



HAL
open science

Regret Analysis of the Stochastic Direct Search Method for Blind Resource Allocation

Juliette Achddou, Olivier Cappe, Aurélien Garivier

► **To cite this version:**

Juliette Achddou, Olivier Cappe, Aurélien Garivier. Regret Analysis of the Stochastic Direct Search Method for Blind Resource Allocation. 2022. hal-03805239

HAL Id: hal-03805239

<https://hal.science/hal-03805239>

Preprint submitted on 10 Oct 2022

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.

Regret Analysis of the Stochastic Direct Search Method for Blind Resource Allocation

Juliette Achddou¹, Olivier Cappé², Aurélien Garivier³

¹ DI ENS, ENS Paris, Université PSL, Numberly

² DI ENS, Inria, ENS Paris, Université PSL, CNRS

³ UMPA, CNRS, Inria, ENS Lyon

Abstract

Motivated by programmatic advertising optimization, we consider the task of sequentially allocating budget across a set of resources. At every time step, a feasible allocation is chosen and only a corresponding random return is observed. The goal is to maximize the cumulative expected sum of returns. This is a realistic model for budget allocation across subdivisions of marketing campaigns, when the objective is to maximize the number of conversions.

We study direct search (aka pattern search) methods for linearly constrained and derivative-free optimization in the presence of noise. Those algorithms are easy to implement and particularly suited to constrained optimization. They have not yet been analyzed from the perspective of cumulative regret. We provide a regret upper-bound of the order of $T^{2/3}$ in the general case. Our mathematical analysis also establishes, as a by-product, time-independent regret bounds in the deterministic, unconstrained case. We also propose an improved version of the method relying on sequential tests to accelerate the identification of descent directions.

Keywords— online learning, regret minimization, derivative-free optimization, direct search.

1 Introduction

1.1 Motivation

In the field of programmatic marketing advertisement, advertisers are given daily budgets that they are required to entirely distribute across a number of predefined subdivisions of a campaign. Their goal is to maximize some notion of cumulative reward on the lifetime of the campaign, corresponding to the number of clicks or purchases generated by the campaign. We posit that the expected reward generated by each subdivision every day is an unknown function of the daily budget allocated to that campaign. We focus on the context in which, when choosing a specific allocation, the advertiser only observes a noisy version of the total reward. The optimization task faced by the advertiser can thus be formalized as a continuous resource allocation problem, under zeroth order and noisy feedback. Indeed, a key operational constraint in this field is the impossibility to directly access higher-order (derivative) information. Furthermore, every noisy evaluation of the objective function has a cost, related to the value of the function, that needs to be accounted for in the performance criterion. In this context, the cumulative reward is thus more relevant than alternative traditional measures of performances based, for instance,

on the distance to the optimum or the norm of the gradient of the objective function reached after some iterations.

1.2 Model

Let us assume that the learner has access to $d + 1 \in \mathbb{N}$ different resources. To keep up with the dominant convention in optimization, we will consider a minimization problem for which the costs may be thought of as minus the rewards. At each round $t \in \{1, \dots, T\}$, the learner is allowed to choose the level of consumption that she makes of each resource, on a continuous scale from 0 to 1. We impose that the consumption levels of all the resources sum to 1. The use of resource $i \in \{1, \dots, d + 1\}$ to a level of $x_t^{(i)}$ generates an expected marginal cost $w_i(x_t^{(i)})$. Overall, the expected one-step cost of the learner is given by $\sum_{i=1}^{d+1} w_i(x_t^{(i)})$, where the set of all possible consumption levels $(x_t^{(1)}, \dots, x_t^{(d+1)})$ correspond to the d dimensional simplex. The goal of the learner is to sequentially minimize the expected cumulative cost over T evaluations of the function, having access only to a noisy version of the expected cost associated to the allocation tried at step t . Not only are the cost functions w_1 to w_{d+1} unknown, but one cannot observe their individual outputs. Although we are primarily motivated by resource allocation, the algorithms studied in this paper may be applied to more general linearly-constrained optimization problems, and their analysis holds in the general case. In the remainder of the paper, we thus consider a generic optimization domain \mathcal{D} that is a subset of \mathbb{R}^d defined by linear constraints: $\mathcal{D} = \{x \in \mathbb{R}^d, A_I x \leq u\}$ with $m \in \mathbb{N}$, $A_I \in \mathbb{R}^{m \times d}$, and $u \in \mathbb{R}^m$. At each round $t \in \{1, \dots, T\}$, the learner selects $x_t \in \mathcal{D}$ and incurs a cost $f(x_t) + \epsilon_t$, where ϵ_t is assumed to be a centered σ -subgaussian noise, with σ known to the learner.

The goal of the learner is to minimize the cumulative cost over T evaluations of the function, or equivalently to *minimize the cumulative regret*

$$R_T = \sum_{t=1}^T f(x_t) - f(x_\star),$$

where we denote by $x_\star \in \mathcal{D} = \arg \min_{x \in \mathcal{D}} f(x)$ the optimal allocation. We make the following regularity assumptions on f .

Assumption 1. *f is continuously differentiable, β -smooth and a -strongly convex on \mathcal{D} .*

We highlight that if f satisfies Assumption 1 then f has a unique minimum and is bounded from below. Assumption 1 is arguably restrictive but common in online optimization due to the difficulty of controlling the cumulative cost without this assumption.

Assuming convexity of f on \mathcal{D} is reasonable for budget allocation, because in programmatic advertising and economics more generally, one commonly observes that the marginal return associated to each resource is diminishing with the level of the resource (following the so-called law of diminishing returns). Indeed, Assumption 1 implies that when f has the form $f(x_t^{(1)}, \dots, x_t^{(d)}) = \sum_{i=1}^d w_i(x_t^{(i)}) + w_{d+1} \left(1 - \sum_{i=1}^d x_t^{(i)}\right)$, the marginal cost functions w_1 to w_{d+1} are also convex and hence satisfy the law of diminishing return (viewing $-w_i$ as the marginal utility associated to the i -th resource).

In this paper, we propose to extend the analysis of direct search methods to the optimization of regular functions. Direct search [Kolda et al., 2003], also known as pattern search, makes use of the well-known fact that if the objective function is continuously differentiable, then at least

one of the directions of a *positive spanning set* (a set that spans the space with non-negative coefficients, abbreviated as PSS in what follows) is a descent direction. It explores the space by testing new points in a number of predefined directions from the current iterate but at a distance that varies with time. The algorithm moves to a new iterate only if it results in a sufficient improvement of the value of the function (there exist other versions of the algorithm where the sufficient decrease condition is replaced by a constraint on the choice of the trial directions). The complexity of such an algorithm has been analyzed in [Vicente, 2013] in the deterministic and unconstrained setting. Lewis and Torczon [2000] study direct search in linearly constrained domains. Handling the constraints in direct search is quite simple, as it consists in testing only the directions in the set of search-directions that are feasible. Gratton et al. [2015] analyzed direct search with random sets of search-directions instead of predefined ones and later extended the analysis to the case of linearly-constrained domains [Gratton et al., 2019]. Dzhahini [2022] extended their work by analyzing a similar algorithm in the presence of noise, but without constraints. Dzhahini [2022] relies on an assumption on the decrease of the noise. In this paper, we will study direct search algorithms that rely on a number of samples at each point to build tight estimates of the function at the trial points, which can be understood as a way to decrease the noise. Dzhahini [2022] analyzes some notion of sample complexity of direct search, which only takes the iterates into account rather than the number of function evaluations needed, which is not fully appropriate in our setting.

1.3 Contributions

Our purpose is to study methods that are suitable for the blind resource allocation model. We choose to study extensions of the direct search algorithm for the optimization of regular functions in the presence of linear constraints and noise. The adaptation of direct search to the noisy case is achieved by performing enough sampling to ensure that the algorithm moves to a new iterate only if it results in a sufficient improvement, with high probability. We propose two ways of doing so: the first method (termed FDS-Plan) simply computes the number of necessary evaluations ahead of time, whereas the second one, FDS-Seq, uses a sequential testing strategy to interrupt sampling as early as possible. The algorithms are specified in Section 2. We analyze the cumulative regret of these algorithms in Section 3, providing an upper bound of their regret of the order of $T^{2/3}$ (up to logarithmic factors), when the optimum is in the interior of the feasible domain. We start Section 3 by the simpler case in which there is neither noise nor constraints, showing that in this basic setup the regret of direct search is bounded by a constant.

2 Algorithms

2.1 Description of the Algorithms

In Algorithm 1 below, we start by describing the most common version of the direct search method, used for deterministic and unconstrained optimization. It requires the setting of an initial point x_0 and an initial parameter α_0 . The learner also specifies a PSS \mathbb{D} , that is, a set of directions that spans \mathbb{R}^d with non-negative coefficients. At each iteration, the algorithm sequentially tests points at a distance α_k from the current iterate and in the directions defined by \mathbb{D} . If none of the test points results in a sufficient decrease of the function's value, the iteration is declared unsuccessful and the trial radius α_k shrinks by a factor $\theta < 1$, otherwise,

the iteration is declared successful and the iterate x_k is moved to the first trial point that results in a sufficient improvement. A decrease is considered to be sufficient if it is larger than some predefined forcing function of α_k , that we take here to be quadratic, with a coefficient that can be set by the learner.

```

1 Choose  $x_0 \in \mathbb{R}^d$ ,  $\alpha_0 > 0$ ,  $\theta \leq 1$ ,  $c > 0$ ,  $\rho(u) = cu^2$  and a PSS  $\mathbb{D}$ 
2 for  $k = 0 \dots K$  do
3   | Set UnsuccessfulSearch  $\leftarrow$  True
4   | for  $v \in \mathbb{D}$  do
5     | Evaluate  $f(x_k + \alpha_k v)$ 
6     | if  $f(x_k) - f(x_k + \alpha_k v) \geq \rho(\alpha_k)$  then
7       |   | Set  $x_{k+1} \leftarrow x_k + \alpha_k v$  and  $\alpha_{k+1} \leftarrow \alpha_k$ 
8       |   | Set UnsuccessfulSearch  $\leftarrow$  False
9       |   | break
10  | if UnsuccessfulSearch then
11  |   | Set  $x_{k+1} \leftarrow x_k$  and  $\alpha_{k+1} \leftarrow \theta \alpha_k$ 
    
```

Algorithm 1: Direct Search with sufficient decrease

The analysis can also be adapted to the presence of a growth factor $\gamma \geq 1$ by which the trial radius α_k expands at successful iterations. To simplify, we choose to focus on the case where $\gamma = 1$, as this parameter does not modify the regret rates obtained in Section 3. Also note that Algorithm 1 is a descent algorithm with respect to the iterates x_k , I.e., the sequence $(f(x_k))_k$ is decreasing.

It is obvious that different choices of PSS result in different trajectories of the algorithm. A common choice of \mathbb{D} as the set of $2d$ vectors of the positive and negative coordinate directions results in the algorithm known as coordinate or compass search. Other frequently considered choices include random directions, as in [Gratton et al., 2015, 2019, Dzahini, 2022].

We apply two sorts of modifications to Algorithm 1 in order to adapt it to the more general model introduced in Section 1.2. The first modification consists in sampling a trial point only if it is feasible. The second modification consists in introducing estimation stages that allow building reliable estimates of f at the trial points. We propose two ways of doing so, that result in two different algorithms. The first algorithm that we study is a plug-in version of Algorithm 1 in which we replace $f(x_k)$ and $f(x_k + \alpha_k v)$ by their empirical estimates, consisting of means computed from $N_k = \frac{32\sigma^2 \log(2/\delta)}{\rho(\alpha_k)^2}$ samples. This number of samples guarantees that with high probability, the estimation gap is smaller than $\rho(\alpha_k)/4$, which in turn ensures that an iteration is declared successful only when it leads to a decrease of $f(x_k)$ by at least $\rho(\alpha_k)/2$ and that an unsuccessful iteration cannot occur if there exists a direction v in \mathbb{D} such that the decrease achieved by moving to $x_k + \alpha_k v$ is larger than $3\rho(\alpha_k)/2$. The resulting algorithm is termed Feasible Direct Search with a planned number of Samples (FDS-Plan) and described in Algorithm 2.

