

FAST NON-NEGATIVE ORTHOGONAL LEAST SQUARES

Mehrdad Yaghoobi and Mike E. Davies

Institute for Digital Communications, the University of Edinburgh, EH9 3JL, UK
 {m.yaghoobi-vaighan, d.wu and mike.davies}@ed.ac.uk

ABSTRACT

An important class of sparse signals is the non-negative sparse signals. While numerous greedy techniques have been introduced for low-complexity sparse approximations, there are few non-negative versions. Among such a large class of greedy techniques, one successful method, which is called the Orthogonal Least Squares (OLS) algorithm, is based on the maximum residual energy reduction at each iteration. However, the basic implementation of the OLS is computationally slow. The OLS algorithm has a fast implementation based on the QR matrix factorisation of the dictionary. The extension of such technique to the non-negative domain is possible. In this paper, we present a fast implementation of the non-negative OLS (NNOLS). The computational complexity of the algorithm is compared with the basic implementation, where the new method is faster with two orders of magnitude. We also show that, if the basic implementation of NNOLS is not computationally feasible for moderate size problems, the proposed method is tractable. We also show that the proposed algorithm is even faster than an *approximate* implementation of the non-negative OLS algorithm.

Index Terms— Non-negative sparse approximations, Orthogonal Least Squares, Efficient Implementations, Non-negative Orthogonal Least Squares and QR Matrix Factorization

1. INTRODUCTION

The signal processing with a low-dimensional model, is a modern technique to explore signal structures. We thus need some low-dimensional signal approximation algorithms for this purpose. Such methods have to be fast, if the dimension of data is large/huge. There are many general purpose sparse approximation techniques with fast implementations. In some applications like hyperspectral data decomposition [1] and DNA microarrays analysis [2], we further know that non-zero coefficients of the sparse approximation are positive. In this setting, non-negative sparse approximation techniques [3], often provide better solutions than the universal techniques. Universal greedy techniques can be adapted to such a setting by confining the solutions to the set of non-negative coefficients. However, extension of the fast implementations of the greedy methods to the non-negative setting

is often not straightforward. An example of such extensions to a greedy method, with a complicated coefficient update *e.g.* Orthogonal Matching Pursuit (OMP), has already been presented in [4]. We here present a fast implementation of the non-negative OLS [5] technique.

Let $\mathbf{y} \in \mathbb{R}^m$ be a signal with the sparse representation using a normalised linear generative model $\Phi \in \mathbb{R}^{m \times n}$, *i.e.* dictionary, where $\mathbf{y} = \Phi \mathbf{x}$ and \mathbf{x} is a sparse vector. At each iteration, OLS selects an element of dictionary, called an atom, which maximally reduces the residual energy of the signal, after adding the atom to the set of previously selected atoms. It then finds the best possible representation of \mathbf{y} using selected atoms, *i.e.* orthogonal projection. In the other words, the algorithm has these steps: i) selecting an atom, *i.e.*

$$i^* = \operatorname{argmin}_i \min_{s_{k+1} := s_k \cup \{i\}} \|\mathbf{y} - \Phi_{s_{k+1}} \Phi_{s_{k+1}}^\dagger \mathbf{y}\|_2, \quad (1)$$

where s_k is the index set of selected atoms up to the iteration k , $\Phi_{s_{k+1}}$ is the subdictionary selected by the index set s_k and \dagger represents the More-Penrose pseudoinverse and ii) using the Least Squares (LS) representation within the selected atoms. As the optimisation problem (1) is computationally expensive, Chen *et al.* [5] presented a technique, which iteratively updates the QR factorisation of the selected subdictionary, *i.e.* similar to the fast QR factorisation based OMP implementation [6]. It has an extra step in which we orthogonalise the rest of the dictionary with respect to the selected subdictionary, and then normalise the new dictionary. We refer the readers to [5], for a more detailed description of the OLS algorithm.

