MITSUBISHI ELECTRIC RESEARCH LABORATORIES
Cambridge, Massachusetts

Recent Advances in Signal Acquisition, Sensing, and Quantization

Petros Boufounos
petrosb@merl.com

CMU
Nov 13, 2009

Evolution of Sensing

~1930s

Evolution of Sensing

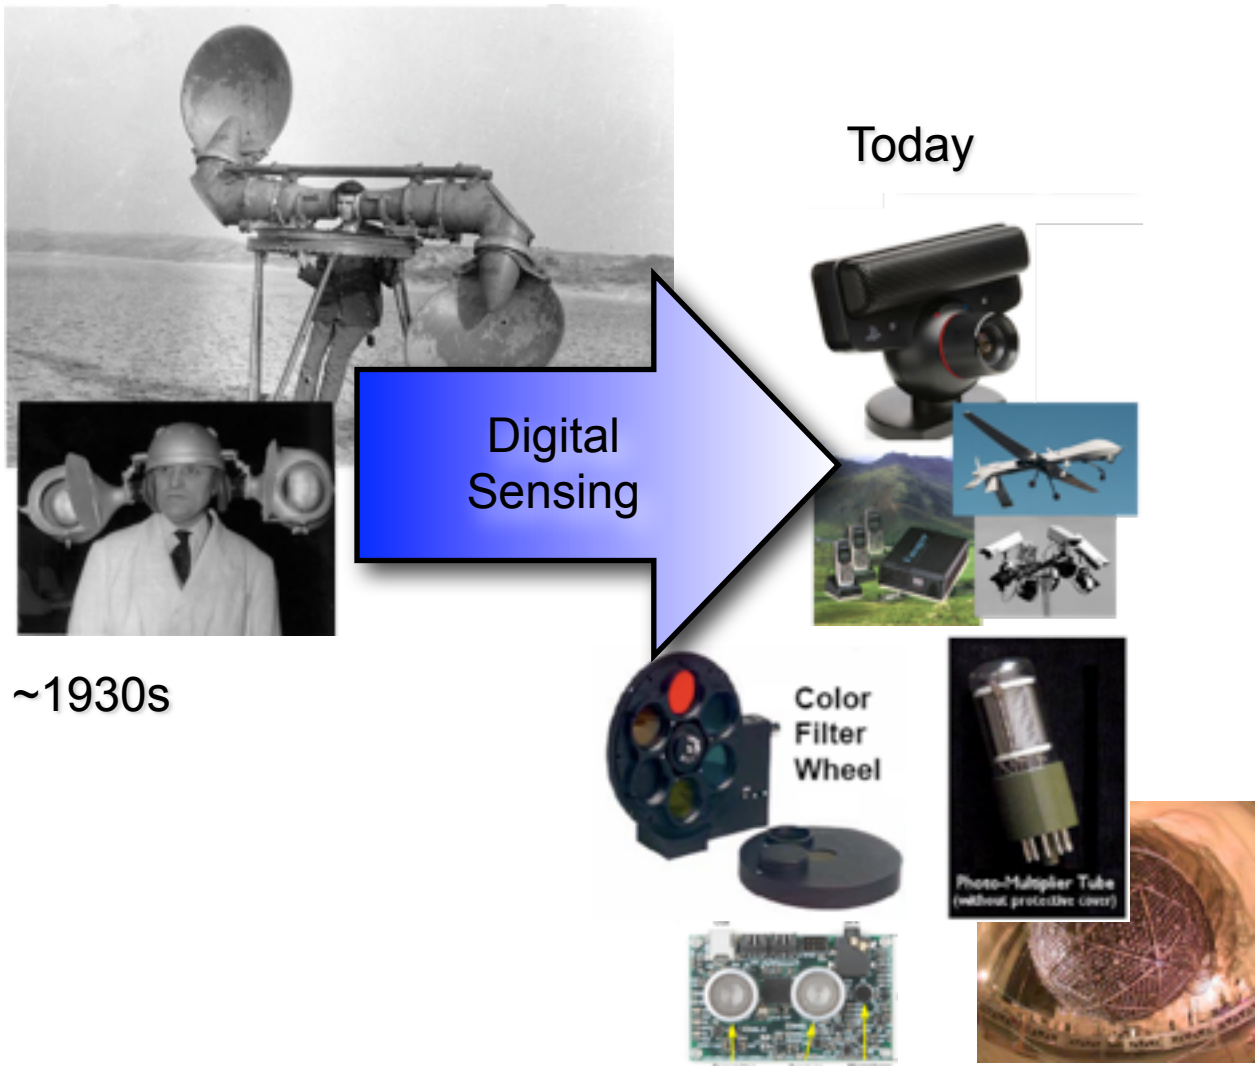

Evolution of Sensing

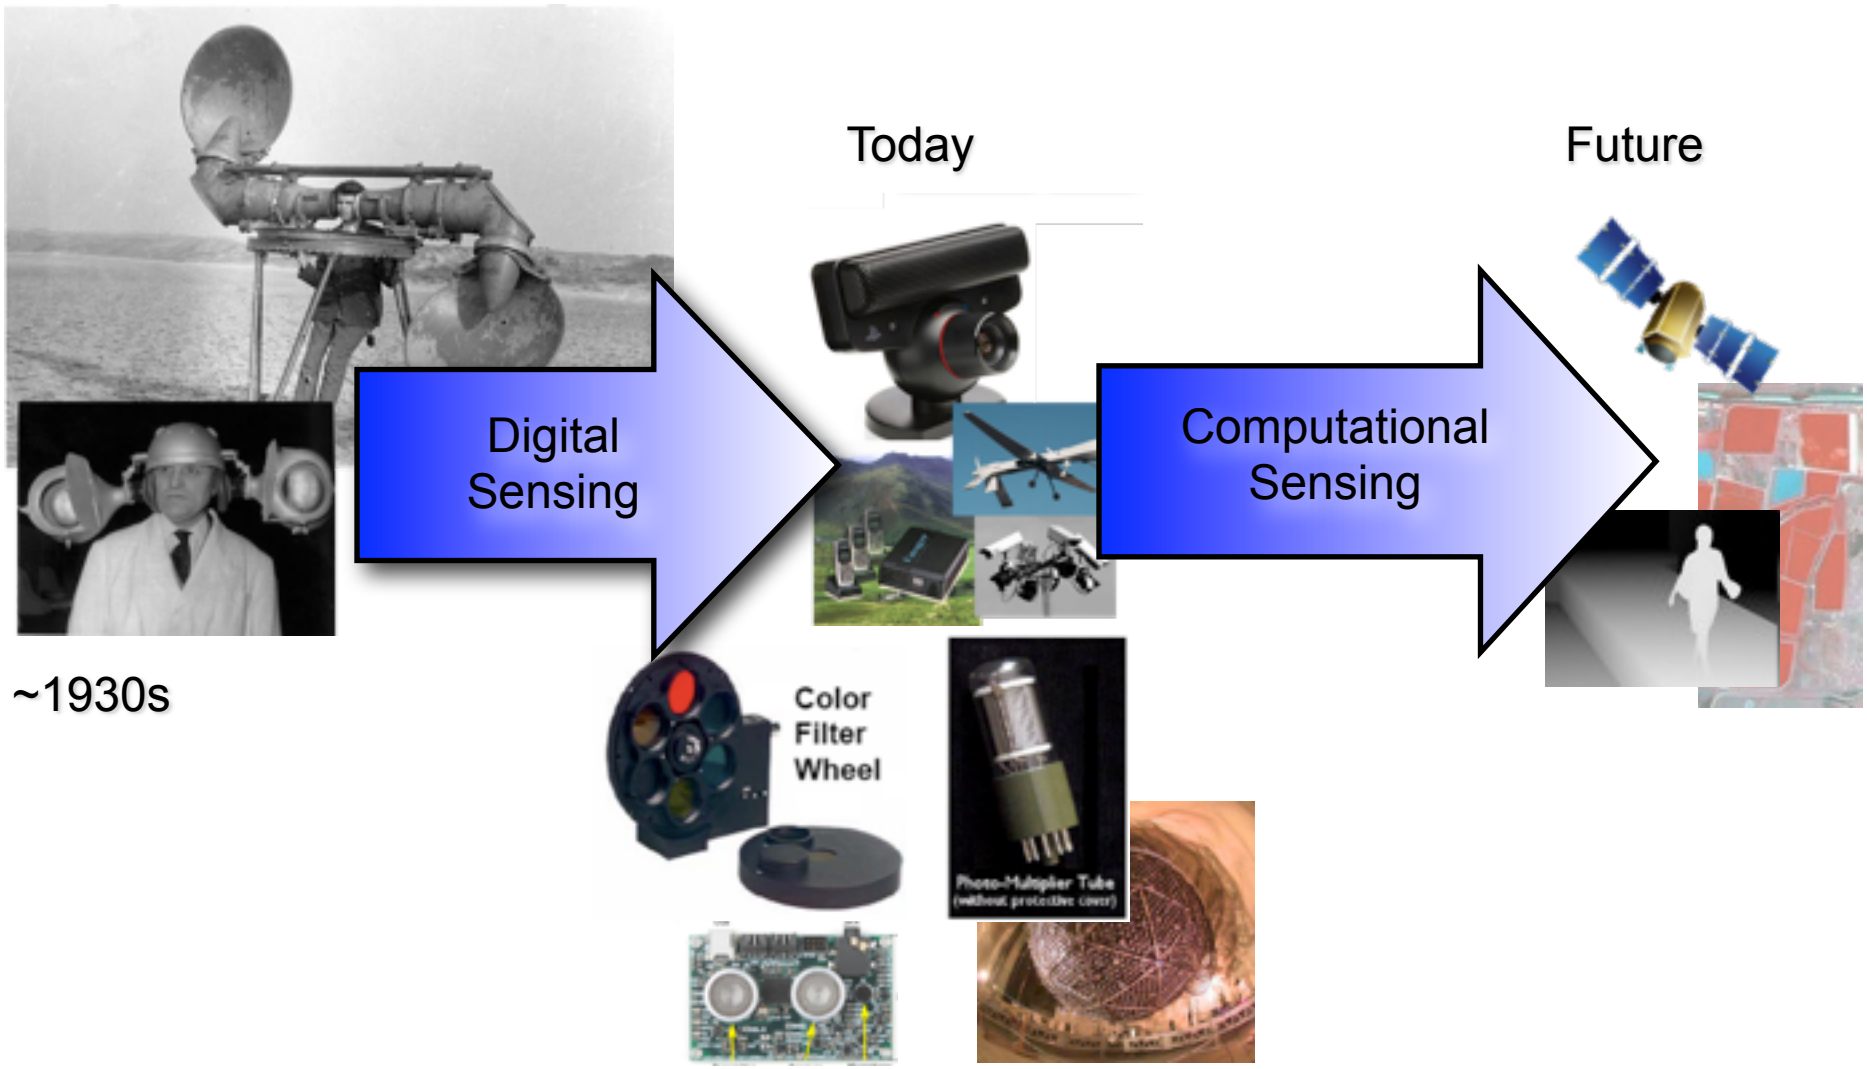

Sensing by Sampling

- Long-established paradigm for digital data acquisition
  - ***sample*** data (A-to-D converter, digital camera, …)
  - ***compress*** data (signal-dependent, nonlinear)
  - ***bottleneck*** to performance of modern acquisition systems

$N \gg K$

f → sample —N→ **compress** —K→ transmit

sparse wavelet transform

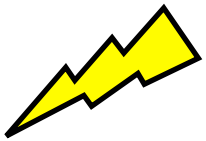

receive —K→ decompress —N→ $\hat{f}$

Compressive Sensing (CS) [Candés, Romberg, Tao; Donoho]

- New signal acquisition method
  - Samples and compresses in **one simple step**
  - Uses **computation** to reconstruct signal

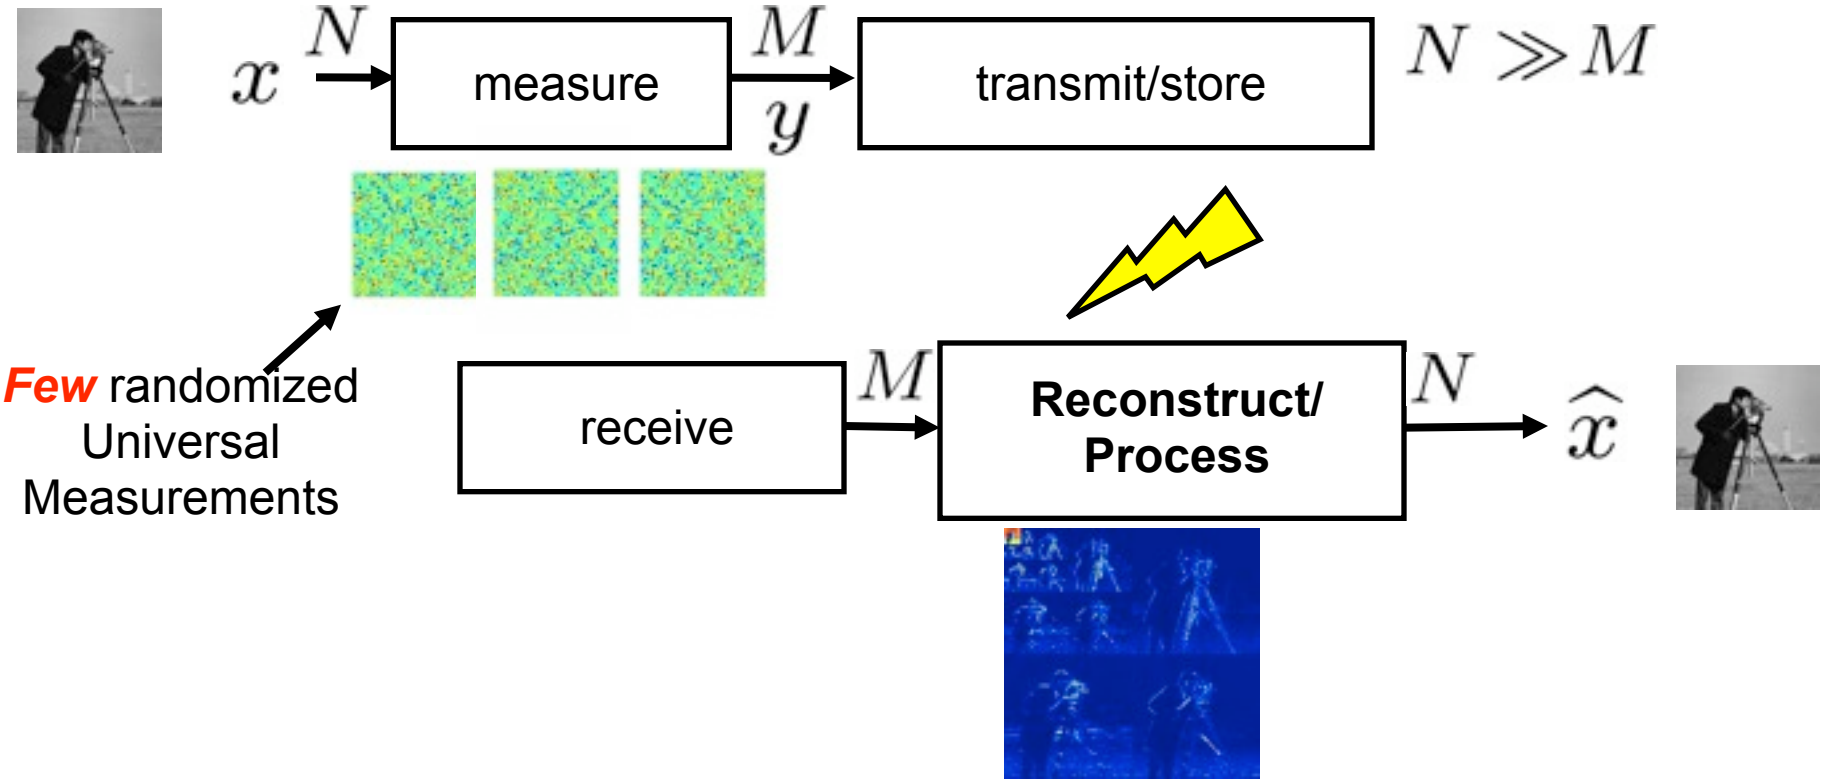

CS AT A GLANCE

Compressed Sensing Measurement Model

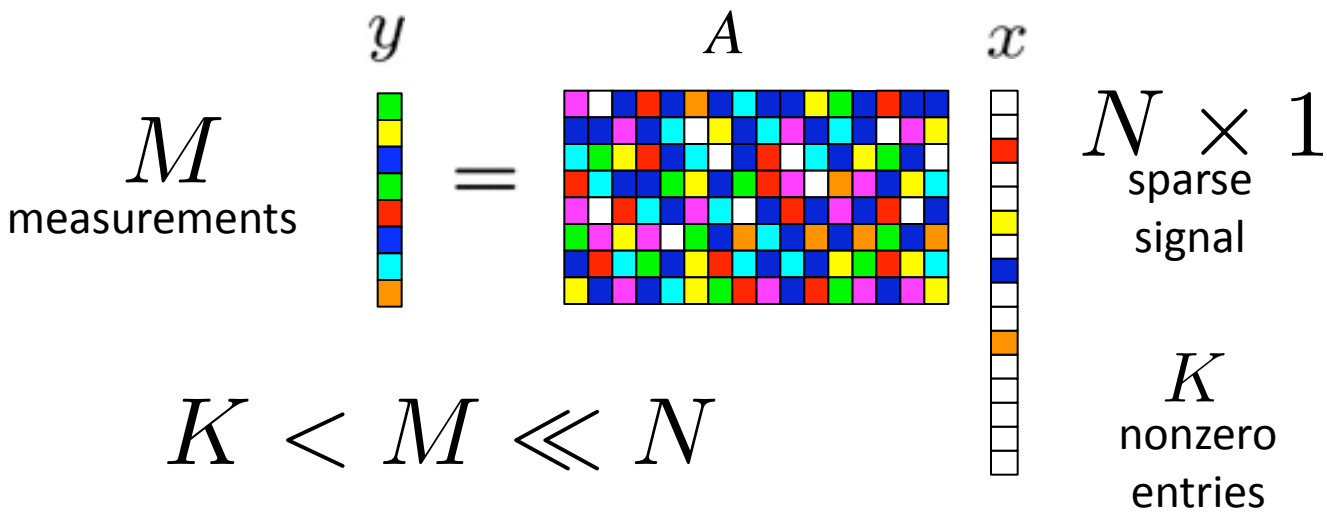

- x is K-sparse or K-compressible
- A **random**, satisfies a *restricted isometry property (RIP)*

  A has RIP of order $2K$ with constant δ

  If there exists δ s.t. for all $2K$-sparse x:

  $$(1-\delta)\|\mathbf{x}\|_2^2 \leq \|\mathbf{A}\mathbf{x}\| \leq (1+\delta)\|\mathbf{x}\|_2^2$$
- M=O(Klog$N$/K)
- A also has small *coherence* $\mu \triangleq \max_{i \neq j} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$

