
The Power of Exploiter: Provable Multi-Agent RL in Large State Spaces

Chi Jin¹ Qinghua Liu¹ Tiancheng Yu²

Abstract

Modern reinforcement learning (RL) commonly engages practical problems with large state spaces, where function approximation must be deployed to approximate either the value function or the policy. While recent progresses in RL theory address a rich set of RL problems with general function approximation, such successes are mostly restricted to the single-agent setting. It remains elusive how to extend these results to multi-agent RL, especially in the face of new game-theoretical challenges. This paper considers two-player zero-sum Markov Games (MGs). We propose a new algorithm that can provably find the Nash equilibrium policy using a polynomial number of samples, for any MG with low *multi-agent Bellman-Eluder dimension*—a new complexity measure adapted from its single-agent version (Jin et al., 2021). A key component of our new algorithm is the exploiter, which facilitates the learning of the main player by deliberately exploiting her weakness. Our theoretical framework is generic, which applies to a wide range of models including but not limited to tabular MGs, MGs with linear or kernel function approximation, and MGs with rich observations.

1. Introduction

Multi-agent reinforcement learning (MARL) systems have recently achieved significant success in many AI challenges including the game Go (Silver et al., 2016; 2017), Poker (Brown & Sandholm, 2019), real-time strategy games (Vinyals et al., 2019; OpenAI, 2018), decentralized controls or multiagent robotics systems (Brambilla et al., 2013), autonomous driving (Shalev-Shwartz et al., 2016), as well as complex social scenarios such as hide-and-seek (Baker et al., 2020). Two crucial components that contribute to these suc-

cesses are *function approximation* and *self-play*. Function approximation is frequently used in modern applications with large state spaces, where either the value function or the policy is approximated by parametric function classes, which are typically deep neural networks. Meanwhile, self-play enables the learner to improve by playing against itself instead of traditional human experts.

Despite the empirical success of MARL, existing theoretical guarantees in MARL only apply to the basic settings where the value functions can be represented by either tables (in cases where the states and actions are discrete) (Bai & Jin, 2020; Bai et al., 2020; Liu et al., 2020) or linear maps (Xie et al., 2020; Chen et al., 2021). While a recent line of works (Jiang et al., 2017; Wang et al., 2020; Yang et al., 2020; Jin et al., 2021; Du et al., 2021) significantly advance our understanding of RL with general function approximation, and provide sample-efficient guarantees for RL with kernels, neural networks, rich observations, and several special cases of partial observability, they are all restricted to the single-agent setting. Distinct from single-agent RL, each agent in MARL are facing not only the unknown environment, but also the opponents that can constantly adapt their strategies in response to the behavior of the learning agent. This additional game-theoretical feature makes it challenging to extend the single-agent general function approximation results to the multi-agent setting. This motivates us to ask the following question:

Can we design sample-efficient MARL algorithms for general function approximation?

By “sample-efficient”, we mean the algorithms provably learn within a polynomial number of samples that is independent of the number of states. This paper provides the first positive answer to this question in the context of two-player zero-sum Markov Games (MGs) (Shapley, 1953; Littman, 1994). In particular, we make the following contributions:

- We design a new self-play algorithm for MARL—GOLF_WITH_EXPLOITER. Our algorithm maintains a main player and an exploiter, where the exploiter facilitates the learning of the main player by deliberately exploiting her weakness. Our algorithm features optimism, and can simultaneously address multi-agent, general function approximation, as well as the trade-off

*Equal contribution ¹Princeton University ²MIT. Correspondence to: Chi Jin <chij@princeton.edu>.

between exploration and exploitation.

- We introduce a new general complexity measure for MARL—*Multi-agent Bellman Eluder (BE) dimension*, which is adapted from its single-agent version (Jin et al., 2021). We prove that our algorithm can learn the Nash equilibrium policies of any MARL problem with low multi-agent BE dimension, using a number of samples that is polynomial in all relevant parameters, but independent of the size of the state space. We further provide an online guarantee for our algorithm when facing against adversarial opponents.
- We show that our theoretical framework applies to a wide range of models including tabular MGs, MGs with linear or kernel function approximation, and MGs with rich observations. Formally, we prove that all examples above admit low multi-agent BE dimension. Particularly, for linear MGs, we obtain a $\tilde{O}(H^2 d^2 / \epsilon^2)$ sample complexity bound for finding ϵ -approximate Nash policies, which improves upon the cubic dependence on dimension d in Xie et al. (2020).

We remark that our algorithm is sample-efficient but *not* computationally efficient. Designing computationally efficient algorithms in the context of general function approximation is an open problem even in the single-agent setting (Jiang et al., 2017; Jin et al., 2021; Du et al., 2021), which we left as an interesting topic for future research.

Technical Challenges One of the most important questions in the design of RL algorithms is the choice of *behavior policies*—the policies that are used to collect samples. For a given general function class \mathcal{F} , the single-agent RL algorithms simply use behavior policy in the form of π_f , which is the greedy policy with respect to certain optimistic value function $f \in \mathcal{F}$, i.e., $\pi_f(s) = \operatorname{argmax}_a f(s, a)$. The choice becomes more sophisticated in the two-player setting. All existing algorithms that can perform provably efficient exploration (Bai & Jin, 2020; Bai et al., 2020; Liu et al., 2020; Xie et al., 2020; Chen et al., 2021) choose the behavior policy as the general-sum Nash equilibrium or coarse correlated equilibrium of two value functions $f, g \in \mathcal{F}$, which are the upper and lower bound of the true value functions respectively. It turns out that this approach is only effective under *optimistic closure*—a strong assumption which is true for tabular and several special linear settings, but does not hold for many practical applications (see Appendix A for more discussions). In contrast, our exploiter framework provides a new approach for the theory community to choose behavior policies—the main agent plays the optimistic Nash policy based on the existing knowledge while the exploiter focuses on exploiting the main agent. Our approach does not rely on strong assumptions such as optimistic closure. We hope our exploiter framework could facilitate the future development of provable MARL algorithms.

Related Works Due to space limit, we defer a detailed survey of related works to Appendix 2.

2. Related Works

There is an extensive literature on empirical MARL, in distributed, cooperative, and competitive settings (see, e.g., Lowe et al. (2017); Rashid et al. (2018); Vinyals et al. (2019); Berner et al. (2019); Li & He (2020); Yu et al. (2021a) and the references therein). Due to space limit, we focus on reviewing the theoretical works in this section.

Markov Game. Markov Game (MG), also known as stochastic game (Shapley, 1953), is a popular model in multi-agent RL (Littman, 1994). Early works have mainly focused on finding Nash equilibria of MGs with known transition and reward (Littman, 2001; Hu & Wellman, 2003; Hansen et al., 2013; Wei et al., 2020), or under strong reachability conditions such as simulators (Jia et al., 2019; Sidford et al., 2020; Zhang et al., 2020a; Wei et al., 2021) where exploration is not needed.

A recent line of works provide non-asymptotic guarantees for learning two-player zero-sum tabular MGs, in the setting that requires strategic exploration. Bai & Jin (2020) and Xie et al. (2020) develop the first provably-efficient learning algorithms in MGs based on optimistic value iteration. Bai et al. (2020) and Liu et al. (2020) improve upon these works and achieve best-known sample complexity for model-free and model-based methods, respectively. Several extensions are also studied, including multi-player general-sum MGs (Liu et al., 2020), unknown games¹ (Tian et al., 2021), vector-valued MGs (Yu et al., 2021b), etc.

Wei et al. (2017) also considers the exploration problem, under the assumption that both players can reach any states within a fixed number steps, regardless of the policy the other player. This reachability condition is necessary for the infinite horizon MGs and can be removed for finite-horizon MGs. They also use an asymmetric policy update as we do.

Beyond tabular MGs, Xie et al. (2020) and Chen et al. (2021) also study learning MGs with linear function approximation. Their techniques heavily rely on “optimistic closure” (see Appendix A for more details), which can not be directly extended to the general function approximation setting with weak assumptions. To our best knowledge, this paper provides the first positive result on learning MGs with general function approximation.

Markov Decision Process. There has been a long line of research studying Markov Decision Process (MDP), which can be viewed as a single-agent version of Markov Game.

¹This terminology stems from game theory, which means the actions of the opponents are not observed.

Tabular MDP has been studied thoroughly in recent years (Brafman & Tennenholtz, 2002; Jaksch et al., 2010; Dann & Brunskill, 2015; Azar et al., 2017; Zanette & Brunskill, 2019; Jin et al., 2018; Zhang et al., 2020b). Particularly, in the episodic setting, the minimax regret or sample complexity is achieved by both model-based (Azar et al., 2017) and model-free (Zhang et al., 2020b) methods, up to logarithmic factors. When the state space is large, the tabular formulation of RL is impractical and function approximation is necessary. The most studied case is linear function approximation (Jin et al., 2020; Wang et al., 2019; Cai et al., 2019; Zanette et al., 2020a;b; Agarwal et al., 2020; Neu & Pike-Burke, 2020).

For general function approximation, there are two common measures of complexity: Eluder dimension (Osband & Van Roy, 2014; Wang et al., 2020) and Bellman rank (Jiang et al., 2017), which are unified by a more generic notion—Bellman-Eluder dimension (Jin et al., 2021). Recently, (Du et al., 2021) develops a new problem class termed bilinear class, which can be considered as an infinite-dimensional extension of low Bellman rank class. Low complexity RL problems under the above criteria are statistically tractable. Polynomial sample complexity guarantees (Jiang et al., 2017; Sun et al., 2019; Jin et al., 2021; Foster et al., 2020), or even regret guarantees (Wang et al., 2019; Dong et al., 2020; Jin et al., 2021; Wang et al., 2020) are established, under additional completeness and realizability conditions. However, most of the above algorithms are not computationally efficient and it remains open how to design computationally efficient algorithms for these settings.

Extensive-form games Finally, we remark that there is another long line of research on MARL based on the model of extensive-form games (EFG, see, e.g., Koller & Megiddo (1992); Gilpin & Sandholm (2006); Zinkevich et al. (2007); Brown & Sandholm (2018; 2019); Celli et al. (2020)). Results on learning EFGs do not directly imply results for learning MGs, since EFGs are naturally tree-structured games, which can not efficiently represent MGs—graph-structured games where a state at the h^{th} step can be the child of multiple states at the $(h - 1)^{\text{th}}$ step.

3. Preliminaries

Markov Games (MGs) are the generalization of standard Markov Decision Processes into the multi-player setting, where each player aims to maximize her own utility. In this paper, we consider *two-player zero-sum episodic* MG, formally denoted by $\text{MG}(H, \mathcal{S}, \mathcal{A}, \mathcal{B}, \mathbb{P}, r)$, where H is the number of steps per episode, \mathcal{S} is the state space, \mathcal{A} and \mathcal{B} are the action space for the max-player and the min-player respectively, $\mathbb{P} = \{\mathbb{P}_h\}_{h \in [H]}$ is the collection of transition measures with $\mathbb{P}_h(\cdot \mid s, a, b)$ defining the

distribution of the next state after taking action (a, b) at state s and step h , $r = \{r_h\}_{h \in [H]}$ is the collection of reward functions with $r_h(s, a, b)$ equal to the reward received after taking action (a, b) at state s and step h .² Throughout this work, we assume the reward is nonnegative, and $\sum_{h=1}^H r_h(s_h, a_h, b_h) \leq 1$ for all possible sequence $(s_1, a_1, b_1, \dots, s_H, a_H, b_H)$.

Each episode starts from a *fixed initial state* s_1 . At each step $h \in [H]$, two players observe $s_h \in \mathcal{S}$ and choose their own actions $a_h \in \mathcal{A}$ and $b_h \in \mathcal{B}$ simultaneously. Then, both players observe the actions of their opponents and receive reward $r_h(s_h, a_h, b_h)$. After that, the environment will transit to $s_{h+1} \sim \mathbb{P}_h(\cdot \mid s_h, a_h, b_h)$. Without loss of generality, we assume the environment always transit to a terminal state s_{end} at step $H + 1$ which marks the end of this episode.

Policy, value function A Markov policy μ of the max-player is a collection of vector-valued functions $\{\mu_h : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}\}_{h \in [H]}$, each mapping a state to a distribution in the probability simplex over \mathcal{A} . We use the a^{th} coordinate of $\mu_h(s)$ to refer to the probability of taking action a at state s and step h . Similarly, we can define a Markov policy ν over \mathcal{B} for the min-player.

We denote $V_h^{\mu, \nu} : \mathcal{S} \rightarrow \mathbb{R}$ as the value function for policy μ and ν at step h , with $V_h^{\mu, \nu}(s)$ equal to the expected cumulative reward received under policy (μ, ν) , starting from state s at step h until the end of the episode:

$$V_h^{\mu, \nu}(s) := \mathbb{E}_{\mu \times \nu} \left[\sum_{t=h}^H r_t(s_t, a_t, b_t) \mid s_h = s \right].$$

Similarly, we denote $Q_h^{\mu, \nu} : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ as the action-value function for policy μ and ν at step h :

$$Q_h^{\mu, \nu}(s, a, b) := \mathbb{E}_{\mu \times \nu} \left[\sum_{t=h}^H r_t(s_t, a_t, b_t) \mid s_h = s, a_h = a, b_h = b \right].$$

With slight abuse of notation, we use \mathbb{P}_h as an operator so that $[\mathbb{P}_h V](s, a, b) := \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot \mid s, a, b)} V(s')$ for any value function V . We also define $[\mathbb{D}_\pi Q](s) := \mathbb{E}_{(a, b) \sim \pi(\cdot, \cdot \mid s)} Q(s, a, b)$ for any policy π and action-value function Q . By the definition of value functions, we have the Bellman equations:

$$Q_h^{\mu, \nu}(s, a, b) = [r_h + \mathbb{P}_h V_{h+1}^{\mu, \nu}](s, a, b),$$

$$V_h^{\mu, \nu}(s) = [\mathbb{D}_{\mu_h \times \nu_h} Q_h^{\mu, \nu}](s),$$

for all $(s, a, b, h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{B} \times [H]$. And for step $H + 1$, we have $V_{H+1}^{\mu, \nu} \equiv 0$.

²In this paper, we study deterministic reward functions for cleaner presentation. All our results immediately generalize to the setting with stochastic reward.

Best response, Nash equilibrium For any policy of the max-player μ , there exists a *best response* policy of the min-player $\nu^\dagger(\mu)$ satisfying $V_h^{\mu, \nu^\dagger(\mu)} = \inf_{\nu} V_h^{\mu, \nu}(s)$ for all (s, h) . For cleaner notation, we denote $V_h^{\mu, \dagger} := V_h^{\mu, \nu^\dagger(\mu)}$. Similarly, we can define $\mu^\dagger(\nu)$ and $V_h^{\dagger, \nu}$. It is known (Filar & Vrieze, 2012) that there exists policies μ^*, ν^* that are optimal against the best responses of the oponents, i.e., for all (s, h)

$$V_h^{\mu^*, \dagger}(s) = \sup_{\mu} V_h^{\mu, \dagger}(s), \quad V_h^{\dagger, \nu^*}(s) = \inf_{\nu} V_h^{\dagger, \nu}(s).$$

We call these optimal policies (μ^*, ν^*) a Nash equilibrium of the Markov game, which are further known to satisfy the following minimax equation for all (s, h)

$$\sup_{\mu} \inf_{\nu} V_h^{\mu, \nu}(s) = V_h^{\mu^*, \nu^*}(s) = \inf_{\nu} \sup_{\mu} V_h^{\mu, \nu}(s).$$

We call the value functions of (μ^*, ν^*) Nash value functions, and abbreviate $V_h^{\mu^*, \nu^*}$ as V_h^* and $Q_h^{\mu^*, \nu^*}$ as Q_h^* . Intuitively speaking, in a Nash equilibrium, no player can benefit by unilaterally deviating from her own policy.

