

Non-Convex Decoding for $\Sigma\Delta$ -Quantized Compressed Sensing

Evan Chou

Courant Institute of Mathematical Sciences, NYU

Email: chou@cims.nyu.edu

Abstract—Recently Güntürk et al. showed that $\Sigma\Delta$ quantization is more effective than memoryless scalar quantization (MSQ) when applied to compressed sensing measurements of sparse signals. MSQ with the l^1 decoder recovers an approximation to the original sparse signal with an error proportional to the quantization step size δ_Q . For an r -th order $\Sigma\Delta$ scheme the reconstruction accuracy can be improved by a factor of $(m/k)^{\alpha(r-1/2)}$ for any $0 < \alpha < 1$ if $m \gtrsim k(\log N)^{1/(1-\alpha)}$, with high probability on the measurement matrix. The method requires a preliminary support recovery stage for which r cannot be too large and δ_Q must be sufficiently small. In this paper, we remove this requirement, showing that the constrained l^0 and l^τ (for sufficiently small τ) minimization problems subject to a $\Sigma\Delta$ -type quantization constraint would approximate the original signal from the $\Sigma\Delta$ quantized measurements with a comparable reconstruction accuracy. We note that these results allow us to achieve root-exponential reconstruction accuracy while using a fixed quantization alphabet.

I. INTRODUCTION

The robust recovery results in compressed sensing, e.g. [3], [6], [11] showed that sparse vectors could be recovered from compressed sensing measurements even when the measurements are perturbed. Quantization of these measurements introduces such a perturbation from which the robust recovery result allows us to recover.

To fix notation, let N be the ambient dimension of the sparse signal x that we wish to recover. Define the sparsity measure $\|x\|_0 := |\{i : x(i) \neq 0\}|$ and let Σ_k^N be the set of all k -sparse vectors in N dimensions $\Sigma_k^N := \{x \in \mathbb{R}^N : \|x\|_0 \leq k\}$. We will use Φ to denote the $m \times N$ measurement matrix, where we wish to recover x from the quantization of the measurements $y = \Phi x$.

Mathematically, a quantizer maps the measurement space \mathbb{R}^m to a finite set, which we will assume to be of the form \mathcal{A}^m , where the quantization alphabet \mathcal{A} is a finite arithmetic progression of step size δ_Q . For memoryless scalar quantization (MSQ), each measurement is simply rounded to the nearest element of \mathcal{A} . For an r -th order $\Sigma\Delta$ scheme, the quantization is found by solving a difference equation

$$y - q = D^r u \quad (1)$$

for $q \in \mathcal{A}^m$ and $u \in \mathbb{R}^m$, where $\|u\|_\infty$ should be bounded independently of m . D^r is the r -th difference operator: in matrix form, D is 1 on the diagonal and -1 on the subdiagonal, with zeros elsewhere. Note that MSQ corresponds to the case $r = 0$.

The authors in [8] investigated the use of $\Sigma\Delta$ quantization for a specific class of compressed sensing matrices: the random $m \times N$ matrices with each entry drawn independently from the standard Gaussian distribution, $\mathcal{N}(0, 1)$. For MSQ, the quantization introduces an error of at most $\delta_Q/2$ per measurement, and the corresponding recovery error is a constant multiple of δ_Q . For r -th order $\Sigma\Delta$, the quantization introduces an error of at most $2^{r-1}\delta_Q$ per measurement, but the error vector is highly structured. Once the support is recovered, for instance via l^1 minimization, the Sobolev-dual approximation (Equation 6) yields an error of at most $\delta_Q(m/k)^{-\alpha(r-1/2)}$ for some $0 < \alpha < 1$, when $m \gtrsim k(\log N)^{1/(1-\alpha)}$. However, the method to recover the support requires that $5\sqrt{2} \cdot 2^r \delta_Q < \min_{i: x_i \neq 0} |x_i|$ [8]. Thus δ_Q needs to be small, and r cannot be too large.

