



HAL
open science

Greedy methods, randomization approaches and multi-arm bandit algorithms for efficient sparsity-constrained optimization

A Rakotomamonjy, S Koço, L Ralaivola

► **To cite this version:**

A Rakotomamonjy, S Koço, L Ralaivola. Greedy methods, randomization approaches and multi-arm bandit algorithms for efficient sparsity-constrained optimization. 2015. hal-01182801v1

HAL Id: hal-01182801

<https://hal.science/hal-01182801v1>

Preprint submitted on 3 Aug 2015 (v1), last revised 18 Aug 2016 (v2)

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 destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Greedy methods, randomization approaches and multi-arm bandit algorithms for efficient sparsity-constrained optimization

A. Rakotomamonjy, S. Koço, L. Ralaivola

Abstract

Several sparsity-constrained algorithms such as Orthogonal Matching Pursuit or the Frank-Wolfe algorithm with sparsity constraints work by iteratively selecting a novel atom to add to the current non-zero set of variables. This selection step is usually performed by computing the gradient and then by looking for the gradient component with maximal absolute entry. This step can be computationally expensive especially for large-scale and high-dimensional data. In this work, we aim at accelerating these sparsity-constrained optimization algorithms by exploiting the key observation that, for these algorithms to work, one only needs the coordinate of the gradient's top entry. Hence, we introduce algorithms based on greedy methods and randomization approaches that aim at cheaply estimating the gradient and its top entry. Another of our contribution is to cast the problem of finding the best gradient entry as a best arm identification in a multi-armed bandit problem. Owing to this novel insight, we are able to provide a bandit-based algorithm that directly estimates the top entry in a very efficient way. Theoretical results stating that the resulting inexact Frank-Wolfe or Orthogonal Matching Pursuit algorithms act, with high probability, similarly to their exact versions are also given. We have carried out several experiments showing that the greedy deterministic and the bandit approaches we propose can achieve an acceleration of an order of magnitude while being as efficient as the exact gradient when used in algorithms such as OMP, Frank-Wolfe or CosAMP.

AR is with University of Rouen, LITIS Lab. Most of this work has been carried while he was visiting LIF at Aix-Marseille University. Email alain.rakoto@insa-rouen.fr

SK is with University of Rouen, LITIS LAB. Email: sokol.koco@univ-rouen.fr

LR is with Aix-Marseille University, LIF. Email: liva.ralaivola@lif.univ-mrs.fr

This work is supported by Agence Nationale de la Recherche, project GRETA 12-BS02-004-01.

Manuscript received August2015; revised XXX.

Index Terms

Sparsity, Orthogonal Matching Pursuit, Frank-Wolfe algorithm, Greedy methods, Best arm identification.

I. INTRODUCTION

Over the last decade, there has been a large interest in inference problems featuring data of very high-dimension and a small number of observations. Such problems occur in a wide variety of application domains ranging from computational biology and text mining to information retrieval and finance. In order to learn from these datasets, statistical models are frequently designed so as to feature some sparsity properties. Hence, fueled by the large interest induced by these application domains, an important amount of research work in the machine learning, statistics and signal processing communities, has been devoted to sparse learning and as such, many algorithms have been developed for yielding models that use only few dimensions of the data. To obtain these models, one typically needs to solve a problem of the form

$$\min_{\mathbf{w}} L(\mathbf{w}) \text{ subject to } \|\mathbf{w}\|_0 \leq K \quad (1)$$

where $L(\mathbf{w})$ is a smooth convex objective function that measures a goodness of fit of the model, $\|\mathbf{w}\|_0$ is the ℓ_0 pseudo-norm that counts the number of non-zero components of the vector \mathbf{w} and K is a parameter that controls the sparsity level.

One usual approach that has been widely considered is the use of a convex and continuous surrogate of the ℓ_0 pseudo-norm, namely an ℓ_1 -norm. The resulting problem is the well-known Lasso problem [28] and a large variety of algorithms for its resolution exist, ranging from homotopy methods [8] to the Frank-Wolfe (FW) algorithm [10]. In the same flavour, non-convex continuous penalties are also common solutions for relaxing the ℓ_0 pseudo-norm [23, 14]. Another possible approach is to consider greedy methods that provide local optimal solution to the problem (1). Once again, in this last context, a flurry of algorithms have been proposed, the most popular ones being the *Matching Pursuit* (MP) and the *Orthogonal Matching Pursuit* (OMP) algorithms [16, 21, 17]. One common point of the aforementioned algorithms for solving problem (1) is that they require, at each iteration, the computation of the objective function's gradient. For large-scale and high-dimensional setting, computing the gradient at each iteration may be very time-consuming.

Stochastic gradient descent (SGD) algorithms are now classical methods for avoiding the computation of the full gradient in large-scale learning problems [32, 24]. Most of these works have been devoted to smooth composite optimization although some efforts addressing ℓ_1 -regularized problems exist [26].

Recently, these SGD algorithms have been further accelerated through the introduction of variance reduction methods for the gradient estimation [12]. In the context of problem (1) where a non-smooth non-convex constraint appears, very few works have envisaged the use of stochastic optimization. Nguyen et al. [19] have proposed stochastic versions of a gradient pursuit algorithm. Following a similar direction, by exploiting inexact gradient information, we address the problem of accelerating sparsity-inducing algorithms such as Matching Pursuit, Orthogonal Matching Pursuit and the Frank-Wolfe algorithm with an ℓ_1 ball constraint. However, unlike stochastic gradient approaches, the acceleration we propose leverages on the fact that at each iteration, these algorithms seek for the gradient's component that has the largest (absolute) value.

Hence, our main contribution in this paper is to propose novel algorithms that allow to efficiently find this top entry of the gradient. By doing so, our objective is to design novel efficient versions of MP, OMP, FW algorithms while keeping intact all the properties of these algorithms for sparse approximation. Indeed, the latter becomes possible if at each iteration of MP, OMP or FW, our inexact gradient-based estimation of the component with largest value is the same as the one obtained with the exact gradient.

We propose two approaches, the first one based on a greedy deterministic algorithm and the second one based on a randomized method, whose aim is to build an inexact estimation of the gradient so that its top entry in the same as the exact one. Next, by casting the problem as a best arm identification multi-armed bandit problem [6], we are able to derive an algorithm that directly estimates the best component of the gradient. Interestingly, these algorithms are supported by theoretical evidences that with high probability, they are able to retrieve the correct component of the gradient. As a consequence, we show that MP, OMP or FW algorithms employing these approaches for spotting the correct component of the gradient behave as their exact counter part with high probability.

The remainder of the paper is organized as follows. Section II presents sparsity-constrained algorithms, formalizes the problem and introduces the key observation on which we build our acceleration strategy. Section III provides the different algorithms we propose for efficiently estimating the extreme gradient component we look for. A discussion with related works is given in Section IV. Experimental results are depicted in Section V while Section VI concludes the work and presents different outlooks.

II. SPARSE LEARNING ALGORITHM WITH EXTREME GRADIENT COMPONENT

In this section, we introduce the sparse learning problem we are interested in, as well as some algorithms that are frequently used for sparse learning. We then point out the prominent common trait of those algorithms, namely *the extreme gradient component* property, and discuss how its estimation can be

employed for accelerating some classes of sparse learning algorithms.

A. Framework

Consider the problem where we want to estimate a relation between a set of n samples gathered in a vector $\mathbf{y} \in \mathbb{R}^n$ and the matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$. In a sparse signal approximation problem, \mathbf{X} would be a matrix whose columns are the elements of a dictionary and \mathbf{y} the target signal, while in a machine learning problem, the i -th row of matrix \mathbf{X} , formalized as \mathbf{x}_i^\top , $\mathbf{x}_i \in \mathbb{R}^d$, depicts the features of the i -th example and y_i is the label or target associated to that example. In the sequel, we denote as $x_{i,j}$ the entry of \mathbf{X} at the i -th row and j -th column.

Our objective is to learn the relation between \mathbf{y} and \mathbf{X} through a linear model of the data denoted $\mathbf{X}\mathbf{w}$ by looking for the vector \mathbf{w} that solves problem (1) when the objective function is of the form

$$L(\mathbf{w}) = \sum_i^n \ell(y_i, g(\mathbf{w}^\top \mathbf{x}_i)),$$

where ℓ is an individual loss function that measures the discrepancy between a true value y and its estimation and $g(\cdot)$ is a given differentiable function that depicts the (potential) non-linear dependence of the loss to \mathbf{w} . Typically, ℓ might be the least-square error function, which leads to $L(\mathbf{w}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2$ or the logistic loss and we have $L(\mathbf{w}) = \sum_{i=1}^n \log(1 + \exp\{-y_i \mathbf{x}_i^\top \mathbf{w}\})$. In the sequel, we present two algorithms that solve problem (1) by a greedy method and by a continuous and a convex relaxation of the ℓ_0 pseudo-norm, respectively.

B. Algorithms

1) *Gradient Pursuit*: This algorithm is a generalization of the greedy algorithm known as *Orthogonal Matching Pursuit* (OMP) to generic losses functions [5]. It can be directly applied to our problem given in Equation (1). Similarly to OMP, the gradient pursuit algorithm is a greedy iterative procedure, which, at each iteration, selects the largest absolute coordinate of the loss gradient. This coordinate is added to the set of already selected elements and the new iterate is obtained as the solution of the original optimization problem restricted to these selected elements while keeping the other ones at 0. The detailed procedure is given in Algorithm 1.

