QUASI-SPARSEST SOLUTIONS FOR QUANTIZED COMPRESSED SENSING BY GRADUATED-NON-CONVEXITY BASED REWEIGHTED ℓ_1 MINIMIZATION

Ines Elleuch, Fatma Abdelkefi*, Mohamed Siala*, Ridha Hamila†, Naofal Al-Dhahir‡*

* MEDIATRON Laboratory, SUP'COM, University of Carthage, Tunisia
† Electrical Engineering Department, Qatar University, Qatar
‡ Electrical Engineering Department, University of Texas at Dallas, USA

ABSTRACT

In this paper, we address the problem of sparse signal recovery from scalar quantized compressed sensing measurements, via optimization. To compensate for compression losses due to dimensionality reduction and quantization, we consider a cost function that is more sparsity-inducing than the commonly used ℓ_1-norm. Besides, we enforce a quantization consistency constraint that naturally handles the saturation issue. We investigate the potential of the recent Graduated-Non-Convexity based reweighted ℓ_1-norm minimization for sparse recovery over polyhedral sets. We demonstrate by simulations, the robustness of the proposed approach towards saturation and its significant performance gain, in terms of reconstruction accuracy and support recovery capability.

***Index Terms*—** Quantized Compressed Sensing, Concave Approximation, Graduated-Non-Convexity, Reweighted ℓ_1, Support Recovery

1. INTRODUCTION AND RELATED WORKS

Recovery of the sparsest signal representation from an underdetermined system of linear equations, which arises in Compressed Sensing (CS) theory, is a well-studied problem for which convex relaxation by ℓ_1-norm minimization provides (1) an exact solution via Basis Pursuit (BP), provided the sparsest solution is unique, and (2) a stable solution via Basis Pursuit De-Noising (BPDN), whenever the measurements are corrupted by noise of bounded energy [1].

When addressing the practical situation of quantized measurements, the saturation phenomenon introduces large and potentially unbounded errors, and the stability property of BPDN does no longer hold. Besides, the convexification approach via the ℓ_1-norm leads to a significantly bad approximation of the ℓ_0 counting "norm". However, in the literature on Quantized Compressed Sensing (QCS), optimization based reconstruction is commonly addressed from an ℓ_1-minimization perspective [2]-[7].

This paper was made possible by grant number NPRP 06-070-2-024 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

Fortunately, the literature on CS provides more sparsity-inducing alternatives to the ℓ_1-norm cost function, enjoying nice properties (differentiability, concavity, etc) and amenable to use tractable reconstruction algorithms. For instance, the Maximally Sparse Convex (MSC) approach in [8], deals with sparsity-penalized least squares using a non-convex sparsity-inducing regularizer that ensures the convexity of the total cost function. The Iterative-Reweighted-Least-Squares method (IRLS) presented in [9] is nothing but a Majorization-Minimization (MM) scheme that employs a logarithmic approximation of the "ℓ_0-norm" upper-bounded by a quadratic surrogate function. The reweighted-ℓ_1 minimization of [10], is also a MM procedure using an ϵ-regularized logarithmic approximation and a linear surrogate function. In [11], the author proposes an exponential concave approximation that more closely approximate the ℓ_0-norm and uses a linear surrogate function, yielding another reweighted-ℓ_1 scheme.

Recently, a *Graduated-Non-Convexity* (GNC) approach [12], also known as *continuation approach*, which is an iterative optimization scheme involving a family of approximation functions with gradual approximation accuracy / smoothness trade-off, has been proposed to solve the sparse recovery problem [13, 14, 15]. The GNC methodology has the merit of converging to a local solution of the ℓ_0-norm itself. In [14], a class of differentiable approximation functions, amenable to use a gradient projection procedure at each iteration, is proposed. More recently, the concave approximation of [11] has been integrated into a GNC-based procedure in [13, 15], to solve the sparse recovery problem under a linear constraint and over a polyhedral set, respectively.

