

Quantized Iterative Hard Thresholding: Bridging 1-bit and High-Resolution Quantized Compressed Sensing

Laurent Jacques*, Kévin Degraux and Christophe De Vleeschouwer*

ICTEAM Institute, ELEN Department, Université catholique de Louvain (UCL)

May 9, 2013

Abstract

In this work, we show that reconstructing a sparse signal from quantized compressive measurement can be achieved in an unified formalism whatever the (scalar) quantization resolution, *i.e.*, from 1-bit to high resolution assumption. This is achieved by generalizing the *iterative hard thresholding* (IHT) algorithm and its binary variant (BIHT) introduced in previous works to enforce the consistency of the reconstructed signal with respect to the quantization model. The performance of this algorithm, simply called *quantized* IHT (QIHT), is evaluated in comparison with other approaches (*e.g.*, IHT, *basis pursuit denoise*) for several quantization scenarios.

1 Introduction

Since the advent of Compressed Sensing (CS) almost 10 years ago [1, 2], many works have treated the problem of inserting this theory into an appropriate quantization scheme. This step is indeed mandatory for transmitting, storing and even processing any compressively acquired information, and more generally for sustaining the embedding of the CS principle in sensor design.

In its most popular version, CS provides uniform theoretical guarantees for stably recovering any sparse (or compressible) signal at a sensing rate proportional to the signal intrinsic dimension (*i.e.*, its *sparsity* level) [1, 2]. In this context, scalar quantization of compressive measurements has been considered along two main directions.

First, under a high-resolution quantization assumption, *i.e.*, when the number of bits allocated to encode each measurement is high, the quantization impact is often modeled as a mere additive Gaussian noise whose variance is adjusted to the quantization ℓ_2 -distortion [3]. In short, under this high-rate model, the CS stability guarantees under additive Gaussian noise, *i.e.*, as derived from the $\ell_2 - \ell_1$ instance optimality [2], are used to bound the reconstruction error obtained from quantized observations. Variants of these works handle quantization saturation [4], prequantization noise [5], ℓ_p -distortion models ($p \geq 2$) for improved reconstruction in oversampled regimes [6, 7], optimize the high-resolution quantization procedure [8] or integrate more evolved $\Sigma\Delta$ -quantization models departing from scalar PCM quantization [9].

Second, and more recently, extreme 1-bit quantization recording only the sign of the compressive measurement, *i.e.*, an information encoded in a single bit, has been considered [10–13]. New guarantees have been developed to tackle the non-linear nature of the sign operation thanks to the replacement of the *restricted isometric property* (RIP) by the quasi-isometric *binary ϵ -stable embedding* (BeSE) [11], or to more general characterization of the binary embedding of sets based on their Gaussian Mean Width [12, 13]. In this context, iterative methods such as the *binary iterative hard thresholding* [11] or linear programming optimization [12] have been introduced for estimating the 1-bit sensed signal.

This work proposes a general procedure for handling the reconstruction of sparse signals observed according to a standard non-uniform scalar quantization of the compressive measurements. The novelty of this scheme is its ability to handle any resolution level, from 1-bit to high-resolution, in a progressive

*LJ and CDV are funded by the Belgian F.R.S-FNRS. Part of this research is supported by the DETROIT project (WIST3), Walloon Region, Belgium. *Acknowledgements:* We thank Prasad Sudhakar (UCL/ICTEAM) and the anonymous reviewers of SAMPTA 2013 for their useful comments. *Note:* This document is a preprint related to another work accepted in Sampta13, Bremen, Germany.

fashion. Conversely to the Bayesian approach of [15], our method relies on a generalization of the *iterative hard thresholding* (IHT) [16] that we simply called *quantized iterative hard thresholding*. Actually, QIHT reduces to BIHT for 1-bit sensing and it converges to IHT at high resolution.

