

Contexts can be Cheap: Solving Stochastic Contextual Bandits with Linear Bandit Algorithms

Osama A. Hanna

Lin F. Yang

Christina Fragouli

University of California, Los Angeles

OHANNA@UCLA.EDU

LINYANG@UCLA.EDU

CHRISTINA.FRAGOULI@UCLA.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

In this paper, we address the stochastic contextual linear bandit problem, where a decision maker is provided a context (a random set of actions drawn from a distribution). The expected reward of each action is specified by the inner product of the action and an unknown parameter. The goal is to design an algorithm that learns to play as close as possible to the unknown optimal policy after a number of action plays. This problem is considered more challenging than the linear bandit problem, which can be viewed as a contextual bandit problem with a *fixed* context. Surprisingly, in this paper, we show that the stochastic contextual problem can be solved as if it is a linear bandit problem. In particular, we establish a novel reduction framework that converts every stochastic contextual linear bandit instance to a linear bandit instance, when the context distribution is known. When the context distribution is unknown, we establish an algorithm that reduces the stochastic contextual instance to a sequence of linear bandit instances with small misspecifications and achieves nearly the same worst-case regret bound as the algorithm that solves the misspecified linear bandit instances. As a consequence, our results imply a $O(d\sqrt{T}\log T)$ high-probability regret bound for contextual linear bandits, making progress in resolving an open problem in (Li et al., 2019b, 2021). Our reduction framework opens up a new way to approach stochastic contextual linear bandit problems, and enables improved regret bounds in a number of instances including the batch setting, contextual bandits with misspecifications, contextual bandits with sparse unknown parameters, and contextual bandits with adversarial corruption.

Keywords: Bandits, contextual linear bandits, online learning

1. Introduction

Linear bandit and contextual linear bandit problems are attracting extensive attention - for example, more than 17,000 papers appear when searching for “linear contextual bandit” on Google Scholar during the last 5 years - as they enable to support impactful active learning applications through elegant formulations. In linear bandits, a learner at each time $t \in [T]$, where T is the time horizon, pulls an arm a_t from a fixed action space \mathcal{A} (that may be continuous or discrete), and receives a reward $r_t = \langle a_t, \theta_\star \rangle + \eta_t$, where θ_\star is an unknown d -dimensional vector of parameters and η_t is random noise. Contextual linear bandits add another layer of complexity by enabling at each round the action space to be different to capture context; in this case, the learner at time t observes an action space (context) \mathcal{A}_t . That is, we can think of linear bandits as single-context contextual bandits, observing $\mathcal{A}_t = \mathcal{A}$ for all t . For example, while linear bandits are used in recommendation systems where the set of actions is fixed and oblivious to the individual the recommendation is addressed to, contextual linear bandits are used in personalized recommendations, where the action space gets tailored to context attributes such as age, gender, income and interests of each individual.

Algorithm	Regret Bound	Type	Assumption/restriction
Abbasi-Yadkori et al. (Li et al., 2021) (Li et al., 2021) Ours	$O(d\sqrt{T} \log T)$ $O(d\sqrt{T} \log T \text{poly}(\log \log T))$ $O(d\sqrt{T} \log T \text{poly}(\log \log T))$ $O(d\sqrt{T} \log T)$	w.h.p. exp w.h.p. w.h.p.	stochastic context
(Ruan et al., 2021) Ours	$O(d\sqrt{T} \log(d) \log(T) \log \log T)$ $O(d\sqrt{T} \log(T) \log \log T)$	exp w.h.p.	batch learning with $O(\log \log T)$ batches
(Foster et al., 2020) Ours	$O(d\sqrt{T} \log T + \epsilon\sqrt{dT})$ $O(d\sqrt{T} \log T + \epsilon\sqrt{dT} \log T)$	exp w.h.p.	misspecified
(Foster et al., 2020) Ours	$\tilde{O}(d^{4.5}\sqrt{T} + d^4C)$ $\tilde{O}(d\sqrt{T} + d^{3/2}C)$	w.h.p. w.h.p.	adversarial corruption
Abbasi-Yadkori et al. Ours	$O(\sqrt{dsT} \log T)$ $O(\sqrt{dsT} \log T)$	w.h.p. w.h.p.	sparse
Ours	$O(\sqrt{dsT} \log T \log \log T)$	w.h.p.	sparse with $O(\log \log T)$ batches

Table 1: Comparison of best known in literature vs. our approach regret bounds. Here, d is the model dimension, T is the time horizon, ϵ is an upper bound on the amount of misspecification, C limits the power of adversary, s is an upper bound on the number of non-zero elements in the unknown parameter, exp indicates a regret bound in expectation, w.h.p. indicates a regret bound that holds with probability at least $1 - 1/T$, and \tilde{O} hides log factors.

It is not surprising that, although more limited in applications, linear bandits are much better understood in theory than contextual linear bandits. Indeed, algorithms for linear bandits often leverage the fixed action space property, and cannot be easily extended to contextual linear bandits. To give a concrete example, the algorithm Phased Elimination (PE) (Lattimore et al., 2020; Valko et al., 2014) leverages the fixed action space by exploring a (small) core set of actions to achieve good estimates of the rewards for all actions. This algorithm achieves a high probability regret bound of $O(d\sqrt{T} \log T)$. Nevertheless, despite several attempts over the last decade (Abbasi-Yadkori et al., 2011; Li et al., 2021, 2019b), the best known regret upper bounds for contextual linear bandits have a log (or iterated log) multiplicative gap over the $O(d\sqrt{T} \log T)$ bound both in high probability and in expectation. Similarly, the best known algorithms for several linear bandits problems (e.g., with misspecification, adversarial corruption, and others (Lattimore et al., 2020; Foster et al., 2020; Bogunovic et al., 2021; Wei et al., 2022; Ruan et al., 2021)), perform better (in the worst-case) than the corresponding algorithms for contextual linear bandits.

1.1. Our Results

THE REDUCTION

We show in this paper the surprising result that, provided the context comes from a distribution \mathcal{D} (stochastic context), contextual linear bandit problems can be reduced to solving (single context) linear bandit problem when the context distribution \mathcal{D} is known, and to linear bandits with $\tilde{O}(1/\sqrt{T})$

misspecification when the distribution \mathcal{D} is unknown. These results are presented in the following informal theorems and their exact statements are given in Theorems 1, 2 and Theorem 4 in App. A.

Informal Statement of Theorem 1. For any contextual linear bandit instance I with known context distribution \mathcal{D} , there exists (constructively) a linear bandit instance L with the same action dimension, and any algorithm solving L solves I with the same worst-case regret bound as L .

Informal Statement of Theorem 2. For any contextual linear bandit instance I with unknown context distribution \mathcal{D} , there exist (constructively) $\log T$ misspecified linear bandit instances $L_1, \dots, L_{\log T}$, where L_i operates on part of the horizon of length T_i , has $\tilde{O}(1/\sqrt{T_i})$ -misspecification and the same action dimension, and any algorithm solving $L_1, \dots, L_{\log T}$ solves I with the same worst-case regret bound as $L_1, \dots, L_{\log T}$.

Informal Statement of Theorem 4. For any contextual linear bandit instance I with unknown context distribution \mathcal{D} but where now the action space has a component-wise product structure, there exists (constructively) a linear bandit instance L with double the action dimension of I , and any algorithm solving L solves I with the same worst-case regret bound as L .

Stochastic contextual bandits encompass practical cases where the context is not selected adversarially; in our example of personalized recommendations, the age, gender, income, come from distributions. Our framework simplifies the contextual linear bandit problem and enables to use any existing (or future) linear bandit algorithms to solve the contextual case. Moreover, our results separate the case of stochastic contexts from the harder case of adversarial contexts and explain why good results, which are not achievable for adversarial contexts, are possible for stochastic contexts.

IMPLICATIONS AND RELATED WORK

The equivalence we proved opens up a new way to approach stochastic contextual linear bandit problems, and results in a number of new results or recovery of existing results in a simpler manner; we next present some of these implications (summarized in Table 1) and their positioning with respect to related work. In the discussion next, d is the model dimension and T is the time horizon. These implications are stated formally and proved in App. B.

- **Tighter Regret Bounds.** To the best of our knowledge, there is a gap in the regret bounds of contextual linear bandits: the state-of-the-art lower bounds are, $\Omega(d\sqrt{T})$ for linear bandits (Lattimore and Szepesvári, 2020) and $\Omega(d\sqrt{T \log T})$ for linear contextual bandit with adversarial contexts (Li et al., 2019b).

Our contribution [Corollary 5 in App. B]. Our approach achieves a regret upper bound $O(d\sqrt{T \log T})$ with high probability even when the action set is infinite. While it is not known if a $\Omega(d\sqrt{T \log T})$ lower bound holds for stochastic contexts, our result improves over state of the art high probability bounds by at least a factor of $\sqrt{\log T}$ and matches the best known upper bound for linear bandits.

Related Work. The best attempts of upper bounds are (Abbasi-Yadkori et al., 2011; Li et al., 2021, 2019b). In particular, (Abbasi-Yadkori et al., 2011) achieves a regret bound of $O(d\sqrt{T \log T})$ with high probability; (Li et al., 2021) achieves $O(d\sqrt{T \log T} \text{poly}(\log \log T))$ in expectation and $O(d\sqrt{T \log T} \text{poly}(\log \log T))$ with high probability; and (Li et al., 2019b) achieves $O(d\sqrt{T \log T} \text{poly}(\log \log(T)))$ in expectation and only when the number of actions is finite and bounded by $2^{d/2}$.

• **Batch Learning.** In batch learning, instead of observing the reward at the end of each round to decide what action to play next, the learning agent is constrained to split the rounds into a fixed number of M batches, use a predetermined policy within each batch, and it can only observe the action outcomes and switch the policy at the end of each batch. This is a central problem in online learning (Gao et al., 2019; Perchet et al., 2016; Ruan et al., 2021; Han et al., 2020) as limited policy adaptivity enables parallelism and facilitates deployment of learning algorithms for large-scale models.

Our Contribution [Theorem 6 in App. B]. Our batch algorithm for contextual linear bandits achieves with *high probability* $O(d\sqrt{T}\log(T)\log\log T)$ regret bound for $O(\log\log T)$ batches, a $\sqrt{\log d}$ better than the *in expectation* regret bound in (Ruan et al., 2021).

Related Work. A number of works have explored batched contextual linear bandits, both in adversarial (Abbasi-Yadkori et al., 2011; Han et al., 2020) and non-adversarial settings (Ruan et al., 2021; Zhang et al., 2021). The breakthrough work in (Ruan et al., 2021) achieved a nearly optimal *in expectation* regret upper bound $O(d\sqrt{T}\log d\log(T)\log\log T)$ using $O(\log\log T)$ batches for the stochastic contexts setting. The near-optimality follows from the result in (Gao et al., 2019) which shows that $\Omega(\log\log T)$ batches are required to achieve a $O(\sqrt{T})$ worst-case regret bound for multi-armed bandits with a finite number of arms (this is a special case of contextual linear bandits). If the contexts are chosen adversarially, $\Omega(\sqrt{T})$ batches are required to achieve a $O(\sqrt{T})$ regret bound (Han et al., 2020).

• **Misspecified Bandits.** Linear bandit algorithms are designed under the assumption that the expected rewards are perfectly linear functions of the actions; misspecified bandits relax this assumption by considering perturbations of the linear model measured by the amount of deviation in the expected rewards (we call the case ϵ -misspecification if the deviation is upper bounded by ϵ). The non-linearity in the model enables to better capture real-world environments and is of high interest in the literature (Du et al., 2019; Crammer and Gentile, 2013; Lattimore et al., 2020; Foster et al., 2020; Ghosh et al., 2017; Foster and Rakhlin, 2020).

Our Contribution [Theorem 7, 8 in App. B]. We provide a regret bound of $O(d\sqrt{T}\log T + \epsilon\sqrt{dT}\log T)$ with high probability for contextual bandits with unknown misspecification, and $O(d\sqrt{T}\log T + \epsilon\sqrt{dT})$ with high probability for ϵ known. To the best of our knowledge, these results offer the first *optimal regret bounds* for ϵ -misspecified contextual linear bandits, and improve over existing literature for unknown ϵ by providing high probability bounds and improved log factors. We also present the first nearly optimal algorithm for misspecified contextual linear bandits with $O(\log T)$ batches.

Related Work. The work in (Lattimore et al., 2020) shows that PE with modified confidence intervals achieves the optimal regret bound of $O(d\sqrt{T}\log T + \epsilon\sqrt{dT})$ with high probability (matching the $\Omega(\epsilon\sqrt{dT})$ lower bound (Lattimore et al., 2020)) for linear bandits with known ϵ misspecification. If ϵ is unknown, the same algorithm was shown to achieve a regret bound of $O(d\sqrt{T}\log T + \epsilon\sqrt{dT}\log T)$ with high probability. The work in (Zanette et al., 2020; Lattimore et al., 2020) proposed variants of LinUCB that achieve a regret bound of $\tilde{O}(d\sqrt{T} + \epsilon\sqrt{dT})$ ¹ with high probability for contextual linear bandits with known ϵ . However, changing action sets with unknown ϵ was left as an open problem. The works in (Pacchiano et al., 2020; Foster et al., 2020) made progress in answering this question by providing a regret bound of $O(d\sqrt{T}\log T + \epsilon\sqrt{dT})$ in expectation (Foster et al., 2020). However,

1. Here \tilde{O} hides log factors.

improving the log factors in the first term and strengthening the result to a high probability bound was unresolved. Due to the challenge that changing action sets impose on the analysis, the techniques of (Foster et al., 2020) cannot be directly extended to provide high probability bounds on the regret. Our result of $O(d\sqrt{T\log T} + \epsilon\sqrt{dT\log T})$ regret removes the $\sqrt{\log T}$ factor from the first term and adds a $\log T$ factor in the second term; depending on the value of ϵ this can lead to a tighter or looser bound.

