1-bit Compressive Sensing

Petros Boufounos
Richard Baraniuk

RICE

Compressive Sensing

1. Signal Model
2. Random linear measurements
3. Non-linear reconstruction

Signal Models

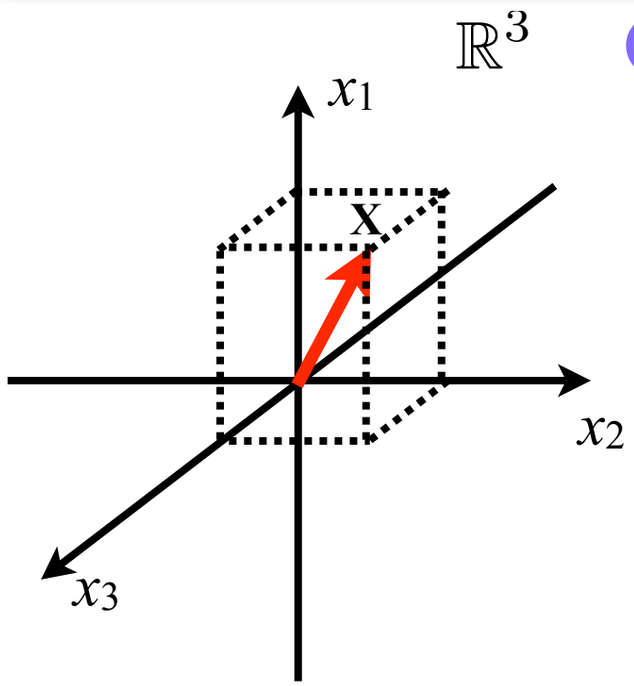

Classical Model: Signal lies in a linear vector space (e.g. bandlimited functions)

Sparse Model: Signals of interest are often sparse or compressible

	Signal	Transform	
Image		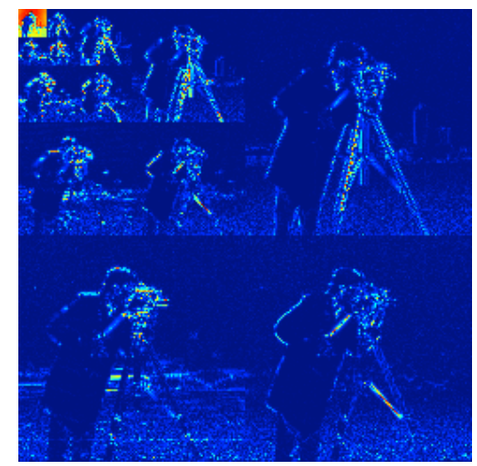	Wavelet
Bat Sonar Chirp	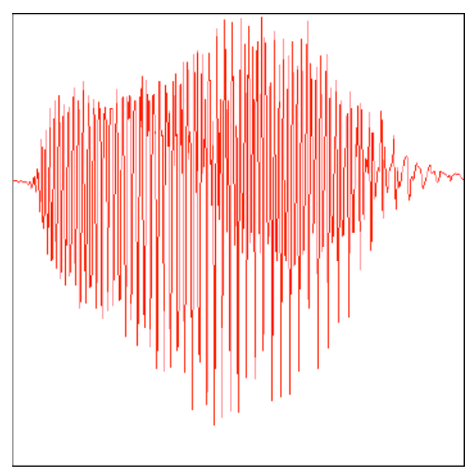	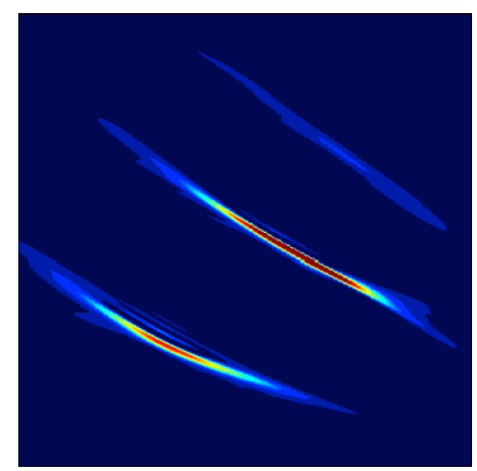	Gabor/ STFT

i.e., very few large coefficients, many close to zero.

Sparse Signal Models

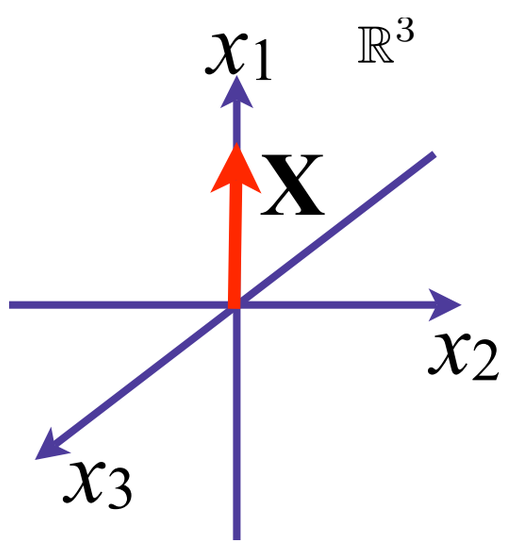

1-sparse

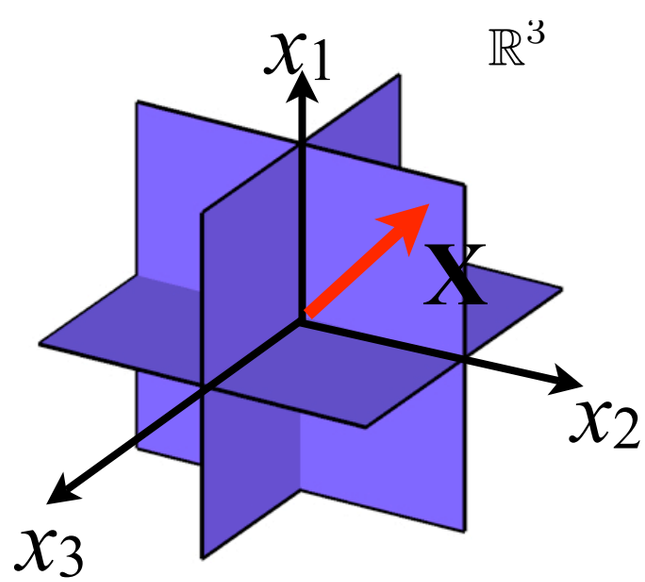

2-sparse

Sparse signals have few non-zero coefficients.