Conventions Most of domain dimensions (*e.g.*, M, N) are denoted by capital roman letters. Vectors and matrices are associated to bold symbols while lowercase light letters are associated to scalar values. The i^{th} component of a vector $\mathbf{u}$ is u_i or $(\mathbf{u})_i$. The identity matrix is $\mathbf{Id}$. The set of indices in $\mathbb{R}^D$ is $[D] = \{1, \dots, D\}$. Scalar product between two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^D$ reads $\mathbf{u}^* \mathbf{v} = \langle \mathbf{u}, \mathbf{v} \rangle$ (using the transposition $(\cdot)^*$), while the Hadamard product $\mathbf{u} \odot \mathbf{v}$ is such that $(\mathbf{u} \odot \mathbf{v})_i = u_i v_i$. For any $p \geq 1$, $\|\cdot\|_p$ represents the ℓ_p -norm such that $\|\mathbf{u}\|_p^p = \sum_i |u_i|^p$ with $\|\mathbf{u}\| = \|\mathbf{u}\|_2$ and $\|\mathbf{u}\|_\infty = \max_i |u_i|$. The ℓ_0 “norm” is $\|\mathbf{u}\|_0 = \#\text{supp } \mathbf{u}$, where $\#$ is the cardinality operator and $\text{supp } \mathbf{u} = \{i : u_i \neq 0\} \subseteq [D]$. For $S \subseteq [D]$, $\mathbf{u}_S \in \mathbb{R}^{\#S}$ (or Φ_S) denotes the vector (resp. the matrix) obtained by retaining the components (resp. columns) of $\mathbf{u} \in \mathbb{R}^D$ (resp. $\Phi \in \mathbb{R}^{D' \times D}$) belonging to $S \subseteq [D]$. The operator $\mathcal{H}_K$ is the hard thresholding operator setting all the coefficients of a vector to 0 but those having the K strongest amplitudes. The set of canonical K -sparse vectors in $\mathbb{R}^N$ is $\Sigma_K = \{\mathbf{v} \in \mathbb{R}^N : \|\mathbf{v}\|_0 \leq K\}$ while $\Sigma_{\mathcal{T}}$ denotes the set of vectors whose support is $\mathcal{T} \subseteq [N]$. Moreover, $\Sigma_K^* = \Sigma_K \cap S^{N-1}$ and $\Sigma_{\mathcal{T}}^* = \Sigma_{\mathcal{T}} \cap S^{N-1}$ with S^{N-1} the $(N-1)$ -sphere in $\mathbb{R}^N$. Finally, $\chi_{\mathcal{I}}$ is the characteristic function on $\mathcal{I} \subset \mathbb{R}$, $\text{sign } \lambda$ equals 1 if λ is positive and -1 otherwise, $(\lambda)_+ = (\lambda + |\lambda|)/2$ and $(\lambda)_- = -(-\lambda)_+$ project λ on $\mathbb{R}_+$ and $\mathbb{R}_-$, respectively, with all these operators being applied component wise onto vectors.

2 Noisy Compressed Sensing Framework

The *iterative hard thresholding* (IHT) algorithm has been introduced for iteratively reconstructing a sparse or compressible signal $\mathbf{x} \in \mathbb{R}^N$ from compressible observations $\mathbf{y} = \Phi \mathbf{x} + \mathbf{n}$, where $\Phi \in \mathbb{R}^{M \times N}$ is the sensing matrix and $\mathbf{n} \in \mathbb{R}^M$ stands for a possible observational noise with bounded energy $\|\mathbf{n}\| \leq \varepsilon$. IHT is an alternative to the *basis pursuit denoise* (BPDN) method [17] which aims at solving a global convex minimization promoting a ℓ_1 -sparse data prior model under the constraint of reproducing the compressive observation.

Assuming that $\mathbf{x}$ is K -sparse in the canonical basis $\Psi = \mathbf{Id}$, *i.e.*, $\mathbf{x} \in \Sigma_K$, the IHT algorithm is designed to approximately solve the (LASSO-type) problem

$$\min_{\mathbf{u} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{u}\|^2 \text{ s.t. } \|\mathbf{u}\|_0 \leq K. \quad (1)$$

It proceeds by computing the following recursion

$$\mathbf{x}^{(n+1)} = \mathcal{H}_K[\mathbf{x}^{(n)} + \mu \Phi^*(\mathbf{y} - \Phi \mathbf{x}^{(n)})], \quad (\text{IHT})$$

where $\mathbf{x}^{(0)} = \mathbf{0}$, and $\mu > 0$ must satisfy $\mu^{-2} > \|\Phi\| := \sup_{\mathbf{u}: \|\mathbf{u}\|=1} \|\Phi \mathbf{u}\|$ for guaranteeing convergence [18].

In other words, at each iteration, starting from the previous estimation $\mathbf{x}^{(n)}$, the fidelity function $\mathcal{E}(\mathbf{u}) := \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{u}\|^2$ is decreased by a gradient descent step with gradient $\nabla \mathcal{E}(\mathbf{x}^{(n)}) = \Phi^*(\Phi \mathbf{x}^{(n)} - \mathbf{y})$, followed by a “projection” on Σ_K accomplished by the hard thresholding $\mathcal{H}_K$.

In [16], it is shown that if Φ respects the *restricted isometry property* (RIP) of order $3K$ with radius $\delta_{3K} < 1/15$, which means that for all $\mathbf{u} \in \Sigma_{3K}$, $(1 - \delta_{3K})\|\mathbf{u}\|^2 \leq \|\Phi \mathbf{u}\|^2 \leq (1 + \delta_{3K})\|\mathbf{u}\|^2$, then, at iteration $n^* = \lceil \log_2 \|\mathbf{x}\|/\varepsilon \rceil$, the reconstruction error satisfies $\|\mathbf{x} - \mathbf{x}^{(n^*)}\| \leq 5\varepsilon$.

3 Quantized Sensing Model

For the sake of simplicity, let us consider a unit K -sparse signal $\mathbf{x}_0 \in \Sigma_K^*$ observed through the following Quantized Compressed Sensing (QCS) model

$$\mathbf{y} = \mathcal{Q}_b[\Phi \mathbf{x}_0], \quad (2)$$

where $\Phi \in \mathbb{R}^{M \times N}$ is the sensing matrix and $\mathcal{Q}_b$ the quantization operator defined at a *resolution* of b -bits per measurement, *i.e.*, with no further encoding treatment, $\mathbf{y}$ requires a total of $\mathfrak{B} = bM$ bits. In this work, we will not consider any prequantization noise in (2).