• **Bandits with Adversarial Corruption.** Linear bandits with adversarial corruption recently attracted significant interest (Lykouris et al., 2021; Chen et al., 2021; Jin et al., 2021; Li et al., 2019a; Bogunovic et al., 2020, 2021; Lee et al., 2021; Wei et al., 2022) due to the vulnerability of online learning applications to attacks. There are multiple adversarial models that are proposed in the literature; here we consider a widely used model that assumes the adversary knows the policy, and observes the history, but does not observe the current action before corrupting the rewards.

Our Contribution [Theorem 9 in App. B]. Our work provides the first algorithm for contextual linear bandits that achieves a regret bound of $\tilde{O}(d\sqrt{T} + d^{3/2}C)$ with high probability for unknown C , which upper bounds the total amount of corruption from the adversary. This improves over the best known $\tilde{O}(d^{4.5}\sqrt{T} + d^4C)$ bound for linear bandits with changing action sets in (Wei et al., 2022). We note that in our regret bound, while the dependency on d in the first term is nearly optimal, the dependency on $d^{3/2}$ in the second term is not. However, we simplify the problem of improving this dependency, as any algorithm that improves it for linear bandits will imply the same improvement for contextual bandits with our reduction.

Related Work. The work in (Bogunovic et al., 2021) considers linear bandits with adversarial corruption and achieves a regret bound of $\tilde{O}(d\sqrt{T} + d^{3/2}C)$ with high probability for known corruption level C and a regret bound of $\tilde{O}(d\sqrt{T} + d^{3/2}C + C^2)$ with high probability for a stronger adversary that observes the current action and unknown corruption level C , while the work in (Wei et al., 2022) achieves a regret bound of $\tilde{O}(d\sqrt{T} + d^{3/2}C)$ with high probability for unknown corruption C . These algorithms have optimal dependency on T, C but it is not known if the $d^{3/2}$ dependency is tight or not. The work in (Lee et al., 2021) improves the dependency on d for linearized corruption, by achieving a regret bound of $\tilde{O}(d\sqrt{T} + C)$ with high probability which is also nearly optimal (Lee et al., 2021). The proposed algorithms and analysis rely on the assumption that the action set is fixed, and as far as we know there are no known generalizations for changing action sets - beyond the work in (Bogunovic et al., 2021) that considers changing action sets but imposes a strong assumption on the context distribution. The work in (He et al., 2022) considers a stronger adversary, that observes the current action, and achieves an $\tilde{O}(d\sqrt{T})$ regret bound for unknown $C \leq \sqrt{T}$ and linear regret otherwise. The first paper to prove a regret bound with nearly optimal dependency on T, C for linear bandits with changing action sets (linear contextual bandits) is (Wei et al., 2022) which achieves a regret bound of $\tilde{O}(d^{4.5}\sqrt{T} + d^4C)$ with high probability for unknown C . While the dependency on T, C is nearly optimal, the dependency on d is clearly not - improving this was left as an open problem. Our results take a step in this direction by removing a factor of $d^{3.5}$ from the first term, by removing a factor of $d^{2.5}$ from the second term and by reducing the problem of further improving the dependency of d in the second term to achieving this improvement over a linear bandit setup.

• **Sparsity.** High dimensional linear bandits with sparsity capture practical cases such as, when there exist a large number of candidate features and limited information on which of them are useful; use cases include personalized medicine and online advertising (Bastani and Bayati, 2020). This setup results in sparsity in the unknown linear bandit parameters, which can be leveraged for more efficient

learning.

Our Contribution [Theorem 10 in App. B]. Our work provides the first $O(\sqrt{dsT \log T})$ regret bound with high probability for contextual linear bandits, improving a factor of $\sqrt{\log T}$ on the state of the art, where s is an upper bound of the nonzero elements in the model.

Related Work. Due to its practical significance, a number of works have examined this setup (Lattimore et al., 2015; Abbasi-Yadkori et al., 2012; Gerchinovitz, 2011; Carpentier and Munos, 2012; Chen et al., 2022; Dai et al., 2022; Jang et al., 2022; Hao et al., 2021, 2020). To the best of our knowledge, the best known regret bound is $O(\sqrt{dsT \log T})$ with high probability (Abbasi-Yadkori et al., 2012). While this is shown to be nearly optimal (Lattimore and Szepesvári, 2020), improving the $\log T$ factor was left as an open problem. Our work resolves this.

1.2. Technical Overview

Our major technique innovation is in the conception and execution of a reduction from a stochastic contextual linear bandit instance to a linear bandit instance. This reduction is made possible by establishing a linear bandit action for each possible parameter θ of the contextual bandit instance. In particular, for a given θ , we establish a new action $g(\theta)$ that is the expected best action (under the distribution of the context \mathcal{A}_t) with respect to the parameter θ , i.e., $g(\theta) = \mathbb{E}_{\mathcal{D}}(\arg \max_{\tilde{a} \in \mathcal{A}_t} \langle \tilde{a}, \theta \rangle)$. Note that $g(\theta)$ may not correspond to any valid action for the present context in the corresponding contextual bandit instance. Yet, we show that if one plays an action $\arg \max_{\tilde{a} \in \mathcal{A}_t} \langle \tilde{a}, \theta \rangle$ that is optimal with respect to θ , then the contextual bandit instance generates a *linear reward* with respect to $g(\theta)$. Moreover, the linear bandit instances share the same optimal parameter θ_* as in the contextual instance. Under standard boundedness assumptions of the contexts and actions, the reward noise in the reduced linear instance and the contextual bandit instance also share a similar sub-Gaussian tail. By mapping the linear bandit action $g(\theta)$ to the contextual bandit $\arg \max_{\tilde{a} \in \mathcal{A}_t} \langle \tilde{a}, \theta \rangle$ for any context \mathcal{A}_t and θ , any algorithm for the linear bandit problem can be immediately applied to solve for the contextual bandit problem and suffer no additional regret in the worst-case.

The reduction becomes more challenging when the context distribution is unknown. One idea is to estimate $g(\theta)$ for all possible θ . Unfortunately, doing so would require a large number of samples resulting in unbounded regret. We resolve this issue by a batched approach where the batches provide increasingly better estimates of $g(\theta)$. In each batch, we estimate $g(\theta)$ using all the contexts generated in the previous batch. Note that this inevitably introduces error which ruins the linearity of the collected reward from the contextual bandit instance. Hence, we can only apply algorithms that are designed for misspecified linear bandits. Luckily, with a carefully designed batch sequence, we show that a linear bandit algorithm that works for all misspecification levels $\epsilon \in [1/\sqrt{T}, 1]$ can be applied to solve the contextual bandit instance. As it is hard to guarantee a good estimate of g for all θ , we restrict our attention to a finite subset of the unknown parameter set Θ (discretization) that is guaranteed to contain a good action. The amount of misspecification is bounded using a union bound argument over the discretization of Θ . While the discretization of Θ may eliminate the optimal arm and the function g is shown to be non-smooth, we show that the function $r(\theta) = \langle g(\theta), \theta_* \rangle$ is smooth on a neighborhood of θ_* . This is sufficient to show that if discretized finely, the discrete set will contain a good arm. The final worst-case regret bound can be controlled by the regret bound of the linear bandit algorithm and the batch lengths.

We next provide a high level explanation on why our reduction enables **tighter regret bounds** for contextual linear bandits with stochastic contexts. Using the Phased Elimination algorithm

(PE) (Lattimore et al., 2020) with modified confidence intervals to solve the misspecified linear bandit instance, we prove a high probability bound of $O(d\sqrt{T\log T})$ on the regret of the contextual bandit problem using our reduction. Existing analysis techniques for changing action sets suffer log factors due to several reasons, such as bounding the radius of an ellipsoid that contains the unknown parameter with high probability, bounding the regret by the radius of the ellipsoid, or summing the regret over multiple episodes; in contrast, the log factors in our analysis appear only from a union bound used to prove concentration of the estimated arm means with high probability.

The reduction also implies improved regret bounds for **misspecified contextual linear bandits** by leveraging the optimal and nearly optimal regret bounds proved in the literature (Lattimore et al., 2020) for fixed action sets. This results in high probability bounds; as opposed to the best known results for changing action sets that add log factors and only hold in expectation. As the reduction introduces a small $\tilde{O}(1/\sqrt{T})$ misspecification, this adds to the possibly unknown misspecification ϵ . There is a little subtlety here; some algorithms adapt to known ϵ better than the unknown case. To avoid adding extra factors to the regret bounds by sub-optimal adaptation to the known $\tilde{O}(1/\sqrt{T})$ part of misspecification, we slightly modify existing algorithms to account for this. This is done by increasing the number of times an arm is explored by a constant factor. Similarly, our results allow to carry over the better regret bounds for linear bandits with **adversarial corruption** to the contextual setting. This is achieved by modifying the algorithms for linear bandits with adversarial corruption to account for the known misspecification added by our reduction.

Our reduction has a byproduct; it limits the size of available actions in each round to the minimum between $|\mathcal{A}|$ and the size of the parameter set Θ (recall that we construct an action $g(\theta)$ for each $\theta \in \Theta$). While in general a discretization of size $T^{\Omega(d)}$ is required to guarantee a good action in the discrete set, if the unknown parameter follows some structure, a discretization of smaller size can be sufficient. For example, for contextual linear bandits with **s-sparse unknown parameter**, we show that a discretization of size $T^{O(s)}$ is sufficient. This directly implies a high probability $O(\sqrt{dsT\log T})$ regret bound as opposed to the best known $O(\sqrt{dsT\log T})$ bound (Abbasi-Yadkori et al., 2012) that loses extra log factors due to solving a linear regression problem over the space of sparse unknown parameters.

As it is enough to make only batch updates to our estimates of the actions $g(\theta)$, our algorithm can be modified to provide improved regret bounds for contextual linear bandits with $O(\log \log T)$ **batches**. We use batch lengths that were introduced in (Gao et al., 2019) which grow as $T^{1-2^{-m}}$. In each batch, as action set is fixed, we utilize elimination algorithms with the G-optimal design. However, at batch m , the gaps of sub-optimal actions depend on the confidence of our estimates in batch $m - 1$, which rely on the G-optimal design using the estimates of g from batch $m - 2$. As a result, the regret in batch m can be at most the ratio between the length of batch m and batch $m - 2$. Considering the growth rate of the batch lengths, this ratio can be large, especially in the first few batches. To fix this, we modify the batch lengths to grow in length only at the batches with odd index, while even batches use the same length as the previous batch. Our algorithm results in an $O(d\sqrt{T\log T}\log \log T)$ regret bound using $O(\log \log T)$ batches, improving a $\sqrt{\log d}$ factor (that appeared due to the distributional G-optimal design proposed in (Ruan et al., 2021)) in the regret bound over the best known result (Ruan et al., 2021). Our result also provides the first high probability bound under the $O(\log \log T)$ batches limitation.

As a consequence of the batch learning result and the fact that our reduction can limit the action set based on the size of Θ , we provide the first algorithm with $O(\log \log T)$ batches for the sparse

setting with a regret bounded by $O(\sqrt{dsT \log T} \log \log T)$ with high probability, where s is the sparsity parameter.

1.3. Paper organization

Section 2 describes our setup and reviews notation; Section 3 describes our reduction and main theorems; App. A describes a spacial case of the reduction; and App. B uses our reduction to prove improved regret bounds for a number of stochastic linear bandit problems.

2. Notation and Models

Notation. We use the following notation throughout the paper. For a vector X we use X_i to denote the i -th element of the vector X . The set $\{1, \dots, i\}$ for $i \in \mathbb{N}, i > 0$ is denoted by $[i]$, where \mathbb{N} is the set of natural numbers. We say that $y = O(f(x))$ if there is x_0 and a constant c such that $y \leq cf(x) \forall x > x_0$; we also use $\tilde{O}(f(x))$ to omit log factors. A δ -net (with respect to norm-2) of a set $\mathcal{A} \subseteq \mathbb{R}^d$ for $\delta > 0$ is any set $\mathcal{B} \subseteq \mathbb{R}^d$ such that for every $a \in \mathcal{A}$, there exists $b \in \mathcal{B}$ with $\|a - b\|_2 \leq \delta$, where \mathbb{R} is the set of real numbers. For a set of sets \mathcal{S} , $\cup \mathcal{S}$ denotes the union of all elements in \mathcal{S} . For a family of sets $\{\mathcal{A}_i\}_{i=1}^n$ we use $\prod_{i=1}^n \mathcal{A}_i = \{(a_1, \dots, a_n) | a_i \in \mathcal{A}_i, \forall i \in [n]\}$ to denote the product set.

