Overview of compressed sensing

Yoshiyuki Kabashima
Tokyo Institute of Technology

Outline

- Introduction of compressed sensing
  - Sparse signal, What is compressed sensing, Background 1：AD conversion and sampling theorem, Background 2：Statistical property and sparse representation
- Algorithms for signal recovery
  - Underdetermined linear system, Orthogonal Matching Pursuit (OMP), Iterative Hard Thresholding (IHT), Basis Pursuit and linear programming, IRLS, Approximate message passing (AMP)
- Methods for performance analysis/guarantee
  - Basic method：Mutual coherence
  - Advanced methods：RIP, Combinatorics, State evolution, Replica method
- Summary and discussion
  - What I talked about, what I did not talk about
  - What I am interested in

INTRODUCTION OF COMPRESSED SENSING

Sparse signal

- Definition: Signal that is composed of many zeros when it is expressed by a certain basis
- “Sparseness” can be useful for various purposes of signal processing

Ex）Noise removal

$y = \cos(2\pi \times \underline{4} \times x/128) + \underline{0.2\eta(x)}$

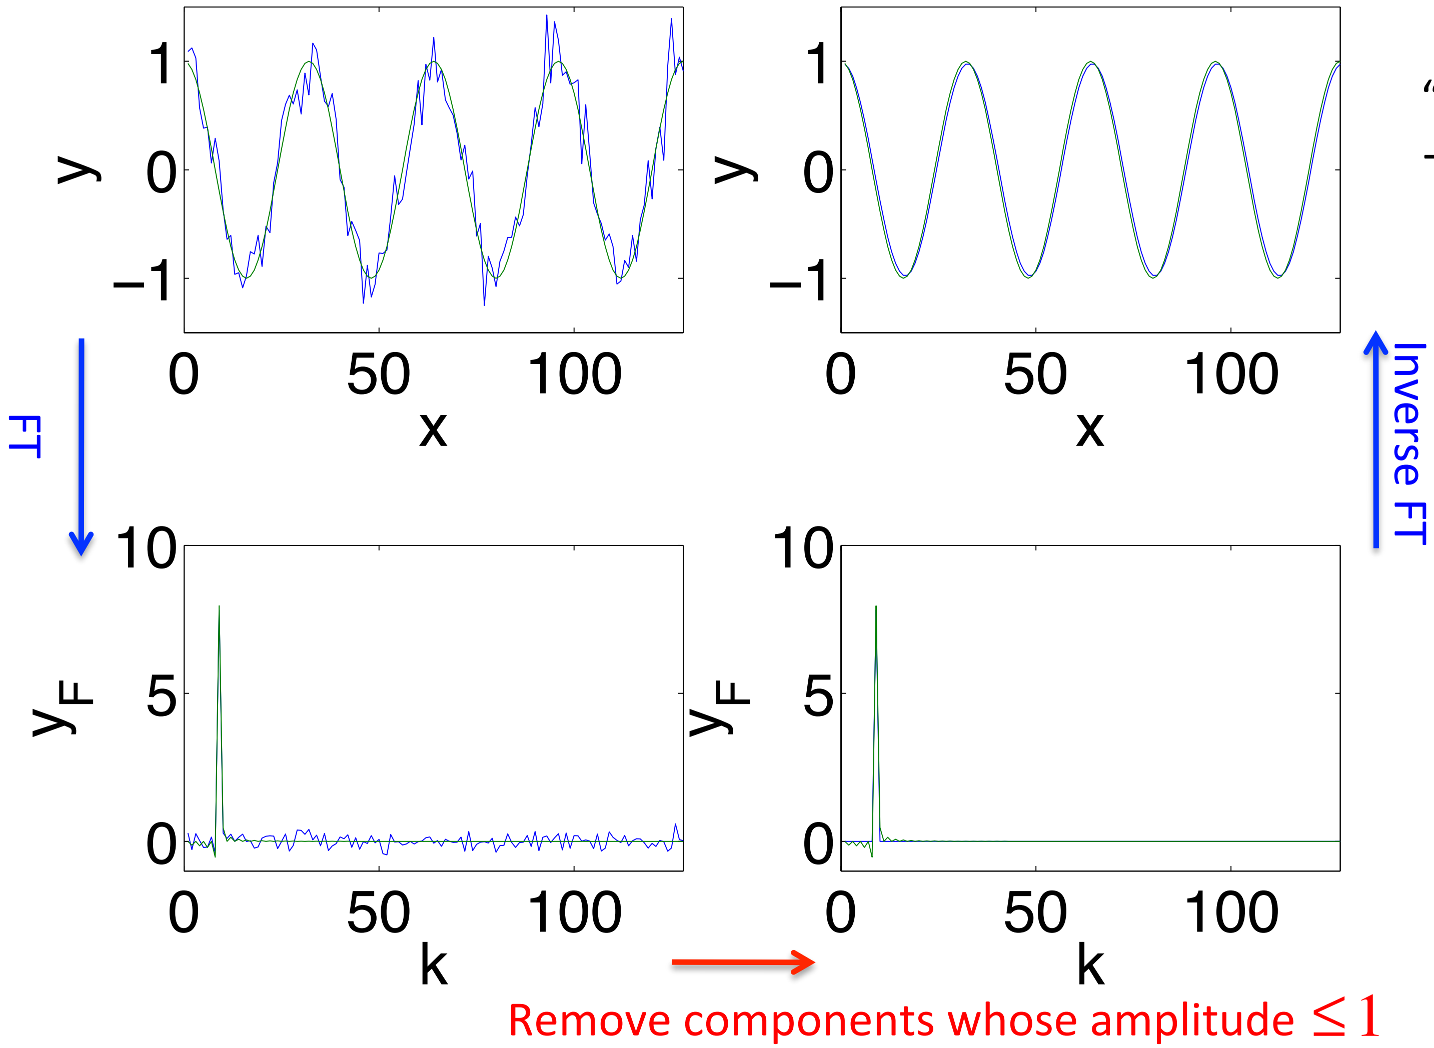

"Sparse"
→ Intensity localized

Small components
→ "Noise"

Remove small components in sparse representation

Noise removal

Ex) Data compression

[Romberg and Wakin, 2007]

Ex) Data compression

[Romberg and Wakin, 2007]

wavelet
transformation

Wavelet coefficients ($\log_{10}$)

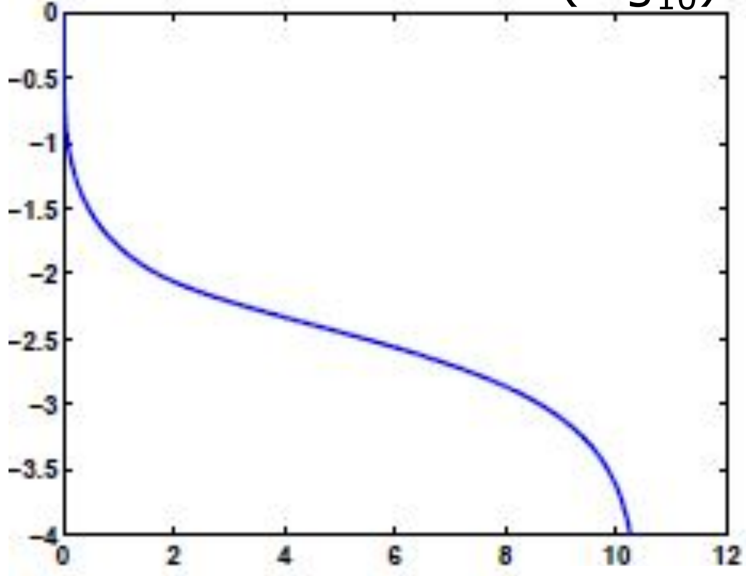

$\times 10^6$

Ex）Data compression

[Romberg and Wakin, 2007]

wavelet transformation

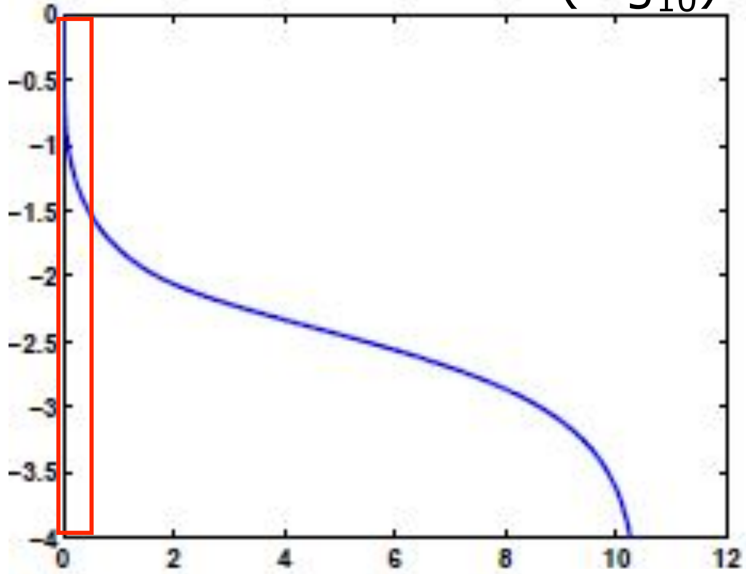

Ex）Data compression

[Romberg and Wakin, 2007]

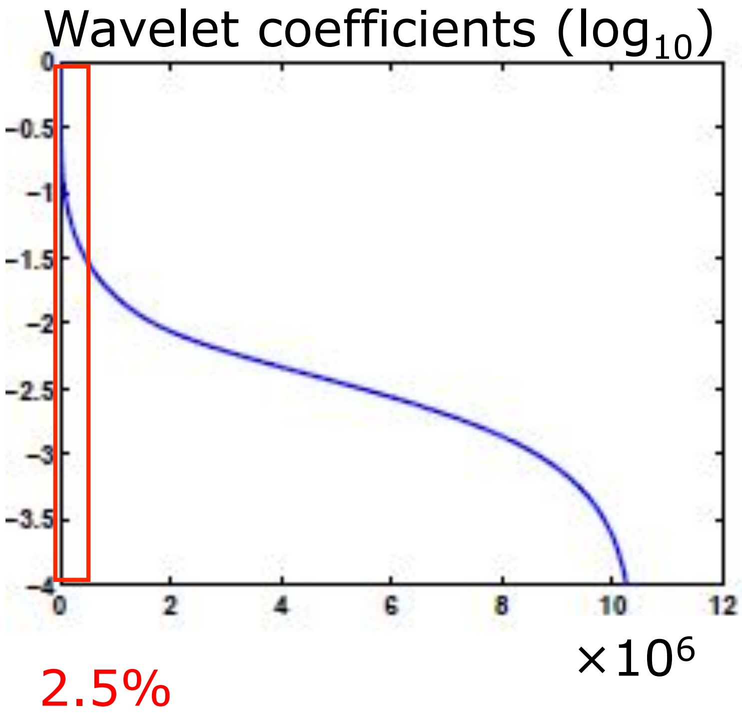

Ex) Data compression

[Romberg and Wakin, 2007]

Original Image

Reconstructed image from 2.5% coefficients

What is compressed sensing?

- A framework that enables signal recovery of the sparse signals from a fewer number of measurements than *conventional theory requires.*

- Application domain
  - Refraction seismic survey (mine examination)
  - Tomography（X-ray CT，MRI）
  - Single pixel camera
  - Noise removal of image
  - Data streaming computing
  - Group testing
  - etc.

Candes-Ronberg-Tao (2006)

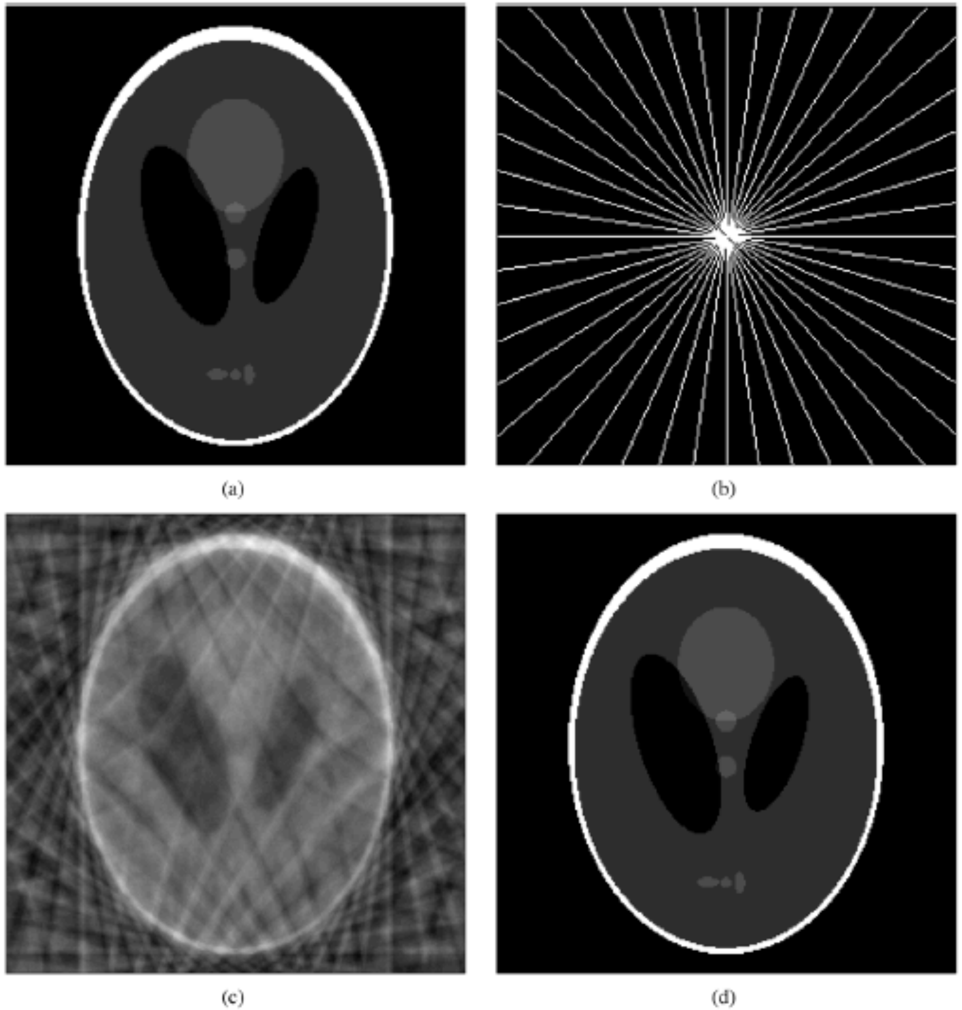

Simulation of tomography

LT : Original(Logan-Shepp Phantom)
#512x512

RT : Sampling 512 points of 2D FT from 22 directions.

LB : Recovery of pseudo-inverse (standard)

RB : Recovery utilizing the “sparseness” of spatial variations. “Original” is perfectly recovered.

Perfect recovery is realized by 1/50 samples of what Nyquist-Shannon’s theory requires.
→Breaking of the conventional limit!

EJ Candes J Romberg and T. Tao, IEEE Trans. IT Vol. 52, 489—502 (2006)より

Single pixel camera

Gang Huang@Bell Lab et al.

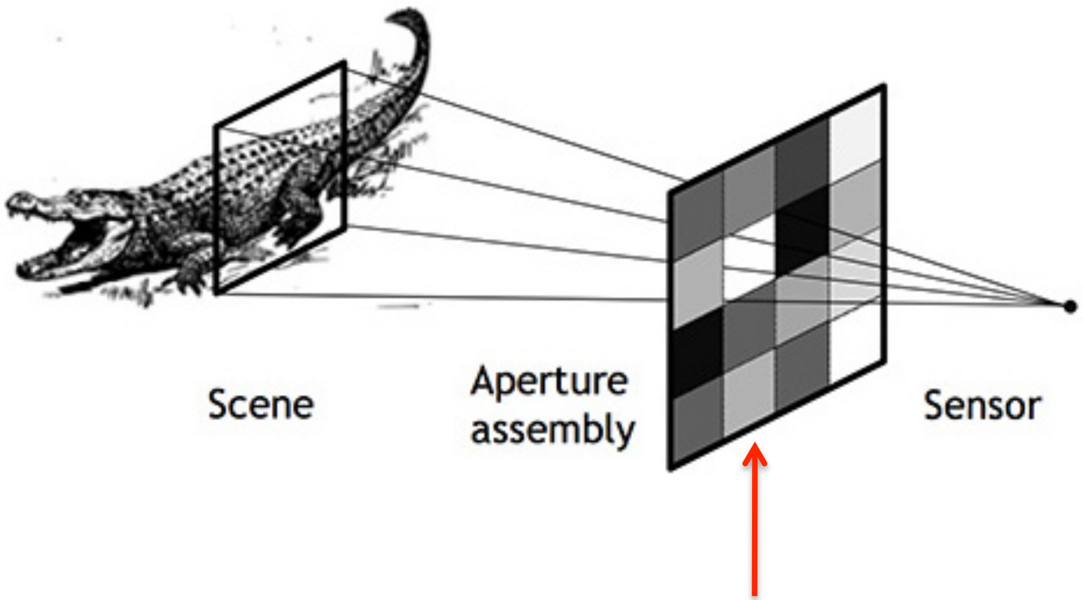

- Collect lights from various directions randomly changing aperture of each window.
- Measure the signals many times by a single sensor. #No need for “calibration”⇒Accurate measurement.

Bottom: The scene can be recovered from 1/4 data of convention.

Background 1: AD conversion and sampling theorem

- AD conversion: Analogue signal $y(t)$ is sampled at a fixed period T_s.
- Question: Can we <u>perfectly recover</u> $y(t)$ from the set of sample values $y(nT_s)$ $(n = 0, \pm 1, \pm 2, \ldots)$?

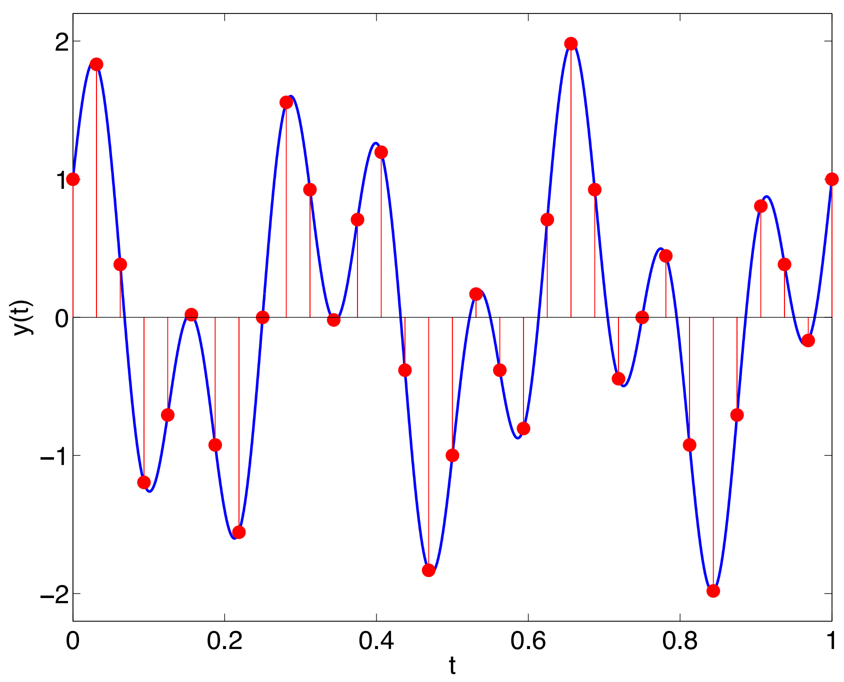

Sampling theorem

- Assumption: The maximum frequency of the original signal is f_m in the Fourier domain.
- **Proposition: The perfect recovery of $y(t)$ from $y(nT_s)$ $(n = 0, \pm 1, \pm 2, \ldots)$ is possible if the sampling frequency $f_s \equiv T_s^{-1}$ satisfies $f_s > 2f_m$.**

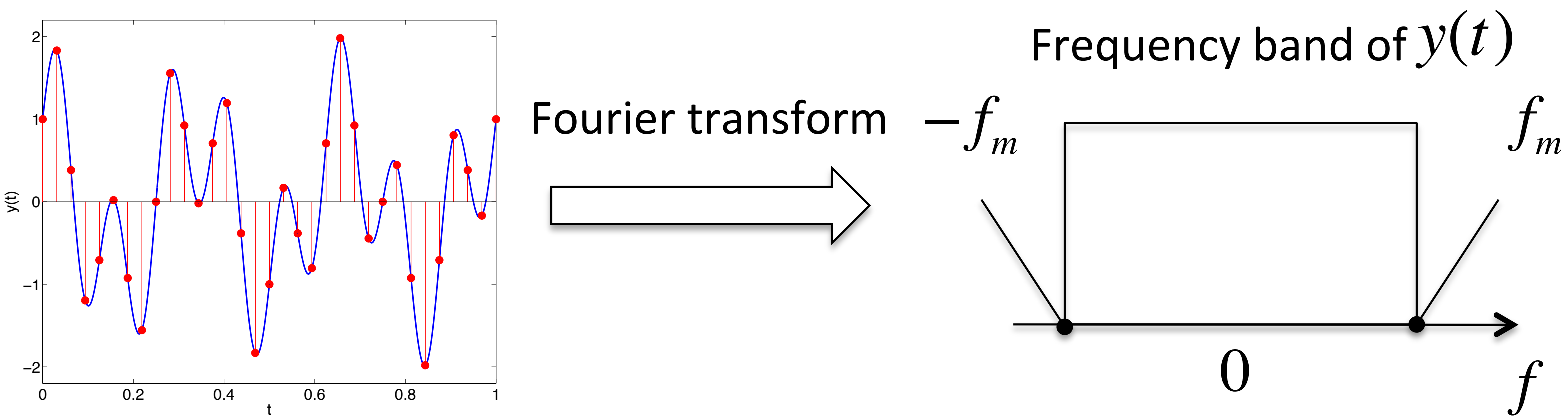

What sampling theorem means

- Let us assume that $y(t)$ is periodic with period T.
  →can be expressed as Fourier series.

