



HAL
open science

Non-Asymptotic Sequential Tests for Overlapping Hypotheses and application to near optimal arm identification in bandit models

Aurélien Garivier, Emilie Kaufmann

► **To cite this version:**

Aurélien Garivier, Emilie Kaufmann. Non-Asymptotic Sequential Tests for Overlapping Hypotheses and application to near optimal arm identification in bandit models. *Sequential Analysis*, 2021, 40 (1), pp.61-96. 10.1080/07474946.2021.1847965 . hal-02123833v2

HAL Id: hal-02123833

<https://hal.science/hal-02123833v2>

Submitted on 17 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Non-Asymptotic Sequential Tests for Overlapping Hypotheses Applied to Near Optimal Arm Identification in Bandit Models

Aurélien Garivier¹ and Emilie Kaufmann²

¹ Unité de Mathématiques Pures et Appliquées et Laboratoire de l'Informatique du Parallélisme
École Normale Supérieure de Lyon, Université de Lyon, Lyon, France

² CNRS & Univ. Lille, Inria SequeL, UMR 9189- CRIStAL, Lille, France

Abstract: In this paper, we study sequential testing problems with *overlapping* hypotheses. We first focus on the simple problem of assessing if the mean μ of a Gaussian distribution is smaller or larger than a fixed $\epsilon > 0$; if $\mu \in (-\epsilon, \epsilon)$, both answers are considered to be correct. Then, we consider PAC-best arm identification in a bandit model: given K probability distributions on \mathbb{R} with means μ_1, \dots, μ_K , we derive the asymptotic complexity of identifying, with risk at most δ , an index $I \in \{1, \dots, K\}$ such that $\mu_I \geq \max_i \mu_i - \epsilon$. We provide non asymptotic bounds on the error of a parallel General Likelihood Ratio Test, which can also be used for more general testing problems. We further propose lower bound on the number of observation needed to identify a correct hypothesis. Those lower bounds rely on information-theoretic arguments, and specifically on two versions of a change of measure lemma (a high-level form, and a low-level form) whose relative merits are discussed.

Keywords: Best arm identification; Generalized Likelihood Ratio test; Multi-armed bandits; Sequential testing.

Subject Classifications: 62L10.

1. INTRODUCTION

In order to extract information on a random source \mathbb{P} , the classical statistical framework relies on a fixed sample X_1, \dots, X_n of a given size n . For example, a classical hypothesis test is a (possibly randomized) function $\hat{\nu}_n$ of the sample. In *sequential statistics*, on the contrary, the statistician is allowed to adaptively decide on the sample size. For example, the sequential hypothesis testing of \mathcal{H}_0 versus \mathcal{H}_1 consists of a stopping time τ with respect to the filtration $(\mathcal{F}_t)_t$, where $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$, and a decision rule $\hat{\nu}_\tau$ which is \mathcal{F}_τ -measurable.

Even more power is devoted to the statistician in *active testing*: the statistician not only chooses when to stop, but also sequentially influences the observations. For example, in the *multi-armed bandit problem*, a set \mathcal{A} of unknown random sources is available on demand, and the statistician chooses one of them $A_t \in \mathcal{A}$ at every time step t , from which she gets a new sample. The choice of A_t is made based only upon past observations: A_t is \mathcal{F}_{t-1} -measurable. The goal of the statistician is typically *best-arm identification*, that is to infer which source has the highest expectation using as few samples as possible. Given a risk parameter δ , a strategy (consisting of a sampling, a stopping and a decision rule) is called δ -correct if the probability that the decision $\hat{\nu}_\tau$ is not the source with highest expectation is smaller than δ . The *sample complexity* is the smallest achievable expectation for the stopping time of a δ -correct strategy.

Address correspondence to E. Kaufmann, Inria Lille Nord-Europe, 40 avenue Halley, 59650 Villeneuve d'Ascq, France; E-mail: emilie.kaufmann@univ-lille.fr

In the case of a finite set of sources in a one-dimensional canonical exponential model, the complexity of the best-arm identification problem is well understood since (Garivier and Kaufmann, 2016; Russo, 2016). A non-asymptotic lower bound has been given on the sample complexity: the expectation of the stopping of any δ -correct strategy is always larger than the product of a model-dependent *characteristic time*, which can be computed as the solution of a convex optimization problem depending on the sources, and a confidence term which scales like $\ln(1/\delta)$. This lower-bound is tight: a strategy asymptotically matching it when the risk δ goes to 0 has been described.

But the problem of δ -correct best-arm identification is not really well-posed: the characteristic time of a model tends to infinity when one of the sources has an expectation almost as high as the best one, and δ -correct strategies never stop (with high probability) if two sources are simultaneously optimal. In fact, a much more natural and useful objective is to search for a *Probably Approximately Correct* (PAC) decision: given a risk parameter δ and a tolerance parameter $\epsilon > 0$, a strategy is called (ϵ, δ) -PAC if for every bandit model of the considered class the probability that the chosen source \hat{a}_τ has a mean smaller than the highest one minus ϵ is smaller than δ . Note that a strategy is δ -correct if it is $(0, \delta)$ -PAC: the latter notion is thus a generalization of the former.

The purpose of this paper is to show to what extent the results on the complexity of δ -correct best-arm identification can be generalized to the PAC setting. An interesting feature of the latter is that several decisions may simultaneously be admissible if several source means are within ϵ of their maximum. In other words, this supposes to deal with *sequential testing of overlapping hypotheses*, a surprisingly little investigated problem of interest on its own.

The paper is organized as follows: Section 2 is devoted to the general problem of sequential testing of overlapping hypotheses. A generic sequential test based on generalized likelihood ratios (called the parallel GLRT) is introduced. Section 3 details the simplest example: testing if the mean of a Gaussian distribution is smaller than ϵ or larger than $-\epsilon$, for which we propose a δ -correct parallel GLRT with optimal sample complexity. We show in Section 4 that parallel GLRTs remain δ -correct for very general active identification problems in multi-armed bandit models. The rest of the paper investigates the sample complexity of an important active identification problem: PAC best-arm identification. Section 5 provides our main result on the complexity of PAC best-arm identification. In a nutshell, it appears that a complexity generalizing the results known for δ -correct procedures can be identified, but only asymptotically, under stronger hypothesis, and using less elegant arguments: high-level information-theoretic reasoning cannot be applied here. In Section 6 we then propose an (ϵ, δ) -PAC strategy that combines the parallel GLRT with a “tracking” sampling rule designed for converging to an optimal allocation. This strategy is matching our lower bound for a large family of regular bandit instances.

2. THE PARALLEL GLRT FOR TESTING OVERLAPPING HYPOTHESES

2.1. From Sequential to Active Testing

In a classical parametric hypothesis testing framework, some data D is collected whose distribution depends on some parameter $\mu \in \mathcal{R}$ and is denoted by \mathbb{P}_μ . Given two hypotheses

$$\mathcal{H}_0 : (\mu \in \mathcal{R}_0) \quad \text{and} \quad \mathcal{H}_1 : (\mu \in \mathcal{R}_1) ,$$

where \mathcal{R}_0 and \mathcal{R}_1 are two *disjoint* subsets of \mathcal{R} , the goal is to build a *test* Φ mapping the data to a decision in $\{0, 1\}$, where $\Phi(D) = 1$ means that we *reject* hypothesis \mathcal{H}_0 based on the observed data. Usually, the data collected consist of a fixed number n of i.i.d. samples from some distribution depending on μ , and the two hypotheses do not play symmetric roles: the risk to reject the null hypothesis H_0 when it is true, called the type I error $\alpha := \sup_{\mu \in \mathcal{R}_0} \mathbb{P}_\mu(\Phi(D) = 1)$, is controlled, and the decision $\Phi(D) = 0$ is not interpreted as

an acceptance of \mathcal{H}_0 , but as a absence of rejection. Under a controlled type I error, the procedure aims at minimizing the type II error $\beta := \sup_{\mu \in \mathcal{R}_1} \mathbb{P}_\mu(\Phi(D) = 0)$, that is at maximizing *power* $1 - \beta$ of the test.

In some cases, however, the hypotheses do play symmetric roles and both type I and type II errors have the same importance. When this happens, we really want to *decide for* \mathcal{H}_0 or \mathcal{H}_1 . This distinction is of importance in what follows, and both cases will be illustrated in our approach.

We focus in this paper on *sequential testing for symmetric hypotheses*. In order to emphasize that there is no preferred (null) hypothesis, the hypotheses are denoted by \mathcal{H}_1 and \mathcal{H}_2 . Given a data stream X_1, X_2, \dots and two hypotheses $\mathcal{H}_1 : (\mu \in \mathcal{R}_1)$ and $\mathcal{H}_2 : (\mu \in \mathcal{R}_2)$, the goal is to stop the data collection process after some (random) number τ of observations, and based on this observation $D_\tau = (X_1, \dots, X_\tau)$ make a decision $\hat{i}_\tau = \Phi(D_\tau) \in \{1, 2\}$, where $\hat{i}_\tau = i$ means that hypothesis \mathcal{H}_i is selected. The duration τ of the test should be minimal, but needs to be long enough so that it is actually possible to discriminate between the two hypotheses. This area was pioneered by [Wald \(1945\)](#), who proposes the Sequential Probability Ratio Test (SPRT) for two simple hypotheses, e.g. in a parametric setting $\mathcal{H}_1 : (\mu = \mu_1)$ and $\mathcal{H}_2 : (\mu = \mu_2)$. Among the sequential tests with prescribed type I and type II errors, the SPRT is proved to have the smallest average duration $\mathbb{E}_{\mu_i}[\tau]$ for $i \in \{1, 2\}$. Later, particular examples of sequential test of *composite* hypotheses (in which \mathcal{R}_1 and \mathcal{R}_2 are not reduced to a singleton) have also been studied, see e.g. ([Robbins and Siegmund, 1974](#); [Lai, 1988](#)).

We introduce below in [Section 2.2](#) a broader sequential testing framework, in which we allow for composite hypotheses that are *possibly overlapping*, allowing for $\mathcal{R}_1 \cap \mathcal{R}_2 \neq \emptyset$. Moreover, we allow to simultaneously test more than two hypotheses. In this setup, the goal is to stop the data collection as quickly as possible in order to identify *one* correct hypothesis. We present a generic sequential test for this problem, which is an extension of the classical Generalized Likelihood Ratio test. Then, we illustrate its efficiency in the context of *active testing*, especially for active identification problems in a multi-armed bandit model.

Active testing can be traced back to the early work of [Chernoff \(1959\)](#): at each round, *an experiment is selected and performed* in order to gain information about which of a finite set of hypotheses is true. An active test is thus made of a *sampling rule* (that indicates which experiment A_t is selected at round t) together with the stopping rule and recommendation rules (τ, \hat{i}) inherent to sequential testing. In a multi-armed bandit model, these experiments consists in sampling one of the K available *arms* (probability distributions) ν_1, \dots, ν_K with respective means μ_1, \dots, μ_K . As explained in [Section 4](#) below, we assume here that the arms belong to a one-parameter exponential family, and as such can be parameterized by their means.

Many works on bandit models in the past decade have focus on the reinforcement learning framework, where the samples collected are viewed as rewards and the goal is to maximize cumulative rewards (see, e.g. [Robbins \(1952\)](#); [Lai and Robbins \(1985\)](#); [Bubeck and Cesa-Bianchi \(2012\)](#)). A broader class of objectives is to *learn something* about the unknown parameter $\mu = (\mu_1, \dots, \mu_K)$ and several active testing problems have been studied, including the ϵ -best arm identification problem ([Even-Dar et al., 2006](#)) where the goal is to find any arm whose expectation is within ϵ of the highest mean. We show in [Section 4](#) that the parallel GLRT test provides a valid stopping rule for very general tests in bandit models. We investigate in [Section 6](#) how it can help reaching the optimal sample complexity for ϵ -best arm identification, when coupled with a good sampling rule.

2.2. Sequential Test of Overlapping Hypotheses

We formally define our framework in a rather general setting in which one sequentially collects samples X_1, X_2, \dots and the joint distribution of these samples depends on some parameter $\mu \in \mathcal{R}$. Let $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_M$ be M regions that cover the parameter space (i.e. such that $\bigcup_{i=1}^M \mathcal{R}_i = \mathcal{R}$) but that do not necessarily form a partition of \mathcal{R} . We consider the corresponding M hypotheses

$$\mathcal{H}_1 : (\mu \in \mathcal{R}_1) \quad \mathcal{H}_2 : (\mu \in \mathcal{R}_2) \quad \dots \quad \mathcal{H}_M : (\mu \in \mathcal{R}_M) .$$

We denote by $(\mathcal{F}_t = \sigma(X_1, \dots, X_t))_{t \geq 1}$ the filtration generated by the observations. A sequential test of the hypotheses $\mathcal{H}_1, \dots, \mathcal{H}_M$ based on the stream of samples $(X_s)_{s \in \mathbb{N}}$ is a pair (τ, \hat{i}_τ) where τ is a *stopping rule* indicating when to stop the test (hence τ is a stopping time with respect to $(\mathcal{F}_t)_t$), and $\hat{i}_\tau \in \{1, \dots, M\}$ is a *recommendation rule* that proposes upon stopping a guess for a correct hypothesis based on the observed data (hence \mathcal{F}_τ -measurable).

Definition 2.1. Let $\delta \in (0, 1)$ be a risk parameter. A sequential test $(\tau_\delta, \hat{i}_{\tau_\delta})$ is δ -correct if for all $\mu \in \mathcal{R}$,

$$\mathbb{P}_\mu(\tau_\delta < \infty, \mu \notin \mathcal{R}_{\hat{i}_{\tau_\delta}}) \leq \delta.$$

In words, the probability that the test stops and recommends a wrong hypothesis \mathcal{H}_i (such that $\mu \notin \mathcal{R}_i$) should be under control. For example, in the presence of two overlapping hypotheses, the test should be such that $\mathbb{P}_\mu(\hat{i}_{\tau_\delta} = 2) \leq \delta$ for all $\mu \in \mathcal{R}_1 \setminus \mathcal{R}_2$ and $\mathbb{P}_\mu(\hat{i}_{\tau_\delta} = 1) \leq \delta$ for all $\mu \in \mathcal{R}_2 \setminus \mathcal{R}_1$; however, for all $\mu \in \mathcal{R}_1 \cap \mathcal{R}_2$ the test may decide for either hypothesis (at any time). Among all δ -correct sequential tests, the preferred one is the one that stops after using as few samples as possible: we want τ_δ to be small in expectation for every $\mu \in \mathcal{R}$.