Contextual Linear Bandits. We consider a contextual linear bandit problem, where a learner interacts with an environment over a time horizon of length T . At time $t \in [T]$, the learner observes a set of admissible arms \mathcal{A}_t representing the context, pulls an arm $a_t \in \mathcal{A}_t$, and receives a reward

$$r_t = \langle a_t, \theta_\star \rangle + \eta_t, \quad (1)$$

where the context \mathcal{A}_t is generated from a distribution \mathcal{D} (independently from other iterations), a_t is a function of the history $\mathcal{H}_t = \{\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_t\}$, θ_\star is an unknown parameter vector of dimension d , and η_t is a random noise. Here, the noise η_t follows an unknown distribution that satisfies $\mathbb{E}[\eta_t | \mathcal{F}_t] = 0$, $\mathbb{E}[\exp(\lambda \eta_t) | \mathcal{F}_t] \leq \exp(\lambda^2/2) \forall \lambda \in \mathbb{R}$ (sub-Gaussian), where $\mathcal{F}_t = \sigma\{\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_t, a_t\}$ is the filtration of all historic information up to time t , and $\sigma(X)$ is the σ -algebra generated by X . We follow the standard assumptions that $\theta_\star \in \Theta \subseteq \{\theta | \|\theta\|_2 \leq 1\}$, $\|a\|_2 \leq 1 \forall a \in \mathcal{A}_t$ and $\forall t \in [T]$ almost surely. The learner adopts a policy π that maps the history up to time t , $(\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_{t-1}, a_{t-1}, r_{t-1}, \mathcal{A}_t)$, to a probability distribution over \mathcal{A}_t ; we denote the policy $\pi(\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_{t-1}, a_{t-1}, r_{t-1}, \mathcal{A}_t)$ at time t by $\pi_t(\mathcal{A}_t)$. The goal of the learner is to minimize the regret defined as

$$R_T = \sum_{t=1}^T \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle - \langle a_t, \theta_\star \rangle. \quad (2)$$

In the next sections we assume for simplicity that for each $\theta \in \Theta$, there is a unique $a_t \in \mathcal{A}_t$ that satisfies $\langle a_t, \theta \rangle = \sup_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ almost surely. This is to avoid dealing with approximations and choice functions (if there are multiple optimal actions) in the description of our algorithms. However, our results do not need this assumption, please see App. C.1.

Batch Learning. In this setting, the learner is allowed to change the policy π_t only at M pre-specified time slots $1 \leq t^{(1)}, \dots, t^{(M)} \leq T$, where M is the number of batches.

Misspecified Linear Bandits. Here pulling an action $a \in \mathcal{A}_t$ generates a reward perturbed as

$$r_t = \langle a, \theta_\star \rangle + \eta_t + f(a), \quad (3)$$

where f is unknown perturbation function, θ_* is the unknown parameter vector, and η_t is a zero-mean noise that is 1-subgaussian conditioned on the history. The amount of deviation in the rewards is upper bounded by $\epsilon = \sup_{a \in \cup \text{Supp}(\mathcal{D})} |f(a)|$, and ϵ is called the amount of misspecification, where \mathcal{D} is the context distribution and $\text{Supp}(\mathcal{D})$ is the support set of \mathcal{D} .

Adversarial Corruption. We assume an adversary that operates as follows at each time t :

- The adversary observes the history of all past contexts $\mathcal{A}_1, \dots, \mathcal{A}_{t-1}$, actions a_1, \dots, a_{t-1} , rewards r_1, \dots, r_{t-1} , and previously corrupted rewards $\tilde{r}_1, \dots, \tilde{r}_{t-1}$, together with the current context \mathcal{A}_t .
- The adversary decides on a corruption function $c_t : \mathcal{A}_t \rightarrow \mathbb{R}$ that determines the amount of corruption for each action.
- The learner observes the history of contexts $\mathcal{A}_1, \dots, \mathcal{A}_t$, actions a_1, \dots, a_{t-1} , and previously corrupted rewards $\tilde{r}_1, \dots, \tilde{r}_{t-1}$.
- The learner pulls arm a_t and observes reward $\tilde{r}_t = r_t + c_t(a_t)$, where $c_t(a_t)$ is the corruption provided by the adversary.

Note that the true reward r_t follows the linear bandits model, while the corrupted rewards do not need to. We assume that the amount of corruption the adversary can inflict is bounded as

$$\sum_{t=1}^T \sup_{a \in \mathcal{A}_t} |c_t(a)| \leq C, \quad (4)$$

where C is the maximum amount of corruption.

Sparsity. We here assume that the d -dimensional parameter vector θ_* in (1) is sparse, namely $\|\theta_*\|_0 \leq s$ for some known $s \in [d]$, where $\|\theta_*\|_0$ denotes the norm-0 or cardinality of the vector θ_* .

3. Reduction from Stochastic Contextual to Linear Bandits

3.1. Reduction for Known Context Distribution \mathcal{D}

We construct a contextual linear bandit algorithm \mathcal{M} that operates at a high level as follows. At each time t , the learner:

Step 1 (plays): observes a set of actions \mathcal{A}_t , uses θ_t (i.e., the current estimate of θ^*) to decide which action $a_t \in \mathcal{A}_t$ to play, and observes the associated reward r_t ;

Step 2 (learns): calculates θ_{t+1} , i.e., an updated estimate of θ^* .

In our reduction, we use a single-context algorithm Λ for learning the parameter θ_* in step 2, i.e., choose θ_{t+1} , and prove that we can achieve the same worst-case regret bound as Λ .

Fixed action space. We provide to the linear bandit algorithm Λ the fixed action space

$$\mathcal{X} = \{g(\theta) | \theta \in \Theta\}, \text{ where } g(\theta) = \mathbb{E}_{\mathcal{A}_t \sim \mathcal{D}} [\arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle | \theta]. \quad (5)$$

That is, for each θ , we create an action $g(\theta)$ that is the expected best action (under the distribution of the context \mathcal{D}) with respect to the parameter θ . We illustrate using a simple example.

Example 1. Assume that we may observe one out of the two following action sets $\{[1], [-1]\}, \{[1]\}$ randomly with probability $1/2$. The function $g(\theta) : \mathbb{R} \rightarrow \mathbb{R}$ can be calculated as follows

$$g(\theta) = \begin{cases} 1 & \text{if } \theta \geq 0 \\ 0 & \text{if } \theta < 0 \end{cases} \quad (6)$$

and thus $\mathcal{X} = \{[0], [1]\}$.

Reduction Algorithm. The algorithm \mathcal{M} proceeds at each time t as follows:

1. The single context algorithm Λ decides to play an action $x_t \in \mathcal{X}$, using the history Λ observed. This action is never actually played. Instead, \mathcal{M} observes what action x_t that Λ selected, and uses a θ_t with $x_t = g(\theta_t)$ as its current estimate of θ_* , where ties are broken arbitrarily.
2. \mathcal{M} observes \mathcal{A}_t , plays the action $a_t = \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle$ and receives reward r_t . It provides this reward to Λ .
3. Λ assumes that the reward r_t it received was generated according to the linear bandit model $r_t = \langle g(\theta_t), \theta_* \rangle + \eta'_t$, and adds the action-reward pair $(g(\theta_t), r_t)$ to its history.

Note that the set of actions \mathcal{X} we created contains actions that may not be part of the original sets \mathcal{A}_t 's; this is fine, since these actions are actually never played - they are used to simulate an environment that enables Λ to correctly update its estimate of θ_* . That is, although all actions played come from the eligible sets \mathcal{A}_t , all learning (updates on θ_t) is derived from the single context algorithm Λ that never explicitly learns \mathcal{A}_t . We also highlight that

Theorem 1 *Let Λ be any algorithm for linear bandits and I be a contextual linear bandit instance with stochastic contexts, unknown parameter θ_* and rewards r_t generated as described in the reduction algorithm described above. It holds that*

- *The reward r_t is generated, by pulling the arm $g(\theta_t)$, from a linear bandit instance L with action set \mathcal{X} , and unknown parameter θ_* .*
- *The reduction results in an algorithm \mathcal{M} for contextual linear bandits such that with probability at least $1 - \delta$ we have*

$$|R_T^{\mathcal{M}}(I) - R_T^{\Lambda}(L)| \leq c\sqrt{T \log(1/\delta)}, \quad (7)$$

where $R_T^{\Lambda}(L)$ is the regret of Λ over the constructed linear bandit instance L , $R_T^{\mathcal{M}}(I)$ is the regret of algorithm \mathcal{M} over the instance I and c is a universal constant.

Proof Outline. The complete proof is provided in App. A; we here provide a brief outline. The basic idea is to show that the action taken by algorithm \mathcal{M} at time t is an unbiased estimate of $g(\theta_t)$, then decompose $R_T^{\mathcal{M}}(I)$ as

$$\begin{aligned} R_T^{\mathcal{M}}(I) &= \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle] \\ &\quad + \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle + \langle g(\theta_*) \rangle, \theta_* \rangle - \langle g(\theta_t) \rangle, \theta_* \rangle, \end{aligned} \quad (8)$$

where the expectation is with respect to the randomness in the context generation \mathcal{A}_t . We then prove and use results such as $\langle g(\theta'), \theta' \rangle = \max_{\theta \in \Theta} \langle g(\theta), \theta' \rangle$, $\forall \theta' \in \Theta$, to arrive at

$$\begin{aligned} |R_T^M(I) - R_T^A(L)| &\leq \left| \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle] \right| \\ &\quad + \left| \sum_{t=1}^T \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle \right|, \end{aligned} \quad (9)$$

Next, we show that the quantity

$$\Sigma_T := \sum_{t=1}^T \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle$$

is a martingale (using that a_t is unbiased estimate of $g(\theta_t)$) with a bounded difference (by boundedness of Θ, \mathcal{A}_t) and apply the Azuma–Hoeffding inequality to bound these terms with high probability. To prove that the rewards r_t come from a linear bandit instance L , we finally show that the reward r_t can be expressed as $r_t = \langle g(\theta_t), \theta_* \rangle + \eta'_t$, where η'_t is a zero mean 1-sub-Gaussian noise conditioned on the filtration of historic information $\theta_1, \dots, \theta_t$ and rewards r_1, \dots, r_{t-1} of the instance L . \square

Algorithm 1 Reduction from stochastic contexts to no context

Input: confidence parameter δ , phase lengths $\{t^{(m)}\}_{m=1}^{M+1}$, and algorithm Λ_ϵ for linear contextual bandits with ϵ misspecification.

Initialize: $g^{(1)} : \Theta' \rightarrow \mathbb{R}^d$ randomly, $\epsilon_1 = 1$, and let $\mathcal{X}_1 = \{g^{(1)}(\theta) | \theta \in \Theta'\}$.

for $m = 1 : M$ **do**

for $t = t^{(m)} + 1, \dots, t^{(m+1)}$ **do**

Let $g^{(m)}(\theta_t) \in \mathcal{X}_m$ be the arm selected by Λ_{ϵ_m} after observing rewards $r_{t^{(m)}+1}, \dots, r_{t-1}$.

Play arm $a_t = \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle$ and receive reward r_t . Provide r_t to Λ_{ϵ_m} .

end

Update: $g^{(m+1)}(\theta) = \frac{1}{t^{(m+1)}} \sum_{t=1}^{t^{(m+1)}} \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle$, $\mathcal{X}_{m+1} = \{g^{(m+1)}(\theta) | \theta \in \Theta'\}$, and $\epsilon_{m+1} = 2\sqrt{\log(M|\Theta'|/\delta)/t^{(m+1)}}$.

end

3.2. Reduction for Unknown Context Distribution

As described in (5), calculating the function $g(\theta)$ requires knowledge of the distribution \mathcal{D} . We do not really need this knowledge; we prove it is sufficient to use empirical estimation of $g(\theta)$. As a result, we prove that any stochastic linear contextual bandit instance, even for unknown context distributions, can be reduced to a linear bandit instance albeit with a small misspecification.