Learning objective We say a policy μ of the max-player is an ϵ -**approximate Nash** policy if $V_1^{\mu, \dagger}(s_1) \geq V_1^*(s_1) - \epsilon$. Suppose an agent interacts with the environment for K episodes, and denote by μ^k the policy executed by the max-player in the k^{th} episode. Then the (cumulative) regret of the max-player is defined as

$$\text{Reg}(K) := \sum_{k=1}^K \left[V_1^*(s_1) - V_1^{\mu^k, \dagger}(s_1) \right].$$

The goal of reinforcement learning is to learn ϵ -approximate Nash policies or to achieve sublinear regret. In this paper, we focus on learning the Nash policies of the max-player. By symmetry, the definitions and techniques directly extend to learning the Nash policies of the min-player.

3.1. Function Approximation

This paper considers reinforcement learning with value function approximation. Formally, the learner is provided with a function class $\mathcal{F} = \mathcal{F}_1 \times \dots \times \mathcal{F}_H$ where each $\mathcal{F}_h \subseteq (\mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1])$ consists of candidate functions to approximate some target action-value functions at step h . Since there is no reward at step $H + 1$, we set $f_{H+1} \equiv 0$ for any $f \in \mathcal{F}$ without loss of generality.

Induced policy and value function For any value function $f \in \mathcal{F}$, we denote by $\mu_f = (\mu_{f,1}, \dots, \mu_{f,H})$ the Nash policy of the max-player induced by f , where

$$\mu_{f,h}(s) = \operatorname{argmax}_{\mu \in \Delta_{\mathcal{A}}} \min_{\nu \in \Delta_{\mathcal{B}}} \mu^\top f_h(s, \cdot, \cdot) \nu \quad \text{for all } (s, h).$$

Moreover, we denote by V_f the Nash value function induced by f , so that

$$V_{f,h}(s) = \max_{\mu \in \Delta_{\mathcal{A}}} \min_{\nu \in \Delta_{\mathcal{B}}} \mu^\top f_h(s, \cdot, \cdot) \nu \quad \text{for all } (s, h).$$

We further generalize the concept of induced value function by additionally introducing a superscript μ (which represents a Markov policy of the max-player) and define V_f^μ so that

$$V_{f,h}^\mu(s) = \min_{\nu \in \Delta_{\mathcal{B}}} \mu_h(s)^\top f_h(s, \cdot, \cdot) \nu \quad \text{for all } (s, h).$$

Based on V_f and V_f^μ , two types of Bellman operators are defined as below

$$\begin{cases} (\mathcal{T}_h f)(s, a, b) := r_h(s, a, b) + \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a, b)} V_{f,h+1}(s') \\ (\mathcal{T}_h^\mu f)(s, a, b) := r_h(s, a, b) + \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a, b)} V_{f,h+1}^\mu(s') \end{cases}$$

which naturally generalize the optimal Bellman operator and the policy Bellman operator from MDPs to MGs. In the remainder of this paper, we will refer to them as the Nash Bellman operator and μ -Bellman operator, respectively.

It is known that RL with function approximation is in general statistically intractable without further assumptions (see, e.g., hardness results in Krishnamurthy et al., 2016; Weisz et al., 2020). Below, we present two assumptions that are generalizations of commonly adopted assumptions in MDP literature.

Assumption 3.1 (ϵ_{real} -Realizability). For any $f \in \mathcal{F}$ and $h \in [H]$, there exist $Q_h, Q'_h \in \mathcal{F}_h$ s.t. $\|Q_h^* - Q_h\|_\infty \leq \epsilon_{\text{real}}$ and $\|Q_h^{\mu_f, \dagger} - Q'_h\|_\infty \leq \epsilon_{\text{real}}$.

Realizability requires the function class to be well-specified so that the value function of the Nash policy and the value function of any induced policy μ_f against its best response lie inside \mathcal{F} approximately.

Let $\mathcal{G} = \mathcal{G}_1 \times \dots \times \mathcal{G}_H$ be an auxiliary function class provided to the learner where each $\mathcal{G}_h \subseteq (\mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1])$. Completeness requires the auxiliary function class \mathcal{G} to be rich enough so that applying Bellman operators to any function in the primary function class \mathcal{F} will end up in \mathcal{G} approximately.

Assumption 3.2 (ϵ_{comp} -Completeness). For any $f, f' \in \mathcal{F}$ and $h \in [H]$, there exist $Q_h, Q'_h \in \mathcal{G}_h$ s.t. $\|\mathcal{T}_h f_{h+1} - Q_h\|_\infty \leq \epsilon_{\text{comp}}$ and $\|\mathcal{T}_h^{\mu_f} f' - Q'_h\|_\infty \leq \epsilon_{\text{comp}}$.

In the single-agent setting (Jin et al., 2021), the realizability and completeness conditions are stated for the special case $\epsilon_{\text{real}} = \epsilon_{\text{comp}} = 0$, while this paper considers a strictly more general condition. This extension makes it possible to handle the misspecified setting, i.e., \mathcal{F} and \mathcal{G} only satisfy these conditions approximately, but not exactly. Even when \mathcal{F} and \mathcal{G} satisfy these conditions exactly, the above

extension can still help us avoid some technical redundancy caused by the possible infinite cardinality of \mathcal{F} and \mathcal{G} , by only considering their finite coverings. To be precise,

Definition 3.3 (ϵ -Covering). The function class $\tilde{\mathcal{F}} \subseteq \mathcal{F}$ is an ϵ -covering of \mathcal{F} if for any $f \in \mathcal{F}$, there exist $\tilde{f} \in \tilde{\mathcal{F}}$ s.t. $\|f_h - \tilde{f}_h\|_\infty \leq \epsilon$, for any step h .

The following lemma implies that we can always restrict our attention to finite coverings, which also approximately satisfy the realizability and completeness conditions up to satisfying accuracy.

Lemma 3.4. *If \mathcal{F} and \mathcal{G} satisfy ϵ_{real} -Realizability and ϵ_{comp} -Completeness, and $\tilde{\mathcal{F}}$ and $\tilde{\mathcal{G}}$ are ϵ -coverings of \mathcal{F} and \mathcal{G} respectively, then $\tilde{\mathcal{F}}$ and $\tilde{\mathcal{G}}$ satisfy $(\epsilon_{\text{real}} + \epsilon)$ -Realizability and $(\epsilon_{\text{comp}} + \epsilon)$ -Completeness.*

As a result, we only need to consider finite function class \mathcal{F} and \mathcal{G} in Section 4. Some of the examples introduced in Section 5 involve infinite function classes. To address this inconsistency, we can first compute the finite coverings of \mathcal{F} and \mathcal{G} , which satisfy the approximate realizability and completeness conditions by Lemma 3.4, and then invoke the algorithms in Section 4 with the coverings instead of the original function classes. It is straightforward to verify that by replacing the cardinality with the covering number, all theoretical guarantees derived in Section 4 still hold.

Finally, we define the L^∞ Projection Operator, which will be frequently used in our technical statements and proofs.

Definition 3.5 (L^∞ Projection Operator). For any function class \mathcal{F} and function g , $\mathcal{P}_{\mathcal{F}}(g) = \operatorname{argmin}_{f \in \mathcal{F}} \|f - g\|_\infty$.

4. Main Results

In this section, we present an optimization-based algorithm `GOLF_WITH_EXPLOITER`, and its theoretical guarantees for any multi-agent RL problem with low Bellman-Eluder dimension.

4.1. Algorithm

We describe `GOLF_WITH_EXPLOITER` in Algorithm 1. At a high level, `GOLF_WITH_EXPLOITER` follows the principle of optimism in face of uncertainty, and maintains a global confidence set for Nash value function Q^* based on local constraints. It extends the single-agent GOLF algorithm (Jin et al., 2021) by introducing an exploiter subroutine—`COMPUTE_EXPLOITER`, which facilitates the learning of the main player (the max player) by continuously exploiting her weakness.

In each episode, `GOLF_WITH_EXPLOITER` performs four main steps:

1. Optimistic planning (Line 3): compute the value func-

tion f^k that induces the most optimistic Nash value from the current confidence set \mathcal{C} , and choose μ^k to be its induced Nash policy of the max-player.

2. Finding exploiter (Line 4): compute an approximate best response policy ν^k for the min-player, by invoking the subroutine `COMPUTE_EXPLOITER` on historical data \mathcal{D} and the policy of the max-player μ^k .
3. Data collection (Line 8-9): we provide two different options for the `SAMPLING` procedure
 - (I) Roll a single trajectory by following (μ^k, ν^k) .
 - (II) For each h , roll a trajectory by following (μ^k, ν^k) in the first $h - 1$ steps and take uniformly random action at step h . This option will roll H trajectories each time, but in the h^{th} trajectory, only the h^{th} transition-reward tuple is augmented into \mathcal{D}_h .
4. Confidence set update (Line 10): update confidence set \mathcal{C} using the renewed dataset \mathcal{D} .

The core components of `GOLF_WITH_EXPLOITER` are the construction of the confidence set and the subroutine `COMPUTE_EXPLOITER`, which we elaborate below in sequence.

Confidence set construction For each $h \in [H]$, `GOLF_WITH_EXPLOITER` maintains a local constraint using the historical data \mathcal{D}_h at this step

$$\mathcal{L}_{\mathcal{D}_h}(f_h, f_{h+1}) \leq \inf_{g \in \mathcal{G}_h} \mathcal{L}_{\mathcal{D}_h}(g, f_{h+1}) + \beta, \quad (1)$$

where the squared loss $\mathcal{L}_{\mathcal{D}_h}$ is defined as

$$\begin{aligned} & \mathcal{L}_{\mathcal{D}_h}(\xi_h, \zeta_{h+1}) \\ := & \sum_{(s,a,b,r,s') \in \mathcal{D}_h} [\xi_h(s, a, b) - r - V_{\zeta, h+1}(s')]^2. \end{aligned} \quad (2)$$

Intuitively speaking, $\mathcal{L}_{\mathcal{D}_h}$ is a proxy to the squared Bellman error under the Nash Bellman operator at step h . Unlike the classic Fitted Q-Iteration algorithm (FQI) (Szepesvári, 2010) where one simply updates $f_h \leftarrow \operatorname{argmin}_{\phi \in \mathcal{F}_h} \mathcal{L}_{\mathcal{D}_h}(\phi, f_{h+1})$, our constraints implement a soft version of minimization so that any function with loss slightly larger than the optimal loss gets preserved. By making this relaxation, we can guarantee the true Nash value function Q^* contained in the confidence set \mathcal{C} with high probability.

Exploiter computation The main challenge of learning MGs compared to learning MDPs lies in the choice of behavior policies as discussed in Section 1. Motivated by the empirical breakthrough work AlphaStar (Vinyals et al., 2019), we adapt the methodology of exploiter. Specifically, we

Algorithm 1 GOLF_WITH_EXPLOITER ($\mathcal{F}, \mathcal{G}, K, \beta$)

- 1: **Initialize:** $\mathcal{D}_1, \dots, \mathcal{D}_H \leftarrow \emptyset, \mathcal{C} \leftarrow \mathcal{F}$.
- 2: **for episode** k from 1 to K **do**
- 3: **Choose policy** $\mu^k = \mu_{f^k}$, where $f^k = \operatorname{argmax}_{f \in \mathcal{C}} V_{f,1}(s_1)$. Let $\bar{V}^k = V_{f^k,1}(s_1)$.
- 4: $(\nu^k, \underline{V}^k) \leftarrow \text{COMPUTE_EXPLOITER}(\mathcal{F}, \mathcal{G}, \beta, \mathcal{D}, \mu^k)$.
- 5: **if** $\bar{V}^k - \underline{V}^k < \Delta$ **then**
- 6: **Output** $\mu^{\text{out}} = \mu^k$.
- 7: **end if**
- 8: **Collect samples**
- Option I** execute policy (μ^k, ν^k) and collect a trajectory $(s_1, a_1, b_1, r_1, \dots, s_H, a_H, b_H, r_H, s_{H+1})$.
- Option II** for all $h \in [H]$, execute policy (μ^k, ν^k) at the first $h - 1$ steps, take uniformly random action at the h^{th} step, and collect a sample $(s_h, a_h, b_h, r_h, s_{h+1})$.
- 9: **Augment** $\mathcal{D}_h = \mathcal{D}_h \cup \{(s_h, a_h, r_h, s_{h+1})\}$ for all $h \in [H]$.
- 10: **Update**
- $$\mathcal{C} = \left\{ f \in \mathcal{F} : \mathcal{L}_{\mathcal{D}_h}(f_h, f_{h+1}) \leq \inf_{g \in \mathcal{G}_h} \mathcal{L}_{\mathcal{D}_h}(g, f_{h+1}) + \beta \text{ for all } h \in [H] \right\}$$
- 11: **end for**

design the COMPUTE_EXPLOITER subroutine that approximately computes a best-response policy ν^k against the policy of the max-player μ^k . By executing ν^k , the min-player exposes the flaws of the max-player, and helps improve her strategy. The pseudocode of COMPUTE_EXPLOITER is given in Algorithm 2, which basically follows the same rationale as Algorithm 1 except that we change the regression target by replacing V_f with V_f^μ , because the confidence set \mathcal{C}^μ is constructed for the best-response value function $Q^{\mu, \dagger}$ instead of the Nash value function Q^* . Formally, Algorithm 2 adapts a new square loss function defined below

$$\begin{aligned} & \mathcal{L}_{\mathcal{D}_h}^\mu(\xi_h, \zeta_{h+1}) \\ := & \sum_{(s,a,b,r,s') \in \mathcal{D}_h} [\xi_h(s, a, b) - r - V_{\zeta, h+1}^\mu(s')]^2. \end{aligned} \quad (3)$$

Finally, we remark that the optimistic planning and exploiter computation steps are computationally inefficient in general. Designing computationally efficient algorithms in the context of general function approximation is an open problem even in the single-agent setting (Jiang et al., 2017; Jin et al., 2021; Du et al., 2021), which we left as an interesting topic for future research.

Algorithm 2 COMPUTE_EXPLOITER ($\mathcal{F}, \mathcal{G}, \beta, \mathcal{D}, \mu$)

- 1: **Construct**
- $$\mathcal{C}^\mu = \left\{ f \in \mathcal{F} : \mathcal{L}_{\mathcal{D}_h}^\mu(w_h, w_{h+1}) \leq \inf_{g \in \mathcal{G}_h} \mathcal{L}_{\mathcal{D}_h}^\mu(g, w_{h+1}) + \beta \text{ for all } h \in [H] \right\}$$
- 2: **Compute** $\tilde{f} \leftarrow \operatorname{argmin}_{f \in \mathcal{C}} V_{f,1}^\mu(s_1)$ and $V = V_{\tilde{f},1}^\mu(s_1)$.
- 3: **Return** (ν^{out}, V) such that for all s, h :

$$\nu_h^{\text{out}}(s) = \operatorname{argmin}_{\nu \in \Delta_{\mathcal{B}}} \mu_h(s)^\top \tilde{f}_h(s, \cdot, \cdot) \nu$$

4.2. Complexity Measure

In this subsection, we introduce our new complexity measure—multi-agent Bellman-Eluder (BE) dimension, which is generalized from its single-agent version (Jin et al., 2021). We will show in Section 5 that the family of low BE dimension problems includes many interesting multi-agent RL settings, such as tabular MGs, MGs with linear or kernel function approximation, and MGs with rich observation.