2.3. The Parallel Generalized Likelihood Ratio Test

The likelihood of the observations collected up to round t under the distribution parameter $\lambda \in \mathcal{R}$ is denoted by $\ell(X_1, \dots, X_t; \lambda)$. To decide between the M symmetric hypotheses, we run in parallel M sequential tests of the following two non-overlapping (and asymmetric) hypotheses

$$\tilde{\mathcal{H}}_0 : (\mu \in \mathcal{R} \setminus \mathcal{R}_i) \quad \text{against} \quad \tilde{\mathcal{H}}_1 : (\mu \in \mathcal{R}_i),$$

for each $i \in \{1, \dots, M\}$. We stop when one of these tests rejects hypothesis $\tilde{\mathcal{H}}_0$. The i -th test causing to stop means that μ is believed to belong to \mathcal{R}_i (this does not exclude that it may also belong to other regions), in which case we set $\hat{i}_\tau = i$.

Our test of $\tilde{\mathcal{H}}_0$ against $\tilde{\mathcal{H}}_1$ is a Generalized Likelihood Ratio Test (GLRT). The Generalized Likelihood Ratio statistic based on t samples is defined as

$$\frac{\max_{\lambda \in \mathcal{R}} \ell(X_1, \dots, X_t; \lambda)}{\max_{\lambda \in \mathcal{R} \setminus \mathcal{R}_i} \ell(X_1, \dots, X_t; \lambda)} = \inf_{\lambda \in \mathcal{R} \setminus \mathcal{R}_i} \frac{\ell(X_1, \dots, X_t; \hat{\mu}(t))}{\ell(X_1, \dots, X_t; \lambda)},$$

where $\hat{\mu}(t)$ is the maximum likelihood estimator (in \mathcal{R}). Large values of this GLR statistic tend to reject $\tilde{\mathcal{H}}_0$. Calibrating the rejection threshold for a GLR based on a fixed sample size is often done by resorting to asymptotic arguments (like the Wilks phenomenon (Wilks, 1938)) describing the limit distribution of the GLR under the null hypothesis; this is however not useful for the finite-confidence bandit analysis that follows. We propose to use a threshold function $\beta(t, \delta)$ that depends on the current number of samples t and on the risk parameter δ . We provide in the cases considered here a valid choice for the threshold $\beta(t, \delta)$ in order to ensure δ -correctness of the corresponding test.

The parallel GLRT using the threshold function $\beta(t, \delta)$ is formally defined in the following way. Given $\delta \in (0, 1)$, the stopping rule is

$$\tau_\delta = \inf \left\{ t \in \mathbb{N} : \max_{i=1, \dots, M} \inf_{\lambda \in \mathcal{R} \setminus \mathcal{R}_i} \ln \frac{\ell(X_1, \dots, X_t; \hat{\mu}(t))}{\ell(X_1, \dots, X_t; \lambda)} > \beta(t, \delta) \right\} \quad (2.1)$$

and the decision rule is

$$\hat{i}_{\tau_\delta} \in \operatorname{argmax}_{i=1, \dots, M} \inf_{\lambda \in \mathcal{R} \setminus \mathcal{R}_i} \ln \frac{\ell(X_1, \dots, X_{\tau_\delta}; \hat{\mu}(\tau_\delta))}{\ell(X_1, \dots, X_{\tau_\delta}; \lambda)} \quad (2.2)$$

(ties can be resolved arbitrarily). Note that the maximum over $i \in \{1, \dots, M\}$ in the definition of τ_δ and \hat{i}_{τ_δ} can be reduced to the set of hypotheses to which $\hat{\mu}(t)$ belongs.

In the remainder of the paper we propose different choices of the threshold function that guarantee δ -correctness for the corresponding parallel GLRT. We first study a simple example in Section 3: based on i.i.d. Gaussian samples, the goal is to decide whether the mean is “rather positive” or “rather negative”. We then generalize it in Section 4 where we propose a parallel GLRT which is δ -correct for any active identification problem in a multi-armed bandit model. In Section 5 and 6, we focus on a particular active identification problem of practical interest: ϵ -best arm identification. In this context, we further show that a parallel GLRT is a natural candidate to minimize the sample complexity, when coupled with an appropriate sampling rule.

3. SEQUENTIAL TEST OF OVERLAPPING HYPOTHESES: AN ARCHETYPICAL EXAMPLE

We focus here in this section on a stream of independent samples X_1, X_2, \dots of the distribution $\mathcal{N}(\mu, \sigma^2)$, where the variance σ^2 is known. The statistical problem is to determine whether μ is positive or negative, but with an indifference region of width 2ϵ : we consider the two hypotheses

$$\mathcal{H}_1 : (\mu < \epsilon) \quad \text{and} \quad \mathcal{H}_2 : (\mu > -\epsilon).$$

The hypotheses \mathcal{H}_1 and \mathcal{H}_2 are not mutually exclusive. In this particular example, one aims at building a stopping rule τ_δ and a recommendation rule \hat{i}_{τ_δ} such that

$$\forall \mu \leq -\epsilon, \mathbb{P}_\mu(\hat{i}_{\tau_\delta} = 2) \leq \delta \quad \text{and} \quad \forall \mu \geq \epsilon, \mathbb{P}_\mu(\hat{i}_{\tau_\delta} = 1) \leq \delta, \quad (3.1)$$

but any answer \hat{i}_{τ_δ} is considered correct when $\mu \in \mathcal{R}_1 \cap \mathcal{R}_2 = (-\epsilon, \epsilon)$.

Such a test can be useful when we simultaneously collect samples from two Gaussian populations with means μ_1 and μ_2 in order to decide which of the two has the largest mean, up to an accuracy of ϵ (that is whether, $\mu_1 > \mu_2 - \epsilon$ or $\mu_2 > \mu_1 - \epsilon$). This corresponds to solving the ϵ -best arm identification problem in a two-armed Gaussian bandit model (see Section 4) based on uniformly sampling the two arms (which turns out to be optimal in this very particular case). The reason we start with this simple example is twofolds: it provides a first concrete example of parallel GLRT that is easy to analyse, and it permit to showcase our tools to study the optimal sample complexity of a test.

3.1. Properties of the Parallel GLRT Test

The expression of the parallel GLRT is easy to compute in this Gaussian case. Observing that

$$\ln \frac{\ell(X_1, \dots, X_t; \hat{\mu}(t))}{\ell(X_1, \dots, X_t; \lambda)} = \frac{t(\hat{\mu}_t - \lambda)^2}{2\sigma^2}$$

where $\hat{\mu}_s = \frac{1}{s} \sum_{i=1}^s X_i$ denote the empirical mean of the observations, and that $\mathcal{R} \setminus \mathcal{R}_1 = [\epsilon, +\infty)$ and $\mathcal{R} \setminus \mathcal{R}_2 = (-\infty, -\epsilon]$, the expressions given in (2.1) and (2.2) particularize to

$$\begin{aligned} \tau_\delta &= \inf \left\{ t \in \mathbb{N} : \frac{t(|\hat{\mu}_t| + \epsilon)^2}{2\sigma^2} > \beta(t, \delta) \right\}, \\ \hat{i}_{\tau_\delta} &= 2 \text{ if and only if } (\hat{\mu}_{\tau_\delta} > 0). \end{aligned} \quad (3.2)$$

Correctness The δ -correctness of the parallel GLRT follows from a more general argument developed in Section 4 for more general active identification problems in bandit models (see Remark 4.1 therein). Lemma 4.1 indeed permits to propose a threshold for which the above test is δ -correct:

$$\beta(t, \delta) = 3 \ln(\ln(t) + 1) + \mathcal{T}(\ln(1/\delta)), \quad (3.3)$$

where the function \mathcal{T} is such that $\mathcal{T}(x) \simeq x + \ln(x)$.

Sample Complexity Under the simple data collection process in this particular case, it is possible to give an upper bound on the expected number of samples used by the parallel GLRT that employs the threshold 3.3 before stopping and making a decision.

We first give a crude, asymptotic analysis. Fix $\mu \in \mathbb{R}$ and let $\alpha \in [0, \epsilon)$.

$$\begin{aligned} \mathbb{E}[\tau_\delta] &\leq \sum_{t=1}^{\infty} \mathbb{P}(t(|\hat{\mu}_t| + \epsilon)^2 \leq 2\sigma^2\beta(t, \delta)) \\ &\leq \sum_{t=1}^{\infty} \mathbb{P}(|\hat{\mu}_t - \mu| > \alpha) + \sum_{t=1}^{\infty} \mathbb{P}(t(|\hat{\mu}_t| + \epsilon)^2 \leq 2\sigma^2\beta(t, \delta), |\hat{\mu}_t - \mu| \leq \alpha) \\ &\leq \sum_{t=1}^{\infty} \mathbb{P}(|\hat{\mu}_t - \mu| > \alpha) + \sum_{t=1}^{\infty} \mathbb{P}(t(|\mu| - \alpha + \epsilon)^2 \leq 2\sigma^2\beta(t, \delta), |\hat{\mu}_t - \mu| \leq \alpha). \end{aligned}$$

The first term is upper bounded by a constant (independent of δ), while the second is upper bounded by

$$T_0(\delta) = \inf \{T \in \mathbb{N}^* : \forall t \geq T, t(|\mu| - \alpha + \epsilon)^2 \leq 2\sigma^2\beta(t, \delta)\}.$$

For $\beta(t, \delta)$ as given in (3.3), Lemma A.1 in Appendix A shows that

$$T_0(\delta) = \frac{2\sigma^2}{(|\mu| - \alpha + \epsilon)^2} \ln \frac{1}{\delta} + o_{\delta \rightarrow 0} \left(\ln \frac{1}{\delta} \right).$$

Letting α go to zero, one obtains, for all $\mu \in \mathbb{R}$,

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \leq \frac{2\sigma^2}{(|\mu| + \epsilon)^2}.$$

The following more precise, non-asymptotic statement is proved in Appendix A.

Theorem 3.1. *Introducing the notation $\ell(\delta) = \mathcal{T}(\ln(1/\delta)) + \frac{3}{e}$, the parallel GLRT with threshold function $\beta(t, \delta)$ given in (3.3) satisfies*

$$\mathbb{E}[\tau_\delta] \leq \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \left[\ell(\delta) + \frac{6}{e} \ln \left(\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) \right) + 8 \sqrt{\ell(\delta) + \frac{6}{e} \ln \left(\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) \right) + 32\delta^{1/8}} \right] + 1.$$

To summarize, the parallel GLRT with threshold function (3.3) is δ -correct and the expected number of samples uses satisfies $\mathbb{E}_\mu[\tau_\delta] \lesssim \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ln \left(\frac{1}{\delta} \right)$ for small values of the risk parameter δ , for every $\mu \in \mathbb{R}$, even in the indifference zone $(-\epsilon, \epsilon)$. We prove in the next section that this is the smallest possible number of samples for a δ -correct test.

3.2. Asymptotic Optimality of the Parallel GLRT

In order to prove the optimality of the parallel GLRT, we need to provide a *lower bound* on the sample complexity of any δ -correct algorithm. Let us emphasize that we are not interested in minimax lower bounds (see e.g. [Tsybakov \(2008\)](#)), but in problem-dependent lower bounds providing, for any (reasonable) algorithm, the minimal number of samples necessary to take a decision with the right precision and risk *for each possible value of the parameters*. Such a result is typically obtained by a *change of distribution* argument.

A change of distribution consists in relating the probability of an event C under two different probability distributions. It is well-known that for any event $C \in \mathcal{F}_t$,

$$\begin{aligned} \mathbb{P}_\lambda(C) &= \mathbb{E}_\lambda[\mathbb{1}_C] = \int \mathbb{1}_C(x_1, \dots, x_t) d\mathbb{P}_\lambda^{X_1, \dots, X_t}(x_1, \dots, x_t) \\ &= \int \mathbb{1}_C(x_1, \dots, x_t) \frac{d\mathbb{P}_\lambda^{X_1, \dots, X_t}(x_1, \dots, x_t)}{d\mathbb{P}_\mu^{X_1, \dots, X_t}(x_1, \dots, x_t)} d\mathbb{P}_\mu^{X_1, \dots, X_t}(x_1, \dots, x_t) \\ &= \int \mathbb{1}_C(x_1, \dots, x_t) \frac{\ell(X_1, \dots, X_t; \lambda)}{\ell(X_1, \dots, X_t; \mu)} d\mathbb{P}_\mu^{X_1, \dots, X_t}(x_1, \dots, x_t) \\ &= \mathbb{E}_\mu \left[\mathbb{1}_C \exp(-L_t(\mu, \lambda)) \right], \end{aligned}$$

where $L_t(\mu, \lambda)$ denotes the log-likelihood ratio of the observations:

$$L_t(\mu, \lambda) := \ln \frac{\ell(X_1, \dots, X_t; \mu)}{\ell(X_1, \dots, X_t; \lambda)}.$$

This elementary change of distribution has been a key ingredient for the derivation of lower bounds in the bandit literature. The most notable example is the famous regret lower bound by Lai and Robbins in the seminal paper ([Lai and Robbins, 1985](#)), and further examples include notably ([Audibert et al., 2010](#)).

But this “low-level” ingredient can also serve as a basis for more elaborate arguments, which lead to more elegant and stronger results as for example illustrated in ([Garivier and Kaufmann, 2016](#)). We propose in Lemma 3.1 below two forms of change of distribution that are useful: the low-level form (3.4), and the high-level form (3.5). The high-level form, of information-theoretic flavor, involves the Kullback-Leibler divergence $\text{KL}(\cdot, \cdot)$ from one distribution to another. For any random variable Z defined on a probabilistic space $(\Omega, \mathcal{F}, \mathbb{P})$, we denote by \mathbb{P}^Z the law of Z .

