
Pure Exploration in Kernel and Neural Bandits

Yinglun Zhu*

Department of Computer Sciences
University of Wisconsin-Madison
Madison, WI 53706
yinglun@cs.wisc.edu

Dongruo Zhou*

Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
drzhou@cs.ucla.edu

Ruoxi Jiang*

Department of Computer Science
University of Chicago
Chicago, IL 60637
roxie62@uchicago.edu

Quanquan Gu

Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
qgu@cs.ucla.edu

Rebecca Willett

Department of Statistics and Computer Science
University of Chicago
Chicago, IL 60637
willett@uchicago.edu

Robert Nowak

Department of Electrical and Computer Engineering
University of Wisconsin-Madison
Madison, WI 53706
rdnowak@wisc.edu

Abstract

We study pure exploration in bandits, where the dimension of the feature representation can be much larger than the number of arms. To overcome the curse of dimensionality, we propose to adaptively embed the feature representation of each arm into a lower-dimensional space and carefully deal with the induced model misspecifications. Our approach is conceptually very different from existing works that can either only handle low-dimensional linear bandits or passively deal with model misspecifications. We showcase the application of our approach to two pure exploration settings that were previously under-studied: (1) the reward function belongs to a possibly infinite-dimensional Reproducing Kernel Hilbert Space, and (2) the reward function is nonlinear and can be approximated by neural networks. Our main results provide sample complexity guarantees that only depend on the effective dimension of the feature spaces in the kernel or neural representations. Extensive experiments conducted on both synthetic and real-world datasets demonstrate the efficacy of our methods.

1 Introduction

Pure exploration in bandits [11, 12, 6] has been extensively studied in machine learning. Consider a set of arms, where each arm is associated with an unknown reward distribution. The goal is to

*Equal contribution

approximately identify the optimal arm using as few samples as possible. Applications of bandit pure exploration range from medical domains [3] to online content recommendation [40].

Despite the popularity of bandit pure exploration, it was previously mainly studied in two relatively restrictive settings: (1) the standard multi-armed bandit setting [22, 19, 9, 18], where the expected rewards among arms are completely unrelated to each other, and (2) the (generalized) linear bandit setting [38, 13, 10, 26], where the expected rewards are assumed to be linearly parameterized by some unknown weight vector. The standard multi-armed bandit setting fails to deal with large arm sets, and the linear bandit setting suffers from both model misspecifications (due to its simplified linear form) and the curse of dimensionality in the high-dimensional setting. Pure exploration is also studied in continuous spaces. However, guarantees therein scale exponentially with dimension [31, 4].

In this paper, we generalize bandit pure exploration to the nonlinear and high-dimensional settings. More specifically, we study the following two settings: (1) the rewards of arms are parameterized by a function belonging to a Reproducing Kernel Hilbert Space (RKHS), and (2) the rewards of arms are nonlinear functions that can be approximated by an overparameterized neural network. Problems in these two settings are often high-dimensional in nature. To overcome the curse of dimensionality, we propose to adaptively embed each arm’s feature representation in a lower-dimensional space and carefully deal with the induced misspecifications. Note that our approach is conceptually very different from all existing work dealing with model misspecifications: they assume the existence of misspecifications and address it in the original space (thus dealing with model misspecifications in a passive way) [28, 7]. On the other hand, we deliberately induce (acceptable) misspecifications to embed arms into lower-dimensional spaces and thus overcome the curse of dimensionality.

1.1 Contribution and Outline

We make the following main contributions:

- In Section 3, we introduce the idea of adaptive embedding to avoid the curse of dimensionality. The induced model misspecifications are carefully handled, which is novel in the bandit pure exploration setting. The sample complexity is theoretically analyzed and we relate the instance-dependent sample complexity to the complexity of a closely-related linear bandit problem without model misspecification. As a by-product, our algorithm can also be applied to constrained high-dimensional linear bandit pure exploration to reduce sample complexity.
- In Section 4, we specialize the adaptive embedding scheme to pure exploration in an RKHS. We construct feature mappings from eigenfunctions and eigenvalues of the associated kernel. The effective dimension of the kernel is analyzed, and we provide sample complexity guarantees in terms of the eigenvalue decay of the associated kernel. We rely on a *known* kernel in this setting.
- In Section 5, we further extend our adaptive embedding scheme to pure exploration with a general nonlinear reward function and model the reward function with an over-parameterized neural network. Sample complexity guarantees are provided with respect to the eigenvalue decay of the associated Neural Tangent Kernel. To the best of our knowledge, this provides the first theoretically founded pure exploration algorithm with a neural network approximation.
- In Section 6, we conduct extensive experiments on both synthetic and real-world datasets to confirm the efficacy of our proposed algorithms. We conclude our paper in Section 7 with open problems.

1.2 Related Work

The bandit pure exploration problem has a long history, dating back to the seminal work by Bechhofer [5], Paulson et al. [32]. One classical objective of pure exploration is Best Arm Identification (BAI), where the goal is to identify the best arm using as few samples as possible [22, 19, 9, 15]. To make it applicable to a large action space, the BAI problem is also extensively studied as the good arm identification problem, where the goal is to identify an ϵ -optimal arm [11, 12, 20, 21, 34, 23, 30].

The pure exploration problem in linear bandits is initially analyzed in Soare et al. [38], where optimal experimental design [27] is applied to guide the allocation of samples. Other approaches dealing with linear bandits, with various sample complexity guarantees, include adaptive sampling [45] and an approach called track-and-stop [10]. Constrained linear bandit pure exploration is also commonly studied with additional assumptions on the reward parameters [41, 10]. We note that the

track-and-stop approach only achieves optimal instance-dependent sample complexity in the regime where the confidence parameter approaches 0, but fails to do so in the moderate confidence regime. Fiez et al. [13] propose an elimination-based algorithm (with optimal design) that achieves (nearly) instance-dependent sample complexity. such algorithm is further generalized to the combinatorial bandit setting [24].

Learning with model misspecifications was recently introduced in bandit learning, with the primary emphasis placed on the regret minimization problem [16, 28, 14]. A very recent independent work studies pure exploration in kernel bandits with misspecifications [7]; both their and our algorithms follow the framework of RAGE [13] and draw inspiration from [28]. Camilleri et al. [7] propose a robust estimator that works in high-dimensional spaces and also explore the project-then-round idea through *regularized least squares*. Our algorithms adaptively embed actions into lower dimensional spaces according to some error tolerances (different embeddings from round to round); our rounding and elimination steps are thus computed only with respect to lower-dimensional embeddings. We additionally study the pure exploration problem with an overparameterized neural network. As mentioned before, our approach is also conceptually different from existing ones: rather than passively dealing with model misspecifications in its original representation, we deliberately and adaptively embed arms into a lower-dimensional space to avoid the curse of dimensionality; the induced model misspecifications are also carefully dealt with in our algorithms.

2 Problem setting

We introduce the general setting and notations for pure exploration in bandits. Consider a set of arms $\mathcal{X} \subseteq \mathbb{R}^D$ where the number of arms $|\mathcal{X}| = K$ is possibly very large. We use an unknown function $h : \mathcal{X} \rightarrow [-1, 1]$ to represent the true reward of each arm. A noisy feedback $h(\mathbf{x}) + \xi$ is observed after each sample arm \mathbf{x} , where the noise ξ is assumed to be 1-sub-Gaussian. The learner is allowed to allocate her samples based on previously collected information, and the goal is to approximately identify an approximately optimal arm using as few samples as possible. Let $\mathbf{x}_* = \arg \max_{\mathbf{x} \in \mathcal{X}} h(\mathbf{x})$ denote the optimal arm among \mathcal{X} . We aim at developing (ϵ, δ) -PAC guarantees: for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the algorithm outputs an ϵ -optimal arm $\hat{\mathbf{x}}$ such that $h(\hat{\mathbf{x}}) \geq h(\mathbf{x}_*) - \epsilon$ using a finite number of samples. The performance of the algorithm is measured by its sample complexity, i.e., the number of samples pulled before it stops and recommends a candidate arm.

Notations. We define $\Delta_{\mathbf{x}} = h(\mathbf{x}_*) - h(\mathbf{x})$ as the sub-optimality gap of arm \mathbf{x} . We use the notations $\mathcal{S}_k := \{\mathbf{x} \in \mathcal{X} : \Delta_{\mathbf{x}} < 4 \cdot 2^{-k}\}$ (with $\mathcal{S}_1 = \mathcal{X}$). We consider feature mappings of the form $\psi_d(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^d$, and define $\psi_d(\mathcal{X}) = \{\psi_d(\mathbf{x}) : \mathbf{x} \in \mathcal{X}\}$. We use $\Lambda_{\mathcal{X}} = \{\lambda \in \mathbb{R}^{|\mathcal{X}|} : \sum_{\mathbf{x} \in \mathcal{X}} \lambda_{\mathbf{x}} = 1, \lambda_{\mathbf{x}} \geq 0\}$ to denote the $(|\mathcal{X}| - 1)$ -dimensional probability simplex over arms in \mathcal{X} ; and set $\mathbf{A}_{\psi_d}(\lambda) = \sum_{\mathbf{x} \in \mathcal{X}} \lambda_{\mathbf{x}} \psi_d(\mathbf{x}) \psi_d(\mathbf{x})^\top$.² We use $\|\mathbf{x}\|_{\mathbf{A}} = \sqrt{\mathbf{x}^\top \mathbf{A} \mathbf{x}}$ to represent the Mahalanobis norm. We also define $\mathcal{Y}(\mathcal{V}) = \{\mathbf{v} - \mathbf{v}' : \mathbf{v}, \mathbf{v}' \in \mathcal{V}\}$ for any set \mathcal{V} . For a matrix $\mathbf{H} \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$, we use $\mathbf{H}(\mathbf{x}, \mathbf{x}')$ to denote the entry of \mathbf{H} which locates at row \mathbf{x} and column \mathbf{x}' .

3 Bandit pure exploration with adaptive embedding

We introduce the idea of bandit pure exploration with adaptive embedding, which can be viewed as an approach that actively trades off sample complexity with accuracy guarantees: we adaptively embed the feature representation into lower-dimensional spaces to avoid the curse of dimensionality, and conduct pure exploration with *misspecified* linear bandits. The embedding dimensions are carefully selected so that we can identify an ϵ -optimal arm.

We formalize the idea as follows. For any $d \in \mathbb{N}$, we assume the existence of a feature mapping $\psi_d : \mathcal{X} \rightarrow \mathbb{R}^d$ and a unknown reward vector $\boldsymbol{\theta}_d \in \mathbb{R}^d$ such that, for any $\mathbf{x} \in \mathcal{X}$,

$$h(\mathbf{x}) = \langle \psi_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle + \eta_d(\mathbf{x}),$$

where $\eta_d(\mathbf{x})$ represents the induced approximation error on arm \mathbf{x} with respect to the low-dimensional embedding $\psi_d(\cdot)$. Without loss of generality, we assume that the action set \mathcal{X} is rich enough so that $\psi_d(\mathcal{X})$ spans \mathbb{R}^d for d considered in this paper. Otherwise, one can always project feature

²A generalized inversion is used for singular matrices. We refer to Appendix A.1 for detailed discussion.

representations $\psi_d(\mathcal{X})$ into an even lower-dimensional space without losing information in the linear component.

We use $\tilde{\gamma} : \mathbb{N} \rightarrow \mathbb{R}$ to represent the misspecification level: an upper bound of the induced approximation error across all arms, i.e., $\max_{\mathbf{x} \in \mathcal{X}} |\eta_d(\mathbf{x})| \leq \tilde{\gamma}(d)$. We define $g(d, \zeta) := (1 + \zeta) \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{X}))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2$, which represents the optimal value of a transductive design among embeddings in \mathbb{R}^d . We define $\gamma(d) := (16 + 8\sqrt{g(d, \zeta)}) \tilde{\gamma}(d)$, which quantifies the sub-optimality gap of the identified arm in the worst case. One can easily show $\gamma(d) \leq O(\tilde{\gamma}(d)\sqrt{d})$ through Kiefer-Wolfowitz theorem [27].

Remark 1. We believe such optimality guarantees are un-improvable in general. In fact, a hard instance is constructed in [28] showing that, even with deterministic feedback, identifying a $o(\tilde{\gamma}(d)\sqrt{d})$ -optimal arm requires sample complexity exponential on d . On the other side, identifying a $\Omega(\tilde{\gamma}(d)\sqrt{d})$ -optimal only requires sample complexity polynomially in d . Such a sharp trade-off between optimality and sample complexity motivates our definition of $\gamma(d)$ (and our sample complexities are polynomially in d).

We assume the knowledge of both the feature mapping $\psi_d(\cdot)$ and the error function $\tilde{\gamma}(\cdot)$. This assumption is mild since one can explicitly construct/analyze $\psi_d(\cdot)$ and $\tilde{\gamma}(\cdot)$ in many cases (as discussed in Section 3.2, Section 4 and Section 5). We further assume that $\gamma(d)$ can be made arbitrarily small for large enough d . Such assumption trivially holds if the rewards are perfectly explained for d large enough, i.e., $\tilde{\gamma}(d) = 0$. We now define the *effective dimension* with respect to $\gamma(d)$ (induced from feature mapping $\psi_d(\cdot)$) as follows.

Definition 1. For any $\epsilon > 0$, we define the effective dimension as $d_{\text{eff}}(\epsilon) := \min\{d \geq 1 : \gamma(d) \leq \epsilon\}$.

In general, the effective dimension $d_{\text{eff}}(\epsilon)$ captures the smallest dimension one needs to explore in order to identify an ϵ -optimal arm. Similar notions have been previously used in regret minimization settings [42, 43]. One can easily see that $d_{\text{eff}}(\epsilon_1) \leq d_{\text{eff}}(\epsilon_2)$ as long as $\epsilon_1 \geq \epsilon_2$.

3.1 Algorithm and analysis

Algorithm 1 follows the framework of RAGE [13] to eliminate arms with sub-optimality gap $\geq O(2^{-k})$ at the k -th iteration. It runs for $n = O(\log(1/\epsilon))$ iterations and identifies an ϵ -optimal arm. We use optimal experimental design to select arms for faster elimination. For any fixed design $\lambda \in \Lambda_{\mathcal{X}}$, with $N \geq r_d(\zeta)$ samples and an approximation factor ζ (with default value $\zeta \in [1/10, 1/4]$), the rounding procedure in \mathbb{R}^d , i.e., $\text{ROUND}(\lambda, N, d, \zeta)$, outputs a discrete allocation $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ satisfying

$$\max_{\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{X}))} \|\mathbf{y}\|_{\left(\sum_{i=1}^N \psi_d(\mathbf{x}_i)\psi_d(\mathbf{x}_i)^\top\right)^{-1}}^2 \leq (1 + \zeta) \max_{\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{X}))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2 / N. \quad (1)$$

Efficient rounding procedures exist with $r_d(\zeta) = \frac{d^2+d+2}{\zeta}$ [33] or $r_d(\zeta) = \frac{180d}{\zeta^2}$ [1, 13]. We refer reads to [13, 33, 1] for detailed rounding algorithms and the associated computational complexities.

Unlike RAGE that directly works in the original high-dimensional space, Algorithm 1 adaptively embeds arms into lower-dimensional spaces and carefully deals with the induced misspecifications. More specifically, the embedding dimension d_k is selected as the smallest dimension such that the induced error term ϵ_k is well controlled, i.e., $\epsilon_k \leq O(2^{-k})$. The embedding is more aggressive at initial iterations due to larger error tolerance; The embedding dimension selected at the last iteration is (roughly) $d_{\text{eff}}(\epsilon)$ to identify an ϵ -optimal arm. The number of samples required for each iteration N_k is with respect to an experimental design in the lower-dimensional space *after embedding*. The ROUND procedure also becomes more efficient due to the embedding. Before stating our main theorem, we introduce the following complexity measure [38, 13, 10], which quantifies the hardness of the pure exploration problem (with respect to mapping $\psi_d(\cdot)$).

$$\rho_d^*(\epsilon) := \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\psi_d(\mathbf{x}_*) - \psi_d(\mathbf{x})\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2}{\max\{h(\mathbf{x}_*) - h(\mathbf{x}), \epsilon\}^2}.$$

Theorem 1. With probability of at least $1 - \delta$, Algorithm 1 correctly outputs an ϵ -optimal arm with sample complexity upper bounded by

$$640 \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \left(\left(k \rho_{d_k}^* (2^{2-k}) \log(k^2 |\mathcal{X}|^2 / \delta) \right) + (r_{d_k}(\zeta) + 1) \right) \leq \tilde{O}(d_{\text{eff}}(\epsilon) \cdot \max\{\Delta_{\min}, \epsilon\}^{-2}),$$

Algorithm 1 Arm Elimination with Adaptive Embedding and Induced Misspecification

Input: Action set \mathcal{X} , confidence parameter δ , accuracy parameter ϵ and rounding approximation factor ζ .

- 1: Set $n = \lceil \log_2(2/\epsilon) \rceil$ and $\widehat{\mathcal{S}}_1 = \mathcal{X}$.
 - 2: **for** $k = 1, 2, \dots, n$ **do**
 - 3: Set $\delta_k = \delta/k^2$, $d_k = d_{\text{eff}}(4 \cdot 2^{-k})$.
 - 4: Select feature representation $\psi_{d_k}(\cdot)$, and calculate the induced misspecification level $\tilde{\gamma}(d_k)$.
Set $r_{d_k}(\zeta) = O(d_k/\zeta^2)$ as the number of samples needed for ROUND in \mathbb{R}^{d_k} .
 - 5: Set λ_k and τ_k be the design and the value of the following optimization problem

$$\inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi_{d_k}(\widehat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_{d_k}(\lambda)}^{-1}}^2.$$
 - 6: Set $\epsilon_k = 2\tilde{\gamma}(d_k) + \tilde{\gamma}(d_k)\sqrt{(1+\zeta)\tau_k}$,
and $N_k = \max\{[(2^{-k} - \epsilon_k)^{-2}2(1+\zeta)\tau_k \log(|\widehat{\mathcal{S}}_k|^2/\delta_k)], r_{d_k}(\zeta)\}$.
 - 7: Get $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\} = \text{ROUND}(\lambda_k, N_k, d_k, \zeta)$.
 - 8: Pull arms $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ and receive rewards $\{y_1, \dots, y_{N_k}\}$.
 - 9: Set $\widehat{\boldsymbol{\theta}}_k = \mathbf{A}_k^{-1}\mathbf{b}_k$, where $\mathbf{A}_k = \sum_{i=1}^{N_k} \psi_{d_k}(\mathbf{x}_i)\psi_{d_k}(\mathbf{x}_i)^\top$ and $\mathbf{b}_k = \sum_{i=1}^{N_k} \psi_{d_k}(\mathbf{x}_i)y_i$.
 - 10: Eliminate arms with respect to criteria

$$\widehat{\mathcal{S}}_{k+1} = \widehat{\mathcal{S}}_k \setminus \{\mathbf{x} \in \widehat{\mathcal{S}}_k : \exists \mathbf{x}' \text{ such that } (\psi_{d_k}(\mathbf{x}') - \psi_{d_k}(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k \geq \omega_k(\psi_{d_k}(\mathbf{x}') - \psi_{d_k}(\mathbf{x}))\},$$
 where $\omega_k(\mathbf{y}) = \epsilon_k + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}}\sqrt{2\log(|\widehat{\mathcal{S}}_k|^2/\delta_k)}$.
 - 11: **end for**
- Output:** Output any arm in $\widehat{\mathcal{S}}_{n+1}$.
-

where $d_k = d_{\text{eff}}(4 \cdot 2^{-k}) \leq d_{\text{eff}}(\epsilon)$ since $4 \cdot 2^{-k} \geq \epsilon$ when $k \leq \lceil \log_2(2/\epsilon) \rceil$.

The rounding term $r_d(\zeta)$ commonly appears in the sample complexity of linear bandits [13, 24]; and our rounding term is with respect to the lower-dimensional space after embedding, which only scales with d_k rather than the ambient dimension. To further interpret the complexity, we define another complexity measure of a closely related linear bandit problem in the low-dimensional space and without model misspecifications.

$$\tilde{\rho}_d^*(\epsilon) := \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\psi_d(\mathbf{x}_*) - \psi_d(\mathbf{x})\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2}{\max\{\langle \psi_d(\mathbf{x}_*) - \psi_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle, \epsilon\}^2},$$

where $\langle \psi_d(\mathbf{x}_*) - \psi_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle$ on the denominator represents the sub-optimality gap characterized by the linear component rather than the true sub-optimality gap $h(\mathbf{x}_*) - h(\mathbf{x})$. The relation between $\rho^*(\epsilon)$ and $\tilde{\rho}^*(\epsilon)$ is discussed as follows.

Proposition 1. *Suppose $\max_{\mathbf{x} \in \mathcal{X}} |h(\mathbf{x}) - \langle \psi_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle| \leq \tilde{\gamma}(d)$. For any $\epsilon \geq \tilde{\gamma}(d)$, we have $\rho_d^*(\epsilon) \leq 9\tilde{\rho}_d^*(\epsilon)$. When $\tilde{\gamma}(d) < \Delta_{\min}/2$, $\tilde{\rho}_d^*(0)$ represents the sample complexity of a closely-related linear bandit problem without model misspecifications, i.e., $\tilde{\mathbf{h}}(\mathbf{x}) = \langle \psi_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle$.*

Remark 2. *When $\tilde{\gamma}(d) < \Delta_{\min}/2$, our sample complexity upper bound is relevant to the sample complexity of closely-related linear bandit problems without model misspecifications in lower-dimensional spaces. In fact, $\tilde{\rho}_d^*(0) \log(1/2.4\delta)$ is the lower bound of the corresponding linear bandit problem in \mathbb{R}^d [38, 13, 10].*

Remark 3. *Although the misspecification levels are generally known for situations considered in this paper, we also provide an algorithm that deals with unknown misspecification levels in Appendix D. Similar sample complexity guarantees are provided, but only in an unverifiably way (due to unknown misspecification levels): the algorithm starts to output ϵ -good arms after N samples, yet it doesn't know when to stop. We refer readers to [23] for details on the unverifiable sample complexity.*

3.2 Application to high-dimensional linear bandits

We apply the idea of adaptive embedding to high-dimensional linear bandits. We consider linear bandit problem of the form $\mathbf{h} = \mathbf{X}\boldsymbol{\theta}_*$ where $\mathbf{X} \in \mathbb{R}^{K \times D}$ and the i -th row of \mathbf{X} represents the

feature vector of arm \mathbf{x}_i . We assume that $\|\boldsymbol{\theta}_*\|_2 \leq C$, which is commonly studied as the constrained linear bandit problem [41, 10].