In this paper, we address the problem of sparse recovery from scalar quantized and possibly saturated CS measurements, under a Quantization Consistency (QC) constraint, by considering the GNC-based reweighted ℓ_1 minimization approach of [13, 15]. The contribution of this paper is threefold. First, the proposed method is essentially a reweighted ℓ_1-minimization technique, thus computationally attractive, and more sparsity-inducing than the commonly used ℓ_1-minimization in QCS. Second, the QC constraint yields a built-in capability to handle the saturation phenomenon, in the same way of [3, 7]. Third, as the proposed algorithm is

978-0-9928-6265-7/16/$31.00 ©2016 IEEE

essentially minimizing the ℓ_0-norm, even though the sparsest solution is not unique, we study the uniqueness of its support, via simulations. The reason behind this motivation is that the feasible solution space defined by the QC constraint is not an arbitrary polyhedron. It should be noted that the problem of support recovery from QCS measurements has been mostly addressed in the case of one-bit CS [16, 17], under additional constraints on the design of the measurement matrix. In [16], the case where the measurements are corrupted by Gaussian noise and outliers is also considered. For the case of multi-bit QCS, the authors in [18] derived lower bounds on the required number of measurements for exact support recovery from scalar quantized and noisy measurements using a Maximum-Likelihood decoder. In [19], support recovery for $\Sigma\Delta$-QCS is addressed using an ℓ_1 minimization approach.

The paper is organized as follows. Section II presents the observation model and the optimization problem to be considered throughout the paper. Section III describes the algorithmic framework to solve the proposed optimization problem. Section IV provides simulations results that demonstrate the significant performance gain of the proposed approach, both in terms of quality of the estimation and support recovery capabilities. Finally, Section V concludes the paper.

2. OBSERVATION MODEL AND PROBLEM FORMULATION

2.1. QCS Observation Model

In this paper, we assume that CS measurements are scalar quantized, yielding the observation model

$$\mathbf{z} = \mathbf{\Phi x}, \quad \mathbf{y} = \mathcal{Q}_b(\mathbf{z}), \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^N$ is a K-sparse signal, $\mathbf{\Phi} \in \mathbb{R}^{M\times N}$ is the measurement matrix with $M < N$, and $\mathcal{Q}_b(\cdot)$ is the quantization function that maps real-valued vectors, element-wise, to a set of 2^b output levels $\{\mathfrak{q}_i, i = 1, \dots, 2^b\}$, according to a set of thresholds $\{\tau_i, i = 1, \dots, 2^{b-1}\}$, with b denoting the quantizer precision in terms of bit-depth. Formally, we have

$$y_m = \begin{cases} \mathfrak{q}_1 & \text{if} \quad z_m < \tau_1, \\ \mathfrak{q}_i & \text{if} \quad z_m \in [\tau_{i-1}, \tau_i),\ i = 2, \dots, 2^{b-1}, \\ \mathfrak{q}_{2^b} & \text{if} \quad z_m \geq \tau_{2^{b-1}}. \end{cases} \tag{2}$$

We will consider a uniform midrise quantizer i.e. the thresholds are equally spaced with step size δ, and the levels are centered between the thresholds. The dynamic range of the quantizer is $[-g, g]$ where $g \triangleq 2^{b-1}\delta$ is the saturation level. Measurements of magnitude larger than $g{-}\delta$ saturate at $\pm(g - \frac{\delta}{2})$.

2.2. Concave Approximation of the ℓ_0-norm

The continuous ℓ_0 approximation function, adopted in this paper, was originally proposed in [11]. By introducing a non-negative parameter σ, we define:

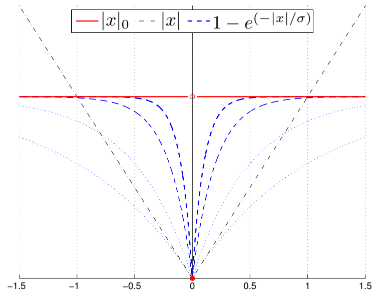

Fig. 1: ℓ_0 approximations. For $f_\sigma(\cdot)$, $\sigma = 0.1, 0.2, 0.5, 1$.

$$F_\sigma(\mathbf{x}) = \sum_{i=1}^{N} f_\sigma(x_i) = \sum_{i=1}^{N}(1 - e^{-|x_i|/\sigma}) \tag{3}$$

The above approximation enjoys the following properties:

- $F_\sigma(\mathbf{x}) \leq \|\mathbf{x}\|_0$,
- $\lim_{\sigma \to 0+} F_\sigma(\mathbf{x}) = \|\mathbf{x}\|_0$,
- $F_\sigma(\mathbf{x}) = F_\sigma(-\mathbf{x})$,
- F_σ is non-decreasing and concave in the positive orthant.

