Robust image recovery via total-variation minimization

Deanna Needell

Claremont McKenna College

CCMS Colloquium, Mar. 2012

Outline

- Compressed Sensing (CS)
  - Motivation
  - Mathematical Formulation & Methods
- Imaging with CS
  - Theoretical possibilities
  - Total Variation
  - Empirical observations
  - New Results

Collaborator

Joint work with Rachel Ward [Univ. of Texas, Austin]

D. Needell and R. Ward. Stable image reconstruction using total variation minimization, Submitted, Mar. 2012.

The Data Deluge

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
    - We need the resources.
    - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
    - Enter the world of: Compressed Sensing (CS)
    - CS gives us efficient compression techniques: "Compressed"
    - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
    - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: “Compressed”
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: “Sensing”
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

The Data Deluge

How can we handle all this data?

- Build hardware that can store and trasmit more data.
  - We need the resources.
  - There are fundamental limitiations to data storage.
- Design more efficient compression methods.
  - Enter the world of: Compressed Sensing (CS)
  - CS gives us efficient compression techniques: "Compressed"
  - More surprisingly, we can acquire the compression without ever having to acquire the entire object!: "Sensing"
  - CS has numerous applications (Radar, Error Correction, Computational Biology (DNA Microarrays), Geophysical Data Analysis, Data Mining, classification, Neuroscience, Imaging ...)

Why is compression possible?

256×256 "Boats" image

Because most practical signals, such as images, contain much less information than their dimension would suggest.

Why is compression possible?

256×256 "Boats" image

Because most practical signals, such as images, contain much less information than their dimension would suggest.

Why is compression possible?

  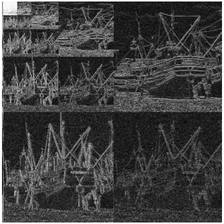

Assume f is s-sparse:

- In the coordinate basis: $\|f\|_0 \stackrel{\text{def}}{=} |\text{supp}(f)| \leq s \ll d$
- In some orthonormal basis: $f = Dx$ where $\|x\|_0 \leq s \ll d$

In practice, we encounter compressible signals.

Why is compression possible?

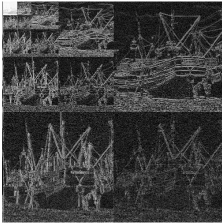

Assume f is s-sparse:

- In the coordinate basis: $\|f\|_0 \stackrel{\text{def}}{=} |\text{supp}(f)| \le s \ll d$
- In some orthonormal basis: $f = Dx$ where $\|x\|_0 \le s \ll d$

In practice, we encounter compressible signals.

Why is compression possible?

  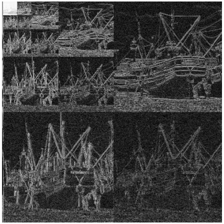

Assume f is s-sparse:

- In the coordinate basis: $\|f\|_0 \stackrel{\text{def}}{=} |\text{supp}(f)| \le s \ll d$
- In some orthonormal basis: $f = Dx$ where $\|x\|_0 \le s \ll d$

In practice, we encounter compressible signals.

Mathematical Formulation

To compress a signal, we take a small number of measurements:

1. Signal of interest $f \in \mathbb{R}^{N \times N}$
2. Measurement operator $A : \mathbb{R}^{N \times N} \to \mathbb{R}^m$ $(m \ll N^2)$
3. Measurements $y = Af$.

$$\begin{bmatrix} y \end{bmatrix} = \begin{bmatrix} & & A & & \end{bmatrix} \begin{bmatrix} f \\ \\ \\ \end{bmatrix}$$

4. y is the compression of f!

Mathematical Formulation

To compress a signal, we take a small number of measurements:

1. Signal of interest $f \in \mathbb{R}^{N \times N}$
2. Measurement operator $A : \mathbb{R}^{N \times N} \to \mathbb{R}^m$ $(m \ll N^2)$
3. Measurements $y = Af$.

$$\begin{bmatrix} y \end{bmatrix} = \begin{bmatrix} & & A & & \end{bmatrix} \begin{bmatrix} f \\ \\ \\ \end{bmatrix}$$

4. y is the compression of f!

Mathematical Formulation

To compress a signal, we take a small number of measurements:

1. Signal of interest $f \in \mathbb{R}^{N \times N}$
2. Measurement operator $A : \mathbb{R}^{N \times N} \to \mathbb{R}^m$ $(m \ll N^2)$
3. Measurements $y = Af$.

$$\begin{bmatrix} y \end{bmatrix} = \begin{bmatrix} & A & \end{bmatrix} \begin{bmatrix} f \\ \\ \\ \end{bmatrix}$$

4. y is the compression of f!

Mathematical Formulation

To compress a signal, we take a small number of measurements:

1. Signal of interest $f \in \mathbb{R}^{N \times N}$
2. Measurement operator $A : \mathbb{R}^{N \times N} \to \mathbb{R}^m$ $(m \ll N^2)$
3. Measurements $y = Af$.

$$\begin{bmatrix} y \end{bmatrix} = \begin{bmatrix} \quad A \quad \end{bmatrix} \begin{bmatrix} f \end{bmatrix}$$

4. y is the compression of f!

Mathematical Formulation

To compress a signal, we take a small number of measurements:

1. Signal of interest $f \in \mathbb{R}^{N\times N}$
2. Measurement operator $A : \mathbb{R}^{N\times N} \to \mathbb{R}^m$ $(m \ll N^2)$
3. Measurements $y = Af + \xi$.

$$\begin{bmatrix} y \end{bmatrix} = \begin{bmatrix} & A & \end{bmatrix} \begin{bmatrix} f \\ \\ \end{bmatrix} + \begin{bmatrix} \xi \end{bmatrix}$$

4. y is the compression of f!
5. And then the measurements get corrupted with noise.