The quantization $\mathcal{Q}_b$ is assumed optimal with respect to the distribution of each component of $\mathbf{z} = \Phi \mathbf{x}_0 \in \mathbb{R}^M$. In particular, by considering only random Gaussian matrices $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$, *i.e.*,




Figure 1: (plain curve) Plot of J as a function of $\nu \in \mathbb{R}$ for $b = 3$ ($\tau_5 = 0$) and $\lambda \in \mathcal{R}_5$. (dashed curve) Plot of $\frac{1}{2}(\nu - q_5)^2$.

where each matrix entry follows $\Phi_{ij} \sim_{\text{iid}} \mathcal{N}(0, 1)$, we have $z_i \sim \mathcal{N}(0, \|\mathbf{x}_0\|^2 = 1)$ and we adjust $\mathcal{Q}_b$ to an optimal b -bits Gaussian Quantizer minimizing the quantization distortion, *e.g.*, using a Lloyd-Max optimization [19]. This provides a set of thresholds $\{\tau_i \in \mathbb{R} : 1 \leq i \leq 2^b + 1\}$ (with $-\tau_1 = \tau_{2^b+1} = +\infty$) defining 2^b quantization bins $\mathcal{R}_i = [\tau_i, \tau_{i+1})$, and a set of levels $\{q_i \in \mathcal{R}_i : 1 \leq i \leq 2^b\}$ such that

$$\mathcal{Q}_b[\lambda] = q_k \Leftrightarrow \lambda \in \mathcal{R}_k,$$

with $2\tau_i = q_{i-1} + q_i$ and $q_i = \mathbb{E}[g_x | g_x \in \mathcal{R}_i]$ with $g_x \sim \mathcal{N}(0, 1)$. Notice that this QCS model includes 1-bit CS scheme since $\mathcal{Q}_1[\lambda] = q_0 \text{sign}(\lambda)$ with $q_0 := q_2 = -q_1 = \sqrt{2/\pi}$.

4 Quantized Iterative Hard Thresholding

In this section, we propose a generalization of the IHT algorithm taking into account the particular nature of the scalar quantization model introduced in Sec. 3. The idea is to enforce the consistency of the iterates with the quantized observations. This is first achieved by defining an appropriate cost measuring deviation from quantization consistency.

Given $\nu, \lambda \in \mathbb{R}$ and using the levels and thresholds associated to $\mathcal{Q}_b$, we first define

$$J(\nu, \lambda) = \sum_{j=2}^{2^b} w_j |(\text{sign}(\lambda - \tau_j)(\nu - \tau_j))_-|, \quad (3)$$

with $w_j = q_j - q_{j-1}$. Equivalently, given $\mathcal{I}(\nu, \lambda) := [\min(\nu, \lambda), \max(\nu, \lambda)]$, $J(\nu, \lambda) = \sum_{j=2}^{2^b} w_j \chi_{\mathcal{I}}(\tau_j) |\nu - \tau_j|$. The non-zero terms are therefore determined by the thresholds lying between λ and ν , *i.e.*, for which $\text{sign}(\lambda - \tau_j) \neq \text{sign}(\nu - \tau_j)$. Interestingly, $J(\nu; \lambda) = J(\nu; \mathcal{Q}_b(\lambda))$ since $\text{sign}(\lambda - \tau_j) = \text{sign}(\mathcal{Q}_b(\lambda) - \tau_j)$ for all $j \in [2^b + 1]$.

Then, our quantization consistency function between two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^M$ reads

$$\mathcal{J}(\mathbf{u}, \mathbf{v}) := \sum_{k=1}^M J(u_k, v_k) = \mathcal{J}(\mathbf{u}, \mathcal{Q}_b(\mathbf{v})). \quad (4)$$

This cost, which is convex with respect to $\mathbf{u}$, has two interesting limit cases. First, for $b = 1$, it reduces to the cost on which relies the *binary iterative hard thresholding* algorithm (BIHT) adapted to 1-bit CS [11]. In this context, the sum in (3) has only one term (for $j = 2$) and $\mathcal{J}(\mathbf{u}, \mathbf{v}) = 2q_0 \|(\text{sign}(\mathbf{v}) \odot \mathbf{u})_-\|_1$. Up to a normalization by $2q_0$, this is the ℓ_1 -sided norm minimized by BIHT which vanishes when $q_0 \text{sign}(\mathbf{u}) = \mathcal{Q}_1(\mathbf{u}) = \mathcal{Q}_1(\mathbf{v}) = q_0 \text{sign}(\mathbf{v})$, with q_0 defined in Sec. 3.

