Compressive Sampling and Redundancy

Deanna Needell

Stanford University

MIT, March 14, 2011

Digital cameras, for example

(Hard work to measure all pixels) (Compression) (Wasteful?)

Digital cameras, for example

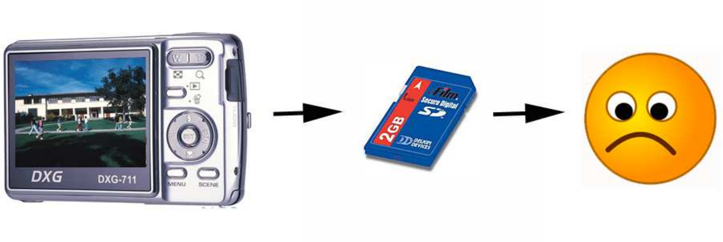

(Hard work to measure all pixels) (Compression) (Wasteful?)

Introducing, CoSa

The issue

Traditional data acquisition is wasteful.

The idea

Combine acquisition and compression.

The solution

Compressive sampling (CoSa) allows us to do this → reconstruct a signal from its (compressed) measurements.

Introducing, CoSa

The issue

Traditional data acquisition is wasteful.

The idea

Combine acquisition and compression.

The solution

Compressive sampling (CoSa) allows us to do this $\rightarrow$ reconstruct a signal from its (compressed) measurements.

Introducing, CoSa

The issue

Traditional data acquisition is wasteful.

The idea

Combine acquisition and compression.

The solution

Compressive sampling (CoSa) allows us to do this $\rightarrow$ reconstruct a signal from its (compressed) measurements.

Numerous Applications

- CoSa single pixel digital camera [Rice]
- Medical Imaging, MRI
- Radar
- Error Correction
- Computational Biology (DNA Microarrays)
- Geophysical Data Analysis
- Data Mining, classification
- ...

Introducing, CoSa

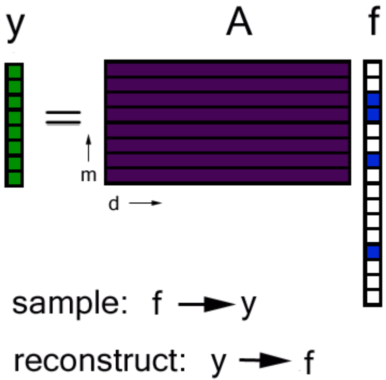

Wait, isn't this impossible?

Without further assumptions, this problem is ill-posed.

Why will this work?

Most signals of interest contain far less information than their dimension d suggests.

Assume f is sparse:

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

In practice, we encounter compressible signals.

Wait, isn't this impossible?

Without further assumptions, this problem is ill-posed.

Why will this work?

Most signals of interest contain far less information than their dimension d suggests.

Assume f is sparse:

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

In practice, we encounter compressible signals.

Wait, isn't this impossible?

Without further assumptions, this problem is ill-posed.

Why will this work?

Most signals of interest contain far less information than their dimension d suggests.

Assume f is sparse:

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

In practice, we encounter compressible signals.

Wait, isn't this impossible?

Without further assumptions, this problem is ill-posed.

Why will this work?

Most signals of interest contain far less information than their dimension d suggests.

Assume f is sparse:

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

In practice, we encounter compressible signals.

CoSa in the real world

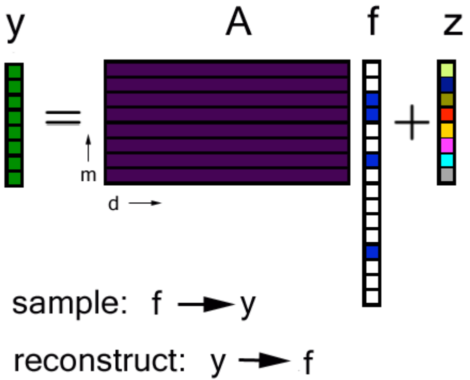

Designing a reconstruction algorithm

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

Important Questions

- What kind(s) of measurement matrices A?
- How many measurements needed?
- Are the guarantees *uniform*?
- Is algorithm *stable*?
- Fast runtime?