Suppose we use r -th order $\Sigma\Delta$ with step size δ_Q to produce $q \in \mathcal{A}^m$ with $\|u\|_\infty \leq \mu$. Rearranging Equation 1 shows that $\|D^{-r}(y - q)\|_2 \leq \sqrt{m}\mu$. We will show in Proposition II.2 that the sparsest solution satisfying this quantization constraint

$$x^{0,\mu} := \underset{\|D^{-r}(\Phi z - q)\|_2 \leq \sqrt{m}\mu}{\text{Argmin}} \|z\|_0, \quad (2)$$

approximates the original sparse vector with the same accuracy up to a constant as the Sobolev-dual approximation. Then we will show in Theorem IV.3 that if we solve the non-convex minimization

$$x^{\tau,\mu} := \underset{\|D^{-r}(\Phi z - q)\|_2 \leq \sqrt{m}\mu}{\text{Argmin}} \|z\|_\tau, \quad (3)$$

where $\|z\|_\tau = \left(\sum_{i=1}^N |z(i)|^\tau\right)^{1/\tau}$, then there is a value of $\tau > 0$ sufficiently small so that the minimizer approximates the original sparse vector with the same accuracy up to a constant as the Sobolev-dual approximation. In Section V we note that given a bit budget R for quantizing the measurements, we can now achieve a reconstruction accuracy of the form $\exp(-c(R/k)^\alpha)$ where c is an absolute constant. Previously, Krahmer et al. showed a similar result for $\Sigma\Delta$ quantization of frame coefficients for specially designed frames [9]. Finally in section VI we briefly discuss approaches for tackling the minimization problems.

II. $\Sigma\Delta$ -QUANTIZATION AND SOBOLEV DUAL RECOVERY

Suppose we quantize the measurements $y = \Phi x$ with r -th order $\Sigma\Delta$, i.e. we solve Equation (1) for $q \in \mathcal{A}^m$ and $u \in \mathbb{R}^m$. We highlight two approaches for accomplishing this, where details can be found in [5], [7]:

- A. The simplest greedy method chooses q_{i+1} which would minimize the corresponding value of $|u_{i+1}|$ in the equation. The result is a solution that requires an alphabet of size $2^r + 2\|y\|_\infty/\delta_Q$ and has the bound $\|u\|_\infty \leq \delta_Q/2$.
- B. An alternative method which can be viewed as a greedy method on a different but related difference equation decreases the required alphabet size to

$$C_1 + 2\|y\|_\infty/\delta_Q \quad (4)$$

for some absolute constant C_1 but increases the bound to

$$\|u\|_\infty \lesssim (C_2 r)^r \delta_Q \quad (5)$$

This method allows us to use a fixed quantization alphabet for all r .

Consider any solution with $\|u\|_\infty \leq \mu$. The difference equation can be rewritten as

$$D^{-r}\Phi x - D^{-r}q = u .$$

We review the results from [2] concerning the Sobolev dual. Suppose that an oracle tell us the support T of x . We can then focus on just Φ_T , the $m \times k$ submatrix with columns corresponding to the index set T . Taking the pseudoinverse,

$$x - (D^{-r}\Phi_T)^\dagger D^{-r}q = (D^{-r}\Phi_T)^\dagger u .$$

Note that if $r = 0$, the quantizer is MSQ, and u is the quantization error vector with norm $\sqrt{m}\delta_Q/2$. Taking the pseudoinverse of Φ_T recovers an approximation $\hat{x}^{(0)} = \Phi_T^\dagger q$ with error

$$\|x - \hat{x}^{(0)}\|_2 \leq \|\Phi_T^\dagger\|_2 \sqrt{m}\delta_Q/2 .$$

From the restricted isometry property, the singular values of every submatrix of Φ with $|T| = k$ columns is concentrated around \sqrt{m} with high probability if the entries of Φ are drawn independently from $\mathcal{N}(0, 1)$; so $\|\Phi_T^\dagger\|_2 \sim 1/\sqrt{m}$ and the error bound is proportional to δ_Q and does not decrease with m , as stressed in [8].