We also propose a faster algorithm, Feasible Direct Search with Sequential Tests (FDS-Seq) described in Algorithm 3. For any $v \in \mathbb{D}$, instead of planning the number of samples at $x_k + \alpha_k v$ ahead of time, it samples at $x_k + \alpha_k v$ and x_k until either

$$\left\{ \begin{array}{l} \left| \hat{f}_{n_{0,k}}(x_k) - \hat{f}_{n_{v,k}}(x_k + \alpha_k v) - \rho(\alpha_k) \right| \leq \sqrt{2\sigma^2 \log(1/\delta) \left(\frac{1}{n_{0,k}} + \frac{1}{n_{v,k}} \right)}, \\ \text{or } (n_{0,k} \geq N_k \text{ and } n_{v,k} \geq N_k), \end{array} \right. \quad (1)$$

```

1 Choose  $x_0 \in \mathbb{R}^d$ ,  $\alpha_0 > 0$ ,  $\theta \leq 1$ ,  $c > 0$ ,  $\rho(u) = cu^2$  and a PSS  $\mathbb{D}$ 
2 for  $k = 0 \dots K$  do
3   | Set UnsuccessfulSearch  $\leftarrow$  True
4   | Set  $N_k = \frac{32\sigma^2 \log(2/\delta)}{\rho(\alpha_k)^2}$ .
5   | Estimate  $f(x_k)$  by making  $N_k$  samples at  $x_k$  and setting
   |    $\hat{f}(x_k) = \frac{1}{N_k} \sum_{j=1}^{N_k} f(x_k) + \epsilon_j$ 
6   | for  $v \in \mathbb{D}$  such that  $x_k + \alpha_k v \in \mathcal{D}$  do
7   |   | Estimate  $f(x_k + \alpha_k v)$  by making  $N_k$  samples at  $x_k + \alpha_k v$  and setting
   |   |    $\hat{f}(x_k + \alpha_k v) = \frac{1}{N_k} \sum_{j=1}^{N_k} f(x_k + \alpha_k v) + \epsilon'_j$ 
8   |   | if  $\hat{f}(x_k) - \hat{f}(x_k + \alpha_k v) \geq \rho(\alpha_k)$  then
9   |   |   | Set  $x_{k+1} \leftarrow x_k + \alpha_k v$  and  $\alpha_{k+1} \leftarrow \alpha_k$ 
10  |   |   | Set UnsuccessfulSearch  $\leftarrow$  False
11  |   |   | break
12  | if UnsuccessfulSearch then
13  |   | Set  $x_{k+1} \leftarrow x_k$  and  $\alpha_{k+1} \leftarrow \theta \alpha_k$ 

```

Algorithm 2: Feasible Direct Search with Planning (FDS-Plan)

where $n_{0,k}$ and $n_{v,k}$ denote the number of samples at x_k and $x_k + \alpha_k v$ and $\hat{f}_{n_{0,k}}(x_k)$ and $\hat{f}_{n_{v,k}}(x_k + \alpha_k v)$ denote the resulting empirical means. Then, successful and unsuccessful iterations are defined as in FDS-Plan, and lead to the same actions.

The sequential stopping rule is designed to achieve early detection of sufficient decrease, but also to detect as early as possible the cases in which the trial point cannot lead to a sufficient decrease. The first test in Condition (1) is a consequence of the fact that the estimation gap at x_k (respectively $x_k + \alpha_k v$) is $\sigma^2/n_{0,k}$ subgaussian (respectively $\sigma^2/n_{v,k}$ subgaussian). The second test of Condition (1) corresponds to a safeguard that prevents from waiting too long in cases where the decrease induced by the trial point is very close to the sufficiency threshold. At worst, the number of evaluations needed is the same as in Algorithm 2. Essentially, the sequential stopping rule reduces the number of evaluations needed for each iteration but maintains the desirable property that with high probability, an iteration is declared successful only if it leads to a decrease of at least $\rho(\alpha_k)/2$ and that an iteration cannot be declared unsuccessful if there exists a direction v in \mathbb{D} such that the decrease achieved by moving to $x_k + \alpha_k v$ is larger than $3\rho(\alpha_k)/2$.

2.2 Illustration

In order to illustrate graphically the behavior of the proposed methods, we show on Figure 1(a) and 1(d) their trajectories in the case where there are 3 resources ($d = 2$) and the loss functions associated to each resource $i \in \{1, \dots, d+1\}$ are of the form $w_i(x) = -\tau_i \frac{\log(1+\gamma x)}{\log(1+\gamma)}$ with $\gamma = 2$, $\tau_1 = 1$, $\tau_2 = 0.45$, and $\tau_3 = 0.95$. We set the horizon to $T = 100,000$ and use a Gaussian noise with standard deviation $\sigma = 0.1$, a realistic value for budget allocation problems. In the symmetric representation of Figure 1, the three vertices correspond to the points where one of the resource is fully saturated (equal to 1) and the edges to linear paths along which one of the resources is set to 0. The contour lines of the target function are materialized by orange lines and the location of the minimum is marked by a black cross. The size of each point is a

```

1 Choose  $x_0 \in \mathbb{R}^d$ ,  $\alpha_0 > 0$ ,  $\delta > 0$ ,  $\theta \geq 1$ ,  $c > 0$ ,  $\rho(u) = cu^2$  and select a PSS  $\mathbb{D}$ 
2 Set  $UnsuccessfulSearch \leftarrow \text{True}$ 
3 for  $k = 0 \dots K$  do
4     for  $v$  in  $\mathbb{D}$  such that  $x_k + \alpha_k v \in \mathcal{D}$  do
5         while Condition (1) is not satisfied do
6             if  $n_{v,k} \leq n_{0,k}$  then
7                 Sample at  $x_k + \alpha_k v$  and update  $n_{v,k}$  and the empirical mean
8                  $\hat{f}_{n_{v,k}}(x_k + \alpha_k v)$ .
9             else
10                Sample at  $x_k$  and update  $n_{0,k}$  and the empirical mean  $\hat{f}_{n_{0,k}}(x_k)$ 
11            if  $\hat{f}_{n_{0,k}}(x_k) - \hat{f}_{n_{v,k}}(x_k + \alpha_k v) \geq \rho(\alpha_k)$  then
12                Set  $x_{k+1} \leftarrow x_k + \alpha_k v$ , and  $\alpha_{k+1} \leftarrow \alpha_k$ .
13                 $UnsuccessfulSearch \leftarrow \text{False}$ 
14                Break.
15 if  $UnsuccessfulSearch$  then
16     Set  $x_{k+1} \leftarrow x_k$  and  $\alpha_{k+1} \leftarrow \theta \alpha_k$ 

```

Algorithm 3: Feasible Direct Search with Sequential Tests (FDS-Seq)

logarithmically growing function of the number of samples made at this point, and its color is a function of the index of the first round at which it has been sampled. Finally, points at which a successful iteration of FDS-Plan (respectively FDS-Seq) occurred are circled in blue. The parameters of both versions of feasible direct search are $\alpha_0 = 0.2$, $c = 5$ and $\theta = 0.7$, and the initial point corresponds to the allocation $x(0) = (1/3, 1/3, 1/3)$ (center of the simplex). The set of directions chosen for these algorithms are the 6 directions that support the edges of the simplex (in both directions).

In a first phase, the algorithm proceeds rapidly by testing directions until a sufficient descent direction is found. Afterwards, when the iterates get closer to the minimizer, the search area has to iteratively shrink and finding descent directions becomes harder. In the first phase, the trajectory is similar to a descent path that would result from a gradient descent algorithm for example, while the second phase is more reminiscent of the behavior of bandit algorithms based on hierarchical partitions, like HOO [Bubeck et al., 2011]. In order to illustrate the differences with such algorithms, we also plot the trajectories of baseline methods, either related to gradient descent or bandits with hierarchical partitioning. The literature related to these methods is reviewed in details in Section 3.3. Before turning to these other algorithms, it is important to note the difference between FDS-Plan and FDS-Seq. Figure 1(d) shows that FDS-Seq is faster than FDS-Plan, as it spends less time on the first iterations, in which it is easy to determine whether the trial points lead to a sufficient decrease. The FDS-Seq algorithm can thus perform more iterations than FDS-Plan.

Let us now comment on Figure 1(b), that represents the trajectory of a version of HOO. It is not straightforward to apply algorithms for \mathcal{X} -armed bandits like HOO on the simplex, because it implies constructing balanced hierarchical partitions of the simplex. In the following, the tree of partitions that we used is built in the following way. We set the parameter ρ of HOO as suggested by Bubeck et al. [2011] to $2^{-2/d}$. A binary tree of depth $H = \frac{\log(1/T)}{2\log(\rho)}$ of partitions of $[0, 1]^d$ is obtained by recursively halving the cells at each depth h of the tree along dimension $h \pmod{2}$. At depth h of the tree, this approach yields a partition formed by rectangular

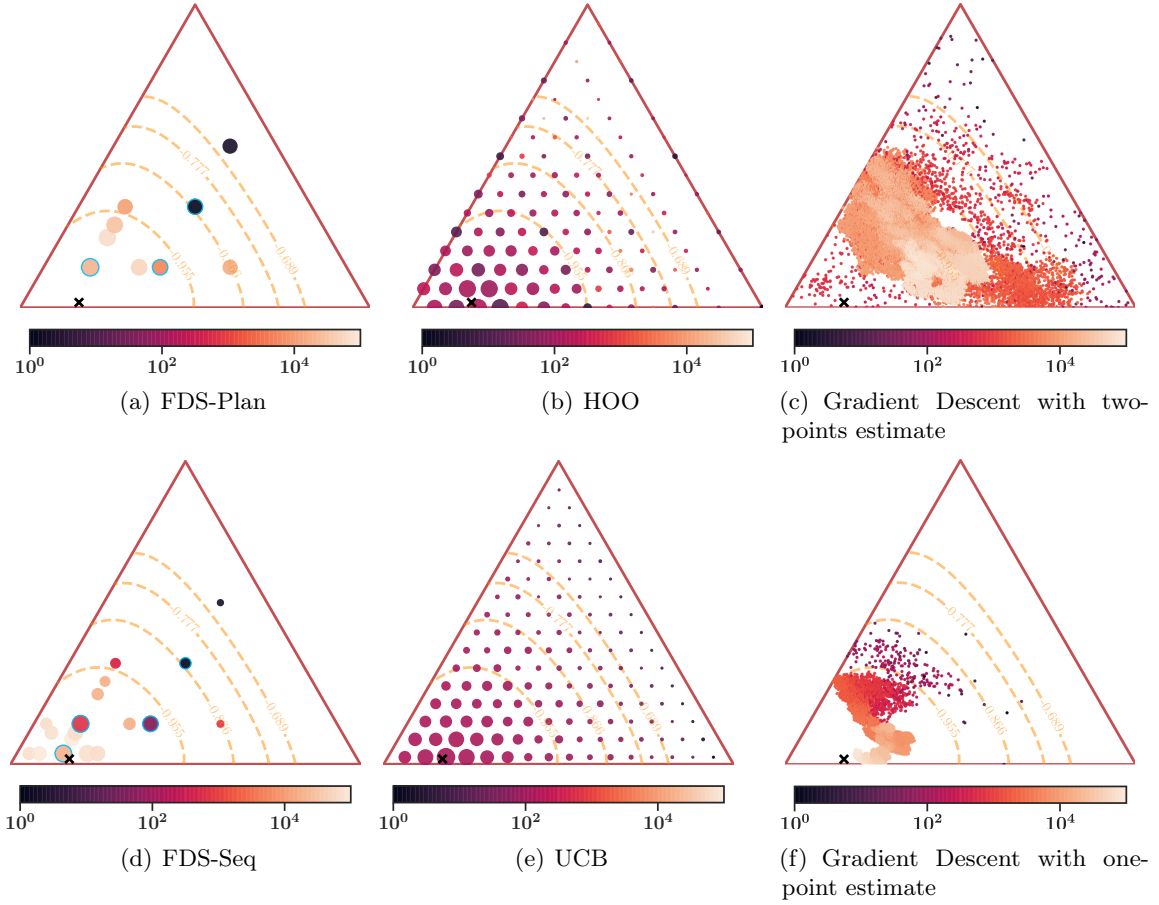


Figure 1: Single trajectories

cells represented by their lower left corner $[a_i, b_i]$. Then, in order to remove unwanted cells, we traverse the tree, starting from the leaves, and remove every cell having an empty intersection with the domain. Due to the geometry of the simplex, knowing if a cell intersects the domain boils down to checking if its representant $[a_i, b_i]$ belongs to it. When the algorithm selects cell (h, i) at time t , the representant of that cell is chosen as a sampling point. In the simulation, the smoothness parameter ν_1 of HOO is set to 16. On Figure 1(b), we observe that this algorithm explores the partition tree in a way that favors the cells close to the optimum while persistently visiting cells that are clearly far from the minimizer. The behavior of algorithms based on direct search can thus be preferred because it makes warm-start possible, in the sense that prior belief on the location of the minimizer can be used for setting the initial point, which is not possible for HOO. The fact that suboptimal points will keep being sampled until the end of the experiment can also be difficult to accept for practitioners such as advertisers for example.

We also illustrate the behavior of UCB on a discretization of the space, which is an interesting strategy, especially in dimension 2. The discretization used for UCB consists of points arranged in a regular grid of $[0, 1]^2$, from which the points lying outside of the feasible domain have been removed. The step parameter of the grid is taken as $T^{-1/4}$ as suggested by Combes and Proutiere [2014]. Although simpler than HOO, this algorithm results in similar sampling patterns, as shown on Figure 1(e), and hence shares some of its drawbacks. The

performance of UCB is good in dimension 2, as the regret can be proved to be of the order of \sqrt{T} with this choice of step-size (see Combes and Proutiere [2014]) but it will worsen in higher dimension due to the difficulty of simultaneously controlling the distance between grid points and the overall number of points in the grid. In fact, the optimal step in this case is of the order of $T^{-1/(d+2)}$ and the resulting regret is of the order of $T^{d/d+2}$.

