Sparse Non-Negative Solution of a Linear System of Equations is Unique

Alfred M. Bruckstein, Michael Elad and Michael Zibulevsky
Computer Science Department – Technion–Israel Institute of Technology
Haifa 32000, Israel, Email:[elad,freddy,mzib]@cs.technion.ac.il

***Abstract*—We consider an underdetermined linear system of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ with non-negative entries of $\mathbf{A}$ and $\mathbf{b}$, and the solution $\mathbf{x}$ being also required to be non-negative. We show that if there exists a sufficiently sparse solution to this problem, it is necessarily unique. Furthermore, we present a greedy algorithm – a variant of the matching pursuit – that is guaranteed to find this sparse solution. The result mentioned above is obtained by extending the existing theoretical analysis of the Basis Pursuit problem, i.e. $\min \|\mathbf{x}\|_1$ s.t. $\mathbf{A}\mathbf{x} = \mathbf{b}$, by studying conditions for perfect recovery of sparse enough solutions. Considering a matrix $\mathbf{A}$ with arbitrary column norms, and an arbitrary monotone element-wise concave penalty replacing the ℓ_1-norm objective function, we generalize known equivalence results, and use those to derive the above uniqueness claim.**

I. Introduction

This paper is devoted to the analysis of underdetermined linear systems of equations of the form $\mathbf{A}\mathbf{x} = \mathbf{b}$, where the entries of $\mathbf{A} \in \mathbb{R}^{n\times k}$ and $\mathbf{b} \in \mathbb{R}^n$ are all non-negative[1], and the desired solution $\mathbf{x} \in \mathbb{R}^k$ is also required to be non-negative. Such problems are frequently encountered in signal and image processing, in handling of multi-spectral data, considering non-negative factorization for recognition, and more (see [3], [19], [24], [27], [28], [31] for representative work).

When considering an underdetermined linear system (i.e. $k > n$), with a full rank matrix $\mathbf{A}$, the removal of the non-negativity requirement $\mathbf{x} \geq 0$ leads to an infinite set of feasible solutions. How is this set reduced when we further require a non-negative solution? How can these solutions be effectively found in practice?

Assuming there could be several possible solutions, the common practice is the definition of an optimization problem of the form

$$(P_f): \quad \min_{\mathbf{x}} \ f(\mathbf{x}) \ \text{ subject to } \ \mathbf{A}\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0, \qquad (1)$$

where $f(\cdot)$ measures the quality of the candidate solutions. Possible choices for this penalty could be various entropy measures, or general ℓ_p-norms for various p in the range $[0, \infty)$. Popular choices are $p = 2$ (minimum ℓ_2-norm), $p = 1$ (minimum ℓ_1-norm), and $p = 0$ (enforcing sparsity). For example, recent work reported in [11]–[13] proved that the $p = 0$ and $p = 1$ choices lead to the same result, provided that this result is sparse enough. This work also provides bounds on the required sparsity that guarantee such equivalence, but in a different setting.

Clearly, if the set of feasible solutions $\{\mathbf{x} | \ \mathbf{A}\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0\}$ contains only one element, then all the above choices of $f(\cdot)$ will lead to the same solution. In such a case, the above-discussed ℓ_0-ℓ_1 equivalence becomes an example of a much wider phenomenon. Surprisingly, this is exactly what happens when a sufficiently sparse solution exists. The main result shown of this paper proves the uniqueness of such a sparse solution, and provides a bound on $\|\mathbf{x}\|_0$ below which such a solution is guaranteed to be unique.

There are several known results reporting an interesting behavior of sparse solutions of a general underdetermined linear system of equations, when minimum of ℓ_1-norm is imposed on the solution (this is the Basis Pursuit algorithm (BP)) [7], [17]. These results assume that the columns of the coefficient matrix have a unit ℓ_2-norm, and state that the minimal ℓ_1-norm solution coincides with the sparsest one, for sparse enough solutions. As mentioned above, a similar claim is made in [11]–[13] for non-negative solutions, leading to stronger bounds.

In this work we extend the BP analysis, presented in [7], [17], to the case of a matrix with arbitrary column norms and an arbitrary monotone element-wise concave penalty replacing the ℓ_1-norm objective function. A generalized theorem of the same flavor is obtained. Using this result, we get conditions of uniqueness of sparse solutions of non-negative system of equations, as mentioned above. Interestingly, there is no need for the ℓ_1 penalty in these cases – non-negativity constraints are sufficient to lead to the unique (and sparsest) solution.

Returning to the practical side of things, and assuming that we are interested in the sparsest (and possibly the only) feasible solution of

$$(P_0^+): \quad \min_{\mathbf{x}} \ \|\mathbf{x}\|_0 \ \text{ subject to } \ \mathbf{A}\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0, \qquad (2)$$

there are several possible numerical methods for solving this problem. We present a variant of the Orthogonal Matching Pursuit (OMP) for this task. We provide a theoretical analysis of this algorithm that shows that it is guaranteed to lead to the desired solution, if it is indeed sparse enough.

