Identifying near-optimal decisions in linear-in-parameter bandit models with continuous decision sets

Abstract

We consider an online optimization problem in a bandit setting in which a learner chooses decisions from a continuous decision set at discrete decision epochs, and receives noisy rewards from the environment in response. While the noise samples are assumed to be independent and sub-Gaussian, the mean reward at each epoch is a fixed but unknown linear function of a feature vector, which depends on the decision through a known (and possibly nonlinear) feature map. We study the problem within the framework of best-arm identification with fixed confidence, and provide a template algorithm for approximately learning the optimal decision in a PAC setting. More precisely, the template algorithm samples the decision space till a stopping condition is met, and returns a subset of decisions such that, with the required confidence, every element of the subset is approximately optimal for the unknown mean reward function. We provide a sample complexity bound for the template algorithm and then specialize it to the case where the mean-reward function is a univariate polynomial of a single decision variable. We provide an implementable algorithm for this case by explicitly instantiating all the steps in the template algorithm. Finally, we provide experimental results to demonstrate the efficacy of our algorithms.

1 INTRODUCTION

Multi-arm bandits have proved to be a fertile setting for studying various aspects of exploration and exploitation in sequential decision-making problems. While the regret minimization setting probes trade-offs between exploration and exploitation [Bubeck and Cesa-Bianchi, 2012], the pure exploration setting examines efficient exploration for maximizing information gain [Even-Dar et al., 2002, Bubeck et al., 2009]. Best arm identification (BAI) is one example of a pure exploration task where the learner seeks to identify the best arm through exploration. BAI is itself studied in two settings, namely, the fixed budget setting and the fixed confidence setting. In the fixed budget setting, the learner seeks to minimize the probability of misidentifying the optimal arm over a fixed number of trials [Audibert et al., 2010]. In contrast, in the fixed confidence setting, the aim of the learner is to minimize the number of trials needed to identify the optimal arm with a given level of confidence [Even-Dar et al., 2003].

Inclusion of additional structure in the reward environment adds a new dimension to the bandit problem. One structured bandit setting that has been widely considered in the literature is that of linear bandits. In a multi-arm linear bandit problem, each arm is associated with a feature vector in a finite-dimensional real vector space, and the mean reward of the arm is an unknown linear function of the feature vector. A more general version of the linear bandit problem results when the set of “arms" is a subset, not necessarily finite, of a real vector space. Unlike in the case of a multi-arm bandit where pulling one arm provides no information about another, the linear structure of the mean reward in a linear bandit problem opens up the possibility of learning optimal decisions even while sampling suboptimal ones.

The linear bandit problem has received significant attention in the regret minimization setting, both the case of finite arms as well as continuous decision sets [Auer, 2002, Dani et al., 2008, Agrawal and Goyal, 2013, Bartlett et al., 2008]. In contrast, the pure exploration setting for linear bandit problems has started gaining attention only relatively recently [Soare et al., 2014, Degenne et al., 2020, Garivier and Kaufmann, 2016, Yang and Tan, 2021, Karnin, 2016, Xu et al., 2018, Tao et al., 2018, Jedra and Proutiere, 2020a]. What is more, except for the specific case of a spherical decision set considered in Jedra and Proutiere [2020a], the literature on pure exploration in linear bandits has so far focused on the case of finite decision sets only.

Submitted to the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022). **To be used for reviewing only**.

In this paper, we consider a bandit problem in which the mean reward is an unknown linear function of a feature vector that depends on the decision through a known, but possibly nonlinear, feature map. Furthermore, we do not assume the decision set to be finite. While BAI algorithms in the finite arm case seek to find the best arm with high confidence, finding the best decision from a continuum of decisions can be prohibitively expensive. Hence we consider a (ε, δ)-PAC formulation, where the goal of the learner is to find a set of points which are ε-optimal with probability at least $1 - \delta$.

By building on the work of Soare et al. [2014], Jedra and Proutiere [2020a], Kaufmann et al. [2016], we provide a lower bound on the sample complexity of (ε, δ)-PAC algorithms in Section 3. Next, we use the notion of volumetric spanners [Hazan and Karnin, 2016] to devise a simple algorithm template, VSBAI, for BAI in our setting in Section 4. We prove VSBAI to be (ε, δ)-PAC, and provide upper bounds on its sample complexity. In Section 5, we consider the case where the mean reward is a polynomial function of a single decision variable. We show that, in this case, a volumetric spanner can be computed using convex optimization, and indicate how the algorithm template VSBAI can be instantiated for BAI under polynomial rewards. Finally, we present experimental results in Section 6.

Before describing the problem setup in Section 2, we introduce some notation used throughout the paper. We use $\mathbb{R}$ and $\mathbb{Z}_+$ to denote the set of real numbers and positive integers, respectively. The 1-norm and 2-norm on $\mathbb{R}^n$ are denoted by $\|\cdot\|_1$ and $\|\cdot\|_2$, respectively. Given a function $g : \mathcal{D} \to \mathbb{R}$ and $\varepsilon > 0$, $s \in \mathcal{D}$ is ε-optimal for g if $g(s') \leq g(s) + \varepsilon$ for all $s' \in \mathcal{D}$. A set $\mathcal{D}' \subseteq \mathcal{D}$ is ε-optimal for g if every element of $\mathcal{D}'$ is ε-optimal for g. Finally, $\|g\|_\infty$ denotes the sup norm of a real-valued function when its domain is clear from the context.

2 PROBLEM SETUP

We consider a bandit optimization setting in which a learner interacts with an environment at discrete decision epochs $t = 1, 2, \ldots$. At each period $t \in \mathbb{Z}_+$, the learner chooses a decision s_t from a compact decision set $\mathcal{D} \in \mathbb{R}^d$ and receives a noisy reward $y_t = \mu^{\mathrm{T}} x_t + \eta_t$, where the f-dimensional feature vector $x_t = \phi(s_t)$ is related to the decision s_t through a continuous feature map $\phi : \mathcal{D} \to \mathbb{R}^f$, $\mu \in \mathbb{R}^f$ is a parameter vector, and $\{\eta_t\}_{t \in \mathbb{Z}_+}$ is a noise sequence. Our reward model is thus given by

$$y_t = g_\mu(s_t) + \eta_t, \; t \in \mathbb{Z}_+, \tag{1}$$

where, for each $\theta \in \mathbb{R}^f$, $g_\theta : \mathcal{D} \to \mathbb{R}$ is defined by $g_\theta(s) = \theta^{\mathrm{T}} \phi(s)$. Our results make use of one or the other of the following two assumptions on the noise sequence.

Assumption 1. The noise sequence $\{\eta_t\}_{t \in \mathbb{Z}_+}$ is a sequence of zero mean i.i.d σ-sub Gaussian random variables for some $\sigma > 0$. Specifically, for each $i \in \mathbb{Z}_+, \eta_i$ satisfies $\mathbb{E}(e^{t\eta_i}) \leq e^{\frac{\sigma^2 t^2}{2}}$ for all $t \in \mathbb{R}$.

Assumption 2. The noise sequence $\{\eta_t\}_{t \in \mathbb{Z}_+}$ is a sequence of zero mean i.i.d. Gaussian random variables with variance σ^2 for some $\sigma > 0$.

We assume that, in the case of either of the two assumptions above, the learner knows σ. In addition, she also has access to the feature map ϕ. However, the parameter vector μ is unknown to the learner. Finally, without any real loss of generality, we assume that $\phi(\mathcal{D})$ is not contained in any proper linear subspace of $\mathbb{R}^f$.

In a best-arm identification setting, the learner's goal is to identify a maximizer s^* of g_μ by using the observations $\{(s_i, y_i)\}_{i=1}^T$ collected over a decision horizon T. However, the presence of noise makes it impossible to identify an optimizer with certainty over a finite horizon. Hence it is standard practice in the literature to seek an algorithm that returns a set that contains the desired optimizer to a high level of confidence. Such an algorithm typically comprises of a sampling rule π that determines the decision $s_t \in \mathcal{D}$ to explore at time t given the history of observations up to $t - 1$, a stopping rule that decides if the exploration conducted so far is sufficient, and an estimation rule that computes a set that contains the desired optimizer to a high level of confidence. We make these ideas more precise in the next section.

3 (ε, δ)-PAC ALGORITHMS AND THEIR SAMPLE COMPLEXITY

Complexity lower bounds on algorithms for best arm identification in a PAC setting have been studied before for the case where the decision set is finite Soare et al. [2014], Jedra and Proutiere [2020a], Degenne et al. [2020], Kaufmann et al. [2016], Xu et al. [2018]. While the analysis we present below follows similar ideas, the continuous nature of the decision set makes it necessary to formally define the elements mentioned above using a little more machinery.

To this end, we note that a sampling rule could also make use of internal randomization in addition to the past history of decisions and rewards. It is easy to see that any randomization scheme requiring n random variables at each decision epoch can be implemented using n i.i.d. samples of a random variable uniformly distributed on the unit interval. Hence, to represent a general sampling rule more formally, we consider the Cartesian product $\mathcal{S} \stackrel{\text{def}}{=} \mathcal{D} \times \mathbb{R} \times [0, 1]^n$, where $n \geq 0$ is a fixed integer. $\mathcal{S}$ is the set of triplets of decision, reward, and a set of n auxiliary quantities used for internal randomization. For each t, we denote by Ω_t the set of sequences in $\mathcal{S}$ of length t, and by Ω the set of all infinite sequences in $\mathcal{S}$. We use $h_t = \{(s_i, y_i, u_i)\}_{i=1}^t$ to denote a general sequence in Ω_t. We assume that $\mathcal{D}$ is a Borel set.

By forming products of the Borel σ-algebras of $\mathcal{D}$, $\mathbb{R}$ and $[0,1]^n$, we obtain a σ-algebra $\mathcal{F}_t$ on Ω_t for each t, as well as a σ-algebra $\mathcal{F}$ on Ω. Moreover, on letting $\mathcal{F}_0$ denote the trivial σ-algebra on Ω, we obtain a filtration $\{\mathcal{F}_t\}_{t=0}^{\infty}$ on Ω.

Next, we define a sampling rule π to be a sequence $\{\pi_t\}_{t\in\mathbb{Z}_+}$ along with a Borel measure λ on $[0,1]^n$, where π_1 is a stochastic kernel on $\mathcal{D}$ given $[0,1]^n$ and, for each $t>1$, π_t is a stochastic kernel on $\mathcal{D}$ given Ω_{t-1} and $[0,1]^n$. In other words, for $t>1$, the following holds: for each $h_{t-1}\in\Omega_{t-1}$ and $u\in[0,1]^n$, $\pi_t(\cdot|h_t,u)$ is a measure on the Borel σ-algebra of $\mathcal{D}$, while for each Borel subset A of $\mathcal{D}$, $\pi_t(A|\cdot,\cdot)$ is a Borel-measurable function on $\Omega_{t-1}\times[0,1]^n$. Informally speaking, λ is the measure used to sample an element of $[0,1]^n$ for any internal randomization used by the sampling rule while, for every $t>1$, the measure $\pi_t(\cdot|h_{t-1},u)$ describes the conditional distribution of the decision sampled at time t given the history $h_{t-1}\in\Omega_{t-1}$ up to time $t-1$ and the randomly sampled $u\in[0,1]^n$. A similar interpretation applies for $t=1$.

Any algorithm used by the learner can be represented by the tuple $\mathcal{A}=(n,\lambda,\pi,\tau,\mathbb{F})$, where n, λ and π are as described above, τ is a stopping time with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^{\infty}$ representing the stopping condition of the algorithm, and $\mathbb{F}$ is a set-valued map that maps each finite history in Ω to a subset of $\mathcal{D}$. The algorithm terminates at the random time τ and returns the set $\mathbb{F}(h_\tau)$ upon terminating.