Compression

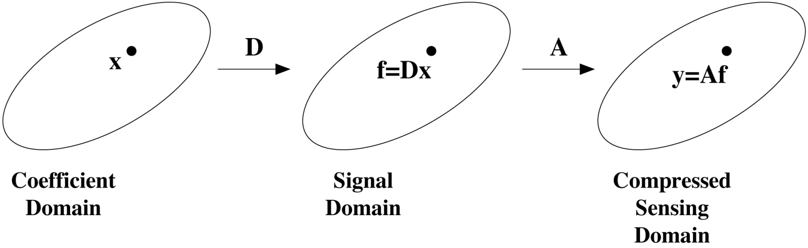

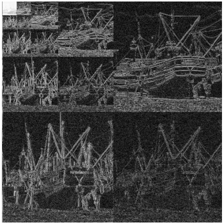

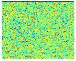

Questions

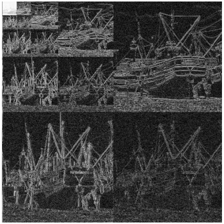

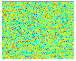

1. What type of measurement operator A can we use?
2. How do we reconstruct the signal f from the compressed measurements y?

Questions

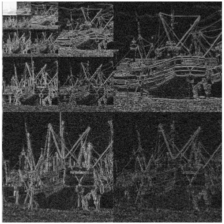

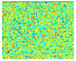

1. What type of measurement operator A can we use?
2. How do we reconstruct the signal f from the compressed measurements y?

Questions

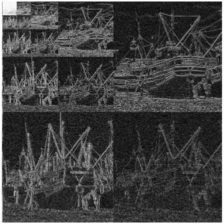  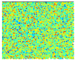

1. What type of measurement operator A can we use?
2. How do we reconstruct the signal f from the compressed measurements y?

Methods for Compressed Sensing

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} (\sum_i |z_i|^p)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} (\sum_i |z_i|^2)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} (\sum_i |z_i|)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} \left(\sum_i |z_i|^p\right)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} \left(\sum_i |z_i|\right)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \operatorname{argmin}_z g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} \left(\sum_i |z_i|^p\right)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} \left(\sum_i |z_i|\right)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} \left(\sum_i |z_i|^p\right)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} \left(\sum_i |z_i|\right)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} (\sum_i |z_i|^p)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} (\sum_i |z_i|^2)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} (\sum_i |z_i|)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z \, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} (\sum_i |z_i|^p)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} (\sum_i |z_i|)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} \left(\sum_i |z_i|^p\right)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} \left(\sum_i |z_i|\right)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} \left(\sum_i |z_i|^p\right)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} \left(\sum_i |z_i|^2\right)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} \left(\sum_i |z_i|\right)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} (\sum_i |z_i|^p)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} (\sum_i |z_i|^2)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} (\sum_i |z_i|)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

Review and Notation

- ℓ_p-norms: $\|z\|_p \stackrel{\text{def}}{=} (\sum_i |z_i|^p)^{1/p}$
- Usual (Euclidean ℓ_2) distance: $\|z\|_2 \stackrel{\text{def}}{=} (\sum_i |z_i|^2)^{1/2}$
- ℓ_1 (Taxicab) distance: $\|z\|_1 \stackrel{\text{def}}{=} (\sum_i |z_i|)$
- The ℓ_2-ball: $\{z : \|z\|_2 \leq 1\}$ (circle/sphere)
- The ℓ_1-ball: $\{z : \|z\|_1 \leq 1\}$ (diamond/octahedron)
- For signal f, f_s (f_s^B) is its best s-sparse representation (in basis B)
- $\hat{f}$ will denote the reconstruction of f
- $h = \text{argmin}_z\, g(z)$ is the argument z which minimizes $g(z)$

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \operatorname*{argmin}_z \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \operatorname*{argmin}_z \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \underset{z}{\operatorname{argmin}} \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \operatorname*{argmin}_z \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \operatorname*{argmin}_z \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

Easy Theorem

Assume A is one-to-one on all s-sparse signals. Assume there is no noise. Reconstruct an s-sparse signal f by:

$$\hat{f} = \operatorname*{argmin}_z \|z\|_0 \quad \text{such that} \quad Az = y.$$

Then we reconstruct f perfectly: $\hat{f} = f$.

Cool, except this problem is NP-Hard!

How should we reconstruct f?

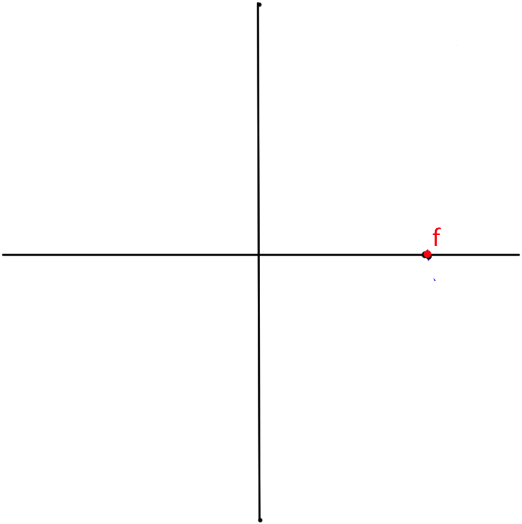

How should we reconstruct f?

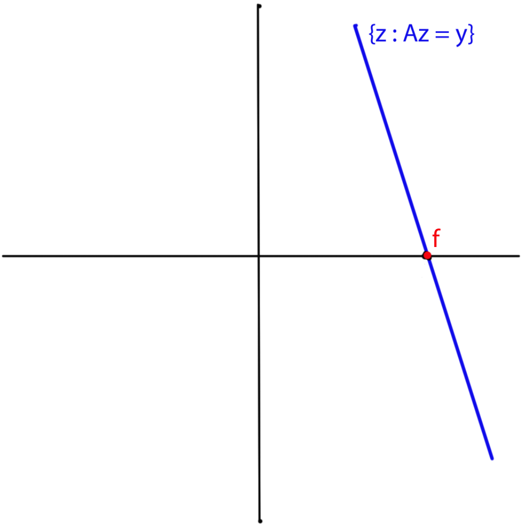

How should we reconstruct f?

