Optimal Gradient-based Algorithms for Non-concave Bandit Optimization

Baihe Huang[1]*, Kaixuan Huang[2]*, Sham M. Kakade[3,4]*, Jason D. Lee[2]*
Qi Lei[2]*, Runzhe Wang[2]*, Jiaqi Yang[5]*

[1]Peking University [2]Princeton University [3]University of Harvard
[4]Microsoft Research [5]Tsinghua University

Abstract

Bandit problems with linear or concave reward have been extensively studied, but relatively few works have studied bandits with non-concave reward. This work considers a large family of bandit problems where the unknown underlying reward function is non-concave, including the low-rank generalized linear bandit problems and two-layer neural network with polynomial activation bandit problem. For the low-rank generalized linear bandit problem, we provide a minimax-optimal algorithm in the dimension, refuting both conjectures in [54, 42]. Our algorithms are based on a unified zeroth-order optimization paradigm that applies in great generality and attains optimal rates in several structured polynomial settings (in the dimension). We further demonstrate the applicability of our algorithms in RL in the generative model setting, resulting in improved sample complexity over prior approaches. Finally, we show that the standard optimistic algorithms (e.g., UCB) are sub-optimal by dimension factors. In the neural net setting (with polynomial activation functions) with noiseless reward, we provide a bandit algorithm with sample complexity equal to the intrinsic algebraic dimension. Again, we show that optimistic approaches have worse sample complexity, polynomial in the extrinsic dimension (which could be exponentially worse in the polynomial degree).

1 Introduction

Bandits [50] are a class of online decision-making problems where an agent interacts with the environment, only receives a scalar reward, and aims to maximize the reward. In many real-world applications, bandit and RL problems are characterized by large or continuous action space. To encode the reward information associated with the action, function approximation for the reward function is typically used, such as linear bandits [20]. Stochastic linear bandits assume the mean reward to be the inner product between the unknown model parameter and the feature vector associated with the action. This setting has been extensively studied, and algorithms with optimal regret are known [20, 50, 2, 13].

However, linear bandits suffer from limited representation power unless the feature dimension is prohibitively large. A comprehensive empirical study [59] found that real-world problems required non-linear models and thus non-concave rewards to attain good performance on a testbed of bandit problems. To take a step beyond the linear setting, it becomes more challenging to design optimal algorithms. Unlike linear bandits, more sophisticated algorithms beyond optimism are necessary. For instance, a natural first step is to look at quadratic [43] and higher-order polynomial [35] reward. In

*Alphabetical order. Correspondence to: Qi Lei, `qilei@princeton.edu`, Jason D. Lee, `jasonlee@princeton.edu`.

35th Conference on Neural Information Processing Systems (NeurIPS 2021).

the context of phase retrieval, which is a special case for the quadratic bandit, people have derived algorithms that achieve minimax risks in the statistical learning setting [9, 51, 15]. However, the straightforward adaptation of these algorithms results in sub-optimal dimension dependency.

In the bandit domain, existing analysis on the nonlinear setting includes eluder dimension [61], subspace elimination [54, 42], etc. Their results also suffer from a larger dimension dependency than the best known lower bound in many settings. (See Table 1 and Section 3 for a detailed discussion of these results.) Therefore in this paper, we are interested in investigating the following question:

What is the optimal regret for non-concave bandit problems, including structured polynomials (low-rank etc.)? Can we design algorithms with optimal dimension dependency?

Contributions: In this paper, we answer the questions and close the gap (in problem dimension) for various non-linear bandit problems.

1. First, we design stochastic zeroth-order gradient-like[2] ascent algorithms to attain minimax regret for a large class of structured polynomials. The class of structured polynomials contains bilinear and low-rank linear bandits and symmetric and asymmetric higher-order homogeneous polynomial bandits with action dimension d. Though the reward is non-concave, we combine techniques from two bodies of work, non-convex optimization and numerical linear algebra, to design robust gradient-based algorithms that converge to global maxima. Our algorithms are also computationally efficient, practical, and easily implementable.
   In all cases, our algorithms attain the optimal dependence on dimension d, which was not previously attainable using existing optimism techniques. As a byproduct, our algorithm refutes[3] the conjecture from [42] on the bilinear bandit, and the conjecture from [54] on low-rank linear bandit by giving an algorithm that attains the optimal dimension dependence.
2. We demonstrate that our techniques for non-concave bandits extend to RL in the generative setting, improving upon existing optimism techniques.
3. When the reward is a general polynomial without noise, we prove that solving polynomial equations achieves regret equal to the intrinsic algebraic dimension of the underlying polynomial class, which is often linear in d for interesting cases. In general, this complexity cannot be further improved.
4. Furthermore, we provide a lower bound showing that all UCB algorithms have a sample complexity of $\Omega(d^p)$, where p is the degree of the polynomial. The dimension of *all homogeneous polynomials* of degree p in dimension d is d^p, showing that UCB is oblivious to the polynomial class and highly sub-optimal even in the noiseless setting.

1.1 Related Work

Linear Bandits. Linear bandit problems and their variants are studied in [20, 50, 1, 13, 2, 20, 60, 34, 33]. The matching upper bound and minimax lower bound achieves $O(d\sqrt{T})$ regret. Structured linear bandits, including sparse linear bandit [3] which developed an online-to-confidence-set technique. This technique yields the optimal $O(\sqrt{sdT})$ rate for sparse linear bandit. However [54] employed the same technique for low-rank linear bandits giving an algorithm with regret $O(\sqrt{d^3\mathrm{poly}(k)T})$ which we improve to $O(\sqrt{d^2\mathrm{poly}(k)T})$, which meets the lower bound given in [54].
Eluder Dimension. [61] proposed the eluder dimension as a general complexity measure for nonlinear bandits. However, the eluder dimension is only known to give non-trivial bounds for linear function classes and monotone functions of linear function classes. For structured polynomial classes, the eluder dimension simply embeds into an ambient linear space of dimension d^p, where d is the dimension and p is the degree. This parallels the linearization/NTK line in supervised learning [70, 30, 7] which show that linearization also incurs a similarly large penalty of d^p sample complexity, and more advanced algorithm design is need to circumvent linearization [10, 17, 25, 71, 27, 58, 29, 56, 36, 67, 19].
Non-concave Bandits. To our knowledge, there is no general study of non-concave bandits, likely due to the difficulty of globally maximizing non-concave functions. A natural starting point of

[2]Our algorithm estimates the gradient, but with some irreducible bias for the tensor case. Note that our algorithms converge linearly despite the bias.

[3]Both papers conjectured regret of the form $O(\sqrt{d^3\mathrm{poly}(k)T})$ based on convincing but potentially misleading heuristics.

studying the non-concave setting are quadratic rewards such as the Rayleigh quotient, or namely bandit PCA [46, 28]. In the bilinear [42] and low-rank linear setting [54] with rank k parameter matrices achieved $\widetilde{O}(\sqrt{d^3\mathrm{poly}(k)T})$ regret. Other literature considered related but different settings that are not comparable to our results [43, 41, 32, 47]. We note that the regret of all previous work is at least $O(\sqrt{d^3T})$. This includes the subspace exploration and refinement algorithms from [42, 54], or from eluder dimension [61]. Recently [35] considers online problem with a low-rank tensor associated with axis-aligned set of arms, which corresponds to finding the largest entry of the tensor. Finally, [46] study the bandit PCA problem in the adversarial setting, attaining regret of $O(\sqrt{d^3T})$. We leave adapting our results to the adversarial setting as an open problem.
Due to space limitation we defer the discussion of some equally important literature to the appendix.

2 Preliminaries

2.1 Setup: Structured Polynomial Bandit

We study structured polynomial bandit problems where the reward function is from a class of structured polynomials (see below). A player plays actions for T rounds; at each round $t \in [T]$, the player chooses one action $\boldsymbol{a}_t$ from the feasible action set $\mathcal{A}$ and receives the reward r_t afterward.

We consider both the stochastic case where $r_t = f_{\boldsymbol{\theta}}(\boldsymbol{a}_t) + \eta_t$ where η_t is the random noise, and the noiseless case $r_t = f_{\boldsymbol{\theta}}(\boldsymbol{a}_t)$. Specifically the function $f_{\boldsymbol{\theta}}$ is unknown to the player, but lies in a known function class $\mathcal{F}$. We use the notation $\mathrm{vec}(\boldsymbol{M})$ to denote the vectorization of a matrix or a tensor $\boldsymbol{M}$, and $\boldsymbol{v}^{\otimes p}$ to denote the p-order tensor product of a vector $\boldsymbol{v}$. For vectors we use $\|\cdot\|_2$ or $\|\cdot\|$ to denote its ℓ_2 norm. For matrices $\|\cdot\|_2$ or $\|\cdot\|$ stands for its spectral norm, and $\|\cdot\|_F$ is Frobenius norm. For integer n, $[n]$ denotes set $\{1, 2, \cdots n\}$. For cleaner presentation, in the main paper we use $\widetilde{O}$, $\widetilde{\Theta}$ or $\widetilde{\Omega}$ to hide universal constants, polynomial factors in p and polylog factors in dimension d, error ε, eigengap Δ, total round number T or failure rate δ.

We now present the outline with the settings considered in the paper:

The stochastic bandit eigenvector case, $\mathcal{F}_{\text{EV}}$, considers action set $\mathcal{A} = \{\boldsymbol{a} \in \mathbb{R}^d : \|\boldsymbol{a}\|_2 \leq 1\}$ as shown in Section 3.1, and

$$\mathcal{F}_{\text{EV}} = \left\{ \begin{array}{l} f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \boldsymbol{a}^T \boldsymbol{M} \boldsymbol{a}, \boldsymbol{M} = \sum_{j=1}^k \lambda_j \boldsymbol{v}_j \boldsymbol{v}_j^\top, \text{ for orthonormal } \boldsymbol{v}_j \\ \boldsymbol{M} \in \mathbb{R}^{d\times d}, 1 \geq \lambda_1 \geq |\lambda_2| \geq \cdots \geq |\lambda_k| \end{array} \right\}.$$

The stochastic low-rank linear reward case, $\mathcal{F}_{\text{LR}}$ considers action sets on bounded matrices $\mathcal{A} = \{\boldsymbol{A} \in \mathbb{R}^{d\times d} : \|\boldsymbol{A}\|_{\text{F}} \leq 1\}$ in Section 3.2, and

$$\mathcal{F}_{\text{LR}} = \left\{ \begin{array}{l} f_{\boldsymbol{\theta}}(\boldsymbol{A}) = \langle \boldsymbol{M}, \boldsymbol{A} \rangle = \mathrm{vec}(\boldsymbol{M})^\top \mathrm{vec}(\boldsymbol{A}), \\ \boldsymbol{M} \in \mathbb{R}^{d\times d}, \mathrm{rank}(\boldsymbol{M}) = k, \boldsymbol{M} = \boldsymbol{M}^\top, \|\boldsymbol{M}\|_{\text{F}} \leq 1 \end{array} \right\}.$$

We illustrate how to apply the established bandit oracles to attain a better sample complexity for RL problems with the simulator in Section 3.2.1.

The stochastic homogeneous polynomial reward case is presented in Section 3.3. For the symmetric case in Section 3.3.1, the action sets are $\mathcal{A} = \{\boldsymbol{a} \in \mathbb{R}^d : \|\boldsymbol{a}\|_2 \leq 1\}$, and

$$\mathcal{F}_{\text{SYM}} = \left\{ \begin{array}{l} f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j (\boldsymbol{v}_j^\top \boldsymbol{a})^p \text{ for orthonormal } \boldsymbol{v}_j, \\ 1 \geq r^* = \lambda_1 > |\lambda_2| \geq \cdots \geq |\lambda_k| \end{array} \right\};$$

in the asymmetric case in Section D.3.1, the action sets are
$\mathcal{A} = \{\boldsymbol{a} = \boldsymbol{a}(1) \otimes \boldsymbol{a}(2) \otimes \cdots \otimes \boldsymbol{a}(p) \in \mathbb{R}^{d^p} : \forall q \in [p], \|\boldsymbol{a}(q)\|_2 \leq 1\}$, and

$$\mathcal{F}_{\text{ASYM}} = \left\{ \begin{array}{l} f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j \prod_{q=1}^p (\boldsymbol{v}_j(q)^\top \boldsymbol{a}(q)) \text{ for orthonormal } \boldsymbol{v}_j(q) \text{ for each } q, \\ 1 \geq r^* = |\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_k| \end{array} \right\}.$$

In the above settings, there is stochastic noise on the observed rewards. We also consider noiseless settings as below:
The noiseless polynomial reward case is deferred to Appendix G.1 due to space limitation. The action sets $\mathcal{A}$ are subsets of $\mathbb{R}^d$, and $\mathcal{F}_{\text{P}} = \{f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \langle \boldsymbol{\theta}, \widetilde{\boldsymbol{a}}^{\otimes p} \rangle : \boldsymbol{\theta} \in \mathcal{V}, \widetilde{\boldsymbol{a}} = [1, \boldsymbol{a}^\top]^\top, \mathcal{V} \subseteq (\mathbb{R}^{d+1})^{\otimes p}$ is an algebraic variety$\}$. Additionally, $\mathcal{F}$ needs to be admissible (Definition G.2). This class includes two-layer neural networks with polynomial activations (i.e. structured polynomials). We study the fundamental limits of all UCB algorithms in Section G.1.2 as they are $\Omega(d^{p-1})$ worse than our algorithm presented in Section G.1.1.

			$\mathcal{F}_{\text{HPS}}$	$\mathcal{F}_{\text{HPA}}$	$\mathcal{F}_{\text{B}}$	$\mathcal{F}_{\text{LL}}$
Upper Bound from LinUCB/eluder			$\widetilde{O}(d^{p+1}k\epsilon^{-2})(k>p)$	$\widetilde{O}(d^{p+1}k\epsilon^{-2})$	$\widetilde{O}(d^3k\epsilon^{-2})$	$\widetilde{O}(d^3k\epsilon^{-2})$
Our Results	NPM	Gap-dependent	$\widetilde{O}(\kappa d^p k\Delta^{-2}\epsilon^{-2})$ (even p)	N/A	$\widetilde{O}(\kappa d^2\Delta^{-2}\epsilon^{-2})$	$\widetilde{O}(d^2k^2\lambda_k^{-2}\epsilon^{-2})$
		Gap-independent	$\widetilde{O}(d^pk\epsilon^{-2})$ (odd p)	$\widetilde{O}((Ck)^pd^p\epsilon^{-2})$	$\widetilde{O}(\min\{d^2\epsilon^{-5}, d^2k^5\epsilon^{-2}\})$	N/A
	Lower bounds		$\Omega(d^p\epsilon^{-2})$	$\Omega(d^p\epsilon^{-2})$	$\Omega(d^2\epsilon^{-2})$	$\Omega(d^2k^2\epsilon^{-2})^*$

Table 1: Baselines and Our Main Results (for stochastic settings). Eigengap $\Delta = \lambda_1 - |\lambda_2|$, $\kappa = \lambda_1/\Delta$. C is a constant that depends on failure probability. The result with $*$ is from [54].

In the above settings, we are concerned with the cumulative regret $\mathfrak{R}(T)$ for T rounds. Let $f^*_{\boldsymbol{\theta}} = \sup_{\boldsymbol{a}\in\mathcal{A}} f_{\boldsymbol{\theta}}(\boldsymbol{a})$, $\mathfrak{R}_{\boldsymbol{\theta}}(T) := \sum_{t=1}^{T}(f^*_{\boldsymbol{\theta}} - f_{\boldsymbol{\theta}}(\boldsymbol{a}_t))$. Since the parameters can be chosen adversarially, we are bounding $\mathfrak{R}(T) = \sup_{\boldsymbol{\theta}} \mathfrak{R}_{\boldsymbol{\theta}}(T)$ in this paper. In all the stochastic settings above, we make the standard assumption on stochasticity that η_t is conditionally zero-mean 1-sub-Gaussian random variable regarding the randomness before t.

2.2 Warm-up: Adapting Existing Algorithms

In all of the above settings, the function class can be viewed as a generalized linear function. Namely, there is fixed feature maps ψ, ϕ so that $f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \psi(\boldsymbol{\theta})^T\phi(\boldsymbol{a})$. Thus it is straightforward to adapt linear bandit algorithms like the renowned LinUCB [52] to our settings. Furthermore, another baseline is given by the eluder dimension argument [61][66] which gives explicit upper bounds for general function classes. We present the best upper bound by adapting these methods as a baseline in Table 1, together with our newly-derived lower bound and upper bound in this paper.
The best-known statistical rates are based on the following result.

Theorem 2.1 (Proposition 4 in [61])**.** *With $\alpha = O(T^{-2})$ appropriately small, given the α-covering-number N (under $\|\cdot\|_\infty$) and the α-eluder-dimension d_E of the function class $\mathcal{F}$, Eluder UCB (Algorithm 2) achieves regret $\widetilde{O}(\sqrt{d_E T\log N})$.*

In the first row of Table 1, we further elaborate on the best results obtained from Theorem 2.1 in individual settings. More details can be found in Appendix B.

3 Main results

We now present our main results. We consider 4 different stochastic settings (see Table 1) and one noiseless setting with structured polynomials.
In the cases of stochastic reward, all our algorithms can be unified as gradient-based optimization. At each stage with a candidate action $\boldsymbol{a}$, we define the estimator $G_n(\boldsymbol{a}) := \frac{1}{n}\sum_{i=1}^n (f_{\boldsymbol{\theta}}((1-\zeta)\boldsymbol{a}+\zeta\boldsymbol{z}_i)+\eta_i)\boldsymbol{z}_i$, with $\boldsymbol{z}_i \sim \mathcal{N}(0, \sigma^2\boldsymbol{I}_d)$ and proper step-size ζ [26]. Therefore $\mathbb{E}_{\boldsymbol{z}}[G(\boldsymbol{a})] = \sigma^2\zeta\nabla f_{\boldsymbol{\theta}}((1-\zeta)\boldsymbol{a}) + O(\zeta^2) = \zeta(1-\zeta)^{p-1}\sigma^2\nabla f_{\boldsymbol{\theta}}(\boldsymbol{a}) + O(\zeta^2)$ for p-th order homogeneous polynomials. Therefore with enough samples, we are able to implement noisy gradient ascent with bias.
In the noiseless setting, our algorithm solves for the parameter $\boldsymbol{\theta}$ with randomly sampled actions $\{\boldsymbol{a}_t\}$ and the noiseless reward $\{f_{\boldsymbol{\theta}}(\boldsymbol{a}_t)\}$, and then determines the optimal action by computing $\arg\max_{\boldsymbol{a}} f_{\boldsymbol{\theta}}(\boldsymbol{a})$.

3.1 Stochastic Eigenvalue Reward ($\mathcal{F}_{\text{EV}}$)

Now consider bandits with stochastic reward $r(\boldsymbol{a}) = \boldsymbol{a}^\top\boldsymbol{M}\boldsymbol{a} + \eta$ with action set $\mathcal{A} = \{\boldsymbol{a}\,|\,\|\boldsymbol{a}\|_2 \le 1\}$. $\boldsymbol{M} = \sum_{i=1}^r \lambda_i \boldsymbol{v}_i\boldsymbol{v}_i^\top$ is symmetric and satisfies $r^* = \lambda_1 > |\lambda_2| \ge |\lambda_3| \cdots \ge |\lambda_r|, \eta \sim \mathcal{N}(0,1)$.
Denote by $\boldsymbol{a}^*$ the optimal action ($\pm\boldsymbol{v}_1$), the leading eigenvector of $\boldsymbol{M}$, $(\boldsymbol{a}^*)^\top\boldsymbol{M}\boldsymbol{a}^* = \lambda_1$. Let $\Delta = \lambda_1 - |\lambda_2| > 0$ be the eigengap and $\kappa := \lambda_1/\Delta$ be the condition number.

Remark 3.1 (Negative leading eigenvalue)**.** *For a symmetric matrix $\boldsymbol{M}$, we will conduct noisy power method to recover its leading eigenvector, and therefore we require its leading eigenvalue λ_1 to be positive. It is straightforward to extend to the setting where the nonzero eigenvalues satisfy: $r^* \equiv \lambda_1 > \lambda_2 \ge \lambda_l > 0 > \lambda_{l+1} \cdots \ge \lambda_k$, and $|\lambda_k| > \lambda_1$. For this problem, we can shift $\boldsymbol{M}$ to get $\boldsymbol{M} + |\lambda_k|\boldsymbol{I}$ and the eigen-spectrum now becomes $\lambda_1 + |\lambda_k|, \lambda_2 + |\lambda_k|, \cdots 0$; therefore, we can still recover the optimal action with dependence on the new condition number $(\lambda_1 + |\lambda_k|)/(\lambda_1 - \lambda_2)$.*

Remark 3.2 (Asymmetric matrix)**.** *Our algorithm naturally extends to the asymmetric setting:* $f(\boldsymbol{a}_1, \boldsymbol{a}_2) = \boldsymbol{a}_1^\top \widetilde{\boldsymbol{M}} \boldsymbol{a}_2$, *where* $\widetilde{\boldsymbol{M}} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$. *This setting can be reduced to the symmetric case via defining*

$$\boldsymbol{M} = \begin{bmatrix} 0 & \widetilde{\boldsymbol{M}}^\top \\ \widetilde{\boldsymbol{M}} & 0 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} \boldsymbol{V} & \boldsymbol{V} \\ \boldsymbol{U} & -\boldsymbol{U} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Sigma} & 0 \\ 0 & -\boldsymbol{\Sigma} \end{bmatrix} \begin{bmatrix} \boldsymbol{V} & \boldsymbol{V} \\ \boldsymbol{U} & -\boldsymbol{U} \end{bmatrix}^\top,$$

which is a symmetric matrix, and its eigenvalues are $\pm\sigma_i(\widetilde{\boldsymbol{M}})$, *the singular values of* $\widetilde{\boldsymbol{M}}$. *Therefore our analysis on symmetric matrices also applies to the asymmetric setting and will equivalently depend on the gap between the top singular values of* $\widetilde{\boldsymbol{M}}$. *A formal asymmetric to symmetric conversion algorithm is presented in Algorithm 1 in [28].*

Algorithm. We note that by conducting zeroth-order gradient estimate $1/n \sum_{i=1}^n (f(\boldsymbol{a}/2 + \boldsymbol{z}_i/2) + \eta_i)\boldsymbol{z}_i$ with step-size $1/2$ [26] and sample size n, we get an estimate for $\mathbb{E}_{\eta,\boldsymbol{z}}[(f(\boldsymbol{a}/2 + \boldsymbol{z}/2) + \eta)\boldsymbol{z}] = \frac{\sigma^2}{2}\boldsymbol{M}\boldsymbol{a}$ when $\boldsymbol{z} \sim \mathcal{N}(0, \sigma^2)$. Therefore we are able to use noisy power method to recover the top eigenvector. We present the complete algorithm in Algorithm 1 and attain a gap-dependent risk bound:

Algorithm 1 Noisy power method for bandit eigenvalue problem.

Input: Quadratic function $f : \mathcal{A} \to \mathbb{R}$ with noisy reward, failure probability δ, error ε.
Initialization: Initial action $\boldsymbol{a}_0 \in \mathbb{R}^d$ randomly sampled on the unit sphere $\mathbb{S}^{d-1}$. We set $\alpha = |\lambda_2/\lambda_1|$, sample size per iteration $n = C_n d^2 \log(d/\delta)(\lambda_1\alpha)^{-2}\varepsilon^{-2}$, sample variance $m = C_m d \log(n/\delta)$, total iteration $L = \lfloor C_L \log(d/\varepsilon) \rfloor + 1$.
for Iteration l from 1 to L **do**
  Sample $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/mI_d), i = 1, 2, \cdots n$. (Re-sample the whole batch if exists $\boldsymbol{z}_i$ with norm greater than 1.)
  Noisy power method:
  Take actions $\widetilde{\boldsymbol{a}}_i = \frac{\boldsymbol{a}_{l-1} + \boldsymbol{z}_i}{2}$ and observe $r_i = f(\boldsymbol{a}_i) + \eta_i, \forall i \in [n]$
  Update normalized action $\boldsymbol{a}_l \leftarrow \frac{m}{n}\sum_{i=1}^n r_i \boldsymbol{z}_i$, and normalize $\boldsymbol{a}_l \leftarrow \boldsymbol{a}_l/\|\boldsymbol{a}_l\|_2$.
Output: $\boldsymbol{a}_L$.

Theorem 3.3 (Regret bound for noisy power method (NPM))**.** *In Algorithm 1, we set* $\varepsilon \in (0, 1/2)$, $\delta = 0.1/(L_0 S)$ *and let* C_L, C_S, C_n, C_m *be large enough universal constants. Then with high probability* 0.9 *we have: the output* $\boldsymbol{a}_L$ *satisfies* $\tan\theta(\boldsymbol{a}^*, \boldsymbol{a}_L) \le \varepsilon$ *and yields* $r^*\varepsilon^2$*-optimal reward; and the total number of actions we take is* $\widetilde{O}(\frac{\kappa d^2}{\Delta^2\varepsilon^2})$. *By explore-then-commit (ETC) the cumulative regret is at most* $\widetilde{O}(\sqrt{\kappa^3 d^2 T})$.

All proofs in this subsection are in Appendix C.1.

Remark 3.4 (Intuition of [42] and how to overcome the conjectured lower bound via the design of adaptive algorithms)**.** *Let us consider the rank* 1 *case of* $r(\boldsymbol{a}) = (\boldsymbol{a}^T\boldsymbol{\theta}^*)^2 + \eta$. *A random action* $\boldsymbol{a} \sim$ *Unif*$(\mathbb{S}^{d-1})$ *has* $f(\boldsymbol{a}) \asymp 1/d^2$, *and the noise has standard deviation* $O(1)$. *Thus the signal-to-noise-ratio is* $O(1/d^2)$ *and the optimal action* θ^* *requires* d *bits to encode. If we were to play non-adaptively, this would require* $O(d^3)$ *queries and result in regret* $\sqrt{d^3 T}$ *which matches the result of [42].*

To go beyond this, we must design algorithms that are adaptive, meaning the information in $f(\boldsymbol{a}) + \eta$ *is* ***strictly larger than*** $\frac{1}{d^2}$. *As an illustration of why this is possible, consider batching the time-steps into* d *stages so that each stage decode* 1 *bit of* $\boldsymbol{\theta}^*$. *At the first stage, random exploration* $\boldsymbol{a} \sim$ *Unif*$(\mathbb{S}^{d-1})$ *gives signal-to-noise-ratio* $O(1/d^2)$. *Suppose* k *bits of* $\boldsymbol{\theta}$ *are decoded at* k*-th stage by* $\widehat{\boldsymbol{\theta}}$, *adaptive algorithms can boost the signal-to-noise-ratio to* $O(k/d^2)$ *by using* $\widehat{\boldsymbol{\theta}}$ *as bootstrap (e.g. exploring with* $\widehat{\boldsymbol{\theta}} \pm a$ *where* a *is random exploration in the unexplored subspace). In this way adaptive algorithms only need* d^2/k *queries in* $(k+1)$*-th stage and so the total number of queries sums up to* $\sum_{k=1}^d d^2/k \approx d^2 \log d$.