It is natural to represent the environment as a stochastic kernel Q^μ on $\mathbb{R}$ given $\mathcal{D}$, such that the measure on $\mathbb{R}$ given by $Q^\mu(\cdot|s)$ describes the conditional distribution of the reward (1) given the decision s. The interaction between the algorithm and the environment induces Borel measures $\mathbb{P}^{\mathcal{A},\mu}$ on Ω and $\mathbb{P}_t^{\mathcal{A},\mu}$ on Ω_t for each $t\in\mathbb{Z}_+$ (see Proposition 7.28 of Bertsekas and Shreve [1996]).

Finally, given $\varepsilon>0$ and $\zeta\in\mathbb{R}^f$, we let $\mathcal{O}_\varepsilon(\zeta)\subseteq\mathcal{D}$ denote the set of decisions that are ε-optimal for the function g_ζ. We are now in a position to precisely define the class of algorithms that we consider as solutions to the problem described above.

Definition 3.1. Given $\varepsilon>0$ and $\delta\in(0,1)$, an algorithm $\mathcal{A}=(n,\lambda,\pi,\tau,\mathbb{F})$ is (ε,δ)-PAC for the environment (1) if the stopping time τ is finite $\mathbb{P}^{\mathcal{A},\mu}$-almost-surely and $\mathbb{P}^{\mathcal{A},\mu}(\{\arg\max_{s\in\mathcal{D}} g_\mu(s)\subseteq\mathbb{F}(h_\tau)\subseteq\mathcal{O}_\varepsilon(\mu)\})\geq 1-\delta$.

The expected sample complexity of an algorithm $\mathcal{A}=(n,\lambda,\pi,\tau,\mathbb{F})$ is the expected number of decisions explored by the algorithm till termination, and is simply given by $\mathbb{E}^{\mathcal{A},\mu}(\tau)$, where $\mathbb{E}^{\mathcal{A},\mu}(\cdot)$ denotes expectation under $\mathbb{P}^{\mathcal{A},\mu}$. Next, we provide a lower bound for the expected sample complexity of a (ε,δ)-PAC algorithm. To do so, we need one more notation. Given $\zeta\in\mathbb{R}^f$ and $\varepsilon>0$, the *ε-alternative of* ζ is the set $\mathrm{Alt}_\varepsilon(\zeta)=\{\zeta'\in\mathbb{R}^f : \mathcal{O}_\varepsilon(\zeta)\cap\mathcal{O}_\varepsilon(\zeta')=\varnothing\}$. We are now ready to state our lower bound. The proof, which builds on ideas given in Soare et al. [2014], Jedra and Proutiere [2020a], Kaufmann et al. [2016], is given in Appendix A.

Theorem 3.2. *Suppose Assumption 2 holds. Let $\varepsilon>0$ and $\delta\in(0,1)$, and suppose $\mathcal{A}=(n,\lambda,\pi,\tau,\mathbb{F})$ is a (ε,δ)-PAC algorithm for (1). Then*

$$\mathbb{E}^{\mathcal{A},\mu}(\tau)\geq\frac{2\sigma^2\ln\left(\frac{1}{2.4\delta}\right)}{\inf\limits_{\zeta\in\mathrm{Alt}_\varepsilon(\mu)}\|g_\mu-g_\zeta\|_\infty}. \tag{2}$$

4 VSBAI: AN ALGORITHM TEMPLATE

In this section, we present VSBAI, a general template for an (ε,δ)-PAC algorithm for the bandit optimization problem described in Section 2, and provide a sample complexity bound for it. We prefer to use the term template rather than an algorithm as some of the steps of the template can only be implemented if $\mathcal{D}$ and ϕ are specified.

VSBAI combinines two ideas, namely,

1. obtaining a ε-optimal set for g_μ from a uniform approximation for g_μ, and
2. efficiently obtaining a uniform approximation of g_μ by regressing the rewards on decisions sampled at points of a suitable exploration basis for $\mathcal{D}$.

We elaborate on each of these two aforementioned ideas next.

4.1 APPROXIMATE OPTIMIZERS FROM UNIFORM APPROXIMATIONS

The intuition behind our algorithm is illustrated in Figure 1 for the case where $d=1$. The thick solid curve in the figure depicts the graph of a uniform approximation $\hat{q}$ of an unknown function q represented by the thin solid curve. Suppose the uniform approximation error does not exceed $\frac{\varepsilon}{4}$ for some $\varepsilon>0$, that is, $\|q-\hat{q}\|_\infty<\frac{\varepsilon}{4}$ holds. The two dashed curves are graphs of the functions $\hat{q}\pm\frac{\varepsilon}{4}$, which serve as upper and lower bounds on the unknown function q. In other words, the graph of q must lie within the region bounded by the two dashed curves. In the figure, the approximation $\hat{q}$ achieves its maximum at $\hat{s}$, while s^* is the maximizer of q. The horizontal segment shown in the figure represents a set $\mathcal{D}'$ such that the approximation $\hat{q}$ does not fall below its maximum value $\hat{q}(\hat{s})$ by more than $\frac{\varepsilon}{2}$ on $\mathcal{D}'$. One can intuitively see from the figure that the set $\mathcal{D}'$ must contain the maximizer s^* of the unknown function q. Moreover, the absolute difference between the values of the unknown function q at any two points in the set $\mathcal{D}'$ cannot exceed ε. In other words, the set $\mathcal{D}'$ is ε-optimal for q.

The next proposition formalizes the intuition reflected in Figure 1. The proof is given in Appendix B.

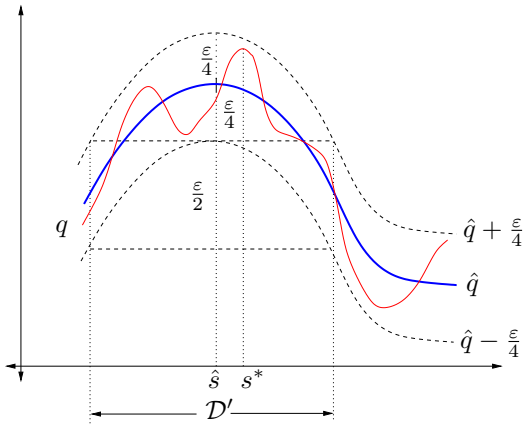

Figure 1: Obtaining ε-optimal points for q from its uniform $\frac{\varepsilon}{4}$-approximation $\hat{q}$

Proposition 4.1. *Let $\varepsilon > 0$, and suppose $q, \hat{q} : \mathcal{D} \to \mathbb{R}$ are such that $\|\hat{q} - q\|_\infty \leq \frac{\varepsilon}{4}$. Let $\hat{s} \in \arg\max_{s\in\mathcal{D}} \hat{q}(s)$. Then the set $\mathcal{D}' \stackrel{\text{def}}{=} \{s \in \mathcal{D} : \hat{q}(s) \geq \hat{q}(\hat{s}) - \frac{\varepsilon}{2}\}$ is ε-optimal for q, and contains $\arg\max_{s\in\mathcal{D}} q(s)$.*

4.2 UNIFORM APPROXIMATION OF THE REWARD FUNCTION

The Cauchy-Schwarz inequality gives $|g_{\hat{\mu}}(s) - g_\mu(s)| \leq \|\phi(s)\|_2 \|\hat{\mu} - \mu\|_2$ for every $s \in \mathcal{D}$ and $\hat{\mu} \in \mathbb{R}^f$. The compactness of $\mathcal{D}$ now shows that any estimate $\hat{\mu}$ of μ yields an uniform approximation $g_{\hat{\mu}}$ of g_μ. Hence, an obvious and popular means of obtaining an approximation to the unknown reward function g_μ is to estimate μ from observed decisions and rewards using least-squares regression. It is not surprising, therefore, that either ordinary least squares (OLS) or regularised least squares forms a part of almost every algorithm available for linear bandit problems in the stochastic as well as adversarial settings with finite arms or continuous arms. We briefly review OLS before proceeding.

At the end of t decision epochs, the learner has access to observations $\{(x_i, y_i)\}_{i=1}^t$, where $x_i = \phi(s_i)$ is the feature vector of the ith decision s_i, and y_i is the corresponding observed reward. Letting $X_t \stackrel{\text{def}}{=} [x_1, \ldots, x_t] \in \mathbb{R}^{f\times t}$ and $y^t \stackrel{\text{def}}{=} [y_1, \ldots, y_t]^{\mathrm{T}} \in \mathbb{R}^t$, the OLS estimate $\hat{\mu}_t$ of μ based on the data $\{(x_i, y_i)\}_{i=1}^t$ is obtained by solving $\min_{\hat{\mu}\in\mathbb{R}^f} \|y^t - X_t^{\mathrm{T}}\hat{\mu}\|_2^2$, and is given by

$$\hat{\mu}_t = (X_t X_t^{\mathrm{T}})^{-1} X_t y^t. \tag{3}$$

For deriving (3), it is assumed that X_t has rank f, which necessarily implies that $t \geq f$.

The parameter error $\hat{\mu}_t - \mu$ clearly depends on the choice of the decisions $s_1, \ldots, s_t$. Indeed, on letting $\eta^t = [\eta_1, \ldots, \eta_t]^{\mathrm{T}}$ denote the vector of noise samples till time t, it is easy to use (1) and (3) to show that

$$\hat{\mu}_t - \mu = (X_t X_t^{\mathrm{T}})^{-1} X_t \eta^t. \tag{4}$$

In a regret minimization setting, the decisions need to be chosen in an adaptive manner so that the required trade-off between exploration and exploitation can be achieved. Even in the pure exploration setting of best-arm identification in a finite multi-arm bandit problem, decisions have to be adaptive so that the exploration budget is diverted away from arms as and when they are revealed to be sub-optimal, since exploring one arm gives no information about another arm. In contrast, in the pure exploration setting that we are considering for the linear bandit problem, each decision that improves the estimate of μ also improves the accuracy of the approximation of g_μ all over the decision domain $\mathcal{D}$. This suggests the possibility of using non-adaptive (that is, deterministic) sampling of the decision space for the purpose of constructing an OLS-based approximation of g_μ.

In this case, it is natural to consider a *volumetric spanner* as a low variance exploration basis (as defined in Hazan and Karnin [2016]) for sampling the reward function. We review the necessary background next.

4.3 VOLUMETRIC SPANNERS

Suppose $L \geq 1$ and $m \geq f$. A (L, m)-*volumetric spanner* for $\phi(\mathcal{D}) \subseteq \mathbb{R}^f$ is a subset $\{x_1, \ldots, x_m\}$ of $\phi(\mathcal{D})$ such that, for every $z \in \mathcal{D}$, there exists $c_1, \ldots, c_m \in \mathbb{R}$ satisfying $z = c_1 x_1 + \cdots + c_m x_m$ and $c_1^2 + \cdots + c_m^2 \leq L^2$. Recall that, for every $z \in \mathbb{R}^f$ and $X \in \mathbb{R}^{f\times m}$ with $m \geq f$, $c = X^{\mathrm{T}}(XX^{\mathrm{T}})^{-1}z$ is the minimum-2-norm solution of the equation $Xc = z$. Hence, it follows from the definition that, if $\{x_1, \ldots x_m\}$ is a (L, m)-volumetric spanner for $\phi(\mathcal{D}) \subseteq \mathbb{R}^f$, then $\|X^{\mathrm{T}}(XX^{\mathrm{T}})^{-1}z\|_2 \leq L$ for all $z \in \phi(\mathcal{D})$, where $X = [x_1, \ldots, x_m] \in \mathbb{R}^{f\times m}$. In particular, if $m = f$, then $\|X^{-1}z\|_2 \leq L$ for every $z \in \phi(\mathcal{D})$. The last observation implies that there exists no (L, f)-volumetric spanner for $\phi(\mathcal{D})$ for $L < 1$. A $(1, m)$-volumetric spanner was called a volumetric spanner in Hazan and Karnin [2016] irrespective of m. Since the cardinality of the volumetric spanner will be required in our algorithm, we choose not to suppress it.

