



Hierarchical Optimistic Region Selection driven by Curiosity

Odalric-Ambrym Maillard

► To cite this version:

Odalric-Ambrym Maillard. Hierarchical Optimistic Region Selection driven by Curiosity. 2012. hal-00740418

HAL Id: hal-00740418

<https://hal.science/hal-00740418>

Preprint submitted on 25 Oct 2012

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.

Hierarchical Optimistic Region Selection driven by Curiosity

Odalric-Ambrym Maillard
Lehrstuhl für Informationstechnologie
Montanuniversität Leoben
Leoben, A-8700, Austria
odalricambrym.maillard@gmail.com

Abstract

This paper aims to take a step forwards making the term “intrinsic motivation” from reinforcement learning theoretically well founded, focusing on curiosity-driven learning. To that end, we consider the setting where, a fixed partition \mathcal{P} of a continuous space \mathcal{X} being given, and a process ν defined on \mathcal{X} being unknown, we are asked to sequentially decide which cell of the partition to select as well as where to sample ν in that cell, in order to minimize a loss function that is inspired from previous work on curiosity-driven learning. The loss on each cell consists of one term measuring a simple worst case quadratic sampling error, and a penalty term proportional to the range of the variance in that cell. The corresponding problem formulation extends the setting known as active learning for multi-armed bandits to the case when each arm is a continuous region, and we show how an adaptation of recent algorithms for that problem and of hierarchical optimistic sampling algorithms for optimization can be used in order to solve this problem. The resulting procedure, called Hierarchical Optimistic Region SElection driven by Curiosity (HORSE.C) is provided together with a finite-time regret analysis.

1 Introduction

In this paper, we focus on the setting of intrinsically motivated reinforcement learning (see Oudeyer and Kaplan [2007], Baranes and Oudeyer [2009], Schmidhuber [2010], Graziano et al. [2011]), which is an important emergent topic that proposes new difficult and interesting challenges for the theorist. Indeed, if some formal objective criteria have been proposed to implement specific notions of intrinsic rewards (see Jung et al. [2011], Martius et al. [2007]), so far, many - and only - experimental work has been carried out for this problem, often with interesting output (see Graziano et al. [2011], Mugan [2010], Konidaris [2011]) but unfortunately no performance guarantee validating a proposed approach. Thus proposing such an analysis may have great immediate consequences for validating some experimental studies.

Motivation. A typical example is the work of Baranes and Oudeyer [2009] about curiosity-driven learning (and later on Graziano et al. [2011], Mugan [2010], Konidaris [2011]), where a precise algorithm is defined together with an experimental study, yet no formal goal is defined, and no analysis is performed as well. They consider a so-called sensory-motor space $\mathcal{X} \stackrel{\text{def}}{=} \mathcal{S} \times \mathcal{M} \subset [0, 1]^d$ where \mathcal{S} is a (continuous) state space and \mathcal{M} is a (continuous) action space. There is no reward, yet one can consider that the goal is to actively select and sample subregions of \mathcal{X} for which a notion of “learning progress” - this intuitively measures the decay of some notion of error when successively sampling into one subregion - is maximal. Two key components are advocated in Baranes and Oudeyer [2009], in order to achieve successful results (despite that success is a fuzzy notion):

- The use of a *hierarchy* of regions, where each region is progressively split into sub-regions.

- Splitting leaf-regions in two based on the optimization of the dissimilarity, amongst the regions, of the learning progress. The idea is to identify regions with a learning complexity that is a globally constant in that region, which also provides better justification for allocating samples between identified regions.

We believe it is possible to go one step towards a full performance analysis of such algorithms, by relating the corresponding active sampling problem to existing frameworks.

Contribution. This paper aims to take a step forwards making the term “intrinsic motivation” from reinforcement learning theoretically well founded, focusing on curiosity-driven learning. We introduce a mathematical framework in which a metric space (which intuitively plays the role of the state-action space) is divided into regions and a learner has to sample from an unknown random function in a way that reduces a notion of error measure the most. This error consists of two terms, the first one is a robust measure of the quadratic error between the observed samples and their unknown mean, the second one penalizes regions with non constant learning complexity, thus enforcing the notion of curiosity. The paper focuses on how to choose the region to sample from, when a partition of the space is provided.

The resulting problem formulation can be seen as a non trivial extension of the setting of active learning in multi-armed bandits (see Carpentier et al. [2011] or Antos et al. [2010]), where the main idea is to estimate the variance of each arm and sample proportionally to it, to the case when each arm is a region as opposed to a point. In order to deal with this difficulty, the maximal and minimal variance inside each region is tracked by means of a hierarchical optimization procedure, in the spirit of the HOO algorithm from Bubeck et al. [2011]. This leads to a new procedure called Hierarchical Optimistic Region SElection driven by Curiosity (HORSE.C) for which we provide a theoretical performance analysis.

Outline. The outline of the paper is the following. In Section 2 we introduce the precise setting and define the objective function. Section 3 defines our assumptions. Then in Section 4 we present the HORSE.C algorithm. Finally in Section 5, we provide the main Theorem 1 that gives performance guarantees for the proposed algorithm.

2 Setting: Robust region allocation with curiosity-inducing penalty.

Let \mathcal{X} assumed to be a metric space and let $\mathcal{Y} \subset \mathbb{R}^d$ be a normed space, equipped with the Euclidean norm $\|\cdot\|$. We consider an unknown \mathcal{Y} -valued process defined on \mathcal{X} , written $\nu : \mathcal{X} \rightarrow \mathfrak{M}_1^+(\mathcal{Y})$, where $\mathfrak{M}_1^+(\mathcal{Y})$ refers to the set of all probability measures on \mathcal{Y} , such that for all $x \in \mathcal{X}$, the random variable $Y \sim \nu(x)$ has mean $\mu(x) \in \mathbb{R}^d$ and covariance matrix $\Sigma(x) \in \mathcal{M}_{d,d}(\mathbb{R})$ assumed to be diagonal. We thus introduce for convenience the notation $\rho(x) \stackrel{\text{def}}{=} \text{trace}(\Sigma(x))$, where trace is the trace operator (this corresponds to the variance in dimension 1). We call \mathcal{X} the input space or sampling space, and \mathcal{Y} the output space or value space.

Intuition Intuitively when applied to the setting of Baranes and Oudeyer [2009], then $\mathcal{X} \stackrel{\text{def}}{=} \mathcal{S} \times \mathcal{A}$ is the space of state-action pairs, where \mathcal{S} is a continuous state space and \mathcal{A} a continuous action space, ν is the transition kernel of an unknown MDP, and finally $\mathcal{Y} \stackrel{\text{def}}{=} \mathbb{R}$. This is the reason why we consider $\mathcal{Y} \subset \mathbb{R}^d$ and not only $\mathcal{Y} \subset \mathbb{R}$ as would seem more natural. One difference is that we assume (see Section 3) that we can sample anywhere in \mathcal{X} , which is a restrictive yet common assumption in the reinforcement learning literature. How to get rid of this assumption is an open and challenging question that is left for future work.

Sampling error and robustness Let us consider a sequential sampling process on \mathcal{X} , i.e. a process that samples at time t a value $Y_t \sim \nu(X_t)$ at point X_t , where $X_t \in \mathcal{F}_{<t}$ is a measurable function of the past inputs and outputs $\{(X_s, Y_s)\}_{s < t}$. It is natural to look at the following quantity, that we call average noise vector η_t :

$$\eta_t \stackrel{\text{def}}{=} \frac{1}{t} \sum_{s=1}^t Y_s - \mu(X_s) \in \mathbb{R}^d.$$

One interesting property is that this is a martingale difference sequence, which means that the norm of this vector enjoys a concentration property. More precisely (see Lemma 1 in the appendix), we have for all deterministic $t > 0$

$$\mathbb{E}[\|\eta_t\|^2] = \frac{1}{t} \mathbb{E}\left[\frac{1}{t} \sum_{s=1}^t \rho(X_s)\right].$$

A similar property holds for a region $R \subset \mathcal{X}$ that has been sampled $n_t(R)$ times, and in order to be robust against a bad sampling strategy inside a region, it is natural to look at the worst case error, that we define as

$$\mathbf{e}_R(n_t) \stackrel{\text{def}}{=} \frac{\sup_{x \in R} \rho(x)}{n_t(R)}.$$

One reason for looking at robustness is that for instance, in the case we work with an MDP, we are generally not completely free to choose the sample $X_s \in \mathcal{S} \times \mathcal{A}$: we can only choose the action and the next state is generally given by Nature. Thus, it is important to be able to estimate this worst case error so as to prevent from bad situations.

Goal Now let \mathcal{P} be a fixed, known partition of the space \mathcal{X} and consider the following game. The goal of an algorithm is, at each time step t , to propose one point x_t where to sample the space \mathcal{X} , so that its allocation of samples $\{n_t(R)\}_{R \in \mathcal{P}}$ (that is, the number of points sampled in each region) minimizes some objective function. Thus, the algorithm is free to sample everywhere in each region, with the goal that the total number of points chosen in each region is optimal in some sense. A simple candidate for this objective function would be the following

$$L_{\mathcal{P}}(n_t) \stackrel{\text{def}}{=} \max \left\{ \mathbf{e}_R(n_t); R \in \mathcal{P} \right\},$$

however, in order to incorporate a notion of curiosity, we would also like to penalize regions that have a variance term ρ that is non homogeneous (i.e. the less homogeneous, the more samples we allocate). Indeed, if a region has constant variance, then we do not really need to understand more its internal structure, and thus its better to focus on an other region that has very heterogeneous variance. For instance, one would like to split such a region in several homogeneous parts, which is essentially the idea behind section C.3 of Baranes and Oudeyer [2009]. We thus add a curiosity-penalization term to the previous objective function, which leads us to define the pseudo-loss of an allocation $n_t \stackrel{\text{def}}{=} \{n_t(R)\}_{R \in \mathcal{P}}$ in the following way:

$$\mathfrak{L}_{\mathcal{P}}(n_t) \stackrel{\text{def}}{=} \max \left\{ \mathbf{e}_R(n_t) + \lambda |R| (\max_{x \in R} \rho(x) - \min_{x \in R} \rho(x)); R \in \mathcal{P} \right\}. \quad (1)$$

Indeed, this means that we do not want to focus just on regions with high variance, but also trade-off with highly heterogeneous regions, which is coherent with the notion of curiosity (see Oudeyer and Kaplan [2007]). For convenience, we also define the pseudo-loss of a region R by

$$\mathcal{L}_R(n_t) \stackrel{\text{def}}{=} \mathbf{e}_R(n_t) + \lambda |R| (\max_{x \in R} \rho(x) - \min_{x \in R} \rho(x)).$$

Regret The regret (or loss) of an allocation algorithm at time T is defined as the difference between the cumulated pseudo-loss of the allocations $n_t = \{n_{R,t}\}_{R \in \mathcal{P}}$ proposed by the algorithm and that of the best allocation strategy $n_t^* = \{n_{R,t}^*\}_{R \in \mathcal{P}}$ at each time steps; we define

$$\mathfrak{R}_T \stackrel{\text{def}}{=} \sum_{t=|\mathcal{P}|}^T \mathfrak{L}_{\mathcal{P}}(n_t) - \mathfrak{L}_{\mathcal{P}}(n_t^*),$$

where an optimal allocation at time t is defined by

$$n_t^* \in \operatorname{argmin} \left\{ \mathfrak{L}_{\mathcal{P}}(n_t); \{n_t(R)\}_{R \in \mathcal{P}} \text{ is such that } \sum_{R \in \mathcal{P}} n_t(R) = t \right\}.$$

Note that the sum starts at $t = |\mathcal{P}|$ for a technical reason, since for $t < |\mathcal{P}|$, whatever the allocation, there is always at least one region with no sample, and thus $\mathfrak{L}_{\mathcal{P}}(n_t) = \infty$.

Example 1 In the special case when $\mathcal{X} = \{1, \dots, K\}$ is finite with $K \ll T$, and when \mathcal{P} is the complete partition (each cell corresponds to exactly one point), the penalization term is canceled. Thus the problem reduces to the choice of the quantities $n_t(i)$ for each arm i , and the loss of an allocation simply becomes

$$\mathfrak{L}(n_t) \stackrel{\text{def}}{=} \max \left\{ \frac{\rho(i)}{n_t(i)}; 1 \leq i \leq K \right\}.$$

This almost corresponds to the already challenging setting analyzed for instance in Carpentier et al. [2011] or Antos et al. [2010]. The difference is that we are interested in the *cumulative* regret of our allocation instead of only the regret suffered for the last round as considered in Carpentier et al. [2011] or Antos et al. [2010]. Also we directly target $\frac{\rho(i)}{n_t(i)}$ whereas they consider the mean sampling error (but both terms are actually of the same order). Thus the setting we consider can be seen as a generalization of these works to the case when each arm corresponds to a continuous sampling domain.

3 Assumptions

In this section, we introduce some mild assumptions. We essentially assume that the unknown distribution is such that it has a sub-Gaussian noise, and a smooth mean and variance functions. These are actually very mild assumptions. Concerning the algorithm, we consider it can use a partition tree of the space, and that this one is essentially not degenerated (a typical binary tree that satisfies all the following assumptions is such that each cell is split in two children of equal volume). Such assumptions on trees have been extensively discussed for instance in Bubeck et al. [2011].

Sampling At any time, we assume that we are able to sample at any point in \mathcal{X} , i.e. we assume we have a generative model¹ of the unknown distribution ν .

Unknown distribution We assume that ν is sub-Gaussian, meaning that for all fixed $x \in \mathcal{X}$

$$\forall \lambda \in \mathbb{R}^d \quad \ln \mathbb{E} \exp[\langle \lambda, Y - \mu(X) \rangle] \leq \frac{\lambda^T \Sigma(x) \lambda}{2},$$

and has diagonal covariance matrix in each point².

