

1-1-2016

Methods for Quantized Compressed Sensing

Hao-Jun Michael Shi

University of California, Los Angeles

Mindy Case

University of California, Los Angeles

Xiaoyi Gu

University of California, Los Angeles

Shenyinying Tu

University of California, Los Angeles

Deanna Needell

Claremont McKenna College

Recommended Citation

H. M. Shi, M. Case, X. Gu, S. Tu, and D. Needell. "Methods for Quantized Compressed Sensing." Proc. Information Theory and Applications (ITA), La Jolla CA, Jan. 2016.

This Conference Proceeding is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

Methods for Quantized Compressed Sensing

Hao-Jun Michael Shi*, Mindy Case*, Xiaoyi Gu*, Shenyinying Tu*, Deanna Needell†

January 1, 2016

Abstract

In this paper, we compare and catalog the performance of various greedy quantized compressed sensing algorithms that reconstruct sparse signals from quantized compressed measurements. We also introduce two new greedy approaches for reconstruction: Quantized Compressed Sampling Matching Pursuit (QCoSaMP) and Adaptive Outlier Pursuit for Quantized Iterative Hard Thresholding (AOP-QIHT). We compare the performance of greedy quantized compressed sensing algorithms for a given bit-depth, sparsity, and noise level.

1 Introduction

Compressed sensing (CS) is an emergent linear sampling framework that enables reconstruction of sparse signals from a small number of linear measurements relative to the total dimension of the signal space. In particular, given the acquired compressed signal $\mathbf{y} \in \mathbb{R}^M$ and measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, one seeks to reconstruct the signal $\mathbf{x} \in \mathbb{R}^N$ by solving the (possibly noisy) underdetermined linear system $\mathbf{y} = \Phi \mathbf{x}$. Candès, et. al. [CRT06] demonstrated that K -sparse signals, i.e. \mathbf{x} that satisfy $\|\mathbf{x}\|_0 = |\text{supp}(\mathbf{x})| \leq K$, may be robustly reconstructed by an ℓ_1 -minimization program if Φ satisfies the restricted isometry property (RIP) [CT05]. Random matrices whose entries are chosen according to an appropriately chosen i.i.d. distribution and random submatrices of structured matrices have been shown to satisfy the RIP with high probability [RV08].

However, classic compressive sensing assumes that the measurements are continuous and real-valued. In practice, real measurements must be *quantized*, or mapped to a discrete value from some finite set. In addition, in real-world applications, severe quantization may be preferred since low-bit measurements tend to be more efficient and inexpensive in acquisition and robust to amplification and other errors. Since compressed sensing seeks to store signal information in a compressed state and memory is often measured in bits, this motivates the rigorous treatment of quantization in compressed sensing.

We define quantization as follows: given some $x \in \mathbb{R}$, our *quantizer* or *quantization function* is defined as

$$f_Q(x) = \begin{cases} m_1 & \text{if } x \in (-\infty, \tau_2) \\ m_i & \text{if } x \in [\tau_i, \tau_{i+1}) \text{ for } i = 2, \dots, Q, \end{cases}$$

where $\{\tau_1 = -\infty, \tau_2, \dots, \tau_{Q+1} = \infty\}$ is a partition of the real line and $\{m_1, \dots, m_Q\}$ consist of quantization values. When f_Q is applied to a vector \mathbf{x} , we simply quantize each component of \mathbf{x} appropriately. This admits the quantized compressed sensing framework:

$$\mathbf{y} = f_Q(\Phi \mathbf{x}),$$

where we would like to solve for the signal $\mathbf{x} \in \mathbb{R}^n$. An extreme case of quantization is the 1-bit quantizer, where

$$f_Q(\Phi \mathbf{x}) = \text{sign}(\Phi \mathbf{x}).$$

* University of California, Los Angeles

† Claremont McKenna College

Since the magnitude of the signal is lost through low-bit or 1-bit quantization, we may restrict our sparse signals to the hypersphere

$$S^{N-1} = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_2 = 1\}.$$

We thus seek to solve the optimization problem:

$$\min_{\mathbf{x} \in S^{N-1}} \|\mathbf{x}\|_0 \quad \text{s.t.} \quad f_Q(\Phi \mathbf{x}) = \mathbf{y}. \quad (1)$$

Equivalently, we can define the *quantization region*

$$\mathcal{R}_{\mathbf{y}} := \mathcal{R}_{y_1} \times \dots \times \mathcal{R}_{y_M}$$

where \mathcal{R}_{y_i} is the quantization region for the i th component of \mathbf{y} , or the interval $f_Q^{-1}(y_i)$.

This gives the equivalent formulation to the above which we consider:

$$\min_{\mathbf{x} \in S^{N-1}} \|\mathbf{x}\|_0 \quad \text{s.t.} \quad \Phi \mathbf{x} \in \mathcal{R}_{\mathbf{y}}. \quad (2)$$

The 1-bit variant of this compressed sensing problem was initially introduced and studied by P. Boufounos, et. al. [BB08], which then led to the development of many subsequent algorithms for 1-bit reconstruction (see e.g. [JLBB11, PV13, BFN⁺14, YYO12]). One of the first was Binary Iterative Hard Thresholding (BIHT) which demonstrated accurate recovery [JLBB11]. Some variants of BIHT were then introduced to account for noise from acquisition and transmission, including Adaptive Outlier Pursuit (AOP) [YYO12]. In a more general setting, methods for quantized compressed sensing have also been introduced (see e.g. [DPM09, DM11, JDDV13]).

Contribution. We study several greedy approaches to solving this problem, and catalog precisely the tradeoff between reconstruction error, bit depth, and number of measurements. We believe such a catalog is useful for guiding practitioners in selecting which method to use in quantized CS, as in [BT15] for the classical case. In addition, we develop new approaches which outperform existing methods in some contexts.

Organization. We introduce existing methods for 1-bit and quantized CS in Sections 2 and 3. Then, in Section 4 we motivate and propose two new adaptations which we show outperform existing methods in certain regimes in Section 5. Also in this section we perform extensive experiments which catalog the reconstruction behavior for our proposed methods and existing methods. We do so in such a way that given a bit-budget, signal size, and sparsity level, one can optimally select the algorithm for best reconstruction error. Lastly, we conclude in Section 6.

2 Quantized Basis Pursuit

In classical compressed sensing, a relaxation from the ℓ_0 -norm to the ℓ_1 -norm is used to obtain the following basis pursuit formulation:

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \Phi \mathbf{x} = \mathbf{y},$$

which has provable robust reconstruction guarantees [CT05, CRT06].