For $r > 0$, we see that $\hat{x}^{(r)} = (D^{-r}\Phi_T)^\dagger D^{-r}q$ recovers an approximation with error

$$\|x - \hat{x}^{(r)}\|_2 \leq \frac{\sqrt{m}\mu}{\sigma_{\min}(D^{-r}\Phi_T)} . \quad (6)$$

Note $(D^{-r}\Phi_T)^\dagger D^{-r}$ is precisely the r -th order Sobolev dual of Φ_T . Here we will restate the relevant result from [8, Theorem 3.8] about the smallest singular value:

Theorem II.1. *Let Φ be an $m \times N$ random matrix whose entries are i.i.d. $\mathcal{N}(0, 1)$. Let $0 < \alpha < 1$ and suppose for some $C_3 = C_3(r)$*

$$\frac{m}{s} \geq C_3(\log N)^{1/(1-\alpha)} .$$

Then there exist constants C_4, C_5 depending only on r such that with probability at least $1 - \exp(-C_4 m^{1-\alpha} s^\alpha)$ on the draw of Φ , every $m \times s$ submatrix E of Φ satisfies

$$\sigma_{\min}(D^{-r}E) \geq C_5 \sqrt{m}(m/s)^{\alpha(r-1/2)}$$

This theorem implies that given the support, the error for the Sobolev dual recovery (6) becomes $C_5(m/k)^{-\alpha(r-1/2)}\mu$.

We now use Theorem II.1 to show that solving (2) will recover the support and have an accuracy matching that of the Sobolev dual.

Proposition II.2. *Let Φ be an $m \times N$ random matrix whose entries are i.i.d. $\mathcal{N}(0, 1)$. Let α, m, r , and $s = 2k$ satisfy the conditions of Theorem II.1. Suppose $x \in \Sigma_k^N$ and let q be the quantization of Φx using r -th order $\Sigma\Delta$. Let $\|u\|_\infty \leq \mu$. Then the minimizer $x^{0,\mu}$ of (2) recovers an approximation of x with error*

$$\|x^{0,\mu} - x\|_2 \leq \frac{2}{C_5} \left(\frac{m}{2k}\right)^{-\alpha(r-1/2)} \mu .$$

Proof: Suppose T is the support of x , and let $x' = x^{0,\mu}$ with support T' . Since x, x' are both feasible and x' is the sparsest feasible point, $|T'| = |x'|_0 \leq |x|_0 \leq k$. Then by Theorem II.1,

$$\|x - x'\|_2 \leq \frac{\|D^{-r}\Phi_{T \cup T'}(x - x')\|_2}{C_5 \sqrt{m} \left(\frac{m}{2k}\right)^{\alpha(r-1/2)}} .$$

Note that $D^{-r}\Phi_{T \cup T'}(x - x') = D^{-r}\Phi(x - x')$. Using the triangle inequality and feasibility conditions,

$$\begin{aligned} \|D^{-r}\Phi(x - x')\|_2 &\leq \|D^{-r}(\Phi x' - q)\|_2 + \|D^{-r}(\Phi x - q)\|_2 \\ &\leq 2\sqrt{m}\mu . \end{aligned}$$

The result follows from substitution. \blacksquare

III. ROBUSTNESS OF l^τ MINIMIZATION

We will follow the approaches of [6], [12] to study $\|x\|_\tau$ minimization as stated in (3). As in [6], we will state our results in terms of the condition numbers of submatrices of the measurement matrix:

Definition III.1. Define $a_s(A)$ to be the largest a and $b_s(B)$ to be smallest b such that the following holds:

$$a\|z\|_2 \leq \|Az\|_2 \leq b\|z\|_2 \text{ for all } z \in \Sigma_s^N .$$

The next result combines ideas from the analysis in [6], [12] which will show that constrained l^τ minimization recovers an approximation with error proportional to $1/a$:

Theorem III.2. *Let A be an $m \times N$ matrix, $0 < \tau \leq 1$, and let $x \in \Sigma_k^N, w \in \mathbb{R}^m$ satisfy $\|Ax - w\|_2 \leq \epsilon$. Define $\rho := k/J$ and $\gamma := \frac{b_J(A)}{a_{k+J}(A)}$. If $\gamma\rho^{1/\tau-1/2} < 1$ holds, then the minimizer*

$$x^\sharp := \underset{\|Az - w\|_2 \leq \epsilon}{\text{Argmin}} \|z\|_\tau$$

satisfies the bound

$$\|x^\sharp - x\|_2 \leq \frac{\sqrt{1 + \frac{1}{2/\tau-1} \left(\frac{k}{k+J}\right)^{2/\tau-1}}}{1 - \gamma\rho^{1/\tau-1/2}} \cdot \frac{2\epsilon}{a_{k+J}(A)} .$$

Proof: Define $\eta := x^\sharp - x$, and T to be the support of x . Using Hölder's inequality and the fact that $x^{\tau,\mu}$ is the l^τ minimizer (see (25) of [12]),

$$\|\eta_{T^c}\|_\tau \leq \|\eta_T\|_\tau \leq k^{1/\tau-1/2} \|\eta_T\|_2 . \quad (7)$$

Block η_{T^c} into disjoint blocks of size J of decreasing magnitudes, i.e. $\eta_{T^c} = \sum_{i=1}^L \eta_{T_i}$ with $|T_i| = J$ and $|\eta_{T_i}(j)| \leq |\eta_{T_{i-1}}(j')|$ for $j \in T_i$, $j' \in T_{i-1}$ and $i > 1$. Using the constraint and singular value conditions,

$$\begin{aligned} \|\eta_{T \cup T_1}\|_2 &\leq \frac{1}{a_{k+J}(A)} \|A\eta_{T \cup T_1}\|_2 \\ &\leq \frac{1}{a_{k+J}(A)} \left(\|A\eta\|_2 + \sum_{i=2}^L \|A\eta_{T_i}\|_2 \right) \\ &\leq \frac{2\epsilon}{a_{k+J}(A)} + \gamma \sum_{i=2}^L \|\eta_{T_i}\|_2. \end{aligned} \quad (8)$$

Using 4.2.II of [12], bound $\|\eta_{T_i}\|_2 \leq J^{\frac{1}{2} - \frac{1}{\tau}} \|\eta_{T_{i-1}}\|_\tau$. Combined with the reversed triangle inequality (for $\tau < 1$ and non-negative vectors), we have $\sum_{i=2}^L \|\eta_{T_i}\|_2 \leq J^{\frac{1}{2} - \frac{1}{\tau}} \|\eta_{T^c}\|_\tau$. Finally using (7),

$$\sum_{i=2}^L \|\eta_{T_i}\|_2 \leq \rho^{1/\tau-1/2} \|\eta_{T \cup T_1}\|_2. \quad (9)$$

Combining equation (30) of [12] with (7) gives

$$\|\eta\|_2 \leq \sqrt{1 + \frac{1}{2/\tau - 1} \rho^{2/\tau-1}} \|\eta_{T \cup T_1}\|_2.$$

The result follows from substituting (9) into (8), solving for $\|\eta_{T \cup T_1}\|_2$ and substituting into the last equation. \blacksquare

