



Adaptive multi-fidelity optimization with fast learning rates

Côme Fiegel, Victor Gabillon, Michal Valko

► To cite this version:

Côme Fiegel, Victor Gabillon, Michal Valko. Adaptive multi-fidelity optimization with fast learning rates. International Conference on Artificial Intelligence and Statistics, 2020, Palermo, Italy. hal-03288879

HAL Id: hal-03288879

<https://inria.hal.science/hal-03288879>

Submitted on 16 Jul 2021

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.

Adaptive multi-fidelity optimization with fast learning rates

Côme Fiegel

École Normale Supérieure, Paris
Inria Lille

Victor Gabillon

Huawei R&D, UK

Michal Valko

Inria Lille

Abstract

In *multi-fidelity optimization*, we have access to *biased* approximations of varying costs of the target function. In this work, we study the setting of optimizing a locally smooth function with a limited budget Λ , where the learner has to make a trade-off between the cost and the bias of these approximations. We first prove lower bounds for the *simple regret* under different assumptions on the fidelities, based on a *cost-to-bias* function. We then present the *Kometo* algorithm which achieves, with additional logarithmic factors, the same rates *without* any knowledge of the function smoothness and fidelity assumptions and improving prior results. Finally, we empirically show that our algorithm outperforms prior multi-fidelity optimization methods without the knowledge of problem-dependent parameters.

1 Introduction

In multi-fidelity optimization (Cutler et al., 2014; Huang et al., 2006; Kandasamy et al., 2016c, 2017), the learner *actively* optimizes a function but only observes, at each of the rounds, biased values of that function. The learner can *pay* to reduce the bias of the observed function values. The smaller the bias the higher the cost, urging the learner to carefully allocate its total cost budget Λ on the fly.

We consider the case of *derivative-free optimization* where no gradient information is available (Matyas, 1965). This is of great interest for the multiple applications in which it is either difficult to access, compute, or even define gradients (Nesterov and Spokoiny, 2017).

Using only zero-order information, derivative-free optimization addresses optimising over functions that are not differentiable, non-continuous, or non-smooth. Moreover, there are known methods that work without knowing the smoothness parameters (ν, ρ) of the function (Auer et al., 2007; Kleinberg et al., 2008; Grill et al., 2015; Valko et al., 2013; Bartlett et al., 2019).

Derivative-free multi-fidelity optimization is useful in particular the hyper-parameters tuning of *complex* machine learning models, where each evaluation of the model is costly such as tokamak simulators. However, the mapping between the hyper-parameter and performance of the learned model can be highly non-convex and non-smooth. Moreover, training a model, given the hyper-parameters can be expensive and time-consuming (Sen et al., 2018). In a situation, where computation or time are constrained by a budget, these constraints prevent us from carefully evaluating the qualities of all the models generated from a continuous set of hyper-parameters. Then, given one fixed set of hyper-parameters, the bias of the estimation of the quality of fully-trained model is a (decreasing) function ζ of the amount of computation resource spent training the model. Ultimately, we would expect this bias to be zero if the model is trained until convergence. However, the bias function ζ is a function that depends on the type of trained models and that is in applications a priori unknown.

The most related approach for the considered setting is the MFPD00 algorithm of Sen et al. (2018). In order to provide theoretical guaranties for MFPD00, ζ is either assumed to be known or some parametric assumptions on ζ are made and the parameters are estimated online. However, knowing ζ or its parametric family is unrealistic.

In this paper, we propose a new method called *Kometo* that adapts to the unknown ζ and the unknown smoothness parameters (ν, ρ) . Our analysis is more general than the analysis of Sen et al. (2018) and provides a broader and finer set of behaviors of the cost-to-bias function. This allows us to provide a characterisation of the complexity of the problem by providing the first

regret lower bounds in multi-fidelity optimization. We also show that *Kometo* obtains rates that match the ones of our lower bounds and improves upon the rates of MFPD00 *while dropping the assumptions of knowing the bias function ζ in advance*.

Related work Among the large work on derivative-free optimization, we focus on algorithms that perform well under *minimal* assumptions as well as minimal knowledge of the function. Under *weak/local* smoothness around one global maximum (Auer et al., 2007; Kleinberg et al., 2008; Bubeck et al., 2011), some algorithms require the knowledge of the local smoothness such as H00 (Bubeck et al., 2011), Zooming (Kleinberg et al., 2008), or D00 (Munos, 2011). Among the work relying on an *unknown* local smoothness, Sequ00L (Bartlett et al., 2019) improves on S00 (Munos, 2011; Kawaguchi et al., 2016) and 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) and GP0 (Shang et al., 2019) provide more general results. Finally, Stroqu00L (Bartlett et al., 2019) matches, up to log factors, the guarantees of Sequ00L and GP0 for deterministic and stochastic feedback respectively, *without requiring the knowledge of the range of the noise b* .

Multi-fidelity optimization is a well studied setting. Here, we address *online* multi-fidelity optimization. Many of approaches rely on Bayesian models, e.g., Gaussian processes. Zhang et al. (2019) relies on entropic search to find the maximum, while Kandasamy et al. (2016a) adapts GP-UCB (Srinivas et al., 2010) to multi-fidelity setting. Most of these methods need an access to a bias function, while Ghosh et al. (2019) use the cost of the approximations to estimate its values. Li et al. (2017) obtains good empirical results by trying a lot of configurations at low fidelities and progressively eliminating the less interesting ones while using higher and higher fidelities. Two prior works adapted algorithms working under local smoothness around one global maximum to multi-fidelity settings. First, Sen et al. (2018) adapted P00 (Grill et al., 2015) to deterministic multi-fidelity settings and later Sen et al. (2019) made it work under stochastic ones.

Main contributions

- We give more general assumptions on the fidelity approximations based on their cost while keeping the smoothness assumption on the target function.
- We prove lower bounds of the simple regret under these more general and different assumptions.
- We provide *Kometo*, an algorithm that, in deter-

ministic settings, without *any* knowledge on the bias function and the smoothness of the target function, achieves minimax optimal rates for simple regret up to logarithmic factors on all considered assumptions on the fidelities. It improves the previously proven guaranteed rates under local smoothness assumptions of Sen et al. (2018), except in the case $\alpha = 1$ of Assumption 2(a),¹ where it has additional logarithmic factors. Our *Kometo* comes with important properties:

- It does *not* assume an access to the target function, only an access to increasingly better approximations, unlike previous algorithms as the ones of Sen et al. (2018).
- It only uses the comparisons of evaluations at the same fidelity level, and not directly the values of the evaluations, which leads to *weaker fidelity assumptions and better empirical results*.
- It works in stochastic settings by changing the number of evaluations at higher fidelities.
- We provide synthetic experiments and a hyperparameter tuning experiment to demonstrate the efficiency of *Kometo*.

2 Problem setting

In this section, we introduce a generalization of the settings presented by Sen et al. (2018).

We want to optimize a target function $f : \mathcal{X} \rightarrow \mathbb{R}$ under a budget $\Lambda \in \mathbb{R}_+$. The evaluation of this target function is done through its fidelity approximations. We thus denote by $Z = [0, 1]$ the fidelity space and by $(f_z)_{z \in Z}$ the fidelity approximations. In particular, $z = 0$ corresponds to the lowest fidelity, while $z = 1$ corresponds to the highest one. We also denote by $\zeta : Z \rightarrow \overline{\mathbb{R}_+}$ the unknown bias function, such that there exists a family $(g_z)_{z \in Z}$ of real-valued strictly increasing function with $\|f - g_z \circ f_z\|_\infty \leq \zeta(z)$ for $z \in Z$; motivations for this assumption are explained below. A known cost function $\lambda : Z \rightarrow \overline{\mathbb{R}_+}$ indicates the budget used at each evaluation for a given fidelity. We also assume that the algorithm can request, for any $c \geq 1$, a fidelity z_c such that $\lambda(z_c) \leq c$, and we define $\Phi : [1, +\infty[\rightarrow \overline{\mathbb{R}_+}$ with $\Phi(c) = \zeta(z_c)$, the cost-to-bias function, which gives for each cost c the minimal bias that one can be guaranteed for an observation of f . Assumptions 2 below are made on this function.

At round t , the algorithm makes an evaluation of the function of a point $x_t \in \mathcal{X}$ and at a fidelity $z_t \in Z$ (or at a cost c_t , see above), as long as $\sum_{s=1}^t \lambda(z_s) \leq \Lambda$.

¹hyperbolic decreasing of the cost-to-bias function

The algorithm observes at round t the value $f_{z_t}(x_t)$ in return. The algorithm must finally output a value x_Λ . We then define the simple regret of a policy π for r_Λ as

$$r_\Lambda(\pi) \triangleq \mathbb{E} \left[\max_{x \in \mathcal{X}} f(x) - f(x_\Lambda) \right],$$

where the expectation is taken over the randomness of the algorithm. In the rest of this paper, we only aim to minimize this regret, without any constraint on time or space complexity.

Problem setting remarks One of the main aspects of our approach is that we *do not assume to have an access to the bias function* ζ . This highlights the fact that our algorithm is fully adaptive, and only needs the cost of each fidelity as an input. Since the bias function is usually unknown in practice, prior works rely on various techniques (e.g., MLE) to guess the values of this function for implementations, but often assume it has a specific form. Given the known results it is therefore surprising we get faster rates, and we do it *without relying on any information on the bias function*.

Moreover we relax the original assumption of Sen et al. (2018) that $\|f - f_z\|_\infty \leq \zeta(z)$ for $z \in Z$ and use instead $\|f - g_z \circ f_z\|_\infty \leq \zeta(z)$ for $z \in Z$. This lets the f_z approximations be potentially arbitrarily biased with respect to f as long as the ordering in f is approximately kept. Indeed, as g_z are increasing functions, we have that $g_z \circ f_z(x_1) \geq g_z \circ f_z(x_2)$ iff $f_z(x_1) \geq f_z(x_2)$ for any $x_1, x_2 \in \mathcal{X}$. This more general model for example fits in cases where evaluating at lower fidelities (with higher bias) has a great impact on individual feedback, but a low impact on how each different points compare to each other at the same fidelity level. This is for example the case on neural network training, where evaluating with less iterations (lower fidelity) may increase the overall error for every set of hyper-parameters at similar rates. Note that theoretical results will simultaneously hold under both assumption as long as two conditions are met. First, the behavior of our algorithms is not based directly on the (estimated) value of the function f_z but only on comparisons of these estimates of f_z . Second the estimates that are compared are computed from evaluations coming from the same fidelity. This is the case of our algorithm *Kometo*. Indeed *Kometo* similarly as SOO (as noted by Munos, 2014) or SequOOL, a *rank-based algorithm*. This means that its behavior is based on the rank of the function evaluations, and not directly on their values. On the contrary the behavior of MFPDOO relies directly on the values in practice when estimating the constant of the parametric model, and would therefore not extend to our general assumption.

Another particularity is that we do not assume that the cost function λ is bounded. We assume quite

generally that $\lambda : Z \rightarrow \overline{\mathbb{R}}_+$ instead of restricting ourself to having $\lambda : Z \rightarrow [0, 1]$ as in Sen et al. (2018). In our scenario, it can happen that some approximations of the function f with low bias are simply too costly for our limited budget. Working under this larger assumption fits better problems in which we can only access feedback from imperfect simulators while the real phenomenon can not be directly evaluated in practice. In such scenario, the MFPDOO Sen et al. (2018) is not usable as it assumes that it directly evaluate the target function f with finite cost during its final cross-validation phase. Our results can also be extended to cases where fidelity space is discrete, by using a piecewise constant cost function.

Finally instead of minimizing the simple regret, the cumulative regret has been also studied in multi-fidelity setting (Kandasamy et al., 2016b), rewarding all accurate evaluations of the target function. However in the present paper we optimize the simple regret as our initial objective is to find the optimum of the target function. The simple regret is adapted to the objectives of *Kometo*, e.g., hyper-parameter optimization, where we wish to spend the entire budget on *pure exploration*.

3 Assumptions

Our algorithm needs two assumptions, one on the target function (which describes its smoothness) and one on the fidelity approximations (which characterizes how well they approximate the function).

Hierarchical partitioning We use the notion of hierarchical partitioning (Munos, 2011). At every depth $h \geq 0$, \mathcal{X} (potentially multi-dimensional) is partitioned into K^h different cells $(\mathcal{P}_{h,i})_{0 \leq i \leq K^h-1}$. All the cells $(\mathcal{P}_{h,i})_{h,i}$ form a tree, where the root is $\mathcal{P}_{0,0} = \mathcal{X}$, and where each cell $\mathcal{P}_{h,i}$ has K children, $(\mathcal{P}_{h+1,Ki+l})_{0 \leq l \leq K-1}$, which form a partition of their parent cell.

