q-ary Compressive Sensing

Youssef Mroueh
LCSL/CBCL - CSAIL, MIT, IIT
Email: ymroueh@mit.edu

Lorenzo Rosasco
DIBRIS, Unige and LCSL - MIT, IIT
Email: lrosasco@mit.edu

***Abstract*—We introduce q-ary compressive sensing, an extension of 1-bit compressive sensing. We propose a novel sensing mechanism and a corresponding recovery procedure. The recovery properties of the proposed approach are analyzed both theoretically and empirically. Results in 1-bit compressive sensing are recovered as a special case. Our theoretical results suggest a tradeoff between the quantization parameter q and the number of measurements m, in controlling the error and robustness to noise of the resulting recovery algorithm.**

I. Introduction

Reconstructing signals from discrete measurements is a classic problem in signal processing. Properties of the signal allow the reconstruction from a minimal set of measurements. The classical Shannon sampling result ensures that band limited signals can be reconstructed by a linear procedure, as long as a number of linear measurements, at least twice the maximum frequency, is available. Modern data analysis typically requires recovering high dimensional signals from few inaccurate measurements. Indeed, the development of Compressed Sensing (CS) and Sparse Approximation [1] shows that this is possible for signals with further structure. For example, d-dimensional, s-sparse signals[1] can be reconstructed with high probability through convex programming, given $m \sim s \log(d/s)$ random linear measurements.

Non-linear measurements have been recently considered in the context of 1-bit compressive sensing [2], [3], [4], [5], [6] (`http://dsp.rice.edu/1bitCS/`). Here, binary (one-bit) measurements are obtained by applying, for example, the "sign" function[2] to linear measurements. More precisely, given $x \in \mathbb{R}^d$, a measurement vector is given by $y = (y_1, \dots, y_m)$, where $y_i = sign(\langle w_i, x\rangle)$ with $w_i \sim \mathcal{N}(0, I_d)$ independent Gaussian random vectors, for $i = 1, \dots, m$. It is possible to prove [4] that, for a signal $x \in K \cap \mathbb{B}^d$ ($\mathbb{B}^d$ is the unit ball in $\mathbb{R}^d$), the solution $\hat{x}_m$ to the problem

$$\max_{x \in K} \sum_{i=1}^{m} y_i \langle w_i, x\rangle \,, \qquad (1)$$

satisfies $\|\hat{x}_m - x\|^2 \leq \frac{\delta}{\sqrt{\frac{2}{\pi}}}$, with probability $1 - 8\exp(-c\delta^2 m)$, $\delta > 0$, as long as $m \geq C\delta^{-2}\omega(K)^2$ [4]. Here, C denotes a universal constant and $\omega(K) = \mathbb{E}\sup_{x\in K-K} \langle w, x\rangle$ the Gaussian mean width of K, which can be interpreted as a complexity measure. If K is a convex set, problem (1) can be solved efficiently.

In this paper, borrowing ideas from signal classification and machine learning, we discuss a novel sensing strategy, based on *q-ary non-linear measurements*, and a corresponding recovery procedure.

II. q-ary Compressive Sensing

A. Sensing and Recovery

The sensing procedure we consider is given by a map C from $K \cap \mathbb{B}^d$ to $\mathcal{F} = \{0, \dots, q-1\}^m$, where $K \subset \mathbb{R}^d$. To define C we need the following definitions.

Definition 1 (Simplex Coding [7])**.** *The simplex coding map is $S : \{0, \dots, q-1\} \to \mathbb{R}^{q-1}$, $S(j) = s_j$, where*
1) $\|s_j\|^2 = 1$,
2)$\langle s_j, s_i\rangle = -\frac{1}{q-1}$, for $i \neq j$,
3) $\sum_{j=0}^{q-1} s_j = 0$.

Definition 2 (q-ary Quantized Measurements)**.** *Let $W \in \mathbb{R}^{q-1,d}$ be a Gaussian random matrix, i.e. $W_{ij} \sim \mathcal{N}(0, 1)$ for all i, j. Then, $Q : K \cap \mathbb{B}^d \to \{0, \dots, q-1\}$,*

$$Q(x) = Q_W(x) = \arg\max_{j=0\dots q-1} \langle s_j, Wx\rangle \,,$$

is called a q-ary quantized measurement.

