

A simple parameter-free and adaptive approach to optimization under a minimal local smoothness assumption

Peter L. Bartlett

University of California, Berkeley, USA

PETER@BERKELEY.EDU

Victor Gabillon

Queensland University of Technology - ACEMS, Australia

VICTOR.GABILLON@QUT.EDU.AU

Michal Valko

Sequel team, INRIA Lille - Nord Europe, France

MICHAL.VALKO@INRIA.FR

Editor: ?

Abstract

We study the problem of optimizing a function under a *budgeted number of evaluations*. We only assume that the function is *locally* smooth around one of its global optima. The difficulty of optimization is measured in terms of 1) the amount of *noise* b of the function evaluation and 2) the local smoothness, d , of the function. A smaller d results in smaller optimization error. We come with a new, simple, and parameter-free approach. First, for all values of b and d , this approach recovers at least the state-of-the-art regret guarantees. Second, our approach additionally obtains these results while being *agnostic* to the values of both b and d . This leads to the first algorithm that naturally adapts to an *unknown* range of noise b and leads to significant improvements in a moderate and low-noise regime. Third, our approach also obtains a remarkable improvement over the state-of-the-art **S00** algorithm when the noise is very low which includes the case of optimization under deterministic feedback ($b = 0$). There, under our minimal local smoothness assumption, this improvement is of exponential magnitude and holds for a class of functions that covers the vast majority of functions that practitioners optimize ($d = 0$). We show that our algorithmic improvement is also borne out in the numerical experiments, where we empirically show faster convergence on common benchmark functions.

Keywords: optimization, tree search, deterministic feedback, stochastic feedback

1. Introduction

In budgeted function optimization, a learner optimizes a function $f : \mathcal{X} \rightarrow \mathbb{R}$ having access to a number of evaluations limited by n . For each of the n evaluations (or rounds), at round t , the learner picks an element $x_t \in \mathcal{X}$ and observes a real number y_t , where $y_t = f(x_t) + \varepsilon_t$, where ε_t is the noise. Based on ε_t , we distinguish two feedback cases:

Deterministic feedback The evaluations are noiseless, that is $\forall t, \varepsilon_t = 0$ and $y_t = f(x_t)$.

Please refer to the work by [de Freitas et al. \(2012\)](#) for a motivation, many applications, and references on the importance of the $b = 0$.

Stochastic feedback The evaluations are perturbed by a noise of range $b \in \mathbb{R}_+$ ¹: At any round, ε_t is a random variable, assumed to be independent of the noise at previous rounds,

$$\mathbb{E}[y_t|x_t] = f(x_t) \quad \text{and} \quad |y_t - f(x_t)| \leq b. \quad (1)$$

The objective of the learner is to return an element $x(n) \in \mathcal{X}$ with largest possible value $f(x(n))$ after the n evaluations. $x(n)$ can be different from the last evaluated element x_n . More precisely, the performance of the algorithm is the loss (or simple regret),

$$r_n = \sup_{x \in \mathcal{X}} f(x) - f(x(n)).$$

We consider the case that the evaluation is costly. Therefore, we minimize r_n as a function of n . We assume that there exists at least one point $x^* \in \mathcal{X}$ such that $f(x^*) = \sup_{x \in \mathcal{X}} f(x)$.

Prior work Among the large work on optimization, we focus on algorithms that perform well under *minimal* assumptions as well as minimal knowledge about the function. Relying on minimal assumptions means that we target functions that are particularly hard to optimize. For instance, we may not have access to the gradients of the function, gradients might not be well defined, or the function may not be continuous. While some prior works assume a *global* smoothness of the function (Pintér, 2013; Strongin and Sergeev, 2013; Hansen and Walster, 2003; Kearfott, 2013), another line of research assumes only a *weak/local* smoothness around one global maximum (Kleinberg et al., 2008; Bubeck et al., 2011a). However, within this latter group, some algorithms require the knowledge of the local smoothness such as H00 (Bubeck et al., 2011a), Zooming (Kleinberg et al., 2008), or D00 (Munos, 2011). Among the works relying on an *unknown* local smoothness, S00 (Munos, 2011; Kawaguchi et al., 2016) represents the state-of-the-art for the deterministic feedback. For the stochastic feedback, StoS00 (Valko et al., 2013) extends S00 for a limited class of functions. P00 (Grill et al., 2015) provides more general results. We classify the most related algorithms in the following table. Note that, for more specific assumptions on the smoothness, some works

smoothness	deterministic	stochastic
known	D00	Zooming, H00
unknown	DiRect, S00, Sequ00L	StoS00, P00, Stroqu00L

study optimization without the knowledge of smoothness: DiRect (Jones et al., 1993) and others (Slivkins, 2011; Bubeck et al., 2011b; Malherbe and Vayatis, 2017) tackle Lipschitz optimization.

Finally, there are algorithms that instead of simple regret, optimize *cumulative regret*, for example, H00 (Bubeck et al., 2011a) or HCT (Azar et al., 2014). However, none of them adapt to the unknown smoothness and compared to them, the algorithms for simple regret that are able to do that, such as P00 or our Stroqu00L, need to explore significantly more, which negatively impacts their cumulative regret (Grill et al., 2015; Locatelli and Carpentier, 2018).

1. Alternatively, we can turn the boundedness assumption into a sub-Gaussianity assumption equipped with a variance parameter equivalent to our range b .

Existing tools **Partitioning and near-optimality dimension:** As in most of the previously mentioned work, the search domain \mathcal{X} is partitioned into cells at different scales (depths), i.e., at a deeper depth, the cells are smaller but still cover all of \mathcal{X} . The objective of many algorithms is to explore the value of f in the cells of the partition and determine at the deepest *depth* possible in which cell is a global maximum of the function. The notion of near-optimality dimension d characterizes the complexity of the optimization task. We adopt the definition of near-optimality dimension given recently by Grill et al. (2015) that unlike Bubeck et al. (2011a), Valko et al. (2013), Munos (2011), and Azar et al. (2014), avoids topological notions and does not artificially attempt to separate the difficulty of the optimization from the partitioning. For each depth h , it simply counts the number of near-optimal cells \mathcal{N}_h , cells whose value is close to $f(x^*)$, and determines how this number evolves with the depth h . The smaller d , the more accurate the optimization should be.

New challenges **Adaptations to different data complexities:** As did Bubeck and Slivkins (2012), Seldin and Slivkins (2014), and De Rooij et al. (2014) in other contexts, we design algorithms that demonstrate near-optimal behavior under data-generating processes of different nature, obtaining the *best of all these possible worlds*. In this paper, we consider the two following data complexities for which we bring new improved adaptation.

- *near-optimality dimension $d = 0$:* In this case, the number of near-optimal cells is simply bounded by a constant that does not depend on h . As shown by Valko et al. (2013), if the function is lower- and upper-bounded by two polynomial envelopes of the same order around a global optimum, then $d = 0$. As discussed in the book by Munos (2014, section 4.2.2), $d = 0$ covers the vast majority of functions that practitioners optimize and the functions with $d > 0$ given as examples in prior work (Bubeck et al., 2011b; Grill et al., 2015; Valko et al., 2013; Munos, 2011) are carefully engineered. Therefore, the case of $d = 0$ is of practical importance. However, even with deterministic feedback, the case $d = 0$ with unknown smoothness has not been known to have a learner with a near-optimal guarantee. In this paper, we also provide that. Our approach not only adapts very well to the case $d = 0$ and $b \approx 0$, it also provides an exponential improvement over the state of the art for the simple regret rate.
- *low or moderate noise regime:* When facing a noisy feedback, most algorithms assume that the noise is of a *known* predefined range, often using $b = 1$ hard-coded in their use of upper confidence bounds. Therefore, they *can't* take advantage of low noise scenarios. Our algorithms have a regret that scales with the range of the noise b , *without a prior knowledge of b* . Furthermore, our algorithms ultimately recover the new improved rate of the deterministic feedback suggested in the precedent case ($d = 0$).

Main results **Improved theoretical results and empirical performance:** We consider the optimization under an unknown *local* smoothness. We design two algorithms, **Sequ00L** for the deterministic case in Section 3, and **Stroqu00L** for the stochastic one in Section 4.

- **Sequ00L** is the first algorithm to obtain a loss $\tilde{\mathcal{O}}(e^{-n})$ under such minimal assumption, with deterministic feedback. The previously known **S00** (Munos, 2011) is only proved to achieve a loss of $\mathcal{O}(e^{-\sqrt{n}})$. Therefore, **Sequ00L** achieves, up to log factors, the result of

DOO that *knows the smoothness*. Note that Kawaguchi et al. (2016) designed a new version of SOO, called LOGO, that gives more flexibility in exploring more local scales but it was still only shown to achieve a loss of $\mathcal{O}(e^{-\sqrt{n}})$ despite the introduction of a new parameter. Achieving exponentially decreasing regret had previously only been achieved in setting with more assumptions (de Freitas et al., 2012; Malherbe and Vayatis, 2017; Kawaguchi et al., 2015). For example, de Freitas et al. (2012) achieves $\tilde{\mathcal{O}}(e^{-n})$ regret assuming several assumptions, for example that the function f is sampled from the Gaussian process with four times differentiable kernel along the diagonal. The consequence of our results is that to achieve $\tilde{\mathcal{O}}(e^{-n})$ rate, none of these strong assumptions is necessary.

- **StroquOOL** recovers, in the stochastic feedback, up to log factors, the results of POO, for the same assumption. However, as discussed later, **StroquOOL** is a simpler approach than POO with also an associated simpler analysis.
- **StroquOOL** adapts naturally to different noise range, i.e., the various values of b .
- **StroquOOL** obtains the *best of both worlds* in the sense that **StroquOOL** also obtains, up to log factors, the new optimal rates reached by **SequOOL** in the deterministic case. **StroquOOL** obtains this result without being aware a priori of the nature of the data, only for an additional log factor. Therefore, if we neglect the additional log factor, we can just have a single algorithm, **StroquOOL**, that performs well in both deterministic and stochastic case, without the knowledge of the smoothness in either one of them.
- In the numerical experiments, **StroquOOL** naturally adapts to lower noise. **SequOOL** obtains an exponential regret decay when $d = 0$ on common benchmark functions.

Algorithmic contributions and originality of the proofs Why does it work?

Both **SequOOL** and **StroquOOL** are simple and parameter-free algorithms. The analysis is also simple and self-contained and does not need to rely on results of other algorithms knowing the smoothness. We now explain the reason behind this combined simplicity and efficiency.

Both **SequOOL** and **StroquOOL** are based on a new core idea that the search for the optimum should progress strictly *sequentially* from an exploration of shallow depths (with large cells) to deeper depths (small and localized cells). This is different from the standard approach in SOO, StoSOO, and the numerous extensions that SOO has inspired (Busoniu et al., 2013; Wang et al., 2014; Al-Dujaili and Suresh, 2018; Qian and Yu, 2016; Kasim and Norreys, 2016; Derbel and Preux, 2015; Preux et al., 2014; Buşoniu and Morărescu, 2014; Kawaguchi et al., 2016). We have identified a bottleneck in SOO (Munos, 2011) and its extensions that open all depths *simultaneously* (their Lemma 2). However, in general, we show that the improved exploration of the shallow depths is beneficial for the deeper depths and therefore, we always complete the exploration of depth h before going to depth $h + 1$. As a result, we design a more *sequential* approach that simplifies our Lemma 2 to the point of being natural and straightforward.

