Sobolev Duals of Random Frames and Sigma-Delta Quantization for Compressed Sensing

Rayan Saab

Department of Mathematics, Duke University

March 7, 2011

Collaborators

Joint work with:

- Sinan Güntürk
- Mark Lammers
- Alex Powell
- Özgür Yılmaz

Notation

- $\Sigma_k^N := \{x \in \mathbb{R}^N : \#\mathrm{supp}(x) \le k\}$
  the set of all "k-sparse" vectors in $\mathbb{R}^N$.

- $A \in \mathbb{R}^{n \times N}, n < N$.
  compressed sensing measurement matrix

- $b = Ax + e, \quad \|e\|_2 \le \epsilon$
  vector of "noisy" compressed sensing measurements

- $\Delta_1 : \mathbb{R}^n \mapsto \mathbb{R}^N$
  $\Delta_1^\epsilon(b) := \arg\min_y \|y\|_1$ subject to $\|b - Ay\|_2 \le \epsilon$.

Compressed sensing

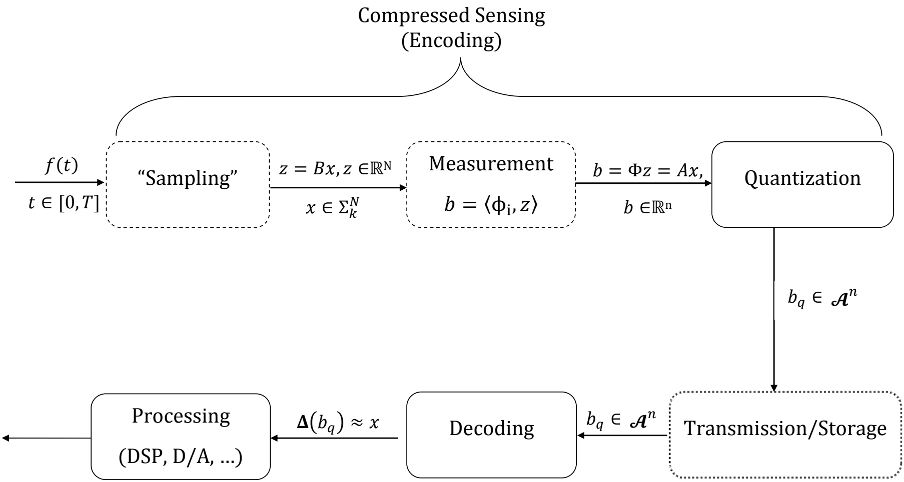

Compressed sensing

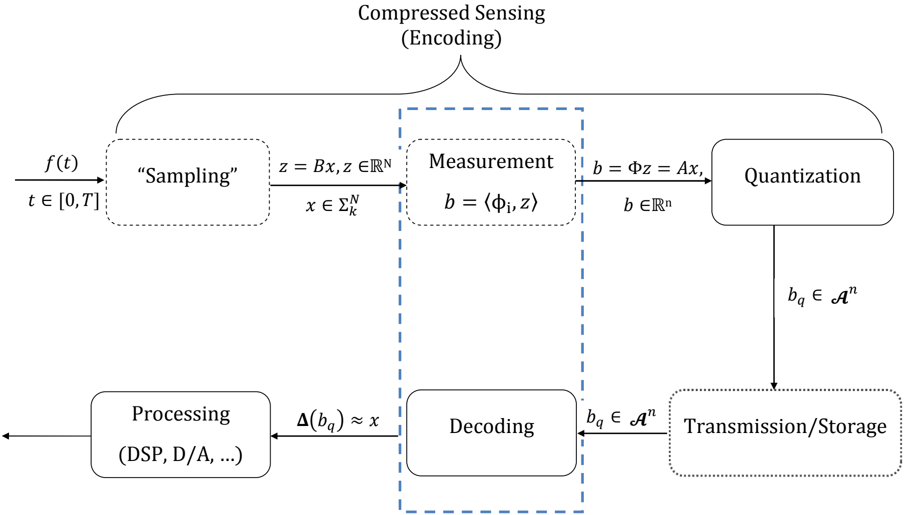

Compressed sensing

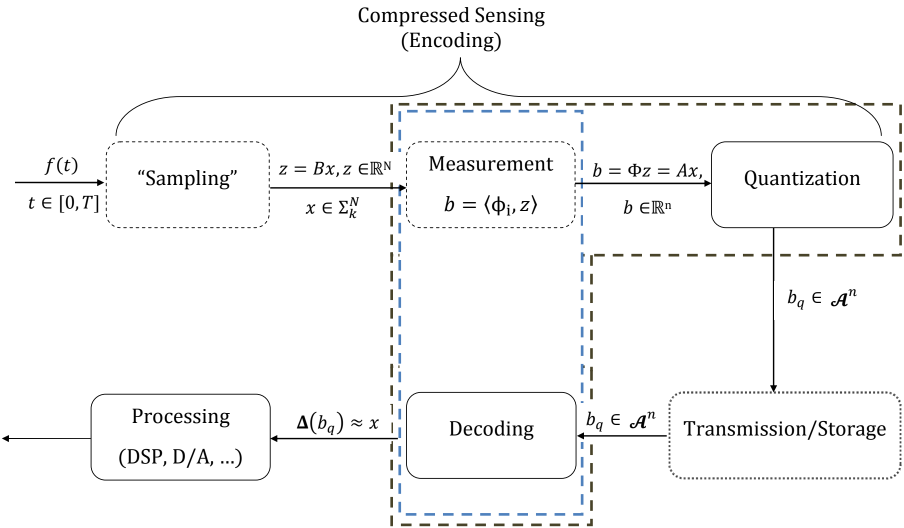

Introduction and Overview

MSQ quantization of CS measurements

ΣΔ quantization of CS measurements

Quantization

- In the context of signal acquisition, we must not only "sample" (or measure) the signal in such a way that we can accurately reconstruct it later (standard compressed sensing results take care of that)
- but we must also quantize the measurements so that we may store/transmit them using digital devices.
- Goal: replace the vector b by a vector whose elements are chosen from a discrete set $\mathcal{A}$, called the quantization alphabet.
- For example,

$$\mathcal{A} = d\mathbb{Z} = \{..., -2d, -d, 0, d, 2d, ...\}.$$

Quantization

- In the context of signal acquisition, we must not only "sample" (or measure) the signal in such a way that we can accurately reconstruct it later (standard compressed sensing results take care of that)
- but we must also quantize the measurements so that we may store/transmit them using digital devices.
- Goal: replace the vector b by a vector whose elements are chosen from a discrete set $\mathcal{A}$, called the quantization alphabet.
- For example,

$$\mathcal{A} = d\mathbb{Z} = \{..., -2d, -d, 0, d, 2d, ...\}.$$

Quantization

- In the context of signal acquisition, we must not only "sample" (or measure) the signal in such a way that we can accurately reconstruct it later (standard compressed sensing results take care of that)
- but we must also quantize the measurements so that we may store/transmit them using digital devices.
- Goal: replace the vector b by a vector whose elements are chosen from a discrete set $\mathcal{A}$, called the quantization alphabet.
- For example,

$$\mathcal{A} = d\mathbb{Z} = \{..., -2d, -d, 0, d, 2d, ...\}.$$

Quantization

- In the context of signal acquisition, we must not only "sample" (or measure) the signal in such a way that we can accurately reconstruct it later (standard compressed sensing results take care of that)