Second, in the high resolution limit when $b \gg 1$, $\mathcal{J}(\mathbf{u}, \mathbf{v})$ tends to $\frac{1}{2}\|\mathbf{u} - \mathbf{v}\|^2$. Indeed, in this case $w_j \ll 1$ and, the sum in (3) tends to

$$J(\nu, \lambda) \simeq \left| \int_{\nu}^{\lambda} (\nu - t) dt \right| = \frac{1}{2}(\nu - \lambda)^2.$$

This asymptotic quadratic behavior of J is illustrated in Fig. 1.

Given the quantization consistency cost $\mathcal{J}$, we can now formulate a generalization of (1) for estimating a K -sparse signal $\mathbf{x}_0$ observed by the model (2):

$$\min_{\mathbf{u} \in \mathbb{R}^N} \mathcal{E}_b(\mathbf{u}) \text{ s.t. } \|\mathbf{u}\|_0 \leq K, \quad (5)$$

with $\mathcal{E}_b(\mathbf{u}) := \mathcal{J}(\Phi \mathbf{u}, \mathbf{y}) = \mathcal{J}(\Phi \mathbf{u}, \mathcal{Q}_b[\Phi \mathbf{x}_0])$.

Following the procedure determining the IHT algorithm from (1) (Sec. 2), our aim is to find an IHT variant which minimizes the quantization inconsistency, as measured by $\mathcal{E}_b$, instead of the quadratic cost $\mathcal{E}$. This is done by first determining a *subgradient* of the convex but non-smooth function $\mathcal{E}_b$ [20].

A quick calculation shows that a subdifferential of $J(\nu, \lambda)$ with respect to ν reads

$$\sum_{j=k_-+1}^{k_+} \frac{w_j}{2} (\text{sign}(\nu - \tau_j) - \text{sign}(\lambda - \tau_j)), \quad (6)$$

where $k_- = \min(k_\nu, k_\lambda)$, $k_+ = \max(k_\nu, k_\lambda)$, and k_ν and k_λ are the bin indices of $\mathcal{Q}_b(\nu)$ and $\mathcal{Q}_b(\lambda)$ respectively. From the definition of the w_j , the sum simplifies to $q_{k_\nu} - q_{k_\lambda}$. Therefore, a subgradient of $\mathcal{J}(\mathbf{u}, \mathbf{v})$ with respect to $\mathbf{u}$ reads simply $\mathcal{Q}_b(\mathbf{u}) - \mathcal{Q}_b(\mathbf{v})$, so that a subgradient of $\mathcal{J}(\Phi\mathbf{u}, \mathbf{y})$ with respect to $\mathbf{u}$ corresponds to $\Phi^*(\mathcal{Q}_b(\Phi\mathbf{u}) - \mathbf{y})$.

Therefore, from this last ingredient, we define the *quantized iterative hard thresholding algorithm* (QIHT) by the recursion

$$\mathbf{x}^{(n+1)} = \mathcal{H}_K[\mathbf{x}^{(n)} + \mu\Phi^*(\mathbf{y} - \mathcal{Q}_b(\Phi\mathbf{x}^{(n)}))], \quad (\text{QIHT})$$

where $\mathbf{x}^{(0)} = \mathbf{0}$ and μ is set hereafter.

5 QIHT analysis

Despite successful simulations of sparse signal recovery from quantized measurements (see Sec. 6), we were not able to prove the stability and the convergence of the QIHT algorithm yet. However, there exist a certain number of promising properties suggesting the existence of such a result. The first one comes from a limit case analysis. Except for the normalizing factor μ , QIHT at 1-bit ($b = 1$) reduces to BIHT [11]. Moreover, when $b \gg 1$, $\mathcal{Q}_b[\mathbf{z}] \simeq \mathbf{z}$ for $\mathbf{z} \in \mathbb{R}^M$ and we recover the IHT algorithm. These limit cases are consistent with the previous observations made above on the asymptotic behaviors of $\mathcal{J}$ in these two cases.

Second, as for the modified Subspace Pursuit algorithm [3], QIHT is designed for improving the quantization consistency of the current iterate with the quantized observations. For the moment, the importance of this improvement can only be understood in 1-bit. Given $0 < \delta < 1$, when $M = O(\delta^{-1}K \log N)$ and with high probability on the drawing of a random Gaussian matrix $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$, $\|\frac{\mathbf{a}}{\|\mathbf{a}\|} - \frac{\mathbf{b}}{\|\mathbf{b}\|}\| \leq \delta$ if $\mathcal{Q}_1(\Phi\mathbf{a}) = \mathcal{Q}_1(\Phi\mathbf{b})$ for all $\mathbf{a}, \mathbf{b} \in \Sigma_K$ [11]. Actually, it is shown in Appendix A that if no more than r components differ between $\mathcal{Q}_1(\Phi\mathbf{a})$ and $\mathcal{Q}_1(\Phi\mathbf{b})$, then, with high probability on Φ ,

$$\|\frac{\mathbf{a}}{\|\mathbf{a}\|} - \frac{\mathbf{b}}{\|\mathbf{b}\|}\| \leq (\frac{K+r}{K})\delta, \quad (7)$$

for $M = O(\delta^{-1}K \log MN)$. We understand then the beneficial impact of any increase of consistency between $\mathcal{Q}_1(\Phi\mathbf{x}^{(n)})$ and $\mathbf{y}$ at each QIHT iteration.

