Introduction to Compressed Sensing

Eric W. Tramel

*“**Biophysique:** de la Mesure au Modèle en biologie”*
École de Physique des Houches
19 October 2015

Warm-Up: 12 Balls

Problem: 12 Balls, 11 of which are of equal weight.
One *outlier* is either heavier or lighter.

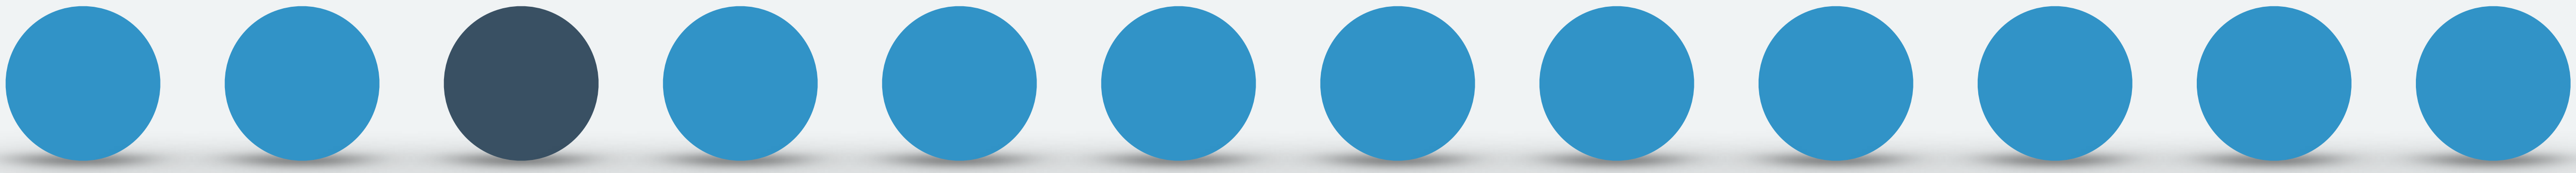

Given a balance, what is the least number of weighings (*tests*) to identify the outlier & know if it is heavier or lighter?

Group-Testing

Context: U.S.A. is drafting soldiers for WW2, but wants to weed out syphilitic recruits.

Problem: Blood tests are expensive, and individual testing is too cost prohibitive.

Wasserman Test

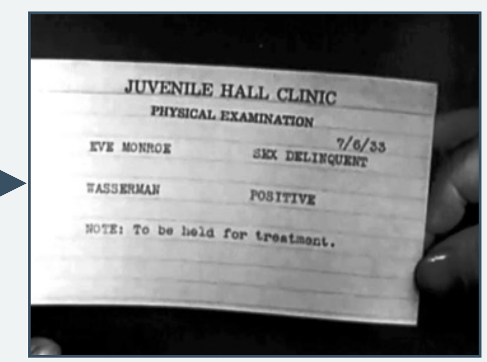

How can one effectively carry out mass-screening?

Group-Testing

[Dorfman, 1943] Test for a positive result in a mixed-sample (*pooling*) and re-test positive pools.

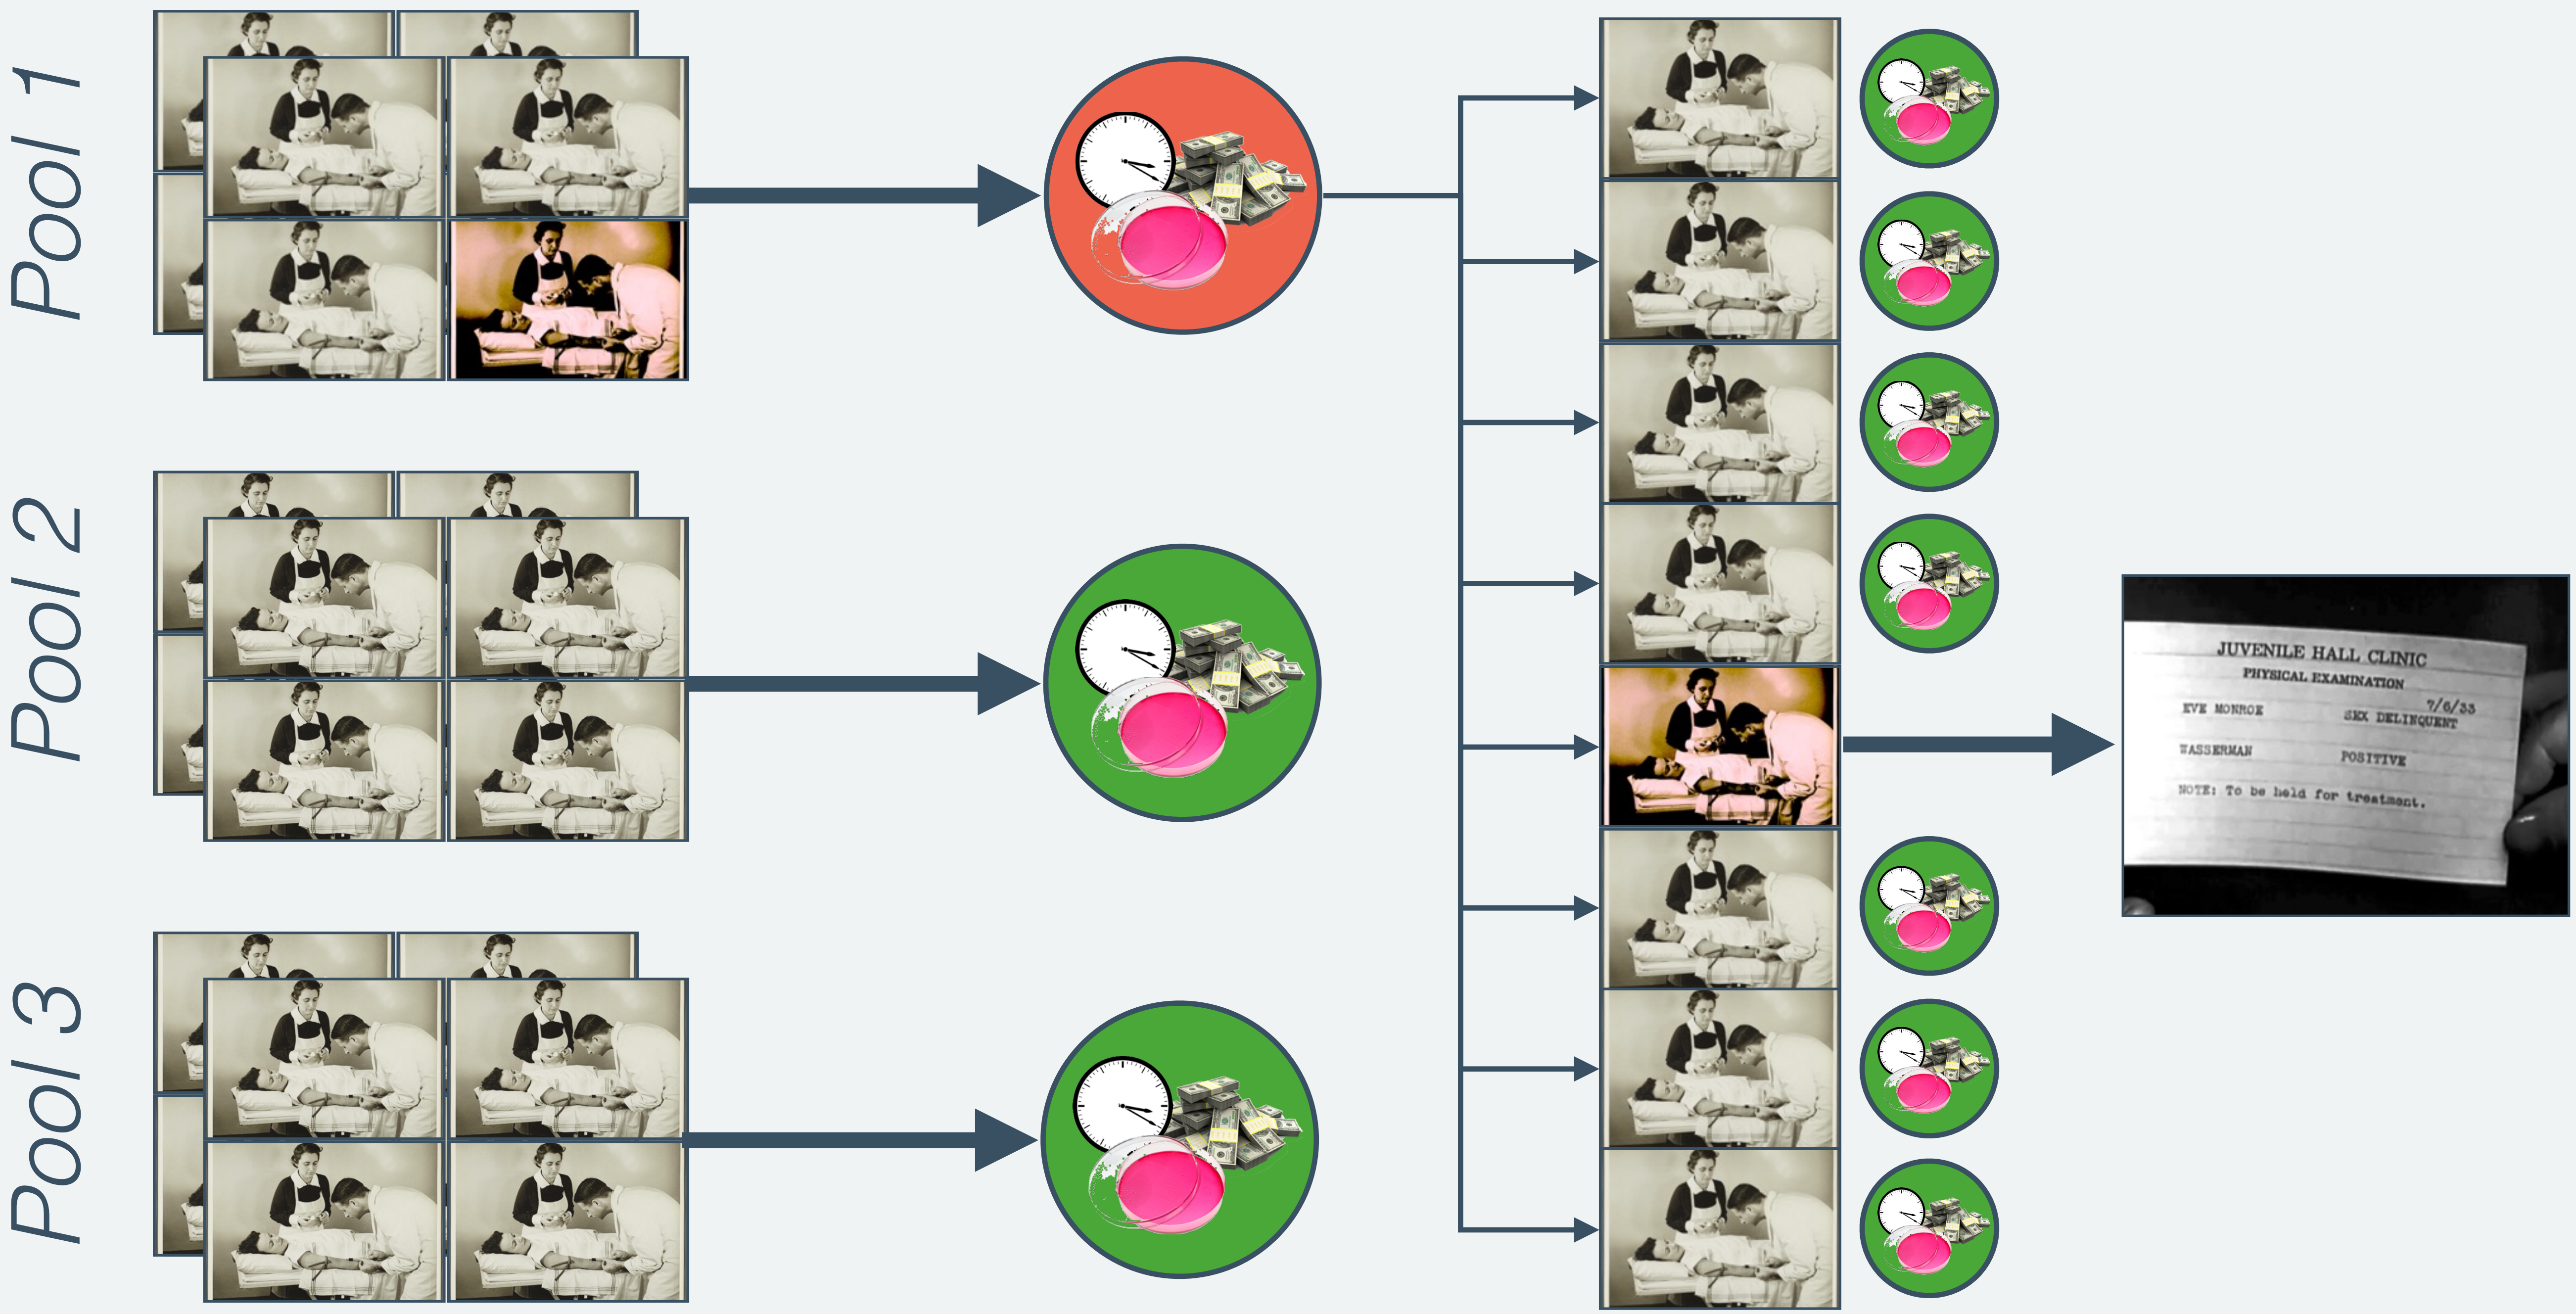

Adaptive v. Non-adaptive

Efficiency: Adaptive testing will always be at least or more efficient (in number of tests) than non-adaptive testing: makes use of intermediate information.

Practicality: What if the “soldier” moves around on assignments? What if our test destroys the sample?

- Often one cannot take advantage of re-testing.

Non-Adaptive Testing: Requires *a priori* pooling design to make the most efficient test schedule that allows accurate inference of faults without re-testing.

Linear Observation

Assuming that we can model our sampling procedure as linear, in *noiseless* setting we have,

$$\mathbf{y} = F\mathbf{x}$$

Problem: Knowing samples and the sampling design, can we know the signal?

Linear Algebra: Only if the system **F** is invertible (*square)*

Nyquist: Only if sampled at rate twice the bandwith of **x**.

Linear Algebra

[Nyquist & LA] For accurate reconstruction of **N** coefficients, one requires as many samples, **M**, as coefficients.

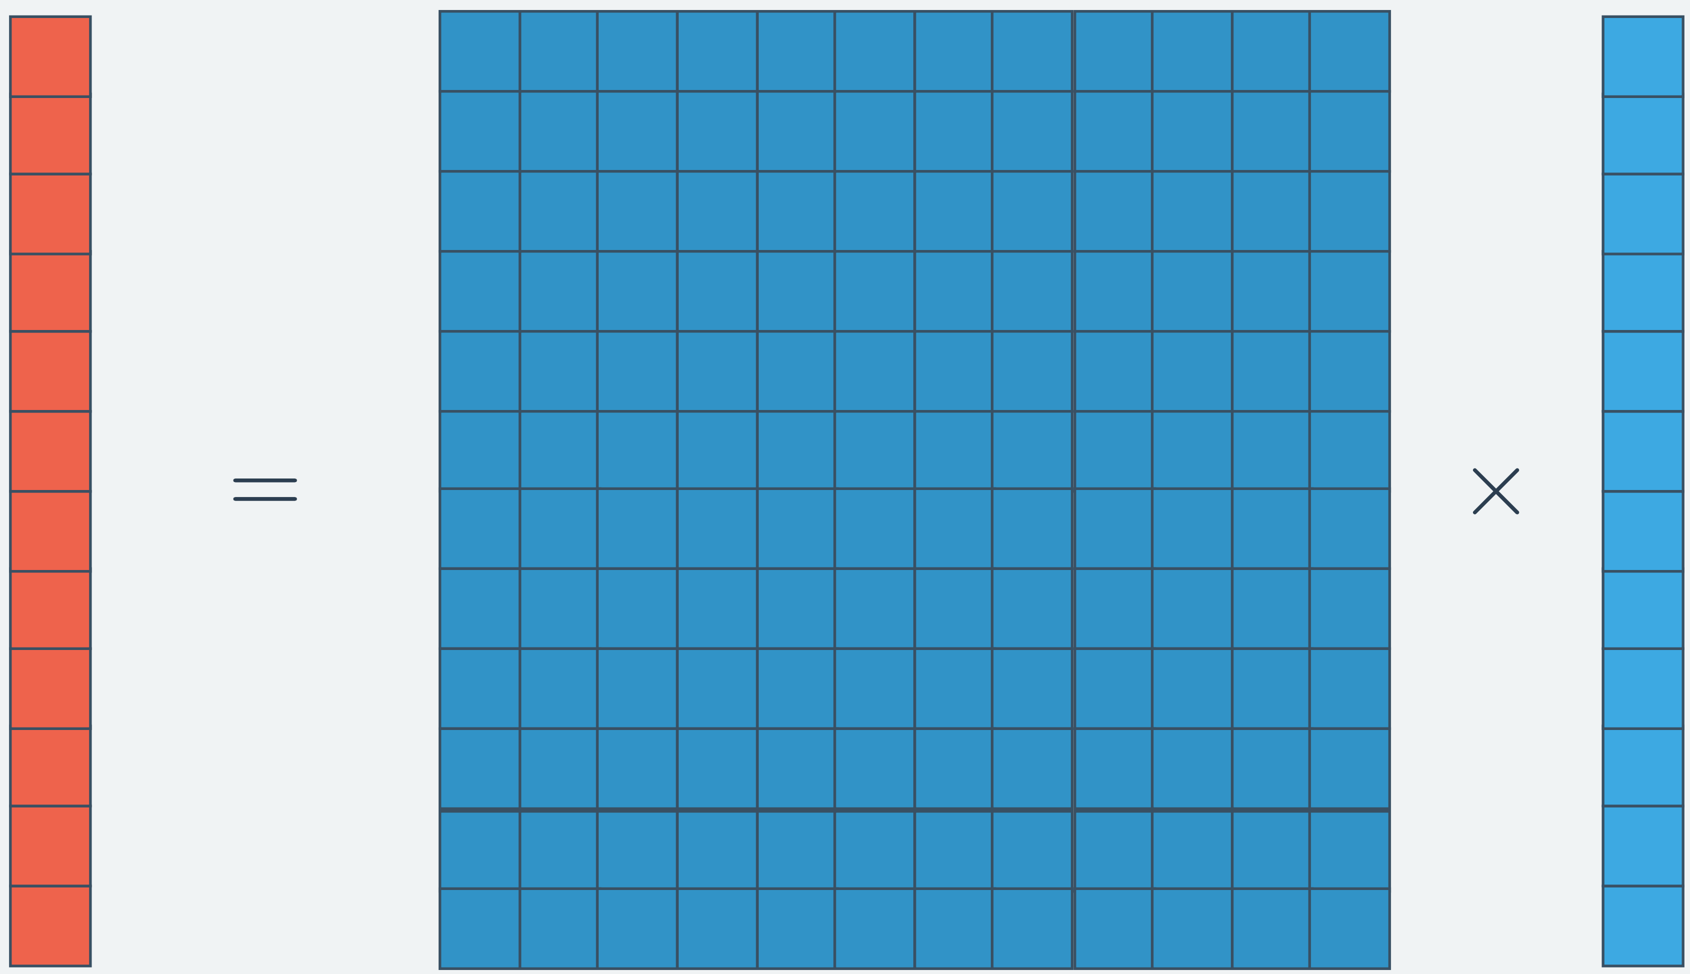

$$\mathbf{y} = F\mathbf{x}$$

Linear Algebra

Undersampling: Our goal is to reduce measurements (**M<N**). Removing measurements from **y**, **F**, makes solving for **x** impossible, *in general*.

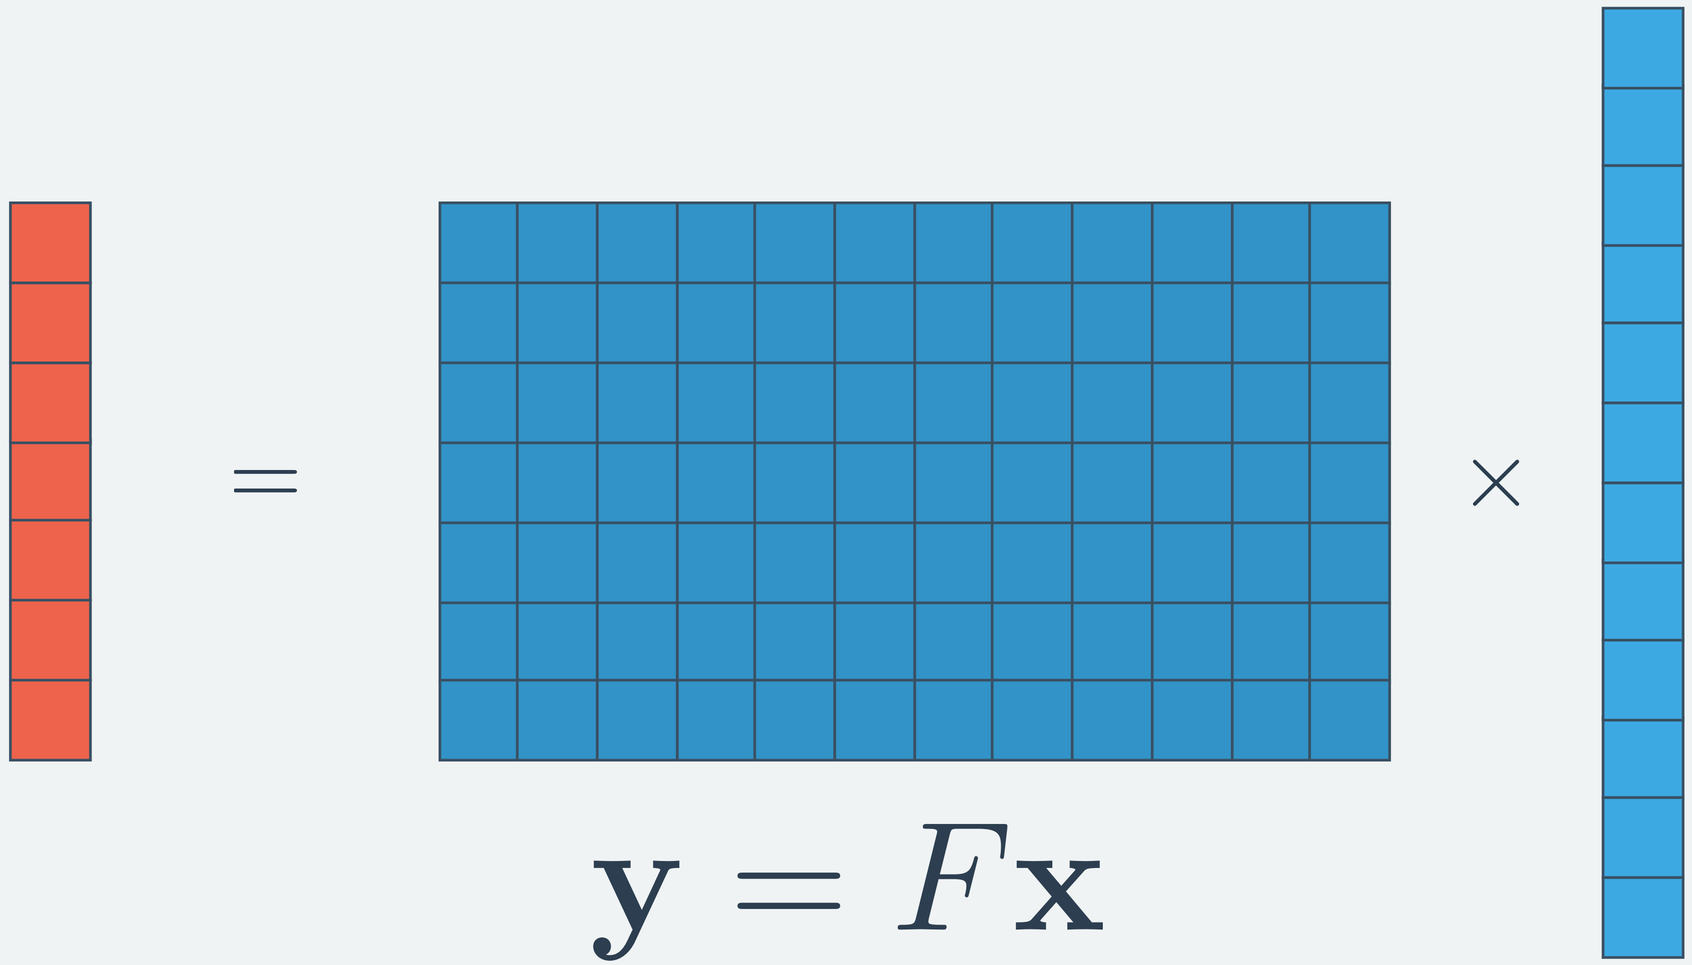

$$\mathbf{y} = F\mathbf{x}$$

An entire space of possible solutions: $\mathbf{y} = F\left(\mathbf{x} + \mathbf{s} \in \mathrm{Null}(F)\right)$

Linear Algebra

Undersampling: Our goal is to reduce measurements.
Removing me ... g for
x impossible, ...

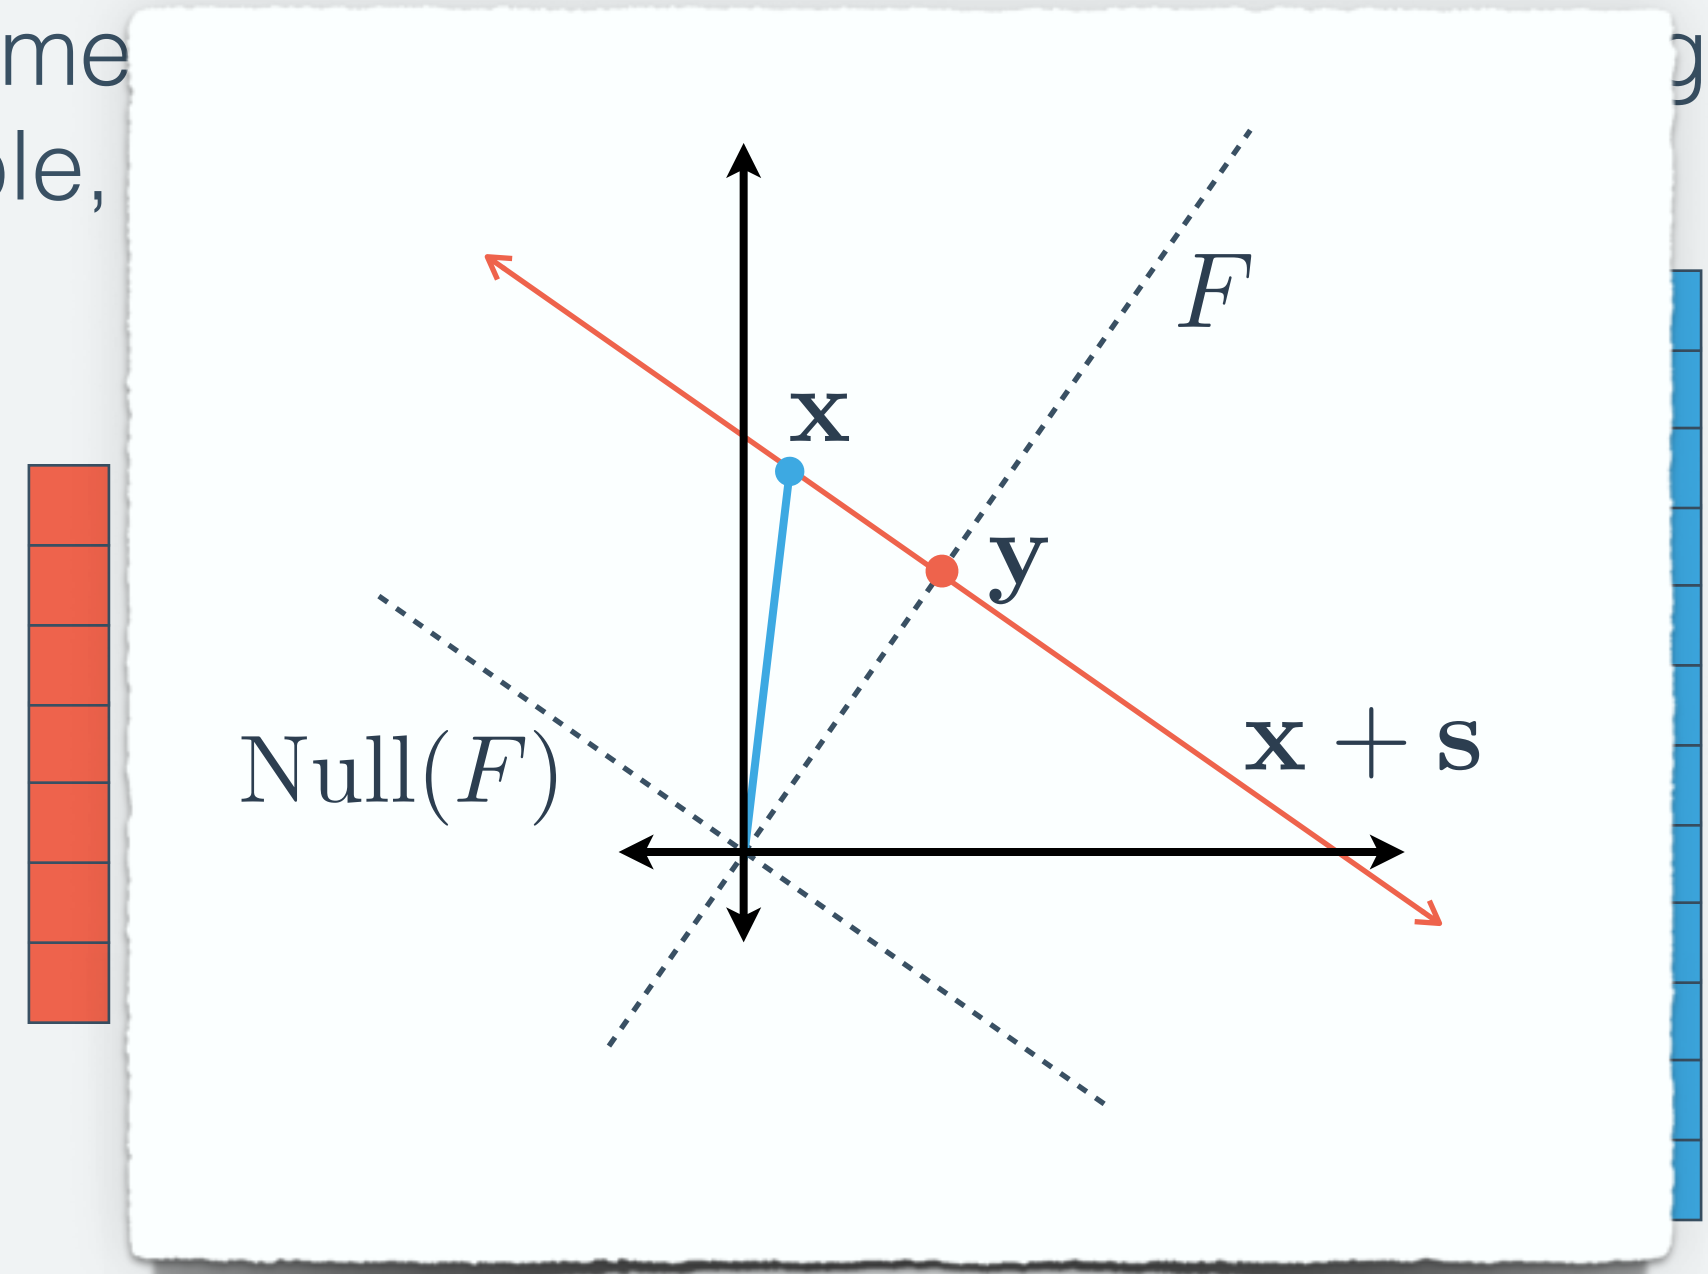

An entire space of possible solutions: $\mathbf{y} = F\left(\mathbf{x} + \mathbf{s} \in \mathrm{Null}(F)\right)$

Linear Algebra

Prior Knowledge: We can't get something for free, but we by imposing what knowledge we have *a prior.*

E.g. Shrinkage… $||\mathbf{x}||_2 \leq \epsilon$

$$\widehat{\mathbf{x}} = \left(F^T F + \lambda I\right)^{-1} F^T \mathbf{y}$$

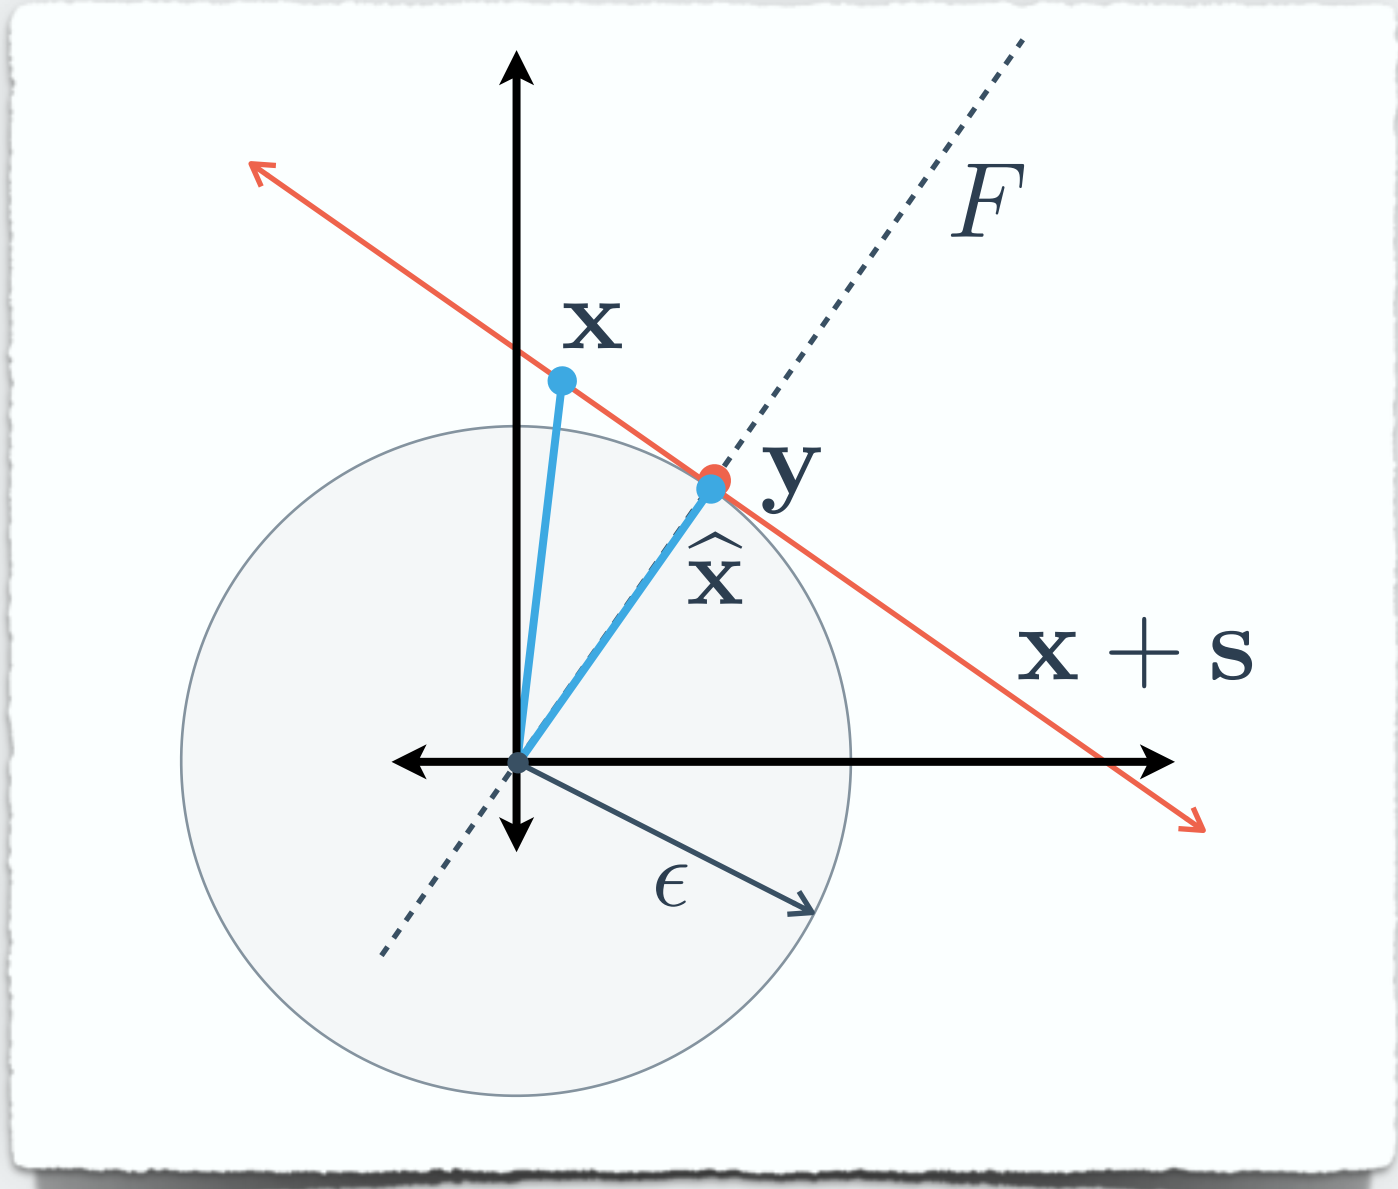

Undersampling of Sparse Signals

Prior Knowledge: A more interesting/useful case, what about a sparse prior?

