Introduction to Compressive Sensing

Aswin Sankaranarayanan

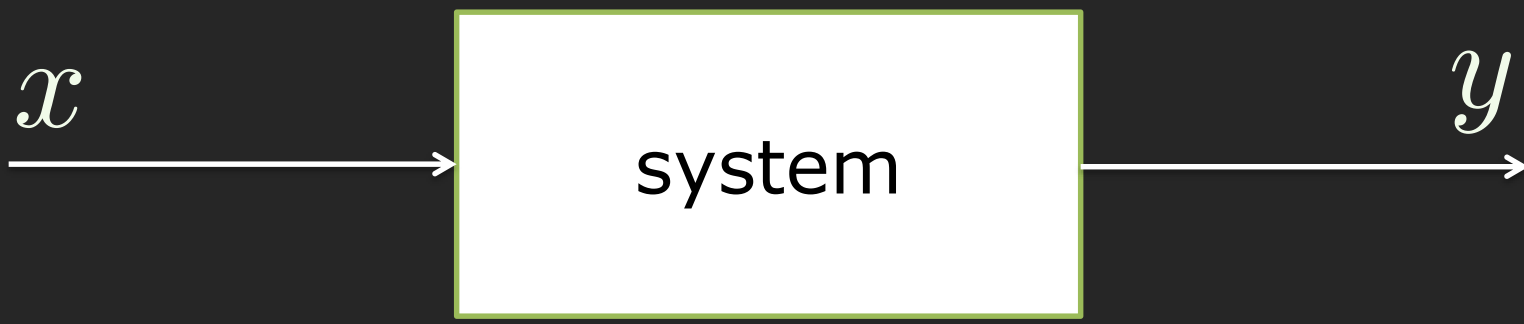

$$y = \Phi x$$

Is this system linear ?

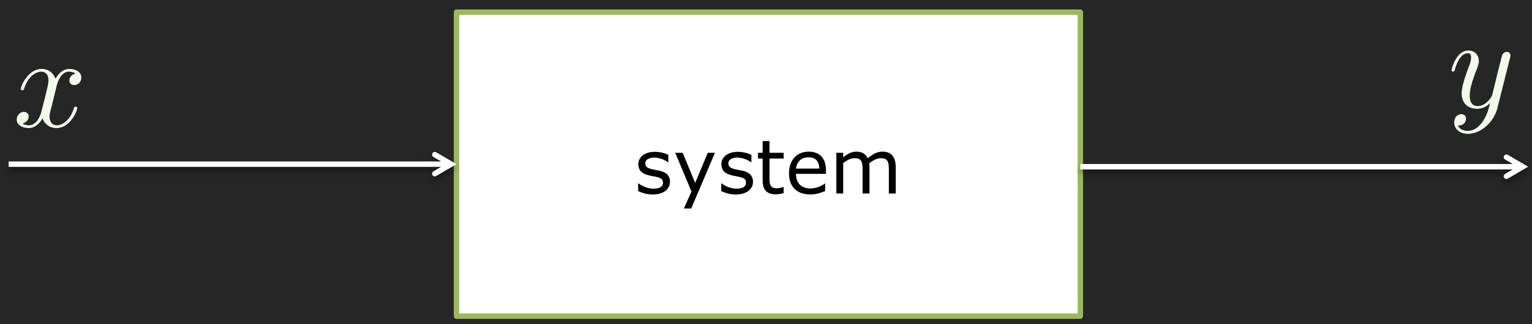

$$y = \Phi x$$

Is this system linear ?

Given *y*, can we recovery *x* ?

