



HAL
open science

Projected Block Coordinate Descent for sparse spike estimation.

Pierre-Jean B enard, Yann Traonmilin, Jean Fran ois Aujol

► **To cite this version:**

Pierre-Jean B enard, Yann Traonmilin, Jean Fran ois Aujol. Projected Block Coordinate Descent for sparse spike estimation.. 2024. hal-04462779

HAL Id: hal-04462779

<https://hal.science/hal-04462779v1>

Preprint submitted on 16 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destin ee au d ep ot et  a la diffusion de documents scientifiques de niveau recherche, publi es ou non,  emanant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv es.

Projected Block Coordinate Descent for sparse spike estimation

Pierre-Jean B enard^{1,*}, Yann Traonmilin¹, Jean-Fran ois Aujol¹

¹Univ. Bordeaux, Bordeaux INP, CNRS, IMB, UMR 5251, F-33400 Talence, France

*contact: pierre-jean.benard@math.u-bordeaux.fr

Abstract—We consider the problem of recovering off-the-grid spikes from linear measurements. The state of the art Over-Parametrized Continuous Orthogonal Matching Pursuit (OP-COMP) with Projected Gradient Descent (PGD) successfully recovers those signals. In most cases, the main computational cost lies in a unique global descent on all parameters (positions and amplitudes). In this paper, we propose to improve this algorithm by accelerating this descent step. We introduce a new algorithm, based on Block Coordinate Descent, that takes advantages of the sparse structure of the problem. Based on qualitative theoretical results, this algorithm shows improvement in calculation times in realistic synthetic microscopy experiments.

Index Terms—spike super-resolution, non-convex optimization, over-parametrization, block-coordinate descent

I. INTRODUCTION

The off-the-grid super-resolution problem consists in recovering spikes from linear measurements in a continuous setting. It has applications in many fields. Among these applications, one can cite Single Molecule Localization Microscopy (SMLM) or Magnetic Resonance Imaging (MRI). In these fields, the signals to recover display properties of sparsity that can be exploited for a faster recovery.

Let x_0 be an off-the-grid sparse signal over \mathbb{R}^d . Such signals can be modeled as a sum of K Dirac measures:

$$x_0 = \sum_{i=1}^K a_i \delta_{t_i} \quad (1)$$

where $a = (a_1, \dots, a_K) \in \mathbb{R}^K$ are the amplitudes and $t = (t_1, \dots, t_K) \in \mathbb{R}^{K \times d}$ are the positions of the spikes. We observe the signal through a linear operator A from the space \mathcal{M} of finite signed measures over \mathbb{R}^d to \mathbb{C}^m with m the number of measurements. The result of this observation is $y \in \mathbb{C}^m$ which can be written as $y = Ax_0$. Note that we consider the noiseless case in this article for the sake of clarity.

A way to recover the true signal x_0 from its observation is to find the minimizer of a non-convex least-square problem:

$$x^* \in \arg \min_{x \in \Sigma_{K,\epsilon}} \|Ax - y\|_2^2 \quad (2)$$

where $\Sigma_{K,\epsilon}$, defined in (5), is a set modeling a separation constraint between spikes. We consider the set of parameters $\Theta_{K,\epsilon}$,

$$\Theta_{K,\epsilon} = \left\{ \theta = (a_1, \dots, a_K, t_1, \dots, t_K) \in \mathbb{R}^{K(d+1)}, \right. \\ \left. \forall i, j \in \{1, \dots, K\}, i \neq j, \|t_i - t_j\|_2 > \epsilon \right\}. \quad (3)$$

We can rewrite (1) with the variable θ :

$$x_0 = \phi(\theta) := \sum_{i=1}^K a_i \delta_{t_i} \in \Sigma_{K,\epsilon}. \quad (4)$$

The unknown x_0 belongs to the low dimensional model

$$\Sigma_{K,\epsilon} = \phi(\Theta_{K,\epsilon}) \quad (5)$$

The problem (2) can then be written as

$$\theta^* \in \arg \min_{\theta \in \Theta_{K,\epsilon}} g(\theta) \quad \text{with} \quad g(\theta) := \|A\phi(\theta) - y\|_2^2. \quad (6)$$