We start by recalling the definition of distributional Eluder (DE) dimension (Jin et al., 2021).

Definition 4.1 (ϵ -independence between distributions). Let \mathcal{G} be a function class defined on \mathcal{X} , and ν, μ_1, \dots, μ_n be probability measures over \mathcal{X} . We say ν is ϵ -independent of $\{\mu_1, \mu_2, \dots, \mu_n\}$ with respect to \mathcal{G} if there exists $g \in \mathcal{G}$ such that $\sqrt{\sum_{i=1}^n (\mathbb{E}_{\mu_i}[g])^2} \leq \epsilon$, but $|\mathbb{E}_\nu[g]| > \epsilon$.

Definition 4.2 (Distributional Eluder (DE) dimension). Let \mathcal{G} be a function class defined on \mathcal{X} , and Π be a family of probability measures over \mathcal{X} . The distributional Eluder dimension $\dim_{\text{DE}}(\mathcal{G}, \Pi, \epsilon)$ is the length of the longest sequence $\{\rho_1, \dots, \rho_n\} \subset \Pi$ such that there exists $\epsilon' \geq \epsilon$ where ρ_i is ϵ' -independent of $\{\rho_1, \dots, \rho_{i-1}\}$ for all $i \in [n]$.

DE dimension generalizes the original Eluder dimension (Russo & Van Roy, 2013) from points to distributions. The main advantage of this generalization is that for RL problems with large state space, it is oftentimes easier to evaluate a function with respect to certain distribution family than to estimate it point-wisely. And for the purpose of this paper, we will focus on the following two specific distribution families.

1. $\mathcal{D}_\Delta := \{\mathcal{D}_{\Delta, h}\}_{h \in [H]}$, where

$$\mathcal{D}_{\Delta, h} = \{\delta_{(s,a,b)}(\cdot) \mid (s, a, b) \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}\},$$

i.e., the collections of probability measures that put measure 1 on a single state-action pair.

2. $\mathcal{D}_{\mathcal{F}} := \{\mathcal{D}_{\mathcal{F}, h}\}_{h \in [H]}$, where $\mathcal{D}_{\mathcal{F}, h}$ denotes the collections of all probability measures over $\mathcal{S} \times \mathcal{A} \times \mathcal{B}$ at the h^{th} step, which can be generated by executing some

$(\mu_f, \nu_{f,g})$ with $f, g \in \mathcal{F}$. Here $\nu_{f,g}$ is the best response to μ_f regarding to Q-value g :

$$\nu_{f,g,h}(s) := \operatorname{argmin}_{\nu \in \Delta_{\mathcal{B}}} \mu_{f,h}(s)^\top g_h(s, \cdot, \cdot) \nu \text{ for all } (s, h).$$

Now, we are ready to introduce our key complexity notion—multi-agent Bellman-Eluder (BE) dimension, which is simply the DE dimension on the function class of Bellman residuals, minimizing over the two aforementioned distribution families, and maximizing over all steps.

Definition 4.3 (Bellman-Eluder dimension). Let $\mathcal{H}_{\mathcal{F}}$ be the function classes of Bellman residuals where $\mathcal{H}_{\mathcal{F},h} := \{f_h - \mathcal{T}_h^{\mu_g} f_{h+1} \mid f, g \in \mathcal{F}\}$. Then the Bellman-Eluder dimension is defined as

$$\dim_{\text{BE}}(\mathcal{F}, \epsilon) = \max_{h \in [H]} \min_{\mathcal{D} \in \{\mathcal{D}_{\Delta}, \mathcal{D}_{\mathcal{F}}\}} \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_h, \epsilon).$$

We remark that the Bellman residual functions used in Definition 4.3 take both state and action as input. By choosing an alternative class of Bellman residuals defined over the state space only, we can similarly define a variant of this notion—V-type BE dimension. For clean presentation, we defer the formal definition to Appendix B.

4.3. Theoretical Guarantees

Now we are ready to present the theoretical guarantees.

Theorem 4.4 (Regret of GOLF_WITH_EXPLOITER). *Under Assumption 3.1 and 3.2, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$, then with probability at least $1 - \delta$, for all $k \in [K]$, Algorithm 1 with Option I satisfies:*

$$\text{Reg}(k) \leq \mathcal{O}(H\sqrt{dk\beta}),$$

where $d = \dim_{\text{BE}}(\mathcal{F}, 1/K)$ is the BE dimension.

Theorem 4.4 claims that GOLF_WITH_EXPLOITER can achieve \sqrt{K} regret on any RL problem with low BE dimension, under the realizability and the completeness assumption. Moreover, the regret scales polynomially with respect to the length of episode, the BE dimension of function class \mathcal{F} , and the log-cardinality of the two function classes. In particular, it is independent of the size of the state space, which is of vital importance for practical RL problems, where the state space is oftentimes prohibitively large.

By pigeonhole principle, we also derive a sample complexity guarantee for GOLF_WITH_EXPLOITER.

Corollary 4.5 (Sample Complexity). *Under Assumption 3.1 and 3.2, there exists absolute constants c, c' such that for any $\epsilon > 0$, choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) +$*

$K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ and $\Delta = c'(H\sqrt{d\beta/K} + \epsilon)$, where $d = \dim_{\text{BE}}(\mathcal{F}, \epsilon/H)$ is the BE dimension of \mathcal{F} , then with probability at least $1 - \delta$, the output condition (Line 5) will be satisfied at least once in the first K episodes.

Furthermore, the output policy μ^{out} of Algorithm 1 with Option I is $\mathcal{O}(\epsilon + H\sqrt{d}(\varepsilon_{\text{real}} + \varepsilon_{\text{comp}}))$ -approximate Nash, if

$$K \geq \Omega((H^2 d / \epsilon^2) \cdot \log(H|\mathcal{F}||\mathcal{G}|d/\epsilon)).$$

Corollary 4.5 asserts that GOLF_WITH_EXPLOITER can find an $\mathcal{O}(\epsilon)$ -approximate Nash policy for the max-player using at most $\tilde{\mathcal{O}}(H^2 \log(|\mathcal{F}||\mathcal{G}|) \dim_{\text{BE}}(\mathcal{F}, \epsilon)/\epsilon^2)$ samples, which scales linearly in the BE dimension and the log-cardinality of the function classes. Since our ultimate goal is to compute an approximate Nash policy, one can early terminate Algorithm 1 once the output condition (Line 5) is satisfied. We remark that by symmetry GOLF_WITH_EXPLOITER can also compute an $\mathcal{O}(\epsilon)$ -approximate Nash policy for the min-player using the same amount of samples.

Remark 4.6 (Generalization to function classes with infinite cardinality). Although Theorem 4.4 and Corollary 4.5 are derived for finite function classes, they can also apply to infinite function classes with bounded ℓ_{∞} -covering number by Lemma 3.4. For example, they hold for linear function classes, with the log-cardinality replaced by the log-covering number.

V-type BE dimension We also derive similar sample complexity guarantee in terms of V-type BE dimension for Algorithm 1 with Option II, which can be found in Appendix B.

4.4. Adversarial Opponents

So far we have been focusing on the self-play setting, where the learner can control both players and the goal is to learn an approximate Nash policy. Another setting of importance is the online setting, where the learner can only control the max-player, and the goal is to achieve high cumulative reward against the adversarial min-player. Intuitively, it is reasonable to expect GOLF_WITH_EXPLOITER could still work in the online setting, because Theorem 4.4 demonstrates that it can achieve sublinear regret competing against the exploiter, which is the strongest possible adversary. This is indeed the case. The only place we need to change in Algorithm 1 is Line 4: instead of computing ν^k using Algorithm 2, we let the adversary pick policy ν^k . Before stating the theoretical guarantee, we introduce an online version of Bellman Eluder dimension.

Definition 4.7 (Online Bellman-Eluder dimension). Let $\mathcal{H}_{\mathcal{F}}$ be the function classes of Bellman residuals where $\mathcal{H}_{\mathcal{F},h} := \{f_h - \mathcal{T}_h f_{h+1} \mid f \in \mathcal{F}\}$. Then the online Bellman-

Eluder dimension is defined as:

$$\dim_{\text{OBE}}(\mathcal{F}, \epsilon) = \max_{h \in [H]} \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F}, h}, \mathcal{D}_{\Delta, h}, \epsilon). \quad (4)$$

Compared to the BE dimension in Definition 4.3, this online version uses a smaller function class that includes only the residuals with respect to the Nash Bellman operator. Another difference is the choice of distribution family is now limited to only \mathcal{D}_{Δ} because the learner cannot control the policy of the min-player in the online setting. Now, we are ready to present the theoretical guarantee.

Theorem 4.8 (Regret against adversarial opponents). *Assuming $\|Q_h^* - \mathcal{P}_{\mathcal{F}_h}(Q_h^*)\| \leq \epsilon_{\text{real}}$ and $\|\mathcal{T}_h f_{h+1} - \mathcal{P}_{\mathcal{G}_h}(\mathcal{T}_h f_{h+1})\| \leq \epsilon_{\text{comp}}$ for all $f \in \mathcal{F}$ and $h \in [H]$, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\epsilon_{\text{comp}}^2 + K\epsilon_{\text{real}}^2)$, then with probability at least $1 - \delta$, for all $k \in [K]$, Algorithm 1 with Option 1 satisfies*

$$\sum_{t=1}^k [V_1^*(s_1) - V_1^{\mu^t, \nu^t}(s_1)] \leq \mathcal{O}(H\sqrt{dk\beta}),$$

where $d = \dim_{\text{OBE}}(\mathcal{F}, 1/K)$ is the online BE dimension.

Theorem 4.8 claims that on any problem of low online BE dimension, GOLF_WITH_EXPLOITER can achieve \sqrt{K} regret against an adversarial opponent. Moreover, the multiplicative factor scales linearly in the horizon length, and sublinearly in the online BE dimension as well as the log-cardinality of the function classes. Comparing to the self-play regret guarantee, Theorem 4.8 holds under a weaker condition, but only guarantees Algorithm 1 can play favorably against the adversary, instead of the best-response policy. This is unavoidable because if the adversary is very weak, then it is impossible to learn Nash policies by playing against it.

5. Examples

In this section, we introduce five concrete multi-agent RL problems with low BE dimension: tabular MGs, MGs with linear function approximation, MGs with kernel function approximation, MGs with rich observation, and feature selection for kernel MGs, which generalize their single-agent versions (Zanette et al., 2020a; Yang et al., 2020; Krishnamurthy et al., 2016; Agarwal et al., 2020) from MDPs to MGs. Except for tabular MGs, all above examples are new and can not be addressed by existing works (Xie et al., 2020; Chen et al., 2021).

Tabular MGs Starting with the simplest scenario, we show that MGs with finite state-action space has BE dimension no larger than the size of its state-action space, up to a logarithmic factor.

Proposition 5.1 (Tabular MGs \subset Low BE dimension). *Consider a tabular MG with state set \mathcal{S} and action set $\mathcal{A} \times \mathcal{B}$. Then any function class $\mathcal{F} \subseteq (\mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1])$ satisfies*

$$\dim_{\text{BE}}(\mathcal{F}, \epsilon) \leq \mathcal{O}(|\mathcal{S} \times \mathcal{A} \times \mathcal{B}| \cdot \log(1 + 1/\epsilon)).$$

Linear function approximation We consider linear function class \mathcal{F} consisting of functions linear in a d -dimensional feature mapping. Specifically, for each $h \in [H]$, we choose $\mathcal{F}_h = \{\phi_h(\cdot, \cdot, \cdot)^\top \theta \mid \theta \in B_d(R)\}$ where $\phi_h : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow B_d(1)$ maps a state-action pair into the d -dimensional unit ball centered at the origin.

Assumption 5.2 (Self-completeness). $\mathcal{T}_h^{\mu^f} \mathcal{F}_{h+1} \subset \mathcal{F}_h$ for any $(f, h) \in \mathcal{F} \times [H]$.

Assumption 5.2 is a special case of Assumption 3.2 (completeness) by choosing auxiliary function class $\mathcal{G} = \mathcal{F}$. Assumption 5.2 also implies Assumption 3.1 (realizability) by backward induction. Below, we show under self-completeness, the BE dimension of \mathcal{F} is upper bounded by the dimension of the feature mappings d up to a logarithmic factor.

Proposition 5.3 (Linear FA \subset Low BE dimension). *Under Assumption 5.2,*

$$\dim_{\text{BE}}(\mathcal{F}, \epsilon) \leq \mathcal{O}(d \cdot \log(1 + R/\epsilon)).$$

We remark that our setting of linear function approximation is strictly more general than linear MGs (Xie et al., 2020) which further assumes both transitions and rewards are linear. The self-completeness assumption is automatically satisfied in linear MGs. Prior works (Xie et al., 2020; Chen et al., 2021) can only address linear MGs or its variant but not the more general setting of linear function approximation presented here. Finally, for linear MGs, by combining Proposition 5.3 and Corollary 4.5, we immediately obtain a $\tilde{\mathcal{O}}(H^2 d^2 / \epsilon^2)$ sample complexity bound for finding ϵ -approximate Nash policies, which improves upon the cubic dependence on dimension d in Xie et al. (2020).

Kernel function approximation This setting extends linear function approximation from d -dimensional Euclidean space \mathbb{R}^d to a decomposable Hilbert space \mathcal{H} . Formally, for each $h \in [H]$, we choose $\mathcal{F}_h = \{\phi_h(\cdot, \cdot, \cdot)^\top \theta \mid \theta \in B_{\mathcal{H}}(R)\}$ where $\phi_h : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow B_{\mathcal{H}}(1)$. Since the ambient dimension of \mathcal{H} can be infinite, we leverage the notion of *effective dimension* and prove the BE dimension of \mathcal{F} is no larger than the effective dimension of certain feature sets in \mathcal{H} .

Proposition 5.4 (Kernel FA \subset Low BE dimension). *Under Assumption 5.2,*

$$\dim_{\text{BE}}(\mathcal{F}, \epsilon) \leq \max_{h \in [H]} d_{\text{eff}}(\mathcal{X}_h, \epsilon / (2R + 1)), \quad (5)$$

where $\mathcal{X}_h = \{\mathbb{E}_\mu[\phi_h(s_h, a_h, b_h)] \mid \mu \in \mathcal{D}_{\mathcal{F}, h}\}$ and d_{eff} is the effective dimension.

The ϵ -effective dimension of a set \mathcal{Z} is the minimum integer $d_{\text{eff}}(\mathcal{Z}, \epsilon) = n$ such that

$$\sup_{z_1, \dots, z_n \in \mathcal{Z}} \frac{1}{n} \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n z_i z_i^\top \right) \leq e^{-1}.$$

For $\mathcal{H} = \mathbb{R}^d$, the effective dimension in the RHS of (5) is upper bounded by $\tilde{O}(d)$ by following the standard ellipsoid potential argument.

MGs with rich observation In this scenario, despite the MG has a large number of states, these states can be categorized into a small number of “effective states”. Formally, there exists an unknown decoding function $q : \mathcal{S} \rightarrow [m]$ such that if two states s and s' satisfy $q(s) = q(s')$, then their transition measures and reward functions are identical, i.e., $\mathbb{P}_h(\cdot \mid s, a, b) = \mathbb{P}_h(\cdot \mid s', a, b)$, $\mathbb{P}_h(s \mid \cdot, a, b) = \mathbb{P}_h(s' \mid \cdot, a, b)$, and $r_h(s, a, b) = r_h(s', a, b)$ for all (h, a, b) . Here, the decoding function q is unknown to the learner, and the size of its codomain m is usually much smaller than that of the state space.