Let $\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ be the singular value decomposition (SVD) of \mathbf{X} , with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$ for some $r \leq \min\{K, D\}$. Let $u_{i,j}$ denote the (i, j) -th entry of matrix \mathbf{U} and $\mathbf{u}_{:,i}$ denote the i -th column of \mathbf{U} (similar notations for \mathbf{V}). We have

$$\mathbf{h} = \mathbf{X}\boldsymbol{\theta}_* = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top\boldsymbol{\theta}_* = \sum_{i=1}^d \sigma_i \mathbf{u}_{:,i} \mathbf{v}_{:,i}^\top \boldsymbol{\theta}_* + \sum_{i=d+1}^D \sigma_i \mathbf{u}_{:,i} \mathbf{v}_{:,i}^\top \boldsymbol{\theta}_* =: \sum_{i=1}^d \sigma_i \mathbf{u}_{:,i} \mathbf{v}_{:,i}^\top \boldsymbol{\theta}_* + \boldsymbol{\eta},$$

where $\|\boldsymbol{\eta}\|_\infty \leq C \sum_{i=d+1}^D \sigma_i$. As a result, for any $d \leq r$, we can construct the feature mapping $\boldsymbol{\psi}_d(\mathbf{x}_i) = [\sigma_1 u_{i,1}, \dots, \sigma_d u_{i,d}]^\top \in \mathbb{R}^d$ such that $h(\mathbf{x}_i) = \langle \boldsymbol{\psi}_d(\mathbf{x}_i), \tilde{\boldsymbol{\theta}}_* \rangle + \boldsymbol{\eta}(\mathbf{x}_i)$, where $\tilde{\boldsymbol{\theta}}_* = [\mathbf{V}^\top \boldsymbol{\theta}_*]_{[1:d]} \in \mathbb{R}^d$ is the associated reward parameter.³ The upper bound of the induced misspecification can be expressed as $\tilde{\gamma}(d) = C \sum_{i=d+1}^D \sigma_i$, which allows us to calculate $\gamma(d)$. We can then apply Algorithm 1 to identify an ϵ -optimal arm. A high-dimensional linear bandit instance is provided in Appendix B.3 showing that: Algorithm 1 takes $\tilde{O}(1/\epsilon^2)$ samples to identify an ϵ -optimal arm, while the sample complexity upper bound of RAGE scales as $\tilde{O}(D/\epsilon^2)$.

4 Pure exploration in RKHS

We consider a kernel function $\mathcal{K} : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ over a compact set \mathcal{Z} ; we assume the kernel function satisfies condition stated in the Mercer's Theorem (see Appendix E.1) and has eigenvalues decay fast enough (see Assumption 1). Let \mathcal{H} be the Reproducing Kernel Hilbert Space (RKHS) induced from \mathcal{K} . We assume $\mathcal{X} \subseteq \mathcal{Z}$ and the true reward of any arm $\mathbf{x} \in \mathcal{X}$ is given by an unknown function $h \in \mathcal{H}$ such that $\|h\|_{\mathcal{H}} \leq 1$.

Let $\{\phi_j\}_{j=1}^\infty$ and $\{\mu_j\}_{j=1}^\infty$ be sequences of eigenfunctions and non-negative eigenvalues associated with kernel \mathcal{K} .⁴ A corollary of Mercer's theorem shows that any $h \in \mathcal{H}$ can be written in the form of $h(\cdot) = \sum_{j=1}^\infty \theta_j \phi_j(\cdot)$ for some $\{\theta_j\}_{j=1}^\infty \in \ell^2(\mathbb{N})$ such that $\sum_{j=1}^\infty \theta_j^2 / \mu_j < \infty$. We also have $\|h\|_{\mathcal{H}}^2 = \sum_{j=1}^\infty \theta_j^2 / \mu_j$. Although functions in RKHS are non-linear in nature, we now can represent them in terms of an infinite-dimensional linear function. We construct feature mappings for the embedding next.

For any $\mathbf{x} \in \mathcal{X}$, we have $h(\mathbf{x}) = \sum_{j=1}^\infty \theta_j \phi_j(\mathbf{x}) = \sum_{j=1}^\infty \frac{\theta_j}{\sqrt{\mu_j}} \sqrt{\mu_j} \phi_j(\mathbf{x})$. Let $C_\phi := \sup_{\mathbf{x} \in \tilde{\mathcal{X}}, j \geq 1} |\phi_j(\mathbf{x})|$. Since $\sum_{j=1}^\infty \theta_j^2 / \mu_j = \|h\|_{\mathcal{H}}^2 \leq 1$ is bounded, for any $d \in \mathbb{N}$, we define feature mapping $\boldsymbol{\psi}_d(\mathbf{x}) = [\sqrt{\mu_1} \phi_1(\mathbf{x}), \dots, \sqrt{\mu_d} \phi_d(\mathbf{x})]^\top \in \mathbb{R}^d$ such that

$$h(\mathbf{x}) = \langle \boldsymbol{\theta}_d, \boldsymbol{\psi}_d(\mathbf{x}) \rangle + \eta_d(\mathbf{x}),$$

where $\boldsymbol{\theta}_d = [\theta_1 / \sqrt{\mu_1}, \dots, \theta_d / \sqrt{\mu_d}]^\top \in \mathbb{R}^d$ and $|\eta_d(\mathbf{x})| \leq \tilde{\gamma}(d) := C_\phi \sqrt{\sum_{j>d} \mu_j}$. We remark here that the constant C_ϕ is calculable and usually mild, e.g., $C_\phi = 1$ for $\phi_j(x) = \sin((2j-1)\pi x/2)$.

We can then construct $\gamma(d)$ and $d_{\text{eff}}(\epsilon)$ as in Section 3 and specialize Algorithm 1 to the kernel setting. Both $\gamma(d)$ and $d_{\text{eff}}(\epsilon)$ depend on eigenvalues of the associated kernel. Fortunately, fast eigenvalue decay are satisfied by most kernel functions, e.g., Gaussian kernel. We quantify such properties through the following assumption.

Assumption 1. We consider kernels with the following eigenvalue decay with some absolute constants C_k and β .

1. Kernel \mathcal{K} is said to have (C_k, β) -polynomial eigenvalue decay (with $\beta > 3/2$) if $\mu_j \leq C_k j^{-\beta}$ for all $j \geq 1$.

³We note that the embeddings and associated quantities can also be constructed on the fly with respect to the set of uneliminated arms.

⁴With a known kernel, the sequence of eigenfunctions and eigenvalues can be analytically calculated or numerically approximated [36, 35]. We assume the knowledge of eigenfunctions and eigenvalues in this paper.

2. Kernel \mathcal{K} is said to have (C_k, β) -exponential eigenvalue decay (with $\beta > 0$) if $\mu_j \leq C_k e^{-\beta j}$ for all $j \geq 1$.

Theorem 2. Suppose Assumption 1 holds. For any $\epsilon > 0$, the following statements hold when we specialize Algorithm 1 to the kernel setting.

1. Suppose \mathcal{K} has (C_k, β) -polynomial eigenvalue decay. We have $d_{\text{eff}}(\epsilon) \leq O(\epsilon^{-2/(2\beta-3)})$, and the sample complexity of identifying an ϵ -optimal arm is upper bounded by $\tilde{O}(\epsilon^{-2/(2\beta-3)} \max\{\Delta_{\min}, \epsilon\}^{-2})$.
2. Suppose \mathcal{K} has (C_k, β) -exponential eigenvalue decay. We have $d_{\text{eff}}(\epsilon) \leq O(\log(1/\epsilon))$, and the sample complexity of identifying an ϵ -optimal arm is upper bounded by $\tilde{O}(\max\{\Delta_{\min}, \epsilon\}^{-2})$.

Remark 4. Our sample complexity guarantees are directly related to the eigenvalue decay of the underlying kernel function, rather than the empirical kernel matrix as studied in previous works [7, 42]. Although one can also provide an instance dependent bound as in Theorem 1, the worst-case sample complexity bound in Theorem 2 provides insightful characterizations of the sample complexity in terms of eigenvalue decay. One should notice that with exponential eigenvalue decay, the sample complexity $\tilde{O}(\epsilon^{-2})$ essentially matches, up to logarithmic factors, the complexity of distinguishing a two-armed bandit up to accuracy ϵ [25].

5 Pure exploration with neural networks

In this section we present a neural network-based pure exploration algorithm in Algorithm 2. Our algorithm is inspired by the recently proposed neural bandits algorithms for regret minimization [47, 46]. At the core of our algorithm is to use a neural network $f(\mathbf{x}; \boldsymbol{\theta})$ to learn the unknown reward function h . Specifically, following [8, 47], we consider a fully connected neural network $f(\mathbf{x}; \boldsymbol{\theta})$ with depth $L \geq 2$

$$f(\mathbf{x}; \boldsymbol{\theta}) = \sqrt{m} \mathbf{W}_L \sigma \left(\mathbf{W}_{L-1} \sigma \left(\cdots \sigma \left(\mathbf{W}_1 \mathbf{x} \right) \right) \right), \quad (2)$$

where $\sigma(x) := \max(x, 0)$ is the ReLU activation function, $\mathbf{W}_1 \in \mathbb{R}^{m \times d}$, $\mathbf{W}_L \in \mathbb{R}^{1 \times m}$, and for $2 \leq i \leq L-1$, $\mathbf{W}_i \in \mathbb{R}^{m \times m}$. Moreover, we denote $\boldsymbol{\theta} = [\text{vec}(\mathbf{W}_1)^\top, \dots, \text{vec}(\mathbf{W}_L)^\top]^\top \in \mathbb{R}^p$, where $p = m + md + m^2(L-2)$ is the number of all the network parameters. We use $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \boldsymbol{\theta})$ to denote the gradient of the neural network output with respect to the weights.

In detail, at k -th iteration, Algorithm 2 firstly applies its current gradient mapping $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})$ over the whole action set \mathcal{X} , and obtains the collection of gradients $\mathbf{G} \in \mathbb{R}^{|\mathcal{X}| \times p}$. Then Algorithm 2 does SVD over \mathbf{G} and constructs a d_k -dimensional feature mapping $\boldsymbol{\psi}_{d_k}$, which can be regarded as the projection of the gradient feature mapping $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})$ to the most informative d_k -dimensional eigenspace. Here we choose d_k such that the summation of the eigenvalues of the remaining eigenspace be upper bounded by some error $\bar{\epsilon}$. Algorithm 2 then computes the optimal design λ_k over $\boldsymbol{\psi}_{d_k}(\mathcal{X})$ and pulls arms $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_k}\}$ based on both the design λ_k and the total number of allocations N_k . Finally, Algorithm 2 trains a new neural network $f(\mathbf{x}; \boldsymbol{\theta}_k)$ using gradient descent starting from the initial parameter $\boldsymbol{\theta}_0$ (details are deferred to Appendix F), then eliminates the arms \mathbf{x} in the current arm set $\widehat{\mathcal{S}}_k$ which are sub-optimal with respect to the neural network function value $f(\mathbf{x}; \boldsymbol{\theta}_k)$.

The main difference between Algorithm 2 and its RKHS counterpart is as follows. Unlike Algorithm 1 which works on known feature mappings $\boldsymbol{\psi}_d$ (derived from a known kernel \mathcal{K}), Algorithm 2 does not have information about the feature mapping, and thus it constructs the feature mapping from the raw high-dimensional up-to-date gradient mapping $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})$. The feature mapping is constructed with respect to a *trained* neural work, which leverages the great representation power of neural networks. This makes Algorithm 2 a more general and flexible algorithm than Algorithm 1.

Now we present the main theorem of Algorithm 2. Let $\mathbf{H}^{|\mathcal{X}| \times |\mathcal{X}|}$ be the Neural Tangent Kernel (NTK)[17] gram matrix over all arms \mathcal{X} (the detailed definition of \mathbf{H} is deferred to Appendix F). We define the effective dimension for the neural version as below. The definition is similar to Definition 1.

Algorithm 2 Neural Arm Elimination

Input: Action set \mathcal{X} , initial parameter θ_0 , neural network $f(\mathbf{x}; \theta)$, gradient mapping $\mathbf{g}(\mathbf{x}, \theta)$, width of the matrix m , parameter of the number of allocations A , approximation parameter ζ , regularization parameter α , error parameter $\bar{\epsilon}$, ϵ , confidence level $\delta_k = \delta/(8k^2)$

- 1: Set $\widehat{\mathcal{S}}_1 = \mathcal{X}$.
- 2: **for** $k = 1, 2, \dots, n$ **do**
- 3: Construct the truncated feature representation $\psi_k(\mathcal{X})$ based on gradient mapping $\mathbf{g}(\mathbf{x}; \theta_{k-1})$. In detail, let $\mathbf{G} \in \mathbb{R}^{|\mathcal{X}| \times p}$ be the collection of gradients such that

$$\mathbf{G} = [\mathbf{g}(\mathbf{x}_1; \theta_{k-1})^\top; \dots; \mathbf{g}(\mathbf{x}_{|\mathcal{X}|}; \theta_{k-1})^\top] / \sqrt{m} \in \mathbb{R}^{|\mathcal{X}| \times p} \quad (3)$$

Let $[U, \Sigma, V]$ be the SVD of \mathbf{G} , where $U = (u_{i,j}) \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$, $\Sigma = [\text{diag}(e_1, \dots, e_{|\mathcal{X}|}), 0] \in \mathbb{R}^{|\mathcal{X}| \times p}$, $V \in \mathbb{R}^{p \times p}$. Get $d_k = \min\{d \in [|\mathcal{X}|] : \sum_{i=d+1}^{|\mathcal{X}|} e_i \leq \bar{\epsilon}\}$, and set $\psi_{d_k}(\mathbf{x}_i) = (e_1 u_{i,1}, \dots, e_{d_k} u_{i,d_k}) \in \mathbb{R}^{d_k}$.

- 4: Set λ_k and τ_k be the experimental design and the value of the following optimization problem

$$\inf_{\lambda \in \Lambda^{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi_{d_k}(\widehat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2. \quad (4)$$

- 5: Set $N_k = \max\{2^{2k} A(1 + \zeta) \log(|\mathcal{X}|^2 / \delta_k), r_{d_k}(\zeta)\}$.
- 6: Get $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\} = \text{ROUND}(\lambda_k, N_k, d_k, \zeta)$.
- 7: Pull arms $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ and receive rewards $\{y_1, \dots, y_{N_k}\}$.
- 8: Using J_k step η_k -step size gradient descent to optimize the following loss function to obtain θ_k ,

$$\theta_k = \arg \min L(\theta) := \sum_{j=1}^{N_k} (f(\mathbf{x}_j; \theta) - y_j)^2 + \frac{m\alpha}{2} \|\theta - \theta_0\|_2^2. \quad (5)$$

- 9: Set $\mathbf{A}_k = \alpha \mathbf{I} + \sum_{i=1}^{N_k} \psi_{d_k}(\mathbf{x}_i) \psi_{d_k}(\mathbf{x}_i)^\top$ and eliminate arms with respect to criteria $\widehat{\mathcal{S}}_{k+1} = \widehat{\mathcal{S}}_k \setminus \{\mathbf{x} \in \widehat{\mathcal{S}}_k : \exists \mathbf{x}' \text{ such that } f(\mathbf{x}'; \theta_k) - f(\mathbf{x}; \theta_k) \geq 2^{-k}/8 + 3\epsilon/8\}$.

10: **end for**

Output: Output any arm in $\widehat{\mathcal{S}}_{n+1}$.

Definition 2 (Neural version). For any $\epsilon > 0$, we define the effective dimension as $d_{\text{eff}}(\epsilon) := \min\{d \in [|\mathcal{X}|] : \sum_{i=d+1}^{|\mathcal{X}|} \lambda_i(\mathbf{H}) \leq \epsilon\}$, where $\lambda_i(\mathbf{H})$ is the i -th eigenvalue of \mathbf{H} .

Next, we make standard assumptions for the initialization of neural networks and the arms $\mathbf{x} \in \mathcal{X}$.

Assumption 2 ([47]). There exists $\lambda_0 > 0$ such that $\mathbf{H} \succeq \lambda_0 \mathbf{I}$. For any $\mathbf{x} \in \mathcal{X}$, the arm \mathbf{x} satisfies $\|\mathbf{x}\|_2 = 1$ and that its j -th coordinate is identical to its $j + d/2$ -th coordinate. Meanwhile, the initial parameter $\theta_0 = [\text{vec}(\mathbf{W}_1)^\top, \dots, \text{vec}(\mathbf{W}_L)^\top]^\top$ is initialized as follows: for $1 \leq l \leq L - 1$, \mathbf{W}_l is set to $\begin{pmatrix} \mathbf{W} & \mathbf{0} \\ \mathbf{0} & \mathbf{W} \end{pmatrix}$, where each entry of \mathbf{W} is generated independently from $\mathcal{N}(0, 4/m)$; \mathbf{W}_L is set to $(\mathbf{w}^\top, -\mathbf{w}^\top)$, where each entry of \mathbf{w} is generated independently from $\mathcal{N}(0, 2/m)$.

We now present our main theorem for pure exploration with neural network approximation. The formal version of our theorem is deferred to Appendix F.

Theorem 3 (Informal). Under Assumption 2, with proper selection of parameters $\alpha, n, \bar{\epsilon}, A, \eta_k, J_k$, then when $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, with probability at least $1 - \delta$, $\widehat{\mathcal{S}}_{K+1}$ only includes arm \mathbf{x} satisfying $\Delta_{\mathbf{x}} \leq \epsilon$, and the total sample complexity of Algorithm 2 is bounded by

$$N = \tilde{O}\left((1 + \zeta) d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)/\epsilon^2 + r_{d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)}(\zeta)\right) = \tilde{O}\left(d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)\epsilon^{-2}\right).$$

Remark 5. For the case where the effective dimension can be well bounded, e.g., $d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|) = O(\log(|\mathcal{X}|/\bar{\epsilon}^2))$, Theorem 7 suggests that Algorithm 2 is able to identify an ϵ -optimal arm within $\tilde{O}(\epsilon^{-2})$ samples. That suggests that our neural network-based algorithm is efficient without constructing a low-dimensional linear approximation of h in prior, like the previous linear or RKHS approaches.

6 Experiments

We conduct four experiments on synthetic and real-world datasets. We specialize our embedding idea to the neural, kernel, and linear regimes, and denote the algorithms as NeuralEmbedding (Algorithm 2), KernelEmbedding (Algorithm 1 with Gaussian kernel), and LinearEmbedding (Algorithm 1 with linear representation), respectively. We compare our algorithms with two baselines: RAGE and ActionElim. RAGE [13] conducts pure exploration in linear bandits and ActionElim [18, 11] ignores all feature representations. The (empirical) sample complexity of each algorithm is calculated as the number of samples needed so that the uneliminated set contains only ϵ -optimal arms. Unsuccessful runs, i.e., those terminate with non- ϵ -optimal arms, are reported as failures. In our experiments, we set $\epsilon = 0.1$ and $\delta = 0.05$. All results are averaged over 50 runs.⁵

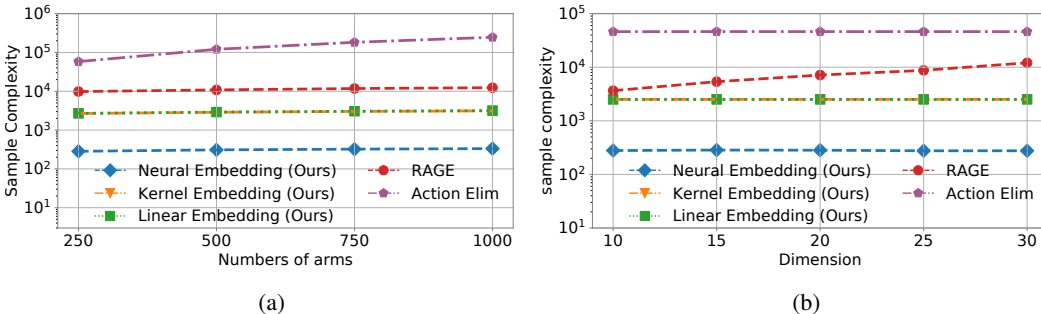


Figure 1: Experiments on synthetic datasets. (a) linear reward function with $D = 20$; (b) nonlinear reward function with $K = 200$.

Synthetic datasets. We first generate the feature matrix $\widetilde{\mathbf{X}} = \mathbf{X} + \mathbf{E} \in \mathbb{R}^{K \times D}$ where \mathbf{X} is constructed as a rank-2 matrix and \mathbf{E} is a perturbation matrix with tiny spectral norm (See Appendix G for details). Each row of $\widetilde{\mathbf{X}}$ represents the feature representation of an arm, and those features can be grouped into two clusters with equal size. In the case with linear rewards, we let θ_* equal to the first row of $\widetilde{\mathbf{X}}$. For nonlinear rewards, the reward of each arm is set as the 2-norm of its feature representation. We vary the number of arms K and the ambient dimension D in our experiments.

Fig. 1 shows experimental results on synthetic datasets. All algorithms successfully identify ϵ -optimal arms with zero empirical failure probability (due to the simplicity of the datasets). In terms of sample complexity, NeuralEmbedding outperforms all other algorithms in most cases, and KernelEmbedding and LinearEmbedding significantly outperform RAGE and Action Elim. The sample complexities of NeuralEmbedding, KernelEmbedding and LinearEmbedding are not affected when increasing number of arms or dimensions since they first identify the important subspace and then conduct elimination. On the other side, the sample complexity of ActionElim gets larger with increasing number of arms and the sample complexity of RAGE gets larger with increasing dimensions.