Orthogonal Matching Pursuit (OMP)

Idea

Noiseless case: When A is (sub)Gaussian, $A^*y = A^*Af$ is a good approximation to f.

Initialize: Set $I = \emptyset$ and $r = y$.
Repeat the following s times:

Identify: Select the largest coordinate λ of $u = A^*r$ in absolute value.

Update: Add the coordinate: $I \leftarrow I \cup \{\lambda\}$, and update the residual:

$$\hat{x} = \operatorname*{argmin}_z \|y - A|_I z\|_2; \qquad r = y - A\hat{x}.$$

Orthogonal Matching Pursuit (OMP)

Idea

Noiseless case: When A is (sub)Gaussian, $A^*y = A^*Af$ is a good approximation to f.

Initialize: Set $I = \emptyset$ and $r = y$.
Repeat the following s times:

Identify: Select the largest coordinate λ of $u = A^*r$ in absolute value.

Update: Add the coordinate: $I \leftarrow I \cup \{\lambda\}$, and update the residual:

$$\hat{x} = \operatorname*{argmin}_z \|y - A|_I z\|_2; \qquad r = y - A\hat{x}.$$

Orthogonal Matching Pursuit (OMP)

Idea

Noiseless case: When A is (sub)Gaussian, $A^*y = A^*Af$ is a good approximation to f.

Initialize: Set $I = \emptyset$ and $r = y$.
Repeat the following s times:

Identify: Select the largest coordinate λ of $u = A^*r$ in absolute value.

Update: Add the coordinate: $I \leftarrow I \cup \{\lambda\}$, and update the residual:

$$\hat{x} = \underset{z}{\operatorname{argmin}} \|y - A|_I z\|_2; \qquad r = y - A\hat{x}.$$

Orthogonal Matching Pursuit (OMP)

Idea

Noiseless case: When A is (sub)Gaussian, $A^*y = A^*Af$ is a good approximation to f.

Initialize: Set $I = \emptyset$ and $r = y$.
Repeat the following s times:

Identify: Select the largest coordinate λ of $u = A^*r$ in absolute value.

Update: Add the coordinate: $I \leftarrow I \cup \{\lambda\}$, and update the residual:

$$\hat{x} = \operatorname*{argmin}_z \|y - A|_I z\|_2; \qquad r = y - A\hat{x}.$$

Compressive Sampling Matching Pursuit [N-Tropp]

Theorem (Gilbert-Tropp)

When A is (sub)Gaussian with $m \gtrsim s \log d$, OMP correctly recovers each fixed s-sparse signal with high probability.

Compressive Sampling Matching Pursuit [N-Tropp]

Answers for OMP

- What kind(s) of measurement matrices A? - Gaussian ✗
- How many measurements needed? - $s \log d$ ✓
- Are the guarantees *uniform*? - No, for fixed signal ✗
- Is algorithm *stable*? - Not known to be robust to noise ✗
- Fast runtime? - Yes ✓

ℓ_0-optimization

The First CoSa Theorem

Let A be one-to-one on $2s$-sparse vectors and set:

$$\hat{f} = \text{argmin}\,\|g\|_0 \quad \text{such that} \quad Ag = y.$$

Then in the noiseless case, we have perfect recovery of all s-sparse signals: $\hat{f} = f$.

End of story?

This is numerically infeasible!

ℓ_0-optimization

The First CoSa Theorem

Let A be one-to-one on $2s$-sparse vectors and set:

$$\hat{f} = \operatorname{argmin} \|g\|_0 \quad \text{such that} \quad Ag = y.$$

Then in the noiseless case, we have perfect recovery of all s-sparse signals: $\hat{f} = f$.

End of story?

This is numerically infeasible!

Just relax: ℓ_1-optimization

Relaxation [Donoho et.al., Candès-Tao]

Let A satisfy the *Restricted Isometry Property* for $2s$-sparse vectors and set:

$$\hat{f} = \operatorname{argmin} \|g\|_1 \quad \text{such that} \quad Ag = y.$$

Then in the noiseless case, we have perfect recovery of all s-sparse signals: $\hat{f} = f$.