This desired simplicity is also achieved by being the first to adequately leverage the reduced and natural set of assumptions introduced in the POO paper (Grill et al., 2015). This adequate and simple leverage should not conceal the fact that our local smoothness assumption is minimal and already way weaker than global Lipschitzness. Second, this

leveraging was absent in the analysis for P00 which additionally relies on the 40 pages proof of H00 (see [Shang et al., 2018](#), for a detailed discussion). Our proofs are succinct² while obtaining performance improvement ($d = 0$) and a new adaptation ($b = 0$). To obtain these, in an original way, our theorems are now based on solving a transcendental equation with the **Lambert W function**. For **Stroqu00L**, a careful discrimination of the parameters of the equation leads to optimal rates both in the deterministic and stochastic case.

Intriguingly, the amount of evaluations allocated to each depth h follows a **Zipf law** ([Powers, 1998](#)), that is, each depth level h is simply pulled inversely proportional to its depth index h . This is a simple but not a straightforward idea. It provides a parameter-free method to explore the depths without knowing the bound C on the number of optimal cells per depth ($\mathcal{N}_h = C \propto n/h$ when $d = 0$) and obtain a maximal optimal depth h^* of order n/C . A Zipf law has been used by [Audibert et al. \(2010\)](#) and [Abbasi-Yadkori et al. \(2018\)](#) in pure-exploration bandit problems but without any notion of depth in the search. In this paper, we introduce the Zipf law to tree-search algorithms.

Another novelty is that of *not using upper bounds* in **Stroqu00L** (unlike **StoS00**, **HCT**, **H00**, **P00**), which results in the contribution of *removing the need to know the noise amplitude*.

2. Partition, tree, assumption, and near-optimality dimension

Partitioning The hierarchical partitioning $\mathcal{P} = \{\mathcal{P}_{h,i}\}_{h,i}$ we consider is similar to the ones introduced in prior work ([Munos, 2011](#); [Valko et al., 2013](#)): For any depth $h \geq 0$ in the tree representation, the set $\{\mathcal{P}_{h,i}\}_{1 \leq i \leq I_h}$ of *cells* (or nodes) forms a partition of \mathcal{X} , where I_h is the number of cells at depth h . At depth 0, the root of the tree, there is a single cell $\mathcal{P}_{0,1} = \mathcal{X}$. A cell $\mathcal{P}_{h,i}$ of depth h is split into children subcells $\{\mathcal{P}_{h+1,j}\}_j$ of depth $h + 1$. As [Grill et al. \(2015\)](#), our work focuses on a notion of near-optimality dimension d that does not directly relate the smoothness property of f to a specific metric ℓ but directly to the hierarchical partitioning \mathcal{P} . Indeed, an interesting fundamental question is to determine a good characterization of the difficulty of the optimization for an algorithm that uses a given hierarchical partitioning of the space \mathcal{X} as its input (see [Grill et al., 2015](#), for a detailed discussion). Given a global maximum x^* of f , i_h^* denotes the index of the unique cell of depth h containing x^* , i.e., such that $x^* \in \mathcal{P}_{h,i_h^*}$. We follow the work by [Grill et al. \(2015\)](#) and state a single assumption on both the partitioning \mathcal{P} and the function f .

Assumption 1 *For any global optimum x^* , there exists $\nu > 0$ and $\rho \in (0, 1)$ such that $\forall h \geq 0, \forall x \in \mathcal{P}_{h,i_h^*}, f(x) \geq f(x^*) - \nu\rho^h$.*

Definition 1 *For any $\nu > 0$ and $\rho \in (0, 1)$, the **near-optimality dimension**³ $d(\nu, \rho)$ of f with respect to the partitioning \mathcal{P} and with associated constant C , is*

$$d(\nu, \rho) \triangleq \inf \left\{ d' \in \mathbb{R}^+ : \exists C > 0, \forall h \geq 0, \mathcal{N}_h(3\nu\rho^h) \leq C\rho^{-dh} \right\},$$

where $\mathcal{N}_h(\varepsilon)$ is the number of cells $\mathcal{P}_{h,i}$ of depth h such that $\sup_{x \in \mathcal{P}_{h,i}} f(x) \geq f(x^*) - \varepsilon$.

2. The proof is even redundantly written twice for **Stroqu00L** and **Sequ00L** for completeness

3. [Grill et al. \(2015\)](#) define $d(\nu, \rho)$ with the constant 2 instead of 3. 3 eases the exposition of our results.

Tree-based learner Tree-based exploration or tree search algorithm is a classical approach that has been widely applied to optimization as well as bandits or planning (Kocsis and Szepesvári, 2006; Coquelin and Munos, 2007; Hren and Munos, 2008), see Munos (2014) for a survey. At each round, the learner selects a cell $\mathcal{P}_{h,i}$ containing a predefined representative element $x_{h,i}$ and asks for its evaluation. We denote its value $f_{h,i} = f(x_{h,i})$. $T_{h,i}$ denotes the total number of evaluations allocated by the learner to the cell $\mathcal{P}_{h,i}$. Our learners collect the evaluations of f and organize them in a tree structure \mathcal{T} that is simply a subset of \mathcal{P} , $\mathcal{T} = \{\mathcal{P}_{h,i} \in \mathcal{P} : T_{h,i} > 0\}$, $\mathcal{T} \subset \mathcal{P}$. We define, specially for the noisy case, the estimated value of the cell $\widehat{f}_{h,i}$. Given the $T_{h,j}$ evaluations $y_1, \dots, y_{T_{h,j}}$, we have $\widehat{f}_{h,i} \triangleq \frac{1}{T_{h,j}} \sum_{s=1}^{T_{h,j}} y_s$, the empirical average of rewards obtained at this cell. We say that the learner *opens* a cell $\mathcal{P}_{h,i}$ with m evaluations if it asks for m evaluations from each of the children cells of cell $\mathcal{P}_{h,i}$. In the deterministic feedback, $m = 1$. For the sake of simplicity, the bounds reported in this paper are in terms of the total number of openings n , instead of evaluations. The number of function evaluations is upper bounded by Kn , where K is the maximum number of children cells of any cell in \mathcal{P} .

The Lambert W function Our results use the **Lambert W function**. Solving for the variable z , the equation $A = ze^z$ gives $z = W(A)$. W is multivalued for $z \leq 0$. However, in this paper, we consider $z \geq 0$ and $W(z) \geq 0$, referred to as the *standard W* . W cannot be expressed in terms of elementary functions. Yet, we have $W(z) = \log(z/\log z) + o(1)$ (Hoorfar and Hassani, 2008). W has applications in physics and applied mathematics (Corless et al., 1996).

Finally, let $[a : c] = \{a, a + 1, \dots, c\}$ with $a, c \in \mathbb{N}$, $a \leq c$, and $[a] = [1 : a]$. \log_d denotes the logarithm in base d , $d \in \mathbb{R}$. Without a subscript, \log is the natural logarithm in base e .

3. Adaptive deterministic optimization and improved rate

3.1 The Sequ00L algorithm

The Sequential Optimistic Optimization algorithm **Sequ00L** is described in Figure 1. **Sequ00L** explores sequentially the depth one by one, going deeper and deeper with a decreasing number of cells opened per depth h : $\lfloor h_{\max}/h \rfloor$ openings at depth h . h_{\max} is the maximal depth that is opened. The analysis of **Sequ00L** shows that it is relevant that $h_{\max} \triangleq \lfloor n/\overline{\log} n \rfloor$, where $\overline{\log} n$ is the n -th harmonic number, $\overline{\log} n \triangleq \sum_{t=1}^n \frac{1}{t}$, with $\overline{\log} n \leq \log n + 1$ for any positive integer n . **Sequ00L** returns the element of the evaluated cell with the highest value, $x(n) = \arg \max_{x_{h,i} : \mathcal{P}_{h,i} \in \mathcal{T}} f_{h,i}$. The budget is set to $n + 1$ to preserve the simplicity of the bounds. **Sequ00L** uses no more openings than that as

Parameters: $n, \mathcal{P} = \{\mathcal{P}_{h,i}\}$

Initialization: Open $\mathcal{P}_{0,1}$. $h_{\max} = \lfloor \frac{n}{\log(n)} \rfloor$.

For $h = 1$ to h_{\max}
 Open $\lfloor \frac{h_{\max}}{h} \rfloor$ cells $\mathcal{P}_{h,i}$ of depth h
 with largest values $f_{h,j}$.

Output $x(n) = \arg \max_{x_{h,i} : \mathcal{P}_{h,i} \in \mathcal{T}} f_{h,i}$.

Figure 1: The **Sequ00L** Algorithm

$$1 + \sum_{h=1}^{h_{\max}} \left\lfloor \frac{h_{\max}}{h} \right\rfloor \leq 1 + h_{\max} \sum_{h=1}^{h_{\max}} \frac{1}{h} = 1 + h_{\max} \overline{\log} h_{\max} \leq n + 1.$$

3.2 Analysis of Sequ00L

For any global optimum x^* , let \perp_h be the depth of the deepest opened node containing x^* at the end of the opening of depth h by Sequ00L (an iteration of the **for** cycle). Note that $\perp_{(\cdot)}$ is increasing. The proofs of the following statements are given in Appendix A.

Lemma 2 *For any global optimum x^* with associated (ν, ρ) as defined in Assumption 1, for any depth $h \in [h_{\max}]$, if $\frac{h_{\max}}{h} \geq C\rho^{-d(\nu, \rho)h}$, we have $\perp_h = h$, while $\perp_0 = 0$.*

Lemma 2 states that as long as Sequ00L opens more cells at depth h than the number of near-optimal cells at depth h , the cell containing x^* is opened at depth h .

Theorem 3 *Let W be the standard Lambert W function (see Section 2). For any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension $d = d(\nu, \rho)$, we have, after n rounds, the simple regret of Sequ00L bounded by*

$$\bullet \text{ If } d = 0, \quad r_n \leq \nu \rho^{\frac{1}{\max(C,1)} \lfloor \frac{n}{\log n} \rfloor}. \quad \bullet \text{ If } d > 0, \quad r_n \leq \nu \rho^{\frac{1}{d \log(1/\rho)} W\left(\frac{d \log(1/\rho)}{\max(C,1)} \lfloor \frac{n}{\log n} \rfloor\right)}.$$

For more readability, Corollary 4 uses a lower bound on W (Hoorfar and Hassani, 2008).