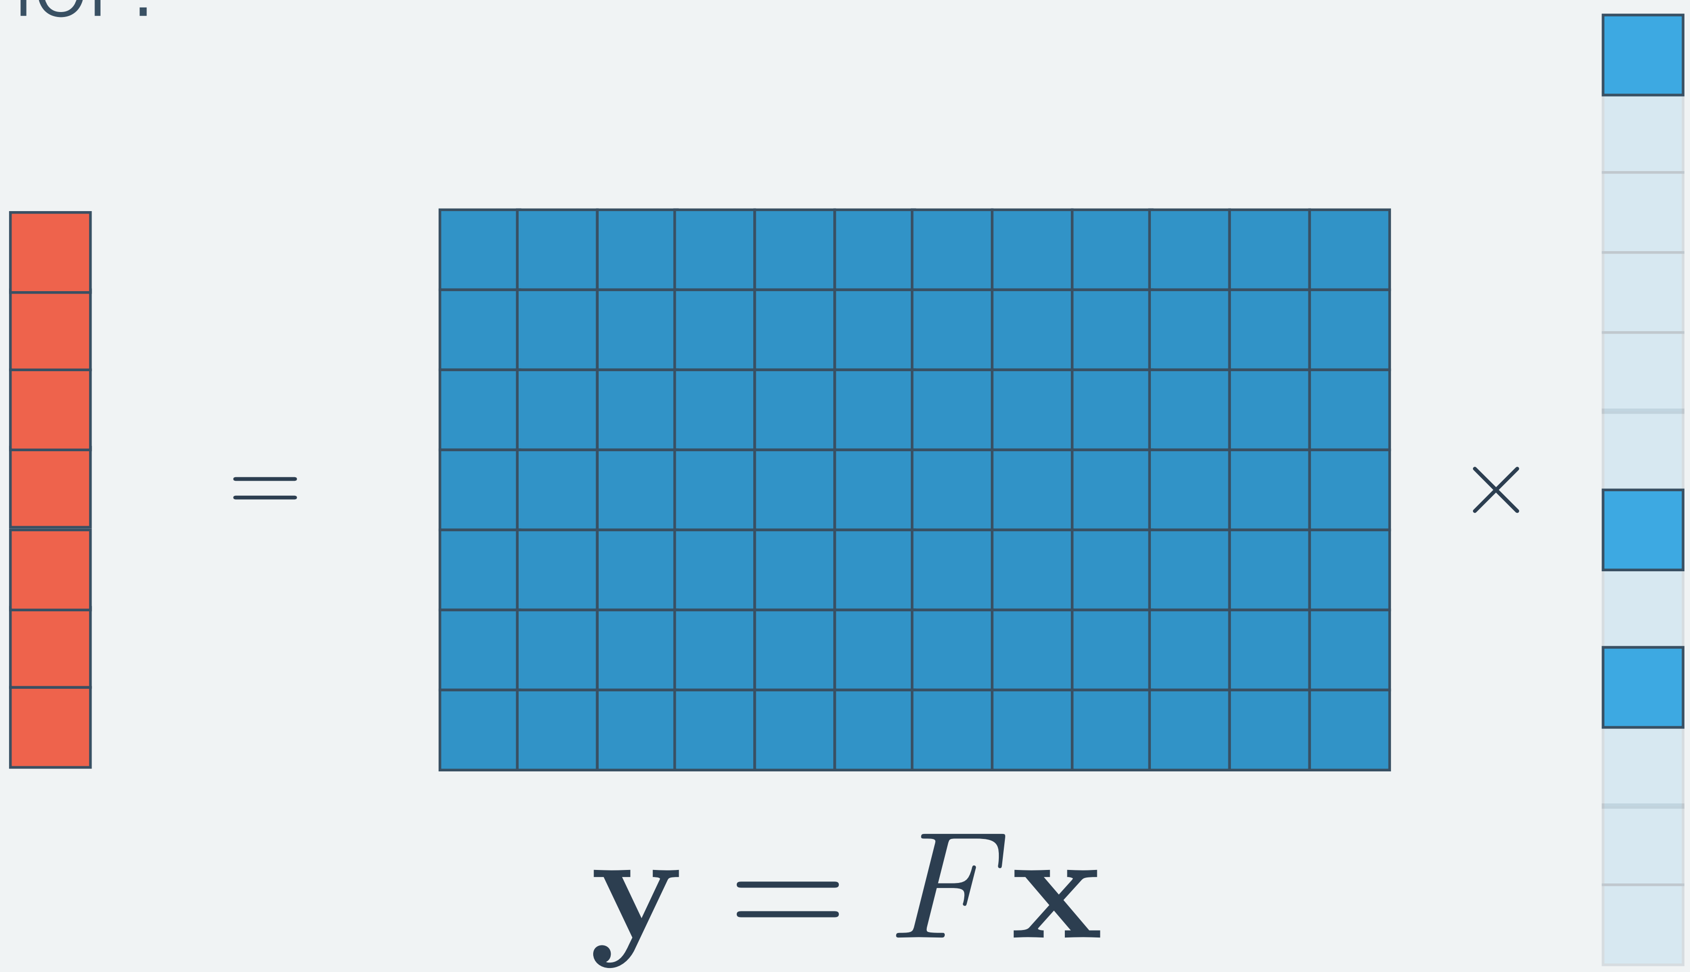

$$\mathbf{y} = F\mathbf{x}$$

K-Sparse: Signal **x** has **K** non-zero elements.

ρ-Dense: Coefficients of **x** are non-zero with probability **ρ**.

Support: Location of non-zero elements.

ASIDE: Sparsity & Information

In (*very*) General: If a signal is interesting or informative, it probably admits a *parsimonious* (simple) description.

- Has some identifiable pattern (*ordered)*.
- Is distinguishable from noise (*order => not max ent)*.

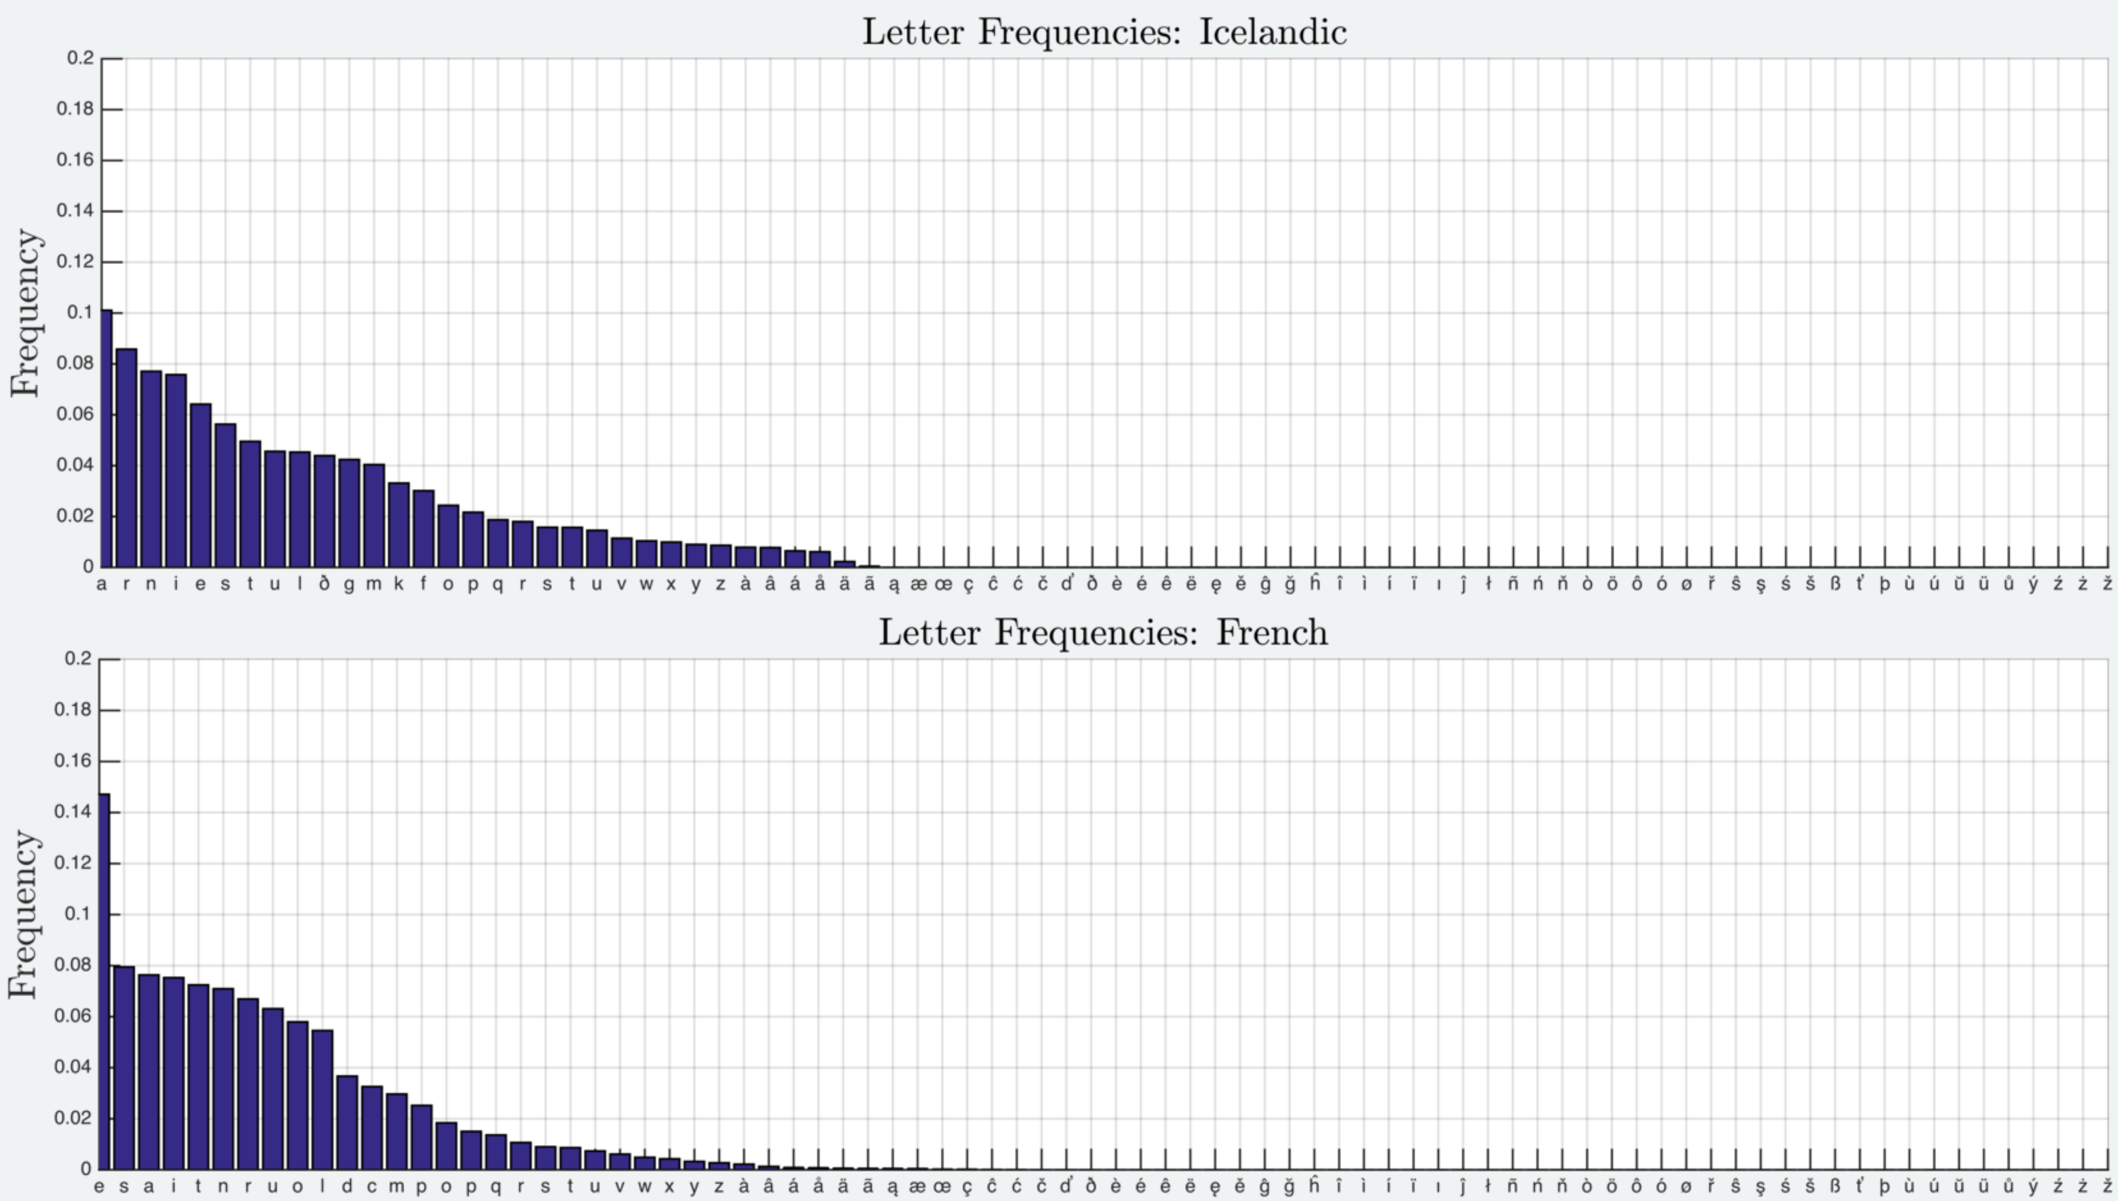

Undersampling of Sparse Signals

If support were known *a priori*, for **M>K,** the system is in fact *overdetermined*, and can be solved <u>exactly</u> in the noiseless setting!

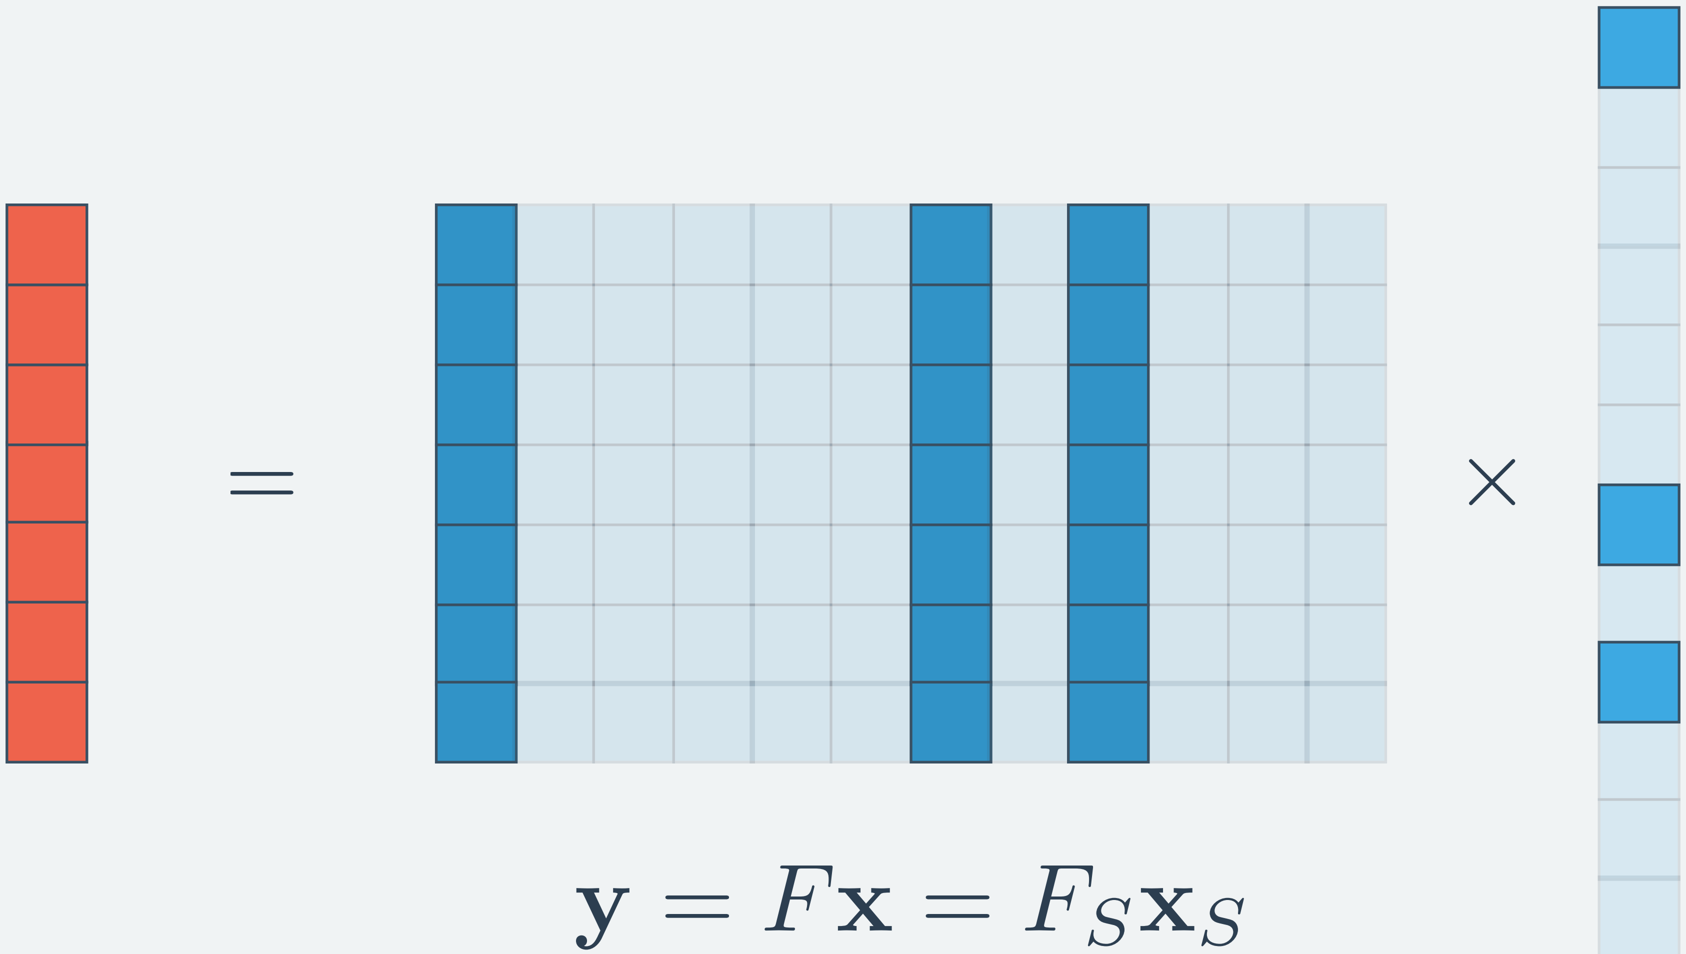

$$\mathbf{y} = F\mathbf{x} = F_S\mathbf{x}_S$$

Undersampling of Sparse Signals

If support were known *a priori*, for **M>K,** the system is in fact *overdetermined*, and can be solved exactly in the noiseless setting!

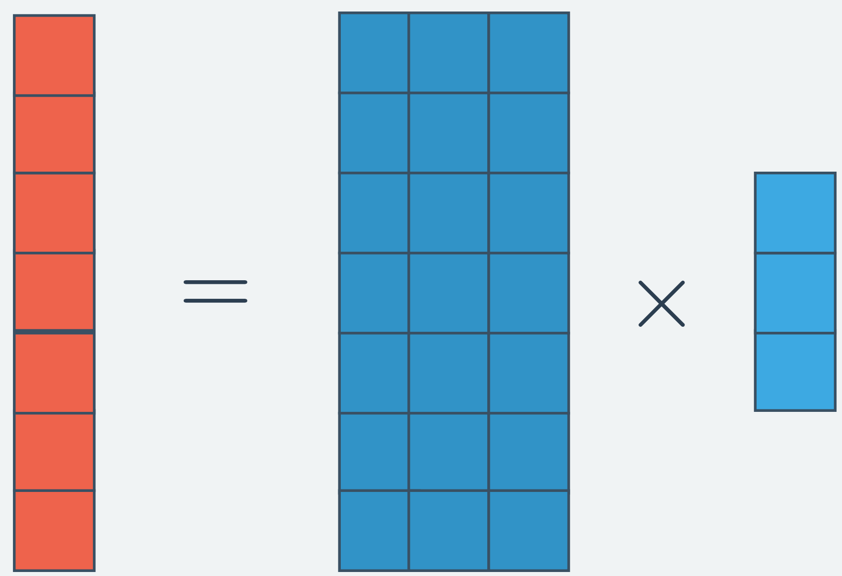

$$\mathbf{y} = F\mathbf{x} = F_S\mathbf{x}_S$$

Big "if" however…need to design **F** such that jointly

* Support can be detected.
* On-support coefficients can be estimated.

Designing Sampling

Perfect Sampling (**M=N**)

$$F = [\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_M]^T = I$$

$$y_j = \langle \mathbf{f}_j, \mathbf{x} \rangle \quad \forall \; j \in [1, M]$$

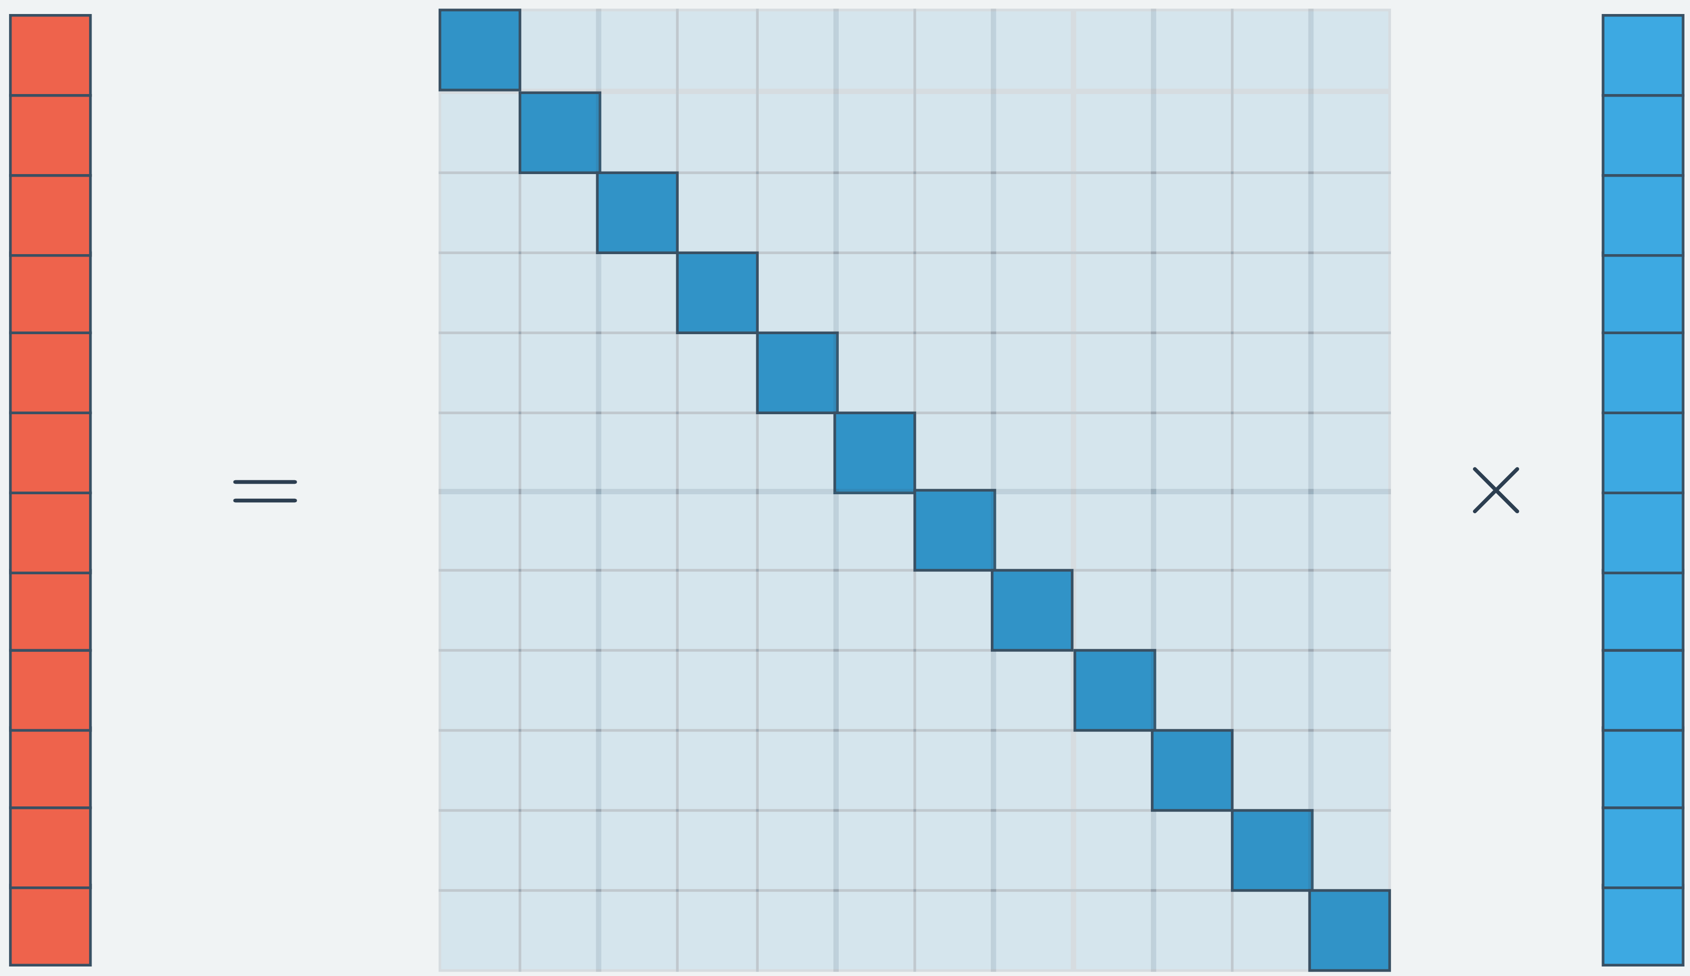

Each row of **F** is a different *measurement* of **x**. Here, a Dirac delta at each dimension of **x**.

$$\mathbf{f}_j = \delta[j - i]$$

Designing Sampling

Undersampling (**M<N**)

$$\mathbf{f}_j = \delta[j - i]$$

$$y_j = \langle \mathbf{f}_j, \mathbf{x} \rangle \quad \forall \; j \in [1, M]$$

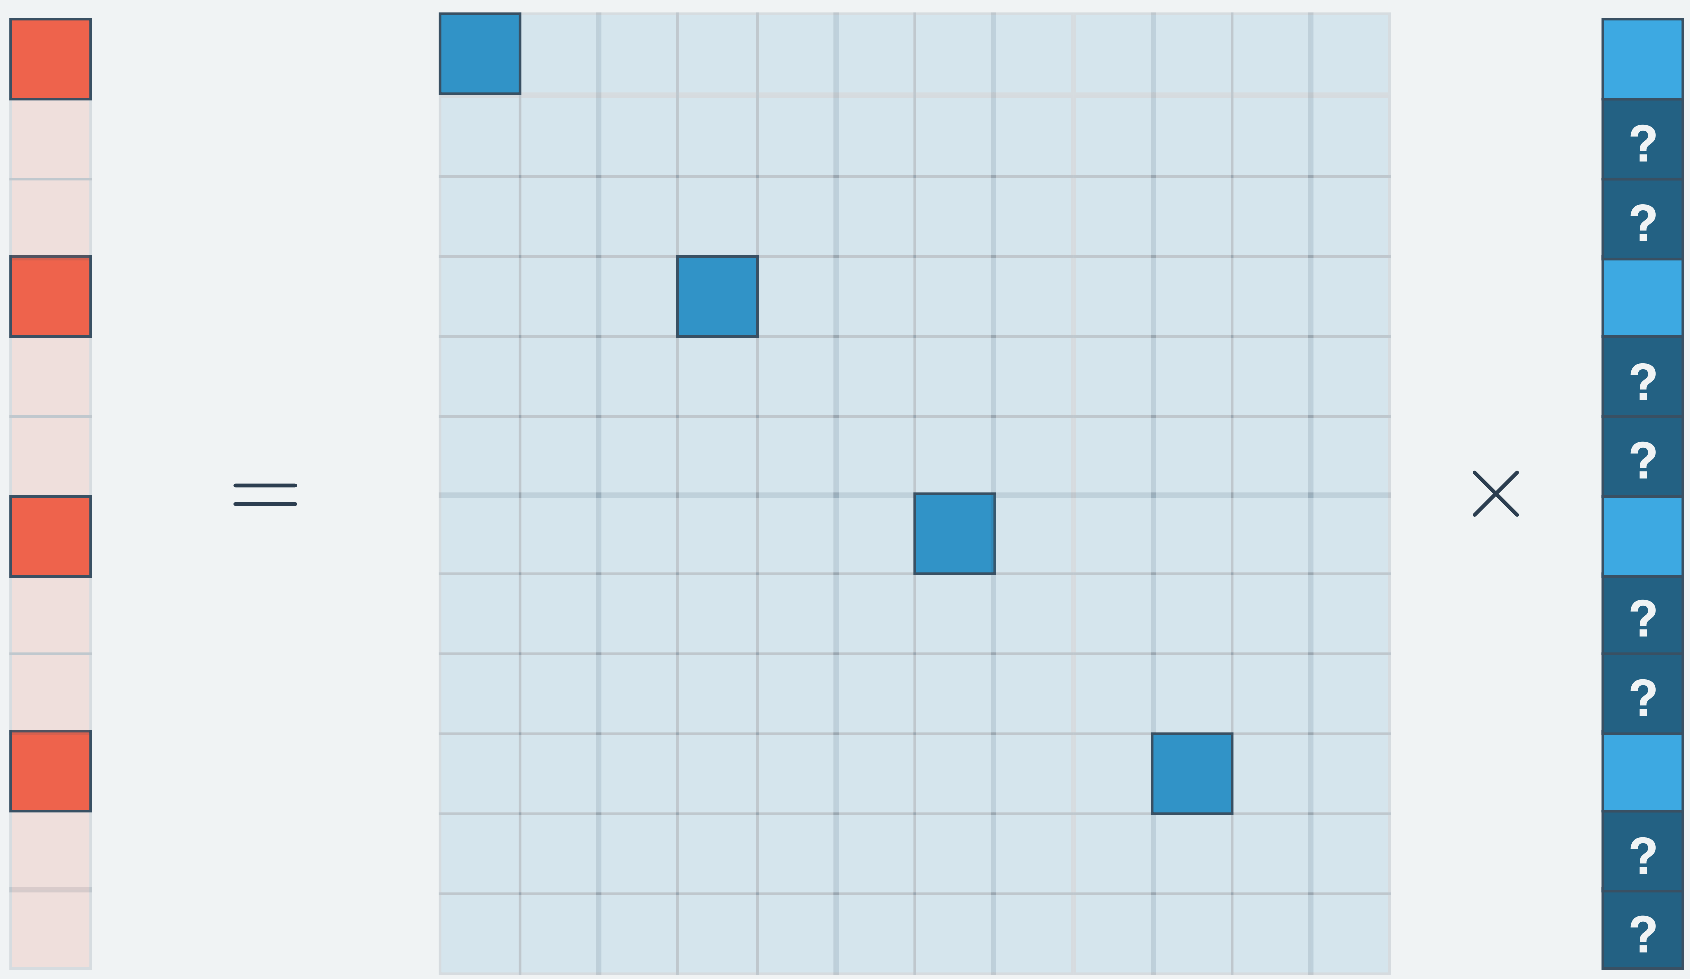

no impact on observations

Since some entries of **x** do not influence measurements, no way to recover them.