While conceptually simple, this algorithm comes with theoretical guarantees on their ability to recover the exact underlying sparsity pattern of the model, for different types of loss functions [21, 1, 15, 27].

Several variations of this algorithm have been proposed ranging from methods exploring new pursuits directions instead of the gradient [4], to methods making only slight changes to the original one that

Algorithm 1 Gradient Pursuit Algorithm

```

1: set  $k = 0$ , initialize  $\mathbf{w}_0$ 
2: for  $k=0,1, \dots$  do
3:    $i^* = \arg \max_i |\nabla L(\mathbf{w}_k)|_i$ 
4:    $\Gamma_k = \Gamma_k \cup \{i^*\}$ 
5:    $\mathbf{w}_{k+1} = \arg \min L(\mathbf{w})$  over  $\Gamma_k$  with  $\mathbf{w}_{\Gamma_k^c} = \mathbf{0}$ 
6: end for

```

Algorithm 2 Frank-Wolfe Algorithm

```

1: set  $k = 0$ , initialize  $\mathbf{w}_0$ 
2: for  $k=0,1, \dots$  do
3:    $\mathbf{s}_k = \arg \min_{\mathbf{s} \in C} \mathbf{s}^\top \nabla L(\mathbf{w}_k)$ 
4:    $\mathbf{d}_k = \mathbf{s}_k - \mathbf{w}_k$ 
5:   set, linesearch or optimize  $\gamma_k \in [0; 1]$ 
6:    $\mathbf{w}_{k+1} = \mathbf{w}_k + \gamma_k \mathbf{d}_k$ 
7: end for

```

have strongly impacted the ability of the algorithm to recover signals. For instance, the *CoSaMP* [18] and GraSP [3] algorithms select the top $2K$ entries in the absolute gradient, K being the desired sparsity pattern, optimize over these entries and the already selected K ones, and prune the resulting estimate to be K -sparse. In addition to being efficient, these algorithms output sparse vectors whose distances from the true sparse optimum are bounded.

2) *Frank-Wolfe algorithm*: The Frank-Wolfe (FW) algorithm aims at solving the following problem

$$\min_{\mathbf{w} \in C} L(\mathbf{w}) \quad (2)$$

with $\mathbf{w} \in \mathbb{R}^d$, L a convex and differentiable function with Lipschitz gradient and C a compact subset of \mathbb{R}^d . The FW algorithm, given in Algorithm 2, is a straightforward procedure that solves problem 2 by first iteratively looking for a search direction and then updating the current iterate. The search direction \mathbf{s}_k is obtained from the following convex optimization problem (line 3 Alg 2)

$$\mathbf{s}_k = \arg \min_{\mathbf{s} \in C} \mathbf{s}^\top \nabla L(\mathbf{w}_k) \quad (3)$$

which may be efficiently solved, depending on the constraint set C . For achieving sparsity, we typically choose C as a ℓ_1 -norm ball, e.g. $C = \{\mathbf{w} : \|\mathbf{w}\|_1 \leq 1\}$, which turns (3) into a linear program. Despite

its simplistic nature, the FW algorithm has been shown to be linearly convergent [9, 10]. Interestingly, it can also be shown that convergence is preserved as long as the following condition holds

$$\mathbf{s}_k^\top \nabla L(\mathbf{w}_k) \leq \min_{\mathbf{s} \in C} \mathbf{s}^\top \nabla L(\mathbf{w}_k) + \epsilon, \quad (4)$$

where ϵ depends on the smoothness of L and the step-size γ_k . In general cases where the minimization problem (3) is expensive to solve, this condition suggests that an approximate solution \mathbf{s}_k may be sufficient provided that the true gradient is available. Similarly, if an inexact gradient information $\hat{\nabla}L(\mathbf{w}_k)$ is available, convergence is still guaranteed under the condition that

$$\mathbf{s}_k^\top \hat{\nabla}L(\mathbf{w}_k) \leq \min_{\mathbf{s} \in C} \mathbf{s}^\top \nabla L(\mathbf{w}_k) + \epsilon$$

and some conditions relating $\hat{\nabla}L(\mathbf{w}_k)$ and $\nabla L(\mathbf{w}_k)$. For instance, if C is a unit-norm ball associated with some norm $\|\cdot\|$ and \mathbf{s}_k a minimizer of $\min_{\mathbf{s} \in C} \mathbf{s}^\top \hat{\nabla}L(\mathbf{w}_k)$, i.e.

$$\mathbf{s}_k = \arg \min_{\mathbf{s} \in C} \mathbf{s}^\top \hat{\nabla}L(\mathbf{w}_k)$$

then in order to ensure convergence, it is sufficient to have \mathbf{s}_k so that [10]

$$\|\hat{\nabla}L(\mathbf{w}_k) - \nabla L(\mathbf{w}_k)\|_\star \leq \epsilon, \quad (5)$$

where $\|\cdot\|_\star$ is the conjugate norm associated with $\|\cdot\|$.

Interestingly, the above equation provides a guarantee on the convergence on an inexact gradient Frank-Wolfe algorithm, but unfortunately, the precision needed on the inexact gradient is given with respect to the exact one. However, while condition ((5)) is unusable in practice, it conveys the important idea that inexact gradient computation may be sufficient for solving sparsity-constrained optimization problems.

C. Leveraging from the extreme gradient component estimation

As stated above, the gradient pursuit algorithm needs to solve at each iteration the following problem

$$i^\star = \arg \max_j |\nabla L(\mathbf{w}_k)|_j,$$

i.e., the goal is to find at each iteration the gradient's component with the largest absolute value.

In the Frank-Wolfe algorithm, similar situations where finding \mathbf{s}_k corresponds to look for an extreme component of the (absolute) gradient, occur. For instance, when the constraint set is the ℓ_1 -norm ball ($C_1 = \{\mathbf{w} : \|\mathbf{w}\|_1 \leq 1\}$) or the positive simplex constraint $C_2 = \{\mathbf{w} : \|\mathbf{w}\|_1 = 1, w_i \geq 0\}$, it can be easily shown that the search directions are respectively given by

$$\arg \min_{\mathbf{s} \in C_1} \mathbf{s}^\top \nabla L(\mathbf{w}_k) = \arg \max_i |\nabla L(\mathbf{w}_k)|_i.$$

and

$$\arg \min_{\mathbf{s} \in C_2} \mathbf{s}^\top \nabla L(\mathbf{w}_k) = \arg \min_i |\nabla L(\mathbf{w}_k)|_i.$$

Hence, as in Matching Pursuit, OMP or the gradient pursuit algorithm, for specific choices of C , we are not exactly interested in the full gradient, nor in its extreme values, but only on the coordinate of the gradient component with the smallest, the largest (absolute) value.

Our objective in this paper is to leverage on these observations for accelerating sparsity-inducing algorithms which look for one extreme component of the gradient. In particular, we are interested in situations where the gradient is expensive to compute and we aim at providing computation-efficient algorithms that produce either approximated gradients whose extreme component of interest is the same as the exact gradient's one, or identify the extreme component without computing the whole (exact) gradient.

III. LOOKING FOR THE EXTREME GRADIENT COMPONENT

This section formalizes the problem of identifying the extreme gradient component and provides different algorithms for its resolution. We first introduce a greedy approach, then a randomized one and finally exhibit how this problem can be cast as best arm identification in a multi-armed bandit framework.

A. The problem

As mentioned in Section II, we are interested in learning problems where the objective function is of the form

$$\sum_i \ell(y_i, g(\mathbf{w}^\top \mathbf{x}_i))$$

The gradient of our objective function is thus given by

$$\nabla L(\mathbf{w}) = \sum_i \ell'(y_i, g(\mathbf{w}^\top \mathbf{x}_i)) g'(\mathbf{w}^\top \mathbf{x}_i) \mathbf{x}_i \quad (6)$$

$$= \mathbf{X}^\top \mathbf{r} \quad (7)$$

where $\mathbf{r} \in \mathbb{R}^n$ is the vector whose i -th component is $\ell'(y_i, g(\mathbf{w}^\top \mathbf{x}_i)) g'(\mathbf{w}^\top \mathbf{x}_i)$. This particular form implies that the gradient can be computed iteratively, obtaining at each iteration an approximate gradient $\hat{\nabla} L$. That is, at iteration t , we have

$$\hat{\nabla} L_t = \hat{\nabla} L_{t-1} + U_t, \quad (8)$$

where U_t is $r_i \mathbf{x}_i$ for a given i .

According to this framework, our objective is then to find an efficient way for computing an approximate gradient $\hat{\nabla}L_t$ for which the desired extreme entry is equal to the one of the exact gradient. We propose and discuss three different approaches in the sequel.

B. Greedy deterministic approach

The first approach we propose is a greedy approach which, at iteration t , looks for the best index i so that $r_i \mathbf{x}_i$ optimizes a criterion depending on $\hat{\nabla}L_t$ and ∇L .