Corollary 4 *If $d > 0$, assumptions in Theorem 3 hold and $\lfloor n/\overline{\log n} \rfloor d \log \frac{1}{\rho} / \max(C, 1) > e$,*

$$r_n \leq \nu (\max(C, 1) / (d \log(1/\rho)))^{\frac{1}{d}} (\log(nd \log(1/\rho) / \max(C, 1)))^{\frac{1}{d}} \lfloor n/\overline{\log n} \rfloor^{-\frac{1}{d}}.$$

3.3 Discussion for the deterministic feedback

Comparison with S00 S00 and Sequ00L both address deterministic optimization without knowledge of the smoothness. The regret guarantees of Sequ00L are an improvement over S00. While when $d > 0$ both algorithms achieve a regret $\tilde{\mathcal{O}}(n^{-1/d})$, when $d = 0$, the regret of S00 is $\mathcal{O}(\rho^{\sqrt{n}})$ while the regret of Sequ00L is $\tilde{\mathcal{O}}(\rho^n)$ which is a significant improvement. As discussed in the introduction and by Valko et al. (2013, Section 5), the case $d = 0$ is very common. As pointed out by Munos (2011, Corollary 2), S00 has to actually know whether $d = 0$ or not to set the maximum depth of the tree as a parameter for S00. Sequ00L is fully adaptive, does not need to know any of this and actually gets a better rate.⁴ The conceptual difference with S00 is that Sequ00L is sequential, for a given depth h , Sequ00L first opens cells at depth h and then at depth $h + 1$ and so on, without coming back to lower depths. Indeed, an opening at depth $h + 1$ is based on the values observed while opening at depth h . Therefore, it is natural and less wasteful to do the opening in a sequential order. Moreover, Sequ00L is more conservative as it opens more the lower depths while S00 opens every depth equally. However from the *depth* perspective, Sequ00L is more aggressive as it opens depth as high as n , while S00 stops at \sqrt{n} .

Comparison with D00 Contrarily to Sequ00L, D00 knows the smoothness of the function. However this knowledge only improves the logarithmic factor in the current upper bound. When $d > 0$, D00 achieves a regret $\mathcal{O}(n^{-1/d})$, when $d = 0$, the loss is $\mathcal{O}(\rho^n)$.

4. A similar behavior is also achieved by combining two S00 algorithms, by running half of the samples for $d = 0$ and half for $d > 0$. However, Sequ00L does this naturally and gets a better rate when $d = 0$.

Lower bounds As discussed by Munos (2014) for $d = 0$, D00 matches the lower bound and it is even comparable to the lower-bound for concave functions. While S00 was not matching the bound of D00, with our result, we now know that, up to a log factor, it is possible to achieve the same performance as D00, *without the knowledge of the smoothness*.

4. Noisy optimization with adaptation to low noise

4.1 The Stroqu00L algorithm

In the presence of noise, it is natural to evaluate the cells multiple times, not just one time as in the deterministic case. The amount of times a cell should be evaluated to differentiate its value from the optimal value of the function depends on the gap between these two values as well as the range of noise. As we do not want to make *any* assumptions on knowing these quantities, our algorithm tries to be robust to any potential values by not making a fixed choice on the number of evaluations. Intuitively, Stroqu00L implicitly uses modified versions of Sequ00L, denoted Sequ00L(m),⁵ where each cell is evaluated m times, $m \geq 1$, while in Sequ00L $m = 1$. On one side, given one instance of Sequ00L(m), evaluating more each cells (m large) leads to a better quality of the mean estimates in each cell. On the other side, as a tradeoff, it implies that Sequ00L(m) is using more evaluations per depth and therefore is not be able to explore deep depths of the partition. The largest depth explored is now $\mathcal{O}(n/m)$. Stroqu00L then *implicitly* performs the same amount of evaluations as it would be performed by $\log n$ instances of Sequ00L(m) each with a number of evaluations of $m = 2^p$, where we have $p \in [0 : \log n]$.

The St(r)ochastic sequential Optimization aLgorithm Stroqu00L is described in Figure 2. Remember that ‘opening’ a cell means ‘evaluating’ its children. The algorithm opens cells by sequentially diving them deeper and deeper from the root node $h = 0$ to a maximal depth of h_{\max} . At depth h , we allocate, in a decreasing fashion, different number of evaluations 2^p to the cells with highest value of that depth, with p starting at $\lfloor \log_2(h_{\max}/h) \rfloor$ down to 0. The best cell that has been evaluated at least $\mathcal{O}(h_{\max}/h)$ times is opened with $\mathcal{O}(h_{\max}/h)$ evaluations, the two next best cells that have been evaluated at least $\mathcal{O}(h_{\max}/(2h))$ times are opened with $\mathcal{O}(h_{\max}/(2h))$ evaluations, the four next best cells that have been evaluated at least $\mathcal{O}(h_{\max}/(4h))$ times are opened with $\mathcal{O}(h_{\max}/(4h))$ evaluations and so on, until some $\mathcal{O}(h_{\max}/h)$ next best cells that have been evaluated at least once are opened with one evaluation. More precisely, given, p and h , we open, with 2^p evaluations, the $\lfloor h_{\max}/(h2^p) \rfloor$ non-previously-opened cells $\mathcal{P}_{h,i}$ with highest values $\hat{f}_{h,i}$ and given

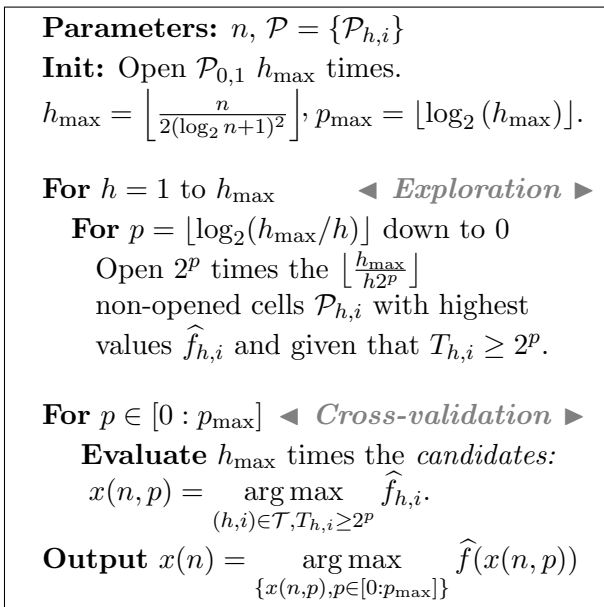


Figure 2: The Stroqu00L Algorithm

5. Again, this is only for the intuition, the algorithm is *not* a meta-algorithm over Sequ00L(m)’s.

that $T_{h,i} \geq 2^p$. The maximum number of evaluations of any cell is $p_{\max} \triangleq \lfloor \log_2(h_{\max}) \rfloor$. For each $p \in [0 : p_{\max}]$, the candidate output $x(n,p)$ is the cell with highest estimated value that has been evaluated at least 2^p times, $x(n,p) \triangleq \arg \max_{(h,i) \in \mathcal{T}, T_{h,i} \geq 2^p} \hat{f}_{h,i}$. We set $h_{\max} \triangleq \lfloor n/(2(\log_2 n + 1)^2) \rfloor$. In Appendix B, we prove that **Stroqu00L** uses less than $n + 1$ openings.

4.2 Analysis of Stroqu00L

The proofs of this section use a similar structure to the ones for the deterministic feedback. Additionally, they take into account the uncertainty created by the noise. The proofs of the following statements are given in Appendix D and E. For any x^* , $\perp_{h,p}$ is the depth of the deepest opened node with at least 2^p evaluations containing x^* at the end of the opening of depth h of **Stroqu00L**.

Lemma 5 *For any global optimum x^* with associated (ν, ρ) (see Assumption 1), with probability at least $1 - \delta$, for all depths $h \in [\lfloor \frac{h_{\max}}{2^p} \rfloor]$, for all $p \in [0 : \lfloor \log_2(h_{\max}/h) \rfloor]$, if $b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \leq \nu\rho^h$ and if $\frac{h_{\max}}{h2^p} \geq C\rho^{-d(\nu,\rho)h}$, we have $\perp_{h,p} = h$ while $\perp_{0,p} = 0$.*

Lemma 5 gives two conditions so that the cell containing x^* is opened at depth h . This holds if (1) **Stroqu00L** opens, with 2^p evaluations, more cells at depth h than the number of near-optimal cells at depth h ($\frac{h_{\max}}{h2^p} \geq C\rho^{-d(\nu,\rho)h}$) and (2) the 2^p evaluations are sufficient to discriminate the empirical average of near-optimal cells from the empirical average of sub-optimal cells ($b\sqrt{\log(4n/\delta)}/2^{p+1} \leq \nu\rho^h$).

To state the next theorems, we introduce \tilde{h} a positive real number satisfying $\frac{h_{\max}\nu^2\rho^{2\tilde{h}}}{hb^2\log(n^{3/2})} = C_1\rho^{-d\tilde{h}}$. We have $\tilde{h} = \frac{1}{(d+2)\log(1/\rho)} \log\left(\frac{\bar{n}}{\log\bar{n}}\right) + o(1)$ with $\bar{n} \triangleq \frac{\nu^2 h_{\max}(d+2)\log(1/\rho)}{C_1 b^2 \log(n^{3/2})}$. The quantity \tilde{h} gives the depth of deepest cell opened by **Stroqu00L** that contains x^* with high probability. Consequently, \tilde{h} also lets us characterize for which regime of the noise range b we recover results similar to the loss of the deterministic case. Discriminating on the noise regime, we now state our results, Theorem 6 for a high noise and Theorem 8 for a low one.

Theorem 6 High-noise regime *After n rounds, for any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension denoted for simplicity $d = d(\nu, \rho)$, if $b \geq \nu\rho^{\tilde{h}}/\sqrt{\log(n^{3/2}/b)}$, the simple regret of **Sequ00L** obeys*

$$r_n \leq \nu\rho^{\frac{1}{(d+2)\log(1/\rho)}} W\left(\left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor \frac{(d+2)\log(1/\rho)\nu^2}{C_1 b^2 \log(n^{3/2}/b)}\right) + 6b\sqrt{\log(n^{3/2}/b) / \left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor},$$

where W is the standard Lambert W function and $C_1 \triangleq \max(C, 1)$.

Corollary 7 *With the assumptions of Theorem 6 and $\bar{n} > e$,*

$$r_n \leq \nu \left(\frac{\log \bar{n}}{\bar{n}}\right)^{\frac{1}{d+2}} + b \sqrt{\frac{18 \log(n^{3/2}/b)}{2 \left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor}}, \quad \text{where } \bar{n} \triangleq \left\lfloor \frac{n/2}{(\log_2 n + 1)^2} \right\rfloor \frac{(d+2)\log(1/\rho)\nu^2}{C_1 b^2 \log(n^{3/2}/b)}.$$

Theorem 8 *Low-noise regime* After n rounds, for any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension denoted for simplicity $d = d(\nu, \rho)$, if $b \leq \nu \rho^{\tilde{h}} / \sqrt{\log(n^{3/2}/b)}$, the simple regret of **Stroqu00L** is bounded as follows

- If $d = 0$, $r_n \leq 7\nu\rho^{\frac{1}{C_1}} \left\lfloor \frac{n/2}{(\log_2(n)+1)^2} \right\rfloor$.
- If $d > 0$, $r_n \leq 7\nu\rho^{\frac{1}{d \log \frac{1}{\rho}}} W\left(\left\lfloor \frac{n/2}{(\log_2 n + 1)^2} \right\rfloor^{\frac{d \log \frac{1}{\rho}}{C_1}}\right)$.

