

WEIGHTED FIDELITY IN NON-UNIFORMLY QUANTIZED COMPRESSED SENSING

Laurent Jacques^{1,3}, David Kenric Hammond², Jalal Fadili³

¹ICTEAM/ELEN, Université catholique de Louvain (UCL), Belgium. ²Neuroinformatics Center, University of Oregon, USA.

³GREYC CNRS-ENSICAEN-Université de Caen, 14050 Caen France

ABSTRACT

Following the Compressed Sensing (CS) paradigm, this paper studies the problem of recovering sparse or compressible signals from (scalar) non-uniformly quantized measurements. We show that a simple adaptation of the Basis Pursuit De-Quantizer introduced earlier, that is, a sign sensitive weighting of their ℓ_p -norm fidelity constraint, yields good SNR improvements in the signal reconstruction. As a good indication of this improvement origin, we prove theoretically that a similar decoder, using a particular *side-position-to-level* oracle, displays a reduction of the reconstruction error when both the number of measurements and the moment p of the constraint increase. This follows the oversampling principle underlined in our previous study for uniformly quantized CS, with an additional gain provided by the non-uniform quantization. We conclude this paper by showing the efficiency of the approach on 1-D and 2-D signal examples.

1. INTRODUCTION

The recent theory of Compressed Sensing (CS) [1, 2] shows how sparse or compressible signals can be reconstructed from few linear measurements compared to the dimension N of the signal space. The gist of this approach relies in the use of a sensing basis sufficiently incoherent with the sparsity basis of the signal. This happens with high probability for a large class of random matrix constructions as soon as the number of measurements M is higher than “few multiples” of the signal sparsity K . For instance, for Random Gaussian matrices, $M = O(K \log N/K)$.

Similarly to recent studies [3–5] in the CS literature, this work is interested in controlling the signal reconstruction stability when the compressive measurements undergo a scalar quantization, possibly non-uniform, of given rate R .

More precisely, given a signal $\mathbf{x} \in \mathbb{R}^N$, we first assume it to be *sparse*, or *sparsely approximable* (compressible), in a certain orthogonal basis $\Psi = (\Psi_1, \dots, \Psi_N) \in \mathbb{R}^{N \times N}$ (e.g., in the wavelet basis or in the pixel domain). In other words, this signal is decomposed as $\mathbf{x} = \Psi \mathbf{c} = \sum_j \Psi_j c_j$ with an approximation error $\|\mathbf{c} - \mathbf{c}_K\|$ (\mathbf{c}_K being the best K term approximation of \mathbf{c}) quickly decreasing when the integer

K increases. For the sake of simplicity, the sparsity basis is assumed to be canonical ($\Psi = \mathbb{1}$) and \mathbf{c} is identified with \mathbf{x} . All the results can be easily extended to the situation $\Psi \neq \mathbb{1}$.

Second, we are interested in the Compressed Sensing of $\mathbf{x} \in \mathbb{R}^N$ with a certain sensing matrix $\Phi \in \mathbb{R}^{M \times N}$ [1, 2]. Each compressed sensing measurement, *i.e.*, each component of the measurement vector $\Phi \mathbf{x}$, undergoes a general scalar (uniform or non-uniform) *quantization* \mathcal{Q} described in Section 2, *i.e.*, our sensing model is

$$\mathbf{y} = \mathcal{Q}[\Phi \mathbf{x}]. \quad (1)$$

Conventions: All space dimensions are denoted by capital letters (e.g., $K, M, N, D \in \mathbb{N}$), vectors and matrices are written in bold symbols. For any vector $\mathbf{u} = (u_1, \dots, u_D)^T \in \mathbb{R}^D$ (with $(\cdot)^T$ the transposition), the ℓ_p -norm ($p \geq 1$) of \mathbf{u} is $\|\mathbf{u}\|_p^p = \sum_i |u_i|^p$, with $\|\mathbf{u}\| = \|\mathbf{u}\|_2$ and $\|\mathbf{u}\|_0 = \#\{i : u_i \neq 0\}$ the ℓ_0 (“not-a”) norm of \mathbf{u} . We denote also $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^D$, and $\mathbb{1}$ for the identity matrix. Given a vector \mathbf{u} , $\mathbf{U} = \text{diag } \mathbf{u}$ is the diagonal matrix such that $U_{ii} = u_i$.