* *Their information is lost in the projection.*

Designing Sampling

Undersampling (**M<N**)

$$\mathbf{f}_j \neq \delta[j-i]$$

$$y_j = \langle \mathbf{f}_j, \mathbf{x} \rangle \quad \forall \ j \in [1, M]$$

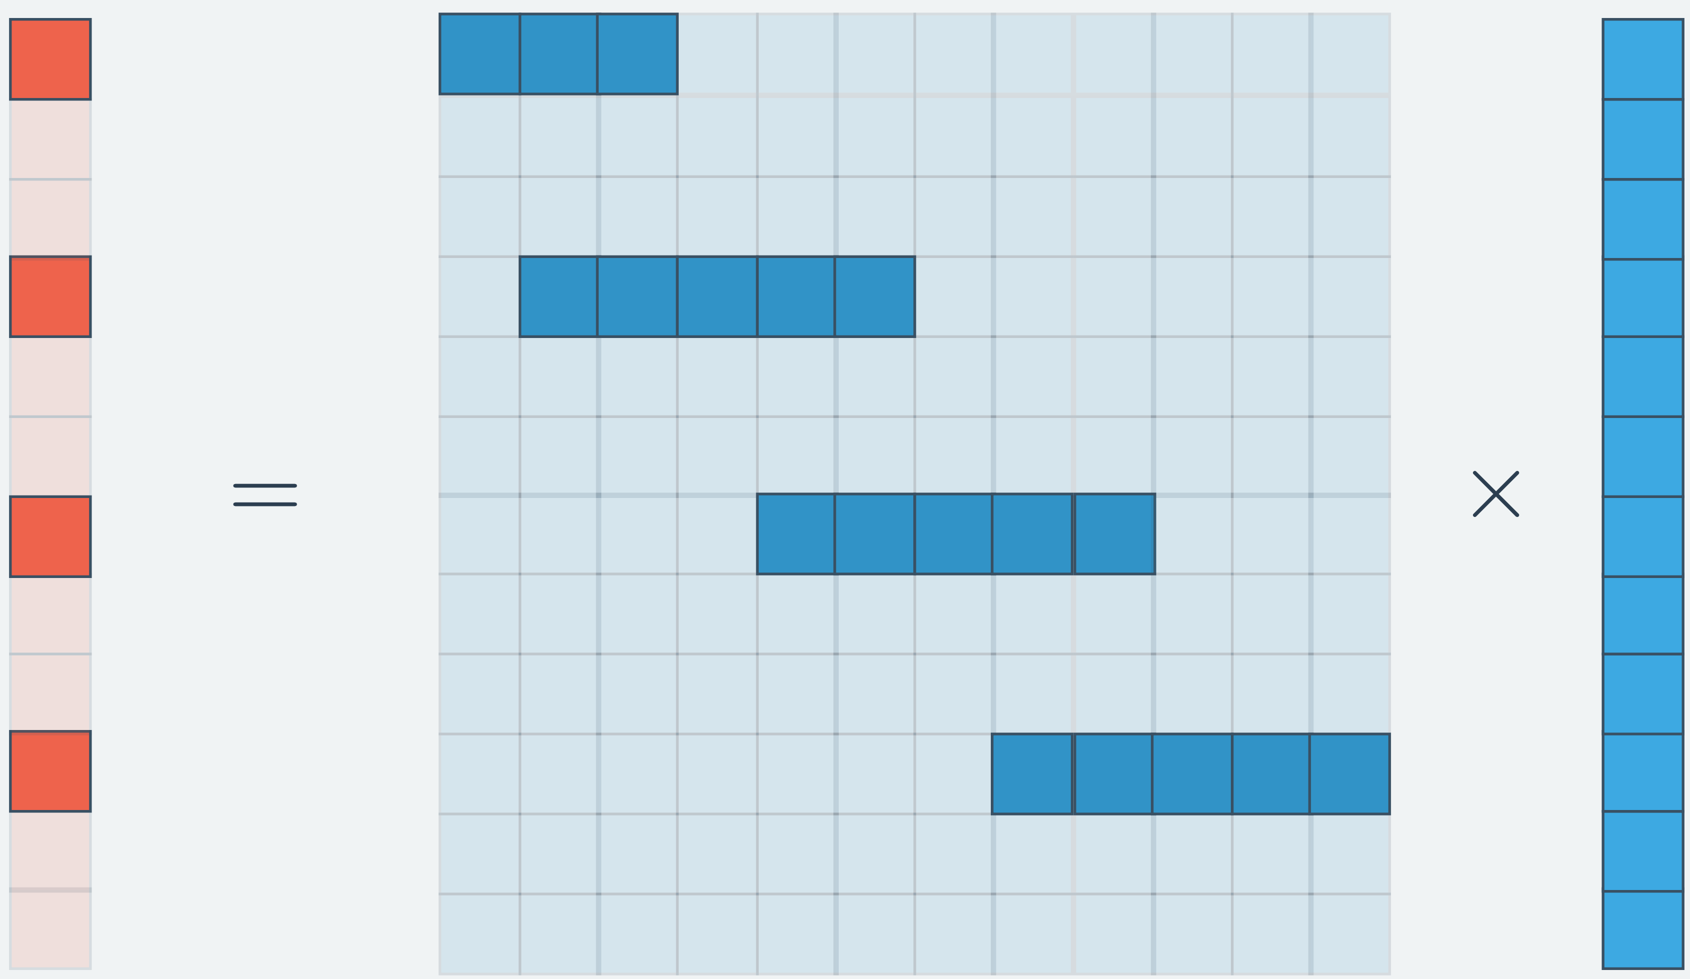

All sampled

If we choose a wider filter for **f**, like a Gaussian or Step function, we ensure all samples contribute to measurements.

Designing Sampling

Related to Anti-Aliasing:
Ex. Downsampling image…

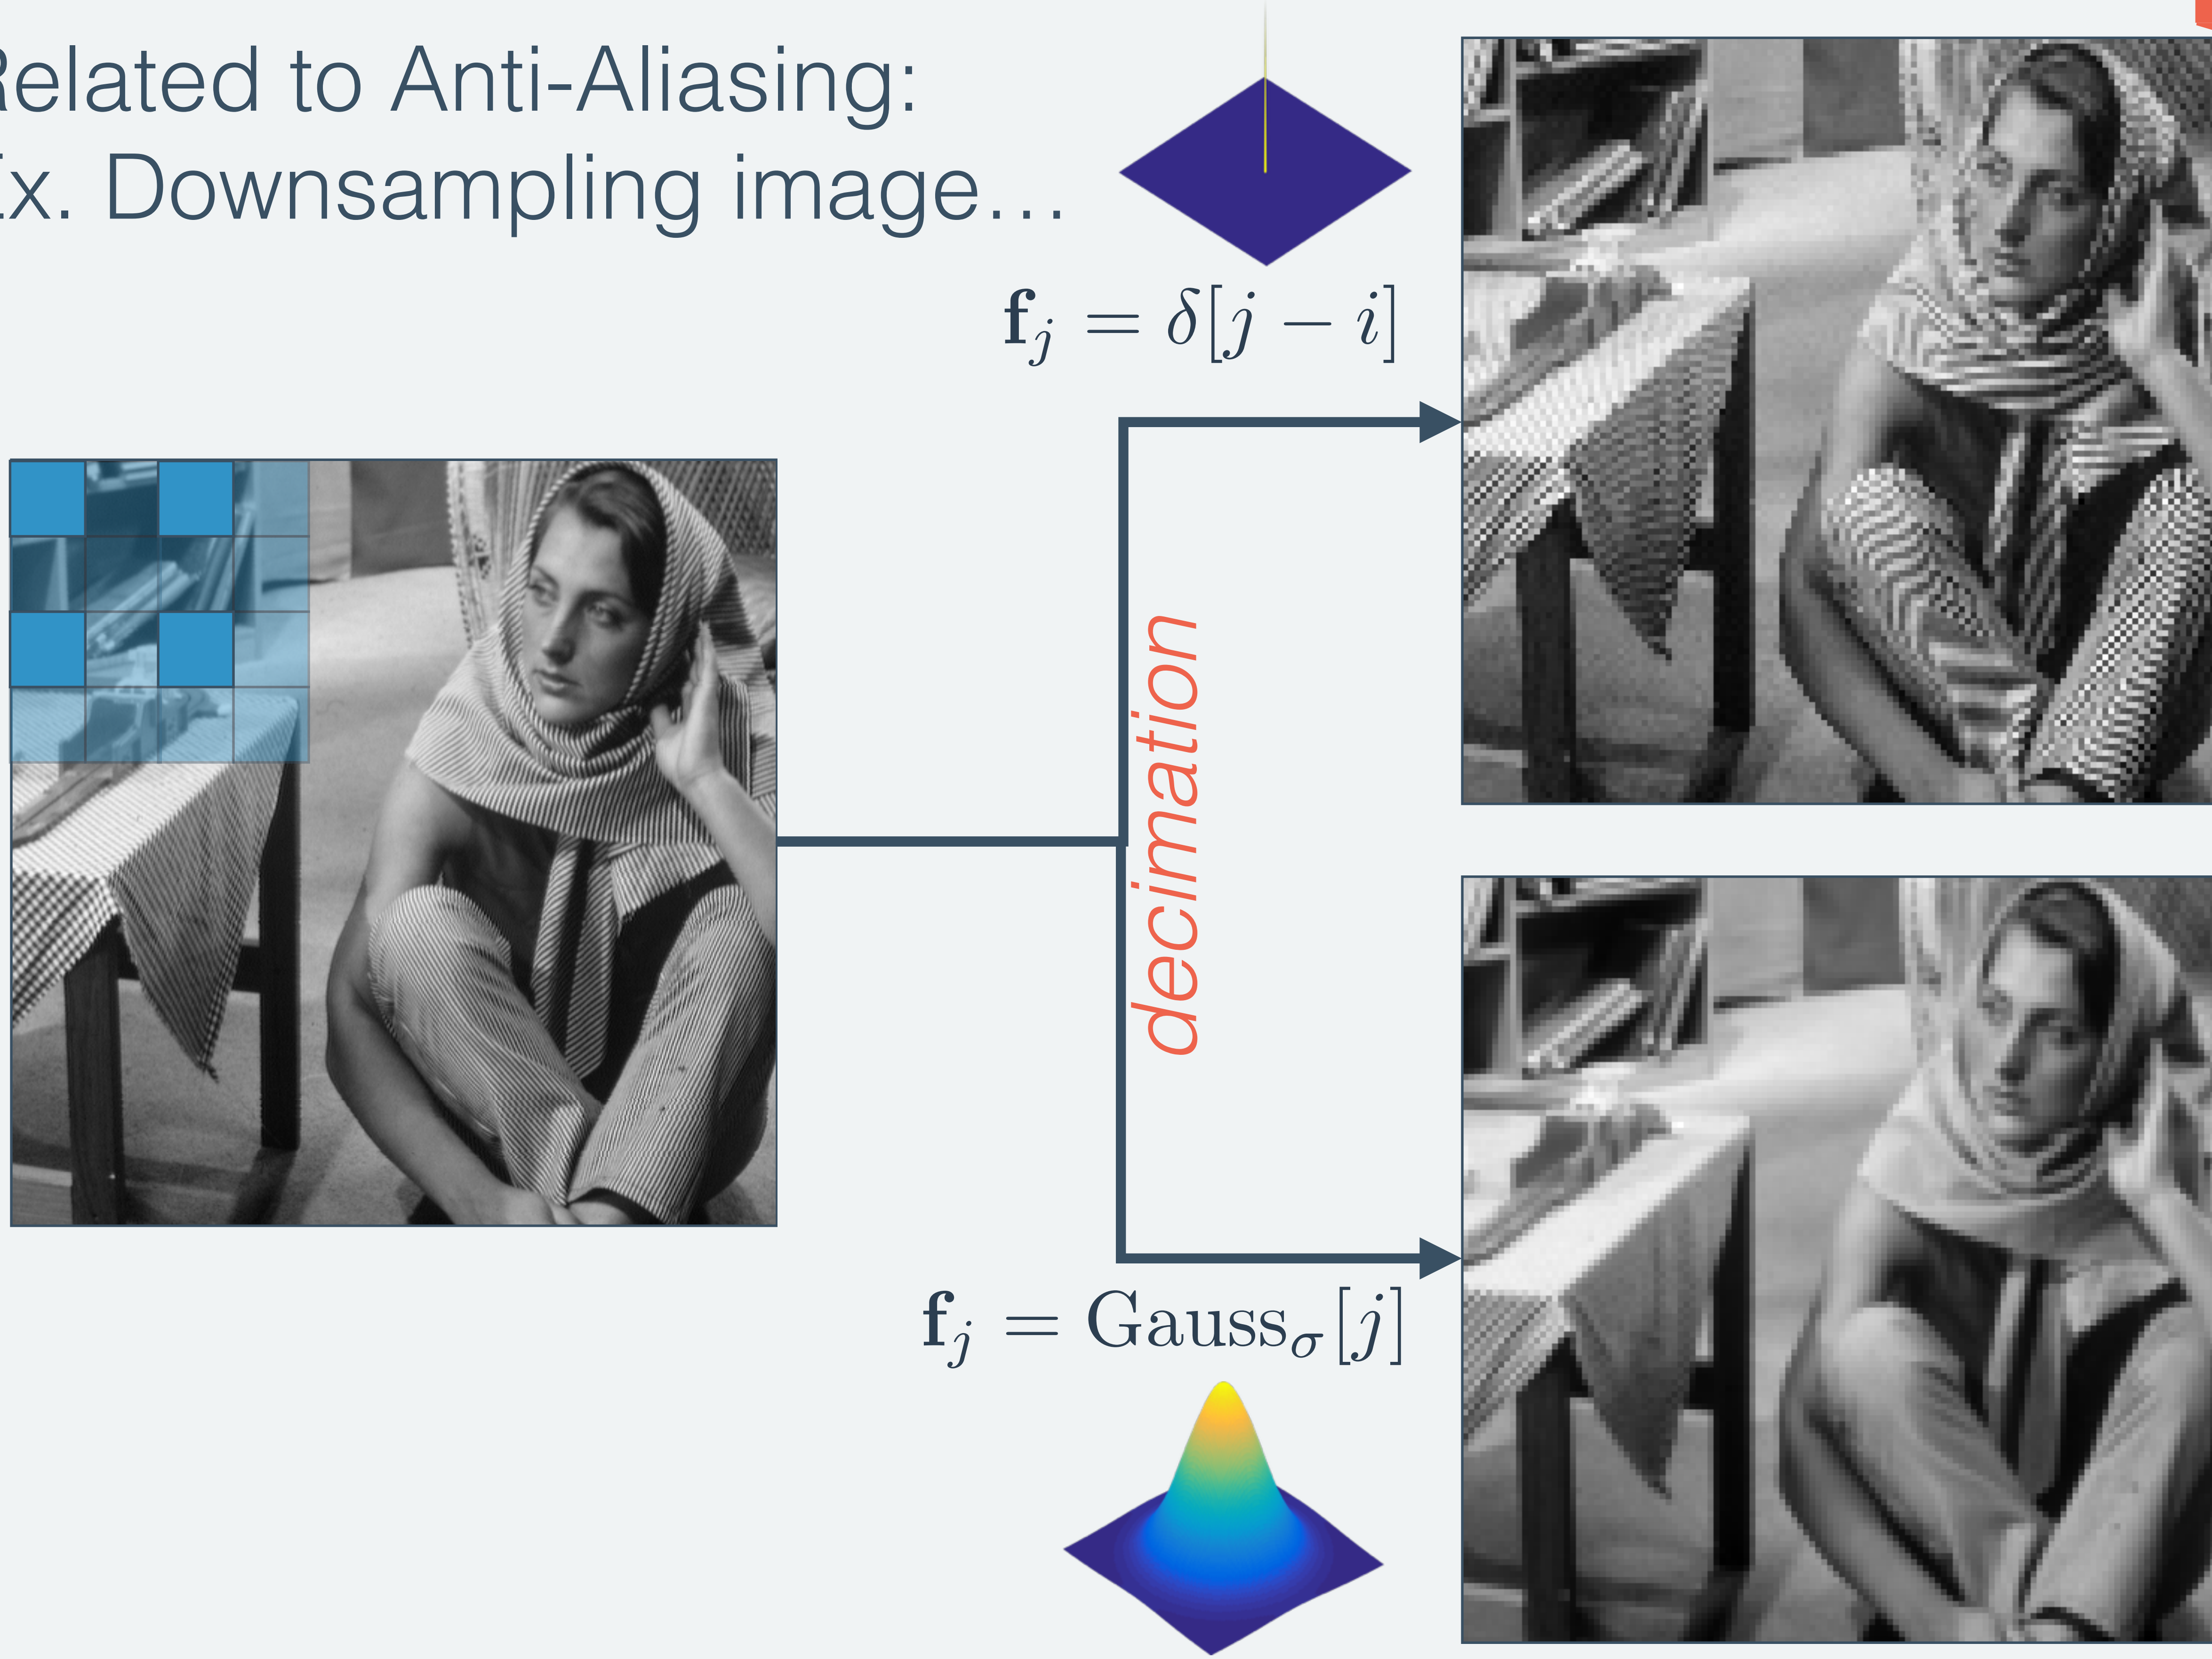

Designing Sampling

Related to Anti-Aliasing:
Ex. Downsampling image…

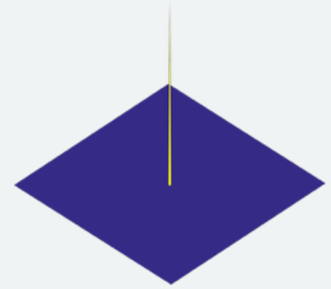

$\mathbf{f}_j = \delta[j - i]$

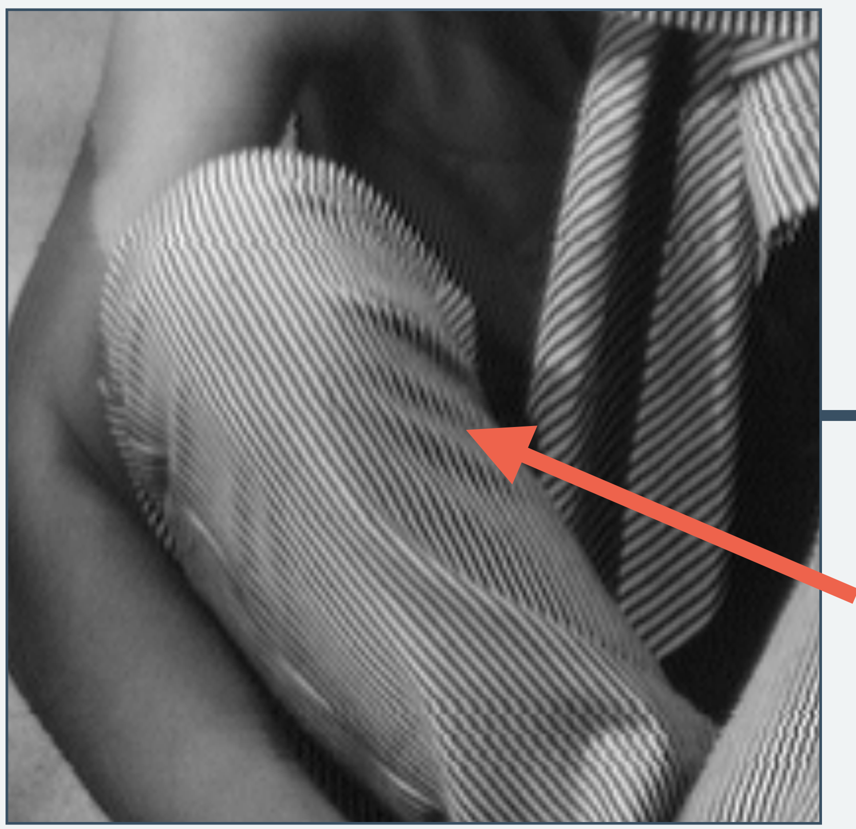

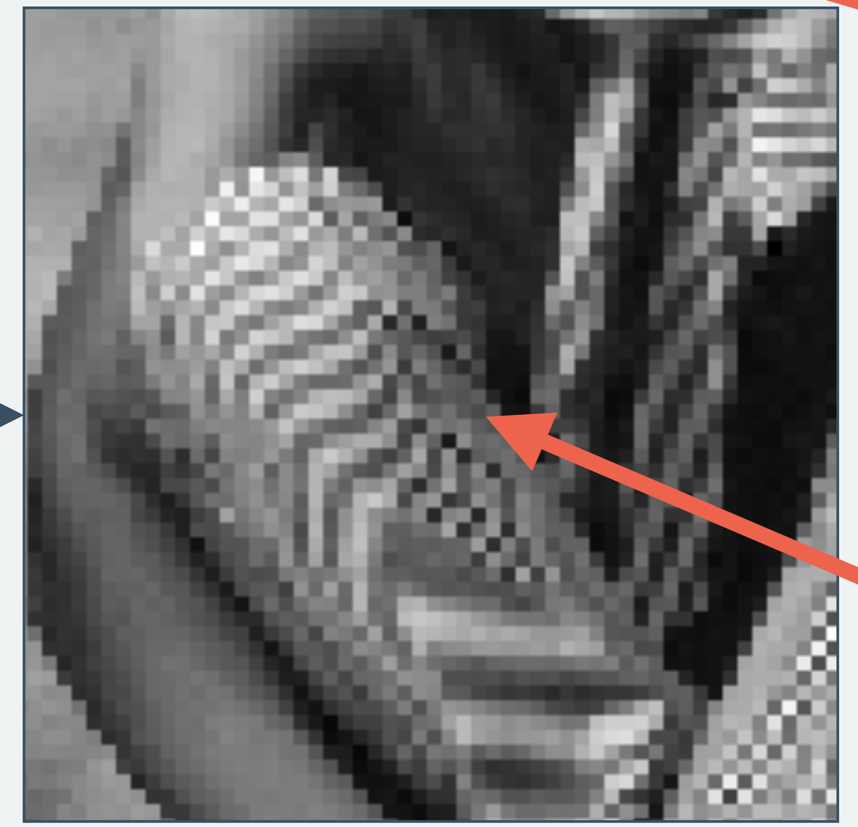

decimation

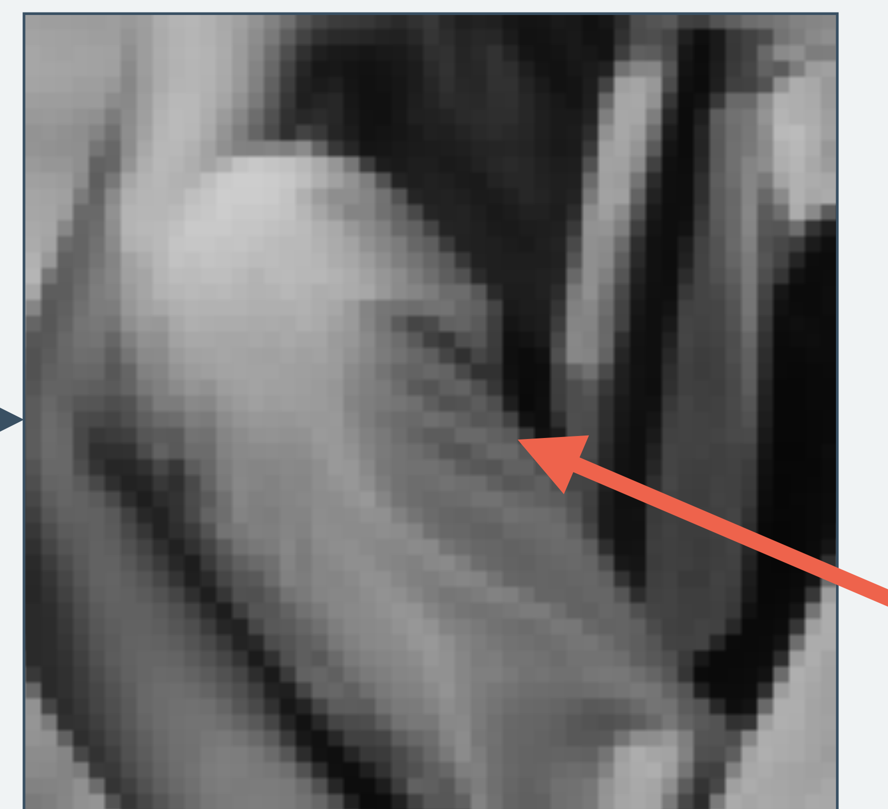

$\mathbf{f}_j = \text{Gauss}_\sigma[j]$

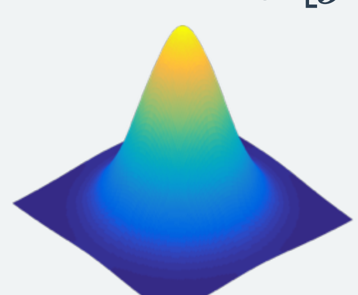

Designing Sampling

Undersampling (**M<N**): Accounting for sparsity

$$y_j = \langle \mathbf{f}_j, \mathbf{x} \rangle$$
$$\forall \;\; j \in [1, M]$$

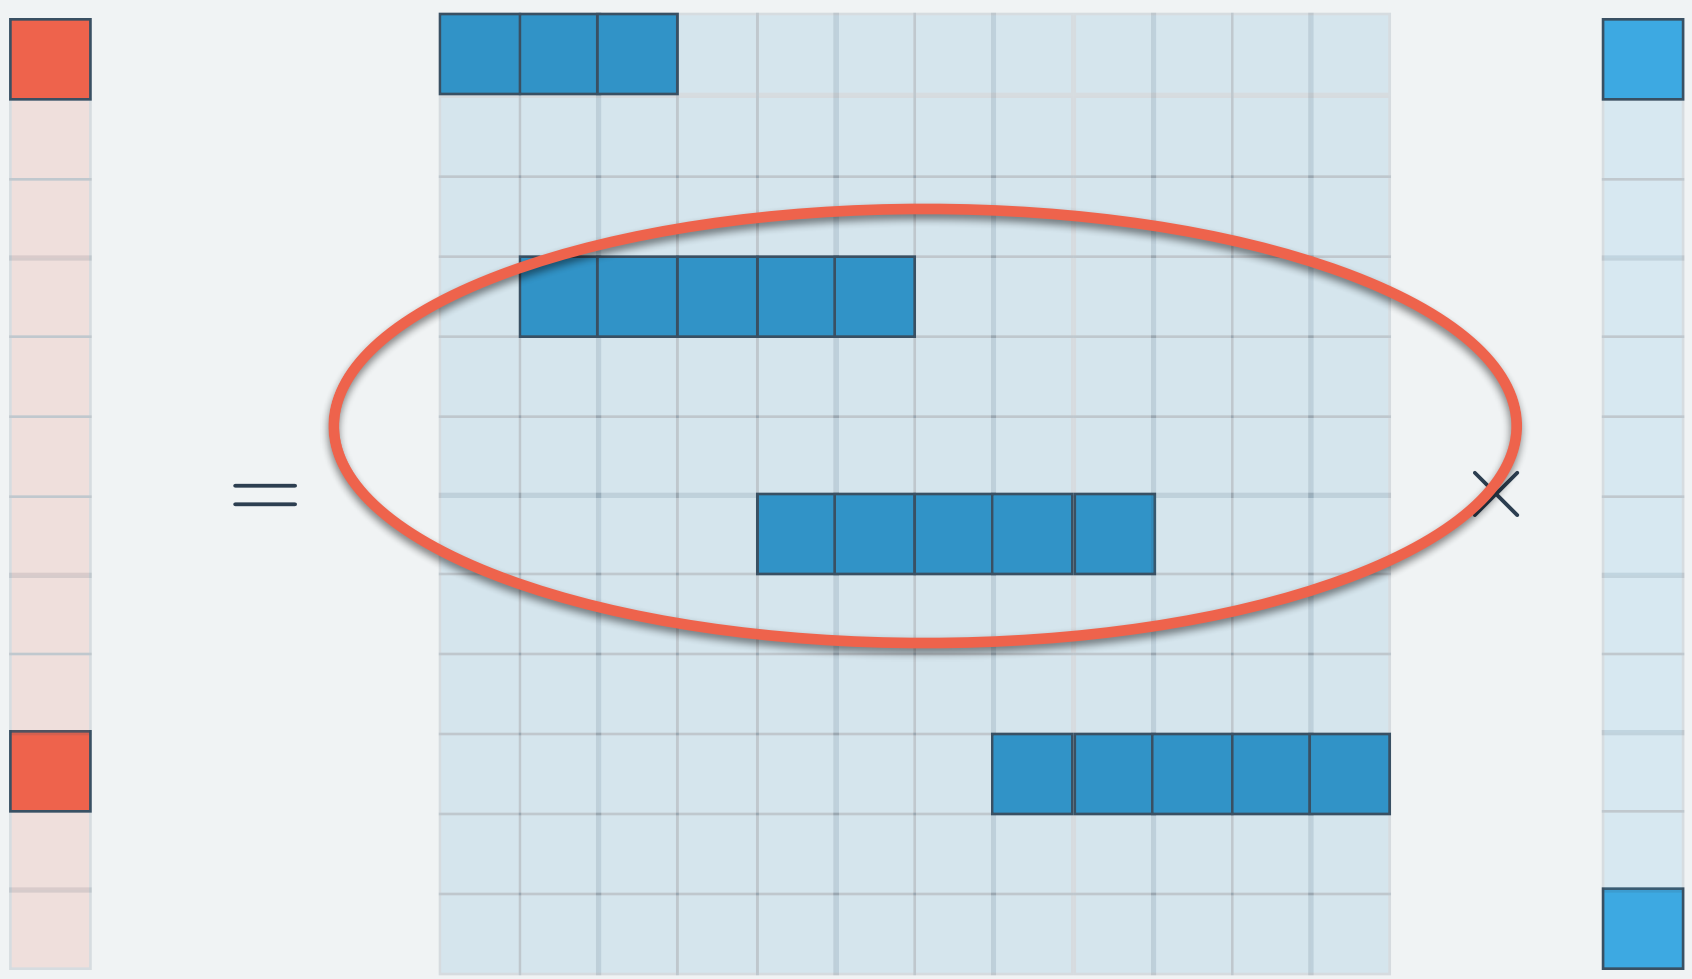

However, for sparse **x**, "localized" filters can miss sparse elements

Designing Sampling

Undersampling (**M<N**): Accounting for sparsity

$y_j = \langle \mathbf{f}_j$

$\forall \;\; j \in [1,$

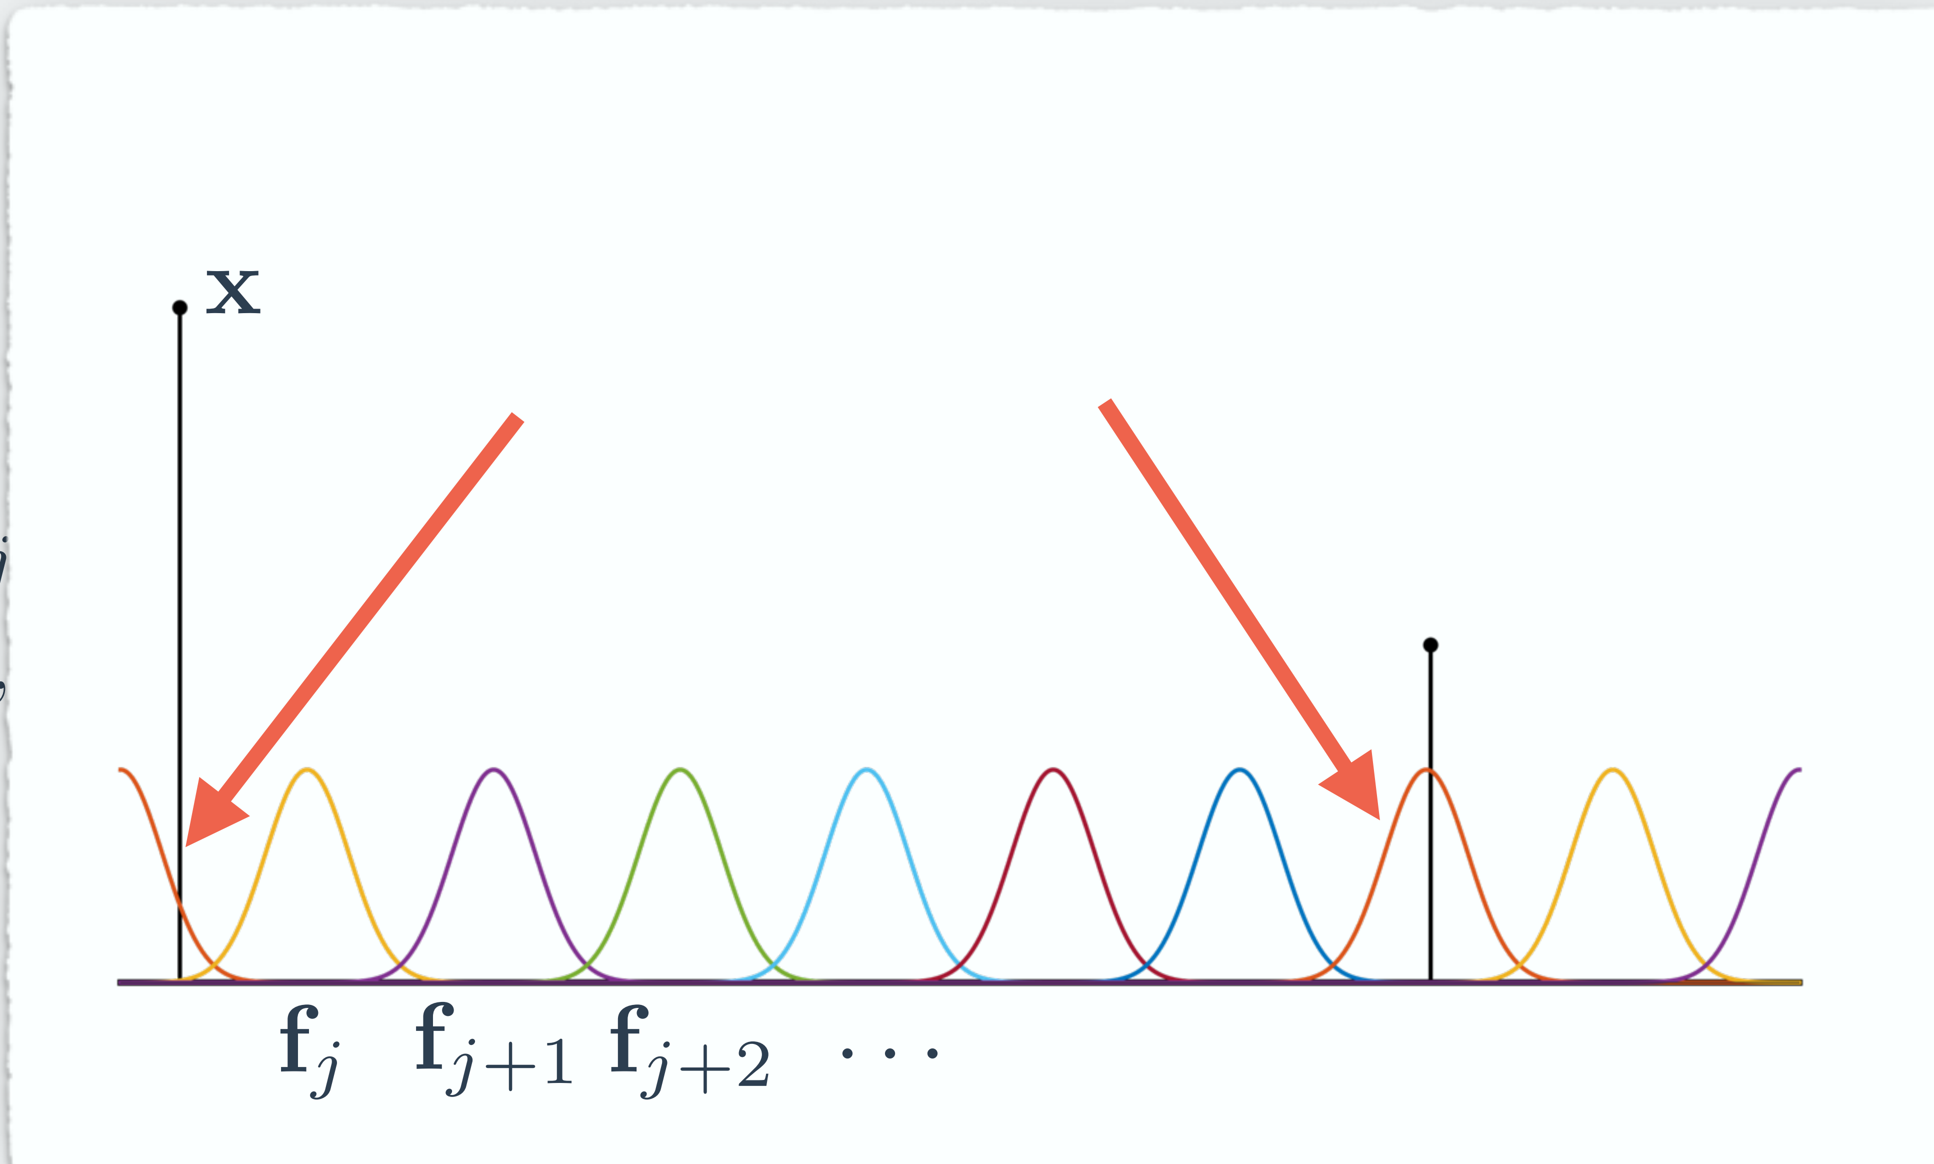

However, for sparse **x**, "localized" filters can miss sparse elements

- Redundancy in measurement from correlation

Designing Sampling

Undersampling (**M<N**): Accounting for sparsity

De-localized (*global*) filters

$$y_j = \langle \mathbf{f}_j, \mathbf{x} \rangle$$
$$\forall \;\; j \in [1, M]$$

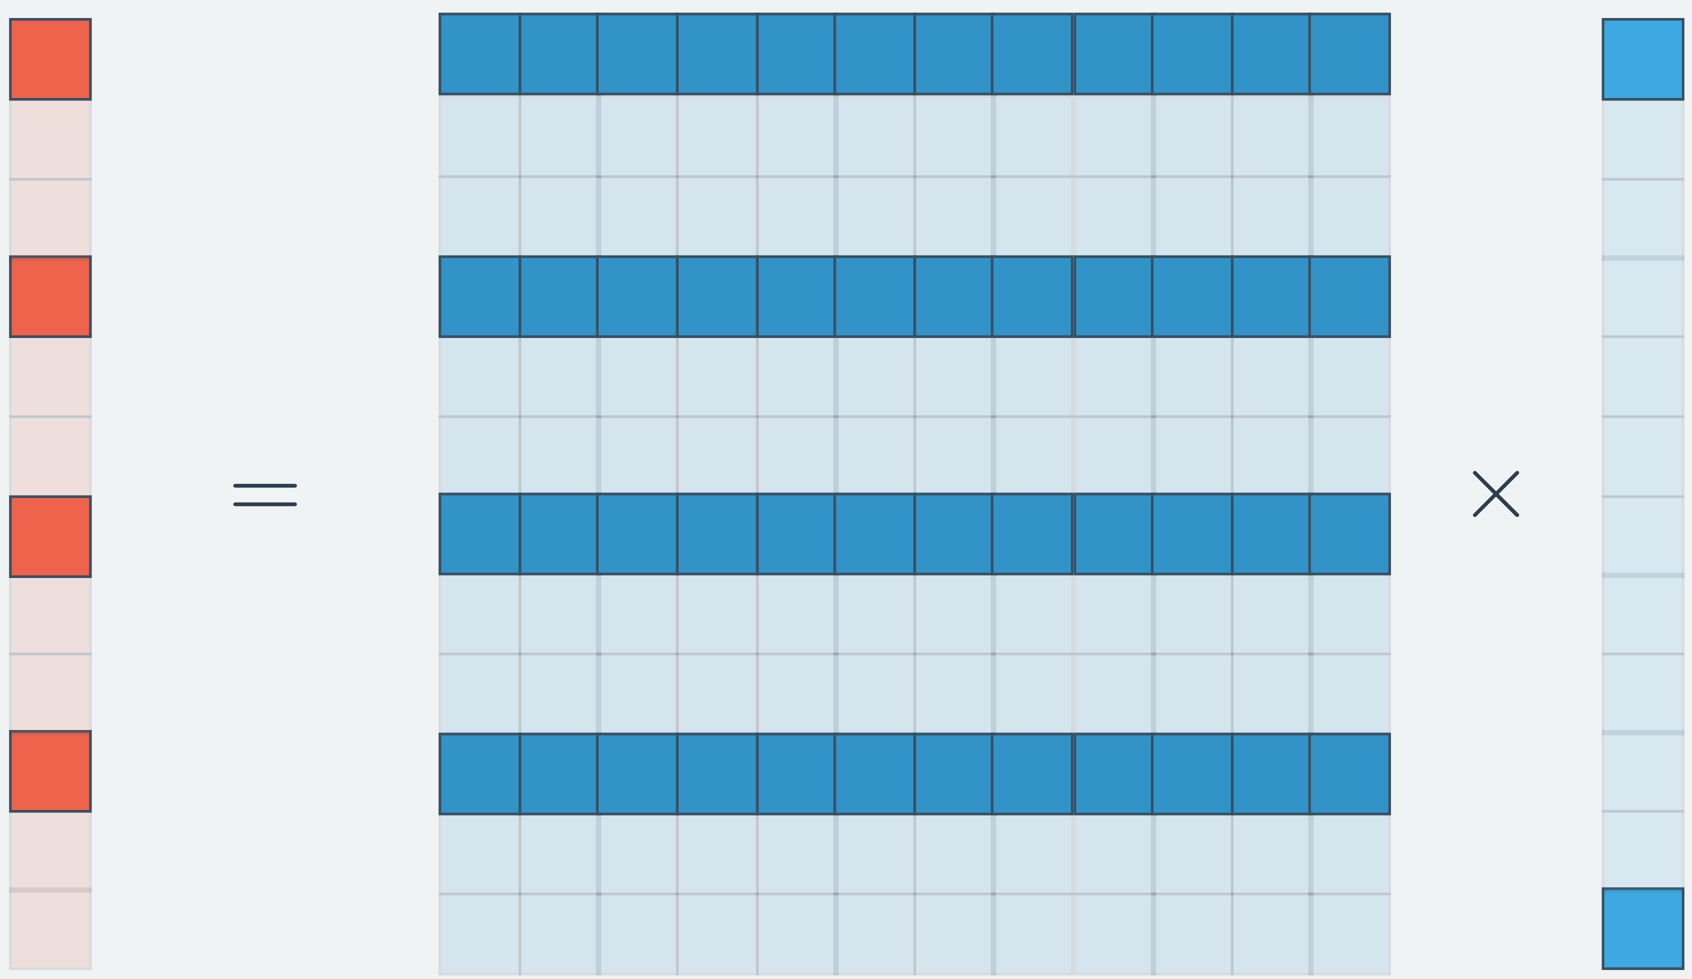

Want:

* Every observation to be informative
* Every observation to tell us something different
* A construction that helps us find the support

2005: Explosion of Compressed Sensing

Decoding by Linear Programming

Practical Signal Recovery from Random Projections

Sparse nonnegative solution of underdetermined linear equations by linear programming

David L. Donoho and Jared Tanner

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

Core contributors…

E. Candès

J. Romberg

T. Tao

J. Tanner

D. Donoho

…and many, many more in subsequent years.