If \mathbf{y} is “non-negative sparse”, *i.e.* the representation \mathbf{x} has the property $\|\mathbf{x}\|_0 \leq k$, $\mathbf{x} \in \mathbb{R}_+^n$, we intend to incorporate the extra information and adapt OLS to the new setting. This adaptation can be done if we combine the two steps of the OLS and enforce the non-negativity of the coefficients, *i.e.*

$$i^* = \operatorname{argmin}_i \min_{\substack{s_{k+1} := s_k \cup \{i\} \\ \alpha \in \mathbb{R}_+^{k+1}}} \|\mathbf{y} - \Phi_{s_{k+1}} \alpha\|_2. \quad (2)$$

We call a direct optimisation of (2), which would be very costly, a Basic NNOLS (BNNOLS) implementation. The extension of the fast implementation of OLS for the non-negative representation is not straightforward, if we want to *exactly* solve NNOLS.

Algorithm 1 Suboptimal Non-Negative OLS

```

1: initialisation:  $s = \emptyset, k = 0, \mathbf{x} = \mathbf{0}, \mathbf{\Omega} = \bar{\Phi}$  and  $\mathbf{r}_0 = \mathbf{y}$ 
2: while  $k < K$  &  $\max(\mathbf{\Omega}^T \mathbf{r}_k) > 0$  do
3:    $i^* \leftarrow \max(\mathbf{\Omega}^T \mathbf{r}_k)$ 
4:    $s \leftarrow s \cup i^*$ 
5:    $\mathbf{x}_s \leftarrow \operatorname{argmin}_{\theta > 0} \|\mathbf{y} - \bar{\Phi}_s \theta\|_2$ 
6:    $\mathbf{r}_{k+1} \leftarrow \mathbf{y} - \bar{\Phi}_s \mathbf{x}_s$ 
7:   Updating  $\mathbf{\Omega} = [\omega_i], \omega_i, \forall i \in s^c$ , from (3)
8:    $k \leftarrow k + 1$ 
9: end while
10:  $\mathbf{x}|_s \leftarrow \mathbf{x}_s$ 

```

We here present a simple technique to extend the fast OLS implementation in section 2, without solving an expensive optimisation problem in the loop. We then show that the new implementation is significantly faster than a basic implementation of NNOLS in section 3.

2. NON-NEGATIVE OLS

Let $\bar{\Phi}_k \in \mathbb{R}^{m \times k}$ be the k selected atoms of $\bar{\Phi}$ at iteration k , in which the columns are ordered based on the iteration number, *i.e.* the j th column is the j th selected atom. We can now generate $\bar{\Phi}_k \in \mathbb{R}^{m \times (n-k)}$ as the rest of the dictionary, *i.e.* including all non-selected atoms. In each iteration of OLS, we need to calculate a matrix $\mathbf{\Omega} = [\omega_i]_{1 \leq i \leq n-k}$, which is the column-normalised version of the orthogonal part of $\bar{\Phi}_k$ to the column span of $\bar{\Phi}_k$. In other words,

$$\omega_i := \frac{\bar{\phi}_i^\perp}{\|\bar{\phi}_i^\perp\|_2}, \quad (3)$$

where $\bar{\phi}_i^\perp := (\mathbf{I} - \bar{\Phi}_k \bar{\Phi}_k^\dagger) \bar{\phi}_i$. We recall that $\bar{\phi}_i$'s with the property $\bar{\phi}_i^\perp = \mathbf{0}$, will be eliminated, as they lie in the span of $\bar{\Phi}_k$ and there is no need to consider them in the following iterations. Let $\mathbf{r}_k = \mathbf{y} - \bar{\Phi}_k \mathbf{x}_k$ be the residual error of the representation at iteration k , and \mathbf{x}_k is the coefficient vector at this iteration. As $\mathbf{\Omega}$ is orthogonal to the span of $\bar{\Phi}_k$ and normalised, we can find the *direction with the maximum residual energy reduction* by simply finding the maximum of $|\mathbf{\Omega}^T \mathbf{r}_k|$, where $|\cdot|$ is the element by element absolute value operator. The optimal direction corresponding to ω_{i^*} provides us the index of an atom in this iteration. If we now consider the non-negative representation, one *suboptimal* approach is to: i) select an atom with positive contribution, *i.e.* $\max_i(\omega_i^T \mathbf{r}_k) > 0$, ii) updating the coefficients using NNLS. A pseudo code for this algorithm, which is called Suboptimal NNOLS (SNNOLS) is presented in Algorithm 1. s^c is the complement set of s . SNNOLS method does not find the solution of iterating (2). The reason is related to the fact that the selection criteria does not consider the fact that representation based on the selected atoms in $\bar{\Phi}_k$ have to be positive. In the case which we enforce an already selected atom to have a zero coefficient value, we may have some suboptimality. We also