2. QUANTIZATION FRAMEWORK

Our operator \mathcal{Q} of interest here is a scalar quantizer of vector components. We do not impose this quantization to be uniform, that is, the quantization bin width is not necessarily constant with respect to the bin level (as in Fig. 1-left).

More precisely, \mathcal{Q} relies on the definition of a set of $B = 2^R$ levels ω_k (coded on $R = \log_2 B$ bits) and of a set of $B + 1$ thresholds $t_k \in \mathbb{R} \cup \{\pm\infty\}$, with $\omega_k < \omega_{k+1}$ and $t_k \leq \omega_k < t_{k+1}$ for all $1 \leq k \leq B$ and given a rate $R \in \mathbb{N}$. The k^{th} quantizer *bin* (or *region*) is $\mathcal{R}_k = [t_k, t_{k+1})$. As illustrated on Fig. 1-left, the quantizer is a mapping between \mathbb{R} and the set of levels $\Omega = \{\omega_k : 1 \leq k \leq B\}$, *i.e.*, $\mathcal{Q}[\lambda] = \omega_k \Leftrightarrow \lambda \in \mathcal{R}_k = \mathcal{Q}^{-1}[\omega_k]$. For vectors $\mathbf{u} \in \mathbb{R}^M$, $\mathcal{Q}[\mathbf{u}] \in \Omega^M$ with $(\mathcal{Q}[\mathbf{u}])_k = \mathcal{Q}[u_k] \in \Omega$, for $1 \leq k \leq M$.

Generally, whatever the definition of the quantization levels and thresholds, many scalar quantization schemes relies on a High Resolution/Rate Assumption (HRA) [6]. They assume that the distortion noise, that is, the difference between the initial and the quantized values, is uniform within each quantization bins, possibly with different widths. Therefore, the creation of the non-uniform quantizer, which can rely on a non-uniformity assumption (like with Lloyd-Max itera-

LJ is funded by the Belgian Science Policy (Return Grant, BELSPO) joined to the Belgian Interuniversity Attraction Pole IAP-VI BCRYPT.

tive method [6]) is distinguished from the uniform distortion model induced by this quantizer.

In this work, we define a stronger variant of the HRA that we dub Asymmetric HRA (AHRA). We will assume that, within each bin, the quantization distortion follows two distinct uniform distributions on the left and on the right of the corresponding level.

3. ASYMMETRIC QUANTIZATION CONSISTENCY

Beyond any stochastic modeling of the quantization distortion, we know that $\mathcal{Q}[t] = \omega_i$ is equivalent to the condition $t \in \mathcal{R}_i = \mathcal{Q}^{-1}[\omega_i] = [t_i, t_{i+1})$. In other words,

$$\mathcal{Q}[t] = \omega_i \Leftrightarrow \begin{cases} t - \omega_i < t_{i+1} - \omega_i & \text{if } t \geq \omega_i \\ t - \omega_i \geq t_i - \omega_i & \text{if } t < \omega_i \end{cases}. \quad (2)$$

This condition can be simplified thanks to the following *asymmetric norm* [7]. Let $\mathbf{S} = \text{diag}(\mathbf{s}) \in \mathbb{R}_+^{2D \times 2D}$ be a diagonal matrix of positive diagonal entries $\mathbf{s} = (\mathbf{s}^+, \mathbf{s}^-)^T \in \mathbb{R}_+^{2D}$, with $\mathbf{s}^\pm \in \mathbb{R}^D$. The sign-dependent weighted ℓ_p norm is given by

$$\|\mathbf{u}\|_{p,\mathbf{s}} \triangleq \left\| \mathbf{S} \begin{pmatrix} (\mathbf{u})_+ \\ (-\mathbf{u})_+ \end{pmatrix} \right\|_p, \quad p \in [1, \infty], \quad (3)$$