Let \mathcal{I}_t be the set of index of the examples chosen in the first t iterations for computing $\hat{\nabla}L$. At iteration $t + 1$, our goal is to find the example i^* that is the solution of the following problem:

$$i^* = \operatorname{argmin}_{i \in \{1, \dots, n\} \setminus \mathcal{I}_t} \|\nabla L - \hat{\nabla}L_t - \mathbf{x}_i r_i\| \quad (9)$$

with $\hat{\nabla}L_{t+1} = \hat{\nabla}L_t + \mathbf{x}_i r_i$. The solution of this problem is thus the vector product $r_i \mathbf{x}_i$ that has minimal norm compared to the current gradient estimation residual $\nabla L - \hat{\nabla}L_t$. While simple, this solution cannot be computed because ∇L is not accessible. Hence, we have to resort to an approximation based on the following equivalent problem:

$$\begin{aligned} i^* &= \operatorname{argmax}_{i \in \{1, \dots, n\} \setminus \mathcal{I}_t} -\|\nabla L - \hat{\nabla}L_{t+1}\| \\ &= \operatorname{argmax}_{i \in \{1, \dots, n\} \setminus \mathcal{I}_t} \|\nabla L - \hat{\nabla}L_t\| - \|\nabla L - \hat{\nabla}L_{t+1}\| \end{aligned} \quad (10)$$

where we use the fact that $\|\nabla L - \hat{\nabla}L_t\|$ does not depend on i . By upper bounding the above objective value, one can derive the best choice of $r_i \mathbf{x}_i$ that achieves the largest variation of the gradient estimation norm residual. Indeed, we have

$$0 \leq \|\nabla L - \hat{\nabla}L_t\| - \|\nabla L - \hat{\nabla}L_{t+1}\| \leq \|-\hat{\nabla}L_t + \hat{\nabla}L_{t+1}\|$$

The right-hand side of this equation can be further simplified

$$\begin{aligned} \|\hat{\nabla}L_{t+1} - \hat{\nabla}L_t\| &= \|\hat{\nabla}L_t - \mathbf{x}_i r_i - \hat{\nabla}L_t\| \\ &= \|\mathbf{x}_i r_i\| = \|\mathbf{x}_i\| |r_i| \end{aligned} \quad (11)$$

Equation (11) suggests that the index i that should be chosen at iteration $t + 1$ is the one with the largest absolute residual weighted by its example norm. Simply put, in the first iteration, the algorithm chooses the index that leads to the largest value $\|\mathbf{x}_i\| |r_i|$, in the second one, it selects the second best, and so forth. That is, the examples are considered in a decreasing order with respect to the weighted value of their residuals.

Note that for this method, at iteration $t = n$ we recover the exact gradient,

$$\hat{\nabla}L_n = \nabla L.$$

Hence, when $t = n$, we are assured to retrieve the correct extreme entry of the gradient at the expense of extra computations needed for running this greedy approach compared to a plain computation of the full gradient. On the other hand, if we stop this gradient estimation procedure before $t = n$, we save computational efforts at the risk of missing the correct extreme entry.

C. Matrix-Vector Product as Expectations

The problem of finding the extreme component of the gradient can also be addressed from the point of view of randomization.

The approach consists in considering the computation of $\mathbf{X}^\top \mathbf{r}$ as the expectation of a given random variable. Recall that \mathbf{X} is composed of the vectors $\{\mathbf{x}_i^\top\}_{i=1}^n$ and $\mathbf{x}_i \in \mathbb{R}^d$. Hence, the matrix-vector product $\mathbf{X}^\top \mathbf{r}$ can be rewritten:

$$\mathbf{X}^\top \mathbf{r} = \sum_{i=1}^n r_i \mathbf{x}_i \quad (12)$$

From now on, given some integer n , Δ_n^* denotes the interior of the probabilistic simplex of size n :

$$\Delta_n^* \doteq \left\{ \mathbf{p} = [p_1 \cdots p_n] : \sum_{i=1}^n p_i = 1, p_i > 0, i = 1, \dots, n \right\}. \quad (13)$$

For any element $\mathbf{p} = (p_i)_{i \in [n]} \in \Delta_n^*$ we introduce a random vector C that takes value in the set

$$\mathcal{C} \doteq \{\mathbf{c}_i \doteq r_i \mathbf{x}_i / p_i : i = 1, \dots, n\}$$

so that $\mathbb{P}(C = \mathbf{c}_i) = p_i$. This way,

$$\mathbb{E}C = \sum_{i=1}^n p_i \mathbf{c}_i = \sum_{i=1}^n p_i r_i \mathbf{x}_i / p_i = \sum_{i=1}^n r_i \mathbf{x}_i = \mathbf{X}^\top \mathbf{r} \quad (14)$$

Hence, if C_1, \dots, C_s are independent copies of C and \hat{C}^s is defined as:

$$\hat{C}^s \doteq \frac{1}{s} \sum_{i=1}^s C_i \quad (15)$$

then

$$\mathbb{E}\hat{C}^s = \mathbf{X}^\top \mathbf{r}. \quad (16)$$

\hat{C}^s is thus an estimator of the matrix-vector product that we are interested in *i.e* the gradient of our objective function.

According to the above, a relevant approach for estimating the extreme component of the gradient is to randomly sample s copies of C , to average them and then to look for the extreme component of this estimated gradient.

Interestingly, this approach based on randomized matrix multiplication can be related to our deterministic approach. Indeed, a result given by [7] (Lemma 4) says that the element $\mathbf{p} \in \Delta_n^*$ that minimizes the variance of \hat{C}^s is such that

$$p_i \propto |r_i| \|\mathbf{x}_i\|. \quad (17)$$

It thus suggests that vectors c_i of large values of $|r_i| \|\mathbf{x}_i\|$ have higher probability to be sampled. This resonates with the greedy deterministic approach in which vectors $\mathbf{x}_i r_i$ are accumulated in the order of the decreasing values of $|r_i| \|\mathbf{x}_i\|$.

D. Best arm identification and multi-armed bandits

The two preceding approaches seek at approximating the gradient, at a given level of accuracy that has yet to be defined, and then at evaluating the coordinate of its extreme component.

Yet, the problem of finding the extreme component coordinate of a gradient vector obtained from matrix-vector multiplication can also be posed as the problem of finding the best arm in a multi-armed bandit problem. In a nutshell, given a slot machine with multiple arms, the goal in the bandit problem is to find the arm that maximizes the expected reward or minimizes the loss. For this, an iterative procedure is used, where at each step, a forecaster selects an arm, based on his previous actions, and receives a reward or observes a loss. Depending on how the reward is obtained, the problem can be stochastic (the reward/loss is drawn from a probability distribution) or non-stochastic. Bubeck et al. [6] propose an extensive review of these methods for various settings.

We cast our problem of finding the extreme gradient component as best arm identification problem as follows. In the remainder of this section, we suppose that we look for a minimum gradient component and thus we look for the arm with minimal loss instead of maximal reward. We consider that the arms are the components of the gradient (we have d arms) and at each pull of a given arm, we observe a loss that is built from a term of the gradient matrix-vector multiplication $\mathbf{X}^\top r$, as made clear in the sequel.

In a stochastic setting, we consider a similar framework as the one we described in Section III-C. We model the loss obtained from the k -th pull of arm $i \in [1, \dots, d]$ as a random variable V , independent of k , that takes value in the set

$$\{v_{i,j} \doteq r_j x_{j,i} / p_j : j = 1, \dots, n\}$$

so that $\mathbb{P}(V = v_{i,j}) = p_j$. From this definition, the expected loss of arm i is

$$\mathbb{E}V = \sum_{j=1}^n p_j v_{i,j} = \sum_{j=1}^n r_j x_{j,i} = (\mathbf{X}^\top \mathbf{r})_i,$$

which is the i -th component of our gradient vector. In this setting, given a certain number of pulls, one pull providing a realization of the random variable V of the chosen arm, the objective of a best arm identification algorithm is to provide recommendation about the arm with minimal expected loss, which in our case is the component that minimizes $\mathbf{X}^\top \mathbf{r}$.

Several algorithms for identifying this best arm have been provided in the literature. Most of them are built around an empirical average loss statistic. This latter can be computed, after s pulls of an arm i , as

$$\hat{V}_{i,s} = \frac{1}{s} \sum_{t=1}^s v_{i,j(t)}$$

where $j(t)$ is the index j drawn at the t -th pull. Some of the most interesting algorithms are the *successive reject* [2] and *successive halving* [13] algorithms which, given a fixed budget of pulls, iteratively discard after some predefined number of pulls (say s) the worst arm or the worst-half arms, respectively, according to the values $\{\hat{V}_{i,s}\}_{i=1}^d$. These approaches are relatively simple and the *successive halving* approach is depicted in detail in Algorithm 3. They directly provide some guarantees on the probability to correctly estimate the extreme component of the gradient. For instance, for a fixed budget of pulls T , under some minor and easily satisfied conditions on $\{v_{i,j}\}$, the *successive halving* algorithm correctly identifies the minimum gradient component with probability at least $1 - 3 \log_2 d \cdot \exp(-\frac{T}{8H_2 \log_2 d})$ where H_2 is a problem-dependent constant (the larger H_2 is, the harder the problem is).

However, they have the drawbacks to work on individual entries $x_{j,i}$ of the matrix \mathbf{X} . Hence, overload due to single memory accesses compared to those needed for accessing chunk of memory may hinder the computational gain obtained by identifying the component with minimal gradient value without computing the full gradient.