MNIST dataset. The MNIST dataset [29] contains hand-written digits from 0 to 9. We view each digit as an arm, and set their rewards according to the confusion matrix of a trained classifier. Digit 7 is chosen as the optimal arm with reward 1; the reward of digits 1, 2 and 9 are set to be 0.8, and all other digits have reward 0.5. In each experiment, we randomly draw 200 samples (20 samples each digit) from the dataset. We project the raw feature matrix $\mathbf{X} \in \mathbb{R}^{200 \times 784}$ into a lower-dimensional space $\widetilde{\mathbf{X}} \in \mathbb{R}^{200 \times 200}$ so that it becomes full rank (but without losing any information) and feasible for RAGE. Our goal is to correctly identify a digit 7.

Yahoo dataset. The Yahoo! User Click Log Dataset R6A⁶ contains users' click-through records. Each record consists of a 36-dimensional feature representation (obtained from an outer product), and a binary outcome stating whether or not a user clicked on the article. We view each record as

⁵All algorithms are elimination-styled for fair comparison. Other implementation details are deferred to Appendix G.

⁶<https://webscope.sandbox.yahoo.com>

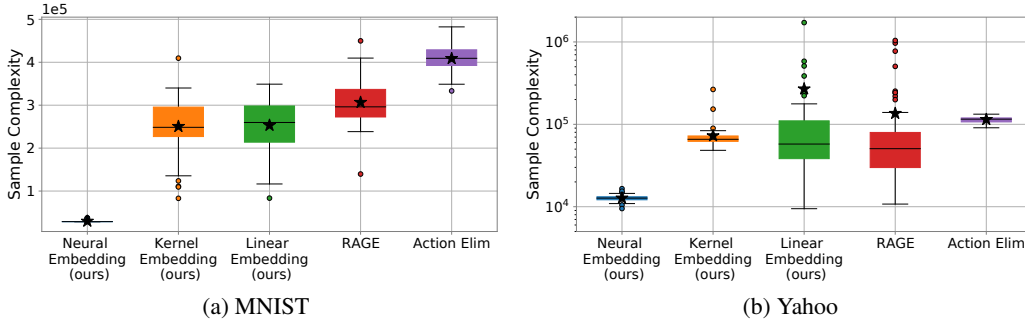


Figure 2: Experiments on real-world datasets. The mean sample complexity is represented by a black star. The mean sample complexities of LinearEmbedding and RAGE are heavily affected by outliers in the Yahoo dataset.

an arm, and set the reward as 0.8 (if clicked) or 0.3 (if not clicked) to makes the problem harder. In each experiment, we randomly draw 200 arms from the dataset, where 5 of them having rewards 0.8 (proportional to true click-through ratio), Our goal is to identify an arm with rewards 0.8. Our experimental setup is similar to the one used in Fiez et al. [13]. However, their true rewards are obtained from a least square regression. We do not enforce linearity in rewards in our experiment.

Box plots in Fig. 2 show the sample complexity of each algorithm on real-world datasets. NeuralEmbedding significantly outperforms all other algorithms thanks to (1) the representation power of neural networks and (2) efficient exploration in low-dimensional spaces. KernelEmbedding and LinearEmbedding have competitive performance on the MNIST dataset. Table 1 shows the success rate of each algorithm. Linear methods such as RAGE and LinearEmbedding have relatively low success rates on the Yahoo dataset (with nonlinear rewards). Our NeuralEmbedding and KernelEmbedding methods have high success rates since they are designed for nonlinear setting.

Table 1: Success rates on real-world datasets

	NeuralEmbedding	KernelEmbedding	LinearEmbedding	RAGE	ActionElim
MNIST	98%	100%	100%	100%	100%
Yahoo	100%	98%	88%	90%	100%

7 Conclusion

We introduce the idea of adaptive embedding in bandit pure exploration. Unlike existing works that passively deal with model misspecifications, we adaptively embed high-dimensional feature representations into lower-dimensional spaces to avoid the curse of dimensionality. The induced misspecifications are carefully dealt with. We further apply our approach to two under-studied settings with the nonlinearity: (1) pure exploration in an RKHS and (2) pure exploration with neural networks. Our sample complexity guarantees depend on the effective dimension of the feature spaces in the kernel or neural representations. We conduct extensive experiments on both synthetic and real-world datasets, and our algorithms greatly outperform existing ones.

Our current analysis with neural representations is in the NTK regime, which can only describe a part of the representation of the neural networks. We leave extending our algorithm to more general settings (beyond the NTK regime) as a future direction.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers and area chair for their helpful comments. YZ and RN are partially supported by AFOSR grant FA9550-18-1-0166. DZ and QG are partially supported by the National Science Foundation CAREER Award 1906169, IIS-1904183 and AWS Machine Learning Research Award. RJ and RW are partially supported by AFOSR FA9550-18-1-0166, NSF OAC-1934637, NSF DMS-2023109, and NSF DGE-2022023. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, Aarti Singh, and Yining Wang. Near-optimal discrete optimization for experimental design: A regret minimization approach. *Mathematical Programming*, pages 1–40, 2020.
- [2] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, 2019.
- [3] Maryam Aziz, Emilie Kaufmann, and Marie-Karelle Riviere. On multi-armed bandit designs for phase i clinical trials. *arXiv preprint arXiv:1903.07082*, 2019.
- [4] Peter L Bartlett, Victor Gabillon, and Michal Valko. A simple parameter-free and adaptive approach to optimization under a minimal local smoothness assumption. In *Algorithmic Learning Theory*, pages 184–206. PMLR, 2019.
- [5] Robert E Bechhofer. A sequential multiple-decision procedure for selecting the best one of several normal populations with a common unknown variance, and its use with various experimental designs. *Biometrics*, 14(3):408–429, 1958.
- [6] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, pages 23–37. Springer, 2009.
- [7] Romain Camilleri, Julian Katz-Samuels, and Kevin Jamieson. High-dimensional experimental design and kernel bandits. *arXiv preprint arXiv:2105.05806*, 2021.
- [8] Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, 2019.
- [9] Lijie Chen and Jian Li. On the optimal sample complexity for best arm identification. *arXiv preprint arXiv:1511.03774*, 2015.
- [10] Rémy Degenne, Pierre Ménard, Xuedong Shang, and Michal Valko. Gamification of pure exploration for linear bandits. In *International Conference on Machine Learning*, pages 2432–2442. PMLR, 2020.
- [11] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270. Springer, 2002.
- [12] Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- [13] Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In *Advances in Neural Information Processing Systems*, pages 10667–10677, 2019.
- [14] Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in contextual bandits. *Advances in Neural Information Processing Systems*, 33, 2020.
- [15] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pages 998–1027. PMLR, 2016.
- [16] Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. Misspecified linear bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017.
- [17] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.

- [18] Kevin Jamieson and Robert Nowak. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2014.
- [19] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. $\text{lil}'\text{ucb}$: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pages 423–439. PMLR, 2014.
- [20] Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pages 655–662, 2012.
- [21] Hideaki Kano, Junya Honda, Kentaro Sakamaki, Kentaro Matsuura, Atsuyoshi Nakamura, and Masashi Sugiyama. Good arm identification via bandit feedback. *Machine Learning*, 108(5): 721–745, 2019.
- [22] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pages 1238–1246. PMLR, 2013.
- [23] Julian Katz-Samuels and Kevin Jamieson. The true sample complexity of identifying good arms. In *International Conference on Artificial Intelligence and Statistics*, pages 1781–1791. PMLR, 2020.
- [24] Julian Katz-Samuels, Lalit Jain, Zohar Karnin, and Kevin Jamieson. An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. *arXiv preprint arXiv:2006.11685*, 2020.
- [25] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1): 1–42, 2016.
- [26] Abbas Kazerouni and Lawrence M Wein. Best arm identification in generalized linear bandits. *Operations Research Letters*, 49(3):365–371, 2021.
- [27] Jack Kiefer and Jacob Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960.
- [28] Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in rl with a generative model. In *International Conference on Machine Learning*, pages 5662–5670. PMLR, 2020.
- [29] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [30] Blake Mason, Lalit Jain, Ardhendu Tripathy, and Robert Nowak. Finding all ϵ -good arms in stochastic bandits. *Advances in Neural Information Processing Systems*, 33, 2020.
- [31] Rémi Munos. From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning, 2014. <https://hal.archives-ouvertes.fr/hal-00747575v4/document>.
- [32] Edward Paulson et al. A sequential procedure for selecting the population with the largest mean from k normal populations. *Annals of Mathematical Statistics*, 35(1):174–180, 1964.
- [33] Friedrich Pukelsheim. *Optimal design of experiments*. SIAM, 2006.
- [34] Sivan Sabato. Epsilon-best-arm identification in pay-per-reward multi-armed bandits, 2019. <https://openreview.net/forum?id=H1xkvNr1LS>.
- [35] Gabriele Santin and Robert Schaback. Approximation of eigenfunctions in kernel-based spaces. *Advances in Computational Mathematics*, 42(4):973–993, 2016.
- [36] Stewart Schlesinger. Approximating eigenvalues and eigenfunctions of symmetric kernels. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):1–14, 1957.

- [37] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [38] Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. *arXiv preprint arXiv:1409.6110*, 2014.
- [39] Gilbert W Stewart. Perturbation theory for the singular value decomposition. Technical report, University of Maryland, 1998. <https://drum.lib.umd.edu/handle/1903/552>.
- [40] Ervin Tanczos, Robert Nowak, and Bob Mankoff. A kl-lucb bandit algorithm for large-scale crowdsourcing. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5896–5905, 2017.
- [41] Chao Tao, Saúl Blanco, and Yuan Zhou. Best arm identification in linear bandits with linear dimension dependency. In *International Conference on Machine Learning*, pages 4877–4886. PMLR, 2018.
- [42] Michal Valko, Nathaniel Korda, Rémi Munos, Ilias Flaounas, and Nelo Cristianini. Finite-time analysis of kernelised contextual bandits. *arXiv preprint arXiv:1309.6869*, 2013.
- [43] Michal Valko, Rémi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, pages 46–54. PMLR, 2014.
- [44] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [45] Liyuan Xu, Junya Honda, and Masashi Sugiyama. A fully adaptive algorithm for pure exploration in linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 843–851. PMLR, 2018.
- [46] Weitong ZHANG, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural thompson sampling. In *International Conference on Learning Representations*, 2020.
- [47] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 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]
 - (b) Did you describe the limitations of your work? [Yes] As discussed in Section 7.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] This paper provides algorithms to conduct pure exploration for kernel and neural bandits. It does not lead to negative societal impacts in the foreseeable future.
 - (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] As stated in the paper.
 - (b) Did you include complete proofs of all theoretical results? [Yes] All proofs are provided in the appendix in supplementary material (full paper).
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)? [Yes] Code, data, and instructions can be found in the supplementary material (code).
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] As stated in Section 6 and Appendix G.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] As reported in Section 6.
 - (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)? [Yes] As stated in the supplementary material (code).
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? [Yes]
 - (b) Did you mention the license of the assets? [Yes] We clarified the source of datasets used in Section 6.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code used to produce our experiments is included in the supplementary material (code).
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] As stated in Appendix G.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] As discussed in Appendix G.
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 Supporting materials

A.1 Matrix inversion and rounding

For possibly singular $\mathbf{A}_{\psi_d(\lambda)}$, pseudo-inverse is used in $\|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2$ if \mathbf{y} belongs to the range of $\mathbf{A}_{\psi_d(\lambda)}$; otherwise, we set $\|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2 = \infty$.

Consider any $\mathcal{S} \subseteq \mathcal{X}$. With our (slightly abused) definition of matrix inversion, the optimal design

$$\inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{S}))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2 \quad (6)$$

will select a design $\lambda^* \in \Lambda_{\mathcal{X}}$ such that every $\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{S}))$ lies in the range of $\mathbf{A}_{\psi_d(\lambda^*)}$.⁷ If $\text{span}(\mathcal{Y}(\psi_d(\mathcal{S}))) = \mathbb{R}^d$, then $\mathbf{A}_{\psi_d(\lambda^*)}$ is positive definite (recall that $\mathbf{A}_{\psi_d(\lambda^*)} = \sum_{\mathbf{x} \in \mathcal{X}} \lambda_{\mathbf{x}} \psi_d(\mathbf{x}) \psi_d(\mathbf{x})^\top$ and $\text{span}(\psi_d(\mathcal{X})) = \mathbb{R}^d$ by assumption). Thus the rounding guarantees (Theorem 2.1 therein, which requires a positive definite design) in Allen-Zhu et al. [1] goes through (with additional simple modifications dealt as in Appendix B of Fiez et al. [13]).

We now consider the case when $\mathbf{A}_{\psi_d(\lambda^*)}$ is singular. Since $\text{span}(\psi_d(\mathcal{X})) = \mathbb{R}^d$, we can always find another λ' such that $\mathbf{A}_{\psi_d(\lambda')}$ is invertible. For any $\zeta_1 > 0$, let $\tilde{\lambda}^* = (1 - \zeta_1)\lambda^* + \zeta_1\lambda'$. We know that $\tilde{\lambda}^*$ leads to a positive definite design. With respect to ζ_1 , we can find another $\zeta_2 > 0$ small enough (e.g., smaller than the smallest eigenvalue of $\zeta_1 \mathbf{A}_{\psi_d(\lambda')}$) such that $\mathbf{A}_{\psi_d(\tilde{\lambda}^*)} \succeq \mathbf{A}_{\psi_d((1 - \zeta_1)\lambda^* + \zeta_2 \mathbf{I})} + \zeta_2 \mathbf{I}$. Since $\mathbf{A}_{\psi_d((1 - \zeta_1)\lambda^* + \zeta_2 \mathbf{I})}$ is positive definite, for any $\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{S}))$, we have

$$\|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\tilde{\lambda}^*)}^{-1}}^2 \leq \|\mathbf{y}\|_{(\mathbf{A}_{\psi_d((1 - \zeta_1)\lambda^* + \zeta_2 \mathbf{I})} + \zeta_2 \mathbf{I})^{-1}}^2.$$

Fix any $\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{S}))$. Since \mathbf{y} lies in the range of $\mathbf{A}_{\psi_d(\lambda^*)}$ (according to the objective and the definition of matrix inversion), we clearly have

$$\|\mathbf{y}\|_{(\mathbf{A}_{\psi_d((1 - \zeta_1)\lambda^* + \zeta_2 \mathbf{I})} + \zeta_2 \mathbf{I})^{-1}}^2 \leq \|\mathbf{y}\|_{\mathbf{A}_{\psi_d((1 - \zeta_1)\lambda^*)}^{-1}}^2 \leq \frac{1}{1 - \zeta_1} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda^*)}^{-1}}^2.$$

To summarize, we have

$$\|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\tilde{\lambda}^*)}^{-1}}^2 \leq \frac{1}{1 - \zeta_1} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda^*)}^{-1}}^2,$$

where ζ_1 can be chosen arbitrarily small. We can thus send the positive definite design $\tilde{\lambda}^*$ to the rounding procedure in Allen-Zhu et al. [1]. We can incorporate the additional $1/(1 - \zeta_1)$ overhead, for $\zeta_1 > 0$ chosen sufficiently small, into the sample complexity requirement $r_d(\zeta)$ of the rounding procedure.

A.2 Supporting lemmas

Lemma 1 (Theorem 10 in Katz-Samuels et al. [24]). *Consider a linear bandit problem with action set $\mathcal{X} \subseteq \mathbb{R}^D$. Suppose $\text{span}(\{\mathbf{x}_* - \mathbf{x}\}_{\mathbf{x} \in \mathcal{X}}) = \mathbb{R}^D$, we then have*

$$\inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \mathbf{x}_*} \|\mathbf{x}_* - \mathbf{x}\|_{\mathbf{A}(\lambda)^{-1}}^2 = \Omega(D).$$

B Omitted proofs for Section 3

B.1 Proof of Theorem 1

Theorem 1. *With probability of at least $1 - \delta$, Algorithm 1 correctly outputs an ϵ -optimal arm with sample complexity upper bounded by*

$$640 \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \left(\left(k \rho_{d_k}^* (2^{2-k}) \log(k^2 |\mathcal{X}|^2 / \delta) \right) + (r_{d_k}(\zeta) + 1) \right) \leq \tilde{O}(d_{\text{eff}}(\epsilon) \cdot \max\{\Delta_{\min}, \epsilon\}^{-2}),$$

where $d_k = d_{\text{eff}}(4 \cdot 2^{-k}) \leq d_{\text{eff}}(\epsilon)$ since $4 \cdot 2^{-k} \geq \epsilon$ when $k \leq \lceil \log_2(2/\epsilon) \rceil$.

⁷If the infimum is not attained, we can simply take a design λ^{**} with associated value $\tau^{**} \leq (1 + \zeta_0) \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi_d(\mathcal{S}))} \|\mathbf{y}\|_{\mathbf{A}_{\psi_d(\lambda)}^{-1}}^2$ for a $\zeta_0 > 0$ arbitrarily small. Our analysis goes through with changes only in constant terms.

Proof. The proof is decomposed into three steps: (1) prove correctness through induction; (2) bound the error probability; (3) upper bound the sample complexity. We first note that $d_{\text{eff}}(\epsilon)$ is well-defined since we assume that $\gamma(d)$ can be made arbitrarily small for d large enough.

Step 1: The induction. We define event

$$\mathcal{E}_k = \{\mathbf{x}_* \in \widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k\},$$

and prove through induction that

$$\mathbb{P}(\mathcal{E}_{k+1} | \cap_{i \leq k} \mathcal{E}_i) \geq 1 - \delta_k,$$

where $\delta_0 := 0$. Recall that $\mathbf{x}_* = \arg \max_{\mathbf{x} \in \mathcal{X}} h(\mathbf{x})$ and $\mathcal{S}_k = \{x \in \mathcal{X} : \Delta_x < 4 \cdot 2^{-k}\}$ (with $\mathcal{S}_1 = \mathcal{X}$).

The base case $\{\mathbf{x}_* \in \widehat{\mathcal{S}}_1 \subseteq \mathcal{S}_1\}$ holds with probability 1 by definition of $\widehat{\mathcal{S}}_1 = \mathcal{S}_1 = \mathcal{X}$. We next analyze event \mathcal{E}_{k+1} conditioned on $\cap_{i \leq k} \mathcal{E}_i$. For simplicity, we'll use notations $\boldsymbol{\psi}(\cdot) := \boldsymbol{\psi}_{d_k}(\cdot)$, $\boldsymbol{\theta}_* := \boldsymbol{\theta}_{d_k}$ and $\eta(\cdot) := \eta_{d_k}(\cdot)$ in the analysis of the k -th iteration.

Step 1.1: Concentration under mis-specifications. Let $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ be the arms pulled at the k -th iteration of Algorithm 1 and $\{y_1, \dots, y_{N_k}\}$ be the corresponding rewards. The least square estimator $\widehat{\boldsymbol{\theta}}_k = \mathbf{A}_k^{-1} \mathbf{b}_k \in \mathbb{R}^{d_k}$ is constructed with $\mathbf{A}_k = \sum_{i=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_i) \boldsymbol{\psi}(\mathbf{x}_i)^\top$ and $\mathbf{b}_k = \sum_{i=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_i) y_i$. Note that $y_i = h(\mathbf{x}_i) + \xi_i = \boldsymbol{\psi}(\mathbf{x}_i)^\top \boldsymbol{\theta}_* + \eta(\mathbf{x}_i) + \xi_i$ are i.i.d. generated with 1-sub-Gaussian noise ξ_i . Fix any $\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k))$, we have

$$\begin{aligned} \left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_* \rangle \right| &= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_j) y_j - \mathbf{y}^\top \boldsymbol{\theta}_* \right| \\ &= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_j) (\boldsymbol{\psi}(\mathbf{x}_j)^\top \boldsymbol{\theta}_* + \eta(\mathbf{x}_j) + \xi_j) - \mathbf{y}^\top \boldsymbol{\theta}_* \right| \\ &= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_j) (\eta(\mathbf{x}_j) + \xi_j) \right| \\ &\leq \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_j) \eta(\mathbf{x}_j) \right| + \left| \sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_j) \xi_j \right|. \end{aligned} \quad (7)$$

We now bound the two terms in Eq. (7) separately as follows. Since $\eta(\mathbf{x}_i) \leq \widetilde{\gamma}(d)$ by construction, for the first term, we have

$$\begin{aligned} \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_j) \eta(\mathbf{x}_j) \right| &\leq \widetilde{\gamma}(d) \sum_{i=1}^{N_k} |\mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i)| \\ &= \widetilde{\gamma}(d_k) \sum_{i=1}^{N_k} \sqrt{(\mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i))^2} \\ &\leq \widetilde{\gamma}(d_k) \sqrt{N_k \sum_{i=1}^{N_k} (\mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i))^2} \end{aligned} \quad (8)$$

$$\begin{aligned} &= \widetilde{\gamma}(d_k) \sqrt{N_k \sum_{i=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i) \boldsymbol{\psi}(\mathbf{x}_i)^\top \mathbf{A}_k^{-1} \mathbf{y}} \\ &= \widetilde{\gamma}(d_k) \sqrt{N_k \|\mathbf{y}\|_{\mathbf{A}_k^{-1}}^2} \\ &\leq \widetilde{\gamma}(d_k) \sqrt{(1 + \zeta) \tau(\mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k)))} \end{aligned} \quad (9)$$

where Eq. (8) comes from Jensen's inequality; Eq. (9) comes from the guarantee of rounding in Eq. (1).

The second term in Eq. (7) is a (weighted) sum of independent random variables ξ_i . Since ξ_i s are 1-sub-Gaussian, then $(\sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i) \xi_i)$ is $\sqrt{\sum_{j=1}^{N_k} (\mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i))^2}$ -sub-Gaussian. Since $\sqrt{\sum_{j=1}^{N_k} (\mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i))^2} = \|\mathbf{y}\|_{\mathbf{A}_k^{-1}}$, a standard Hoeffding's inequality leads to the following bound

$$\mathbb{P} \left(\left| \sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i) \xi_i \right| \geq \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)} \right) \leq \frac{\delta_k}{|\widehat{\mathcal{S}}_k|^2 / 2}$$

Since there are at most $|\widehat{\mathcal{S}}_k|^2 / 2$ directions in $\mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k))$, taking a union bound over all possible directions, we have

$$\mathbb{P} \left(\forall \mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k)) : \left| \sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i) \xi_i \right| \geq \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)} \right) \leq \delta_k. \quad (10)$$

Plugging Eq. (9) and Eq. (10) into Eq. (7), we have that for any $\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k))$,

$$\left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}^* \rangle \right| \leq \tilde{\gamma}(d_k) \sqrt{(1 + \zeta) \tau(\mathcal{Y}(\boldsymbol{\psi}(\widehat{\mathcal{S}}_k)))} + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)} =: \tilde{\omega}_k(\mathbf{y}) \quad (11)$$

with probability at least $1 - \delta_k$.