Gradient descent and power method offer a computationally efficient and seamless way to implement the above intuition. For every iterate action $\boldsymbol{a}$, *we estimate* $\boldsymbol{M}\boldsymbol{a}$ *from noisy observations and take it as our next action* $\boldsymbol{a}^+$. *With* $d^2/(\Delta^2\varepsilon^2)$ *samples, noisy power method enjoys linear progress*

$\tan\theta(\boldsymbol{a}^+,\boldsymbol{a}^*) \le \max\{c\tan(\boldsymbol{a},\boldsymbol{a}^*),\varepsilon\}$, *where* $c<1$ *is a constant that depends on* λ_1,λ_2, *and* ε. *Therefore even though every step costs* d^2 *samples, overall we only need logarithmic (in* $d,\varepsilon,\lambda_1,\lambda_2$*) iterations to find an* ε*-optimal action.*

Remark 3.5 (Connection to phase retrieval and eluder dimension)**.** *For rank-1 case* $\boldsymbol{M}=\boldsymbol{x}\boldsymbol{x}^\top$, *the bilinear bandits can be viewed as phase retrieval, where one observes* $y_r=(\boldsymbol{a}_r^\top\boldsymbol{x})^2=\boldsymbol{a}_r^\top\boldsymbol{M}\boldsymbol{a}_r$ *plus some noise* $\eta_r\sim\mathcal{N}(0,\sigma^2)$. *The optimal (among non-adaptive algorithms) sample complexity to recover* $\boldsymbol{x}$ *is* $\sigma^2 d/\epsilon^2$ *[16, 15] where they play* $\boldsymbol{a}$ *from random Gaussian* $\mathcal{N}(0,\boldsymbol{I})$. *However, in bandit, we need to set the variance of* $\boldsymbol{a}$ *to at most* $1/d$ *to ensure* $\|\boldsymbol{a}\|\le 1$. *Our problem is equivalent to observing* y_r/d *where their* $\boldsymbol{a}/\sqrt{d}\sim\mathcal{N}(0,1/d\boldsymbol{I})$ *and noise level* $\eta_r/d\sim\mathcal{N}(0,1)$, *i.e.,* $\sigma^2=d^2$. *Therefore one gets* d^3/ϵ^2 *even for the rank-1 problems. On the other hand, for all rank-1* $\boldsymbol{M}$ *the condition number* $\kappa=1$*; and thus our results match the lower bound (see Section 3.3.2) up to logarithmic factors, and also have fundamental improvements for phase retrieval problems by leveraging adaptivity.*

For UCB algorithms based on eluder dimension, the regret upper bound is $O(\sqrt{d_E\log(N)T})=O(\sqrt{d^3T})$ *as presented in Theorem 2.1, where the dependence on* d *is consistent with [42, 54] and is non-optimal.*

The previous result depends on the eigen-gap. When the matrix is ill-conditioned, i.e., λ_1 is very close to λ_2, we can obtain gap-free versions with a modification: The first idea stems from finding higher reward instead of recovering the optimal action. Therefore when λ_1 and λ_2 are very close (gap being smaller than desired accuracy ε), it is acceptable to find any direction in the span of $(\boldsymbol{v}_1,\boldsymbol{v}_2)$. More formally, we care about the convergence speed of identifying any action in the space spanned by any top eigenvectors (whose associated eigenvalues are higher than $\lambda_1-\varepsilon$). Therefore the convergence speed will depend on ε instead of $\lambda_1-\lambda_2$.

Corollary 3.6 (Gap-free regret bound)**.** *For positive semi-definite matrix* $\boldsymbol{M}$, *by setting* $\alpha=1-\varepsilon^2/2$ *in Algorithm 1 and performing ETC afterwards, one can obtain cumulative regret of* $\widetilde{O}(\lambda_1^{3/5}d^{2/5}T^{4/5})$.

Again, the PSD assumption is not essential. For general symmetric matrices with $\lambda_1\ge\lambda_2\ge\cdots>0>\cdots\lambda_k$. We can still conduct shifted power method on $\boldsymbol{M}-\lambda_k\boldsymbol{I}$, yielding a cumulative regret of $(\lambda_1+|\lambda_k|)^{3/5}d^{2/5}T^{4/5}$.

Another novel gap-free algorithm requires to identify any top eigenspace $\boldsymbol{V}_{1:l}, l\in[k]$: $\boldsymbol{V}_{1:l}$ is the column span of $\{\boldsymbol{v}_1,\boldsymbol{v}_2,\cdots\boldsymbol{v}_l\}$. Notice that in traditional subspace iteration, the convergence rate of recovering $\boldsymbol{V}_l$ depends on the eigengap $\Delta_l:=|\lambda_l|-|\lambda_{l+1}|$. Meanwhile, since $\sum_{l=1}^k\Delta_l=\lambda_1$, at least one eigengap is larger or equal to λ_1/k. Suppose $\Delta_{l^*}\ge\lambda_1/k$; we can, therefore, set $\alpha=1-1/k$ and recover the top l^* subspace up to $\lambda_1\epsilon$ error, which will give an ϵ-optimal reward in the end. We don't know l^* beforehand and will try recovering the top subspace $\boldsymbol{V}_1,\boldsymbol{V}_2,\cdots\boldsymbol{V}_k$ respectively, which will only lose a k factor. With the existence of l^*, at least one trial (on recovering $\boldsymbol{V}_{l^*}$) will be successful with the parameters of our choice, and we simply output the best action among all trials.

Theorem 3.7 (Informal statement: (gap-free) subspace iteration)**.** *By running subspace iteration (Algorithm 3) with proper choices of parameters, we attain a cumulative regret of* $\widetilde{O}(\lambda_1^{1/3}k^{4/3}(dT)^{2/3})$. *Algorithm 3) with another set of parameters can also recover the whole eigenspace, and achieve cumulative regret of* $\widetilde{O}((\lambda_1 k)^{1/3}(\widetilde{\kappa}dT)^{2/3})$, *where* $\widetilde{\kappa}=\lambda_1/|\lambda_k|$.

3.2 Stochastic Low-rank Linear Bandits ($\mathcal{F}_{\mathbf{LR}}$)

In the low-rank linear bandit, the reward function is $f(\boldsymbol{A})=\langle\boldsymbol{A},\boldsymbol{M}\rangle$, with noisy observations $r_t=f(\boldsymbol{A})+\eta_t$, and the action space is $\{\boldsymbol{A}\in\mathbb{R}^{d\times d}:\|\boldsymbol{A}\|_F\le 1\}$. Without loss of generality we assume k, the rank of $\boldsymbol{M}$ satisfies $k\le\frac{d}{2}$, since when k is of the same order as d, the known upper and lower bound are both $\sqrt{d^2k^2T}=\Theta(\sqrt{d^4T})$ [54] and there is no room for improvement. We write $r^*=\|\boldsymbol{M}\|_F\le 1$, therefore the optimal action is $\boldsymbol{A}^*=\boldsymbol{M}/r^*$. In this section, we write $\boldsymbol{X}(s)$ to be the s-th column of any matrix $\boldsymbol{X}$.

As presented in Algorithm 4, we conduct noisy subspace iteration to estimate the right eigenspace of $\boldsymbol{M}$. Subspace iteration requires calculating $\boldsymbol{M}\boldsymbol{X}_t$ at every step. This can be done by considering

$\mathcal{F}_{\text{EV}}$	LB ($k=1$)	[42]	NPM	Gap-free NPM	Subspace Iteration
Regret	$\sqrt{d^2T}$	$\sqrt{d^3k\lambda_k^{-2}T}$	$\sqrt{\kappa^3d^2T}$	$d^{2/5}T^{4/5}$	$\min(k^{4/3}(dT)^{2/3}, k^{1/3}(\widetilde{\kappa}dT)^{2/3})$

$\mathcal{F}_{\text{LR}}$	LB ([54])	UB ([54])	Subspace Iteration
Regret	$\Omega(\sqrt{d^2k^2T})$	$\sqrt{d^3kT}^*$ or $\sqrt{d^3k\lambda_k^{-2}T}$	$\min(\sqrt{d^2k\lambda_k^{-2}T}, (dkT)^{2/3})$

Table 2: Summary of results for quadratic reward. All red expressions are our results. LB, UB, NPM stands for lower bound, upper bound, and noisy power method respectively. $\Delta = \lambda_1 - |\lambda_2|$ is the eigengap, and $\kappa = \lambda_1/\Delta$, $\widetilde{\kappa} = \lambda_1/|\lambda_k|$ are the condition numbers. The result with $*$ is not computationally tractable. For low-rank setting in this table, we treat $\|\boldsymbol{M}\|_F$ as a constant for simplicity and leave its dependence in the theorems. Our upper bounds match the lower bound in terms of dimension and substantially improve over existing algorithms that are computationally efficient.

a change of variable of $g(\boldsymbol{X}) := f(\boldsymbol{X}\boldsymbol{X}^\top) = \langle \boldsymbol{X}\boldsymbol{X}^\top, \boldsymbol{M}\rangle$ whose gradient[4] is $\nabla g(\boldsymbol{X}) = 2\boldsymbol{M}\boldsymbol{X}$. The zeroth-order gradient estimator can then be employed to stochastically estimate $\boldsymbol{M}\boldsymbol{X}$. We instantiate the analysis with symmetric $\boldsymbol{M}$ while extending to asymmetric setting is straightforward since the problem can be reduced to symmetric setting (suggested in Remark 3.2).

With stochastic observations and randomly sampled actions, we achieve the next iterate $\boldsymbol{Y}_l$ that satisfies $\boldsymbol{M}\boldsymbol{X}_{l-1} \equiv \mathbb{E}[\boldsymbol{Y}_l]$ in Algorithm 4. Let $\boldsymbol{M} = \boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{V}^\top$. With proper concentration bounds presented in the appendix, we can apply the analysis of noisy power method [37] and get:

Theorem 3.8 (Informal statement; sample complexity for low-rank linear reward)**.** *With Algorithm 4, $\boldsymbol{X}_L$ satisfies $\|(\boldsymbol{I} - \boldsymbol{X}_L\boldsymbol{X}_L^\top)\boldsymbol{V}\| \leq \varepsilon/4$, and output $\widehat{\boldsymbol{A}}$ satisfies $\|\widehat{\boldsymbol{A}} - \boldsymbol{A}^*\|_F \leq \varepsilon\|\boldsymbol{M}\|_F$ with sample size $\widetilde{O}(d^2k\lambda_k^{-2}\varepsilon^{-2})$.*

We defer the proofs for low-rank linear reward in Appendix C.2.

Corollary 3.9 (Regret bound for low-rank linear reward)**.** *We first call Algorithm 4 with $\varepsilon^4 = \widetilde{\Theta}(d^2k\lambda_k^{-2}T^{-1})$ to obtain $\widehat{\boldsymbol{A}}$; we then play $\widehat{\boldsymbol{A}}$ for the remaining steps. The cumulative regret satisfies*

$$\mathfrak{R}(T) \leq \widetilde{O}(\sqrt{d^2k(r^*)^2\lambda_k^{-2}T}),$$

with high probability 0.9.

To be more precise, we need $T \geq \widetilde{\Theta}(d^2k/\lambda_k^2)$ for Algorithm 4 to take sufficient actions; however, the conclusion still holds for smaller T. Since simply playing 0 for all T actions will give a sharper bound of $\mathfrak{R}(T) \leq r^*T \leq \widetilde{O}(r^*\sqrt{d^2k\lambda_k^{-2}T})$). For cleaner presentation, we won't stress this for every statement.

Notice that $k \leq \|\boldsymbol{M}\|_F^2/\lambda_k^2 \leq k\widetilde{\kappa}^2$ is order k. Thus for well-conditioned matrices $\boldsymbol{M}$, our upper bound of $\widetilde{O}(\sqrt{d^2k\|\boldsymbol{M}\|_F^2/\lambda_k^2T})$ matches the lower bound $\sqrt{d^2k^2T}$ except for logarithmic factors.

In the previous setting with the bandit eigenvalue problem, estimating $\boldsymbol{M}$ up to an ϵ-error (measured by **operator norm**) gives us an ϵ-optimal reward. Therefore the sample complexity for eigenvalue reward with similar subspace iteration is $\|\boldsymbol{M}\|_2^2d^2k/(\lambda_k^2\epsilon^2)$. In this section, on the other hand, we need **Frobenius norm** bound $\|\boldsymbol{A}_L - \boldsymbol{A}^*\|_F \leq \epsilon$; naturally the complexity becomes $\|\boldsymbol{M}\|_F^2d^2k/(\lambda_k^2\epsilon^2)$.

Theorem 3.10 (Regret bound for low-rank linear reward: gap-free case)**.** *Set $\varepsilon^6 = \Theta(\frac{d^2k^2}{(r^*)^2T})$, $n = \widetilde{\Theta}(\frac{d^2k^2}{(r^*)^2\varepsilon^4})$, $L = \Theta(\log(d/\varepsilon))$ and $k' = 2k$ in Algorithm 4 and get $\widehat{\boldsymbol{A}}$. Then we play it for the remaining steps, the cumulative regret satisfies:*

$$\mathfrak{R}(T) \leq \widetilde{O}((dkT)^{2/3}(r^*)^{1/3}).$$

We summarize all our results and prior work for quadratic reward in Table 2.

[4]Directly performing projected gradient descent on $f(\boldsymbol{A})$ would not work, since this is not an adaptive algorithm as the gradient of a linear function is constant. This would incur regret $\sqrt{d^3\text{poly}(k)T}$.

3.2.1 RL with Simulator: Q-function is Quadratic and Bellman Complete

In this section we demonstrate how our results for non-concave bandits also apply to reinforcement learning. Let $\mathcal{T}_h$ be the Bellman operator applied to the Q-function Q_{h+1} defined as:

$$\mathcal{T}_h(Q_{h+1})(s,a) = r_h(s,a) + \mathbb{E}_{s' \sim \mathbb{P}(\cdot|s,a)}[\max_{a'} Q_{h+1}(s',a')].$$

Definition 3.11 (Bellman complete)**.** Given MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$, function class $\mathcal{F}_h : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}, h \in [H]$ is called Bellman complete if for all $h \in [H]$ and $Q_{h+1} \in \mathcal{F}_{h+1}$, $\mathcal{T}_h(Q_{h+1}) \in \mathcal{F}_h$.

Assumption 3.12 (Bellman complete for low-rank quadratic reward)**.** We assume the function class $\mathcal{F}_h = \{f_{\boldsymbol{M}} : f_{\boldsymbol{M}}(s,a) = \phi(s,a)^\top \boldsymbol{M} \phi(s,a), \operatorname{rank}(\boldsymbol{M}) \le k$, and $0 < \lambda_1(\boldsymbol{M})/\lambda_{\min}(\boldsymbol{M}) \le \widetilde{\kappa}.\}$ is a class of quadratic function and the MDP is Bellman complete. Here $f_{\boldsymbol{M}}(\phi(s,a)) = \sum_{j=1}^k \lambda_j (\boldsymbol{v}_j^\top \phi(s,a))^2$ when $\boldsymbol{M} = \sum_{j=1}^k \lambda_j \boldsymbol{v}_j \boldsymbol{v}_j^\top$. Write $Q_h^* = f_{\boldsymbol{M}_h^*} \in \mathcal{F}_h$.

Observation: When querying s_{h-1}, a_{h-1}, we observe $s_h' \sim \mathbb{P}(\cdot|s_{h-1}, a_{h-1})$ and reward $r_{h-1}(s_{h-1}, a_{h-1})$.
Oracle to recover parameter $\widehat{\boldsymbol{M}}$: Given $n \ge \widetilde{\Theta}(d^2k^2\widetilde{\kappa}^2(\boldsymbol{M})/\varepsilon^2)$, if one can play $\ge n$ samples $\boldsymbol{a}_i$ and observe $y_i \sim \boldsymbol{a}_i^\top \boldsymbol{M} \boldsymbol{a}_i + \eta_i, i \in [n]$ with 1-sub-gaussian and mean-zero noise η, we can recover $\widehat{\boldsymbol{M}} = \widehat{\boldsymbol{M}}(\{(\boldsymbol{a}_i, y_i)\})$ such that $\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_2 \le \varepsilon$. This oracle is implemented via our analysis from the bandit setting.

With the oracle, at time step H, we can estimate $\widehat{\boldsymbol{M}}_H$ that is ϵ/H close to $\boldsymbol{M}_H^*$ in spectral norm through noisy observations from the reward function with $\widetilde{O}(\widetilde{\kappa}^2 d^2 H^2/\epsilon^2)$ samples. Next, for each time step $h = H-1, H-1, \cdots, 1$, sample $s_i' \sim \mathbb{P}(\cdot|s,a)$, we define $\eta_i = \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s_i', a') - \mathbb{E}_{s' \sim \mathbb{P}(\cdot|s,a)} \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s', a')$. η_i is mean-zero and $O(1)$-sub-gaussian since it is bounded. Denote $\boldsymbol{M}_h$ as the matrix that satisfies $f_{\boldsymbol{M}_h} := \mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}}$, which is well-defined due to Bellman completeness. We estimate $\widehat{\boldsymbol{M}}_h$ from the noisy observations $y_i = r_h(s,a) + \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s_i', a') = \mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}} + \eta_i =: f_{\boldsymbol{M}_h} + \eta_i$. Therefore with the oracle, we can estimate $\widehat{\boldsymbol{M}}_h$ such that $\|\widehat{\boldsymbol{M}}_h - \boldsymbol{M}_h\|_2 \le \epsilon/H$ with $\Theta(\widetilde{\kappa}^2 d^2 k^2 H^2/\epsilon^2)$ bandits. More details are deferred to Algorithm 5 and the Appendix. We state the theorem on sample complexity of finding ϵ-optimal policy here:

Theorem 3.13. *Suppose $\mathcal{F}$ is Bellman complete associated with parameter $\widetilde{\kappa}$. With probability $1-\delta$, Algorithm 5 learns an ϵ-optimal policy π with $\widetilde{\Theta}(d^2k^2\widetilde{\kappa}^2H^3/\epsilon^2)$ samples.*

Existing approaches require $O(d^3H^2/\epsilon^2)$ trajectories, or equivalently $O(d^3H^3/\epsilon^2)$ samples, though they operate in the online RL setting [73, 23, 40], which is worse by a factor of d.

It is an open problem to attain d^2 sample complexity in the online RL setting. The quadratic Bellman complete setting can also be easily extended to any of the polynomial settings of Section 3.3.

3.3 Stochastic High-order Homogeneous Polynomial Reward

Next we move on to homogeneous high-order polynomials.

3.3.1 The symmetric setting

Let reward function be a p-th order stochastic polynomial function $f : \mathcal{A} \to \mathbb{R}$, where the action set $\mathcal{A} := \{B_1^d = \{\boldsymbol{a} \in \mathbb{R}^d, \|\boldsymbol{a}\| \le 1.\}$. $f(\boldsymbol{a}) = \boldsymbol{T}(\boldsymbol{a}^{\otimes p}), r_t = f(\boldsymbol{a}) + \eta_t$, where $\boldsymbol{T} = \sum_{j=1}^k \lambda_j \boldsymbol{v}_j^{\otimes p}$ is an orthogonally decomposable rank-k tensor. $\{\boldsymbol{v}_1, \cdots \boldsymbol{v}_k\}$ form an orthonormal basis. Optimal reward r^* satisfies $1 \ge r^* = \lambda_1 \ge |\lambda_2| \cdots \ge |\lambda_k|$. Noise $\eta_t \sim \mathcal{N}(0,1)$.

In this setting, the problem is fundamentally more challenging than quadratic reward functions. On the one hand, it has a higher noise-to-signal ratio with larger p. One can tell from the rank-1 setting where $\boldsymbol{T} = \lambda_1 \boldsymbol{v}_1^{\otimes p}$. For a randomly generated action $\boldsymbol{a}$ on the unit ball, $\mathbb{E}[\|\boldsymbol{a}^\top \boldsymbol{v}_1\|^2] = 1/d$. Therefore on average the signal strength is only $(\boldsymbol{a}^\top \boldsymbol{v}_1)^p \sim d^{-p/2}$, much smaller than the noise level 1. Intuitively this demonstrates why higher complexity is needed for high-order polynomials. On the other hand, it is also technically more challenging. Unlike the matrix case, the expected zeroth-order update is no longer equal to any tensor product. Therefore existing tensor decomposition arguments do not apply. Fortunately, we prove that zeroth-order optimization still pushes the iterated actions toward the optimal action with linear convergence, given a good initialization. We show the bandit optimization procedure in Algorithm 6 and present the result in Theorem 3.14:

Theorem 3.14 (Staged progress)**.** *For each stage s, with high probability $\mathcal{A}_s$ is not empty; and at least one action $\boldsymbol{a} \in \mathcal{A}_s$ satisfies: $\tan\theta(\boldsymbol{a}, \boldsymbol{a}^*) \leq \widetilde{\varepsilon}_s = 2^{-s}$.*

We defer the proofs together with formal statements to Appendix D.1.1 and D.1.

Remark 3.15 (Choice of step-size)**.** *Here we choose step-size $\zeta = 1/2p$. Note that $\mathbb{E}[\boldsymbol{y}] = \zeta(1-\zeta)^p\sigma^2\nabla_{\boldsymbol{a}} f(\boldsymbol{a})$. The scaling $(1-\zeta)^p \geq 1/\sqrt{e}$ ensures the signal to noise ratio not too small. The choice of weighted action is a delicate balance between making progress in optimization and controlling the noise to signal ratio.*

Corollary 3.16 (Regret bound for tensors)**.** *Algorithm 6 yields an regret of: $\mathfrak{R}(T) \leq \widetilde{O}\left(\sqrt{kd^pT}\right)$, with high probability.*

Corollary 3.17 (Regret bound with burn-in period)**.** *In Algorithm 6, we first set $\varepsilon = 1/p, n_s = \widetilde{\Theta}(d^p/\lambda_1^2)$. We can first estimate the action $\boldsymbol{a}$ such that $\boldsymbol{v}_1^\top \boldsymbol{a} \geq 1 - 1/p$ with $\widetilde{O}(kd^p/\lambda_1^2)$ samples. Next we change $\varepsilon = \widetilde{\Theta}(k^{1/4}d^{1/2}\lambda_1^{-1/2}T^{-1/4})$, and $n_s = \widetilde{\Theta}(d^2\varepsilon^{-2}\lambda_1^{-2})$ in Algorithm 6. This procedure suffices to find $\lambda_1\varepsilon^2$-optimal reward with $\widetilde{O}(kd^2/\lambda_1^2\varepsilon^2)$ samples in the candidate set with size at most $\widetilde{O}(k)$. Finally with the UCB algorithm altogether we have a regret bound of: $\widetilde{O}(\frac{kd^p}{\lambda_1} + \sqrt{kd^2T})$.*

In Section 3.3.2 we also demonstrate the necessity of the sample complexity in this burn-in period.

3.3.2 Lower Bounds for Stochastic Polynomial Bandits

In this section we show lower bound for stochastic polynomial bandits.

$$f(\boldsymbol{a}) = \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}) + \eta, \text{ where } \eta \sim N(0,1), \|\boldsymbol{a}\|_2 \leq 1 \text{ and } f(\boldsymbol{a}) \leq 1. \tag{1}$$

Theorem 3.18. *Define minimax regret as follow*

$$\mathfrak{R}(d,p,T) = \inf_{\pi} \sup_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)} \mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\left[T\max_{\boldsymbol{a}} \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}) - \sum_{t=1}^{T}\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}^{(t)})\right].$$

For all algorithms $\mathcal{A}$ that adaptively interact with bandit (Eq (1)) for T rounds, we have $\mathfrak{R}(d,p,T) \geq \Omega(\sqrt{d^pT}/p^p)$.

From the theorem, we can see even when the problem is rank-1, any algorithm incurs at least $\Omega(\sqrt{d^pT}/p^p)$ regret. This further implies algorithm requires sample complexity of $\Omega((d/p^2)^p/\epsilon^2)$ to attain ϵ-optimal reward. This means our regret upper bound obtained in Corollary 3.16 is optimal in terms of dependence on d. In Appendix G we also show a $\Omega(\sqrt{d^pT})$ lower bound for asymmetric actions setting, which our upper bound up to poly-logarithmic factors.

We note that with burn-in period, our algorithm also obtains a cumulative regret of $kd^p/\lambda_1 + \sqrt{kd^2T}$, as shown in Corollary 3.17. Here for a fixed λ_1 and very large T, this is better result than the previous upper bound. We note that there is no contradiction with the lower bound above, since this worst case is achieved with a specific relation between T and λ_1.

The burn-in period requires kd^p/λ_1^2 samples to get a constant of r^*. We want to investigate whether our dependence on $\lambda_1 \equiv r^*$ is optimal. Next we show a gap-dependent lower bound for finding an arm that is close to the optimal arm by a constant factor.

Theorem 3.19. *For all algorithms $\mathcal{A}$ that adaptively interact with bandit (Eq (1)) for T rounds and output a vector $\boldsymbol{a}^{(T)} \in \mathbb{R}^d$, it requires at least $T = \Omega(d^p/\|\boldsymbol{\theta}\|^{2p})$ rounds to find an arm $\boldsymbol{a}^{(T)} \in \mathbb{R}^d$ such that $\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}^{(T)}) \geq \frac{3}{4} \cdot \max_{\boldsymbol{a}} \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a})$.*

In the rank-1 setting $r^* = \|\boldsymbol{\theta}\|^p$ and the lower bound for achieving a constant approximation for the optimal reward is $\Omega(d^p/(r^*)^2)$. Therefore our burn-in sample complexity is also optimal in the dependence on d and r^*.

Acknowledgment

JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0303, the Sloan Research Fellowship, NSF CCF 2002272, and an ONR Young Investigator Award. QL is supported by NSF 2030859 and the Computing Research Association for the CIFellows Project. SK acknowledges funding from the NSF Award CCF-1703574 and the ONR award N00014-18-1-2247. The authors would like to thank Qian Yu for numerous conversations regarding the lower bound in Theorem 3.23. JDL would like to thank Yuxin Chen and Anru Zhang for several conversations on tensor power iteration, Simon S. Du and Yangyi Lu for explaining to him to the papers of [42, 54], and Max Simchowitz and Chao Gao for several conversations regarding adaptive lower bounds.

References

[1] Yasin Abbasi-Yadkori, Peter L Bartlett, and Alan Malek. Linear programming for large-scale Markov decision problems. *arXiv preprint arXiv:1402.6763*, 2014.

[2] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *NIPS*, volume 11, pages 2312–2320, 2011.

[3] Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2012.

[4] Alekh Agarwal, Dean P Foster, Daniel Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. *arXiv preprint arXiv:1107.1744*, 2011.

[5] Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep*, 2019.

[6] Zeyuan Allen-Zhu and Yuanzhi Li. First efficient convergence for streaming k-pca: a global, gap-free, and near-optimal rate. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 487–492. IEEE, 2017.

[7] Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels? *arXiv preprint arXiv:1905.10337*, 2019.

[8] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of machine learning research*, 15:2773–2832, 2014.

[9] Ery Arias-Castro, Emmanuel J Candes, and Mark A Davenport. On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory*, 59(1):472–481, 2012.

[10] Yu Bai and Jason D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. In *International Conference on Learning Representations*, 2020.

[11] Maria-Florina Balcan, Simon Shaolei Du, Yining Wang, and Adams Wei Yu. An improved gap-dependency analysis of the noisy power method. In *Conference on Learning Theory*, pages 284–309. PMLR, 2016.

[12] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real algebraic geometry*, volume 36. Springer Science & Business Media, 2013.

[13] Sebastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1), 2012.

[14] Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 72–85, 2017.

[15] T Tony Cai, Xiaodong Li, Zongming Ma, et al. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *The Annals of Statistics*, 44(5):2221–2251, 2016.

[16] Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.

[17] Minshuo Chen, Yu Bai, Jason D Lee, Tuo Zhao, Huan Wang, Caiming Xiong, and Richard Socher. Towards understanding hierarchical learning: Benefits of neural representations. *Neural Information Processing Systems (NeurIPS)*, 2020.