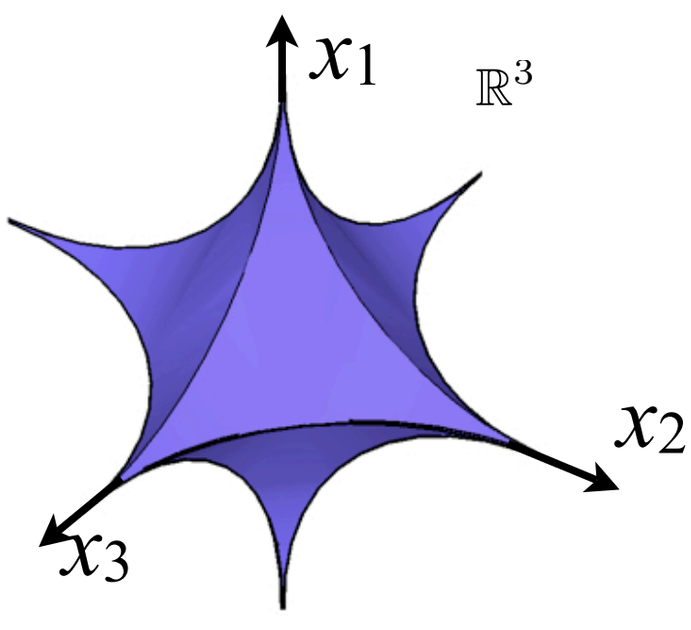

Compressible (ℓ_p ball, $p<1$)

Compressible signals have few significant coefficients. The coefficients decay as a power law.

Compressive Sensing in a Nutshell

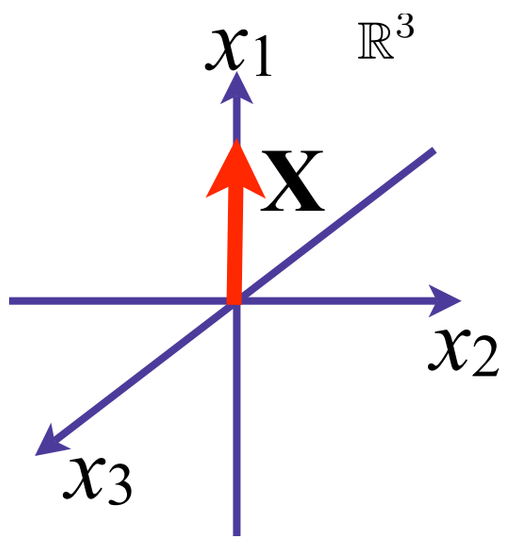

1-sparse

If a signal is sparse, do not waste effort sampling the empty space.

Instead, use fewer samples and allow ambiguity.

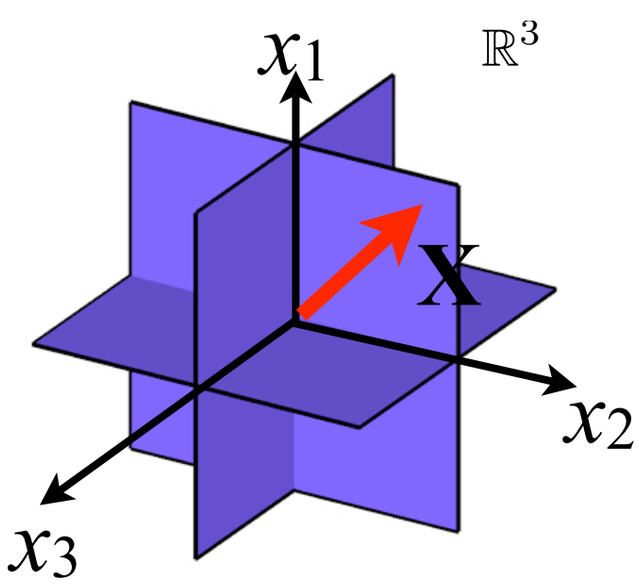

2-sparse

Use the sparsity model to reconstruct and uniquely resolve the ambiguity.