Step 1.2: Correctness. We prove the result under the event described in Eq. (11), which holds with (conditional) probability at least $1 - \delta_k$. We know that $\mathbf{x}_* \in \widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k$ holds conditioned on $\cap_{i \leq k} \mathcal{E}_i$. We now prove it for iteration $k + 1$. Recall the elimination criteria is

$$\widehat{\mathcal{S}}_{k+1} = \widehat{\mathcal{S}}_k \setminus \{\mathbf{x} \in \widehat{\mathcal{S}}_k : \exists \mathbf{x}' \text{ such that } (\boldsymbol{\psi}(\mathbf{x}') - \boldsymbol{\psi}(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k \geq \omega_k(\boldsymbol{\psi}(\mathbf{x}') - \boldsymbol{\psi}(\mathbf{x}))\}.$$

Step 1.2.1: $\mathbf{x}_* \in \widehat{\mathcal{S}}_{k+1}$. Let $\widehat{\mathbf{x}} \in \widehat{\mathcal{S}}_k$ be any arm such that $\widehat{\mathbf{x}} \neq \mathbf{x}_*$. We have

$$\begin{aligned} (\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*))^\top \widehat{\boldsymbol{\theta}}_k &\leq (\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*))^\top \boldsymbol{\theta}^* + \tilde{\omega}_k(\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*)) \\ &= h(\widehat{\mathbf{x}}) - \eta(\widehat{\mathbf{x}}) - (h(\mathbf{x}_*) - \eta(\mathbf{x}_*)) + \tilde{\omega}_k(\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*)) \\ &\leq -\Delta(\widehat{\mathbf{x}}) + (2\tilde{\gamma}(d_k) + \tilde{\omega}_k(\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*))) \\ &< (2\tilde{\gamma}(d_k) + \tilde{\omega}_k(\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*))) \\ &= \omega_k(\boldsymbol{\psi}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}(\mathbf{x}_*)). \end{aligned}$$

As a result, the optimal arm \mathbf{x}_* remains in $\widehat{\mathcal{S}}_{k+1}$.

Step 1.2.2: $\widehat{\mathcal{S}}_{k+1} \subseteq \mathcal{S}_{k+1}$. For any $\mathbf{x} \in \mathcal{S}_{k+1}^c \cap \widehat{\mathcal{S}}_k$, we have $\Delta(\mathbf{x}) = h(\mathbf{x}_*) - h(\mathbf{x}) \geq 4 \cdot 2^{-(k+1)} = 2 \cdot 2^{-k}$ by definition. Since $\mathbf{x}_* \in \widehat{\mathcal{S}}_k$, we have

$$\begin{aligned} (\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k &\geq (\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x}))^\top \boldsymbol{\theta}^* - \tilde{\omega}_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x})) \\ &= h(\mathbf{x}_*) - \eta(\mathbf{x}_*) - (h(\mathbf{x}) - \eta(\mathbf{x})) - \tilde{\omega}_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x})) \\ &\geq \Delta(\mathbf{x}) - (2\tilde{\gamma}(d_k) + \tilde{\omega}_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x}))) \\ &= \Delta(\mathbf{x}) - \omega_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x})) \\ &\geq \omega_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x})), \end{aligned} \quad (12)$$

where Eq. (12) comes from the fact that $\omega_k(\boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x})) \leq 2^{-k}$ by the selection of N_k (as discussed later) and the guarantees from the rounding procedure Eq. (1). As a result, any $\mathbf{x} \in \mathcal{S}_{k+1}^c \cap \widehat{\mathcal{S}}_k$ will be eliminated and $\widehat{\mathcal{S}}_{k+1} \subseteq \mathcal{S}_{k+1}$.

To summarize, we prove the induction at iteration $k + 1$, i.e.,

$$\mathbb{P}(\mathcal{E}_{k+1} | \cap_{i < k+1} \mathcal{E}_i) \geq 1 - \delta_k.$$

To show the selection of N_k guarantees Eq. (12), it suffices to show that $2^{-k} - \epsilon_k \geq 2^{-k}/2$. We first notice the following inequalities:

$$\begin{aligned}
(1 + \zeta) \tau(\mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))) &= (1 + \zeta) \inf_{\lambda \in \Lambda^{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2 \\
&\leq (1 + \zeta) \inf_{\lambda \in \Lambda^{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi(\mathcal{X}))} \|\mathbf{y}\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2 \\
&= g(d, \zeta).
\end{aligned} \tag{13}$$

Recall that $\gamma(d) = (16 + 8\sqrt{g(d, \zeta)})\tilde{\gamma}(d)$. At each the k -th iteration with $k \leq n := \lceil \log_2(2/\epsilon) \rceil$, we have

$$\epsilon_k \leq \gamma(d_k)/8 \tag{14}$$

$$\leq 2^{-k}/2, \tag{15}$$

where Eq. (14) comes from Eq. (13), and Eq. (15) comes from the definition of $d_k = d_{\text{eff}}(4 \cdot 2^{-k})$. This implies that $2^{-k} - \epsilon_k \geq 2^{-k}/2$ and thus N_k is well-defined.

Step 2: The error probability. From the analysis in Step 1, we have

$$\mathbb{P}(\mathcal{E}_{k+1} | \mathcal{E}_k \cap \dots \cap \mathcal{E}_1) \geq 1 - \delta_k.$$

Let $\mathcal{E} = \bigcap_{k=1}^{n+1} \mathcal{E}_k$, we then have

$$\begin{aligned}
\mathbb{P}(\mathcal{E}) &= \prod_{k=1}^{n+1} \mathbb{P}(\mathcal{E}_k | \mathcal{E}_{k-1} \cap \dots \cap \mathcal{E}_1) \\
&= \prod_{k=1}^{n+1} (1 - \delta_{k-1}) \\
&\geq \prod_{k=1}^{\infty} (1 - \delta/k^2) \\
&= \frac{\sin(\pi\delta)}{\pi\delta} \\
&\geq 1 - \delta,
\end{aligned} \tag{16}$$

where we use the fact that $\sin(\pi\delta)/\pi\delta \geq 1 - \delta$ for any $\delta \in (0, 1)$ in Eq. (16). We analyze the following steps under the good event \mathcal{E} .

Step 3: Sample complexity upper bound. As one can see, the total sample complexity N is upper bounded by

$$\begin{aligned}
N &\leq \sum_{k=1}^n N_k \\
&= \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \max \left\{ \left[(2^{-k} - \epsilon_k)^{-2} 2(1 + \zeta) \tau_k \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k) \right], r_{d_k}(\zeta) \right\} \\
&\leq \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \left(\left((2^{-k} - \epsilon_k)^{-2} 2(1 + \zeta) \tau_k \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k) \right) + (r_{d_k}(\zeta) + 1) \right) \\
&\leq \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \left(\left((2^{2k}) 10 \tau_k \log(k^2 |\mathcal{X}|^2 / \delta) \right) + (r_{d_k}(\zeta) + 1) \right).
\end{aligned}$$

We analyze $\rho_{d_k}^*(2^{2-k})$ in the following to obtain an instance dependent sample complexity upper bound.

$$\begin{aligned}
\rho_{d_k}^*(2^{2-k}) &= \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{\max\{\Delta(\mathbf{x}), 4 \cdot 2^{-k}\}^2} \\
&= \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{i \leq k} \sup_{\mathbf{x} \in \mathcal{S}_i \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{\max\{\Delta(\mathbf{x}), 4 \cdot 2^{-k}\}^2} \\
&\geq \sup_{i \leq k} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_i \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{\max\{\Delta(\mathbf{x}), 4 \cdot 2^{-k}\}^2} \\
&\geq \sup_{i \leq k} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_i \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{\max\{4 \cdot 2^{-i}, 4 \cdot 2^{-k}\}^2} \\
&\geq \frac{1}{k} \sum_{i=1}^k \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_i \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{16 \cdot 2^{-2i}} \\
&\geq \frac{1}{k} \sum_{i=1}^k \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_i} \frac{\|\psi(\mathbf{x}) - \psi(\mathbf{x}')\|_{\mathbf{A}_{\psi}(\lambda)}^2}{16 \cdot 2^{-2i}} \tag{17} \\
&= \frac{1}{k} \sum_{i=1}^k (2^{2i}) \tau(\mathcal{Y}(\psi(\mathcal{S}_i)))/64 \\
&\geq \frac{1}{64k} (2^{2k}) \tau(\mathcal{Y}(\psi(\mathcal{S}_k))),
\end{aligned}$$

where Eq. (17) comes from the fact that we can write $\psi(\mathbf{x}) - \psi(\mathbf{x}') = \psi(\mathbf{x}) - \psi(\mathbf{x}_*) + \psi(\mathbf{x}_*) - \psi(\mathbf{x}')$ and a triangle inequality.

We have $\tau_k = \tau(\mathcal{Y}(\psi_{d_k}(\widehat{\mathcal{S}}_k))) \leq \tau(\mathcal{Y}(\psi_{d_k}(\mathcal{S}_k)))$ since $\widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k$. Thus, we have

$$(2^{2k}) \tau_k \leq 64k \rho_{d_k}^*(2^{2-k}).$$

As a result, we have

$$\begin{aligned}
N &\leq 640 \sum_{k=1}^{\lceil \log_2(2/\epsilon) \rceil} \left((k \rho_{d_k}^*(2^{2-k}) \log(k^2 |\mathcal{X}|^2 / \delta)) + (r_{d_k}(\zeta) + 1) \right) \\
&\leq \widetilde{O}(d_{\text{eff}}(\epsilon) \cdot \max\{\Delta_{\min}, \epsilon\}^{-2}), \tag{18}
\end{aligned}$$

where Eq. (18) is discussed as follows.

Since we clearly have $2^{2-k} \geq \epsilon$ when $k \leq n := \lceil \log_2(2/\epsilon) \rceil$, we have

$$\begin{aligned}
\rho_{d_k}^*(2^{2-k}) &= \rho_{d_k}^*(\epsilon) \\
&= \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2}{\max\{h(\mathbf{x}_*) - h(\mathbf{x}), \epsilon\}^2} \\
&\leq \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \|\psi_{d_k}(\mathbf{x}_*) - \psi_{d_k}(\mathbf{x})\|_{\mathbf{A}_{\psi_{d_k}}(\lambda)}^2 \cdot \max\{\Delta_{\min}, \epsilon\}^{-2} \\
&\leq 4d_k \cdot \max\{\Delta_{\min}, \epsilon\}^{-2} \tag{19} \\
&= 4d_{\text{eff}}(4 \cdot 2^{-k}) \cdot \max\{\Delta_{\min}, \epsilon\}^{-2} \\
&\leq 4d_{\text{eff}}(\epsilon) \cdot \max\{\Delta_{\min}, \epsilon\}^{-2}, \tag{20}
\end{aligned}$$

where Eq. (19) comes from Kiefer–Wolfowitz Theorem [27] and Eq. (20) comes from the fact that $2^{2-k} \geq \epsilon$ and the definition of $d_{\text{eff}}(\epsilon)$. Similarly, we have $r_{d_k}(\zeta) \leq O(d_k/\zeta^2) \leq O(d_{\text{eff}}(\epsilon))$. \square

B.2 Proof of Proposition 1

Proposition 1. Suppose $\max_{\mathbf{x} \in \mathcal{X}} |h(\mathbf{x}) - \langle \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle| \leq \tilde{\gamma}(d)$. For any $\epsilon \geq \tilde{\gamma}(d)$, we have $\rho_d^*(\epsilon) \leq 9\tilde{\rho}_d^*(\epsilon)$. When $\tilde{\gamma}(d) < \Delta_{\min}/2$, $\tilde{\rho}_d^*(0)$ represents the sample complexity of a closely-related linear bandit problem without model misspecifications, i.e., $\tilde{\mathbf{h}}(\mathbf{x}) = \langle \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle$.

Proof. Recall that

$$\rho_d^*(\epsilon) = \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x})\|_{\mathbf{A}_{\boldsymbol{\psi}_d(\lambda)}^{-1}}^2}{\max\{\Delta_{\mathbf{x}}, \epsilon\}^2},$$

$$\tilde{\rho}_d^*(\epsilon) = \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x})\|_{\mathbf{A}_{\boldsymbol{\psi}_d(\lambda)}^{-1}}^2}{\max\{\langle \boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle, \epsilon\}^2},$$

and

$$\Delta_{\mathbf{x}} = h(\mathbf{x}_*) - h(\mathbf{x}) = \langle \boldsymbol{\psi}(\mathbf{x}_*) - \boldsymbol{\psi}(\mathbf{x}), \boldsymbol{\theta}_d \rangle + \eta(\mathbf{x}_*) - \eta(\mathbf{x}), \quad (21)$$

where $|\eta(\mathbf{x})| \leq \gamma$ for any $\mathbf{x} \in \mathcal{X}$.

To relate $\rho^*(\epsilon)$ with $\tilde{\rho}^*(\epsilon)$, we only need to relate $\max\{\Delta_{\mathbf{x}}, \epsilon\}$ with $\max\{\langle \boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle, \epsilon\}$. From Eq. (21) and the fact that $\epsilon \geq \gamma$, we know that

$$\langle \boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle \leq \Delta_{\mathbf{x}} + 2\gamma \leq \Delta_{\mathbf{x}} + 2\epsilon \leq 3 \max\{\Delta_{\mathbf{x}}, \epsilon\},$$

and thus

$$\max\{\langle \boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle, \epsilon\} \leq 3 \max\{\Delta_{\mathbf{x}}, \epsilon\}.$$

As a result, we have $\rho^*(\epsilon) \leq 9\tilde{\rho}^*(\epsilon)$.

When $\gamma < \Delta_{\min}/2$, we have

$$\langle \boldsymbol{\psi}_d(\mathbf{x}_*) - \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle \geq \Delta_{\mathbf{x}} - 2\gamma > 0,$$

which implies that \mathbf{x}_* is still the optimal arm in the perfect linear bandit problem $\tilde{\mathbf{h}}(\mathbf{x}) = \langle \boldsymbol{\psi}_d(\mathbf{x}), \boldsymbol{\theta}_d \rangle$ in \mathbb{R}^d . Thus, $\tilde{\rho}^*(0) \log(1/2.4\delta)$ gives the lower bound for best arm identification in the perfect linear bandit model [38, 13, 10]; and $\tilde{\rho}^*(\epsilon)$ can be viewed as a complexity measure with ϵ relaxation. \square

B.3 Example for high-dimensional linear bandits

Example. Let $\mathcal{X}_1 \oplus \mathcal{X}_2 = \{\mathbf{x}_i + \mathbf{x}_j : \mathbf{x}_i \in \mathcal{X}_1, \mathbf{x}_j \in \mathcal{X}_2\}$. Consider a high dimensional linear bandit problem in \mathbb{R}^D with an action set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_1 \oplus \{\eta \mathbf{e}_i\}_{i=1}^D, \mathbf{x}_2 \oplus \{\eta \mathbf{e}_i\}_{i=1}^D\}$. For any $\epsilon \in (0, 1/4)$, we select $\mathbf{x}_1, \mathbf{x}_2$ and $\eta > 0$ such that: (1) $\|\mathbf{x}_1\|_2 = 1$; (2) $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle = 1 - 2\epsilon$; (3) $\langle \mathbf{x}_1, \mathbf{e}_i \rangle < 0$; and (4) $8\sqrt{2}(2 + \sqrt{5D})D\eta \leq \epsilon$. We select $\boldsymbol{\theta}_* = \mathbf{x}_1$ (thus $\|\boldsymbol{\theta}_*\|_2 = 1$), which implies that arm $\mathbf{x}_* = \mathbf{x}_1$ is the optimal arm, arms in $\{\mathbf{x}_1 \oplus \{\eta \mathbf{e}_i\}_{i=1}^D\}$ are ϵ -optimal and arms in $\{\mathbf{x}_2 \oplus \{\eta \mathbf{e}_i\}_{i=1}^D\}$ have sub-optimality gap $> \epsilon$ but $< 3\epsilon$ (note that we must have $\eta < \epsilon$ be construction). As an example, one can select $\mathbf{x}_1 = -\mathbf{1}/\sqrt{D}$ ($\mathbf{1} \in \mathbb{R}^D$ represents the all 1 vector) and $\mathbf{x}_2 = (1 - 2\epsilon)\mathbf{x}_1$ with a sufficiently small η .

Analysis. Let $\mathbf{X} \in \mathbb{R}^{(2D+1) \times D}$ be the data matrix whose rows are $\mathbf{x} \in \mathcal{X}$; and let $\tilde{\mathbf{X}} \in \mathbb{R}^{(2D+1) \times D}$ be the matrix such that its first $(D+1)$ rows are all \mathbf{x}_1 and its last D rows are all \mathbf{x}_2 . One can clearly see that $\tilde{\mathbf{X}}$ is (at most) rank 2, and $\mathbf{X} = \tilde{\mathbf{X}} + \mathbf{E}$ where $\mathbf{E} = [\mathbf{0}^\top; \eta \mathbf{I}_D; \eta \mathbf{I}_D]$ (and thus $\|\mathbf{E}\|_2 = \sqrt{2}\eta$). Let $\{\sigma_i\}_{i=1}^D$ represents the singular values of matrix \mathbf{X} and let $\{\tilde{\sigma}_i\}_{i=1}^D$ represents the singular values of matrix $\tilde{\mathbf{X}}$. Perturbation theory on SVD [39] implies that

$$|\sigma_i - \tilde{\sigma}_i| \leq \|\mathbf{E}\|_2 = \sqrt{2}\eta, \forall i \in [D].$$

Since $\tilde{\sigma}_i = 0$ for $i > 2$, we have

$$\tilde{\gamma}(2) = \|\boldsymbol{\theta}_*\|_2 \cdot \left(\sum_{i>2} \sigma_i \right) < \sqrt{2}D\eta.$$

The fourth requirement in the construction implies that $d_{\text{eff}}(\epsilon) \leq 2$. To identify a ϵ -optimal arm, Algorithm 3 will first project the action set into a two-dimensional subspace and then perform arm elimination. Invoking Theorem 4, we know that the sample complexity is upper bounded by $\tilde{O}(1/\epsilon^2)$, which is independent of D (besides logarithmic dependence on $|\mathcal{X}|$).

We now analyze the sample complexity of RAGE [13]. To identifying an ϵ -optimal arm, we would run RAGE for $n = \lceil \log_2(2/\epsilon) \rceil$ iterations in \mathbb{R}^D , and its sample complexity upper bound scales as

$$\sum_{k=1}^n \max \left\{ 2(1 + \zeta) 2^{2k} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\mathcal{S}_k)} \|\mathbf{y}\|_{\mathbf{A}(\lambda)^{-1}}^2 \log(k^2 |\mathcal{X}|^2 / \delta), r_D(\zeta) \right\}.$$

Consider the round $k = \lfloor \log_2(1/\epsilon) \rfloor$. The sample complexity bound at that round is larger than

$$2^{2 \lfloor \log_2(1/\epsilon) \rfloor} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\mathcal{S}_{\lfloor \log_2(1/\epsilon) \rfloor})} \|\mathbf{y}\|_{\mathbf{A}(\lambda)^{-1}}^2 \geq \frac{1}{4\epsilon^2} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\mathcal{X})} \|\mathbf{y}\|_{\mathbf{A}(\lambda)^{-1}}^2,$$

since $\mathcal{S}_{\lfloor \log_2(1/\epsilon) \rfloor} = \mathcal{X}$ by construction (note that all arms in \mathcal{X} are 4ϵ -optimal). Since $\text{span}(\{\mathbf{x}_* - \mathbf{x}\}_{\mathbf{x} \in \mathcal{X}}) = \mathbb{R}^D$ by construction, applying Lemma 1 in Appendix A shows that

$$\frac{1}{4\epsilon^2} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\mathcal{X})} \|\mathbf{y}\|_{\mathbf{A}(\lambda)^{-1}}^2 = \Omega(D/\epsilon^2).$$

C Arm elimination with a fixed embedding

We only consider the case when $\tilde{\gamma}(d) > 0$. Otherwise, one can simply apply RAGE to identify the best arm (when it is unique) or to identify an ϵ -optimal arm (by stopping RAGE after $O(\log(1/\epsilon))$ rounds).

C.1 The algorithm

Algorithm 3 Arm Elimination with Fixed Embedding and Induced Misspecification

Input: Action set \mathcal{X} , confidence parameter δ , embedding dimension d and rounding approximation factor ζ .

- 1: Construct the compressed feature representation $\psi(\mathcal{X}) = \psi_d(\mathcal{X})$ with the induced misspecification level $\tilde{\gamma}(d)$.
- 2: Set $\hat{\mathcal{S}}_1 = \mathcal{S}_1$, $\gamma(d) = (16 + 8\sqrt{g(d, \zeta)}) \tilde{\gamma}(d)$, and $n = \lceil \log_2(2/\gamma(d)) \rceil$.
- 3: **for** $k = 1, 2, \dots, n$ **do**
- 4: Set $\delta_k = \delta/k^2$.
- 5: Set λ_k and τ_k be the design and the value of the following optimization problem

$$\inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi(\hat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\psi}(\lambda)^{-1}}^2.$$

- 6: Set $\epsilon_k = 2\tilde{\gamma}(d) + \tilde{\gamma}(d)\sqrt{(1 + \zeta)\tau_k}$,
and set $N_k = \max\{\lceil (2^{-k} - \epsilon_k)^{-2} 2(1 + \zeta)\tau_k \log(|\hat{\mathcal{S}}_k|^2/\delta_k) \rceil, r_d(\zeta)\}$.
- 7: Get $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\} = \text{ROUND}(\lambda_k, N_k, d, \zeta)$.
- 8: Pull arms $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ and receive rewards $\{y_1, \dots, y_{N_k}\}$.
- 9: Set $\hat{\boldsymbol{\theta}}_k = \mathbf{A}_k^{-1} \mathbf{b}_k$, where $\mathbf{A}_k = \sum_{i=1}^{N_k} \psi(\mathbf{x}_i) \psi(\mathbf{x}_i)^\top$ and $\mathbf{b}_k = \sum_{i=1}^{N_k} \psi(\mathbf{x}_i) y_i$.
- 10: Eliminate arms with respect to criteria

$$\hat{\mathcal{S}}_{k+1} = \hat{\mathcal{S}}_k \setminus \{\psi(\mathbf{x}) \in \mathcal{S}_k : \exists \mathbf{x}' \text{ such that } (\psi(\mathbf{x}') - \psi(\mathbf{x}))^\top \hat{\boldsymbol{\theta}}_k \geq \omega_k(\psi(\mathbf{x}') - \psi(\mathbf{x}))\},$$

$$\text{where } \omega_k(\mathbf{y}) = \epsilon_k + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\hat{\mathcal{S}}_k|^2/\delta_k)}.$$

11: **end for**

Output: Output any arm in $\hat{\mathcal{S}}_{n+1}$.

