Decoding by Linear Programming
and other Perhaps (not so) Surprising Phenomena

Emmanuel Candès, California Institute of Technology

Workshop on Sparse Representations in Redundant Systems
CSCAMM, University of Maryland, May 2005

Collaborators: Justin Romberg (Caltech), Terence Tao (UCLA)

The Error Correction Problem

- We wish to transmit a "plaintext" $\boldsymbol{f} \in \mathbf{R}^n$ reliably
- Frequently discussed approach: encoding, e.g. generate a "ciphertext" $\boldsymbol{Af}$, where $\boldsymbol{A} \in \mathbf{R}^{m \times n}$ is a coding matrix
- Assume a fraction of the entries of $\boldsymbol{Af}$ are corrupted $\longrightarrow \boldsymbol{y}$

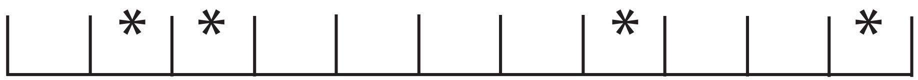

y

  - Corruption is arbitrary
  - We do not know which entries are corrupted
  - We do not know how the corrupted entries are affected
- Is it possible to recover the plaintext exactly from the corrupted ciphertext?

What is Possible?

- If the fraction of corrupted entries is too large, there iws no hope or reconstructing the plaintext.
- Example: $n \ll m$; consider two distinct vectors $f_1, f_2 \in \mathbf{R}^n$ and

$$y = \begin{pmatrix} A_1 f_1 \\ A_2 f_2 \end{pmatrix} \qquad A = \begin{pmatrix} A_1 \\ A_2 \end{pmatrix}.$$

  - $y = A_1 f_1$ with at most half of its entries corrupted
  - $y = A_2 f_2$ with at most half of its entries corrupted

  Cannot distinguish between f_1 and f_2.
- Common assumption: fraction of corrupted entries is not too large
- Ultimate limit of performance: fraction is less than $\frac{1-r}{2}$, $r = \frac{n}{m}$.

Fundamental Questions

- For which fractions ρ is accurate decoding possible?
- Interested in practical algorithms

Decoding by Linear Programming

- ℓ_1-norm

$$\|x\|_{\ell_1} := \sum_{i=1}^{n} |x_i|$$

- To recover f from corrupted data $y = Af + e$, simply solve the ℓ_1-minimization problem

$$(P_1) \qquad \min_{g \in \mathbf{R}^n} \|y - Ag\|_{\ell_1}.$$

- Equivalent linear program:

$$\min \sum_{i=1}^{m} t_i, \quad \text{subject to} \quad -t \leq y - Ag \leq t$$

optimization variables: $t \in R^m, g \in \mathbf{R}^n$

Surprise!

$$(P_1) \qquad \min_{g \in \mathbf{R}^n} \|y - Ag\|_{\ell_1}.$$

Under suitable conditions on the coding matrix A, the input f is the unique solution to (P_1), provided that the fraction of corrupted entries is not too large, i.e. does not exceed some strictly positive constant $\rho^*(A)$

- Minimizing ℓ_1 recovers all signals regardless of the corruption pattern
- Size of corruption does not matter
- There is nothing a clever opponent can do to corrupt the ciphertext, and fool the LP decoder.

Peek at the Results

- Random Gaussian coding matrix A: A_{ij} i.i.d. $N(0, 1)$
- With overwhelming probability, if the fraction of the corrupted entries does not exceed ρ^*, the solution to (P_1) is unique and equal to f.
- Universal: probability that A allows exact decoding of all plaintexts is at least $1 - O(e^{-\alpha m})$
- See also very recent work of Vershynin and Rudelson (2005).

The Importance of ℓ_1

- Minimize instead the ℓ_2-distance

$$\min_{g \in \mathbf{R}^n} \|y - Ag\|_{\ell_2}$$

- Solution given by least-squares

$$g^\star = (A^T A)^{-1} A^T y = f + (A^T A)^{-1} A^T e$$

- Error term:
  - No reason to vanish
  - Goes to infinity as $\|e\|_{\ell_2}$ goes to infinity.

Practical Performance, I

- A_{ij} i.i.d. $N(0,1)$
- $f \in \mathbf{R}^n$
- Corruption: flip the sign of randomly selected entries

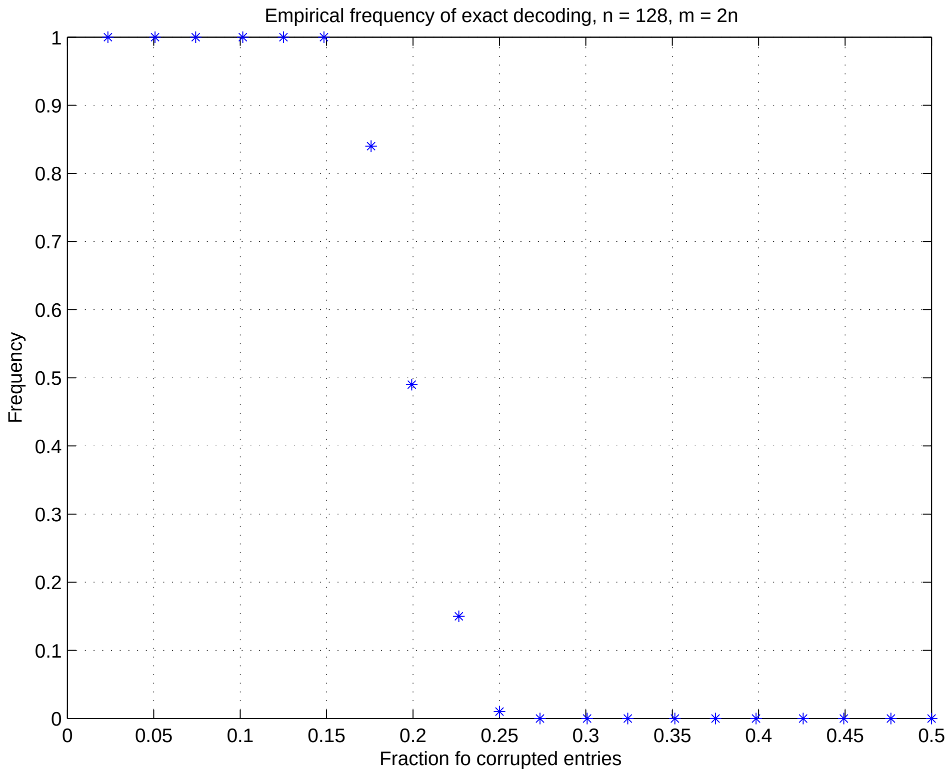

$n = 128$, $m = 2n$

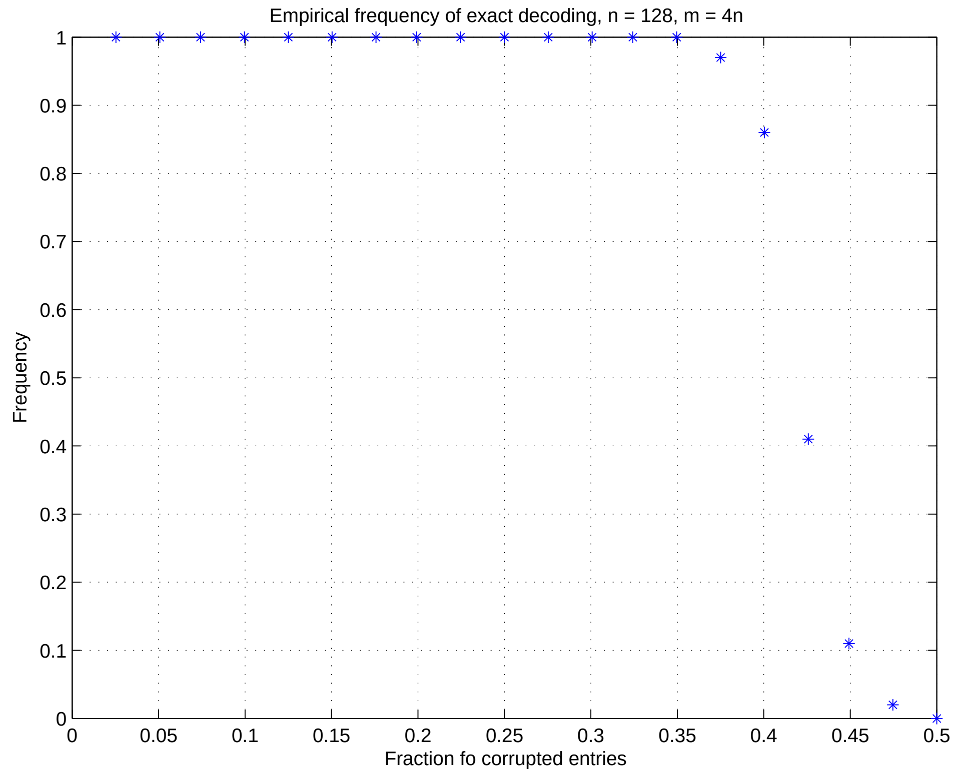

$n = 128$, $m = 4n$

Practical Performance, II

- A_{ij} i.i.d. $\mathbf{P}(A_{ij} = \pm 1) = 1/2$
- $f \in 0, 1^n$
- Corruption: flip the sign of randomly selected entries
- Solve $\min_{g \in \mathbf{R}^n} \|y - Ag\|_{\ell_1}$ subject to $0 \leq g \leq 1$, and round up.

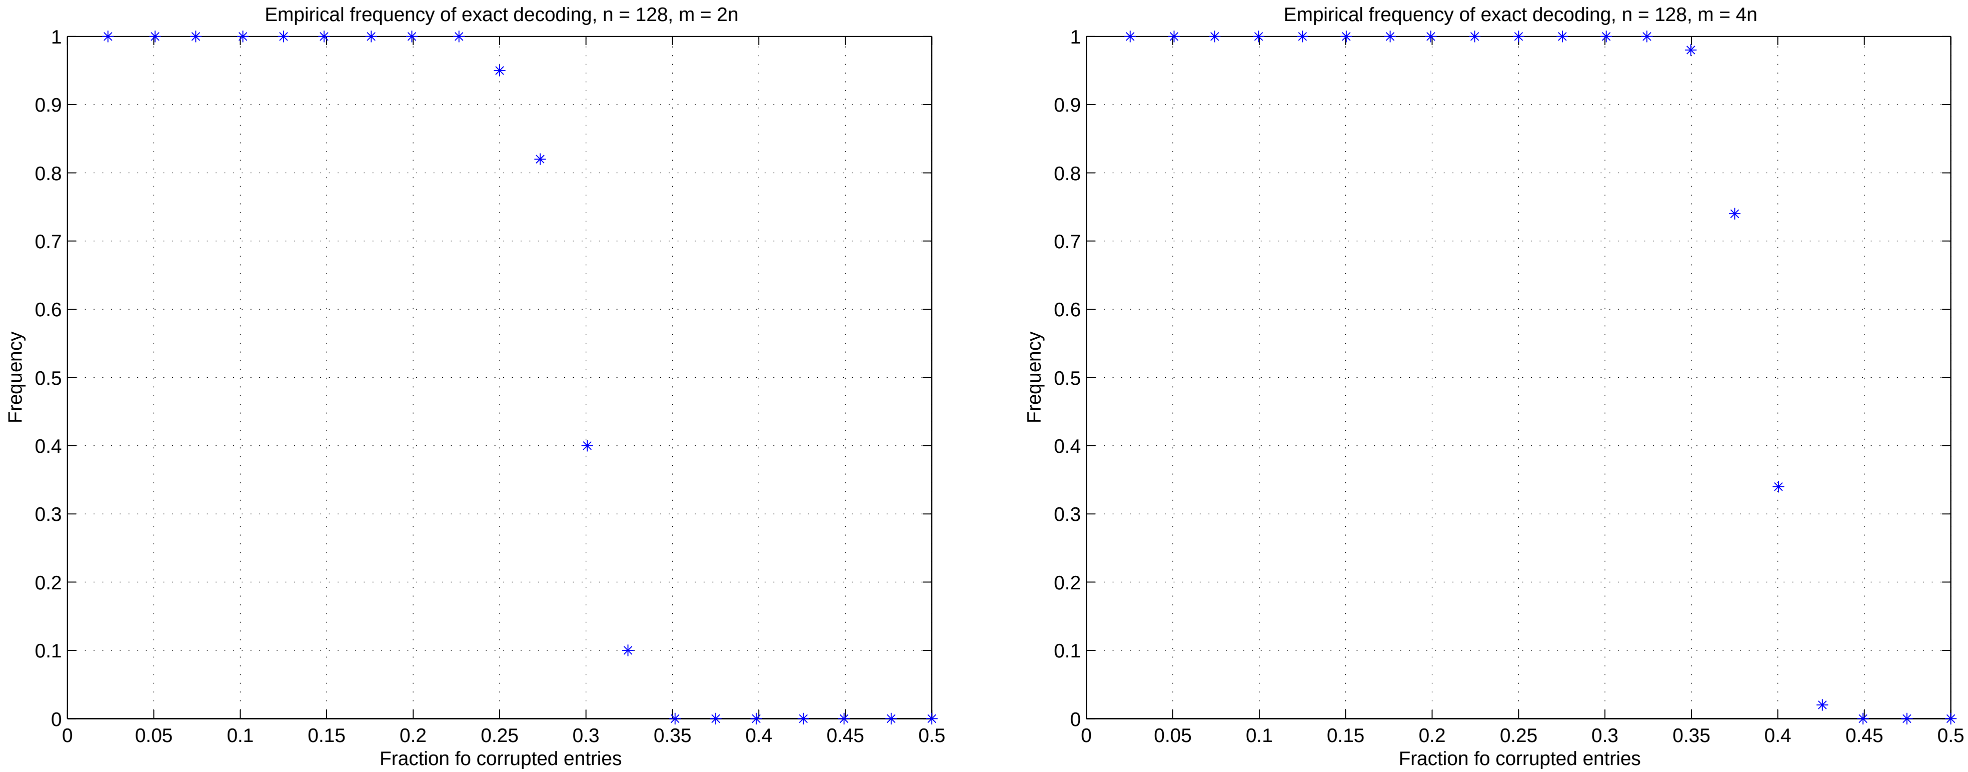

$n = 128$, $m = 2n$ $\qquad$ $n = 128$, $m = 4n$

Understanding this Phenomenon

- Corrupted ciphertext: $y = Af + e$
- $(m-n) \times n$ matrix B which annihilates A on the left, i.e. such that $BA = 0$

$$\tilde{y} = By = B(Af + e) = Be$$

- Equivalent problem: recover e from $\tilde{y}$
- Need to solve an *underdetermined system of linear equations*
- Useful equivalence: set $g = f + h$

$$\begin{aligned}
(P_1) \quad \min_{g \in \mathbf{R}^n} \|y - Ag\|_{\ell_1}, \quad &\Leftrightarrow \quad \min_{h \in \mathbf{R}^n} \|e - Ah\|_{\ell_1}, \\
&\Leftrightarrow \quad \min \|d\|_{\ell_1}, \qquad d = e - Ah
\end{aligned}$$

Observe that $d = e - Ah \Leftrightarrow Bd = Be$, i.e.