End of story?

This is numerically feasible!

Just relax: ℓ_1-optimization

Relaxation [Donoho et.al., Candès-Tao]

Let A satisfy the *Restricted Isometry Property* for $2s$-sparse vectors and set:

$$\hat{f} = \operatorname{argmin} \|g\|_1 \quad \text{such that} \quad Ag = y.$$

Then in the noiseless case, we have perfect recovery of all s-sparse signals: $\hat{f} = f$.

End of story?

This is numerically feasible!

Let's talk geometry...

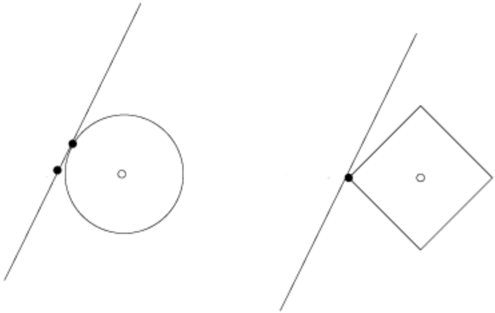

Make some noise: ℓ_1-optimization

The ℓ_1-optimization method succeeds for *arbitrary* vectors with *noisy* samples.

Stability [Candès-Romberg-Tao]

Let A satisfy the Restricted Isometry Property as before and set:

$$\hat{f} = \operatorname{argmin} \|g\|_1 \quad \text{such that} \quad \|Ag - y\|_2 \le \varepsilon.$$

Then we have optimal recovery error:

$$\|\hat{f} - f\|_2 \le C\left(\varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}\right).$$

Make some noise: ℓ_1-optimization

The ℓ_1-optimization method succeeds for *arbitrary* vectors with *noisy* samples.

Stability [Candès-Romberg-Tao]

Let A satisfy the Restricted Isometry Property as before and set:

$$\hat{f} = \operatorname{argmin} \|g\|_1 \quad \text{such that} \quad \|Ag - y\|_2 \leq \varepsilon.$$

Then we have optimal recovery error:

$$\|\hat{f} - f\|_2 \leq C\left(\varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}\right).$$

Restricted Isometry Property

- The sth restricted isometry constant of $(m \times d)$ A is the smallest δ_s such that

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

- For Gaussian or Bernoulli measurement matrices, with high probability

$$\delta_s \leq c < 1 \quad \text{when } m \gtrsim s \log d.$$

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

Restricted Isometry Property

- The s^{th} restricted isometry constant of $(m \times d)$ A is the smallest δ_s such that

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

- For Gaussian or Bernoulli measurement matrices, with high probability

$$\delta_s \leq c < 1 \quad \text{when } m \gtrsim s \log d.$$

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

Restricted Isometry Property

- The sth restricted isometry constant of $(m \times d)$ A is the smallest δ_s such that

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

- For Gaussian or Bernoulli measurement matrices, with high probability
$$\delta_s \leq c < 1 \quad \text{when } m \gtrsim s \log d.$$

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

ℓ_1-optimization

Answers for ℓ_1-optimization

- What kind(s) of measurement matrices A? - Many (RIP)
- How many measurements needed? - $s \log d$
- Are the guarantees *uniform*? - Yes, the RIP holds
- Is algorithm *stable*? - Yes, optimal error bounds
- Fast runtime? - Not bad, but not ideal...

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^\dagger y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^{\dagger}y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^{\dagger} y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^\dagger y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^\dagger y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit (CoSaMP)

Idea

- When A satisfies the RIP, $A^*y = A^*Af$ is a good approximation to f.
- At each iteration, select many components of A^*y to be in support. Estimate f, and prune.

- Initialize: Set $a^0 = 0$, $v = y$, $k = 0$. Repeat the following steps:
- Signal Proxy: Set $u = A^*v$, $\Omega = \operatorname{supp} u_{2s}$ and merge: $T = \Omega \cup \operatorname{supp} a^{k-1}$.
- Signal Estimation: Set $b|_T = A_T^\dagger y$ and $b|_{T^c} = 0$.
- Prune: Set $a^k = b_s$.
- Sample Update: Update $v = y - Aa^k$.

