

Multi Type Mean Field Reinforcement Learning

Sriram Ganapathi Subramanian*
 University of Waterloo
 Waterloo, Ontario
 s2ganapa@uwaterloo.ca

Matthew E. Taylor
 Borealis AI
 Edmonton, Alberta
 matthew.taylor@borealisai.com

Pascal Poupart
 Borealis AI
 Waterloo, Ontario
 pascal.poupart@borealisai.com

Nidhi Hegde
 Borealis AI
 Edmonton, Alberta
 nidhi.hegde@borealisai.com

ABSTRACT

Mean field theory provides an effective way of scaling multiagent reinforcement learning algorithms to environments with many agents that can be abstracted by a virtual mean agent. In this paper, we extend mean field multiagent algorithms to multiple types. The types enable the relaxation of a core assumption in mean field games, which is that all agents in the environment are playing almost similar strategies and have the same goal. We conduct experiments on three different testbeds for the field of many agent reinforcement learning, based on the standard MAgents framework. We consider two different kinds of mean field games: a) Games where agents belong to predefined types that are known a priori and b) Games where the type of each agent is unknown and therefore must be learned based on observations. We introduce new algorithms for each type of game and demonstrate their superior performance over state of the art algorithms that assume that all agents belong to the same type and other baseline algorithms in the MAgent framework.

KEYWORDS

Mean Field Methods; Multiagent Systems; Reinforcement Learning; Many-Agent Learning

ACM Reference Format:

Sriram Ganapathi Subramanian, Pascal Poupart, Matthew E. Taylor, and Nidhi Hegde. 2020. Multi Type Mean Field Reinforcement Learning. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

Multiagent Reinforcement Learning (MARL) is a quickly growing field with lots of recent research pushing its boundaries [6, 14]. Yet, scaling the multiagent algorithms to environments with a large number of learning agents continues to be a problem [5]. Research advances in the field of MARL [2, 4] deal with only a limited number of agents and the proposed methods cannot be easily generalized to more complex scenarios with many agents. Some recent research has used the concept of mean field theory for enabling the use of MARL approaches to environments with many agents [12, 17]. Yet,

*Sriram did this work while he was an intern at Borealis AI.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

the current algorithms implemented for many agents require some strong assumptions about the game environment to perform adequately. The important mean field approximation reduces a many agent problem into a simplified two agent problem where all the other competing agents are approximated as a single mean field. This mean field approximation, however, would be valid only for scenarios where all the agents in the environment can be considered similar to each other in objectives and abilities. Real world applications often have a set of agents that are diverse and therefore it is virtually impossible to aggregate them into a single mean field.

In this paper, we introduce a concept of multiple types to model agent diversity in the mean field approximation for many agent reinforcement learning. The types are groupings applied to other agents in the environment in such a way that all members of a particular type play approximately similar strategies and have similar overall goals. Now, the modelling agent can consider each type to be a distinct agent, which has to be modelled separately. Within each type, the mean field approximation should still be reasonable as the agents within one particular type are more related to each other than other agents from different types. Thus, the many agent interaction is effectively reduced to N agent interactions where N is the number of types. Note that this is more complex than the simple two agent interaction considered by previous research and this can approximate a real world dynamic environment in a much better way. Most real world applications for many agent reinforcement learning can be broadly classified into two categories. The first category involves applications where we have predefined types and the type of each agent is known a priori. Some common applications include games with multiple teams (like quiz competitions), multiple party elections with coalitions, airline price analysis with multiple airlines forming an alliance beforehand, etc. We call these applications predefined *known* type scenarios. The other category are applications that involve a large number of agents that may have different policies due to differences in their rewards, actions or observations. Common examples are demand and supply scenarios, stock trading scenarios, etc., where one type of agent may be risk averse while another type may be risk seeking. We call these applications pre-defined *unknown* type scenarios, since there are no true underlying types like the previous case and a suitable type therefore must be assigned through observations. Another important aspect in this paper will be the notion of *neighbourhood* since each individual agent may be impacted more by agents whose states are “closer” (according to some distance measure) to the state

of an agent. For instance, in battle environments, nearby agents pose a greater threat than far away agents.

Using the open source test environment for many agent games – MAgents [18] – we consider three testbeds that involve many strategic agents in the environment. Two of these testbeds correspond to the known type and the third one corresponds to the unknown type. We introduce two different algorithms for the known and unknown type scenarios. We demonstrate the superior performance of these algorithms in comparison to previous algorithms that assume a single type of agents in the testbeds.

2 BACKGROUND

Reinforcement Learning: Single agent reinforcement learning (RL) [15] is the most common form of reinforcement learning in the literature [1]. Here the problem is modelled in the framework of a Markov Decision Process (MDP) where the MDP is composed of $\langle S, A, P, R \rangle$, where S denotes the set of states that the agent can move into, A denotes the set of actions that the agent can take, P denotes the transition distribution ($P(s'|s, a)$) and R denotes the reward function ($R(s, a)$). The agents are allowed to explore the environment during the process of training and collect experience tuples $\langle s, a, s', r \rangle$. An agent learns a policy $\pi : S \rightarrow A$, which is a mapping from states to actions where the goal is to maximize the expected cumulative rewards $\sum_t \gamma^t R(s_t, a_t)$, where $\gamma \in [0, 1]$ is the discount factor. An optimal policy obtains the highest cumulative rewards among all possible policies.

In Multiagent Reinforcement Learning (MARL), there is a notion of stochastic games [3], where the state and action space are defined as Cartesian products of individual states and actions of different agents in the environment. A stochastic game can be considered a special type of normal form game [9], where a particular iteration of the game depends on previous game(s) played and the experiences of all the agents in the previous game(s). A stochastic game can be defined as a tuple $\langle S, N, \mathbf{A}, P, R \rangle$ where S is a finite set of states (assumed to be the same for all agents), N is a finite set of n agents, $\mathbf{A} = A^1 \times \dots \times A^n$ where A^j denotes the actions of agent j . P is the transition distribution $P(s'|s, \mathbf{a})$ where $\mathbf{a} = (a^1, \dots, a^n)$ and $R_j(s, \mathbf{a}) = r^j$ is the reward function with r^j denoting the reward received by agent j . Here each agent is trying to learn a policy that maximizes its return upon consideration of opponent behaviours. Agents are typically self interested and the combined system moves towards a Nash equilibrium [11]. Scalability in MARL environments is often a bottleneck. Important research efforts are aimed at handling only up to a handful of agents and the solutions or algorithms considered become intractable in many agent scenarios.

Mean Field Games: Mean field theory approximates many agent interactions in a multiagent environment into two agent interactions [10] where the second agent corresponds to the mean effect of all the other agents. This allows domains with many agents that were previously considered intractable to be revisited and scalable approximate solutions to be devised [12, 17]. A mean field game is defined to be a special kind of stochastic game. In the paper by Yang et al. [17], the Q function for mean field games is decomposed additively into local Q functions that capture pairwise interactions:

$$Q^j(s, \mathbf{a}) = \frac{1}{n^j} \sum_{k \in \eta(j)} Q^j(s, a^j, a^k). \quad (1)$$

Here n^j is the number of neighbours of agent j and $\eta(j)$ is the index set of neighbouring agents. Yang et al. [17] showed that this decomposition can be well approximated by the mean field Q -function $Q^j(s, \mathbf{a}) \approx Q_{MF}^j(s, a^j, \bar{a}^j)$ under certain conditions. The mean action \bar{a}^j on the neighbourhood $\eta(j)$ of agent j is expressed as $\bar{a}^j = \frac{1}{n^j} \sum_{k \in \eta(j)} a^k$ where a^k is the action of each neighbour k . In the case of discrete actions, a^k is a one-hot vector encoding and \bar{a}^j is a vector of fractions corresponding to the probability that each action may be executed by an agent at random.