The structure of this paper is as follows: In Section 2 we extend the BP analysis to the case of arbitrary monotone element-wise concave penalty and matrix $\mathbf{A}$ with arbitrary column norms. This analysis relies on a special definition of coherence measure of the matrix $\mathbf{A}$. In Section 3 we develop the main theoretical result in this paper, claiming that a sufficiently sparse solution of $\{\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq 0\}$ is unique.

[1] This non-negativity requirement could be substantially relaxed, as we shall see in Section 3.

We also introduce preconditioning that improves the above-mentioned coherence. Section 4 presents the OMP variant for the non-negative problem, along with empirical and theoretical analysis of its performance. e note that a longer and more detailed version of this paper is available in [1].

II. Basis Pursuit: An Extended Result

In this section we develop a theorem claiming that a sufficiently sparse solution of a general underdetermined linear system $\mathbf{Dz} = \mathbf{b}$ is necessarily a minimizer of a separable concave function[2]. It extends Theorem 7 in [7] in several ways:

- It does not assume normalization of the columns in $\mathbf{D}$ (and thus it is more general);
- It relies on a different feature of the matrix $\mathbf{D}$ – a one-sided coherence measure;
- The objective function is more general than the ℓ_1-norm used in [7]. In fact, it is similar to the one proposed by Gribonval and Nielsen in [18], but due to the above changes, the analysis is different.

The results presented in this section form the grounds for the analysis of non-negative linear systems in Section 3.

For an arbitrary $n \times k$ matrix $\mathbf{D}$ with columns $\mathbf{d}_i$ we define its one-sided coherence as

$$\rho(\mathbf{D}) = \max_{i,j; j\neq i} \frac{|\mathbf{d}_i^T \mathbf{d}_j|}{\|\mathbf{d}_i\|_2^2}. \tag{3}$$

Defining the Gram matrix $\mathbf{G} = \mathbf{D}^T\mathbf{D}$, its elements satisfy

$$\frac{|G_{ij}|}{G_{ii}} \le \rho(\mathbf{D}) \qquad \forall i, j \neq i. \tag{4}$$

This measure tends to behave like $1/\sqrt{n}$ for random matrices, very much like the regular *mutual-coherence* as defined in [7].

Lemma 1: Any vector δ from the null-space of $\mathbf{D}$ satisfies

$$\|\delta\|_\infty \le t_{\mathbf{D}} \|\delta\|_1, \tag{5}$$

where we denote

$$t_{\mathbf{D}} = \frac{\rho(\mathbf{D})}{1+\rho(\mathbf{D})}. \tag{6}$$

Proof: Multiplying the null-space condition $\mathbf{D}\delta = 0$ by $\mathbf{D}^T$, and using $\mathbf{G} = \mathbf{D}^T\mathbf{D}$, we get $\mathbf{G}\delta = 0$. The i-th row of this equation,

$$G_{ii}\delta_i + \sum_{j\neq i} G_{ij}\delta_j = 0, \quad i = 1, 2, \ \dots\ , k, \tag{7}$$

gives us

$$\delta_i = -\sum_{j\neq i} \frac{G_{ij}}{G_{ii}} \delta_j, \quad i = 1, 2, \ \dots\ , k. \tag{8}$$

Taking absolute value of both sides, we obtain

$$|\delta_i| = \left| \sum_{j\neq i} \frac{G_{ij}}{G_{ii}} \delta_j \right| \le \sum_{j\neq i} \left| \frac{G_{ij}}{G_{ii}} \right| |\delta_j| \le \rho(\mathbf{D}) \sum_{j\neq i} |\delta_j|, \tag{9}$$

[2] We introduce here a different notation for the linear system: $\mathbf{Dz} = \mathbf{b}$, instead of $\mathbf{Ax} = \mathbf{b}$, for reasons to be clarified in the next Section.

where the last inequality is due to (4). Adding a term $\rho(\mathbf{D})|\delta_i|$ to both sides, we get

$$|\delta_i| \le \frac{\rho(\mathbf{D})}{1+\rho(\mathbf{D})} \|\delta\|_1 = t_{\mathbf{D}} \|\delta\|_1, \quad i = 1, 2, \ \dots\ , k. \tag{10}$$

Thus, $\|\delta\|_\infty \le t_{\mathbf{D}}\|\delta\|_1$, as the Lemma claims. □

Theorem 1: Consider the following optimization problem:

$$\min_{\mathbf{z}} \ \sum_{i=1}^{k} \varphi(|z_i|) \quad \text{subject to} \quad \mathbf{Dz} = \mathbf{b} \tag{11}$$

with a scalar function $\varphi : \mathbb{R}_+ \to \mathbb{R}$ satisfying:

- φ is concave (see [1]);
- φ is increasing semi-monotonic: $\varphi(t_2) \ge \varphi(t_1)$ for all $t_2 > t_1 > 0$; and
- $\varphi(t) > \varphi(0) \quad \forall t > 0$ (in order to avoid $\varphi(t) \equiv 0$).