IV. l^τ MINIMIZATION WITH $\Sigma\Delta$ AND COMPRESSED SENSING

Finally we put together our two main observations and state the known bounds and conditions for recovery. We state precisely the results concerning the singular value of submatrices of $D^{-r}\Phi$ to use with Theorem III.2. We already know that Theorem II.1 covers the smallest singular values of the submatrices. For the largest singular values, we can first use Gershgorin's circle theorem for eigenvalues on $(D^{-1})^T D^{-1}$, to show that $\sigma_{\max}(D^{-1}) \leq \sqrt{m + \frac{(m-1)m}{2}} \leq m$. Then using the bound $\sigma_{\max}(AB) \leq \sigma_{\max}(A)\sigma_{\max}(B)$,

$$\sigma_{\max}(D^{-r}E) \leq m^r \sigma_{\max}(E). \quad (10)$$

The standard restricted isometry property allows us to bound the largest singular value of submatrices of Φ , which has a simple proof in [1]:

Theorem IV.1 (e.g. Theorem 5.2 of [1]). *Let Φ be an $m \times N$ matrix whose entries are i.i.d. $\mathcal{N}(0, 1)$, and suppose*

$$\frac{m}{s} \geq C_6 \log(N/s)$$

for some absolute constant C_6 . Then there exists an absolute constant C_7 such that $b_s(\Phi) < 2\sqrt{m}$ with probability $\geq 1 - 2e^{-C_7 m}$.

We can now combine these two results to obtain upper and lower singular value bounds:

Corollary IV.2. *Let Φ be an $m \times N$ random matrix whose entries are i.i.d. $\mathcal{N}(0, 1)$. Let $0 < \alpha < 1$ and suppose that*

$$\frac{m}{s} \geq C_8 (\log N)^{\frac{1}{1-\alpha}} \quad (11)$$

for $C_8 = \max(C_3, C_6)$ from Theorems II.1 and IV.1. Then for some constant C_9 depending on r , with probability $\geq 1 - 3\exp(-C_9 m^{1-\alpha} s^\alpha)$ the following holds: For all $z \in \Sigma_s^N$,

$$C_5 \sqrt{m} (m/s)^{\alpha(r-1/2)} \|z\|_2 \leq \|D^{-r}\Phi_T z\|_2 \leq 2m^{r+1/2} \|z\|_2$$

where C_5 is from Theorem II.1. In other words,

$$\begin{aligned} a_s(D^{-r}\Phi) &\geq C_5 \sqrt{m} (m/s)^{\alpha(r-1/2)} \\ b_s(D^{-r}\Phi) &\leq 2m^{r+1/2}. \end{aligned}$$

Proof: Note that the condition (11) implies that the conditions for both Theorems II.1 and IV.1 are satisfied. Then with the union bound, both conclusions hold with probability $\geq 1 - \exp(-C_4 m^{1-\alpha} s^\alpha) - 2\exp(-C_7 m)$. Since $m \geq m^{1-\alpha} s^\alpha$, we can bound this by $\geq 1 - 3\exp(-C_9 m^{1-\alpha} s^\alpha)$ with $C_9 = \min(C_4, C_7)$. The conclusion of Theorem II.1 gives the lower inequality (a_s), and the conclusion of Theorem IV.1 along with observation (10) gives the upper inequality (b_s). \blacksquare

The following result is then immediate from Theorem III.2 and Corollary IV.2 using $A = D^{-r}\Phi$, $w = D^{-r}q$, $J = 2k$ and $\epsilon = \sqrt{m}\mu$:

Theorem IV.3. *Let Φ be an $m \times N$ matrix whose entries are i.i.d. $\mathcal{N}(0, 1)$, and let $0 < \alpha < 1$. Suppose for k and $0 < \tau \leq 1$, the following conditions are satisfied:*

- i.
$$\frac{m}{k} \geq 3C_8 (\log N)^{\frac{1}{1-\alpha}}$$
- ii.
$$\frac{1}{\tau} > \frac{1}{2} + \log_2(2/C_5) + r \log_2 m$$

Then with probability $\geq 1 - \exp(-3^\alpha C_9 m^{1-\alpha} k^\alpha)$, the following holds:

For every $x \in \Sigma_k^N$, if r -th order $\Sigma\Delta$ is used to quantize Φx , with q being the quantization and $\|u\|_\infty < \mu$ in the corresponding difference equation, then the minimizer $x^{\tau, \mu}$ of (3) satisfies the bound

$$\|x^{\tau, \mu} - x\|_2 \leq C_{10} \mu \left(\frac{m}{k}\right)^{-\alpha(r-1/2)}$$

for some r -dependent constant C_{10} .