Similarly, W. Dai, et. al. [DM11] proposed a quantized basis pursuit approach. Given a quantizer f_Q , one defines vectors \mathbf{b}_1 and \mathbf{b}_2 consisting of the thresholds for each coordinate's quantization region. That is, we may solve

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \mathbf{b}_1 \leq \Phi \mathbf{x} \leq \mathbf{b}_2.$$

Since both of these problems are linear programs, one can use traditional interior-point, simplex methods, or even Bregman algorithms to solve them. However, these algorithms are typically slower, motivating the development of greedy approaches.

3 Current Greedy Methods

We first present several current greedy algorithms for quantized compressed sensing found in the literature.

3.1 Quantized Subspace Pursuit (QSP)

The Subspace Pursuit (SP) algorithm was developed by W. Dai, et. al. [DM09] for classic compressed sensing reconstruction. In particular, SP has provable reconstruction capability similar to basis pursuit methods. It was adapted for quantized compressed sensing in [DM11] as we describe here.

Let $T \subset \{1, \dots, N\}$ be an index set and denote Φ_T and \mathbf{x}_T as the truncated matrix consisting of the columns of Φ indexed by T and entries of \mathbf{x} indexed by T , respectively.

We can define the set

$$\mathcal{Q} := \{(\mathbf{x}', \mathbf{y}') \in \mathbb{R}^{|T|} \times \mathcal{R}_y : \|\mathbf{y}' - \Phi_T \mathbf{x}'\|_2 \text{ is minimized}\},$$

the set of minimizers for $\|\mathbf{y}' - \Phi_T \mathbf{x}'\|_2$. We also let

$$(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \underset{(\mathbf{x}', \mathbf{y}') \in \mathcal{Q}}{\operatorname{argmin}} \|\mathbf{y}' - \mathbf{y}\|_2.$$

Then we define the functions:

$$\operatorname{resid}(\mathbf{y}, \Phi_T) := \tilde{\mathbf{y}} - \Phi_T \tilde{\mathbf{x}}, \quad \operatorname{pcoeff}(\mathbf{y}, \Phi_T) := \tilde{\mathbf{x}},$$

which give the residual and projected coefficients, respectively.

These functions may be interpreted as projection operations onto the quantization region of \mathbf{y} , i.e. \mathcal{R}_y . Note that $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ may be uniquely determined by first solving the quadratic optimization problem,

$$\tilde{\mathbf{y}} = \underset{\mathbf{y}' \in \mathcal{R}_y}{\operatorname{argmin}} \|\mathbf{y}' - \mathbf{y}\|_2,$$

which gives a unique solution for $\tilde{\mathbf{y}}$, then solving for $\tilde{\mathbf{x}}$ by

$$\tilde{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\tilde{\mathbf{y}} - \Phi_T \mathbf{x}\|_2,$$

which also admits a unique solution.

Note that this problem is computationally tractable since the constraint $\mathbf{y} \in \mathcal{R}_y$ is a set of linear inequalities. Though no theory has been developed for this algorithm, it has been shown to work well empirically. The algorithm is summarized below. Here and throughout, we write $\operatorname{supp}_K(\mathbf{x})$ to denote the set of indices corresponding to the largest K entries of \mathbf{x} in magnitude, $\operatorname{supp}_K(\mathbf{x}) = \max_{|T| \leq K} \|\mathbf{x}_T - \mathbf{x}\|_2$.

Algorithm 1 Quantized Subspace Pursuit (QSP)

Input: sparsity level K , measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, compressed quantized signal $\mathbf{y} \in \mathbb{R}^M$

Initialize: $T^0 = \operatorname{supp}_K(\Phi^* \mathbf{y})$, $\mathbf{y}_r^0 = \operatorname{resid}(\mathbf{y}, \Phi_{T^0})$

repeat

$$\tilde{T}^l = T^{l-1} \cup \operatorname{supp}_K(\Phi^* \mathbf{y}_r^{l-1}).$$

$$\mathbf{x}_p = \operatorname{pcoeff}(\mathbf{y}, \Phi_{\tilde{T}^l}) \text{ and } T^l = \operatorname{supp}_K(\mathbf{x}_p)$$

$$\mathbf{y}_r^l = \operatorname{resid}(\mathbf{y}, \Phi_{T^l})$$

until $\|\mathbf{y}_r^l\|_2 > \|\mathbf{y}_r^{l-1}\|_2$

$$T^l = T^{l-1}$$

Output: $\hat{\mathbf{x}} / \|\hat{\mathbf{x}}\|_2$ where $\hat{\mathbf{x}}_{\{1, \dots, N\} - T^l} = 0$ and $\mathbf{x}_{T^l} = \Phi_{T^l}^\dagger \mathbf{y}$

3.2 Quantized Iterative Hard Thresholding (QIHT)

The Quantized Iterative Hard Thresholding (QIHT) [JDDV13] algorithm is based on the Iterative Hard Thresholding (IHT) [BD09] and Binary Iterative Hard Thresholding (BIHT) [JLBB11] algorithms. IHT was introduced for iteratively reconstructing a sparse signal in classical compressed sensing. The algorithm may be interpreted as solving:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{x}\|_0 \leq K.$$

IHT solves this problem by iteratively computing

$$\mathbf{x}^{l+1} = \eta_K \left(\mathbf{x}^l + \Phi^* \left(\mathbf{y} - \Phi \mathbf{x}^l \right) \right),$$

where $\eta_K(\mathbf{x})$ thresholds \mathbf{x} by maintaining the K largest entries in magnitude of \mathbf{x} ($\text{supp}_K(\mathbf{x})$) and setting the rest to zero. \mathbf{x} is initialized at $\mathbf{x}^0 = \mathbf{0}$. This algorithm was shown to converge for $\|\Phi\|_2 < 1$ in [BD09].

We may interpret the IHT algorithm as taking a gradient step to minimize the consistency-enforcing objective $\frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|_2^2$ then giving the best K -term approximation by hard thresholding.

Similarly, in the 1-bit setting, BIHT modifies the gradient step in IHT and iteratively computes the following:

$$\mathbf{x}^{l+1} = \eta_K \left(\mathbf{x}^l + \mu \Phi^* \left(\mathbf{y} - \text{sign}(\Phi \mathbf{x}^l) \right) \right).$$

where μ is a scalar that controls the gradient step-size.

This may be interpreted as attempting to minimize the following objective:

$$\min_{\mathbf{x} \in \mathbb{S}^{N-1}} \mu \|\mathbf{y} \odot (\Phi \mathbf{x})\|_- \quad \text{s.t.} \quad \|\mathbf{x}\|_0 \leq K, \quad (3)$$

where $\mu \in \mathbb{R}$ is the chosen gradient step size, \odot is component-wise multiplication, and $[\cdot]_-$ is the projection of each component to the negative real line, i.e. for each component of $\mathbf{x} \in \mathbb{R}^N$,

$$[x_i]_- = \begin{cases} x_i & \text{if } x_i < 0 \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Note that the first term of the objective function enforces the inequality

$$\mathbf{y} \odot (\Phi \mathbf{x}) \geq \mathbf{0},$$

which ensures the consistency of signs of \mathbf{y} and $\Phi \mathbf{x}$.

