

Sparse linear regression with compressed and low-precision data via concave quadratic programming

Original

Sparse linear regression with compressed and low-precision data via concave quadratic programming / Cerone, V.; Fosson, S.; Regruto, D.. - ELETTRONICO. - (2019), pp. 6971-6976. (Intervento presentato al convegno IEEE Conference on Decision and Control (CDC) tenutosi a Nice (France) nel 11-13 Dicembre 2019) [10.1109/CDC40024.2019.9030257].

Availability:

This version is available at: 11583/2785501 since: 2021-01-28T13:45:29Z

Publisher:

IEEE

Published

DOI:10.1109/CDC40024.2019.9030257

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Sparse linear regression with compressed and low-precision data via concave quadratic programming

V. Cerone, S. M. Fosson*, D. Regruto

September 10, 2019

Abstract

We consider the problem of the recovery of a k -sparse vector from compressed linear measurements when data are corrupted by a quantization noise. When the number of measurements is not sufficiently large, different k -sparse solutions may be present in the feasible set, and the classical ℓ_1 approach may be unsuccessful. For this motivation, we propose a non-convex quadratic programming method, which exploits prior information on the magnitude of the non-zero parameters. This results in a more efficient support recovery. We provide sufficient conditions for successful recovery and numerical simulations to illustrate the practical feasibility of the proposed method.

1 Introduction

Sparse linear regression is the recovery of a sparse vector $x \in \mathbb{R}^n$ from linear measurements $y = Ax \in \mathbb{R}^m$, with $A \in \mathbb{R}^{m,n}$. A vector is sparse if it has few non-zero components; more precisely, we call k -sparse a vector with $k \ll n$ non-zero components. The interest for sparse solutions has different motivations. In machine learning and system identification, a purpose is to build models as simple as possible from large datasets. Indeed, we know that in many cases the true number of parameters of a system is much smaller than the global dimensionality of the problem, and sparsity supports the removal of redundant parameters. In the recent literature, the identification of linear systems under sparsity constraints is considered in [21, 25, 24, 9]. Furthermore, in the last decade, sparsity has been attracting a lot of attention due to the theory of compressed sensing (CS, [6, 12]), which states that a sparse vector can be recovered from compressed linear measurements, that is, when $m < n$, under suitable conditions on the structure of A .

In real applications, observations may be affected by different sources of noise. The most common one is measurement noise, which is usually modeled as an unknown additive term $\delta_y \in \mathbb{R}^m$, i.e., $y = Ax + \delta_y$. Linear regression in the presence of measurement noise is usually solved via least squares; if x is desired to be sparse, an ℓ_1 regularizer is added, which promotes sparsity and keeps the problem convex. More precisely, the problem can be formulated as $\min_{x \in \mathbb{R}^n} \|x\|_1$ s. t. $\|Ax - y\|_p \leq \epsilon$, where $\|\cdot\|_p$ is a suitable norm and $\epsilon > 0$ is

a known bound; this is referred to as Basis Pursuit Denoising (BPDN_p). In the literature, BPDN₂ is very popular, in particular in signal processing and CS, see, e.g., [12]. The choice $p = 2$ bounds the mean energy of the error, hence BPDN₂ is tolerant to possible outliers. The case $p = \infty$, instead, is considered to deal with quantized or low-precision data in [7, 14, 26, 4]. When y is quantized, a bound on each component is given, and no outliers occur. This makes the ℓ_∞ description more suitable than the ℓ_2 one.

In many applications, both the measurements and the matrix A are affected by noise. In particular, in linear systems identification, this is analyzed by the errors-in-variables (EIV) model, where both the input and the output are corrupted by noise; we refer the reader to [22] for a complete overview on EIV models. In [3, 5], the identification of the feasible set for EIV models is tackled. Perturbations on A are more tricky to deal with than measurement noise, as they make the problem non-convex. Furthermore, in CS, the recovery of sparse vectors from compressed measurements is strongly compromised by perturbations on A , which are tolerated only if sufficiently small or with particular structures. In [13], robustness to perturbation in CS is investigated, providing information on the amount of perturbation on A that can be tolerated by BPDN₂. In [27], the perturbation is assumed to be n -dimensional, instead of nm -dimensional, and bounded in the ℓ_∞ norm. Alternating minimization is used to find a minimum, which is global when the components of x are known to be non-negative, see [27, Section IV.A].

