Non-convex compressed sensing using partial support information

Navid Ghadermarzy

University of British Columbia
SINDBAD consortium Meeting 2012

December 1, 2012

Collaborators

Joint work with:

- Hassan Mansour
- Özgür Yılmaz

Outline

- Introduction and overview
- Recovery by weighted ℓ_p

Outline

- Introduction and overview
- Recovery by weighted ℓ_p
- Experimental results

Outline

- Introduction and overview
- Recovery by weighted ℓ_p
- Experimental results
- Weighted ℓ_p minimization for seismic data interpolation

Outline

- Introduction and overview
- Recovery by weighted ℓ_p
- Experimental results
- Weighted ℓ_p minimization for seismic data interpolation

Notation

Consider a signal $z \in R^N$ s.t. $z = Dx$ where D is a transform matrix and x is a k-sparse vector.

We want to recover x, given n linear and noisy measurements $y = \Psi Dx + e$ where $n \ll N$ and $||e|| < \epsilon$.

Notation

Consider a signal $z \in R^N$ s.t. $z = Dx$ where D is a transform matrix and x is a k-sparse vector.

We want to recover x, given n linear and noisy measurements $y = \Psi Dx + e$ where $n \ll N$ and $||e|| < \epsilon$.

Assuming $A = \Psi D$ we should solve the underdetermined system $y = Ax + e$.

Notation

Consider a signal $z \in R^N$ s.t. $z = Dx$ where D is a transform matrix and x is a k-sparse vector.

We want to recover x, given n linear and noisy measurements $y = \Psi Dx + e$ where $n \ll N$ and $||e|| < \epsilon$.

Assuming $A = \Psi D$ we should solve the underdetermined system $y = Ax + e$.

Problem formulation

Theorem

The following optimization problem can approximately recover x from the measurements y if $k < \frac{n}{2}$ and A is in general position:

$$minimize_{z \in R^N} ||z||_0 \quad subject\ to \ \ ||Az - y||_2 \leq \epsilon$$

.

This problem is a combinatorial problem, so now different optimization problems is used to approximate x.

Problem formulation

Theorem

The following optimization problem can approximately recover x from the measurements y if $k < \frac{n}{2}$ and A is in general position:

$$minimize_{z \in R^N} ||z||_0 \quad subject\ to\ \ ||Az - y||_2 \leq \epsilon$$

.

This problem is a combinatorial problem, so now different optimization problems is used to approximate x.

Recovery by ℓ_1 minimization

Candes, Romberg and Tao showed that if A is sufficiently incoherent, solving the following convex optimization problem recovers x from measurements $y = Ax + e$:

$$\text{minimize}_{z \in R^N} ||z||_1 \quad \text{subject to} \quad ||Az - y||_2 \leq \epsilon.$$

Assuming x^* as the solution and x_k as the best k-term approximation of x, then:

$$||x^* - x||_2 \leq C_1^{\ell_1} \cdot \epsilon + C_2^{\ell_1} \cdot \frac{||x - x_k||_1}{\sqrt{k}}.$$

Recovery by ℓ_1 minimization

Candes, Romberg and Tao showed that if A is sufficiently incoherent, solving the following convex optimization problem recovers x from measurements $y = Ax + e$:

$$\text{minimize}_{z \in R^N} ||z||_1 \quad \text{subject to} \quad ||Az - y||_2 \leq \epsilon.$$

Assuming x^* as the solution and x_k as the best k-term approximation of x, then:

$$||x^* - x||_2 \leq C_1^{\ell_1} \cdot \epsilon + C_2^{\ell_1} \cdot \frac{||x - x_k||_1}{\sqrt{k}}.$$

Remark

If the measurement matrix A is a random Gaussian matrix then the sufficient condition would be $k \lesssim \frac{n}{log(\frac{N}{n})}$ *which is much worse than the* ℓ_0 *sufficient condition* $k < \frac{n}{2}$.

Recovery by ℓ_1 minimization

Candes, Romberg and Tao showed that if A is sufficiently incoherent, solving the following convex optimization problem recovers x from measurements $y = Ax + e$:

$$\text{minimize}_{z \in R^N} ||z||_1 \quad \text{subject to} \quad ||Az - y||_2 \leq \epsilon.$$

Assuming x^* as the solution and x_k as the best k-term approximation of x, then:

$$||x^* - x||_2 \leq C_1^{\ell_1} \cdot \epsilon + C_2^{\ell_1} \cdot \frac{||x - x_k||_1}{\sqrt{k}}.$$

Remark

If the measurement matrix A is a random Gaussian matrix then the sufficient condition would be $k \lesssim \frac{n}{log(\frac{N}{n})}$ *which is much worse than the* ℓ_0 *sufficient condition* $k < \frac{n}{2}$.

Recovery by ℓ_p minimization

Solving the following non-convex problem also estimates x with weaker sufficient conditions on A than ℓ_1:

$$\text{minimize}_{z \in R^N} ||z||_p \quad \text{subject to} \quad ||Az - y||_2 \leq \epsilon.$$

Theorem

(SY) Assuming x^* as the solution and x_k as the best k-term approximation of x, then if A is sufficiently incoherent, we have:

$$||x^* - x||_2^p \leq C_1^{\ell_p} \cdot \epsilon^p + C_2^{\ell_p} \cdot \frac{||x - x_k||_p^p}{k^{1-\frac{p}{2}}}.$$

Recovery by ℓ_p minimization

Solving the following non-convex problem also estimates x with weaker sufficient conditions on A than ℓ_1:

$$\text{minimize}_{z \in R^N} ||z||_p \quad \text{subject to} \quad ||Az - y||_2 \leq \epsilon.$$

Theorem

(SY) Assuming x^* as the solution and x_k as the best k-term approximation of x, then if A is sufficiently incoherent, we have:

$$||x^* - x||_2^p \leq C_1^{\ell_p} \cdot \epsilon^p + C_2^{\ell_p} \cdot \frac{||x - x_k||_p^p}{k^{1-\frac{p}{2}}}.$$

Phase-diagrams for reconstruction via ℓ_p minimization

This diagram shows the success rate of recovering S-sparse signals using ℓ_p minimization for a Gaussian matrix $A \in \mathbb{R}^{100\times 300}$.

The light-shaded areas show the pairs (p,S) that we have guaranteed recovery.

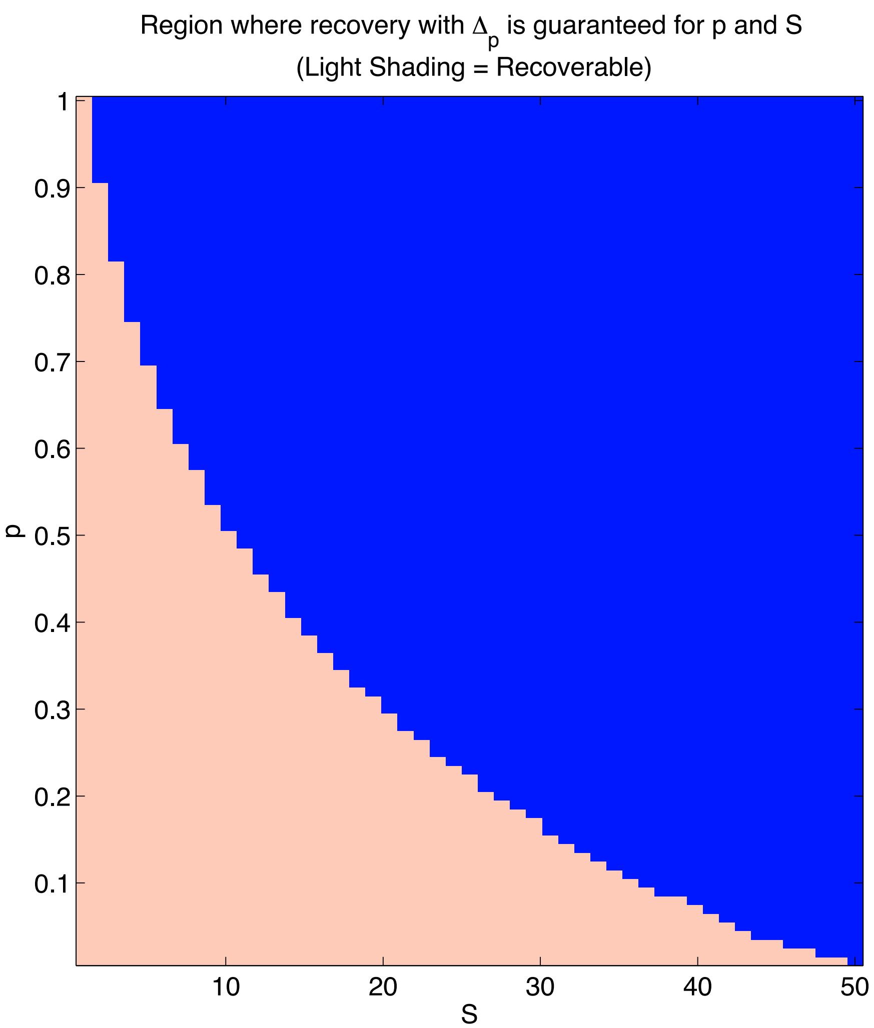

(a) theoretical results

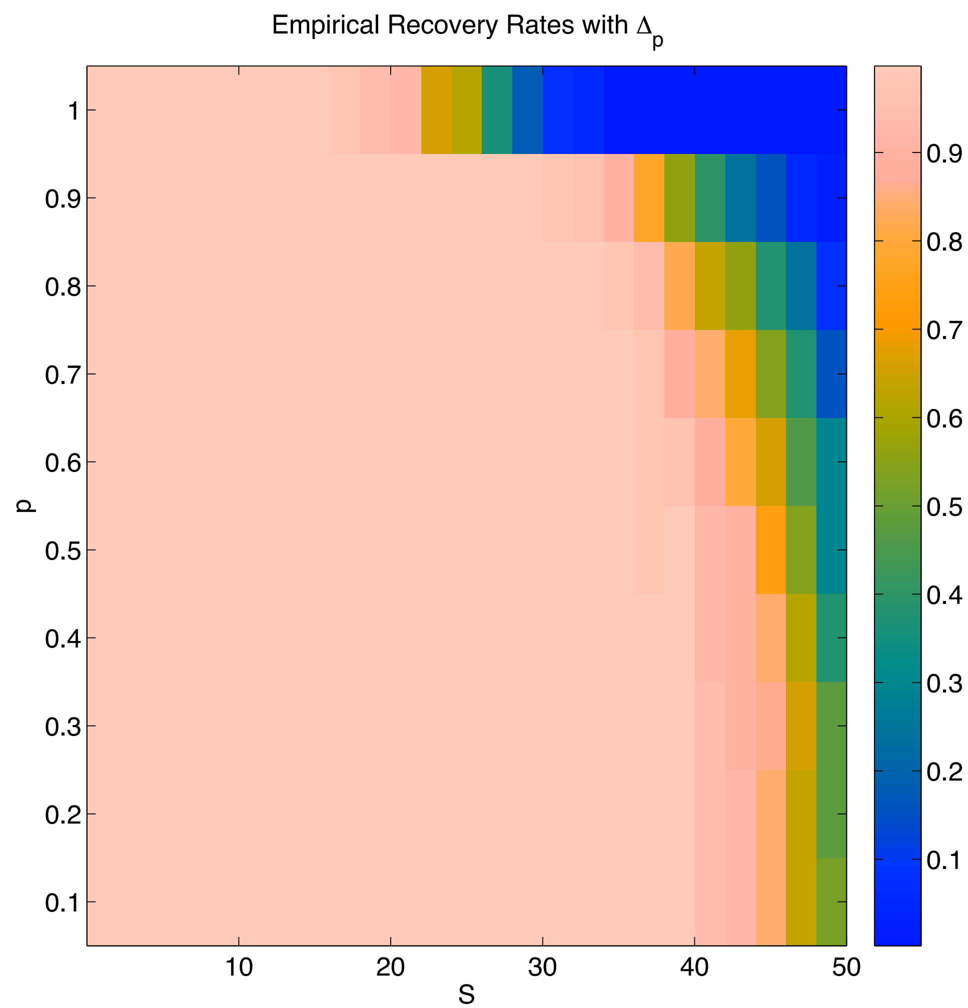

(b) empirical results

Recovery by weighted ℓ_1 minimization

Mansour et al. used a new method to recover x using prior information about it. Assume x is a k-sparse vector which has its support on set T_0 and we estimate the support to be on the set $\widetilde{T}$ which is partially correct.

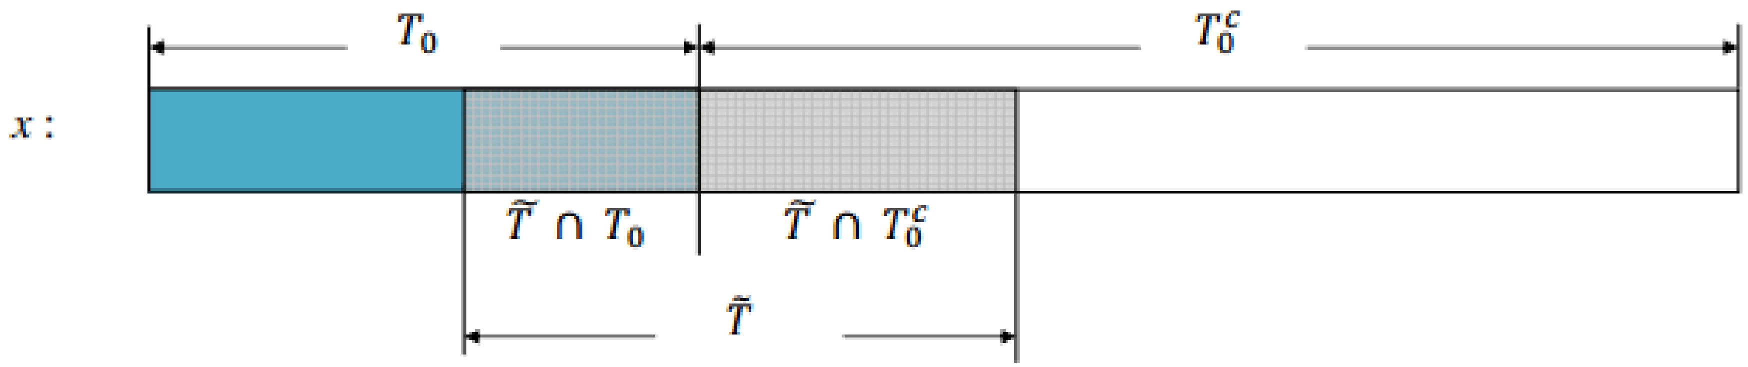

Then minimizing the following weighted ℓ_1 optimization gives us better recovery when we have a good estimate by using the weighted ℓ_1 norm, $||z||_{1,w} = \Sigma_i w_i |z_i|$ instead of the ℓ_1 norm.

$$\text{minimize}_{z \in R^N} ||z||_{1,w} \text{ subject to } ||Az - y||_2 \le \epsilon \; \textit{with} \; w_i = \begin{cases} 1, & \text{if } i \in \widetilde{T}^c \\ w < 1, & \text{if } i \in \widetilde{T} \end{cases}.$$

Recovery by weighted ℓ_1 minimization

Mansour et al. used a new method to recover x using prior information about it. Assume x is a k-sparse vector which has its support on set T_0 and we estimate the support to be on the set $\widetilde{T}$ which is partially correct.