We make an assumption on the target function f and the hierarchical partitioning \mathcal{P} , identical to the settings of Sen et al. (2018). This following assumption is way weaker than global Lipschitzness and as explained by Grill et al. (2015) is simpler and weaker than assumptions made in previous works (Auer et al., 2007; Munos, 2011).

Assumption 1. [Assumption on the target function] For one of the global optimum x^* of f , there exists $\nu > 0$ and $\rho \in]0, 1[$ such that $\forall h \in \mathbb{N}, \forall x \in \mathcal{P}_{h,i_h^*}, f(x) \geq f(x^*) - \nu \rho^h$, where \mathcal{P}_{h,i_h^*} is the cell of depth h containing x^*

We now define a notion of near-optimality dimension that only depends on the hierarchical partitioning \mathcal{P} ,

and not on a metric.

Definition 1. Near-optimality dimension: For any $\nu > 0$ and $\rho \in]0, 1[$, we say that $d \in \mathbb{R}_+$ is a near optimality dimension of f with respect to the partitioning \mathcal{P} and the smoothness parameters (ν, ρ) if

$$\exists C > 1, \forall h \in \mathbb{N}, N_h(3\nu\rho^h) \leq C\rho^{-dh},$$

where $N_h(\varepsilon)$ is the number of cells $\mathcal{P}_{h,i}$ such that

$$\sup_{x \in \mathcal{P}_{h,i}} f(x) \geq f(x^*) - \varepsilon.$$

We then define

$$S(\mathcal{P}, \nu, \rho, d, C) \triangleq \{f : X \rightarrow \mathbb{R} \mid f \text{ has smoothness parameters } (\nu, \rho) \text{ for } \mathcal{P} \text{ and } d \text{ is a near-optimality dimension with associated constant } C\}$$

Note on the near-optimality dimension definition: Grill et al. (2015) define the near-optimality dimension as the infimum of the set of d that satisfies this definition (with ν and ρ fixed). However, they then assume that this infimum also satisfies this definition, which is not necessarily true (the set can be of the form $\mathbb{R}_{>0}$ for example). Bartlett et al. (2019) solve this issue by adding an extra dependence on the constant C to get a closed set (fixed parameters are then ν, ρ and C , instead of just ν and ρ). To avoid this extra dependence, we chose to define d as a near-optimality dimension, rather than the near optimality dimension.

We can notice that a function with smoothness parameters (ν, ρ) has necessarily an associated constant

$$C \geq C_{\min} \triangleq \left(\frac{K}{\rho^{-d}} \right)^{\left\lceil \frac{\log 3}{\log \frac{1}{\rho}} \right\rceil}$$

since all the cells at depth $h_0 \triangleq \lfloor (\log 3)/(\log 1/\rho) \rfloor$ are near-optimal because of Assumption 1. Indeed, it guarantees that $\forall x \in X, f(x) \geq f(x^*) - \nu$, which implies that $\forall x \in X, f(x) \geq f(x^*) - 3\nu\rho^{h_0}$

We also have that $d_{\max} \triangleq (\log K)/(\log 1/\rho)$ is always a near-optimality dimension of the function, because of the bound K^h on the number of cell of depth h . This emphasizes the fact that the near-optimality dimension of a function is a way to characterize the complexity of optimizing the function, and not an assumption. The case $d = 0$ allows for faster rates the best empirical results. As explained by Munos (2014), the case $d = 0$ is the most relevant in practice and covers most of the real-world setups.

We now state three new different assumptions on the rate at which the cost-to-bias function Φ is decreasing, namely polynomially, exponentially, or by a constant.

Assumption 2 (Assumption on the fidelities).

- (a) There exist $A, \alpha > 0$ such that $\Phi(c) \leq A/c^\alpha$.
- (b) There exist $B, \sigma, \beta > 0$ such that $\Phi(c) \leq Be^{-\frac{c^\beta}{\sigma}}$.
- (c) There exists $a \geq 1$ such that $\Phi(c) = 0$ for all $c \geq a$.

Definition 2. For Asm being one of the three Assumptions 2 (either Assumption 2(a), Assumption 2(b), Assumption 2(c), with its specific parameters depending on the case), we define $F(\text{Asm}, f, \lambda) = \{(f_z)_{z \in Z} \mid \text{there exists a function } \zeta \text{ such that assumption Asm holds on } f \text{ and } (f_z)_{z \in Z}, \text{ with } \lambda \text{ as a cost function and } \zeta \text{ as a bias function}\}$.

The above assumptions describe realistic rates for the cost-to-bias function. Assumption 2(a) generalizes Assumption 3 of Sen et al. (2018) which is equivalent to the case $\alpha \geq 1$. Assumption 2(b) generalizes Assumption 2 of Sen et al. (2018) which corresponds to the case $\beta = 1$. Assumption 2(c) is relevant when a minimal cost to get a perfectly accurate estimation is needed but unknown. It is also useful to link our results (especially the theorem below) to works using single-fidelity optimization, since the settings are then equivalent to deterministic single-fidelity settings.

4 Lower bound

We provide the first lower bounds for the assumptions on the fidelities considered. Theorem 1 gives, for assumptions 2(a), 2(b) and 2(c), bounds on the achievable theoretical performance of an algorithm working under these assumptions.

Theorem 1 (Lower bounds on simple regret). *Let \mathcal{P} a partitioning of a space \mathcal{X} , (ν, ρ) some smoothness parameters, $d \in [0, d_{\max}]$ a near-optimality dimension with associated constant $C \geq C_{\min}$ and Asm one of the three Assumptions 2 with associated parameters. Then, for any budget Λ large enough, for any (deterministic or random) policy π , there exist a target function $f \in S(\mathcal{P}, \nu, \rho, d, C)$, a cost function λ , and fidelity approximations $(f_z)_{z \in Z} \in F(\text{Asm}, f, \lambda)$ such that: Under Assumption 2(a) ($\Phi(c) \leq A/c^\alpha$):*

$$r_\Lambda(\pi) \geq D_1 \Lambda^{\frac{-1}{d+\frac{1}{\alpha}}}$$

Under Assumption 2(b) ($\Phi(c) \leq Be^{-\frac{c^\beta}{\sigma}}$):

$$r_\Lambda(\pi) \geq \begin{cases} e^{-D_2 \Lambda^{\frac{\beta}{1+\beta}}}, & \text{when } d = 0 \\ D_3 \Lambda^{\frac{-1}{d}}, & \text{when } d > 0 \end{cases}$$

Under Assumption 2(c) ($\Phi(c) = 0$ for all $c \geq a$):

$$r_\Lambda(\pi) \geq \begin{cases} e^{-D_4 \Lambda}, & \text{when } d = 0 \\ D_5 \Lambda^{\frac{-1}{d}}, & \text{when } d > 0, \end{cases}$$

where $D_1, D_2, D_3, D_4, D_5 > 0$ are constants that do not depend on Λ and π .

Ideas of the proof The proof is in the appendix. It is based on the construction of a target function and its approximations, such that the algorithm π may not reach a certain depth h and open a near-optimal cell at depth h . The construction of the target function is done thanks to a tree, whose leaves are cells of the partitioning, and which reflects which cells are near-optimal for the target function. The approximations are made such that we can lower bound the cost that π has to invest to get precise enough information.

We thus have to construct this tree, which is the tricky part of the proof. This implies choosing near-optimal cells that π is unlikely to open. We then get that this depth h , which depends on the parameters of the problem and on the budget, may not be reached by π with a certain fixed probability. We can use this to lower bound the regret.

Link with the upper bounds In Section 5 we give an algorithm that, without any knowledge on ν, ρ, d, C , and Φ , achieves these rates with additional constants and logarithmic factors. This means that these lower bounds are close to the optimal rate for policies working with these assumptions, both with and without knowledge of these parameters.

The only previous work using hierarchical partitioning optimization with multi-fidelity model and deterministic feedback worked with narrower assumptions as said above. It obtained, under Assumption 2(a) with $\alpha \geq 1$, a regret of $\mathcal{O}((\frac{\Lambda}{\log \Lambda})^{\frac{1}{\alpha+1}})$, which is only optimal (ignoring constant and log factors) when $\alpha=1$. Under Assumption 2(b), MPFDOO gets, assuming $\beta = 1$ a regret of $\mathcal{O}((\frac{\Lambda}{\log \Lambda})^{\frac{1}{\alpha+\varepsilon}})$, for any $\varepsilon > 0$, provided the budget is large enough (with the threshold having a dependence on ε), which does not show that this lower bound was reached.

Assumption 2(c) let us extend our results to single-fidelity algorithm with deterministic feedback. A true exponential decay for $d = 0$ (and thus optimal up to a constant) was first achieved by DOO (Munos, 2011, but required the knowledge of the smoothness. SequOOL (Bartlett et al., 2019) then managed to achieve an exponential decay without the knowledge of the smoothness, but with a logarithmic factor in the exponent. We however realized it is possible to get a true exponential decay without the knowledge of the smoothness parameters by changing the number of opened cells at each depth h of SequOOL, to either $\lfloor 2\sqrt{n/h} \rfloor$ up to depth n , or $\lfloor n/(h \log(n/h)^2) \rfloor$ up to depth $\lfloor n/e^2 \rfloor$.

5 Algorithm

In this section we propose a new algorithm for multi-fidelity optimization called **Kometo**. We start with some helpful notation.

Cell evaluations: Cell evaluations are done through a single representant of each cell $\mathcal{P}_{h,i}$, denoted $x_{h,i}$. $T_{h,i,j}$ denote the number of evaluation potentially done for the cell $\mathcal{P}_{h,i}$ at fidelity level j .

For **Kometo**, the fidelity level j , with j a non-negative integer, is defined as z_{e^j} . At each fidelity level, at most one evaluation can be done for each cell, which means that $T_{h,i,j}$ is equal to either 0 or 1. We hence denote as $f_{h,i,j}$ the result of the potential evaluation, when $T_{h,i,j} = 1$. We can notice that, for any j , because of how the cells are opened, $\{\mathcal{P}_{h,i}, T_{h,i,j}=1\}$ is always a tree.

We also slightly modify the usual definition of a cell opening to make it work with our multi-fidelity settings.

Multi-Fidelity Cell Opening: Opening a cell at fidelity level j means that, for each of its children $\mathcal{P}_{h+1,i}$, the $T_{h+1,i,u}$ for $0 \leq u \leq j$ are set to 1.

This means that the values $f_{h+1,i,u}$, equal to $f_{z_{e^u}}(x_{h+1,i})$, with $x_{h,i}$ the representative element of the cell, can be requested and hence the evaluations can be performed. With this definition, the opening at fidelity level j of a cell can not induce a total cost of more than $\frac{K e^{j+1}}{e-1}$.

Kometo explanations: **Kometo** is detailed in Algorithm 1. The algorithm presented is inspired by StroquOOL (Bartlett et al. (2019)). Its main feature is that, using Zipf sampling (which means, opening up to $\tilde{\Lambda}$ cells at $h=1$, up to $\tilde{\Lambda}/2$ cells at $h=2$ and so on) it manages to reach the optimal rate up to logarithmic factor without the knowledge of the smoothness. This is done, in the exploration part, by opening a decreasing number of cell at each depth, and at a given depth, gradually decreasing the fidelity at which cells are opened. The intuition behind this idea is that, for each depth h , and each $0 \leq j_h \leq j_{\max}$, the number of cell opened at fidelity level j_h or higher will decrease with j_h . If this j_h is too low, the precision might also be too low for the choices to be relevant, but if j_h is too high, not enough cells will be opened. Cross-validation is then used by the algorithm in order to choose the best cell regardless of depth and fidelity level. It ensures that the choice of a particular j_h is not needed.

Algorithm 1 KOMETO

1: **Parameters:** $(f_z)_{z \in Z}$, \mathcal{P} , Λ , λ

2: **Init:**
 $\tilde{\Lambda} \leftarrow \frac{(e-1)\Lambda}{2Ke(\log \Lambda + 1)^2}$, $j_{\max} \leftarrow \lfloor \log \tilde{\Lambda} \rfloor$.
 Open with budget $\tilde{\Lambda}$ the cell $P_{0,0}$.

For $h = 1$ to $\lfloor \tilde{\Lambda} \rfloor$ **◀ Exploration ▶**
 For $m = 1$ to $\lfloor \tilde{\Lambda}/h \rfloor$
 $j \leftarrow \lfloor \log \frac{\tilde{\Lambda}}{hm} \rfloor$
 Open at fidelity level j the non-opened
 cell $P_{h,i}$ with the highest value $f_{h,i,j}$,
 given that $T_{h,i,j} = 1$

For $j = 0$ to j_{\max} **◀ Cross-validation ▶**
 Evaluate at cost $\tilde{\Lambda}$ the candidates

$x_j^c \leftarrow \arg \max_{(h,i) \in \mathcal{T}, T_{h,i,j}=1} f_{h,i,j}$.