Compressive Measurements

$$\mathbf{y} = \Phi\mathbf{x}$$
$$y_i = \langle \phi_i, \mathbf{x} \rangle$$

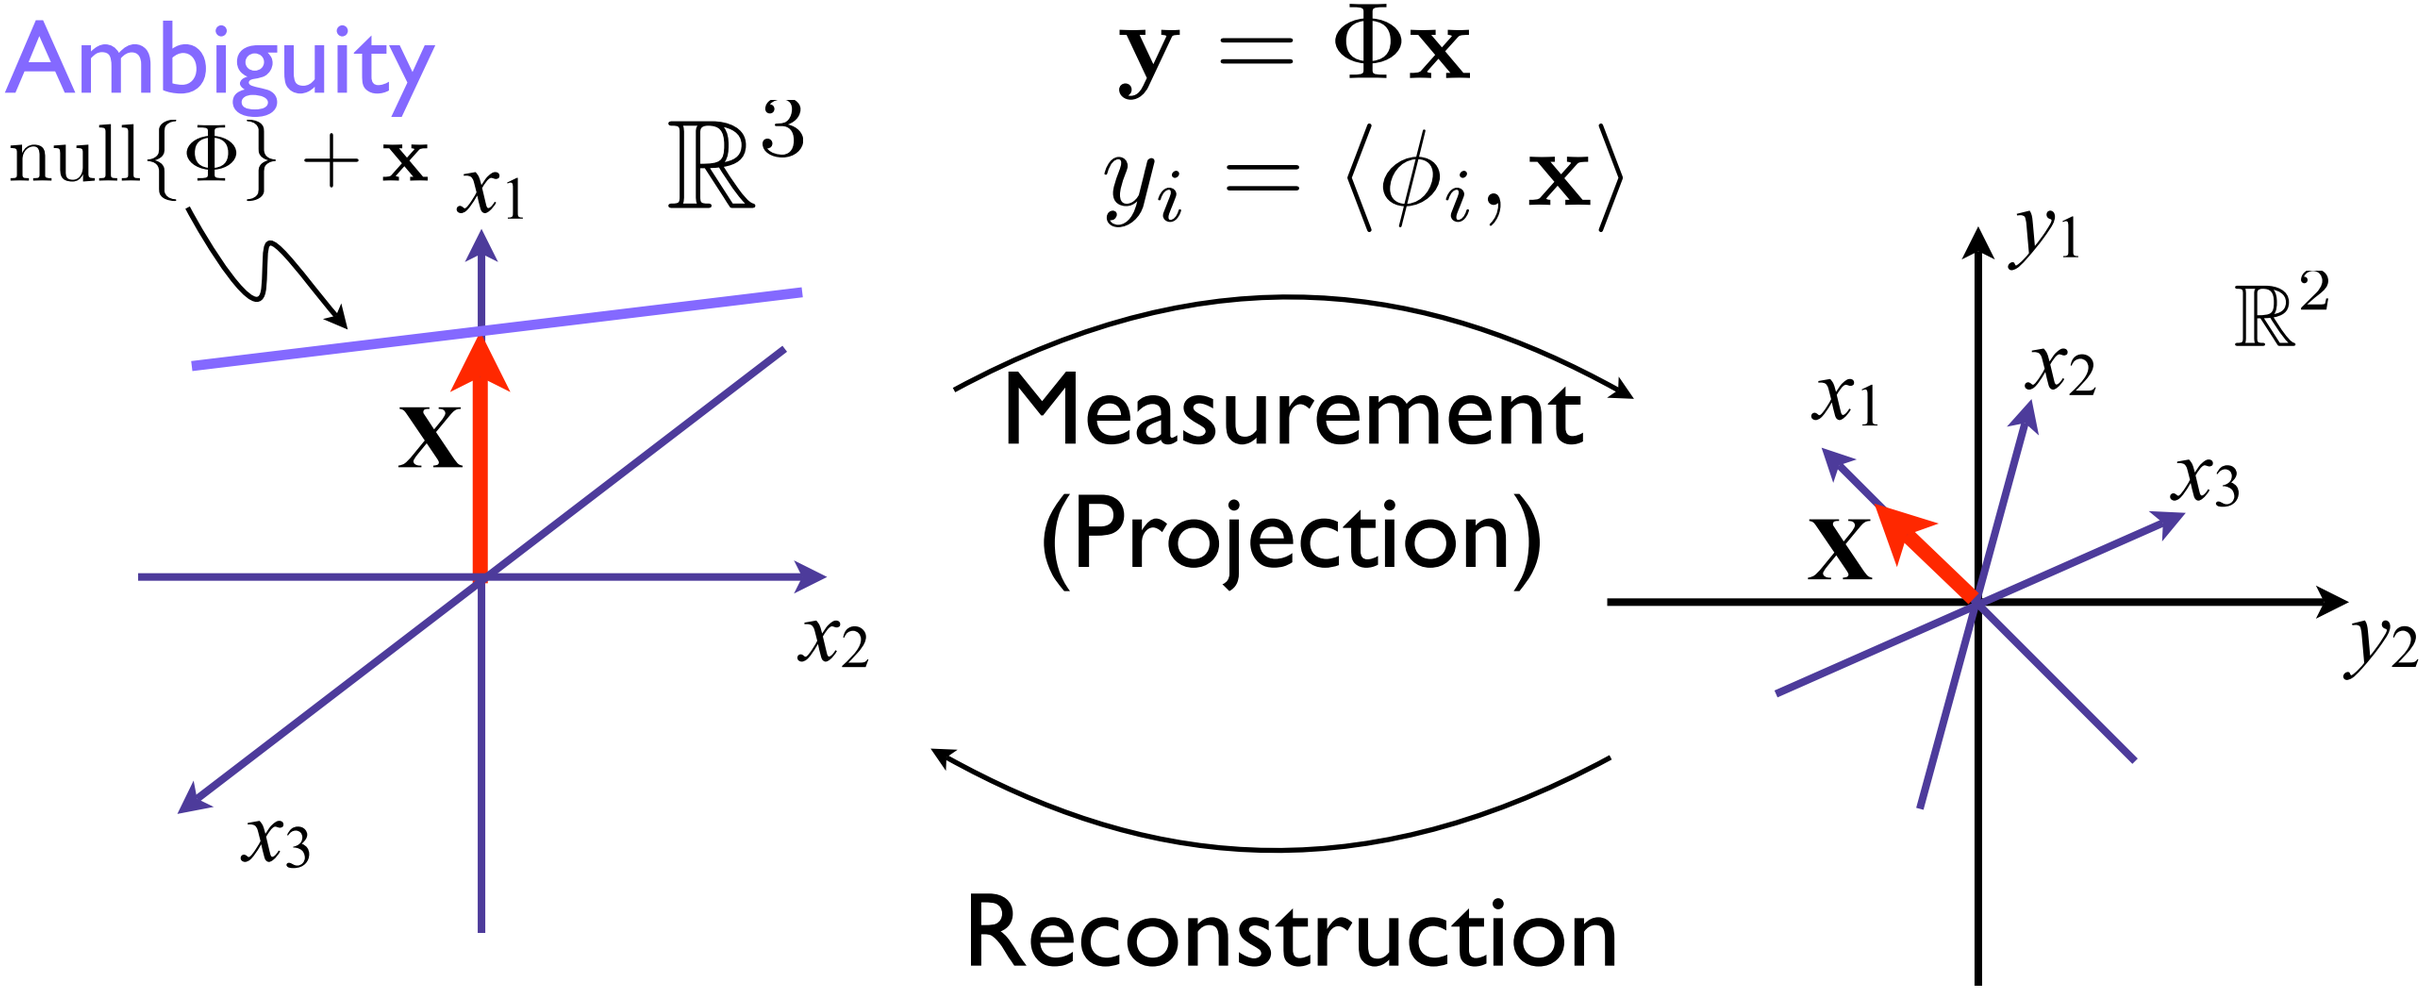

Φ has rank $M \ll N$

Φ is usually random w/ $M = O(K \log N/K)$

N = Signal dimensionality

K = Signal sparsity

M = Number of measurements (dimensionality of $\mathbf{y}$)