The mean field Q function can be updated in a recurrent manner as follows:

$$Q_{t+1}^j(s, a^j, \bar{a}^j) = (1 - \alpha)Q_t^j(s, a^j, \bar{a}^j) + \alpha[r^j + \gamma v_t^j(s')] \quad (2)$$

where r^j is the reward obtained. The s, s' are the old and new states respectively. α_t is the learning rate. The value function $v_t^j(s')$ for agent j at time t is given by:

$$v_t^j(s') = \sum_{a^j} \pi_t^j(a^j | s', \bar{a}^j) \mathbb{E}_{a_{-j} \sim \pi_{-j}} Q_t^j(s', a^j, \bar{a}^j) \quad (3)$$

Here, the term \bar{a}^j denotes the mean action of all the other agents apart from j . The mean action for all the agents is first calculated using the relation, $\bar{a}^j = \frac{1}{N^j} \sum_k a^k$, $a^k \sim \pi_t^k(\cdot | s, \bar{a}_-^k)$, where π_t^k is the policy of agent k and \bar{a}_-^k represents the previous mean action for the neighbours of agent k . N^j denotes the total number of agents in the neighbourhood of j . Then, the Boltzmann policy for each agent j is calculated using β , which is the Boltzmann softmax parameter

$$\pi_t^j(a^j | s, \bar{a}^j) = \frac{\exp(-\beta Q_t^j(s, a^j, \bar{a}^j))}{\sum_{a^j \in A^j} \exp(-\beta Q_t^j(s, a^j, \bar{a}^j))} \quad (4)$$

3 MEAN FIELD MARL AND TYPES

We consider environments where there are M types that the neighbouring agents can be classified into. In this paper, we assume that the Q function decomposes additively according to a partition of the agents into X^j subsets that each include one agent of each type. This decomposition can be viewed as a generalization of pairwise decompositions to multiple types since each term depends on a representative from each type.

Let the standard Q function be $Q^j(s, \mathbf{a})$.

$$Q^j(s, \mathbf{a}) = \frac{1}{X^j} \sum_{i=1}^{X^j} [Q^j(s, a^j, a_1^{k_i}, a_2^{k_i}, \dots, a_M^{k_i})] \quad (5)$$

There are a total of M types and a_m^k denotes the action of agent k belonging to type m in the neighbourhood of agent j . Notice that this representation of the Q function includes the interaction with each one of the types and is not a simple pairwise interaction as done by [17]. Let us assume that we have a scheme in which we can classify each agent into one of these subsets. Note that we do not need each group to contain an equal number of agents as we can always make a new subset that contains one agent of a type and other agents to be place holder agents (dead agents) of other types. We will relax this requirement of decomposition shortly and in practice we do not need to make these subsets at all.

3.1 Mean Field Approximation

We also assume discrete action spaces and use a one hot representation of the actions as in [17]. The one hot action of each agent k belonging to type m in the neighbourhood of agent j is represented as $a_m^{k,m} = \bar{a}_m^j + \hat{\delta}^{j,k,m}$ where \bar{a}_m^j is the mean action of all agents in the neighbourhood of agent j belonging to type m and $\hat{\delta}^{j,k,m}$ is the deviation between the action of an agent and the mean action of its type.

Let $\delta^{j,k_i} = [\hat{\delta}^{j,k_1}; \hat{\delta}^{j,k_2}; \dots; \hat{\delta}^{j,k_M}]$ be a vector obtained by the concatenation of all such deviations of the agents in the neighbourhood of agent j belonging to each of the M types (all agents of a single subset). Similar to [17], we apply Taylor's Theorem to expand the Q function in Equation 5:

$$\begin{aligned} Q^j(s, \mathbf{a}) &= \frac{1}{X^j} \sum_{i=1}^{X^j} Q^j(s, a^j, a_1^{k_i}, a_2^{k_i}, \dots, a_M^{k_i}) \\ &= \frac{1}{X^j} \sum_{i=1}^{X^j} [Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \\ &\quad + \nabla_{\bar{a}_1^j, \dots, \bar{a}_M^j} Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \cdot \delta^{j,k_i} \\ &\quad + \frac{1}{2} \delta^{j,k_i} \cdot \nabla_{\bar{a}_1^j, \dots, \bar{a}_M^j}^2 Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \cdot \delta^{j,k_i}] \\ &= Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \\ &\quad + \nabla_{\bar{a}_1^j, \dots, \bar{a}_M^j} Q^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j) \cdot [\frac{1}{X^j} \sum_{i=1}^{X^j} \delta^{j,k_i}] \\ &\quad + \frac{1}{2X^j} \sum_{i=1}^{X^j} [\delta^{j,k_i} \cdot \nabla_{\bar{a}_1^j, \dots, \bar{a}_M^j}^2 Q^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j) \cdot \delta^{j,k_i}] \\ &= Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) + \frac{1}{2X^j} \sum_{i=1}^{X^j} [R_{s,a^j}^j(a^{k_i})] \approx Q^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \end{aligned}$$

where $R_{s,a^j}^j(a^k) \triangleq \delta^{j,k} \cdot \nabla_{\bar{a}_1^j, \dots, \bar{a}_M^j}^2 Q^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j) \cdot \delta^{j,k}$

which is the Taylor polynomial remainder. The summation term $[\frac{1}{X^j} \sum_{i=1}^{X^j} \delta^{j,k_i}]$ sums out to 0. Finally, if we ignore the remainder terms R_{s,a^j}^j , we obtain the following approximation:

$$Q^j(s, \mathbf{a}) \approx Q_{MTMF}^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \quad (6)$$

The magnitude of this approximation depends on the deviation $\hat{\delta}^{j,k,m}$ between each action $a_m^{k,m}$ and its mean field approximation \bar{a}_m^j . More precisely, we can quantify the overall effect of the mean field approximation by the average deviation $\sum_k \|\hat{\delta}_k\|_2 / N$. Theorems 3.1 and 3.2 show that the average deviation is reduced as we increase the number of types and Theorem 3.3 provides an explicit bound on the approximation in Eq. 6 based on the average deviation.

THEOREM 3.1. *When there are two types in the environment, but they have been considered to be the same type, the average deviation induced by the mean field approximation is bounded as follows:*

$$\frac{\sum_k \|\hat{\delta}_k\|_2}{N} \leq \frac{K_1}{N} \epsilon_1 + \frac{K_2}{N} \epsilon_2 + \frac{K_1}{N} \alpha_1 + \frac{K_2}{N} \alpha_2 \quad (7)$$

where N denotes the total number of agents in the environment; K_1 and K_2 denote the total number of agents of types 1 and 2, respectively; ϵ_1 and ϵ_2 are bounds on the average deviation for agents of types 1 and 2 respectively.

$$\frac{1}{K_1} (\sum_{k_1} \|a_{k_1} - \bar{a}_1\|) \leq \epsilon_1; \frac{1}{K_2} (\sum_{k_2} \|a_{k_2} - \bar{a}_2\|) \leq \epsilon_2$$

Similarly, α_1 and α_2 denote the errors induced by using a single type (instead of two types) in the mean field approximation.

$$\alpha_1 = \|\bar{a}_1 - \bar{a}\|; \alpha_2 = \|\bar{a}_2 - \bar{a}\| \quad (8)$$

Furthermore, a_{k_1} denotes an action of an agent belonging to type 1 and a_{k_2} denotes an action of an agent belonging to type 2. Here \bar{a} denotes the mean field action of all the agents, \bar{a}_1 denotes the mean action of all agents of type 1 and \bar{a}_2 denotes the mean action of all agents of type 2.