[18] Sitan Chen and Raghu Meka. Learning polynomials in few relevant dimensions. In *Conference on Learning Theory*, pages 1161–1227. PMLR, 2020.

[19] Alex Damian, Tengyu Ma, and Jason Lee. Label noise sgd provably prefers flat global minimizers. *arXiv preprint arXiv:2106.06530*, 2021.

[20] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. *The 21st Annual Conference on Learning Theory*, 2008.

[21] Kefan Dong, Jiaqi Yang, and Tengyu Ma. Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature. *arXiv preprint arXiv:2102.04168*, 2021.

[22] Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.

[23] Simon S Du, Sham M Kakade, Jason D Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in RL. *arXiv preprint arXiv:2103.10897*, 2021.

[24] Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685, 2019.

[25] Cong Fang, Jason D Lee, Pengkun Yang, and Tong Zhang. Modeling from features: a mean-field framework for over-parameterized deep neural networks. *arXiv preprint arXiv:2007.01452*, 2020.

[26] Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*, 2004.

[27] Ruiqi Gao, Tianle Cai, Haochuan Li, Liwei Wang, Cho-Jui Hsieh, and Jason D Lee. Convergence of adversarial training in overparametrized networks. *Neural Information Processing Systems (NeurIPS)*, 2019.

[28] Dan Garber, Elad Hazan, and Tengyu Ma. Online learning of eigenvectors. In *International Conference on Machine Learning*, pages 560–568. PMLR, 2015.

[29] Rong Ge, Jason D Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. *International Conference on Learning Representations (ICLR)*, 2018.

[30] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *arXiv preprint arXiv:1904.12191*, 2019.

[31] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021.

[32] Aditya Gopalan, Odalric-Ambrym Maillard, and Mohammadi Zaki. Low-rank bandits with latent mixtures. *arXiv preprint arXiv:1609.01508*, 2016.

[33] Botao Hao, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Online sparse reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 316–324. PMLR, 2021.

[34] Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. *arXiv preprint arXiv:2011.04020*, 2020.

[35] Botao Hao, Jie Zhou, Zheng Wen, and Will Wei Sun. Low-rank tensor bandits. *arXiv preprint arXiv:2007.15788*, 2020.

[36] Jeff Z. HaoChen, Colin Wei, Jason D. Lee, and Tengyu Ma. Shape matters: Understanding the implicit bias of the noise covariance. *arXiv preprint arXiv:2006.08680*, 2020.

[37] Moritz Hardt and Eric Price. The noisy power method: A meta algorithm with applications. *Advances in neural information processing systems*, 27:2861–2869, 2014.

[38] Elad Hazan and Yuanzhi Li. An optimal algorithm for bandit convex optimization. *arXiv preprint arXiv:1603.04350*, 2016.

[39] Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In *NeurIPS*, 2018.

[40] Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *arXiv preprint arXiv:2102.00815*, 2021.

[41] Nicholas Johnson, Vidyashankar Sivakumar, and Arindam Banerjee. Structured stochastic linear bandits. *arXiv preprint arXiv:1606.05693*, 2016.

[42] Kwang-Sung Jun, Rebecca Willett, Stephen Wright, and Robert Nowak. Bilinear bandits with low-rank structure. In *International Conference on Machine Learning*, pages 3163–3172. PMLR, 2019.

[43] Sumeet Katariya, Branislav Kveton, Csaba Szepesvari, Claire Vernade, and Zheng Wen. Stochastic rank-1 bandits. In *Artificial Intelligence and Statistics*, pages 392–401. PMLR, 2017.

[44] Joe Kileel, Matthew Trager, and Joan Bruna. On the expressive power of deep polynomial neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

[45] Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. *Advances in Neural Information Processing Systems*, 17:697–704, 2004.

[46] Wojciech Kotłowski and Gergely Neu. Bandit principal component analysis. In *Conference On Learning Theory*, pages 1994–2024. PMLR, 2019.

[47] Sahin Lale, Kamyar Azizzadenesheli, Anima Anandkumar, and Babak Hassibi. Stochastic linear bandits with hidden low rank structure. *arXiv preprint arXiv:1901.09490*, 2019.

[48] Tor Lattimore. Improved regret for zeroth-order adversarial bandit convex optimisation, 2020.

[49] Tor Lattimore and Botao Hao. Bandit phase retrieval. *arXiv preprint arXiv:2106.01660*, 2021.

[50] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

[51] Guillaume Lecué and Shahar Mendelson. Minimax rate of convergence and the performance of erm in phase recovery. *arXiv preprint arXiv:1311.5024*, 2013.

[52] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.

[53] Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.

[54] Yangyi Lu, Amirhossein Meisami, and Ambuj Tewari. Low-rank generalized linear bandit problems. In *International Conference on Artificial Intelligence and Statistics*, pages 460–468. PMLR, 2021.

[55] James S. Milne. Algebraic geometry (v6.02), 2017. Available at www.jmilne.org/math/.

[56] Edward Moroshko, Suriya Gunasekar, Blake Woodworth, Jason D Lee, Nathan Srebro, and Daniel Soudry. Implicit bias in deep linear classification: Initialization scale vs training accuracy. *Neural Information Processing Systems (NeurIPS)*, 2020.

[57] Cameron Musco and Christopher Musco. Randomized block krylov methods for stronger and faster approximate singular value decomposition. *arXiv preprint arXiv:1504.05477*, 2015.

[58] Mor Shpigel Nacson, Suriya Gunasekar, Jason Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In *International Conference on Machine Learning*, pages 4683–4692. PMLR, 2019.

[59] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling, 2018.

[60] Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.

[61] Dan Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pages 2256–2264, 2013.

[62] Maziar Sanjabi, Sina Baharlouei, Meisam Razaviyayn, and Jason D Lee. When does non-orthogonal tensor decomposition have no spurious local minima? *arXiv preprint arXiv:1911.09815*, 2019.

[63] Igor R Shafarevich. *Basic Algebraic Geometry 1: Varieties in Projective Space*. Springer Science & Business Media, 2013.

[64] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

[65] Michal Valko, Nathaniel Korda, Rémi Munos, Ilias Flaounas, and Nelo Cristianini. Finite-time analysis of kernelised contextual bandits. *arXiv preprint arXiv:1309.6869*, 2013.

[66] Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33:6123–6135, 2020.

[67] Xiang Wang, Chenwei Wu, Jason D Lee, Tengyu Ma, and Rong Ge. Beyond lazy training for over-parameterized tensor decomposition. *Neural Information Processing Systems (NeurIPS)*, 2020.

[68] Yang Wang and Zhiqiang Xu. Generalized phase retrieval: measurement number, matrix recovery and beyond. *Applied and Computational Harmonic Analysis*, 47(2):423–446, 2019.

[69] Yining Wang and Animashree Anandkumar. Online and differentially-private tensor decomposition. *arXiv preprint arXiv:1606.06237*, 2016.

[70] Colin Wei, Jason D Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9709–9721, 2019.

[71] Blake Woodworth, Suriya Genesekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Kernel and deep regimes in overparametrized models. In *Conference on Learning Theory (COLT)*, 2019.

[72] Pan Xu, Zheng Wen, Handong Zhao, and Quanquan Gu. Neural contextual bandits with deep representation and shallow exploration. *arXiv preprint arXiv:2012.01780*, 2020.

[73] Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent Bellman error. In *International Conference on Machine Learning*, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Our abstract and introduction accurately reflect our main contribution, which is also summarized in Table 1.
   (b) Did you describe the limitations of your work? [Yes] In the Preliminary section we clearly state the settings we considered and when our results apply.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Our work is theoretical and generally will not have negative social impacts.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] The assumptions are clearly included in every theorem or lemma.
   (b) Did you include complete proofs of all theoretical results? [Yes] All the proofs can be found in the appendix
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Omitted Related Work

Neural Kernel Bandits. [65] initiated the study of kernelized linear bandits, showing regret dependent on the information gain. [72] specialized this to the Neural Tangent Kernel (NTK) [53, 24, 39, 27, 10], where the algorithm utilizes gradient descent but remains close to initialization and thus remains a kernel class. Furthermore NTK methods require d^p samples to express a degree p polynomial in d dimensions [31], similar to eluder dimension of polynomials, and so lack the inductive biases necessary for real-world applications of decision-making problems [59].

Concave Bandits. There has been a rich line of work on concave bandits starting with [26, 45]. [4] attained the first $\sqrt{T}$ regret algorithm for concave bandits though with a large $\mathrm{poly}(d)$ dependence. In the adversarial setting, a line of work [38, 14, 48] have attained polynomial-time algorithms with $\sqrt{T}$ regret with increasingly improved dimension dependence. The sharp dimension dependence remains unknown.

Noiseless Bandits. In the noiseless setting, there is some investigation in phase retrieval borrowing the tools from algebraic geometry (see e.g. [68]). In this paper, we will study the bandit problem with more general reward functions: neural nets with polynomial activation (structured polynomials) including phase retrieval. [44] study similar structured polynomials, also using tools from algebraic geometry, but they only study the expressivity of those polynomials and do not consider the learning problems. [21] study noiseless bandits with bounded Sequential Rademacher Complexity, but focus on attaining local optimality.

Concurrent work. [49] address the phase retrieval bandit problem which is equivalent to a symmetric rank 1 variant of the bilinear bandit of [42] and attain $\widetilde{O}(\sqrt{d^2 T})$ regret. Our work in Section 3.1 specialized to the rank 1 case attains the same regret.

Matrix/Tensor Power Method. Our analysis stems from noisy power methods for matrix/tensor decomposition problems. Robust power method, subspace iteration, and tensor decomposition that tolerate noise first appeared in [37, 8]. Follow-up work attained the optimal rate for both gap-dependence and gap-free settings for matrix decomposition [57, 6]. An improvement on the problem dimension for tensor power method is established in [69]. [62] considers the convergence of tensor power method in the non-orthogonal case.

B Additional Preliminaries

In this section we show that adapting the eluder UCB algorithms from [61] would yield the sample complexity in Theorem 2.1. Especially we give the rates in Table 1 for our stochastic settings.

Algorithm 2 Eluder UCB

1: **Input:** Function class $\mathcal{F}$, failure probability δ, parameters α, N, C.
2: <u>**Initialization:**</u> $\mathcal{F}_0 \leftarrow \mathcal{F}$.
3: **for** t from 1 to T **do**
4: <u>**Select Action:**</u>
5: $\boldsymbol{a}_t \in \arg\max_{\boldsymbol{a}\in\mathcal{A}} \sup_{f_{\boldsymbol{\theta}}\in\mathcal{F}_{t-1}} f_{\boldsymbol{\theta}}(\boldsymbol{a})$
6: Play action $\boldsymbol{a}_t$ and observe reward r_t
7: <u>**Update Statistics:**</u>
8: $\widehat{\boldsymbol{\theta}}_t \in \arg\min_{\boldsymbol{\theta}} \sum_{s=1}^{t}(f_{\boldsymbol{\theta}}(\boldsymbol{a}_s) - r_s)^2$
9: $\beta_t \leftarrow 8\log(N/\delta) + 2\alpha t(8C + \sqrt{8\ln(4t^2/\delta)})$
10: $\mathcal{F}_t \leftarrow \{f_{\boldsymbol{\theta}} : \sum_{s=1}^{t}(f_{\boldsymbol{\theta}} - f_{\widehat{\boldsymbol{\theta}}_t})^2(\boldsymbol{a}_s) \leq \beta_t\}$

The algorithm [61] consider Algorithm 2 for the stochastic generalized linear bandit problem. Assume that $\boldsymbol{\theta}^*$ is the true parameter of the reward model. The reward is $r_t = f_{\boldsymbol{\theta}^*}(\boldsymbol{a}_t) + \eta_t$ for

$f_{\boldsymbol{\theta}^*} \in \mathcal{F}$. Let N be the α-covering-number (under $\|\cdot\|_\infty$) of $\mathcal{F}$, d_E be the α-eluder-dimension of $\mathcal{F}$ (see Definition 3,4 in [61]). Let $C = \sup_{f\in\mathcal{F}, a\in\mathcal{A}} |f(a)|$. We set $\alpha = \frac{1}{T^2}$ in the algorithm.

The regret analysis Choosing $\alpha = 1/T^2$, proposition 4 in [61] state that with probability $1-\delta$, for some universal constant C, the total regret $\mathfrak{R}(T) \le \frac{1}{T} + C\min\{d_E, T\} + 4\sqrt{d_E\beta_T T} \le 1 + C\sqrt{d_E T} + 4\sqrt{d_E\beta_T T} = O(\sqrt{d_E(1+\beta_T)T})$. In our settings with $\alpha = 1/T^2$, $\beta_T = 8\log(N/\delta) + 2(8C + \sqrt{8\ln(4T^2/\delta)})/T = O(\log(N/\delta))$ where $\log(N) = \Omega(1)$ for our action sets, and thus

$$\mathfrak{R}(T) = \widetilde{O}(\sqrt{d_E T \log N}).$$

Applications in our settings We show that in our settings Theorem 2.1 will obtain the rates listed in Table 1.

The covering numbers

Lemma B.1. *The log-covering-number (of radius α with $\alpha \ll 1$, under $\|\cdot\|_\infty$) of the function classes are: $\log N(\mathcal{F}_{SYM}) = O(dk\log\frac{k}{\alpha})$, $\log N(\mathcal{F}_{ASYM}) = O(dk\log\frac{k}{\alpha})$, $\log N(\mathcal{F}_{EV}) = O(dk\log\frac{k}{\alpha})$, and $\log N(\mathcal{F}_{LR}) = O(dk\log\frac{k}{\alpha})$.*

Proof. Let S_ξ^d denote a minimal ξ-covering of $\mathbb{S}^{d-1}$ (under $\|\cdot\|_2$) for $0 < \xi < \frac{1}{10}$, and $|S_\xi^d| = O(d\log 1/\xi)$ (see for example [61]). Then we can construct the coverings in our settings from S_ξ^d:

- $\mathcal{F}_{\text{SYM}}$: let $\xi = \frac{\alpha}{kp}$, and for k copies of S_ξ^d, we can construct a covering of $\mathcal{F}_{\text{SYM}}$ with size $|S_\xi^d|^k$. Specifically, let the covering be $S_{\text{SYM}} = \{g(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j(\boldsymbol{u}_j^\top \boldsymbol{a})^p : (\boldsymbol{u}_1, \boldsymbol{u}_2, \cdots, \boldsymbol{u}_k) \in S_\xi^d \times S_\xi^d \times \cdots \times S_\xi^d\}$, then for each $f(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^p \in \mathcal{F}_{\text{SYM}}$, as we can find $\boldsymbol{u}_j \in S_\xi^d$ that $\|\boldsymbol{u}_j - \boldsymbol{v}_j\|_2 \le \xi$,
$$\sup_{\boldsymbol{a}}[f(\boldsymbol{a}) - g(\boldsymbol{a})] \le \sup_{\boldsymbol{a}}[\sum_{j=1}^k |\lambda_j||\boldsymbol{u}_j^\top \boldsymbol{a} - \boldsymbol{v}_j^\top \boldsymbol{a}||\sum_{q=0}^{p-1}(\boldsymbol{u}_j^\top \boldsymbol{a})^q(\boldsymbol{v}_j^\top \boldsymbol{a})^{p-q-1}|] \le pk\xi = \alpha;$$
- $\mathcal{F}_{\text{ASYM}}$: let $\xi = \frac{\alpha}{kp}$, and for kp copies of S_ξ^d, let the covering be $S_{\text{ASYM}} = \{g(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j \prod_{q=1}^p (\boldsymbol{u}_j(q)^\top \boldsymbol{a}(q)) : (\boldsymbol{u}_1(1), \boldsymbol{u}_1(2), \cdots, \boldsymbol{u}_1(p), \boldsymbol{u}_2(1), \cdots, \boldsymbol{u}_k(p)) \in S_\xi^d \times S_\xi^d \times \cdots \times S_\xi^d\}$ with size $|S_\xi^d|^{kp}$. Then for each $f(\boldsymbol{a}) = \sum_{j=1}^k \lambda_j \prod_{q=1}^p (\boldsymbol{v}_j(q)^\top \boldsymbol{a}(q)) \in \mathcal{F}_{\text{ASYM}}$, as we can find $\boldsymbol{u}_j(q) \in S_\xi^d$ that $\|\boldsymbol{u}_j(q) - \boldsymbol{v}_j(q)\|_2 \le \xi$,
$$\begin{aligned}\sup_{\boldsymbol{a}}[f(\boldsymbol{a}) - g(\boldsymbol{a})] &\le \sup_{\boldsymbol{a}}[\textstyle\sum_{j=1}^k |\lambda_j| \sum_{q=1}^p |\boldsymbol{u}_j(q)^\top \boldsymbol{a} - \boldsymbol{v}_j(q)^\top \boldsymbol{a}| \cdot \\ &\quad |\textstyle\prod_{r<q}(\boldsymbol{u}_j(r)^\top \boldsymbol{a}) \prod_{r>q}(\boldsymbol{v}_j(r)^\top \boldsymbol{a})|] \\ &\le pk\xi = \alpha;\end{aligned}$$
- $\mathcal{F}_{\text{EV}}$: the construction follows that of $\mathcal{F}_{\text{SYM}}$ by taking $p = 2$;
- $\mathcal{F}_{\text{LR}}$: taking the construction of $\mathcal{F}_{\text{SYM}}$ with $p = 2$ and $\xi = \frac{\alpha}{2k}$, for $\boldsymbol{N} = \sum_{j=1}^k \lambda_j \boldsymbol{u}_j \boldsymbol{u}_j^\top$ and $\boldsymbol{M} = \sum_{j=1}^k \lambda_j \boldsymbol{v}_j \boldsymbol{v}_j^\top$ with $\|\boldsymbol{u}_j - \boldsymbol{v}_j\|_2 \le \xi$, we know $\|\boldsymbol{N} - \boldsymbol{M}\|_{\text{F}} \le \left\|\boldsymbol{N} - \sum_{j=1}^k \lambda_j \boldsymbol{u}_j \boldsymbol{v}_j^\top\right\|_{\text{F}} + \left\|\sum_{j=1}^k \lambda_j \boldsymbol{u}_j \boldsymbol{v}_j^\top - \boldsymbol{M}\right\|_{\text{F}} \le \sum_{j=1}^k 2|\lambda_j|\xi \le \alpha$. Then $\sup_{\boldsymbol{A}}[f_{\boldsymbol{M}}(\boldsymbol{A}) - f_{\boldsymbol{N}}(\boldsymbol{A})] \le \sup_{\boldsymbol{A}} \|\boldsymbol{M} - \boldsymbol{N}\|_{\text{F}} \cdot \|\boldsymbol{A}\|_{\text{F}} \le \alpha$.

Then we can bound the covering numbers in Theorem 2.1. Notice that in the settings the log-covering numbers are only different by constant factors. □

The eluder dimensions

Lemma B.2. *The ϵ-eluder-dimension ($\epsilon < 1$) d_E of the function classes are: $d_E(\mathcal{F}_{SYM}) = \widetilde{\Theta}(d^p)$ (for $k \ge p$), $d_E(\mathcal{F}_{ASYM}) = \widetilde{\Theta}(d^p)$, $d_E(\mathcal{F}_{EV}) = \widetilde{\Theta}(d^2)$, and $d_E(\mathcal{F}_{LR}) = \widetilde{\Theta}(d^2)$. In the settings WLOG we assume the top eigenvalue is $r^* = \lambda_1 = 1$ as we are mostly interested in the cases where $r^* > \epsilon$.*

Proof. The upper bounds for the eluder dimension can be given by the linear argument. [61] show that the d-dimension linear model $\{f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \boldsymbol{\theta}^\top \boldsymbol{a}\}$ has ϵ-eluder-dimension $O(d \log \frac{1}{\epsilon})$. In all of these settings, we can find feature maps ϕ and ψ so that $\mathcal{F} = \{f_{\boldsymbol{\theta}}(\boldsymbol{a}), f_{\boldsymbol{\theta}}(\boldsymbol{a}) = \phi(\boldsymbol{\theta})^\top \psi(\boldsymbol{a}), \|\phi(\boldsymbol{\theta})\|_2 \leq k, \|\psi(\boldsymbol{a})\|_2 \leq k\}$. Then the eluder dimensions will be bounded by the corresponding linear dimension as an original ϵ-independent sequence $\{\boldsymbol{a}_i\}$ will induce an ϵ-independent sequence $\{\psi(\boldsymbol{a}_i)\}$ in the linear model. Therefore for matrices ($\mathcal{F}_{\text{LR}}$ and $\mathcal{F}_{\text{EV}}$) the eluder dimension is $O(d^2 \log \frac{k}{\epsilon})$ and for the tensors ($\mathcal{F}_{\text{SYM}}$ and $\mathcal{F}_{\text{ASYM}}$) it is $O(d^p \log \frac{k}{\epsilon})$.

Then we consider the lower bounds. We provide the following example of $O(1)$-independent sequences to bound the eluder dimension in our settings up to a $\log$ factor.

- $\mathcal{F}_{\text{SYM}}$: the sequence is $\{\boldsymbol{a}_i = (\boldsymbol{e}_{i_1}, \boldsymbol{e}_{i_2}, \cdots, \boldsymbol{e}_{i_p}) : i = (i_1, i_2, \cdots, i_p) \in [d]^p\}$. For $f_j(\boldsymbol{a}) = \prod_{q=1}^p \boldsymbol{e}_{j_q}^\top \boldsymbol{a}(q)$, $f_j(\boldsymbol{a}_i)$ is only 1 when $i = j$ and 0 otherwise. Then each $\boldsymbol{a}_i$ is 1-independent to the predecessors on f_i and zero, and thus the eluder dimension is lower bounded by d^p.
- $\mathcal{F}_{\text{ASYM}}$: for $p \leq d$ and $k \geq p$, the sequence is $\{\boldsymbol{a}_i = \frac{1}{\sqrt{p}}(\boldsymbol{e}_{i_1} + \boldsymbol{e}_{i_2} + \cdots + \boldsymbol{e}_{i_p}) : i = (i_1, i_2, \cdots, i_p) \in [d]^p, i_1 < i_2 < \cdots < i_p\}$. There are tensors f_j and g_j of CP-rank k that $(f_j - g_j)(\boldsymbol{a}) = \prod_{q=1}^p (\boldsymbol{e}_{j_q}^\top \boldsymbol{a})$ where $j_1 < j_2 < \cdots < j_p$, $(f_j - g_j)(\boldsymbol{a}_i)$ is only 1 when $i = j$ and 0 otherwise. Then each $\boldsymbol{a}_i$ is 1-independent to the predecessors on f_i and g_i, and thus the eluder dimension is lower bounded by $\binom{d}{p}$.
- $\mathcal{F}_{\text{EV}}$: the sequence is $\{\boldsymbol{a}_i = \frac{1}{\sqrt{2}}(\boldsymbol{e}_{i_1} + \boldsymbol{e}_{i_2}) : i = (i_1, i_2) \in [d]^2, i_1 \leq i_2\}$. For $f_j(\boldsymbol{a}) = \frac{1}{2}\boldsymbol{a}^\top (\boldsymbol{e}_{j_1} + \boldsymbol{e}_{j_2})(\boldsymbol{e}_{j_1} + \boldsymbol{e}_{j_2})^\top \boldsymbol{a}$ and $g_j(\boldsymbol{a}) = \frac{1}{2}\boldsymbol{a}^\top (\boldsymbol{e}_{j_1} - \boldsymbol{e}_{j_2})(\boldsymbol{e}_{j_1} - \boldsymbol{e}_{j_2})^\top \boldsymbol{a}$ with $j_1 \leq j_2$, $(f_j - g_j)(\boldsymbol{a}_i)$ is only 1 when $i = j$ and 0 otherwise. Then each $\boldsymbol{a}_i$ is 1-independent to the predecessors on f_i and g_i, and thus the eluder dimension is lower bounded by $\binom{d}{2}$.
- $\mathcal{F}_{\text{LR}}$: the sequence is $\{\boldsymbol{A}_i = \frac{1}{2}\boldsymbol{e}_{i_1}\boldsymbol{e}_{i_2}^T + \boldsymbol{e}_{i_1}\boldsymbol{e}_{i_2}^T : i = (i_1, i_2) \in [d]^2, i_1 \leq i_2\}$. For $f_j(\boldsymbol{A}) = \langle \frac{1}{2}(\boldsymbol{e}_{j_1}\boldsymbol{e}_{j_2}^T + \boldsymbol{e}_{j_2}\boldsymbol{e}_{j_1}^T), \boldsymbol{A}\rangle$ with $j_1 \leq j_2$, $f_j(\boldsymbol{A}_i)$ is only 1 when $i = j$ and 0 otherwise. Then each $\boldsymbol{A}_i$ is 1-independent to the predecessors on f_i and zero, and thus the eluder dimension is lower bounded by $\binom{d}{2}$.

$\square$

Then we are all set for the results in the first line of 1. Notice that when we choose $\alpha = O(1/T^2)$ and $\epsilon = O(1/T^2)$ in our analysis of Algorithm 2, the regret upper bound would only expand by $\log(T)$ factors.

C Omitted Proofs for Quadratic Reward

In this section we include all the omitted proof of the theorems presented in the main paper.

C.1 Omitted Proofs of Main Results for Stochastic Bandit Eigenvector Problem

Proof of Theorem 3.3. Notice in Algorithm 1, for each iterate $\boldsymbol{a}$, its next iterate $\boldsymbol{y}$ satisfies

$$\begin{aligned}\boldsymbol{y} =& \frac{1}{n_s}\sum_{i=1}^{n_s}(\boldsymbol{a}/2 + \boldsymbol{z}_i/2)^\top \boldsymbol{M}(\boldsymbol{a}/2 + \boldsymbol{z}_i/2)\boldsymbol{z}_i + \eta_i \boldsymbol{z}_i \\ =& \frac{m_s}{n_s}\sum_{i=1}^{n_s}(\frac{1}{4}\boldsymbol{a}^\top \boldsymbol{M}\boldsymbol{a} + \frac{1}{2}\boldsymbol{a}^\top \boldsymbol{M}\boldsymbol{z}_i + \eta_i)\boldsymbol{z}_i.\end{aligned}$$

Therefore $\mathbb{E}[\boldsymbol{y}] = \frac{1}{2}\boldsymbol{M}\boldsymbol{a}$. We can write $2\boldsymbol{y} = \boldsymbol{M}\boldsymbol{a} + \boldsymbol{g}$ where $\boldsymbol{g} := \frac{m_s}{n_s}\sum_{i=1}^{n_s}(\frac{1}{2}\boldsymbol{a}^\top \boldsymbol{M}\boldsymbol{a} + 2\eta_i)\boldsymbol{z}_i$. With Claim D.12 and Claim D.11 we get that $\|\boldsymbol{g}\| \leq C\sqrt{\frac{m_s \log^2(n/\delta)\log(d/\delta)d}{n_s}}$. Therefore with our choice of $n_s \geq \widetilde{\Theta}(\frac{d^2}{\varepsilon_s^2(\lambda_1 - |\lambda_2|)^2})$ we guarantee $\|\boldsymbol{g}\| \leq \varepsilon_s(\lambda_1 - |\lambda_2|)$. Therefore it satisfies the requirements for noisy power method, and by applying Corollary C.4, we have with $L =$

$O(\kappa \log(d/\varepsilon))$ iterations we will be able to find $\|\widehat{\boldsymbol{a}}-\boldsymbol{a}^*\| \leq \varepsilon$. By setting $\delta < 0.1/L$ in the algorithm we can guarantee the whole process succeed with high probability. Altogether it is sufficient to take $Ln_s = \widetilde{O}(\kappa d^2/(\varepsilon\Delta)^2)$ actions to get an ε-optimal arm.