Compressed Sensing Theory

Tool: Restricted Isometry

[Candès & Tao, 2005] A matrix F satisfies the restricted isometry property (RIP) of order K if there exists some small, bounded constant δ_K such that

$$(1-\delta_K)||\mathbf{x}||_2^2 \leq ||F\mathbf{x}||_2^2 \leq (1+\delta_K)||\mathbf{x}||_2^2$$

holds for all K-sparse $\mathbf{x}$,

$$\mathbf{x} \in \{\mathbf{x} : ||\mathbf{x}||_0 \leq K\}\,.$$

Essentially:

* If **F** obeys RIP-**K**, then it is approximately orthonormal for all K-sparse vectors.
* If **F** obeys RIP-**2K**, then it approximate preserves distance relationships of K-sparse vectors.

An Aside for Lp Norms

Supposing some vector $\mathbf{x}$ of dimensionality N, we define the ℓ_p norm as,

$$||\mathbf{x}||_p \triangleq \left(\sum_{i=1}^{N} |x_i|^p \right)^{\frac{1}{p}} .$$

Hence,

- $||\mathbf{x}||_2 = \sqrt{x_1^2 + x_2^2 + \cdots + x_N^2}$
- $||\mathbf{x}||_1 = |x_1| + |x_2| + \cdots + |x_N|$
- $||\mathbf{x}||_0 = \text{Count}(x_i \neq 0; \forall i \in [1, N])$ (*semi-norm*)
- $||\mathbf{x}||_\infty = \max_{i \in [1,N]} \quad |x_i|$

ASIDE: Lp Norms

To ℓ_0

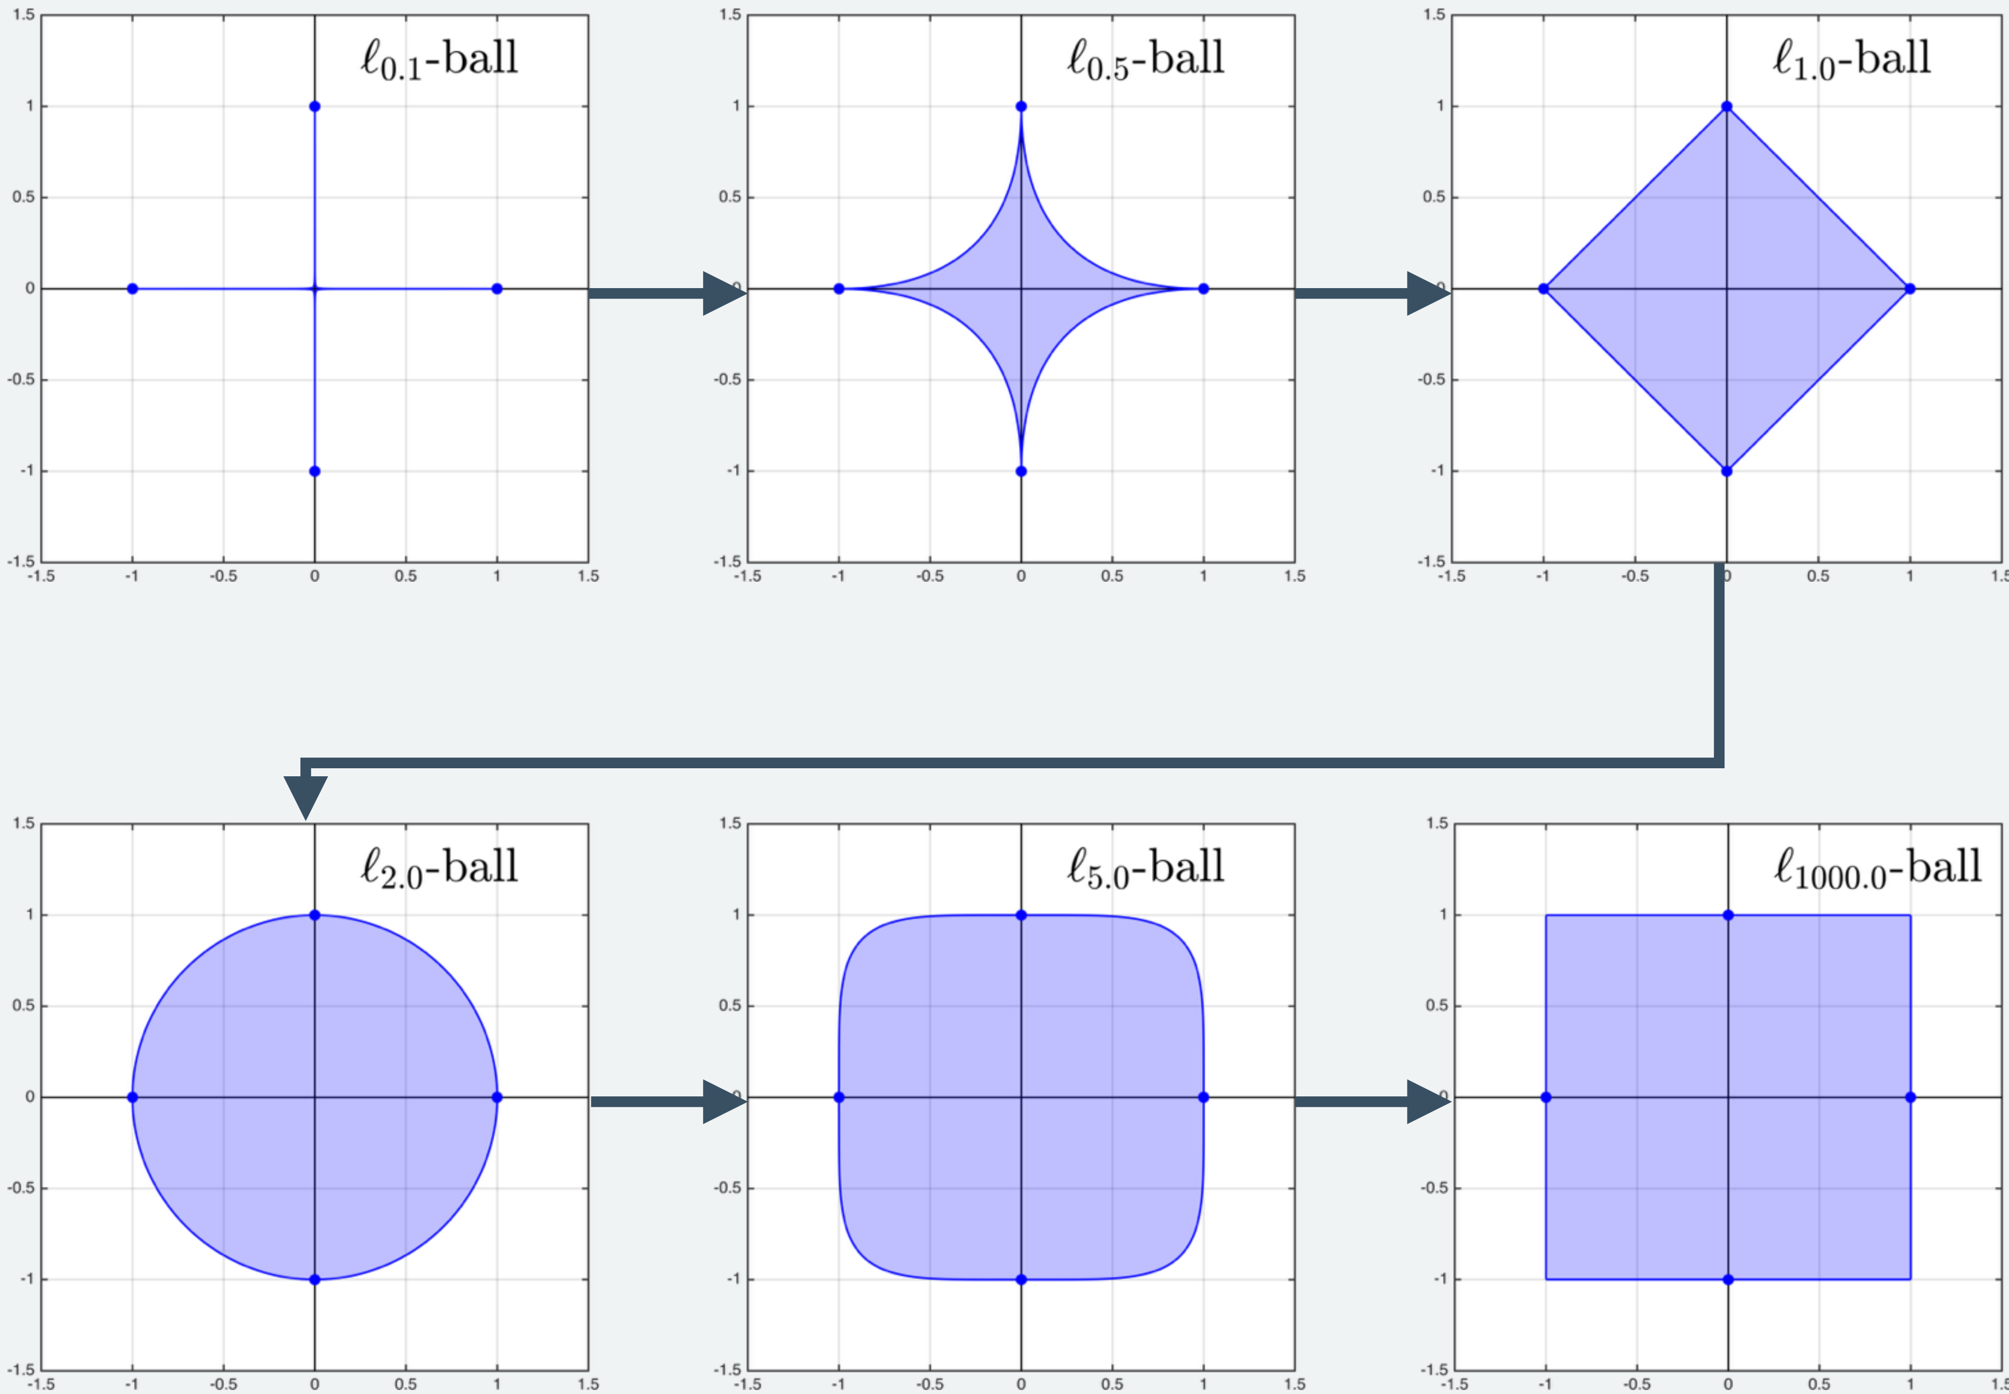

To ℓ_∞

Compressed Sensing Theory

Result: Existence of Unique Solution

[Candès & Tao, 2005] Suppose F satisfies the RIP for $\delta_{2K} < 1$ for some $K \geq 1$. For some support set T with $|T| \leq K$, let

$$\mathbf{y} \triangleq F_T \mathbf{c}$$

for some arbitrary $|T|$ dimensional vector $\mathbf{c}$.

- The set T and the coefficients $(c_j)_{j \in t}$ can be reconstructed *uniquely* from knowledge of $\mathbf{y}$ and F.

Essentially:

* If we have a RIP-**2K** satisfying **F**, the sparsest solution in the feasible set is the true one.
* Only implies existence, search algorithm over **T** is **NP-Hard**.

Compressed Sensing Theory

Result: Efficient Algorithm Exists

[Candès & Tao, 2005] Suppose F satisfies the stronger RIP,

$$\delta_K + \delta_{2K} + \delta_{3K} < \frac{1}{4},$$

and $\mathbf{c}$ is a real vector with support T obeying $|T| \leq K$. Let $\mathbf{y} = F\mathbf{c}$. Then, $\mathbf{c}$ is the unique minimizer of

$$\min_{\mathbf{d}} \quad ||\mathbf{d}||_1 \quad s.t. \quad F\mathbf{d} = \mathbf{y}.$$

Essentially:

* Given a stricter RIP-**3K** on **F**, the true solution is unique and can be found efficient via *a convex optimization!*

Compressed Sensing Theory

Result: Ef

[Candès &

and **c** is a re

the unique r

c. Then, **c** is

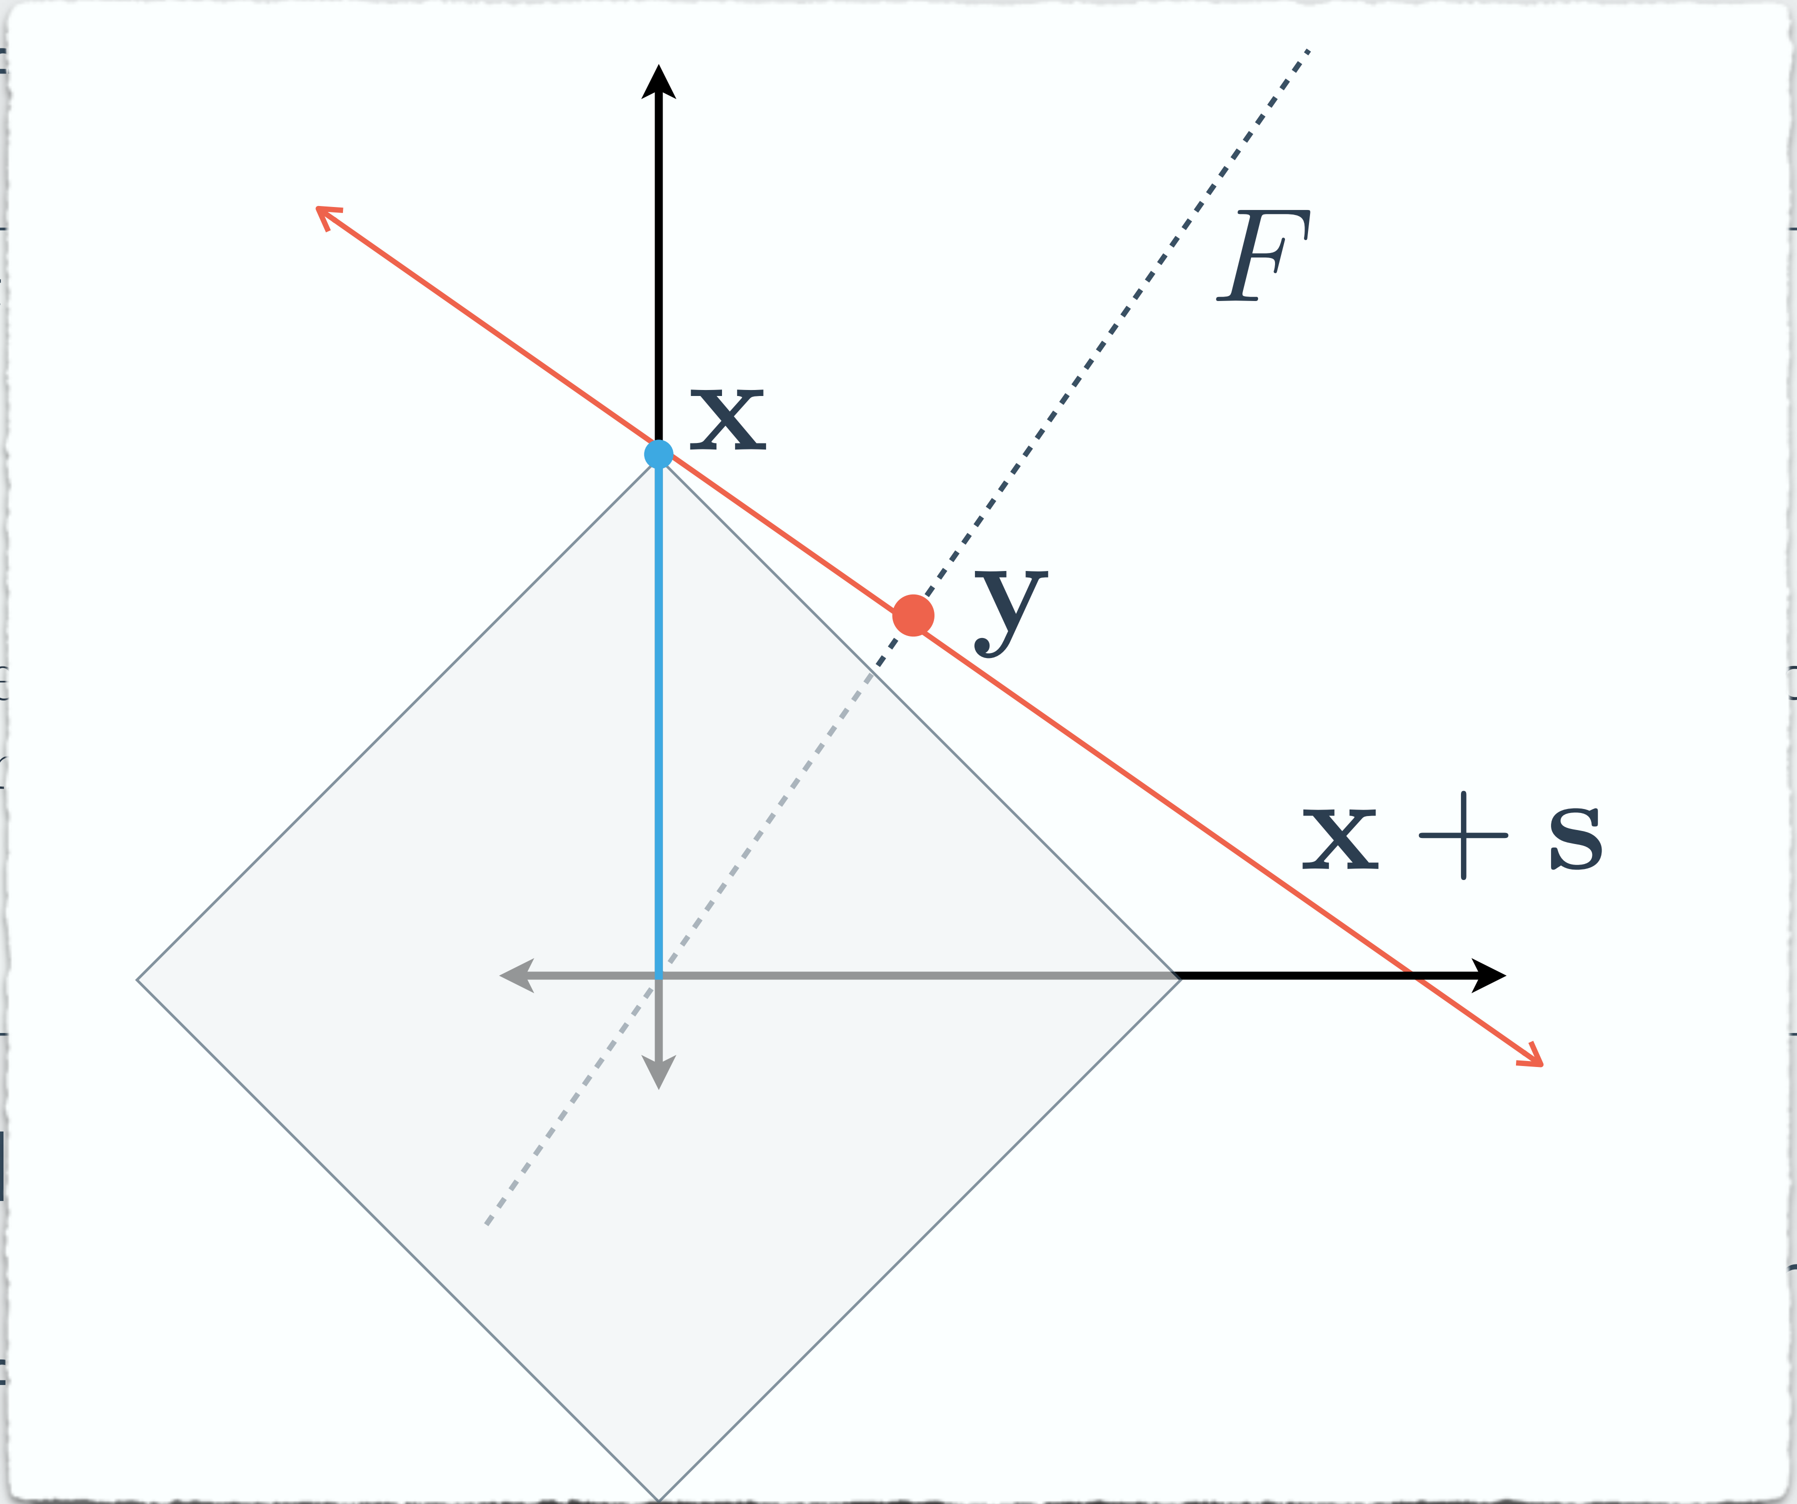

Essentiall

* Given a

can be f

nique and

Compressed Sensing Theory

Result: Efficient Algorithm Exists

[**Candès & Tao, 2005**] Suppose F satisfies the stronger RIP,

$$\delta_K + \delta_{2K} + \delta_{3K} < \frac{1}{4},$$

and $\mathbf{c}$ is a real vector with support T obeying $|T| \leq K$. Let $\mathbf{y} = F\mathbf{c}$. Then, $\mathbf{c}$ is the unique minimizer of

$$\min_{\mathbf{d}} \quad ||d||_1 \quad s.t. \quad F\mathbf{d} = \mathbf{y}.$$

Essentially:

* Given a stricter RIP-**3K** on **F**, the true solution is unique and can be found efficient via *a convex optimization*.

However, RIP verification of a matrix is NP-Hard, so deterministic design is intractable!

Compressed Sensing Theory

Result: Approximately Sparse Signals

[Candès, Romberg, & Tao, 2006] If F obeys a RIP for $\delta_{2K} < \sqrt{2} - 1$, then the ℓ_1 recovered solution

$$\mathbf{x}^* = \arg\min_{\mathbf{a}} \quad ||\mathbf{a}||_1 \quad s.t. \quad F\mathbf{a} = \mathbf{y}$$

has an error bounded by,

$$||\mathbf{x}^* - \mathbf{x}||_2 \leq ||\mathbf{x} - \mathbf{x}_K||_1,$$

where $\mathbf{x}_K$ is equal to the true solution $\mathbf{x}$ for the K largest components and 0 everywhere else.

Effectively: We can recover compressible signals (ones with *power-law decay*) up to their nearest K-sparse approximation.

Compressed Sensing Theory

Result: Ap

[**Candès, R** $\sqrt{2}-1$, then
the ℓ_1 recove

has an error

where $\mathbf{x}_K$ is nents and 0
everywhere

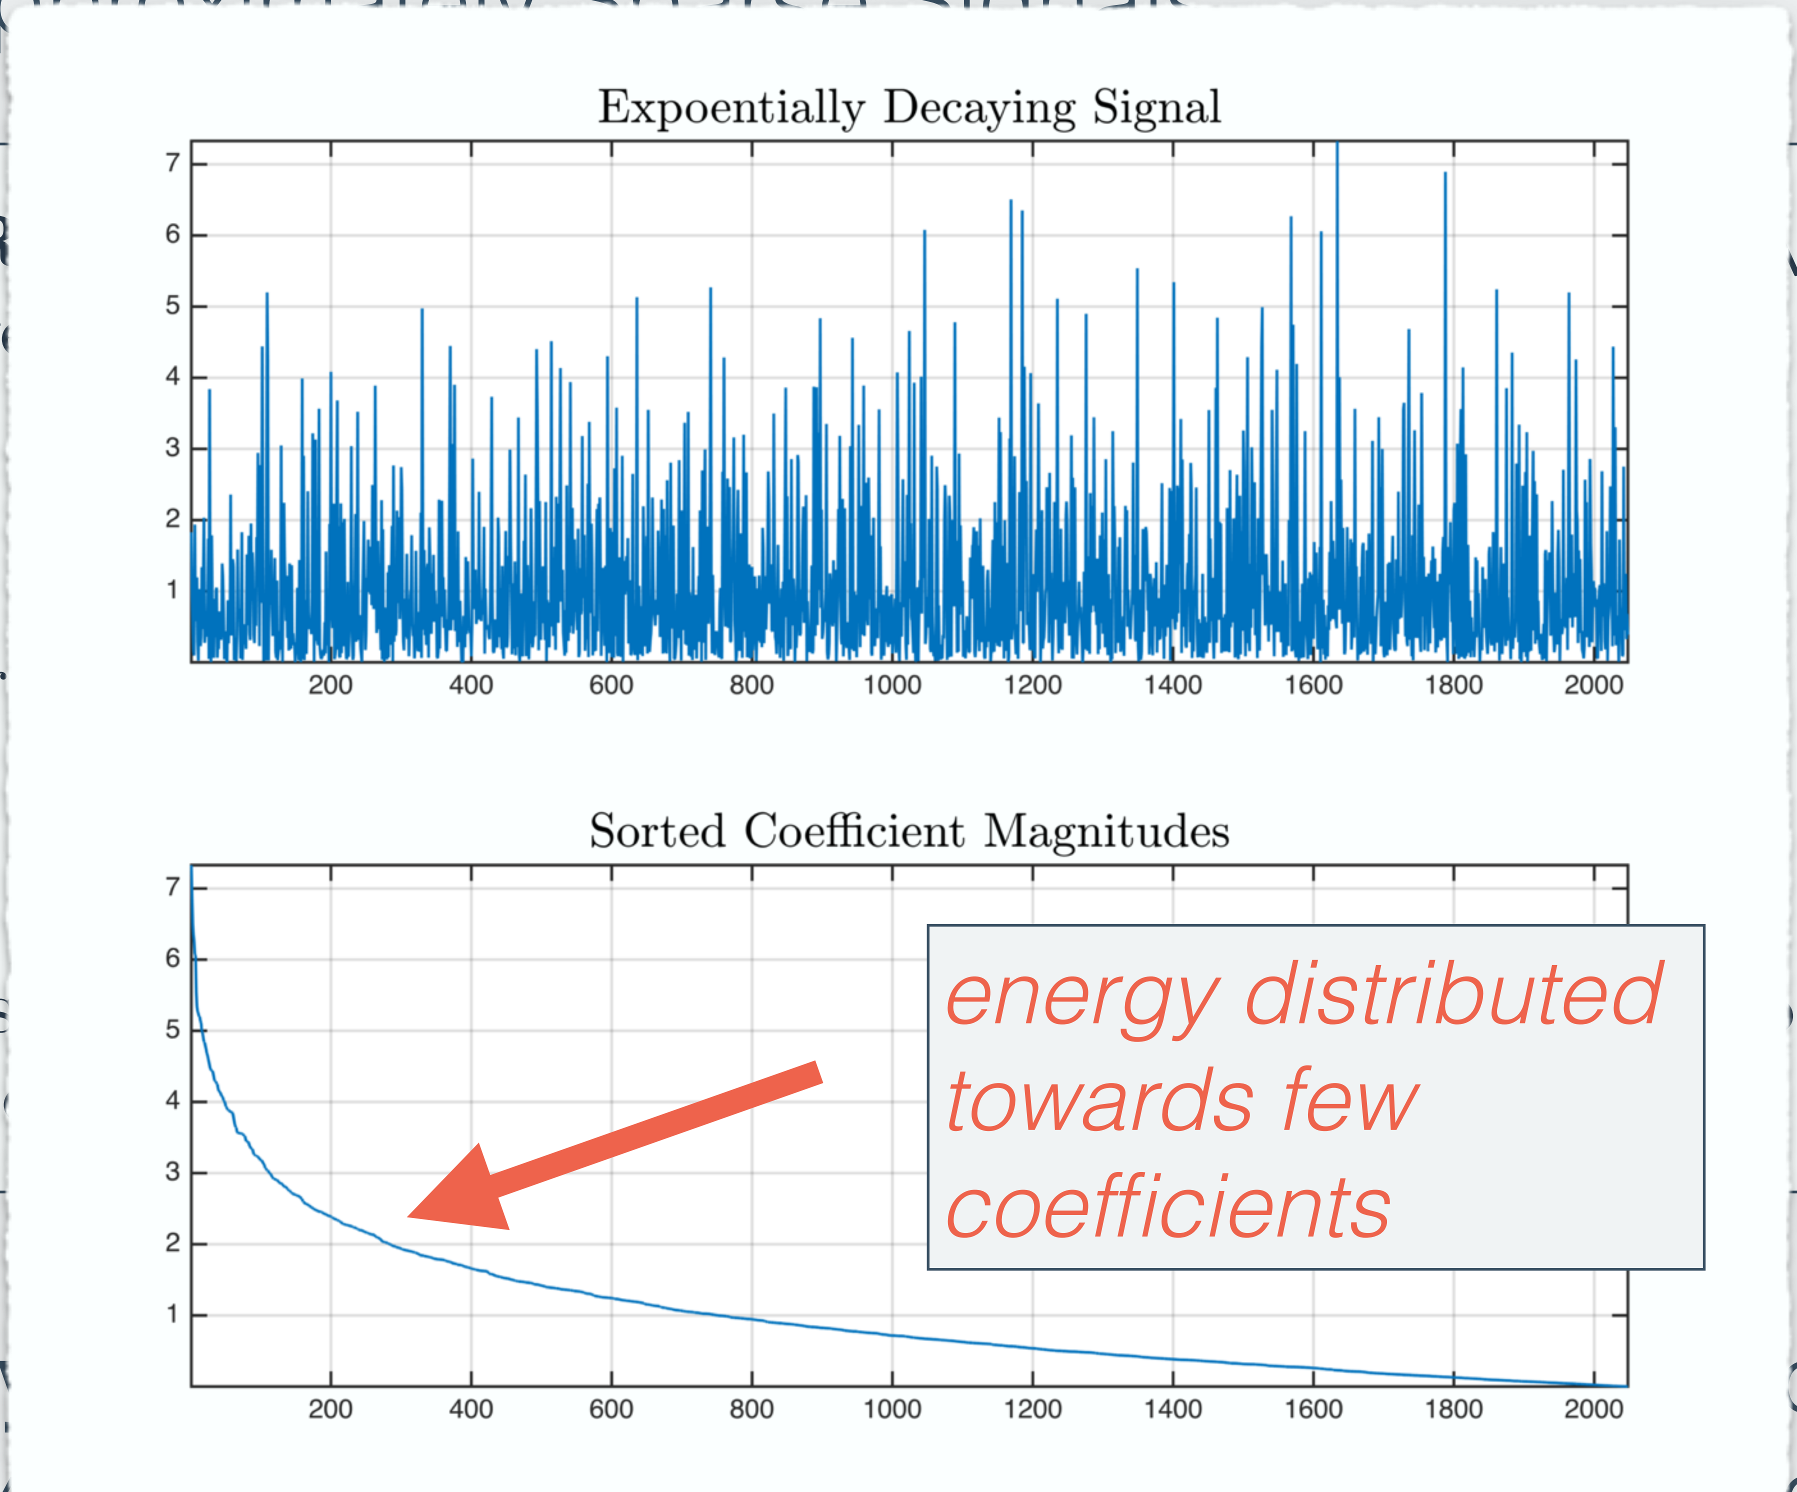

Effectivel ones with
power-law oximation.

Compressed Sensing Theory

Final Piece of the Puzzle: Randomness

[**Candès & Tao, 2005**] Assume $M \leq N$ and let F be an $M \times N$ matrix whose entries are **i.i.d. Gaussian** with zero mean and variance $\frac{1}{M}$. Then, unique ℓ_1 recoverability holds with overwhelming probability for sufficiently small ratio K/N.