The function μ is assumed to be Lipschitz w.r.t a metric ℓ_1 , i.e. it satisfies

$$\forall x, x' \in \mathcal{X} \quad \|\mu(x) - \mu(x')\| \leq \ell_1(x, x').$$

Similarly, the function ρ is assumed to be Lipschitz w.r.t a metric ℓ_2 , i.e. it satisfies

$$\forall x, x' \in \mathcal{X} \quad |\rho(x) - \rho(x')| \leq \ell_2(x, x').$$

Hierarchy We assume that \mathcal{Y} is a convex and compact subset of $[0, 1]^d$. We consider an infinite binary tree \mathcal{T} whose nodes correspond to regions of \mathcal{X} . A node is indexed by a pair (h, i) , where $h \geq 0$ is the depth of the nodes in \mathcal{T} and $0 \leq i < 2^h$ is the position of the node at depth h . We write $R(h, i) \subset \mathcal{X}$ the region associated with node (h, i) . The regions are fixed in advance, are all assumed to be measurable with positive measure, and must satisfy that for each $h \geq 1$, $\{R(h, i)\}_{0 \leq i < 2^h}$ is a partition of \mathcal{X} that is compatible with depth $h - 1$, where $R(0, 0) \stackrel{\text{def}}{=} \mathcal{X}$; in particular for all $h \geq 0$, for all $0 \leq i < 2^h$, then

$$R(h, i) = R(h + 1, 2i) \cup R(h + 1, 2i + 1).$$

In dimension d , a standard way to define such a tree is to split each parent node in half along the largest side of the corresponding hyper-rectangle, see Bubeck et al. [2011] for details.

For a finite sub-tree \mathcal{T}_t of \mathcal{T} , we write $\text{Leaf}(\mathcal{T}_t)$ for the set of all leaves of \mathcal{T}_t . For a region $(h, i) \in \mathcal{T}_t$, we denote by $\mathcal{C}_t(h, i)$ the set of its children in \mathcal{T}_t , and by $\mathcal{T}_t(h, i)$ the subtree of \mathcal{T}_t starting with root node (h, i) .

Algorithm and partition The partition \mathcal{P} is assumed to be such that each of its regions R corresponds to one region $R(h, i) \in \mathcal{T}$; equivalently, there exists a finite sub-tree $\mathcal{T}_0 \subset \mathcal{T}$ such that $\text{Leaf}(\mathcal{T}_0) = \mathcal{P}$. An algorithm is only allowed to expand one node of \mathcal{T}_t at each time step t . In the sequel, we write indifferently $\mathcal{P} \in \mathcal{T}$ and $(h, i) \in \mathcal{T}$ or \mathcal{P} and $R(h, i) \subset \mathcal{X}$ to refer to the partition or one of its cell.

Exponential decays Finally, we assume that the ℓ_1 and ℓ_2 diameters of the region $R(h, i)$ as well as its volume $|R(h, i)|$ decay at exponential rate in the sense that there exists positive constants $\gamma, \gamma_1, \gamma_2$ and c, c_1, c_2 such that for all $h \geq 0$, then $|R(h, i)| \leq c\gamma^h$,

$$\max_{x', x \in R(h, i)} \ell_1(x, x') \leq c_1 \gamma_1^h \quad \text{and} \quad \max_{x', x \in R(h, i)} \ell_2(x, x') \leq c_2 \gamma_2^h.$$

Similarly, we assume that there exists positive constants $c' \leq c, c'_1 \leq c_1$ and $c'_2 \leq c_2$ such that for all $h \geq 0$, then $|R(h, i)| \geq c'\gamma^h$,

$$\max_{x', x \in R(h, i)} \ell_1(x, x') \geq c'_1 \gamma_1^h \quad \text{and} \quad \max_{x', x \in R(h, i)} \ell_2(x, x') \geq c'_2 \gamma_2^h.$$

This assumption is made to avoid degenerate trees and for general purpose only. It actually holds for any reasonable binary tree.

¹using the standard terminology in Reinforcement Learning.

²this assumption is only here to make calculations easier and avoid nasty technical considerations that anyway do not affect the order of the final regret bound but only concern second order terms.

4 Allocation algorithm

In this section, we now introduce the main algorithm of this paper in order to solve the problem considered in Section 2. It is called Hierarchical Optimistic Region SElection driven by Curiosity. Before proceeding, we need to define some quantities.

4.1 High-probability upper-bound and lower-bound estimations

Let us consider the following (biased) estimator

$$\hat{\sigma}_t^2(R) \stackrel{\text{def}}{=} \frac{1}{N_t(R)} \sum_{s=1}^t \|Y_s\|^2 \mathbb{I}\{X_s \in R\} - \left\| \frac{1}{N_t(R)} \sum_{s=1}^t Y_s \mathbb{I}\{X_s \in R\} \right\|^2.$$

Apart from a small multiplicative biased by a factor $\frac{N_t(R)-1}{N_t(R)}$, it has more importantly a positive bias due to the fact that the random variables do not share the same mean; this phenomenon is the same as the estimation of the average variance for independent but non i.i.d variables with different means $\{\mu_i\}_{i \leq n}$, where the bias would be given by $\frac{1}{n} \sum_{i=1}^n [\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j]^2$ (see Lemma 5). In our case, it is thus always non negative, and under the assumption that μ is Lipschitz w.r.t the metric ℓ_1 , it is fortunately bounded by $d_1(R)^2$, the diameter of R w.r.t the metric ℓ_1 .

We then introduce the two following key quantities, defined for all $x \in R$ and $\delta \in [0, 1]$ by

$$U_t(R, x, \delta) \stackrel{\text{def}}{=} \hat{\sigma}_t^2(R) + (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2N_t(R)}} + \frac{d \ln(2d/\delta)}{2N_t(R)} + \frac{1}{N_t(R)} \sum_{s=1}^t \ell_2(X_s, x) \mathbb{I}\{X_s \in R\},$$

$$L_t(R, x, \delta) \stackrel{\text{def}}{=} \hat{\sigma}_t^2(R) - (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2N_t(R)}} - d_1(R)^2 - \frac{1}{N_t(R)} \sum_{s=1}^t \ell_2(X_s, x) \mathbb{I}\{X_s \in R\}.$$

Note that we would have preferred to replace the terms involving $\ln(2d/\delta)$ with a term depending on the empirical variance, in the spirit of Carpentier et al. [2011] or Antos et al. [2010]. However, contrary to the estimation of the mean, extending the standard results valid for i.i.d data to the case of a *martingale difference sequence* is non trivial for the estimation of the variance, especially due to the additive bias resulting from the fact that the variables may not share the same mean, but also to the absence of such results for U-statistics (up to the author's knowledge). For that reason such an extension is left for future work.

The following results (the proof of which is provided in Appendix A.3) show that $U_t(R, x, \delta)$ is a high probability upper bound on $\rho(x)$ while $L_t(R, x, \delta)$ is a high probability lower bound on $\rho(x)$.

Proposition 1 *Under the assumptions that \mathcal{Y} is a convex subset of $[0, 1]^d$, ν is sub-Gaussian, ρ is Lipschitz w.r.t. ℓ_2 and $R \subset \mathcal{X}$ is compact and convex, then*

$$\mathbb{P}(\forall x \in \mathcal{X}; U_t(R, x, \delta) \leq \rho(x)) \leq t\delta.$$

Similarly, under the same assumptions, then

$$\mathbb{P}(\forall x \in \mathcal{X}; L_t(R, x, \delta) \leq \rho(x) - b(x, R, N_t(R), \delta)) \leq t\delta,$$

where we introduced for convenience the quantity

$$b(x, R, n, \delta) \stackrel{\text{def}}{=} 2 \max_{x' \in R} \ell_2(x, x') + d_1(R)^2 + 2(1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n}.$$

Now on the other other hand, we have that following inequalities (see the proof in Appendix A.3.)

Proposition 2 *Under the assumptions that \mathcal{Y} is a convex subset of $[0, 1]^d$, ν is sub-Gaussian, μ is Lipschitz w.r.t. ℓ_1 , ρ is Lipschitz w.r.t. ℓ_2 and $R \subset \mathcal{X}$ is compact and convex, then*

$$\mathbb{P}(\forall x \in \mathcal{X}; U_t(R, x, \delta) \geq \rho(x) + b(x, R, N_t(R), \delta)) \leq t\delta.$$

Similarly, under the same assumptions, then

$$\mathbb{P}(\forall x \in \mathcal{X}; L_t(R, x, \delta) \geq \rho(x)) \leq t\delta.$$

4.2 Hierarchical Optimistic Region SElection driven by Curiosity (HORSE.C).

The pseudo-code of the HORSE.C algorithm is presented in Figure 1 below. This algorithm relies on the estimation of the quantities $\max_{x \in R} \rho(x)$ and $\min_{x \in R} \rho(x)$ in order to define which point X_{t+1} to sample at time $t + 1$. It is chosen by expanding a leaf of a hierarchical tree $\mathcal{T}_t \subset \mathcal{T}$, in an optimistic way, starting with a tree \mathcal{T}_0 with leaves corresponding to the partition \mathcal{P} .

The intuition is the following: let us consider a node (h, i) of the tree \mathcal{T}_t expanded by the algorithm at time t . The maximum value of ρ in $R(h, i)$ is thus achieved for one of its children node $(h', i') \in \mathcal{C}_t(h, i)$. Thus if we have computed an upper bound on the maximal value of ρ in each child, then we have an upper bound on the maximum value of ρ in $R(h, i)$. Proceeding in a similar way for the lower bound, this motivates the following two recursive definitions:

$$\hat{\rho}_t^+(h, i; \delta) \stackrel{\text{def}}{=} \min \left\{ \max_{x \in R(h, i)} U_t(R(h, i), x, \delta), \max \{ \hat{\rho}_t^+(h', i'; \delta); (h', i') \in \mathcal{C}_t(h, i) \} \right\},$$

$$\hat{\rho}_t^-(h, i; \delta) \stackrel{\text{def}}{=} \max \left\{ \min_{x \in R(h, i)} L_t(R(h, i), x, \delta), \min \{ \hat{\rho}_t^-(h', i'; \delta); (h', i') \in \mathcal{C}_t(h, i) \} \right\}.$$

These values are used in order to build an optimistic estimate of the quantity $\mathcal{L}_{R(h, i)}(N_t)$ in region (h, i) (step 4), and then to select in which cell of the partition we should sample (step 5). Then the algorithm chooses where to sample in the selected region so as to improve the estimations of $\hat{\rho}_t^+$ and $\hat{\rho}_t^-$. This is done by alternating (step 6.) between expanding a leaf following a path that is optimistic according to $\hat{\rho}_t^+$ (step 7,8,9), or according to $\hat{\rho}_t^-$ (step 11.)

Thus, at a high level, the algorithm performs on each cell $(h, i) \in \mathcal{P}$ of the given partition two hierarchical searches, one for the maximum value of ρ in region $R(h, i)$ and one for its minimal value. This can be seen as an adaptation of the algorithm HOO from Bubeck et al. [2011] with the main difference that we target the variance and not just the mean (this is more difficult). On the other hand, there is a strong link between step 5, where we decide to allocate samples between regions $\{R(h, i)\}_{(h, i) \in \mathcal{P}}$, and the CH-AS algorithm from Carpentier et al. [2011].

5 Performance analysis of the HORSE.C algorithm

In this section, we are now ready to provide the main theorem of this paper, i.e. a regret bound on the performance of the HORSE.C algorithm, which is the main contribution of this work. To this end, we make use of the notion of near-optimality dimension, introduced in Bubeck et al. [2011], and that measures a notion of intrinsic dimension of the maximization problem.

Definition (Near optimality dimension) For $c > 0$, the c -optimality dimension of ρ restricted to the region R with respect to the pseudo-metric ℓ_2 is defined as

$$\max \left\{ \limsup_{\epsilon \rightarrow 0} \frac{\ln(\mathcal{N}(R_{c\epsilon}, \ell_2, \epsilon))}{\ln(\epsilon^{-1})}, 0 \right\} \text{ where } R_{c\epsilon} \stackrel{\text{def}}{=} \left\{ x \in R; \rho(x) \geq \max_{x \in R} \rho(x) - \epsilon \right\},$$

and where $\mathcal{N}(R_{c\epsilon}, \ell_2, \epsilon)$ is the ϵ -packing number of the region $R_{c\epsilon}$.

Let $d^+(h_0, i_0)$ be the c -optimality dimension of ρ restricted to the region $R(h_0, i_0)$ (see e.g. Bubeck et al. [2011]), with the constant $c \stackrel{\text{def}}{=} 4(2c_2 + c_1^2)/c_2'$. Similarly, let $d^-(h_0, i_0)$ be the c -optimality dimension of $-\rho$ restricted to the region $R(h_0, i_0)$. Let us finally define the biggest near-optimality dimension of ρ over each cell of the partition \mathcal{P} to be

$$d_\rho \stackrel{\text{def}}{=} \max \left\{ \max \{ d^+(h_0, i_0), d^-(h_0, i_0) \}; (h_0, i_0) \in \mathcal{P} \right\}.$$

Theorem 1 (Regret bound for HORSE.C) *Under the assumptions of Section 3 and if moreover $\gamma_1^2 \leq \gamma_2$, then for all $\delta \in [0, 1]$, the regret of the Hierarchical Optimistic Region SElection driven by Curiosity procedure parameterized with δ is bounded with probability higher than $1 - 2\delta$ as follows.*

$$\mathfrak{R}_T \leq \sum_{t=|P|}^T \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{1}{n_t^*(h_0, i_0)} + 2\lambda c \gamma^{h_0} \right) B(h_0, n_t^*(h_0, i_0), \delta_t),$$

Algorithm 1 The HORSE.C algorithm.

Require: An infinite binary tree \mathcal{T} , a partition $\mathcal{P} \subset \mathcal{T}$, $\delta \in [0, 1]$, $\lambda \geq 0$

1: Let \mathcal{T}_0 be such that $Leaf(\mathcal{T}_0) = \mathcal{P}$, and $\delta_{i,t} = \frac{6\delta}{\pi^2 i^2 (2t+1)|\mathcal{P}|t^3}$.

2: Let $t := 0$.

3: **while true do**

4: define for each region $(h, i) \in \mathcal{T}_t$ the estimated loss

$$\hat{\mathcal{L}}_t(h, i) \stackrel{\text{def}}{=} \frac{\hat{\rho}_t^+(h, i; \delta)}{N_t(R(h, i))} + \lambda |R(h, i)| (\hat{\rho}_t^+(h, i; \delta) - \hat{\rho}_t^-(h, i; \delta)),$$

where $\delta = \delta_{N_t(R(h, i)), t}$, where by convention $\hat{\mathcal{L}}_t(h, i)$ if it is undefined.

