

A Regularized Opponent Model with Maximum Entropy Objective

Zheng Tian^{1*}, Ying Wen^{1*}, Zhichen Gong¹, Faiz Punakkath¹, Shihao Zou² and Jun Wang¹

¹University College London

²University of Alberta

{zheng.tian, ying.wen, jun.wang}@cs.ucl.ac.uk

Abstract

In a single-agent setting, reinforcement learning (RL) tasks can be cast into an inference problem by introducing a binary random variable o , which stands for the “optimality”. In this paper, we redefine the binary random variable o in multi-agent setting and formalize multi-agent reinforcement learning (MARL) as probabilistic inference. We derive a variational lower bound of the likelihood of achieving the optimality and name it as Regularized Opponent Model with Maximum Entropy Objective (ROMMEO). From ROMMEO, we present a novel perspective on opponent modeling and show how it can improve the performance of training agents theoretically and empirically in cooperative games. To optimize ROMMEO, we first introduce a tabular Q-iteration method ROMMEO-Q with proof of convergence. We extend the exact algorithm to complex environments by proposing an approximate version, ROMMEO-AC. We evaluate these two algorithms on the challenging iterated matrix game and differential game respectively and show that they can outperform strong MARL baselines.

1 Introduction

Casting decision making and optimal control as an inference problem have a long history, which dates back to [Kalman, 1960] where the Kalman smoothing is used to solve optimal control in linear dynamics with quadratic cost. Bayesian methods can capture the uncertainties regarding the transition probabilities, the rewards functions in the environment or other agents’ policies. This distributional information can be used to formulate a more structured exploration/exploitation strategy than those commonly used in classical RL, e.g. ϵ -greedy. A common approach in many works [Toussaint and Storkey, 2006; Rawlik *et al.*, 2013; Levine and Koltun, 2013; Abdolmaleki *et al.*, 2018] for framing RL as an inference problem is by introducing a binary random variable o which represents “optimality”. By this way, RL problems are able to lend itself to powerful inference tools [Levine, 2018]. However, the Bayesian approach in a multi-agent environment is

less well studied.

In many single-agent works, maximizing entropy is part of a training agent’s objective for resolving ambiguities in inverse reinforcement learning [Ziebart *et al.*, 2008], improving the diversity [Florensa *et al.*, 2017], robustness [Fox *et al.*, 2015] and the compositionality [Haarnoja *et al.*, 2018a] of the learned policy. In Bayesian RL, it often presents in the evidence lower bound (ELBO) for the log likelihood of optimality [Haarnoja *et al.*, 2017; Schulman *et al.*, 2017; Haarnoja *et al.*, 2018b], commonly known as maximum entropy objective (MEO), which encourages the optimal policy to maximize the expected return and long term entropy.

In MARL, there is more than one agent interacting with a stationary environment. In contrast with the single agent environment, an agent’s reward not only depends on the current environment state and the agent’s action but also on the actions of others. The existence of other agents increases the uncertainty in the environment. Therefore, the capability of reasoning about other agents’ belief, private information, behavior, strategy, and other characteristics is crucial. A reasoning model can be used in many different ways, but the most common case is where an agent utilize its reasoning model to help its self decision making [Brown, 1951; Heinrich and Silver, 2016; He *et al.*, 2016; Raileanu *et al.*, 2018; Wen *et al.*, 2019; Tian *et al.*, 2018]. In this work, we use the word “opponent” when referring to another agent in the environment irrespective of the environment’s cooperative or adversarial nature.

In our work, we reformulate the MARL problem into Bayesian inference and derive a multi-agent version of MEO, which we call the regularized opponent model with maximum entropy objective (ROMMEO). Optimizing this objective with respect to one agent’s opponent model gives rise to a new perceptive on opponent modeling. We present two off-policy RL algorithms for optimizing ROMMEO in MARL. ROMMEO-Q is applied in discrete action case with proof of convergence. For the complex and continuous action environment, we propose ROMMEO Actor-Critic (ROMMEO-AC), which approximates the former procedure and extend itself to continuous problems. We evaluate these two approaches on the matrix game and the differential game against strong baselines and show that our methods can outperform all the baselines in terms of the overall performance and speed of convergence.

*The first two authors contributed equally.