Compressed Sensing Theory

Final Piece of the Puzzle: Randomness

[Candès & Wakin, 2008] If F is constructed by

- Randomly sampling columns as unit vectors from $\mathbb{R}^M$,
- Randomly sampling i.i.d. entries from $\mathcal{N}(0, \frac{1}{M})$,
- Randomly sampling from and some orthonormal basis and normalizing,
- Randomly sampling i.i.d. $\pm\frac{1}{\sqrt{M}}$ Bernoulli entries,

then unique ℓ_1 recoverability holds for K-sparse $\mathbf{x}$ for the **nearly-optimal** bound

$$M \geq C \cdot K \log\left(\frac{N}{K}\right).$$

Compressed Sensing Theory

Final Piece of the Puzzle: Randomness

[Candès & Wakin, 2008] If F is constructed by

- Randomly sampling columns as unit vectors from $\mathbb{R}^M$,
- Randomly sampling i.i.d. entries from $\mathcal{N}(0, \frac{1}{M})$,
- Randomly sampling from and some orthonormal basis and normalizing,
- Randomly sampling i.i.d. $\pm\frac{1}{\sqrt{M}}$ Bernoulli entries,

then unique ℓ_1 recoverability holds for K-sparse $\mathbf{x}$ for the **nearly-optimal** bound

$$M \geq C \cdot K \log\left(\frac{N}{K}\right).$$

Vital for practical implementation of CS for real sensing problems.

CS: Restricted Isometry

Monte Carlo: Set N=2048 and test stability of K-sparse subsets of projection matrix. (Here, 20 realizations.)

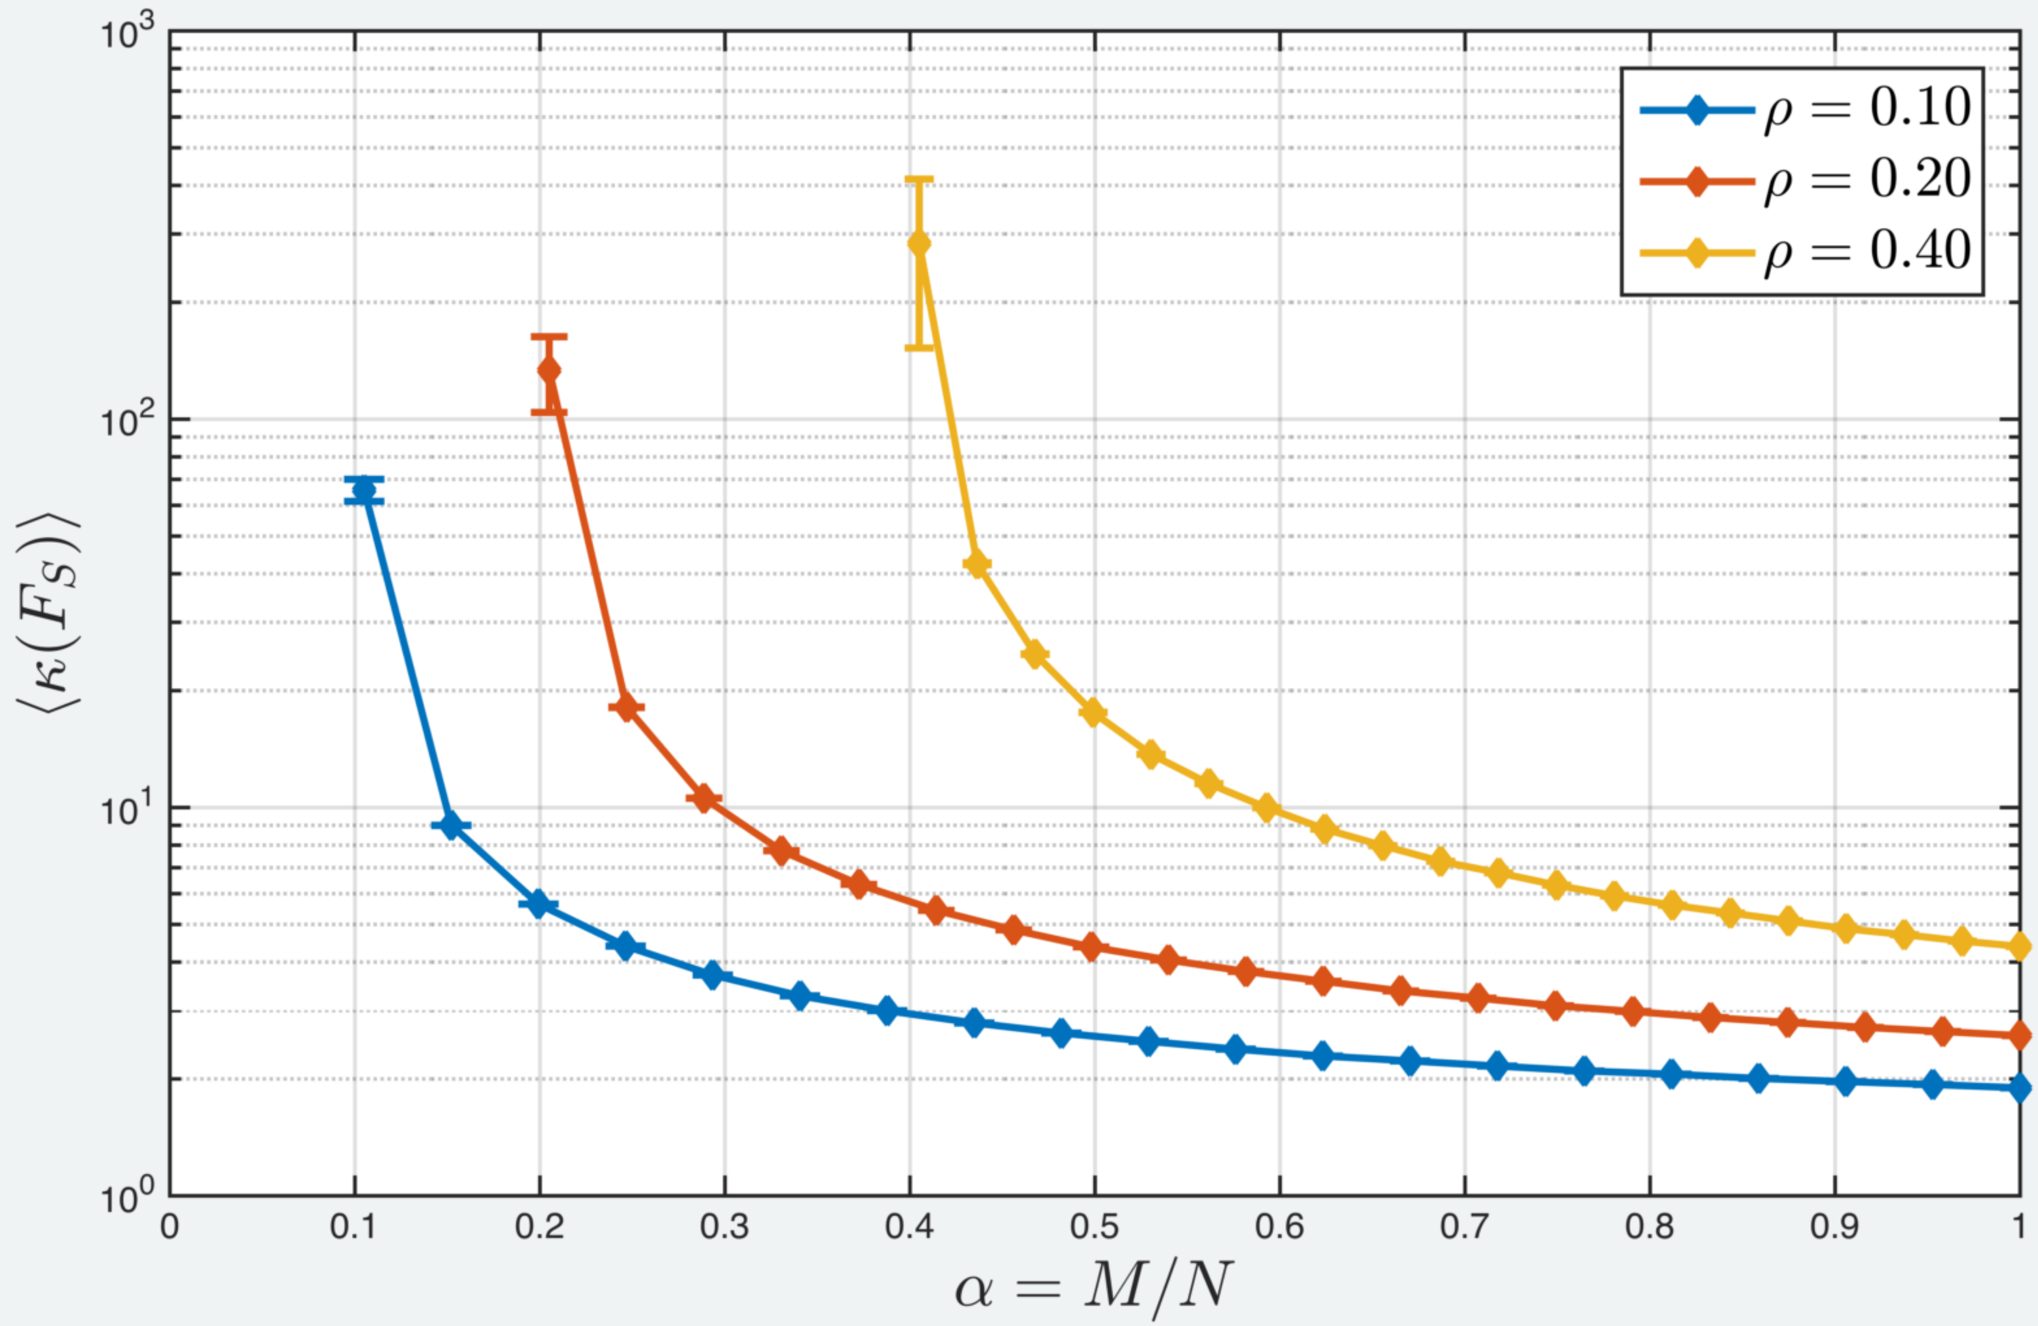

Compressed Sensing Theory

ENS

Result: Sparse Bases & Mutual Incoherence

Assume that a signal $\mathbf{x}$ has a *sparse representation basis*, Ψ, such that

$$\mathbf{x} = \Psi^{-1}\theta$$

where θ is K-sparse. One may then write the measurements as

$$\mathbf{y} = F\mathbf{x} = A\theta,$$

where $A = F\Psi^{-1}$, and solve

$$\begin{aligned} \theta^* &= \arg\min_{\nu} \quad ||\nu||_1 \quad s.t. \quad A\nu = \mathbf{y}, \\ \mathbf{x}^* &= \Psi\theta^* \end{aligned}$$

ASIDE: Sparse Bases in 2D

I. Discrete 2D Fourier Basis

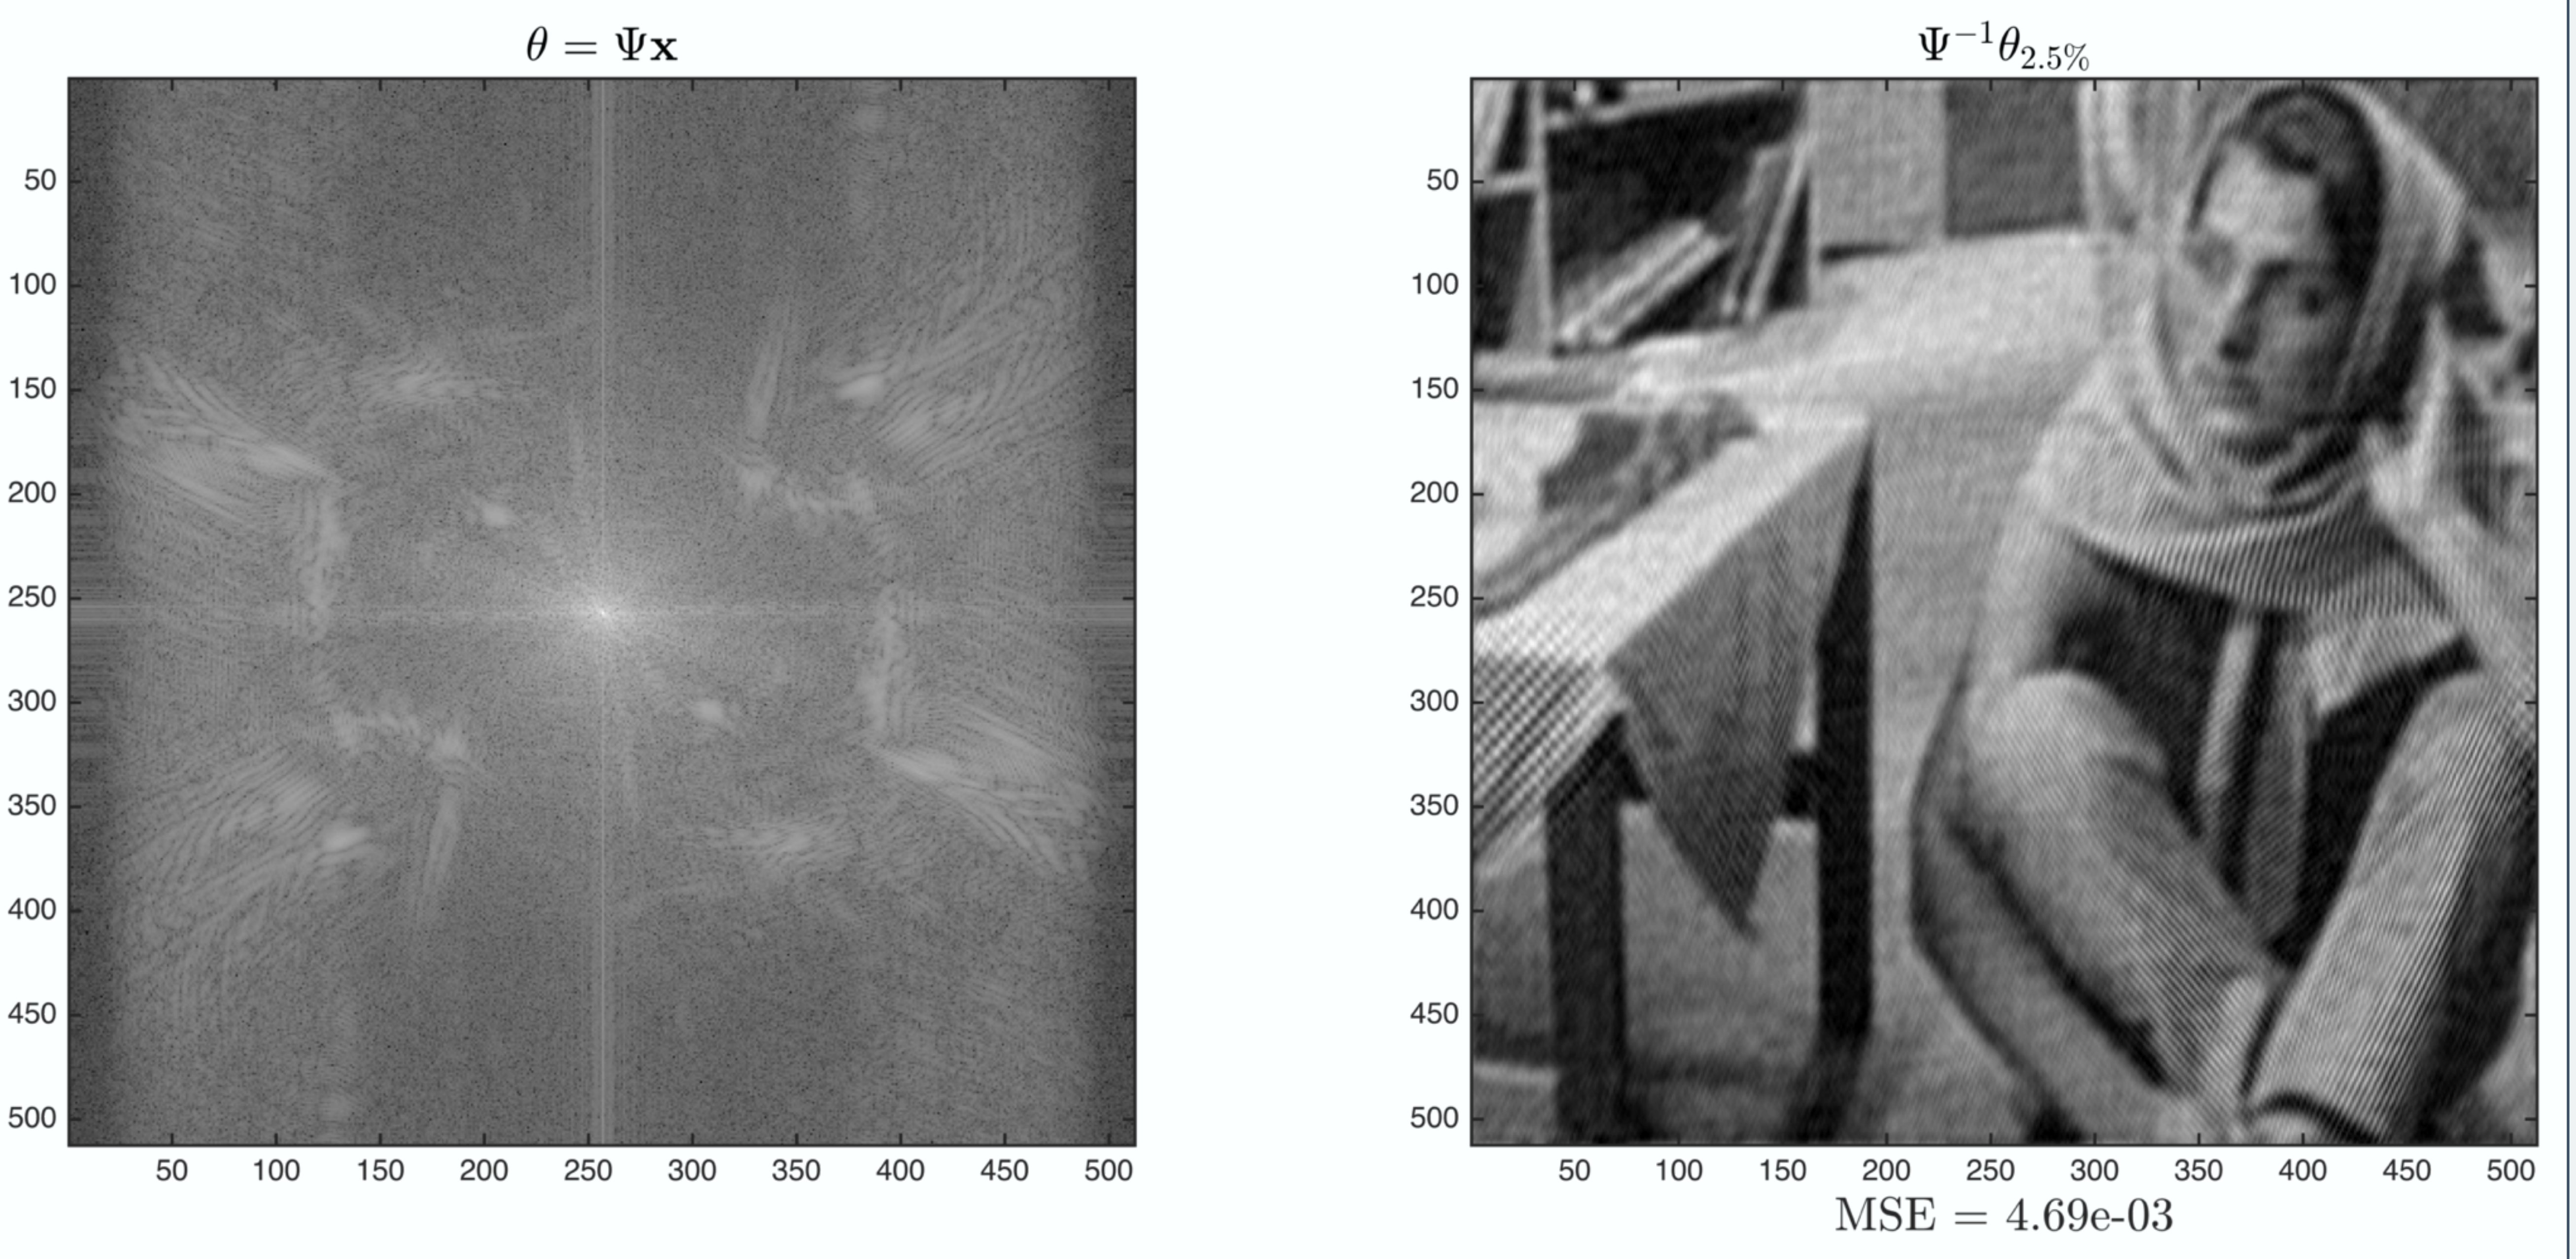

ASIDE: Sparse Bases in 2D

II. 2D Discrete Cosine Transform

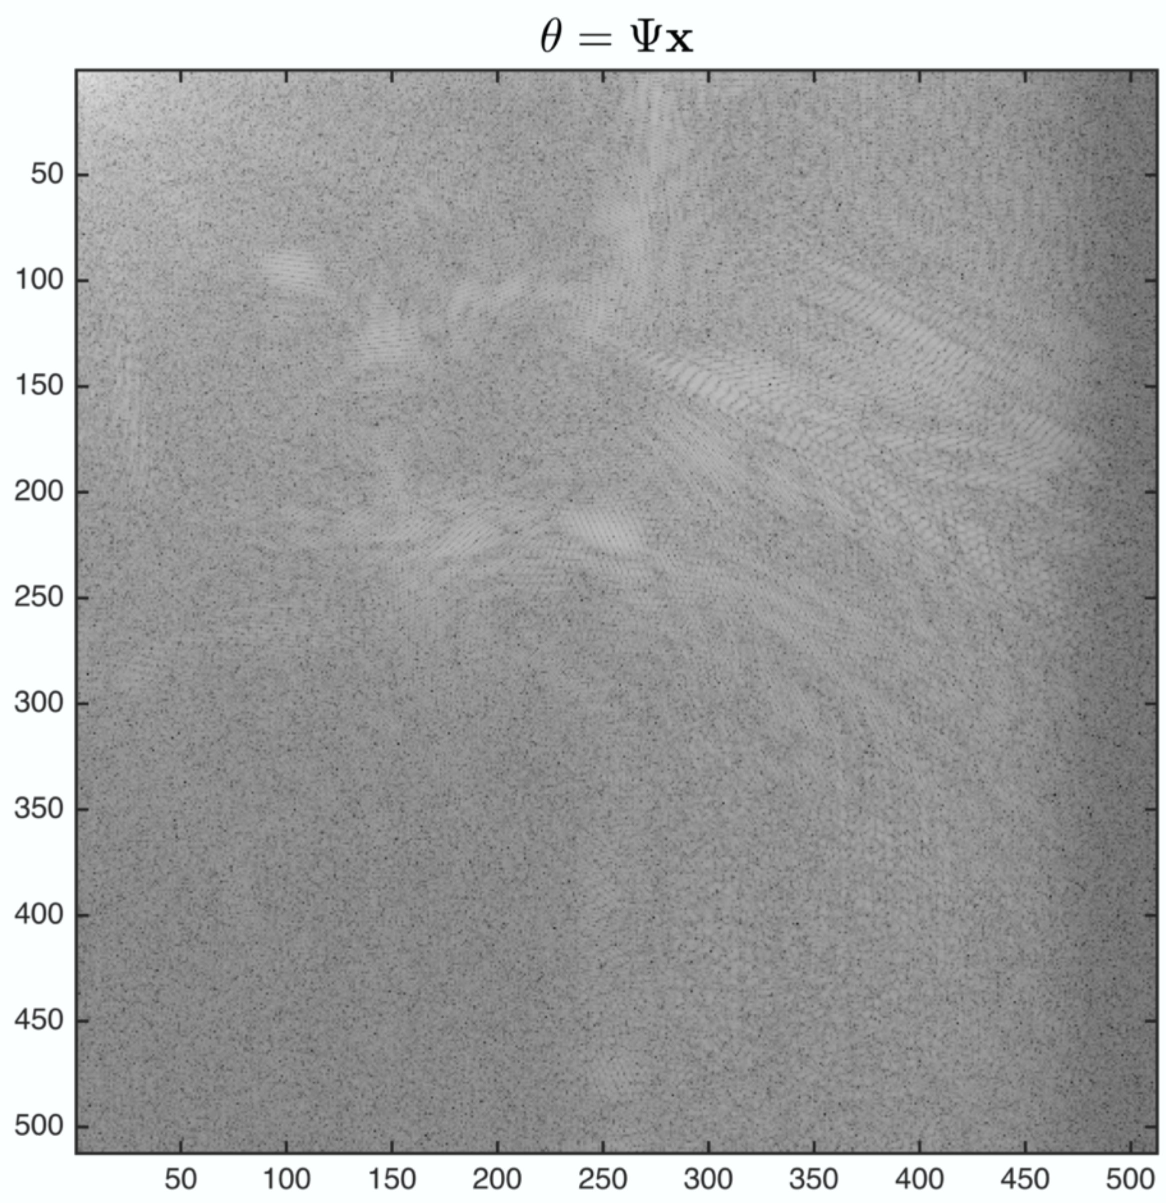

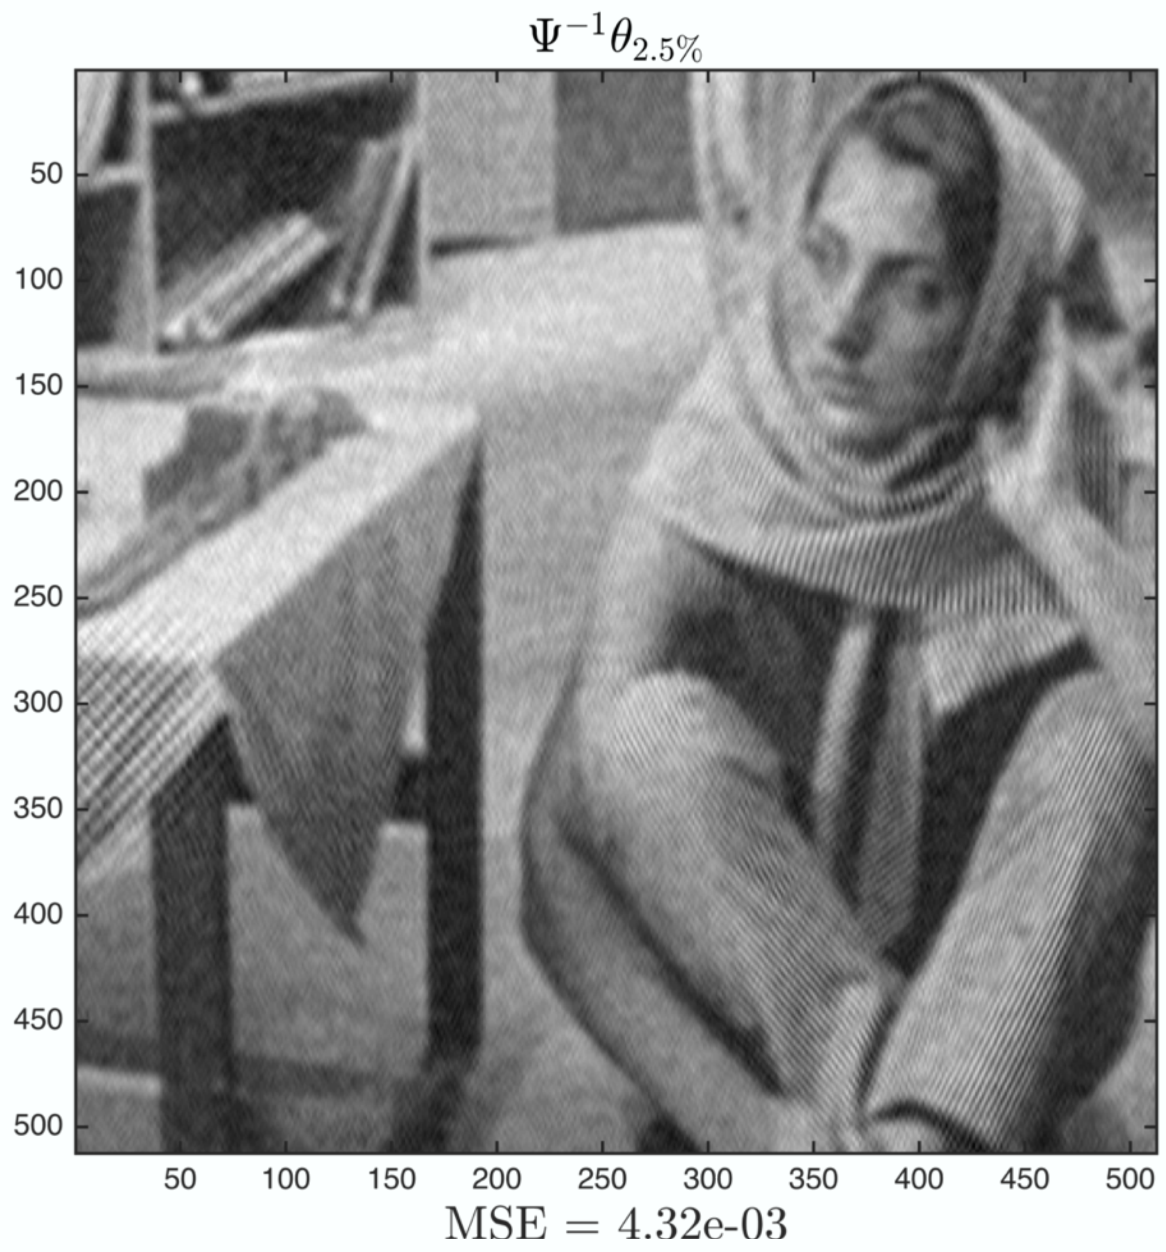

ASIDE: Sparse Bases in 2D

III. 2D Haar Wavelets

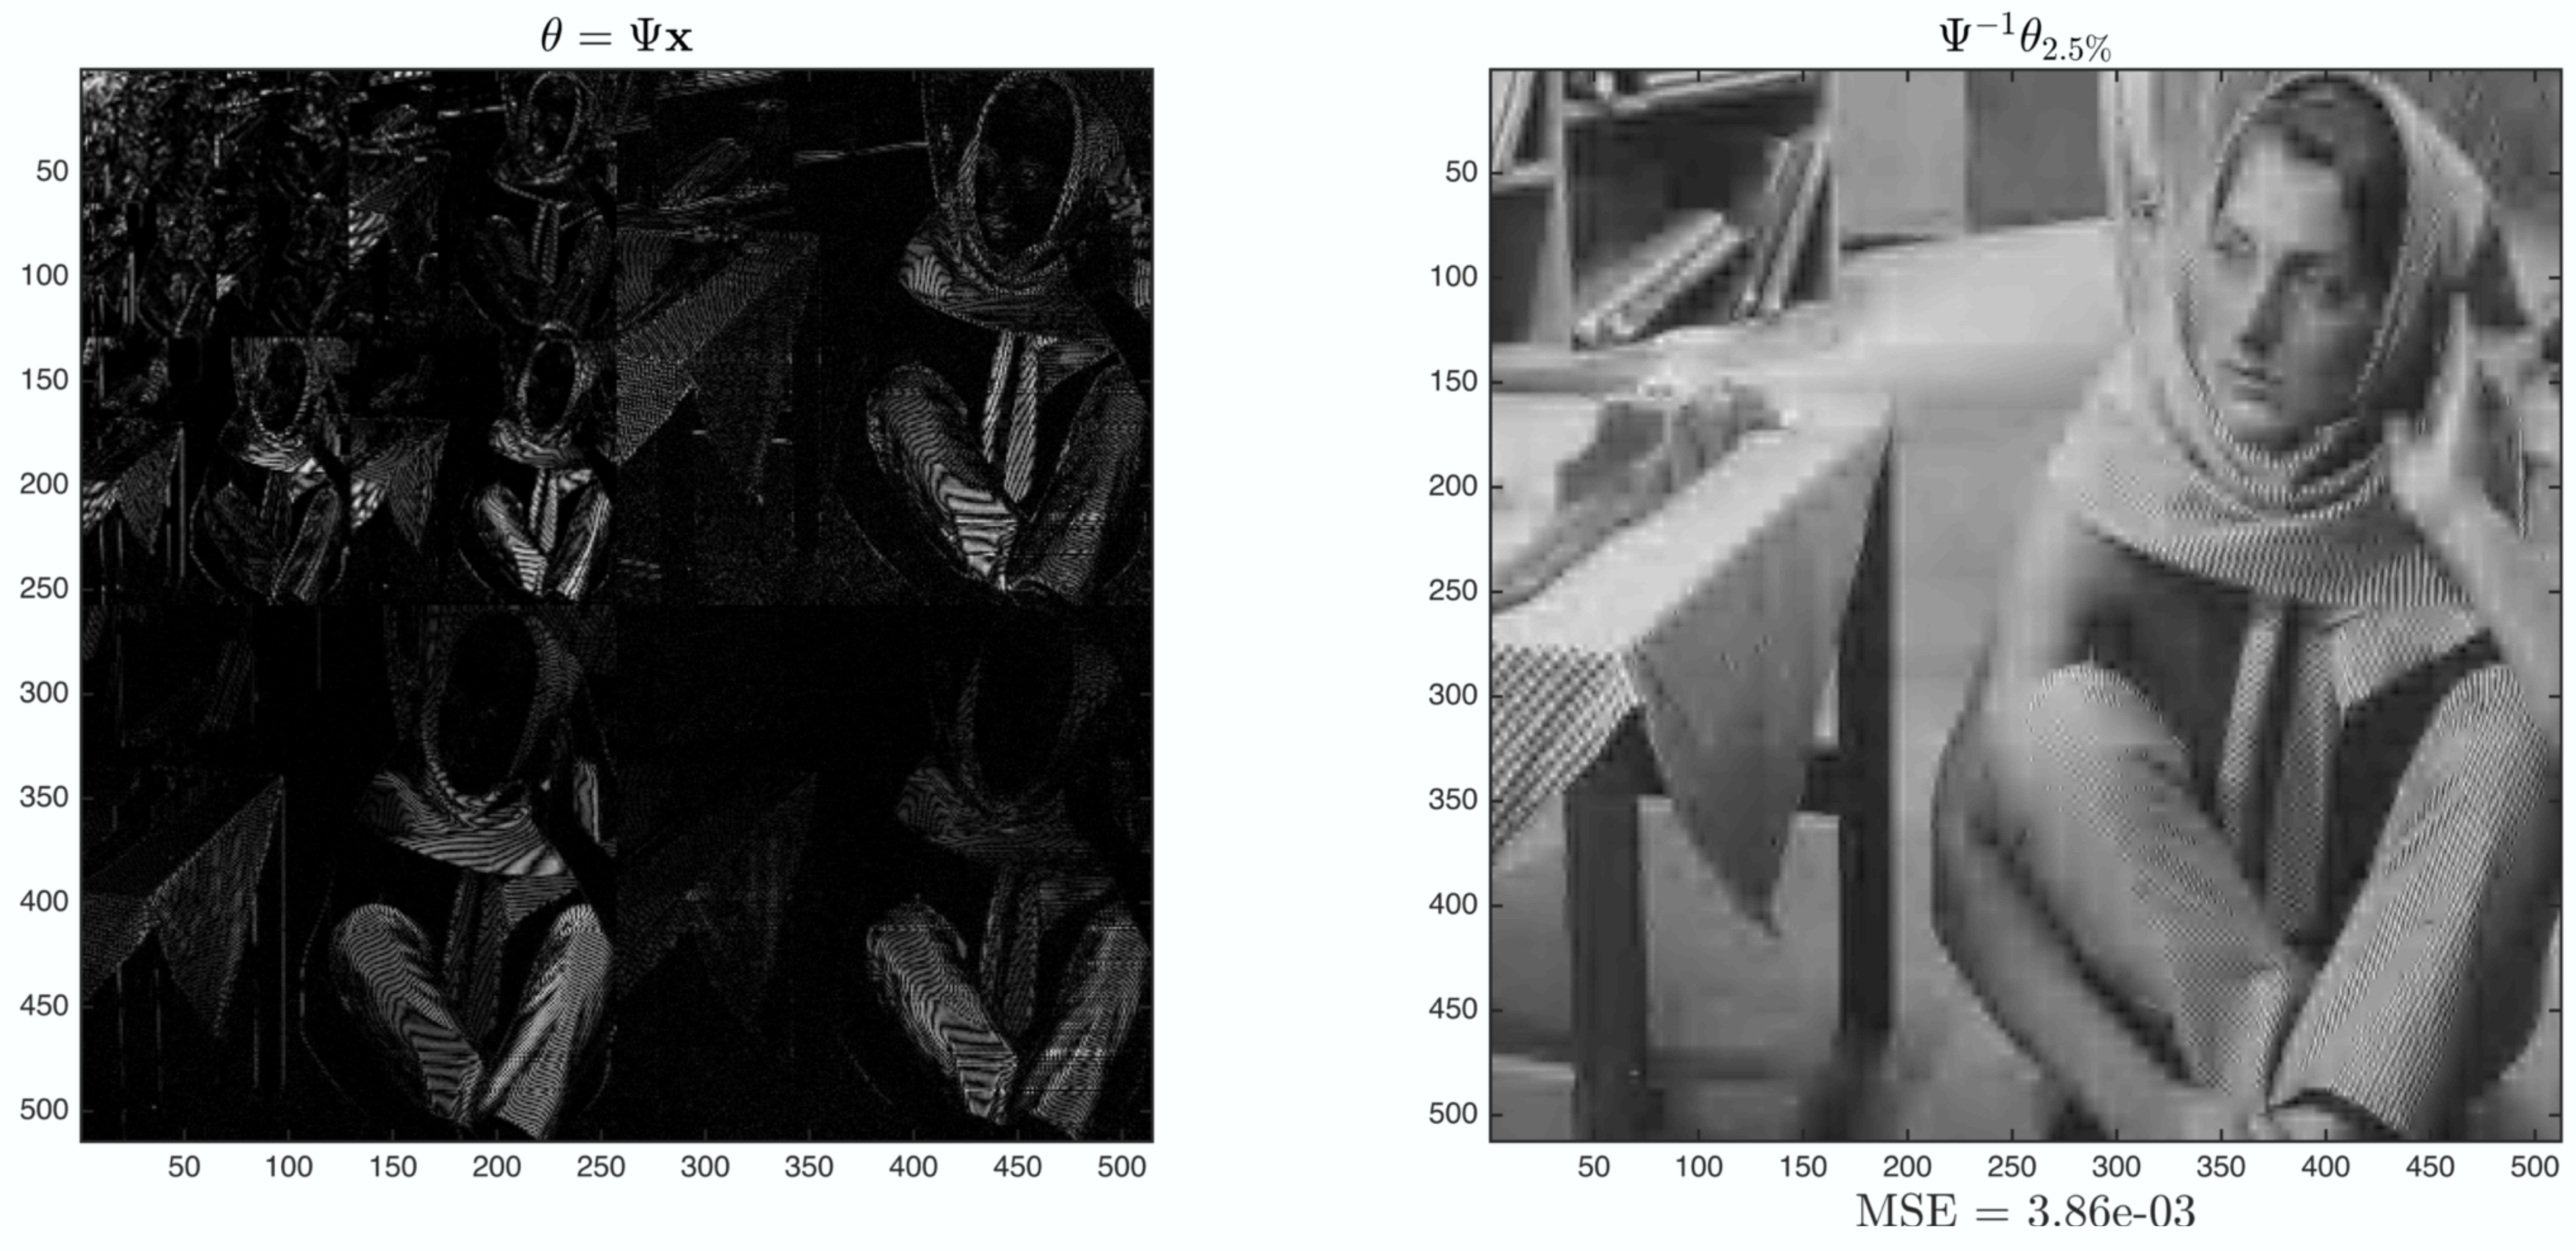

Compressed Sensing Theory

Result: Sparse Bases & Mutual Incoherence

[Donoho, Elad, & Temlyakov, 2006], **[Candès & Romberg, 2007]** Given two orthobases of $\mathbb{R}^N$, F and Ψ, the *mutual coherence* between the orthobases is defined to be

$$\mu(F, \Psi) \triangleq \max_{i,j} |\langle F_j, \Psi_i \rangle|,$$

where F_j and Ψ_i refer to the columns of the matrices F and Ψ, respectively.