$(\mathbf{u})_+$ being the non-negative vector such that $((\mathbf{u})_+)_i = (u_i)_+ = \max(u_i, 0)$. It is easy to check that $\|\cdot\|_{p,\mathbf{s}}$ is an asymmetric norm but not a norm¹. Moreover, for $p \geq 1$, $\|\lambda \cdot\|_{p,\mathbf{s}}$ is convex.

Let us define $\Delta_i^+ = t_{i+1} - \omega_i$, $\Delta_i^- = \omega_i - t_i$, and $k(y_i)$, the quantization bin index of the i^{th} component of \mathbf{y} , that is, such that $y_i \in \mathcal{R}_{k(y_i)}$. Our asymmetric weighting matrix is $\mathbf{S}(\mathbf{y}) = \text{diag}(\mathbf{s}(\mathbf{y})) \in \mathbb{R}^{2M \times 2M}$, with $\mathbf{s} = \mathbf{s}(\mathbf{y}) = (\mathbf{s}^+(\mathbf{y}), \mathbf{s}^-(\mathbf{y}))^T$ and $s_i^\pm(\mathbf{y}) = 1/\Delta_{k(y_i)}^\pm$. Therefore, from (2),

$$\mathcal{Q}[\Phi \mathbf{x}] = \mathbf{y} \Leftrightarrow \|\Phi \mathbf{x} - \mathbf{y}\|_{\infty, \mathbf{s}} \leq 1. \quad (\text{QC})$$

This is the *Quantization Consistency* (QC) constraint that any reconstructed vector \mathbf{x}^* should respect in order to be consistent with the quantized measurement \mathbf{y} . Notice that keeping the same \mathbf{S} , $\lim_{p \rightarrow \infty} \|\Phi \mathbf{x} - \mathbf{y}\|_{p,\mathbf{s}} = \|\Phi \mathbf{x} - \mathbf{y}\|_{\infty, \mathbf{s}} \leq 1$.

4. RECONSTRUCTION AND WEIGHTED FIDELITY

There is a straightforward way for estimating a signal \mathbf{x} sensed through the model (1). We can indeed alter the Basis Pursuit DeQuantizer introduced in [5] by replacing their ℓ_p -norm fidelity constraint by the weighted asymmetric norm introduced in (3), that is,

$$\arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \|\Phi \mathbf{u} - \mathbf{y}\|_{p,\mathbf{s}} \leq \epsilon, \quad (\text{WBPDQ})$$

¹There is no positive homogeneity: $\|\lambda \mathbf{u}\|_{p,\mathbf{s}} \neq |\lambda| \|\mathbf{u}\|_{p,\mathbf{s}}$ for $\lambda < 0$.

with \mathbf{s} defined as in Sec. 3. Notice that for $p = 2$ and $\mathbf{s}^\pm = \mathbf{1}$, WBPDQ reduces to the Basis Pursuit DeNoiser (BPDN) [1, 2]. As for the BPDQ decoders, we are going to discover which moment p minimizes the reconstruction error; and the answer is not necessarily $p = \infty$ despite the QC relation!