5: choose the next region of the current partition $\mathcal{P} \subset \mathcal{T}$ to sample

$$(H_{t+1}, I_{t+1}) \stackrel{\text{def}}{=} \operatorname{argmax} \left\{ \hat{\mathcal{L}}_t(h, i); (h, i) \in \mathcal{P} \right\}.$$

6: **if** $N_t(R(H_{t+1}, I_{t+1})) = n$ is odd **then**

7: sequentially select a path of children of (H_{t+1}, I_{t+1}) in \mathcal{T}_t defined by $(H_{t+1}^0, I_{t+1}^0) \stackrel{\text{def}}{=} (H_{t+1}, I_{t+1})$, and

$$(H_{t+1}^{j+1}, I_{t+1}^{j+1}) \stackrel{\text{def}}{=} \operatorname{argmax} \left\{ \hat{\rho}_t^+(h, i; \delta_{n,t}); (h, i) \in \mathcal{C}_t(H_{t+1}^j, I_{t+1}^j) \right\},$$

until $j = j_{t+1}$ is such that $(H_{t+1}^{j_{t+1}}, I_{t+1}^{j_{t+1}}) \in Leaf(\mathcal{T}_t)$.

8: expand the node $(H_{t+1}^{j_{t+1}}, I_{t+1}^{j_{t+1}})$ in order to define \mathcal{T}_{t+1} and then define the candidate child

$$(h_{t+1}, i_{t+1}) \stackrel{\text{def}}{=} \operatorname{argmax} \left\{ \hat{\rho}_t^+(h, i; \delta_{n,t}); (h, i) \in \mathcal{C}_{t+1}(H_{t+1}^{j_{t+1}}, I_{t+1}^{j_{t+1}}) \right\}.$$

9: sample at point X_{t+1} and receive the value $Y_{t+1} \sim \nu(X_{t+1})$, where

$$X_{t+1} \stackrel{\text{def}}{=} \operatorname{argmax} \left\{ U_t(R(h_{t+1}, i_{t+1}), x, \delta_{n,t}); x \in R(h_{t+1}, i_{t+1}) \right\},$$

10: **else**

11: proceed similarly than steps 6,7,8 with $\hat{\rho}_t^+$ replaced with $\hat{\rho}_t^-$.

12: **end if**

13: $t := t + 1$.

14: **end while**

where δ_t is a shorthand notation for the quantity $\delta_{n_t^*(h_0, i_0), t-1}$, where $n_t^*(h_0, i_0)$ is the optimal allocation at round t for the region $(h_0, i_0) \in \mathcal{P}$ and where

$$B(h_0, k, \delta_{k,t}) \stackrel{\text{def}}{=} \min_{h_0 \leq h} \left\{ 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N_{h_0}(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N_{h_0}(h, k)} \right\},$$

in which we have used the following quantity

$$N_{h_0}(h, k) \stackrel{\text{def}}{=} \frac{1}{C(c_2'\gamma_2^h)^{-d_\rho}} \left(k - 2^{h-h_0} [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2} \right).$$

Note that the assumption $\gamma_1^2 \leq \gamma_2$ is only here so that d_ρ can be defined w.r.t the metric ℓ_2 only. We can remove it at the price of using instead a metric mixing ℓ_1 and ℓ_2 together and of much more technical considerations. Similarly, we could have expressed the result using the local values $d^+(h_0, i_0)$ instead of the less precise d_ρ (neither those, nor d_ρ need to be known by the algorithm). The full proof of this theorem is reported in the appendix. The main steps of the proof are as follows. First we provide upper and lower confidence bounds for the estimation of the quantities $U_t(R, x, \delta)$ and $L_t(R, x, \delta)$. Then, we lower-bound the depth of the subtree of each region $(h_0, i_0) \in \mathcal{P}$ that contains a maximal point $\operatorname{argmax}_{x \in R(h_0, i_0)} \rho(x)$, and proceed similarly for a minimal point. This uses the near-optimality dimension of ρ and $-\rho$ in the region $R(h_0, i_0)$, and enables to provide an upper bound on $\hat{\rho}_t^+(h, i; \delta)$ as well as a lower bound on $\hat{\rho}_t^-(h, i; \delta)$. This then enables us to deduce

bounds relating the estimated loss $\hat{\mathcal{L}}_t(h, i)$ to the true loss $\mathfrak{L}_{R(h, i)}(N_t)$. Finally, we relate the true loss of the current allocation to the one using the optimal one $n_{t+1}^*(h_0, i_0)$ by discussing whether a region has been over or under sampled. This final part is closed in spirit to the proof of the regret bound for CH-AS in Carpentier et al. [2011].

In order to better understand the gain in Theorem 1, we provide the following corollary that gives more insights about the order of magnitude of the regret.

Corollary 1 *Let $\beta \stackrel{\text{def}}{=} 1 + \ln(\max\{2, \gamma_2^{-d_\rho}\})$. Under the assumptions of Theorem 1, assuming that the partition \mathcal{P} of the space \mathcal{X} is well behaved, i.e. that for all $(h_0, i_0) \in \mathcal{P}$, then $n_{t+1}^*(h_0, i_0)$ grows at least at speed $O\left(\ln(t)\left(\frac{1}{\gamma_2}\right)^{2h_0\beta}\right)$, then for all $\delta \in [0, 1]$, with probability higher than $1 - 2\delta$ we have*

$$\mathfrak{R}_T = O\left(\sum_{t=|\mathcal{P}|}^T \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{1}{n_t^*(h_0, i_0)} + 2\lambda c \gamma^{h_0}\right) \left(\frac{\ln(t)}{n_t^*(h_0, i_0)}\right)^{\frac{1}{2\beta}}\right).$$

This regret term has to be compared with the typical range of the cumulative loss of the optimal allocation strategy, that is given by

$$\sum_{t=|\mathcal{P}|}^T \mathfrak{L}_{\mathcal{P}}(n_t^*) = \sum_{t=|\mathcal{P}|}^T \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{\rho_{(h_0, i_0)}^+}{n_t^*(h_0, i_0)} + 2\lambda c \gamma^{h_0} (\rho_{(h_0, i_0)}^+ - \rho_{(h_0, i_0)}^-)\right),$$

where $\rho_{(h_0, i_0)}^+ \stackrel{\text{def}}{=} \max_{x \in R(h_0, i_0)} \rho(x)$, and similarly $\rho_{(h_0, i_0)}^- \stackrel{\text{def}}{=} \min_{x \in R(h_0, i_0)} \rho(x)$. Thus, this shows that, after normalization, the relative regret on each cell (h_0, i_0) is roughly of order $\frac{1}{\rho^+(h_0, i_0)} \left(\frac{\ln(t)}{n_t^*(h_0, i_0)}\right)^{\frac{1}{2\beta}}$, i.e. decays at speed $n_t^*(h_0, i_0)^{-\frac{1}{2\beta}}$. This shows that we are not only able to compete with the performance of the best allocation strategy, but we actually achieve the exact same performance with multiplicative factor 1, up to a second order term. Note also that, when specified to the case of Example 1, the order of this regret is competitive with the standard results from Carpentier et al. [2011].

The loss of the variance term $\rho^+(h_0, i_0)^{-1}$ (that is actually a constant) here comes from the fact that we are only able to use Hoeffding's like bounds for the estimation of the variance. In order to remove it, one would need empirical Bernstein's bounds for variance estimation in the case of martingale difference sequences. This is postponed to future work.

6 Discussion

In this paper, we have provided an algorithm together with a regret analysis for a problem of online allocation of samples in a fixed partition, where the objective is to minimize a loss that contains a penalty term that is driven by a notion of curiosity. A very specific case (finite state space) already corresponds to a difficult question known as *active learning in the multi-armed bandit setting* and has been previously addressed in the literature (e.g. Antos et al. [2010], Carpentier et al. [2011]). We have considered an extension of this problem to a continuous domain where a fixed partition of the space as well as a generative model of the unknown dynamic are given, using our curiosity-driven loss function as a measure of performance. Our main result is a regret bound for that problem, that shows that our procedure is first order optimal, i.e. achieves the same performance as the best possible allocation (thus with multiplicative constant 1).

We believe this result contributes to filling the important gap that exists between existing algorithms for the challenging setting of intrinsic reinforcement learning and a theoretical analysis of such, the HORSE.C algorithm being related in spirit to, yet simpler and less ambitious the RIAC algorithm from Baranes and Oudeyer [2009]. Indeed, in order to achieve the objective that tries to address RIAC, one should first remove the assumption that the partition is given: One trivial solution is to run the HORSE.C algorithm in episodes of doubling length, starting with the trivial partition, and to select at the end of each a possibly better partition based on computed confidence intervals, however making efficient use of previous samples and avoiding a blow-up of candidate partitions happen to be a challenging question; then one should relax the generative model assumption (i.e. that we can sample wherever we want), a question that shares links with a problem called *autonomous exploration*. Thus, even if the regret analysis of the HORSE.C algorithm is already a strong, new result that is interesting independently of such difficult specific goals and of the reinforcement learning framework (no MDP structure is required), those questions are naturally left for future work.

References

- Andr s Antos, Varun Grover, and Csaba Szepesv ri. Active learning in heteroscedastic noise. *Theoretical Computer Science*, 411(29-30):2712–2728, 2010.
- A. Baranes and P.-Y. Oudeyer. R-IAC: Robust Intrinsically Motivated Exploration and Active Learning. *IEEE Transactions on Autonomous Mental Development*, 1(3):155–169, October 2009.
- S bastien Bubeck, R mi Munos, Gilles Stoltz, and Csaba Szepesv ri. X-armed bandits. *Journal of Machine Learning Research*, 12:1655–1695, 2011.
- Alexandra Carpentier, Alessandro Lazaric, Mohammad Ghavamzadeh, R mi Munos, and Peter Auer. Upper-confidence-bound algorithms for active learning in multi-armed bandits. In Jyrki Kivinen, Csaba Szepesv ri, Esko Ukkonen, and Thomas Zeugmann, editors, *Algorithmic Learning Theory*, volume 6925 of *Lecture Notes in Computer Science*, pages 189–203. Springer Berlin / Heidelberg, 2011.
- Vincent Graziano, Tobias Glasmachers, Tom Schaul, Leo Pape, Giuseppe Cuccu, J. Leitner, and J. Schmidhuber. Artificial Curiosity for Autonomous Space Exploration. *Acta Futura (in press)*, (1), 2011.
- Tobias Jung, Daniel Polani, and Peter Stone. Empowerment for continuous agent-environment systems. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 19(1): 16–39, 2011.
- G.D. Konidaris. *Autonomous robot skill acquisition*. PhD thesis, University of Massachusetts Amherst, 2011.
- Georg Martius, J. Michael Herrmann, and Ralf Der. Guided self-organisation for autonomous robot development. In *Proceedings of the 9th European conference on Advances in artificial life*, ECAL’07, pages 766–775, Berlin, Heidelberg, 2007. Springer-Verlag.
- Jonathan Mugan. *Autonomous Qualitative Learning of Distinctions and Actions in a Developing Agent*. PhD thesis, University of Texas at Austin, 2010.
- Pierre-Yves Oudeyer and Frederic Kaplan. What is Intrinsic Motivation? A Typology of Computational Approaches. *Frontiers in neurorobotics*, 1(November):6, January 2007.
- J. Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (19902010). *Autonomous Mental Development, IEEE Transactions on*, 2(3):230–247, 2010.

A Technical details: Concentration results

In the sequel, we use the notation y_k in order to refer to the k -th component of vector y . The norm of some element $y \in \mathcal{Y}$ is denoted by $\|y\| \stackrel{\text{def}}{=} \left(\sum_{k=1}^d y_k^2 \right)^{1/2}$.

A.1 Expectations

Lemma 1 *Let us consider a process on \mathcal{X} that samples at time t a value $Y_t \sim \nu(X_t)$ at point X_t , where $X_t \in \mathcal{F}_{< t}$ is a measurable function of the past inputs and outputs $\{(X_s, Y_s)\}_{s < t}$. Then the noise vector η_t defined by*

$$\eta_t \stackrel{\text{def}}{=} \frac{1}{t} \sum_{s=1}^t Y_s - \mu(X_s) \in \mathbb{R}^d,$$

satisfies the following equality for all deterministic $t > 0$

$$\mathbb{E}[\|\eta_t\|^2] = \frac{1}{t} \mathbb{E}\left[\frac{1}{t} \sum_{s=1}^t \rho(X_s)\right].$$

Proof: For a vector $y \in \mathbb{R}^d$, let us write $y_k \in \mathbb{R}$ is k -th component. With this notation, we thus have the following rewriting,

$$\|\eta_t\|^2 = \frac{1}{t^2} \sum_{k=1}^d \sum_{s=1}^t (Y_{k,s} - \mu_k(x_s))^2 + \frac{2}{t^2} \sum_{k=1}^d \sum_{s=1}^t \sum_{s'=1}^{s-1} (Y_{k,s} - \mu_k(x_s))(Y_{k,s'} - \mu_k(x_{s'})),$$

from which we deduce the result immediately by taking expectation, using the assumption that $\Sigma(x)$ is diagonal and using the martingale structure. □

A.2 Mean estimation

Let us first provide the following results for independent random variables, in order to provide more intuition about the more involved case of a martingale difference sequence.

Lemma 2 *Let $Y_1, \dots, Y_n \in \mathbb{R}^d$ be n random variables, independent and all sub-Gaussian in the sense that for all $i \leq n$, the random variable Y_i has mean μ_i , covariance matrix Σ_i and satisfies that*

$$\forall \lambda \in \mathbb{R}^d \quad \ln \mathbb{E} \exp[\langle \lambda, Y_i - \mu_i \rangle] \leq \frac{\lambda^T \Sigma_i \lambda}{2}.$$

Assume moreover that each Σ_i is diagonal. Then, the sum of these random variables has the following concentration property. For all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\left\|\sum_{i=1}^n (Y_i - \mu_i)\right\| \geq \sqrt{2n\bar{\rho}_n \ln(2d/\delta)}\right) \leq \delta,$$

where $\bar{\rho}_n \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \text{trace}(\Sigma_i)$ is the average trace of the covariance matrices.