- Subsequently, $\mu(F, \Psi) \in \left[1, \sqrt{N}\right]$
- Maximal *incoherence* at $\mu(F, \Psi) = 1$, e.g. time (spike) and frequency (Fourier) bases.

Effectively: A measure of the similarity between two domains.

Compressed Sensing Theory

Result: Sparse Bases & Mutual Incoherence

[**Candès & Romberg, 2007**] Given *random* sampling matrix F and that the representation θ of $\mathbf{x}$ in the basis Ψ is K-sparse, if

$$M \geq C \cdot \mu^2(F, \Psi) \cdot K \cdot \log N,$$

then the ℓ_1 recovered solution is exact with overwhelming probability.

- Desire maximally incoherent pairs (F, Ψ)

[**Candès & Wakin, 2008**] Random matrices are largely incoherent with any fixed basis Ψ. For random orthobasis F,

$$\mu(F, \Psi) = \sqrt{2 \log N} \quad \text{w.h.p.}$$

Compressed Sensing: Two Parts

I. Random Sampling

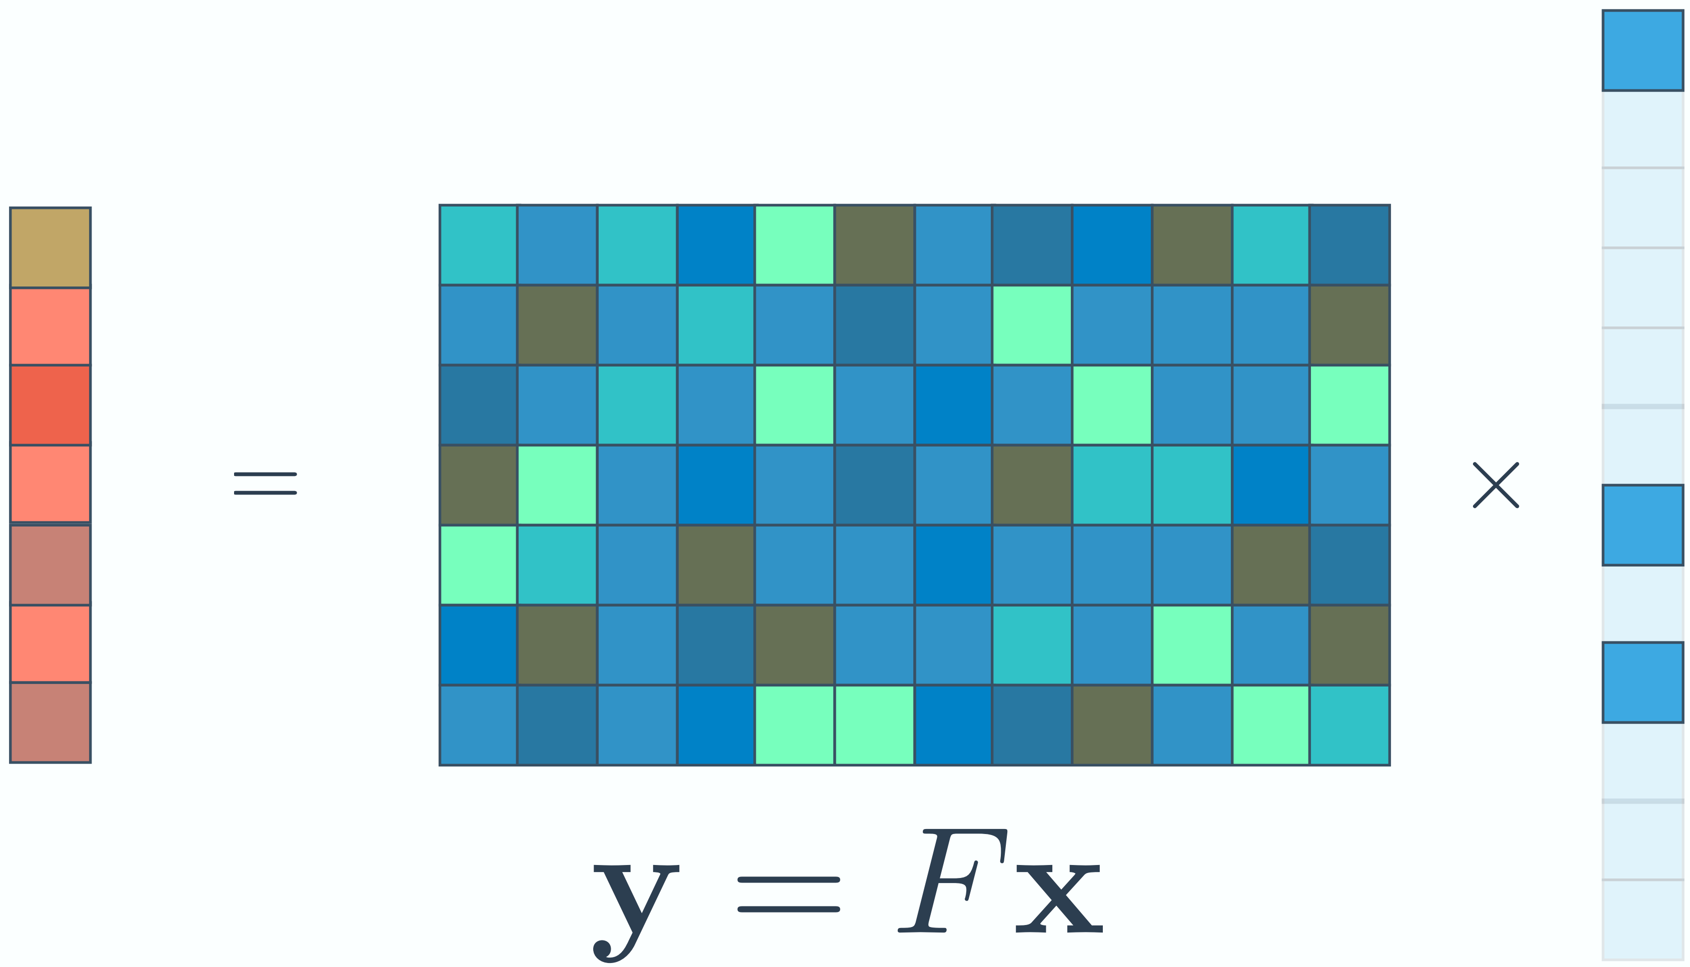

$$\mathbf{y} = F\mathbf{x}$$

II. Sparse Reconstruction

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{a}} \quad ||\mathbf{a}||_1 \quad s.t. \quad \mathbf{y} = F\mathbf{a}$$

Perspective: Universal Encoder

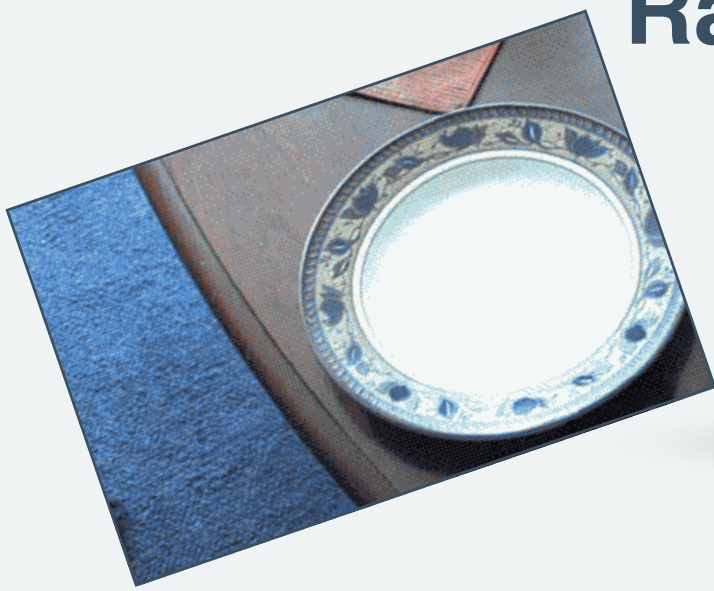

Raw (Massive) Data

High-Res Senors

JPEG/J2K, H.264/5, …

Encoder
(Compression & Quantization)

Heavy & Slow

Decoder
(Reconstruction/Decompression)

Light & Fast

Perspective: Universal Encoder

Raw (Massive) Data

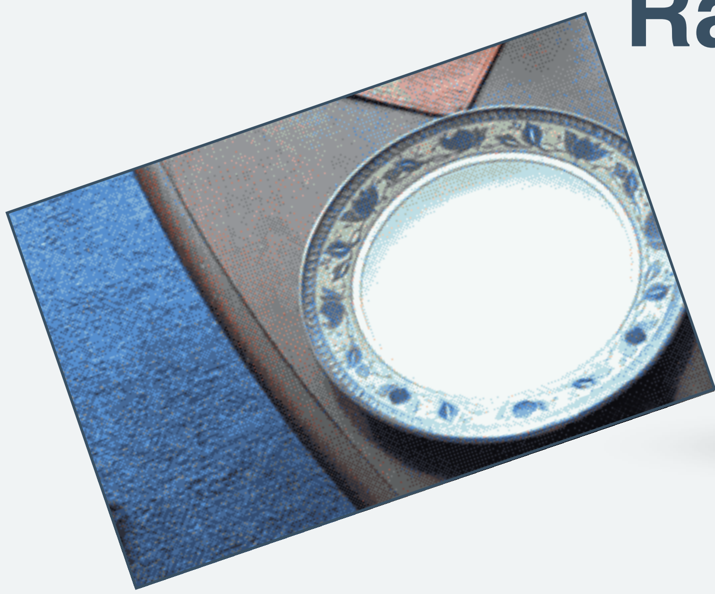

JPEG/J2K, H.264/5, …

Encoder
(Compression & Quantization)

1% of bits

99% of bits

Why did we need so many bits in the first place?

Perspective: Universal Encoder

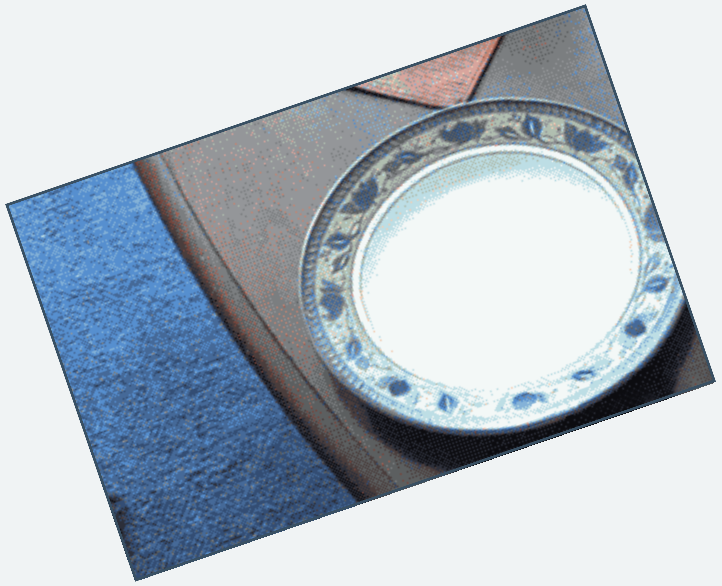

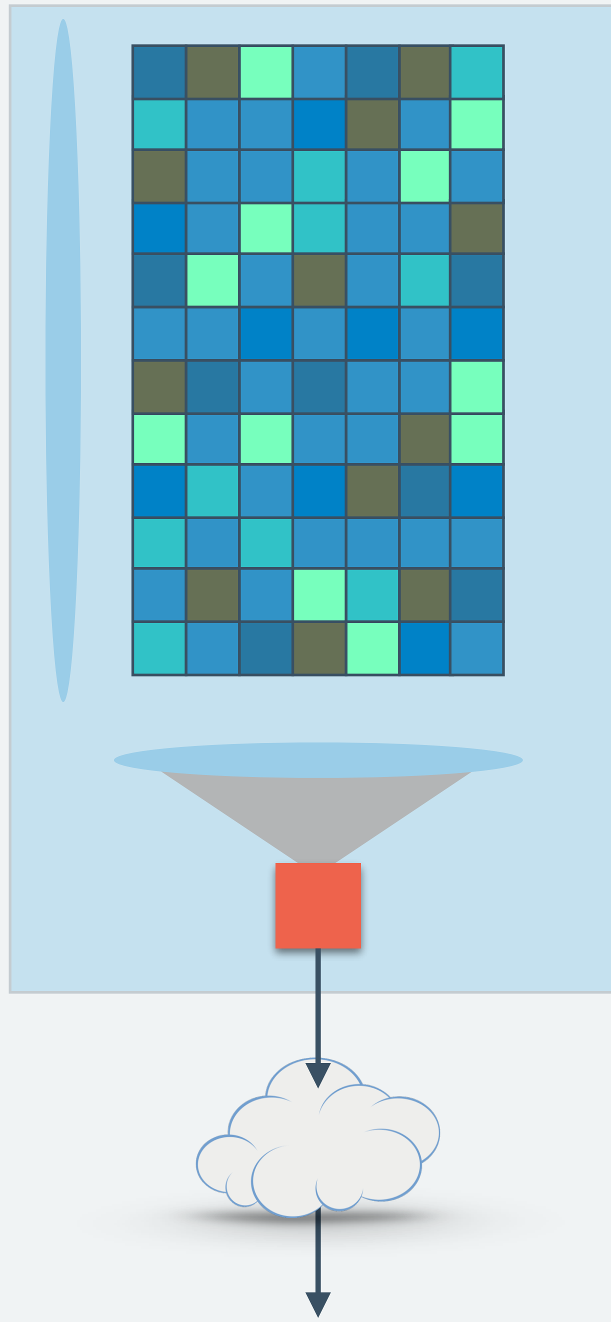

Light & Fast
(Instantaneous?)

Low-Res (*single?*) Sensor

(potentially)
Heavy & Slow

Decoder
(Reconstruction/Decompression)

Perspective: Universal Encoder

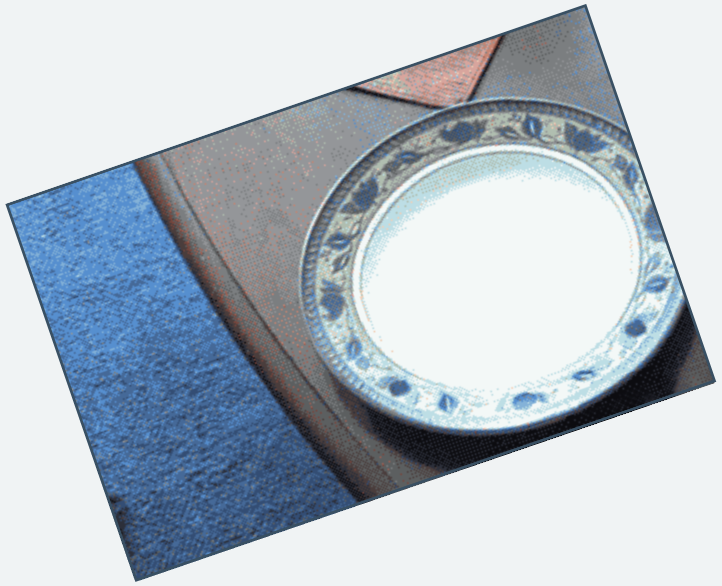

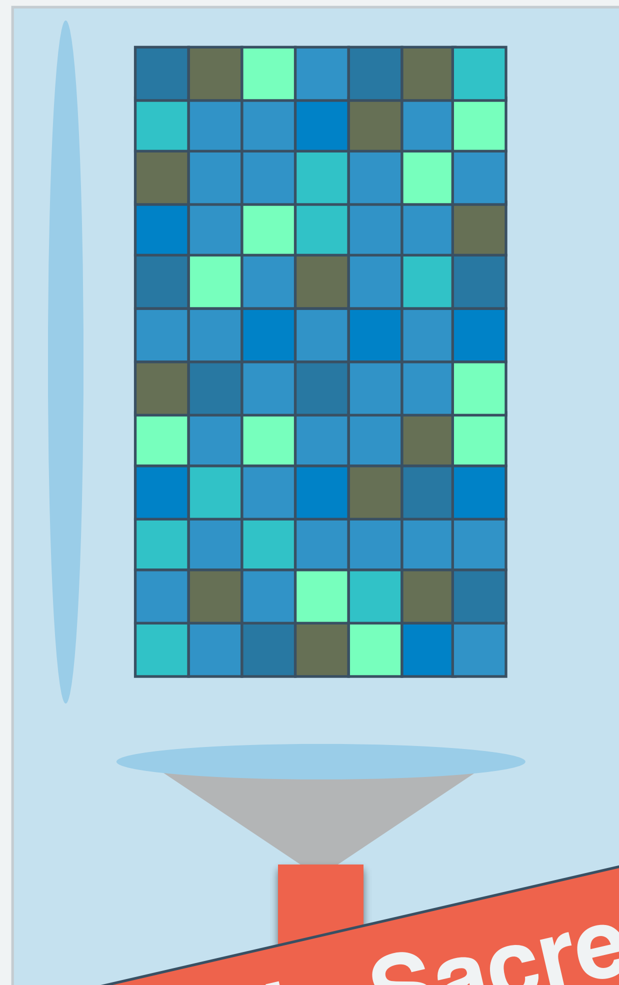

Light & Fast
(Instantaneous?)

ow-Res (*single?*) Sensor

Every Bit is Sacred!

Every Bit is Meaningful!

(potentially)
Heavy & Slow

Decoder
(Reconstruction/Decompression)

Perspective: Universal Encoder

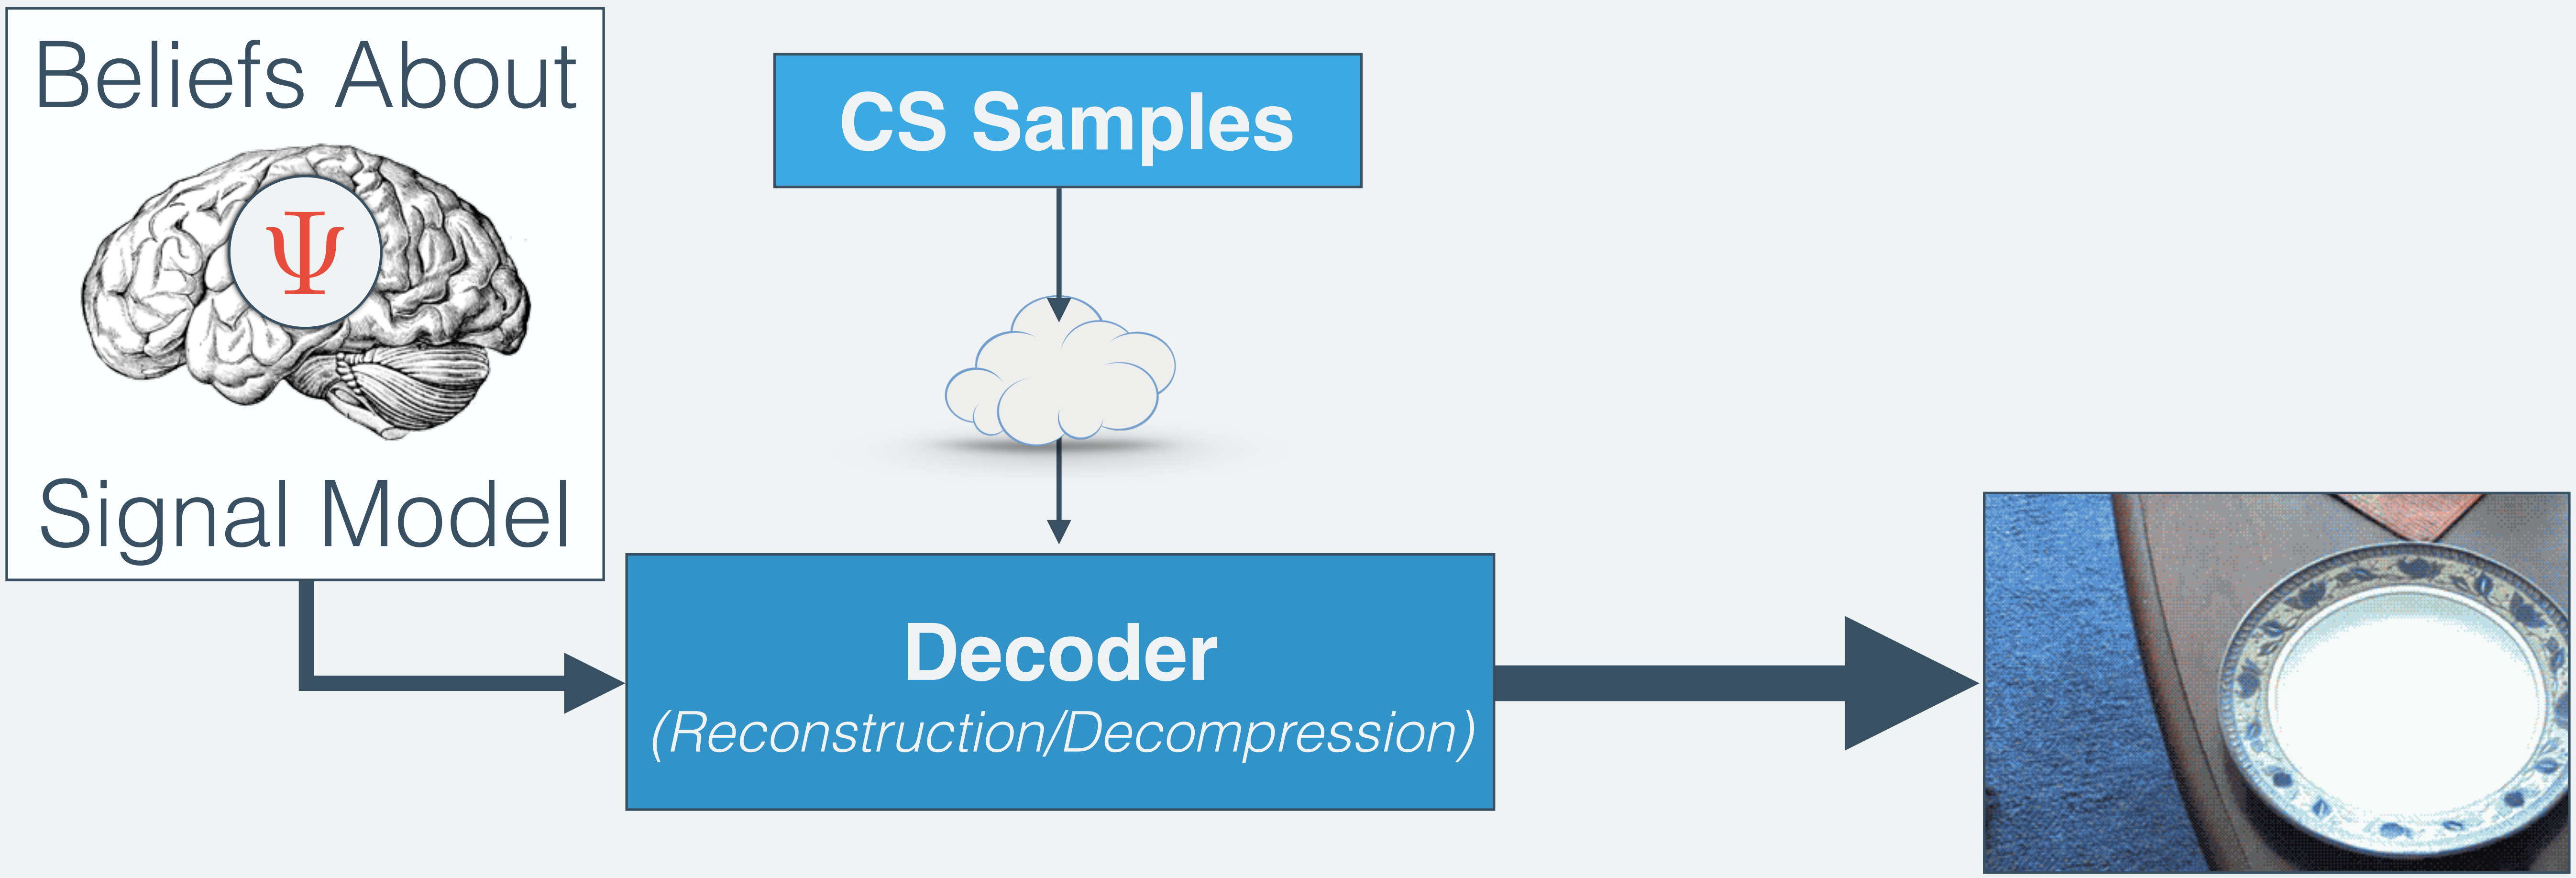

Priors: For fixed $\mathbf{M}$, the information we can bring to the table about $\mathbf{x}$ *a priori*, controls the degree to which we can recover the signal.

Perspective: Universal Encoder

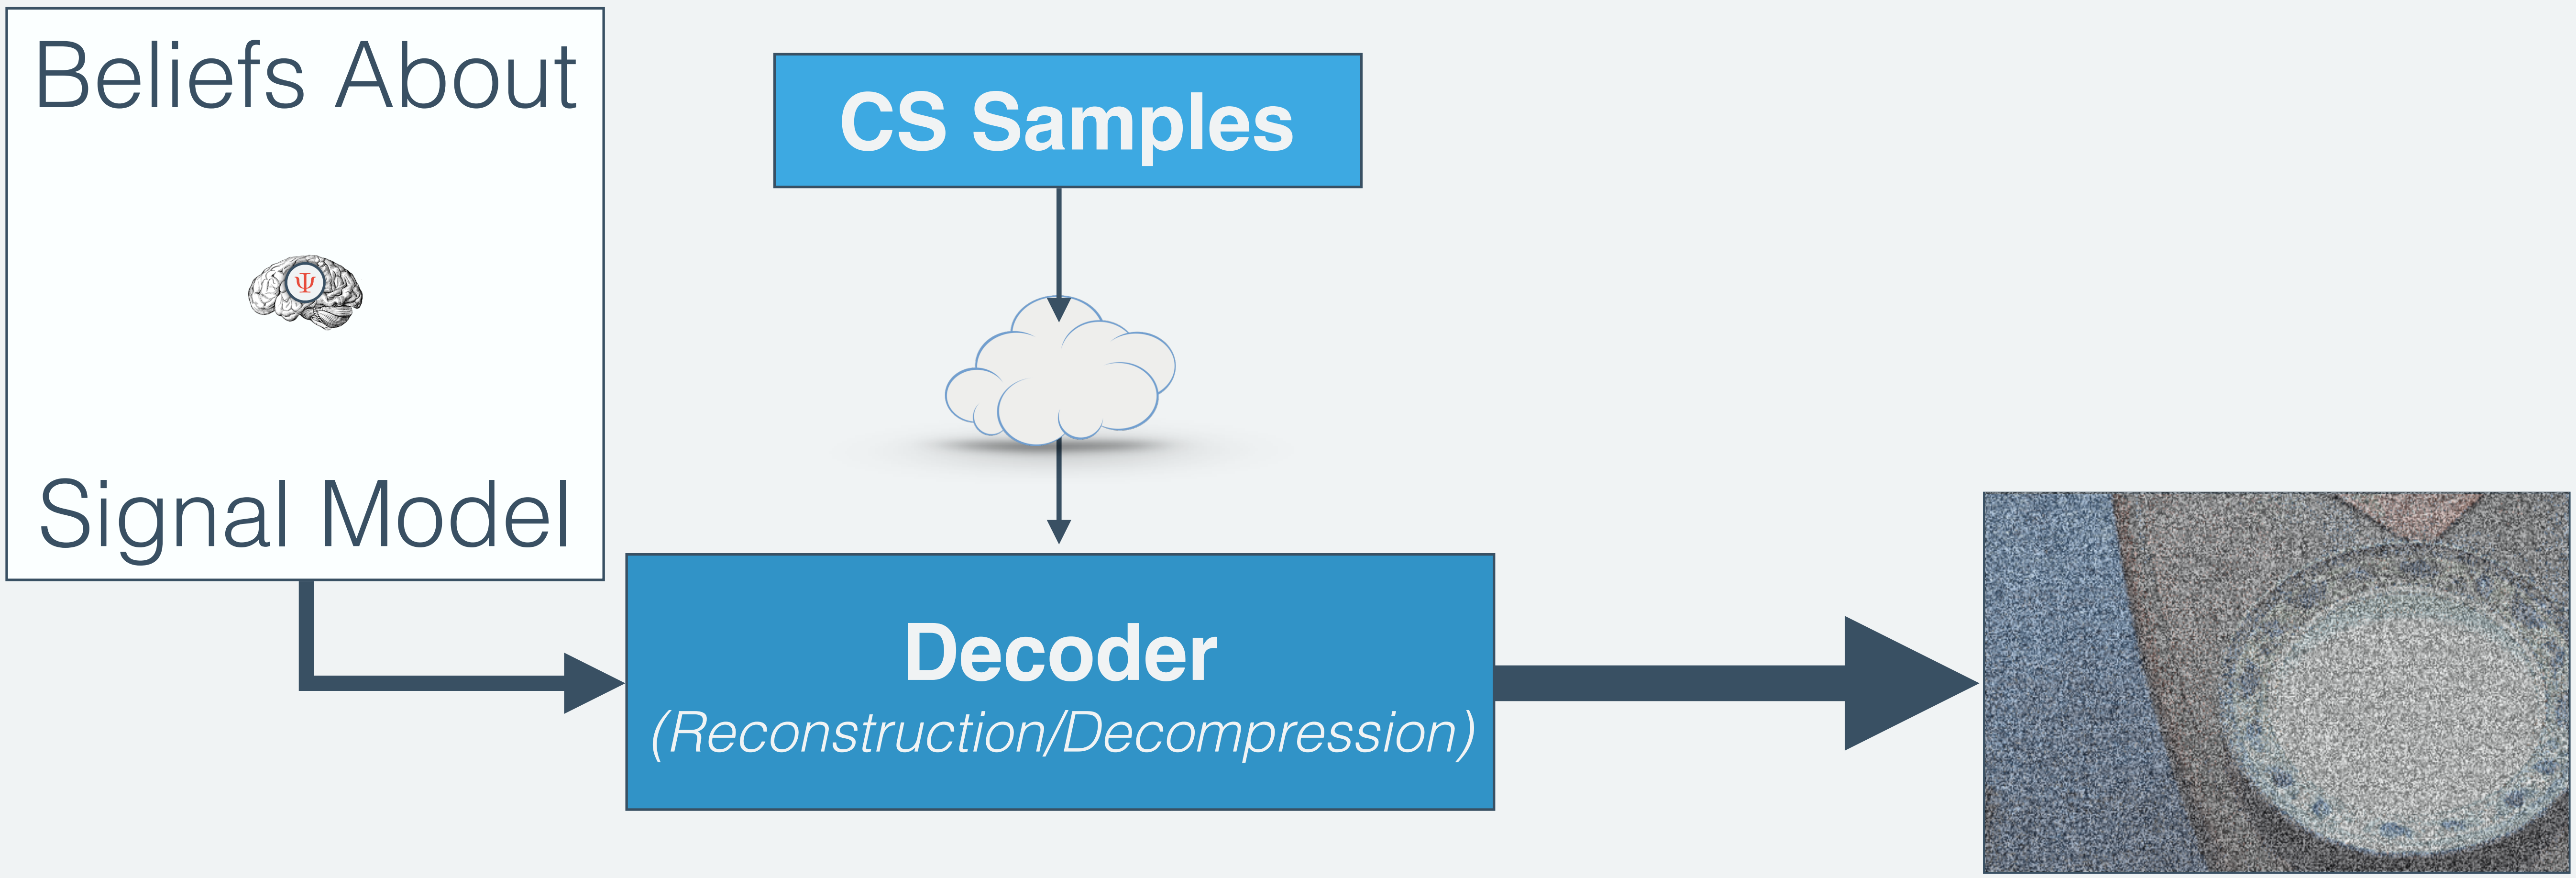

Priors: For fixed **M**, the information we can bring to the table about **x** *a priori*, controls the degree to which we can recover the signal.