Theoretical guarantees for the recovery of x_0 with (2) have been given in [13]. As an example, when the linear operator A models random Fourier measurements at frequencies drawn with a Gaussian distribution and $\dim(y) \gtrsim \mathcal{O}(K^2 d \text{polylog}(K, d))$, we have $x^* = x_0$. In practice, Sliding Continuous Orthogonal Matching Pursuit (SCOMP) [14] has been successful at minimizing (2). Recovery guarantees are given in [11], [12], [5]. We can also cite the Sliding Frank-Wolfe (SFW) algorithm [10], that aims to solve a regularized version of (2). Although state-of-the-art methods can fully recover signals in this setting, they lack of performance in computation time. The Over-Parametrized COMP + Projected Gradient Descent (OP-COMP + PGD) algorithm has been introduced to address these performance issues [27], [4]. This method initializes an over-parametrized signal without a costly descent step (OP-COMP). Then, it performs a single descent on all parameters while performing a projection to reduce the number of spikes (PGD). It has been proven that a good initialization of the signal is possible in K steps [5]. This means that each initialized spike is close to a true spike, on the conditions that one has a sufficient separation ϵ and a well-behaved operator A . However, these conditions are not always verified. In these cases, while OP-COMP + PGD is still faster than SFW and SCOMP, it lacks in efficiency during the descent on all parameters in θ . The large number of gradients to compute is the main limiting factor of OP-COMP + PGD [5].

For gradients descent methods, an idea to accelerate computation time is the Block Coordinate descent (BCD). It consists in only performing iterations of the descent on a subset of the coordinates of the parameters that we estimate.

The idea of an underlying structure to separate the signal in blocks is common in many applications (see e.g. [25], [29]). Among many applications, there are Non-Negative Matrix Factorization [31], Sparse Logistic Regression [30].

The LASSO problem has also been solved with BCD methods [15], [29].

The definition of Block Coordinate Descent relies on the blocking strategy, the block selection rule and the block update rule [19].

Blocking strategy. The blocks can be either fixed beforehand or variable and recomputed at each iteration. Fixed blocks permit better control and separation in sub-problems [24] when the structure of the problem is known before computations [16]. Variable blocks allow more freedom, e.g. by creating blocks based on greedy methods [28] such as Gauss-Southwell.

Block selection rule. As the choice of the block to update may lead to different minima, it is a main area of focus for BCD methods [25]. Basic ideas such as the cyclic variants for index choice are common as they are deterministic methods with proven convergence [3]. Implementation-wise, they are simple and do not require computation overhead. Instead of going through all coordinates in a cyclic order, one can use randomized methods by choosing coordinates based on a uniform distribution [18], weighted distribution [1] or just from arbitrary rules [23]. Greedy methods are also used in the block selection process. They choose the most suitable block at each iteration by solving an optimization problem [19]. A well known greedy method, for its efficiency both theoretically and practically and its simplicity, is the Gauss-Southwell (GS) rule which select a block based on the norm of its gradient,

$$b_k \in \arg \max_{b \in \mathcal{B}} \|\nabla g(\theta^k)\|_b \quad (7)$$

where \mathcal{B} is the set of all blocks. Other greedy methods such as Maximum Improvement [9] or GS-Lipschitz [20] are also valid choices for block selection rules. Greedy methods perform better than cyclic or random methods but at the cost of a more expensive computation [20].

Block update rule. The role of this rule is to make the chosen block maximally decrease the objective function (like (2)). It is important to use a block update rule suited with the problem. Starting with the classical gradient descent, one can set the step-size in different ways (fixed, approximate line-search or optimal step-size) to improve performance [6]. Many more descent methods used in convex and non-convex optimization are used in BCD. Among the many methods introduced, one can use proximal update [2], stochastic gradient [17] depending, for example, on the geometry of the objective function. In our case, we wish to improve upon the Projected Gradient Descent from OP-COMP + PGD using a FISTA Restart method for the block updating rule. An early use of FISTA for BCD method can be found in [22]. While there are no theoretical guarantees for the convergence of FISTA in BCD methods, it works well in various cases [8].