Compressive Sampling Matching Pursuit

Theorem (N-Tropp)

When A satisfies the RIP, CoSaMP recovers an approximation to f with optimal error:

$$\|\hat{f} - f\|_2 \leq C\left(\varepsilon + \frac{\|f - f_s\|_1}{\sqrt{s}}\right).$$

Compressive Sampling Matching Pursuit [N-Tropp]

Answers for CoSaMP

- What kind(s) of measurement matrices A? - Many (RIP) ✓
- How many measurements needed? - $s \log d$ ✓
- Are the guarantees *uniform*? - Yes, the RIP holds ✓
- Is algorithm *stable*? - Yes, optimal error bounds ✓
- Fast runtime? - Yes, roughly same cost as applying A ✓

The News

Good News

Many methods hold for signals f which are sparse in the coordinate basis or in some other orthonormal basis (ONB).

Bad News

There are many applications for which the signal f is sparse not in an ONB, but in some overcomplete dictionary! This means that $f = Dx$ where D is a redundant dictionary. When D is not an ONB, AD is not at all likely to satisfy the RIP (or be incoherent).

The News

Good News

Many methods hold for signals f which are sparse in the coordinate basis or in some other orthonormal basis (ONB).

Bad News

There are many applications for which the signal f is sparse not in an ONB, but in some overcomplete dictionary! This means that $f = Dx$ where D is a redundant dictionary. When D is not an ONB, AD is not at all likely to satisfy the RIP (or be incoherent).

The News

Good News

Many methods hold for signals f which are sparse in the coordinate basis or in some other orthonormal basis (ONB).

Bad News

There are many applications for which the signal f is sparse not in an ONB, but in some overcomplete dictionary! This means that $f = Dx$ where D is a redundant dictionary. When D is not an ONB, AD is not at all likely to satisfy the RIP (or be incoherent).

Big Picture

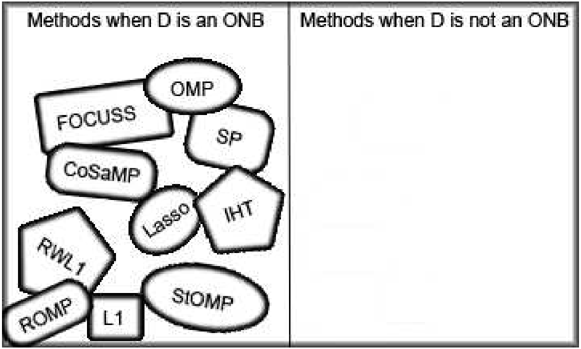
Methods when D is an ONB
Methods when D is not an ONB
OMP
FOCUSS
SP
CoSaMP
Lasso
IHT
RWL1
ROMP
L1
StOMP

Example: Oversampled DFT

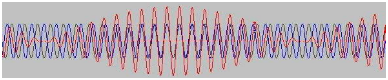

- $n \times n$ DFT: $d_k(t) = \frac{1}{\sqrt{n}} e^{-2\pi ikt/n}$
- Sparse in the DFT = superpositions of sinusoids with frequencies in the lattice.
- Instead, use the oversampled DFT: frequencies may be over even smaller intervals or intervals of varying length.
- Then D is an overcomplete frame with highly coherent columns $\rightarrow$ current CS does not apply.

Example: Oversampled DFT

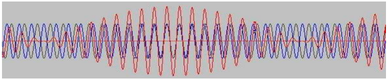

- $n \times n$ DFT: $d_k(t) = \frac{1}{\sqrt{n}} e^{-2\pi i k t/n}$
- Sparse in the DFT $=$ superpositions of sinusoids with frequencies in the lattice.
- Instead, use the oversampled DFT: frequencies may be over even smaller intervals or intervals of varying length.
- Then D is an overcomplete frame with highly coherent columns $\rightarrow$ current CS does not apply.