$$y(t)=a_0+\sum_{k=1}^{k_m}\left(a_k\cos\left(\frac{2\pi kt}{T}\right)+b_k\sin\left(\frac{2\pi kt}{T}\right)\right)\quad(k_m:\ \text{Max. wave number})$$

- Perfectly identify $y(t)$
  ⇄Perfectly identify Fourier coefficients $\{a_0, a_1, b_1, a_2, b_2, \ldots, a_{k_m}, b_{k_m}\}$

$2k_m+1$ coefficients

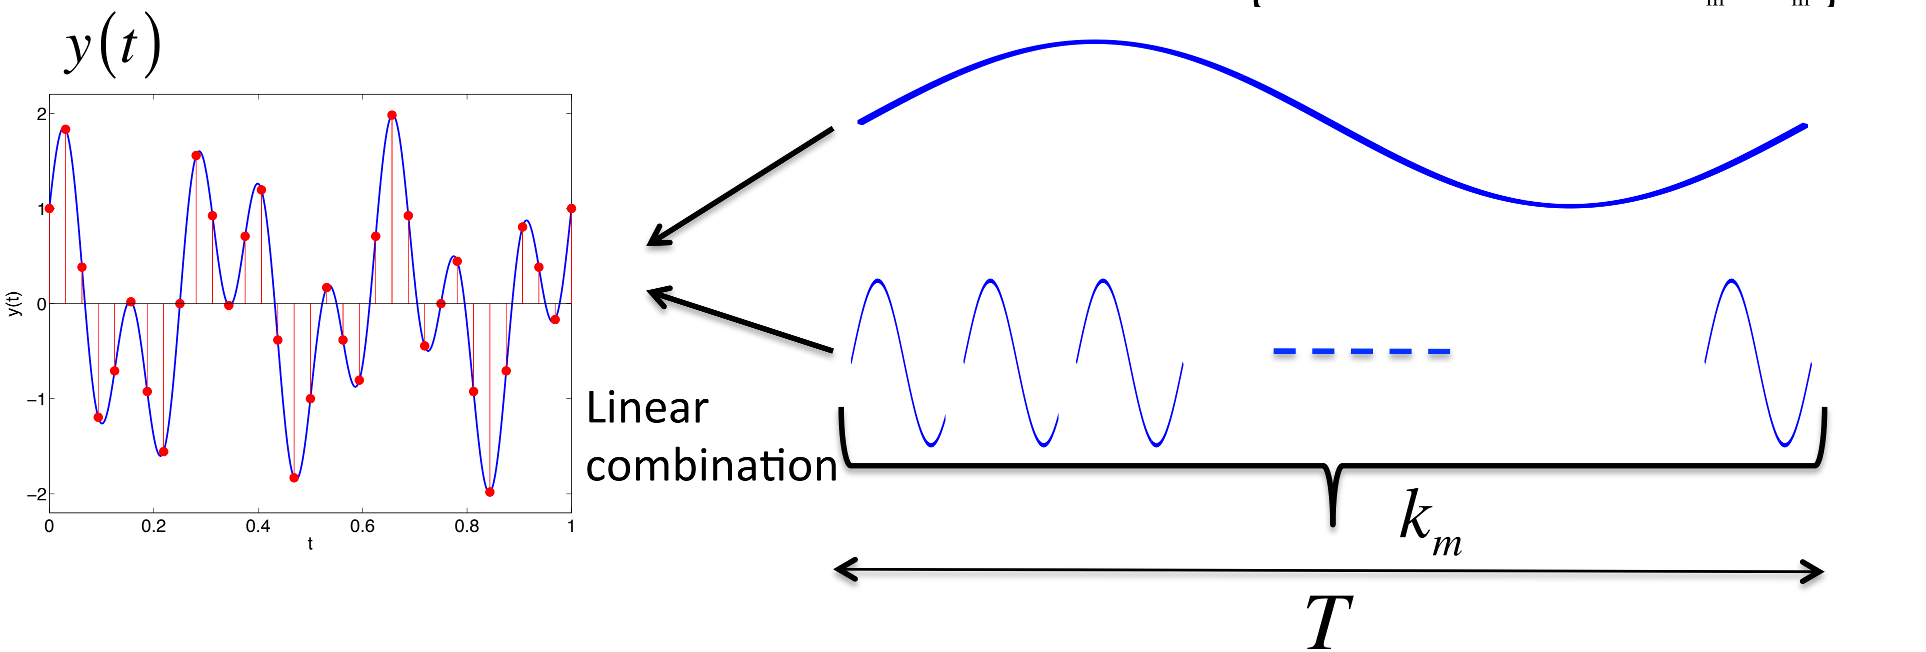

What sampling theorem means

- Sampling of a fixed period provides a set of independent linear equations for determining the Fourier coefficients.

$$y(nT_s) = a_0 + \sum_{n=1}^{k_m}\left(a_k \cos\left(\frac{2\pi knT_s}{T}\right) + b_k \sin\left(\frac{2\pi knT_s}{T}\right)\right)$$

- #Unknown variables $= 2k_m + 1$.

- This indicates that the perfect recovery is possible if the number of sampling in period T is greater than $2k_m + 1$.

Condition for getting a unique solution

#Equations $\geq$ #Unknown variables

Matrix expression

Trigonometric function of a fixed frequency

$$\begin{pmatrix} y\left(\frac{T}{2k_m+1}\right) \\ y\left(\frac{2T}{2k_m+1}\right) \\ y\left(\frac{3T}{2k_m+1}\right) \\ y\left(\frac{4T}{2k_m+1}\right) \\ \vdots \\ y\left(\frac{(2k_m+1)T}{2k_m+1}\right) \end{pmatrix} = \begin{pmatrix} 1 & \cos\left(\frac{2\pi}{2k_m+1}\right) & \sin\left(\frac{2\pi}{2k_m+1}\right) & \cos\left(\frac{4\pi}{2k_m+1}\right) & \cdots & \sin\left(\frac{2k_m\pi\cdot 1}{2k_m+1}\right) \\ 1 & \cos\left(\frac{2\pi\cdot 2}{2k_m+1}\right) & \sin\left(\frac{2\pi\cdot 2}{2k_m+1}\right) & \cos\left(\frac{4\pi\cdot 2}{2k_m+1}\right) & \cdots & \sin\left(\frac{2k_m\pi\cdot 2}{2k_m+1}\right) \\ 1 & \cos\left(\frac{2\pi\cdot 3}{2k_m+1}\right) & \sin\left(\frac{2\pi\cdot 3}{2k_m+1}\right) & \cos\left(\frac{4\pi\cdot 3}{2k_m+1}\right) & \cdots & \sin\left(\frac{2k_m\pi\cdot 3}{2k_m+1}\right) \\ 1 & \cos\left(\frac{2\pi\cdot 4}{2k_m+1}\right) & \sin\left(\frac{2\pi\cdot 4}{2k_m+1}\right) & \cos\left(\frac{4\pi\cdot 4}{2k_m+1}\right) & \cdots & \sin\left(\frac{2k_m\pi\cdot 4}{2k_m+1}\right) \\ 1 & \vdots & \vdots & \vdots & \cdots & \vdots \\ 1 & \cos\left(\frac{2\pi\cdot(2k_m+1)}{2k_m+1}\right) & \sin\left(\frac{2\pi\cdot(2k_m+1)}{2k_m+1}\right) & \cos\left(\frac{4\pi\cdot(2k_m+1)}{2k_m+1}\right) & \cdots & \sin\left(\frac{2k_m\pi\cdot(2k_m+1)}{2k_m+1}\right) \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ b_1 \\ a_2 \\ \vdots \\ b_{k_m} \end{pmatrix}$$

$2k_m+1$ sample values

$(2k_m+1)\times(2k_m+1)$ matrix

(invertible)

$2k_m+1$ unknown variables

(Fourier coefficients)

To get sampling theorem

Expression by sampling rate

$$N_{\text{sample}} \geq 2k_m + 1$$

Divide the both sides by T

$$f_s \equiv \frac{N_{\text{sample}}}{T} > \frac{2k_m + 1}{T} = 2\frac{k_m}{T} + \frac{1}{T}$$

Right figure means $f_m = \dfrac{k_m}{T}$

Take $T \to \infty$ for considering arbitrary signals

$$f_s > 2f_m + \frac{1}{T} \xrightarrow{T \to \infty} 2f_m$$

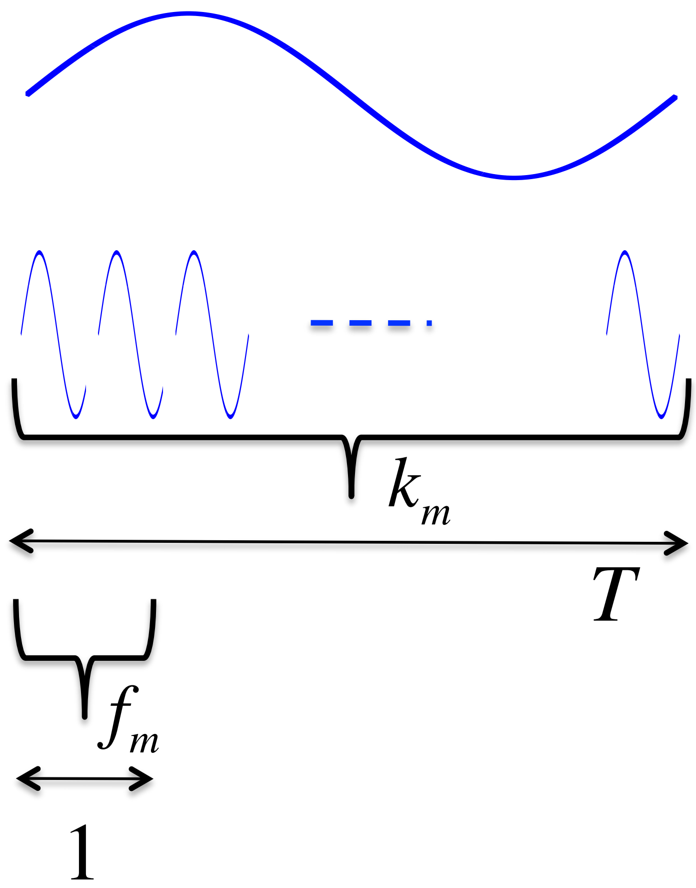

Origin of “2” is 2 varieties of “sin” and “cos”.

“sin, cos” or “+, -”

- Euler’s formula: $e^{i\theta} = \cos\theta + i\sin\theta$

$$\cos\theta = \frac{e^{i\theta} + e^{-i\theta}}{2} \qquad \sin\theta = \frac{e^{i\theta} - e^{-i\theta}}{2i}$$

- This enables us to interpret that “sin and cos” is a combination of “+ and –” of imaginary exponentials of an identical frequency. Since this interpretation is convenient in calculation, we often express the degree of freedom of signals using frequency band that symmetrically spreads in “+ and –” directions.

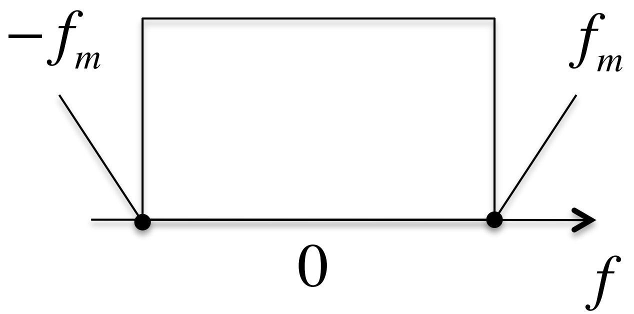

Consideration

- Sampling theorem gives the condition that the Fourier coefficients of the objective signal is *uniquely* determined.

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

Sample values (measured) → $\mathbf{y}$

Fourier basis (known) → F

Fourier coefficients (what we want to know) → $\mathbf{x}$

Consideration

- <u>The followings are not essential.</u>
  - Sampling of a fixed period
    - Only need #independent samples $\geq$ #<u>unknown Fourier coefficients.</u>
  - Fourier basis
    - One can multiply any invertible matrix to the both sides.
    - Any basis is OK. Even the orthogonality is not necessary for the basis.
    - More generally, Matrix = Measurement x Basis.

Measurement result

Measurement matrix

$$\mathbf{y}' = M\mathbf{y} = M\left(F\mathbf{x}\right) = A\mathbf{x}$$

- <u>From now on, we mathematically express the measurement problem as</u>

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

Background 2 :
Statistical property and sparse representation

- Keep in mind that the Sampling Theorem is a *worst case condition.*
- We can easily show an example to which the perfect recovery is possible by a fewer samples than that S.T. requires.

$$y(nT_s) = \sum_{n=k_0}^{k_m} \left(a_k \cos\left(\frac{2\pi knT_s}{T}\right) + b_k \sin\left(\frac{2\pi knT_s}{T}\right)\right)$$

For example, introduce “min. “ wave number

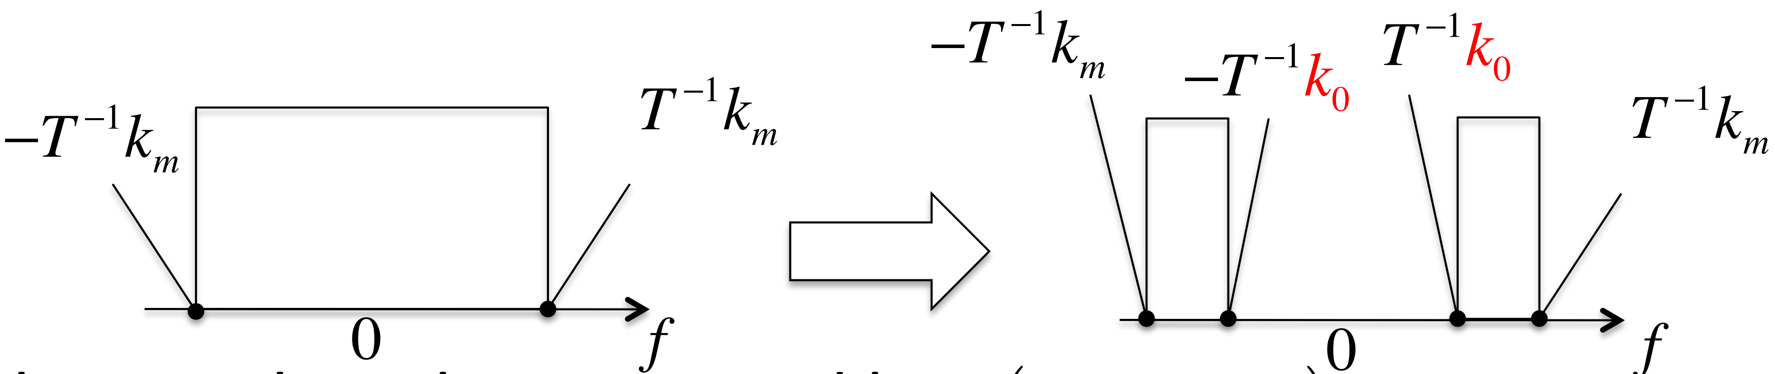

The signal can be recovered by $2(k_m - k_0 + 1)$ measurements.

Expression by equations

- This corresponds to a situation in which low frequency components are guaranteed to be zero.

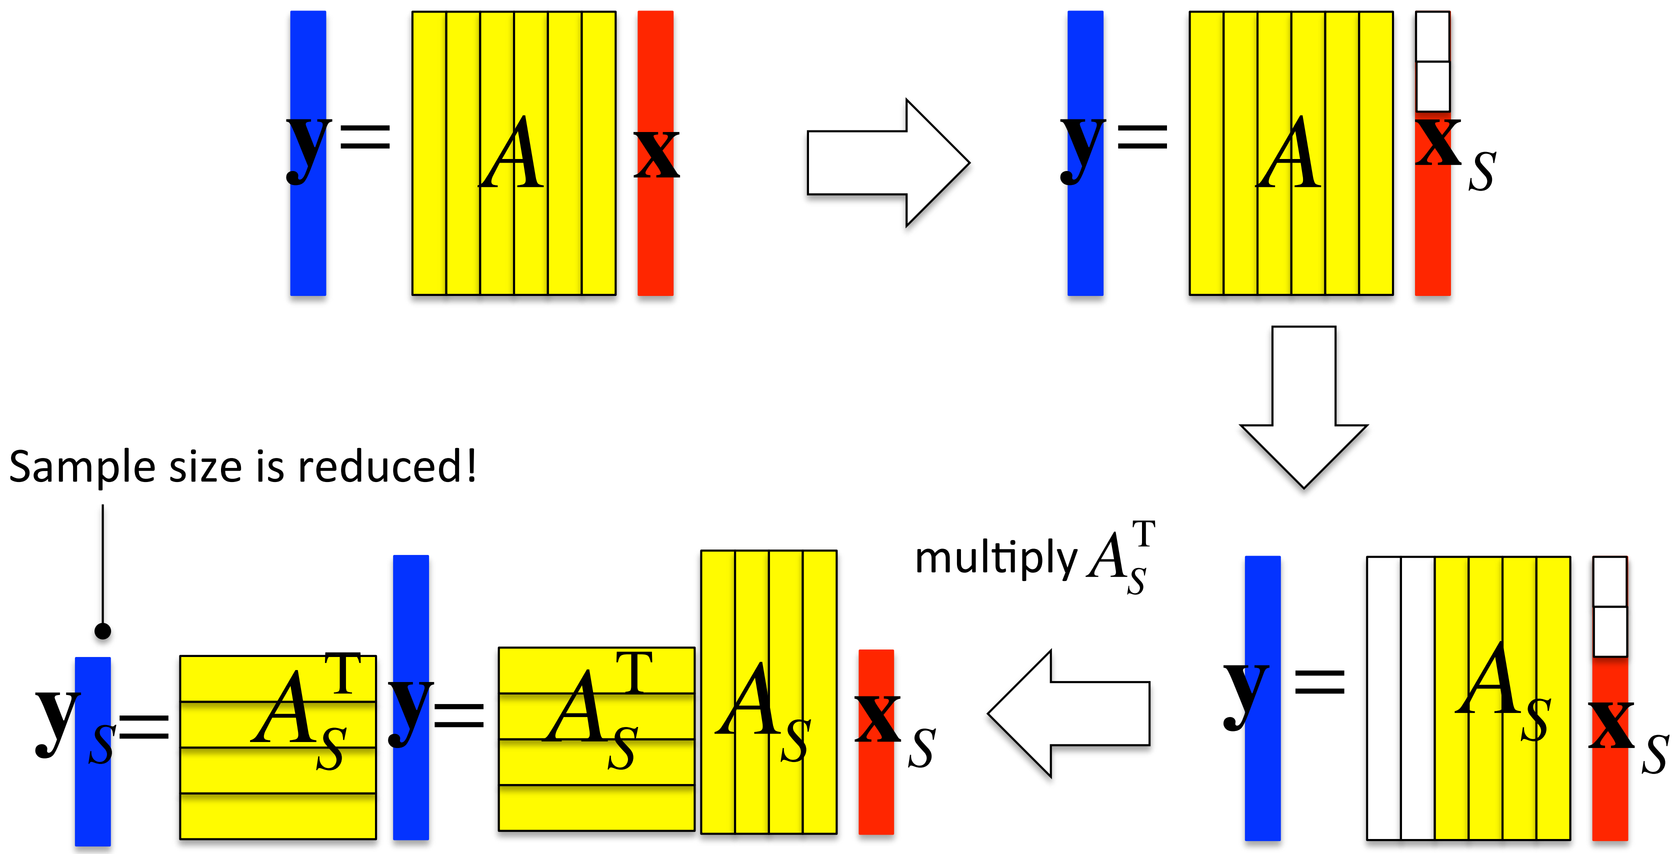

Generalization

- Cutoff of low frequencies may not be common in practice.
- But one can still empirically find many zeros below the max (Nyquist) frequency f_{m} for real world signals.
- Such cases, however, can be reduced to the previous example by reordering the positions if the zero positions are known.

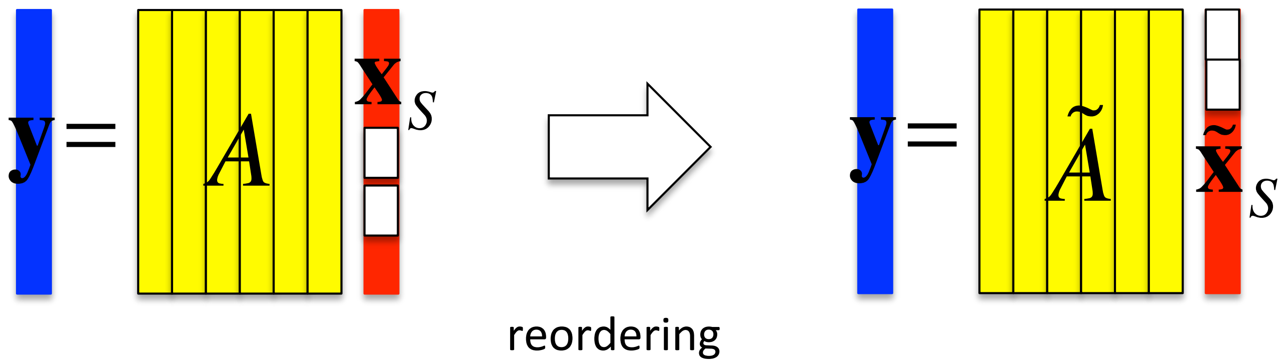

Compressed sensing

- Then, what can we do if the positions of zeros are *unknown*?
  - *Compressed sensing*
- Recovery problem of compressed sensing
  - Recover a sparse signal in which the positions of zeros are unknown from a fewer samples than S.T. requires.

