
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.

2020). Two crucial components that contribute to these successes 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 and 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:

Contents

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 and 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.,

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

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, this paper makes 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 between exploration and exploitation.

¹Princeton University, ²Massachusetts Institute of Technology. Correspondence to: Qinghua Liu <qinghual@princeton.edu>.

- 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.

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 and 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 C for more discussions). In contrast, our exploiter framework provides a completely new approach to choose behavior policies, and does not rely on such strong assumptions.

Paper Organization Due to space limit, we postpone related works and preliminaries to Appendix A and B, respectively. We introduce BE dimension and our main algorithm in Section 2. Examples of low BE dimension are presented in Appendix E.

2. 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.

2.1. 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 function 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 5-6): 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 7): 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

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

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.

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 $f^k = \operatorname{argmax}_{f \in \mathcal{C}} V_{f,1}(s_1)$.
- 4: $\nu^k \leftarrow \text{COMPUTE_EXPLOITER}(\mathcal{F}, \mathcal{G}, \beta, \mathcal{D}, \mu^k)$.
- 5: **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 $(s_h, a_h, b_h, r_h, s_{h+1})$.
- 6: **Augment** $\mathcal{D}_h = \mathcal{D}_h \cup \{(s_h, a_h, r_h, s_{h+1})\}$ for all $h \in [H]$.
- 7: **Update**
 $\mathcal{C} = \left\{ f \in \mathcal{F} : \mathcal{L}_{\mathcal{D}_h}(f_h, f_{h+1}) \leq \inf_{g \in \mathcal{G}_h} \mathcal{L}_{\mathcal{D}_h}(g, f_{h+1}) + \beta \text{ for all } h \in [H] \right\}$,
- 8: **Output** μ^{out} , the random policy defined by sampling uniformly at random from $\{\mu^k\}_{k=1}^K$.

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 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

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

2.2. 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 E 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.

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** $\tilde{f} \leftarrow \operatorname{argmin}_{f \in \mathcal{C}} V_{f,1}^\mu(s_1)$.
- 3: **Return** ν^{out} such that
 $\nu_h^{\text{out}}(s) = \operatorname{argmin}_{\nu \in \Delta_{\mathcal{B}}} \mu_h(s)^\top \tilde{f}_h(s, \cdot, \cdot) \nu$ for all s, h .

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

Definition 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 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 and 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 :

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

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 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*

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

We remark that the Bellman residual functions used in Definition 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 D.

2.3. Theoretical Guarantees

Now we are ready to present the theoretical guarantees.

Theorem 4 (Regret of Algorithm 1). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$, 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 $d = \dim_{BE}(\mathcal{F}, 1/K)$ is the BE dimension.*

Theorem 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 standard online-to-batch reduction (Lattimore and Szepesvári, 2020), we also derive a sample complexity guarantee for GOLF_WITH_EXPLOITER.

Corollary 5 (Sample complexity of Algorithm 1). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\epsilon > 0$, choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|)$, the output policy μ^{out} of Algorithm 1 with Option I is $\mathcal{O}(\epsilon)$ -approximate Nash with probability at least $1/2$, if $K \geq \Omega((H^2 d / \epsilon^2) \cdot \log(H|\mathcal{F}||\mathcal{G}|d/\epsilon))$, where $d = \dim_{BE}(\mathcal{F}, \epsilon/H)$ is the BE dimension of \mathcal{F} .*

Corollary 5 asserts that GOLF_WITH_EXPLOITER can find an $\mathcal{O}(\epsilon)$ -approximate Nash policy for the max-player using at most $\tilde{\mathcal{O}}(H^2 \log(|\mathcal{F}||\mathcal{G}|) \dim_{BE}(\mathcal{F}, \epsilon)/\epsilon^2)$ samples, which scales linearly in the BE dimension and the log-cardinality of the function classes. 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 D.

2.4. 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 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 6 (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:*

$$\dim_{OBE}(\mathcal{F}, \epsilon) = \max_{h \in [H]} \dim_{DE}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_{\Delta,h}, \epsilon). \quad (2)$$

Compared to the BE dimension in Definition 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 7 (Regret against adversarial opponents). *Assuming $Q_h^* \in \mathcal{F}_h$ and $\mathcal{T}_h \mathcal{F}_{h+1} \subset \mathcal{G}_h$ for all $h \in [H]$, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$, then with probability at least $1 - \delta$, for all $k \in [K]$, Algorithm 1 with Option I 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 $d = \dim_{OBE}(\mathcal{F}, 1/K)$ is the online BE dimension.*

Theorem 7 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 7 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.