As shown in Fig 1, this function better approximates the ℓ_0-norm and is likely to be more sparsity-inducing than the commonly used ℓ_1-norm.

2.3. Optimization Problems Formulation

We propose to find a sparsest consistent solution by considering the QC-constrained optimization problem:

$$\min_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{x}\|_0 \quad s.t. \quad \mathcal{Q}_b(\mathbf{\Phi x}) = \mathbf{y}, \tag{4}$$

denoted ℓ_0-QC. As we are dealing with a scalar quantization scheme, the QC constraint is obviously a linear inequality constraint. To be precise, let $\mathbf{\Phi}_+$ and $\mathbf{\Phi}_-$ denote row submatrices of $\mathbf{\Phi}$, corresponding to the positively saturated and the negatively saturated measurements, respectively and let $\mathbf{\Phi}_*$ be its row submatrix corresponding to the remaining measurements. Then, it can be easily verified that the QC constraint is equivalent to the following linear inequality:

$$\mathcal{Q}_b(\mathbf{\Phi x}) = \mathbf{y} \quad \Leftrightarrow \quad \overline{\mathbf{\Phi}}\mathbf{x} \leq \overline{\mathbf{y}}, \tag{5}$$

where $\overline{\mathbf{\Phi}} = [\mathbf{\Phi}_*^T, -\mathbf{\Phi}_*^T, -\mathbf{\Phi}_+^T, \mathbf{\Phi}_-^T]^T \in \mathbb{R}^{(2M-S)\times N}$ with S being the number of saturated measurements, and $\overline{\mathbf{y}}$ is the vector of thresholds with respect to the measurements.
Then, the ℓ_0-QC problem in (4) could be written as:

$$\min_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{x}\|_0 \quad s.t. \quad \overline{\mathbf{\Phi}}\mathbf{x} \leq \overline{\mathbf{y}}, \tag{6}$$

where the problem at hand is to recover a sparsest solution over a polyhedral set.

By considering the approximation function F_σ, the ℓ_0-QC problem could be approximated by :

$$\min_{\mathbf{x}\in\mathbb{R}^N} F_\sigma(|\mathbf{x}|) \quad s.t. \quad \overline{\mathbf{\Phi}}\mathbf{x} \leq \overline{\mathbf{y}}, \tag{7}$$

3. GRADUATED-NON-CONVEXITY BASED REWEIGHTED ℓ_1 MINIMIZATION FOR QCS

3.1. QCS by Reweighted ℓ_1 Minimization

In order to take advantage of the concavity of F_σ in the positive orthant, we introduce the variable conversion suggested in [15], $\mathbf{u} = [\mathbf{x}^+, \mathbf{x}^-] \in \mathbb{R}_+^{2N}$, where $\mathbf{x}^+ = \max(\mathbf{x}, \mathbf{0})$ and $\mathbf{x}^- = -\min(\mathbf{x}, \mathbf{0})$. Then, it can be verified that the approximate problem (7) can be equivalently formulated as:

$$\min_{\mathbf{u}\in\mathbb{R}^{2N}} F_\sigma(\mathbf{u}) \quad s.t. \quad \mathbf{u} \in \mathcal{P}, \tag{8}$$

where the objective function is concave and bounded from below over the feasible solution set

$$\mathcal{P} = \{\mathbf{u} \in \mathbb{R}^{2N}, \quad [\overline{\mathbf{\Phi}}, -\overline{\mathbf{\Phi}}]\mathbf{u} \leq \bar{\mathbf{y}}, \quad \mathbf{u} \succeq \mathbf{0}\}.$$

A common way to solve a concave optimization problem, is to apply a Majorize-Minimize (MM) iterative procedure where the surrogate function is constructed from the linearization of the objective function around the current iterate. More precisely, to solve (8), we proceed iteratively by majorizing $F_\sigma(\mathbf{u})$ by $\langle \nabla F_\sigma(\mathbf{u}^k), \mathbf{u} - \mathbf{u}^k \rangle$, based on its concavity. Hence solving (8) amounts to iteratively solve the linear program:

$$\mathbf{u}^{k+1} = \underset{\mathbf{u}\in\mathbb{R}^{2N}}{\operatorname{argmin}} \langle \nabla F_\sigma(\mathbf{u}^k), \mathbf{u} \rangle \quad s.t. \quad \mathbf{u} \in \mathcal{P}. \tag{9}$$