Motivated by this optimization problem, [JDDV13] formulated a similar problem for multiple quantization values and thresholds by considering the following objective:

$$\min_{\mathbf{x}} \sum_{k=1}^M \sum_{j=2}^{2^B} w_j |\text{sign}((\Phi \mathbf{x})_k - \tau_j)(y_k - \tau_j)| \quad \text{s.t.} \quad \|\mathbf{x}\|_0 \leq K, \quad (5)$$

where $w_j = m_j - m_{j-1}$. This optimization objective may be interpreted as the sum of the BIHT objective over all possible quantization thresholds. Calculating the subgradient of this function gives the update

$$\mathbf{a}^{l+1} = \mathbf{x}^l + \mu \Phi^* \left(\mathbf{y} - f_Q(\Phi \mathbf{x}^l) \right).$$

which we then threshold by

$$\mathbf{x}^{l+1} = \eta_K(\mathbf{a}^{l+1}).$$

QIHT is summarized below in Algorithm 2 [JDDV13].

Algorithm 2 Quantized Iterative Hard Thresholding (QIHT)

Input: measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, compressed quantized signal $\mathbf{y} \in \mathbb{R}^M$, quantization function f_Q , $\mu > 0$ step size, stopping criterion

Initialize: $\mathbf{x}^0 = \frac{\Phi^* \mathbf{y}}{\|\Phi^* \mathbf{y}\|}$

while not converged **do**

$$\mathbf{a}^{l+1} = \mathbf{x}^l + \mu \Phi^* (\mathbf{y} - f_Q(\Phi \mathbf{x}^l))$$

$$\mathbf{x}^{l+1} = \eta_K(\mathbf{a}^{l+1})$$

end while

Output: $\hat{\mathbf{x}} = \mathbf{x}^k / \|\mathbf{x}^k\|_2$

Intuitively, the algorithm takes a subgradient step with step-size μ then projects the signal back to the $K - \ell_0$ sphere. Though the stability and convergence of QIHT have not yet been proven, numerical results and a limited case analysis suggest accurate empirical performance.

Note that when QIHT uses 1-bit measurements, it reduces to BIHT. When QIHT uses extremely fine measurements, the IHT method is recovered.

Like any method, QIHT is robust to noise that does not change the quantized values. To address noise that changes the quantization values of $\Phi \mathbf{x}$, we modify QIHT using Adaptive Outlier Pursuit, which we describe in the next section.

4 New Adaptations

We present two new algorithm variants for reconstruction from quantized measurements.

4.1 Quantized CoSaMP (QCoSaMP)

Motivated by the adaptation of QSP, we consider the adaptation of a similar method, CoSaMP [NT09] to the quantized setting. Quantized Compressive Sampling Matching Pursuit (QCoSaMP) is described below.

Algorithm 3 Quantized Compressive Sampling Matching Pursuit (QCoSaMP)

Input: measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, quantized compressed signal $\mathbf{y} = \Phi \mathbf{x}$, sparsity level $K > 0$, maximum number of iterations I
Initialize: $\mathbf{a}^0 = \mathbf{0}$, $\mathbf{v} = \mathbf{y}$, $k = 0$
while $k < I$ **do**
 Set $\mathbf{u} = \Phi \mathbf{v}$, $\Omega = \text{supp}_{2K}(\mathbf{u})$
 Merge: $T = \Omega \cup \text{supp}(\mathbf{a}^{k-1})$
 Projection: $\mathbf{x} = \arg \min_{\mathbf{u} \in \mathcal{R}_{\mathbf{y}}} \|\mathbf{y} - \mathbf{u}\|_2$
 $\mathbf{b}_T = \arg \min_{\mathbf{x}} \|\mathbf{u} - \Phi_T \mathbf{x}\|_2$ and $\mathbf{b}_{T^c} = \mathbf{0}$
 Set $\mathbf{a}^k = \mathbf{b}$
 Update: $\mathbf{v} = \mathbf{y} - f_Q(\Phi \mathbf{a}^k)$
 $k = k + 1$.
end while
Output: $\hat{\mathbf{x}} = \mathbf{a}^k / \|\mathbf{a}^k\|_2$

The primary difference between QCoSaMP and CoSaMP is the projection step; instead of estimating the signal using least squares, $\Phi_T^\dagger \mathbf{y}$, we project \mathbf{y} into the quantization region $\mathcal{R}_{\mathbf{y}}$, and then apply least squares, as described in the algorithm above. Also, when computing the sample update, we include the quantizer f_Q in computing the residual.

We will see later that empirically QCoSaMP accurately reconstructs a signal from its quantized measurements. Like the other approaches, theoretical guarantees may also be possible; we leave this for future work.

4.2 Quantized Adaptive Outlier Pursuit (AOP-QIHT)

In many applications, unintended noise may be added to the measurements during acquisition and transmission, creating both pre-quantization and post-quantization noise. Since QIHT's performance is best demonstrated when measurements are quantized into the correct corresponding bin, this motivates the use of Adaptive Outlier Pursuit following [Y YO12] to detect potential quantization errors.

Mathematically, we can model our pre-quantization noise as \mathbf{n} where $(\Phi \mathbf{x})_i + n_i$ for some fixed number of i . If n_i is large enough in magnitude for some i , then it may cause $(\Phi \mathbf{x})_i$ to be quantized in a different bin. Post-quantization noise may be modeled as sparse noise in which a measurement is quantized to a completely

random quantization value. In either case, we can model both forms of noise as some sparse vector \mathbf{n} such that its non-zero components satisfy

$$f_Q((\Phi \mathbf{x})_i + n_i) \neq f_Q((\Phi \mathbf{x})_i).$$

Motivated by ideas from 1-bit compressive sensing using Adaptive Outlier Pursuit [YYO12] and QIHT [JDDV13], we define

$$\phi(x, y) = \sum_{j=2}^{2^Q} \left| [(x - \tau_j)(y - \tau_j)]_- \right|.$$

Given x and y , $\phi(x, y)$ penalizes the case where x and y do not lie in the same quantization region.

We may then formulate our problem as

$$\min_{\mathbf{x} \in S^{N-1}, \mathbf{n}} \sum_{k=1}^M \phi((\Phi \mathbf{x} + \mathbf{n})_k, y_k) \quad \text{s.t.} \quad \|\mathbf{n}\|_0 \leq L, \|\mathbf{x}\|_0 \leq K. \quad (6)$$

where we penalize over all components of $\Phi \mathbf{x} + \mathbf{n}$ and \mathbf{y} .

