Top-m identification for linear bandits

Clémence Réda[1,*] **Emilie Kaufmann**[2] **Andrée Delahaye-Duriez**[1,3,4]

[1] Université de Paris, Inserm UMR 1141 NeuroDiderot, F-75019, Paris, France, *`clemence.reda@inria.fr`
[2] Université Lille, CNRS, Inria, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France
[3] Université Sorbonne Paris Nord, UFR SMBH, F-93000, Bobigny, France & [4] Assistance Publique des Hôpitaux de Paris, Hôpital Jean Verdier, Service d'Histologie-Embryologie-Cytogénétique, F-93140, Bondy, France

Abstract

Motivated by an application to drug repurposing, we propose the first algorithms to tackle the identification of the $m \geq 1$ arms with largest means in a linear bandit model, in the fixed-confidence setting. These algorithms belong to the generic family of Gap-Index Focused Algorithms (GIFA) that we introduce for Top-m identification in linear bandits. We propose a unified analysis of these algorithms, which shows how the use of features might decrease the sample complexity. We further validate these algorithms empirically on simulated data and on a simple drug repurposing task.

1 INTRODUCTION

The multi-armed bandit setting, in which an agent sequentially gathers samples from K unknown probability distributions called arms, is a powerful framework for sequential resource allocation tasks. While a large part of the literature focuses on the reinforcement learning problem in which the samples are viewed as reward that the agent seeks to maximize (Bubeck and Cesa-Bianchi, 2012), *pure-exploration* objectives have also received a lot of attention (Bubeck et al., 2009; Degenne and Koolen, 2019). In this paper, we focus on *Top-m identification* in which the goal is to identify the $m < K$ arms with the largest expected rewards. While several Top-m identification algorithms have been given, no algorithm has been specifically designed to tackle the challenging *linear bandit* setting (Auer, 2003). This paper aims at filling this gap.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130.

In a linear bandit model, the mean μ_a of each arm a is assumed to depend linearly on a known feature vector $x_a \in \mathbb{R}^N$ associated to the arm: $\mu_a = \theta^\top x_a$ for some vector $\theta \in \mathbb{R}^N$. In contrast, the so-called classical bandit model does not make any assumption on the means $(\mu_a)_{a\in[K]}$[1]. A Top-m identification algorithm outputs a subset of size m as a guess for the m arms with largest means. Several objectives exist: in the fixed-budget setting, the goal is to minimize the probability that the guess is wrong after a pre-specified total number of samples from the arms (Bubeck et al., 2013). In this paper, we focus on the *fixed-confidence setting*, in which the error probability should be guaranteed to be smaller than a given risk parameter $\delta \in (0, 1)$ while minimizing the *sample complexity*, that is, the total number of samples needed to output the guess. This choice is motivated by a real-life application to drug repurposing, in which we would like to control the failure rate in our predictions.

Drug repurposing is a field of research aimed at discovering new indications for drugs which are already approved for marketing, and may contribute to solve the problem of ever increasing research budget need for drug discovery (Hwang et al., 2016). For a given disease, we are interested in identifying a subset of drugs that may have a therapeutic interest. Providing a group of 5 or 10 drugs, rather than a single one, can ease the decision of further investigation, as many leads are provided. We believe sequential methods could be of interest, when considering a drug repurposing method called "signature reversion" (Musa et al., 2018). In this context, drugs recommended for repurposing are the ones minimizing the difference in gene activity between treated patients and healthy individuals. A possible way to measure this difference is to build a simulator that evaluates the genewise impact of a given drug on patients. However, this simulator may be stochastic and computationally expen-

[1] For $n \in \mathbb{N}^*$ we use the shorthand $[n] = \{1, \ldots, n\}$.

sive hence the need for sequential queries, that can be modeled as sampling of arms. As the arms (drugs) can be characterized by a real-valued feature vector which represents genewise activity change due to treatment, we resort to linear bandits to tackle this Top-m identification problem.

Related work Two types of fixed-confidence algorithms have been proposed for Top-m identification in a classical bandit: those based on adaptive sampling such as LUCB (Kalyanakrishnan et al., 2012) or UGapE (Gabillon et al., 2012) or those based on uniform sampling and eliminations (Kaufmann and Kalyanakrishnan, 2013; Chen et al., 2017). For linear bandits, to the best of our knowledge, the only efficient algorithms have been proposed for the *best arm identification* (BAI) problem, which corresponds to $m = 1$. This setting, first investigated by Soare et al. (2014), recently received a lot of attention: an efficient adaptive sampling algorithm called LinGapE was proposed by Xu et al. (2018) and subsequent work such as Fiez et al. (2019) sought to achieve the minimal sample complexity. In particular, the LinGame algorithm of Degenne et al. (2020) is proved to exactly achieve the problem-dependent sample complexity lower bound for linear BAI in a regime in which δ goes to zero.

We note that, in principle, LinGame can be used for any pure exploration problem in a linear bandit, which includes Top-m identification for $m > 1$. However, this algorithm uses a game theoretic formalism which needs the computation of a best response for Nature in response to the player's selection; a computable expression of this strategy is not available to our knowledge for Top-m ($m > 1$). Besides, computing the information-theoretic lower bound for Top-m identification is also computationally hard. These remarks led us to investigate efficient adaptive sampling algorithms for general Top-m identification ($m \geq 1$) in linear bandit, which are still missing in the literature, instead of trying to propose asymptotically optimal algorithms, as done in linear BAI.

Contributions First, by carefully looking at known adaptive sampling bandits for classical Top-m, we propose a generic algorithm structure based on *Gap Indices*, called GIFA, which encompasses existing adaptive algorithms for classical Top-m identification and linear BAI. This structure allows a higher order and modular understanding of the learning process, and correctness properties can readily be inferred from a partially specified bandit algorithm. It allows us to define two interesting new algorithms, called m-LinGapE and LinGIFA. In Section 4, we present a unified sample complexity analysis of a subclass of GIFA algorithms which comprises existing methods, which shows that the use of features can help decreasing the sample complexity in some cases. Finally, we show in Section 5 that m-LinGapE and LinGIFA perform better than their counterparts for classical bandits, both on artificially generated linear bandit instances and on a simple instance of our drug repurposing application, that is described in detail in Appendix B.

Notation We let $\overset{[m]}{\max}$, $\overset{m}{\max}$ (resp. $\overset{[m]}{\min}$, $\overset{m}{\min}$) be the operator returning the m, m^{th} greatest (resp. smallest) value(s), $\|x\|_M = \sqrt{x^\top M x}$ where M is positive definite, and $\|x\| = \sqrt{x^\top x}$.