Unfortunately, in spite of good numerical results (see Sec. 7), we did not find a convincing way for bounding theoretically the WBPDQ approximation error. We were instead able to characterize the behavior of the following auxiliary program helped with the *side-position-to-level* (SPTL) oracle $\sigma = \text{sign}(\Phi \mathbf{x} - \mathbf{y}) \in \{\pm 1\}^M$:

$$\mathbf{x}^* = \arg \min_{\mathbf{u} \in \mathbb{R}^N} \|\mathbf{u}\|_1 \text{ s.t. } \begin{cases} \|\Phi \mathbf{u} - \mathbf{y}\|_{p,\mathbf{w}} \leq \epsilon, \\ \text{sign}(\Phi \mathbf{u} - \mathbf{y}) = \sigma, \end{cases} \quad (4)$$

with $\|\cdot\|_{p,\mathbf{w}} \triangleq \|\text{diag}(\mathbf{w}) \cdot\|_p$, $\mathbf{w} = \mathbf{w}(\mathbf{s}, \sigma) \in \mathbb{R}^M$ being the weighting vector such that $w_i = s_i^+(\sigma_i)_+ + s_i^-(-\sigma_i)_+$.

Similarly to what we discovered in [5], the stability of (4) depends on the good behavior of Φ in the normed space $\ell_{p,\mathbf{w}}^M = (\mathbb{R}^M, \|\cdot\|_{p,\mathbf{w}})$. In particular, we say that, given a weight vector $\mathbf{w}' \in \mathbb{R}_+^M$, a matrix $\Phi \in \mathbb{R}^{M \times N}$ satisfies the *Restricted Isometry Property from $\ell_{p,\mathbf{w}'}^M$ to ℓ_2^N* at order $K \in \mathbb{N}$, radius $0 \leq \delta < 1$ and for a normalization $\mu = \mu(p, M, N) > 0$, if for all $\mathbf{x} \in \Sigma_K = \{\mathbf{u} \in \mathbb{R}^N : \|\mathbf{u}\|_0 \leq K\}$,

$$(1 - \delta)^{1/2} \|\mathbf{x}\| \leq \frac{1}{\mu} \|\Phi \mathbf{x}\|_{p,\mathbf{w}'} \leq (1 + \delta)^{1/2} \|\mathbf{x}\|, \quad (5)$$

We will write shortly that Φ is $\text{RIP}_{\ell_{p,\mathbf{w}'}, \ell_2^N}^M(K, \delta, \mu)$. Of course, the common RIP and the $\text{RIP}_{p,2}$ [5] are obtained with $\mathbf{w}' = \mathbf{1}$ and $p \geq 2$. We found that a small modification of the proof of Proposition 1 in [5] allows us to check that, with very high (controllable) probability, a Standard Gaussian Random (SGR) matrix $\Phi \in \mathbb{R}^{M \times N}$ with $\Phi_{ij} \sim \mathcal{N}(0, 1)$ is $\text{RIP}_{\ell_{p,\mathbf{w}'}, \ell_2^N}^M(K, \delta, \mu)$ as soon as

$$M^{2/p} = O(K \log N / K) \text{ and } \mu = \mathbb{E} \|\xi\|_{p,\mathbf{w}},$$

for $\xi \in \mathbb{R}^M$ a SGR vector

Obviously, if Φ is $\text{RIP}_{\ell_{p,\mathbf{w}'}, \ell_2^N}^M$, then $\Phi' = \text{diag}(\mathbf{w}) \Phi$ is $\text{RIP}_{\ell_p^M, \ell_2^N}^M$, while $\|\Phi \mathbf{u} - \mathbf{y}\|_{p,\mathbf{w}} = \|\Phi' \mathbf{u} - \text{diag}(\mathbf{w}) \mathbf{y}\|_p$. Therefore, despite the complementary constraint on the sign σ , the stability proved in [5] holds for (4).