It will be convenient to define $p_1, \ldots, p_m \in \mathcal{D}$ to be (L, m)-*volumetric points* for the pair $(\phi, \mathcal{D})$ if $\{\phi(p_1), \ldots, \phi(p_m)\}$ is a (L, m)-volumetric spanner for $\phi(\mathcal{D})$.

Since a volumetric spanner forms a critical component of the algorithm that we present in the next subsection, it is important to consider the existence of such spanners as well as algorithms for computing them. We start with an easy observation. Our assumption that the set $\phi(\mathcal{D})$ of feature vectors is not contained in a proper linear subspace of $\mathbb{R}^f$ implies that $\phi(\mathcal{D})$ contains a set of f linearly indepen-

dent vectors. Since $\phi(\mathcal{D})$ is compact, it is easy to see that any linearly independent subset of cardinality f will serve as a (L, f) volumetric spanner for sufficiently large L. It is separately known that, being compact, $\phi(\mathcal{D})$ possesses a $(1, m)$-volumetric spanner for some $m \leq 12f$. In addition, if $\phi(\mathcal{D})$ is finite, then the aforementioned volumetric spanner can be constructed in polynomial time (see Theorem 3 of Hazan and Karnin [2016] for both facts above as well as additional details). We will see later that the sample complexity of our algorithm can be improved if a (L, m)-volumetric spanner with lower values of L and m is chosen. In this context, the following bound proved in Appendix B is of interest.

Lemma 4.2. *If $C \subset \mathcal{D}$ is a (L, m)-volumetric spanner for $\phi(\mathcal{D})$ for some $L \geq 1$ and $m \geq f$, then $L^2 m \geq f$.*

The lower bound in Lemma 4.2 is achieved by a $(1, f)$-volumetric spanner.

4.4 VSBAI: DESCRIPTION AND ANALYSIS

The template algorithm VSBAI that we present requires a set of (L, m)-volumetric points $\{p_1, \ldots, p_m\}$ for the pair $(\phi, \mathcal{D})$, for some $L \geq 1$ and $m \geq f$. The algorithm proceeds in rounds with each round consisting of m decision epochs. In each round, the algorithm picks the points $\{p_1, \ldots, p_m\}$ in sequence as the decisions for that round. In the notation of section 3, the template algorithm is given by a tuple $\mathcal{A}^* = (n^*, \lambda^*, \pi^*, \tau^*, \mathbb{F}^*)$ whose sampling rule π^* is defined by

$$\pi_t^*(\cdot|u) = \delta_{p_i}(\cdot),\ i = 1 + (t \bmod m), \tag{5}$$

for every $t \in \mathbb{Z}_+$ and $u \in [0, 1]^{n^*}$, where $\delta_s(\cdot)$ denotes the Dirac measure at $s \in \mathcal{D}$. Note that the sampling rule π^* is deterministic, and hence the choices of n^* and λ^* are immaterial.

As described in subsection 4.2, our template algorithm $\mathcal{A}^*$ involves obtaining successively better uniform approximations of g_μ using a sequence of OLS estimates of μ obtained through (3). Our next result gives a high probability bound on the uniform error with which the estimate $g_{\hat{\mu}_{km}}$ obtained after k rounds approximates g_μ. The proof is given in Appendix C.

Proposition 4.3. *Consider an algorithm $\mathcal{A}^*$ whose sampling rule is described by (5). Let $k \in \mathbb{Z}_+$ and $\varepsilon > 0$. If Assumption 1 holds, then*

$$\mathbb{P}^{\mathcal{A}^*,\mu}(\|g_{\hat{\mu}_{km}} - g_\mu\|_\infty > \varepsilon) \leq \beta\left(k, \frac{\varepsilon}{L}\right), \tag{6}$$

where

$$\beta(k, \varepsilon) = \beta_{\mathrm{SG}}(k, \varepsilon) \stackrel{\mathrm{def}}{=} 2^{\frac{f}{2}} \exp\left(-\frac{k\varepsilon^2}{4\sigma^2}\right). \tag{7}$$

On the other hand, if Assumption 2 holds, then (6) holds with

$$\beta(k, \varepsilon) = \beta_{\mathrm{G}}(k, \varepsilon) \stackrel{\mathrm{def}}{=} 2 \exp\left(-\frac{k\varepsilon^2}{2\sigma^2}\right). \tag{8}$$

Propositions 4.1 and 4.3 immediately suggest the stopping criterion that yields an (ε, δ)-PAC algorithm under the sampling rule described by (5). Indeed, by Proposition 4.3, choosing

$$\tau^* = \inf\left\{km : \beta\left(k, \frac{\varepsilon}{4L}\right) < \delta\right\} \tag{9}$$

ensures that, with probability at least $1 - \delta$, the uniform approximation condition required by Proposition 4.1 holds with $q = g_\mu$ and $\hat{q} = g_{\hat{\mu}_{\tau^*}}$. Letting $\mathcal{D}_{\tau^*} = \mathbb{F}(h_{\tau^*})$ to be the set $\mathcal{D}'$ in Proposition 4.1 then ensures that $\mathcal{D}_{\tau*}$ is ε-optimal for g_μ with the same probability. The resulting algorithm is given as Algorithm 1 below.

In Algorithm 1, β is taken to be given by (7) or (8) accordingly as assumption 1 or 2 holds. Also, the steps at lines 10 and 12 in the algorithm come from (21) in Appendix C.

Algorithm 1 VSBAI

1: **Input:** $\varepsilon > 0$, $\delta \in (0, 1)$, sub-Gaussianity parameter σ, (L, m)-volumetric points $p_1, \ldots, p_m$ for $(\phi, \mathcal{D})$
2: Set $B_{L,m} = [\phi(p_1), \ldots, \phi(p_m)]$
3: Initialize $k \leftarrow 1$, $r \leftarrow 0$
4: Set STOP = False
5: **while** STOP == False **do**
6: Initialize reward vector $\bar{y}^k = []$
7: **for** $t = 1, \ldots, m$, **do**
8: Apply decision $s_{(k-1)m+t} \leftarrow p_t$
9: Observe reward $y_{(k-1)m+t}$
10: Augment reward vector $\bar{y}^k \leftarrow [(\bar{y}^k)^{\mathrm{T}}; y_{(k-1)f+t}]^{\mathrm{T}}$
11: **end for**
12: Update total reward vector $r \leftarrow r + \bar{y}^k$
13: **if** $\beta(k, \frac{\varepsilon}{4L}) < \delta$ **then**
14: STOP = True
15: **else**
16: $k = k + 1$
17: **end if**
18: **end while**
19: $\tau^* \leftarrow km$
20: $\hat{\mu}_{\tau^*} \leftarrow \frac{1}{k}(B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}B_{L,m}r$
21: Pick $\hat{s} \in \arg\max_{s \in \mathcal{D}} g_{\hat{\mu}_{\tau^*}}(s)$.
22: $\mathcal{D}_{\tau^*} = \{s \in \mathcal{D} : g_{\hat{\mu}_{\tau^*}}(s) \geq g_{\hat{\mu}_{\tau^*}}(\hat{s}) - \frac{\varepsilon}{2}\}$
23: **Output:** $\mathcal{D}_{\tau^*}$

The main result of this section given below states that VSBAI is (ε, δ)-PAC.

Theorem 4.4. *If Assumption 1 holds, then Algorithm 1 terminates in at most $\tau^* \leq m[1 - 64L^2\sigma^2\varepsilon^{-2}\ln(2^{-\frac{f}{2}}\delta)]$ decision epochs. If Assumption 2 holds, then Algorithm 1 terminates in at most $\tau^* \leq m[1 - 32L^2\sigma^2\varepsilon^{-2}\ln(2^{-1}\delta)]$ decision epochs. In either case, with $\mathbb{P}^{\mathcal{A}^*,\mu}$-probability at least $1-\delta$, the set $\mathcal{D}_{\tau^*}$ returned by the algorithm is ε-optimal for g_μ and contains all the maximizers of g_μ. In particular, Algorithm 1 is (ε,δ)-PAC.*

Proof. Let τ^* be as computed by Algorithm 1, and let $k = \tau^*/m$. The upper bound for τ^* comes from using (7) or (8) in the stopping condition (9). Next, consider the event $\mathcal{E} = \{\|g_{\hat{\mu}_\tau} - g_\mu\|_\infty > \frac{\varepsilon}{4}\}$. By Proposition 4.3 and the definition (9) of τ^*, it follows that $\mathbb{P}^{\mathcal{A}^*,\mu}(\mathcal{E}) < \beta(k, \frac{\varepsilon}{4L}) < \delta$. Proposition 4.1 now implies that, on the complement of the event $\mathcal{E}$, $\mathcal{D}_{\tau^*}$ is ε-optimal for g_μ and contains all the maximizers of g_μ. This completes the proof. □

Note that three critical steps in the algorithm depend on the pair $(\phi, \mathcal{D})$, namely, computation of the (L, m) volumetric points used as inputs to the algorithm, computation of an optimizer $\hat{s}$ for the approximation $g_{\hat{\mu}_{\tau^*}}$ at line 21, and computation of the set $\mathcal{D}_{\tau^*}$ at line 22 of the algorithm. Hence we view the algorithm more as a template requiring the two aforementioned steps to be worked out for specific problem instances. We present a simple example considered in Jedra and Proutiere [2020a] to illustrate these steps.

4.5 LINEAR BANDIT ON THE UNIT SPHERE

Let $f > 1$, and choose $\mathcal{D}$ to be the unit sphere $\mathrm{S}^{f-1} \stackrel{\text{def}}{=} \{s \in \mathbb{R}^f : \|s\|_2 = 1\}$. Let $\phi : \mathrm{S}^{f-1} \to \mathbb{R}^f$ be the inclusion map. Then the reward function in (1) becomes $g_\mu(s) = \mu^{\mathrm{T}} s$.

Any set of f orthonormal vectors is seen to be a set of $(1, f)$-volumetric points for the pair (ϕ, S^{f-1}). For every non-zero $\theta \in \mathbb{R}^f$, $\arg\max_{s\in\mathrm{S}^{f-1}} g_\theta$ equals $\{\|\theta\|_2^{-1}\theta\}$. Line 21 of Algorithm 1 thus returns $\hat{s} = \|\hat{\mu}_{\tau^*}\|_2^{-1}\hat{\mu}_{\tau^*}$, while the set $\mathcal{D}_{\tau^*}$ at line 22 of Algorithm 1 is given by the "spherical cap" $\{s \in \mathrm{S}^{f-1} : \hat{s}^{\mathrm{T}} s \geq 1 - \frac{\varepsilon}{2\|\hat{\mu}_\tau\|_2}\}$.

Under Assumption 2, Theorem 4 of Jedra and Proutiere [2020a] gives a lower bound for the sample complexity of any (ε,δ)-PAC algorithm $\mathcal{A}$ for the case of the unit sphere considered here. On using inequality (3) of Kaufmann et al. [2016], the lower bound given by Jedra and Proutiere [2020a] may be written as

$$\mathbb{E}^{\mathcal{A},\mu}(\tau) \geq \frac{\sigma^2(f-1)}{20\varepsilon\|\mu\|_2}\ln\left(\frac{1}{2.4\delta}\right) \tag{10}$$

for $\varepsilon < \|\mu\|_2/5$. Jedra and Proutiere [2020a] also provide an algorithm for this case, and show that the sample complexity of their algorithm recovers the dependence on ε, f and δ seen in the lower bound (10) asymptotically as $\delta \to 0$ (see Theorem 5 of Jedra and Proutiere [2020a]). Interestingly, the sampling rule given by Jedra and Proutiere [2020a] for their algorithm involves choosing f orthogonal vectors in a round-robin manner just as mentioned above. However, their stopping rule is more intricate.

On using $L = 1$ and $m = f$, the upper bound provided by Theorem 4.4 for Algorithm 1 under Assumption 2 reduces to

$$\tau^* \leq f\left[1 + \frac{32\sigma^2}{\varepsilon^2}\ln\left(\frac{2}{\delta}\right)\right]. \tag{11}$$

On comparing (10) and (11), we see that while the dependence of the sample complexity of Algorithm 1 on f and δ compares favourably with the lower bound (10), the dependence on ε does not, at least for small values of ε. This could indicate that either the lower bound is conservative (for $\delta > 0$), or that Algorithm 1 is sub-optimal. Closing this gap remains an open problem.