Corollary 9 With the assumptions of Theorem 8, if $d > 0$, then

$$r_n \leq 7\nu \left(\frac{\log(\tilde{n})}{\tilde{n}} \right)^{\frac{1}{d}} \quad \text{with} \quad \tilde{n} \triangleq \left\lfloor \frac{n/2}{(\log_2 n + 1)^2} \right\rfloor \frac{d \log(1/\rho)}{C_1} \quad \text{and} \quad \tilde{n} > e.$$

4.3 Discussion for the stochastic feedback

Worst-case comparison with P00 and StoS00 When b is large and known: **Stroqu00L** is an algorithm designed for the noisy feedback while adapting to the smoothness of the function. Therefore, it can be directly compared to P00 and StoS00 that both tackle the same problem. The results for **Stroqu00L**, like the ones for P00, hold for $d \geq 0$, while the theoretical guarantees of StoS00 are only for the case $d = 0$. The general rate of **Stroqu00L** in Corollary 7⁶ is similar to the ones of P00 (for $d \geq 0$) and StoS00 (for $d = 0$) as their loss is $\tilde{\mathcal{O}}(n^{-\frac{1}{d+2}})$. More precisely, looking at the log factors, we can first notice an improvement over StoS00 when $d = 0$. We have $r_n^{\text{Stroqu00L}} = \mathcal{O}\left(\frac{\log^{3/2} n}{\sqrt{n}}\right) \leq r_n^{\text{StoS00}} = \mathcal{O}\left(\frac{\log^2 n}{\sqrt{n}}\right)$. Comparing with P00, we obtain a worse logarithmic factor, as $r_n^{\text{P00}} = \mathcal{O}\left(\left(\frac{\log^2 n}{n}\right)^{\frac{1}{d+2}}\right) \leq r_n^{\text{Stroqu00L}} = \mathcal{O}\left(\left(\frac{\log^3 n}{n}\right)^{\frac{1}{d+2}}\right)$. Despite having this (theoretically) slightly worse logarithmic factor compared to P00, **Stroqu00L** has two nice new features. First, our algorithm is conceptually simple, parameter-free, and does not need to call a sub-algorithm: P00 repetitively calls different instances of H00 which makes it a heavy meta-algorithm. Second, our algorithm, as we detail in next paragraphs, naturally adapts to low noise and, even more, recovers the rates of **Sequ00L** in the deterministic case, leading to exponentially decreasing loss when $d = 0$. We do not know if this deterioration of the logarithmic factor from P00 to **Stroqu00L** is the unavoidable price to pay to obtain an adaptation to the deterministic feedback case.

Comparison with oracle H00 H00 is also designed for the noisy optimization setting. However H00 *knows the smoothness* of f , i.e., (ν, ρ) are input parameters of H00. Using this extra knowledge H00 is only able to improve the logarithmic factor to achieve a regret of $r_n^{\text{H00}} = \mathcal{O}\left(\left(\frac{\log n}{n}\right)^{\frac{1}{d+2}}\right)$.

Adaptation to the range of the noise b without a prior knowledge A favorable feature of our bound in Corollary 7 is that it characterizes how the range of the noise b affects the rate of the regret for all $d \geq 0$. Considering the common case of $d = 0$, the regret in Corollary 7 scales linearly with the range of the noise b leading to potential large improvement for small b . Note that b is any real non-negative number and it is unknown by **Stroqu00L**. H00, P00, and StoS00, on the other hand, would only obtain a regret scaling with b when b is known to them as they directly encode a confidence bound that must include b , in the definition of their code. To achieve this result, and contrarily to H00, StoS00,

6. Note that the second term in our bound has at most the same rate as the first one.

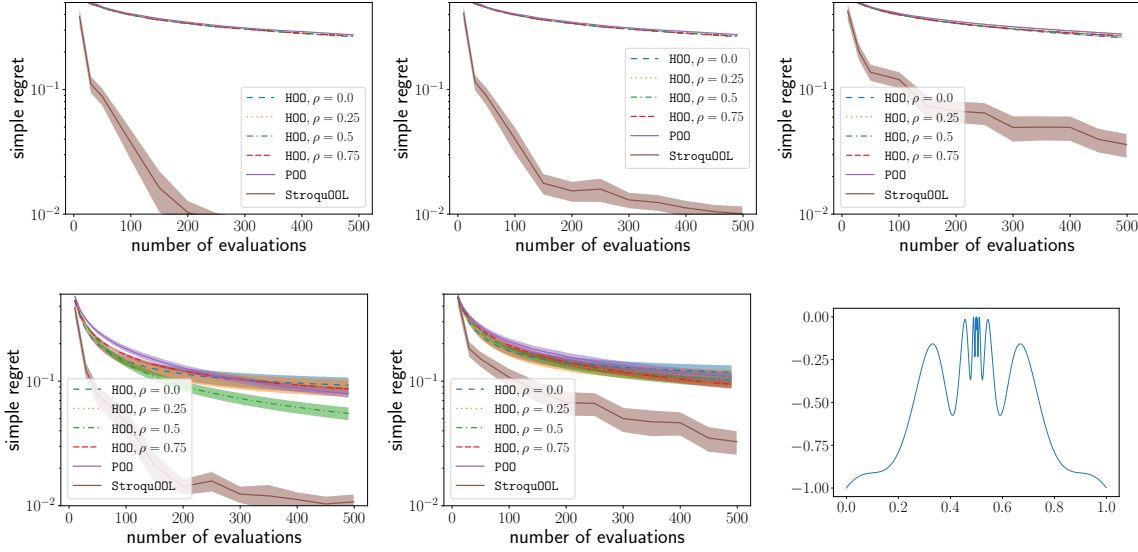


Figure 3: *Bottom right:* **Wrapped-sine** function ($d > 0$). The true range of the noise is b and the range used by HOO and P00 is \tilde{b} . *Top:* $b = 0, \tilde{b} = 1$ left — $b = 0.1, \tilde{b} = 1$ middle — $b = \tilde{b} = 1$ right. *Bottom:* $b = \tilde{b} = 0.1$ left — $b = 1, \tilde{b} = 0.1$ middle.

or P00, we designed **Stroqu00L** *without using upper-confidence bounds* (UCBs). Indeed, UCB approaches are overly conservative as they use hard-coded (and often overestimated) upper-bound on b . Finally, note that using UCB approaches with empirical estimation of the variance *would not* achieve the best of both worlds: a result that is discussed in the next paragraph. Indeed, an assumption on the noise b is still used in these approaches. This *prevents* having $\tilde{\mathcal{O}}(e^{-n})$ when $d = 0$ and $b \approx 0$.

Adaptation to the deterministic case and $d=0$ When the noise is very low, i.e., when $b \leq \nu \rho^{\tilde{h}} / \sqrt{\log(n^{3/2}/b)}$, which includes the deterministic feedback, in Theorem 8 and Corollary 9, **Stroqu00L** recovers the same rate as D00 and **Sequ00L** up to logarithmic factors. Remarkably, **Stroqu00L** obtains an exponentially decreasing regret when $d = 0$ while P00, **StoS00** or HOO only guarantee a regret of $\tilde{\mathcal{O}}(\sqrt{1/n})$ when unaware of the range b . Therefore, up to log factors, **Stroqu00L** achieves naturally the *best of both worlds* without being aware of the nature of the feedback (either stochastic or deterministic). Again, this is a behavior that one *cannot* expect from HOO, P00, and **StoS00** as they explicitly use confidence intervals in their algorithm assuming the range of noise is $b = 1$ which limits the maximum depth that can be explored.

5. Experiments

We empirically demonstrate how **Sequ00L** and **Stroqu00L** adapt to the complexity of the data and compare them to S00, P00, and HOO. We use two functions used by prior work as testbeds for optimization of difficult function without the knowledge of smoothness. The first one is the **wrapped-sine** function ($S(x)$, [Grill et al., 2015](#), Figure 3, bottom right) with

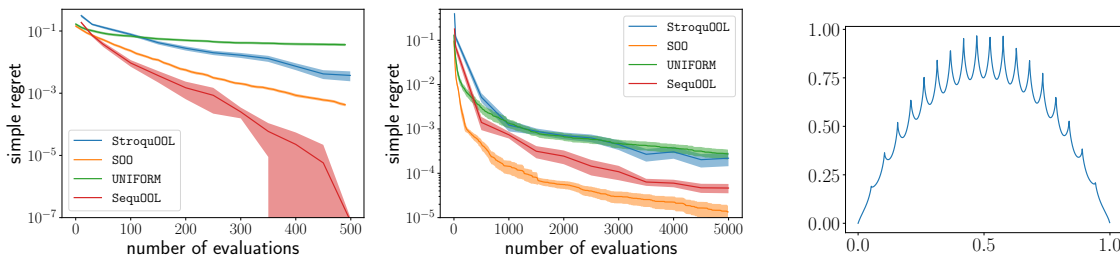


Figure 4: *Left* \mathcal{E} *center*: Deterministic feedback. *Right*: **Garland** function for which $d = 0$.

$S(x) \triangleq \frac{1}{2}(\sin(\pi \log_2(2|x - \frac{1}{2}|)) + 1)((2|x - \frac{1}{2}|)^{-\log .8} - (2|x - \frac{1}{2}|)^{-\log .3}) - (2|x - \frac{1}{2}|)^{-\log .8}$. This function has $d > 0$ for the standard partitioning (Grill et al., 2015). The second is the **garland** function ($G(x)$, Valko et al., 2013, Figure 4, bottom right) with $G(x) \triangleq 4x(1-x)(\frac{3}{4} + \frac{1}{4}(1 - \sqrt{|\sin(60x)|}))$. Function G has $d = 0$ for the standard partitioning (Valko et al., 2013). Both functions are in one dimension, $\mathcal{X} = \mathbb{R}$. We remark that our algorithms work in any dimension, but with the current computational power they would *not* scale beyond a thousand dimensions.

Stroqu00L outperforms P00 and H00 and adapts to lower noise. In Figure 3, we report the results of Stroqu00L, P00, and H00 for different values of ρ . As detailed in the caption, we vary the range of noise b and the range of noise \tilde{b} used by H00 and P00. In all our experiments, Stroqu00L outperforms P00 and H00. Stroqu00L adapts to low noise, its performance improves when b diminishes. To see that, compare top-left ($b = 0$), top-middle ($b = .1$), and top-right ($b = 1$) subfigures. On the other hand, P00 and H00 do not naturally adapt to the range of the noise: For a given parameter $\tilde{b} = 1$, the performance is unchanged when the range of the real noise varies as seen by comparing again top-left ($b = 0$), top-middle ($b = .1$), and top-right ($b = 1$). However, note that P00 and H00 *can* adapt to noise and perform empirically well if they have a good estimate of the range $b = \tilde{b}$ as in bottom-left, or if they underestimate the range of the noise, $\tilde{b} \ll b$, as in bottom-middle. In Appendix F, we report similar results on the garland function. Finally, Stroqu00L demonstrates its adaptation to both worlds in Figure 4 (left), where it achieves exponential decreasing loss in the case $d = 0$ and deterministic feedback.

Regrets of Sequ00L and Stroqu00L have exponential decay when $d = 0$. In Figure 4, we test in the deterministic feedback case with Sequ00L, Stroqu00L, S00 and the uniform strategy on the garland function (left) and the wrap-sine function (middle). Interestingly, for the garland function, where $d = 0$, Sequ00L outperforms S00 and displays a truly exponential regret decay (y-axis is in log scale). S00 appears to have the regret of $e^{-\sqrt{n}}$. Stroqu00L which is expected to have a regret $e^{-n/\log^2 n}$ lags behind S00. Indeed, $n/\log^2 n$ exceeds \sqrt{n} for $n > 10000$, for which the result is beyond the numerical precision. In Figure 4 (middle), we used the wrapped-sine. While all algorithms have similar theoretical guaranties since here $d > 0$, S00 outperforms the other algorithms.