Proof: For a vector $y \in \mathbb{R}^d$, let us write $y_k \in \mathbb{R}$ is k -th component. First, by using the definition of the norm and a union bound, we get that for all $\{\epsilon_k\}_{k \leq d}$ such that $\sum_{k=1}^d \epsilon_k^2 = \epsilon^2$, then

$$\begin{aligned} \mathbb{P}\left(\left\|\sum_{i=1}^n (Y_i - \mu_i)\right\| \geq \epsilon\right) &= \mathbb{P}\left(\sum_{k=1}^d \left[\sum_{i=1}^n (Y_{k,i} - \mu_{k,i})\right]^2 \geq \epsilon^2\right) \\ &\leq \sum_{k=1}^d \mathbb{P}\left(\sum_{i=1}^n (Y_{k,i} - \mu_{k,i}) \geq \epsilon_k\right) + \sum_{k=1}^d \mathbb{P}\left(\sum_{i=1}^n (\mu_{k,i} - Y_{k,i}) \geq \epsilon_k\right), \end{aligned}$$

where we used the fact that the event $|Y_{k,i} - \mu_{k,i}| \leq \epsilon_k$ is equivalent to having at the same time both $Y_{k,i} - \mu_{k,i} \leq \epsilon_k$ and $\mu_{k,i} - Y_{k,i} \leq \epsilon_k$.

Then, for all positive $\lambda_k > 0$ we have by Markov's inequality that

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n (Y_{k,i} - \mu_{k,i}) \geq \epsilon_k\right) &= \mathbb{P}\left(\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_{k,i})\right) \geq \exp(\lambda_k \epsilon_k)\right) \\ &\leq \exp(-\lambda_k \epsilon_k) \mathbb{E}\left[\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_{k,i})\right)\right] \\ &= \exp(-\lambda_k \epsilon_k) \prod_{i=1}^n \mathbb{E}\left[\exp\left(\lambda_k (Y_{k,i} - \mu_{k,i})\right)\right], \end{aligned}$$

where we used the independence of the random variables in the last line. Now using the sub-Gaussian assumption with $\lambda = (0, \dots, 0, \lambda_k, 0, \dots, 0)^T \in \mathbb{R}^d$ where λ_k is at position k , then we deduce that

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n (Y_{k,i} - \mu_{k,i}) \geq \epsilon_k\right) &\leq \exp(-\lambda_k \epsilon_k) \prod_{i=1}^n \exp\left(\frac{\lambda_k^2}{2} \Sigma_{k,k,i}\right) \\ &= \exp(-\lambda_k \epsilon_k + \frac{\lambda_k^2}{2} \sum_{i=1}^n \Sigma_{k,k,i}), \end{aligned}$$

where $\Sigma_{k,k,i}$ is the value of element (k, k) of the matrix Σ_i . After optimizing this quantity with the choice $\lambda_k \stackrel{\text{def}}{=} \frac{\epsilon_k}{\sum_{i=1}^n \Sigma_{k,k,i}}$ and proceeding in a similar way for the symmetric term, we get

$$\mathbb{P}\left(\left\|\sum_{i=1}^n (Y_i - \mu_i)\right\| \geq \epsilon\right) \leq 2 \sum_{k=1}^d \exp\left(-\frac{\epsilon_k^2}{2 \sum_{i=1}^n \Sigma_{k,k,i}}\right).$$

Now let us choose $\epsilon_k^2 \stackrel{\text{def}}{=} 2 \sum_{i=1}^n \Sigma_{k,k,i} \ln(2d/\delta)$ so that with the notation of the lemma, we have exactly $\epsilon^2 = 2n\bar{\rho}_n \ln(2d/\delta)$. We conclude from the fact that with our definition of ϵ_k^2 , then

$$2 \sum_{k=1}^d \exp\left(-\frac{\epsilon_k^2}{2 \sum_{i=1}^n \Sigma_{k,k,i}}\right) = \delta.$$

□

In the case of a martingale difference sequence, we replace the average trace of the covariance matrices with the supremum of ρ over the domain of the variables.

Lemma 3 *Let us consider a sequential sampling process on \mathcal{X} that samples at time $i \leq n$ a value $Y_i \sim \nu(X_i)$ at point X_i , where $X_i \in \mathcal{F}_{< i}$ is a measurable function of the past inputs and outputs $\{(X_j, Y_j)\}_{j < i}$. Assume that for all i , Y_i is sub-Gaussian with conditional mean $\mu(X_i)$ and conditional covariance matrix $\Sigma(X_i)$, in the sense that*

$$\forall \lambda \in \mathbb{R}^d \quad \ln \mathbb{E}\left(\exp[\langle \lambda, Y_i - \mu(X_i) \rangle] \mid X_i\right) \leq \frac{\lambda^T \Sigma(X_i) \lambda}{2} \text{ a.s.}$$

Assume moreover that each $\Sigma(X_i)$ is diagonal. Then, the sum of these random variables has the following concentration property. For all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\left\|\sum_{i=1}^n (Y_i - \mu(X_i))\right\| \geq \sqrt{2nd\rho^* \ln(2d/\delta)}\right) \leq \delta,$$

where $\rho^* = \sup_{x \in \mathcal{X}} \text{trace}(\Sigma(x))$ is the supremum of the trace.

Note that we loose here an additional factor \sqrt{d} compared to the independent case. This additional factor would not appear if we replace the ℓ_2 norm with the ℓ_1 norm, however we do not know whether it can be removed for the ℓ_2 norm or whether this is a restriction due to our analysis.

Proof: For a vector $y \in \mathbb{R}^d$, let us write $y_k \in \mathbb{R}$ is k -th component. First, by using the definition of the norm and a union bound, we get that for all $\{\epsilon_k\}_{k \leq d}$ such that $\sum_{k=1}^d \epsilon_k^2 = \epsilon^2$, then

$$\begin{aligned} \mathbb{P}\left(\left\|\sum_{s=1}^t (Y_s - \mu(X_s))\right\| \geq \epsilon\right) &= \mathbb{P}\left(\sum_{k=1}^d \left[\sum_{s=1}^t (Y_{k,s} - \mu_k(X_s))\right]^2 \geq \epsilon^2\right) \\ &\leq \sum_{k=1}^d \mathbb{P}\left(\sum_{s=1}^t (Y_{k,s} - \mu_k(X_s)) \geq \epsilon_k\right) + \sum_{k=1}^d \mathbb{P}\left(\sum_{s=1}^t (\mu_k(X_s) - Y_{k,s}) \geq \epsilon_k\right), \end{aligned}$$

where we used the fact that the event $|Y_{k,s} - \mu_k(X_s)| \leq \epsilon_k$ is equivalent to having at the same time both $Y_{k,s} - \mu_k(X_s) \leq \epsilon_k$ and $\mu_k(X_s) - Y_{k,s} \leq \epsilon_k$.

Then, for all positive $\lambda_k > 0$ we have by Markov's inequality that

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n (Y_{k,i} - \mu_k(X_i)) \geq \epsilon_k\right) &= \mathbb{P}\left(\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))\right) \geq \exp(\lambda_k \epsilon_k)\right) \\ &\leq \exp(-\lambda_k \epsilon_k) \mathbb{E}\left[\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))\right)\right]. \end{aligned}$$

Now by using the conditional independence property, and the sub-Gaussian assumption, we get

$$\begin{aligned} \mathbb{E}\left[e^{\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))}\right] &= \mathbb{E}\left[\mathbb{E}\left[e^{\lambda_k (Y_{k,n} - \mu_k(X_n))} \mid X_n, \mathcal{F}_{<n}\right] \mathbb{E}\left[\prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))} \mid X_t, \mathcal{F}_{<t}\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[e^{\lambda_k (Y_{k,n} - \mu_k(X_n))} \mid X_n\right] \prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{\lambda_k^2}{2} \Sigma_{k,k}(X_n)\right) \prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right], \end{aligned}$$

where $\Sigma_{k,k}(X_n)$ is the value of element (k, k) of the matrix $\Sigma(X_n)$. At this point we need a uniform control on $\Sigma_{k,k}(X_n)$ due to the dependency of X_n with the past random variables in order to derive a useful bound. We thus introduce the quantity $\rho^* \stackrel{\text{def}}{=} \max_{x \in \mathcal{X}} \text{trace}(\Sigma(x))$. Since Σ is a positive matrix, then for any $x \in \mathcal{X}$, $\Sigma_{k,k}(x) \leq \text{trace}(\Sigma(x)) \leq \rho^*$. Thus we get by applying the same kind of transformations as above recursively

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n (Y_{k,i} - \mu_k(X_i)) \geq \epsilon_k\right) &\leq \exp(-\lambda_k \epsilon_k) \exp\left(\frac{\lambda_k^2}{2} \Sigma_{k,k}^*\right) \mathbb{E}\left[\prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right] \\ &\leq \exp\left(-\lambda_k \epsilon_k + \frac{\lambda_k^2}{2} n \rho^*\right). \end{aligned}$$

After optimizing this quantity with the choice $\lambda_k \stackrel{\text{def}}{=} \frac{\epsilon_k}{n \rho^*}$ and proceeding in a similar way for the symmetric term, we get

$$\mathbb{P}\left(\left\|\sum_{i=1}^n (Y_i - \mu(X_i))\right\| \geq \epsilon\right) \leq 2 \sum_{k=1}^d \exp\left(-\frac{\epsilon_k^2}{2n \rho^*}\right).$$

Finally we choose $\epsilon_k^2 \stackrel{\text{def}}{=} 2n \rho^* \ln(2d/\delta)$ so that $\epsilon^2 = 2nd \rho^* \ln(2d/\delta)$.

□

Finally, we provide a standard Azuma's inequality for bounded martingale difference sequence with d -dimensional variables.

Lemma 4 *Let us consider a sequential sampling process on \mathcal{X} that samples at time $i \leq n$ a value $Y_i \sim \nu(X_i)$ at point X_i , where $X_i \in \mathcal{F}_{<i}$ is a measurable function of the past inputs and outputs*

$\{(X_j, Y_j)\}_{j < i}$. Assume that for all t , $Y_i \in [0, 1]^d$ with conditional mean $\mu(X_i)$ and diagonal conditional covariance matrix $\Sigma(X_i)$. Then, the sum of these random variables has the following concentration property. For all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i)\right\| \geq \sqrt{\frac{2d \ln(2d/\delta)}{n}}\right) \leq \delta.$$

Proof: For a vector $y \in \mathbb{R}^d$, let us write $y_k \in \mathbb{R}$ is k -th component. First, by using the definition of the norm and a union bound, we get that for all $\{\epsilon_k\}_{k \leq d}$ such that $\sum_{k=1}^d \epsilon_k^2 = \epsilon^2$, then

$$\begin{aligned} \mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i)\right\| \geq \epsilon\right) &= \mathbb{P}\left(\sum_{k=1}^d \left[\frac{1}{n} \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))\right]^2 \geq \epsilon^2\right) \\ &\leq \sum_{k=1}^d \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i)) \geq \epsilon_k\right) + \sum_{k=1}^d \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n (\mu_k(X_i) - Y_{k,i}) \geq \epsilon_k\right), \end{aligned}$$

where we used the fact that the event $|Y_{k,i} - \mu_k(X_i)| \leq \epsilon_k$ is equivalent to having at the same time both $Y_{k,i} - \mu_k(X_i) \leq \epsilon_k$ and $\mu_k(X_i) - Y_{k,i} \leq \epsilon_k$.

Then, for all positive $\lambda_k > 0$ we have by Markov's inequality that

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i)) \geq \epsilon_k\right) &= \mathbb{P}\left(\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))\right) \geq \exp(\lambda_k n \epsilon_k)\right) \\ &\leq \exp(-\lambda_k n \epsilon_k) \mathbb{E}\left[\exp\left(\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))\right)\right]. \end{aligned}$$

Now since by the bounded assumption $Y_{k,n} \in [0, 1]$, we have that $Y_{k,n} - \mu_k(X_n) \in [-1, 1]$, then by an application of Hoeffding's Lemma we get

$$\ln \mathbb{E}\left[e^{\lambda_k (Y_{k,n} - \mu_k(X_n))} \mid X_n\right] \leq \frac{\lambda_k^2}{2},$$

and thus we deduce by using the conditional independence property that

$$\begin{aligned} \mathbb{E}\left[e^{\lambda_k \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i))}\right] &= \mathbb{E}\left[\mathbb{E}\left[e^{\lambda_k (Y_{k,n} - \mu_k(X_n))} \mid X_n, \mathcal{F}_{<n}\right] \mathbb{E}\left[\prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))} \mid X_n, \mathcal{F}_{<n}\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[e^{\lambda_k (Y_{k,n} - \mu_k(X_n))} \mid X_n\right] \prod_{i=1}^{n-1} \mathbb{E}\left[e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right]\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{\lambda_k^2}{2}\right) \prod_{i=1}^{n-1} \mathbb{E}\left[e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right]\right]. \end{aligned}$$

Thus we get by applying the same kind of transformations as above recursively

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n (Y_{k,i} - \mu_k(X_i)) \geq \epsilon_k\right) &\leq \exp(-\lambda_k n \epsilon_k) \exp\left(\frac{\lambda_k^2}{2}\right) \mathbb{E}\left[\prod_{i=1}^{n-1} e^{\lambda_k (Y_{k,i} - \mu_k(X_i))}\right] \\ &\leq \exp\left(-\lambda_k n \epsilon_k + \frac{\lambda_k^2}{2} n\right). \end{aligned}$$

After optimizing this quantity with the choice $\lambda_k \stackrel{\text{def}}{=} \epsilon_k$ and proceeding in a similar way for the symmetric term, we get

$$\mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^n (Y_i - \mu(X_i))\right\| \geq \epsilon\right) \leq 2 \sum_{k=1}^d \exp\left(-\frac{n \epsilon_k^2}{2}\right).$$

Finally we choose $\epsilon_k^2 \stackrel{\text{def}}{=} \frac{2 \ln(2d/\delta)}{n}$ so that $\epsilon^2 = \frac{2d \ln(2d/\delta)}{n}$.