To solve this problem, we introduce a new binary variable, $\Lambda \in \{0, 1\}^M$, such that

$$\Lambda_k = \begin{cases} 1, & \text{if } y_k = f_Q((\Phi \mathbf{x} + \mathbf{n})_k) \\ 0, & \text{otherwise.} \end{cases}$$

and redefine our problem as the following

$$\min_{\mathbf{x} \in S^{N-1}, \Lambda \in \{0, 1\}^M} \sum_{k=1}^M \Lambda_k \phi((\Phi \mathbf{x})_k, y_k) \quad \text{s.t.} \quad \sum_{k=1}^M (1 - \Lambda_k) \leq L, \|\mathbf{x}\|_0 \leq K.$$

We note a few observations: Firstly, our sparsity constraint, $\|\mathbf{n}\|_0 \leq L$, is equivalent to the constraint $\sum_{k=1}^M (1 - \Lambda_k) \leq L$ by definition of Λ .

Secondly, the objective for AOP-QIHT also differs slightly from QIHT's objective. Instead of weighing each inconsistent term by the same weight w_j , we allow the distance from each threshold to influence the weight of each term. Larger differences from each quantization threshold correspond to higher values of $\phi(y_i, (\Phi \mathbf{x})_i)$, allowing for better detection of outliers.

Lastly, if L is unknown, which is true in many real-world applications, we must apply heuristics to choose L . If L is chosen too small or too large, AOP-QIHT will perform worse. If L is smaller, then corrupted measurements will stay in the measurements. If L is larger, uncorrupted measurements will also be classified as outliers leading to the loss of information.

Following M. Yan, et. al.'s solution for Adaptive Outlier Pursuit for BIHT [YYO12], we apply the alternating minimization method:

1. Fix Λ and solve for \mathbf{x} :

$$\min_{\mathbf{x} \in S^{N-1}} \sum_{k=1}^M \Lambda_k \phi((\Phi \mathbf{x})_k, y_k) \quad \text{s.t.} \quad \|\mathbf{x}\|_0 \leq K. \quad (7)$$

We can solve this by performing a QIHT update:

$$\mathbf{x}^{l+1} = \eta_K(\mathbf{x}^l + \mu \Phi^*(\mathbf{y} - f_Q(\Phi \mathbf{x}^l))). \quad (8)$$

2. Fix \mathbf{x} and solve for Λ :

$$\min_{\Lambda \in \{0, 1\}^M} \sum_{k=1}^M \Lambda_k \phi((\Phi \mathbf{x})_k, y_k) \quad \text{s.t.} \quad \sum_{k=1}^M (1 - \Lambda_k) \leq L. \quad (9)$$

We can solve this by:

$$\Lambda_k = \begin{cases} 0 & \text{if } \phi((\Phi \mathbf{x})_k, y_k) \geq M \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

where M is the L th largest component of $\{\phi((\Phi \mathbf{x})_k, y_k)\}_{k=1}^M$.

Intuitively, AOP adaptively removes outliers from the data that it assumes has been corrupted by noise by removing the L largest terms from the objective. These terms are then not used in the subsequent update of the algorithm. Our proposed method is described as follows.

Algorithm 4 Adaptive Outlier Pursuit for Quantized Iterative Hard Thresholding (AOP-QIHT)

Input: measurement matrix $\Phi \in \mathbb{R}^{M \times N}$, quantized compressed signal \mathbf{y} , sparsity level $K > 0$, number of wrongly detected measurements $L > 0$, $\mu > 0$ step size, maximum number of iterations $I > 0$

Initialize: $\mathbf{x}^0 = \frac{\Phi^* \mathbf{y}}{\|\Phi^* \mathbf{y}\|}$, $k = 0$, $\Lambda = \mathbf{1} \in \mathbf{R}^M$, $T = \{1, \dots, M\}$, $\text{tol} = \infty$, $\text{TOL} = \infty$

while $l \leq I$ and $L \leq \text{tol}$ **do**

Compute $\mathbf{a}^{l+1} = \mathbf{x}^k + \mu \Phi_T^* (\mathbf{y} - f_Q(\Phi_T \mathbf{x}^k))$

Update $\mathbf{x}^{l+1} = \eta_K(\mathbf{a}^{l+1})$

Set $\text{tol} = \|\mathbf{y} - f_Q(\Phi \mathbf{x}^{l+1})\|_0$

if $\text{tol} \leq \text{TOL}$ **then**

Compute Λ as above

Update: $T = \text{supp}(\Lambda)$

Set $\text{TOL} = \text{tol}$.

end if

$l = l + 1$.

end while

Output: $\hat{\mathbf{x}} = \mathbf{x}^k / \|\mathbf{x}^k\|_2$

5 Numerical Experiments

In this section, we perform several numerical experiments to demonstrate the effectiveness of the approaches. Furthermore, we perform an in-depth comparison and analysis of all the methods to choose a preferred bit-depth and algorithm for given noise, bit budget, and sparsity.

To setup these experiments, we generate a measurement matrix $\Phi \in \mathbb{R}^{M \times N}$ whose elements follow an i.i.d. Gaussian distribution, a K -sparse signal \mathbf{x}^* whose non-zero entries are drawn from a standard Gaussian distribution which is then normalized to have unit norm. We compute our compressed signal $\mathbf{y} = f_Q(\Phi \mathbf{x}^*)$.

5.1 QCoSaMP Experiments

In the first experiment, we set the signal size to be $N = 1000$ with sparsity $K = 10$, set $M = 1000$ and vary the total bit budget T_B and the bit-depth B (bits per measurement, $B \approx \log(Q)$). We compute the reconstruction SNR¹ (RSNR) over 40 trials and compare the average for each method. The results are shown in Figure 1. This experiment demonstrates that QIHT outperforms the other greedy approaches for extremely low bit-depth. However, at higher bit-depths, QCoSaMP yields better recovery in the low bit budget regime. A theoretical understanding of this phenomena is interesting future work.

¹We use $\text{SNR} = 10 \log \left(\frac{\|\mathbf{x}^*\|_2^2}{\|\mathbf{x} - \mathbf{x}^*\|_2^2} \right)$.