C.2 Sample complexity analysis

Theorem 4. *With probability of at least $1 - \delta$, Algorithm 3 correctly outputs an $\gamma(d)$ -optimal arm with sample complexity upper bounded by*

$$N \leq 640 \rho_d^*(\gamma(d)) \log_2 \left(\frac{4}{\gamma(d)} \right) \log \left(\frac{\log_2(4/\gamma(d)) |\mathcal{X}|^2}{\delta} \right) + (r_d(\zeta) + 1) \log_2 \left(\frac{4}{\gamma(d)} \right). \quad (22)$$

Remark 6. *Algorithm 3 can be viewed as a modified version of Algorithm 1 with a fixed embedding. When we choose $d = d_{\text{eff}}(\epsilon)$ as the input, Algorithm 3 will output an ϵ -optimal arm.*

Proof. The proof is very similar to the proof of Theorem 1 (we re-define $\psi(\cdot) := \psi_d(\cdot)$ with the fixed d).

We first notice that

$$\begin{aligned} \epsilon_k &\leq (2 + \sqrt{g(d, \zeta)}) \tilde{\gamma}(d) \\ &= \gamma(d)/8 \\ &\leq 2^{-\lceil \log_2(2/\gamma(d)) \rceil} / 2 \\ &\leq 2^{-k} / 2, \end{aligned}$$

for any $1 \leq k \leq n = \lceil \log_2(2/\gamma(d)) \rceil$. Thus, N_k is well-defined and Step 1 and 2 from the proof of Theorem 1 hold true.

We now analyze the sample complexity. As one can see, the total sample complexity N is upper bounded by

$$\begin{aligned} N &\leq \sum_{k=1}^n N_k \\ &= \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} \max \left\{ \left[(2^{-k} - \epsilon_k)^{-2} 2(1 + \zeta) \tau(\mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))) \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k) \right], r_d(\zeta) \right\} \\ &\leq \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} \left((2^{-k} - \epsilon_k)^{-2} 2(1 + \zeta) \tau(\mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))) \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k) \right) + (r_d(\zeta) + 1) \log_2(4/\gamma(d)) \\ &\leq \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} (2^{2k}) 10 \tau(\mathcal{Y}(\psi(\mathcal{S}_k))) \log(k^2 |\mathcal{X}|^2 / \delta) + (r_d(\zeta) + 1) \log_2(4/\gamma(d)), \end{aligned} \quad (23)$$

where we use the fact that $\epsilon_k \leq 2^{-k}/2$, $\widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k \subseteq \mathcal{X}$, $\delta_k = \delta/k^2$ and $\zeta \leq 1/4$ in Eq. (23).

We analyze the instance dependent sample complexity $\rho^*(\gamma(d))$ next.

$$\begin{aligned}
\rho^*(\gamma(d)) &= \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}} \frac{\|\psi(\mathbf{x}_*) - \psi(\mathbf{x})\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2}{\max\{\Delta(\mathbf{x}), \gamma(d)\}^2} \\
&= \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{k \leq \lceil \log_2(2/\gamma(d)) \rceil} \sup_{\mathbf{x} \in \mathcal{S}_k \setminus \{\mathbf{x}_*\}} \frac{\|\psi(\mathbf{x}_*) - \psi(\mathbf{x})\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2}{\max\{\Delta(\mathbf{x}), \gamma(d)\}^2} \\
&\geq \sup_{k \leq \lceil \log_2(2/\gamma(d)) \rceil} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_k \setminus \{\mathbf{x}_*\}} \frac{\|\psi(\mathbf{x}_*) - \psi(\mathbf{x})\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2}{\max\{\Delta(\mathbf{x}), \gamma(d)\}^2} \\
&\geq \sup_{k \leq \lceil \log_2(2/\gamma(d)) \rceil} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_k \setminus \{\mathbf{x}_*\}} \frac{\|\psi(\mathbf{x}_*) - \psi(\mathbf{x})\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2}{\max\{4 \cdot 2^{-k}, \gamma(d)\}^2} \\
&\geq \frac{1}{\lceil \log_2(2/\gamma(d)) \rceil} \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x} \in \mathcal{S}_k \setminus \{\mathbf{x}_*\}} \frac{\|\psi(\mathbf{x}_*) - \psi(\mathbf{x})\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2}{\max\{4 \cdot 2^{-k}, \gamma(d)\}^2} \\
&\geq \frac{1}{\lceil \log_2(2/\gamma(d)) \rceil} \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_k} \frac{\|\psi(\mathbf{x}) - \psi(\mathbf{x}')\|_{\mathbf{A}_{\psi(\lambda)}^{-1}}^2/4}{\max\{4 \cdot 2^{-k}, \gamma(d)\}^2} \quad (24) \\
&= \frac{1}{\lceil \log_2(2/\gamma(d)) \rceil} \sum_{k=1}^{\lceil \log_2(2/\gamma(d)) \rceil} (2^{2k}) \tau(\mathcal{Y}(\psi(\mathcal{S}_k)))/64, \quad (25)
\end{aligned}$$

where Eq. (24) comes from the fact that we can write $\psi(\mathbf{x}) - \psi(\mathbf{x}') = \psi(\mathbf{x}) - \psi(\mathbf{x}_*) + \psi(\mathbf{x}_*) - \psi(\mathbf{x}')$ and apply triangle inequality; Eq. (25) comes from the definition of $\tau(\mathcal{Y}(\psi(\mathcal{S}_k)))$ and the fact that $4 \cdot 2^{-k} \geq \gamma(d)$ when $1 \leq k \leq \lceil \log_2(2/\gamma(d)) \rceil$.

By comparing terms in Eq. (23) and terms in Eq. (25), we now relate the complexity bound N with the instance-dependent complexity $\rho^*(\gamma(d))$ as follows:

$$N \leq 640 \rho^*(\gamma(d)) \log_2\left(\frac{4}{\gamma(d)}\right) \log\left(\frac{\log_2(4/\gamma(d)) |\mathcal{X}|^2}{\delta}\right) + (r_d(\zeta) + 1) \log_2\left(\frac{4}{\gamma(d)}\right).$$

□

D Arm elimination with unknown misspecifications

We consider of version of Algorithm 3 with unknown mis-specification level $\tilde{\gamma}(\cdot)$. As explained in Appendix C, we only consider the case when $\tilde{\gamma}(d) > 0$. We (re-)define some notations for this section as follows.

We set $n(d)$ as a critical value for the number of iteration

$$n(d) = \max\left\{\tilde{k} \in \mathbb{N} : \forall k \leq \tilde{k}, \left(2 + \sqrt{(1 + \zeta)\tau(\mathcal{Y}(\psi_d(\mathcal{S}_k)))}\right) \tilde{\gamma}(d) \leq 2^{-k}/2\right\},$$

which captures the largest number of iterations such that the induced misspecification error is still well-controlled. Note that $n(d)$ crucially depends on the action set $\psi_d(\mathcal{X})$ and cannot be calculated due to unknown identities of $\{\mathcal{S}_k\}$. We then define $\bar{\gamma}(d)$, which quantifies the optimality of the identified arm as $\bar{\gamma}(d) = 2 \cdot 2^{-n(d)}$. We also introduce the notation $g_k(d, \zeta) = (1 + \zeta)\tau(\mathcal{Y}(\psi_d(\mathcal{S}_k)))$.

Proposition 2. *The following inequalities hold.*

$$n(d) \geq \lfloor \log_2(1/(2(2 + \sqrt{g(d, \zeta)})\tilde{\gamma}(d))) \rfloor \quad \text{and} \quad \bar{\gamma}(d) \leq 8(2 + \sqrt{g(d, \zeta)})\tilde{\gamma}(d) = \gamma(d).$$

Proof. Since $\tau(\mathcal{Y}(\psi_d(\mathcal{S}_k))) \leq \tau(\mathcal{Y}(\psi_d(\mathcal{X})))$ according to Kiefer-Wolfowitz theorem [27] and $\zeta \in [1/10, 1/4]$, we have

$$2 + \sqrt{(1 + \zeta)\tau(\mathcal{Y}(\psi_d(\mathcal{S}_k)))} \leq 2 + \sqrt{g(d, \zeta)},$$

which implies that

$$n(d) \geq \lfloor \log_2(1/(2(2 + \sqrt{g(d, \zeta)})\tilde{\gamma}(d))) \rfloor.$$

The upper bound on $\bar{\gamma}(d)$ immediately follows. □

Algorithm 4 Arm Elimination with Fixed Embedding and Unknown Misspecification

Input: Action set \mathcal{X} , confidence parameter δ , embedding dimension d and rounding approximation factor ζ .

- 1: Construct the compressed feature representation $\psi(\mathcal{X}) = \psi_d(\mathcal{X})$ with the induced misspecification level $\tilde{\gamma}(d)$.
- 2: Set $\widehat{\mathcal{S}}_1 = \mathcal{X}$, and randomly select $\widehat{\mathbf{x}}_* \in \widehat{\mathcal{S}}_1$ as the recommendation.
- 3: **for** $k = 1, 2, \dots$ **do**
- 4: Set $\delta_k = \delta/k^2$.
- 5: Set λ_k and τ_k be the design and the value of the following optimization problem

$$\inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\psi}(\lambda)}^{-1}.$$

- 6: Set $N_k = \max\{\lceil 2^{2k} 8(1 + \zeta) \tau_k \log(|\widehat{\mathcal{S}}_k|^2/\delta_k) \rceil, r_d(\zeta)\}$.
- 7: Set $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\} = \text{ROUND}(\lambda_k, N_k, d, \zeta)$.
- 8: Pull arms $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_k}\}$ and receive rewards $\{y_1, \dots, y_{N_k}\}$.
- 9: Set $\widehat{\boldsymbol{\theta}}_k = \mathbf{A}_k^{-1} \mathbf{b}_k$, where $\mathbf{A}_k = \sum_{i=1}^{N_k} \psi(\mathbf{x}_i) \psi(\mathbf{x}_i)^\top$ and $\mathbf{b}_k = \sum_{i=1}^{N_k} \psi(\mathbf{x}_i) y_i$.
- 10: Update $\widehat{\mathbf{x}}_* = \arg \max_{\mathbf{x} \in \widehat{\mathcal{S}}_k} \langle \psi(\mathbf{x}), \widehat{\boldsymbol{\theta}}_k \rangle$, and eliminate arms with respect to criteria

$$\widehat{\mathcal{S}}_{k+1} = \widehat{\mathcal{S}}_k \setminus \{\psi(\mathbf{x}) \in \widehat{\mathcal{S}}_k : (\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k \geq 2^{-k}\}.$$

11: **end for**

Theorem 5. Suppose $\tilde{\gamma}(d)$ is unknown. With probability of at least $1 - \delta$, Algorithm 4 starts to output a $\bar{\gamma}(d)$ -optimal arm after at most N samples where

$$N \leq 640 \rho^*(\bar{\gamma}(d)) \log_2(n(d)) \log\left(\frac{(n(d)|\mathcal{X}|)^2}{\delta}\right) + (r(\zeta) + 1) n(d).$$

Remark 7. An arm with slightly smaller optimality gap $\bar{\gamma}(d)$ is identified, with slightly larger sample complexity, in the unverifiable way by Algorithm 4. In fact, the same can be achieved by Algorithm 3 if an upper bound on the number of iteration k is not set. The optimality gap $\bar{\gamma}(d)$, however, cannot be quantified without knowing $h(\cdot)$. That's the reason for us to set an upper bound on the number of iterations in Algorithm 3.

Proof. The proof is similar to the proof of Theorem 1 and Theorem 4, we provide it here for completeness. We decompose the proof into three steps: (1) prove correctness through induction; (2) bound the error probability; and (3) upper bound the sample complexity. For simplicity, we'll use notations $\psi(\cdot) := \psi_d(\cdot)$, $\boldsymbol{\theta}_* := \boldsymbol{\theta}_d$ and $\eta(\cdot) = \eta_d(\cdot)$ in the proof.

Step 1: The induction. We define event

$$\mathcal{E}_k = \{\mathbf{x}_* \in \widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k\},$$

and prove through induction that

$$\mathbb{P}(\mathcal{E}_{k+1} | \cap_{i \leq k} \mathcal{E}_i) \geq 1 - \delta_k,$$

where $\delta_0 := 0$ and $k \leq n(d)$. Recall that $\mathcal{S}_k = \{x \in \mathcal{X} : \Delta_x < 4 \cdot 2^{-k}\}$ (with $\mathcal{S}_1 = \mathcal{X}$), and thus $\mathcal{S}_{n(d)+1}$ contains the set of $\bar{\gamma}(d)$ -optimal arms. Note that as long as $\widehat{\mathcal{S}}_{k+1} \subseteq \mathcal{S}_{k+1}$ holds true, Algorithm 4 will have $\widehat{\mathcal{S}}_n \subseteq \widehat{\mathcal{S}}_{k+1}$ for $n > k + 1$ due to the nature of the elimination-styled algorithm.

The base case $\{\mathbf{x}_* \in \widehat{\mathcal{S}}_1 \subseteq \mathcal{S}_1\}$ holds by definition of $\widehat{\mathcal{S}}_1 = \mathcal{S}_1 = \mathcal{X}$. We next analyze event \mathcal{E}_{k+1} conditioned on $\cap_{i \leq k} \mathcal{E}_i$.

Step 1.1: Concentration under mis-specifications. This step is almost the same as the Step 1.1 as analyzed in the proof of Theorem 1 since the unknown mis-specification level $\tilde{\gamma}(d)$ is only used in the analysis. As a result, with probability at least $1 - \delta_k$, for any $\mathbf{y} \in \mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))$,

$$\begin{aligned} \left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_* \rangle \right| &\leq \tilde{\gamma}(d) \sqrt{(1 + \zeta) \tau(\mathcal{Y}(\psi(\widehat{\mathcal{S}}_k)))} + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2/\delta_k)} \\ &\leq \sqrt{g_k(d, \zeta)} \tilde{\gamma}(d) + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2/\delta_k)}, \end{aligned} \quad (26)$$

where Eq. (26) comes from the fact that $\tau(\mathcal{Y}(\psi(\widehat{\mathcal{S}}_k))) \leq \tau(\mathcal{Y}(\psi(\mathcal{S}_k)))$. In the following, we define

$$\omega_k(\mathbf{y}) := \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)}.$$

Step 1.2: Correctness. We prove the result under the event described in Eq. (26), which holds with (conditional) probability at least $1 - \delta_k$. We know that $\mathbf{x}_* \in \widehat{\mathcal{S}}_k \subseteq \mathcal{S}_k$ holds conditioned on $\cap_{i \leq k} \mathcal{E}_i$. We now prove it for iteration $k + 1$. Recall the elimination criteria is

$$\widehat{\mathcal{S}}_{k+1} = \widehat{\mathcal{S}}_k \setminus \{\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x})\}^\top \widehat{\boldsymbol{\theta}}_k \geq 2^{-k}\},$$

where $\widehat{\mathbf{x}}_* = \arg \max_{\mathbf{x} \in \widehat{\mathcal{S}}_k} \langle \psi(\mathbf{x}), \widehat{\boldsymbol{\theta}}_k \rangle$.

Step 1.2.1: $\mathbf{x}_* \in \widehat{\mathcal{S}}_{k+1}$. We trivially have $\mathbf{x}_* \in \widehat{\mathcal{S}}_{k+1}$ if $\widehat{\mathbf{x}}_* = \mathbf{x}_*$. Suppose $\widehat{\mathbf{x}}_* \neq \mathbf{x}_*$, we have

$$\begin{aligned} (\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*))^\top \widehat{\boldsymbol{\theta}}_k &\leq (\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*))^\top \boldsymbol{\theta}^* + \sqrt{g_k(d, \zeta)} \widetilde{\gamma}(d) + \omega_k(\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*)) \\ &= h(\widehat{\mathbf{x}}_*) - \eta(\widehat{\mathbf{x}}_*) - (h(\mathbf{x}_*) - \eta(\mathbf{x}_*)) + \sqrt{g_k(d, \zeta)} \widetilde{\gamma}(d) + \omega_k(\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*)) \\ &\leq -\Delta(\widehat{\mathbf{x}}_*) + \left(2\widetilde{\gamma}(d) + \sqrt{g_k(d, \zeta)} \widetilde{\gamma}(d) + \omega_k(\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*))\right) \end{aligned} \quad (27)$$

$$\begin{aligned} &< 2^{-k}/2 + 2^{-k}/2 \\ &= 2^{-k}, \end{aligned} \quad (28)$$

where Eq. (28) comes from the definition of $n(d)$ and the fact that $\omega_k(\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*)) \leq 2^{-k}/2$ (due to the selection of N_k and the guarantee of the rounding procedure in Eq. (1)). As a result, the optimal arm \mathbf{x}_* remains in $\widehat{\mathcal{S}}_{k+1}$.

Step 1.2.2: $\widehat{\mathcal{S}}_{k+1} \subseteq \mathcal{S}_{k+1}$. For any $\mathbf{x} \in \mathcal{S}_{k+1}^c \cap \widehat{\mathcal{S}}_k$, we have $\Delta(\mathbf{x}) = h(\mathbf{x}_*) - h(\mathbf{x}) \geq 4 \cdot 2^{-(k+1)} = 2 \cdot 2^{-k}$ by definition. Since $(\psi(\widehat{\mathbf{x}}_*) - \psi(\mathbf{x}_*))^\top \widehat{\boldsymbol{\theta}}_k \geq 0$, we only need to lower bound $(\psi(\mathbf{x}_*) - \psi(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k$.

$$\begin{aligned} (\psi(\mathbf{x}_*) - \psi(\mathbf{x}))^\top \widehat{\boldsymbol{\theta}}_k &\geq (\psi(\mathbf{x}_*) - \psi(\mathbf{x}))^\top \boldsymbol{\theta}^* - \sqrt{g_k(d, \zeta)} \widetilde{\gamma}(d) - \omega_k(\psi(\mathbf{x}_*) - \psi(\mathbf{x})) \\ &= h(\mathbf{x}_*) - \eta(\mathbf{x}_*) - (h(\mathbf{x}) - \eta(\mathbf{x})) - \sqrt{g_k(d, \zeta)} \widetilde{\gamma}(d) - \omega_k(\psi(\mathbf{x}_*) - \psi(\mathbf{x})) \\ &\geq 2 \cdot 2^{-k} - (2 + \sqrt{g_k(d, \zeta)}) \widetilde{\gamma}(d) - \omega_k(\psi(\mathbf{x}_*) - \psi(\mathbf{x})) \\ &\geq 2 \cdot 2^{-k} - 2^{-k}/2 - 2^{-k}/2 \\ &\geq 2^{-k}, \end{aligned} \quad (29)$$

where Eq. (29) comes from a similar analysis as in Eq. (28). As a result, any $\mathbf{x} \in \mathcal{S}_{k+1}^c \cap \widehat{\mathcal{S}}_k$ will be eliminated and we have $\widehat{\mathcal{S}}_{k+1} \subseteq \mathcal{S}_{k+1}$.

To summarize, we prove the induction at iteration $k + 1$, i.e.,

$$\mathbb{P}(\mathcal{E}_{k+1} | \cap_{i < k+1} \mathcal{E}_i) \geq 1 - \delta_k.$$

Step 2: The error probability. Using exactly the same argument as appeared in Step 2 in the proof of Theorem 4, we have

$$\mathbb{P}(\cap_{i=1}^{k+1} \mathcal{E}_i) \geq 1 - \delta.$$

Step 3: Sample complexity upper bound. Up to iteration $k \leq n(d)$, the sample complexity of Algorithm 4 can be upper bounded by

$$\begin{aligned} N &= \sum_{k=1}^{n(d)} N_k \\ &= \sum_{k=1}^{n(d)} \max \left\{ \left\lceil 2^{2k} 8(1 + \zeta) \tau_k \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k) \right\rceil, r_d(\zeta) \right\} \\ &\leq \sum_{k=1}^{n(d)} \left(2^{2k} 10 \tau(\mathcal{Y}(\psi(\mathcal{S}_k))) \log(k^2 |\mathcal{X}|^2 / \delta) \right) + (r_d(\zeta) + 1) n(d). \end{aligned} \quad (30)$$

Following a similar analysis as in Eq. (25) (note that $4 \cdot 2^{-k} \geq \bar{\gamma}(d)$ for $k \leq n(d)$ by definition), we bound the instance dependent sample complexity as follows.

$$N \leq 640 \rho^*(\bar{\gamma}(d)) \log_2(n(d)) \log\left(\frac{(n(d)|\mathcal{X}|)^2}{\delta}\right) + (r(\zeta) + 1) n(d).$$

□

E Omitted materials for Section 4

E.1 Mercer's theorem and corollary

Let \mathbb{P} be a non-negative measure over the compact metric space \mathcal{Z} , we define the kernel integral operator as follows

$$T_{\mathcal{K}}(f)(\mathbf{x}) = \int_{\mathcal{Z}} \mathcal{K}(\mathbf{x}, \mathbf{z}) f(\mathbf{z}) d\mathbb{P}(\mathbf{z}).$$

Theorem 6 (Mercer's theorem, see Wainwright [44]). *Suppose \mathcal{Z} is compact, the kernel function \mathcal{K} is continuous, positive semi-definite and satisfies the Hilbert-Schmidt condition. Then there exist a sequence of eigenfunctions $\{\phi_j\}_{j=1}^{\infty}$ that form an orthonormal basis of $L^2(\mathcal{Z}; \mathbb{P})$, and non-negative eigenvalues $\{\mu_j\}_{j=1}^{\infty}$ such that*

$$T_{\mathcal{K}}(\phi_j) = \mu_j \phi_j \quad \text{for } j = 1, 2, \dots$$

Moreover, the kernel function has the expansion

$$\mathcal{K}(\mathbf{x}, \mathbf{z}) = \sum_{j=1}^{\infty} \mu_j \phi_j(\mathbf{x}) \phi_j(\mathbf{z}).$$

Corollary 1 (Mercer's corollary, see Wainwright [44]). *Consider a kernel satisfying the conditions of Mercer's theorem with associated eigenfunctions $\{\phi_j\}_{j=1}^{\infty}$ and non-negative eigenvalues $\{\mu_j\}_{j=1}^{\infty}$. It induces the reproducing kernel Hilbert space*

$$\mathcal{H} := \left\{ h = \sum_{j=1}^{\infty} \theta_j \phi_j : \text{for some } \{\theta_j\}_{j=1}^{\infty} \in \ell^2(\mathbb{N}) \text{ with } \sum_{j=1}^{\infty} \theta_j^2 / \mu_j < \infty \right\},$$

along with inner product

$$\langle h_1, h_2 \rangle_{\mathcal{H}} := \sum_{j=1}^{\infty} \frac{\langle h_1, \phi_j \rangle \langle h_2, \phi_j \rangle}{\mu_j},$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product in $L^2(\mathcal{Z}; \mathbb{P})$.