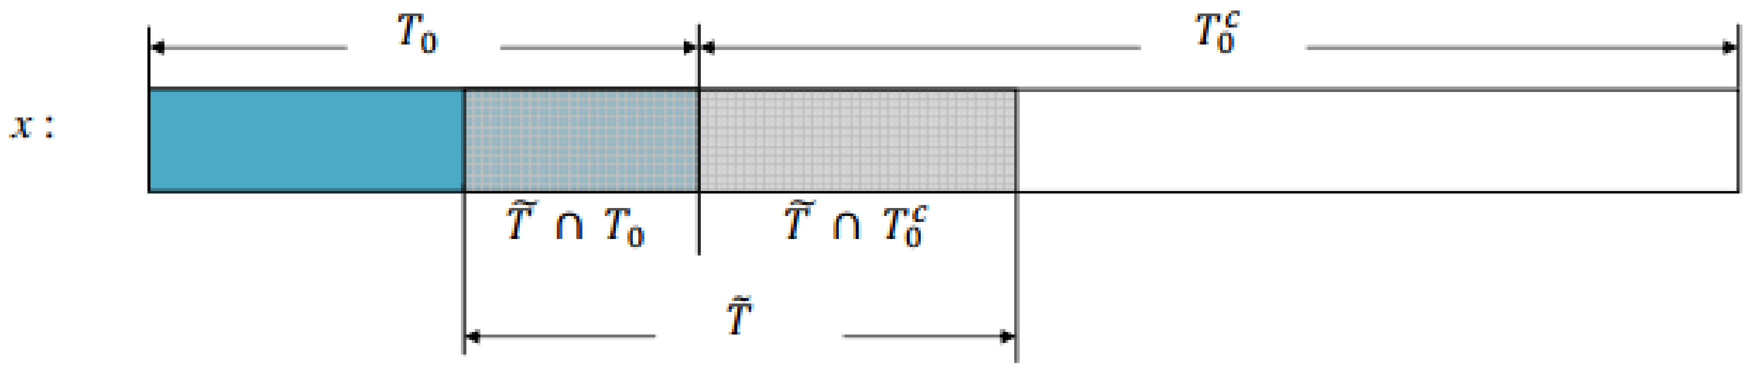

Then minimizing the following weighted ℓ_1 optimization gives us better recovery when we have a good estimate by using the weighted ℓ_1 norm, $||z||_{1,w} = \Sigma_i w_i |z_i|$ instead of the ℓ_1 norm.

$$\text{minimize}_{z \in R^N} ||z||_{1,w} \text{ subject to } ||Az - y||_2 \leq \epsilon \;\; with \;\; w_i = \begin{cases} 1, & \text{if } i \in \widetilde{T}^c \\ w < 1, & \text{if } i \in \widetilde{T} \end{cases} .$$

Stable and robust recovery guarantees of weighted ℓ_1

Let $|\widetilde{T}| = \rho k$ and $\alpha = \frac{|T_0 \cap \widetilde{T}|}{|\widetilde{T}|}$, where ρ defines size of the support estimate and α determines the accuracy of the estimate.

Theorem

(FMSY) If A is sufficiently incoherent, the solution x^* to weighted ℓ_1 obeys:

$$\|x^* - x\|_2 \leq C_1^{w\ell_1} \cdot \epsilon + C_2^{w\ell_1} k^{\frac{-1}{2}} (w\|x - x_k\|_1 + (1-w)\|x_{\widetilde{T}^c \cap T_0^c}\|_1).$$

Stable and robust recovery guarantees of weighted ℓ_1

Let $|\widetilde{T}| = \rho k$ and $\alpha = \frac{|T_0 \cap \widetilde{T}|}{|\widetilde{T}|}$, where ρ defines size of the support estimate and α determines the accuracy of the estimate.

Theorem

(FMSY) If A is sufficiently incoherent, the solution x^* to weighted ℓ_1 obeys:

$$\|x^* - x\|_2 \leq C_1^{w\ell_1} \cdot \epsilon + C_2^{w\ell_1} k^{\frac{-1}{2}} (w\|x - x_k\|_1 + (1-w)\|x_{\widetilde{T}^c \cap T_0^c}\|_1).$$

If α, the accuracy of our estimate is better than 50% then weighted ℓ_1 recovers better than ℓ_1 in terms of sufficient recovery conditions and error bounds.

Stable and robust recovery guarantees of weighted ℓ_1

Let $|\widetilde{T}| = \rho k$ and $\alpha = \frac{|T_0 \cap \widetilde{T}|}{|\widetilde{T}|}$, where ρ defines size of the support estimate and α determines the accuracy of the estimate.

Theorem

(FMSY) If A is sufficiently incoherent, the solution x^* to weighted ℓ_1 obeys:

$$\|x^* - x\|_2 \leq C_1^{w\ell_1} \cdot \epsilon + C_2^{w\ell_1} k^{\frac{-1}{2}} (w\|x - x_k\|_1 + (1-w)\|x_{\widetilde{T}^c \cap T_0^c}\|_1).$$

If α, the accuracy of our estimate is better than 50% then weighted ℓ_1 recovers better than ℓ_1 in terms of sufficient recovery conditions and error bounds.

Stable and robust recovery guarantees of weighted ℓ_1

This plot shows the results of using weighted ℓ_1, recovering 40-sparse signal when $N = 500$.

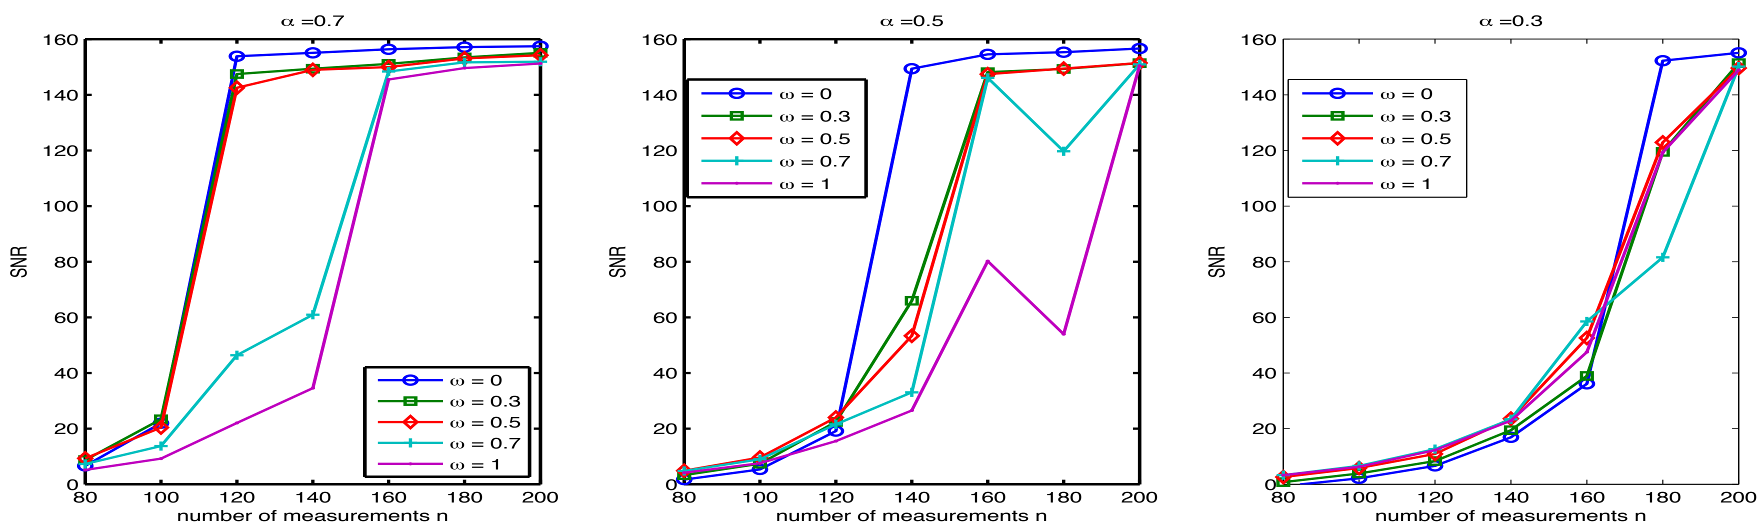

(c) *no noise*

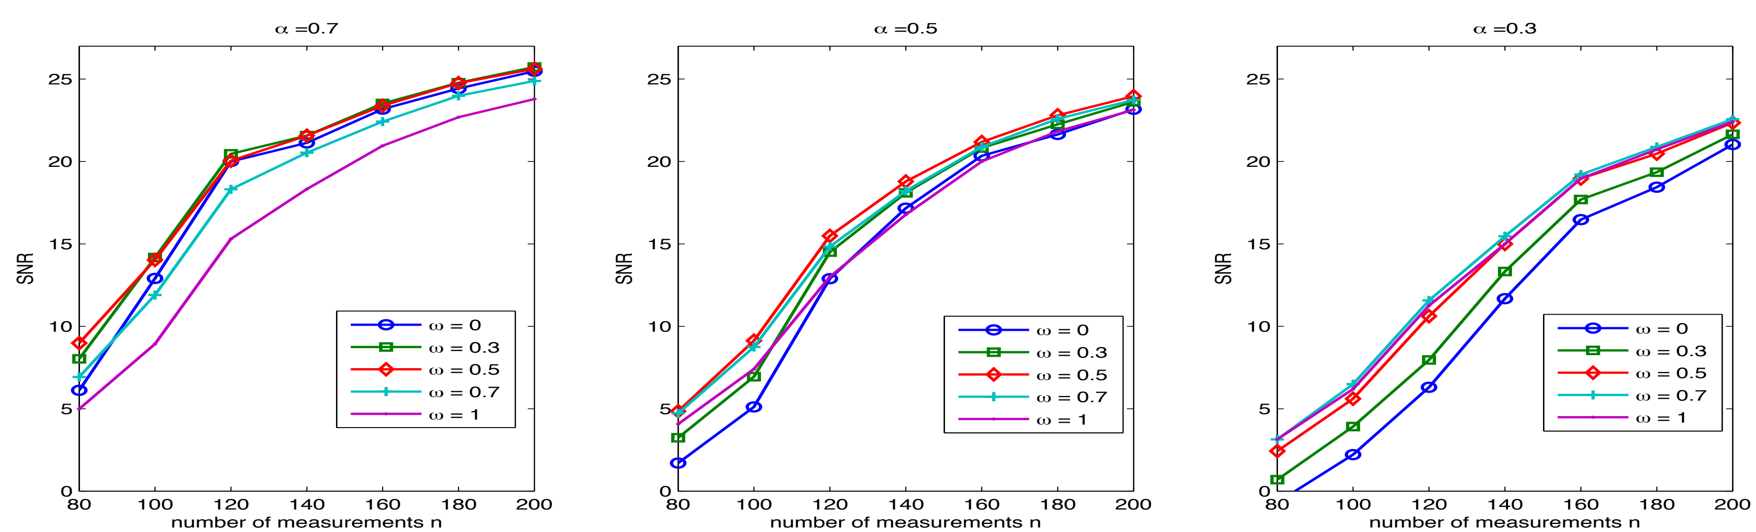

(d) 5% *noise*

Monday, 3 December, 12

Motivation

Using ℓ_p minimization recovers the true signal for a wider range of measurement matrices compared to ℓ_1.

If we can guess a support estimate which is at least 50% accurate, then using weighted ℓ_1 minimization guarantees recovery with weaker RIP conditions and smaller recovery error bounds compared to ℓ_1.

Motivation

Using ℓ_p minimization recovers the true signal for a wider range of measurement matrices compared to ℓ_1.

If we can guess a support estimate which is at least 50% accurate, then using weighted ℓ_1 minimization guarantees recovery with weaker RIP conditions and smaller recovery error bounds compared to ℓ_1.

By using weighted ℓ_p minimization we can enjoy the advantages of both.

Motivation

Using ℓ_p minimization recovers the true signal for a wider range of measurement matrices compared to ℓ_1.

If we can guess a support estimate which is at least 50% accurate, then using weighted ℓ_1 minimization guarantees recovery with weaker RIP conditions and smaller recovery error bounds compared to ℓ_1.

By using weighted ℓ_p minimization we can enjoy the advantages of both.

Weighted ℓ_p minimization

We estimate x from measurements y by solving the following optimization problem:

$$\text{minimize}_{z \in R^N} ||z||_{p,w} \ \text{ subject to } \ ||Az - y||_2 \leq \epsilon \ \ \textit{with} \ \ w_i = \begin{cases} 1, & \text{if } i \in \widetilde{T}^c \\ w < 1, & \text{if } i \in \widetilde{T} \end{cases},$$

where $0 \leq w \leq 1$ and $||z||_{p,w} = (\Sigma_i |w_i z_i|^p)^{\frac{1}{p}}$.

Stable and robust recovery guarantees of weighted ℓ_p

Recall that $|\widetilde{T}| = \rho k$ and $\alpha = \frac{|T_0 \cap \widetilde{T}|}{|\widetilde{T}|}$.

Theorem

(G, Mansour, Yılmaz) If A satisfies some sufficient conditions which are weaker than the analogous sufficient conditions of ℓ_p and weighted ℓ_1, then the solution x^* to weighted ℓ_p obeys:

$$\|x^* - x\|_2^p \leq C_1 \cdot \epsilon^p + C_2 k^{\frac{p}{2}-1}(w^p\|x - x_k\|_p^p + (1 - w^p)\|x_{\widetilde{T}^c \cap T_0^c}\|_p^p).$$

Stable and robust recovery guarantees of weighted ℓ_p

Recall that $|\widetilde{T}| = \rho k$ and $\alpha = \frac{|T_0 \cap \widetilde{T}|}{|\widetilde{T}|}$.