In [4], the same non-negativity assumption is leveraged to formulate the sparse linear regression problem, with bounded noise both on y and A , as a linear programming problem. Specifically, the proposed problem in [4] is the minimization of the ℓ_1 norm within the feasible set, which, under the non-negativity assumption, is a convex polytope. The approach is proved to be robust, in the sense that the distance between the solution and the true vector is proportional to the noise and to the sparsity level of the true signal.

The ℓ_1 approach may fail if the feasible set does not delineate the sparsity pattern. More precisely, many k -sparse solutions might belong to the feasible set, though with different support. When several k -sparse solutions are feasible, the ℓ_1 minimizer selects the one with less energy, which may not correspond to the desired one. We highlight that, particularly in the case of compressed measurements, the support recovery is the key point for the vector recovery. If the support is known, and if we assume sparsity level $k \leq m$, the problem is not underdetermined, and the solution can be found by inversion, see, e.g., [11]. We also remark that, in CS, conditions on A are provided that guarantee a unique k -sparse solution for BPDN₂, see, e.g., [12]; nevertheless, these conditions are actually limited to specific classes of random matrices, and perturbations on A are not envisaged.

In this work, we tackle the above described problem by proposing a novel solution. Specifically, we propose a non-convex, polynomial approach which is oriented to improve the recovery performance, particularly in terms of sparsity pattern, by leveraging a prior information on the magnitude of the non-zero components. Sufficient conditions that guarantee the recovery of the correct support are proven, and numerical simulations illustrate the effectiveness of the approach in different settings.

The paper is organized as follows. In Section 2, we present the background and a motivating example, which highlights the flaws of the ℓ_1 approach. In

Section 3, we formally state the proposed problem, and discuss the considered assumptions. In Section 4, we explain how to solve the proposed problem. In Section 5, we prove sufficient conditions that guarantee successful recovery, while Section 6 is devoted to numerical experiments. Finally, some conclusions are drawn in Section 7.

2 Background and motivating example

Let us consider an over-dimensional vector of parameters $\tilde{x} \in \mathbb{R}^n$, i.e., only $k \ll n$ components are different from zero, and let us assume that the system output is given by $y = A\tilde{x}$, $A \in \mathbb{R}^{m,n}$. The goal is to recover \tilde{x} under the following assumptions: (a) $m < n$, i.e., compressed output observations are available; (b) both y and A are not exactly known. More precisely, low-precision (or quantized) versions of A and y are known, which are denoted by $Q(A) \in \mathbb{R}^{m,n}$ and $Q(y) \in \mathbb{R}^m$, respectively. The corresponding perturbations are denoted by $\delta_A = Q(A) - A$ and $\delta_y = Q(y) - y$. Moreover, we assume to know the maximum perturbation magnitudes, denoted by Δ_A and Δ_y , respectively. For example, if we consider data quantized over a fixed number of bits, we assume to know the number of bits.

Let us represent by \mathcal{D} the feasible set. The natural approach to estimate \tilde{x} , which is sparse, is by looking for the sparsest solution in \mathcal{D} . This can be performed by minimizing the ℓ_1 -norm over $\mathcal{D} \subset \mathbb{R}^n$, as illustrated in [4]. Specifically, the problem is formulated as follows: given $Q(y) \in \mathbb{R}^m$, $Q(A) \in \mathbb{R}^{m,n}$, $\Delta_A > 0$, and $\Delta_y > 0$,

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{s. t. } y &= Ax \\ Q(y) &= y + \delta_y \\ Q(A) &= A + \delta_A \\ \|\delta_y\|_\infty &\leq \Delta_y \\ \|\delta_A\|_\infty &\leq \Delta_A. \end{aligned} \tag{1}$$

In [4], the assumption $x_i \geq 0$, for any $i = 1, \dots, n$, is considered, which makes the problem solvable via linear programming. Actually, the same result is obtained if each component is known to be either non-positive or non-negative: non-negativity is assumed for simplicity without loss of generality. However, we notice that the non-negative setting has itself a number of applications, that range from sparse localization problems, see [1], to sparse sensor selection, see [2].