$$(P_1) \quad \Leftrightarrow \quad \min \|d\|_{\ell_1}, \qquad Bd = Be$$

Sparse Solutions to Underdetermined Systems

- Equivalent problem

$$(P_1) \quad \min \|d\|_{\ell_1}, \qquad Bd = Be$$

- Also known as *Basis Pursuit* (Chen, Donoho, Saunders, 1996)
- Ability to decode accurately $\Leftrightarrow$ ability to find sparse solutions to underdetermined systems

Agenda: Finding sparse solutions to underdetermined systems

- Error correction
- Signal recovery from incomplete measurements
- Uniform uncertainty principles
- Stability
- Implications for information/coding theory
- Numerical evidence

Model 1D Problem

$f \in \mathbf{R}^N$ is a superposition of $|T|$ spikes ($|T|$ nonzero components)

$$f = \sum_{t \in T} f(t) \delta_t$$

(δ_t is a spike at t) and with Discrete Fourier Transform (DFT)

$$\hat{f}(\omega) = \sum_{t=0}^{N-1} f(t) e^{-i 2\pi \omega t / N}$$

Observe $\hat{f}(\omega)$ on Ω, $K := |\Omega| \ll N$

Sparse Spike Train

Sparse sequence of $|T|$ spikes

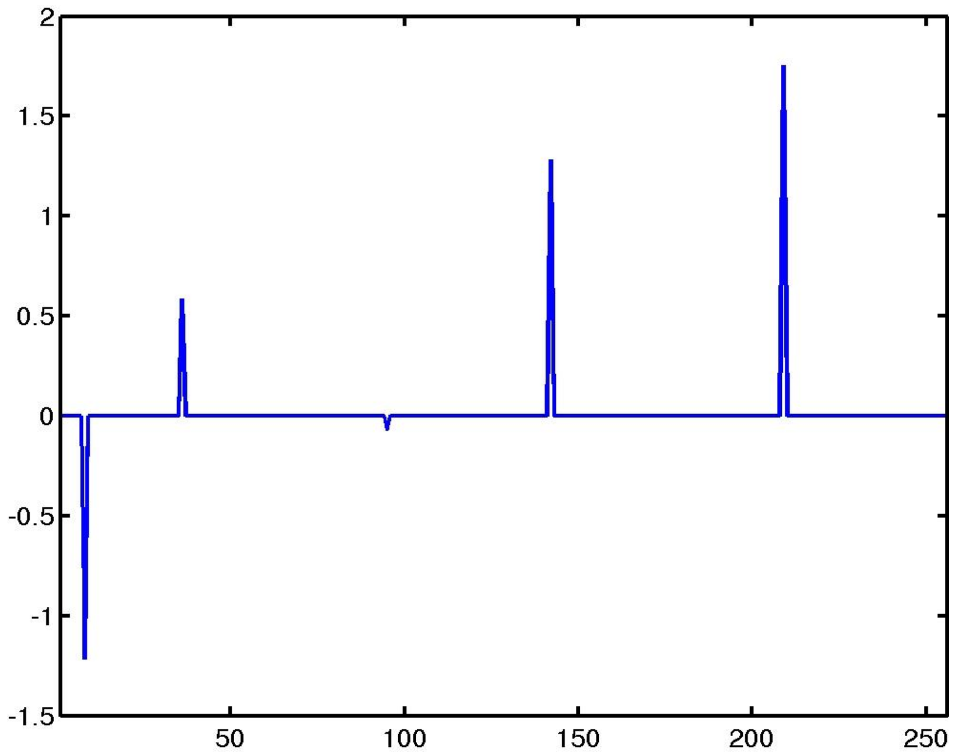

Observe $|\Omega|$ Fourier coefficients

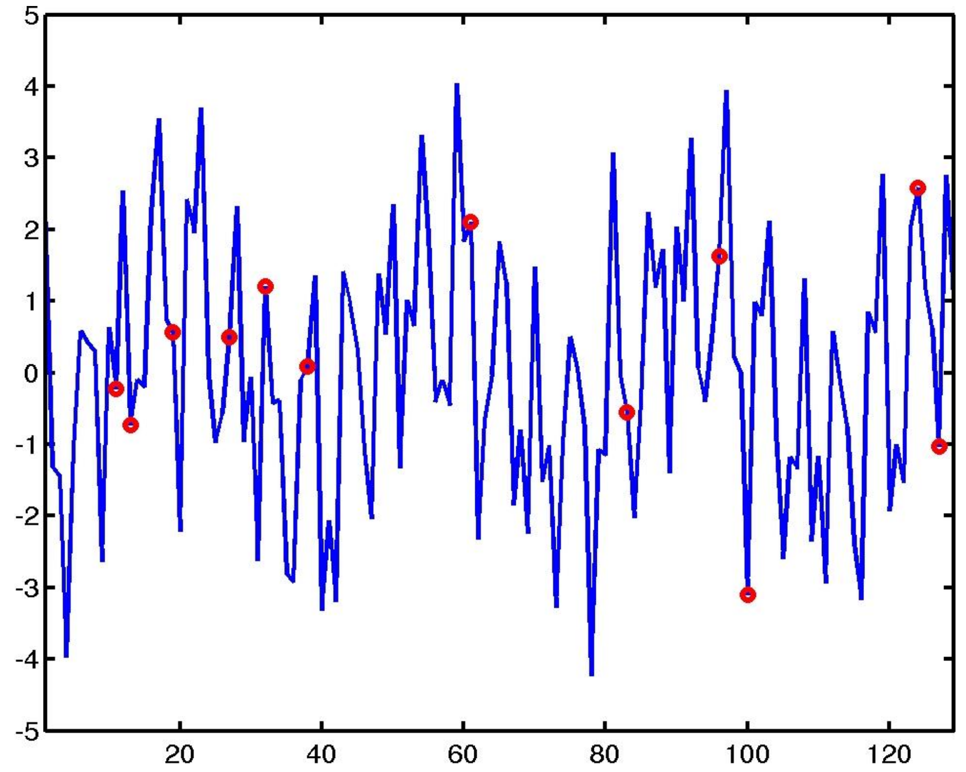

ℓ_1 Reconstruction

Reconstruct by solving

$$\min_{g} \|g\|_{\ell_1} := \sum_{t} |g(t)| \quad \text{s.t.} \quad \hat{g}(\omega) = \hat{f}(\omega), \;\; \omega \in \Omega$$

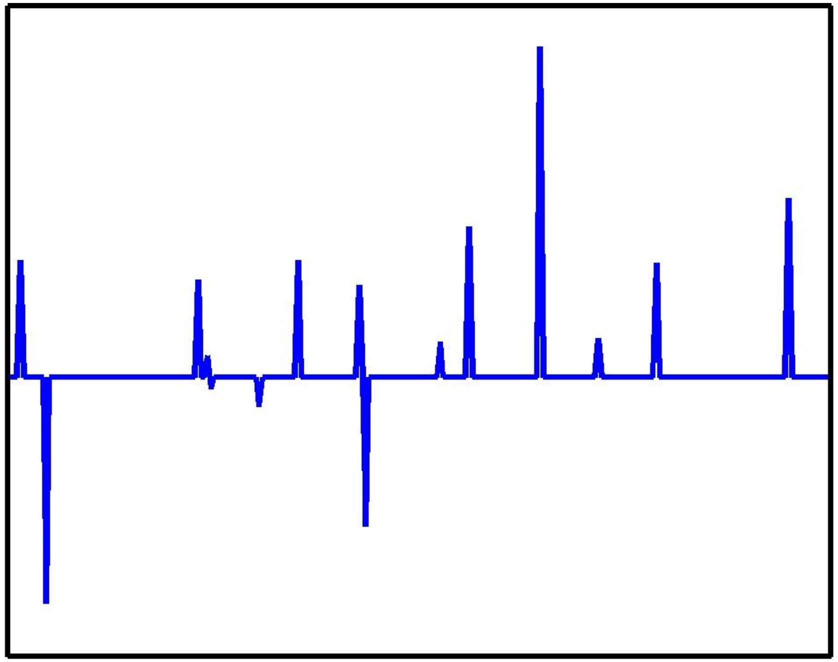

original

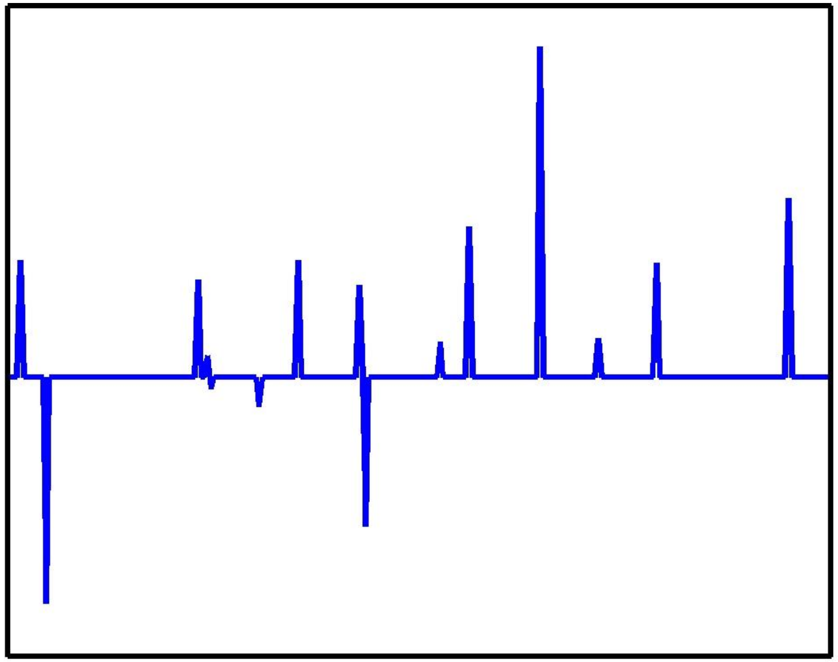

recovered from 30 Fourier samples

A First Recovery Theorem

$$(P_1) \qquad \min_{g \in \mathbf{R}^N} \|g\|_{\ell_1}, \qquad \hat{g}(\omega) = \hat{f}(\omega),\ \omega \in \Omega$$

Theorem 1 (C., Romberg, Tao) *Suppose*

- f *supported on set* T
- *Observations selected at random with*

$$|\Omega| \geq C \cdot |T| \log N.$$

Minimizing ℓ_1 *reconstructs exactly with overwhelming probability.*

- Unimprovable
- In theory, $C \approx 20$
- In practice, $C \log N \approx 2$
- (Very) hard stuff