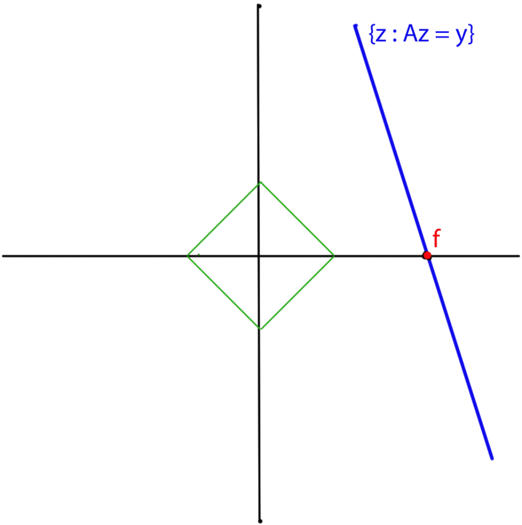

How should we reconstruct f?

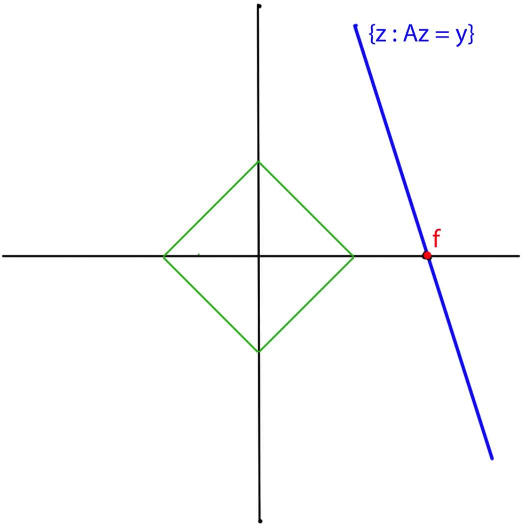

How should we reconstruct f?

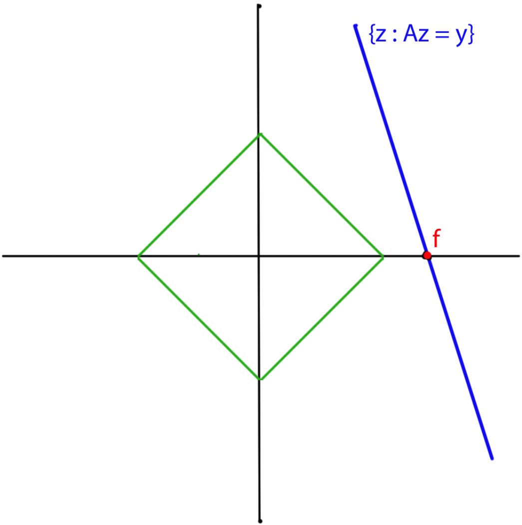

How should we reconstruct f?

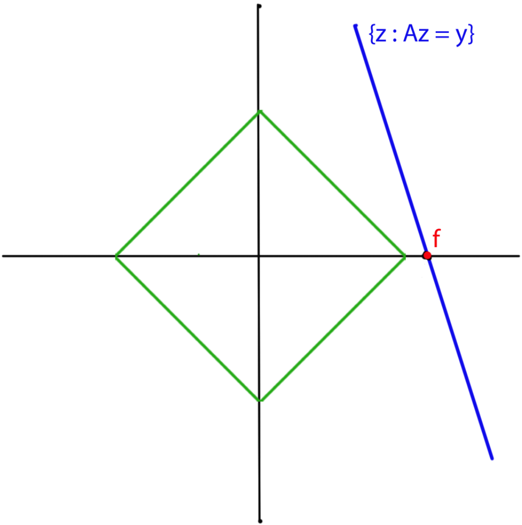

How should we reconstruct f?

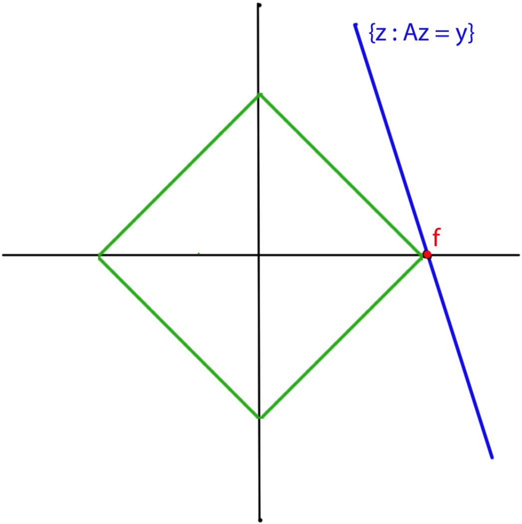

Was that contrived?

Wait, did I cheat?

Was that contrived?

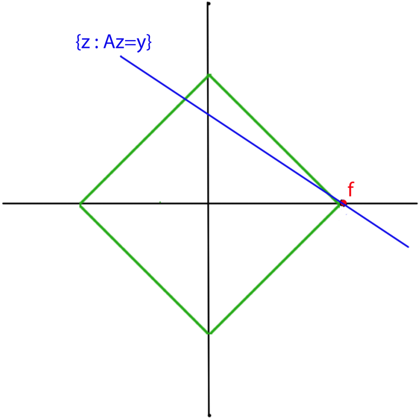

But in higher dimensions, for “sufficiently random” operators A, this picture happens with extremely low probability!

Was that contrived?

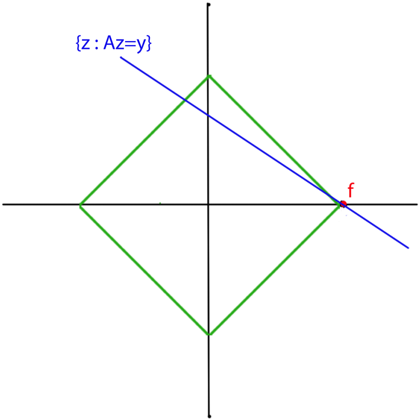

But in higher dimensions, for "sufficiently random" operators A, this picture happens with extremely low probability!

Okay, but what about noise?

Recall $y = Af + \xi$.

