RECOVERY OF EXACT SPARSE REPRESENTATIONS IN THE PRESENCE OF NOISE.

Jean-Jacques FUCHS

IRISA/Université de Rennes I
Campus de Beaulieu - 35042 Rennes Cedex - France
fuchs@irisa.fr

ABSTRACT

The purpose of this contribution is to extend some recent results on sparse representations of signals in redundant bases developed in the noise-free case to the case of noisy observations. The type of questions addressed so far is : given a (n,m)-matrix A with $m > n$ and a vector $b = Ax$, find a sufficient condition for b to have an unique sparsest representation as a linear combination of the columns of A. The answer is a bound on the number of non-zero entries of say x_o, that guaranties that x_o is the unique and sparsest solution of $Ax = b$ with $b = Ax_o$. We consider the case $b = Ax_o + e$ where x_o satisfies the sparsity conditions requested in the noise-free case and seek conditions on e, a vector of additive noise or modeling errors, under which x_o can be recovered from b in a sense to be defined.

1. INTRODUCTION

Sparse approximation is the problem of finding a representation of a signal as a linear combination of a small number of elements from an over-complete set of vectors or signals often called a dictionary or a redundant basis. Indeed several problems are of interest depending on the context. One may seek (*) the sparsest exact representation of the signal in terms of the elements or (*) the representation of a given complexity that minimizes a certain approximation error or (*) the sparsest representation that yields an approximation error smaller than a specified threshold.

Recently, some theoretical results concerning the first of these questions have been obtained. Given a (n,m)-matrix A with $m > n$ and a vector b that indeed admits an exact sparse representation, say $b = Ax_o$, it has been shown that if the number of nonzero entries in x_o is smaller than a given bound, then x_o is the unique sparsest representation. Since searching for the sparsest representation is a NP hard problem [1] that can only be solved by exhaustive search, one is tempted to replace the true search for the sparsest solution :

$$\min_x \|x\|_0 \quad s.t. \; : Ax = b \qquad (\ell_o)$$

with $\|x\|_0$ the number of nonzero entries in x, by the easy-to-solve linear program :

$$\min_x \|x\|_1 \quad s.t. \; : Ax = b \qquad (LP)$$

with $\|x\|_1 = \sum_1^m |x_j|$, i.e. to minimize the ℓ_1 norm of x instead of the sparsity itself. The problem is then to determine sufficient conditions for the two criteria to have the same unique solution. This problem has been studied in [2] and later been refined in [3] in the special case where A is the union of two orthonormal matrices.

These results have been extended in [4] and [5] to matrices with an arbitrary number of unit norm columns.

We consider here an extension of this problem. We assume the signal to have an exact sparse representation but we observe it in additive noise : $b = Ax_o + e$ with e a vector of additive Gaussian noise $N(0, \sigma^2 I)$ and seek conditions under which x_o can be recovered from the observation of b by solving a convex optimization problem such as a linear or a quadratic program. If one seeks an x that leads to an exact reconstruction of b, it will have generically at least n nonzero components. To get a sparse representation one therefore has to allow for reconstruction errors and the best one can expect is that the optimum of the program and x_o have their nonzero components at the same locations, with the same signs but of course slightly different values. That is we want to recover the true sparse expansion with slightly biased weights. The bias converging to zero as the variance of the noise diminishes.

In the linear regression context, this is known as the subset selection or selection of variables problem and is indeed a difficult detection problem for which only ad hoc solution have been proposed [7].

This paper is organized as follows. The known results are described in section 2, the problem is formalized in section 3 where an mainly unrealistic (because too time consuming) maximum likelihood approach is proposed. In order to mimick the maximum likelihood approach we introduce and analyze a more adequate criterion in section 4. In section 5, we analyze the recovery problem for the traditional noise free case as well as the noisy case and we conclude in section 6.

2. PREVIOUS RESULTS

Consider a set of m n-dimensional vectors a_j with $m > n$ and denote A the (n, m) matrix having these vectors as columns. Any linear combinations b of these m vectors can then be written as : $b = Ax$ with x a m-dimensional vector of weights. In the sequel we assume that the columns of A are normalized : $\|a_i\|_2 = 1, \; \forall i$.