Output $x_{\Lambda} \leftarrow \arg \max_{\{j \in [0:j_{\max}]\}} f_{z_{\tilde{\Lambda}}}(x_j^c)$

Budget optimization: With a given budget Λ , we can actually initialize the $\tilde{\Lambda}$ constant with a way higher value than $\frac{(e-1)\Lambda}{2Ke(\log(\Lambda+1))^2}$, for multiple reasons:

- The actual cost used for a cell opening is rounded down to $e^{\lfloor \log c \rfloor}$.
- The total budget mentioned for a cell opening assumes that all the evaluation at different fidelities will be requested for the children, which is not the case.
- The number of opened cells at each depth is bounded by K^h
- For some partitioning, it is possible to use the parent evaluations for one of its children.

Since the budget used can be predicted using only $\tilde{\Lambda}$ and the partitioning, and increase with $\tilde{\Lambda}$, it is possible to quickly calculate the optimal initial value of $\tilde{\Lambda}$ using dichotomy. However, these previous optimization can only increase $\tilde{\Lambda}$ by a multiplicative constant. Even if the budget needs to be set in advance for this algorithm, since we optimize the simple regret, we can obtain anytime guarantees which only differ by a multiplicative constant using the doubling trick.

6 Theoretical guarantees

We first state a simple proposition which asserts, with the initial value of $\tilde{\Lambda}$, the budget condition is respected.

Proposition 2 (Budget use). *The budget used by Kometo does not exceed Λ .*

Our upper bounds use the Lambert function, evaluated at positive real values. This function is defined as the inverse of the function $f(z) = ze^z$. With the first two terms of its asymptotic expansion, we get, when z goes to infinity, that $W(x) = \log x - \log \log x + o(1)$.

We now state the main results of our analysis, using the same context as Theorem 1 on lower bounds. The proof is given in appendix:

Theorem 3 (Upper bounds on the simple regret). *Let \mathcal{P} be a partitioning of a space \mathcal{X} , (ν, ρ) some smoothness parameters, $d \in [0, d_{\max}]$ a near-optimality dimension with associated constant $C \geq C_{\min}$ and Asm being one of the three Assumption 2 with its associated parameters.*

Then, for any budget $\Lambda \geq 1$, target function $f \in S(\mathcal{P}, \nu, \rho, d, C)$, cost function λ , and fidelity approximations $(f_z)_{z \in Z} \in F(\text{Asm}, f, \lambda)$ provided to Kometo,

Under Assumption 2(a) ($\Phi(c) \leq \frac{A}{c^\alpha}$): We first define two values, then state the regret

Value of h_1
$\frac{1}{(d + \frac{1}{\alpha}) \log \frac{1}{\rho}} W \left(\frac{\tilde{\Lambda} \nu^{\frac{1}{\alpha}} (d + \frac{1}{\alpha}) \log \frac{1}{\rho}}{4CeA^{\frac{1}{\alpha}}} \right)$

Value of h_2	
When $d = 0$	$\frac{\tilde{\Lambda}}{4C}$
When $d > 0$	$\frac{1}{d \log \frac{1}{\rho}} W \left(\frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4C} \right)$

Regret	
High budget ($\nu \rho^{h_1} \leq e^\alpha A$)	$r_{\Lambda} \leq \frac{3\nu}{\rho} \rho^{h_1} + 2 \frac{A}{\Lambda^\alpha}$
Low budget ($\nu \rho^{h_1} > e^\alpha A$)	$r_{\Lambda} \leq \frac{3\nu}{\rho} \rho^{h_2} + 2 \frac{A}{\Lambda^\alpha}$

Under Assumption 2(b) ($\Phi(c) \leq Be^{\frac{-c^\beta}{\sigma}}$): We also define $a_{b,\nu} = \max(\frac{1}{2\sigma}, \log(\frac{B}{\nu}))$

Value of h_1	
When $d = 0$	$\left(\frac{\tilde{\Lambda}}{4Che}\right)^{\frac{\beta}{\beta+1}} \left(\frac{1}{2\sigma \log \frac{1}{\rho}}\right)^{\frac{1}{\beta+1}}$
When $d > 0$	$\frac{\beta+1}{\beta d \log \frac{1}{\rho}} W\left(\frac{\beta}{\beta+1} d \log \frac{1}{\rho} \left(\frac{\tilde{\Lambda}}{4Che}\right)^{\frac{\beta}{\beta+1}}\right) \left(\frac{1}{2\sigma \log \frac{1}{\rho}}\right)^{\frac{1}{\beta+1}}$

Value of h_2	
When $d = 0$	$\frac{\tilde{\Lambda}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}}$
When $d > 0$	$\frac{1}{d \log \frac{1}{\rho}} W\left(\frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}}\right)$

Regret	
High budget ($h_1 \geq \frac{a_{b,\nu}}{\log \frac{1}{\rho}}$)	$r_{\Lambda} \leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{\frac{-\tilde{\Lambda}\beta}{\sigma}}$
Low budget ($h_1 < \frac{a_{b,\nu}}{\log \frac{1}{\rho}}$)	$r_{\Lambda} \leq \frac{3\nu}{\rho} \rho^{h_2} + 2Be^{\frac{-\tilde{\Lambda}\beta}{\sigma}}$

Under Assumption 2(c) ($\Phi(c) = 0$ for all $c \geq a$):

Value of h		Regret
When $d = 0$	$\frac{\tilde{\Lambda}}{4Ca\epsilon}$	
When $d > 0$	$\frac{1}{d \log \frac{1}{\rho}} W\left(\frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4Ca\epsilon}\right)$	$r_{\Lambda} \leq \frac{\nu}{\rho} \rho^h$

Corollary 4 (Regret decreasing rates). *Following Theorem 3 (the exact upper bounds used for the rates are given in appendix):*

Assumption 2(a)	High budget	Low budget
When $d = 0$	$\tilde{\mathcal{O}}(\Lambda^{-\alpha})$	
When $d > 0$	$\tilde{\mathcal{O}}(\Lambda^{\frac{-1}{d+1/\alpha}})$	$\tilde{\mathcal{O}}(\Lambda^{\frac{-1}{d}} + \Lambda^{-\alpha})$

Assumption 2(b)	High budget	Low budget
When $d = 0$	$e^{\tilde{\mathcal{O}}(-\Lambda^{\frac{\beta}{1+\beta}})}$	$e^{\tilde{\mathcal{O}}(-\Lambda^{\beta})} + e^{\tilde{\mathcal{O}}(-\Lambda)}$
When $d > 0$	$\tilde{\mathcal{O}}(\Lambda^{\frac{-1}{d}})$	

Assumption 2(c)	
When $d = 0$	$e^{\tilde{\mathcal{O}}(-\Lambda)}$
When $d > 0$	$\tilde{\mathcal{O}}(\Lambda^{\frac{-1}{d}})$

As explained in the next paragraph, in practice and for asymptotic comparisons only the results for high budget settings are relevant.

We can notice that the rates of decreasing are better until the threshold for high budget. This is because,

until the threshold, the algorithm does not have to focus on increasing the fidelity cost to improve the result, since the improvements in the regret it can make by exploring more cells is vastly superior to the improvements it can make with more precise analysis (which involves more precise evaluations: a higher fidelity). This explains why the rates are close to the one obtained on single-fidelity optimisation (or, similarly, on Assumption 2(c), which materializes this case). However, the low budget case actually requires a very low budget (or very accurate fidelities) so these rates are not really relevant in practice. This dichotomy was similarly noticed, by Bartlett et al. (2019) for the StroquOOL algorithm, in a stochastic case: using only one evaluation was enough as long as the noise did not exceed the potential regret that could be obtained.

7 Empirical results

We chose to do the same synthetic and practical deterministic experiments that the one done in Sen et al. (2018), and used their code for fair comparisons. The algorithm to which Kometo is compared are MFPDOO (Sen et al. (2018)), POO (Grill et al. (2015)) and SequOOL (Bartlett et al. (2019)). We directly used Kometo without any tweaking. This shows Kometo adaptability, which only needed the cost function and the space \mathcal{X} in order to work.

Experiments explanation Five of them are synthetic deterministic experiments of different, but always low, dimensions. The budget is expressed in terms of the number of multiple of the highest fidelity cost $\lambda(1)$. Note that these experiences may easily be unfair toward non multi-fidelity algorithms, because the results of the multi-fidelity algorithms heavily depends on how useful the low fidelities are, which is arbitrary on synthetic experiments. Therefore, since non multi-fidelity algorithms have no access to low fidelities and thus have less information, synthetic experiences should not be used to directly compare the efficiency of a multi-fidelity and a non multi-fidelity algorithm.

The last experience aims to measure the efficiency of the algorithms in practical settings. It involves tuning two hyperparameters for text classification, with the number of samples used to obtain 5-fold cross-validation accuracy determined by the fidelity. The budget is, for this experience, determined by the time used by the algorithm to return its result, reflecting simultaneously the actual time used for the algorithm execution and the cost of computing the accuracies.

Details about the experiments, along with comparisons to other multi-fidelity algorithms, can be found in Sen et al. (2018).

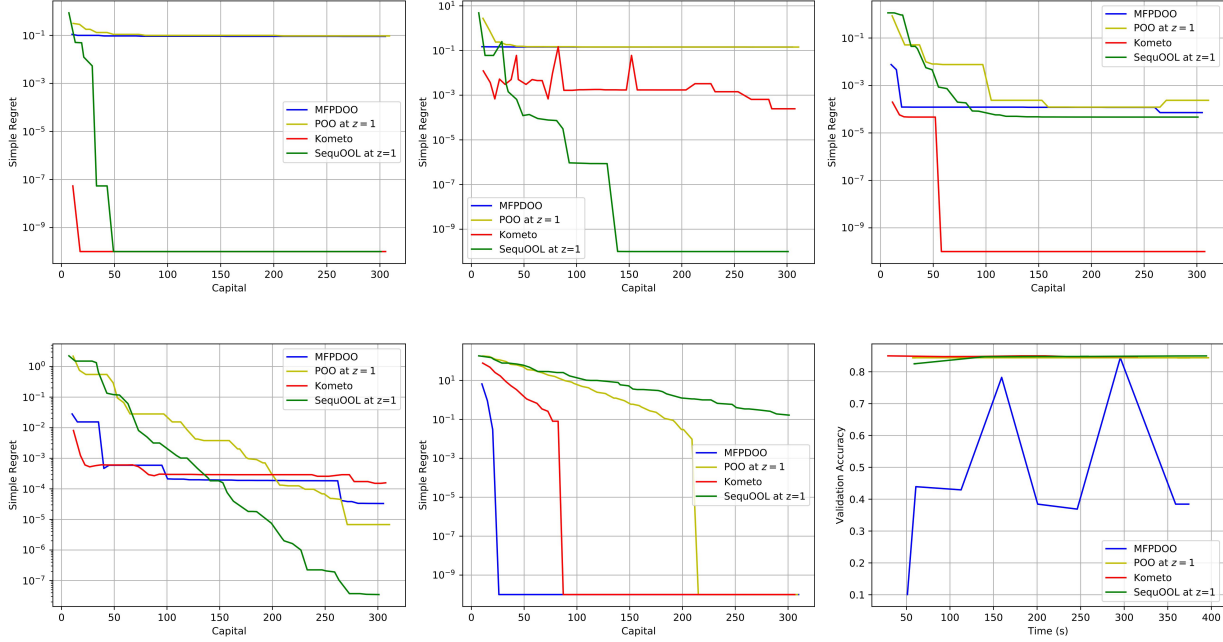


Figure 1: **(a) top-left:** Curin 2-dimensions, **(b) top-right:** Branin 2-dimensions **(c) top-left:** Hartman3d 3-dimensions **(d) bottom-right:** Hartman6d 6-dimensions **(e) bottom-right:** Borehole 8-dimensions **(f) bottom-right:** SVM 2-dimensions. Experiments are composed of five synthetic experiments, from (a) to (e), and one real-world, (f). The multi-fidelity algorithms can use all fidelities, while the non multi-fidelity algorithms only request at fidelity $z = 1$. The x-axis gives the budget effectively used by the algorithm, to reflect algorithm which exceed the attributed budget. The y-axis denotes the regret for the synthetic experiments (the lower the better), and the accuracy for the SVM experiment (the higher the better). For readability, the graph only plot the regret down to 10^{-10} .

Experiments analysis We can notice that *Kometo* largely outperforms MFPDOO on three of the synthetic experiments (Branin, Curin and Hartman3d) and on the practical experiment. It however gets beaten by MFPDOO on Borehole and Hartman6d by a relatively small margin. For the Hartman3d and Curin experiments, the better results of *Kometo* could be explained with its rank-based property, low fidelities may give highly accurate information on the way close points compare each other on the target function.