Compressed Sensing Measurement Model

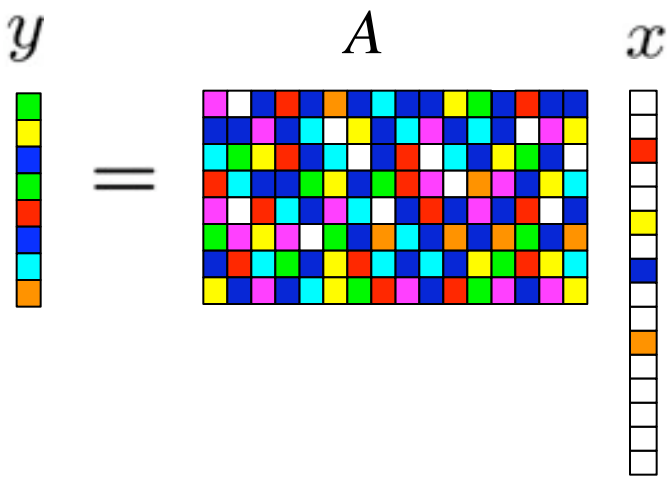

- x is K-sparse or K-compressible
- A **random**, satisfies a *restricted isometry property (RIP)*

  A has RIP of order $2K$ with constant δ

  If there exists δ s.t. for all $2K$-sparse x:

  $$(1-\delta)\|\mathbf{x}\|_2^2 \leq \|\mathbf{A}\mathbf{x}\| \leq (1+\delta)\|\mathbf{x}\|_2^2$$
- M=O(Klog$N$/K)
- A also has small *coherence* $\mu \triangleq \max_{i \neq j} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$

Compressed Sensing Measurement Model

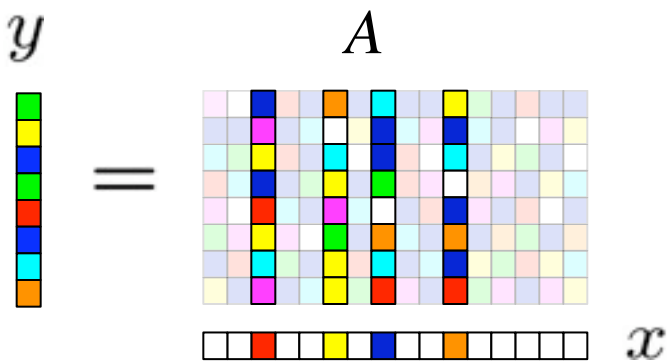

- x is K-sparse or K-compressible
- A **random**, satisfies a *restricted isometry property (RIP)*

A has RIP of order $2K$ with constant δ

If there exists δ s.t. for all $2K$-sparse x:

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

- $M = O(K \log N/K)$
- A also has small *coherence* $\mu \triangleq \max_{i \neq j} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$

CS RECONSTRUCTION

CS Reconstruction

- Reconstruction using **sparse approximation**:
  - Find sparsest $\mathbf{x}$ such that $\mathbf{y} \approx \mathbf{A}\mathbf{x}$
- **Convex optimization** approach:
  - Minimize l_1 norm: e.g.,

$$\widehat{\mathbf{x}} = \arg\max_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{s.t. } \mathbf{y} = \mathbf{A}\mathbf{x}$$

- **Greedy algorithms** approach:
  - MP, OMP, ROMP, StOMP, CoSaMP, …
  - **AndrewMP**, PYAMP (Pick Your Acronym Matching Pursuit)
- If coherence μ or RIP δ is ***small***: Exact reconstruction

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

$$\min \ \| \mathbf{x} \|_0 \ \text{s.t.} \ \ \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

$$\min \|\mathbf{x}\|_0 \quad \text{s.t.} \quad \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Number of non-zeros
(sparsity measure)

Data Fidelity
(approximation quality)

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

$$\min \|\mathbf{x}\|_0 \quad \text{s.t.} \quad \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Number of non-zeros
(sparsity measure)

Data Fidelity
(approximation quality)

Combinatorial complexity.
Very hard problem!

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

$$\min \|\mathbf{x}\|_{\cancel{0}} \text{ s.t. } \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Convex Relaxation

$$\min \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Optimization (Basis Pursuit)

Sparse approximation:

Minimize non-zeros in representation
s.t.: representation is **close** to signal

$$\min \ \| \mathbf{x} \|_{\cancel{0}} \ \text{s.t.} \ \ \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Convex Relaxation

$$\min \ \| \mathbf{x} \|_{1} \ \text{s.t.} \ \ \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

Ploynomial complexity.
Solved using linear programming.

Why l_1 relaxation works

$$\min \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

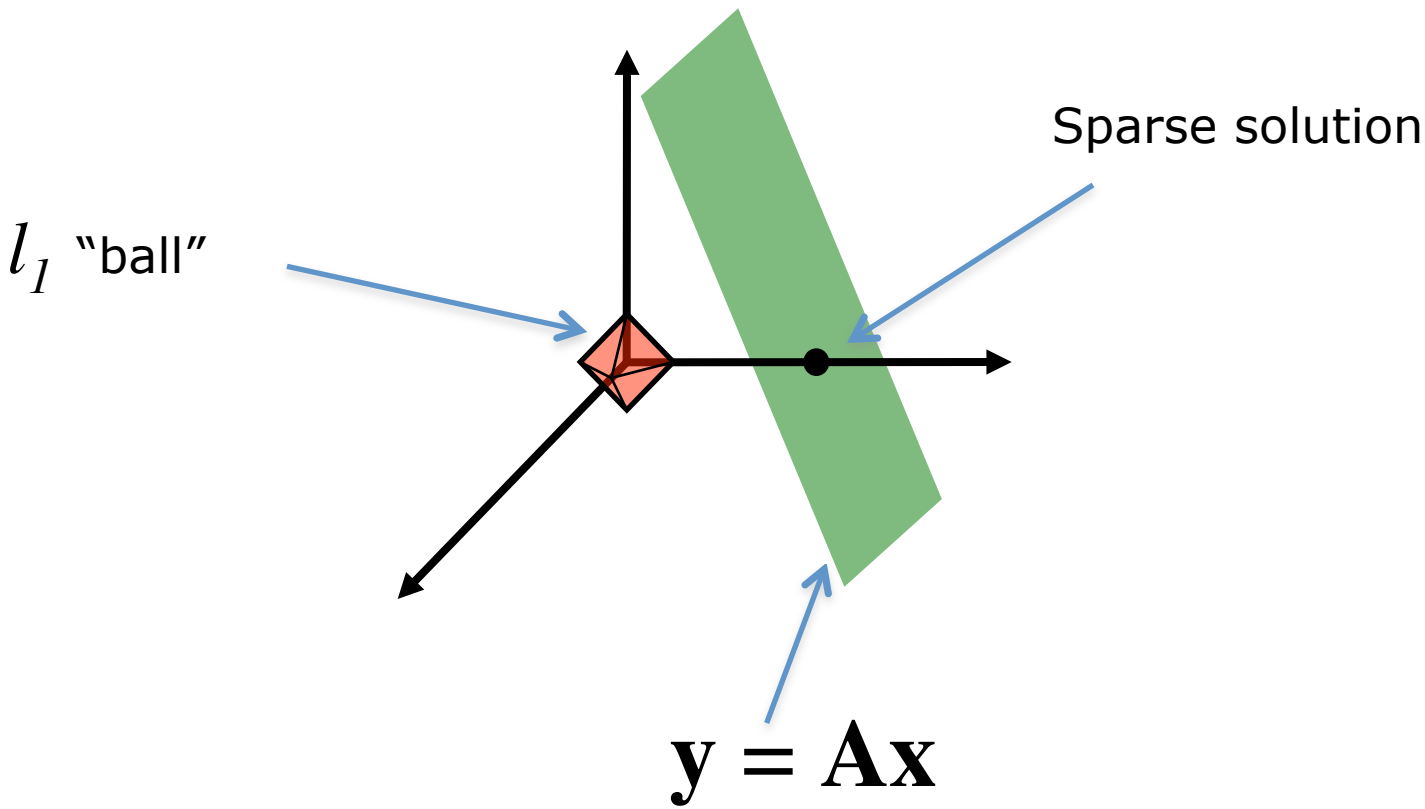

Why l_1 relaxation works

$$\min \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} \approx \mathbf{A}\mathbf{x}$$

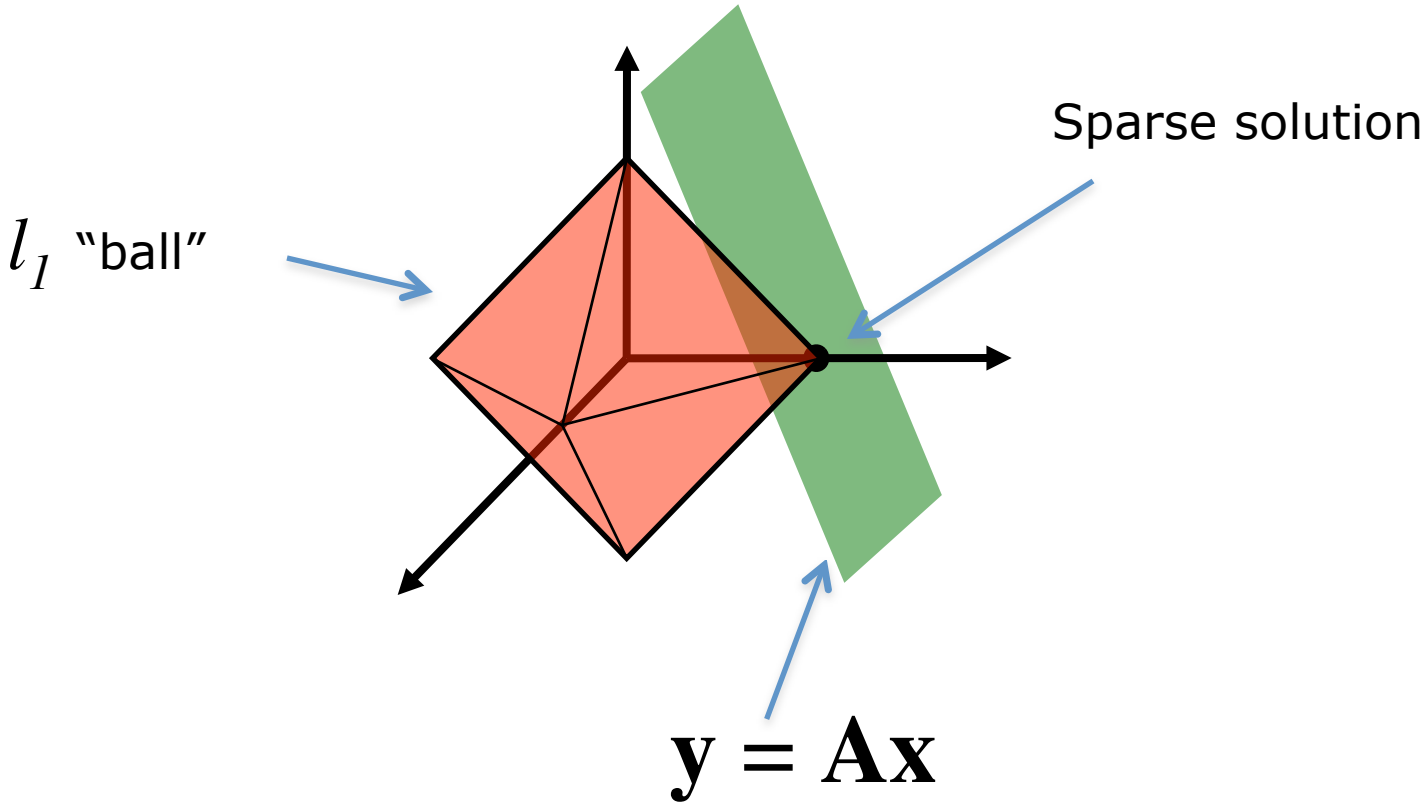

Greedy Pursuits Core Idea

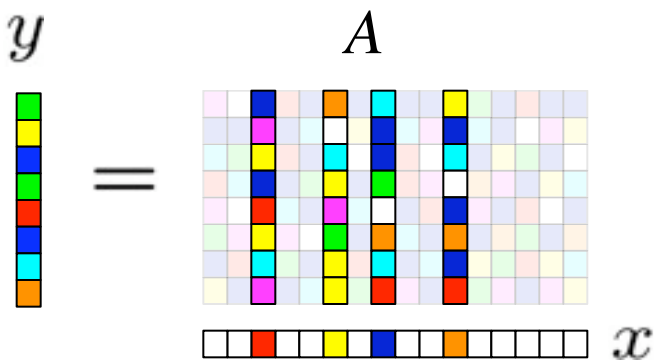

- $\boldsymbol{y}$ highly correlated with $\boldsymbol{A}$ at locations where $\boldsymbol{x}$ is high
- $\boldsymbol{A}^T\boldsymbol{y}$ provides a good idea of these locations
  - This is why low coherence is important

$$\mu \triangleq \max_{i \neq j} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$$

  - $\boldsymbol{A}^T\boldsymbol{y}$ referred to as *proxy* for $\boldsymbol{x}$
- General Strategy:
  - Identify locations
  - Invert the system only on those locations

CoSaMP (Compressive Sampling MP) [Needell and Tropp]

Correlate residual
with dictionary
→ signal proxy

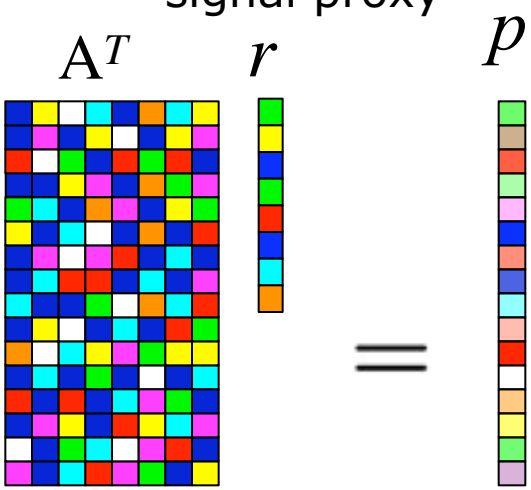

$$\langle a_k, r \rangle = p_k$$

Signal estimate: $\hat{x}$

Support estimate: T

Residual: $r = y - A\hat{x}$

CoSaMP (Compressive Sampling MP) [Needell and Tropp]

Correlate residual
with dictionary
→ signal proxy

Select location
of largest
2K correlations

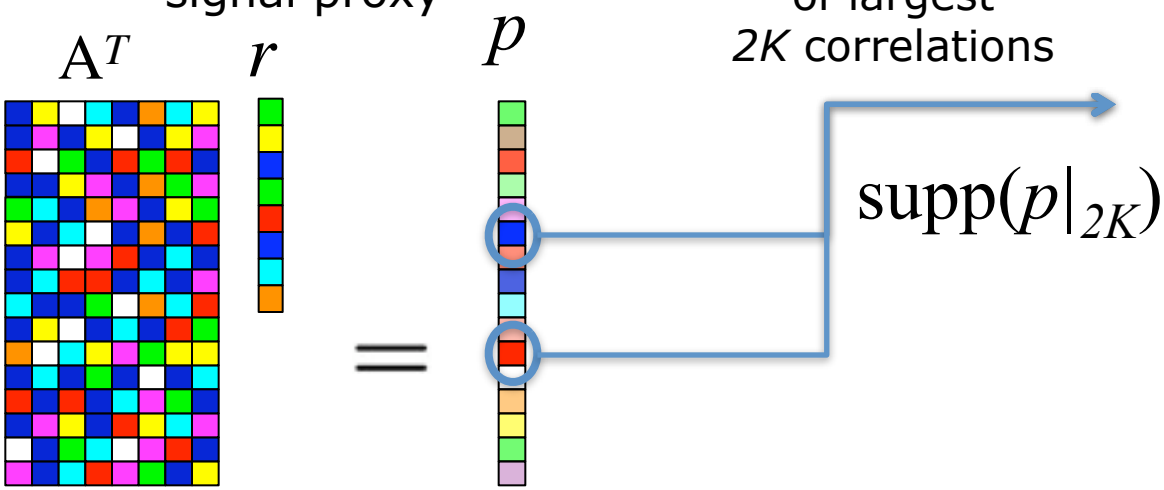

$$\langle a_k, r \rangle = p_k$$

Signal estimate: $\hat{x}$

Support estimate: T

Residual: $r = y - A\hat{x}$

CoSaMP (Compressive Sampling MP) [Needell and Tropp]

Correlate residual with dictionary → signal proxy

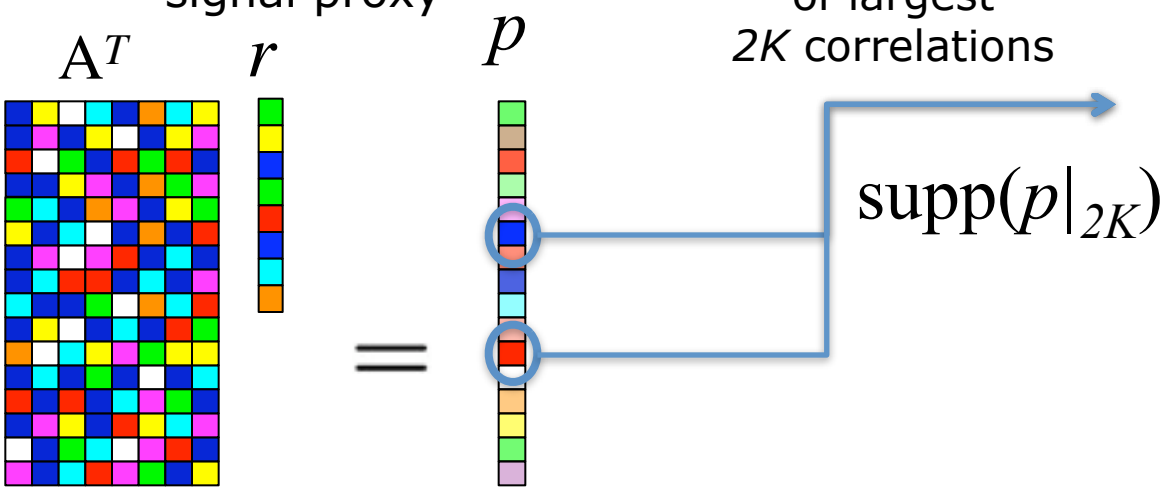

$$\langle a_k, r \rangle = p_k$$

Select location of largest *2K* correlations

Add to support set

$$\Omega = \mathrm{supp}(p|_{2K}) \cup T$$

Invert over support

$$b = A_\Omega^\dagger y$$

Truncate and compute residual

$$T = \mathrm{supp}(b|_K)$$

$$\widehat{x} = b|K$$

$$r \leftarrow y - A\widehat{x}$$

Signal estimate: $\hat{x}$

Support estimate: T

Residual: $r=y-A\hat{x}$

CoSaMP (Compressive Sampling MP) [Needell and Tropp]