If $b = Ax_o$, with x_o having just a few non-zero components it may be possible to recover x_o from the knowledge of b [2], [3] though this may be a difficult task. We mean that while $Ax = Ax_o$ has generically an infinite number of solutions in x, adding a sparsity requirement may single out x_o. On the other hand even if the *true* solution x_o is the unique sparsest solution, solving (ℓ_o) can in general only be achieved using a combinatorial approach that is essentially unfeasible since too time consuming.

A more realistic problem is therefore to seek conditions under which x_o is the unique solution to :

$$\min_x \|x\|_1 \quad s.t. \; : Ax = b \qquad (LP)$$

with $b = Ax_o$. Minimizing the ℓ_1 norm of the components of the representation is much simpler than minimizing the sparsity, (ℓ_o). It leads to a convex program that can be transformed into an easy-to-solve linear program. To state the sufficient sparsity conditions to be satisfied by x_o, we define the coherence M of the set of vectors $\{a_j\}$ to be [8]:

$$M = \sup_{i \neq j} |a_i^T a_j|, \qquad (1)$$

i.e. the maximal inner product between two distinct vectors. In [2] this problem is considered when A is the concatenation of two orthonormal matrices and it is shown that if :

$$\|x_o\|_0 \leq \frac{1}{2}(1 + \frac{1}{M}) \qquad (2)$$

then x_o is the unique minimum point of (LP) with $b = Ax_o$. This also establishes that x_o is the optimum of (ℓ_o), since if a sparser solution existed it would also satisfy the condition above and one would arrive at a contradiction. This result has been slightly improved in [3] but still for the same limited class of A-matrices.

More recently, in [4] and [5] it has proven that the bound (2) is also valid for A matrices with an arbitrary number m of columns or the concatenation of any number of orthogonal matrices. Of course when m increases, M increases as well and the bound in (2) decreases.

It is worth noting that (2) is independent of the magnitudes of the nonzero entries of x_o. Being able to x_o appears to be only a matter of structure, of angles between vectors. In the sequel where we will take : $b = Ax_o + e$ with e is a Gaussian vector with mean zero and covariance matrix $\sigma^2 I$, this will no longer be true, the relative magnitudes of the nonzero entries of x_o and σ will play a major role. It is only for σ much smaller than the smallest entry of x_o that one can expect to recover the expansion, in a sense to be be defined.

3. PROBLEM FORMULATION

3.1. Model and optimal solution

One observes $b = Ax_o + e$ with $e \in N(0, \sigma^2 I)$ and one knows that the number of nonzero entries in x_o, denoted $\|x_o\|_0$, satisfies (2). In the absence of noise, for $\sigma = 0$, the program (LP) has then an unique solution point x_o. In the presence of noise, the recovery of the true model can only be expected for high signal to noise ratios, i.e. small σ or large non-zero entries in x_o. Let us formalize this observation.

One has to decide upon both the number of nonzero components in x and upon their values. A maximum likelihood approach applied to this problem would consist in the following steps.

For each p, the potential number of components, going from 1 to the integer part of $\frac{1}{2}(1 + \frac{1}{M})$, select the p columns that realize :

$$\min_x \|b - \sum_{k=1}^{p} a_{i_k} x_k\|_2^2 \qquad (3)$$

$$\forall \{i_1, i_2, ..., i_p\} \in \{1, 2, ..., m\}$$

where the minimum is taken over all possible subsets of p columns. Denote S_p^* the minimum of (3) and retain the order p^* that minimizes for instance Akaike's criterion :

$$p^* = \arg\min_p \frac{1}{2} S_p^* + p\sigma^2 \qquad (4)$$

where we assume to know σ^2. This procedure is extremely time consuming. Its complexity is similar to solving (ℓ_o) in the noise-free case, a possibility that we excluded when we switched to the more realistic ℓ_1-norm-"sparsity". We therefore seek a similar approach in the noisy case.

3.2. The LP solution

In the presence of noise, when $b = Ax_o + e$, the optimum of (LP) has generically n nonzero entries since no sparser solution exists generically. One can however expect that for σ small enough, the $\|x_o\|_0$ largest, in absolute value, components correspond to the nonzero entries of x_o, while the others are induced by the noise. Applying some pruning to the components of the optimum to null those that are below a threshold could thus lead to the recovery of the true expansion. In the linear regression analysis literature this is known as *elimination of variables*.

