
Breaking the Curse of Many Agents: Provable Mean Embedding Q -Iteration for Mean-Field Reinforcement Learning

Lingxiao Wang¹ Zhuoran Yang² Zhaoran Wang¹

Abstract

Multi-agent reinforcement learning (MARL) achieves significant empirical successes. However, MARL suffers from the curse of many agents. In this paper, we exploit the symmetry of agents in MARL. In the most generic form, we study a mean-field MARL problem. Such a mean-field MARL is defined on mean-field states, which are distributions that are supported on continuous space. Based on the mean embedding of the distributions, we propose MF-FQI algorithm, which solves the mean-field MARL and establishes a non-asymptotic analysis for MF-FQI algorithm. We highlight that MF-FQI algorithm enjoys a “blessing of many agents” property in the sense that a larger number of observed agents improves the performance of MF-FQI algorithm.

1

1. Introduction

Reinforcement learning (RL) (Sutton & Barto, 2018) searches for the optimal policy for sequential decision making through interacting with environments and learning from experiences. Multi-agent reinforcement learning (MARL) (Bu et al., 2008) generalizes RL to multi-agent systems. For competitive tasks such as zero-sum game and general-sum game, various MARL algorithms (Hu & Wellman, 2003; Littman, 1994; Wang & Sandholm, 2003) are proposed in search for the Nash equilibrium (Nash, 1951). Meanwhile, for cooperative tasks, MARL searches for the optimal policy that maximizes the social welfare (Ng, 1975), i.e., the expected total reward obtained by all agents (Claus & Boutilier,

¹Department of Industrial Engineering and Management Sciences, Northwestern University, USA ²Department of Operations Research and Financial Engineering, Princeton University, USA. Correspondence to: Lingxiao Wang <lingxiaoawang2022@u.northwestern.edu>.

1998; Džeroski et al., 2001; Guestrin et al., 2002; Kar et al., 2013; Lauer & Riedmiller, 2000; Panait & Luke, 2005; Tan, 1993; Wang & Sandholm, 2003; Zambaldi et al., 2018). Combined with the breakthrough in deep learning, MARL achieves significant empirical successes in both settings, e.g., autonomous driving (Shalev-Shwartz et al., 2016), Go (Silver et al., 2016; 2017), esports (OpenAI, 2018; Vinyals et al., 2019), and robotics (Yang & Gu, 2004).

Despite its empirical successes, MARL remains challenging in the “many-agent” setting, as the capacity of state-action space grows exponentially in the number of agents, which hinders the learning of value function and policy due to the curse of dimensionality. Such a challenge is named as the “curse of many agents”. One way to break such a curse is through mean-field approximation, which exploits the symmetry of homogeneous agents and summarizes them as a population. In the most general form, such a population is represented by a distribution over the state space of individual agents, while the reward and transition are parametrized as functionals of distributions (Acciaio et al., 2018). Although mean-field MARL demonstrates remarkable efficiency in applications such as large-scale fleet management (Lin et al., 2018) and ridesharing order dispatching (Li et al., 2019), its theoretical analysis remains scarce. In particular, despite significant progress (Guo et al., 2019; Jiachen et al., 2017; Jiang & Lu, 2018; Yang et al., 2018), we still lack a principled model-free algorithm that allows individual agents to have continuous states, which requires approximating nonlinear functionals of infinite-dimensional mean-field states, e.g., value function and policy.

In this paper, we study mean-field MARL in the collaborative setting, where the mean-field states are distributions over a continuous space \mathcal{S} . Here \mathcal{S} denotes the state space of individual agents. In particular, we consider the setting with a centralized controller, which has a finite action space \mathcal{A} . Such a setting is extensively studied in the analysis of societal-scale systems (Gomes et al., 2015; Guéant et al., 2011; Moll et al., 2019). As a simplified example, the central bank or the central government decides whether to raise the interest rate or reduce the fiscal budget, respectively, both with the goal of maximizing social welfare. In such an example, the action space only contains two actions. How-

ever, the action taken by the centralized controller affects the dynamics of billions of individuals. Such a setting can be viewed as centralized mean-field control (Carmona et al., 2013; Fornasier & Solombrino, 2014; Huang et al., 2012) with an infinite number of homogeneous agents, which faces two challenges: (i) learning the value function and policy is intractable as they are functionals of distributions, which are infinite dimensional as \mathcal{S} is continuous, and (ii) the mean-field state is only accessible through the observation of a finite number of agents, which only provides partial information. To tackle these challenges, we resort to the mean embedding of mean-field states into a reproducing kernel Hilbert space (RKHS) (Gretton et al., 2007; Smola et al., 2007; Sriperumbudur et al., 2010), which allows us to parametrize value functions as nonlinear functionals over the RKHS. Based on value function approximation, we propose the mean-field fitted Q-iteration algorithm (MF-FQI), which provably attains the optimal value function at a linear rate of convergence. In particular, we show that MF-FQI breaks the curse of many agents in the sense that its computational complexity only scales linearly in the number of observed agents, and moreover, the statistical accuracy enjoys a “blessing of many agents”, that is, a larger number of observed agents improves the statistical accuracy. Moreover, we characterize the phase transition in the statistical accuracy in terms of the batch size in fitted Q-iteration and the number of observed agents.

Our Contribution. Our contribution is three-fold: (i) We propose the first model-free mean-field MARL algorithm, namely MF-FQI, that allows for continuous support with provable guarantees. (ii) We prove that MF-FQI breaks the curse of many agents by establishing its nonasymptotic computational and statistical rates of convergence. (iii) We motivate a principled framework for exploiting the invariance in MARL, e.g., exchangeability, via mean embedding.

Related Works. Our work is related to mean-field games and mean-field control. The study of mean-field games focuses on the search of the Nash equilibrium (Carmona & Delarue, 2018; Guéant et al., 2011; Huang et al., 2003; 2007; Lasry & Lions, 2006a;b; 2007), whereas the goal of mean-field control is to optimally control a McKean-Vlasov process (Acciaio et al., 2018; Andersson & Djehiche, 2011; Bensoussan et al., 2013; Buckdahn et al., 2009; 2011; Carmona et al., 2015; Meyer-Brandis et al., 2012). Most of these works focus on the continuous-time setting and require the knowledge of the transition model. In contrast, we consider the model-free and discrete-time setting.

Our work falls in the study of mean-field MARL, which generalizes finite-agent MARL by incorporating the notion of mean fields. Previous works investigate MARL in both cooperative and competitive settings (Guo et al., 2019; Jiachen et al., 2017; Jiang & Lu, 2018; Lin et al., 2019; Suttle et al.,

2019; Yang et al., 2018; Zhang et al., 2018a;b; 2019a). See (Zhang et al., 2019b) for the survey. In (Jiachen et al., 2017), a similar setting of mean-field MARL is studied, where the mean-field states are supported on a discrete space and the transition is linear in the state and action. In contrast, our work is model-free and allows for continuous support. In (Guo et al., 2019; Yang et al., 2018), mean-field MARL algorithms are proposed in the competitive setting, which have provable guarantees when the support is discrete. In comparison, we consider the cooperative setting with continuous support and establish nonasymptotic guarantees.

Our work exploits the exchangeability of agents via mean embedding. See e.g., (Fukumizu et al., 2008; Gretton et al., 2009; 2012; Smola et al., 2007; Sriperumbudur et al., 2010; Tolstikhin et al., 2017) and references therein for the study of mean embedding. Our work is closely related to various statistical models that exploit invariance, such as set kernels (Gärtner et al., 2002; Haussler, 1999; Kondor & Jebara, 2003) and deep sets (Zaheer et al., 2017). We refer to (Bloem-Reddy & Teh, 2019) for a detailed survey on learning with invariance.