Correlate residual with dictionary → signal proxy

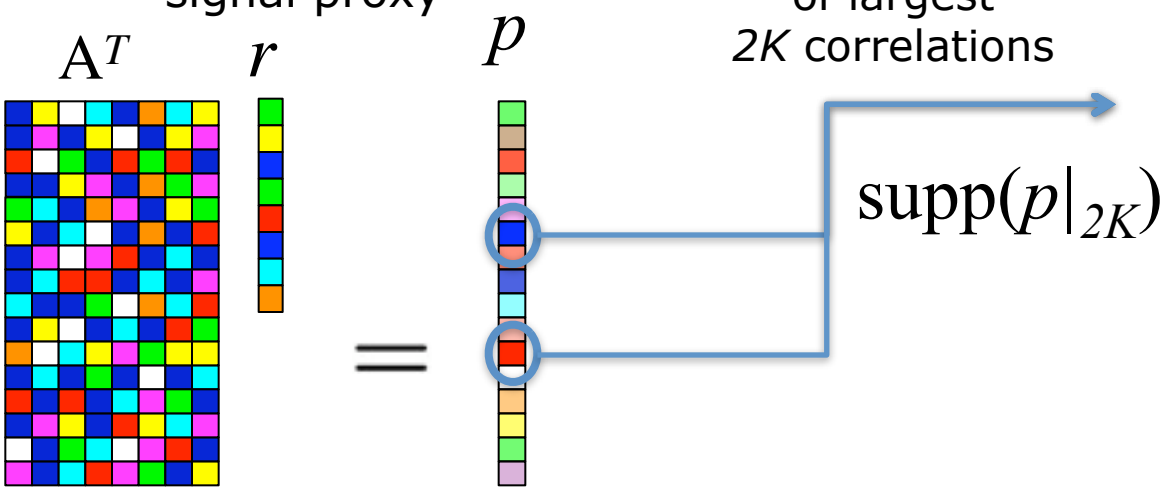

Select location of largest *2K* correlations

Add to support set

$$\Omega = \mathrm{supp}(p|_{2K}) \cup T$$

Invert over support

$$b = A_\Omega^\dagger y$$

Truncate and compute residual

$$T = \mathrm{supp}(b|_K)$$

$$\widehat{x} = b|K$$

$$r \leftarrow y - A\widehat{x}$$

$$\langle a_k, r \rangle = p_k$$

Signal estimate: $\hat{x}$

Support estimate: T

Residual: $r = y - A\hat{x}$

Iterate using residual

QUANTIZATION

Classical Sampling and Quantization

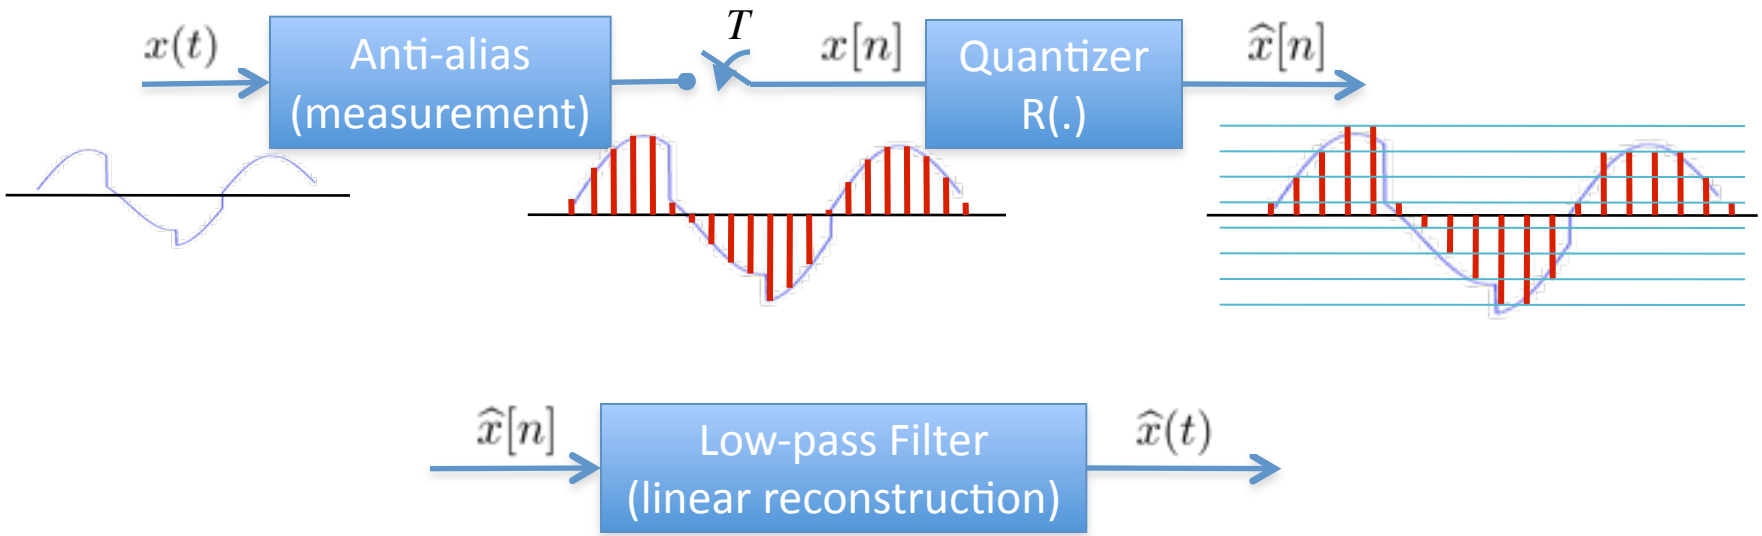

Sampling: discretization in time
Lossless at the Nyquist rate

Quantization: discretization in amplitude
Always lossy

Need **both** for digital data acquisition

Quantizer Design

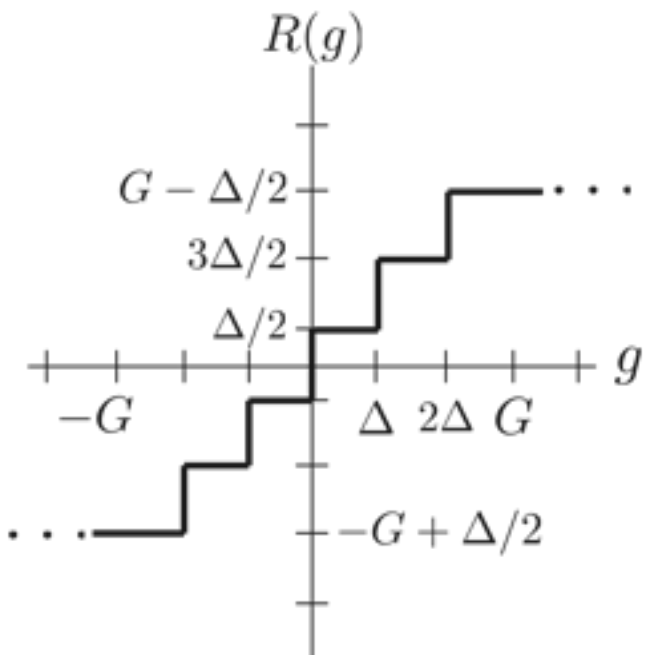

- Finite Range Quantizer
  - Finite range necessary
  - Hardware constraints
  - Bit-length constraints
- Design parameters:
  - G: Saturation level (threshold)
  - Δ: Quantization interval
  - B: Bit-rate (per coefficient)
- Design goals:
  - Smaller interval Δ: Less quantization error
  - Larger threshold G: Less saturation error
  - Trade-off: $\Delta=2^{-B}G$

Classical Quantizer Design

- Linear reconstruction penalizes saturation significantly

- Universal design principle:
  reduce or eliminate saturation

  Given: **bit budget** B bits/sample
  Given: **maximum** signal **amplitude** A
  Set: **threshold** $G \leq A$, i.e., $\Delta \geq 2^{-B+1}A$

- Drawback: **amplitude** A often **unknown**
  - In practice: Automatic Gain Control
  - Setting can be too **conservative**

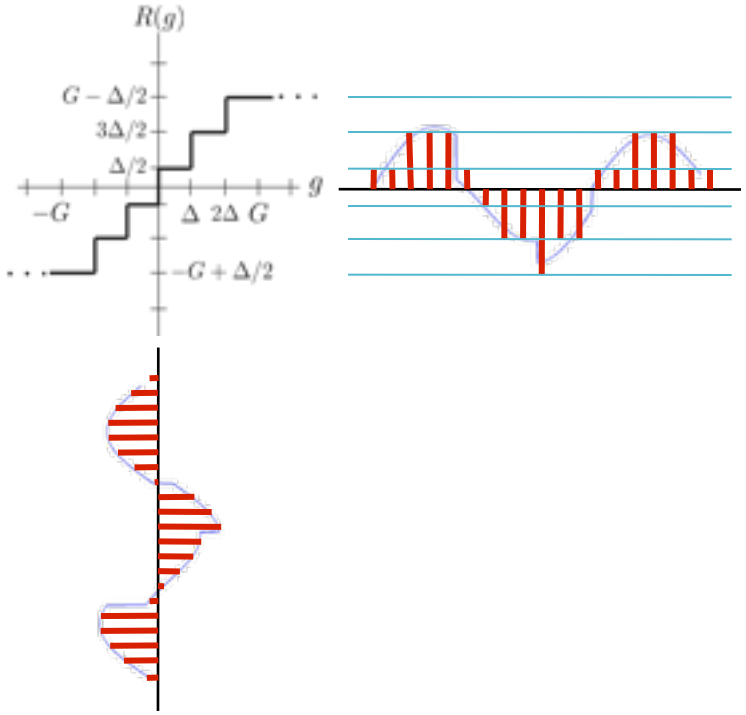

CS and Quantization [Laska, Boufounos, Davenport, Baraniuk]

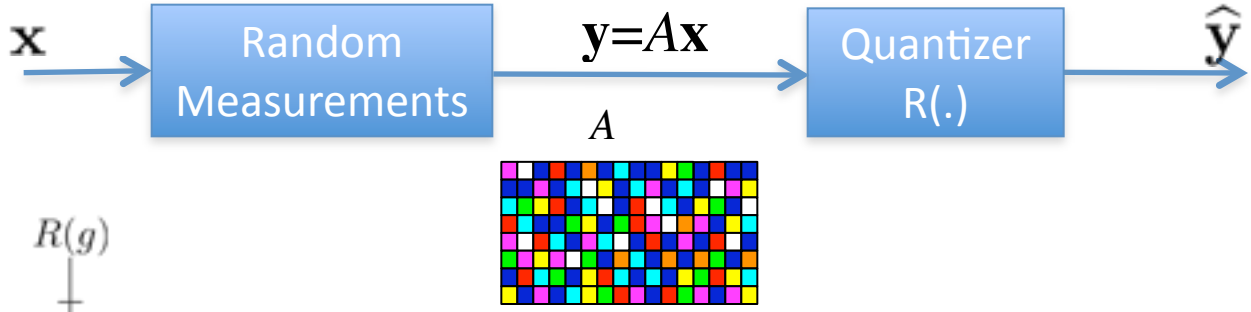

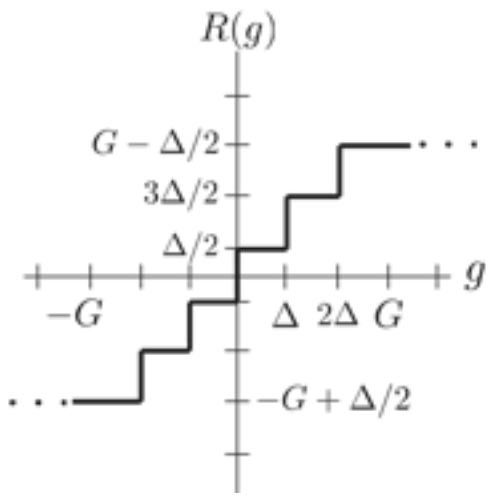

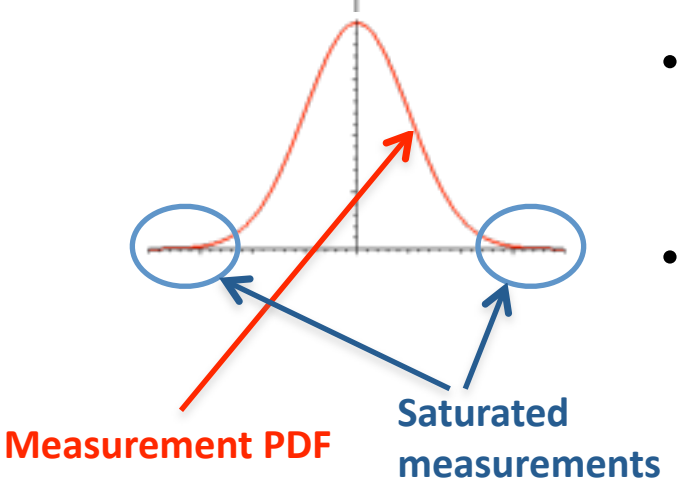

- Given: **Bit budget** B bits/sample, **Signal norm** $\|\mathbf{x}\|_2$
- Set quantization **threshold** G
  - Implicitly sets quantization interval $\Delta = 2^{-B+1}G$
  - Implicitly sets saturation rate at $2Q(G/\|\mathbf{x}\|_2)$
- Classical heuristic: Set G large (avoid saturation)
- Note:
  - equivalent to fixing G and varying signal amplification
  - $Q(\cdot)$ denotes the tail of the Gaussian distribution

CS and Quantization [Laska, Boufounos, Davenport, Baraniuk]

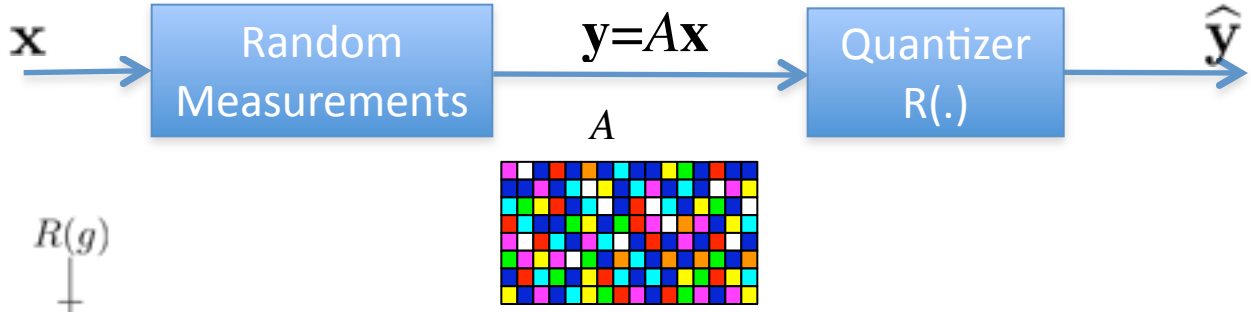

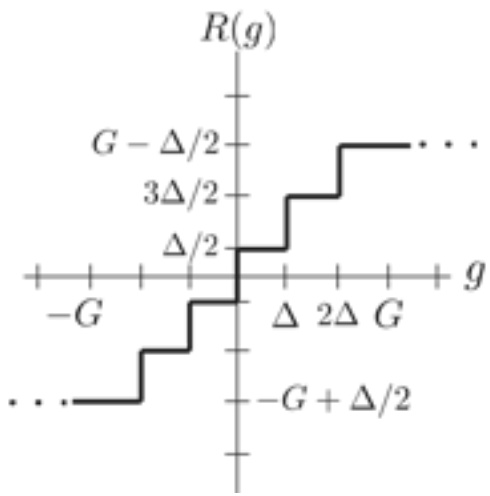

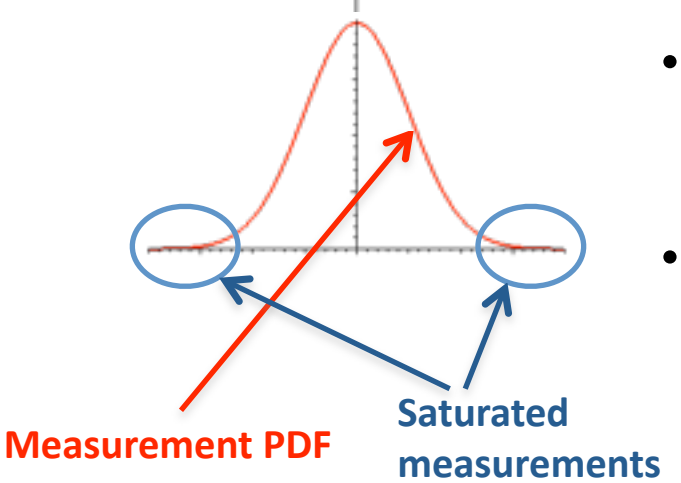

- Given: **Bit budget** B bits/sample, **Signal norm** $\|\mathbf{x}\|_2$
- Set quantization **threshold** G
  - Implicitly sets quantization interval $\Delta=2^{-B+1}G$
  - Implicitly sets saturation rate at $2Q(G/\|\mathbf{x}\|_2)$
- ~~Classical heuristic: Set G large (avoid saturation)~~

  Wrong! Will revisit!
- Note:
  - equivalent to fixing G and varying signal amplification
  - $Q(\cdot)$ denotes the tail of the Gaussian distribution

Reconstruction

- Convex optimization:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\widetilde{\mathbf{y}} - \widetilde{A}\mathbf{x}\|_2 \leq \epsilon$$

- ϵ bounds the quantization error

- Question: What should ϵ be set to?
  - "Straightforward" if no saturation
    - upper bound on quantization error
    - use a quantization noise model
    - l_p, $p>2$ constraints instead [e.g. Jaques et al. '09, others]
  - Saturation can make error **large** or **unbounded**

- How to deal with saturated measurements?
  - Treat them as **any other measurement**
  - Treat them as missing data and **ignore them**
  - Treat them as providing **different information**