Remark IV.4. Note that for any fixed δ_Q and order r , there is a τ sufficiently small for which the conditions for recovery hold, and in the recovery error μ will typically have a linear dependence on δ_Q .

V. ROOT-EXPONENTIAL ACCURACY

Suppose we impose a bit budget of R bits for quantizing measurements from unit-norm vectors in Σ_k^N . Define $R_{\text{eff}} := R/k$, the effective bit-rate per sparse dimension. We will also work with a fixed quantization alphabet \mathcal{A} of spacing δ_Q . This requires that we use the quantization method (II.B) and that the measurements be bounded independently from the number

of measurements. Unfortunately, Gaussian measurements do not satisfy this criteria [8], but there is also ongoing work that would allow us to use alternative matrix ensembles which are bounded, such as the Bernoulli- $\{\pm 1\}$ matrices [10]. For what follows we will assume the bound $\|y\|_\infty \leq M$ for some absolute constant M .

Also, by inspecting the proofs in [8] we can expand the r -dependent constants in the paper so that in Theorem IV.3, C_8 does not actually depend on r and $C_{10} \leq (C_{11}r)^r$ where C_{11} is now an absolute constant. Substituting (5) for μ in the conclusion of Theorem IV.3 gives a reconstruction accuracy of

$$\delta_Q (C_{12}r^2)^r (3k/m)^{\alpha(r-1/2)}$$

with $C_{12} = C_2 C_{11}$, and the number of bits needed for quantization is $R_{\text{eff}} = C_{13} \frac{m}{k}$ with $C_{13} = \log_2(C_1 + 2M/\delta_Q)$ from (4). Solving for m/k in the rate and substituting, the accuracy becomes

$$\delta_Q (R_{\text{eff}}/C_{13})^{\alpha/2} (C_{14}r^2/R_{\text{eff}}^\alpha)^r$$

with $C_{14} = C_{12}(3C_{13})^\alpha$. Then optimizing over r , or choosing $r = \sqrt{R_{\text{eff}}^\alpha/(eC_{14})}$ gives

$$\delta_Q (R_{\text{eff}}/C_{13})^{\alpha/2} \exp(-C_{15}R_{\text{eff}}^{\alpha/2})$$

with $C_{15} = 1/\sqrt{eC_{14}}$.

VI. ALGORITHMS

Solving the constrained l^τ minimization problem (3) is tricky given the non-convexity of the $\|\cdot\|_\tau$, but there are several approaches. In [11], Saab et al. use a modification to iterative reweighted least squares with encouraging numerical results. If we want a weight that encourages minimization of the sparsity measure $\|x\|_0$ instead, [13] mentions a weighting scheme that is non-separable which could potentially be used in this situation. Other approaches involve projected gradient, and different regularizations of the l^τ norm [4].

We conclude with a sample plot from the approach of [11], which uses the iteration

$$w_i^{(n)} = (|\hat{x}_i^{(n)}|^2 + \epsilon_w)^{\tau/2-1}$$

$$\hat{x}^{(n+1)} = W^{-1}A'(AW^{-1}A' + \lambda I)^{-1}D^{-r}q$$

where $W = W^{(n)}$ is diagonal with entries $w_i^{(n)}$, and $w^{(0)} \equiv 1$. Fixing $\epsilon_w = 10^{-10}$ and $\lambda = 1$, we start with $\tau = 1$ and decrease τ to 0.1. With $N = 200$ and $k = 3$, we generate a k -sparse signal and a 180×200 Bernoulli random matrix. For a range of m , we take the first m measurements, quantize and recover, recording the resulting error. In figure 1 we plot the result, comparing the iterative method with l^1 minimization and with the Sobolev dual (assuming a support oracle). What we observe in many cases is that after a certain number of measurements, the error starts tracking that of the Sobolev dual. In fact, if $w_i^{(n)} \rightarrow \infty$ for $i \notin \text{supp}(x)$, $\hat{x}^{(n+1)}$ converges to a small perturbation of the Sobolev dual reconstruction. Thus the success of the method hinges on a reweighting scheme that can detect the support of the source signal. We