References

- Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *arXiv preprint arXiv:2102.00815*, 2021.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587): 484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michael Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354, 2019.
- OpenAI. Openai five. <https://blog.openai.com/openai-five/>, 2018.
- Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Bowen Baker, Ingmar Kanitscheider, Todor Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. Emergent tool use from multi-agent autocurricula. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkxpxJBKwS>.
- Yu Bai and Chi Jin. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR, 2020.
- Yu Bai, Chi Jin, and Tiancheng Yu. Near-optimal reinforcement learning with self-play. *arXiv preprint arXiv:2006.12007*, 2020.
- Qinghua Liu, Tiancheng Yu, Yu Bai, and Chi Jin. A sharp analysis of model-based reinforcement learning with self-play. *arXiv preprint arXiv:2010.01604*, 2020.
- Qiaomin Xie, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In *Conference on Learning Theory*, pages 3674–3682. PMLR, 2020.
- Zixiang Chen, Dongruo Zhou, and Quanquan Gu. Almost optimal algorithms for two-player markov games with linear function approximation. *arXiv preprint arXiv:2102.07404*, 2021.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- Ruosong Wang, Ruslan Salakhutdinov, and Lin F Yang. Provably efficient reinforcement learning with general value function approximation. *arXiv preprint arXiv:2005.10804*, 2020.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. Bridging exploration and general function approximation in reinforcement learning: Provably efficient kernel and neural value iterations. *arXiv preprint arXiv:2011.04622*, 2020.
- Simon S Du, Sham M Kakade, Jason D Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in rl. *arXiv preprint arXiv:2103.10897*, 2021.
- Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- Csaba Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pages 2256–2264, 2013.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed

- cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*, 2017.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 4295–4304. PMLR, 2018.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Hepeng Li and Haibo He. Multi-agent trust region policy optimization. *arXiv preprint arXiv:2010.07916*, 2020.
- Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of mappo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021a.
- Michael L Littman. Friend-or-foe q-learning in general-sum games. In *ICML*, volume 1, pages 322–328, 2001.
- Junling Hu and Michael P Wellman. Nash q-learning for general-sum stochastic games. *Journal of machine learning research*, 4(Nov):1039–1069, 2003.
- Thomas Dueholm Hansen, Peter Bro Miltersen, and Uri Zwick. 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.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. *arXiv e-prints*, pages arXiv–2006, 2020.
- Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Online reinforcement learning in stochastic games. *arXiv preprint arXiv:1712.00579*, 2017.
- Zeyu Jia, Lin F Yang, and Mengdi Wang. Feature-based q-learning for two-player stochastic games. *arXiv preprint arXiv:1906.00423*, 2019.
- Aaron Sidford, Mengdi Wang, Lin Yang, and Yinyu Ye. Solving discounted stochastic two-player games with near-optimal time and sample complexity. In *International Conference on Artificial Intelligence and Statistics*, pages 2992–3002. PMLR, 2020.
- Kaiqing Zhang, Sham M Kakade, Tamer Başar, and Lin F Yang. Model-based multi-agent rl in zero-sum markov games with near-optimal sample complexity. *arXiv preprint arXiv:2007.07461*, 2020a.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive markov games. *arXiv preprint arXiv:2102.04540*, 2021.
- Yi Tian, Yuanhao Wang, Tiancheng Yu, and Suvrit Sra. On-line learning in unknown markov games. *arXiv preprint arXiv:2010.15020*, 2021.
- Tiancheng Yu, Yi Tian, Jingzhao Zhang, and Suvrit Sra. Provably efficient algorithms for multi-objective competitive rl. *arXiv preprint arXiv:2102.03192*, 2021b.
- Ronen I Brafman and Moshe Tennenholtz. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3 (Oct):213–231, 2002.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. *arXiv preprint arXiv:1703.05449*, 2017.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *arXiv preprint arXiv:2004.10019*, 2020b.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143, 2020.
- Yining Wang, Ruosong Wang, Simon S Du, and Akshay Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.

- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*, 2019.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. *arXiv preprint arXiv:2003.00153*, 2020a.
- Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in Neural Information Processing Systems*, 33, 2020.
- Gergely Neu and Ciara Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- Ian Osband and Benjamin Van Roy. Model-based reinforcement learning and the eluder dimension. In *Advances in Neural Information Processing Systems*, pages 1466–1474, 2014.
- Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*, pages 2898–2933, 2019.
- Dylan J Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020.
- Kefan Dong, Jian Peng, Yining Wang, and Yuan Zhou. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. In *Conference on Learning Theory*, pages 1554–1557. PMLR, 2020.
- Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and economic behavior*, 4(4):528–552, 1992.
- Andrew Gilpin and Tuomas Sandholm. Finding equilibria in large sequential games of imperfect information. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 160–169, 2006.
- Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20:1729–1736, 2007.
- Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- Andrea Celli, Alberto Marchesi, Gabriele Farina, and Nicola Gatti. No-regret learning dynamics for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems*, 33, 2020.
- Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.
- Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. *arXiv preprint arXiv:1602.02722*, 2016.
- Gellert Weisz, Philip Amortila, and Csaba Szepesvári. Exponential lower bounds for planning in mdps with linearly-realizable optimal action-value functions. *arXiv preprint arXiv:2010.01374*, 2020.

A. 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 and 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 and Wellman, 2003](#); [Hansen et al., 2013](#); [Wei et al., 2020](#)), or under strong reachability conditions such as simulators ([Wei et al., 2017](#); [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 and 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.

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 C 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. Tabular MDP has been studied thoroughly in recent years ([Brafman and Tennenholtz, 2002](#); [Jaksch et al., 2010](#); [Dann and Brunskill, 2015](#); [Azar et al., 2017](#); [Zanette and 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 and Pike-Burke, 2020](#)).

For general function approximation, there are two common measures of complexity: Eluder dimension ([Osband and 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 and Megiddo, 1992](#); [Gilpin and Sandholm, 2006](#); [Zinkevich et al., 2007](#); [Brown and 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.

B. 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

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

$\mathbb{P}_h(\cdot \mid s, a, b)$ 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:

$$Q_h^{\mu, \nu}(s, a, b) = [r_h + \mathbb{P}_h V_{h+1}^{\mu, \nu}](s, a, b), \quad V_h^{\mu, \nu}(s) = [\mathbb{D}_{\mu_h \times \nu_h} Q_h^{\mu, \nu}](s)$$

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

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 and Vrieze, 2012) that there exists policies μ^*, ν^* that are optimal against the best responses of the opponents, i.e.,

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

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

$$\sup_{\mu} \inf_{\nu} V_h^{\mu, \nu}(s) = V_h^{\mu^*, \nu^*}(s) = \inf_{\nu} \sup_{\mu} V_h^{\mu, \nu}(s), \quad \text{for all } (s, h).$$

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.

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

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.

B.1. Function Approximation

This paper considers reinforcement learning with value function approximation. Formally, the learner is provided with a function class $\mathcal{F} = \mathcal{F}_1 \times \cdots \times \mathcal{F}_H$ 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

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

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

$$\begin{cases} (\mathcal{T}_h f)(s, a, b) := r_h(s, a, b) + \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a, b)} V_{f,h+1}(s') \\ (\mathcal{T}_h^\mu f)(s, a, b) := r_h(s, a, b) + \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a, b)} V_{f,h+1}^\mu(s') \end{cases} \quad (3)$$

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 8 (Realizability). $Q_h^* \in \mathcal{F}_h$ and $Q_h^{\mu_f, \dagger} \in \mathcal{F}_h$ for any $f \in \mathcal{F}$ and $h \in [H]$.

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} .

Let $\mathcal{G} = \mathcal{G}_1 \times \cdots \times \mathcal{G}_H$ 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} .

Assumption 9 (Completeness). $\mathcal{T}_h \mathcal{F}_{h+1} \subset \mathcal{G}_h$ and $\mathcal{T}_h^{\mu_f} \mathcal{F}_{h+1} \subset \mathcal{G}_h$ for any $f \in \mathcal{F}$ and $h \in [H]$.

C. 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 and 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

$$[\mathbb{D}_{\mu^k \times \nu^k}]f^k(s) \geq \max_{f \in \mathcal{C}^k, \mu} [\mathbb{D}_{\mu \times \nu^k} f](s), \quad [\mathbb{D}_{\mu^k \times \nu^k}]g^k(s) \leq \min_{g \in \mathcal{C}^k, \nu} [\mathbb{D}_{\mu^k \times \nu} g](s). \quad (4)$$