2 SETTING

In this section, we introduce the *probably approximately correct* (PAC) fixed-confidence linear Top-m problem. From now on, we identify each arm by an integer in $[K]$ such that $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_m > \mu_{m+1} \geq \mu_{m+2} \geq \cdots \geq \mu_K$. This ordering is of course unknown to the learner. In the PAC formalism, we are willing to relax the optimality constraint on arms using slack variable $\varepsilon > 0$ and we denote by $\mathcal{S}_m^{*,\varepsilon} \triangleq \{a \in [K] : \mu_a \geq \mu_m - \varepsilon\}$ the set of best arms up to ε. Since $\mu_m > \mu_{m+1}$, we use the shorthand $\mathcal{S}_m^\star \triangleq \mathcal{S}_m^{*,0} = [m]$ for the set of m arms with largest means. The set of $K - m$ arms worst arms is its complementary $(\mathcal{S}_m^\star)^c = [K] \backslash \mathcal{S}_m^\star$. A linear bandit model is parameterized by an unknown vector $\theta \in \mathbb{R}^N$, such that there are constants $L, S > 0$ which satisfy $\|\theta\| \leq S \in \mathbb{R}^{*+}$ and such that the mean of arm a is $\mu_a = \theta^\top x_a$, where the feature vector x_a satisfies $\|x_a\| \leq L \in \mathbb{R}^{*+}$. In each round $t \geq 1$, a learner selects an arm $a_t \in [K]$ and observes a sample

$$r_t = \theta^\top x_{a_t} + \eta_t,$$

where η_t is a zero-mean sub-Gaussian noise with variance σ^2 –that is, $\mathbb{E}[e^{\lambda \eta_t}] \leq \exp\left((\lambda^2\sigma^2)/2\right)$ for all $\lambda \in \mathbb{R}$ – which is independent from past observations.

For $m \in [K]$, an algorithm for Top-m identification consists of a *sampling rule* $(a_t)_{t \in \mathbb{N}^*}$, where a_t is $\mathcal{F}_{t-1}$-measurable, and $\mathcal{F}_t = \sigma(a_1, r_1, \ldots, a_t, r_t)$ is the σ-algebra generated by the observations up to round t ; a *stopping rule* τ, which is a stopping time with respect to $\mathcal{F}_t$ that indicates when the learning process is over; and a *recommendation rule* $\hat{S}_m^\tau$, which is a $\mathcal{F}_\tau$-measurable set of size m that provides a guess for the m arms with the largest means. Then, given a slack variable $\varepsilon \geq 0$ and the failure rate $\delta \in (0, 1)$, an (ε, m, δ)-PAC algorithm is such that

$$\mathbb{P}\left(\hat{S}_m^\tau \subseteq \mathcal{S}_m^{*,\varepsilon}\right) \geq 1 - \delta.$$

In the fixed confidence setting, the goal is to design a (ε, m, δ)-PAC algorithm with a small *sample complexity* τ. In order to do so, the learner needs to estimate the unknown parameter $\theta \in \mathbb{R}^N$, which can be done with a (regularized) least-squares estimator. For any arm $a \in [K]$, we let $N_a(t) \triangleq \sum_{s=1}^{t} \mathbb{1}(a_s = a)$ be the number of times that arm a is sampled up to time t, and define the λ-regularized design matrix and least-squares estimate:

$$\hat{V}_t^\lambda \triangleq \lambda I_N + \sum_{a=1}^{K} N_a(t) x_a x_a^\top$$
$$\text{and } \hat{\theta}_t^\lambda \triangleq \left(\hat{V}_t^\lambda\right)^{-1} \left(\sum_{s=1}^{t} r_s x_{a_s}\right) .$$

We let $\hat{\Sigma}_t^\lambda \triangleq \sigma^2 \left(\hat{V}_t^\lambda\right)^{-1}$, which can be interpreted as the posterior covariance in a Bayesian linear regression model in which the covariance of the prior is $(\sigma^2/\lambda) I_N$.

The algorithms that we present in the next section crucially rely on estimating the gaps $\Delta_{i,j} \triangleq \mu_i - \mu_j$ between pairs of arms (i, j), and building *upper confidence bounds* (UCBs) on these quantities. For this purpose, we introduce the empirical mean of each arm a, $\hat{\mu}_a(t) \triangleq (\hat{\theta}_t^\lambda)^\top x_a$ and define the empirical gap $\hat{\Delta}_{i,j}(t) \triangleq \hat{\mu}_i(t) - \hat{\mu}_j(t)$. A first option to build UCBs on $\Delta_{i,j}$ consists in building individual confidence intervals on the mean of each arm a, that are of the form $L_a(t) = \hat{\mu}_a(t) - W_t(a)$ for the lower-confidence bound, and $U_a(t) = \hat{\mu}_a(t) + W_t(a)$ for the upper confidence bound, where $W_t(a) \triangleq C_{t,\delta} \|x_a\|_{\hat{\Sigma}_t^\lambda}$ for some threshold function $C_{t,\delta}$ to be specified later. Clearly,

$$B_{i,j}^{\text{ind}}(t) = U_i(t) - L_j(t) = \hat{\Delta}_{i,j}(t) + W_t(i) + W_t(j)$$

is an upper bound on $\Delta_{i,j}$ if $L_j(t) \le \mu_j$ and $\mu_i \le U_i(t)$. Yet, using the linear model, one can also directly build an UCB on the difference via

$$B_{i,j}^{\text{pair}}(t) = \hat{\Delta}_{i,j}(t) + C_{t,\delta} \|x_i - x_j\|_{\hat{\Sigma}_t^\lambda}.$$

Both constructions lead to symmetrical bounds[2] of the form $B_{i,j}(t) = \hat{\Delta}_{i,j}(t) + W_t(i,j)$, where $W_t(i,j) = C_{t,\delta}\|x_i - x_j\|_{\hat{\Sigma}_t^\lambda}$ for paired UCBs and $W_t(i,j) = C_{t,\delta}\left(\|x_i\|_{\hat{\Sigma}_t^\lambda} + \|x_j\|_{\hat{\Sigma}_t^\lambda}\right)$ for individual UCBs. In both cases, $W_t(i,j) = W_t(j,i) \ge 0$. The fact that these quantities are indeed upper confidence bounds (for some pairs $(i,j) \in [K]^2$) will be justified in Section 3.

Remark 1. *Observe that the linear bandit setting subsumes the classical bandit model, choosing arm features $(x_a)_{a\in[K]}$ to be vectors of the canonical basis of $\mathbb{R}^K$ ($N = K$) and $\theta = \mu = [\mu_1, \mu_2, \ldots, \mu_K]^\top$. Then, $\lambda = 0$ yields standard UCBs in this model: $\hat{\mu}_a(t)$ reduces to the empirical average of the rewards gathered from arm a and $\|x_a\|_{\hat{\Sigma}_t^\lambda} = \sigma/\sqrt{N_a(t)}$.*

[2] Note also that there exist algorithms with non symmetrical bounds, such the KL version of LUCB (Kaufmann and Kalyanakrishnan, 2013).

3 GAP-INDEX FOCUSED ALGORITHMS (GIFA)

3.1 Generic GIFA Algorithms

Looking at the literature for Top-m identification and BAI in (linear) bandits, existing adaptive sampling designs (Kalyanakrishnan et al., 2012; Gabillon et al., 2012; Xu et al., 2018) have many ingredients in common. We formalize in this section a generic structure that encompasses these algorithms, and under which we propose two new algorithms. We introduce the notion of "Gap-Index Focused" Algorithms (GIFA), which associates to each *pair of arm* (i, j) an index $B_{i,j}(t)$ at time t (called "gap index").

The idea in GIFA is to estimate in each round t a set of candidate m best arms, denoted by $J(t)$, and to select the two most ambiguous arms: $b_t \in J(t)$, which can be viewed as a guess for the m-best arm, and a *challenger* $c_t \notin J(t)$. c_t is defined as a potentially misassessed m-best arm, with largest possible gap to b_t: $c_t = \arg\max_{c\in[K]} B_{c,b_t}(t)$. The idea of using two ambiguous arms goes back to LUCB (Kalyanakrishnan et al., 2012) for Top-m identification in classical bandits. Then the final arm a_t selected by the algorithm should help discriminate between b_t and c_t. A naive idea is to either draw b_t or c_t, but alternative selection rules will be discussed later. At the end of the learning phase, at stopping time τ, the final set $J(\tau)$ is recommended.