Remark 8. *The RKHS induced by any kernel is unique.*

E.2 Proof of Theorem 2

Theorem 2. *Suppose Assumption 1 holds. For any $\epsilon > 0$, the following statements hold when we specialize Algorithm 1 to the kernel setting.*

1. *Suppose \mathcal{K} has (C_k, β) -polynomial eigenvalue decay. We have $d_{\text{eff}}(\epsilon) \leq O(\epsilon^{-2/(2\beta-3)})$, and the sample complexity of identifying an ϵ -optimal arm is upper bounded by $\tilde{O}(\epsilon^{-2/(2\beta-3)} \max\{\Delta_{\min}, \epsilon\}^{-2})$.*
2. *Suppose \mathcal{K} has (C_k, β) -exponential eigenvalue decay. We have $d_{\text{eff}}(\epsilon) \leq O(\log(1/\epsilon))$, and the sample complexity of identifying an ϵ -optimal arm is upper bounded by $\tilde{O}(\max\{\Delta_{\min}, \epsilon\}^{-2})$.*

Proof. We only need calculate the $d_{\text{eff}}(\epsilon)$ under different assumptions, and the sample complexity bounds follow from Theorem 1. Since $g(d, \zeta) \leq 5d$, we have

$$d_{\text{eff}}(\epsilon) \leq \min \left\{ d \geq 1 : (16 + 8\sqrt{5d})C_\phi \sqrt{\sum_{j>d} \mu_j} \leq \epsilon \right\}.$$

We analyze an upper bound of $d_{\text{eff}}(\epsilon)$ in the following.

Case 1: Polynomial eigenvalue decay. When \mathcal{K} has (C_k, β) -polynomial eigenvalue decay, we have, with $\beta > 3$, $\mu_j \leq C_k j^{-\beta}$ for all $j \geq 1$. Since

$$\sqrt{\sum_{j>d} \mu_j} \leq C_k^{1/2} \sqrt{\int_{x=d}^{\infty} x^{-\beta} dx} \leq \frac{C_k^{1/2}}{\beta-1} d^{1-\beta},$$

we then have

$$\begin{aligned} (16 + 8\sqrt{5d})C_\phi \sqrt{\sum_{j>d} \mu_j} &\leq \frac{(16 + 8\sqrt{5d})C_\phi C_k^{1/2} d^{1-\beta}}{\beta-1} \\ &\leq \frac{16\sqrt{5}C_\phi C_k^{1/2} d^{3/2-\beta}}{\beta-1}, \end{aligned} \quad (31)$$

where we use the fact that $2 \leq \sqrt{5d}$ for any $d \geq 1$ in Eq. (31). Since Eq. (31) is decaying in d (recall $\beta > 3/2$), it suffices to set

$$d = \left\lceil \left(\frac{16\sqrt{5}C_\phi C_k^{1/2}}{\epsilon(\beta/2-1)} \right)^{2/(2\beta-3)} \right\rceil = O\left(\epsilon^{-2/(2\beta-3)}\right)$$

to make it smaller than or equal to ϵ . This also gives an upper bound of $d_{\text{eff}}(\epsilon)$.

Case 2: Exponential eigenvalue decay. When \mathcal{K} has (C_k, β) -exponential eigenvalue decay, we have, with $\beta > 0$, $\mu_j \leq C_k \exp(-\beta j)$ for all $j \geq 1$. Since

$$\sqrt{\sum_{j>d} \mu_j} \leq \frac{C_k^{1/2} e^{-\beta d}}{\beta},$$

we then have (with $2 \leq \sqrt{5d}$ for any $d \geq 1$)

$$\begin{aligned} (16 + 8\sqrt{5d})C_\phi \sqrt{\sum_{j>d} \mu_j} &\leq \frac{(16 + 8\sqrt{5d})C_\phi C_k^{1/2} e^{-\beta d}}{\beta} \\ &\leq \frac{16\sqrt{5}C_\phi C_k^{1/2} \sqrt{d} e^{-\beta d}}{\beta} \\ &= \frac{16\sqrt{5}C_\phi C_k^{1/2}}{\beta} e^{1/2 \cdot \log d - \beta d} \end{aligned}$$

To have $(16 + 8\sqrt{5d})C_\phi \sqrt{\sum_{j>d} \mu_j} \leq \epsilon$, it suffices to have

$$e^{\beta d - 1/2 \cdot \log d} \geq \frac{1}{\tau} := \frac{16\sqrt{5}C_\phi C_k^{1/2}}{\epsilon\beta}.$$

Since $x \geq 2a \log a \implies x \geq a \log x$ for any $a > 0$ [37]. We know that $\beta d - 1/2 \cdot \log d \geq \beta d/2$ when $d \geq 2/\beta \cdot \log(1/\beta)$. Thus, it suffices to have

$$\begin{aligned} d &\geq \left\lceil \max \left\{ \frac{2 \log(1/\beta)}{\beta}, \frac{2 \log \left(\frac{16\sqrt{5}C_\phi C_k^{1/2}}{\epsilon\beta} \right)}{\beta} \right\} \right\rceil \\ &= \left\lceil 2 \log \left(\max \left\{ \frac{16\sqrt{5}C_\phi \sqrt{C_k}}{\epsilon\beta}, \frac{1}{\beta} \right\} \right) / \beta \right\rceil = O(\log(1/\epsilon)). \end{aligned}$$

This gives an upper bound of $d_{\text{eff}}(\epsilon)$. □

F Details and proofs for Section 5

F.1 Omitted details for Section 5

Details of Gradient Descent For the completeness, we present the gradient descent algorithm in Algorithm 5, which is originally used in [47] as a subroutine for the NeuralUCB algorithm.

Algorithm 5 Gradient Descent

- 1: **Input:** Regularization parameter α , step size η , number of gradient descent steps J , network width m , contexts $\{\mathbf{x}_i\}_{i=1}^N$, rewards $\{r_i\}_{i=1}^N$, initial parameter $\boldsymbol{\theta}^{(0)}$.
 - 2: Define $\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^N (f(\mathbf{x}_i; \boldsymbol{\theta}) - r_i)^2/2 + m\alpha \|\boldsymbol{\theta} - \boldsymbol{\theta}^{(0)}\|_2^2/2$.
 - 3: **for** $j = 0, \dots, J-1$ **do**
 - 4: $\boldsymbol{\theta}^{(j+1)} = \boldsymbol{\theta}^{(j)} - \eta \nabla \mathcal{L}(\boldsymbol{\theta}^{(j)})$
 - 5: **end for**
 - 6: **Return** $\boldsymbol{\theta}^{(J)}$.
-

Definition of Neural Tangent Kernel We define the Neural Tangent Kernel (NTK) gram matrix as follows.

Definition 3 ([17, 47]). Given arm set \mathcal{X} , define $\mathbf{H}^{(l)}, \widetilde{\mathbf{H}}^{(l)}, \boldsymbol{\Sigma}^{(l)} \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ as follows:

$$\begin{aligned} \widetilde{\mathbf{H}}^{(1)}(\mathbf{x}, \mathbf{x}') &= \boldsymbol{\Sigma}^{(1)}(\mathbf{x}, \mathbf{x}') = \langle \mathbf{x}, \mathbf{x}' \rangle, \mathbf{A}^{(l)}(\mathbf{x}, \mathbf{x}') = \begin{pmatrix} \boldsymbol{\Sigma}^{(l)}(\mathbf{x}, \mathbf{x}) & \boldsymbol{\Sigma}^{(l)}(\mathbf{x}, \mathbf{x}') \\ \boldsymbol{\Sigma}^{(l)}(\mathbf{x}, \mathbf{x}') & \boldsymbol{\Sigma}^{(l)}(\mathbf{x}', \mathbf{x}') \end{pmatrix}, \\ \boldsymbol{\Sigma}^{(l+1)}(\mathbf{x}, \mathbf{x}') &= 2\mathbb{E}_{(u,v) \sim N(\mathbf{0}, \mathbf{A}^{(l)}(\mathbf{x}, \mathbf{x}'))} \max\{u, 0\} \max\{v, 0\}, \\ \widetilde{\mathbf{H}}^{(l+1)}(\mathbf{x}, \mathbf{x}') &= 2\widetilde{\mathbf{H}}^{(l)}(\mathbf{x}, \mathbf{x}') \mathbb{E}_{(u,v) \sim N(\mathbf{0}, \mathbf{A}^{(l)}(\mathbf{x}, \mathbf{x}'))} \mathbb{1}(u \geq 0) \mathbb{1}(v \geq 0) + \boldsymbol{\Sigma}^{(l+1)}(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

Then, $\mathbf{H} = (\widetilde{\mathbf{H}}^{(L)} + \boldsymbol{\Sigma}^{(L)})/2$ is called the neural tangent kernel matrix on the arm set \mathcal{X} .

The gram matrix \mathbf{H} is defined over all arms \mathcal{X} , and $\mathbf{H}(\mathbf{x}, \mathbf{x}')$ can be regarded as the limitation of $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)^\top \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0)/m$ when m goes infinity [17]. \mathbf{H} plays an important role in our theoretical analysis of Algorithm 2.

Formal Version of Theorem 3 Next we present the formal version of Theorem 3.

Theorem 7. Under Assumption 2, let $\mathbf{h} := [h(\mathbf{x})] \in \mathbb{R}^{|\mathcal{X}|}$ denote the vector consisting of all rewards from arm set \mathcal{X} . Let S be some constant satisfying $S \geq \sqrt{\mathbf{h}^\top \mathbf{H}^{-1} \mathbf{h}}$. Then there exist constants $C_1, C_2 > 0$ such that for any ϵ , if we set the parameters in Algorithm 2 as follows:

$$\begin{aligned} \alpha &= \min\{1, \log(|\mathcal{X}|^2)/S\}, \quad n = \lceil \log(1/\epsilon) \rceil, \\ \bar{\epsilon}^2 &= \min\{\alpha^2/(r_{|\mathcal{X}|}^2(\zeta)L), \epsilon^2/((1+\zeta)|\mathcal{X}|S), \epsilon^7 \alpha^3/((1+\zeta)^3 |\mathcal{X}|^3 \log^3(|\mathcal{X}|^2/\delta_n)L^2)\}, \\ A &= d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|), \quad \eta_k = C_1(m\alpha + N_k mL)^{-1}, \quad J_k = \log(C_2 \epsilon^2 \alpha / (N_k L)) / (\eta_k mL), \end{aligned}$$

then when $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, with probability at least $1 - \delta$, $\widehat{\mathcal{S}}_{n+1}$ only includes arm \mathbf{x} satisfying $\Delta_{\mathbf{x}} \leq \epsilon$, and the total sample complexity of Algorithm 2 is bounded by

$$N = \widetilde{O}\left((1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)/\epsilon^2 + r_{d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)}(\zeta)\right) = \widetilde{O}\left(d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)\epsilon^{-2}\right).$$

Remark 9. Note that $N_k \sim 4^k$, therefore m actually depends on ϵ polynomially given the selection of m .

Remark 10. It is worth noting that the parameter selection depends on S , which is the upper bound of the quantity $\sqrt{\mathbf{h}^\top \mathbf{H}^{-1} \mathbf{h}}$. Since $\sqrt{\mathbf{h}^\top \mathbf{H}^{-1} \mathbf{h}} \leq \|\mathbf{h}\|_2 \lambda_{\min}(\mathbf{H})^{-1/2} \leq \sqrt{|\mathcal{X}|/\lambda_0}$, $S = \sqrt{|\mathcal{X}|/\lambda_0}$ is always a valid upper bound. Moreover, if the reward function h belongs to the RKHS space spanned by NTK with norm $\|h\|_{\mathcal{H}}$, then we can set a tighter upper bound $S = \|h\|_{\mathcal{H}}$ which ensures $S \geq \sqrt{\mathbf{h}^\top \mathbf{H}^{-1} \mathbf{h}}$ [47].

E.2 Proof of Theorem 7

To prove Theorem 7, we need the following lemmas. For simplicity, we first denote $\bar{\mathbf{Z}}_k, \mathbf{Z}_k \in \mathbb{R}^{p \times p}, \bar{\mathbf{b}}_k \in \mathbb{R}^p, \mathbf{A}_k \in \mathbb{R}^{d_k \times d_k}, \mathbf{b}_k, \hat{\boldsymbol{\theta}}_k \in \mathbb{R}^{d_k}$ as follows.

$$\begin{aligned} \mathbf{Z}_k &= \alpha \mathbf{I} + \sum_{i=1}^{N_k} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_{k-1}) \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_{k-1})^\top / m, \quad \bar{\mathbf{Z}}_k = \alpha \mathbf{I} + \sum_{i=1}^{N_k} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0)^\top / m, \\ \bar{\mathbf{b}}_k &= \sum_{i=1}^{N_k} y_i \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m}, \\ \mathbf{A}_k &= \alpha \mathbf{I} + \sum_{i=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_i) \boldsymbol{\psi}_{d_k}(\mathbf{x}_i)^\top, \quad \mathbf{b}_k = \sum_{i=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_i) y_i, \quad \hat{\boldsymbol{\theta}}_k = \mathbf{A}_k^{-1} \mathbf{b}_k, \end{aligned} \quad (32)$$

where $\{\mathbf{x}_i\}_{i=1}^{N_k}$ are arms selected at round k . We first present several lemmas which provide error bounds between neural network functions and their linear estimations.

Lemma 2 (Theorem 3.1, [2]). *Fix $\epsilon > 0$ and $\delta \in (0, 1)$. Suppose that*

$$m = \Omega\left(\frac{L^6 |\mathcal{X}|^{16} \log(|\mathcal{X}|^2 L / \delta)}{\epsilon^8}\right),$$

then with probability at least $1 - \delta$, for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we have

$$|\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| \leq \epsilon^2 / (2|\mathcal{X}|^4). \quad (33)$$

Lemma 3 (Lemma 4.1, [8]). *There exist positive constants $\bar{C}_1, \bar{C}_2, C_f$ such that for any $\delta > 0$, if τ satisfies that*

$$\bar{C}_1 m^{-3/2} L^{-3/2} [\log(|\mathcal{X}| L^2 / \delta)]^{3/2} \leq \tau \leq \bar{C}_2 L^{-6} [\log m]^{-3/2},$$

then with probability at least $1 - \delta$, for all $\tilde{\boldsymbol{\theta}}, \hat{\boldsymbol{\theta}}$ satisfying $\|\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \tau, \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \tau$ and $\mathbf{x} \in \mathcal{X}$ we have

$$\left| f(\mathbf{x}; \tilde{\boldsymbol{\theta}}) - f(\mathbf{x}; \hat{\boldsymbol{\theta}}) - \langle \mathbf{g}(\mathbf{x}; \hat{\boldsymbol{\theta}}), \tilde{\boldsymbol{\theta}} - \hat{\boldsymbol{\theta}} \rangle \right| \leq C_f \tau^{4/3} L^3 \sqrt{m \log m}.$$

Lemma 4 (Lemma B.3, [8]). *There exist positive constants $\bar{C}_1, \bar{C}_2, C_{g,1}, C_{g,2}$ such that for any $\delta > 0$, if τ satisfies that*

$$\bar{C}_1 m^{-3/2} L^{-3/2} [\log(|\mathcal{X}| L^2 / \delta)]^{3/2} \leq \tau \leq \bar{C}_2 L^{-6} [\log m]^{-3/2},$$

then with probability at least $1 - \delta$, for any $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \tau$ and $\mathbf{x} \in \mathcal{X}$ we have

$$\|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta})\|_2 \leq C_{g,1} \sqrt{mL}, \quad \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}) - \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)\|_2 \leq C_{g,2} \tau^{1/3} \sqrt{m \log m} L^{7/2} / 2.$$

Lemma 5 (Lemma 5.1, [47]). *There exists a positive constant \bar{C} such that for any $\delta \in (0, 1)$, if $m \geq \bar{C} |\mathcal{X}|^4 L^6 \log(|\mathcal{X}|^2 L / \delta) / \lambda_0^4$, then with probability at least $1 - \delta$, there exists a $\boldsymbol{\theta}_* \in \mathbb{R}^p$ such that for any $\mathbf{x} \in \mathcal{X}$,*

$$h(\mathbf{x}) = \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_* - \boldsymbol{\theta}_0 \rangle, \quad \sqrt{m} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}_0\|_2 \leq \sqrt{2} S. \quad (34)$$

Lemma 6 (Lemma B.2, [47]). *There exist constants $\{\bar{C}_i\}_{i=1}^5 > 0$ such that for any $\delta > 0$, if for all N_k, η, m satisfy*

$$\begin{aligned} 2\sqrt{N_k / (m\alpha)} &\geq \bar{C}_1 m^{-3/2} L^{-3/2} [\log(|\mathcal{X}| L^2 / \delta)]^{3/2}, \\ 2\sqrt{N_k / (m\alpha)} &\leq \bar{C}_2 \min \{ L^{-6} [\log m]^{-3/2}, (m(\alpha\eta)^2 L^{-6} N_k^{-1} (\log m)^{-1})^{3/8} \}, \\ \eta &\leq \bar{C}_3 (m\alpha + N_k mL)^{-1}, \\ m^{1/6} &\geq \bar{C}_4 \sqrt{\log m} L^{7/2} N_k^{7/6} \alpha^{-7/6} (1 + \sqrt{N_k / \alpha}), \end{aligned}$$

then with probability at least $1 - \delta$, we have that $\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_0\|_2 \leq 2\sqrt{N_k / (m\alpha)}$ and

$$\begin{aligned} &\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_0 - \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_k / \sqrt{m}\|_2 \\ &\leq (1 - \eta m \alpha)^{J/2} \sqrt{N_k / (m\alpha)} + \bar{C}_5 m^{-2/3} \sqrt{\log m} L^{7/2} N_k^{5/3} \alpha^{-5/3} (1 + \sqrt{N_k / \alpha}). \end{aligned}$$

Lemma 7 (Lemma B.3, [47]). *There exist constants $\{\bar{C}_i\}_{i=1}^5 > 0$ such that for any $\delta > 0$, if m satisfies that*

$$\bar{C}_1 m^{-3/2} L^{-3/2} [\log(|\mathcal{X}|L^2/\delta)]^{3/2} \leq 2\sqrt{N_k/(m\lambda)} \leq \bar{C}_2 L^{-6} [\log m]^{-3/2},$$

then with probability at least $1 - \delta$, we have

$$\|\bar{\mathbf{Z}}_k - \mathbf{Z}_k\|_F \leq \bar{C}_3 m^{-1/6} \sqrt{\log m} L^4 N_k^{7/6} \alpha^{-1/6},$$

Based on previous lemmas, we can define the random event $\mathcal{E}_k^{\text{neural}}$ as follows:

$$\begin{aligned} \mathcal{E}_k^{\text{neural}} := & \left\{ \exists \text{ positive constants } \{C_i\}_{i=1}^5, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \right. \\ & |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| \leq \epsilon^2 / (2|\mathcal{X}|^4), \\ & \left| f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_k - \boldsymbol{\theta}_0 \rangle \right| \leq C_1 N_k^{2/3} m^{-1/6} \alpha^{-2/3} L^3 \sqrt{\log m}, \\ & \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)\|_2, \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})\|_2 \leq C_2 \sqrt{mL}, \\ & \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}) - \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)\|_2 \leq C_3 N_{k-1}^{1/6} \alpha^{-1/6} m^{1/6} \sqrt{\log m} L^{7/2}, \\ & \exists \boldsymbol{\theta}_* \in \mathbb{R}^p, h(\mathbf{x}) = \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_* - \boldsymbol{\theta}_0 \rangle, \sqrt{m} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}_0\|_2 \leq \sqrt{2}S, \\ & \|\boldsymbol{\theta}_k - \boldsymbol{\theta}_0 - \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_k / \sqrt{m}\|_2 \\ & \leq (1 - \eta m \alpha)^{J/2} \sqrt{N_k / (m\alpha)} + C_4 m^{-2/3} \sqrt{\log m} L^{7/2} N_k^{5/3} \alpha^{-5/3} (1 + \sqrt{N_k / \alpha}), \\ & \left. \|\bar{\mathbf{Z}}_k - \mathbf{Z}_k\|_F \leq C_5 m^{-1/6} \sqrt{\log m} L^4 N_k^{7/6} \alpha^{-1/6}. \right\} \end{aligned} \quad (35)$$

It can be verified that when $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, $\mathbb{P}(\mathcal{E}_k^{\text{neural}}) \geq 1 - \delta_k$.

Suppose event $\mathcal{E}_k^{\text{neural}}$ holds, then we have the following lemmas.