Under-determined problems

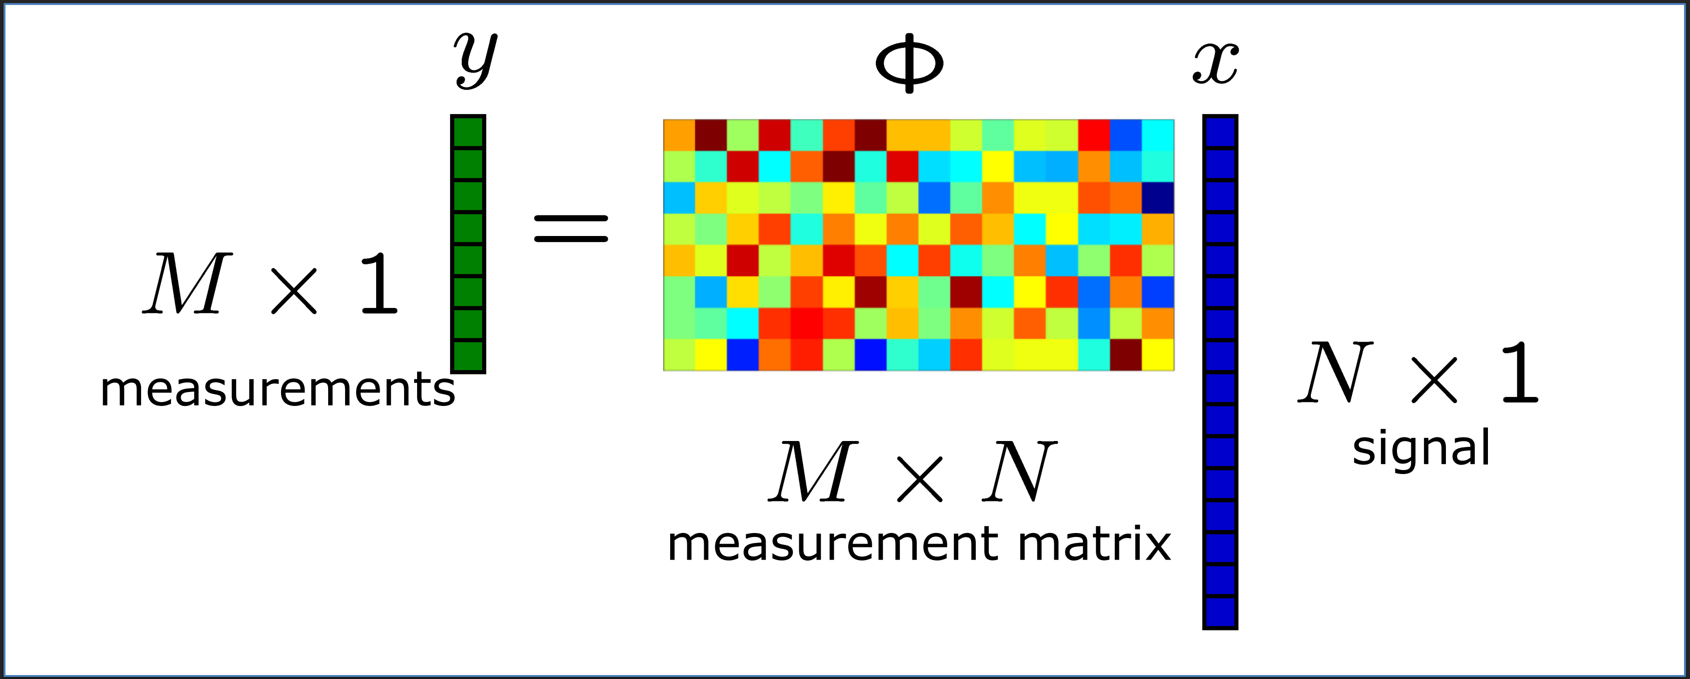

If M < N, then the system is information lossy

Image credit
Graeme Pope

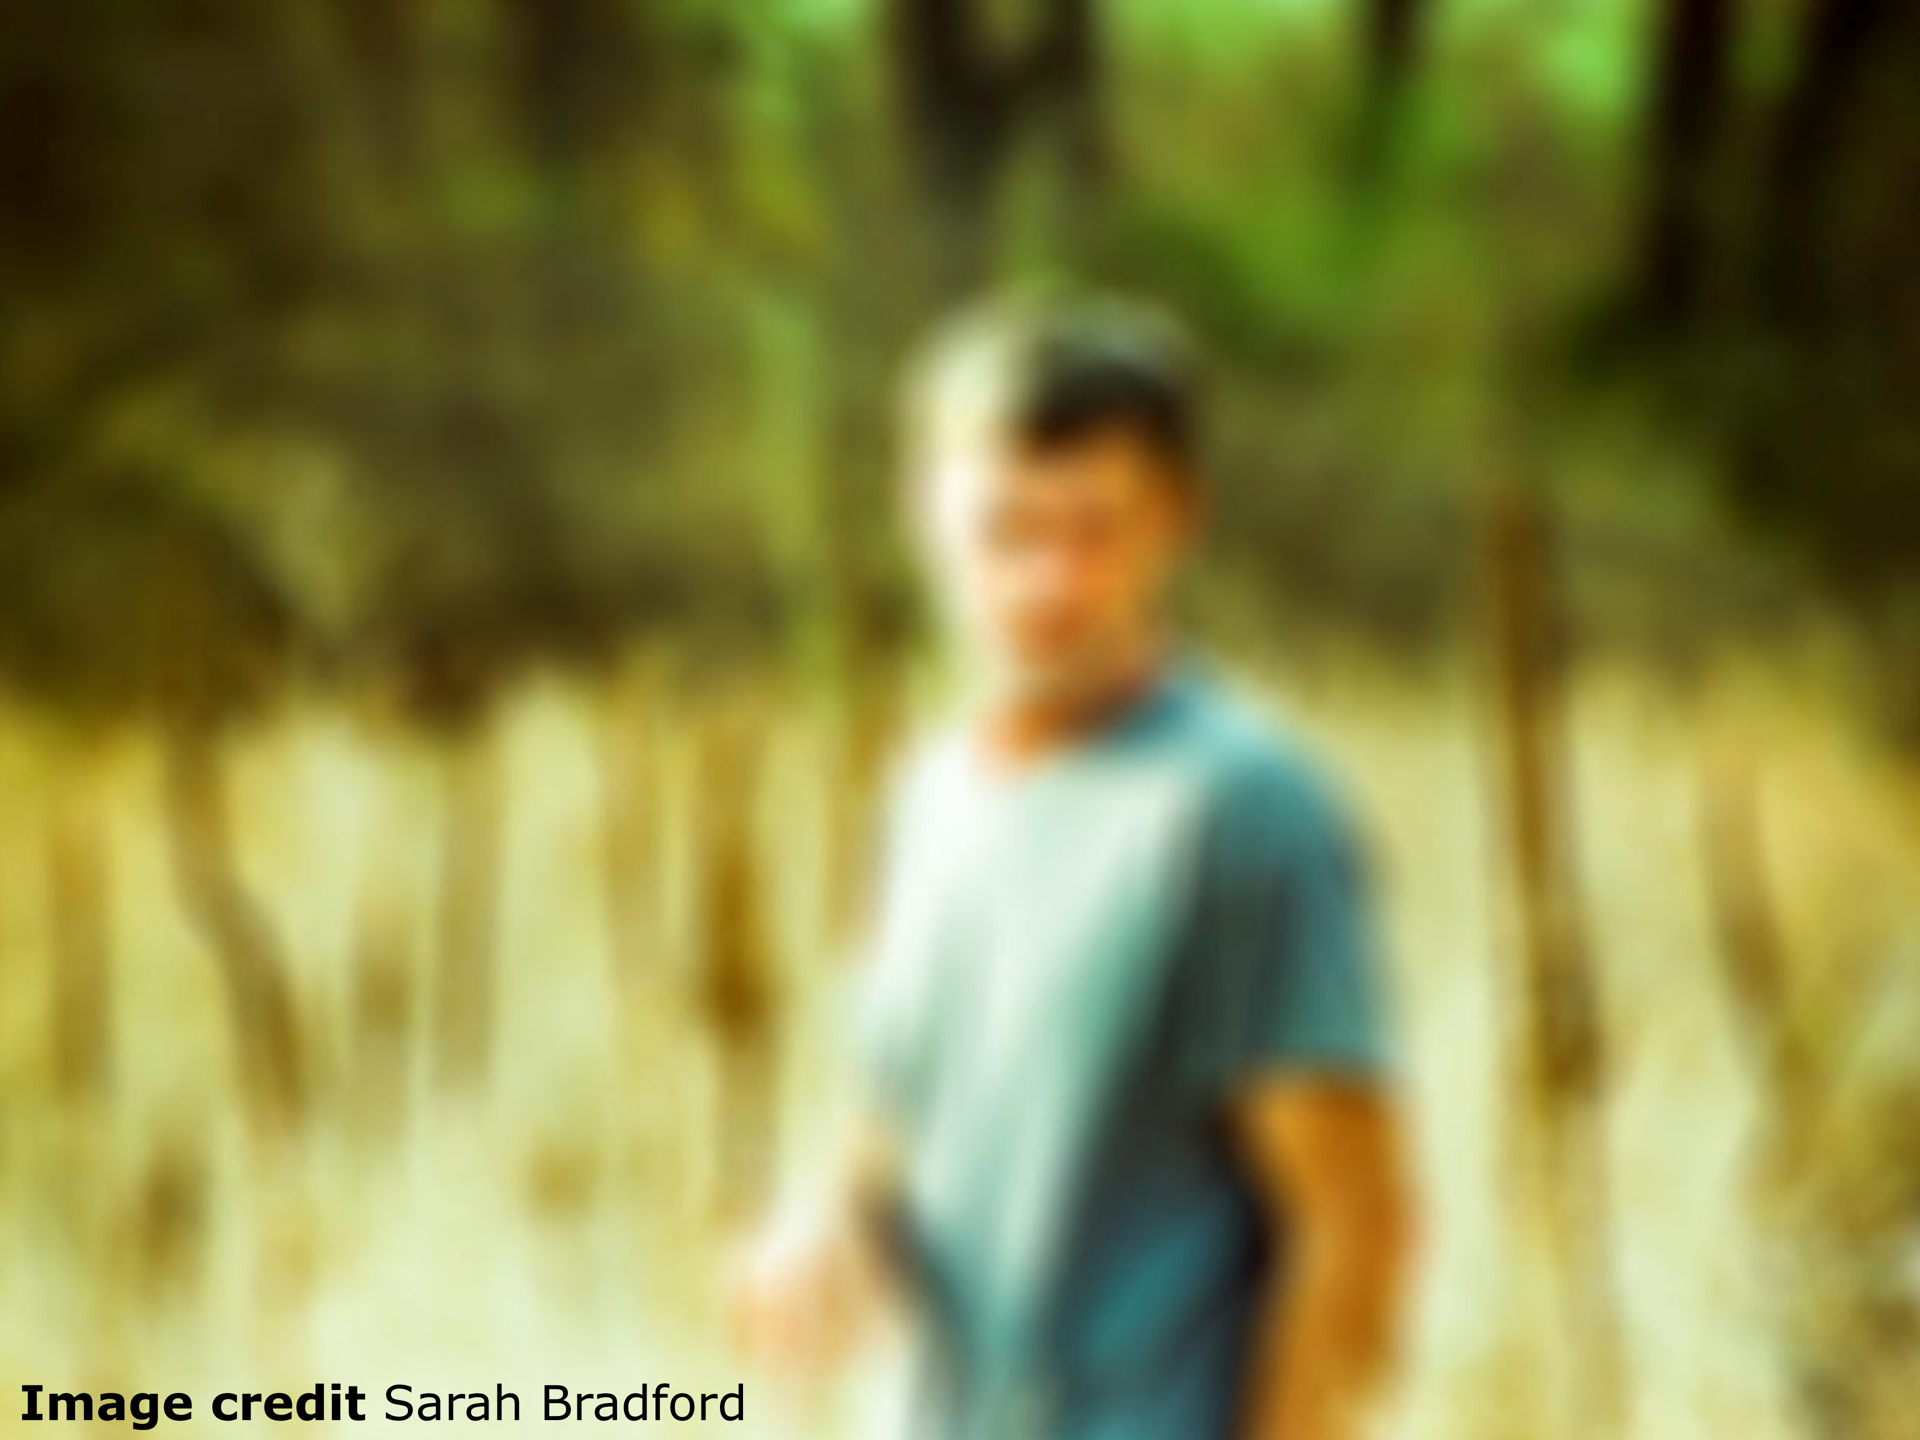

Image credit Sarah Bradford

Super-resolution

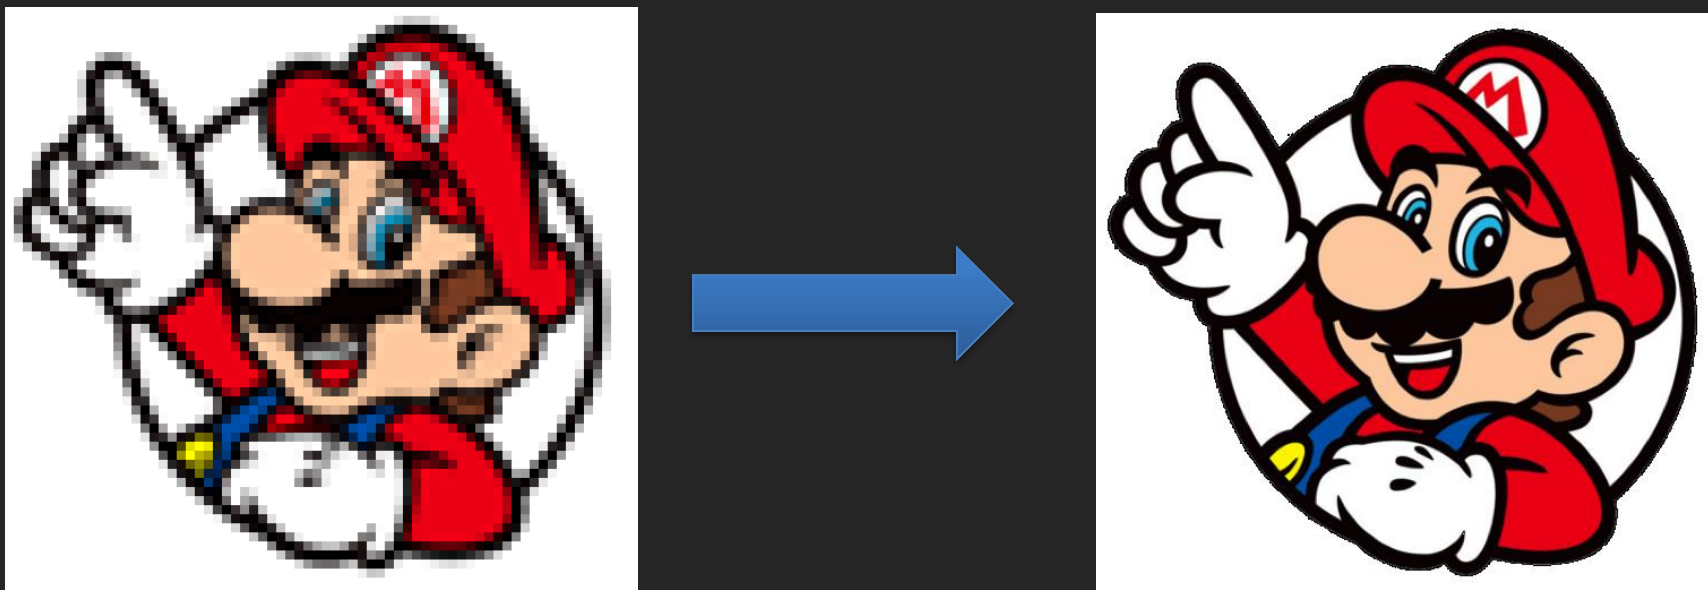

Can we increase the resolution of this image ?