Algorithm 2 Fast Non-Negative Orthogonal Least Squares

```

1: initialisation:  $s = z_0 = \emptyset, k = 0, \mathbf{\Omega} = \bar{\Phi}$  and  $\mathbf{r}_0 = \mathbf{y}$ 
2: while  $k < K$  &  $\max(\mathbf{\Omega}^T \mathbf{r}_k) > 0$  do
3:    $(\zeta, \iota) \leftarrow \operatorname{sort}_\downarrow(\mathbf{\Omega}^T \mathbf{r}_k)$ 
4:    $p \leftarrow 1$ 
5:    $p^c \leftarrow 1$ 
6:    $z^c \leftarrow 0$ 
7:   while  $\sim$  Terminate &  $p < N$  do
8:      $z^t$  from (6)
9:      $z \leftarrow \psi_{\iota(p)}^T \mathbf{r}_k : \psi_{\iota(p)} = \bar{\Phi}_k^\perp / \|\bar{\Phi}_k^\perp\|_2,$   

        $\bar{\Phi}_k^\perp = (\mathbf{I} - \bar{\Psi}_k \bar{\Psi}_k^T) \bar{\phi}_{\iota(p)}$ 
10:    Update based on Table 1
11:    end while
12:     $s \leftarrow s \cup \iota(p)$ 
13:    Update  $\bar{\Psi}, \bar{\Phi}, \mathbf{\Omega}$  and  $\mathbf{R}^{-1}$ 
14:     $\mathbf{z}_{k+1} \leftarrow [\mathbf{z}_k; z_{k+1}]$ 
15:     $\mathbf{r}_{k+1} \leftarrow \mathbf{r}_k - z_{k+1} \psi_{k+1}$ 
16:     $k \leftarrow k + 1$ 
17: end while
18: output:  $\mathbf{x}|_s \leftarrow \mathbf{R}^{-1} \mathbf{z}_K$ 

```

If	Then
$0 < z \leq z^t, z > z^c$	$z_{k+1} \leftarrow z$, Terminate
$0 < z \leq z^t, z \leq z^c$	$z_{k+1} \leftarrow z^c, p \leftarrow p^c$, Terminate
$z > z^c \geq z^t$	$p \leftarrow p + 1$
$z^c \geq z > z^t$	$z_{k+1} \leftarrow z^c, p \leftarrow p^c$, Terminate
$z > z^t > z^c$	$z^c \leftarrow z^t, p^c \leftarrow p, p \leftarrow p + 1$
$z < 0$	Terminate

Table 1. Decision rules to guarantee positivity of the coefficients.

need to consider the computational cost of SNNOLS, as it iteratively needs to solve NNLS. The solution which we propose here is to combine the selection and coefficient update steps to avoid such suboptimality and reduce the computational cost by not directly solving NNLS.

2.1. Fast Non-Negative OLS

If $\bar{\Psi}_k \in \mathbb{R}^{m \times k}$ is iteratively generated using the selected atoms of $\mathbf{\Omega}$, up to the k th iteration, we can QR factorise $\bar{\Phi}_k$ by $\bar{\Psi}_k \mathbf{R}_k$, where $\bar{\Psi}_k$ is orthonormal and $\mathbf{R}_k \in \mathbb{R}^{k \times k}$ is an upper-diagonal matrix. The reason for such a structure for \mathbf{R} is the fact that $\bar{\Phi}_j$ is orthogonal to ψ_p 's, for $\forall p : j < p$. This QR factorisation is updating throughout the iterations as follows: i) adding ω_{i^*} as the new column of $\bar{\Psi}$, *i.e.*

$$\bar{\Psi}_{k+1} = [\bar{\Psi}_k \ \omega_{i^*}] \quad (4)$$

and ii) updating \mathbf{R}_k using the following equation,

$$\mathbf{R}_{k+1} = \begin{bmatrix} \mathbf{R}_k & \nu \\ \mathbf{0} & \mu \end{bmatrix}, \quad (5)$$