It is easy to see from Theorem 4.4 that the best sample complexity for Algorithm 1 results when $L = 1$ and $m = f$, that is, when a set of $(1, f)$-volumetric points is available for the pair $(\phi, \mathcal{D})$. The unit sphere example considered in this subsection provided a simple setting in which a set of $(1, f)$-volumetric points is available. In the next section, We will see a nontrivial setting where such a set of volumetric points exists, and can be easily computed.

5 UNIVARIATE DECISION VARIABLE WITH POLYNOMIAL REWARD

As a concrete instance of the general problem setup described in Section 2, we consider the case where the reward function g_μ in (1) is a univariate polynomial of degree $f-1 > 0$ on an interval $[p_{\min}, p_{\max}] \subset \mathbb{R}$ for some $p_{\max} > p_{\min}$. To cast this case of polynomial rewards in our general setup, we let $\mathcal{D} \stackrel{\text{def}}{=} [p_{\min}, p_{\max}]$ and define $\phi : [p_{\min}, p_{\max}] \to \mathbb{R}^f$ by $\phi(s) \stackrel{\text{def}}{=} [1, s, \ldots, s^{f-1}]^{\mathrm{T}}$. Then, for each $\theta \in \mathbb{R}^f$, g_θ is the univariate polynomial in s of degree $f-1$ with coefficients given by the parameter vector θ. Our next result shows that a set of $(1, f)$ volumetric points for the pair $\phi(\mathcal{D})$ exists. The proof is given in Appendix D.

Proposition 5.1. *Suppose $p_{\min} \leq p_1 \leq \cdots \leq p_f \leq p_{\max}$. Then the following two statements are equivalent.*

1. *The points $p_1, \ldots, p_f \in \mathcal{D}$ are $(1, f)$ volumetric points for the pair $(\phi, \mathcal{D})$.*
2. *The points $p_1, \ldots, p_f$ satisfy $p_{\min} = p_1 < p_2 \cdots < p_f = p_{\max}$ and*

$$\sum_{1\leq j\leq f, j\neq i} \frac{1}{p_i - p_j} = 0,\ i = 2, \ldots, f-1. \tag{12}$$

Equations (12) also appear in Amballa et al. [2021], where it is shown that (12) provide necessary and sufficient conditions for the points $\phi(p_1), \ldots, \phi(p_f)$ to form a barycentric spanner for the set $\phi(\mathcal{D})$. Proposition 5.1 above thus implies that, in the case of univariate polynomial reward functions, a barycentric spanner is also a volumetric spanner. Amballa et al. [2021] also show that the equations (12) possess a unique solution, and this solution may be computed efficiently either by numerically solving the algebraic equations (12) or by solving a convex optimization problem. Furthermore, volumetric points for the general case $\mathcal{D} = [p_{\min}, p_{\max}]$ can be easily recovered from volumetric points for the special case $\mathcal{D} = [0, 1]$. This means that, effectively, the solution of (12) needs to be computed just once for a given f.

Proposition 5.1 enables us to implement the initialisation step on line 1 of Algorithm 1. The optimization $\arg\max_{s \in \mathcal{D}} g_{\hat{\mu}_{\tau^*}}(s)$ at line 21 of the algorithm may be performed by finding the roots of the derivative of the polynomial $g_{\hat{\mu}_{\tau^*}}$ and picking the maximizer of $g_{\hat{\mu}_{\tau^*}}$ among them by evaluation. Note that the set $\mathcal{D}_{\tau^*}$ at line 22 may be a disjoint union of multiple closed intervals. The endpoints of these intervals may be found by numerically computing roots of the polynomial $s \mapsto g_{\hat{\mu}_{\tau^*}}(s) - g_{\hat{\mu}_{\tau^*}}(\hat{s}) + \frac{\varepsilon}{2}$. A sequence of easy checks can then be used to pair the roots to yield the actual intervals that make up $\mathcal{D}_{\tau^*}$. Thus, VSBAI can be implemented rather easily for the case where the mean reward is a polynomial function of a single decision variable. The algorithm VSBAI-Poly in Appendix E provides an instantiation of VSBAI for the case of polynomial rewards and a single decision variable.

6 EXPERIMENTAL RESULTS

In this section we present experiments comparing VSBAI with other recent algorithms in various settings described below. We first consider the toy example considered in Fiez et al. [2019] and Jedra and Proutiere [2020a], and compare the sample complexities along with run-times in different scenarios. We also present in Appendix G some results for the polynomial setting described in Section 5. All the results that we present were computed on an AMD Ryzen 5 2500U CPU with Radeon Vega mobile gfx × 8 with 12GB memory.

6.1 MULTI-ARM SETTING

We consider the "finitely many arms with moderate gaps" example first presented in Fiez et al. [2019] and further used in Jedra and Proutiere [2020a]. The decision set is a finite collection of n 2-dimensional unit vectors given by $\mathcal{D} = \{[0, 1]^{\mathrm{T}}, [\cos(3\pi/4), \sin(\pi/4)]^{\mathrm{T}}\} \cup \{[\cos(\pi/4 + \phi_i), \sin(\pi/4 + \phi_i)]^{\mathrm{T}} : i = 3, \ldots, n\}$, where $n \geq 3$. Each choice of the angles $\{\phi_i\}_{i=3}^n$ represents a problem instance. In order to examine robustness across different problem instances, our experiments involve randomly sampling sets of these angles to generate different problem instances. The results we present below use $\mathcal{N}(0, .09)$ for generating the angles $\{\phi_i\}_{i=3}^n$. We also report the results from using the uniform distribution on the interval $[0, 0.1]$ in Appendix G. Typical arm configurations obtained by sampling the angles are depicted in figures 3 and 4 in Appendix F.

The reward is given by (1) with $\mu = [1, 0]^{\mathrm{T}}$. Also, Assumption 2 holds with $\sigma = 1$. To implement VSBAI on a problem instance, we first find the index j of the arm which has the least inner product with the arm $[\cos(3\pi/4), \sin(\pi/4)]^{\mathrm{T}}$. We then find a value of L such that the arms $[\cos(3\pi/4), \sin(\pi/4)]^{\mathrm{T}}$ and j form a $(L, 2)$-volumetric spanner for the decision set. Furthermore, VSBAI was run with $\varepsilon = 0.1$ and the Gaussian tail bound given in (8).

For drawing a comparison, we consider the LAZYTS (averaged) algorithm proposed in algorithm 1 in Jedra and Proutiere [2020a], the RAGE algorithm given as algorithm 1 in Fiez et al. [2019], and the ORACLE algorithm given by equations (4) and (5) of Soare et al. [2014]. For each choice of the size of the decision set, we generate 20 instances of the problem by sampling as many sets of the angles using either the normal distribution or the uniform distribution as described above. In addition to comparing sample complexities, we also compare run times as a measure of efficiency. The results that we present below for sample complexity and run time were obtained by averaging these quantities over all 20 problem instances for each algorithm.

Table 1 gives the sample complexities of the three baselines along with the VSBAI algorithm as the number of arms increase. We observe that, for all the baselines, the sample complexity grows with the number of arms, but the sample complexity of VSBAI remains almost constant. This is not surprising. While the other algorithms need to know the number of arms, VSBAI is independent of the number of arms. We also note the standard deviation (over the randomly generated problem instances) of the sample complexity for VSBAI decreases as the number of arms increase. In contrast, it increases for the baselines. This can be explained by observing that the value of L used by VSBAI can be expected to be closer to 1 as the number of arms increase.

In Table 2, we present the run-times of the algorithms given in Table 1 averaged over 20 problem instances as above. We note that VSBAI takes roughly a constant time to terminate whereas the run time of all the other algorithms increases as the number of arms increase. This is because all the three baselines attempt to find the best arm among all the arms. As a consequence, they can end up sampling the best two arms a large number of times in a scenario where the best two arms are very close to each other. VSBAI does

Algorithm	LazyTS		Rage		Oracle		VSBAI	
Arms	Mean	Std	Mean	Std	Mean	Std	Mean	Std
10	3490.05	1121.99	7617.4	2989.33	3470.05	1102.36	24460.4	243.93
20	72081.1	65078.96	103903.1	85734.65	47876.4	41692.63	24038.4	113.51
100	146331.55	64260.81	623143.05	366464.09	217162.25	111605.07	23691.1	22.22
1000	1218591.27	39881.14	16235680.31	5974249.14	7500331.73	2882866.47	23620.4	1.93

Table 1: Expected sample complexity for the setting described in 6.1

Algorithm	LazyTS		Rage		Oracle		VSBAI	
Arms	Mean	Std	Mean	Std	Mean	Std	Mean	Std
10	1.75	0.48	0.27	0.05	0.01	0	0.79	0.03
20	26.79	23.58	0.81	0.2	0.2	0.18	0.85	0.03
100	63.38	27.19	2.34	0.3	2.02	0.86	0.8	0.03
1000	39141.2	1270.31	120.92	6.01	116.56	38.37	0.76	0.03

Table 2: Run-time in sec for the setting described in 6.1

not suffer from this drawback as it seeks to find the best arm only to a certain degree of approximation, and this is a task that does not increase in difficulty with the number of arms. Also, in situations where run time is of importance, VSBAI makes it possible to use ε as an additional tuning parameter to balance accuracy and speed.

6.2 POLYNOMIAL SETTING

Next, we present results for case of polynomial rewards considered in section 5. As described in that section, the algorithm template VSBAI specializes to VSBAI-Poly, which is given as Algorithm 2 in Appendix E.

To implement VSBAI-Poly, we compute the a $(1, f)$-volumetric points this problem using (12) of Proposition 5.1 and the numerical technique suggested in Amballa et al. [2021]. We ran the VSBAI-Poly for various degrees ranging from 3 to 10. Although the noise sequence used was Gaussian (that is, satisfying Assumption 2) with $\sigma = 10$, VSBAI-Poly was run using both the Gaussian tail bound (8) as well as the sub-Gaussian tail bound (7). The tolerance ε was fixed to be 6.

Figure 2 represents the run-time of VSBAI-Poly using the Gaussian and sub-Gaussian tail bounds as the degree of the polynomial increases . The plots show the run time averaged over 20 polynomials having their maximum values around 350, but otherwise chosen randomly. As expected, the run time increases with the degree, and is higher when the sub-Gaussian tail bound is used in the algorithm.

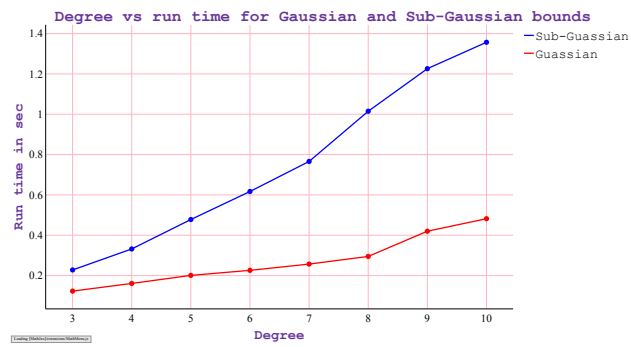

Figure 2: Arms sampled with Gaussian noise

7 CONCLUSION

We have considered a bandit problem in which the mean reward is a linearly parametrized (but possibly nonlinear) function on a continuous decision set. We have used a (ε, δ)-PAC formulation in which the goal is to find a set of points that are ε-optimal with probability at least $1 - \delta$. We have given a lower bound on the sample complexity of (ε, δ)-PAC algorithms. We have used the notion of volumetric spanners to devise a simple (ε, δ)-PAC algorithm template and provided an upper bound on its sample complexity. As a special case of our general setting, we have also considered the case where the mean reward is a polynomial function of a single decision variable, and indicated how all the problem-specific steps in VSBAI can be instantiated to apply to this case. VSBAI showed advantages in experiments in terms of run time and sampling complexity when compared to recent algorithms proposed for the BAI problem in linear bandits with finite arms.