Algorithm 1 GIFA for (ε, δ)-PAC Top-m

Input: K arms of means $\mu_1, \ldots, \mu_K$; $\varepsilon \ge 0$; $m \le K$; $\delta \in (0, 1)$.
Output: $\hat{S}_m^{\tau_\delta}$ estimated m ε-optimal arms at final time τ_δ
$t \leftarrow 1$
`initialization`()
while $\neg$ `stopping_rule`$((B_{i,j}(t))_{i,j\in[K]}; \varepsilon, t)$ **do**
  // *$J(t)$: estimated m best arms at t*
  $J(t) \leftarrow$ `compute_Jt`$((B_{i,j}(t))_{i,j\in[K]}, (\hat{\mu}_i(t))_{i\in[K]})$
  // *b_t: estimated m-best arm at t*
  $b_t \leftarrow$ `compute_bt`$(J(t), (B_{i,j}(t))_{i,j\in[K]})$
  // *c_t: challenger to b_t at t*
  $c_t \leftarrow \arg\max_{a\notin J(t)} B_{a,b_t}(t)$
  // *selecting and pulling arms*
  $a_t \leftarrow$ `selection_rule`$(b_t, c_t; \hat{V}_{t-1}^\lambda)$
  $r_t \leftarrow \text{pulling}(a_t)$
  Update design matrix $\hat{V}_t^\lambda$, means $(\hat{\mu}_i(t))_{i\in[K]}$
  Update gap indices $(B_{i,j}(t+1))_{i,j\in[K]}$
  $t \leftarrow t + 1$
end while
$\hat{S}_m^{\tau_\delta} \leftarrow J(\tau_\delta)$
return $\hat{S}_m^{\tau_\delta}$

Table 1: Adaptive samplings for Top-m (our proposals are in bold type ; except for LUCB and UGapE, all algorithms use indices $B = B^{\text{pair}}$ instead of individual indices B^{ind}).

Algorithm	`compute_Jt`	`compute_bt`	`selection_rule`	`stopping_rule`
LUCB	$\mathop{\arg\max}\limits^{[m]}_{j\in[K]} \hat{\mu}_j(t)$	$\arg\max\limits_{j\in J(t)} \max\limits_{i\notin J(t)} B_{i,j}(t)$	largest variance	$B_{c_t,b_t}(t) \leq \varepsilon$
UGapE	$\mathop{\arg\min}\limits^{[m]}_{j\in[K]} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t)$	$\arg\max\limits_{j\in J(t)} \max\limits_{i\notin J(t)} B_{i,j}(t)$	largest variance	$\max\limits_{j\in J(t)} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t) \leq \varepsilon$
LinGapE	$\arg\max\limits_{j\in[K]} \hat{\mu}_j(t)$	$J(t) \quad (\lvert J(t)\rvert = m = 1)$	greedy, optimized	$B_{c_t,b_t}(t) \leq \varepsilon$
***m*-LinGapE** ($m \geq 1$)	$\mathop{\arg\max}\limits^{[m]}_{j\in[K]} \hat{\mu}_j(t)$	$\arg\max\limits_{j\in J(t)} \max\limits_{i\notin J(t)} B_{i,j}(t)$	largest variance, greedy, optimized	$B_{c_t,b_t}(t) \leq \varepsilon$
LinGIFA	$\mathop{\arg\min}\limits^{[m]}_{j\in[K]} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t)$	$\arg\max\limits_{j\in J(t)} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t)$	largest variance, greedy	$\max\limits_{j\in J(t)} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t) \leq \varepsilon$

The GIFA structure is presented in Algorithm 1. `initialization` is an optional phase where all arms are sampled once. We assume that ties are randomly broken. Degrees of freedom in designing the bandit algorithm lie in the choice of the rules `compute_Jt`, `compute_bt`, `selection_rule` and `stopping_rule`, some of which may also rely on the gap indices. For example, for the stopping rule, we restrict our attention to two stopping rules already proposed for the LUCB and UGapE algorithms, respectively:

$$\tau^{LUCB} \triangleq \inf \left\{ t \in \mathbb{N}^* : B_{c_t,b_t}(t) \leq \varepsilon \right\} ,$$
$$\tau^{UGapE} \triangleq \inf \left\{ t \in \mathbb{N}^* : \max_{j\in J(t)} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t) \leq \varepsilon \right\} .$$

3.2 Existing GIFA algorithms

We summarize in Table 1 existing and new algorithms that fit the GIFA framework.

LUCB and UGapE The only algorithms for Top-m identification in the fixed-confidence setting using adaptive sampling –namely, LUCB (Kalyanakrishnan et al., 2012) and UGapE (Gabillon et al., 2012)– were proposed for classical bandits. Both can be cast in the GIFA framework with the individual gap indices $B^{\text{ind}}_{i,j}(t) = U_i(t) - L_j(t)$ presented in Section 2. Indeed, with this choice of indices, it can be shown that the ambiguous arms $b_t = \arg\min_{j\in J(t)} L_j(t)$ and $c_t = \arg\max_{i\notin J(t)} U_i(t)$ used by both LUCB and UGapE can be rewritten as $\arg\max_{j\in J(t)} \min_{i\notin J(t)} B_{i,j}(t)$ and $\arg\max_{i\notin J(t)} B_{i,b_t}(t)$, respectively. This rewriting is crucial to propose extensions for the linear case. However, LUCB and UGapE differ by the set $J(t)$ they use and their stopping rule, as can be seen in Table 1. LUCB uses the m arms with largest empirical means for the set $J(t)$ while UGapE uses the m arms with smallest values of $\mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t)$. This choice is motivated by Gabillon et al. (2012) by the fact that, with high probability, the quantity $\mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t)$ is an upper bound on $\mu_m - \mu_j$. This is also true for more general gap indices, see Lemma 9 in Appendix D.

This observation also justifies the UGapE stopping rule: when $\max_{j\in J(t)} \mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t) \leq \varepsilon$, with high probability $\mu_m - \mu_j \leq \varepsilon$ for all $j \in J(t)$, hence this set is likely to be included in $\mathcal{S}^{*,\varepsilon}_m$. LUCB stops when $\max_{j\in J(t)} \max_{i\notin J(t)} B_{i,j}(t) \leq \varepsilon$. Interestingly, it follows from Lemma 1 below that, for fixed gap indices $(B_{i,j}(t))_{i,j\in[K]^2,t>0}$, UGapE will always stop before LUCB. Still, we chose to investigate the two stopping rules, as the sample complexity analysis of UGapE uses the upper bound $\tau^{LUCB} \geq \tau^{UGapE}$. We provide correctness guarantees for both stopping rules in Section 4.1 and investigate their practical impact in Section 5.

Lemma 1. *For all $t > 0$, for any subset J of size m, for all $j \in J$, $\mathop{\max}\limits^{m}_{i\neq j} B_{i,j}(t) \leq \max_{i\notin J} B_{i,j}(t)$.*

(proof in Appendix D)

Regarding the selection rule, UGapE selects the least sampled arm among b_t and c_t, which coincides with the largest variance rule that we propose for the general linear setting:

- "largest variance": $\arg\max_{a\in\{b_t,c_t\}} W_t(a)$.

In the original version of LUCB, both arms b_t and c_t are sampled at time t, but the analysis that we propose in this paper obtains similar guarantees for LUCB using the largest variance rule, so we only consider this selection rule in the remainder of the paper.