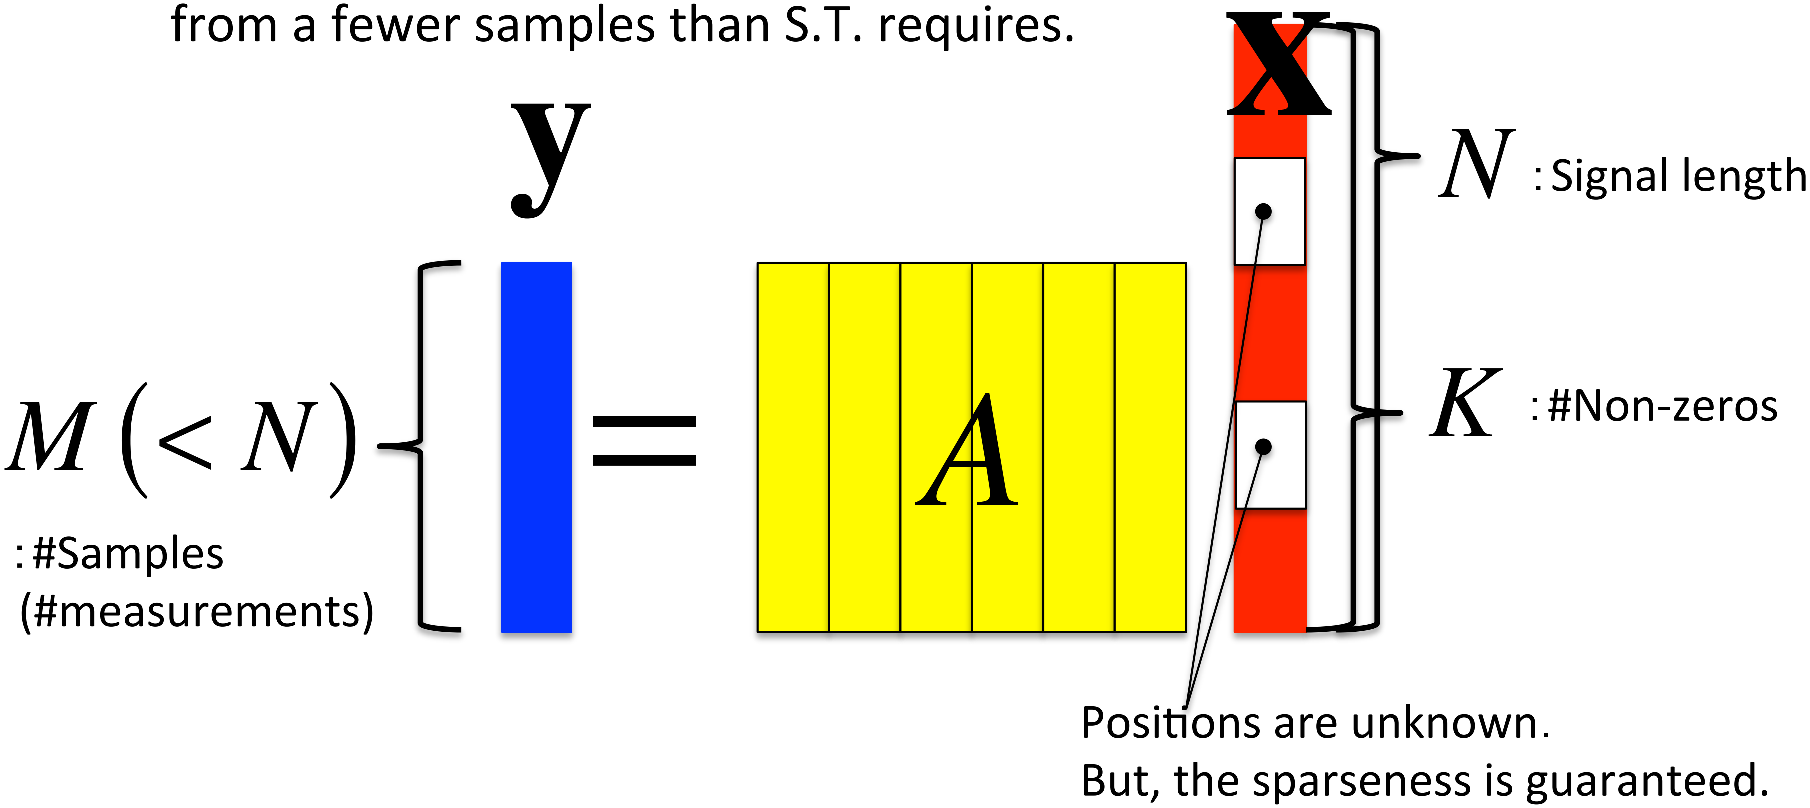

Practical motivation

- In practical signals in real world, many components usually vanish or nearly vanish statistically even below f_{m}.
  - Sound, image, seismic wave, ….
  - It is a waste of resources to take all the frequencies below f_{m} into account.

256x256

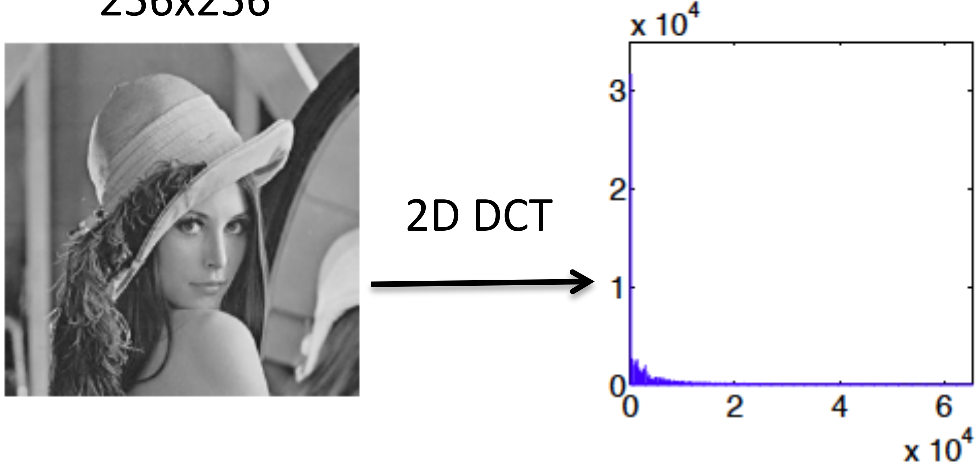

- This motivates us to efficiently recover the signal by identifying the zero components. ⇒ *Compressed sensing*

Two major points at issue

- Algorithm development
  - Develop practical procedures for doing it.
- Performance analysis/guarantee
  - Clarify the recovery condition for the developed algorithms.

ALGORITHMS FOR SIGNAL RECOVERY

Underdetermined linear system

- Compressed sensing problem is expressed as an underdetermined linear system.
  - For theoretical simplicity, we mainly focus on the case of *i.i.d. random matrices* from now on.

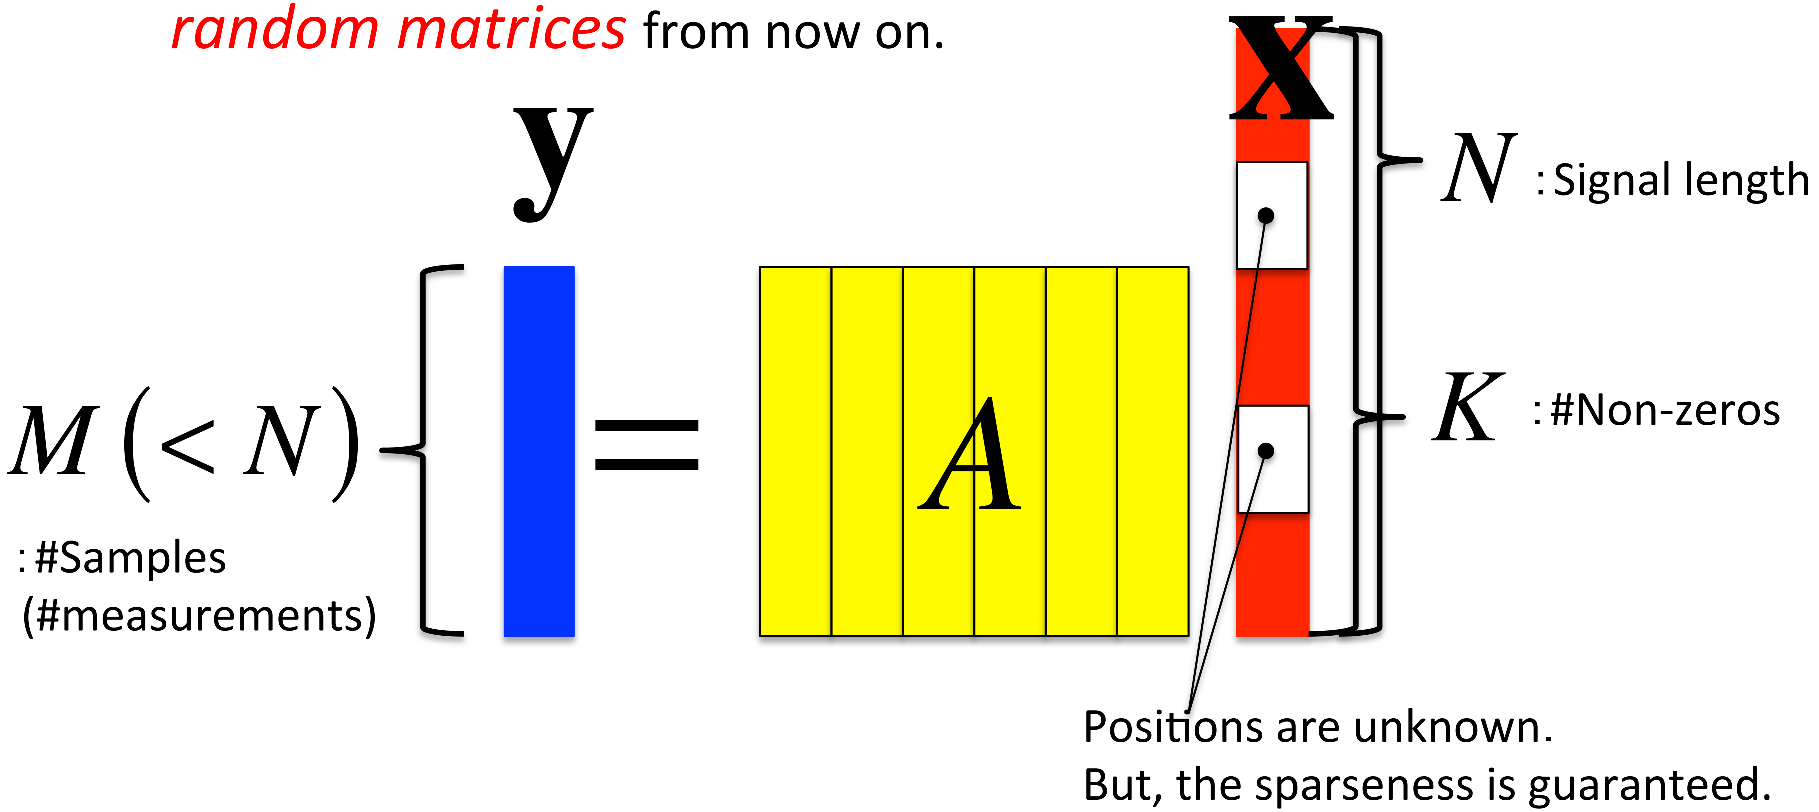

Regularization

- A standard method to solve such problems is to minimize an appropriate cost function $J(\mathbf{x})$ under the measurement constraint.

$$\left(P_J\right): \quad \min_{\mathbf{x}} J(\mathbf{x}) \ \text{subj. to } \mathbf{y} = A\mathbf{x}$$

  – What $J(\mathbf{x})$ should we use?

l_0-recovery

- Answer : The choice of $J(\mathbf{x}) = \|\mathbf{x}\|_0 \,(= \#\,\text{non-zeros in } \mathbf{x})$ is optimal <u>if we need not care about computational cost.</u>

  – Reason
    - $M > K \rightarrow$ Correct solution is always found.
    - $M \leq K \rightarrow$ Correct solution cannot be found *in principle* even if the positions of zeros are known.

l_0-recovery algorithm

#Assume that any combinations of $M-1$
#columns are linearly independent.

Try the following for all combinations S of $M-1$ columns.

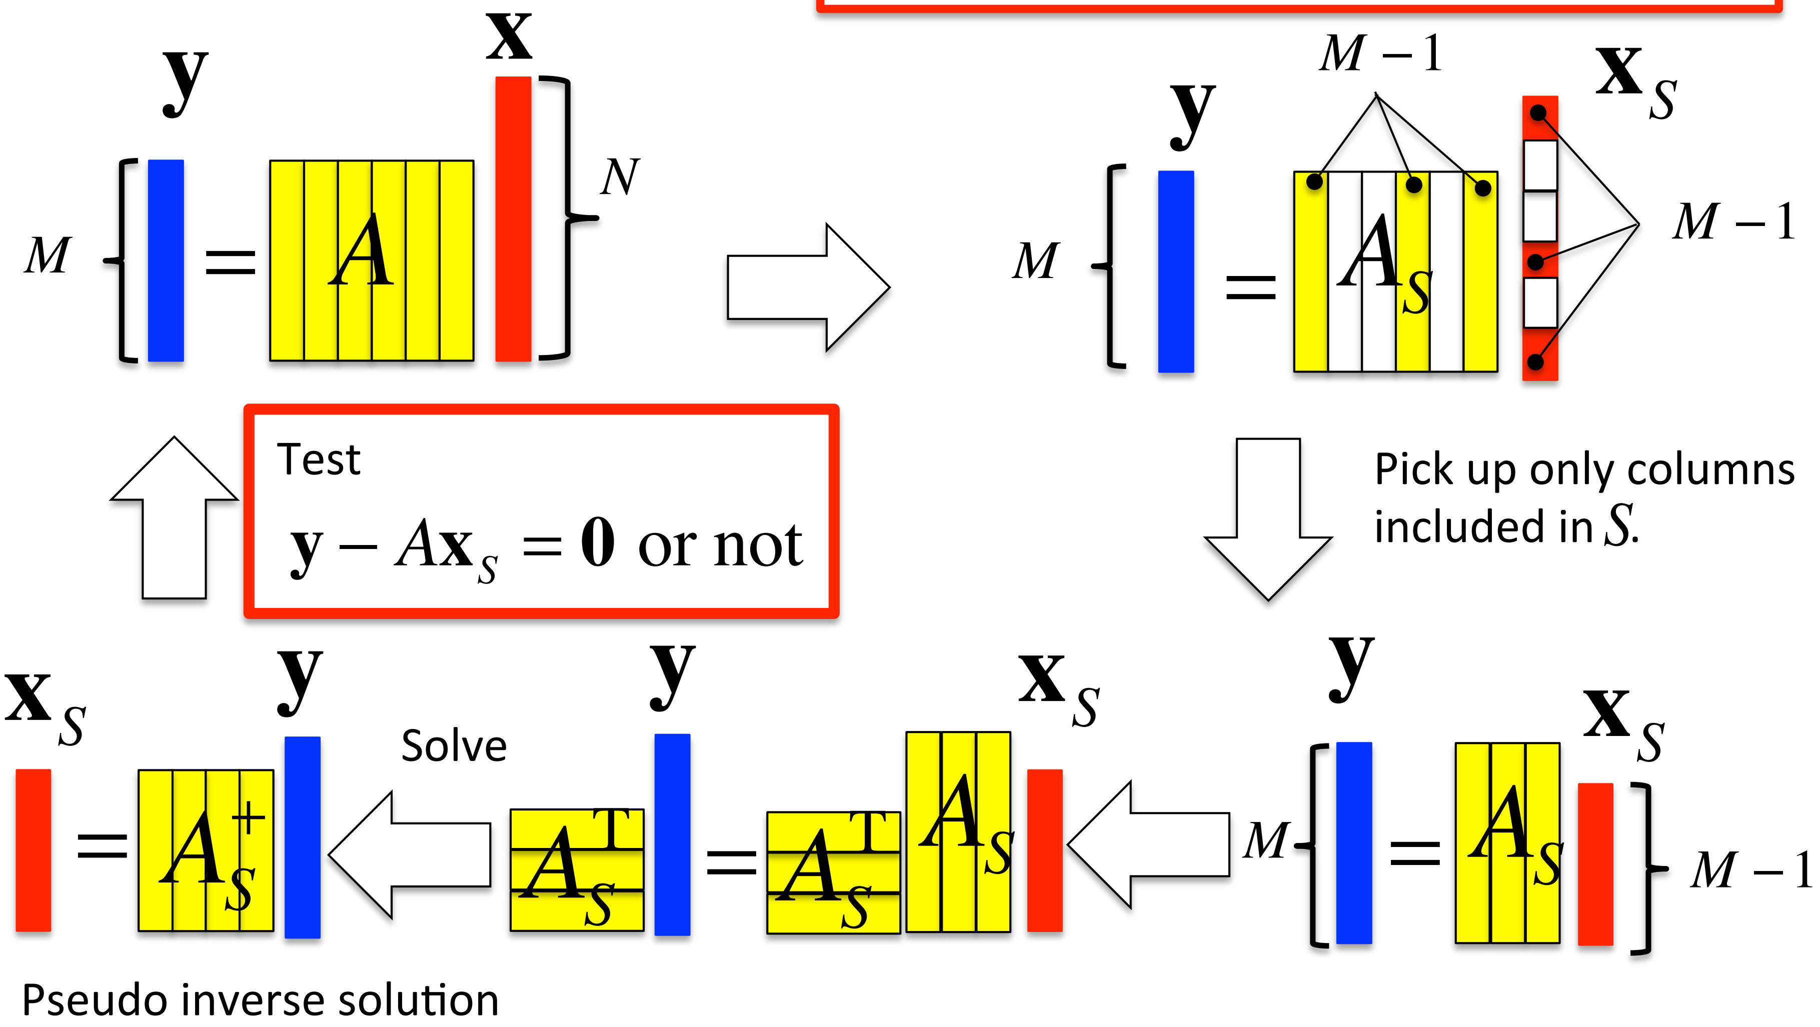

Consideration

- $K \leq M-1 \rightarrow$ Correct solution is provided by certain combinations.
  - Correct solution passes the test.
  - Otherwise, fails.

  } *Correct solution is the one that passes the test.*

⇨ $M > K \rightarrow$ l_0–recovery can always find the solution.

Many ($_{N-K}C_{M-1-K}$) combinations correspond to the solution.

- Unfortunately, this recovery becomes computationally difficult as the number of non-zeros increases.
  - Computational cost for l_0-recovery

$$O\left({}_{N}C_{M-1}\right) \sim O\left(\exp\left(\gamma N\right)\right) \ \text{if } M \propto N$$

Practical solutions

- Greedy algorithms
  - Greedy construction of "support"(=column combination) by adding one-by-one/best choice at each iteration.
  - **Orthogonal Matching Pursuit (OMP)**, MP, Weak MP, LS-OMP, **Iterative Hard Thresholding**, …
- Convex relaxation
  - Approximation of the cost by convex functions.
  - **Basis Pursuit (BP)(=l_1-recovery), Iterated-Reweighted-Least-Squares (IRLS)**, …
- Probabilistic inference
  - (Approximate) Employment of probabilistic inference.
  - **Approximate message passing (AMP), EM-BP**, …

Greedy algorithms

- Basic idea
  - Sample vector $\mathbf{y}$ stands for a linear combination of columns $\mathbf{a}_i$ of A.

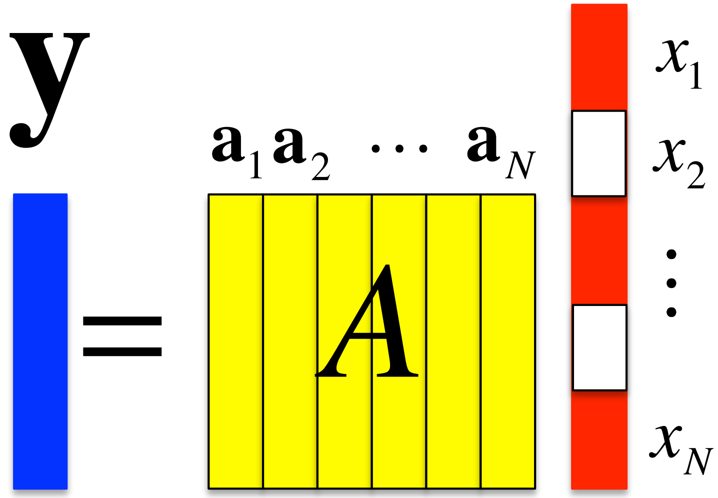

$$= x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \ldots + x_N\mathbf{a}_N$$

  - Construct an appropriate set of columns whose coefficients are non-zero, which is termed *support*, in a *greedy* manner.

Approximation by a single column

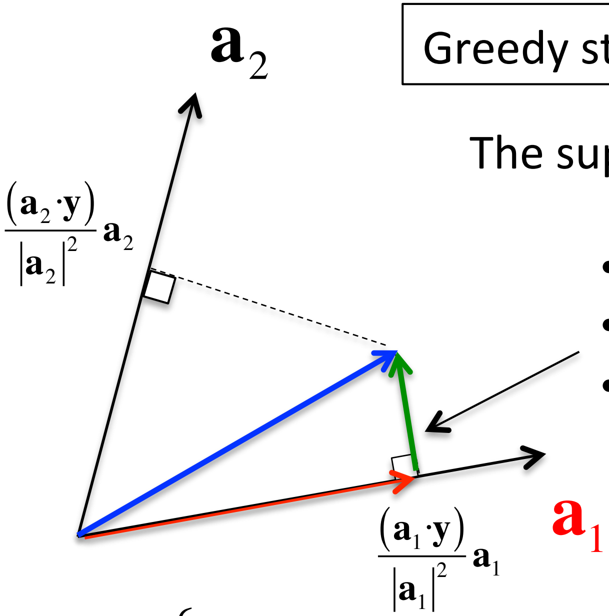

Greedy strategy

The support size is increased one-by-one.