In [4], Problem (1) is shown to provide encouraging results if compared to state-of-the-art algorithms for the recovery of sparse signals from compressed and low-precision data. However, a drawback of the ℓ_1 approach is that it may fail the recovery of the right solution when several k -sparse vectors are present in \mathcal{D} , with different supports. We illustrate this occurrence with a simple motivating example. Let us consider $A = [0.2131, 1.2414]$, $\tilde{x} = [1, 0]^T$. Then we have $y = A\tilde{x} = 0.2131$. Data are subject to quantization with step 0.1, thus we observe $Q(A) = [0.2, 1.2]$ and $Q(y) = 0.2$. If we assume non-negativity, the feasible set \mathcal{D} , which is shown in Fig. 1, is the intersection between the first orthant and the area between the two lines that represent the




Figure 1: $A = [0.2131, 1.2414]$, $\tilde{x} = [1, 0]^T$ (blue point); $\mathcal{Q}(A) = [0.2, 1.2]$, $\mathcal{Q}(y) = 0.2$. The feasible set \mathcal{D} is the light blue polytope. The solution of Problem (1) is $[0, 0.0769]^T$ (red point). The support is inverted with respect to \tilde{x} .

constraints. The solution of problem (1) is $[0, 0.0769]^T$, depicted in Fig. 1 with a red point. This vector is as sparse as \tilde{x} , while it has the opposite support. Moreover, this solution has smaller non-zero value than the true one, which is not a natural choice in practice. In fact, in real applications, the null entries of \tilde{x} are not exactly zero, instead they take small values; from this point of view, the support discriminates between significant values and non-significant ones. In our example, 0.0769 is rather close to zero if compared to 1, therefore a suitable algorithm should estimate it as a zero, and search for a more significant non-zero value to put in the support.

To tackle this occurrence, intuitively, one could assume that an interval $[\alpha, \beta]$, $\alpha > 0$, $\beta \geq \alpha$, of significant magnitudes is known, thus the constraint $x_i \in \{0\} \cup [\alpha, \beta]$, $i = 1, \dots, n$, might be added to Problem (1). Nevertheless, this would compromise the convexity of \mathcal{D} and make the problem definitely more complex. For this motivation, in this work we propose a different solution.

3 Problem Statement

In this section, we present a novel formulation of the sparse linear regression problem with compressed measurements and low-precision data, under the assumption that the significant values are in $[\alpha, \beta]$. We start by studying the case of non-negative parameters, while we leave for future work more general sign settings.

Assumption 1 For any $i = 1, \dots, n$, $x_i \in \{0\} \cup [\alpha, \beta]$.

Then, we define $d := \frac{\alpha + \beta}{2}$, and we formulate the problem as follows.

Problem 1

Given $\mathcal{Q}(y) \in \mathbb{R}^m$, $\mathcal{Q}(A) \in \mathbb{R}^{m,n}$, $\Delta_A > 0$, and $\Delta_y > 0$,

$$\begin{aligned} \min_{x \in [0,d]^n} d\|x\|_1 - \|x\|_2^2 \quad \text{s. t. } y = Ax \\ \mathcal{Q}(y) = y + \delta_y \\ \mathcal{Q}(A) = A + \delta_A \\ \|\delta_y\|_\infty \leq \Delta_y \\ \|\delta_A\|_\infty \leq \Delta_A. \end{aligned} \tag{2}$$

Intuitively, Problem 1 is more accurate than (1) since the concave penalty $d\|x\|_1 - \|x\|_2^2$ is closer to $\|x\|_0$ than $\|x\|_1$, see, e.g., [8] for more details. The rest of the paper is devoted to prove this enhancement in the proposed setting.

First of all, we observe that if $x \in [0, d]^n$, then the minimum of the objective functional is zero, which is achieved for any $x \in \{0, d\}^n$. Since the true non-zero parameters are in $[\alpha, \beta]$, this approach does not exactly evaluate the non-zero parameters. However, our main goal is to obtain the right support. As mentioned before, once the right support is obtained, if $k \leq m$, the problem of evaluating the non-zero parameters is not underdetermined and can be obtained by inversion. Nevertheless, in this work, we assume $[\alpha, \beta]$ small enough such that d itself is a good approximation; the case of large intervals is left for future work.

In the next section, we illustrate how to practically tackle Problem 1 through polynomial optimization.

4 Polynomial optimization approach

Thanks to the following result, we show that, under Assumption 1, Problem 1 can be solved through non-convex quadratic programming. In the following, given two vectors $a, b \in \mathbb{R}^n$, we write $a \succeq b$ to denote $a_i \geq b_i$ for each $i = 1, \dots, n$. We denote by $I_n \in \mathbb{R}^{n,n}$ the identity matrix. Moreover, $\mathbf{1}_n := (1, 1, \dots, 1)^T \in \mathbb{R}^n$.