[(Link: Depixelizing pixel art)]

Under-determined problems

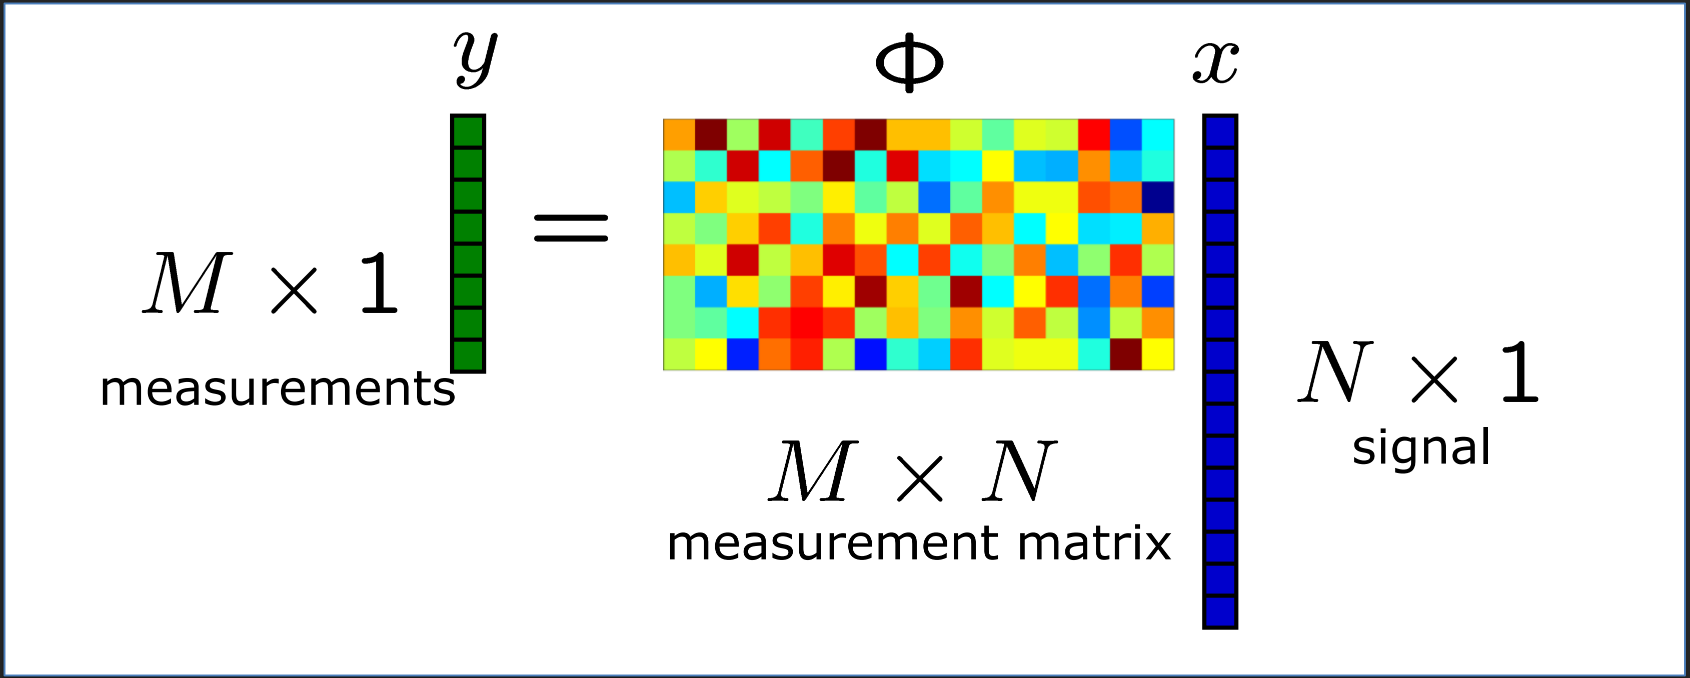

Fewer knowns than unknowns!

Under-determined problems

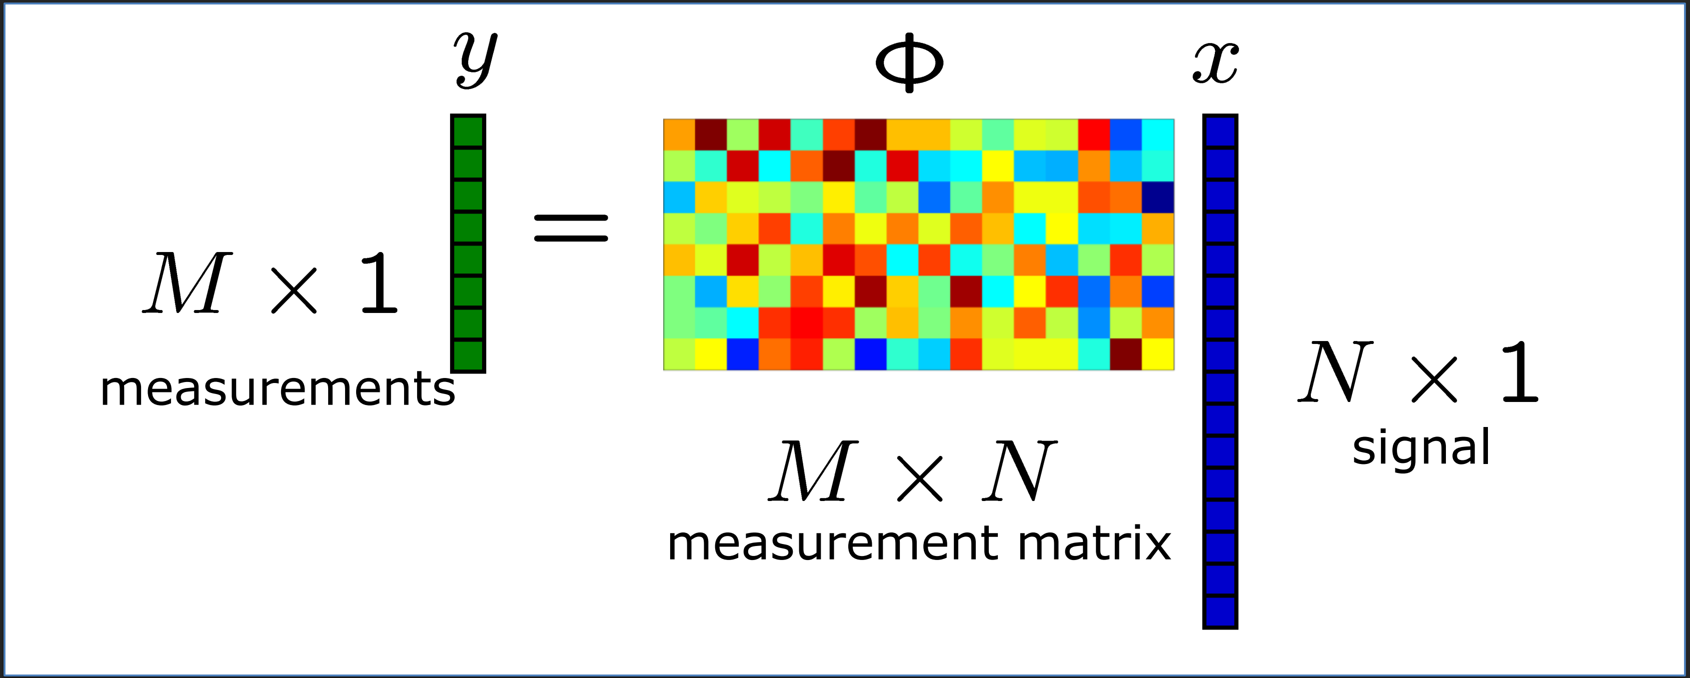

Fewer knowns than unknowns!

An infinite number of solutions to such problems

Credit: Rob Fergus and Antonio Torralba

Credit: Rob Fergus and Antonio Torralba

Is there anything we can do about this ?

Complete the sentences

I cnt blv I m bl t rd ths sntnc.

Wntr s cmng, n .. Wntr s hr

Hy, I m slvng n ndr-dtrmnd lnr systm.

how: ?

Complete the matrix

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

how: ?

Complete the image

Model ?

Dictionary of visual words

I cnt blv I m bl t rd ths sntnc.

Shrlck s th vc f th drgn

Hy, I m slvng n ndr-dtrmnd lnr systm.