- Choose the column that has the largest projection.
- Remove the approximated part.
- Iteratively employ the same procedure to the *residual.*

$$
\begin{cases}
\hat{x}_i = \arg\min_{x_i} |\mathbf{y} - x_i \mathbf{a}_i|^2 = \dfrac{(\mathbf{a}_i \cdot \mathbf{y})}{|\mathbf{a}_i|^2} \\
\epsilon(i) = \min_{x_i} |\mathbf{y} - x_i \mathbf{a}_i|^2 = |\mathbf{y}|^2 - \dfrac{(\mathbf{a}_i \cdot \mathbf{y})^2}{|\mathbf{a}_i|^2}
\end{cases}
$$

Approx. = projection

Quality of approximation

Orthogonal Matching Pursuit (OMP)

- **Initialization:** Initialize $k = 0$, and set

$$\mathbf{x}^0 = \mathbf{0},\ \mathbf{r}^0 = \mathbf{y} - A\mathbf{x}^0 = \mathbf{y},\ S^0 = \phi$$

- **Main iteration**: Increment k by 1 and perform the followings:
  - **Sweep:**

$$\epsilon(i) = \min_{x_i} \left| x_i \mathbf{a}_i - \mathbf{r}^{k-1} \right|^2 = \left| \mathbf{r}^{k-1} \right|^2 - \frac{\left(\mathbf{a}_i \cdot \mathbf{r}^{k-1} \right)^2}{\left| \mathbf{a}_i \right|^2}$$

Rating of columns

  - **Update Support:**

$$i_0 = \arg\min_{i \notin S^{k-1}} \left\{ \epsilon\left(i\right) \right\},\ S^k = S^{k-1} \cup \left\{ i_0 \right\}$$

  - **Update Provisional Solution:**

$$\hat{\mathbf{x}}^k = \arg\min_{\mathbf{x}_{S^k}} \left| \mathbf{y} - A_{S^k} \mathbf{x}_{S^k} \right|^2$$

Best approximation by the support of the moment

  - **Update Residual:**

$$\mathbf{r}^k = \mathbf{y} - A_{S^k} \hat{\mathbf{x}}^k$$

  - **Stopping Rule:** Stop if $\left| \mathbf{r}^k \right| < \epsilon_0$ holds. Otherwise, apply another iteration.

Consideration

- Computational cost

$$k_0 : \text{ final support size}$$

OMP $\quad\quad$ Exact enumeration

$$O(MNk_0) \longleftarrow O(MN^{k_0} k_0^2)$$

Drastic reduction

(when $k_0 \sim O(1)$)

Varieties

- **LS-OMP:** In **Sweep**, approximation error $\epsilon(i)$ is evaluated for $S^{k-1} \cup \{i\}$.
  - More accurate evaluation although cost increases.

- **MP:** In **Update Provisional Solution,** $\hat{\mathbf{x}}^k = \hat{\mathbf{x}}^{k-1} + \frac{\left(\mathbf{a}_{i_0} \cdot \mathbf{r}^{k-1}\right)}{\left|\mathbf{a}_{i_0}\right|^2} \mathbf{a}_{i_0}$
  - Get lazy in approximating solution for saving cost.

- **Weak-MP:** In **Sweep**, get lazy in optimizing $\epsilon(i)$

  Stop the sweep when $\frac{\left(\mathbf{a}_i \cdot \mathbf{r}^{k-1}\right)^2}{\left|\mathbf{a}_i\right|^2} \geq t \left|\mathbf{r}^{k-1}\right|^2 \quad \left(t \in [0,1]\right)$
  - Stop the searching column if a certain condition is satisfied.

Thresholding Algorithm

- Given a support size k. Support is fixed by the result of <u>the first quality evaluation.</u>
  - **Quality Evaluation:**

$$\epsilon(i) = \min_{x_i} \left| x_i \mathbf{a}_i - \mathbf{y} \right|^2 = \left| \mathbf{y} \right|^2 - \frac{\left(\mathbf{a}_i \cdot \mathbf{y} \right)^2}{\left| \mathbf{a}_i \right|^2}$$

  - **Update Support:**

$$S: \text{ set of indices of } k \text{ lowest values of } \epsilon(i)$$

  - **Update Provisional Solution:**

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}_S} \left| A_S \mathbf{x}_S - \mathbf{y} \right|^2$$

  - **Output:** Output $\hat{\mathbf{x}}$

Performance comparison

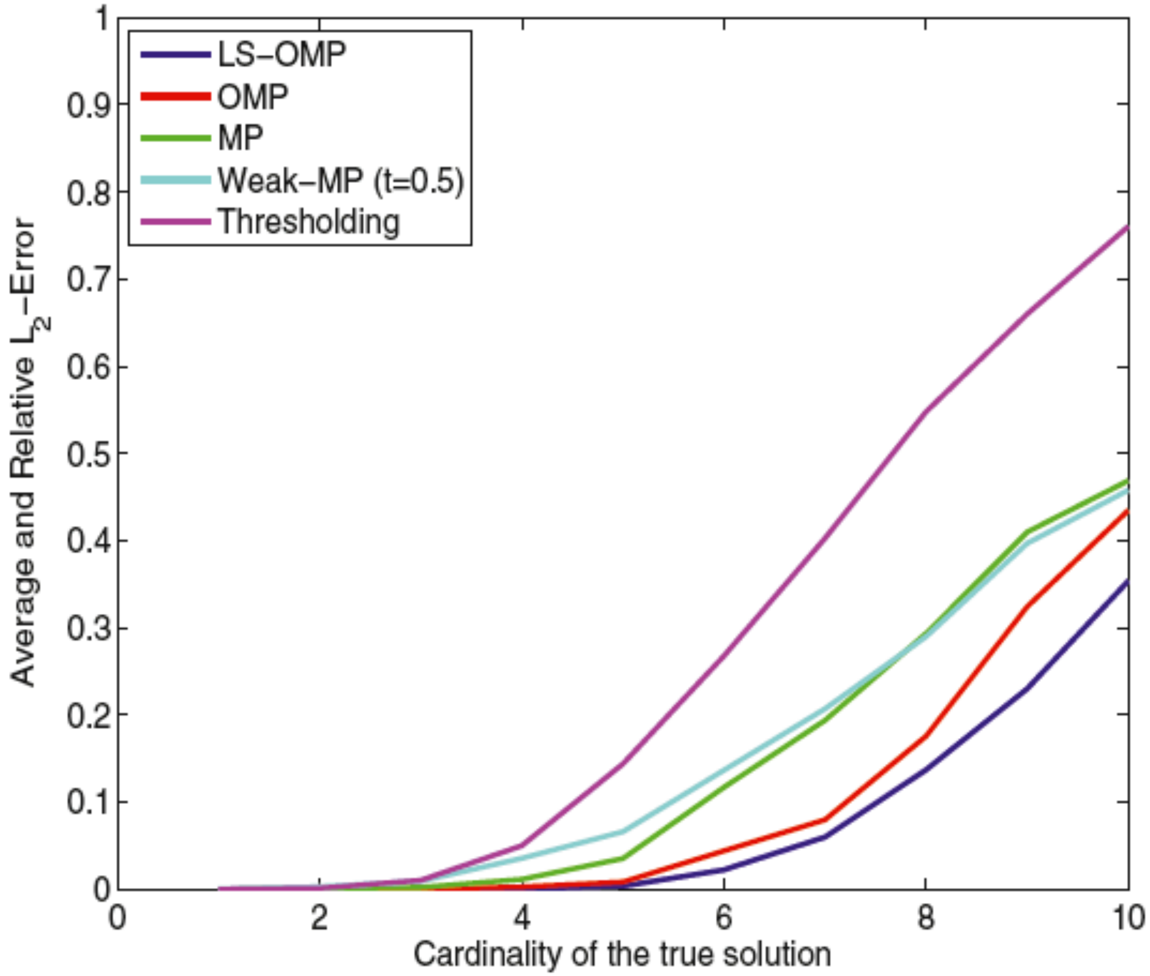

l_2-Error with true sol.

$M = 30$
$N = 50$

$A:$ i.i.d. Gaussian

$x_i \in [-2,-1] \cup [1,2]$
if $x_i \neq 0$

M. Elad (2010) Sparse and Redundant Representations (NY: Springer)

Performance comparison

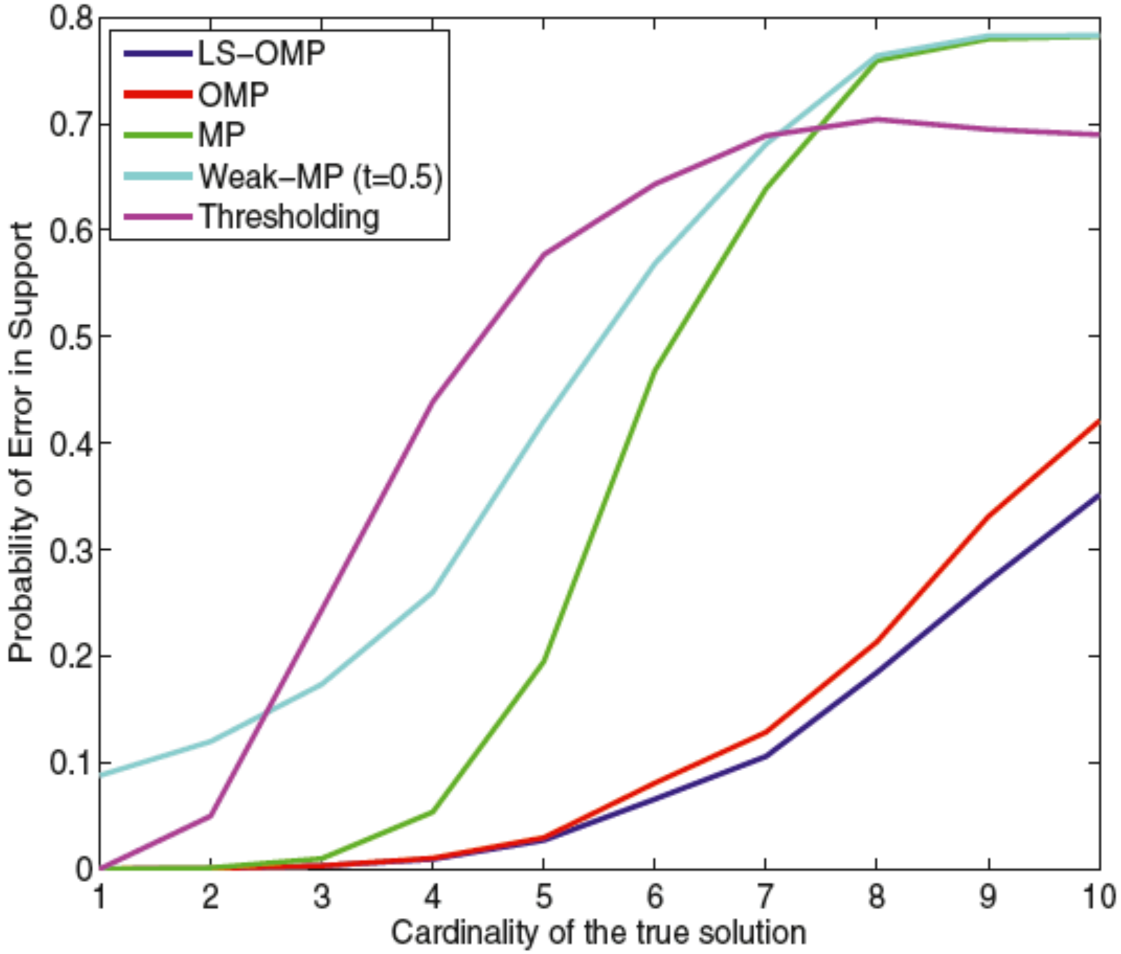

Support error with true sol.

$$dist(\hat{S},S) = \frac{\max\left\{|\hat{S}|,|S|\right\} - |\hat{S}\cap S|}{\max\left\{|\hat{S}|,|S|\right\}}$$

$$M = 30$$
$$N = 50$$

$$A: \text{ i.i.d. Gaussian}$$

$$x_i \in [-2,-1]\cup[1,2] \text{ if } x_i \neq 0$$

M. Elad (2010) Sparse and Redundant Representations (NY: Springer)

Iterative Hard Thresholding

- Given a support size k. Iterate the following until convergence.

$$\mathbf{x}^{t+1} = H_k\left(\mathbf{x}^t + A^{\mathrm{T}}\left(\mathbf{y} - A\mathbf{x}^t\right)\right)$$

$H_k(\cdot)$: Set all but k largest (in amplitude) components to zero

Comput. cost: $O(MN) / iteration$

- Guarantee for convergence to <u>local minimum of l_0-cost</u>
  - Blumensath and Davies (2009)

$$\|A\|_2 = \sqrt{\sum_{\mu,i} |A_{\mu i}|^2} < 1$$

$\mathbf{x}^t$ converges to a local minimum of

$$\mathbf{x}^* = \min_{\mathbf{x}:\|\mathbf{x}\|_0 \le k} \|\mathbf{y} - A\mathbf{x}\|_2$$

Performance comparison

- A sufficient condition of correct recovery is obtained by RIP theory that is mentioned later.
- But, experimentally observed performance of the naïve IHT is not so good as that for Basis Pursuit (l_1-recovery).

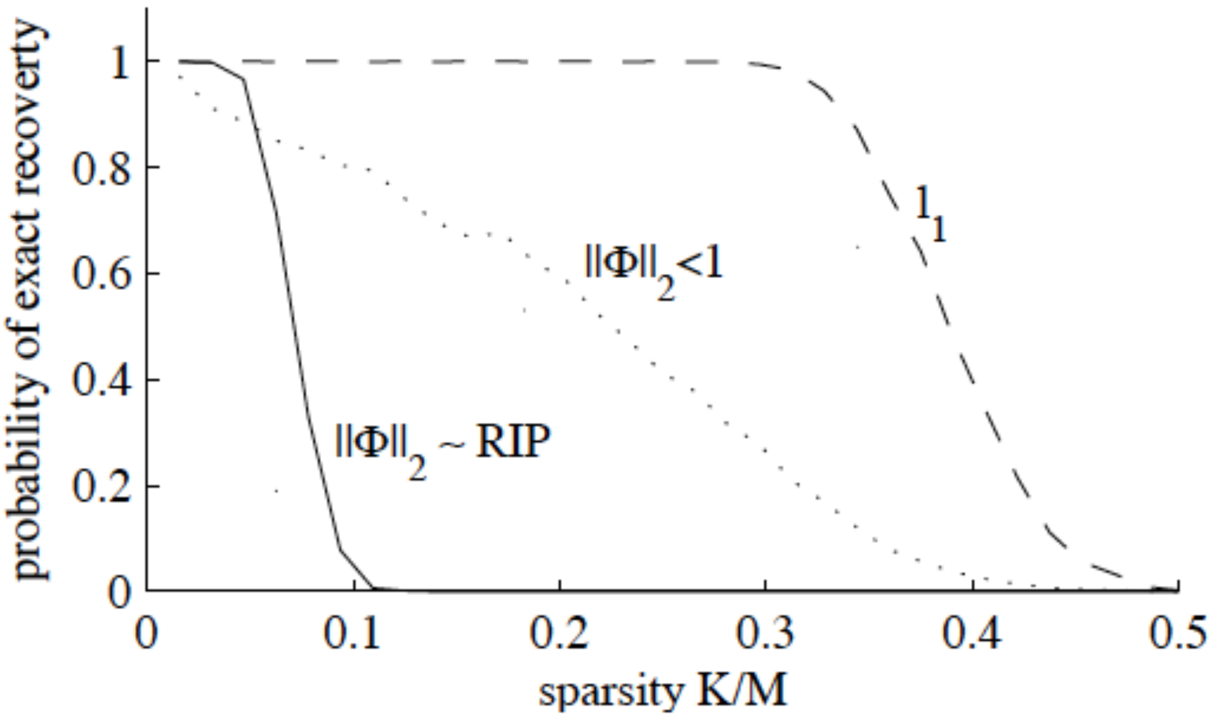

Normalized Iterative Hard Thresholding

- Given a support size k. Iterate the following until convergence.

$$\mathbf{x}^{t+1} = H_k\left(\mathbf{x}^t + \mu^t A^{\mathrm{T}}\left(\mathbf{y} - A\mathbf{x}^t\right)\right) \qquad \mu^t : \text{ step size}$$

$H_k(\cdot)$: Set all but k largest (in amplitude) components to zero

Comput. cost: $O(MN)/iteration$

- Step size μ^t is controlled using the information of deviation

$$\mathbf{y} - A\mathbf{x}^t$$

and A, appropriately.
(Blumensath and Davies (2010))

Performance comparison

- A sufficient condition of correct recovery is obtained by RIP theory for the *normalized* version as well.
- Experimentally observed performance is as good as l_1-recovery.

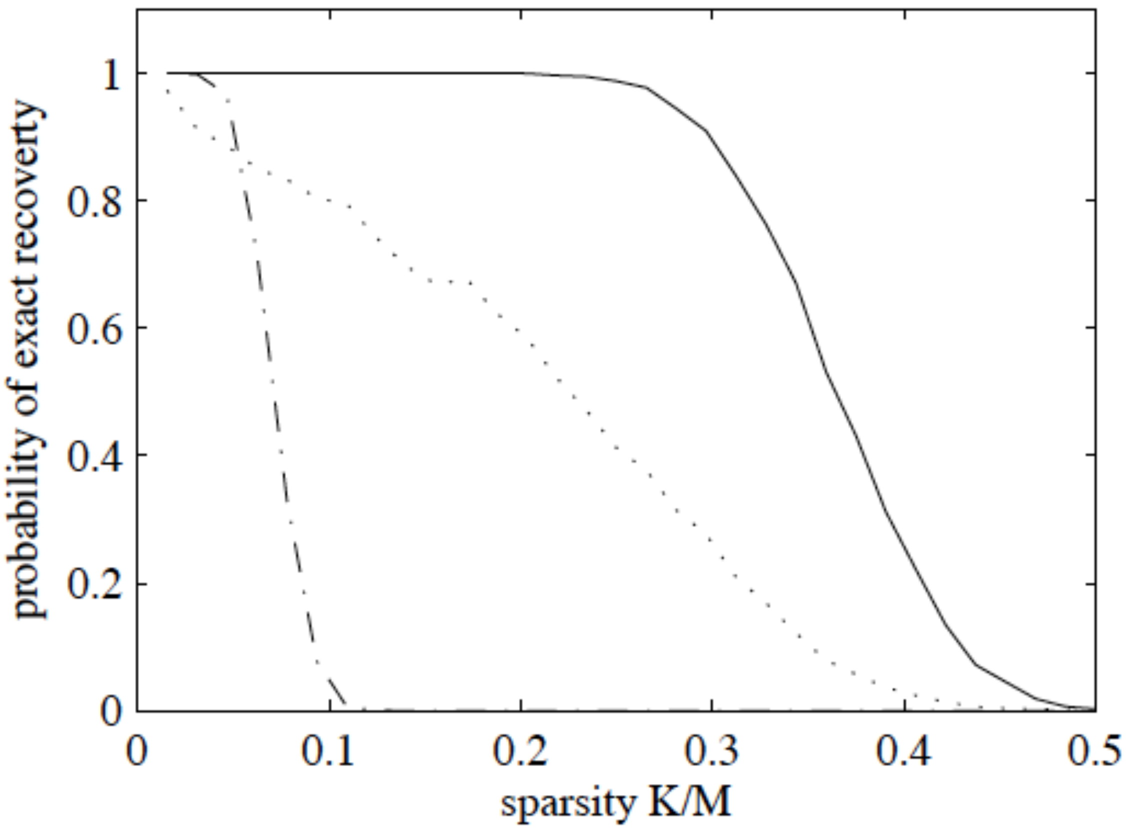

Advantageous as the comput. cost is lower than l_1-recovery.