Lemma 3.1. *Let μ and λ be two parameters.*

1. **High-level form:** *for any stopping time τ and any event $C \in \mathcal{F}_\tau$,*

$$\mathbb{E}_\mu[L_\tau(\mu, \lambda)] = \text{KL} \left(\mathbb{P}_\mu^{X_1, \dots, X_\tau}, \mathbb{P}_\lambda^{X_1, \dots, X_\tau} \right) \geq \text{kl}(\mathbb{P}_\mu(C), \mathbb{P}_\lambda(C)), \quad (3.4)$$

where $\text{kl}(x, y) = x \ln(x/y) + (1-x) \ln((1-x)/(1-y))$ is the binary relative entropy.

2. **Low-level form:** *for all $x \in \mathbb{R}$, $n \in \mathbb{N}^*$, for all event $C \in \mathcal{F}_n$,*

$$\mathbb{P}_\lambda(C) \geq e^{-x} \left[\mathbb{P}_\mu(C) - \mathbb{P}_\mu(L_n(\mu, \lambda) \geq x) \right]. \quad (3.5)$$

Proof. The data-processing inequality states that, if P and Q are some probability distributions on the same measurable space \mathcal{X} and if $f : \mathcal{X} \rightarrow \mathcal{Y}$ is a measurable function, then $\text{KL}(P, Q) \geq \text{KL}(P^f, Q^f)$, where P^f (resp Q^f) denotes the push-forward measure of P (resp. Q) by f . It is a simple consequence of Jensen’s

inequality applied to the function $x \mapsto x \log(x)$. The high-level form is a consequence of the data-processing inequality:

$$\mathbb{E}_\mu[L_\tau(\mu, \lambda)] = \text{KL}\left(\mathbb{P}_\mu^{X_1, \dots, X_\tau}, \mathbb{P}_\lambda^{X_1, \dots, X_\tau}\right) \geq \text{KL}\left(\mathbb{P}_\mu^{\mathbb{1}_C}, \mathbb{P}_\lambda^{\mathbb{1}_C}\right) = \text{kl}\left(\mathbb{P}_\mu(C), \mathbb{P}_\lambda(C)\right),$$

see (Garivier et al., 2019) for details. The low-level form is even more elementary:

$$\begin{aligned} \mathbb{P}_\lambda(C) &= \mathbb{E}_\mu\left[\mathbb{1}_C \exp(-L_t(\mu, \lambda))\right] \geq \mathbb{E}_\mu\left[\mathbb{1}_C \mathbb{1}_{(L_t(\mu, \lambda) < x)} \exp(-L_t(\mu, \lambda))\right] \\ &\geq e^{-x} \mathbb{P}_\mu\left(C \cap (L_t(\mu, \lambda) < x)\right) \\ &\geq e^{-x} \left[\mathbb{P}_\mu(C) - \mathbb{P}_\mu(L_n(\mu, \lambda) \geq x)\right]. \end{aligned}$$

□

The sample complexity lower bounds provided below illustrate the strengths and weaknesses of each form, that lead to the two statements in Theorem 3.2 below. This result permits to identify the sample complexity of the testing problem with overlapping hypotheses studied in this section. Indeed, Theorem 3.2 shows that any δ -correct sequential procedure satisfies $\mathbb{E}_\mu[\tau_\delta] \geq \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ln\left(\frac{1}{\delta}\right)$ for every μ , in a regime of small values of δ . In Section 3.1 we exhibited a δ -correct parallel GLRT that uses not more than this number of samples for small values of δ . This proves that the sample complexity of the testing problem, that is the minimal number of samples needed by a δ -correct test, is close to $\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ln\left(\frac{1}{\delta}\right)$ for small values of the risk δ , and justifies that the parallel GLRT, which attains this sample complexity, is asymptotically optimal.

It is notable that the lower bound of Theorem 3.2 is non-asymptotic for all μ that belong to a single hypothesis, which is obtained using the high-level form in Lemma 3.1, but only asymptotic when μ is in the intersection $(-\epsilon, \epsilon)$, which requires the low-level form, as can be seen in the proof.

Theorem 3.2. *Any δ -correct sequential test satisfies*

$$\begin{aligned} &\bullet \forall \mu \notin (-\epsilon, \epsilon), \quad \mathbb{E}_\mu[\tau_\delta] \geq \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \text{kl}(\delta, 1 - \delta), \\ &\bullet \forall \mu \in \mathbb{R}, \quad \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq \frac{2\sigma^2}{(|\mu| + \epsilon)^2}. \end{aligned}$$

Proof. We first treat the case $\mu \notin (-\epsilon, \epsilon)$. In the particular case considered in this section the samples X_1, \dots, X_t are i.i.d. from a Gaussian distribution. Assuming that $\mathbb{E}_\mu[\tau] < \infty$ and using Wald's lemma, the expected log-likelihood ratio takes the simple form

$$\mathbb{E}_\mu[L_\tau(\mu, \lambda)] = \mathbb{E}_\mu[\tau] \text{KL}(\mathcal{N}(\mu, \sigma^2), \mathcal{N}(\lambda, \sigma^2)) = \mathbb{E}_\mu[\tau] \frac{(\mu - \lambda)^2}{2\sigma^2}.$$

Hence, the high-level change of measure inequality (3.4) translates into the following inequality, which obviously also holds when $\mathbb{E}_\mu[\tau] = \infty$:

$$\mathbb{E}_\mu[\tau] \frac{(\mu - \lambda)^2}{2\sigma^2} \geq \text{kl}(\mathbb{P}_\mu(C), \mathbb{P}_\lambda(C)). \quad (3.6)$$

For $\mu < -\epsilon$, choosing $\lambda = \epsilon$ and $C = (\hat{\nu}_{\tau_\delta} = 2)$, which is such that $\mathbb{P}_\mu(C) \leq \delta$ and $\mathbb{P}_\lambda(C) \geq 1 - \delta$, directly yields

$$\mathbb{E}_\mu[\tau] \frac{(|\mu| + \epsilon)^2}{2\sigma^2} \geq \text{kl}(\delta, 1 - \delta).$$

Similarly, for $\mu > \epsilon$, we use $\lambda = -\epsilon$ and $C = (\hat{i}_{\tau_\delta} = 1)$ so as to obtain the same inequality.

We now turn our attention to the (more interesting) case $\mu \in (-\epsilon, \epsilon)$. For every $\beta > 0$, we prove that

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu \left(\tau_\delta \leq \frac{2\sigma^2(1-\beta)}{(|\mu| + \epsilon)^2} \ln(1/\delta) \right) = 0, \quad (3.7)$$

which entails that

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq \frac{2\sigma^2(1-\beta)}{(|\mu| + \epsilon)^2}.$$

Letting β tend to zero gives the second statement of Theorem 3.2 for $\mu \in (-\epsilon, \epsilon)$. Note that this statement is also true for $\mu \notin (-\epsilon, \epsilon)$, as a consequence of the non-asymptotic lower bound obtained above and the fact that $\text{kl}(\delta, 1-\delta) \sim \ln(1/\delta)$ when δ goes to zero.

We now prove (3.7) for a fixed $\beta > 0$. Introducing

$$n_\delta := \left\lfloor \frac{2\sigma^2(1-\beta)}{(|\mu| + \epsilon)^2} \ln \frac{1}{\delta} \right\rfloor,$$

and the event $C_\delta = (\tau_\delta \leq n_\delta)$, we need to prove that $\mathbb{P}_\mu(C_\delta) \rightarrow 0$ when $\delta \rightarrow 0$. To do so, we write

$$\mathbb{P}_\mu(C_\delta) = \mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 1) + \mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 2),$$

and prove that the two terms in the RHS tend to zero when δ goes to zero, using the low-level form of Lemma 3.1 with two different changes of distribution.

$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 1) = 0$ Choosing $\lambda = \epsilon$ yields that $\mathbb{P}_\lambda(C_\delta, \hat{i}_{\tau_\delta} = 1) \leq \mathbb{P}_\lambda(\hat{i}_{\tau_\delta} = 1) \leq \delta$. Since τ_δ is a stopping time, the event $C_\delta \cap (\hat{i}_{\tau_\delta} = 1)$ belongs to \mathcal{F}_{n_δ} . Hence, by Inequality (3.5), for all $x \in \mathbb{R}$,

$$\delta \geq e^{-x} [\mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 1) - \mathbb{P}_\mu(L_{n_\delta}(\mu, \epsilon) \geq x)],$$

which can be rewritten as

$$\mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 1) \leq \delta e^x + \mathbb{P}_\mu(L_{n_\delta}(\mu, \epsilon) \geq x).$$

The choice $x = (1 - \beta/2) \ln(1/\delta)$ yields

$$\mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 1) \leq \delta^{\frac{\beta}{2}} + \mathbb{P}_\mu \left(\frac{L_{n_\delta}(\mu, \epsilon)}{n_\delta} \geq \frac{1 - \beta/2}{1 - \beta} \frac{(|\mu| + \epsilon)^2}{2\sigma^2} \right). \quad (3.8)$$

By the law of large numbers and the fact that $n_\delta \rightarrow +\infty$ when $\delta \rightarrow 0$,

$$\frac{L_{n_\delta}(\mu, \epsilon)}{n_\delta} \xrightarrow{\delta \rightarrow 0} \mathbb{E}_\mu \left[\ln \frac{\ell(X_1; \mu)}{\ell(X_1; \epsilon)} \right] = \text{KL}(\mathbb{P}_\mu^{X_1}, \mathbb{P}_\epsilon^{X_1}) = \frac{(\mu - \epsilon)^2}{2\sigma^2} < \frac{1 - \beta/2}{1 - \beta} \frac{(|\mu| + \epsilon)^2}{2\sigma^2},$$

and the RHS of Equation (3.8) goes to zero.

$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 2) = 0$ Using the exact same reasoning as above, this time with $\lambda = -\epsilon$ which is such that $\mathbb{P}_\lambda(C_\delta, \hat{i}_{\tau_\delta} = 2) \leq \delta$, yields

$$\mathbb{P}_\mu(C_\delta, \hat{i}_{\tau_\delta} = 2) \leq \delta^{\frac{\beta}{2}} + \mathbb{P}_\mu \left(\frac{L_{n_\delta}(\mu, -\epsilon)}{n_\delta} \geq \frac{1 - \beta/2}{1 - \beta} \frac{(|\mu| + \epsilon)^2}{2\sigma^2} \right). \quad (3.9)$$

Writing, thanks to the law of large numbers,

$$\frac{L_{n_\delta}(\mu, -\epsilon)}{n_\delta} \xrightarrow{\delta \rightarrow 0} \frac{(\mu + \epsilon)^2}{2\sigma^2} < \frac{1 - \beta/2}{1 - \beta} \frac{(|\mu| + \epsilon)^2}{2\sigma^2},$$

proves that the RHS in (3.9) goes to zero.

□

4. A δ -CORRECT PARALLEL GLRT IN A BANDIT MODEL

In this section we formally introduce the multi-armed bandit model, under which we can study a broad family of *active* tests of possibly overlapping hypotheses.

In this work, we consider bandit models where all arms belong to a canonical one-parameter exponential family

$$\mathcal{F} = \left\{ \nu_\theta \text{ of density } f_\theta(x) = \exp(\theta x - b(\theta)) \text{ w.r.t. } \xi; \theta \in \Theta \right\},$$

where ξ is some reference measure, $\Theta \subseteq \mathbb{R}$ is an interval of natural parameters and $b : \Theta \rightarrow \mathbb{R}$ is convex and twice differentiable. For appropriate choices of b and ξ , the class \mathcal{F} can be the set of Bernoulli distributions, Gaussian distribution with a known variance, Poisson, Exponential distributions, etc. We refer the reader to (Cappé et al., 2013) for more details about the use of these exponential families for bandit problems. We just recall that distributions in a one-parameter exponential family can be alternatively parameterized by their means $\mu = \dot{b}(\theta)$. Hence, if $\mathcal{I} = \dot{b}(\Theta)$, an exponential bandit model can be parameterized by $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K) \in \mathcal{I}^K$. Letting ν^μ be the unique distribution that has mean $\mu \in \mathcal{I}$, the Kullback-Leibler divergence between the distribution of mean μ and than of mean μ' is

$$d(\mu, \mu') := \text{KL}(\nu^\mu, \nu^{\mu'}) = \mu(\dot{b}^{-1}(\mu) - \dot{b}^{-1}(\mu')) - b(\dot{b}^{-1}(\mu)) + b(\dot{b}^{-1}(\mu')).$$

Remarkable special cases are the family of Gaussian distributions with variance σ^2 , with $d(\mu, \mu') = (\mu - \mu')^2 / (2\sigma^2)$, and the family of Bernoulli distributions associated to the binary relative entropy $d(\mu, \mu') = \mu \ln \frac{\mu}{\mu'} + (1 - \mu) \ln \frac{1-\mu}{1-\mu'}$.

The accumulation of observations in a bandit model is governed by a *sampling rule* which specifies, at each round t , the arm A_t that is selected; a sample X_t from the distribution $\nu^{\mu_{A_t}}$ is subsequently collected. The sampling rule is sequential in that A_t is only allowed to depend on the past observations $A_1, X_1, \dots, A_{t-1}, X_{t-1}$ and possibly on some exogenous (independent) source of randomness.