A feasible solution $\bar{\mathbf{z}}$ (i.e. $\mathbf{D}\bar{\mathbf{z}} = \mathbf{b}$) is a unique global optimum of (11) if

$$\|\bar{\mathbf{z}}\|_0 < \frac{1}{2t_{\mathbf{D}}} \ , \tag{12}$$

where $t_{\mathbf{D}}$ is given by (6).

Proof: Adding a constant to the objective function does not change the solution. Therefore, without loss of generality, we shall assume hereafter that

$$\varphi(0) = 0.$$

We intend to show that under the conditions of the theorem, any feasible non-zero deviation vector $\delta \in \mathbb{R}^k$ of $\bar{\mathbf{z}}$ (i.e. $\bar{\mathbf{z}} + \delta$) necessarily leads to an increase in the objective function, namely

$$\sum_{i=1}^{k} \varphi(|\bar{z}_i + \delta_i|) > \sum_{i=1}^{k} \varphi(|\bar{z}_i|). \tag{13}$$

Feasibility of the deviation vector means

$$\mathbf{D}(\bar{\mathbf{z}} + \delta) = \mathbf{b}, \tag{14}$$

implying $\mathbf{D}\delta = \mathbf{0}$. Thus, by Lemma 1 we can state that

$$|\delta_i| \le t_{\mathbf{D}}\|\delta\|_1 \equiv \delta_{\text{tol}}, \quad i = 1, \ \dots\ , k. \tag{15}$$

Among all vectors δ of fixed total amplitude $\|\delta\|_1 = \delta_{\text{tol}}/t_{\mathbf{D}}$ satisfying inequalities (15), we shall try to compose one that reduces the objective function. For this, we separate the summation of the objective function into two parts – the on-support elements (i.e. those with $|\bar{z}_i| > 0$), and the off-support ones (where $\bar{z}_i = 0$). We denote the support of $\bar{\mathbf{z}}$ as Γ and write

$$\sum_{i=1}^{k} \varphi(|\bar{z}_i + \delta_i|) = \sum_{i\in\Gamma} \varphi(|\bar{z}_i + \delta_i|) + \sum_{i\notin\Gamma} \varphi(|\delta_i|). \tag{16}$$

We consider first the on-support term. Taking into account monotonicity and concavity of $\varphi(\cdot)$, a maximal decrease of the objective function would be possible for the choice

$$\delta_i = \begin{cases} -\delta_{\text{tol}} \cdot \text{sign}(\bar{z}_i) & |\bar{z}_i| \ge \delta_{\text{tol}} \\ -\bar{z}_i & 0 < |\bar{z}_i| < \delta_{\text{tol}} \end{cases} . \tag{17}$$

For this assignment, the descent in the objective function is given by

$$\begin{aligned} E_\Gamma &= \sum_{i\in\Gamma}\varphi(|\bar{z}_i|) - \sum_{i\in\Gamma}\varphi(|\bar{z}_i+\delta_i|) \\ &= \sum_{i\in\Gamma}\varphi(|\bar{z}_i|) - \sum_{i\in\Gamma}\varphi(|\bar{z}_i|-\delta_i) \\ &= \sum_{i\in\Gamma}\varphi(|\bar{z}_i|) - \sum_{i\in\Gamma}\varphi\left(\max\{|\bar{z}_i|-\delta_{\text{tol}},0\}\right) \\ &\le |\Gamma|\cdot\varphi(\delta_{\text{tol}}). \end{aligned} \tag{18}$$

The last inequality is a direct consequence of the zero-bias ($\varphi(0)=0$), monotonicity, and the concavity properties of the function $\varphi(\cdot)$.

Turning to the off-support term in Equation (16), any assignment of $\delta_i \neq 0$ implies an ascent. Of the original deviation vector δ, we are left with a total amplitude of at least $\|\delta\|_1 - |\Gamma|\cdot\delta_{\text{tol}} = \|\delta\|_1\cdot(1-t_{\mathbf{D}}|\Gamma|)$, to be assigned to the off-support elements. Again, due to the concavity of $\varphi(\cdot)$, this remaining energy leads to the smallest possible ascent if the assignment chosen is δ_{tol} to as few as possible elements. Thus, the obtained ascent becomes

$$\begin{aligned} E_{\bar{\Gamma}} &= \frac{\|\delta\|_1 - |\Gamma|\cdot\delta_{\text{tol}}}{\delta_{\text{tol}}}\cdot\varphi(\delta_{\text{tol}}) \\ &= \frac{\|\delta\|_1\cdot(1-t_{\mathbf{D}}|\Gamma|)}{\|\delta\|_1\cdot t_{\mathbf{D}}|\Gamma|}\cdot\varphi(\delta_{\text{tol}}) = \frac{1-t_{\mathbf{D}}|\Gamma|}{t_{\mathbf{D}}|\Gamma|}\cdot\varphi(\delta_{\text{tol}}). \end{aligned} \tag{19}$$