Theorem 1. *Let $\mathbf{x} \in \mathbb{R}^N$ be a signal with a K -term ℓ_1 -approximation error $e_0(K) = K^{-\frac{1}{2}} \|\mathbf{x} - \mathbf{x}_K\|_1$, for $0 \leq K \leq N$. Let Φ be a $\text{RIP}_{\ell_{p,\mathbf{w}'}, \ell_2^N}^M(s, \delta_s, \mu)$ matrix for $s \in \{K, 2K, 3K\}$ and $2 \leq p < \infty$. If \mathbf{x} is a feasible point of the constraints in (4), then*

$$\|\mathbf{x}^* - \mathbf{x}\|_2 \leq A_p e_0(K) + B_p \mu^{-1} \epsilon, \quad (6)$$

for values $A_p(\Phi, K) = \frac{2(1+C_p-\delta_{2K})}{1-\delta_{2K}-C_p}$, $B_p(\Phi, K) = \frac{4\sqrt{1+\delta_{2K}}}{1-\delta_{2K}-C_p}$, $C_p = O(\sqrt{(\delta_{2K} + \delta_{3K})(p-2)})$ as $p \gg 2$ and $C_p = \delta_{3K} + O(p-2)$ as $p \rightarrow 2$.

²The precise definition of C_p is given in [5].




Fig. 1: Left: Quantizing a Gaussian distribution with Lloyd-Max quantizer with $R = 3$. Right: Estimation of ρ_p^{\min} for $M = 320$, $R = 4$, and $p \in [2, 20]$. Lloyd-Max quantizer (plain); optimal uniform quantizer (dashed).

5. OVERSAMPLING EFFECT

An interesting effect occurs in (6) when the sensing is oversampled, that is, when M/K is large enough to allow us to select high moment p in (4) and still have the RIP $_{\ell_p^M, \ell_2^N}$ of the SGR sensing matrix Φ . Indeed, we are going to show that the error term $\mu^{-1}\epsilon$ in (6) decreases as $1/\sqrt{p+1}$ when p increases, and it is further reduced when the weights w are well adjusted.

We first need to find an estimator for ϵ in the sensing model (1). This is done by observing when, with high probability, x is a feasible point of the constraints in (4). We notice that $\|\Phi x - y\|_{p,w}^p = \sum_i w_i^p(\sigma_i) |(\Phi x)_i - y_i|^p$. From the AHRA, each $|w_i(\sigma_i)| |(\Phi x)_i - y_i|$ can be modeled as a uniform random variable on $[0, 1]$. We can then take $\alpha = 2$ in Lemma 3 of [5] and determine that

$$\|\Phi x - y\|_{p,w} \leq \epsilon \triangleq \epsilon_p(M) = ((p+1)^{-1}M + \kappa\sqrt{M})^{1/p}$$

holds with probability higher than $1 - e^{-2\kappa^2}$.

Second, we must lower bound the RIP $_{\ell_p^M, \ell_2^N}$ normalization μ . For this purpose, let us assume that the weights $w = G(M) \in \mathbb{R}_+^M$ have been generated by a particular weight generator $G(M)$ (for instance as an output of a Lloyd-Max quantizer) respecting the following property.

Definition 1. *The generator G (and by extension $w = G(M)$) has the Converging Moments (CM) property if, for any $p \geq 1$, there is a M_0 such that*

$$\rho_p^{\min} \leq M^{-1/p} \|G(M)\|_p \leq \rho_p^{\max}, \quad \forall M \geq M_0, \quad (7)$$

where $\rho_p^{\min} > 0$ and $\rho_p^{\max} > 0$ are, respectively, the biggest and the smallest values such that (7) holds.

The CM property makes sense for instance if all the weights $\{w_i\}_{1 \leq i \leq M}$ are taken (with repetition) inside a finite set of values (of size independent of M).

