"
  - keep the detector fps rate unchanged
  - and add high rate *coding* of the shutter!

(Reddy, Veeraraghavan, Chellappa, ...)

Highspeed Coded Strobing Imaging

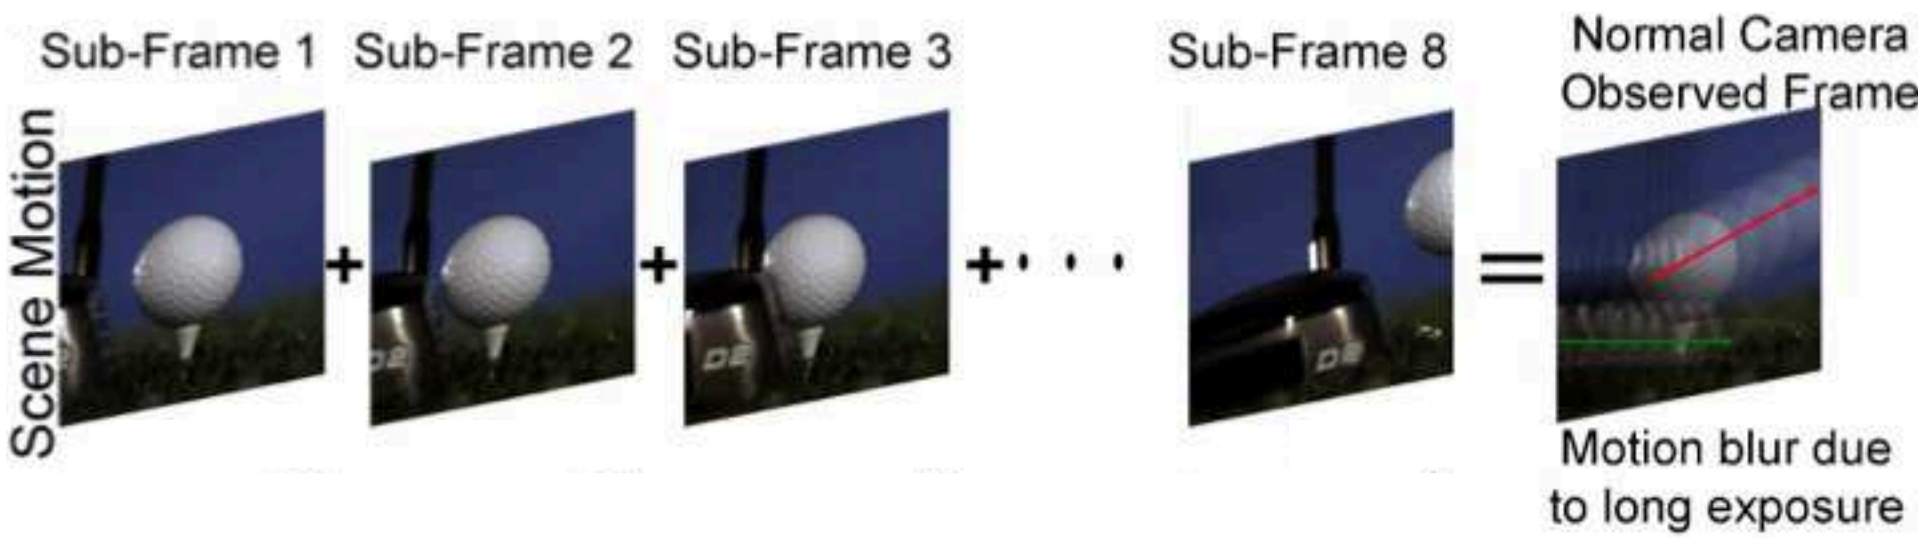

R. Dikpal, A. Veeraraghavan, and R. Chellappa. "P2C2: Programmable pixel compressive camera for high speed imaging" Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. IEEE, 2011.