In order for $\bar{\mathbf{z}}$ to be a unique global minimizer of the problem posed in Equation (11), the change of the objective function should be positive – i.e., $E_\Gamma < E_{\bar{\Gamma}}$, implying

$$\frac{1-t_{\mathbf{D}}|\Gamma|}{t_{\mathbf{D}}|\Gamma|}\cdot\varphi(\delta_{\text{tol}}) > |\Gamma|\cdot\varphi(\delta_{\text{tol}}). \tag{20}$$

which is always satisfied if $|\Gamma| \equiv \|\bar{\mathbf{z}}\|_0 < \frac{1}{2t_{\mathbf{D}}}$, as claimed. □

III. Non-Negative Systems of Equations

We now turn to the main result of this paper, showing that a sufficiently sparse solution of a linear system is necessarily the only non-negative possible solution. Afterwards, we show how to preconditioning of the linear system can be used to strengthen this theorem.

Suppose that we are given a system of linear of equations with non-negativity constraints

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \quad \mathbf{x} \ge 0, \tag{21}$$

with non-negative $\mathbf{A} \in \mathbb{R}^{n\times k}$ and $\mathbf{b} \in \mathbb{R}^n$. In order to simplify the exposition, we re-scale the problem to have unit column sums of the coefficients. Let $\mathbf{W}$ be a diagonal matrix with the entries $W_{jj} = \sum_i \mathbf{A}_{ij}$. We assume that there are no zero columns of in $\mathbf{A}$, and thus $\mathbf{W}$ is invertible. The equivalent system is

$$\mathbf{A}\mathbf{W}^{-1}\mathbf{W}\mathbf{x} = \mathbf{b}, \quad \mathbf{x} \ge 0. \tag{22}$$

Denoting $\mathbf{D} \equiv \mathbf{A}\mathbf{W}^{-1}$ and $\mathbf{z} \equiv \mathbf{W}\mathbf{x}$, we get the normalized system with $\sum_i \mathbf{D}_{ij} = 1$:

$$\mathbf{D}\mathbf{z} = \mathbf{b}, \quad \mathbf{z} \ge 0. \tag{23}$$

Denote $\mathbf{1}_n$ a column vector of n ones. Multiplying the last equation with $\mathbf{1}_n^T$, and using the fact that $\mathbf{1}_n^T\mathbf{D} = \mathbf{1}_k^T$, we obtain

$$\mathbf{1}_k^T\mathbf{z} = \mathbf{1}_n^T\mathbf{b} = c, \tag{24}$$

where c denotes the sum of the entries in $\mathbf{b}$.

Theorem 2: Suppose that we are given a system of linear equations with non-negativity constraints $\mathbf{D}\mathbf{z} = \mathbf{b}, \quad \mathbf{z} \ge 0$, such that all its solutions satisfy $\mathbf{1}^T\mathbf{z} = c$, where c is some constant. If a vector $\bar{\mathbf{z}}$ is a sparse solution of this system with $\|\bar{\mathbf{z}}\|_0 < \frac{1}{2t_{\mathbf{D}}}$, then it is a unique solution of this problem[3].

Proof: Taking into account non-negativity of $\mathbf{z}$, we can rewrite the condition $\mathbf{1}^T\mathbf{z} = c$ differently, as

$$\|\mathbf{z}\|_1 = c. \tag{25}$$

The vector $\bar{\mathbf{z}}$ is a sparse (with less than $1/2t_{\mathbf{D}}$ non-zeros) feasible solution of the linear programming problem

$$\min_{\mathbf{z}} \ \|\mathbf{z}\|_1 \quad \text{subject to} \quad \mathbf{D}\mathbf{z} = \mathbf{b}. \tag{26}$$

Notice that we do not specify the constraint $\mathbf{1}_k^T\mathbf{z} = c$ because any feasible solution of this problem must satisfy this condition anyhow. Also, we do not add a non-negativity constraint – the problem is defined as described above, and we simply observe that $\bar{\mathbf{z}}$ is a feasible solution.