A simple modification of Lemma 1 in [5] shows that, if $M \geq 2\beta^{-1} (2/\rho_p^{\min})^p$ and if w is CM, then

$$\mu \geq c \rho_p^{\min} \sqrt{p+1} (1+\beta)^{\frac{1}{p}-1} M^{1/p}, \quad (8)$$

with $c = (8\sqrt{2})/(9\sqrt{e})$. Therefore, using this relation, simple calculations show that, with a probability higher

than $1 - e^{-2\kappa^2}$,

$$\mu^{-1}\epsilon_p(M) \leq C (\rho_p^{\min} \sqrt{p+1})^{-1}. \quad (9)$$

with $C < 2.17$ as soon as $M > (\frac{p+1}{p}\kappa)^2$. Interestingly, this last inequality reduces to the bound of ϵ/μ for the uniform quantization of bin width $\alpha > 0$ obtained in [5]. Indeed, this case is equivalent to $w = \frac{2}{\alpha}\mathbf{1}$, so that $\rho_p^{\text{unif}} = \rho_p^{\min} = \rho_p^{\max} = \frac{2}{\alpha}$. For non-uniform quantization, in addition to the division of the error by $\sqrt{p+1}$ due to the oversampled sensing, a new effect occurs: an error reduction due to ρ_p^{\min} . As estimated in Fig. 1-right with $\rho_p^{\min} \simeq M^{-1/p}\|w\|_p$, this value can be higher for non-uniform quantization than for uniform one.

6. NUMERICAL IMPLEMENTATION

We have implemented the convex WBPDQ decoder thanks to the BPDQ toolbox³. Briefly, this toolbox solves the unweighted problem ($s^\pm = 1$) where the optimization constraint reduces to a simple ℓ_p norm. It proceeds by using the Douglas Rachford splitting [9]. This iterative methods combines the proximal operator of the ℓ_1 norm (a soft-thresholding) and the orthogonal projector onto a ℓ_p ball of radius ϵ centered at the measurement vector y . We have adapted this toolbox in order to solve WBPDQ by modifying this last operation with a projection on the ball associated to the asymmetric norm in (3). As described in [5], we used for that a Newton's method solving the related KKT system which is reminiscent of the Sequential Quadratic Programming (SQP) method.

7. EXPERIMENTS

Given the quantized sensing model (1), the purpose of this section is to observe how the WBPDQ reconstruction quality evolves for different set of parameters N , K , M , R and moments p . In all our experiments, a *spread-spectrum* sensing matrix has been used [10]. More precisely, $\Phi = RF \text{diag}(h)$, where $R \in \mathbb{R}^{M \times N}$ is a restriction matrix picking M values uniformly at random in $\{1, \dots, N\}$, $F \in \mathbb{R}^{N \times N}$ is the DCT basis, and $h \in \mathbb{R}^N$ is a random ± 1 Bernoulli sequence. This sensing matrix is not proved to be as optimal as the SGR matrix (or Random Gaussian Ensemble). However, we observed that the measurement vector Φx still follows a Gaussian distribution, and the multiplication Φu or $\Phi^T v$ is very fast ($O(N \log_2 N)$). Moreover, it is a tight frame with $\Phi \Phi^T = \mathbf{1}$ which makes the Douglas Rachford (DR) faster [5]. For all experiments, 300 iterations were sufficient to observe a convergence of the DR algorithm.

Sparse 1-D signals: For this experiment, we have randomly generated K -sparse signals in \mathbb{R}^N by picking uniformly at random their support in $\{1, \dots, N\}$, their K non-zero values following a Normal distribution. Each signal has been

³<http://wiki.epfl.ch/bpdq>




Fig. 2: Left: King's House, Brussels; Middle: BPDN (25.65dB); Right: WBDPQ ($p = 6$, 26.44dB).




Fig. 3: SNR Gain. $N = 1024$, $K = 16$, $R = 4$, averaged over 100 trials.