Contributions. In this paper, our main contribution is an accelerated method for recovering sparse signals based on OP-COMP + PGD. We propose a Projected Block Coordinate Descent method that performs the descent step on spikes that have not already converged. The main ingredient for this BCD is a block selection strategy based on the specific

structure of the sparse spike model to decompose the problem in small blocks. We define the three ingredients defining BCD within the OP-COMP + PGD method and justify qualitatively our choices. We provide synthetic experiments in the context of microscopy that show an improvement up to 35% in calculation time compared to PGD.

II. THE PROJECTED BLOCK COORDINATE DESCENT ALGORITHM

In this section, we motivate the idea of using a block system to divide our problem in a sum of smaller and simpler problems. Moreover, we describe our method: the projected Block Coordinate Descent. This descent only operates on parameters of non-converged spikes of the estimated signal.

A. A result to guide block selection

The Block Coordinate Descent (BCD) is a way to fasten the computation time of descent algorithms by using the underlying structure of the problem. In our case, the more spikes there are in the signal to recover, the more parameters there are to compute in the descent process. As our model contains a separation constraint, every spike to be recovered is estimated by one or more initialized spikes independently of the other true spikes. This leads to two cases: either there is one initialized spike estimating a true spike, or a cluster of initialized spikes. For one spike estimations, the convergence to their associated true spike is fast. On the contrary for the clusters of initialized spikes, their convergence is much slower toward a true spike [5]. We propose to only compute the gradients of spikes that have not yet converged. Next, we show qualitatively why we can use our model to separate the recovery of each spike using our separation constraint with an idealized initialized signal.

Let $x_0 = \sum_{i=1}^K a_i \delta_{t_i} \in \Sigma_{K,\epsilon}$ the signal to recover and $x = \sum_{j=1}^K b_j \delta_{s_j}$ the initialized signal obtained with COMP. Suppose that x is well initialized i.e. for all $i \in \{1, \dots, K\}$, $\|t_i - s_i\|_2 < \frac{\epsilon}{3}$. By computing their difference, we get

$$x_0 - x = \sum_{i=1}^K (a_i \delta_{t_i} - b_i \delta_{s_i}) = \sum_{i=1}^K \nu_i \quad (8)$$

where $\nu_i = a_i \delta_{t_i} - b_i \delta_{s_i}$ is called a $\frac{\epsilon}{3}$ -dipole. From [5], we recall the definition of a $\frac{\epsilon}{3}$ -dipole,

Definition 1 ((ϵ) -Dipole). *An ϵ -dipole is a measure $\nu = a\delta_t - b\delta_s$ where $\|t - s\|_2 \leq \epsilon$.*

Let us also recall the definition of $\frac{\epsilon}{3}$ -separated dipoles,

Definition 2 ((ϵ) -Separation of dipoles). *Two dipoles $\nu_1 = a_1\delta_{t_1} - b_1\delta_{s_1}$ and $\nu_2 = a_2\delta_{t_2} - b_2\delta_{s_2}$ are ϵ -separated if their support is strictly ϵ -separated, i.e. if $\|t_1 - t_2\|_2 > \epsilon$ and $\|t_1 - s_2\|_2 > \epsilon$ and $\|s_1 - t_2\|_2 > \epsilon$ and $\|s_1 - s_2\|_2 > \epsilon$.*

The following lemma gives us bounds to better interpret the block structure of our model with our initialization x .

Lemma 1. *Let $x_0 = \sum_{i=1}^K a_i \delta_{t_i} \in \Sigma_{K,\epsilon}$ and $x = \phi(\theta) = \sum_{i=1}^K b_i \delta_{s_i}$ with $\theta = (b_1, \dots, b_K, s_1, \dots, s_K)$ such that for all*

$i \in \{1, \dots, K\}$, $\|t_i - s_i\|_2 < \frac{\epsilon}{3}$. Let A be a linear operator from the space of finite measures over \mathbb{R}^d to \mathbb{C}^m . Suppose that for two $\frac{\epsilon}{3}$ -separated $\frac{\epsilon}{3}$ dipoles ν_1, ν_2 , there exists $\mu_A > 0$ such that A follows the assumption, □

$$\langle A\nu_1, A\nu_2 \rangle \leq \mu_A \|A\nu_1\|_2^2 \|A\nu_2\|_2^2. \quad (9)$$