Notations. For a topological space X , we denote by $\mathcal{C}_B(X)$ the set of bounded and continuous real functions on X . We denote by $\mathcal{M}(X)$ the space of all the probability distributions supported on X . For $x \in X$, we denote by $\delta_x \in \mathcal{M}(X)$ the point mass at x . For a real-valued function f defined on X , we denote by $\|f\|_{p,\nu}$ the $L_p(\nu)$ norm for $p \geq 1$, where $\nu \in \mathcal{M}(X)$. We write $\|f\|_\nu = \|f\|_{2,\nu}$ for notational simplicity.

2. Mean-Field MARL via Mean Embedding

In this section, we first motivate mean-field MARL by an example of N -player control with invariance. We then introduce the problem setup of mean-field MARL and the mean embedding of distributions. Finally, we propose the MF-FQI algorithm, which solves mean-field MARL based on the mean embedding.

2.1. Exchangeability in MARL

We first consider an N -player control problem with a centralized controller in discrete time. Such a setting is extensively studied in the analysis of societal-scale systems (Gomes et al., 2015; Guéant et al., 2011; Moll et al., 2019), such as the example of central bank or central government in §1. At each time step t , the central controller takes an action $a_t \in \mathcal{A}$ based on the current joint state $s_t = (s_{1,t}, \dots, s_{N,t})$, where $s_{i,t} \in \mathcal{S}$ is the state of the i -th agent at time t . The immediate reward r_t follows a distribution that depends on the current state $s_t \in \mathcal{S}^N$ and action $a_t \in \mathcal{A}$. The transition of the joint state follows a distribution, which is determined by the current state $s_t \in \mathcal{S}^N$ and action $a_t \in \mathcal{A}$. In summary,

it holds that

$$r_t \sim r(s_t, a_t), \quad S_{t+1} \sim P(\cdot | s_t, a_t). \quad (2.1)$$

The process defined by the tuple $(\mathcal{S}, \mathcal{A}, P, r)$ is a Markov decision process (MDP). We define a policy $\pi : \mathcal{S}^N \mapsto \mathcal{M}(\mathcal{A})$ as a mapping that maps a joint state $s \in \mathcal{S}^N$ to a probability distribution $\pi(\cdot | s)$ over \mathcal{A} . We define the value function corresponding to the policy π as

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(S_t, A_t) \mid S_0 = s \right], \quad (2.2)$$

where $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \mathbb{P}(\cdot | S_t, A_t)$, and $\gamma \in (0, 1)$ is the discount factor. Similarly, we define the action-value function corresponding to the policy π as follows,

$$Q^\pi(s, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(S_t, A_t) \mid S_0 = s, A_0 = a \right], \quad (2.3)$$

where $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \mathbb{P}(\cdot | S_t, A_t)$. We define the Bellman operator T^π as follows,

$$T^\pi Q(s, a) = \mathbb{E} [r(s, a) + \gamma \cdot Q(S', A')], \quad (2.4)$$

where $S' \sim P(\cdot | s, a)$ and $A' \sim \pi(\cdot | s)$. Our goal is to find the optimal policy that maximizes the expected total reward as follows,

$$Q^*(s, a) = \sup_{\pi} Q^\pi(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (2.5)$$

We denote by π^* the optimal solution of (2.5). It can be shown that $Q^* = Q^{\pi^*}$, and the following Bellman optimality equation holds,

$$Q^*(s, a) = TQ^*(s, a) = \mathbb{E} \left[r(s, a) + \gamma \cdot \max_{a \in \mathcal{A}} Q^*(S', a) \right], \quad (2.6)$$

where $S' \sim P(\cdot | s, a)$. Here we call T the Bellman optimality operator.

Curse of Many Agents: The learning of the optimal action-value function Q^* under the N -player setting suffers from the curse of many agents. More specifically, as N increases, the capacity of the joint state space \mathcal{S}^N grows exponentially in N and incurs intractability in the learning of the action-value function Q . To address such a curse, we exploit the exchangeability of the MDP in (2.1). More specifically, we assume that the MDP is exchangeable in the sense that

$$r(s_t, a_t) \stackrel{d}{=} r(\sigma(s_t), a_t), \quad P(s_{t+1} | s_t, a_t) = P(\sigma(s_{t+1}) | \sigma(s_t), a_t) \quad (2.7)$$

which holds for any $s_t, s_{t+1} \in \mathcal{S}^N$, $a_t \in \mathcal{A}$, and $\sigma \in \mathbb{S}_N$. Here σ is a block-wise permutation of the vector $s_t \in \mathcal{S}^N$, and \mathbb{S}_N is the permutation group of order N . Under the exchangeability defined in (2.7), the following proposition shows that the optimal policy is invariant to permutations of the joint state.

Proposition 2.1 (Invariance of Q^*). If (2.7) holds, then it holds for any $\sigma \in \mathbb{S}_N$, $s \in \mathcal{S}^N$, and $a \in \mathcal{A}$ that

$$Q^*(s, a) = Q^*(\sigma(s), a). \quad (2.8)$$

Moreover, it holds for any $\sigma \in \mathbb{S}_n$, $s \in \mathcal{S}^N$, and $a \in \mathcal{A}$ that $\pi^*(a | s) \stackrel{d}{=} \pi^*(a | \sigma(s))$.

Proof. See §B.1 for a detailed proof. \square

Meanwhile, the following proposition proves that the action-value function Q^π is invariant to permutations of the joint states.

Proposition 2.2 (Invariant Representation). Let π be invariant to permutations such that $\pi(\cdot | s) \stackrel{d}{=} \pi(\cdot | \sigma(s))$. If (2.7) holds, then it holds for some $g : \mathcal{M}(\mathcal{S}) \times \mathcal{A} \mapsto \mathbb{R}$ that

$$Q^\pi(s, a) = g(M_s, a). \quad (2.9)$$

Here $M_s(\cdot)$ is the empirical measure supported on the set $\{s_i\}_{i \in [N]}$ corresponding to $s = (s_1, \dots, s_N)$, which takes the form of

$$M_s = \frac{1}{N} \sum_{i=1}^N \delta_{s_i}, \quad (2.10)$$

where recall that δ_{s_i} is the point mass at $s_i \in \mathcal{S}$ for all $i \in [N]$.

Proof. See §B.2 for a detailed proof. \square

By Propositions 2.1 and 2.2, the optimal action-value function Q^* is related to the joint state through the empirical state distribution M_s defined in (2.10). When the number of agents N goes to infinity, the empirical state distribution M_s converges to a limiting continuous distribution. To capture such a limiting dynamics of infinite agents with exchangeability, we define an MDP with $\mathcal{M}(\mathcal{S})$, the space of probability measures supported on \mathcal{S} , as the mean-field state space as follows.

Mean-Field MARL. We define the discounted mean-field MDP by the tuple $(\mathcal{M}(\mathcal{S}), \mathcal{A}, P, r, \gamma)$. Here \mathcal{A} is the action space and $\mathcal{M}(\mathcal{S})$ is the mean-field state space, which is the space of all the distributions supported on \mathcal{S} . Given a mean-field state $p_s \in \mathcal{M}(\mathcal{S})$ and an action $a \in \mathcal{A}$, the immediate reward follows the distribution $r(a, p_s)$, where $r : \mathcal{A} \times \mathcal{M}(\mathcal{S}) \mapsto \mathcal{M}(\mathbb{R})$. The Markov kernel $P(\cdot | \cdot)$ maps the

action-state pair (a, p_s) to a distribution on $\mathcal{M}(\mathcal{S})$, which is the distribution of the mean-field state after transition from the action-state pair (a, p_s) . For a policy $\pi(\cdot | p_s)$ that maps from $\mathcal{M}(\mathcal{S})$ to $\mathcal{M}(\mathcal{A})$, we define the action-value function as

$$Q^\pi(a, p_s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(P_{s,t}, A_t) \mid P_{s,0} = p_s, A_0 = a \right], \quad (2.11)$$