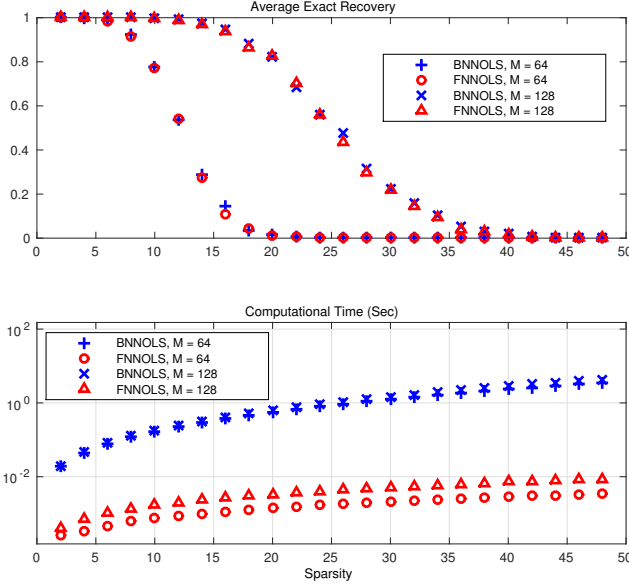


Fig. 1. Exact recovery and computation time. $N = 256$, M is 64 or 128 and sparsity K is changing. The proposed technique FNNOLS is roughly 100 times faster than BNNOLS!

where $\nu = \Psi_k^T \bar{\phi}_{i^*}$, $\mu = \|\bar{\phi}_{i^*}^\perp\|_2$ and $\bar{\phi}_{i^*}^\perp$ can be easily calculated using $\bar{\phi}_{i^*}^\perp = (\mathbf{I} - \Psi_k \Psi_k^T) \bar{\phi}_{i^*}$, *i.e.* no need to calculate the pseudoinverse of Φ_k , comparing to the definition after (3). This process can be interpreted as a type of Gram-Schmidt (GS) orthogonalisation process, which is a well-known QR factorisation technique [7]. With the new setting, we can rewrite the OLS algorithm as follows: In the first iteration, we do not need any orthogonalisation and we have $\phi_1 = \psi_1$ and $\mathbf{R}_1 = [1]$. In the k th iteration, $k \geq 1$, let the best approximation of \mathbf{y} , using Φ_k , be $\sum_{i=1}^k x_i \phi_i = \sum_{i=1}^k z_i \psi_i$. In the $k+1$ iteration, we have,

$$\begin{aligned} \sum_{i=1}^{k+1} z_i \psi_i &= \sum_{i=1}^k z_i \psi_i + z_{k+1} \psi_{k+1} \\ &= \sum_{i=1}^k x_i \phi_i + z_{k+1} \psi_{k+1}. \end{aligned}$$

We recall that ψ_i 's are orthogonal and choosing a new elementary function ψ_{k+1} does not change the optimal value of the best representation z_i , $\forall i : i < k+1$, calculated up to the previous iteration k . This is obviously not true for x_i 's as Φ_{k+1} is not orthogonal.

As ψ_{k+1} lives in the span of the non-redundant set $\{\phi_j\}_{j \in [1, k+1]}$, $\psi_{k+1} = \sum_{j=1}^{k+1} \gamma_j \phi_j$ for some unique γ_j 's.




Fig. 2. Computation time for the fixed $N = 256$ and $K = 24$ and 32.