Then we obtain

$$\|y - Ax\|_2^2 \leq \sum_{i=1}^K \|A\nu_i\|_2^2 + \mu_A \sum_{\substack{i,j=1 \\ i \neq j}}^K \|A\nu_i\|_2^2 \|A\nu_j\|_2^2. \quad (10)$$

Moreover, let $\partial_{i,r}g(\theta)$ be the partial derivative of $g(\theta)$ with respect to the r -th coordinate of s_i . Then we get

$$\left| \partial_{i,r}\|y - Ax\|_2^2 - \partial_{i,r}\|A\nu_i\|_2^2 \right| \leq 2a_i(K-1)\mu_A \|\partial_{i,r}A\nu_i\|_2^2 \max_{j \in \{1, \dots, K\}} \|A\nu_j\|_2^2. \quad (11)$$

Proof. Since we suppose that for all $i \in \{1, \dots, K\}$, $\|t_i - s_i\|_2 < \frac{\epsilon}{3}$, we can write $x_0 - x = \sum_{i=1}^K (a_i \delta_{t_i} - b_i \delta_{s_i})$ as a sum of $\frac{\epsilon}{3}$ -dipoles $\sum_{i=1}^K \nu_i$. To get (10), we first start by expanding $\|y - Ax\|_2^2$,

$$\|y - Ax\|_2^2 = \|A(x_0 - x)\|_2^2 = \left\| A \sum_{i=1}^K \nu_i \right\|_2^2 \quad (12)$$

$$= \sum_{i=1}^K \|A\nu_i\|_2^2 + \sum_{\substack{i,j=1 \\ i \neq j}}^K \langle A\nu_i, A\nu_j \rangle \quad (13)$$

We use our assumption (9) so we can bound $\langle A\nu_i, A\nu_j \rangle$,

$$\|y - Ax\|_2^2 \leq \sum_{i=1}^K \|A\nu_i\|_2^2 + \mu_A \sum_{\substack{i,j=1 \\ i \neq j}}^K \|A\nu_i\|_2^2 \|A\nu_j\|_2^2. \quad (14)$$

Now, we prove (11). Using Proposition 2.1 in [26], we have

$$\partial_{i,r}\|y - Ax\|_2^2 = 2a_i \mathcal{R}e \left\langle A\partial_{i,r}\delta_{s_i}, A \sum_{j=1}^K \nu_j \right\rangle \quad (15)$$

$$= 2a_i \mathcal{R}e \langle A\partial_{i,r}\delta_{s_i}, A\nu_i \rangle + 2a_i \mathcal{R}e \left\langle A\partial_{i,r}\delta_{s_i}, A \sum_{\substack{j=1 \\ j \neq i}}^K \nu_j \right\rangle. \quad (16)$$

Using our assumption (9), we get (using the fact that this identity is also verified for Dirac derivatives that are generalized dipoles [26]):

$$\begin{aligned} \partial_{i,r}\|y - Ax\|_2^2 &\leq 2a_i \mathcal{R}e \langle A\partial_{i,r}\delta_{s_i}, A\nu_i \rangle \\ &\quad + 2a_i \mu_A \|\partial_{i,r}A\delta_{s_i}\|_2^2 \sum_{\substack{j=1 \\ j \neq i}}^K \|A\nu_j\|_2^2. \end{aligned} \quad (17)$$

We bound each $\|A\nu_j\|_2^2$ by their maximum $\max_{j \in \{1, \dots, K\}} \|A\nu_j\|_2^2$. Moreover, we have that $2a_i \mathcal{R}e \langle A\partial_{i,r}\delta_{s_i}, A\nu_i \rangle = \partial_{i,r}\|A\nu_i\|_2^2$. Finally, we obtain

$$\left| \partial_{i,r}\|y - Ax\|_2^2 - \partial_{i,r}\|A\nu_i\|_2^2 \right| \leq 2a_i \mu_A (K-1) \|\partial_{i,r}A\nu_i\|_2^2 \max_{j \in \{1, \dots, K\}} \|A\nu_j\|_2^2. \quad (18)$$

Using the result (10) from Lemma 1, we get that in order to minimize $\|y - Ax\|_2^2$, a sufficient condition is to minimize individually each $\|A\nu_i\|_2^2$, i.e. recover K distinct 1-spike signals.