Problem (9) could be cast into the following reweighted ℓ_1-norm minimization:

$$\mathbf{x}^{k+1} = \underset{\mathbf{x}\in\mathbb{R}^{N}}{\operatorname{argmin}} \|\mathbf{W}\mathbf{x}\|_1 \quad s.t. \quad \overline{\mathbf{\Phi}}\mathbf{x} \leq \bar{\mathbf{y}}, \tag{10}$$

where $\mathbf{W} = \operatorname{diag}(\nabla F_\sigma(|\mathbf{x}^k|))$.

3.2. QCS by GNC-based Optimization

Solving the ℓ_0-QC problem by considering a single approximation function F_σ is not very judicious. Indeed, the choice of σ would either favor approximation accuracy at the price of extra local minima or vice versa.

3.2.1. GNC General Framework

Graduated-Non-Convexity is a deterministic global optimization methodology that addresses non-convex cost functions by considering a family of approximation functions scaled by a parameter σ that monotonically controls the tradeoff between approximation accuracy and smoothness. When σ grows towards ∞, the family should embed a convex approximation function. Then, as σ decreases, it leads continuously to the original cost function with an increasing "non-convexity" rate i.e. an increasing number of local minima. The motivation behind such a construction is essentially to avoid these annoying local minima by: (1) gradually decreasing the scale parameter so that the global minimizers of two consecutive problems are sufficiently close, more precisely, the former should fall within the locally convex region around the latter, the so-called basin of attraction, and (2) each approximation problem is initialized with the minimizer of the previous one, so that a descent method would converge to one of its global minimizer. Clearly, the initialization issue and the decreasing rule for σ are crucial for the effectiveness of such an optimization methodology.

3.2.2. GNC-based scheme for ℓ_0-QC

F_σ minimization is equivalent to the convex ℓ_1-minimization for σ sufficiently large and to the ℓ_0-minimization for σ sufficiently small [13]. Then, ℓ_0-QC could be solved using a GNC-based procedure involving the family of approximation functions F_σ. Moreover, a solution of ℓ_1-QC is potentially a good starting point. The decreasing rule for σ is inspired from the exposition suggested in [13], with respect to the behaviour of the function f_σ. To be more precise,

$$f_\sigma(x) \simeq \begin{cases} 1, & \text{if} \quad x \gg \sigma, \\ 0, & \text{if} \quad x \ll \sigma. \end{cases} \tag{11}$$

Let $\mathbf{x}^j$ be the current iterate, i.e a solution of $F_{\sigma^{j-1}}$ minimization, and let $I = \{i \mid |x_i^j| < \sigma^{j-1}\}$. It is reasonable to take I as an estimation of the zero components of $\mathbf{x}^j$. To ensure a gradual decrease of σ, yielding a small number of GNC iterations, we propose the following adaptive update rule:

$$\sigma^j = \operatorname{mean}(\mathbf{x}_I) + \operatorname{std}(\mathbf{x}_I), \tag{12}$$

where std denotes the standard deviation, so as to ask whether the most significant components of $\mathbf{x}_I$ are really zero components. The QCS reconstruction scheme, involving a GNC-based outer loop and a reweighted ℓ_1 minimization inner loop, is summarized in Algorithm 1.