2 Method

2.1 Stochastic Games

For an n -agent stochastic game [Shapley, 1953], we define a tuple $(\mathcal{S}, \mathcal{A}^1, \dots, \mathcal{A}^n, R^1, \dots, R^n, p, \mathcal{T}, \gamma)$, where \mathcal{S} denotes the state space, p is the distribution of the initial state, γ is a discount factor, \mathcal{A}^i and $R^i = R^i(s, a^i, a^{-i})$ are the action space and the reward function for agent $i \in \{1, \dots, n\}$ respectively. States are transitioned according to $\mathcal{T} : \mathcal{S} \times \mathcal{A}$, where $\mathcal{A} = \{\mathcal{A}^1, \dots, \mathcal{A}^n\}$. Agent i chooses its action $a^i \in \mathcal{A}^i$ according to the policy $\pi_{\theta^i}^i(a^i|s)$ parameterized by θ^i conditioning on some given state $s \in \mathcal{S}$. Let us define the joint policy as the collection of all agents' policies π_{θ} with θ representing the joint parameter. It is convenient to interpret the joint policy from the perspective of agent i such that $\pi_{\theta} = (\pi_{\theta^i}^i(a^i|s), \pi_{\theta^{-i}}^{-i}(a^{-i}|s))$, where $a^{-i} = (a^j)_{j \neq i}$, $\theta^{-i} = (\theta^j)_{j \neq i}$, and $\pi_{\theta^{-i}}^{-i}(a^{-i}|s)$ is a compact representation of the joint policy of all complementary agents of i . At each stage of the game, actions are taken simultaneously.

In fully cooperative games, different agents have the same reward function $R^i(s, a^i, a^{-i}) = R^{-i}(s, a^i, a^{-i}), \forall i \in 1, \dots, n$. Therefore, each agent's objective is to maximize the shared expected return:

$$\max \eta^i(\pi_{\theta}) = \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^t R(s_t, a_t^i, a_t^{-i}) \right], \quad (1)$$

with (a_t^i, a_t^{-i}) sample from $(\pi_{\theta^i}^i, \pi_{\theta^{-i}}^{-i})$.

2.2 A Variational Lower Bound for Multi-Agent Reinforcement Learning Problems

We transform the control problem into an inference problem by introducing a binary random variable o_t^i which serves as the indicator for "optimality" for each agent i at each time step t . Recall that in single agent problem, reward $R(s_t, a_t)$ is bounded, but the achievement of the maximum reward given the action a_t is unknown. Therefore, in the single-agent case, o_t indicates the optimality of achieving the bounded maximum reward r_t^* . It thus can be regarded as a random variable and we have $P(o_t = 1|s_t, a_t) \propto \exp(R(s_t, a_t))$. Intuitively, this formulation dictates that higher rewards reflect a higher likelihood of achieving optimality, i.e., the case when $o_t = 1$. However, the definition of "optimality" in the multi-agent case is subtly different from the one in the single-agent situation.

In cooperative multi-agent reinforcement learning (CMARL), to define agent i 's optimality, we first introduce the definition of optimum and optimal policy:

Definition 1. In cooperative multi-agent reinforcement learning, optimum is a strategy profile $(\pi^{1*}, \dots, \pi^{n*})$ such that:

$$\begin{aligned} & \mathbb{E}_{s \sim p_s, a_t^{i*} \sim \pi^{i*}, a_t^{-i*} \sim \pi^{-i*}} \left[\sum_{t=1}^{\infty} \gamma^t R(s_t, a_t^{i*}, a_t^{-i*}) \right] \\ & \geq \mathbb{E}_{s \sim p_s, a_t^i \sim \pi^i, a_t^{-i} \sim \pi^{-i}} \left[\sum_{t=1}^{\infty} \gamma^t R(s_t, a_t^i, a_t^{-i}) \right] \forall \pi, \quad (2) \end{aligned}$$

where $\pi = (\pi^i, \pi^{-i})$ and Agent i 's optimal policy is π^{i*} .