Our basic approach follows the reduction in Section 3.1 but uses a sequence of functions $g^{(m)}$ that approximate g increasingly well (as m increases). To do so, as it is hard to guarantee a good estimate of g for all θ , we restrict our attention to a finite subset of Θ that is large enough to include a good action. In particular, instead of considering actions in a continuous space Θ , we only consider a finite subset of actions $\Theta' \subseteq \Theta$, where Θ' is an $1/T$ -net for Θ according to the norm-2 distance. We

divide the time horizon T into M epochs, each of duration $T_m = t^{(m+1)} - t^{(m)}$, $m = 1 \dots M$. For each epoch, we construct an empirical estimate of g as $g^{(m)}(\theta) = \frac{1}{t^{(m)}} \sum_{t=1}^{t^{(m)}} (\arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle)$, and calculate the set of actions $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$. We then use at each epoch a single context algorithm as before, but we now provide at epoch m the fixed set of actions \mathcal{X}_m for the algorithm to choose from. As a result, the regret of the linear bandit problem is defined as

$$R_T^{\Lambda_\epsilon} = \sum_{t=1}^T \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle. \quad (10)$$

Our algorithm relies on Λ_ϵ , an algorithm for linear bandits with ϵ misspecification. The misspecification reflects our confidence in our estimate of the function g , hence, decreases each epoch. We start with a large value of the misspecification parameter $\epsilon_1 = 1$ and a random initialization of the function g (which we call $g^{(1)}$), and hence random initialization of \mathcal{X} denoted \mathcal{X}_1 . At time slot t of epoch m , the algorithm asks Λ_{ϵ_m} for an action to play $g^{(m)}(\theta_t) \in \mathcal{X}_m$ given the history of action and rewards $\{(g^{(m)}(\theta_i), r_i)\}_{i=t^{(m)+1}}^{t-1}$ in epoch m only. The algorithm pulls the action $a_t = \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle$ and observes r_t . The reward r_t is then passed to the algorithm Λ_{ϵ_m} . At the end of each epoch, the misspecification parameter, estimates $g^{(m)}$, and action set \mathcal{X}_m are updated. The pseudo-code of our reduction is provided in Algorithm 1, and the proof in App. B. To achieve nearly optimal regret bounds, the misspecification ϵ need to be $\tilde{O}(1/\sqrt{T})$ (recall the $\Omega(\sqrt{d\epsilon T})$ regret lower bound (Lattimore et al., 2020)). Attempting to use $M = 2$ to first estimate g and then learn in the second epoch, would require the length of the first epoch to be $\Omega(T)$, to ensure the $\tilde{O}(1/\sqrt{T})$ misspecification, resulting in linear regret. Instead, as we clarify next, we use exponentially increasing epoch lengths to mix the learning with a gradual estimation of g resulting in misspecification that is effectively $\tilde{O}(1/\sqrt{T})$.

Theorem 2 *Let Λ_ϵ be an algorithm for linear bandits with ϵ misspecification and I be a contextual linear bandit instance with stochastic contexts, unknown parameter θ_\star and rewards r_t are generated as described in Algorithm 1. The following holds:*

- *Conditioned on $\mathcal{H}_{t^{(m)}} = \sigma(\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_{t^{(m)}}, a_{t^{(m)}}, r_{t^{(m)}})$: in epoch m , the rewards r_t are generated, by pulling arm $g^{(m)}(\theta_t)$, from a misspecified linear bandit instance L_m for $t = t^{(m)} + 1, \dots, t^{(m+1)}$, action set $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$, unknown parameter θ_\star , mean rewards $\langle g(\theta), \theta_\star \rangle$, and unknown misspecification ϵ'_m .*
- *The misspecification ϵ'_m is bounded by the known quantity ϵ_m with probability at least $1 - c\delta/M$.*
- *With probability at least $1 - \delta$ we have that $|R_T(I) - \sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)| \leq c\sqrt{T \log(1/\delta)}$, where $R_T(I)$ is the regret of Algorithm 1 over the instance I , $R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$ is the regret of algorithm Λ_{ϵ_m} over the bandit instance L_m in phase m , $T_m = t^{(m+1)} - t^{(m)}$, and c is a universal constant.*

As a consequence, we prove the following corollary in App. B.1.

Corollary 3 *For Algorithm 1 with $t^{(m)} = 2^{m-1}$ and Λ_ϵ to be PE with modified confidence intervals (Lattimore et al., 2020), it holds that with probability at least $1 - c\delta$ we have that $R_T = O(d\sqrt{T \log(T/\delta)})$.*

In App. A we describe an important special case where the reduction can be performed without approximating g .

3.3. Computational Complexity

In this subsection, we show that our reduction can be made *computationally-efficient* with the help of optimization oracles and linear regression oracles (Awerbuch and Kleinberg, 2008; Zhu et al., 2022), which are standard assumptions when computation is a concern in bandit problems.

In particular, we are able to

- Efficiently construct a linear optimization oracle for the linear bandit instance L_m given a linear optimization oracle for the original contextual bandit instance I . Together with the linear regression oracle, this suffices to construct the linear bandit policy (Awerbuch and Kleinberg, 2008; Dani et al., 2008; Bubeck et al., 2012; Hazan and Karnin, 2016; Ito et al., 2019).
- Efficiently compute an inverse of the function $g^{(m)}$, which is performed in our reduction to translate the linear bandit action to an action in the original instance.

Recall that a linear optimization oracle of the contextual bandit instance states that $\arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ can be efficiently solved for any context \mathcal{A}_t . However, in the reduced bandit instance, we are solving an optimization problem related to the reduced arms, $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$. In Lemma 16 in App. E, we provide an efficient (approximate) optimization oracle for solving $\arg \max_{x \in \mathcal{X}_m} \langle x, \theta \rangle$ by applying a careful discretization of the set Θ . By storing $g^{(m)}(\theta)$ for all the arms pulled by the linear bandit algorithm, we can then apply the (efficient) regression oracle for the original contextual bandit instance to get a regression oracle in the reduced instance.

If all actions played by the linear bandit algorithm are the output of the linear optimization oracle for some θ (Awerbuch and Kleinberg, 2008; Dani et al., 2008; Ito et al., 2019), then from Lemma 16 in App. E, the inversion of $g^{(m)}(\theta)$ for the actions pulled by the linear bandit algorithm can be performed by storing θ whenever the corresponding action $g^{(m)}(\theta)$ is stored. This increases both the space and time complexity only by a constant factor.

4. Conclusions

We presented a novel reduction from stochastic contextual linear bandit problems to (fixed-context) linear bandit problems; our reduction explains why results, that are not achievable for adversarial contexts, are possible for stochastic contextual bandits, and offers a framework that can be leveraged to gain new algorithms and bounds for contextual linear bandit problems. We illustrate the power of our approach by applying it to achieve improved bounds over a number of cases; this is not an exhaustive list, and we hope that our approach will be a useful tool to researchers in this field.

Acknowledgments

This work was partially supported by the NSF grants #2007714 and #2221871, the DARPA grant #HR00112190130, and by the Army Research Laboratory under the Cooperative Agreement W911NF-17-2-0196.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2012.
- Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.
- Ilija Bogunovic, Andreas Krause, and Jonathan Scarlett. Corruption-tolerant gaussian process bandit optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 1071–1081. PMLR, 2020.
- Ilija Bogunovic, Arpan Losalka, Andreas Krause, and Jonathan Scarlett. Stochastic linear bandits robust to adversarial attacks. In *International Conference on Artificial Intelligence and Statistics*, pages 991–999. PMLR, 2021.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, and Sham M Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *Conference on Learning Theory*, pages 41–1. JMLR Workshop and Conference Proceedings, 2012.
- Alexandra Carpentier and Rémi Munos. Bandit theory meets compressed sensing for high dimensional stochastic linear bandit. In *Artificial Intelligence and Statistics*, pages 190–198. PMLR, 2012.
- Yi Chen, Yining Wang, Ethan X Fang, Zhaoran Wang, and Runze Li. Nearly dimension-independent sparse linear bandit over small action spaces via best subset selection. *Journal of the American Statistical Association*, (just-accepted):1–31, 2022.
- Yifang Chen, Simon Du, and Kevin Jamieson. Improved corruption robust algorithms for episodic reinforcement learning. In *International Conference on Machine Learning*, pages 1561–1570. PMLR, 2021.
- Koby Crammer and Claudio Gentile. Multiclass classification with bandit feedback using adaptive regularization. *Machine learning*, 90(3):347–383, 2013.
- Yan Dai, Ruosong Wang, and Simon S Du. Variance-aware sparse linear bandits. *arXiv preprint arXiv:2205.13450*, 2022.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. 2008.
- Simon S Du, Sham M Kakade, Ruosong Wang, and Lin F Yang. Is a good representation sufficient for sample efficient reinforcement learning? *arXiv preprint arXiv:1910.03016*, 2019.

- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in contextual bandits. *Advances in Neural Information Processing Systems*, 33:11478–11489, 2020.
- Zijun Gao, Yanjun Han, Zhimei Ren, and Zhengqing Zhou. Batched multi-armed bandits problem. *Advances in Neural Information Processing Systems*, 32, 2019.
- Sébastien Gerchinovitz. Sparsity regret bounds for individual sequences in online linear regression. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 377–396. JMLR Workshop and Conference Proceedings, 2011.
- Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. Misspecified linear bandits. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Yanjun Han, Zhengqing Zhou, Zhengyuan Zhou, Jose Blanchet, Peter W Glynn, and Yinyu Ye. Sequential batch learning in finite-action linear contextual bandits. *arXiv preprint arXiv:2004.06321*, 2020.
- Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. *Advances in Neural Information Processing Systems*, 33:10753–10763, 2020.
- Botao Hao, Tor Lattimore, and Wei Deng. Information directed sampling for sparse linear bandits. *Advances in Neural Information Processing Systems*, 34:16738–16750, 2021.
- Elad Hazan and Zohar Karnin. Volumetric spanners: an efficient exploration basis for learning. *Journal of Machine Learning Research*, 2016.
- Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. Nearly optimal algorithms for linear contextual bandits with adversarial corruptions. *arXiv preprint arXiv:2205.06811*, 2022.
- Shinji Ito, Daisuke Hatano, Hanna Sumita, Kei Takemura, Takuro Fukunaga, Naonori Kakimura, and Ken-Ichi Kawarabayashi. Oracle-efficient algorithms for online linear optimization with bandit feedback. *Advances in Neural Information Processing Systems*, 32, 2019.
- Kyoungseok Jang, Chicheng Zhang, and Kwang-Sung Jun. Popart: Efficient sparse regression and experimental design for optimal sparse linear bandits. *arXiv preprint arXiv:2210.15345*, 2022.
- Tiancheng Jin, Longbo Huang, and Haipeng Luo. The best of both worlds: stochastic and adversarial episodic mdps with unknown transition. *Advances in Neural Information Processing Systems*, 34: 20491–20502, 2021.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Tor Lattimore, Koby Crammer, and Csaba Szepesvári. Linear multi-resource allocation with semi-bandit feedback. *Advances in Neural Information Processing Systems*, 28, 2015.

- 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.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, Mengxiao Zhang, and Xiaojin Zhang. Achieving near instance-optimality and minimax-optimality in stochastic and adversarial linear bandits simultaneously. In *International Conference on Machine Learning*, pages 6142–6151. PMLR, 2021.
- Yingkai Li, Edmund Y Lou, and Liren Shan. Stochastic linear optimization with adversarial corruption. *arXiv preprint arXiv:1909.02109*, 2019a.
- Yingkai Li, Yining Wang, and Yuan Zhou. Nearly minimax-optimal regret for linearly parameterized bandits. In *Conference on Learning Theory*, pages 2173–2174. PMLR, 2019b.
- Yingkai Li, Yining Wang, Xi Chen, and Yuan Zhou. Tight regret bounds for infinite-armed linear contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 370–378. PMLR, 2021.
- Thodoris Lykouris, Max Simchowitz, Alex Slivkins, and Wen Sun. Corruption-robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, pages 3242–3245. PMLR, 2021.
- Aldo Pacchiano, My Phan, Yasin Abbasi Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. *Advances in Neural Information Processing Systems*, 33:10328–10337, 2020.
- Vianney Perchet, Philippe Rigollet, Sylvain Chassang, and Erik Snowberg. Batched bandit problems. *The Annals of Statistics*, pages 660–681, 2016.
- Yufei Ruan, Jiaqi Yang, and Yuan Zhou. Linear bandits with limited adaptivity and learning distributional optimal design. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 74–87, 2021.
- 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.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Chen-Yu Wei, Christoph Dann, and Julian Zimmert. A model selection approach for corruption robust reinforcement learning. In *International Conference on Algorithmic Learning Theory*, pages 1043–1096. PMLR, 2022.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.
- Zihan Zhang, Xiangyang Ji, and Yuan Zhou. Almost optimal batch-regret tradeoff for batch linear contextual bandits. *arXiv preprint arXiv:2110.08057*, 2021.

Yinglun Zhu, Dylan J Foster, John Langford, and Paul Mineiro. Contextual bandits with large action spaces: Made practical. In *International Conference on Machine Learning*, pages 27428–27453. PMLR, 2022.

Appendix A. Reduction When Context is a Product Set

Algorithm 2 Reduction from stochastic contexts to no context for product context sets

Input: an algorithm Λ for linear contextual bandits with action set

$$\mathcal{X} = \{a' \in \{0, 1\}^{2d} \mid a'_{2i-1} + a'_{2i} = 1 \forall i \in [d]\} \quad (11)$$

for $t = 1 : T$ **do**

 Ask Λ for an arm to play $a'_t \in \mathcal{X}$ given the history r_1, \dots, r_{t-1} .

 Play arm a_t with $(a_t)_i = \max\{a \in \mathcal{A}_t^{(i)}\}$ if $a'_{2i-1} = 1$ and $(a_t)_i = \min\{a \in \mathcal{A}_t^{(i)}\}$ otherwise.

 Receive reward r_t .

end

We here consider the special case where the action distribution \mathcal{D} is unknown, but the action space has a coordinate-wise product structure, i.e.,

Assumption 1 $\mathcal{A}_t = \prod_{i=1}^d \mathcal{A}_t^{(i)}$, where $\mathcal{A}_t^{(i)} \subset \mathbb{R}$.

This is an important hard instance that appears in many lower bounds (Lattimore and Szepesvári, 2020). We will show that in this case, the d -dimensional stochastic contextual bandits can be reduced to a linear bandit problem with no misspecification, although the distribution \mathcal{D} is unknown, but where now the parameter vector is over $2d$ dimensions.