To distinguish converged spikes from non-converged spikes, we will compare their gradients using the result (11) from Lemma 1. Indeed, the absolute difference between $\partial_{i,r}\|y - Ax\|_2^2$ and $\partial_{i,r}\|A\nu_i\|_2^2$ is small: For common observation operators (e.g. Gaussian filter), we have that the mutual coherence parameter μ_A is small and the terms $\|\partial_{i,r}A\nu_i\|_2^2$ and $\sup_{j \in \{1, \dots, K\}} \|A\nu_j\|_2^2$ are bounded. In practice, the over-parametrization of the initialized signal yields more spikes than dipoles, however the simpler case given by Lemma 1 suggests that blocks of spikes with large gradients should be considered for the selection strategy.

B. Description of the algorithm

Our new algorithm: the Projected Block Coordinate Descent (Algorithm 1) is an improvement of the Projected Gradient Descent [27], which performs a descent on all parameters while enforcing a separation constraint on the positions of the estimated spikes. We keep this main feature and we integrate our Block Coordinate Descent acceleration. Our method is derived from the Gauss-Southwell rule. We choose our blocks from the set of all blocks $\mathcal{B}_{\text{super-res}}$, where each block represents a spike. In each iteration of the main loop, we compute the norm of the gradients ($\|\nabla[g(\theta)]_i\|_2$ for all $i \in \{1, \dots, K\}$) associated to each spike of the estimated signal with $\nabla[g(\theta)]_i := \nabla g((a_i, t_i))$ i.e. the gradients for the amplitude and position of the spike i . Then, we create two sets I and J . The set $I = \{i, \|\nabla[g(\theta)]_i\|_2 \geq \text{threshold} \times \max_j \|\nabla[g(\theta)]_j\|_2\}$ represents all spikes that have not yet converged. In practice, we set manually the value of `threshold` and we discuss this choice in more details in the next section. We set $J = \{1, \dots, K\} \setminus I$. We note $x_I = \sum_{i \in I} a_i \delta_{s_i}$ and $x_J = \sum_{j \in J} a_j \delta_{s_j}$. With only the selected spikes from I , we perform a projected gradient descent on x_I for a fixed number of iteration m_{it} (or until convergence). Note that the observation in this partial projected descent y_I is $y - Ax_J$ i.e. the original observation updated with the non-selected spikes from this iteration. If a projection occurs during this descent, we stop the projected gradient descent. We merge both signals x_I and x_J and we compute again I and J . We continue this process until convergence of our estimated signal x . To perform the descent, we choose the FISTA Restart algorithm [21] as in practice it outperforms a classic Gradient Descent and it provides fast convergence.

III. EXPERIMENTS

In this section, we compare our new Projected Block Coordinate Descent with the state-of-the-art Projected Gradient Descent. The signal to recover in every experiment is composed of 50 Dirac measures for MA-TIRF in 3 dimensions [10]. As these methods aim to recover a signal from an initial estimation, we use Over-Parametrized Continuous Orthogonal

Algorithm 1 Projected Block Coordinate Descent.

```

procedure BCD( $y, A, \theta^{(0)}, n_{it}, m_{it}, \text{threshold}$ )
   $n \leftarrow 0$ 
   $k \leftarrow \# \text{spikes}$ 
  while  $n < n_{it}$  do
     $I = \{i, \|\nabla[g(\theta^{(n)})]_i\|_2 \geq$ 
       $\text{threshold} \times \max_j \|\nabla[g(\theta^{(n)})]_j\|_2\}$ 
     $J = \{1, \dots, k\} \setminus I$ 
     $\tilde{\theta}^{(n)} \leftarrow \{\theta_i^{(n)}, i \in I\}, \bar{\theta}^{(n)} \leftarrow \{\theta_j^{(n)}, j \in J\}$ 
     $\tilde{y} \leftarrow y - A\phi(\bar{\theta}^{(n)})$ 
     $\tilde{\theta}^{(n^*)} \leftarrow \text{projected\_descent}(\tilde{y}, \tilde{\theta}^{(n)}, m_{it})$ 
     $\theta^{(n^*)} \leftarrow \{\tilde{\theta}_i^{(n^*)} \text{ if } i \in I \text{ else } \bar{\theta}_i^{(n)}\}_{i=1, \dots, k}$ 
     $\theta^{(n)} \leftarrow \theta^{(n^*)}$ 
     $k \leftarrow \# \text{spikes}$ 
  end while
  return  $\theta^{(n_{it})}$ 
end procedure

```

Matching Pursuit (OP-COMP) for the initialization of the signal.