PROOF. Since we have considered every agent to belong to only one type, the deviation between the agent's action and the overall mean action is the error estimate. Hence, we will have the following,

$$\frac{\sum_k \|\hat{\delta}_k\|_2}{N} = \frac{\sum_k \|a^j - \bar{a}^j\|_2}{N}$$

$$\frac{\sum_k \|\hat{\delta}_k\|_2}{N} = \frac{1}{N} (\sum_{k_1} \|a_{k_1} - \bar{a}\| + \sum_{k_2} \|a_{k_2} - \bar{a}\|)$$

The superscript j has been dropped for simplicity.

$$\begin{aligned} &= \frac{1}{N} (\sum_{k_1} \|a_{k_1} - \bar{a}_1 + \bar{a}_1 - \bar{a}\| + \sum_{k_2} \|a_{k_2} - \bar{a}_2 + \bar{a}_2 - \bar{a}\|) \\ &\leq \frac{1}{N} (\sum_{k_1} \|a_{k_1} - \bar{a}_1\| + \sum_{k_1} \|\bar{a}_1 - \bar{a}\| + \sum_{k_2} \|a_{k_2} - \bar{a}_2\| + \sum_{k_2} \|\bar{a}_2 - \bar{a}\|) \\ &= \frac{1}{N} (\sum_{k_1} \|a_{k_1} - \bar{a}_1\| + K_1 \|\bar{a}_1 - \bar{a}\| + \sum_{k_2} \|a_{k_2} - \bar{a}_2\| + K_2 \|\bar{a}_2 - \bar{a}\|) \\ &\leq \frac{K_1}{N} \epsilon_1 + \frac{K_2}{N} \epsilon_2 + \frac{K_1}{N} \alpha_1 + \frac{K_2}{N} \alpha_2 \end{aligned} \quad (9)$$

□

THEOREM 3.2. *When there are two types in the environment and they have been considered to be different types, the average deviation induced by the mean field approximation is bounded as follows:*

$$\frac{\sum_k \|\hat{\delta}_k\|_2}{N} \leq \frac{K_1}{N} \epsilon_1 + \frac{K_2}{N} \epsilon_2 \quad (10)$$

The variables have the same meaning as in Theorem 3.1.

PROOF. In this scenario we will have,

$$\frac{\sum_k \|\hat{\delta}_k\|_2}{N} = \frac{1}{N} (\sum_{k_1} \|a_{k_1} - \bar{a}_1\| + \sum_{k_2} \|a_{k_2} - \bar{a}_2\|) \quad (11)$$

$$= \frac{K_1}{N} \frac{\sum_{k_1} \|a_{k_1} - \bar{a}_1\|}{K_1} + \frac{K_2}{N} \frac{\sum_{k_2} \|a_{k_2} - \bar{a}_2\|}{K_2} \leq \frac{K_1}{N} \epsilon_1 + \frac{K_2}{N} \epsilon_2$$

□

We presented Theorems 3.1 and 3.2 to demonstrate the reduction in the bound on the average deviation as we increase the number of types from 1 to 2. Similar derivations can be performed to demonstrate that the bounds on the average deviation decrease as we increase the number of types (regardless of the true number of types).

Let ϵ be a bound on the average deviation achieved based on a certain number of types: $\frac{\sum_{k=1}^X \|\delta\|_2}{X} \leq \epsilon$. The following theorem bounds the error of the approximate mean field Q -function as a function of ϵ and the smoothness L of the exact Q function.

THEOREM 3.3. *When the Q -function is additively decomposable according to Equation 5 and it is L -smooth, then the multi-type mean field Q function provides a good approximation bounded by*

$$|Q^j(s, \mathbf{a}) - Q_{MTMF}^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j)| \leq \frac{1}{2} L \epsilon \quad (12)$$

PROOF. We rewrite the expression for the Q function as $Q(a) \triangleq Q^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j)$. Suppose that Q is L -smooth, where its gradient ∇Q is Lipschitz-continuous with constant L such that for all a, \bar{a}

$$\|\nabla Q(a) - \nabla Q(\bar{a})\|_2 \leq L\|a - \bar{a}\|_2 \quad (13)$$

where $\|\cdot\|_2$ indicates the l_2 -norm. Note that all the eigenvalues of $\nabla^2 Q$ can be bounded in the symmetric interval $[-L, L]$. As the Hessian matrix $\nabla^2 Q$ is real symmetric and hence diagonalizable, there exists an orthogonal matrix U such that $U^T[\nabla^2 Q]U = \Lambda \triangleq \text{diag}[\lambda_1, \dots, \lambda_D]$. It then follows that:

$$\delta a \cdot \nabla^2 Q \cdot \delta a = [U\delta a]^T \Lambda [U\delta a] = \sum_{i=1}^D \lambda_i [U\delta a]_i^2 \quad (14)$$

with

$$-L\|U\delta a\|_2 \leq \sum_{i=1}^D \lambda_i [U\delta a]_i^2 \leq L\|U\delta a\|_2 \quad (15)$$

Let, $\hat{\delta} a = a_m^j - \bar{a}_m^j$, consistent with the previous definition. Recall that the term δ is then $\delta a = \hat{\delta}_1 a + \hat{\delta}_2 a + \dots + \hat{\delta}_M a$, where a is the one-hot encoding for D actions, and \bar{a} is a D -dimensional categorical distribution. Then, it can be shown that

$$\|U(\delta a)\|_2 = \|\delta a\|_2 \quad (16)$$

Consider the term $\frac{1}{2X} \sum_{k=1}^X R_{s, a^k}^j(a^k)$. Since L is the maximum eigenvalue, from Equation 14 (with a slight abuse of notation):

$$R = \delta a \cdot \nabla^2 Q \cdot \delta a = [U\delta a]^T \Lambda [U\delta a] = \sum_i \lambda_i [U\delta a]_i^2 \leq L\|U\delta a\|_2^2 \quad (17)$$

Therefore, from Equation 16:

$$R \leq L\|U\delta a\|_2 = L\|\delta a\|_2 \quad (18)$$

Thus,

$$\frac{1}{2X} \sum_k R \leq \frac{L}{2} \sum_k \frac{\|\delta a\|_2}{X} \leq \frac{L\epsilon}{2} \quad (19)$$

□

Thus, in this paper, we modify the mean field Q function to include a finite number of types that each have a corresponding mean field. The Q function then considers a finite number of interactions across types.