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 (4)? A helpful condition is optimistic closure, which indeed holds for tabular and linear MGs (Bai and 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 (4), 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 (4) 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 :

$$\begin{aligned} M_1 &= \begin{pmatrix} 0 & 1.1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_2 = \begin{pmatrix} 0 & 1 & -1.1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_3 = \begin{pmatrix} 0 & 1 & -1 \\ -1.1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \\ M_4 &= \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1.1 \\ 1 & -1 & 0 \end{pmatrix}, \quad M_5 = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1.1 & -1 & 0 \end{pmatrix}, \quad M_6 = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1.1 & 0 \end{pmatrix}. \end{aligned}$$

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 $(\overline{M}, \underline{M})$ and a pair of policies (μ, ν) which solves sub-problem (4). If we can prove such (μ, ν) must be deterministic, then we get into a contradiction, since whatever $(\overline{M}, \underline{M})$ 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 (4) has no solution, it suffices to prove (μ, ν) must be deterministic.

Since μ is a solution of sub-problem (4), we must have

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

Since we are maximizing the above quantity, \overline{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 $(\mu')^T \overline{M} \nu$, if it is not deterministic, there will be at least two entries in $\overline{M} \nu$ that take the same highest value, say $(\overline{M} \nu)[1] = (\overline{M} \nu)[2] \geq (\overline{M} \nu)[3]$.

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

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

Then consider two cases:

- If $\nu[3] > 0$, we can choose $\overline{M} = M_4$ to make $(\overline{M} \nu)[2] = (M^* \nu)[2] + 0.1\nu[3] > (M^* \nu)[2]$. Therefore, the value of $\max(\mu')^T \overline{M} \nu$ 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 \overline{M} cannot be M_1 . Similarly, \overline{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 (4) can have no solution.

D. 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.

D.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 10 (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*

$$\dim_{VBE}(\mathcal{F}, \epsilon) = \max_{h \in [H]} \min_{\mathcal{D} \in \{\mathcal{D}_{\Delta}, \mathcal{D}_{\mathcal{F}}\}} \dim_{DE}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_h, \epsilon). \quad (5)$$

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.

D.2. Theoretical Guarantees

Now we are ready to present the theoretical guarantees.

Theorem 11 (Regret of GOLF_WITH_EXPLOITER). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$, 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 $d = \dim_{VBE}(\mathcal{F}, 1/K)$ is the VBE dimension.

Theorem 11 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 online-to-batch argument (Lattimore and Szepesvári, 2020) to derive the sample complexity of GOLF_WITH_EXPLOITER.

Corollary 12 (Sample complexity of GOLF_WITH_EXPLOITER). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\epsilon > 0$, choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|)$, the output policy μ^{out} of Algorithm 1 with Option II is $\mathcal{O}(\epsilon)$ -approximate Nash with probability at least $1/2$, if $K \geq \Omega((H^2|\mathcal{A}||\mathcal{B}|d/\epsilon^2) \cdot \log(H|\mathcal{A}||\mathcal{B}||\mathcal{F}||\mathcal{G}|d/\epsilon))$, where $d = \dim_{VBE}(\mathcal{F}, \epsilon/H)$ is the VBE dimension of \mathcal{F} .*

Corollary 12 claims that GOLF_WITH_EXPLOITER can find an $\mathcal{O}(\epsilon)$ -approximate Nash policy for the max-player using at most $\tilde{\mathcal{O}}(H^2|\mathcal{A}||\mathcal{B}| \log(|\mathcal{F}||\mathcal{G}|) \dim_{BE}(\mathcal{F}, \epsilon)/\epsilon^2)$ samples, which scales linearly in the VBE dimension, the cardinality of the two action sets, and the log-cardinality of the function classes.

E. 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 13 (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*

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

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 14 (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 14 is a special case of Assumption 9 (completeness) by choosing auxiliary function class $\mathcal{G} = \mathcal{F}$. Assumption 14 also implies Assumption 8 (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 15 (Linear FA \subset Low BE dimension). *Under Assumption 14,*

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

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 15 and Corollary 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).⁴

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 16 (Kernel FA \subset Low BE dimension). *Under Assumption 14,*

$$\dim_{BE}(\mathcal{F}, \epsilon) \leq \max_{h \in [H]} d_{\text{eff}}(\mathcal{X}_h, \epsilon/(2R + 1)), \quad (6)$$

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 (6) is upper bounded by $\tilde{\mathcal{O}}(d)$ 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 17 (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*

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

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

⁴Although Corollary 5 is derived for finite function classes, it can be easily extended to the case where the function classes are infinitely large but have bounded ℓ_∞ -covering number. Particularly, it holds for linear function classes, with the log-cardinality replaced by the log-covering number which is $\tilde{\mathcal{O}}(d)$.