Acknowledgements

We would like to thank Jean-Bastien Grill for sharing his code. We gratefully acknowledge the support of the NSF through grant IIS-1619362 and of the Australian Research Council through an Australian Laureate Fellowship (FL110100281) and through the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS). The research presented was also supported by European CHIST-ERA project DELTA, French Ministry of Higher Education and Research, Nord-Pas-de-Calais Regional Council, Inria and Otto-von-Guericke-Universität Magdeburg associated-team north-european project Allocate, and French National Research Agency projects ExTra-Learn (n.ANR-14-CE24-0010-01) and BoB (n.ANR-16-CE23-0003).

References

- Yasin Abbasi-Yadkori, Peter Bartlett, Victor Gabillon, Alan Malek, and Michal Valko. [Best of both worlds: Stochastic & adversarial best-arm identification](#). In *Conference on Learning Theory*, 2018.
- Abdullah Al-Dujaili and S. Suresh. [Multi-objective simultaneous optimistic optimization](#). *Information Sciences*, 424:159–174, 2018.
- Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. [Best arm identification in multi-armed bandits](#). In *Conference on Learning Theory*, pages 41–53, 2010.
- Mohammad Gheshlaghi Azar, Alessandro Lazaric, and Emma Brunskill. [Online stochastic optimization under correlated bandit feedback](#). In *International Conference on Machine Learning*, 2014.
- Sébastien Bubeck and Aleksandrs Slivkins. [The best of both worlds: stochastic and adversarial bandits](#). In *Conference on Learning Theory*, pages 42–1, 2012.
- Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. [X-armed bandits](#). *Journal of Machine Learning Research*, 12:1587–1627, 2011a.
- Sébastien Bubeck, Gilles Stoltz, and Jia Yuan Yu. [Lipschitz Bandits without the Lipschitz Constant](#). In *Algorithmic Learning Theory*, 2011b.
- Lucian Buşoniu and Irinel-Constantin Morărescu. [Consensus for black-box nonlinear agents using optimistic optimization](#). *Automatica*, 50(4):1201–1208, 2014.
- Lucian Busoniu, Alexander Daniels, Rémi Munos, and Robert Babuska. [Optimistic planning for continuous-action deterministic systems](#). In *Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2013 IEEE Symposium on*, pages 69–76. IEEE, 2013.
- Pierre-Arnaud Coquelin and Rémi Munos. [Bandit algorithms for tree search](#). In *Uncertainty in Artificial Intelligence*, 2007.
- Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. [On the lambert w function](#). *Advances in Computational mathematics*, 5(1):329–359, 1996.

- Nando de Freitas, Alex Smola, and Masrour Zoghi. [Exponential regret bounds for Gaussian process bandits with deterministic observations](#). In *International Conference on Machine Learning*, 2012.
- Steven De Rooij, Tim Van Erven, Peter D Grünwald, and Wouter M Koolen. [Follow the leader if you can, hedge if you must](#). *The Journal of Machine Learning Research*, 15(1): 1281–1316, 2014.
- Bilel Derbel and Philippe Preux. [Simultaneous optimistic optimization on the noiseless BBOB testbed](#). In *IEEE Congress on Evolutionary Computation, CEC 2015, Sendai, Japan, May 25-28, 2015*, pages 2010–2017, 2015.
- Jean-Bastien Grill, Michal Valko, and Rémi Munos. [Black-box optimization of noisy functions with unknown smoothness](#). In *Advances in Neural Information Processing Systems*, pages 667–675, 2015.
- Eldon Hansen and G William Walster. *Global optimization using interval analysis: revised and expanded*, volume 264. CRC Press, 2003.
- Abdolhossein Hoorfar and Mehdi Hassani. [Inequalities on the lambert w function and hyperpower function](#). *Journal of Inequalities in Pure and Applied Mathematics (JIPAM)*, 9(2):5–9, 2008.
- Jean-Francois Hren and Rémi Munos. [Optimistic Planning of Deterministic Systems](#). In *European Workshop on Reinforcement Learning*, 2008.
- David Jones, Cary Perttunen, and Bruce Stuckman. [Lipschitzian optimization without the Lipschitz constant](#). *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.
- Muhammad F Kasim and Peter A Norreys. [Infinite dimensional optimistic optimisation with applications on physical systems](#). *arXiv preprint arXiv:1611.05845*, 2016.
- Kenji Kawaguchi, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. [Bayesian optimization with exponential convergence](#). In *Advances in neural information processing systems*, pages 2809–2817, 2015.
- Kenji Kawaguchi, Yu Maruyama, and Xiaoyu Zheng. [Global continuous optimization with error bound and fast convergence](#). *Journal of Artificial Intelligence Research*, 56:153–195, 2016.
- R Baker Kearfott. *Rigorous global search: continuous problems*, volume 13. Springer Science & Business Media, 2013.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. [Multi-armed bandits in metric spaces](#). In *ACM Symposium on Theory of Computing (STOC)*, pages 681–690. ACM, 2008.
- Levente Kocsis and Csaba Szepesvári. [Bandit-based Monte-Carlo planning](#). In *European Conference on Machine Learning*, 2006.

- Andrea Locatelli and Alexandra Carpentier. [Adaptivity to Smoothness in X-armed bandits](#). In *Conference on Learning Theory*, 2018.
- Cédric Malherbe and Nicolas Vayatis. [Global optimization of lipschitz functions](#). In *Proceedings of the 34th International Conference on Machine Learning*, pages 2314–2323, 2017.
- Rémi Munos. [Optimistic optimization of a deterministic function without the knowledge of its smoothness](#). In *Advances in Neural Information Processing Systems*, pages 783–791, 2011.
- Rémi Munos. [From bandits to Monte-Carlo tree search: The optimistic principle applied to optimization and planning](#). *Foundations and Trends in Machine Learning*, 7(1):1–130, 2014.
- János D Pintér. *Global optimization in action: Continuous and lipschitz optimization: Algorithms, implementations and applications*. Nonconvex Optimization and Its Applications. Springer US, 2013.
- David Powers. [Applications and explanations of Zipf’s law](#). In *New methods in language processing and computational natural language learning*. Association for Computational Linguistics, 1998.
- Philippe Preux, Rémi Munos, and Michal Valko. [Bandits attack function optimization](#). In *Evolutionary Computation (CEC), 2014 IEEE Congress on*, pages 2245–2252. IEEE, 2014.
- Hong Qian and Yang Yu. [Scaling simultaneous optimistic optimization for high-dimensional non-convex functions with low effective dimensions](#). In *AAAI*, pages 2000–2006, 2016.
- Yevgeny Seldin and Aleksandrs Slivkins. [One practical algorithm for both stochastic and adversarial bandits](#). In *International Conference on Machine Learning*, pages 1287–1295, 2014.
- Xuedong Shang, Emilie Kaufmann, and Michal Valko. [Adaptive black-box optimization got easier: HCT needs only local smoothness](#). In *European Workshop on Reinforcement Learning*, 2018.
- Aleksandrs Slivkins. [Multi-armed bandits on implicit metric spaces](#). In *Neural Information Processing Systems*, 2011.
- Roman G Strongin and Yaroslav D Sergeyev. *Global optimization with non-convex constraints: Sequential and parallel algorithms*, volume 45. Springer Science & Business Media, 2013.
- Michal Valko, Alexandra Carpentier, and Rémi Munos. [Stochastic simultaneous optimistic optimization](#). In *International Conference on Machine Learning*, pages 19–27, 2013.
- Ziyu Wang, Babak Shakibi, Lin Jin, and Nando de Freitas. [Bayesian Multi-Scale Optimistic Optimization](#). In *International Conference on Artificial Intelligence and Statistics*, 2014.

Appendix A. Regret analysis of Sequ00L for deterministic feedback

Lemma 10 *For any global optimum x^* with associated (ν, ρ) as defined in Assumption 1, for any depth $h \in [h_{\max}]$, if $\frac{h_{\max}}{h} \geq C\rho^{-d(\nu, \rho)h}$, we have $\perp_h = h$, while $\perp_0 = 0$.*

Proof We prove Lemma 2 by induction over depth h .

1° For $h = 0$, we trivially have $\perp_h \geq 0$.

2° Now consider $h > 0$ and assume that $\frac{h_{\max}}{h} \geq C\rho^{-d(\nu, \rho)h}$. We want to show that $\perp_{h-1} = h - 1$. If $h = 1$ we already know $\perp_0 \geq 0$ and if $h > 1$, we have that for all $h' \in [h - 1]$,

$$\frac{h_{\max}}{h'} \geq \frac{h_{\max}}{h} \geq C\rho^{-d(\nu, \rho)h} \geq C\rho^{-d(\nu, \rho)h'},$$

which means, assuming that the proposition of the lemma is true for $h' = h - 1$ that $\perp_{h-1} = h - 1$. Therefore, at the end of the processing of depth $h - 1$, during which we were opening the cells of depth $h - 1$ we managed to open the cell $(h - 1, i_{h-1}^*)$ the optimal node of depth $h - 1$ (i.e., such that $x^* \in \mathcal{P}_{h-1, i_{h-1}^*}$). During phase h , the $\lfloor \frac{h_{\max}}{h} \rfloor$ cells from $\{\mathcal{P}_{h, i}\}_i$ with highest values $\{f_{h, i}\}_i$ are opened. For the purpose of contradiction, let us assume that \mathcal{P}_{h, i_h^*} is not one of them. This would mean that there exist at least $\lfloor \frac{h_{\max}}{h} \rfloor$ cells from $\{\mathcal{P}_{h, i}\}_i$, distinct from \mathcal{P}_{h, i_h^*} , satisfying $f_{h, i} \geq f_{h, i_h^*}$. As $f_{h, i_h^*} \geq f(x^*) - \nu\rho^h$ by Assumption 1, this means we have $\mathcal{N}_h(3\nu\rho^h) \geq \lfloor \frac{h_{\max}}{h} \rfloor + 1$ (the +1 is for \mathcal{P}_{h, i_h^*}). However by assumption of the lemma we have $\lfloor \frac{h_{\max}}{h} \rfloor \geq \lfloor C\rho^{-d(\nu, \rho)h} \rfloor$. It follows that $\mathcal{N}_h(3\nu\rho^h) > \lfloor C\rho^{-d(\nu, \rho)h} \rfloor$. This leads to a contradiction with f being of near-optimality dimension $d(\nu, \rho)$ with associated constant C as defined in Definition 1. Indeed the condition $\mathcal{N}_h(3\nu\rho^h) \leq C\rho^{-dh}$ in Definition 1 is equivalent to the condition $\mathcal{N}_h(3\nu\rho^h) \leq \lfloor C\rho^{-dh} \rfloor$ as $\mathcal{N}_h(3\nu\rho^h)$ is an integer. \blacksquare