We do not investigate this approach as little can be said about the way the degenerate optimum $x^* = x_o$ of (LP) in the noise-free case evolves as b drifts away from Ax_o. This is due to the fact that for $e = 0$ the optimum of the dual of (LP) is undetermined and it is thus impossible to know a priori which additional components in the optimum point of (LP), the primal, will become non-zero.

3.3. A more realistic solution

The natural extension to the present context of the ℓ_1-norm-"sparsity" approach in the noise-free case is indeed :

$$\min_x \|x\|_1 \quad s.t. : \|Ax - b\|_2^2 \leq B_2 \qquad (5)$$

with $\|x\|_2^2 = \sum_i x_i^2$, in the noisy case. The constraint in (5) is closely related to the maximum likelihood criterion (3) and as we will point it out below the ℓ_1-norm has much to do with the penalty term in Akaike's criterion (4).

This is a convex optimization problem that is equivalent to :

$$\min_x \frac{1}{2}\|Ax - b\|_2^2 + h\|x\|_1 \qquad (QP)$$

for an adequately chosen parameter $h > 0$. Indeed if h is taken equal the inverse of twice the Lagrange multiplier of the constraint in (5) at the optimum, then both problems have the same optimum. Now (QP) is in turn identical to :

$$\min_x \|Ax\|_2^2 \quad s.t. : \|A^T(Ax - b)\|_\infty \leq h \qquad (DQP)$$

with $\|x\|_\infty = \max_i |x_i|$. The two programs (QP) and (DQP) can be transformed into quadratic programs that are equivalent by duality [9].

In the sequel, we mainly concentrate on (QP) which is the most convenient for our purpose. It is worth noting the analogy between (QP) and the combination of (3) and (4), the regularization term in (QP) : $h\|x\|_1$ replacing $p\sigma^2$ in (4).

4. THE QUADRATIC PROGRAMMING APPROACH

Before we use (QP) to possibly recover a sparse solution x_o satisfying (2) from the observation of $b = Ax_o + e$, let us first discuss its use in the noise free case.

When $e = 0$, one can take $B_2 = 0$ in (5) which is then identical to (LP), this corresponds to taking $h = 0^+$ in (QP) and

(DQP) as can be deduced from the results obtained below. In the noise-free case, for $h = 0^+$, the optimum of (QP) is thus x_o just like the optimum of (LP) provided x_o satisfies (2). One can thus expect that for $h > 0$ small enough, this continues to hold. This has indeed been established in [5].

For non-zero h, the optimum point, say $\hat{x}_o$, of (QP) achieves a compromise between the two terms of the criterion $\|Ax - b\|_2^2$ and $\|x\|_1$. At the optimum one therefore has : $A\hat{x}_o \neq b$, and the best one can expect is that $\hat{x}_o$ and x_o have their non-zero components at the same locations and with the same signs. If this holds we say that *(QP) allows to recover x_o*.

This is also the best one can expect in the presence of noise. Then the discrepancy between $\hat{x}_o$ and x_o is due to both h being positive and the presence of the noise e.

In order to investigate the conditions under which recovery of x_o is possible we need to detail the optimality conditions of (QP).

They could be obtained in a quite traditional way by first transforming the problem into a quadratic program (by doubling the number of unknowns) and writing the first order necessary (Kuhn Tucker) conditions that are also sufficient since the problem is convex. We adopt a more direct path. Since the criterion is non-smooth at zero because of the presence of $\|x\|_1$, we introduce the sub-differential of $\|x\|_1$ [10], a set of vectors called the sub-gradients, denoted $\partial \|x\|_1$:

$$\partial \|x\|_1 = \{u | u^T x = \|x\|_1,\ \|u\|_\infty \leq 1\}$$
$$= \{u | u_i = \text{sign}(x_i) \text{ if } x_i \neq 0 \text{ and } |u_i| \leq 1 \text{ otherwise}\}$$

where $\text{sign}(x_i) = 1$ when $x_i > 0$ and $\text{sign}(x_i) = -1$ when $x_i < 0$. A necessary and sufficient condition (NSC) for x^* to be a global minimum of (QP) is that the vector zero is a sub-gradient of the criterion at x^* [10] :

$$\exists\ u \in \partial \|x\|_1 \quad \ni \quad A^T (Ax^* - b) + hu = 0 \qquad (NSC)$$