In CMARL, a single agent's "optimality" o_t^i cannot imply that it obtains the maximum reward because the reward depends on the joint actions of all agents (a^i, a^{-i}) . Therefore, we define $o_t^i = 1$ only indicates that agent i 's policy at time step t is optimal. The posterior probability of agent i 's optimality given its action a_t^i is the probability that the action is sampled from the optimal policy:

$$P(o_t^i|a_t^i) = P(a_t^i \sim \pi^{i*}|a_t^i) = \pi^{i*}(a_t^i). \quad (3)$$

We also assume that given other players actions a_t^{-i} , the posterior probability of agent i 's optimality is proportional to its exponential reward:

$$P(o_t^i = 1|s_t, a_t^i, a_t^{-i}) \propto \exp(R(s_t, a_t^i, a_t^{-i})). \quad (4)$$

The conditional probabilities of "optimality" in both of CMARL and single-agent case have similar forms. However, it is worth mentioning that the "optimality" in CMARL has a different interpretation to the one in single-agent case.

For cooperative games, if all agents play optimally, then agents can receive the maximum rewards, which is the optimum of the games. Therefore, given the fact that other agents are playing their optimal policies $o^{-i} = 1$, the probability that agent i also plays its optimal policy $P(o^i = 1|o^{-i} = 1)$ is the probability of obtaining the maximum reward from agent i 's perspective. Therefore, we define agent i 's objective as:

$$\max \mathcal{J} \triangleq \log P(o_{1:T}^i = 1|o_{1:T}^{-i} = 1) \quad (5)$$

As we assume no knowledge of the optimal policies and the model of the environment, we treat them as latent variables. To optimize the observed evidence defined in Eq. 5, therefore, we use variational inference (VI) with an auxiliary distribution over these latent variables $q(a_{1:T}^i, a_{1:T}^{-i}, s_{1:T}|o_{1:T}^i = 1, o_{1:T}^{-i} = 1)$. Without loss of generality, we here derive the solution for agent i . We factorize $q(a_{1:T}^i, a_{1:T}^{-i}, s_{1:T}|o_{1:T}^i = 1, o_{1:T}^{-i} = 1)$ so as to capture agent i 's conditional policy on the current state and opponents actions, and beliefs regarding opponents actions. This way, agent i will learn optimal policy, while also possessing the capability to model opponents actions a^{-i} . Using all modelling assumptions, we may factorize $q(a_{1:T}^i, a_{1:T}^{-i}, s_{1:T}|o_{1:T}^i = 1, o_{1:T}^{-i} = 1)$ as:

$$\begin{aligned} & q(a_{1:T}^i, a_{1:T}^{-i}, s_{1:T}|o_{1:T}^i = 1, o_{1:T}^{-i} = 1) \\ & = P(s_1) \prod_t P(s_{t+1}|s_t, a_t) q(a_t^i|a_t^{-i}, s_t, o_t^i = o_t^{-i} = 1) \\ & \quad \times q(a_t^{-i}|s_t, o_t^i = o_t^{-i} = 1) \\ & = P(s_1) \prod_t P(s_{t+1}|s_t, a_t) \pi(a_t^i|s_t, a_t^{-i}) \rho(a_t^{-i}|s_t), \end{aligned}$$

where we have assumed the same initial and states transitions as in the original model. With this factorization, we derive a

lower bound on the likelihood of optimality of agent i :

$$\begin{aligned}
 & \log P(o_{1:T}^i = 1 | o_{1:T}^{-i} = 1) \\
 & \geq \mathcal{J}(\pi, \rho) \triangleq \sum_t \mathbb{E}_{(s_t, a_t^i, a_t^{-i}) \sim q} [R(s_t, a_t^i, a_t^{-i}) \\
 & + H(\pi(a_t^i | s_t, a_t^{-i})) - D_{\text{KL}}(\rho(a_t^{-i} | s_t) || P(a_t^{-i} | s_t))] \quad (6) \\
 & = \sum_t \mathbb{E}_{s_t} \underbrace{\left[\mathbb{E}_{a_t^i \sim \pi, a_t^{-i} \sim \rho} [R(s_t, a_t^i, a_t^{-i}) + H(\pi(a_t^i | s_t, a_t^{-i}))] \right]}_{\text{MEO}} \\
 & - \underbrace{\mathbb{E}_{a_t^{-i} \sim \rho} [D_{\text{KL}}(\rho(a_t^{-i} | s_t) || P(a_t^{-i} | s_t))] }_{\text{Regularizer of } \rho}. \quad (7)
 \end{aligned}$$