Proposition 5.5 (MGs with rich observation \subset Low BE dimension). *Consider an MG with decoding function $q : \mathcal{S} \rightarrow [m]$. Then any function class $\mathcal{F} \subseteq (\mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1])$ satisfies*

$$\dim_{\text{VBE}}(\mathcal{F}, \epsilon) \leq \mathcal{O}(m \cdot \log(1 + 1/\epsilon)),$$

where \dim_{VBE} denotes the V-type BE dimension defined in Appendix B.

Kernel feature selection To begin with, we introduce kernel MGs, which is a special case of kernel function approximation introduced earlier in this section by additionally assuming that both transitions and rewards are linear in RKHS. Concretely, in a kernel MG, for each step $h \in [H]$, there exist feature mappings $\phi_h : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{H}$ and $\psi_h : \mathcal{S} \rightarrow \mathcal{H}$ where \mathcal{H} is a decomposable Hilbert space, so that the transition measure can be represented as the inner product of features, i.e., $\mathbb{P}_h(s' \mid s, a, b) = \langle \phi_h(s, a, b), \psi_h(s') \rangle_{\mathcal{H}}$. Besides, the reward function is linear in ϕ , i.e., $r_h(s, a, b) = \langle \phi_h(s, a, b), \theta_h^r \rangle_{\mathcal{H}}$ for some $\theta_h^r \in \mathcal{H}$. Here, ϕ , ψ and θ^r are *unknown* to the learner. Moreover, a kernel MG satisfies the following regularization conditions: for all h ,

- $\|\theta_h^r\|_{\mathcal{H}} \leq 1$ and
- $\|\sum_{s \in \mathcal{S}} \mathcal{V}(s) \psi_h(s)\|_{\mathcal{H}} \leq 1$ for any function $\mathcal{V} : \mathcal{S} \rightarrow [0, 1]$.

In kernel feature selection, the learner is provided with a feature class Φ satisfying: (a) $\phi \in \Phi$, and (b) $\|\tilde{\phi}_h(s, a, b)\|_2 \leq$

1 for all (s, a, b, h) and $\tilde{\phi} \in \Phi$. Therefore, it is natural to consider the union of all linear function classes induced by Φ . Specifically, let $\mathcal{F} := \mathcal{F}_1 \times \dots \times \mathcal{F}_H$ where $\mathcal{F}_h = \{\langle \tilde{\phi}_h(\cdot, \cdot, \cdot), \theta \rangle_{\mathcal{H}} \mid \theta \in B_{\mathcal{H}}(R), \tilde{\phi} \in \Phi\}$. Below, we show the V-type BE dimension of \mathcal{F} is upper bounded by the effective dimension of certain feature sets induced by \mathcal{F} .

Proposition 5.6 (Kernel feature selection \subset Low BE dimension). *Let \mathcal{M} be a kernel MG. Then*

$$\dim_{\text{VBE}}(\mathcal{F}, \epsilon) \leq \max_{h \in [H]} d_{\text{eff}}(\mathcal{X}_h, \epsilon/(2R + 1)),$$

where $\mathcal{X}_h = \{\mathbb{E}_\mu[\phi_h(s_h, a_h, b_h)] \mid \mu \in \mathcal{D}_{\mathcal{F}, h}\}$.

6. Conclusion

This paper presents the first line of sample-efficient results for Markov games with large state space and general function approximation. We propose a new complexity measure—multiagent Bellman-Eluder dimension, and design a new algorithm that can sample-efficiently learn any MGs with low BE dimension. At the heart of our algorithm is the exploiter, which facilitates the learning of the main player by continuously exploiting her weakness. Our generic framework applies to a wide range of new problems including MGs with linear or kernel function approximation, MGs with rich observations, and kernel feature selection, all of which can not be addressed by existing works.

Acknowledgements

This work was partially supported by Office of Naval Research Grant N00014-22-1-2253. TY is partially supported by NSF CCF-2112665 (TILOS AI Research Institute).

References

- Agarwal, A., Kakade, S., Krishnamurthy, A., and Sun, W. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in Neural Information Processing Systems*, 33, 2020.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. *arXiv preprint arXiv:1703.05449*, 2017.
- Bai, Y. and Jin, C. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pp. 551–560. PMLR, 2020.
- Bai, Y., Jin, C., and Yu, T. Near-optimal reinforcement learning with self-play. *arXiv preprint arXiv:2006.12007*, 2020.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from

- multi-agent autotutorials. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkxpxJBKwS>.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Brafman, R. I. and Tennenholtz, M. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct): 213–231, 2002.
- Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- Brown, N. and Sandholm, T. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- Brown, N. and Sandholm, T. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*, 2019.
- Celli, A., Marchesi, A., Farina, G., and Gatti, N. No-regret learning dynamics for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems*, 33, 2020.
- Chen, Z., Zhou, D., and Gu, Q. Almost optimal algorithms for two-player markov games with linear function approximation. *arXiv preprint arXiv:2102.07404*, 2021.
- Dann, C. and Brunskill, E. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2818–2826, 2015.
- Dong, K., Peng, J., Wang, Y., and Zhou, Y. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. In *Conference on Learning Theory*, pp. 1554–1557. PMLR, 2020.
- Du, S. S., Kakade, S. M., Lee, J. D., Lovett, S., Mahajan, G., Sun, W., and Wang, R. Bilinear classes: A structural framework for provable generalization in rl. *arXiv preprint arXiv:2103.10897*, 2021.
- Filar, J. and Vrieze, K. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.
- Foster, D. J., Rakhlin, A., Simchi-Levi, D., and Xu, Y. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020.
- Gilpin, A. and Sandholm, T. Finding equilibria in large sequential games of imperfect information. In *Proceedings of the 7th ACM conference on Electronic commerce*, pp. 160–169, 2006.
- Hansen, T. D., Miltersen, P. B., and Zwick, U. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. *Journal of the ACM (JACM)*, 60(1):1–16, 2013.
- Hu, J. and Wellman, M. P. Nash q-learning for general-sum stochastic games. *Journal of machine learning research*, 4(Nov):1039–1069, 2003.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Jia, Z., Yang, L. F., and Wang, M. Feature-based q-learning for two-player stochastic games. *arXiv preprint arXiv:1906.00423*, 2019.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143, 2020.
- Jin, C., Liu, Q., and Miryoosefi, S. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *arXiv preprint arXiv:2102.00815*, 2021.
- Koller, D. and Megiddo, N. The complexity of two-person zero-sum games in extensive form. *Games and economic behavior*, 4(4):528–552, 1992.
- Krishnamurthy, A., Agarwal, A., and Langford, J. Pac reinforcement learning with rich observations. *arXiv preprint arXiv:1602.02722*, 2016.
- Li, H. and He, H. Multi-agent trust region policy optimization. *arXiv preprint arXiv:2010.07916*, 2020.
- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pp. 157–163. Elsevier, 1994.
- Littman, M. L. Friend-or-foe q-learning in general-sum games. In *ICML*, volume 1, pp. 322–328, 2001.

- Liu, Q., Yu, T., Bai, Y., and Jin, C. A sharp analysis of model-based reinforcement learning with self-play. *arXiv preprint arXiv:2010.01604*, 2020.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*, 2017.
- Neu, G. and Pike-Burke, C. A unifying view of optimism in episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- OpenAI. Openai five. <https://blog.openai.com/openai-five/>, 2018.
- Osband, I. and Van Roy, B. Model-based reinforcement learning and the eluder dimension. In *Advances in Neural Information Processing Systems*, pp. 1466–1474, 2014.
- Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., and Whiteson, S. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 4295–4304. PMLR, 2018.
- Russo, D. and Van Roy, B. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pp. 2256–2264, 2013.
- Shalev-Shwartz, S., Shammah, S., and Shashua, A. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Shapley, L. S. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- Sidford, A., Wang, M., Yang, L., and Ye, Y. Solving discounted stochastic two-player games with near-optimal time and sample complexity. In *International Conference on Artificial Intelligence and Statistics*, pp. 2992–3002. PMLR, 2020.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Sun, W., Jiang, N., Krishnamurthy, A., Agarwal, A., and Langford, J. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*, pp. 2898–2933, 2019.
- Szepesvári, C. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Tian, Y., Wang, Y., Yu, T., and Sra, S. Online learning in unknown markov games. *arXiv preprint arXiv:2010.15020*, 2021.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Wang, R., Salakhutdinov, R., and Yang, L. F. Provably efficient reinforcement learning with general value function approximation. *arXiv preprint arXiv:2005.10804*, 2020.
- Wang, Y., Wang, R., Du, S. S., and Krishnamurthy, A. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.
- Wei, C.-Y., Hong, Y.-T., and Lu, C.-J. Online reinforcement learning in stochastic games. *arXiv preprint arXiv:1712.00579*, 2017.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. Linear last-iterate convergence in constrained saddle-point optimization. *arXiv e-prints*, pp. arXiv–2006, 2020.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive markov games. *arXiv preprint arXiv:2102.04540*, 2021.
- Weisz, G., Amortila, P., and Szepesvári, C. Exponential lower bounds for planning in mdps with linearly-realizable optimal action-value functions. *arXiv preprint arXiv:2010.01374*, 2020.
- Xie, Q., Chen, Y., Wang, Z., and Yang, Z. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In *Conference on Learning Theory*, pp. 3674–3682. PMLR, 2020.
- Yang, Z., Jin, C., Wang, Z., Wang, M., and Jordan, M. I. Bridging exploration and general function approximation in reinforcement learning: Provably efficient kernel and neural value iterations. *arXiv preprint arXiv:2011.04622*, 2020.
- Yu, C., Velu, A., Vinitsky, E., Wang, Y., Bayen, A., and Wu, Y. The surprising effectiveness of mappo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021a.

- Yu, T., Tian, Y., Zhang, J., and Sra, S. Provably efficient algorithms for multi-objective competitive rl. *arXiv preprint arXiv:2102.03192*, 2021b.
- Zanette, A. and Brunskill, E. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.
- Zanette, A., Lazaric, A., Kochenderfer, M., and Brunskill, E. Learning near optimal policies with low inherent bellman error. *arXiv preprint arXiv:2003.00153*, 2020a.
- Zanette, A., Lazaric, A., Kochenderfer, M. J., and Brunskill, E. Provably efficient reward-agnostic navigation with linear value iteration. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Zhang, K., Kakade, S. M., Başar, T., and Yang, L. F. Model-based multi-agent rl in zero-sum markov games with near-optimal sample complexity. *arXiv preprint arXiv:2007.07461*, 2020a.
- Zhang, Z., Zhou, Y., and Ji, X. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *arXiv preprint arXiv:2004.10019*, 2020b.
- Zinkevich, M., Johanson, M., Bowling, M., and Piccione, C. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20:1729–1736, 2007.

A. Discussion on Technical Challenges

In this section, we discuss some technical challenges faced when designing provably efficient algorithms for MGs with general function approximation, which explains why direct extension of existing algorithmic solutions is not enough.

The algorithmic solutions designed for tabular MGs (Bai & Jin, 2020) and linear MGs (Xie et al., 2020) can be viewed as special cases of solving the following sub-problem: in each episode k , given the confidence set of possible functions $f \in \mathcal{C}^k$, find function pair (f^k, g^k) and policy pair (μ^k, ν^k) s.t. for any state s

$$[\mathbb{D}_{\mu^k \times \nu^k}]f^k(s) \geq \max_{f \in \mathcal{C}^k, \mu} [\mathbb{D}_{\mu \times \nu^k} f](s), \quad [\mathbb{D}_{\mu^k \times \nu^k}]g^k(s) \leq \min_{g \in \mathcal{C}^k, \nu} [\mathbb{D}_{\mu^k \times \nu} g](s). \quad (6)$$

Here we omit the dependence on h by considering a single-step special case. This is similar to the contextual bandits problem, but a game version.

Once the above sub-problem is solved, using the fact that the true value function f^* is contained in \mathcal{C}^k , we can bound the duality gap by

$$(V^{\dagger, \nu^k} - V^{\mu^k, \dagger})(s) = [\mathbb{D}_{\mu^{\dagger}(\nu^k) \times \nu^k} f^*](s) - [\mathbb{D}_{\mu^k \times \nu^{\dagger}(\mu^k)} f^*](s) \leq [\mathbb{D}_{\mu^k \times \nu^k} (f^k - g^k)](s),$$

where the summation of the RHS over k can be further bounded by pigeon-hole type of arguments.

Then a natural question is: how can we solve this sub-problem (6)? A helpful condition is optimistic closure, which indeed holds for tabular and linear MGs (Bai & Jin, 2020; Xie et al., 2020). For the max-player version, this means the pointwise upper bound defined by $\bar{f}^k(s, a, b) := \max_{f \in \mathcal{C}^k} f(s, a, b)$ remains in the function class, i.e., $\bar{f}^k \in \mathcal{C}^k \subseteq \mathcal{F}$. The min-player version is similar.

Under this condition, the function pair $(\bar{f}^k, \underline{f}^k)$ is clearly the maximizer/minimizer to the subproblem, and therefore it reduces to finding (μ^k, ν^k) so that

$$\mathbb{D}_{\mu^k \times \nu^k} \bar{f}^k(s) \geq \max_{\mu} \mathbb{D}_{\mu \times \nu^k} \bar{f}^k(s), \quad \mathbb{D}_{\mu^k \times \nu^k} \underline{f}^k(s) \leq \min_{\nu} \mathbb{D}_{\mu^k \times \nu} \underline{f}^k(s).$$

By definition, the solution is a coarse correlated equilibrium.

However, the optimistic closure may not hold in general. Even worse, no solution may exist for the sub-problem (6), as we will see below in a concrete example. In that case, the existing techniques fail and it becomes unclear how to bound the two-sided duality gap. As a result, the general function approximation problem is significantly harder than the special cases mentioned before.

Now we describe a concrete example where sub-problem (6) has no solution when *optimistic closure condition does not hold*. We consider rock-paper-scissors, with one state and three actions for both players, i.e., $|\mathcal{S}| = 1$ and $|\mathcal{A}| = |\mathcal{B}| = 3$. Then each function $f \in \mathcal{F} : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [-2, 2]^3$ is essentially a 3×3 matrix and we will use M to represent such matrices. The matrix corresponding to Nash value Q^* is

$$M^* = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

which describes the true reward associated with different actions. Assume there are 6 other matrices in our confidence set \mathcal{C}^k :

$$M_1 = \begin{pmatrix} 0 & 1.1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_2 = \begin{pmatrix} 0 & 1 & -1.1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_3 = \begin{pmatrix} 0 & 1 & -1 \\ -1.1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix},$$

$$M_4 = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1.1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_5 = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1.1 & -1 & 0 \end{pmatrix}, \quad M_6 = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1.1 & 0 \end{pmatrix}.$$

³We assign the reward $[-2, 2]$ instead of $[0, 1]$ to make the example looks simpler. Everything stated here still hold after scaling.

Now suppose we find a pair of matrices (\bar{M}, \underline{M}) and a pair of policies (μ, ν) which solves sub-problem (6). If we can prove such (μ, ν) must be deterministic, then we get into a contradiction, since whatever (\bar{M}, \underline{M}) is, it is impossible for deterministic policies μ and ν to be the best response of each other. Therefore, in order to show sub-problem (6) has no solution, it suffices to prove (μ, ν) must be deterministic.