To write these NSC in a more usable way, we distinguish between the non-zero components and the zero components of x^*. We denote $\bar{x}^*$ the reduced dimensional vector built upon the non-zero components of x^*. Similarly $\bar{A}$ denotes the associated columns in A. One then has e.g. $Ax^* = \bar{A}\bar{x}^*$. For the rows in (NSC) associated with the non-zero entries $\bar{x}^*$, the sub-gradient is unique, equal to the gradient and known, for the other rows the sub-gradient takes any value in $[-1, 1]$. The necessary and sufficient conditions become :

$$\bar{A}^T (b - \bar{A}\bar{x}^*) = h\ \text{sign}\bar{x}^* \qquad (NSC_1)$$
$$|a_j^T (b - \bar{A}\bar{x}^*)| \leq h,\ \forall\ a_j \notin \bar{A} \qquad (NSC_2)$$

One can further establish that if $\bar{A}$ is full rank and the inequalities in (NSC_2) satisfied strictly, then x^* is the *unique* minimum point of (QP). The first of these relations then leads to :

$$\bar{x}^* = \bar{A}^+ b - h\ (\bar{A}^T \bar{A})^{-1}\ \text{sign}\bar{x}^* \qquad (6)$$

with $\bar{A}^+ = (\bar{A}^T \bar{A})^{-1} \bar{A}^T$ a pseudo-inverse of $\bar{A}$.

The optimum of (QP) can only be obtained through an iterative search and indeed the relation (6) is implicit since $\bar{x}^*$ is present on both sides. This relation however allows to compute $\bar{x}^*$ and thus x^* if one knows beforehand the indices and the signs of the non-zero components of x^*. Together with (NSC_2), it is thus helpful to test if a candidate point is indeed the optimum.

The last term in (6) is a bias term induced by the regularization term in (QP), this bias can be removed if desired since its analytical form is known.

5. RECOVERY CONDITIONS

5.1. The noise-free case

We are now ready to write the conditions under which the optimum of (QP) allows to recover x_o, when $b = Ax_o$.

We denote $\hat{x}_o$ this optimum that allows to recover x_o in the sense defined above and introduce $\bar{A}_o$ and $\bar{x}_o$ such that $b = Ax_o = \bar{A}_o\bar{x}_o$ and similarly $\hat{\bar{x}}_o$ and $\tilde{\bar{A}}_o$ that is identical to $\bar{A}_o$. The NS conditions at $\hat{\bar{x}}_o$, $\bar{A}_o$ yield :

$$\bar{A}_o^T (b - \bar{A}_o\hat{\bar{x}}_o) = h\ \text{sign}\bar{x}_o \qquad (NSC_1)$$
$$|a_j^T (b - \bar{A}_o\hat{\bar{x}}_o)| \leq h\ \text{ for } a_j \notin \bar{A}_o \qquad (NSC_2)$$

since $b = \bar{A}_o\bar{x}_o$, the first relation gives :

$$\hat{\bar{x}}_o = \bar{x}_o - h\ (\bar{A}_o^T \bar{A}_o)^{-1}\ \text{sign}\bar{x}_o \qquad (7)$$

Replacing now $\hat{\bar{x}}_o$ by its value (7) in (NSC_2) one gets :

$$|a_j^T d\ | < 1,\ \forall\ a_j \notin \bar{A}_o\ \text{ and } d = \bar{A}_o^{+T}\text{sign}\bar{x}_o \qquad (8)$$

Since $\bar{A}_o$ is full rank, if (7) and (8) are satisfied, the unique optimum of (QP) allows to recover x_o. As far as (7) is concerned, one has to check that the sign of the resulting $\hat{\bar{x}}_o$ is indeed the same as the sign of the true $\bar{x}_o$. This is always possible by taking h small enough since all the components in $\bar{x}_o$ are strictly different from 0.