where $A_t \sim \pi(\cdot | P_{s,t})$ and $P_{s,t+1} \sim P(\cdot | a, p_s, t)$. Correspondingly, we define the Bellman evaluation operator T^π as follows,

$$T^\pi Q(a, p_s) = \mathbb{E} [r(a, p_s) + \gamma \cdot Q(A', P_{s'})], \quad (2.12)$$

where $P_{s'} \sim P(\cdot | a, p_s)$ and $A' \sim \pi(\cdot | a, p_s)$. We define the optimal action-value function as $Q^* = \sup_\pi Q^\pi$. The Bellman optimality equation then takes the form of

$$Q^*(a, p_s) = TQ^*(a, p_s) = \mathbb{E} \left[r(a, p_s) + \gamma \cdot \max_{a \in \mathcal{A}} Q^*(a, P_{s'}) \right], \quad (2.13)$$

where $P_{s'} \sim P(\cdot | a, p_s)$. We write $\Omega = \mathcal{S} \times \mathcal{A}$ and define the space of state-action configurations $\widetilde{\mathcal{M}}(\Omega)$ as follows,

$$\widetilde{\mathcal{M}}(\Omega) = \{ \omega_{a,p_s} = \delta_a \times p_s : a \in \mathcal{A}, p_s \in \mathcal{M}(\mathcal{S}) \}. \quad (2.14)$$

Here we denote by $\delta_a \in \mathcal{M}(\mathcal{A})$ the point mass at action $a \in \mathcal{A}$ and by $\delta_a \times p_s$ the product measure on $\Omega = \mathcal{S} \times \mathcal{A}$ induced by δ_a and p_s . See §A for the definition of a topological structure that allows us to define distributions on $\widetilde{\mathcal{M}}(\Omega)$. Note that the transition kernel $P(\cdot | a, p_s)$ equivalently defines a Markov kernel from $\widetilde{\mathcal{M}}(\Omega)$ to $\mathcal{M}(\mathcal{S})$. With a slight abuse of notations, we denote by $P(\cdot | \omega_{a,p_s})$ such a Markov kernel and do not distinguish between them. Similarly, we denote by $r(\omega)$ and $Q(\omega)$ the immediate reward and action-value function defined on the state-action configuration $\omega \in \widetilde{\mathcal{M}}(\Omega)$, respectively. We assume that the action set \mathcal{A} is finite, and the immediate reward is upper bounded by a positive absolute constant R_{\max} . It then holds that the action-value functions are upper bounded by $Q_{\max} = R_{\max}/(1 - \gamma)$.

In a multi-agent environment with infinite homogeneous agents and continuous state space \mathcal{S} , we cannot access the mean-field state p_s directly. Instead, we assume that we observe the states of N agents that follows the mean-field state p_s . In what follows, we construct an algorithm that solves the mean-field MARL via such a finite observation for each mean-field state. In the sequel, we denote by $\widehat{Q}_\kappa^\lambda$ and π_κ the outputs of MF-FQI.

We highlight that the mean-field MARL setting faces two challenges: (i) learning the value function and policy is

intractable as they are functionals of distributions, which are infinite dimensional as \mathcal{S} is continuous, and (ii) the mean-field state is only accessible through the observation of a finite number of agents, which only provides partial information. In what follows, we tackle these challenges via mean embedding.

2.2. Mean Embedding

To learn the optimal action-value function Q^* defined on $\widetilde{\mathcal{M}}(\Omega)$, which is a space of distributions, we introduce mean embedding, which embeds the space of distributions to a reproducing kernel Hilbert space (RKHS). We denote by $\mathcal{H}(k)$ the RKHS with reproducing kernel $k : \Omega \times \Omega \mapsto \mathbb{R}$. For any state-action configuration $\omega \in \widetilde{\mathcal{M}}(\Omega)$, the mean embedding $\mu_\omega(\cdot)$ of ω into the RKHS $\mathcal{H}(k)$ is defined as follows (Gretton et al., 2007; Smola et al., 2007; Sriperumbudur et al., 2010),

$$\mu_\omega(x) = \int_{\Omega} k(x, t) d\omega(t) \in \mathcal{H}(k). \quad (2.15)$$

Let $\mathcal{X} = \{ \mu_\omega : \omega \in \widetilde{\mathcal{M}}(\Omega) \} \subseteq \mathcal{H}(k)$. To tackle challenge (i), we introduce another reproducing kernel $K : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$. Such a kernel generates an RKHS $\mathcal{H}(K)$ that include functions defined on \mathcal{X} . Our idea is then to approximate Q^* using functions in $\mathcal{H}(K)$. Note that upon a proper selection of kernel $K(\cdot, \cdot)$, the corresponding RKHS $\mathcal{H}(K)$ captures a rich family of functions defined on \mathcal{X} . As an example, for universal kernels such as the radial basis function (RBF) kernel, the corresponding RKHS is dense in $\mathcal{C}(\mathcal{X})$. To further regulate the behavior of mean embedding and RKHS, we introduce the following regularity conditions on kernels $k(\cdot, \cdot)$ and $K(\cdot, \cdot)$.

Assumption 2.3 (Regularity Condition of Kernels). We assume that the kernel $k(\cdot, \cdot)$ and $K(\cdot, \cdot)$ are continuous and bounded as follows,

$$k(u, u) \leq \varrho, \quad \forall u \in \Omega, \quad K(\mu_\omega, \mu_\omega) \leq \varsigma, \quad \forall \omega \in \widetilde{\mathcal{M}}(\Omega), \quad (2.16)$$

where ϱ and ς are positive absolute constants. We assume that $k(\cdot, \cdot)$ is universal and the mean embeddings $\mu_\omega(\cdot)$ are continuous for any $\omega \in \widetilde{\mathcal{M}}(\Omega)$. We further assume that $K(\cdot, \mu_\omega)$ is Hölder continuous such that for any $x, y \in \widetilde{\mathcal{M}}(\Omega)$, it holds that

$$\|K(\cdot, \mu_x) - K(\cdot, \mu_y)\|_{\mathcal{H}(K)} \leq L \cdot \|\mu_x - \mu_y\|_{\mathcal{H}(k)}^h, \quad (2.17)$$

where L and h are positive absolute constants.

The assumption on the boundedness of the kernels in (2.16) is a standard assumption in the learning with kernel embedding (Caponnetto & De Vito, 2007; Lin et al., 2017; Muandet et al., 2012; Szabó et al., 2015). The universality assumption on $k(\cdot, \cdot)$ ensures that each mean embedding uniquely

characterizes a distribution (Gretton et al., 2007; 2012). The continuous assumption on the embeddings μ_ω is a mild regularity condition, which holds if the kernel $k(\cdot, \cdot)$ is universal and the domain Ω is compact (Sriperumbudur et al., 2010). Meanwhile, the Hölder continuity of $K(\cdot, \cdot)$ in Assumption 2.3 is a mild regularity condition. Such an assumption holds for a rich family of commonly used reproducing kernels, such as the linear kernel $K(\mu_x, \mu_y) = \langle \mu_x, \mu_y \rangle_{\mathcal{H}(k)}$ and the RBF kernel $K(\mu_x, \mu_y) = \exp(-\|\mu_x - \mu_y\|_{\mathcal{H}(k)}^2 / \sigma^2)$.

We highlight that the Hölder continuity of $K(\cdot, \cdot)$ allows for an approximation of mean embedding μ_p based on the empirical approximation \widehat{p} of the distribution $p \in \widetilde{\mathcal{M}}(\Omega)$ with finite observations, which further allows for an approximation of the action-value function with finite observations, thus tackles the challenge (ii). For an empirical approximation of state-action configuration $\omega_{a, \widehat{p}_s} = \delta_a \times \widehat{p}_s$, where \widehat{p}_s is the empirical distribution of the observed states $\{s_i\}_{i \in N}$, the mean embedding takes the following form,

$$\mu_{\omega_{a, \widehat{p}_s}}(\cdot) = \frac{1}{N} \sum_{i=1}^N k(\cdot, (a, s_i)), \quad (2.18)$$

which is invariant to the permutation of states $\{s_i\}_{i \in [N]}$. Such an invariance is also exploited by the neural-network based approach named deep sets (Zaheer et al., 2017). In what follows, we connect our mean embedding approach to the invariant deep reinforcement learning under the framework of overparametrized two-layer neural networks (Arora et al., 2019; Jacot et al., 2018; Neyshabur et al., 2018; Zhang et al., 2016).