Then, we can define the q-ary sensing strategy induced by non-linear quantized measurements.

Definition 3 (q-ary Sensing)**.** *Let $W_1, \dots, W_m$, be independent Gaussian random matrices in*

[1] A d-dimensional signal, that is a vector in $\mathbb{R}^d$, is s-sparse if only s of its components are different from zero.

[2] More generally, any function $\theta : \mathbb{R} \to [-1, 1]$, such that $\mathbb{E}(g\theta(g) > 0)$ can be used.

$\mathbb{R}^{q-1,d}$ *and* $Q_{W_i}(x), i = 1, \dots, m$ *as in Def. 2. The q-ary sensing is* $C : K \cap \mathbb{B}^d \to \{0, \dots, q-1\}^m$,

$$C(x) = (Q_{W_1}(x), \dots Q_{W_m}(x)),$$

$\forall x \in K \cap \mathbb{B}^d$.

Before describing the recovery strategy we consider, we discuss the connection to 1-bit CS and binary embeddings [8] [6] .

Remark 1 (Connection to 1-bit CS)**.** *If* $q = 2$, *W reduces to a Gaussian random vector, and* $2Q(x) - 1 = sign(Wx)$, *so that the q-ary quantized measurements become equivalent to those considered in* 1*-bit CS.*

Remark 2 (Sensing and Embeddings)**.** *It can be shown that C defines an ϵ-isometric embedding of* $(K, \|\cdot\|)$ *into* $(\mathcal{F}, d_H)$ *– up-to a bias term. Here d_H is the (normalized) Hamming distance,* $d_H(u, v) = \frac{1}{m}\sum_{i=1}^m \mathbb{1}_{u_i \neq v_i}, u, v \in \mathcal{F}$. *This analysis is deferred to the long version of this paper.*

In this paper, we are interested in provably (and efficiently) recovering a signal x from its q-ary measurements $y = (y_1, \dots, y_m) = C(x)$. Following [4], we consider the recovery strategy $D : \{0, \dots, q-1\}^m \to K \cap \mathbb{B}^d$ defined by,

$$D(y) = \arg\max_{u \in K \cap \mathbb{B}^d} \frac{1}{m}\sum_{i=1}^m \langle s_{y_i}, W_i u\rangle . \quad (2)$$

The above problem is convex if K is convex and can be solved efficiently, see Section III-A. In the next section, we prove that the solution to Problem (2) has good recovery guarantees both in noiseless and noisy settings.

Remark 3 (Connection to Classification)**.** *The inspiration for considering q-ary CS stems from an analogy between* 1*-bit compressed sensing and binary classification in machine learning. In this view, Definition* (3) *is related to the approach proposed for multi category classification in [7]. Following these ideas, we can extend the recovery strategy* (2) *by considering*

$$D_V(y) = \arg\min_{u \in K \cap \mathbb{B}^d} \frac{1}{m}\sum_{i=1}^m V(-\langle s_{y_i}, W_i u\rangle), \quad (3)$$

where V is a convex, Lipschitz, non-decreasing loss function $V : \mathbb{R} \to \mathbb{R}^+$. *Problem* (2) *corresponds to the choice* $V(x) = x$. *Other possible choices include* $V(x) = \max(1+x, 0)$, $V(x) = \log(1 + e^x)$, *and* $V(x) = e^x$. *Each of these loss functions can be seen as convex relaxations of the* 0-1 *loss function, defined as* $V(x) = 0$ *if* $x \leq 0$, *and* 1 *otherwise. The* $0 - 1$ *loss defines the misclassification risk, which corresponds to Hamming distance in CS, and is a natural measure of performance while learning classification rules.*

Remark 4 (Recovery of Distorted Signals)**.** *We note that the q-ary approach could be of particular interest in situations where the signals can undergo unknown non-linear distortions, because of the robustness of the maximum in the definition of the q-ary measurements.*

B. Recovery guarantees: Noiseless Case

The following theorem describes the recovery guarantees for the proposed procedure for signals in a set K of Gaussian mean width $w(K)$. We first consider a noiseless scenario.