3.2 Mean Field Update

The mean action \bar{a}_i^j represents the mean action of the neighbours of agent j belonging to type i . As in the paper by Yang et al. [17], the mean field Q function can be updated in a recurrent manner.

$$Q_{t+1}^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j) = (1 - \alpha)Q_t^j(s, a^j, \bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j) + \alpha[r^j + \gamma v_t^j(s')] \quad (20)$$

where r^j is the reward obtained. The s and s' are the old and new states respectively. α_t is the learning rate. The value function $v_t^j(s')$ for agent j at time t is given by:

$$v_t^j(s') = \sum_{a^j} \pi_t^j(a^j|s', \bar{a}_1^j, \dots, \bar{a}_M^j) \mathbb{E}_{\bar{a}_i^j \sim \pi_i^j} [Q_t^j(s', a^j, \bar{a}_1^j, \dots, \bar{a}_M^j)] \quad (21)$$

Here, the term \bar{a}_i^j denotes the mean action of all the other agents apart from j belonging to type i . In all of our implementations, the mean action for all the types is first calculated using the relation

$$\bar{a}_i^j = \frac{1}{N_i^j} \sum_k a_i^k, a_i^k \sim \pi_i^k(\cdot|s, \bar{a}_{i-}^k, \dots, \bar{a}_{M-}^k) \quad (22)$$

where π_i^k is the policy of agent k (in j 's neighbourhood) and \bar{a}_{i-}^k represents the previous mean action for the neighbours of agent k belonging to type i . N_i^j is the total number of agents of type i in j 's neighbourhood. Then, the Boltzmann policy for each agent j is

$$\pi_i^j(a^j|s, \bar{a}_1^j, \dots, \bar{a}_M^j) = \frac{\exp(\beta Q_t^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j))}{\sum_{a^{j'} \in \mathcal{A}^j} \exp(\beta Q_t^j(s, a^{j'}, \bar{a}_1^j, \dots, \bar{a}_M^j))} \quad (23)$$

where β is the Boltzmann softmax parameter.

By iterating through Equations 21, 22 and 23, the mean actions and respective policies of all agents keep improving. We prove in Theorem 3.4 that this approach converges to a fixed point within a small bounded distance of the Nash equilibrium. In Appendix A, we give a specific example for a case in which the Multi Type Mean Field algorithm does better than the simple Mean Field method.

We first make three mild assumptions about the Multi Type Mean Field update and then state two lemmas before giving Theorem 3.4.

Assumption 1: In the Multi Type Mean Field update, each action-value pair is visited infinitely often, and the reward is bounded by some constant.

Assumption 2: The agents policies are Greedy in the Limit with Infinite Exploration (GLIE). In the case of the Boltzmann policy, the policy becomes greedy w.r.t. the Q -function in the limit as the temperature decays asymptotically to zero.

Assumption 3: For each stage game $[Q_t^1(s), \dots, Q_t^N(s)]$ at time t and in state s in training, for all $t, s, j \in 1, \dots, N$ the Nash equilibrium $\pi_* = [\pi_*^1, \dots, \pi_*^N]$ is recognized either as a global optimum or a saddle point as expressed as:

$$1. \mathbb{E}_{\pi_*} [Q_t^j(s)] \geq \mathbb{E}_{\pi} [Q_t^j(s)], \forall \pi \in \Omega(\Pi_k \setminus \mathcal{A}^k) \quad (24)$$

$$2. \mathbb{E}_{\pi_*} [Q_t^j(s)] \geq \mathbb{E}_{\pi_*^j} \mathbb{E}_{\pi_*^{-j}} [Q_t^j(s)], \forall \pi^j \in \Omega(\mathcal{A}^k) \quad (25)$$

$$3. \mathbb{E}_{\pi_*} [Q_t^j(s)] \leq \mathbb{E}_{\pi_*^j} \mathbb{E}_{\pi_*^{-j}} [Q_t^j(s)], \forall \pi^{-j} \in \Omega(\Pi_k \neq j \setminus \mathcal{A}^k) \quad (26)$$

LEMMA 1. Under Assumption 3, the Nash operator \mathcal{H}^{Nash} forms a contraction mapping under the complete metric space from Q to Q with the fixed point being the Nash Q value of the entire game. $\mathcal{H}^{Nash} Q_* = Q_*$

PROOF. Refer to Theorem 17 in [7] for a detailed proof. □

We define a new operator \mathcal{H}^{MTMF} which is the Multi Type Mean Field operator. We differentiate this from the Nash operator defined above. This operator is defined as:

$$\mathcal{H}^{MTMF} \mathbf{Q}(s, \mathbf{a}) = \mathbb{E}_{s' \sim p} [\mathbf{r}(s, \mathbf{a}) + \gamma \mathbf{v}^{MTMF}(s')] \quad (27)$$

The Multi Type Mean Field value function is defined as $v^{MTMF} \triangleq [v^1(s), \dots, v^N(s)]$. This is the value function obtained from Equation 21. Now using the same principle of Lemma 1 on the multi type mean field operator, we can show that the Multi Type Mean Field operator also forms a contraction mapping (additionally refer to Proposition 1 in [17]).

LEMMA 2. The random process Δ_t defined in \mathcal{R} as

$$\Delta_{t+1}(x) = (1 - \alpha)\Delta_t(x) + \alpha F_t(x) \quad (28)$$

converges to a constant S with probability 1 (w.p.t 1) when

$$1) \quad 0 \leq \alpha \leq 1, \sum_t \alpha = \infty, \sum_t \alpha^2 < \infty \quad (29)$$

$$2) \quad x \in \mathcal{X}; |\mathcal{X}| < \infty \quad (30)$$

where \mathcal{X} is the set of possible states,

$$3) \quad \|\mathbb{E}[F_t(x)|\mathcal{F}_t]\|_w \leq \gamma \|\Delta_t\|_w + K \quad (31)$$

where $\gamma \in [0, 1)$ and K is finite.

$$4) \quad \text{var}[F_t(x)|\mathcal{F}_t] \leq K_2(1 + \|\Delta_t\|_w^2) \quad (32)$$

with constant $K_2 > 0$ and finite.

Here \mathcal{F}_t denotes the filtration of an increasing sequence of σ -fields including the history of processes; $\alpha_t, \Delta_t, F_t \in \mathcal{F}_t$ and $\|\cdot\|_w$ is a weighted maximum norm. The value of this constant $S = \frac{\psi C_1 + \alpha |K|}{\alpha \beta_0}$ where $\psi \in (0, 1)$ and C_1 is the value with which the scale invariant iterative process is bounded. β_0 is the scale factor applied to the original process.

PROOF. This lemma follows from Theorem 1 in [8]. We provide the complete proof of this lemma in the Appendix C, highlighting the changes from the Theorem 1 in [8]. \square

THEOREM 3.4. When updating $Q_{MTMF}^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j)$ according to Equations 20, 21, 22 and 23, $Q_{MTMF}^j(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j)$ will converge to a bounded distance of the Nash- Q equilibrium under Assumptions 1, 2 and 3:

$$Q_*(s, \mathbf{a}) - Q_t(s, \mathbf{a}) \leq D - S$$

where $S = \frac{\psi C_1 + \alpha \gamma |D|}{\alpha \beta_0}$. Here $D = \frac{1}{2} L \epsilon$, from Theorem 3.3.

PROOF. The proof of convergence of this theorem is structurally similar to that discussed by the authors in [7] and [17]. We provide the proof here, with changes necessitated by the multiple type case.

Note that in Assumption 3, Equation 24 corresponds to the global optimum and Equations 25 and 26 correspond to the saddle point. These assumptions are the same as those considered in [7]. Also note that the authors in [7] and [17] mention that Assumption 3 is a strong assumption to impose, which is needed to show the theoretical convergence, but this is not required to impose in practice.