For the *successive halving* algorithm, one way of overcoming this issue is to consider *non i.i.d* sampling of the arm's loss. As such, we consider that at each iteration, the losses generated by pulling the remaining arms come from the same component j of the residual. Hence, if we denote \mathcal{A} the set of components j that have been drawn after s pulls, then the loss for arm i can be defined as

$$\hat{V}_{i,s} = \frac{1}{s} \sum_{j=1}^s v_{i,\mathcal{A}_j}, \quad (18)$$

where \mathcal{A}_j is the same for all arms. According to this equation, the empirical losses of all the arms are dependent through the set \mathcal{A} . However, it is easy to see that each empirical loss is still a relevant estimation of the expected loss of each arm.

Algorithm 3 Successive Halving algorithm for finding the minimum gradient component

- 1: **inputs:** \mathbf{X} , \mathbf{r} , set B the budget of pulls
 - 2: Initialize $S_0 = [d]$
 - 3: **for** $\ell=0,1, \dots, \lceil \log_2(d) \rceil - 1$ **do**
 - 4: Pull each arm in S_ℓ for $r_\ell = \lfloor \frac{B}{|S_\ell| \log_2(n)} \rfloor$ additional times and compute resulting loss (stochastic) \hat{V}_{\cdot, R_ℓ} or (non-stochastic) v_{\cdot, R_ℓ} with $R_\ell = \sum_{j=0}^{\ell} r_j$
 - 5: Sort arms in S_ℓ by increasing value of losses
 - 6: Define $S_{\ell+1}$ as the $\lfloor |S_\ell|/2 \rfloor$ indices of arms with smallest values
 - 7: **end for**
 - 8: **output:** index of the best arm
-

Non-stochastic best arm identification has been barely studied and only a very recent work has addressed the problem [11]. In this work, the main hypothesis about the non-stochasticity of the loss is that they are assumed to be generated before the game starts. This is exactly our situation since losses are given by $x_{j,i}r_j$ which can be computed beforehand. In this non-stochastic setting [11], the framework is that the k -th pull of an arm i provides a loss $v_{i,k}$ and the objective of the bandit algorithm is to identify

$$\arg \min_i \lim_{k \rightarrow \infty} v_{i,k}$$

assuming that such limits exist for all i . Again we can fit our problem of finding the extreme gradient component (here the minimum) into this framework by defining the loss for a given arm at pull k as

$$v_{i,k} = \begin{cases} 0 & \text{if } k = 0 \\ v_{i,k-1} + r_{\tau_k} x_{\tau_k, i} & \text{if } 1 \leq k \leq n \\ v_{i,k-1} & \text{if } k > n \end{cases} \quad (19)$$

where τ is a predefined or random permutation of the rows of the vector \mathbf{r} and the matrix \mathbf{X} , that needs not to be the same for all the arms. According to this definition, we have

$$\lim_{k \rightarrow \infty} v_{i,k} = \sum_{k=1}^n r_{\tau_k} x_{\tau_k, i} = \sum_{k=1}^n r_k x_{k, i} = (\mathbf{X}^\top \mathbf{r})_i$$

Hence, an algorithm that recommends the best arm after a given number of pulls, will return the index of the minimum component in our gradient. Interestingly, the algorithm proposed by Jamieson et al. [11] for solving the non-stochastic best arm identification problem is also the one used in the stochastic setting namely the *successive halving* algorithm (Alg. 3).

As already mentioned, for a fixed budget T of pulls, these bandit algorithms both come with theoretical guarantee in their ability to identify the best arm. In the stochastic case, the probability of success depends on the number of arms, the number of pulls T and on a parameter denoting the hardness of the problem (see Theorem 4.1 in [13]). In the non-stochastic case, the budget of pulls needed for guaranteeing the correct recovery of the best arm essentially depends on a function $\gamma_i(t)$ such that

$$|v_{i,k} - \lim_{k \rightarrow \infty} v_{i,k}| \leq \gamma_i(t)$$

and on a parameter denoting the gap between any of $(\mathbf{X}^\top \mathbf{r})_i$ and the smallest component of $\mathbf{X}^\top \mathbf{r}$ which is not accessible unless we compute the exact gradient.

One may have noted that the *non-iid* bandit strategy we have proposed bear strong resemblance with the matrix-vector product approximation. In addition, in the non-stochastic bandit setting, we can recover the greedy deterministic approach if we assume that the permutation τ defines a re-ordering of $\|\mathbf{x}_i\| r_i$ in decreasing order, then the accumulation given in Equation (19) is exactly the one given in the greedy deterministic approach. Multi-armed bandit framework and the gradient approximation approaches use thus similar ways for computing the criteria used for estimating the best arm. The main difference resides in the fact that with multi-armed bandit, one is directly provided with the estimation of the best arm.

E. Stopping criteria

In the greedy deterministic and randomized methods introduced in this section, we have no clues on how many elements $r_j \mathbf{x}_j$ have to be accumulated in order to achieve a sufficient approximation of the gradient or in the multi-armed bandit approach how many pulls we need to draw. Here, we discuss two possible stopping criteria for the non-bandit algorithms: one that holds for any approach and a second one that holds only for the Frank-Wolfe algorithm in the deterministic sampling case. Discussion on the budgets that need to be allocated to the bandit problem is also provided.

1) *Stability condition:* For the sake of simplicity, we limit the exposition to the search of the smallest component of the gradient, although the approach can be generalized to other cases.

Denote by j^* the coordinate such that $j^* = \arg \min_j \nabla L(\mathbf{w}_k)|_j$ and let T_s be the maximal number of iterations or samplings allowed for computing the inexact gradient (for instance, in the greedy deterministic approach, $T_s = n$). Our objective is to estimate j^* with the fewest number t of iterations. For this to be possible, we make the hypothesis that there exists an iteration t_1 , $t_1 \leq T_s$, and

$$j^* = \arg \min_j \hat{\nabla} L_t(\mathbf{w}_k)|_j \quad \forall t : t_1 \leq t \leq T_s$$

in other words, we suppose that starting from a given number of iterations t_1 , the gradient approximation is sufficiently accurate so that the updates of the gradient will leave the minimum coordinate unchanged. Formally, this condition means that $\forall t \in [t_1, \dots, n]$, we have

$$\left[\hat{\nabla}L_t + \sum_{i=0}^{T_s-t} U_{t+i} \right]_{j^*} \leq \left[\hat{\nabla}L_t + \sum_{i=0}^{T_s-t} U_{t+i} \right]_j \quad \forall j \in [1, \dots, d].$$

However, checking the above condition is as expensive as computing the full gradient, thus we propose an estimation of j^* based on an approximation of this inequality, by truncating the sum to few iterations on each side. Basically, this consists in evaluating j^* at each iteration and checking whether this index has changed over the last N_s iterations. We refer to this criterion as the *stability criterion*, parametrized by N_s .

2) *Error bound criterion*: In Section II-B2, we discussed that the convergence of the Frank-Wolfe inexact algorithm can be guaranteed as long as the norm difference between the approximate gradient and the exact one could be upper-bounded by some quantity ϵ . Formally, this means that the iterations over gradient approximation can be stopped as soon as

$$\|\hat{\nabla}L_t - \nabla L\|_\infty \leq \epsilon, \quad (20)$$

where ϵ depends on the curvature of the function $L(\mathbf{w})$ [10]. In practice, the criterion $\|\hat{\nabla}L_t - \nabla L\|_\infty$ cannot be computed as it depends on the exact gradient but it can be upper-bounded by a term that is accessible. For the greedy deterministic approach, by norm equivalence, we have

$$\|\hat{\nabla}L_t - \nabla L\|_\infty \leq \|\hat{\nabla}L_t - \nabla L\|_1$$

and then, through simple algebras, we have

$$\begin{aligned} \|\hat{\nabla}L_t - \nabla L\|_1 &= \sum_{i=1}^d |\hat{\nabla}L_t - \nabla L|_i = \sum_i \left| \sum_{j \notin \mathcal{I}_t} \mathbf{x}_j r_j \right|_i \\ &\leq \sum_i \sum_{j \notin \mathcal{I}_t} |\mathbf{x}_j|_i |r_j| = \sum_{j \notin \mathcal{I}_t} \|\mathbf{x}_j\|_1 |r_j| \end{aligned}$$

which leads to the following criterion in the greedy deterministic approach,

$$\sum_{j \notin \mathcal{I}_t} \|\mathbf{x}_j\|_1 |r_j| \leq \epsilon$$

Hence if the norms $\{\|\mathbf{x}_j\|_1\}$ have been precomputed beforehand, this criterion can be easily computed at each gradient update iteration.

3) *Pull budget for the bandit*: In multi-armed bandit algorithms, one typically specifies the number of pulls B available for estimating the best arm. As such, T can be considered as hyperparameter of the algorithm. A possible strategy for removing the dependency of the bandit algorithm to this pull budget, is to use the *doubling trick* [11], which consists in running the algorithm with a small value of T and then repeatedly doubling it until some stopping criterion is met. The algorithm can then be adapted so as to exploit the loss computations that have already been carried from iteration to iteration. However, this strategy needs a stopping criterion for the *doubling trick*. According to Theorem 1 in [11], there exists a lower bound of pulls for which the algorithm is guaranteed to return the best arm. Hence, the following heuristic can be considered: if T' and $2T'$ number of pulls return the same best arm, then we conjecture that the proposed arm is indeed the best one.

IV. DISCUSSIONS

This section provides comments and discussions of the approaches we proposed compared to existing works.