By Theorem 1, the vector $\bar{\mathbf{z}}$ is necessarily a unique global minimizer of (26), i.e. any other feasible vector $\mathbf{z} : \mathbf{D}\mathbf{z} = \mathbf{b}$ has a larger value of $\|\mathbf{z}\|_1$; hence, being non-negative, it cannot satisfy $\mathbf{1}^T\mathbf{z} = c$, and therefore it can not be a solution of $\mathbf{D}\mathbf{z} = \mathbf{b}, \quad \mathbf{z} \ge 0$. □

We add the following two comments:

- Assume that a very sparse vector $\bar{\mathbf{z}}$ has been found to be a feasible solution of $\mathbf{D}\mathbf{z} = \mathbf{b}, \quad \mathbf{z} \ge 0$. At least locally, if we aim to find other feasible solutions, we must use a deviation vector that lies in the null-space of $\mathbf{D}$, i.e., $\mathbf{D}\delta = \mathbf{0}$. Positivity of the alternative solution $\bar{\mathbf{z}}+\delta$ forces us to require that at the off-support of $\bar{\mathbf{z}}$, all entries of δ are non-negative. Thus, the above theorem is parallel to the claim that such constrained vector is necessarily the trivial zero one.
- We started the discussion in this section by requiring that $\mathbf{A}$ is non-negative. The only place that used this property is the normalization by $\mathbf{W}$ in Equation (22), requiring that $\mathbf{W}$ is invertible. Thus, any matrix $\mathbf{A}$ that leads to an invertible positive weight matrix $\mathbf{W}$ is adequate for our analysis.

The problem (11) can be rewritten in an equivalent form

$$\min_{\mathbf{z}} \ \sum_i \varphi(|z_i|) \quad \text{subject to} \quad \mathbf{P}\mathbf{D}\mathbf{z} = \mathbf{P}\mathbf{b}, \tag{27}$$

where the "preconditioner" matrix $\mathbf{P}$ is any invertible $n \times n$ matrix. Therefore the statement of the Theorem 1 remains valid if we change $\rho(\mathbf{D})$ and $t_{\mathbf{D}}$ by $\rho(\mathbf{PD})$ and $t_{\mathbf{PD}}$. This gives us a useful degree of freedom in the analysis of problem

[3]Recall that t_D is given by (3) and (6).

(11): In order to relax the requirement on the number of non-zeros, one can try to find such $\mathbf{P}$ that reduces $\rho(\mathbf{PD})$. The same is valid for problem posed in (23).

In the case of positive matrix $\mathbf{D}$, an efficient preconditioning can be obtained just by subtracting the mean of each column $\mathbf{d_i}$ from its entries:

$$\mathbf{P}\mathbf{d}_i = \mathbf{d}_i - \mathrm{mean}(\mathbf{d}_i)\mathbf{1}_n = \left(\mathbf{I} - \frac{1}{n}\mathbf{E}\right)\mathbf{d_i}, \qquad (28)$$

where $\mathbf{E}$ is $n \times n$ matrix of ones and $\mathbf{I}$ is the identity matrix of the same size. Such preconditioner typically reduces correlation between columns. For example, when the columns are normalized, $\sum_k d_{kl} = 1$,

$$(\mathbf{P}\mathbf{d}_i)^T(\mathbf{P}\mathbf{d}_j) = \mathbf{d_i}^T\left(\mathbf{I} - \frac{1}{n}\mathbf{E}\right)\mathbf{d_i} \quad = \quad \mathbf{d}_i^T\mathbf{d}_j - \frac{1}{n}, \quad i,j = 1,...,k.$$

This usually leads to a smaller coherence constant $\rho(\mathbf{PD})$. Note that the matrix $(\mathbf{I} - \frac{1}{n}\mathbf{E})$ is singular (and thus non-invertible); therefore we use $(\mathbf{I} - \frac{1-\epsilon}{n}\mathbf{E})$ instead, with ϵ being a small positive constant $0 < \epsilon << 1$.

We should note again that this preconditioning does not change the solution of the original linear system; it just improves our worst-case forecast of uniqueness versus sparsity. On the other hand, as we shall see in Section 4, it improves significantly the behavior of the OMP [21], [23], targeting the above problem.

IV. Orthogonal Matching Pursuit Performance

A. Approximation Algorithm

We have defined an optimization task of interest, (P_0^+) (2) but this problem is very hard to solve in general. We could replace the ℓ_0-norm by an ℓ_1, and solve a linear programming problem. This is commonly done in a quest for sparse solutions of general linear systems, with good theoretical foundations. In fact, based on Theorem 2, one could replace the ℓ_0 with any other norm, or just solve a non-negative feasibility problem, and still get the same result, if it is indeed sparse enough. However, when this is not the case, we may deviate strongly from the desired solution of (P_0^+).