Written out in full, $\rho(a_t^{-i} | s_t, o_t^{-i} = 1)$ is agent i 's opponent model estimating *optimal* policies of its opponents, $\pi(a_t^i | s_t, a_t^{-i}, o_t^i = 1, o_t^{-i} = 1)$ is the agent i 's conditional policy at optimum ($o_t^i = o_t^{-i} = 1$) and $P(a_t^{-i} | s_t, o_t^{-i} = 1)$ is the prior of optimal policy of opponents. In our work, we set the prior $P(a_t^{-i} | s_t, o_t^{-i} = 1)$ equal to the observed empirical distribution of opponents' actions given states. As we are only interested in the case where ($o_t^i = 1, o_t^{-i} = 1$), we drop them in π, ρ and $P(a_t^{-i} | s_t)$ here and thereafter. $H(\cdot)$ is the entropy function. Eq. 6 is a variational lower bound of $\log P(o_{1:T}^i = 1 | o_{1:T}^{-i} = 1)$ and the derivation is deferred to Appendix B.1.

2.3 The Learning of Opponent Model

We can further expand Eq. 6 into Eq. 7 and we find that it resembles the maximum entropy objective in single-agent reinforcement learning [Kappen, 2005; Todorov, 2007; Ziebart *et al.*, 2008; Haarnoja *et al.*, 2017]. We denote agent i 's expectation of reward $R(s_t, a_t^i, a_t^{-i})$ plus entropy of the conditional policy $H(\pi(a_t^i | s_t, a_t^{-i}))$ as agent i 's *maximum entropy objective (MEO)*. In the multi-agent version, however, it is worthy of noting that optimizing the MEO will lead to the *optimization of ρ* . This can be counter-intuitive at first sight as opponent behaviour models are normally trained with only past state-action data (s, a^{-i}) to predict opponents' actions.

However, recall that $\rho(a_t^{-i} | s_t, o_t^{-i} = 1)$ is modelling opponents' *optimal* policies in our work. Given agent i 's policy π^i being fixed, optimizing MEO with respect to ρ updates agent i 's opponent model in the direction of the higher shared reward $R(s, a^i, a^{-i})$ and the more stochastic conditional policy $\pi^i(a^i | s, a^{-i})$, making it closer to the real optimal policies of the opponents. Without any regularization, at iteration d , agent i can freely learn a new opponent model ρ_{d+1}^i which is the closest to the optimal opponent policies π^{-i*} from its perspective given $\pi_d^i(a^i | s, a^{-i})$. Next, agent i can optimize the lower bound with respect to $\pi_{d+1}^i(a^i | s, a^{-i})$ given ρ_{d+1}^i . Then we have an EM-like iterative training and can show it monotonically increases the probability that the opponent model ρ is optimal policies of the opponents. Then, by acting optimally to the converged opponent model $\rho^{i\infty}$, we can recover agent i 's optimal policy π^{i*} .

However, it is unrealistic to learn such an opponent model. As the opponents have no access to agent i 's conditional policy $\pi_d^i(a^i | s, a^{-i})$, the learning of its policy can be different

from the one of agent i 's opponent model. Then the actual opponent policies π_{d+1}^{-i} can be very different from agent i 's converged opponent model $\rho^{i\infty}$ learned in the above way given agent i 's conditional policy $\pi_d^i(a^i | s, a^{-i})$. Therefore, acting optimally to an opponent model far from the real opponents' policies can lead to poor performance.

The last term in Eq. 7 can prevent agent i building an unrealistic opponent model. The Kullback-Leibler (KL) divergence between opponent model and a prior $D_{\text{KL}}(\rho(a_t^{-i} | s_t) || P(a_t^{-i} | s_t))$ can act as a regularizer of ρ . By setting the prior to the empirical distribution of opponent past behaviour, the KL divergence penalizes ρ heavily if it deviates from the empirical distribution too much. As the objective in Eq. 7 can be seen as a *Maximum Entropy objective* for one agent's policy and opponent model with *regularization on the opponent model*, we call this objective as *Regularized Opponent Model with Maximum Entropy Objective (ROMMEO)*.