Third, the adjustment of μ , which is decisive for QIHT efficiency, leads also to some interesting observations. Extensive simulations not presented here pointed us that, for $\Phi \sim \mathcal{N}^{M \times M}(0, 1)$, $\mu \propto 1/M$ seems to be a universal rule of efficiency at any bit rate. Interestingly, this setting was already characterized for IHT where $\mu \simeq 1/(1 + \delta_{2K})$ if the sensing matrix respects the RIP property with radius δ_{2K} [18]. Since $\Phi/\sqrt{M}$ is RIP for $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$ as soon as $M = O(K \log N/K)$ this is equivalent to impose $\mu \simeq 1/M$.

At the other extreme, the rule $\mu \propto 1/M$ is also consistent with the following 1-bit analysis. In [13], it is shown that the mapping $\mathbf{u} \rightarrow \text{sign}(\Phi\mathbf{u})$ respects an interesting property that we arbitrary call *sign product embedding*¹ (SPE):

Proposition 1. *Given $0 < \delta < 1$, there exist two constants $c, C > 0$ such that, if $M \geq C\delta^{-6}K \log N/K$, then, with a probability higher than $1 - 8 \exp(-c\delta^2 M)$, $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$ satisfies*

$$|\mu^*(\text{sign}(\Phi\mathbf{u}), \Phi\mathbf{v}) - \langle \mathbf{u}, \mathbf{v} \rangle| \leq \delta, \quad \forall \mathbf{u}, \mathbf{v} \in \Sigma_K^*, \quad (8)$$

with $\mu^* = 1/(q_0 M)$. When $\mathbf{u}$ is fixed, the condition on M is relaxed to $M \geq C\delta^{-2}K \log N/K$.

When Φ respects (8), we simply write that Φ is SPE(Σ_K^*, δ). When $\mathbf{u}$ is fixed, we say that Φ is locally SPE(Σ_K^*, δ) on $\mathbf{u}$. This SPE property leads to an interesting phenomenon.

¹In [13], more general embeddings than this of Σ_K are studied.




Figure 2: Comparison between (from left to right) BPDN, IHT and QIHT for several quantization scenarios. The SNR is expressed in dB as a function of the bit budget $\mathfrak{B}$ and the number of bits b used to quantize each measurement.

Proposition 2. *Given $\mathbf{x} \in \Sigma_K^*$ and let $\Phi \in \mathbb{R}^{M \times N}$ be a matrix respecting the local SPE(Σ_{2K}^*, δ) on $\mathbf{x}$ for some $0 < \delta < 1$. Then, given $\mathbf{y} = \mathcal{Q}_1[\Phi\mathbf{x}] = q_0 \text{sign}(\Phi\mathbf{x})$, the vector*

$$\hat{\mathbf{x}} := \frac{1}{q_0^2 M} \mathcal{H}_K(\Phi^* \mathbf{y}),$$

satisfies $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq 2\delta$.

Proof. Let us define $\mathcal{T}_0 = \text{supp } \mathbf{x}$, $\mathcal{T} = \mathcal{T}_0 \cup \text{supp } \hat{\mathbf{x}}$, and $\mathbf{a} = \frac{1}{q_0^2 M} \Phi^* \mathbf{y} = \mu^* \Phi^* \text{sign}(\Phi\mathbf{x})$ with $\hat{\mathbf{x}} = \mathcal{H}_K(\mathbf{a})$. Then $\hat{\mathbf{x}}$ is also the best K -term approximation $\mathbf{a}_{\mathcal{T}} = \Phi_{\mathcal{T}}^* \mathbf{y}$, so that $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq \|\mathbf{x} - \mathbf{a}_{\mathcal{T}}\| + \|\hat{\mathbf{x}} - \mathbf{a}_{\mathcal{T}}\| \leq 2\|\mathbf{x} - \mathbf{a}_{\mathcal{T}}\|$. Therefore, since $\|\mathbf{x} - \mathbf{a}_{\mathcal{T}}\| = \sup_{\mathbf{w} \in \Sigma_{\mathcal{T}}^*} \langle \mathbf{w}, \mathbf{x} - \mathbf{a}_{\mathcal{T}} \rangle$ and Φ is SPE(Σ_{2K}^*, δ), $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq 2 \sup_{\mathbf{w} \in \Sigma_{\mathcal{T}}^*} (\langle \mathbf{w}, \mathbf{x} \rangle - \mu^* \langle \Phi \mathbf{w}, \text{sign}(\Phi\mathbf{x}) \rangle) \leq 2 \sup_{\mathbf{w} \in \Sigma_{\mathcal{T}}^*} (\langle \mathbf{w}, \mathbf{x} \rangle - \langle \mathbf{w}, \mathbf{x} \rangle + \delta) = 2\delta$, using $\text{supp}(\mathbf{x} - \mathbf{a}_{\mathcal{T}}) \subseteq \mathcal{T}$ with $\#\mathcal{T} \leq 2K$. $\square$

This proposition shows that a single hard thresholding of $\frac{1}{q_0^2 M} \Phi^* \mathbf{y}$ already provides a good estimation of $\mathbf{x}$. Actually, from the condition on M for reaching the local SPE, we deduce that $\|\mathbf{x} - \hat{\mathbf{x}}\| = O(\sqrt{K/M})$. This is quite satisfactory for such a simple $\mathbf{x}$ estimation and it suggests setting $\mu \propto 1/M$ in QIHT for $b = 1$ where $\hat{\mathbf{x}}$ is related to $\mathbf{x}^{(1)}$.