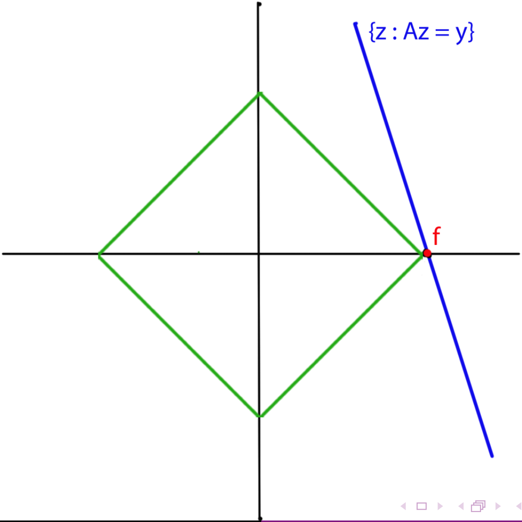

Okay, but what about noise?

Recall $y = Af + \xi$.

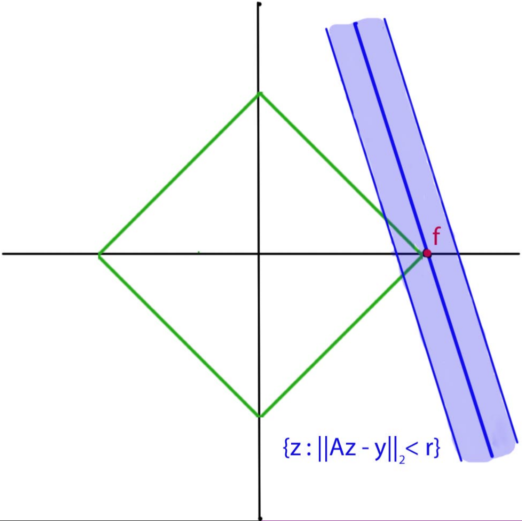

Idea behind a method

From our geometric intuition, we can reconstruct the signal f from its measurements $y = Af + \xi$:

1. If the measurement operator A is "well-behaved"
2. We can reconstruct our image f by solving

$$\hat{f} = \underset{z}{\operatorname{argmin}} \|z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

   where r bounds the noise term: $\|\xi\|_2 \leq r$.
3. If f is sparse with respect to some orthonormal basis B, meaning, $f = Bx$ for sparse x,

$$\hat{f} = \underset{z}{\operatorname{argmin}} \|B^{-1}z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

4. We call these methods the ℓ_1-minimization method, which are easily solved by convex programming methods.

Idea behind a method

From our geometric intuition, we can reconstruct the signal f from its measurements $y = Af + \xi$:

1. If the measurement operator A is "well-behaved"
2. We can reconstruct our image f by solving

$$\hat{f} = \underset{z}{\text{argmin}} \|z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

   where r bounds the noise term: $\|\xi\|_2 \leq r$.

3. If f is sparse with respect to some orthonormal basis B, meaning, $f = Bx$ for sparse x,

$$\hat{f} = \underset{z}{\text{argmin}} \|B^{-1}z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

4. We call these methods the ℓ_1-minimization method, which are easily solved by convex programming methods.

Idea behind a method

From our geometric intuition, we can reconstruct the signal f from its measurements $y = Af + \xi$:

1. If the measurement operator A is "well-behaved"
2. We can reconstruct our image f by solving

$$\hat{f} = \operatorname*{argmin}_z \|z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

   where r bounds the noise term: $\|\xi\|_2 \leq r$.
3. If f is sparse with respect to some orthonormal basis B, meaning, $f = Bx$ for sparse x,

$$\hat{f} = \operatorname*{argmin}_z \|B^{-1}z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

4. We call these methods the ℓ_1-minimization method, which are easily solved by convex programming methods.

Idea behind a method

From our geometric intuition, we can reconstruct the signal f from its measurements $y = Af + \xi$:

1. If the measurement operator A is "well-behaved"
2. We can reconstruct our image f by solving

$$\hat{f} = \operatorname*{argmin}_z \|z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \le r,$$

   where r bounds the noise term: $\|\xi\|_2 \le r$.
3. If f is sparse with respect to some orthonormal basis B, meaning, $f = Bx$ for sparse x,

$$\hat{f} = \operatorname*{argmin}_z \|B^{-1}z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \le r,$$

4. We call these methods the ℓ_1-minimization method, which are easily solved by convex programming methods.

Idea behind a method

From our geometric intuition, we can reconstruct the signal f from its measurements $y = Af + \xi$:

1. If the measurement operator A is "well-behaved"
2. We can reconstruct our image f by solving

$$\hat{f} = \underset{z}{\operatorname{argmin}} \|z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

   where r bounds the noise term: $\|\xi\|_2 \leq r$.
3. If f is sparse with respect to some orthonormal basis B, meaning, $f = Bx$ for sparse x,

$$\hat{f} = \underset{z}{\operatorname{argmin}} \|B^{-1}z\|_1 \quad \text{such that} \quad \|Az - y\|_2 \leq r,$$

4. We call these methods the ℓ_1-minimization method, which are easily solved by convex programming methods.

How do we actually reconstruct the signal f from measurements y?

ℓ_1-minimization [Candès-Romberg-Tao '05]

Let A satisfy the *Restricted Isometry Property* and suppose $\hat{f}$ is the solution to the ℓ_1-minimization problem, from measurements $y = Af + \xi$ (with $\|\xi\|_2 \leq \varepsilon$). Then we can stably recover the signal f:

$$\|f - \hat{f}\|_2 \lesssim \varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}.$$

Thus, the reconstruction error is proportional to the noise level and the tail of the compressible signal. This error bound is optimal.

How do we actually reconstruct the signal f from measurements y?

ℓ_1-minimization [Candès-Romberg-Tao '05]

Let A satisfy the *Restricted Isometry Property* and suppose $\hat{f}$ is the solution to the ℓ_1-minimization problem, from measurements $y = Af + \xi$ (with $\|\xi\|_2 \leq \varepsilon$). Then we can stably recover the signal f:

$$\|f - \hat{f}\|_2 \lesssim \varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}.$$

Thus, the reconstruction error is proportional to the noise level and the tail of the compressible signal. This error bound is optimal.

How do we actually reconstruct the signal f from measurements y?

