



HAL
open science

Regret analysis of the Piyavskii-Shubert algorithm for global Lipschitz optimization

Clément Bouttier, Tommaso R Cesari, Mélanie Ducoffe, Sébastien Gerchinovitz

► **To cite this version:**

Clément Bouttier, Tommaso R Cesari, Mélanie Ducoffe, Sébastien Gerchinovitz. Regret analysis of the Piyavskii-Shubert algorithm for global Lipschitz optimization. 2020. hal-02466084v2

HAL Id: hal-02466084

<https://hal.science/hal-02466084v2>

Preprint submitted on 15 Mar 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 Piyavskii-Shubert algorithm for global Lipschitz optimization

Clément Bouttier^a Tommaso Cesari^{b,c}, Mélanie Ducoffe^d, and Sébastien Gerchinovitz^{d,e}

^aAURA AERO, Toulouse, France; ^bUniversità degli Studi di Milano, Milano, Italy; ^cToulouse School of Economics, Toulouse, France; ^dIRT Saint Exupéry, Toulouse, France; ^eInstitut de Mathématiques de Toulouse, Toulouse, France.

ARTICLE HISTORY

Compiled March 15, 2022

ABSTRACT

We consider the problem of maximizing a non-concave Lipschitz multivariate function over a compact domain by sequentially querying its (possibly perturbed) values. We study a natural algorithm designed originally by Piyavskii and Shubert in 1972, for which we prove new bounds on the number of evaluations of the function needed to reach or certify a given optimization accuracy. Our analysis uses a bandit-optimization viewpoint and solves an open problem from Hansen et al. (1991) by bounding the number of evaluations to certify a given accuracy with a near-optimal sum of packing numbers.

KEYWORDS

Sequential optimization; Lipschitz; non-convex

1. Introduction

In this paper, $f: \mathcal{X} \rightarrow \mathbb{R}$ denotes an unknown function defined on a known compact subset \mathcal{X} of \mathbb{R}^d containing at least two points (if $|\mathcal{X}| \leq 1$, the problem becomes trivial). At a high level, we consider the following global optimization problem: with only black-box access to f , find an approximate maximizer of f at prescribed accuracy ε , with as little evaluations of f as possible. We make the setting more precise below.

1.1. Setting: Lipschitz black-box optimization, with or without certificates

We make the following weak Lipschitz assumption,¹ where the norm $\|\cdot\|$ and the constant L are known to the learner. For some of our results, we will instead assume f to be globally L -Lipschitz.

Assumption 1.1 (Lipschitzness around a maximizer). The function f attains its maximum at some $\mathbf{x}^* \in \mathcal{X}$, and there exist a constant $L > 0$ and a norm $\|\cdot\|$ such that,

CONTACT Tommaso Cesari. Email: tommaso.cesari@tse-fr.eu

¹Assumption 1.1 is sometimes referred to as *local smoothness* [20] and close in spirit to *calmness* [24].

for all $\mathbf{x} \in \mathcal{X}$,

$$f(\mathbf{x}) \geq f(\mathbf{x}^*) - L\|\mathbf{x}^* - \mathbf{x}\| .$$

Online learning protocol. We study the case in which f is black-box, i.e., except for some *a priori* knowledge on its smoothness, we can only access f by sequentially querying its values at a family $\mathbf{x}_1, \mathbf{x}_2, \dots \in \mathcal{X}$ of points of our choice. At every round $k \geq 1$, the query point \mathbf{x}_k can be chosen as a deterministic function of the values $f(\mathbf{x}_1), \dots, f(\mathbf{x}_{k-1})$ observed so far. At the end of round k , using the values $f(\mathbf{x}_1), \dots, f(\mathbf{x}_k)$, the learner outputs a recommendation $\mathbf{x}_k^* \in \mathcal{X}$ with the goal of minimizing the *optimization error* (or *simple regret*): $\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - f(\mathbf{x}_k^*) = f(\mathbf{x}^*) - f(\mathbf{x}_k^*)$.

The learner may also output an *error certificate*, depending on whether the algorithm is certified or not. More precisely, in all the sequel, we consider two different types of algorithms:

- *Non-certified algorithms* only output a recommendation $\mathbf{x}_k^* \in \mathcal{X}$ at the end of each time step k .
- *Certified algorithms* output an *error certificate* $\xi_k \geq 0$ together with the recommendation \mathbf{x}_k^* at the end of each time step k . The error certificate has to satisfy

$$f(\mathbf{x}^*) - f(\mathbf{x}_k^*) \leq \xi_k$$

for any function f that is L -Lipschitz around \mathbf{x}^* . In other words, the error certificate guarantees that the recommendation \mathbf{x}_k^* is ξ_k -optimal.

These two types of algorithms are summarized in Online Protocol 1 below.

Online Protocol 1 Non-certified (/certified) setting

for $k = 1, 2, \dots$ **do**
 pick the next query point $\mathbf{x}_k \in \mathcal{X}$
 observe the value $f(\mathbf{x}_k) \in \mathbb{R}$
 output a recommendation $\mathbf{x}_k^* \in \mathcal{X}$
 output an error certificate $\xi_k \geq 0$ (*certified setting only*)
end for

The Piyavskii–Shubert algorithm. Several algorithms were designed for these settings in the past (see related works in Section 5). In this paper, we focus on arguably one of the most natural ones, first appearing in the one-dimensional works of Piyavskii [23] and Shubert [25]. The principle is simple: at each round k , the query point \mathbf{x}_k is chosen as a (possibly approximate) maximizer of the piecewise-conic proxy function

$$\begin{aligned} \widehat{f}_{k-1}: \mathcal{X} &\rightarrow \mathbb{R} \\ \mathbf{x} &\mapsto \min_{i \in \{1, \dots, k-1\}} \{f(\mathbf{x}_i) + L\|\mathbf{x}_i - \mathbf{x}\|\} \end{aligned}$$

(see Figure 1), which satisfies $\widehat{f}_{k-1}(\mathbf{x}^*) \geq f(\mathbf{x}^*)$ under Assumption 1.1, and is the tightest upper bound on f given the available information when f is L -globally Lipschitz.

At the end of round k , the error certificate is chosen as $\xi_k = \max_{\mathbf{x} \in \mathcal{X}} \widehat{f}_k(\mathbf{x}) - f(\mathbf{x}_k^*)$, which upper bounds $f(\mathbf{x}^*) - f(\mathbf{x}_k^*)$ by construction.

In this paper, we study the sample complexity of both the non-certified and certified versions of the Piyavskii–Shubert algorithm. That is, for any target accuracy $\varepsilon > 0$, we bound the number n of evaluations of f after which $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \varepsilon$ (i.e., the recommendation of the algorithm is ε -optimal) or after which $\xi_n \leq \varepsilon$ (i.e., the algorithm can certify from the observed data that the recommendation is ε -optimal). As will become apparent from our bounds, certifying the accuracy of a recommendation is sometimes much harder than just outputting an approximate maximizer with correct (but unknown) accuracy (see comment after Theorem 3.6).

Open questions. Though the Piyavskii–Shubert algorithm is very natural, little is known about its sample complexity. In the non-certified setting, we are only aware of the worst-case bound $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \lesssim n^{-1/d}$, which corresponds to the fact that $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \varepsilon$ whenever $n \gtrsim (1/\varepsilon)^d$. In the certified setting, Hansen et al. [8] proved in dimension $d = 1$ that $\xi_n \leq \varepsilon$ whenever $n \gtrsim \int_0^1 (f(x^*) - f(x) + \varepsilon)^{-1} dx$, but the authors claimed that “*extending the results of this paper to the multivariate [case] appears to be difficult*”. We refine or extend these previous bounds in the sequel. See below and the related works in Section 5 for more details.

1.2. Outline of the paper and main contributions

In Section 1.3, we present some recurring notation. We recall the definition of the Piyavskii–Shubert algorithm in Section 2. We then make the following contributions:

- In Sections 3.1 and 3.2 we derive optimization error bounds for the Piyavskii–Shubert algorithm in any dimension $d \geq 1$, both in the non-certified and the certified settings (Theorems 3.5 and 3.6). We bound the sample complexity in terms of packing numbers of sets of suboptimal points at different scales. As mentioned above, error bounds for the Piyavskii–Shubert algorithm were previously known only for worst or special cases. Moreover, the absence of any underlying discrete structure (as opposed to, e.g., the DOO algorithm [19, 22], or, to some extent, the Zooming algorithm [13]) makes the Piyavskii–Shubert algorithm and its analysis quite simple. We hope this helps better understand the role of packing numbers in such types of results.
- In Sections 3.3 and 3.4 we reinterpret the error bounds of Section 3.2 in view of the very recent results of Bachoc et al. [2].² The latter results have two consequences: (i) our error bound in the certified case is nearly instance-optimal; (ii) this bound is proportional to the integral $\int_{\mathcal{X}} d\mathbf{x} / (f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d$ for any dimension $d \geq 1$, which solves the question left open by Hansen et al. [8] three decades ago.
- In Section 4 we investigate the robustness of the Piyavskii–Shubert algorithm to deterministic perturbations and imperfect maximization of the proxy function. We prove optimization error bounds that generalize those of Section 3.2, by allowing both perturbed evaluations of f and imperfect maximization of the proxy function.

²For contingent reasons, [2] was written after but submitted before the present paper. Since there are some dependencies between the results of these two papers, we clarify in Section 5 which result belongs to which paper, and that no logic loops arise from cross-citing.

Related works from the global optimization and bandit optimization literatures are discussed in details in Section 5. Finally, although the main proofs appear in the main body, some technical details are postponed to the appendix, where we also recall some useful results on packing and covering numbers.

1.3. Recurring notation

We begin by recalling the definition of packing numbers, which is of utmost importance for our analysis. In Appendix A, we also recall a few known inequalities about packing (and covering) numbers that will prove useful throughout the paper.

Let $\|\cdot\|$ be a norm. For any bounded set $A \subset \mathbb{R}^d$ and $r > 0$, the r -packing number of A (with respect to $\|\cdot\|$) is the largest number of r -separated points contained in A , i.e.,

$$\mathcal{N}_{\|\cdot\|}(A, r) := \sup \left\{ k \in \{1, 2, \dots\} : \exists \mathbf{x}_1, \dots, \mathbf{x}_k \in A, \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\| > r \right\},$$

if A is nonempty, zero otherwise.

Our results will be expressed in terms of packing numbers of near-optimal points and layers, defined as follows. Let $\mathcal{X} \subset \mathbb{R}^d$, $f: \mathcal{X} \rightarrow \mathbb{R}$, and \mathbf{x}^* be a maximizer of f . For all $\varepsilon > 0$, we denote the set $\{\mathbf{x} \in \mathcal{X} : f(\mathbf{x}) \geq f(\mathbf{x}^*) - \varepsilon\}$ of all ε -optimal points of f by \mathcal{X}_ε and its complement (i.e., the set of ε -suboptimal points) by $\mathcal{X}_\varepsilon^c$. For all $0 \leq a < b$, we denote the $(a, b]$ -layer $\mathcal{X}_a^c \cap \mathcal{X}_b$ (i.e., the set of points that are a -suboptimal but b -optimal) by $\mathcal{X}_{(a, b]}$.

Note that if $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer with respect to a norm $\|\cdot\|$, then every point in \mathcal{X} is ε_0 -optimal, where $\varepsilon_0 = \varepsilon_0(\mathcal{X}, L, \|\cdot\|) := L \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$. In other words, $\mathcal{X}_{\varepsilon_0} = \mathcal{X}$. For this reason, without loss of generality we will only consider values of ε smaller than ε_0 .