Example: Oversampled DFT

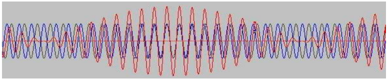

- $n \times n$ DFT: $d_k(t) = \frac{1}{\sqrt{n}} e^{-2\pi i k t/n}$
- Sparse in the DFT = superpositions of sinusoids with frequencies in the lattice.
- Instead, use the oversampled DFT: frequencies may be over even smaller intervals or intervals of varying length.
- Then D is an overcomplete frame with highly coherent columns → current CS does not apply.

Example: Oversampled DFT

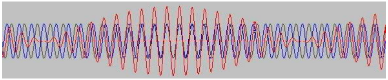

- $n \times n$ DFT: $d_k(t) = \frac{1}{\sqrt{n}} e^{-2\pi i k t/n}$
- Sparse in the DFT = superpositions of sinusoids with frequencies in the lattice.
- Instead, use the oversampled DFT: frequencies may be over even smaller intervals or intervals of varying length.
- Then D is an overcomplete frame with highly coherent columns $\rightarrow$ current CS does not apply.

Example: Gabor frames

- Gabor frame: $G_k(t) = g(t - k_2 a)e^{2\pi i k_1 bt}$
- Radar, sonar, and imaging system applications use Gabor frames and wish to recover signals in this basis.
- Then D is an overcomplete frame with possibly highly coherent columns $\rightarrow$ current CS does not apply.

Example: Gabor frames

- Gabor frame: $G_k(t) = g(t - k_2 a)e^{2\pi i k_1 bt}$
- Radar, sonar, and imaging system applications use Gabor frames and wish to recover signals in this basis.
- Then D is an overcomplete frame with possibly highly coherent columns $\rightarrow$ current CS does not apply.

Example: Gabor frames

- Gabor frame: $G_k(t) = g(t - k_2 a)e^{2\pi i k_1 bt}$
- Radar, sonar, and imaging system applications use Gabor frames and wish to recover signals in this basis.
- Then D is an overcomplete frame with possibly highly coherent columns $\rightarrow$ current CS does not apply.

Example: Curvelet frames

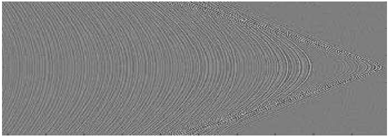

- A Curvelet frame has some properties of an ONB but is overcomplete.
- Curvelets approximate well the curved singularities in images and are thus used widely in image processing.
- Again, this means D is an overcomplete dictionary $\rightarrow$ current CS does not apply.

Example: Curvelet frames

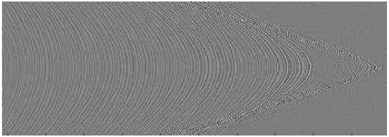

- A Curvelet frame has some properties of an ONB but is overcomplete.
- Curvelets approximate well the curved singularities in images and are thus used widely in image processing.
- Again, this means D is an overcomplete dictionary $\rightarrow$ current CS does not apply.

Example: Curvelet frames

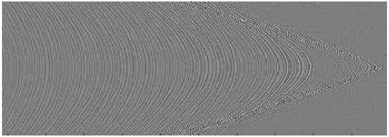

- A Curvelet frame has some properties of an ONB but is overcomplete.
- Curvelets approximate well the curved singularities in images and are thus used widely in image processing.
- Again, this means D is an overcomplete dictionary $\rightarrow$ current CS does not apply.

Example: UWT

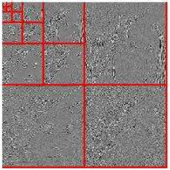

- The undecimated wavelet transform has a translation invariance property that is missing in the DWT.
- The UWT is overcomplete and this redundancy has been found to be helpful in image processing.
- Again, this means D is a redundant dictionary $\rightarrow$ current CS does not apply.

Example: UWT

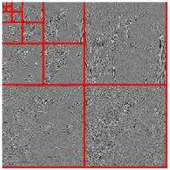

- The undecimated wavelet transform has a translation invariance property that is missing in the DWT.
- The UWT is overcomplete and this redundancy has been found to be helpful in image processing.
- Again, this means D is a redundant dictionary $\rightarrow$ current CS does not apply.