Lemma 8. *Let $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, then under event $\mathcal{E}_k^{\text{neural}}$, we have*

$$\begin{aligned} & |f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k \rangle| \\ & \leq C[(1 - \eta m \alpha)^{J/2} \sqrt{N_k L / \alpha} + 4N_k^3 \bar{\epsilon}^2 L^2 / \alpha^3 + N_k^2 L \bar{\epsilon}^2 / \alpha^2 + \bar{\epsilon}^2 N_k / \alpha]. \end{aligned} \quad (36)$$

Proof. See Appendix F.3.1. □

Lemma 9. *Let $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, then under event $\mathcal{E}_k^{\text{neural}}$, there exists an $\hat{\boldsymbol{\theta}}_k^* \in \mathbb{R}^{d_k}$ such that $|h(\mathbf{x}) - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k^* \rangle| \leq 2S\bar{\epsilon}$, $\|\hat{\boldsymbol{\theta}}_k^*\|_2 \leq \sqrt{2}S$.*

Proof. See Appendix F.3.2. □

We use $\hat{\boldsymbol{\theta}}_k^*$ to denote the underlying vector, and we avoid using double subscripts for the simplicity of notions.

Lemma 10. *Let $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, then under event $\mathcal{E}_k^{\text{neural}}$, we have $d_k \leq d_{\text{eff}}(\bar{\epsilon}^2 / |\mathcal{X}|)$.*

Proof. See Appendix F.3.3. □

Now we begin to prove Theorem 7.

Proof of Theorem 7. We follow the main steps of the proof of Theorem 4.

Step 1.1: By Lemma 9, we know that there exists $\widehat{\boldsymbol{\theta}}_k^*$ such that for any $\mathbf{x} \in \mathcal{X}$, $h(\mathbf{x}) = \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}), \widehat{\boldsymbol{\theta}}_k^* \rangle + \eta(\mathbf{x})$, where $|\eta(\mathbf{x})| \leq 2S\bar{\epsilon}$. Therefore, for any $\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k))$, we have

$$\begin{aligned}
\left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \widehat{\boldsymbol{\theta}}_k^* \rangle \right| &= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) y_j - \mathbf{y}^\top \widehat{\boldsymbol{\theta}}_k^* \right| \\
&= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) \left(\boldsymbol{\psi}_{d_k}(\mathbf{x}_j)^\top \widehat{\boldsymbol{\theta}}_k^* + \eta(\mathbf{x}_j) + \xi_j \right) - \mathbf{y}^\top \widehat{\boldsymbol{\theta}}_k^* \right| \\
&= \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) (\eta(\mathbf{x}_j) + \xi_j) - \alpha \mathbf{y}^\top \mathbf{A}_k^{-1} \widehat{\boldsymbol{\theta}}_k^* \right| \\
&\leq \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) \eta(\mathbf{x}_j) \right| + \left| \sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) \xi_j \right| + \alpha \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \widehat{\boldsymbol{\theta}}_k^* \right|.
\end{aligned} \tag{37}$$

Following the proof of Theorem 4, the first and second terms in (37) can be bounded as follows, with probability at least $1 - \delta_k$ for any $\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k))$:

$$\begin{aligned}
\left| \mathbf{y}^\top \mathbf{A}_k^{-1} \sum_{j=1}^{N_k} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) \eta(\mathbf{x}_j) \right| &\leq 2S\bar{\epsilon} \sqrt{(1 + \zeta) \tau(\mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k)))} \\
\left| \sum_{j=1}^{N_k} \mathbf{y}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}_{d_k}(\mathbf{x}_j) \xi_j \right| &\leq \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)}
\end{aligned} \tag{38}$$

The last term in (37) can be bounded as

$$\alpha \left| \mathbf{y}^\top \mathbf{A}_k^{-1} \widehat{\boldsymbol{\theta}}_k^* \right| \leq \sqrt{\alpha} \|\widehat{\boldsymbol{\theta}}_k^*\|_2 \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \leq \sqrt{2\alpha} S \|\mathbf{y}\|_{\mathbf{A}_k^{-1}}.$$

Therefore, we have

$$\begin{aligned}
&\left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \widehat{\boldsymbol{\theta}}_k^* \rangle \right| \\
&\leq 2S\bar{\epsilon} \sqrt{(1 + \zeta) \tau(\mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k)))} + \|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \left(\sqrt{2 \log(|\widehat{\mathcal{S}}_k|^2 / \delta_k)} + \sqrt{2\alpha} S \right) \\
&:= \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\mathbf{x}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}')).
\end{aligned} \tag{39}$$

Similar to the proof of Theorem 4, since by the property of rounding procedure, we have

$$\|\mathbf{y}\|_{\mathbf{A}_k^{-1}} \leq \max_{\mathbf{y}' \in \mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k))} \|\mathbf{y}'\|_{\mathbf{A}_k^{-1}} \leq \sqrt{(1 + \zeta)/N_k} \|\mathbf{y}'\|_{\mathbf{A}_{\boldsymbol{\psi}_{d_k}(\lambda_k)}^{-1}} \leq 2\sqrt{(1 + \zeta)d_k/N_k}. \tag{40}$$

Meanwhile, according to Kiefer–Wolfowitz Theorem, we have

$$\tau(\mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k))) = \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}_{d_k}(\widehat{\mathcal{S}}_k))} \|\mathbf{y}\|_{\mathbf{A}_{\boldsymbol{\psi}_{d_k}(\lambda)}^{-1}}^2 \leq \inf_{\lambda \in \Lambda_{\mathcal{X}}} \sup_{\mathbf{y} \in \mathcal{Y}(\boldsymbol{\psi}_{d_k}(\mathcal{X}))} \|\mathbf{y}\|_{\mathbf{A}_{\boldsymbol{\psi}_{d_k}(\lambda)}^{-1}}^2 \leq 4d_k. \tag{41}$$

Submitting (40) and (41) into (39), we have

$$\left| \langle \mathbf{y}, \widehat{\boldsymbol{\theta}}_k - \widehat{\boldsymbol{\theta}}_k^* \rangle \right| \leq 4S\bar{\epsilon} \sqrt{(1 + \zeta)d_k} + 2\sqrt{(1 + \zeta)d_k/N_k} \left(\sqrt{2 \log(|\mathcal{X}|^2 / \delta_k)} + \sqrt{2\alpha} S \right), \tag{42}$$

Step 1.2: We show that $\mathbf{x}_* \in \widehat{\mathcal{S}}_{k+1}$ for any $0 \leq k \leq n - 1$ first. For any $\widehat{\mathbf{x}} \in \widehat{\mathcal{S}}_k$, we have

$$\begin{aligned}
&f(\widehat{\mathbf{x}}) - f(\mathbf{x}_*) \\
&\leq \langle \boldsymbol{\psi}_{d_k}(\widehat{\mathbf{x}}), \widehat{\boldsymbol{\theta}}_k \rangle - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}_*), \widehat{\boldsymbol{\theta}}_k \rangle + \underbrace{|f(\widehat{\mathbf{x}}) - \langle \boldsymbol{\psi}_{d_k}(\widehat{\mathbf{x}}), \widehat{\boldsymbol{\theta}}_k \rangle| + |f(\mathbf{x}_*) - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}_*), \widehat{\boldsymbol{\theta}}_k \rangle|}_{I_1} \\
&\leq \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}_*)) + h(\widehat{\mathbf{x}}) - \eta(\widehat{\mathbf{x}}) - h(\mathbf{x}_*) + \eta(\mathbf{x}_*) + I_1 \\
&\leq 4S\bar{\epsilon} + \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\widehat{\mathbf{x}}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}_*)) + I_1,
\end{aligned} \tag{43}$$

To further bound (43), we have

$$\begin{aligned}
& 4S\bar{\epsilon} + \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\hat{\mathbf{x}}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}_*)) + I_1 \\
& \leq 4S\bar{\epsilon} \\
& \quad + 4S\bar{\epsilon}\sqrt{(1+\zeta)d} + 2\sqrt{(1+\zeta)d/N_k}\left(\sqrt{2\log(|\mathcal{X}|^2/\delta_k)} + \sqrt{2\alpha S}\right) \\
& \quad + 2C[(1-\eta m\alpha)^{J/2}\sqrt{N_k L/\alpha} + 4N_k^3\bar{\epsilon}^2 L^2/\alpha^3 + N_k^2 L\bar{\epsilon}^2/\alpha^2 + \bar{\epsilon}^2 N_k/\alpha] \\
& \leq 8S\bar{\epsilon}\sqrt{(1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)} + 4\sqrt{(1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)/N_k}\sqrt{\log(|\mathcal{X}|^2/\delta_k)} \\
& \quad + 2C[(1-\eta m\alpha)^{J/2}\sqrt{N_k L/\alpha} + 6N_k^3\bar{\epsilon}^2 L^2/\alpha^3] \tag{44}
\end{aligned}$$

where the first inequality holds due to Lemma 8 and (42), the second inequality holds since $\alpha \leq \min\{1, \log(|\mathcal{X}|^2/S)\}$. Now, according to the choice of N_k , we have

$$4\sqrt{(1+\zeta)\log(|\mathcal{X}|^2/\delta_k)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)/N_k} \leq 2^{-k}/8, \tag{45}$$

where the inequality holds due to the selection of N_k . According to the choice of $\bar{\epsilon}$, we have

$$8S\bar{\epsilon}\sqrt{(1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)} \leq \epsilon/8,$$

According to the choice of $N_k, \bar{\epsilon}$, we have

$$\begin{aligned}
& 12CN_k^3\bar{\epsilon}^2 L^2/\alpha^3 \\
& \leq 12C \cdot 1024^3(1+\zeta)^3 \log^3(|\mathcal{X}|^2/\delta_k)d^3(\bar{\epsilon}^2/|\mathcal{X}|)\epsilon^{-6}\bar{\epsilon}^2 L^2/\alpha^3 + 12Cr_d(\zeta)^3\bar{\epsilon}^2 L^2/\alpha^3 \\
& \leq \epsilon/8, \tag{46}
\end{aligned}$$

According to the choice of J , we have

$$2C(1-\eta m\alpha)^{J/2}\sqrt{N_k L/\alpha} \leq \epsilon/8. \tag{47}$$

Therefore, substituting these bounds into (44), we have

$$4S\bar{\epsilon} + \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\hat{\mathbf{x}}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}_*)) + I_1 \leq 2^{-k}/8 + 3\epsilon/8, \tag{48}$$

Therefore, $\mathbf{x}_* \in \hat{\mathcal{S}}_{k+1}$.

Step 1.3: Next, we show that for any $k = n$, any $\mathbf{x} \in \hat{\mathcal{S}}_k$ satisfying $\Delta(\mathbf{x}) \geq \epsilon$ will be eliminated. In other words, all the arms \mathbf{x}' remaining in $\hat{\mathcal{S}}_{k+1}$ satisfy $\Delta(\mathbf{x}') \leq \epsilon$. Suppose $\Delta(\mathbf{x}) \geq \epsilon$ and $\mathbf{x} \in \hat{\mathcal{S}}_k$, then we have $\Delta(\mathbf{x}) = \epsilon \geq 4 \cdot 2^{-(k+1)} = 2 \cdot 2^{-k}$ by the selection of k . Now we have

$$\begin{aligned}
f(\mathbf{x}_*) - f(\mathbf{x}) & \geq \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}_*), \hat{\boldsymbol{\theta}}_k \rangle - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k \rangle - \underbrace{|f(\mathbf{x}) - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k \rangle| + |f(\mathbf{x}_*) - \langle \boldsymbol{\psi}_{d_k}(\mathbf{x}_*), \hat{\boldsymbol{\theta}}_k \rangle|}_{I_2} \\
& \geq h(\mathbf{x}_*) - h(\mathbf{x}) - \eta(\mathbf{x}_*) + \eta(\mathbf{x}) - \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\mathbf{x}_*) - \boldsymbol{\psi}_{d_k}(\mathbf{x})) - I_2 \\
& \geq \Delta(\mathbf{x}) - (4S\bar{\epsilon} + \tilde{\omega}_k(\boldsymbol{\psi}_{d_k}(\mathbf{x}) - \boldsymbol{\psi}_{d_k}(\mathbf{x}_*)) + I_2), \tag{49}
\end{aligned}$$

Similar to (48), (49) can be bounded by $2^{-k}/8 + 3\epsilon/8$. Therefore, when $k > n = \log(1/\epsilon)$, we have $f(\mathbf{x}_*) - f(\mathbf{x}) \geq \epsilon - \epsilon/2 = \epsilon/2 > 2^{-k}/8 + 3\epsilon/8$, which suggests that \mathbf{x} will be eliminated.

Step 1.4: The probability for all the events hold simultaneously in round $k = 1, \dots, n$ can be bounded by a union bound, the same as (16).

Step 1.5: Finally we bound the total sample complexity, which is

$$\begin{aligned}
N & = \sum_{k=1}^n N_k \\
& = \tilde{O}\left((1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|) \sum_{k=1}^n (4^k + r_{d_k}(\zeta))\right) \\
& = \tilde{O}\left((1+\zeta)d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)/\epsilon^2 + r_{d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)}(\zeta)\right).
\end{aligned}$$

□

F.3 Proof of Lemmas

In this section we propose the proofs of lemmas in Appendix F.2. We first propose the following proposition, which can be directly derived by the construction rule of ψ_{d_k} introduced in Algorithm 2.

Proposition 3. *For any $\mathbf{x} \in \mathcal{X}$, we have*

$$\psi_{d_k}(\mathbf{x}) = \sum_{j=1}^d a_{\mathbf{x},j} e_j \mathbf{e}_j, \quad \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}) / \sqrt{m} = \sum_{j=1}^{|\mathcal{X}|} a_{\mathbf{x},j} e_j \mathbf{b}_j = \mathbf{B}_{\mathcal{X}}^\top [\boldsymbol{\psi}(\mathbf{x}); \boldsymbol{\epsilon}_{\mathbf{x}}], \quad \|\boldsymbol{\epsilon}_{\mathbf{x}}\|_2 \leq \bar{\epsilon},$$

where $\mathbf{b}_j \in \mathbb{R}^p$ are orthogonal unit vectors in p -dimensional space, $\mathbf{B}_{\mathcal{X}} = [\mathbf{b}_1^\top; \dots, \mathbf{b}_{|\mathcal{X}|}^\top] \in \mathbb{R}^{|\mathcal{X}| \times p}$.

In the following proofs, for simplicity, let $\tilde{\psi}(\cdot)$ denote $\psi_{d_k}(\cdot)$.

F.3.1 Proof of Lemma 8

In detail, to prove Lemma 8, we need the following lemmas.

Lemma 11. *Let $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, then under event $\mathcal{E}_k^{\text{neural}}$, we have*

$$|f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_t / \sqrt{m} \rangle| \leq C(1 - \eta m \lambda)^{J/2} \sqrt{N_k L / \alpha}. \quad (50)$$

Proof of Lemma 11. By Lemma 3, we have for all $\mathbf{x} \in \mathcal{X}$,

$$\begin{aligned} \left| f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_k - \boldsymbol{\theta}_0 \rangle \right| &= \left| f(\mathbf{x}; \boldsymbol{\theta}_k) - f(\mathbf{x}; \boldsymbol{\theta}_0) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_k - \boldsymbol{\theta}_0 \rangle \right| \\ &\leq 3C_f N_k^{2/3} m^{-1/6} \alpha^{-2/3} L^3 \sqrt{\log m}, \end{aligned}$$

where we use the fact that $f(\mathbf{x}_j; \boldsymbol{\theta}_0) = 0$ due to the initialization scheme adapted by Algorithm 2, and $\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_0\|_2 \leq 2\sqrt{N_k/(m\alpha)}$ from Lemma 6. Therefore, we have

$$\begin{aligned} &|f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_t / \sqrt{m} \rangle| \\ &\leq |f(\mathbf{x}; \boldsymbol{\theta}_k) - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_k - \boldsymbol{\theta}_0 \rangle| + \|\boldsymbol{\theta}_k - \boldsymbol{\theta}_0 - \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_t / \sqrt{m}\|_2 \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)\|_2 \\ &\leq 3C_f N_k^{2/3} m^{-1/6} \alpha^{-2/3} L^3 \sqrt{\log m} \\ &\quad + [(1 - \eta m \alpha)^{J/2} \sqrt{N_k/(m\alpha)} + \bar{C}_5 m^{-2/3} \sqrt{\log m} L^{7/2} N_k^{5/3} \alpha^{-5/3} (1 + \sqrt{N_k/\alpha})] C_{g,1} \sqrt{mL} \\ &\leq C(1 - \eta m \alpha)^{J/2} \sqrt{N_k L / \alpha}, \end{aligned}$$

where the first inequality holds due to triangle inequality, the second one holds due to Lemma 4 and Lemma 6, the last one holds since $m = \text{poly}(L, N_{k-1}, \alpha, \log(1/\delta))$. \square

We also need the following lemma.

Lemma 12. *Suppose $\alpha \geq 2C_{g,1} N_k \bar{\epsilon} \sqrt{L}$ and $m = \text{poly}(L, N_{k-1}, \alpha, \log(1/\delta))$, then with probability at least $1 - \delta$, for all $\mathbf{x} \in \mathcal{X}$, we have*

$$\langle \boldsymbol{\psi}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k \rangle - \langle \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_t / \sqrt{m}, \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0) \rangle \leq C_\psi (N_k L \alpha / \lambda_k^2 + \bar{\epsilon}^2 / \lambda_k^2 N_k^2 L + \bar{\epsilon}^2 N_k / \alpha). \quad (51)$$

Proof of Lemma 12. For each $\mathbf{x} \in \mathcal{X}$, we have

$$\begin{aligned} &\langle \boldsymbol{\psi}(\mathbf{x}), \hat{\boldsymbol{\theta}}_k \rangle - \langle \bar{\mathbf{Z}}_k^{-1} \bar{\mathbf{b}}_t / \sqrt{m}, \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0) \rangle \\ &= \left| \left\langle \boldsymbol{\psi}(\mathbf{x}), \mathbf{A}_k^{-1} \sum_{i=1}^{N_k} \boldsymbol{\psi}(\mathbf{x}_i) y_i \right\rangle - \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \sum_{i=1}^{N_k} \sum_{i=1}^{N_k} y_i \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle \right| \\ &\leq \sum_{i=1}^{N_k} \left| \left\langle \boldsymbol{\psi}(\mathbf{x}), \mathbf{A}_k^{-1} \boldsymbol{\psi}(\mathbf{x}_i) \right\rangle - \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle \right|, \end{aligned} \quad (52)$$

where the last line holds due to triangle inequality. Then, to bound (52), we first decompose it to two terms by triangle inequality:

$$\begin{aligned}
& \left| \left\langle \boldsymbol{\psi}(\boldsymbol{x}), \mathbf{A}_k^{-1} \boldsymbol{\psi}(\boldsymbol{x}_i) \right\rangle - \left\langle \frac{\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_0)}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \boldsymbol{g}(\boldsymbol{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle \right| \\
& \leq \underbrace{\left| \left\langle \boldsymbol{\psi}(\boldsymbol{x}), \mathbf{A}_k^{-1} \boldsymbol{\psi}(\boldsymbol{x}_i) \right\rangle - \left\langle \frac{\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \boldsymbol{g}(\boldsymbol{x}_i; \boldsymbol{\theta}_{k-1}) / \sqrt{m} \right\rangle \right|}_{I_1} \\
& \quad + \underbrace{\left| \left\langle \frac{\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_0)}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \boldsymbol{g}(\boldsymbol{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle - \left\langle \frac{\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \boldsymbol{g}(\boldsymbol{x}_i; \boldsymbol{\theta}_{k-1}) / \sqrt{m} \right\rangle \right|}_{I_2}. \quad (53)
\end{aligned}$$