For all $x \in \mathbb{R}$, we denote the smallest integer greater than or equal to x by $\lceil x \rceil$. We also write $[k] = \{1, \dots, k\}$ for any integer $k \geq 1$. For any norm $\|\cdot\|$, radius $r > 0$, and center $\mathbf{x} \in \mathbb{R}^d$, we denote by $B_{\|\cdot\|}(r, \mathbf{x})$ the ball in $(\mathbb{R}^d, \|\cdot\|)$ of radius r centered at \mathbf{x} . For any norm $\|\cdot\|$ and set E , we denote the diameter $\sup_{\mathbf{x}, \mathbf{y} \in E} \|\mathbf{x} - \mathbf{y}\|$ of E (with respect to $\|\cdot\|$) by $\text{diam}_{\|\cdot\|}(E)$. We denote the Lebesgue measure of a (Lebesgue-measurable) set $E \subset \mathbb{R}^d$ by $\text{vol}(E)$ and we simply refer to it as its *volume*.

2. The Piyavskii–Shubert algorithm

The Piyavskii–Shubert algorithm (Algorithm 2, illustrated in Figure 1) works by maintaining a proxy \widehat{f}_k of the objective function f that upper bounds f at the maximizer \mathbf{x}^* (see Lemma 3.1). The next query point \mathbf{x}_{k+1} is chosen as a maximizer of \widehat{f}_k . The intuition is that \mathbf{x}_{k+1} is either an approximate maximizer of f (which reduces the optimization error immediately) or belongs to a largely unexplored suboptimal region (which helps reduce the regret in the future, see Lemmas 3.3 and 3.4). Note the similarity with the “optimism in the face of uncertainty” principle in stochastic multiarmed bandits [15]. There, a proxy for the real mean of an arm is given by an upper confidence bound, and arms can have high upper confidence bounds only if they are either actually good, or they belong to an under-explored set of arms. In Section 3.1, we will show that the certified version of the Piyavskii–Shubert algorithm is indeed a certified




Figure 1. First four iterations of the Piyavskii–Shubert algorithm in dimension $d = 1$. In thick blue the function f . In black the proxy function \hat{f}_4 .

algorithm.

Algorithm 2 Piyavskii–Shubert, for the non-certified (*/certified*) setting

input: Lipschitz constant $L > 0$, norm $\|\cdot\|$, initial guess $\mathbf{x}_1 \in \mathcal{X}$

for $k = 1, 2, \dots$ **do**

pick the next query point \mathbf{x}_k

observe the value $f(\mathbf{x}_k)$

output the recommendation $\mathbf{x}_k^* \leftarrow \operatorname{argmax}_{\mathbf{x} \in \{\mathbf{x}_1, \dots, \mathbf{x}_k\}} f(\mathbf{x})$

update the proxy function

$$\hat{f}_k(\cdot) \leftarrow \min_{i \in [k]} \left\{ f(\mathbf{x}_i) + L \|\mathbf{x}_i - (\cdot)\| \right\}$$

output the error certificate $\xi_k \leftarrow \hat{f}_k^* - f_k^*$, where

$$\hat{f}_k^* \leftarrow \max_{\mathbf{x} \in \mathcal{X}} \hat{f}_k(\mathbf{x}), \quad f_k^* \leftarrow \max_{i \in [k]} f(\mathbf{x}_i) \quad (\text{certified setting only})$$

let $\mathbf{x}_{k+1} \in \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \hat{f}_k(\mathbf{x})$

end for

It is worth noting that the Piyavskii–Shubert algorithm might not be computationally efficient in high dimensions. Indeed, for a budget of n evaluations of f , the problem of optimizing f is replaced by the optimization of the n proxy functions $\hat{f}_1, \dots, \hat{f}_n$, which could be a demanding task on its own. To see this, note that finding the maximum of \hat{f}_k is related to the problem of determining a Voronoi diagram for a set of k points, which is known to have a computational complexity that is exponential in the dimension of the ambient space [1, Theorem 4.5, Section 4.3.2 “Power diagrams and convex hulls”]. Some algorithms get around these computational issues by replacing the piecewise-conic proxies of the Piyavskii–Shubert algorithm with looser but simpler

functions. The DOO algorithm [19] (see also [22]), for example, uses piecewise-constant functions, with pieces that correspond to a predetermined hierarchical partition of \mathcal{X} .

Beyond pedagogical reasons and the fact that it appears to be one of the most natural methods in our setting, the Piyavskii–Shubert algorithm can be applied to low-dimensional real-life problems (e.g., hyperparameter tuning), but also to maximization problems in which the computational cost is driven by the evaluation of the objective f rather than the dimension of the ambient space. Indeed, when measuring algorithms by their sample complexity (number of evaluations of f) rather than their running time (number of elementary operations aside evaluations of f), we conjecture that the Piyavskii–Shubert algorithm is better (by non-negligible multiplicative constants) than simplified algorithms such as DOO. This is supported by some low-dimensional experiments carried out in [10].

3. Error bounds

We begin this section by showing some key properties of the Piyavskii–Shubert algorithm (Section 3.1). We then prove upper bounds on the number of queries that the Piyavskii–Shubert algorithm needs before being able to reach or certify a certain accuracy (Section 3.2). Afterwards, we discuss optimality (Section 3.3), and give a compact integral representation of our upper bound in the certified setting (Section 3.4).

3.1. Useful lemmas

We begin this section by proving two important properties of the proxy functions \widehat{f}_k that hold whenever the objective function f is Lipschitz around a maximizer \mathbf{x}^* . First, the proxy functions are always greater than or equal to the objective function at the maximizer \mathbf{x}^* . (Note that when f is globally Lipschitz, the lower bound $\widehat{f}_k(\mathbf{x}) \geq f(\mathbf{x})$ is well known and true not only for $\mathbf{x} = \mathbf{x}^*$ but for all $\mathbf{x} \in \mathcal{X}$.) Second, the proxy functions are always lower than or equal to the objective function at all past query points (so, in particular, the same is true for all recommendations).

Lemma 3.1. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1) and the Piyavskii–Shubert algorithm (Algorithm 2) is run with input $L, \|\cdot\|, \mathbf{x}_1$. Then, for all $k \geq 1$, the proxy function $\widehat{f}_k(\cdot) = \min_{i \in [k]} \{f(\mathbf{x}_i) + L\|\mathbf{x}_i - (\cdot)\|\}$ is L -Lipschitz and satisfies*

$$\widehat{f}_k(\mathbf{x}^*) \geq f(\mathbf{x}^*) \quad \text{and} \quad \widehat{f}_k(\mathbf{x}_j) \leq f(\mathbf{x}_j), \quad \forall j \in [k].$$

Proof. Fix any $k \geq 1$. The fact that \widehat{f}_k is L -Lipschitz is straightforward. Moreover, by the L -Lipschitzness of f around \mathbf{x}^* , we get

$$\begin{aligned} \widehat{f}_k(\mathbf{x}^*) &= \min_{i \in [k]} \left\{ f(\mathbf{x}_i) + L\|\mathbf{x}_i - \mathbf{x}^*\| \right\} \\ &\geq \min_{i \in [k]} \left\{ (f(\mathbf{x}^*) - L\|\mathbf{x}_i - \mathbf{x}^*\|) + L\|\mathbf{x}_i - \mathbf{x}^*\| \right\} = f(\mathbf{x}^*). \end{aligned}$$

Furthermore, for all $j \in [k]$, the definition of $\widehat{f}_k(\mathbf{x}_j)$ implies immediately

$$\widehat{f}_k(\mathbf{x}_j) = \min_{i \in [k]} \left\{ f(\mathbf{x}_i) + L\|\mathbf{x}_i - \mathbf{x}_j\| \right\} \leq f(\mathbf{x}_j) + L\|\mathbf{x}_j - \mathbf{x}_j\| = f(\mathbf{x}_j). \quad \square$$

A first immediate consequence of Lemma 3.1 is that the certified version of the Piyavskii–Shubert algorithm is indeed a certified algorithm, as defined in Section 1.1.

Lemma 3.2. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1) and the Certified Piyavskii–Shubert algorithm (Algorithm 2) is run with input $L, \|\cdot\|, \mathbf{x}_1$. Then, for all $n \geq 1$, $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \xi_n$.*

Proof. By Lemma 3.1 and the definitions of f_n^* , \widehat{f}_n^* , and ξ_n , for all $n \geq 1$, we get $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \widehat{f}_n(\mathbf{x}^*) - f_n^* \leq \xi_n$. \square

Another important implication of Lemma 3.1 is that if the Piyavskii–Shubert algorithm observes f at a Δ -suboptimal point \mathbf{x}_i , then all the next query points \mathbf{x}_j are at least (Δ/L) -distant from \mathbf{x}_i . In other words, the algorithm does not explore too much in suboptimal regions.

Lemma 3.3. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1) and the Piyavskii–Shubert algorithm (Algorithm 2) is run with input $L, \|\cdot\|, \mathbf{x}_1$. Fix any $\Delta > 0$ and assume that there exists $i \geq 1$ such that the i -th query point \mathbf{x}_i belongs to \mathcal{X}_Δ^c . Then, for all $j > i$, the j -th query point satisfies*

$$\|\mathbf{x}_j - \mathbf{x}_i\| > \frac{\Delta}{L}.$$

Proof. Fix any $j > i$. Then, using the fact that \mathbf{x}_j has been selected as the j -th query point (in the first inequality), Lemma 3.1 (in the second and fourth), and $\mathbf{x}_i \in \mathcal{X}_\Delta^c$ (in the third), we get

$$\widehat{f}_{j-1}(\mathbf{x}_j) \geq \widehat{f}_{j-1}(\mathbf{x}^*) \geq f(\mathbf{x}^*) > f(\mathbf{x}_i) + \Delta \geq \widehat{f}_{j-1}(\mathbf{x}_i) + \Delta.$$

Since \widehat{f}_{j-1} is L -Lipschitz (by Lemma 3.1), we can thus conclude that $L\|\mathbf{x}_j - \mathbf{x}_i\| \geq |\widehat{f}_{j-1}(\mathbf{x}_j) - \widehat{f}_{j-1}(\mathbf{x}_i)| > \Delta$. \square

As a final corollary to Lemma 3.1, we show that the Certified Piyavskii–Shubert algorithm automatically adapts its exploration to all possible approximation scales at the same time. In words, for any accuracy ε , until the algorithm can certify that the optimization error is ε or better (i.e., until the error certificate is smaller than or equal to ε), it only queries points that are at least (ε/L) -distant from all past query points. The strength of this statement comes from the fact that the algorithm *does not* take an accuracy ε as a parameter and the result is consequently not proved for a specific and known ε : it holds for all ε simultaneously.

Lemma 3.4. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1) and the Certified Piyavskii–Shubert algorithm (Algorithm 2) is run with input $L, \|\cdot\|, \mathbf{x}_1$. Fix any $\varepsilon > 0$ and let τ be the first time*

where the algorithm returns an error certificate $\xi_\tau \leq \varepsilon$. Then, for all distinct $i, j \in [\tau]$,

$$\|\mathbf{x}_i - \mathbf{x}_j\| > \frac{\varepsilon}{L}.$$

Proof. Without loss of generality, assume $1 \leq i < j \leq \tau$. Then, by the definition of \widehat{f}_{j-1}^* (first equality), that of ξ_{j-1} together with the assumption $\xi_{j-1} > \varepsilon$ (first inequality), the definition of f_{j-1}^* (second), and Lemma 3.1 (last), we get

$$\widehat{f}_{j-1}(\mathbf{x}_j) = (\widehat{f}_{j-1}^* - f_{j-1}^*) + f_{j-1}^* > \varepsilon + f_{j-1}^* \geq \varepsilon + f(\mathbf{x}_i) \geq \varepsilon + \widehat{f}_{j-1}(\mathbf{x}_i).$$

Since \widehat{f}_{j-1} is L -Lipschitz (by Lemma 3.1), we can thus conclude that $L\|\mathbf{x}_j - \mathbf{x}_i\| \geq |\widehat{f}_{j-1}(\mathbf{x}_j) - \widehat{f}_{j-1}(\mathbf{x}_i)| > \varepsilon$. \square

The previous result shows that ξ_n serves a synergistic double-role in the Certified Piyavskii–Shubert algorithm. When it is large, it guarantees that the queries are sufficiently sparse, and as soon as it becomes small, it certifies that the recommendation \mathbf{x}_n is sufficiently good.