Theorem

(G, Mansour, Yılmaz) If A satisfies some sufficient conditions which are weaker than the analogous sufficient conditions of ℓ_p and weighted ℓ_1, then the solution x^* to weighted ℓ_p obeys:

$$\|x^* - x\|_2^p \leq C_1 \cdot \epsilon^p + C_2 k^{\frac{p}{2}-1}(w^p\|x - x_k\|_p^p + (1 - w^p)\|x_{\widetilde{T}^c \cap T_0^c}\|_p^p).$$

Comparison of sufficient recovery conditions

This plot compares the sufficient recovery conditions on the measurement matrix A, using weighted ℓ_1 and weighted ℓ_p.

When $\alpha = 0.5$ the sufficient recovery conditions of weighted minimizers is the same as using regular minimizers.

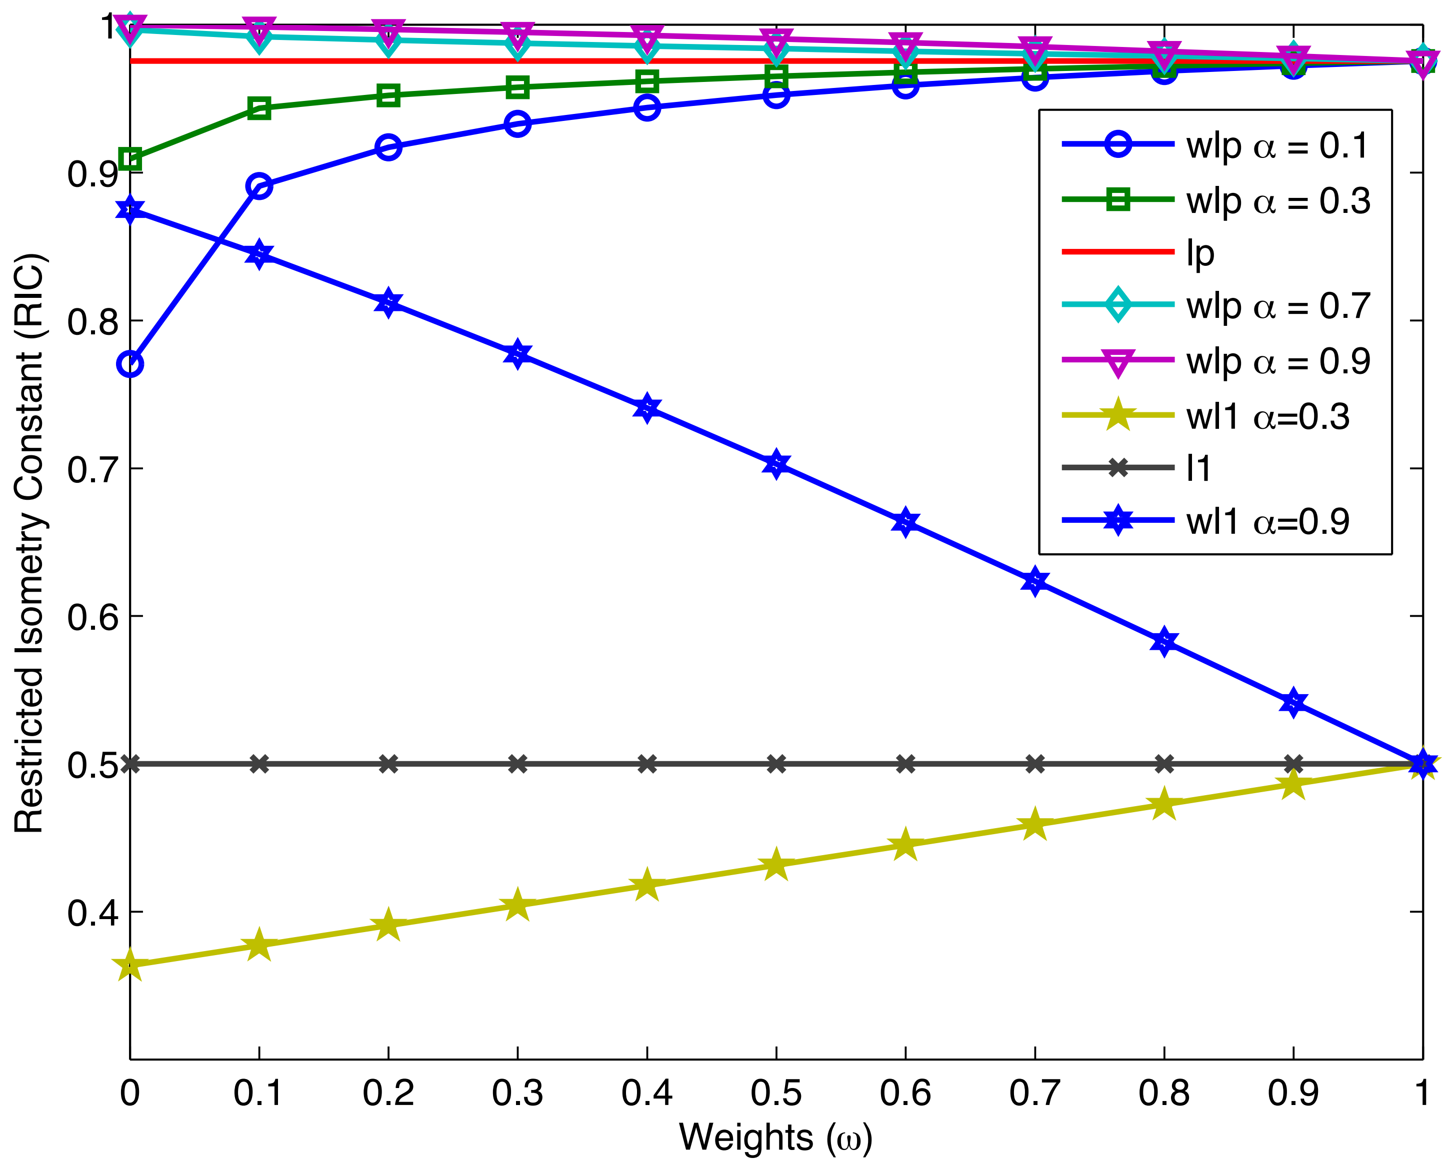

Comparison of constants

As the following plot shows when $\alpha > 0.5$ using smaller weights results in better error bounds.

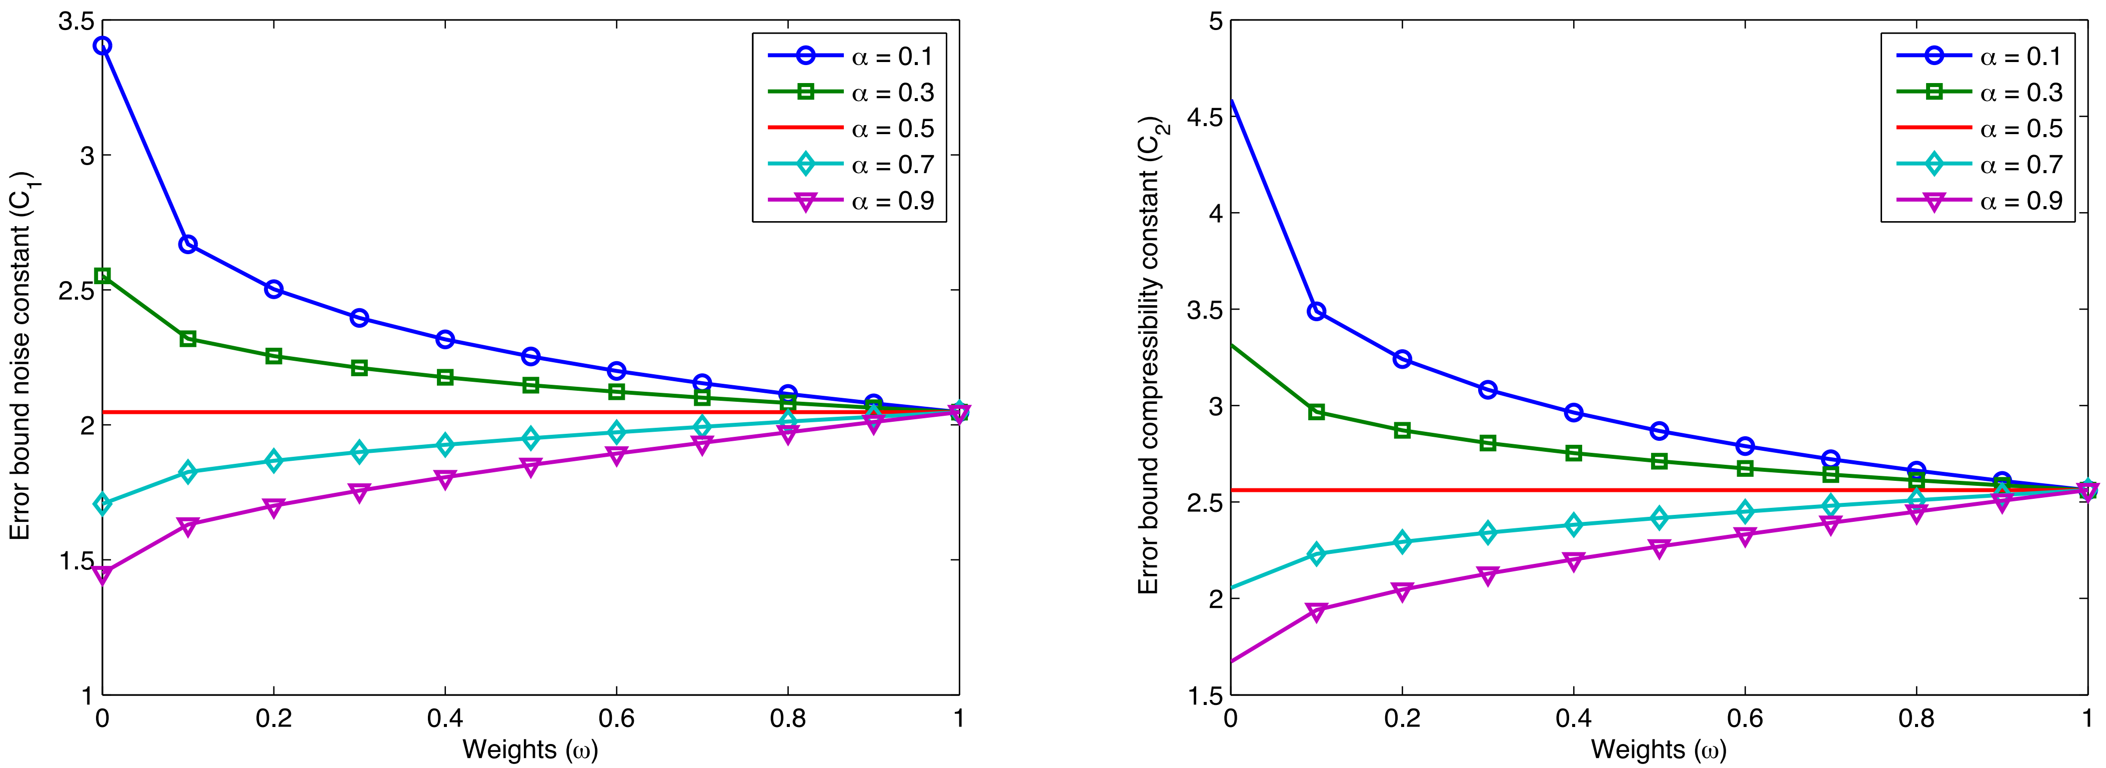

(e) Measurement noise constant (f) Signal compressibility constant

Example 1: Regular ℓ_1 vs Weighted ℓ_1

We generate a 40-sparse random vector x and try to recover x by measurements $b = Ax$ where A is a 80×500 random Gaussian matrix.

The following plot shows the $x - x_{\text{recovered}}$ when we use regular ℓ_1 and weighted ℓ_1 when $\alpha = 0.7$.

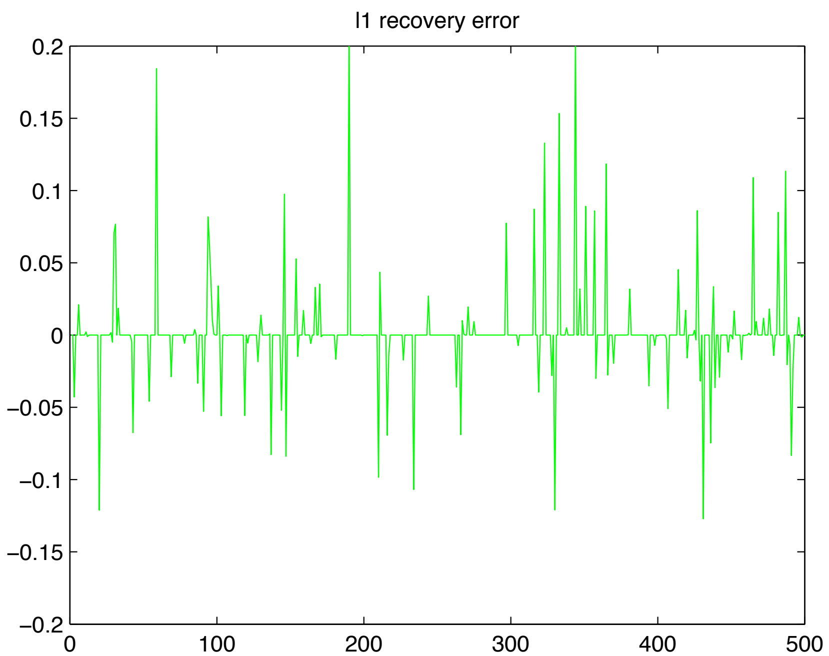

Figure: Recovery error by regular ℓ_1

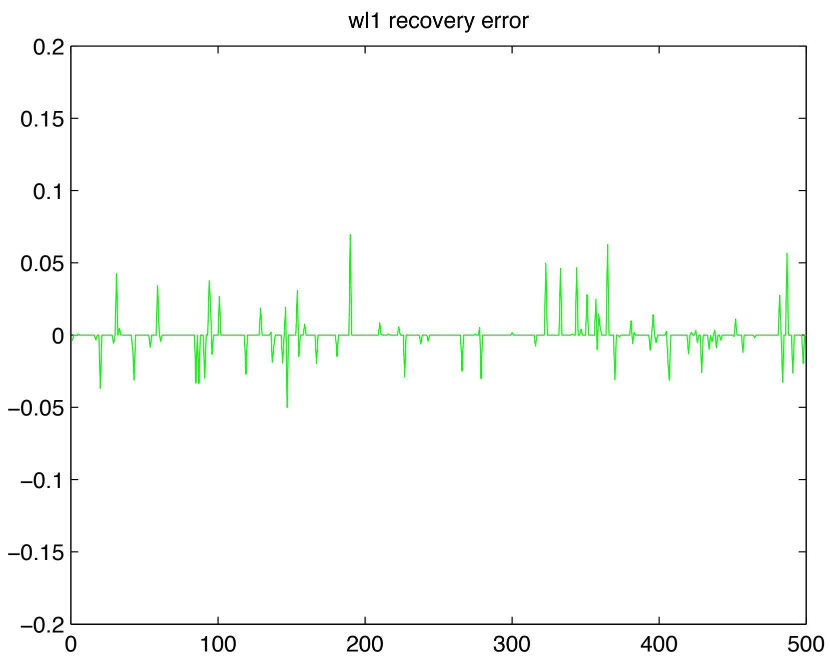

Figure: Recovery error by weighted ℓ_1

Example 1: Regular ℓ_p vs Weighted ℓ_p

Following plot shows the result when we use regular ℓ_p and weighted ℓ_p to recover the same signal when $\alpha = 0.7$ and $p = 0.5$.