The main idea of the reduction is that $g(\theta)$ can be factored into a coordinate-wise product between an unknown vector that only depends on the context distribution and a known vector that only depends on θ . The unknown vector can be incorporated with θ_* reducing the contextual instance to a linear bandit instance but now with a different unknown parameter θ'_* . In particular, we can write $\langle g(\theta), \theta_* \rangle = \langle a'(\theta), \theta'_* \rangle$, where $a'(\theta)$ is a vector in \mathbb{R}^{2d} that does not depend on the distribution, and equals

$$(a'(\theta))_{2i} = \begin{cases} 1 & \text{if } (\theta)_i < 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and } (a'(\theta))_{2i-1} = 1 - (a'(\theta))_{2i} \quad \forall i = 1, \dots, d. \quad (12)$$

Thus, we can follow the same reduction algorithm as in Section 3.1, but where now we call a $2d$ -dimensional linear bandit algorithm Λ and provide Λ with the fixed action set

$$\mathcal{X} = \{a' \in \{0, 1\}^{2d} \mid a'_{2i-1} + a'_{2i} = 1 \forall i \in [d]\} \quad (13)$$

as the pseudocode Algorithm 2 describes.

Theorem 4 *Let Λ be any algorithm for linear bandits, I be a contextual linear bandit instance with stochastic contexts that satisfy Assumption 1 with unknown parameter θ_* , and r_t be the rewards generated as described in Algorithm 2. It holds that*

- The rewards r_t are generated, by pulling arm a'_t , from a linear bandit instance L with action set \mathcal{X} in (11), and unknown parameter $\theta'_* \in \mathbb{R}^{2d}$ with $\|\theta'_*\|_2 \leq 2$.
- With probability at least $1 - \delta$ it holds that

$$|R_T(I) - R_T^\Lambda(L)| \leq c\sqrt{T \log(1/\delta)}, \quad (14)$$

where $R_T(I)$ is the regret of Algorithm 2 over the instance I and $R_T^\Lambda(L)$ is the regret of algorithm Λ over the instance L , and c is a universal constant.

Proof The proof follows from Theorem 1 using the observation that

$$(g(\theta))_i = \begin{cases} \mathbb{E}_{\mathcal{A} \sim \mathcal{D}}[\max_{a \in \mathcal{A}}(a)_i] & \text{if } (\theta)_i > 0 \\ \mathbb{E}_{\mathcal{A} \sim \mathcal{D}}[\min_{a \in \mathcal{A}}(a)_i] & \text{if } (\theta)_i < 0 \end{cases}. \quad (15)$$

and thus, given (13), we can define $\theta'_\star \in \mathbb{R}^{2d}$ as

$$(\theta'_\star)_i = \begin{cases} \mathbb{E}[\max_{a \in \mathcal{A}}(a)_{(i+1)/2}](\theta_\star)_{(i+1)/2} & \text{if } i \text{ is odd} \\ \mathbb{E}[\min_{a \in \mathcal{A}}(a)_{\lfloor (i+1)/2 \rfloor}](\theta_\star)_{\lfloor (i+1)/2 \rfloor} & \text{if } i \text{ is even.} \end{cases} \quad (16)$$

By definition of $a'(\theta), \theta'_\star$, we have that $\langle g(\theta), \theta_\star \rangle = \langle a'(\theta), \theta'_\star \rangle$. To see that $\|\theta'_\star\| \leq 2$, we observe that

$$\sum_{i=1}^d |(\theta'_\star)_{2i-1}| = \sum_{i=1}^d |\mathbb{E}[\max_{a \in \mathcal{A}}(a)_i](\theta_\star)_i| \leq \langle |g(\mathbf{1}/\sqrt{d})|, |\theta_\star| \rangle \leq 1, \quad (17)$$

where absolute value of a vector is defined as a vector with $(|\theta_\star|)_i = |(\theta_\star)_i|$. Similarly, $\sum_{i=1}^d |(\theta'_\star)_{2i}| \leq 1$. Hence, $\|\theta'_\star\|_2 \leq \|\theta'_\star\|_1 \leq 2$.

As before, to construct the linear bandit instance, we need rewards that follow the stochastic linear bandits model. The result follows from Proposition 12 and the fact that $\langle g(\theta_t), \theta_\star \rangle = \langle a_t, \theta'_\star \rangle$. ■

Appendix B. Implications

B.1. Tighter Regret Bound for Contextual Linear Bandits

We here show that our reduction leads to the first $O(d\sqrt{T \log T})$ high probability upper bound for linear bandits with changing action sets. We recall that in all our results, we only consider contextual linear bandits with unknown context distribution as described in Section 2. We rely on the Phased Elimination (PE) (Lattimore et al., 2020) as our linear bandit algorithm Λ_ϵ . PE is the same as Algorithm 3, that we will use next in the batched setting, except that \mathcal{X}, g are fixed (recall that we apply PE within an epoch that fixes the estimate of g). The parameter γ_m in Algorithm 3 is called the confidence interval; we will specify its value in our theorems. In Algorithm 1, setting $t^{(m)} = 2^{m-1}$ and Λ_ϵ to be PE with modified confidence intervals (Lattimore et al., 2020) to account for the misspecification we get the following corollary.

Corollary 5 *For Algorithm 1 with $t^{(m)} = 2^{m-1}$ and Λ_{ϵ_m} to be PE with $\gamma_m = 6\sqrt{d \log(T|\Theta'|/\delta)/t^{(m)}}$ it holds that with probability at least $1 - c\delta$ we have that $R_T \leq c\sqrt{dT \log(T|\Theta'|/\delta)}$.*

Proof Let the length of phase m be $T_m = t^{(m+1)} - t^{(m)} = 2^m - 2^{m-1}$. Conditioned on the event that the misspecification in phase m is bounded by ϵ_m , PE with modified confidence intervals achieves a regret $R_{T_m}^{\Lambda_{\epsilon_m}}$ that is upper bounded by $O(\sqrt{dT_m \log(T_m M |\Theta'|/\delta)}) + \sqrt{dT_m \epsilon_m}$ with probability at least $1 - \delta/M$. Hence, by Proposition 14 and the union bound we have that it holds with probability at least $1 - c\delta$ that

$$R_{T_m}^{\Lambda_{\epsilon_m}} \leq c(\sqrt{dT_m \log(T|\Theta'|/\delta)} + \sqrt{dT_m \epsilon_m}) \forall m \in [M], \quad (18)$$

where we used the fact that $T_m \leq T, M \leq T$. Substituting the value of ϵ_m we get that the following holds with probability at least $1 - c\delta$

$$\begin{aligned} R_{T_m}^{\Lambda_{\epsilon_m}} &\leq c(\sqrt{dT_m \log(T|\Theta'|/\delta)} + \sqrt{dT_m} \sqrt{\frac{\log(M|\Theta'|/\delta)}{t^{(m)}}}) \\ &= c(\sqrt{dT_m \log(T|\Theta'|/\delta)} + \sqrt{dT_m \log(M|\Theta'|/\delta)}) \forall m \in [M]. \end{aligned} \quad (19)$$

Hence, using $M \leq T$, we get that with probability at least $1 - c\delta$ we have

$$R_{T_m}^{\Lambda_{\epsilon_m}} \leq c\sqrt{dT_m \log(T|\Theta'|/\delta)} \quad (20)$$

Substituting in Theorem 2, we get that with probability at least $1 - c\delta$ it holds that

$$\begin{aligned} R_T &\leq c\sqrt{T \log T} + c\sqrt{d \log(T|\Theta'|/\delta)} \sum_{m=1}^{\log T} \sqrt{T_m} \\ &\leq c'\sqrt{dT \log(T|\Theta'|/\delta)} \sum_{m=1}^{\log T} \sqrt{2^{m-\log T}} \\ &\leq c'\sqrt{dT \log(T|\Theta'|/\delta)} \sum_{i=0}^{\infty} \sqrt{2^{-i}} \\ &\leq \frac{c'}{1 - 1/\sqrt{2}} \sqrt{dT \log(T|\Theta'|/\delta)}. \end{aligned} \quad (21)$$

■

It is well known that if $\Theta \subseteq \{a \in \mathbb{R}^d \mid \|a\|_2 \leq 1\}$, then there is $1/T$ -net of Θ , Θ' , such that $|\Theta'| \leq (6T)^d$, $\Theta' \subseteq \Theta$. This directly implies that $R_T = O(d\sqrt{T \log(T)})$ with probability at least $1 - 1/T$. This improves a factor of $\sqrt{\log T} \log \log T$ over (Li et al., 2021) and a factor of $\sqrt{\log T}$ over (Abbasi-Yadkori et al., 2011).

B.2. Batch Learning

WE here show that our reduction can be applied to improve the result of (Ruan et al., 2021) for contextual linear bandits with stochastic contexts by providing a regret upper bound of $O(d\sqrt{T \log T} \log \log T)$ with high probability as opposed to the *in expectation*

$O(d\sqrt{T \log d \log T} \log \log T)$ regret bound in (Ruan et al., 2021). This can be achieved by replacing Λ_{ϵ_m} with the G -optimal design policy constructed using \mathcal{X}_m . To compute $t^{(i)}$ we first define

$$u_m = T^{1-2^{-m}}, m = 1, \dots, M/2. \quad (22)$$

We let

$$t^{(m)} = \lfloor u_{m//2+1} \rfloor \forall m \in [M], t^{(M+1)} = T. \quad (23)$$

where $//$ denotes integer division. For completeness, we include the pseudo-code in Algorithm 3. The following result follows using Theorem 2.

Theorem 6 For Algorithm 1 with $M = 2 \log \log T$ (corresponds to $2 \log \log T + 1$ batches), $t^{(m)}$ given in (23) and Λ_{ϵ_m} replaced by the G-optimal design policy constructed using \mathcal{X}_m , it holds that there is a constant c such that with probability at least $1 - c\delta$, we have

$$R_T \leq c \sqrt{dT \log\left(\frac{M|\Theta'|}{\delta}\right) \log \log T}. \quad (24)$$

Proof We first notice that for $M = \log \log T$, we have that $u_{m+1}/\sqrt{u_m} = \sqrt{T} \forall m \in [M-1]$, $u_{M+1}/\sqrt{u_M} = \sqrt{2T}$. Hence, we have that

$$u_{m+1}/\sqrt{u_m} \leq \sqrt{2T} \forall m \in [M]. \quad (25)$$

We also have that $u_{m+1} \geq u_m \forall m \geq 1$ (since $u_m \leq T$). The proof follows from the properties of the G-optimal design together with the properties of $g^{(m)}$ in the proof of Theorem 2. The G-optimal design ensures that for any $\theta \in \Theta'$ the following holds with probability at least $1 - \delta$

$$|\langle g^{(m)}(\theta), \hat{\theta}_m - \theta_* \rangle| \leq 2\epsilon'_m \sqrt{d} + \sqrt{\frac{4d}{T_m} \log 1/\delta}, \quad (26)$$

where $\epsilon'_m = \sup_{\theta \in \Theta'} |\langle g^{(m)}(\theta) - g(\theta), \theta_* \rangle|$. By the triangle inequality, we have that

$$\begin{aligned} |\langle g^{(m)}(\theta), \hat{\theta}_m \rangle - \langle g(\theta), \theta_* \rangle| &\leq 2\epsilon'_m \sqrt{d} + \sqrt{\frac{4d}{T_m} \log 1/\delta} + |\langle g^{(m)}(\theta) - g(\theta), \theta_* \rangle| \\ &\leq \epsilon'_m (2\sqrt{d} + 1) + \sqrt{\frac{4d}{T_m} \log 1/\delta} \end{aligned} \quad (27)$$

By Proposition 14 we have that $\epsilon'_m \leq \epsilon_m \forall m \in [M]$ with probability at least $1 - \delta$. Hence, by the union bound we have that the following holds with probability at least $1 - \delta$

$$|\langle g^{(m)}(\theta), \hat{\theta}_m \rangle - \langle g(\theta), \theta_* \rangle| \leq \epsilon_m (2\sqrt{d} + 1) + \sqrt{\frac{4d}{T_m} \log\left(\frac{M|\Theta'|}{\delta}\right)} \forall \theta \in \Theta' \forall m \in [M]. \quad (28)$$

Hence, with probability at least $1 - \delta$, the best arm is not eliminated and the arms that are not eliminated at the end of batch m , will have a gap that is at most twice the confidence interval in (28), otherwise, they satisfy the elimination criterion with θ_* . Hence, the sum regret $\sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$ is bounded as follows with probability at least $1 - \delta$

$$\begin{aligned} \sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m) &\leq \sqrt{T} + c' \sum_{m=1}^M \epsilon_m (2\sqrt{d} + 1) T_{m+1} + c' \sqrt{\frac{4d}{T_m} \log\left(\frac{M|\Theta'|}{\delta}\right)} T_{m+1} \\ &= \sqrt{T} + c \sqrt{d \log\left(\frac{M|\Theta'|}{\delta}\right)} \sum_{m=1}^M \frac{T_{m+1}}{\sqrt{T_{m-1}}} + \frac{T_{m+1}}{\sqrt{T_m}} \\ &\stackrel{(i)}{\leq} \sqrt{T} + 2c \sqrt{d \log\left(\frac{M|\Theta'|}{\delta}\right)} \sum_{m=1}^M \frac{T_{m+1}}{\sqrt{T_{m-1}}} + \sqrt{T} \\ &\stackrel{(ii)}{\leq} \sqrt{T} + 2c \sqrt{d \log\left(\frac{M|\Theta'|}{\delta}\right)} \sum_{m=1}^M \sqrt{T} + \sqrt{T} \end{aligned}$$

$$\leq c'' \sqrt{dT \log\left(\frac{M|\Theta'|}{\delta}\right)} \log \log T, \quad (29)$$

where (i) uses (25) and the fact that either $T_{m+1} \leq \sqrt{2T_m}$ or $T_{m+1} = T_m \leq \sqrt{2TT_m}$, and (ii) follows from the fact that either ($T_m = T_{m-1}$ and $T_{m+1} \leq \sqrt{2TT_m}$) or ($T_m \leq \sqrt{2TT_{m-1}}$ and $T_{m+1} = T_m$), hence, in both cases we have $T_{m+1} \leq \sqrt{2TT_{m-1}}$. The results follow by Theorem 2 and the union bound. \blacksquare

As in App. B.1, this implies that $R_T = O(d\sqrt{T \log(T)} \log \log T)$ with probability at least $1 - 1/T$, which strengthens the result in (Ruan et al., 2021) to high probability result and also improves a factor of $\log d$ in the regret bound. Moreover, as will be seen in the following appendices, if θ_* is s -sparse, Theorem 6 implies a regret bound of $O(\sqrt{dsT \log T} \log \log T)$ with probability at least $1 - 1/T$. To the best of our knowledge, this is the first nearly optimal algorithm for sparse unknown parameters that uses $O(\log \log T)$ batches.