A. Relation and gains compared to OMP and variants

Several recent works on sparse approximation have proposed theoretically founded algorithms. These works include OMP [21, 29], greedy pursuit [5, 30], CosAmp [18] and several others like [25]. Most of these algorithms make use of the top absolute entry of the gradient vector at each iteration. The work presented in this paper is strongly related to these ones as we share the same quest for the top entry. Indeed, the proposed methodology provides tools that can be applied to many sparse approximation algorithms including the aforementioned ones. What makes our work novel and compelling is that at each iteration, the gradient is computed with as few information as possible. If the stopping criterion for estimating this gradient is based on a maximal number of samples — *e.g.* we are interested in constructing the best approximation of the gradient from only 20% of the samples—, our approach can be interpreted as a method for computing the gradient on a limited budget. Hence, the proposed method allows to obtain a gain in the computational time needed for the estimation of the gradient. On the downside, if other stopping criteria are used (alone or jointly with the budget criterion), this gain may be partly impaired by further computations needed for their estimation. As an example, the *stability* criterion induces a $O(d)$ overhead at each iteration due to the max computation.

B. About stochasticity

The randomization approach for approximating the gradient, introduced in Section III-C, involve random sampling of the columns. In the extreme situation where only a single column i is sampled, we thus have :

$$\hat{\nabla}L = \mathbf{x}_i r_i$$

and the method we propose boils down to a stochastic gradient method. In the context of sparse greedy approximation, the first work devoted to stochastic gradient approximation has been recently released [19]. Nguyen et al. [19] show that their stochastic version of the iterative hard thresholding algorithm, or the gradient matching pursuit algorithm which aim at greedily solving a sparse approximation problem with arbitrary loss functions behave properly in expectation.

The randomized approach we propose in this work go beyond the stochastic gradient method for greedy approximation, since it also provides a novel approach for computing stochastic gradient. Indeed, we differ from their setting in several important aspects:

- First, in our stochastic gradient approximation, we always consider a number of samples larger than 1. As such, we are essentially using a stochastic mini-batch gradient.
- Second, the size of the mini-batch is variable (depending on the stopping criterion considered) and it depends on some heuristics that estimates on the fly the ability of the approximate gradient to retrieve the top entry of the true gradient.
- Finally, one important component of our approach is the importance sampling used in the stochastic mini-batch sampling. This component theoretically helps in reducing the variance of the gradient estimation [20]. In the context of matrix multiplication approximation we used for developing the randomized approach in Section III-C, theoretical results of Drineas et al. [7] have also shown there exists an importance sampling that minimizes the variance of the matrix multiplication estimation. Our experiments confort the results showing that, compared to a uniform sampling, importance sampling clearly enhances the efficiency in retrieving the top entry of the true gradient.

All these differences make our randomized algorithm not only clearly distinguishable from stochastic gradient approaches, but also harder to analyze. We thus defer for further work the theoretical analyses of such a stochastic adaptive-size mini-batch gradient coupled with the importance sampling approach.

C. Some theoretical considerations

Although a complete theoretical analysis of the algorithm is out of the scope of this paper, an interesting property deserves to be mentioned here. Note that, unlike stochastic gradient approaches, our algorithm is built upon inexact gradients that hopefully have the same minimum component as the true gradient. If this latter fact occurs along the iterations of the FW or OMP algorithms, then all the properties (*e.g.* *linear convergence*, *exact recovery property*, ...) of these algorithms apply. Based on the probability of recovering an exact minimum component of the gradient at each iteration, we prove below a bound on the probability of our algorithm to recover at a given iteration K of OMP or FW, the same sequence of minimum components as the one obtained with exact gradient.

Theorem 1: Suppose that at each iteration t of OMP or FW, our algorithm for estimating the minimum component correctly identifies this component with a probability at least $1 - P_t$ and there exists \bar{P} so that $P_t > \bar{P}$, $\forall t \leq K$, then the probability of identifying the correct sequence of minimum component is at least

$$1 - K\bar{P}.$$

Proof. Denote as $B(t) = \{I^1 = i_e^1, I^2 = i_e^2, \dots, I^t = i_e^t\}$, the event that our algorithm outputs the exact sequence of minimum component up to iteration t , I^t and i_e^t being the coordinates retrieved with the inexact and exact gradient. Similarly, we note $A(t) = \{I^t = i_e^t\}$ the event of retrieving, at iteration t , the correct component of the gradient. We assume that $\mathbb{P}(B(0)) = 1$. Formally, we are interested in lower-bounding the probability of $B(K)$. By definition, we naturally have

$$\begin{aligned} \mathbb{P}(B(K)) &= \mathbb{P}(A(K)|B(K-1))\mathbb{P}(B(K-1)) \\ &= \mathbb{P}(B(0))\prod_{t=1}^K \mathbb{P}(A(t)|B(t-1)) \end{aligned}$$

and $\mathbb{P}(A(t)|B(t-1)) \geq 1 - P_t$. Accordingly, we have

$$\begin{aligned} \mathbb{P}(B(K)) &\geq \prod_{t=1}^K (1 - P_t) \\ &\geq (1 - \bar{P})^K \\ &\geq (1 - K\bar{P}) \end{aligned}$$

where in the last inequality, we used the fact that $(1 - u)^K \geq (1 - Ku)$ for $0 \leq u \leq 1$. ■

For instance, in the *successive halving* algorithm, we have $P_t = 3 \log_2 d \cdot \exp\left(-\frac{T}{8H_2(t)\log_2 d}\right)$, where $H_2(t)$ is a iteration-dependent constant [13] and T the number of pulls. Thus, if we define \bar{H} so that

$\forall t, \bar{H} \geq H_2(t)$, we have $\bar{P} = 3 \log_2 d \cdot \exp\left(-\frac{T}{8\bar{H}\log_2 d}\right)$ and

$$P(B(K)) \geq 1 - 3K \log_2 d \cdot \exp\left(-\frac{T}{8\bar{H}\log_2 d}\right).$$

We can see that the probability of our OMP or FW having the same behaviour as their exact counterpart decreases with the number K of iterations and the number d of dimensions of the problems and increases with the number of pulls. By rephrasing this last equation, we also get the following property. For $\delta \in [0, 1]$, if the number T of pulls is set so that at each iteration t ,

$$T \geq \left(\log \log_2 \frac{d}{\delta} + \log \frac{K}{\delta} + \log \frac{3}{\delta}\right) 8\bar{H} \log_2 d$$

then, when using the *successive halving* algorithm for retrieving the extremum gradient component, an inexact OMP or FW algorithm behaves like the exact OMP or FW with probability $1 - \delta$. This last property is another emphasis on the strength of our inexact gradient method compared to stochastic gradient descent approaches as it shows that with high probability, all the theoretical properties of OMP or FW (*e.g* convergence, exact recovery of sparse signal) apply.

V. NUMERICAL EXPERIMENTS

In this section, we describe the experimental studies we have carried out for illustrating the computational benefits of using inexact gradient for sparsity-constraint optimization problems.

A. Experimental setting

In order to illustrate the benefit of using inexact gradient for sparse learning or sparse approximation, we have set up a simple sparse approximation problem which focuses on the computational gain, and for which a sparse signal has to be recovered by the Frank-Wolfe, OMP or CosAMP algorithm. The target sparse signals are built as follows. For a given value of the dictionary size d and a number k of active elements in the dictionary, the true coefficient vector \mathbf{w}^* is obtained as follows. The k non-zero positions are chosen randomly, and their values are drawn from a zero-mean unit variance Gaussian distribution, to which we added ± 0.1 according to the sign of the values. The columns of the regression design matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ are drawn uniformly from the surface of a unit hypersphere of dimension n . Finally, the target vector is obtained as

$$\mathbf{y} = \mathbf{X}\mathbf{w}^* + \mathbf{e}$$

where \mathbf{e} is a random noise vector drawn from a Gaussian distribution with zero-mean and variance σ_e^2 determined from a given signal-to-noise as

$$\sigma_e^2 = \frac{1}{n} \|\mathbf{X}\mathbf{w}^*\|^2 \cdot 10^{-SNR/10}.$$

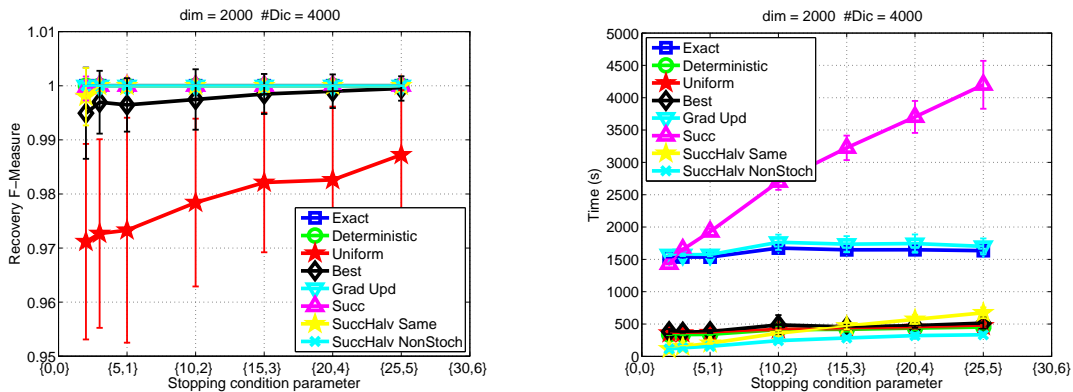


Fig. 1. Comparing vanilla Frank-Wolfe and inexact FW algorithms with different ways for computing the inexact gradient. Performances are compared with increasing precision on the inexact computations.