Fixed sampling rule: A sequential testing problem We consider M regions $\mathcal{R}_1, \dots, \mathcal{R}_M$ covering \mathcal{I}^K : $\mathcal{R} = \bigcup_{i=1}^M \mathcal{R}_i = \mathcal{I}^K$ (but that do not necessarily form a partition). For any given sampling rule, we are interested in building a sequential test $(\tau_\delta, \hat{\nu}_{\tau_\delta})$ that δ -correctly tests the M possibly overlapping hypotheses

$$\mathcal{H}_1 : (\boldsymbol{\mu} \in \mathcal{R}_1) \quad \mathcal{H}_2 : (\boldsymbol{\mu} \in \mathcal{R}_2) \quad \dots \quad \mathcal{H}_M : (\boldsymbol{\mu} \in \mathcal{R}_M) \quad (4.1)$$

based on the stream of observations $A_1, X_1, A_2, X_2, \dots$. We denote by $N_a(t) = \sum_{s=1}^t \mathbb{1}_{(A_s=a)}$ the number of selections of arm a up to round t and use the notation $\hat{\mu}_a(t) = \frac{1}{N_a(t)} \sum_{s=1}^t X_s \mathbb{1}_{(A_s=a)}$ for the empirical mean of those observations.

Letting $\hat{\boldsymbol{\mu}}(t) = (\hat{\mu}_1(t), \dots, \hat{\mu}_K(t))$ be the vector of empirical means of the arms, for every $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K) \in \mathcal{R}$ the log-likelihood ratio can be expressed in the following way:

$$\begin{aligned} \ln \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\ell(X_1, \dots, X_t; \boldsymbol{\lambda})} &= \sum_{a=1}^K N_a(t) \left[\hat{\mu}_a(t) (\dot{b}^{-1}(\hat{\mu}_a(t)) - \dot{b}^{-1}(\lambda_a)) - b(\dot{b}^{-1}(\hat{\mu}_a(t))) + b(\dot{b}^{-1}(\lambda_a)) \right] \\ &= \sum_{a=1}^K N_a(t) d(\hat{\mu}_a(t), \lambda_a). \end{aligned} \quad (4.2)$$

Hence, the parallel GLRT for testing (4.1) is defined by

$$\tau_\delta = \inf \left\{ t \in \mathbb{N} : \max_{a=1, \dots, M} \inf_{\boldsymbol{\lambda} \in \mathcal{R} \setminus \mathcal{R}_a} \sum_{j=1}^K N_j(t) d(\hat{\mu}_j(t), \lambda_j) > \beta(t, \delta) \right\}, \quad (4.3)$$

$$\hat{\nu}_{\tau_\delta} = \operatorname{argmax}_{a=1, \dots, M} \inf_{\boldsymbol{\lambda} \in \mathcal{R} \setminus \mathcal{R}_a} \sum_{j=1}^K N_j(\tau_\delta) d(\hat{\mu}_j(\tau_\delta), \lambda_j).$$

In Lemma 4.1 below, we provide a choice of the threshold function $\beta(t, \delta)$ for which the parallel GLRT is δ -correct, whatever the sampling rule. The definition of β requires to introduce the function

$$\mathcal{T}(x) = 2\tilde{h}\left(\frac{h^{-1}(1+x) + \ln(2\zeta(2))}{2}\right) \quad (4.4)$$

where, for $u \geq 1$, $h(u) = u - \ln u$ and for any $x \geq 0$

$$\tilde{h}(x) = \begin{cases} e^{1/h^{-1}(x)} h^{-1}(x) & \text{if } x \geq h^{-1}(1/\ln(3/2)), \\ (3/2)(x - \ln \ln(3/2)) & \text{otherwise.} \end{cases}$$

The function \mathcal{T} is easy to compute numerically. Its use for the construction of anytime, self-normalized confidence intervals is detailed in (Kaufmann and Koolen, 2018), where the following approximations are derived: $\mathcal{T}(x) \simeq x + 4 \ln(1 + x + \sqrt{2x})$ for $x \geq 5$ and $\mathcal{T}(x) \sim x$ when x is large. It can be noted that in the particular case of Gaussian bandits Lemma 4.1 holds for a slightly smaller function \mathcal{T} that satisfies $\mathcal{T}(x) \simeq x + \ln(x)$.

Lemma 4.1. *For any sampling rule, the parallel GLRT test $(\tau_\delta, \hat{v}_{\tau_\delta})$ using the threshold function*

$$\beta(t, \delta) = 3K \ln(1 + \ln t) + K\mathcal{T}\left(\frac{\ln(1/\delta)}{K}\right)$$

is δ -correct: for all $\boldsymbol{\mu} \in \mathcal{R}$, $\mathbb{P}_{\boldsymbol{\mu}}(\tau_\delta < \infty, \boldsymbol{\mu} \notin \mathcal{R}_{\hat{v}_{\tau_\delta}}) \leq \delta$.

Proof. Let $\boldsymbol{\mu} \in \mathcal{R}$. The probability of error $\mathbb{P}_{\boldsymbol{\mu}}(\tau_\delta < \infty, \boldsymbol{\mu} \notin \mathcal{R}_{\hat{v}_{\tau_\delta}})$ can be upper-bounded by

$$\begin{aligned} & \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t, \exists i : \boldsymbol{\mu} \notin \mathcal{R}_i, \inf_{\boldsymbol{\lambda} \in \mathcal{R} \setminus \mathcal{R}_i} \ln \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\ell(X_1, \dots, X_t; \boldsymbol{\lambda})} > \beta(t, \delta)\right) \\ & \leq \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t, \exists i : \boldsymbol{\mu} \in \mathcal{R} \setminus \mathcal{R}_i, \ln \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\ell(X_1, \dots, X_t; \boldsymbol{\mu})} > \beta(t, \delta)\right) \\ & \leq \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t, \ln \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\ell(X_1, \dots, X_t; \boldsymbol{\mu})} > \beta(t, \delta)\right). \end{aligned}$$

Hence, to prove the δ -correctness of the parallel GLRT, it is sufficient to control the deviations of the quantity $\ln \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\ell(X_1, \dots, X_t; \boldsymbol{\mu})}$ uniformly over time. For this purpose, in the case of an exponential family bandit model, one can use the deviations bounds recently obtained by Kaufmann and Koolen (2018). More precisely, thanks to Equation (4.2),

$$\begin{aligned} \mathbb{P}_{\boldsymbol{\mu}}(\tau_\delta < \infty, \boldsymbol{\mu} \notin \mathcal{R}_{\hat{v}_{\tau_\delta}}) & \leq \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t \in \mathbb{N} : \sum_{a=1}^K N_a(t) d(\hat{\boldsymbol{\mu}}_a(t), \boldsymbol{\mu}_a) > \beta(t, \delta)\right) \\ & \leq \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t \in \mathbb{N} : \sum_{a=1}^K N_a(t) d(\hat{\boldsymbol{\mu}}_a(t), \boldsymbol{\mu}_a) > 3 \sum_{a=1}^K \ln(1 + \ln N_a(t)) + K\mathcal{T}\left(\frac{\ln(1/\delta)}{K}\right)\right), \end{aligned}$$

which is upper bounded by δ according to Theorem 14 in (Kaufmann and Koolen, 2018). In the Gaussian case, Corollary 10 in the same paper justifies the use of a slightly smaller function \mathcal{T} . □

Remark 4.1. The sequential testing problem studied in Section 3 can be viewed as a one-armed bandit problem with a Gaussian arm, for which $d(x, y) = (x - y)^2 / (2\sigma^2)$. It is easy to check that the expression in (4.3) coincides with that in (3.2) for the test studied in Section 3. Therefore, Lemma 4.1 can be applied with $K = 1$ to justify the chosen threshold.

Active testing: Optimizing the test duration Lemma 4.1 provides a universal choice of threshold for which the parallel GLRT is δ -correct, whatever the sampling rule. Obviously, however, in order to minimize the number of samples τ_δ needed before stopping the test, the sampling rule also plays a crucial role. For example, if for $K = 3$ the question is whether $\mu_2 > \mu_3$, there is no point in sampling from ν_1 .

In an *active* identification problem, the goal is precisely to optimize the stopping rule *and* the sampling rule in order to reach the minimal number of samples, while preserving the δ -correctness of the decision. The sample complexity of some active identification problems has been exhibited when the different regions form a *partition* (Garivier and Kaufmann, 2016; Kaufmann et al., 2018; Kaufmann and Koolen, 2018; Juneja and Krishnasamy, 2019) but the case of overlapping hypotheses is known to be more complex (Degegne and Koolen, 2019).

In the rest of the paper, we focus on a particular example of active identification: ϵ -Best Arm Identification for which we investigate the optimal sample complexity. As we show below, a *lower-bound* analysis of the sample complexity indirectly leads to an efficient sampling rule.

5. SAMPLE COMPLEXITY LOWER BOUNDS FOR ACTIVE TESTING: ϵ -BEST ARM IDENTIFICATION

ϵ -Best-Arm Identification (BAI) was first introduced by Mannor and Tsitsiklis (2004); Even-Dar et al. (2006) in the bandit literature. Given a risk parameter $\delta \in (0, 1]$ and an accuracy parameter $\epsilon > 0$, the goal is to find as quickly as possible an element of the set $\mathcal{A}_\epsilon(\boldsymbol{\mu}) := \{a : \mu_a \geq \max_i \mu_i - \epsilon\}$ of ϵ -optimal arms, with probability at least $1 - \delta$, by adaptively sampling the arms. A slightly different formulation with an *indifference zone* has been studied in the ranking and selection community (see, e.g., Kim and Nelson (2001)) in which correctness guarantees are provided only when there is a unique ϵ -best arm (i.e. for $\boldsymbol{\mu}$ such that $|\mathcal{A}_\epsilon(\boldsymbol{\mu})| = 1$). In contrast, we highlight that in ϵ -BAI, in the presence of multiple ϵ -best arms (i.e., overlapping hypotheses), our goal is to identify one of those.

A possible motivation to study ϵ -best arm identification stems from its application to A/B/C Testing: In this context, each arm models some version of a company's website. For each visitor, the company adaptively selects one of the versions and checks whether a conversion occurs. The goal is to identify one version with the highest probability of conversion. In the presence of many near optimal versions, it makes sense to relax the problem of finding a near-optimal version, as it leads to shorter tests.

ϵ -best arm identification can be regarded as a particular example of *active testing* of the hypotheses

$$\mathcal{H}_1 = \left(\mu_1 > \max_i \mu_i - \epsilon \right) \quad \dots \quad \mathcal{H}_K = \left(\mu_K > \max_i \mu_i - \epsilon \right)$$

based on adaptively sampling the marginals of $f_{\boldsymbol{\mu}} = f_{\mu_1} \otimes \dots \otimes f_{\mu_K}$. A strategy for ϵ -BAI not only consists in a sequential test for those hypotheses, that is a pair (τ, \hat{i}) of stopping and recommendation rule, but also on a sampling rule $(A_t)_t$ indicating which arm is selected at round t . As explained in Section 4, we consider the case where the arms belongs to a one-dimension canonical exponential family and their distribution is therefore parameterized by their means that belong to some interval $\mathcal{I} = (\mu^-, \mu^+)$.

Definition 5.1. A (ϵ, δ) -PAC strategy for ϵ -best arm identification is a triple $(A_t^\delta, \tau_\delta, \hat{i}_{\tau_\delta})$, where A_t^δ is \mathcal{F}_{t-1} -measurable for each $t \geq 1$, τ is a stopping time relative to the filtration $(\mathcal{F}_t)_t$, and \hat{i}_{τ_δ} is a $\mathcal{F}_{\tau_\delta}$ -measurable random variable, such that

$$\forall \boldsymbol{\mu} \in \mathcal{I}^K, \quad \mathbb{P}_{\boldsymbol{\mu}}(\tau_\delta < \infty, \hat{i}_{\tau_\delta} \notin \mathcal{A}_\epsilon(\boldsymbol{\mu})) \leq \delta.$$

An ϵ -PAC strategy is a family of strategies $\pi_\delta = (A_t^\delta, \tau_\delta, \hat{i}_{\tau_\delta})$ such that, for all δ , π_δ is (ϵ, δ) -PAC.

The goal in ϵ -BAI is to build an (ϵ, δ) -PAC strategy such that on every bandit instance $\boldsymbol{\mu}$, the expected number of samples needed before stopping $\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]$, is as close as possible to the *sample complexity*, which is the minimal number of samples needed by any (ϵ, δ) -PAC algorithm before stopping. In this section, we provide lower bounds on the sample complexity of ϵ -Best Arm Identification.

5.1. The Characteristic Time of ϵ -Best Arm Identification

All the lower bounds that we propose feature a quantity $T_\epsilon^*(\boldsymbol{\mu})$ which we term the characteristic time of $\boldsymbol{\mu}$.

Definition 5.2. The characteristic time $T_\epsilon^*(\boldsymbol{\mu})$ is defined by:

$$T_\epsilon^*(\boldsymbol{\mu})^{-1} = \sup_{w \in \Sigma_K} \max_{a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})} \min_{b \neq a} \inf_{(\lambda_a, \lambda_b): \lambda_a \leq \lambda_b - \epsilon} \left[w_a d(\mu_a, \lambda_a) + w_b d(\mu_b, \lambda_b) \right],$$

where $\Sigma_K = \{w \in [0, 1]^K : \sum_i w_i = 1\}$ is the set of probability vectors.

It can be observed that the quantity $T_0^*(\boldsymbol{\mu})$ coincides with the quantity $T^*(\boldsymbol{\mu})$ introduced by [Garivier and Kaufmann \(2016\)](#) as the characteristic time for the best arm identification problem.