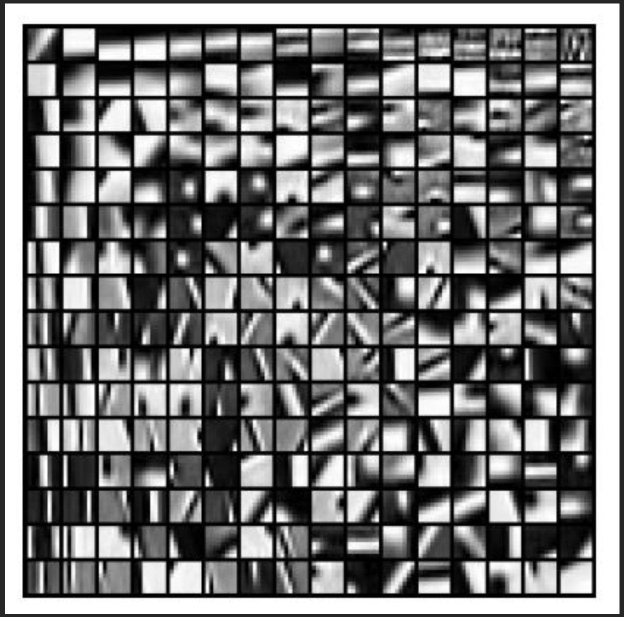

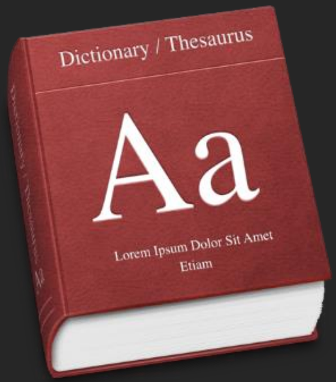

Dictionary of visual words

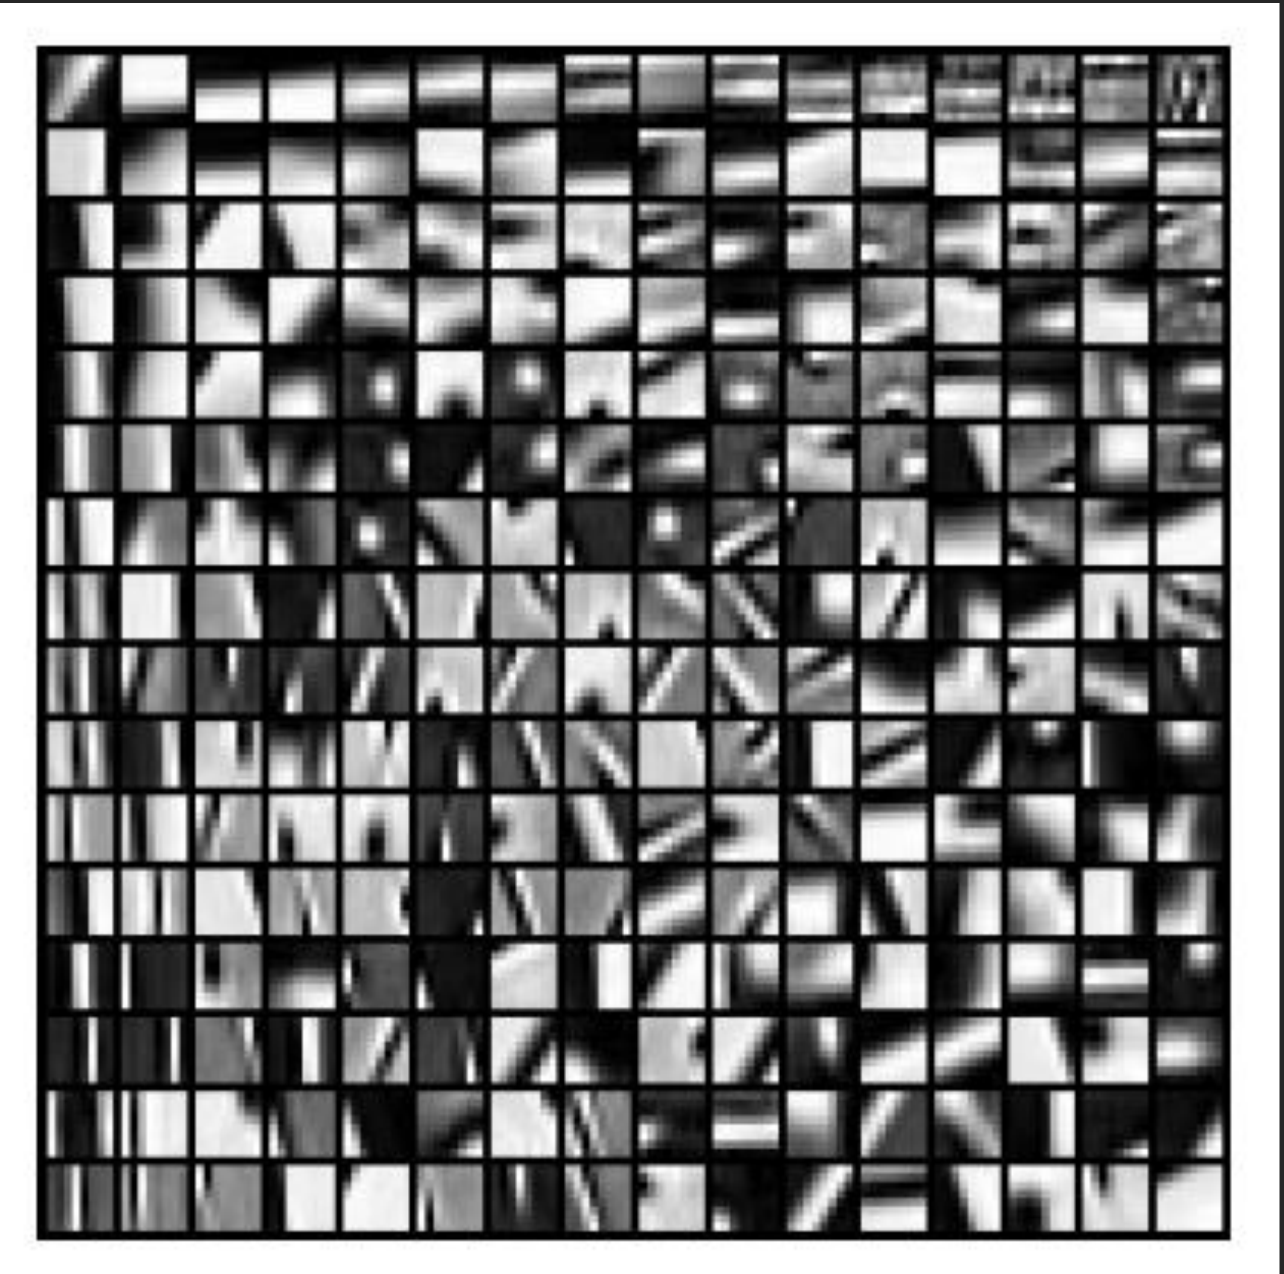


Image credit
Graeme Pope


Image credit
Graeme Pope

Result
Studer, Baraniuk, ACHA 2012

Compressive Sensing

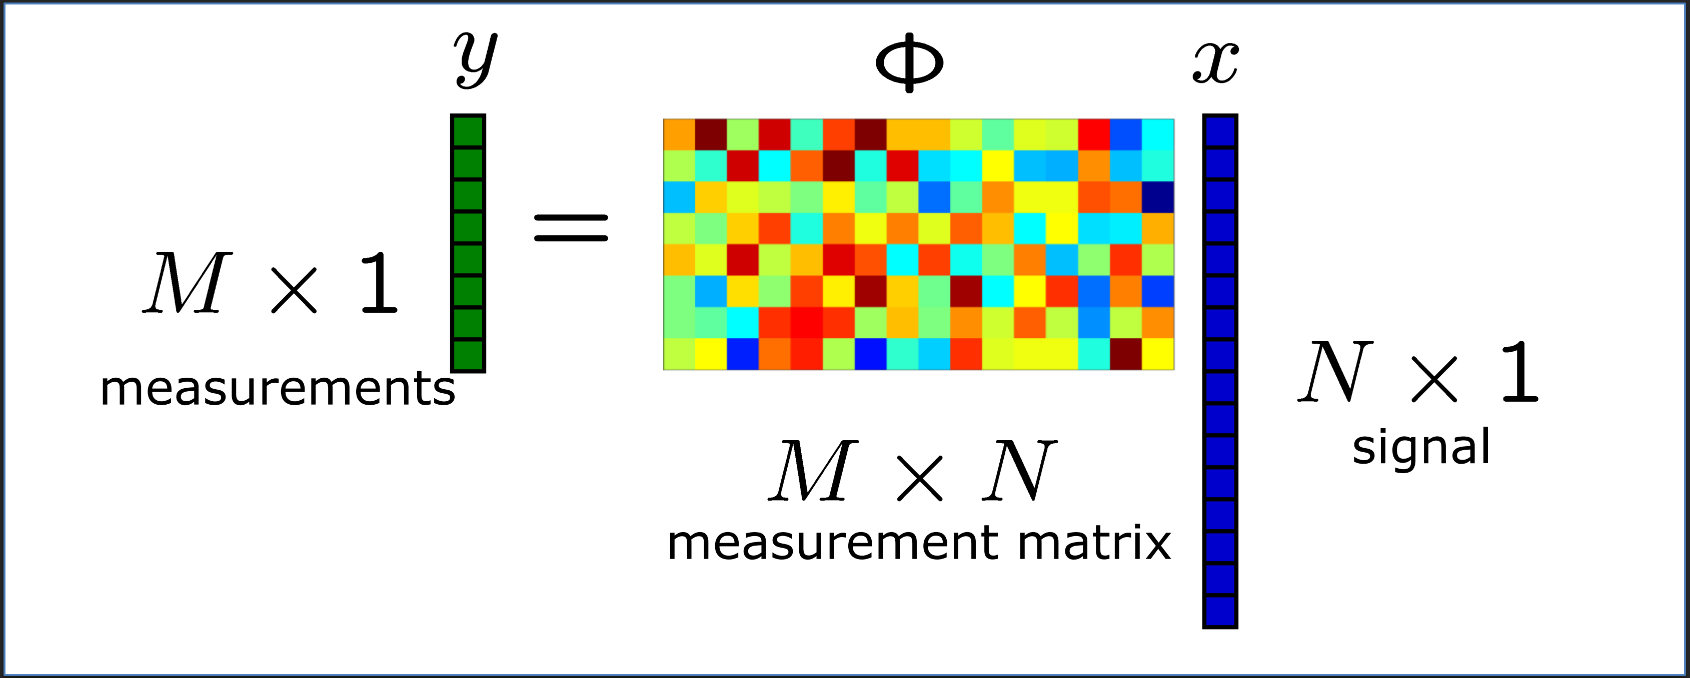

A toolset to solve under-determined systems by exploiting additional structure/models on the signal we are trying to recover.

modern sensors are linear systems!!!