Lastly, we illustrate the behavior of two methods related to stochastic gradient descent. The first method is related to that proposed by Akhavan et al. [2020]. Without constraints, this method would estimate the gradient of the function at x_t , by evaluating the function at $y_t^+ = x_t + h_t Z_t$ and $y_t^- = x_t - h_t Z_t$, where Z_t is a random vector of the sphere of radius 1, and use $\frac{f(y_t^+) - f(y_t^-)}{h_t} Z_t$ as an estimation of the gradient. The method in itself does not take constraints into account, but a slight modification results in an algorithm that is feasible in the presence of constraints. This modification consists in performing a homothetic perturbation [Bravo et al., 2018] on the evaluation points y_t^+ and y_t^- : instead of using these points, the algorithm evaluates the function at $\tilde{y}_t = y_t + h_t/r(c - x_t)$, where c is a point in the interior of the simplex such that $\mathcal{B}(c, r) \subset \mathcal{D}$. We use $\tilde{y}_t^+ = y_t^+ + h_t/r(c - x_t)$ and $\tilde{y}_t^- = y_t^- - h_t/r(c - x_t)$, where c is a point in the interior of the simplex such that $\mathcal{B}(c, r) \subset \mathcal{D}$. This ensures that the evaluation point belongs to the constrained domain, provided that $x_t \in \mathcal{D}$, but adds a bias which is proportional to h_t , under suitable regularity assumptions on f . To ensure that $x_t \in \mathcal{D}$, we also project the result of the gradient descent step on \mathcal{D} . We use an estimation step h_t equal to $(t/2)^{-1/3}$ and a learning rate decreasing as $1/(2.5t)$. This choice is justified by the following reasoning: with this choice of value for the learning rate and h_t set to $t^{-1/4}$, $\sum_{t=1}^T f(x_t) - f(x^*)$ would be bounded by $O(T^{1/2})$, thanks to the analysis of Akhavan et al. [2020]; but we have to add a term related to the sum of evaluation steps $\sum_{t=1}^{T/2} (f(y_t^+) + f(y_t^-) - 2f(x^*))$ to bound the actual regret, which under suitable assumptions is of order $\sum_{t=1}^T h_t$; so that setting h_t to $(t/2)^{-1/3}$ allows to bound both terms by $O(T^{2/3})$. On Figure 1(c), we see one trajectory of this method when the starting iterate is on the center of the simplex. The convergence speed is rather fast at the beginning but the speed is limited by the homothetic perturbation.

The second method is inspired by Flaxman et al. [2004]. This paper proposes to use gradient descent with a one-point gradient estimation. In order to evaluate the gradient of the function at x_t , the algorithm evaluates the function at $y_t = x_t + h_t Z_t$, where Z_t is a random vector of the sphere of radius 1, and uses $f(y_t)/h_t z_t$ as an estimation of the gradient. As for the previous method, we apply an homothetic perturbation to y_t . We use an estimation step h_t equal to $t^{-1/3}$ and a learning rate decreasing as $1/(2.5t)$. The trajectory that we see on Figure 1(f) is not very indicative of the average performance of the algorithm, since this method comes with a very high variance. We see that the algorithm generates a trajectory that roughly gets closer to the minimizer, but that is far from being a descent path because of the poor estimation of the gradient. This method, which has been designed for adversarially evolving objective functions, is clearly not advisable for static objectives with stochastic perturbations.

3 Regret Analysis

As discussed in the introduction, the regret criterion takes into account the number T of function evaluations instead of focusing on the number K of iterations, as in the more traditional analysis. In the noiseless case of Algorithm 1, T and K differ by a factor of at most $|\mathbb{D}|+1$ and this is not an issue. However, in the case of Algorithms 2 and 3, the situation is very different as the number of function evaluations per iteration is stochastic and typically increases as the

algorithm converges. In this case, it is not possible to predict in advance the evolution of T as a function of K because it depends on the function and starting point. A significant part of the analysis is indeed devoted to quantifying this phenomenon. In practice, it means that in order to comply with a preset number T of function evaluations, the algorithms are run without a fixed number of rounds K , instead they are run until the number of function evaluations reaches T .

In the following, we will analyze the proposed algorithms and show that in the constrained and noisy set-up of interest, FDS-Plan and FDS-Test have a regret of the order of $(\log T)^{2/3}T^{2/3}$ under some further assumptions on f and on the chosen direction set \mathbb{D} , provided that the optimal point lies in the interior of the feasible domain.

To provide some intuition on the proofs, we start with the analysis of Algorithm 1 in the unconstrained and deterministic setting, thereby providing the first regret bound of any direct search algorithm.

3.1 Warm-up: the Unconstrained and Deterministic Setting

The choice of \mathbb{D} is decisive for the performance of the algorithm. In the sequel, we make the following assumption on the chosen PSS.

Assumption 2. *The vectors of \mathbb{D} have unit norm and the cosine measure of \mathbb{D} is lower-bounded, i.e., there exists $\kappa > 0$ such that*

$$cm(\mathbb{D}) := \min_{u \in \mathbb{R}^d, u \neq 0} \max_{v \in \mathbb{D}} \frac{u^T v}{\|u\| \|v\|} > \kappa.$$

Assumption 2 guarantees in particular that at each iteration k , the cosine similarity between at least one direction in \mathbb{D} and $-\nabla f(x_k)$ is larger than κ . Note that Assumption 2 is common and that as long as \mathbb{D} is a PSS then there exists a κ such that it is satisfied.

Theorem 1. *Under Assumptions 1 and 2, the cumulative regret of Algorithm 1 satisfies*

$$R_T \leq (|\mathbb{D}|+1) \left[\left(\frac{1}{c} \left(\frac{1}{1-\theta^2} (f(x_0) - f(x_*) + \rho(\alpha_0)) \right) \right) \left(\left(1 + \frac{\eta}{a}\right) \eta + \beta \right) + \frac{f(x_0) - f(x_*)}{\rho(\alpha_0)} \left(\frac{\beta}{a} \|\nabla f(x_0)\| \alpha_0 + \beta \alpha_0^2 \right) \right],$$

where $\eta := \frac{\beta}{a} \frac{1}{\kappa \theta} (c + \frac{\beta}{2})$.

This result shows that under Assumption 1, the asymptotic behavior of the regret of direct search can be compared to that of the more traditional gradient descent algorithm, whose regret is also bounded under this assumption [see Theorem 3.6 of Bubeck et al., 2008].

3.1.1 Elements of Proof

The proof of Theorem 1 combines two well-known properties of direct search and the fact that under Assumptions 1, the norm of the gradient evaluated at x_k cannot increase by more than a factor of $\frac{\beta}{a}$ as stated in the following Lemma, which is true for any descent algorithm.

Lemma 1. *If f satisfies Assumption 1, then $\forall k' > k, \|\nabla f(x_{k'})\| \leq \frac{\beta}{a} \|\nabla f(x_k)\|$.*

In the proof of Theorem 1, Lemma 1 is used in conjunction with the following well-known property [see e.g. Vicente, 2013] of direct search.

Lemma 2. *If f satisfies Assumption 1 and 2 and iteration k corresponds to a unsuccessful iteration then $\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + \frac{\rho(\alpha_k)}{\alpha_k} \right) = \frac{1}{\kappa} \left(\frac{\beta}{2} + c \right) \alpha_k$.*

The above lemma follows from the definition of the cosine measure of \mathbb{D} , that results in a bound of $v_k^T(-\nabla f(x_k))$, where v_k is the direction in \mathbb{D} maximizing the gap $f(x_k) - f(x_k - \alpha_k v)$, and from the smoothness assumption on f . Thanks to Lemma 1, this lemma also means that when iteration k is unsuccessful, we can bound all subsequent gradients by $\frac{\beta}{\kappa a} \left(\frac{\beta}{2} + c \right) \alpha_k$. We can deduce that for any k' following the first unsuccessful iteration, $\|\nabla f(x_{k'})\| \leq \eta \alpha_{k'}$, where $\eta := \frac{\beta}{a} \frac{1}{\kappa \theta} \left(c + \frac{\beta}{2} \right)$. Indeed, if k' is the index of an unsuccessful iteration, Lemma 1 suffices to prove $\|\nabla f(x_{k'})\| \leq \eta \theta \alpha_{k'}$. In contrast, when k' is the index of a successful iteration, one should consider the last unsuccessful iteration k . Since $\alpha_{k'} \geq \theta \alpha_k$, there has been at most one reduction of the step-size since iteration k and

$$\|\nabla f(x_{k'})\| \leq \frac{\beta}{\kappa a \theta} \left(\frac{\beta}{2} + c \right) \alpha_{k'} = \eta \alpha_{k'} . \quad (2)$$

The following general argument on direct search is the final key element of the proof of Theorem 1, that links the sum of the squared search-radius to the initial sub-optimality gap $f(x_0) - f(x_*)$.

Lemma 3. *If f satisfies Assumptions 1 and 2,*

$$\sum_{k=0}^{\infty} \rho(\alpha_k) = \sum_{k=0}^{\infty} c \alpha_k^2 \leq \frac{1}{1 - \theta^2} (f(x_0) - f(x_*) + \rho(\alpha_0)) .$$

Lemma 3 can be explained by the fact that $\rho(\alpha_k)$ decreases geometrically by a ratio θ^2 between two successive successful iterations, so that the contribution to the sum of these iterations boils down to multiplying the remainder of the sum by a factor of $\frac{1}{1 - \theta^2}$. The sum on successful iterations cannot be too large, because by definition of successful iterations, $f(x_k) - f(x_*)$ is lower bounded by this sum plus the initial sub-optimality gap $f(x_0) - f(x_*)$. Bringing Lemma 3 and the bound of Equation (2) together results in a bound on the squared norm of the gradients $\|\nabla f(x_k)\|^2$ after the first unsuccessful iteration. Using the regularity conditions of Assumption 1 suffices to relate the regret to the squared norm of the gradients, which in turn yields Theorem 1. The complete proof can be found in Appendix A.

3.2 The Constrained and Noisy Setting

We now turn to the noisy and constrained case, described in Section 1.2. We further impose the following assumptions on the domain.

Assumption 3. *\mathcal{D} is contained in a ball of radius b .*

This assumption, together with assumption 1, implies that $f(x) - f(x_*)$ is bounded. For the constrained case, we need an additional regularity assumption, also required by Gratton et al. [2019].

Assumption 4. *∇f is bounded in norm by a constant B on the feasible set \mathcal{D} .*

While in the unconstrained case, the chosen PSS \mathbb{D} only needed to satisfy Assumption 2, a stronger assumption is required in the presence of linear constraints. Indeed, Assumption 2 was a way to ensure that there was at least one trial direction v in \mathbb{D} satisfying $\frac{-\nabla f(x_k)^T v}{\|v\| \|\nabla f(x_k)\|} \geq \kappa$. This property is not sufficient in the unconstrained case, because in this case, the directions of interest at iteration k in \mathbb{D} are those that are feasible. A problem that might arise for example, is that a sufficient descent direction is not detected even in a situation where $x_k - \alpha_k \nabla f(x_k)$ is feasible, because the set of feasible directions in \mathbb{D}_k does not positively span the feasible region. To avoid such cases, we impose a constraint on \mathbb{D} that involves the notion of approximate tangent cones. Approximate tangent cones at a point x are the polar cones of the cones that are generated by the α -binding constraints at x as defined by Kolda et al. [2007].

Let a_i^T be the i -th row of the constraint matrix A_I and let $\mathcal{C}_i = \{y, \text{ s.t. } a_i^T y = u_i\}$ denote the sets where the i -th constraint are binding. If there exists a point of \mathcal{C}_i at a distance smaller than α from x , then the i -th constraint is said to be α -binding. The indices of α -binding constraints are x are denoted $I(x, \alpha) = \{i, \text{ dist}(x, \mathcal{C}_i) \leq \alpha\}$, where dist is induced by the Euclidian distance. We define the approximate normal cone $N(x, \alpha)$ to be the cone generated by the set $\{a_i, \text{ s.t. } i \in I(x, \alpha)\} \cup \{0\}$. The approximate tangent cone $T(x, \alpha)$ is the polar of $N(x, \alpha)$, which means that $T(x, \alpha) = \{v : y^T v \leq 0, \forall y \in N(x, \alpha)\}$. Informally $T(x, \alpha)$ is the cone inside of the boundaries generated by the α -binding constraints at x . We highlight that since the number of constraints m is finite, there can only be a finite number, smaller than 2^m , of tangent cones. Consequently, Assumption 5 is rather mild.

Assumption 5. \mathbb{D} contains a set \mathcal{G}_k of positively generating directions of $T(x, \alpha)$ included in $T(x, \alpha)$, for any $x \in \mathcal{D}$ and $\alpha \in \mathbb{R}_+$.