To bound I_1 : Let \boldsymbol{x}' denote \boldsymbol{x}_i . First we write $\langle \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1}), \mathbf{Z}_k^{-1} \boldsymbol{g}(\boldsymbol{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m$ as products of block matrices, as follows:

$$\begin{aligned}
& \left\langle \frac{\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \frac{\boldsymbol{g}(\boldsymbol{x}'; \boldsymbol{\theta}_{k-1})}{\sqrt{m}} \right\rangle \\
& = \left\langle \mathbf{B}_{\mathcal{X}}^\top [\boldsymbol{\psi}(\boldsymbol{x}); \boldsymbol{\epsilon}_{\boldsymbol{x}}], \left(\alpha \mathbf{I} + \mathbf{B}_{\mathcal{X}}^\top \sum_{j=1}^{N_k} [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}] [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}]^\top \mathbf{B}_{\mathcal{X}} \right)^{-1} \mathbf{B}_{\mathcal{X}}^\top [\boldsymbol{\psi}(\boldsymbol{x}'); \boldsymbol{\epsilon}_{\boldsymbol{x}'}] \right\rangle \\
& = \left\langle [\boldsymbol{\psi}(\boldsymbol{x}); \boldsymbol{\epsilon}_{\boldsymbol{x}}], \left(\alpha \mathbf{I} + \sum_{j=1}^{N_k} [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}] [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}]^\top \right)^{-1} [\boldsymbol{\psi}(\boldsymbol{x}'); \boldsymbol{\epsilon}_{\boldsymbol{x}'}] \right\rangle, \quad (54)
\end{aligned}$$

where the first equality holds due to Proposition 3. To further bound (54), we give the matrix inverse an explicit expression. Denote $\mathbf{M}, \mathbf{N}, \mathbf{P}, \mathbf{Q}$ as follows:

$$\begin{aligned}
\mathbf{M} &= \sum_{j=1}^{N_k} \boldsymbol{\psi}(\boldsymbol{x}_j) \boldsymbol{\epsilon}_{\boldsymbol{x}_j}^\top, \quad \mathbf{N} = \sum_{j=1}^{N_k} \boldsymbol{\epsilon}_{\boldsymbol{x}_j} \boldsymbol{\epsilon}_{\boldsymbol{x}_j}^\top, \\
\mathbf{P} &= \mathbf{A}_k - \mathbf{M}(\mathbf{N} + \alpha \mathbf{I})^{-1} \mathbf{M}^\top, \quad \mathbf{Q} = (\mathbf{N} + \alpha \mathbf{I})^{-1} + (\mathbf{N} + \alpha \mathbf{I})^{-1} \mathbf{M}^\top \mathbf{P}^{-1} \mathbf{M} (\mathbf{N} + \alpha \mathbf{I})^{-1},
\end{aligned}$$

then we have

$$\begin{aligned}
\left(\alpha \mathbf{I} + \sum_{j=1}^{N_k} [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}] [\boldsymbol{\psi}(\boldsymbol{x}_j); \boldsymbol{\epsilon}_{\boldsymbol{x}_j}]^\top \right)^{-1} &= \begin{pmatrix} \mathbf{A}_k & \mathbf{M} \\ \mathbf{M}^\top & \mathbf{N} + \alpha \mathbf{I} \end{pmatrix}^{-1} \\
&= \begin{pmatrix} \mathbf{P}^{-1} & -\mathbf{A}_k^{-1} \mathbf{M} \mathbf{P}^{-1} \\ \mathbf{P}^{-1} \mathbf{M}^\top \mathbf{A}_k^{-1} & \mathbf{Q} \end{pmatrix}, \quad (55)
\end{aligned}$$

where the second equality holds due to the block matrix inverse formula. Then substituting (55) into (54) and considering the difference between $\langle \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1}), \mathbf{Z}_k^{-1} \boldsymbol{g}(\boldsymbol{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m$ and $\langle \boldsymbol{\psi}(\boldsymbol{x}), \mathbf{A}_k^{-1} \boldsymbol{\psi}(\boldsymbol{x}') \rangle$, we have

$$\begin{aligned}
& \left| \left\langle \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1}), \mathbf{Z}_k^{-1} \boldsymbol{g}(\boldsymbol{x}'; \boldsymbol{\theta}_{k-1}) \right\rangle / m - \left\langle \boldsymbol{\psi}(\boldsymbol{x}), \mathbf{A}_k^{-1} \boldsymbol{\psi}(\boldsymbol{x}') \right\rangle \right| \\
& = \left| \boldsymbol{\psi}(\boldsymbol{x})^\top (\mathbf{P}^{-1} - \mathbf{A}_k^{-1}) \boldsymbol{\psi}(\boldsymbol{x}') + \boldsymbol{\epsilon}_{\boldsymbol{x}}^\top \mathbf{P}^{-1} \mathbf{M}^\top \mathbf{A}_k^{-1} \boldsymbol{\psi}(\boldsymbol{x}') \right. \\
& \quad \left. + \boldsymbol{\psi}(\boldsymbol{x}) \mathbf{A}_k^{-1} \mathbf{M} \mathbf{P}^{-1} \boldsymbol{\epsilon}_{\boldsymbol{x}'} + \boldsymbol{\epsilon}_{\boldsymbol{x}}^\top \mathbf{Q} \boldsymbol{\epsilon}_{\boldsymbol{x}'} \right| \\
& \leq \|\boldsymbol{\psi}(\boldsymbol{x}')\|_2 \|\boldsymbol{\psi}(\boldsymbol{x})\|_2 \|\mathbf{P}^{-1} - \mathbf{A}_k^{-1}\|_2 + \|\mathbf{Q}\|_2 \|\boldsymbol{\epsilon}_{\boldsymbol{x}}\|_2 \|\boldsymbol{\epsilon}_{\boldsymbol{x}'}\|_2 \\
& \quad + \|\mathbf{P}^{-1}\|_2 \|\mathbf{M}\|_2 \|\mathbf{A}_k^{-1}\|_2 (\|\boldsymbol{\epsilon}_{\boldsymbol{x}}\|_2 \|\boldsymbol{\psi}(\boldsymbol{x}')\|_2 + \|\boldsymbol{\epsilon}_{\boldsymbol{x}'}\|_2 \|\boldsymbol{\psi}(\boldsymbol{x})\|_2), \quad (56)
\end{aligned}$$

where we use triangle inequality in the inequality. To bound (56), we have the following inequalities:

$$\forall \boldsymbol{x} \in \mathcal{X}, \quad \|\boldsymbol{\psi}(\boldsymbol{x})\|_2 \leq \|\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1})\|_2 / \sqrt{m} \leq C_{g,1} \sqrt{L}, \quad (57)$$

where the first inequality holds since $\boldsymbol{\psi}(\boldsymbol{x})$ is the truncation of $\boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}_{k-1}) / \sqrt{m}$, the second one holds due to Lemma 4. Use the fact that $\alpha \geq 2C_{g,1} N_k \bar{\epsilon} \sqrt{L}$, we also have

$$\mathbf{A}_k \succeq \mathbf{P} \succeq \mathbf{A}_k - \|\mathbf{M}\|_2^2 / \alpha \succeq \mathbf{A}_k - C_{g,1}^2 N_k^2 \bar{\epsilon}^2 L / \alpha \succeq \mathbf{A}_k - \alpha / 2 \cdot \mathbf{I}, \quad (58)$$

and

$$\|Q\|_2 \leq \|(N + \alpha I)^{-1}\|_2 + \|(N + \alpha I)^{-1}\|_2^2 \|M\|_2^2 \|P^{-1}\|_2 \leq \alpha^{-1} + \alpha^{-3} N_k^2 C_{g,1}^2 L \bar{\epsilon}^2 \leq 2/\alpha. \quad (59)$$

Lastly, we have

$$\|A_k^{-1} - P^{-1}\|_2 = \|P^{-1}(P - A_k)A_k^{-1}\|_2 \leq \|A_k^{-1}\|_2 \|P^{-1}\|_2 \|A_k - P\|_2 \leq C_{g,1}^2 N_k^2 \bar{\epsilon}^2 L / \alpha^3, \quad (60)$$

where the bounds of $\|A_k^{-1}\|_2$, $\|P^{-1}\|_2$ and $\|A_k - P\|_2$ come from (58) and (59). Finally, substituting (57), (58), (59), (60) into (56), we can bound I_1 as

$$I_1 \leq C_{g,1}^4 N_k^2 \bar{\epsilon}^2 L^2 / \alpha^3 + 4\bar{\epsilon}^2 / \alpha^2 N_k L + 2\bar{\epsilon}^2 / \alpha. \quad (61)$$

To bound I_2 : To bound I_2 , we have

$$\begin{aligned} I_2 &\leq \left| \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0)}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle - \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle \right| \\ &\quad + \left| \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \bar{\mathbf{Z}}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle - \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle \right| \\ &\quad + \left| \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0) / \sqrt{m} \right\rangle - \left\langle \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})}{\sqrt{m}}, \mathbf{Z}_k^{-1} \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_{k-1}) / \sqrt{m} \right\rangle \right| \\ &\leq \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0) - \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})\|_2 / \sqrt{m} \|\bar{\mathbf{Z}}_k^{-1}\|_2 \|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0)\|_2 / \sqrt{m} \\ &\quad + \|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_{k-1})\|_2 / \sqrt{m} \|\bar{\mathbf{Z}}_k^{-1} - \mathbf{Z}_k^{-1}\|_2 \|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_0)\|_2 / \sqrt{m} \\ &\quad + \|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}_{k-1})\|_2 / \sqrt{m} \|\mathbf{Z}_k^{-1}\|_2 \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0) - \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})\|_2 / \sqrt{m} \\ &\leq 2C_{g,1} C_{g,2} N_k^{1/6} \alpha^{-7/6} m^{-1/3} \sqrt{\log m} L^4 + C_{g,1}^2 L^5 \cdot \bar{C}_4 m^{-1/6} \sqrt{\log m} N_k^{7/6} \alpha^{-13/6}, \quad (62) \end{aligned}$$

where the first inequality holds due to triangle inequality, the third one holds due to Lemma 4 and the fact

$$\|\bar{\mathbf{Z}}_k^{-1} - \mathbf{Z}_k^{-1}\|_2 \leq \|\bar{\mathbf{Z}}_k^{-1}\|_2 \|\mathbf{Z}_k^{-1}\|_2 \|\bar{\mathbf{Z}}_k - \mathbf{Z}_k\|_2 \leq \bar{C}_4 L^4 m^{-1/6} \sqrt{\log m} N_k^{7/6} \alpha^{-13/6}, \quad (63)$$

where the last inequality holds due to the fact $\bar{\mathbf{Z}}_k, \mathbf{Z}_k \succeq \alpha I$ and Lemma 7.

Final bound: Substituting (61) and (62) into (52) ends our proof. \square

Then we begin our proof, which is a direct combination of Lemma 11 and Lemma 12.

Proof of Lemma 8. Adding (50) with (51) finishes our proof. \square

F.3.2 Proof of Lemma 9

Proof of Lemma 9. First, by Lemma 5, we know that there exists $\boldsymbol{\theta}_* \in \mathbb{R}^p$ such that

$$h(\mathbf{x}) = \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_* - \boldsymbol{\theta}_0 \rangle, \sqrt{m} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}_0\|_2 \leq \sqrt{2}S. \quad (64)$$

By Proposition 3, we have

$$\boldsymbol{\psi}(\mathbf{x}) = \sum_{j=1}^d u_{\mathbf{x},j} e_j e_j, \quad \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}) / \sqrt{m} = \sum_{j=1}^{|\mathcal{X}|} u_{\mathbf{x},j} e_j \mathbf{b}_j = \mathbf{B}_{\mathcal{X}}^{\top} [\boldsymbol{\psi}(\mathbf{x}); \boldsymbol{\epsilon}_{\mathbf{x}}], \quad \|\boldsymbol{\epsilon}_{\mathbf{x}}\|_2 \leq \bar{\epsilon}. \quad (65)$$

Therefore, suppose $\mathbf{B}_{\mathcal{X}} \sqrt{m} (\boldsymbol{\theta}_* - \boldsymbol{\theta}_0) = [\mathbf{r}; \mathbf{t}]$, $\mathbf{r} \in \mathbb{R}^d$, $\mathbf{t} \in \mathbb{R}^{|\mathcal{X}|-d}$, set $\hat{\boldsymbol{\theta}}_* = \mathbf{r}$, we have $\|\hat{\boldsymbol{\theta}}_*\|_2 \leq \sqrt{m} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}_0\|_2 \leq \sqrt{2}S$, and

$$\begin{aligned} |\langle \boldsymbol{\psi}(\mathbf{x}), \hat{\boldsymbol{\theta}}_* \rangle - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \boldsymbol{\theta}_* - \boldsymbol{\theta}_0 \rangle| &= |\boldsymbol{\psi}(\mathbf{x})^{\top} \mathbf{r} - [\boldsymbol{\psi}(\mathbf{x})^{\top}, \boldsymbol{\epsilon}_{\mathbf{x}}^{\top}] \mathbf{B}_{\mathcal{X}} \sqrt{m} (\boldsymbol{\theta}_* - \boldsymbol{\theta}_0)| \\ &= |\boldsymbol{\psi}(\mathbf{x})^{\top} \mathbf{r} - \boldsymbol{\psi}(\mathbf{x})^{\top} \mathbf{r} - \boldsymbol{\epsilon}_{\mathbf{x}}^{\top} \mathbf{t}| \\ &\leq \bar{\epsilon} \|\sqrt{m} (\boldsymbol{\theta}_* - \boldsymbol{\theta}_0)\|_2 \\ &\leq \sqrt{2}S \bar{\epsilon}, \quad (66) \end{aligned}$$

where the last inequality holds due to (64). Therefore, we have

$$\begin{aligned}
& |h(\mathbf{x}) - \langle \phi(\mathbf{x}), \hat{\boldsymbol{\theta}}_\star \rangle| \\
&= |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \boldsymbol{\theta}_\star - \boldsymbol{\theta}_0 \rangle - \langle \phi(\mathbf{x}), \hat{\boldsymbol{\theta}}_\star \rangle| \\
&\leq \|\boldsymbol{\theta}_\star - \boldsymbol{\theta}_0\|_2 \|\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0) - \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1})\|_2 + |\langle \psi(\mathbf{x}), \hat{\boldsymbol{\theta}}_\star \rangle - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \boldsymbol{\theta}_\star - \boldsymbol{\theta}_0 \rangle| \\
&\leq \sqrt{2} S C_{g,2} N_{k-1}^{1/6} \alpha^{-1/6} m^{-1/3} \sqrt{\log mL}^{7/2} + \sqrt{2} S \bar{\epsilon} \\
&\leq 2S\bar{\epsilon},
\end{aligned}$$

where at the last line we use the fact that $m = \text{poly}(L, N_{k-1}, \alpha)$. \square

F.3.3 Proof of Lemma 10

Proof of Lemma 10. First we bound $|\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m|$. We have

$$\begin{aligned}
& |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m| \\
&\leq |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m| \\
&\quad + |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m| \\
&\leq 2C_{g,1} C_{g,2} N_{k-1}^{1/6} \alpha^{-1/6} m^{-1/3} \sqrt{\log mL}^4, \tag{67}
\end{aligned}$$

where the first inequality holds due to triangle inequality, the second one holds due to (35). Therefore, when $m = \text{poly}(|\mathcal{X}|, L, \lambda_0^{-1}, \log(|\mathcal{X}|/\delta_k), N_k, \alpha, \bar{\epsilon}^{-1})$, we have

$$\begin{aligned}
& |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| \\
&\leq |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| \\
&\quad + |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m| \\
&\leq |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_0) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| + 2C_{g,1} C_{g,2} N_{k-1}^{1/6} \alpha^{-1/6} m^{-1/3} \sqrt{\log mL}^4 \\
&\leq \bar{\epsilon}^2 / (2|\mathcal{X}|^4). \tag{68}
\end{aligned}$$

Next, by the definition of $d_{\text{eff}}(\epsilon)$, we know that

$$\sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \lambda_i(\mathbf{H}) \leq \bar{\epsilon}^2 / |\mathcal{X}|, \tag{69}$$

which suggests that

$$\begin{aligned}
\sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \lambda_i(\mathbf{G}\mathbf{G}^\top) &\leq \sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \lambda_i(\mathbf{H}) + \sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \|\mathbf{G}\mathbf{G}^\top - \mathbf{H}\|_2 \\
&\leq \bar{\epsilon}^2 / (2|\mathcal{X}|) + |\mathcal{X}|^3 \cdot \max_{\mathbf{x}, \mathbf{x}'} |\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}_{k-1}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}_{k-1}) \rangle / m - \mathbf{H}(\mathbf{x}, \mathbf{x}')| \\
&\leq \bar{\epsilon}^2 / (2|\mathcal{X}|) + \bar{\epsilon}^2 / (2|\mathcal{X}|) \\
&= \bar{\epsilon}^2 / |\mathcal{X}|. \tag{70}
\end{aligned}$$

Finally, we have

$$\sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \lambda_i(\mathbf{G}) = \sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \sqrt{\lambda_i(\mathbf{G}\mathbf{G}^\top)} \leq \sqrt{|\mathcal{X}| \sum_{i=d_{\text{eff}}(\bar{\epsilon}^2/|\mathcal{X}|)+1}^{|\mathcal{X}|} \lambda_i(\mathbf{G}\mathbf{G}^\top)} \leq \bar{\epsilon}. \tag{71}$$

\square

G Implementation details

We provide further details of experiments presented in Section 6. We ran all experiments on a Xeon Gold 6130 CPU, except for NeuralEmbedding where we used an Nvidia GeForce RTX2080Ti GPU. All the datasets we are using are either synthetic or publicly available and do not contain personally identifiable information or offensive content.

G.1 Experimental setups

All algorithms are run at confidence level of $\delta = 0.05$, and a misspecification level of $\epsilon = 0.1$. For each algorithm that requires computing a design λ (or the associated value τ of the minimax optimization problem), we employed a Frank-Wolfe algorithm and set the step-size and convergence rate in the same way as in [13]. Since our algorithms are all elimination-styled, the (empirical) sample complexity is calculated as the number of samples pulled when the set of uneliminated arms only consists of ϵ -optimal arms. All algorithms were forced to stop when the number of pulls reaches 10^7 (and such a case was reported as a failure). We choose 10^7 as the stopping criteria since continuing running an algorithm after pulling a large number of arms (10^7) is likely to raise memory issues on the running machine. For a fair comparison, each experiment is repeated 50 times and we report sample complexity results based on successful instances runs.

KernelEmbedding was instantiated with Gaussian kernel $\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma\|\mathbf{x}_i - \mathbf{x}_j\|_2^2)$ over action set \mathcal{X} . To decide the value of γ in KernelEmbedding as well as C (in Section 3.2) in KernelEmbedding and LinearEmbedding, for each dataset, we ran the method with several different seeds, used a grid search over $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$, and monitored the interested metrics (empirical sample complexity and the success rate). For KernelEmbedding and LinearEmbedding, we implemented Algorithm 1 for synthetic datasets; for Yahoo and MNIST datasets, due to the difficulty of the data, e.g., nonlinearity, we found that Algorithm 1 with adaptive embedding might compress arms into undesired inseparable subspaces in earlier rounds, which worsens the empirical performance. As a result, we implemented a slightly modified version of Algorithm 1 (i.e., Algorithm 3; see details in Remark 6) for Yahoo and MNIST datasets for better empirical performance (sample complexity and success rate). To save computation, we use $4(1 + \zeta)d$ as an approximation of $g(d, \zeta)$ when computing d_k for KernelEmbedding and LinearEmbedding (recall that we always have $g(d, \zeta) \leq 4(1 + \zeta)d$).⁸

The model used in Neural Embedding is a three-layer (i.e., two hidden layers) fully connected neural network, where each hidden layer has 128 nodes. We used a Rectified Linear Unit (ReLU) as our activation function. To overcome overfitting on real datasets, we add one dropout layer with 50% dropout rate. The learning rate is set as 10^{-4} and the maximum training iteration number is set as 6,000. A grid search over $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ is used to determine the $\bar{\epsilon}$ parameter (in Algorithm 2).

G.2 Detailed descriptions of datasets

Synthetic dataset with linear rewards. We tested all methods with varying the number of arms (K) and a fixed dimension ($D = 20$). To construct a set of arms \mathcal{X} , we first randomly generate $\mathbf{x}_1 = \boldsymbol{\theta}_*$ so that $\langle \boldsymbol{\theta}_*, \mathbf{x}_1 \rangle = 0.8$; we then generate a \mathbf{x}_2 such that $\langle \boldsymbol{\theta}_*, \mathbf{x}_2 \rangle \approx 0.4$. After generating the two principal arms \mathbf{x}_1 and \mathbf{x}_2 , we constructed an action set \mathcal{X} by adding randomly perturbed arms, i.e., $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_1 \oplus \{\eta_i \mathbf{e}_{i_j}\}_{i=1}^{\frac{K}{2}-1}, \mathbf{x}_2 \oplus \{\eta_i \mathbf{e}_{i_j}\}_{i=1}^{\frac{K}{2}-1}\}$, where i_j is randomly chosen from $[D]$ and $\eta_i \sim \mathcal{N}(0, (10^{-5})^2)$. We ensure $\text{span}(\mathcal{X}) = \mathbb{R}^D$ in the construction. We used Bernoulli reward with $\mathbf{x}^\top \boldsymbol{\theta}_*$ success probability in this experiment.

Synthetic dataset with nonlinear rewards. We tested all algorithms with increasing dimensions (D) and a fixed number of arms ($K = 200$). The reward of an arm is set as the 2-norm of its feature representation. We randomly generate \mathbf{x}_1 and \mathbf{x}_2 such that $h(\mathbf{x}_1) = 0.8$ and $h(\mathbf{x}_2) = 0.4$. We construction our action set \mathcal{X} by adding randomly perturbed arms, i.e., $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_1 \oplus \{\eta_i \mathbf{e}_{i_j}\}_{i=1}^{\frac{K}{2}-1}, \mathbf{x}_2 \oplus \{\eta_i \mathbf{e}_{i_j}\}_{i=1}^{\frac{K}{2}-1}\}$, where i_j is randomly chosen from $[D]$ and $\eta_i \sim \mathcal{N}(0, (10^{-5})^2)$. We ensure $\text{span}(\mathcal{X}) = \mathbb{R}^D$ in the construction. We used Bernoulli reward with $h(\mathbf{x})$ success probability in this experiment.

Yahoo dataset. We used Yahoo! User Click Log Dataset R6A in this experiment.⁹ The dataset contains the users' click-through records from the Today news module on Yahoo! front page between May 1st, 2009 and May 10th, 2009. Each user click log record consists of 6 article features and 6 user features, along with a binary variable stating whether or not a user clicked on the article.

⁸One can also determine d_k in a binary search manner if monotonicity of (an upper bound of) $\gamma(d)$ can be guaranteed, e.g., see examples in the proof of Theorem 2.

⁹<https://webscope.sandbox.yahoo.com>

To process the data, we considered the records collected in the 1st day and randomly selected 200 records from it. We construct an arm set $\mathbf{X} \in \mathbb{R}^{200 \times 36}$ by taking the flattened outer product of the user and the corresponding article feature vector. And in our further examination, we found that the rank of \mathbf{X} is 35 instead of 36 (full rank). To ensure full rank, we preprocess \mathbf{X} by projecting it into a lower-dimensional space $\widetilde{\mathbf{X}} \in \mathbb{R}^{200 \times 35}$. The reward of each arm is determined as in Section 6. Noise in the observed rewards was generated from a standard normal distribution.

MNIST dataset. The dataset is described as in Section 6. Noise in the observed rewards was generated from a standard normal distribution.

G.3 Empirical effective dimension d_k

For implementations with a fixed embedding (i.e., Algorithm 3), d_k is fixed over all iterations. For implementations with an adaptive embedding (i.e., Algorithm 1 and 2), d_k is increasing with iteration index k since we decrease tolerance to misspecification levels in later iterations. For experimental results of synthetic dataset in Figure 1, due to the simplicity of the data, NeuralEmbedding is able to compress arms in \mathbb{R}^2 and completed the elimination in one round, and Alg. 1-KernelEmbedding and Alg. 1-LinearEmbedding are able to compress arms in \mathbb{R}^2 and completed the elimination in two rounds. For experimental results of real-world datasets in Figure 2, our KernelEmbedding and LinearEmbedding are implemented with fixed embedding (as discussed in Appendix G). For NeuralEmbedding, at each iteration k , we calculated the averaged empirical value of d_k ; we also use s_k to denote the percentage of runs that have successfully completed the elimination process at round k . We summarized the pair value (d_k, s_k) in Table 2.

Table 2: Empirical effective dimension and success rate on real-world datasets

NeuralEmbedding	MNIST	Yahoo
(d_1, s_1)	(190.22, 82%)	(72.70, 86%)
(d_2, s_2)	(192.11, 98%)	(167.57, 100%)