Reconstructed perfectly from 30 Fourier samples

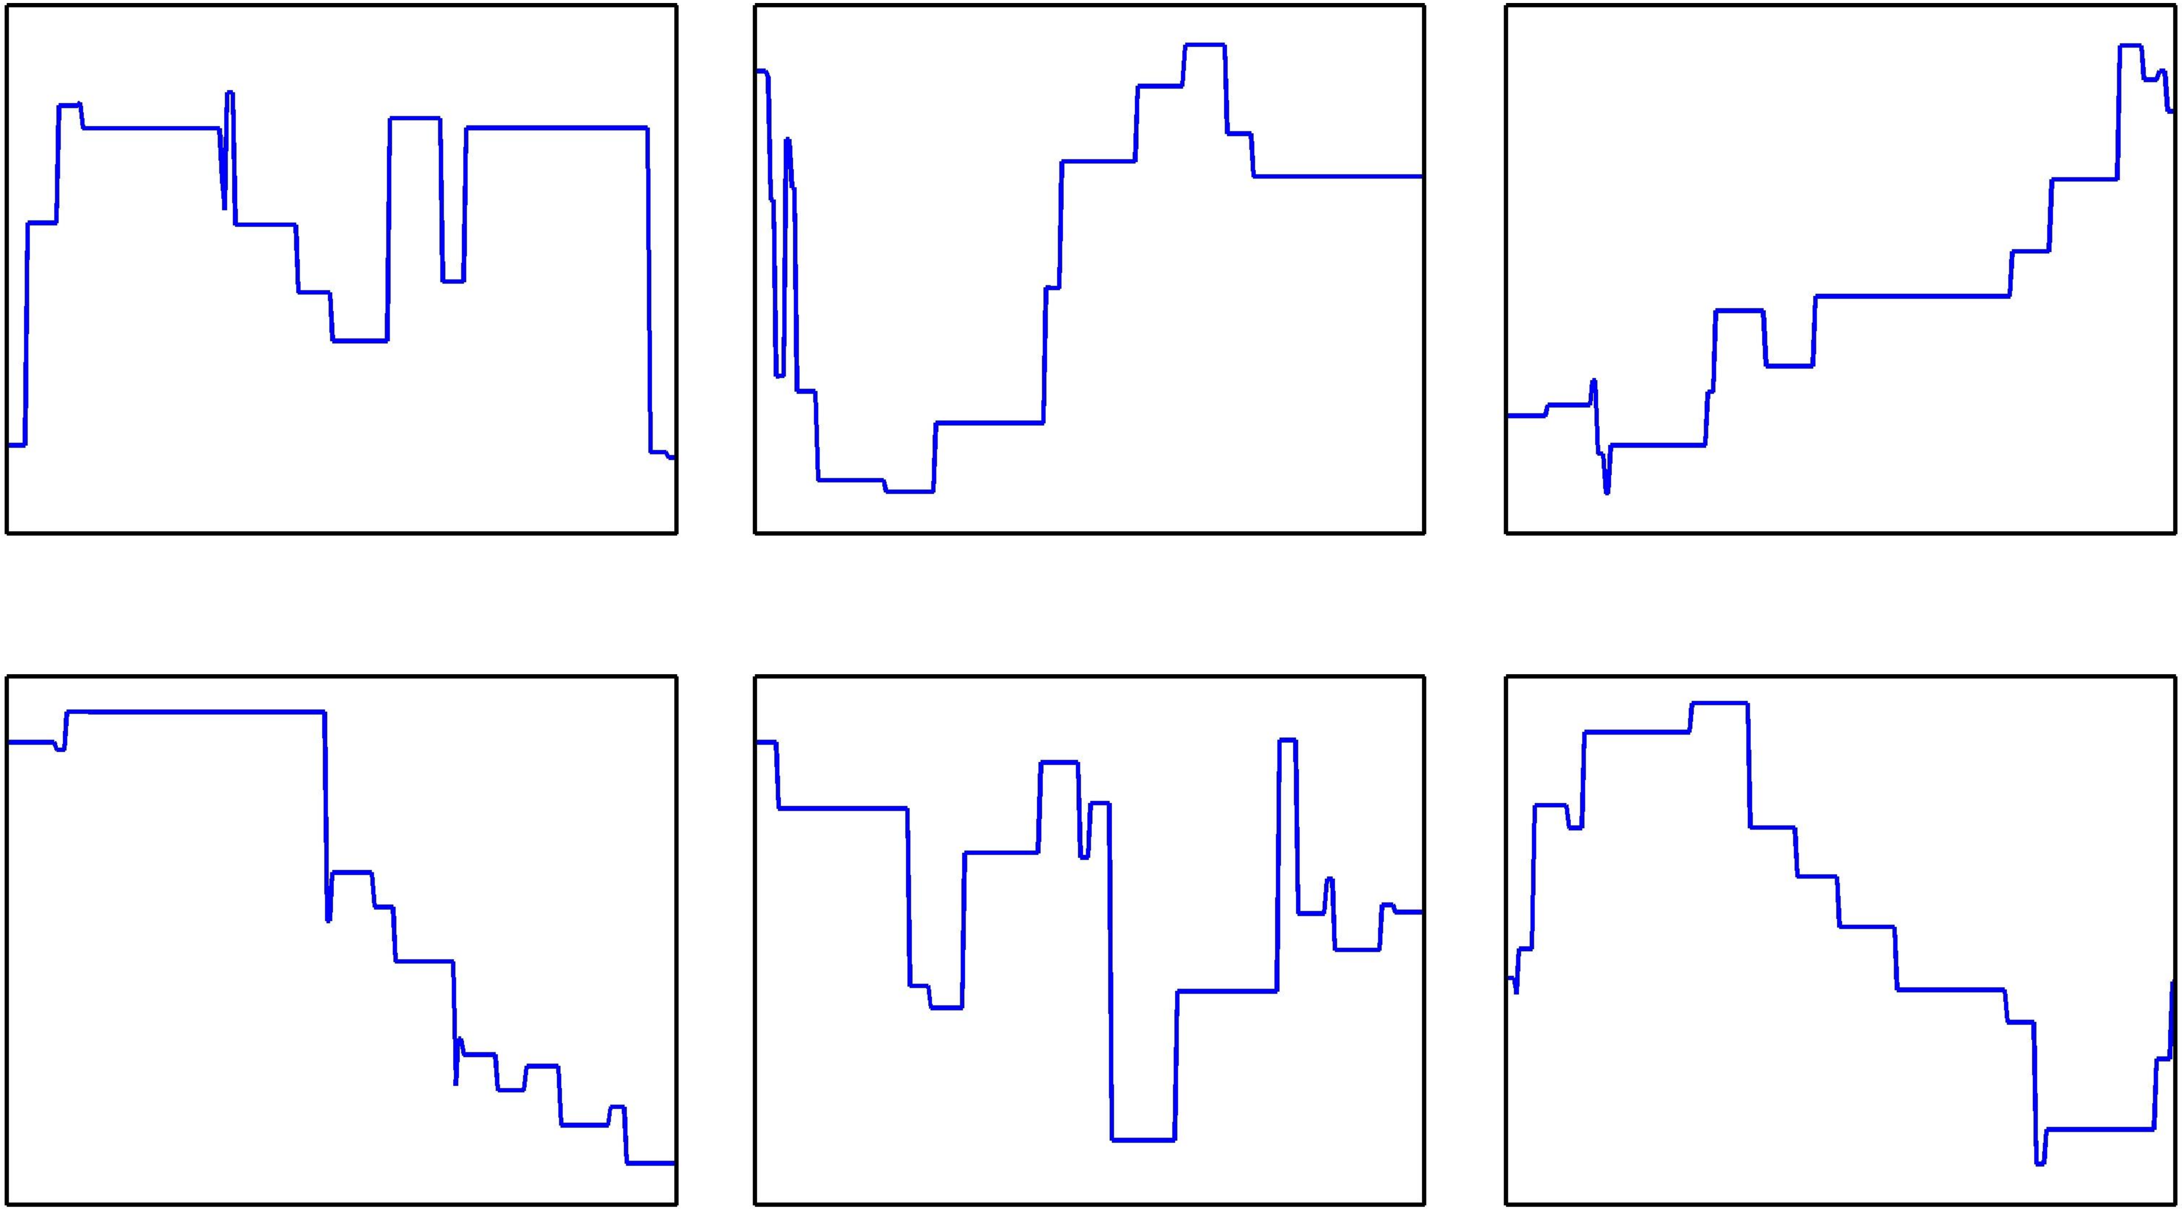

Nonlinear Sampling Theorem

- Switch roles of time and frequency:
  - $\hat{f}$ supported on set Ω in freq domain
  - sample on set T in time domain
- Shannon sampling theorem:
  - Ω is a connected set of size B
  - we can reconstruct from B equally spaced time-domain samples
  - linear reconstruction by sinc interpolation
- Nonlinear sampling theorem:
  - Ω is an *arbitrary* set of size B
  - we can reconstruct from $\sim B \log N$ *randomly* placed samples
  - nonlinear reconstruction by convex programming

A Second Recovery Theorem

- Gaussian random matrix

$$F(k,t) = X_{k,t}, \qquad X_{k,t} \ i.i.d. \ N(0,1)$$

- This will be called the *Gaussian ensemble*

$$(P_1) \qquad \min_{g \in \mathbf{R}^N} \|g\|_{\ell_1} \qquad Fg = Ff.$$

Theorem 2 (C., Tao) *Suppose*

- f *supported on set* T
- K *observations (random projection) with*

$$K \geq C \cdot |T| \log N.$$

Minimizing ℓ_1 *reconstructs exactly with overwhelming probability.*

Gaussian Random Measurements and Random Projections

$$y = Ff, \qquad y_k = \langle f, X \rangle, \quad X_t \ i.i.d. \ N(0,1),$$

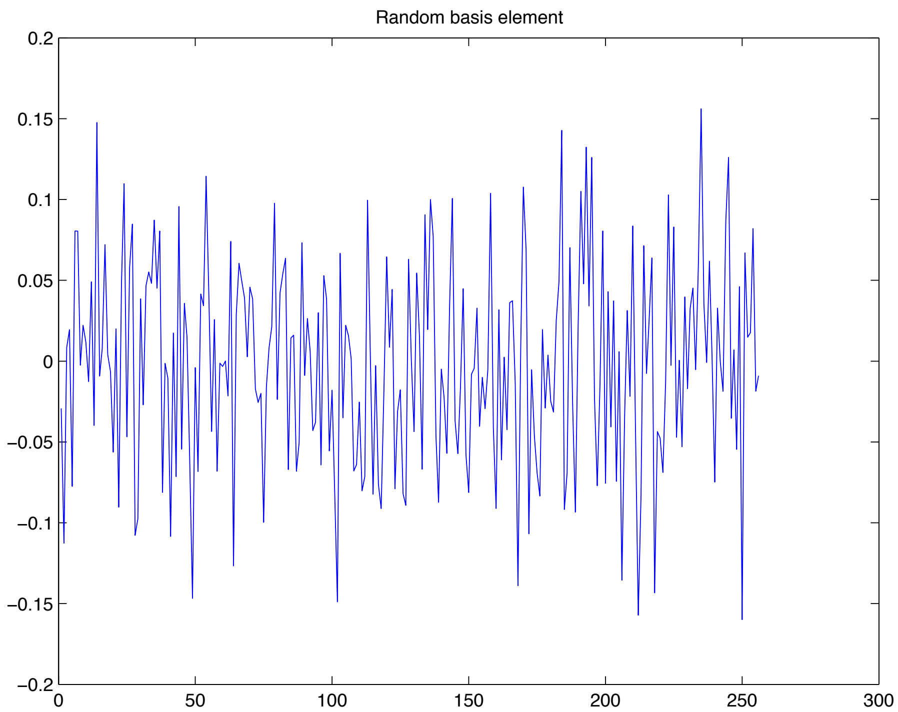

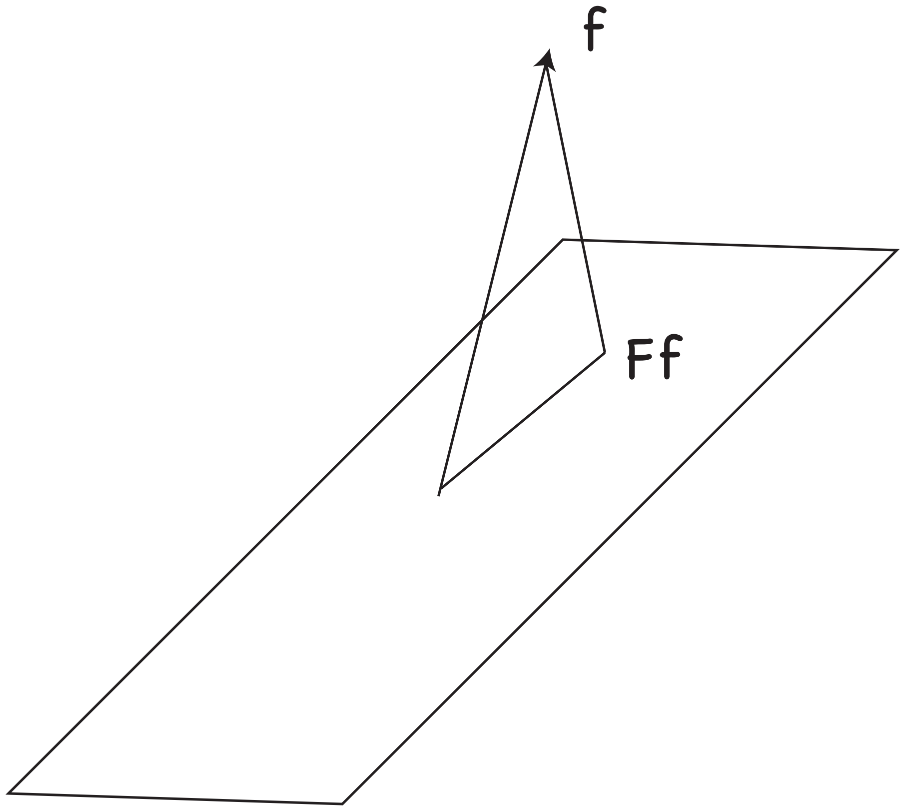

Previous Work

- ℓ_1 reconstruction in widespread use
  - Santosa and Symes (1986), and others in Geophysics (Claerbout)
  - Donoho and Stark
- Sparse decompostions via Basis Pursuit
  - Chen, Donoho, Saunders (1996)
  - Donoho, Huo, Elad, Gribonval, Nielsen, Fuchs, Tropp (2001-2005)
- Novel sampling theorems
  - Bresler and Feng (2002)
  - Vetterli and others (2002-2004)
- Fast algorithms
  - Gilbert, Strauss, et al. (2002-2005)

Numerical Results

- Signal length $N = 1024$
- Randomly place $|T|$ spikes, observe K random frequencies
- Measure $\%$ recovered perfectly
- white = always recovered, black = never recovered

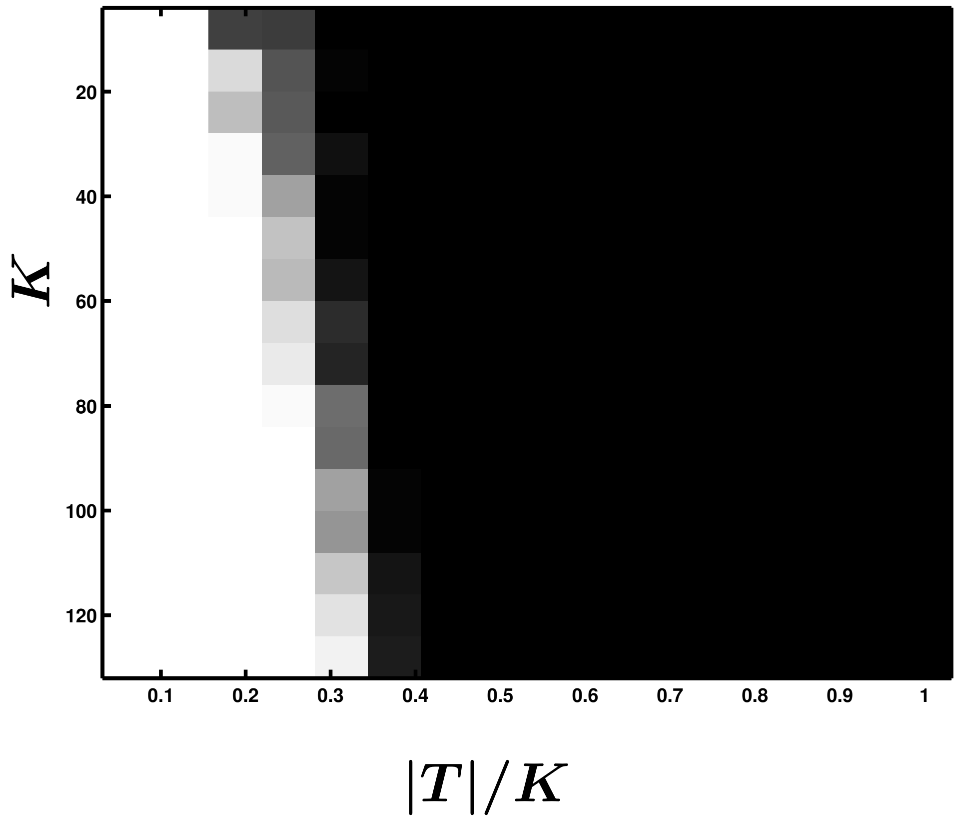

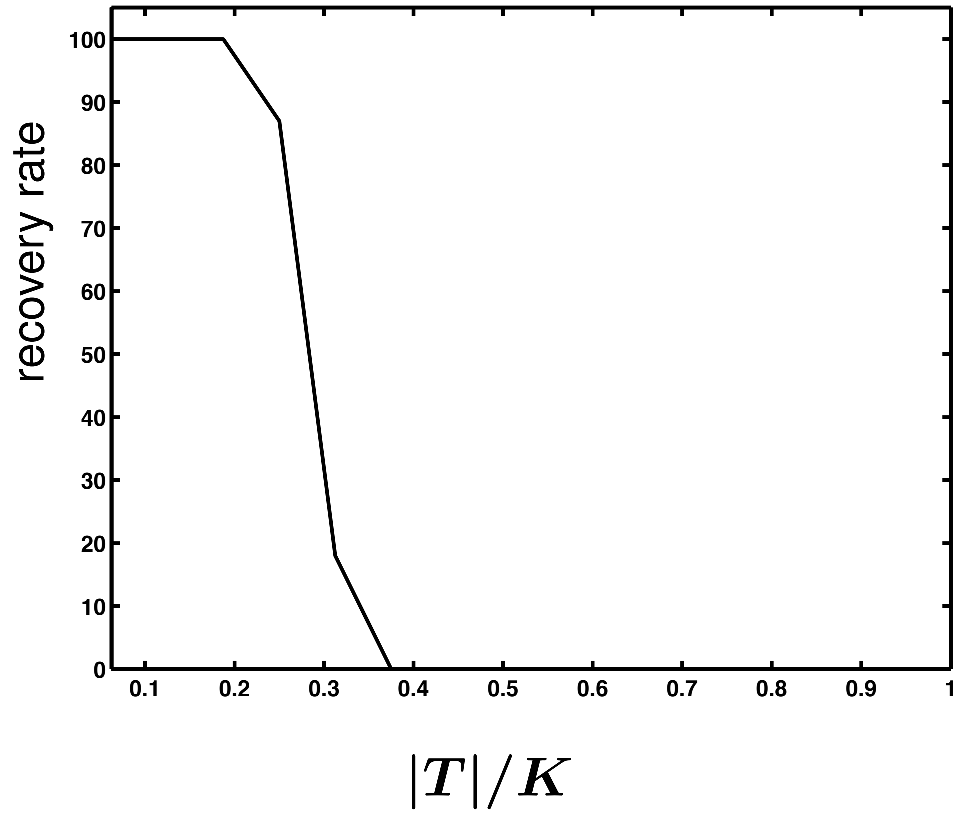

Key to Recovery: Uncertainty Principles

Weyl-Heisenberg Uncertainty Principle

W. Heisenberg, 1901-1976

Weyl-Heisenberg

- f 'lives' on an interval of length Δt
- $\hat{f}$ 'lives' on an interval of length $\Delta \omega$

$$\Delta t \cdot \Delta \omega \geq 1$$

Restricted Isometries

- Measurement matrix $\boldsymbol{F}$, $\boldsymbol{F} \in \mathbf{R}^{\boldsymbol{K}\times\boldsymbol{N}}$; $\boldsymbol{F_T}$ columns of $\boldsymbol{F}$ corresponding to $\boldsymbol{T}$, $\boldsymbol{F_T} \in \mathbf{R}^{\boldsymbol{K}\times|\boldsymbol{T}|}$.