3.2. Upper bounds

We now prove an upper bound on the number of evaluations of f that are needed before the recommendations of the Piyavskii–Shubert algorithm become near-optimal. To this end, we use a peeling technique in which the input space \mathcal{X} is partitioned in terms of the output values of f and leverage the fact that highly suboptimal regions $\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$ (with $\varepsilon_k \gg \varepsilon$) need not be explored too much thanks to the results we showed in Section 3.1.

Theorem 3.5. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), $\varepsilon_0 := L \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$, $\varepsilon \in (0, \varepsilon_0)$, $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon} := \varepsilon$, $\varepsilon_k := \varepsilon_0 2^{-k}$ for all $k < m_\varepsilon$, and*

$$S_{\text{NC}}(f, \varepsilon) := \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L} \right).$$

If $n \geq 1 + S_{\text{NC}}(f, \varepsilon)$, then the n -th recommendation \mathbf{x}_n^ of the Piyavskii–Shubert algorithm (Algorithm 2, run with input $L, \|\cdot\|, \mathbf{x}_1$) satisfies $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \varepsilon$.*

We can draw immediate corollaries from Theorem 3.5 if f is such that $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_r, r) \leq C^*/r^{d^*}$, for some $C^* > 0$, $d^* \in [0, d]$, and all $r \in (\varepsilon, \varepsilon_0]$ (i.e., if the *near-optimality dimension* is d^* —see the related works in Section 5 for more details). In this case, $S_{\text{NC}}(f, \varepsilon) \lesssim \log(1/\varepsilon)$ if $d^* = 0$, and $S_{\text{NC}}(f, \varepsilon) \lesssim 1/\varepsilon^{d^*}$ if $d^* > 0$. Examples of these functions are linear functions of the norms $f(\mathbf{x}) = a - b\|\mathbf{x} - \mathbf{x}_0\|$, where $d^* = 0$, or quadratic functions of the norms $f(\mathbf{x}) = a - b\|\mathbf{x} - \mathbf{x}_0\|^2$, where $d^* = d/2$.

Importantly, our bound in terms of $S_{\text{NC}}(f, \varepsilon)$ is more general than the previous corollaries. It accommodates cases in which f has different of such d^* 's at different scales r , e.g., the function $f(\mathbf{x}) = (1/4 - \|\mathbf{x}\|^2)\mathbb{I}\{\|\mathbf{x}\| \leq 1/2\} + (1/2 - \|\mathbf{x}\|)\mathbb{I}\{\|\mathbf{x}\| > 1/2\}$, for which $S_{\text{NC}}(f, \varepsilon) \lesssim \log(1/\varepsilon)$ for large values of ε (linear regime if $\varepsilon \geq 1/2$) but

$S_{\text{NC}}(f, \varepsilon) \lesssim 1/\varepsilon^{d/2}$ for small values of ε (quadratic regime if $\varepsilon < 1/2$). Many different examples can be designed this way.

Proof. By definition of ε_k , we can partition the set $\mathcal{X}_\varepsilon^c = \mathcal{X}_{(\varepsilon, \varepsilon_0]}$ of ε -suboptimal points into m_ε layers

$$\mathcal{X}_\varepsilon^c = \bigcup_{k=1}^{m_\varepsilon} \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$$

Therefore, for any $i \geq 1$, if $\mathbf{x}_i \in \mathcal{X}_\varepsilon^c$, then there exists a unique $k \in [m_\varepsilon]$ such that $\mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$. Since the Piyavskii–Shubert algorithm does not explore too much in suboptimal regions, Lemma 3.3 and the definition of packing numbers imply that for all $k \in [m_\varepsilon]$, the maximum number of times that a query point \mathbf{x}_i is chosen in $\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]} \subset \mathcal{X}_\varepsilon^c$ is upper bounded by $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L})$. Hence

$$\begin{aligned} \left| \{i \in \{1, 2, \dots\} : \mathbf{x}_i \in \mathcal{X}_\varepsilon^c\} \right| &= \sum_{k=1}^{m_\varepsilon} \left| \{i \in \{1, 2, \dots\} : \mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}\} \right| \\ &\leq \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L} \right) = S_{\text{NC}}(f, \varepsilon). \end{aligned}$$

Therefore, if $n \geq 1 + S_{\text{NC}}(f, \varepsilon)$, then there exists $i \in [n]$ such that $\mathbf{x}_i \in \mathcal{X}_\varepsilon$, which implies (by definition of \mathbf{x}_n^*) that $f(\mathbf{x}_n^*) \geq f(\mathbf{x}_i) \geq f(\mathbf{x}^*) - \varepsilon$ and concludes the proof. \square

We now prove the analogous (in the certified setting) of the previous result. Its proof combines the ideas in Theorem 3.5 with the synergistic double-role of the error certificate presented in Lemma 3.4.

Notably, when $n \geq S_{\text{C}}(f, \varepsilon)$, not only $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \varepsilon$ but the Certified Piyavskii–Shubert algorithm can certify that this bound holds at run-time based only on the available data. See below the proof for further comments.

Theorem 3.6. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), $\varepsilon_0 := L \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$, $\varepsilon \in (0, \varepsilon_0)$, $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon} := \varepsilon$, $\varepsilon_k := \varepsilon_0 2^{-k}$ for all $k < m_\varepsilon$, and*

$$S_{\text{C}}(f, \varepsilon) := \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L} \right) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L} \right).$$

If $n \geq S_{\text{C}}(f, \varepsilon)$, then the n -th error certificate ξ_n of the Certified Piyavskii–Shubert algorithm (Algorithm 2, run with input $L, \|\cdot\|, \mathbf{x}_1$) satisfies $\xi_n \leq \varepsilon$.

Proof. By definition of ε_k , we can partition \mathcal{X} into the $m_\varepsilon + 1$ sets

$$\mathcal{X} = \mathcal{X}_\varepsilon \cup \bigcup_{k=1}^{m_\varepsilon} \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$$

Let τ be the first time where the algorithm returns an error certificate $\xi_\tau \leq \varepsilon$. Then, for any $i \in [\tau]$, either $\mathbf{x}_i \in \mathcal{X}_\varepsilon$ or there exists a unique $k \in [m_\varepsilon]$ such that $\mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$.

Since the Certified Piyavskii–Shubert algorithm automatically adapts its exploration to all possible approximation scales at the same time, Lemma 3.4 and the definition of packing numbers imply that the maximum number of times that a query point \mathbf{x}_i is chosen in \mathcal{X}_ε (up to and including time τ) is upper bounded by $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L})$. Similarly, since the Certified Piyavskii–Shubert algorithm does not explore too much in suboptimal regions, Lemma 3.3 and the definition of packing numbers imply that for all $k \in [m_\varepsilon]$, the maximum number of times that a query point \mathbf{x}_i is chosen in $\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1})} \subset \mathcal{X}_{\varepsilon_k}^c$ is upper bounded by $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1})}, \frac{\varepsilon_k}{L})$. Hence

$$\begin{aligned} \tau &= \left| \{i \in [\tau] : \mathbf{x}_i \in \mathcal{X}_\varepsilon\} \right| + \sum_{k=1}^{m_\varepsilon} \left| \{i \in [\tau] : \mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1})}\} \right| \\ &\leq \mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1})}, \frac{\varepsilon_k}{L}) = S_C(f, \varepsilon). \end{aligned}$$

Therefore, noting that $k \mapsto \xi_k$ is non-increasing (because $\xi_k = \widehat{f}_k^* - f_k^*$, with $k \mapsto \widehat{f}_k^*$ non-increasing and $k \mapsto f_k^*$ non-decreasing), if $n \geq S_C(f, \varepsilon)$, then $n \geq \tau$, which in turn yields $\xi_n \leq \xi_\tau \leq \varepsilon$. \square

Though the two bounds in Theorems 3.5 and 3.6 look very similar, they differ by the term $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L})$. This difference is not negligible in general and can be explained with a simple example. Indeed, consider a constant function $c: [0, 1]^d \rightarrow \mathbb{R}$. Any non-certified algorithm outputs an ε -maximizer of c after 1 single query at any scale ε , because all points in the domain are maxima. However, the only way to *certify* that the output is ε -optimal is essentially to perform a grid-search of $[0, 1]^d$ with step-size roughly ε/L , so as to be sure there is no hidden bump of height more than ε . This is reflected in the term $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L})$, which is of order $(L/\varepsilon)^d$ for constant functions. At a high level, the more “constant” a function is, the easier it is to recommend an ε -optimal point, but the harder it is to certify that such recommendation is actually a good recommendation.

3.3. Optimality

We now discuss the optimality of the two bounds (Theorems 3.5 and 3.6) we derived in the previous section.

Non-certified optimality. In this paper, we focus on f -dependent bounds. By its nature, no meaningful lower bounds can be stated in the non-certified setting (because for each given function there is an algorithm that outputs one of its maximizers after only 1 query). For the interested reader, there exist minimax lower bounds over specific classes of functions that match the corresponding upper bounds obtained from $S_{\text{NC}}(f, \varepsilon)$, such as when the near-optimality dimension is d [21] or $d/2$ [11] (see the related works in Section 5 for more details).

Certified optimality. One of the interesting features of the certified setting is that it makes it possible to write non-trivial f -dependent lower bounds, as shown by [2].

In this paragraph, we focus on the set \mathcal{F}_L of globally L -Lipschitz functions f with $L > \text{Lip}(f)$, where $\text{Lip}(f)$ is the smallest Lipschitz constant for f . In other words, we

assume that the global Lipschitz constant of the unknown objective function is not known exactly. (We will later briefly discuss the unlikely case in which $\text{Lip}(f)$ is known precisely.)

For any certified algorithm A , function $f \in \mathcal{F}_L$, and accuracy ε , we denote by $\sigma(A, f, \varepsilon)$ the smallest number of queries n that are needed to get $\xi_n \leq \varepsilon$, i.e., to find and certify an ε -maximizer. The following result holds.

Theorem 3.7 (Bachoc et al. [2, Theorem 2]). *For any certified algorithm A , globally L -Lipschitz function $f \in \mathcal{F}_L$, and accuracy $\varepsilon \in (0, \varepsilon_0]$, it holds that*

$$\sigma(A, f, \varepsilon) > \frac{c}{1 + m_\varepsilon} S_C(f, \varepsilon),$$

where $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$ and $c := 2^{-2}(2^{-7}(1 - \text{Lip}(f)/L))^d$.

The previous result states that unless the smallest Lipschitz constant of the black-box function f is known exactly, no algorithm can find and certify an ε -maximizer with less than $S_C(f, \varepsilon)$ queries (up to logarithmic terms).

In the unlikely scenario in which $\text{Lip}(f)$ is known exactly, it is known that $S_C(f, \varepsilon)$ remains a tight lower bound (up to logarithmic terms) in dimension $d = 1$ [2, Proposition 2]. Finding a tight f -dependent lower bound for this boundary case in dimension $d \geq 2$ is still an open problem.