Theorem 3 *Let W be the standard Lambert W function (see Section 2). For any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension $d = d(\nu, \rho)$, we have, after n rounds, the simple regret of Sequ00L bounded by*

$$\bullet \text{ If } d = 0, \quad r_n \leq \nu\rho^{\frac{1}{\max(C, 1)} \lfloor \frac{n}{\log n} \rfloor}. \quad \bullet \text{ If } d > 0, \quad r_n \leq \nu\rho^{\frac{1}{d \log(1/\rho)} W\left(\frac{d \log(1/\rho)}{\max(C, 1)} \lfloor \frac{n}{\log n} \rfloor\right)}.$$

Proof Let x^* be a global optimum with associated (ν, ρ) . For simplicity, let $d = d(\nu, \rho)$. We have

$$f(x(n)) \stackrel{\text{(a)}}{\geq} f_{\perp_{h_{\max}} + 1, i^*} \geq f(x^*) - \nu\rho^{\perp_{h_{\max}} + 1}.$$

where (a) is because $x(\perp_{h_{\max}} + 1, i^*) \in \mathcal{T}$ and $x(n) = \arg \max_{x \in \mathcal{X}} f(x)$. Note that the tree has depth $h_{\max} + 1$ in the end. From the previous inequality we have $r_n = \sup_{x \in \mathcal{X}} f(x) - f(x(n)) \leq \nu\rho^{\perp_{h_{\max}} + 1}$. For the rest of the proof, we want to lower bound $\perp_{h_{\max}}$. Lemma 2 provides a sufficient condition on h to get lower bounds. This condition is an inequality in which as h gets larger (more depth) the condition is more and more likely not to hold. For our bound on the regret of Stroqu00L to be small, we want a quantity h so that the inequality holds but having h as large as possible. So it makes sense to see when the inequality flips signs which is when it turns to equality. This is what we solve next. We solve

Equation 2 and then verify that it gives a valid indication of the behavior of our algorithm in term of its optimal h . We denote \bar{h} the positive real number satisfying

$$\frac{h_{\max}}{\bar{h}} = C_1 \rho^{-d\bar{h}} \quad (2)$$

where $C_1 = \max(C, 1) \geq 1$. As $\rho < 1$, $d \geq 0$ and $\bar{h} \geq 0$ we have $\rho^{-d\bar{h}} \geq 1$. This gives $C_1 \rho^{-d\bar{h}} \geq 1$. Finally as $\frac{h_{\max}}{\bar{h}} = C_1 \rho^{-d\bar{h}}$, we have $\bar{h} \leq h_{\max}$.

If $d = 0$ we have $\bar{h} = h_{\max}/C_1$. If $d > 0$ we have $\bar{h} = \frac{1}{d \log(1/\rho)} W(h_{\max} d \log(1/\rho)/C_1)$ where W is the standard Lambert W function. Using standard properties of the $\lfloor \cdot \rfloor$ function, we have

$$\frac{h_{\max}}{\lfloor \bar{h} \rfloor} \geq \frac{h_{\max}}{\bar{h}} = C_1 \rho^{-d\bar{h}} \geq C_1 \rho^{-d \lfloor \bar{h} \rfloor}. \quad (3)$$

We always have $\perp_{h_{\max}} \geq 0$. If $\bar{h} \geq 1$, as discussed above $\lfloor \bar{h} \rfloor \in [h_{\max}]$, therefore $\perp_{h_{\max}} \geq \perp_{\lfloor \bar{h} \rfloor}$, as \perp_{\cdot} is increasing. Moreover $\perp_{\bar{h}} = \bar{h}$ because of Lemma 2 which assumptions are verified because of Equation 3 and $\lfloor \bar{h} \rfloor \in [0 : h_{\max}]$. So in general we have $\perp_{h_{\max}} \geq \lfloor \bar{h} \rfloor$. If $d = 0$ we have, $r_n \leq \nu \rho^{\perp_{h_{\max}}+1} \leq \nu \rho^{\lfloor \bar{h} \rfloor+1} = \nu \rho^{\lfloor \frac{h_{\max}}{C_1} \rfloor+1} \leq \nu \rho^{\frac{1}{C_1}} = \nu \rho^{\frac{1}{\lfloor \frac{n}{\log n} \rfloor}}$. If $d > 0$ $r_n \leq \nu \rho^{\perp_{h_{\max}}+1} \leq \nu \rho^{\frac{1}{d \log(1/\rho)} W\left(\frac{h_{\max} d \log(1/\rho)}{C_1}\right)}$. $W(x)$ verifies for $x \geq e$, $W(x) \geq \log\left(\frac{x}{\log x}\right)$ (Hoorfar and Hassani, 2008). Therefore, if $h_{\max} d \log(1/\rho)/C_1 > e$ we have, denoting $d_\rho = d \log(1/\rho)$,

$$\frac{r_n}{\nu} \leq \rho^{\frac{1}{d_\rho} \left(\log\left(\frac{h_{\max} d_\rho / C_1}{\log(h_{\max} d_\rho / C_1)}\right) \right)} = e^{\frac{1}{d \log(1/\rho)} \left(\log\left(\frac{h_{\max} d_\rho / C_1}{\log\left(\frac{h_{\max} d_\rho}{C_1}\right)}\right) \right) \log(\rho)} = \left(\frac{h_{\max} d_\rho / C_1}{\log\left(\frac{h_{\max} d_\rho}{C_1}\right)} \right)^{-\frac{1}{d}}.$$

■

Appendix B. **Stroqu00L** is not using a budget larger than $n + 1$

Notice, for any given depth $h \in [1 : h_{\max}]$, **Stroqu00L** never uses more openings than $(p_{\max} + 1) \frac{h_{\max}}{h}$ as

$$\sum_{p=0}^{\lfloor \log_2(h_{\max}/h) \rfloor} \left\lfloor \frac{h_{\max}}{h 2^p} \right\rfloor 2^p \leq (\lfloor \log_2(h_{\max}/h) \rfloor + 1) \frac{h_{\max}}{h}$$

Summing over the depths, **Stroqu00L** never uses more openings than the budget $n + 1$ during its depth exploration as

$$\begin{aligned} 1 + (p_{\max} + 1) \sum_{h=1}^{h_{\max}} \left\lfloor \frac{h_{\max}}{h} \right\rfloor &\leq 1 + (p_{\max} + 1) h_{\max} \sum_{h=1}^{h_{\max}} \frac{1}{h} \\ &= 1 + h_{\max} \overline{\log}(h_{\max}) (p_{\max} + 1) \leq 1 + h_{\max} (p_{\max} + 1)^2 \leq n/2 + 1. \end{aligned}$$

We need to add the additional openings for the evaluation at the end,

$$\sum_{p=0}^{p_{\max}} \left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor \leq \frac{n}{2}.$$

Therefore, in total the budget is not more than $\frac{n}{2} + \frac{n}{2} + 1 = n + 1$. Again notice we use the budget of $n + 1$ only for the notational convenience, we could also use $n/4$ for the evaluation in the end to fit under n (it's important that the amount of openings is *linear* in n).

Appendix C. Lower bound on the probability of event ξ

In this section, we define and consider event ξ and prove it holds with high probability.

Lemma 11 *Let \mathcal{C} be the set of cells evaluated by **Stroqu00L** during one of its runs. \mathcal{C} is a random quantity. Let ξ be the event under which all average estimates in the cells receiving at least one evaluation from **Stroqu00L** are within their classical confidence interval, then $P(\xi) \geq 1 - \delta$, where*

$$\xi \triangleq \left\{ \forall \mathcal{P}_{h,i} \in \mathcal{C}, p \in [0 : p_{\max}] : \text{if } T_{h,j} = 2^p, \text{ then } \left| \widehat{f}_{h,j} - f_{h,j} \right| \leq b \sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \right\}.$$

Proof The idea of the proof of this lemma follows the similar line as the proof of the equivalent statement given for **StoS00** (Valko et al., 2013). The crucial point is that while we have potentially exponentially many combinations of cells that can be evaluated, given any particular execution we need to consider only a polynomial number of estimators for which we can use Chernoff-Hoeffding concentration inequality.

The identity of the set \mathcal{C} of the cells evaluated by **Stroqu00L**, $\mathcal{C} = \{\mathcal{P}_{h,i} : T_{h,j} > 0\}$, is random and can change at every run of **Stroqu00L**. However, no cells with a depth larger than $n+2$ are evaluated. Therefore, given n , the number of possible sets of cells \mathcal{C} associated with any run of **Stroqu00L** is finite. Let us denote the set of all such possible sets of cells as $\{\mathcal{C}\}_c$. Given any given set of cells \mathcal{C}_c , that **Stroqu00L** could open we denote S_c the event when **Stroqu00L** opens exactly all the cells in \mathcal{C}_c and define the related event ξ_c ,

$$\xi_c \triangleq \left\{ \forall \mathcal{P}_{h,i} \in \mathcal{C}_c, \forall p \in [0 : p_{\max}], \text{if } T_{h,j} = 2^p \text{ then } \left| \widehat{f}_{h,j} - f_{h,j} \right| \leq b \sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \mid S_c \right\}.$$

Given any \mathcal{C}_c , the number of cells in \mathcal{C}_c is bounded by n and the number of the possible values of p is bounded by $\log_2(n)$. For any \mathcal{C}_c , using a union bound on those $n \log_2(n) \leq 2n$ values, as well as using Chernoff-Hoeffding bound we have that given a fixed set of cells \mathcal{C}_c and associated event ξ_c , $P(\xi_c) \geq 1 - \delta$. Note that we can use a union bound on the cells in \mathcal{C}_c as this set is a set of fixed cells because we have conditioned on the event S_c which gives $\mathcal{C} = \mathcal{C}_c$. Also notice, that we could not do the union bound over all possible events $\{S\}_c$ because that would entangle too much failing probability. However, as $\{S\}_c$ is a set of pairwise disjoint events whose union is the entire sample space, using the law of total probability we have

$$P(\xi) = \sum_c P(\xi|S_c)P(S_c) = \sum_c P(\xi_c)P(S_c) \geq \sum_c (1 - \delta)P(S_c) = (1 - \delta) \sum_c P(S_c) = 1 - \delta.$$

■

Appendix D. Proof of Lemma 5

Lemma 12 *For any global optimum x^* with associated (ν, ρ) (see Assumption 1), with probability at least $1 - \delta$, for all depths $h \in [\lfloor \frac{h_{\max}}{2^p} \rfloor]$, for all $p \in [0 : \lfloor \log_2(h_{\max}/h) \rfloor]$, if $b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \leq \nu\rho^h$ and if $\frac{h_{\max}}{h2^p} \geq C\rho^{-d(\nu,\rho)h}$, we have $\perp_{h,p} = h$ while $\perp_{0,p} = 0$.*

Proof We place ourselves on event ξ defined in Lemma 11 and for which we proved that $P(\xi) \geq 1 - \delta$. We fix p . We prove the statement of the lemma, given that event ξ holds, by induction on the depth h .

1° For $h = 0$, we trivially have that $\perp_{h,p} \geq 0$.