The Nash operator \mathcal{H}^{Nash} is defined by

$$\mathcal{H}^{Nash} = \mathbb{E}_{s' \sim p}[r(s, \mathbf{a}) + \gamma v^{Nash}(s')] \quad (33)$$

where $Q \triangleq [Q^1, \dots, Q^N]$ and $\mathbf{r}(s, \mathbf{a}) \triangleq [r^1(s, \mathbf{a}), \dots, r^N(s, \mathbf{a})]$. The Nash value function is $v^{Nash}(s) \triangleq [v_{\pi_*}^1(s), \dots, v_{\pi_*}^N(s)]$. Here the Nash policy is represented as π_* . The Nash value function is calculated with the assumption that all agents are following π_* from the initial state s .

We need to apply Lemma 2 to prove Theorem 3.4. By subtracting $Q_*(s, \mathbf{a})$ on both sides of Equation 20 and in relation to Equation 28:

$$\begin{aligned} \Delta_t(x) &= Q_t(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) - Q_*(s, \mathbf{a}) \\ F_t(x) &= r_t + \gamma v_t^{MTMF}(s_{t+1}) - Q_*(s_t, \mathbf{a}_t) \end{aligned} \quad (34)$$

where $x \triangleq (s_t, \mathbf{a}_t)$ denotes the visited state-action pair at time t .

In Theorem 3.3, we proved a bound for the actual Q function and the multi type mean field Q function. We apply that in Equation 34, to get the following equation for Δ .

$$\begin{aligned} \Delta_t(x) &= Q_t(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) - Q_*(s, \mathbf{a}) \\ \Delta_t(x) &= Q_t(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) - Q_t(s, \mathbf{a}) - Q_t(s, \mathbf{a}) - Q_*(s, \mathbf{a}) \\ \Delta_t(x) &\leq |Q_t(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) - Q_t(s, \mathbf{a})| + Q_t(s, \mathbf{a}) - Q_*(s, \mathbf{a}) \\ \Delta_t(x) &\leq Q_t(s, \mathbf{a}) - Q_*(s, \mathbf{a}) + D \end{aligned} \quad (35)$$

where $D = \frac{1}{2} L \epsilon$.

The aim is to prove that the four conditions of Lemma 2 hold and that Δ in Equation 35 converges to a constant S according to Lemma 2 and thus the MTMF Q function in Equation 35 converges to a point whose distance to the Nash Equilibrium is bounded. In Equation 28, $\alpha(t)$ refers to the learning rate and hence the first condition of Lemma 2 is automatically satisfied. The second condition is also true as we are dealing with finite state and action spaces.

Let \mathcal{F}_t be the σ -field generated by all random variables in the history time $t - (s_t, \alpha_t, \mathbf{a}_t, r_{t-1}, \dots, s_1, \alpha_1, \mathbf{a}_1, \mathbf{Q}_0)$. Thus, Q_t is a random variable derived from the historical trajectory up to time t .

To prove the third condition of Lemma 2, from Equation 34:

$$\begin{aligned} F_t(s_t, \mathbf{a}_t) &= r_t + \gamma v_t^{MTMF} - Q_*(s_t, \mathbf{a}_t) \\ &= r_t + \gamma v_t^{Nash} - Q_*(s_t, \mathbf{a}_t) + \gamma [v_t^{MTMF}(s_{t+1}) - v_t^{Nash}(s_{t+1})] \\ &= (r_t + \gamma v_t^{Nash} - Q_*(s_t, \mathbf{a}_t)) + C_t(s_t, \mathbf{a}_t) = F_t^{Nash}(s_t, \mathbf{a}_t) + C_t(s_t, \mathbf{a}_t) \end{aligned} \quad (36)$$

From Lemma 1, F_t^{Nash} forms a contraction mapping with the norm $\|\cdot\|_\infty$ being the maximum norm on \mathbf{a} . Thus from Equation 35,

$$\|\mathbb{E}[F_t^{Nash}(s_t, \mathbf{a}_t)|\mathcal{F}_t]\|_\infty \leq \gamma \|Q_* - Q_t\|_\infty \leq \gamma \|D - \Delta_t\|_\infty \quad (37)$$

Now, applying Equation 37 in Equation 36.

$$\begin{aligned} \|\mathbb{E}[F_t(s_t, \mathbf{a}_t)|\mathcal{F}_t]\|_\infty &= \|F_t^{Nash}(s_t, \mathbf{a}_t)|\mathcal{F}_t\|_\infty + \|C_t(s_t, \mathbf{a}_t)|\mathcal{F}_t\|_\infty \\ &\leq \gamma \|D - \Delta_t\|_\infty + \|C_t(s_t, \mathbf{a}_t)|\mathcal{F}_t\|_\infty \\ &\leq \gamma \|\Delta_t\|_\infty + \|C_t(s_t, \mathbf{a}_t)|\mathcal{F}_t\|_\infty + \gamma \|D\|_\infty \leq \gamma \|\Delta_t\|_\infty + \gamma |D| \end{aligned} \quad (38)$$

Since we are taking the max norm, the last two terms in the right hand side of Equation 38 are both positive and finite. We can prove that the term $\|C_t(s_t, \mathbf{a}_t)\|$ converges to 0 with probability one. The proof involves the use of Assumption 3 (refer to Theorem 1 in [17]). We use this fact in the last term of Equation 38. Hence, the third condition of Lemma 2 is proved. The value of constant $K = \gamma |D| = \gamma \frac{1}{2} L \epsilon$.

For the fourth condition we use the fact that the MTMF operator \mathcal{H}^{MTMF} forms a contraction mapping. Hence, $\mathcal{H}^{MTMF} Q_* = Q_*$ and it follows that:

$$\begin{aligned} \text{var}[F_t(s_t, \mathbf{a}_t)|\mathcal{F}_t] &= E[(r_t + \gamma v_t^{MTMF}(s_{t+1}) - Q_*(s_t, \mathbf{a}_t))^2] \\ &= E[(r_t + \gamma v_t^{MTMF}(s_{t+1}) - \mathcal{H}^{MTMF}(Q_*)^2] \\ &= \text{var}[r_t + \gamma v_t^{MTMF}(s_{t+1})|\mathcal{F}_t] \leq K_2(1 + \|\Delta_t\|_w^2) \end{aligned} \quad (39)$$

In the last step, the left side of the equation contains the reward and the value function as the variables. The reward is bounded by Assumption 1 and the value function is also bounded by being updated recursively by Equation 21 (MTMF is a contraction operator). So we can choose a positive, finite K_2 such that the inequality holds.

Finally, with all conditions met, it follows from Lemma 2 that Δ_t converges to constant S with probability 1. The value of this constant is $S = \frac{\psi C_1 + \alpha \gamma |D|}{\alpha \beta_0}$ from Lemma 2 and using the value of K derived above. Therefore, from Equation 35 we get

$$Q_*(s, \mathbf{a}) - Q_t(s, \mathbf{a}) \leq D - S \leq \frac{1}{2} L \epsilon - S \quad (40)$$

Hence the Q function converges to a point within a bounded distance from the real Nash equilibrium of the game. The distance is a function of the error in the type classification and the closeness of resemblance of each agent to its type. \square

4 IMPLEMENTATION

We propose two algorithms based on Q-learning to estimate the multi-type mean field Q-function for known and unknown types. The first algorithm, MTMFQ (Multi-Type Mean Field Q-learning), trains an agent j to minimize the loss function $\mathcal{L}(\phi^j) = (y^j - Q_{\phi^j}(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j))^2$. Here $y^j = r^j + \gamma v_{\phi^j}^{MTMF}(s')$ is the target value used to calculate the TD error using the weights ϕ^j . Here the $v^{MTMF}(s) \triangleq [v^1(s), \dots, v^N(s)]$. Now, gradient can be taken as,

$$\nabla_{\phi^j} \mathcal{L}(\phi^j) = 2(Q_{\phi^j}(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) - y^j) \times \nabla_{\phi^j} Q_{\phi^j}(s, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j) \quad (41)$$

Algorithm 1 describes the Multi Type Mean Field Q learning (MTMFQ) algorithm when agent types are known. In this algorithm, different groups of agents in the environment are considered as types. An agent models its relation to each type separately and ultimately chooses the action that provides maximum benefit in the face of competition against the different types. This is referred to as version 1 of MTMFQ. This algorithm deals with multiple types in contrast to MFQ described in the paper by Yang et al. [17]. In Line 8, each agent is chosen and its neighbours are considered. The neighbours are classified into different types and in each type a new mean action is calculated (Line 9). In Lines 14 – 19, the Q networks are updated as done in common practice [13] for all the agents in the environment. At Line 12, the current actions are added to a buffer containing previous mean actions.