In order to illustrate the dependency of $T_\epsilon^*(\boldsymbol{\mu})$ on the means $\boldsymbol{\mu}$, we first consider the Gaussian case. Introducing the notation $\boldsymbol{\mu}^{a,\epsilon}$ for the bandit instance such that $\mu_i^{a,\epsilon} = \mu_i$ for all $i \neq a$ and $\mu_a^{a,\epsilon} = \mu_a + \epsilon$, it can be shown (see [Appendix B.1](#)) that

$$T_\epsilon^*(\boldsymbol{\mu}) = \min_{a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})} T_0^*(\boldsymbol{\mu}^{a,\epsilon}) = T_0^*(\boldsymbol{\mu}^{a^*,\epsilon}), \quad (5.1)$$

where $a^* \in \operatorname{argmax}_a \mu_a$ is an optimal arm. Hence it is sufficient to be able to compute the characteristic time for $\epsilon = 0$, for which [Garivier and Kaufmann \(2016\)](#) provide an efficient algorithm. Besides, [Garivier and Kaufmann \(2016\)](#) propose an approximation for $T_0^*(\boldsymbol{\mu})$ from which it directly follows that

$$\sum_{a \neq a^*}^K \frac{2\sigma^2}{(\mu_{a^*} + \epsilon - \mu_a)^2} \leq T_\epsilon^*(\boldsymbol{\mu}) \leq 2 \times \sum_{a \neq a^*}^K \frac{2\sigma^2}{(\mu_{a^*} + \epsilon - \mu_a)^2}.$$

Finally, there is a simple expression in the two armed-case: $T_\epsilon^*(\boldsymbol{\mu}) = \frac{8\sigma^2}{(|\mu_1 - \mu_2| + \epsilon)^2}$. Explicit expressions for general (non-Gaussian) two-armed bandits can be found in [Section 5.4](#). However, beyond 2 arms there is no closed-form expression for the characteristic time $T_\epsilon^*(\boldsymbol{\mu})$ and we refer to [Section 6](#) for a discussion on its efficient computation.

5.2. A First Non-Asymptotic Lower Bound

To illustrate the difficulty of extending the lower bound technique of [Garivier and Kaufmann \(2016\)](#) to ϵ -BAI, we begin with a sub-optimal but simple sample complexity lower bound, obtained using the high-level change of distribution argument presented in [Lemma 3.1](#).

Theorem 5.1. For any (ϵ, δ) -PAC strategy, one has

$$\mathbb{E}_\boldsymbol{\mu}[\tau_\delta] \geq T_\epsilon^*(\boldsymbol{\mu}) \left[\frac{1 - \delta}{|\mathcal{A}_\epsilon(\boldsymbol{\mu})|} \ln\left(\frac{1}{\delta}\right) - \ln(2) \right].$$

Proof. Using the δ -PAC property and the pigeonhole principle, there exists $a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})$ such that

$$\mathbb{P}_\boldsymbol{\mu}(\hat{i}_{\tau_\delta} = a) \geq \frac{1 - \delta}{|\mathcal{A}_\epsilon(\boldsymbol{\mu})|}.$$

By definition of $T_\epsilon^*(\boldsymbol{\mu})$, for the above choice of $a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})$, there exists $b \neq a$ and $\bar{\lambda}$ such that $\bar{\lambda}_a < \bar{\lambda}_b - \epsilon$:

$$T_\epsilon^*(\boldsymbol{\mu})^{-1} \geq \frac{\mathbb{E}_\boldsymbol{\mu}[N_a(\tau_\delta)]}{\mathbb{E}_\boldsymbol{\mu}[\tau_\delta]} d(\mu_a, \bar{\lambda}_a) + \frac{\mathbb{E}_\boldsymbol{\mu}[N_b(\tau_\delta)]}{\mathbb{E}_\boldsymbol{\mu}[\tau_\delta]} d(\mu_b, \bar{\lambda}_b). \quad (5.2)$$

Indeed, fixing $\bar{w}_a := \mathbb{E}_\mu[N_a(\tau_\delta)]/\mathbb{E}_\mu[\tau_\delta]$ and defining

$$F_a(w, \boldsymbol{\mu}) = \min_{b \neq a} \inf_{\lambda_a < \lambda_b - \epsilon} [w_a d(\mu_a, \lambda_a) + w_b d(\mu_b, \lambda_b)],$$

observe that

$$\sup_{w \in \Sigma_K} F_a(w, \boldsymbol{\mu}) \geq F_a(\bar{w}, \boldsymbol{\mu}).$$

Letting $b, \bar{\lambda}_a, \bar{\lambda}_b$ be the minimizer in the definition of $F_a(\bar{w}, \boldsymbol{\mu})$ yields

$$\max_{a \in \mathcal{A}_\epsilon} \sup_{w \in \Sigma_K} F_a(w, \boldsymbol{\mu}) \geq \bar{w}_a d(\mu_a, \bar{\lambda}_a) + \bar{w}_b d(\mu_b, \bar{\lambda}_b),$$

and the left hand side of this inequality is $T_\epsilon^{-1}(\boldsymbol{\mu})$.

We now assume $\mathbb{E}[\tau_\delta] < \infty$ (otherwise any lower bound on the number of samples of that algorithm obviously holds). Using Wald's Lemma, the expression of the expected log-likelihood ratio between $\boldsymbol{\mu}$ and $\bar{\boldsymbol{\lambda}}$ is

$$\mathbb{E}_\mu [L_{\tau_\delta}(\boldsymbol{\mu}, \bar{\boldsymbol{\lambda}})] = \mathbb{E}_\mu [N_a(\tau_\delta)] d(\mu_a, \bar{\lambda}_a) + \mathbb{E}_\mu [N_b(\tau_\delta)] d(\mu_b, \bar{\lambda}_b).$$

Therefore, one can use Inequality (5.2) and the high-level form of Lemma 3.1 to write

$$\begin{aligned} \frac{\mathbb{E}_\mu[\tau_\delta]}{T_\epsilon^*(\boldsymbol{\mu})} &\geq \mathbb{E}_\mu [L_{\tau_\delta}(\boldsymbol{\mu}, \bar{\boldsymbol{\lambda}})] \geq \text{kl}(\mathbb{P}_\mu(\hat{i}_{\tau_\delta} = a), \mathbb{P}_{\bar{\boldsymbol{\lambda}}}(\hat{i}_{\tau_\delta} = a)) \\ &\geq \text{kl}\left(\frac{1-\delta}{|\mathcal{A}_\epsilon(\boldsymbol{\mu})|}, \delta\right) \geq \frac{1-\delta}{|\mathcal{A}_\epsilon|} \ln\left(\frac{1}{\delta}\right) - \ln(2), \end{aligned}$$

where the last inequality follows from the inequality $\text{kl}(x, y) \geq x \ln(1/y) - \ln(2)$. □

5.3. A Tighter Lower Bound for Converging Strategies

The lower bound obtained in Theorem 5.1 could incorrectly suggest that the sample complexity would be of order $\frac{T_\epsilon^*(\boldsymbol{\mu})}{|\mathcal{A}_\epsilon(\boldsymbol{\mu})|} \ln\left(\frac{1}{\delta}\right)$. As we will see, the factor $|\mathcal{A}_\epsilon(\boldsymbol{\mu})|$ in the denominator is not correct: this bound is tight only for values of $\boldsymbol{\mu}$ with a unique ϵ -best arm. We now provide a tighter lower bound, which comes with two shortcomings. First, it is asymptotic: it relies on the low-level form of change of distribution presented in Lemma 3.1. Second, it requires an additional assumption on the sampling rule, that needs to be *converging*.

Definition 5.3. A sampling rule A_t^δ is said to be *converging* if it does not depend on δ (that is, $A_t^\delta = A_t$) and for every $\boldsymbol{\mu}$, there exists $\bar{w}(\boldsymbol{\mu}) \in \Sigma_K$ such that, \mathbb{P}_μ -almost surely,

$$\forall a = 1, \dots, K, \quad \frac{N_a(t)}{t} \xrightarrow[t \rightarrow \infty]{} \bar{w}_a(\boldsymbol{\mu}).$$

Theorem 5.2. For any ϵ -PAC family of converging strategies, one has

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_\epsilon^*(\boldsymbol{\mu}).$$

Proof. To prove Theorem 5.2, it is sufficient to prove that for every $\beta > 0$,

$$\lim_{\delta \rightarrow 0} \mathbb{P}\left(\tau_\delta \leq (1-\beta)T_\epsilon^*(\boldsymbol{\mu}) \ln\left(\frac{1}{\delta}\right)\right) = 0. \quad (5.3)$$

Indeed, the conclusion follows by Markov inequality and letting β go to zero. We now fix $\beta > 0$ and introduce

$$n_\delta := \left\lceil (1 - \beta) T_\epsilon^*(\boldsymbol{\mu}) \ln \left(\frac{1}{\delta} \right) \right\rceil .$$

Defining $\mathcal{C}_\delta = (\tau_\delta \leq n_\delta)$, we establish that $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu(\mathcal{C}_\delta) = 0$.

First, using that the strategy is δ -correct, one can write

$$\begin{aligned} \mathbb{P}_\mu(\mathcal{C}_\delta) &= \mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} \in \mathcal{A}_\epsilon(\boldsymbol{\mu})) + \mathbb{P}(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} \notin \mathcal{A}_\epsilon(\boldsymbol{\mu})) \\ &\leq \sum_{a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})} \mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) + \delta . \end{aligned}$$

To conclude, it is sufficient to prove that $\forall a \in \mathcal{A}_\epsilon(\boldsymbol{\mu}), \lim_{\delta \rightarrow 0} \mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) = 0$.

Fix $a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})$. Just like in the proof of Theorem 5.1, we use that by definition of $T_\epsilon^*(\boldsymbol{\mu})$, there exists $b \neq a$ and an alternative model $\boldsymbol{\lambda}$ in which only arm a and b are changed and such that $\lambda_a < \lambda_b - \epsilon$ that satisfies

$$T_\epsilon^*(\boldsymbol{\mu})^{-1} \geq \bar{w}_a(\boldsymbol{\mu}) d(\mu_a, \lambda_a) + \bar{w}_b(\boldsymbol{\mu}) d(\mu_a, \lambda_b) , \quad (5.4)$$

where we recall that the weights vector $\bar{w}(\boldsymbol{\mu})$ contains the limit fraction of samples allocated to each arm by the converging sampling rule (A_t) : for all i , it holds that $\mathbb{P}_\mu(\lim_{t \rightarrow \infty} N_i(t)/t = \bar{w}_i(\boldsymbol{\mu})) = 1$.

As arm a is not ϵ -optimal in the bandit instance $\boldsymbol{\lambda}$, one has $\mathbb{P}_\lambda(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) \leq \mathbb{P}_\lambda(\hat{i}_{\tau_\delta} = a) \leq \delta$. Moreover, the event $\mathcal{C}_\delta \cap (\hat{i}_{\tau_\delta} = a)$ belongs to \mathcal{F}_{n_δ} . Hence, using the low-level Inequality (3.5) in Lemma 3.1, for all $x \in \mathbb{R}$,

$$\delta \geq e^{-x} [\mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) - \mathbb{P}_\mu(L_{n_\delta}(\boldsymbol{\mu}, \boldsymbol{\lambda}) \geq x)]$$

which can be rewritten as

$$\mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) \leq \delta e^x + \mathbb{P}_\mu(L_{n_\delta}(\boldsymbol{\mu}, \boldsymbol{\lambda}) \geq x) .$$

The choice $x = (1 - \beta/2) \ln(1/\delta)$ yields

$$\mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) \leq \delta^{\frac{\beta}{2}} + \mathbb{P}_\mu \left(\frac{L_{n_\delta}(\boldsymbol{\mu}, \boldsymbol{\lambda})}{n_\delta} \geq \frac{1 - \beta/2}{1 - \beta} T_\epsilon^*(\boldsymbol{\mu})^{-1} \right) . \quad (5.5)$$

Because of the converging assumption, it holds \mathbb{P}_μ -a.s. (independently of δ) that

$$\frac{L_t(\boldsymbol{\mu}, \boldsymbol{\lambda}_a)}{t} \xrightarrow[t \rightarrow \infty]{} \bar{w}_a(\boldsymbol{\mu}) d(\mu_a, \lambda_a) + \bar{w}_b(\boldsymbol{\mu}) d(\mu_b, \lambda_b) \leq T_\epsilon^*(\boldsymbol{\mu})^{-1} < \frac{1 - \beta/2}{1 - \beta} T_\epsilon^*(\boldsymbol{\mu})^{-1} ,$$

where the first inequality holds by definition of $\boldsymbol{\lambda}$, which satisfies (5.4). As $n_\delta \rightarrow \infty$ when $\delta \rightarrow 0$, it follows that

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu \left(\frac{L_{n_\delta}(\boldsymbol{\mu}, \boldsymbol{\lambda})}{n_\delta} \geq \frac{1 - \beta/2}{1 - \beta} T_\epsilon^*(\boldsymbol{\mu})^{-1} \right) = 0$$

and $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu(\mathcal{C}_\delta, \hat{i}_{\tau_\delta} = a) = 0$, which concludes the proof. \square

5.4. An improved lower bound for two arms

ϵ -best arm identification in a two-armed bandit model $\boldsymbol{\mu} = (\mu_1, \mu_2)$ coincides with an active testing procedure to select one of the hypotheses

$$\mathcal{H}_1 : (\mu_1 > \mu_2 - \epsilon) \quad \text{or} \quad \mathcal{H}_2 : (\mu_2 > \mu_1 - \epsilon)$$

and an (ϵ, δ) -PAC strategy should satisfy $\mathbb{P}(\hat{i}_{\tau_\delta} = 2) \leq \delta$ for all $\boldsymbol{\mu}$ such that $\mu_1 > \mu_2 + \epsilon$ and $\mathbb{P}(\hat{i}_{\tau_\delta} = 1) \leq \delta$ for all $\boldsymbol{\mu}$ such that $\mu_2 > \mu_1 + \epsilon$. In that case, a first observation is that the complexity term $T_\epsilon^*(\boldsymbol{\mu})$ can be made slightly more explicit. The proof of Lemma 5.1 can be found in Appendix B.2.