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., $\mathbb{P}_h(s' | s, a, b) = \langle \phi_h(s, a, b), \psi_h(s') \rangle_{\mathcal{H}}$. 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 , (a) $\|\theta_h^r\|_{\mathcal{H}} \leq 1$, and (b) $\|\sum_{s \in \mathcal{S}} \mathcal{V}(s) \psi_h(s)\|_{\mathcal{H}} \leq 1$ 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 18 (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 $\mathcal{X}_h = \{\mathbb{E}_{\mu}[\phi_h(s_h, a_h, b_h)] \mid \mu \in \mathcal{D}_{\mathcal{F}, h}\}$.

F. Proofs for Section 2

In this section, we present the proof of Theorem 4, Corollary 5, and Theorem 7.

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

Lemma 19 (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 $\omega > 0$,*

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

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

Lemma 20 (value difference). *For any function f s.t. $f_{H+1} = 0$ and policy π ,*

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

F.1. Proof of the Online Guarantee

We begin with the proof of Theorem 7. The following two concentration lemmas help us upper bound the Bellman residuals. We defer their proofs to Appendix F.3.1 and F.3.2.

Lemma 21. *Assuming $Q_h^* \in \mathcal{F}_h$ and $\mathcal{T}_h \mathcal{F}_{h+1} \subset \mathcal{G}_h$ for all $h \in [H]$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$ 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 22. *Under the same condition of Lemma 21, with probability at least $1 - \delta$, we have $Q^* \in \mathcal{C}^k$ for all $k \in [K]$.*

Proof of Theorem 7. By Lemma 22, $Q^* \in \mathcal{C}^k$. Since we choose f^k optimistically,

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

where (i) is by value difference (Lemma 20), (ii) is due to $\mu^k = \mu_{f^k}$, (iii) is by Azuma-Hoeffding, and (iv) is by incurring Lemma 19 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 6), $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and Lemma 21. \square

F.2. Proof of the Self-play Guarantee

Proving Theorem 4 requires more work because we need to develop guarantees for the sub-routine COMPUTE_EXPLOITER. Similar to Lemma 21 and 22, we establish two concentration lemmas below, whose proofs are deferred to Section F.3.3 and F.3.4.

Lemma 23. *Under Assumption 8 and 9, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$ 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 24. *Under the same condition of Lemma 23, with probability at least $1 - \delta$, we have $Q^{\mu^k, \dagger} \in \mathcal{C}^{\mu^k}$ for all $k \in [K]$.*

Proposition 25 (approximate best-response regret). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$, then with probability at least $1 - \delta$, for all $k \in [K]$,*

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

Proof. Again by optimism and value difference (Lemma 24 and 20),

$$\begin{aligned}
 & \sum_{t=1}^k \left(V_1^{\mu^t, \nu^t}(s_1) - V_1^{\mu^t, \dagger}(s_1) \right) \\
 & \leq \sum_{t=1}^k \left(V_1^{\mu^t, \nu^t}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^t \times \nu} \tilde{f}_1^t(s_1) \right) \\
 & = - \sum_{t=1}^k \sum_{h=1}^H \mathbb{E}_{\pi^t} \left[(\tilde{f}_h^t - \mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h, a_h, b_h) \right] \\
 & \leq - \sum_{h=1}^H \sum_{t=1}^k (\tilde{f}_h^t - \mathcal{T}_h^{\mu^t} \tilde{f}_{h+1}^t)(s_h^t, a_h^t, b_h^t) + \sqrt{kH \log(KH/\delta)},
 \end{aligned} \tag{7}$$

where \tilde{f}^t denotes the optimistic function computed in sub-routine COMPUTE_EXPLOITER in the t^{th} episode.

By Lemma 23 (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 19 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$, $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and obtain

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

Plugging this inequality back into (7) gives us the first upper bound.

We can also invoke Lemma 23 (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 19 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 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 (7) gives us the second upper bound.

Combining the two upper bounds concludes the proof. \square

Now we are ready to prove the regret guarantee.

Proof of Theorem 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 7, 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 25, (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 5. We proceed as in the proof of Theorem 4 but take $\omega = \epsilon/H$ (instead of $\omega = 1/K$) every time we incur Lemma 19, we can get

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

where $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|)$.