□

A.3 Variance estimation

We now analyze an upper-bound estimation of the variance. Note that there is an important difference between the iid case and the case of a martingale difference sequence. The main reason is that the standard estimator of the variance is biased when the mean of the variables considered do not have the same mean, which happens for a martingale difference sequence.

Lemma 5 *Let $\{X_i\}_{1 \leq i \leq n}$ be independent real-valued random variables, with mean μ_i and variance σ_i^2 . Let the following estimator*

$$\hat{\sigma}_n^2 \stackrel{\text{def}}{=} \frac{1}{n-1} \sum_{i=1}^n \left(X_i - \frac{1}{n} \sum_{j=1}^n X_j \right)^2,$$

satisfies that

$$\mathbb{E}[\hat{\sigma}_n^2] = \frac{1}{n} \sum_{i=1}^n \sigma_i^2 + \frac{1}{n-1} \sum_{i=1}^n \left(\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j \right)^2.$$

Proof: By construction, we have that

$$\begin{aligned} \mathbb{E}[\hat{\sigma}_n^2] &= \mathbb{E} \left[\frac{n}{n-1} \left(\frac{1}{n} \sum_{i=1}^n (X_i - \mu_i) - \frac{1}{n} \sum_{j=1}^n (X_j - \mu_j) \right)^2 \right] + \frac{1}{n-1} \sum_{i=1}^n \left(\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j \right)^2 \\ &= \mathbb{E} \left[\frac{n}{n-1} \left(\frac{1}{n} \sum_{i=1}^n (X_i - \mu_i)^2 - \left(\frac{1}{n} \sum_{j=1}^n (X_j - \mu_j) \right)^2 \right) \right] + \frac{1}{n-1} \sum_{i=1}^n \left(\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j \right)^2 \\ &= \frac{n}{n-1} \left(\frac{1}{n} \sum_{i=1}^n \sigma_i^2 - \frac{1}{n^2} \sum_{i=1}^n \sigma_i^2 \right) + \frac{1}{n-1} \sum_{i=1}^n \left(\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j \right)^2 \\ &= \frac{1}{n} \sum_{i=1}^n \sigma_i^2 + \frac{1}{n-1} \sum_{i=1}^n \left(\mu_i - \frac{1}{n} \sum_{j=1}^n \mu_j \right)^2. \end{aligned}$$

□

A.3.1 Hoeffding-like upper bounds

The next lemma provides a result “à la Hoeffding” for the variance estimation of a martingale difference sequence.

Lemma 6 *Let us consider a sequential sampling process on $R \subset \mathcal{X}$ that samples at time $i \leq n$ a value $Y_i \sim \nu(X_i)$ at point $X_i \in R$, where $X_i \in \mathcal{F}_{<i}$ is a measurable function of the past inputs and outputs $\{(X_j, Y_j)\}_{j < i}$. Assume that for all t , $Y_i \in [0, 1]^d$ with conditional mean $\mu(X_i)$ and diagonal conditional covariance matrix $\Sigma(X_i)$. Then the following estimator (note that it is biased)*

$$\hat{\sigma}_s^2 = \frac{1}{n} \sum_{i=1}^n \|Y_i\|^2 - \left\| \frac{1}{n} \sum_{i=1}^n Y_i \right\|^2.$$

satisfies the property that for all $\delta \in [0, 1]$,

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n \rho_i - \hat{\sigma}_n^2 \geq (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n} \right) \leq \delta.$$

Proof: Since $\rho_i = \mathbb{E}[\|Y_i\|^2 | X_i] - \|\mu(X_i)\|^2$, we deduce the following rewriting

$$\frac{1}{n} \sum_{i=1}^n \rho_i - \hat{\sigma}_n^2 = \left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|Y_i\|^2 | X_i] - \|\mu(X_i)\|^2 \right) - \frac{1}{n} \sum_{i=1}^n \|\mu(X_i)\|^2 + \left\| \frac{1}{n} \sum_{j=1}^n Y_j \right\|^2.$$

Now, we use the fact that on the one hand, we have the equality

$$\left\| \frac{1}{n} \sum_{j=1}^n Y_j \right\|^2 = \left\| \frac{1}{n} \sum_{j=1}^n Y_j - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|^2 = \left\| \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|^2 + 2 \left\langle \frac{1}{n} \sum_{i=1}^n Y_i, \frac{1}{n} \sum_{k=1}^n \mu(X_k) \right\rangle,$$

and on the other hand we have the equality

$$\frac{1}{n} \sum_{i=1}^n \|\mu(X_i)\|^2 = \frac{1}{n} \sum_{i=1}^n \|\mu(X_i) - \frac{1}{n} \sum_{j=1}^n \mu(X_j)\|^2 + \|\frac{1}{n} \sum_{i=1}^n \mu(X_i)\|^2,$$

to deduce the following bound:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \rho_i - \hat{\sigma}_n^2 &= \left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|Y_i\|^2 | X_i] - \|Y_i\|^2 \right) + \left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|^2 \\ &\quad + 2 \left\langle \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i), \frac{1}{n} \sum_{k=1}^n \mu(X_k) \right\rangle - \frac{1}{n} \sum_{i=1}^n \|\mu(X_i) - \frac{1}{n} \sum_{j=1}^n \mu(X_j)\|^2 \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|Y_i\|^2 | X_i] - \|Y_i\|^2 \right) + \left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|^2 + 2\sqrt{d} \left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|, \end{aligned}$$

where we used in the last line Cauchy-Schwarz's inequality and the fact that $\left\| \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\| \leq \sqrt{d}$ since $\mathcal{Y} \subset [0, 1]^d$ by assumption.

We now take care of these three terms. By an application of Hoeffding-Azuma's inequality for the random variables $\{\|Y_i\|^2\}_{i \leq n}$ that belongs to the interval $[0, d]$ since by assumption $\mathcal{Y} \subset [0, 1]^d$, we first deduce that for all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|Y_i\|^2 | X_i] - \|Y_i\|^2 \geq \sqrt{\frac{d \ln(1/\delta)}{2n}}\right) \leq \delta.$$

Similarly, by application of Lemma 4, then for all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\| \geq \sqrt{\frac{d \ln(2d/\delta)}{2n}}\right) \leq \delta.$$

Thus, by using a union bound over the two events, we deduce that with probability higher than $1 - 2\delta$, then

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \rho_i - \hat{\sigma}_n^2 &\leq \sqrt{\frac{d \ln(1/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n} + 2d \sqrt{\frac{\ln(2d/\delta)}{2n}} \\ &\leq (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n}. \end{aligned}$$

□

This Lemma enables us to prove the following upper confidence bound on the value $\rho^*(R)$ of a region R .

Proposition 1 *Under the assumptions that \mathcal{Y} is a convex subset of $[0, 1]^d$, ν is sub-Gaussian, ρ is Lipschitz w.r.t. ℓ_2 and $R \subset \mathcal{X}$ is compact and convex, then*

$$\mathbb{P}\left(\forall x \in \mathcal{X}; U_t(R, x, \delta) \leq \rho(x)\right) \leq t\delta$$

The proof of the bound for the quantity $L_t(R, x, \delta)$ results from a straightforward adaptation of this proof.

Proof: Let us introduce the following notation

$$b(n, \delta) \stackrel{\text{def}}{=} (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n}.$$

Since ρ is Lipschitz, then we deduce that for all $x \in \mathcal{X}$

$$\sum_{s=1}^t (\rho(X_s) + \ell_2(X_s, x) - \rho(x)) \mathbb{I}\{X_s \in R\} \geq \sum_{s=1}^t (-\ell_2(X_s, x) + \ell_2(X_s, x_R)) \mathbb{I}\{X_s \in R\} \geq 0.$$

Thus, using this remark and the above notation, we deduce the following bound

$$\begin{aligned}
& \mathbb{P}\left(\forall x \in \mathcal{X}; U_t(R, x, \delta) \leq \rho(x) \text{ and } N_t(R) \geq 1\right) \\
&= \mathbb{P}\left(\forall x \in \mathcal{X}; \hat{\sigma}_t^2(R) + b(N_t(R), \delta) + \frac{1}{N_t(R)} \sum_{s=1}^t \ell_2(X_s, x) \mathbb{I}\{X_s \in R\} \leq \rho(x) \text{ and } N_t(R) \geq 1\right) \\
&= \mathbb{P}\left(N_t(R) \hat{\sigma}_t^2(R) + \left[\sum_{s=1}^t (\ell_2(X_s, x) - \rho^*(x)) \mathbb{I}\{X_s \in R\}\right] \leq -b(N_t(R), \delta) N_t(R) \text{ and } N_t(R) \geq 1\right) \\
&\leq \mathbb{P}\left(\sum_{s=1}^t \rho(X_s) \mathbb{I}\{X_s \in R\} - N_t(R) \hat{\sigma}_t^2(R) \geq b(N_t(R), \delta) N_t(R) \text{ and } N_t(R) \geq 1\right).
\end{aligned}$$

We handle this last term thanks to a concentration result for sequences of martingale difference. Let us introduce the stopping times $\{\tau_i\}_{i \geq 1}$ corresponding to the times when an arm falls into R , defined by

$$\tau_i \stackrel{\text{def}}{=} \min\{t : N_t(R) = i\}.$$

We then denote $\tilde{X}_i \stackrel{\text{def}}{=} X_{\tau_i}$ the i^{th} samples in region R , and its associated value $\tilde{Y}_i \stackrel{\text{def}}{=} Y_{\tau_i}$. First, note that using this definition, we have on the event $\{N_t(R) = s\}$ the rewriting

$$\hat{\sigma}_t^2(R) = \frac{1}{s} \sum_{i=1}^s \|\tilde{Y}_i\|^2 - \left\| \frac{1}{s} \sum_{j=1}^s \tilde{Y}_j \right\|^2,$$

that we define to be $\tilde{\sigma}_s^2(R)$ for convenience. Then, we plug this equality in the main term, and using the fact that $N_t(R) \leq t$ we get

$$\begin{aligned}
& \mathbb{P}\left(\sum_{s=1}^t \rho(X_s) \mathbb{I}\{X_s \in R\} - N_t(R) \hat{\sigma}_t^2(R) \geq b(N_t(R), \delta) N_t(R) \text{ and } N_t(R) \geq 1\right) \\
&= \mathbb{P}\left(\sum_{i=1}^{N_t(R)} \rho(\tilde{X}_i) - N_t(R) \hat{\sigma}_t^2(R) \geq b(N_t(R), \delta) N_t(R) \text{ and } N_t(R) \geq 1\right) \\
&= \sum_{s=1}^t \mathbb{P}\left(\sum_{i=1}^s \rho(\tilde{X}_i) - s \tilde{\sigma}_s^2(R) \geq sb(s, \delta)\right) \\
&\leq \sum_{s=1}^t \delta,
\end{aligned}$$

where we used in the last line an application of Lemma 6, thanks to an application of Azuma's inequality for bounded martingale difference sequence and of Lemma 4. □

A.3.2 Hoeffding-like lower bounds

Lemma 7 *Let us consider a sequential sampling process on $R \subset \mathcal{X}$ that samples at time $i \leq n$ a value $Y_i \sim \nu(X_i)$ at point $X_i \in R$, where $X_i \in \mathcal{F}_{< i}$ is a measurable function of the past inputs and outputs $\{(X_j, Y_j)\}_{j < i}$. Assume that for all t , $Y_i \subset [0, 1]^d$ with conditional mean $\mu(X_i)$ and diagonal conditional covariance matrix $\Sigma(X_i)$. Then the following estimator (note that it is biased)*

$$\hat{\sigma}_s^2 = \frac{1}{n} \sum_{i=1}^n \|Y_i\|^2 - \left\| \frac{1}{n} \sum_{i=1}^n Y_i \right\|^2.$$

satisfies the property that for all $\delta \in [0, 1]$,

$$\mathbb{P}\left(\hat{\sigma}_n^2 - \frac{1}{n} \sum_{i=1}^n \rho_i \geq (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + d_1(R)^2\right) \leq \delta.$$

Proof: We start as for the proof of Lemma 6 and thus directly consider derive the following inequality

$$\begin{aligned}
\hat{\sigma}_n^2 - \frac{1}{n} \sum_{i=1}^n \rho_i &= \left(\frac{1}{n} \sum_{i=1}^n \|Y_i\|^2 - \mathbb{E}[\|Y_i\|^2 | X_i] \right) - \left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\|^2 \\
&\quad + 2 \left\langle \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i), \frac{1}{n} \sum_{k=1}^n \mu(X_k) \right\rangle + \frac{1}{n} \sum_{i=1}^n \|\mu(X_i) - \frac{1}{n} \sum_{j=1}^n \mu(X_j)\|^2 \\
&\leq \left(\frac{1}{n} \sum_{i=1}^n \|Y_i\|^2 - \mathbb{E}[\|Y_i\|^2 | X_i] \right) + 2\sqrt{d} \left\| \frac{1}{n} \sum_{i=1}^n Y_i - \frac{1}{n} \sum_{i=1}^n \mu(X_i) \right\| + d_1(R)^2,
\end{aligned}$$

where we used in the last line the fact that by Cauchy-Schwarz's inequality, by the assumption that μ is Lipschitz w.r.t ℓ_1 , and since $X_i \in R$ for all $i \leq n$,

$$\begin{aligned}
\left\| \mu(X_i) - \frac{1}{n} \sum_{j=1}^n \mu(X_j) \right\|^2 &\leq \frac{1}{(n-1)^2} \sum_{j,k \neq i=1}^n \|\mu(X_i) - \mu(X_j)\| \|\mu(X_i) - \mu(X_k)\| \\
&\leq d_1(R)^2.
\end{aligned}$$

We conclude similarly to Lemma 6

□

We finally prove the following lower bound for a region $R \subset \mathcal{T}_t$.