LinGapE A striking example of adaptive sampling in linear BAI (thus, only for $m = 1$) is LinGapE (Xu et al., 2018). The first novelty in LinGapE compared to UGapE and LUCB is the use of paired gap indices $B^{\text{pair}}_{i,j}(t)$ exploiting the arm features. Paired indices may indeed increase performance. First, it follows from the triangular inequality for the Mahalanobis norm $\|\cdot\|_{\hat{\Sigma}^\lambda_t}$ that $||x_i - x_j||_{\hat{\Sigma}^\lambda_t} \leq ||x_i||_{\hat{\Sigma}^\lambda_t} + ||x_j||_{\hat{\Sigma}^\lambda_t}$ for all pairs of arms (i, j) ; therefore, if both types of gap indices use the same threshold $C_{t,\delta}$, paired indices are smaller: $B^{\text{pair}}_{i,j}(t) \leq B^{\text{ind}}_{i,j}(t)$. Moreover, Lemma 2 below (proved in Appendix D) implies that paired or individual indices using arm features always yield smaller bounds on the gaps than individual indices without arm features, that take the form

$$\hat{\Delta}_{i,j}(t) + C_{t,\delta}\left[\frac{1}{\sqrt{N_i(t)}} + \frac{1}{\sqrt{N_j(t)}}\right] .$$

Lemma 2. $\forall t > 0, \forall a \in [K], \forall y \in \mathbb{R}^N$,

$$||y||_{(\hat{V}^\lambda_t)^{-1}} \leq ||y||/\left(\sqrt{N_a(t)||x_a||^2 + \lambda}\right).$$

The second novelty in LinGapE is the proposed selection rules. Let X be the matrix which columns are the arm contexts, and $\|\cdot\|_1 : v \mapsto \sum_i |v_i|$. Instead of selecting either b_t or c_t, Xu et al. (2018) propose two selection rules to possibly sample another arm that would reduce the variance of the estimate $\theta^\top(x_{b_t} - x_{c_t})$:

- "greedy": $\arg\min_{a\in[K]} ||x_{b_t} - x_{c_t}||_{(\hat{V}^\lambda_{t-1}+x_a x_a^\top)^{-1}}$,
- "optimized": if $w^*(b_t, c_t) = \underset{w\in\mathbb{R}^K : x_{b_t}-x_{c_t}=wX^T}{\arg\min} \|w\|_1$,

$$\underset{a\in[K]:w^*_a(b_t,c_t)>0}{\arg\max} N_a(t)\frac{||w^*(b_t, c_t)||_1}{|w^*_a(b_t, c_t)|}. \tag{1}$$

As explained by Xu et al. (2018), these two rules are meant to bring the empirical proportions of selections close to an optimal design $(\lambda_1, \ldots, \lambda_K)$ asymptotically minimizing $\|x_{b_t} - x_{c_t}\|_{(\sum_a \lambda_a x_a x_a^\top)^{-1}}$.

3.3 m-LinGapE and LinGIFA

Building on the interesting features of existing algorithms, we now propose two new algorithms for Top-m identification in linear bandits, called m-LinGapE and LinGIFA and described in Table 1.

m-LinGapE is an extension of LinGapE (Xu et al., 2018) to linear Top-m, which coincides with the original algorithm for $m = 1$. It may also be viewed as an extension of LUCB using paired indices. We investigate three possible selection rules for this algorithm: largest variance, greedy and optimized. LinGIFA is inspired by UGapE (Gabillon et al., 2012), and also uses paired gap indices. Note that LinGIFA has a unique `compute_bt` rule $b_t = \arg\max_{j\in J(t)} \overset{m}{\underset{i\neq j}{\max}}\ B_{i,j}(t)$, which does not coincide with any of the previously mentioned rules. In Table 1, we considered the stopping rules of the original algorithms they derive from. As LinGapE is one of the most performant algorithms for linear BAI (see for instance experiments in Degenne et al. (2020)), we expect m-LinGapE to work well for Top-m. However, the strength of LinGIFA is that it is completely defined in terms of gap indices, and, as such, can easily be tuned for performance by deriving tighter bounds on the gaps. Moreover, LinGIFA relies on a stopping rule that is more aggressive than that used by m-LinGapE, as discussed above.

We emphasize that provided that the regularizing constant λ in the design matrix is positive, both $m-$LinGapE and LinGIFA can be run without an initialization phase (as done in the experiments for LinGIFA). This permits to avoid the initial sampling cost when the number of arms is large that was noticed by Fiez et al. (2019).

4 THEORETICAL GUARANTEES

In this section, we present an analysis of m-LinGapE and LinGIFA. The fact that these algorithms are (ε, m, δ)-PAC is a consequence of generic correctness guarantees that can be obtained for GIFA algorithms (even not fully specified) and is presented in Section 4.1. We further analyze in Section 4.2 the sample complexity of LUCB-like algorithms, which comprise m-LinGapE.

4.1 Correctness of GIFA instances

We justify that the two stopping rules introduced above lead to (ε, m, δ)-PAC algorithms, provided a condition on the gap indices given in Definition 1.

Definition 1 (Good gap indices)**.** Let us denote

$$\mathcal{E}^{GIFA}_m \triangleq \bigcap_{t>0} \bigcap_{j\in(\mathcal{S}^{*,\varepsilon}_m)^c} \bigcap_{k\in\mathcal{S}^\star_m} (B_{k,j}(t) \geq \mu_k - \mu_j) .$$

A *good choice* of gap indices $(B_{i,j}(t))_{i,j\in[K],t>0}$ satisfies $\mathbb{P}(\mathcal{E}^{GIFA}_m) \geq 1 - \delta$.

First, we observe that on the event $\mathcal{E}^{GIFA}_m$ introduced in Definition 1 both stopping rules τ^{UGapE} and τ^{LUCB} output an (ε, δ) correct answer.

Theorem 1. *On the event* $\mathcal{E}^{GIFA}_m$*:* **(i)** *any GIFA algorithm using* $b_t = \arg\max_{j\in J(t)} \max_{i\notin J(t)} B_{i,j}(t)$ *sat-*

isfies $J(\tau^{LUCB}) \subseteq \mathcal{S}_m^{*,\varepsilon}$ and **(ii)** *any GIFA algorithm using $b_t \in J(t)$ satisfies $J(\tau^{UGapE}) \subseteq \mathcal{S}_m^{*,\varepsilon}$.*

Proof. Let us assume by contradiction that there is an arm $b \in J(t) \cap (\mathcal{S}_m^{*,\varepsilon})^c$ at stopping time t. We first prove that in both cases there exists $c \in \mathcal{S}_m^\star$ such that $B_{c,b}(t) \leq \varepsilon$ (*). Assuming that (*) does not hold, observe that $\mathcal{S}_m^\star \subseteq \{a \neq b, B_{a,b}(t) > \varepsilon\}$ **(a)** and then $\arg\overset{m}{\max}_{a\neq b} B_{a,b}(t) \in \{a \neq b, B_{a,b}(t) > \varepsilon\}$ **(b)** since $b \notin \mathcal{S}_m^\star$ and $|\mathcal{S}_m^\star| = m$. Depending on the definition of b_t and on the stopping rule, we now split the proof into two parts:

(i). Using the definition of τ^{LUCB}, c_t and b_t, it holds that $\forall b \in J(t), \forall c \notin J(t), B_{c,b}(t) \leq \varepsilon$. Then, using **(a)**, $J(t)^c \subseteq \{a \neq b, B_{a,b}(t) \leq \varepsilon\}$, which means $\mathcal{S}_m^\star \cap (J(t))^c = \emptyset$. Thus $\mathcal{S}_m^\star = J(t)$ (because $|\mathcal{S}_m^\star| = |J(t)| = m$) whereas $b \in J(t) \cap (\mathcal{S}_m^\star)^c$, which is a contradiction. Hence (*) holds.