Since μ is a solution of sub-problem (6), we must have

$$\mu = \operatorname{argmax}_{\mu'} \max_{M \in \mathcal{C}^k} (\mu')^T M \nu.$$

Since we are maximizing the above quantity, \bar{M} can only take M_1, M_4 or M_5 because the others are dominated. Direct computation gives

$$\begin{aligned} M_1 \nu &= M^* \nu + [0.1\nu[2], 0, 0]^T, \\ M_4 \nu &= M^* \nu + [0, 0.1\nu[3], 0]^T, \\ M_5 \nu &= M^* \nu + [0, 0, 0.1\nu[1]]^T. \end{aligned}$$

Now that μ maximizes $(\mu')^T \bar{M} \nu$, if it is not deterministic, there will be at least two entries in $\bar{M} \nu$ that take the same highest value, say $(\bar{M} \nu)[1] = (\bar{M} \nu)[2] \geq (\bar{M} \nu)[3]$.

If $\bar{M} = M_1$, the above condition is just

$$(\bar{M} \nu)[2] = (M^* \nu)[2] = (M^* \nu)[1] + 0.1\nu[2] = (\bar{M} \nu)[1].$$

Then consider two cases:

- If $\nu[3] > 0$, we can choose $\bar{M} = M_4$ to make $(\bar{M} \nu)[2] = (M^* \nu)[2] + 0.1\nu[3] > (M^* \nu)[2]$. Therefore, the value of $\max(\mu')^T \bar{M} \nu$ is increased, which is contradictory to the maximization condition.
- If $\nu[3] = 0$, the condition $(M^* \nu)[2] = (M^* \nu)[1] + 0.1\nu[2]$ is just $\nu[2] + 0.1\nu[2] = -\nu[1]$, which implies $\nu[2] = \nu[1] = 0$. This is impossible since ν is a probability distribution.

As a result, we prove \bar{M} cannot be M_1 . Similarly, \bar{M} cannot be M_4 or M_5 . We obtain a contradiction! Therefore, μ must be deterministic and by symmetrical argument so is ν . Putting everything together completes the proof. That is, when *optimistic closure condition does not hold*, sub-problem (6) can have no solution.

B. V-type Bellman Eluder Dimension

In this section, we define V-type Bellman-Eluder (VBE) dimension, and provide the corresponding theoretical guarantee for Algorithm 1 with Option II.

B.1. Complexity Measure

With slight abuse of notation, we redefine the following distribution families for VBE dimension. The only difference is that now we use distributions over \mathcal{S} instead of $\mathcal{S} \times \mathcal{A} \times \mathcal{B}$.

1. $\mathcal{D}_\Delta := \{\mathcal{D}_{\Delta, h}\}_{h \in [H]}$, where $\mathcal{D}_{\Delta, h} = \{\delta_s(\cdot) \mid s \in \mathcal{S}\}$, i.e., the collections of probability measures that put measure 1 on a single state.
2. $\mathcal{D}_\mathcal{F} := \{\mathcal{D}_{\mu, h}\}_{h \in [H]}$, where $\mathcal{D}_{\mu, h}$ denotes the collections of all probability measures over \mathcal{S} at the h^{th} step, which can be generated by executing some $(\mu_f, \nu_{f, g})$ with $f, g \in \mathcal{F}$. Here, $\nu_{f, g}$ is the best response to μ_f regarding to Q-value g :

$$\nu_{f, g, h}(s) := \operatorname{argmin}_{\nu \in \Delta_{\mathcal{B}}} \mu_{f, h}(s)^\top g_h(s, \cdot, \cdot) \nu \quad \text{for all } (s, h).$$

To proceed, we introduce the following Bellman residual function defined over the state space \mathcal{S}

$$[\mathcal{E}_h(f, g, w)](s) := \mathbb{E}[(f_h - \mathcal{T}_h^{\mu^g} f_{h+1})(s, a, b) \mid (a, b) \sim \mu_{g,h} \times \nu_{g,w,h}].$$

Equipped with the new definition, we can define V-type BE dimension.

Definition B.1 (V-type Bellman-Eluder dimension). Let $\mathcal{H}_{\mathcal{F}}$ be the function classes of Bellman residual where $\mathcal{H}_{\mathcal{F},h} := \{\mathcal{E}_h(f, g, w) \mid f, g, w \in \mathcal{F}\}$. Then the V-type Bellman-Eluder (VBE) dimension is defined as

$$\dim_{\text{VBE}}(\mathcal{F}, \epsilon) = \max_{h \in [H]} \min_{\mathcal{D} \in \{\mathcal{D}_{\Delta}, \mathcal{D}_{\mathcal{F}}\}} \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_h, \epsilon). \quad (7)$$

We comment that we do not have the online version of V-type BE dimension, because the uniform sampling step in Option II is in general not compatible with the policy of the opponent.

B.2. Theoretical Guarantees

Now we are ready to present the theoretical guarantees.

Theorem B.2 (Regret of GOLF_WITH_EXPLOITER). *Under Assumption 3.1 and 3.2, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$, then with probability at least $1 - \delta$, for all $k \in [K]$, Algorithm 1 with Option II satisfies:*

$$\sum_{t=1}^k V_1^*(s_1) - V_1^{\mu^t, \dagger}(s_1) \leq \mathcal{O}(H\sqrt{|\mathcal{A}||\mathcal{B}|dk\beta}),$$

where $d = \dim_{\text{VBE}}(\mathcal{F}, 1/K)$ is the VBE dimension.

Theorem B.2 provides a \sqrt{K} pseudo-regret guarantee in terms of VBE dimension for Algorithm 1 with Option II. The reason for calling it pseudo-regret is that in each episode, the samples are collected following a combination of (μ^k, ν^k) and uniform sampling, instead of only (μ^k, ν^k) as in Option I. Still, this is enough for applying the pigeonhole principle to derive the sample complexity of GOLF_WITH_EXPLOITER.

Corollary B.3 (Sample complexity of GOLF_WITH_EXPLOITER). *Under Assumption 3.1 and 3.2, there exists an absolute constant c such that for any $\epsilon > 0$, choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ and $\Delta = c'(H\sqrt{|\mathcal{A}||\mathcal{B}|d\beta/K} + \epsilon)$, where $d = \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H)$ is the VBE dimension of \mathcal{F} , then with probability at least $1 - \delta$, the output condition (Line 5) will be satisfied at least once in the first K episodes. Furthermore, the output policy μ^{out} of Algorithm 1 with Option II is $\mathcal{O}(\epsilon + H\sqrt{d(\varepsilon_{\text{real}} + \varepsilon_{\text{comp}})})$ -approximate Nash, if $K \geq \Omega((H^2d/\epsilon^2) \cdot \log(H|\mathcal{F}||\mathcal{G}|d/\epsilon))$.*

Corollary B.3 claims that GOLF_WITH_EXPLOITER can find an $\mathcal{O}(\epsilon)$ -approximate Nash policy for the max-player using at most $\tilde{\mathcal{O}}(H^2|\mathcal{A}||\mathcal{B}| \log(|\mathcal{F}||\mathcal{G}|) \dim_{\text{BE}}(\mathcal{F}, \epsilon)/\epsilon^2)$ samples, which scales linearly in the VBE dimension, the cardinality of the two action sets, and the log-cardinality of the function classes.

C. Proofs for Section 4

In this section, we present the proof of Theorem 4.4, Corollary 4.5, and Theorem 4.8.

The following auxiliary lemma (Lemma 26 in Jin et al. (2021)) will be useful.

Lemma C.1 (Pigeon-hole). *Given a function class \mathcal{G} defined on \mathcal{X} with $|g(x)| \leq C$ for all $(g, x) \in \mathcal{G} \times \mathcal{X}$, and a family of probability measures Π over \mathcal{X} . Suppose sequence $\{g_k\}_{k=1}^K \subset \mathcal{G}$ and $\{\mu_k\}_{k=1}^K \subset \Pi$ satisfy that for all $k \in [K]$, $\sum_{t=1}^{k-1} (\mathbb{E}_{\mu_t}[g_k])^2 \leq \beta$. Then for all $k \in [K]$ and $\omega > 0$,*

$$\sum_{t=1}^k |\mathbb{E}_{\mu_t}[g_t]| \leq \mathcal{O}\left(\sqrt{\dim_{\text{DE}}(\mathcal{G}, \Pi, \omega)\beta k} + \min\{k, \dim_{\text{DE}}(\mathcal{G}, \Pi, \omega)\}C + k\omega\right).$$

Another useful decomposition is the value difference lemma (Lemma 1 in (Jiang et al., 2017)).

Lemma C.2 (value difference). *For any function f s.t. $f_{H+1} = 0$ and policy π ,*

$$V_{f,h}(s) - V_h^\pi(s) = \sum_{h'=h}^H \mathbb{E}_\pi [V_{f,h'}(s_{h'}) - r_{h'}(s_{h'}, a_{h'}, b_{h'}) - V_{f,h'+1}(s_{h'+1}) \mid s_h = s].$$

C.1. Proof of the Online Guarantee

We begin with the proof of Theorem 4.8. The following two concentration lemmas help us upper bound the Bellman residuals. We defer their proofs to Appendix C.3.1 and C.3.2.

Lemma C.3. *Assuming $\|Q_h^* - \mathcal{P}_{\mathcal{F}_h}(Q_h^*)\| \leq \varepsilon_{\text{real}}$ and $\|\mathcal{T}_h f_{h+1} - \mathcal{P}_{\mathcal{G}_h}(\mathcal{T}_h f_{h+1})\| \leq \varepsilon_{\text{comp}}$ for all $f \in \mathcal{F}$ and $h \in [H]$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ with some large absolute constant c in Algorithm 1 with Option I, then with probability at least $1 - \delta$, for all $(k, h) \in [K] \times [H]$, we have*

$$(a) \sum_{i=1}^{k-1} \mathbb{E} \left[\left(f_h^k(s_h, a_h, b_h) - (\mathcal{T}_h f_{h+1}^k)(s_h, a_h, b_h) \right)^2 \mid s_h, a_h, b_h \sim \pi^i \right] \leq \mathcal{O}(\beta),$$

$$(b) \sum_{i=1}^{k-1} \left(f_h^k(s_h^i, a_h^i, b_h^i) - (\mathcal{T}_h f_{h+1}^k)(s_h^i, a_h^i, b_h^i) \right)^2 \leq \mathcal{O}(\beta),$$

where $(s_1^i, a_1^i, b_1^i, \dots, s_H^i, a_H^i, b_H^i)$ denotes the trajectory sampled by following $\pi^i = (\mu^i, \nu^i)$ in the i^{th} episode.

Lemma C.4. *Under the same condition of Lemma C.3, with probability at least $1 - \delta$, we have $\mathcal{P}_{\mathcal{F}}(Q^*) \in \mathcal{C}^k$ for all $k \in [K]$.*

Proof of Theorem 4.8. By Lemma C.4, $\mathcal{P}_{\mathcal{F}}(Q^*) \in \mathcal{C}^k$. Since we choose f^k optimistically and \mathcal{F} satisfies $\varepsilon_{\text{real}}$ -approximate realizability,

$$\begin{aligned} & \sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) \\ & \leq \sum_{k=1}^K \left(V_{\mathcal{P}_{\mathcal{F}}(Q^*), 1}(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) + K\varepsilon_{\text{real}} \\ & \leq \sum_{k=1}^K \left(V_{f^k, 1}(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) + K\varepsilon_{\text{real}} \\ & \stackrel{(i)}{=} \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[V_{f^k, h}(s_h) - r_h(s_h, a_h, b_h) - V_{f^k, h+1}(s_{h+1}) \right] + K\varepsilon_{\text{real}} \\ & \stackrel{(ii)}{=} \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[\min_{\nu} \mathbb{D}_{\mu_h^k \times \nu} f_h^k(s_h) - r_h(s_h, a_h, b_h) - V_{f^k, h+1}(s_{h+1}) \right] + K\varepsilon_{\text{real}} \\ & \leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[\mathbb{D}_{\mu_h^k \times \nu_h^k} f_h^k(s_h) - r_h(s_h, a_h, b_h) - V_{f^k, h+1}(s_{h+1}) \right] + K\varepsilon_{\text{real}} \\ & = \sum_{h=1}^H \sum_{k=1}^K \mathbb{E}_{\pi^k} \left[f_h^k(s_h, a_h, b_h) - r_h(s_h, a_h, b_h) - V_{f^k, h+1}(s_{h+1}) \right] + K\varepsilon_{\text{real}} \\ & \stackrel{(iii)}{\leq} \sum_{h=1}^H \sum_{k=1}^K \left(f_h^k - \mathcal{T}_h f_{h+1}^k \right)(s_h^k, a_h^k, b_h^k) + \mathcal{O} \left(\sqrt{KH \log(KH/\delta)} \right) + K\varepsilon_{\text{real}} \\ & \stackrel{(iv)}{\leq} \mathcal{O} \left(H \sqrt{\beta \cdot \dim_{\text{OBE}}(\mathcal{F}, 1/K) \cdot K} \right) + K\varepsilon_{\text{real}} \leq \mathcal{O} \left(H \sqrt{\beta \cdot \dim_{\text{OBE}}(\mathcal{F}, 1/K) \cdot K} \right). \end{aligned}$$

where (i) is by value difference (Lemma C.2), (ii) is due to $\mu^k = \mu_{f^k}$, (iii) is by Azuma-Hoeffding, and (iv) is by incurring Lemma C.1 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 4.7), $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and Lemma C.3. The final inequality follows from the choice of β . \square

C.2. Proof of the Self-play Guarantee

Proving Theorem 4.4 requires more work because we need to develop guarantees for the sub-routine COMPUTE_EXPLOITER. Similar to Lemma C.3 and C.4, we establish two concentration lemmas below, whose proofs are deferred to Section C.3.3 and C.3.4.

Lemma C.5. Under Assumption 3.1 and 3.2, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ with some large absolute constant c in Algorithm 1 with Option I, then with probability at least $1 - \delta$, for all $(k, h) \in [K] \times [H]$, we have

$$(a) \sum_{i=1}^{k-1} \mathbb{E} \left[\left(\tilde{f}_h^k(s_h, a_h, b_h) - (\mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h, a_h, b_h) \right)^2 \mid s_h, a_h, b_h \sim \pi^i \right] \leq \mathcal{O}(\beta).$$

$$(b) \sum_{i=1}^{k-1} \left(\tilde{f}_h^k(s_h^i, a_h^i, b_h^i) - (\mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^i, a_h^i, b_h^i) \right)^2 \leq \mathcal{O}(\beta),$$

where $(s_1^i, a_1^i, b_1^i, \dots, s_H^i, a_H^i, b_H^i, s_{H+1}^i)$ denotes the trajectory sampled by following $\pi^i = (\mu^i, \nu^i)$ in the i^{th} episode, and \tilde{f}^k denotes the optimistic function computed in sub-routine COMPUTE_EXPLOITER in the k^{th} episode.

Lemma C.6. Under the same condition of Lemma C.5, with probability at least $1 - \delta$, we have $\mathcal{P}_{\mathcal{F}}(Q^{\mu^k, \dagger}) \in \mathcal{C}^{\mu^k}$ for all $k \in [K]$.