Proposition 2 *Under the assumptions that \mathcal{Y} is a convex subset of $[0, 1]^d$, ν is sub-Gaussian, μ is Lipschitz w.r.t. ℓ_1 , ρ is Lipschitz w.r.t. ℓ_2 and $R \subset \mathcal{X}$ is compact and convex, then*

$$\mathbb{P}\left(\forall x \in \mathcal{X}; U_t(R, x, \delta) \geq \rho(x) + 2 \max_{x' \in R} \ell_2(x, x') + d_1(R)^2 + 2(1+2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2N_t(R)}} + \frac{d \ln(2d/\delta)}{2N_t(R)}\right) \leq t\delta$$

The proof of the bound for the quantity $L_t(R, x, \delta)$ results from a straightforward adaptation of this proof.

Proof: Indeed, we first get the following bound by definition and by the assumption that ρ Lipschitz w.r.t. ℓ_2 , for all $x \in \mathcal{X}$

$$\begin{aligned}
U_t(R, x, \delta) &= \hat{\sigma}_t^2(R) + b(N_t(R), \delta) + \frac{1}{N_t(R)} \sum_{s=1}^t \ell_2(X_s, x) \mathbb{I}\{X_s \in R\} \\
&= \hat{\sigma}_t^2(R) - \frac{1}{N_t(R)} \sum_{s=1}^t \rho(X_s) \mathbb{I}\{X_s \in R\} + b(N_t(R), \delta) \\
&\quad + \frac{1}{N_t(R)} \sum_{s=1}^t \rho(X_s) \mathbb{I}\{X_s \in R\} + \frac{1}{N_t(R)} \sum_{s=1}^t \ell_2(X_s, x) \mathbb{I}\{X_s \in R\} \\
&\leq \hat{\sigma}_t^2(R) - \frac{1}{N_t(R)} \sum_{s=1}^t \rho(X_s) \mathbb{I}\{X_s \in R\} + b(N_t(R), \delta) + \rho(x) + 2 \max_{x' \in R} \ell_2(x, x'),
\end{aligned}$$

where we remind that by definition

$$b(n, \delta) \stackrel{\text{def}}{=} (1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2n}} + \frac{d \ln(2d/\delta)}{2n}.$$

We now proceed as for the proof Proposition 1, by introducing the stopping times and the variables \tilde{X}_i and \tilde{Y}_i . With these notations, we thus deduce that

$$\mathbb{P}\left(\forall x \in \mathcal{X}; U_t(R, x, \delta) \geq \rho(x) + 2 \max_{x' \in R} \ell_2(x, x') + d_1(R)^2 + 2(1+2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta)}{2N_t(R)}} + \frac{d \ln(2d/\delta)}{2N_t(R)} \text{ and } N_t(R) \geq 1\right)$$

$$\begin{aligned}
&\leq \sum_{s=1}^t \mathbb{P} \left(s\tilde{\sigma}_s^2(R) - \sum_{i=1}^s \rho(\tilde{X}_i) \geq -sb(s, \delta) + sd_1(R)^2 + 2(1 + 2\sqrt{d})\sqrt{ds \ln(2d/\delta)/2} + ds \ln(2d/\delta)/2 \right) \\
&\leq \sum_{s=1}^t \mathbb{P} \left(\tilde{\sigma}_s^2(R) - \frac{1}{s} \sum_{i=1}^s \rho(\tilde{X}_i) \geq (1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta)}{2s}} + d_1(R)^2 \right) \\
&\leq \sum_{s=1}^t \delta.
\end{aligned}$$

□

B Main regret bound

B.1 Hierarchical Optimistic Sampling

In this subsection, we now provide properties of the tree expanded by the HORSE.C algorithm.

Let us consider $(h_0, i_0) \in \mathcal{P}$ and let us define for all $(h, i) \in \mathcal{T}(h_0, i_0)$ in the sub tree of \mathcal{T}_t with root (h_0, i_0) the following quantity called respectively the max-gap and the min gap

$$\Delta_{h_0, i_0}^+(h, i) \stackrel{\text{def}}{=} \max_{x \in R(h_0, i_0)} \rho(x) - \max_{x \in R(h, i)} \rho(x) \quad \text{and} \quad \Delta_{h_0, i_0}^-(h, i) \stackrel{\text{def}}{=} \min_{x \in R(h, i)} \rho(x) - \min_{x \in R(h_0, i_0)} \rho(x).$$

A node (h, i) is called (h_0, i_0) -max-optimal if it holds that $\Delta_{h_0, i_0}^+(h, i) = 0$, and similarly it is called (h_0, i_0) -min-optimal in the case when $\Delta_{h_0, i_0}^-(h, i) = 0$.

Lemma 8 *Under the assumptions of Section 3, and for the tree \mathcal{T}_t grown by the HORSE.C algorithm, for all $(h, i) \in \mathcal{T}_t(h_0, i_0)$ such that $\Delta_{h_0, i_0}^+(h, i) > 2c_2\gamma_2^h + c_1^2\gamma_1^{2h}$, then there exists an event $\Omega(h, i)$ of probability higher than $1 - \sum_{j=1}^{k-1} (2j+1)^2 \delta_{2j+1, t}$ such that on $\Omega(h, i) \cap \{N_t(h_0, i_0) = 2k+1\}$, then*

$$N_t(h, i) \leq [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{2k+1, t})/2}]^2 \frac{d \ln(2d/\delta_{2k+1, t})}{2(\Delta_{h_0, i_0}^+(h, i) - 2c_2\gamma_2^h - c_1^2\gamma_1^{2h})^2}.$$

Proof: The proof follows the same steps as in Lemma 14 of Bubeck et al. [2011]. Let us consider some path of (h_0, i_0) -optimal nodes starting from the root $\{(h, i_h^*)\}_{h \geq h_0}$. Let $(h, i) \in \mathcal{T}(h_0, i_0)$ where $h = h_0 + k$ be such that $\Delta_{h_0, i_0}^+(h, i) > 2c_2\gamma_2^h + c_1^2\gamma_1^{2h}$, and $h_0 \leq h^* \leq h-1$ with $k^* = h^* - h_0$ be the largest depth such that $(h^*, i_{h^*}^*)$ is on the path from (h_0, i_0) to (h, i) .

Step 1. Let us consider a round $t+1$ such that $(H_{t+1}, I_{t+1}) = (h_0, i_0)$ and the event $N_t(R(h_0, i_0)) = n = 2n_0 + 1$ for some integers n, n_0 . Thus $N_t(R(h_0, i_0))$ is odd. Now if $(h_{t+1}, i_{t+1}) \in \mathcal{T}(h, i)$, then this is because $\hat{\rho}_t^+(H_{t+1}^{k^*+1}, I_{t+1}^{k^*+1}; \delta_{n, t}) \geq \hat{\rho}_t^+(h^*+1, i_{h^*+1}^*; \delta_{n, t})$. Moreover, on the path $\{(H_{t+1}^j, I_{t+1}^j)\}_{j \geq 0}$, the value of $\hat{\rho}_t^+$ can only increase. Indeed, we have by definition of the algorithm

$$\begin{aligned}
\hat{\rho}_t^+(H_{t+1}^j, I_{t+1}^j; \delta_{n, t}) &\leq \max\{\hat{\rho}_t^+(h, i; \delta_{n, t}); (h, i) \in \mathcal{C}_t(H_{t+1}^j, I_{t+1}^j)\} \\
&\leq \hat{\rho}_t^+(H_{t+1}^{j+1}, I_{t+1}^{j+1}; \delta_{n, t}).
\end{aligned}$$

Thus, we deduce that $\hat{\rho}_t^+(h^*+1, i_{h^*+1}^*; \delta_{n, t}) \leq \hat{\rho}_t^+(h, i; \delta_{n, t}) \leq \max_{x \in R(h, i)} U_t(R(h, i), x, \delta_{n, t})$, by using the definition of $\hat{\rho}_t^+(h, i; \delta_{n, t})$. Now by property of the optimal path, we also have the useful recursive equality

$$\hat{\rho}_t^+(h^*+1, i_{h^*+1}^*; \delta_{n, t}) = \min\left\{ \max_{x \in R(h^*+1, i_{h^*+1}^*)} U_t(R(h^*+1, i_{h^*+1}^*), x, \delta_{n, t}), \hat{\rho}_t^+(h^*+2, i_{h^*+2}^*; \delta_{n, t}) \right\}.$$

Then, at round $t + 1$ we know that no more than $N_t(h_0, i_0) = 2n_0 + 1$ nodes have been played in $\mathcal{T}(h_0, i_0)$, thus the node (n, i_n^*) has not been sampled and $\hat{\rho}_t^+(n, i_n^*; \delta_{n,t}) = \infty$. This proves that

$$\begin{aligned} & \mathbb{P}\left((h_{t+1}, i_{t+1}) \in \mathcal{T}(h, i) \text{ and } N_t(h_0, i_0) = 2n_0 + 1\right) \leq \\ & \sum_{h=1}^{2n_0} \mathbb{P}\left(\max_{x \in R(h, i_h^*)} U_t(R(h, i_h^*), x, \delta_{n,t}) \leq \max_{x \in R(h_0, i_0)} \rho(x) \text{ and } N_t(h_0, i_0) = 2n_0 + 1\right) \\ & + \mathbb{P}\left(\max_{x \in R(h, i)} U_t(R(h, i), x, \delta_{n,t}) > \max_{x \in R(h_0, i_0)} \rho(x) \text{ and } N_t(h_0, i_0) = 2n_0 + 1\right) \end{aligned}$$

Step 2. Now note that since for all $h \leq N_t(h_0, i_0) - 1 = 2n_0$, we have that (h, i_h^*) is an (h_0, i_0) -optimal node and that $N_t(h, i_h^*) \leq N_t(h_0, i_0) = n$, then we deduce that

$$\mathbb{P}\left(\max_{x \in R(h, i_h^*)} U_t(R(h, i_h^*), x, \delta_{n,t}) \leq \max_{x \in R(h_0, i_0)} \rho(x) \text{ and } N_t(h_0, i_0) = 2n_0 + 1\right) \leq n\delta_{n,t},$$

Since $\Delta_{h_0, i_0}^+(h, i) > 2c_2\gamma_2^h + c_1^2\gamma_1^{2h}$, let us define some integer $u_n > 0$ such that

$$u_n \geq [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{n,t})/2}]^2 \frac{d \ln(2d/\delta_{n,t})}{2(\Delta_{h_0, i_0}^+(h, i) - 2c_2\gamma_2^h - c_1^2\gamma_1^{2h})^2}.$$

Using the fact that $\sqrt{x} \leq x$ for $x \leq 1$, such an integer satisfies that

$$\Delta_{h_0, i_0}^+(h, i) - 2c_2\gamma_2^h - c_1^2\gamma_1^{2h} \geq 2(1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta_{n,t})}{2u_n}} + \frac{d \ln(2d/\delta_{n,t})}{2u_n}.$$

Now let us introduce as a shorthand the notation $U_t \stackrel{\text{def}}{=} \max_{x \in R(h, i)} U_t(R(h, i), x, \delta_{n,t})$ as well as $E_n \stackrel{\text{def}}{=} \{N_t(h_0, i_0) = n = 2n_0 + 1\}$. With these notations, we deduce that for $k \geq n$, we have

$$\begin{aligned} & \mathbb{P}\left(U_t > \max_{x \in R(h_0, i_0)} \rho(x) \text{ and } E_n \text{ and } N_t(R(h, i)) > u_k\right) \\ & = \mathbb{P}\left(U_t > \Delta_{h_0, i_0}^+(h, i) + \max_{x \in R(h, i)} \rho(x) \text{ and } E_n \text{ and } N_t(R(h, i)) > u_k\right) \\ & \leq \mathbb{P}\left(U_t > \max_{x \in R(h, i)} \rho(x) + 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta_{k,t})}{2u_k}} \right. \\ & \quad \left. + \frac{d \ln(2d/\delta_{k,t})}{2u_k} \text{ and } E_n \text{ and } N_t(R(h, i)) > u_k\right) \\ & \leq \mathbb{P}\left(U_t > \max_{x \in R(h, i)} \rho(x) + 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta_{n,t})}{2N_t(R(h, i))}} + \frac{d \ln(2d/\delta_{n,t})}{2N_t(R(h, i))} \text{ and } E_n\right) \\ & \leq n\delta_{n,t}. \end{aligned}$$

Thus we have shown that for any $k \geq n$, then

$$\begin{aligned} \mathbb{P}\left((h_{t+1}, i_{t+1}) \in \mathcal{T}(h, i) \text{ and } N_t(h_0, i_0) = 2n_0 + 1 \text{ and } N_t(h, i) > u_k\right) & \leq n\delta_{n,t} + (n - 1)n\delta_{n,t} \\ & = n^2\delta_{n,t}. \end{aligned}$$

Step 3.