- but we must also quantize the measurements so that we may store/transmit them using digital devices.
- Goal: replace the vector b by a vector whose elements are chosen from a discrete set $\mathcal{A}$, called the quantization alphabet.
- For example,

$$\mathcal{A} = d\mathbb{Z} = \{..., -2d, -d, 0, d, 2d, ...\}.$$

MSQ quantization for compressed sensing

- Memoryless Scalar Quantization (MSQ): "Standard", simple approach.
- In MSQ, we replace each measurement $b(\ell)$, $\ell \in \{1, ..., n\}$ by its nearest neighbor $q_{\mathsf{MSQ}}(\ell) \in \mathcal{A}$.
- Preliminary analysis:

$$|b(\ell) - q_{\mathsf{MSQ}}(\ell)| \leq d/2 \implies \|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n}.$$

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A \sim \mathcal{N}(0,1/n)]{robustness} \|\Delta_1^{\epsilon}(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd\sqrt{n}.$$

- Issues:
  1. The error bound increases as we take more measurements.
  2. The normalization depends on the number of measurements.
- It is more reasonable to use a different normalization.

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A \sim \mathcal{N}(0,1)]{robustness} \|\Delta_1^{\epsilon}(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd.$$

MSQ quantization for compressed sensing

- Memoryless Scalar Quantization (MSQ): "Standard", simple approach.
- In MSQ, we replace each measurement $b(\ell)$, $\ell \in \{1, ..., n\}$ by its nearest neighbor $q_{\mathsf{MSQ}}(\ell) \in \mathcal{A}$.
- Preliminary analysis:

$$|b(\ell) - q_{\mathsf{MSQ}}(\ell)| \leq d/2 \implies \|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n}.$$

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1/n)]{robustness} \|\Delta_1^\epsilon(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd\sqrt{n}.$$

- Issues:
  1. The error bound increases as we take more measurements.
  2. The normalization depends on the number of measurements.
- It is more reasonable to use a different normalization.

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1)]{robustness} \|\Delta_1^\epsilon(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd.$$

MSQ quantization for compressed sensing

- Memoryless Scalar Quantization (MSQ): "Standard", simple approach.
- In MSQ, we replace each measurement $b(\ell)$, $\ell \in \{1, ..., n\}$ by its nearest neighbor $q_{\text{MSQ}}(\ell) \in \mathcal{A}$.
- Preliminary analysis:

$$|b(\ell) - q_{\text{MSQ}}(\ell)| \leq d/2 \implies \|b - q_{\text{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n}.$$

$$\|b - q_{\text{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1/n)]{robustness} \|\Delta_1^{\epsilon}(q_{\text{MSQ}}) - x\|_2 \leq Cd\sqrt{n}.$$

- Issues:
  1. The error bound increases as we take more measurements.
  2. The normalization depends on the number of measurements.
- It is more reasonable to use a different normalization.