This assumption was already necessary in [Kolda et al., 2003] and required by Lewis and Torczon [2000], while in [Gratton et al., 2019], the descent set at iteration k is assumed to be contained in $T(x, \alpha)$ and to generate it. We explain the purpose of this assumption in the following, While the purpose of Assumption 2 was to ensure that the maximal cosine similarity of a vector in \mathbb{D} with $-\nabla f(x_k)$ was bounded away from 0, we focus on a different kind of measure of similarity to $-\nabla f(x_k)$ given by

$$\begin{cases} \max_{v \in \mathcal{G}_k} \frac{-\nabla f(x_k)^T v}{\|P_{T(x, \alpha)}(-\nabla f(x_k))\| \|v\|} & \text{if } P_{T(x, \alpha)}(-\nabla f(x_k)) \neq 0, \\ 1 & \text{otherwise.} \end{cases}$$

If $P_{T(x, \alpha)}(-\nabla f(x_k))$ gets close to 0, this measure of similarity to $-\nabla f(x_k)$ necessarily becomes small, although $-\nabla f(x_k)^T v$ is small for any v in \mathcal{G}_k . In order to bound this measure of similarity between a vector of \mathcal{G}_k and $-\nabla f(x_k)$ away from 0, we define the following approximate cosine measure:

$$cm_{T(x_k, \alpha_k)}(\mathcal{G}_k) := \inf_{u \in \mathbb{R}^d, P_{T(x_k, \alpha_k)}(u) \neq 0} \max_{v \in \mathcal{G}_k} \frac{u^T v}{\|P_{T(x_k, \alpha_k)}(u)\| \|v\|}.$$

As proved by Lewis and Torczon [2000] and recalled by Gratton et al. [2019], if \mathcal{C} is a set of cones c_j that are respectively positively generated from a set of vectors $G(c_j)$, then

$$\lambda(\mathcal{C}) := \min_{c_j \in \mathcal{C}} \left\{ \inf_{u \in \mathbb{R}^d, P_{c_j}(u) \neq 0} \max_{v \in G(c_j)} \frac{u^T v}{\|P_{c_j}(u)\| \|v\|} \right\} > 0,$$

which guarantees that $cm_{T(x_k, \alpha_k)}(\mathcal{G}_k)$ is bounded by some constant $\kappa_{min} > 0$. Under the above assumptions, we can bound the regret of FDS-Plan.

Theorem 2. *Under Assumptions 1, 3, 4, and 5, and if $x_* \in \text{int}(\mathcal{D})$, the cumulative regret R_T of FDS-Plan (respectively FDS-Seq) after the first T evaluations of f satisfies*

$$\mathbb{E}(R_T) = O(\log(T)^{2/3}T^{2/3})$$

for the choice $\delta = T^{-4/3}$ (respectively $\delta = T^{-10/3}$ for FDS-Seq).

In the absence of a lower bound, the optimality of such a regret rate is unsure. It is difficult to compare it to other known bounds, as the performance of related algorithms is often not evaluated in the same way. In particular, the performance of the version of stochastic gradient descent proposed by Akhavan et al. [2020] is analyzed with respect to a different notion of regret, \tilde{R}_T , that does not take into account the samples needed for the estimation of each gradient. Their analysis yields $\tilde{R}_T = O(\sqrt{T})$. It is important to note that the algorithm by Akhavan et al. [2020] takes advantage of the fact that in the setting of the latter paper, sampling points outside of the feasible domain is possible. When using homothetic perturbation as we did for the illustration in Figure 1(c), the regret of such a method is of the order of $T^{2/3}$, as explained in Section 2.2.

Black-box algorithms such as HOO or StoOO [Bubeck et al., 2011, Munos, 2014] require the domain to be bounded. When given a balanced hierarchical partition of \mathcal{X} , and the smoothness of the function around its optimum, these algorithms would incur a regret of the order of \sqrt{T} . The regret rate of FDS-Plan appears to be larger than that of HOO instantiated with the right parameters. However, HOO relies on a partition of the feasible domain that is computationally difficult to achieve with arbitrary linearly constrained domains.

In the case of resource allocation, the feasible domain is the simplex and one way of satisfying Assumption 5 is to set \mathbb{D} to the set of vectors corresponding to edges of the simplex, which amounts to $2d(d-1)$ directions. The regret bound is proportional to the number of directions contained in \mathbb{D} (see the proof of Theorem 2, in Appendix C.2), so that the dependence of the regret with respect to d is quadratic. This is rather detrimental to FDS algorithms and could be improved by selecting random subsets of directions, as in [Gratton et al., 2019].

3.2.1 Elements of Proof

After some finite number of iterations that depends on $\Delta := \min_{i \in 1 \dots m} \text{dist}(x_*, \mathcal{C}_i)$, the distance from x_k to the boundaries of \mathcal{D} is smaller than $\Delta/4$ with high probability, thanks to the analysis of Gratton et al. [2019]. Waiting for another number of iterations, α_k gets small enough for the approximate tangent cone $T(x_k, \alpha_k)$ to describe the whole space \mathbb{R}_d . Then, the trajectory of the algorithm is the same as in the unconstrained setting. In the unconstrained setting, the following elements provide an intuition of why the regret is of the order of $T^{2/3}$. With similar arguments to those of the proof of Theorem 1, i.e Lemmas 1, 2 and 3, it is easy to see that the instantaneous regret incurred at iteration k of the algorithm is proportional to the sum of α_k^{-2} (up to logarithmic factors), whereas it was proportional to α_k^2 in the deterministic case: indeed, iteration k now involves N_k times more evaluations than in the noiseless case and N_k is proportional to α_k^{-4} (up to logarithmic factors). Thanks to Lemma 3, we know that α_k^2 is summable. Then, thanks to Hölder's inequality applied to the sum of α_k^{-2} written as $(\alpha_k)^{2/3} (\alpha_k)^{-2/3-2}$, the regret is proportional to the total number of evaluations to the power of $2/3$, up to logarithmic factors. The complete proof can be found in Appendix C.

3.3 Related Works

The discrete counterpart of the resource allocation model in which the resources can only be used up to discrete consumption levels, is a celebrated model of operations research with multiple applications. Its properties have first been discussed by Koopman [Koopman, 1953] who proposed the first algorithmic solution for this problem. Koopman’s works have further been extended by Gross [1956], Katoh et al. [1979] who propose more efficient algorithms under specific assumptions on the number of resources and the total consumption budget. The range of applications is wide, including experimental design, load management in an industrial context, computer scheduling and, more recently, the *adwords* problem introduced by Mehta et al. [2007]. Recently, Agrawal and Devanur [2015] studied online and offline resource allocation, motivated by the latter task. With the same motivation, Fontaine et al. [2020] focus on the online and continuous version of resource allocation in which the learner accesses the derivatives. The method studied by Fontaine et al. [2020] extends the bisection method in dimension $d > 1$.

The more general problem of derivative-free optimization in noisy environments has been considered by researchers coming from different horizons. A relevant stream of works originates from the bandit community, which considered this task as an extension of the more traditional multi-armed bandit problem [see e.g Auer et al., 2002]. The class of \mathcal{X} -armed bandits models focuses on the case where a learner can select actions in a generic measurable space and the mean-payoff function is regular. In [Bubeck et al., 2011], for example, the mean payoff function is supposed to be locally Lipschitz with respect to some dissimilarity measure. Bubeck et al. [2011] and Munos [2014] adopt the approach of hierarchical optimization, in which the optimization domain is iteratively partitioned, resulting in finer and finer partitions, that are required to be balanced in some sense. The learner maintains an upper confidence bound of the goal function that is constant on each cell defined by the finest partition. The algorithm proposed by Bubeck et al. [2011] achieves a regret of the order of \sqrt{T} when the learner knows the exact order of the smoothness at the optimum. However, partitioning the domain in a hierarchical and balanced way is relatively easy when the domain is an hypercube, but is a computational problem in itself when the domain has a more complex form. We also mention that knowing the smoothness is considered a challenge most of the time in black-box optimization, so that several methods have been introduced that are adaptive to the smoothness (Locatelli and Carpentier [2018], Valko et al. [2013], Shang et al. [2019]). We mention that concurrently to HOO based on hierarchical partitioning, Agarwal et al. [2011] has also proposed a different strategy for \mathcal{X} -armed bandits, but this time convex, with ellipsoid methods, that also result in $O(\sqrt{T})$ regret.

The extension of more traditional first-order optimization methods has also been considered. In the deterministic case in which the function evaluation is not perturbed by any noise, Nesterov and Spokoiny [2017] consider random gradient descent based on finite differences to estimate the gradient. Flaxman et al. [2004] later considers a version of stochastic gradient descent with a one-point estimate of the gradient for the adversarial setting introduced by Zinkevich [2003] in which at each time step, a new goal function is chosen by an adversary, making it impossible to rely on a two-point estimate of the gradient. In this setting, Flaxman et al. [2004] show an adversarial regret bound of the order of $T^{5/6}$. Later on, Hazan and Levy [2014], Hazan and Li [2016], Bubeck et al. [2017] propose new methods for the same setting, but with adversarial convex or strongly-convex functions, showing improved regret bounds, as low as \sqrt{T} . In a stochastic setting that is closer to ours, Akhavan et al. [2020], Bach and

Perchet [2016] consider a version of stochastic gradient descent with unbiased estimates of the gradient, obtained by finite differences. While they provide an analysis in term of the regret, they focus on a restricted notion of regret that is different from the one considered in this work. The algorithm that they propose relies on a number of samples used to estimate the gradient at each iteration of the gradient descent algorithm. But the regret only accounts for the cost incurred by the iterates of the gradient descent algorithm and ignores the regret incurred by the samples used for the estimation of the gradient. Moreover, the constraints also do not apply to those samples, meaning that the algorithm is allowed to get samples outside of the feasible domain in order to estimate the gradient, which is impossible in the application of interest. The authors prove an upper bound on their version of the regret, which is of the order of \sqrt{T} when Assumption 1 is satisfied.

Finally, a number of heuristics and algorithms have been studied for derivative-free optimization without noise and constraints. To cite a few, evolution strategies [Beyer and Schwefel, 2002] mimic evolution dynamics and include the class of genetic algorithms; the Nelder-Mead technique [Lagarias et al., 1998] is a heuristic search method that maintains a set of $d + 1$ trial points arranged as a simplex and relies on the extrapolation of the objective function measured at these trial points; simulated annealing [Chibante, 2010] has a more probabilistic flavor, since it decides on the next move in a probabilistic fashion.

4 Conclusion

We have studied feasible direct search algorithms designed for linearly-constrained zeroth-order optimization in a stochastic setting. We have shown that these algorithms, though being fairly simple, suffer a regret of the order of $T^{2/3}$, which is quite satisfactory when compared to other algorithms, like those inspired by bandit algorithms, or by gradient descent schemes. Even though continuously-armed bandit approaches seem to yield better performances when implementable, direct search algorithms have the advantage of allowing the learner to set the initial point and the initial step-size, which can be particularly interesting when the learner has a some prior knowledge on the minimizer of the objective function.

References

- Alekh Agarwal, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. *Advances in Neural Information Processing Systems*, 24, 2011.
- Shipra Agrawal and Nikhil R. Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, 2015.
- Arya Akhavan, Massimiliano Pontil, and Alexandre Tsybakov. Exploiting higher order smoothness in derivative-free optimization and continuous bandits. *Advances in Neural Information Processing Systems*, 33:9017–9027, 2020.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.

- Francis Bach and Vianney Perchet. Highly-smooth zero-th order online optimization. In *Conference on Learning Theory*, pages 257–283. PMLR, 2016.
- Hans-Georg Beyer and Hans-Paul Schwefel. Evolution strategies—a comprehensive introduction. *Natural computing*, 1(1):3–52, 2002.
- Mario Bravo, David Leslie, and Panayotis Mertikopoulos. Bandit learning in concave n-person games. *Advances in Neural Information Processing Systems*, 31, 2018.
- Sébastien Bubeck, Gilles Stoltz, Csaba Szepesvári, and Rémi Munos. Online optimization in x-armed bandits. *Advances in Neural Information Processing Systems*, 21, 2008.
- Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 72–85, 2017.
- Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12(5v), 2011.
- Rui Chibante. *Simulated annealing: theory with applications*. BoD—Books on Demand, 2010.
- Richard Combes and Alexandre Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pages 521–529. PMLR, 2014.
- Kwassi Joseph Dzahini. Expected complexity analysis of stochastic direct-search. *Computational Optimization and Applications*, 81(1):179–200, 2022.
- Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*, 2004.
- Xavier Fontaine, Shie Mannor, and Vianney Perchet. An adaptive stochastic optimization algorithm for resource allocation. In *Algorithmic Learning Theory*, pages 319–363. PMLR, 2020.
- Serge Gratton, Clément W Royer, Luís Nunes Vicente, and Zaikun Zhang. Direct search based on probabilistic descent. *SIAM Journal on Optimization*, 25(3):1515–1541, 2015.
- Serge Gratton, Clément W Royer, Luís Nunes Vicente, and Zaikun Zhang. Direct search based on probabilistic feasible descent for bound and linearly constrained problems. *Computational Optimization and Applications*, 72(3):525–559, 2019.
- O Gross. A class of discrete-type minimization problems. Technical report, 1956.
- Elad Hazan and Kfir Levy. Bandit convex optimization: Towards tight bounds. *Advances in Neural Information Processing Systems*, 27, 2014.
- Elad Hazan and Yuanzhi Li. An optimal algorithm for bandit convex optimization. *arXiv preprint arXiv:1603.04350*, 2016.
- Naoki Katoh, Toshihide Ibaraki, and Hisashi Mine. A polynomial time algorithm for the resource allocation problem with a convex objective function. *Journal of the Operational Research Society*, 30(5):449–455, 1979.