Version 2 of MTMFQ (see Algorithm 2) deals with the second type of scenario, where the type of each agent is unknown. An additional step doing K-means clustering is introduced to determine the types. Once we recognize the type of each agent, the algorithm is very similar to Algorithm 1. The clustering does not necessarily recognize the types correctly and the overall process is not as efficient as the known type case. But we will show that this way of approximate type determination is still better than using only a single mean field for all the agents. For the implementation we use neural networks but it can be done without neural networks too. We only need a way to recursively update Equations 21, 22 and 23.

Note that the total number of types in the environment is unknown and the agent does not need to guess the correct number of types. Generally, the more types are used, the closer the approximate Q-function may be to the exact Nash Q-function as shown by the bounds in Theorems 3.1, 3.2, 3.3 and 3.4. There is no risk of overfitting when using more types. In the limit, when there is one type per agent, we recover the exact multi-agent Q-function. The only drawback is an increase in computational complexity. The code repository will appear at <https://github.com/BorealisAI/mtmfql>.

5 EXPERIMENTS AND RESULTS

We report results with three games designed within the MAgents framework: the Multi Team Battle, Battle-Gathering and the Predator Prey domains. In the first two games, Multi Team Battle and the Battle-Gathering game, the conditions are such that the different groups are fully known upfront. In the third game, the conditions are such that the types of the agents are initially unknown. Hence, the agents must also learn the identity of the opponent agents during game play. Multi Team Battle and Gathering are two separately existing MAgent games, which have been combined to obtain the

Algorithm 1 Multi Type Mean Field Q Learning for known types

```

1: Initialize the number of types  $M$  and total number of agents  $N$ .
2: Initialize the  $Q$  functions (parameterised by weights)  $Q_{\phi^j}, Q_{\phi^j_-}$ , for all agents  $j$ .
3: Initialize the mean action for each type  $\bar{a}_1^j, \bar{a}_2^j, \dots, \bar{a}_M^j$  for each agent  $j \in 1, \dots, N$ 
4: Initialize the total number of steps (T) and total number of episodes (E)
5: while Episode < E do
6:   while Step < T do
7:     while m = 1 ... M do
8:       For each agent  $j$  take action  $a^j$  from  $Q_{\phi^j}$  according to Eq. 23 with the
       current mean action for each type  $\bar{a}_1^j, \dots, \bar{a}_M^j$  and the exploration rate
        $\beta$ .
9:       For each agent  $j$ , compute the new mean action for each type
        $\bar{a}_1^j, \dots, \bar{a}_M^j$  according to Eq. 22.
10:    end while
11:    Execute the joint action  $\mathbf{a} = [a^1, \dots, a^N]$ . Observe the rewards  $\mathbf{r} = [r^1, \dots, r^N]$  and the next state  $s'$ .
12:    Store  $(s, \mathbf{a}, \mathbf{r}, s', \bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_M)$  in replay buffer  $D$ , where  $\bar{\mathbf{a}}_i$  is the mean
    action for type  $i$  in the neighbourhood. The  $\mathbf{a}$  captures all the  $N$  agents.
13:    end while
14:    while  $j = 1$  to  $N$  do
15:      Sample a minibatch of  $K$  experiences  $(s, \mathbf{a}, \mathbf{r}, s', \bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_M)$  from  $D$ .
16:      Sample action  $a^j_-$  from  $Q_{\phi^j_-}$  with  $\bar{a}_i^- \leftarrow \bar{a}_i^j$  for each type  $i$ .
17:      Set  $y^j = r^j + \gamma v_{\phi^j_-}^{MTMF}(s')$  according to Eq. 21.
18:      Update the Q network by minimizing the loss  $L(\phi^j) = \frac{1}{k} \sum (y^j - Q_{\phi^j}(s^j, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j))^2$ 
19:    end while
20:    Update the parameters of the target network for each agent  $j$  with learning
    rate  $\tau$ ;  $\phi^j_- \leftarrow \tau \phi^j + (1 - \tau) \phi^j_-$ 
21: end while

```

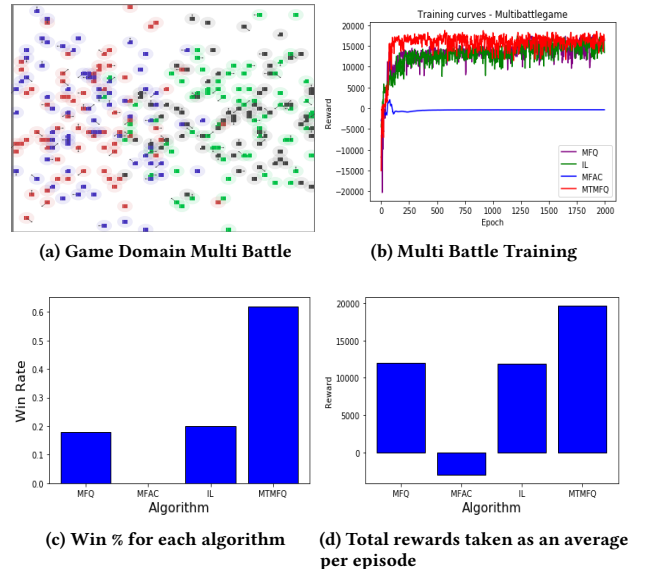


Figure 1: Multi Battle Game Results.

Battle-Gathering game used in this paper. Predator Prey domain is also obtained from combining Multi Team Battle with another existing MAgent game (Pursuit). In preparation for each game, agents train for 2000 episodes of game play against different groups training using the same algorithm, which is referred to as the first stage. Next, in the second stage, they enter into a faceoff against other agents trained by other algorithms where they fight each other for