3.4. Integral representation of $S_C(f, \varepsilon)$

In this section, we give an integral representation to the quantity $S_C(f, \varepsilon)$ we presented in Theorem 3.6. This result was first proven in [2, Theorem 1]. We state it and prove it again below, but with tighter multiplicative constants.

The theorem holds under a mild geometrical assumption on \mathcal{X} that roughly states that the set is “everywhere d -dimensional”, in the sense that all points $\mathbf{x} \in \mathcal{X}$ are the center of a sufficiently small ball, a constant portion of which is included in \mathcal{X} . This natural assumption has already appeared in the past (e.g., [2, 12]) and is weaker than another classical assumption in the statistics literature (the rolling ball assumption, [6, 28]).

Assumption 3.8. Let a norm $\|\cdot\|$ on \mathbb{R}^d and a constant $\gamma \in (0, 1]$ be both known to the learner. For all positive $r < \frac{1}{2} \text{diam}_{\|\cdot\|}(\mathcal{X})$,

$$\text{vol}(B_{\|\cdot\|}(r, \mathbf{x}) \cap \mathcal{X}) \geq \gamma \text{vol}(B_{\|\cdot\|}(r, \mathbf{0})).$$

The previous assumption holds for all convex sets with non-empty interior, finite unions of them, and more generally for any set \mathcal{X} that does not contain isolated points, cusps, or lower dimensional regions (such as segments in \mathbb{R}^2). Not also that γ may depend on d , e.g., $\gamma = 2^{-d}$ if \mathcal{X} is a hypercube.

The following intermediate result relates packing numbers with volumes. It first appeared in [2, Proposition 4], but we include a slightly updated proof for it in Appendix B.

Proposition 3.9. *Fix any norm $\|\cdot\|$, $\gamma \in (0, 1]$, and $L > 0$. If $\mathcal{X} \subset \mathbb{R}^d$ satisfies Assumption 3.8, then, for any function $f: \mathcal{X} \rightarrow \mathbb{R}$ that is globally L -Lipschitz with respect to $\|\cdot\|$ and for any $0 < w < u < \varepsilon_0 := L \text{diam}_{\|\cdot\|}(\mathcal{X})$, the following inequalities*

hold:

$$\mathcal{N}_{\|\cdot\|}\left(\mathcal{X}_u, \frac{u}{L}\right) \leq \frac{1}{\gamma} \frac{\text{vol}(\mathcal{X}_{(3/2)u})}{\text{vol}\left(\frac{u}{2L} B_{\|\cdot\|}(1, \mathbf{0})\right)}, \quad \mathcal{N}_{\|\cdot\|}\left(\mathcal{X}_{(w,u)}, \frac{w}{L}\right) \leq \frac{1}{\gamma} \frac{\text{vol}(\mathcal{X}_{(w/2, 3u/2)})}{\text{vol}\left(\frac{w}{2L} B_{\|\cdot\|}(1, \mathbf{0})\right)}.$$

The next key lemma shows that sums of volumes of non-overlapping sets of suboptimal points can be lower-bounded by a sum of volumes of overlapping ones. Compared to the original result in [2, Lemma 3], we slightly refine the multiplicative constants. The proof is deferred to Appendix B.

Lemma 3.10. *Fix any norm $\|\cdot\|$ and $L > 0$. Then, for any function $f: \mathcal{X} \rightarrow \mathbb{R}$ that is globally L -Lipschitz with respect to $\|\cdot\|$ and for all positive $\varepsilon < \varepsilon_0 := L \text{diam}_{\|\cdot\|}(\mathcal{X})$, letting $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon} := \varepsilon$, and $\varepsilon_k := \varepsilon_0 2^{-k}$ for all $k < m_\varepsilon$, the following inequality holds:*

$$\frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1})})}{\varepsilon_{k-1}^d} \leq \frac{15}{8} \cdot 4^d \left(\frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1})})}{\varepsilon_{k-1}^d} \right).$$

We can now state our integral characterization of the quantity $S_C(f, \varepsilon)$ appearing in Theorem 3.6, that controls the number of queries that the Certified Piyavskii–Shubert algorithm needs to output a certified ε -maximizer.

Theorem 3.11. *Fix any norm $\|\cdot\|$, $\gamma \in (0, 1]$, and $L > 0$. Then there exist two positive constants $c, C > 0$ depending only on $d, \gamma, L, \|\cdot\|$ (e.g., $c := L^d / \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))$ and $C := (15/8)(16L)^d / \gamma \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))$) such that:*

- (1) For any subset \mathcal{X} of \mathbb{R}^d that satisfies Assumption 3.8,
- (2) For any function $f: \mathcal{X} \rightarrow \mathbb{R}$ that is globally L -Lipschitz with respect to $\|\cdot\|$,

and for all positive $\varepsilon < \varepsilon_0 := L \text{diam}_{\|\cdot\|}(\mathcal{X})$, the following inequalities hold:³

$$c \left(\frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \int_{\mathcal{X}_\varepsilon} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}))^d} \right) \leq S_C(f, \varepsilon) \leq C \left(\frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \int_{\mathcal{X}_\varepsilon} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}))^d} \right).$$

The above result can be restated as $S_C(f, \varepsilon) = \Theta\left(\int_{\mathcal{X}} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d}\right)$ (for a direct proof of this claim, see Theorem B.1 in Appendix B). As noted in Section 1.2, combining this result with Theorem 3.6 resolves a long-standing conjecture of Hansen et al. [8].

Proof. Fix any $\mathcal{X} \subset \mathbb{R}^d$ and $f: \mathcal{X} \rightarrow \mathbb{R}$ satisfying the assumptions of the theorem, and let $\varepsilon < \varepsilon_0$. Let also $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon} := \varepsilon$, and $\varepsilon_k := \varepsilon_0 2^{-k}$ for all $k < m_\varepsilon$. Partition the domain of integration \mathcal{X} into the following $m_\varepsilon + 1$ sets: the set of ε -optimal points \mathcal{X}_ε and the m_ε layers $\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$, for $k \in [m_\varepsilon]$.

We begin by proving the first inequality. To this end, for any bounded set $A \subset \mathbb{R}^d$ and $r > 0$, we consider the r -covering number $\mathcal{M}_{\|\cdot\|}(A, r)$ of A (with respect to $\|\cdot\|$) defined as the smallest cardinality of an r -covering of A , i.e.,

$$\mathcal{M}_{\|\cdot\|}(A, r) := \min\{k \in \{1, 2, \dots\} : \exists \mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^d, \forall \mathbf{x} \in A, \exists i \in [k], \|\mathbf{x} - \mathbf{x}_i\| \leq r\}.$$

³We actually prove a stronger result. The first inequality holds more generally for any f that is L -Lipschitz around a maximizer with respect to $\|\cdot\|$ and Lebesgue-measurable, and does not require \mathcal{X} to be approximately d -dimensional.

if A is nonempty, zero otherwise. Then,

$$\begin{aligned}
\frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \int_{\mathcal{X}_\varepsilon^c} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}))^d} &\leq \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_k^d} \\
&\leq \frac{\mathcal{M}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) \cdot \text{vol}(B_{\|\cdot\|}(1, \mathbf{0})) (\frac{\varepsilon}{L})^d}{\varepsilon^d} \\
&\quad + \sum_{k=1}^{m_\varepsilon} \frac{\mathcal{M}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}) \cdot \text{vol}(B_{\|\cdot\|}(1, \mathbf{0})) (\frac{\varepsilon_k}{L})^d}{\varepsilon_k^d} \\
&\leq \frac{\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{L^d} \left(\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}) \right) \\
&= \frac{\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{L^d} S_C(f, \varepsilon),
\end{aligned}$$

where the first inequality follows by lower bounding $f(\mathbf{x}^*) - f$ with its infimum on the partition, the second one by upper bounding the volume of a set with the volume of the balls of a smallest ε_k/L -covering, and the last one by the fact that covering numbers are always upper bounded by packing numbers (we recall this known result in Appendix A, eq. (A1)). This proves the first part of the theorem.

For the second one, we have

$$\begin{aligned}
\frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \int_{\mathcal{X}_\varepsilon^c} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}))^d} &\geq \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \\
&\geq \frac{8}{15} \frac{1}{4^d} \left(\frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \right) \\
&\geq \frac{8}{15} \frac{\gamma \text{vol}(\frac{\varepsilon}{2L} B_{\|\cdot\|}(1, \mathbf{0})) \mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L})}{(4\varepsilon)^d} \\
&\quad + \frac{8}{15} \sum_{k=1}^{m_\varepsilon} \frac{\gamma \text{vol}(\frac{\varepsilon_k}{2L} B_{\|\cdot\|}(1, \mathbf{0})) \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L})}{(4\varepsilon_{k-1})^d} \\
&\geq \frac{8}{15} \frac{\gamma \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{(8L)^d} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) \\
&\quad + \frac{8}{15} \frac{\gamma \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{(16L)^d} \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}),
\end{aligned}$$

where the first inequality follows by upper bounding $f(\mathbf{x}^*) - f$ with its supremum on the partition, the second one by lower bounding the sum of disjoint layers with that of overlapping ones (Lemma 3.10), the third one by the elementary inclusions $\mathcal{X}_{2\varepsilon} \supseteq \mathcal{X}_{\frac{3}{2}\varepsilon}$ and $\mathcal{X}_{(\frac{1}{2}\varepsilon_k, 2\varepsilon_{k-1}]} \supseteq \mathcal{X}_{(\frac{1}{2}\varepsilon_k, \frac{3}{2}\varepsilon_{k-1}]}$ (for all $k \in [m_\varepsilon]$) followed by a relationship between packing numbers and volumes (Proposition 3.9) that holds under Assumption 3.8, and the last one by $\varepsilon_{k-1} \leq 2\varepsilon_k$ (for all $k \in [m_\varepsilon]$). \square

4. Robustness

In this section, we consider a harder setting in which, for each $k \geq 1$, querying the k -th value of f at a point \mathbf{x}_k reveals a perturbed evaluation $f(\mathbf{x}_k) + \alpha_k$, where $\alpha_k \in [-\alpha, \alpha]$ is the k -th element of an adversarial sequence of perturbations. Both the sequence $(\alpha_k)_{k \geq 1}$ and $\alpha \geq 0$ are assumed to be unknown in the non-certified setting, while in order to output a certificate, α has to be known in the certified setting. Furthermore, we now only assume that \mathcal{X} is bounded, lifting the assumption that \mathcal{X} is closed.

4.1. Approximate Piyavskii–Shubert algorithm

In this section, we present a natural extension of the Piyavskii–Shubert algorithm to our setting with perturbations and a bounded domain. Algorithm 3 proceeds similarly to Algorithm 2, but it only approximately maximizes the proxy function after each update. This is not only necessary when \mathcal{X} is not closed, but it could prove useful in general whenever the task of optimizing \widehat{f}_k is computationally expensive, which, as discussed in Section 2, is often the case (except in small dimensions).