References

Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.

Chaitanya Amballa, Manu K. Gupta, and Sanjay P. Bhat. Computing an efficient exploration basis for learning with univariate polynomial features. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(8):6636–6643, May 2021.

Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *COLT*, pages 41–53. Citeseer, 2010.

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.

Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Proceedings of the 21st Annual Conference on Learning Theory-COLT 2008*, pages 335–342, 2008.

Dennis S. Bernstein. *Scalar, Vector, and Matrix Mathematics: Theory, Facts, and Formulas*. Princeton University Press, revised edition, 2018. ISBN 9780691151205.

Dimitri P Bertsekas and Steven E Shreve. *Stochastic optimal control: the discrete-time case*, volume 5. Athena Scientific, 1996. ISBN 9781886529038.

Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, pages 23–37. Springer, 2009.

Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012. ISSN 1935-8237. doi: 10.1561/2200000024.

Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory-COLT 2008*, pages 355–366, 2008.

Rémy Degenne, Pierre Ménard, Xuedong Shang, and Michal Valko. Gamification of pure exploration for linear bandits. In *International Conference on Machine Learning*, pages 2432–2442. PMLR, 2020.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270. Springer, 2002.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for reinforcement learning. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 162–169, 2003.

Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. *Advances in neural information processing systems*, 32, 2019.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pages 998–1027. PMLR, 2016.

Elad Hazan and Zohar Karnin. Volumetric spanners: an efficient exploration basis for learning. *The Journal of Machine Learning Research*, 17(1):4062–4095, 2016.

Yassir Jedra and Alexandre Proutiere. Optimal best-arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 33:10007–10017, 2020a.

Yassir Jedra and Alexandre Proutiere. Optimal best-arm identification in linear bandits. *arXiv preprint arXiv:2006.16073*, 2020b.

Zohar S Karnin. Verification based solution for structured mab problems. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *J. Mach. Learn. Res.*, 17(1):1–42, jan 2016. ISSN 1532-4435.

Jack Kiefer and Jacob Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960.

Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020. doi: 10.1017/9781108571401.

Jacques Neveu. *Mathematical foundations of the calculus of probability*. Holden-day, 1965.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 27:828–836, 2014.

Chao Tao, Saúl Blanco, and Yuan Zhou. Best arm identification in linear bandits with linear dimension dependency. In *International Conference on Machine Learning*, pages 4877–4886. PMLR, 2018.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019. doi: 10.1017/9781108627771.

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, 09–11 Apr 2018.

Junwen Yang and Vincent YF Tan. Towards minimax optimal best arm identification in linear bandits. *arXiv preprint arXiv:2105.13017*, 2021.

A PROOF OF THEOREM 3.2

We begin by recalling properties of the measures $\mathbb{E}^{\mathcal{A},\mu}$ and $\mathbb{E}_t^{\mathcal{A},\mu}$ introduced in Section 3. By Proposition 7.28 in Bertsekas and Shreve [1996], the measures $\mathbb{P}_t^{\mathcal{A},\mu}$, $t \in \mathbb{Z}_+$, and $\mathbb{P}^{\mathcal{A},\mu}$ satisfy the following properties (see also Proposition V.1.1 of Neveu [1965]):

1. For every real-valued function q that is integrable on $(\Omega_1, \mathcal{F}_1, \mathbb{P}_1^{\mathcal{A},\mu})$, we have
$$\int_{\Omega_1} q(h_1)\mathbb{P}_1^{\mathcal{A},\mu}(\mathrm{d}h_1) = \int_{[0,1]^n}\int_{\mathcal{D}}\int_{\mathbb{R}} q(u,s,y) \times Q^{\mu}(\mathrm{d}y|s)\pi_1(\mathrm{d}s|u)\lambda(\mathrm{d}u) \tag{13}$$

2. For every $t > 1$ and every integrable real-valued function q on $(\Omega_t, \mathcal{F}_t, \mathbb{P}_t^{\mathcal{A},\mu})$, we have
$$\int_{\Omega_t} q(h_t)\mathbb{P}_t^{\mathcal{A},\mu}(\mathrm{d}h_t) = \int_{\Omega_{t-1}}\int_{[0,1]^n}\int_{\mathcal{D}}\int_{\mathbb{R}} q(h_{t-1},u,s,y) \times Q^{\mu}(\mathrm{d}y|s)\pi_1(\mathrm{d}s|u)\lambda(\mathrm{d}u)\mathbb{P}_{t-1}^{\mathcal{A},\mu}(\mathrm{d}h_{t-1}). \tag{14}$$

3. For every $t \in \mathbb{Z}_+$ and every Borel subset A of Ω_t, we have $\mathbb{P}_t^{\mathcal{A},\mu}(A) = \mathbb{P}^{\mathcal{A},\mu}(A \times \mathcal{S} \times \mathcal{S} \times \cdots)$.

Next, suppose Assumption 2 holds, and let $\rho_{z,\nu}(\cdot)$ denote the Gaussian density on $\mathbb{R}$ with mean $z \in \mathbb{R}$ and standard deviation $\nu > 0$. Then, for every $\zeta \in \mathbb{R}^f$ and $s \in \mathcal{D}$, the measure $Q^{\zeta}(\cdot|s)$ is a Gaussian measure on $\mathbb{R}$ having density $\rho_{g_\zeta(s),\sigma}(\cdot)$ w.r.t. the Lebesgue measure on $\mathbb{R}$. Consequently, for every $\zeta \in \mathbb{R}^f$ and $s \in \mathcal{D}$, the measures $Q^{\zeta}(\cdot|s)$ and $Q^{\mu}(\cdot|s)$ are mutually absolutely continuous, and
$$\left.\frac{\mathrm{d}Q^{\mu}(\cdot|s)}{\mathrm{d}Q^{\zeta}(\cdot|s)}\right|_y = \frac{\rho_{g_\mu(s),\sigma}(y)}{\rho_{g_\zeta(s),\sigma}(y)}. \tag{15}$$

We are now ready to begin the proof of Theorem 3.2.

Proof of Theorem 3.2 Consider an alternative reward model given by $\zeta \in \mathrm{Alt}_\varepsilon(\mu)$. For each $t \in \mathbb{Z}_+$, define the log-likelihood ratio $\mathcal{L}_t^{\mu,\zeta} \stackrel{\mathrm{def}}{=} \ln \frac{\mathrm{d}\mathbb{P}_t^{\mathcal{A},\mu}}{\mathrm{d}\mathbb{P}_t^{\mathcal{A},\zeta}}$, and note that $\mathcal{L}_t^{\mu,\zeta}$ is a random variable on $(\Omega_t, \mathcal{F}_t, \mathbb{P}_t^{\mathcal{A},\mu})$. It is now easy to see from (13), (14) and (15) that, for each $t \in \mathbb{Z}_+$ and each $h_t = \{(s_i, y_i, u_i)\}_{i=1}^t \in \Omega_t$, we have
$$\mathcal{L}_t^{\mu,\zeta}(h_t) = \sum_{i=1}^t [\ln \rho_{g_\mu(s_i),\sigma}(y_i) - \ln \rho_{g_\zeta(s_i),\sigma}(y_i)]. \tag{16}$$

Next, define the event $\mathcal{E}^{\mu} \stackrel{\mathrm{def}}{=} \{\arg\max_{s\in\mathcal{D}} g_\mu(s) \subseteq \mathbb{F}(h_\tau) \subseteq \mathcal{O}_\varepsilon(\mu)\}$, and note that $\mathcal{E}^{\mu}$ is contained in the σ-algebra $\mathcal{F}_\tau$ generated by the stopping time τ. It follows from Lemma 19 of Kaufmann et al. [2016] that
$$\mathbb{E}^{\mathcal{A},\mu}(\mathcal{L}_\tau^{\mu,\zeta}) \geq \mathrm{kl}(\mathbb{P}^{\mathcal{A},\mu}(\mathcal{E}^{\mu}), \mathbb{P}^{\mathcal{A},\zeta}(\mathcal{E}^{\mu})), \tag{17}$$
where $\mathrm{kl}(\nu_1, \nu_2)$ is the KL-divergence between two Bernoulli distributions having parameters $\nu_1, \nu_2 \in [0,1]$.

Next, define the event $\mathcal{E}^{\zeta}$ by replacing μ in the definition of $\mathcal{E}^{\mu}$ with ζ. Since $\mathcal{A}$ is a (ε,δ)-PAC algorithm, we have $\mathbb{P}^{\mathcal{A},\mu}(\mathcal{E}^{\mu}) \geq 1-\delta$ and $\mathbb{P}^{\mathcal{A},\zeta}(\mathcal{E}^{\zeta}) \geq 1-\delta$. By our choice of ζ, we have $\mathcal{E}^{\mu} \cap \mathcal{E}^{\zeta} \subseteq \{\mathbb{F}(h_\tau) \subseteq \mathcal{O}_\varepsilon(\mu) \cap \mathcal{O}_\varepsilon(\zeta)\} = \varnothing$. As a result, we infer that $\mathbb{P}^{\mathcal{A},\zeta}(\mathcal{E}^{\mu}) < \delta$. Monotonicity properties of the KL divergence now imply that $\mathrm{kl}(\mathbb{P}^{\mathcal{A},\mu}(\mathcal{E}^{\mu}), \mathbb{P}^{\mathcal{A},\zeta}(\mathcal{E}^{\mu})) \geq \mathrm{kl}(\delta, 1-\delta)$ By inequality (3) in Kaufmann et al. [2016], we further have $\mathrm{kl}(\delta, 1-\delta) \geq \ln(1/2.4\delta)$. Using this in (17), we get
$$\mathbb{E}^{\mathcal{A},\mu}(\mathcal{L}_\tau^{\mu,\zeta}) \geq \ln\left(\frac{1}{2.4\delta}\right). \tag{18}$$

Combining (18) with (19) from Lemma A.1 below yields
$$\frac{1}{2\sigma^2}\mathbb{E}^{\mathcal{A},\mu}(\tau)\|g_\mu - g_\zeta\|_\infty^2 \geq \ln\left(\frac{1}{2.4\delta}\right).$$

Inequality (2) now follows by taking an infimum over $\zeta \in \mathrm{Alt}_\varepsilon(\mu)$ on the left hand side in the inequality above and rearranging the resulting inequality. □

Lemma A.1. *Let algorithm $\mathcal{A}$ and $\mu, \zeta \in \mathbb{R}^f$ be as in the proof of Theorem (4.2). Suppose $\{\mathcal{L}_t^{\mu,\zeta}\}_{t=1}^\infty$ is defined as in (16), and let τ be a stopping time with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. Then we have*
$$\mathbb{E}^{\mathcal{A},\mu}(\mathcal{L}_\tau^{\mu,\zeta}) \leq \frac{1}{2\sigma^2}\mathbb{E}^{\mathcal{A},\mu}(\tau)\|g_\mu - g_\zeta\|_\infty^2. \tag{19}$$

Proof. For each $t \in \mathbb{Z}_+$, denote $\ell_t = \ln \rho_{g_\mu(s_t),\sigma}(y_t) - \ln \rho_{g_\zeta(s_t),\sigma}(y_t)$, and let $\mathcal{G}_t$ denote the σ-algebra on Ω generated by (h_{t-1}, u_t, s_t). Note that, for each $t \in \mathbb{Z}_+$, ℓ_t is a $\mathcal{F}_t$-measurable random variable, while $\mathcal{F}_{t-1} \subseteq \mathcal{G}_t \subseteq \mathcal{F}_t$.