$$N \gg M \gtrsim K$$

Non-linear Reconstruction

Ambiguity

$\text{null}\{\Phi\} + \mathbf{x}$

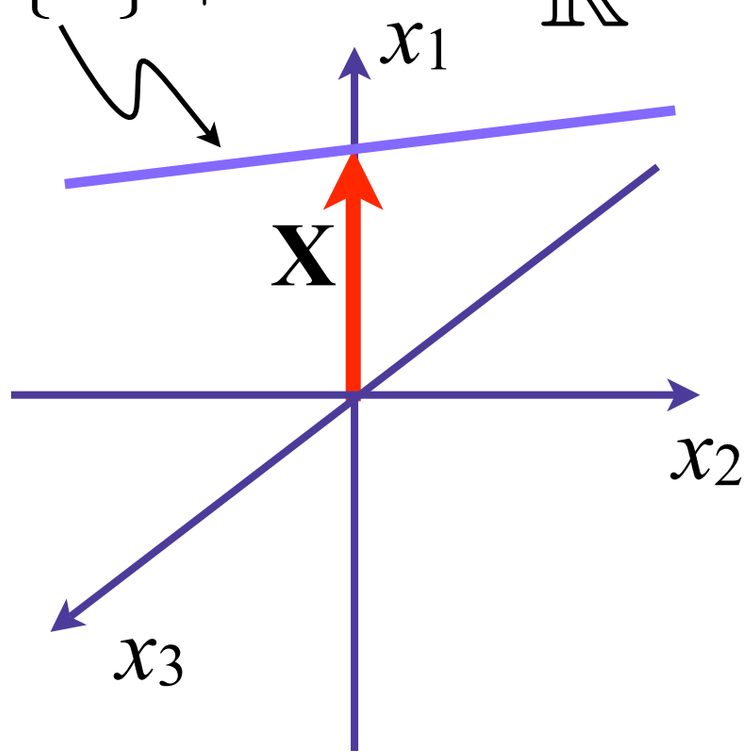

Reconstruction should be:

1. Consistent with measurements:

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

2. Consistent with the model:

$\mathbf{x}$ is as sparse as possible

Sparsity measure

Expensive!

$$\cancel{\min_{\mathbf{x}} \|\mathbf{x}\|_0 \text{ subject to } \mathbf{y} = \Phi\mathbf{x}}$$

$$\Updownarrow$$

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ subject to } \mathbf{y} = \Phi\mathbf{x}$$

Compressive Sensing

1. Signal Model
2. Random linear measurements
3. Non-linear reconstruction

Beyond Linear Measurements:
1-bit Quantization

1-Bit Compressive Sensing

Q: Can we quantize measurements to 1-bit:

$$\mathbf{y} = \mathrm{sign}(\Phi\mathbf{x})$$
$$y_i = \mathrm{sign}(\langle \phi_i, \mathbf{x} \rangle)$$

and recover the signal (within a positive scaling factor)?

1-bit measurements are inexpensive.

Focus on bits rather than measurements.

Exact recovery is not possible.

Reconstruction from 1-bit Measurements

Sign information from 1-bit measurements:

$$y_i = \text{sign}(\Phi\mathbf{x})_i \Leftrightarrow y_i \cdot (\Phi\mathbf{x})_i \geq 0$$

Reconstruction should enforce model.

Reconstruction should be consistent with measurements.

Reconstruction should enforce a non-trivial solution.

$$\begin{aligned} \widehat{\mathbf{x}} &= \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \\ \text{subject} \quad \text{to} \quad & y_i \cdot (\Phi\mathbf{x})_i \geq 0 \\ \text{and} \quad & \|\mathbf{x}\|_2 = 1 \end{aligned}$$