Algorithm 3 Approximate Piyavskii–Shubert, for the non-certified (*/certified*) setting

input: Lipschitz constant $L > 0$, norm $\|\cdot\|$, initial guess $\mathbf{x}_1 \in \mathcal{X}$, tolerance $\eta \geq 0$ (*and magnitude of the perturbations $\alpha \geq 0$, certified setting only*)

for $k = 1, 2, \dots$ **do**

observe the perturbed value $g(\mathbf{x}_k) \leftarrow f(\mathbf{x}_k) + \alpha_k$

output the recommendation $\mathbf{x}_k^* \leftarrow \operatorname{argmax}_{\mathbf{x} \in \{\mathbf{x}_1, \dots, \mathbf{x}_k\}} g(\mathbf{x})$

update the proxy function

$$\widehat{f}_k(\cdot) \leftarrow \min_{i \in [k]} \left\{ g(\mathbf{x}_i) + L \|\mathbf{x}_i - (\cdot)\| \right\}$$

pick as the next query point any $\mathbf{x}_{k+1} \in \mathcal{X}$ such that $\widehat{f}_k(\mathbf{x}_{k+1}) \geq \sup_{\mathbf{x} \in \mathcal{X}} \widehat{f}_k(\mathbf{x}) - \eta$

output the error certificate $\xi_k \leftarrow \widehat{f}_k^* - f_k^* + (2\alpha + \eta)$, where

$$\widehat{f}_k^* \leftarrow \widehat{f}_k(\mathbf{x}_{k+1}), \quad f_k^* \leftarrow g(\mathbf{x}_k^*) \quad (\text{certified setting only})$$

end for

4.2. Upper bounds

In this section, we show that if $\alpha, \eta \ll 1$, the number of queries that the Approximate Piyavskii–Shubert algorithm needs to reach (or certify) a given accuracy are comparable to those of the Piyavskii–Shubert algorithm in the simpler setting without perturbations. In other words, the paradigm imposed by the Piyavskii–Shubert algorithm is robust to small perturbations or approximation errors. The proofs of the results of this section require more technical work but follow the same high level ideas of the corresponding ones in Sections 3.1 and 3.2. For this reason, we defer all of them to Appendix C.

We begin by stating an upper bound on the number of evaluations of f that are needed before the recommendations of the Approximate Piyavskii–Shubert algorithm

become near-optimal.

Theorem 4.1. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some unknown $\alpha \geq 0$, $\eta \geq 0$, $\varepsilon_0 := L \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$, $\varepsilon \in (0, \varepsilon_0)$, $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon}^{\alpha, \eta} := \varepsilon + 2\alpha + \eta$, $\varepsilon_k^{\alpha, \eta} := \varepsilon_0 2^{-k} + 2\alpha + \eta$ for all $k < m_\varepsilon$, and*

$$S_{\text{NC}}^{\alpha, \eta}(f, \varepsilon) := \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha, \eta}, \varepsilon_{k-1}^{\alpha, \eta}]}, \frac{\varepsilon_k^{\alpha, \eta} - 2\alpha - \eta}{L} \right).$$

If $n \geq 1 + S_{\text{NC}}^{\alpha, \eta}(f, \varepsilon)$, then the n -th recommendation \mathbf{x}_n^ of the Approximate Piyavskii–Shubert algorithm (Algorithm 3, run with input $L, \|\cdot\|, \mathbf{x}_1, \eta$) satisfies $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \varepsilon + 4\alpha + \eta$.*

Similarly to the previous result, we now generalize Theorem 3.6 showing that when the evaluations of f are perturbed by an adversarial corruption of magnitude at most α , we can control the number of queries n that the Approximate Certified Piyavskii–Shubert Algorithm needs to guarantee that the error certificate ξ_n is smaller than $\varepsilon + 2\alpha + \eta$.

Theorem 4.2. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some known $\alpha \geq 0$, $\eta \geq 0$, $\varepsilon_0 := L \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$, $\varepsilon \in (0, \varepsilon_0)$, $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon}^{\alpha, \eta} := \varepsilon + (2\alpha + \eta)$, $\varepsilon_k^{\alpha, \eta} := \varepsilon_0 2^{-k} + (2\alpha + \eta)$ for all $k < m_\varepsilon$, and*

$$S_{\text{C}}^{\alpha, \eta}(f, \varepsilon) := \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha, \eta}}, \frac{\varepsilon_{m_\varepsilon}^{\alpha, \eta} - 2\alpha}{L} \right) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha, \eta}, \varepsilon_{k-1}^{\alpha, \eta}]}, \frac{\varepsilon_k^{\alpha, \eta} - 2\alpha - \eta}{L} \right).$$

If $n \geq S_{\text{C}}^{\alpha, \eta}(f, \varepsilon)$, then the n -th error certificate ξ_n of the Approximate Certified Piyavskii–Shubert Algorithm (Algorithm 3, run with input $L, \|\cdot\|, \mathbf{x}_1, \eta, \alpha$) satisfies $\xi_n \leq \varepsilon + 2\alpha + \eta$.

5. Related Works

We now discuss related works, both about the Piyavskii–Shubert algorithm itself and close variants from the global optimization and bandit optimization literature.

Known bounds for the Piyavskii–Shubert algorithm (non-certified setting).

Several papers studied the Piyavskii–Shubert algorithm in the eighties and nineties (see, e.g., [8, 17, 18] or the survey by [9]). Despite this literature, little was known about the rate of convergence of its optimization error $f(\mathbf{x}^*) - f(\mathbf{x}_n^*)$ as a function of the number n of evaluations of f . A crude bound of the form $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) = \mathcal{O}(n^{-1/d})$, where d is the ambient dimension, can be obtained from [18, Theorem 4.2] when f is globally Lipschitz. The authors showed that the Piyavskii–Shubert algorithm is minimax optimal among all algorithms, and therefore superior to a uniform grid search ensuring an optimization error of the order of $n^{-1/d}$.

Known bounds for the Piyavskii–Shubert algorithm (certified setting). In dimension $d = 1$ an f -dependent bound on the sample complexity was derived by Hansen et al. [8] for the certified version of the Piyavskii–Shubert algorithm. Rephrasing their main result into our own terms, they proved that ε -accuracy is guaranteed (that is, $\xi_n \leq \varepsilon$) after a number of iterations at most proportional to $\int_0^1 (f(x^*) - f(x) + \varepsilon)^{-1} dx$. The authors relied heavily on the one-dimensional setting to study the proxy functions \widehat{f}_k in an explicit manner, and claimed that “*Extending the results of this paper to the multivariate case appears to be difficult*”. To the best of our knowledge, we provide the first such bound in higher dimensions $d \geq 2$.

Other packing-based bounds (with other non-certified bandit algorithms).

The non-certified version of our global Lipschitz optimization problem has been well studied in the bandit optimization literature over the past decades; see the textbooks by Munos [20] and Slivkins [26]. Several algorithms have been developed, with various f -dependent error bounds involving sets of ε -optimal points \mathcal{X}_ε or layers $\mathcal{X}_{(a,b]}$ of f . For instance, for the DOO algorithm, which can be seen as a discretized version of the Piyavskii–Shubert algorithm, Munos [19, Theorem 1] proved a bound roughly of the form $\sum_{i=1}^{\delta^{-1}(\varepsilon)} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{\delta(i)}, c\delta(i))$ for some constant $c > 0$ independent of f , and where the sequence $i \mapsto \delta(i)$ is decreasing. In the special case where $\delta(i) \approx \gamma^i$ with $\gamma \in (0, 1)$, and for (weakly) Lipschitz functions f such that $\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, c\varepsilon) \lesssim (1/\varepsilon)^{d^*}$ with $d^* \in [0, d]$ (we say that f has near-optimality dimension smaller than d^*), this sample complexity bound implies a bound roughly of $\log(1/\varepsilon)$ if $d^* = 0$ or $(1/\varepsilon)^{d^*}$ if $d^* > 0$. A similar bound was proved by Perevozchikov [22] for a branch-and-bound algorithm close to the DOO algorithm, with an assumption on the volume of \mathcal{X}_ε (instead of a packing number), or by Malherbe and Vayatis [16] for a stochastic version of the Piyavskii–Shubert algorithm under a stronger assumption on the shape of f around its maximizer. Rather than maximizing the proxy function \widehat{f}_{k-1} , the latter algorithm (called LIPO) queries the k -th point uniformly at random in the set of potential maximizers $\{x \in \mathcal{X} : \widehat{f}_{k-1}(x) \geq \max_{i=1, \dots, k-1} f(x_i)\}$.

The matching worst-case lower bounds of Nesterov [21, Theorem 1.1.2] when $d^* = d$, of Horn [11] when $d^* = d/2$, and the lower bounds of Bubeck [4] for any d^* but in the stochastic case suggest that this bound is worst-case optimal over the class of functions with given near-optimality dimension d^* . However, the bounds mentioned above have several limitations. First, as noted earlier by Kleinberg et al. [14, remark after Theorem 4.4], bounds involving a single notion of dimension such as the near-optimality or zooming dimension might be too crude. Indeed, functions that feature different shapes at different scales (such as, e.g., different values of d^* for different ranges of ε) are better analyzed in a non-asymptotic framework through sums of packing numbers at different scales. Second, considering packing numbers of the near-optimal sets $\mathcal{X}_{\delta(i)}$ as in the first bound mentioned above can also be quite suboptimal. For instance, for constant functions f , the sets $\mathcal{X}_{\delta(i)}$ are equal to the whole domain \mathcal{X} (maximal packing number) while f is optimized by any algorithm after only a single evaluation of f . The second limitation was overcome in general metric spaces by Kleinberg et al. [13, 14] by considering packing numbers of layers $\mathcal{X}_{(a,b]}$ instead of near-optimal sets \mathcal{X}_ε .

Our analysis of the Piyavskii–Shubert algorithm indicates that this very natural algorithm can also handle multiple shapes at different scales, with a sample complexity bound also involving layers of f , and that this also holds in the certified setting.

Relationship with Bachoc et al. [2]. The very recent paper [2], about optimal instance-dependent sample complexity bounds in the certified setting, was mainly written after but published before the present paper. Taking inspiration from our S_C bound and analysis, the authors prove a similar bound for a certified version of the more practical DOO algorithm. They also provide a nearly matching instance-dependent lower bound, as well as an equivalent integral representation of S_C . For the sake of completeness, we recalled these two complementary results in Sections 3.3 and 3.4, while improving the multiplicative constants in the integral representation.

Adaptivity to L . In the context of global optimization, there is a line of research devoted to maximizing L -Lipschitz functions *without* knowing L . Algorithms for these problems typically sacrifice some of the available evaluations of f to compute an estimate of the Lipschitz constant L and use the remaining ones to optimize the objective. Such estimates \hat{L} of L can be defined (up to small margins) as $\hat{L} = \max_{i \neq j} \frac{f(x_i) - f(x_j)}{x_i - x_j}$, where the maximum is over all pairs of distinct values x_i, x_j that have been queried so far. Bubeck et al. [5] present a two-phase algorithm in which the estimation of the Lipschitz constant occurs in the first phase. Their results hold for twice differentiable objectives for which the eigenvalues of the Hessian are bounded everywhere by a constant M . Notably, the bounds they obtain are meaningless unless the time horizon is bigger than a function of M (i.e., in order to apply the result in practice one would need prior knowledge of M).

Malherbe and Vayatis [16] introduce AdaLIPO, an adaptive version of LIPO (see above). Their algorithm flips a biased coin at each round. Depending on the outcome, AdaLIPO either performs a LIPO step (using the current best estimate of the Lipschitz constant) or it evaluates the function at a point drawn uniformly in the domain of f . These uniform draws are used to obtain increasingly more accurate approximations of the Lipschitz constant. Again, the bounds they obtain become meaningless if the probability of getting a good estimate of the Lipschitz constant with only two uniform queries to f (what the authors call Γ , in Proposition 18) is too small. In some papers the Lipschitzness assumption is replaced by some other regularity assumption (e.g., [3, 7] use a notion of smoothness related to hierarchical partitioning that is comparable to our Assumption 1.1).

The idea of investing some of the available evaluations to approximate the Lipschitz constant could also be applied to the Piyavskii–Shubert algorithm in the non-certified setting. However, missing any prior quantitative information on the smoothness of f has a crucial limitation. It makes it *impossible* to design algorithms that certify the accuracy of their recommendations (without any other inputs). To see this, note that *all* algorithms will fail when the objective function is the “spike” $x \mapsto \max\{0, a - L\|x - x_0\|\}$ for some values of a, L , and x_0 that depend on the algorithm (typically a and L should be chosen large enough to force the algorithm to miss a tall and narrow spike).