Interestingly, SequOOL outperforms *Kometo* on the Branin and Hartman6d experiments. This happens because, for these experiments, a lot of high-fidelity evaluations are needed to minimize the regret. Since *Kometo* keeps an important portion of its budget for low-fidelity evaluations, it is late compared to SequOOL which only does high-fidelity evaluations. This is materialized in the theoretical guarantees by the fact that *Kometo* has additional logarithmic factors compared to SequOOL under Assumption 2(c).

8 Discussion

Possible stochastic settings Our algorithm works in deterministic settings. However, our hypothesis of a

bounded bias can be replaced with an hypothesis of a noise (potentially biased), with the same bounds. Our algorithm can therefore work in stochastic settings, the guarantees being given instead at high probability with a cost-to-bias function changed accordingly.

However, in cases where the noise does not naturally decrease to 0 at higher fidelities, the Φ function will not decrease to 0 either although required by 2. This issue can be resolved by gradually increasing the number of evaluations at higher fidelities, to get a Φ function that would converges to 0. Indeed using concentration inequalities, we could then have Assumption 2 true with high probability, which could bound the regret.

Cumulative regret in adaptive multi-fidelity optimization Locatelli and Carpentier (2018) states that the minimax optimal cumulative regret with the knowledge of the smoothness cannot be attained by single-fidelity algorithm without the knowledge of the smoothness of the function. We wonder if this result remains true in multi-fidelity settings using adapted cumulative regret definitions.

References

- Auer, P., Ortner, R., and Szepesvári, C. (2007). Improved rates for the stochastic continuum-armed bandit problem. In *Conference on Learning Theory*.
- Bartlett, P. L., Gabillon, V., and Valko, M. (2019). A simple parameter-free and adaptive approach to optimization under a minimal local smoothness assumption. In *Algorithmic Learning Theory*.
- Bubeck, S., Munos, R., Stoltz, G., and Szepesvári, C. (2011). X-armed bandits. *Journal of Machine Learning Research*, 12:1587–1627.
- Cutler, M., Walsh, T. J., and How, J. P. (2014). Reinforcement learning with multi-fidelity simulators. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3888–3895. IEEE.
- Ghosh, S., Kristensen, J., Zhang, Y., Subber, W., and Wang, L. (2019). A Strategy for Adaptive Sampling of Multi-fidelity Gaussian Process to Reduce Predictive Uncertainty. *preprint*.
- Grill, J.-B., Valko, M., and Munos, R. (2015). Black-box optimization of noisy functions with unknown smoothness. In *Neural Information Processing Systems*.
- Hoorfar, A. and Hassani, M. (2008). Inequalities on the Lambert W function and hyperpower function. *Journal of Inequalities in Pure and Applied Mathematics*, 9(2):5–9.
- Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006). Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32(5):369–382.
- Kandasamy, K., Dasarthy, G., Oliva, J. B., Schneider, J., and Poczos, B. (2016a). Gaussian process bandit optimisation with multi-fidelity evaluations. *30th Conference on Neural Information Processing Systems*.
- Kandasamy, K., Dasarthy, G., Oliva, J. B., Schneider, J., and Poczos, B. (2016b). Multi-fidelity Gaussian Process Bandit Optimisation. *Journal of Artificial Intelligence Research*.
- Kandasamy, K., Dasarthy, G., Poczos, B., and Schneider, J. (2016c). The multi-fidelity multi-armed bandit. In *Advances in Neural Information Processing Systems*, pages 1777–1785.
- Kandasamy, K., Dasarthy, G., Schneider, J., and Poczos, B. (2017). Multi-fidelity bayesian optimisation with continuous approximations. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1799–1808. JMLR. org.
- Kawaguchi, K., Maruyama, Y., and Zheng, X. (2016). Global continuous optimization with error bound and fast convergence. *Journal of Artificial Intelligence Research*, 56:153–195.
- Kleinberg, R., Slivkins, A., and Upfal, E. (2008). Multi-armed bandit problems in metric spaces. In *Symposium on Theory Of Computing*.
- Li, L., Jamieson, K., DeSalvo, G., and Talwalkar, A. R. A. (2017). Hyperband: Bandit-based configuration evaluation for hyperparameter optimization. In *International Conference on Learning Representations*.
- Locatelli, A. and Carpentier, A. (2018). Adaptivity to Smoothness in X-armed bandits. In *Conference on Learning Theory*.
- Matyas, J. (1965). Random optimization. *Automation and Remote control*, 26(2):246–253.
- Munos, R. (2011). Optimistic optimization of deterministic functions without the knowledge of its smoothness. In *Neural Information Processing Systems*.
- Munos, R. (2014). From Bandits to Monte-Carlo Tree Search: The Optimistic Principle Applied to Optimization and Planning. *Foundations and Trends in Machine Learning*.
- Nesterov, Y. and Spokoiny, V. (2017). Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566.
- Sen, R., Kandasamy, K., and Shakkottai, S. (2018). Multi-Fidelity Black-Box Optimization with Hierarchical Partitions. *International Conference on Machine Learning*.
- Sen, R., Kandasamy, K., and Shakkottai, S. (2019). Noisy Blackbox Optimization with Multi-Fidelity Queries: A Tree Search Approach. *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*.
- Shang, X., Kaufmann, E., and Valko, M. (2019). General parallel optimization without metric. In *Algorithmic Learning Theory*.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. *International Conference on Machine Learning*.
- Valko, M., Carpentier, A., and Munos, R. (2013). Stochastic simultaneous optimistic optimization. In *International Conference on Machine Learning*.
- Zhang, Y., Hoang, T. N., Kian, B., Low, H., and Kankanhalli, M. (2019). Information-Based Multi-Fidelity Bayesian Optimization. *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*.

A Proof of Theorem 1

In the following proof, we set λ a bijective cost function, and $\zeta(z) = \Phi(x_z)$, where Φ is the upper bound of the cost-to-bias function given by the assumptions.

We first give a way to construct target and fidelity function for the lower bounds. Let T_a a truncated tree of \mathcal{P} (more precisely, a tree with the same root as \mathcal{P} , and included in \mathcal{P}). T_a nodes are therefore sub spaces of \mathcal{X} .

We then define:

- For all $h \geq 0$, T_a^h , the union of all the sub spaces associated to nodes of T_a of depth h . We can notice that $(T_a^h)_{h \geq 0}$ is a decreasing sequence for the inclusion, with $T_a^0 = \mathcal{P}_{0,0} = \mathcal{X}$.
- f^{T_a} by $f^{T_a}(x) = \sup_{h \geq 0} \{-\nu \rho^h | x \in T_a^h\}$, the target function
- $(f_z^{T_a})_{z \in Z}$ by $f_z^{T_a}(x) = \min\{f^{T_a}(x), -\zeta(z)\}$, the fidelities

We now state Lemma 5, related to the constructions above. It shows that the previous target and fidelity functions can be used as counter-examples in order to show the lower bounds.

Lemma 5. *Let T_a a truncated tree of \mathcal{P} such that, for every depth $h \geq 0$, T_a has between 1 and $\lfloor \rho^{-dh} \rfloor$ cells of depth h .*

Then:

- $f^{T_a} \in S(\mathcal{P}, \nu, \rho, d, C)$
- $(f_z^{T_a})_{z \in Z} \in F(\text{Asm}, f^{T_a}, \lambda)$

Proof. We start by showing that $f^{T_a} \in S(\mathcal{P}, \nu, \rho, d, C)$.

As T_a has an infinite depth, there exists a sequence $(i_h)_{h \in \mathbb{N}} \in \mathbb{N}$ of indexes such that $\forall h \in \mathbb{N}$, $\{\mathcal{P}_{h, i_h}\} \in T_a$. By taking $(x_h)_{h \in \mathbb{N}}$, such that for all $h \in \mathbb{N}$ $x_h \in \mathcal{P}_{h, i_h}$, we have that $\sup_{h \in \mathbb{N}} f^{T_a}(x_h) = 0$, which means $\sup_{x \in X} f^{T_a}(x) = 0$

Furthermore, we have that, for all $h \in \mathbb{N}$, for all $x \in \mathcal{P}_{h, i_h}$, $f^{T_a}(x) \geq -\nu \rho^h$ by definition of f^{T_a} , which is equivalent to $f^{T_a}(x) \geq \sup_{x' \in X} f^{T_a}(x') - \nu \rho^h$. We therefore have that Assumption 1 is true for f^{T_a} for the smoothness parameters (ν, ρ) .

We now show that d is a near-optimality dimension of f^{T_a} for these smoothness parameters. Let $h_o = \left\lfloor \frac{\log 3}{\log \frac{1}{\rho}} \right\rfloor$

(we defined C_{\min} with $C_{\min} = \left(\frac{K}{\rho^{-d}}\right)^{h_o}$). Let $h \in \mathbb{N}$, let's show that the related assumption is true at depth h : If $h \leq h_o$, then the number of near-optimal cells at depth h is simply bounded by the number of cells of depth h , K^h . As

$$K^h = \frac{K^{h_o}}{K^{h_o-h}} \leq \frac{K^{h_o}}{\rho^{-d(h_o-h)}} = \left(\frac{K}{\rho^{-d}}\right)^{h_o} \rho^{-dh} \leq C_{\min} \rho^{-dh},$$

the hypothesis is true at depth h .

If $h \geq h_o$, then thanks to hypothesis (2), there are at most $\lfloor \rho^{-d(h-h_o)} \rfloor$ cells of height $h - h_o$ such that $\sup_{x \in \mathcal{P}_{h-h_o, i}} f^{T_a}(x) > -\nu \rho^{h-h_o-1}$. By taking for each of these cells its K^{h_o} grandchildren, we get that there are at

most $\lfloor K^{h_o} \rho^{-d(h-h_o)} \rfloor$ cells of depth h such that $\sup_{x \in \mathcal{P}_{h, i}} f^{T_a}(x) > -\nu \rho^{h-h_o-1}$.

As

$$\left\lfloor K^{h_o} \rho^{-d(h-h_o)} \right\rfloor \leq C_{\min} \rho^{-dh}$$

and

$$-\nu\rho^{h-h_o-1} = -\nu\rho^{h-\left\lfloor \frac{\log 3}{\log \frac{1}{\rho}} + 1 \right\rfloor} \geq -\nu\rho^{h-\frac{\log 3}{\log \frac{1}{\rho}}} = -3\nu\rho^h = \sup_{x' \in X} f^{T_a}(x') - 3\nu\rho^h,$$

there are at most $C_{\min}\rho^{-dh}$ near-optimal cells at depth h , which concludes.

We finally want to show that $(f_z^{T_a})_{z \in Z} \in F(\text{Asm}, f^{T_a}, \lambda)$. As we defined $\zeta(z) = \Phi(x_z)$, where Φ is the upper bound given by Assumption 1, we just need to show, by taking $g_z = Id$ for all $z \in Z$, that $\|f^{T_a} - f_z^{T_a}\|_\infty \leq \zeta(z)$ for all $z \in Z$.

Taking $x \in X$, we either have

- $f^{T_a}(x) \leq -\zeta(z)$, and thus $f_z^{T_a}(x) = f^{T_a}(x)$, or
- $f^{T_a}(x) > -\zeta(z)$ which implies $|f^{T_a}(x) - f_z^{T_a}(x)| = f^{T_a}(x) + \zeta(z) \leq \zeta(z)$ as the function f^{T_a} is non-positive by definition.

The distinction between these two cases gives us that the bound on the infinity norm is true, and thus concludes. \square

Lemma 6. *By choosing appropriate target functions and fidelity approximations, we can get two lower bounds of the regret with π :*

Lower bound a: $r_\Lambda \geq \sup \left\{ r \in]0, \frac{1}{2}\nu\rho] \mid \Lambda \leq \frac{1}{K} \left(\frac{2r}{\nu\rho} \right)^{-d} \inf \Phi^{-1} \left(\left[0, \frac{2r}{\rho} \right] \right) \right\}$

Lower bound b: $r_\Lambda \geq \sup \left\{ r \in]0, \frac{1}{4}\nu\rho^6] \mid \Lambda \leq \left(\frac{\log \frac{\nu}{4r}}{4 \log \frac{1}{\rho}} - 2 \right) \inf \Phi^{-1} \left(\left[0, \frac{\nu}{\rho} \left(\frac{4r}{\nu} \right)^{\frac{1}{4}} \right] \right) \right\}$

Proof. We will denote by $(X_l^{T_a})_{l \in \mathbb{N}}$ the family of random variables equal to the successive points requested by π which has been given the fidelity approximations $(f_z^{T_a})$ (the values are null when l exceeds the number of evaluations). We also denote, for a given depth h , by $S_h^{T_a}$ the (also random) set of indexes l such that the fidelity z associated to the request of the point $X_l^{T_a}$ respects $\zeta(z) \leq \nu\rho^h$.