Information in 1-bit Measurements

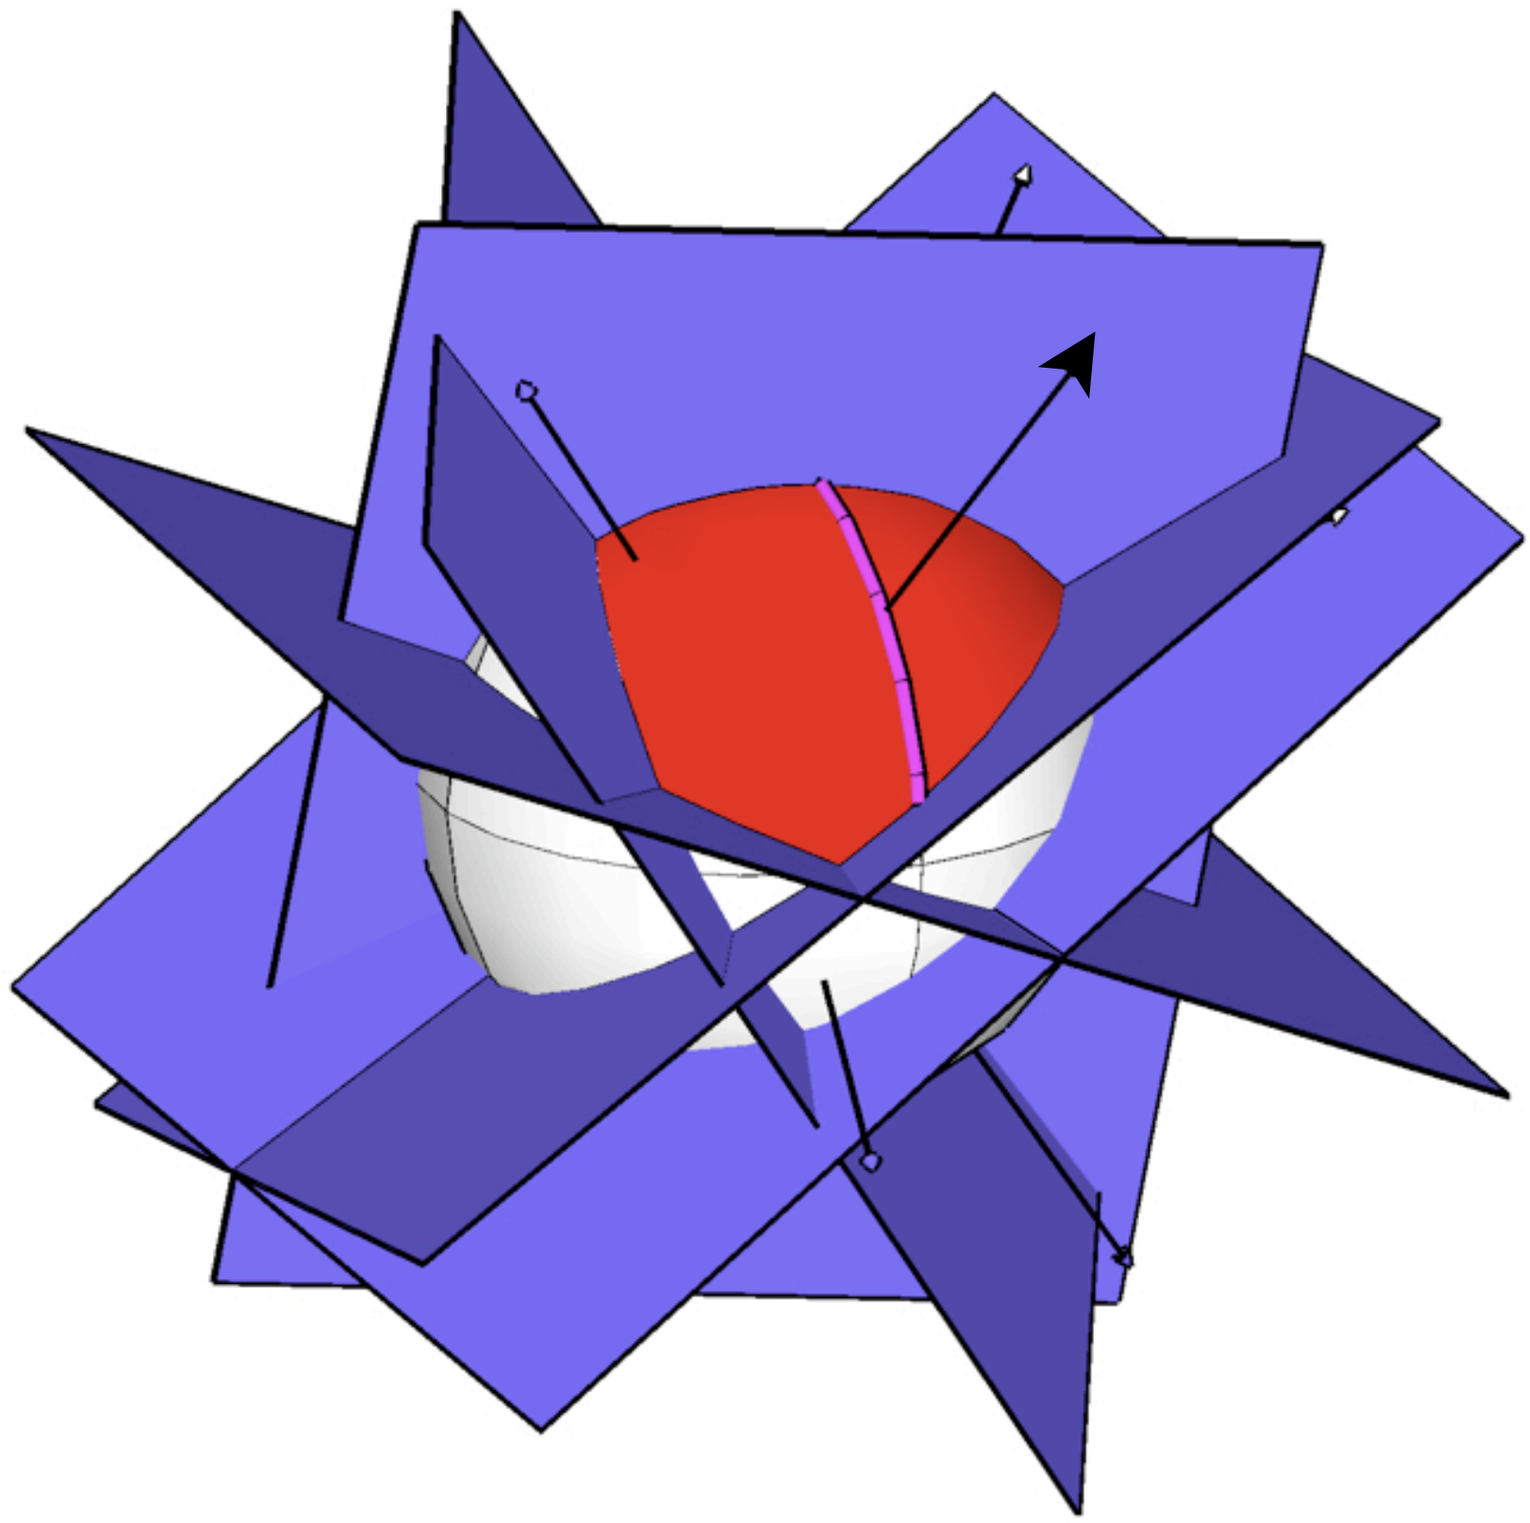

Constraint Relaxation

$$\begin{aligned} \widehat{\mathbf{x}} &= \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \\ \text{subject} \quad \text{to} \quad & y_i \cdot (\Phi\mathbf{x})_i \geq 0 \\ \text{and} \quad & \|\mathbf{x}\|_2 = 1 \end{aligned}$$

We relax the inequality constraints:

$$\begin{aligned} \widehat{\mathbf{x}} &= \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 + \frac{\lambda}{2} \sum_i f\left(y_i \cdot (\Phi\mathbf{x})\right) \\ & \text{subject to } \|\mathbf{x}\|_2 = 1 \end{aligned}$$

where $f(x)$ is a one sided quadratic:

$$f(x) = \begin{cases} x^2 & x \leq 0 \\ 0 & x > 0 \end{cases}$$

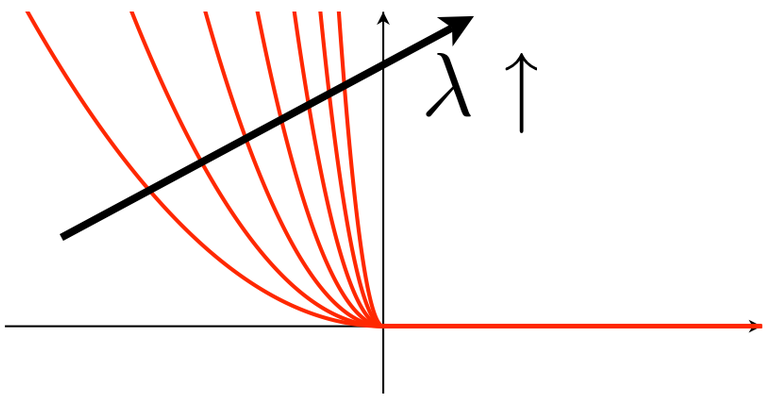

Fixed point equilibrium

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 + \frac{\lambda}{2}\sum_i f\left(y_i \cdot (\Phi\mathbf{x})\right)$$

subject to $\|\mathbf{x}\|_2 = 1$

Unconstrained minimization:

$$\begin{aligned}
\mathrm{Y} &\equiv \mathrm{diag}(\mathbf{y}) \\
\mathrm{Cost}(\mathbf{x}) &= g(\mathbf{x}) + \frac{\lambda}{2} f(\mathrm{Y}\Phi\mathbf{x}) \\
\mathrm{Cost}'(\mathbf{x}) &= g'(x) + \frac{\lambda}{2}(\mathrm{Y}\Phi)^T f(\mathrm{Y}\Phi\mathbf{x})
\end{aligned}$$

$$(g'(\mathbf{x}))_i = \begin{cases} -1 & x_i < 0 \\ [-1,1] & x_i = 0 \\ +1 & x_i > 0 \end{cases} \quad \text{and} \quad \left(\frac{f'(\mathbf{x})}{2}\right)_i = \begin{cases} -x_i & x_i \leq 0 \\ 0 & x_i > 0 \end{cases}$$

No change if gradients are projected on unit sphere.