Finally to get the cumulative regret bound, we apply Claim D.7 with $A = \frac{d^2\kappa}{\Delta^2}$ and $a = 2$. Therefore we set $\varepsilon = A^{1/4}T^{-1/4} = \frac{d^{1/2}\kappa^{1/4}}{\Delta^{1/2}T^{1/4}}$ and get:

$$\mathrm{Reg}(T) \lesssim T^{1/2}A^{1/2}r^* = \sqrt{\frac{d^2\kappa}{\Delta^2}T}r^* = \sqrt{d^2\kappa^3 T}.$$

□

Corollary C.1 (Formal statement for Corollary 3.6)**.** *In Algorithm 1, by setting $\alpha = 1 - \varepsilon^2/2$, one can get ε-optimal reward with a total of $\widetilde{O}(d^2\lambda_1^2/\varepsilon^4)$ total samples to get $\boldsymbol{a}$ such that $r^* - f(\boldsymbol{a}) \leq \varepsilon$. Therefore one can get an accumulative regret of $\widetilde{O}(\lambda_1^{3/5}d^{2/5}T^{4/5})$.*

Proof of Lemma 3.6. In order to find an arm with $\lambda_1\varepsilon^2$-optimal reward, one will want to recover an arm that is $\varepsilon/2$-close (meaning to find an $\boldsymbol{a}$ such that $\tan\theta(\boldsymbol{V}_l, \boldsymbol{a}) \leq \varepsilon/2$) to the top eigenspace $\mathrm{span}(\boldsymbol{v}_1, \cdots \boldsymbol{v}_l)$, where l satisfies $\lambda_l \geq \lambda_1 - \widetilde{\varepsilon}$ and $\lambda_{l+1} \leq \lambda_1 - \widetilde{\varepsilon}$. Here we set $\widetilde{\varepsilon} := \lambda_1\varepsilon^2/2$. We first show 1) this is sufficient to get an $\lambda_1\varepsilon$-optimal reward, and next show 2) how to set parameter to achieve this.

To get 1), we write $\boldsymbol{V}_l = [\boldsymbol{v}_1, \cdots \boldsymbol{v}_l] \in \mathbb{R}^{d\times l}$ and $\boldsymbol{V}_l^\perp = [\boldsymbol{v}_{l+1}, \cdots \boldsymbol{v}_k]$. When $\tan\theta(\boldsymbol{V}_l, \boldsymbol{a}_T) = \|\boldsymbol{V}^\perp\boldsymbol{a}\|/\|\boldsymbol{V}\boldsymbol{a}\| \leq \varepsilon/2$, from the proof of Claim D.6, we get $r^* - f(\boldsymbol{a}) \leq \min\{\lambda_1, \lambda_1 2(\varepsilon/2)^2 + \widetilde{\varepsilon}\} = \lambda_1\varepsilon^2$.

Now to get 2), we note that in each iteration we try to conduct the power iteration to find an action $\tan\theta(\boldsymbol{V}_l, \widehat{\boldsymbol{a}}) \leq \varepsilon/2$ and with eigengap $\geq \widetilde{\varepsilon} := \lambda_1\varepsilon^2/2$. Therefore it is sufficient to let $\|\boldsymbol{g}\| \leq 0.1\widetilde{\varepsilon}\varepsilon$ and $|\boldsymbol{v}_1^\top\boldsymbol{g}| \leq 0.1\widetilde{\epsilon}\frac{1}{\sqrt{d}}$, and thus $n_s \geq \widetilde{\Theta}(\frac{d^2}{\varepsilon^2\widetilde{\varepsilon}^2}) \leq \widetilde{\Theta}(d^2/\lambda_1^2\varepsilon^6)$. Together we need $\lambda_1/\widetilde{\epsilon}\log(2d/\varepsilon)n_s = \widetilde{\Theta}(d^2/\lambda_1^2\varepsilon^8)$ samples to get an $\lambda_1\varepsilon^2$-optimal reward. Namely we get $\widetilde{\varepsilon}$-optimal reward with $\widetilde{O}(d^2\lambda_1^2/\widetilde{\varepsilon}^4)$ samples.

Finally by applying Claim D.7 we get:

$$\mathfrak{R}(T) \lesssim (d^2\lambda_1^2)^{\frac{1}{5}}T^{\frac{4}{5}}\lambda_1^{\frac{1}{5}} \leq \widetilde{O}(\lambda_1^{3/5}d^{2/5}T^{4/5}).$$

□

Algorithm 3 Gap-free Subspace Iteration for Bilinear Bandit

1: **Input:** Quadratic reward $f : \mathcal{X} \to \mathbb{R}$ generating noisy reward, failure probability δ, error ε.
2: **Initialization:** Set $k' = 2k$. Initial candidate matrix $\boldsymbol{X}_0 \in \mathbb{R}^{d\times k'}$, $\boldsymbol{X}_0(j) \in \mathbb{R}^d, j = 1, 2, \cdots k'$ is the j-th column of X_0 and are i.i.d sampled on the unit sphere $\mathbb{S}^{d-1}$ uniformly. Sample variance m, # sample per iteration n, total iteration L.
3: **for** Iteration l from 1 to L **do**
4:     **for** s from 1 to k' **do**
5:         **Noisy subspace iteration:**
6:         Sample $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/mI_d), i = 1, 2, \cdots n_s$.
7:         Calculate tentative rank-1 arms $\widetilde{\boldsymbol{a}}_i = \frac{1}{2}(\boldsymbol{X}_{l-1}(s) + \boldsymbol{z}_i)$.
8:         Conduct estimation $\boldsymbol{Y}_l(s) \leftarrow 4m/n\sum_{i=1}^n (f(\widetilde{\boldsymbol{a}}_i) + \eta_i)\boldsymbol{z}_i$. ($\boldsymbol{Y}_l \in \mathbb{R}^{d\times k'}$)
9:     Let $\boldsymbol{Y}_l = \boldsymbol{X}_l\boldsymbol{R}_l$ be a QR-factorization of $\boldsymbol{Y}_l$
10:     Update target arm $\boldsymbol{a}_l \leftarrow \arg\max_{\|\boldsymbol{a}\|=1} \boldsymbol{a}^\top\boldsymbol{Y}_l\boldsymbol{X}_{l-1}^\top\boldsymbol{a}$.
11: **Output:** $\boldsymbol{a}_L$.

Theorem C.2 (Formal statement of Theorem 3.7)**.** *In Algorithm 3, if we set $n = \widetilde{\Theta}(\frac{d^2\lambda_1^2}{\varepsilon^2\lambda_k^2}), m = d\log(n/\delta), L = \Theta(\log(d/\varepsilon)), \delta = 0.1/L$, we will be able to identify an action $\widehat{\boldsymbol{a}}$ that yield at most ε-regret with probability 0.9. Therefore by applying the standard PAC to regret conversion* *as discussed in Claim D.7 we get a cumulative regret of* $\widetilde{O}(\lambda_1^{1/3}k^{1/3}(\widetilde{\kappa}dT)^{2/3})$ *for large enough* T*, where* $\widetilde{\kappa} = \lambda_1/|\lambda_k|$.

On the other hand, we set $n = \widetilde{\Theta}(\frac{d^2k^2}{\varepsilon^2})$ *and keep the other parameters. If we play Algorithm 3 k times by setting* $k' = 2, 4, 6, \cdots 2k$ *and select the best output among them, we can get a gap-free cumulative regret of* $\widetilde{O}(\lambda_1^{1/3}k^{4/3}(dT)^{2/3})$ *for large enough* T *with high probability.*

Proof of Theorem 3.7. First we show the first setting identify an ε-optimal reward with $\widetilde{O}(\widetilde{\kappa}^2d^2k\epsilon^{-2})$ samples.

Similarly as Theorem 3.8, when setting $n \geq \widetilde{\Theta}(d^2/(\sigma_k^2\widetilde{\epsilon}^2))$, we can find $\boldsymbol{X}_L$ that satisfies $\|(\boldsymbol{X}_L\boldsymbol{X}_L^\top - I)\boldsymbol{U}\| \leq \widetilde{\epsilon}$, and therefore we recover an $\boldsymbol{Y}_L = \boldsymbol{M}\boldsymbol{X}_{L-1} + \boldsymbol{G}_L$ with $\|\boldsymbol{G}_L\| \leq \sigma_k\widetilde{\varepsilon}$ and $\|\boldsymbol{Y}_L\boldsymbol{X}_{L-1}^\top - \boldsymbol{M}\|_2 = \|\boldsymbol{M}\boldsymbol{X}_{L-1}\boldsymbol{X}_{L-1}^\top - \boldsymbol{M} + \boldsymbol{G}_L\boldsymbol{X}_{L-1}^\top\|_2 \leq (\lambda_1 + |\lambda_k|)\widetilde{\varepsilon}$. Therefore by definition of $\boldsymbol{a}_L$, $\boldsymbol{a}_L^\top\boldsymbol{Y}_L\boldsymbol{X}_{L-1}^\top\boldsymbol{a}_L = \max_{\|\boldsymbol{a}\|=1}\boldsymbol{a}^\top(\boldsymbol{M}\boldsymbol{X}_{L-1}\boldsymbol{X}_{L-1}^\top + \boldsymbol{G}_L\boldsymbol{X}_{L-1}^\top)\boldsymbol{a} \geq \lambda_1 - (\lambda_1 + |\lambda_k|)\widetilde{\varepsilon}$. Therefore $\boldsymbol{a}_L^\top\boldsymbol{M}\boldsymbol{a}_L \geq \lambda_1 - 2(\lambda_1 + |\lambda_k|)\widetilde{\varepsilon}$. Therefore we set $2(\lambda_1 + |\lambda_k|)\widetilde{\varepsilon} = \epsilon$, i.e., $\widetilde{\varepsilon} = 0.5\epsilon/(\lambda_1 + |\lambda_k|)$ which will get a total sample of $T = \widetilde{\Theta}(kn) = \widetilde{\Theta}(d^2\widetilde{\kappa}^2k\varepsilon^{-2})$. Then by applying Claim D.7 we get the cumulative regret bound.

Next we show how to estimate the action with $\widetilde{O}(d^2k^4\varepsilon^{-2})$ samples. To achieve this result, we need to slightly alter Algorithm 3 where we respectively set $k' = 2, 4, 6, \cdots 2k$ and keep the best arm among the k outputs. We argue that among all the choices of k', at least for one $l \in [k], k' = 2l$, we have $|\lambda_l| - |\lambda_{l+1}| \geq \lambda_1/k$. Notice with similar argument as above, when we set $n = \widetilde{\Theta}(d^2\lambda_l^{-2}\widetilde{\varepsilon}^{-2}) \leq \widetilde{\Theta}(d^2k^2\lambda_1^{-2}\widetilde{\varepsilon}^{-2})$ we can get $\|\boldsymbol{G}\| \leq \widetilde{\varepsilon}\lambda_l$ as required by Corollary C.4, the total number of iterations $L = O(\sigma_l/(\sigma_l - \sigma_{l+1})\log(2d/\epsilon) = \widetilde{O}(k)$. Finally by setting $\widetilde{\varepsilon} = \epsilon/(4\lambda_1)$ we get the overall samples we required is $\widetilde{O}(k^2n) = \widetilde{O}(d^2k^4\epsilon^{-2})$.

For both settings, directly applying our arguments in the PAC to regret conversion: Claim D.7 will finish the proof. $\square$

C.2 Omitted Details of Main Results of Low-Rank Linear Reward

Algorithm 4 Subspace Iteration Exploration for Low-rank Linear Reward.

Input: Quadratic function $f : \mathcal{A} \to \mathbb{R}$ with noisy reward, failure probability δ, error ε.
Initialization: Set $k' = 2k$. Initial candidate matrix $\boldsymbol{X}_0 \in \mathbb{R}^{d\times k'}$, $\boldsymbol{X}_0(j) \in \mathbb{R}^d, j = 1, 2, \cdots k'$ is the j-th column of $\boldsymbol{X}_0$ and are i.i.d sampled on the unit sphere $\mathbb{S}^{d-1}$ uniformly. Sample variance m, # sample per iteration n, total iteration L.
for Iteration l from 1 to L **do**
    Sample $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/mI_d), i = 1, 2, \cdots n$.
    for s from 1 to k' **do**
        Noisy subspace iteration:
        Calculate tentative rank-1 actions $\widetilde{\boldsymbol{A}}_i = \boldsymbol{X}_{l-1}(s)\boldsymbol{z}_i^\top$.
        Conduct estimation $\boldsymbol{Y}_l(s) \leftarrow m/n\sum_{i=1}^n(\langle\boldsymbol{M}, \widetilde{\boldsymbol{A}}_i\rangle + \eta_{i,s})\boldsymbol{z}_i$. ($\boldsymbol{Y}_l \in \mathbb{R}^{d\times k'}$)
    Let $\boldsymbol{Y}_l = \boldsymbol{X}_l\boldsymbol{R}_l$ be a QR-factorization of $\boldsymbol{Y}_l$
    Update target action $\boldsymbol{A}_l \leftarrow \boldsymbol{Y}_l\boldsymbol{X}_l^\top$.
Output: $\widehat{\boldsymbol{A}} = \boldsymbol{A}_L/\|\boldsymbol{A}_L\|_F$

Theorem C.3 (Formal statement of Theorem 3.8)**.** *In Algorithm 4, for large enough constants* C_n, C_L, C_m*, let* $n = C_nd^2\log^2(d/\delta)\sigma_k^{-2}\varepsilon^{-2}$*,* $m = C_md\log(n/\delta)$*, and* $L = C_L\log(d/\varepsilon)$*,* $\boldsymbol{X}_L$ *satisfies* $\|(\boldsymbol{I} - \boldsymbol{X}_L\boldsymbol{X}_L^\top)\boldsymbol{V}\| \leq \varepsilon/4$*, and the output* $\widehat{\boldsymbol{A}}$ *satisfies* $\|\widehat{\boldsymbol{A}} - \boldsymbol{A}^*\|_F \leq \|\boldsymbol{M}\|_F\varepsilon$*. Altogether to get an* ε*-optimal action, it is sufficient to have total sample complexity of* $T \leq \widetilde{O}(d^2k\lambda_k^{-2}\varepsilon^{-2})$.

Proof of Theorem 3.8 . Let $\boldsymbol{M} = \boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{V}^\top$. From Claim C.6 we get that for each noisy subspace iteration step we get $\boldsymbol{Y}_l = \boldsymbol{M}\boldsymbol{X}_l + \boldsymbol{G}_l$ with $5\|\boldsymbol{G}_l\| \leq \varepsilon\sigma_k$ and $\|\boldsymbol{V}^\top\boldsymbol{G}\| \leq \sigma_k\sqrt{k}/3\sqrt{d} \leq$

$\sigma_k(\sqrt{2k}-\sqrt{k})/2\sqrt{d}$. Therefore we can apply Corollary C.4, and get $\|\boldsymbol{V}(\boldsymbol{X}_L\boldsymbol{X}_L^\top - \boldsymbol{I})\| \le \varepsilon/4$ with $O(\log 2d/\epsilon)$ steps. Therefore we have:

$$\begin{aligned}
\|\boldsymbol{A}_L - \boldsymbol{M}\|_F =& \|(\boldsymbol{M}\boldsymbol{X}_L + \boldsymbol{G}_L)\boldsymbol{X}_L^\top - \boldsymbol{M}\|_F = \|\boldsymbol{V}^\top\boldsymbol{\Sigma}\boldsymbol{V}(\boldsymbol{X}_L\boldsymbol{X}_L^\top - \boldsymbol{I})) + \boldsymbol{G}_L\boldsymbol{X}_L^\top\|_F \\
\le& \|\boldsymbol{M}\|_F\|\boldsymbol{V}(\boldsymbol{X}_L\boldsymbol{X}_L^\top - \boldsymbol{I})\| + \|\boldsymbol{G}_L\|\|\boldsymbol{X}_L\|_F \\
\le& (\|\boldsymbol{M}\|_F + \sigma_k)\varepsilon/4 < \|\boldsymbol{M}\|_F\varepsilon/2.
\end{aligned}$$

Meanwhile, notice $\|\boldsymbol{A}^*\|_F = 1, \|\boldsymbol{M}\|_F = r^*$ and $\|\widehat{\boldsymbol{A}}\|_F = 1$. $\|\boldsymbol{A}_L/r^* - \boldsymbol{A}^*\|_F \le \varepsilon/2$. $\|\widehat{\boldsymbol{A}} - \boldsymbol{A}^*\|_F = \|\boldsymbol{A}_L/\|\boldsymbol{A}_L\|_F - \boldsymbol{A}^*\|_F = \|\mathrm{vec}(\boldsymbol{A}_L)/\|\mathrm{vec}(\boldsymbol{A}_L)\|_2 - \mathrm{vec}(\boldsymbol{A}^*)\|_2$.

Write $\theta_A := \theta(\mathrm{vec}(\boldsymbol{A}_L), \mathrm{vec}(\boldsymbol{A}^*)$. The worst case that makes $\|\mathrm{vec}(\widehat{\boldsymbol{A}}) - \mathrm{vec}(\boldsymbol{A}^*)\|$ to be larger than $\|\mathrm{vec}(\boldsymbol{A}_L/r^*) - \mathrm{vec}(\boldsymbol{A}^*)\|$ is when $\|\mathrm{vec}(\boldsymbol{A}_L/r^*) - \mathrm{vec}(\boldsymbol{A}^*)\| = \sin\theta_A$ and $\|\mathrm{vec}(\widehat{\boldsymbol{A}}) - \mathrm{vec}(\boldsymbol{A}^*)\|$ is always $2\sin(\theta_A/2)$. Notice trivially $2\sin(\theta_A/2) \le 2\sin(\theta_A)$ Therefore we could get $\|\widehat{\boldsymbol{A}} - \boldsymbol{A}^*\|_F \le 2\|\boldsymbol{A}_L/r^* - \boldsymbol{A}^*\|_F \le \varepsilon$.

□

Proof of Corollary 3.9. The corollary uses a special property of the strongly convex action set that ensures: $\boldsymbol{A}^* = \boldsymbol{M}/r^*$. With $\widehat{\boldsymbol{A}}$ that satisfies $\|\widehat{\boldsymbol{A}}\|_F = 1$, we have

$$\begin{aligned}
r^* - f_{\boldsymbol{M}}(\boldsymbol{A}) =& r^* - \langle\widehat{\boldsymbol{A}}, \boldsymbol{M}\rangle = r^* - \langle\widehat{\boldsymbol{A}}, r^*\boldsymbol{A}^*\rangle \\
=& \frac{r^*}{2}(2 - 2\langle\widehat{\boldsymbol{A}}, \boldsymbol{A}^*\rangle) = \frac{r^*}{2}(\|\widehat{\boldsymbol{A}}\|_F^2 + \|\boldsymbol{A}^*\|_F^2 - \langle\widehat{\boldsymbol{A}}, \boldsymbol{A}^*\rangle) \\
=& \frac{r^*}{2}\|\widehat{\boldsymbol{A}} - \boldsymbol{A}^*\|_F^2 \le \frac{r^*\varepsilon^2}{2} \qquad (2)
\end{aligned}$$

Therefore, with first $T_1 = \widetilde{O}(d^2k\lambda_k^{-2}\varepsilon^{-2})$ exploratory samples we get $r^* - f(\widehat{\boldsymbol{A}}) \le r^*\varepsilon^2/2 = r^*\sqrt{\frac{d^2k}{\lambda_k^2T}} = \sqrt{\frac{(r^*)^2d^2k}{\lambda_k^2T}}$. Together we have:

$$\begin{aligned}
\mathfrak{R}(T) =& \sum_{t=1}^{T_1} r^* - f(\boldsymbol{A}_t) + \sum_{t=T_1+1}^{T} r^* - f(\widehat{\boldsymbol{A}}) \\
<& r^*T_1 + Tr^*\varepsilon^2 \\
\le& \widetilde{O}(\sqrt{d^2k(r^*)^2\lambda_k^{-2}T}).
\end{aligned}$$

□

Proof of Theorem 3.10. We find an l to be the smallest integer such that $\sum_{i=l+1}^k \sigma_i^2 \le \epsilon^2\|\boldsymbol{M}\|_F^2$. Then we have $\sigma_l \ge \epsilon/\sqrt{k-l} > \epsilon/\sqrt{k}$.

Notice that in Algorithm 4, we set $n \ge \widetilde{\Theta}(\frac{d^2k}{(r^*)^2\varepsilon^4})$ large enough such that $\|\boldsymbol{G}\|_2 \le O(\|\boldsymbol{M}\|_F\epsilon^2/\sqrt{k}) \lesssim \epsilon(\sigma_l - 0)$ and $\|\boldsymbol{U}^\top\boldsymbol{G}\|_2 \le \|\boldsymbol{M}\|_F\epsilon/\sqrt{k}\frac{\sqrt{k'}-\sqrt{k-1}}{2\sqrt{d}}$. (This comes from the argument proved in Claim C.6.)

Therefore by conducting noisy power method we get with $O(nk) = \widetilde{O}(\frac{d^2k^2}{(r^*)^2\varepsilon^4})$ samples we can get an action $\widehat{\boldsymbol{A}}$ that satisfies:

$$\|\boldsymbol{M} - \boldsymbol{X}_L\boldsymbol{X}_L^\top\boldsymbol{M}\|_F^2 \le \sum_{i=l+1}^k \sigma_i^2 + l\epsilon^2\sigma_l^2 \le 2\|\boldsymbol{M}\|_F^2\epsilon^2.$$

Therefore we could get $\|\boldsymbol{A}^* - \widehat{\boldsymbol{A}}\| \le 2\epsilon$, and with similar argument as (2) we have $r^* - f(\widehat{\boldsymbol{A}}) \le \|\boldsymbol{M}\|_F\epsilon^2$.

Therefore if we want to take a total of T actions, we will set $\epsilon^6 = \widetilde{\Theta}(\frac{d^2k^2}{(r^*)^2T})$ and we get:

$$\begin{aligned}\mathfrak{R}(T) =& \sum_{t=1}^{T_1} r^* - f(\boldsymbol{A}_t) + \sum_{t=T_1+1}^{T} r^* - f(\widehat{\boldsymbol{A}})\\ <& r^*T_1 + Tr^*\varepsilon^2\\ \leq& \widetilde{O}(d^{2/3}k^{2/3}(r^*)^{1/3}T^{2/3}).\end{aligned}$$

□

C.3 Technical Details for Quadratic Reward

Noisy Power Method.

Corollary C.4 (Adapted from Corollary 1.1 from [37]). *Let $k' \geq l$. Let $\boldsymbol{U} \in \mathbb{R}^{d\times l}$ represent the top l singular vectors of $\boldsymbol{M}$ and let $\sigma_1 \geq \cdots \geq \sigma_k > 0$ denote its singular values. Suppose X_0 is an orthonormal basis of a random k'-dimensional subspace. Further suppose that at every step of NPM we have*

$$\begin{aligned}5\|\boldsymbol{G}\| \leq& \epsilon(\sigma_l - \sigma_{l+1}),\\ \text{and } 5\|\boldsymbol{U}^\top \boldsymbol{G}\| \leq& (\sigma_l - \sigma_{l+1})\frac{\sqrt{k'} - \sqrt{l-1}}{2\sqrt{d}}\end{aligned}$$

for some fixed parameter $\epsilon < 1/2$. Then with all but $2^{-\Omega(k'+1-l)} + e^{\Omega(d)}$ probability, there exists an $L = O(\frac{\sigma_l}{\sigma_l - \sigma_{l+1}}\log(2d/\epsilon))$ so that after L steps we have that $\|(\boldsymbol{I} - \boldsymbol{X}_L\boldsymbol{X}_L^\top)\boldsymbol{U}\| \leq \epsilon$.

Theorem C.5 (Adapted from Theorem 2.2 from [11]). *Let $\boldsymbol{U}_l \in \mathbb{R}^{d\times l}$ represent the top l singular vectors of $\boldsymbol{M}$ and let $\sigma_1 \geq \cdots \geq \sigma_k > 0$ denote its singular values. Naturally $l \leq k$. Suppose X_0 is an orthonormal basis of a random k'-dimensional subspace where $k' \geq k$. Further suppose that at every step of NPM we have*

$$\begin{aligned}\|\boldsymbol{G}\| \leq& O(\epsilon\sigma_l),\\ \text{and } \|\boldsymbol{U}_k^\top \boldsymbol{G}\|_2 \leq& O(\sigma_l\frac{\sqrt{k'} - \sqrt{k-1}}{2\sqrt{d}})\end{aligned}$$

for small enough ϵ. Then with all but $2^{-\Omega(k'+1-k)} + e^{\Omega(d)}$ probability, there exists an $L = O(\log(2d/\epsilon))$ so that after L steps we have that $\|(\boldsymbol{I} - \boldsymbol{X}_L\boldsymbol{X}_L^\top)\boldsymbol{U}_l\| \leq \epsilon$. Furthermore:

$$\|\boldsymbol{M} - \boldsymbol{X}_L\boldsymbol{X}_L^\top\boldsymbol{M}\|_F^2 \leq \sum_{i=l+1}^{k}\sigma_i^2 + l\sigma^2\sigma_l^2.$$

Concentration Bounds.

Claim C.6. *Write the eigendecomposition for $\boldsymbol{M}$ as $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top$. In Algorithm 4, when $n \geq \widetilde{\Theta}(d^2/(\lambda_k^2\varepsilon^2))$, the noisy subspace iteration step can be written as: $\boldsymbol{Y}_l = \boldsymbol{M}\boldsymbol{X}_{l-1} + \boldsymbol{G}_l$, where the noise term satisfies:*

$$\begin{aligned}5\|\boldsymbol{G}_l\| \leq& \varepsilon|\lambda_k|\\ 5\|\boldsymbol{U}^\top\boldsymbol{G}_l\| \leq& \varepsilon|\lambda_k|\frac{\sqrt{k}}{3\sqrt{d}}.\end{aligned}$$

with high probability for our choice of n.

Proof. For compact notation, write vector $\boldsymbol{\eta}_i := [\eta_{i,1}, \eta_{i,2}, \cdots \eta_{i,k'}]^\top \in \mathbb{R}^{k'}$. We have:

$$\begin{aligned}\boldsymbol{G}_l(s) =& \frac{m}{n}\sum_{i=1}^{n}(\boldsymbol{z}_i^\top \boldsymbol{M}\boldsymbol{X}_l(s))\boldsymbol{z}_i + \frac{m}{n}\sum_{i=1}^{n}\eta_{i,s}\boldsymbol{z}_i - \boldsymbol{M}\boldsymbol{X}_l(s), \text{ therefore}\\ \boldsymbol{G}_l =& (\frac{m}{n}\sum_{i=1}^{n}[\boldsymbol{z}_i\boldsymbol{z}_i^\top] - I)\boldsymbol{M}\boldsymbol{X}_l + \frac{m}{n}\sum_{i=1}^{n}\boldsymbol{z}_i\boldsymbol{\eta}_i^\top.\end{aligned}$$

First note that for orthogonal matrix $\boldsymbol{X}_l$, $\|\boldsymbol{M}\boldsymbol{X}_l\| \leq \lambda_1$, and $\|\frac{m}{n}\sum_{i=1}^{n}[\boldsymbol{z}_i\boldsymbol{z}_i^\top] - I\| \leq O(\sqrt{\frac{d+\log(1/\delta)}{n}})$. The bottleneck is from the second term and we will use Matrix Bernstein to concentrate it. Write $\boldsymbol{S}_i = \frac{m}{n}\boldsymbol{z}_i\boldsymbol{\eta}_i^\top$. We have $\|\boldsymbol{S}_i\| \le O(\frac{\sqrt{mk'}\log(n/\delta)}{n})$ with probability $1-\delta$ and $\mathbb{E}[\sum_i \boldsymbol{S}_i\boldsymbol{S}_i^\top] = \frac{mk'}{n}I_d$ and $\mathbb{E}[\sum_i \boldsymbol{S}_i^\top \boldsymbol{S}_i] = \frac{md}{n}I_{k'}$. Therefore with matrix Bernstein we can get that $\|\sum_i \boldsymbol{S}_i\|_i \le O(\sqrt{\frac{md}{n}}\log(d/\delta))$ with probability $1-\delta$.