Result 1 *Under Assumption 1, Problem 1 can be equivalently written as the following polynomial problem:*

$$\begin{aligned} \min_{x \in [0,d]^n} d\|x\|_1 - \|x\|_2^2 \quad \text{s. t. } Cx \preceq g \\ \text{where} \\ C = \begin{pmatrix} \mathcal{Q}(A) - \Delta_A \mathbf{1}_m \mathbf{1}_n^T \\ -\mathcal{Q}(A) - \Delta_A \mathbf{1}_m \mathbf{1}_n^T \end{pmatrix} \in \mathbb{R}^{2m,n} \\ g = \begin{pmatrix} \mathcal{Q}(y) + \Delta_y \mathbf{1}_m \\ -\mathcal{Q}(y) + \Delta_y \mathbf{1}_m \end{pmatrix} \in \mathbb{R}^{2m}. \end{aligned} \tag{3}$$

Result 1 is obtained by observing that Problem 1 can be rewritten as follows:

$$\begin{aligned} \min_{x \in [0,d]^n} d\|x\|_1 - \|x\|_2^2 \quad \text{s. t. } \|\mathcal{Q}(y) - (\mathcal{Q}(A) - \delta_A)x\|_\infty \leq \Delta_y \\ \|\delta_A\|_\infty \leq \Delta_A \end{aligned} \tag{4}$$

and by applying the results on bounded EIV identification proved in [3], under Assumption 1.

We notice that Problem 1, as written in (3), is non-convex and semi-algebraic, with linear constraints. In particular, when x is known to have only non-negative entries, the ℓ_1 term is a sum, and the problem is a concave quadratic programming problem. To solve Problem 1, we can apply polynomial optimization results by [15], which state that the global minimum of a constrained polynomial problem can be achieved by a hierarchy of relaxed semidefinite programming (SDP) problems. As the order of relaxation increases, the solutions of these SDP converge to the global optimal solution of the original problem. The hierarchy has finite convergence generically, as illustrated in [19, 20].

The SDP relaxation technique is computationally intense in case of large dimensional data. However, in [16] it is shown that if the polynomial problem has a sparse structure, the dimension of the SDP relaxations can be reduced. This polynomial sparsity arises, for example, when the variables are decoupled, which is the case of the functional $d\|x\|_1 - \|x\|_2^2$. Moreover, since the functional is quadratic, also the relaxation order can be kept low, as described in the numerical experiments in Section 6.

We remark that the generalization to a formulation that envisages unknown signs can be derived by exploiting results in [3]. This generalization produces a feasible set which is not overall convex, though convex over each orthant. The solution is then feasible, while computationally intensive for large dimensional problem. The analysis of this problem is left for future work.

5 Analysis

In this section, we provide conditions under which Problem 1 is successful when the non-zero parameters of the system are known to be in the range $[\alpha, \beta]$, $\alpha > 0$, $\beta \geq \alpha$. We start from the simplified case where $\alpha = \beta = d$, i.e., $\tilde{x} \in \{0, d\}^n$.

The following proposition analyses conditions under which \tilde{x} is the unique feasible solution in $\{0, d\}^n$.

Proposition 1 *Let $A\tilde{x} = y$, with $\tilde{x} \in \{0, d\}^n$. Let A_i be the i -th column of A . If*

$$\left\| \sum_{i=1}^n \gamma_i A_i \right\|_{\infty} > \frac{2\Delta_y}{d}, \text{ for any non-null } \gamma \in \{0, \pm 1\}^n, \quad (5)$$

then \tilde{x} is the unique solution of Problem 1.

Proof: Let us consider any $z \in \{0, d\}^n$, $z \neq \tilde{x}$. Then, $w := \tilde{x} - z \in \{0, \pm d\}^n$. If $\left\| \sum_{i=1}^n \gamma_i A_i \right\|_{\infty} > \frac{2\Delta_y}{d}$ for any non-null $\gamma_i \in \{0, \pm 1\}$, then

$$\|Aw\|_{\infty} > 2\Delta_y. \quad (6)$$

This implies:

$$\begin{aligned} & \|Aw \pm \delta_y\|_{\infty} > 2\Delta_y \\ \Rightarrow & \|Aw + \delta_y\|_{\infty} + \|\delta_y\|_{\infty} > 2\Delta_y \\ \Rightarrow & \|Aw + \delta_y\|_{\infty} > 2\Delta_y - \|\delta_y\|_{\infty} \geq \Delta_y. \end{aligned} \quad (7)$$