- Restricted isometry constants $\boldsymbol{\delta_S}$

$$(1 - \delta_S)\,\mathrm{Id} \le F_T^* F_T \le (1 + \delta_S)\,\mathrm{Id}, \qquad \forall T,\ |T| \le S.$$

- $\boldsymbol{F}$ obeys a *uniform uncertainty principle* for sets of size $\boldsymbol{F}$ if $\boldsymbol{\delta_S} \le \mathbf{1/2}$, say.

- Uniform because must hold for *all* $\boldsymbol{T}$'s.

Why Do We Call This an Uncertainty Principle?

- $\boldsymbol{F_\Omega}$, rows of the DFT isometry (corresponding to $\boldsymbol{\Omega}$)
- $\boldsymbol{F_{\Omega T}}$, columns of $\boldsymbol{F_\Omega}$ (corresponding to $\boldsymbol{T}$)
- UUP

$$(1-\delta_S)\,\frac{|\Omega|}{N}\cdot\|f_T\|^2 \le \|F_{\Omega T} f_T\|^2 \le (1+\delta_S)\,\frac{|\Omega|}{N}\cdot\|f_T\|^2$$

- Implications
  - $\boldsymbol{f}$ supported on $\boldsymbol{T}$, $|\boldsymbol{T}| \le \boldsymbol{S}$
  - If UUP holds, then

$$(1-\delta_S)\,\frac{|\Omega|}{N} \le \|\hat{f}\cdot 1_\Omega\|^2/\|\hat{f}\|^2 \le (1+\delta_S)\,\frac{|\Omega|}{N}$$

Sparse/Compressible Signals

- Sparse signal: f is sparse if f is supported on a "small" set T
- In real life, signals of interest may not be sparse but compressible
- Compressible signal: f is compressible if it is well-approximated by a sparse signal.
- Frequently discussed model of compressible signals: rearrange the entries in decreasing order $|f_{(1)|} \geq |f_{(2)}| \geq \ldots \geq |f_{(N)}|$

$$|f|_{(k)} \leq C \cdot k^{-s}, \forall k$$

- Implications: f_T truncated vector corresponding to the $|T|$ largest entries of $f \in \mathbf{R}^N$

$$\|f - f_T\|_{\ell_2} \leq C \cdot |T|^{-(s-1/2)}$$

- This is what makes transform coders work (sparse coding)

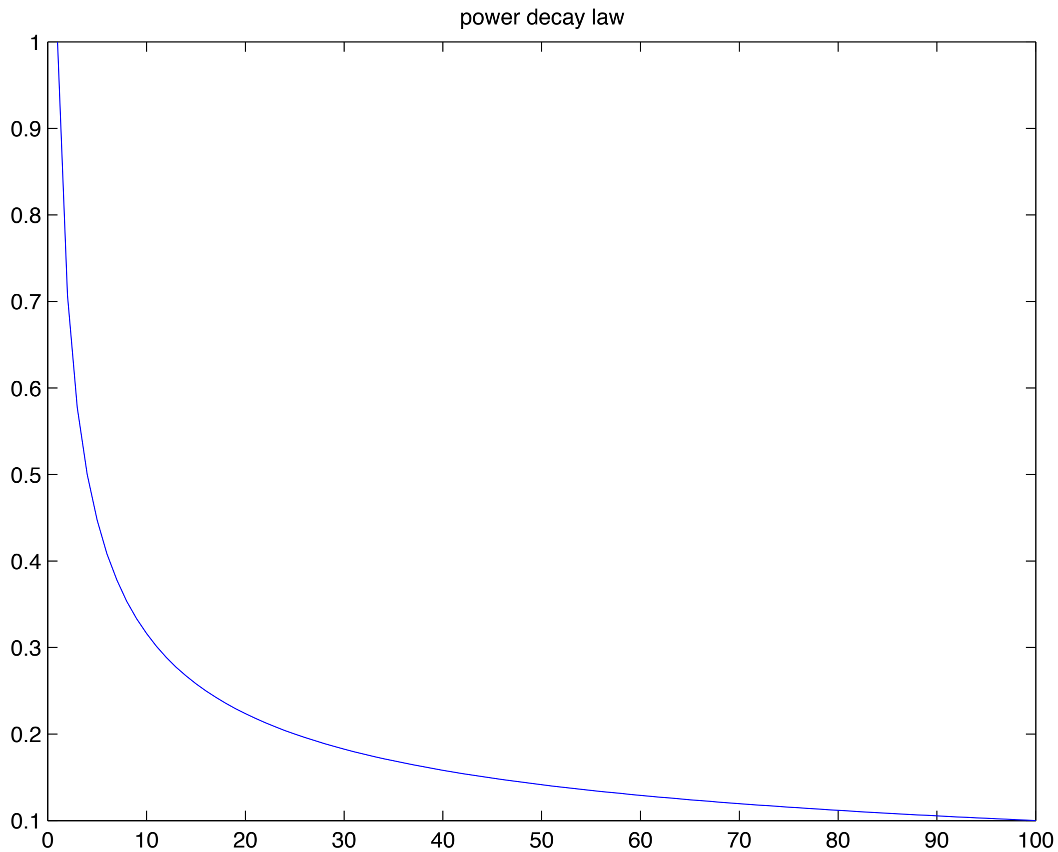
power decay law
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
10
20
30
40
50
60
70
80
90
100

Compressible Signals I: Wavelets in 1D

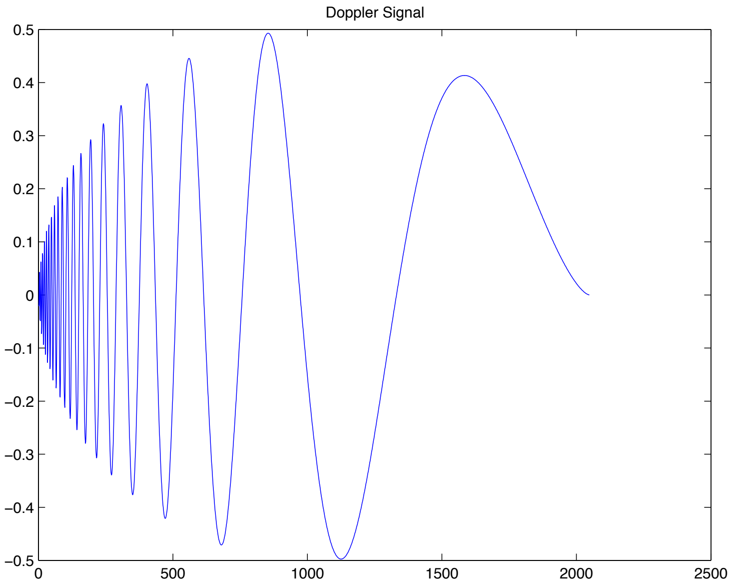

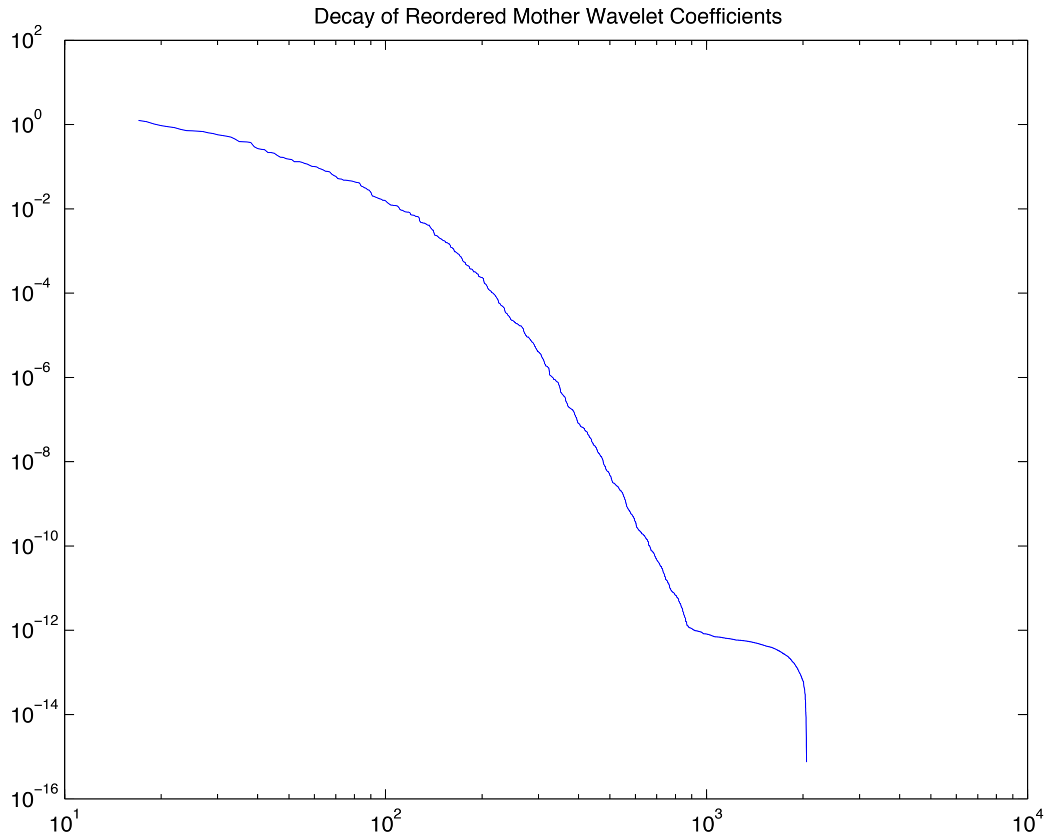

Compressible Signals II: Wavelets in 2D

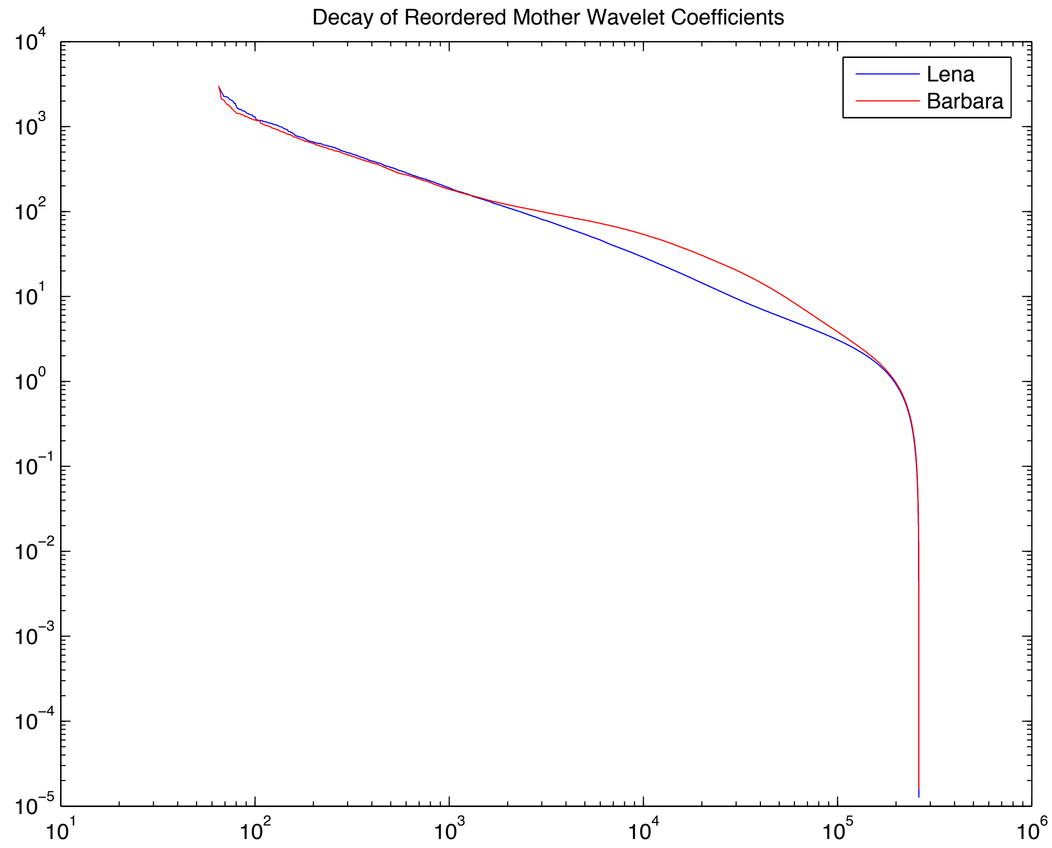

UUP and Signal Recovery from Undersampled Data

$$(P_1) \qquad \min_{g \in \mathbf{R}^N} \|g\|_{\ell_1}, \qquad Ag = Af.$$

Theorem 3 (C., Tao, 2004) *Assume $\delta_{3S} + 3\delta_{4S} < 2$ (UUP holds).*

- *If f supported on any set T, $|T| \le S$, then the recovery is exact.*
- *For all $f \in \mathbf{R}^N$*

$$\|f - f^\sharp\|_{\ell_2} \le 8 \, \frac{\|f - f_S\|_{\ell_1}}{\sqrt{S}}$$

This is a purely deterministic statement. Nothing is random here!

- If f is sufficiently sparse, the recovery is exact
- If f is compressible

$$\|f - f^\sharp\|_{\ell_2} \le 8 \cdot S^{-(s-1/2)}$$

- Useful if S is large

Examples

- Gaussian ensemble A_{ij} i.i.d. $N(0, 1/K)$ obeys UUP with

$$S \lesssim K/\log[N/K]$$

- Binary ensemble A_{ij} i.i.d. $P(A_{ij} = \pm 1/\sqrt{K}) = 1/2$ obeys UUP with

$$S \lesssim K/\log[N/K]$$

- Fourier ensemble (K random rows) obeys UUP with

$$S \lesssim K/(\log N)^6$$

  Probably true with $\log^4 N$ (C., Tao and Vershynin and Rudelson)

All with overwhelming probability.

UUP for General Orthonormal Systems

- f is sparse in an orthogonal basis Ψ

$$f(t) = \sum_{m \in \mathcal{I}} \alpha_m \psi_m(t)$$

- Measurements in different orthogonal basis Φ

$$y_k = \langle f, \phi_k \rangle \quad k \in \Omega \quad y = \Phi_\Omega f.$$

- Recover via

$$\min \|\alpha\|_{\ell_1} \qquad \Phi_\Omega \Psi^* \alpha = y$$

- General orthogonal ensemble $\Phi\Psi^*$ (random rows) obeys UUP with

$$S \lesssim K/[\mu^2 (\log N)^6]$$

- Incoherence $\mu = \sqrt{N}\ \max |\langle \phi_\omega, \psi_m \rangle|$.