Unless specified, the SNR ratio has been set to 3. For each setting, the results are averaged over 20 trials, and \mathbf{X} , \mathbf{w}^* and \mathbf{e} are resampled at each trial.

The criteria used for evaluating and comparing the proposed approaches are the running time of the algorithms and their ability to recover the true non-zero elements of \mathbf{w}^* . The latter is computed through the F-measure between the support of the true coefficient vector \mathbf{w}^* and the estimated one $\hat{\mathbf{w}}$:

$$\text{F-meas} = 2 \frac{|\text{supp}_\gamma(\mathbf{w}^*) \cup \text{supp}_\gamma(\hat{\mathbf{w}})|}{|\text{supp}_\gamma(\mathbf{w}^*)| + |\text{supp}_\gamma(\hat{\mathbf{w}})|}$$

where, $\text{supp}_\gamma(\mathbf{w}) = \{j : |w_j| > \gamma\}$ is the support of vector \mathbf{w} and γ is a threshold used to neglect some non-zero coefficients that may be obliterated by the noise. In all our experiments, we have set $\gamma = 0.001$ which is small compared to the minimal absolute value of a non-zero coefficient (0.1).

All the algorithms (Frank-Wolfe, OMP and CosAMP) and exact and inexact gradient codes have been written in Matlab except for the *successive reject* bandit which as been written in C and transformed in a mex file. All computations have been run on a Linux machine with 24 Gb of memory on each single core of an Intel Xeon E5530 processor clocked at 2.4 GHz.

B. Sparse learning using a Frank-Wolfe algorithm

For this experiment, the constraint set \mathcal{C} is the ℓ_1 unit-ball and the loss function is $L(\mathbf{w}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2$. Our objectives are is two-folded:

- 1) analyze the capability of inexact gradient approaches to recover the true support and compare them to the FW algorithm with,

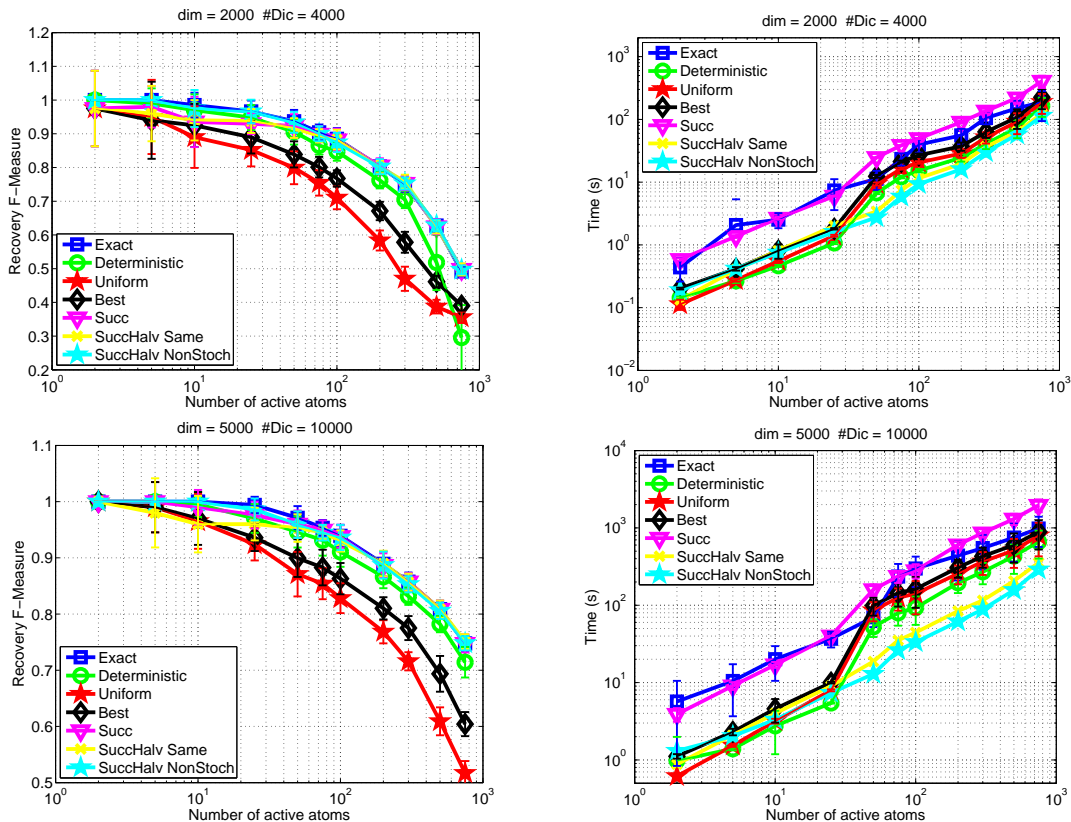


Fig. 2. Comparing OMP algorithm with different ways for computing the inexact gradient. The comparison holds for dictionary size with (top) $n = 2000$, $d = 4000$. (bottom) $n = 5000$, $d = 10000$ and for an exact recovery criterion (left) and running time (right).

- 2) compare the two stopping criteria for computing the inexact gradient : the *stability condition* and the *error bound condition*.

In the latter, while this *error bound condition* provides an adaptive condition for stopping — recall that the parameter ϵ in Equation (20) is determined automatically through the data and the related curvature of the loss function —, the *stability condition* needs a user-defined parameter N_s for stopping the accumulation of the partial gradient. Similarly, we use a fixed pre-defined budget of pulls in the best-arm identification problem. This budget is given as a ratio of $n \times d$. The exact gradient is computed using the accumulation strategy as given in Equation (8) so as to make all running times comparable. The maximum number of iterations for FW is set to 5000.

Figure 1 presents the results obtained for $n = 2000$ samples, $d = 4000$ dictionary elements and $k = 50$ active atoms. We give the running time and recovery abilities of the Frank-Wolfe algorithm with an exact

gradient (**exact**), a greedy deterministic gradient sampling computation with a stability stopping criterion (**deterministic**) and an error bound stopping criterion (**grad upb**), the randomized approach with a uniform sampling (**uniform**), and with a best probability sampling as given in Equation (17) (**best**), the successive reject bandit (**succ**), the *non-iid* successive halving approach with losses computed as given in Equation (18) (**SuccHalvSame**) and the non-stochastic successive halving approach with losses computed as given in Equation (18) (**SuccHalvNonStoch**).

The figures depict the performances with respect to stopping condition parameter N_s of the stability criterion (the first value in the bracket) and the sampling budget of the bandit approach ($\frac{n \times d}{10}z$ where z is the second value in the bracket). First, we can note that the deterministic approach used with any stopping criterion and the non-stochastic successive halving approaches are able to perfectly recover the exact support of the true vector \mathbf{w}^* , regardless of the considered stopping criterion’s value. Randomized approaches with uniform and best probability sampling nearly achieve perfect recovery with an average F-measure of 0.975 with the stability criterion N_s equal to 5. As expected when N_s increases, the performances of these two approaches increase also but still fail to achieve perfect recovery.

From a running time point of view, the proposed approaches based on greedy deterministic and randomized sampling strategies with stability criterion and the successive halving strategies are faster than all methods. They achieve a gain in running time of a factor 2 to 4 with respect to the exact Frank-Wolfe algorithm. Interestingly, this gain is achieved without compromise on the recovery performance for the greedy deterministic approach and the successive halving approaches. For the randomized strategies, increasing the stability parameter N_s leads to a very slight increase of running time, hence for these methods, a trade-off can eventually be found. When comparing bandit approaches, one can note the substantial gain in performances that can be obtained by the halving strategy, the *non-iid* and *non-stochastic* strategies. We conjecture that this higher computational running time of the *successive reject* algorithm is essentially due to computational overhead needed for accessing each single matrix entry $\mathbf{X}_{i,j}$ in memory while all other methods use slices of this matrix (through the samples \mathbf{x}_i) and thus they can leverage on the chunk of memory access.

When comparing the *stability* and the *error bound* stopping criteria, the latter one is rather inefficient. While grounded on theoretical analysis, the bound is loose enough to be non-informative. Indeed, a careful inspection shows that most inexact gradient computations needed the accumulation of roughly 98% of the sample’s residuals. In addition, other computational overheads necessary for the bound estimation, make the approach less efficient than the exact Frank-Wolfe algorithm.

In summary, from this experiment, we can conclude that the non-stochastic bandit approach is the

most efficient one. It can achieve a gain in computation of about an order of magnitude (the left most point in the Figure 1’s right panel) without compromising accuracy. The greedy deterministic approach with stability criterion performs also very well but it is slightly less efficient. We can remark that these two best methods both use the same strategy of gradient accumulation.

C. Sparse Approximation with OMP

Here, we evaluate the usefulness of using inexact gradient in a greedy framework like OMP. The toy problem is similar as the one used above except that we analyze the performance of the algorithm for an increasing number k of active atoms and two sizes of dictionary matrix \mathbf{X} have been considered.

Several ways for computing the inexact gradient are evaluated and compared in terms of efficiency and correctness to the true gradient in an OMP algorithm. For all sampling approaches, the stopping criterion for gradient accumulation is based on the stability criterion with the parameter N_s adaptively set at 1% of the dictionary size. For the successive reject bandit approach, the sampling budget has been limited to 20% of the number of entries in the matrix \mathbf{X} . In all cases, the stopping criterion for the OMP algorithm is based on a fixed number of iterations and this number is the desired sparsity k .