The linear operator used for the observations is the MA-TIRF operator [10]. It is used in microscopy for recovering cells in 3D images. The positions t are in a domain of size $6.4\text{mm} \times 6.4\text{mm} \times 0.8\text{mm}$. The observation is composed of 4 planes which represent different focal angles. Each plane is composed of 64×64 pixel. In total, $\dim(y) = 64 \times 64 \times 4 = 16384$. The amplitudes follow a uniform distribution $U([1, 2])$. The code for these experiments is available for download at [7].

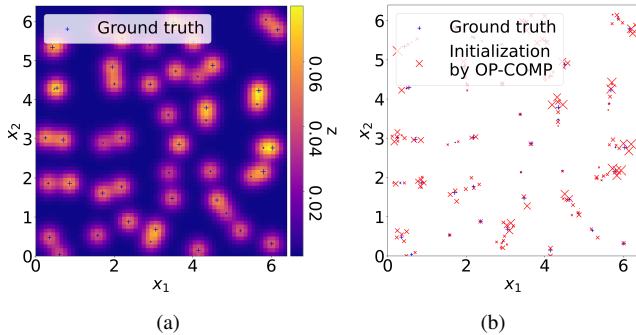


Figure 1. (a) Ground truth and a plane of its observation through MA-TIRF, (b) Ground truth and the initialized signal by OP-COMP.

In Fig. 1(a), we display a focal plane of the observation. It corresponds to an image taken at a certain angle. Though it is in 2 dimensions, the combination of multiple angles is useful to recover spikes in 3D. We see in this figure that most of the spikes are well separated. In Fig. 1(b), we have the true signal and the initialized signal by OP-COMP. Note that every spike from the ground truth has at least one initialized spike close to itself. Moreover, the over-parametrization is clearly visible as the number of initialized spikes is larger than the number of spikes to recover.

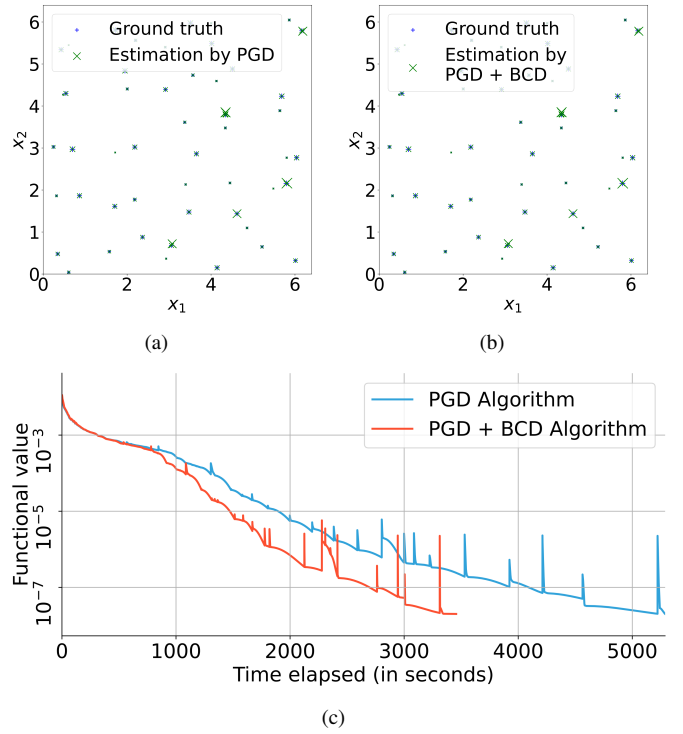


Figure 2. (a) Ground truth and the estimated signal by PGD with MA-TIRF, (b) Ground truth and the estimated signal by BCD with MA-TIRF, (c) Comparison norms of residues by PGD and BCD in function of time.