Lemma 5.1. *If $\mu_a \geq \mu_b - \epsilon$, define $\mu_\epsilon^*(\mu_a, \mu_b)$ as the unique solution in $\lambda \in (\mu^-, \mu^+ - \epsilon)$ to*

$$d(\mu_a, \lambda) = d(\mu_b, \lambda + \epsilon).$$

Then

$$T_\epsilon^*(\boldsymbol{\mu})^{-1} := \begin{cases} d(\mu_1, \mu_\epsilon^*(\mu_1, \mu_2)) & \text{if } \mu_1 > \mu_2 + \epsilon, \\ \max[d(\mu_1, \mu_\epsilon^*(\mu_1, \mu_2)), d(\mu_2, \mu_\epsilon^*(\mu_2, \mu_1))] & \text{if } \mu_1 - \mu_2 \in (-\epsilon, \epsilon), \\ d(\mu_2, \mu_\epsilon^*(\mu_2, \mu_1)) & \text{if } \mu_2 > \mu_1 + \epsilon. \end{cases}$$

Note that the quantity $\mu_\epsilon^*(\mu_a, \mu_b)$ defined in Lemma 5.1 is indeed well-defined when $\mu_a \geq \mu_b - \epsilon$, as the mapping $g : \mu \mapsto d(\mu_a, \mu) - d(\mu_b, \mu + \epsilon)$ is decreasing on $[\mu_b - \epsilon, \mu_a]$ and satisfies $g(\mu_b - \epsilon) = d(\mu_a, \mu_b - \epsilon) > 0$ and $g(\mu_a) = -d(\mu_b, \mu_a + \epsilon) < 0$.

For two armed bandits, we provide below an asymptotic sample complexity lower bound that relax the converging assumption required for Theorem 5.2. Theorem 5.3 given below still requires the sampling rule A_t^δ to be independent from δ (we write A_t instead of A_t^δ), but doesn't need the fraction of selection of each arm to converge to a certain proportion under that sampling rule. Its proof is given in Appendix C.

Theorem 5.3. *Every ϵ -PAC family of anytime strategies (i.e., for which the sampling rule (A_t) is independent from δ) satisfies, for all $\boldsymbol{\mu}$,*

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\ln(1/\delta)} \geq T_\epsilon^*(\boldsymbol{\mu}).$$

6. ASYMPTOTICALLY OPTIMAL STRATEGIES FOR ϵ -BEST ARM IDENTIFICATION

The strategy described in this section hinges on the same ideas than the Track-and-Stop strategy proposed by Garivier and Kaufmann (2016) for the case $\epsilon = 0$. The sampling rule builds on the knowledge of ‘‘optimal weights’’ under which the arms should be sampled, whereas the stopping is a parallel GLRT test with an appropriate choice of threshold. We prove that the resulting strategy, called ϵ -Track-and-Stop is asymptotically optimal for some instances of ϵ -best arm identification.

6.1. Optimal Weights and their Computation

From the proof of our lower bounds, it appears that a converging strategy matching the lower bound should have its empirical proportions of draws converge to a vector that belongs to the set

$$\mathcal{W}_\epsilon^*(\boldsymbol{\mu}) = \operatorname{argmax}_{w \in \Sigma_K} \max_{a \in \mathcal{A}_\epsilon} \min_{b \neq a} \inf_{(\lambda_a, \lambda_b) : \lambda_a \leq \lambda_b - \epsilon} [w_a d(\mu_a, \lambda_a) + w_b d(\mu_b, \lambda_b)].$$

To ease the notation, we assume that μ_1 is an optimal arm, i.e. $\mu_1 \geq \mu_b$ for all $b \in \{1, \dots, K\}$. We define

$$T_\epsilon^{*,a}(\boldsymbol{\mu})^{-1} := \sup_{w \in \Sigma_K} \min_{b \neq a} \inf_{(\lambda_a, \lambda_b) : \lambda_a \leq \lambda_b - \epsilon} [w_a d(\mu_a, \lambda_a) + w_b d(\mu_b, \lambda_b)].$$

Observe that for a such that there exists $b \neq a$ with $\mu_a = \mu_b - \epsilon$, then $T_\epsilon^{*,a}(\boldsymbol{\mu}) = +\infty$. For a such that $\mu_a > \mu_b - \epsilon$ for all $b \neq a$, we argue below that the argmax in w is unique and define

$$w_\epsilon^{*,a}(\boldsymbol{\mu}) := \operatorname{argmax}_{w \in \Sigma_K} \min_{b \neq a} \inf_{(\lambda_a, \lambda_b): \lambda_a \leq \lambda_b - \epsilon} [w_a d(\mu_a, \lambda_a) + w_b d(\mu_b, \lambda_b)] . \quad (6.1)$$

With this notation, for $\epsilon > 0$, one has

$$\mathcal{W}_\epsilon^*(\boldsymbol{\mu}) = \left\{ w_\epsilon^{*,a}(\boldsymbol{\mu}) \text{ for } a \in \operatorname{argmin}_{i \in \mathcal{A}_\epsilon(\boldsymbol{\mu}): \mu_i > \max_{j \neq i} \mu_j - \epsilon} T_\epsilon^{*,i}(\boldsymbol{\mu}) \right\} . \quad (6.2)$$

For $\epsilon > 0$, this set is always non-empty as arm 1 satisfies $\mu_1 > \min_{j \neq 1} \mu_j - \epsilon$. For $\epsilon = 0$, in the presence of a unique optimal arm, $\mathcal{W}_\epsilon(\boldsymbol{\mu}) = \{w_0^{*,1}(\boldsymbol{\mu})\}$, but in the presence of multiple arms with mean μ_1 , the optimal weights are not well defined. In this degenerate case, we define $\mathcal{W}_\epsilon(\boldsymbol{\mu}) = \{|\mathcal{A}_0(\boldsymbol{\mu})|^{-1} \mathbb{1}_{\mathcal{A}_0(\boldsymbol{\mu})}\}$.

We now fix $a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})$ such that $\mu_a > \max_{b \neq a} \mu_b - \epsilon$ and explain why $w_\epsilon^{*,a}(\boldsymbol{\mu})$ is well-defined and how to compute this probability vector. Introducing the interval

$$\mathcal{I}_{a,b}^\epsilon = [\mu^- \vee (\mu_b - \epsilon), \mu_a \wedge (\mu^+ - \epsilon)] ,$$

it follows from Lemma B.1 stated in Appendix B that

$$\begin{aligned} w_\epsilon^{*,a}(\boldsymbol{\mu}) &= \operatorname{argmax}_{w \in \Sigma_K} \min_{b \neq a} \inf_{\lambda \in \mathcal{I}_{a,b}^\epsilon} [w_a d(\mu_a, \lambda) + w_b d(\mu_b, \lambda + \epsilon)] \\ &= \operatorname{argmax}_{w \in \Sigma_K} \min_{b \neq a} w_a g_b \left(\frac{w_b}{w_a} \right) \end{aligned}$$

where for $b \neq a$ the application g_b (also depending on a , $\boldsymbol{\mu}$ and ϵ) is defined by

$$g_b(x) = \inf_{\lambda \in \mathcal{I}_{a,b}^\epsilon} [d(\mu_a, \lambda) + x d(\mu_b, \lambda + \epsilon)] . \quad (6.3)$$

In the particular case in which $\mu_a > \mu^+ - \epsilon$ and $\mu_b = \mu^+$, the interval $\mathcal{I}_{a,b}^\epsilon$ is reduced to the point $\mu^+ - \epsilon$ and the mapping $g_b(x) = d(\mu_a, \mu^+ - \epsilon)$ is constant. Besides this degenerate case, $\mathcal{I}_{a,b}^\epsilon$ is a non-degenerate compact interval, and one can leverage this property to show that the mapping $x \mapsto g_b(x)$ is increasing. Moreover, it satisfies $g_b(0) = d(\mu_a, \mu_a \wedge (\mu^+ - \epsilon))$ and $\lim_{x \rightarrow \infty} g_b(x) = d(\mu_a, (\mu_b - \epsilon) \vee \mu^-)$. Therefore, one can define the inverse mapping

$$x_b(y) = g_b^{-1}(y) \text{ for } y \in [d(\mu_a, \mu_a \wedge (\mu^+ - \epsilon)), d(\mu_a, (\mu_b - \epsilon) \vee \mu^-)] .$$

We furthermore define the constant mapping $x_a(y) = 1$ for all $y \in [d(\mu_a, \mu_a \wedge (\mu^+ - \epsilon)), d(\mu_a, (\max_{b \neq a} \mu_b - \epsilon) \vee \mu^-)]$.

To state Theorem 6.1 below, which provides a more explicit expression of the weights $w_\epsilon^{*,a}(\boldsymbol{\mu})$, we also introduce the notation $\lambda_b(x)$ for the minimizer in (6.3):

$$g_b(x) = d(\mu_a, \lambda_b(x)) + x d(\mu_b, \lambda_b(x) + \epsilon)$$

The proof of Theorem 6.1, given in Appendix D, hinges on the fact that at the optimum all the values of $g_b(w_b/w_a)$ are identical, which allows to re-parameterize the optimization problem by their common value.

Theorem 6.1. Fix a bandit instance $\boldsymbol{\mu}$ and $a \in \mathcal{A}_\epsilon(\boldsymbol{\mu})$ and let

$$\mathcal{I}_a = \left[d(\mu_a, \mu_a \wedge (\mu^+ - \epsilon)), d(\mu_a, (\max_{b \neq a} \mu_b - \epsilon) \vee \mu^-) \right] .$$

1. For degenerate instances $(\boldsymbol{\mu}, \epsilon)$ such that $\mu_a > \mu^+ - \epsilon$ and there exists $b \neq a$ such that $\mu_b = \mu^+$,

$$(w_{\epsilon}^{*,a}(\boldsymbol{\mu}))_b = \mathbb{1}_{(b=a)}.$$

2. For other instances $(\boldsymbol{\mu}, \epsilon)$, let y^* be the unique solution on \mathcal{I}_a to the equation

$$\sum_{b \neq a} \frac{d(\mu_a, \lambda_b(x_b(y)))}{d(\mu_b, \lambda_b(x_b(y))) + \epsilon} = 1. \quad (6.4)$$

Then the argmax in (6.1) is unique and is defined by

$$(w_{\epsilon}^{*,a}(\boldsymbol{\mu}))_b = \frac{x_b(y^*)}{\sum_{b=1}^k x_b(y^*)}.$$

In practice, the optimal weights can be computed by solving (6.4) using binary search. Each computation of the left-hand side requires to compute $\lambda_b(x_b(y))$. For each x , $\lambda_b(x)$ is the minimizer of a twice differentiable convex function and can be approximated numerically using for example Newton's method. In the Gaussian case, we may also use the closed form $\lambda_b(x) = \frac{\mu_a + x(\mu_b - \epsilon)}{1+x}$. Then the value of the inverse function $x_b(y)$ can be obtained by solving $g_b(x) = y$ using binary search.

6.2. The ϵ -Tracking rule

The Tracking sampling rule was originally proposed by [Garivier and Kaufmann \(2016\)](#) for best arm identification ($\epsilon = 0$), for which the set of optimal weights $\mathcal{W}_0(\boldsymbol{\mu})$ is reduced to a single vector, $w^*(\boldsymbol{\mu})$. Building on our ability to compute those optimal weights, the Tracking rule is a mechanism that forces the empirical proportions of arm selections to converge to this oracle allocation, by relying on plug-in estimates $w^*(\hat{\boldsymbol{\mu}}(t))$ and ensuring enough exploration.

We present below an extension of this rule, called ϵ -Tracking, that can be used in any ϵ -best arm identification problem. The difference with Tracking is that ties need to be broken when $\mathcal{W}_*(\hat{\boldsymbol{\mu}}(t))$ contains multiple vectors. We do by relying on a fixed ordering of the arms, introducing the notation

$$w_{\epsilon}^*(\boldsymbol{\mu}) = w_{\epsilon}^{*,a^*}(\boldsymbol{\mu}) \quad \text{where } a^* = \min \{a \in \{1, \dots, K\} : w_{\epsilon}^{*,a}(\boldsymbol{\mu}) \in \mathcal{W}_{\epsilon}(\boldsymbol{\mu})\}.$$

At round $t + 1$, ϵ -Tracking first computes the set $U_t = \{a : N_a(t) < \sqrt{t} - K/2\}$ of arms that are currently under-sampled. Then it selects

$$A_{t+1} \in \begin{cases} \operatorname{argmin}_{a \in U_t} N_a(t) & \text{if } U_t \neq \emptyset \quad (\text{forced exploration}), \text{ or otherwise} \\ \operatorname{argmax}_{1 \leq a \leq K} [t \times (w_{\epsilon}^*)_a(\hat{\boldsymbol{\mu}}(t)) - N_a(t)] & (\text{tracking the plug-in estimate}). \end{cases}$$

Observe that this sampling rule does not depend on δ . Other Tracking rules exist, for example tracking the cumulative proportions (C-Tracking), which forces $N_a(t)$ to be close to $\sum_{s=1}^t w_a^*(\hat{\boldsymbol{\mu}}_s(t))$. Their relative merits is discussed in [\(Degegne and Koolen, 2019\)](#).

Definition 6.1. An instance $(\boldsymbol{\mu}, \epsilon)$ is regular if the set $\mathcal{W}_{\epsilon}^*(\boldsymbol{\mu})$ is of cardinality one. For a regular instance, $\mathcal{W}_{\epsilon}^*(\boldsymbol{\mu}) = \{w_{\epsilon}^*(\boldsymbol{\mu})\}$.

Using the same arguments as those used by [Garivier and Kaufmann \(2016\)](#) for the Tracking rule, one can show that for regular instances ϵ -Tracking is a converging strategy (in the sense of [Definition 5.3](#)), which converges to the optimal proportions.