Reconstruction of Piecewise Polynomials, I

- Randomly select a few jump discontinuities
- Randomly select cubic polynomial in between jumps
- Observe about 500 random coefficients
- Minimize ℓ_1 norm of wavelet coefficients

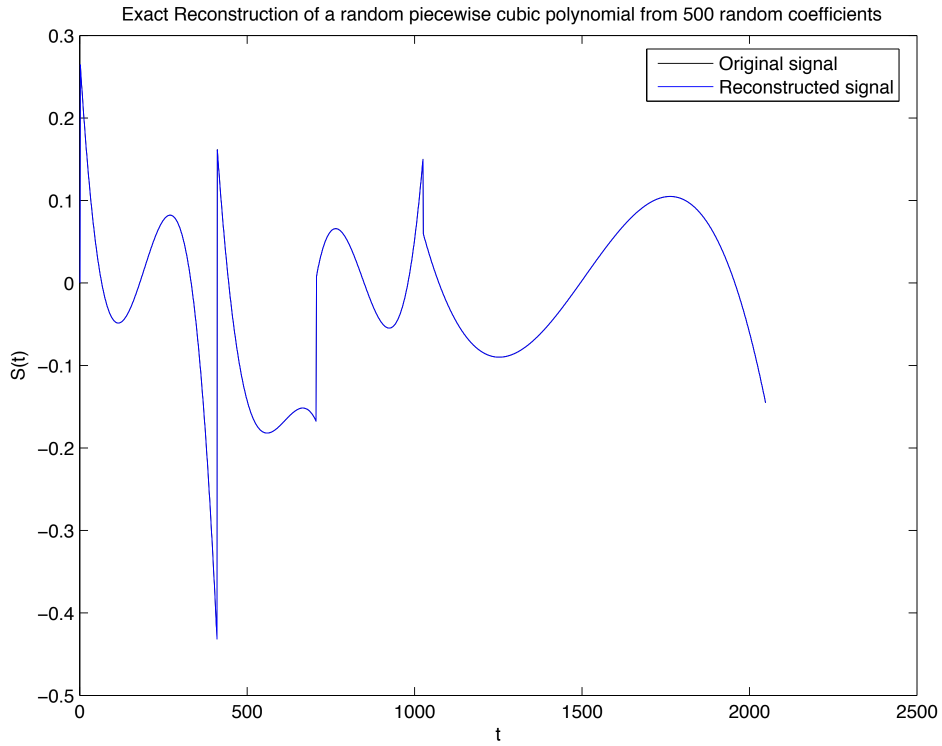

Reconstructed signal

Reconstruction of Piecewise Polynomials,II

- Randomly select 8 jump discontinuities
- Randomly select cubic polynomial in between jumps
- Observe about 200 Fourier coefficients at random

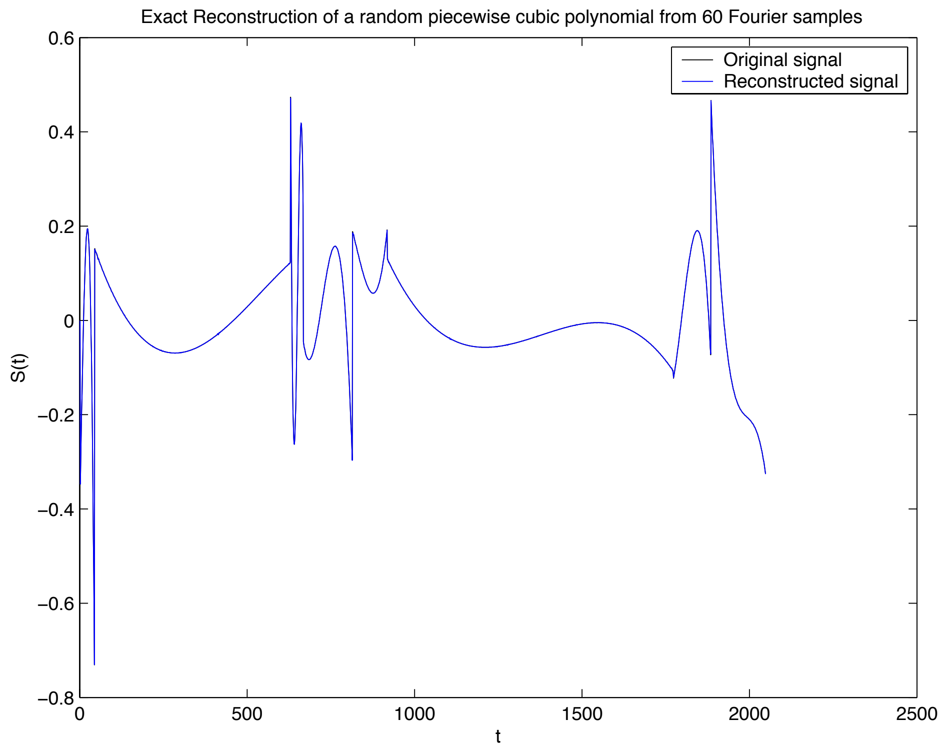

Reconstructed signal

Reconstruction of Piecewise Polynomials, III

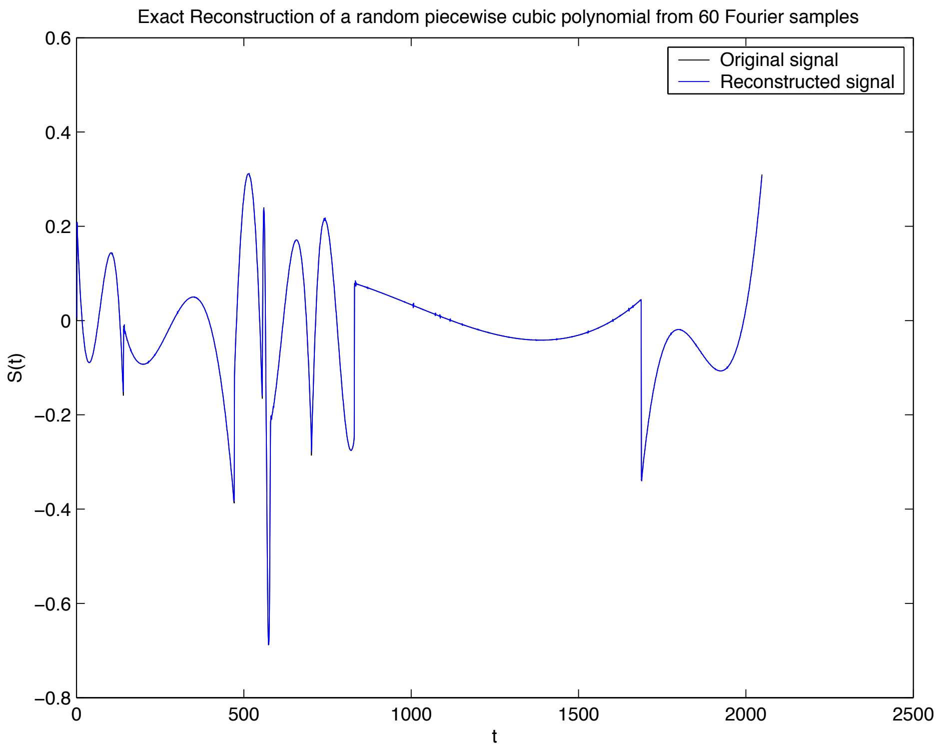

Reconstructed signal

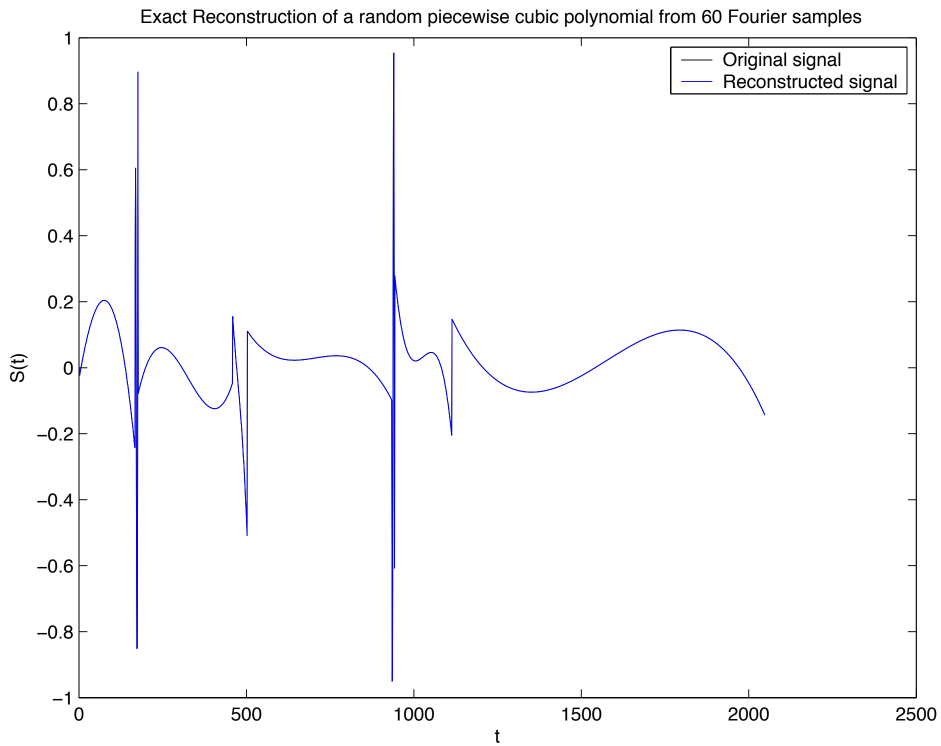

Reconstructed signal

About 200 Fourier coefficients only!

Recovery of Sparse/Compressible Signals

- How many measurements to recover f to within precision $\epsilon = K^{-(s-1/2)}$?
- Intuition: at least K, probably many more.

Where Are the Largest Coefficients?

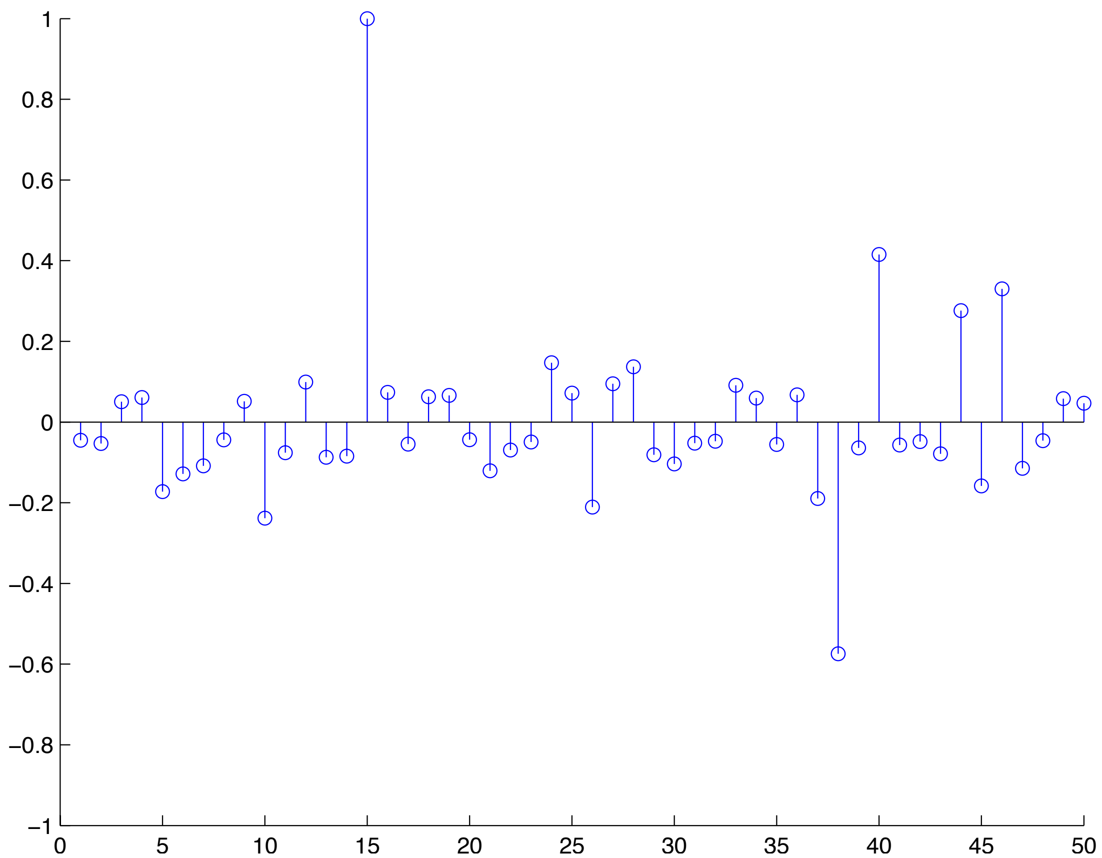

Implications of Approximate Recovery

- Gaussian ensemble: $A_{i,j}$ i.i.d. $N(0, 1/K)$
- Random projection on a K-dimensional plane ($y = Af$)
- f compressible

$$\|f - f^\sharp\|_{\ell_2} \leq C \cdot (K/\log[N/K])^{-(s-1/2)}.$$

- See also recent work by D. Donoho (2004)

Another Surprise!

Want to know an object up to an error ϵ; e.g. an object whose wavelet coefficients are sparse.

- *Strategy 1:* Oracle tells exactly (or you collect all N wavelet coefficients) which K coefficients are large and measure those

$$\|f - f_K\| \asymp \epsilon$$

- *Strategy 2:* Collect $K \log[N/K]$ random coefficients and reconstruct using ℓ_1.

Surprising claim

- Same performance but with only $K \log[N/K]$ coefficients!
- Performance is achieved by solving an LP.

Optimality

- Can you do with fewer than $K \log[N/K]$ for accuracy $K^{-(s-1/2)}$?
- Simple answer: NO
- Connected with theory of Gelfand widths
- Connected with information theory (rate-distortion curve of compressible signals)

Stable Recovery?

- In real applications, data are corrupted
- Better model: $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{f} + \boldsymbol{e}$, where $\boldsymbol{e}$ may be stochastic, deterministic.
- Recall most of the singular values of $\boldsymbol{A}$ are zero
- Hopeless?