The key to the following proof is that, for any depth $h \geq 0$, for any tree T_a , if T'_a is the tree T_a whom nodes of depth $h' \geq h+1$ have been cut, then for any $l \geq 0$, $X_l^{T_a} \mid (\forall m \in S_h^{T_a}, m < l \implies X_m^{T_a} \notin T_a^{h+1})$ has the same distribution as $X_l^{T'_a} \mid (\forall m \in S_h^{T'_a}, m < l \implies X_m^{T'_a} \notin T_a^{h+1})$. This is because, in this instance, all the previous evaluations were either at a fidelity z too low and were equal to $-\zeta(z) \leq -\nu\rho^h$, or were evaluated outside of the sub space of a node of T_a of depth $h+1$. In both cases, the returned values were unaffected by the cutting of the nodes of depth $h' \geq h+1$ of T'_a , which implies that the behavior of the algorithm will remain the same for the request of the point $X_l^{T'_a}$.

We thus define the idea of opening a cell C as, when C is of depth h , having a $l \in S_h^{T_a}$ such that $X_l^{T_a} \in C$.

This means requesting a point of this cell with a high enough fidelity to differentiate the cells of T_a^{h+1} from the rest of the cells.

To simplify the proof, we will assume that the output of the algorithm is done with a last free evaluation at the highest fidelity. This evaluation will thus belong to $S_h^{T_a}$ for any $h \geq 0$, and the lemma will be proved by showing that this last evaluation has a certain chance not to belong to $T_a^{h'}$ for some h' when $(f_z^{T_a})$ is given to π .

We finally define, for any $h \geq 0$, $c_h = \inf \Phi^{-1}([0, \nu\rho^h])$ and $N_h = \left\lfloor \frac{\Lambda}{c_h} \right\rfloor + 1$. c_h is the minimum cost necessary to get an evaluation of bias $b \leq \nu\rho^h$, while N_h is the maximum number of such evaluations an algorithm can get, including the last free one. This gives, for any truncated tree T_a and $h \geq 0$, $\#S_h^{T_a} \leq N_h$.

Reminder: If $P_{h,i}$ is a cell of \mathcal{X} , its children are the K cells $P_{h+1,Ki}, \dots, P_{h+1,Ki+K-1}$.

We then prove the two different lower bounds:

Lower bound a (width limitation):

This lowerbound is based on the idea that, at a given depth h , the algorithm may have the budget to only open a fraction of the children of the near optimal cells (remember that the idea of opening, previously defined, also takes into account the fidelity level). As the behavior of the algorithm is potentially random, we choose the cell C that is the least likely to be opened as the only one potentially containing the optimum (ie $T_a^{h+1} = C$), and then bound the probability of its opening.

We first assume that there exists a non-negative integer h such that $\frac{\rho^{-dh} c_h}{\Lambda} \geq K$. Then, there also exists an integer s such that $\frac{\Lambda}{c_h} < K^s \leq \rho^{-dh}$. Because $\rho^{-dh} \geq 1$, we can assume s to be non-negative. Since K^s is an integer, we get that $K^s \geq \left\lfloor \frac{\Lambda}{c_h} \right\rfloor + 1 = N_h$. We also get that $K^s \leq \rho^{-dh} \leq K^h$ because of $d \leq \frac{\log K}{\log \frac{1}{\rho}}$, which implies $h - s \geq 0$.

We now define the trees we will use, along with Lemma 5, to lower bound the regret:

$$\text{We start with } T = \left(\bigcup_{h'=0}^{h-s-1} \{\mathcal{P}_{h',0}\} \right) \cup \left(\bigcup_{h'=h-s}^h \bigcup_{i=0}^{K^{h'-(h-s)}-1} \{\mathcal{P}_{h',i}\} \right)$$

We can first notice that T partially satisfies the hypothesis of lemma 5. Indeed, the first $h - s$ depths of T only have one node. For $h - s \leq h' \leq h$, there are $K^{h'-(h-s)}$ nodes of depth h' . Since $K^s \leq \rho^{-dh} \leq \rho^{-dh'} K^{h-h'}$, we indeed have $K^{h'-(h-s)} \leq \rho^{-dh'}$. However, the theorem also requires at least one node per depth.

We now define, for all integer i such that $0 \leq i \leq K^{s+1}$,

$$T_i = T \cup \left(\bigcup_{h'=h+1}^{\infty} \{\mathcal{P}_{h',K^{h'-(h+1)}i}\} \right)$$

We can notice that, for all i , T_i is a truncated tree of \mathcal{P} , has at least one node per depth, and the upper bound is still verified, as only one node was added for each empty depth was added. It thus verifies the hypothesis of the lemma.

We also define $p_i = \mathbb{P}(\exists l \in S_h^{T_i}, X_l^{T_i} \in \mathcal{P}_{h+1,i})$. It corresponds to the probability of the cell $\mathcal{P}_{h',i}$ being

opened when the algorithm is given $(f_z^{T_i})_{z \in Z}$. We then try to upper bound one of the p_i

$$\begin{aligned}
 \sum_{i=0}^{K^{s+1}-1} p_i &= \sum_{i=0}^{K^{s+1}-1} \mathbb{P} \left(\bigcup_{l \in \mathbb{N}} (l \in S_h^{T_i}, X_l^{T_i} \in \mathcal{P}_{h+1,i}) \right) \\
 &= \sum_{i=0}^{K^{s+1}-1} \mathbb{P} \left(\bigsqcup_{l \in \mathbb{N}} (l \in S_h^{T_i}, X_l^{T_i} \in \mathcal{P}_{h+1,i} \text{ and } (\forall m \in S_h^{T_i}, m < l \implies X_m^{T_i} \notin \mathcal{P}_{h+1,i})) \right) \\
 &= \sum_{i=0}^{K^{s+1}-1} \sum_{l \in \mathbb{N}} \mathbb{P} \left(l \in S_h^{T_i}, X_l^{T_i} \in \mathcal{P}_{h+1,i} \text{ and } (\forall m \in S_h^{T_i}, m < l \implies X_m^{T_i} \notin \mathcal{P}_{h+1,i}) \right) \\
 &= \sum_{i=0}^{K^{s+1}-1} \sum_{l \in \mathbb{N}} \mathbb{P} (l \in S_h^T, X_l^T \in \mathcal{P}_{h+1,i} \text{ and } (\forall m \in S_h^T, m < l \implies X_m^T \notin \mathcal{P}_{h+1,i})) \\
 &\leq \sum_{i=0}^{K^{s+1}-1} \sum_{l \in \mathbb{N}} \mathbb{P} (l \in S_h^T, X_l^T \in \mathcal{P}_{h+1,i}) \\
 &= \sum_{l \in \mathbb{N}} \sum_{i=0}^{K^{s+1}-1} \mathbb{P} (l \in S_h^T, X_l^T \in \mathcal{P}_{h+1,i}) \\
 &\leq \sum_{l \in \mathbb{N}} \mathbb{P} (l \in S_h^T) \\
 &= \mathbb{E}(\#S_h^T) \\
 &\leq N_h \\
 &\leq K^s
 \end{aligned}$$

We therefore have the existence of $i \leq K^{s+1} - 1$ such that $p_i \leq \frac{1}{K} \leq \frac{1}{2}$.

By giving π the fidelity approximations $(f_z^{T_i})_{z \in Z}$, since T_i respect the hypothesis of Lemma 5, we can lower bound the minimax regret with $\frac{1}{2}\nu\rho^h$

This result needed that $\frac{\rho^{-dh} c_h}{\Lambda} \geq K$ with h a positive integer, which can be rewrote $\frac{\rho^{-dh} \inf \Phi^{-1}([0, \nu\rho^h])}{K} \geq \Lambda$. Since this condition remains true if we take instead $h' \geq h$, we can just suppose that \tilde{h} is a positive real number such that $\frac{\rho^{-d\tilde{h}} \inf \Phi^{-1}([0, \nu\rho^{\tilde{h}}])}{K} \geq \Lambda$, and in this case get $r_\Lambda \geq \frac{1}{2}\nu\rho^{\lfloor \tilde{h} \rfloor + 1} \geq \frac{1}{2}\nu\rho^{\tilde{h}+1}$. By replacing \tilde{h} with $\frac{1}{2}\nu\rho^{\tilde{h}+1}$, this is the same thing as assuming that there exists $r \in]0, \frac{1}{2}\nu\rho]$ such that $\frac{(\frac{2r}{\nu\rho})^{-d} \inf \Phi^{-1}([0, \frac{2r}{\nu\rho}])}{K} \geq \Lambda$ to get $r_\Lambda \geq r$, which concludes.

Lower bound b (depth limitation):

In this second lower bound, the idea is that, after a depth h , the depth an algorithm can consistently reach when exploring a branch is at most proportional to the number of opening at depth h . We first take h such that $h \geq N_h$. We define recursively for $h' \geq h$, $(p_{h',i})_{0 \leq i \leq K-1}$, $i_{h'}$ and $T_{h'}$ with

$$\text{When } h' = h: \begin{cases} p_{h,i} = 1 \\ i_h = 0 \\ T_h = \bigcup_{h_p=0}^h \{\mathcal{P}_{h_p,0}\} \end{cases}$$

$$\text{When } h' > h \quad \left\{ \begin{array}{l} p_{h',i} = \mathbb{P}\left(\exists l \in S_h^{T_{h'-1}}, X_l^{T_{h'-1}} \in \mathcal{P}_{h',K i_{h'-1}+i} \text{ and } \left(\forall m \in S_h^{T_{h'-1}}, m < l \implies X_m^{T_{h'-1}} \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right) \\ i_{h'} = K i_{h'-1} + \arg \min_i p_{h',i} \\ T_{h'} = T_{h'-1} \cup \{\mathcal{P}_{h',i_{h'}}\} \end{array} \right.$$

We also define $T = \bigcup_{h'=h}^{\infty} T_{h'}$

Since T only has one infinite branch, we have that T verifies the hypothesis of Lemma 5.

Because, for every depth $h' \geq h+1$, the $p_{h',i}$ are K probability of disjoint events, we necessarily have that the probability associated to $i_{h'}$ is upper bounded by $\frac{1}{K} \leq \frac{1}{2}$.

We now define, for $h' \geq h+1$ the events $E_{h'} = \left(\exists l \in S_h^T, X_l^T \in \mathcal{P}_{h',i_{h'}} \text{ and } \left(\forall m \in S_h^T, m < l \implies X_m^T \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right)$, and the random variable $M = \sum_{h'=h+1}^{4h} \mathbb{1}_{E_{h'}}$. Let's first bound M with a certain probability, using Markov inequality.

$$\begin{aligned} \mathbb{E}(M) &= \sum_{h'=h+1}^{4h} \mathbb{E}(\mathbb{1}_{E_{h'}}) \\ &= \sum_{h'=h+1}^{4h} \mathbb{P}\left(\exists l \in S_h^T, X_l^T \in \mathcal{P}_{h',i_{h'}} \text{ and } \left(\forall m \in S_h^T, m < l \implies X_m^T \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right) \\ &= \sum_{h'=h+1}^{4h} \mathbb{P}\left(\exists l \in S_h^{T_{h'-1}}, X_l^{T_{h'-1}} \in \mathcal{P}_{h',i_{h'}} \text{ and } \left(\forall m \in S_h^{T_{h'-1}}, m < l \implies X_m^{T_{h'-1}} \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right) \\ &= \sum_{h'=h+1}^{4h} p_{h',i_{h'}-K i_{h'-1}} \\ &\leq \sum_{h'=h+1}^{4h} \frac{1}{2} \\ &= \frac{3h}{2} \end{aligned}$$

We then get $\mathbb{P}(M \geq 2h) \leq \frac{\mathbb{E}(M)}{2h} \leq \frac{3}{4}$

We now define the event $B = (\forall l \in S_h^T, X_l^T \notin \mathcal{P}_{4h,i_{4h}})$ and show that $(M < 2h) \subset B$.

We denote by j_1, \dots, j_t the different elements of S_h^T ranked from lowest to highest (with $t = \#S_h^T$), and we define the $(a_r)_{0 \leq r \leq t}$ with $a_0 = h$ and $a_r = \max \{h' \in \llbracket h ; 4h \rrbracket \mid \exists m \leq r, X_{j_m}^T \in \mathcal{P}_{h',i_{h'}}\} \cup \{h\}$ when $r > 0$. This definition ensures that a_r is superior or equal to the deepest opened depth after r requests of points at fidelities z of bias $\zeta(z) \leq \nu \rho^h$. We especially have that the sequence $(a_r)_{0 \leq r \leq t}$ is non-decreasing. Note that all these values are random variables, and that $(a_t < 4h)$ is exactly the event B .