Next, define the process $\{M_t\}_{t=0}^\infty$ by $M_0 = 0$ and $M_t = \sum_{i=1}^t [\ell_i - \mathbb{E}^{\mathcal{A},\mu}(\ell_i|\mathcal{G}_i)]$ for each $t \in \mathbb{Z}_+$. The inclusions $\mathcal{F}_{t-1} \subseteq \mathcal{G}_t \subseteq \mathcal{F}_t$ along with the tower property of conditional expectations show that the process $\{M_t\}_{t=0}^\infty$ is adapted to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$ and is a martingale under the measure $\mathbb{P}^{\mathcal{A},\mu}$. The optional stopping theorem now implies that $\mathbb{E}^{\mathcal{A},\mu}(M_\tau) = \mathbb{E}^{\mathcal{A},\mu}(M_0) = 0$. This immediately yields
$$\mathbb{E}^{\mathcal{A},\mu}\left[\sum_{i=1}^\tau \ell_i\right] = \mathbb{E}^{\mathcal{A},\mu}\left[\sum_{i=1}^\tau \mathbb{E}^{\mathcal{A},\mu}(\ell_i|\mathcal{G}_i)\right]. \tag{20}$$

Substituting the expression for a Gaussian density in the expression for ℓ_i yields $2\sigma^2\ell_i = 2y_i[g_\mu(s_i) - g_\zeta(s_i)] +$

$[g_\zeta(s_i)]^2 - [g_\mu(s_i)]^2$ for each $i \in \mathbb{Z}_+$. Using the fact that $\mathbb{E}^{\mathcal{A},\mu}(y_i) = g_\mu(s_i)$ gives $2\sigma^2\mathbb{E}^{\mathcal{A},\mu}(\ell_i|\mathcal{G}_i) = [g_\zeta(s_i) - g_\mu(s_i)]^2 \leq \max_{s\in\mathcal{D}}[g_\mu(s) - g_\zeta(s)]^2$ for each $i \in \mathbb{Z}_+$. Using the last inequality in (20) and recognizing the left hand side of (20) to be $\mathbb{E}^{\mathcal{A},\mu}(\mathcal{L}_\tau^{\mu,\zeta})$ yields (19). □

B PROOFS FOR SUBSECTIONS 4.1 AND 4.3

Proof of Proposition 4.1. Choose $s^* \in \arg\max_{s\in\mathcal{D}} q(s)$, and consider $s \in \mathcal{D}'$. We have $q(s^*) \leq \hat{q}(s^*) + \frac{\varepsilon}{4} \leq \hat{q}(\hat{s}) + \frac{\varepsilon}{4} \leq \hat{q}(s) + \frac{3\varepsilon}{4} \leq q(s) + \varepsilon$, where the first and last inequalities follow from $\|\hat{q} - q\|_\infty \leq \frac{\varepsilon}{4}$, the second inequality follows from the definition of $\hat{s}$, and the third follows from the definition of $\mathcal{D}'$ and our choice $s \in \mathcal{D}'$. We have thus shown that every $s \in \mathcal{D}'$ is ε-optimal for q.

Next, we have $\hat{q}(s^*) \geq q(s^*) - \frac{\varepsilon}{4} \geq q(\hat{s}) - \frac{\varepsilon}{4} \geq \hat{q}(\hat{s}) - \frac{\varepsilon}{2}$, where the first and the last inequalities follow from $\|\hat{q} - q\|_\infty \leq \frac{\varepsilon}{4}$ while the second inequality follows from the fact that s^* is a maximizer of q. We have thus shown that $s^* \in \mathcal{D}'$. Since $s^* \in \arg\max_{s\in\mathcal{D}} q(s)$ was chosen arbitrarily, the last assertion of the result follows. □

Proof of Lemma 4.2. First, we recall a definition from the theory of optimal designs. A *design* is a probability measure on the Borel σ-algebra of $\mathcal{D}$. Given a design ξ on $\mathcal{D}$, we denote $V_\xi = \int_\mathcal{D} \phi(s)\phi^{\mathrm{T}}(s)\xi(\mathrm{d}s)$. Note that the integral is defined since $\mathcal{D}$ is compact and ϕ is continuous.

Next, suppose $C = \{\phi(p_1), \ldots, \phi(p_m)\}$ is a (L, m)-volumetric spanner for some $L \geq 1$ and $m \geq f$. Consider the design ξ which places mass $1/m$ at each of the points of C, and let $X = [\phi(p_1), \ldots, \phi(p_m)] \in \mathbb{R}^{f\times m}$. Note that $XX^{\mathrm{T}} = mV_\xi$.

By the definition of a (L, m)-volumetric spanner, we have $\max_{z\in\phi(\mathcal{D})}\|X^{\mathrm{T}}(XX^{\mathrm{T}})^{-1}z\|_2^2 \leq L^2$. A simple calculation shows that, for each $z \in \phi(\mathcal{D})$, we have $\|X^{\mathrm{T}}(XX^{\mathrm{T}})^{-1}z\|_2^2 = z^{\mathrm{T}}(XX^{\mathrm{T}})^{-1}z = m^{-1}z^{\mathrm{T}}V_\xi z$. The Keifer-Wolfowitz theorem [Kiefer and Wolfowitz, 1960], [Lattimore and Szepesvári, 2020, Thm. 21.1] implies that $\max_{z\in\phi(\mathcal{D})} x^{\mathrm{T}}V_\xi x \geq f$. Putting everything together, we have $L^2 \geq m^{-1}\max_{z\in\phi(\mathcal{D})} z^{\mathrm{T}}V_\xi z \geq f/m$. This completes the proof. □

C PROOF OF PROPOSITION 4.3

By way of preparation for the proof of Proposition 4.3, we will find it convenient to rewrite (3) and (4) by grouping together observations made during each round. To this end, let $B_{L,m} = [\phi(p_1), \ldots, \phi(p_m)] \in \mathbb{R}^{f\times m}$ and, for each $j \in \mathbb{Z}_+$, let $\bar{y}^j = [y_{(j-1)m+1}, \ldots, y_{jm}]^{\mathrm{T}} \in \mathbb{R}^m$ and $\bar{\eta}^j = [\eta_{(j-1)m+1}, \ldots, \eta_{jm}]^{\mathrm{T}} \in \mathbb{R}^m$ denote the vectors of rewards and noise samples, respectively, encountered in the jth round. The decision epoch at the end of $k > 0$ rounds is $t = km$. In the notation of (3), we have

$$X_t = [\underbrace{B_{L,m}|\cdots|B_{L,m}}_{k \text{ times}}].$$

Equations (3)-(4) now become

$$\hat{\mu}_{km} = (B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}B_{L,m}\left[\frac{1}{k}\sum_{j=1}^{k}\bar{y}^j\right], \quad (21)$$

$$\hat{\mu}_{km} - \mu = (B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}B_{L,m}\left[\frac{1}{k}\sum_{j=1}^{k}\bar{\eta}^j\right]. \quad (22)$$

The proof of the sub-Gaussian part of Proposition 4.3 essentially applies to the right hand side of (22) the tail concentration inequality below for the norm of the average of k random vectors having independent σ-sub-Gaussian components. The proof is given later in this appendix.

Proposition C.1. *Suppose $\xi^1, \ldots, \xi^k$ are f-dimensional random vectors such that the random variables $\{\xi_i^j : i = 1, \ldots, f, j = 1, \ldots, k\}$ are independent and σ-sub Gaussian. Let $S_k = \left(\xi^1 + \cdots + \xi^k\right)$. Then the following statements hold.*

1. *$\exp\left(\lambda\|S_k\|_2^2\right)$ is integrable for each $\lambda \in \left(0, 1/2\sigma^2k\right)$.*
2. *For every $\varepsilon > 0$, we have $\mathbb{P}\left(\frac{1}{k}\|S_k\|_2 > \varepsilon\right) \leq \beta(k, \varepsilon)$, where β is given by (7).*

We are now ready to prove Proposition 4.3.

Proof of Proposition 4.3. First, suppose Assumption 1 holds. We have

$$\begin{aligned}\|g_{\hat{\mu}_{km}} - g_\mu\|_\infty &= \max_{s\in\mathcal{D}}|\phi^{\mathrm{T}}(s)(\hat{\mu}_{km} - \mu)| \\ &= \max_{s\in\mathcal{D}}\left|\phi^{\mathrm{T}}(s)(B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}B_{L,m}\left[\frac{1}{k}S_k\right]\right|, \quad (23)\end{aligned}$$

where $S_k = \sum_{j=1}^k \bar{\eta}^j$, and the last equality uses (22). Since the columns of $B_{L.m}$ form a (L, m)-volumetric spanner for $\phi(\mathcal{D})$, it follows that $\|B_{L,m}^{\mathrm{T}}(B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}\phi(s)\|_2 \leq L$ for all $s \in \mathcal{D}$. Using this fact along with the Cauchy-Schwarz inequality in (23) gives $\|g_{\hat{\mu}_{km}} - g_\mu\|_\infty \leq \frac{L}{k}\|S_k\|_2$. The first part of the proposition now follows immediately from Proposition C.1.

Next, we prove the second part of the proposition under Assumption 2. For convenience, define $\varphi(s) = B_{L,m}^{\mathrm{T}}(B_{L,m}B_{L,m}^{\mathrm{T}})^{-1}\phi(s)$ for every $s \in \mathcal{D}$, and recall that $\|\varphi(s)\|_2 \leq L$ for all $s \in \mathcal{D}$. Define $G : \mathbb{R}^f \to \mathbb{R}$ by $G(z) = \max_{s\in\mathcal{D}}|\varphi^{\mathrm{T}}(s)z|$, for every $z \in \mathbb{R}^f$.

We claim that G is Lipschitz with respect to the Euclidean norm with Lipschitz constant 1. To prove the claim, consider $z_1, z_2 \in \mathbb{R}^f$. We assume without loss of generality

that $G(z_1) \geq G(z_2)$, and let $s' \in \arg\max_{s\in\mathcal{D}}|\varphi^{\mathrm{T}}(s)z_1|$. Then we have $|G(z_1) - G(z_2)| = G(z_1) - G(z_2) = \max_{s\in\mathcal{D}}|\varphi^{\mathrm{T}}(s)z_1| - \max_{s\in\mathcal{D}}|\varphi^{\mathrm{T}}(s)z_2| \leq |\varphi^{\mathrm{T}}(s')z_1| - |\varphi^{\mathrm{T}}(s')z_2| \leq |(\varphi^{\mathrm{T}}(s')(z_1 - z_2)| \leq \|\varphi(s')\|_2\|z_1 - z_2\|_2 \leq \|z_1 - z_2\|_2$. This proves the claim.

Next, note that under Assumption 2 invoked in the second part of the proposition, $Z \stackrel{\text{def}}{=} \frac{1}{k}S_k$ is an average of k independent zero-mean f-dimensional Gaussian random vectors, each having covariance matrix $\sigma^2 I$, where I is the $f\times f$ identity matrix. Hence Z is a zero-mean Gaussian random vector with covariance matrix $\frac{\sigma^2}{k}I$. We now observe from (23) that $\|g_{\hat{\mu}_{km}} - g_\mu\|_\infty = G(Z)$. Thus $\|g_{\hat{\mu}_{km}} - g_\mu\|_\infty$ is a Lipschitz function of the Gaussian vector Z having i.i.d. components. Applying a slightly extended version of Theorem 2.26 of Wainwright [2019] (which is stated for the special case where the components of Z are standard normal) gives (8). □

For proving Proposition C.1, we begin by recollecting a few preliminary results in preparation for the proof of Proposition C.1. Though these results are known, we state them to make the constants explicit, and provide proofs for easy reference.