Lemma 6.1. Let $(\boldsymbol{\mu}, \epsilon)$ be a regular instance of almost optimal best arm identification. Then, under the ϵ -Tracking sampling rule,

$$\mathbb{P}_{\boldsymbol{\mu}} \left(\forall a = 1, \dots, K, \lim_{t \rightarrow \infty} \frac{N_a(t)}{t} = (w_{\epsilon}^*(\boldsymbol{\mu}))_a \right) = 1.$$

6.3. The Parallel GLRT Stopping Rule for ϵ -Best Arm Identification

Introducing for all a, b the statistic

$$\hat{Z}_{a,b}^\epsilon(t) = \inf_{\lambda: \lambda_a < \lambda_b - \epsilon} [N_a(t)d(\hat{\mu}_a(t), \lambda_a) + N_b(t)d(\hat{\mu}_b(t), \lambda_b)]$$

and $\hat{\mathcal{A}}_\epsilon(t) = \{a : \hat{\mu}_a(t) \geq \max_j \hat{\mu}_j(t) - \epsilon\}$, the parallel GLRT stopping rule can be written

$$\tau_\delta = \inf \left\{ t \in \mathbb{N} : \max_{a \in \hat{\mathcal{A}}_\epsilon(t)} \min_{b \neq a} \hat{Z}_{a,b}^\epsilon(t) > \beta(t, \delta) \right\} \quad (6.5)$$

and the associated recommendation rule is

$$\hat{i}_{\tau_\delta} \in \operatorname{argmax}_{a \in \hat{\mathcal{A}}_\epsilon(\tau_\delta)} \min_{b \neq a} \hat{Z}_{a,b}^\epsilon(\tau_\delta). \quad (6.6)$$

In the Gaussian case, a more explicit expression of $\hat{Z}_{a,b}^\epsilon(t)$ can be given for a and b such that $\hat{\mu}_a(t) \geq \hat{\mu}_b(t) - \epsilon$ (which holds when $a \in \hat{\mathcal{A}}_\epsilon(t)$):

$$\hat{Z}_{a,b}^\epsilon(t) = \frac{1}{2\sigma^2} \frac{N_a(t)N_b(t)}{N_a(t) + N_b(t)} (|\hat{\mu}_a(t) - \hat{\mu}_b(t)| + \epsilon)^2,$$

Beyond the Gaussian case, when $\hat{\mu}_a(t) \geq \hat{\mu}_b(t) - \epsilon$ one can write

$$\hat{Z}_{a,b}^\epsilon(t) = \inf_{\lambda \in]\mu^-, \mu^+ - \epsilon[} [N_a(t)d(\hat{\mu}_a(t), \lambda) + N_b(t)d(\hat{\mu}_b(t), \lambda + \epsilon)],$$

and the solution of this minimization problem has to be numerically approximated.

Lemma 4.1 ensures that the parallel GLRT stopping rule defined in (6.5) based on the threshold

$$\beta(t, \delta) = 3K \ln(1 + \ln t) + K\mathcal{T} \left(\frac{\ln(1/\delta)}{K} \right)$$

yields an (ϵ, δ) -PAC strategy when combined with any sampling rule. Using a more refined analysis tailored to particular structure of the parallel GLRT test for ϵ -best arm identification permit to guarantee the (ϵ, δ) -PAC property for a smaller threshold. Note that the threshold $\beta(t, \delta) = \log \left(\frac{2(K-1)t}{\delta} \right)$ proposed by Garivier and Kaufmann (2016) for $\epsilon = 0$ (and that rely on different deviation inequalities) can also be used here, but the one given in Lemma 6.2 can be smaller for large values of t .

Lemma 6.2. *For any sampling rule, the parallel GLRT stopping rule with threshold*

$$\beta(t, \delta) = 6 \ln(1 + \ln t) + 2\mathcal{T} \left(\frac{\ln((K-1)/\delta)}{2} \right) \quad (6.7)$$

satisfies $\mathbb{P}(\tau_\delta < \infty, \hat{i}_{\tau_\delta} \notin \mathcal{A}_\epsilon) \leq \delta$.

Proof.

$$\begin{aligned} \mathbb{P}_\mu(\mu_{\hat{a}_\tau} < \mu_1 - \epsilon) &\leq \sum_{b \notin \mathcal{A}_\epsilon} \mathbb{P}_\mu(\hat{a}_\tau = b, \tau < \infty) \\ &\leq \sum_{b \notin \mathcal{A}_\epsilon} \mathbb{P}_\mu \left(\exists t \in \mathbb{N} : \inf_{\mu'_b - \mu'_1 < -\epsilon} [N_1(t)d(\hat{\mu}_1(t), \mu'_1) + N_b(t)d(\hat{\mu}_b(t), \mu'_b)] > \beta(t, \delta) \right) \\ &\leq \sum_{b \notin \mathcal{A}_\epsilon} \mathbb{P}_\mu(\exists t \in \mathbb{N} : N_1(t)d(\hat{\mu}_1(t), \mu_1) + N_b(t)d(\hat{\mu}_b(t), \mu_b) > \beta(t, \delta)) \\ &\leq \sum_{b \notin \mathcal{A}_\epsilon} \frac{\delta}{K-1} = \delta, \end{aligned}$$

where the last inequality relies on Theorem 14 in (Kaufmann and Koolen, 2018).

□

6.4. The ϵ -Track-and-Stop Strategy

We define the ϵ -Track-and-Stop strategy with parameter δ as the following strategy for ϵ -Best Arm Identification:

- its sampling rule is the ϵ -Tracking rule
- its stopping rule is the parallel GLRT with threshold $\beta(t, \delta)$ given in (6.7)
- its recommendation rule is the one associated to the parallel GLRT, that is \hat{i}_{τ_δ} defined in (6.6).

This strategy, that we denote by ϵ -TaS(δ) has the following properties.

Theorem 6.2. *Fix $\delta \in (0, 1]$. Then ϵ -TaS(δ) is (ϵ, δ) -PAC. Moreover, for every regular instance $(\boldsymbol{\mu}, \epsilon)$, if τ_δ denotes the stopping rule of ϵ -TaS(δ),*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\ln(1/\delta)} \leq T_\epsilon^*(\boldsymbol{\mu}).$$

This result shows that the family of ϵ -TaS(δ) strategies for δ in $(0, 1)$, which is a family of converging strategies by Lemma 6.1, is ϵ -PAC and is matching the lower bound given in Theorem 5.2, for regular instances.

Beyond regular instances, the picture is intriguingly more complex, and was recently studied in details by Degenne and Koolen (2019), who consider a general active testing framework with overlapping hypotheses. They show that ϵ -Tracking can fail to converge to an element in $\mathcal{W}_\epsilon^*(\boldsymbol{\mu})$ and propose a “sticky” variant fixing this problem. This alternative sampling rule turns out not to be a converging strategy anymore, as under two different runs it may converge to a different element in $\mathcal{W}_\epsilon^*(\boldsymbol{\mu})$. However, the authors manage to extend the lower bound of Theorem 5.2 to *any* strategy (using a nice game-theoretic argument). Hence the quantity $T_\epsilon^*(\boldsymbol{\mu})$ is indeed the right characteristic number of samples for ϵ -best arm identification.

Sketch of proof To understand why the parallel GLRT attains the lower bound of Theorem 5.2 if coupled with a good sampling rule, we first rewrite its stopping rule

$$\tau_\delta = \inf \{t \in \mathbb{N}^* : \hat{Z}(t) \geq \beta(t, \delta)\},$$

where we introduce

$$\begin{aligned} \hat{Z}(t) &= \max_{a \in \hat{\mathcal{A}}_\epsilon(t)} \min_{b \neq a} \inf_{\lambda \in (\mu^-, \mu^+ - \epsilon)} [N_a(t)d(\hat{\mu}_a(t), \lambda) + N_b(t)d(\hat{\mu}_b(t), \lambda + \epsilon)] \\ &= t \times \left(\max_{a \in \hat{\mathcal{A}}_\epsilon(t)} \min_{b \neq a} \inf_{\lambda \in (\mu^-, \mu^+ - \epsilon)} \left[\frac{N_a(t)}{t} d(\hat{\mu}_a(t), \lambda) + \frac{N_b(t)}{t} d(\hat{\mu}_b(t), \lambda + \epsilon) \right] \right). \end{aligned}$$

By Lemma 6.1, under a regular instance for which $\mathcal{W}_\epsilon^*(\boldsymbol{\mu}) = \{w_\epsilon^*(\boldsymbol{\mu})\}$ the ϵ -Tracking sampling rule satisfies that, for every arm a , $\hat{\mu}_a(t)$ converges almost surely to μ_a and the sampling proportion $N_a(t)/t$ converges almost surely to the target $(w_\epsilon^*(\boldsymbol{\mu}))_a$. Using the fact that $\boldsymbol{\mu} \mapsto w_\epsilon^*(\boldsymbol{\mu})$ is continuous in any regular instance $\boldsymbol{\mu}$, for large values of t it holds that

$$\begin{aligned} \hat{Z}(t) &\simeq t \times \left(\max_{a \in \mathcal{A}_\epsilon} \min_{b \neq a} \inf_{\lambda \in (\mu^-, \mu^+ - \epsilon)} [(w_\epsilon^*(\boldsymbol{\mu}))_a d(\mu_a, \lambda) + (w_\epsilon^*(\boldsymbol{\mu}))_b d(\mu_b, \lambda + \epsilon)] \right) \\ &= t \times (T_\epsilon^*(\boldsymbol{\mu}))^{-1}. \end{aligned}$$

Therefore one obtains

$$\tau_\delta \simeq \inf \{t \in \mathbb{N}^* : t \geq T_\epsilon^*(\boldsymbol{\mu})\beta(t, \delta)\} . \quad (6.8)$$

If the threshold function were chosen to be $\beta(t, \delta) = \ln(1/\delta)$ one would immediately get the upper bound $T_\epsilon^*(\boldsymbol{\mu}) \ln(1/\delta)$. However, a slightly larger threshold is needed to ensure the (ϵ, δ) -PAC property of ϵ -TaS(δ). Still, the threshold defined in (6.7) satisfies $\beta(t, \delta) \leq \ln(1/\delta) + \ln(t)$ for large values of t hence Lemma A.1 can be applied to show that the right hand side of (6.8) can be upper bounded by

$$T_\epsilon^*(\boldsymbol{\mu}) \ln\left(\frac{1}{\delta}\right) + 2T_\epsilon^*(\boldsymbol{\mu}) \ln\left(T_\epsilon^*(\boldsymbol{\mu}) \ln\left(\frac{1}{\delta}\right)\right) .$$

This argument can be made rigorous to prove that on every regular instance $\boldsymbol{\mu}$,

$$\mathbb{P}_\boldsymbol{\mu} \left(\lim_{\delta \rightarrow 0} \frac{\tau_\delta}{\ln(1/\delta)} \leq T_\epsilon^*(\boldsymbol{\mu}) \right) = 1 ,$$

in the spirit of the proof of Proposition 13 in (Garivier and Kaufmann, 2016). The proof of Theorem 6.2 is a bit more complex as it requires to control the expectation $\mathbb{E}[\tau_\delta]$. It can be obtained by following similar arguments as the proof of Theorem 14 in (Garivier and Kaufmann, 2016).

7. NUMERICAL EXPERIMENTS

We report results of some experiments comparing ϵ -Track-and-Stop (TaS) to some state-of-the-art algorithms for ϵ -best arm identification for Bernoulli distributed arms. These competitors are KL-LUCB, KL-Racing (Kaufmann and Kalyanakrishnan, 2013) and UGapE (Gabillon et al., 2012), briefly described below. For instances having a unique optimal arm, we also include 0-TaS in our study (designed for finding the unique best arm) in order to assess the gain in sample complexity achieved when studying the ϵ relaxation.

We experiments with two regular instances,

$$\begin{aligned} \boldsymbol{\mu}_1 &= [0.2 \ 0.4 \ 0.5 \ 0.55 \ 0.7] & \epsilon_1 &= 0.1 \\ \boldsymbol{\mu}_2 &= [0.4 \ 0.5 \ 0.6 \ 0.7 \ 0.75 \ 0.8] & \epsilon_2 &= 0.15 , \end{aligned}$$

and one non-regular instance $\boldsymbol{\mu}_3 = [0.2 \ 0.3 \ 0.45 \ 0.55 \ 0.6 \ 0.6] \ \epsilon_3 = 0.1$.

In the experiments we set $\delta = 0.1$ and perform $N = 1000$ independent replications in order to estimate the probability of error $\mathbb{P}_{\boldsymbol{\mu}_i}(\boldsymbol{\mu}_{\hat{\tau}_\delta} \notin \mathcal{A}_{\epsilon_i}(\boldsymbol{\mu}_i))$, and the expected number of samples $\mathbb{E}_{\boldsymbol{\mu}_i}[\tau_\delta]$ for each algorithm and problem instance. First, Table 1 permits to check that all algorithms are indeed δ -correct: their empirical probability of error is (much) smaller than the maximal value $\delta = 0.1$ which is prescribed.

	ϵ -TaS	KL-LUCB	UGapE	KL-Racing	0-TaS
$\boldsymbol{\mu}_1, \epsilon_1 = 0.1$	0.015	0.003	0.002	0.003	0.003
$\boldsymbol{\mu}_2, \epsilon_2 = 0.15$	0	0	0	0	0
$\boldsymbol{\mu}_3, \epsilon_3 = 0.1$	0.001	0	0	0	-

Table 1. Empirical error probabilities for various algorithms estimated on $N = 1000$ repetitions.

In those experiments, the stopping rule of Track-and-Stop is the parallel GLRT with threshold $\beta(t, \delta) = \ln((1 + \ln(t))/\delta)$. This threshold is an approximation of the theoretical threshold given in Lemma 6.2 that we recommend to use in practice. All the other algorithms rely on confidence bounds, and we implement them with upper and lower confidence bounds based on the Kullback-Leibler divergence

$$\begin{aligned} u_a(t) &= \max \{q \in [0, 1] : N_a(t)d(\hat{\mu}_a(t), q) \leq \beta(t, \delta)\} \quad \text{and} \\ \text{and } \ell_a(t) &= \min \{q \in [0, 1] : N_a(t)d(\hat{\mu}_a(t), q) \leq \beta(t, \delta)\} , \end{aligned}$$

for the same choice of $\beta(t, \delta) = \ln((1 + \ln(t))/\delta)$. The stopping rule of KL-LUCB and UGapE consists in waiting for the lower confidence bound of the candidate best arm to overlap by at most ϵ with the upper-confidence bound of all other arms. KL-Racing is an algorithm based on eliminations that also exploits those confidence intervals: arms in a list of candidate arms are selected in a round robin fashion and when some arm a satisfies $\ell_b(t) > u_a(t) - \epsilon$ for some other arm b , it is removed from the list of candidate arms. The algorithm stops when all candidate arms are empirically ϵ -optimal. Those confidence-based stopping or elimination rules seem to be far too conservative compared to the parallel GLRT stopping rule, as the sample complexity results now reveal.