Results are reported in Figure 2. They globally follow the same trend as those obtained for the Frank-Wolfe algorithm. First, note that in terms of support recovery, when the number of active atoms is small, the greedy deterministic approach seems to perform better than the randomized sampling strategy. Bandit approaches, except the non-stochastic one, perform slightly worse than the greedy deterministic method. As the number of active atoms increases, the bandit approaches succeed better in recovering the extreme component of the gradient while the deterministic approach is less accurate. Note that for any value of k , the randomized strategies suffer more than the other strategies for recovering the true vector \mathbf{w}^* support. From a running time point of view, again, we note that deterministic and randomized strategies are the most efficient ones, when k is small. As the number of active atoms increases, successive halving bandit approaches become the most efficient methods. The gap in running time compared to other approaches is significant especially for large dimension dictionaries.

D. Sparse Approximation with CosAMP

To the best of our knowledge, there are very few greedy algorithms that are able to leverage from stochastic gradients. One of these algorithms has been introduced in [19]. In this experiment, we essentially want to evaluate the efficiency gain achieved by our inexact gradient approach compared to this stochastic greedy algorithm. Our objective is to show that the approach we propose is empirically

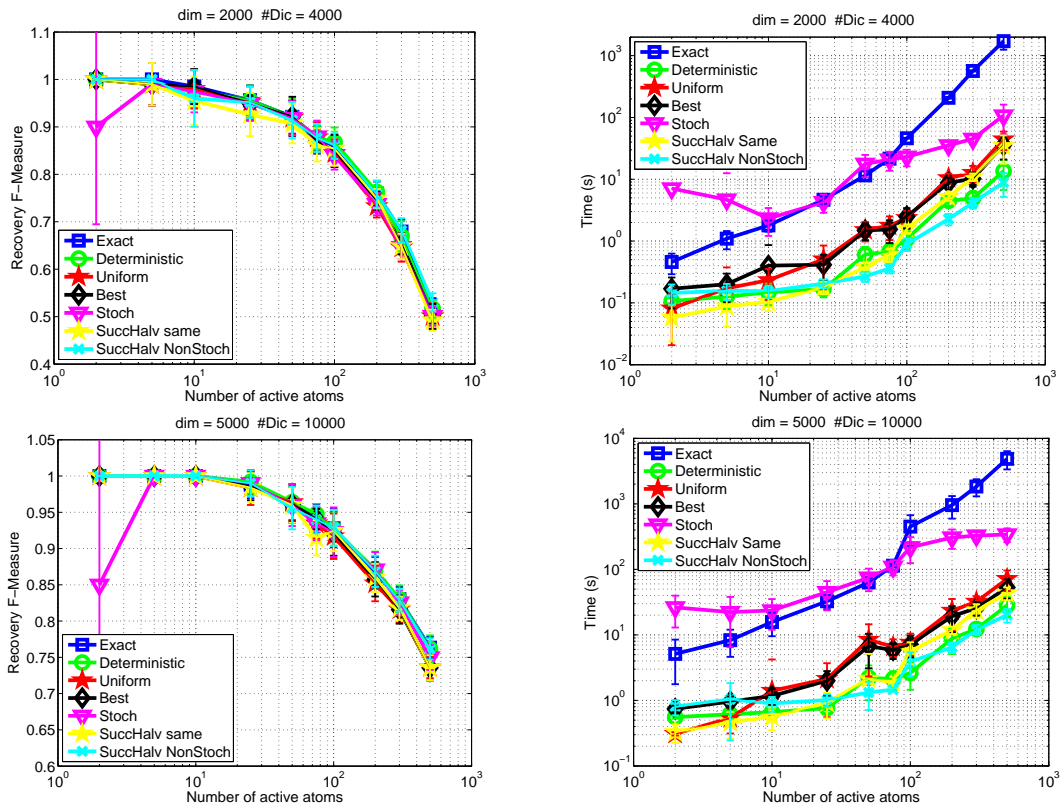


Fig. 3. Comparing CosAMP and CosAMP algorithms with different ways for computing the inexact gradient. The comparison holds for dictionary size with (top) $n = 2000$, $d = 4000$. (bottom) $n = 5000$, $d = 10000$ and for an exact recovery criterion (left) and running time (right).

significantly faster than a pure stochastic gradient approach. For the different versions of the CosAMP algorithm, we have set the stopping criterion as follows. For the CosAMP with exact gradient approach, which serves only as a baseline for computing the exact solution, the number of iteration is set to the level of sparsity k of the target signal. A tolerance of the residual norm is also used as a stopping criterion which should be below $1e^{-3}$. Next, for the stochastic and the inexact gradient CosAMP versions, the algorithms were stopped when the norm of the residual ($y - Xw$) became smaller than 1.01 times the one obtained by the exact CosAMP or when a maximal number of iterations. Regarding gradient accumulation, the stopping criterion we choose is based on the stability condition with the parameter N_S set dynamically at 1% of the dictionary size. For the bandit approaches, we have fixed the budget of pulls at $0.2n \times d$.

Note that for the CosAMP algorithm, we do not look for the top entry of the gradient vector but for

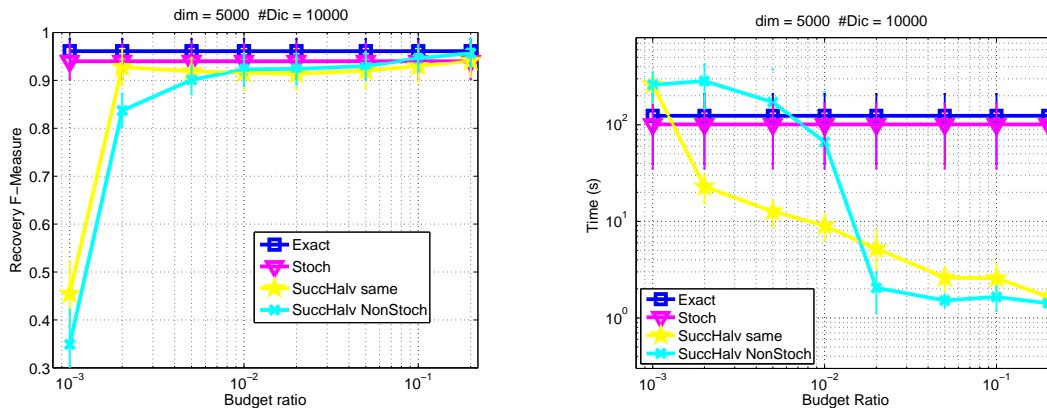


Fig. 4. Analyzing the effect of the pull budget on the successive halving algorithm (left) recovery F-measure and (right) computational running time. The pull budget is defined as the budget ratio times $n \cdot d$.

the top $2k$ entries as such, we have thus straightforwardly adapted the *successive halving* algorithm to handle such a situation.

Figure 3 presents the observed results. Regarding support recovery, we remark that all approaches achieve performances similar to the approach which uses exact gradient computations. When few active atoms are in play, we can note that sometimes, the stochastic approach of [19] fails to recover the support of \mathbf{w}^* . This occurs seldom but it happens regardless of the dictionary size we have experimented with.

From a running time perspective, the results show that the proposed approaches are highly more efficient than the exact gradient approach and interestingly, they are faster than a pure gradient stochastic approach. One or two orders of magnitude can be gained depending on the level of sparsity of the signal to be recovered. This observation clearly depicts the trade-off that occurs in sparsity-constrained optimization problems in which the gradient computation and an approximation problem on a limited number of atoms are the major computational burdens (Lines 3 and 5 of Algo 1). Indeed in a stochastic gradient approach, inexact gradient computations are very cheap but more approximation problems to be solved may be needed for achieving a desired accuracy. In the approaches we propose, the inexact gradient computation is slightly more expensive but we somehow “ensure” that it provides the correct information needed by the CosAMP algorithm. Hence, our approaches need less approximation problems to be solved, making them more efficient than a stochastic gradient approach.

When comparing the efficiency of the proposed algorithm, the approach based on non-stochastic successive halving and the greedy deterministic approach are the most efficient especially as the number of active atoms grow.

In the second experiment, we analyze the effect of the budget of pulls on the running time and on the recovery performance. We consider the following setting of the problem: $n = 5000$, $d = 10000$ and $k = 50$. Results are given in Figure 4. We note that a budget ratio varying from 0.02 and 0.2 allows a good compromise between ability of recovering the true vector and gain in computational time, particularly for the non-stochastic successive halving method. As the budget of pulls decreases, both algorithms fail more frequently in recovery and in addition, the computational gain vanishes. This experiment suggests that one should not be too greedy and should allow sufficient amount of pulls (about a budget ratio of 0.1 for the bandit algorithms).

E. Application to audio data

We have compared the efficiency of the approaches we propose on a real signal processing application. The audio dataset we use is the one considered by Yaghoobi et al. [31]. This dataset is composed of an audio sample recorded from a BBC radio session which plays classical music. From that audio sample, 8192 pieces of signal have been extracted, each being composed of 1024 time samples. Details about the dataset can be found in [31]. From this dataset, we have learned 2048 dictionary atoms using the approach described in [22]. In what follows, we have evaluated the runtime and the approximation quality of a CoSaMP algorithm using an exact gradient computation (**Exact**), a stochastic gradient CoSaMP algorithm (**Stock k**) and the CoSaMP variants with inexact gradient computations as we propose. For all algorithms, the number of CoSaMP iterations is fixed to the sparsity pattern, here $k = 10$. Note that for the stochastic gradient approach, we have also considered a version with more iterations (**Stock $3k$**).