Since $w = \tilde{x} - z$ and $\mathcal{Q}(y) = A\tilde{x} + \delta_y$, we obtain:

$$\|\mathcal{Q}(y) - Az\|_{\infty} > \Delta_y. \quad (8)$$

The last inequality follows from the fact that $\|\delta_y\|_\infty \leq \Delta_y$. This proves that z does not belong to the feasible set. Therefore, \tilde{x} is the unique feasible solution. \square Concerning the condition $\|\sum_{i=1}^n \gamma_i A_i\|_\infty > \frac{2\Delta_y}{d}$ for any non-null $\gamma \in \{0, \pm 1\}^n$, we highlight that, as discussed in [23], in many applications the columns of A are in general position, i.e., $\sum_{i=1}^n \gamma_i A_i \neq 0$ for any non-null $\gamma \in \{0, \pm 1\}^n$. This definition is rather technical; however, it is proven to hold in many common cases, for instance, when A has random entries generated from a continuous distribution, see Lemmas 3-4 in [23]. Under the general position assumption, there exists a $\sigma > 0$ such that $\|\sum_{i=1}^n \gamma_i A_i\|_\infty \geq \sigma$. Therefore, if d is sufficiently large with respect to Δ_y , we can expect $\frac{2\Delta_y}{d} < \sigma$. On the other hand, σ can be increased by increasing m .

Now, we extend Proposition 1 to the case of non-null \tilde{x}_i belonging to $[\alpha, \beta]$.

Proposition 2 *Let $A\tilde{x} = y$, with $\tilde{x}_i \in \{0\} \cup [\alpha, \beta]$ for any $i = 1, \dots, n$. Let A_i be the i -th column of A . If*

$$\left\| \sum_{i=1}^n \gamma_i A_i \right\|_\infty > 2\Delta_y \text{ for any non-null } \gamma \in Q^n, \quad (9)$$

where

$$Q := \{-d\} \cup \left[\frac{\alpha - \beta}{2}, \frac{\beta - \alpha}{2} \right] \cup [\alpha, \beta]$$

and $d = \frac{\alpha + \beta}{2}$, then Problem 1 has a unique solution $x^* \in \{0, d\}^n$, and x^* has the same support of \tilde{x} .

Proof: We denote by $S(v)$ the support of any $v \in \mathbb{R}^n$. Let us consider any $z \in \{0, d\}^n$, $S(z) \neq S(\tilde{x})$. Let $w := \tilde{x} - z$. Then, $w \in Q^n$.

If $\|\sum_{i=1}^n \gamma_i A_i\|_\infty > 2\Delta_y$ for any $\gamma \in Q^n$, then

$$\|Aw\|_\infty > 2\Delta_y. \quad (10)$$

By computations similar to (7) and (8), we can conclude that z is not in the feasible set. Therefore, the solution of Problem 1 is the unique $x^* \in \{0, d\}^n \cap \mathcal{D}$ such that $S(x^*) = S(\tilde{x})$. \square

We finally notice that in Proposition 1 and Proposition 2, the condition $\|\sum_{i=1}^n \gamma_i A_i\|_\infty > \sigma$ for some $\sigma > 0$ is defined on A , which is not known. However, even if unknown, it makes sense to assume that some information on the properties of A is available. In particular, the fact that the general position of the columns holds can be assumed in many applications. Nevertheless, in the case that σ has to be assessed more precisely, one can state the condition on the observed $\mathcal{Q}(A)$, and the following sufficient condition can be proven.

Proposition 3 *Let $A\tilde{x} = y$, with $\tilde{x}_i \in \{0\} \cup [\alpha, \beta]$ for any $i = 1, \dots, n$. Let $\mathcal{Q}(A) = A + \delta_A$, and let $\mathcal{Q}(A)_i$ be the i -th column of $\mathcal{Q}(A)$. If*

$$\left\| \sum_{i=1}^n \gamma_i \mathcal{Q}(A)_i \right\|_\infty > 2\Delta_y + \Delta_A \beta n \text{ for any non-null } \gamma \in Q^n \quad (11)$$

where $Q = \{-d\} \cup \left[\frac{\alpha - \beta}{2}, \frac{\beta - \alpha}{2} \right] \cup [\alpha, \beta]$, $d = \frac{\alpha + \beta}{2}$, then, Problem 1 has a unique solution $x^* \in \{0, d\}^n$, and x^* has the same support of \tilde{x} .