Connection to Invariant Deep Reinforcement Learning.

In what follows, we assume that $(a, s_i) \in \mathbb{R}^d$ and write $x_i = (a, s_i)$. We define the feature mappings $\{\phi_j(\cdot)\}_{j \in [m]}$ and $\{\Phi_\ell(\cdot)\}_{\ell \in [M]}$ as follows,

$$\begin{aligned} \phi_j(x) &= \frac{1}{\sqrt{m}} \cdot b_j \cdot \mathbb{1}\{w_j^\top x > 0\} \cdot x, \\ \Phi_r(q) &= \frac{1}{\sqrt{\ell}} \cdot b'_r \cdot \mathbb{1}\{W_r'^\top q > 0\} \cdot q, \end{aligned} \quad (2.19)$$

where $b_j, b'_r \sim \text{Unif}\{-1, 1\}$, $w_j \sim N(0, I_{d\ell}/(d\ell))$, and $W_r' \sim N(0, I_\ell/\ell)$. Correspondingly, we define the kernels $k_m(x, y) = \sum_{j=1}^m \phi_j(x)^\top \phi_j(y)$ and $K_\ell(p, q) = \sum_{r=1}^\ell \Phi_r(p)^\top \Phi_r(q)$.

Note that the mean embedding of a point mass δ_{x_i} by the kernel k_m is an array $\mu_i = [\phi_1(x_i)^\top, \dots, \phi_m(x_i)^\top]^\top \in \mathbb{R}^{md}$. For a mean embedding $\mu \in \mathbb{R}^{md}$, we consider the parametrization of action-value functions $Q(\mu) = \widetilde{Q}(D^\top \mu)$, where $\widetilde{Q} \in \mathcal{H}(K_\ell)$ and $D = [D_1, \dots, D_m] \in \mathbb{R}^{m \times \ell}$ with $D_j \in \mathbb{R}^{d \times \ell}$ ($j \in [m]$). Let μ_p be the mean embedding of the empirical measure supported on $\{x_i\}_{i \in [N]}$.

It then holds for some $\{\alpha_r\}_{r \in [\ell]} \subseteq \mathbb{R}^d$ that

$$\begin{aligned} Q(\mu_p) &= \sum_{r=1}^\ell \alpha_r^\top \Phi_r(D^\top \mu_p) \\ &= \frac{1}{\sqrt{\ell}} \sum_{r=1}^\ell b'_r \cdot \mathbb{1}\{W_r'^\top \rho > 0\} \cdot \alpha_r^\top (D^\top \mu_p) \\ &= f(D^\top \mu_p). \end{aligned} \quad (2.20)$$

Note that if α_r is sufficiently close to W_r' , then $Q(\mu_p)$ is close to $\widetilde{f}(D^\top \mu_p)$, where

$$\widetilde{f}(\rho) = \frac{1}{\sqrt{\ell}} \sum_{r=1}^\ell b'_r \cdot \mathbb{1}\{\alpha_r^\top \rho > 0\} \cdot \alpha_r^\top(\rho), \quad \forall \rho \in \mathbb{R}^\ell. \quad (2.21)$$

Here $\widetilde{f}(\cdot)$ is a two-layer neural network with parameters b'_r, α_r ($r \in [\ell]$) and ReLU activation function. Meanwhile, it holds that

$$\begin{aligned} D^\top \mu_p &= \frac{1}{N} \sum_{i=1}^N \frac{1}{\sqrt{m}} \sum_{j=1}^m b_j \cdot \mathbb{1}\{w_j^\top x_i > 0\} \cdot D_j^\top x_i \\ &= \frac{1}{N} \sum_{i=1}^N \psi(x_i). \end{aligned} \quad (2.22)$$

Similarly, if D_j is close to w_j , then $\psi(x_i)$ is close to $\widetilde{\psi}(x_i)$ defined as follows,

$$\widetilde{\psi}(x_i) = \frac{1}{\sqrt{m}} \sum_{j=1}^m b_j \cdot \mathbb{1}\{D_j^\top x_i > 0\} \cdot D_j^\top x_i, \quad \forall x_i \in \mathbb{R}^d. \quad (2.23)$$

Here $\widetilde{\psi}(x_i)$ is a two-layer neural network with input x_i , parameters b'_j ($j \in [m]$) and D , and ReLU activation function. Note that for the functions f and ψ defined in (2.20) and (2.22) to approximately take the form of neural networks, we requires the parameters α and D to be sufficiently close to W and w , respectively. Such a requirement is formally characterized by the study of overparametrized neural networks. More specifically, if the widths m and ℓ of neural networks are sufficiently large (which depends on the deviation of parameters α and D from their respective initializations W and w), then the functions f and ψ well approximates the neural networks \widetilde{f} and $\widetilde{\psi}$ defined in (2.21) and (2.23), respectively.

In conclusion, under the mean embedding with the feature mappings defined in (2.19), the parameterization of action-value function takes the form of $Q(\mu_p) = f(1/N \cdot \sum_{i=1}^N \psi(x_i))$, where f and ψ are approximations of two-layer neural networks \widetilde{f} and $\widetilde{\psi}$, respectively. Hence, the action-value function $Q(\mu_p)$ approximately takes the form of deep sets (Zaheer et al., 2017) with $\{x_i\}_{i \in [N]}$ as the set input.

2.3. Mean-Field Fitted Q -Iteration

In what follows, we establish a value-based algorithm that solves mean-field MARL problem in §2.1 based on fitted Q -iteration (Ernst et al., 2005). More specifically, we propose an algorithm that learns the optimal action-value function Q^* by the sample $\{\delta_{a_i} \times p_{i,s}\}_{i \in [n]}$ that follows a sampling distribution ν over the space of state-action configurations $\widetilde{\mathcal{M}}(\Omega)$. For each state-action configuration $\delta_{a_i} \times p_{i,s}$, the mean-field state $p_{i,s} \in \mathcal{M}(\mathcal{S})$ is available to us through the observed states $\{s_{i,j}\}_{j \in [N]}$, which are sampled independently from the mean-field state $p_{i,s}$. We further observe the immediate reward r_i and the states $\{s'_{i,j}\}_{j \in [N]}$ that are independently sampled from the mean-field state $p_{i,s'}$. Here $p_{i,s'} \sim P(\cdot | a_i, p_{i,s})$ is the mean-field state after transition from the state-action configuration $(a_i, p_{i,s})$. Given the batch of data $\{(\{s_{i,j}\}_{j \in [N]}, a_i, r_i, \{s'_{i,j}\}_{j \in [N]})\}_{i \in [n]}$, the mean-field fitted Q -iteration (MF-FQI) sequentially computes

$$\widehat{y}_{i,k} = r_i + \max_{a \in \mathcal{A}} \widehat{Q}_k^\lambda(\mu_{\omega_a, \widehat{p}_{i,s'}}) \quad (2.24)$$

at the k -th iteration. Here $\mu_{\omega_a, \widehat{p}_{i,s'}}$ is the mean embedding of the distribution $\omega_a, \widehat{p}_{i,s'} = \delta_a \times \widehat{p}_{i,s'}$, $\widehat{p}_{i,s'}$ is the empirical distribution supported on the set $\{s'_{i,j}\}_{j \in [N]}$, and Q_k^λ is the approximation of the optimal action-value function at the k -th iteration of MF-FQI. Upon computing $\{\widehat{y}_{i,k}\}_{i \in [n]}$ according to (2.24), MF-FQI then updates the approximation of the optimal action-value function in the RKHS $\mathcal{H}(K)$ by solving the following optimization problem,

$$\begin{aligned} Q_{k+1}^\lambda &= \operatorname{argmin}_{f \in \mathcal{H}(K)} \frac{1}{n} \sum_{i=1}^n (f(\mu_{\widehat{\omega}_i}) - \widehat{y}_{i,k})^2 + \lambda \cdot \|f\|_{\mathcal{H}(K)}^2, \\ \widehat{Q}_{k+1}^\lambda &= \min\{Q_{k+1}^\lambda, Q_{\max}\}, \end{aligned} \quad (2.25)$$

where $\widehat{\omega}_i = \delta_a \times \widehat{p}_{i,s}$ and $\widehat{p}_{i,s}$ is the empirical distribution supported on the set $\{s_{i,j}\}_{j \in [N]}$. We summarize MF-FQI defined by (2.24) and (2.25) in Algorithm 1. We highlight that MF-FQI has a linear computational complexity in terms of the number of observed agents N . Therefore, MF-FQI is computationally tractable even for a large number of observed agents N .