Therefore for $n \ge \widetilde{\Omega}(d^2/(\lambda_k^2\varepsilon^2)$, we can get that $5\|\boldsymbol{G}_l\| \le \varepsilon|\lambda_k|$.

Similarly since $\boldsymbol{U}^\top \boldsymbol{z}_i \sim \mathcal{N}(0, \frac{1}{m}I_{k'})$, with the same argument one can easily get that $\|\boldsymbol{U}^\top \boldsymbol{G}_l\| \le O(\sqrt{\frac{mk'}{n}}\log(d/\delta))$. Therefore with the same lower bound for n one can get $15\|\boldsymbol{U}^\top \boldsymbol{G}_l\| \le \varepsilon|\lambda_k|\sqrt{\frac{k}{d}}$. □

C.4 Omitted Proof for RL with Quadratic Q function

Algorithm 5 Learn policy complete polynomial with simulator.

1: **Initialize:** Set $n = \widetilde{\Theta}(\widetilde{\kappa}^2 d^2 H^3/\varepsilon^2)$, Oracle to estimate $\widehat{T}_h$ from noisy observations.
2: **for** $h = H, \dots 1$ **do**
3: Sample $\phi(s_h^i, a_h^i), i \in [n]$ from standard Gaussian $N(0, I_d)$
4: **for** $i \in [n]$ **do**
5: Query (s_h^i, a_h^i) and use $\pi_{h+1}, \dots, \pi_H$ as the roll-out to get estimation $\widehat{Q}_h^{\pi_{h+1},\dots,\pi_H}(s_h^i, a_h^i)$
6: Retrieve $\widehat{\boldsymbol{M}}_h$ from estimation $\widehat{Q}_h^{\pi_{h+1},\dots,\pi_H}(s_h^i, a_h^i), i \in [n]$
7: Set $\widehat{Q}_h(s,a) \leftarrow f_{\widehat{T}_h}$
8: Set $\pi_h(s) \leftarrow \arg\max_{a\in\mathcal{S}} \widehat{Q}_h(s,a)$
9: **Return** $\pi_1, \dots, \pi_H$

Proof of Theorem 3.13. With the oracle, at horizon H, we can estimate $\widehat{\boldsymbol{M}}_H$ that is ϵ/H close to $\boldsymbol{M}_H^*$ in spectral norm through noisy observations from the reward function with $\widetilde{O}(\widetilde{\kappa}^2 d^2 H^2/\varepsilon^2)$ samples. Next, for each horizon $h = H-1, H-1, \cdots, 1$, sample $s_i' \sim \mathbb{P}(\cdot|s,a)$, we define $\eta_i = \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s_i', a') - \mathbb{E}_{s'\sim\mathbb{P}(\cdot|s,a)} \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s', a')$. η_i is mean-zero and $O(1)$-sub-gaussian since it is bounded. Denote $\boldsymbol{M}_h$ as the matrix that satisfies $f_{\boldsymbol{M}_h} := \mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}}$, which is well-defined due to Bellman completeness. We estimate $\widehat{\boldsymbol{M}}_h$ from the noisy observations $y_i = r_h(s,a) + \max_{a'} f_{\widehat{\boldsymbol{M}}_{h+1}}(s_i', a') = \mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}} + \eta_i =: f_{\boldsymbol{M}_h} + \eta_i$. Therefore with the oracle, we can estimate $\widehat{\boldsymbol{M}}_h$ such that $\|\widehat{\boldsymbol{M}}_h - \boldsymbol{M}_h\|_2 \le \epsilon/H$ with $\Theta(\widetilde{\kappa}^2 d^2 k^2 H^2/\epsilon^2)$ bandits. Together we have:

$$
\begin{aligned}
\|f_{\widehat{\boldsymbol{M}}_h} - f_{\widehat{\boldsymbol{M}}_h^*}\|_\infty =& \|\widehat{\boldsymbol{M}}_h - \boldsymbol{M}_h^*\| \\
\le& \|\widehat{\boldsymbol{M}}_h - \boldsymbol{M}_h\| + \|\boldsymbol{M}_h - \boldsymbol{M}_h^*\| \\
\le& \epsilon/H + \|\mathcal{T} f_{\widehat{\boldsymbol{M}}_{h+1}} - \mathcal{T} f_{\boldsymbol{M}_{h+1}^*}\|_\infty \\
\le& \epsilon/H + \|f_{\widehat{\boldsymbol{M}}_{h+1}} - f_{\boldsymbol{M}_{h+1}^*}\|_\infty \\
\le& 2\epsilon/H + \|f_{\widehat{\boldsymbol{M}}_{h+2}} - f_{\boldsymbol{M}_{h+2}^*}\|_\infty \\
\le& \cdots \\
\le& (H-h)\epsilon/H.
\end{aligned}
$$

Finally for $h=1$ we have $\|\widehat{\boldsymbol{M}}_1 - \boldsymbol{M}^*\| \le \epsilon$ if we sample $n = \widetilde{\Theta}(\widetilde{\kappa}^2 d^2 k^2 H^2/\epsilon^2)$ for each $h \in [H]$. Therefore for all the H timesteps we need $\Theta(\widetilde{\kappa}^2 d^2 k^2 H^3/\epsilon^2)$.

□

Algorithm 6 Phased elimination with zeroth order exploration.

Input: Function $f : \mathcal{A} \to \mathbb{R}$ of polynomial degree p generating noisy reward, failure probability δ, error ε.
Initialization: $L_0 = C_L k \log(1/\delta)$; Total number of stages $S = C_S\lceil \log(1/\varepsilon)\rceil + 1$, $\mathcal{A}_0 = \{\boldsymbol{a}_0^{(1)}, \boldsymbol{a}_0^{(2)}, \cdots \boldsymbol{a}_0^{(L_0)}\}$ where each $\boldsymbol{a}_0^{(l)}$ is uniformly sampled on the unit sphere $\mathbb{S}^{d-1}$. $\widetilde{\varepsilon}_0 = 1$.
for s from 1 to S **do**
  $\widetilde{\varepsilon}_s \leftarrow \widetilde{\varepsilon}_{s-1}/2$, $n_s \leftarrow C_n d^p \log(d/\delta)/\lambda_1^2 \widetilde{\varepsilon}_s^2)$,$n_s \leftarrow n_s \cdot \log^3(n_s/\delta)$, $m_s \leftarrow C_m d \log(n_s/\delta)$, $\mathcal{A}_s = \varnothing$.
  for l from 1 to L_{s-1} **do**
    Zeroth-order optimization:
    Locate current action $\widetilde{\boldsymbol{a}} = \boldsymbol{a}_{s-1}^{(l)}$.
    for $\lceil (1/(1-\alpha)) \log(2d)\rceil$ times **do**
      Sample $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/m_s I_d)$, $i = 1, 2, \cdots n_s$.
      Take actions $\boldsymbol{a}_i = (1 - \frac{1}{2p})\widetilde{\boldsymbol{a}} + \frac{1}{2p}\boldsymbol{z}_i$ and observe $r_i = \boldsymbol{T}(\boldsymbol{a}_i) + \eta_i, i \in [n_s]$; Take actions $\frac{1}{2p}\boldsymbol{z}_i$ and observe $r_i' = \boldsymbol{T}(\frac{1}{2p}\boldsymbol{z}_i) + \eta_i', i \in [n_s]$.
      Conduct estimation $\boldsymbol{y} \leftarrow 1/n_s \sum_{i=1}^{n_s}(r_i - r_i')\boldsymbol{z}_i$.
      Update the current action $\widetilde{\boldsymbol{a}} \leftarrow \boldsymbol{y}/\|\boldsymbol{y}\|$.
    Estimate the expected reward for $\widetilde{\boldsymbol{a}}$ through n_s samples: $r_n(\widetilde{\boldsymbol{a}}) = 1/n_s \sum_{i=1}^{n_s}(\boldsymbol{T}(\widetilde{\boldsymbol{a}})+\eta_i)$.
    Candidate Elimination:
    if $r_n \geq \lambda_1(1 - p\widetilde{\varepsilon}_s^2)$ **then**
      Keep the action $\mathcal{A}_s \leftarrow \mathcal{A}_s \cup \{\widetilde{\boldsymbol{a}}\}$
  Label the actions: $L_s = |\mathcal{A}_s|, \mathcal{A}_s =: \{\boldsymbol{a}_s^{(1)}, \cdots \boldsymbol{a}_s^{(L_s)}\}$.
Run UCB (Algorithm 7) with the candidate set $\mathcal{A}_S$.

D Technical details for General Tensor Reward

D.1 Technical Details for Symmetric Setting

Lemma D.1 (Zeroth order optimization for noiseless setting)**.** *For $p \geq 3$, suppose $0.5\boldsymbol{a}^\top \boldsymbol{v}_1 > |\boldsymbol{a}^\top \boldsymbol{v}_j|$ for all $j \geq 2$, we have:*

$$\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1) \leq \frac{1}{2}\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1).$$

Proof. We first simplify $G(\boldsymbol{a}) = \sum_{j=1}^r \lambda_j \boldsymbol{v}_j \cdot S_j$, where

$$\begin{aligned}
G(\boldsymbol{a}) &= \sum_{s=0}^{\lfloor (p-3)/2\rfloor} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s}\binom{p}{2s+1} T(I^{\otimes s+1}\otimes \boldsymbol{a}^{\otimes p-2s-1}) \\
&= \sum_{s=0}^{\lfloor (p-3)/2\rfloor} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s}\binom{p}{2s+1} \sum_{j=1}^k \lambda_j (\boldsymbol{v}_j^\top \boldsymbol{a})^{p-2j-1}\boldsymbol{v}_j \\
&= \sum_{j=1}^k v_j \cdot \lambda_j \overbrace{\sum_{s=0}^{\lfloor (p-3)/2\rfloor} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s}\binom{p}{2s+1}(\boldsymbol{v}_j^\top \boldsymbol{a})^{p-2s-1}}^{S_j :=} \\
&= \sum_{j=1}^k S_j \boldsymbol{v}_j.
\end{aligned}$$

Notice for even p,

$$
\begin{aligned}
S_j =& \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^3 \cdot \sum_{s=0}^{p/2-2} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s} \binom{p}{2s+1} (\boldsymbol{v}_j^\top \boldsymbol{a})^{p-2s-4} \\
=& \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^3 \cdot \sum_{r=0}^{p/2-2} \frac{(1-\frac{1}{2p})^{2r+3}(\frac{1}{2p})^{p-3-2r}}{m^{p/2-2-r}} \binom{p}{p-2r-3} (\boldsymbol{v}_j^\top \boldsymbol{a})^{2r}. \\
& \text{(let } 2r = p-4-2s\text{)}
\end{aligned}
$$

$$
\begin{aligned}
\frac{S_j}{\lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^3} =& \sum_{r=0}^{p/2-2} \frac{(1-\frac{1}{2p})^{2r+3}(\frac{1}{2p})^{p-3-2r}}{m^{p/2-2-r}} \binom{p}{p-2r-3} (\boldsymbol{v}_j^\top \boldsymbol{a})^{2r} \\
& \text{(Divide both sides by } \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^3\text{)} \\
\leq & \sum_{r=0}^{p/2-2} \frac{(1-\frac{1}{2p})^{2r+3}(\frac{1}{2p})^{p-3-2r}}{m^{p/2-2-r}} \binom{p}{p-2r-3} (\boldsymbol{v}_1^\top \boldsymbol{a})^{2r}. \\
& \text{(Since the first term is constant and } |\boldsymbol{v}_j^\top \boldsymbol{a}| \leq \boldsymbol{v}_1^\top \boldsymbol{a} \text{ for } r \geq 1\text{)} \\
=& \frac{S_1}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^3}.
\end{aligned}
$$

Therefore for even $p \geq 4$:

$$
|S_j| \leq \frac{|\lambda_j|}{\lambda_1} \frac{|\boldsymbol{v}_j^\top \boldsymbol{a}|^3}{|\boldsymbol{v}_1^\top \boldsymbol{a}|^3} S_1 \leq \frac{1}{4} \frac{|\boldsymbol{v}_j^\top \boldsymbol{a}|}{|\boldsymbol{v}_1^\top \boldsymbol{a}|} S_1, \forall j \geq 2. \tag{3}
$$

Similarly for odd p, we have:

$$
\begin{aligned}
S_j =& \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^2 \cdot \sum_{s=0}^{(p-3)/2} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s} \binom{p}{2s+1} (v_j^\top a)^{p-2s-3} \\
=& \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^2 \cdot \sum_{r=0}^{(p-3)/2} \frac{(1-\frac{1}{2p})^{2r+2}(\frac{1}{2p})^{p-2-2r}}{m^{(p-3)/2-r}} \binom{p}{p-2-2r} (\boldsymbol{v}_j^\top \boldsymbol{a})^{2r}, \\
& \text{(Let } r = (p-3)/2 - s\text{)}
\end{aligned}
$$

$$
\begin{aligned}
\frac{S_j}{\lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^2} =& \sum_{r=0}^{(p-3)/2} \frac{(1-\frac{1}{2p})^{2r+2}(\frac{1}{2p})^{p-2-2r}}{m^{(p-3)/2-r}} \binom{p}{p-2-2r} (\boldsymbol{v}_j^\top \boldsymbol{a})^{2r} \\
& \text{(Divide both sides by } \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})^2\text{)} \\
\leq & \sum_{r=0}^{(p-3)/2} \frac{(1-\frac{1}{2p})^{2r+2}(\frac{1}{2p})^{p-2-2r}}{m^{(p-3)/2-r}} \binom{p}{p-2-2r} (\boldsymbol{v}_1^\top \boldsymbol{a})^{2r} \\
& \text{(Since the first term is constant and } |\boldsymbol{v}_j^\top \boldsymbol{a}| \leq \boldsymbol{v}_1^\top \boldsymbol{a} \text{ for } r \geq 1\text{)} \\
=& \frac{S_1}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^2}.
\end{aligned}
$$

Therefore for odd p we have:

$$
|S_j| \leq \frac{|\lambda_j|}{\lambda_1} \frac{|\boldsymbol{v}_j^\top \boldsymbol{a}|^2}{|\boldsymbol{v}_1^\top \boldsymbol{a}|^2} S_1 \leq \frac{1}{2} \frac{|\boldsymbol{v}_j^\top \boldsymbol{a}|}{|\boldsymbol{v}_1^\top \boldsymbol{a}|} S_1, \forall j \geq 2. \tag{4}
$$

Write $\boldsymbol{V} = [\boldsymbol{v}_2, \boldsymbol{v}_3, \cdots, \boldsymbol{v}_k] \in \mathbb{R}^{d \times k}$ be the complement for $\boldsymbol{v}_1$. Therefore for any $\boldsymbol{x}$ without normalization, one can conveniently represent $|\tan\theta(\boldsymbol{x}, \boldsymbol{v}_1)|$ as $\|\boldsymbol{V}^\top \boldsymbol{x}\|_2 / |\boldsymbol{v}_1^\top x|$.

$$
\begin{aligned}
\|\boldsymbol{V}^\top G(\boldsymbol{a})\|^2 =& \sum_{j=2}^{k} S_j^2 && (5)\\
\leq & \sum_{j=2}^{k} \frac{|\boldsymbol{v}_j^\top \boldsymbol{a}|^2}{4|\boldsymbol{v}_1^\top \boldsymbol{a}|^2} S_1^2 && \text{(from (4),(3))}\\
=& \frac{1}{4}\tan^2\theta(\boldsymbol{v}_1, \boldsymbol{a})(\boldsymbol{v}_1^\top G(\boldsymbol{a}))^2. && (6)
\end{aligned}
$$

Therefore for $p \geq 3$, $\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1) \leq \frac{1}{2}\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)$. □

D.1.1 Proof Sketch of Theorem 3.14

Definition D.2 (Zeroth order gradient function)**.** For some scalar m, we define an empirical operator $G_n : \mathcal{A} \to \mathcal{A}$ that is similar to the zeroth-order gradient of f through n samples:

$$
G_n(\boldsymbol{a}) := \frac{m}{n}\sum_{i=1}^{n}\left(T\left(\left((1-\frac{1}{2p})\boldsymbol{a} + \frac{1}{2p}\boldsymbol{z}_i\right)^{\otimes p}\right) - \boldsymbol{T}(\frac{1}{2p}\boldsymbol{z}_i)\right)\boldsymbol{z}_i + (\eta_i - \eta_i')\boldsymbol{z}_i.
$$

where $\boldsymbol{z}_i \sim \mathcal{N}(0, \frac{1}{m}\boldsymbol{I})$ and η_i, η_i' are independent zero-mean 1-sub-Gaussian noise. Therefore we have:

$$
\begin{aligned}
\mathbb{E}[G_n(\boldsymbol{a})] =& m\mathbb{E}[\sum_{l=0}^{p-1}\binom{p}{l}\boldsymbol{T}((1-\frac{1}{2p})^{p-l}\boldsymbol{a}^{\otimes(p-l)} \otimes (\frac{1}{2p})^{l}\boldsymbol{z}^{\otimes l})\boldsymbol{z}]\\
& \text{(Due to symmetry of Gaussian only for odd } l =: 2s+1 \text{ expectation is nonzero)}\\
=& (1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}[\sum_{s=0}^{\lfloor p/2-1\rfloor} m^{-s}\binom{p}{2s+1}\boldsymbol{T}(\boldsymbol{a}^{\otimes(p-2s-1)} \otimes \boldsymbol{I}^{\otimes s+1})]
\end{aligned}
$$

Note that for even p the last term (when $s = p/2-1$) is $\boldsymbol{T}(\boldsymbol{a}\otimes\boldsymbol{I}^{\otimes p/2}) = \sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)\boldsymbol{v}_j$. While all other terms will push the iterate towards the optimal action at a superlinear speed, the last term perform a matrix multiplication and the convergence speed will depend on the eigengap. Therefore for $p \geq 4$ we will remove the extra bias in the last term that is orthogonal to $\boldsymbol{v}_1$ and will treat it as noise. (Notice for quadratic function $s = 0 = p/2-1$ is the only term in $\mathbb{E}[G_n(\boldsymbol{a})]$. This is the distinction between $p = 2$ and larger p, and why its convergence depends on eigengap.)

We further define $G(\boldsymbol{a})$ as the population version of $G_n(\boldsymbol{a})$ by removing this undesirable bias term that will be treated as noise:

$$
\begin{aligned}
G(\boldsymbol{a}) =& \begin{cases} \mathbb{E}[G_n] - \frac{(\frac{1}{2p})^{p-1}(1-\frac{1}{2p})p}{m^{p/2-1}}\sum_{j=2}^{k}\lambda_j(\boldsymbol{v}_j^\top\boldsymbol{a})\boldsymbol{v}_j, & \text{when } p \text{ is even}\\ \mathbb{E}[G_n], & \text{when } p \text{ is odd.}\end{cases}\\
=& \sum_{s=0}^{\lfloor (p-3)/2\rfloor}\frac{(\frac{1}{2p})^{2s+1}}{m^s}\binom{p}{2s+1}\boldsymbol{T}(\boldsymbol{I}^{\otimes s+1}\otimes((1-\frac{1}{2p})\boldsymbol{a})^{\otimes p-2s-1})\\
=& \frac{1}{2}(1-\frac{1}{2p})^{p-1}\boldsymbol{T}(\boldsymbol{I}, \boldsymbol{a}^{\otimes p-1}) + O(1/m).
\end{aligned}
$$

We define $G(\boldsymbol{a})$ to push the action $\boldsymbol{a}$ towards the $\boldsymbol{v}_1$ direction with at least linear convergence rate. More precisely, their angle $\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)$ will converge linearly to 0 for proper initialization with the dynamics $\boldsymbol{a} \to G(\boldsymbol{a})$. An easy way to see that is when $p = 2$ or 3, G is conducting (3-order tensor) power iteration. For higher-order problems, this operation G is equivalent to the summation of $p, p-2, p-4, \cdots$-th order tensor product and hence the linear convergence.

The estimation error $G_n(\boldsymbol{a}) - G(\boldsymbol{a})$ will be treated as noise (which is not mean zero when p is even but will be small enough: $O((2p)^{-p}m^{-(p-1)/2})$). Therefore the iterative algorithm with $\boldsymbol{a} \to G_n(\boldsymbol{a})$ will converge to a small neighborhood of $\boldsymbol{v}_1$ depending on the estimation error. This estimation error is controlled by the choice of sample size n in each iteration. We now provide the proof sketch:

Lemma D.3 (Initialization for $p \geq 3$; Corollary C.1 from [69])**.** *For any $\eta \in (0, 1/2)$, with $L = \Theta(k\log(1/\eta))$ samples $\mathcal{A} = \{\boldsymbol{a}^{(1)}, \boldsymbol{a}^{(2)}, \cdots \boldsymbol{a}^{(L)}\}$ where each $\boldsymbol{a}^{(l)}$ is sampled uniformly on the* *sphere* $\mathbb{S}^{d-1}$. *At least one sample* $\boldsymbol{a} \in \mathcal{A}$ *satisfies*

$$\max_{j\neq 1} |\boldsymbol{v}_j^\top \boldsymbol{a}| \le 0.5|\boldsymbol{v}_1^\top \boldsymbol{a}|, \text{ and } |\boldsymbol{v}_1^\top \boldsymbol{a}| \ge 1/\sqrt{d}. \tag{7}$$

with probability at least $1-\eta$.

Lemma D.4 (Iterative progress)**.** *Let* $\alpha = 1/2$ *for* $p \ge 3$ *in Algorithm 6. Consider noisy operation* $\boldsymbol{a}^+ \to G(\boldsymbol{a}) + \boldsymbol{g}$. *If the error term* $\boldsymbol{g}$ *satisfies:*

$$\begin{aligned}\|\boldsymbol{g}\| \le &\min\{\frac{0.025}{p}\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2}, 0.1\lambda_1\widetilde{\varepsilon}\} \\ &+ 0.03\lambda_1|\sin\theta(\boldsymbol{v}_1,\boldsymbol{a})|(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2}, \\ |\boldsymbol{v}_1^\top \boldsymbol{g}| \le &0.05\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}.\end{aligned}$$

Suppose $\boldsymbol{a}$ *satisfies* $0.5|\boldsymbol{v}_1^\top \boldsymbol{a}| \ge \max_{j\ge 2}|\boldsymbol{v}_j^\top \boldsymbol{a}|$, *we have:*

$$\tan\theta(\boldsymbol{a}^+,\boldsymbol{v}_1) \le 0.8\tan\theta(\boldsymbol{a},\boldsymbol{v}_1) + \widetilde{\varepsilon}.$$

We can also bound $\boldsymbol{g}$ by standard concentration plus an additional small bias term.

Lemma D.5 (Estimation error bound for G)**.** *For fixed value* $\delta \in (0,1)$ *and large enough universal constant* c_1, c_2, c_m, c_n, *when* $m = c_m d\log(n/\delta), n \ge c_n d\log(d/\delta)$, *we have*

$$\begin{aligned}\|\boldsymbol{g}\| \equiv \|G_n(\boldsymbol{a}) - G(\boldsymbol{a})\| \le &c_1\sqrt{\frac{d^2\log^3(n/\delta)\log(d/\delta)}{n}} + e\lambda_2|\sin\theta(\boldsymbol{a},\boldsymbol{v}_1)|, \\ |\boldsymbol{v}_1^\top \boldsymbol{g}| \equiv |\boldsymbol{v}_1^\top G_n(\boldsymbol{a}) - \boldsymbol{v}_1^\top G(\boldsymbol{a})| \le &c_2\sqrt{\frac{d\log^3(n/\delta)\log(d/\delta)}{n}}.\end{aligned}$$

with probability $1-\delta$. $e = 0$ *for odd p and* $e = (2p)^{-(p-1)}m^{-(p/2-1)}$ *for even p.*

Together we are able to prove Theorem 3.14:

Proof of Theorem 3.14. Initially with high probability there exists an $\boldsymbol{a}_0 \in \mathcal{A}_0$ such that Eqn. (7) holds, i.e., $\boldsymbol{v}_1^\top \boldsymbol{a}_0 \ge 1/\sqrt{d}$ and $\boldsymbol{v}_1^\top \boldsymbol{a}_0 \ge 2|\boldsymbol{v}_j^\top \boldsymbol{a}_0|, \forall j \ge 2$.

Next, from Lemma D.5, the extra bias term is bounded by $e\lambda_2|\sin\theta(\boldsymbol{a},\boldsymbol{v}_1)|$ $\le 0.03\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2}|\sin\theta(\boldsymbol{a},\boldsymbol{v}_1)|$ since $e = (2p)^{-p+1}m_s^{-p/2+1}$ and with our choice of variance $m_s \ge d \ge (\boldsymbol{v}_1^\top \boldsymbol{a})^{-2}$, plus $p \ge 3$. Next with our setting of $n_s = \widetilde{\Theta}(d^p/(\lambda_1^2\widetilde{\varepsilon}_t^2))$, the error term $\|\mathbb{E}[G(\boldsymbol{a})] - G_n(\boldsymbol{a})\|$ is upper bounded by $\widetilde{O}(\sqrt{\frac{d^2}{n}}) \le 0.025\lambda_1 d^{-(p-2)/2}\widetilde{\varepsilon}_s/p + 0.1\lambda_1\widetilde{\varepsilon}_s$. Meanwhile $|\boldsymbol{v}_1^\top \boldsymbol{g}| \le \widetilde{O}(\sqrt{\frac{d}{n_s}}) \le 0.05\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}$.

This meets the requirements for Theorem D.4 and therefore $\tan\theta(G_n(\boldsymbol{a}_0),\boldsymbol{v}_1) \le 0.8\tan\theta(\boldsymbol{a}_0,\boldsymbol{v}_1) + 0.1\lambda_1\widetilde{\varepsilon}_s$. Therefore after l steps will have

$$\begin{aligned}\tan\theta(G_n^l(\boldsymbol{a}_0),\boldsymbol{v}_1) \le &0.8^l\tan\theta(\boldsymbol{a}_0,\boldsymbol{v}_1) + \sum_{i=1}^{l} 0.8^i \cdot 0.1\widetilde{\varepsilon}_s \\ \le &0.8^l\tan\theta(\boldsymbol{a}_0,\boldsymbol{v}_1) + 0.5\widetilde{\varepsilon}_s.\end{aligned}$$

Notice initially $\tan\theta(\boldsymbol{a}_0,\boldsymbol{v}_1) \le 1/(\boldsymbol{v}_1^\top \boldsymbol{a}_0) \le \sqrt{d}$. Therefore after at most $l = O(\log_2(\tan\theta(\boldsymbol{a}_0,\boldsymbol{v}_1))) \le O(\log_2(d))$ steps, we will have $\tan\big(G_n^l(\boldsymbol{a}_0),\boldsymbol{v}_1\big) \le \widetilde{\varepsilon}_0/2 = \widetilde{\varepsilon}_1$. With the same argument, the progress also holds for $s > 0$ with even smaller l. ☐

Proof of Lemma D.5. We first estimate $G_n(\boldsymbol{a}) - \mathbb{E}[G_n(\boldsymbol{a})]$, which is want we want for even p. For odd p we will need to analyze an extra bias term that is orthogonal to $\boldsymbol{v}_1$, $\boldsymbol{e} := \frac{(\frac{1}{2p})^{p-1}(1-\frac{1}{2p})p}{m^{p/2-1}}\sum_{j=2}^k \lambda_j(\boldsymbol{v}_j^\top \boldsymbol{a})\boldsymbol{v}_j$; and we have $G_n(\boldsymbol{a}) - \mathbb{E}[G_n(\boldsymbol{a})] = G_n(\boldsymbol{a}) - G(\boldsymbol{a}) + \boldsymbol{e}$.