We can then notice that, with these definitions, for any $h' \in \llbracket h+1 ; 4h \rrbracket$,

$$\begin{aligned} \left(\exists r \in \llbracket 0 ; t-1 \rrbracket, X_{j_{(r+1)}}^T \in \mathcal{P}_{h',i_{h'}} \text{ and } \left(\forall m < r+1, X_{j_m}^T \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right) = \\ \left(\exists l \in S_h^T, X_l^T \in \mathcal{P}_{h',i_{h'}} \text{ and } \left(\forall m \in S_h^T, m < l \implies X_m^T \notin \mathcal{P}_{h'-1,i_{h'-1}}\right)\right) = E_{h'} \quad (1) \end{aligned}$$

Using this equality, we can see that the existence of $r \in \llbracket 0 ; t-1 \rrbracket$ such that $h' \in \llbracket a_r+2 ; a_{r+1} \rrbracket$ implies $E_{h'}$.

Counting the h' then let us get that $\# \bigcup_{r=0}^{t-1} \llbracket a_r + 2 ; a_{r+1} \rrbracket \leq \sum_{h'=h+1}^{4h} \mathbb{1}_{E_{h'}}$.

Since all the sets $\llbracket a_r + 2 ; a_{r+1} \rrbracket$ are disjoint (because $(a_r)_{0 \leq r \leq t}$ is non-decreasing), we have $\sum_{r=0}^{t-1} (a_{r+1} - a_r - 1) \leq M$, which gives $a_t \leq a_0 + t + M$. As $h \geq N_h$ by hypothesis and $a_0 = h$, we have $a_t \leq 2h + M$. Thus, $(M < 2h)$ implies $(a_t < 4h)$, ie $(M < 2h)$ implies B.

Finally, with $\mathbb{P}(B) \geq \mathbb{P}(M < 2h) \geq \frac{1}{4}$, we can bound the regret with $r_\Lambda \geq \frac{1}{4}\nu\rho^{4h}$ by giving π the fidelity approximations $(f_z^T)_{z \in Z}$, since, as mentioned above, T respects the hypothesis of Lemma 5.

To conclude, we do the same thing as for Lower bound a. The hypothesis was the existence of h such that $h \geq N_h$ to get a bound $\nu\rho^{4h}$. It is especially the case when $(h+1) \inf \Phi^{-1}([0, \nu\rho^{h+1}]) \geq \Lambda$. This can similarly be changed to the existence of a real number $\tilde{h} \geq 1$ such that $\tilde{h} \inf \Phi^{-1}([0, \nu\rho^{\tilde{h}}]) \geq \Lambda$ to get a bound $\frac{1}{4}\nu\rho^{4(\tilde{h}+2)}$. With $r = \frac{1}{4}\nu\rho^{4(\tilde{h}+2)}$, we need $r \leq \frac{1}{4}\nu\rho^6$ and $\left(\frac{\log \frac{\nu}{4r}}{4 \log \frac{1}{\rho}} - 2\right) \inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right) \geq \Lambda$

□

Proof of Theorem 1: We now apply Lemma 6 to get the wanted lower bounds in the different cases.

Under Assumption 2(a):

We first try to solve $y = \Phi(c)$ for $c \geq 1$. We have $y = \frac{A}{c^\alpha}$, which means that $c = \left(\frac{y}{A}\right)^{-\frac{1}{\alpha}}$ when $A \geq y$. This implies that, if $r \leq \frac{\rho A}{2}$, $\inf \Phi^{-1}\left([0, \frac{2r}{\rho}]\right) = \left(\frac{2r}{\rho A}\right)^{-\frac{1}{\alpha}}$.

We thus try, for $r \leq r_{\min}^1 \triangleq \min\{\frac{1}{2}\rho A, \frac{1}{2}\nu\rho\}$, to solve the equation $\Lambda_r = \frac{1}{K}\left(\frac{2r}{\nu\rho}\right)^{-d} \inf \Phi^{-1}\left([0, \frac{2r}{\rho}]\right)$ in order to apply Lemma 6.a. We have $\Lambda_r = \frac{1}{K}\left(\frac{2}{\nu\rho}\right)^{-d}\left(\frac{2}{\rho A}\right)^{-\frac{1}{\alpha}} r^{-d-\frac{1}{\alpha}}$, which is equivalent to $K\left(\frac{2}{\nu\rho}\right)^d\left(\frac{2}{\rho A}\right)^{\frac{1}{\alpha}} \Lambda_r = r^{-d-\frac{1}{\alpha}}$.

With $D_1 \triangleq K^{\frac{-1}{d+\frac{1}{\alpha}}}\left(\frac{2}{\nu\rho}\right)^{\frac{-d}{d+\frac{1}{\alpha}}}\left(\frac{2}{\rho A}\right)^{\frac{-1}{1+d\alpha}}$, we have $D_1 \Lambda_r^{\frac{-1}{d+\frac{1}{\alpha}}} = r$.

Finally, using Lemma 6.a, if $\Lambda \geq \Lambda_{r_{\min}^1}$, we have the bound $r_\Lambda \geq D_1 \Lambda^{\frac{-1}{d+\frac{1}{\alpha}}}$. Otherwise, if $\Lambda < \Lambda_{r_{\min}^1}$ we only get that $r_\Lambda \geq r_{\min}^1$.

Under Assumption 2(b)

d = 0 : Similarly, we solve $y = \Phi(c)$ for $c \geq 1$. Since $\Phi(c) = B e^{-\frac{c^\beta}{\sigma}}$, we get $c = \left(\sigma \log\left(\frac{B}{y}\right)\right)^{\frac{1}{\beta}}$ when $B \geq y$. Using this, we have that if $\frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}} \leq B$, ie $r \leq \left(\frac{\rho B}{\nu}\right)^4 \frac{\nu}{4}$, then $\inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right) = \left(\sigma \log\left(\frac{\rho B}{\nu}\left(\frac{\nu}{4r}\right)^{\frac{1}{4}}\right)\right)^{\frac{1}{\beta}}$.

To use Lemma 6.b, we then try to solve, for $r \leq \min\{\frac{1}{4}\nu\rho^6, \left(\frac{\rho B}{\nu}\right)^4 \frac{\nu}{4}\}$, the equation $\Lambda_r = \left(\frac{\log \frac{\nu}{4r}}{4 \log \frac{1}{\rho}} - 2\right) \inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right)$. This equation is equivalent to $4^{1+\frac{1}{\beta}} \log \frac{1}{\rho} \sigma^{\frac{-1}{\beta}} \Lambda_r = \left(\log \frac{\nu}{4r} - 8 \log \frac{1}{\rho}\right) \left(\log\left(\left(\frac{B}{\nu}\right)^4 \frac{\nu}{4r}\right)\right)^{\frac{1}{\beta}}$. We define $L_1 \triangleq \log \frac{\nu}{4} - 8 \log \frac{1}{\rho}$, $L_2 \triangleq \log\left(\left(\frac{B}{\nu}\right)^4 \frac{\nu}{4}\right)$, and $D \triangleq 4^{1+\frac{1}{\beta}} \log \frac{1}{\rho} \sigma^{\frac{-1}{\beta}}$ to get it in the form $D \Lambda_r = (L_1 + \log \frac{1}{r})(L_2 + \log \frac{1}{r})^{\frac{1}{\beta}}$. We now also assume $L_1 \geq \frac{-1}{2} \log \frac{1}{r}$ and $L_2 \geq \frac{-1}{2} \log \frac{1}{r}$ (equivalent to $r \leq e^{2L_1}$ and $r \leq e^{2L_2}$), to get $D \Lambda_r \geq \left(\frac{1}{2} \log \frac{1}{r}\right) \left(\frac{1}{2} \log \frac{1}{r}\right)^{\frac{1}{\beta}}$, ie $2D^{\frac{\beta}{1+\beta}} \Lambda_r^{\frac{\beta}{1+\beta}} \geq \log\left(\frac{1}{r}\right)$. This gives $e^{-D_2} \Lambda_r^{\frac{\beta}{1+\beta}} \leq r$ with $D_2 \triangleq 2D^{\frac{\beta}{1+\beta}}$.

We now conclude like above, this time using Lemma 6.b. If we define $r_{\min}^2 \triangleq \min\{\frac{1}{4}\nu\rho^6, \left(\frac{\rho B}{\nu}\right)^4 \frac{\nu}{4}, e^{2L_1}, e^{2L_2}\}$,

for all $\Lambda \geq \Lambda_{\min}^2$, we can get the bound $r_\Lambda \geq e^{-D_2 \Lambda^{\frac{\beta}{1+\beta}}}$, and $r_\Lambda \geq r_{\min}^2$ otherwise.

d > 0 : In this case, we simply use that $\inf \Phi^{-1}\left([0, \frac{2r}{\rho}]\right) \geq 1$ by definition of Φ . We try to apply the Lemma 6.a, and solve for $r \leq \frac{1}{2}\nu\rho$, $\Lambda_r = \frac{1}{K}\left(\frac{2r}{\nu\rho}\right)^{-d} \inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right)$. We here have $\Lambda_r \geq \frac{1}{K}\left(\frac{2r}{\nu\rho}\right)^{-d}$, which means $\frac{\nu\rho}{2} K^{-\frac{1}{d}} \Lambda_r^{-\frac{1}{d}} \leq r$. With $D_3 = \frac{\nu\rho}{2} K^{-\frac{1}{d}}$, we have $D_3 \Lambda_r^{-\frac{1}{d}} \leq r$, and we conclude like above with $r_{\min}^3 = \frac{1}{2}\nu\rho$.

Under Assumption 2(c)

d=0: We try to apply the Lemma 6.b, and solve, for $r \leq r_{\min}^4 = \frac{1}{4}\nu\rho^6$, $\Lambda_r \leq \left(\frac{\log \frac{\nu}{4r}}{4 \log \frac{1}{\rho}} - 2\right) \inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right)$. Since the required cost to get any information on the function is a , we have that $\inf \Phi^{-1}\left([0, \frac{\nu}{\rho}\left(\frac{4r}{\nu}\right)^{\frac{1}{4}}]\right) = a$, which means that $\Lambda_r = \left(\frac{\log \frac{\nu}{4r}}{4 \log \frac{1}{\rho}} - 2\right)a$, ie $\log \frac{4}{\nu} + 8 \log \frac{1}{\rho} + 4 \log \frac{1}{\rho} \frac{\Lambda_r}{a} = \log \frac{1}{r}$. If we assume $\Lambda_r \geq \frac{a}{4 \log \frac{1}{\rho}} \left(\log \frac{4}{\nu} + 8 \log \frac{1}{\rho}\right)$, we have $8 \log \frac{1}{\rho} \frac{\Lambda_r}{a} \geq \log \frac{1}{r}$, and with $D_4 = 8 \log \frac{1}{\rho} \frac{1}{a}$, we finally get $e^{-D_4 \Lambda_r} \leq r$.

We then conclude like before by applying the Lemma 6.b, and then, for any $\Lambda \geq \max \left\{ \Lambda_{\min}^4, \frac{a}{4 \log \frac{1}{\rho}} \left(\log \frac{4}{\nu} + 8 \log \frac{1}{\rho}\right) \right\}$, we can get the bound $r_\Lambda \geq e^{-D_4 \Lambda_r}$.

d > 0 : The reasoning and the wanted bounds (with $D_3 = D_5$) are exactly the same as the case $d > 0$ of Assumption 2(b). □

B Upper bounds

Proof of Proposition 2. When a cell is opened at fidelity j , the maximum budget used for this opening can not exceed Ke^j . We can therefore upper bound the budget used for all of the cell opening, ignoring the initial cell opening of $\mathcal{P}_{0,0}$

$$\begin{aligned}
 \sum_{h=1}^{\lfloor \tilde{\Lambda} \rfloor} \sum_{m=1}^{\lfloor \frac{\tilde{\Lambda}}{h} \rfloor} \sum_{j=0}^{\lfloor \log \frac{\tilde{\Lambda}}{hm} \rfloor} Ke^j &\leq \sum_{h=1}^{\lfloor \tilde{\Lambda} \rfloor} \sum_{m=1}^{\lfloor \frac{\tilde{\Lambda}}{h} \rfloor} \frac{Ke^{\lfloor \log \frac{\tilde{\Lambda}}{hm} \rfloor + 1}}{e-1} \\
 &\leq \sum_{h=1}^{\lfloor \tilde{\Lambda} \rfloor} \sum_{m=1}^{\lfloor \frac{\tilde{\Lambda}}{h} \rfloor} \frac{Ke\tilde{\Lambda}}{(e-1)hm} \\
 &\leq \sum_{h=1}^{\lfloor \tilde{\Lambda} \rfloor} \frac{Ke\tilde{\Lambda}}{(e-1)h} \left(\log \frac{\tilde{\Lambda}}{h} + 1 \right) \\
 &\leq \frac{Ke\tilde{\Lambda}}{(e-1)} \left(\log \tilde{\Lambda} + 1 \right) \left(\log \frac{\tilde{\Lambda}}{h} + 1 \right) \\
 &\leq \frac{Ke\tilde{\Lambda}}{(e-1)} (\log \Lambda + 1)^2 \\
 &\leq \frac{\Lambda}{2}
 \end{aligned}$$

The budget used for the initial opening and for the cross-validation can be bounded by

$$\left(K + \log \tilde{\Lambda} \right) \tilde{\Lambda} \leq K(\log \Lambda + 1) \tilde{\Lambda} \leq \frac{\Lambda}{2}$$

This shows that the total budget can be bounded by Λ . □

Lemma 7. *Let Ψ a non-increasing function upper bounding Φ , j a non-negative integer and \tilde{h} a positive real number such that*

1. $\Psi(e^j) \leq \nu \rho^{\tilde{h}}$, and
2. $\frac{\tilde{\Lambda}}{4he^j} \geq C\rho^{-d\tilde{h}}$.