An alternative to the ℓ_1 measure is a greedy algorithm. We present this option in this section and study is performance, both empirically and theoretically. Specifically, we consider the use of the orthogonal matching pursuit (OMP) algorithm [21], [23], finding the sparsest and non-negative solution to $\mathbf{Ax} = \mathbf{b}$ one atom at a time. As we have shown before, rather than operate on the original linear system, we can use the ℓ_1-normalized version $\mathbf{Dz} = \mathbf{b}$, or even the centered version $\mathbf{PDz} = \mathbf{Pb}$.

The algorithm is is a modified version of the regular OMP that takes into account the non-negativity of the sought solution. The positivity is imposed in two locations: (i) When searching for the next atom to join, we consider only positive inner products; and (ii) When updating the residual based on the accumulated atoms, the least-squares solver should impose a positivity on the coefficients.

An important observation we have already mentioned is that one can apply the above-described OMP in the very same way on the original problem (P_0^+), or a centered version of it. The solutions of both problems are equivalent. How well will this algorithm perform in the two described cases? In the following subsections we offer two kinds of answers - an empirical one and a theoretical one. We start with an empirical evidence.

B. Experimental Results

We start with a (random) non-negative and ℓ^1-normalized dictionary $\mathbf{D}$ of size 100×200. We generate 1000 random sparse representations with varying cardinalities in the range $1 - 40$, and then check the performance of the OMP in recovering them. We test both the regular OMP and the centered version. Matlab's lsqnonneg instruction is used within this algorithm for computing the non-negative least-squares solution for the current set of chosen atoms.

For comparison, we also test the BP, solving the problem $\min \|\mathbf{z}\|_1$ subject to $\mathbf{Dz} = \mathbf{b}$. When there exist only one solution, this method necessarily finds it exactly, thus, expected to outperform the OMP (both its versions). On the other hand, when there are several possible solutions, it is not necessarily finding the sparsest one, thus leading to sometimes to errors. Note that this alternative requires many more computations, as its complexity is much[4] higher. Also, preconditioning of the form discussed here does not affect its solutions.

Figure 1 shows the relative average number of wrong atoms detected in the tested algorithms. Figure 2 shows the average representation error. As can be seen in both graphs, the centered OMP performs much better. We also see, as expected, that BP outperforms both greedy options and yielding very low error rate, with the obvious added cost in complexity. Notice that the BP's representation error is zero simply because the BP always finds a solution to satisfy $\mathbf{Dz} = \mathbf{b}$, whereas the OMP is operated with a fixed (assumed to be known) number of atoms.

C. Theoretical Study

First, let us remind the reader of the two-sided (un-centered) coherence, as defined and used in [7], [17]:

$$\mu(\mathbf{D}) = \max_{1 \le i,j \le k, i \ne j} \frac{\left|\mathbf{d}_i^T \mathbf{d}_j\right|}{\|\mathbf{d}_i\|_2 \cdot \|\mathbf{d}_j\|_2}. \qquad (29)$$

Using this definition, we state the following result.

Theorem 3: For the linear system of equations $\mathbf{Dz} = \mathbf{b}$ (where $\mathbf{D} \in \mathbb{R}^{n \times k}$), if a non-negative solution exists such that

$$\|\mathbf{z}\|_0^0 < \frac{1}{2}\left(1 + \frac{1}{\mu\{\mathbf{D}\}}\right), \qquad (30)$$

then the non-negative OMP is guaranteed to find it exactly.

We omit the proof as it appears in details in [1]. Furthermore, it is a variation on a similar result that appears in [9].

The above theorem is practically useless for handling of (P_0^+) directly, since $\mu\{\mathbf{D}\}$ tends to be too high, implying that

[4] As an example, the Matlab run-time ratio BP-versus-OMP for 1000 examples was found to be roughly $500/\|\mathbf{z}\|_0$.

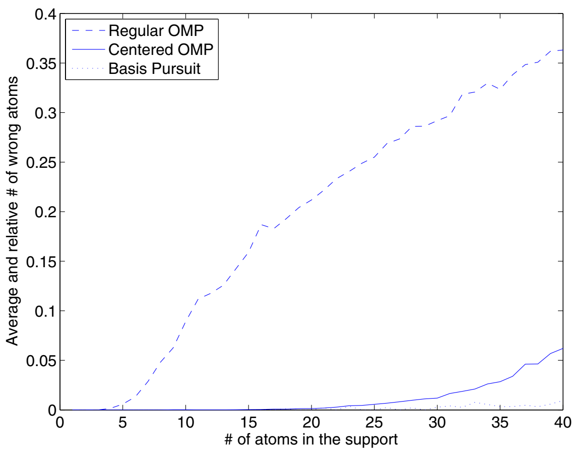

Fig. 1. Performance comparison between regular and centered OMP. This graph shows the relative and average number of wrongly detected atoms as a function of the original cardinality.