	$T^*(\boldsymbol{\mu}) \ln(1/\delta)$	ϵ -TaS	KL-LUCB	UGapE	KL-Racing	0-TaS
$\boldsymbol{\mu}_1, \epsilon_1 = 0.1$	97	171 (104)	322 (137)	324 (143)	372 (159)	433 (277)
$\boldsymbol{\mu}_2, \epsilon_2 = 0.15$	108	162 (83)	345 (135)	344 (141)	402 (146)	2659 (1863)
$\boldsymbol{\mu}_3, \epsilon_3 = 0.1$	531	501 (261)	1236 (403)	1199 (414)	1348 (436)	-

Table 2. Estimated values of $\mathbb{E}_{\mu_i}[\tau_\delta]$ based on $N = 1000$ repetitions for different instances and algorithms (standard deviation indicated in parenthesis).

Indeed, it can be seen in Table 2 that ϵ -TaS requires about half the number of samples needed by other methods before stopping, and is sometimes getting very close to the ideal sample complexity $T^*(\boldsymbol{\mu}) \ln(1/\delta)$. In the first instance, in which there is a single ϵ -best arm, one can see that ϵ -TaS slightly improves over 0-TaS due to the ϵ -relaxation. This improvement is much more flagrant in the second instance, which contains 3 ϵ -best arms.

For the two instances with 3 ϵ -best arms $\boldsymbol{\mu}_2$ and $\boldsymbol{\mu}_3$, we now report in Table 3 the optimal weights in $\mathcal{W}_\epsilon(\boldsymbol{\mu})$. There is one such weight for $\boldsymbol{\mu}_2$ and two for $\boldsymbol{\mu}_3$ in which the fraction allocation to the two best arms are exchanged. We also report in this table the empirical fraction of selection of each arm by ϵ -TaS.

$(\boldsymbol{\mu}_2)_a$	0.4	0.5	0.6	0.7	0.75	0.8
$(w_{\epsilon_2}^*(\boldsymbol{\mu}_2))_a$	0.024	0.036	0.060	0.136	0.275	0.469
$\mathbb{E}_{\mu_2}[N_a(\tau_\delta)/\tau_\delta]$	0.079	0.077	0.099	0.156	0.235	0.353
$(\boldsymbol{\mu}_3)_a$	0.2	0.3	0.45	0.55	0.6	0.6
$(w_{\epsilon_3}^{*,5}(\boldsymbol{\mu}_3))_a$	0.008	0.0133	0.035	0.114	0.436	0.393
$(w_{\epsilon_3}^{*,4}(\boldsymbol{\mu}_3))_a$	0.008	0.0133	0.035	0.114	0.393	0.436
$\mathbb{E}_{\mu_3}[N_a(\tau_\delta)/\tau_\delta]$	0.049	0.049	0.081	0.194	0.317	0.309

Table 3. Optimal proportions and empirical proportions achieved by ϵ -TaS (ϵ -optimal arms in bold)

For the regular instance $\boldsymbol{\mu}_2$, according to Lemma 6.1 the empirical proportions converge to the true proportions. However, as can be seen from Table 3 the convergence is slow and did not yet occur at the stopping time τ_δ (all sub-optimal arms are slightly over sampled which may be a consequence of the forced exploration scheme). For the non-regular instance $\boldsymbol{\mu}_3$ in which there are two candidate optimal weights, ϵ -TaS has no convergence guarantee towards one of these weights (and Degenne and Koolen (2019) propose a fix). Just like for $\boldsymbol{\mu}_2$, we see that all sub-optimal arms are slightly over-sampled under $\boldsymbol{\mu}_3$. On average over many simulations, the strategy spends about the same time on the two optimal arms, which could indicate that “convergence” occurs towards $w_{\epsilon_3}^{*,4}(\boldsymbol{\mu}_3)$ or $w_{\epsilon_3}^{*,5}(\boldsymbol{\mu}_3)$ half of the time. Even without proper convergence guarantees, ϵ -TaS performs very well on $\boldsymbol{\mu}_3$ as can be noted from Table 2.

8. CONCLUSION

We proposed a non-asymptotic way to analyze parallel Generalized Likelihood Ratio Tests, which can be used in various adaptive testing settings, from sequential to active testing, even in the presence of overlapping hypotheses. We also presented a way to derive lower bounds providing the exact asymptotic sample complexity of some testing problems. Remarkably, for problem instances that belong to multiple hypotheses, non-asymptotic lower bounds are still out of reach, maybe because of the use of the low-level change of low formula. An interesting open question is thus to know whether lower bounds can always be derived using only the high-level form, in a way that provides non-asymptotic results. So far, we were not able to achieve this goal, nor to show that it is impossible.

Natural extensions of our approach can be obtained for other sequential and active testing problems. For example, in bandit models, our approach is likely to apply also for thresholding bandits, or for best-arm identification with a structure on the set of arms. Moreover, if a regularity constraint like $\mu_{i+1} - \mu_i \leq L$ could be handled, this would open the way to the study of the sample complexity of optimizing a Lipschitz continuous function. In that context, one can hope to exhibit a optimal density of sampling (corresponding to the optimal weights in the finite case), depending on the function only, which an optimal optimization procedure would need to track. All these developments are left for future investigations.

ACKNOWLEDGEMENTS

Emilie Kaufmann acknowledges the support of the Agence Nationale de la Recherche (ANR) under grants ANR-16-CE40-0002 (BADASS project), ANR-19-CE23-0026-04 (BOLD project) and the European CHIST-ERA project DELTA. Aurélien Garivier acknowledges the support of the Project IDEXLYON of the University of Lyon, in the framework of the Programme Investissements d’Avenir (ANR-16-IDEX-0005).

A. NON-ASYMPTOTIC SAMPLE COMPLEXITY ANALYSIS FOR A GAUSSIAN TEST

We provide a general sample complexity analysis of the parallel GLRT test using a generic threshold function $\beta(t, \delta)$. Introducing the function

$$f(t, \delta) := \frac{|\mu| + \epsilon}{\sigma} \sqrt{t} - \sqrt{2\beta(t, \delta)},$$

we make the following assumptions:

1. $\beta(t, \delta) \leq \ell(\delta) + c \ln(t)$ where $\ell(\delta) \geq \ln(1/\delta)$. This permits to define $T_0(\delta)$ to be any constant such that

$$\forall t \geq T_0(\delta), \quad f(t, \delta) \geq 0.$$

2. For all $\delta \in (0, 1]$, $\forall t \geq T_0(\delta)$,

$$\frac{\partial}{\partial t}(\sqrt{2\beta(t, \delta)}) \leq \frac{1}{2} \frac{\partial}{\partial t} \sqrt{t \frac{(|\mu| + \epsilon)^2}{\sigma^2}} = \frac{|\mu| + \epsilon}{4\sigma^2} \frac{1}{\sqrt{t}}. \quad (\text{A.1})$$

A consequence of Assumption 2 is that $t \mapsto f(t, \delta)$ is non-decreasing, which permits to upper bound the expectation of T_δ as follows, introducing (Z_t) a sequence of standard normal random variables.

$$\begin{aligned} \mathbb{E}[T_\delta] &\leq \sum_{t=1}^{\infty} \mathbb{P} \left(t \frac{(|\hat{\mu}_t| + \epsilon)^2}{2\sigma^2} \leq \beta(t, \delta) \right) \\ &\leq \sum_{t=1}^{\infty} \mathbb{P} \left(\sqrt{t} \left(\left| \mu + \frac{\sigma Z_t}{\sqrt{t}} \right| + \epsilon \right) \leq \sqrt{2\sigma^2 \beta(t, \delta)} \right) \\ &\leq \sum_{t=1}^{\infty} \mathbb{P} \left(\sqrt{t} (|\mu| - |\sigma Z_t / \sqrt{t}| + \epsilon) \leq \sqrt{2\sigma^2 \beta(t, \delta)} \right) \\ &\leq \sum_{t=1}^{\infty} \mathbb{P} \left(|Z_t| \geq \frac{|\mu| + \epsilon}{\sigma} \sqrt{t} - \sqrt{2\beta(t, \delta)} \right) \\ &\leq [T_0(\delta)] + \sum_{t=[T_0(\delta)]+1}^{\infty} \mathbb{P} (|Z_t| \geq f(t, \delta)) dt \\ &\leq [T_0(\delta)] + \int_{T_0(\delta)}^{\infty} \mathbb{P} (|Z_t| \geq f(t, \delta)) dt. \end{aligned}$$

Using assumption (A.1) one can write

$$\begin{aligned} f(t, \delta) &= f(T_0(\delta), \delta) + \int_{T_0(\delta)}^t f'(s, \delta) ds \\ &\geq 0 + \frac{1}{2} \int_{T_0(\delta)}^t \left(\sqrt{t \frac{(|\mu| + \epsilon)^2}{\sigma^2}} \right)' ds \\ &= \int_{T_0(\delta)}^t \frac{|\mu| + \epsilon}{4\sigma} \frac{1}{\sqrt{s}} ds. \end{aligned}$$

For $t \in [T_0(\delta), 2T_0(\delta)]$, observe that

$$f(t, \delta) \geq (t - T_0(\delta)) \frac{|\mu| + \epsilon}{4\sqrt{2}\sigma} \frac{1}{\sqrt{T_0(\delta)}},$$

while for $t > 2T_0(\delta)$,

$$f(t, \delta) \geq (t - T_0(\delta)) \frac{|\mu| + \epsilon}{4\sigma} \frac{1}{\sqrt{t}} \geq \sqrt{t - T_0(\delta)} \frac{|\mu| + \epsilon}{4\sqrt{2}\sigma}.$$

Hence

$$\begin{aligned} \mathbb{E}[T_\delta] &\leq [T_0(\delta)] + \int_{T_0(\delta)}^{2T_0(\delta)} 2\bar{\Phi}(f(t, \delta)) dt + \int_{2T_0(\delta)}^{\infty} 2\bar{\Phi}(f(t, \delta)) dt \\ &\leq [T_0(\delta)] + 2 \int_{T_0(\delta)}^{\infty} \bar{\Phi} \left((t - T_0(\delta)) \frac{|\mu| + \epsilon}{4\sqrt{2}\sigma} \frac{1}{\sqrt{T_0(\delta)}} \right) dt + 2 \int_{2T_0(\delta)}^{\infty} \bar{\Phi} \left(\sqrt{t - T_0(\delta)} \frac{|\mu| + \epsilon}{4\sqrt{2}\sigma} \right) dt \\ &\leq [T_0(\delta)] + \frac{8\sqrt{2}\sigma}{|\mu| + \epsilon} \sqrt{T_0(\delta)} \int_0^{\infty} \bar{\Phi}(u) du + \frac{64\sigma^2}{(|\mu| + \epsilon)^2} \int_{T_0(\delta)}^{\infty} \frac{(\frac{|\mu| + \epsilon}{8\sigma^2})^2}{\sqrt{u}} \bar{\Phi}(\sqrt{u}) du \\ &\leq [T_0(\delta)] + \frac{8\sqrt{2}\sigma}{|\mu| + \epsilon} \sqrt{T_0(\delta)} \frac{1}{\sqrt{2\pi}} + \frac{64\sigma^2}{(|\mu| + \epsilon)^2} \int_{T_0(\delta)}^{\infty} \frac{(\frac{|\mu| + \epsilon}{8\sigma^2})^2}{\sqrt{u}} e^{-\frac{u}{2}} du \\ &\leq 1 + T_0(\delta) + \frac{8\sqrt{2}\sigma}{|\mu| + \epsilon} \sqrt{T_0(\delta)} \frac{1}{\sqrt{2\pi}} + \frac{64\sigma^2}{(|\mu| + \epsilon)^2} e^{-T_0(\delta)} \frac{(\frac{|\mu| + \epsilon}{8\sqrt{2}\sigma^2})^2}{\sqrt{2\pi}}. \end{aligned}$$

If the exploration function β satisfies

$$\beta(t, \delta) \leq \ell(\delta) + c \ln(t),$$

Lemma A.1 below permits to prove that one can pick

$$T_0(\delta) = \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) + \frac{4c\sigma^2}{(|\mu| + \epsilon)^2} \ln \left(\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) \right).$$

In that case,

$$\begin{aligned} \mathbb{E}[T_\delta] &\leq \frac{2\sigma^2}{(|\mu| + \epsilon)^2} \left[\ell(\delta) + 2c \ln \left(\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) \right) \right. \\ &\quad \left. + 8\sqrt{\ell(\delta) + 2c \ln \left(\frac{2\sigma^2}{(|\mu| + \epsilon)^2} \ell(\delta) \right) + 32\delta^{1/8}} \right] + 1. \end{aligned}$$

Theorem 3.1 follows from observing that the threshold (3.3) for which the parallel GLRT is δ -correct satisfies

$$\beta(t, \delta) \leq \mathcal{T}(\ln(1/\delta)) + 3 \frac{1 + \ln(t)}{e},$$

and thus Assumption 1 is satisfied with $c = 3/e$ and $\ell(\delta) = \mathcal{T}(\ln(1/\delta)) + 3/e$. It can be checked by calculus that Assumption 2 is also satisfied.

Lemma A.1. Fix $\alpha \geq 0$ and $\gamma \geq 1 + \alpha$. Then for all $t > 0$,

$$t \geq \gamma + 2\alpha \ln(\gamma) \Rightarrow t \geq \gamma + \alpha \ln(t).$$

Proof. Let $f(t) = t - \gamma - \alpha \ln(t)$. Then $f'(t) = 1 - \alpha/t \geq 0$ for $t \geq \alpha$. Hence, for all $t \geq t_0 := \gamma + 2\alpha \ln(\gamma) \geq \alpha$,

$$\begin{aligned} f(t) &\geq f(t_0) = \gamma + 2\alpha \ln(\gamma) - \gamma - \alpha \ln(\gamma + 2\alpha \ln(\gamma)) \\ &= \alpha \ln(\gamma) - \alpha \ln \left(1 + \frac{2\alpha}{\gamma} \ln(\gamma) \right) \end{aligned}$$