Sampling

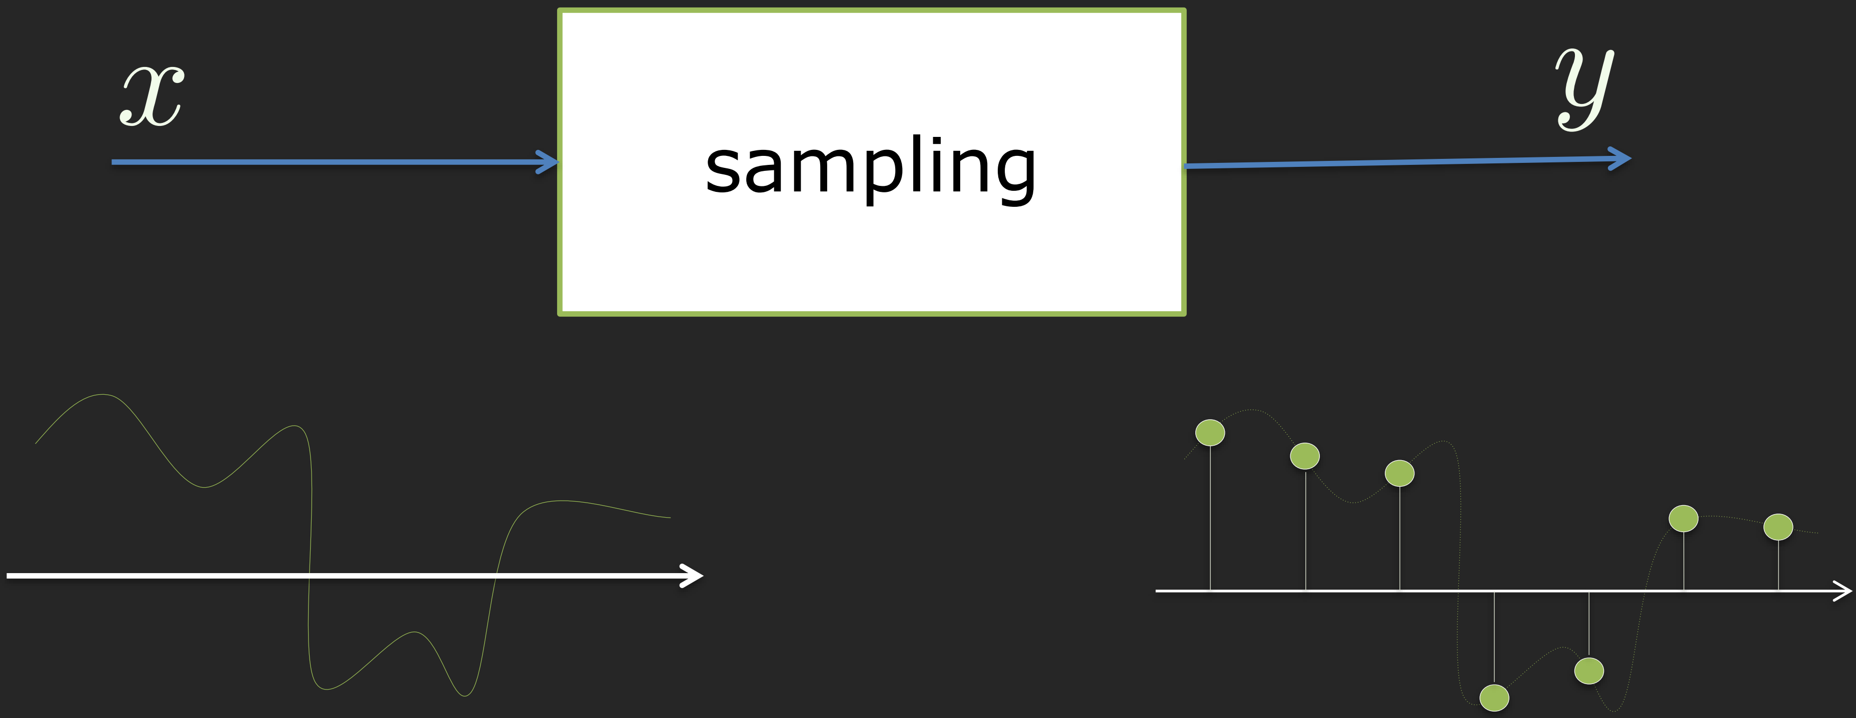

Can we recover the analog signal from its discrete time samples ?

Nyquist Theorem

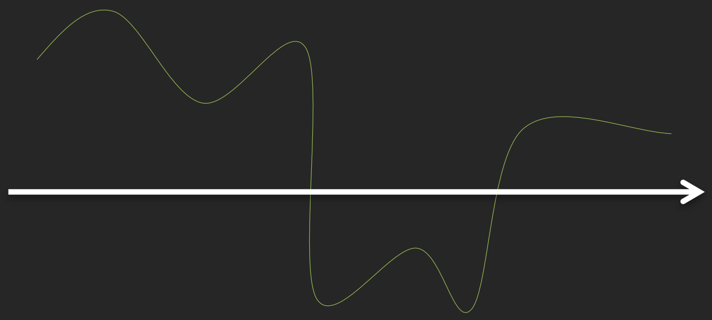

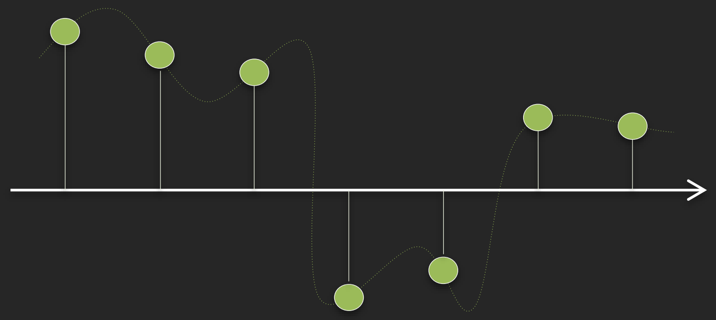

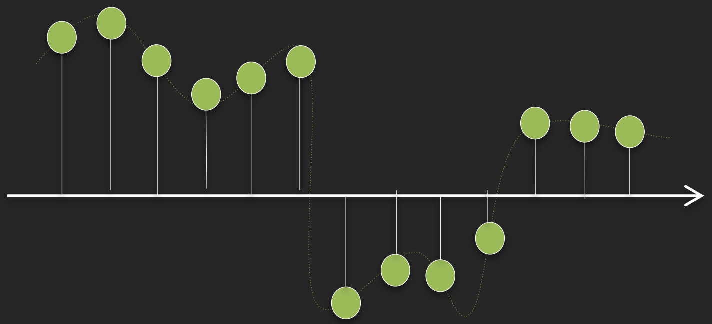

An analog signal can be reconstructed perfectly from discrete samples *provided you sample it densely*.

The Nyquist Recipe

sample faster

sample denser

the more you sample,
the more detail is preserved

The Nyquist Recipe

sample faster

sample denser

the more you sample,
the more detail is preserved

But what happens if you do not follow the Nyquist recipe ?

Credit: Rob Fergus and Antonio Torralba

Image credit: Boston.com

The Nyquist Recipe

sample faster

sample denser

the more you sample,
the more detail is preserved

But what happens if you do not follow the Nyquist recipe ?

What you must learn is that these rules are no different than the rules of a computer system. Some of them can be bent. Others can be broken.

Breaking resolution barriers

- Observing a 2000 fps spinning tool with a 25 fps camera

Normal Video: 25fps

Compressively obtained video: 25fps

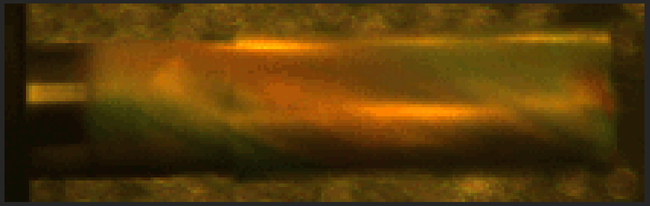

Recovered Video: 2000fps

Slide/Image credit: Reddy et al. 2011

Compressive Sensing

Use of **motion flow-models** in the context of compressive video recovery

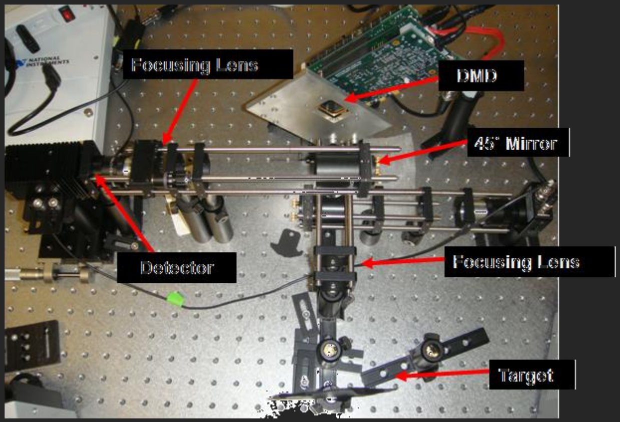

single pixel camera

128x128 images sensed at 61x comp.