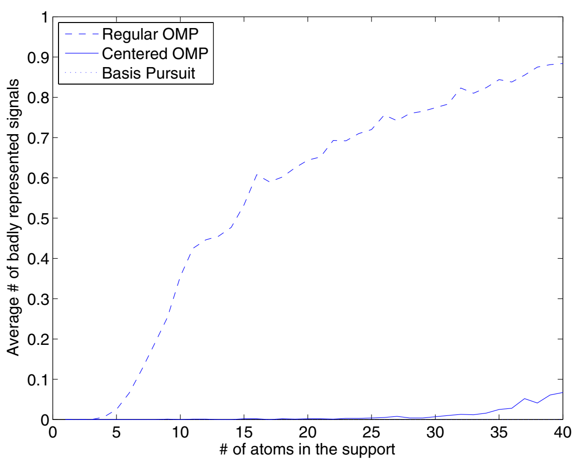

Fig. 2. Performance comparison between regular and centered OMP. This graph shows the average representation error for the test set.

we can handle only empty supports well. This of-course is a meaningless claim. However, when applying the very same analysis to the centered problem, we obtain a parallel result (with the very same proof) of the form

Theorem 4: For the linear system of equations $\mathbf{PDz} = \mathbf{Pb}$ (where $\mathbf{D} \in \mathbb{R}^{n \times k}$, $\mathbf{P} \in \mathbb{R}^{n \times n}$ is invertible,) if a non-negative solution exists such that

$$\|\mathbf{z}\|_0^0 < \frac{1}{2}\left(1 + \frac{1}{\mu\{\mathbf{PD}\}}\right), \tag{31}$$

then OMP is guaranteed to find it exactly.

As we saw in the experimental results, the centered OMP is indeed performing much better, and theoretically we see that this is an expected phenomenon. However, as mentioned in past work on the analysis of pursuit algorithms, we should note that the bounds we provide here are far worse compared to the actual (empirical) performance, as they tend to be over-pessimistic. In the experiments reported in the previous section we have $\mu\{\mathbf{D}\} = 0.858$ and $\mu\{\mathbf{PD}\} = 0.413$, implying that at best one can recover supports of cardinality $T = 1$. Clearly, the OMP succeeds far beyond this point.

V. Conclusions

Non-negative linear systems of equations come up often in many applications in signal and image processing. Solving such systems is usually done by adding conditions such as minimal ℓ_2 length, maximal entropy, maximal sparsity, and so on. In this work we have shown that if a sparse *enough* solution exists, then it is the only one, implying that all the mentioned measures lead to the same solution. We have also proposed an effective preconditioning for improving the chances of such linear system to be handled well by greedy algorithms.

In addition, in this work we have introduced an extended analysis of general BP to the case of arbitrary monotone element-wise concave penalty, and a matrix $\mathbf{A}$ having arbitrary column norms and being preconditioned. The obtained result generalizes the equivalence claims found in [7], [17].

There are several directions in which this work should/could be extended, and several intriguing questions that form the grounds for such extended work. Is the positivity of the matrix $\mathbf{A}$ important for the claims that are developed here? This question requires a closer inspection, as this requirement seems to be redundant. What about optimal choice of the preconditioning? this option seems to encompass stronger potential in the developed bounds. One might follow the algorithm proposed in [15] for designing the optimal $\mathbf{P}$. Also, it should be clear that the positivity requirement may be replaced by any sign-pattern requirement, with the same effect. How such generalization should affect the requirements on $\mathbf{A}$? Finally, there is a clear gap between the proved bound and the empirical behavior as reported in Section 4. Can we strengthen the bounds by relying on probabilistic analysis? All these questions and more promise a fruitful path for more work on this topic.

References

[1] A.M. Bruckstein, M. Elad, and M. Zibulevsky, A Non-negative and sparse enough solution of an underdetermined linear system of equations is unique, submitted to *IEEE Trans. on Information Theory*.

[2] S.S. Chen, D.L. Donoho, and M.A. Saunders, Atomic decomposition by basis pursuit, *SIAM Review*, Vol. 43, pp. 129–159, Jan. 2001.

[3] M. Chu and R. Plemmons, Nonnegative matrix factorization and applications. *IMAGE – Bulletin of the International Linear Algebra Society*, Vol. 34, pp. 2–7, July 2005.

[4] R.R. Coifman and M.V.Wickerhauser,Entropy-based algorithms for best basis selection, *IEEE Trans. on Information Theory*, Vol. 38, pp. 713–718, Mar. 1992.

[5] D. De-Vries and A.J. Berkhout, Velocity analysis based on minimum entropy Geophysics, *Geophysics*, Vol 49, pp. 2132–2142, Dec. 1984.

[6] D. Donoho, On minimum entropy segmentation, In *Wavelets: Theory, Algorithms, and Applications*, edited by C.K.Chui et. al., pp. 233–269, Academic Press, 1994.