3 Multi-Agent Soft Actor Critic

To optimize the ROMMEO in Eq. 7 derived in the previous section, we propose two off-policy algorithms. We first introduce an exact tabular Q-iteration method with proof of convergence. For practical implementation in a complex continuous environment, we then propose the ROMMEO actor critic ROMMEO-AC, which is an approximation to this procedure.

3.1 Regularized Opponent Model with Maximum Entropy Objective Q-Iteration

In this section, we derive a multi-agent version of Soft Q-iteration algorithm proposed in [Haarnoja *et al.*, 2017] and we name our algorithm as ROMMEO-Q. The derivation follows from a similar logic to [Haarnoja *et al.*, 2017], but the extension of Soft Q-learning to MARL is still nontrivial. From this section, we slightly modify the objective in Eq. 7 by adding a weighting factor α for the entropy term and the original objective can be recovered by setting $\alpha = 1$.

We first define multi-agent soft Q-function and V-function respectively. Then we can show that the conditional policy and opponent model defined in Eq. 10 and 11 below are optimal solutions with respect to the objective defined in Eq. 7:

Theorem 1. *We define the soft state-action value function of agent i as*

$$\begin{aligned}
 Q_{soft}^{\pi^*, \rho^*}(s_t, a_t^i, a_t^{-i}) &= r_t + \mathbb{E}_{(s_{t+1}, a_{t+1}^i, a_{t+1}^{-i}, \dots) \sim q} \left[\sum_{l=1}^{\infty} \gamma^l (r_{t+l} \right. \\
 & \left. + \alpha H(\pi^*(a_{t+l}^i | a_{t+l}^{-i}, s_{t+l})) - D_{\text{KL}}(\rho^*(a_{t+l}^{-i} | s_{t+l}) || P(a_{t+l}^{-i} | s_{t+l})) \right], \quad (8)
 \end{aligned}$$

and soft state value function as

$$V^*(s) = \log \sum_{a^{-i}} P(a^{-i} | s) \left(\sum_{a^i} \exp(\frac{1}{\alpha} Q_{soft}^{\pi^*, \rho^*}(s, a^i, a^{-i})) \right)^\alpha, \quad (9)$$

Then the optimal conditional policy and opponent model for Eq. 6 are

$$\pi^*(a^i | s, a^{-i}) = \frac{\exp(\frac{1}{\alpha} Q_{soft}^{\pi^*, \rho^*}(s, a^i, a^{-i}))}{\sum_{a^i} \exp(\frac{1}{\alpha} Q_{soft}^{\pi^*, \rho^*}(s, a^i, a^{-i}))}, \quad (10)$$

and

$$\rho^*(a^{-i}|s) = \frac{P(a^{-i}|s) \left(\sum_{a^i} \exp(\frac{1}{\alpha} Q_{soft}^*(s, a^i, a^{-i})) \right)^\alpha}{\exp(V^*(s))}. \quad (11)$$

Proof. See Appendix C.2. \square

Following from Theorem 1, we can find the optimal solution of Eq. 7 by learning the soft multi-agent Q-function first and recover the optimal policy π^* and opponent model ρ^* by Equations 10 and 11. To learn the Q-function, we show that it satisfies a Bellman-like equation, which we name it as multi-agent soft Bellman equation:

Theorem 2. We define the soft multi-agent Bellman equation for the soft state-action value function $Q_{soft}^{\pi, \rho}(s, a^i, a^{-i})$ of agent i as

$$Q_{soft}^{\pi, \rho}(s, a^i, a^{-i}) = r_t + \gamma \mathbb{E}_{(s_{t+1})} [V_{soft}^*(s_{t+1})]. \quad (12)$$

Proof. See Appendix C.3. \square

With this Bellman equation defined above, we can derive a solution to Eq. 12 with a fixed point iteration, which we call ROMMEO Q-iteration (ROMMEO-Q). Additionally, We can show that it can converge to the optimal Q_{soft}^* and V_{soft}^* with certain restrictions as stated in [Wen *et al.*, 2019]:

Theorem 3. ROMMEO Q-iteration. In a symmetric game with only one global optimum, i.e. $\mathbb{E}_{\pi^*} [Q_t^i(s)] \geq \mathbb{E}_{\pi} [Q_t^i(s)]$, where π^* is the optimal strategy profile. Let $Q_{soft}(\cdot, \cdot, \cdot)$ and $V_{soft}(\cdot)$ be bounded and assume

$$\sum_{a^{-i}} P(a^{-i}|s) \left(\sum_{a^i} \exp(\frac{1}{\alpha} Q_{soft}^*(s, a^i, a^{-i})) \right)^\alpha < \infty$$

and that $Q_{soft}^* < \infty$ exists. Then the fixed-point iteration

$$Q_{soft}(s_t, a_t^i, a_t^{-i}) \leftarrow r_t + \gamma \mathbb{E}_{(s_{t+1})} [V_{soft}(s_{t+1})], \quad (13)$$

where $V_{soft}(s_t) \leftarrow \log \sum_{a_t^{-i}} P(a_t^{-i}|s_t) \times \left(\sum_{a_t^i} \exp(\frac{1}{\alpha} Q_{soft}(s_t, a_t^i, a_t^{-i})) \right)^\alpha$, $\forall s_t, a_t^i, a_t^{-i}$, converges to Q_{soft}^* and V_{soft}^* respectively.

Proof. See Appendix C.3. \square

3.2 Regularized Opponent Model with Maximum Entropy Objective Actor Critic

The ROMMEO-Q assumes we have the model of the environment and is impractical to implement in high-dimensional continuous problems. To solve these problems, we propose the ROMMEO actor critic (ROMMEO-AC) which is a model-free method. We use neural networks (NNs) as function approximators for the conditional policy, opponent model and Q-function and learn these functions by stochastic gradient. We parameterize the Q-function, conditional policy and opponent model by $Q_\omega(s, a^i, a^{-i})$, $\pi_\theta(a_t^i|s_t, a_t^{-i})$ and $\rho_\phi(a_t^{-i}|s_t)$ respectively.

Without access to the environment model, we first replace the Q-iteration with Q-learning. Therefore, we can train ω to minimize:

$$\mathcal{J}_Q(\omega) = \mathbb{E}_{(s_t, a_t^i, a_t^{-i}) \sim \mathcal{D}} \left[\frac{1}{2} (Q_\omega(s_t, a_t^i, a_t^{-i}) - R(s_t, a_t^i, a_t^{-i}) - \gamma \mathbb{E}_{s_{t+1} \sim p_s} [\bar{V}(s_{t+1})])^2 \right], \quad (14)$$

with

$$\begin{aligned} \bar{V}(s_{t+1}) &= Q_{\bar{\omega}}(s_{t+1}, a_{t+1}^i, \hat{a}_{t+1}^{-i}) - \log \rho_\phi(\hat{a}_{t+1}^{-i}|s_{t+1}) \\ &\quad - \alpha \log \pi_\theta(a_{t+1}^i|s_{t+1}, \hat{a}_{t+1}^{-i}) + \log P(\hat{a}_{t+1}^{-i}|s_{t+1}), \end{aligned} \quad (15)$$

where $Q_{\bar{\omega}}$ are target functions for providing relatively stable target values. We use \hat{a}_t^{-i} denoting the action sampled from agent i 's opponent model $\rho(a_t^{-i}|s_t)$ and it should be distinguished from a_t^{-i} which is the real action taken by agent i 's opponent. Eq. 15 can be derived from Eq. 10 and 11.

To recover the optimal conditional policy and opponent model and avoid intractable inference steps defined in Eq. 10 and 11 in complex problems, we follow the method in [Haarnoja *et al.*, 2018b] where θ and ϕ are trained to minimize the KL-divergence:

$$\begin{aligned} \mathcal{J}_\pi(\theta) &= \mathbb{E}_{s_t \sim D, a_t^{-i} \sim \rho} \\ &\quad \left[D_{\text{KL}} \left(\pi_\theta(\cdot|s_t, \hat{a}_t^{-i}) \left\| \frac{\exp(\frac{1}{\alpha} Q_\omega(s_t, \cdot, \hat{a}_t^{-i}))}{Z_\omega(s_t, \hat{a}_t^{-i})} \right\| \right) \right], \end{aligned} \quad (16)$$

$$\begin{aligned} \mathcal{J}_\rho(\phi) &= \mathbb{E}_{(s_t, a_t^i) \sim D} \\ &\quad \left[D_{\text{KL}} \left(\rho(\cdot|s_t) \left\| \frac{P(\cdot|s_t) \left(\frac{\exp(\frac{1}{\alpha} Q(s_t, a_t^i, \cdot))}{\pi_\theta(a_t^i|s_t, \cdot)} \right)^\alpha}{Z_\omega(s_t)} \right\| \right) \right]. \end{aligned} \quad (17)$$