- Tamara G Kolda, Robert Michael Lewis, and Virginia Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482, 2003.
- Tamara G Kolda, Robert Michael Lewis, and Virginia Torczon. Stationarity results for generating set search for linearly constrained optimization. *SIAM Journal on Optimization*, 17(4):943–968, 2007.
- Bernard O Koopman. The optimum distribution of effort. *Journal of the Operations Research Society of America*, 1(2):52–63, 1953.
- Jeffrey C Lagarias, James A Reeds, Margaret H Wright, and Paul E Wright. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on optimization*, 9(1):112–147, 1998.
- Robert Michael Lewis and Virginia Torczon. Pattern search methods for linearly constrained minimization. *SIAM Journal on Optimization*, 10(3):917–941, 2000.
- Andrea Locatelli and Alexandra Carpentier. Adaptivity to smoothness in x-armed bandits. In *Conference on Learning Theory*, pages 1463–1492. PMLR, 2018.
- Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.
- Rémi Munos. From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning. 2014.
- Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.
- Xuedong Shang, Emilie Kaufmann, and Michal Valko. General parallel optimization a without metric. In *Algorithmic Learning Theory*, pages 762–788. PMLR, 2019.
- Michal Valko, Alexandra Carpentier, and Rémi Munos. Stochastic simultaneous optimistic optimization. In *International Conference on Machine Learning*, pages 19–27. PMLR, 2013.
- Luís Nunes Vicente. Worst case complexity of direct search. *EURO Journal on Computational Optimization*, 1(1):143–153, 2013.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936, 2003.

Supplementary Material

Outline. We prove in Appendix A all the results pertaining to the noiseless, unconstrained case. In Appendix B, we provide an additional result for the noisy but unconstrained case. The analysis of direct search in the latter case paves the way for the proof of Theorem 2 whose proof is deferred to Appendix C.

A Deterministic and unconstrained set-up

A.1 Preliminary Results

Lemma 1. *If f satisfies Assumption 1,*

$$\forall k' > k, \|\nabla f(x_{k'})\| \leq \frac{\beta}{a} \|\nabla f(x_k)\|$$

Proof. First observe that because of strong convexity $(\nabla f(x_k) - \nabla f(x_*))^\top (x_k - x_*) \geq a\|x_k - x_*\|^2$, $a\|x_k - x_*\|^2 \leq \|\nabla f(x_k)\| \|x_k - x_*\|$ which implies that

$$a\|x_k - x_*\| \leq \|\nabla f(x_k)\|. \quad (3)$$

Hence,

$$\begin{aligned} \|\nabla f(x_{k'})\| &\leq \beta \|x_{k'} - x_*\| \\ &\leq \frac{\beta}{\sqrt{a}} \sqrt{f(x_{k'}) - f(x_*)} \leq \frac{\beta}{\sqrt{a}} \sqrt{f(x_k) - f(x_*)} \\ &\leq \frac{\beta}{\sqrt{a}} \sqrt{\nabla f(x_k)^\top (x_k - x_*)} \leq \frac{\beta}{\sqrt{a}} \sqrt{\|\nabla f(x_k)\| \|x_k - x_*\|} \\ &\leq \frac{\beta}{a} \|\nabla f(x_k)\|, \end{aligned}$$

where the first inequality comes from the smoothness (in fact $\|\nabla f(x_k)\| = \|\nabla f(x_k) - \nabla f(x_*)\| \leq \beta \|x_k - x_*\|$), the second inequality comes from the strong convexity, the third one comes from the fact that the algorithm is a descent, the fourth one comes from convexity and the fifth one comes from the strong convexity property of Equation (3). \square

Lemma 2. *If f satisfies Assumption 1 and iteration k corresponds to an unsuccessful iteration then*

$$\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + \frac{\rho(\alpha_k)}{\alpha_k} \right) = \frac{1}{\kappa} \left(\frac{\beta}{2} + c \right) \alpha_k.$$

This lemma is already well-known [see e.g. Vicente, 2013], we only prove it here for completeness.

Proof. Since $cm(\mathbb{D}) := \min_{v \in \mathbb{R}^d} \max_{v \in \mathbb{D}} \frac{v^\top v}{\|v\| \|v\|} > \kappa$, there exists $v \in \mathbb{D}$ such that

$$-f(x_k)^\top v \geq \kappa \|\nabla f(x_k)\|.$$

Since the iteration is an unsuccessful iteration, we have $f(x_k) - f(x_k + \alpha_k v) \leq \rho(\alpha_k) = c\alpha_k^2$. Then

$$\begin{aligned} \kappa\alpha_k\|\nabla f(x_k)\| - \rho(\alpha_k) &\leq -f(x_k)^\top v + f(x_k + \alpha_k v) - f(x_k) \\ &\leq \int_0^{\alpha_k} \nabla f(x_k + uv)^\top v - \nabla f(x_k)^\top v du \\ &\leq \int_0^{\alpha_k} \|\nabla f(x_k + uv) - \nabla f(x_k)\| \|v\| du \\ &\leq \beta \int_0^{\alpha_k} u du \leq \frac{\beta}{2} \alpha_k^2, \end{aligned}$$

which yields $\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + \frac{\rho(\alpha_k)}{\alpha_k} \right) = \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + c\alpha_k \right)$. \square

Lemma 3.

$$\sum_{k=0}^{\infty} \rho(\alpha_k) \leq \frac{1}{1 - \theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)).$$

This lemma is also a common element of the analysis of direct search analysis [see e.g. Gratton et al., 2019], we only prove it here for completeness.

We assume that there are infinitely many successful iterations as it is trivial to adapt the argument otherwise. Let k_i be the index of the i th successful iteration ($i \geq 1$). Define $k_0 = -1$ and $\alpha_{-1} = \alpha_0$ for convenience. Let us rewrite $\sum_{k=0}^{\infty} \rho(\alpha_k)$ as $\sum_{i=0}^{\infty} \sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k)$ and study first $\sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k)$. Thanks to the definition of the update on a successful iteration and on unsuccessful iterations,

$$\sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k) = \sum_{k=k_i+1}^{k_{i+1}} \rho(\theta^i \alpha_{k_i}) = \sum_{k=k_i+1}^{k_{i+1}} c\theta^{2i} \rho(\alpha_{k_i}) \leq \frac{1}{1 - \theta^2} \rho(\alpha_{k_i}).$$

Since on successes $\rho(\alpha_{k_i}) \leq f(x_{k_i}) - f(x_{k_i} + \alpha_{k_i}) = f(x_{k_i}) - f(x_{k_{i+1}})$,

$$\sum_{i=1}^{\infty} \rho(\alpha_{k_i}) \leq f(x_1) - f(x_*).$$

Hence

$$\sum_{k=0}^{\infty} \rho(\alpha_k) \leq \frac{1}{1 - \theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)).$$

Lemma 4. *The index k_f of the first unsuccessful iteration satisfies:*

$$k_f \leq \frac{f(x_0) - f(x_*)}{\rho(\alpha_0)}.$$

Another version of this lemma is due to Gratton et al. [2015].

Proof. Before the first unsuccessful iteration, $\alpha_k = \alpha_0$. So by definition of a successful iteration, $\forall 0 < k \leq k_f$

$$f(x_{k-1}) - f(x_k) \geq \rho(\alpha_0).$$

By summing,

$$f(x_{k_0}) - f(x_{k_f}) \geq k_f \rho(\alpha_0).$$

The left hand-side of this inequality is upper-bounded by $f(x_0) - f(x_*)$, which suffices to conclude the proof. \square

A.2 Regret Bound

In this section, we prove in Theorem 3 below a result involving the regret at iteration K of the algorithm instead of the regret after T function evaluations. As $T \leq K(|\mathbb{D}|+1)$, Theorem 3 directly implies Theorem 1.

Consider

$$\tilde{R}_K = \sum_{k=0}^K \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right)$$

which is an upper bound of the cumulative regret suffered by the algorithm at iteration K . \tilde{R}_K can be bounded as follows.

Theorem 3. *Under Assumptions 1 and 2,*

$$\begin{aligned} \tilde{R}_K \leq & (|\mathbb{D}|+1) \left[\left(\frac{1}{c} \left(\frac{1}{1-\theta^2} (f(x_0) - f(x_*) + \rho(\alpha_0)) \right) \right) \left(\left(1 + \frac{\eta}{a}\right) \eta + \beta \right) \right. \\ & \left. + \frac{f(x_0) - f(x_*)}{\rho(\alpha_0)} \left(\frac{\beta}{a} \|\nabla f(x_0)\| \alpha_0 + \beta \alpha_0^2 \right) \right], \end{aligned}$$

where $\eta := \frac{\beta}{a} \frac{1}{\kappa\theta} (c + \frac{\beta}{2})$.

Proof. We decompose the regret as

$$\begin{aligned} \tilde{R}_K &= \sum_{k=0}^K \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\leq \sum_{k=0}^K \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) + f(x_k) - f(x_*) \right) \\ &\leq \sum_{k=0}^K ((|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=0}^K \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \\ &\leq \sum_{k=0}^{k_f} ((|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=0}^{k_f} \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \\ &\quad + \sum_{k=k_f}^K ((|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=k_f}^K \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right), \end{aligned} \quad (4)$$

where k_f is the iteration of the first unsuccessful iteration. The third inequality provides a decomposition of the regret in a first term that involves the suboptimality of the iterate, and a second term that involves the difference between values of f at the iterate and at the trial points. As is usual for direct search algorithm, the behavior of the algorithm before the first unsuccessful iteration has to be studied separately, which explains the use of the decomposition of the fourth inequality. We bound the regret due to the rounds preceding k_f by:

Lemma 5. *The regret due to the rounds preceding k_f is bounded by*

$$\sum_{k=0}^{k_f} ((|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=0}^{k_f} \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \leq C_1,$$

where we denote by k_f is the index of the first unsuccessful iteration and by $C_1 = \frac{f(x_0) - f(x_*)}{c\alpha_0^2} \left((f(x_0) - f(x_*) + \frac{\beta}{a} \|\nabla f(x_0)\| \alpha_0 + \beta \alpha_0^2) \right)$.

Proof. As until k_f , $f(x_k) \leq f(x_0)$ and $\alpha_k = \alpha_0$, it holds that

$$\begin{aligned} & \sum_{k=0}^{k_f} ((|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=0}^{k_f} \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \\ & \leq \sum_{k=0}^{k_f} ((|\mathbb{D}|+1)(f(x_0) - f(x_*))) + \sum_{k=0}^{k_f} \left(\sum_{v \in \mathbb{D}} \|\nabla f(x_k)\| \alpha_k + \beta \alpha_k^2 \right) \\ & \leq k_f ((|\mathbb{D}|+1)(f(x_0) - f(x_*))) + (|\mathbb{D}|+1) k_f \left(\frac{\beta}{a} \|\nabla f(x_0)\| \alpha_0 + \beta \alpha_0^2 \right) \\ & \leq \frac{f(x_0) - f(x_*)}{c\alpha_0^2} \left((|\mathbb{D}|+1)(f(x_0) - f(x_*) + \frac{\beta}{a} \|\nabla f(x_0)\| \alpha_0 + \beta \alpha_0^2) \right) \\ & = C_1, \end{aligned}$$

where Lemma 1 is used for the second inequality and Lemma 4 is used for the third inequality. The first inequality uses the following property of convex and β -smooth functions: $f(y) - f(x) \leq \nabla f(y)^T(x - y) \leq \|\nabla f(x)^T(x - y)\| + \beta \|x - y\|^2$, applied to $x_k + \alpha_k v$ and x_k . \square

Lemma 6. After k_f ,

$$\sum_{k=k_f}^K f(x_k + \alpha_k v) - f(x_k) \leq C_2$$

where k_f is the index of the first unsuccessful iteration and

$$C_2 = \frac{1}{c} (\eta + \beta) \left(\frac{1}{1-\theta^2} (f(x_0) - f(x_*) + \rho(\alpha_0)) \right).$$

Proof. We take $k > k_f$. Using the property of convex and β -smooth functions that $f(y) - f(x) \leq \|\nabla f(x)^T(x - y)\| + \beta \|x - y\|^2$, applied to $x_k + \alpha_k v$ and x_k , as in Lemma 5, we get

$$f(x_k + \alpha_k v) - f(x_k) \leq \alpha_k \|\nabla f(x_k)\| + \beta \alpha_k^2.$$

We note that if k is the index of an unsuccessful iteration,

$$\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(c + \frac{\beta}{2} \right) \alpha_k = \frac{1}{L'_1} \alpha_k,$$

by Lemma 2, if $\alpha_k \leq 1$. If k is the index of a successful iteration, we can come back to the last unsuccessful iteration k' , since

$$\|\nabla f(x_k)\| \leq \frac{\beta}{a} \|\nabla f(x_{k'})\| \leq \frac{\beta}{a} \frac{1}{\kappa} \left(c + \frac{\beta}{2} \right) \alpha_{k'} \leq \frac{\beta}{a} \frac{1}{\kappa} \left(c + \frac{\beta}{2} \right) \frac{\alpha_k}{\theta}.$$

where the first inequality comes from Lemma 1, and the third from the fact that $\alpha_k \geq \theta \alpha_{k'}$. Hence for any $k > k_f$,

$$\|\nabla f(x_k)\| \leq \eta \alpha_k,$$

and

$$f(x_k + \alpha_k v) - f(x_k) \leq (\eta + \beta) \alpha_k .$$

Hence

$$\sum_{k=0}^K f(x_k + \alpha_k v) - f(x_k) \leq \sum_{k=0}^K (\eta + \beta) \alpha_k^2 .$$

Consequently,

$$\begin{aligned} \sum_{k=0}^K f(x_k + \alpha_k v) - f(x_k) &\leq (\eta + \beta) \frac{1}{c} \left(\sum_{k=0}^{\infty} \rho(\alpha_k) \right) \\ &\leq \frac{1}{c} (\eta + \beta) \left(\frac{1}{1 - \theta^2} (f(x_0) - f(x_*) + \rho(\alpha_0)) \right) . \end{aligned}$$

□

Lemma 7.

$$\sum_{k=k_f}^K (f(x_k) - f(x_*)) \leq \frac{1}{ac} \eta^2 \left(\frac{1}{1 - \theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right) := C_3$$

Proof. Take $k > k_f$. Thanks to the convexity of f ,

$$\begin{aligned} f(x_k) - f(x_*) &\leq \nabla f(x_k)^\top (x_k - x_*) \\ &\leq \frac{1}{a} \|\nabla f(x_k)\|^2 , \end{aligned}$$

where the second inequality stems from Equation (3), which itself come from strong convexity.