Example: UWT

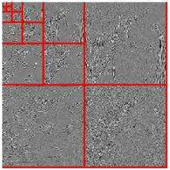

- The undecimated wavelet transform has a translation invariance property that is missing in the DWT.
- The UWT is overcomplete and this redundancy has been found to be helpful in image processing.
- Again, this means D is a redundant dictionary $\rightarrow$ current CS does not apply.

Example: Concatenations

$$\begin{bmatrix} I & F \end{bmatrix}$$

- In many applications, a signal may be sparse in several ONBs.
- Correlations between the bases mean current CS techniques do not apply.

Example: Concatenations

$$\begin{bmatrix} I & F \end{bmatrix}$$

- In many applications, a signal may be sparse in several ONBs.
- Correlations between the bases mean current CS techniques do not apply.

ℓ_1-Analysis

Proposed Method

It has been observed (empirically) that ℓ_1-analysis often succeeds:

$$\hat{f} = \underset{\tilde{f} \in \mathbb{R}^n}{\text{argmin}} \|D^* \tilde{f}\|_1 \quad \text{subject to} \quad \|A\tilde{f} - y\|_2 \leq \varepsilon.$$

Condition on A?

Let Σ_s be the union of all subspaces spanned by all subsets of s columns of D.

D-RIP

We say that the measurement matrix A obeys the *restricted isometry property adapted to D* (D-RIP) with constant δ_s if

$$(1-\delta_s)\|v\|_2^2 \le \|Av\|_2^2 \le (1+\delta_s)\|v\|_2^2$$

holds for all $v \in \Sigma_s$.

Similarly to the RIP, Gaussian, subgaussian, and Bernoulli matrices satisfy the D-RIP with $m \approx s\log(d/s)$. Matrices with a fast multiply (DFT with random signs) also satisfy the D-RIP with m approximately of this order.

Condition on A?

Let Σ_s be the union of all subspaces spanned by all subsets of s columns of D.

D-RIP

We say that the measurement matrix A obeys the *restricted isometry property adapted to D* (D-RIP) with constant δ_s if

$$(1-\delta_s)\|v\|_2^2 \leq \|Av\|_2^2 \leq (1+\delta_s)\|v\|_2^2$$

holds for all $v \in \Sigma_s$.

Similarly to the RIP, Gaussian, subgaussian, and Bernoulli matrices satisfy the D-RIP with $m \approx s\log(d/s)$. Matrices with a fast multiply (DFT with random signs) also satisfy the D-RIP with m approximately of this order.

Condition on A?

Let Σ_s be the union of all subspaces spanned by all subsets of s columns of D.

D-RIP

We say that the measurement matrix A obeys the *restricted isometry property adapted to D* (D-RIP) with constant δ_s if

$$(1-\delta_s)\|v\|_2^2 \leq \|Av\|_2^2 \leq (1+\delta_s)\|v\|_2^2$$

holds for all $v \in \Sigma_s$.

Similarly to the RIP, Gaussian, subgaussian, and Bernoulli matrices satisfy the D-RIP with $m \approx s\log(d/s)$. Matrices with a fast multiply (DFT with random signs) also satisfy the D-RIP with m approximately of this order.

Main Result

Theorem (Candès-Eldar-N-Randall)

Let D be an arbitrary tight frame and let A be a measurement matrix satisfying D-RIP (with δ_{2s} small). Then the solution $\hat{f}$ to ℓ_1-analysis satisfies

$$\|\hat{f} - f\|_2 \leq C_0 \varepsilon + C_1 \frac{\|D^* f - (D^* f)_s\|_1}{\sqrt{s}},$$

where the constants C_0 and C_1 may only depend on δ_{2s}.

Implications

In other words,

Our result says that ℓ_1-analysis is very accurate when D^*f has rapidly decaying coefficients. This is the case in applications using the Oversampled DFT, Gabor frames, Undecimated WT, and Curvelet frames (and many others).