Acknowledgements

The authors would like to thank Jean-Baptiste Hiriart-Urruty, Marcel Mongeau, Edouard Pauwels, and Michal Valko for their helpful discussions and bibliographic insights.

Funding

Sébastien Gerchinovitz gratefully acknowledges the support of the DEEL project.⁴ Tommaso Cesari gratefully acknowledges the support of the European Network for Game Theory (GAMENET) and the project BOLD from the French national research agency (ANR). This project has also received funding from the French “Investing for the Future – PIA3” program under the Grant agreement ANITI ANR-19-P3IA-0004.

References

- [1] F. Aurenhammer and R. Klein, *Voronoi diagrams*, Handbook of computational geometry 5 (2000), pp. 201–290.
- [2] F. Bachoc, T. Cesari, and S. Gerchinovitz, *Instance-Dependent Bounds for Zeroth-order Lipschitz Optimization with Error Certificates*, in *Advances in Neural Information Processing Systems 34 proceedings (NeurIPS 2021)*. 2021.
- [3] P.L. Bartlett, V. Gabillon, and M. Valko, *A simple parameter-free and adaptive approach to optimization under a minimal local smoothness assumption*, in *Algorithmic Learning Theory*. 2019, pp. 184–206.
- [4] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári, *X-armed bandits*, Journal of Machine Learning Research 12 (2011), pp. 1655–1695.
- [5] S. Bubeck, G. Stoltz, and J.Y. Yu, *Lipschitz bandits without the Lipschitz constant*, in *International Conference on Algorithmic Learning Theory*. Springer, 2011, pp. 144–158.
- [6] A. Cuevas, R. Fraiman, and B. Pateiro-López, *On statistical properties of sets fulfilling rolling-type conditions*, Advances in Applied Probability 44 (2012), pp. 311–329.
- [7] J.B. Grill, M. Valko, and R. Munos, *Black-box optimization of noisy functions with unknown smoothness*, in *Advances in Neural Information Processing Systems*. 2015, pp. 667–675.
- [8] P. Hansen, B. Jaumard, and S.H. Lu, *On the number of iterations of Piyavskii’s global optimization algorithm*, Mathematics of Operations Research 16 (1991), pp. 334–350.
- [9] P. Hansen, B. Jaumard, and S.H. Lu, *Global optimization of univariate lipschitz functions: I. survey and properties*, Mathematical Programming 55 (1992), pp. 251–272.
- [10] P. Hansen, B. Jaumard, and S.H. Lu, *Global optimization of univariate lipschitz functions: II. new algorithms and computational comparison*, Mathematical programming 55 (1992), pp. 273–292.
- [11] M. Horn, *Optimal algorithms for global optimization in case of unknown Lipschitz constant*, Journal of Complexity 22 (2006), pp. 50–70.
- [12] Y. Hu, N. Kallus, and X. Mao, *Smooth Contextual Bandits: Bridging the Parametric and Non-differentiable Regret Regimes*, in *Proceedings of Thirty Third Conference on Learning Theory*, J. Abernethy and S. Agarwal, eds., Proceedings of Machine Learning Research Vol. 125, 09–12 Jul. PMLR, 2020, pp. 2007–2010.
- [13] R. Kleinberg, A. Slivkins, and E. Upfal, *Multi-armed bandits in metric spaces*, in *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM, 2008, pp. 681–690.
- [14] R. Kleinberg, A. Slivkins, and E. Upfal, *Bandits and experts in metric spaces*, Journal of the ACM 66 (2019).
- [15] T. Lattimore and C. Szepesvári, *Bandit algorithms*, Cambridge University Press, 2018, In press.
- [16] C. Malherbe and N. Vayatis, *Global optimization of Lipschitz functions*, in *Proceedings of the 34th International Conference on Machine Learning*, Vol. PMLR 70. 2017, pp. 2314–2323.

⁴<https://www.deel.ai/>

- [17] D.Q. Mayne and E. Polak, *Outer approximation algorithm for nondifferentiable optimization problems*, Journal of Optimization Theory and Applications 42 (1984), pp. 19–30.
- [18] R.H. Mladineo, *An algorithm for finding the global maximum of a multimodal, multivariate function*, Mathematical Programming 34 (1986), pp. 188–200.
- [19] R. Munos, *Optimistic Optimization of a Deterministic Function without the Knowledge of its Smoothness*, in *Advances in Neural Information Processing Systems 24 (NIPS 2011)*. 2011, pp. 783–791.
- [20] R. Munos, *From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning*, Foundations and Trends® in Machine Learning 7 (2014), pp. 1–129.
- [21] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*, Vol. 87, Springer Science & Business Media, 2003.
- [22] A. Perevozchikov, *The complexity of the computation of the global extremum in a class of multi-extremum problems*, USSR Computational Mathematics and Mathematical Physics 30 (1990), pp. 28–33.
- [23] S. Piyavskii, *An algorithm for finding the absolute extremum of a function*, USSR Computational Mathematics and Mathematical Physics 12 (1972), pp. 57–67.
- [24] R.T. Rockafellar, *Extensions of subgradient calculus with applications to optimization*, Nonlinear Analysis: Theory, Methods & Applications 9 (1985), pp. 665 – 698.
- [25] B.O. Shubert, *A sequential method seeking the global maximum of a function*, SIAM Journal on Numerical Analysis 9 (1972), pp. 379–388.
- [26] A. Slivkins, *Introduction to multi-armed bandits*, arXiv preprint arXiv:1904.07272 (2019).
- [27] M.J. Wainwright, *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2019.
- [28] G. Walther, *Granulometric smoothing*, The Annals of Statistics 25 (1997), pp. 2273–2299.

Appendix A. Standard inequalities about packing and covering numbers

We recall that for any norm $\|\cdot\|$, bounded set $A \subset \mathbb{R}^d$, and $r > 0$:

- The r -packing number of A (with respect to $\|\cdot\|$) is the largest number of r -separated points contained in A , i.e.,

$$\mathcal{N}_{\|\cdot\|}(A, r) := \sup \left\{ k \geq 1 : \exists \mathbf{x}_1, \dots, \mathbf{x}_k \in A, \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\| > r \right\},$$

if A is nonempty, zero otherwise

- The r -covering number of A (with respect to $\|\cdot\|$) is the smallest number of points needed to cover A with balls with radius r centered at those points, i.e.,

$$\mathcal{M}_{\|\cdot\|}(A, r) := \min \{ k \geq 1 : \exists \mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^d, \forall \mathbf{x} \in A, \exists i \in [k], \|\mathbf{x} - \mathbf{x}_i\| \leq r \}.$$

if A is nonempty, zero otherwise

Covering and packing numbers are closely related. In particular, the following well-known inequalities hold—see, e.g., [27, Lemmas 5.5 and 5.7, with permuted notations for $\mathcal{M}_{\|\cdot\|}$ and $\mathcal{N}_{\|\cdot\|}$].⁵

⁵The definition of r -covering number of a subset A of \mathbb{R}^d implied by [27, Definition 5.1] is slightly stronger than the one used in our paper, because elements x_1, \dots, x_N of r -covers belong to A rather than just \mathbb{R}^d .

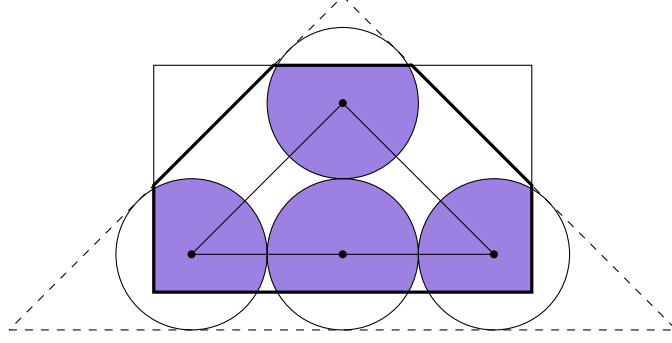


Figure B1. An illustration that for any $\eta > 0$ (the radii of the circles) and $A \subset \mathcal{X}$ (inner triangle and rectangle respectively), the balls of radius $\eta/2$ centered at the elements of an η -packing of A (the circles centered at the dots on the perimeter of A) intersected with \mathcal{X} are all disjoint and included in $(A + B_{\|\cdot\|}(\eta/2, \mathbf{0})) \cap \mathcal{X}$ (the outer dashed triangle intersected with the rectangle, which give the bold hexagon containing the colored figures).

Lemma A.1. For any bounded set $A \subset \mathbb{R}^d$ and any real number $r > 0$,

$$\mathcal{N}_{\|\cdot\|}(A, 2r) \leq \mathcal{M}_{\|\cdot\|}(A, r) \leq \mathcal{N}_{\|\cdot\|}(A, r). \quad (\text{A1})$$

Appendix B. Missing proofs of Section 3.4

In this section, we include all missing details from Section 3.4.

Proof of Proposition 3.9. Fix any $0 < w < u < \varepsilon_0$. Let $\eta_1 := \frac{w}{L}$, $\eta_2 := \frac{w}{L}$, $E_1 := \mathcal{X}$, $E_2 := \mathcal{X}_w^c$, and $i \in [2]$. Note that for any $\eta > 0$ and $A \subset \mathcal{X}$, the balls of radius $\eta/2$ centered at the elements of an η -packing of A intersected with \mathcal{X} are all disjoint and included in $(A + B_{\|\cdot\|}(\eta/2, \mathbf{0})) \cap \mathcal{X}$ (see Figure B1). Thus, letting P_i be a set of η_i -separated points included in $A_i := \mathcal{X}_u \cap E_i$ with cardinality $|P_i| = \mathcal{N}_{\|\cdot\|}(A_i, \eta_i)$, we have

$$\begin{aligned} \text{vol}\left((A_i + B_{\|\cdot\|}(\eta_i/2, \mathbf{0})) \cap \mathcal{X}\right) &\geq \sum_{\mathbf{x} \in P_i} \text{vol}(B_{\|\cdot\|}(\eta_i/2, \mathbf{x}) \cap \mathcal{X}) \\ &\geq \gamma \text{vol}(B_{\|\cdot\|}(\eta_i/2, \mathbf{0})) \mathcal{N}_{\|\cdot\|}(A_i, \eta_i), \end{aligned}$$

where the second inequality follows by Assumption 3.8. We now further upper bound the left-hand side. Take an arbitrary point $\mathbf{x}_i \in (\mathcal{X}_u \cap E_i + B_{\|\cdot\|}(\eta_i/2, \mathbf{0})) \cap \mathcal{X}$. By definition of Minkowski sum, there exists $\mathbf{x}'_i \in \mathcal{X}_u \cap E_i$ such that $\|\mathbf{x}_i - \mathbf{x}'_i\| \leq \eta_i/2$. Hence $f(\mathbf{x}^*) - f(\mathbf{x}_i) \leq f(\mathbf{x}^*) - f(\mathbf{x}'_i) + |f(\mathbf{x}'_i) - f(\mathbf{x}_i)| \leq u + L(\eta_i/2) \leq (3/2)u$. This implies that $\mathbf{x}_i \in \mathcal{X}_{(3/2)u}$, which proves the first inequality. For the second one, note that \mathbf{x}_2 satisfies $f(\mathbf{x}^*) - f(\mathbf{x}_2) \geq f(\mathbf{x}^*) - f(\mathbf{x}'_2) - |f(\mathbf{x}'_2) - f(\mathbf{x}_2)| \geq w - L(\eta_2/2) = (1/2)w$. \square

Proof of Lemma 3.10. Fix any \mathcal{X} , f , and ε as in the statement of the lemma. To avoid clutter, we denote m_ε simply by m .