Let us introduce the stopping times s_j corresponding to the first round $\tau + 1$ such that $N_\tau(h_0, i_0) = 2j + 1$, and note that we thus have $(H_{s_j+1}, I_{s_j+1}) = (h_0, i_0)$ at this round. Let us use the notation $n_t \stackrel{\text{def}}{=} \lceil \frac{N_t(h_0, i_0) - 1}{2} \rceil$. Thus, we have the following rewriting

$$\begin{aligned} N_t(h, i) & = \sum_{s=0}^{t-1} \mathbb{I}\{(H_{s+1}, I_{s+1}) = (h_0, i_0)\} \mathbb{I}\{(h_{s+1}, i_{s+1}) \in \mathcal{T}(h, i)\} \\ & = \sum_{j=0}^{n_t-1} \mathbb{I}\{(h_{s_j+1}, i_{s_j+1}) \in \mathcal{T}(h, i)\}. \end{aligned}$$

Now for any $u > 0$, since $N_{s_j}(h, i) + 1 \leq N_{s_{j+1}}(h, i)$, on the one hand we deduce that

$$\sum_{j=0}^{n_t-1} \mathbb{I}\{(h_{s_j+1}, i_{s_j+1}) \in \mathcal{T}(h, i)\} \mathbb{I}\{N_{s_j}(h, i) \leq u\} \leq u,$$

and on the other hand, since $N_{s_{j+1}}(h, i) \leq N_{s_j}(h, i) + 2$, we deduce that

$$\sum_{j=0}^{n_t-1} \mathbb{I}\{(h_{s_j+1}, i_{s_j+1}) \in \mathcal{T}(h, i)\} \mathbb{I}\{N_{s_j}(h, i) > u\} \leq \sum_{j=\lceil \frac{u+1}{2} \rceil}^{n_t-1} \mathbb{I}\{(h_{s_j+1}, i_{s_j+1}) \in \mathcal{T}(h, i)\} \mathbb{I}\{N_{s_j}(h, i) > u\}.$$

Moreover, using the result of Step 2, we know that for each $j \leq k$, there exists an event Ω_j of probability higher than $1 - (2j+1)^2 \delta_{2j+1,t}$ such that on this event, then

$$\mathbb{I}\{(h_{s_j+1}, i_{s_j+1}) \in \mathcal{T}(h, i) \text{ and } N_{s_j}(h_0, i_0) = 2j+1\} \mathbb{I}\{N_{s_j}(h, i) > u_k\} = 0.$$

Thus, by combining those result with a union bound, we deduce that there exists an event $\Omega(h, i)$ of probability higher than $1 - \sum_{j=1}^{k-1} (2j+1)^2 \delta_{2j+1,t}$ such that on $\Omega(h, i) \cap \{N_t(h_0, i_0) = 2k+1\}$, then

$$N_t(h, i) \leq [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{2k+1,t})/2}]^2 \frac{d \ln(2d/\delta_{2k+1,t})}{2(\Delta_{h_0, i_0}^+(h, i) - 2c_2\gamma_2^h - c_1^2\gamma_1^{2h})^2}.$$

□

Lemma 9 Let $I_{h_0, i_0}^+(h)$ be the set of all nodes at depth $h \geq h_0$ satisfying

$$\Delta_{h_0, i_0}^+(h, i) \leq 4c_2\gamma_2^h + 2c_1^2\gamma_1^{2h}.$$

Then the number of elements of $I_{h_0, i_0}^+(h)$ is bounded, for all $d' \geq d^+(h_0, i_0)$ where $d^+(h_0, i_0)$ is the near-optimality dimension of ρ in region $R(h_0, i_0)$ has

$$|I_{h_0, i_0}^+(h)| \leq C(c_2'\gamma_2^h)^{-d'}.$$

where C is a constant depending only on $\gamma_1, \gamma_2, c_1, c_2, c_1', c_2'$ and $d^+(h_0, i_0)$.

Proof: First, using a similar argument as for Lemma 3 of Bubeck et al. [2011], then we have that all arms in (h, i) are $2(4c_2\gamma_2^h + 2c_1^2\gamma_1^{2h})$ optimal w.r.t node (h_0, i_0) . Now by the assumption that the diameter in ℓ_2 is lower bounded and that $\gamma_1^2 \leq \gamma_2$ then we deduce, following second step of the proof of Theorem in Bubeck et al. [2011] that there exists a constant depending only on $\gamma_1, \gamma_2, c_1, c_2, c_1', c_2'$ and $d(h_0, i_0)$ such that for all $d' \geq d^+(h_0, i_0)$ where $d^+(h_0, i_0)$ is the near-optimality dimension of ρ in region $R(h_0, i_0)$, then

$$|I_{h_0, i_0}^+(h)| \leq C(c_2'\gamma_2^h)^{-d'}.$$

□

Lemma 10 For all $(h, i) \in \mathcal{P}$, for all $k = 2k_o + 1 \geq 0$, there exists an event $\Omega(h_0, i_0)$ of probability higher than $1 - (2k+1) \sum_{j=1}^{k_o-1} (2j+1)^2 \delta_{2j+1,t}$ such that on the event $\Omega(h_0, i_0) \cap \{N_t(h_0, i_0) = k\}$, then the following inequalities hold

$$\hat{\rho}_t^+(h_0, i_0; \delta_{k,t}) - \max_{x \in R(h_0, i_0)} \rho(x) \leq \min_{h_0 \leq h \leq h^+(k)} \left\{ 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^+(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^+(h, k)} \right\}$$

where

$$N_{(h_0, i_0)}^+(h, k) \stackrel{\text{def}}{=} \frac{1}{|I_{h_0, i_0}^+(h)|} \left(k - 2^h [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2} \right).$$

and $h^+(k) \stackrel{\text{def}}{=} \max \left\{ h \geq h_0; N_{(h_0, i_0)}^+(h, k) \geq 1 \right\}.$

Proof: Step 1. Since $I_{h_0, i_0}^+(h)$ and its complementary in $\mathcal{T}(h_0, i_0)$ form a partition of the region $R(h_0, i_0)$, then we first have the property that for all t

$$\sum_{(h, i) \in I_{h_0, i_0}^+(h)} N_t(h, i) + \sum_{(h, i) \in \overline{I_{h_0, i_0}^+(h)} \cap \mathcal{T}(h_0, i_0)} N_t(h, i) = N_t(h_0, i_0).$$

Thus, using the result of Lemma 8 on the $|\mathcal{T}_t(h_0, i_0)| = 2(2k+1) + 1$ nodes of $\mathcal{T}_t(h_0, i_0)$ on the event $\{N_t(h_0, i_0) = k = 2k_o + 1\}$, we deduce that there exists an event Ω of probability higher than $1 - (2k+1) \sum_{j=1}^{k_o-1} (2j+1)^2 \delta_{2j+1, t}$ such that on $\Omega \cap \{N_t(h_0, i_0) = k = 2k_o + 1\}$, then

$$|I_{h_0, i_0}^+(h)| \max_{(h, i) \in I_{h_0, i_0}^+(h)} N_t(h, i) \geq k - (2^{h-h_0} - |I_{h_0, i_0}^+(h)|) [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k, t})/2}]^2 \frac{d \ln(2d/\delta_{k, t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2},$$

and thus, on this event, we have that

$$\max_{(h, i) \in I_{h_0, i_0}^+(h)} N_t(h, i) \geq N_{(h_0, i_0)}^+(h, k) \stackrel{\text{def}}{=} \frac{1}{|I_{h_0, i_0}^+(h)|} \left(k - 2^{h-h_0} [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k, t})/2}]^2 \frac{d \ln(2d/\delta_{k, t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2} \right).$$

Moreover, note that by construction of the tree \mathcal{T}_t , then each leaf of \mathcal{T}_t , and thus of any $\mathcal{T}_t(h, i)$ contains at most one sample. Indeed this holds for the tree \mathcal{T}_0 , now at round $t+1$, if each leaf of \mathcal{T}_t contains at most one sample, let us consider the leaf $(H_{t+1}^{j_{t+1}}, I_{t+1}^{j_{t+1}}) \in \mathcal{T}_t$. If it contains no sample, then it is split in two, one of which is sampled, thus each leaf of \mathcal{T}_{t+1} contains at most one sample. Now if it already contains one sample, it is split in two, and the chosen leaf $(h_{t+1}, i_{t+1}) \in \mathcal{T}_{t+1}$ is the one that contains no sample, since it satisfies that $\hat{\rho}_t^+(h, i) = \infty$ or $\hat{\rho}_t^-(h, i) = -\infty$. Thus each leaf of \mathcal{T}_{t+1} contains at most one sample.

Now, as long as the term on the right-hand side of the previous inequality is positive, this means that there exists one near-optimal node that has been sampled at depth h . Thus, we introduce the following depth defined by

$$h^+(k) \stackrel{\text{def}}{=} \max \left\{ h \geq h_0 ; N_{(h_0, i_0)}^+(h, k) \geq 1 \right\}.$$

Step 2.

For the corresponding node $(h^+(k), i(k))$, we can now consider the path $\{(h, i_h)\}_{h_0 \leq h \leq h^+(k)}$ leading to it from the root (h_0, i_0) at the last round $\tau+1$ when it is sampled. From the previous step, we deduce that on the previous event $\Omega \cap \{N_\tau(h_0, i_0) = k = 2k_o + 1\}$, then

$$\begin{aligned} \hat{\rho}_t^+(h_0, i_0; \delta_{k, t}) &\leq \min \{ \hat{\rho}_t^+(h, i_h; \delta_{k, t}); h_0 \leq h \leq h^+(k) \} \\ &\leq \min \left\{ \max_{x \in R(h, i)} U_t(R(h, i_h), x, \delta_{k, t}); h_0 \leq h \leq h^+(k) \right\}, \end{aligned}$$

from which we deduce, using the property that for $(h, i) \in \mathcal{T}_t(h_0, i_0)$ then we have $\max_{x \in R(h, i)} \rho(x) \leq \max_{x \in R(h_0, i_0)} \rho(x)$, that

$$\hat{\rho}_t^+(h_0, i_0; \delta_{k, t}) - \max_{x \in R(h_0, i_0)} \rho(x) \leq \min_{h_0 \leq h \leq h^+(k)} \left\{ 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k, t})}{2N_{(h_0, i_0)}^+(h, k)}} + \frac{d \ln(2d/\delta_{k, t})}{2N_{(h_0, i_0)}^+(h, k)} \right\}$$

□

The following result provides precise upper and lower bounds on the quantity $\hat{\mathcal{L}}_t(R(h, i))$ defined in the algorithm in Figure 1, depending on the number of times $N_t(R(h, i))$ the region $R(h, i)$ has been sampled up to time t . It is a simple consequence of Proposition 10 and its counter part for $\hat{\rho}_t^-(h_0, i_0; \delta_{k, t})$.

Proposition 3 Let $d_\rho \stackrel{\text{def}}{=} \max \{ \max \{ d^+(h_0, i_0), d^-(h_0, i_0) \}; (h_0, i_0) \in \mathcal{P} \}$ be the biggest near-optimality dimension of ρ over each cell of the partition \mathcal{P} . Under the assumptions of Section 3 and

if moreover $\gamma_1^2 \leq \gamma_2$, then there exists an event Ω of probability higher than $1 - 2\delta$ such that on this event, simultaneously for all $t \leq T$, for all $(h, i) \in \mathcal{T}_t$ we have

$$\hat{\mathcal{L}}_t(R(h_0, i_0)) \geq \mathcal{L}_{R(h_0, i_0)}(N_t(R(h_0, i_0))),$$

and also simultaneously for all $|\mathcal{P}| - 1 \leq t \leq T$, for all $(h, i) \in \mathcal{T}_t$ we have if $N_t(R(h_0, i_0)) \geq k$

$$\hat{\mathcal{L}}_t(R(h_0, i_0)) \leq \mathcal{L}_{R(h_0, i_0)}(k) + \left(\frac{1}{k} + 2\lambda c \gamma^{h_0}\right) B(h_0, k, \delta_{k,t}).$$

where we introduced

$$B(h_0, k, \delta_{k,t}) \stackrel{\text{def}}{=} \min_{h_0 \leq h} \left\{ 2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h} + 2(1 + 2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N(h, k)} \right\}$$

$$\text{where } N(h, k) \stackrel{\text{def}}{=} \frac{1}{C(c_2' \gamma_2^h)^{-d_\rho}} \left(k - 2^{h-h_0} [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h})^2} \right).$$

Proof: Step 1. Union bounds For each region $(h_0, i_0) \in \mathcal{P}$, using a union bound for the max-gap and the min-gap, we look at the quantity

$$(2(2k+1)+1) \sum_{j=1}^{k-1} (2j+1)^2 \delta_{2j+1,t} + (2(2k)+1) \sum_{j=1}^{k-1} (2j)^2 \delta_{2j,t} \leq (2n+1) \sum_{p=1}^n p^2 \delta_{p,t}.$$

Now, note that $N_t(R(h_0, i_0)) \leq t$, so that by a union bound over all $(h_0, i_0) \in \mathcal{P}$ and all $t \leq T$ we finally get that there exists an event Ω of probability higher than $1 - \sum_{t=1}^T \sum_{i=1}^t i^2 \delta_{i,t} (2t+1) |\mathcal{P}|$, then we simultaneously have that for all $(h_0, i_0) \in \mathcal{P}$ and all $t \leq T$, if $N_t(R(h_0, i_0)) \geq k$, then

$$\hat{\rho}_t^+(h_0, i_0; \delta_{N_t(R(h_0, i_0)), t}) - \rho^+(h_0, i_0) \leq \min_{h_0 \leq h \leq h^+(k)} \left\{ 2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h} + 2(1+2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^+(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^+(h, k)} \right\}$$

on the one hand, and on the other hand

$$\rho^-(h_0, i_0) - \hat{\rho}_t^-(h_0, i_0; \delta_{N_t(R(h_0, i_0)), t}) \leq \min_{h_0 \leq h \leq h^-(k)} \left\{ 2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h} + 2(1+2\sqrt{d}) \sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^-(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N_{(h_0, i_0)}^-(h, k)} \right\}$$

where we introduced $\rho^+(h_0, i_0) \stackrel{\text{def}}{=} \max_{x \in R(h_0, i_0)} \rho(x)$ as well as the quantity $\rho^-(h_0, i_0) \stackrel{\text{def}}{=} \min_{x \in R(h_0, i_0)} \rho(x)$ as shorthand notations, and where we used the following two quantities

$$N_{(h_0, i_0)}^+(h, k) \stackrel{\text{def}}{=} \frac{1}{|I_{h_0, i_0}^+(h)|} \left(k - 2^{h-h_0} [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h})^2} \right),$$

$$N_{(h_0, i_0)}^-(h, k) \stackrel{\text{def}}{=} \frac{1}{|I_{h_0, i_0}^-(h)|} \left(k - 2^{h-h_0} [2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2 \gamma_2^h + c_1^2 \gamma_1^{2h})^2} \right).$$

Now, by definition of $\hat{\mathcal{L}}_t(R(h_0, i_0))$, we have the following bound

$$\begin{aligned} \hat{\mathcal{L}}_t(R(h_0, i_0)) - \mathcal{L}_{R(h_0, i_0)}(N_t(R(h_0, i_0))) &= \frac{\hat{\rho}_t^+(h_0, i_0; \delta_{N_t(R(h_0, i_0)), t}) - \rho^+(h_0, i_0)}{N_t(R(h_0, i_0))} \\ &+ \lambda |R(h_0, i_0)| \left(\hat{\rho}_t^+(h_0, i_0; \delta_{N_t(R(h_0, i_0)), t}) - \rho^+(h_0, i_0) + \rho^-(h_0, i_0) - \hat{\rho}_t^-(h_0, i_0; \delta_{N_t(R(h_0, i_0)), t}) \right), \end{aligned}$$