We can then have,

$$\begin{aligned} \sum_{i=1}^{k+1} z_i \psi_i &= \sum_{i=1}^k x_i \phi_i + \sum_{j=1}^{k+1} z_{k+1} \gamma_j \phi_j \\ &= \sum_{i=1}^k (x_i + z_{k+1} \gamma_i) \phi_i + z_{k+1} \gamma_{k+1} \phi_{k+1}. \end{aligned}$$

As $z_{k+1} \gamma_{k+1}$ is always positive, we only need to assure that $x_i + z_{k+1} \gamma_i \geq 0$ or

$$z_{k+1} \leq z^t := \begin{cases} \min_{i, \gamma_i < 0} \left| \frac{x_i}{\gamma_i} \right| & \exists i, \gamma_i < 0 \\ +\infty & \text{Otherwise.} \end{cases} \quad (6)$$

In OLS, we only need to keep z_i 's in the memory and calculate x_i 's at the end of algorithm [8]. To assure that x_i 's are all non-negative, z_i 's should comply the condition of (6). We then choose the atom that the corresponding z_{k+1} , or shrunk by the upper-bound of (6), has the largest value. We therefore need to track the record of the best possible value for z , if $z = \omega_{i^*}^T \mathbf{r}_k$ does not comply with (6). If $(\zeta, \iota) = \text{sort}_\downarrow(\Omega^T \mathbf{r}_k)$, where sort_\downarrow is the sorting operator in a descent order, ζ is the sorted coefficients with corresponding indices ι . A pseudocode of the fast non-negative orthogonal least square is presented in Algorithm 2. In this algorithm, z is the current candidate, starting with $z = \zeta(p)$, $p = 1$, in an internal loop in the k th iteration, we make the decision based on the rules of Table 1, to guarantee that the coefficients x_i 's will remain non-negative throughout iterations. We also need to calculate \mathbf{R}^{-1} at the last iteration, which would be expensive. We thus can iteratively generate \mathbf{R}^{-1} using the following update rule,

$$\mathbf{R}_{k+1}^{-1} = \begin{bmatrix} \mathbf{R}_k^{-1} & -\frac{\mathbf{R}_k^{-1} \nu}{\mu} \\ \mathbf{0} & \frac{1}{\mu} \end{bmatrix}. \quad (7)$$

where ν and μ were defined after (5). For the fast implementation of this algorithm, we have to also efficiently calculate

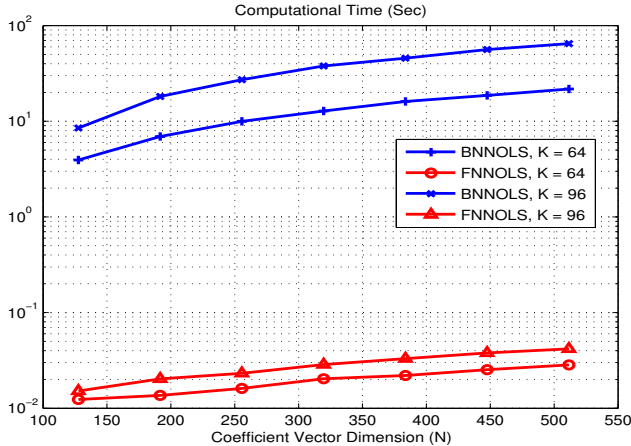


Fig. 3. Computation time for the fixed $M = 128$ and $K = 64$ and 96 .




Fig. 4. Exact recovery and computation time. Comparison between SNNOLS and FNNOLS. $N = 1024$, $M = 512$ sparsity K is changing.

γ . It is not difficult to show that γ is the last column of \mathbf{R}_{k+1}^{-1} , i.e., $\gamma = [-\frac{\mathbf{R}_{k+1}^{-1}\nu}{\mu}; \frac{1}{\mu}]$, which is kept in the memory.

3. SIMULATIONS

We investigate the runtime of the algorithms on a single core of an Intel core 2.6 GHz processor. In the first experiment, we randomly generated Φ with 256 atoms and 64 or 128 rows using *i.i.d.* Gaussian random variables, followed by column normalisation. \mathbf{y} was generated using a Gaussian-Bernoulli model, i.e. uniformly random support and Normal distribution for the non-zero coefficients. We repeated the simulations 100 times, using proposed FNNOLS and BNNOLS and different sparsity levels. The correct recovery of the support and computational time are shown in Figure 1. While the exact recoveries are very similar, proposed method is signif-

icantly faster. The difference is more obvious for large k 's, where we can get a computation gain of 100.