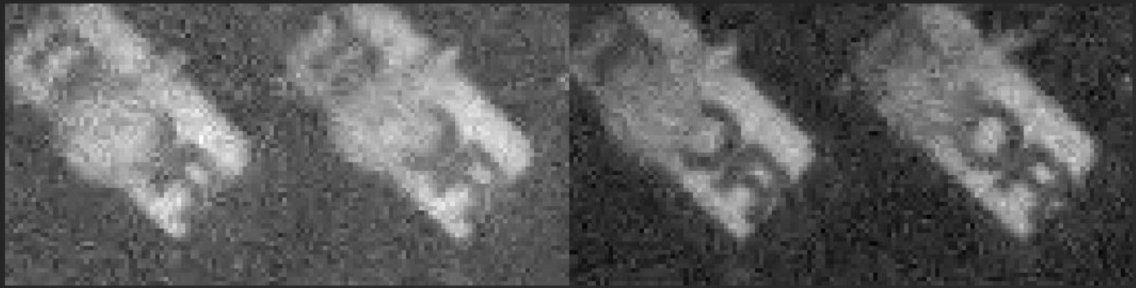

Naïve frame-to-frame recovery

CS-MUVI at 61x compression

Sankaranarayanan et al. ICCP 2012, SIAM J. Imaging Sciences, 2015*

Compressive Imaging Architectures

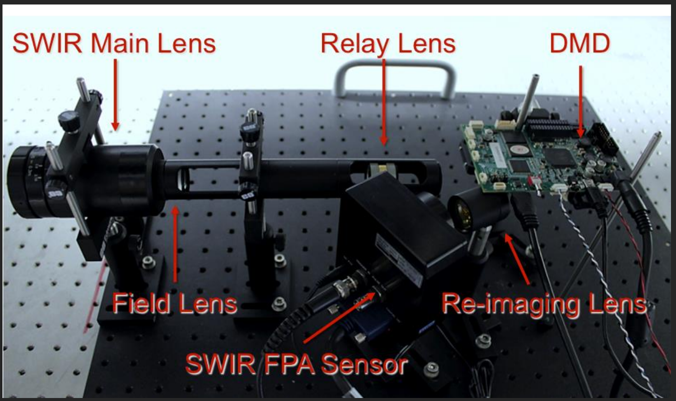

Scalable imaging architectures that deliver videos at **mega-pixel resolutions** in infrared

visible image

SWIR image

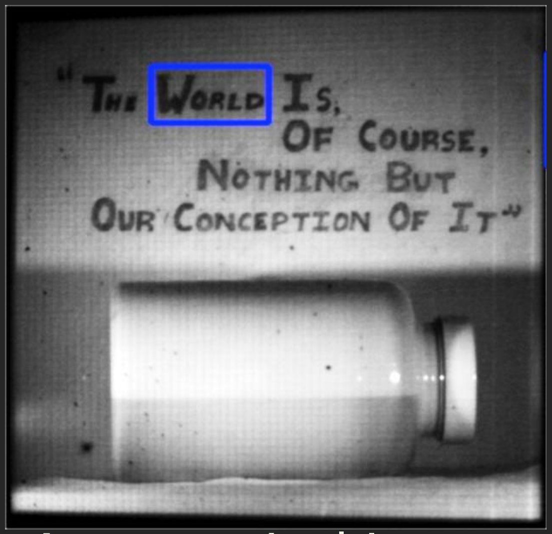

A mega-pixel image obtained from a 64x64 pixel array sensor

Chen et al. CVPR 2015, Wang et al. ICCP 2015

Advances in Compressive Imaging

Linear Inverse Problems

- Many classic problems in computer can be posed as linear inverse problems

- Notation
  - **Signal** of interest $x \in \mathbb{R}^N$
  - **Observations** $y \in \mathbb{R}^M$
  - Measurement model $y = \Phi x + e$

measurement matrix

measurement noise

- Problem definition: given y, recover x

Linear Inverse Problems

$$y = \Phi x$$

$M \times 1$
measurements

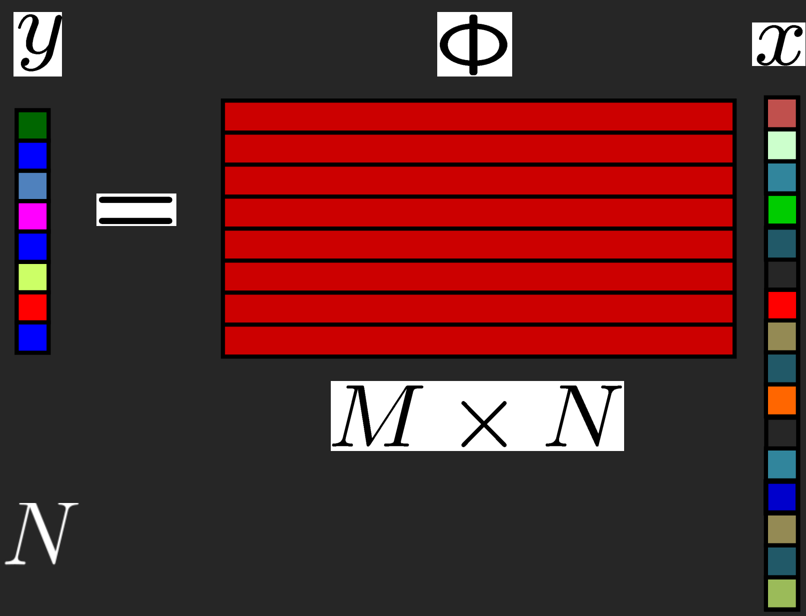

$N \times 1$
signal

$$M < N$$

Measurement matrix has a *(N-M)* dimensional **null-space**

Solution is no longer **unique**

Sparse Signals

$$y = \Phi x$$

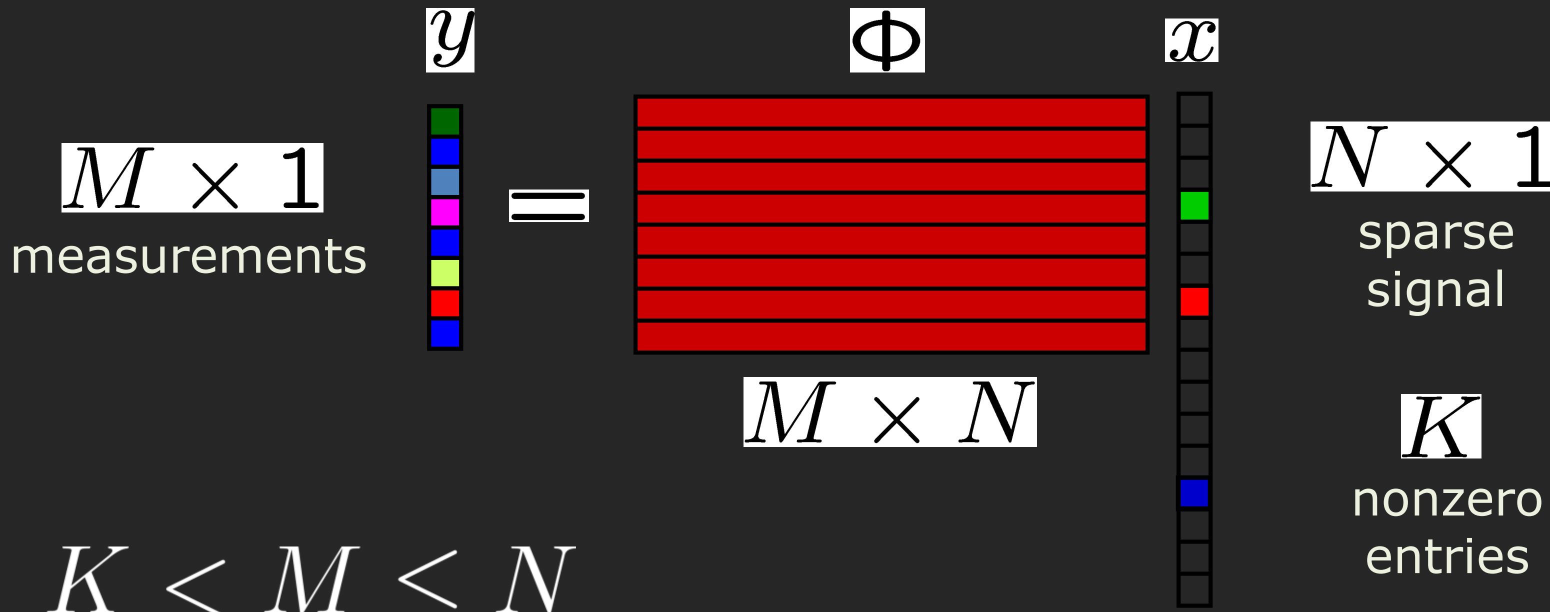

$$K < M \leq N$$

How Can It Work?