Step 2. Near optimality Let us introduce $d_\rho \stackrel{\text{def}}{=} \max\{\max\{d^+(h_0, i_0), d^-(h_0, i_0)\}; (h_0, i_0) \in \mathcal{P}\}$, the biggest near-optimality dimension over each cell of the partition \mathcal{P} . We now make use of the assumption that $\gamma_1^2 \leq \gamma_2$ that simplifies the bound on the $|I_{h_0, i_0}^+(h)|$. Indeed, in that case it can be shown that for all $d' \geq d_\rho$, there exist a constant C such that

$$\max\{\max\{|I_{h_0, i_0}^+(h)|, |I_{h_0, i_0}^-(h)|\}; (h_0, i_0) \in \mathcal{P}\} \leq C(c_2' \gamma_2^h)^{-d'}$$

Thus, if we naturally introduce the notation

$$B(h_0, k, \delta_{k,t}) \stackrel{\text{def}}{=} \min_{h_0 \leq h} \left\{ 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N(h, k)} \right\}$$

$$\text{where } N(h, k) \stackrel{\text{def}}{=} \frac{1}{C(c_2'\gamma_2^h)^{-d'}} \left(k - 2^{h-h_0}[2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2} \right),$$

then we deduce that there exists an event Ω of probability higher than $1 - \sum_{t=1}^T \sum_{i=1}^t i^2 \delta_{i,t} (2t+1)|\mathcal{P}|$, such that simultaneously for all $(h_0, i_0) \in \mathcal{P}$ and all $t \leq T$, if $N_t(R(h_0, i_0)) \geq k$, then

$$\hat{\mathcal{L}}_t(R(h_0, i_0)) \leq \mathcal{L}_{R(h_0, i_0)}(k) + \left(\frac{1}{k} + 2\lambda c\gamma^{h_0}\right) B(h_0, k, \delta_{k,t}).$$

We conclude by the fact that

$$\sum_{t=1}^T \sum_{i=1}^t i^2 \delta_{i,t} (2t+1)|\mathcal{P}| = \sum_{t=1}^T \frac{6\delta}{\pi^2 t^2} \leq \delta$$

□

B.2 Proof of the main Theorem

Theorem 1 (Regret bound for HORSE.C) *Under the assumptions of Subsection 3, then the regret of the Hierarchical Optimistic Region Selection driven by Curiosity procedure is bounded with probability higher than $1 - 2\delta$ as follows.*

$$\mathfrak{R}_T \leq \sum_{t=|\mathcal{P}|-1}^{T-1} \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{1}{n_{t+1}^*(h_0, i_0)} + 2\lambda c\gamma^{h_0} \right) B(h_0, n_{t+1}^*(h_0, i_0), \delta_{n_{t+1}^*(h_0, i_0), t})$$

where $n_{t+1}^*(h_0, i_0)$ is the optimal allocation at round $t+1$ for the region $(h_0, i_0) \in \mathcal{P}$ and where

$$B(h_0, k, \delta_{k,t}) \stackrel{\text{def}}{=} \min_{h_0 \leq h} \left\{ 2c_2\gamma_2^h + c_1^2\gamma_1^{2h} + 2(1 + 2\sqrt{d})\sqrt{\frac{d \ln(2d/\delta_{k,t})}{2N(h, k)}} + \frac{d \ln(2d/\delta_{k,t})}{2N(h, k)} \right\}$$

in which we have used the following quantity

$$N(h, k) \stackrel{\text{def}}{=} \frac{1}{C(c_2'\gamma_2^h)^{-d_\rho}} \left(k - 2^{h-h_0}[2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2}]^2 \frac{d \ln(2d/\delta_{k,t})}{2(2c_2\gamma_2^h + c_1^2\gamma_1^{2h})^2} \right).$$

Proof: Step 1. The regret of the algorithm at time T is given by definition by

$$\mathfrak{R}_T = \sum_{t=|\mathcal{P}|-1}^{T-1} \max_{R \in \mathcal{P}} \mathcal{L}_R(N_{t+1}(R)) - \max_{R \in \mathcal{P}} \mathcal{L}_R(n_{t+1}^*(R; \mathcal{P})).$$

For each $t+1 \in [1, T]$, let $\bar{R}_{t+1} \stackrel{\text{def}}{=} \operatorname{argmax}\{\mathcal{L}_R(N_{t+1}(R)); R \in \mathcal{P}\}$ be a region of highest loss. Then, either this region is over-pulled or it is under-pulled at the end of round $t+1$. We control the contribution to the regret in the first case in Step 2, and in the second case in Step 3 below.

Step 2 \bar{R}_{t+1} is over-pulled. In that case, then $N_{t+1}(\bar{R}_{t+1}) > n_{t+1}^*(\bar{R}_{t+1}; \mathcal{P})$, thus by definition of the loss, this means that

$$\mathcal{L}_{\bar{R}_{t+1}}(N_{t+1}(\bar{R}_{t+1})) < \mathcal{L}_{\bar{R}_{t+1}}(n_{t+1}^*(\bar{R}_{t+1}; \mathcal{P})) \leq \max_{R \in \mathcal{P}} \mathcal{L}_R(n_{t+1}^*(R; \mathcal{P})),$$

Thus, the contribution to the regret is negative in that case.

Step 3 \bar{R}_{t+1} is under-pulled. In that case, then $N_{t+1}(\bar{R}_{t+1}) < n_{t+1}^*(\bar{R}_{t+1}; \mathcal{P})$, thus by definition of the loss, this means that

$$\mathcal{L}_{\bar{R}_{t+1}}(N_{t+1}(\bar{R}_{t+1})) > \mathcal{L}_{\bar{R}_{t+1}}(n_{t+1}^*(\bar{R}_{t+1}; \mathcal{P})).$$

Since \bar{R}_{t+1} is under-pulled, from the constraint that $\sum_{R \in \mathcal{P}} N_{t+1}(R) = t + 1$, we deduce that there exists a region $R' \in \mathcal{P}$ that is over-pulled, i.e. satisfies $N_{t+1}(R') > n_{t+1}^*(R'; \mathcal{P})$. Let $\tau(R') + 1$ be the last round when a point in R' was pulled; we have precisely $N_{\tau(R')+1}(R') = N_{t+1}(R')$ and $N_{\tau(R')}(R') = N_{t+1}(R') - 1$.

Now at the end of round $t + 1 \geq |\mathcal{P}|$, then all regions R' have been pulled at least once thus satisfies $\tau(R') + 1 \geq 1$. In that case, in the corresponding round $\tau(R') + 1$ we have for all $R \in \mathcal{P}$ first by definition of the algorithm and then on an event Ω of probability higher than $1 - 2\delta$ that

$$\hat{\mathcal{L}}_{\tau(R')}(R') \geq \hat{\mathcal{L}}_{\tau(R')}(R) \geq \mathcal{L}_R(N_{\tau(R')}(R)).$$

In particular, this happens for R being the region \bar{R}_{t+1} . Now, for all $t' \geq \tau(R')$ and all region R we have the property that $l(R, N_{\tau(R')}(R)) \geq l(R, N_{t'}(R))$ and thus we deduce that on Ω

$$\hat{\mathcal{L}}_{\tau(R')}(R') \geq \mathcal{L}_{\bar{R}_{t+1}}(N_{t+1}(\bar{R}_{t+1})).$$

On the other hand, if we now introduce for a region $R = R(h, i)$ the notation

$$\epsilon_{t+1}(R(h, i), n) = \left(\frac{1}{n} + 2\lambda c \gamma^h \right) B(h_0, n, \delta_{n,t}),$$

then on Ω we have the following property

$$\begin{aligned} \hat{\mathcal{L}}_{\tau(R')}(R') &\leq \mathcal{L}_{R'}(N_{\tau(R')}(R')) + \epsilon_{\tau(R')+1}(R', N_{\tau(R')}(R')) \\ &= \mathcal{L}_{R'}(N_{t+1}(R') - 1) + \epsilon_{t+1}(R', N_{t+1}(R') - 1) \\ &\leq \mathcal{L}_{R'}(n_{t+1}^*(R'; \mathcal{P}_i)) + \epsilon_{t+1}(R', n_{t+1}^*(R'; \mathcal{P}_i)), \end{aligned}$$

where we used in the last line the fact that $N_{t+1}(R') \geq n_{t+1}^*(R'; \mathcal{P}_i) + 1$.

By combining the last two inequalities together, we thus proved that on the event Ω , then

$$\mathcal{L}_{\bar{R}_{t+1}}(N_{t+1}(\bar{R}_{t+1})) - \max_{R \in \mathcal{P}_i} \mathcal{L}_R(n_{t+1}^*(R; \mathcal{P}_i)) \leq \max_{R \in \mathcal{P}_i} \epsilon_{t+1}(R, n_{t+1}^*(R; \mathcal{P}_i)). \quad (2)$$

Step 4. Summary So far, by the results of Step 2 and Step 3 we thus have proved that the following upper bound on the regret holds on the event Ω of probability higher than $1 - 2\delta$

$$\mathfrak{R}_T \leq \sum_{t=|\mathcal{P}|-1}^{T-1} \max_{R \in \mathcal{P}} \epsilon_{t+1}(R, n_{t+1}^*(R; \mathcal{P})).$$

We now use the definition of the ϵ_{t+1} and the result of Proposition 3 in order to get

$$\mathfrak{R}_T \leq \sum_{t=|\mathcal{P}|-1}^{T-1} \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{1}{n_{t+1}^*(h_0, i_0)} + 2\lambda c \gamma^{h_0} \right) B(h_0, n_{t+1}^*(h_0, i_0), \delta_{n_{t+1}^*(h_0, i_0), t})$$

□

Corollary 1 Let $\beta \stackrel{\text{def}}{=} 1 + \ln(\max\{2, \gamma_2^{-d_\rho}\})$. Under the assumptions of Theorem 1, assuming that the partition \mathcal{P} of the space \mathcal{X} is well behaved, i.e. that for all $(h_0, i_0) \in \mathcal{P}$, then $n_{t+1}^*(h_0, i_0)$ grows at least as speed $O(\ln(t)(\frac{1}{\gamma_2})^{2h_0\beta})$, then with probability higher than $1 - 2\delta$ we have

$$\mathfrak{R}_T = O\left(\sum_{t=|\mathcal{P}|}^T \max_{(h_0, i_0) \in \mathcal{P}} \left(\frac{1}{n_t^*(h_0, i_0)} + 2\lambda c \gamma^{h_0} \right) \left(\frac{\ln(t)}{n_t^*(h_0, i_0)} \right)^{\frac{1}{2\beta}} \right)$$

Proof: First, note that under the assumption that $\gamma_1^2 \leq \gamma_2$, then

$$B(h_0, k, \delta_{k,t}) \leq \min_{h_0 \leq h} \left\{ (2c_2 + c_1^2) \gamma_2^h + \sqrt{\frac{C_d(k, t)}{N_{h,k}}} \right\}$$

where we introduce for convenience the quantity

$$C_d(k, t) \stackrel{\text{def}}{=} (2 + 4\sqrt{d} + \sqrt{d \ln(2d/\delta_{k,t})/2})^2 d \ln(2d/\delta_{k,t})/2,$$

and where we remind that

$$N_{h,k} \stackrel{\text{def}}{=} \frac{1}{C(c'_2 \gamma_2^h)^{-d_\rho}} \left(k - 2^{h-h_0} \frac{C_d(k, t)}{(2c_2 + c_1^2)^2 \gamma_2^{2h}} \right).$$

With these notations, we note that the inequality $(2c_2 + c_1^2) \gamma_2^h \geq \sqrt{\frac{C_d(k, t)}{N_{h,k}}}$ is equivalent to having the condition,

$$k \frac{(2c_2 + c_1^2)^2}{C_d(k, t)} \gamma_2^{2h} \geq 2^{h-h_0} + C(c'_2 \gamma_2^h)^{-d_\rho},$$

that is satisfied as soon as $h \geq h_0$ is such that

$$k \frac{(2c_2 + c_1^2)^2}{C_d(k, t)} \geq \left(\frac{1}{\gamma_2^2} \max\{2, \gamma_2^{-d_\rho}\} \right)^h \max\{2^{-h_0}, C c_2'^{-d_\rho}\} \quad \text{i.e.} \quad h \leq h^+ \stackrel{\text{def}}{=} \frac{\ln \left(k \frac{(2c_2 + c_1^2)^2}{C_d(k, t) \max\{2^{-h_0}, C c_2'^{-d_\rho}\}} \right)}{\ln \left(\frac{1}{\gamma_2^2} \max\{2, \gamma_2^{-d_\rho}\} \right)}.$$

For such a h^+ , then (up to rounding effects), we thus deduce that the following order holds

$$\begin{aligned} \gamma_2^{h^+} &= \left(k \frac{(2c_2 + c_1^2)^2}{C_d(k, t) \max\{2^{-h_0}, C c_2'^{-d_\rho}\}} \right)^{-\frac{1}{2} \frac{\ln(1/\gamma_2^2)}{\ln \left(\frac{1}{\gamma_2^2} \max\{2, \gamma_2^{-d_\rho}\} \right)}} \\ &= O \left(\left(\frac{\ln(t)}{k} \right)^{\frac{1}{2} \frac{1}{1 + \ln(\max\{2, \gamma_2^{-d_\rho}\})}} \right). \end{aligned}$$

Thus, from the condition that $h^+ \geq h_0$ and the definition of B , we deduce that as soon as

$$n_{t+1}^*(h_0, i_0) \geq c \ln(t) \left(\frac{1}{\gamma_2} \right)^{2h_0(1 + \ln(\max\{2, \gamma_2^{-d_\rho}\}))},$$

for some constant c then we have the following upper bound that concludes the proof

$$B(h_0, n_{t+1}^*(h_0, i_0), \delta_{n_{t+1}^*(h_0, i_0), t}) = O \left(\left(\frac{\ln(t)}{n_{t+1}^*(h_0, i_0)} \right)^{\frac{1}{2} \frac{1}{1 + \ln(\max\{2, \gamma_2^{-d_\rho}\})}} \right).$$

□