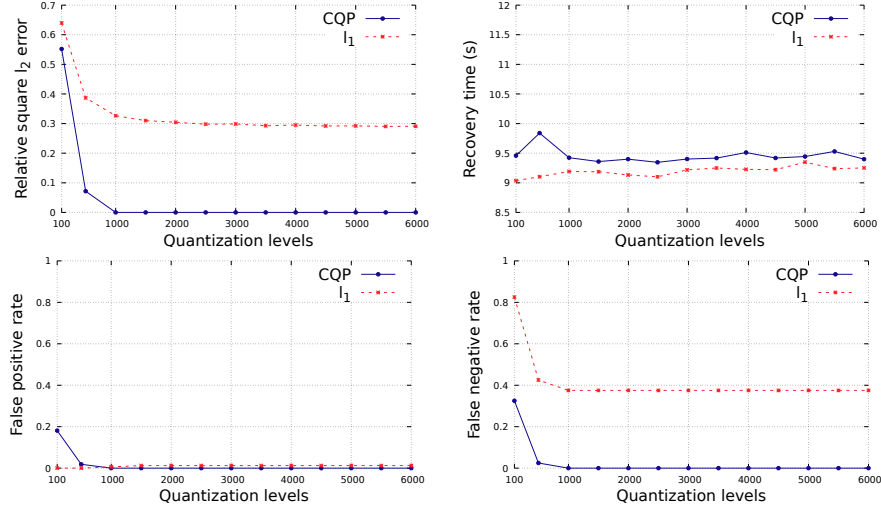


Figure 2: Experiment 1: comparison between the proposed method, denoted by CQP, and the ℓ_1 method [4], under the assumption that $\tilde{x} \in \{0, d\}^n$.

The proof is similar to that of Proposition 2, thus omitted for brevity. As expected, the sufficient condition in Proposition 3 is more restrictive than that of Proposition 2, while the two tend to coincide when the perturbation on A tends to zero.

If we consider the motivating example illustrated in Section 2, condition (5) is fulfilled if $d = 1$, since $|A\gamma| \geq 0.2131 > 0.2 = 2\Delta_y$. Therefore, our approach provides the right solution. Furthermore, condition (9) is fulfilled, for example, with $[\alpha, \beta] = [0.9, 1.1]$ and $\Delta_y = 0.03$.

6 Numerical results

In this section, we present numerical results, which illustrate the effectiveness of the proposed approach with respect to the ℓ_1 approach in [4]. We show two experiments. In the first one, we consider the ideal setting where the true vector of parameters is $\tilde{x} \in \{0, d\}^n$; thus, in this case, we expect to recover exactly both the support and the non-zero values, under suitable conditions. We specify that, despite its peculiarity, the binary setting is relevant in widespread applications, such as localization and image processing, see [8, 10] for a complete overview. In the second experiment, instead, the non-zero components are in $[\alpha, \beta]$, and $d = \frac{\alpha+\beta}{2}$; in this case, we expect to obtain the right support and a biased estimation of the non-zero values.

For both experiments, we propose the following setting. A system acquires a sparse vector $\tilde{x} \in \mathbb{R}^n$, with $n = 10$, through a regression matrix $A \in \mathbb{R}^{m,n}$, $m = 4$. The sparsity level of \tilde{x} is $k = 2$, i.e., only two parameters are non-null. The knowledge of k is not required by the considered recovery algorithms. A is generated according to a Gaussian distribution $\mathcal{N}(0, \frac{1}{m})$; the support of \tilde{x} is generated uniformly at random, while the non-zero entries are uniformly distributed in $[\alpha, \beta]$. We assume that data are uniformly quantized: given a certain