Proposition C.7 (approximate best-response regret). Under Assumption 3.1 and 3.2, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$, then with probability at least $1 - \delta$, for all $k \in [K]$,

$$\sum_{t=1}^k \left(V_1^{\mu^t, \nu^t}(s_1) - V_1^{\mu^t, \dagger}(s_1) \right) \leq \mathcal{O}(H\sqrt{k\beta \cdot \dim_{BE}(\mathcal{F}, 1/K)}).$$

Proof. By Lemma C.6, $\mathcal{P}_{\mathcal{F}}(Q^{\mu^t, \dagger}) \in \mathcal{C}^t$. Since we choose f^t optimistically and \mathcal{F} satisfies $\varepsilon_{\text{real}}$ -approximate realizability,

$$\begin{aligned} & \sum_{t=1}^k \left(V_1^{\mu^t, \nu^t}(s_1) - V_1^{\mu^t, \dagger}(s_1) \right) \\ & \leq \sum_{t=1}^k \left(V_1^{\mu^t, \nu^t}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^t \times \nu} \tilde{f}_1^t(s_1) \right) + k\varepsilon_{\text{real}} \\ & = - \sum_{t=1}^k \sum_{h=1}^H \mathbb{E}_{\pi^t} \left[(\tilde{f}_h^t - \mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h, a_h, b_h) \right] + k\varepsilon_{\text{real}} \\ & \leq - \sum_{h=1}^H \sum_{t=1}^k (\tilde{f}_h^t - \mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h^t, a_h^t, b_h^t) + \mathcal{O}(\sqrt{kH \log(KH/\delta)}) + k\varepsilon_{\text{real}}, \end{aligned} \tag{8}$$

where \tilde{f}^t denotes the optimistic function computed in sub-routine COMPUTE_EXPLOITER in the t^{th} episode and the equality follows from value difference Lemma C.2.

By Lemma C.5 (b), we have for all $k \in [K]$

$$\sum_{t=1}^{k-1} \left[\tilde{f}_h^k(s_h^t, a_h^t, b_h^t) - (\mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^t, a_h^t, b_h^t) \right]^2 \leq \mathcal{O}(\beta).$$

So we can apply Lemma C.1 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$, $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and obtain

$$\sum_{t=1}^k \left| \tilde{f}_h^t(s_h^t, a_h^t, b_h^t) - (\mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h^t, a_h^t, b_h^t) \right| \leq \sqrt{\beta \cdot \dim_{DE}(\mathcal{H}_{\mathcal{F}, h}, \mathcal{D}_{\Delta, h}, 1/K)} \cdot k.$$

Plugging this inequality back into (8) gives us the first upper bound.

We can also invoke Lemma C.5 (a) with Jensen's inequality, and obtain for all $k \in [K]$

$$\sum_{t=1}^{k-1} \left(\mathbb{E}_{\pi^t} \left[(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h, a_h, b_h) \right] \right)^2 \leq \mathcal{O}(\beta).$$

Again, we can apply Lemma C.1 with $\mathcal{G} = \mathcal{H}_{\mathcal{F},h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 4.3), $\Pi = \mathcal{D}_{\mathcal{F},h}$, $\epsilon = 1/K$, and obtain

$$\sum_{t=1}^k \left| \mathbb{E}_{\pi^t} \left[(\tilde{f}_h^t - \mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h^t, a_h^t, b_h^t) \right] \right| \leq \sqrt{\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_{\mathcal{F},h}, 1/K) \cdot k}.$$

Plugging this inequality back into (8) gives us the second upper bound.

Combining the two upper bounds and noticing that $k\varepsilon_{\text{real}} \leq \mathcal{O}(H\sqrt{k\beta \cdot \dim_{\text{BE}}(\mathcal{F}, 1/K)})$ (because of the choice of β) conclude the proof. \square

Now we are ready to prove the regret guarantee.

Proof of Theorem 4.4. The regret can be decomposed into two terms

$$\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) = \underbrace{\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right)}_{(A)} + \underbrace{\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right)}_{(B)}.$$

By Theorem 4.8, where the opponent's policy $\{\nu^k\}_{k=1}^K$ can be arbitrary, (A) can be upper bounded by

$$\sum_{k=1}^K V_1^*(s_1) - V_1^{\pi^k}(s_1) \leq \mathcal{O}\left(H\sqrt{K\beta \cdot \dim_{\text{BE}}(\mathcal{F}, 1/K)}\right).$$

By Proposition C.7, (B) can be upper bounded by

$$\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \leq \mathcal{O}\left(H\sqrt{K\beta \cdot \dim_{\text{BE}}(\mathcal{F}, 1/K)}\right).$$

Combining the two inequalities concludes the proof. \square

By the standard online-to-batch reduction, we can also derive the sample complexity guarantee.

Proof of Corollary 4.5. We proceed as in the proof of Theorem 4.4 but take $\omega = \epsilon/H$ (instead of $\omega = 1/K$) every time we incur Lemma C.1.

Theorem 4.8 essentially shows

$$\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) \leq \sum_{k=1}^K \left(V_{f^k, 1}(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) \leq \mathcal{O}(H\sqrt{K\beta \cdot \dim_{\text{BE}}(\mathcal{F}, \epsilon/H)} + K\epsilon),$$

where we have used the fact that $\dim_{\text{OBE}}(\mathcal{F}, \epsilon/H) \leq \dim_{\text{BE}}(\mathcal{F}, \epsilon/H)$.

Proposition C.7 essentially shows

$$\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \leq \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1) \right) \leq \mathcal{O}(H\sqrt{K\beta \cdot \dim_{\text{BE}}(\mathcal{F}, \epsilon/H)} + K\epsilon).$$

Comparing with the form in Theorem 4.4, now we have an additional $\mathcal{O}(K\epsilon)$ term, since we take $\omega = \epsilon/H$ when incurring Lemma C.1.

Putting them together and noticing by definition $\bar{V}^k = V_{f^k,1}(s_1)$ and $\underline{V}^k = \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1)$, we can get

$$\frac{1}{K} \sum_{k=1}^K [\bar{V}^k - \underline{V}^k] \leq \mathcal{O}\left(\frac{H}{K} \sqrt{\dim_{\text{BE}}(\mathcal{F}, \epsilon/H) K \beta} + \epsilon\right),$$

with probability at least $1 - \delta$, where $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$.

By pigeonhole prinple, there must exist some k s.t. $\bar{V}^k - \underline{V}^k \leq \Delta = c'(H\sqrt{\dim_{\text{BE}}(\mathcal{F}, \epsilon/H)\beta/K} + \epsilon)$. Therefore, the output condition must be satisfied by some $k \in [K]$.

To make the right hand side order $\mathcal{O}(\epsilon + H\sqrt{d}(\varepsilon_{\text{real}} + \varepsilon_{\text{comp}}))$, it suffices to take

$$K \geq \Omega((H^2 d / \epsilon^2) \cdot \log(H|\mathcal{F}||\mathcal{G}|d/\epsilon)),$$

where $d = \dim_{\text{BE}}(\mathcal{F}, \epsilon/H)$. □

C.3. Proofs of the Concentration Arguments

C.3.1. PROOF OF LEMMA C.3

Proof. We prove inequality (b) first.

Consider a fixed (k, h, f) tuple. For notational simplicity, we denote $z_h^t := (s_h^t, a_h^t, b_h^t)$. Let

$$X_t(h, f) := (f_h(z_h^t) - r_h^t - V_{f,h+1}(s_{h+1}^t))^2 - (\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h f_{h+1})(z_h^t) - r_h^t - V_{f,h+1}(s_{h+1}^t))^2$$

and $\mathfrak{F}_{t,h}$ be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_h^t\}$. We have

$$\begin{aligned} \mathbb{E}[X_t(h, f) \mid \mathfrak{F}_{t,h}] &= [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 - [(\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h f_{h+1}) - \mathcal{T}_h f_{h+1})(z_h^t)]^2 \\ &\geq [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 - \varepsilon_{\text{comp}}^2 \end{aligned}$$

and

$$\begin{aligned} \text{Var}[X_t(h, f) \mid \mathfrak{F}_{t,h}] &\leq \mathbb{E}[(X_t(h, f))^2 \mid \mathfrak{F}_{t,h}] \leq 36[(\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h f_{h+1}) - f_h)(z_h^t)]^2 \\ &\leq \mathcal{O}([(\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h f_{h+1}) - \mathcal{T}_h f_{h+1})(z_h^t)]^2 + [(\mathcal{T}_h f_{h+1} - f_h)(z_h^t)]^2) \\ &\leq \mathcal{O}([(\mathcal{T}_h f_{h+1} - f_h)(z_h^t)]^2 + \varepsilon_{\text{comp}}^2). \end{aligned}$$

By Freedman's inequality, we have, with probability at least $1 - \delta$,

$$\begin{aligned} &\sum_{t=1}^k [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 - k\varepsilon_{\text{comp}}^2 - \sum_{t=1}^k X_t(h, f) \\ &\leq \mathcal{O}\left(\sqrt{\log(1/\delta) \left(\sum_{t=1}^k [(\mathcal{T}_h f_{h+1} - f_h)(z_h^t)]^2 + k\varepsilon_{\text{comp}}^2\right)} + \log(1/\delta)\right). \end{aligned}$$

Now taking a union bound for all $(k, h, f) \in [K] \times [H] \times \mathcal{F}$, we obtain that with probability at least $1 - \delta$, for all $(k, h, f) \in [K] \times [H] \times \mathcal{F}$

$$\begin{aligned} &\sum_{t=1}^k [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 - k\varepsilon_{\text{comp}}^2 - \sum_{t=1}^k X_t(h, f) \\ &\leq \mathcal{O}\left(\sqrt{\iota \left(\sum_{t=1}^k [(\mathcal{T}_h f_{h+1} - f_h)(z_h^t)]^2 + k\varepsilon_{\text{comp}}^2\right)} + \iota\right). \end{aligned} \tag{9}$$

where $\iota = \log(HK|\mathcal{F}|/\delta)$. From now on, we will do all the analysis conditioning on this event being true.

Consider an arbitrary pair $(h, k) \in [H] \times [K]$. By the definition of \mathcal{C}^k and Assumption 3.2

$$\begin{aligned}
 & \sum_{t=1}^{k-1} X_t(h, f^k) \\
 &= \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 - \sum_{t=1}^{k-1} (\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h f_{h+1}^k)(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 \\
 &\leq \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 - \inf_{g \in \mathcal{G}_h} \sum_{t=1}^{k-1} (g(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 \leq \beta.
 \end{aligned} \tag{10}$$

Putting (9) and (10) together, we obtain

$$\sum_{t=1}^{k-1} [(f_h^k - \mathcal{T}_h f_{h+1}^k)(z_h^t)]^2 \leq \mathcal{O}(\beta + \iota + k\varepsilon_{\text{comp}}^2),$$

which concludes the proof of inequality (b).

To prove inequality (a), we only need to redefine $\mathfrak{F}_{t,h}$ to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1}$ and then repeat the arguments above verbatim. \square

C.3.2. PROOF OF LEMMA C.4

Proof. We will continue to use the notations defined in the proof of Lemma C.3. To further simplify notations, we denote

$$\hat{Q} := \mathcal{P}_{\mathcal{F}}(Q^*) \quad \text{and}$$

Consider a fixed tuple $(k, h, g) \in [K] \times [H] \times \mathcal{G}$. Let

$$W_t(h, g) := (g_h(z_h^t) - r_h^t - V_{\hat{Q}, h+1}(s_{h+1}^t))^2 - (\hat{Q}_h(z_h^t) - r_h^t - V_{\hat{Q}, h+1}(s_{h+1}^t))^2$$

and $\mathfrak{F}_{t,h}$ be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_h^t\}$. We have

$$\begin{aligned}
 \mathbb{E}[W_t(h, g) \mid \mathfrak{F}_{t,h}] &= \left[(g_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2 - \left[(\hat{Q}_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2 \\
 &\geq \left[(g_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2 - \varepsilon_{\text{real}}^2
 \end{aligned}$$

and

$$\begin{aligned}
 \text{Var}[X_t(h, g) \mid \mathfrak{F}_{t,h}] &\leq \mathbb{E}[(X_t(h, g))^2 \mid \mathfrak{F}_{t,h}] \leq 36[(\hat{Q}_h - g_h)(z_h^t)]^2 \\
 &\leq \mathcal{O} \left([(\hat{Q}_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t)]^2 + [(\mathcal{T}_h \hat{Q}_{h+1} - g_h)(z_h^t)]^2 \right) \\
 &\leq \mathcal{O} \left([(\mathcal{T}_h \hat{Q}_{h+1} - g_h)(z_h^t)]^2 + \varepsilon_{\text{comp}}^2 \right).
 \end{aligned}$$

By Freedman's inequality, we have, with probability at least $1 - \delta$,

$$\begin{aligned}
 & - \sum_{t=1}^k W_t(h, g) + \sum_{t=1}^k \left[(g_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2 - k\varepsilon_{\text{real}}^2 \\
 &\leq \mathcal{O} \left(\sqrt{\log(1/\delta) \left(\sum_{t=1}^k [(\mathcal{T}_h \hat{Q}_{h+1} - g_h)(z_h^t)]^2 + k\varepsilon_{\text{comp}}^2 \right)} + \log(1/\delta) \right).
 \end{aligned}$$

Now taking a union bound for all $(k, h, g) \in [K] \times [H] \times \mathcal{G}$, we obtain that with probability at least $1 - \delta$, for all $(k, h, g) \in [K] \times [H] \times \mathcal{G}$

$$\begin{aligned} & - \sum_{t=1}^k W_t(h, g) + \sum_{t=1}^k \left[(g_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2 - k\varepsilon_{\text{real}}^2 \\ & \leq \mathcal{O} \left(\sqrt{\iota \left(\sum_{t=1}^k [(\mathcal{T}_h \hat{Q}_{h+1} - g_h)(z_h^t)]^2 + k\varepsilon_{\text{comp}}^2 \right)} + \iota \right), \end{aligned} \quad (11)$$

where $\iota = \log(HK|\mathcal{G}|/\delta)$. From now on, we will do all the analysis conditioning on this event being true.

Since $\sum_{t=1}^k \left[(g_h - \mathcal{T}_h \hat{Q}_{h+1})(z_h^t) \right]^2$ is nonnegative, by Cauchy-Schwarz inequality, (11) implies for all $(k, h, g) \in [K] \times [H] \times \mathcal{G}$

$$\sum_{t=1}^k W_t(h, g) \geq -\mathcal{O}(\iota + k\varepsilon_{\text{comp}}^2).$$

Plugging in the definition of $W_t(h, g)$ and the choice of β completes the proof. \square

C.3.3. PROOF OF LEMMA C.5

Proof. Recall μ_f denotes the Nash policy of the max-player induced by f . If there exist more than one induced Nash policies, we can break the tie arbitrarily so that μ_f is uniquely defined for each $f \in \mathcal{F}$. Denote $\Pi_{\mathcal{F}} := \{\mu_f \mid f \in \mathcal{F}\}$. We have $|\Pi_{\mathcal{F}}| \leq |\mathcal{F}|$.