Results are gathered in Table I and they are obtained as the averaged performance when approximating the first 4500 pieces of audio signal in the dataset. We can see that the inexact approaches we introduce lead to the best compromise between approximation error and running time. The stochastic gradient CoSaMP approaches are efficient but lack in properly approximating target signals. An example of signal approximation is provided in Figure 5.

VI. CONCLUSIONS

The methodologies proposed in this paper aim at accelerating sparsity-constrained optimization algorithms. This is made possible thanks to the key observation that, at each iteration, only the component of the gradient with smallest or largest entry is needed, instead of the full gradient. Exploiting this insight, we proposed greedy algorithms, randomized approaches and bandit-based best arm identification methods

TABLE I

APPROXIMATION PERFORMANCE RESULTS AND RUNNING TIME FOR CoSAMP AND VARIANTS. \mathbf{y} AND $\hat{\mathbf{y}}$ RESPECTIVELY DEPICTS THE SIGNAL AND ITS RESULTING APPROXIMATION. RESULTS ARE AVERAGED OVER THE APPROXIMATION OF 4500 SIGNALS. *Stoch 3k* DENOTES THE STOCHASTIC GRADIENT ALGORITHM THAT USED $3k$ ITERATIONS INSTEAD OF k .

| Approaches | $\frac{\ \mathbf{y}-\hat{\mathbf{y}}\ }{\ \mathbf{y}\ }$ | Time (s) |
|------------------|--|------------------|
| Exact | 0.390 ± 0.21 | 0.476 ± 0.01 |
| Stoch k | 0.663 ± 0.14 | 0.044 ± 0.00 |
| Stoch $3k$ | 0.572 ± 0.16 | 0.123 ± 0.01 |
| Deterministic | 0.370 ± 0.21 | 0.297 ± 0.01 |
| Uniform | 0.354 ± 0.19 | 0.194 ± 0.01 |
| SuccHalvSame | 0.385 ± 0.21 | 0.101 ± 0.00 |
| SuccHalvNonStoch | 0.388 ± 0.21 | 0.116 ± 0.00 |

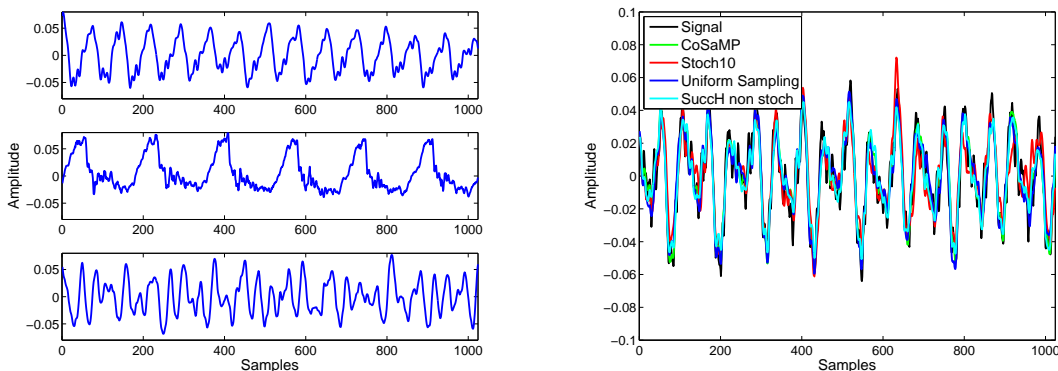


Fig. 5. (left) Examples of dictionary atoms learned from [22] and (right) comparing a signal and its different approximations obtained from the CosAMp algorithm with a sparsity pattern of $k = 10$ and different methods for finding the $2T$ most correlated atoms .

for estimating efficiently this top entry. Our experimental results show that the bandit and the greedy approaches seem to be the most efficient methods for this estimation.

Future works will be geared towards gaining further theoretical understandings on the good behaviour of the greedy approach, analyze the role of the importance sampling in the randomized methods. In addition, we plan to explore how this work can be extended to an online and/or distributed computation setting.

REFERENCES

- [1] Aleksandr Aravkin, Aurelie Lozano, Ronny Luss, and Prabhajan Kambadur. Orthogonal matching pursuit for sparse quantile regression. In *Data Mining (ICDM), 2014 IEEE International Conference on*, pages 11–19. IEEE, 2014.
- [2] Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *COLT-23th Conference on Learning Theory-2010*, pages 13–20, 2010.
- [3] Sohail Bahmani, Bhiksha Raj, and Petros T Boufounos. Greedy sparsity-constrained optimization. *The Journal of Machine Learning Research*, 14(1):807–841, 2013.
- [4] Thomas Blumensath and Michael E Davies. Gradient pursuits. *Signal Processing, IEEE Transactions on*, 56(6):2370–2382, 2008.
- [5] Thomas Blumensath and Mike E Davies. Gradient pursuit for non-linear sparse signal modelling. In *European Signal Processing Conference (EUSIPCO)*, pages 25–29, 2008.
- [6] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *Algorithmic Learning Theory*, pages 23–37. Springer, 2009.
- [7] Petros Drineas, Ravi Kannan, and Michael W Mahoney. Fast monte carlo algorithms for matrices i: Approximating matrix multiplication. *SIAM Journal on Computing*, 36(1):132–157, 2006.
- [8] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression (with discussion). *Annals of statistics*, 32(2):407–499, 2004.
- [9] J. Guélat and P. Marcotte. Some comments on wolfe’s away step. *Mathematical Programming*, 35(1), 1986.
- [10] M. Jaggi. Revisiting frank-wolfe : Projection free sparse convex optimization. In *Proceedings of the International Conference on Machine Learning*, 2013.
- [11] Kevin Jamieson and Amee Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. *arXiv preprint arXiv:1502.07943*, 2015.
- [12] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- [13] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1238–1246, 2013.
- [14] L. Laporte, R. Flamary, S. Canu, S. Dejean, and J. Mothe. Nonconvex regularizations for feature selection in ranking with sparse svm. *Neural Networks and Learning Systems, IEEE Transactions*

- on, 25(6):1118–1130, 2014.
- [15] Aurélie C Lozano, Grzegorz Swirszcz, and Naoki Abe. Group orthogonal matching pursuit for logistic regression. In *International Conference on Artificial Intelligence and Statistics*, pages 452–460, 2011.
- [16] S. Mallat and Z. Zhang. Matching pursuit with time-frequency dictionaries. *IEEE Trans Signal Processing*, 41(12):3397–3415, 1993.
- [17] D. Merhej, C. Diab, M. Khalil, and R. Prost. Embedding prior knowledge within compressed sensing by neural networks. *Neural Networks, IEEE Transactions on*, 22(10):1638–1649, Oct 2011.
- [18] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [19] N Nguyen, Deanna Needell, and T. Woolf. Linear convergence of stochastic iterative greedy algorithms with sparse constraints. <http://arxiv.org/abs/1407.0088>, 2014.
- [20] Zhao P and Zhang T. Stochastic optimization with importance sampling. <http://arxiv.org/abs/1401.2753>, 2014.
- [21] Y.C. Pati, R. Rezaifar, and P. Krishnaprasad. Orthogonal matching pursuit : Recursive function approximation with applications to wavelet decomposition. In *Proc. of the 27th Annual Asilomar Conference on Signals, Systems and Computers*, 1993.
- [22] A. Rakotomamonjy. Direct optimization of the dictionary learning problem. *IEEE Trans. on Signal Processing*, 61(12):5495–5506, 2013.
- [23] Alain Rakotomamonjy, Remi Flamary, and Gilles Gasso. Dc proximal newton for nonconvex optimization problems. *IEEE Trans. on Neural Networks and Learning Systems*, 1(1):1–13, 2015.
- [24] S. Shalev-Shwartz, Y. Singer, and N. Srebro. Pegasos : Primal estimated subgradient solver for svm. In *Proceedings of the International Conference on Machine Learning*, pages 807–814, 2007.
- [25] S. Shalev-Shwartz, N. Srebro, and T. Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *Siam Journal on Optimization*, 20, 2010.
- [26] Shai Shalev-Shwartz and Ambuj Tewari. Stochastic methods for l_1 -regularized loss minimization. *The Journal of Machine Learning Research*, 12:1865–1892, 2011.
- [27] Vikas Sindhwani and Aurélie C Lozano. Non-parametric group orthogonal matching pursuit for sparse learning with multiple kernels. In *Advances in Neural Information Processing Systems*, pages 2519–2527, 2011.
- [28] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58(1):267–288, 1996.

- [29] J. Tropp and A. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Information Theory*, 53(12):4655–4666, 2007.
- [30] J. Tropp, A. Gilbert, and M. Strauss. Algorithms for simultaneous sparse approximation. part I: Greedy pursuit. *Signal Processing*, 86:572–588, 2006.
- [31] M. Yaghoobi, T. Blumensath, and M. Davies. Dictionary learning for sparse approximations with the majorization method. *IEEE Transaction on Signal Processing*, 57(6):2178–2191, 2009.
- [32] Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, pages 116–123. ACM, 2004.