Algorithm 3 Batched algorithm for linear bandits with stochastic context

Input: confidence parameter δ , and phase lengths $\{t^{(m)}\}_{m=1}^M$.

Initialize: $g^{(1)} : \Theta' \rightarrow \mathbb{R}^d$ randomly, $\Theta_1 = \Theta'$, and let $\mathcal{X}_1 = \{g^{(1)}(\theta) | \theta \in \Theta'\}$.

for $m = 1 : M$ **do**

Find design $\rho : \mathcal{X}_m \rightarrow [0, 1]$ with $\max_{a \in \text{Supp}(\rho)} \|a\|_{G^{-1}(\rho)}^2 \leq 2d$, $|\text{Supp}(\rho)| \leq 4d \log \log d + 16$, where $G(\rho) = \sum_{a \in \text{Supp}(\rho)} \rho(a) aa^T$.

Compute $u(x) = \lceil \rho(x) T_m \rceil$ and $u = \sum_{x \in \text{Supp}(\rho)} u(x)$.

Use the policy described by $a_t = \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ for $u(g^{(m)}(\theta))$ times for each $\theta \in \Theta_m$ with $g^{(m)}(\theta) \in \text{Supp}(\rho)$.

$\hat{\theta}_m = (\sum_{g \in \text{Supp}(\rho)} u(g) gg^T)^{-1} \sum_{i=1}^u r_i g^{(m)}(\theta_i)$.

Update: $\Theta_{m+1} = \left\{ \theta \in \Theta_m \mid \max_{\theta' \in \Theta_m} \langle \hat{\theta}_m, g^{(m)}(\theta') - g^{(m)}(\theta) \rangle \leq \gamma_m \right\}$, $\gamma_m = 10 \sqrt{\frac{d}{T_{m-1}} \log(M|\Theta'|/\delta)}$.

Update: $g^{(m+1)}(\theta) = \frac{1}{t^{(m+1)}} \sum_{t=1}^{t^{(m+1)}} \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle$, and $\mathcal{X}_{m+1} = \{g^{(m+1)}(\theta) | \theta \in \Theta'\}$.

end

B.3. Misspecified Contextual Linear Bandits

Our reduction framework can be used to improve regret bounds for misspecified contextual linear bandits. By noticing that for ϵ misspecified contextual linear bandits, the total amount of misspecification in epoch m is bounded by $\epsilon_m + \epsilon$ with high probability, the following result follows directly from (Lattimore et al., 2020).

Theorem 7 For contextual linear bandits with known misspecification that is bounded by ϵ , Algorithm 1 with Λ_{ϵ_m} being PE with $\gamma_m = 6d\sqrt{\log(T)/T_m} + \epsilon\sqrt{d}$ achieves a regret bound

$$R_T \leq c(d\sqrt{T \log(T/\delta)} + \epsilon\sqrt{dT}) \quad (30)$$

with probability at least $1 - c\delta$.

If ϵ is unknown, (Lattimore et al., 2020) showed that PE achieves a regret upper bounded by $O(d\sqrt{T \log(T)} + \epsilon\sqrt{dT} \log T)$ with high probability. We notice that the dependency on ϵ has an extra $\log T$ factor as compared to the known misspecification result. To avoid adding the $\log T$ factor in the term containing ϵ_m , we exploit the knowledge of ϵ_m to modify the confidence intervals in PE. We note that following the same analysis in (Lattimore et al., 2020) by replacing ϵ with $\epsilon + \epsilon_m$ results in the same concentration of estimated means except for constant factors, which implies the following result.

Theorem 8 *For contextual linear bandits with unknown misspecification that is bounded by ϵ , Algorithm 1 with Λ_{ϵ_m} being PE with $\gamma_m = 6d\sqrt{\log(T)}/T_m$ achieves a regret bound*

$$R_T \leq c(d\sqrt{T \log(T/\delta)} + \epsilon\sqrt{dT} \log T) \quad (31)$$

with probability at least $1 - c\delta$.

Proof The proof follows similar steps as in the proof of Proposition 5.1 in (Lattimore et al., 2020) with different constants due to the change in γ_m . \blacksquare

We note that as PE is performed in each epoch, the number of batches is $\Omega(\log^2 T)$. However, we constructed the algorithm this way only for simplicity. It is possible to perform PE once and update the estimates of g at the end of each batch similar to Algorithm 3.

B.4. Bandits With Adversarial Corruption

Our reduction directly extends the results for linear bandits with adversarial corruption to the contextual setting while maintaining the same regret bounds up to log factors. This leads to $\tilde{O}(d\sqrt{T} + d^{1.5}C)$ high probability regret bound as opposed to the best known $\tilde{O}(d^{4.5}\sqrt{T} + d^4C)$ regret bound. It was shown in (Bogunovic et al., 2021) that PE with modified confidence intervals achieves an $\tilde{O}(d\sqrt{T} + d^{1.5}C)$ high probability regret bound for unknown corruption C . A model selection based approach in (Foster et al., 2020) that uses PE as a subroutine is shown to generalize this result for unknown C without changing the regret bound except for log factors. By adapting the confidence intervals of PE to account for the known ϵ_m misspecification the $\tilde{O}(d\sqrt{T} + d^{1.5}C)$ high probability regret bound extends for a model with $O(\sqrt{d \log T/\bar{T}})$ misspecification and known corruption C . We will modify the PE confidence interval to be

$$\gamma_m = 8d\sqrt{\log(T)/t^{(m)}} + \frac{2C(4d \log \log d + 18)}{T_m} \sqrt{8d}. \quad (32)$$

The model selection approach in (Foster et al., 2020) implies the following result.

Theorem 9 *For contextual linear bandits with unknown corruption that is bounded by C , Algorithm 1 with $t^{(m)} = 2^{m-1}$ and Λ_{ϵ_m} being G-COBE in (Foster et al., 2020) with PE, that uses γ_m in (32), as a subroutine, achieves*

$$R_T = \tilde{O}(d\sqrt{T} + d^{1.5}C) \quad (33)$$

with probability at least $1 - c/T$.

Proof The proof follows by proving a regret bound for PE with the chosen γ_m for known C . Then use the model selection result in (Foster et al., 2020) to prove a regret bound for unknown C and finally apply our main theorem to bound the final regret.

We notice that

$$|\langle g^{(m)}(\theta), \hat{\theta}_m \rangle - \langle g(\theta), \theta_\star \rangle| \leq |\langle g^{(m)}(\theta), \hat{\theta}_m - \theta_\star \rangle| + |\langle g^{(m)}(\theta) - g(\theta), \theta_\star \rangle|. \quad (34)$$

Using Lemma 1 in (Bogunovic et al., 2021) and equation (26) to bound the first term and Proposition 14 to bound the second term, we obtain that $|\langle g^{(m)}(\theta), \hat{\theta}_m \rangle - \langle g(\theta), \theta_\star \rangle| \leq \gamma_m/2$ with probability at least $1 - c/T$. A standard calculation in Theorem 1 in (Bogunovic et al., 2021) implies a regret bound of $\tilde{O}(d\sqrt{T} + d^{1.5}C)$ with high probability. Hence, Theorem 4 in (Wei et al., 2022) implies a regret bound of $\tilde{O}(d\sqrt{T} + d^{1.5}C)$ with high probability for G-COBE with unknown C . Applying Theorem 2 and proceeding as in equation (21) concludes the proof. \blacksquare

B.5. Sparsity

For linear contextual bandits with s -sparse unknown parameter, our reduction can be used to prove $O(\sqrt{dsT \log(T)})$ regret bound with high probability as opposed to the best known $O(\sqrt{dsT \log^2 T})$ regret bound. It is not hard to show that there is $1/T$ cover with size at most $(6T)^{2s+1}$ in that case; also proved below for completeness. The following result directly follows from Corollary 5.

Theorem 10 *For Algorithm 1 with $t^{(m)} = 2^{m-1}$ and Λ_{ϵ_m} to be PE with $\gamma_m = 6\sqrt{2ds \log(T/\delta)/t^{(m)}}$ it holds that with probability at least $1 - c\delta$ we have that $R_T = O(\sqrt{dsT \log(T/\delta)})$.*

Proof By Corollary 5, we only need to show that Θ contains an $1/T$ -net with size at most $(6T)^{2s+1}$. We have that there is $1/2T$ -net for $\{\phi \in \mathbb{R}^s \mid \|\phi\|_2 \leq 1\}$ with size at most $(6T)^s$. To construct s -sparse vectors in \mathbb{R}^d , there is $\sum_{i=1}^s \binom{d}{i}$ ways to choose the non-zero entries. This implies that the set $\{\theta \in \mathbb{R}^d \mid \|\theta\|_2 \leq 1, \|\theta\|_0 \leq s\}$ has $1/2T$ -net \mathcal{N} of size at most

$$|\mathcal{N}| \leq (6T)^s \sum_{i=1}^s \binom{d}{i} \leq (6T)^s \sum_{i=1}^s d^i \leq (6T)^s d^{s+1} \leq (6T)^{2s+1}. \quad (35)$$

We next construct an $1/T$ -net that is subset of Θ (recall that this is required for Theorem 2). For every $x \in \mathcal{N}$, let $\mathcal{N}_x = \{\theta \in \Theta \mid \|\theta - x\| \leq 1/2T\}$. Let $\alpha : \mathcal{N} \rightarrow |\mathcal{N}|$ be an ordering of the set \mathcal{N} . And let $\mathcal{N}' \subseteq \mathcal{N}$ be defined as $\mathcal{N}' = \{x \in \mathcal{N} \mid \mathcal{N}_x \not\subseteq \cup_{y: \alpha(y) \leq \alpha(x)} \mathcal{N}_y\}$. Hence, $\{\mathcal{N}_x \mid x \in \mathcal{N}'\}$ is a set of pairwise disjoint sets. By the axiom of choice, there is a set \mathcal{N}'' such that \mathcal{N}'' contains exactly one element of each set $\mathcal{N}_x \forall x \in \mathcal{N}'$. By construction of \mathcal{N}'' we have that $\mathcal{N}'' \subseteq \Theta$.

We also have that for each $\theta \in \Theta$, there is $x \in \mathcal{N}$ with $\|x - \theta\|_2 \leq 1/2T$. Then for each $y \in \mathcal{N}_x$, we have by the triangle inequality that $\|y - \theta\|_2 \leq 1/T$. By the definition of \mathcal{N}' , we have that $\cup \mathcal{N}' = \cup_{z \in \mathcal{N}} \mathcal{N}_z$. Then, by construction of \mathcal{N}'' , there is $z \in \mathcal{N}''$ such that $z \in \mathcal{N}_x$. Hence, $\|z - \theta\|_2 \leq 1/T$. This implies that $\mathcal{N}'' \subseteq \Theta$ is an $1/T$ -net of Θ . We also have by construction that also that $|\mathcal{N}''| \leq |\mathcal{N}| \leq (6T)^{2s+1}$. \blacksquare

B.6. Structured Unknown Parameters

In some cases, the dimension d can be large but the unknown parameter is mapped from a small space of dimension s . In particular let $f : \mathbb{R}^s \rightarrow \mathbb{R}^d$ be such that $\Theta \subseteq \{f(\phi) \mid \|\phi\|_2 \leq 1\}$ and

$$\|f(\phi) - f(\phi')\|_2 \leq c\|\phi - \phi'\|_2, \quad (36)$$

where c is a universal constant. Finding an $1/T$ -net for $\Phi = \{\phi \mid f(\phi) \in \Theta\}$ implies a c/T -net for Θ with the same size. Since $\Phi \subseteq \{\phi \in \mathbb{R}^s \mid \|\phi\|_2 \leq 1\}$, we have that there is an $1/T$ -net for Φ that is contained in Φ with size at most $(6T)^s$. The following result immediately follows from Corollary 5.