UUP and Stable Recovery from Undersampled Data

ℓ_1-based regularization

$$\min \|g\|_{\ell_1} \qquad \|y - Ag\|_{\ell_2} \leq \|e\|_{\ell_2}$$

Theorem 4 (C., Romberg, Tao) *Assume* $\delta_{3S} + 3\delta_{4S} < 2$.

$$\|f - f^\sharp\|_{\ell_2} \leq 8 \cdot \left(\frac{\|f - f_S\|_{\ell_1}}{\sqrt{S}} + \|e\|_{\ell_2} \right)$$

- No blow up!
- Reconstruction within the noise level
- Nicely degrades as noise level increases

Geometric Intuition

- f feasible $\Rightarrow f^\sharp$ inside the diamond
- $f^\sharp$ obeys the constraint $\Rightarrow f^\sharp$ inside the slab (tube)

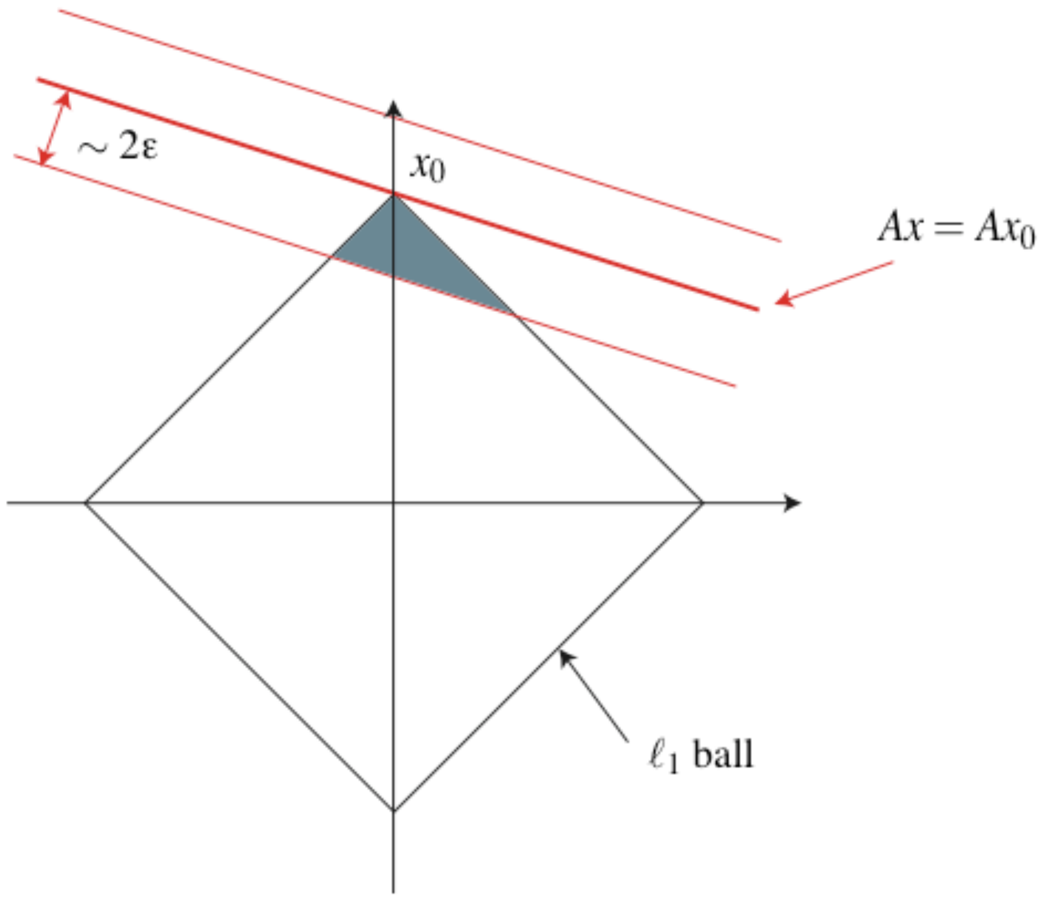

Reconstruction from 100 Random Coefficients

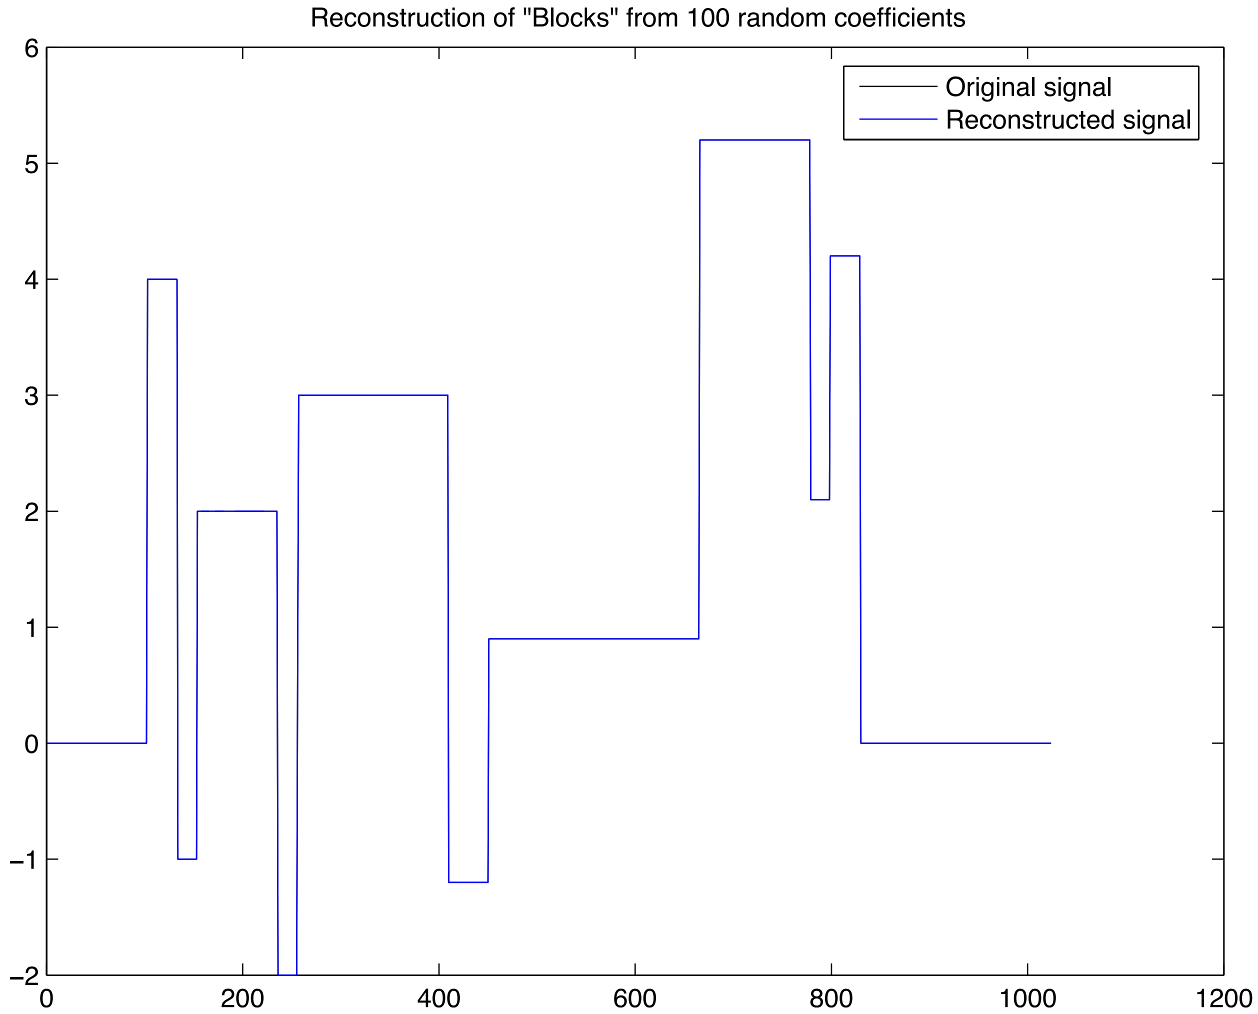

'Compressed' signal

Reconstruction from Random Coefficients

Minimize TV subject to random coefficients + ℓ_1-norm of wavelet coefficients.