We prove inequality (b) first. Consider a fixed tuple $(k, h, f, \mu) \in [K] \times [H] \times \mathcal{F} \times \Pi_{\mathcal{F}}$. Again we denote $z_h^t := (s_h^t, a_h^t, b_h^t)$. Let

$$X_t(h, f, \mu) := \left(f_h(z_h^t) - r_h^t - V_{f, h+1}^\mu(s_{h+1}^t) \right)^2 - \left(\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h^\mu f_{h+1})(z_h^t) - r_h^t - V_{f, h+1}^\mu(s_{h+1}^t) \right)^2$$

and $\mathfrak{F}_{t, h}$ be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_h^t\}$. We have

$$\begin{aligned} \mathbb{E}[X_t(h, f, \mu) \mid \mathfrak{F}_{t, h}] &= [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 - [(\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h^\mu f_{h+1}) - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 \\ &\geq [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 - \varepsilon_{\text{comp}}^2 \end{aligned}$$

and by Cauchy-Schwarz inequality

$$\begin{aligned} \text{Var}[X_t(h, f, \mu) \mid \mathfrak{F}_{t, h}] &\leq \mathbb{E}[(X_t(h, f, \mu))^2 \mid \mathfrak{F}_{t, h}] \\ &\leq 36 [(f_h - \mathcal{P}_{\mathcal{G}}(\mathcal{T}_h^\mu f_{h+1}))(z_h^t)]^2 \leq \mathcal{O} \left([(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 + \varepsilon_{\text{comp}}^2 \right). \end{aligned}$$

By Freedman's inequality, we have, with probability at least $1 - \delta$,

$$\begin{aligned} & \sum_{t=1}^k [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 - k\varepsilon_{\text{comp}}^2 - \sum_{t=1}^k X_t(h, f, \mu) \\ & \leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \left([(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 + \varepsilon_{\text{comp}}^2 \right)} + \log(1/\delta) \right). \end{aligned}$$

Now taking a union bound for all $(k, h, f, \mu) \in [K] \times [H] \times \mathcal{F} \times \Pi_{\mathcal{F}}$, we obtain that with probability at least $1 - \delta$, for all $(k, h, f, \mu) \in [K] \times [H] \times \mathcal{F} \times \Pi_{\mathcal{F}}$

$$\begin{aligned} & \sum_{t=1}^k [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 - k\varepsilon_{\text{comp}}^2 - \sum_{t=1}^k X_t(h, f, \mu) \\ & \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k \left([(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 + \varepsilon_{\text{comp}}^2 \right)} + \iota \right), \end{aligned} \quad (12)$$

where $\iota = \log(HK|\mathcal{F}|/\delta)$. From now on, we will do all the analysis conditioning on this event being true.

Consider an arbitrary pair $(h, k) \in [H] \times [K]$. By the definition of \mathcal{C}^k and Assumption 3.2

$$\begin{aligned}
 & \sum_{t=1}^{k-1} X_t(h, f^k, \mu^k) \\
 &= \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 - \sum_{t=1}^{k-1} (\mathcal{P}_{\mathcal{G}}(\mathcal{T}_h^{\mu^k} f_{h+1}^k)(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 \\
 &\leq \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 - \inf_{g \in \mathcal{G}_h} \sum_{t=1}^{k-1} (g(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 \leq \beta.
 \end{aligned} \tag{13}$$

Putting (12) and (13) together, we obtain

$$\sum_{t=1}^{k-1} [(f_h^k - \mathcal{T}_h^{\mu^k} f_{h+1}^k)(z_h^t)]^2 \leq \mathcal{O}(\beta + \iota + k\varepsilon_{\text{comp}}^2),$$

which concludes the proof of inequality (b).

To prove inequality (a), we only need to redefine $\mathfrak{F}_{t,h}$ to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1}$ and then repeat the arguments above verbatim. \square

C.3.4. PROOF OF LEMMA C.6

Proof. Recall μ_f denotes the Nash policy of the max-player induced by f . If there exist more than one induced Nash policies, we can break the tie arbitrarily so that μ_f is uniquely defined for each $f \in \mathcal{F}$. Denote $\Pi_{\mathcal{F}} := \{\mu_f \mid f \in \mathcal{F}\}$. We have $|\Pi_{\mathcal{F}}| \leq |\mathcal{F}|$.

Consider a fixed tuple $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$. Let

$$W_t(h, g, \mu) := \left(g_h(z_h^t) - r_h^t - Q_{h+1}^{\mu, \dagger}(s_{h+1}^t) \right)^2 - \left(\mathcal{P}_{\mathcal{F}}(Q_h^{\mu, \dagger})(z_h^t) - r_h^t - Q_{h+1}^{\mu, \dagger}(s_{h+1}^t) \right)^2$$

and $\mathfrak{F}_{t,h}$ be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_h^t\}$. We have

$$\begin{aligned}
 \mathbb{E}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] &= \left[\left(g_h - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 - \left[\left(\mathcal{P}_{\mathcal{F}}(Q_h^{\mu, \dagger}) - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 \\
 &\geq \left[\left(g_h - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 - \varepsilon_{\text{real}}^2
 \end{aligned}$$

and by Cauchy-Schwarz inequality

$$\begin{aligned}
 \text{Var}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] &\leq \mathbb{E}[(W_t(h, g, \mu))^2 \mid \mathfrak{F}_{t,h}] \\
 &\leq 36 \left[\left(g_h - \mathcal{P}_{\mathcal{F}}(Q_h^{\mu, \dagger}) \right) (z_h^t) \right]^2 \leq \mathcal{O} \left(\left[\left(g_h - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 + \varepsilon_{\text{real}}^2 \right).
 \end{aligned}$$

By Freedman's inequality, we have, with probability at least $1 - \delta$,

$$\begin{aligned}
 & - \sum_{t=1}^k W_t(h, g, \mu) + \sum_{t=1}^k \left[\left(g_h - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 - k\varepsilon_{\text{real}}^2 \\
 &\leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \left(\left[\left(g_h - Q_h^{\mu, \dagger} \right) (z_h^t) \right]^2 + \varepsilon_{\text{real}}^2 \right)} + \log(1/\delta) \right).
 \end{aligned}$$

Now taking a union bound for all $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$, we obtain that with probability at least $1 - \delta$, for all $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$

$$\begin{aligned} & - \sum_{t=1}^k W_t(h, g, \mu) + \sum_{t=1}^k \left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2 - k\varepsilon_{\text{real}}^2 \\ & \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k \left(\left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2 + \varepsilon_{\text{real}}^2 \right)} + \iota \right), \end{aligned} \quad (14)$$

where $\iota = \log(HK|\mathcal{G}||\mathcal{F}|/\delta)$. From now on, we will do all the analysis conditioning on this event being true.

Since $\sum_{t=1}^k [(g_h - Q_h^{\mu, \dagger})(z_h^t)]^2$ is nonnegative, (14) implies for all $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$

$$\sum_{t=1}^k W_t(h, g, \mu) \geq -\mathcal{O}(\iota + k\varepsilon_{\text{real}}^2).$$

By choosing $\mu = \mu^k$, we have for all $(k, h) \in [K] \times [H]$

$$\sum_{t=1}^{k-1} \left(\mathcal{P}_{\mathcal{F}}(Q_h^{\mu^k, \dagger})(z_h^t) - r_h^t - Q_{h+1}^{\mu^k, \dagger}(s_{h+1}^t) \right)^2 \leq \inf_{g \in \mathcal{G}_h} \sum_{t=1}^{k-1} \left(g(z_h^t) - r_h^t - Q_{h+1}^{\mu^k, \dagger}(s_{h+1}^t) \right)^2 + \mathcal{O}(\iota + k\varepsilon_{\text{real}}^2).$$

We conclude the proof by recalling $\beta = \Theta(\log(HK|\mathcal{G}||\mathcal{F}|/\delta) + k\varepsilon_{\text{real}}^2 + k\varepsilon_{\text{comp}}^2)$. \square

D. Proofs for Section 5

In this section, we will first generalize Bellman rank (Jiang et al., 2017) to the setting of Markov Game. Then we show any problems of low Bellman rank also have low BE dimension. Finally, we prove all the examples in Section 5 have low Bellman rank, and thus low BE dimension.

Denote $\Pi_{\mathcal{F}} := \{(\mu_f, \nu_{f,g}) \mid f, g \in \mathcal{F}\}$, which is exactly the policy class that induces $\mathcal{D}_{\mathcal{F}}$.

Definition D.1 (Q-type ϵ -effective Bellman rank). The ϵ -effective Bellman rank is the minimum integer d so that

- There exists $\phi_h : \Pi_{\mathcal{F}} \rightarrow \mathcal{H}$ and $\psi_h : \mathcal{F} \rightarrow \mathcal{H}$ for each $h \in [H]$ where \mathcal{H} is a separable Hilbert space, such that for any $\pi \in \Pi$, $f, g \in \mathcal{F}$, the average Bellman error

$$\mathbb{E}_{\pi}[(f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(s_h, a_h, b_h)] = \langle \phi_h(\pi), \psi_h(f, g) \rangle_{\mathcal{H}},$$

where $\|\psi_h(f, g)\|_{\mathcal{H}} \leq 1$.

- $d = \max_{h \in [H]} d_{\text{eff}}(\mathcal{X}_h(\phi, \mathcal{F}), \epsilon)$ where $\mathcal{X}_h(\phi, \mathcal{F}) = \{\phi_h(\pi) : \pi \in \Pi_{\mathcal{F}}\}$.

Proposition D.2 (low Q-type effective Bellman rank \subset low Q-type BE dimension). *If an MG with function class \mathcal{F} has Q-type ϵ -effective Bellman rank d , then*

$$\dim_{BE}(\mathcal{F}, \epsilon) \leq d.$$

Proof. Assume there exists $\pi_1, \dots, \pi_n \in \Pi_{\mathcal{F}}$ and $(f^1, g^1), \dots, (f^n, g^n)$ such that for all $t \in [n]$, $\sqrt{\sum_{i=1}^{t-1} (\mathbb{E}_{\pi_i}[f_h^t - \mathcal{T}_h^{\mu_{g^i}} f_{h+1}^t])^2} \leq \epsilon$ and $|\mathbb{E}_{\pi_t}[f_h^t - \mathcal{T}_h^{\mu_{g^t}} f_{h+1}^t]| > \epsilon$. By the definition of effective Bellman rank, this is equivalent to: $\sqrt{\sum_{i=1}^{t-1} (\langle \phi_h(\pi^i), \psi_h(f^i, g^i) \rangle)^2} \leq \epsilon$ and $|\langle \phi_h(\pi^t), \psi_h(f^t, g^t) \rangle| > \epsilon$ for all $t \in [n]$.

For notational simplicity, define $x_i = \phi_h(\pi^i)$ and $\theta_i = \psi_h(f^i, g^i)$. Then

$$\begin{cases} \sum_{i=1}^{t-1} (x_i^{\top} \theta_t)^2 \leq \epsilon^2, & t \in [n] \\ |x_t^{\top} \theta_t| \geq \epsilon, & t \in [n]. \end{cases} \quad (15)$$

Define $\Sigma_t = \sum_{i=1}^{t-1} x_i x_i^\top + \epsilon^2 \cdot \mathbf{I}$. We have

$$\|\theta_t\|_{\Sigma_t} \leq \sqrt{2}\epsilon \implies \epsilon \leq |x_t^\top \theta_t| \leq \|\theta_t\|_{\Sigma_t} \cdot \|x_t\|_{\Sigma_t^{-1}} \leq \sqrt{2}\epsilon \|x_t\|_{\Sigma_t^{-1}}, \quad t \in [n]. \quad (16)$$

As a result, we should have $\|x_t\|_{\Sigma_t}^2 \geq 1/2$ for all $t \in [n]$. Now we can apply the standard log-determinant argument,

$$\sum_{t=1}^n \log(1 + \|x_t\|_{\Sigma_t}^2) = \log \left(\frac{\det(\Sigma_{n+1})}{\det(\Sigma_1)} \right) = \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n x_i x_i^\top \right),$$

which implies

$$0.5 \leq \min_{t \in [n]} \|x_t\|_{\Sigma_t}^2 \leq \exp \left(\frac{1}{n} \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n x_i x_i^\top \right) \right) - 1. \quad (17)$$

Choose $n = d_{\text{eff}}(\mathcal{X}_h(\phi, \mathcal{F}), \epsilon)$ that is the minimum positive integer satisfying

$$\sup_{x_1, \dots, x_n \in \mathcal{X}_h(\phi, \mathcal{F})} \frac{1}{n} \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n x_i x_i^\top \right) \leq e^{-1}. \quad (18)$$

This leads to a contradiction because $0.5 > e^{-1} - 1$. So we must have $n \leq d_{\text{eff}}(\mathcal{X}_h(\psi, \mathcal{F}), \epsilon)$. \square

Similarly, we can define V-type Bellman rank, and prove it is upper bounded by V-type BE dimension.

Definition D.3 (V-type ϵ -effective Bellman rank). The ϵ -effective Bellman rank is the minimum integer d so that

- There exists $\phi_h : \Pi_{\mathcal{F}} \rightarrow \mathcal{H}$ and $\psi_h : \mathcal{F} \rightarrow \mathcal{H}$ for each $h \in [H]$ where \mathcal{H} is a separable Hilbert space, such that for any $\pi \in \Pi$, $f, g, w \in \mathcal{F}$, the average Bellman error

$$\mathbb{E}[(f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(s_h, a_h, b_h) \mid s_h \sim \pi, (a_h, b_h) \sim \mu_{g,h} \times \nu_{g,w,h}] = \langle \phi_h(\pi), \psi_h(f, g, w) \rangle_{\mathcal{H}},$$

where $\|\psi_h(f, g, w)\|_{\mathcal{H}} \leq 1$.

- $d = \max_{h \in [H]} d_{\text{eff}}(\mathcal{X}_h(\phi, \mathcal{F}), \epsilon)$ where $\mathcal{X}_h(\phi, \mathcal{F}) = \{\phi_h(\pi) : \pi \in \Pi_{\mathcal{F}}\}$.

Proposition D.4 (low V-type effective Bellman rank \subset low V-type BE dimension). *If an MG with function class \mathcal{F} has V-type ϵ -effective Bellman rank d , then*

$$\dim_{VBE}(\mathcal{F}, \epsilon) \leq d.$$

We omit the proof here because it is basically the same as Proposition D.2.

D.1. Proofs for Examples

Below, we prove the problems introduced in Section 5 have either low Q-type or low V-type Bellman rank. Therefore, by Proposition D.2 and D.4, they also have low Q-type or low V-type BE dimension.