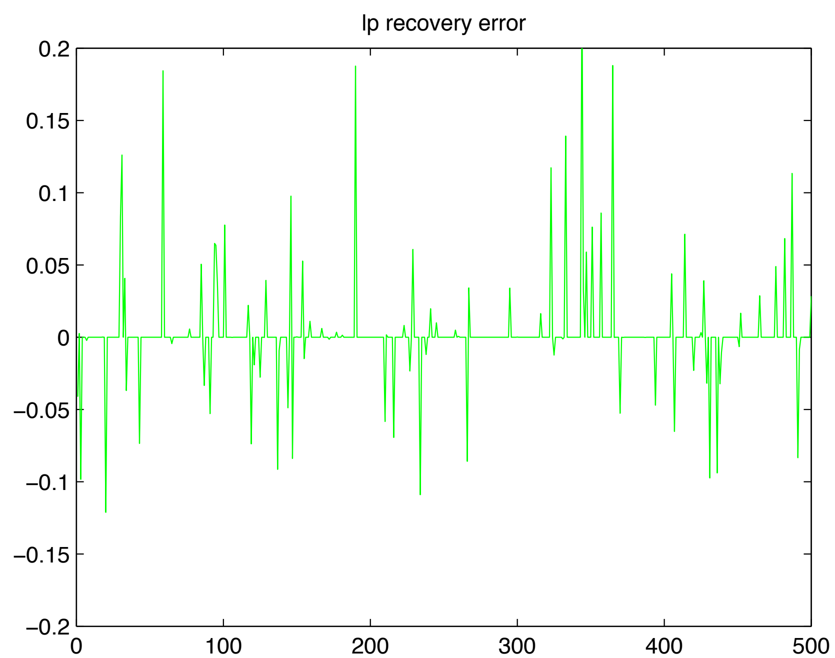

Figure: Recovery error by regular ℓ_p

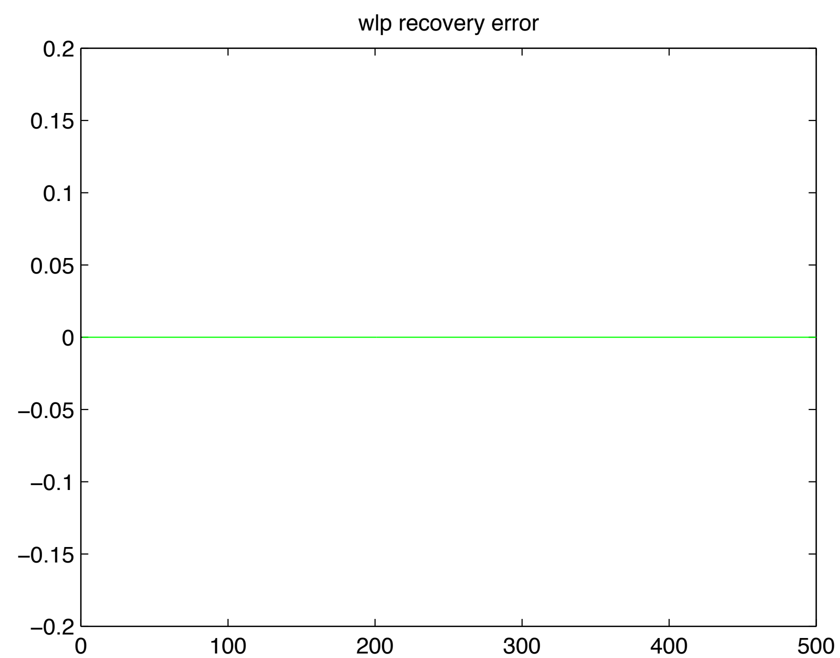

Figure: Recovery error by weighted ℓ_p

Example 2: ℓ_1 vs ℓ_p

This time we generate a 40-sparse random vector x and try to recover x by measurements $b = Ax$ where A is a 100×500 random Gaussian matrix.

The following plot shows the $x - x_{\text{recovered}}$ when we use regular ℓ_1 and regular ℓ_p when $p = 0.5$.

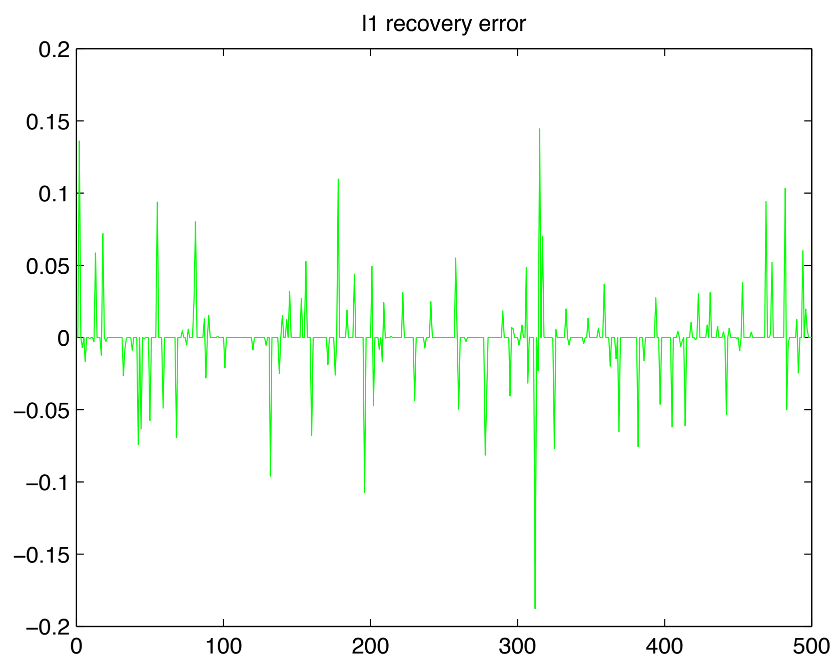

Figure: Recovery error by ℓ_1

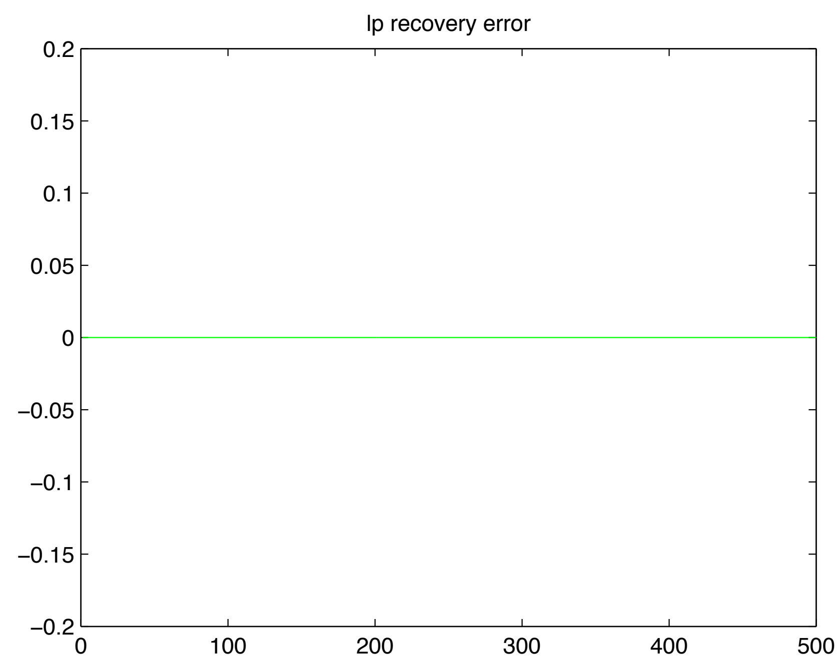

Figure: Recovery error by ℓ_p

Monday, 3 December, 12

Example 2: weighted ℓ_1 vs Weighted ℓ_p

Following plot shows the result when we use weighted ℓ_1 and weighted ℓ_p to recover the same signal when $\alpha = 0.5$ and $p = 0.5$.

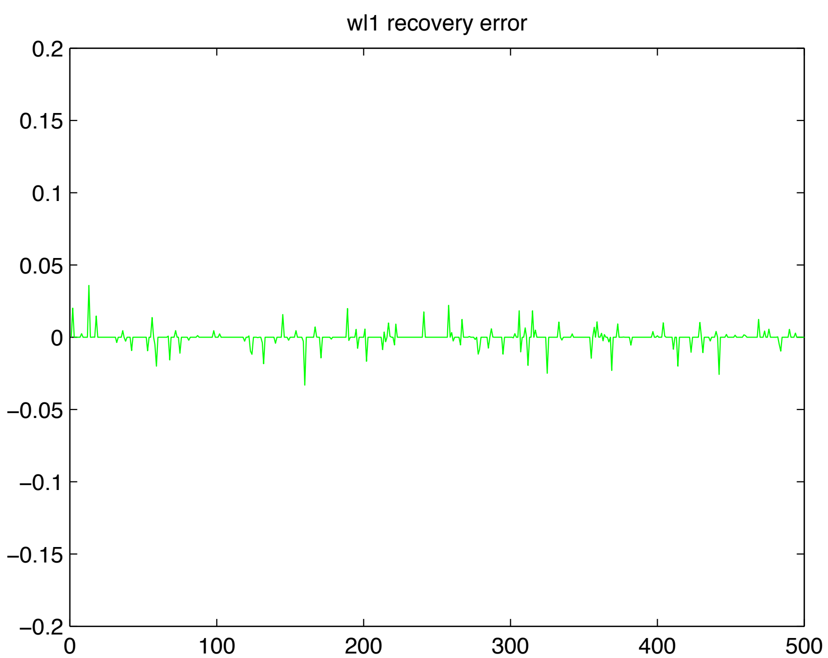

Figure: Recovery error by weighted ℓ_1

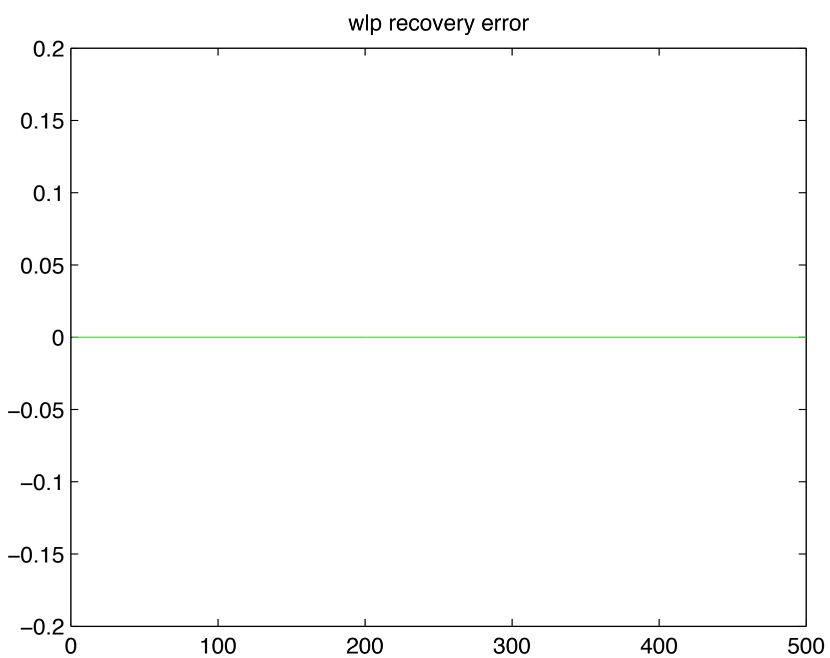

Figure: Recovery error by weighted ℓ_p

Recovery of sparse signals

We take the averaged SNR over 10 experiments for k-sparse signals x with $k = 40$, and $N = 500$ and $p = 0.5$ for variable weights and α.
The noise free case:

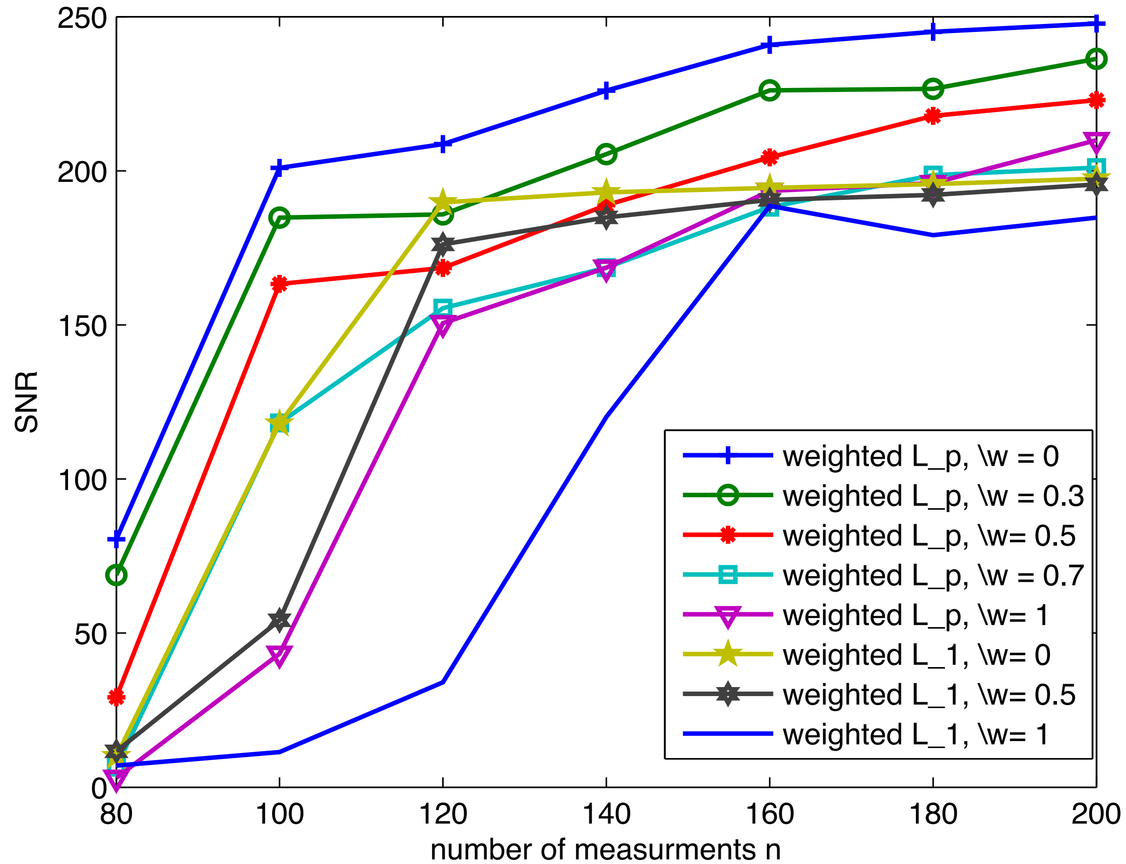

(a) $\alpha = 0.7$

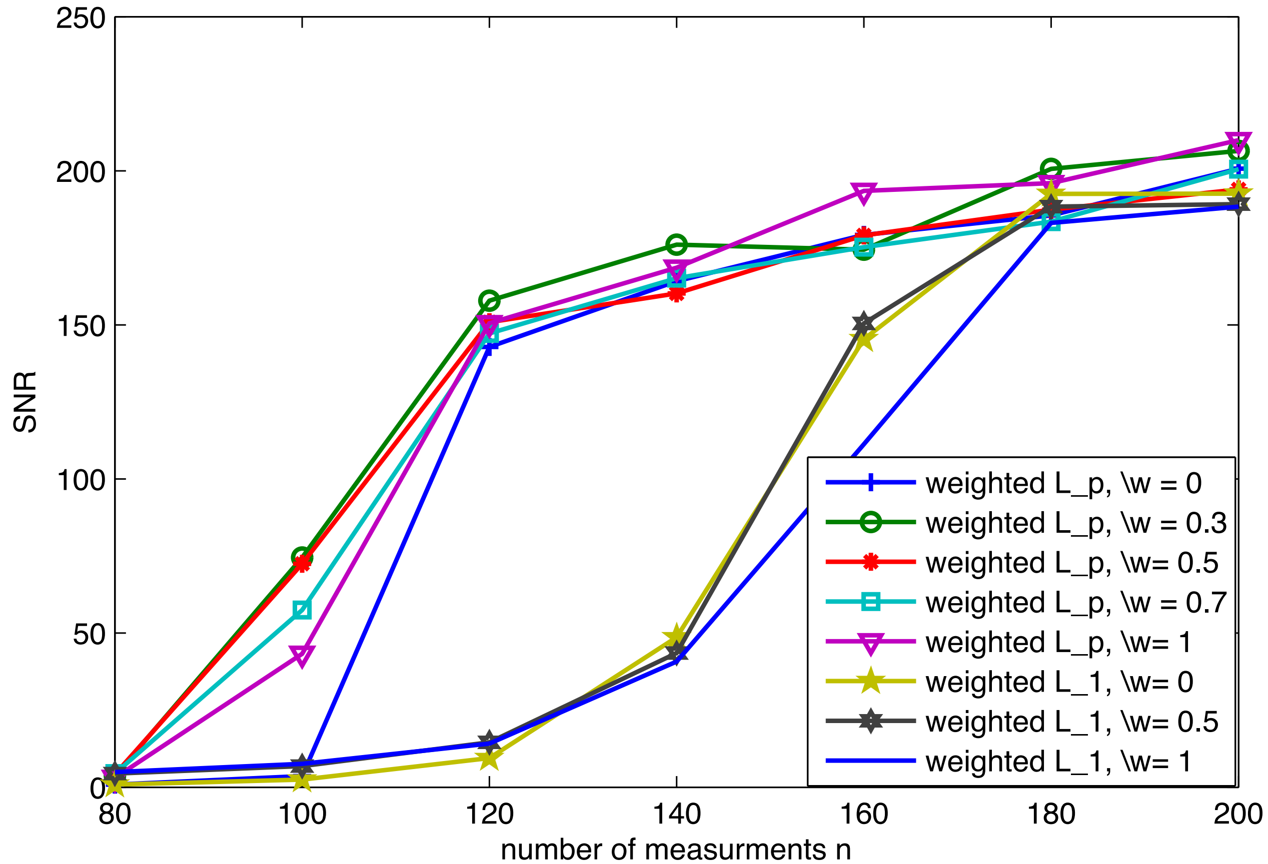

(b) $\alpha = 0.3$

Recovery of sparse signals

We take the averaged SNR over 10 experiments for k-sparse signals x with $k = 40$, and $N = 500$ and $p = 0.5$ for variable weights and α.
The noisy case:

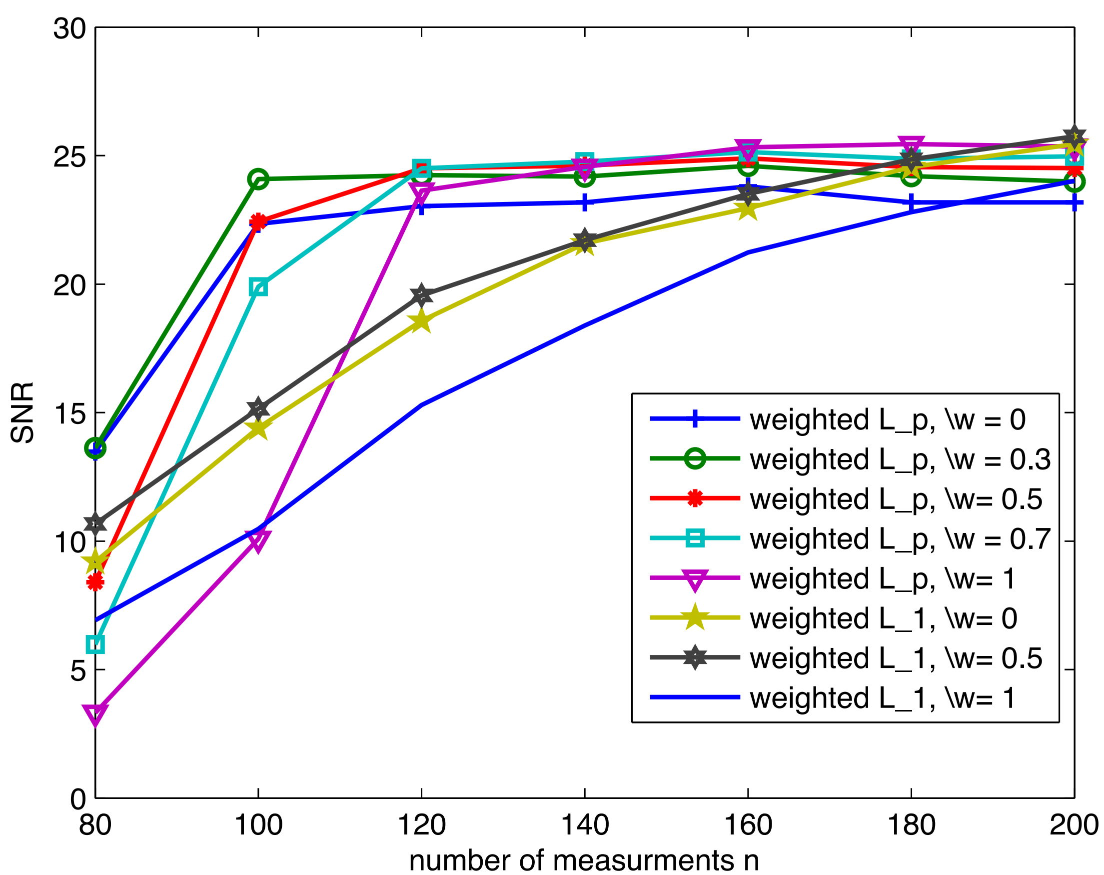

(c) $\alpha = 0.7$

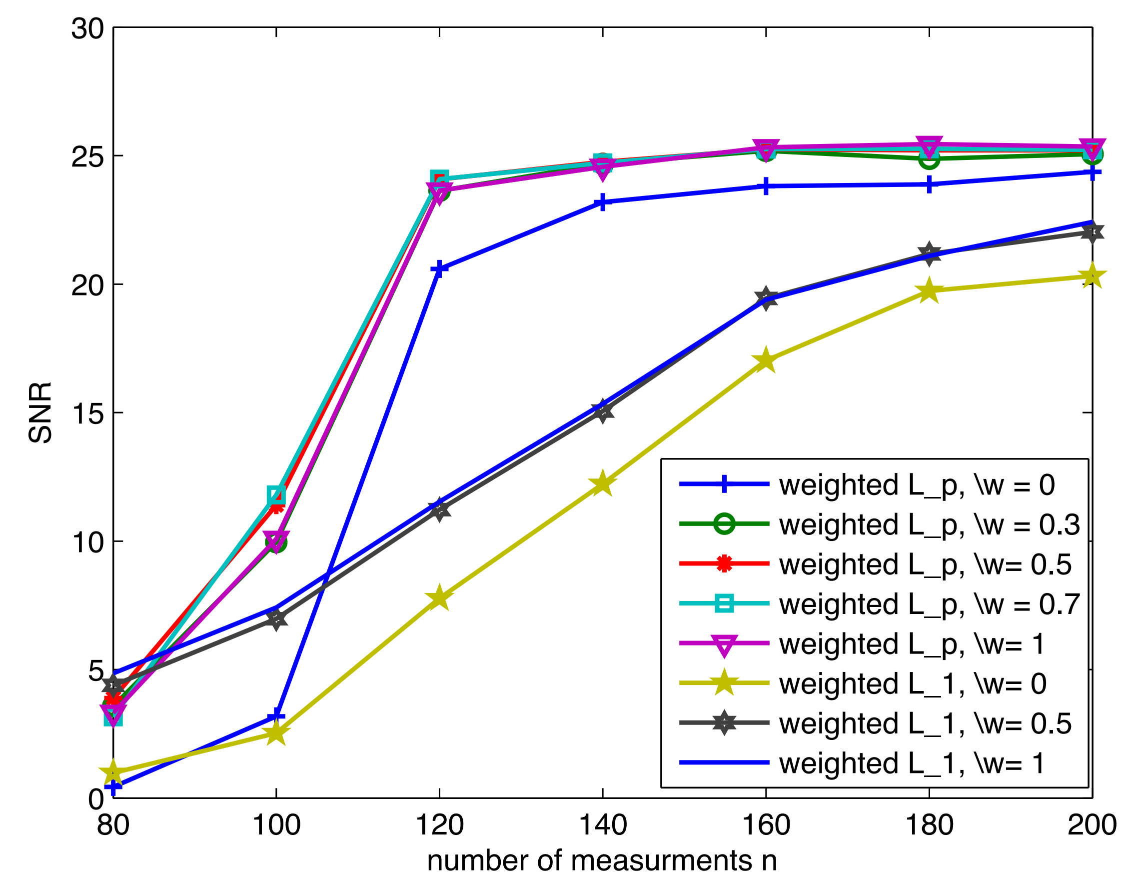

(d) $\alpha = 0.3$

Recovery of Compressible signals

We take the averaged SNR over 10 experiments for signals x whose coefficients decay like j^{-d} where $j \in 1, 2, ..., N$ with $d = 1.5$, $p = 0.5$, $N = 500$ and $n = 100$ for variable weights and support estimate size. The accuracy of the support estimate, α is calculated with respect to the best $k = 20$ term approximation. No noise case:

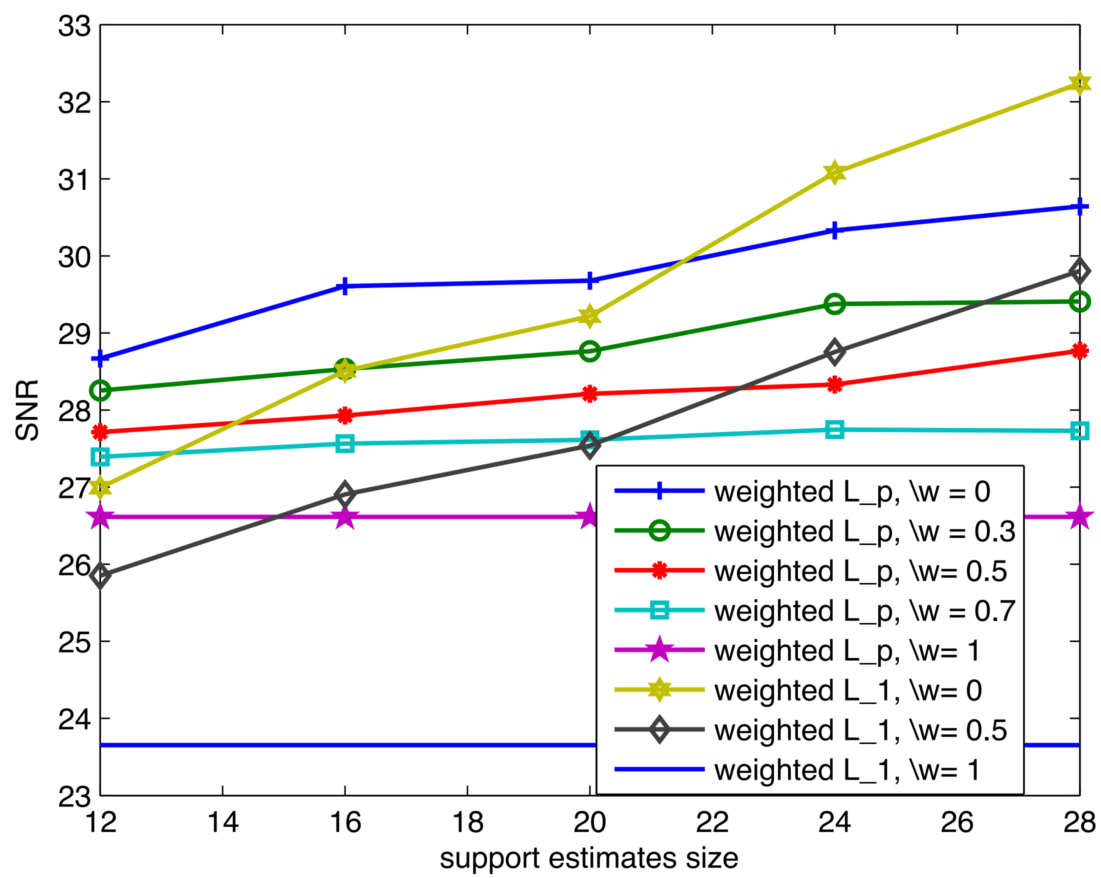

(e) $\alpha = 0.7$

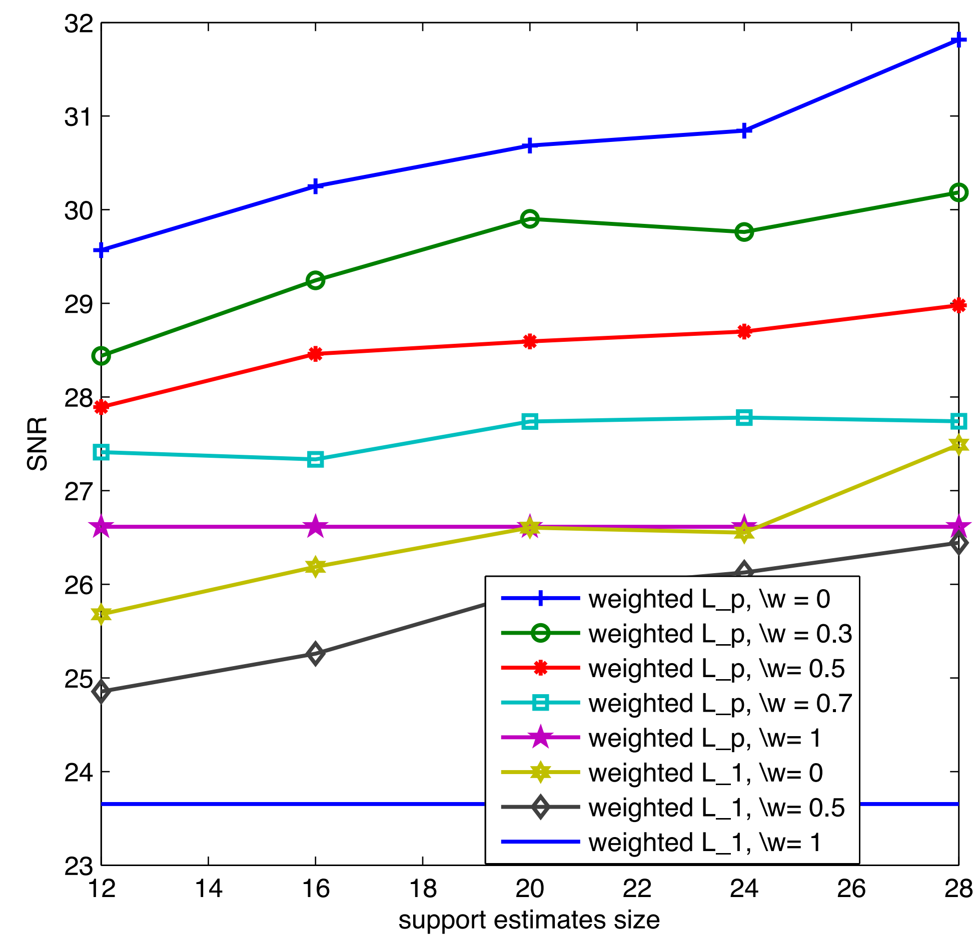

(f) $\alpha = 0.3$

Recovery of Compressible signals

We take the averaged SNR over 10 experiments for signals x whose coefficients decay like j^{-d} where $j \in 1, 2, ..., N$ with $d = 1.5$, $p = 0.5$, $N = 500$ and $n = 100$ for variable weights and support estimate size. The accuracy of the support estimate, α is calculated with respect to the best $k = 20$ term approximation.
Noisy case:

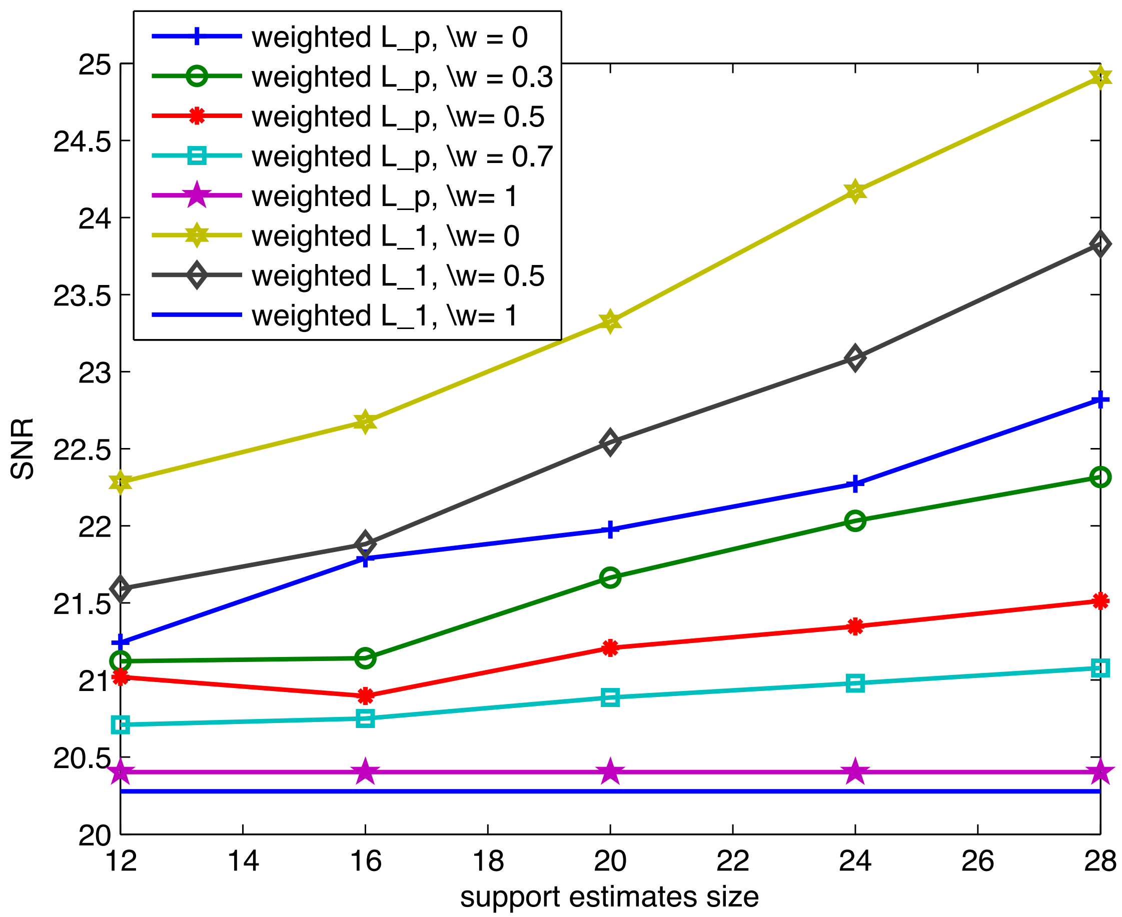

(g) $\alpha = 0.7$

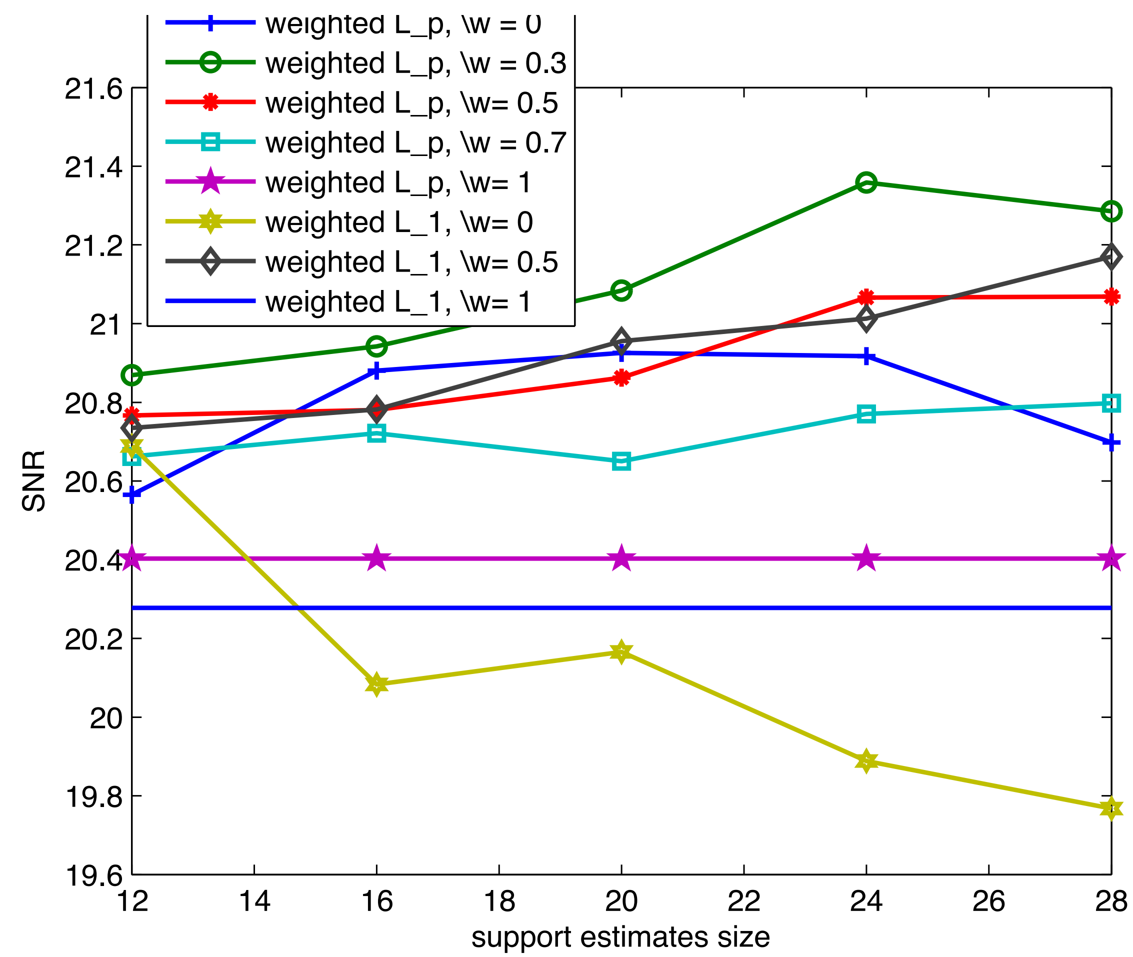

(h) $\alpha = 0.3$

Monday, 3 December, 12

Which p and w to use? sparse signals, $n = 120$

We take the averaged SNR over 10 experiments for k-sparse signals x with $k = 40$, and $N = 500$ and $\alpha = 0.5$ for variable weights and p.
$n = 120$:

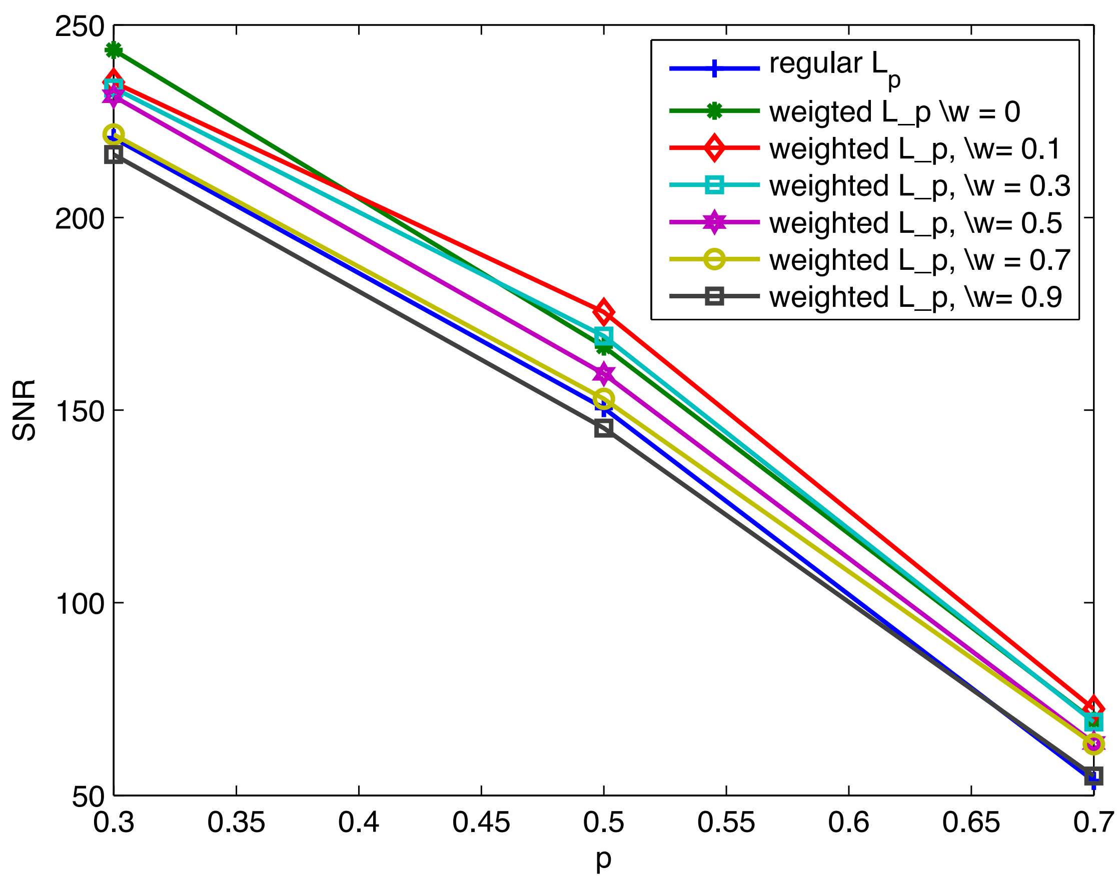

(i) *no noise*

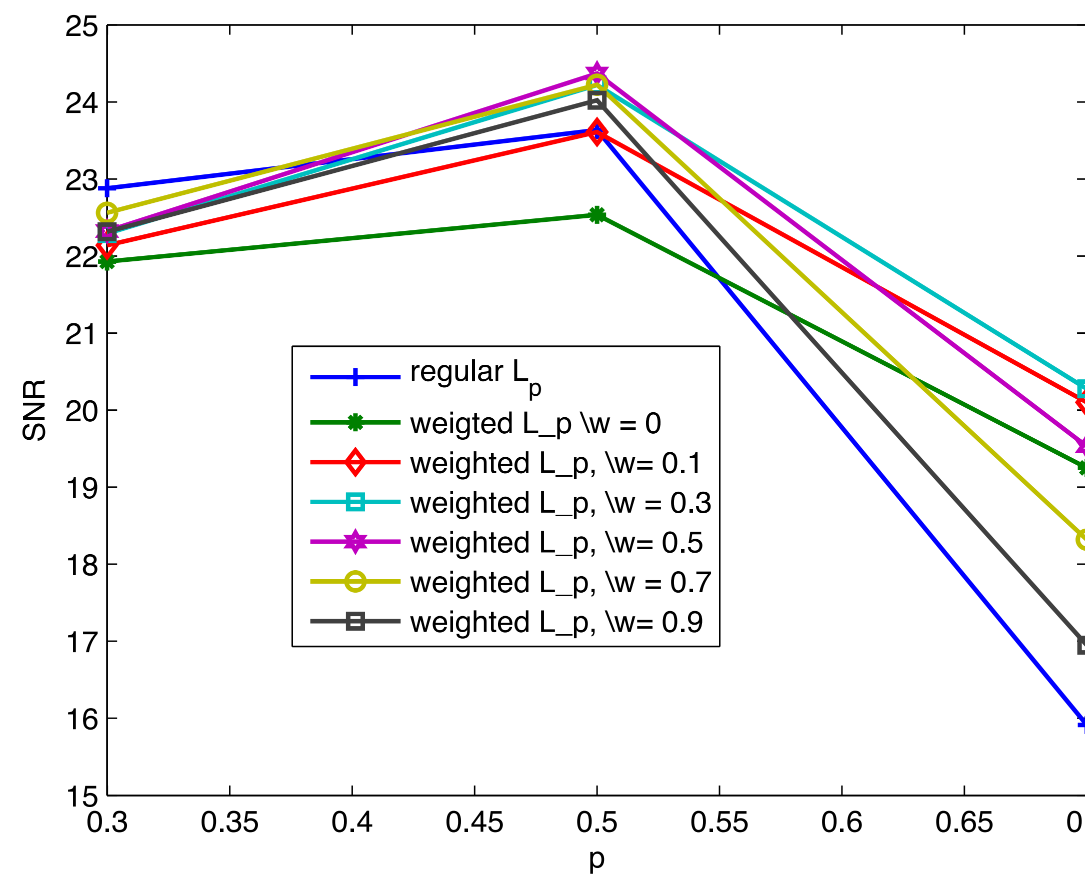

(j) 5%*noise*

Which p and w to use? compressible signals, $n = 120$

We take the averaged SNR over 10 experiments for k-sparse signals x with $k = 40$, and $N = 500$ and $\alpha = 0.5$ for variable weights and p.
$n = 120$:

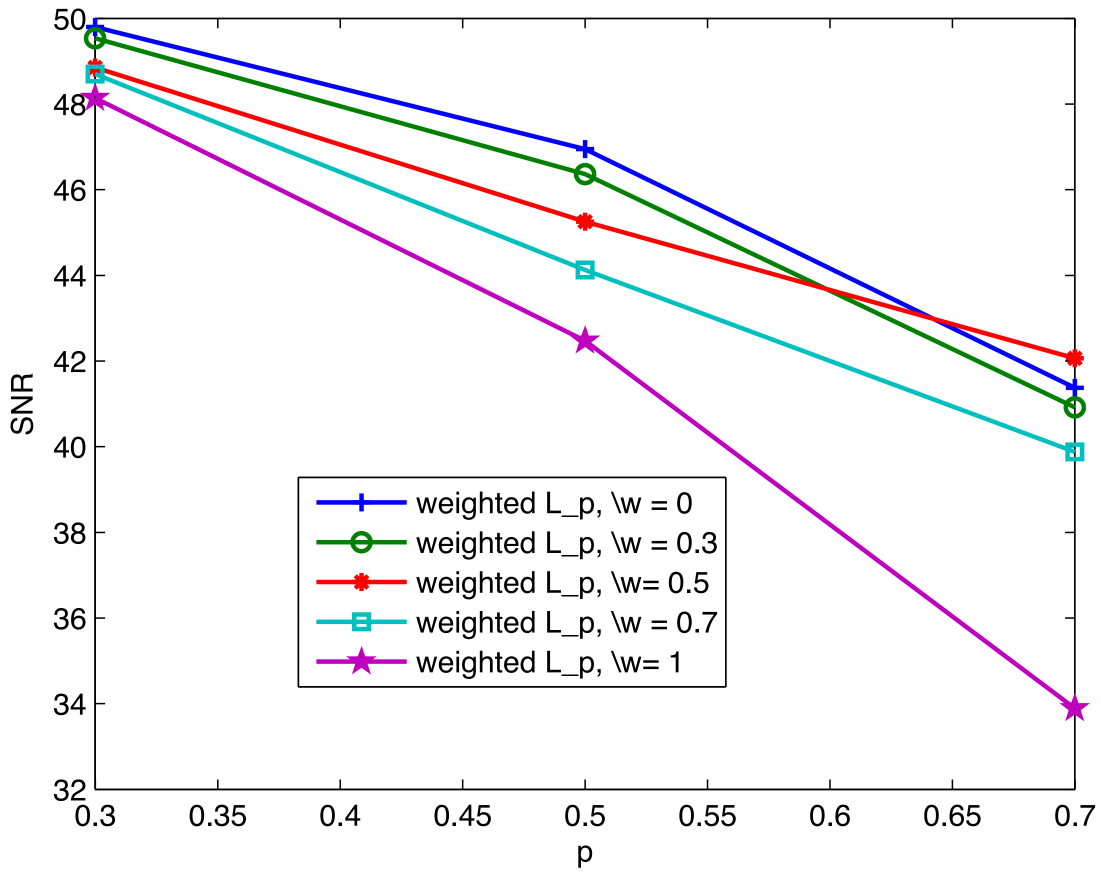

(k) *no noise*

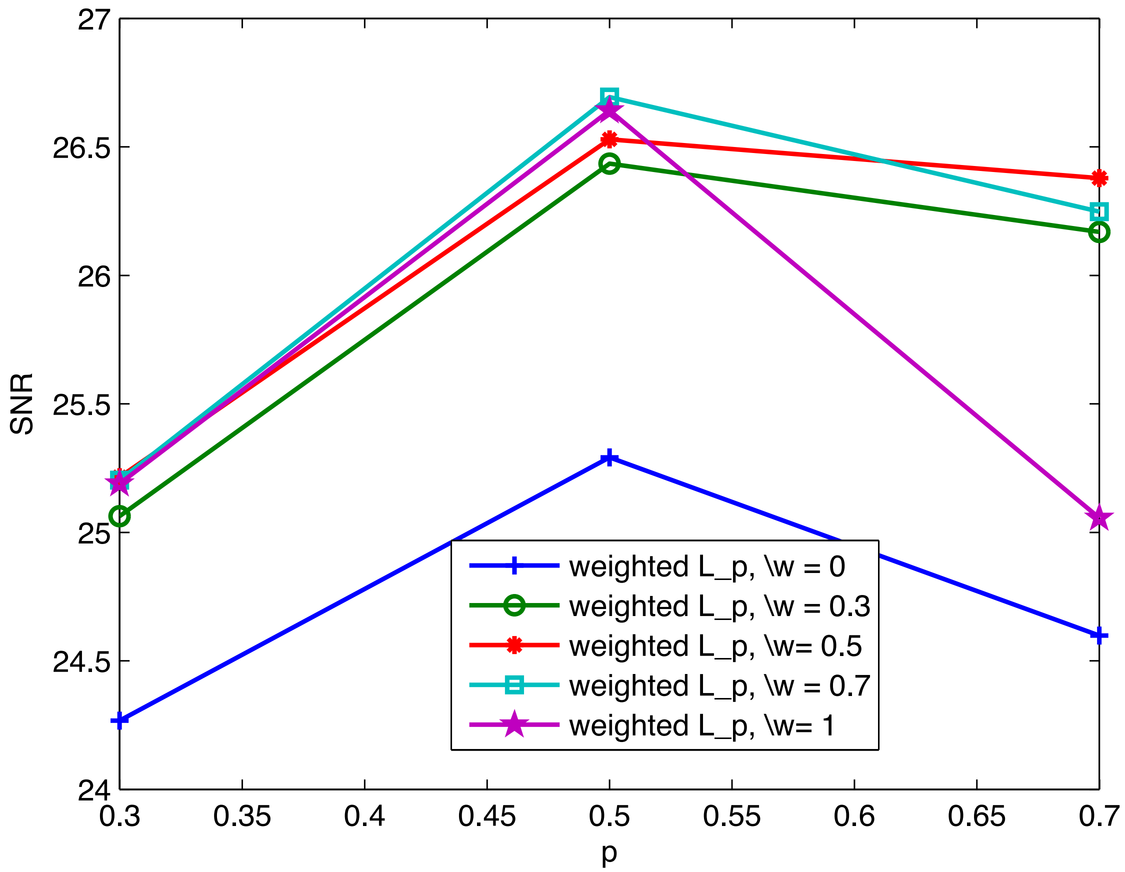

(l) 5%*noise*

Why weighted minimization is a good idea

Seismic data organized in a seismic line exhibit continuity in the time/frequency dimension as well as continuity across the offset/azimuth directions.

As a result there is high correlation between the support sets of adjacent time/frequency slices and of adjacent common-offset/common-azimuth slices.

Why weighted minimization is a good idea

Seismic data organized in a seismic line exhibit continuity in the time/frequency dimension as well as continuity across the offset/azimuth directions.

As a result there is high correlation between the support sets of adjacent time/frequency slices and of adjacent common-offset/common-azimuth slices.

We can use the support set of each slice as an estimate for the support of the next slice.

Why weighted minimization is a good idea

Seismic data organized in a seismic line exhibit continuity in the time/frequency dimension as well as continuity across the offset/azimuth directions.

As a result there is high correlation between the support sets of adjacent time/frequency slices and of adjacent common-offset/common-azimuth slices.

We can use the support set of each slice as an estimate for the support of the next slice.

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

The sensing matrix used at each partition is defined by $A^{(j)} = R^{(j)}S^H$ where $R^{(j)}$ is the subsampling operator restricted to j-th partition and S is the 2-D curvelet transform.

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

The sensing matrix used at each partition is defined by $A^{(j)} = R^{(j)}S^H$ where $R^{(j)}$ is the subsampling operator restricted to j-th partition and S is the 2-D curvelet transform.

First we use the regular ℓ_1 optimization to solve the first partition.

$$\tilde{x}^{(1)} = \arg\min_u ||u||_1 \quad \text{subject to} \quad A^{(1)}u = b^{(1)}.$$

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

The sensing matrix used at each partition is defined by $A^{(j)} = R^{(j)}S^H$ where $R^{(j)}$ is the subsampling operator restricted to j-th partition and S is the 2-D curvelet transform.

First we use the regular ℓ_1 optimization to solve the first partition.

$$\tilde{x}^{(1)} = \arg\min_u ||u||_1 \quad \text{subject to} \quad A^{(1)}u = b^{(1)}.$$

Then at each partition j we use the solution to the previous partition to construct a support estimate for the next iteration and solve the following weighted ℓ_p problem.

$$\tilde{x}^{(j)} = \arg\min_u ||u||_{p,w} \quad \text{subject to} \quad A^{(j)}u = b^{(j)}.$$

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

The sensing matrix used at each partition is defined by $A^{(j)} = R^{(j)}S^H$ where $R^{(j)}$ is the subsampling operator restricted to j-th partition and S is the 2-D curvelet transform.

First we use the regular ℓ_1 optimization to solve the first partition.

$$\tilde{x}^{(1)} = \arg\min_u ||u||_1 \quad \text{subject to} \quad A^{(1)}u = b^{(1)}.$$

Then at each partition j we use the solution to the previous partition to construct a support estimate for the next iteration and solve the following weighted ℓ_p problem.

$$\tilde{x}^{(j)} = \arg\min_u ||u||_{p,w} \quad \text{subject to} \quad A^{(j)}u = b^{(j)}.$$

Problem formulation

We want to recover a high dimensional seismic data volume f by interpolating between a smaller number of measurements $b = RMf$.

To deal with the problem of dimensionality we partition the data along some dimension and then sequentially recover each partition.

The sensing matrix used at each partition is defined by $A^{(j)} = R^{(j)} S^H$ where $R^{(j)}$ is the subsampling operator restricted to j-th partition and S is the 2-D curvelet transform.