Dropping Measurements

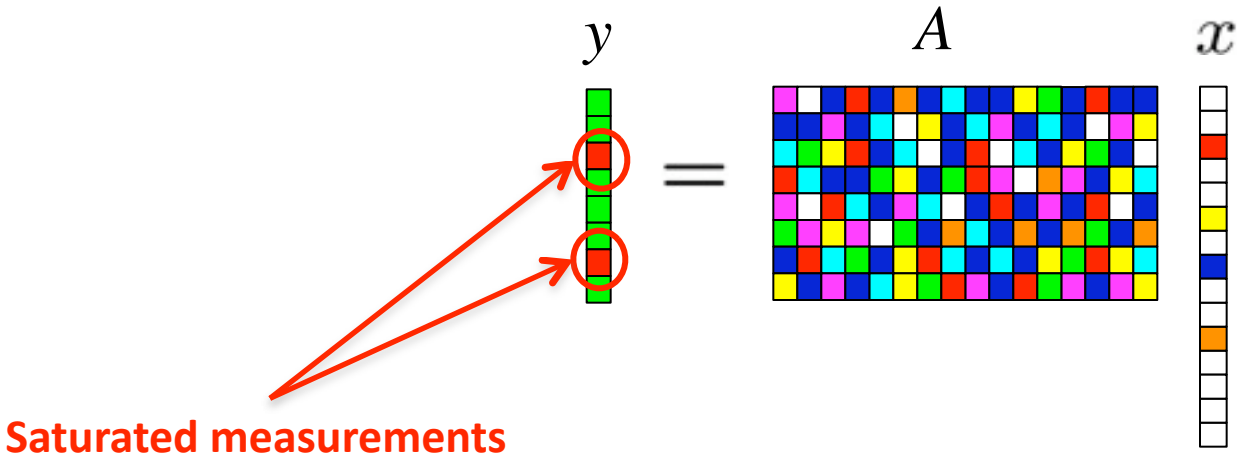

Dropping Measurements

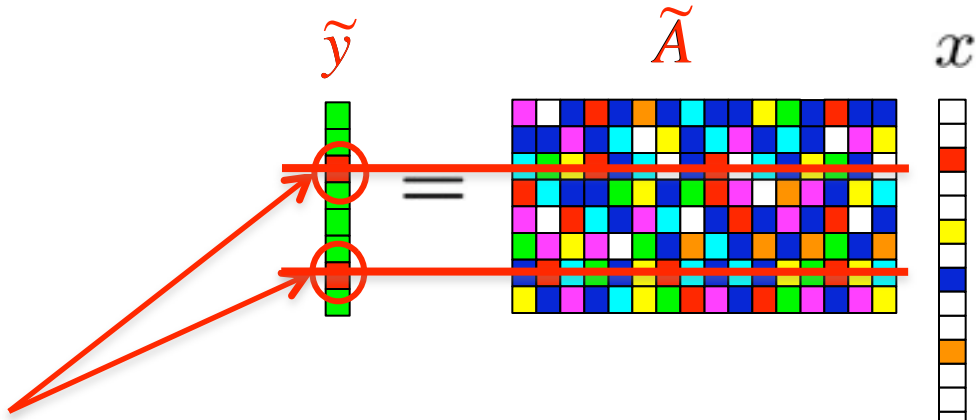

Saturated measurements

- Measurements are **democratic** [Davenport, Laska, Boufounos, Baraniuk]
  - They are all equally important
  - We can **drop** the **saturated** ones

- Run the standard optimization on the remaining ones

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\widetilde{\mathbf{y}} - \widetilde{A}\mathbf{x}\|_2 \leq \epsilon$$

- The remaining $\widetilde{A}$ still satisfies RIP (as long as we don't drop too many)

Quantization Trade-Off

- **Increasing** saturation level G:
  - We **keep more measurements**, i.e., better reconstruction
  - We have **larger quantization interval**, i.e., worse reconstruction
- **Decreasing** saturation level G:
  - We **drop more measurements**, i.e., worse reconstruction
  - We have **smaller quantization interval**, i.e., better reconstruction
- A trade-off exists, **contrary** to the classical heuristics

Example (experimental)

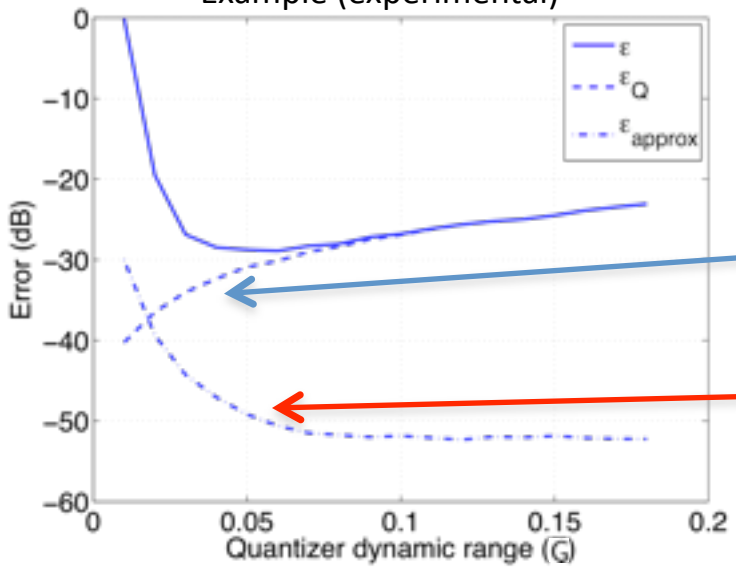

Quantization error

Approximation error
(due to dropping measurements)

Experimental Results

Conventional approach (treat saturated data as measurements)

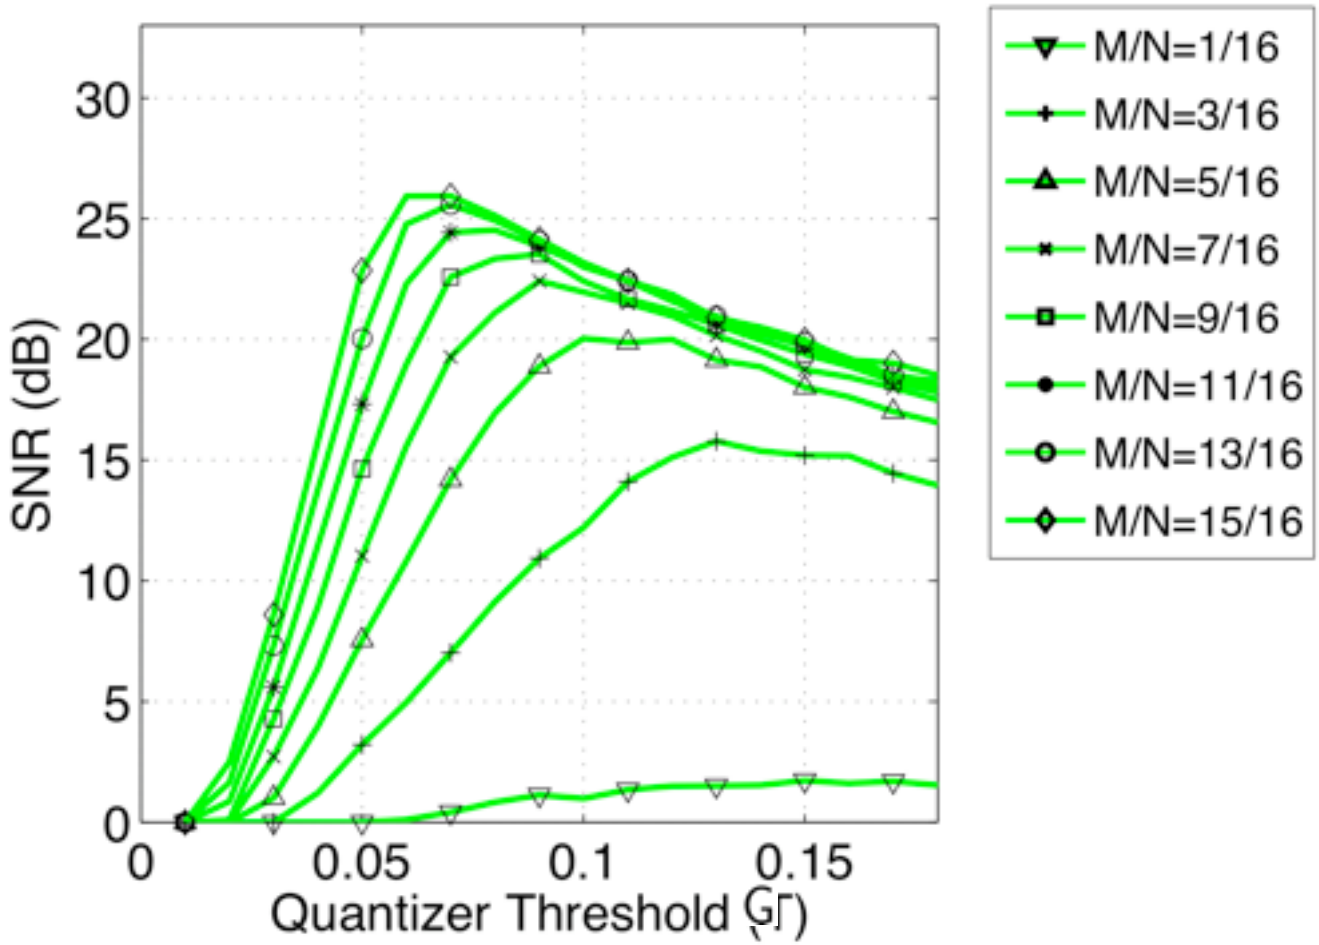

(a)

Experimental Results

New approach (discard saturated measurements)

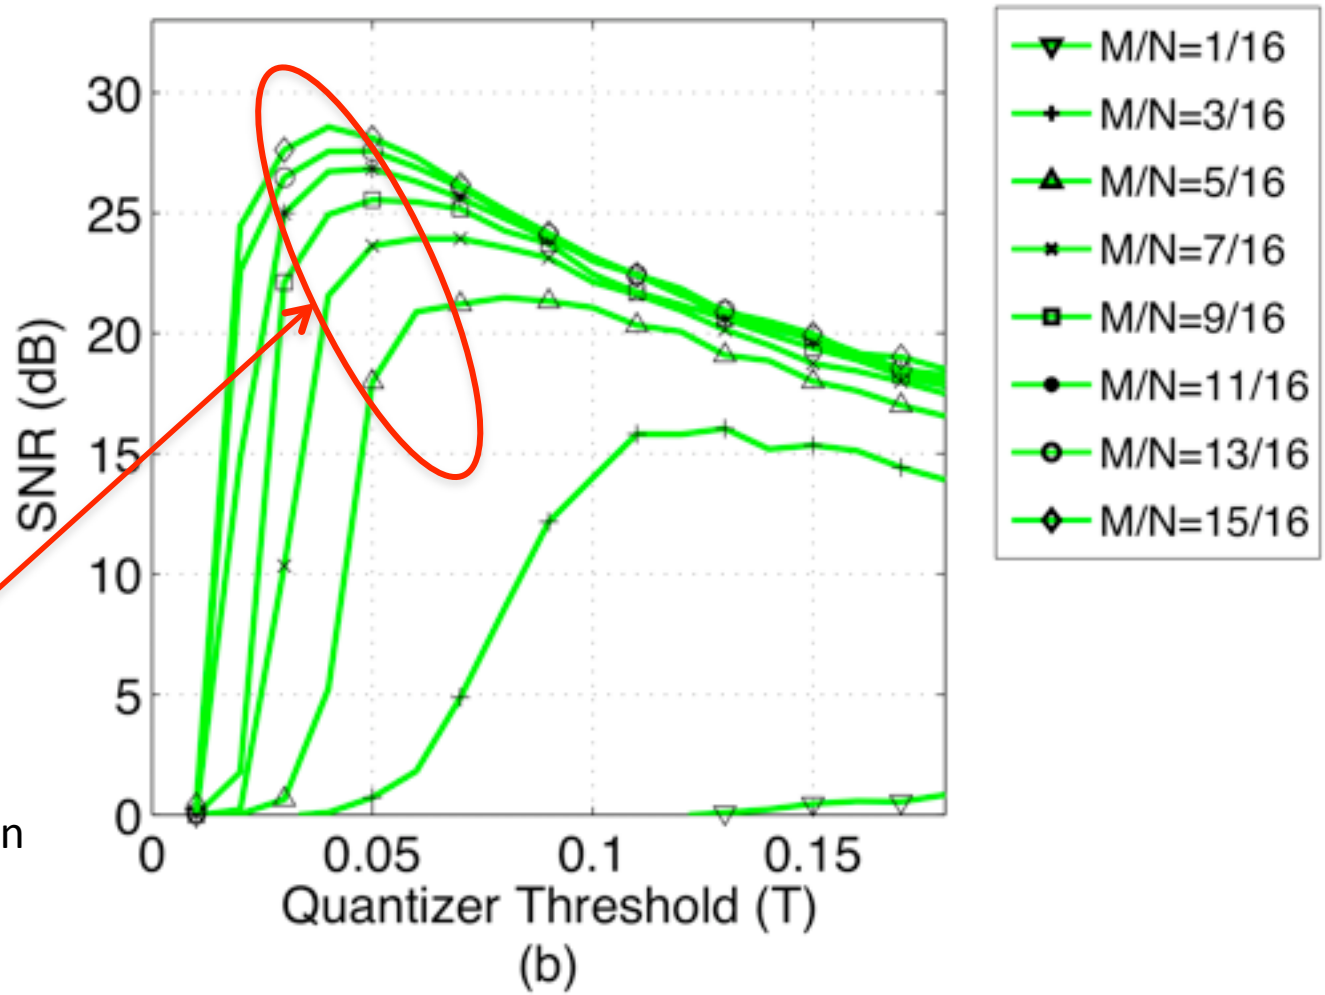

Performance improves
Best performance when
saturation occurs

Experimental Results

Side by side comparison

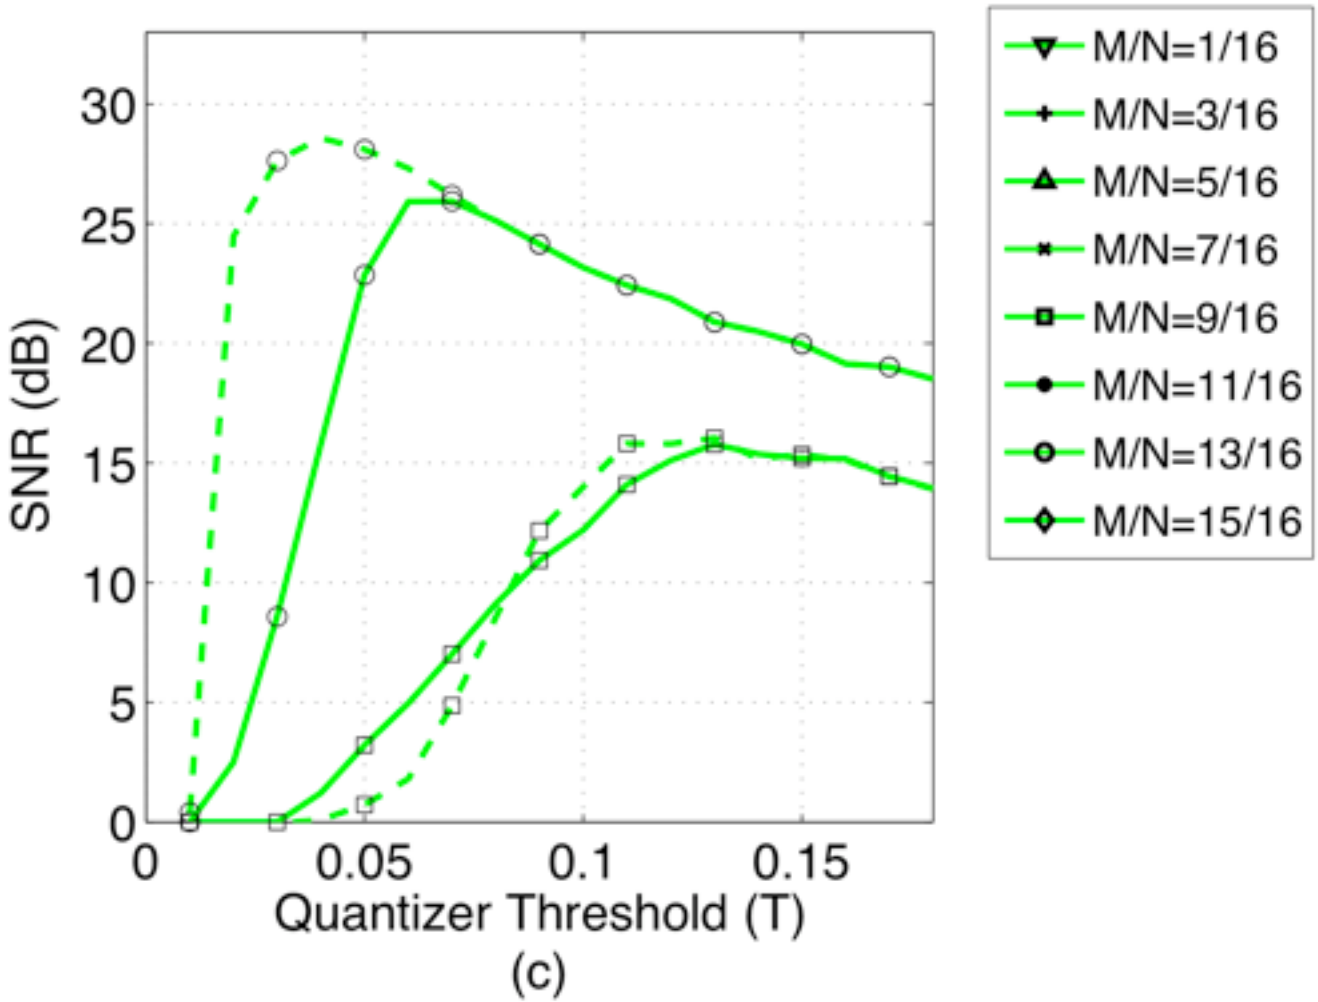

Alternative: Exploit Implicit Information

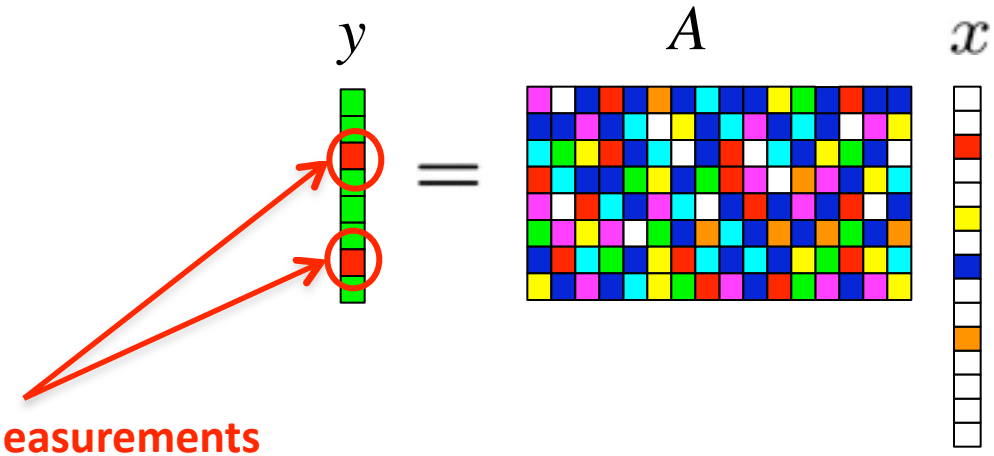

Saturation provides information:
The measurement is **larger** than G.