Corollary 11 *Under the considered structured unknown parameters assumption, Algorithm 1 with $t^{(m)} = 2^{m-1}$ and Λ_{ϵ_m} to be PE with $\gamma_m = 6\sqrt{ds \log(T/\delta)}/t^{(m)}$ satisfies that with probability at least $1 - c\delta$ we have that $R_T = O(\sqrt{dsT \log(T/\delta)})$.*

Appendix C. Proof of Theorem 1

Theorem [Restatement of Theorem 1] *Let Λ be any algorithm for linear bandits and I be a contextual linear bandit instance with stochastic contexts, unknown parameter θ_* and rewards r_t generated as described in the reduction in Section 3.1. It holds that*

- The reward r_t is generated, by pulling the arm $g(\theta_t)$, from a linear bandit instance L with action set \mathcal{X} , and unknown parameter θ_* .
- The reduction results in an algorithm \mathcal{M} for contextual linear bandits such that with probability at least $1 - \delta$ we have

$$|R_T^{\mathcal{M}}(I) - R_T^{\Lambda}(L)| \leq c\sqrt{T \log(1/\delta)}, \quad (37)$$

where $R_T^{\Lambda}(L)$ is the regret of Λ over the constructed linear bandit instance L , $R_T^{\mathcal{M}}(I)$ is the regret of algorithm \mathcal{M} over the instance I and c is a universal constant.

Proof Following the reduction described in the section, we start by showing that $R_T^{\Lambda}(L)$, the regret of the algorithm Λ on a linear bandit instance, is at most $\tilde{O}(\sqrt{T})$ away from $R_T^{\mathcal{M}}(I)$ with high probability. Recall that the regret $R_T^{\Lambda}(L)$ is defined as

$$R_T^{\Lambda}(L) = \sum_{t=1}^T \max_{\theta \in \Theta} \langle g(\theta), \theta_* \rangle - \langle g(\theta_t), \theta_* \rangle. \quad (38)$$

We notice that the function g depends on the context distribution \mathcal{D} . In the following we assume for simplicity that for each $\theta \in \Theta$, there is a unique $a_t \in \mathcal{A}_t$ that satisfies $\langle a_t, \theta \rangle = \sup_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ almost surely. We discuss how to remove this assumption at the end of the proof.

The regret $R_T^{\mathcal{M}}(I)$ of the contextual algorithm can be decomposed as

$$\begin{aligned} R_T^{\mathcal{M}}(I) &= \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle \\ &= \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle, \theta_* \rangle - \langle \mathbb{E}[\arg \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle], \theta_* \rangle \\ &\quad + \langle \mathbb{E}[\arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle | \theta_t], \theta_* \rangle - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_* \rangle \end{aligned}$$

$$\begin{aligned}
 & + \langle \mathbb{E}[\arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle], \theta_\star \rangle - \langle \mathbb{E}[\arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle | \theta_t], \theta_\star \rangle \\
 = & \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle] \\
 & + \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle \\
 & + \langle g(\theta_\star), \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle, \tag{39}
 \end{aligned}$$

where the expectation is with respect to the randomness in the context generation \mathcal{A}_t . In the following we will first show that $\langle g(\theta_\star), \theta_\star \rangle = \max_{\theta \in \Theta} \langle g(\theta), \theta_\star \rangle$.

Indeed, we observe that $\forall \theta', \theta'' \in \Theta$ we have

$$\begin{aligned}
 \max_{\theta \in \Theta} \langle g(\theta), \theta' \rangle & \geq \langle g(\theta'), \theta' \rangle = \mathbb{E}[\max_{a \in \mathcal{A}_t} \langle a, \theta' \rangle] \\
 & \geq \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta'' \rangle, \theta' \rangle] = \langle g(\theta''), \theta' \rangle. \tag{40}
 \end{aligned}$$

The above inequalities have to be met with equality since we can select $\theta'' = \arg \max_{\theta \in \Theta} \langle g(\theta), \theta' \rangle$ making the first and last terms equal. Hence, we have proved that

$$\langle g(\theta'), \theta' \rangle = \max_{\theta \in \Theta} \langle g(\theta), \theta' \rangle, \quad \forall \theta' \in \Theta. \tag{41}$$

Substituting in the last line of (39) using the triangle inequality, we get that

$$\begin{aligned}
 |R_T^{\mathcal{M}}(I) - R_T^{\Lambda}(L)| & \stackrel{(i)}{=} \left| \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle] \right. \\
 & \quad \left. + \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle \right| \\
 & \leq \left| \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle] \right| \\
 & \quad + \left| \sum_{t=1}^T \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle \right|, \tag{42}
 \end{aligned}$$

where (i) follows by definition of $R_T^{\Lambda}(L)$ and (41).

We next bound the quantity

$$\Sigma_T := \sum_{t=1}^T \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \theta_t] - \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle. \tag{43}$$

Let $\mathcal{F}'_t = \sigma\{\theta_1, r_1, \dots, \theta_t\}$ be the filtration of all historic information of the linear bandit problem up to time t . we notice that

$$\begin{aligned}
 \mathbb{E}[\Sigma_t | \mathcal{F}'_t] & = \mathbb{E}[\Sigma_{t-1} | \mathcal{F}'_t] + \mathbb{E}[\mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \theta_t] | \mathcal{F}'_t] - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_t \rangle, \theta_\star \rangle | \mathcal{F}'_t] \\
 & = \Sigma_{t-1}. \tag{44}
 \end{aligned}$$

Hence, Σ_t is a martingale with a bounded difference (by boundedness of Θ, \mathcal{A}_t). By Azuma–Hoeffding inequality (Wainwright, 2019), we have that $|\Sigma_T| \leq c\sqrt{T \log 1/\delta}$ with probability at least $1 - \frac{\delta}{2}$. For completeness, we state a special case of the Azuma-Hoeffding inequality at the end of our proof. The summation $\Sigma'_T = \sum_{t=1}^T \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle - \mathbb{E}[\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta_\star \rangle, \theta_\star \rangle]$ can be bounded similarly. Hence, by (42), we have that with probability at least $1 - \delta$

$$|R_T^M(I) - R_T^A(L)| \leq c\sqrt{T \log 1/\delta}. \quad (45)$$

We have shown that the regret of Algorithm 1 over the instance I is $O(\sqrt{T \log 1/\delta})$ away from $R_T^A(L)$ with probability at least $1 - \delta$. It remains to show that the rewards r_t generated by the described interaction with the instance I , are generated from L by the interaction of algorithm Λ .

Proposition 12 *The reward r_t can be rewritten as*

$$r_t = \langle g(\theta_t), \theta_\star \rangle + \eta'_t, \quad (46)$$

where $\mathbb{E}[\eta'_t | \mathcal{F}'_t] = 0$, $\mathbb{E}[\exp(\lambda \eta'_t) | \mathcal{F}'_t] \leq \exp(2\lambda^2) \quad \forall \lambda \in \mathbb{R}$, and where $\mathcal{F}'_t = \sigma\{\theta_1, r_1, \dots, \theta_t\}$ is the filtration of historic information up to time t .

Proof We have that

$$\begin{aligned} r_t &= \langle a_t, \theta_\star \rangle + \eta_t = \langle \arg \max_{b \in \mathcal{A}_t} \langle b, \theta_t \rangle, \theta_\star \rangle + \eta_t \\ &= \langle g(\theta_t), \theta_\star \rangle + \left(\eta_t + \langle \arg \max_{b \in \mathcal{A}_t} \langle b, \theta_t \rangle, \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle \right). \end{aligned} \quad (47)$$

We let $\eta'_t = \eta_t + \langle \arg \max_{b \in \mathcal{A}_t} \langle b, \theta_t \rangle, \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle$. The proof that $\mathbb{E}[\eta'_t | \mathcal{F}_t] = 0$ follows similarly to (44)

$$\begin{aligned} \mathbb{E}[\eta'_t | \mathcal{F}_t] &= \mathbb{E}[\eta_t | \mathcal{F}_t] + \mathbb{E}[\langle \arg \max_{b \in \mathcal{A}_t} \langle b, \theta_t \rangle, \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle | \mathcal{F}_t] \\ &= \mathbb{E}[\langle \arg \max_{b \in \mathcal{A}_t} \langle b, \theta_t \rangle, \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle | \theta_t] = 0. \end{aligned} \quad (48)$$

Lastly, $\mathbb{E}[\exp(\lambda \eta'_t) | \mathcal{F}_t] \leq \exp(2\lambda^2) \forall \lambda \in \mathbb{R}$ follows by boundedness of η'_t which follows by boundedness of $\eta_t, \Theta, \mathcal{A}_t$. \blacksquare

This concludes the proof.

Lemma 13 [Azuma's Inequality (Wainwright, 2019)] *Let $\Sigma_0, \Sigma_1, \dots$ be a martingale with respect to filtration $\mathcal{F}_0, \mathcal{F}_1, \dots$ such that $|\Sigma_i - \Sigma_{i-1}| \leq c$ almost surely. Then for all $\epsilon > 0$, we have that*

$$\mathbb{P}[|\Sigma_n - \Sigma_0| > \epsilon] \leq 2 \exp\left(-\frac{\epsilon^2}{2nc^2}\right). \quad (49)$$

\blacksquare

C.1. When $\sup_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ is not Unique

One solution is to choose a_t as any action $\langle a_t, \theta \rangle \geq \sup_{a \in \mathcal{A}_t} \langle a, \theta \rangle - \delta$ for some $\delta > 0$. The error arising from performing this step can be controlled by δ , e.g., by setting $\delta = \exp(-T)$. Our proofs will follow by choosing a_t as described above, using any deterministic or random choice function, as long as the action a_t is measurable.

Appendix D. Proof of Theorem 2

Theorem [Restatement of Theorem 2] Let Λ_ϵ be an algorithm for linear bandits with ϵ misspecification and I be a contextual linear bandit instance with stochastic contexts, unknown parameter θ_\star and rewards r_t are generated as described in Algorithm 1. The following holds:

- Conditioned on $\mathcal{H}_{t^{(m)}} = \sigma(\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_{t^{(m)}}, a_{t^{(m)}}, r_{t^{(m)}})$: in epoch m , the rewards r_t are generated, by pulling arm $g^{(m)}(\theta_t)$, from a misspecified linear bandit instance L_m for $t = t^{(m)} + 1, \dots, t^{(m+1)}$, action set $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$, unknown parameter θ_\star , mean rewards $\langle g(\theta), \theta_\star \rangle$, and unknown misspecification ϵ'_m .
- The misspecification ϵ'_m is bounded by the known quantity ϵ_m with probability at least $1 - c\delta/M$.
- With probability at least $1 - \delta$ we have that $|R_T(I) - \sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)| \leq c\sqrt{T \log(1/\delta)}$, where $R_T(I)$ is the regret of Algorithm 1 over the instance I , $R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$ is the regret of algorithm Λ_{ϵ_m} over the bandit instance L_m in phase m , $T_m = t^{(m+1)} - t^{(m)}$, and c is a universal constant.

Proof Let L_m be a bandit instance with actions $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$ indexed by the set Θ' , mean rewards $\langle g(\theta), \theta_\star \rangle \forall \theta \in \Theta'$, and $t \in \{t^{(m)} + 1, \dots, t^{(m+1)}\}$. Let

$$\mathcal{H}_{t^{(m)}} = \sigma\{\mathcal{A}_1, a_1, r_1, \dots, \mathcal{A}_{t^{(m)}}, a_{t^{(m)}}, r_{t^{(m)}}\}$$

be the filtration of all historic information before epoch m . Note that $g^{(m)}$ is defined in line 7 of Algorithm 1 and is the empirical estimate of g using history $\mathcal{H}_{t^{(m)}}$, hence, $g^{(m)}$ is $\mathcal{H}_{t^{(m)}}$ -predictable. Conditioned on $\mathcal{H}_{t^{(m)}}$, we have that L_m is a misspecified linear bandit instance with misspecification, $\epsilon'_m = \sup_{\theta \in \Theta'} \langle g(\theta) - g^{(m)}(\theta), \theta_\star \rangle$. And the regret of the algorithm Λ_{ϵ_m} over L_m is the random quantity

$$R_{T_m}^{\Lambda_{\epsilon_m}}(L_m) := \sum_{t=t^{(m)}+1}^{t^{(m+1)}} \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle.$$

As we have shown in Proposition 12, conditioned on $\mathcal{H}_{t^{(m)}}$, r_t is generated from L_m by pulling arm $g^{(m)}(\theta_t)$. Define the event $\mathcal{E}_m = \{\epsilon'_m > \epsilon_m\}$ be the bad event that the random quantity ϵ'_m is greater than ϵ_m defined in Algorithm 1. We will show in Proposition 14 that $\sum_{m=1}^M \mathbb{P}[\mathcal{E}_m] \leq \delta$.