ℓ_1-minimization [Candès-Romberg-Tao '05]

Let A satisfy the *Restricted Isometry Property* and suppose $\hat{f}$ is the solution to the ℓ_1-minimization problem, from measurements $y = Af + \xi$ (with $\|\xi\|_2 \leq \varepsilon$). Then we can stably recover the signal f:

$$\|f - \hat{f}\|_2 \lesssim \varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}.$$

Thus, the reconstruction error is proportional to the noise level and the tail of the compressible signal. This error bound is optimal.

How do we actually reconstruct the signal f from measurements y?

ℓ_1-minimization [Candès-Romberg-Tao '05]

Let A satisfy the *Restricted Isometry Property* and suppose $\hat{f}$ is the solution to the ℓ_1-minimization problem, from measurements $y = Af + \xi$ (with $\|\xi\|_2 \leq \varepsilon$). Then we can stably recover the signal f:

$$\|f - \hat{f}\|_2 \lesssim \varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}.$$

Thus, the reconstruction error is proportional to the noise level and the tail of the compressible signal. This error bound is optimal.

Restricted Isometry Property

- A satisfies the Restricted Isometry Property (RIP) when there is $\delta < c$ such that

$$(1-\delta)\|f\|_2 \leq \|Af\|_2 \leq (1+\delta)\|f\|_2 \quad \text{whenever } \|f\|_0 \leq s.$$

- Gaussian or Bernoulli measurement matrices satisfy the RIP with high probability when

$$m \gtrsim s \log N.$$

- Random Fourier and others with fast multiply have similar property: $m \gtrsim s \log^4 N$.

Restricted Isometry Property

- A satisfies the Restricted Isometry Property (RIP) when there is $\delta < c$ such that

$$(1-\delta)\|f\|_2 \leq \|Af\|_2 \leq (1+\delta)\|f\|_2 \quad \text{whenever } \|f\|_0 \leq s.$$

- Gaussian or Bernoulli measurement matrices satisfy the RIP with high probability when

$$m \gtrsim s \log N.$$

- Random Fourier and others with fast multiply have similar property: $m \gtrsim s \log^4 N$.

Restricted Isometry Property

- A satisfies the Restricted Isometry Property (RIP) when there is $\delta < c$ such that

$$(1-\delta)\|f\|_2 \le \|Af\|_2 \le (1+\delta)\|f\|_2 \quad \text{whenever } \|f\|_0 \le s.$$

- Gaussian or Bernoulli measurement matrices satisfy the RIP with high probability when

$$m \gtrsim s \log N.$$

- Random Fourier and others with fast multiply have similar property: $m \gtrsim s \log^4 N$.

Imaging via Compressed Sensing

Image sparsity

Recall, some images are sparse:

Imaging via compressed sensing

Results in compressed sensing [CRT '06, etc.] imply:

- if an image $f \in \mathbb{R}^{N \times N}$ is s-sparse
- if the measurement operator satisfies the RIP
- then using traditional ℓ_1-minimization,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

Results in compressed sensing [CRT '06, etc.] imply:

- if an image $f \in \mathbb{R}^{N\times N}$ is s-sparse
- if the measurement operator satisfies the RIP
- then using traditional ℓ_1-minimization,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

Results in compressed sensing [CRT '06, etc.] imply:

- if an image $f \in \mathbb{R}^{N\times N}$ is s-sparse
- if the measurement operator satisfies the RIP
- then using traditional ℓ_1-minimization,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

Results in compressed sensing [CRT '06, etc.] imply:

- if an image $f \in \mathbb{R}^{N\times N}$ is s-sparse
- if the measurement operator satisfies the RIP
- then using traditional ℓ_1-minimization,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

Recall, some images are sparse with respect to some orthonormal basis, like the Haar wavelet basis:

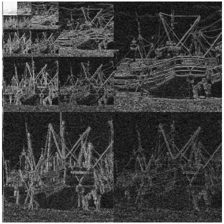

Figure: Haar basis functions

Imaging via compressed sensing

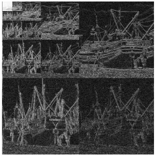

Results in compressed sensing [CRT '06, etc.] imply:

- if $f \in \mathbb{R}^{N\times N}$ is s-sparse in an orthonormal basis B
- if the measurement operator satisfies the RIP
- then using ℓ_1-minimization with basis B,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s^B\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

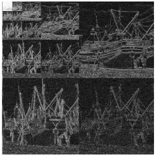

Results in compressed sensing [CRT '06, etc.] imply:

- if $f \in \mathbb{R}^{N \times N}$ is s-sparse in an orthonormal basis B
- if the measurement operator satisfies the RIP
- then using ℓ_1-minimization with basis B,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s^B\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

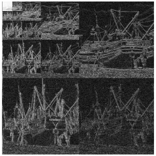

Results in compressed sensing [CRT '06, etc.] imply:

- if $f \in \mathbb{R}^{N \times N}$ is s-sparse in an orthonormal basis B
- if the measurement operator satisfies the RIP
- then using ℓ_1-minimization with basis B,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s^B\|_1}{\sqrt{s}} + \varepsilon$$

Imaging via compressed sensing

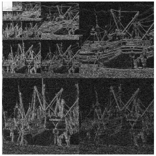

Results in compressed sensing [CRT '06, etc.] imply:

- if $f \in \mathbb{R}^{N \times N}$ is s-sparse in an orthonormal basis B
- if the measurement operator satisfies the RIP
- then using ℓ_1-minimization with basis B,

$$\|f - \hat{f}\|_2 \lesssim \frac{\|f - f_s^B\|_1}{\sqrt{s}} + \varepsilon$$

Other notions of sparsity for images

256×256 "Boats" image

Natural images

Images are compressible in the *discrete gradient*.

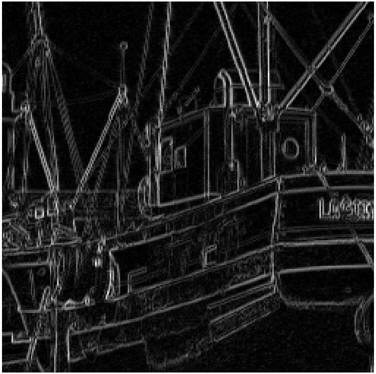

Natural images

Images are compressible in the *discrete gradient*.