We decompose $G_n(\boldsymbol{a})$ as $G_n(\boldsymbol{a}) = \sum_{s=1}^k G_n^{(s)} + N$, where $G_n^{(s)} := \frac{m}{n}\sum_{i=1}^n \binom{p}{s} T(((1-0.5/p)\boldsymbol{a})^{\otimes p-s} \otimes (\boldsymbol{z}_i/(2p))^{\otimes s})\boldsymbol{z}_i$. The noise term $N := \frac{m}{n}\sum \epsilon_i \boldsymbol{z}_i$.

$$
\begin{aligned}
G_n^{(s)} :=& \frac{m}{n}\sum_{i=1}^{n}\binom{p}{s}T(((1-0.5/p)\boldsymbol{a})^{\otimes p-s}\otimes(\boldsymbol{z}_i/(2p))^{\otimes s})\boldsymbol{z}_i \\
=& \frac{m}{n}(1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}\binom{p}{s}\sum_{i=1}^{n}\sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)^{p-s}(\boldsymbol{z}_i^\top\boldsymbol{v}_j)^s\boldsymbol{z}_i. \\
\mathbb{E}[G_n^{(s)}] =& m(1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}\binom{p}{s}\sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)^{p-s}\mathbb{E}[(\boldsymbol{z}^\top\boldsymbol{v}_j)^s\boldsymbol{z}] \\
=& \begin{cases} (1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}m\binom{p}{s}\sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)^{p-s}\frac{1}{m^{(s+1)/2}}(s)!!\boldsymbol{v}_j, & \text{for odd } s, \\ 0, & \text{for even } s \end{cases} \\
=& \begin{cases} (1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}\frac{s!!}{m^{(s-1)/2}}\binom{p}{s}\sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)^{p-s}\boldsymbol{v}_j, & \text{for odd } s, \\ 0, & \text{for even } s \end{cases}
\end{aligned}
$$

$$
G_n^{(s)}-\mathbb{E}[G_n^{(s)}] = m(1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}\binom{p}{s}\sum_{j=1}^{k}\lambda_j(\boldsymbol{a}^\top\boldsymbol{v}_j)^{p-s}\boldsymbol{g}_{n,s}(j),
$$

where $\boldsymbol{g}_{n,s}(j) := \frac{1}{n}\sum_{i=1}^{n}(\boldsymbol{z}_i^\top\boldsymbol{v}_j)^s\boldsymbol{z}_i - \mathbb{E}[(\boldsymbol{z}^\top\boldsymbol{v}_j)^s\boldsymbol{z}]$.

Notice the scaling in each $G_n^{(s)}$ is $(1-\frac{1}{2p})^{p-s}(\frac{1}{2p})^{s}\binom{p}{s} \le (\frac{1}{2p})^s p^s/(s!) < 2^{-s}$ decays exponentially. In Claim D.12 we give bounds for $g_{n,s}(j)$. We note the bound for each $g_{n,s}$ also decays with s.

Therefore the bottleneck of the upper bound mostly depend on $\boldsymbol{g}_{n,0}$ and $\boldsymbol{v}_1^\top\boldsymbol{g}_{n,0}$, and we get:

$$
\begin{aligned}
\|G_n-\mathbb{E}[G_n]\| \le& C_1\lambda_1\sqrt{\frac{(d+\log(1/\delta))d\log(n/\delta)}{n}}+N, \\
|\boldsymbol{v}_1^\top G_n-\mathbb{E}[\boldsymbol{v}_1^\top G_n]| \le& C_2\lambda_1\sqrt{\frac{d\log(n/\delta)(1+\log(1/\delta))}{n}}+\boldsymbol{v}_1^\top N.
\end{aligned}
$$

Next from Claim D.11, the noise term

$$
\begin{aligned}
N \le C_3\sqrt{\frac{m\log(n/\delta)(d+\log(n/\delta))\log(d/\delta)}{n}}, \\
|v_1^\top N| \le C_4\sqrt{m\frac{\log^2(n/\delta)\log(d/\delta)}{n}},
\end{aligned}
$$

Finally $\boldsymbol{e}$ is very small: $\|\boldsymbol{e}\| \le \frac{1}{m^{p/2-1}(2p)^(p-1)}\lambda_2\|\boldsymbol{V}^\top\boldsymbol{a}\| = \lambda_2\frac{1}{m^{p/2-1}(2p)^(p-1)}\sin\theta(\boldsymbol{a},\boldsymbol{v}_1)$. $|\boldsymbol{e}^\top\boldsymbol{v}_1| = 0$.

Together we can bound $G_n(\boldsymbol{a})-G(\boldsymbol{a})$ and finish the proof. $\square$

Proof of Lemma D.4. From Lemma D.1 we have: $|\tan\theta(G(\boldsymbol{a}),\boldsymbol{v}_1)| \le 1/2|\tan\theta(\boldsymbol{a},\boldsymbol{v}_1)|$. Let $\boldsymbol{V} = [\boldsymbol{v}_2,\cdots\boldsymbol{v}_k]$. For any $p \ge 2$, we have:

$$
\begin{aligned}
|\tan\theta(\boldsymbol{a}^+,\boldsymbol{v}_1)| =& \frac{\|\boldsymbol{V}^\top\boldsymbol{a}^+\|_2}{|\boldsymbol{v}_1^\top\boldsymbol{a}^+|} \\
=& \frac{\|\boldsymbol{V}^\top(G(\boldsymbol{a})+\boldsymbol{g})\|}{|\boldsymbol{v}_1^\top(G(\boldsymbol{a})+\boldsymbol{g})|} \\
\le& \frac{\|\boldsymbol{V}^\top G(\boldsymbol{a})\|+\|\boldsymbol{V}^\top\boldsymbol{g}\|}{|\boldsymbol{v}_1^\top G(\boldsymbol{a})|-|\boldsymbol{v}_1^\top\boldsymbol{g}|} \\
\le& \frac{1/2|\tan\theta(\boldsymbol{a},\boldsymbol{v}_1)||\boldsymbol{v}_1^\top G(\boldsymbol{a})|+\|\boldsymbol{g}\|}{|\boldsymbol{v}_1^\top G(\boldsymbol{a})|-|\boldsymbol{v}_1^\top\boldsymbol{g}|} \\
=& \alpha|\tan\theta(\boldsymbol{a},\boldsymbol{v}_1)|\frac{S_1}{S_1-\|\boldsymbol{v}_1^\top\boldsymbol{g}\|}+\frac{\|\boldsymbol{g}\|}{S_1-\|\boldsymbol{v}_1^\top\boldsymbol{g}\|},
\end{aligned}
$$

where $S_1 := \boldsymbol{v}_1^\top G(\boldsymbol{a}) = \boldsymbol{v}_1^\top G(\boldsymbol{a}) = \lambda_1 \sum_{s=0}^{\lfloor (p-3)/2 \rfloor} \frac{(1-\frac{1}{2p})^{p-2s-1}(\frac{1}{2p})^{2s+1}}{m^s} \binom{p}{2s+1} (\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2s-1} \geq \lambda_1(1-\frac{1}{2p})^{p-1}(\frac{1}{2p})p(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1} \geq \frac{\lambda_1}{4}(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}$. The inequality comes from keeping only the first term where $s=0$. With the assumption that $|\boldsymbol{v}_1^\top \boldsymbol{g}| \leq 0.05\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}$, we have $|\boldsymbol{v}_1^\top \boldsymbol{g}| \leq 0.2 S_1$. Therefore

$$
\begin{aligned}
|\tan\theta(\boldsymbol{a}^+, \boldsymbol{v}_1)| \leq & 1.25/2|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)| + \frac{\|\boldsymbol{g}\|}{S_1 - |\boldsymbol{v}_1^\top \boldsymbol{g}|} \\
\leq & 1.25/2|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)| + 5/4\frac{\|\boldsymbol{g}\|}{S_1} \\
\leq & 1.25/2|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)| + 5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}}.
\end{aligned}
$$

Notice when $5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}} \leq \max\{0.125|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)|, \widetilde{\epsilon}\}$, which will ensure $|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| \leq (1.25/2 + 0.125)|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| + \widetilde{\epsilon}$. (We will prove this condition is satisfied when $\|\boldsymbol{g}\| \leq \min\{\frac{0.025}{p}\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2}, 0.1\lambda_1\widetilde{\varepsilon}\}$. We will handle the additional term in the upper bound of $\|\boldsymbol{g}\|$ later.) We divide this requirement into the following two cases. On one hand, when $|\boldsymbol{v}_1^\top \boldsymbol{a}| \leq 1 - 1/(p-1)$, $\|\boldsymbol{V}^\top \boldsymbol{a}\| \geq \sqrt{1-(1-1/(p-1))^2} > 1/p$, therefore $|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)| \geq 1/p|\boldsymbol{v}_1^\top \boldsymbol{a}|$. Therefore

$$
\begin{aligned}
& 5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}} \leq 0.125|\tan\theta(\boldsymbol{a}, \boldsymbol{v}_1)| \\
\Leftarrow & 5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}} \leq 0.125/\left(p|\boldsymbol{v}_1^\top \boldsymbol{a}|\right) \\
\Leftrightarrow & \|\boldsymbol{g}\| \leq 0.025\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-2}/p.
\end{aligned}
$$

On the other hand, when $|\boldsymbol{v}_1^\top \boldsymbol{a}| \geq 1 - 1/(p-1)$, $|\boldsymbol{v}_1^{\boldsymbol{a}}|^{p-1} \geq 1/4$ when $p=3$. Therefore $5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}} \leq 20\|\boldsymbol{g}\|/\lambda_1$. Therefore we will need $\|\boldsymbol{g}\| \leq 0.05\lambda_1\widetilde{\epsilon}$, and then the requirement that $5\frac{\|\boldsymbol{g}\|}{\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}} \leq \widetilde{\epsilon}$ is satisfied.

Altogether in both cases we have: $|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| \leq 0.75|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| + \widetilde{\epsilon}$. Finally if we additionally increase $\|\boldsymbol{g}\|$ by $0.05\lambda_1(\boldsymbol{v}_1^\top \boldsymbol{a})^{p-1}$ we will have: $|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| \leq 0.8|\tan\theta(G(\boldsymbol{a}), \boldsymbol{v}_1)| + \widetilde{\epsilon}$. □

Proof of Corollary 3.16 . As shown in Theorem 3.14 at least one action $\boldsymbol{a}$ in $\mathcal{A}_S$, $|\mathcal{A}_S| \leq \widetilde{O}(k)$ satisfies $\tan\theta(\boldsymbol{a}, \boldsymbol{a}^*) \leq \varepsilon$ with a total of $\widetilde{O}(\frac{d^p k}{\lambda_1^2 \varepsilon^2})$ steps. Therefore with Claim D.6 we have to get $\widetilde{\varepsilon}$-optimal reward we need $\widetilde{O}(\frac{d^p k}{\lambda_1 \widetilde{\varepsilon}})$ steps. Notice the eluder dimension for symmetric polynomials is d^p and the size of $\mathcal{A}_S$ is at most $\widetilde{O}(k)$. Then by applying Corollary D.9 we get that the total regret is at most $\widetilde{O}(\sqrt{d^p kT} + \sqrt{|\mathcal{A}_S|T}) = \widetilde{O}(\sqrt{d^p kT})$. □

D.2 PAC to Regret Bound Relation.

Claim D.6 (Connecting angle to regret)**.** *When $0 < \tan\theta(\boldsymbol{a}, \boldsymbol{v}_1) \leq \zeta$, we have regret $r^* - r(\boldsymbol{a}) \leq r^* \min\{2, p\zeta^2\}$.*

Proof.

$$\begin{aligned}
|\cos\theta(\boldsymbol{a},\boldsymbol{v}_1)| =& |\boldsymbol{a}^\top \boldsymbol{v}_1| =: b,\\
|\tan\theta(\boldsymbol{a},\boldsymbol{v}_1)| =& \frac{\sqrt{1-b^2}}{b} \le \zeta \Leftrightarrow b \ge \frac{1}{\sqrt{\zeta^2+1}}.\\
\Rightarrow r^* - r(\boldsymbol{a}) \le& \lambda_1 - \lambda_1 b^p\\
\le& \lambda_1(1-(\zeta^2+1)^{-p/2})\\
=& \lambda_1 \frac{(\zeta^2+1)^{p/2}-1}{(\zeta^2+1)^{p/2}}\\
\le& \lambda_1((\zeta^2+1)^{p/2}-1) \qquad \text{(since denominator } (\zeta^2+1)^{p/2}\ge 1)\\
\le& \lambda_1 p\zeta^2, \text{ when } \zeta^2 \le 1/p.
\end{aligned}$$

Additionally by definition $r^* - r(\boldsymbol{a}) \le \lambda_1 - (-\lambda_1) = 2\lambda_1$ and thus $r^* - r(a) \le \lambda_1 \min\{2, p\zeta^2\}$. We now derive the last inequality. When $\zeta \ge 1/p$ it is trivially true. When $\zeta \le 1/p$, we have $(1+\zeta^2)^{p/2} \le 1 + p\zeta^2$ for any $p \ge 2$. Since the LHS is a convex function for ζ when $p \ge 2$ and when $\zeta = 0$ LHS=RHS and when $\zeta^2 = 1/p$ LHS is always smaller than RHS (=2).

Notice the argument is straightforward to extend to the setting where the angle is between $\boldsymbol{a}$ and subspace V_1 that satisfies $\forall \boldsymbol{v} \in \boldsymbol{V}_1, T(\boldsymbol{v}) \ge \lambda_1 - \epsilon$, then one also get $r^* - r(\boldsymbol{a}) \le \lambda_1 - (\lambda_1 - \epsilon)b^p \le \min\{\lambda_1, \lambda_1 p\zeta^2 + \epsilon b^p\} \le \min\{\lambda_1, \lambda_1 p\zeta^2 + \epsilon\}$. □

Claim D.7 (Connecting PAC to Cumulative Regret)**.** *Suppose we have an algorithm alg(ζ) that finds ζ-optimal action $\widehat{\boldsymbol{a}}$ that satisfies $0 < \tan\theta(\boldsymbol{a},\boldsymbol{v}_1) \le \zeta$ by taking $A\zeta^{-a}$ actions. Here A can depend on any parameters such as d, λ_1, probability error δ, etc., that are not ζ. Then for large enough T, by calling alg with $\zeta = A^{\frac{1}{a+2}} T^{-\frac{1}{a+2}} p^{-\frac{1}{a+2}}$ and playing its output action $\widehat{\boldsymbol{a}}$ for the remaining actions, one can get a cumulative regret of:*

$$\mathfrak{R}(T) \lesssim T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^*.$$

Similarly, if an oracle finds ε-optimal action $\widehat{\boldsymbol{a}}$ that satisfies $r^* - r(\boldsymbol{a}) \le \varepsilon$ with $B\varepsilon^{-b}$ samples, then by setting $\varepsilon = (Br^*/T)^{\frac{1}{1+b}}$, and playing the output arm for the remaining actions, one can get cumulative regret of:

$$\mathfrak{R}(T) \lesssim B^{\frac{1}{1+b}} T^{\frac{b}{1+b}} r^{\frac{1}{1+b}}.$$

Proof. For the chosen ζ, write $T_1 = A\zeta^{-a}$ be the number of actions that finds ζ-optimal action. Therefore $T_1 = A^{\frac{2}{a+2}} T^{\frac{a}{a+2}} p^{\frac{a}{a+2}}$. First, when $T \ge Ap^{a/2}$, $\zeta^2 \le 1/p$, namely $r^* - r(\boldsymbol{a}) \le r^* p\zeta^2$. We have:

$$\begin{aligned}
\mathfrak{R}(T) \le& \sum_{t=1}^{T_1} 2r^* + \sum_{t=T_1+1}^{T} r^* p\zeta^2\\
\le& 2r^* T_1 + Tr^* p\zeta^2\\
\le& 3T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^*.
\end{aligned}$$

When $T < Ap^{a/2}$, it trivially holds that $\mathfrak{R}(T) \le 2r^*T < 2T^{\frac{a}{a+2}} p^{\frac{a}{a+2}} A^{\frac{2}{a+2}} r^*$. □

Algorithm 7 UCB (Algorithm 1 in Section 5 of [5])

1: **Input:** Stochastic reward function f, failure probability δ, action set $\mathcal{A}$ with finite size K.
2: **for** t from 1 to $T-1-K$ **do**
3: Execute arm $I_t = \arg\max_{i\in[K]} \left(\widehat{\mu}^t(i) + \sqrt{\frac{\log(TK/\delta)}{N^t(i)}}\right)$. Here $N^t(\boldsymbol{a}) = 1 + \sum_{i=1}^t \mathbf{1}\{I_i = \boldsymbol{a}\}$; and $\widehat{\mu}^t(\boldsymbol{a}) = \frac{1}{N^t(\boldsymbol{a})}\left(r_a + \sum_{i=1}^t \mathbf{1}\{I_i = \boldsymbol{a}\} r_i\right)$.
4: Observe r_{I_t}

Theorem D.8 (Theorem 5.1 from [5]). *With UCB algorithm on action set with size K, we have with probability $1-\delta$,*

$$\mathfrak{R}(T) = \widetilde{O}(\min\{\sqrt{KT}\} + K).$$

Corollary D.9. *With the same setting of Claim D.7, except that now the algorithm alg(ε) finds **a set** $\mathcal{A}$ of size S where at least one action $\boldsymbol{a} \in \mathcal{A}$ satisfies $r^* - f(\boldsymbol{a}) \leq \varepsilon$. Then all argument in Claim D.7 still hold by adding $\widetilde{O}(\sqrt{ST})$ on the RHS of each regret bound.*

Proof. Suppose we run *alg* for T_1 steps and achieve ε-optimal reward.

Let $r_\varepsilon := \max_{\boldsymbol{a}\in\mathcal{A}} f(\boldsymbol{a})$. Therefore with UCB on mutiarm bandit we have: $\sum_{t=T_1+1}^{T} r_\varepsilon - f(\boldsymbol{a}_t) \leq \widetilde{O}(\sqrt{ST})$ by Theorem D.8.

From the statement $r_\varepsilon \geq r^* - \zeta$. Therefore $\sum_{t=T_1+1}^{T} r^* - f(\boldsymbol{a}_t) \leq \widetilde{O}(\sqrt{ST}) + \varepsilon(T-T_1)$. Therefore

$$\mathfrak{R}(T) \leq \sum_{t=1}^{T_1} 2r^* + \varepsilon(T-T_1) + \widetilde{O}(\sqrt{ST}).$$

With the same choices of T_1 in Claim D.7, the same conclusion still holds with an additional term of $\widetilde{O}(\sqrt{ST})$. ☐

For symmetric tensor problems the set size is $\widetilde{O}(k)$ and therefore we will have an additional $\sqrt{kT}$ term which will be subsumed in our regret bound.

D.3 Variance and Noise Concentration

Lemma D.10 (Vector Bernstein; adapted from Theorem 7.3.1 in [64]). *Consider a finite sequence $\{\boldsymbol{x}_k\}_{k=1}^n$ be i.i.d randomly generated samples, $x_k \in \mathbb{R}^d$, and assume that $\mathbb{E}[\boldsymbol{x}_k] = 0$, $\|\boldsymbol{x}_k\| \leq L$, and covariance matrix of x_k is Σ. Then it satisfies that when $n \geq \log d/\delta$, we have:*

$$\left\|\frac{\sum_{i=1}^n \boldsymbol{x}_i}{n}\right\| \leq C\sqrt{\frac{(\|\Sigma\| + L^2)\log d/\delta}{n}},$$

with probability $1-\delta$.

Claim D.11 (Noise concentration). *Let independent samples $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/mI_d)$ and $\epsilon_i \sim \mathcal{N}(0,1)$. With probability $1-\delta, \delta \in (0,1)$:*

$$\begin{aligned}\left\|\frac{m}{n}\sum_{i=1}^n \epsilon_i \boldsymbol{z}_i\right\| \leq & C\sqrt{\frac{m\log(n/\delta)(d+\log(n/\delta))\log(d/\delta)}{n}} \\ \left|\frac{m}{n}\sum_{i=1}^n \epsilon_i \boldsymbol{z}_i^\top \boldsymbol{v}_1\right| \leq & C'\sqrt{\frac{m\log^2(n/\delta)\log(d/\delta)}{n}}.\end{aligned}$$

Proof. We use the Vector Bernstein Lemma D.10. The covariance matrix for $\boldsymbol{x}_i = \epsilon_i \boldsymbol{z}_i$ satisfies $\mathbb{E}[\boldsymbol{x}_i \boldsymbol{x}_i^\top] = 1/mI_d$. $\boldsymbol{x}_i$ is mean zero. $\|\epsilon_i \boldsymbol{z}_i\|^2 = \epsilon_i^2 \|\boldsymbol{z}_i\|^2$. Notice $\epsilon_i^2 \sim \chi(1) \lesssim 1 + \log(1/\delta)$ and $m\boldsymbol{z}_i^\top \boldsymbol{z}_i \sim \chi(d) \lesssim d + \log(1/\delta)$. Therefore by directly applying Vector Bernstein $\|\epsilon_i \boldsymbol{z}_i\| \leq c\sqrt{\frac{(1+\log(1/\delta))(d+\log(1/\delta))}{m}}$ with probability $1-\delta$. By union bound we have: for all i, $\|\epsilon_i \boldsymbol{z}_i\| \leq c\sqrt{\frac{\log(n/\delta)(d+\log(n/\delta))}{m}}$ with probability $1-\delta$. Therefore

$$\left\|\frac{1}{n}\sum_{i=1}^n \epsilon_i \boldsymbol{z}_i\right\| \leq C\sqrt{\frac{\log(n/\delta)(d+\log(n/\delta))\log(d/\delta)}{mn}},$$

with probability $1-\delta$. Similarly

$$\begin{aligned}\left|\frac{1}{n}\sum_{i=1}^n \epsilon_i \boldsymbol{z}_i^\top \boldsymbol{v}_1\right| \leq & C\sqrt{\frac{\log(n/\delta)(1+\log(n/\delta))\log(d/\delta)}{mn}} \\ = & C'\sqrt{\frac{\log^2(n/\delta)\log(d/\delta)}{mn}},\end{aligned}$$

$\square$

Claim D.12. *Let $\{\boldsymbol{z}_i\}_{i=1}^n$ be i.i.d samples from $\mathcal{N}(0, 1/mI_d)$. Let $g_{n,s}(j) := \frac{1}{n}\sum_{i=1}^n (\boldsymbol{z}_i^\top \boldsymbol{v}_j)^s \boldsymbol{z}_i - \mathbb{E}[(\boldsymbol{z}^\top \boldsymbol{v}_j)^s \boldsymbol{z}]$. We have:*

$$
\begin{aligned}
\|g_{n,0}(j)\| &\lesssim \sqrt{\frac{d+\log(1/\delta)}{nm}}, \\
|\boldsymbol{v}_1^\top g_{n,0}(j)| &\lesssim \sqrt{\frac{1+\log(1/\delta)}{nm}}, \\
|\boldsymbol{v}_1^\top g_{n,1}(j)| \le \|g_{n,1}(j)\| &\lesssim \sqrt{\frac{d+\log(1/\delta)}{m^2 n}}, \text{ when } n \ge d\log(1/\delta), \\
|\boldsymbol{v}_1^\top g_{n,s}(j)| \le \|g_{n,s}(j)\| &\lesssim \sqrt{\frac{\log(d/\delta)}{d^s n}}, \text{ when } n \ge \log(d/\delta), m \ge c_0 d \log(n/\delta), s \ge 2.
\end{aligned}
$$

For any $j \in [k]$.

We mostly care about the correct concentration for smaller s. For larger s a very loose bound will already suffice our requirement.

Proof of Claim D.12. For $s = 0$, $nm\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\|^2 \sim \chi(d)$, therefore $\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i\| \lesssim \sqrt{\frac{d+\log(1/\delta)}{nm}}$. $nm(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i^\top \boldsymbol{v}_1)^2 \sim \chi(1)$. Therefore $|\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i^\top \boldsymbol{v}_1| \lesssim \sqrt{\frac{1+\log(1/\delta)}{nm}}$.

For $s = 1$, due to standard concentration for covariance matrices (see e.g. [64, 22]), we have:

$$
m\|(\frac{1}{n}\sum_{i=1}^n \boldsymbol{z}_i \boldsymbol{z}_i^\top - \mathbb{E}[\boldsymbol{z}\boldsymbol{z}^\top])\| \le \max\{\sqrt{\frac{d+\log(2/\delta)}{n}}, \frac{d+\log(2/\delta)}{n}\}.
$$

Therefore when $n \ge d\log(1/\delta)$, both results

$$
\begin{aligned}
\|g_{n,1}(j)\| &\lesssim \sqrt{\frac{d+\log(1/\delta)}{m^2 n}} \|\boldsymbol{v}_j\|, \\
&= \sqrt{\frac{d+\log(1/\delta)}{m^2 n}}, \text{ and} \\
\|\boldsymbol{v}_1^\top g_{n,1}(j)\| &\lesssim \sqrt{\frac{d+\log(1/\delta)}{m^2 n}} \|\boldsymbol{v}_1\|\|\boldsymbol{v}_j\| \\
&= \sqrt{\frac{d+\log(1/\delta)}{m^2 n}}
\end{aligned}
$$

hold.

For larger $s \ge 2$, with probability $1-\delta$, $|\boldsymbol{z}_i^\top \boldsymbol{v}_j| \le C\sqrt{\log(n/\delta)/m} = Cc_0/\sqrt{d} \le 1/\sqrt{d}$. When $m \ge c_0 d \log(n/\delta)$, for small enough c_0 we have $|\boldsymbol{z}_i^\top \boldsymbol{v}_j| \le 1/\sqrt{d}$ and $\|\boldsymbol{z}_i\| \le 1$ for all $i \in [n]$. Therefore $\|(\boldsymbol{z}_i^\top \boldsymbol{v}_j)^s \boldsymbol{z}_i\| \le d^{-s/2}$ We can use vector Bernstein, i.e., Lemma D.10 to get:

$$
\|g_{n,s}(j)\| \le C_1 \sqrt{\frac{\log(d/\delta)}{d^s n}}.
$$

Therefore we have:

$$
|g_{n,s}(j)^\top \boldsymbol{v}_1| \le C_1 \sqrt{\frac{\log(d/\delta)}{d^s n}}.
$$

$\square$

D.3.1 The asymmetric setting

Now we consider the asymmetric tensor problem with reward $f : \mathcal{A} \to \mathbb{R}$. The input space $\mathcal{A}$ consists of p vectors in a unit ball: $\vec{\boldsymbol{a}} = (\boldsymbol{a}(1), \boldsymbol{a}(2), \cdots \boldsymbol{a}(p)) \in \mathcal{A}, \|\boldsymbol{a}(s)\| \le 1, \forall s \in [p]$. $f(\vec{\boldsymbol{a}}) = \boldsymbol{T}(\otimes_{s=1}^p \boldsymbol{a}(s)) + \eta$. Tensor $\boldsymbol{T} = \sum_{j=1}^k \lambda_j \boldsymbol{v}_j(1) \otimes \boldsymbol{v}_j(2) \cdots \otimes \boldsymbol{v}_j(p)$. For each $s \in [p]$, $\{\boldsymbol{v}_1(s), \boldsymbol{v}_2(s), \cdots \boldsymbol{v}_k(s)\}$ are orthonormal vectors. We order the eigenvalues such that $\lambda_1 \ge |\lambda_2| \cdots \ge |\lambda_k|$. Therefore the optimal reward is λ_1 and can be achieved by $\boldsymbol{a}^*(s) = \boldsymbol{v}_1(s), s \in [p]$.