Reconstruction

I. Basis Pusuit

$$\arg\min_{\mathbf{x}} \quad ||\mathbf{x}||_1 \quad s.t. \quad F\mathbf{x} = \mathbf{y}$$

Linear Program: Can be solved efficiently using any number of methods, including,

- Interior-point methods (e.g. *path-following primal-dual*)
- Simplex methods

Implementations: See the original L1-Magic Toolbox,
http://users.ece.gatech.edu/justin/l1magic/

Reconstruction

II. Basis Pursuit Denoising (BPDN), Lasso

$$\arg\min_{\mathbf{x}} \quad ||\mathbf{y} - F\mathbf{x}||_2^2 \quad s.t. \quad ||\mathbf{x}||_1 \leq K$$

$$\arg\min_{\mathbf{x}} \quad ||\mathbf{x}||_1 \quad s.t. \quad ||\mathbf{y} - F\mathbf{x}||_2^2 \leq \epsilon$$

$$\arg\min_{\mathbf{x}} \quad ||\mathbf{y} - F\mathbf{x}||_2^2 + \lambda||\mathbf{x}||_1$$

Realistic: Accounts for noisy measurements.

Second Order Cone Program: Solvable via log-barrier.

Lasso: Solvable via any number of methods,(*Least Angle Regression, Gauss-Siedel, Shooting, Block Coordinate, Active Set…*), but also Iterative Soft Thresholding (**see:** TwIST, FISTA, NESTA)

Reconstruction

III. Relaxed L0

$$\arg\min_{\mathbf{x}} \quad ||\mathbf{y} - F\mathbf{x}||_2^2 \quad s.t. \quad ||\mathbf{x}||_0 \leq K$$

Why return to non-convex? Requires greedy techniques…

- Easy-to-implement solvers
- Relaxed RIP requirements, potentially lower requirements on **M**
- Generally computationally/memory efficient
- Robust to inconsistencies/pathologies of **F**

Solvable via: Orthogonal Matching Pursuit (OMP), Stagewise OMP, Compressed Sampling MP, Iterative Hard Thresholding.

Reconstruction

IV. Probabilistic

$$P(\mathbf{x}|F,\mathbf{y}) \propto P_0(\mathbf{x})P(\mathbf{y}|F,\mathbf{x})$$

$$\arg\max_{\mathbf{x}} \quad P(\mathbf{x}|F,\mathbf{y})$$

$$\arg\max_{\mathbf{x}} \quad \int \mathrm{d}\mathbf{x} \;\; \mathbf{x} \cdot P(\mathbf{x}|F,\mathbf{y})$$

Powerful Analytics: Can use all the tools of statistical mechanics to study CS.

Powerful Performance: Bayes-optimal recovery thresholds, but conditions are brittle.

Solve via: relaxed-Belief Propagation, Approximate Message Passing, Expectation Propagation.

Sampling Design Examples

Magnetic Resonance Imaging (MRI)

M. Lustig, D. Donoho, and J. M. Pauly, “Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging,” Magnetic Resonance in Medicine, vol. 58, no. 6, 2007.

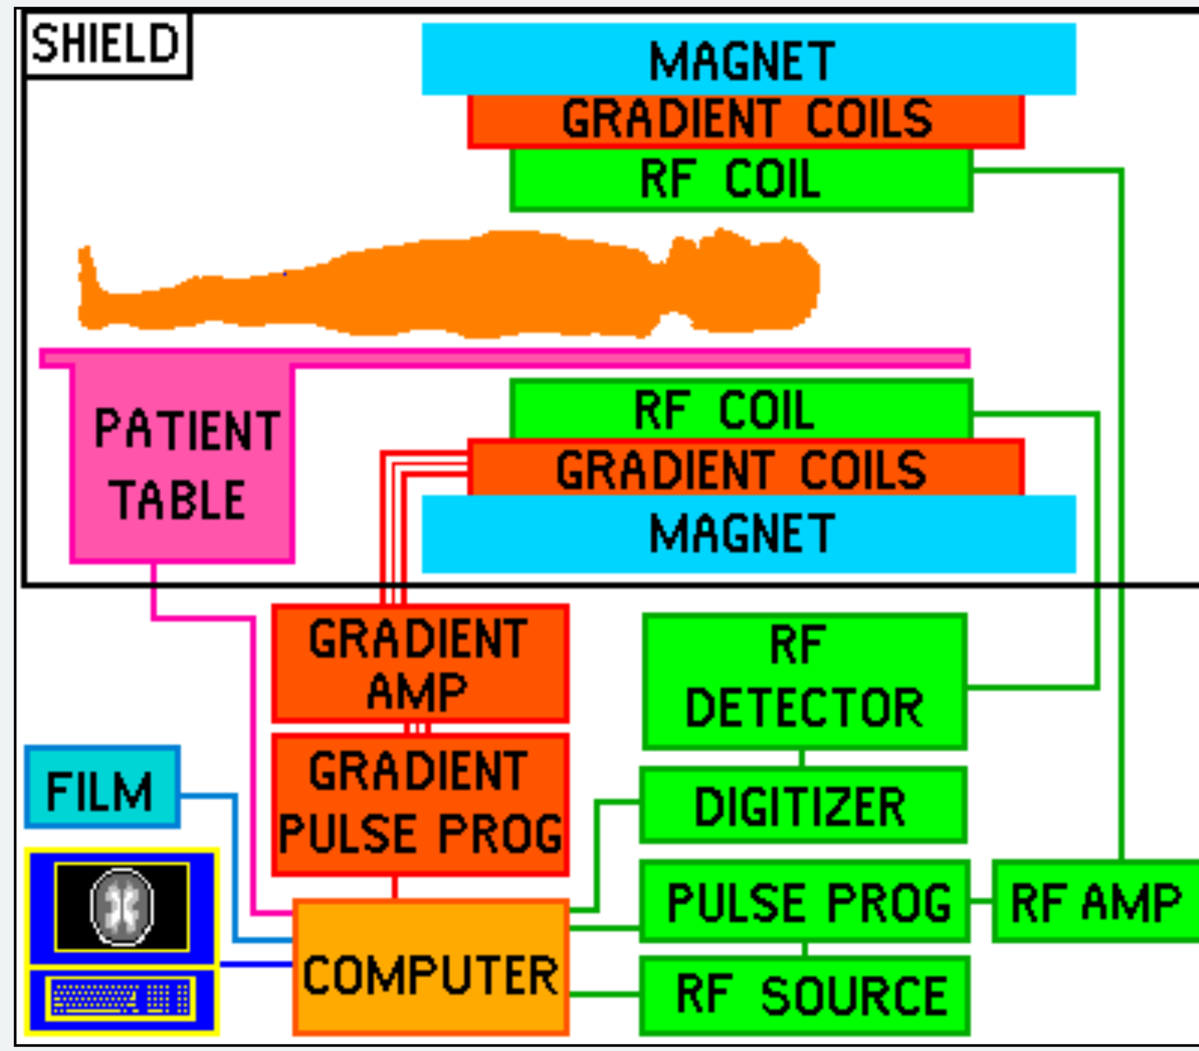

Figures from paper.

Sampling Design Examples

Magnetic Resonance Imaging (MRI)

M. Lustig, D. Donoho, and J. M. Pauly, “Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging,” Magnetic Resonance in Medicine, vol. 58, no. 6, 2007.

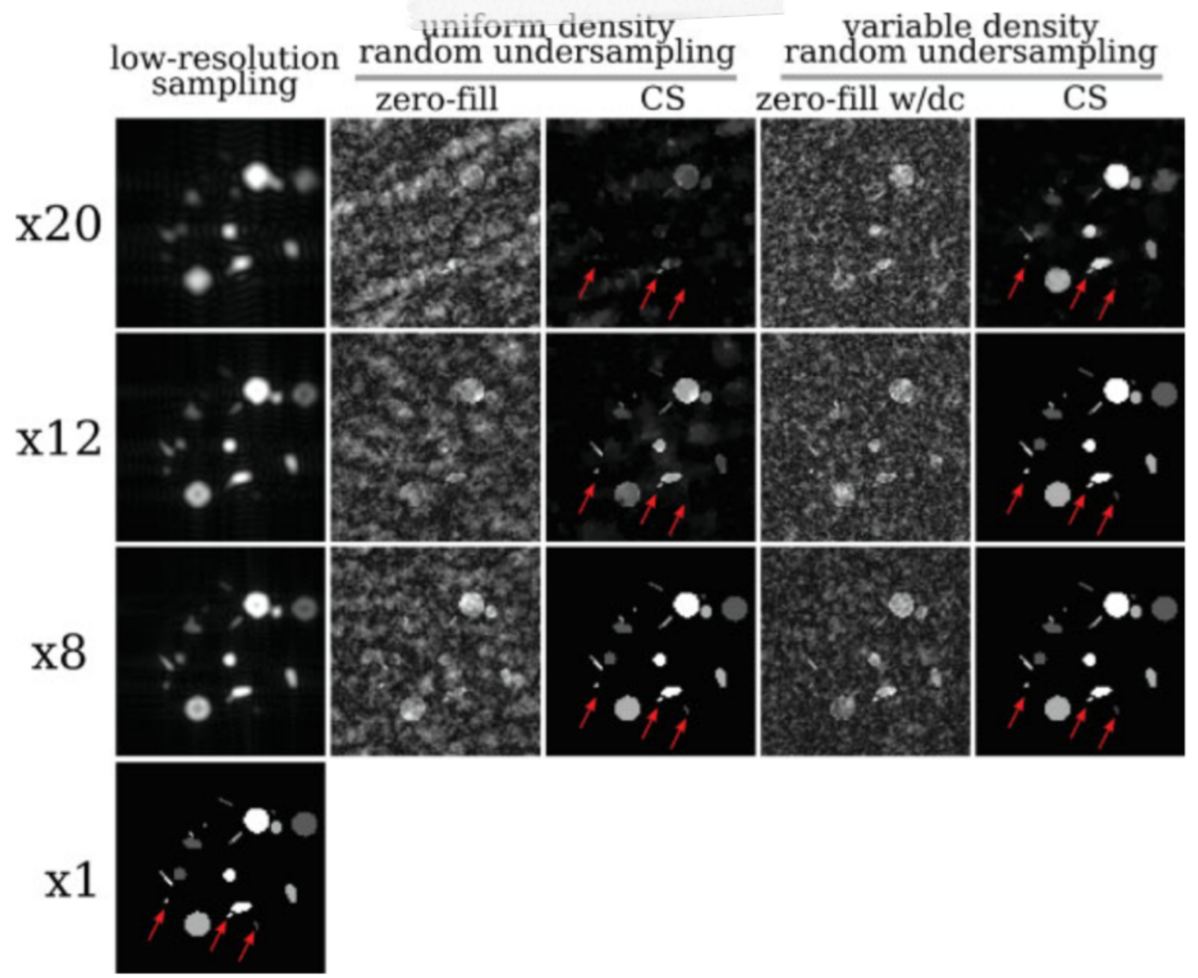

FIG. 6. Simulation: Reconstruction artifacts as a function of acceleration. The LR reconstructions exhibit diffused boundaries and loss of small features. The ZF-w/dc reconstructions exhibit an significant increase of apparent noise due to incoherent aliasing, the apparent noise appears more “white” with variable density sampling. The CS reconstructions exhibit perfect reconstruction at 8- and 12-fold (only var. dens.) accelerations. With increased acceleration there is loss of low-contrast features and not the usual loss of resolution. The reconstructions from variable density random undersampling significantly outperforms the reconstructions from uniform density random undersampling. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

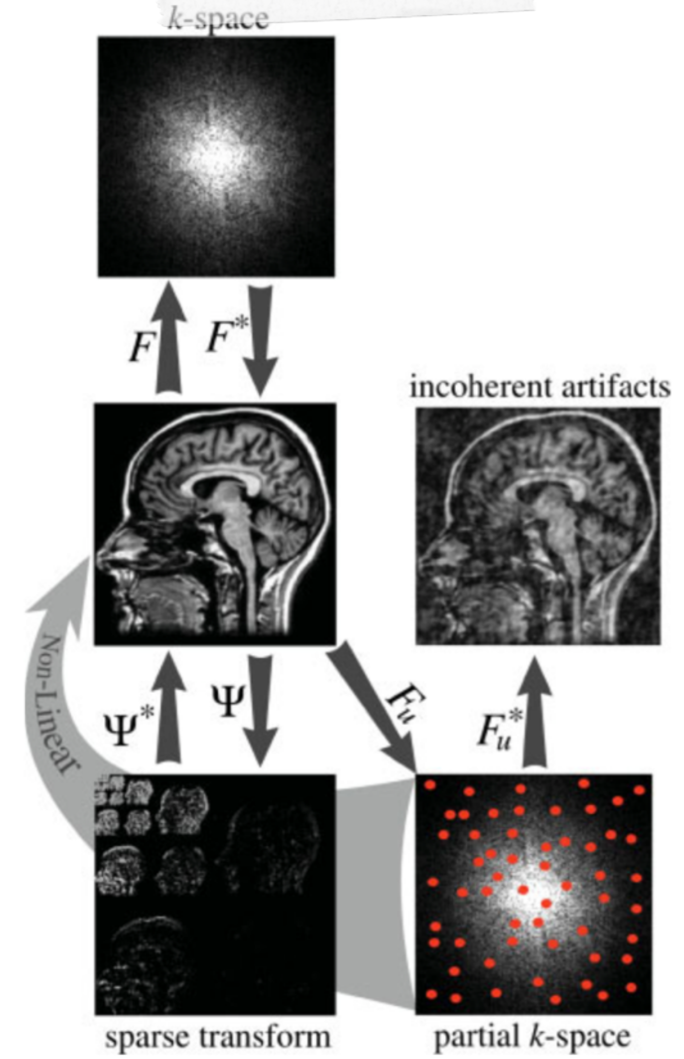

FIG. 1. Illustration of the domains and operators used in the paper as well as the requirements of CS: sparsity in the transform domain, incoherence of the undersampling artifacts, and the need for non-linear reconstruction that enforces sparsity. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figures from paper.

Sampling Design Examples

Single Pixel Camera

M. Duarte et al, “Single-Pixel Imaging via Compressive Sampling,” Signal Processing Magazine, vol. 25, no. 2, 2008.

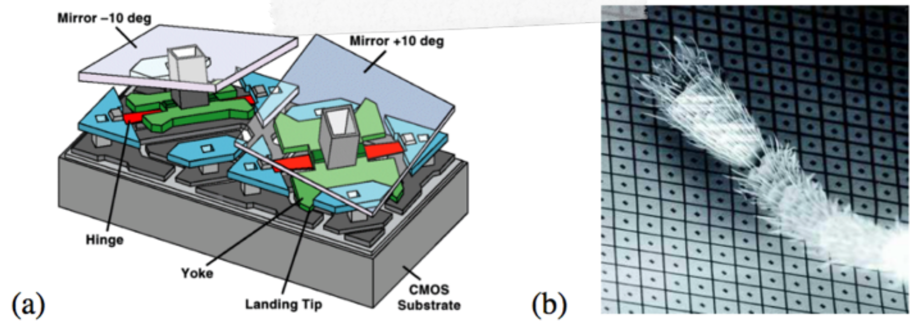

Fig. 6. *(a) Schematic of two mirrors from a Texas Instruments digital micromirror device (DMD). (b) A portion of an actual DMD array with an ant leg for scale. (Image provided by DLP Products, Texas Instruments.)*

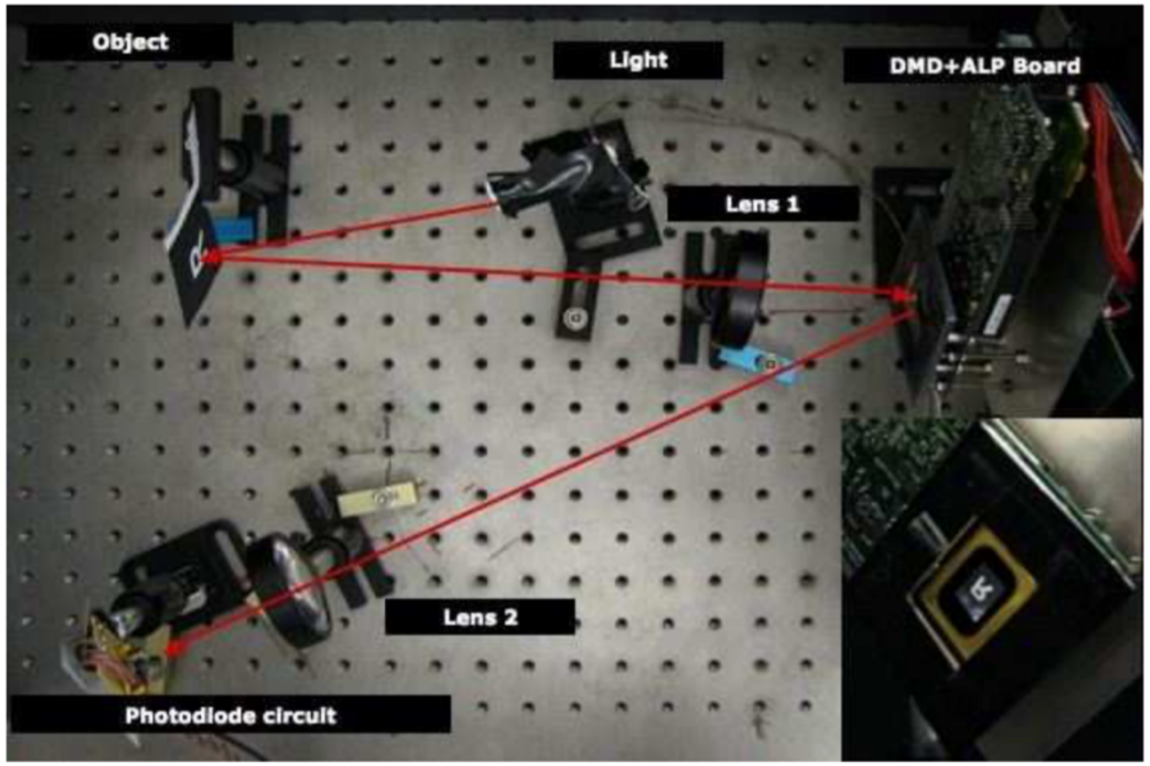

Fig. 1. *Aerial view of the single-pixel compressive sampling (CS) camera in the lab [5].*

Fig. 2. *Single-pixel photo album. (a) 256×256 conventional image of a black-and-white R. (b) Single-pixel camera reconstructed image from $M = 1300$ random measurements ($50\times$ sub-Nyquist). (c) 256×256 pixel color reconstruction of a printout of the Mandrill test image imaged in a low-light setting using a single photomultiplier tube sensor, RGB color filters, and $M = 6500$ random measurements.*

Figures from paper.

Sampling Design Examples

Structured Illumination and Fluorescence Microscopy

V. Studer et al, “Compressive Fluorescence Microscopy for Biological and Hyperspectral Imaging,” PNAS, vol. 109, no. 26, 2012.

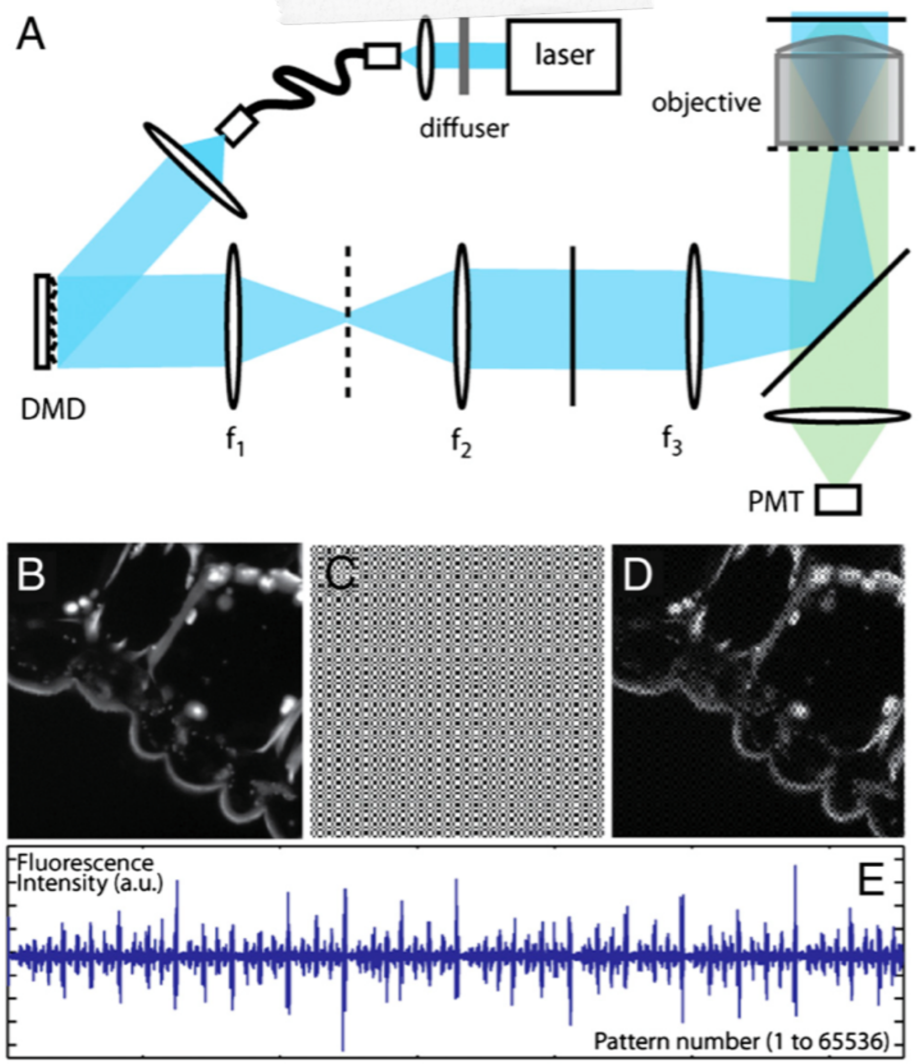

Fig. 1. (*A*) Experimental setup. The dotted and plain segments correspond to planes respectively conjugated to the pupil and sample planes. (*B*) Slice of lily anther (endogenous fluorescence with epifluorescence microscopy image recorded on a CCD camera). (*C*) Projection of a Hadamard pattern on a uniform fluorescent sample. (*D*) Projection of the same Hadamard pattern on the biological sample. (*E*) Fluorescence intensity during an acquisition sequence.

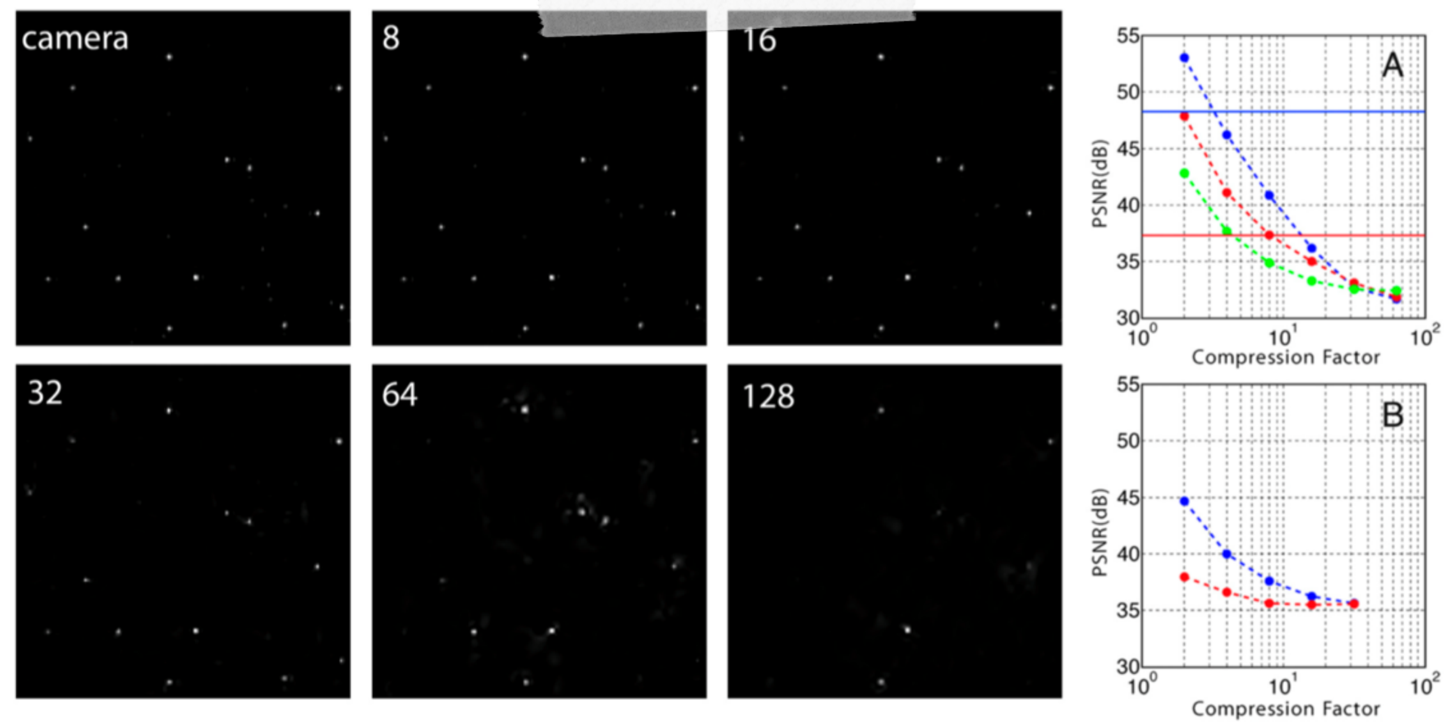

Fig. 2. Top left to bottom right: camera snapshot and reconstructed 256-by-256 bead images for values of the undersampling ratio equal to 8, 16, 32, 64, and 128. (*A*) Plot of the PSNR (see text) for a nominal illumination level (blue curve) and for the same level reduced by a factor 10 (red curve) and a factor of 100 (green curve) (simulated data). The solid lines correspond to the PSNR in raster scan for the same surfacic illumination. (*B*) Same as (*A*) for the experimental data.

Sampling Design Examples

Random Lens Imager

R. Fergus, A. Torralba, and W. T. Freeman, “Random Lens Imaging,” Tech. Report, MIT, no. MIT-CSAIL-2006-058, September, 2006.

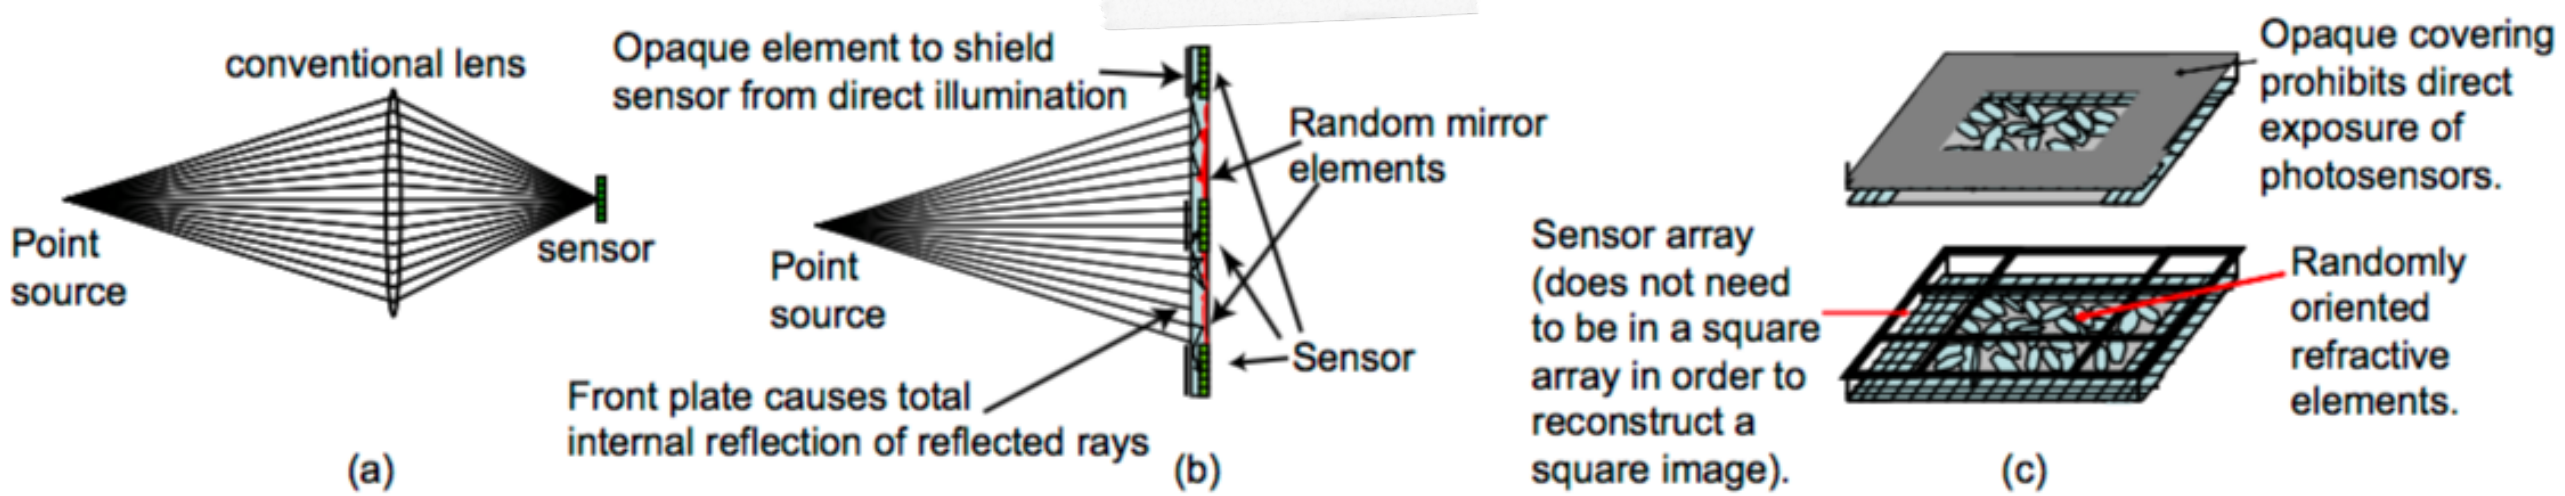

Figure 2: Candidate physical designs. (a) Conventional lens. (b) Random lens using reflective elements, (c) Random lens using refractive elements.

Figures from paper.

Sampling Design Examples

Random Lens Imager

R. Fergus, A. Torralba, and W. T. Freeman, “Random Lens Imaging,” Tech. Report, MIT, no. MIT-CSAIL-2006-058, September, 2006.

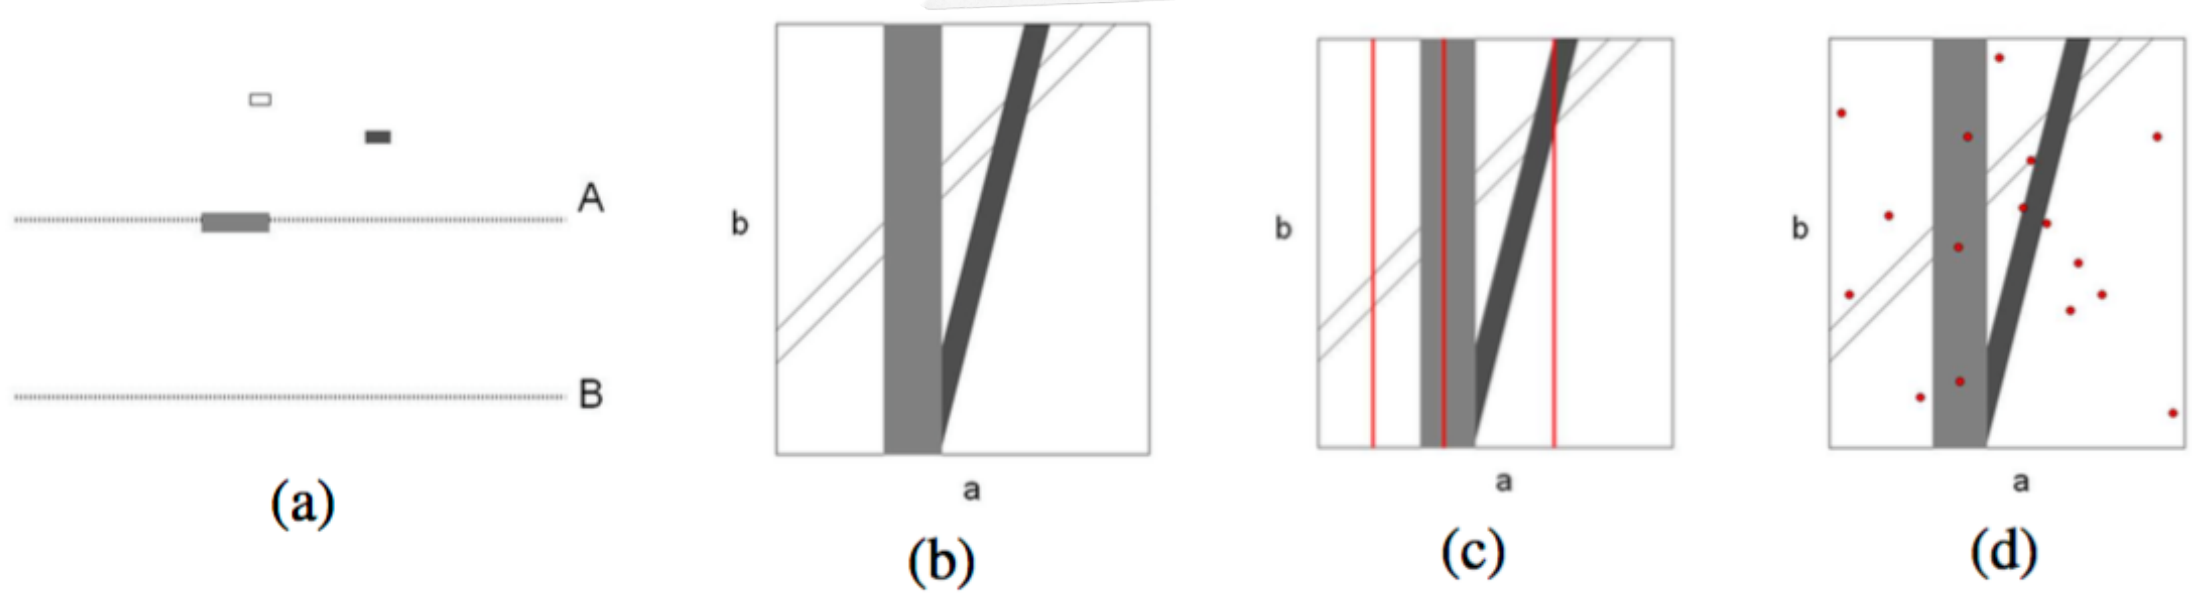

Figure 3: (a) Consider these 3 Lambertian objects in our 2-d world. (b) The resulting lightfield, or intensity of each ray (a,b). Under most conditions, the lightfield exhibits extraordinary structure and redundancy. (c) Conventional lens, focussed at A, integrates at each sensor position along vertical slices of this lightfield, like the 3 integral lines shown in red. (d) A random lens sensor element integrates over a pseudo-random set of lightfield points.

Figures from paper.

Sampling Design Examples

Random Lens Imager

R. Fergus, A. Torralba, and W. T. Freeman, "Random Lens Imaging," Tech. Report, MIT, no. MIT-CSAIL-2006-058, September, 2006.