Then, $r_\Lambda \leq \frac{3\nu}{\rho} \rho^{\tilde{h}} + 2\Psi(\tilde{\Lambda})$.

Proof. We denote by $\mathcal{P}_{x^*, h'}$ the cell of depth h' containing x^* , i^* the associated index and $h = \lfloor \tilde{h} \rfloor$.

Case $h > 0$: To prove this lemma, we first want to show that the cell \mathcal{P}_{x^*} , is opened with a cost c_h with $c \geq e^j$. We do it by induction, and show this is true for any h' such that $h' \leq h$

When $h' = 0$, the proposition is trivial since there is only one cell of depth 0, opened by the algorithm with $c_0 = \tilde{\Lambda} \geq e^j$ because of (2) and $\tilde{h} > 1$.

When $0 < h' \leq h$, we assume that $\mathcal{P}_{x^*, h'-1}$ is opened, and $c_{h'-1} \geq e^j$. We want to show that at least $C\rho^{-dh'}$ cells are opened at a cost superior or equal to e^j and strictly inferior to e^{j+1} . This number $n_{h'}$ of cells is equal to $n_{h'} = \#\{m \in \llbracket 1 ; \frac{\tilde{\Lambda}}{h'} \rrbracket | e^{j+1} > \frac{\tilde{\Lambda}}{h'm} \geq e^j\} = \#\{m \in \mathbb{N}^* | \frac{\tilde{\Lambda}}{h'e^{j+1}} < m \leq \frac{\tilde{\Lambda}}{h'e^j}\} \geq \frac{\tilde{\Lambda}}{h'e^j} - \frac{\tilde{\Lambda}}{h'e^{j+1}} - 1 \geq \frac{\tilde{\Lambda}}{2h'e^j} - 1$. Since $\frac{\tilde{\Lambda}}{2h'e^j} \geq \frac{\tilde{\Lambda}}{2he^j} \geq 2C\rho^{-d\tilde{h}} \geq 2$, we conclude that $n_{h'} \geq \frac{\tilde{\Lambda}}{4h'e^j} \geq \frac{\tilde{\Lambda}}{4he^j} \geq C\rho^{-d\tilde{h}} \geq C\rho^{-dh'}$.

We now suppose that the cell $\mathcal{P}_{x^*, h'-1}$ was not opened with $c_{h'} \geq e^j$. Then at least $n_{h'}$ cells of depth h' were such that $f_{h', i, j} \geq f_{h', i^*, j}$, with i denoting the index of any of these cells. This means that, using Assumption 1 and hypothesis (1) of the lemma,

$$\sup_{x \in \mathcal{P}_{h', i}} f(x) + \nu \rho^{h'} \geq f(x_{h', i}) + \Psi(e^j) \geq f_{z_{e^j}}(x_{h', i}) = f_{h', i, j} \geq f_{h', i^*, j} = f_{z_{e^j}}(x_{h', i^*}) \geq f(x_{h', i^*}) - \Psi(e^j) \geq f(x_{h', i^*}) - \nu \rho^{h'} \geq f(x^*) - 2\nu \rho^{h'} (\Delta)$$

absurd since $N_{h'} \leq C\rho^{-dh'} \leq n_{h'}$. We thus have that $\mathcal{P}_{x^*, h'-1}$ was opened with $c_{h'} \geq e^j$

We have shown that the cell \mathcal{P}_{x^*} , has been opened with a cost $c_h \geq e^j$. This means that the cell has been evaluated at least once for a fidelity of cost $e^{j'}$ with $j' \geq j$. Using the same chain of inequality (Δ) for h , and $\Psi(e^{j'}) \leq \Psi(e^j)$, we get that the cross-validation candidate $x_j^{c'}$ is such that $f(x_j^c) \geq f(x^*) - 3\nu \rho^h$. Then, if $x_{j^*}^c$ is the returned value, we have that

$$f(x_{j^*}^c) + \Psi(\tilde{\Lambda}) \geq f_{z_{\tilde{\Lambda}}}(x_{j^*}^c) \geq f_{z_{\tilde{\Lambda}}}(x_{j'}^c) \geq f(x_{j'}^c) - \Psi(\tilde{\Lambda}) \geq f(x^*) - \Psi(\tilde{\Lambda}) - 3\nu \rho^h$$

, and we can conclude that $r_\Lambda \leq 3\frac{\nu}{\rho} \rho^{-d\tilde{h}}$ since $\tilde{h} \geq h - 1$

Case $h=0$: Since any returned x^c belongs to $\mathcal{P}_{x^*, 0}$, we have that $r_\Lambda \leq \nu \rho^0 \leq \frac{\nu}{\rho} \rho^{\tilde{h}} \leq \frac{3\nu}{\rho} \rho^{\tilde{h}} + 2\Psi(\tilde{\Lambda})$ □

Proof of Theorem 3. The proofs for the first two hypothesis are similar since they use the same techniques to apply the lemma

Under Assumption 2(a) ($\Phi(c) \leq \frac{\Lambda}{c^\alpha}$):

We first try to solve $\frac{\tilde{\Lambda} \nu^{\frac{1}{\alpha}} \rho^{\frac{h}{\alpha}}}{4eA^{\frac{1}{\alpha}} h} = C\rho^{-dh}$ with h unknown. Since this equality is equivalent to $\frac{\tilde{\Lambda} \nu^{\frac{1}{\alpha}} (d + \frac{1}{\alpha}) \log \frac{1}{\rho}}{4CeA^{\frac{1}{\alpha}} h} = (d + \frac{1}{\alpha}) \log \frac{1}{\rho} h e^{(d + \frac{1}{\alpha}) \log \frac{1}{\rho} h}$, there exists a single positive real number h , which we will

name h_1 , such that the equality is respected, with $h_1 = \frac{1}{(d+\frac{1}{\alpha})\log \frac{1}{\rho}} W(\frac{\tilde{\Lambda}\nu^{\frac{1}{\alpha}}(d+\frac{1}{\alpha})\log \frac{1}{\rho}}{4CeA^{\frac{1}{\alpha}}})$. We define the same way h_2 as the only solution of the equation $\frac{\tilde{\Lambda}}{4h} = C\rho^{-dh}$ (equal to $\frac{\tilde{\Lambda}}{4C}$ if $d = 0$ and $\frac{1}{d\log \frac{1}{\rho}} W(\frac{\tilde{\Lambda}d\log \frac{1}{\rho}}{4C})$ if $d > 0$). Note that h_1 and h_2 are both increasing with the budget. We also define Ψ with $\Psi(x) = \frac{A}{x^\alpha}$. We then discriminate between the two cases:

First case (High Budget): $\nu\rho^{h_1} \leq e^\alpha A$

We set $j_1 = \lfloor \log(\frac{eA^{\frac{1}{\alpha}}}{\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}}) \rfloor$. Thanks to the hypothesis, we have that j_1 is a positive integer.

We then try to apply Lemma 7 on Ψ , h_1 and j_1 . Since $j_1 > \log(\frac{eA^{\frac{1}{\alpha}}}{\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}}) - 1 = \log(\frac{A^{\frac{1}{\alpha}}}{\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}})$, we have $\frac{A}{e^{\alpha j_1}} < \nu\rho^{h_1}$ which means $\Psi(e^{j_1}) < \nu\rho^{h_1}$, hypothesis (1).

We also have $j_1 \leq \log(\frac{eA^{\frac{1}{\alpha}}}{\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}})$, ie $\frac{1}{e^{j_1}} \geq \frac{\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}}{eA^{\frac{1}{\alpha}}}$, which gives $C\rho^{-dh_1} = \frac{\tilde{\Lambda}\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}}{4eA^{\frac{1}{\alpha}}h_1} \leq \frac{\tilde{\Lambda}}{4h_1e^{j_1}}$, hypothesis (2). Applying the lemma then let us obtain the claimed results.

Second case (low budget): $\nu\rho^{h_1} > e^\alpha A$

We set $j_2 = 0$ and here try to apply Lemma 7 on Ψ , h_2 and j_2 . We have $h_1\rho^{-dh_1} = \frac{\tilde{\Lambda}\nu^{\frac{1}{\alpha}}\rho^{\frac{h_1}{\alpha}}}{4eA^{\frac{1}{\alpha}}C} > \frac{\tilde{\Lambda}}{4C} = h_2\rho^{-dh_2}$, which means $h_1 > h_2$ since $h \mapsto h\rho^{-dh}$ is increasing on \mathbb{R}_+ . This gives $\Psi(e^{j_2}) = A < e^\alpha A < \nu\rho^{h_1} < \nu\rho^{h_2}$ (1). The definition of h_2 immediately gives $\frac{\tilde{\Lambda}}{4h_2e^{j_2}} \geq C\rho^{-dh_2}$ (2), and we conclude with the lemma.

Under Assumption 2(b) ($\Phi(c) \leq Be^{-\frac{c^\beta}{\sigma}}$):

We first define $a_{b,\nu}$ as equal to $\max(\frac{1}{2\sigma}, \log(\frac{B}{\nu}))$. Similarly to assumption 2.a, we then define h_1 as the only real positive number h such that $\frac{\tilde{\Lambda}}{4he(2\sigma h \log \frac{1}{\rho})^{\frac{1}{\beta}}} = C\rho^{-dh}$, and h_2 the only real positive number h such that

$$\frac{\tilde{\Lambda}}{4he(2\sigma a_{b,\nu})^{\frac{1}{\beta}}} = C\rho^{-dh}.$$

Value of h_1 : We have that $\frac{\tilde{\Lambda}}{4Ce(2\sigma \log \frac{1}{\rho})^{\frac{1}{\beta}}} = h_1^{\frac{1+\beta}{\beta}} \rho^{-dh_1}$, which gives $(\frac{\tilde{\Lambda}}{4Che})^{\frac{\beta}{\beta+1}} (\frac{1}{2\sigma \log \frac{1}{\rho}})^{\frac{1}{\beta+1}} = h_1 \rho^{-\frac{\beta}{\beta+1}dh}$.

When $d = 0$, we then get $h_1 = (\frac{\tilde{\Lambda}}{4Che})^{\frac{\beta}{\beta+1}} (\frac{1}{2\sigma \log \frac{1}{\rho}})^{\frac{1}{\beta+1}}$. When $d > 0$, we have $\frac{\beta}{\beta+1} d \log \frac{1}{\rho} (\frac{\tilde{\Lambda}}{4Che})^{\frac{\beta}{\beta+1}} (\frac{1}{2\sigma \log \frac{1}{\rho}})^{\frac{1}{\beta+1}} = (\frac{\beta}{\beta+1} d \log \frac{1}{\rho}) h_1 e^{\frac{\beta}{\beta+1} d \log \frac{1}{\rho} h_1}$. We finally get $h_1 = \frac{\beta+1}{\beta d \log \frac{1}{\rho}} W(\frac{\beta}{\beta+1} d \log \frac{1}{\rho} (\frac{\tilde{\Lambda}}{4Che})^{\frac{\beta}{\beta+1}} (\frac{1}{2\sigma \log \frac{1}{\rho}})^{\frac{1}{\beta+1}})$

Value of h_2 : The definition gives $\frac{\tilde{\Lambda}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}} = h_2 \rho^{-dh_2}$. If $d = 0$, we have $h_2 = \frac{\tilde{\Lambda}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}}$. If $d > 0$, we get, like for a) $h_2 = \frac{1}{d \log \frac{1}{\rho}} W(\frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}})$.