$$\|b - q_{\text{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1)]{robustness} \|\Delta_1^{\epsilon}(q_{\text{MSQ}}) - x\|_2 \leq Cd.$$

MSQ quantization for compressed sensing

- Memoryless Scalar Quantization (MSQ): “Standard”, simple approach.
- In MSQ, we replace each measurement $b(\ell)$, $\ell \in \{1, ..., n\}$ by its nearest neighbor $q_{\mathsf{MSQ}}(\ell) \in \mathcal{A}$.
- Preliminary analysis:

$$|b(\ell) - q_{\mathsf{MSQ}}(\ell)| \le d/2 \implies \|b - q_{\mathsf{MSQ}}\|_2 \le \frac{d}{2}\sqrt{n}.$$

$$\|b - q_{\mathsf{MSQ}}\|_2 \le \frac{d}{2}\sqrt{n} \xrightarrow[A \sim \mathcal{N}(0,1/n)]{robustness} \|\Delta_1^{\epsilon}(q_{\mathsf{MSQ}}) - x\|_2 \le Cd\sqrt{n}.$$

- Issues:
  1. The error bound increases as we take more measurements.
  2. The normalization depends on the number of measurements.
- It is more reasonable to use a different normalization.

$$\|b - q_{\mathsf{MSQ}}\|_2 \le \frac{d}{2}\sqrt{n} \xrightarrow[A \sim \mathcal{N}(0,1)]{robustness} \|\Delta_1^{\epsilon}(q_{\mathsf{MSQ}}) - x\|_2 \le Cd.$$

MSQ quantization for compressed sensing

- Memoryless Scalar Quantization (MSQ): "Standard", simple approach.
- In MSQ, we replace each measurement $b(\ell)$, $\ell \in \{1, ..., n\}$ by its nearest neighbor $q_{\mathsf{MSQ}}(\ell) \in \mathcal{A}$.
- Preliminary analysis:

$$|b(\ell) - q_{\mathsf{MSQ}}(\ell)| \leq d/2 \implies \|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n}.$$

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1/n)]{robustness} \|\Delta_1^\epsilon(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd\sqrt{n}.$$

- Issues:
  1. The error bound increases as we take more measurements.
  2. The normalization depends on the number of measurements.
- It is more reasonable to use a different normalization.

$$\|b - q_{\mathsf{MSQ}}\|_2 \leq \frac{d}{2}\sqrt{n} \xrightarrow[A\sim\mathcal{N}(0,1)]{robustness} \|\Delta_1^\epsilon(q_{\mathsf{MSQ}}) - x\|_2 \leq Cd.$$

Limitations of MSQ quantization

- Problem: the MSQ error (bound) does not decrease with n.

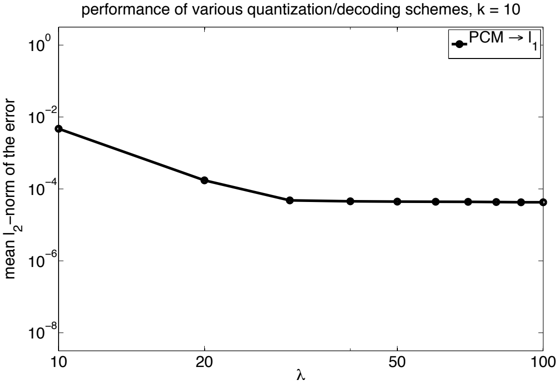

- Possible Remedy: use different decoders to reconstruct from MSQ quantized measurements. However...

Limitations of MSQ quantization

Theorem (Goyal et al.)

Let E be an $n \times k$ real matrix, and let K be a bounded set in $\mathbb{R}^k$. For $x \in K$, suppose we obtain q_{MSQ} by quantizing the entries of $b = Ex$ using MSQ with alphabet $\mathcal{A} = d\mathbb{Z}$. Let Δ_{opt} be an optimal decoder. Then,

$$\left[\mathbb{E}\,\|x - \Delta_{opt}(q_{MSQ}(x))\|_2^2\right]^{1/2} \gtrsim \frac{k}{n} d$$

- Above, the expectation is with respect to a probability measure on x that is, for example, absolutely continuous.
- $\implies$ alternative reconstruction algorithms from MSQ-quantized compressed sensing measurements offer limited improvement.

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.
- 1st order $\Sigma\Delta$ scheme with alphabet $\mathcal{A}$ (greedy rule):

```
Initialize u_0 = 0
for i = 1 to n do
    q_i = Q(u_{i-1} + b_i)
    u_i = u_{i-1} + b_i - q_i.
end for
```

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.
- 1st order $\Sigma\Delta$ scheme with alphabet $\mathcal{A}$ (greedy rule):

$$\begin{cases} q_i & = & Q(u_{i-1} + b_i) \\ (\Delta u)_i & := & u_i - u_{i-1} = b_i - q_i \end{cases}$$

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.
- 2nd order $\Sigma\Delta$ scheme with alphabet $\mathcal{A}$ (greedy rule):