(ii). Using the definition of τ^{UGapE}, for any $b \in J(t)$, $\overset{m}{\underset{a\neq b}{\max}}\ B_{a,b}(t) \leq \epsilon$ holds. However, **(b)** means $\overset{m}{\underset{a\neq b}{\max}}\ B_{a,b}(t) > \epsilon$, which is a contradiction: (*) holds.

Then, at stopping time t, there exists $b \in J(t) \cap (\mathcal{S}_m^{*,\varepsilon})^c$ and $c \in \mathcal{S}_m^\star$ such that $B_{c,b}(t) \leq \varepsilon$. Using successively the definition of the event $\mathcal{E}_m^{GIFA}$, and the fact that $c \in \mathcal{S}_m^\star$, this means that there exists $b \in J(t)^c \cap (\mathcal{S}_m^\star)^c$ such that $\mu_b \geq \mu_m - \varepsilon$. Then $b \in (\mathcal{S}_m^{*,\varepsilon})^c \cap \mathcal{S}_m^{*,\varepsilon}$, which is absurd. Hence this proves that $J(t) \subseteq \mathcal{S}_m^{*,\varepsilon}$ on the event $\mathcal{E}_m^{GIFA}$ at stopping time t. □

It easily follows from Theorem 1 that if m-LinGapE or LinGIFA are based on good gap indices in the sense of Definition 1, both algorithms are (ε, m, δ)-PAC. We exhibit below a threshold for which the corresponding paired indices $B_{i,j}^{\text{pair}}(t)$ and individual indices $B_{i,j}^{\text{ind}}(t)$ (defined in Section 2) are good gap indices.

Lemma 3. *For indices of the form $B_{i,j}(t) = \hat{\Delta}_{i,j}(t) + C_{\delta,t} W_t(i,j)$ with $W_t(i,j) = \|x_i - x_j\|_{\hat{\Sigma}_t^\lambda}$ (paired) or $W_t(i,j) = ||x_i||_{\hat{\Sigma}_t^\lambda} + ||x_j||_{\hat{\Sigma}_t^\lambda}$ (individual) with*

$$C_{\delta,t} = \sqrt{2\ln\left(\frac{1}{\delta}\right) + N\ln\left(1 + \frac{(t+1)L^2}{\lambda^2 N}\right)} + \frac{\sqrt{\lambda}}{\sigma}S, \tag{2}$$

where we recall that $\max_{a\in[K]} \|x_a\| \leq L$ and $\|\theta\| \leq S$, we have $\mathbb{P}(\mathcal{E}_m^{GIFA}) \geq 1 - \delta$.

Proof. The proof follows from the fact that

$$\left\{\forall t \in \mathbb{N}^* : \|\hat{\theta}_t^\lambda - \theta\|_{(\hat{\Sigma}_t^\lambda)^{-1}} \leq C_{\delta,t}\right\} \subseteq \mathcal{E}_m^{GIFA}$$

together with Lemma 4.1 in Kaufmann (2014) which yields $\mathbb{P}\left(\forall t \in \mathbb{N}^* : \|\hat{\theta}_t^\lambda - \theta\|_{(\hat{\Sigma}_t^\lambda)^{-1}} \leq C_{\delta,t}\right) \geq 1-\delta$. For paired indices, the inclusion follows from the fact that

$$|(\hat{\mu}_i(t) - \hat{\mu}_j(t)) - (\mu_i - \mu_j)| \leq \|\hat{\theta}_t^\lambda - \theta\|_{(\hat{\Sigma}_t^\lambda)^{-1}} \|x_i - x_j\|_{\hat{\Sigma}_t^\lambda},$$

where we use that $x^\top y \leq \|x\|_{\Sigma^{-1}} \|y\|_\Sigma$ for any $x, y \in \mathbb{R}^N$ and positive definite $\Sigma \in \mathbb{R}^{N\times N}$. For individual indices, we further need the triangular inequality $\|x_i - x_j\|_{\hat{\Sigma}_t^\lambda} \leq \|x_i\|_{\hat{\Sigma}_t^\lambda} + \|x_j\|_{\hat{\Sigma}_t^\lambda}$ to prove the inclusion. □

4.2 Sample complexity results

We derive below a high-probability upper bound on the sample complexity of a subclass of GIFA algorithms which comprise m-LinGapE, combined with different selection rules. More precisely, we upper bound the sample complexity on

$$\mathcal{E} \triangleq \bigcap_{t>0} \bigcap_{i,j\in[K]} \Big(\mu_i - \mu_j \in [-B_{j,i}(t), B_{i,j}(t)]\Big),$$

an event which is trivially included in $\mathcal{E}^{\text{GIFA}}$ and which holds therefore with probability larger than $1-\delta$ with the choice of threshold (2). To state our results, we define the true gap of an arm k as $\Delta_k \triangleq \mu_k - \mu_{m+1}$ if $k \in \mathcal{S}_m^\star$, $\mu_m - \mu_k$ otherwise ($\Delta_k \geq 0$ for any $k \in [K]$).

Theorem 2. *For m-LinGapE, on event $\mathcal{E}$ on which algorithm $\mathcal{A}$ is (ε, m, δ)-PAC, stopping time τ_δ satisfies $\tau_\delta \leq \inf\{u \in \mathbb{R}^{*+} : u > 1 + H^\varepsilon(\mathcal{A},\mu) C_{\delta,u}^2 + \mathcal{O}(K)\}$, where:*

(i). *for $\mathcal{A} = m$-LinGapE with the largest variance selection rule[3] :*

$$H^\varepsilon(\mathcal{A},\mu) \triangleq 4\sigma^2 \sum_{a\in[K]} \max\left(\varepsilon, \frac{\varepsilon + \Delta_a}{3}\right)^{-2},$$

(ii). *for $\mathcal{A} = m$-LinGapE with the optimized selection rule in Xu et al. (2018) :*

$$H^\varepsilon(\mathcal{A},\mu) \triangleq \sigma^2 \sum_{a\in[K]} \max_{i,j\in[K]} \frac{|w_a^*(i,j)|}{\max\left(\varepsilon, \frac{\varepsilon+\Delta_i}{3}, \frac{\varepsilon+\Delta_j}{3}\right)^2},$$

where $w^*(i,j)$ satisfies Equation 1 applied to the arm pair (i,j).

Upper bounds for all analyzed algorithms and classical counterparts are shown in Table 2.

4.3 Sketch of proof

The proof of Theorem 2 generalizes and extends the proofs for classical Top-m and linear BAI (Xu et al., 2018), with paired or individual gap indices. We sketch below the sample complexity analysis of a specific subclass of GIFA algorithms, postponing the proof of some

[3] or pulling both arms in $\{b_t, c_t\}$ at time t

Table 2: Sample complexity results for linear and classical Top-m algorithms. w^* is defined as in Equation (1). The additive term in K is due to the initialization phase. Our proposal is in bold type.