Highspeed Coded Strobing Imaging

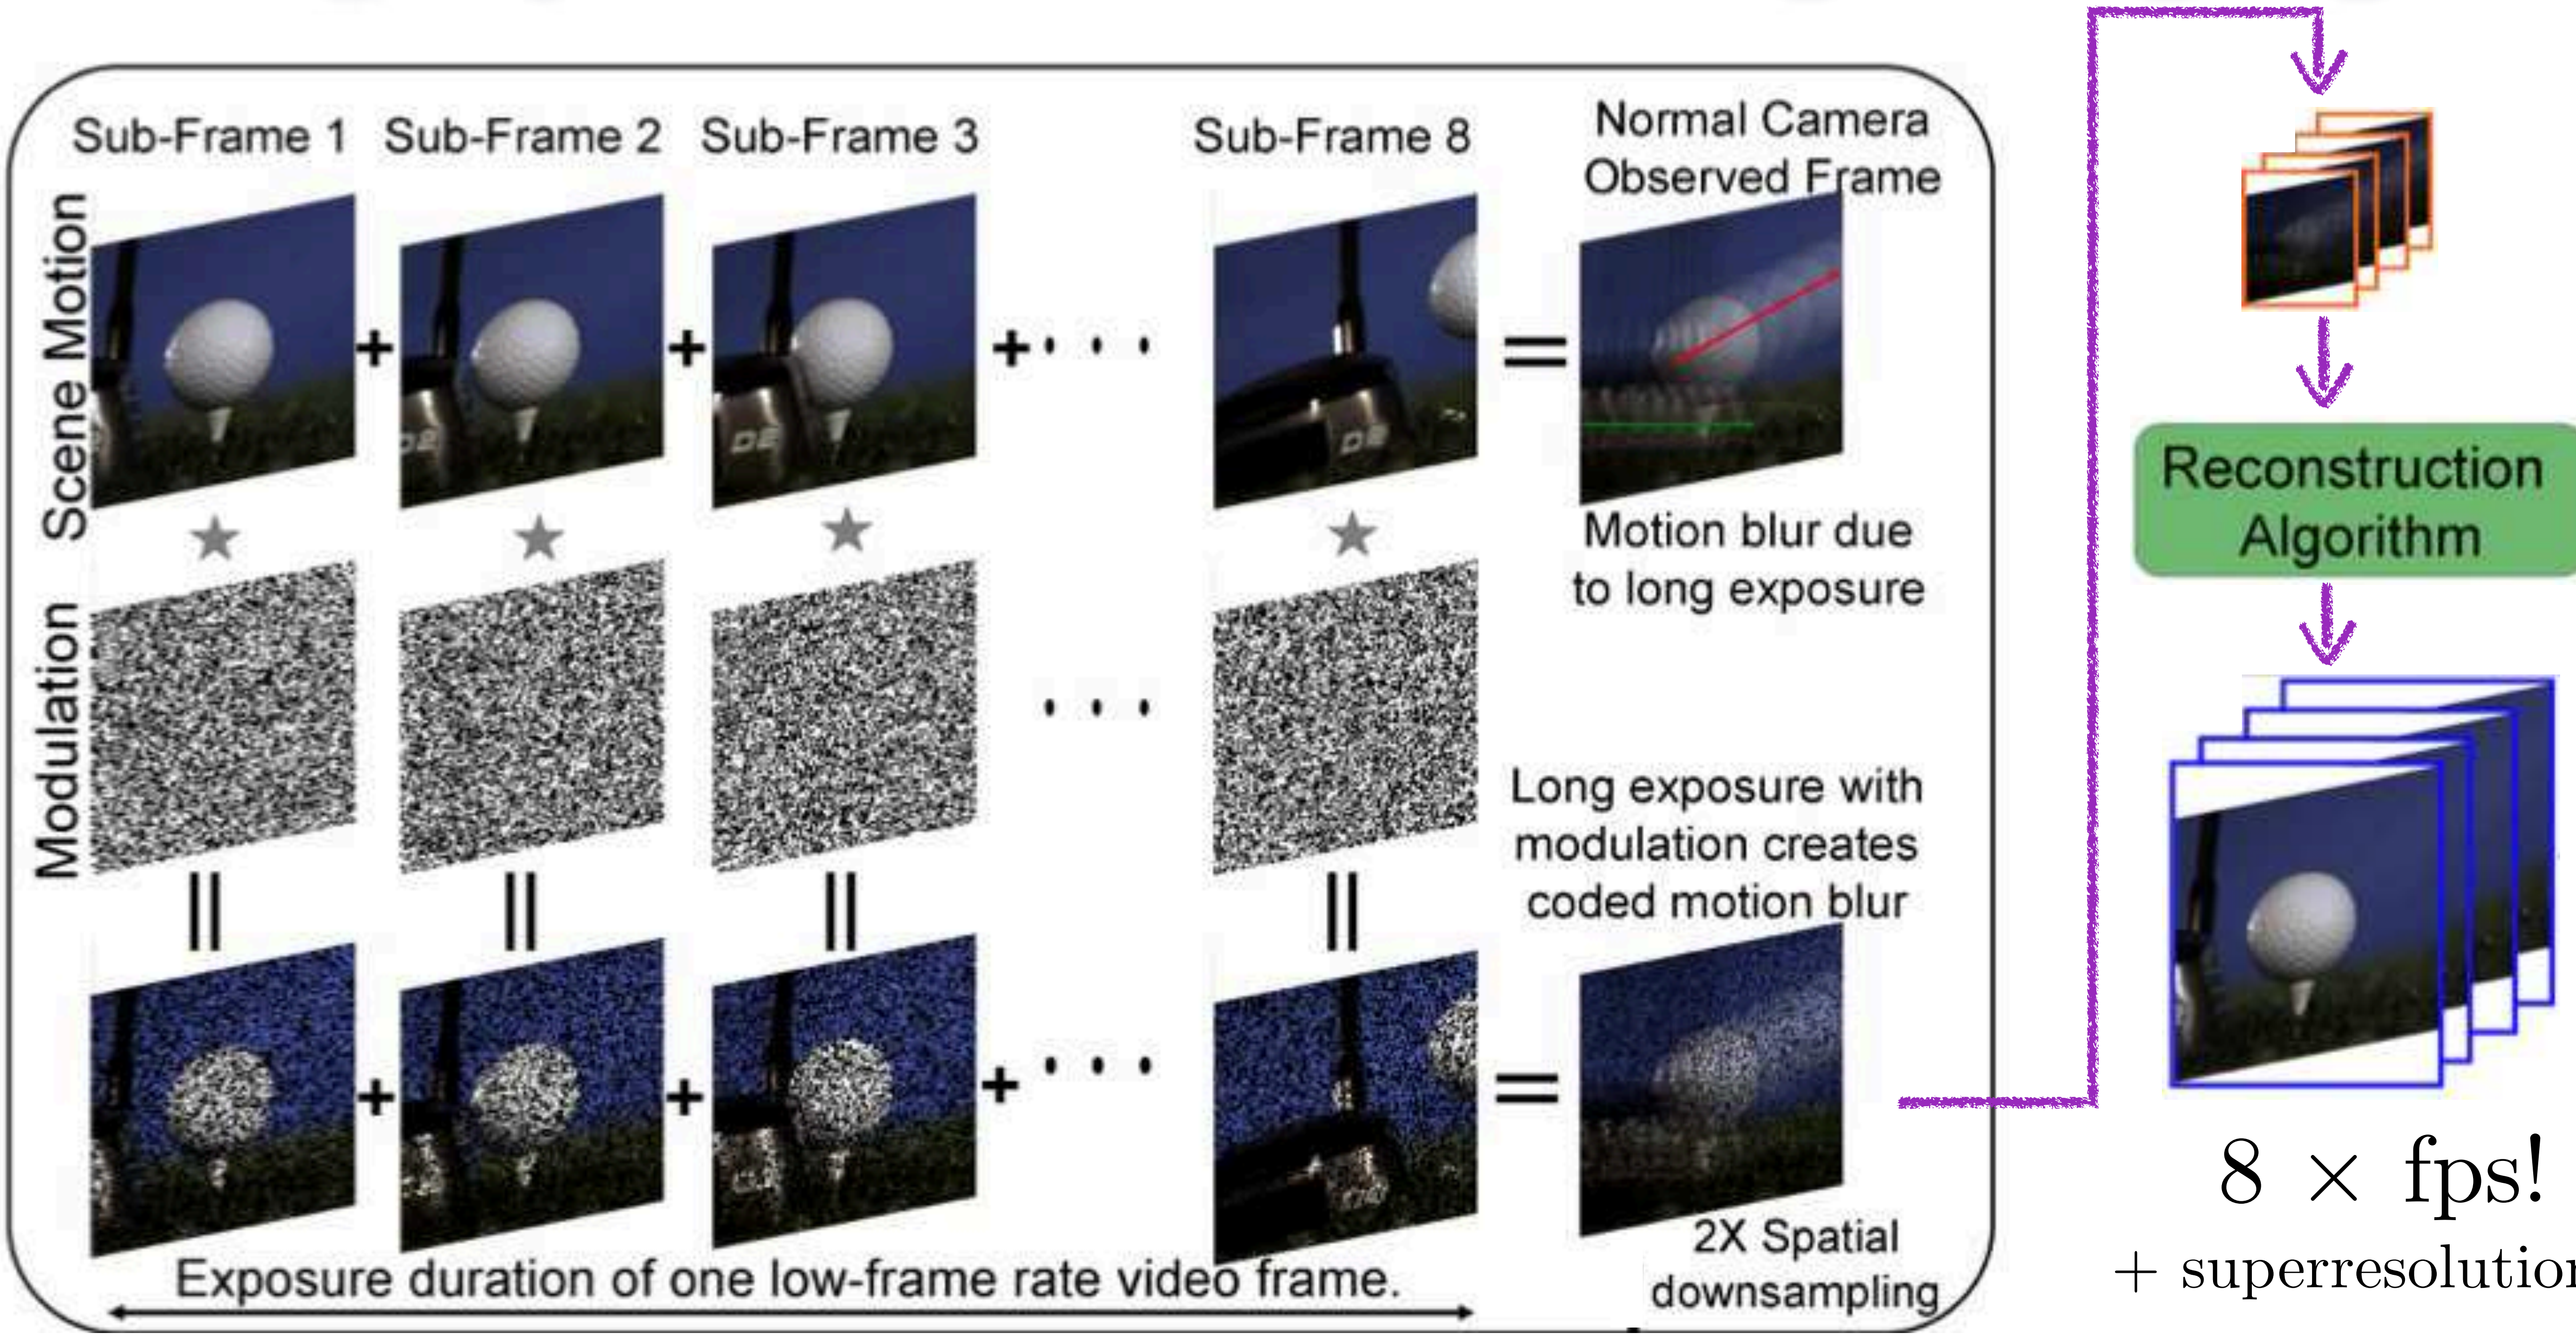

$8 \times$ fps!

+ superresolution

R. Dikpal, A. Veeraraghavan, and R. Chellappa. "P2C2: Programmable pixel compressive camera for high speed imaging" Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. IEEE, 2011.

Highspeed Coded Strobing Imaging

Optically:

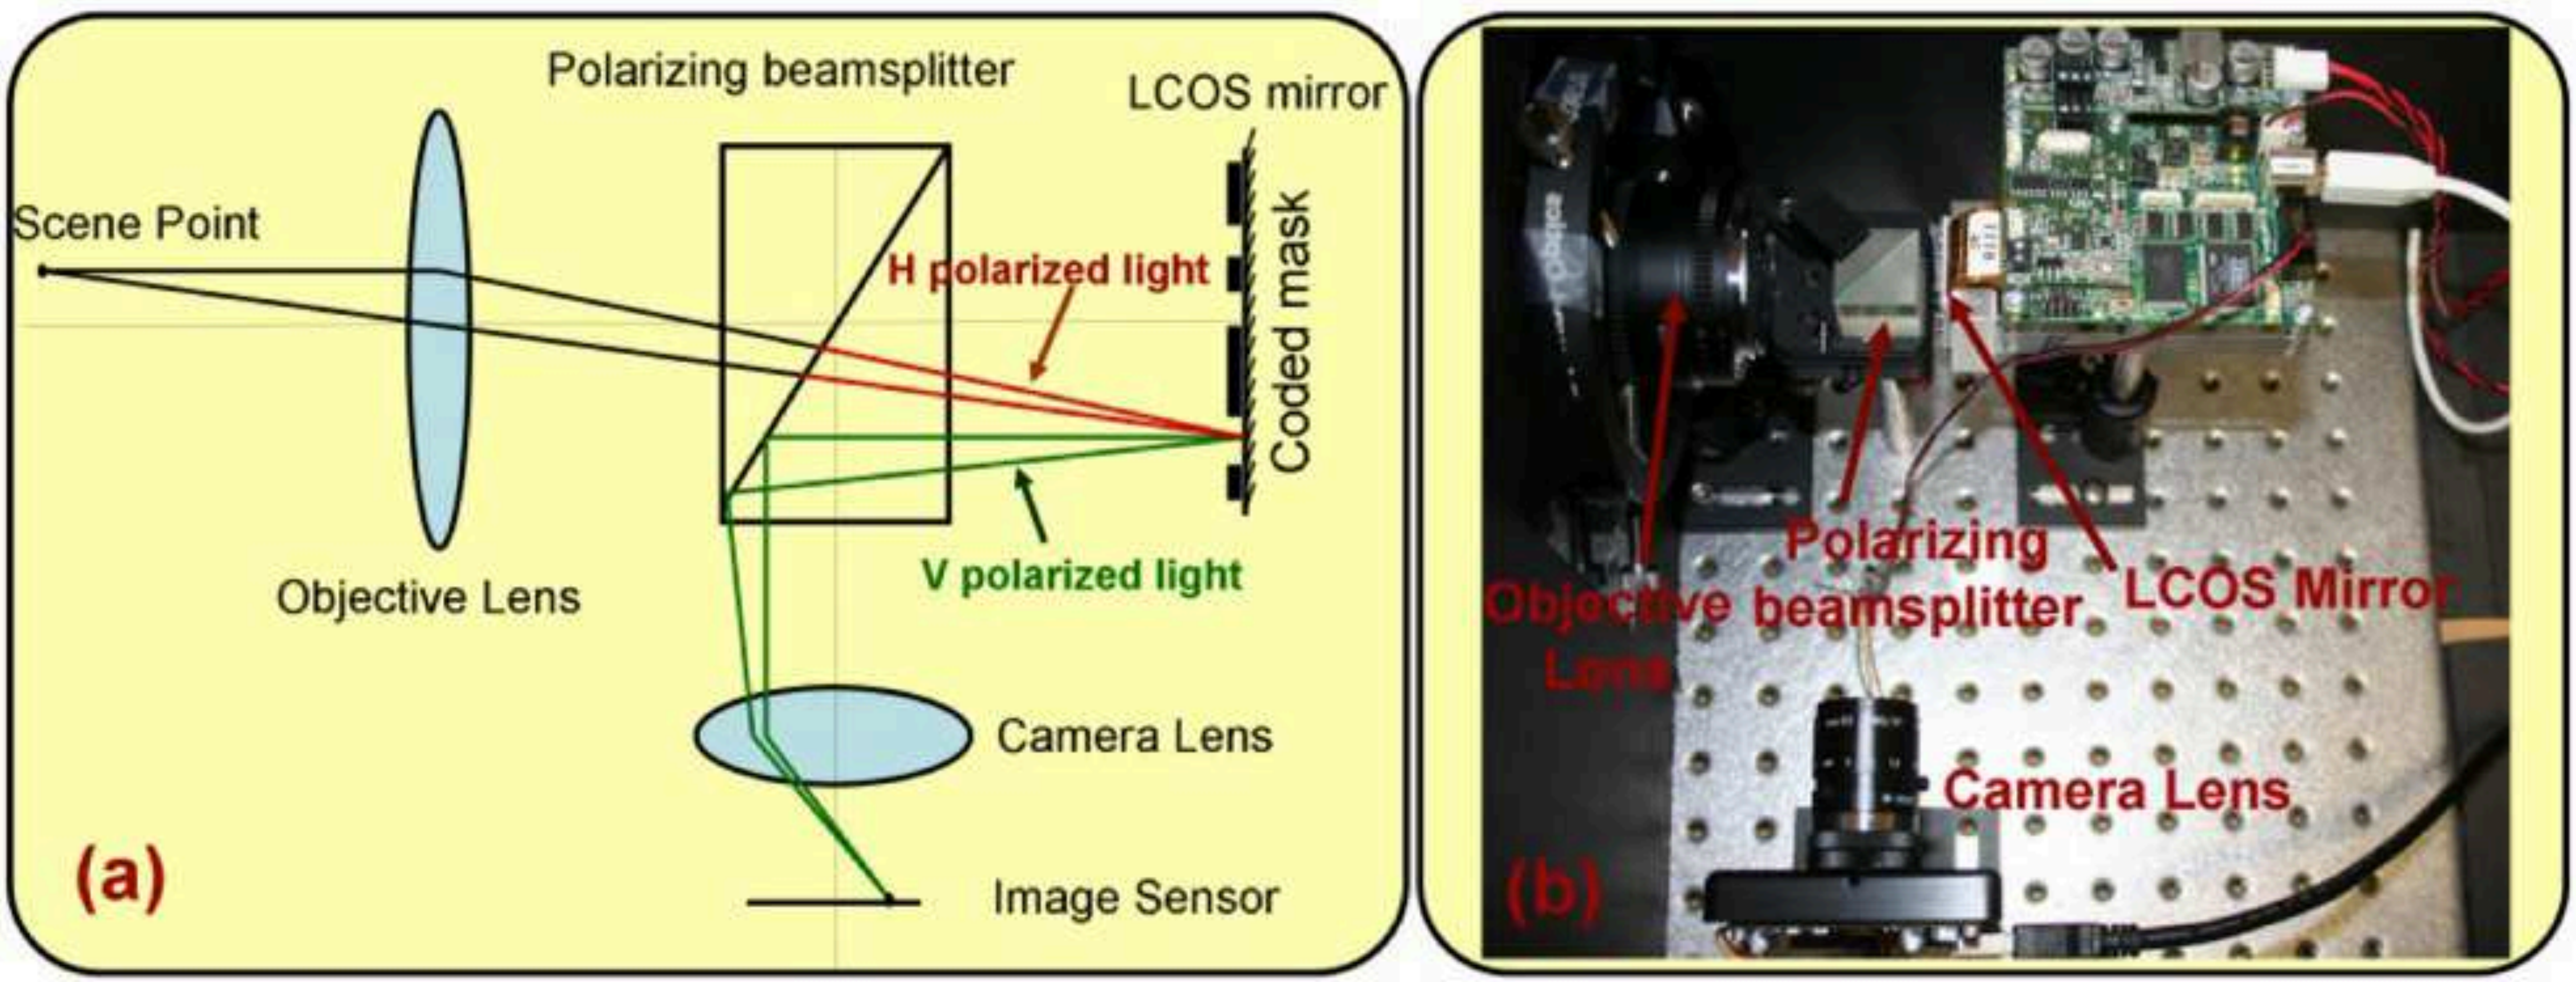

R. Dikpal, A. Veeraraghavan, and R. Chellappa. "P2C2: Programmable pixel compressive camera for high speed imaging" Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. IEEE, 2011.