Initialize $u_0^{(1)} = 0, u_0^{(2)} = 0$
for $i = 1$ to n **do**
$$q_i = Q\left(\sum_{j=1}^{2} u_{i-1}^{(j)} + b_i\right)$$
$$u_i^{(1)} = u_{i-1}^{(1)} + b_i - q_i.$$
$$u_i^{(2)} = u_{i-1}^{(2)} + u_i^{(1)}.$$
end for

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.
- 2nd order $\Sigma\Delta$ scheme with alphabet $\mathcal{A}$ (greedy rule):

$$\text{Define } u := u^{(2)}, \text{ then } \begin{cases} q_i & = & Q\left(2u_{i-1} - u_{i-2} + b_i\right) \\ (\Delta^2 u)_i & = & b_i - q_i \end{cases}$$

$\Sigma\Delta$ quantization

- Alternative quantization scheme.
- Used, for example, in quantizing bandlimited signals.
- Define $Q(v) := \arg\min_{q \in \mathcal{A}} |v - q|$.
- rth order $\Sigma\Delta$ scheme with alphabet $\mathcal{A}$ (greedy rule):

$$\begin{cases} q_i & = & Q\left(\sum_{j=1}^{r} (-1)^{j-1} \binom{r}{j} u_{i-j} + b_i\right) \\ (\Delta^r u)_i & = & b_i - q_i \end{cases}$$

Numerical experiments:

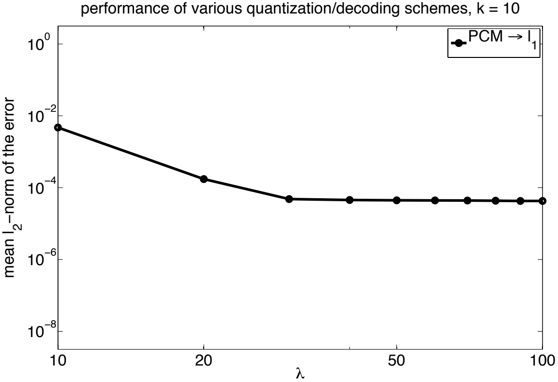

Numerical experiments:

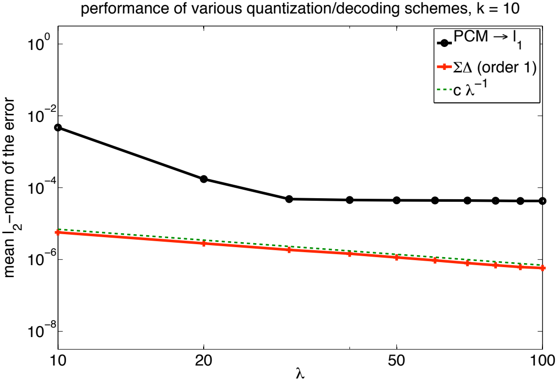

Numerical experiments:

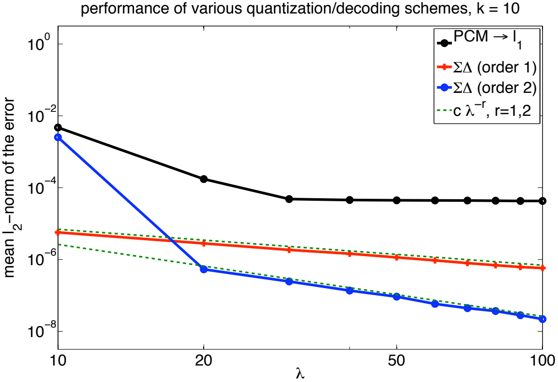

Compressed sensing: Undersampled or oversampled?

Consider

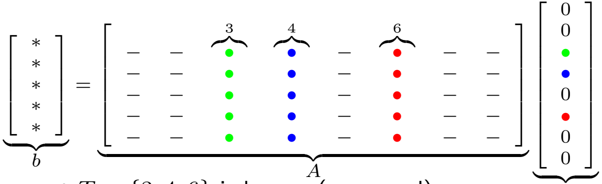

If (once) the support $T = \{3, 4, 6\}$ is known (recovered)

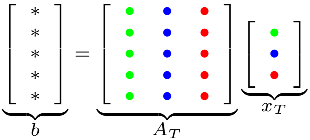

- Rows of A_T are a frame for $\mathbb{R}^k$ with $n > k$ vectors.
- Measurements b_j are associated frame coefficients.
- **When the support is known, this is a redundant frame quantization problem!**

$\Sigma\Delta$ for finite frame expansions

- For the moment we will assume that the support of the signal is known and rely on frame quantization.
- So, we can work with: $x \in \mathbb{R}^k$, $E \in \mathbb{R}^{n\times k}$(with $n > k$), and $F \in \mathbb{R}^{k\times n}$ with $FE = I$ (for now, F is any left-inverse of E).
- Now, suppose $b = Ex$ and quantize b using an rth order $\Sigma\Delta$ scheme to obtain $q_{\Sigma\Delta}$. How well can we do?
- In particular, lets estimate x from $q_{\Sigma\Delta}$ via $\hat{x} = Fq_{\Sigma\Delta}$ using some carefully chosen left-inverse F.