Convex relaxation

- Basic idea
  - l_0-recovery is difficult since it is formulated as discrete optimization problem → How about *approximating* the cost by continuous function?

l_0-norm

$$\|\mathbf{x}\|_0 = \lim_{p \to +0} \sum_{i=1}^{N} |x_i|^p$$

l_p-norm

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

$|x|^p$ profiles

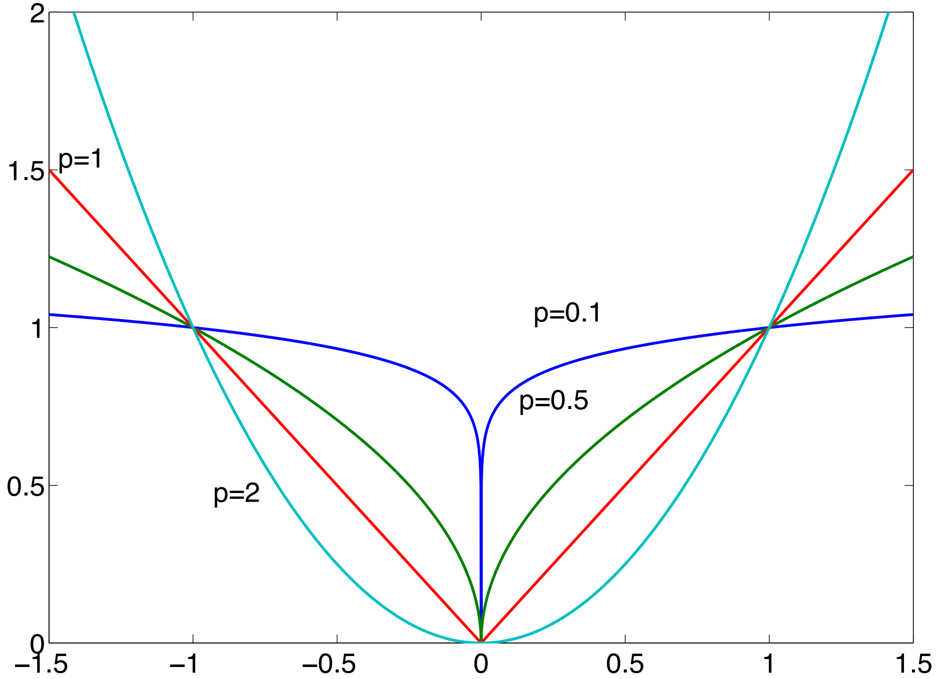

Convex relaxation

- Among them, convex functions are preferred. $\rightarrow p \geq 1$
  - Optimization is easy
    - Uniqueness guarantee of solution
    - Various versatile packages
- Necessity for producing sparse solutions
  $\rightarrow p \leq 1$
- Promising candidate

l_1-norm

$$\|\mathbf{x}\|_1 = \sum_{i=1}^{N} |x_i|$$

$|x|^p$ profiles

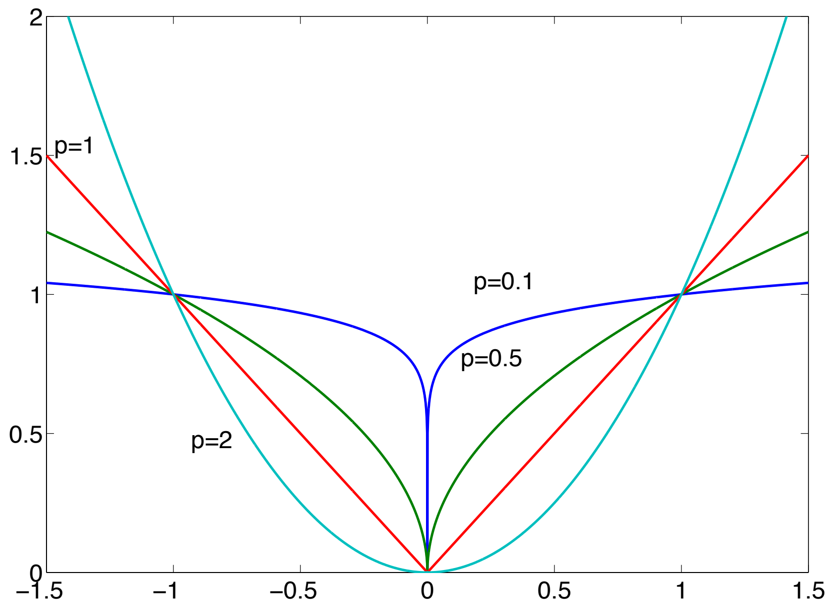

Norm and solution

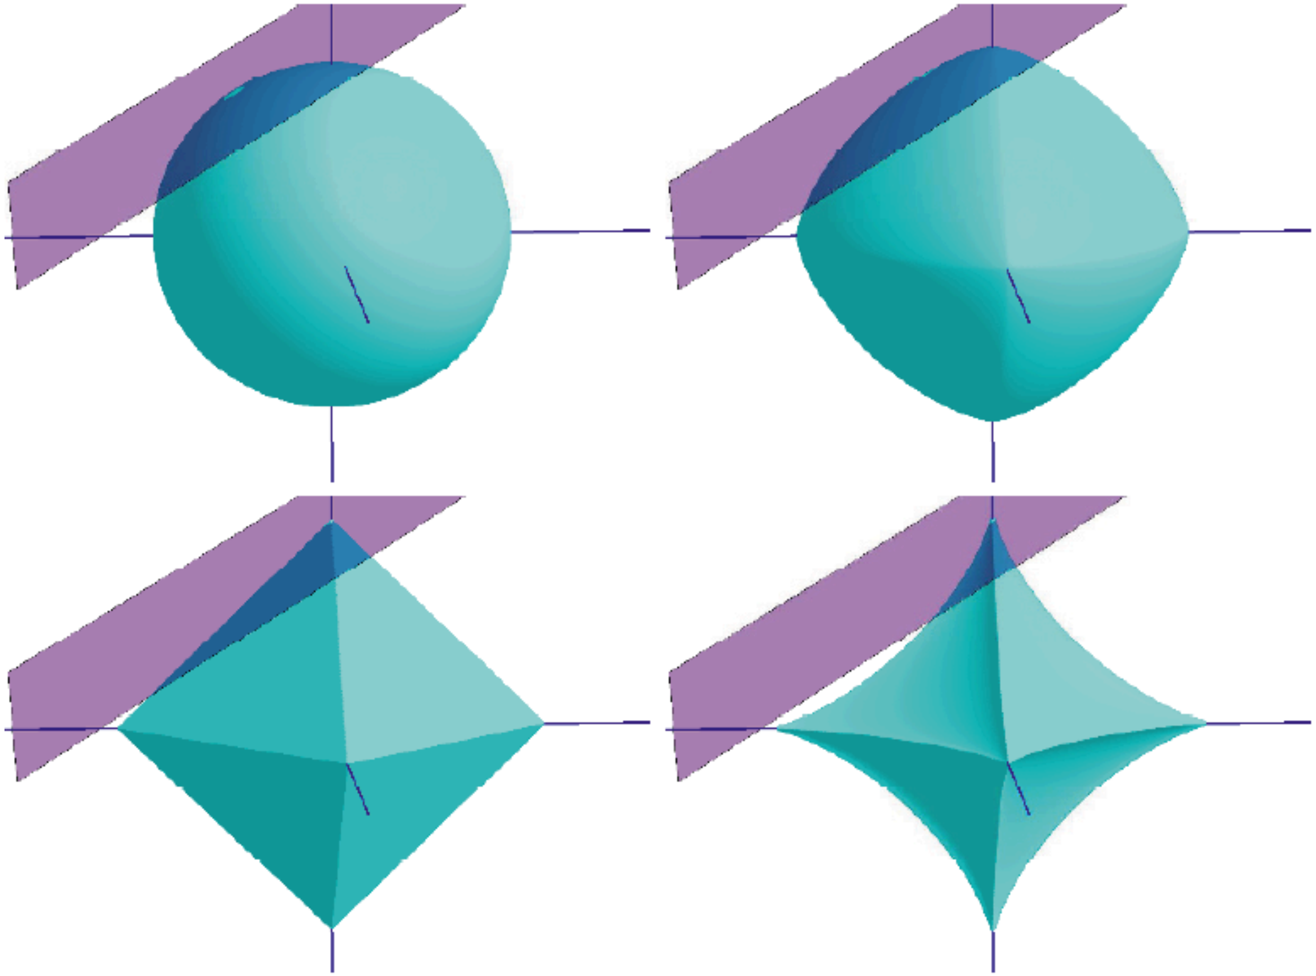

Purple plane:

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

Green figures:

$$\|\mathbf{x}\|_p = const$$

(l_p-ball)

LRT: $p > 1$

LB: $p = 1$

RB: $p < 1$

M. Elad (2010) Sparse and Redundant Representations (NY: Springer)

l_1-recovery = Basis Pursuit (BP)

$$\left(P_0\right): \quad \min_{\mathbf{x}} \|\mathbf{x}\|_0 \ \text{ subj. to } \mathbf{y} = A\mathbf{x}$$

$$\left(P_1\right): \quad \min_{\mathbf{x}} \|\mathbf{x}\|_1 \ \text{ subj. to } \mathbf{y} = A\mathbf{x}$$

Conversion to linear programing

$$\mathbf{x} = \mathbf{u} - \mathbf{v} \begin{cases} \mathbf{u} : \text{In } \mathbf{x}, \\ \mathbf{v} : \text{In } \mathbf{x}, \end{cases}$$

$\mathbf{u}$: In $\mathbf{x}$, Keep positive entities as they are, Set negative entities to zeros

$\mathbf{v}$: In $\mathbf{x}$, Set positive entities to zeros Set negative entities to their amplitudes

$$\mathbf{z} = \left[\mathbf{u}^{\mathrm{T}}, \mathbf{v}^{\mathrm{T}}\right]^{\mathrm{T}} \in \mathbb{R}^{2N}$$

⇩

$$\|\mathbf{x}\|_1 = \mathbf{1}^{\mathrm{T}}(\mathbf{u} + \mathbf{v}),\ A\mathbf{x} = A(\mathbf{u} - \mathbf{v}) = [A, -A]\mathbf{z}$$

⇩

$$(P_1) \Leftrightarrow \underbrace{\min_{\mathbf{z}} \ \mathbf{1}^{\mathrm{T}}\mathbf{z} \text{ subj. to } \mathbf{y} = [A, -A]\mathbf{z} \text{ and } \mathbf{z} \geq \mathbf{0}}_{\text{Linear Programming}}$$

Linear Programming

How to carry out l_1-recovery

- Utilize various algorithms and packages for convex optimization.
  - Algorithms
    - Simplex method, interior point method, homotopy method, etc.
  - Packages
    - LAPACK, GPLK, l1-magic, CVX, L1-LS, Sparselab, etc.
- Computational cost is guaranteed to be $O(N^3)$.
  - Proof for interior point method
  - Simplex method is also efficient empirically for many problems.
- How is the solution related to the l_0-solution?
  - Sufficient conditions are known for the coincidence of the l_1- and l_0-solutions.

l_p-recovery

- Relaxation to *l_p-norm* $(0 < p < 1)$ often leads to good performance although it is not formulated as convex optimization.

$$\left(P_p\right): \ \min_{\mathbf{x}} \sum_{i=1}^{N} \left|x_i\right|^p \ \text{subj. to } \mathbf{y} = A\mathbf{x}$$

- But, performing the optimization is a nontrivial task.

Core idea

- Exceptionally, one can *analytically* solve the optimization of *quadratic functions.*

$$B = \mathrm{diag}(b_i),\ \left(B^+\right)_i = \begin{cases} b_i^{-1}, & (b_i \neq 0) \\ 0, & (b_i = 0) \end{cases}$$

$$(M):\ \min_{\mathbf{x}} \sum_{i=1}^{N} \left(B^+\right)_i |x_i|^2 \text{ subj. to } \mathbf{y} = A\mathbf{x}$$

Weighted quadratic norm

$$\mathbf{x} = BA^{\mathrm{T}}(ABA^{\mathrm{T}})^{+}\mathbf{y}$$

- Employ this formula for *iterative optimization* of l_p-cost function.

Iterated-Reweighted-Least Squares (IRLS)

- $\mathbf{x}^k$: Solution obtained at the k th iteration

$$X_k = \mathrm{diag}\left(\left|x_i^k\right|^{2-p}\right), \quad \left(X_k^+\right)_i = \begin{cases} \left|x_i^k\right|^{p-2} & \left(x_i^k \neq 0\right) \\ 0 & \left(x_i^k = 0\right) \end{cases}$$

Iterate the following until convergence

$$\left(M_k\right): \ \min_{\mathbf{x}} \sum_{i=1}^{N} \left(X_{k-1}^+\right)_i \left|x_i\right|^2 \ \text{subj. to } \mathbf{y} = A\mathbf{x}$$

$$\mathbf{x}^k = X_{k-1} A^{\mathrm{T}} (A X_{k-1} A^{\mathrm{T}})^+ \mathbf{y}$$

Comput. cost: *O(N³)*/iter

⇒ Converges to a local minimum of $\left(P_p\right)$

Performance comparison

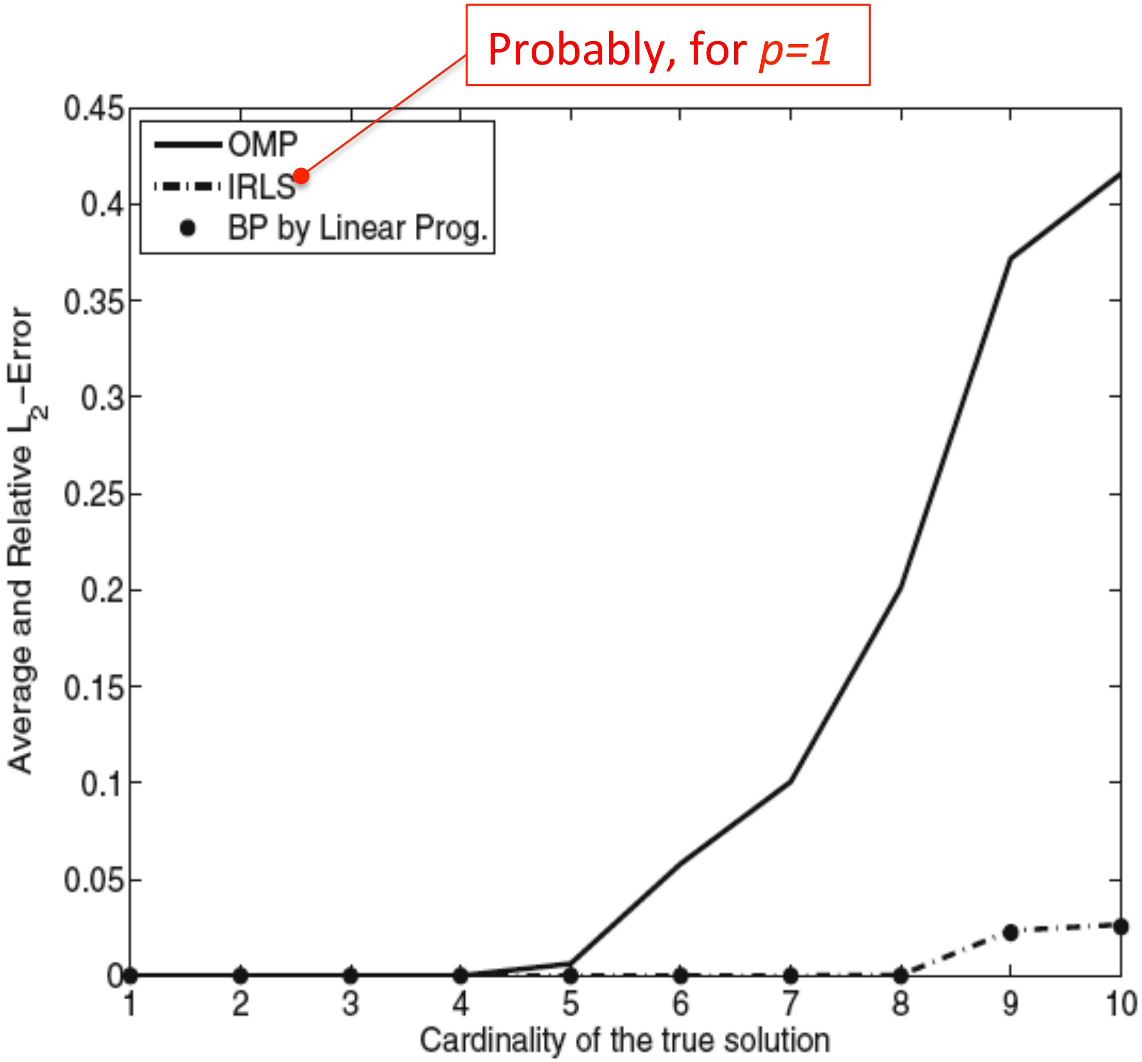

l_2-Error with true sol.

$M = 30$
$N = 50$

$A:$ i.i.d. Gaussian

$x_i \in [-2,-1] \cup [1,2]$
if $x_i \neq 0$

Comput. cost
BP, IRLS >> OMP

M. Elad (2010) Sparse and Redundant Representations (NY: Springer)

Performance comparison

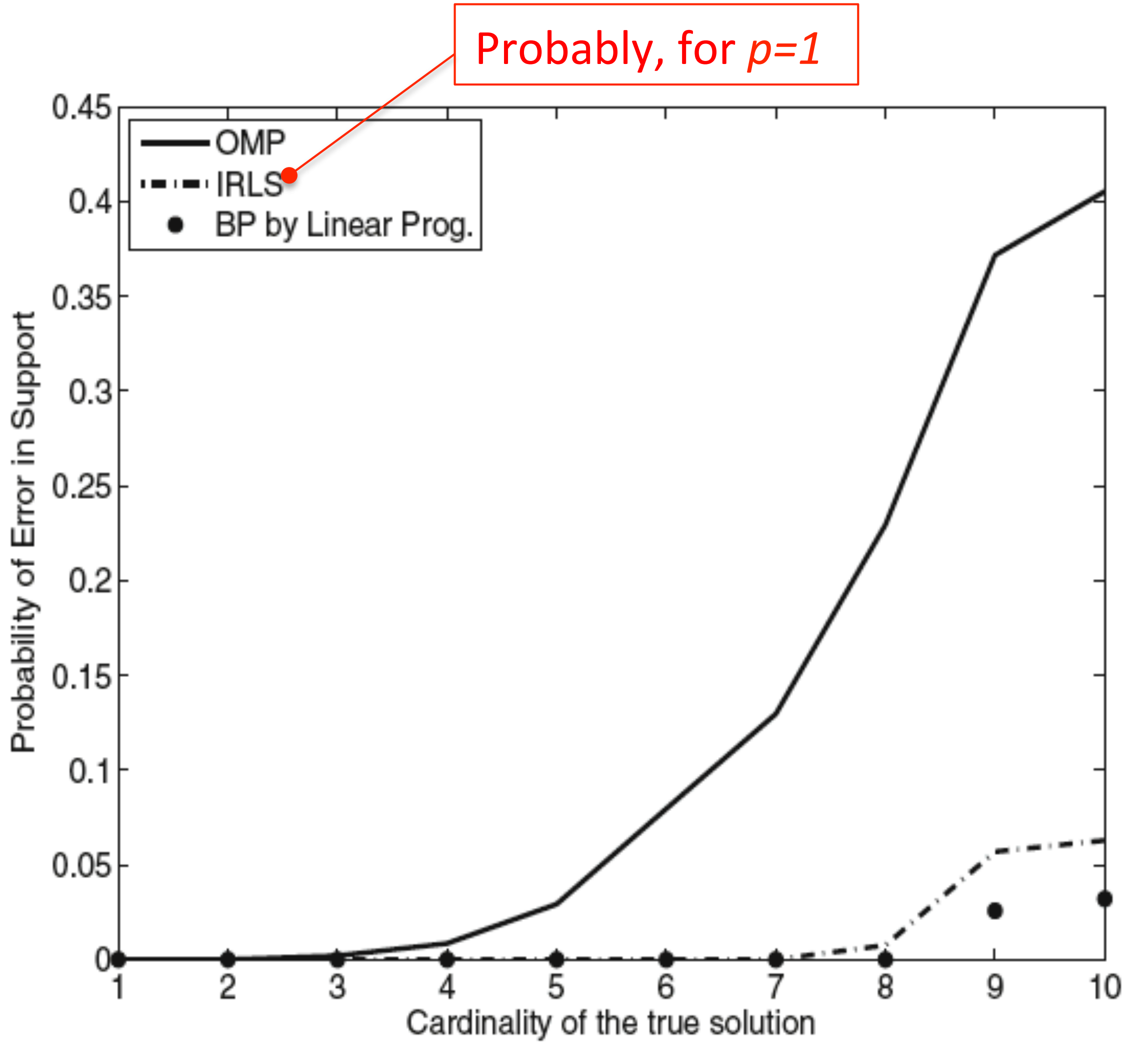

Support error with true sol.

$$dist(\hat{S}, S) = \frac{\max\left\{|\hat{S}|, |S|\right\} - |\hat{S} \cap S|}{\max\left\{|\hat{S}|, |S|\right\}}$$

$$M = 30$$
$$N = 50$$

$$A: \text{ i.i.d. Gaussian}$$

$$x_i \in [-2,-1] \cup [1,2]$$
$$\text{if } x_i \neq 0$$

Comput. cost
BP, IRLS >> OMP

M. Elad (2010) Sparse and Redundant Representations (NY: Springer)

Probabilistic inference

- Basic idea
  - Signal recovery = Probabilistic inference from measurements
  - Knowledge of *generative model* may lead to better performance

Bayes formula

$$\underbrace{P(\mathbf{x} \mid \mathbf{y}, A)}_{\text{Posterior}} = \frac{1}{Z}\delta(\mathbf{y} - A\mathbf{x})\underbrace{P(\mathbf{x})}_{\text{Sparse prior}}$$

  - Unfortunately, extracting information from large scale distributions is computationally difficult in general.
  - Utilize various approximate inference algorithm studied in machine learning/artificial intelligence.

l_p-recovery as probabilistic inference

Prior

$$P(\mathbf{x}) \propto \exp\left(-\beta \|x\|_p^p\right)$$

Noisy measurement model

$$P\left(\mathbf{y} \mid \mathbf{x}, A, \sigma^2\right) = \frac{1}{\left(\sqrt{2\pi\sigma^2}\right)^M} \exp\left(-\frac{1}{2\sigma^2}|\mathbf{y} - A\mathbf{x}|^2\right)$$

Posterior

$$P\left(\mathbf{x} \mid \mathbf{y}, A, \sigma^2\right) \propto \exp\left(-\frac{1}{2\sigma^2}|\mathbf{y} - A\mathbf{x}|^2 - \beta\|\mathbf{x}\|_p^p\right)$$

$$= \exp\left(-\frac{1}{\sigma^2}\left(\underline{\frac{1}{2}|\mathbf{y} - A\mathbf{x}|^2 + \beta\sigma^2\|\mathbf{x}\|_p^p}\right)\right)$$

Maximum a posteriori (MAP) inference $=l_p$-recovery

$$\left(Q_p\right): \ \min_{\mathbf{x}}\left\{\frac{1}{2}|\mathbf{y} - A\mathbf{x}|^2 + \beta\sigma^2\|\mathbf{x}\|_p^p\right\} \xrightarrow{\sigma^2 \to 0} \left(P_p\right): \ \min_{\mathbf{x}}\|\mathbf{x}\|_p^p \ \text{subj. to } \mathbf{y} = A\mathbf{x}$$

How to solve *p=1* case

Subderivative

$$\frac{\partial}{\partial \mathbf{x}}\left\{\frac{1}{2}|\mathbf{y}-A\mathbf{x}|^2+\beta\sigma^2\|\mathbf{x}\|_1\right\}=\mathbf{0} \rightarrow A^{\mathrm{T}}\underbrace{(A\mathbf{x}-\mathbf{y})}_{-\mathbf{z}}+\underbrace{\beta\sigma^2}_{k}\,\mathrm{sgn}(\mathbf{x})=\mathbf{0}$$

$$\mathbf{x}-\underbrace{\left(A^{\mathrm{T}}\mathbf{z}+\mathbf{x}\right)}_{h}=-k\,\mathrm{sgn}(\mathbf{x})$$

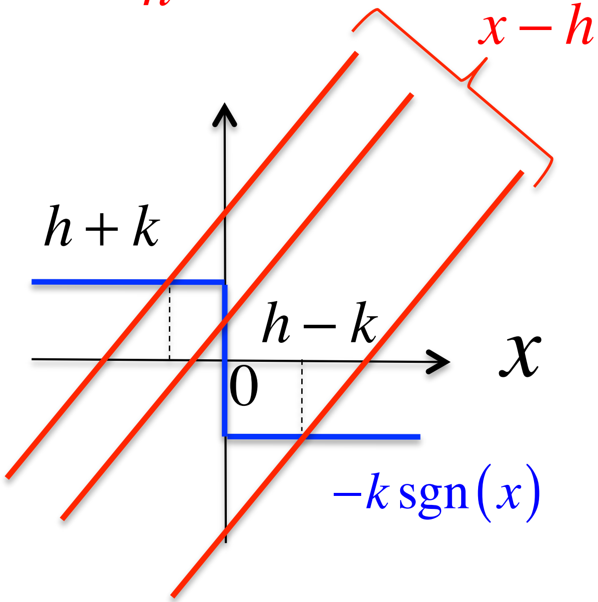

$$\begin{cases}\mathbf{x}=\eta_k\left(A^{\mathrm{T}}\mathbf{z}+\mathbf{x}\right)\\ \mathbf{z}=\mathbf{y}-A\mathbf{x}\end{cases}$$

$$\eta_k(h)=\left(h-k\,\mathrm{sgn}(h)\right)\theta\left(|h|-k\right)$$

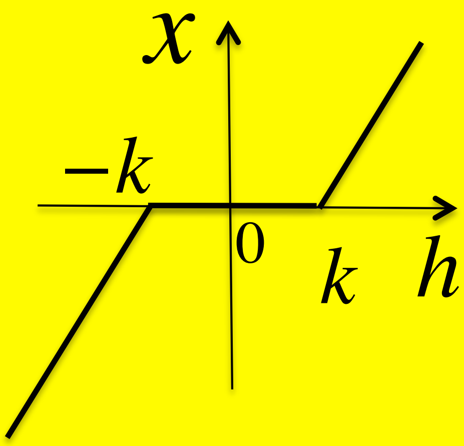

Naïve iteration does not work well

Each column of A is assumed to be normalized to unity.

$$\begin{cases} \mathbf{x} = \eta_k\left(A^{\mathrm{T}}\mathbf{z} + \mathbf{x}\right) \\ \mathbf{z} = \mathbf{y} - A\mathbf{x} \end{cases} \quad \Rightarrow \quad \begin{cases} \mathbf{x}^{t+1} = \eta_k\left(A^{\mathrm{T}}\mathbf{z}^t + \mathbf{x}^t\right) \\ \mathbf{z}^t = \mathbf{y} - A\mathbf{x}^t \end{cases}$$

iteration

Doesn't work well

Reason

Self-feedback between *t* and *t+1* is cancelled. But, that for *t-1* and *t+1* is too large.