In this section we only consider $p \geq 3$ and leave the quadratic and low-rank matrix setting to the next section.

Theorem D.13. *For $p \geq 3$, by conducting alternating power iteration, one can get a ε-optimal reward with a total $\widetilde{O}\left((2k)^p \log^p(p/\delta) d^p \lambda_1^{-1} \varepsilon^{-1}\right)$ actions; therefore the regret bound is at most $\widetilde{O}(\sqrt{k^p d^p T})$.*

This setting is actually much easier than the symmetric setting. Notice by replacing one slice of $\vec{\boldsymbol{a}}$ by random Gaussian $\boldsymbol{z}_i \sim \mathcal{N}(0, 2/d \log(d/\delta))$, one directly gets $\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p))$ on each slice with $1/n \sum_i f(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{z}_i, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)) \boldsymbol{z}_i$ which is tensor product. We defer the proof to Appendix D.4.

D.4 Omitted Details for Asymmetric Tensors

Algorithm 8 Phased elimination with alternating tensor product.

Input: Stochastic reward $r : (B_1^d)^{\otimes p} \to \mathbb{R}$ of polynomial degree p, failure probability δ, error ε.
Initialization: $L_0 = C_L k \log(1/\delta)$; Total number of stages $S = C_S \lceil \log(1/\varepsilon) \rceil + 1$, $\mathcal{A}_0 = \{\boldsymbol{a}_0^{(1)}, \boldsymbol{a}_0^{(2)}, \cdots \boldsymbol{a}_0^{(L_0)}\} \subset (B_1^d)^{\otimes p}$ where each $\boldsymbol{a}_0^{(l)}(j), j \in [p]$ is uniformly sampled on the unit sphere $\mathbb{S}^{d-1}$. $\widetilde{\varepsilon}_0 = 1$.
for s from 1 to S **do**
  $\widetilde{\varepsilon}_s \leftarrow \widetilde{\varepsilon}_{s-1}/2$, $n_s \leftarrow C_n d^p \log(d/\delta)/\widetilde{\varepsilon}_s^2$, $n_s \leftarrow n_s \cdot \log^3(n_s/\delta)$, $m_s \leftarrow C_m d \log(n/\delta)$, $\mathcal{A}_s = \varnothing$.
  for l from 1 to L_{s-1} **do**
    Tensor product update:
    Locate current arm $\widetilde{\boldsymbol{a}} = \boldsymbol{a}_{s-1}^{(l)}$.
    for $\lceil (\lambda_1/\Delta) \log(2d) \rceil$ times **do**
      for j from 1 to p **do**
        Sample $\boldsymbol{z}_i \sim \mathcal{N}(0, 1/m_s I_d), i = 1, 2, \cdots n_s$.
        Calculate tentative arm $\boldsymbol{a}_i \leftarrow \widetilde{\boldsymbol{a}}, \boldsymbol{a}_i(j) = (1 - \widetilde{\varepsilon}_s)\widetilde{\boldsymbol{a}}(j) + \widetilde{\varepsilon}_s \boldsymbol{z}_i$
        Conduct estimation $\boldsymbol{y} \leftarrow 1/n_s \sum_{i=1}^{n_s} r_{\epsilon_i}(\boldsymbol{a}_i)\boldsymbol{z}_i$.
        Update the current arm $\widetilde{\boldsymbol{a}}(j) \leftarrow \boldsymbol{y}/\|\boldsymbol{y}\|$.
    Estimate the expected reward for $\widetilde{\boldsymbol{a}}$ through n_s samples: $r_n = 1/n_s \sum_{i=1}^{n_s} r_{\epsilon_i}(\widetilde{\boldsymbol{a}})$.
    Candidate Elimination:
    if $r_n \geq \lambda_1(1 - p\widetilde{\varepsilon}_s)$ **then**
      Keep the arm $\mathcal{A}_s \leftarrow \mathcal{A}_s \cup \{\widetilde{\boldsymbol{a}}\}$
  Label the arms: $L_s = |\mathcal{A}_s|, \mathcal{A}_t =: \{\boldsymbol{a}_s^{(1)}, \cdots \boldsymbol{a}_s^{(L_s)}\}$.
Run Algorithm 7 with $\mathcal{A}_S$.

Lemma D.14 (Asymmetric Tensor Initialization)**.** *With probability $1 - \delta$, with $L = \widetilde{\Theta}((2k)^p \log^p(p/\delta))$ random initializations $\mathcal{A}_0 = \{\boldsymbol{a}_0^{(0)}, \boldsymbol{a}_0^{(1)}, \cdots \boldsymbol{a}_0^{(L)}\}$, there exists an initialization $\boldsymbol{a}_0 \in \mathcal{A}_0$ that satisfies:*

$$\alpha \boldsymbol{a}_0(s)^\top \boldsymbol{v}_1^{(s)} \geq |\boldsymbol{a}_0(s)^\top \boldsymbol{v}_j^{(s)}|, \forall j \geq 2 \& j \in [k], \forall s \in [p], \qquad (8)$$
$$\boldsymbol{a}_0(s)^\top \boldsymbol{v}_1^{(s)} \geq 1/\sqrt{d}.$$

with some constant $\alpha < 1$.

Proof. This lemma simply comes from applying Lemma D.3 for p times and we need $\geq 2k \log_2(p\delta)$ to ensure the condition for each $\boldsymbol{a}_0(s), s \in [p]$ holds. Therefore together we will need $(2k \log_2(p/\delta))^p$ samples. $\square$

Lemma D.15 (Asymmetric tensor progress)**.** *For each a that satisfies Eqn.* (8) *with constant $\alpha < 1$, we have:*

$$\tan \theta(\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)), \boldsymbol{v}_1^{(s)}) \leq \alpha \tan \theta(\boldsymbol{a}_j, \boldsymbol{v}_1^{(j)}),$$

for any j that is in $[p]$ but is not s. When $n \geq \Theta(d^p \log(d/\delta) \log^3(n/\delta)/\widetilde{\varepsilon}^2)$ and $m = \Theta(d \log(n/\delta))$, we have:

$$\begin{aligned}&\tan\theta(\boldsymbol{T}(\boldsymbol{a}(1), \cdots \boldsymbol{a}(s-1), \boldsymbol{I}, \boldsymbol{a}(s+1), \cdots \boldsymbol{a}(p)), \boldsymbol{v}_1^{(s)})\\ \leq&(1+\alpha)/2 \tan\theta(\boldsymbol{a}_j, \boldsymbol{v}_1^{(j)}) + \widetilde{\varepsilon}, \forall j \in [p] \& j \neq s.\end{aligned}$$

The remaining proof is a simpler version for the symmetric tensor setting on conducting noisy power method with the good initialization and iterative progress.

Finally due to the good initialization that satisfies (8) and together with Lemma D.15 we can finish the proof for Theorem D.13.

E Proof of Theorem G.3

E.1 Additional Notations

Here, we briefly introduce complex and real algebraic geometry. This section is based on [55, 63, 12, 68].

An (affine) **algebraic variety** is the common zero loci of a set of polynomials, defined as $V = Z(S) = \{\boldsymbol{x} \in \mathbb{C}^n : f(\boldsymbol{x}) = 0, \forall f \in S\} \subseteq \mathbb{A}^n = \mathbb{C}^n$ for some $S \subseteq \mathbb{C}[x_1, \cdots, x_n]$. A **projective variety** U is a subset of $\mathbb{P}^n = (\mathbb{C}^{n+1} \setminus \{0\})/\sim$, where $(x_0, \cdots, x_n) \sim k(x_0, \cdots, x_n)$ for $k \neq 0$ and S is a set of homogeneous polynomials of $(n+1)$ variables.

For an affine variety V, its **projectivization** is the variety $\mathbb{P}(V) = \{[\boldsymbol{x}] : x \in V\} \subseteq \mathbb{P}^{n-1}$, where $[\boldsymbol{x}]$ is the line corresponding to $\boldsymbol{x}$.

The **Zariski topology** is the topology generated by taking all varieties to be the closed sets.

A set is **irreducible** if it is not the union of two proper closed subsets.

A **variety is irreducible** if and only if it is irreducible under the Zariski topology.

The **algebraic dimension** $d = \dim V$ of a variety V is defined as the length of the longest chain $V_0 \subset V_1 \subset \cdots \subset V_d = V$, such that each V_i is irreducible.

A variety V is said to be **admissible** to a set of linear functions $\{\ell_\alpha : \mathbb{C}^d \to \mathbb{C}\}_{\alpha \in I}\}$, if for every ℓ_α, we have $\dim(V \cap \{\boldsymbol{x} \in \mathbb{C}^d : \ell_\alpha(\boldsymbol{x}) = 0\}) < \dim V$.

A map $f = (f_1, \cdots, f_m) : \mathbb{A}^n \to \mathbb{A}^m$ is **regular** if each f_i is a polynomial.

A **algebraic set** is the common real zero loci of a set of polynomials.

For a complex variety $V \subseteq \mathbb{A}^n$, its real points form a algebraic set $V_\mathbb{R}$.

For an algebraic set $V_\mathbb{R}$, its real dimension $d = \dim_\mathbb{R} V_\mathbb{R}$ is the maximum number d such that $V_\mathbb{R}$ is locally semi-algebraically homeomorphic to the unit cube $(0,1)^d$, details can be found in [12].

E.2 Proof of Sample Complexity

Lemma E.1 ([68], Theorem 3.2)**.** *For $i = 1, \ldots, T$, let $L_i : \mathbb{C}^n \times \mathbb{C}^m \to \mathbb{C}$ be bilinear functions and V_i be varieties given by homogeneous polynomials in $\mathbb{C}^n$. Let $V = V_1 \times \cdots \times V_T \subseteq (\mathbb{C}^n)^N$. Let $W \subseteq \mathbb{C}^m$ be a variety given by homogeneous polynomials. In addition, we assume V_i is admissible with respect to the linear functions $\{f^{\boldsymbol{w}}(\cdot) = L_i(\cdot, \boldsymbol{w}) : \boldsymbol{w} \in W \setminus \{0\}\}$. When $T \geq \dim W$, let $\delta = T - \dim W + 1 \geq 1$. Then there exists a subvariety $Z \subseteq V$ with $\dim Z \leq \dim V - \delta$ such that for any $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T) \in V \setminus Z$ and $\boldsymbol{w} \in W$, if $L_1(\boldsymbol{x}_1, \boldsymbol{w}) = \cdots = L_T(\boldsymbol{x}_T, \boldsymbol{w}) = 0$, then $\boldsymbol{w} = 0$.*

Lemma E.2 ([68], Lemma 3.1)**.** *Let V be an algebraic variety in $\mathbb{C}^d$. Then $\dim_\mathbb{R} V_\mathbb{R} \leq \dim V$.*

Lemma E.3. *Let W be a vector space. For vectors $\boldsymbol{x}_1, \cdots, \boldsymbol{x}_T$, if the map $f : \boldsymbol{w} \mapsto (\langle \boldsymbol{x}_1, w \rangle, \ldots, \langle \boldsymbol{x}_T, \boldsymbol{w} \rangle)$ is not injective over $W - W := \{\boldsymbol{w}_1 - \boldsymbol{w}_2 : \boldsymbol{w}_1, \boldsymbol{w}_2 \in W\}$, then there exists $\boldsymbol{v} \in W$ such that $f(\boldsymbol{v}) = 0$.*

Proof. Suppose $f(\boldsymbol{w}_1) = f(\boldsymbol{w}_2)$. Let $\boldsymbol{v} = \boldsymbol{w}_1 - \boldsymbol{w}_2$. Then $\boldsymbol{v} \in W - W$ and $f(\boldsymbol{v}) = f(\boldsymbol{w}_1) - f(\boldsymbol{w}_2) = 0$. $\square$

Definition E.4 (Tensorization)**.** Let f be a polynomial of $x_1, \cdots, x_d$ with degree $\deg f \le p$. Then every p-tensor $\mathbf{W}_f$ satisfying $\langle \mathbf{W}_f, \mathbf{X}_{\boldsymbol{x}} \rangle = f(\boldsymbol{x})$ is said to be a *tensorization* of the polynomial f, where $\mathbf{X}_{\boldsymbol{x}}$ is the tensorization of $\boldsymbol{x}$ itself:

$$\mathbf{X}_{\boldsymbol{x}} = \begin{pmatrix} 1 \\ \boldsymbol{x} \end{pmatrix}^{\otimes p}. \tag{9}$$

Let $\mathcal{F}$ be a class of polynomials. A variety of tensorization of $\mathcal{F}$ is defined to be an irredicuble closed variety defined by homogeneous polynomials W, such that for every $f \in \mathcal{F}$, there is a tensorization $\mathbf{W}_f$ of f, such that $W \ni \mathbf{W}_f$ contains its tensorization. Note that neither tensorization of f nor variety of tensorization of $\mathcal{F}$ is unique.

We define the variety of tensorization of $\boldsymbol{x}$ as follows. (Note that this is uniquely defined.) Consider the regular map

$$\varphi_1 : \mathbb{C}^d \to \mathbb{C}^{(d+1)^p}, \qquad \boldsymbol{x} \mapsto \begin{pmatrix} 1 \\ \boldsymbol{x} \end{pmatrix}^{\otimes p}, \tag{10}$$

the tensorization of $\boldsymbol{x}$ is defined as $V_i = \mathbb{P}(\overline{\operatorname{Im} \varphi_1})$.

Note that V_i is irreducible because φ_1 is regular and $\mathbb{C}^d$ is irreducible. By [55, Theorem 9.9], its dimension is given by

$$\dim V_i \le \dim \overline{\operatorname{Im} \varphi_1} + 1 \le \dim \mathbb{C}^d + 1 = d + 1. \tag{11}$$

Lemma E.5. *For any non-zero polynomial $f \neq 0$ with $\deg f \le p$. Let W_f be a tensorization of f. Then V_i is admissible with respect to $\{L_i(\cdot) = \langle \cdot, W_f \rangle\}$.*

Proof. Since V_i is irredicuble and L_i is a linear function, it suffices to verify that $\langle \mathbf{X}_{\boldsymbol{x}}, W_f \rangle \neq 0$ [68]. But according to Definition E.4, $\langle \mathbf{X}_{\boldsymbol{x}}, W_f \rangle \neq 0$ is equivalent to

$$f(\boldsymbol{x}) = \left\langle \mathbf{W}_f, \begin{pmatrix} 1 \\ \boldsymbol{x} \end{pmatrix}^{\otimes p} \right\rangle \neq 0. \tag{12}$$

Since $f \neq 0$, we must have $f(\boldsymbol{x}) \neq 0$ for some $\boldsymbol{x}$, which gives a non-zero $\mathbf{X}_{\boldsymbol{x}} \neq 0$ for the above equation: $\langle \mathbf{X}_x, W_f \rangle \neq 0$, and we conclude that V_i is admissible. □

Lemma E.6. *Let $V \subset \mathbb{C}^n$ be a (Zariski) closed proper subset, $V \neq \mathbb{C}^n$. Then V is a null set, i.e. it has (Lebesgue) measure zero.*

Proof. Suppose $V = Z(S)$ is the vanishing set for some $S \subseteq \mathbb{C}[x_1, \cdots, x_n]$. Since $V \neq \mathbb{C}^n$, let $f \in S$, we have $V \subseteq Z(f)$, so it suffices to show $\operatorname{Leb}(Z(f)) = 0$, which is because $Z(f) = f^{-1}(0), \operatorname{Leb}(\{0\}) = 0, f$ is a continuous function (under Euclidean topology), and $\operatorname{Leb}(\{\boldsymbol{x} : \nabla f(\boldsymbol{x}) = 0\}) = 0$. □

Theorem E.7. *Assume that the reward function class is a class of polynomials $\mathcal{F}$. Let W be (one of) its variety of tensorization. If we sample $T \ge \dim W$ times, and the sample points satisfying $(\boldsymbol{x}_1, \cdots, \boldsymbol{x}_T) \in (\mathbb{C}^d)^T \setminus Z$ for some null set Z. Then we can uniquely determine the reward function f from the observed rewards $(f(\boldsymbol{x}_1), \cdots, f(\boldsymbol{x}_T))$.*

Proof. Let $n = m = (d+1)^p, L_i(\boldsymbol{x}, \boldsymbol{w}) = \langle \boldsymbol{x}, \boldsymbol{w} \rangle, V = V_1 \times \cdots \times V_T$, where V_i is as in Definition E.4. By [63, Example 1.33], we have $\dim V \le (d+1)T$. Since W is a vareity of tensorization, by Lemma E.5, V_i is admissible with respect to $\{L_i(\cdot, \mathbf{W}) : \mathbf{W} \in W\}$.

We are now ready to apply Lemma E.1, which gives that when $T \ge \dim W$, there exists subvariety $Z \subset V$ with $\dim Z < \dim V \le rT$, and for any $(\mathbf{X}_1, \cdots, \mathbf{X}_T) \in V \setminus Z$ and any $\mathbf{W} \in W$, if $\langle \mathbf{X}_1, \mathbf{W} \rangle = \cdots = \langle \mathbf{X}_T, \mathbf{W} \rangle = 0$, then $\mathbf{W} = 0$. By Lemma E.3, we have for every $(\mathbf{X}_1, \cdots, \mathbf{X}_T) \in V \setminus Z$, the map $\mathbf{W} \mapsto (\langle \mathbf{X}_1, \mathbf{W} \rangle, \cdots, \langle \mathbf{X}_T, \mathbf{W} \rangle)$ is injective, so $\mathbf{W}_f$ and thus f can be uniquely recovered from the observed rewards.

Finally, we show that $(\varphi_1^{-1} \times \cdots \times \varphi_1^{-1})(Z)$ is a null set, where φ_1 is as in (10). According to the proof of Lemma E.1 by [68], we find that Z is also defined by homogeneous polynomials. We take the slice $Z' = \{\boldsymbol{x} \in Z : x_{11} = \cdots = x_{T1} = 1\}, V' = \{\boldsymbol{x} \in V : x_{11} = \cdots = x_{T1} = 1\}$, (here x_{ij} is the j-th coordinate of $\boldsymbol{x}_i$), then Z', V' are varieties. Since $\dim Z < \dim V$, we have $\dim Z' = \dim Z - T < \dim V - T = \dim V'$ and $Z' \subset V'$.

Now consider the regular map $\varphi_1' : V' \to (\mathbb{C}^d)^T$,

$$\left(\begin{pmatrix}1\\ \boldsymbol{x}_1\end{pmatrix}^{\otimes p}, \cdots, \begin{pmatrix}1\\ \boldsymbol{x}_1\end{pmatrix}^{\otimes p}\right) \mapsto (\boldsymbol{x}_1, \cdots, \boldsymbol{x}_T). \tag{13}$$

Then $\varphi_1'(Z'), \varphi_1'(V')$ are both varieties. By [55, Lemma 9.9], we have $\dim \overline{\varphi_1'(Z')} \leq \dim Z$. Since $\varphi_1'(V) = (\mathbb{C}^d)^T$ and $\dim \varphi_1'(V') \leq \dim V' = \dim V - T \leq (d+1)T - T$, we have $\dim V = \dim \varphi_1'(V) = dT$ and as a result, $\dim \overline{\varphi_1'(Z)} \leq \dim Z < \dim V = dT$. By Lemma E.6, $\overline{\varphi_1'(Z)}$ is a null set. Since $(\boldsymbol{x}_1, \cdots, \boldsymbol{x}_T) \notin \overline{\varphi_4'(Z)}$ implies that $(\varphi_1(\boldsymbol{x}_1), \cdots, \varphi_1(\boldsymbol{x}_T)) \notin Z$, we conclude the proof. □

Theorem E.7 is stated for complex sample points. Next we extend it to the real case.

Lemma E.8. *In Lemma E.1, if we assume in addition that* $\dim_{\mathbb{R}} V_{\mathbb{R}} = \dim V$*, then the conclusion can be enhanced to ensure that* Z *is a real subvariety and* $\dim_{\mathbb{R}} Z < \dim_{\mathbb{R}} V_{\mathbb{R}}$.

Lemma E.9. *Let* $V \subset \mathbb{R}^n$ *be a (Zariski) closed proper subset,* $V \neq \mathbb{R}^n$*. Then* V *is a null set.*

The proof of Lemma E.9 is the same as that of Lemma E.6.

Theorem E.10. *We can additionally assume* $\boldsymbol{x}_i \in \mathbb{R}^d$ *in Theorem E.7.*

Proof. We verify that $\dim V = \dim_{\mathbb{R}} V_{\mathbb{R}}$, where V is defined in the proof of Theorem E.7, but this follows clearly by [12, Corollary 2.8.2]. We conclude the proof by applying Lemma E.8. □

Finally, we apply Theorem E.10 to two concrete classes of polynomials, namely Examples G.4 and G.5. For Example G.4, we construct its variety of tensorization of $\mathcal{R}_{\mathcal{V}}$ as follows. We first construct the tensorization of each polynomial. We define

$$\mathbf{W}_f = \sum_{i=1}^{r} a_i \begin{pmatrix}1\\ \boldsymbol{w}_i\end{pmatrix}^{\otimes p_i} \otimes \begin{pmatrix}1\\ 0\end{pmatrix}^{\otimes (p-p_i)}. \tag{14}$$

Next we construct the variety of tensorization W. Consider the map $\varphi_2 : (\mathbb{C}^d)^r \to \mathbb{C}^{(d+1)^p}$,

$$\varphi_2(\boldsymbol{w}_1, \cdots, \boldsymbol{w}_r) = \sum_{i=1}^{r} \begin{pmatrix}1\\ \boldsymbol{w}_i\end{pmatrix}^{\otimes p_i} \otimes \begin{pmatrix}1\\ 0\end{pmatrix}^{\otimes (p-p_i)}, \tag{15}$$

and let $Y = \mathbb{P}(\overline{\operatorname{Im} \varphi_2})$. Similar to V_i, we can prove that Y is an irredicible closed variety defined by homogeneous polynomials with $\dim Y \leq dr + 1$. Next consider the map $\varphi_2' : (\mathbb{C}^d)^{2r} \to \mathbb{C}^{(d+1)^p}$,

$$\varphi_2'(\boldsymbol{w}_1, \cdots, \boldsymbol{w}_{2r}) = \varphi_2(\boldsymbol{w}_1, \cdots, \boldsymbol{w}_r) - \varphi_2(\boldsymbol{w}_{r+1}, \cdots, \boldsymbol{w}_{2r}) \tag{16}$$

and let $W = \mathbb{P}(\overline{\operatorname{Im} \varphi_2'})$. Similar to Y, we can prove that W is an irredicible closed variety defined by homogeneous polynomials with $\dim W \leq 2dr + 1$. Together with Theorem E.10, we can conlude that the optimal action for Example G.4 can be uniquely determined using at most $2dr + 1$ samples.

For Example G.5, we construct W as follows. Let

$$\boldsymbol{U} = (\boldsymbol{w}_1 \quad \cdots \quad \boldsymbol{w}_k), \qquad q = \sum_{I \subseteq [k] : |I| \leq p} a_I x^I,$$

then we construct the tensorization of each polynomial by

$$\mathbf{W}_f = \sum_{I \subseteq [k] : |I| \leq p} a_I \bigotimes_{i \in I} \begin{pmatrix}1\\ \boldsymbol{w}_i\end{pmatrix} \otimes \begin{pmatrix}1\\ 0\end{pmatrix}^{\otimes (p-|I|)}. \tag{17}$$

Then we have $f(\boldsymbol{x}) = \langle \mathbf{W}_f, \mathbf{X}_{\boldsymbol{x}} \rangle$. To reduce the dimension of W and get better sample complexity bound, we construct in a manner slightly different from what we did for Example G.4. Consider the map $\varphi_3 : (\mathbb{C}^d)^k \times \mathbb{C}^{(k+1)^p} \to \mathbb{C}^{(d+1)^p}$,

$$(\boldsymbol{w}_1, \cdots, \boldsymbol{w}_k) \times (a_I : I \subseteq [k], |I| \leq p) \mapsto \mathbf{W}_f, \tag{18}$$

where $\mathbf{W}_f$ is as defined in (17). Let $Y = \mathbb{P}(\overline{\operatorname{Im} \varphi_3})$ and $W = \mathbb{P}(\overline{\operatorname{Im} \varphi_3 - \operatorname{Im} \varphi_3})$. We end up with $\dim Y = \leq dk + (k+1)^p + 1, \dim W \leq 2(dk + (k+1)^p) + 1$. So we conlude that the optimal action for Example G.5 can be uniquely determined using at most $2dk + 2(k+1)^p + 1$ samples.

F Omitted Proof for Lower Bounds with UCB Algorithms

In this section, we provide the proof for the lower bounds for learning with UCB algorithms in Subsection G.1.2.

Notation Recall that we use Λ to denote the subset of the p-th multi-indices $\Lambda = \{(\alpha_1, \dots, \alpha_p) | 1 \leq \alpha_1 < \cdots < \alpha_p \leq d\}$. For an $\alpha = (\alpha_1, \dots, \alpha_p) \in \Lambda$, denote $\boldsymbol{M}_\alpha = \boldsymbol{e}_{\alpha_1} \otimes \cdots \otimes \boldsymbol{e}_{\alpha_p}$, $\boldsymbol{A}_\alpha = (\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p})^{\otimes p}$. The model space $\mathcal{M}$ is a subset of rank-1 p-th order tensors, which is defined as $\mathcal{M} = \left\{ \boldsymbol{M}_\alpha | \alpha \in \Lambda \right\}$. We define the core action set $\mathcal{A}_0$ as $\mathcal{A}_0 = \{\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p} | \alpha \in \Lambda\}$. The action set $\mathcal{A}$ is the convex hull of $\mathcal{A}_0$: $\mathcal{A} = \operatorname{conv}(\mathcal{A}_0)$. Assume that the ground-truth parameter is $\boldsymbol{M}^* = \boldsymbol{M}_{\alpha^*} \in \mathcal{M}$. At round t, the algorithm chooses an action $\boldsymbol{a}_t \in \mathcal{A}$, and gets the **noiseless** reward $r_t = r(\boldsymbol{M}^*, \boldsymbol{a}_t) = \langle \boldsymbol{M}^*, (\boldsymbol{a}_t)^{\otimes p} \rangle = \prod_{i=1}^p \langle \boldsymbol{e}_{\alpha_i^*}, \boldsymbol{a}_t \rangle$.

F.1 Proof for Theorem G.6

We introduce a lemma showing that if the action set is **restricted** to the core action set $\mathcal{A}_0$, then at least $|\mathcal{A}_0| - 1 = \binom{d}{p} - 1$ actions are needed to identify the ground-truth.