We next bound the regret of Algorithm 1 in terms of the random quantities $\{R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)\}_{m=1}^M$. By choosing δ sufficiently small, it will be enough to bound $R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$ conditioned on $\mathcal{H}_{t^{(m)}}$ and the good event $\mathcal{G}_m = \{\epsilon'_m \leq \epsilon_m\}$. Let us define the random quantity $R_T^{\Lambda_\epsilon}(L_\epsilon) = \sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$. We show that $R_T^{\Lambda_\epsilon}(L_\epsilon)$ is at most $\tilde{O}(\sqrt{T})$ away from R_T^M with high probability. By definition of $R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$ we have that

$$R_T^{\Lambda_\epsilon}(L_\epsilon) = \sum_{t=1}^T \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle. \quad (50)$$

As in the proof of Theorem 1, we assume for simplicity that for each $\theta \in \Theta$, there is a unique $a_t \in \mathcal{A}_t$ that satisfies $\langle a_t, \theta \rangle = \sup_{a \in \mathcal{A}_t} \langle a, \theta \rangle$ almost surely. This can be relaxed in the same way as in Theorem 1.

Recall that L is the linear bandit instance in Theorem 1 with access to the function g , and actions in $\mathcal{X} = \{g(\theta) | \theta \in \Theta\}$. The regret $R_T^M(I)$ of the contextual algorithm can be bounded as

$$|R_T^M(I) - R_T^{\Lambda_\epsilon}(L_\epsilon)| \leq |R_T^M(I) - R_T^\Lambda(L)| + |R_T^\Lambda(L) - R_T^{\Lambda_\epsilon}(L_\epsilon)|, \quad (51)$$

where $R_T^\Lambda(L)$ is defined as follows

$$R_T^\Lambda(L) = \sum_{t=1}^T \max_{\theta \in \Theta} \langle g(\theta), \theta_\star \rangle - \langle g(\theta_t), \theta_\star \rangle,$$

and $\{\theta_t\}$ are the actions played by $\{\Lambda_{\epsilon_m}\}$. The first term in (51) is bounded in Theorem 1. In the following, we focus on bounding $|R_T^\Lambda(L) - R_T^{\Lambda_\epsilon}(L_\epsilon)|$. To that end, we first show that

$$|\langle g(\theta_\star), \theta_\star \rangle - \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle| \leq 2/T. \quad (52)$$

We recall from (41) that

$$\langle g(\theta'), \theta' \rangle = \max_{\theta \in \Theta} \langle g(\theta), \theta' \rangle, \quad \forall \theta' \in \Theta. \quad (53)$$

From $1/T$ -net properties, we also have that there exists $\phi \in \Theta'$ such that $\|\theta_\star - \phi\|_2 \leq 1/T$. Hence,

$$\begin{aligned} \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle &\stackrel{(i)}{\leq} \max_{\theta \in \Theta} \langle g(\theta), \theta_\star \rangle \stackrel{(ii)}{=} \langle g(\theta_\star), \theta_\star \rangle \\ &\stackrel{(iii)}{\leq} \langle g(\theta_\star), \phi \rangle + 1/T \\ &\leq \max_{\theta \in \Theta} \langle g(\theta), \phi \rangle + 1/T \\ &\stackrel{(iv)}{=} \langle g(\phi), \phi \rangle + 1/T \leq \langle g(\phi), \theta_\star \rangle + 2/T \\ &\stackrel{(v)}{\leq} \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle + 2/T, \end{aligned} \quad (54)$$

where (i) follows from $\Theta' \subseteq \Theta$, (ii) follows from (41), (iii) follows from $\langle g(\theta_\star), \theta_\star - \phi \rangle \leq \|g(\theta_\star)\|_2 \|\theta_\star - \phi\|_2 \leq 1/T$, (iv) follows from (41), (v) follows as in (iii), and (vi) follows from $\phi \in \Theta'$. Eq. (52) follows. Note that in this part it is important to have $\Theta' \subseteq \Theta$.

As a result $|R_T^\Lambda(L) - R_T^{\Lambda_\epsilon}(L_\epsilon)|$ can be bounded as

$$\begin{aligned} |R_T^\Lambda(L) - R_T^{\Lambda_\epsilon}(L_\epsilon)| &= \left| \sum_{t=1}^T \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle - \max_{\theta \in \Theta} \langle g(\theta), \theta_\star \rangle \right| \\ &\leq \sum_{t=1}^T \left| \max_{\theta \in \Theta'} \langle g(\theta), \theta_\star \rangle - \max_{\theta \in \Theta} \langle g(\theta), \theta_\star \rangle \right| \stackrel{(i)}{\leq} 2, \end{aligned} \quad (55)$$

where (i) follows uses (52) and (41).

Hence, by (51), Theorem 1, (55) and union bound, we have that with probability at least $1 - \delta$

$$|R_T^M(I) - R_T^{\Lambda_\epsilon}(L_\epsilon)| \leq c\sqrt{T \log 1/\delta}. \quad (56)$$

Since $R_T^{\Lambda_\epsilon}(L_\epsilon) = \sum_{m=1}^M R_{T_m}^{\Lambda_{\epsilon_m}}(L_m)$, we have proved the second part of the theorem. It remains to show that in each epoch m the rewards r_t are generated from the for linear bandit instance L_m and to bound the amount of misspecification. By Proposition 12, it suffices to show the following.

Proposition 14 For each $m \in [M]$, we have that with probability at least $1 - \delta/M$ it holds that

$$|\langle g(\theta), \theta' \rangle - \langle g^{(m)}(\theta), \theta' \rangle| \leq 2\sqrt{\frac{\log(2M|\Theta'|/\delta)}{t^{(m)}}} \forall \theta \in \Theta', \theta' \in \Theta. \quad (57)$$

Proof

Since for a fixed θ' we have that $\langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle, \theta' \rangle$ is $1/4$ -subgaussian with mean $\langle g(\theta), \theta' \rangle$, we have that with probability at least $1 - \frac{\delta}{M|\Theta'|^2}$ it holds that

$$\begin{aligned} |\langle g(\theta), \theta' \rangle - \langle g^{(m)}(\theta), \theta' \rangle| &= |\langle g(\theta), \theta' \rangle - \frac{1}{t^{(m+1)}} \sum_{t=1}^{t^{(m+1)}} \langle \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle, \theta' \rangle| \\ &\leq 2\sqrt{\frac{\log(2M|\Theta'|/\delta)}{t^{(m)}}}. \end{aligned} \quad (58)$$

By the union bound, the following holds with probability at least $1 - \delta/M$

$$|\langle g(\theta), \theta' \rangle - \langle g^{(m)}(\theta), \theta' \rangle| \leq 2\sqrt{\frac{\log(2M|\Theta'|/\delta)}{t^{(m)}}} \forall \theta, \theta' \in \Theta'. \quad (59)$$

Let us pick arbitrary $\theta \in \Theta', \theta' \in \Theta$. We have that there is $\phi' \in \Theta'$ such that $\|\theta' - \phi'\|_2 \leq 1/T$. Hence, by Cauchy-Schwartz and the triangle inequality, we have that the following holds with probability at least $1 - \delta/M$

$$\begin{aligned} |\langle g(\theta), \theta' \rangle - \langle g^{(m)}(\theta), \theta' \rangle| &\leq |\langle g(\theta), \phi' \rangle - \langle g^{(m)}(\theta), \phi' \rangle| + \frac{2}{T} \\ &\leq 2\sqrt{\frac{\log(2M|\Theta'|/\delta)}{t^{(m)}}} + \frac{2}{T} \forall \theta \in \Theta', \theta' \in \Theta. \end{aligned} \quad (60)$$

■

This concludes the proof. ■

Appendix E. Computational Complexity

In this appendix, we construct an efficient approximate linear optimization oracle over the set $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$, defined in Algorithm 1. This construction relies on a linear optimization oracle over the sets \mathcal{A}_t , which we define next.

Definition 15 A **linear optimization oracle** for a set \mathcal{A} is a function $\mathcal{O}(\mathcal{A}; \cdot)$ which takes as input $\theta \in \{a \in \mathbb{R}^d | \|a\|_2 \leq 1\}$ and outputs $\mathcal{O}(\mathcal{A}; \theta) \in \mathcal{A}$ with $\langle \mathcal{O}(\mathcal{A}; \theta), \theta \rangle = \sup_{a \in \mathcal{A}} \langle a, \theta \rangle$.

We consider the following discretization of the set Θ that simplifies the construction of the linear optimization oracle over \mathcal{X}_m . The set Θ is discretized to $\Theta' = \{[\theta]_q | \theta \in \Theta\}$, where $[\theta]_q = q\lceil \theta\sqrt{d}/q \rceil / \sqrt{d}$ and q is the discretization parameter. Note that for $\Theta \subseteq \{a \in \mathbb{R}^d | \|a\|_2 \leq 1\}$, the size of the set Θ' can be bounded as $|\Theta'| \leq Cd(1/q)^{4d+2}$, where C is a universal constant. The following lemma shows that $g^{(m)}([\theta/\|\theta\|_2]_q)$ for sufficiently small q can be used as an approximate linear optimization oracle over the set \mathcal{X}_m . We observe that $g^{(m)}$ can be calculated by invoking the linear optimization oracle of the contextual bandit instance at most T times.

Lemma 16 Consider a given $m \in [M]$ and let $g^{(m)}(\theta) = \frac{1}{t^{(m)}} \sum_{t=1}^{t^{(m)}} \arg \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle$, $\mathcal{X}_m = \{g^{(m)}(\theta) | \theta \in \Theta'\}$, where $t^{(m)}$ is the length of phase m and $\Theta' = \{[\theta]_q | \theta \in \Theta\}$ is a discretization of Θ , $[\theta]_q = q \lfloor \theta \sqrt{d} / q \rfloor / \sqrt{d}$ and q is the discretization parameter. For any $\theta \in \{a \in \mathbb{R}^d | \|a\|_2 \leq 1\}$, $\epsilon \in \mathbb{R}^+$, if $q \leq \epsilon/2$, we have that

$$\langle g^{(m)}([\theta / \|\theta\|_2]_q), \theta \rangle \geq \sup_{a \in \mathcal{X}_m} \langle a, \theta \rangle - \epsilon \|\theta\|_2. \quad (61)$$

Proof We first observe that

$$0 \leq \theta - [\theta]_q = \theta - \frac{\lfloor \theta \sqrt{d} / q \rfloor}{\sqrt{d} / q} \leq q / \sqrt{d} \mathbf{1}. \quad (62)$$

It follows that $\|\theta - [\theta]_q\|_2 \leq q$. We notice that

$$\begin{aligned} \langle g^{(m)}(\theta), \theta \rangle &= \frac{1}{t^{(m)}} \sum_{t=1}^{t^{(m)}} \langle \mathcal{O}(\mathcal{A}_t; \theta), \theta \rangle \geq \frac{1}{t^{(m)}} \sum_{t=1}^{t^{(m)}} \langle \mathcal{O}(\mathcal{A}_t; \theta'), \theta \rangle \\ &= \langle g^{(m)}(\theta'), \theta \rangle, \quad \forall \theta' \in \Theta. \end{aligned} \quad (63)$$

Hence,

$$\langle g^{(m)}(\theta), \theta \rangle \geq \sup_{\theta' \in \Theta'} \langle g^{(m)}(\theta'), \theta \rangle. \quad (64)$$

We also have that

$$\begin{aligned} \langle g^{(m)}([\theta]_q), \theta \rangle &= \langle g^{(m)}([\theta]_q), [\theta]_q \rangle + \langle g^{(m)}([\theta]_q), \theta - [\theta]_q \rangle \\ &\geq \langle g^{(m)}([\theta]_q), [\theta]_q \rangle - \|g^{(m)}([\theta]_q)\|_2 \|\theta - [\theta]_q\|_2 \\ &\geq \langle g^{(m)}([\theta]_q), [\theta]_q \rangle - q \\ &\stackrel{(i)}{\geq} \langle g^{(m)}(\theta), [\theta]_q \rangle - q \\ &\geq \langle g^{(m)}(\theta), \theta \rangle - \|g^{(m)}(\theta)\|_2 \|\theta - [\theta]_q\|_2 - q \\ &\geq \langle g^{(m)}(\theta), \theta \rangle - 2q \stackrel{(ii)}{\geq} \sup_{\theta' \in \Theta'} \langle g^{(m)}(\theta'), \theta \rangle - 2q = \sup_{a \in \mathcal{X}_m} \langle a, \theta \rangle - 2q, \end{aligned} \quad (65)$$

where (i) and (ii) follow from (64). As (65) holds $\forall \theta \in \mathbb{R}^d$, we get that

$$\langle g^{(m)}([\theta / \|\theta\|_2]_q), \theta / \|\theta\|_2 \rangle \geq \sup_{a \in \mathcal{X}_m} \langle a, \frac{\theta}{\|\theta\|_2} \rangle - 2q = \frac{\sup_{a \in \mathcal{X}_m} \langle a, \theta \rangle}{\|\theta\|_2} - 2q. \quad (66)$$

It follows that

$$\langle g^{(m)}([\theta / \|\theta\|_2]_q), \theta \rangle \geq \sup_{a \in \mathcal{X}_m} \langle a, \theta \rangle - 2q \|\theta\|_2. \quad (67)$$

■