Algorithm	Complexity constant $\mathrm{H}^\varepsilon(\cdot,\mu)$	Upper bound on τ_δ
LUCB	$2\sum_{a\in[K]}\max\left(\frac{\varepsilon}{2},\Delta_a\right)^{-2}$	$\inf_{u>0}\left\{u>1+\mathrm{H}^\varepsilon(\text{LUCB},\mu)C^2_{\delta,u}+K\right\}$
UGapE	$2\sum_{a\in[K]}\max\left(\varepsilon,\frac{\varepsilon+\Delta_a}{2}\right)^{-2}$	$\inf_{u>0}\left\{u>1+\mathrm{H}^\varepsilon(\text{UGapE},\mu)C^2_{\delta,u}+K\right\}$
m-LinGapE (1) (largest variance)	$4\sigma^2\sum_{a\in[K]}\max\left(\varepsilon,\frac{\varepsilon+\Delta_a}{3}\right)^{-2}$	$\inf_{u>0}\left\{u>1+\mathrm{H}^\varepsilon(m\text{-LinGapE}(1),\mu)C^2_{\delta,u}\right\}$
m-LinGapE (2) (optimized)	$\sigma^2\sum_{a\in[K]}\max_{i,j\in[K]}\frac{\lvert w^*_a(i,j)\rvert}{\max\left(\varepsilon,\frac{\varepsilon+\Delta_i}{3},\frac{\varepsilon+\Delta_j}{3}\right)^2}$	$\inf_{u>0}\left\{u>1+\mathrm{H}^\varepsilon(m\text{-LinGapE}(2),\mu)C^2_{\delta,u}\right\}$

auxiliary lemmas to Appendix C. We focus on GIFA algorithms that use

$$J(t)\triangleq \underset{i\in[K]}{\arg\max}^{[m]}\ \hat{\mu}_i(t) \ \text{ and } \ b_t\triangleq \underset{j\in J(t)}{\arg\max}\ \max_{i\notin J(t)} B_{i,j}(t)\ .$$

This includes LinGapE, m-LinGapE, but also LUCB. The key ingredient in the proof is the following lemma, which holds for any gap indices of the form $B_{i,j}(t)\triangleq\hat{\mu}_i(t)-\hat{\mu}_j(t)+W_t(i,j)$ for $i,j\in[K]^2$.

Lemma 4. *On the event $\mathcal{E}$, for all $t>0$,*

$$B_{c_t,b_t}(t)\leq\min(-(\Delta_{b_t}\vee\Delta_{c_t})+2W_t(b_t,c_t),0)+W_t(b_t,c_t).$$

This result is a counterpart of Lemma 4 in Xu et al. (2018), but does not require $|J(t)|=1$ at a given time $t>0$, notably by noticing that, by definition of b_t and c_t, $B_{c_t,b_t}(t)=\max_{j\in J(t)}\max_{i\notin J(t)}B_{i,j}(t)$. In order to get the upper bound in Theorem 2 for m-LinGapE using the optimized rule, one can straightforwardly apply Lemma 1 in Xu et al. (2018) to the inequality stemming from Lemma 4. For the selection rules which either select both b_t and c_t, or the arm in $\{b_t,c_t\}$ with the largest variance term, by combining Lemma 4 with the definition of the stopping rule τ^{LUCB}, we obtain the following upper bound on $N_{a_t}(t)$, where a_t is (one of) the pulled arm(s) at time $t<\tau_\delta$.

Lemma 5. $\forall t>0,\tau_\delta>t,N_{a_t}(t)\leq T^*(a_t,\delta,t)$, *where a_t is a pulled arm at time t, and*

$$T^*(a_t,\delta,t)=4\sigma^2C^2_{\delta,t}\max\left(\varepsilon,\frac{\varepsilon+\Delta_{a_t}}{3}\right)^{-2}.$$

Finally, the upper bound in Theorem 2 is a consequence of the following result.

Lemma 6. *Let $T^*:[K]\times(0,1)\times\mathbb{N}^*\to\mathbb{R}^{*+}$ be a function that is nondecreasing in t, and $\mathcal{I}_t$ the set of pulled arms at time t. Let $\mathcal{E}$ be an event such that for all $t<\tau_\delta,\delta\in(0,1)$, $\exists a_t\in\mathcal{I}_t, N_{a_t}(t)\leq T^*(a_t,\delta,t)$. Then it holds on the event $\mathcal{E}$ that $\tau_\delta\leq T(\mu,\delta)$ where*

$$T(\mu,\delta)\triangleq\inf\left\{u\in\mathbb{R}^{*+}:u>1+\sum_{a=1}^K T^*(a,\delta,u)\right\}.$$

4.4 Discussion

If $\mathrm{H}^\varepsilon(\mathcal{A},\mu)$ is the complexity constant associated with bandit algorithm $\mathcal{A}$ and bandit instance $\mu\in\mathbb{R}^K$, then, by looking at Table 2, we can first notice that

$$\mathrm{H}^\varepsilon(\text{LUCB},\mu)\geq\mathrm{H}^\varepsilon(m\text{-LinGapE}(1),\mu)\geq\mathrm{H}^\varepsilon(\text{UGapE},\mu).$$

Note that the quantity $w^*_a(i,j)$ featured in the complexity quantity associated to m-LinGapE with the optimized selection rule (m-LinGapE(2)) is not *a priori* bounded by a constant – see Equation (1). Authors in (Xu et al., 2018) have shown that in specific instances, $\mathrm{H}^\varepsilon(m\text{-LinGapE}(2),\mu)\leq\frac{9}{8}\mathrm{H}^\varepsilon(\text{UGapE},\mu)$. To further compare these two complexity quantities, we designed the following experiment to estimate how many times $\mathrm{H}^\varepsilon(m\text{-LinGapE}(2),\mu)\leq\mathrm{H}^\varepsilon(\text{UGapE},\mu)$: $K\times N$ values are randomly sampled from a Gaussian distribution $N(0,D)$. From these we build matrix $X'\in\mathbb{R}^{N\times K}$, which is then column-normalized (norm $\|\cdot\|$) to yield feature matrix X. We use $\theta=e_1$, where $e_1=[1,0,\ldots,0]^\top\in\mathbb{R}^K$ and $\|\theta\|=1$. Then we compute both constants $\mathrm{H}^\varepsilon(m\text{-LinGapE}(2),\mu)$ and $\mathrm{H}^\varepsilon(\text{UGapE},\mu)$ where $\mu=X\theta$ and $m=\left[\frac{K}{3}\right]$, check whether $\mathrm{H}^\varepsilon(m\text{-LinGapE}(2),\mu)\leq\mathrm{H}^\varepsilon(\text{UGapE},\mu)$. We perform this experiment $1,000$ times and report in Table 4 (in Appendix A) the fraction of times this condition holds, for multiple values of K (number of arms), N (dimension) and D (variance of the Gaussian distribution). We observe that, in these artificially generated instances, the condition $\mathrm{H}^{\varepsilon}(m\text{-LinGapE}(2), \mu) \leq \mathrm{H}^{\varepsilon}(\mathrm{UGapE}, \mu)$ is seldom verified.

Hence, our theoretical analysis suggests that, in these cases –which usually turn out to be really hard linear instances– LinGIFA, which structure is similar to that of UGapE, may perform better than m-LinGapE. However, as we will see below, the practical story is different and m-LinGapE with the optimized selection rule can be as efficient, and often better than LinGIFA using the "largest variance" selection rule.