We also define Ψ with $\Psi(x) = Be^{-\frac{x^\beta}{\sigma}}$. We again discriminate between the two cases.

First case (High Budget): $h_1 \geq \frac{a_{b,\nu}}{\log \frac{1}{\rho}}$

We set $j_1 = \lfloor \frac{1}{\beta} \log(2\sigma h_1 \log \frac{1}{\rho}) \rfloor + 1$. Since $2\sigma h_1 \log \frac{1}{\rho} \geq 2\sigma a_{b,\nu} \geq 1$, $j_1 \geq 1$. We thus try to apply Lemma 7 on Ψ , h_1 and j_1 . Thanks to

$$e^{\beta j_1} \geq 2\sigma h_1 \log \frac{1}{\rho} \geq \sigma(h_1 \log \frac{1}{\rho} + \log \frac{B}{\nu})$$

we have $Be^{-\frac{e^{\beta j_1}}{\sigma}} \leq \nu\rho^{h_1}$, ie $\Psi(e^{j_1}) \leq \nu\rho^{h_1}$, hypothesis (1).

Since $e^{j_1} \leq e(2\sigma h_1 \log \frac{1}{\rho})^{\frac{1}{\beta}}$, we have

$$\frac{\tilde{\Lambda}}{4h_1e^{j_1}} \geq \frac{\tilde{\Lambda}}{4h_1e(2\sigma h_1 \log \frac{1}{\rho})^{\frac{1}{\beta}}} = C\rho^{-dh_1} \quad (2)$$

We then get the wanted results by applying the lemma.

Second case (Low Budget): $h_1 < \frac{a_{b,\nu}}{\log \frac{1}{\rho}}$

We set $j_2 = \lfloor \frac{1}{\beta} \log(2\sigma a_{b,\nu}) \rfloor + 1 \geq 1$ and apply Lemma 7 on Ψ, h_2 and j_2 .

The proof is really similar to the previous case, but we first need to show that $h_1 \geq h_2$. This is due to

$$h_1 \rho^{-dh_1} = \frac{\tilde{\Lambda}}{4Ce(2\sigma h_1 \log \frac{1}{\rho})^{\frac{1}{\beta}}} \geq \frac{\tilde{\Lambda}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}} = h_2 \rho^{-dh_2}$$

which implies $h_1 \geq h_2$. We then use this inequality to show

$$e^{\beta j_2} \geq 2\sigma a_{b,\nu} \geq \sigma(h_1 \log \frac{1}{\rho} + \log \frac{B}{\nu}) \geq \sigma(h_2 \log \frac{1}{\rho} + \log \frac{B}{\nu})$$

and we conclude like before with $\Psi(e^{j_2}) \leq \nu \rho^{h_2}$ (1).

Like in the previous case, we have $e^{j_2} \leq e(2\sigma a_{b,\nu})^{\frac{1}{\beta}}$ and thus

$$\frac{\tilde{\Lambda}}{4h_2 e^{j_2}} \geq \frac{\tilde{\Lambda}}{4h_2 e(2\sigma a_{b,\nu})^{\frac{1}{\beta}}} = C \rho^{-dh_2} \quad (2)$$

which let us conclude with the lemma.

Under Assumption 2(c) ($\Phi(c) = 0$ for all $c \geq a$)

We define h_1 as the only real solution to the equation $\frac{\tilde{\Lambda}}{4ae h} = C \rho^{-dh}$. Similarly to the two previous cases, we get $h_1 = \frac{\tilde{\Lambda}}{4Cae}$ when $d=0$ and $h_1 = \frac{1}{d \log \frac{1}{\rho}} W(\frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4Cae})$ when $d > 0$.

If $\tilde{\Lambda} \geq a$:

We now define Ψ with $\Psi(x) = \begin{cases} +\infty & \text{if } x < a \\ 0 & \text{if } x \geq a \end{cases}$, j_1 with $j_1 = \lfloor \log(ea) \rfloor$ and we try to apply Lemma 7 on

Ψ, h_1 and j_1 .

Since $e^{j_1} \geq a$, $\Psi(e^{j_1}) \leq \nu \rho^{h_1}$ (1). Using $e^{j_1} \leq ea$, we have

$$\frac{\tilde{\Lambda}}{4h_1 e^{j_1}} \geq \frac{\tilde{\Lambda}}{4h_1 ea} = C \rho^{-dh_1} \quad (2)$$

. We then use the lemma (and $\Psi(\tilde{\Lambda}) = 0$, but in its proof, since we can upper bound $\Psi(e^{j'})$ by 0 in this specific case, we can loose a 3 factor in the result and just keep $\frac{\nu}{\rho} \rho^{h_1}$

If $\tilde{\Lambda} < a$:

We then have $h_1 < 1$. Since $r_\Lambda \leq \nu$, as explained in the lemma proof, we have $r_\Lambda \leq \frac{\nu}{\rho} \rho^{h_1}$ which concludes. \square

Proof of Corollary 4. For the proof of the different bounds, we use that $W(x) \geq \log x - \log \log x$ for $x \geq e$, as shown in [Hoorfar and Hassani \(2008\)](#). We can also notice that for $x \leq e$, we have $W(x) \geq x/e$.

We remind that the value of $\tilde{\Lambda}$ is bounded by

$$\frac{(e-1)\Lambda}{2Ke(\log \Lambda + 1)^2}$$

Under Assumption 2(a) ($\Phi(c) \leq \frac{A}{c^\alpha}$):

High Budget: We define $\tilde{\Lambda}_c = \frac{\tilde{\Lambda} \nu^{\frac{1}{\alpha}} (d + \frac{1}{\alpha}) \log \frac{1}{\rho}}{4C e A^{\frac{1}{\alpha}}}$.

If $\tilde{\Lambda}_c \geq e$ then

$$\begin{aligned} r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{(d+\frac{1}{\alpha}) \log \frac{1}{\rho}}} W(\tilde{\Lambda}_c) + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &\leq \frac{3\nu}{\rho} \rho^{\frac{1}{(d+\frac{1}{\alpha}) \log \frac{1}{\rho}} \log \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)^{\frac{-1}{d+\frac{1}{\alpha}}} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \end{aligned}$$

If $\tilde{\Lambda}_c \leq e$, then

$$\begin{aligned} r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{(d+\frac{1}{\alpha}) \log \frac{1}{\rho}}} W(\tilde{\Lambda}_c) + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{(d+\frac{1}{\alpha}) \log \frac{1}{\rho}} \frac{\tilde{\Lambda}_c}{e}} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \end{aligned}$$

Low Budget: If $d = 0$, the exponential upper bound is given by

$$r_\Lambda \leq \frac{3\nu}{\rho} \rho^{h_2} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \leq \frac{3\nu}{\rho} \rho^{\frac{\tilde{\Lambda}}{4C}} + 2 \frac{A}{\tilde{\Lambda}^\alpha}$$

If $d > 0$, we define $\tilde{\Lambda}_c = \frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4C}$

Then, if $\tilde{\Lambda}_c \geq e$

$$\begin{aligned} r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}}} W(\tilde{\Lambda}_c) + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &\leq \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \log \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\ &= \frac{3\nu}{\rho} \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)^{\frac{-1}{d}} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \end{aligned}$$

Else, if $\tilde{\Lambda}_c < e$

$$\begin{aligned}
 r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\
 &= \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}}} W(\tilde{\Lambda}_c) + 2 \frac{A}{\tilde{\Lambda}^\alpha} \\
 &\leq \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \frac{\tilde{\Lambda}_c}{e}} + 2 \frac{A}{\tilde{\Lambda}^\alpha}
 \end{aligned}$$

Under Assumption 2(b) ($\Phi(c) \leq Be^{\frac{-c^\beta}{\sigma}}$):

High Budget: If $d = 0$, then we directly get

$$r_\Lambda \leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \leq \frac{3\nu}{\rho} \rho^{\left(\frac{\tilde{\Lambda}}{4Ch e}\right)^{\frac{\beta}{\beta+1}} \left(\frac{1}{2\sigma \log \frac{1}{\rho}}\right)^{\frac{1}{\beta+1}}} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}}$$

If $d > 0$, we define $\tilde{\Lambda}_c = \left(\frac{\beta}{\beta+1} d \log \frac{1}{\rho} \left(\frac{1}{2\sigma \log \frac{1}{\rho}}\right)^{\frac{1}{\beta+1}}\right)^{\frac{\beta+1}{\beta}} \frac{\tilde{\Lambda}}{4Ch e}$

Then, if $\tilde{\Lambda}_c \geq e$

$$\begin{aligned}
 r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \\
 &= \frac{3\nu}{\rho} \rho^{\frac{\beta+1}{\beta d \log \frac{1}{\rho}} W(\tilde{\Lambda}_c^{\frac{\beta}{\beta+1}})} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \\
 &\leq \frac{3\nu}{\rho} \rho^{\frac{\beta+1}{\beta d \log \frac{1}{\rho}} \log \left(\frac{\tilde{\Lambda}_c^{\frac{\beta}{\beta+1}}}{\log \tilde{\Lambda}_c}\right)} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \\
 &= \frac{3\nu}{\rho} \left(\frac{\tilde{\Lambda}_c}{(\log \tilde{\Lambda}_c)^{\frac{\beta+1}{\beta}}}\right)^{\frac{-1}{d}} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}}
 \end{aligned}$$

Else, if $\tilde{\Lambda}_c < e$

$$\begin{aligned}
 r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \\
 &= \frac{3\nu}{\rho} \rho^{\frac{\beta+1}{\beta d \log \frac{1}{\rho}} W(\tilde{\Lambda}_c^{\frac{\beta}{\beta+1}})} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \\
 &\leq \frac{3\nu}{\rho} \rho^{\frac{\beta+1}{\beta d \log \frac{1}{\rho}} \frac{\tilde{\Lambda}_c^{\frac{\beta}{\beta+1}}}{e}} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}}
 \end{aligned}$$

Low Budget: If $d = 0$, we get

$$r_\Lambda \leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}} \leq \frac{3\nu}{\rho} \rho^{\frac{\tilde{\Lambda}}{4Ce(2\sigma a_{b,\nu})^{\frac{1}{\beta}}}} + 2Be^{\frac{-\tilde{\Lambda}^\beta}{\sigma}}$$

If $d > 0$, we define $\tilde{\Lambda}_c = \frac{\beta}{\beta+1} d \log \frac{1}{\rho} \left(\frac{\tilde{\Lambda}}{4Che} \right)^{\frac{\beta}{\beta+1}} \left(\frac{1}{2\sigma \log \frac{1}{\rho}} \right)^{\frac{1}{\beta+1}}$

Then, if $\tilde{\Lambda}_c \geq e$

$$\begin{aligned} r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} W(\tilde{\Lambda}_c)} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \\ &\leq \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \log \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \\ &= \frac{3\nu}{\rho} \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)^{\frac{-1}{d}} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \end{aligned}$$

Else, if $\tilde{\Lambda}_c < e$

$$\begin{aligned} r_\Lambda &\leq \frac{3\nu}{\rho} \rho^{h_1} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \\ &= \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} W(\tilde{\Lambda}_c)} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \\ &\leq \frac{3\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \frac{\tilde{\Lambda}_c}{e}} + 2Be^{-\frac{\tilde{\Lambda}^\beta}{\sigma}} \end{aligned}$$

Under Assumption 2(c) ($\Phi(c) = 0$ for all $c \geq a$)

If $d = 0$, we get

$$r_\Lambda \leq \frac{\nu}{\rho} \rho^{h_1} \leq \frac{\nu}{\rho} \rho^{\frac{\tilde{\Lambda}}{4Cae}}$$

If $d > 0$, we define $\tilde{\Lambda}_c = \frac{\tilde{\Lambda} d \log \frac{1}{\rho}}{4Cae}$

Then, if $\tilde{\Lambda}_c \geq e$

$$\begin{aligned} r_\Lambda &\leq \frac{\nu}{\rho} \rho^{h_1} \\ &= \frac{\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} W(\tilde{\Lambda}_c)} \\ &\leq \frac{\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \log \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)} \\ &= \frac{\nu}{\rho} \left(\frac{\tilde{\Lambda}_c}{\log \tilde{\Lambda}_c} \right)^{\frac{-1}{d}} \end{aligned}$$

Else, if $\tilde{\Lambda}_c < e$

$$\begin{aligned} r_\Lambda &\leq \frac{\nu}{\rho} \rho^{h_1} \\ &= \frac{\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}}} W(\tilde{\Lambda}_c) \\ &\leq \frac{\nu}{\rho} \rho^{\frac{1}{d \log \frac{1}{\rho}} \frac{\tilde{\Lambda}_c}{e}} \end{aligned}$$

□