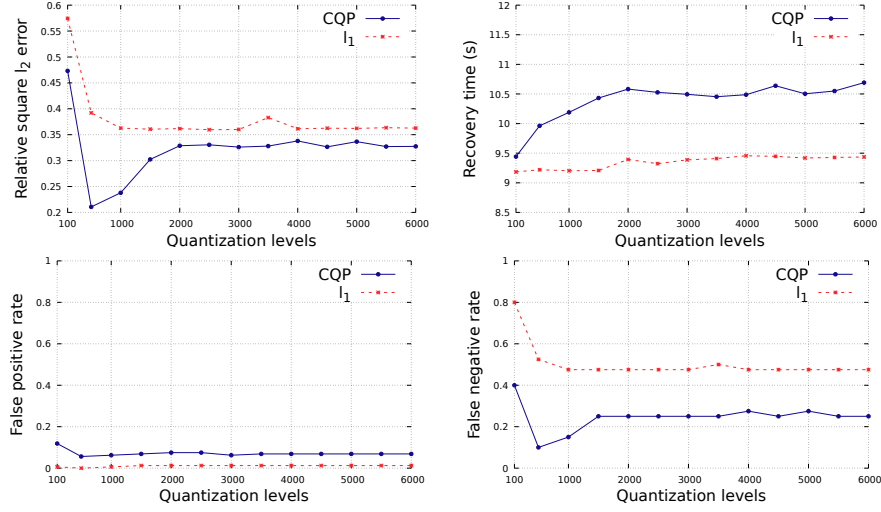


Figure 3: Experiment 2: comparison between the proposed method, denoted by CQP, and the ℓ_1 method [4], under the assumption that the non-zero values of \tilde{x} are in $[0.8, 1.2]$.

number of equidistant quantization levels, each entry of A and y is approximated with the closest point in the quantization codebook. The quantization range is considered sufficiently large to avoid saturation problems. The considered quantization levels range from 100 to 6000, both for A and y .

In the first experiment, depicted in Fig. 2 we consider the case $d = \alpha = \beta = 1$; in the second experiment, illustrated in Fig. 3, $[\alpha, \beta] = [0.8, 1.2]$. The parameters are supposed to be non-negative. The proposed method, based on concave quadratic programming (CQP in the figures) is compared to the ℓ_1 approach of [4] in terms of different performance metrics: the relative square ℓ_2 error, defined as $\|\hat{x} - \tilde{x}\|_2^2 / \|\tilde{x}\|_2^2$, where $\hat{x} \in \mathbb{R}^n$ is the estimate; the false positive rate, defined as the number of events where $\hat{x}_i \neq 0$ while $\tilde{x}_i = 0$, over $n - k$; the false negative rate, defined as the number of events where $\hat{x}_i = 0$ while $\tilde{x}_i \neq 0$, over k ; the run time (in seconds). Experiments are performed on Matlab R2016b, on a 2.3GHz multicore processor. The results are averaged over 20 random runs. The solution of CQP is obtained via SDP relaxation according to [15, 16], using the software YALMIP, with MOSEK solver, see [17, 18]. Concerning the SDP relaxation, we observe the relaxation order equal to 2 is generally sufficient to achieve the minimum and the minimizer of the problem.

In Fig. 2, we observe a significant improvement obtained with the proposed CQP with respect to the ℓ_1 minimization. In particular, when the quantization is sufficiently fine, i.e., more than 1000 quantization levels, the CQP recovery is always exact, while the ℓ_1 approach slightly improves with the decreasing of the perturbation. CQP is more accurate, at the price of a slight increase of run time: 9.5 seconds versus 9.2 seconds on average. We also notice that the false positive rate is small for both methods, which is expected by methods that encourage sparsity in the solution. Thus, it is more significant to observe the false negative rate, which shows whether the significant parameters are selected.

In Fig. 2, we see that the ℓ_1 approach never achieves a false negative rate below 0.375, which means that the erasure of significant parameters is quite frequent.

From Fig. 3, we see that when \tilde{x} is not binary, the estimate provided by CQP is still favorable with respect to the ℓ_1 approach, in terms of relative square ℓ_2 error and false negative rate. In particular, we notice that the ℓ_1 approach does not reach a false negative rate below $\frac{1}{2}$, which is particularly critical. CQP instead achieves values below $\frac{1}{4}$. On the other hand, the false positive rate is a bit unfavorable for CQP, but generally below $\frac{1}{10}$. We also remark that the false negative rate is more relevant for those applications where erasing significant entries is more damaging than preserving a few more parameters. In this second experiment, the run time of CQP is slightly increased, passing from 9.5 seconds to 10.5 seconds, which is however acceptable.

7 Conclusions

In this paper, we address the problem of sparse linear regression from few measurements, when both the regression matrix and the measurements are known in low-precision, or quantized. In the literature, the problem is tackled via ℓ_1 minimization to promote sparsity on the solution; however, this approach is not reliable when different sparse solutions are present in the feasible set. In this paper, we propose a concave quadratic programming approach which, under suitable conditions, is efficient in the support recovery as well as in the estimation of the non-zero parameters. Sufficient conditions that guarantee the exact recovery of the support are provided. Numerical results show enhancement with respect to the ℓ_1 approach, at the price of a slightly increased computational complexity. Future work will envisage a deeper analysis of the performance and the refinement of the sufficient conditions.