Figure 5: A closeup of the random reflective surface and camera setup used in our experiments. The schematic diagram on the right shows the light path to the sensor.

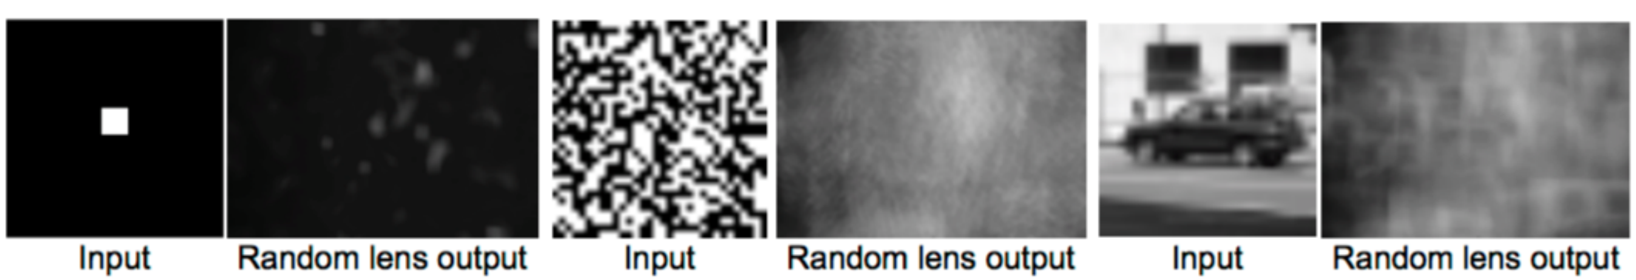

Figure 6: Examples of pictures taken with our random lens camera. Each pair shows an image projected on the wall, and the output of the camera.

1. Calibrate
2. Recovery

Figures from paper.

Sampling Design Examples

Multiply Scattering Media

A. Liutkus et al, "Imaging with Nature: Compressive Imaging Using a Multiply Scattering Medium," Scientific Reports 4, 2014.

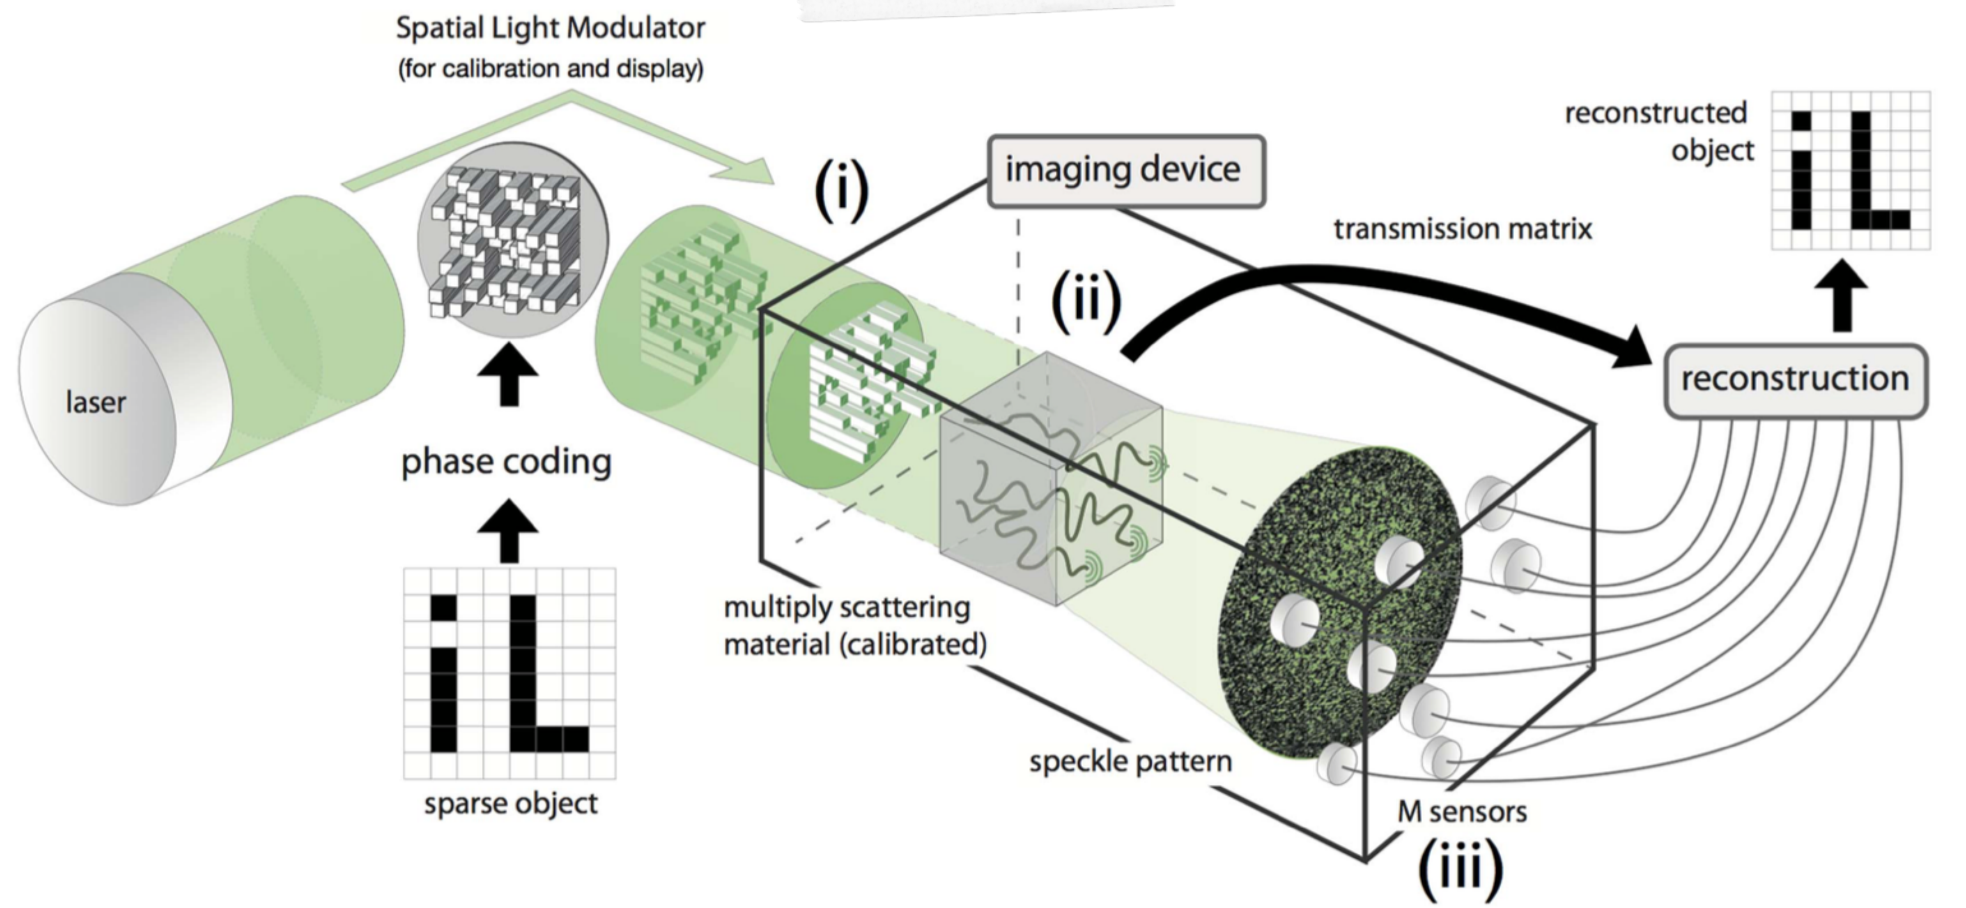

Figure 2 | Experimental setup for compressive imaging using multiply scattering medium. Within the imaging device, waves coming from the object (i) go through a scattering material (ii) that efficiently multiplexes the information to all *M* sensors (iii). Provided the transmission matrix of the material has been estimated beforehand, reconstruction can be performed using only a limited number of sensors, potentially much lower than without the scattering material. In our optical scenario, the light coming from the object is displayed using a spatial light modulator.

Figures from paper.

Sampling Design Examples

Multiply Scattering Media

A. Liutkus et al, "Imaging with Nature: Compressive Imaging Using a Multiply Scattering Medium," Scientific Reports 4, 2014.

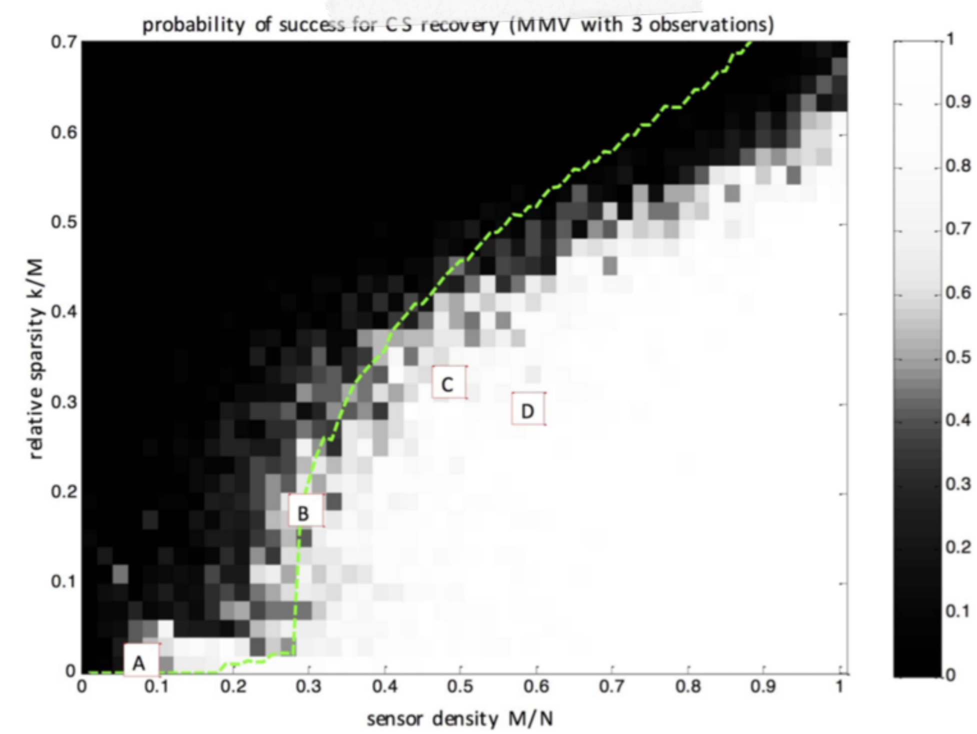

Figure 5 | Probability of success for CS recovery. Experimental probability of successful recovery (between 0 and 1) for a k-sparse image of N pixels via M measurements. On the x-axis is displayed the sensor density ratio M/N. A ratio of 1 corresponds to the Nyquist rate, meaning that all correct reconstructions found in this figure beat traditional sampling. On the y-axis is displayed the relative sparsity ratio k/M. A clear phase transition between failure and success is observable, which is close to that obtained by simulations (dashed line), where exactly the same experimental protocol was conducted with simulated noisy observations both for calibration and imaging. Boxes A, B, C and D locate the corresponding examples of Fig. 4. Each point in this 50×50 grid is the average performance over approximately 50 independent measurements. This figure hence summarizes the results of more than 10^5 actual physical experiments.

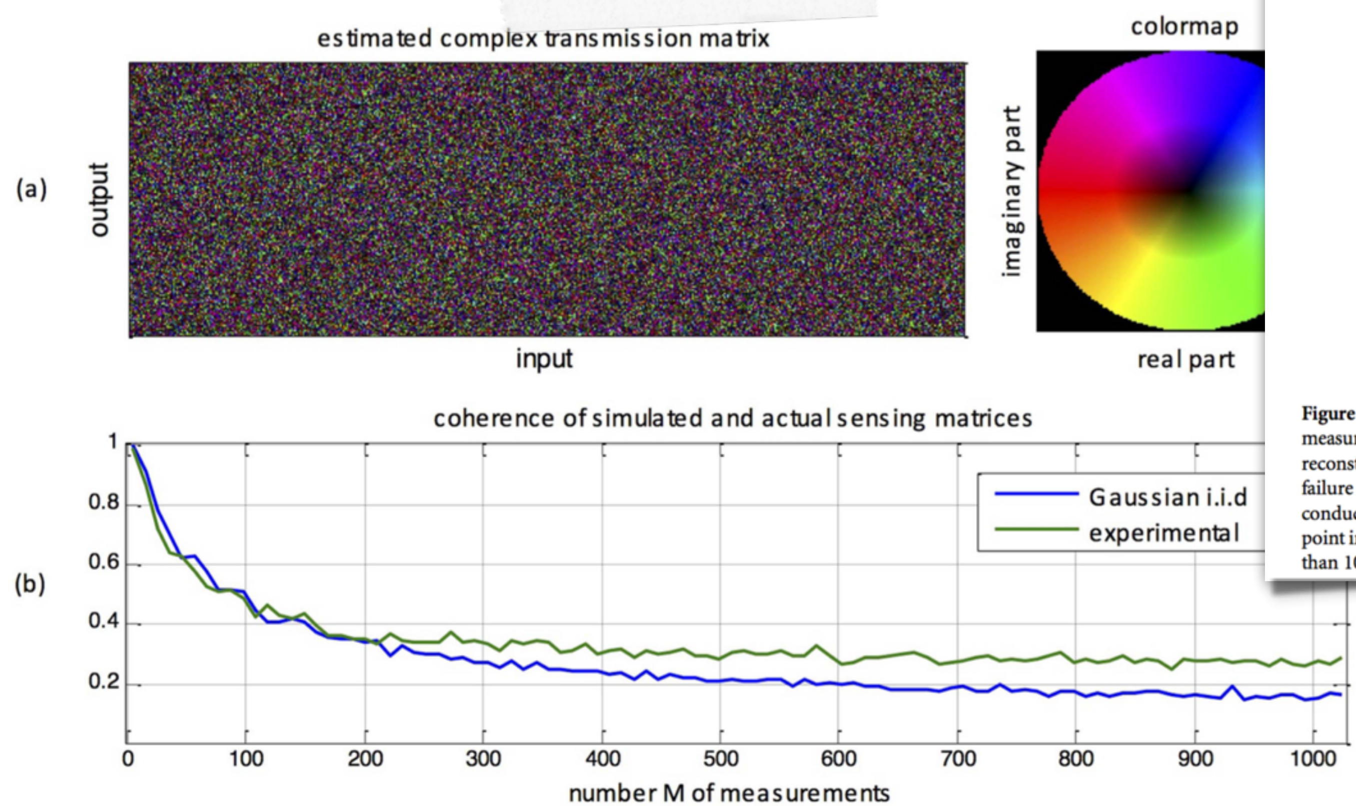

Figure 3 | Experimentally measured Transmission Matrix (TM). (a) TM for a multiply scattering material as obtained in our experimental study. (b) Coherence of sensing matrices as a function of their number M of rows, for both a randomly generated Gaussian i.i.d. matrix, and an actual experimental TM. Coherence gives the maximal colinearity between the columns of a matrix. The lower, the better is the matrix for CS.

Figures from paper.

Sampling Design Examples

Coded Aperture Snapshot Spectral Imaging (CASSI)

A. Wagadarikar et al, “Single Disperser Design for Coded Aperture Snapshot Spectral Imaging,” Applied Optics, vol 47, no. 10, 2008.

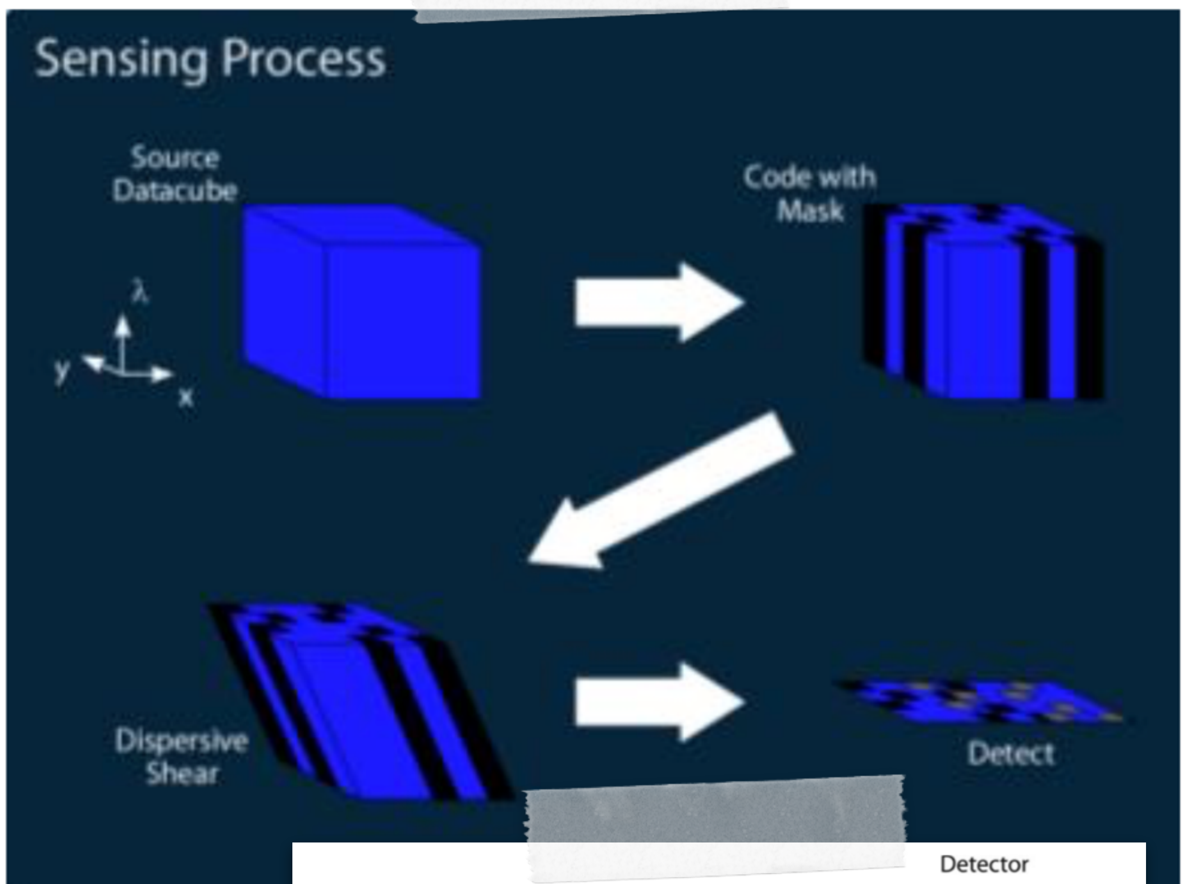

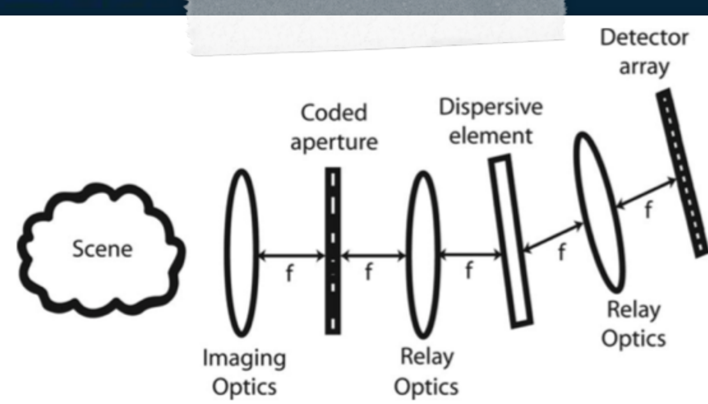

Fig. 1. Schematic of a SD CASSI. The imaging optics image the scene onto the coded aperture. The relay optics relay the image from the plane of the coded aperture to the detector through the dispersive element.

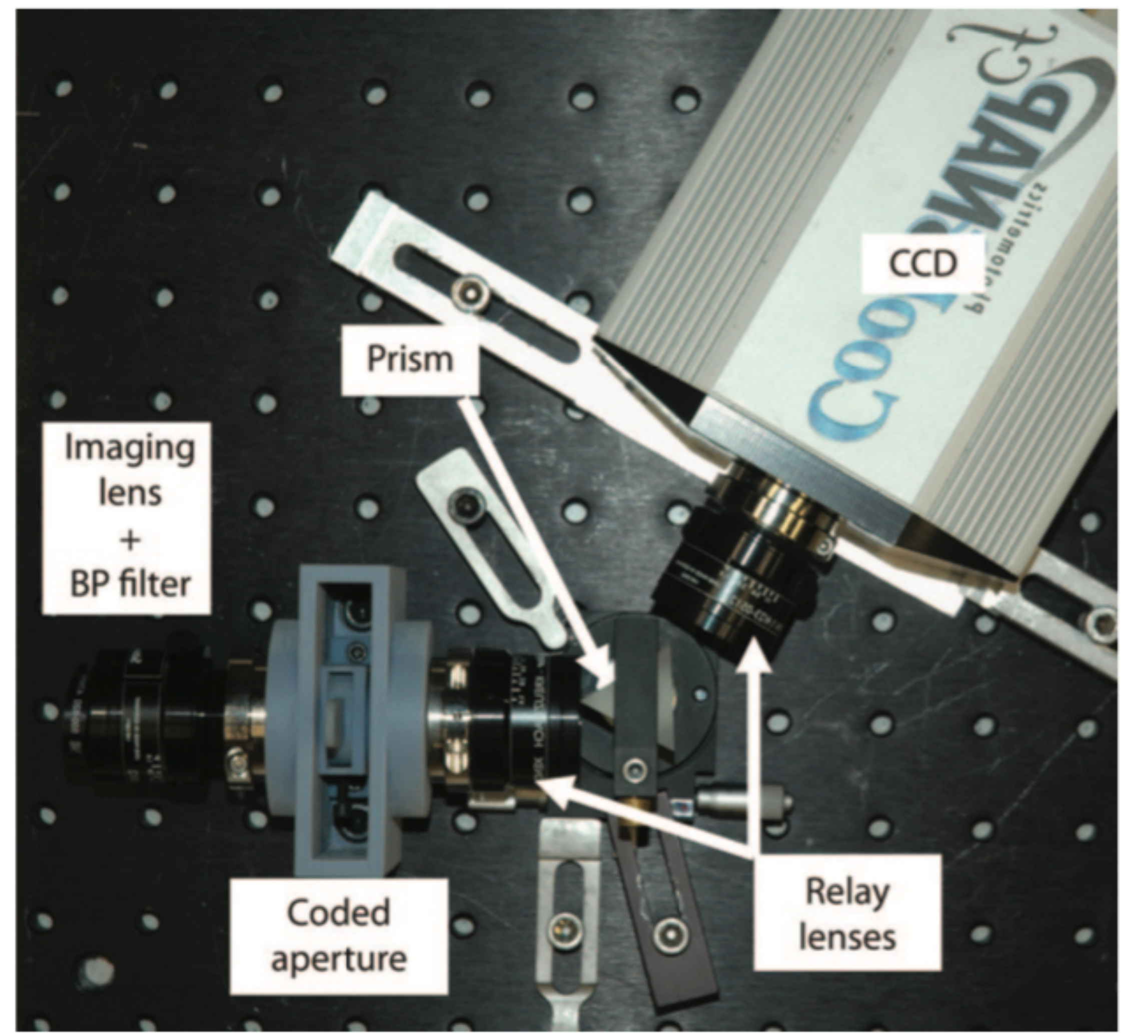

Fig. 2. (Color online) Experimental prototype of the SD CASSI.

Figures from paper.

Sampling Design Examples

ENS

Coded Aperture Snapshot Spectral Imaging (CASSI)

A. Wagadarikar et al, "Single Disperser Design for Coded Aperture Snapshot Spectral Imaging," Applied Optics, vol 47, no. 10, 2008.

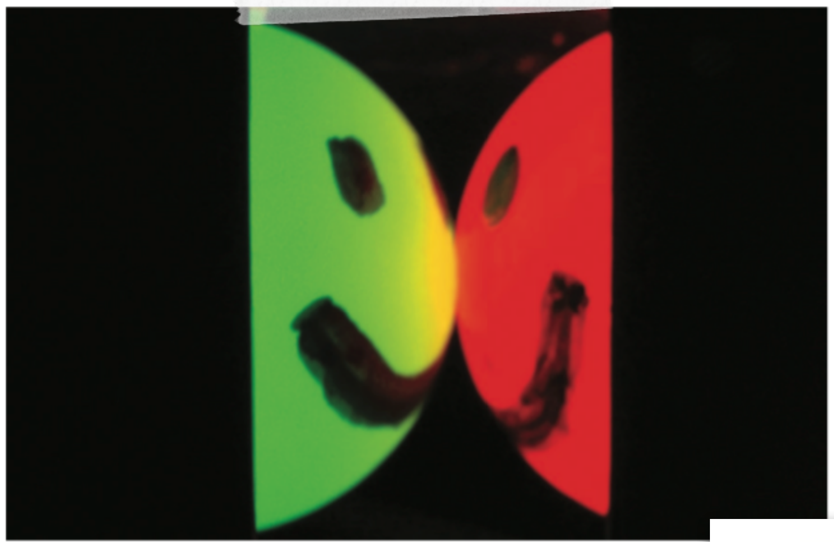

Fig. 4. (Color online) Scene consisting of a Ping-Pong ... nated by a 543 nm green laser and a white light sour... a 560 nm narrowband filter (left), and a red Ping-P... minated by a white light source (right).

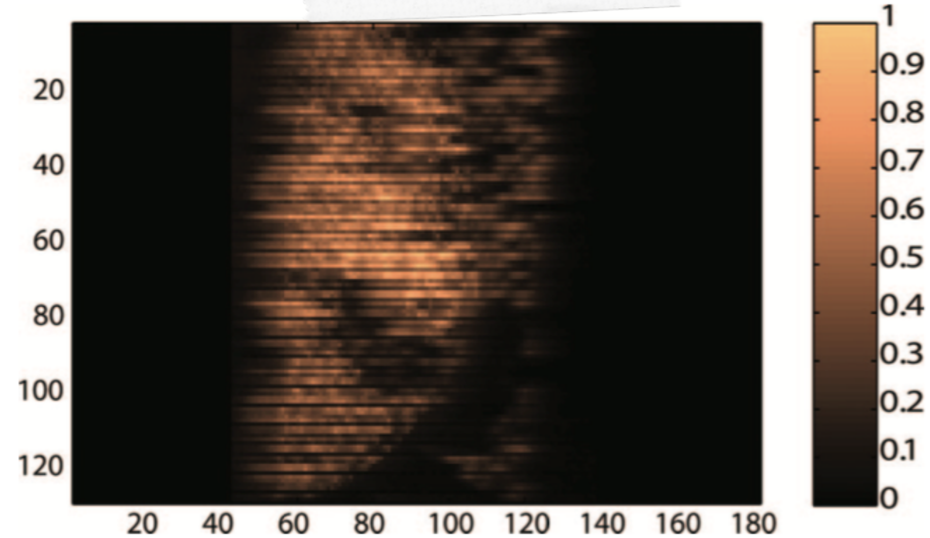

Fig. 5. (Color online) Detector measurement of the scene consisting of the two Ping-Pong balls. Given the low linear dispersion of the prism, there is spatiospectral overlap of the aperture code-modulated images of each ball.

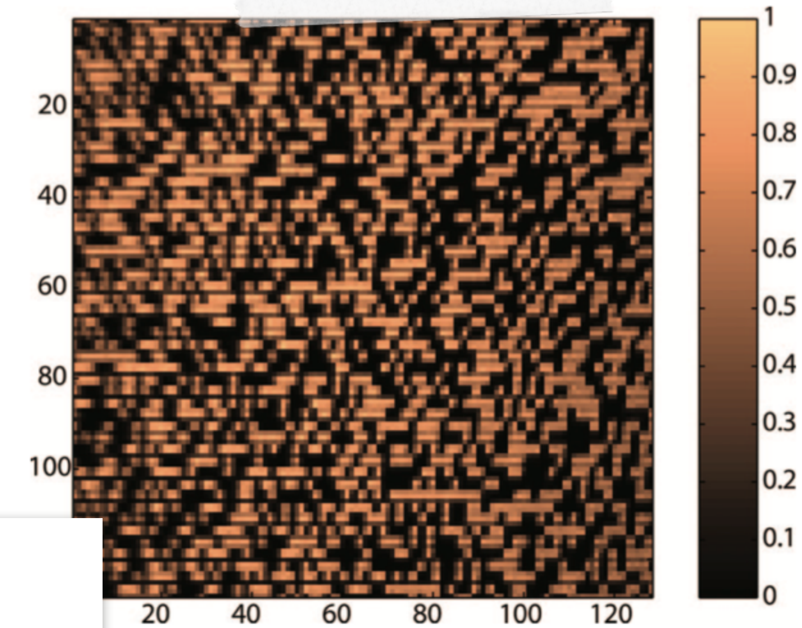

...Color online) Aperture code pattern used by the recon... algorithm to generate an estimate of the data cube.

Figures from paper.

Sampling Design Examples

Coded Aperture Snapshot Spectral Imaging (CASSI)

A. Wagadarikar et al, “Single Disperser Design for Coded Aperture Snapshot Spectral Imaging,” Applied Optics, vol 47, no. 10, 2008.

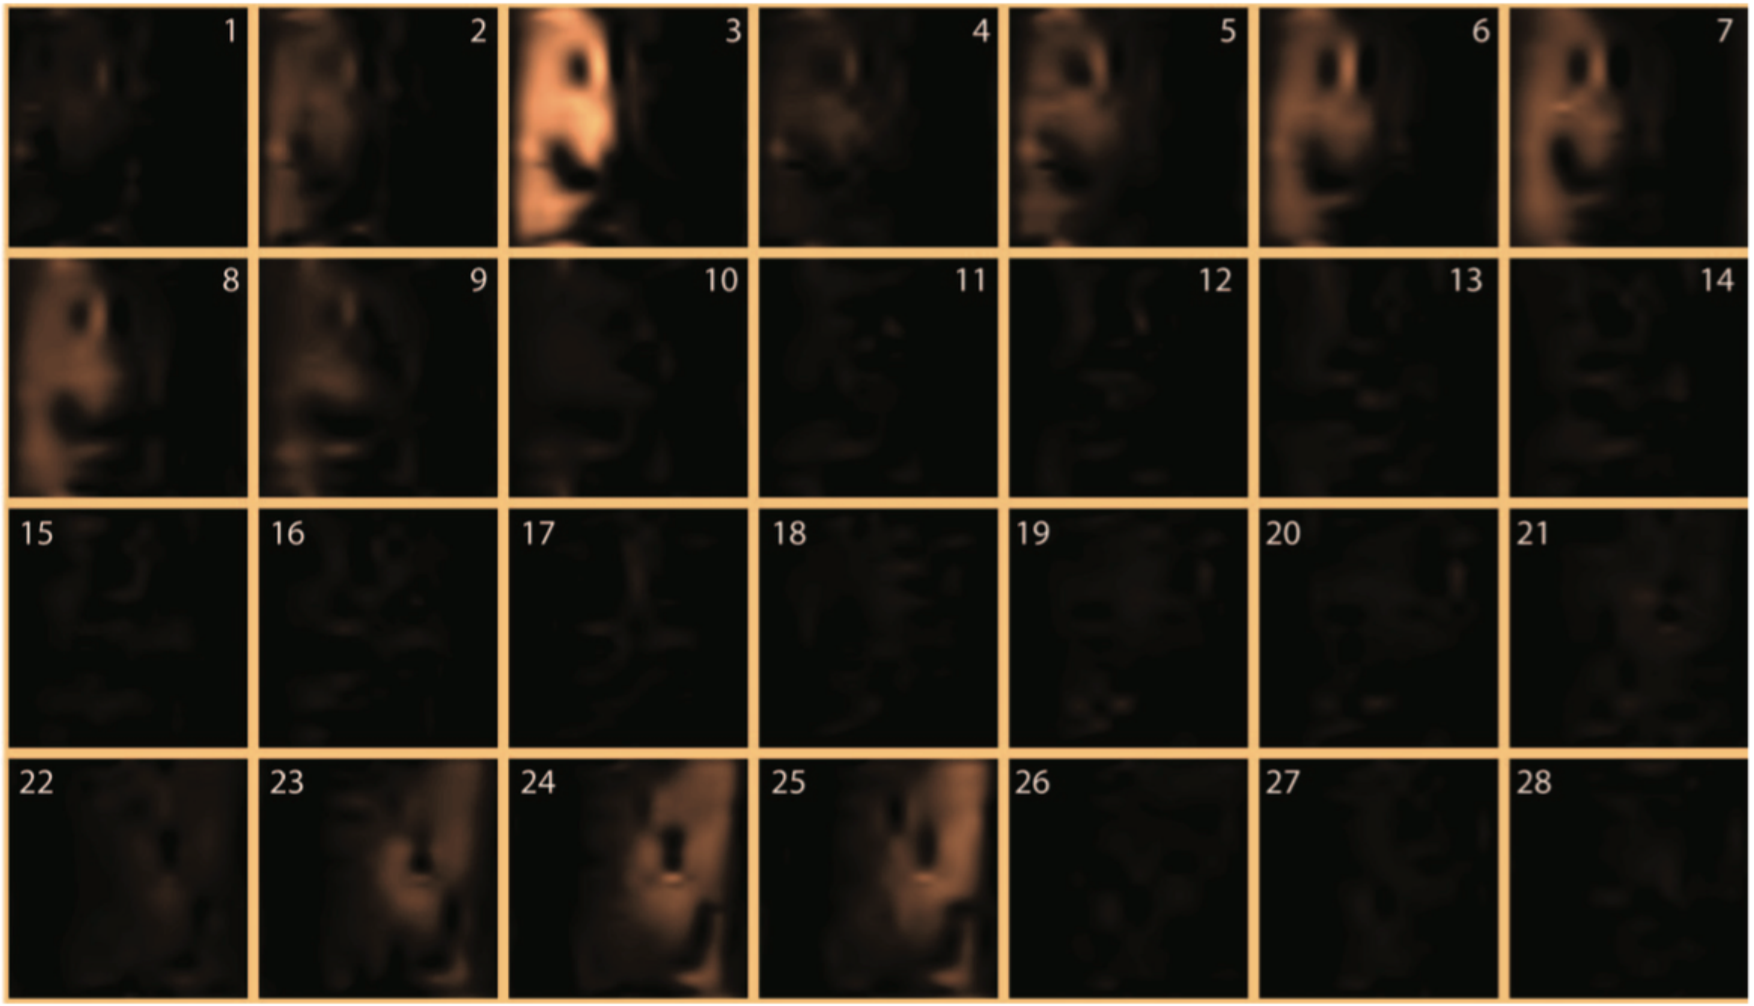

Fig. 6. (Color online) Spatial content of the scene in each of 28 spectral channels between 540 and 640 nm. The green ball can be seen in channels 3, 4, 5, 6, 7, and 8; the red ball can be seen in channels 23, 24, and 25.

Figures from paper.

Can CS Apply to My Problem?

I. Think About Sampling

- Can your sampling be re-designed to take advantage of randomness in the sampling procedure?
- Do you have a manner of efficiently imposing random projections in analog?
- Does this new procedure require sequential measurements? Is your signal time-varying?
- Does knowledge of **F** require careful calibration?

Can CS Apply to My Problem?

ENS

II. Think About Reconstruction

- Is your signal sparse in the ambient domain?
- If not, does there exist a sparse basis for which it is?
- If not, do you have enough data to infer one? (*Dictionary Learning*)
- Is the support of your signal correlated?
  - *E.g. wavelet-trees, etc.*
- What reconstruction methods are best suited for your signal dimensionality?
  - *Trade-off in accuracy and efficiency…*
- Is your noise Gaussian? If not, does a reconstruction method exist for your noise model?

SPHINX @ENS

Statistical **PH**ysics of **IN**formation e**X**traction

«OU»

Statistical **PH**ysics of **IN**verse comple**X** sysems

ENS

1794

Questions?

Merci!