Noticeably, it has been recently observed in [14] that $\hat{\mathbf{x}}' := \hat{\mathbf{x}}/\|\hat{\mathbf{x}}\|$ is actually solution of

$$\arg\max_{\mathbf{u} \in \mathbb{R}^N} \langle \mathbf{y}, \Phi \mathbf{u} \rangle \quad \text{s.t.} \quad \|\mathbf{u}\|_0 \leq K,$$

for which there exists the weaker error bound $\|\mathbf{x} - \hat{\mathbf{x}}'\|^2 = O(\sqrt{K/M})$ when $\mathbf{x}$ is fixed [13].

6 Experiments

An extensive set of simulations has been designed for evaluating the efficiency of QIHT in comparison with two other methods more suited to high-resolution quantization, namely, IHT and BPDN. Our objective is to show that QIHT provides better quality results at least at small quantization levels. For all experiments, we set $N = 1024$, $K = 16$ and the K -sparse signals were generated by choosing their supports uniformly at random amongst the $\binom{N}{K}$ available ones, while their non-zero coefficients were drawn uniformly at random on the sphere $S^{K-1} \subseteq \mathbb{R}^K$. For each algorithm, 100 initial such sparse vectors were generated and the reconstruction method was tested for $1 \leq b \leq 5$ and for $\mathfrak{B} = bM \in \{64, 128, \dots, 1280\}$, *i.e.*, approximately fixing $M = \lfloor \mathfrak{B}/b \rfloor$. For each experimental condition, the quantized M -dimensional measurement vectors $\mathbf{y}_b$ was generated as in (2) with a random sensing matrix $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$ and according to an optimal Lloyd-Max b -bits Quantizer $\mathcal{Q}_b$ (Sec. 3). IHT and QIHT iterations were both stopped at step n as soon as $\|\mathbf{x}^{(n+1)} - \mathbf{x}^{(n)}\| \|\mathbf{x}^{(n+1)}\|^{-1} < 10^{-4}$ or if $n = 1000$. The BPDN algorithm was solved with the SPGL1 MATLAB toolbox [21]. In IHT and QIHT, signal sparsity K was assumed known and both were set with $\mu = \frac{1}{M} (1 - \sqrt{\frac{2K}{M}})$. This fits the IHT condition $\mu < 1/(1 + \delta_{2K})$ mentioned in Sec. 5 by assuming that the RIP radius δ_{2K} behaves like $\sqrt{2K/M}$, which is a common assumption in CS. For BPDN, the noise energy was given by an oracle installing BPDN in the best reconstruction scenario, *i.e.*, $\epsilon = \|\Phi \mathbf{x}_0 - \mathbf{y}\|_2$. Whatever the reconstruction method, given an initial signal $\mathbf{x}_0 \in \Sigma_K^*$ and its reconstruction $\mathbf{x}^*$, the reconstruction quality was measured

by $\text{SNR}(\mathbf{x}_0, \mathbf{x}^*) = -20 \log_{10} \|\mathbf{x}_0 - \|\mathbf{x}^*\|^{-1} \mathbf{x}^*\|$. In other words, we focus here on a good “angular” estimation of the signals, adopting therefore a common metric for $b > 1$ and for $b = 1$, where amplitude information is lost. Finally, for each method and each couple of (M, b) , the SNR was averaged over the 100 test signals and expressed in dB.

Fig. 2 gathers the SNR performances of the 3 methods as a function of $\mathfrak{B}$. QIHT outperforms both BPDN and IHT for the selected scenarios, especially for low bit quantizers. At high resolution, the gain between QIHT and IHT decreases as expected from the limit case analysis of QIHT. We can also notice that, first, there is almost no quality difference between QIHT at $b = 1$ and $b = 2$. This could be due to a non-optimality of the Lloyd-Max quantizer with respect to QIHT reconstruction error minimization. Second, BPDN and IHT asymptotically present the “6dB per bit” gain, while QIHT hardly exhibits such behavior only when $b = 4 \rightarrow 5$.

Finally, in order to test Prop. 2, the SNR reached by the single thresholding solution $\hat{\mathbf{x}}$ is plotted in dashed in Fig 2-right. Despite its poor behavior compared to QIHT at $b = 1$, it outperforms BPDN at high $\mathfrak{B} = M$ with a $\text{SNR} \geq 10\text{dB}$ at $M = N = 1024$. A curve fitting (no shown here) shows that this SNR increases a bit faster than $20 \log_{10} \sqrt{K/M} + O(1)$.