Treat measurement as a **constraint**!

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1$$

$$\text{s.t. } \|\widetilde{\mathbf{y}} - \widetilde{\Phi}\mathbf{x}\|_2 \leq \epsilon \qquad \text{Unsaturated}$$

$$\Phi^{+}\mathbf{x} \geq G - \Delta/2 \qquad \text{Positive Saturation}$$

$$\Phi^{-}\mathbf{x} \leq -G + \Delta/2 \qquad \text{Negative Saturation}$$

Experimental Results

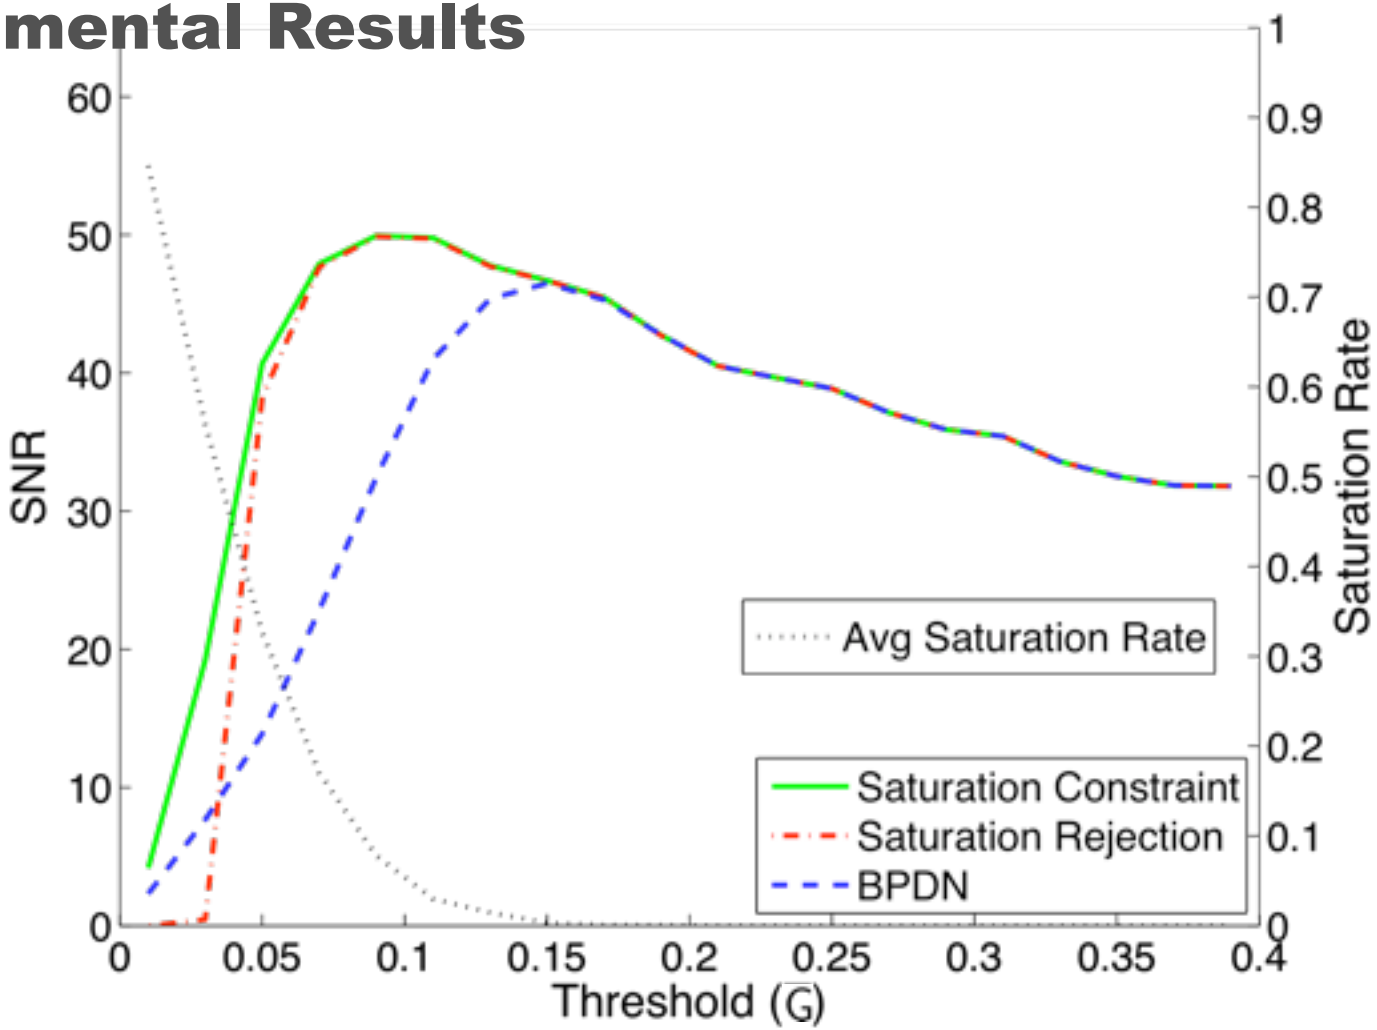

Experimental Results

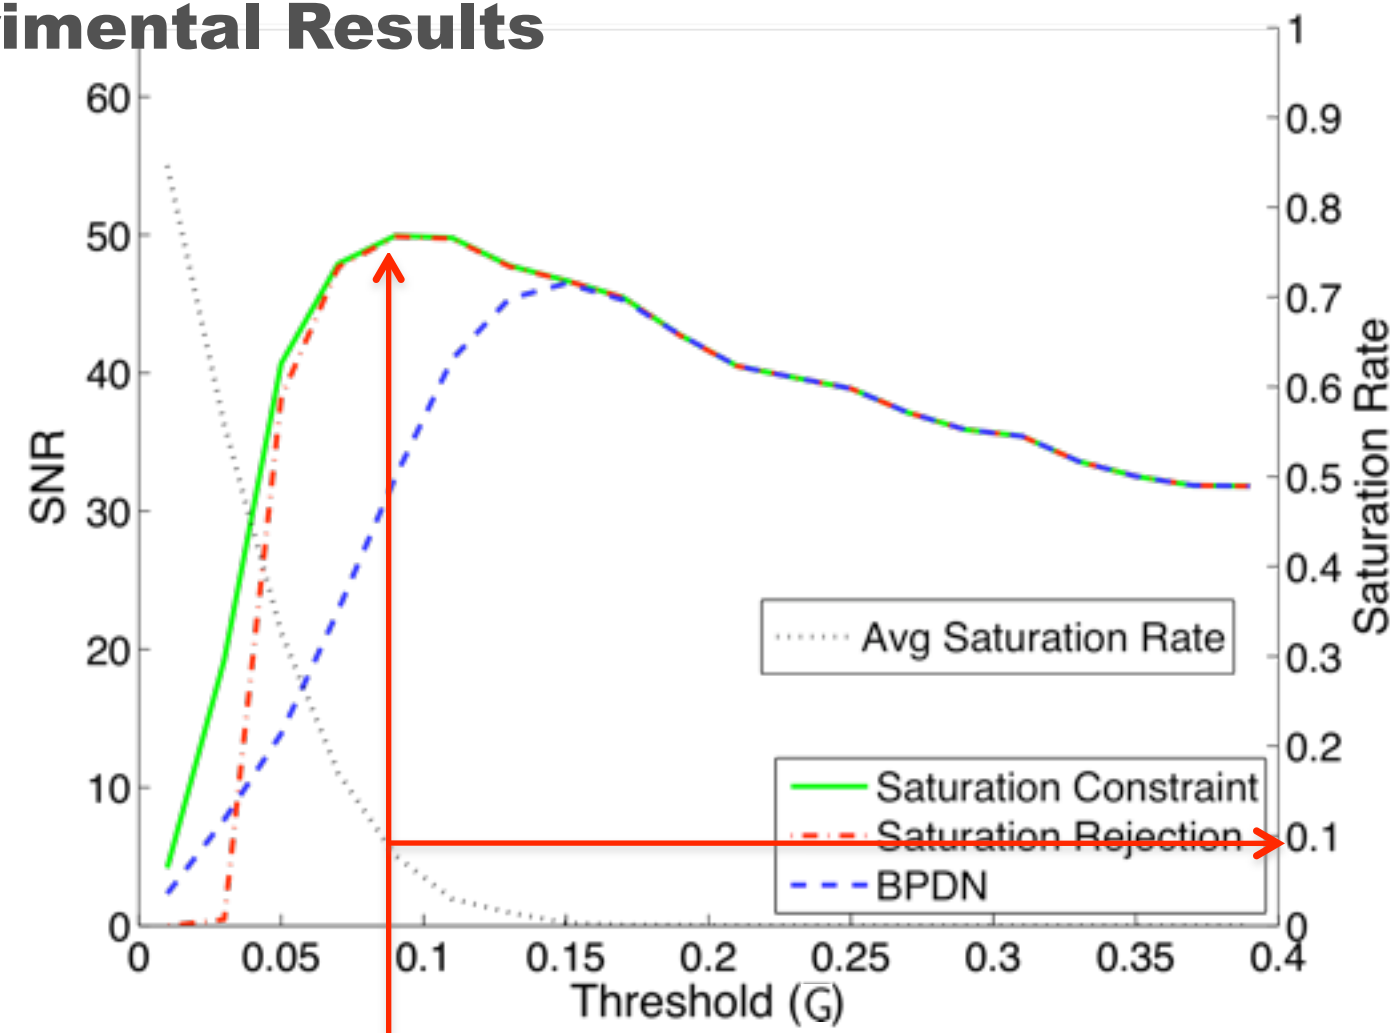

Note: optimal performance **requires** 10% saturation

Experimental Results

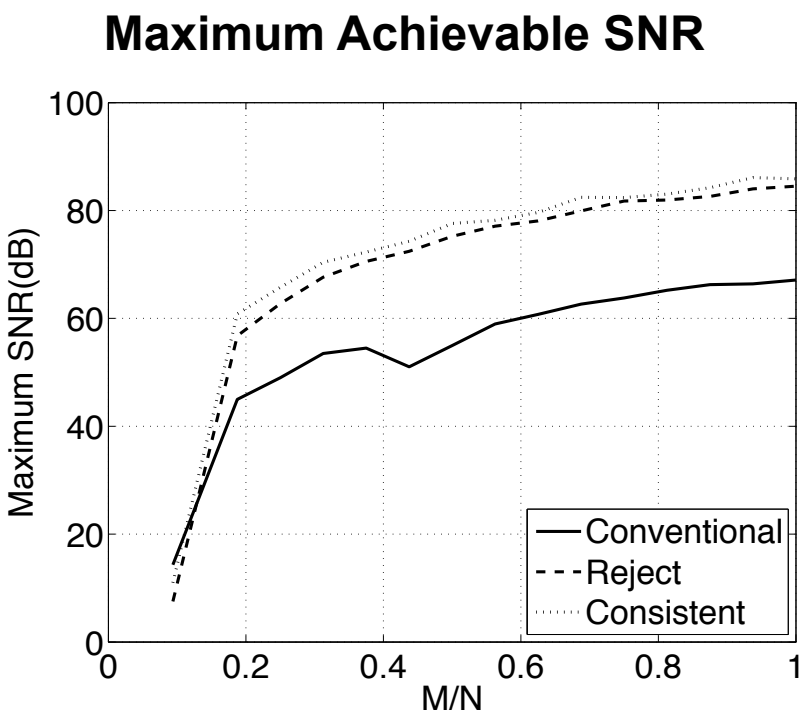

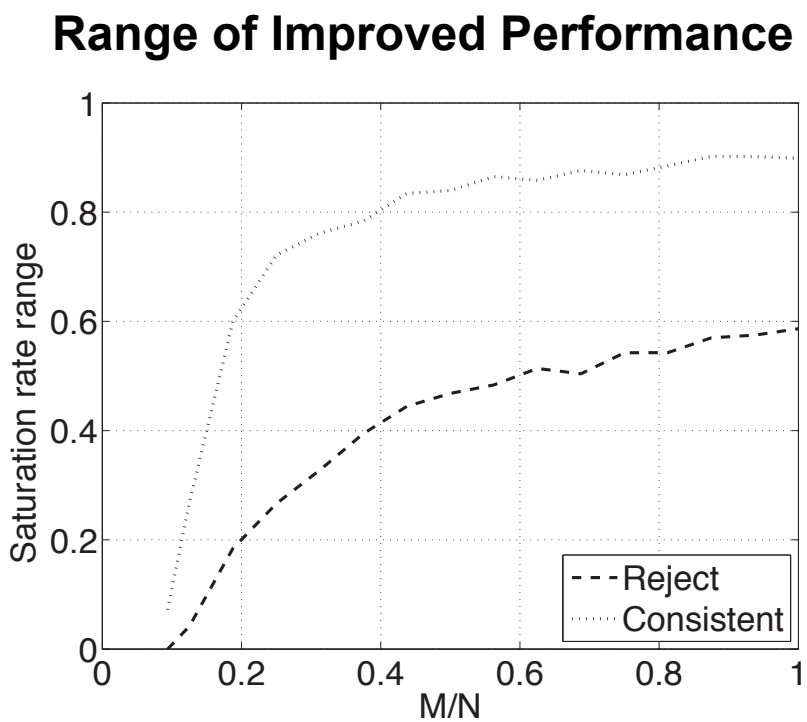

Consistent Reconstruction:

Only **slightly better** SNR performance

More **robust** to saturation variability

1-BIT COMPRESSED SENSING

1-bit Compressive Sensing

Q: Can we quantize the measurements to 1-bit?

$$y = \mathrm{sign}(Ax)$$

and recover using a greedy algorithm?

Motivation:

1-bit quantization is easy, fast and cheap

Key issues:

Need to identify the cost function

Can only recover within a positive scaling factor

$$\mathrm{sign}(Ax) = \mathrm{sign}(Acx), \;\; c \geq 0$$

1-Bit Information

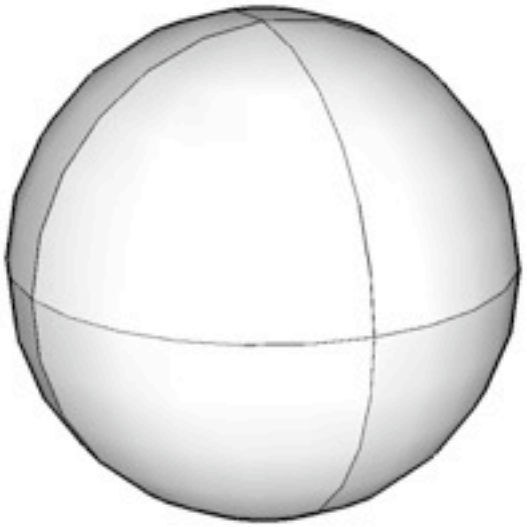

1-Bit Information

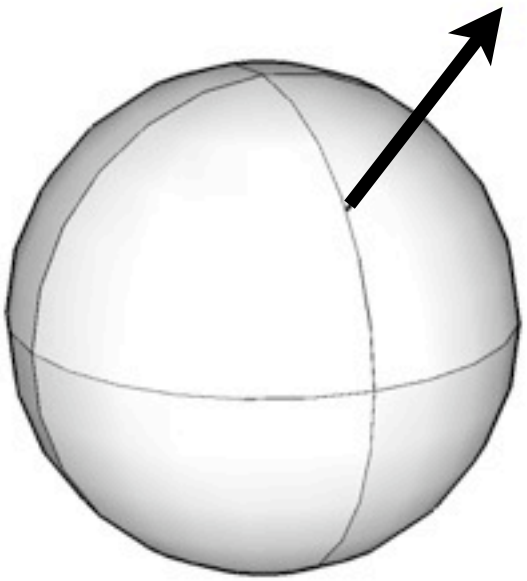

1-Bit Information

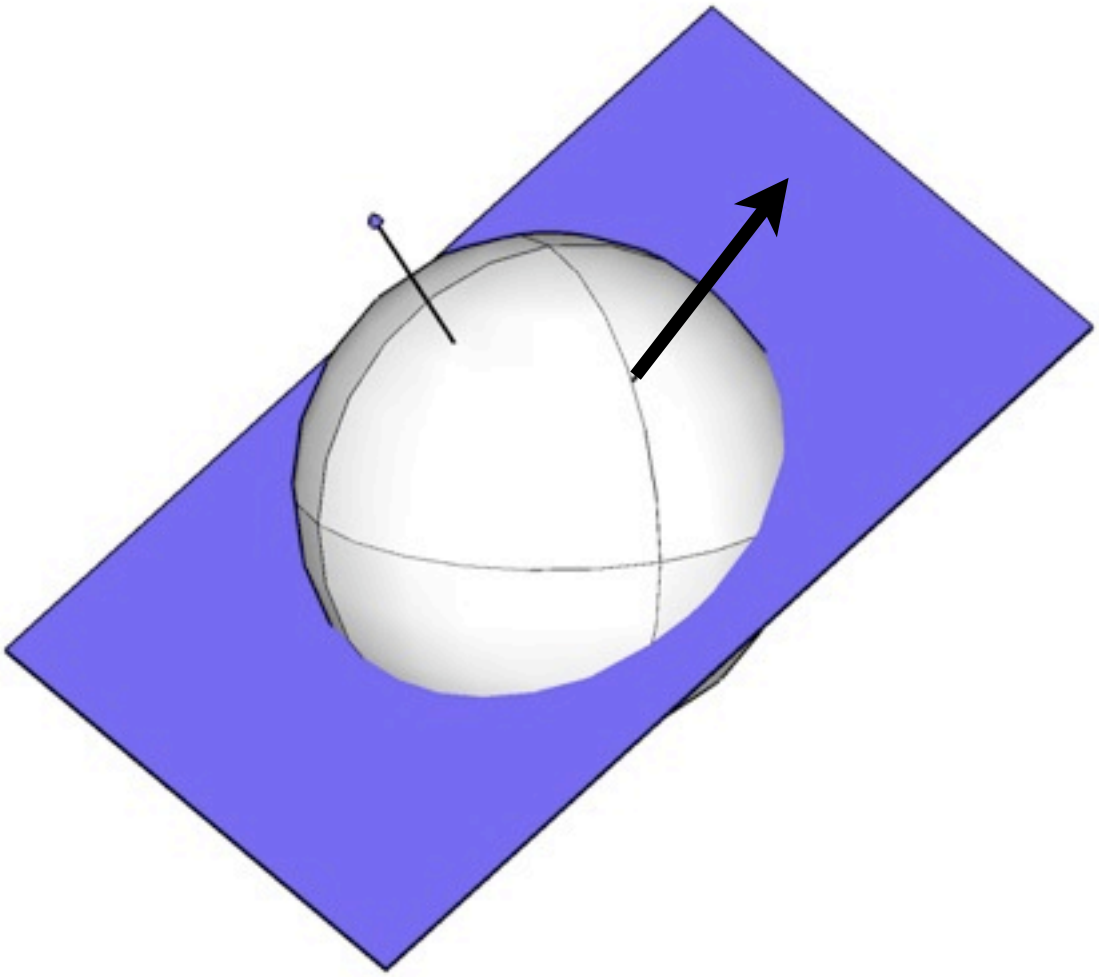

1-Bit Information

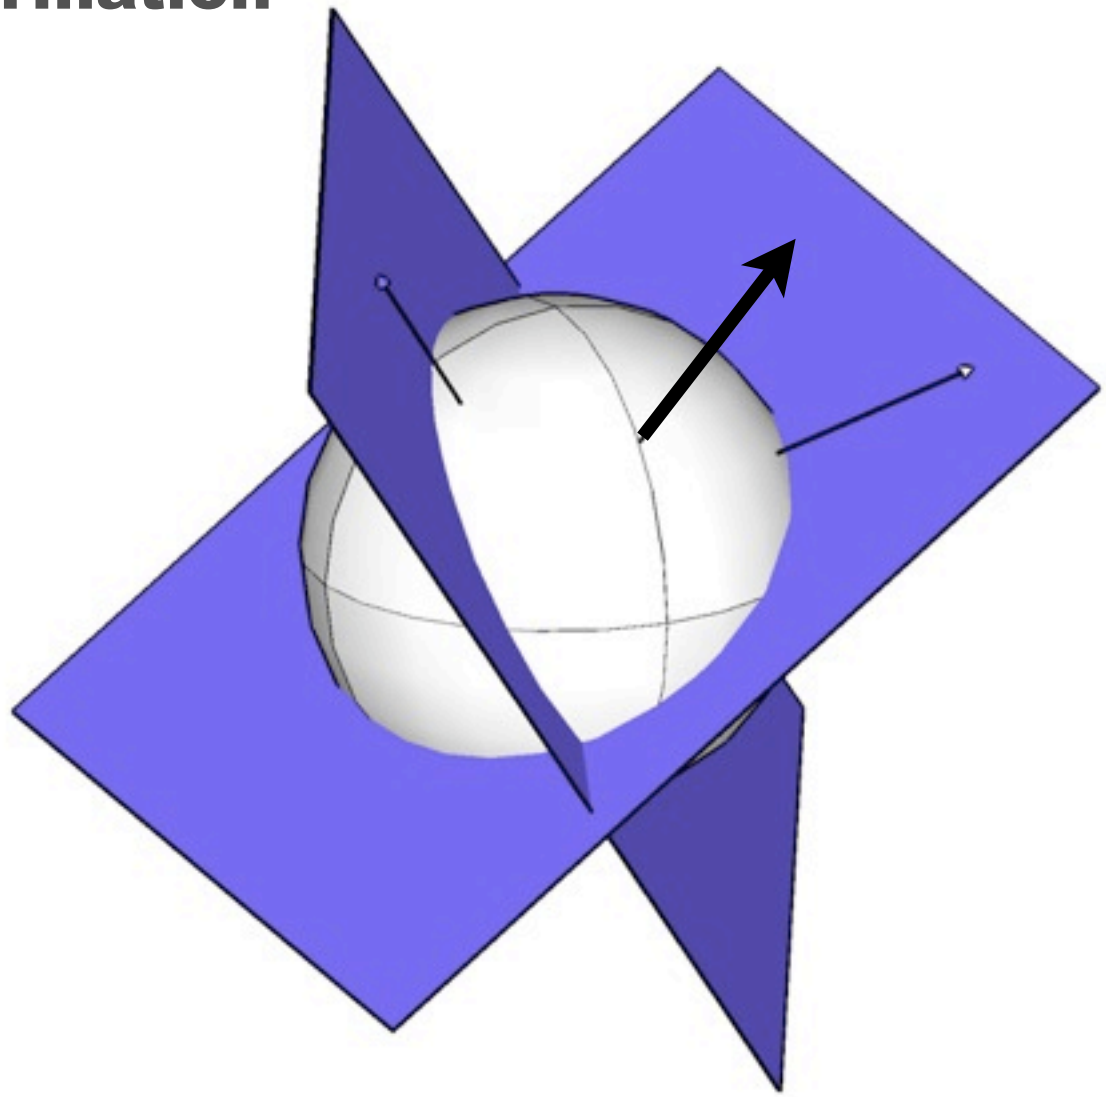

1-Bit Information

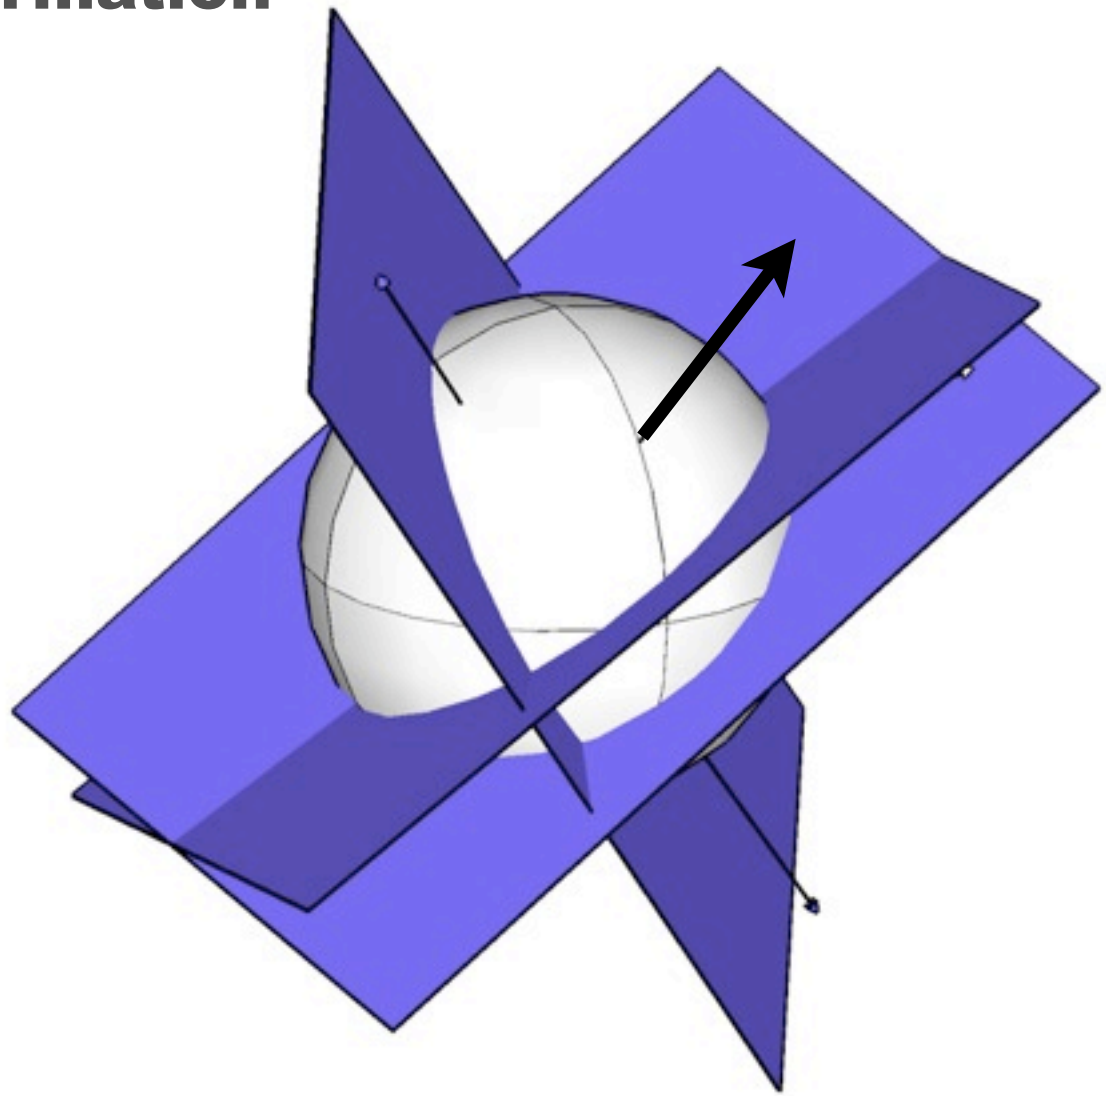

1-Bit Information

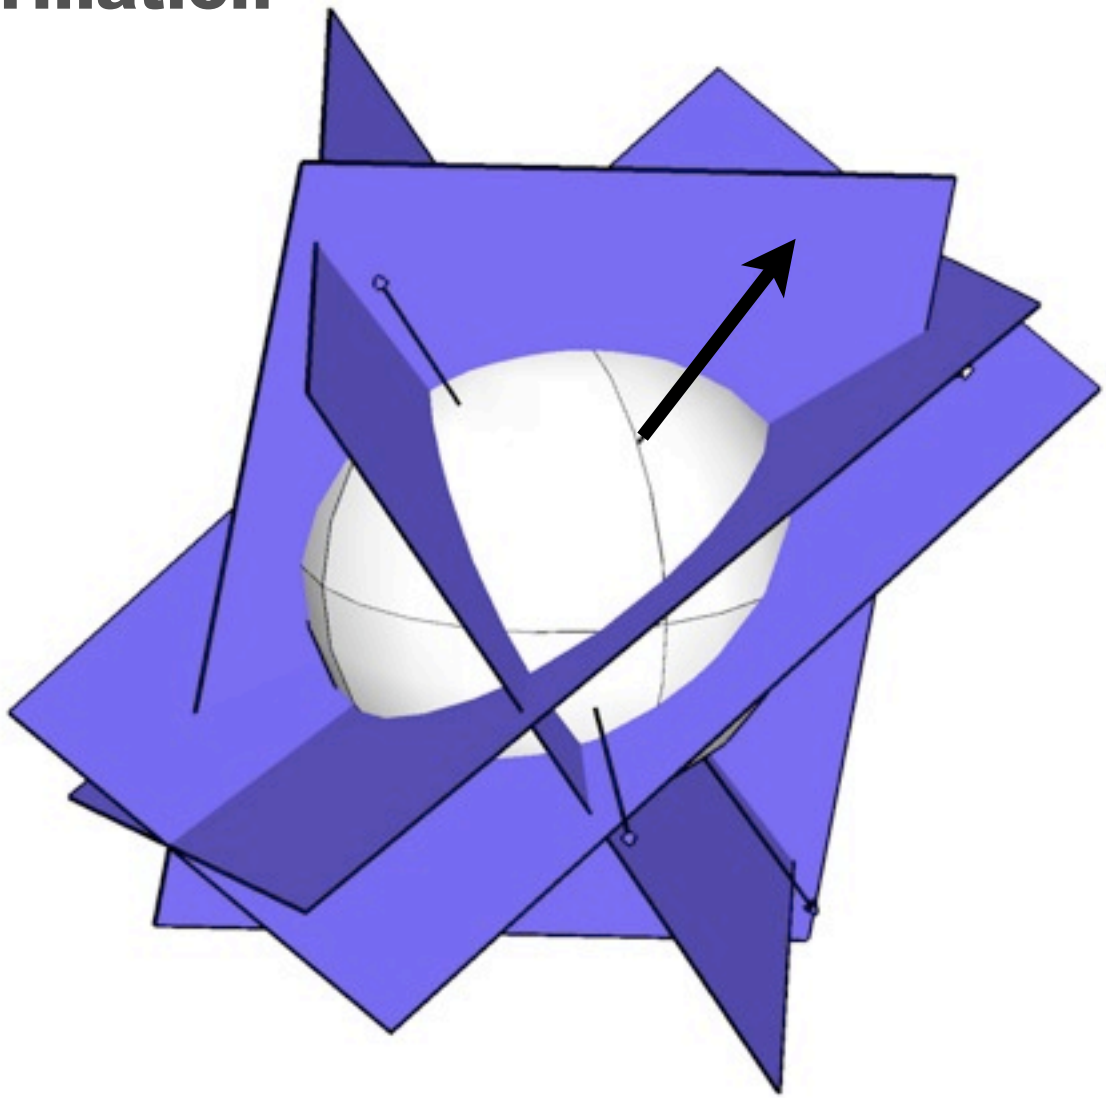

1-Bit Information

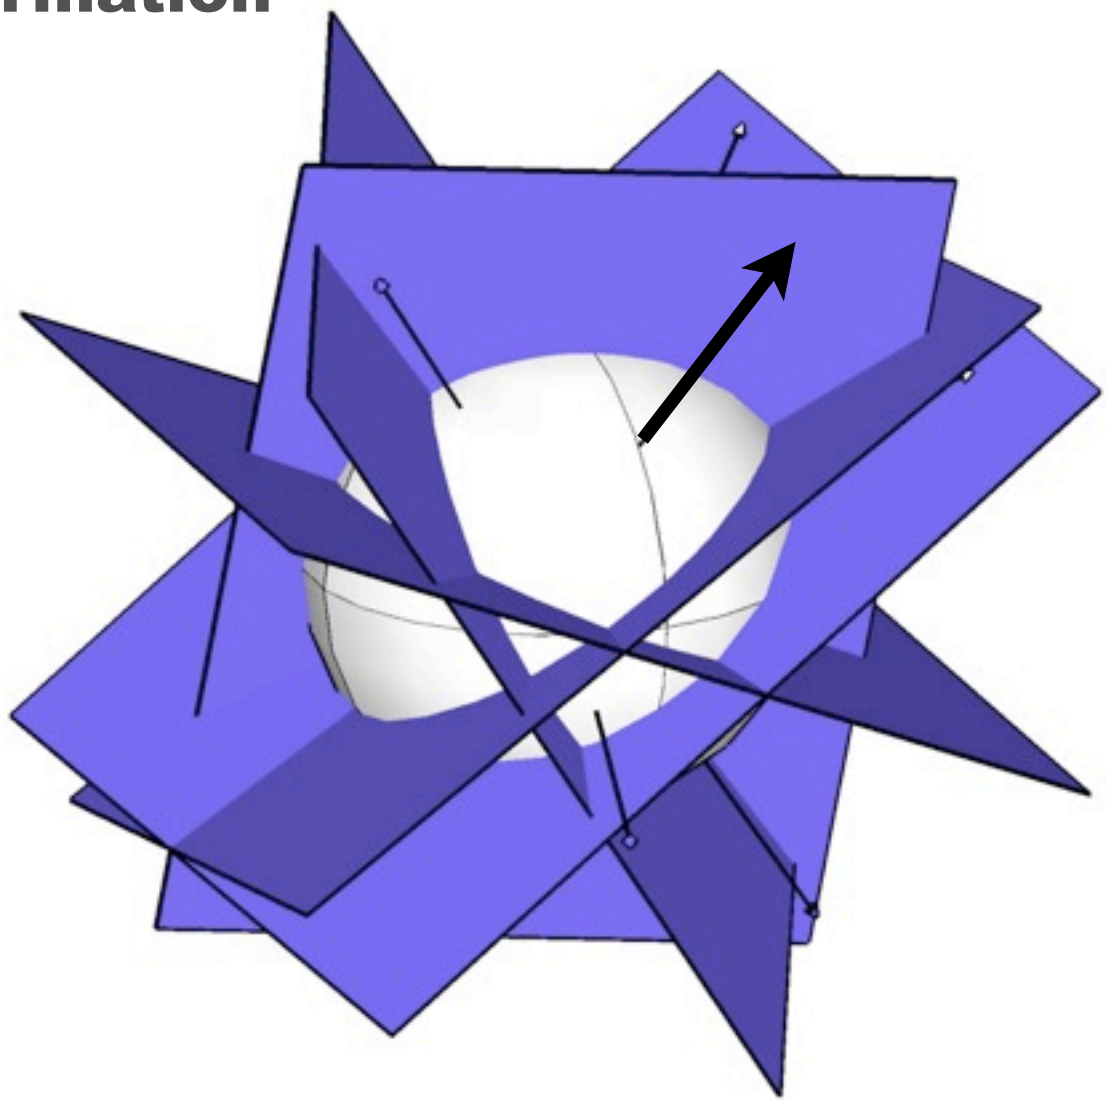

1-Bit Information

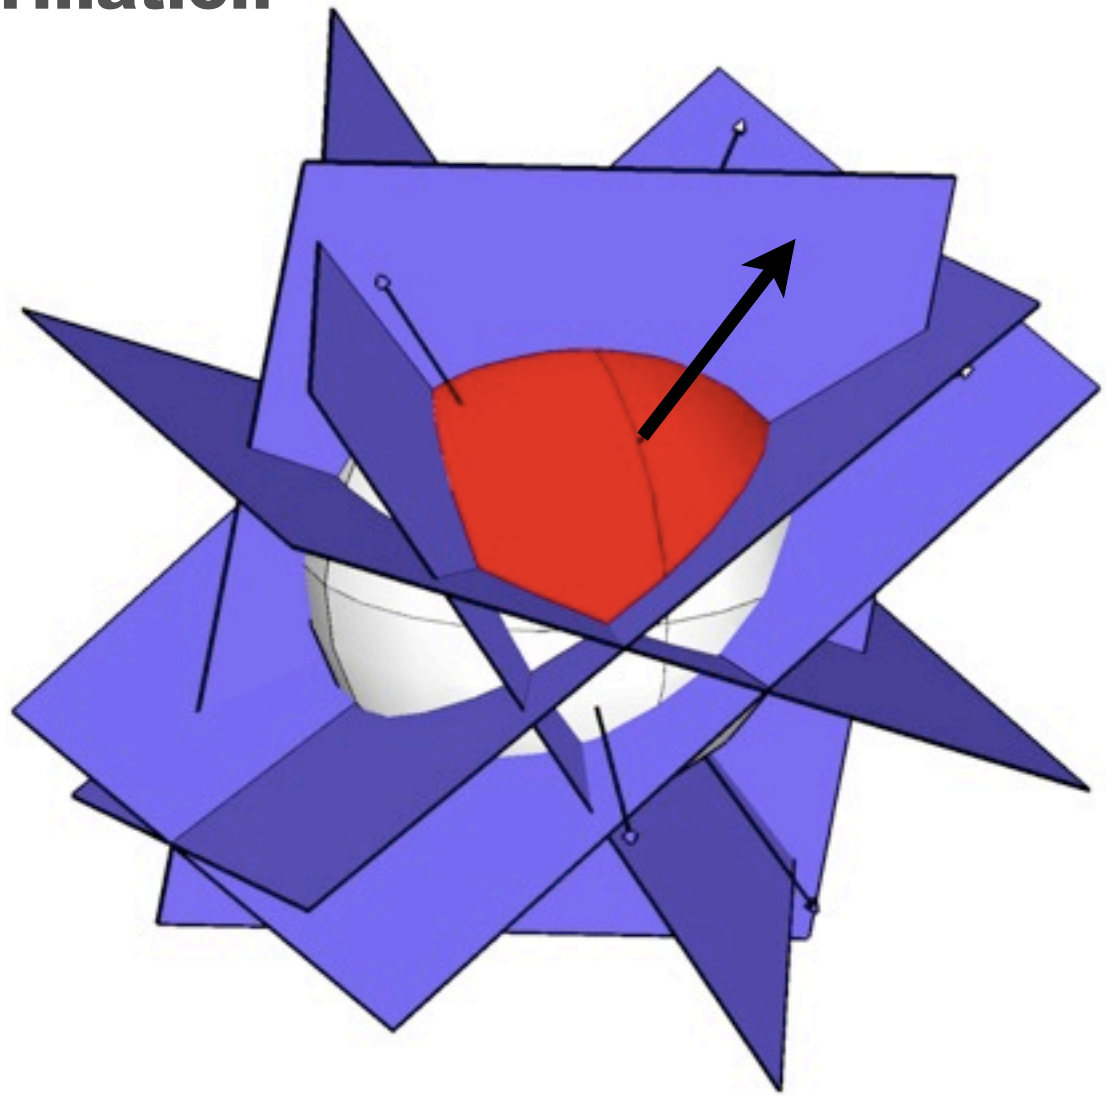

1-Bit Information

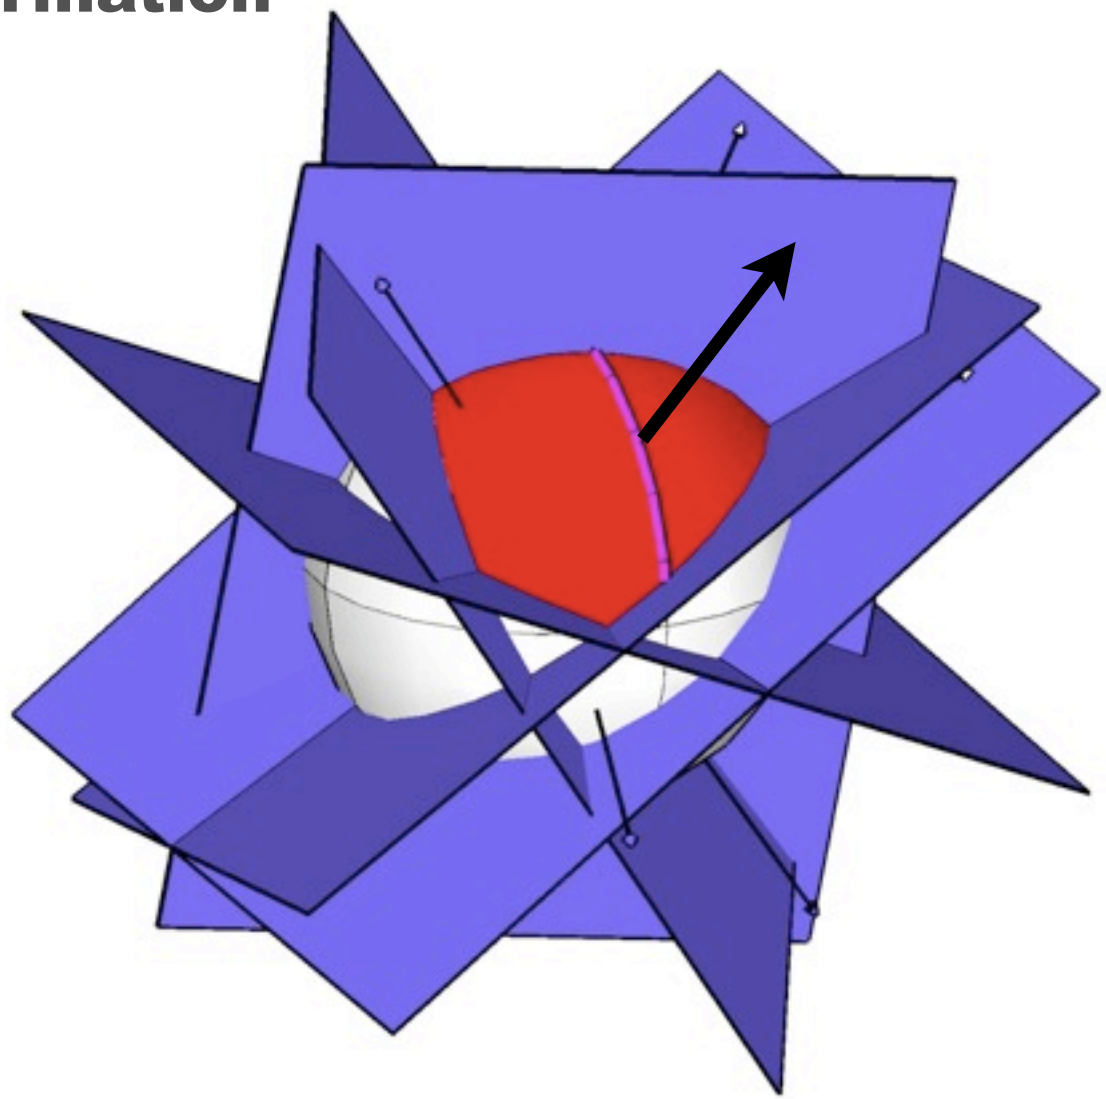

1-bit CS Cost Function

- Signs of solutions must match with data:

$$\mathrm{sign}(A\widehat{x}) = \mathrm{sign}(Ax) = y$$

$$\text{i.e. } y_i \times (A\widehat{x})_i \geq 0$$

- Penalize mismatches quadratically:

$$\sum_i f\left(y_i \times (A\widehat{x})_i\right)$$

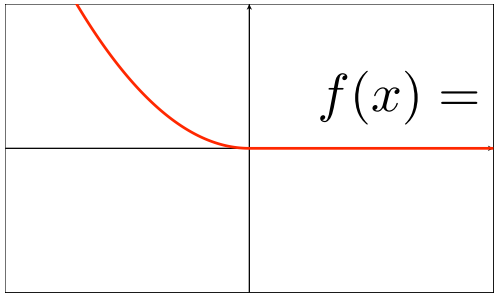

$$f(x) = \begin{cases} x^2 & x \leq 0 \\ 0 & x > 0 \end{cases}$$