Lemma C.2. *Suppose X is σ-sub Gaussian for some $\sigma > 0$. If $\lambda \in (0, 1/2\sigma^2)$, then $\exp\left(\lambda X^2\right)$ is integrable, and $\mathbb{E}\left[\exp\left(\lambda X^2\right)\right] \leq 2^{2\sigma^2\lambda}(1-2\sigma^2\lambda)^{-1}$.*

Proof. Let $\lambda \in (0, 1/2\sigma^2)$. Since X is σ-sub Gaussian, we have $\mathbb{P}\left(|X| > t\right) \leq 2\exp\left(-\frac{t^2}{2\sigma^2}\right)$ for all $t > 0$ Next, note that $\exp\left(\lambda X^2\right) \geq 1$. Let $s \geq 1$. Then

$$\begin{aligned}\mathbb{P}\left(\exp\left(\lambda X^2\right) > s\right) &= \mathbb{P}\left(|X| > \sqrt{\frac{\ln s}{\lambda}}\right)\\ &\leq 2\exp\left(-\frac{1}{2\sigma^2}\frac{\ln s}{\lambda}\right) = 2s^{-\frac{1}{2\sigma^2\lambda}}.\end{aligned}$$

Thus, we conclude that

$$\mathbb{P}\left(\exp\left(\lambda X^2\right) > s\right) \leq \begin{cases}1, & \text{if } s \leq 2^{2\sigma^2\lambda},\\ 2s^{-\frac{1}{2\sigma^2\lambda}}, & \text{if } s > 2^{2\sigma^2\lambda}.\end{cases} \quad (24)$$

Since $2\sigma^2\lambda < 1$, the integral $\int_0^\infty \mathbb{P}\left(\exp\left(\lambda x^2\right) > s\right)\mathrm{d}s$ exists. Indeed, (24) implies that

$$\begin{aligned}&\int_0^\infty \mathbb{P}\left(\exp\left(\lambda x^2\right) > s\right)\mathrm{d}s\\ &\leq \int_0^{2^{2\sigma^2\lambda}} 1\mathrm{d}s + \int_{2^{2\sigma^2\lambda}}^\infty 2s^{-\frac{1}{2\sigma^2\lambda}}\mathrm{d}s\\ &= 2^{2\sigma^2\lambda}(1-2\sigma^2\lambda)^{-1}.\end{aligned}$$

Since $\exp\left(\lambda x^2\right)$ is a non-negative random variable, it follows that $\mathbb{E}\left[\exp\left(\lambda x^2\right)\right] = \int_0^\infty \mathbb{P}\left(\exp\left(\lambda x^2\right) > s\right)\mathrm{d}s$, and the result follows. □

Lemma C.3. *Suppose ζ is a random vector of dimension f such that $\zeta_1, \ldots, \zeta_f$ are independent σ-sub Gaussian random variables. Then $\mathbb{E}\left[\exp\left(\lambda x^{\mathrm{T}}\zeta\right)\right] \leq \exp\left(\frac{\lambda^2\|x\|_2^2\sigma^2}{2}\right)$ for all $x \in \mathbb{R}^f$ and $\lambda \in \mathbb{R}$. Furthermore, if $\lambda \in (0, 1/2\sigma^2)$, then $\exp\left(\lambda\|\zeta\|_2^2\right)$ is integrable, and $\mathbb{E}\left[\exp\left(\lambda\|\zeta\|_2^2\right)\right] \leq \left(\frac{2^{2\sigma^2\lambda}}{1-2\sigma^2\lambda}\right)^f$.*

Proof. By independence and σ-sub Gaussianity, we have

$$\begin{aligned}\mathbb{E}\left[\exp\left(\lambda x^{\mathrm{T}}\zeta\right)\right] &= \prod_{i=1}^f \mathbb{E}\left[\exp\left(\lambda x_i\zeta_i\right)\right]\\ &\leq \prod_{i=1}^f \exp\left(\frac{\lambda^2 x_i^2\sigma^2}{2}\right) = \exp\left(\frac{\lambda^2\|x\|_2^2\sigma^2}{2}\right).\end{aligned}$$

This proves the first assertion. To prove the second assertion, let $\lambda \in \left(0, 1/2\sigma^2\right)$. By Lemma C.2, $\exp\left(\lambda\zeta_i^2\right)$ is integrable for each i. Hence it follows by independence that $\exp\left(\lambda\|\zeta\|_2^2\right)$ is also integrable, and $\mathbb{E}\left[\exp\left(\lambda\|\zeta\|_2^2\right)\right] = \prod_{i=1}^f \mathbb{E}\left[\exp\left(\lambda\zeta_i^2\right)\right] \leq \left(\frac{2^{2\sigma^2\lambda}}{1-2\sigma^2\lambda}\right)^f$. □

The next lemma, which we state without proof, is a conditional version of the first part of Lemma C.3.

Lemma C.4. *Suppose ζ is a random vector of dimension f such that $\zeta_1, \ldots, \zeta_f$ are independent, σ-sub Gaussian random variables. Let Y be a $\mathcal{G}$-measurable f-dimensional random vector, where $\mathcal{G}$ is a σ-algebra such that ζ is independent of $\mathcal{G}$. Then*

$$\mathbb{E}\left[\exp\left(\lambda Y^{\mathrm{T}}\zeta\right)|\mathcal{G}\right] \leq \exp\left(\frac{\lambda^2\|Y\|_2^2\sigma^2}{2}\right) \text{ a.s.}$$

The proof of Proposition C.1 follows next.

Proof of Proposition C.1. The ith component of S_k is a sum of k independent σ-sub Gaussian random variables. Applying the first part of Lemma C.3 with $\zeta = \left[\xi_i^1, \ldots, \xi_i^k\right]$ and $x = [1, \ldots, 1]$ lets us conclude that the ith element of S_k is $\sigma\sqrt{k}$-sub Gaussian. Applying the second part of Lemma C.3 with $\zeta = S_k$ shows that $\exp\left(\lambda\|S_k\|_2^2\right)$ is integrable for $\lambda \in \left(0, 1/2\sigma^2 k\right)$. This proves the first assertion.

To prove the second statement, choose $x \in \mathbb{R}^f$, and define the process $\{M_j(x)\}_{j=0}^k$ by $M_0(x) = 1$ and

$$M_j(x) = \exp\left(x^{\mathrm{T}}S_j - j\frac{\sigma^2}{2}\|x\|_2^2\right), j = 1, \ldots, k,$$

where $S_j = \xi^1 + \cdots + \xi^j$ for each j. It follows from the first part of Lemma C.3 that $M_j(x)$ is integrable for each j. Next, let $\mathcal{G}_j$ denote the σ-algebra generated by $\xi^1, \ldots, \xi^j$, with $\mathcal{G}_0$ denoting the trivial σ-algebra, and note that $M_j(x)$

is $\mathcal{F}_j$-measurable. For each $j = 1, \ldots, k$, we have

$$\begin{aligned}
&\mathbb{E}\left[M_j(x)|\mathcal{G}_{j-1}\right] \\
&= \mathbb{E}\left[M_{j-1}(x)\exp\left(x^{\mathrm{T}}\xi^j - \frac{\sigma^2}{2}\|x\|_2^2\right)\middle|\mathcal{G}_{j-1}\right] \\
&= M_{j-1}(x)\mathbb{E}\left[\exp\left(x^{\mathrm{T}}\xi^j - \frac{\sigma^2}{2}\|x\|_2^2\right)\right] \\
&\leq M_{j-1}(x),
\end{aligned}$$

where the second equality follows from the $\mathcal{G}_{j-1}$-measurability of $M_{j-1}(x)$ and the $\mathcal{G}_{j-1}$-independence of ξ^j (see Lemma C.4), while the last inequality follows by applying the first part of Lemma C.3 with $\zeta = \xi^j$. Thus, $\{M_j(x)\}_{j=0}^k$ is a supermartingale with respect to the filtration $\{\mathcal{G}_j\}_{j=0}^k$.

Next, define $\{\bar{M}_j\}_{j=0}^k$ by

$$\bar{M}_j = \left(\frac{k\sigma^2}{2\pi}\right)^{f/2}\int_{\mathbb{R}^f} M_j(x)\exp\left(\frac{-k\sigma^2}{2}x^{\mathrm{T}}x\right)\mathrm{d}x, \tag{25}$$

and note that $\bar{M}_0 = 1$. Substituting for $M_j(x)$ in (25), completing the square in the exponent, rearranging terms yields

$$\bar{M}_j = \left[\left(\frac{k}{(j+k)}\right)^{\frac{f}{2}}\exp\left(\frac{\|S_j\|_2^2}{2(j+k)\sigma^2}\right)\right]\times J, \tag{26}$$

where J is the integral over x of the F-dimensional Gaussian density over x with mean $[(j+k)\sigma^2]^{-1}S_j$ and covariance matrix $[(j+k)\sigma^2]^{-1}I$, with I denoting the $f \times f$ identity matrix. Thus, J evaluates to 1. Next, S_j is a random vector with independent $\sigma\sqrt{j}$-sub Gaussian components. Also, $\frac{1}{2(j+k)\sigma^2} < \frac{1}{2j\sigma^2}$. Hence, by Lemma C.3, $\bar{M}_j$ is integrable. In addition, it follows from Lemma 20.3 in Lattimore and Szepesvári [2020] that $\{\bar{M}_j\}_{j=1}^k$ is a submartingale.

Letting $j = k$ in (26) gives $\bar{M}_k = 2^{-\frac{f}{2}}\exp(\frac{\|S_k\|_2^2}{4k\sigma^2}$. By Ville's maximal inequality (see Theorem 3.9 in Lattimore and Szepesvári [2020]), we have

$$\begin{aligned}
\mathbb{P}\left(\|S_k\|_2 > \epsilon\right) &= \mathbb{P}\left(\bar{M}_k > \frac{1}{2^{f/2}}\exp\left(\frac{\epsilon^2}{4k\sigma^2}\right)\right) \\
&\leq \mathbb{P}\left(\max_j \bar{M}_j > \frac{1}{2^{f/2}}\exp\left(\frac{\epsilon^2}{4k\sigma^2}\right)\right) \\
&\leq \frac{\mathbb{E}\left[\bar{M}_0\right]}{\frac{1}{2^{f/2}}\exp\left(\frac{\epsilon^2}{4k\sigma^2}\right)} \\
&= 2^{f/2}\exp\left(\frac{-\epsilon^2}{4k\sigma^2}\right).
\end{aligned}$$

Replacing ε by $k\varepsilon$ in the last inequality completes the proof of the second assertion. □

D PROOF OF PROPOSITION 5.1

The proof of Proposition 5.1 uses the following lemma.

Lemma D.1. *Let $s \in [p_{\min}, p_{\max}]$ and suppose $p_1, \ldots, p_f \in [p_{\min}, p_{\max}]$ are such that $p_i \neq p_j$ for all $i \neq j$. Then $c_1, \ldots, c_{n+1} \in \mathbb{R}$ satisfy*

$$c_1\phi(p_1) + \cdots + c_{n+1}\phi(p_f) = \phi(s) \tag{27}$$

if and only if $c_i = l_i(s, \mathbf{p})$ for each $i = 1, \ldots, f$ where $\mathbf{p} = [p_1, \ldots, p_f]^{\mathrm{T}}$, and $l_i(\cdot, \mathbf{p})$ is the ith Lagrange basis polynomial for the points $\{p_1, p_2, \ldots, p_f\}$ given by

$$l_i(s, \mathbf{p}) \stackrel{\text{def}}{=} \frac{\prod_{j\neq i}(s - p_j)}{\prod_{j\neq i}(p_i - p_j)}. \tag{28}$$

Proof. Equation (27) may be rewritten as

$$V(\mathbf{p})c(s) = \phi(s), \tag{29}$$

where $V(\mathbf{p}) \stackrel{\text{def}}{=} [\phi(p_1), \ldots, \phi(p_f)] \in \mathbb{R}^{f\times f}$. Note that $V(\mathbf{p})$ is a Vandermonde matrix, and its determinant is given by (see Fact 7.18.5 from Bernstein [2018])

$$\det(V(\mathbf{p})) = \prod_{1\leq i<j\leq f}(p_j - p_i), \tag{30}$$

The determinant of $V(\mathbf{p})$ in (30) is nonzero since $p_i \neq p_j$ for $j \neq i$. Equation (29) thus has a unique solution. Applying Cramer's rule (see Fact 3.16.12 from Bernstein [2018]) gives this solution to be

$$c_i = \frac{\det(V(\mathbf{p}_i^s))}{\det(V(\mathbf{p}))} \tag{31}$$

where $\mathbf{p}_i^s$ is the vector obtained by replacing the ith element of $\mathbf{p}$ by s. Using (30) to expand the determinants of the two Vandermonde matrices in (31) and canceling common terms gives $c_i = l_i(s, \mathbf{p})$. □

The proof of Proposition 5.1 follows.