Assume first that $m \geq 3$. Then

$$\begin{aligned}
& \frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \leq \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon^d} \\
& \quad + \sum_{k=1}^{m-2} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_{k+1}, \varepsilon_k]}) + \text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{k-1}, \varepsilon_{k-2}]})}{\varepsilon_{k-1}^d} \\
& \quad + \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-2}, \varepsilon_{m-3}]})}{\varepsilon_{m-2}^d} \\
& \quad + \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon_{m-1}^d} \\
& \leq \left(2 + \frac{1}{2^d}\right) \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \left(2^d + 1 + \frac{1}{2^d}\right) \frac{\text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]})}{\varepsilon_{m-1}^d} \\
& \quad + (4^d + 2^d + 1) \frac{\text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon_{m-2}^d} \\
& \quad + \frac{1}{2^d} \sum_{k=2}^{m-1} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} + \sum_{k=1}^{m-2} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} + 2^d \sum_{k=1}^{m-2} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \\
& = \left(2 + \frac{1}{2^d}\right) \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + 4^d \frac{\text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon_{m-2}^d} \\
& \quad + \frac{1}{2^d} \sum_{k=2}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} + 2^d \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \\
& \leq \left(2 + \frac{1}{2^d}\right) \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \left(4^d + \frac{1}{2^d} + 1 + 2^d\right) \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d}
\end{aligned}$$

where we applied several times the definitions of and the relations between the ε_k 's. The bound then follows after noting that $2 + 2^{-d} \leq 4^d + 2^{-d} + 1 + 2^d \leq (15/8) \cdot 4^d$.

Assume now that $m = 2$. Then

$$\begin{aligned}
& \frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \\
& \leq \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon^d} \\
& \quad + \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon_{m-2}^d} \\
& \quad + \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]}) + \text{vol}(\mathcal{X}_{(\varepsilon_{m-1}, \varepsilon_{m-2}]})}{\varepsilon_{m-1}^d} \\
& \leq \left(2 + \frac{1}{2^d}\right) \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \left(4^d + 1 + 2^d\right) \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d}
\end{aligned}$$

and the bound follows from $2 + 2^{-d} \leq 4^d + 1 + 2^d \leq (7/4) \cdot 4^d \leq (15/8) \cdot 4^d$.

Finally, assume that $m = 1$. Then

$$\begin{aligned} \frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} &\leq \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]})}{\varepsilon^d} \\ &\quad + \frac{\text{vol}(\mathcal{X}_\varepsilon) + \text{vol}(\mathcal{X}_{(\varepsilon_m, \varepsilon_{m-1}]})}{\varepsilon_{m-1}^d} \\ &\leq 2 \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + 2^d \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \end{aligned}$$

and the bound follows from $2 \leq 2^d \leq (15/8) \cdot 4^d$. This concludes the proof. \square

Theorem B.1. Fix any norm $\|\cdot\|$, $\gamma \in (0, 1]$, and $L > 0$. Then there exist two positive constants $c, C > 0$ depending only on $d, \gamma, L, \|\cdot\|$ (e.g., $c := L^d/\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))$ and $C := (15/8)(32L)^d/\gamma\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))$) such that:

- (1) For any subset \mathcal{X} of \mathbb{R}^d that is compact and γ -approximately d -dimensional with respect to $\|\cdot\|$,
- (2) For any function $f: \mathcal{X} \rightarrow \mathbb{R}$ that is globally L -Lipschitz with respect to $\|\cdot\|$,

and for all positive $\varepsilon < \varepsilon_0 := L\text{diam}_{\|\cdot\|}(\mathcal{X})$, the following inequalities hold:⁶

$$c \int_{\mathcal{X}} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d} \leq S_C(f, \varepsilon) \leq C \int_{\mathcal{X}} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d}.$$

Proof. Fix any $\mathcal{X} \subset \mathbb{R}^d$ and $f: \mathcal{X} \rightarrow \mathbb{R}$ satisfying the assumptions of the theorem, and let $\varepsilon < \varepsilon_0$. Let also $m_\varepsilon := \lceil \log_2(\varepsilon_0/\varepsilon) \rceil$, $\varepsilon_{m_\varepsilon} := \varepsilon$, and $\varepsilon_k := \varepsilon_0 2^{-k}$ for all $k < m_\varepsilon$. Partition the domain of integration \mathcal{X} into the following $m_\varepsilon + 1$ sets: the set of ε -optimal points \mathcal{X}_ε and the m_ε layers $\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}$, for $k \in [m_\varepsilon]$.

We begin by proving the first inequality. To this end, for any bounded $A \subset \mathbb{R}^d$ and $r > 0$, we define the r -covering number $\mathcal{M}_{\|\cdot\|}(A, r)$ of A (with respect to $\|\cdot\|$) as the smallest cardinality of an r -covering of A , i.e.,

$$\mathcal{M}_{\|\cdot\|}(A, r) := \min\{k \in \{1, 2, \dots\} : \exists \mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^d, \forall \mathbf{x} \in A, \exists i \in [k], \|\mathbf{x} - \mathbf{x}_i\| \leq r\}.$$

⁶We actually prove a stronger result. The first inequality holds more generally for any f that is L -Lipschitz around a maximizer with respect to $\|\cdot\|$ and Lebesgue-measurable, and does not require \mathcal{X} to be approximately d -dimensional.

if A is nonempty, zero otherwise. Then,

$$\begin{aligned}
\int_{\mathcal{X}} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d} &\leq \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{(\varepsilon_k + \varepsilon)^d} \\
&\leq \frac{\mathcal{M}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) \cdot \text{vol}(B_{\|\cdot\|}(1, \mathbf{0})) (\frac{\varepsilon}{L})^d}{\varepsilon^d} \\
&\quad + \sum_{k=1}^{m_\varepsilon} \frac{\mathcal{M}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}) \cdot \text{vol}(B_{\|\cdot\|}(1, \mathbf{0})) (\frac{\varepsilon_k}{L})^d}{\varepsilon_k^d} \\
&\leq \frac{\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{L^d} \left(\mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}) \right) \\
&= \frac{\text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{L^d} S_C(f, \varepsilon),
\end{aligned}$$

where the first inequality follows by lower bounding $f(\mathbf{x}^*) - f$ with its infimum on the partition, the second one by dropping $\varepsilon > 0$ from the second denominator and upper bounding the volume of a set with the volume of the balls of a smallest ε_k/L -covering, and the last one by the fact that covering numbers are always smaller than packing numbers (we recall this known result in Appendix A, eq. (A1)). This proves the first part of the theorem.

For the second one, we have

$$\begin{aligned}
\int_{\mathcal{X}} \frac{d\mathbf{x}}{(f(\mathbf{x}^*) - f(\mathbf{x}) + \varepsilon)^d} &\geq \frac{\text{vol}(\mathcal{X}_\varepsilon)}{(\varepsilon + \varepsilon)^d} + \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{(\varepsilon_{k-1} + \varepsilon)^d} \\
&\geq \frac{1}{2^d} \frac{\text{vol}(\mathcal{X}_\varepsilon)}{\varepsilon^d} + \frac{1}{2^d} \sum_{k=1}^{m_\varepsilon} \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \\
&\geq \frac{8}{15} \frac{1}{8^d} \left(\frac{\text{vol}(\mathcal{X}_{2\varepsilon})}{\varepsilon^d} + \sum_{k=1}^m \frac{\text{vol}(\mathcal{X}_{(\varepsilon_k/2, 2\varepsilon_{k-1}]})}{\varepsilon_{k-1}^d} \right) \\
&\geq \frac{8}{15} \frac{\gamma \text{vol}(\frac{\varepsilon}{2L} B_{\|\cdot\|}(1, \mathbf{0})) \mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L})}{(8\varepsilon)^d} + \\
&\quad \frac{8}{15} \sum_{k=1}^{m_\varepsilon} \frac{\gamma \text{vol}(\frac{\varepsilon_k}{2L} B_{\|\cdot\|}(1, \mathbf{0})) \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L})}{(8\varepsilon_{k-1})^d} \\
&\geq \frac{8}{15} \frac{\gamma \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{(16L)^d} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_\varepsilon, \frac{\varepsilon}{L}) \\
&\quad + \frac{8}{15} \frac{\gamma \text{vol}(B_{\|\cdot\|}(1, \mathbf{0}))}{(32L)^d} \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|}(\mathcal{X}_{(\varepsilon_k, \varepsilon_{k-1}]}, \frac{\varepsilon_k}{L}),
\end{aligned}$$

where the first inequality follows by upper bounding $f(\mathbf{x}^*) - f$ with its supremum on the partition, the second one by $\varepsilon \leq \varepsilon_{k-1}$ (for all $k \in [m_\varepsilon + 1]$), the third one by lower bounding the sum of disjoint layers with that of overlapping ones (proved in the appendix, Lemma 3.10), the fourth one by the elementary inclusions $\mathcal{X}_{2\varepsilon} \supseteq \mathcal{X}_{\frac{3}{2}\varepsilon}$ and $\mathcal{X}_{(\frac{1}{2}\varepsilon_k, 2\varepsilon_{k-1})} \supseteq \mathcal{X}_{(\frac{1}{2}\varepsilon_k, \frac{3}{2}\varepsilon_{k-1})}$ (for all $k \in [m_\varepsilon]$) followed by a relationship between

packing numbers and volumes (proved in the appendix, Proposition 3.9) that holds under Assumption 3.8, and the last one by $\varepsilon_{k-1} \leq 2\varepsilon_k$ (for all $k \in [m_\varepsilon]$). \square

Appendix C. Missing proofs of Section 4.2

In this section, we include all missing details from Section 4.2.

C.1. Useful lemmas: general version

In this section, we extend the lemmas we proved in Section 3.1 to the case of the Approximate Piyavskii–Shubert algorithm run in a setting with perturbed observations.

Lemma C.1. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some $\alpha \geq 0$, and the Approximate Piyavskii–Shubert algorithm (Algorithm 3) is run with input $L, \|\cdot\|, \mathbf{x}_1, \eta$. Then, for all $k \geq 1$, the proxy function $\widehat{f}_k(\cdot) = \min_{i \in [k]} \{g(\mathbf{x}_i) + L\|\mathbf{x}_i - (\cdot)\|\}$ is L -Lipschitz and satisfies*

$$\widehat{f}_k(\mathbf{x}^*) \geq f(\mathbf{x}^*) - \alpha \quad \text{and} \quad \widehat{f}_k(\mathbf{x}_j) \leq f(\mathbf{x}_j) + \alpha, \quad \forall j \in [k].$$

Proof. Fix any $k \geq 1$. The fact that \widehat{f}_k is L -Lipschitz is straightforward. Moreover, by the L -Lipschitzness of f around \mathbf{x}^* , we get

$$\begin{aligned} \widehat{f}_k(\mathbf{x}^*) &= \min_{i \in [k]} \left\{ f(\mathbf{x}_i) + \alpha_i + L\|\mathbf{x}_i - \mathbf{x}^*\| \right\} \\ &\geq \min_{i \in [k]} \left\{ (f(\mathbf{x}^*) - L\|\mathbf{x}_i - \mathbf{x}^*\|) - \alpha + L\|\mathbf{x}_i - \mathbf{x}^*\| \right\} = f(\mathbf{x}^*) - \alpha. \end{aligned}$$

Furthermore, for all $j \in [k]$, the definition of $\widehat{f}_k(\mathbf{x}_j)$ implies immediately

$$\widehat{f}_k(\mathbf{x}_j) = \min_{i \in [k]} \left\{ f(\mathbf{x}_i) + \alpha_i + L\|\mathbf{x}_i - \mathbf{x}_j\| \right\} \leq f(\mathbf{x}_j) + \alpha + L\|\mathbf{x}_j - \mathbf{x}_j\| = f(\mathbf{x}_j) + \alpha. \quad \square$$

A first immediate consequence of Lemma C.1 is that the certified version of the Piyavskii–Shubert algorithm is a certified algorithm.