This will not necessarily be the case when using concatenations of two ONBs $\rightarrow$ ℓ_1-analysis not the right method.

Implications

In other words,

Our result says that ℓ_1-analysis is very accurate when D^*f has rapidly decaying coefficients. This is the case in applications using the Oversampled DFT, Gabor frames, Undecimated WT, and Curvelet frames (and many others).

This will not necessarily be the case when using concatenations of two ONBs $\rightarrow$ ℓ_1-analysis not the right method.

Implications

In other words,

Our result says that ℓ_1-analysis is very accurate when D^*f has rapidly decaying coefficients. This is the case in applications using the Oversampled DFT, Gabor frames, Undecimated WT, and Curvelet frames (and many others).

This will not necessarily be the case when using concatenations of two ONBs $\rightarrow$ ℓ_1-analysis not the right method.

Experimental Setup

$n = 8192, m = 400, d = 491,520$
A: $m \times n$ Gaussian, D: $n \times d$ Gabor

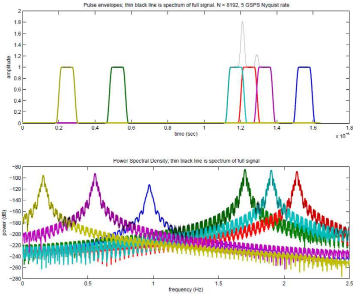

Figure: The signal is a superposition of 6 radar pulses, each of which being about 200 ns long, and with frequency carriers distributed between 50 MHz and 2.5 GHz (top plot). As can be seen, three of these pulses overlap in the time domain.

Experimental Results

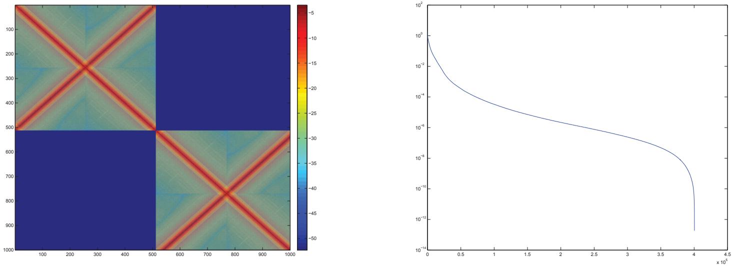

Figure: Portion of the matrix $D^* D$, in log-scale (left). Sorted analysis coefficients (in absolute value) of the signal from Figure 1 (right).

Experimental Results

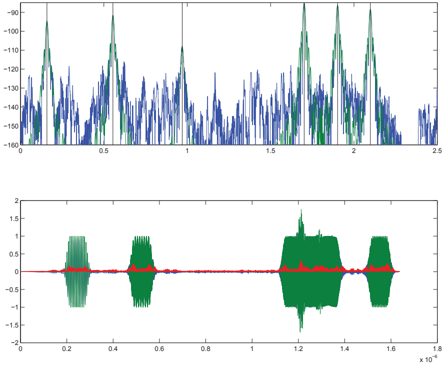

Figure: Recovery in both the time (below) and frequency (above) domains by ℓ_1-analysis. Blue denotes the recovered signal, green the actual signal, and red the difference between the two.

Experimental Results

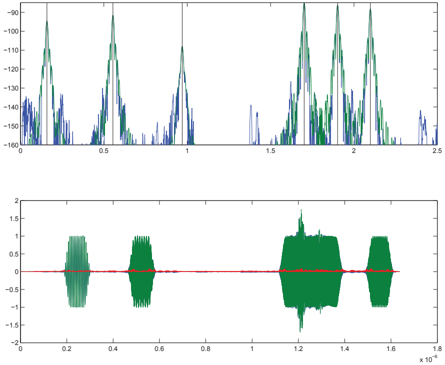

Figure: Recovery in both the time (below) and frequency (above) domains by ℓ_1-analysis after one reweighted iteration. Blue denotes the recovered signal, green the actual signal, and red the difference.