2° Now consider $h > 0$ and assume that $\frac{h_{\max}}{h2^p} \geq C\rho^{-d(\nu,\rho)h}$ and $b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \leq \nu\rho^h$. We need to show that $\perp_{h,p} = h$. As a first step, we show that $\perp_{h-1,p} = h - 1$. If $h = 1$ we already know $\perp_{0,p} \geq 0$ and if $h > 1$, we have that for all $h' \in [h - 1]$,

$$\frac{h_{\max}}{h'2^p} \geq \frac{h_{\max}}{h2^p} \geq C\rho^{-d(\nu,\rho)h} \geq C\rho^{-d(\nu,\rho)h'} \quad \text{and} \quad b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \leq \nu\rho^h \leq \nu\rho^{h'},$$

which means, assuming that the proposition of the lemma is true for $h' = h - 1$ and p that $\perp_{h-1,p} = h - 1$. Therefore, at the end of the processing of depth $h - 1$, during which we were opening the cells of depth $h - 1$ we managed to open the cell $\mathcal{P}_{h-1, i_{h-1}^*}$ with at least 2^p evaluations. $\mathcal{P}_{h-1, i_{h-1}^*}$ is the optimal node of depth $h - 1$ (i.e., such that $x^* \in \mathcal{P}_{h-1, i^*}$). During phase h , the $\lfloor \frac{h_{\max}}{h2^p} \rfloor$ cells from $\{\mathcal{P}_{h,i}\}$ with highest values $\{f(x_{h,i})\}_{h,i}$ are opened. For the purpose of contradiction, let us assume that \mathcal{P}_{h, i_h^*} is not one of them. This would mean that there exist at least $\lfloor \frac{h_{\max}}{h2^p} \rfloor$ cells from $\{\mathcal{P}_{h,i}\}$, distinct from \mathcal{P}_{h, i_h^*} , satisfying $\hat{f}_{h,i} \geq \hat{f}_{h, i_h^*}$. This means that, for these cells we have $f_{h,i} + \nu\rho^h \stackrel{\text{(a)}}{\geq} f_{h,i} + b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \stackrel{\text{(b)}}{\geq} \hat{f}_{h,i} \geq \hat{f}_{h, i_h^*} \stackrel{\text{(b)}}{\geq} f_{h, i_h^*} - b\sqrt{\frac{\log(4n/\delta)}{2^{p+1}}} \stackrel{\text{(a)}}{\geq} f_{h, i_h^*} - \nu\rho^h$, where **(a)** is by assumption of the lemma, **(b)** is because ξ holds. As $f_{h, i_h^*} \geq f(x^*) - \nu\rho^h$ by Assumption 1, this means we have $\mathcal{N}_h(3\nu\rho^h) \geq \lfloor \frac{h_{\max}}{h2^p} \rfloor + 1$ (the +1 is for \mathcal{P}_{h, i_h^*}). However by assumption of the lemma we have $\lfloor \frac{h_{\max}}{h2^p} \rfloor \geq \lfloor C\rho^{-d(\nu,\rho)h} \rfloor$. It follows that $\mathcal{N}_h(3\nu\rho^h) > \lfloor C\rho^{-d(\nu,\rho)h} \rfloor$. This leads to having a contradiction with the function f being of near-optimality dimension $d(\nu, \rho)$ with associated constant C as defined in Definition 1. Indeed, the condition $\mathcal{N}_h(3\nu\rho^h) \leq C\rho^{-dh}$ in Definition 1 is equivalent to the condition $\mathcal{N}_h(3\nu\rho^h) \leq \lfloor C\rho^{-dh} \rfloor$ as $\mathcal{N}_h(3\nu\rho^h)$ is an integer. Reaching the contradiction proves the claim of the lemma. ■

Appendix E. Proof of Theorem 6 and Theorem 8

Theorem 6 *High-noise regime* After n rounds, for any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension denoted for simplicity $d =$

$d(\nu, \rho)$, if $b \geq \nu \rho^{\tilde{h}} / \sqrt{\log(n^{3/2}/b)}$, the simple regret of **Sequ00L** obeys

$$r_n \leq \nu \rho^{\frac{1}{(d+2)\log(1/\rho)}} W\left(\left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor \frac{(d+2)\log(1/\rho)\nu^2}{C_1 b^2 \log(n^{3/2}/b)}\right) + 6b \sqrt{\log(n^{3/2}/b)} \Big/ \left\lfloor \frac{n}{2(\log_2 n + 1)^2} \right\rfloor,$$

where W is the standard Lambert W function and $C_1 \triangleq \max(C, 1)$.

Theorem 8 *Low-noise regime* After n rounds, for any function f and one of its global optima x^* with associated (ν, ρ) , and near-optimality dimension denoted for simplicity $d = d(\nu, \rho)$, if $b \leq \nu \rho^{\tilde{h}} / \sqrt{\log(n^{3/2}/b)}$, the simple regret of **Stroqu00L** is bounded as follows

- If $d = 0$, $r_n \leq 7\nu \rho^{\frac{1}{C_1} \lfloor \frac{n/2}{(\log_2(n)+1)^2} \rfloor}$.
- If $d > 0$, $r_n \leq 7\nu \rho^{\frac{1}{d \log \frac{1}{\rho}} W\left(\left\lfloor \frac{n/2}{(\log_2 n + 1)^2} \right\rfloor \frac{d \log \frac{1}{\rho}}{C_1}\right)}$.

Proof [Proof of Theorem 6 and Theorem 8] We first place ourselves on the event ξ defined in Lemma 11 and where it is proven that $P(\xi) \geq 1 - \delta$. We bound the simple regret of **Stroqu00L** on ξ . We chose $\delta = \frac{4b}{\sqrt{n}}$ for the bound. We consider a global optimum x^* with associated (ν, ρ) . For simplicity we write $d = d(\nu, \rho)$. We have for all $p \in [0 : p_{\max}]$

$$\begin{aligned} f(x(n)) + b \sqrt{\frac{\log(n^{3/2}/b)}{2h_{\max}}} &\stackrel{\text{(a)}}{\geq} \hat{f}(x(n)) \stackrel{\text{(c)}}{\geq} \hat{f}(x(n, p)) \stackrel{\text{(b)}}{\geq} \hat{f}(x(\perp_{h_{\max}, p} + 1, i^*)) \\ &\stackrel{\text{(a)}}{\geq} f(x(\perp_{h_{\max}, p} + 1, i^*)) - b \sqrt{\frac{\log(n^{3/2}/b)}{2h_{\max}}} \stackrel{\text{(b)}}{\geq} f(x^*) - \nu \rho^{\perp_{h_{\max}, p} + 1} - b \sqrt{\frac{\log(n^{3/2}/b)}{2h_{\max}}} \end{aligned}$$

where **(a)** is because the $x(n, p)$ are evaluated h_{\max} times at the end of **Stroqu00L** and because ξ holds, **(b)** is because $x_{\perp_{h_{\max}, p} + 1, i^*} \in \{(h, i) \in \mathcal{T}, T_{h, i} \geq 2^p\}$ and $x(n, p) = \arg \max_{\mathcal{P}_{h, i} \in \mathcal{T}, T_{h, i} \geq 2^p} \hat{f}_{h, i}$ **(c)** is because $x(n) = \arg \max_{\{x(n, p), p \in [0 : p_{\max}]\}} \hat{f}(x(n, p))$.

From the previous inequality we have $r_n = f(x^*) - f(x(n)) \leq \nu \rho^{\perp_{h_{\max}, p} + 1} + 2b \sqrt{\frac{\log(n^{3/2}/b)}{2h_{\max}}}$, for $p \in [0 : p_{\max}]$.

For the rest of proof we want to lower bound $\max_{p \in [0 : p_{\max}]} \perp_{h_{\max}, p}$. Lemma 5 provides some sufficient conditions on p and h to get lower bounds. These conditions are inequalities in which as p gets smaller (fewer samples) or h gets larger (more depth) these conditions are more and more likely not to hold. For our bound on the regret of **Stroqu00L** to be small, we want quantities p and h where the inequalities hold but using as few samples as possible (small p) and having h as large as possible. Therefore we are interested in determining when the inequalities flip signs which is when they turn to equalities. This is what we solve next. We denote \tilde{h} and \tilde{p} the real numbers satisfying

$$\frac{h_{\max} \nu^2 \rho^{2\tilde{h}}}{\tilde{h} b^2 \log(n^{3/2})} = C_1 \rho^{-d\tilde{h}} \quad \text{and} \quad b \sqrt{\frac{\log(n^{3/2}/b)}{2\tilde{p}}} = \nu \rho^{\tilde{h}}. \quad (4)$$

Our approach is to solve Equation 4 and then verify that it gives a valid indication of the behavior of our algorithm in term of its optimal p and h . We have

$$\tilde{h} = \frac{1}{(d+2)\log(1/\rho)} W\left(\frac{\nu^2 h_{\max} (d+2)\log(1/\rho)}{C_1 b^2 \log(n^{3/2})}\right)$$

where standard W is the Lambert W function.

However after a close look at the Equation 4, we notice that it is possible to get values $\tilde{p} < 0$ which would lead to a number of evaluations $2^{\tilde{p}} < 1$. This actually corresponds to an interesting case when the noise has a small range and where we can expect to obtain an improved result, that is: obtain a regret rate close to the deterministic case. This low range of noise case then has to be considered separately.

Therefore, we distinguish two cases which corresponds to different noise regimes depending on the value of b . Looking at the equation on the right of (4), we have that $\tilde{p} < 0$ if $\frac{\nu^2 \rho^{2\tilde{h}}}{b^2 \log(n^{3/2}/b)} > 1$. Based on this condition we now consider the two cases. However for both of them we define some generic \tilde{h} and \tilde{p} .

High-noise regime $\frac{\nu^2 \rho^{2\tilde{h}}}{b^2 \log(n^{3/2}/b)} \leq 1$: In this case, we denote $\tilde{h} = \tilde{h}$ and $\tilde{p} = \tilde{p}$. As $\frac{1}{2^{\tilde{p}}} = \frac{\nu^2 \rho^{2\tilde{h}}}{b^2 \log(n^{3/2}/b)} \leq 1$ by construction, we have $\tilde{p} \geq 0$. Using standard properties of the $\lfloor \cdot \rfloor$ function, we have

$$b\sqrt{\frac{\log(n^{3/2}/b)}{2^{\lfloor \tilde{p} \rfloor + 1}}} \leq b\sqrt{\frac{\log(n^{3/2}/b)}{2^{\tilde{p}}}} = \nu \rho^{\tilde{h}} \leq \nu \rho^{\lfloor \tilde{h} \rfloor} \quad (5)$$

$$\text{and, } \frac{h_{\max}}{\lfloor \tilde{h} \rfloor^{2^{\lfloor \tilde{p} \rfloor}}} \geq \frac{h_{\max}}{\lfloor \tilde{h} \rfloor^{2^{\tilde{p}}}} = \frac{h_{\max} \nu^2 \rho^{2\tilde{h}}}{\lfloor \tilde{h} \rfloor \log(n^{3/2}/b)} \geq \frac{h_{\max} \nu^2 \rho^{2\tilde{h}}}{\tilde{h} \log(n^{3/2}/b)} \geq C_1 \rho^{-d\tilde{h}} \geq C_1 \rho^{-d\lfloor \tilde{h} \rfloor}. \quad (6)$$