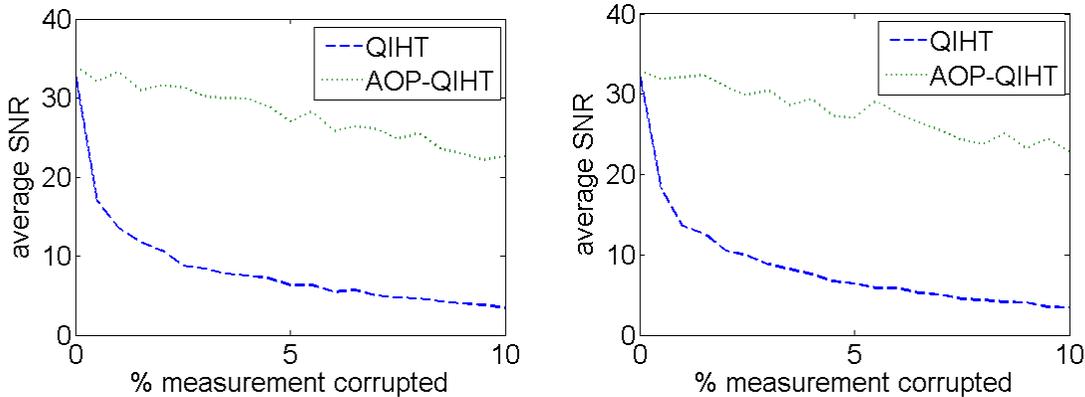


Figure 2: Comparison of QIHT and AOP-QIHT on corrupted data with different noise levels. In these experiments, $N = M = 1000$ and $K = 10$. The top uses 1-bit measurements and the bottom uses 4-bit measurements. We plot the percentage of corrupted measurements against the average SNR over 40 trials.

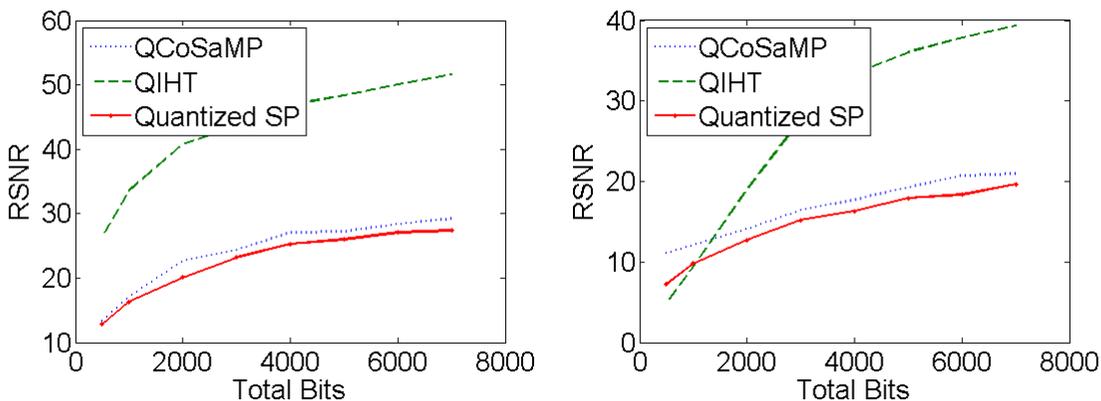


Figure 1: Comparison of QCoSaMP, Quantized SP, and QIHT. We graph total bits against RSNR with $N = M = 1000$ and $K = 10$, averaged over 40 trials. The left uses 1-bit measurements and the right uses 4-bit measurements.

5.2 AOP-QIHT Experiments

In this experiment we keep the parameters as above, but in each trial we select a few measurements and flip their sign to corrupt the measurement’s quantization. We vary the corruption level to be between 0% and 10% of the total number of measurements and record the average SNR over 40 trials. Figure 2 demonstrates that AOP-QIHT outperforms QIHT for all noise levels, especially when noise affects a higher percentage of measurements. We observe that AOP-QIHT performs significantly better than QIHT especially for $> 2\%$ measurement corruption. For future work, we can also examine how this algorithm performs for different quantization schemes and different types of noise.

5.3 Quantized CS Algorithm Comparison and Analysis

In this set of experiments, we compare all the greedy algorithms and examine which algorithms and bit-depths perform best for a given sparsity, total bits, and noise level. We consider sparsity $K = \{2, 4, 6, 8, 10, 12, 14, 16\}$, total bits $T_B = \{500, 1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000, 10000\}$, input signal-to-noise ratio $ISNR = \{35, 20, 10\}$, and bit depth $B = \{1, 2, 3, 4\}$. We average over 20 trials. We show results for selected parameter values, with omitted results that vary as one expects in between those we show here. In addition, we summarize all results compactly in Figures 5 – 7. These figures catalog which algorithm exhibits the best performance for various

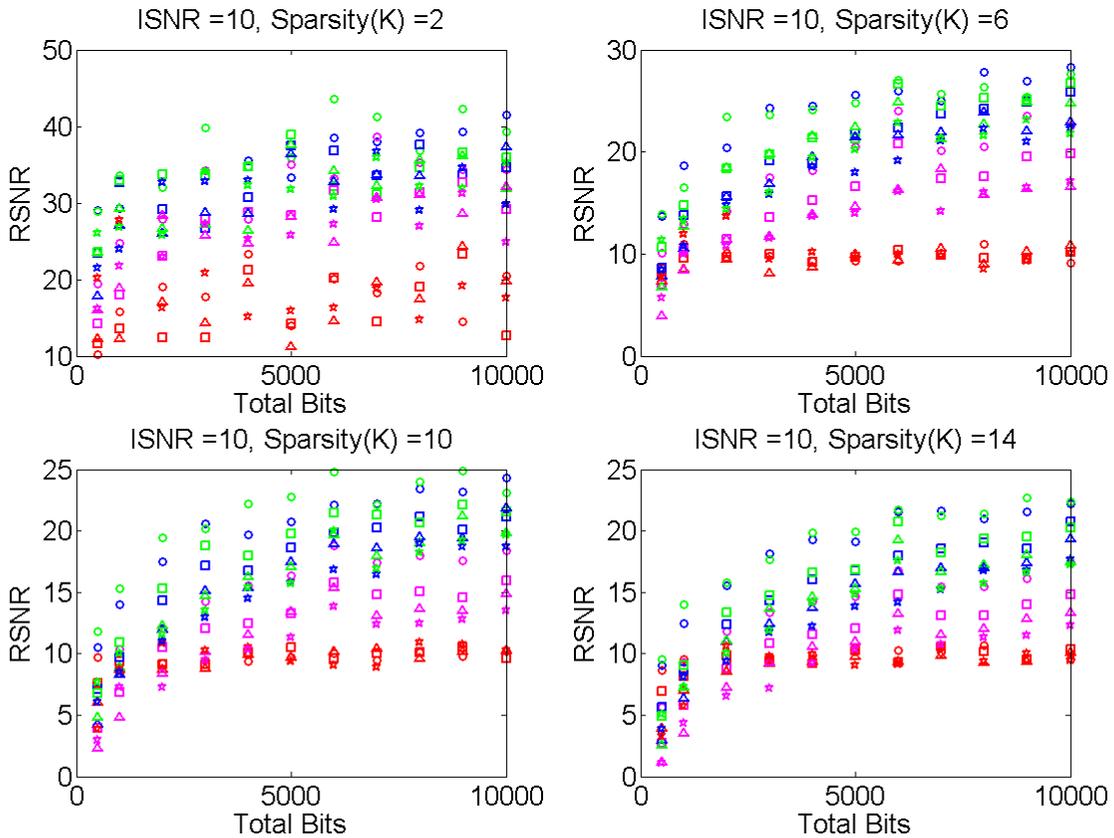


Figure 3: Comparison of QIHT (red), AOP-QIHT (magenta), QCoSaMP (green), and Quantized SP (blue) with bit-depths 1 (circle), 2 (square), 3 (triangle), and 4 (star). These figures graph total bits against RSNR for ISNR = 10 and various sparsity levels.