As in the proof of Lemma 6 we have for any $k > k_f$,

$$\|\nabla f(x_k)\| \leq \eta \alpha_k ,$$

so that for any $k > k_f$,

$$f(x_k) - f(x_*) \leq \frac{1}{a} (\eta)^2 \left(\frac{\alpha_k}{\theta} \right)^2 .$$

Thanks to Lemma 3, we have $\sum_{k=0}^{\infty} \alpha_k^2 \leq \left(\frac{1}{c} \frac{1}{1 - \theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right)$.

Eventually,

$$f(x_k) - f(x_*) \leq \frac{1}{a} \eta^2 \left(\frac{1}{c} \frac{1}{1 - \theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right) = C_3 .$$

□

Using the regret decomposition of Equation (4) together with Lemmas 4, 6, and 7 completes the proof of Theorem 3. □

B Noisy and unconstrained set-up

Before considering the constrained setting, we analyze the algorithms described in Section 2 (Algorithms 2 and 3) when there are no constraints, that is, $\mathcal{D} = \mathbb{R}^d$.

B.1 Presentation of the main Result

Theorem 4. *Assume that f is lower bounded and upper bounded on \mathbb{R}^d , so that there exists U , $f(x) - f(x_*) \leq U$, $\forall x \in \mathbb{R}^d$. Also assume that the region \mathcal{X} where $f(x) < f(x_0)$ is convex and that f is a -strongly convex and β -smooth on \mathcal{X} . Let R_T be the cumulative regret on the T first evaluations of f made by FDS-Plan. Set $\delta = T^{-4/3}$. Then*

$$\mathbb{E}(R_T) = O(\log(T)^{2/3}T^{2/3})$$

This regret bound is also valid for FDS-Seq under the same Assumptions, with $\delta = T^{-10/7}/2$. In the following, we give a proof of the regret bound for FDS-Plan. Note that Sections B.3 and B.2 refer to FDS-Plan, and Section B.4 deals with FDS-Seq.

The regularity assumption in Theorem 4 requires that f is bounded and satisfies a local version of Assumption 1. The initial point x_0 should not be chosen too far from x^* , nor should α_0 be too large. This assumption is not unreasonable, since for every x_0 and α_0 , it is naturally satisfied by bounded and strictly convex functions in \mathcal{C}^2 for some choice of a and β . We stress that under the alternative assumption 1, the same kind of regret bound could still be proved, but with a smaller choice of δ , resulting in higher confidence bonuses and the multiplication of the regret by some constant factor. Indeed, in this case, estimating f incorrectly at each round can lead to a trajectory that always deviates from x^* , which is highly detrimental to the regret rate; meanwhile, under the assumption required by Theorem 4, f is bounded by U , so that deviating from x^* contributes to the regret by at most UT .

In the following, we will use the following additional notations.

Notation. We define v_k to be

$$v_k := \begin{cases} \arg \max_{v \in \mathbb{D}} f(x_k) - f(x_k - \alpha_k v) & \text{if iteration } k \text{ is unsuccessful} \\ \text{the chosen direction} & \text{otherwise.} \end{cases}$$

B.2 Intermediate results

Lemma 8. *We call \mathcal{E}_k the event*

$$\mathcal{E}_k = \{|f(x_k + \alpha_k v) - \hat{f}(x_k + \alpha_k v)| \leq c/4(\alpha_k)^2\}, \forall v \in \mathbb{D} \cup \{0\}.$$

The probability of \mathcal{E}_k is lower bounded by

$$\mathbb{P}(\mathcal{E}_k | \mathcal{F}_{k-1}) \geq 1 - \delta(|\mathbb{D}|+1)$$

where \mathcal{F}_{k-1} is the σ -field representing the history.

Proof. Let $v \in \mathbb{D} \cup \{0\}$ $f(x_k + \alpha_k v) - \hat{f}(x_k + \alpha_k v) = \sum_{i=1}^{N_k} \epsilon_j$ with ϵ_j independent Gaussian variables with variance σ^2 and we have

$$\left| \sum_{i=1}^{N_j} \epsilon_j \right| \leq \sqrt{2\sigma^2 \frac{\log(2/\delta)}{N_j}} \leq \sqrt{\frac{2\sigma^2 \log(2/\delta)}{32\sigma^2 \log(2/\delta) / \rho(\alpha_k)^2}} \leq \frac{\rho(\alpha_k)}{4}$$

with probability $1 - \delta$, when knowing N_k . By a union bound, $\mathbb{P}(\mathcal{E}_k | \mathcal{F}_{k-1}) \geq 1 - \delta(|\mathbb{D}|+1)$ where \mathcal{F}_{k-1} is the σ -field representing the history. \square

The following lemma characterizes unsuccessful iterations and successes when \mathcal{E}_k occurs.

Lemma 9. *On \mathcal{E}_k , if k is an unsuccessful iteration then $f(x_k) - f(x_k + \alpha_k v_k) \leq 3c/2(\alpha_k)^2$ and if k is a successful iteration then $f(x_k) - f(x_k + \alpha_k v_k) \geq c/2(\alpha_k)^2$.*

Lemma 9 implies that if \mathcal{E}_k occurs for all k , then each iteration of the algorithm results in a descent.

Lemma 10. *On $\cap_{k \leq K} \mathcal{E}_k$, the algorithm is a descent algorithm. In particular, $x_k \in \mathcal{X}$, $\forall k \in 1 \dots K$.*

Lemma 11. *If f satisfies the assumptions of Theorem 4 and on $\cap_{k \leq K} \mathcal{E}_k$ then,*

$$\forall k' > k, \|\nabla f(x_{k'})\| \leq \frac{\beta}{a} \|\nabla f(x_k)\|.$$

Proof. The proof of Lemma 1 applies verbatim thanks to Lemma 10. \square

Lemma 12. *If f satisfies the assumptions of Theorem 4 and the iteration k corresponds to an unsuccessful iteration then on \mathcal{E}_k ,*

$$\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + \frac{3\rho(\alpha_k)}{2\alpha_k} \right) = \frac{1}{2\kappa} (\beta\alpha_k + 3c\alpha_k^2).$$

Proof. We reproduce the proof of Lemma 2 by using Lemma 9.

Since $cm(\mathbb{D}) := \min_{v \in \mathbb{R}^d} \max_{v \in \mathbb{D}} \frac{v^T v}{\|v\| \|v\|} > \kappa$, there exists $v \in \mathbb{D}$ such that

$$-f(x_k)^\top v \geq \kappa \|\nabla f(x_k)\|.$$

Since the iteration is an unsuccessful iteration, we have $f(x_k) - f(x_k + \alpha_k v) \leq \frac{3}{2}\rho(\alpha_k) = \frac{3}{2}c\alpha_k^2$, thanks to Lemma 9. Then

$$\kappa\alpha_k \|\nabla f(x_k)\| - \rho(\alpha_k) \leq \frac{\beta}{2}\alpha_k^2,$$

exactly as in the proof of Lemma 2, which yields $\|\nabla f(x_k)\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2}\alpha_k + \frac{3}{2}\frac{\rho(\alpha_k)}{\alpha_k} \right) = \frac{1}{\kappa} \left(\frac{\beta}{2}\alpha_k + \frac{3}{2}c\alpha_k \right)$. \square

Lemma 13. *If f satisfies the assumptions of Theorem 4 and on $\cap_{k \leq K} \mathcal{E}_k$,*

$$\sum_{k=0}^K \rho(\alpha_k) \leq \frac{2}{1-\theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)).$$

Assume that $\cap_{k \leq K} \mathcal{E}_k$ holds. Let k_i be the index of the i -th successful iteration ($i \geq 1$). Define $k_0 = -1$ and $\alpha_{-1} = \alpha_0$, and $\alpha_k = 0$, $\forall k > K$ for convenience. Define K_I the number of successes until K . We rewrite $\sum_{k=0}^{K_I} \rho(\alpha_k)$ as $\sum_{i=0}^{K_I} \sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k)$ and study first $\sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k)$.

Thanks to the definition of the update on a successful iteration and on unsuccessful iterations,

$$\sum_{k=k_i+1}^{k_{i+1}} \rho(\alpha_k) \leq \frac{1}{1-\theta^2} \rho(\alpha_{k_i}).$$

exactly as in the proof of Lemma 3. Since on successes,

$$\frac{1}{2} \rho(\alpha_{k_i}) \leq f(x_{k_i}) - f(x_{k_i} + \alpha_{k_i}) = f(x_{k_i}) - f(x_{k_{i+1}}),$$

we have

$$\frac{1}{2} \sum_{i=1}^{K_I} \rho(\alpha_{k_i}) \leq f(x_1) - f(x_*).$$

Hence

$$\sum_{k=0}^{K_I} \rho(\alpha_k) \leq \frac{2}{1-\theta^2} (f(x_1) - f(x_*) + \rho(\alpha_0)).$$

Lemma 14. *On $\cap_{k \leq K} \mathcal{E}_k$, the first iteration that results in an unsuccessful iteration occurs at round k_f , satisfying:*

$$k_f \leq 2 \frac{f(x_0) - f(x_*)}{\rho(\alpha_0)}.$$

Before the first unsuccessful iteration, $\alpha_k = \alpha_0$. So by Lemma 9, $\forall 0 < k \leq k_f$

$$f(x_{k-1}) - f(x_k) \geq \frac{1}{2} \rho(\alpha_0).$$

By summing,

$$f(x_{k_0}) - f(x_{k_f}) \geq \frac{1}{2} k_f \rho(\alpha_0).$$

The left hand-side of this inequality is upper-bounded by $f(x_0) - f(x_*)$, which suffices to conclude the proof.

Lemma 15. *If f satisfies the assumptions of Theorem 4 and on $\cap_{k \leq K} \mathcal{E}_k$, for any k after the first unsuccessful iteration,*

$$\|\nabla f(x_k)\| \leq \eta_2 \alpha_k,$$

where we denote by $\eta_2 = \frac{\beta}{2a\kappa\theta} (\beta\alpha_k + 3c\alpha_k^2)$.

Proof. For unsuccessful iterations,

$$\|\nabla f(x_k)\| \leq \frac{1}{2\kappa} (3c + \beta) \alpha_k,$$

thanks to Lemma 12. If k is the index of a successful iteration, we can come back to the last unsuccessful iteration k' , since

$$\|\nabla f(x_k)\| \leq \frac{\beta}{a} \|\nabla f(x_{k'})\| \leq \frac{\beta}{a} \frac{1}{2\kappa} (3c + \beta) \alpha_{k'} \leq \frac{\beta}{a} \frac{1}{2\kappa} (3c + \beta) \left(\frac{\alpha_k}{\theta} \right),$$

where the first inequality comes from Lemma 11, the second from Lemma 12 and the third from the fact that $\alpha_k \geq \theta \alpha_{k'}$. \square

B.3 Regret Analysis of FDS-Plan

Lemma 16. *If f satisfies the assumptions of Theorem 4 and on $\cap_{k \leq K} \mathcal{E}_k$,*

$$\tilde{R}_K \leq C_4 \log(2/\delta) + C_5 \log(2/\delta) \left(\sum_{k=1}^K N_k \right)^{2/3},$$

where $C_4 = \frac{32}{c^2} C_1 \alpha_0^{-4} (|\mathbb{D}|+1)$, and

$$C_5 = \frac{32}{c^2} (|\mathbb{D}|+1) \left(\frac{c}{32} \frac{1}{(1-\theta^2)} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right)^{1/3} \left(\frac{1}{a} \eta_2^2 + \eta_2 + \beta \right).$$

Proof. In the following we study the case where $\cap_{k \leq K} \mathcal{E}_k$ holds true.