⇧

Knowledge of prob. inference

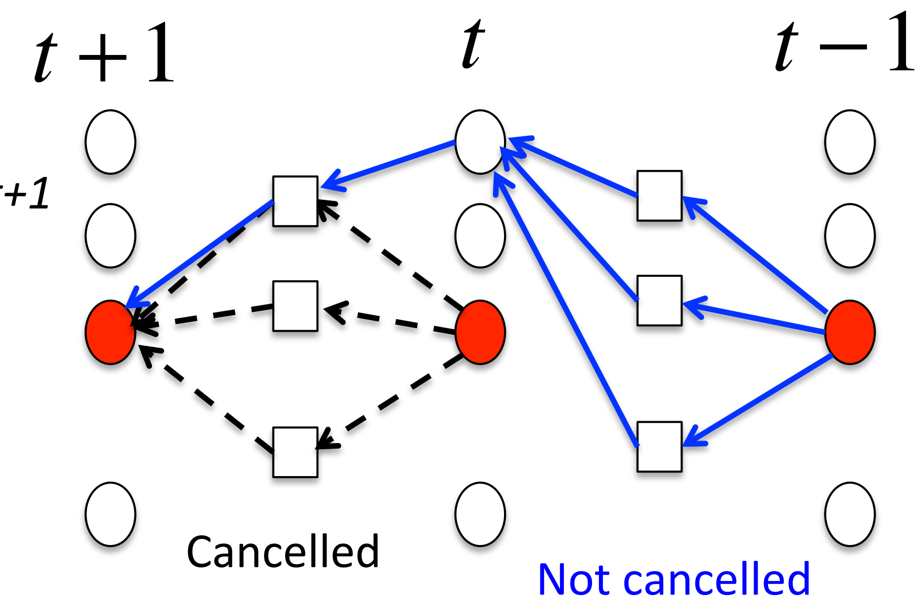

Approximate Message Passing (AMP)

- In *t+1*st update, remove self-feedback effect from *t-1st* update.
  - Analytical evaluation is possible if entries of A independently obey a distribution of zero mean and variance $O(M^{-1})$.

$$
\begin{cases}
\mathbf{x}^{t+1} = \eta_k\left(A^{\mathrm{T}}\mathbf{z}^t + \mathbf{x}^t\right) \\
\mathbf{z}^t = \mathbf{y} - A\mathbf{x}^t + \underline{\frac{1}{\delta}\mathbf{z}^{t-1}\left\langle \eta_k^{t-1}\left(A^{\mathrm{T}}\mathbf{z}^{t-1} + \mathbf{x}^{t-1}\right)\right\rangle}
\end{cases}
$$

Where $\delta = \frac{M}{N}, \ \langle \mathbf{u} \rangle \equiv \frac{1}{N}\sum_{i=1}^{N} u_i$

- Appropriately reduce the thresholding parameter k as update number *t* grows.

Donoho, Maleki and Montanari, PNAS 109, 18914 (2009)

Consideration

- Approximation of l_1-recovery, which is computationally feasible originally. The gain is reduction of comput. cost.

$$O\left(N^3\right) \rightarrow O(N^2)$$

- The obtained approximate solution is guaranteed to typically converge to the exact solution as $N \longrightarrow \infty$.
  - Performance analysis by the State Evolution (SE).

Performance

Donoho, Makeli and Montanari, PNAS 109, 18914 (2009)

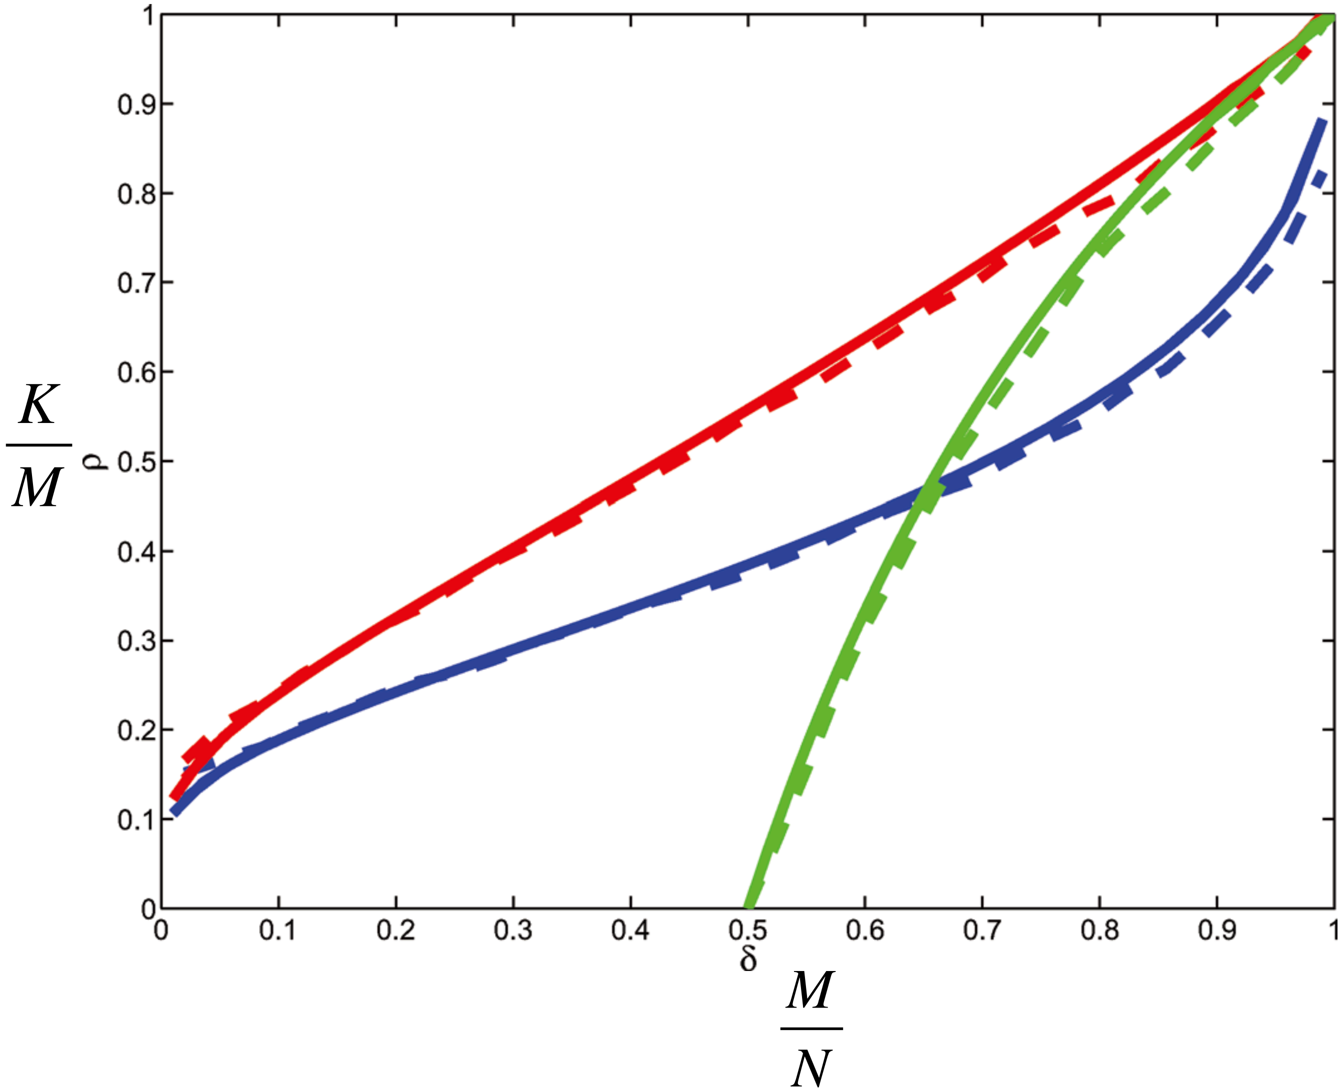

Full curves : Theoretical prediction

Broken curves : Experimental results of AMP for N=1000

For performance improvement

- If we need to construct an approximation, it is reasonable to do it for the ***best possible*** inference scheme.
- Best possible scheme in Bayesian framework
  ＝Inference using the true generative model

Generative model

$$\mathbf{x} \sim P(\mathbf{x})$$: Sparse prior

$$\mathbf{y} \sim P(\mathbf{y} \mid \mathbf{x}, A) = \delta(\mathbf{y} - A\mathbf{x})$$: Measurement process

Performance measure

Result by an arbitrary recovery scheme

$$\mathrm{MSE} = \left\langle \left\| \hat{\mathbf{x}}(\mathbf{y}, A) - \mathbf{x} \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}}$$: Mean Squared Error (MSE)

Bayes optimal scheme

- The following inequality holds.

$$\mathrm{MSE} \geq \left\langle \|\mathbf{x}\|^2 \right\rangle_{\mathbf{x},\mathbf{y}} - \left\langle \left\| \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \right\|^2 \right\rangle_{\mathbf{y}}$$

- The lower bound is achieved by the Bayes optimal recovery (MMSE recovery, in this case).

$$\hat{\mathbf{x}}^{\mathrm{opt}}(\mathbf{y}, A) = \sum_{\mathbf{x}} \mathbf{x} P(\mathbf{x} \mid \mathbf{y}, A) = \frac{\sum_{\mathbf{x}} \mathbf{x} P(\mathbf{y} \mid \mathbf{x}, A) P(\mathbf{x})}{\sum_{\mathbf{x}'} P(\mathbf{y} \mid \mathbf{x}', A) P(\mathbf{x}')}$$

$= \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}}$: conditional mean for given measurement $\mathbf{y}$.

Proof

Independent of true signal $\mathbf{x}$

$$\mathrm{MSE} = \left\langle \left\| \hat{\mathbf{x}}(\mathbf{y}, A) \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}} - 2 \left\langle \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \cdot \hat{\mathbf{x}}(\mathbf{y}, A) \right\rangle_{\mathbf{y}} + \left\langle \left\| \mathbf{x} \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}}$$

$$= \left\langle \left\| \hat{\mathbf{x}}(\mathbf{y}, A) \right\|^2 - 2 \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \cdot \hat{\mathbf{x}}(\mathbf{y}, A) \right\rangle_{\mathbf{y}} + \left\langle \left\| \mathbf{x} \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}}$$

$$= \left\langle \left\| \hat{\mathbf{x}}(\mathbf{y}, A) - \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \right\|^2 - \left\| \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \right\|^2 \right\rangle_{\mathbf{y}} + \left\langle \left\| \mathbf{x} \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}}$$

$$\geq \left\langle \left\| \mathbf{x} \right\|^2 \right\rangle_{\mathbf{x},\mathbf{y}} - \left\langle \left\| \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}} \right\|^2 \right\rangle_{\mathbf{y}}$$

Completing square

Minimizing the quadratic function w.r.t. $\hat{\mathbf{x}}(\mathbf{y}, A)$

$$\hat{\mathbf{x}}^{\mathrm{opt}}(\mathbf{y}, A) = \langle \mathbf{x} \rangle_{\mathbf{x}|\mathbf{y}}$$

Gaussian-Bernoulli prior

- We assume that the sparse prior can be well approximated by a mixture distribution of $\delta(x)$(=always returns "zero") and Gaussian.
  - Bernoulli-Gaussian prior

$$P(\mathbf{x},\rho,\mu,\sigma^2)=\prod_{i=1}^{N}\left((1-\rho)\delta(x_i)+\rho\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{(x_i-\mu)^2}{2\sigma^2}\right)\right)$$

- Noise free case

$$P(\mathbf{y}\mid\mathbf{x},A)=\delta(\mathbf{y}-A\mathbf{x})$$

Need for determining hyper-parameters

- Posterior distribution

$$P(\mathbf{x}\mid\mathbf{y},A,\rho,\sigma^2)=\frac{1}{Z}P(y\mid x,A)P(\mathbf{x},\rho,\mu,\sigma^2)$$

EM-BP

- Optimal recovery in the Bayesian framework

$$\hat{\mathbf{x}}^{\mathrm{opt}} = \langle \hat{\mathbf{x}} \rangle_{\mathbf{x}|\mathbf{y}} = \int d\mathbf{x} P\left(\mathbf{x} \mid \mathbf{y}, A, \rho, \mu, \sigma^2\right)\mathbf{x}$$

Carried out by a variant of AMP

- Hyper-parameter estimation for ρ, μ, σ^2

iteration

$$\left(\hat{\rho}, \hat{\mu}, \hat{\sigma}^2\right) = \underset{\rho, \sigma^2}{\arg\max} \left\{ \log Z\left(\rho, \mu, \sigma^2, A.\mathbf{y}\right) \right\}$$

$$= \underset{\rho, \mu, \sigma^2}{\arg\max} \left\{ \log\left(\int d\mathbf{x} P(\mathbf{y} \mid \mathbf{x}, A) P\left(\mathbf{x}, \rho, \mu, \sigma^2\right) \right) \right\}$$

Carried out by EM algorithm

Vila and Schniter (2011),
Krzakala et. al., Phys. Rev. X 2, 021005 (2012)

Consideration

- Comput. cost (similar to original AMP)

$$O\left(N^3\right) \to O(N^2)$$

- Better performance than l_1-recovery.
- By introducing *spatial-coupling* design for the measurement matrix, the theoretical recovery limit $\alpha_c(\rho) = \rho$ is achieved for $N \longrightarrow \infty$.
  - Talk by Marc Mezard.

Performance

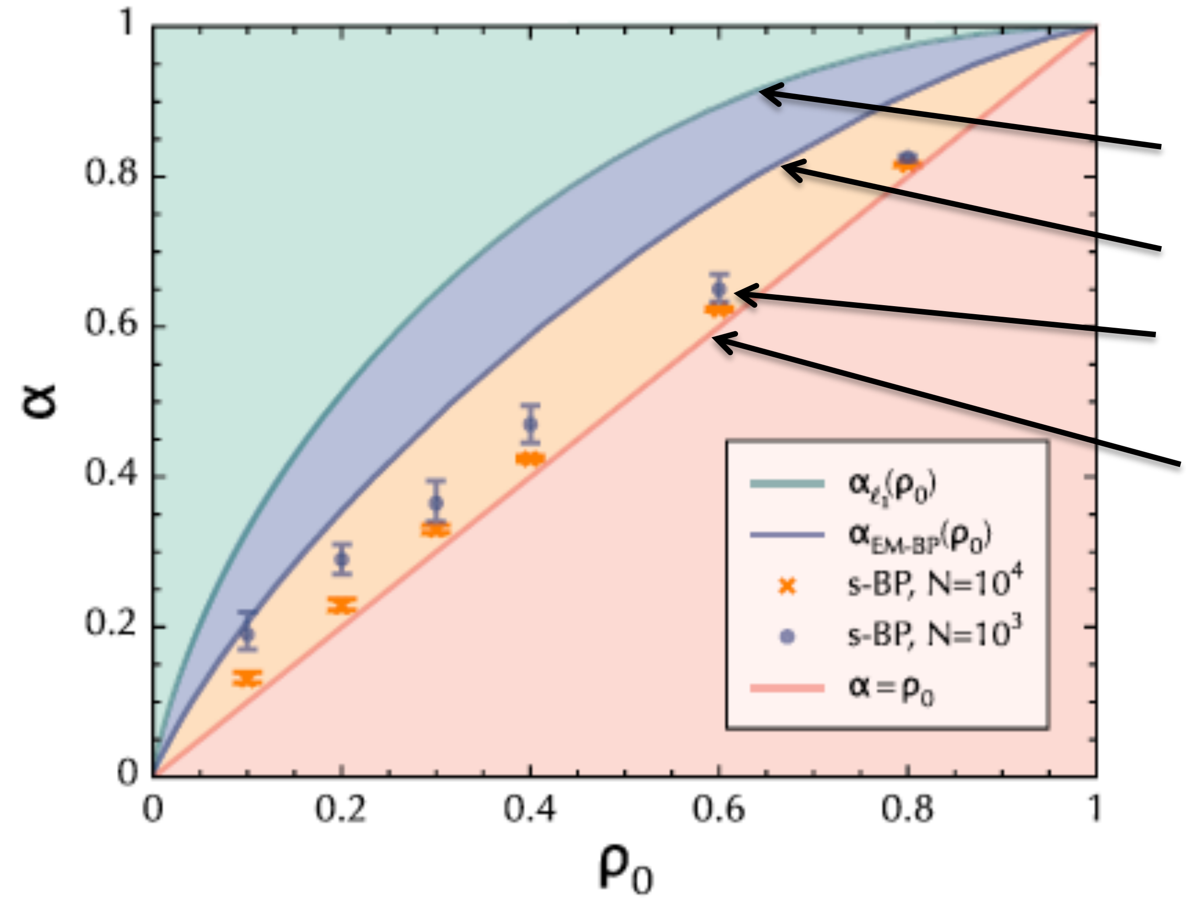

$$\left(\frac{M}{N}=\right)\alpha=\rho_0\left(=\frac{K}{N}\right)$$

Krzakala et. al., Phys. Rev. X 2, 021005 (2012)

METHODS FOR PERFORMANCE ANALYSIS/GUARANTEE

Basic method

Naïve consideration

- Measurement provides a linear combination of columns of the used matrix A.

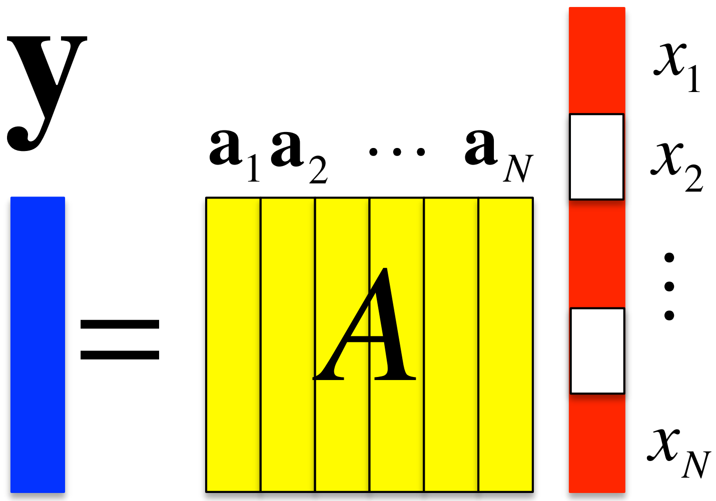

$$= x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \ldots + x_N\mathbf{a}_N$$

- If a certain pair of columns are identical (in direction), recovery of $\mathbf{y} \longrightarrow \mathbf{x}$ is impossible.

Any pairs of columns in A should be as different as possible.

Mutual coherence

- A measure for quantifying the difference of columns (in direction) for given matrix A.
- Evaluation is computationally feasible.
  - *$O(N^2M)$*

$$\mu(A) = \max_{1\le i,j\le N, i\ne j} \frac{\left|\mathbf{a}_i^{\mathrm{T}} \mathbf{a}_j\right|}{\left\|\mathbf{a}_i\right\|_2 \left\|\mathbf{a}_j\right\|_2}$$

Max. of direction cosine between two columns in a matrix A.