3. Main Results

In this section, we establish the theoretical guarantee of MF-FQI defined in Algorithm 1. In the sequel, we denote by $\widehat{Q}_\kappa^\lambda$ and π_κ the outputs of MF-FQI. Our goal is to establish an upper bound for $\|Q^* - Q^{\pi_\kappa}\|_{1,\mu}$, where μ is the measurement distribution over $\widetilde{\mathcal{M}}(\Omega)$.

We first introduce the definition of concentration coefficients. For a policy π_1 , we define $E^{\pi_1}\nu$ as the distribution of $\Lambda_1 = \delta_{A_1} \times P_{1,s}$, where $P_{1,s} \sim P(\cdot | \nu)$ and

Algorithm 1 Mean-Field Fitted Q -Iteration (MF-FQI)

- 1: **Input:** Batch of data $\{(\{s_{i,j}\}_{j \in [N]}, a_i, r_i, \{s'_{i,j}\}_{j \in [N]})\}_{i \in [n]}$, reproducing kernels $k(\cdot, \cdot)$ and $K(\cdot, \cdot)$, number of iterations κ , parameter λ , initial action-value function \widehat{Q}_0^λ .
- 2: For all $i \in [n]$, compute mean embeddings $\mu_{\widehat{\omega}_i}$ and $\mu_{a, \widehat{p}_{i,s'}}$ for all $a \in \mathcal{A}$ as follows,

$$\mu_{\widehat{\omega}_i}(\cdot) = \frac{1}{N} \sum_{j=1}^N k(\cdot, (a_i, s_{i,j})),$$

$$\mu_{a, \widehat{p}_{i,s'}}(\cdot) = \frac{1}{N} \sum_{j=1}^N k(\cdot, (a, s'_{i,j})), \quad \forall a \in \mathcal{A}.$$

- 3: **for** $k = 0, 1, \dots, \kappa - 1$ **do**
- 4: Compute $\widehat{y}_{i,k} = r_i + \max_{a \in \mathcal{A}} \widehat{Q}_k^\lambda(\mu_{a, \widehat{p}_{i,s'}})$ for all $i \in [n]$.
- 5: Update the action-value function $\widehat{Q}_{k+1}^\lambda$ by

$$Q_{k+1}^\lambda = \operatorname{argmin}_{f \in \mathcal{H}(K)} \frac{1}{n} \sum_{i=1}^n (f(\mu_{\widehat{\omega}_i}) - \widehat{y}_{i,k})^2 + \lambda \cdot \|f\|_{\mathcal{H}(K)}^2,$$

$$\widehat{Q}_{k+1}^\lambda = \min\{Q_{k+1}^\lambda, Q_{\max}\}.$$

- 6: **end for**
- 7: **Output:** $\{\widehat{Q}_i^\lambda\}_{i \in [\kappa]}$.

$A_1 \sim \pi_1(\cdot | P_{1,s})$. Similarly, for policies $\{\pi_i\}_{i \in [\ell]}$, we recursively define $E^{\pi_\ell} \circ E^{\pi_{\ell-1}} \circ \dots \circ E^{\pi_1}\nu$ as the distribution of $\Lambda_\ell = \delta_{A_\ell} \times P_{\ell,s}$, where $P_{\ell,s} \sim P(\cdot | E^{\pi_{\ell-1}} \circ \dots \circ E^{\pi_1}\nu)$ and $A_\ell \sim \pi_\ell(\cdot | P_{\ell,s})$. In what follows, we define the concentration coefficients that measures the difference between the sampling distribution ν and the measurement distribution μ on $\widetilde{\mathcal{M}}(\Omega)$.

Assumption 3.1. (Concentration Coefficients) Let ν be the sampling distribution on $\widetilde{\mathcal{M}}(\Omega)$. Let μ be the measurement distribution on $\widetilde{\mathcal{M}}(\Omega)$. We assume that for any policies $\{\pi_i\}_{i \in [\ell]}$, the distribution $E^{\pi_\ell} \circ E^{\pi_{\ell-1}} \circ \dots \circ E^{\pi_1}\nu$ is absolutely continuous with respect to μ . We define the ℓ -th concentration coefficients between ν and μ as follows,

$$\phi(\ell; \mu, \nu) \quad (3.1)$$

$$= \sup_{\pi_1, \dots, \pi_\ell} \left(\mathbb{E}_\mu \left[\left(\frac{dE^{\pi_\ell} \circ E^{\pi_{\ell-1}} \circ \dots \circ E^{\pi_1}\nu}{d\mu} \right)^2 \right] \right)^{1/2}.$$

We assume that $\phi(\ell; \mu, \nu) < +\infty$ for any $\ell \in \mathbb{N}$. We further assume that there exist a positive absolute constant $\Phi(\mu, \nu)$ such that

$$\sum_{\ell=1}^{\infty} \gamma^{\ell-1} \cdot \ell \cdot \phi(\ell; \mu, \nu) \leq \Phi(\mu, \nu) / (1 - \gamma)^2. \quad (3.2)$$

Assumption 3.2 is a standard assumption in the theoretical analysis of reinforcement learning (Antos et al., 2008; Chen & Jiang, 2019; Farahmand et al., 2010; 2016; Lazaric et al., 2016; Munos & Szepesvári, 2008; Scherrer, 2013; Scherrer et al., 2015; Szepesvári & Munos, 2005; Yang et al., 2019). Under Assumption 3.1, the following theorem upper bounds the error $\|Q^* - Q^{\pi_\kappa}\|_{1,\mu}$ of MF-FQI.

Theorem 3.2 (Error Propagation). Let $\{\widehat{Q}_i^\lambda\}_{i \in [\kappa]}$ be the output of Algorithm 1. Let π_κ be the greedy policy corresponding to $\widehat{Q}_\kappa^\lambda$. Under Assumption 3.1, it holds that

$$\begin{aligned} \|Q^* - Q^{\pi_\kappa}\|_{1,\mu} &\leq \underbrace{\frac{2\gamma \cdot \Phi(\mu, \nu)}{(1-\gamma)^2} \cdot \max_{i \in [\kappa]} \|\widehat{Q}_i^\lambda - T\widehat{Q}_{i-1}^\lambda\|_\nu}_{(a)} \\ &\quad + \underbrace{\frac{4\gamma^{\kappa+1} \cdot Q_{\max}}{1-\gamma}}_{(b)}. \end{aligned} \quad (3.3)$$

Proof. See §B.3 for a detailed proof. \square

Following from Theorem 3.2, the error of MF-FQI is upper bounded by the sum of the two terms on the right-hand side of (3.3). Here term (b) characterizes the algorithmic error that hinges on the number of iterations κ . Meanwhile, term (a) characterizes the one-step approximation error that hinges on the approximation \widehat{Q}_i^λ of $T\widehat{Q}_{i-1}^\lambda$. In the sequel, we upper bound the one-step approximation error characterized by term (a). To this end, we first impose the following regularity condition on the Bellman optimality operator T and the RKHS $\mathcal{H}(K)$.