Minimization algorithm

Big Picture: Gradient descent until equilibrium.

Initialization parameters: $\widehat{\mathbf{x}}, \ \tau$

1. Compute quadratic gradient: $\mathbf{h} = (\Upsilon\Phi)^T f'(\Upsilon\Phi\mathbf{x})$
2. Project onto sphere: $\mathbf{h}_p = \mathbf{h} - \langle \widehat{\mathbf{x}}, \mathbf{h} \rangle$
3. Quadratic gradient descent: $\widehat{\mathbf{x}} \leftarrow \widehat{\mathbf{x}} - \tau \mathbf{h}_p$
4. Shrink (ℓ_1 gradient descent):

$$\widehat{x}_i \quad \leftarrow \quad \mathrm{sign}(\widehat{x}_i) \max \left\{ |\widehat{x}_i| - \frac{\tau}{\lambda}, 0 \right\}$$

5. Normalize: $\widehat{\mathbf{x}} \quad \leftarrow \quad \dfrac{\widehat{\mathbf{x}}}{\|\widehat{\mathbf{x}}\|}$
6. Iterate until equilibrium.

Reconstruction Error (N=512)

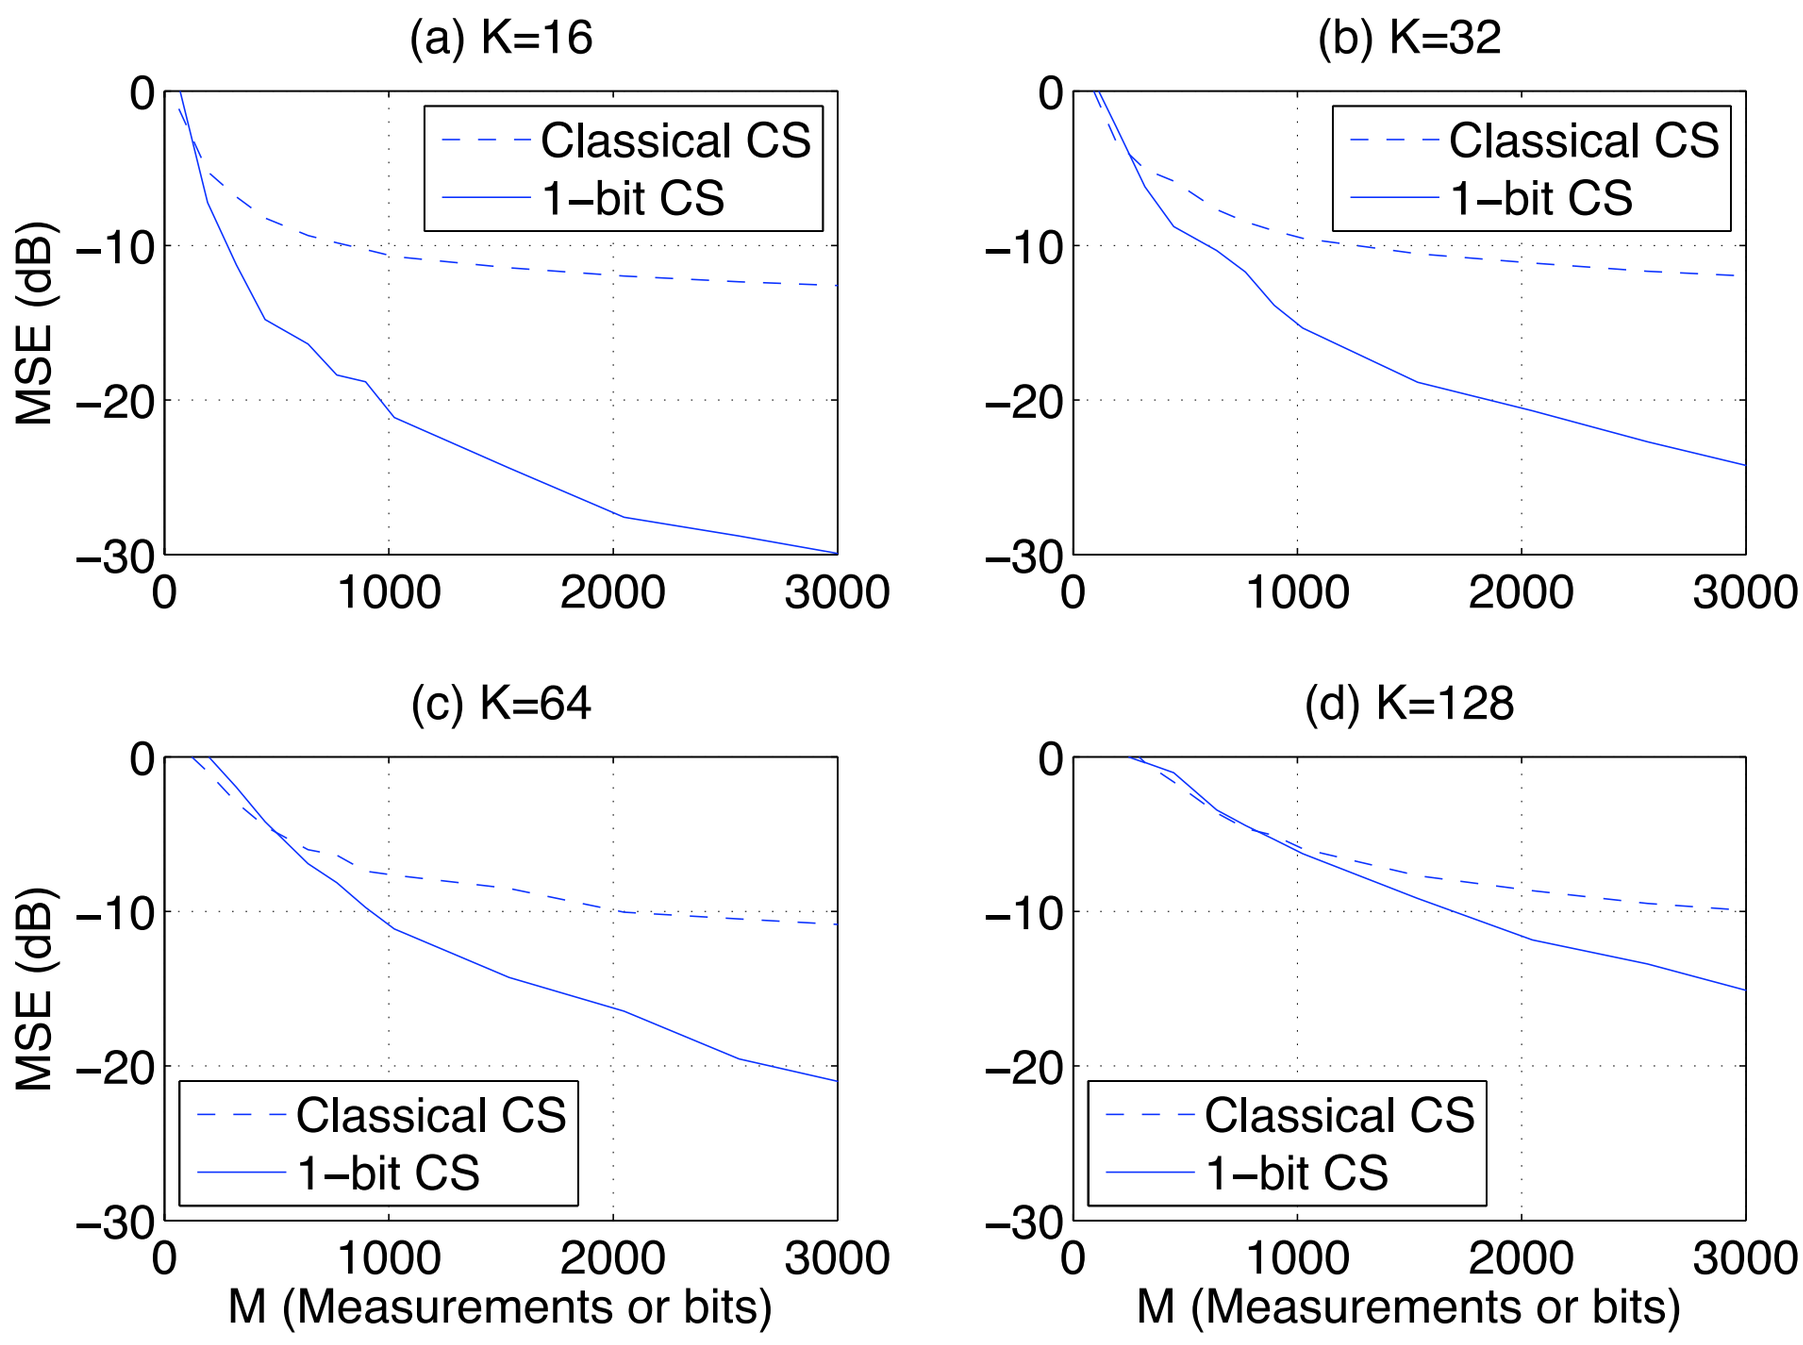

1-bit Sampling of Images

If the signal is an image, we have more information!
(i.e., a better signal model)

Images are sparse in wavelets and positive:

$$\begin{aligned} \mathbf{x} &= \mathbf{W}\alpha \\ x_i &\geq 0 \end{aligned}$$

and α is sparse

Incorporate better model in the reconstruction:

$$\hat{\alpha} = \arg\min_{\alpha} \|\alpha\|_1$$

subject to $y_i \times (\Phi\mathbf{W})_i \geq 0$

and $(\mathbf{W}\alpha)_i \geq 0$

and $\|\alpha\|_2 = 1$

Results

Reconstruction on unit sphere
1 bit per pixel

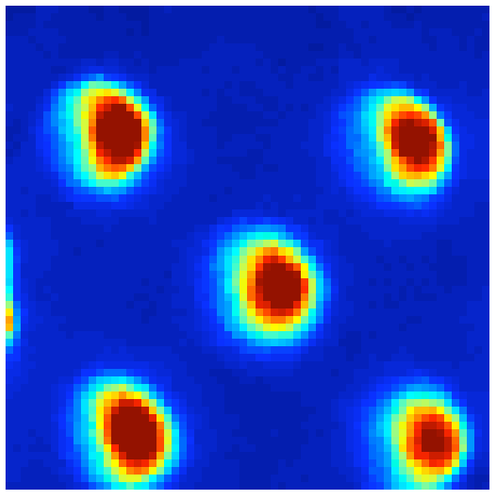

Original Image
4096 pixels
256 levels

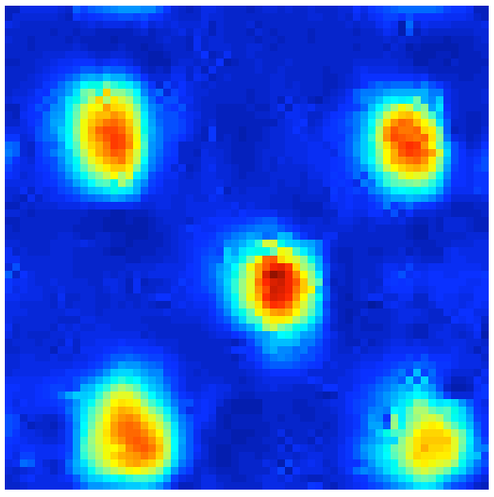

4096 measurements
1 bit per measurement

Classical Compressive Sensing, 1 bit per pixel

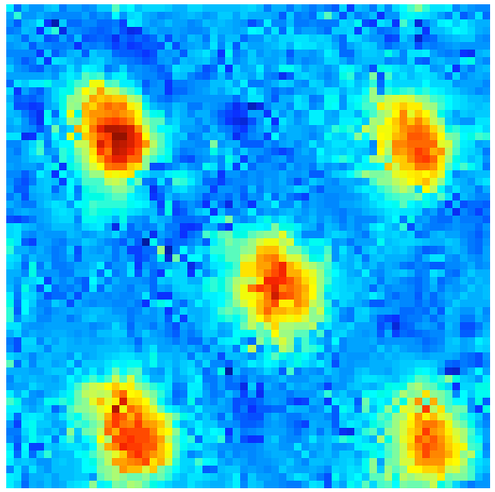

4096 measurements
1 bit per measurement

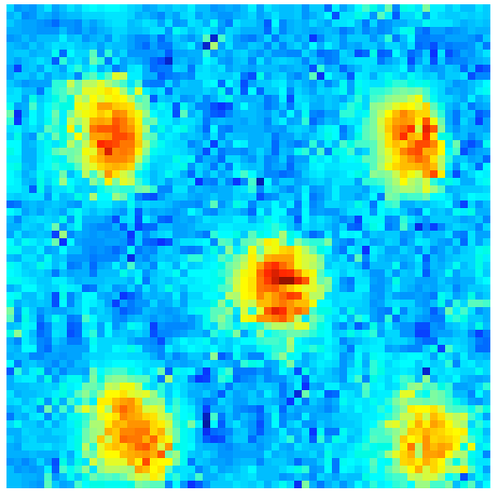

2048 measurements
2 bits per measurement

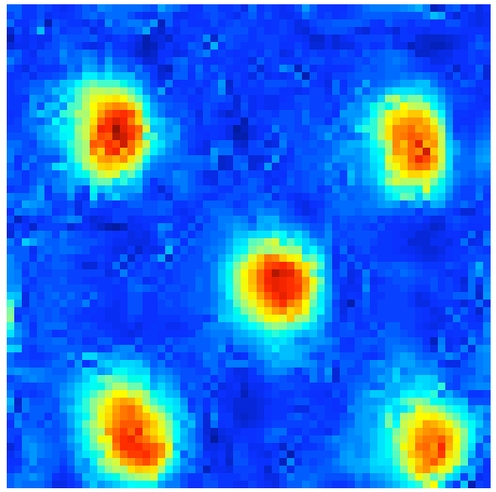

1024 measurements
4 bits per measurement

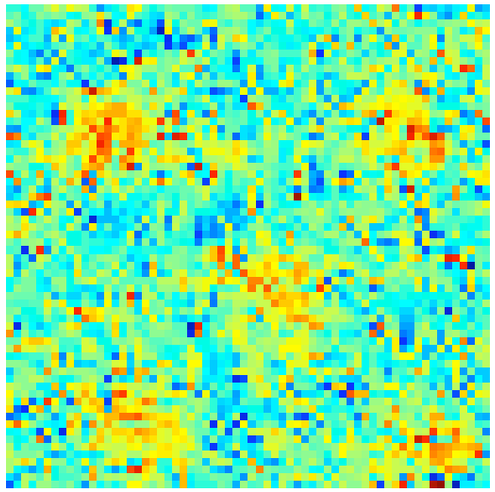

512 measurements
8 bits per measurement

Results

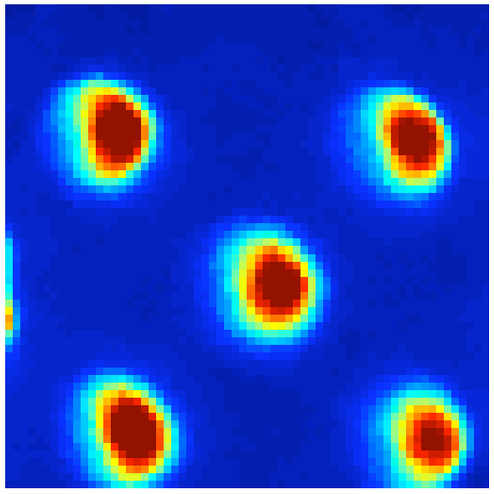

Original Image
4096 pixels
256 levels

Reconstruction on unit sphere

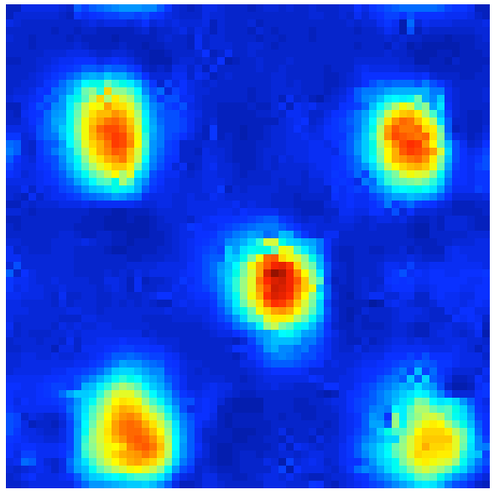

4096 measurements
1 bit per measurement
4096 bits (1 bit per pixel)

Classical Compressive Sensing

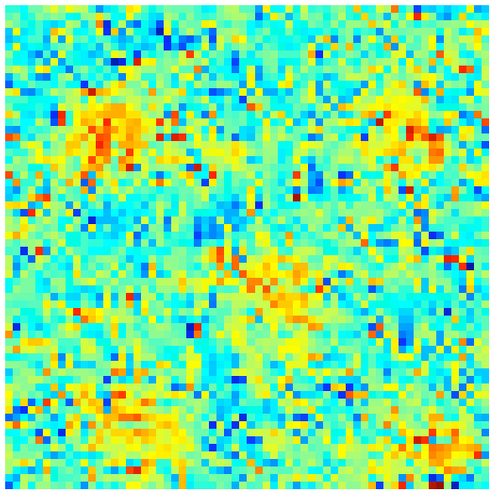

512 measurements
8 bits per measurement

Reconstruction on unit sphere

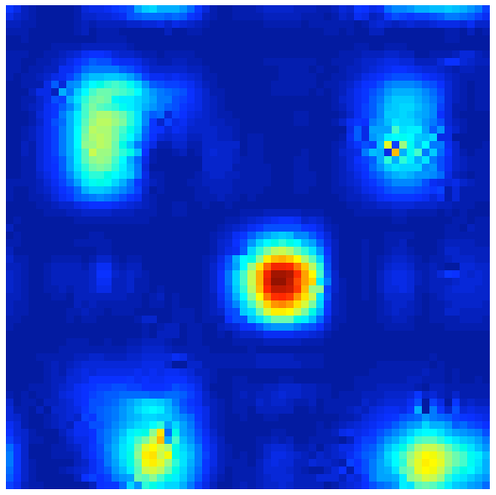

512 measurements
1 bit per measurement
512 bits (0.125 bits per pixel)

1-Bit Compressive Sensing

Q: Can we quantize measurements to 1-bit:

$$\mathbf{y} = \mathrm{sign}(\Phi\mathbf{x})$$
$$y_i = \mathrm{sign}(\langle \phi_i, \mathbf{x} \rangle)$$

and recover the signal (within a positive scaling factor)?

YES:

- 1-bit measurements only provide sign information
- We treat measurements as constraints

- We do not try to recover amplitude information
- We enforce reconstruction on the unit sphere

- Better signal model provides dramatic improvements