In Fig. 2(a), we observe that the spikes are well recovered by PGD as expected. In Fig. 2(c) (in blue), we show the norm of the residue $\|y - Ax\|_2^2$ through the iterations for both PGD (in blue) and BCD (in red). Note that we stop the descent when the norm of the residue gets below 2×10^{-8} . For PGD, the norm of the residue decreases and reaches its minimum at 2×10^{-8} . The sudden jumps in the norm of the residue are due to the projection used.

In Fig. 2(b), we obtain an estimated signal by BCD similar to the one obtained with PGD. In Fig. 2(c) (in red), as for PGD, we observe the norm of the residue through the iterations of BCD. We note that for the same norm of the residue, BCD converges much faster with an acceleration of almost 35% in computation time.

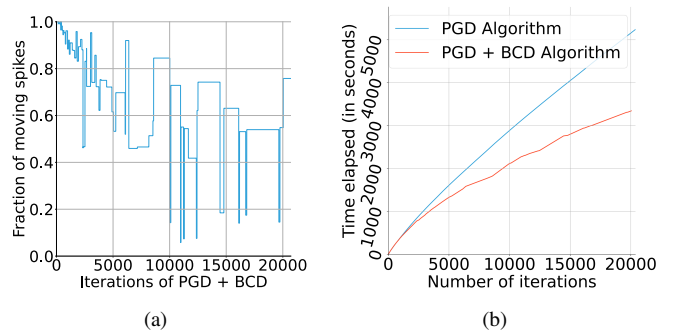


Figure 3. (a) Fraction of spikes whose gradients are computed during BCD, (b) Time elapsed for PGD and BCD in function of the number of iterations

While the start of both methods takes the same time for about 2000 iterations, we observe a clear acceleration as the time passes. Indeed, as seen in Figure 3(a), the fraction of gradients computed at each iteration of BCD decreases. A good empirical value for the choice of `threshold` in Algorithm 1 is 10^{-3} . This value is used for all of our experiments. It provides a good trade-off in terms of minimization of g (see Equation (6)) and computation time. Indeed, this threshold limits the number of blocks used for each `projected_descent` by selecting only non-converged spikes. At the start of the descent, almost all gradients are still computed at each iteration, but as the iterations pass, the fraction of computed gradients decreases, hence the computation time gained over PGD. Both Figures 3(a) and 3(b) reveal that by computing fewer gradients, we can gain time on the total computation.

After being tested on 10 similar signals composed of 50 spikes with a high separation between spikes, the results stay the same and we have that in general, BCD is 35% faster than PGD for MA-TIRF. Other metrics used to compare both estimations are the Jaccard, Precision and Recall. For these metrics, there is no particular difference between the two methods PGD and BCD.

We note that for other kind of linear operators (such as the Fourier Transform), performances of the BCD can vary, but it always improves the computation speed.

IV. DISCUSSION/CONCLUSION

To conclude, we observe that the block structure for the sparse spike resolution permits to improve computation time of descent algorithms. Indeed, in a realistic microscopy setting, the time gained by BCD is up to 35% when compared to previous methods. For future works, we intend to try other greedy methods for a better block selection. Moreover, one can parallelize the descent on the different blocks as it is a major benefit of Block Coordinate Descent methods.

V. ACKNOWLEDGMENTS

This work was supported by the French National Research Agency (ANR) under reference ANR-20-CE40-0001 (EF-FIREG project). Experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d'Aquitaine (see <https://www.plafrim.fr>).