$\Sigma\Delta$ for finite frame expansions

- For the moment we will assume that the support of the signal is known and rely on frame quantization.
- So, we can work with: $x \in \mathbb{R}^k$, $E \in \mathbb{R}^{n\times k}$(with $n > k$), and $F \in \mathbb{R}^{k\times n}$ with $FE = I$ (for now, F is any left-inverse of E).
- Now, suppose $b = Ex$ and quantize b using an rth order $\Sigma\Delta$ scheme to obtain $q_{\Sigma\Delta}$. How well can we do?
- In particular, lets estimate x from $q_{\Sigma\Delta}$ via $\hat{x} = Fq_{\Sigma\Delta}$ using some carefully chosen left-inverse F.

$\Sigma\Delta$ for finite frame expansions

- For the moment we will assume that the support of the signal is known and rely on frame quantization.
- So, we can work with: $x \in \mathbb{R}^k$, $E \in \mathbb{R}^{n\times k}$(with $n > k$), and $F \in \mathbb{R}^{k\times n}$ with $FE = I$ (for now, F is any left-inverse of E).
- Now, suppose $b = Ex$ and quantize b using an rth order $\Sigma\Delta$ scheme to obtain $q_{\Sigma\Delta}$. How well can we do?
- In particular, lets estimate x from $q_{\Sigma\Delta}$ via $\hat{x} = Fq_{\Sigma\Delta}$ using some carefully chosen left-inverse F.

$\Sigma\Delta$ for finite frame expansions

- For the moment we will assume that the support of the signal is known and rely on frame quantization.
- So, we can work with: $x \in \mathbb{R}^k$, $E \in \mathbb{R}^{n\times k}$(with $n > k$), and $F \in \mathbb{R}^{k\times n}$ with $FE = I$ (for now, F is any left-inverse of E).
- Now, suppose $b = Ex$ and quantize b using an rth order $\Sigma\Delta$ scheme to obtain $q_{\Sigma\Delta}$. How well can we do?
- In particular, lets estimate x from $q_{\Sigma\Delta}$ via $\hat{x} = Fq_{\Sigma\Delta}$ using some carefully chosen left-inverse F.

$\Sigma\Delta$ for finite frame expansions

- Recall that $b - q_{\Sigma\Delta} = D^r u$, where

$$D := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ & \ddots & \ddots & \ddots & \\ 0 & \cdots & 0 & -1 & 1 \end{bmatrix}$$

$\Sigma\Delta$ for finite frame expansions

- Recall that $b - q_{\Sigma\Delta} = D^r u$, where

$$D := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ & \ddots & \ddots & \ddots & \\ 0 & \cdots & 0 & -1 & 1 \end{bmatrix}$$

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \leq \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.

$\Sigma\Delta$ for finite frame expansions

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \leq \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

$\Sigma\Delta$ for finite frame expansions

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \leq \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

$\Sigma\Delta$ for finite frame expansions

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \le \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

$\Sigma\Delta$ for finite frame expansions

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \leq \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

ΣΔ for finite frame expansions

- Consequently,

$$x - \hat{x} = FD^r u$$

- and

$$\|x - \hat{x}\|_2 \le \|FD^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy ΣΔ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|FD^r\|_2$: From among all left inverses of E, choose F to minimize $\|FD^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r}E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

$\Sigma\Delta$ for finite frame expansions

- Consequently,

$$x - \hat{x} = F D^r u$$

- and

$$\|x - \hat{x}\|_2 \leq \|F D^r\|_2 \|u\|_\infty \sqrt{n}.$$

- The greedy $\Sigma\Delta$ scheme guarantees that $\|u\|_\infty$ is bounded nicely (by $C_r d$),
- so we are left with controlling $\|F D^r\|_2$: From among all left inverses of E, choose F to minimize $\|F D^r\|_2$! $\implies$ Sobolev duals!
- The Sobolev dual is given by the expression $F := (D^{-r} E)^\dagger D^{-r}$.
- Results from frame theory (Blum, Lammers, Powell, Yılmaz) show that if E obeys a "smoothness" condition, then reconstruction via Sobolev duals yields favorable error guarantees:

$$\|x - \hat{x}\|_2 \lesssim \left(\frac{n}{k}\right)^{-r} d.$$

- Gaussian frames are not smooth! Nevertheless...