Assumption 3.3 (Regularity Condition of T and $\mathcal{H}(K)$). We define the integral operator \mathcal{C} as follows,

$$\mathcal{C}f(x) = \int_{\widetilde{\mathcal{M}}(\Omega)} K(x, \mu_\omega) f(\mu_\omega) d\nu(\omega). \quad (3.4)$$

We assume that the eigenvalues $\{t_n\}_{n \in \mathbb{N}}$ of \mathcal{C} is bounded such that $\alpha \leq n^b t_n \leq \beta$ for all $n \in \mathbb{N}$, where α, β , and $b > 1$ are positive absolute constants. We further assume that for any output $Q^\lambda \in \mathcal{H}(K)$ of the regression problem defined in (2.25), it holds for some $g \in \mathcal{H}(K)$ that

$$Q_{\mathcal{H},T} = \mathcal{C}^{(c-1)/2} g, \quad \|g\|_{\mathcal{H}(K)} \leq R. \quad (3.5)$$

Here $R > 0$ and $c \in [1, 2]$ are absolute constants, and $Q_{\mathcal{H},T}$ is defined as follows,

$$\begin{aligned} \widehat{Q}^\lambda &= \min\{Q^\lambda, Q_{\max}\}, \\ Q_{\mathcal{H},T} &= \Pi_{\mathcal{H}(K)}(T\widehat{Q}^\lambda) = \operatorname{argmin}_{f \in \mathcal{H}(K)} \|f - T\widehat{Q}^\lambda\|_\nu, \end{aligned} \quad (3.6)$$

where we denote by $\Pi_{\mathcal{H}(K)}$ the projection onto $\mathcal{H}(K)$ with respect to the norm $\|\cdot\|_\nu$.

Assumption 3.3 is a mild regularity assumption on the RKHS $\mathcal{H}(K)$ and the Bellman optimality operator T . Similar assumptions arises in the analysis of kernel ridge regression (Caponnetto & De Vito, 2007; Lin et al., 2017; Szabó et al., 2015). The parameters b and c in Assumption 3.3 define a prior space $\mathcal{P}(b, c)$ in the context of kernel ridge regression (Caponnetto & De Vito, 2007). Intuitively, the parameter c controls the smoothness of $Q_{\mathcal{H},T}$ defined in (3.6), and the parameter b controls the size of $\mathcal{H}(K)$. Under Assumption 3.3, the following theorem characterizes the one-step approximation error of MF-FQI defined in Algorithm 1.

Theorem 3.4 (One-step Approximation Error). Let η, τ be two constants such that $0 < \eta + \tau < 1$. Let $C(\eta) = 32 \log^2(6/\eta)$. Under Assumptions 2.3 and 3.3, for

$$\begin{aligned} N &\geq 2\varrho \cdot (1 + \sqrt{\log(|\mathcal{A}| \cdot n/2\tau)})^2 \cdot (64L^2\varsigma^2/\lambda^2)^{1/h}, \\ n &\geq \frac{2C(\eta)\varsigma\beta b}{(b-1)\lambda^{1+1/b}}, \quad \lambda \leq \|\mathcal{C}\|_{\mathcal{H}(K)}, \end{aligned} \quad (3.7)$$

it holds with probability at least $1 - \eta - \tau$ that

$$\|\widehat{Q}_k^\lambda - T\widehat{Q}_{k-1}^\lambda\|_\nu^2 \leq \mathcal{G}_1 + \mathcal{G}_2 + \psi_T^2, \quad \forall k \in [\kappa], \quad (3.8)$$

where

$$\begin{aligned} \mathcal{G}_1 &= \left(\frac{8L^2 Q_{\max}^2 (1 + \sqrt{\log(|\mathcal{A}| \cdot n/2\tau)})^{2h} \cdot (2\varrho)^h}{\lambda \cdot N^h} \right) \\ &\quad \cdot \left(1 + \frac{5\varsigma^2}{\lambda^2} \right), \\ \mathcal{G}_2 &= C(\eta) \cdot \left(R\lambda^c + \frac{\varsigma^2 R}{\lambda^{2-c} n^2} + \frac{\varsigma R \lambda^{c-1}}{4n} \right. \\ &\quad \left. + \frac{\varsigma M^2}{\lambda n^2} + \frac{\Sigma^2 \beta b}{(b-1)n\lambda^{1/b}} \right), \end{aligned} \quad (3.9)$$

and the term ψ_T is defined as follows,

$$\psi_T = \sup_{k \in [\kappa]} \|T\widehat{Q}_k^\lambda - \Pi_{\mathcal{H}(K)}(TQ_k^\lambda)\|_\nu. \quad (3.10)$$

Here M and Σ are positive absolute constants, the parameters ς, ϱ, h are defined in Assumption 2.3, and the parameters \mathcal{C}, b, c are defined in Assumption 3.3, and $\Pi_{\mathcal{H}(K)}$ the projection onto $\mathcal{H}(K)$ with respect to the norm $\|\cdot\|_\nu$.

Proof. See §B.4 for a detailed proof. \square

Following from Theorem 3.4, the one-step approximation error is upper bounded by the sum of the three terms, $\mathcal{G}_1, \mathcal{G}_2$, and ψ_T , on the right-hand side of (3.8). Here the term \mathcal{G}_1 characterizes the error by approximating the mean-field state via N observed agents, the term \mathcal{G}_2 characterizes the error by estimating TQ with the batch of size n , and the term ψ_T characterizes the error in of approximating TQ by

functions from the RKHS $\mathcal{H}(K)$. If we further assume that $TQ \in \mathcal{H}(K)$ for all $Q \in \mathcal{H}(K)$ with $Q \leq Q_{\max}$, then term ψ_T vanishes and $Q^* \in \mathcal{H}(K)$ is the unique fixed point of T in $\mathcal{H}(K)$. Combining the error propagation in Theorem 3.2 and the one-step approximation error in Theorem 3.4, we obtain the following theorem that upper bounds the error of MF-FQI defined in Algorithm 1.

Theorem 3.5 (Theoretical Guarantee of MF-FQI). Let π_{κ} be the output of MF-FQI and $n = N^a$. Under the assumptions imposed by Theorem 3.2 and 3.4, for

$$\|\mathcal{C}\|_{\mathcal{H}(K)} \geq \max\{(a \cdot \log N/N)^{h/(c+3)}, 1/N^{ab/(bc+1)}\},$$

it holds with probability at least $1 - \eta - \tau$ that

$$\begin{aligned} \|Q^* - Q^{\pi_{\kappa}}\|_{1,\mu} &\leq \frac{2\gamma \cdot \Phi(\mu, \nu)}{(1-\gamma)^2} \cdot \left(\underbrace{C \cdot \Xi}_{(i)} + \underbrace{\psi_T}_{(ii)} \right) \\ &\quad + \underbrace{\frac{4\gamma^{\kappa+1} \cdot Q_{\max}}{1-\gamma}}_{(iii)}, \end{aligned} \quad (3.11)$$

where C is a positive absolute constant, ψ_T is defined in (3.10) of Theorem 3.4, and

$$\Xi = \max\left\{ \left(\frac{\log(|\mathcal{A}| \cdot N/\tau)}{N} \right)^{\frac{hc}{2(c+3)}}, 1/N^{\frac{abc}{2(bc+1)}} \right\}.$$

Here the integral operator \mathcal{C} is defined in (3.4) of Assumption 3.3, the parameter h is defined in Assumption 2.3, and parameters b, c are defined in Assumption 3.3.