7 Conclusion

We have introduced the QIHT algorithm as a generalization of the BIHT and IHT algorithms aiming at enforcing consistency with quantized observations at any bit resolution. In particular, we showed that the almost obvious inclusion of the quantization operator in the IHT recursion is actually related to the implicit minimization of a particular inconsistency cost $\mathcal{E}_b$. This function generalizes the one-sided ℓ_1 cost of BIHT and asymptotically converges to the quadratic fidelity minimized by IHT. There is still a hard work to be performed in order to prove QIHT convergence and stability. However, the different ingredients defining it, as $\mathcal{E}_b$, deserve independent analysis extending previous 1-bit embeddings developed in [11–13].

A Proximity of almost 1-bit consistent sparse vectors

The relation (7) is induced by the following theorem and by its subsequent Corollary 1. These use the normalized *Hamming distance* between two strings $\mathbf{a}, \mathbf{b} \in \{-1, +1\}^M$ defined by $d_H(\mathbf{a}, \mathbf{b}) = \frac{1}{M} \sum_{i=1}^M a_i \oplus b_i$, where $\oplus$ is the XOR operation such that $a_i \oplus b_i$ equals 0 if $a_i = b_i$ and 1 otherwise. For shortening the notations, we define also $\varphi(\mathbf{u}) := \text{sign}(\Phi \mathbf{u}) \in \{-1, +1\}^M$ for $\mathbf{u} \in \mathbb{R}^N$.

Theorem 1. *Let $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$. Fix $r \leq M/2$, $0 \leq \eta \leq 1$ and $0 < \delta < 1$. If the number of measurements M satisfies*

$$M - r \geq \frac{2}{\delta} \left(2K \log(N) + r \log(M) + 4K \log\left(\frac{17}{\delta}\right) + \log \frac{2e}{\eta} \right), \quad (9)$$

then,

$$\forall \mathbf{a}, \mathbf{b} \in \Sigma_K^*, \quad d_H(\varphi(\mathbf{a}), \varphi(\mathbf{b})) \leq \frac{r}{M} \quad \Rightarrow \quad \|\mathbf{a} - \mathbf{b}\| \leq \delta$$

with probability exceeding $1 - \eta$.

This improves the previous theorem proved in [11].

Proof. First, notice that if $M d_H(\varphi(\mathbf{a}), \varphi(\mathbf{b})) \leq r$, there exists a $\mathcal{T} \subset [M]$ with $|\mathcal{T}| \geq M - r$ such that $\varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b})$. Let $[M]_r$ be the set of subsets of $[M]$ whose size is bigger than $M - r$. Using a union bound argument, we have

$$\begin{aligned} & \mathbb{P} \left[\exists \mathcal{T} \subset [M]_r, \exists \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| > \delta \right] \\ & \leq \bigcup_{\mathcal{T} \subset [M]_r} \mathbb{P} \left[\exists \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| > \delta \right]. \end{aligned} \quad (10)$$

We know from [11, Theorem 2] that, as soon as

$$M' \geq \frac{2}{\delta} \left(2K \log(N) + 4K \log\left(\frac{17}{\delta}\right) + \log \frac{1}{\eta} \right),$$

the random generation of $\Phi' \sim \mathcal{N}^{M' \times N}(0, 1)$ fullfils

$$\mathbb{P} \left[\exists \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi'(\mathbf{a}) = \varphi'(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| > \delta \right] \leq \eta,$$

with $\varphi'(\cdot) = \text{sign}(\Phi' \cdot)$. Therefore, for a given $\mathcal{T} \subset [M]_r$ and by setting $\Phi' = (\mathbf{Id}_{\mathcal{T}})^T \Phi$, *i.e.*, the matrix obtained by restricting Φ to the rows indexed in $\mathcal{T}$, we have $\varphi_{\mathcal{T}} = \varphi'$, $M' = |\mathcal{T}| \geq M - r$ and

$$\mathbb{P}[\exists \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| > \delta] \leq \eta$$

if $M \geq r + \frac{2}{\delta} (2K \log(N) + 4K \log(\frac{17}{\delta}) + \log \frac{1}{\eta})$.

Under the same condition on M , and observing that, for $r \leq \lfloor M/2 \rfloor$, $|[M]_r| = \sum_{k=0}^r \binom{M}{M-k} \leq (r+1) \binom{M}{r} \leq (r+1)(eM/r)^r$, (10) provides

$$\mathbb{P}[\exists \mathcal{T} \subset [M]_r, \exists \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| > \delta] \leq (r+1)(\frac{eM}{r})^r \eta.$$

Analyzing the complementary event and redefining $\eta \leftarrow (r+1)(eM/r)^r \eta$, we get finally

$$\begin{aligned} & \mathbb{P}[\forall \mathbf{a}, \mathbf{b} \in \Sigma_K^* : M d_H(\varphi_{\mathcal{T}}(\mathbf{a}), \varphi_{\mathcal{T}}(\mathbf{b})) \leq r, \|\mathbf{a} - \mathbf{b}\| \leq \delta] \\ &= \mathbb{P}[\forall \mathcal{T} \subset [M]_r, \forall \mathbf{a}, \mathbf{b} \in \Sigma_K^* : \varphi_{\mathcal{T}}(\mathbf{a}) = \varphi_{\mathcal{T}}(\mathbf{b}), \|\mathbf{a} - \mathbf{b}\| \leq \delta] \geq 1 - \eta, \end{aligned}$$

as soon as

$$M \geq r + \frac{2}{\delta} (2K \log(N) + r \log(M) + 4K \log(\frac{17}{\delta}) + \log \frac{2e}{\eta}).$$