5 EXPERIMENTAL STUDY

In this section, we report the results of experiments on simulated data and on a simple drug repurposing instance. In all experiments, we use $\sigma = 1/2$, $\varepsilon = 0$, $\delta = 5\%$. For all artificially generated experimental settings, we use Gaussian arms, that is, at time t, $r_t \sim N(\mu_{a_t}, \sigma^2)$. For the tuning of our algorithm we set the regularization parameter to $\lambda = \sigma/20$, following Hoffman et al. (2014). In an effort to respect the trade-off between identification performance and speed, we set $C_{\delta,t} = \sqrt{2\ln\left((\ln(t)+1)/\delta\right)}$ ("heuristic" value) if the empirical error remained below δ instead of the theoretically valid threshold (2) that guarantees all algorithms to be $(\varepsilon, m, \delta)-$PAC. Otherwise, we use (2), and for classical bandits, we use the threshold used in Kalyanakrishnan et al. (2012). We consider the following experimental settings.

Classic instances We consider K arms such that, for $a \in [m]$, $x_a = e_1 + e_a$, $x_{m+1} = \cos(\omega)e_1 + \sin(\omega)e_{m+1}$, and for any $a > m+1$, $x_a = e_{a-1}$, and $\theta = e_1$, where $(e_i)_{i \in [K-1]}$ are the canonical basis of $\mathbb{R}^{K-1}$. This construction extends the usual "hard" instance for BAI in linear bandits ($m = 1$), originally proposed by Soare et al. (2014) where one considers three arms, of respective feature vectors $x_1 = [1,0]^\top, x_2 = [0,1]^\top$ and $x_3 = [\cos(\omega), \sin(\omega)]^\top$, and $\theta = e_1$, where $\omega \in (0, \frac{\pi}{2})$. Arm 1 is the best arm, and as ω decreases, it becomes harder to discriminate between arms 1 and 3, and as such, pulling suboptimal arm 2 might be useful.

Drug repurposing instance Our drug-scoring function relies on the simulation of the treatment effect on gene activity via a Boolean network (Kauffman, 1969). The details are laid out in Appendix B. When a drug is selected, a reward is generated by applying the simulator to the gene activity profile of a patient chosen at random, and computing the cosine score between the post-treatment gene activity profile and the healthy gene activity profile. The higher this score is, the most similar the final treated patient and healthy samples are. We consider as drug features the treated gene activity profiles, which belong to $\mathbb{R}^{71}$.

Results Each figure reports the empirical distribution of the sample complexity (estimated over 500 runs) for different algorithms. In all the experiments, the empirical error is always below 5%, and is reported in Table 3.

Figure 1(a) reports the results for a classic problem with $m = 2, K = 4$ and $N = 3$. We compare two algorithms that are not using the arm features, LUCB and UGapE, with GIFA algorithms based on contextual indices. We investigate the use of individual and paired gap indices, different selection rules, and the two stopping rules τ^{LUCB} and τ^{UGapE} in both m-LinGapE and LinGIFA. We observe that there is almost no difference in the use of either stopping rules, but using paired indices leads to noticeably better sample complexity than individual ones. As noticed in Xu et al. (2018) for $m = 1$, m-LinGapE with the optimized and the greedy sampling rules have similar performance. In addition, using the greedy or optimized rules leads to slightly better performance compared to the largest variance rule. More importantly, we observe that linear bandit algorithms largely outperform their classical counterparts, even on a rather easy instance ($\omega = \frac{\pi}{6}$, hence $\mu_m - \mu_{m+1} \approx 0.13$).

Figure 1(b) shows the empirical sample complexity on a hard Top-1 classic problem ($\omega = 0.1$, $\mu_1 - \mu_2 \approx 0.005$) which allows to compare our algorithms to the state-of-the-art LinGame algorithm. $m-$LinGapE and LinGIFA have a better performance than LinGame, which is encouraging, even if our algorithms do not have any optimality guarantees in theory.

Figure 1(c) displays the result for a drug repurposing instance for epilepsy, including $K = 10$ arms (drugs) with 5 anti-epileptic and 5 pro-convulsant drugs, hence the choice $m = 5$ (a close-up is shown in Appendix B without LUCB). Although the linear dependency between rewards and features may be violated in this real-world example, taking into account features still helps in considerably reducing the sample complexity (approximately by a factor $\frac{1}{2}$).

6 DISCUSSION

To our knowledge, we have provided the first unified framework and fully adaptive algorithms for linear Top-m. Our theoretical analysis shows our algorithms do not perform any worse than their clas-

Table 3: Error frequencies rounded to 5 decimal places, for each Top-m and BAI algorithm (averaged across 500 runs). Proposed algorithms' names are in bold type. Each column corresponds to one figure.

Algorithm	(1)a	(1)b	(1)c
m-LinGapE (greedy)	0.0	0.044	0.0
m-LinGapE (optimized)	0.0	-	0.0
LinGIFA (largest var.)	0.0	-	0.0
LinGIFA (τ^{LUCB}, largest var.)	0.0	-	-
LinGIFA (τ^{LUCB}, greedy)	0.0	-	-
LinGIFA (greedy)	0.0	0.0	0.0
LinGIFA (individual indices)	0.0	-	-
LUCB (largest var.)	0.0	-	0.0
LUCB (sampling both arms)	-	-	0.0
UGapE	0.0	-	-
LinGame	-	0.0	-

sical counterparts. Code is publicly available at `https://github.com/clreda/linear-top-m`. However, in real life and in our drug repurposing instance, the linear dependency between features and rewards does not hold. A future direction of our work would be dealing with such model misspecification. Another perspective would be the analysis of the greedy sampling rule. Indeed, this sampling rule leads to more efficient algorithms in our linear experiments.

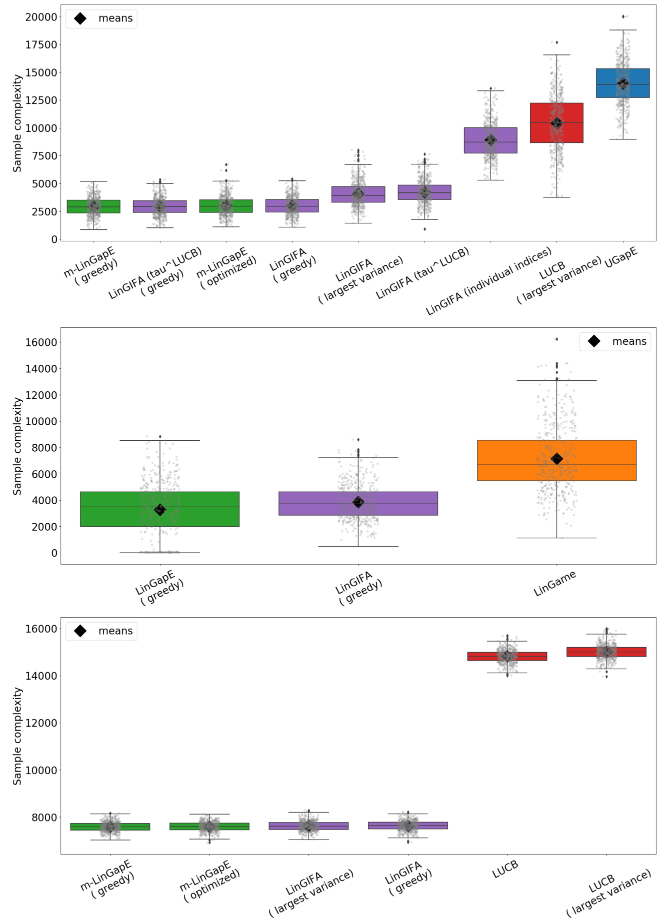

Figure 1: From top to bottom: classic instances (a) $K = 4, \omega = \frac{\pi}{6}, m = 2$; (b) $K = 3, \omega = 0.1, m = 1$; (c) drug repurposing instance $K = 10$, $m = 5$. Lines are quantiles and jittered individual outcomes are plotted in grey.