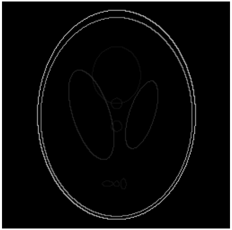

The discrete directional derivatives of an image $f \in \mathbb{R}^{N \times N}$ are

$$f_x : \mathbb{R}^{N \times N} \to \mathbb{R}^{(N-1) \times N}, \qquad (f_x)_{j,k} = f_{j,k} - f_{j-1,k},$$

$$f_y : \mathbb{R}^{N \times N} \to \mathbb{R}^{N \times (N-1)}, \qquad (f_y)_{j,k} = f_{j,k} - f_{j,k-1},$$

the discrete gradient or *total variation* operator is

$$\nabla[f] = (f_x, f_y)$$

Natural images

Images are compressible in the *discrete gradient*.

 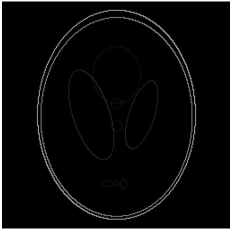

The discrete directional derivatives of an image $f \in \mathbb{R}^{N \times N}$ are

$$f_x : \mathbb{R}^{N\times N} \to \mathbb{R}^{(N-1)\times N}, \qquad (f_x)_{j,k} = f_{j,k} - f_{j-1,k},$$
$$f_y : \mathbb{R}^{N\times N} \to \mathbb{R}^{N\times (N-1)}, \qquad (f_y)_{j,k} = f_{j,k} - f_{j,k-1},$$

the discrete gradient or *total variation* operator is

$$\nabla[f] = (f_x, f_y)$$

Natural Notation

Images are compressible in *discrete gradient*.


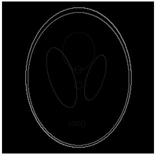

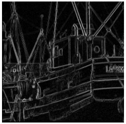

- "Phantom": $\|\nabla[f]\|_0 = .03N^2$
- "Boats": $\|\nabla[f] - \nabla[f]_s\|_2$ decays quickly in s

Natural Notation

Images are compressible in *discrete gradient*.

 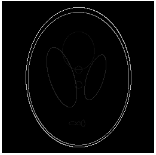  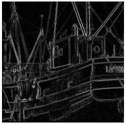

- "Phantom": $\|\nabla[f]\|_0 = .03N^2$
- "Boats": $\|\nabla[f] - \nabla[f]_s\|_2$ decays quickly in s

Natural Notation

Images are compressible in *discrete gradient*.


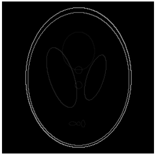

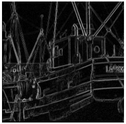

- "Phantom": $\|\nabla[f]\|_0 = .03N^2$
- "Boats": $\|\nabla[f] - \nabla[f]_s\|_2$ decays quickly in s

Total Variation Image Recovery

Comparison of two compressed sensing reconstruction algorithms

Haar-minimization (L_1-Haar)

$$\hat{f}_{Haar} = \text{argmin}_Z \|H(Z)\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon$$

Total Variation minimization (TV)

$$\hat{f}_{TV} = \text{argmin}_Z \|\nabla[Z]\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon,$$ where

$\|Z\|_{TV} = \|\nabla[Z]\|_1$ is the *total-variation norm*.

The mapping $Z \to \nabla[Z]$ is not orthonormal, stable image recovery via (TV) is not mathematically justified!

Comparison of two compressed sensing reconstruction algorithms

Haar-minimization (L_1-Haar)

$\hat{f}_{Haar} = \text{argmin}_Z \|H(Z)\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon$

Total Variation minimization (TV)

$\hat{f}_{TV} = \text{argmin}_Z \|\nabla[Z]\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon$, where

$\|Z\|_{TV} = \|\nabla[Z]\|_1$ is the *total-variation norm*.

The mapping $Z \to \nabla[Z]$ is not orthonormal, stable image recovery via (TV) is not mathematically justified!

Comparison of two compressed sensing reconstruction algorithms

Haar-minimization (L_1-Haar)

$$\hat{f}_{Haar} = \text{argmin}_Z \|H(Z)\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon$$

Total Variation minimization (TV)

$$\hat{f}_{TV} = \text{argmin}_Z \|\nabla[Z]\|_1 \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon,$$ where

$\|Z\|_{TV} = \|\nabla[Z]\|_1$ is the *total-variation norm*.

The mapping $Z \to \nabla[Z]$ is not orthonormal, stable image recovery via (TV) is not mathematically justified!

Imaging via compressed sensing

(a) Original

(b) TV

(c) L_1-Haar

Figure: Reconstruction using $m = .2N^2$

Imaging via compressed sensing

(a) Original

(b) TV

(c) L_1-Haar

Figure: Reconstruction using $m = .2N^2$ measurements

Imaging via compressed sensing

(a) Original

(b) TV

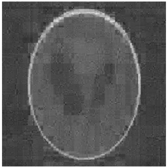

(c) L_1-Haar

Figure: Reconstruction using $m = .2N^2$ measurements

Imaging via compressed sensing

(a) (Quantization)

(b) TV

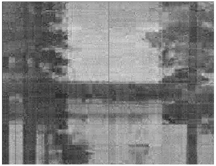

(c) L_1-Haar

Figure: Reconstruction using $m = .2N^2$ measurements

Imaging via compressed sensing

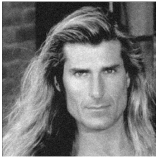

(a) (Gaussian)

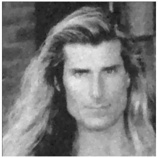

(b) TV

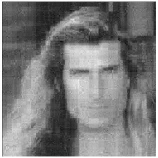

(c) L_1-Haar

Figure: Reconstruction using $m = .2N^2$ measurements

Imaging via compressed sensing

InView (Austin TX)