References

- [1] A. Bay, D. Carrera, S. M. Fosson, P. Fragneto, M. Grella, C. Ravazzi, and E. Magli. Block-sparsity-based localization in wireless sensor networks. *EURASIP J. Wirel. Commun. Netw.*, 2015(182):1–15, 2015.
- [2] M. Calvo-Fullana, J. Matamoros, C. Antón-Haro, and S. M. Fosson. Sparsity-promoting sensor selection with energy harvesting constraints. In *IEEE ICASSP*, pages 3766–3770, 2016.
- [3] V. Cerone. Feasible parameter set for linear models with bounded errors in all variables. *Automatica*, 29(6):1551 – 1555, 1993.
- [4] V. Cerone, S. M. Fosson, and D. Regruto. A linear programming approach to sparse linear regression with quantized data. In *Proc. American Control Conf. (ACC)*, pages 2990–2995, 2019.
- [5] V. Cerone, D. Piga, and D. Regruto. Set-membership error-in-variables identification through convex relaxation techniques. *IEEE Trans. Autom. Control*, 57(2):517–522, 2012.
- [6] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.

- [7] David L. Donoho and Michael Elad. On the stability of the basis pursuit in the presence of noise. *Signal Processing*, 86(3):511 – 532, 2006.
- [8] S. M. Fossion. Non-convex approach to binary compressed sensing. In *Proc. Asilomar Conf. Signals Syst. Comput.*, pages 1959–1963, 2018.
- [9] S. M. Fossion. Online optimization in dynamic environments: a regret analysis for sparse problems. In *Proc. IEEE Conf. Dec. Contr. (CDC)*, pages 7225–7230, 2018.
- [10] S. M. Fossion and M. Abuabiah. Recovery of binary sparse signals from compressed linear measurements via polynomial optimization. *IEEE Signal Process. Lett.*, 26(7):1070–1074, 2019.
- [11] S. M. Fossion, J. Matamoros, C. Antón-Haro, and E. Magli. Distributed support detection of jointly sparse signals. In *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, pages 6434–6438, 2014.
- [12] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer, New York, 2013.
- [13] M. A. Herman and T. Strohmer. General deviants: An analysis of perturbations in compressed sensing. *IEEE J. Sel. Top. Sign. Proces.*, 4(2):342–349, 2010.
- [14] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk. Democracy in action: Quantization, saturation, and compressive sensing. *Appl. Comput. Harmon. Anal.*, 31(3):429 – 443, 2011.
- [15] J. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.
- [16] J. B. Lasserre. Convergent SDP-relaxations in polynomial optimization with sparsity. *SIAM J. Optim.*, (17):822–843, 2006.
- [17] J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proc. CACSD Conference*, volume 3, 2004.
- [18] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual, Version 8.1.*, 2017.
- [19] J. Nie. Certifying convergence of lasserre’s hierarchy via flat truncation. *Math. Program.*, 142(1):485–510, 2013.
- [20] J. Nie. Optimality conditions and finite convergence of Lasserre’s hierarchy. *Math. Program.*, 146(1):97–121, 2014.
- [21] B. M. Sanandaji, T. L. Vincent, M. B. Wakin, and R. Tóth. Compressive system identification of LTI and LTV ARX models. In *Proc. IEEE Conf. Dec. Contr. (CDC)*, pages 791–798, 2011.
- [22] T. Söderström. *Errors-in-Variables Methods in System Identification*. Springer, 2018.
- [23] Ryan J. Tibshirani. The Lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.

- [24] R. Toth, H. Hjalmarsson, and C.R. Rojas. Sparse estimation of rational dynamical models. In *Proc. IFAC SYSID*, pages 983–988, 2012.
- [25] R. Tóth, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Compressive system identification in the linear time-invariant framework. In *Proc. IEEE Conf. Dec. Contr. (CDC)*, pages 783–790, 2011.
- [26] D. Valsesia, G. Coluccia, and E. Magli. Graded quantization for multiple description coding of compressive measurements. *IEEE Trans. Commun.*, 63(5):1648–1660, 2015.
- [27] Z. Yang, C. Zhang, and L. Xie. Robustly stable signal recovery in compressed sensing with structured matrix perturbation. *IEEE Trans. Signal Process.*, 60(9):4658–4671, 2012.