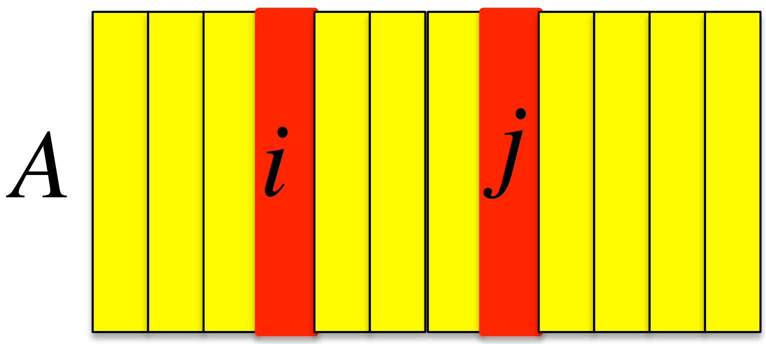

Guarantee for OMP

- Suppose that true signal $\mathbf{x}^0$ satisfies the following condition. Then, OMP is guaranteed to search it by $\left\|\mathbf{x}^0\right\|_0$ steps.

$$\left\|\mathbf{x}^0\right\|_0 \leq \frac{1}{2}\left(1+\frac{1}{\mu(A)}\right)$$

#From now on, we assume that each column in A is normalized to 1.

$$\left|\mathbf{a}_i\right|^2 = 1 \ \left(1 \leq {}^{\forall} i \leq N\right) \qquad G = A^{\mathrm{T}}A = \begin{pmatrix} 1 & \leq \mu(A) & \leq \mu(A) & \leq \mu(A) \\ \leq \mu(A) & 1 & \leq \mu(A) & \leq \mu(A) \\ \leq \mu(A) & \leq \mu(A) & \ddots & \leq \mu(A) \\ \leq \mu(A) & \leq \mu(A) & \leq \mu(A) & 1 \end{pmatrix}$$

Orthogonal Matching Pursuit (OMP)

- **Initialization:** Initialize $k=0$, and set

$$\mathbf{x}^0=\mathbf{0},\ \mathbf{r}^0=\mathbf{y}-A\mathbf{x}^0=\mathbf{y},\ S^0=\phi$$

- **Main iteration**: Increment k by 1 and perform the followings:
  - **Sweep:**

$$\epsilon(i)=\min_{x_i}\left|x_i\mathbf{a}_i-\mathbf{r}^{k-1}\right|^2=\left|\mathbf{r}^{k-1}\right|^2-\frac{\left(\mathbf{a}_i\cdot\mathbf{r}^{k-1}\right)^2}{\left|\mathbf{a}_i\right|^2}$$

Rating of columns

  - **Update Support:**

$$i_0=\arg\min_{i\notin S^{k-1}}\left\{\epsilon\left(i\right)\right\},\ S^k=S^{k-1}\cup\left\{i_0\right\}$$

  - **Update Provisional Solution:**

$$\hat{\mathbf{x}}^k=\arg\min_{\mathbf{x}_{S^k}}\left|\mathbf{y}-A_{S^k}\mathbf{x}_{S^k}\right|^2$$

Best approximation by the support of the moment

  - **Update Residual:**

$$\mathbf{r}^k=\mathbf{y}-A_{S^k}\hat{\mathbf{x}}^k$$

  - **Stopping Rule:** Stop if $\left|\mathbf{r}^k\right|<\epsilon_0$ holds. Otherwise, apply another iteration.

Proof

Let us suppose $k_0 = \|\mathbf{x}^0\|_0$. Without loss of generality, we can assume

$$\left|x_1{}^0\right| \geq \left|x_2{}^0\right| \geq \ldots \geq \left|x_{k_0}{}^0\right| > 0,\ x_{k_0+1} = x_{k_0+2} = \ldots = x_N = 0.$$

Then, the measurement is expressed as

$$\mathbf{y} = A\mathbf{x}^0 = \sum_{t=1}^{k_0} x_t{}^0 \mathbf{a}_t .$$

A sufficient condition that a column indexed by $1 \leq {}^{\exists}i \leq k_0$ is chosen in the first **Sweep** step

$$\left|\mathbf{a}_1{}^{\mathrm{T}}\mathbf{y}\right| \geq \left|\mathbf{a}_j{}^{\mathrm{T}}\mathbf{y}\right| \ \left({}^{\forall} j > k_0\right).$$

Proof

Lower bound of LHS

$$\left|\mathbf{a}_1^{\mathrm{T}}\mathbf{y}\right| = \left|\mathbf{a}_1^{\mathrm{T}}\sum_{t=1}^{k_0} x_t^0 \mathbf{a}_t\right| = \left|x_1^0 + \sum_{t=2}^{k_0} x_t^0\left(\mathbf{a}_1^{\mathrm{T}}\mathbf{a}_t\right)\right| \geq \left|x_1^0\right| - \sum_{t=2}^{k_0}\left|x_t^0\right|\left|\left(\mathbf{a}_1^{\mathrm{T}}\mathbf{a}_t\right)\right|$$
$$\geq \left|x_1^0\right| - \left|x_1^0\right|\left(k_0-1\right)\mu\left(A\right) = \left|x_1^0\right|\left(1-\left(k_0-1\right)\mu\left(A\right)\right)$$

Upper bound of RHS

$$\left|\mathbf{a}_j^{\mathrm{T}}\mathbf{y}\right| = \left|\mathbf{a}_j^{\mathrm{T}}\sum_{t=1}^{k_0} x_t^0 \mathbf{a}_t\right| \leq \sum_{t=1}^{k_0}\left|x_t^0\right|\left|\left(\mathbf{a}_j^{\mathrm{T}}\mathbf{a}_t\right)\right| \leq \left|x_1^0\right| k_0 \mu\left(A\right)$$

Lower bound of LHS > Upper bound of RHS

$$\left|x_1^0\right|\left(1-\left(k_0-1\right)\mu\left(A\right)\right) > \left|x_1^0\right| k_0\mu\left(A\right)$$

$$\Rightarrow k_0 < \frac{1}{2}\left(1+\frac{1}{\mu\left(A\right)}\right).$$

Then, the first step successfully finds a column of $1 \leq {}^{\exists}i \leq k_0$ in the true support.

Proof

After the first step

Proved by method of mathematical induction.

Suppose that columns of the true support were correctly chosen up to k-th step.
Then,

- **Update Residual** guarantees that the residual $\mathbf{r}^k$ is also a linear combination of the true support columns.
- **Update Provisional Solution** guarantees that once a column was chosen, the column has no possibility of being chosen again.

Therefore, a column in the true support is chosen at $k+1$-th step.

These mean that OMP finds all k_0 columns in the true support by k_0 steps, which guarantees the correct recovery.

Q.E.D.

Guarantee for l_1-recovery

- Suppose that true signal $\mathbf{x}^0$ satisfies the below condition. Then, l_1-recovery (basis pursuit) is guaranteed to search it.

$$\left\|\mathbf{x}^0\right\|_0 \leq \frac{1}{2}\left(1+\frac{1}{\mu(A)}\right)$$

  - Complicated proof. (skip)
  - The same expression to that for OMP. But, keep in mind that these are sufficient conditions. Actual performance is different between the two methods and objective signals.

Guarantee for Thresholding

- Suppose that true signal $\mathbf{x}^0$ satisfies the below condition. Then, Thresholding is guaranteed to search it.

$$\left\|\mathbf{x}^0\right\|_0 \leq \frac{1}{2}\left(1+\frac{1}{\mu(A)}\frac{\left|x_{\min}^0\right|}{\left|x_{\max}^0\right|}\right)$$

  – Complicated proof. (skip)
  – As this is a poor-man's version of OMP, the recovery bound becomes lower than that of OMP.

Advanced methods

Necessity for advanced methods

- The method of mutual coherence is relatively easy to follow.
- Unfortunately, its evaluation is not so accurate even in terms of order estimation.

Ex) Random matrix

$$\mu(A) = \max_{1 \le i, j \le N, i \ne j} \frac{\left| \mathbf{a}_i^{\mathrm{T}} \mathbf{a}_j \right|}{\left\| \mathbf{a}_i \right\|_2 \left\| \mathbf{a}_j \right\|_2} \sim O(M^{-1/2})$$

$$\left\| \mathbf{x}^0 \right\|_0 \le \frac{1}{2}\left(1 + \frac{1}{\mu(A)} \right) \sim O\left(M^{1/2} \right) \sim O(N^{1/2})$$

⇨ No guarantee for *O(1)* non-zero density

- For overcoming this drawback, several advanced methods have been developed so far.

Method by Candes and Tao

- Candes and Tao (2006) developed a method introducing a novel notion termed restricted isometric property (RIP) to matrix A.

Restricted Isometric Property: RIP

#We assume that each column in A is normalized to 1.

- If the following holds for any *S*-sparse vector

$$\left(1-\delta_S\right)\|\mathbf{x}\|_2^2 \le \|A\mathbf{x}\|_2^2 \le \left(1+\delta_S\right)\|\mathbf{x}\|_2^2$$

  matrix A is said to satisfy the *S*-restricted isometric property (RIP) with the RIP constant δ_S.
  - Characterization of how much the length of *S*-sparse vectors can be modified by A.

- One can show that the solution of l_1-recovery accords with the correct solution if $\delta_{2S} < \sqrt{2} - 1.$

RIP and max/min eigenvalues

- In general, the change of the vector length is characterized by max/min eigenvalues.

$$\lambda_{\min}\left(A^{\mathrm{T}}A\right)\|\mathbf{x}\|_2^2 \le \|A\mathbf{x}\|_2^2 \le \lambda_{\max}\left(A^{\mathrm{T}}A\right)\|\mathbf{x}\|_2^2$$

- When $\mathbf{x}$ is restricted to *S*-sparse vector whose non-zero positions are fixed, this inequality is reduced to that for a *submatrix* of A, A_T.

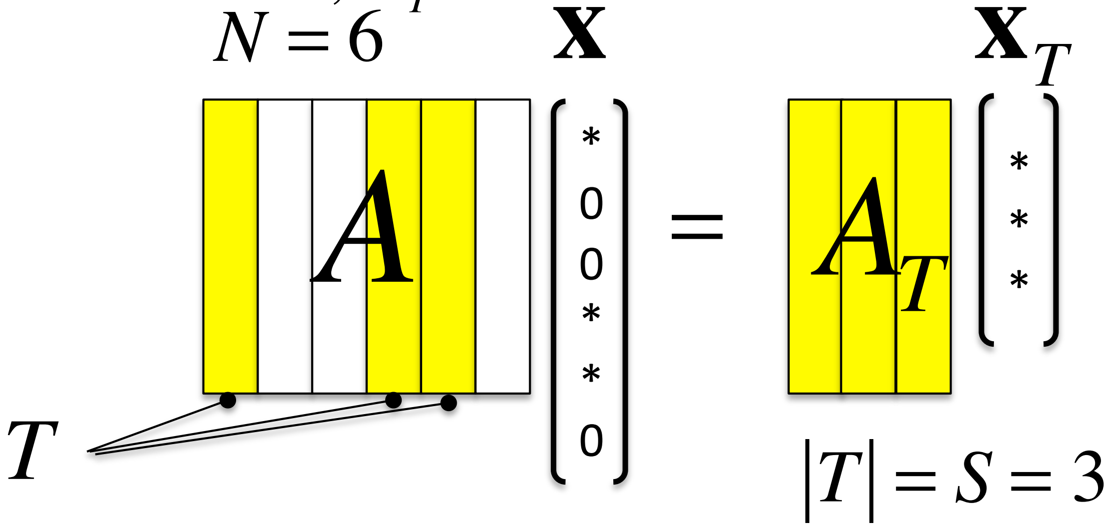

Evaluation of RIP constant

- Using the relation, we have an explicit formula for evaluating RIP constant as

$$\delta_S = \max\left\{1 - \min_{T:T\subset\{1,2,\ldots,N\},|T|=S} \lambda_{\min}\left(A_T^{\mathrm{T}} A_T\right), \max_{T:T\subset\{1,2,\ldots,N\},|T|=S} \lambda_{\max}\left(A_T^{\mathrm{T}} A_T\right) - 1\right\}$$

- But, unfortunately, this evaluation is *computationally difficult.*

Statistical evaluation

- However, one can still <u>*statistically*</u> evaluate it for ensembles of matrices.
  - For several types of random matrices, it has been shown that the l_1-recovery condition <u>*typically holds*</u> for large systems if

$$M \geq const \times S \log(N/S)$$

    is satisfied.
    - Employment of large deviation analysis of the largest/smallest eigenvalues of random matrices by Davidson and Szarek (2001).

- Guarantee for correct reconstruction of l_1-recovery for *O(1)* non-zero density.

Comments on CT method

- *Advantage*
  - Mathematically rigorous
  - Wide applicability
    - Applicable to *matrix completion* and analysis of *other recovery schemes* as well.

- *Drawback*
  - The assessment is not so good in terms of the estimation of the actual critical condition although the order estimation is accurate.
    - The bound is generally far from a threshold suggested by experimental observations.

Method by Donoho and Tanner

- Idea: Follows geometrical argument on convex (linear) programming
- Ex) $N=2, M=1, S=1$
  - l_1-reconstruction is successful with a prob. of 1/2 if $F_{\mu i} \sim N(0,1/N)$

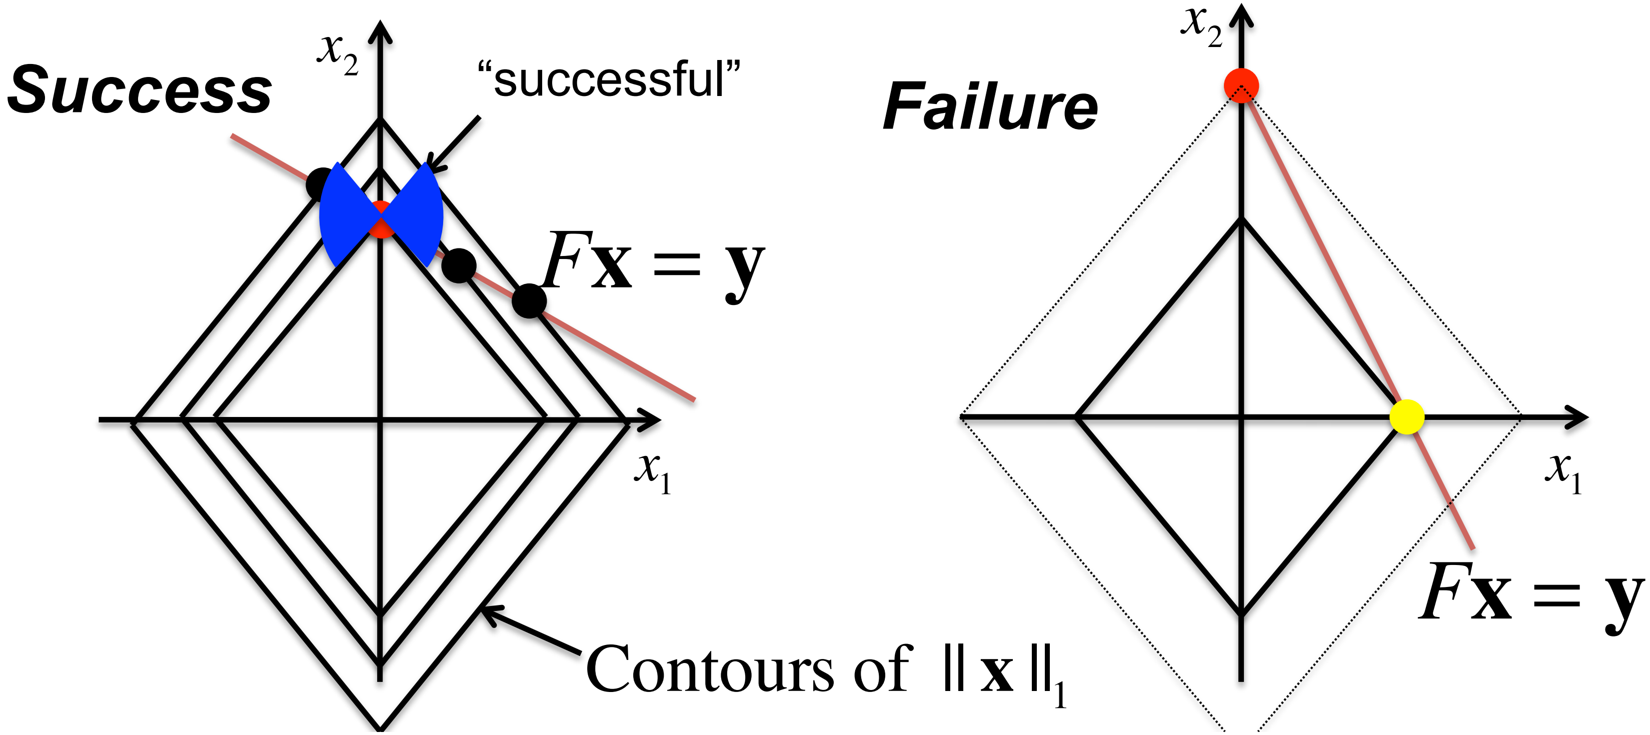

In large system limit

- Surprisingly, the success prob. can also be evaluated in higher-dimensional cases in a similar manner
- The problem is finally reduced to that of *counting the number of S-dimensional faces for M-dimensional random projection of N-dimensional polytopes*

$$P_S \sim \frac{[f_S(AQ_N)]_A}{f_S(Q_N)} \to P_S\left(\frac{M}{N}, \frac{S}{N}\right) \begin{cases} Q_N : N\text{-dimensional polytope} \\ FQ_N : M\text{-dim proj. of } Q_N \text{ by } A \\ f_S(Q_N) : \#\text{ of } S\text{-dim. faces of } Q_N \end{cases}$$

- As $M, N, S \to \infty, \alpha = M/N \sim O(1), \rho = S/N \sim O(1)$
  P_S exhibits a "phase transition"

$$\alpha \begin{cases} > \alpha_c(\rho) : & \text{success} \\ < \alpha_c(\rho) : & \text{failure} \end{cases}$$

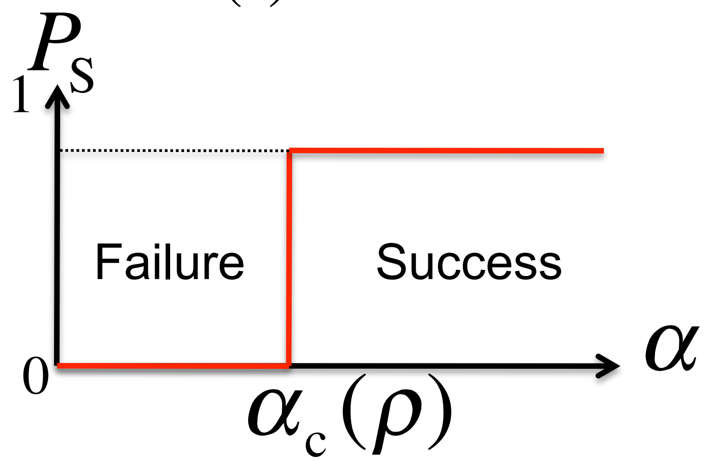

Donoho and Tanner (2009)

$$\rho = \frac{S}{M}$$

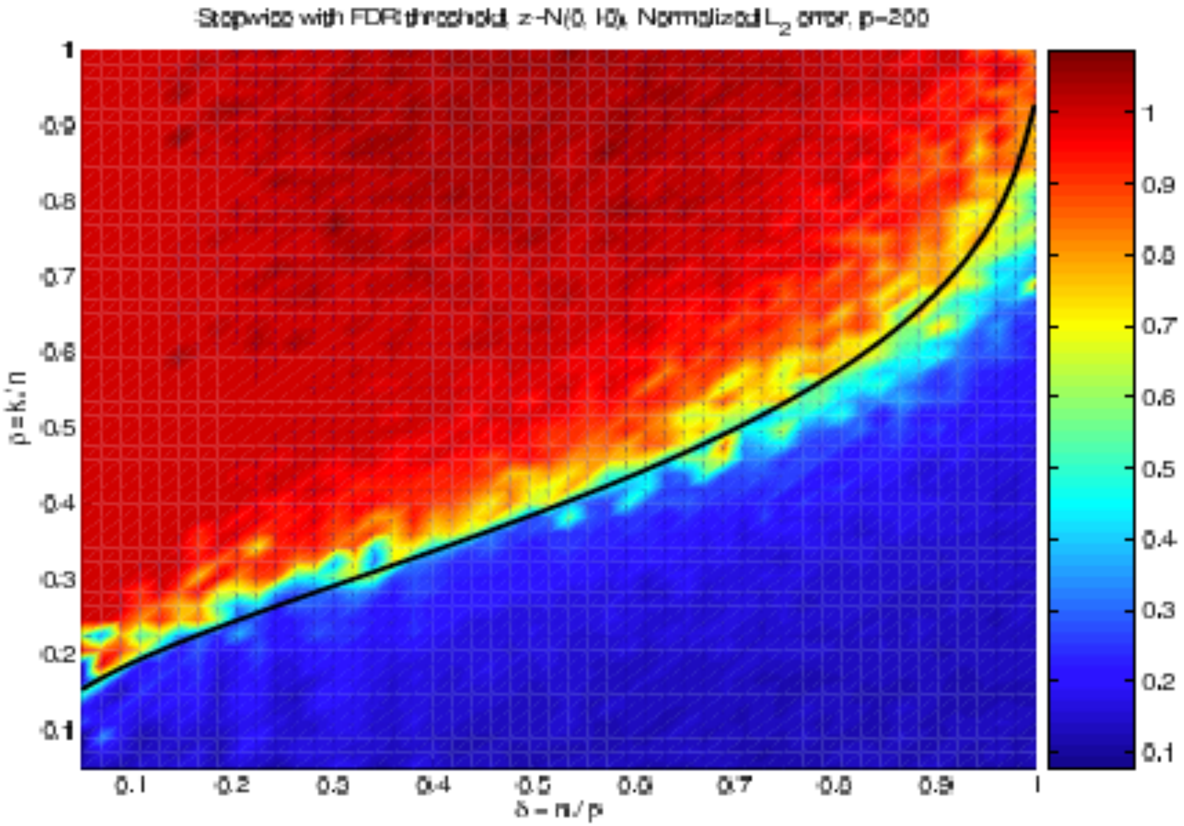

$$\delta = \frac{M}{N}$$

Curve :
Critical relation
(theory)

Red:
Low success prob.

Blue:
High success prob.