- Matrix Φ not full rank…

  $M < N$

  … and so loses information in general

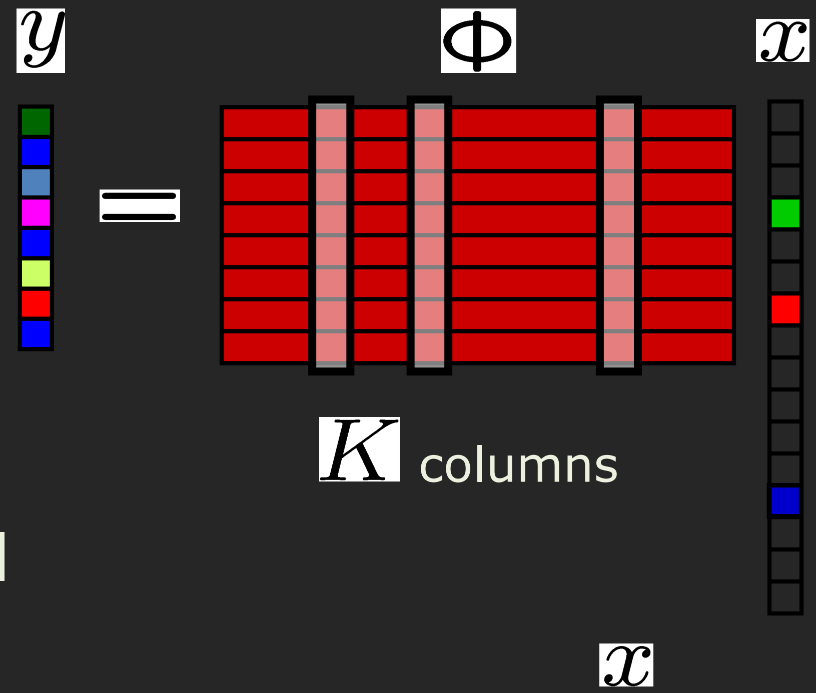

- But we are only interested in ***sparse*** vectors

Restricted Isometry Property (RIP)

- Preserve the structure of sparse/compressible signals

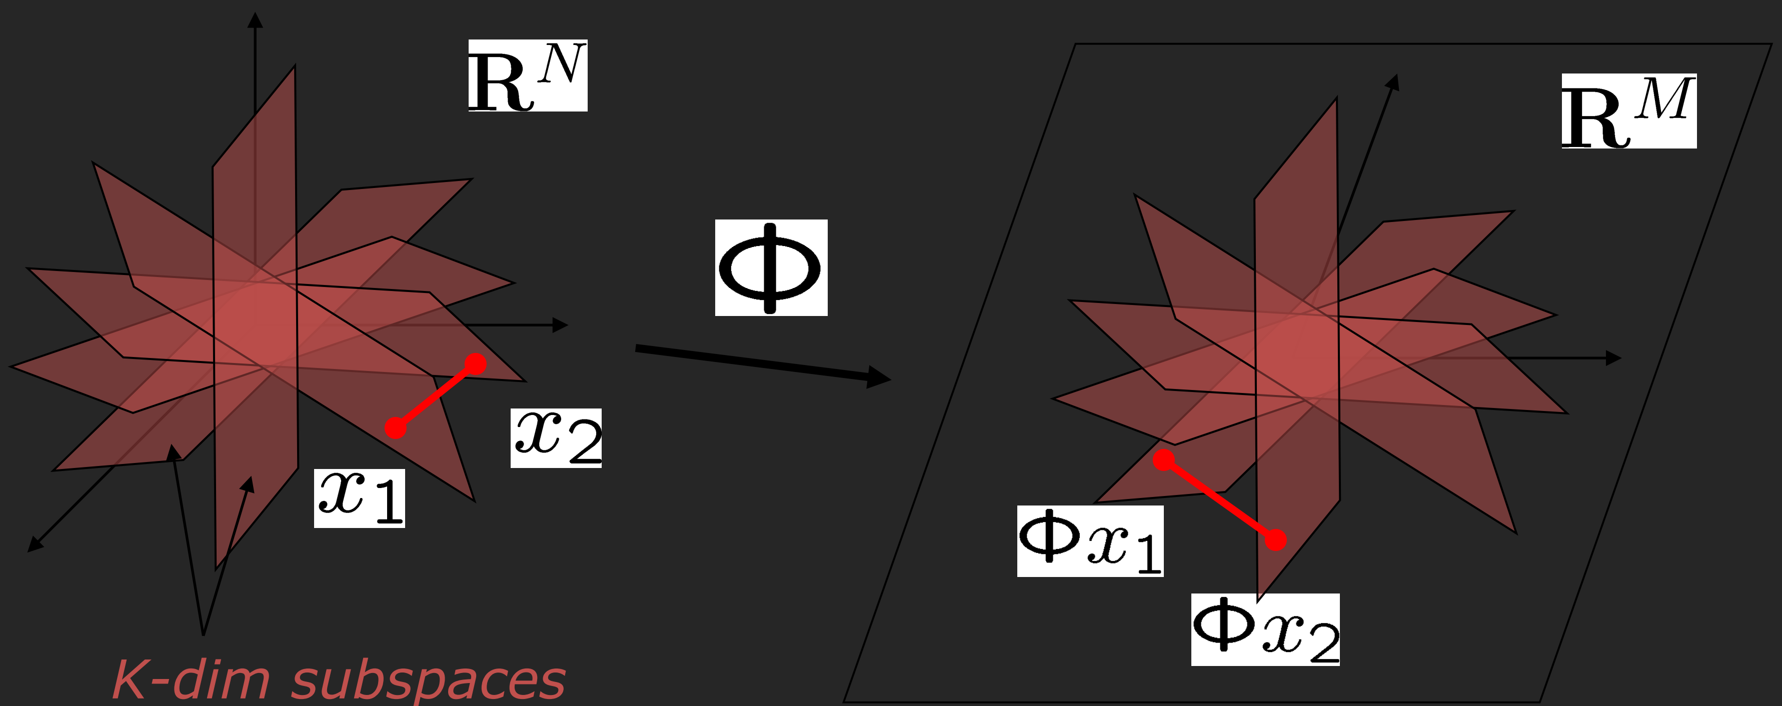

Restricted Isometry Property (RIP)

- RIP of order *2K* implies: for all *K*-sparse x_1 and x_2

$$(1-\delta_{2K}) \leq \frac{\|\Phi x_1 - \Phi x_2\|_2^2}{\|x_1 - x_2\|_2^2} \leq (1+\delta_{2K})$$

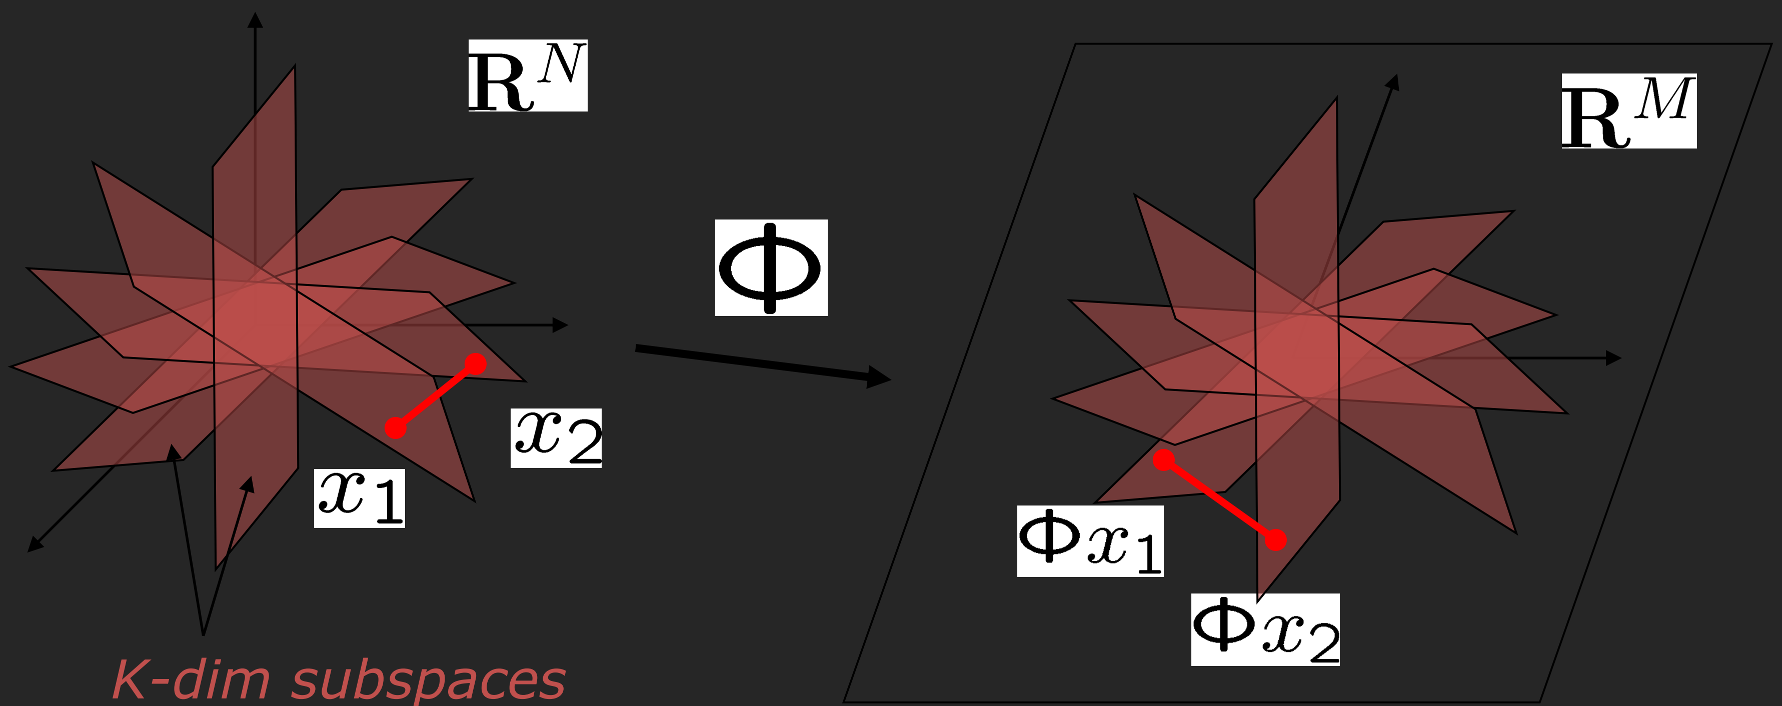

How Can It Work?