Algorithm 1 GNC-based Reweighted ℓ_1 for QCS

Require: $\mathbf{\Phi}, \mathbf{y}, \epsilon_1, \epsilon_2, \sigma^0, \sigma_{min}$.
compute $\overline{\mathbf{\Phi}}$ according to (5).
$\mathbf{x}^0 = \operatorname{argmin}_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{x}\|_1 \; s.t. \; \overline{\mathbf{\Phi}}\mathbf{x} \leq \bar{\mathbf{y}}$.
$j = 0$.
while $d_1 > \epsilon_1$ and $\sigma^j > \sigma_{min}$ **do** ▷ GNC outer loop
  $j = j + 1, k = 0, \widetilde{\mathbf{x}}^0 = \mathbf{x}^j$.
  while $d_2 > \epsilon_2$ **do** ▷ reweighted ℓ_1 inner loop
    $k = k + 1$.
    $\mathbf{W} = \operatorname{diag}(\nabla F_{\sigma^{j-1}}(|\widetilde{\mathbf{x}}^{k-1}|))$.
    $\widetilde{\mathbf{x}}^k = \operatorname{argmin}_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{W}\mathbf{x}\|_1 \quad s.t. \quad \overline{\mathbf{\Phi}}\mathbf{x} \leq \bar{\mathbf{y}}$.
    $d_2 = \|\widetilde{\mathbf{x}}^k - \widetilde{\mathbf{x}}^{k-1}\| / \|\widetilde{\mathbf{x}}^{k-1}\|$.
  end while
  $\mathbf{x}^j = \widetilde{\mathbf{x}}^k, d_1 = \|\mathbf{x}^j - \mathbf{x}^{j-1}\| / \|\mathbf{x}^{j-1}\|$.
  $I = \{i \mid |x_i^j| < \sigma^{j-1}\}, \sigma^j = \operatorname{mean}(\mathbf{x}_I) + \operatorname{std}(\mathbf{x}_I)$.
end while

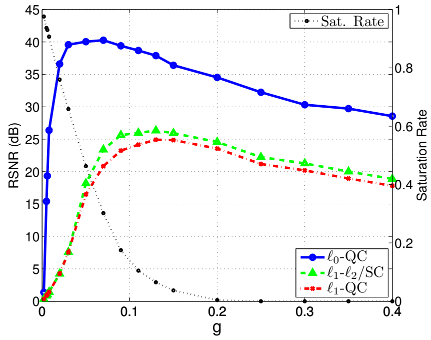

Fig. 2: Average RSNR versus the saturation level g, with $K = 10$, $M = 250$, $N = 500$ and $b = 4$.

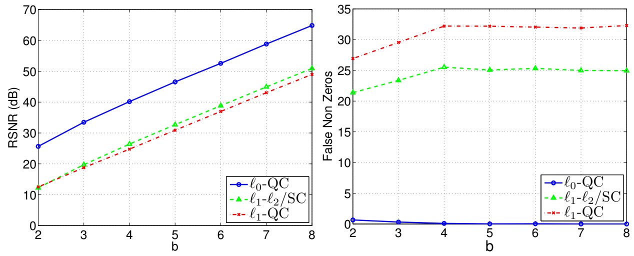

(a) RSNR (b) False Non-Zeros

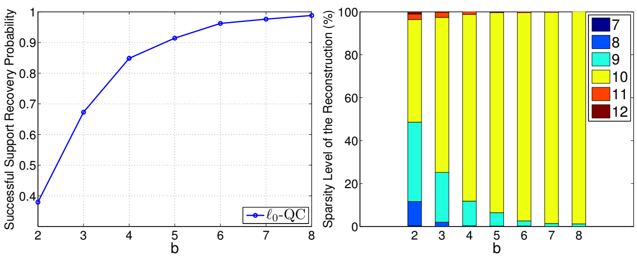

(c) Support Recovery (d) Support Size Recovery (ℓ_0-QC)

Fig. 3: Reconstruction performance versus the bit-depth b, with $K = 10$, $M = 250$ and $N = 500$.

4. SIMULATION RESULTS

In this section, we demonstrate the performance gain of the proposed ℓ_0-QC recovery approach over two benchmark ℓ_1 recovery approaches that handle saturation errors. More precisely, we consider the ℓ_1-QC method [3] and the Saturation Consistency (SC) method, denoted here ℓ_1-ℓ_2/SC, [6]. The ℓ_1-ℓ_2/SC method decouples the measurements and uses a SC constraint on the saturated measurements and an ℓ_2-norm constraint over the remaining ones.