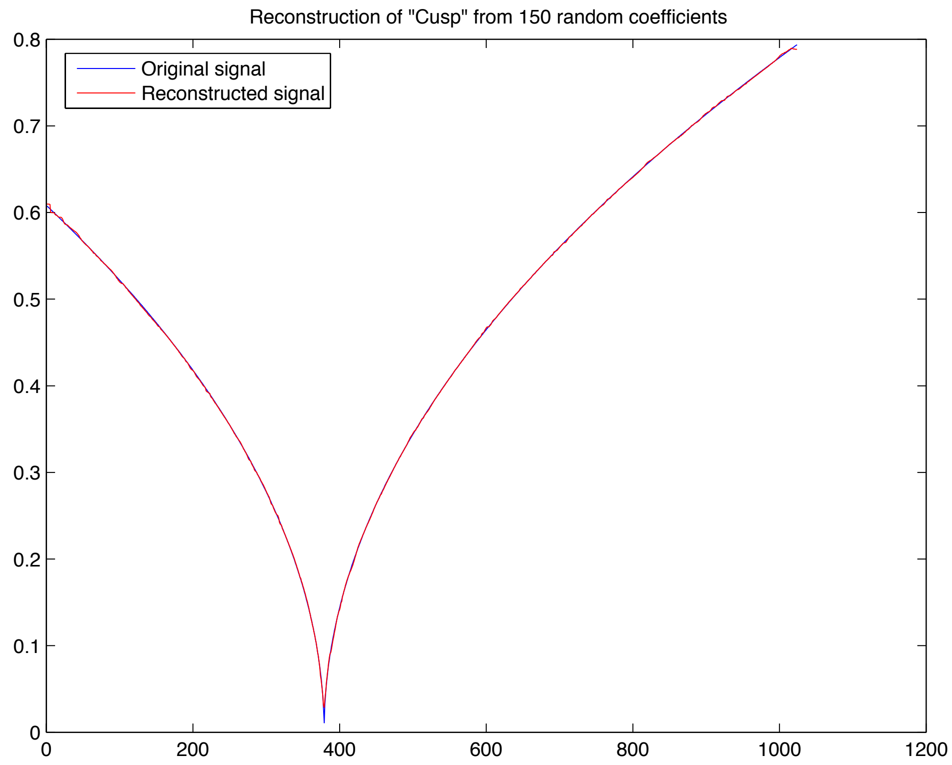

Reconstruction from Random Coefficients

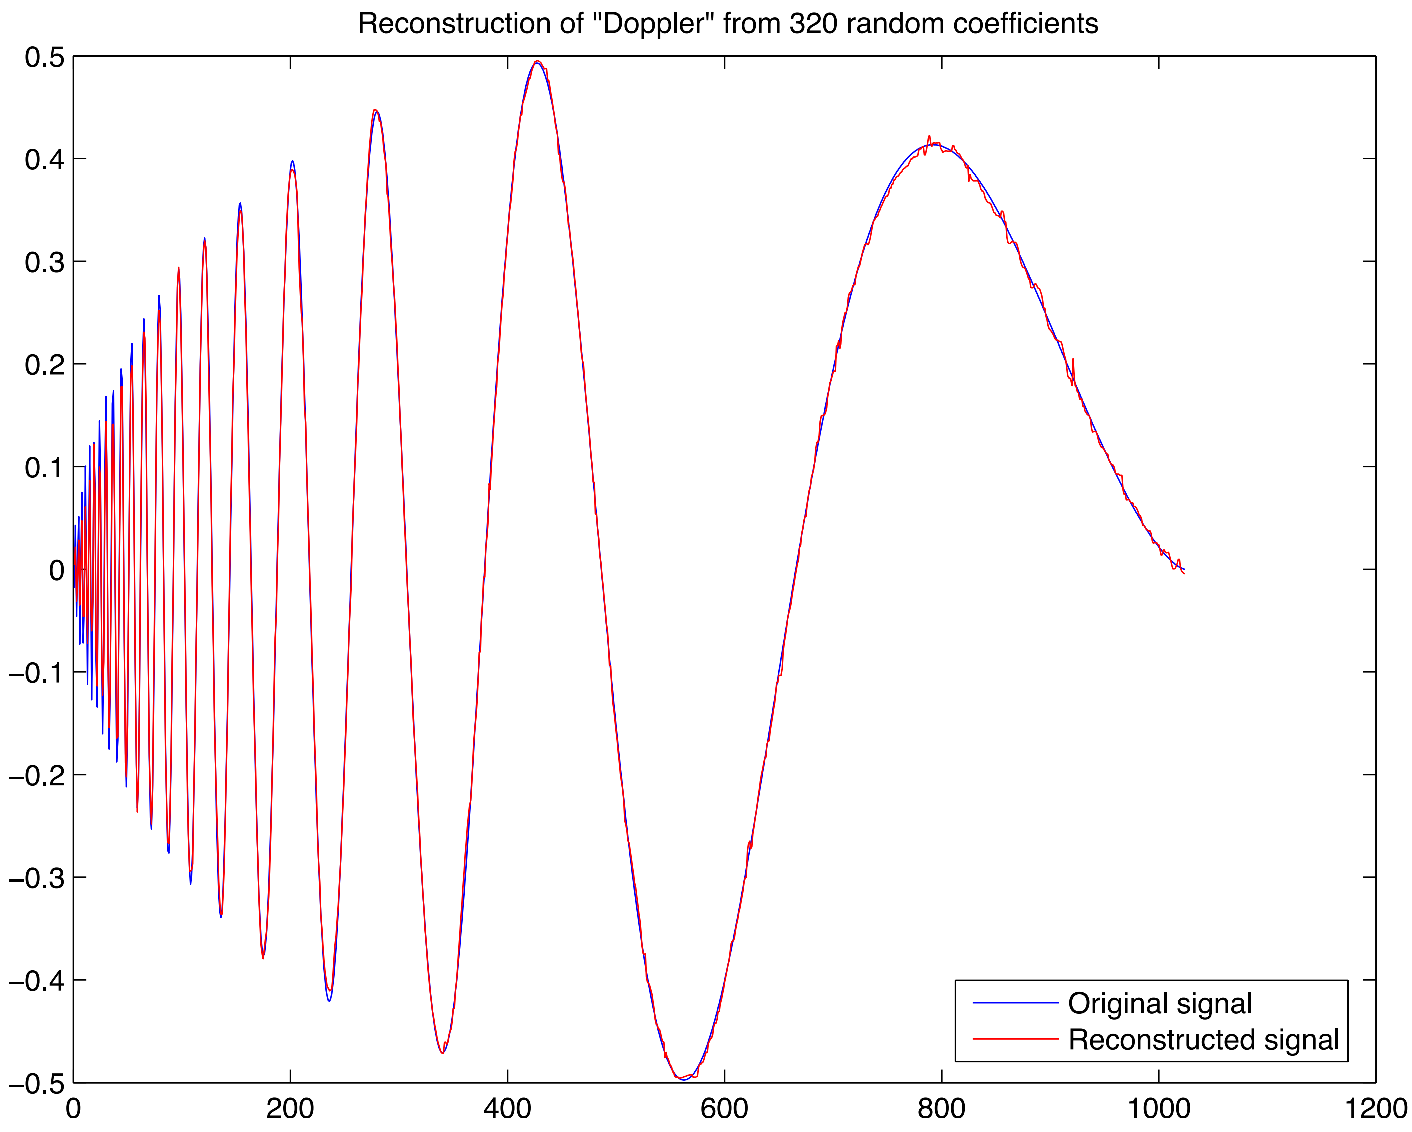

original, 65k pixels

wavelet 7207-term approx

recovery from 20k proj

⇓ zoom

⇓ zoom

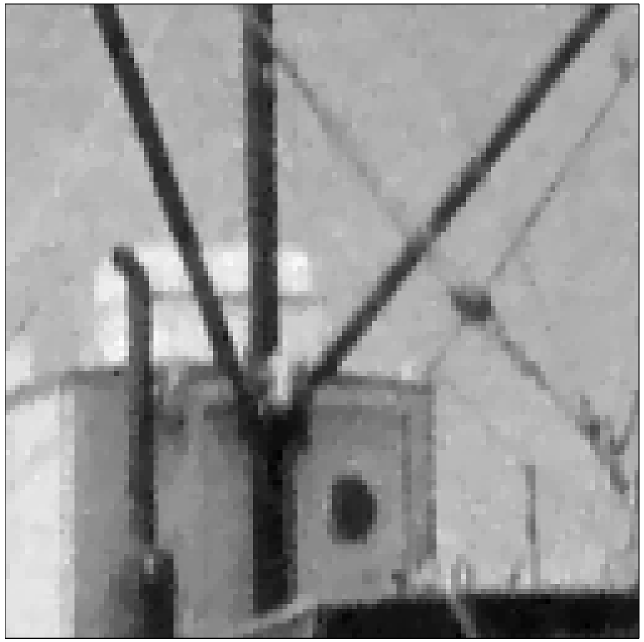

original

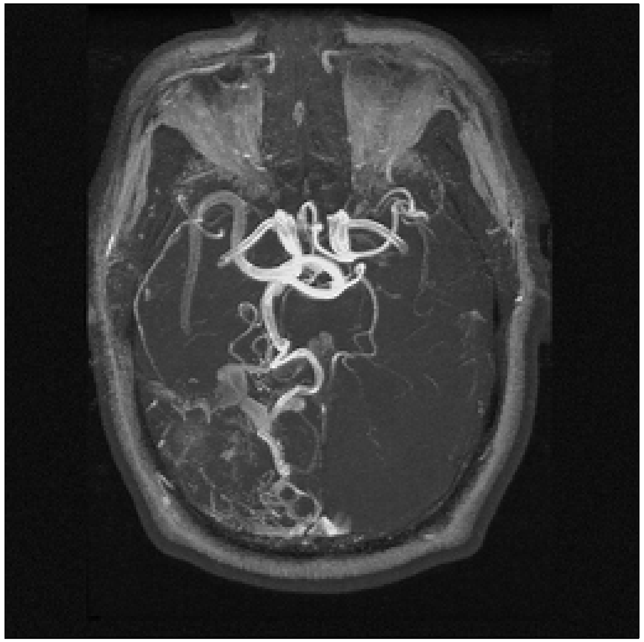

backprojection

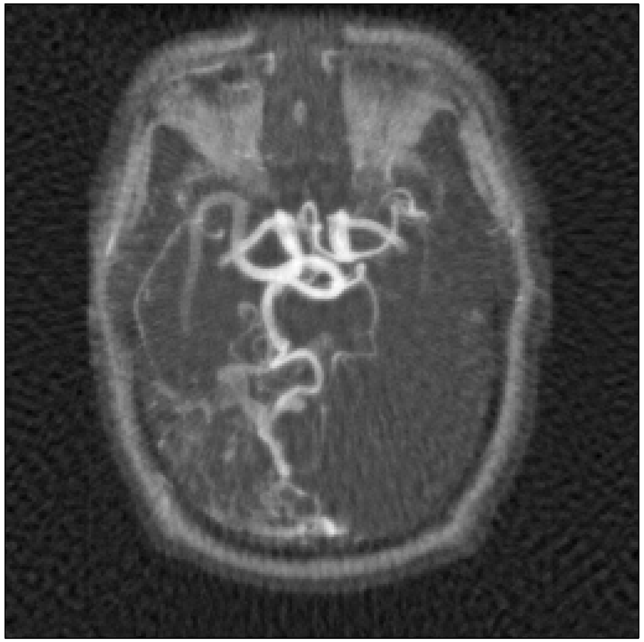

min TV

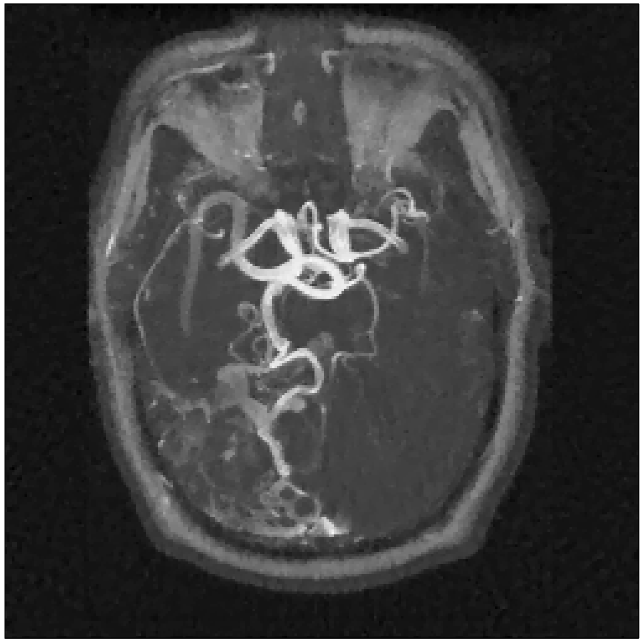

$\Omega \approx 29\%$ of samples

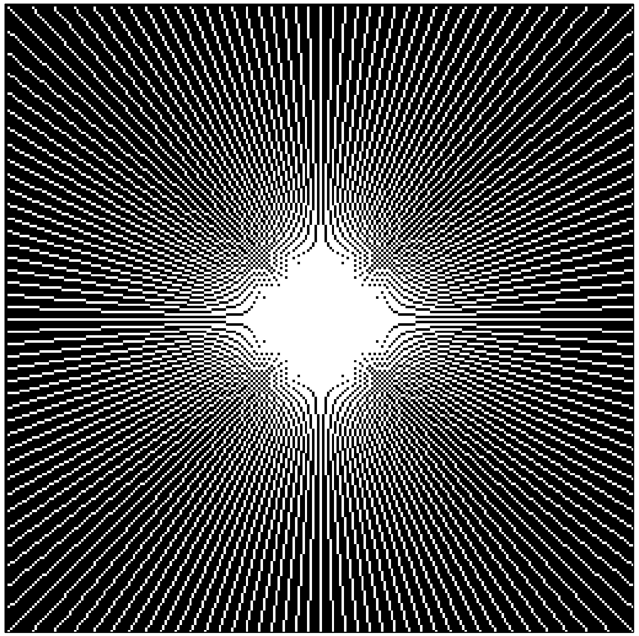

$\Downarrow$ zoom

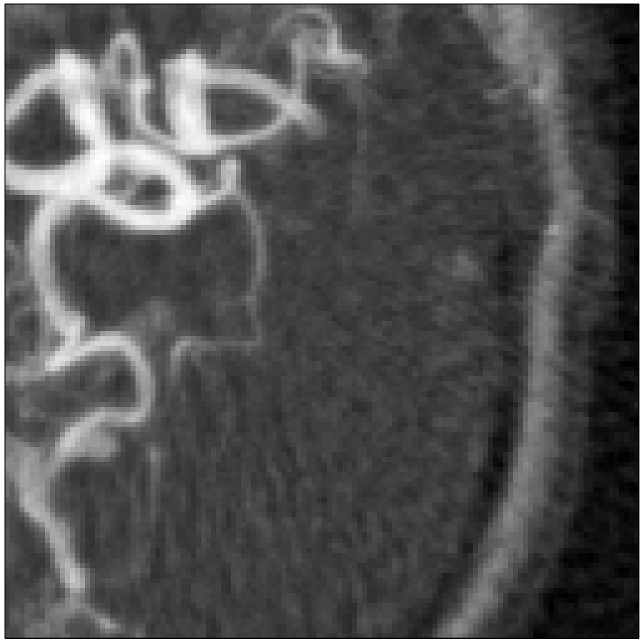

$\Downarrow$ zoom

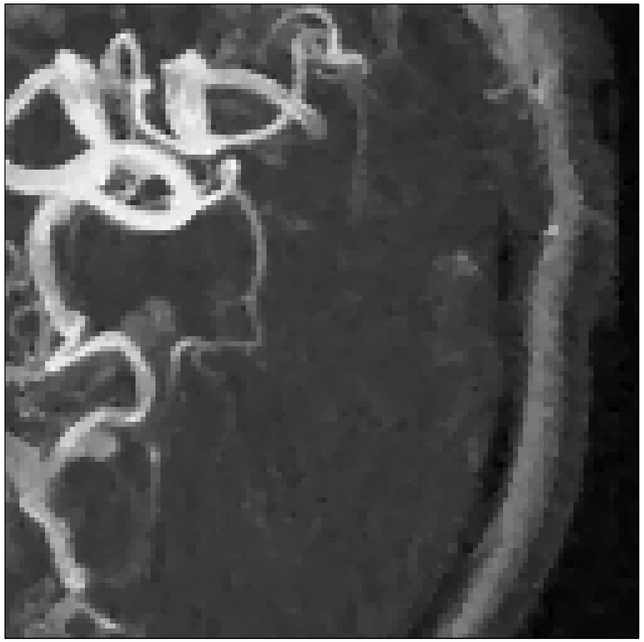

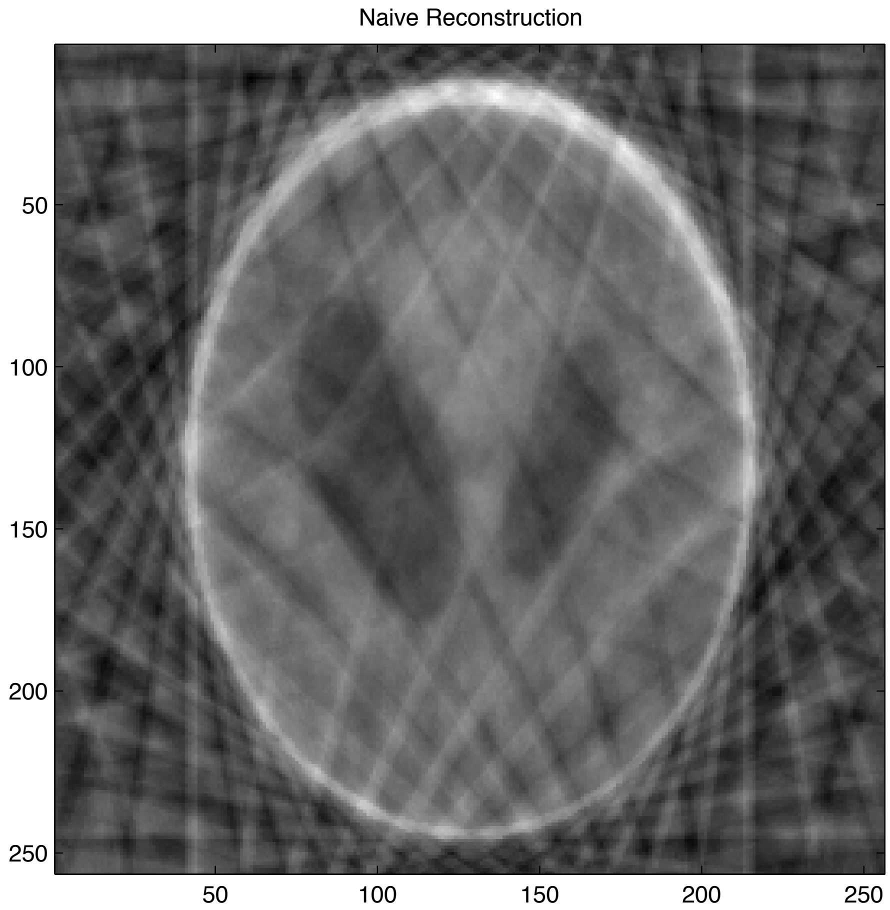

min ℓ_2

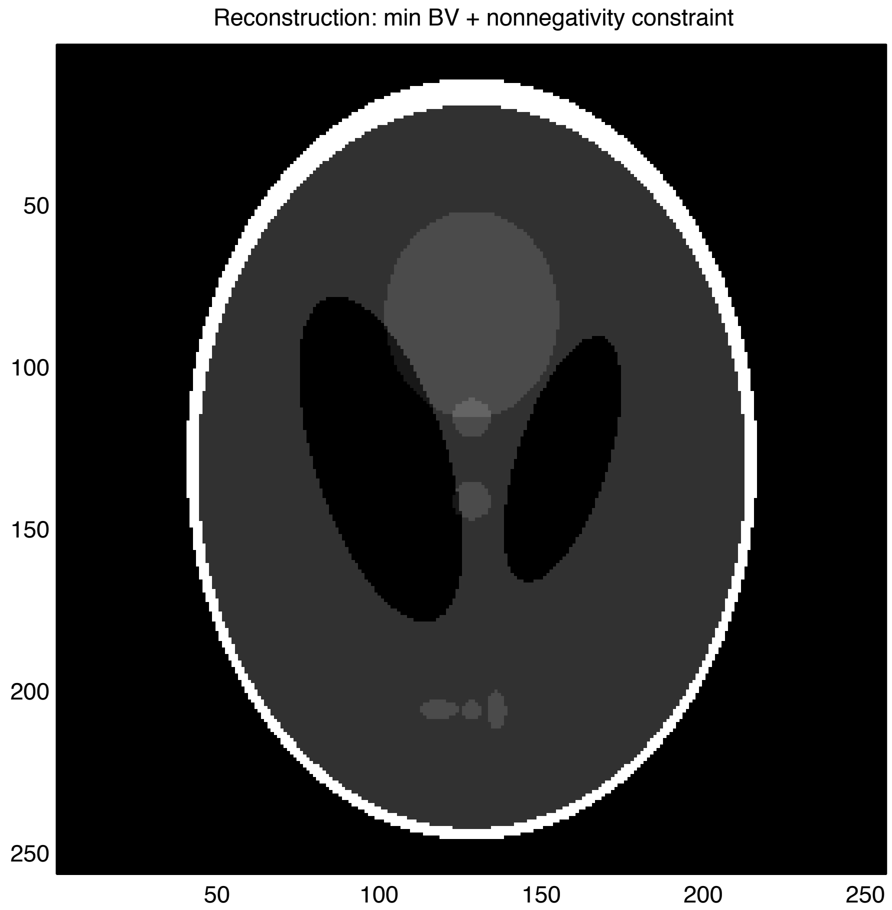

min TV – Exact!