- Also need to normalize after each iteration
  - Makes algorithm non-convex

Matching Sign Pursuit

- Iterative greedy algorithm, similar to CoSaMP
- Maintains running signal estimate and its support T.
- MSP Iteration:
  - Identify **sign violations** $r = (\mathrm{diag}(y)A\widehat{x})^-$
  - Compute **proxy** $p = A^T r$
  - Identify **support** $\Omega = \mathrm{supp}(p|_{2K}) \cup T$
  - **Consistent Reconstruction** over support

$$b|_\Omega = \arg\min_x \left\| (\mathrm{diag}(y)Ax)^- \right\|_2^2,$$

$$\text{s.t. } \|x\|_2 = 1 \text{ and } x|_{T^c} = 0$$

  - Truncate, normalize, and **update** estimate $\widehat{x} \leftarrow \frac{b|_K}{\|b|_K\|_2}$

Matching Sign Pursuit (MSP)

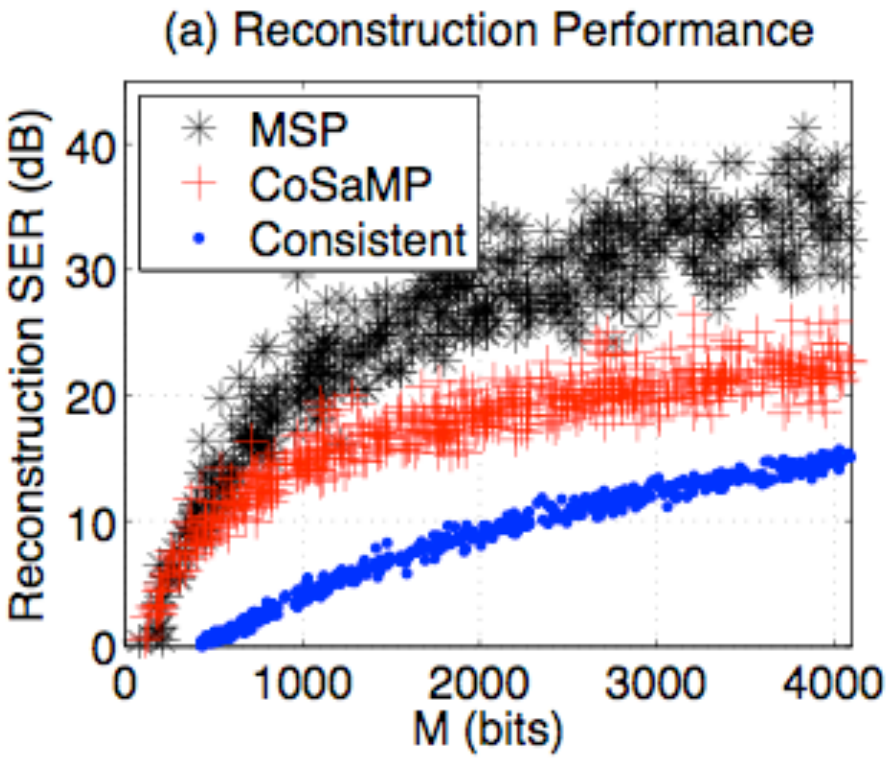

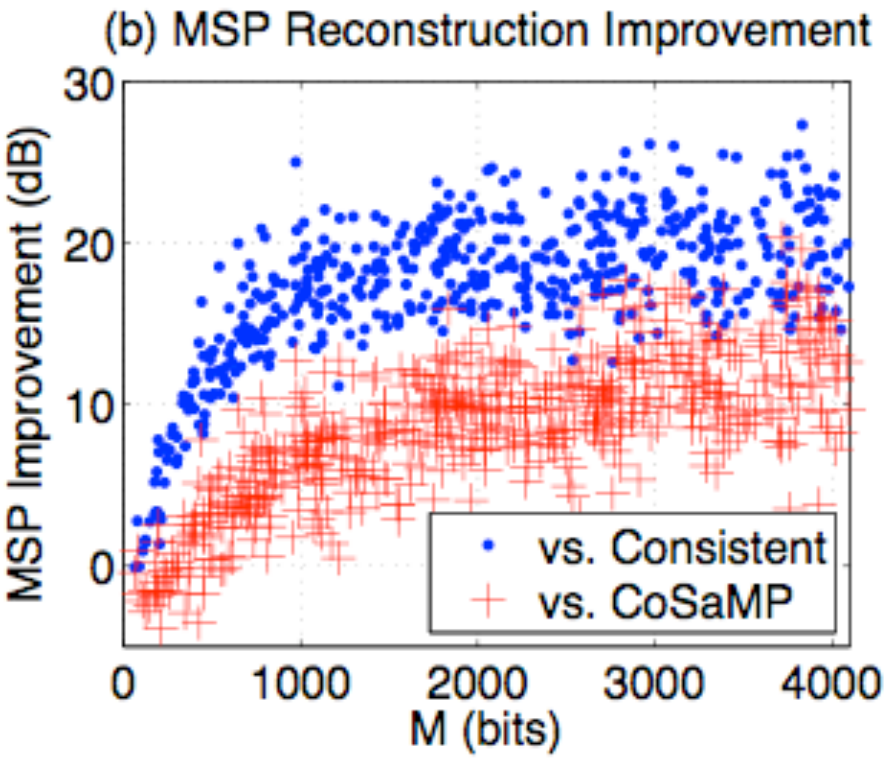

NON-LINEAR DISTORTION

CS with Distorted Measurements

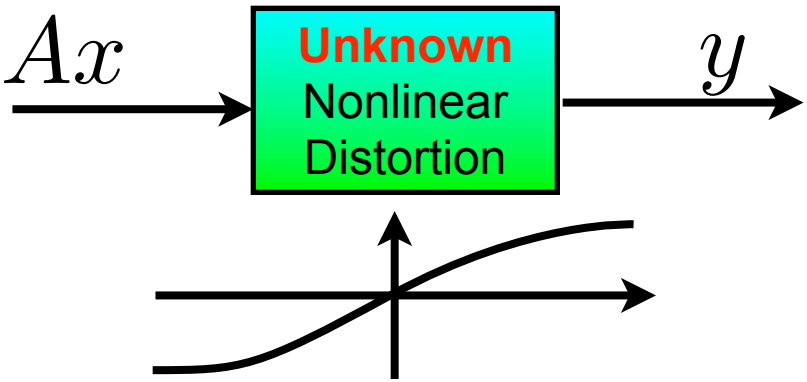

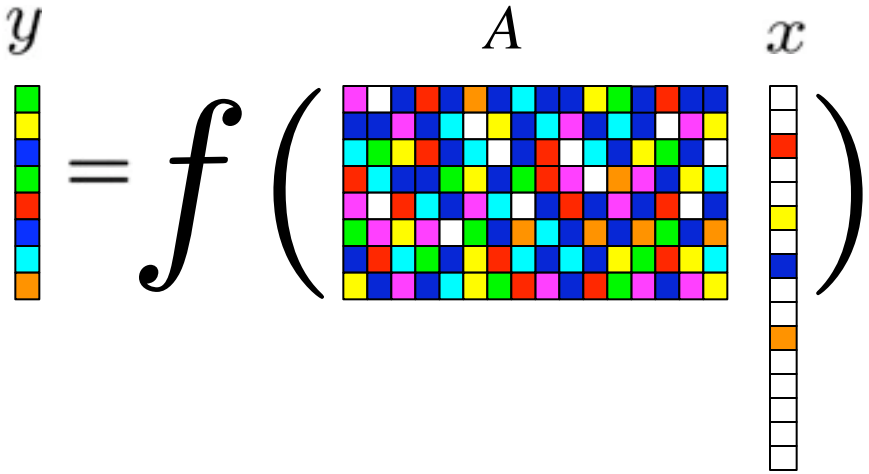

- Distortion is **unknown** but monotonic
- Monotonicity **preserves** the **order** of the measurements