Proof. See §B.5 for a detailed proof. \square

By Theorem 3.5, the approximation error of the action-value function attained by MF-FQI is characterized by the three terms on the right-hand side of (3.11). Here term (i) characterizes the statistical error, which is small for a sufficiently large number of observed agents N and batch size $n = N^a$. Term (ii) is the bias ψ_T defined in (3.10) of Theorem 3.4, which vanishes if the Bellman optimality operator T is closed in the RKHS $\mathcal{H}(K)$. Term (iii) characterizes the algorithmic error, which is small for a sufficiently large number of iterations κ .

In the sequel, we assume that T is closed in $\mathcal{H}(K)$ and thus $\psi_T = 0$ for simplicity. Note that the algorithmic error characterized by term (ii) has a linear rate of convergence, which is negligible comparing with the statistical error characterized by term (i) if the iteration number is sufficiently large. More specifically, if it holds for some positive absolute constant C that

$$\kappa \geq C \cdot \max\left\{ \frac{hc}{2(c+3) \cdot \log(1/\gamma)} \cdot \log\left(\frac{N}{\log(|\mathcal{A}| \cdot N)}\right), \frac{abc}{2(bc+1) \cdot \log(1/\gamma)} \cdot \log N \right\}, \quad (3.12)$$

then the dominating term on the right-hand side of (3.11) in Theorem 3.5 is term (i) that characterizes the statistical error.

Phase Transition. Note that MF-FQI involves two stage of sampling, where the first stage samples n mean-field state from the sampling distribution ν , and the second stage samples N states from each mean-field state. In what follows, we discuss the connection between the performance of MF-FQI and the sample complexity of the two-stage sampling involved. More specifically, we discuss the phase transition in the statistical error of MF-FQI when $a = \log n / \log N$ transits from zero to infinity. We categorize the phase transition into the following regimes in terms of a .

1. For $a > h \cdot (c + 1/b)/(c + 3)$, the rate of convergence of the statistical error takes the form of $\mathcal{O}((\log(|\mathcal{A}| \cdot N)/N)^{hc/(2c+6)})$. In this regime, increasing the number of observations N for each mean-field state improves the performance of MF-FQI, whereas increasing the batch size n of mean-field state cannot improve the performance of MF-FQI.
2. For $0 < a < h \cdot (c + 1/b)/(c + 3)$, the rate of convergence of the statistical error takes the form of $\mathcal{O}(1/n^{bc/(2bc+2)})$. In this regime, increasing batch size n of mean-field state improves the performance of MF-FQI, whereas increasing the number of observations N for each mean-field state cannot improve the performance of MF-FQI.

In conclusion, under regularity conditions, MF-FQI approximately achieves the optimal policy for a sufficiently large number of iteration κ , batch size n , and number of observations N . We highlight that MF-FQI enjoys a ‘‘blessing of many agents’’ property. More specifically, for a sufficiently large batch size n of the mean-field state, a larger number N of observed agents improves the learning of Q^* .

4. Limitation and Future Work

In this work, we propose MF-FQI, which tackles mean-field reinforcement learning problem with symmetric agents and a centralized controller. Such a setting is extensively studied in the analysis of societal-scale systems (Gomes et al., 2015; Guéant et al., 2011; Moll et al., 2019), such as the example of central bank or central government in §1. MF-FQI tackles the ‘‘curse of many agents’’ via mean embedding of the mean-field state for the (inexact) policy evaluation step, which approximately calculates the action-value function for the greedy policy. Based on the action-value function, we obtain a greedy policy, which corresponds to the policy improvement step. Such an approach is intractable when the action space also suffers from the ‘‘curse of many agents’’, as Q-learning requires taking the maximum over the action

space at each iteration, which can be combinatorially large if each agent takes its own action. However, the mean embedding technique is still applicable for the policy evaluation step even if each agent takes its own action, which can be coupled with other policy optimization methods, such as policy gradient (Sutton & Barto, 2018) and proximal policy optimization (Schulman et al., 2015; 2017). By replacing the greedy policy improvement step with other policy optimization methods, we are able to tackle the “curse of many agents” of both the state space and the action space, which is left as our future research.

References

- Acciaio, B., Backhoff-Veraguas, J., and Carmona, R. Extended mean field control problems: stochastic maximum principle and transport perspective. *arXiv preprint arXiv:1802.05754*, 2018.
- Altun, Y. and Smola, A. Unifying divergence minimization and statistical inference via convex duality. In *International Conference on Computational Learning Theory*, pp. 139–153. Springer, 2006.
- Andersson, D. and Djehiche, B. A maximum principle for SDEs of mean-field type. *Applied Mathematics & Optimization*, 63(3):341–356, 2011.
- Antos, A., Szepesvári, C., and Munos, R. Learning near-optimal policies with Bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- Bensoussan, A., Frehse, J., Yam, P., et al. *Mean field games and mean field type control theory*, volume 101. Springer, 2013.
- Bloem-Reddy, B. and Teh, Y. W. Probabilistic symmetry and invariant neural networks. *arXiv preprint arXiv:1901.06082*, 2019.
- Bu, L., Babu, R., De Schutter, B., et al. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- Buckdahn, R., Djehiche, B., Li, J., Peng, S., et al. Mean-field backward stochastic differential equations: a limit approach. *The Annals of Probability*, 37(4):1524–1565, 2009.
- Buckdahn, R., Djehiche, B., and Li, J. A general stochastic maximum principle for SDEs of mean-field type. *Applied Mathematics & Optimization*, 64(2):197–216, 2011.
- Caponnetto, A. and De Vito, E. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Carmona, R. and Delarue, F. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer, 2018.
- Carmona, R., Delarue, F., and Lachapelle, A. Control of McKean–Vlasov dynamics versus mean field games. *Mathematics and Financial Economics*, 7(2):131–166, 2013.
- Carmona, R., Delarue, F., et al. Forward–backward stochastic differential equations and controlled McKean–Vlasov dynamics. *The Annals of Probability*, 43(5):2647–2700, 2015.
- Chen, J. and Jiang, N. Information-theoretic considerations in batch reinforcement learning. *arXiv preprint arXiv:1905.00360*, 2019.
- Claus, C. and Boutilier, C. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI*, 1998:746–752, 1998.
- Džeroski, S., De Raedt, L., and Driessens, K. Relational reinforcement learning. *Machine Learning*, 43(1-2):7–52, 2001.
- Ernst, D., Geurts, P., and Wehenkel, L. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.
- Farahmand, A.-m., Szepesvári, C., and Munos, R. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*, pp. 568–576, 2010.
- Farahmand, A.-m., Ghavamzadeh, M., Szepesvári, C., and Mannor, S. Regularized policy iteration with nonparametric function spaces. *Journal of Machine Learning Research*, 17(1):4809–4874, 2016.
- Fornasier, M. and Solombrino, F. Mean-field optimal control. *ESAIM: Control, Optimisation and Calculus of Variations*, 20(4):1123–1152, 2014.
- Fukumizu, K., Gretton, A., Sun, X., and Schölkopf, B. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems*, pp. 489–496, 2008.
- Gärtner, T., Flach, P. A., Kowalczyk, A., and Smola, A. J. Multi-instance kernels. In *International Conference on Machine Learning*, volume 2, pp. 7, 2002.