By using the reparameterization trick: $\hat{a}_t^{-i} = g_\phi(\epsilon_t^{-i}; s_t)$ and $a_t^i = f_\theta(\epsilon_t^i; s_t, \hat{a}_t^{-i})$, we can rewrite the objectives above as:

$$\begin{aligned} \mathcal{J}_\pi(\theta) &= \mathbb{E}_{s_t \sim D, \epsilon_t^i \sim N, \hat{a}_t^{-i} \sim \rho} [\alpha \log \pi_\theta(f_\theta(\epsilon_t^i; s_t, \hat{a}_t^{-i})) \\ &\quad - Q_\omega(s_t, f_\theta(\epsilon_t^i; s_t, \hat{a}_t^{-i}), \hat{a}_t^{-i})], \end{aligned} \quad (18)$$

$$\begin{aligned} \mathcal{J}_\rho(\phi) &= \mathbb{E}_{(s_t, a_t) \sim D, \epsilon_t^{-i} \sim N} [\log \rho_\phi(g_\phi(\epsilon_t^{-i}; s_t)|s_t) \\ &\quad - \log P(\hat{a}_t^{-i}|s_t) - Q(s_t, a_t^i, g_\phi(\epsilon_t^{-i}; s_t)) \\ &\quad + \alpha \log \pi_\theta(a_t^i|s_t, g_\phi(\epsilon_t^{-i}; s_t))]. \end{aligned} \quad (19)$$

The gradient of Eq. 14, 18 and 19 with respect to the corresponding parameters are listed as below:

$$\begin{aligned} \nabla_\omega \mathcal{J}_Q(\omega) &= \nabla_\omega Q_\omega(s_t, a_t^i, a_t^{-i}) (Q_\omega(s_t, a_t^i, a_t^{-i}) \\ &\quad - R(s_t, a_t^i, a_t^{-i}) - \gamma \bar{V}(s_{t+1})), \end{aligned} \quad (20)$$

$$\begin{aligned} \nabla_\theta \mathcal{J}_\pi(\theta) &= \nabla_\theta \alpha \log \pi_\theta(a_t^i|s_t, \hat{a}_t^{-i}) + \nabla_\theta f_\theta(\epsilon_t^i; s_t, \hat{a}_t^{-i}) \\ &\quad (\nabla_{a_t^i} \alpha \log \pi_\theta(a_t^i|s_t, \hat{a}_t^{-i}) - \nabla_{a_t^i} Q_\omega(s_t, a_t^i, \hat{a}_t^{-i})), \end{aligned} \quad (21)$$

$$\begin{aligned} \nabla_\phi \mathcal{J}_\rho(\phi) &= \nabla_\phi \log \rho_\phi(\hat{a}_t^{-i}|s_t) + (\nabla_{\hat{a}_t^{-i}} \log \rho_\phi(\hat{a}_t^{-i}|s_t) \\ &\quad - \nabla_{\hat{a}_t^{-i}} \log P(\hat{a}_t^{-i}|s_t) - \nabla_{\hat{a}_t^{-i}} Q_\omega(s_t, a_t^i, \hat{a}_t^{-i}) \\ &\quad + \nabla_{\hat{a}_t^{-i}} \alpha \log \pi_\theta(a_t^i|s_t, \hat{a}_t^{-i})) \nabla_\phi g_\phi(\epsilon_t^{-i}; s_t). \end{aligned} \quad (22)$$

We list the pseudo-code of ROMMEO-Q and ROMMEO-AC in Appendix A.




Figure 1: (a): Learning curves of ROMMEO and baselines on ICG over 100 episodes. (b): Probability of convergence to the global optimum for ROMMEO and baselines on ICG over 100 episodes. The vertical axis is the joint probability of taking actions A for both agents. (c): Probability of taking A estimated by agent i 's opponent model ρ^i and observed empirical frequency P^i in one trail of training, $i \in \{1, 2\}$.