Proof of Proposition 5.1. To show 1) implies 2), suppose $p_1, \ldots, p_f \in \mathcal{D}$ are $(1, f)$-volumetric points for the pair $(\phi, \mathcal{D})$. Choose $s \in \mathcal{D} = [p_{\min}, p_{\max}]$ arbitrarily. Applying the definition of $(1, f)$-volumetric points, it follows that there exist $c_1, \ldots, c_f \in \mathbb{R}$ such that $c_1\phi(p_1) + \cdots + c_f\phi(p_f) = \phi(s)$ and $c_1^2 + \ldots + c_f^2 \leq 1$. Clearly, $|c_i| \leq 1$ for all $i = 1, \ldots, f$. Since $s \in \mathcal{D}$ was chosen arbitrarily, it follows that $\{\phi(p_1), \ldots, \phi(p_f)\}$ is a barycentric spanner for $\phi(\mathcal{D})$ (see Amballa et al. [2021] for a definition). Theorem 1 of Amballa et al. [2021] now implies that 2) holds.

To prove that 2) implies 1), suppose $p_{\min} = p_1 \leq p_2 \leq \cdots \leq p_f = p_{\max}$ satisfy (12). Define $\mathbf{p}$ as in Lemma D.1. The Lagrange polynomials defined in Lemma D.1 satisfy

$$\begin{aligned}
l_i(p_i, \mathbf{p}) &= 1,\ i = 1, \ldots, f, && (32)\\
l_i(p_j, \mathbf{p}) &= 0,\ i, j = i = 1, \ldots, f,\ i \neq j, && (33)\\
\frac{\mathrm{d}l_i}{\mathrm{d}s}(p_i, \mathbf{p}) &= 0,\ i = 2, \ldots, f-1, && (34)
\end{aligned}$$

$$\frac{\mathrm{d}l_1}{\mathrm{d}s}(p_1, \mathbf{p}) < 0 < \frac{\mathrm{d}l_f}{\mathrm{d}s}(p_f, \mathbf{p}). \tag{35}$$

Equations (32), (33) and the inequalities in (35) follow by substituting appropriate values for s in (28), while (34) follows by differentiating (28) with respect to s, substituting appropriately for s, and then using (12).

Next, define the function $G : \mathcal{D} \to \mathbb{R}$ by $G(s) \stackrel{\text{def}}{=} l_1^2(s) + \cdots + l_f^2(s) - 1$. We claim that $G(s) \leq 0$ for all $s \in \mathcal{D}$. In light of Lemma D.1 and the definition of $(1, f)$-volumetric points, our claim implies that 1) holds. Hence, to complete the proof, it is sufficient to prove our claim.

To prove our claim, note that G is a polynomial of degree $2(f-1)$. Also, we observe from (32), (34) and (35) that p_1 and p_f are roots of G of multiplicity 1, while each p_i is a root of G of multiplicity at least 2 for $i \neq 1, f$. Thus, the polynomial $H(s) \stackrel{\text{def}}{=} (s-p_1)(s-p_2)^2 \cdots (s-p_{f-1})^2(s-p_f)$ divides G. Since H also clearly has degree $2(f-1)$, it follows that there exists $K \in \mathbb{R}$ such that $G(s) = KH(s)$ for all $s \in \mathcal{D}$. The value of K may be computed as $K = \frac{G'(p_1)}{H'(p_1)}$, where $'$ indicates the derivative. It is easy to use (32) and (33) to verify that $G'(p_1) = 2\frac{\mathrm{d}l_1}{\mathrm{d}s}(p_1, \mathbf{p})$, which is negative by (35). An easy calculation also yields $H'(p_1) = (p_1 - p_2)^2 \cdots (p_1 - p_{f-1})^2(p_1 - p_f)$ which is negative since $p_1 < p_f$. These arguments show that $K > 0$. Our claim now follows by noting that H takes only non-positive values on $\mathcal{D}$. This completes the proof. $\square$

E ALGORITHM 2

Algorithm 2 VSBAI-Poly: Best Arm Identification for Polynomial Rewards

Input: $\varepsilon > 0$, $\delta \in (0, 1)$, sub-Gaussianity parameter σ, $(1, f)$-volumetric points $p_1, \ldots, p_f$ for $(\phi, \mathcal{D})$
Set $B_{1,f} = [\phi(p_1), \ldots, \phi(p_f)]$
Initialize $k \leftarrow 1$, $r \leftarrow 0$
Set STOP = False
while STOP==False **do**
  $\bar{y}^k$ = []
  for $t = 1, \ldots, f$ **do**
    $y_{(k-1)m+t} = g_\mu(p_t) + \eta_t$
    $\bar{y}^k \leftarrow [(\bar{y}^k)^{\mathrm{T}}; y_{(k-1)f+t}]^{\mathrm{T}}$
  end for
  $r = r + \bar{y}^k$
  if $\beta(k, \frac{\varepsilon}{4L}) < \delta$ **then**
    STOP = True
  else
    $k = k + 1$
  end if
end while
$\tau^* = kd$
$\hat{\mu}_{\tau^*} = \frac{1}{k} {B_{1,f}}^{-\mathrm{T}} r$
$\hat{s}$ = global_optimizer($\hat{\mu}_{\tau^*}, p_{\min}, p_{\max}$)
$\mathcal{D}_{\tau^*}$ = get_dtau($\hat{\mu}_{\tau^*}, \hat{s}, p_{\min}, p_{\max}, \varepsilon$)
Output: $\mathcal{D}_{\tau^*}$

Function global_optimizer($\hat{\mu}_{\tau^*}, p_{\min}, p_{\max}$)
$\hat{\mu}'_{\tau^*}$ = differentiate($\hat{\mu}_{\tau^*}$)
roots = find_roots($\hat{\mu}'_{\tau^*}$)
roots.add($p_{\min}, p_{\max}$)
values = $g_{\hat{\mu}_{\tau^*}}$(roots)
opt_value = argmax(values)
return opt_value

Function get_dtau($\hat{\mu}_{\tau^*}, \hat{s}, p_{\min}, p_{\max}, \varepsilon$)
d_tau = []
find_roots($g_{\hat{\mu}_{\tau^*}}(s) - g_{\hat{\mu}_{\tau^*}}(\hat{s}) + \varepsilon/2$)
roots.add($p_{\min}$, $p_{\max}$)
roots = sort(roots)
root_left = get_closest_left_root_to_$\hat{s}$(roots, $\hat{s}$)
root_right = get_closest_right_root_to_$\hat{s}$(roots, $\hat{s}$)
d_tau.add(root_left, root_right)
d_tau.add(every_pair_to_the_left_of_root_left)
d_tau.add(every_pair_to_the_right_of_root_right)
return d_tau

F MULTI-ARM SETTING CONFIGURATIONS

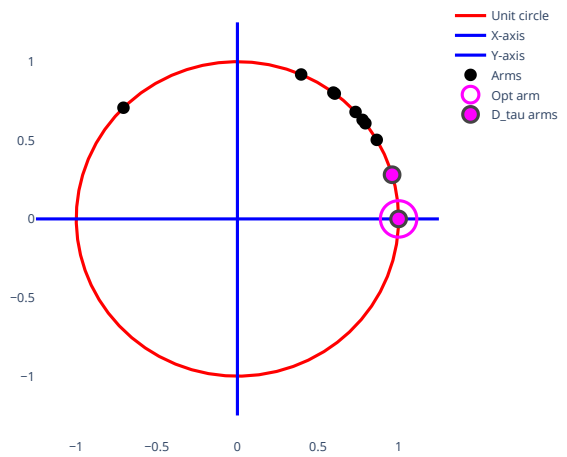

Figure 3: 10 arm setting when the angles ϕ of the arms (3 to 10) are sampled from $\mathcal{N}(0, .09)$. $(\varepsilon, \delta) = (0.1, 0.05)$ for the VSBAI algorithm

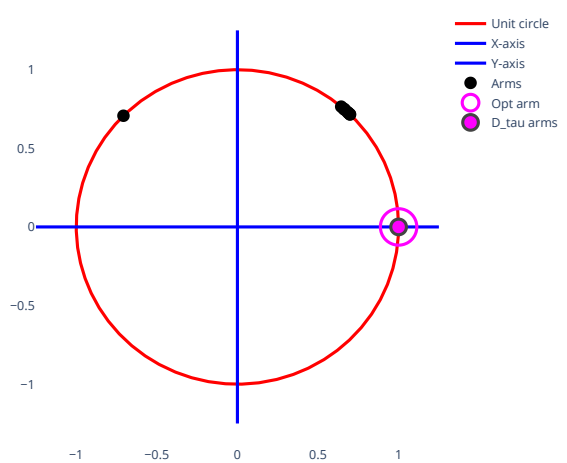

Figure 4: 10 arm setting when the angles ϕ of the arms (3 to 10þ) are sampled uniformly from $[0, 0.1]$. $(\varepsilon, \delta) = (0.1, 0.05)$ for the VSBAI algorithm

Algorithm	LazyTS		Rage		Oracle		VSBAI	
Arms	Mean	Std	Mean	Std	Mean	Std	Mean	Std
10	335.1	22.71	524.1	33.84	347.1	32.22	23847.2	52.61
20	423.05	28.92	683.05	92.07	356.15	32.31	23712.7	20.67
100	421.75	32.11	1038.75	148.01	426.4	39.76	23636.2	6.71
1000	419.65	28.43	1152.7	50.23	476.45	40.74	23611.8	0.63
2500	446.15	29.06	1447.3	150.46	481.8	41.33	23610.20	0.13
5000	431.65	32.23	1546.9	160.17	510.05	48.87	23610.4	0.18

Table 3: Expected sample complexity for the uniform setting described below

Algorithm	LazyTS		Rage		Oracle		VSBAI	
Arms	Mean	Std	Mean	Std	Mean	Std	Mean	Std
10	0.27	0.01	0.05	0.001	0.001	0.0	0.81	0.03
20	0.33	0.02	0.06	0.00	0.001	0.00	0.79	0.03
100	0.39	0.02	1.09	0.07	0.02	0.00	0.78	0.03
1000	34.78	1.06	27.61	0.33	0.69	0.02	0.79	0.03
2500	211.24	8.22	335.82	2.46	0.65	0.02	0.85	0.03
5000	422.35	22.45	884.32	3.85	0.89	0.03	1.58	0.03

Table 4: Run-time in sec for the uniform setting described below

G OTHER EXPERIMENTS

We consider the similar setting outlined in section 6.1 and present results for a different configuration of the problem instances. We first note that the implementation of the baseline algorithms presented in Fiez et al. [2019], Jedra and Proutiere [2020a], and Soare et al. [2014] for the setting in 6.1 is true when the angles ϕ_i for $i = 3, \ldots, n$ are sampled from a uniform distribution [0, 0.1]. So, we present the results for this uniform setting and compare with the VSBAI algorithm. Note that we were able reproduce the results (look Table 2 and Table 3 of Jedra and Proutiere [2020b]).

We observe that the sample complexity of VSBAI is greater than the other baselines. However, we argue that the instances generated in this setting are simple and in situations where it is difficult to separate out the best-arm from the next best (like when the angles ϕ_i of the arms are sampled from Gaussian Gaussian setting in 6.1), all these baselines suffer from huge sample complexities and run-times. In other words VSBAI is independent of the way the instances are generated but on the other hand all the other baselines are not robust, hence the potential to perform very bad in adversarial environments. In Table 4, we give the run-times for this setting. The results are averaged over 20 seeds (20 different problem instances are considered for each number of arms).