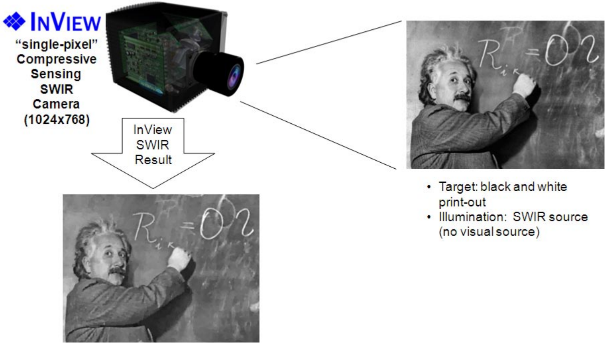

Figure: SWIR Reconstruction using $m = .5N^2$ measurements

Pediatric MRI

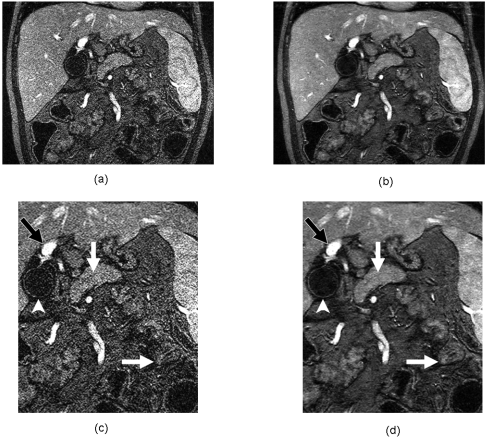

(a–d) Submillimeter near-isotropic–resolution contrast-enhanced T1-weighted MR images in 8-year-old boy. (a, c) Standard and (b, d) compressed sensing reconstruction images. (c, d) Zoomed images show improved delineation of the pancreatic duct (vertical arrow), bowel (horizontal arrow), and gallbladder wall (arrowhead), and equivalent definition of portal vein (black arrow) with L1 SPIR-iT reconstruction.

(Caffey Award : Faster Pediatric MRI Via Compressed Sensing - Shreyas Vasanawala et.al. (Stanford University))

Empirical → Theoretical?

TV Works

Empirically, it has been well known that

$$\hat{f}_{TV} = \text{argmin} \|Z\|_{TV} \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon, \qquad (TV)$$

provides quality, stable image recovery.

No provable stability guarantees.

Empirical → Theoretical?

TV Works

Empirically, it has been well known that

$$\hat{f}_{TV} = \text{argmin} \|Z\|_{TV} \quad \text{subject to} \quad \|AZ - y\|_2 \leq \varepsilon, \qquad (TV)$$

provides quality, stable image recovery.

No provable stability guarantees.

Stable signal recovery using total-variation minimization

Theorem (N-Ward '12)

From $m \gtrsim s\log(N)$ linear RIP measurements, for any $f \in \mathbb{C}^{N\times N}$,

$$\hat{f} = \operatorname{argmin} \|Z\|_{TV} \quad \textit{such that} \quad \|A(Z) - y\|_2 \leq \varepsilon,$$

satisfies

$$\|f - \hat{f}\|_{TV} \lesssim \|\nabla[f] - \nabla[f]_s\|_1 + \sqrt{s}\varepsilon \qquad \text{(gradient error)}$$

and

$$\|f - \hat{f}\|_2 \lesssim \log(N) \cdot \left[\frac{\|\nabla[f] - \nabla[f]_s\|_1}{\sqrt{s}} + \varepsilon\right] \qquad \text{(signal error)}$$

This error guarantee is optimal up to the $\log(N)$ factor

Stable signal recovery using total-variation minimization

Theorem (N-Ward '12)

From $m \gtrsim s\log(N)$ linear RIP measurements, for any $f \in \mathbb{C}^{N\times N}$,

$$\hat{f} = \text{argmin} \|Z\|_{TV} \quad \textit{such that} \quad \|A(Z) - y\|_2 \leq \varepsilon,$$

satisfies

$$\|f - \hat{f}\|_{TV} \lesssim \|\nabla[f] - \nabla[f]_s\|_1 + \sqrt{s}\varepsilon \qquad \text{(gradient error)}$$

and

$$\|f - \hat{f}\|_2 \lesssim \log(N) \cdot \left[\frac{\|\nabla[f] - \nabla[f]_s\|_1}{\sqrt{s}} + \varepsilon\right] \qquad \text{(signal error)}$$

This error guarantee is optimal up to the $\log(N)$ factor

Stable signal recovery using total-variation minimization

Theorem (N-Ward '12)

From $m \gtrsim s\log(N)$ linear RIP measurements, for any $f \in \mathbb{C}^{N\times N}$,

$$\hat{f} = \operatorname{argmin} \|Z\|_{TV} \quad \textit{such that} \quad \|A(Z) - y\|_2 \leq \varepsilon,$$

satisfies

$$\|f - \hat{f}\|_{TV} \lesssim \|\nabla[f] - \nabla[f]_s\|_1 + \sqrt{s}\varepsilon \qquad \text{(gradient error)}$$

and

$$\|f - \hat{f}\|_2 \lesssim \log(N) \cdot \left[\frac{\|\nabla[f] - \nabla[f]_s\|_1}{\sqrt{s}} + \varepsilon\right] \qquad \text{(signal error)}$$

This error guarantee is optimal up to the $\log(N)$ factor

Stable signal recovery using total-variation minimization

Theorem (N-Ward '12)

From $m \gtrsim s\log(N)$ linear RIP measurements, for any $f \in \mathbb{C}^{N\times N}$,

$$\hat{f} = \operatorname{argmin} \|Z\|_{TV} \quad \textit{such that} \quad \|A(Z) - y\|_2 \le \varepsilon,$$

satisfies

$$\|f - \hat{f}\|_{TV} \lesssim \|\nabla[f] - \nabla[f]_s\|_1 + \sqrt{s}\varepsilon \qquad \text{(gradient error)}$$

and

$$\|f - \hat{f}\|_2 \lesssim \log(N) \cdot \left[\frac{\|\nabla[f] - \nabla[f]_s\|_1}{\sqrt{s}} + \varepsilon\right] \qquad \text{(signal error)}$$

This error guarantee is optimal up to the $\log(N)$ factor

Stable signal recovery using total-variation minimization

Theorem (N-Ward '12)