Sobolev duals of Gaussian frames

Sobolev duals of Gaussian frames

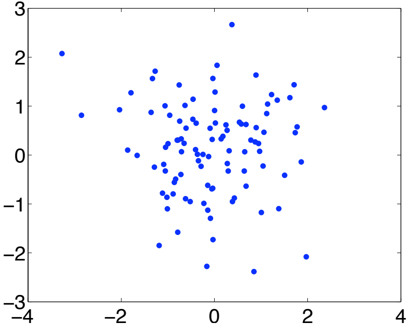

Gaussian frame in $\mathbb{R}^2$

Sobolev duals of Gaussian frames

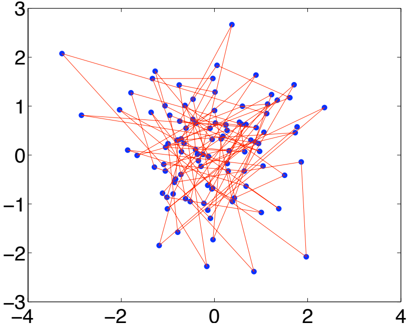

Gaussian frame in $\mathbb{R}^2$

Sobolev duals of Gaussian frames

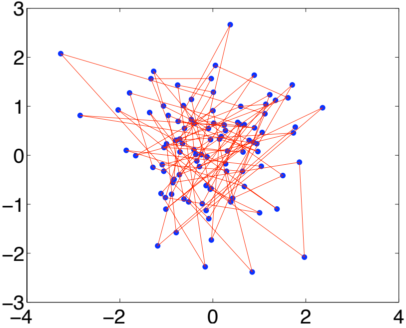

Gaussian frame in $\mathbb{R}^2$

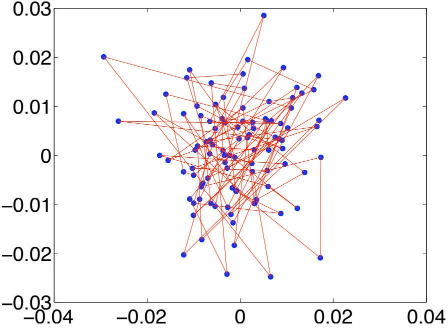

Its canonical dual

Sobolev duals of Gaussian frames

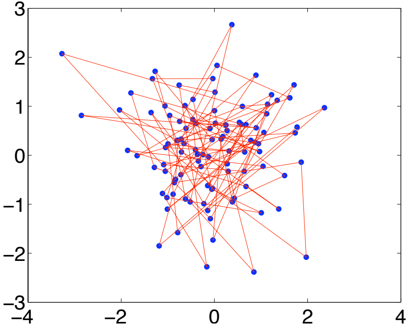

Gaussian frame in $\mathbb{R}^2$

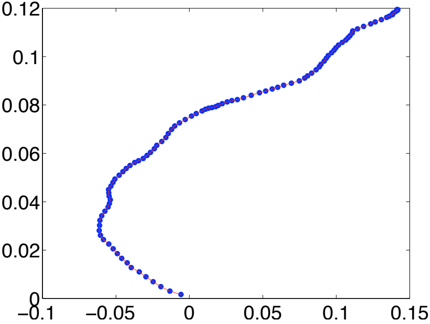

Its Sobolev dual of order $r = 1$

Sobolev duals of Gaussian frames

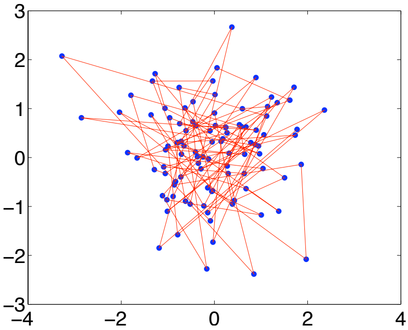

Gaussian frame in $\mathbb{R}^2$

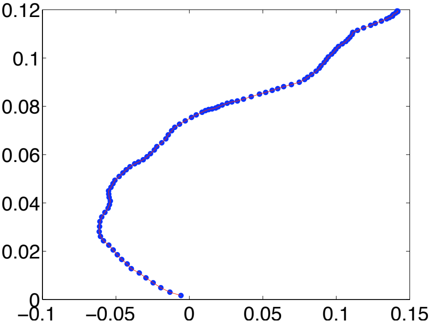

Its Sobolev dual of order $r = 1$

$\Sigma\Delta$ quantization and random frame expansions

Theorem (Theorem 1)

Let E be an $n \times k$ random matrix whose entries are i.i.d. $\mathcal{N}(0,1)$. For any $\alpha \in (0,1)$, if $\lambda \geq c(\log n)^{1/(1-\alpha)}$, then with probability at least $1 - \exp(-c' n^{1-\alpha} k^{\alpha})$,

$$\sigma_{\min}(D^{-r}E) \gtrsim_r (n/k)^{\alpha(r-\frac{1}{2})}\sqrt{n}, \tag{1}$$

which yields the reconstruction error bound

$$\|x - \hat{x}_{\Sigma\Delta}\|_2 \lesssim_r \left(\frac{n}{k}\right)^{-\alpha(r-\frac{1}{2})} d.$$

Proof of Theorem 1

proof outline:

- $\|x-\hat{x}\|_2 \le \|FD^r\|\|u\|_2 = \|(D^{-r}E)^\dagger\|\|u\|_2 = \frac{\|u\|_2}{\sigma_{\min}(D^{-r}E)}$
- $\sigma_{\min}(D^{-r}E) = \sigma_{\min}(U\Sigma V^* E)$
- Weyl's inequality for the singular value estimates, in particular to estimate the singular values of D^{-r} (from the singular values of D^{-1}).
- Unitary invariance of the i.i.d. Gaussian measure: Reduces the problem to estimating $\sigma_{\min}(\Sigma E)$ where Σ is diagonal with Σ_{ii} are estimated as described above.
- Concentration of measure for ΣE: estimate (for a fixed x)

$$\mathbb{P}\{\gamma\|x\|_2 \le \|\Sigma E x\|_2 \le \theta\|x\|_2\}.$$

- Pass to the singular values of ΣE by using a standard net argument.