Theorem 1. *Let* $\delta > 0$, $m \geq C\delta^{-2}w(K)^2$. *Then with probability at least* $1 - 8\exp(-c\delta^2 m)$, *the solution* $\hat{x}_m = D(y)$ *of problem* (2) *satisfies,*

$$\|\hat{x}_m - x\|^2 \leq \frac{\delta}{\sqrt{\log(q)}}. \quad (4)$$

A proof sketch of the above result is given in Section II-D, while the complete proof is deferred to the long version of the paper. Here, we add four comments. First, it can be shown the the above result bound is derived from an error bound,

$$||\hat{x}_m - x||^2 \leq C(\frac{w(K)}{\sqrt{\log(q)m}} + t), \quad (5)$$

with probability at least, $1 - 4\exp(-2t^2), t > 0$.

Second, Inequalities (4), (5) can be compared to results in 1-bit CS. For the same number of measurements, $m \geq C\delta^{-2}w(K)^2$, the error for q-ary CS is $\frac{\delta}{\sqrt{\log(q)}}$, in contrast with $\frac{\delta}{\sqrt{\frac{2}{\pi}}}$ in the 1-bit CS [4], at the expense of a more demanding sensing procedure. Also note that, for $q = 2$, we recover the result in 1-bit CS as a special case. Third, we see that for a given accuracy our results highlight a trade-off between the number of q-ary measurements m and the quantization parameter q. To achieve an error ϵ with a memory budget of ℓ bits, one can choose m and q so that $\epsilon = O(\frac{1}{\sqrt{m\log(q)}})$, and $m\log_2(q) = \ell$ (see also section III-B). Finally, in the following we will be interested in K being the set of s-sparse signals. Following again [4], it is interesting to consider in Problem (2) the relaxation

$$K_1 = \{u \in \mathbb{R}^d : \|u\|_1 \leq \sqrt{s}, \|u\|_2 \leq 1\}.$$

With this choice, it it possible to prove that $w(K_1) \leq C\sqrt{s\log(\frac{2d}{s})}$, and that for $m \geq C\delta^{-2}s\log(\frac{2d}{s})$, the solution of the convex program (2) on K_1 satisfies, $\|\hat{x}_m - x\|^2 \leq \frac{\delta}{\sqrt{\log(q)}}$. We end noting that other choices of K are possible, for example in [9] the set of group sparse signals (and its Gaussian width) is studied.

C. Recovery Guarantees: Noisy Case

Next we discuss the q-ary approach in two noisy settings, related to those considered in [4].
Noise before quantization. For $i = 1, \ldots, m$, let

$$y_i = \arg \max_{j=0\ldots q-1} \{\langle s_j, W_i x\rangle + g_j\}, \quad (6)$$

with g_j independent Gaussian random variables of variance σ^2. In this case, it is possible to prove that, for $m \geq C\delta^{-2}w(K)^2$,

$$\|\hat{x}_m - x\|^2 \leq \frac{\delta\sqrt{1+\sigma^2}}{\sqrt{\log(q)}},$$

with probability at least $1-8\exp(-c\delta^2 m)$. The quantization level q can be chosen to adjust to the noise level σ for a more robust recovery of x. This result can be viewed in the perspective of the *bit-depth versus measurement-rates* perspective studied in [10], where it is shown that 1-bit CS outperforms conventional scalar quantization. In this view, $q-$ary CS provides a new way to adjust the quantization parameter to the noise level.
Inexact maximum. For $i = 1, \ldots, m$, let $y_i = Q_{W_i}(x)$, with probability p, and $y_i = r$ with probability $1-p$, with r drawn uniformly at random from $\{0, \ldots, q-1\}$. In this case, it is possible to prove that, for $m \geq C\delta^{-2}w(K)^2$,

$$\|\hat{x}_m - x\|^2 \leq \frac{\delta}{\sqrt{\log(q)}(2p-1)}.$$

with probability at least $1-8\exp(-c\delta^2 m)$. The signal x can be recovered even if *nearly half* of the q-ary bits are flipped.

D. Elements of the proofs

We sketch the main steps in proving our results. The proof of Theorem 1 is based on: 1) deriving a bound in expectation, and 2) deriving a concentration result. The proof of the last step uses Gaussian concentration inequality extending the proof strategy in [4]. Step 1) gives the bound

$$\mathbb{E}\left(||\hat{x}_m - x||^2\right) \leq \frac{w(K)}{C\sqrt{\log(q)m}},$$

the proof of which is based on the following proposition.