[7] D.L. Donoho and M. Elad, Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization, *Proc. Of The National Academy Of Sciences Of The United States Of America*, Vol. 100, pp. 2197–2202, Mar. 2003.

[8] D.L. Donoho and M. Elad, On the stability of the basis pursuit in the presence of noise, *Signal Processing*, Vol. 86, pp. 511–532, Mar. 2006.

[9] D.L. Donoho, M. Elad, and V.N. Temlyakov, Stable recovery of sparse overcomplete representations in the presence of noise, *IEEE Trans. on Information Theory*, Vol. 52 pp. 6–18, Jan. 2006.

[10] D.L. Donoho and X. Huo, Uncertainty principles and ideal atomic decomposition, *IEEE Trans. on Information Theory*, Vol. 47, pp. 2845–2862, Nov. 2001.

[11] D.L. Donoho and J. Tanner, Sparse nonnegative solutions of underdetermined linear equations by linear programming, *Proc. of the National Academy of Sciences*, Vol. 102, No. 27, pp. 9446–9451, 2005.

[12] D.L. Donoho and J. Tanner, Thresholds for the recovery of sparse solutions via L1 minimization, *Proc. of the Conference on Information Sciences and Systems*, March 2006.

[13] D.L. Donoho and J. Tanner, Neighborliness of randomly-projected simplices in high dimensions, *Proc. of the National Academy of Sciences*, Vol. 102(27), 9452–9457, 2005.

[14] D.L. Donoho, and P.B. Stark, Uncertainty principles and signal recovery, *SIAM J. Applied Math.*, Vol. 49/3, pp. 906–931, June 1989.

[15] M. Elad, Optimized projections for compressed-sensing, to appear in *IEEE Trans. on Signal Processing*.

[16] M. Elad and A.M. Bruckstein, A generalized uncertainty principle and sparse representation in pairs of bases,*IEEE Trans. on Information Theory*, Vol. 48 pp. 2558–2567, Sep. 2002.

[17] R. Gribonval and M. Nielsen, Sparse representations in unions of bases, *IEEE Trans. on Information Theory*, Vol. 49, pp. 3320–3325, Dec. 2003.

[18] R. Gribonval and M. Nielsen, On the strong uniqueness of highly sparse representations from redundant dictionaries, *ICA 2004: Lecture Notes In Computer Science*, Vol. 3195, pp. 201–208, 2004.

[19] T. Hazan, S. Polak and A. Shashua, Sparse image coding using a 3d non-negative tensor factorization, *International Conference on Computer Vision (ICCV)*, Beijing, China, October, 2005.

[20] K. Kreutz-Delgado, B. D. Rao, K. Engan, T.W. Lee and T. Sejnowski, Convex/Schur-convex (csc) log-priors and sparse coding, Proceedings of the *Sixth Joint Symposium on Neural Computation*, Pasadena, CA, My 1999.

[21] S. Mallat, and Z. Zhang, Matching pursuit with time-frequency dictionaries, *IEEE Trans. Signal Processing*, Vol. 41, pp. 3397–3415, Dec. 1993.

[22] M. Nafie, M. Ali, and A.H. Tewfik, Optimal subset selection for adaptive signal representation, Proceedings of the *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1996.

[23] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition, in Conf. Rec. 27th Asilomar Conf. Signals, Syst. Comput., Vol. 1, 1993.

[24] P. Pauca, J. Piper, and R. Plemmons, Nonnegative matrix factorization for spectral data analysis, *Linear Algebra and Applications*, 2006.

[25] S. Raboi, *On the recovery of missing samples via sparsity conditions*, M.Sc. Thesis, The Electrical Engineering department, the Technion - Israel Institute of Technology, February 2007.

[26] B.D. Rao and K. Kreutz-Delgado, Deriving algorithms for computing sparse solutions to linear inverse problems, *Asilomar Conference on Signals, Systems and Computers*, Vol. 1, pp. 955–959, Monterey, CA, November 1997.

[27] A. Shashua and T. Hazan, Non-negative tensor factorization with applications to statistics and computer vision, *International Conference on Machine Learning (ICML)*, August 2005.

[28] A. Shashua, R. Zass and T. Hazan, Multi-way clustering using super-symmetric non-negative tensor factorization, Proc. of the *European Conference on Computer Vision (ECCV)*, May 2006, Graz, Austria.

[29] J.A. Tropp, Greed is good: Algorithmic results for sparse approximation, *IEEE Trans. on Information Theory*, Vol. 50, pp. 2231–2242, Oct. 2004.

[30] J.A. Tropp, Just relax: Convex programming methods for identifying sparse signals in noise, *IEEE Trans. on Information Theory*, Vol. 52, pp. 1030–1051, Mar. 2006.

[31] R. Zass and A. Shashua, Nonnegative sparse PCA, *Neural Information and Processing Systems (NIPS)*, December, 2006.