As in the deterministic case, we decompose the regret as

$$\begin{aligned} \tilde{R}_K &= \sum_{k=0}^{k_f} N_k \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\quad + \sum_{k=k_f}^K N_k \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right). \end{aligned}$$

We start by dealing with the cumulative regret before k_f . We write

$$\begin{aligned} &\sum_{k=0}^{k_f} N_k \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\leq N_0 \sum_{k=0}^{k_f} \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\leq \frac{32}{c^2} \alpha_0^{-4} \log(2/\delta) \sum_{k=0}^{k_f} \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\leq \frac{32}{c^2} \alpha_0^{-4} \log(2/\delta) \times 2C_1 \\ &= C_4 \log(2/\delta), \end{aligned}$$

where the last inequality is obtained exactly as in the proof of Lemma 5 with the help of Lemma 14 instead of Lemma 4. By using the above inequality and the decomposition of the

regret, we get

$$\begin{aligned}
 & \tilde{R}_K - C_4 \log(2/\delta) \\
 & \leq \sum_{k=k_f}^K N_k \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\
 & \leq \sum_{k=k_f}^K N_k \left(f(x_k) - f(x_*) + \sum_{v \in \mathbb{D}} f(X_k + \alpha_k v) - f(x_k) + f(x_k) - f(x_*) \right) \\
 & \leq \sum_{k=k_f}^K N_k \left((|\mathbb{D}|+1)(f(x_k) - f(x_*)) + \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \\
 & \leq (|\mathbb{D}|+1) \sum_{k=k_f}^K N_k \left(\frac{1}{a} \|\nabla f(x_k)\|^2 + \|\nabla f(x_k)\| \alpha_k + \beta \alpha_k^2 \right),
 \end{aligned}$$

where $C_4 = \frac{32}{c^2} C_1 \alpha_0^{-4} (|\mathbb{D}|+1)$. The fourth inequality comes from the regularity assumptions required for Theorem 2 together with Lemma 10. We use Lemma 15 to get that for any $k > k_f$,

$$\|\nabla f(x_k)\| \leq \eta_2 \alpha_k.$$

Then

$$\frac{1}{a} \|\nabla f(x_k)\|^2 + \|\nabla f(x_k)\| \alpha_k + \beta \alpha_k^2 \leq \frac{1}{a} \eta_2^2 \alpha_k^2 + \eta_2 \alpha_k^2 + \beta \alpha_k^2.$$

We get

$$\begin{aligned}
 \sum_{k=k_f}^K N_k \left(\frac{1}{a} \|\nabla f(x_k)\|^2 + \|\nabla f(x_k)\| \alpha_k \right) & \leq C_6 \log(2/\delta) \sum_{k=k_f}^K (\alpha_k)^{-2} \\
 & \leq C_6 \log(2/\delta) \sum_{k=0}^K (\alpha_k)^{-2}.
 \end{aligned}$$

where $C_6 = (\frac{1}{a} \eta_2^2 + \eta_2 + \beta) \frac{32}{c^2}$. Consequently

$$\tilde{R}_K \leq C_4 \log(2/\delta) + C_6 \log(2/\delta) \sum_{k=0}^K (\alpha_k)^{-2}.$$

The number of function evaluations is defined as

$$\sum_{k=0}^K N_k = \frac{32 \log(2/\delta)}{c} \sum_{k=0}^K \alpha_k^{-4}.$$

Thanks to Lemma 13,

$$\sum_{k=0}^K (\alpha_k)^2 \leq \frac{2}{c(1-\theta^2)} (f(x_1) - f(x_*) + \rho(\alpha_0)).$$

By Hölder's inequality, we get

$$\begin{aligned}
 \sum_{k=0}^K (\alpha_k)^{-2} &= \sum_{k=0}^K (\alpha_k)^{2/3} (\alpha_k)^{-2/3-2} \\
 &\leq \left(\sum_{k=0}^K (\alpha_k)^{2/3 \times 3} \right)^{1/3} \left(\sum_{k=0}^K (\alpha_k)^{-32/3 \times 3/2} \right)^{2/3} \\
 &\leq \left(\sum_{k=0}^K (\alpha_k)^2 \right)^{1/3} \left(\sum_{k=0}^K (\alpha_k)^{-4} \right)^{2/3}.
 \end{aligned}$$

And thus

$$\begin{aligned}
 &\tilde{R}_K - C_4 \log(2/\delta) \\
 &\leq C_6 \log(2/\delta) \left(\frac{2}{c(1-\theta^2)} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right)^{1/3} \left(\frac{c}{8 \log(2/\delta)} \sum_{k=1}^K N_k \right)^{2/3} \\
 &\leq C_5 \log(2/\delta) \left(\frac{1}{\log(2/\delta)} \sum_{k=1}^K N_k \right)^{2/3} \\
 &\leq C_5 \log(2/\delta)^{1/3} \left(\sum_{k=1}^K N_k \right)^{2/3},
 \end{aligned}$$

where $C_5 = (|\mathbb{D}|+1)C_6 \left(\frac{c}{32(1-\theta^2)} (f(x_1) - f(x_*) + \rho(\alpha_0)) \right)^{1/3}$.

□

Proof. of Theorem 4 We note K_T the last round reached by the algorithm with T evaluations. Lemma 16 proves that on $\cap_{k \leq K_T} \mathcal{E}_k$,

$$\tilde{R}_{K_T} \leq C_4 \log(2/\delta) + C_5 \left(\frac{1}{(|\mathbb{D}|+1)} \right)^{2/3} \log(2/\delta)^{2/3} (T)^{2/3}.$$

Thanks to Lemma 8,

$$\mathbb{P} \left(\cup_{k=1}^{K_T} \mathcal{E}_k^C \right) \leq \sum_{k=1}^T \mathbb{P}(\mathcal{E}_k^C) \leq (|\mathbb{D}|+1) \sum_{t=1}^T T^{-4/3} / 2 \leq (|\mathbb{D}|+1) T^{-1/3}, \quad (5)$$

when taking $\delta = T^{-4/3}$, since $K_T \leq T$. Hence,

$$\begin{aligned}
 \mathbb{E}(R_T) &\leq \frac{4}{3} C_4 \log(2T) + \frac{4}{3} \left(\frac{1}{(|\mathbb{D}|+1)} \right)^{2/3} C_5 \log(2T)^{2/3} T^{2/3} + (|\mathbb{D}|+1) U T^{2/3} \\
 &= O((\log T)^{2/3} T^{2/3}).
 \end{aligned}$$

□

B.4 Regret Analysis of FDS-Seq

Instead of considering \mathcal{E}_k as in the previous section, we need to consider $\mathcal{E}'_k = \{(f(x_k) - f(x_k + \alpha_k v_k) \leq 3c/2(\alpha_k)^2 \text{ and } k \text{ is an unsuccessful iteration}) \text{ or } (k \text{ is a successful iteration and } f(x_k) - f(x_k + \alpha_k v_k) \geq c/2(\alpha_k)^2)\}$. Instead of Lemma 8 we prove the following result.

Lemma 17. *The probability of \mathcal{E}'_k is lower bounded by*

$$\mathbb{P}(\mathcal{E}'_k | \mathcal{F}_{k-1}) \geq 1 - \delta \times N_k^2 \times (|\mathbb{D}|+1),$$

where \mathcal{F}_{k-1} is the σ -field representing the history.

Proof. Fix $v \in \mathbb{D}$. First assume that $f(x_k) > f(x_k + \alpha_k v) + 3/2\rho(\alpha_k)$. In particular, $f(x_k) > f(x_k + \alpha_k v) + \rho(\alpha_k)$. We denote $n_{0,k}^\tau$ and $n_{v,k}^\tau$ the values of $n_{0,k}$ and $n_{v,k}$ at the end of the while loop of FDS-Seq. Observe that

$$\begin{aligned} \mathbb{P}(\mathcal{E}'_k | \mathcal{F}_{k-1}, n_{0,k}^\tau = n_{v,k}^\tau = N_k) \\ \geq 1 - \delta \times (|\mathbb{D}|+1). \end{aligned}$$

Hence we only need to focus on the case when the first part of Condition (1) is first satisfied. In this case, knowing $n_{0,k}^\tau, n_{v,k}^\tau$, the probability that $\hat{f}_{n_{0,k}^\tau}(x_k) < \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v) + \rho(\alpha_k)$ when the first row of Condition (1) is first satisfied is bounded as follows. We have

$$\begin{aligned} \text{and } 0100 \text{ and } 0101 \\ \mathbb{P}\left(f(x_k) - \hat{f}_{n_{0,k}^\tau}(x_k) - (f(x_k + \alpha_k v) - \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v)) \geq \right. \\ \left. \sqrt{2\sigma^2 \log(1/\delta)} \sqrt{\frac{1}{n_{0,k}^\tau} + \frac{1}{n_{v,k}^\tau}} \middle| \mathcal{F}_k, n_{0,k}^\tau, n_{v,k}^\tau\right) \leq \delta. \end{aligned}$$

Since $f(x_k) > f(x_k + \alpha_k v) + \rho(\alpha_k)$,

$$\begin{aligned} f(x_k) - \hat{f}_{n_{0,k}^\tau}(x_k) - (f(x_k + \alpha_k v) - \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v)) \\ \geq -\rho(\alpha_k) - \hat{f}_{n_{0,k}^\tau}(x_k) + \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v). \end{aligned}$$

So that the above deviation bound results in:

$$\begin{aligned} \text{and } 0100 \text{ and } 0101 \\ \mathbb{P}\left(\hat{f}_{n_{0,k}^\tau}(x_k) - \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v) - \rho(\alpha_k) \leq \right. \\ \left. -\sqrt{2\sigma^2 \log(1/\delta)} \sqrt{\frac{1}{n_{0,k}^\tau} + \frac{1}{n_{v,k}^\tau}} \middle| \mathcal{F}_k, n_{0,k}^\tau, n_{v,k}^\tau\right) \end{aligned}$$

Finally, we apply a union bound. Since $n_{0,k}^\tau$ and $n_{v,k}^\tau$ both belong to $[0, N_k]$ and cannot be simultaneously equal to N_k :

$$\begin{aligned} \text{and } 0100 \text{ and } 0101 \\ \mathbb{P}\left(\hat{f}_{n_{0,k}^\tau}(x_k) - \hat{f}_{n_{v,k}^\tau}(x_k + \alpha_k v) - \rho(\alpha_k) \leq \right. \\ \left. -\sqrt{2\sigma^2 \log(1/\delta)} \sqrt{\frac{1}{n_{0,k}^\tau} + \frac{1}{n_{v,k}^\tau}} \middle| \mathcal{F}_k, \text{not}(n_{0,k}^\tau = n_{v,k}^\tau = N_k)\right) \leq (N_k^2 - 1)\delta. \end{aligned}$$

This amounts to a bound of the probability of k being an unsuccessful iteration and thus of \mathcal{E}'_k , when the first part of condition 1 is satisfied. Summing with the probability of \mathcal{E}'_k in the other case, we obtain $\mathbb{P}(\mathcal{E}'(k)) \geq N_k^2 \delta$.

The case $f(x_k) > f(x_k + \alpha_k v) + 3/2\rho(\alpha_k)$ can be treated in the exact same way. \square

To adapt the proof of Theorem 4 to FDS-Seq (with a different choice of $\delta = T^{-10/3}$), it suffices to replace Equation (5) by

$$\begin{aligned} \mathbb{P}\left(\bigcup_{k=1}^{K_T} \mathcal{E}'_k\right) &\leq (|\mathbb{D}|+1) \sum_{k=1}^T \mathbb{P}\left(\mathcal{E}'_k\right) \leq (|\mathbb{D}|+1) \sum_{k=1}^{K_T} N_k^2 T^{-10/3} \\ &\leq (|\mathbb{D}|+1) \sum_{k=1}^T T^2 T^{-10/3} / 2 \leq (|\mathbb{D}|+1) T^{-1/3}. \end{aligned}$$

The regret is hence

$$\begin{aligned} \mathbb{E}(R_T) &\leq \frac{10}{3} C_4 \log(2T) + \frac{10}{3} \left(\frac{1}{(|\mathbb{D}|+1)}\right)^{2/3} C_5 \log(2T)^{2/3} T^{2/3} + (|\mathbb{D}|+1) U T^{2/3} \\ &= O((\log T)^{2/3} T^{2/3}). \end{aligned}$$

C Noisy and Constrained Set-Up

In this section we analyze the behavior of the algorithms in the presence of linear constraints. The complexity of feasible direct search with linear constraints in the noiseless case has been studied by Gratton et al. [2019]. In this paper, instead of studying the speed at which the gradient converges to 0 as is usual in the unconstrained case, the authors study the convergence of a lower bound of the gradient $\chi(x) := \max_{x+v \in \mathcal{D}, \|v\| \leq 1} -\nabla f(x)^T v$ to 0. Indeed, the convergence of the gradient to 0 might be unachievable when the optimum lies on the boundaries, but $\chi(x)$ is equal to 0 if and only if x is optimal. The paper proves that the first iteration K_ϵ at which $\chi(x_k)$ is smaller than ϵ is of the order of ϵ^{-2} , like in the unconstrained case.

In the following, we denote by U the global upper bound of $f(x) - f(x_*)$ on the domain.

C.1 Intermediate Results

We recall that \mathcal{E}_k denotes the event

$$\mathcal{E}_k = \{|f(x_k + \alpha_k v) - \hat{f}(x_k + \alpha_k v)| \leq c/4(\alpha_k)^2\}, \forall v \in \mathbb{D} \cup \{0\}.$$

Lemmas 8, 9, 10 are left unchanged by the transition to constrained domains. A version of Lemma 3.4. of [Gratton et al., 2019] reads :

Lemma 18. *On $\cap_{k \leq K} \mathcal{E}_k$, and if f satisfies Assumptions 4, and 1, then the following holds: if the k -th iteration is unsuccessful, then*

$$\chi(x_k) \leq \left(\frac{\beta}{2\kappa} + \frac{Bg}{\eta_{\min}}\right) \alpha_k + \frac{3\rho(\alpha_k)}{2\kappa\alpha_k} := L_2 \alpha_k,$$

where $\eta_{\min} := \lambda(N)$ where N is the set of all possible approximate normal cones $N(x, \alpha)$, $\forall \in \mathcal{D}$, $\alpha \in \mathbb{R}^d$.