Lemma F.1. *If the actions are restricted to $\mathcal{A}_0$, then for the noiseless degree-p polynomial bandits, any algorithm needs to play at least $\binom{d}{p} - 1$ actions to determine $\boldsymbol{M}^*$ in the worst case. Furthermore, the worst-case cumulative regret at round T can be lower bounded by*

$$\mathfrak{R}(T) \geq \min\{T, \binom{d}{p} - 1\}.$$

proof of Lemma F.1. For any α and α', the reward of playing $\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p}$ when the ground-truth model is $\boldsymbol{M}'_\alpha$ is

$$\begin{aligned} \langle \boldsymbol{M}'_\alpha, (\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p})^{\otimes p} \rangle &= \prod_{i=1}^p \langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p} \rangle \\ &= \prod_{i=1}^p \mathbb{I}\{\alpha'_i \in \alpha\} \\ &= \begin{cases} 1, & \text{if } \alpha = \alpha' \\ 0, & \text{otherwise .} \end{cases} \end{aligned}$$

Hence, no matter how the algorithm adaptively chooses the actions, in the worst case $\binom{d}{p} - 1$ actions are needed to determine $\boldsymbol{M}^*$. Also notice that the reward for $\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p}$ is zero if $\alpha \neq \alpha^*$. Therefore the regret lower bound follows. □

Next, we show that even when the action set is unrestricted, any UCB algorithm fails to explore in an unrestricted way. This is because the optimistic mechanism forbids the algorithm to play an informative action that is known to be low reward for all models in the confidence set. We first recall the definition of UCB algorithms.

UCB Algorithms The UCB algorithms sequentially maintain a confidence set $\mathcal{C}_t$ after playing actions $\boldsymbol{a}_1, \dots, \boldsymbol{a}_t$. Then UCB algorithms play

$$\boldsymbol{a}_{t+1} \in \arg\max_{\boldsymbol{a} \in \mathcal{A}} \operatorname{UCB}_t(\boldsymbol{a}),$$

where

$$\operatorname{UCB}_t(\boldsymbol{a}) = \max_{\boldsymbol{M} \in \mathcal{C}_t} \langle \boldsymbol{M}, (\boldsymbol{a})^{\otimes p} \rangle.$$

proof of Theorem G.6. We prove that even if the action set is unrestrcited, the optimistic mechanism in the UCB algorithm above forces it to choose actions in the restricted action set $\mathcal{A}_0$.

Assume $\boldsymbol{M}^* = \boldsymbol{M}_{\alpha^*}$. Next we show that for all $\boldsymbol{a} \in \mathcal{A} - \mathcal{A}_0$ (where the minus sign should be understood as set difference), we have

$$\operatorname{UCB}_t(\boldsymbol{a}) < 1.$$

For all $\boldsymbol{a} \in \mathcal{A}$, since $\mathcal{A} = \mathrm{conv}(\mathcal{A}_0)$, we can write

$$\boldsymbol{a} = \sum_{\alpha \in \Lambda} p_\alpha (\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p}),$$

where $\sum_{\alpha \in \Lambda} p_\alpha = 1$ and $p_\alpha \geq 0$. Therefore,

$$\begin{aligned}
\mathrm{UCB}_t(\boldsymbol{a}) &= \max_{\boldsymbol{M} \in \mathcal{C}_t} \langle \boldsymbol{M}, (\boldsymbol{a})^{\otimes p} \rangle \\
&\leq \max_{\boldsymbol{M} \in \mathcal{M}} \langle \boldsymbol{M}, (\boldsymbol{a})^{\otimes p} \rangle \\
&= \max_{\alpha'} \langle \boldsymbol{M}_{\alpha'}, (\boldsymbol{a})^{\otimes p} \rangle \\
&= \max_{\alpha'} \prod_{i=1}^{p} \langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{a} \rangle.
\end{aligned}$$

Plug in the expression of $\boldsymbol{a}$, we have

$$\begin{aligned}
\langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{a} \rangle &= \sum_{\alpha} p_\alpha \langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p} \rangle \\
&= \sum_{\alpha} p_\alpha \mathbb{I}\{\alpha'_i \in \alpha\} \\
&\leq \sum_{\alpha} p_\alpha = 1.
\end{aligned}$$

Therefore, for any fixed $\alpha' = (\alpha'_1, \ldots, \alpha'_p)$,

$$\begin{aligned}
\prod_{i=1}^{p} \langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{a} \rangle &= \Big(\sum_{\alpha} p_\alpha \mathbb{I}\{\alpha'_1 \in \alpha\} \Big) \cdots \Big(\sum_{\alpha} p_\alpha \mathbb{I}\{\alpha'_p \in \alpha\} \Big) \\
&\leq 1,
\end{aligned}$$

where the equality holds if and only if for any $p_\alpha > 0$, $\alpha = \alpha'$, which is equivalent to $\boldsymbol{a} = \boldsymbol{e}_{\alpha'_1} + \cdots + \boldsymbol{e}_{\alpha'_p}$. Therefore, if $\boldsymbol{a} \in \mathcal{A} - \mathcal{A}_0$, for any $\alpha' \in \Lambda$, we have $\prod_{i=1}^{p} \langle \boldsymbol{e}_{\alpha'_i}, \boldsymbol{a} \rangle < 1$. This means

$$\mathrm{UCB}_t(\boldsymbol{a}) < 1.$$

Meanwhile, we can see that for the action $\boldsymbol{a}^* = \boldsymbol{e}_{\alpha_1^*} + \cdots + \boldsymbol{e}_{\alpha_p^*} \in \mathcal{A}_0$,

$$\begin{aligned}
\mathrm{UCB}_t(\boldsymbol{a}^*) &= \max_{\boldsymbol{M} \in \mathcal{C}_t} \langle \boldsymbol{M}, (\boldsymbol{a}^*)^{\otimes p} \rangle \\
&\geq \langle \boldsymbol{M}^*, (\boldsymbol{a}^*)^{\otimes p} \rangle & (\boldsymbol{M}^* \in \mathcal{C}_t) \\
&= \langle \boldsymbol{M}^*, \boldsymbol{A}_{\alpha^*} \rangle = 1.
\end{aligned}$$

Therefore, we see that $(\mathcal{A} - \mathcal{A}_0) \cap \arg\max_{\boldsymbol{a} \in \mathcal{A}} \mathrm{UCB}_t(\boldsymbol{a}) = \varnothing$, which means $\boldsymbol{a}_{t+1} \in \mathcal{A}_0$ for all $t \geq 0$. Therefore, by Lemma F.1, the theorem holds.

$\square$

F.2 $O(d)$ Actions via Solving Polynomial Equations

Firstly, we verify that the model falls into the category of Example G.5 with $k = p$. For every $\alpha \in \Lambda$, the reward of playing $\boldsymbol{a}$ when the ground-truth model is $\boldsymbol{M}_\alpha$ is

$$\langle \boldsymbol{M}_\alpha, (\boldsymbol{a})^{\otimes p} \rangle = \prod_{i=1}^{p} \langle \boldsymbol{e}_{\alpha_i}, \boldsymbol{a} \rangle,$$

which can be written as $q_0(\boldsymbol{U}_\alpha \boldsymbol{a})$, where $q_0(x_1, \ldots, x_p) = x_1 x_2 \cdots x_p$ and $\boldsymbol{U}_\alpha \in \mathbb{R}^{p \times d}$ is a matrix with $\boldsymbol{e}_{\alpha_i}$ as the i-th row.

Secondly, we show that since the ground-truth model is p-homogenous, we can extend the action set to $\mathrm{conv}(\mathcal{A}, \boldsymbol{0})$. This is because for every action of the form $c\boldsymbol{a}$, where $0 \leq c \leq 1$ and $\boldsymbol{a} \in \mathcal{A}$, the reward is c^p times the reward at $\boldsymbol{a}$. Therefore, to get the reward at $c\boldsymbol{a}$, we only need to play at $\boldsymbol{a}$ and multiply the reward by c^p.

Notice that $\mathrm{conv}(\mathcal{A}, \boldsymbol{0})$ is of positive Lebesgue measure. By Theorem G.3, we know that only $2(dk + (p+1)^p) = O(d)$ actions are needed to determine the optimal action almost surely.

G Proof of Section 3.3.2

We present the proof of Theorem 3.18 in the following.

Proof. We overload the notation and use $[d]$ to denote the set $\{e_1, e_2, \ldots, e_d\}$. The hard instances are chosen in $\Delta \cdot [d]^p$, i.e. $(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) = \Delta \cdot (\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p)$ where $(\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p) \in [d]^p$. For a group of vectors $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p \in [d]$, we use

$$\operatorname{supp}(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) := (\max_{i \in [p]} (\boldsymbol{\theta}_i)_1, \ldots, \max_{i \in [p]} (\boldsymbol{\theta}_i)_d) \in \{0,1\}^d$$

to denote the support of these vectors. We use $\boldsymbol{a}^{(t)} \in \mathbb{R}^d$ to denote the action in t-th episode.

We use $\mathbb{P}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)}$ to denote the measure on outcomes induced by the interaction of the fixed policy and the bandit paramterised by $r = \prod_{i=1}^p (\boldsymbol{\theta}_i^\top \boldsymbol{a}) + \epsilon$. Specifically, We use $\mathbb{P}_0$ to denote the measure on outcomes induced by the interaction of the fixed policy and the pure noise bandit $r = \epsilon$.

$$\begin{aligned}
&\mathfrak{R}(d,p,T) \\
&\geq \frac{1}{d^p} \sum_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) \in \Delta \cdot [d]^p} \mathbb{E}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)} \left[T\Delta^p / p^{p/2} - \sum_{t=1}^T \prod_{i=1}^p (\boldsymbol{\theta}_i^\top \boldsymbol{a}^{(t)}) \right] \\
&= \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) \in \Delta \cdot [d]^p} \left(T/p^{p/2} - \mathbb{E}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)} \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right] \right) \\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) \in \Delta \cdot [d]^p} \left(T/p^{p/2} - \mathbb{E}_0 \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right] - T \|\mathbb{P}_0 - \mathbb{P}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)}\|_{\mathrm{TV}} \right) \\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) \in \Delta \cdot [d]^p} \left(T/p^{p/2} - \mathbb{E}_0 \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right] - T \sqrt{D_{\mathrm{KL}}(\mathbb{P}_0 || \mathbb{P}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)})} \right) \\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) \in \Delta \cdot [d]^p} \left(T/p^{p/2} - \mathbb{E}_0 \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right] - T \sqrt{\Delta^{2p} \mathbb{E}_0 \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)})^2 \right]} \right) \\
&\geq \frac{\Delta^p}{d^p} \left(\frac{d^p T}{p^{\frac{p}{2}}} - \mathbb{E}_0 \left[\sum_{(\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p) \in [d]^p} \sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right] - T d^{\frac{p}{2}} \Delta^p \sqrt{\mathbb{E}_0 \left[\sum_{(\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p) \in [d]^p} \sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)})^2 \right]} \right)
\end{aligned}$$

where the first step comes from

$$\text{Regret} \geq \mathbb{E}_{(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)} \left[T\Delta^p / p^{p/2} - \sum_{t=1}^T \prod_{i=1}^p (\boldsymbol{\theta}_i^\top \boldsymbol{a}^{(t)}) \right]$$

(the optimal action in hindsight is $\boldsymbol{a} = \operatorname{supp}(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p)/\sqrt{p}$); the second step comes from $(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) = \Delta \cdot (\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p)$ and algebra; the third step comes from $\left| \sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)}) \right| \leq T$; the fourth step comes from Pinsker's inequality; the fifth step comes from

$$\begin{aligned}
D_{\mathrm{KL}}(\mathbb{P}_0 || \mathbb{P}_{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p}) &= \mathbb{E}_0 \left[\sum_{t=1}^T D_{\mathrm{KL}} \left(N(0,1) || N(\prod_{i=1}^p (\boldsymbol{\theta}_i^\top \boldsymbol{a}^{(t)}), 1) \right) \right] \\
&= \Delta^{2p} \mathbb{E}_0 \left[\sum_{t=1}^T \prod_{i=1}^p (\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)})^2 \right]
\end{aligned}$$

and the final step comes from Jensen's inequality and algebra.

Notice that

$$
\begin{aligned}
\mathbb{E}_0\left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)})\right] &= \mathbb{E}_0\left[\sum_{(j_1,\ldots,j_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}(\boldsymbol{a}_{j_i}^{(t)})\right] \\
&= \mathbb{E}_0\left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\sum_{j=1}^{d}\boldsymbol{a}_j^{(t)})\right] \\
&\leq \mathbb{E}_0\left[\sum_{t=1}^{T}\prod_{i=1}^{p}\|\boldsymbol{a}^{(t)}\|_1\right] \\
&\leq d^{p/2}T
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{E}_0\left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}^{(t)})^2\right] &= \mathbb{E}_0\left[\sum_{(j_1,\ldots,j_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}(\boldsymbol{a}_{j_i}^{(t)})^2\right] \\
&= \mathbb{E}_0\left[\sum_{t=1}^{T}\prod_{i=1}^{p}\|\boldsymbol{a}^{(t)}\|_2^2\right] \\
&\leq T
\end{aligned}
$$

where we used $\|\boldsymbol{a}^{(t)}\|_2 \leq 1, \forall t \in [T]$. Therefore plugging back we have

$$
\mathfrak{R}(d,p,T) \geq \frac{\Delta^p}{d^p}\left(\frac{d^pT}{p^{\frac{p}{2}}} - d^{\frac{p}{2}}T - Td^{\frac{p}{2}}\Delta^p\sqrt{T}\right)
$$

and finally letting $\Delta^p = \sqrt{\frac{d^p}{4Tp^p}}$ leads to

$$
\mathfrak{R}(d,p,T) \geq O(\sqrt{d^pT}/p^p).
$$

$\square$

Remark G.1. *Better result $O(\sqrt{d^pT})$ holds for bandits $r = \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}_i) + \epsilon$ where $\boldsymbol{a}_i \in \mathbb{R}^d, \|\boldsymbol{a}_i\|_2 \leq 1$.*

For completeness, we show the proof of the above remark.

Proof. We overload the notation and use $[d]$ to denote the set $\{e_1, e_2, \ldots, e_d\}$. The hard instances are chosen in $\Delta \cdot [d]^p$, i.e. $(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p) = \Delta \cdot (\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p)$ where $(\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_p) \in [d]^p$. We use $\boldsymbol{a}_i^{(t)} \in \mathbb{R}^d$ to denote the i-th action in t-th episode, where $i \in [p], t \in [T]$.

We use $\mathbb{P}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}$ to indicate the measure on outcomes induced by the interaction of the fixed policy and the bandit paramterised by $r = \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}_i) + \epsilon$. Specifically, We use $\mathbb{P}_0$ to indicate the measure on outcomes induced by the interaction of the fixed policy and the pure noise bandit $r = \epsilon$.

$$
\begin{aligned}
&\mathfrak{R}(d,p,T)\\
&\geq \frac{1}{d^p} \sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p} \mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)} \left[T\Delta^p - \sum_{t=1}^{T}\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}_i^{(t)}) \right]\\
&= \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p} \left(T - \mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)} \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] \right)\\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p} \left(T - \mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] - T\|\mathbb{P}_0 - \mathbb{P}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\|_{\mathrm{TV}} \right)\\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p} \left(T - \mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] - T\sqrt{D_{\mathrm{KL}}(\mathbb{P}_0||\mathbb{P}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)})} \right)\\
&\geq \frac{\Delta^p}{d^p} \sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p} \left(T - \mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] - T\sqrt{\Delta^{2p}\mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)})^2 \right]} \right)\\
&\geq \frac{\Delta^p}{d^p} \left(d^pT - \mathbb{E}_0 \left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p} \sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] - Td^{\frac{p}{2}}\Delta^p \sqrt{\mathbb{E}_0 \left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p} \sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)})^2 \right]} \right)
\end{aligned}
$$

where the first step comes from

$$
\text{Regret} \geq \mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)} \left[T\Delta^p - \sum_{t=1}^{T}\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}_i^{(t)}) \right]
$$

(the optimal action in hindsight is $\boldsymbol{a}_i = \widehat{\boldsymbol{\theta}}_i$); the second step comes from $(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p) = \Delta \cdot (\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)$ and algebra; the third step comes from $\left|\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)})\right| \leq T$; the fourth step comes from Pinsker's inequality; the fifth step comes from

$$
\begin{aligned}
D_{\mathrm{KL}}(\mathbb{P}_0||\mathbb{P}_{\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p}) &= \mathbb{E}_0 \left[\sum_{t=1}^{T} D_{\mathrm{KL}} \left(N(0,1)||N(\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top \boldsymbol{a}_i^{(t)}),1) \right) \right]\\
&= \Delta^{2p}\mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)})^2 \right]
\end{aligned}
$$

and the final step comes from Jensen's inequality and algebra.

Notice that

$$
\begin{aligned}
\mathbb{E}_0 \left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p} \sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)}) \right] &= \mathbb{E}_0 \left[\sum_{(j_1,\ldots,j_p)\in[d]^p} \sum_{t=1}^{T}\prod_{i=1}^{p} \left((\boldsymbol{a}_i^{(t)})_{j_i} \right) \right]\\
&= \mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p} \left(\sum_{j=1}^{d} (\boldsymbol{a}_i^{(t)})_j \right) \right]\\
&\leq \mathbb{E}_0 \left[\sum_{t=1}^{T}\prod_{i=1}^{p} \|\boldsymbol{a}_i^{(t)}\|_1 \right]\\
&\leq d^{p/2}T
\end{aligned}
$$

and

$$\mathbb{E}_0\left[\sum_{(\widehat{\boldsymbol{\theta}}_1,\ldots,\widehat{\boldsymbol{\theta}}_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}(\widehat{\boldsymbol{\theta}}_i^\top \boldsymbol{a}_i^{(t)})^2\right] = \mathbb{E}_0\left[\sum_{(j_1,\ldots,j_p)\in[d]^p}\sum_{t=1}^{T}\prod_{i=1}^{p}\left((\boldsymbol{a}_i^{(t)})_{j_i}\right)^2\right]$$
$$= \mathbb{E}_0\left[\sum_{t=1}^{T}\prod_{i=1}^{p}\|\boldsymbol{a}_i^{(t)}\|_2^2\right]$$
$$\leq T$$

where we used $\|\boldsymbol{a}_i^{(t)}\|_2 \leq 1, \forall t \in [T]$. Therefore plugging back we have

$$\mathfrak{R}(d,p,T) \geq \frac{\Delta^p}{d^p}\left(d^pT - d^{\frac{p}{2}}T - Td^{\frac{p}{2}}\Delta^p\sqrt{T}\right)$$

and finally letting $\Delta^p = \sqrt{\frac{d^p}{4T}}$ leads to

$$\mathfrak{R}(d,p,T) \geq O(\sqrt{d^pT}).$$

□

We present the proof of Theorem 3.19 in the following.

Proof. Denote the optimal action in hindsight as $\boldsymbol{a}^* = \mathrm{supp}(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)/\sqrt{p}$. From the proof of Theorem 3.18 we know that if $T \leq \frac{1}{4p^p}\cdot\frac{d^p}{\Delta^{2p}}$, then

$$\frac{1}{d^p}\sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p}\mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\left[\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top\boldsymbol{a}^*) - \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top\boldsymbol{a}^{(t)})\right]$$
$$\geq \frac{\Delta^p}{d^p}\left(\frac{d^p}{p^{\frac{p}{2}}} - d^{\frac{p}{2}} - d^{\frac{p}{2}}\Delta^p\sqrt{T}\right)$$
$$\geq \frac{\Delta^p}{4p^{\frac{p}{2}}}$$
$$\geq \frac{1}{4}\cdot\frac{1}{d^p}\sum_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)\in\Delta\cdot[d]^p}\mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\left[\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top\boldsymbol{a}^*)\right]$$

which indicates the following

$$\inf_{\pi}\sup_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\mathbb{E}_{(\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_p)}\left[\frac{3}{4}\cdot\prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top\boldsymbol{a}^*) - \prod_{i=1}^{p}(\boldsymbol{\theta}_i^\top\boldsymbol{a}^{(t)})\right] \geq 0.$$

□

G.1 Noiseless Polynomial Reward

In this subsection, we study the regret bounds for learning bandits with **noiseless** polynomial rewards. First we present the definition of admissible polynomial families.

Definition G.2 (Admissible Polynomial Family)**.** For $\boldsymbol{a}\in\mathbb{R}^d$, define $\widetilde{\boldsymbol{a}} = [1,\boldsymbol{a}^\top]^\top$. For a algebraic variety $\mathcal{V}\subseteq(\mathbb{R}^{d+1})^{\otimes p}$, define $\mathcal{R}_{\mathcal{V}} := \{r_{\boldsymbol{\theta}}(\boldsymbol{a}) = \langle\boldsymbol{\theta},\widetilde{\boldsymbol{a}}^{\otimes p}\rangle : \boldsymbol{\theta}\in\mathcal{V}\}$ as the polynomial family with parameters in $\mathcal{V}$. We define the dimension of the family $\mathcal{R}_{\mathcal{V}}$ as the algebraic dimension of $\mathcal{V}$. Next, define $\mathcal{X} := \left\{\widetilde{\boldsymbol{a}}^{\otimes p} : \boldsymbol{a}\in\mathbb{R}^d\right\}$. An polynomial family $\mathcal{R}_{\mathcal{V}}$ is said to be admissible[5] w.r.t. $\mathcal{X}$ if for any $\boldsymbol{\theta}\in\mathcal{V}$, $\dim(\mathcal{X}\cap\{X\in\mathcal{X} : \langle X,\boldsymbol{\theta}\rangle = 0\rangle\}) < \dim(\mathcal{X}) = d$.

G.1.1 Upper Bounds via Solving Polynomial Equations

We show that if the action set $\mathcal{A}$ is of positive measure with respect to the Lebesgue measure μ, then by playing actions randomly, we can uniquely solve for the ground-truth reward function $r_{\boldsymbol{\theta}}(\boldsymbol{a})$ almost surely with samples of size that scales with the intrinsic algebraic dimension of $\mathcal{V}$, provided that $\mathcal{V}$ is an admissible algebraic variety.

[5] Intuitively, admissibility means the dimension of $\mathcal{X}$ decreases by one when there is an additional linear constraint $\langle\boldsymbol{\theta}, X\rangle = 0$

Theorem G.3. *Assume that the reward function class is an admissible polynomial family $\mathcal{R}_{\mathcal{V}}$, and the maximum reward is upper bounded by 1. If $\mu(\mathcal{A}) > 0$, where μ is the Lebesgue meaure, then by randomly sample actions $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_T$ from $\mathbb{P}_{\boldsymbol{a}\sim\mathcal{N}(0,I_d)}(\cdot|\boldsymbol{a} \in \mathcal{A})$, when $T \geq 2\dim(\mathcal{V})$, we can uniquely solve for the ground-truth $\boldsymbol{\theta}$ and thus determine the optimal action almost surely. Therefore, the cumulative regret at round T can be bounded as*

$$\mathfrak{R}(T) \leq \min\{T, 2\dim(\mathcal{V})\}.$$

We state two important examples of admissible polynomial families with $O(d)$ dimensions.

Example G.4 (low-rank polynomials)**.** The function class $\mathcal{R}_{\mathcal{V}}$ of possibly inhomogeneous degree-p polynomials with k summands $\mathcal{R}_{\mathcal{V}} = \{r(\boldsymbol{a}) = \sum_{i=1}^{k} \lambda_i \langle \boldsymbol{v}_i, \boldsymbol{a}\rangle^{p_i} \mid \lambda_i \in \mathbb{R}, \boldsymbol{v}_i \in \mathbb{R}^d\}$ is admissible with $\dim(\mathcal{R}_{\mathcal{V}}) \leq dk$, where $p = \max\{p_i\}$. Neural network with monomial/polynomial activation functions are low-rank polynomials.

Example G.5 ([18])**.** The function class $\mathcal{R}_{\mathcal{V}} = \{r(\boldsymbol{a}) = q(\boldsymbol{U}\boldsymbol{a}) \mid \boldsymbol{U} \in \mathbb{R}^{k\times d}, \deg q(\cdot) \leq p\}$ is admissible with $\dim(\mathcal{V}) \leq dk + (k+1)^p$.

G.1.2 Lower Bounds with UCB Algorithms

In this subsection, we construct a hard bandit problem where the rewards are noiseless degree-p polynomial, and show that any UCB algorithm needs at least $\Omega(d^p)$ actions to learn the optimal action. On the contrary, Theorem G.3 shows that by playing actions randomly, we only need $2(dk + (p+1)^p) = O(d)$ actions.

Hard Case Construction Let $\boldsymbol{e}_i$ denotes the i-th standard orthonormal basis of $\mathbb{R}^d$, i.e., $\boldsymbol{e}_i$ has only one 1 at the i-th entry and 0's for other entries. We define a p-th multi-indices set Λ as $\Lambda = \{(\alpha_1, \ldots, \alpha_p)|1 \leq \alpha_1 < \cdots < \alpha_p \leq d\}$. For an $\alpha = (\alpha_1, \ldots, \alpha_p) \in \Lambda$, denote $\boldsymbol{M}_\alpha = \boldsymbol{e}_{\alpha_1} \otimes \cdots \otimes \boldsymbol{e}_{\alpha_p}$. Then the model space $\mathcal{M}$ is defined as $\mathcal{M} = \left\{\boldsymbol{M}_\alpha | \alpha \in \Lambda\right\}$, which is a subset of rank-1 p-th order tensors. The action set $\mathcal{A}$ is defined as $\mathcal{A} = \mathrm{conv}(\{\boldsymbol{e}_{\alpha_1} + \cdots + \boldsymbol{e}_{\alpha_p} | \alpha \in \Lambda\})$. Assume that the ground-truth parameter is $\boldsymbol{M}^* = \boldsymbol{M}_{\alpha^*} \in \mathcal{M}$. The noiseless reward $r_t = r(\boldsymbol{M}^*, \boldsymbol{a}_t) = \langle \boldsymbol{M}^*, (\boldsymbol{a}_t)^{\otimes p}\rangle = \prod_{i=1}^{p} \langle \boldsymbol{e}_{\alpha_i^*}, \boldsymbol{a}_t\rangle$ is a polynomial of $\boldsymbol{a}_t$ and falls into the case of Example G.5.

UCB Algorithms The UCB algorithms sequentially maintain a confidence set $\mathcal{C}_t$ after playing actions $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_t$. Then UCB algorithms play $\boldsymbol{a}_{t+1} \in \arg\max_{\boldsymbol{a}\in\mathcal{A}} \mathrm{UCB}_t(\boldsymbol{a})$, where $\mathrm{UCB}_t(\boldsymbol{a}) = \max_{\boldsymbol{M}\in\mathcal{C}_t} \langle \boldsymbol{M}, (\boldsymbol{a})^{\otimes p}\rangle$.

Theorem G.6. *Assume that for each $t \geq 0$, the confidence set $\mathcal{C}_t$ contains the ground-truth model, i.e., $\boldsymbol{M}^* \in \mathcal{C}_t$. Then for the noiseless degree-p polynomial bandits, any UCB algorithm needs to play at least $\binom{d}{p} - 1$ actions to distinguish models in $\mathcal{M}$. Furthermore, the worst-case cumulative regret at round T can be lower bounded by*

$$\mathfrak{R}(T) \geq \min\{T, \binom{d}{p} - 1\}.$$

Theorem G.6 shows the failure of the optimistic mechanism, which forbids the algorithm to play an informative action that is known to be of low reward for all models in the confidence set. On the contrary, the reward function class falls into the form of $q(\boldsymbol{U}\boldsymbol{a})$, therefore, by playing actions randomly[6], we only need $O(d)$ actions as Theorem G.3 suggests.

[6] Careful readers may notice that $\mathcal{A}$ is of measure zero in this setting. However, since the reward function is a homogenous polynomial of degree p, we can actually obtain the rewards on $\mathrm{conv}(\mathcal{A}, \mathbf{0})$, which is of positive measure.