It is (8) that is the true limiting condition. It can be called a separability condition - since defining two hyperplanes : $H_\pm = \{a | a^T d = \pm 1\}$ - condition (8) says that the vectors $a_j \notin \bar{A}_o$ must lie strictly in between these hyperplanes while the other vectors $a_j \in \bar{A}_o$ lie in these hyperplanes, by the very definition of d. In [5], we have shown, using results from numerical analysis, that if (2) is satisfied then the vector d defined in (8) verifies the constraints in (8).

This means that if x_o a solution of $Ax = b$ is sparse enough to satisfy (2) then it is the unique solution point of (QP) for h small enough and thus also the unique sparsest solution of $Ax = b$.

5.2. The noisy case

We proceed as in the previous section and replace in the NS conditions b by $Ax_o + e$ where we assume that x_o satisfies (2). By continuity one can expect that for e small enough and h adequately fixed, the optimum $\hat{x}_o$ of (QP) will have the same non-zero components as x_o with the same signs, since this is is true for $e = 0$.

With the same notation as above, the first part of the NS conditions (7) becomes :

$$\begin{aligned} \hat{\bar{x}}_o &= \bar{A}_o^+ b - h\ (\bar{A}_o^T \bar{A}_o)^{-1}\ \text{sign}\bar{x}_o \\ &= \bar{x}_o + \bar{A}_o^+ e - h\ (\bar{A}_o^T \bar{A}_o)^{-1}\ \text{sign}\bar{x}_o \qquad (9) \end{aligned}$$

and the second part (NSC_2) :

$$|a_j^T (b - \bar{A}_o\hat{\bar{x}}_o)| \leq h,\ \forall\ a_j \notin \bar{A}_o$$

can be transformed, by substituting $\hat{\bar{x}}_o$ (9), into :

$$|a_j^T ((I - \bar{A}_o\bar{A}_o^+)e + h\bar{A}_o(\bar{A}_o^T \bar{A}_o)^{-1}\ \text{sign}\bar{x}_o)| \leq h$$
$$|a_j^T ((I - \bar{A}_o\bar{A}_o^+)e + hd)| \leq h$$
$$|a_j^T (e^\perp + hd)| \leq h,\ \forall a_j \notin \bar{A}_o \qquad (10)$$

with $d = \bar{A}_o^{+T}\mathrm{sgn}\bar{x}_o$ and $e^\perp = (I - \bar{A}_o\bar{A}_o^+)e$.

The two conditions in (9) and (10) are again of quite different nature. In (9) provided e is small enough for $\mathrm{sign}(\bar{x}_o + \bar{A}_o^+ e) = \mathrm{sign}\,\bar{x}_o$ to hold, one can always choose h for the sign of $\hat{\bar{x}}_o$ to be equal to the sign $\bar{x}_o$. For admissible e the relation fixes an upper-bound on h.

The conditions in (10) are satisfied if :

$$|a_j^T e^\perp| < (1 - |a_j^T d|)h, \ \ \forall a_j \notin \bar{A}_o. \qquad (11)$$

Since $|a_j^T d| < 1$ because x_o satisfies (2) and thus (8), (11) holds if h is taken large enough. Condition (11) fixes a lower-bound for h.

To be sure that the two conditions in (9) and (10) can be met simultaneously, it is important to observe that as e decreases, the upper-bound on h - implicitly defined by (9) - increases and the lower bound on h - implicitly defined by (11) - decreases. For e small enough it is thus always possible to find a domain in h in which (QP) allows to recover x_o. We have only been considering the possibility to decrease the noise e but the conclusions are the same if instead of decreasing e, one increases the magnitudes of the non-zero weights in $\bar{x}_o$. It is of course the signal to noise ratio that is important and that has to be large enough.

We have thus established that if the signal to noise ratio is large enough the optimum of (QP) allows to recover the true x_o, i.e. to select the true subset of vectors, provided that h is adequately chosen. This justifies the preference that should be given to (QP) over (LP).