Algorithm 2 Multi Type Mean Field Q Learning for unknown types

```

1: Initialize the number of types  $M$  and total number of agents  $N$ .
2: Initialize the  $Q$  functions (parameterised by weights)  $Q_{\phi^j}$ ,  $Q_{\phi_-^j}$ , for all agents  $j$ .
3: Initialize the mean action for each type  $\bar{a}_1^j, \bar{a}_2^j \dots \bar{a}_M^j$  for each agent  $j \in 1, \dots, N$ 
4: Initialize the total number of steps (T) and total number of episodes (E)
5: Initialize every agent to a type at random. Initialize an array  $A$  containing the previous action of all agents.
6: Maintain a buffer  $B$  for storing the last  $C$  actions of all agents.  $C$  is determined by the conditions of the environment. Initialize all values to 0.
7: while Episode < E do
8:   while steps < T do
9:     while  $m = 1 \dots M$  do
10:      For each agent  $j$  take action  $a^j$  from  $Q_{\phi^j}$  according to Eq. 23 with the current mean action for each type  $\bar{a}_1^j, \dots, \bar{a}_M^j$  and the exploration rate  $\beta$ .
11:      For each agent  $j$ , compute the new mean action for each type  $\bar{a}_1^j, \dots, \bar{a}_M^j$  according to Eq. 22.
12:    end while
13:    Execute the joint action  $\mathbf{a} = [a^1, \dots, a^N]$ . Observe the rewards  $\mathbf{r} = [r^1, \dots, r^N]$  and the next state  $s'$ .
14:    Store  $\langle s, \mathbf{a}, \mathbf{r}, s', \bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_M \rangle$  in replay buffer  $D$ , where  $\bar{\mathbf{a}}_i$  is the mean action for type  $i$  in the neighbourhood. The  $\bar{\mathbf{a}}_i$  captures all the  $N$  agents.
15:    Store each action  $[a^1, \dots, a^N]$  in the Array  $A$ . Update the Buffer  $B$  with the last action taken by all agents.
16:    Perform a K means clustering on  $B$  with the number of clusters equal to the number of types  $M$ .
17:    Reassign the agents to different types based on the cluster in K means.
18:  end while
19:  while  $j = 1$  to  $N$  do
20:    Sample a minibatch of K experiences  $\langle s, \mathbf{a}, \mathbf{r}, s', \bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_M \rangle$  from  $D$ .
21:    Sample action  $a_-^j$  from  $Q_{\phi_-^j}$  with  $\bar{a}_i^j \leftarrow \bar{a}_i^j$  for each type  $i$ .
22:    Set  $y^j = r^j + \gamma v_{\phi_-^j}^{MTMF}(s')$  according to Eq. 21.
23:    Update the Q network by minimizing the loss  $L(\phi^j) = \frac{1}{k} \sum (y^j - Q_{\phi^j}(s^j, a^j, \bar{a}_1^j, \dots, \bar{a}_M^j))^2$ 
24:  end while
25:  Update the parameters of the target network for each agent  $j$  with learning rate  $\tau$ :  $\phi_-^j \leftarrow \tau \phi^j + (1 - \tau) \phi_-^j$ 
26: end while

```

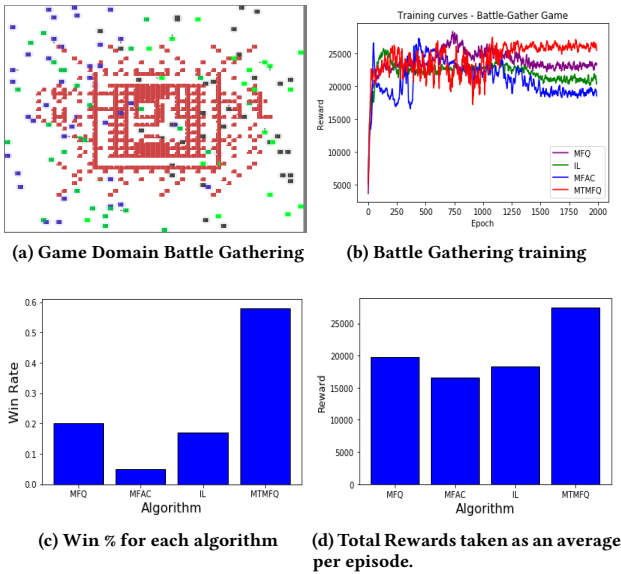


Figure 2: Battle-Gathering Game Results.

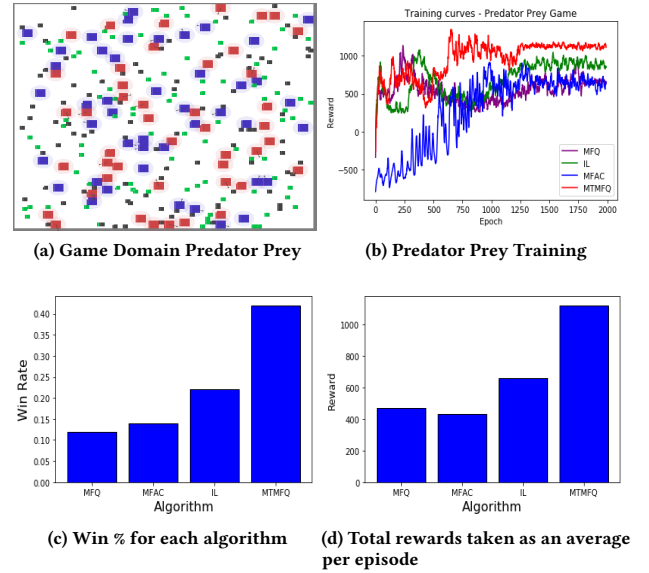


Figure 3: Predator Prey Game Results.

1000 games. We report the results for both stages. We repeat all experiments 50 times and report the averages. As can be seen from the nature of the experiments, variances can be quite large across individual experimental runs.

In the first game (Multi Team Battle, see Figure 1(a)), there are four teams (different colors in Figure 1(a)) fighting against each other to win a battle. Here, agents belonging to one team are competing against agents from other teams and cooperating with agents of the same team. During the first stage, a team does not know how the other teams will play during the faceoff – so it clones itself to three additional teams with slightly different rewards in order to induce different strategies. This is similar to self play. Each agent receives a reward equal to the sum of the local rewards attributed to each agent in the team. The reward function is defined in such a way that it encourages local cooperation in killing opponents and discourages getting attacked and dying in the game. The reward function for different agent groups are also subtly different (refer to Appendix B for the details). Let us call each of these Groups: Group A, Group B, Group C, and Group D. We maintain a notion of favourable opponents and each Group gets slightly higher rewards for killing the favourable opponents than the others. Four algorithms, namely MFQ [17], MFAC [17], Independent Q Learning (IL) [16], and MTMFQ are trained separately where each of these groups trains its own separate network using the same algorithm. We start all the battles with 72 agents of each group for training and testing. Group A from MTMFQ, Group B from IL, Group C from MFQ, and Group D from MFAC enter the faceoff stage where they fight each other for 1000 games. Each game (or episode) has a maximum of 500 steps and a game is considered to be won by the group/groups that have the most number of agents alive at the end of the game.

The results of the first training stage are reported in Figure 1(a). We report the cumulative rewards from all agents in Group A for

each algorithm. These rewards are for Group A against the other 3 groups. Since the different groups are just clones of each other the reward curves for other groups are similar to that of Group A. In the training stage, the teams trained by different algorithms did not play against each other, but simply against the cloned teams trained by the same algorithm. At the beginning of training, for about 100 episodes the agents are still exploring and their strategies are not well differentiated yet. As a result, MTMFQ’s performance is still comparable to the performance of other algorithms. At this stage, the assumption of a single type is fine. As training progresses, each group begins to identify the favorable opponents and tries to make a targeted approach in the battle. When such differences exist across a wide range of agents, the MTMFQ algorithm shows a better performance than the other techniques as it explicitly considers the presence of different types in the game. Overall, we observe that MTMFQ has a faster convergence than all other algorithms and it also produces higher rewards at the end of the complete training. This shows that MTMFQ identifies favorable opponents early, but the other algorithms struggle longer to learn this condition. The MFAC algorithm is the worst overall. This is consistent with the observation by Yang et al. [17], where the authors give particular reasons for this bad performance, including the Greedy in the limit with infinite exploration (GLIE) assumption and a positive bias of Q learning algorithms. This algorithm is not able to earn an overall positive reward upon the complete training. Figure 1(c) shows the win rate of the teams trained by each algorithm (MTMFQ, MFQ, MFAC, and IL) in a direct faceoff of 1000 episodes. In the face off, each group is trained by a different algorithm and with a different reward function that induces different strategies. The only algorithm that handles different strategies among the opponent teams is MTMFQ and therefore it dominates the other algorithms with a win rate of 60%. Figure 1(d) reinforces this domination.