To make the right hand side smaller than $\mathcal{O}(\epsilon)$, it suffices to take

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

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

We conclude the proof by recalling the definition of μ^{out} and Markov's inequality. \square

F.3. Proofs of the Concentration Arguments

F.3.1. PROOF OF LEMMA 21

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

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

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

$$\mathbb{E}[X_t(h, f) \mid \mathfrak{F}_{t, h}] = [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2$$

and

$$\text{Var}[X_t(h, f) \mid \mathfrak{F}_{t, h}] \leq \mathbb{E}[(X_t(h, f))^2 \mid \mathfrak{F}_{t, h}] \leq 36[(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 = 36\mathbb{E}[X_t(h, f) \mid \mathfrak{F}_{t, h}].$$

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

$$\left| \sum_{t=1}^k X_t(h, f) - \sum_{t=1}^k \mathbb{E}[X_t(h, f) \mid \mathfrak{F}_{t, h}] \right| \leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \mathbb{E}[X_t \mid \mathfrak{F}_{t, h}] + \log(1/\delta)} \right).$$

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 $(k, h, f) \in [K] \times [H] \times \mathcal{F}$

$$\left| \sum_{t=1}^k X_t(h, f) - \sum_{t=1}^k [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 \right| \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k [(f_h - \mathcal{T}_h f_{h+1})(z_h^t)]^2 + \iota} \right), \quad (8)$$

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 9

$$\begin{aligned}
 & \sum_{t=1}^{k-1} X_t(h, f^k) \\
 &= \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 - \sum_{t=1}^{k-1} ((\mathcal{T}_h f_{h+1}^k)(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 \\
 &\leq \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 - \inf_{g \in \mathcal{G}_h} \sum_{t=1}^{k-1} (g(z_h^t) - r_h^t - V_{f^k, h+1}(s_{h+1}^t))^2 \leq \beta.
 \end{aligned} \tag{9}$$

Putting (8) and (9) together, we obtain

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

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

F.3.2. PROOF OF LEMMA 22

Proof. We will continue to use the notations defined in the proof of Lemma 21.

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

$$W_t(h, g) := (g_h(z_h^t) - r_h^t - V_{h+1}^*(s_{h+1}^t))^2 - (Q_h^*(z_h^t) - r_h^t - V_{h+1}^*(s_{h+1}^t))^2$$

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

$$\mathbb{E}[W_t(h, g) \mid \mathfrak{F}_{t,h}] = [(g_h - Q_h^*)(z_h^t)]^2$$

and

$$\text{Var}[W_t(h, g) \mid \mathfrak{F}_{t,h}] \leq \mathbb{E}[(W_t(h, g))^2 \mid \mathfrak{F}_{t,h}] \leq 36[(g_h - Q_h^*)(z_h^t)]^2 = 36\mathbb{E}[W_t(h, g) \mid \mathfrak{F}_{t,h}].$$

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

$$\left| \sum_{t=1}^k W_t(h, g) - \sum_{t=1}^k \mathbb{E}[W_t(h, g) \mid \mathfrak{F}_{t,h}] \right| \leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \mathbb{E}[W_t \mid \mathfrak{F}_{t,h}] + \log(1/\delta)} \right).$$

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 $(k, h, g) \in [K] \times [H] \times \mathcal{G}$

$$\left| \sum_{t=1}^k W_t(h, g) - \sum_{t=1}^k [(g_h - Q_h^*)(z_h^t)]^2 \right| \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k [(g_h - Q_h^*)(z_h^t)]^2 + \iota} \right), \tag{10}$$

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

Since $\sum_{t=1}^k [(g_h - Q_h^*)(z_h^t)]^2$ is nonnegative, (10) implies for all $(k, h, g) \in [K] \times [H] \times \mathcal{G}$

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

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

F.3.3. PROOF OF LEMMA 23

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

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

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

$$\mathbb{E}[X_t(h, f, \mu) \mid \mathfrak{F}_{t,h}] = [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2$$

and

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

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

$$\begin{aligned} &\left| \sum_{t=1}^k X_t(h, f, \mu) - \sum_{t=1}^k \mathbb{E}[X_t(h, f, \mu) \mid \mathfrak{F}_{t,h}] \right| \\ &\leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \mathbb{E}[X_t(h, f, \mu) \mid \mathfrak{F}_{t,h}] + \log(1/\delta)} \right). \end{aligned}$$

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 $(k, h, f, \mu) \in [K] \times [H] \times \mathcal{F} \times \Pi_{\mathcal{F}}$

$$\left| \sum_{t=1}^k X_t(h, f, \mu) - \sum_{t=1}^k [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 \right| \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k [(f_h - \mathcal{T}_h^\mu f_{h+1})(z_h^t)]^2 + \iota} \right), \quad (11)$$