- Matrix Φ
  not full rank…

  $M < N$

  … and so
  loses information in general

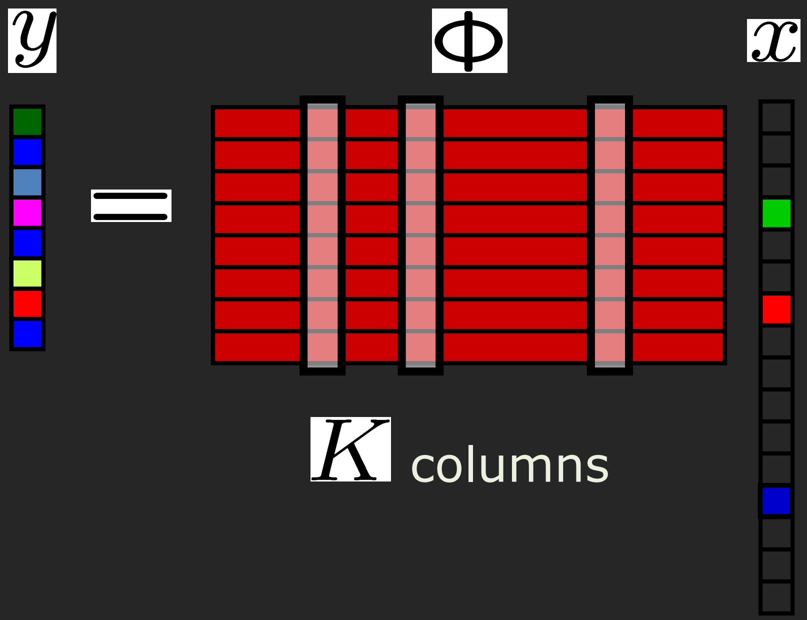

- **Design** Φ so that each of its *M*x2*K* submatrices are full rank (RIP)

How Can It Work?

- Matrix Φ not full rank…

  $M < N$

  … and so loses information in general

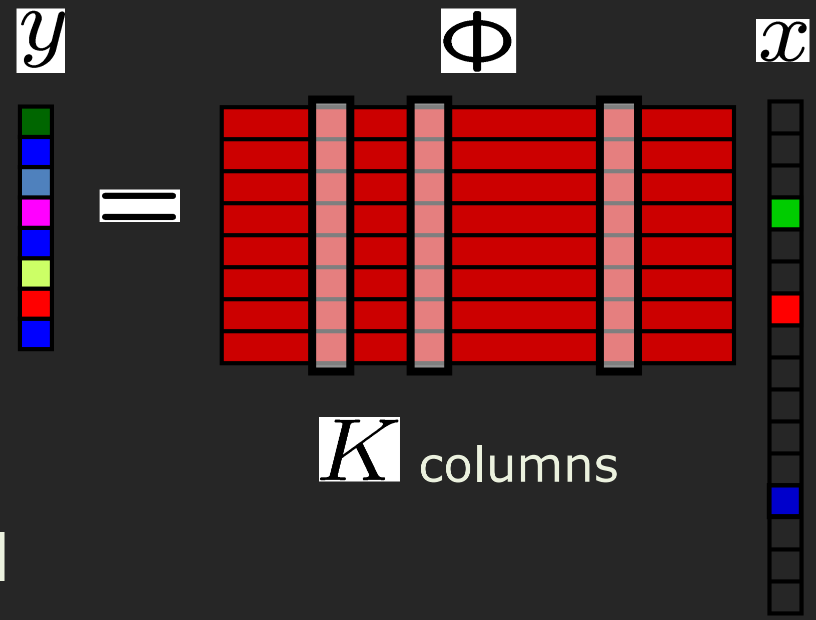

- **Design** Φ so that each of its *M*x2*K* submatrices are full rank (RIP)

- Random measurements provide RIP with

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

CS Signal Recovery

- Random projection Φ not full rank
- Recovery problem: given $y = \Phi x$ find x
- **Null space**
- Search in null space for the "sparsest" x

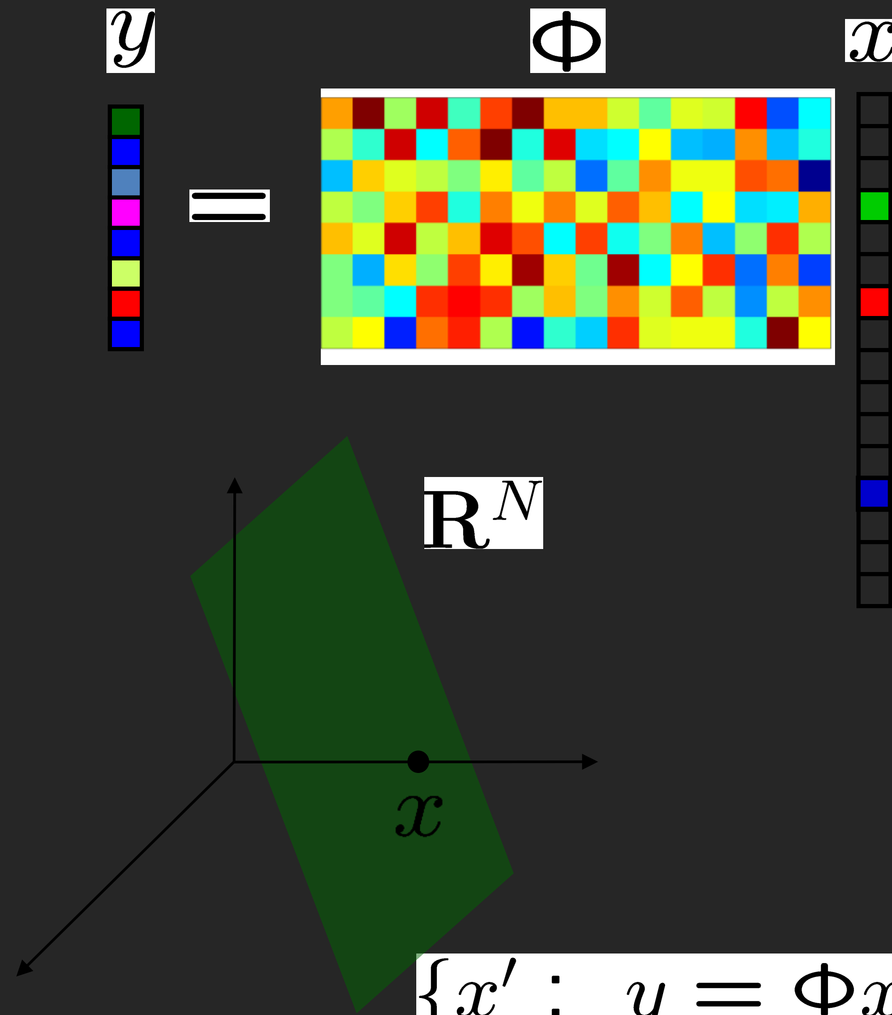

$\{x' : \; y = \Phi x'\}$

(N-M)-dim hyperplane
*at **random angle***

ℓ_1 Signal Recovery

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

- Optimization: $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_1$

- **Convexify** the ℓ_0 optimization

Candes

Romberg

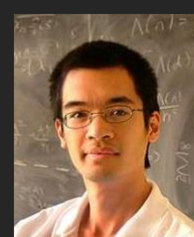

Tao

Donoho

ℓ_1 Signal Recovery

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

- Optimization: $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_1$

- **Convexify** the ℓ_0 optimization

- **Polynomial time** alg
  (linear programming)

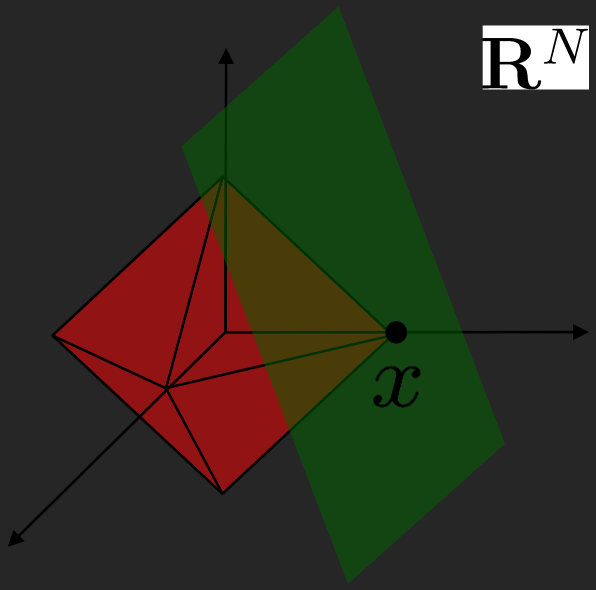

Compressive Sensing

Let. $y = \Phi x_0 + e$

$$\widehat{x} = \arg\min_x \|x\|_1 \quad s.t. \quad \|y - \Phi x\|_2 \leq \|e\|$$

If Φ satisfies RIP with $\delta_{2K} \leq \sqrt{2} - 1$,

Then

$$\|\hat{x} - x_0\|_1 \leq C_1 \|e\|_2 + C_2 \|x_0 - x_{0,K}\|_2 / \sqrt{K}$$

Best K-sparse approximation