Other Phantoms

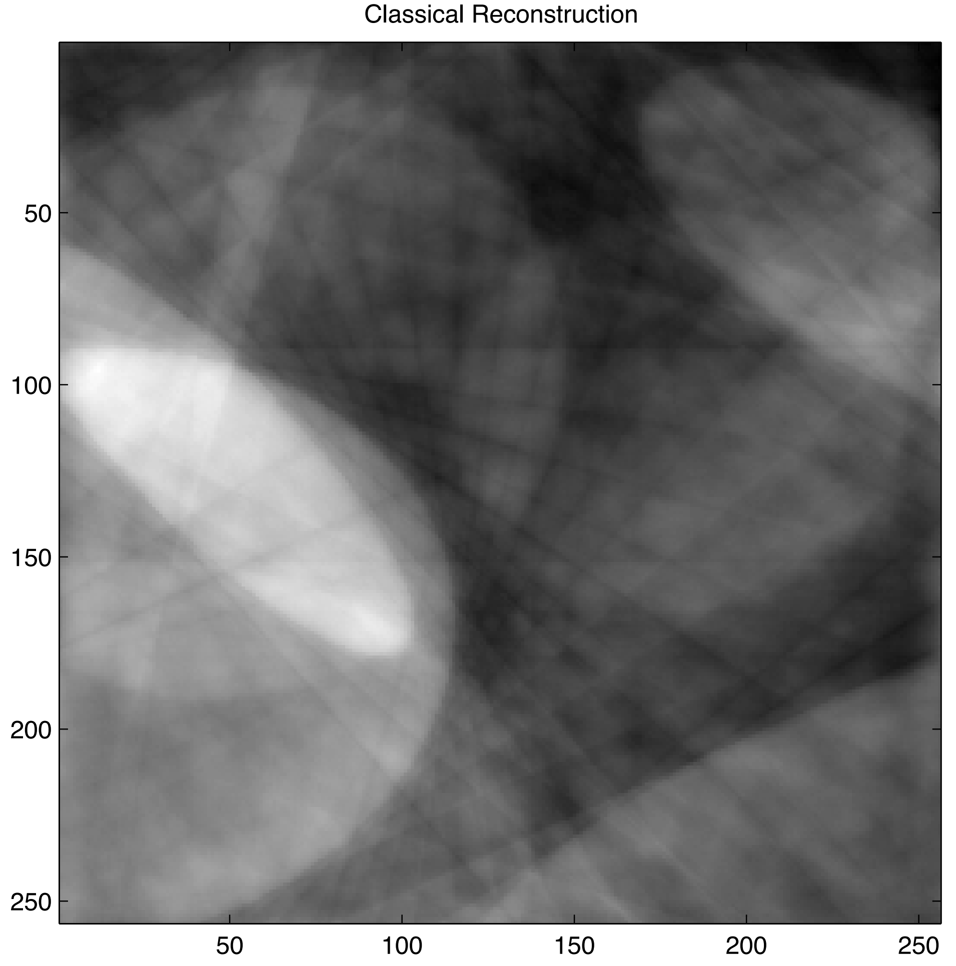

min ℓ_2

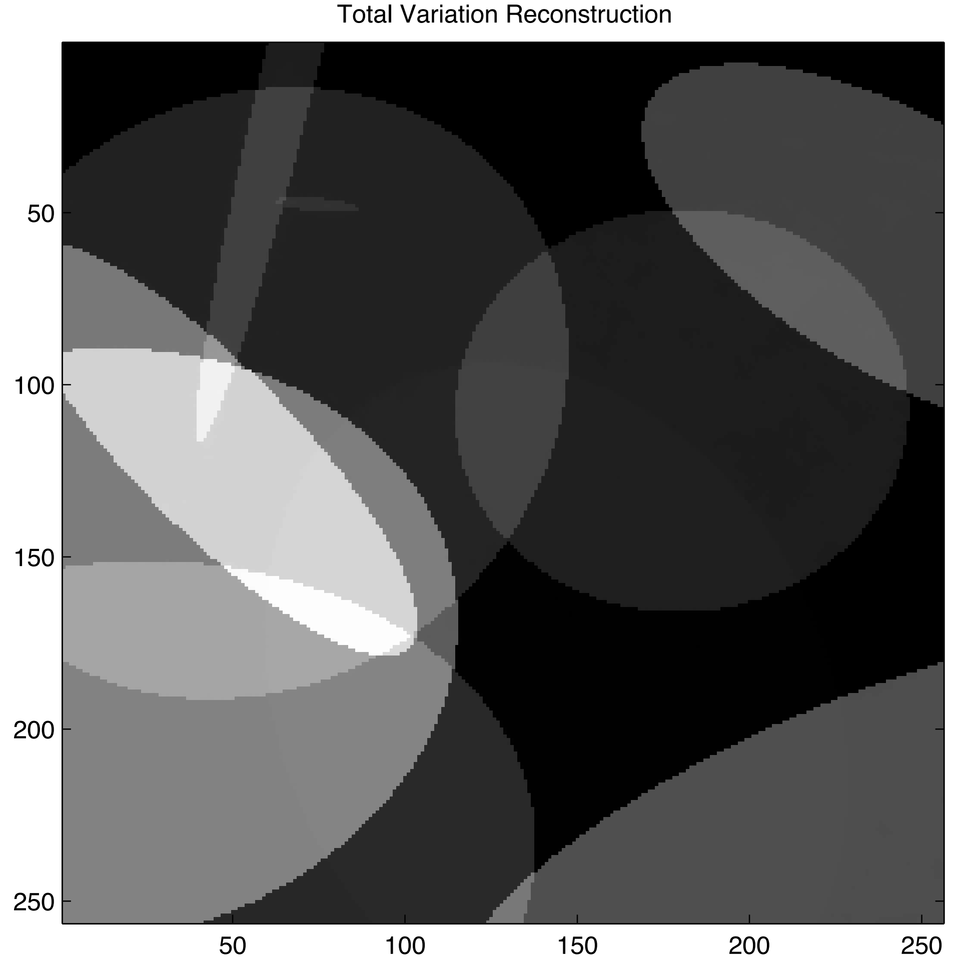

min TV – Exact!

Error Correction: Epilogue

- Gaussian coding matrix $A \in \mathbf{R}^{m \times n}$, A_{ij} i.i.d. $N(0,1)$
- Corrupted entries $y = Af + e$.
- Annihilator: $BA = 0$, B random projection on a plane of co-dimension n
- Can think of the entries of $B \in \mathbf{R}^{(m-n)\times m}$ i.i.d. $N(0,1)$.
- Decoding by LP is exact if e is the unique solution to

$$\min \|d\|_{\ell_1}, \qquad Bd = Be$$

- Need $\delta_{3S} + 3\delta_{4S} < 2$

Theorem 5 (C., Tao, 2004) *Exact decoding occurs for all corruption patterns and all plaintexts (with overwhelming probability) if the fraction ρ of error obeys*

$$\rho \lesssim \frac{1}{\log(\frac{m}{m-n})} \lesssim \frac{1}{\log(\frac{1}{1-n/m})} = \rho^*$$

For finite values of n/m (rate), the constant also matters! See Donoho (2004, 2005).

Summary

- Possible to reconstruct a sparse/compressible signal from very few measurements
- Need to solve an LP (or SOCP)
- Tied to new uncertainty principles
- Many applications
  - Finding sparse decompositions in overcomplete dictionaries
  - Decoding of linear codes
  - Biomedical imagery
- Extraordinary opportunities
  - New A/D devices
  - New paradigms for sensor networks

Connections with Information Theory (Mostly Speculative)

Universal Codes

Want to compress sparse signals

- *Encoder*. To encode a discrete signal $\boldsymbol{f}$, the encoder simply calculates the coefficients $\boldsymbol{y_k = \langle f, X_k \rangle}$ and quantizes the vector $\boldsymbol{y}$.
- *Decoder*. The decoder then receives the quantized values and reconstructs a signal by solving the linear program $(\boldsymbol{P_1})$.

Claim: Asymptotically nearly achieves the information theoretic limit.

Information Theoretic Limit: Example

- Want to encode the unit-ℓ_1 ball: $f \in \mathbf{R}^N : \sum_t |f(t)| \leq 1$.
- Want to achieve distortion D

$$\|f - f^\sharp\|^2 \leq D$$

- How many bits? Lower bounded by entropy of the unit-ℓ_1 ball:

$$\#\text{ bits } \geq C \cdot D \cdot (\log(N/D) + 1)$$

- How many bits does the universal encoder need? Up to possibly a $\log(1/D)$ factor

$$\#\text{ bits } \sim D \cdot (\log(N/D) + 1)$$

  - Same as the number of measurements
  - Robustness vis a vis quantization

Robustness

- Say with K coefficients

$$\|f - f^\sharp\|^2 \asymp 1/K$$

- Say we loose half of the bits (packet loss). How bad is the reconstruction?

$$\|f - f^\sharp_{50\%}\|^2 \asymp 2/K$$

- Democratic!

Why Does This Work? Geometric Viewpoint

Suppose $f \in \mathbf{R^2}$, $f = (\mathbf{0}, \mathbf{1})$.

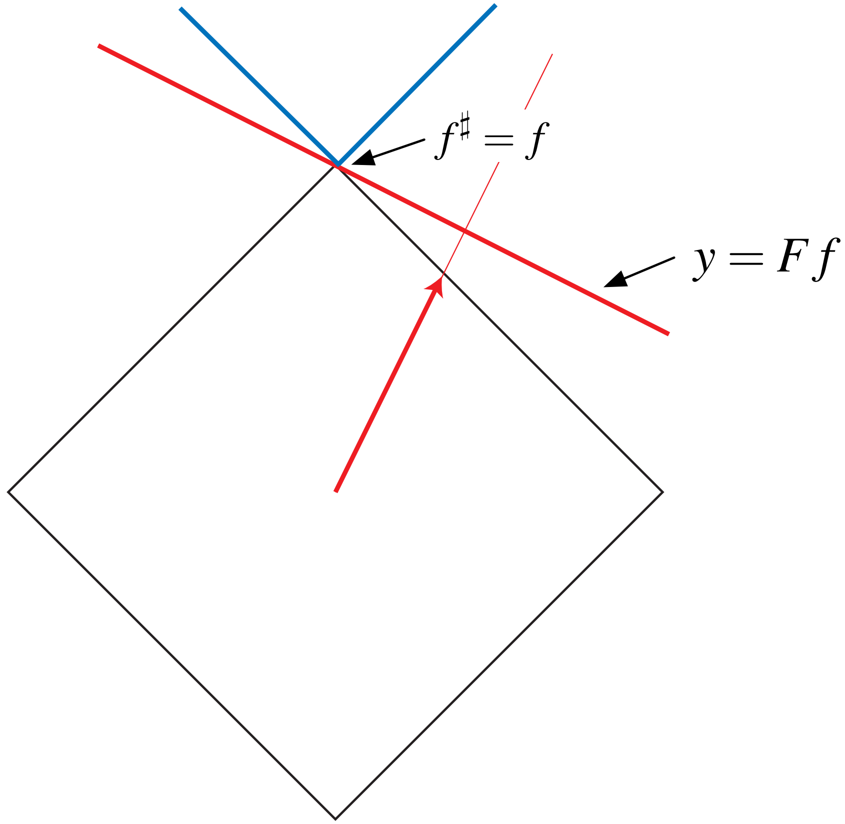

Exact

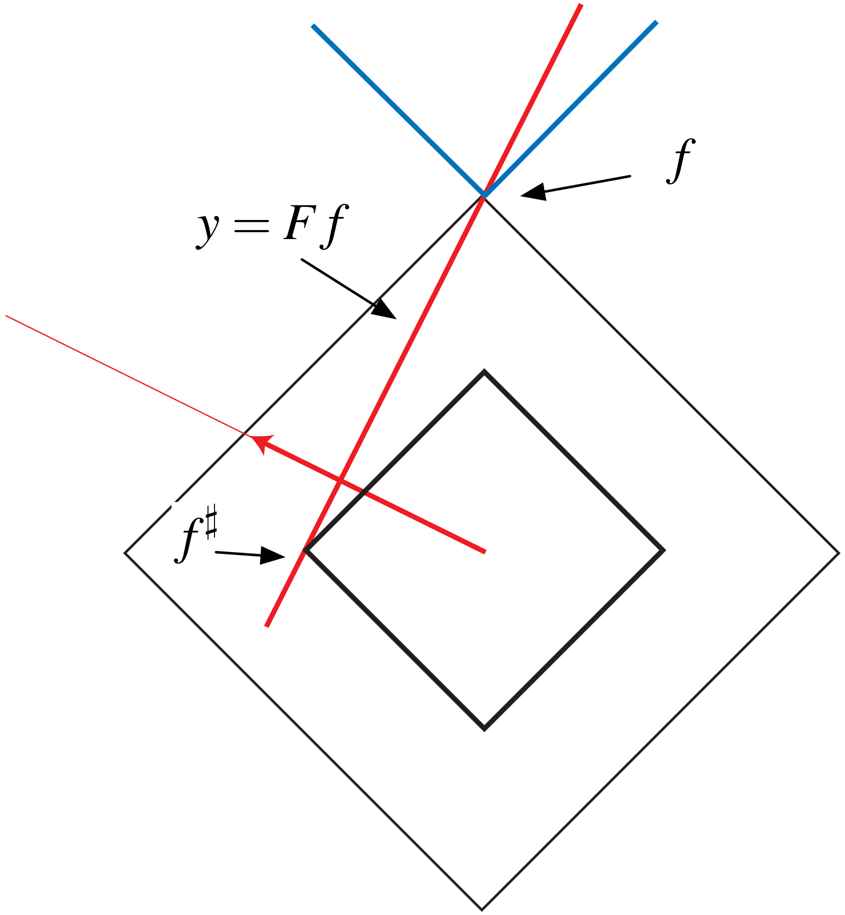

Miss

Higher Dimensions

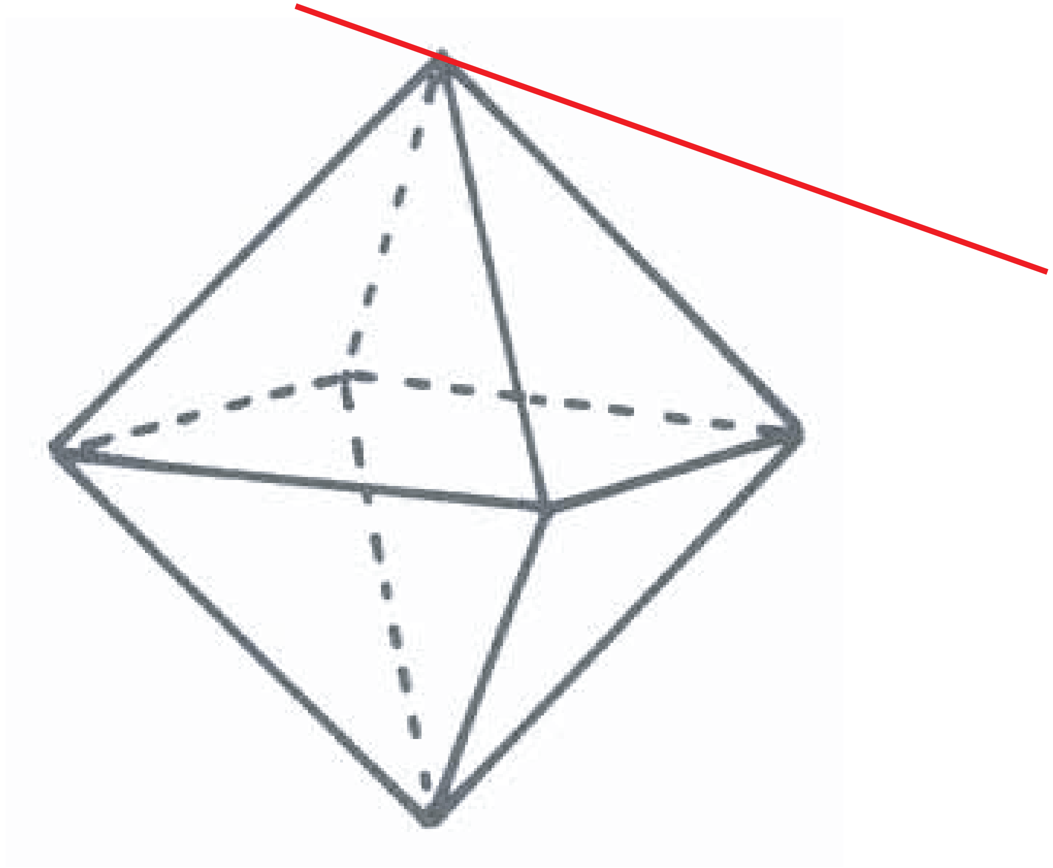

Equivalence

- Combinatorial optimization problem

$$(P_0) \qquad \min_g \|g\|_{\ell_0} := \#\{t,\, g(t) \neq 0\}, \qquad Fg = Ff$$

- Convex optimization problem (LP)

$$(P_1) \qquad \min_g \|g\|_{\ell_1}, \qquad Fg = Ff$$

- Equivalence:

For $K \asymp |T| \log N$, the solutions to (P_0) and (P_1) are unique and are the same!