First we use the regular ℓ_1 optimization to solve the first partition.

$$\tilde{x}^{(1)} = \arg\min_u ||u||_1 \quad \text{subject to} \quad A^{(1)} u = b^{(1)}.$$

Then at each partition j we use the solution to the previous partition to construct a support estimate for the next iteration and solve the following weighted ℓ_p problem.

$$\tilde{x}^{(j)} = \arg\min_u ||u||_{p,w} \quad \text{subject to} \quad A^{(j)} u = b^{(j)}.$$

Partitioning in the time/frequency domain

Consider the following time slice in the source-receiver domain:

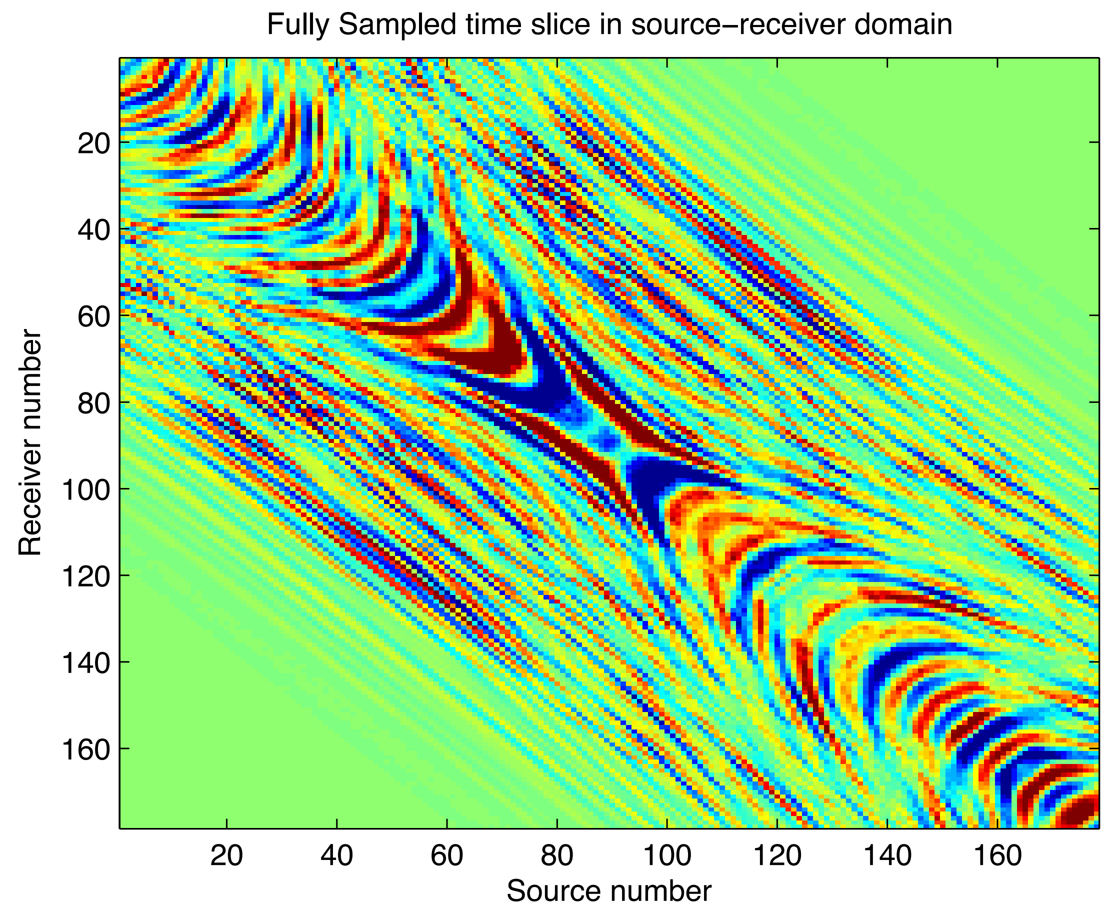

Partitioning in the time/frequency domain

We use the mask shown in the left to get the subsampled time slice in right. The subsampling ratio is 50%.

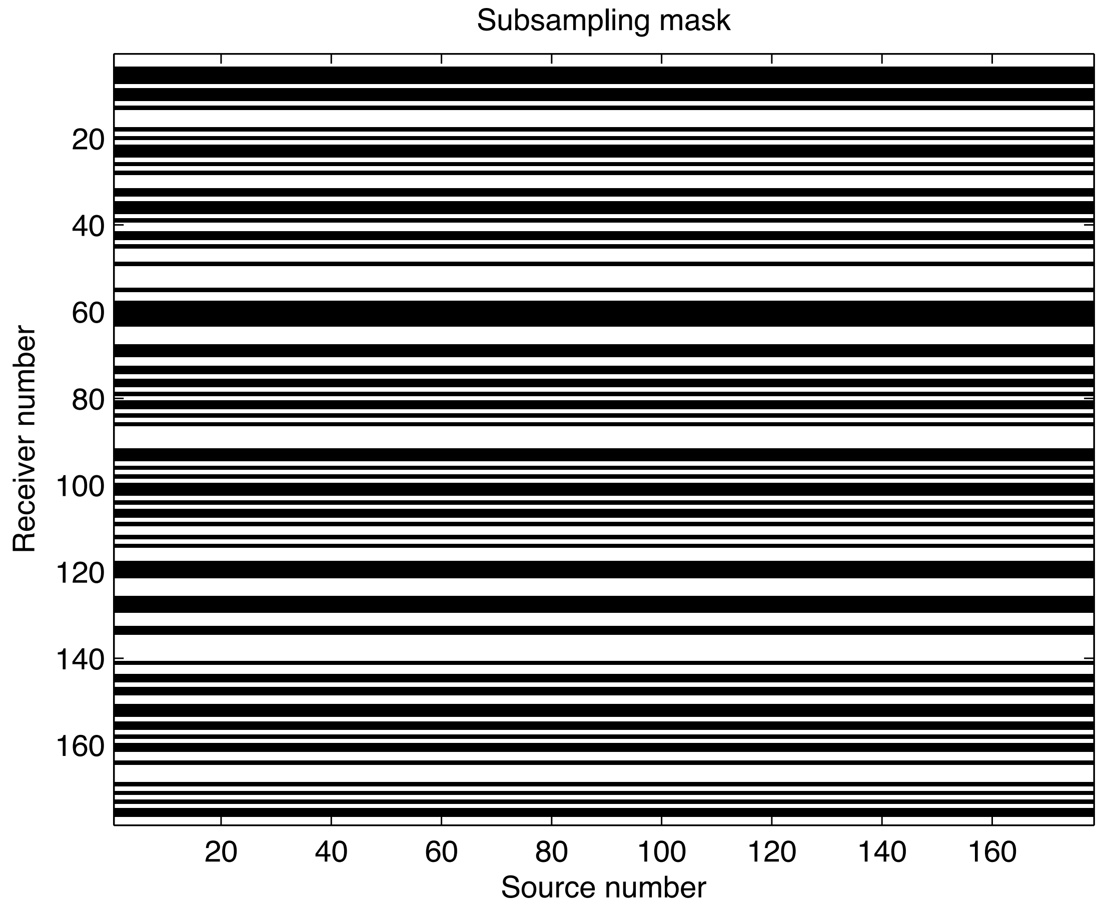

(m) subsampling mask

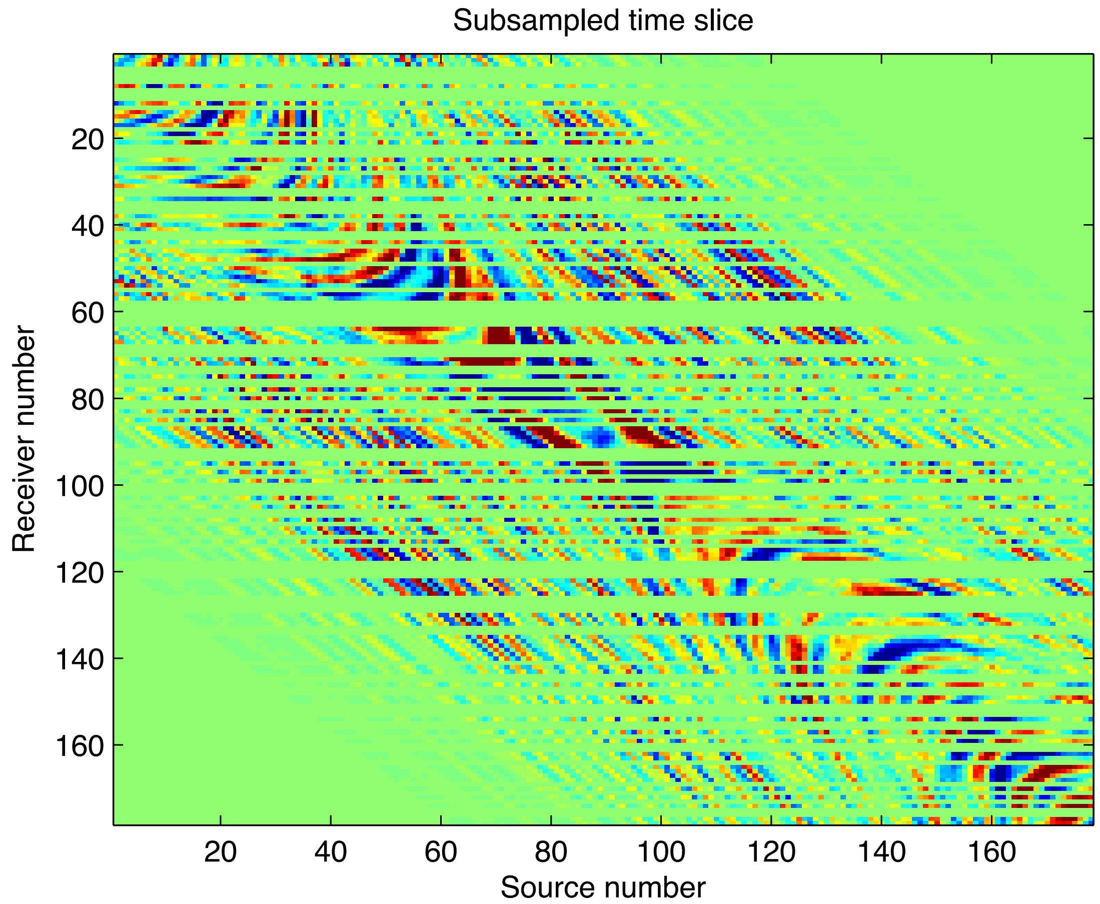

(n) subsampled time slice

Results of partitioning in the time/frequency domain

We used weighted ℓ_p minimization for recovering a seismic line from the Gulf of Suez with 50% randomly subsampled receivers using the mask shown above. The seismic line at full resolution has $N_s = 178$ sources, $N_r = 178$ receivers with a sample distance of 12.5 meters, and $N_t = 500$ time samples acquired with a sampling interval of 4 milliseconds. Here the data are organized in the frequency-source-receiver domain.

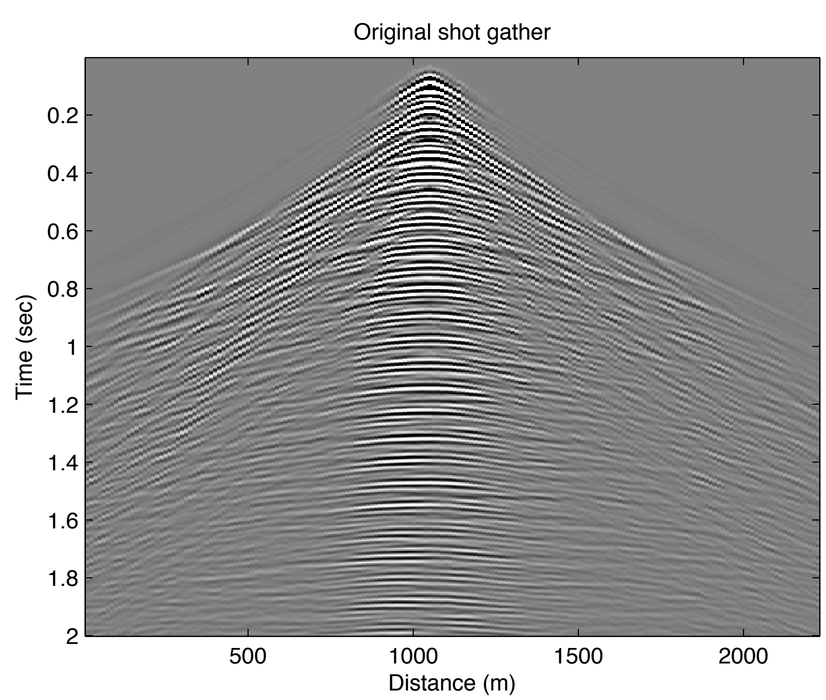

(o) original shot gather

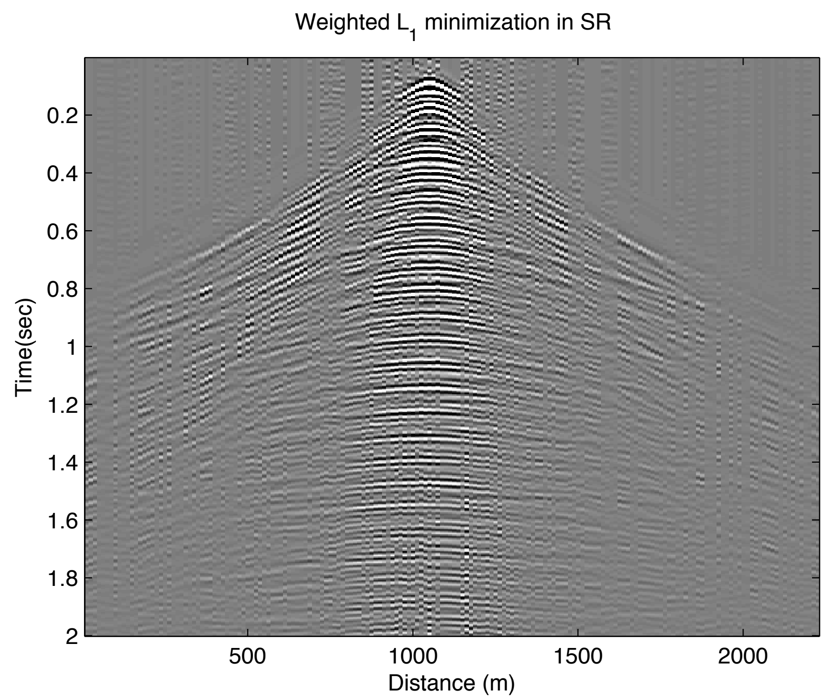

(p) weighted ℓ_p minimization in source-receiver

Conclusion

- We derived stability and robustness guarantees for the recovery of a signal x with partial support estimate $\widetilde{T}$ using weighted ℓ_p minimization.
- We showed that by using ℓ_p recovery and partial support we get much better recovery guarantees in terms of RIP condition and error bounds.

Conclusion

- We derived stability and robustness guarantees for the recovery of a signal x with partial support estimate $\widetilde{T}$ using weighted ℓ_p minimization.
- We showed that by using ℓ_p recovery and partial support we get much better recovery guarantees in terms of RIP condition and error bounds.

Acknowledgement

We thank Rayan Saab for sharing his ℓ_p solver.

This work was in part financially supported by the Collaborative Research and Development Grant DNOISE II (375142-08). This research was carried out as part of the SINBAD II project with support from the following organizations:
BG Group, BGP, BP, Chevron, ConocoPhillips, Petrobras, PGS, Total SA, and WesternGeco.