Remark :

The analysis has been qualitative making it quantitative seems to be a difficult task. Indeed in conditions such as (11) :

$$|a_j^T e^\perp| < (1 - |a_j^T d|)h, \ \ \forall a_j \notin \bar{A}_o.$$

while the random variables $a_j^T e^\perp$ are Gaussian with mean zero and computable variance, the fact that they are not independent for distinct values of j makes any evaluation of the probability of these conditions to be satisfied simultaneously impossible, even though a common lower bound for the right sides can be computed. The same holds true for (9) :

$$\hat{\bar{x}}_o = \bar{x}_o + \bar{A}_o^+ e - h\,(\bar{A}_o^T\bar{A}_o)^{-1}\,\mathrm{sign}\bar{x}_o$$

where a preliminary condition is that $\mathrm{sign}(\bar{x}_o + \bar{A}_o^+ e) = \mathrm{sign}\,\bar{x}_o$. The probability of this being true for a given variance σ^2 of the components in e together with the conditions in (11) seems again extremely difficult to evaluate.

Indeed it seems that only the case where one assumes to known that only one component in x_o is non-zero seems to be amenable to a complete analysis. Without loss of generality one can take this unique component equal to one. The maximum likelihood approach amounts to select :

$$j^* = \arg\max_i |a_i^T b|,$$

the maximum output of the matched filter. The probability of detection, i.e. of j^* to be the index of the *true* unique non-zero component of x_o and the probability of false alarm are then computable, at least in theory. There is no need for a more elaborate algorithm such as (QP) though it is quite easy to show that if h is tuned for the optimum of (QP) to have an unique non-zero entry, then this entry has index j^*.

6. CONCLUSIONS

We have extended a deterministic result on recovery of sparse representations to a stochastic context by replacing the usual linear program minimizing the ℓ_1-norm of the weights by a quadratic program depending upon a regularization parameter h that allows to adapt the criterion to the signal to noise ratio and encompasses the linear program as a special case.

The problem amounts to consider a regression model in which one has to choose a very small subset of regressors out of a very large set of possible candidate regressors. We have shown that provided the *true* number of regressors satisfies condition (2) and the signal to noise ratio is large enough, solving the quadratic program for an adequately tuned regularization parameter leads to the *true* subset of regressors.

For this quite general problem we have obtained qualitative feasibility results. Further investigations on more specific problems can be conducted in a number of directions. If the *true* number of regressors is known a priori, it may be possible to perform a more precise analysis and to investigate how far a well tuned (QP) algorithm is from the *optimal* but cumbersome maximum likelihood approach. The case where $\|x_o\|_0 = 1$ is feasible along the lines described above but of little interest, but the case $\|x_o\|_0 = 2$ is already challenging. We have considered e to be Gaussian, hence the ℓ_2 norm and the (QP) criterion but other models for e may be of interest as well.

7. REFERENCES

[1] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24, 21, 227-234, april 1995.

[2] D.L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. on I.T.*, 47, 11, 2845-2862, nov; 2001.

[3] M. Elad an A.M. Bruckstein. A generalized uncertainty principle and sparse representation in pairs of bases. *IEEE Trans. on I.T.*, 48, 9, 2558-256, sept. 2002.

[4] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *Technical Report INRIA*, 1499, nov. 2002 submitted to *IEEE Trans. on I.T.* nov. 2002.

[5] J.J. Fuchs. More on sparse representations in arbitrary bases. *13th Ifac-Ifors symposium on Identification and system Parameter Estimation*, pp. 1357-1362, Rotterdam, august 2003, submitted to *IEEE Trans. on I.T.* dec. 2002.

[6] A. Feuer and A. Nemirovski. On sparse representations in pairs of bases. *IEEE Trans. on I.T.*, 49, 6, 1579-1581, june 2003.

[7] N.R. Draper and H. Smith. Applied Regression. *John Wiley and Sons*, 1966.

[8] A. Gilbert, S. Muthukrishnan and M.J. Strauss. Approximation of functions over redundant dictionaries using coherence. *14th ACM-SIAM Symposium on Discrete Algorithms*, (SODA'03), Jan. 2003.

[9] J.J. Fuchs. On the application of the global matched filter to DOA estimation with uniform circular arrays. *IEEE-T-SP*, vol. 49, p. 702–709, avr. 2001.

[10] R. Fletcher. Practical methods of optimization. *John Wiley and Sons*, 1987.