$\Sigma\Delta$ quantization for compressed sensing: recovery guarantees

The previous theorem states that if the support T of a k-sparse signal is known, the Sobolev dual of A_T can be used in the reconstruction.

- If $\forall j \in T$, $|x_j| > Cd$, then using a robust decoder, support recovery is guaranteed.

Proposition

Let $\|x - x^{\#}\|_2 \leq \eta$, $T = supp(x)$ and $k = |T|$. For any $k' \in \{k, ..., N-1\}$, let T' be the support of the k' largest entries of $x^{\#}$. If $|x_j| > \gamma\eta$ for all $j \in T$, where $\gamma := \left(1 + \frac{1}{k'-k+1}\right)^{1/2}$, then $T' \supset T$.

$\Sigma\Delta$ quantization for compressed sensing: recovery guarantees

The previous theorem states that if the support T of a k-sparse signal is known, the Sobolev dual of A_T can be used in the reconstruction.

- If $\forall j \in T$, $|x_j| > Cd$, then using a robust decoder, support recovery is guaranteed.

Proposition

Let $\|x - x^\#\|_2 \leq \eta$, $T = supp(x)$ and $k = |T|$. For any $k' \in \{k, ..., N-1\}$, let T' be the support of the k' largest entries of $x^\#$. If $|x_j| > \gamma\eta$ for all $j \in T$, where $\gamma := \left(1 + \frac{1}{k'-k+1}\right)^{1/2}$, then $T' \supset T$.

$\Sigma\Delta$ quantization for compressed sensing: recovery guarantees

In light of this we propose the following two-stage algorithm:

1. **Coarse recovery:** any robust decoder applied to $q_{\Sigma\Delta}$ yields an initial, "coarse" approximation $x^{\#}$ of x, and in particular, the exact (or approximate) support T of x.
2. **Fine recovery:** The rth order Sobolev dual of the frame A_T applied to $q_{\Sigma\Delta}$ yields a finer approximation $\hat{x}_{\Sigma\Delta}$ of x.

$\Sigma\Delta$ quantization for compressed sensing: recovery guarantees

Theorem (Theorem 2)

- *A: $n \times N$ matrix whose entries are i.i.d. according to $\mathcal{N}(0,1)$.*
- *$n \geq ck(\log N)^{1/(1-\alpha)}$ where $\alpha \in (0,1)$ and $c = c(r,\alpha)$.*
- $x \in \Sigma_k^N$, $\min_{j \in \mathrm{supp}(x)} |x_j| \geq Cd$

Then with probability at least $1 - \exp(-c' n^{1-\alpha} k^{\alpha})$ on the draw of A:

$$\|x - \hat{x}_{\Sigma\Delta}\|_2 \lesssim_r \left(\frac{n}{k}\right)^{-\alpha(r-\frac{1}{2})} d.$$

Here, c' and C depend only on r.

ΣΔ quantization for compressed sensing

Pros

1. **More accurate** than any known quantization scheme in this setting (even when sophisticated recovery algorithms are employed).
2. **Modular:** If the fine recovery stage is not available or impractical, then the standard (coarse) recovery procedure is applicable as is.
3. **Progressive:** If new measurements arrive (in any given order), noise shaping can be continued on these measurements as long as the state of the system (r real values for an rth order scheme) has been stored.
4. **Universal:** It uses no information about the measurement matrix or the signal.

"Cons"

More computation at the decoder. Extensions to handle non-quantization noise and compressible signals (in progress).

$\Sigma\Delta$ quantization for compressed sensing

Pros

1. **More accurate** than any known quantization scheme in this setting (even when sophisticated recovery algorithms are employed).
2. **Modular:** If the fine recovery stage is not available or impractical, then the standard (coarse) recovery procedure is applicable as is.
3. **Progressive:** If new measurements arrive (in any given order), noise shaping can be continued on these measurements as long as the state of the system (r real values for an rth order scheme) has been stored.
4. **Universal:** It uses no information about the measurement matrix or the signal.

"Cons"

More computation at the decoder. Extensions to handle non-quantization noise and compressible signals (in progress).

ΣΔ quantization for compressed sensing

Pros

1. **More accurate** than any known quantization scheme in this setting (even when sophisticated recovery algorithms are employed).

2. **Modular:** If the fine recovery stage is not available or impractical, then the standard (coarse) recovery procedure is applicable as is.

3. **Progressive:** If new measurements arrive (in any given order), noise shaping can be continued on these measurements as long as the state of the system (r real values for an rth order scheme) has been stored.

4. **Universal:** It uses no information about the measurement matrix or the signal.

"Cons"

More computation at the decoder. Extensions to handle non-quantization noise and compressible signals (in progress).

ΣΔ quantization for compressed sensing

Pros

1. **More accurate** than any known quantization scheme in this setting (even when sophisticated recovery algorithms are employed).
2. **Modular:** If the fine recovery stage is not available or impractical, then the standard (coarse) recovery procedure is applicable as is.
3. **Progressive:** If new measurements arrive (in any given order), noise shaping can be continued on these measurements as long as the state of the system (r real values for an rth order scheme) has been stored.
4. **Universal:** It uses no information about the measurement matrix or the signal.