REFERENCES

- [1] Z. Allen-Zhu, Z. Qu, P. Richtárik, and Y. Yuan. Even faster accelerated coordinate descent using non-uniform sampling. In *ICML*, volume 48, pages 1110–1119. PMLR, 2016.
- [2] A. Auslender. Asymptotic properties of the fenchel dual functional and applications to decomposition problems. *Journal of Optim. Theory and Applications*, 73:427–449, 1992.
- [3] A. Beck and L. Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optim.*, 23(4):2037–2060, 2013.
- [4] P.-J. Bénard, Y. Traonmilin, and J.-F. Aujol. Fast off-the-grid sparse recovery with over-parametrized projected gradient descent. In *2022 30th EUSIPCO*, pages 2206–2210. IEEE, 2022.
- [5] P.-J. Bénard, Y. Traonmilin, J.-F. Aujol, and E. Soubies. Estimation of off-the-grid sparse spikes with over-parametrized projected gradient descent: theory and application. working paper or preprint, Sept. 2023.
- [6] D. P. Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3):334–334, 1997.
- [7] P.-J. Bénard, Y. Traonmilin, and J.-F. Aujol. Code of the experiments. https://github.com/pjbenard/opCOMP_PGD_BCD, 2024. [Online].
- [8] A. Chambolle and C. H. Dossal. On the convergence of the iterates of "fista". *Journal of Optim. Theory and Applications*, 166(3):25, 2015.
- [9] B. Chen, S. He, Z. Li, and S. Zhang. Maximum block improvement and polynomial optim. *SIAM Journal on Optim.*, 22(1):87–107, 2012.
- [10] Q. Denoyelle, V. Duval, G. Peyré, and E. Soubies. The Sliding Frank-Wolfe Algorithm and its Application to Super-Resolution Microscopy. *Inverse Problems*, 2019.
- [11] C. Elvira, R. Gribonval, C. Soussen, and C. Herzet. OMP and continuous dictionaries: Is k-step recovery possible? In *ICASSP*, pages 5546–5550. IEEE, 2019.
- [12] C. Elvira, R. Gribonval, C. Soussen, and C. Herzet. When does OMP achieve exact recovery with continuous dictionaries? *Applied and Comp. Harmonic Analysis*, 51:39, 2021.
- [13] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin. Compressive Statistical Learning with Random Feature Moments. *Mathematical Statistics and Learning*, 3(2):113–164, 2021.
- [14] N. Keriven, N. Tremblay, Y. Traonmilin, and R. Gribonval. Compressive k-means. In *ICASSP*, pages 6369–6373. IEEE, 2017.
- [15] Y. Li and S. Osher. Coordinate descent optimization for l1 minimization with application to compressed sensing; a greedy algorithm. *Inverse Problems and Imaging*, 3(3):487–503, 2009.
- [16] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 70(1):53–71, 2008.
- [17] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optim.*, 19(4):1574–1609, 2009.
- [18] Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optim.*, 22(2):341–362, 2012.
- [19] J. Nutini, I. Laradji, and M. Schmidt. Let's make block coordinate descent converge faster: Faster greedy rules, message-passing, active-set complexity, and superlinear convergence. *Journal of Machine Learning Research*, 23(131):1–74, 2022.
- [20] J. Nutini, M. Schmidt, I. Laradji, M. Friedlander, and H. Koepke. Coordinate descent converges faster with the gauss-southwell rule than random selection. In *ICML*, pages 1632–1641. PMLR, 2015.
- [21] B. O'donoghue and E. Candes. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15:715–732, 2015.
- [22] Z. Qin, K. Scheinberg, and D. Goldfarb. Efficient block-coordinate descent algorithms for the group lasso. *Mathematical Programming Computation*, 5(2):143–169, 2013.
- [23] P. Richtárik and M. Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144(1-2):1–38, 2014.
- [24] S. Sardy, A. G. Bruce, and P. Tseng. Block coordinate relaxation methods for nonparametric wavelet denoising. *Journal of computational and graphical statistics*, 9(2):361–379, 2000.
- [25] H.-J. M. Shi, S. Tu, Y. Xu, and W. Yin. A primer on coordinate descent algorithms. *arXiv preprint arXiv:1610.00040*, 2016.
- [26] Y. Traonmilin and J.-F. Aujol. The basins of attraction of the global minimizers of the non-convex sparse spike estimation problem. *Inverse Problems*, 2020.
- [27] Y. Traonmilin, J.-F. Aujol, and A. Leclaire. Projected gradient descent for non-convex sparse spike estimation. *IEEE Signal Processing Letters*, 27:1110–1114, 2020.
- [28] P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117:387–423, 2009.
- [29] S. J. Wright. Coordinate descent algorithms. *Mathematical programming*, 151(1):3–34, 2015.
- [30] T. T. Wu and K. Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, 2(1):224, 2008.
- [31] Y. Xu and W. Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.