DL Donoho and J Tanner, arXiv:0906.2530 (2009)

Comments on DT Method

- *Advantage*
  - Mathematically rigorous
  - Accurate
    - The threshold is excellently in accordance with experimental behavior.
    - Even finite size effects are exactly evaluated!

- *Drawback*
  - Application to advanced settings is technically difficult.
    - Even a slight change of setting makes it difficult to follow the DT method

Method by Monatnari et al

- Idea: Reduce AMP to a macroscopic dynamics, which is termed *state evolution*, and prove that it holds exactly for large system limit $N \to \infty$.

AMP(=microscopic dynamics)

Holds exactly for $N \to \infty$

SE(=macroscopic dynamics)

$$\begin{cases} \mathbf{x}^{t+1} = \eta_k\left(A^{\mathrm{T}}\mathbf{z}^t + \mathbf{x}^t\right) \\ \mathbf{z}^t = \mathbf{y} - A\mathbf{x}^t + \frac{1}{\delta}\mathbf{z}^{t-1}\left\langle \eta_k^{t-1}\left(A^{\mathrm{T}}\mathbf{z}^{t-1} + \mathbf{x}^{t-1}\right)\right\rangle \end{cases}$$

⇒

Described by dynamics of MSE σ_t^2

$$\sigma_{t+1}^2 = \Psi\left(\sigma_t^2\right)$$

$$\Psi\left(\sigma^2\right) \equiv \mathbb{E}\left\{\left[\eta\left(X + \frac{\sigma}{\eta}Z; \lambda\sigma\right) - X\right]^2\right\}$$

Where $\delta = \frac{M}{N}$, $\langle \mathbf{u} \rangle \equiv \frac{1}{N}\sum_{i=1}^{N} u_i$

Fixed point analysis

- Criticality of convergence to $\sigma^2 = 0$ exactly accords to DT condition for i.i.d. matrices.

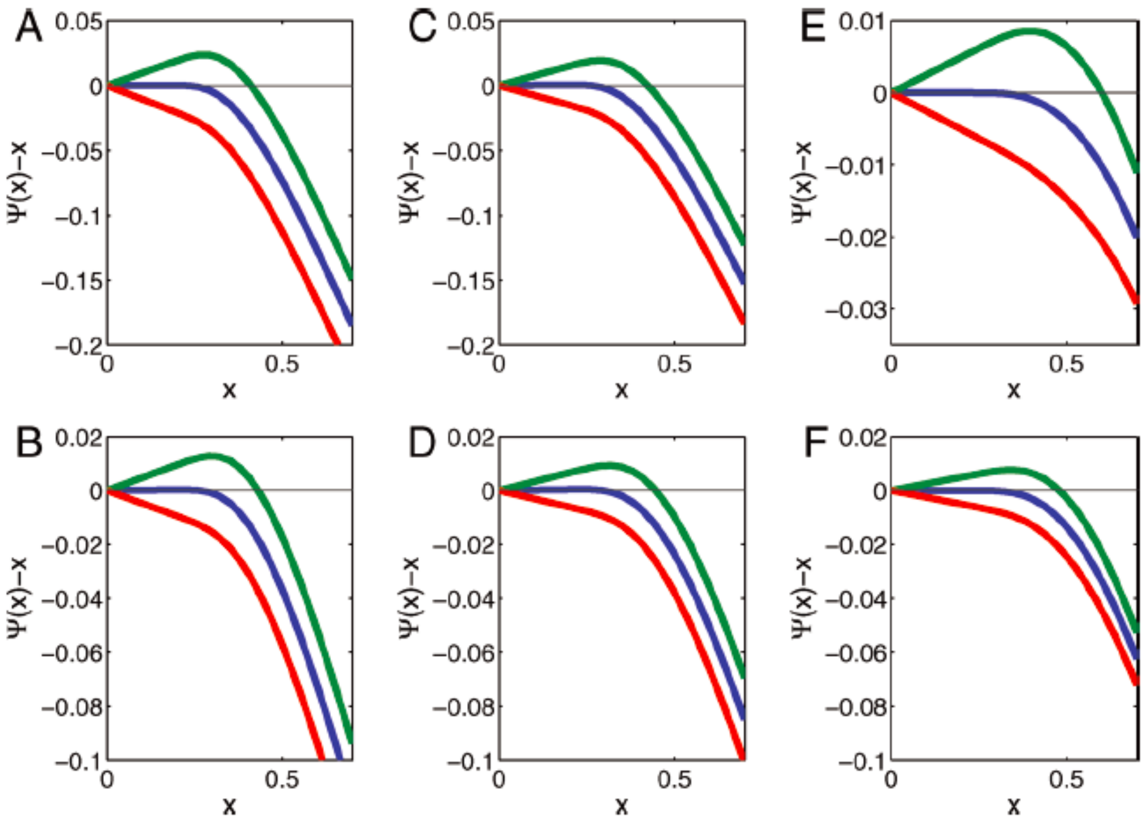

+Gauss signal   Gauss signal   +1,-1 signal

Red: Unique solution $\sigma^2 = 0$

Blue: Critical

Green: Two solutions

Donoho, Makeli and Montanari, PNAS 109, 18914 (2009)

Comments on SE

- *Advantage*
  - Mathematically rigorous
  - Accurate
    - The threshold is excellently in accordance with experimental behavior.
  - Wide applicability
    - Applicable to EM-BP
    - Applicable to non-uniform nonzero signal density
- *Drawback*
  - Application has been limited to i.i.d. matrices so far.
    - Employment for orthogonal matrices, which are practically relevant, is non-trivial due to weak correlations among entries.
      - Very recently, a group from Hong Kong posted a paper on the generalization of SE under three assumptions to orthogonal matrices to arXiv.
        Ma et al, "Turbo compressed sensing with partial DFT sensing matrix", arXiv:1408.3904

Replica method

- Idea: Follow the Bayesian framework
- Assumptions
  - Data generation (*correct* prior)

$$P(\mathbf{x}^0) = \prod_{i=1}^{N} \left((1-\rho)\delta(x_i^0) + \rho f(x_i^0) \right)$$

$\left(\begin{array}{r} f(x): \text{ arbitrary dist.} \\ \text{unit variance} \end{array} \right)$

  - Matrix ensemble (Gaussian i.i.d.)

$$A_{\mu i} \sim N(0, N^{-1})$$

Just for simplicity
One can generalized the analysis
to more advanced ensembles.

  - Large system limit

$$M, N \to \infty, \alpha = M/N \sim O(1)$$

Bayesian formulation

- *Model* (mismatched) prior based on l_p-cost

$$P_\beta(\mathbf{x}) \propto \exp\left(-\beta \|\mathbf{x}\|_p\right) \qquad (\beta > 0: \text{ inverse temp.})$$

- Posterior distribution

$$P_\beta(\mathbf{x} \mid \mathbf{x}^0, A) = \frac{\exp(-\beta \|\mathbf{x}\|_p)\delta(A\mathbf{x} - \mathbf{y}(= A\mathbf{x}^0))}{Z_\beta(\mathbf{x}^0, A)}$$

- l_p-reconstruction corresponds to the maximum a posteriori (MAP) estimator for $\forall \beta > 0$

Performance measure

- MAP accords to the posterior mean of $\beta \to \infty$
- Typical mean square error (per element)

$$\mathrm{mse} = \frac{1}{N}\left[\left\langle \left|\mathbf{x} - \mathbf{x}^0\right|^2 \right\rangle_{\beta\to\infty}\right]_{\mathbf{x}^0,F} = Q - 2m + \rho$$

$$Q \equiv \frac{1}{N}\left[\left\langle |\mathbf{x}|^2 \right\rangle_{\beta\to\infty}\right]_{\mathbf{x}^0,F} \qquad m \equiv \frac{1}{N}\left[\mathbf{x}^0 \cdot \langle \mathbf{x} \rangle_{\beta\to\infty}\right]_{\mathbf{x}^0,F}$$

- Criterion of successful reconstruction

$$\mathrm{mse} \begin{cases} =0: & \text{success} \\ >0: & \text{failure} \end{cases}$$

Negligible errors in $N \to \infty$ are allowed (Different from CT(2006) and DT(2006))

Difficulty and solution

- Exact evaluation of the nested average $\left[\langle O(\mathbf{x})\rangle\right]$ is technically difficult in general.
- For large system limit $N \to \infty$, one can do it utilizing the saddle point method under the assumption that the following analytical continuation of $\phi(n,\beta)$ is allowed (= Replica method).

$$\phi(n,\beta)=\frac{1}{N}\log\left[\left(Z_\beta(\mathbf{x}^0,A)\right)^n\right]\quad (n\in\mathbb{N})$$

⇩ Analytical continuation

$$\phi(n,\beta)=\frac{1}{N}\log\left[\left(Z_\beta(\mathbf{x}^0,A)\right)^n\right]\quad (n\in\mathbb{R})$$

- The critical relation is evaluated by monitoring "$\mathrm{mse}=0 \to >0$".

Results

YK, Wadayama and Tanaka, J. Stat. Mech. (2009) L09003

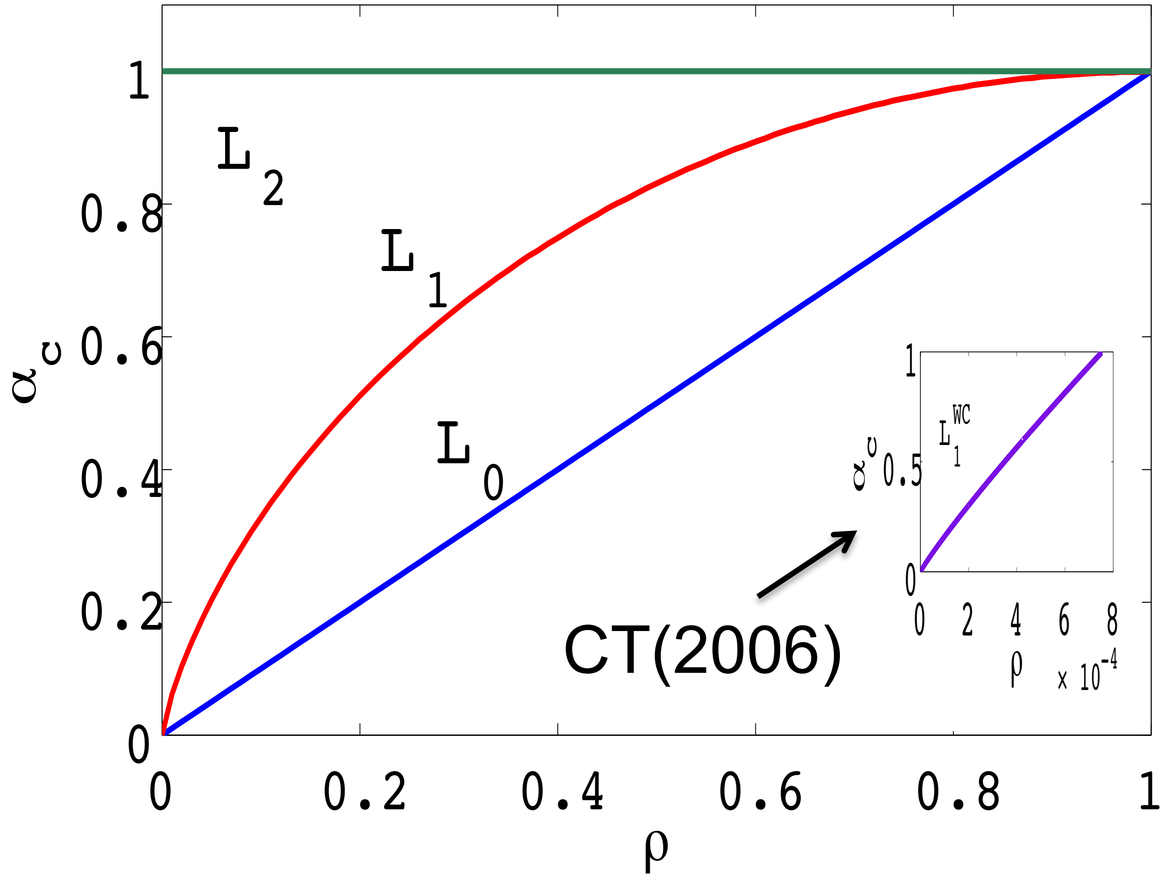

l_0: $\alpha_c(\rho) = \rho$
Correct

l_1: Numerical
Correct
(Identical to DT(2006))

l_2: $\alpha_c(\rho) = 1$
Correct

Experiments for l_1-reconst.

- Employed CVX (Grant and Boyd (2009))
  - Excellent agreement between theory and experiments

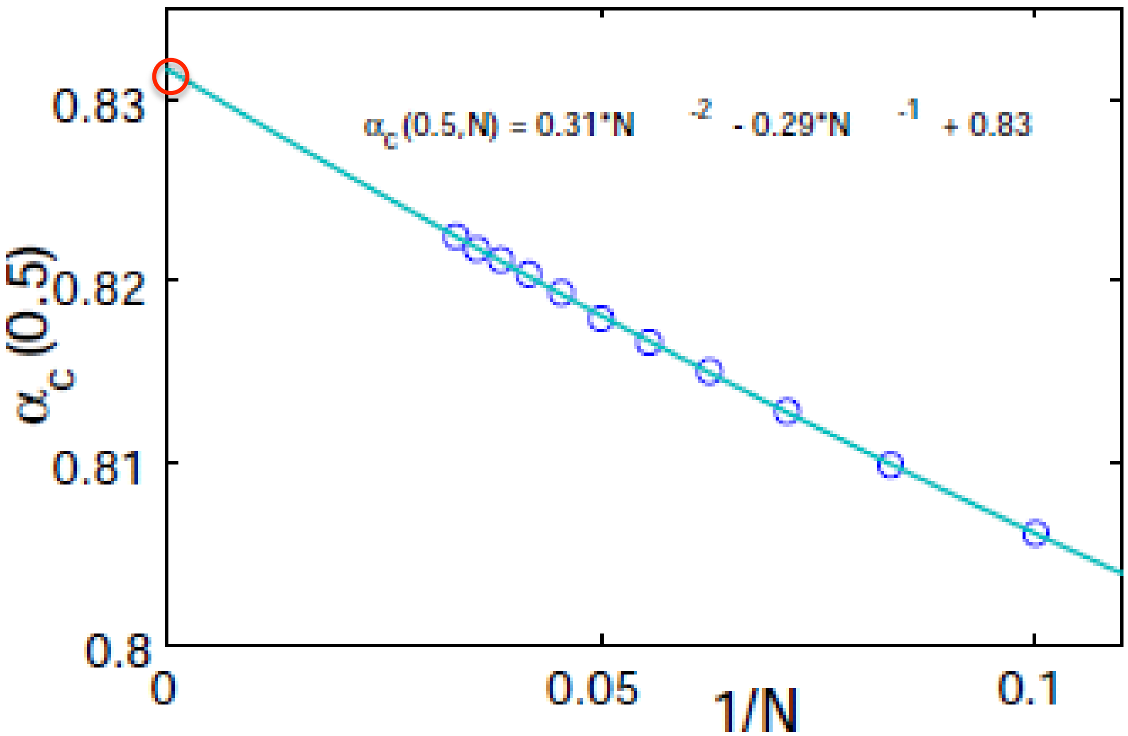

YK, Wadayama and Tanaka, J. Stat. Mech. (2009) L09003

ρ=0.5
10^6 experiments for each of N=10,12,…,30.
$\alpha_c(\rho, N\to\infty)$ is inferred by the quadratic fitting with respect to 1/N

Theoretical prediction

$\alpha_c = \underline{0.831}29\ldots$

Experimental value

$\alpha_c = \underline{0.831}65\ldots$

Comments on replica method

- *Advantage*
  - Accurate
    - The threshold is excellently in accordance with experimental behavior.
  - Wide applicability
    - Applicable to wide classes of matrix ensembles.

Ex) Random orthogonal matrices

$$A = UDV^{\mathrm{T}}$$

D : *M*x*N* diagonal matrix whose diagonal entries asymptotically follow a certain dist.

V : Sample from uniform dist. of *N*x*N* orthogonal matrices

YK, Vehkapera and Chatterjee (2012)
Vehkapera, YK and Chatterjee (2014)
YK and Vehkapera (2014)

- *Drawback*
  - Mathematically non-rigorous

Comparison among the advanced methods

Characteristics of each method may be summarized as follows.

	Math. Validity	Accuracy	Applicability
Candes-Tao (RIP)	○	×	◎
Donoho-Tanner (combinatorics)	○	○	×
State evolution	○	○	○
Replica method	×	○	◎

SUMMARY AND DISCUSSION

Summary

- What I talked about
  - Problem setup of *noiseless* compressed sensing
  - Introduction of several signal recovery algorithms
  - Introduction of several methods for performance analysis/guarantee

- What I did not talk about
  - *Many*
    - ✓ Noisy compressed sensing
    - ✓ Advanced matrix ensembles (union of random orthogonal matrices)
    - ✓ Expectation consistency scheme for signal recovery
    - ✓ *Matrix completion*
    - ✓ Dictionary learning, matrix factorization
    - ✓ …

Summary

- What I currently am interested in
  1. Analysis of *greedy* algorithms
     - ✓ No accurate analysis yet. But, excellent performance is observed in experiment.
     - ✓ Other two strategies, *convex relaxation* and *probabilistic inference,* are a good match with equilibrium stat. mech. because cost (energy) function is provided. But, analysis of greedy algorithms is nontrivial.
     - ✓ Dynamical theory may be required. Formalisms of Path integral, dynamical replica, statistical neurodynamics, …, may be useful.
  2. Matrix completion
     - ✓ RIP analysis was applied to the critical condition of l_1-recovery (Recht et al. (2010),(2011)).
     - ✓ But, little quantitatively accurate analysis.
     - ✓ Replica theory may be able to provide more accurate estimation.

Matrix completion

- Inference of missing entries of a matrix from existing entries
  - Important for recommendation system and so on.

Ex) Matrix of cinema rating

Evaluator/ Cinema	Alice	Bob	Charles	David
Star Trek	2		4	
Rain Man		1		3
Godzilla	3	4		
Psycho			3	
Titanic	5		2	3

Mathematical formulation

"Signal"= matrix: $\mathbf{X} = \left(X_{ij} \right) \in \mathbb{R}^{N_1 \times N_2}$

Set of "measurement matrices": $\left\{ A^1, A^2, \ldots, A^{\mu} \right\}$ (Measurement of $\left(i_{\mu}, j_{\mu} \right)$ entry of $\mathbf{X}$ $\Leftrightarrow A^{\mu} = \left(A_{ij}{}^{\mu} \right) = \left(\delta_{i, i_{\mu}} \times \delta_{j, j_{\mu}} \right)$)

Measurement: $Y_{\mu} = \mathrm{Tr}\left(\left(A^{\mu} \right)^{\mathrm{T}} \mathbf{X} \right) = \sum_{i,j} A_{ij}^{\mu} X_{ij} \quad \left(\mu = 1, 2, \ldots, M \right)$

Signal density of $\mathbf{X}$: $\mathrm{rank}\left(\mathbf{X} \right)$

"l_p-norm": $\left\| \mathbf{X} \right\|_p = \left(\sum_{i=1}^{\min\left(N_1, N_2 \right)} \left| \sigma_i \left(\mathbf{X} \right) \right|^p \right)^{1/p}$ (Schatten *p*-norm)

Problem to solve and convex relaxation

- l_0-recovery
  - Computationally hard

$$\min_{\mathbf{X}} \operatorname{rank}(\mathbf{X}) \text{ subj. to } \mathrm{Y}_\mu = \operatorname{Tr}\left(\left(A^\mu\right)^{\mathrm{T}} \mathbf{X}\right) \quad (\mu = 1, 2, \ldots, M)$$

- l_1-recovery
  - Computationally feasible

$$\min_{\mathbf{X}} \|\mathbf{X}\|_1 \text{ subj. to } \mathrm{Y}_\mu = \operatorname{Tr}\left(\left(A^\mu\right)^{\mathrm{T}} \mathbf{X}\right) \quad (\mu = 1, 2, \ldots, M)$$

$$\Leftrightarrow \min_{\mathbf{Y},\mathbf{Z}} \begin{pmatrix} \mathbf{Y} & O \\ O & \mathbf{Z} \end{pmatrix} \text{ subj. to } \begin{pmatrix} \mathbf{Y} & \mathbf{X} \\ \mathbf{X}^{\mathrm{T}} & \mathbf{Z} \end{pmatrix} \succeq 0,\ \mathrm{Y}_\mu = \operatorname{Tr}\left(\left(A^\mu\right)^{\mathrm{T}} \mathbf{X}\right) \quad (\mu = 1, 2, \ldots, M)$$

Soluble by semi-definite programming by convex packages

Greedy algorithm and probabilistic inference may also be candidates (maybe).

RIP analysis

- Linear measurement

$$\mathbf{Y} = \mathbf{A}\mathbf{X} \Leftrightarrow Y_\mu = \mathrm{Tr}\left(\left(A^\mu\right)^\mathrm{T} \mathbf{X}\right) = \sum_{i,j} A_{ij}^\mu X_{ij} \quad \left(\mu = 1, 2, \ldots, M\right)$$

- RIP constant δ_r for operator $\mathbf{A}$
  - Min of δ that satisfies the following for any $\mathbf{X}$ whose rank $\leq r$.

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

- Condition for l_1-recovery for i.i.d. Gaussian $\mathbf{A}$
  - $\delta_{5r} < 1/10 \Rightarrow$ Matrix whose rank $\leq r$ can be recovered by l_1-recovery.
  - This condition typically holds for large systems if the number of measurements satisfies the following.

$$M \geq O\left(r\left(N_1 + N_2\right)\right)$$ [Recht et al. (2010),(2011)]

Last slide

- Conventional principle of information processing
  ⇒"Least square method"
  - Why? → Analytically feasible.
    - Unique feasible option before computers were developed.

- Background of compressed sensing
  - Computers have become fast and cheap.
    → Now, we do not necessarily rely on analytically feasible methods.

- Consequence of *game change*
  - "Least square" → "Sparsity"
  - Revolution in the level of *principle*.
    This research trend probably continues substantially long.

Thanks for your attention