All the methods were implemented using the general-purpose convex optimization package CVX [20], which is an interior-point method solver. For each trial, a new $M \times N$ measurement matrix $\mathbf{\Phi}$ is randomly generated with its entries drawn from a zero-mean, $\frac{1}{M}$-variance Gaussian distribution $\mathcal{N}(0, \frac{1}{M})$. For each realization of the K-sparse signal $\mathbf{x}$, the support is selected uniformly at random from all possible supports of cardinality K, the non-zero entries are drawn independently from $\mathcal{N}(0, 1)$, and $\mathbf{x}$ is scaled to unit ℓ_2-norm. For the sake of comparison, we consider the ϵ-support of the reconstructed signal $\widehat{\mathbf{x}}$, defined as $\epsilon\text{-supp}(\widehat{\mathbf{x}}) = \{i \mid |x_i| > \epsilon\}$ where ϵ is a small value used to prune out negligible non-zeros, and we consider the following performance metrics:

- The Reconstruction Signal-to-Noise-Ratio RSNR $\triangleq -20\log_{10}(\|\mathbf{x} - \widehat{\mathbf{x}}\|_2)$, where $\widehat{\mathbf{x}}$ is the reconstructed signal;
- The number of False Non-Zeros (FNZ) and the number of False Zeros (FZ);
- The probability of successful support recovery i.e. the probability of the event SE $= 0$, where SE $\triangleq$ FZ + FNZ stands for Support Error.

For all experiments, we set $\epsilon = 10^{-7}$, $\epsilon_1 = 10^{-4}$, $\epsilon_2 = 10^{-3}$, $\sigma^0 = 2\max(|\mathbf{x}|)$ and $\sigma_{min} = 10^{-7}$. We use an oracle bound for the ℓ_2-norm constraint within the ℓ_1-ℓ_2/SC method, namely $\|\mathbf{y}_{\mathcal{T}} - \mathbf{z}_{\mathcal{T}}\|_2$, where $\mathcal{T}$ is the support of non-saturated measurements i.e. $\mathcal{T} = \{i, |y_i| < g - \frac{\delta}{2}\}$.

In the first experiment, we set $K = 10$, $M = 250$, $N = 500$, $b = 4$ and we tune the saturation rate by varying the saturation level g over the range $]0, 0.4]$. For the proposed ℓ_0-QC approach, the maximal inner iterations number is set to 10, in order to tradeoff computational complexity at sub-optimal saturation rates. Figure 2 depicts the RSNR performance and the saturation rate, indicated on the left vertical axis and the right vertical axis, respectively. Results are averaged over 100 trials. As shown, all the methods reach their optimal performances at non-zero saturation rates. The proposed approach provides a substantial reconstruction performance gain (15 dB) with a higher saturation rate at its optimal operating point.

In the second experiment, we examine the impact of the bit-depth on the reconstruction performance. Figure 3 reports the average performance over 500 trials, at the optimal saturation rate for each method, in terms of (a) RSNR and (b) the number of FNZ. For the proposed ℓ_0-QC method, we also depicts (c) the successful support recovery rate and (d) the histograms of support sizes (sparsity levels) of the recovered signals. As expected, the proposed approach shows an interesting support recovery capability. Indeed, the search space induced by the QC constraint is a N-dimensional polyhedron included within the intersection of many slabs (a slab is the region between two parallel hyperplanes) defined by the unsaturated measurements. For a given bit-depth, a high saturation rate induces a small step size δ, and consequently thinner slabs are involved in the definition of the polyhedral solution set. Thus, even though the polyhedron is unbounded, it is likely to intersect a single K-dimensional subspace, given a small δ. Of course, the step size δ could not be arbitrary small to ensure a sufficient number of unsaturated measurements and equivalently a sufficient number of slabs. As illustrated in Figure 3d, the proposed approach provide a quasi-sparsest consistent solution of support size K ($\mathbf{x}$ is a solution of ℓ_0-QC) with an increasing probability with the bit-precision. Support sizes slightly greater than K $(11, 12)$ reflect local minima caused by sub-optimal σ-initialization and update rule, and stopping rules. Support sizes less than K $(7 - 9)$ are rather related to very small non-zero components in $\mathbf{x}$, with comparaison to δ. It is also to be mentioned that the proposed method requires around 3.2 GNC iterations on average, for all bit-depth settings.