Proof of Proposition 5.1. We have

$$\begin{aligned} & \mathbb{E}[(f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(s_h, a_h, b_h) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \langle \mathbb{P}^\pi[(s_h, a_h, b_h) = \cdot], (f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(\cdot) \rangle, \end{aligned} \quad (19)$$

where the LHS only depends on π , and the RHS only depends on f, g . \square

Proof of Proposition 5.3. Consider two arbitrary $\theta_h, \theta_{h+1} \in B_d(R)$ and $g \in \mathcal{F}$. By self-completeness, there exists $\theta_h^g \in B_d(R)$ so that $\mathcal{T}_h^{\mu_g}(\phi_{h+1}^\top \theta_{h+1}) = \phi_h^\top \theta_h^g$. Therefore, we have

$$\begin{aligned} & \mathbb{E}[(\phi_h^\top \theta_h - \mathcal{T}_h^{\mu_g}(\phi_{h+1}^\top \theta_{h+1}))(s_h, a_h, b_h) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \mathbb{E}[\phi_h(s_h, a_h, b_h)^\top (\theta_h - \theta_h^g) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \langle \mathbb{E}_\pi[\phi_h(s_h, a_h, b_h)], \theta_h - \theta_h^g \rangle, \end{aligned} \quad (20)$$

where the LHS only depends on π , and the RHS only depends on $\theta_h, \theta_{h+1}, g$. \square

Proof of Proposition 5.4. Consider two arbitrary $\theta_h, \theta_{h+1} \in B_{\mathcal{H}}(R)$ and $g \in \mathcal{F}$. By self-completeness, there exists $\theta_h^g \in B_{\mathcal{H}}(R)$ so that $\mathcal{T}_h^{\mu^g}(\phi_{h+1}^\top \theta_{h+1}) = \phi_h^\top \theta_h^g$. Therefore, we have

$$\begin{aligned} & \mathbb{E}[(\phi_h^\top \theta_h - \mathcal{T}_h^{\mu^g}(\phi_{h+1}^\top \theta_{h+1}))(s_h, a_h, b_h) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \mathbb{E}[\phi_h(s_h, a_h, b_h)^\top (\theta_h - \theta_h^g) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \langle \mathbb{E}_\pi[\phi_h(s_h, a_h, b_h)], \theta_h - \theta_h^g \rangle, \end{aligned} \quad (21)$$

where the LHS only depends on π , and the RHS only depends on $\theta_h, \theta_{h+1}, g$. \square

Proof of Proposition 5.5. Let $d(i)$ denote the distribution of state s given $q(s) = i$. For any policy π and $f, g, w \in \mathcal{F}$

$$\begin{aligned} & \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h \sim \pi, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \\ &= \langle \mathbb{P}_h^\pi[q(s_h) = \cdot], \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h \sim d(\cdot), (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \rangle, \end{aligned}$$

where the LHS only depends on π while the RHS only depends on f, g, w . Both of them are m -dimensional. \square

Proof of Proposition 5.6. The case $h = 1$ is trivial. We only need to consider $h \geq 2$. For any policy π and $f, g, w \in \mathcal{F}$

$$\begin{aligned} & \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h \sim \pi, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \\ &= \sum_{z \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}} \sum_s \mathbb{P}^\pi[(s_{h-1}, a_{h-1}, b_{h-1}) = z] \times \mathbb{P}[s_h = s \mid (s_{h-1}, a_{h-1}, b_{h-1}) = z] \\ & \quad \times \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h = s, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \\ &= \sum_{z \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}} \sum_s \mathbb{P}^\pi[(s_{h-1}, a_{h-1}, b_{h-1}) = z] \times \langle \phi_{h-1}(z), \psi_{h-1}(s) \rangle \\ & \quad \times \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h = s, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \\ &= \left\langle \sum_{z \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}} \mathbb{P}^\pi[(s_{h-1}, a_{h-1}, b_{h-1}) = z] \times \phi_{h-1}(z), \right. \\ & \quad \left. \sum_{s \in \mathcal{S}} \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h = s, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \times \psi_{h-1}(s) \right\rangle \\ &= \left\langle \mathbb{E}_\pi[\phi_{h-1}(s_{h-1}, a_{h-1}, b_{h-1})], \right. \\ & \quad \left. \sum_{s \in \mathcal{S}} \mathbb{E}[(f_h - \mathcal{T}^{\mu^g} f_{h+1})(s_h, a_h, b_h) \mid s_h = s, (a_h, b_h) \sim \mu_g \times \nu_{g,w}] \times \psi_{h-1}(s) \right\rangle, \end{aligned}$$

where the LHS only depends on π while the RHS only depends on f, g, w . By the regularization condition of Kernel MGs, the norm of the RHS is bounded by $2R + 1$. \square

E. Proofs for Appendix B

In this section, we present the proof of Theorem B.2 and Corollary B.3. The techniques are basically the same as those in Appendix C. Please notice that whenever coming across $\mathcal{D}_{\mathcal{F}}$ and \mathcal{D}_{Δ} in this section, we use their definitions introduced in Appendix B.

To begin with, we have the following concentration lemma (akin to Lemma C.5 and C.6) for the sub-routine COMPUTE_EXPLOITER under the samples collected with Option II.

Lemma E.1. *Under Assumption 3.1 and 3.2, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ with some large absolute constant c in Algorithm 1 with Option II, then with probability at least $1 - \delta$, for all $(k, h) \in [K] \times [H]$, we have*

$$(a) \sum_{i=1}^{k-1} \mathbb{E} \left[\left((\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h, a, b) \right)^2 \mid s_h \sim \pi^i, (a, b) \sim \text{Uniform}(\mathcal{A} \times \mathcal{B}) \right] \leq \mathcal{O}(\beta),$$

$$(b) \sum_{i=1}^{k-1} \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} \left((\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^i, a, b) \right)^2 \leq \mathcal{O}(|\mathcal{A}||\mathcal{B}|\beta),$$

$$(c) \mathcal{P}_{\mathcal{F}}(Q^{\mu^k, \dagger}) \in \mathcal{C}^{\mu^k},$$

where (i) $\pi^i = (\mu^i, \nu^i)$; (ii) s_h^i denotes the state at step h , collected by following π^i until step h , in the i^{th} outer iteration; (iii) \tilde{f}^k denotes the optimistic function computed in sub-routine COMPUTE_EXPLOITER in the k^{th} outer iteration.

Proof. To prove (a), we only need to redefine the filtration $\mathfrak{F}_{t,h}$ in Appendix C.3.3 to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1}$ where $z_h^i = (s_h^i, a_h^i, b_h^i)$, and repeat the arguments therein verbatim. Similarly, for (b), we only need to redefine $\mathfrak{F}_{t,h}$ in Appendix C.3.3 to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_{h-1}^t, r_{h-1}^t, s_h^t\}$. And the proof of (c) is the same as that of Lemma C.6 in Appendix C.3.4. \square

Proposition E.2 (approximate best-response regret). *Under Assumption 3.1 and 3.2, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$, then with probability at least $1 - \delta$*

$$\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \leq \mathcal{O} \left(H \sqrt{|\mathcal{A}||\mathcal{B}|k\beta \cdot \dim_{\text{VBE}}(\mathcal{F}, 1/K)} \right).$$

Proof. By Lemma E.1, $\mathcal{P}_{\mathcal{F}}(Q^{\mu^k, \dagger}) \in \mathcal{C}^{\mu^k}$. Since COMPUTE_EXPLOITER chooses \tilde{f}^k optimistically, value difference lemma (Lemma C.2) gives

$$\begin{aligned} & \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \\ & \leq \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1) \right) + k\varepsilon_{\text{real}} \\ & = - \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right] + k\varepsilon_{\text{real}} \\ & \leq - \sum_{h=1}^H \sum_{k=1}^K \mathbb{E} \left[(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^k, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right] + \sqrt{KH \log(KH/\delta)} + k\varepsilon_{\text{real}}, \end{aligned}$$

where the expectation in the second last line is over $s_h \sim \pi^k$ and $(a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h}$, the expectation in the last line is only over $(a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h}$ conditioning on s_h being fixed as s_h^k , and the definition of $\nu_{f,g}$ is introduced in Section 4.

By Jensen's inequality and Lemma E.1 (b), we have for all $k \in [K]$

$$\begin{aligned} & \sum_{t=1}^k \left(\mathbb{E} \left[(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^t, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right] \right)^2 \\ & \leq \sum_{t=1}^k \mathbb{E} \left[\left((\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^t, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right)^2 \right] \\ & \leq \sum_{t=1}^k \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} \left(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k \right)^2 (s_h^t, a, b) \leq \mathcal{O}(|\mathcal{A}||\mathcal{B}|\beta). \end{aligned}$$

So we can apply Lemma C.1 with $\mathcal{G} = \mathcal{H}_{\mathcal{F},h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition B.1), $\Pi = \mathcal{D}_{\Delta,h}$, $\epsilon = 1/K$ and obtain

$$\begin{aligned} & \sum_{k=1}^K \left| \mathbb{E} \left[(\tilde{f}_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h^k, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right] \right| \\ & \leq \mathcal{O} \left(\sqrt{|\mathcal{A}||\mathcal{B}|\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_{\Delta,h}, 1/K) \cdot K} \right). \end{aligned}$$

Similarly by using Lemma E.1 (a), we can show

$$\begin{aligned} & \sum_{k=1}^K \mathbb{E}_{\pi^k} \left[(f_h^k - \mathcal{T}_h^{\mu^k} \tilde{f}_{h+1}^k)(s_h, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \tilde{f}^k, h} \right] \\ & \leq \mathcal{O} \left(\sqrt{|\mathcal{A}||\mathcal{B}|\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F}, h}, \mathcal{D}_{\mathcal{F}, h}, 1/K) \cdot K} \right). \end{aligned}$$

Putting all relations together as in Proposition C.7 and noticing that

$$k\varepsilon_{\text{real}} \leq \mathcal{O} \left(H \sqrt{|\mathcal{A}||\mathcal{B}|k\beta \cdot \dim_{\text{VBE}}(\mathcal{F}, 1/K)} \right)$$

completes the proof. \square

Equipped with the regret guarantee for COMPUTE_EXPLOITER, we are ready to bound the pseudo-regret for the main algorithm GOLF_WITH_EXPLOITER. To begin with, we have the following concentration lemma (akin to Lemma C.3 and C.4) under the samples collected with Option II.

Lemma E.3. *Under Assumption 3.1 and 3.2, if we choose $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$ with some large absolute constant c in Algorithm 1, then with probability at least $1 - \delta$, for all $(k, h) \in [K] \times [H]$, we have*

$$(a) \sum_{i=1}^{k-1} \mathbb{E}[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h, a, b)^2 \mid s_h \sim \pi^i, (a, b) \sim \text{Uniform}(\mathcal{A} \times \mathcal{B})] \leq \mathcal{O}(\beta),$$

$$(b) \sum_{i=1}^{k-1} \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} ((f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h^i, a, b))^2 \leq \mathcal{O}(|\mathcal{A}||\mathcal{B}|\beta),$$

$$(c) \mathcal{P}_{\mathcal{F}}(Q^*) \in \mathcal{C}^k,$$

where s_h^i denotes the state at step h collected following π^i until step h in the i^{th} outer iteration.

Proof. To prove (a), we only need to redefine the filtration $\mathfrak{F}_{t,h}$ in Appendix C.3.1 to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1}$ where $z_h^i = (s_h^i, a_h^i, b_h^i)$, and repeat the arguments therein verbatim. Similarly, for (b), we only need to redefine $\mathfrak{F}_{t,h}$ in Appendix C.3.1 to be the filtration induced by $\{z_1^i, r_1^i, \dots, z_H^i, r_H^i\}_{i=1}^{t-1} \cup \{z_1^t, r_1^t, \dots, z_{h-1}^t, r_{h-1}^t, s_h^t\}$. And the proof of (c) is the same as that of Lemma C.6 in Appendix C.3.2. \square

Proof of Theorem B.2. The regret can be decomposed into two terms

$$\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) = \underbrace{\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right)}_{(A)} + \underbrace{\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right)}_{(B)}.$$

By Proposition E.2, (B) can be upper bounded by with high probability,

$$\sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \leq \mathcal{O} \left(H \sqrt{|\mathcal{A}||\mathcal{B}|K\beta \cdot \dim_{\text{VBE}}(\mathcal{F}, 1/K)} \right).$$

So it remains to control (A). By minicking the proof of Theorem 4.8, we have

$$\begin{aligned}
 & \sum_{k=1}^K [V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1)] \\
 & \leq \sum_{k=1}^K [V_{f^k, 1}(s_1) - V_1^{\mu^k, \nu^k}(s_1)] + K\varepsilon_{\text{real}} \\
 & = \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} [V_{f^k, h}(s_h) - r_h(s_h, a_h, b_h) - V_{f^k, h+1}(s_{h+1})] + K\varepsilon_{\text{real}} \\
 & = \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[\min_{\nu} \mathbb{D}_{\mu_h^k \times \nu} f_h^k(s_h) - \mathcal{T}_h f_{h+1}^k(s_h, a_h, b_h) \right] + K\varepsilon_{\text{real}} \\
 & \leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} [f_h^k(s_h, a_h, b_h) - \mathcal{T}_h f_{h+1}^k(s_h, a_h, b_h)] + K\varepsilon_{\text{real}} \\
 & = \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{\pi^k} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] + K\varepsilon_{\text{real}} \\
 & \leq \sum_{h=1}^H \sum_{k=1}^K \mathbb{E} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h^k, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] + \mathcal{O} \left(\sqrt{KH \log(KH/\delta)} \right) + K\varepsilon_{\text{real}},
 \end{aligned}$$

where the expectation in the second last line is over $s_h \sim \pi^k$ and $(a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h}$, the expectation in the last line is only over $(a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h}$ conditioning on s_h being fixed as s_h^k , and the definition of $\nu_{f, g}$ is introduced in Section 4.

By Jensen's inequality and Lemma E.3 (b), we have for all $k \in [K]$

$$\begin{aligned}
 & \sum_{t=1}^k \left(\mathbb{E} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h^t, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] \right)^2 \\
 & \leq \sum_{t=1}^k \mathbb{E} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)^2(s_h^t, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] \\
 & \leq \sum_{t=1}^k \sum_{(a, b) \in \mathcal{A} \times \mathcal{B}} (f_h^k - \mathcal{T}_h f_{h+1}^k)^2(s_h^t, a, b) \leq \mathcal{O}(|\mathcal{A}||\mathcal{B}|).
 \end{aligned}$$

So we can apply Lemma C.1 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition B.1), $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and obtain

$$\begin{aligned}
 & \sum_{k=1}^K \left| \mathbb{E} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h^k, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] \right| \\
 & \leq \sqrt{|\mathcal{A}||\mathcal{B}|\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F}, h}, \mathcal{D}_{\Delta, h}, 1/K) \cdot K}.
 \end{aligned}$$

Similarly by using Lemma E.3 (a), we can show

$$\begin{aligned}
 & \sum_{k=1}^K \left| \mathbb{E}_{\pi^k} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] \right| \\
 & \leq \sqrt{|\mathcal{A}||\mathcal{B}|\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F}, h}, \mathcal{D}_{\mathcal{F}, h}, 1/K) \cdot K}.
 \end{aligned}$$

Putting all relations together completes the proof. \square

By an standard online-to-batch reduction, we can also prove the sample complexity guarantee.

Proof of Corollary B.3. We proceed as in the proof of Theorem B.2 but take $\omega = \epsilon/H$ (instead of $\omega = 1/K$) every time we incur Lemma C.1.

Theorem B.2 essentially shows

$$\sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) \leq \sum_{k=1}^K \left(V_{f^{k,1}}(s_1) - V_1^{\mu^k, \nu^k}(s_1) \right) \leq \mathcal{O}(H\sqrt{K\beta \cdot |\mathcal{A}||\mathcal{B}| \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H)} + K\epsilon).$$

Proposition E.2 essentially shows

$$\begin{aligned} & \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \\ & \leq \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1) \right) \\ & \leq \mathcal{O}(H\sqrt{K\beta \cdot |\mathcal{A}||\mathcal{B}| \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H)} + K\epsilon). \end{aligned}$$

Comparing with the form in Theorem 4.4, now we have an additional $\mathcal{O}(K\epsilon)$ term, since we take $\omega = \epsilon/H$ when incurring Lemma C.1.

Putting them together and noticing by definition $\bar{V}^k = V_{f^{k,1}}(s_1)$ and $\underline{V}^k = \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1)$, we can get

$$\frac{1}{K} \sum_{k=1}^K [\bar{V}^k - \underline{V}^k] \leq \mathcal{O}\left(\frac{H}{K} \sqrt{|\mathcal{A}||\mathcal{B}| \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H) K\beta} + \epsilon\right),$$

with probability at least $1 - \delta$, where $\beta = c \cdot (\log(KH|\mathcal{F}||\mathcal{G}|/\delta) + K\varepsilon_{\text{comp}}^2 + K\varepsilon_{\text{real}}^2)$.

By pigeonhole prinple, there must exist some k s.t.

$$\bar{V}^k - \underline{V}^k \leq \Delta = c'(H\sqrt{|\mathcal{A}||\mathcal{B}| \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H)\beta/K} + \epsilon).$$

Therefore, the output condition must be satisfied by some $k \in [K]$.

To make the right hand side order $\mathcal{O}(\epsilon + H\sqrt{d}(\varepsilon_{\text{real}} + \varepsilon_{\text{comp}}))$, it suffices to take

$$K \geq \Omega((H^2 d / \epsilon^2) \cdot \log(H|\mathcal{F}||\mathcal{G}|d/\epsilon)),$$

where $d = \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H)$. □