emphasize that with such a coarse quantization step size, l^1 minimization generally will fail to detect the support, a crucial requirement for the results in [8].

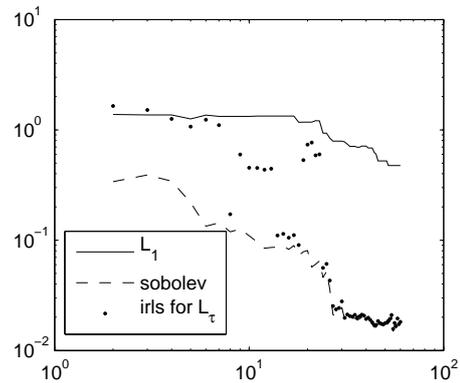


Fig. 1. Log-log plot comparing accuracy vs oversampling ratio m/k for a fixed k -sparse signal and Bernoulli measurements for $r = 1$ and $\delta_Q = 2$. In this example, $N = 200$, and $k = 3$.

ACKNOWLEDGEMENT

The author would like to thank Sinan Güntürk and the reviewers for the valuable discussion and input for this paper.

REFERENCES

- [1] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [2] J. Blum, M. Lammers, A.M. Powell, and Ö. Yılmaz. Sobolev duals in frame theory and sigma-delta quantization. *Journal of Fourier Analysis and Applications*, 16(3):365–381, 2010.
- [3] E.J. Candes, J.K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.
- [4] R. Chartrand. Fast algorithms for nonconvex compressive sensing: Mri reconstruction from very few data. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI '09. IEEE International Symposium on*, pages 262–265, 28 2009-july 1 2009.
- [5] P. Deift, F. Krahermer, and C.S. Güntürk. An optimal family of exponentially accurate one-bit sigma-delta quantization schemes. *Communications on Pure and Applied Mathematics*, 64(7):883–919, 2011.
- [6] S. Foucart and M.J. Lai. Sparsest solutions of underdetermined linear systems via lq-minimization for $0 < q \leq 1$. *Applied and Computational Harmonic Analysis*, 26(3):395–407, 2009.
- [7] C.S. Güntürk. One-bit sigma-delta quantization with exponential accuracy. *Communications on Pure and Applied Mathematics*, 56(11):1608–1630, 2003.
- [8] S. Güntürk, A. Powell, R. Saab, and Ö. Yılmaz. Sobolev duals for random frames and sigma-delta quantization of compressed sensing measurements. *FoCM*, 13(1):1–36, 2013.
- [9] F. Krahermer, R. Saab, and R. Ward. Root-exponential accuracy for coarse quantization of finite frame expansions. *Information Theory, IEEE Transactions on*, 58(2):1069–1079, 2012.
- [10] F. Krahermer, R. Saab, and Ö. Yılmaz. Personal communication, 2013.
- [11] R. Saab, R. Chartrand, and Ö. Yılmaz. Stable sparse approximations via nonconvex optimization. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 3885–3888. IEEE, 2008.
- [12] R. Saab and Ö. Yılmaz. Sparse recovery by non-convex optimization—instance optimality. *Applied and Computational Harmonic Analysis*, 29(1):30–48, 2010.
- [13] D. Wipf, S. Nagarajan, et al. Solving sparse linear inverse problems: Analysis of reweighted l1 and l2 methods. In *SPARS-Signal Processing with Adaptive Sparse Structured Representations*, 2009.