Proposition 1. *Let $\mathcal{E}_x(u) = \mathbb{E}_W(\langle s_\gamma, Wu\rangle)$, where $\gamma = Q_W(x)$. Then, $\forall u \in \mathbb{B}^d$, we have,*

$$\frac{1}{2}\|u - x\|^2 \leq \frac{1}{\lambda(q)}\left(\mathcal{E}_x(x) - \mathcal{E}_x(u)\right),$$

where $\lambda(q) = \mathbb{E}_{\bar{\gamma},g}(\langle s_{\bar{\gamma}}, g\rangle)$, and $g \sim \mathcal{N}(0, I_{q-1})$, and $\bar{\gamma} = \arg\max_{j=0\ldots q-1}\langle s_j, g\rangle$.

Using results in empirical process theory it possible to show that

$$|\mathcal{E}_x(x) - \mathcal{E}_x(\hat{x}_m)| \leq C\frac{w(K)}{\sqrt{m}}.$$

The bound on the expected recovery follows combining the above inequality and Proposition 1 with the inequality,

$$\lambda(q) \geq C\sqrt{\log(q)},$$

which is proved using Slepian's inequality and Sudakov minoration.

Results in the noisy settings follow from suitable estimates of $\lambda(q)$. Indeed, for the *noise before quantization* case it can be proved that $\lambda(q) \geq C\sqrt{\frac{\log(q)}{1+\sigma^2}}$. For the *inexact maximum* case one has

$$\begin{aligned}\lambda(q) \quad &= \mathbb{E}_{y,g}(\langle s_y, g\rangle) = \\ &p\mathbb{E}(\max_{j=1\ldots q}\langle s_j, g\rangle) + (1-p)\mathbb{E}(\langle s_r, g\rangle) \geq \\ &Cp\sqrt{\log(q)} + (1-p)\mathbb{E}(\min_{j=1\ldots q}\langle s_j, g\rangle) \geq \\ &(2p-1)C\sqrt{\log(q)}.\end{aligned}$$

III. Experimental Validation

A. An Algorithm for Sparse recovery

In our experiments, we considered the following variation of problem (2), Let $\xi_i = s_{y_i}^\top W_i \in \mathbb{R}^d, i = 1\ldots m$.

$$\max_{u, \|u\|_2 \leq 1} \frac{1}{m}\sum_{i=1}^m \langle \xi_i, u\rangle - \eta\|u\|_1, \quad (7)$$

where $\eta > 0$. The above problem can be solved efficiently using Proximal Methods [11]. Indeed, a solution can be computed via the iteration,

$$\begin{aligned} u_{t+1} &= u_t + \frac{\nu_t}{m}\sum_{i=1}^m \xi_i, \\ u_{t+1} &= Prox_\eta(u_{t+1}), \\ u_{t+1} &= u_{t+1}\min(\frac{1}{\|u_{t+1}\|_2}, 1). \end{aligned}$$

Where ν_t is the gradient step size, and $Prox_\eta$ acts component-wise as $\max(1-\frac{\eta}{|u_i|}, 0)u_i$. The iteration is initialized randomly to a unit vector.

Remark 5. *The computational complexity of the sensing process depends on both m and q, while, once computed ξ_i, that of the recovery algorithm depends only on m, and is the same as in 1-bit CS. In this sense, given a bit rate, the same precision can achieved by 1-bit CS and q-ary CS, with a better computational complexity for the decoding in the q-ary case.*

B. Sparse Recovery

We tested our approach for recovering a signal from from its q-ary measurements. We considered sparse signals of dimension d generated via a Gauss-Bernoulli model. In Figure 1.(a), we see that the reconstruction error $\hat{x}_m$ (in blue), for varying m and q fixed, follows the theoretical bound $\frac{1}{\sqrt{m}}$ (in red). In Figure 1.(b), we see that the reconstruction error of $\hat{x}_m$ (in blue), for varying q and m fixed, follows the theoretical bound $\frac{1}{\sqrt{\log(q)}}$ (in red). Figures 1.(c), and 1.(d) highlight the tradeoff between the number of measurements and the quantization parameter. For a precision ϵ, and a memory budget 2^B, one can choose an operating point (m, q), according to the theoretical bound $\frac{1}{\sqrt{m\log(q)}}$.