"Cons"

More computation at the decoder. Extensions to handle non-quantization noise and compressible signals (in progress).

$\Sigma\Delta$ quantization for compressed sensing

Pros

1. **More accurate** than any known quantization scheme in this setting (even when sophisticated recovery algorithms are employed).
2. **Modular:** If the fine recovery stage is not available or impractical, then the standard (coarse) recovery procedure is applicable as is.
3. **Progressive:** If new measurements arrive (in any given order), noise shaping can be continued on these measurements as long as the state of the system (r real values for an rth order scheme) has been stored.
4. **Universal:** It uses no information about the measurement matrix or the signal.

"Cons"

More computation at the decoder. Extensions to handle non-quantization noise and compressible signals (in progress).

Numerical experiments

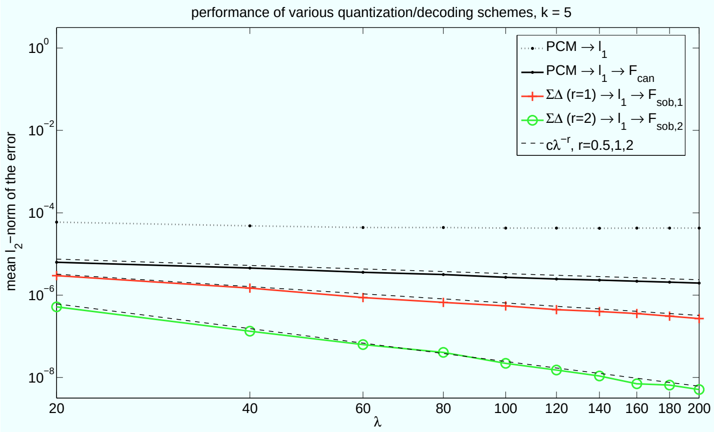

Figure: The average performance of the proposed $\Sigma\Delta$ quantization and reconstruction schemes for $k = 5$. For this experiment the non-zero entries of x are i.i.d. $\mathcal{N}(0,1)$, $N = 2000$ and $d = 10^{-4}$.

Numerical experiments

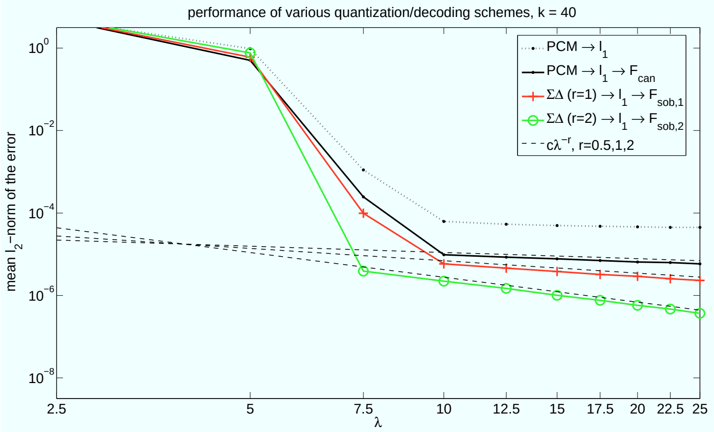

Figure: The average performance of the proposed $\Sigma\Delta$ quantization and reconstruction schemes for $k = 40$. For this experiment the non-zero entries of x are i.i.d. $\mathcal{N}(0,1)$, $N = 2000$ and $d = 10^{-4}$.

Numerical experiments

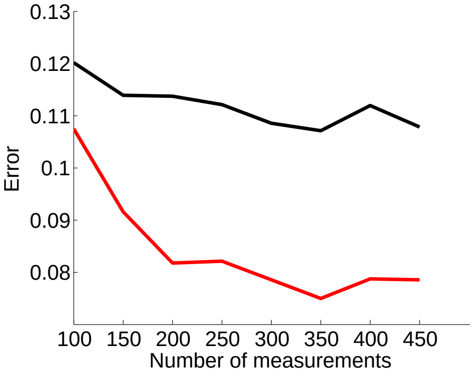

Figure: (Work in progress) The average performance of the proposed $\Sigma\Delta$ quantization and reconstruction scheme for a compressible signal in the presence of non-quantization noise.

Numerical experiments

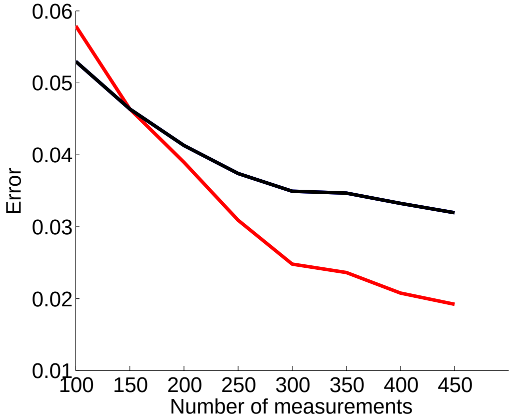

Figure: (Work in progress) The average performance of the proposed $\Sigma\Delta$ quantization and reconstruction schemes for compressible signals (no noise).