sensed with the sensing model (1) according to two quantization scenario. An optimal uniform quantizer and a non-uniform Lloyd-Max quantizer both tuned to inputs with Gaussian distributions. For the first quantizer, the Basis Pursuit Denoiser (BPDN) has reconstructed the signals, while for the second, the WBDPQ decoder has been selected with weights computed as in Sec. 4. The parameter N , K and R have been set to 1024, 16 and 4 bits respectively, while $M/K \in [8, 64]$ and $p \in [2, 10]$. Fig. 3-left shows the averaged reconstruction quality evolution (in SNR) over 100 trials with respect to M/K for different p . For comparison, the BPDN quality curve is shown with a dashed line. It can be clearly observed that as soon as the oversampling factor M/K is sufficiently high, taking a p higher than 2 yields significant improvement with up to 4 dB at the highest oversampling for $p = 10$. However, Fig. 3-right, which displays the averaged SNR gain relatively to WBDPQ with $p = 2$, confirms that the moment p leading to the highest gain for a given ratio oversampling ratio increases with M/K , while the performance decays if higher moments are selected.

Compressible 2-D image: In this second experiment, we have challenged the WBDPQ decoder on the reconstruction of a 512×512 image ($N = 262144$) compressible in a Daubechies 7/9 wavelet basis and sensed with (1) for $M = \frac{3}{4}N$ and 4 bits per measurements. The reconstruction results are shown in Fig. 2. Despite a weak visual improvement between the WBDPQ ($p = 6$, on Lloyd-Max quantized measurements) and the BPDN reconstructions (on uniformly quantized measurements), the PSNR gain for WBDPQ is about 0.8 dB. This is not a large value which is mainly due to the compressible nature of the image; the reduction

of the quantization error is masked by the importance of the compressible error in (6). However, we measured that the WBDPQ solution satisfies the quantization consistency, while for the BPDN solution, $\|\Phi x^* - y\|_{\infty, (2/\alpha)\mathbf{1}} = 1.963$.

8. CONCLUSION

The objective of this work was to show how to improve the reconstruction of sparse or compressible signals sensed through non-uniformly quantized compressive measurements. We have shown that a weighted asymmetric ℓ_p norm allows us to incorporate the quantization consistency of the measurements in the signal decoding, that is, in the WBDPQ program. Finally, theoretical indications and numerical experiments confirm the possibility to reduce the quantization impact in the reconstructed signals.

9. REFERENCES

- [1] D. Donoho, "Compressed Sensing," *IEEE Tr. Inf. Th.*, **52**(4), pp. 1289–1306, 2006.
- [2] E. Candès and J. Romberg, "Quantitative Robust Uncertainty Principles and Optimally Sparse Decompositions," *Found. Comp. Math.*, **6**(2), pp. 227–254, 2006.
- [3] W. Dai, H. V. Pham, and O. Milenkovic, "Distortion-rate functions for quantized compressive sensing," *IEEE Inf. Th. Work. (ITW)*, pp. 171–175, 2009.
- [4] J. Laska, P. Boufounos, and R. Baraniuk, "Finite-range scalar quantization for compressive sensing," in *Conf. Samp. Th. App. (SampTA)*, 2009.
- [5] L. Jacques, D. Hammond, and M. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Tr. Inf. Th.*, **57**(1), pp. 559–571, 2011.
- [6] R. Gray and D. Neuhoff, "Quantization," *IEEE Tr. Inf. Th.*, **44**(6), pp. 2325–2383, 1998.
- [7] L. García Raffi and S. Pérez, "Asymmetric norms and optimal distance points in linear spaces," *Topol. & Appl.*, **155**(13), pp. 1410–1419, 2008.
- [8] E. Candès, "The restricted isometry property and its implications for compressed sensing," *Compte Rendus Acad. Sciences, Paris, Serie I*, **346**, pp. 589–592, 2008.
- [9] P. Combettes and J. Pesquet, "A Douglas–Rachford Splitting Approach to Nonsmooth Convex Variational Signal Recovery," *IEEE J. Sel. Top. Sig. Proc.*, **1**(4), pp. 564–574, 2007.
- [10] Y. Wiaux, G. Puy, Y. Boursier, and P. Vandergheynst, "Spread spectrum for imaging techniques in radio interferometry," *Mon. Not. R. Astron. Soc.*, **400**(2), pp. 1029–1038, 2009.