parameter choices. We hope that this serves as a useful tool for practitioners when deciding which method to use for their application.

We observe that for lower noise levels, QIHT and AOP-QIHT tend to perform better while QCoSaMP and Quantized SP perform better for higher noise. This may be explained by the hard thresholding that QIHT and AOP-QIHT performs. AOP-QIHT also performs better than QIHT for higher noise. Most graphs also indicate that 1-bit is best.

6 Conclusion

Compressed sensing has furnished a rigorous theory and class of methods for signal reconstruction from compressed measurements. However, for practical use one must also consider how quantization affects the reconstruction error. Moreover, extreme quantization – where each measurement is captured by only a few bits – can be useful in its own right due to the efficiency of implemented hardware.

In this paper, we propose two novel robust greedy algorithms for reconstructing quantized compressed signals. QCoSaMP modifies CoSaMP by projecting onto the quantization region by solving a set of computationally tractable optimization problems, then considering quantization in the residual. AOP-QIHT iteratively detects falsely quantized measurements and recovers signals from the “correct” measurements by applying QIHT.

Allowing the bit-depth and bit-budget to be additional parameters in the CS framework, one now needs to select a method from the existing approaches that is most in line with the desired parameter regime. For that reason, in this paper we compare the performance of four greedy algorithms, Quantized SP, QCoSaMP, QIHT, and AOP-QIHT over multiple settings of parameters for normalized signals. We show that the one-bit QIHT and AOP-QIHT algorithms tend to perform best for low-noise cases while QCoSaMP and QSP perform better

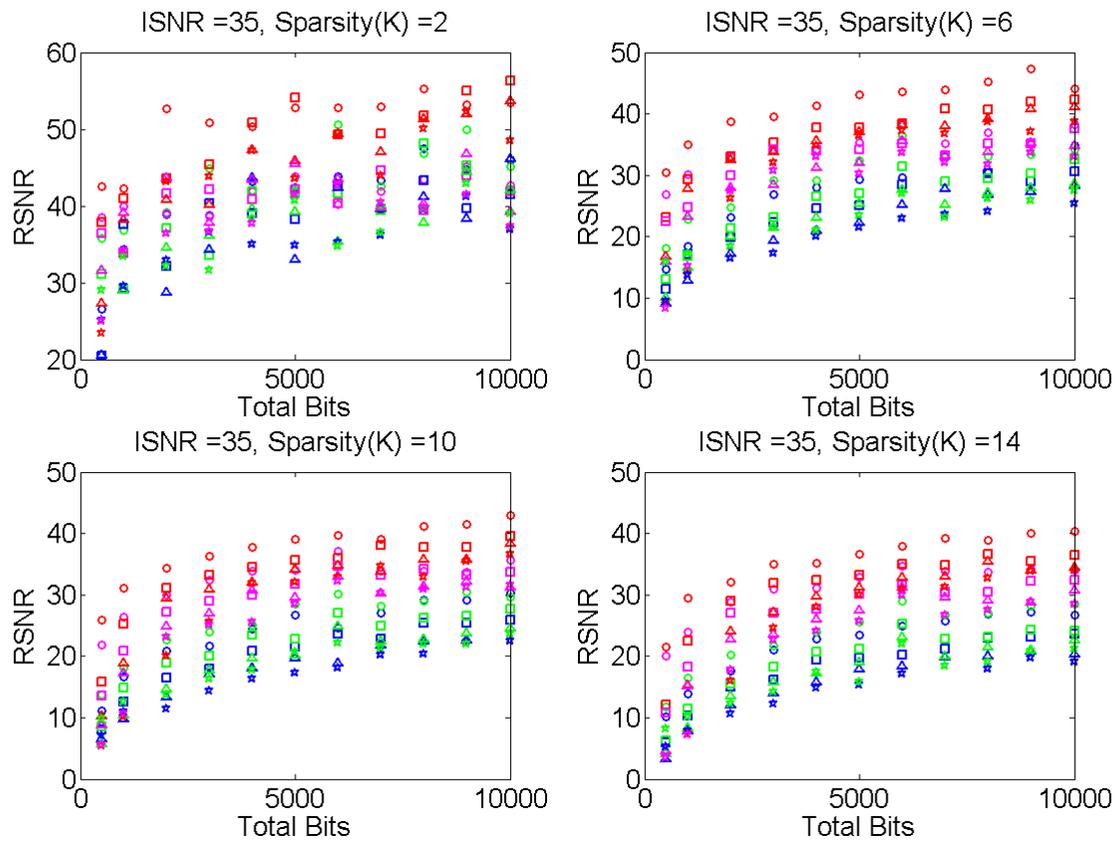


Figure 4: Comparison of QIHT (red), AOP-QIHT (magenta), QCoSaMP (green), and Quantized SP (blue) with bit-depths 1 (circle), 2 (square), 3 (triangle), and 4 (star). These figures graph total bits against RSNR for ISNR = 35 and various sparsity levels.