Lemma C.2. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some known $\alpha \geq 0$, and the Approximate Certified Piyavskii–Shubert Algorithm (Algorithm 3) is run with input $L, \|\cdot\|, \mathbf{x}_1, \eta, \alpha$. Then, for all $n \geq 1$, $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq \xi_n$.*

Proof. By Lemma C.1 and the definitions of $f_n^* := g(\mathbf{x}_n^*)$, $\widehat{f}_n^* := \widehat{f}_n(\mathbf{x}_{n+1}) \geq \sup \widehat{f}_n - \eta$, and $\xi_n := \widehat{f}_n^* - f_n^* + (2\alpha + \eta)$, for all $n \geq 1$, we get $f(\mathbf{x}^*) - f(\mathbf{x}_n^*) \leq (\widehat{f}_n(\mathbf{x}^*) + \alpha) - (f_n^* - \alpha) \leq \xi_n$. \square

Another important implication of Lemma C.1 is that if the Piyavskii–Shubert algorithm observes f at a Δ -suboptimal point \mathbf{x}_i , then all the next query points \mathbf{x}_j are at

least (Δ/L) -distant from \mathbf{x}_i . In other words, the algorithm does not explore too much in suboptimal regions.

Lemma C.3. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some $\alpha \geq 0$, and the Approximate Piyavskii–Shubert algorithm (Algorithm 3) is run with input $L, \|\cdot\|, \mathbf{x}_1, \eta$. Fix any $\Delta > 0$ and assume that there exists $i \geq 1$ such that the i -th query point \mathbf{x}_i belongs to \mathcal{X}_Δ^c . Then, for all $j > i$, the j -th query point satisfies*

$$\|\mathbf{x}_j - \mathbf{x}_i\| > \frac{\Delta - 2\alpha - \eta}{L}.$$

Proof. Fix any $j > i$. Then, using the fact that \mathbf{x}_j has been selected as the j -th query point (in the first inequality), Lemma C.1 (in the second and fourth), and $\mathbf{x}_i \in \mathcal{X}_\Delta^c$ (in the third), we get

$$\widehat{f}_{j-1}(\mathbf{x}_j) \geq \widehat{f}_{j-1}(\mathbf{x}^*) - \eta \geq f(\mathbf{x}^*) - \alpha - \eta > f(\mathbf{x}_i) + \Delta - \alpha - \eta \geq \widehat{f}_{j-1}(\mathbf{x}_i) + \Delta - 2\alpha - \eta.$$

Since \widehat{f}_{j-1} is L -Lipschitz (by Lemma C.1), we can thus conclude that $L\|\mathbf{x}_j - \mathbf{x}_i\| \geq |\widehat{f}_{j-1}(\mathbf{x}_j) - \widehat{f}_{j-1}(\mathbf{x}_i)| > \Delta - 2\alpha - \eta$. \square

As a final corollary to Lemma C.1, we show that the Approximate Certified Piyavskii–Shubert Algorithm automatically adapts its exploration to all possible approximation scales at the same time. In words, for any accuracy ε , until the algorithm can certify that the approximation error is $\varepsilon + 2\alpha + \eta$ or better (i.e., until the error certificate is smaller than or equal to $\varepsilon + 2\alpha + \eta$), it only queries points that are at least $(\varepsilon - 2\alpha)/L$ -distant from all past query points. The strength of this statement comes from the fact that the algorithm *does not* take an accuracy ε as a parameter and the result is consequently not proved for a specific and known ε : it holds for all ε simultaneously.

Lemma C.4. *Assume that $f: \mathcal{X} \rightarrow \mathbb{R}$ is L -Lipschitz around a maximizer \mathbf{x}^* with respect to a norm $\|\cdot\|$ (Assumption 1.1), queries are corrupted by additive perturbations with absolute value bounded by some known $\alpha \geq 0$, and the Certified Piyavskii–Shubert algorithm (Algorithm 3) is run with input $L, \|\cdot\|, \mathbf{x}_1, \alpha$. Fix any $\varepsilon > 0$ and let τ be the first time where the algorithm returns an error certificate $\xi_\tau \leq \varepsilon + 2\alpha + \eta$. Then, for all distinct $i, j \in [\tau]$,*

$$\|\mathbf{x}_i - \mathbf{x}_j\| > \frac{\varepsilon - 2\alpha}{L}.$$

Proof. Without loss of generality, assume $1 \leq i < j \leq \tau$. Then, by the definition of \widehat{f}_{j-1}^* (first equality), that of ξ_{j-1} together with the assumption $\xi_{j-1} > \varepsilon + 2\alpha + \eta$ (first inequality), the definition of f_{j-1}^* (second), and Lemma C.1 (last), we get

$$\widehat{f}_{j-1}(\mathbf{x}_j) = (\widehat{f}_{j-1}^* - f_{j-1}^*) + f_{j-1}^* > \varepsilon + f_{j-1}^* \geq \varepsilon + f(\mathbf{x}_i) - \alpha \geq \varepsilon + \widehat{f}_{j-1}(\mathbf{x}_i) - 2\alpha.$$

Since \widehat{f}_{j-1} is L -Lipschitz (by Lemma C.1), we can thus conclude that $L\|\mathbf{x}_j - \mathbf{x}_i\| \geq |\widehat{f}_{j-1}(\mathbf{x}_j) - \widehat{f}_{j-1}(\mathbf{x}_i)| > \varepsilon - 2\alpha$. \square

The previous result shows that ξ_n serves a synergistic double-role in the Approximate Certified Piyavskii–Shubert Algorithm. When it is large, it guarantees that the queries are sufficiently sparse, and as soon as it becomes small, it certifies that the recommendation \mathbf{x}_n is sufficiently good.

C.2. Proofs of Theorems 4.1 and 4.2

Proof of Theorem 4.1. By definition of $\varepsilon_k^{\alpha,\eta}$, we can partition the set $\mathcal{X}_{\varepsilon+(2\alpha+\eta)}^c = \mathcal{X}_{(\varepsilon+(2\alpha+\eta), \varepsilon_0]}$ of $(\varepsilon + 2\alpha + \eta)$ -suboptimal points into m_ε layers

$$\mathcal{X}_{\varepsilon+(2\alpha+\eta)}^c = \bigcup_{k=1}^{m_\varepsilon} \mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}$$

Therefore, for any $i \geq 1$, if $\mathbf{x}_i \in \mathcal{X}_{\varepsilon+(2\alpha+\eta)}^c$, then there exists a unique $k \in [m_\varepsilon]$ such that $\mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}$. Since the Piyavskii–Shubert algorithm does not explore too much in suboptimal regions, Lemma C.3 and the definition of packing numbers imply that for all $k \in [m_\varepsilon]$, the maximum number of times that a query point \mathbf{x}_i is chosen in $\mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]} \subset \mathcal{X}_{\varepsilon_k^{\alpha,\eta}}^c$ is upper bounded by $\mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}, \frac{\varepsilon_k^{\alpha,\eta} - 2\alpha - \eta}{L} \right)$. Hence

$$\begin{aligned} \left| \{i \in \{1, 2, \dots\} : \mathbf{x}_i \in \mathcal{X}_{\varepsilon+(2\alpha+\eta)}^c\} \right| &= \sum_{k=1}^{m_\varepsilon} \left| \{i \in \{1, 2, \dots\} : \mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}\} \right| \\ &\leq \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}, \frac{\varepsilon_k^{\alpha,\eta} - 2\alpha - \eta}{L} \right) = S_{\text{NC}}(f, \varepsilon). \end{aligned}$$

Therefore, if $n \geq 1 + S_{\text{NC}}^{\alpha,\eta}(f, \varepsilon)$, then there exists $i \geq 1$ such that $\mathbf{x}_i \in \mathcal{X}_{\varepsilon+(2\alpha+\eta)}$, which implies (by definition of \mathbf{x}_n^*) that $f(\mathbf{x}_n^*) \geq f(\mathbf{x}_i) - 2\alpha \geq f(\mathbf{x}^*) - \varepsilon - (2\alpha + \eta) - 2\alpha$ and concludes the proof. \square

Proof of Theorem 4.2. By definition of $\varepsilon_k^{\alpha,\eta}$, we can partition \mathcal{X} into the $m_\varepsilon + 1$ sets

$$\mathcal{X} = \mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha,\eta}} \cup \bigcup_{k=1}^{m_\varepsilon} \mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}$$

Let τ be the first time where the algorithm returns an error certificate $\xi_\tau \leq \varepsilon + 2\alpha + \eta$. Then, for any $i \in [\tau]$, either $\mathbf{x}_i \in \mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha,\eta}}$ or there exists a unique $k \in [m_\varepsilon]$ such that $\mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}$. Since the Certified Piyavskii–Shubert algorithm automatically adapts its exploration to all possible approximation scales at the same time, Lemma C.4 and the definition of packing numbers imply that the maximum number of times that a query point \mathbf{x}_i is chosen in $\mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha,\eta}}$ (up to and including time τ) is upper bounded by $\mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha,\eta}}, \frac{\varepsilon_{m_\varepsilon}^{\alpha,\eta} - 2\alpha}{L} \right)$. Similarly, since the Certified Piyavskii–Shubert algorithm does not explore too much in suboptimal regions, Lemma C.3 and the definition of packing numbers imply that for all $k \in [m_\varepsilon]$, the maximum number of times that a query point \mathbf{x}_i is chosen in $\mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}$ is upper bounded by $\mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha,\eta}, \varepsilon_{k-1}^{\alpha,\eta}]}, \frac{\varepsilon_k^{\alpha,\eta} - 2\alpha - \eta}{L} \right)$.

Hence

$$\begin{aligned} \tau &= \left| \{i \in [\tau] : \mathbf{x}_i \in \mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha, \eta}}\} \right| + \sum_{k=1}^{m_\varepsilon} \left| \{i \in [\tau] : \mathbf{x}_i \in \mathcal{X}_{(\varepsilon_k^{\alpha, \eta}, \varepsilon_{k-1}^{\alpha, \eta})}\} \right| \\ &\leq \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{\varepsilon_{m_\varepsilon}^{\alpha, \eta}}, \frac{\varepsilon_{m_\varepsilon}^{\alpha, \eta} - 2\alpha}{L} \right) + \sum_{k=1}^{m_\varepsilon} \mathcal{N}_{\|\cdot\|} \left(\mathcal{X}_{(\varepsilon_k^{\alpha, \eta}, \varepsilon_{k-1}^{\alpha, \eta})}, \frac{\varepsilon_k^{\alpha, \eta} - 2\alpha - \eta}{L} \right) = S_C^{\alpha, \eta}(f, \varepsilon). \end{aligned}$$

Therefore, noting that $k \mapsto \xi_k$ is non-increasing (because $\xi_k + 2\alpha + \eta = \widehat{f}_k^* - f_k^* + 2\alpha + \eta$, with $k \mapsto \widehat{f}_k^*$ non-increasing and $k \mapsto f_k^*$ non-decreasing), if $n \geq S_C(f, \varepsilon)$, then $n \geq \tau$, which in turn yields $\xi_n \leq \xi_\tau \leq \varepsilon + 2\alpha + \eta$. \square