Low-noise regime $\frac{\nu^2 \rho^{2\tilde{h}}}{b^2 \log(n^{3/2}/b)} > 1$: In this case, we can reuse arguments close to the argument used in the deterministic feedback case in the proof of [Sequ00L](#) (Theorem 3), we denote $\tilde{h} = \bar{h}$ and $\tilde{p} = \bar{p}$ where \bar{h} and \bar{p} verify,

$$\frac{h_{\max}}{\bar{h}} = C_1 \rho^{-d\bar{h}} \quad \text{and} \quad \bar{p} = 0. \quad (7)$$

If $d = 0$ we have $\bar{h} = h_{\max}/C_1$. If $d > 0$ we have $\bar{h} = \frac{1}{d \log(1/\rho)} W\left(\frac{h_{\max} d \log(1/\rho)}{C_1}\right)$ where standard W is the standard Lambert W function. Using standard properties of the $\lfloor \cdot \rfloor$ function, we have

$$b\sqrt{\frac{\log(n^{3/2}/b)}{2^{\lfloor \tilde{p} \rfloor + 1}}} \leq b\sqrt{\log(n^{3/2}/b)} < \nu \rho^{\tilde{h}} \stackrel{\text{(a)}}{\leq} \nu \rho^{\bar{h}} \leq \nu \rho^{\lfloor \bar{h} \rfloor} \quad (8)$$

where (a) is because of the following reasoning. As we have $\frac{h_{\max} \nu^2 \rho^{2\tilde{h}}}{hb^2 \log(n^{3/2}/b)} = C_1 \rho^{-d\tilde{h}}$ and $\frac{\nu^2 \rho^{2\tilde{h}}}{b^2 \log(n^{3/2}/b)} > 1$, then, $\frac{h_{\max}}{h} < C_1 \rho^{-d\tilde{h}}$. From the inequality $\frac{h_{\max}}{h} < C_1 \rho^{-d\tilde{h}}$ and the fact that \bar{h} corresponds to the case of equality $\frac{h_{\max}}{h} = C_1 \rho^{-d\bar{h}}$, we deduce that $\bar{h} \leq \tilde{h}$, since the left term of the inequality decreases with h while the right term increases. Having $\bar{h} \leq \tilde{h}$ gives $\rho^{\bar{h}} \geq \rho^{\tilde{h}}$.

Given these particular definitions of \ddot{h} and \ddot{p} in two distinct cases we now bound the regret.

As $\rho < 1$, $d \geq 0$ and $\ddot{h} \geq 0$ we have $\rho^{-d\ddot{h}} \geq 1$. This gives $C_1 \rho^{-d\ddot{h}} \geq 1$. Finally as $\frac{h_{\max}}{h 2^{\ddot{p}}} = C_1 \rho^{-d\ddot{h}}$, we have $\ddot{h} \leq h_{\max}/2^{\ddot{p}}$. Note also that from the previous equation we have that if $\ddot{h} \geq 1$, $\ddot{p} \leq \log_2(h_{\max})$. Finally in both regimes we already proved that $\ddot{p} \geq 0$.

We always have $\perp_{\lfloor h_{\max}/2^{\ddot{p}} \rfloor, \lfloor \ddot{p} \rfloor} \geq 0$. If $\ddot{h} \geq 1$, as discussed above $\lfloor \ddot{h} \rfloor \in \llbracket \lfloor h_{\max}/2^{\ddot{p}} \rfloor \rrbracket$, therefore $\perp_{\lfloor h_{\max}/2^{\ddot{p}} \rfloor, \lfloor \ddot{p} \rfloor} \geq \perp_{\lfloor \ddot{h} \rfloor, \lfloor \ddot{p} \rfloor}$, as $\perp_{\cdot, \lfloor p \rfloor}$ is increasing for all $p \in [0, p_{\max}]$. Moreover on event ξ , $\perp_{\lfloor \ddot{h} \rfloor, \lfloor \ddot{p} \rfloor} = \lfloor \ddot{h} \rfloor$ because of Lemma 5 which assumptions on $\lfloor \ddot{h} \rfloor$ and $\lfloor \ddot{p} \rfloor$ are verified because of Equations 5 and 6 in the high-noise regime and because of Equations 7 and 8 in the low-noise regime, and, in general, $\lfloor \ddot{h} \rfloor \in \llbracket \lfloor h_{\max}/2^{\ddot{p}} \rfloor \rrbracket$ and $\lfloor \ddot{p} \rfloor \in [0 : p_{\max}]$. So in general we have $\perp_{\lfloor h_{\max}/2^{\ddot{p}} \rfloor, \lfloor \ddot{p} \rfloor} \geq \lfloor \ddot{h} \rfloor$.

We bound the regret now discriminating on whether or not the event ξ holds. We have

$$\begin{aligned} r_n &\leq (1 - \delta) \left(\nu \rho^{\perp_{h_{\max}, \ddot{p}} + 1} + 2b \sqrt{\frac{\log(n^{3/2})}{2h_{\max}}} \right) + \delta \times 1 \leq \nu \rho^{\perp_{h_{\max}, \ddot{p}} + 1} + 2b \sqrt{\frac{\log(n^{3/2})}{2h_{\max}}} + \frac{4b}{\sqrt{n}} \\ &\leq \nu \rho^{\perp_{h_{\max}, \ddot{p}} + 1} + 6b \sqrt{\frac{\log(n^{3/2})}{h_{\max}}}. \end{aligned}$$

We can now bound the regret in the two regimes.

High-noise regime In general, we have

$$r_n \leq \nu \rho^{\frac{1}{(d+2)\log(1/\rho)}} W\left(\frac{\nu^2 h_{\max}(d+2)\log(1/\rho)}{C_1 \log(n^{3/2})}\right) + 6b \sqrt{\frac{\log(n^{3/2})}{h_{\max}}}.$$

While in the deterministic feedback case, the regret was scaling with d when $d \geq 0$, in the stochastic feedback case, the regret scale with $d + 2$. This is because the uncertainty due to the presence of noise diminishes as $n^{-\frac{1}{2}}$ when we collect n observations.

Moreover, as proved by [Hoorfar and Hassani \(2008\)](#), the Lambert $W(x)$ function verifies for $x \geq e$, $W(x) \geq \log\left(\frac{x}{\log x}\right)$. Therefore, if $\frac{\nu^2 h_{\max}(d+2)\log(1/\rho)}{C_1 \log(n)} > e$ we have, denoting $d' = (d + 2)\log(1/\rho)$,

$$\begin{aligned} r_n - 6b \sqrt{\frac{\log(n^{3/2})}{h_{\max}}} &\leq \nu \rho^{\frac{1}{d'}} \left(\log\left(\frac{\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2}/b)}}{\log\left(\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2}/b)}\right)}\right) \right) \\ &= \nu e^{\frac{1}{(d+2)\log(1/\rho)} \left(\log\left(\frac{\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2}/b)}}{\log\left(\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2}/b)}\right)}\right) \right)} \log(\rho) = \nu \left(\frac{\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2})}}{\log\left(\frac{h_{\max} d' \nu^2}{C_1 \log(n^{3/2})}\right)} \right)^{-\frac{1}{d+2}}. \end{aligned}$$

Low-noise regime We have $6b \sqrt{\frac{\log(n^{3/2})}{h_{\max}}} \leq 6 \frac{\nu \rho^{\tilde{h}}}{\sqrt{\log(n^{3/2}/b)}} \sqrt{\frac{\log(n^{3/2}/b)}{h_{\max}}} \leq 6\nu \rho^{\tilde{h}} \leq 6\nu \rho^{\bar{h}}$.

Therefore $r_n \leq \nu \rho^{\perp_{h_{\max}, \bar{p}} + 1} + 6b \sqrt{\frac{\log(n^{3/2})}{h_{\max}}} \leq 7\nu \rho^{\bar{h}}$. Discriminating between $d = 0$ and $d > 0$

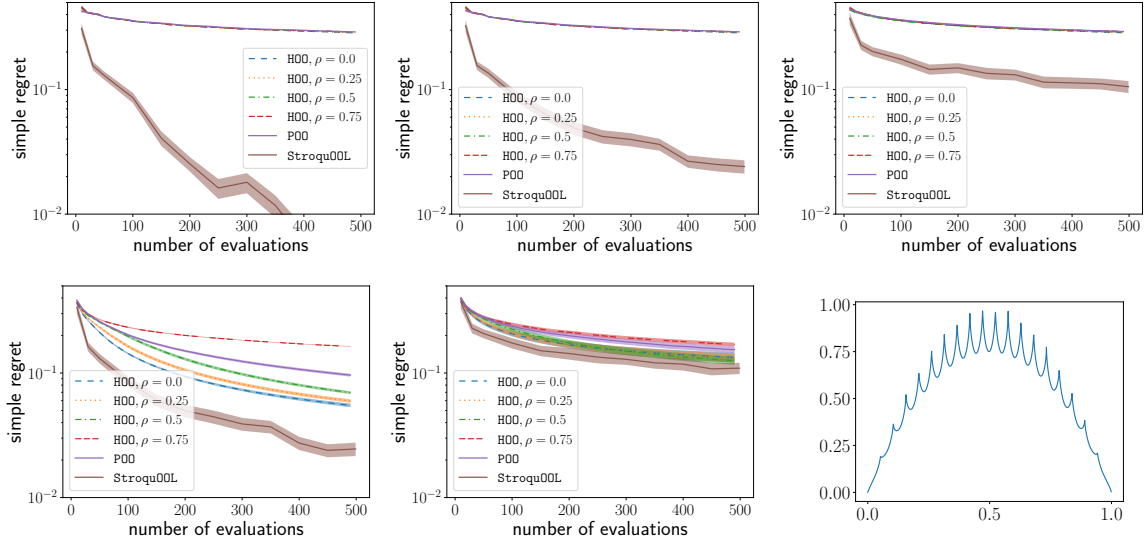


Figure 5: **Garland** function: The true range of the noise is b and the range of noise used by HOO and P00 is \tilde{b} and they are set as top: $b = 0, \tilde{b} = 1$ left — $b = 0.1, \tilde{b} = 1$ middle — $b = 1, \tilde{b} = 1$ right bottom: $b = 0.1, \tilde{b} = 0.1$ left — $b = 1, \tilde{b} = 0.1$ middle

leads to the claimed results. ■

Appendix F. More experiments with the garland function

We report more experiments in Figure 5. As mentioned in the main paper, the observations are very similar to those in Figure 3.

Appendix G. Use of the budget

Remark 13 *The algorithm can be made anytime and unaware of n using the classic ‘doubling trick’.*

Remark 14 (More efficient use of the budget) *Because of the use of the floor functions $\lfloor \cdot \rfloor$, the budget used in practice, $1 + \sum_{h=1}^{h_{\max}} \lfloor \frac{h_{\max}}{h} \rfloor$, can be significantly smaller than n . While this only affects numerical constants in the bounds, in practice, it can influence the performance noticeably. Therefore one should consider having h_{\max} replaced by $c \times h_{\max}$ with $c \in \mathbb{R}$ and $c = \max\{c' \in \mathbb{R} : 1 + \sum_{h=1}^{h_{\max}} \lfloor \frac{h_{\max}}{h} \rfloor \leq n\}$. Additionally, the use the budget n could be slightly optimized by taking into account that the necessary number of pulls at depth h is actually $\min(\lfloor h_{\max}/h \rfloor, K^h)$.*