Experimental Results

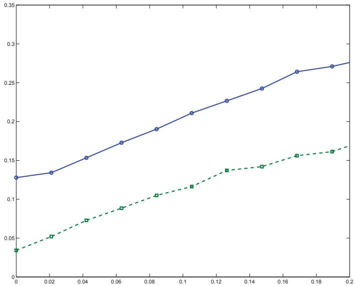

Figure: Relative recovery error of ℓ_1-analysis as a function of the (normalized) noise level, averaged over 5 trials. The solid line denotes standard ℓ_1-analysis, and the dashed line denotes ℓ_1-analysis with 3 reweighted iterations. The x-axis is the relative noise level $\sqrt{m}\sigma/\|Af\|_2$ while the y-axis is the relative error $\|\hat{f} - f\|_2/\|f\|_2$.

Experimental Results

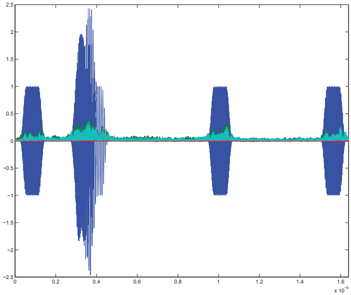

Figure: Relative error $\|\hat{f} - f\|_2 / \|f\|_2$ of a compressible signal. Blue denotes the actual signal, while green, red, and cyan denote the recovery error from ℓ_1-analysis, reweighted ℓ_1-analysis, and ℓ_1-synthesis, respectively.

Big Picture

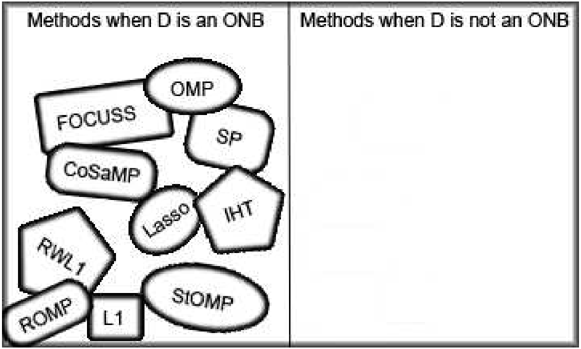
Methods when D is an ONB
Methods when D is not an ONB
OMP
FOCUSS
SP
CoSaMP
Lasso
IHT
RWL1
ROMP
L1
StOMP

Just recently

Candès-Edlar-N-Randall proved that a method called ℓ_1-analysis recovers signals sparse in arbitrary overcomplete dictionaries.

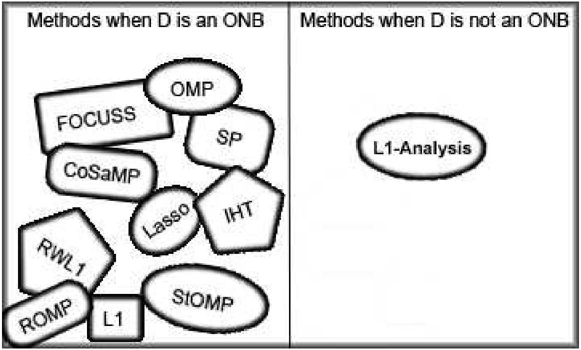

Future work

- Proof for another method for overcomplete dictionaries: ℓ_1-synthesis
- Is RIP the right theory?
- What about concatenations?

Future work

- Proof for another method for overcomplete dictionaries: ℓ_1-synthesis
- Is RIP the right theory?
- What about concatenations?

Future work

- Proof for another method for overcomplete dictionaries: ℓ_1-synthesis
- Is RIP the right theory?
- What about concatenations?

For more information

E-mail:

- dneedell@stanford.edu

Web:

- www-stat.stanford.edu/~dneedell

References:

- E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl.Math.*, 59(8):12071223, 2006.
- D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmonic Anal.*, 26(3):301-321, 2008.
- E. J. Candès, Y. C. Eldar, D. Needell, and P. Randall, "Compressed sensing with coherent and redundant dictionaries," *Appl. Comp. Harmonic Anal.*, to appear.