Proof. It is straightforward to prove

$$\|P_{T(x_k, \alpha_k)}(-\nabla f(x_k))\| \leq \frac{1}{\kappa} \left(\frac{\beta}{2} \alpha_k + \frac{3\rho(\alpha_k)}{2\alpha_k} \right) = \frac{1}{2\kappa} (\beta\alpha_k + 3c\alpha_k^2),$$

with the same elements as in the proof Lemma 11, by noticing that \mathbb{D} contains \mathcal{G}_k , that generates $T(x_k, \alpha_k)$. To prove the bound on $\chi(x_k)$, we use the Moreau decomposition, stating that any vector $v \in \mathbb{R}^d$ can be decomposed as $v = P_{T_k}[v] + P_{N_k}[v]$ with $N_k = N(x_k, \alpha_k)$ and $P_{T_k}[v] \perp P_{N_k}[v] = 0$, and write

$$\begin{aligned} \chi(x_k) &= \max_{x+v \in \mathcal{D}, \|v\| \leq 1} (P_{T_k}[-\nabla f(x_k)] + (P_{T_k}[v]v^T + P_{N_k}[v])^T P_{N_k}[-\nabla f(x_k)]) \\ &\leq \max_{x+v \in \mathcal{D}, \|v\| \leq 1} (v^T P_{T_k}[-\nabla f(x_k)] + P_{N_k}[v]^T P_{N_k}[-\nabla f(x_k)]) \\ &\leq \max_{x+v \in \mathcal{D}, \|v\| \leq 1} \|P_{T_k}[-\nabla f(x_k)]\| + \|P_{N_k}[v]\| \|P_{N_k}[-\nabla f(x_k)]\|. \end{aligned} \quad (6)$$

The first term of the right hand side of Equation (6) is bounded in the following way

$$\|P_{T(x_k, \alpha_k)}(-\nabla f(x_k))\| \leq \frac{1}{2\kappa} (\beta\alpha_k + 3c\alpha_k^2)$$

consequently.

Lemma 19 (Proposition B.1 of [Lewis and Torczon, 2000]). *Let $x \in \mathcal{D}$ and $\alpha > 0$. Then, for any vector v such that $x + v \in \mathcal{D}$, one has*

$$\|P_{N(x, \alpha)}[v]\| \leq \frac{\alpha}{\eta_{\min}}.$$

This in turn provides a bound of the second term of the right hand side of Equation (6):

$$\|P_{N_k}[v]\| \|P_{N_k}[-\nabla f(x_k)]\| \leq \frac{\alpha}{\eta_{\min}} B_g,$$

which suffices to conclude the proof. \square

Lemma 20. *Assume that $\cap_{k \leq K} \mathcal{E}_k$ holds, and f satisfies Assumptions 4, and 1. Set $\epsilon > 0$. Let h denote the mapping from ϵ to*

$$h(\epsilon) = \left(\frac{2L_2^2 U}{c\theta^1} \right) \epsilon^{-2} + \frac{\log(\frac{\alpha_0 L_2}{\theta})}{\log(1/\theta)} + \frac{2U}{\alpha_0^2} = E_1 \epsilon^{-2} + E_2,$$

where $E_1 = \left(\frac{2L_2^2 U}{c\theta^1} \right)$ and $E_2 = \frac{\log(\frac{\alpha_0 L_2}{\theta})}{\log(1/\theta)} + \frac{2U}{\alpha_0^2}$. Denote by $k(\epsilon)$ the first iteration of the algorithm where $\chi(x_k) \leq \epsilon$. If $h(\epsilon) \leq K$, then

$$k(\epsilon) \leq h(\epsilon).$$

The proof is a mere adaptation of the proof of Theorem 1 of [Gratton et al., 2015], with different constants (we use Lemma 18 and 13).

Lemma 21. *If f satisfies Assumption 1,*

$$a\|x_k - x_*\| \leq \chi(x_k).$$

Proof.

$$a\|x_k - x_*\| \leq \frac{f(x_k) - f(x_*)}{\|x_k - x_*\|} \leq \frac{-\nabla f(x_k)(x_* - x_k)}{\|x_k - x_*\|} \leq \chi(x_k),$$

by definition of $\chi(x_k)$. □

Lemma 22. *If f satisfies Assumption 1 and if x_* is in the interior of \mathcal{D} and the algorithm achieves descent at each iteration, then*

$$\forall k' > k, \|\chi(x_{k'})\| \leq \left(\frac{\beta}{a}\right)^{3/2} \|\chi(x_k)\|.$$

Proof. Like in the proof of Lemma 1, we get

$$a\|x_k - x_*\| \leq \|\nabla f(x_k)\|.$$

Hence,

$$\begin{aligned} \beta\|x_{k'} - x_*\| &\leq \frac{\beta}{\sqrt{a}} \sqrt{f(x_{k'}) - f(x_*)} \leq \frac{\beta}{\sqrt{a}} \sqrt{f(x_k) - f(x_*)} \\ &\leq \frac{\beta}{\sqrt{a}} \sqrt{\beta(x_k - x_*)^2} \leq \frac{\beta^{3/2}}{\sqrt{a}} \sqrt{\frac{\chi(x_k)^2}{a}} \\ &\leq \left(\frac{\beta}{a}\right)^{3/2} \chi(x_k), \end{aligned}$$

where the first inequality comes from the strong convexity, and the second one comes from the fact that the algorithm is a descent, the third one comes from Lemma 21.

Now let $v = \arg \max_{x+v \in \mathcal{D}, \|v\| \leq 1} -v^T \nabla f(x_{k'})$.

$$-v^T \nabla f(x_{k'}) = -v^T \nabla f(x_{k'}) + v^T \nabla f(x_*) \leq \|\nabla f(x_{k'}) - \nabla f(x_*)\| \leq \beta\|x_k - x_*\|,$$

because $\nabla f(x_*) = 0$. This concludes the proof. □

C.2 Regret Analysis when the Optimum is in the interior of \mathcal{D}

Let us assume that x_* is in the interior of \mathcal{D} . Let us denote by Δ the distance from x_* to the closest boundary, and by $r = \Delta/4$.

Lemma 23. *If $\chi(x) \leq ar$ then $\|x - x_*\| \leq r$.*

If $\|x - x_*\| \geq r$ then $r \leq \frac{1}{a}\chi(x)$ thanks to Lemma 21.

Lemma 24. *For any $k \geq k\left((a/\beta)^{3/2} ar\right)$, $\|x_k - x_*\| \leq r$.*

Thanks to Lemma 1, after $k\left((a/\beta)^{3/2} ar\right)$ iterations, $\chi(x_k) \leq ar$. And thanks to the previous lemma, we thus have $\|x_k - x_*\| \leq r$.

Lemma 25. Set k_s the index of the first successful iteration following $k \left((a/\beta)^{3/2} ar \right)$ where $\alpha_k \leq 2\Delta$. After iteration k_s , $T(x_k, \alpha_k)$ spans all directions in \mathbb{R}^d , so that the instantaneous regret is the same as that of the algorithm in the unconstrained case with initial point x_{k_s} and initial step-size α_{k_s} . The iteration of this first successful iteration comes before $k_i := k \left((a/\beta)^{3/2} ar \right) + \frac{\log(\alpha_0/\Delta)}{\log 1/\theta}$.

Proof. If k is a successful iteration $\alpha_k \leq 2r = \Delta/(2)$, since $\|x_k - x_*\| \leq r$ and $\|x_{k+1} - x_*\| \leq r$. And if k is an unsuccessful iteration, it comes after one of those successes and a sequence of unsuccessful iterations, which yields $\alpha_k \leq \Delta/(2)$. \square

Lemma 26. On $\cap_{k \leq K} \mathcal{E}_k$,

$$\tilde{R}_K \leq C_7 \log(2/\delta) (1/\theta)^{-4C_f} + C_5 \log(2/\delta)^{1/3} \left(\sum_{k=1}^K N_k \right)^{2/3},$$

where $C_7 = (|\mathbb{D}|+1)U \frac{32}{c^2 \alpha_0^4} \frac{1}{(1/\theta)-1}$

and

$$C_f \leq E_1 \left(\left(\frac{a}{\beta} \right)^{3/2} ar \right)^{-2} + E_2 + \frac{\log(\alpha_0/\Delta)}{\log 1/\theta} + \frac{\beta \Delta^2}{\alpha_0}.$$

Proof. In the proof of Lemma 16, we isolated the steps preceding the first unsuccessful iteration. Similarly here, we treat the iterations before the first unsuccessful iteration after k_i , denoted by k'_f , separately from other iterations.

$$\begin{aligned} \tilde{R}_K &\leq \sum_{k=0}^{k'_f} N_k (f(x_k) - f(x_*)) + \sum_{k=0}^{k'_f} \left(N_k \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\quad + \sum_{k=k'_f}^K (N_k (|\mathbb{D}|+1) (f(x_k) - f(x_*))) + \sum_{k=k'_f}^K N_k \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right). \end{aligned}$$

Because $f(x) - f(x_*)$ is bounded by U ,

$$\begin{aligned} &\sum_{k=0}^{k'_f} N_k (f(x_k) - f(x_*)) + \sum_{k=0}^{k'_f} \left(N_k \sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_*) \right) \\ &\leq \sum_{k=0}^{k'_f} (|\mathbb{D}|+1) N_k U. \end{aligned}$$

By rewriting N_k ,

$$\begin{aligned} \sum_{k=0}^{k'_f} (|\mathbb{D}|+1) N_k U &\leq \sum_{k=0}^{k'_f} (|\mathbb{D}|+1) U \frac{32 \log(1/\delta)}{c^2 \alpha_k^4} \\ &\leq \sum_{k=0}^{k'_f} (|\mathbb{D}|+1) U \frac{32 \log(1/\delta)}{c^2 \alpha_0^4 \theta^{4k}} \\ &\leq (|\mathbb{D}|+1) U \frac{32 \log(1/\delta)}{c^2 \alpha_0^4} \frac{(1/\theta)^{-4k'_f}}{(1/\theta) - 1}. \end{aligned}$$

Also on $\cap_{k \leq K} \mathcal{E}_k$,

$$k'_f \leq k_i + \frac{f(x_{k_i}) - f(x_*)}{\alpha_0} \leq k_i + \frac{\beta \Delta^2}{\alpha_0} \leq k \left(\left(\frac{a}{\beta} \right)^{3/2} ar \right) + \frac{\log(\alpha_0/\Delta)}{\log 1/\theta} + \frac{\beta \Delta^2}{\alpha_k},$$

where the first inequality comes from the same argument used to prove Lemma 4, the second inequality comes from the smoothness of f and the third one comes from the definition of k_i .

On $\cap_{k \leq K} \mathcal{E}_k$,

$$k \left(\left(\frac{a}{\beta} \right)^{3/2} ar \right) \leq E_1 \left(\left(\frac{a}{\beta} \right)^{3/2} ar \right)^{-2} + E_2,$$

with E_1, E_2 defined in [Gratton et al., 2019]. Finally, we focus on the part of the regret accumulated before k'_f . On $\cap_{k \leq K} \mathcal{E}_k$

$$\begin{aligned} \sum_{k=k'_f}^K (N_k(|\mathbb{D}|+1)(f(x_k) - f(x_*))) + \sum_{k=k'_f}^K N_k \left(\sum_{v \in \mathbb{D}} f(x_k + \alpha_k v) - f(x_k) \right) \\ \leq C_5 \log(2/\delta)^{1/3} \left(\sum_{k=1}^K N_k \right)^{2/3}, \end{aligned}$$

by following exactly the same steps as those needed to bound the regret in the unconstrained case. \square

Theorem 2. *Under Assumptions 1, 3, 4, and 5, and if $x_* \in \text{int}(\mathcal{D})$, the cumulative regret R_T of FDS-Plan (respectively FDS-Seq) after the first T evaluations of f , satisfies*

$$\mathbb{E}(R_T) = O(\log(T)^{2/3} T^{2/3})$$

for the choice $\delta = T^{-4/3}$ (respectively $\delta = T^{-10/3}$ for FDS-Seq).

Proof. for FDS-Plan. We denote by K_T the last round reached by the algorithm with T evaluations.. Lemma 26 proves that on the event $\cap_{k \leq K_T} \mathcal{E}_k$,

$$\tilde{R}_K \leq C_7 \log(2/\delta)(1/\theta)^{-4C_f} + C_5 \log(2/\delta)^{1/3} \left(\sum_{k=1}^K N_k \right)^{2/3}.$$

Thanks to Lemma 8,

$$\mathbb{P} \left(\cup_{k=1}^{K_T} \mathcal{E}_k^C \right) \leq (|\mathbb{D}|+1) \mathbb{P} \left(\cup_{k=1}^T \mathcal{E}_k^C \right) \leq (|\mathbb{D}|+1) \sum_{t=1}^T T^{-4/3}/2 \leq (|\mathbb{D}|+1) T^{-1/3}$$

since $K_T \leq T$. Hence,

$$\begin{aligned} \mathbb{E}(R_T) &\leq \frac{4}{3} C_7 (1/\theta)^{-4C_f} \log(2T) + \frac{4}{3} \left(\frac{1}{(|\mathbb{D}|+1)} \right)^{2/3} C_5 \log(2T)^{2/3} T^{2/3} \\ &\quad + (|\mathbb{D}|+1) U T^{2/3} \\ &= O((\log T)^{2/3} T^{2/3}). \end{aligned}$$

\square

Adaptation of the proof for FDS-Seq The way of adapting the proof of FDS-Plan to the case of FDS-Seq of Section B.4 applies verbatim.