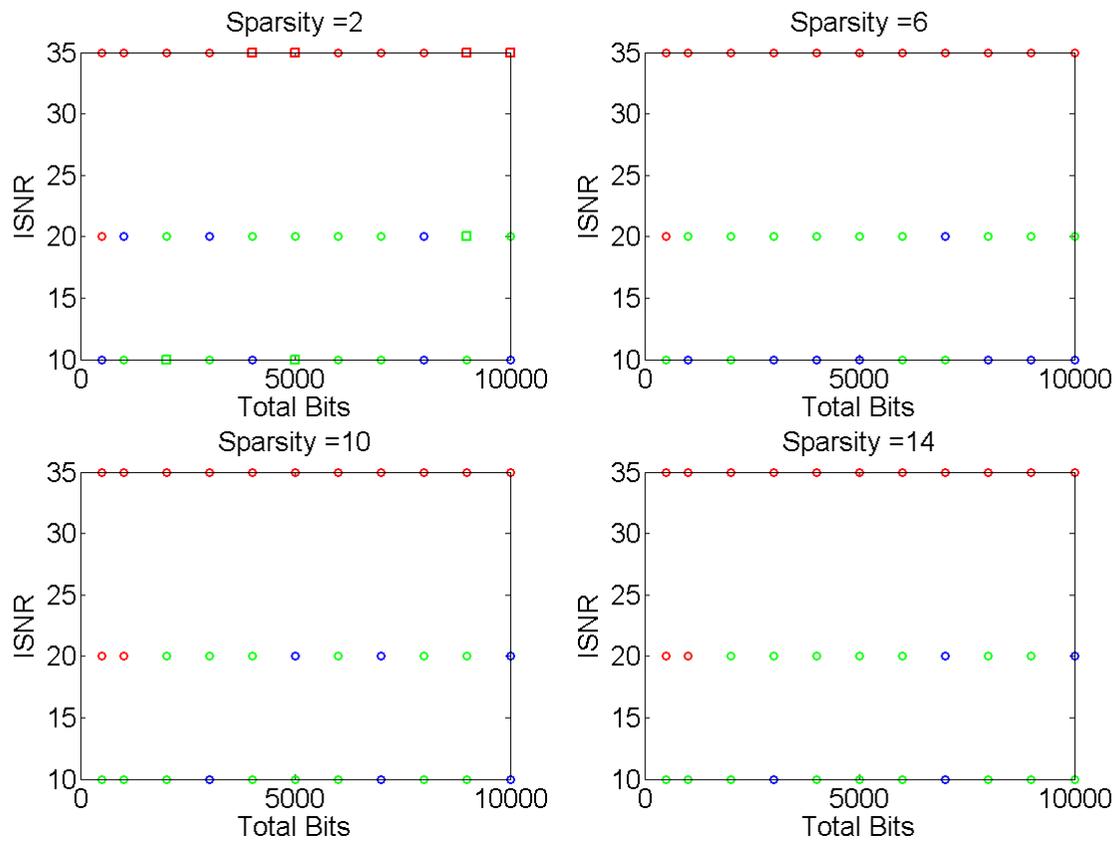


Figure 5: Comparison of QIHT (red), AOP-QIHT (magenta), QCoSaMP (green), and Quantized SP (blue) with bit-depths 1 (circle), 2 (square), 3 (triangle), and 4 (star). These figures graph the best algorithm and bit-depth for given total bits and ISNR for various sparsities.

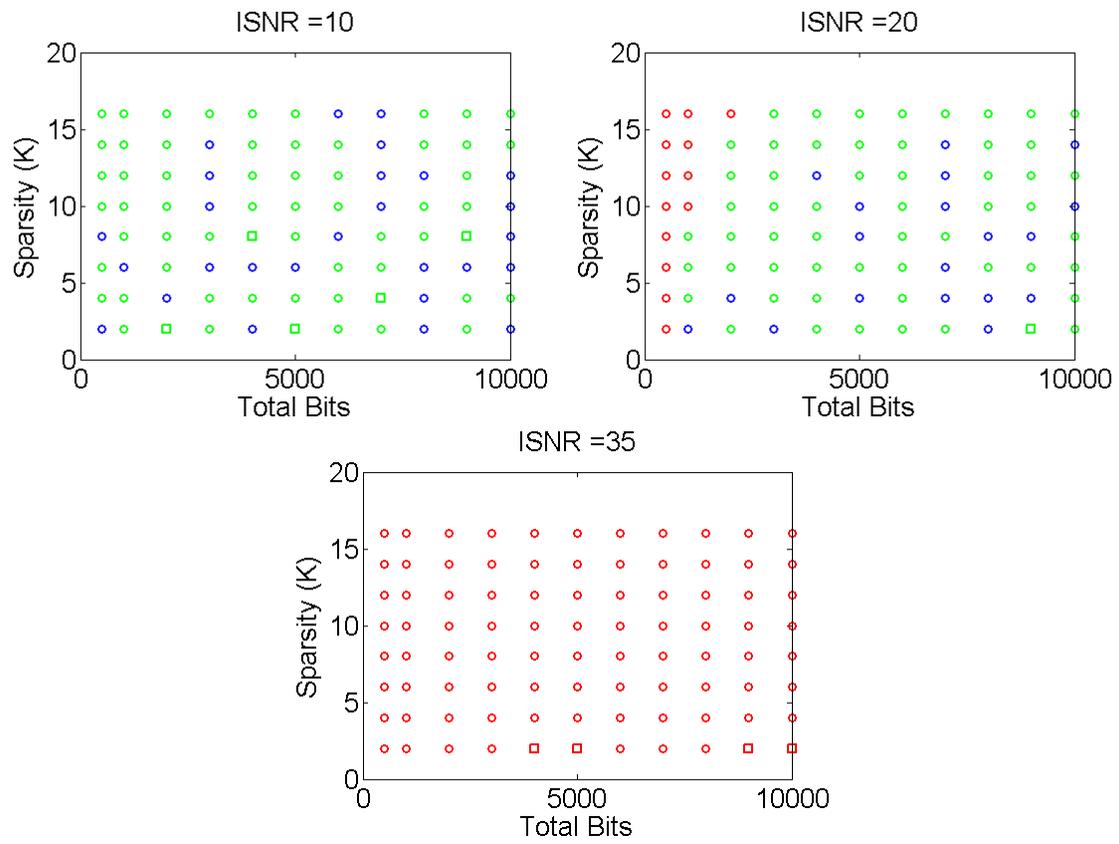


Figure 6: Comparison of QIHT (red), AOP-QIHT (magenta), QCoSaMP (green), and Quantized SP (blue) with bit-depths 1 (circle), 2 (square), 3 (triangle), and 4 (star). These figures graph the best algorithm and bit-depth for given total bits and sparsity for various ISNR.