□

Corollary 1. *Let $\Phi \sim \mathcal{N}^{M \times N}(0, 1)$. Fix $r \leq M/2$, $0 \leq \eta \leq 1$ and $0 < \delta < 1$. If the number of measurements M satisfies*

$$M \geq \frac{2}{\delta} (2K \log(\max(N, M)) + 4K \log(\frac{17}{\delta}) + \log \frac{2e}{\eta}),$$

then,

$$\forall \mathbf{a}, \mathbf{b} \in \Sigma_K^*, \quad d_H(A(\mathbf{a}), A(\mathbf{b})) \leq \frac{r}{M} \quad \Rightarrow \quad \|\mathbf{a} - \mathbf{b}\| \leq \frac{K+r}{K} \delta$$

with probability exceeding $1 - \eta$.

Proof. The proof is obtained from Theorem 1 by redefining $\delta \leftarrow \frac{K+r}{K} \delta$, observing that $\frac{2}{\delta} \log M \geq 1$ if $M > 1$ and by slightly enforcing the condition (9) on M . □

References

- [1] D. L. Donoho, “Compressed Sensing,” *IEEE Trans. Inf. Th.*, **52**(4):1289–1306, 2006.
- [2] E. J. Candès, J. Romberg and T. Tao. “Stable signal recovery from incomplete and inaccurate measurements,” *Comm. Pure Appl. Math.*, **59**(8):1207–1223, 2006
- [3] W. Dai, H. V. Pham, and O. Milenkovic, “Information theoretical and algorithmic approaches to quantized compressive sensing,” *IEEE Trans. Comm.*, **59**(7):1857–1866, 2011.
- [4] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing,” *App. Comp. Harm. Anal.*, **31**(3): 429–443, 2011.
- [5] A. Zymnis, S. Boyd, and E. Candès, “Compressed sensing with quantized measurements,” *IEEE Sig. Proc. Let.*, **17**(2): 149–152, 2010.
- [6] L. Jacques, D. K. Hammond, and M. J. Fadili, “Dequantizing Compressed Sensing: When Over-sampling and Non-Gaussian Constraints Combine,” *IEEE Trans. Inf. Th.*, **57**(1): 559–571, 2011.
- [7] L. Jacques, D. K. Hammond, and M. J. Fadili, “Stabilizing nonuniformly quantized compressed sensing with scalar companders,” *arXiv:1206.6003*, 2012.
- [8] J. Z. Sun and V. K. Goyal, “Optimal quantization of random measurements in compressed sensing,” in *Int. Symp. Inf. Th. (ISIT)*, June 2009.
- [9] C. Sinan Güntürk, M Lammers, A. M. Powell, R. Saab, and Ö. Yılmaz, “Sobolev duals for random frames and $\Sigma\Delta$ quantization of compressed sensing measurements,” *Found. Comp. Math.*, **13**(1):1–36, 2013.

- [10] P. Boufounos and R. Baraniuk, “1-bit compressive sensing,” in *Proc. Conf. Inform. Sc. Sys. (CISS)*, Princeton, NJ, Mar. 2008.
- [11] 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 Trans. Inf. Th.*, **59**(4):2082-2102, 2013.
- [12] Y. Plan and R. Vershynin, “One-bit compressed sensing by linear programming,” *Comm. Pure Appl. Math.*, 2013.
- [13] Y. Plan and R. Vershynin, “Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach,” *IEEE Trans. Inf. Th.*, **59**(1):482-494, 2013.
- [14] S. Bahmani, P. T. Boufounos and B. Raj, “Robust 1-bit Compressive Sensing via Gradient Support Pursuit,” *arXiv:1304.6627*, 2013.
- [15] Z. Yang, L. Xie, and C. Zhang, “Unified framework and algorithm for quantized compressed sensing,” *arXiv:1203.4870*, 2012.
- [16] T. Blumensath and M.E. Davies, “Iterative hard thresholding for compressed sensing,” *App. Comp. Harm. Anal.*, **27**(3):265-274, 2009.
- [17] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic Decomposition by Basis Pursuit,” *SIAM J. Sc. Comp.*, **20**(1):33-61, 1998.
- [18] T. Blumensath, “Accelerated iterative hard thresholding,” *Sig. Proc.*, **92**(3):752-756, 2012.
- [19] R. M. Gray and D. L. Neuhoff, “Quantization,” *IEEE Trans. Inf. Th.*, **44**(6):2325-2383, 1998.
- [20] R.T. Rockafellar, *Convex analysis*, vol. 28, Princeton Univ. Press, 1970.
- [21] 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>.