5. CONCLUSION

We presented an efficient method to recover a sparse signal from scalar QCS measurements, by considering the set of sparsest consistent solutions. A quasi-sparsest solution was estimated using an iterative GNC-based reweighted ℓ_1 minimization approach that consider a concave approximation of the ℓ_0-norm, with an increasing accuracy. We demonstrated the performance gain of the proposed approach, with comparison to two benchmark ℓ_1 minimization based QCS sparse recovery methods, both in terms of reconstruction accuracy and probability of successful support recovery, via simulations.

REFERENCES

[1] E. J. Candes, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.

[2] A. Zymnis, S. Boyd, and E. Candes, "Compressed sensing with quantized measurements," *IEEE Sig. Proces. Letters*, vol. 17, no. 2, pp. 149–152, 2010.

[3] W. Dai and O. Milenkovic, "Information theoretical and algorithmic approaches to quantized compressive sensing," *IEEE trans. Commun.*, vol. 59, no. 7, pp. 1857–1866, 2011.

[4] L. Jacques, D. Hammond, and J. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Trans. Inf. Theory*, vol. 57, no. 1, pp. 559–571, Jan. 2011.

[5] Z. Yang, L. Xie, and C. Zhang, "Variational bayesian algorithm for quantized compressed sensing," *IEEE Trans. Sig. Proces.*, vol. 61, no. 11, pp. 2815–2824, 2013.

[6] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing," *Appl. Comput. Harmon. Anal.*, vol. 31, no. 3, pp. 429–443, 2011.

[7] A. Moshtaghpour, L. Jacques, V. Cambareri, K. Degraux, and C. D. Vleeschouwer, "Consistent basis pursuit for signal and matrix estimates in quantized compressed sensing," *IEEE Sig. Process. Letters*, vol. 23, no. 1, pp. 25–29, 2016.

[8] I. W. Selesnick and I. Bayram, "Sparse signal estimation by maximally sparse convex optimization," *IEEE Trans. Sig. Process.*, vol. 62, no. 5, pp. 1078–1092, 2014.

[9] B. D. Rao and K. Kreutz-Delgado, "An affine scaling methodology for best basis selection," *IEEE Trans. Signal Processing*, vol. 47, no. 1, pp. 187–200, 1999.

[10] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier analysis and applications*, vol. 14, no. 5-6, pp. 877–905, 2008.

[11] O. Mangasarian, "Machine learning via polyhedral concave minimization," in *Applied Mathematics and Parallel Computing*. Springer, 1996, pp. 175–188.

[12] A. Blake and A. Zisserman, *Visual reconstruction*. MIT press Cambridge, 1987, vol. 2.

[13] T. Migot and M. Haddou, "A smoothing method for sparse optimization over polyhedral sets," in *Model. Comput. Opt. Inf. Syst. Manag. Sci.* Springer, 2015, pp. 369–379.

[14] H. Mohimani, M. Babaie-Zadeh, and C. Jutten, "A fast approach for overcomplete sparse decomposition based on smoothed norm," *IEEE Trans. Signal Processing*, vol. 57, no. 1, pp. 289–301, 2009.

[15] M. Malek-Mohammadi, A. Koochakzadeh, M. Babaie-Zadeh, M. Jansson, and C. R. Rojas, "Successive concave sparsity approximation: Near-oracle performance in a wide range of sparsity levels," *arXiv preprint arXiv:1505.06841*, 2015.

[16] J. Haupt and R. Baraniuk, "Robust support recovery using sparse compressive sensing matrices," in *45th Annual Conf. Inf. Sci. Syst. (CISS)*, 2011.

[17] S. Gopi, P. Netrapalli, P. Jain, and A. Nori, "One-bit compressed sensing: Provable support and vector recovery," in *30th Inter. Conf. Mach. Learn. (ICML)*, 2013.

[18] T. Wimalajeewa and P. K. Varshney, "Performance bounds for sparsity pattern recovery with quantized noisy random projections," *IEEE Journal of Selected Topics in Signal Processing*, vol. 6, no. 1, pp. 43–57, 2012.

[19] C. S. Güntürk, M. Lammers, A. M. Powell, R. Saab, and Ö. Yılmaz, "Sobolev duals for random frames and $\Sigma\Delta$ quantization of compressed sensing measurements," *Foundations of Computational mathematics*, vol. 13, no. 1, pp. 1–36, 2013.

[20] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.0," available online at http://cvxr.com/cvx, Aug. 2012.