- Gomes, D. A., Nurbekyan, L., and Pimentel, E. Economic models and mean-field games theory. In *30th Brazilian Mathematics Colloquium*, 2015.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. J. A kernel method for the two-sample problem. In *Advances in Neural Information Processing Systems*, pp. 513–520, 2007.
- Gretton, A., Smola, A., Huang, J., Schmittfull, M., Borgwardt, K., and Schölkopf, B. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4): 5, 2009.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- Guéant, O., Lasry, J.-M., and Lions, P.-L. Mean field games and applications. In *Paris-Princeton lectures on mathematical finance 2010*, pp. 205–266. Springer, 2011.
- Guestrin, C., Lagoudakis, M., and Parr, R. Coordinated reinforcement learning. In *ICML*, volume 2, pp. 227–234. Citeseer, 2002.
- Guo, X., Hu, A., Xu, R., and Zhang, J. Learning mean-field games. *arXiv preprint arXiv:1901.09585*, 2019.
- Haussler, D. Convolution kernels on discrete structures. Technical report, Department of Computer Science, University of California, 1999.
- Hu, J. and Wellman, M. P. Nash Q-learning for general-sum stochastic games. *Journal of Machine Learning Research*, 4(Nov):1039–1069, 2003.
- Huang, M., Caines, P. E., and Malhamé, R. P. Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions. In *42nd IEEE International Conference on Decision and Control*, volume 1, pp. 98–103. IEEE, 2003.
- Huang, M., Caines, P. E., and Malhamé, R. P. Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized ε -Nash equilibria. *IEEE Transactions on Automatic Control*, 52(9):1560–1571, 2007.
- Huang, M., Caines, P. E., and Malhamé, R. P. Social optima in mean field LQG control: centralized and decentralized strategies. *IEEE Transactions on Automatic Control*, 57(7):1736–1751, 2012.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.
- Jiachen, Y., Xiaojing, Y., Rakshit, T., Huan, X., and Hongyuan, Z. Learning deep mean field games for modeling large population behavior. *arXiv preprint arXiv:1711.03156*, 2017.
- Jiang, J. and Lu, Z. Learning attentional communication for multi-agent cooperation. In *Advances in Neural Information Processing Systems*, pp. 7254–7264, 2018.
- Kar, S., Moura, J. M., and Poor, H. V. QD-learning: A collaborative distributed strategy for multi-agent reinforcement learning through Consensus + Innovations. *IEEE Transactions on Signal Processing*, 61(7):1848–1862, 2013.
- Kondor, R. and Jebara, T. A kernel between sets of vectors. In *International Conference on Machine Learning*, pp. 361–368, 2003.
- Lasry, J.-M. and Lions, P.-L. Jeux à champ moyen. i–le cas stationnaire. *Comptes Rendus Mathématique*, 343(9): 619–625, 2006a.
- Lasry, J.-M. and Lions, P.-L. Jeux à champ moyen. ii–horizon fini et contrôle optimal. *Comptes Rendus Mathématique*, 343(10):679–684, 2006b.
- Lasry, J.-M. and Lions, P.-L. Mean field games. *Japanese journal of mathematics*, 2(1):229–260, 2007.
- Lauer, M. and Riedmiller, M. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *International Conference on Machine Learning*. Citeseer, 2000.
- Lazaric, A., Ghavamzadeh, M., and Munos, R. Analysis of classification-based policy iteration algorithms. *Journal of Machine Learning Research*, 17(1):583–612, 2016.
- Li, M., Jiao, Y., Yang, Y., Gong, Z., Wang, J., Wang, C., Wu, G., Ye, J., et al. Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. *arXiv preprint arXiv:1901.11454*, 2019.
- Lin, K., Zhao, R., Xu, Z., and Zhou, J. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1774–1783. ACM, 2018.
- Lin, S.-B., Guo, X., and Zhou, D.-X. Distributed learning with regularized least squares. *Journal of Machine Learning Research*, 18(1):3202–3232, 2017.
- Lin, Y., Zhang, K., Yang, Z., Wang, Z., Basar, T., Sandhu, R., and Liu, J. A communication-efficient multi-agent actor-critic algorithm for distributed reinforcement learning. *arXiv preprint arXiv:1907.03053*, 2019.

- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings*, pp. 157–163. Elsevier, 1994.
- Meyer-Brandis, T., Øksendal, B., and Zhou, X. Y. A mean-field stochastic maximum principle via Malliavin calculus. *Stochastics An International Journal of Probability and Stochastic Processes*, 84(5-6):643–666, 2012.
- Moll, B., Rachel, L., and Restrepo, P. Uneven growth: automation’s impact on income and wealth inequality. *Manuscript, Princeton University*, 2019.
- Muandet, K., Fukumizu, K., Dinuzzo, F., and Schölkopf, B. Learning from distributions via support measure machines. In *Advances in Neural Information Processing Systems*, pp. 10–18, 2012.
- Munos, R. and Szepesvári, C. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9 (May):815–857, 2008.
- Nash, J. Non-cooperative games. *Annals of Mathematics*, pp. 286–295, 1951.
- Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.
- Ng, Y.-K. Bentham or Bergson? Finite sensibility, utility functions and social welfare functions. *The Review of Economic Studies*, 42(4):545–569, 1975.
- OpenAI. Openai Five. <https://blog.openai.com/openai-five/>, 2018.
- Panait, L. and Luke, S. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.
- Scherrer, B. On the performance bounds of some policy search dynamic programming algorithms. *arXiv preprint arXiv:1306.0539*, 2013.
- Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B., and Geist, M. Approximate modified policy iteration and its application to the game of Tetris. *Journal of Machine Learning Research*, 16:1629–1676, 2015.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning*, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shalev-Shwartz, S., Shammah, S., and Shashua, A. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354, 2017.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. A Hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, pp. 13–31. Springer, 2007.
- Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. R. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11(Apr):1517–1561, 2010.
- Suttle, W., Yang, Z., Zhang, K., Wang, Z., Basar, T., and Liu, J. A multi-agent off-policy actor-critic algorithm for distributed reinforcement learning. *arXiv preprint arXiv:1903.06372*, 2019.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- Szabó, Z., Gretton, A., Póczos, B., and Sriperumbudur, B. Two-stage sampled learning theory on distributions. In *Artificial Intelligence and Statistics*, pp. 948–957, 2015.
- Szepesvári, C. and Munos, R. Finite time bounds for sampling based fitted value iteration. In *International Conference on Machine Learning*, 2005.
- Tan, M. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *International Conference on Machine Learning*, pp. 330–337, 1993.
- Tolstikhin, I., Sriperumbudur, B. K., and Muandet, K. Minimax estimation of kernel mean embeddings. *Journal of Machine Learning Research*, 18(1):3002–3048, 2017.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., Dudzik, A., Huang, A., Georgiev, P., Powell, R., Ewalds, T., Horgan, D., Kroiss, M., Danihelka, I., Agapiou, J., Oh, J., Dalibard, V., Choi, D., Sifre, L., Sulsky, Y., Vezhnevets, S., Molloy, J., Cai, T., Budden, D., Paine, T., Gulcehre, C., Wang, Z., Pfaff, T., Pohlen, T., Wu, Y., Yogatama, D., Cohen, J., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Apps, C., Kavukcuoglu, K., Hassabis, D., and Silver,

- D. AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>, 2019.
- Wang, X. and Sandholm, T. Reinforcement learning to play an optimal Nash equilibrium in team Markov games. In *Advances in Neural Information Processing Systems*, pp. 1603–1610, 2003.
- Yang, E. and Gu, D. Multiagent reinforcement learning for multi-robot systems: A survey. Technical report, University of Essex, 2004.
- Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W., and Wang, J. Mean field multi-agent reinforcement learning. *arXiv preprint arXiv:1802.05438*, 2018.
- Yang, Z., Xie, Y., and Wang, Z. A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*, 2019.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in Neural Information Processing Systems*, pp. 3391–3401, 2017.
- Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E., et al. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*, 2018.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Zhang, K., Yang, Z., Liu, H., Zhang, T., and Başar, T. Finite-sample analyses for fully decentralized multi-agent reinforcement learning. *arXiv preprint arXiv:1812.02783*, 2018a.
- Zhang, K., Yang, Z., Liu, H., Zhang, T., and Başar, T. Fully decentralized multi-agent reinforcement learning with networked agents. *arXiv preprint arXiv:1802.08757*, 2018b.
- Zhang, K., Yang, Z., and Başar, T. Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. In *Advances in Neural Information Processing Systems*, pp. 11598–11610, 2019a.
- Zhang, K., Yang, Z., and Başar, T. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *arXiv preprint arXiv:1911.10635*, 2019b.