In the second experiment, we used the same method to generate the dictionary, but we fixed $N = 256$, to investigate the computation time as a function of M . Here, $K = 24$ or 32 and M is between 32 and 196 . The computational time is plotted in Figure 2. The computational time is slowly increasing with M . We can observe that the difference in computations is not significant for different K 's using the fast technique, while it is noticeable using the basic implementation.

We then fixed $M = 128$ and changed N from 128 to 512 and similarly repeated each simulation 100 times. The sparsity was set to be 64 or 96 , and the results are shown in Figure 3. The computational cost is slowly increasing *w.r.t.* the parameter N in FNNOLS, while it is increasing faster for the BNNOLS. However, the difference between the computational time of the fast and basic implementations are considerable, where we yield a gain of 1000 times acceleration, using the FNNOLS algorithm!

We also presented an approximate NNOLS technique, called the SNNOLS, which can be seen as an intuitive extension of the fast OLS implementation to the non-negative domain. As we can see, the computational time of SNNOLS is also low, which means that we can compare it with FNNOLS in a larger dimension setting. We then chose $M = 512$, $N = 1024$ and changed the sparsity level K from 32 to 80 and averaged the simulations, with the random settings which explained before, over 500 trials. We have plotted the exact support recovery (top panel) and the algorithms run-time (bottom panel) in Figure 4. While the exact recoveries are roughly the same, the computational times are slightly different. For the example, the computational time of SNNOLS is 25% higher than FNNOLS for $K = 80$. The similarity between the exact recovery of two algorithms is an interesting fact, which needs more investigation.

4. CONCLUSION

A significantly faster implementation for the non-negative orthogonal least square algorithm was proposed here. The new implementation is based on the recursive generation of a QR factorisation of the dictionary and it allows us to use FNNOLS in a broader application. We also presented an algorithm which approximately solve the OLS problem in a non-negative setting, which is called SNNOLS. This algorithm performs similarly in the exact recovery of the sparsity support, when the dictionaries are randomly generated, but it is slightly slower than FNNOLS. The performance of the proposed techniques, i.e. FNNOLS and SNNOLS, have to be investigated, using other types of dictionaries. The computational complexity analysis of the FNNOLS is also an interesting subject, which we left for the future work.

Acknowledgment

The authors acknowledge the support of EPSRC and the MOD University Defence Research Collaboration in Signal Processing, through the grant number EP/K014277/1.

REFERENCES

- [1] Y Qian, S Jia, J Zhou, and A Robles-Kelly, “Hyperspectral unmixing via sparsity-constrained nonnegative matrix factorization,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 49, no. 11, pp. 4282–4297, 2011.
- [2] F. Parvaresh, H. Vikalo, S. Misra, and B. Hassibi, “Recovering sparse signals using sparse measurement matrices in compressed DNA microarrays,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 3, pp. 275–285, 2008.
- [3] PO Hoyer, “Non-negative sparse coding,” in *Neural Networks for Signal Processing, 2002. Proceedings of the 2002 12th IEEE Workshop on*. 2002, pp. 557–565, IEEE.
- [4] M. Yaghoobi, D. Wu, and M.E. Davies, “Fast non-negative orthogonal matching pursuit,” *IEEE Signal Processing Letters*, vol. 22, no. 9, 2015.
- [5] S. Chen, S.A. Billings, and W. Luo, “Orthogonal least squares methods and their application to non-linear system identification,” *International Journal of Control*, vol. 50, no. 5, pp. 1873–1896, 1989.
- [6] B.K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM Journal of Comput.*, vol. 24, no. 2, pp. 227–234, 1995.
- [7] G.H. Golub and C.F. Van Loan, *Matrix computations*, Johns Hopkins University Press Baltimore, 1996.
- [8] T Blumensath and ME Davies, “On the difference between orthogonal matching pursuit and orthogonal least squares,” Tech. Rep., 2007.