Highspeed Coded Strobing Imaging

- Reconstruction: regularized with optical flow

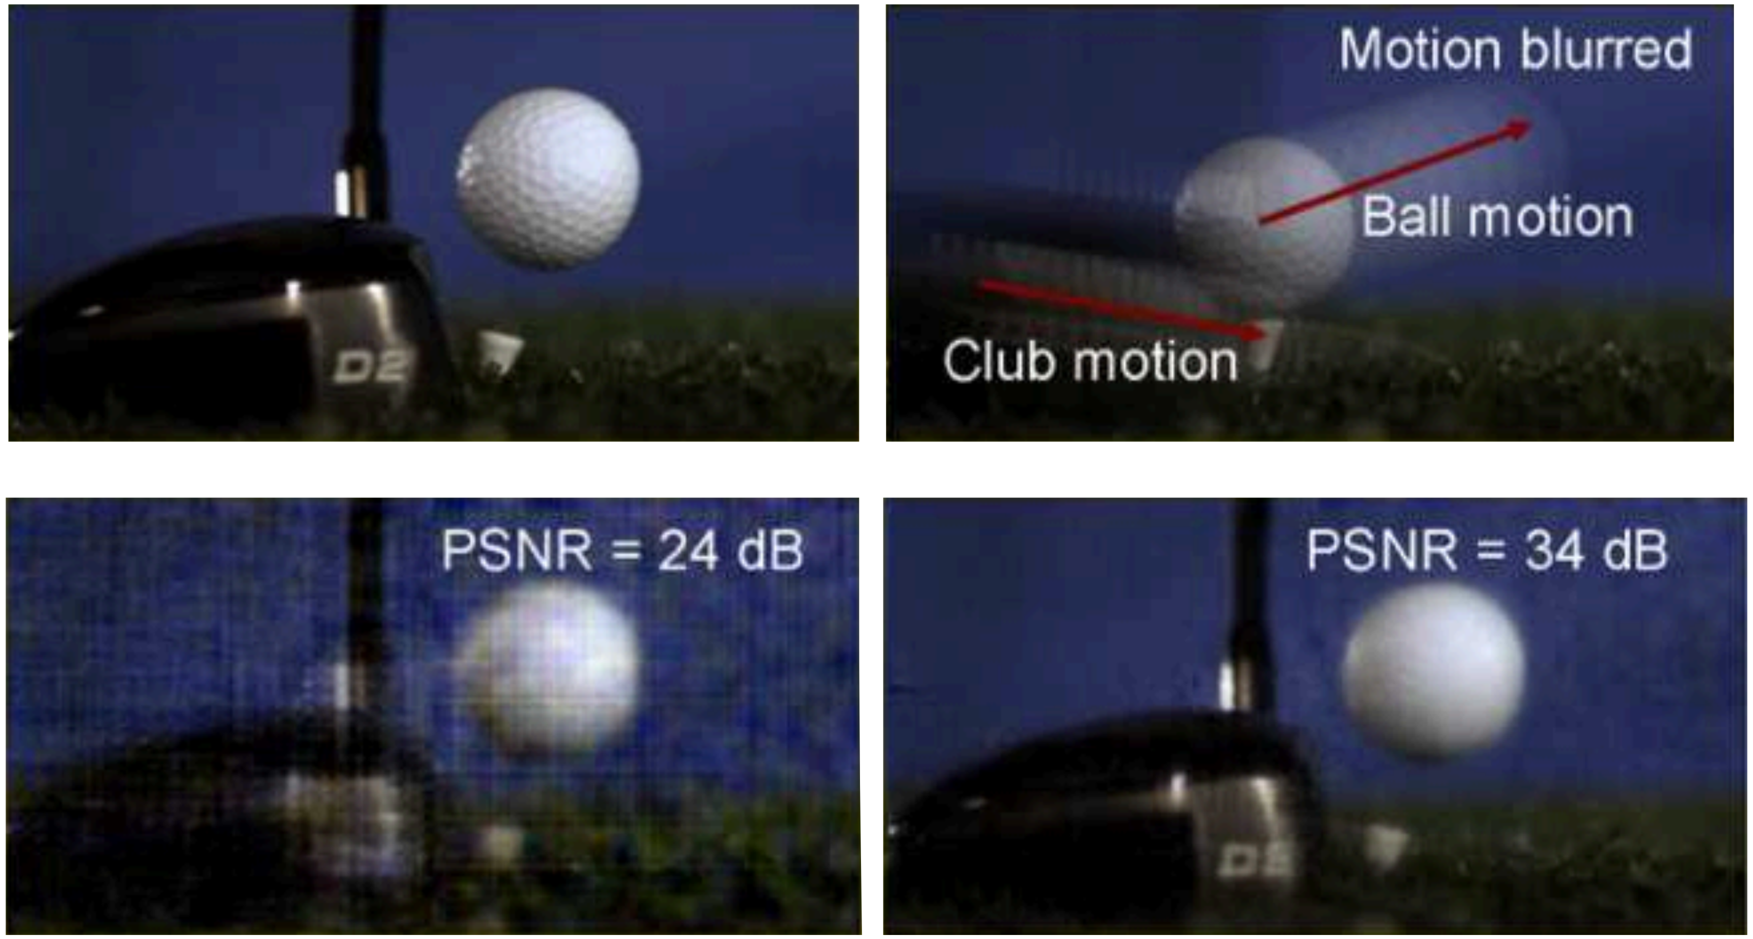

with wavelet prior   + optical flow reg.

R. Dikpal, A. Veeraraghavan, and R. Chellappa. "P2C2: Programmable pixel compressive camera for high speed imaging" Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. IEEE, 2011.

Conclusions

Conclusion

- Sparsity prior involves new sensing methods:
  e.g., Compressed Sensing, Compressive Imaging.
- Future:
  - More sensing examples: http://nuit-blanche.blogspot.com
    hyperspectral, network, GPR, Lidar, ... (explosion)
  - Better sparsity prior:
    structured, model-based, mixed-norm (Cevher, Bach, ...)
    co-sparsity/analysis model (Gribonval, Nam, Davies, Elad, Candes)
  - Non-linear sensing models ?
    1-bit CS is one instance, phase recovery (Candès),
    polychromatic CT, ...

Links (Science 2.0.)

- Rice CS Resources page:
  http://www-dsp.rice.edu/cs

- Igor Carron's
  "*Nuit Blanche*" blog:
  http://nuit-blanche.blogspot.com

  1 CS post/day!

Thank you!