The second domain is the Battle-Gathering game (Figure 2(a)). In this game, all the agent groups compete for food resources that are limited (red denotes food particles and other colours are competing agents) in addition to killing its opponents as in the Multi Team Battle game. Hence, this game is harder than the first one. All the training and competition are similar to the first game.

Figure 2(b) reports the results of training in the Battle-Gathering game. Like the Multi Battle Game, we plot the rewards obtained by Group A while fighting other groups for each algorithm. Again, MTMFQ shows the strongest performance in comparison to the other three algorithms. The MFQ technique performs better than both MFAC and IL. In this game, MTMFQ converges in around 1500 episodes, while the other algorithms take around 1800 episodes to converge. The win rates shown in Figure 2(c) and the total rewards reported in Figure 2(d) also show the dominance of MTMFQ.

The third domain is the multiagent Predator Prey (Figure 3(a)). Here we have two distinct types of agents – predator and prey, each having completely different characteristics. The prey are faster than the predators and their ultimate goal is to escape the predators. The predators are slower than the prey, but they have higher attacking abilities than the prey. So the predators try to attack more and kill more prey. We model this game as an unknown type scenario where the types of the other agents in the competition are not known before hand (refer to Appendix B for more details). The MTMFQ algorithm plays the version with unknown types (Algorithm 2).

Here we have four groups with the first two groups (Groups A and B) being predators and the other two groups (Groups C and D) being prey. Each algorithm will train two kinds of predator agents and two kinds of prey agents. All these agents are used in the playoff stage. In the playoff stage we have 800 games where we change the algorithm of predator and prey at every 200 games to maintain a fair competition. For the first 200 games, MTMFQ plays Group A, MFQ plays Group B, IL plays Group C, and MFAC plays Group D. In the next 200 games, MFAC plays Group A, MTMFQ plays Group B, MFQ plays Group C and IL plays Group D, and so on. We start all training and testing episodes with 90 prey and 45 predators for each group. Winning a game in the playoff stage is defined in the same way as the previous two games. Notice that this makes it more fair, as predators have to kill a lot more prey to win the game (as we start with more prey than predators) and prey have to survive longer. In this setup, the highly different types of agents make type identification easier for MTMFQ (as the types are initially unknown). The prey execute more move actions while the predators execute more attack actions. This can be well differentiated by clustering.

The results of the first training stage are reported in Figure 2(b). MTMFQ has comparable or even weaker performance than other algorithms in the first 600 episodes of the training and the reasoning is similar to the reasoning in the Multi Battle game (the agent strategies are not sufficiently differentiated for multiple types to be useful). Notice that for this game, MTMFQ takes many more episodes than the earlier two games to start dominating. This is because of the inherent hardness of this domain compared to the other domains (very different and unknown types). Similar to observations in the other domains, MTMFQ converges earlier (after around 1300 episodes as opposed to 1700 for the other algorithms). This shows its robustness to the different kinds of opponents in the environment. MTMFQ also gains higher cumulative rewards than the other algorithms. Win rates in Figure 3(c) show that MTMFQ still wins more games than the other algorithms, but the overall number of games won is less than the other domains. Thus, irrespective of the difficulty of the challenge we can see that MTMFQ has an upper hand. The lead of MTMFQ is also observed in Figure 3(d).

6 CONCLUSION

In this paper, we extended the notion of mean field theory to multiple types in MARL. We demonstrate that reducing many agent interactions to simple two agent interactions does not give very accurate solutions in environments where there are clearly different teams/types playing different strategies. We perform suitable experiments using MAgents and demonstrate superior performances using a type based approach. We hope that this paper will provide a different dimension to the mean field theory based MARL research.

One limitation of our approach is that it is computationally more expensive than the Mean Field Method without types. If we really have only one type in the environment, then our method would add more compute time and not necessarily produce a better result. As future work we would like to extend this work for completely heterogeneous agents with different action spaces as well. StarCraft is one example of such a domain. Our work would be well suited for this scenario as clustering would be even easier. Another approach would be to consider sub types, further dividing types.

REFERENCES

- [1] K. Arulkumar, M. P. Deisenroth, M. Brundage, and A. A. Bharath. 2017. Deep Reinforcement Learning: A Brief Survey. *IEEE Signal Processing Magazine* 34, 6 (Nov 2017), 26–38. <https://doi.org/10.1109/MSP.2017.2743240>
- [2] Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. 2015. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research* 53 (2015), 659–697.
- [3] Michael Bowling and Manuela Veloso. 2000. *An analysis of stochastic game theory for multiagent reinforcement learning*. Technical Report. Carnegie-Mellon Univ Pittsburgh Pa School of Computer Science.
- [4] Lucian Bu, Robert Babu, Bart De Schutter, et al. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 2 (2008), 156–172.
- [5] Lucian Busoniu, Robert Babuska, and Bart De Schutter. 2010. *Multi-agent Reinforcement Learning: An Overview*. Vol. 310. Delft University of Technology, 183–221. https://doi.org/10.1007/978-3-642-14435-6_7
- [6] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. 2019. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 33, 6 (01 Nov 2019), 750–797. <https://doi.org/10.1007/s10458-019-09421-1>
- [7] Junling Hu and Michael P Wellman. 2003. Nash Q-learning for general-sum stochastic games. *Journal of machine learning research* 4, Nov (2003), 1039–1069.
- [8] Tommi Jaakkola, Michael I. Jordan, and Satinder P. Singh. 1994. On the Convergence of Stochastic Iterative Dynamic Programming Algorithms. *Neural Computation* 6, 6 (1994), 1185–1201. <https://doi.org/10.1162/neco.1994.6.6.1185>
- [9] James S Jordan. 1991. Bayesian learning in normal form games. *Games and Economic Behavior* 3, 1 (1991), 60–81.
- [10] Jean-Michel Lasry and Pierre-Louis Lions. 2007. Mean field games. *Japanese journal of mathematics* 2, 1 (2007), 229–260.
- [11] Eric Maskin. 1999. Nash equilibrium and welfare optimality. *The Review of Economic Studies* 66, 1 (1999), 23–38.
- [12] David Mguni, Joel Jennings, and Enrique Munoz de Cote. 2018. Decentralised learning in systems with many, many strategic agents. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [14] Thanh Thi Nguyen, Ngoc Duy Nguyen, and Saeid Nahavandi. 2018. Deep reinforcement learning for multi-agent systems: a review of challenges, solutions and applications. *arXiv preprint arXiv:1812.11794* (2018).
- [15] Richard S Sutton and Andrew G Barto. 1998. *Introduction to reinforcement learning*. Vol. 135. MIT press Cambridge.
- [16] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*. 330–337.
- [17] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean Field Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, Stockholm, Stockholm Sweden, 5571–5580. <http://proceedings.mlr.press/v80/yang18d.html>
- [18] Lianmin Zheng, Jiacheng Yang, Han Cai, Ming Zhou, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Magent: A many-agent reinforcement learning platform for artificial collective intelligence. In *Thirty-Second AAAI Conference on Artificial Intelligence*.