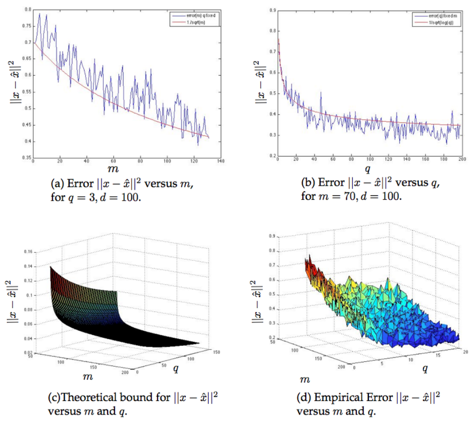

Fig. 1. q-ary Compressive Sensing: Quantization/Number of measurements tradeoff.

C. Image Reconstruction

We considered the problem of recovering an image from q-ary measurements. We used the $8-$bit grayscale boat image of size 64×64 pixels shown in Figure 2(a). We extracted and thresholded the wavelet coefficients to get a sparse signal. We normalized the resulting vector of wavelets coefficients of dimension $d = 3840$ to obtain a unit vector. Then, we performed sensing and recovery with $q = 2^5$ (5-bit compressive sensing) and $q = 2$ (1-bit compressive sensing) for the same $m = 2048 < d$. We compared the SNR of the corresponding reconstructed images in noiseless (Figures 2(b)-(c)), and noisy settings (noise before quantization model (6), with $\sigma = 0.8$), Figures 2(d)-(e). Note that in this setting we are comparing 1-bit CS and q-ary CS, for the same decoding time (same m). The results confirm our theoretical findings: higher quantization improves the SNR, as well as robustness to noise of q-ary CS.

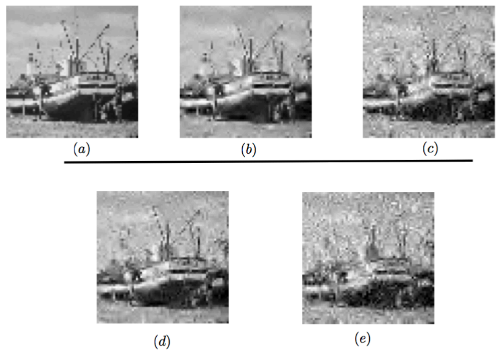

Fig. 2. Image recovery with q-ary CS. (a) Original image. (b) Reconstruction with no-noise: $q = 2^5$, SNR $= 20.2$ dB. (c) Reconstruction with no-noise: $q = 2$, SNR $= 16.2$ dB. (d) Reconstruction with noise: $q = 2^5$, SNR $= 18.3$ dB. (e) Reconstruction with noise: $q = 2$, SNR $= 15$ dB.

REFERENCES

[1] Y.C. Eldar, and G. Kutyniok Eds, Compressed Sensing. Cambridge University Press, 2012.
[2] P. T. Boufounos and R. G. Baraniuk. 1-bit Compressive Sensing, in Proceedings of Conference on Information Science and Systems (CISS), Princeton, NJ, March 2008.
[3] J. Z. Sun and V. K. Goyal, Optimal Quantization of Random Measurements in Compressed Sensing, Proc. IEEE Int. Symp. Information Theory (Seoul, Korea), June-July 2009.
[4] Y. Plan and R. Vershynin, Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach, IEEE Transactions on Information Theory, to appear.
[5] Plan, Y. and Vershynin, R. (2013), One-bit compressed sensing by linear programming. Comm. Pure Appl. Math.
[6] 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. Info. Theory, v. 59, no. 4, April, 2013.
[7] Y. Mroueh, T. Poggio, L. Rosasco, and J.J Slotine, Multiclass learning with simplex coding, NIPS 2012.
[8] Y. Plan and R. Vershynin, Dimension reduction by random hyperplane tessellations, submitted 2011.
[9] N. Rao, R. Nowak and B. Recht, Tight Measurement Bounds for Exact Recovery of Structured Sparse Signals, Arxiv preprint,2011.
[10] J. Laska, and R G. Baraniuk, Regime Change: Bit-Depth versus Measurement-Rate in Compressive Sensing, Arxiv preprint, available online at http://arxiv.org/abs/1110.3450, 2011.
[11] S. Boyd and L. Vandenberghe. Convex Optimization. Cambridge University Press, March 2004. ISBN 0521833787.