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 9

$$\begin{aligned} &\sum_{t=1}^{k-1} X_t(h, f^k, \mu^k) \\ &= \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 - \sum_{t=1}^{k-1} ((\mathcal{T}_h^{\mu^k} f_{h+1}^k)(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 \\ &\leq \sum_{t=1}^{k-1} (f_h^k(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 - \inf_{g \in \mathcal{G}_h} \sum_{t=1}^{k-1} (g(z_h^t) - r_h^t - V_{f^k, h+1}^{\mu^k}(s_{h+1}^t))^2 \leq \beta. \end{aligned} \quad (12)$$

Putting (11) and (12) together, we obtain

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

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

F.3.4. PROOF OF LEMMA 24

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

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

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

$$\mathbb{E}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] = \left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2$$

and

$$\text{Var}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] \leq \mathbb{E}[(W_t(h, g, \mu))^2 \mid \mathfrak{F}_{t,h}] \leq 36 \left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2 = 36 \mathbb{E}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}].$$

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

$$\begin{aligned} & \left| \sum_{t=1}^k W_t(h, g, \mu) - \sum_{t=1}^k \mathbb{E}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] \right| \\ & \leq \mathcal{O} \left(\sqrt{\log(1/\delta) \sum_{t=1}^k \mathbb{E}[W_t(h, g, \mu) \mid \mathfrak{F}_{t,h}] + \log(1/\delta)} \right). \end{aligned}$$

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 $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$

$$\left| \sum_{t=1}^k W_t(h, g, \mu) - \sum_{t=1}^k \left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2 \right| \leq \mathcal{O} \left(\sqrt{\iota \sum_{t=1}^k \left[(g_h - Q_h^{\mu, \dagger})(z_h^t) \right]^2} + \iota \right), \quad (13)$$

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, (13) implies for all $(k, h, g, \mu) \in [K] \times [H] \times \mathcal{G} \times \Pi_{\mathcal{F}}$

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

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

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

We conclude the proof by recalling $\beta = \Theta(\log(HK|\mathcal{G}||\mathcal{F}|/\delta))$. □

G. Proofs for Section E

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 E 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 26 (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 27 (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

$$\begin{cases} \sum_{i=1}^{t-1} (x_i^\top \theta_t)^2 \leq \epsilon^2, & t \in [n] \\ |x_t^\top \theta_t| \geq \epsilon, & t \in [n]. \end{cases} \quad (14)$$

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

$$\|\theta_t\|_{\Sigma_t} \leq \sqrt{2}\epsilon \implies \epsilon \leq |x_t^\top \theta_t| \leq \|\theta_t\|_{\Sigma_t} \cdot \|x_t\|_{\Sigma_t^{-1}} \leq \sqrt{2}\epsilon \|x_t\|_{\Sigma_t^{-1}}, \quad t \in [n]. \quad (15)$$

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

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

which implies

$$0.5 \leq \min_{t \in [n]} \|x_t\|_{\Sigma_t^{-1}}^2 \leq \exp \left(\frac{1}{n} \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n x_i x_i^\top \right) \right) - 1. \quad (16)$$

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

$$\sup_{x_1, \dots, x_n \in \mathcal{X}_h(\phi, \mathcal{F})} \frac{1}{n} \log \det \left(\mathbf{I} + \frac{1}{\epsilon^2} \sum_{i=1}^n x_i x_i^\top \right) \leq e^{-1}. \quad (17)$$

This leads to a contradiction because $0.5 > e^{e^{-1}} - 1$. 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 28 (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 29 (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*

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

We omit the proof here because it is basically the same as Proposition 27.

G.1. Proofs for Examples

Below, we prove the problems introduced in Section E have either low Q-type or low V-type Bellman rank. Therefore, by Proposition 27 and 29, they also have low Q-type or low V-type BE dimension.