Measurement Ordering

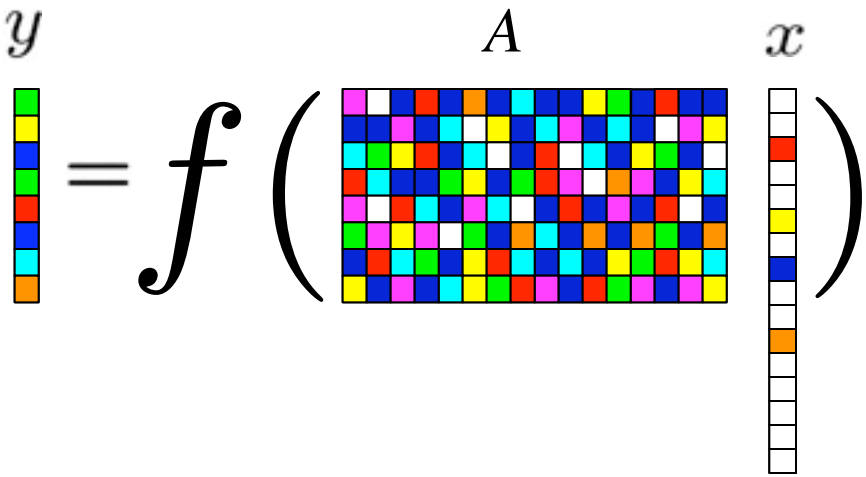

Measurement Ordering

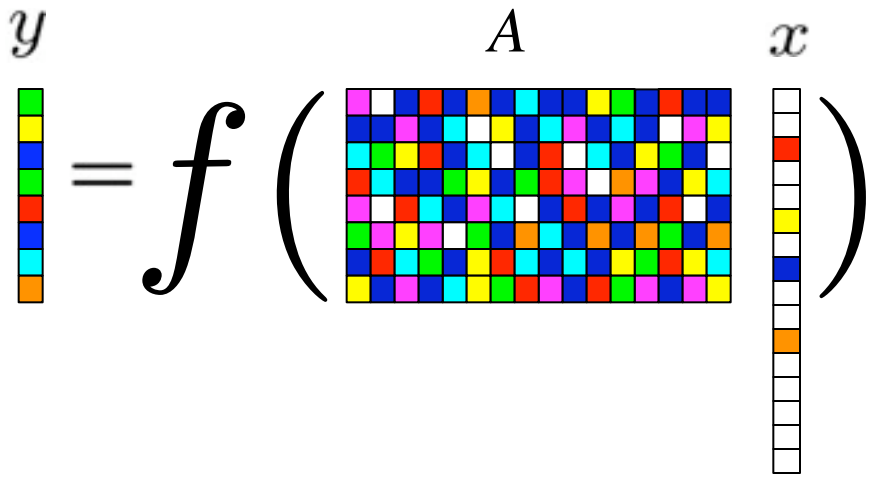

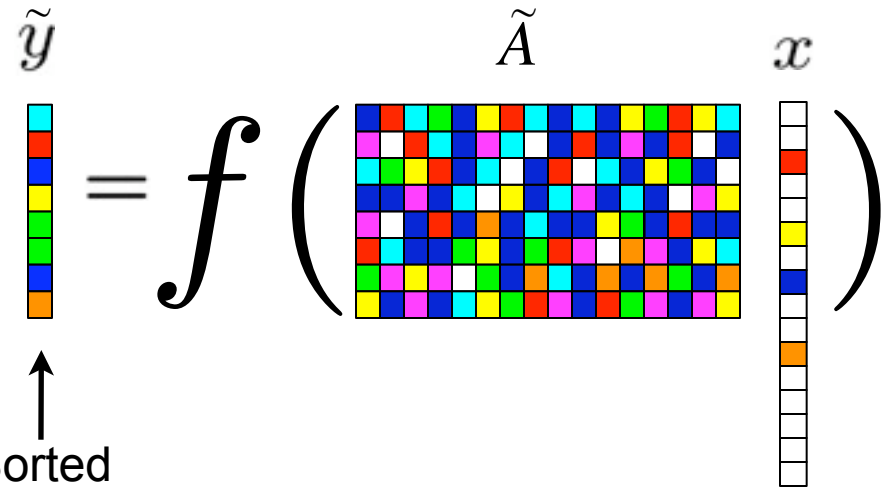

Sorted
Measurements

Method 1: MSP on Measurement Difference

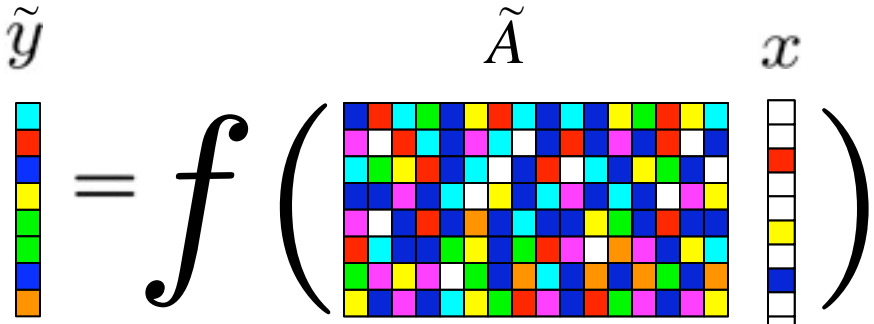

Sorted Measurements

Method 1: MSP on Measurement Difference

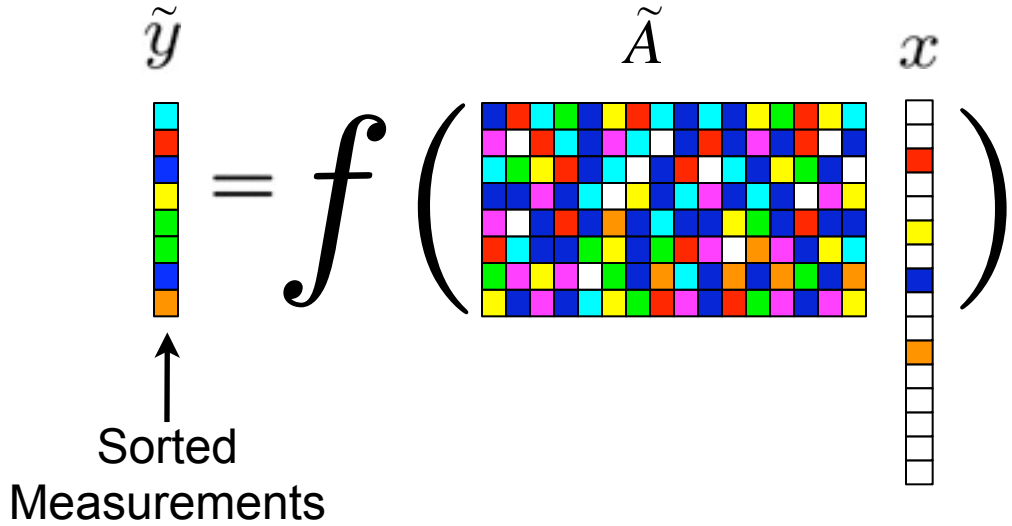

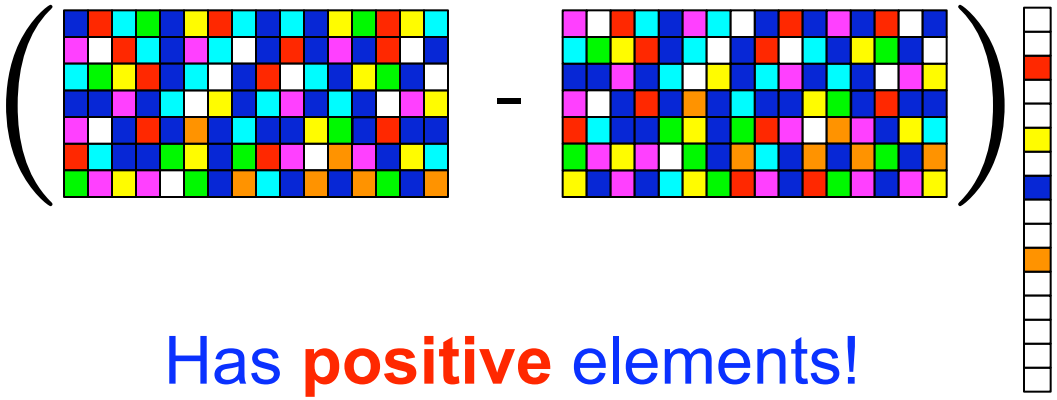

Has **positive** elements!

Use **MSP** (1-bit CS) to estimate x.

Alternative Approach: Measurement Substitution

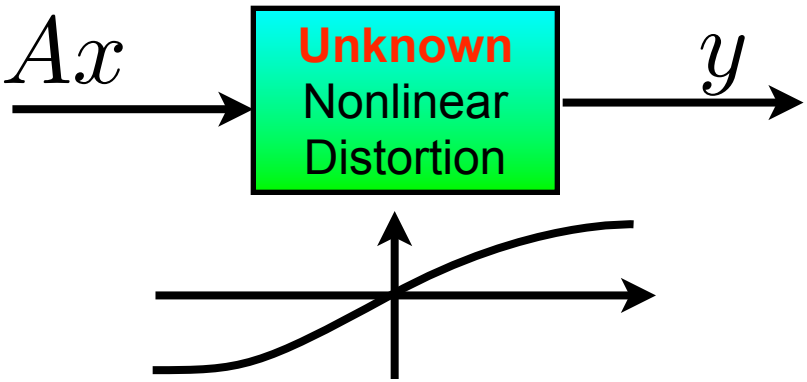

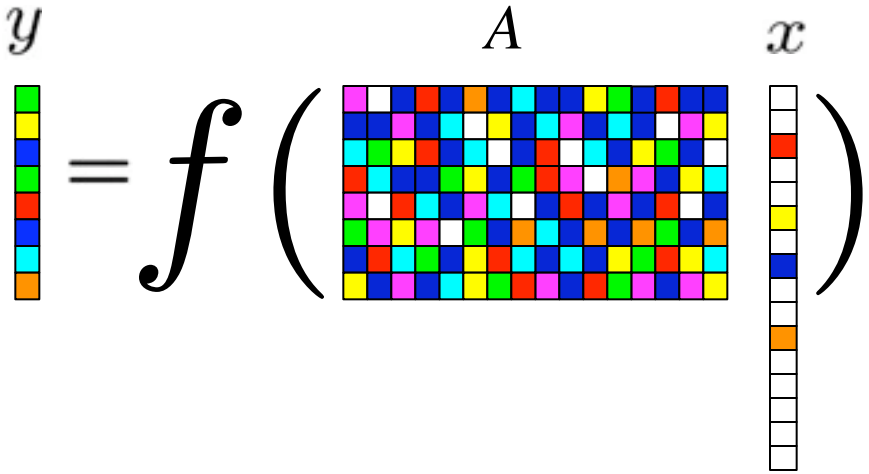

- A is random; Ax is (approximately) normally distributed
- Use order statistics to estimate the measurements
- Method 2:
  - Order measurements
  - Replace measurements with estimate
  - Use standard algorithms

Measurement Ordering

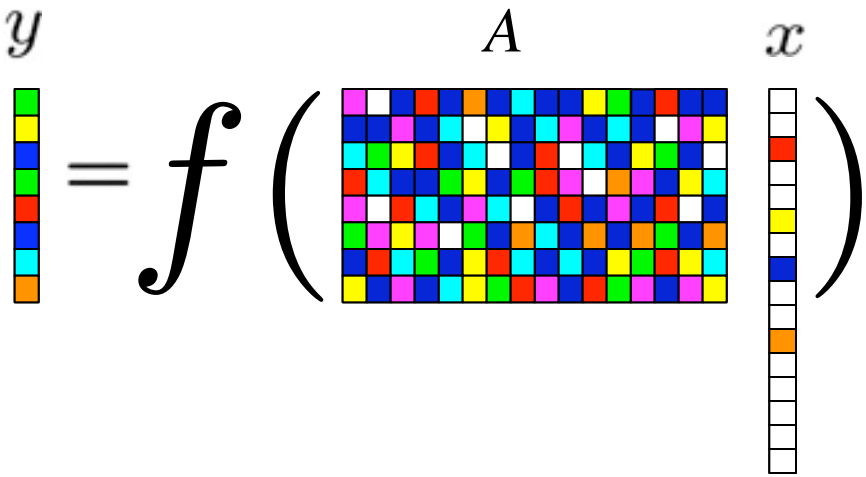

Measurement Ordering

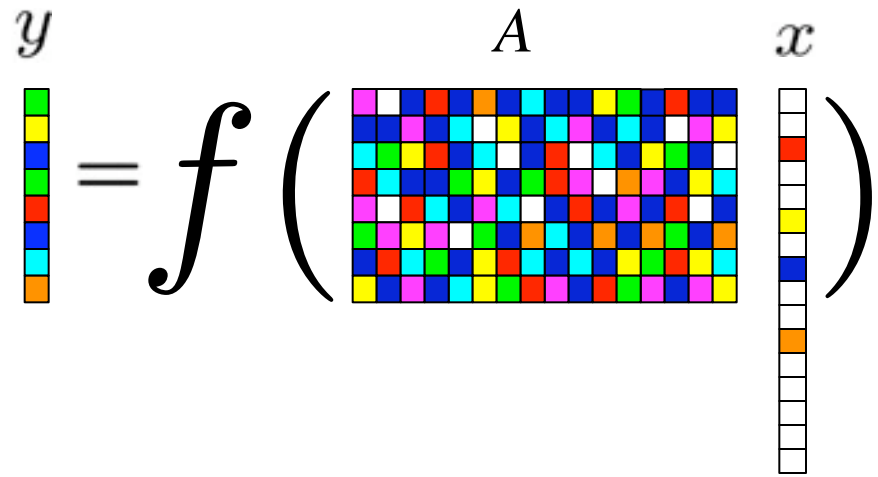

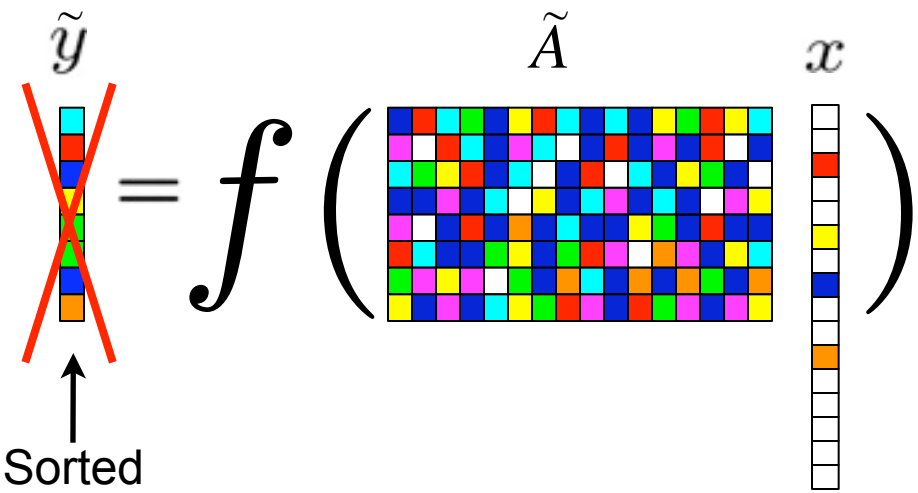

Sorted
Measurements
Ignore their value!

Measurement Substitution: Order Statistics

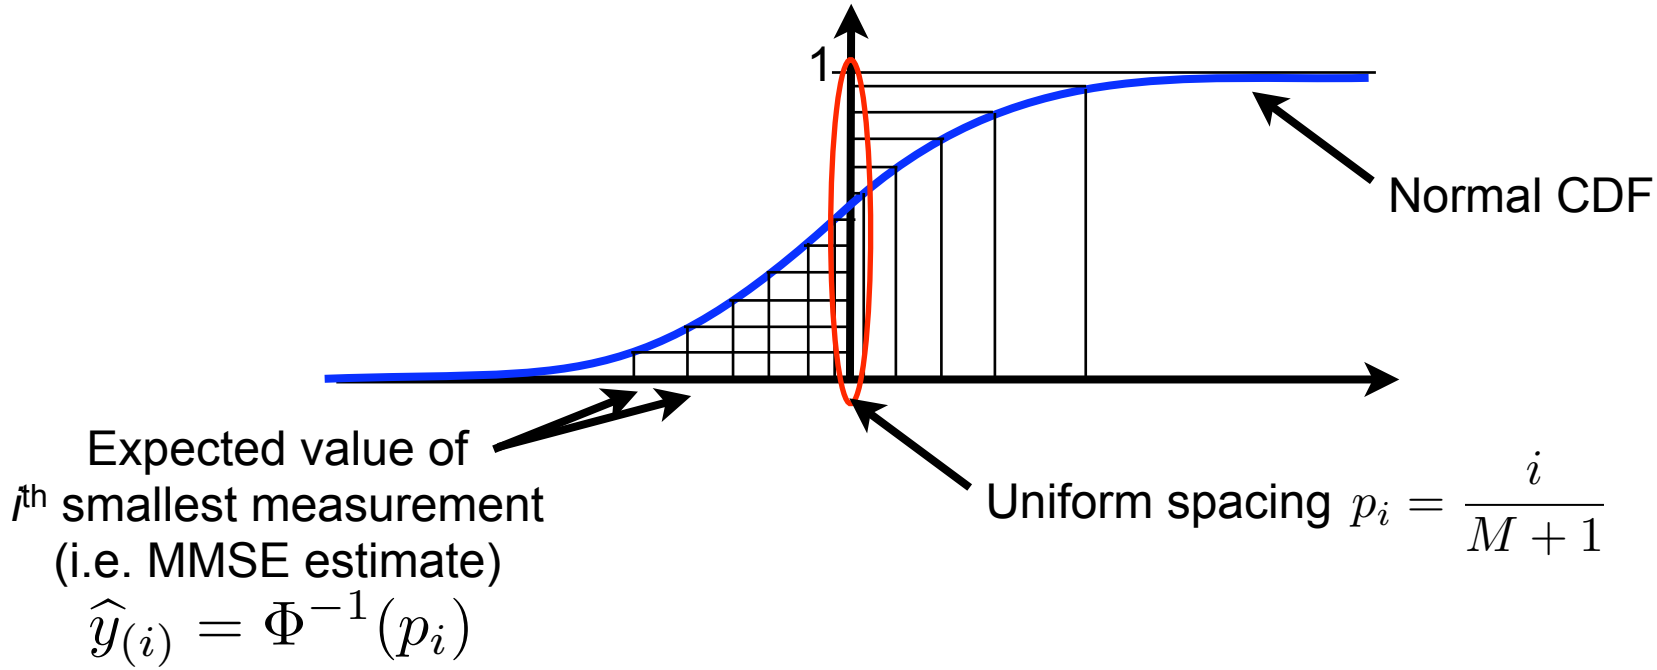

- Solve system for $\widehat{y} = \widetilde{A}x$
- Uses standard CS solvers (convex)
- Can be used in non-CS problems
- Slightly worse perfromance

Performance Evaluation

$$y = \tanh(Ax) + \text{noise}$$

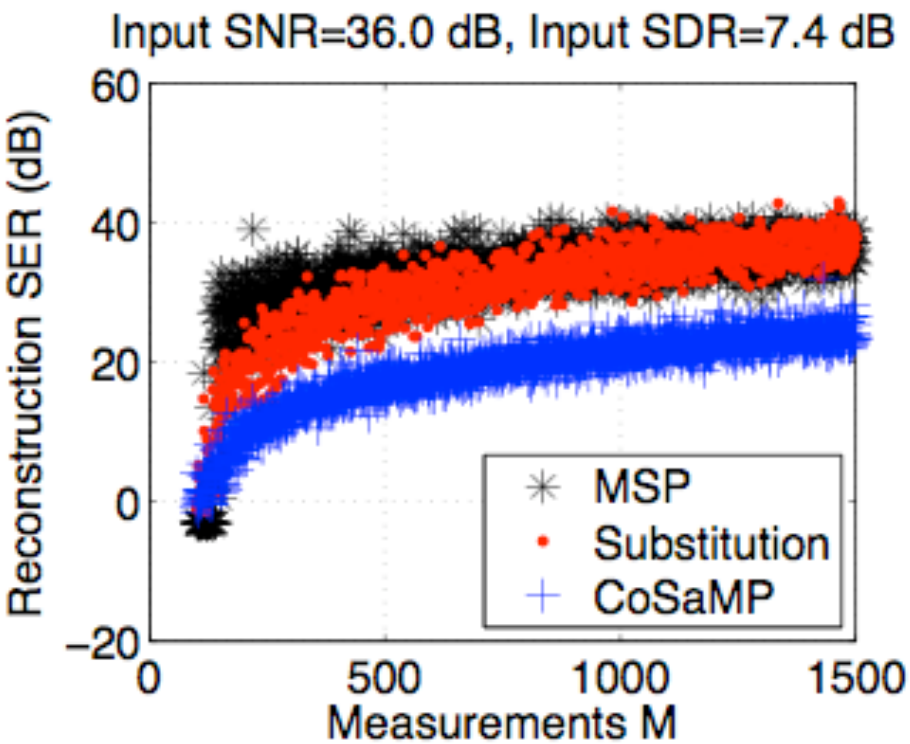

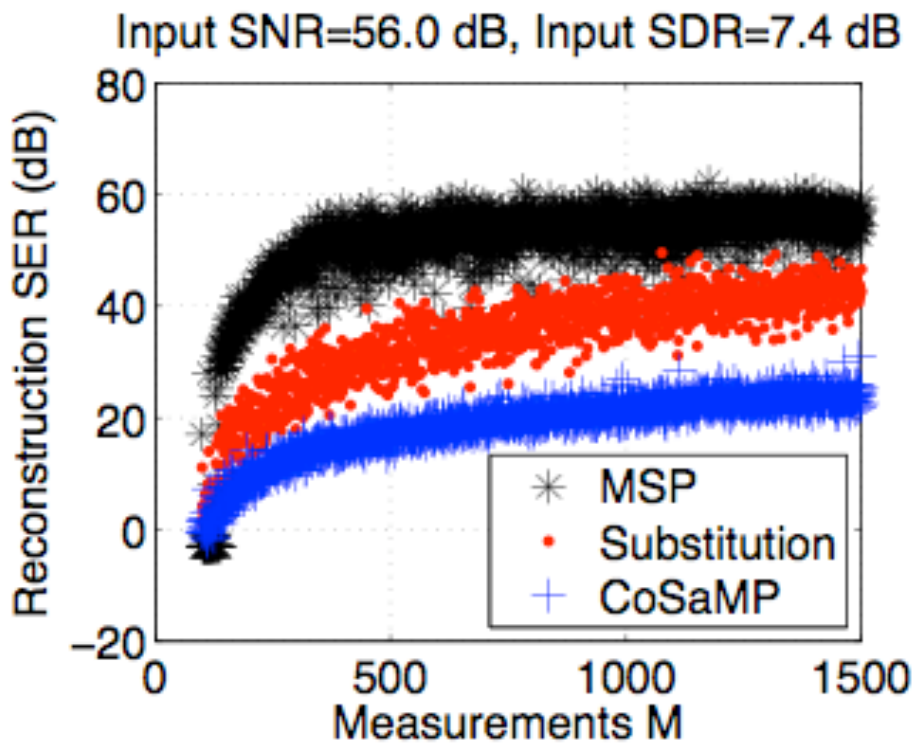

Wrap Up

- Lessons from Compressive Sensing
  - Non-linear reconstruction allows us to **incorporate models**
  - **Signal model**: sparsity; we can extend that further (another talk)
  - **Measurement model**: Quantization, Saturation, Nonlinearities
  - **Randomness helps!** Makes everything well behaved.
  - Conventional wisdom might not hold anymore...

- To dig more:
  - Rice Compressive Sensing repository: dsp.rice.edu/cs
  - CS blog: nuit-blanche.blogspot.com
  - E-mail me: *petrosb@merl.com*