ACKNOWLEDGEMENTS

The authors thank the anonymous reviewers for their feedback, Xuedong Shang for his valuable input for the implementation of LinGame, and Baptiste Porte for his help on building the drug repurposing instance applied to epilepsy. This work was supported by the Institut National pour la Santé et la Recherche Médicale (INSERM, France), the Centre National de la Recherche Scientifique (CNRS, France), the Université Sorbonne Paris Nord, the Université de Paris, France, the French Ministry of Higher Education and Research [ENS.X19RDTME-SACLAY19-22] (C. R.), the French National Research Agency [ANR-18-CE17-0009-01 and ANR-18-CE37-0002-03] (A. D.-D.), [ANR-19-CE23-0026-04] (E. K.), and by the "Digital health challenge" Inserm CNRS joint program (C. R., E. K. and A. D.-D.).

References

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.*, 3(null):397–422, 2003. ISSN 1532-4435.

Adam S Brown and Chirag J Patel. A standard database for drug repositioning. *Scientific data*, 4:170029, 2017.

Séastian Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Fondations and Trends in Machine Learning*, 5(1):1–122, 2012.

Sébastian Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 258–265, Atlanta, Georgia, USA, 2013. PMLR. URL http://proceedings.mlr.press/v28/bubeck13.html.

Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *Proceedings of the 20th International Conference on Algorithmic Learning Theory*, ALT'09, page

23–37, Berlin, Heidelberg, 2009. Springer-Verlag. ISBN 3642044131.

Lijie Chen, Jian Li, and Mingda Qiao. Nearly Instance Optimal Sample Complexity Bounds for Top-k Arm Selection. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 101–110, Fort Lauderdale, FL, USA, 2017. PMLR. URL http://proceedings.mlr.press/v54/chen17a.html.

Lijun Cheng and Lang Li. Systematic quality control analysis of lincs data. *CPT: pharmacometrics & systems pharmacology*, 5(11):588–598, 2016.

Neil R Clark, Kevin S Hu, Axel S Feldmann, Yan Kou, Edward Y Chen, Qiaonan Duan, and Avi Ma'ayan. The characteristic direction: a geometrical approach to identify differentially expressed genes. *BMC bioinformatics*, 15(1):79, 2014.

Rémy Degenne and Wouter M Koolen. Pure exploration with multiple correct answers. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 14591–14600. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/60cb558c40e4f18479664069d9642d5a-Paper.pdf.

Rémy Degenne, Pierre Ménard, Xuedong Shang, and Michał Valko. Gamification of pure exploration for linear bandits. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2432–2442. PMLR, 13–18 Jul 2020. URL http://proceedings.mlr.press/v119/degenne20a.html.

Andrée Delahaye-Duriez, Prashant Srivastava, Kirill Shkura, Sarah R Langley, Liisi Laaniste, Aida Moreno-Moral, Bénédicte Danis, Manuela Mazzuferi, Patrik Foerch, Elena V Gazina, et al. Rare and common epilepsies converge on a shared gene regulatory network providing opportunities for novel antiepileptic drug discovery. *Genome biology*, 17(1): 1–18, 2016.

Qiaonan Duan, St Patrick Reid, Neil R Clark, Zichen Wang, Nicolas F Fernandez, Andrew D Rouillard, Ben Readhead, Sarah R Tritsch, Rachel Hodos, Marc Hafner, et al. L1000cds 2: Lincs l1000 characteristic direction signatures search engine. *NPJ systems biology and applications*, 2:16015, 2016.

Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 10667–10677. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/8ba6c657b03fc7c8dd4dff8e45defcd2-Paper.pdf.

Victor Gabillon, Mohammad Ghavamzadeh, and Alessandro Lazaric. Best arm identification: A unified approach to fixed budget and fixed confidence. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 3212–3220. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper/2012/file/8b0d268963dd0cfb808aac48a549829f-Paper.pdf.

Matthew Hoffman, Bobak Shahriari, and Nando Freitas. On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In Samuel Kaski and Jukka Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 365–374, Reykjavik, Iceland, 2014. PMLR. URL http://proceedings.mlr.press/v33/hoffman14.html.

Thomas J Hwang, Daniel Carpenter, Julie C Lauffenburger, Bo Wang, Jessica M Franklin, and Aaron S Kesselheim. Failure of investigational drugs in late-stage clinical development and publication of trial results. *JAMA internal medicine*, 176(12):1826–1833, 2016.

Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, ICML'12, page 227–234, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.

Stuart A Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of theoretical biology*, 22(3):437–467, 1969.

Émilie Kaufmann. *Analyse de stratégies bayésiennes et fréquentistes pour l'allocation séquentielle de ressources.* PhD thesis, 2014. URL http://www.theses.fr/2014ENST0056. Thèse de doctorat dirigée par Cappé, Olivier et Garivier, Aurélien, Signal et images Paris, ENST 2014.

Émilie Kaufmann and Shivaram Kalyanakrishnan. Information complexity in bandit subset selection. In Shai Shalev-Shwartz and Ingo Steinwart, editors, *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 228–251, Princeton, NJ, USA, 2013. PMLR. URL http://proceedings.mlr.press/v30/Kaufmann13.html.

Aliyu Musa, Laleh Soltan Ghoraie, Shu-Dong Zhang, Galina Glazko, Olli Yli-Harja, Matthias Dehmer, Benjamin Haibe-Kains, and Frank Emmert-Streib. A review of connectivity map and computational approaches in pharmacogenomics. *Briefings in bioinformatics*, 19(3):506–523, 2018.

Clémence Réda and Wilczyński. Automated inference of gene regulatory networks using explicit regulatory modules. *Journal of Theoretical Biology*, 486:110091, 2020. ISSN 0022-5193. doi: https://doi.org/10.1016/j.jtbi.2019.110091. URL http://www.sciencedirect.com/science/article/pii/S0022519319304606.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, pages 828–836. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper/2014/file/f387624df552cea2f369918c5e1e12bc-Paper.pdf.

Aravind Subramanian, Rajiv Narayan, Steven M Corsello, David D Peck, Ted E Natoli, Xiaodong Lu, Joshua Gould, John F Davis, Andrew A Tubelli, Jacob K Asiedu, et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. *Cell*, 171(6):1437–1452, 2017.

Damian Szklarczyk, John H Morris, Helen Cook, Michael Kuhn, Stefan Wyder, Milan Simonovic, Alberto Santos, Nadezhda T Doncheva, Alexander Roth, Peer Bork, et al. The string database in 2017: quality-controlled protein–protein association networks, made broadly accessible. *Nucleic acids research*, page gkw937, 2016.

Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160, 1972.

Liyuan Xu, Junya Honda, and Masashi Sugiyama. A fully adaptive algorithm for pure exploration in linear bandits. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 843–851. PMLR, 2018. URL http://proceedings.mlr.press/v84/xu18d.html.

Boyan Yordanov, Sara-Jane Dunn, Hillel Kugler, Austin Smith, Graziano Martello, and Stephen Emmott. A method to identify and analyze biological programs through automated reasoning. *NPJ systems biology and applications*, 2:16010, 2016.