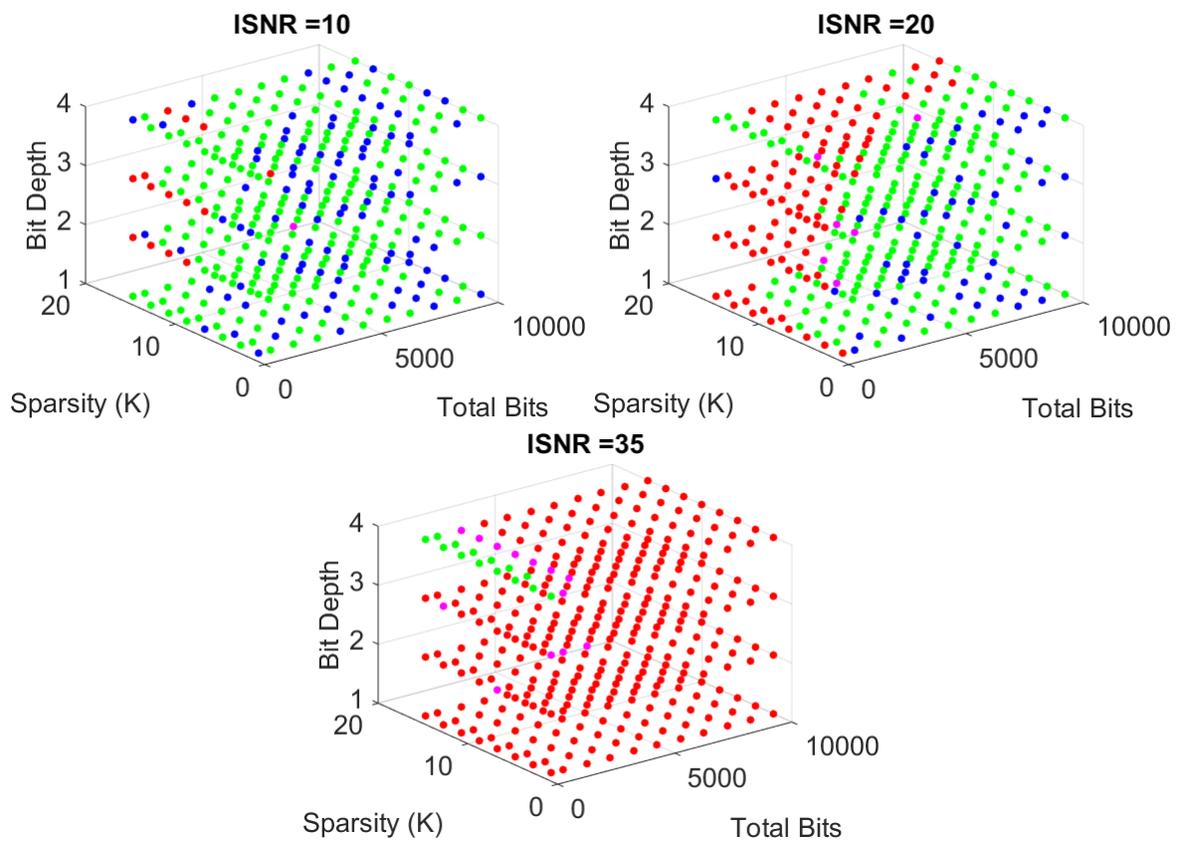


Figure 7: Comparison of QIHT (red), AOP-QIHT (magenta), QCoSaMP (green), and Quantized SP (blue). These figures graph the best algorithm for given total bits, sparsity, and bit-depth for various ISNR.

for higher noise. We believe this is a useful tool that will help guide practitioners when navigating the existing approaches.

Acknowledgements

The work of Hao-Jun Michael Shi and Mindy Case was supported in part by the California Research Training Program for Computational and Applied Mathematics 2015 under NSF Grant DMS #1045536. The work of Xiaoyi Gu and Shenyinying Tu was supported in part by the California Research Training Program for Computational and Applied Mathematics 2015 under NSF CAREER #1348721. The work of Deanna Needell was supported under NSF CAREER #1348721. The authors would also like to thank Prof. Andrea Bertozzi for hosting the summer program in which this work originated.

References

- [BB08] P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *42nd Annual Conference on Information Sciences and Systems (CISS)*, pages 16–21. IEEE, 2008.
- [BD09] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- [BFN⁺14] R. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wootters. Exponential decay of reconstruction error from binary measurements of sparse signals. *arXiv preprint arXiv:1407.8246*, 2014.
- [BT15] J. D. Blanchard and J. Tanner. Performance comparisons of greedy algorithms in compressed sensing. *Numerical Linear Algebra*, 22(2):254–282, 2015.
- [BV82] W. G. Bath and V. D. Vandelinde. Robust memoryless quantization for minimum signal distortion. *IEEE Transactions on Information Theory*, 28(2), 1982.
- [CRT06] E. J. Candès, 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.
- [CT05] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51:4203–4215, 2005.
- [DM09] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE Transactions on Information Theory*, 55(5):2230–2249, 2009.
- [DM11] W. Dai and O. Milenkovic. Information theoretical and algorithmic approaches to quantized compressive sensing. *IEEE transactions on communications*, 59(7):1857–1866, 2011.
- [DPM09] W. Dai, H. V. Pham, and O. Milenkovic. Distortion-rate functions for quantized compressive sensing. Citeseer, 2009.
- [GB08] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008.
- [GB14] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1, March 2014.
- [GN98a] R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(6):2325–2383, 1998.

- [GN98b] R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(3), 1998.
- [JDDV13] L. Jacques, K. Degraux, and C. De Vleeschouwer. Quantized iterative hard thresholding: Bridging 1-bit and high-resolution quantized compressed sensing. *arXiv preprint arXiv:1305.1786*, 2013.
- [JLBB11] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, 2011.
- [Kab09] P. Kabal. Quantizers. 2009.
- [LB11] J. N. Laska and R. G. Baraniuk. Regime change: Bit-depth versus measurement-rate in compressive sensing. 2011.
- [Llo82] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, IT-28:129–137, 1982.
- [MV74] J. M. Morris and V. D. Vandelinde. Robust quantization of discrete-time signals with independent samples. *IEEE Transactions on Communication*, 22(12), 1974.
- [Nee09] D. Needell. Topics in compressed sensing, 2009.
- [NT09] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [PD51] P. F. Panter and W. Dite. Quantization distortion in pulse-count modulation with nonuniform spacing of levels. *Proceedings of the IRE*, 39(1):40–48, 1951.
- [PV13] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66(8):1275–1297, 2013.
- [PV15] Y. Plan and R. Vershynin. The generalized lasso with non-linear observations. *arXiv preprint arXiv:1502.04071*, 2015.
- [RV08] M. Rudelson and R. Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Communications on Pure and Applied Mathematics*, 61:1025–1045, 2008.
- [San15] B. Santhanam. Non-uniform quantization. 2015.
- [vdBF07] E. van den Berg and M. P. Friedlander. SPGL1: A solver for large-scale sparse reconstruction, June 2007. <http://www.cs.ubc.ca/labs/scl/spgl1>.
- [vdBF08] E. van den Berg and M. P. Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.
- [YYO12] M. Yan, Y. Yang, and S. Osher. Robust 1-bit compressive sensing using adaptive outlier pursuit. *IEEE Transactions on Signal Processing*, 60(7):3868–3875, 2012.