From $m \gtrsim s\log(N)$ linear RIP measurements, for any $f \in \mathbb{C}^{N\times N}$,

$$\hat{f} = \operatorname{argmin} \|Z\|_{TV} \quad \text{such that} \quad \|A(Z) - y\|_2 \leq \varepsilon,$$

satisfies

$$\|f - \hat{f}\|_{TV} \lesssim \|\nabla[f] - \nabla[f]_s\|_1 + \sqrt{s}\varepsilon \qquad \text{(gradient error)}$$

and

$$\|f - \hat{f}\|_2 \lesssim \log(N) \cdot \left[\frac{\|\nabla[f] - \nabla[f]_s\|_1}{\sqrt{s}} + \varepsilon\right] \qquad \text{(signal error)}$$

This error guarantee is optimal up to the $\log(N)$ factor

Stable signal recovery using total-variation minimization

Method of proof:

1. First prove stable *gradient* recovery
2. Translate stable *gradient* recovery to stable *signal* recovery using a (nontrivial) Sobolev inequality which shows that Haar coefficients of functions of bounded variation are in weak-ℓ_1 space.

Stable signal recovery using total-variation minimization

Method of proof:

1. First prove stable *gradient* recovery
2. Translate stable *gradient* recovery to stable *signal* recovery using a (nontrivial) Sobolev inequality which shows that Haar coefficients of functions of bounded variation are in weak-ℓ_1 space.

Open questions

1. Remove the log factor?
2. The relationship between Haar compressibiity and total variation norm doesn't hold in one-dimension. What about stable (1D) signal recovery?
3. [Patel, Maleh, Gilbert, Chellappa '11] Images are even sparser in individual directional derivatives f_x, f_y. If we minimize separately over directional derivatives, can we still prove stable recovery?
4. Extend results to higher dimensions...movies? [In preparation.]

Open questions

1. Remove the log factor?
2. The relationship between Haar compressibiity and total variation norm doesn't hold in one-dimension. What about stable (1D) signal recovery?
3. [Patel, Maleh, Gilbert, Chellappa '11] Images are even sparser in individual directional derivatives f_x, f_y. If we minimize separately over directional derivatives, can we still prove stable recovery?
4. Extend results to higher dimensions...movies? [In preparation.]

Open questions

1. Remove the log factor?
2. The relationship between Haar compressibiity and total variation norm doesn't hold in one-dimension. What about stable (1D) signal recovery?
3. [Patel, Maleh, Gilbert, Chellappa '11] Images are even sparser in individual directional derivatives f_x, f_y. If we minimize separately over directional derivatives, can we still prove stable recovery?
4. Extend results to higher dimensions...movies? [In preparation.]

Open questions

1. Remove the log factor?
2. The relationship between Haar compressibiity and total variation norm doesn't hold in one-dimension. What about stable (1D) signal recovery?
3. [Patel, Maleh, Gilbert, Chellappa '11] Images are even sparser in individual directional derivatives f_x, f_y. If we minimize separately over directional derivatives, can we still prove stable recovery?
4. Extend results to higher dimensions...movies? [In preparation.]

Movies are very sparse!

- Movies are very sparse in all three dimensions
- Silicon Retina (Institute of Neuroinformatics) is the design of a camera that mimics retinas
  - Explanation of how the brain and eye communicate?
  - Really cool video cameras! (lower cost, lower power consumption, portable, continuous processing, real-time data acquisition)
  - Dynamic Vision Sensor (DVS) from Silicon Retina, Institute of Neuroinformatics

Movies are very sparse!

- Movies are very sparse in all three dimensions
- Silicon Retina (Institute of Neuroinformatics) is the design of a camera that mimics retinas
  - Explanation of how the brain and eye communicate?
  - Really cool video cameras! (lower cost, lower power consumption, portable, continuous processing, real-time data acquisition)
  - Dynamic Vision Sensor (DVS) from Silicon Retina, Institute of Neuroinformatics

Movies are very sparse!

- Movies are very sparse in all three dimensions
- Silicon Retina (Institute of Neuroinformatics) is the design of a camera that mimics retinas
  - Explanation of how the brain and eye communicate?
  - Really cool video cameras! (lower cost, lower power consumption, portable, continuous processing, real-time data acquisition)
  - Dynamic Vision Sensor (DVS) from Silicon Retina, Institute of Neuroinformatics

Movies are very sparse!

- Movies are very sparse in all three dimensions
- Silicon Retina (Institute of Neuroinformatics) is the design of a camera that mimics retinas
  - Explanation of how the brain and eye communicate?
  - Really cool video cameras! (lower cost, lower power consumption, portable, continuous processing, real-time data acquisition)
  - Dynamic Vision Sensor (DVS) from Silicon Retina, Institute of Neuroinformatics

Movies are very sparse!

- Movies are very sparse in all three dimensions
- Silicon Retina (Institute of Neuroinformatics) is the design of a camera that mimics retinas
  - Explanation of how the brain and eye communicate?
  - Really cool video cameras! (lower cost, lower power consumption, portable, continuous processing, real-time data acquisition)
  - Dynamic Vision Sensor (DVS) from Silicon Retina, Institute of Neuroinformatics

Fast vision in bad lighting

Figure: ("RoboGoalie", Silicon Retina, Institute of Neuroinformatics)

Fluid Particle Tracking Velocimetry

Figure: (“PTV”, Silicon Retina, Institute of Neuroinformatics)

Mobile Robotics

Figure: ("Robotic Driver", Silicon Retina, Institute of Neuroinformatics)

Sleep disorder research

Figure: ("Sleeping Mouse", Silicon Retina, Institute of Neuroinformatics)

Thank you!

E-mail:

- dneedell@cmc.edu

Web:

- www.cmc.edu/pages/faculty/DNeedell

References:

- E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. Communications on Pure and Applied Mathematics, 59(8):12071223, 2006.
- E. J. Candès, Y. C. Eldar, D. Needell and P. Randall. Compressed sensing with coherent and redundant dictionaries. Applied and Computational Harmonic Analysis, 31(1):59-73, 2010.
- V. Patel, R. Maleh, A. Gilbert, and R. Chellappa. Gradient-based image recovery methods from incomplete Fourier measurements, IEEE T. Image Process., 21(1), 2012.
- D. Needell and R. Ward. Stable image reconstruction using total variation minimization. Submitted.