Proof of Proposition 13. We have

$$\begin{aligned} & \mathbb{E}[(f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(s_h, a_h, b_h) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \langle \mathbb{P}^\pi[(s_h, a_h, b_h) = \cdot], (f_h - \mathcal{T}_h^{\mu_g} f_{h+1})(\cdot) \rangle, \end{aligned} \quad (18)$$

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

Proof of Proposition 15. 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

$$\begin{aligned} & \mathbb{E}[(\phi_h^\top \theta_h - \mathcal{T}_h^{\mu_g}(\phi_{h+1}^\top \theta_{h+1}))(s_h, a_h, b_h) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \mathbb{E}[\phi_h(s_h, a_h, b_h)^\top (\theta_h - \theta_h^g) \mid (s_h, a_h, b_h) \sim \pi] \\ &= \langle \mathbb{E}_\pi[\phi_h(s_h, a_h, b_h)], \theta_h - \theta_h^g \rangle, \end{aligned} \quad (19)$$

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

Proof of Proposition 16. 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

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

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

Proof of Proposition 17. Let $d(i)$ denote the distribution of state s given $q(s) = i$. For any policy π and $f, g, w \in \mathcal{F}$

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

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

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

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

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

H. Proofs for Appendix D

In this section, we present the proof of Theorem 11 and Corollary 12. The techniques are basically the same as those in Appendix F. Please notice that whenever coming across $\mathcal{D}_\mathcal{F}$ and \mathcal{D}_Δ in this section, we use their definitions introduced in Appendix D.

To begin with, we have the following concentration lemma (akin to Lemma 23 and 24) for the sub-routine COMPUTE_EXPLOITER under the samples collected with Option II.

Lemma 30. *Under Assumption 8 and 9, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$ 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*

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

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 F.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 F.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 24 in Appendix F.3.4. \square

Proposition 31 (approximate best-response regret). *Under Assumption 8 and 9, there exists an absolute constant c such that for any $\delta \in (0, 1]$, $K \in \mathbb{N}$, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$, 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 30, $Q^{\mu^k, \dagger} \in \mathcal{C}^{\mu^k}$. Since COMPUTE_EXPLOITER chooses \tilde{f}^k optimistically, value difference lemma (Lemma 20) gives

$$\begin{aligned}
 & \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - V_1^{\mu^k, \dagger}(s_1) \right) \\
 & \leq \sum_{k=1}^K \left(V_1^{\mu^k, \nu^k}(s_1) - \min_{\nu} \mathbb{D}_{\mu_1^k \times \nu} \tilde{f}_1^k(s_1) \right)
 \end{aligned}$$

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

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 2.

By Jensen's inequality and Lemma 30 (b), we have for all $k \in [K]$

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

So we can apply Lemma 19 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 10), $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and obtain

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

Similarly by using Lemma 30 (a), we can show

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

Putting all relations together as in Proposition 25 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 21 and 22) under the samples collected with Option II.

Lemma 32. *Under Assumption 8 and 9, if we choose $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|/\delta)$ 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) $Q^* \in \mathcal{C}^k,$

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 F.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 F.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 24 in Appendix F.3.2. \square

Proof of Theorem 11. 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 31, (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 7, we have

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

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 2.

By Jensen's inequality and Lemma 32 (b), we have for all $k \in [K]$

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

So we can apply Lemma 19 with $\mathcal{G} = \mathcal{H}_{\mathcal{F}, h}$ (here, $\mathcal{H}_{\mathcal{F}}$ refers to the one introduced in Definition 10), $\Pi = \mathcal{D}_{\Delta, h}$, $\epsilon = 1/K$ and obtain

$$\sum_{k=1}^K \left| \mathbb{E} \left[(f_h^k - \mathcal{T}_h f_{h+1}^k)(s_h^k, a, b) \mid (a, b) \sim \mu_h^k \times \nu_{f^k, \bar{f}^k, h} \right] \right|$$

$$\leq \sqrt{|\mathcal{A}||\mathcal{B}|\beta \cdot \dim_{\text{DE}}(\mathcal{H}_{\mathcal{F},h}, \mathcal{D}_{\Delta,h}, 1/K) \cdot K}.$$

Similarly by using Lemma 32 (a), we can show

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

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 12. We proceed as in the proof of Theorem 11 but take $\omega = \epsilon/H$ (instead of $\omega = 1/K$) every time we incur Lemma 19, we can get

$$\frac{1}{K} \sum_{k=1}^K [V_1^*(s_1) - V_1^{\mu^{\text{out}}, \dagger}(s_1)] \leq \mathcal{O}\left(\frac{H}{K} \sqrt{|\mathcal{A}||\mathcal{B}| \dim_{\text{VBE}}(\mathcal{F}, \epsilon/H) K \beta} + \epsilon\right),$$

where $\beta = c \cdot \log(KH|\mathcal{F}||\mathcal{G}|)$.

To make the right hand side smaller than $\mathcal{O}(\epsilon)$, it suffices to take

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

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

We conclude the proof by recalling the definition of μ^{out} and Markov's inequality. \square