4 Related Works

In early works, the maximum entropy principle has been used in policy search in linear dynamics [Todorov, 2010; Toussaint, 2009; Levine and Koltun, 2013] and path integral control in general dynamics [Kappen, 2005; Theodorou *et al.*, 2010]. Recently, off-policy methods [Haarnoja *et al.*, 2017; Schulman *et al.*, 2017; Nachum *et al.*, 2017] have been proposed to improve the sample efficiency in optimizing MEO. To avoid complex sampling procedure, training a policy in supervised fashion is employed in [Haarnoja *et al.*, 2018b]. Our work is closely related to this series of recent works because ROMMEO is an extension of MEO to MARL.

A few related works to ours have been conducted in multi-agent soft Q-learning [Wen *et al.*, 2019; Wei *et al.*, 2018; Grau-Moya *et al.*, 2018], where variants of soft Q-learning are applied for solving different problems in MARL. However, unlike previous works, we do not take the soft Q-learning as given and apply it to MARL problems with modifications. In our work, we first establish a novel objective ROMMEO and ROMMEO-Q is only an off-policy method we derive with complete convergence proof, which can optimize the objective. There are other ways of optimizing ROMMEO, for example, the on-policy gradient-based methods, but they are not included in the paper.

There has been substantial progress in combining RL with probabilistic inference. However, most of the existing works focus on the single-agent case. The literature of Bayesian methods in MARL is limited. Among these are methods performing on cooperative games with prior knowledge on distributions of the game model and the possible strategies of others [Chalkiadakis and Boutilier, 2003] or policy parameters and possible roles of other agents [Wilson *et al.*, 2010]. In our work, we assume very limited prior knowledge of the environment model, optimal policy, opponents or the observations during the play. In addition, our algorithms are fully *decentralized at training and execution*, which is more challenging than problems from the centralized training [Foerster *et al.*, 2017; Lowe *et al.*, 2017; Rashid *et al.*, 2018].

In our work, we give a new definition of optimality in CMARL and derive a novel objective ROMMEO. We provide two off-policy RL algorithms for optimizing ROMMEO and the exact version comes with convergence proof. In addition, we provide a natural perspective on opponent modeling in co-

ordination problems: biasing one's opponent model towards the optimum from its perspective but regularizing it with the empirical distribution of opponent's real behavior.

5 Experiments

5.1 Iterated Matrix Games

We first present the proof-of-principle result of ROMMEO-Q¹ on iterated matrix games where players need to cooperate to achieve the shared maximum reward. To this end, we study the iterated climbing games (ICG) which is a classic purely cooperative two-player stateless iterated matrix games.

Climbing game (CG) is a fully cooperative game proposed in [Claus and Boutilier, 1998] whose payoff matrix is summarized as follows:

$$R = \begin{matrix} & \begin{matrix} A & B & C \end{matrix} \\ \begin{matrix} A \\ B \\ C \end{matrix} & \begin{bmatrix} (11, 11) & (-30, -30) & (0, 0) \\ (-30, -30) & (7, 7) & (6, 6) \\ (0, 0) & (0, 0) & (5, 3) \end{bmatrix} \end{matrix}. \quad \text{It is}$$

a challenging benchmark because of the difficulty of convergence to its global optimum. There are two Nash equilibrium (A, A) and (B, B) but one global optimal (A, A) . The punishment of miscoordination by choosing a certain action increases in the order of $C \rightarrow B \rightarrow A$. The safest action is C and the miscoordination punishment is the most severe for A . Therefore it is very difficult for agents to converge to the global optimum in ICG.

We compare our method to a series of strong baselines in MARL, including Joint Action Learner (JAL) [Claus and Boutilier, 1998], WoLF Policy Hillclimbing (WoLF-PHC) [Bowling and Veloso, 2001], Frequency Maximum Q (FMQ) [Kapetanakis and Kudenko, 2002] and Probabilistic Recursive Reasoning (PR2) [Wen *et al.*, 2019]. ROMMEO-Q-EMP is an ablation study to evaluate the effectiveness of our proposed opponent model learning process, where we replace our opponent model with empirical frequency. Fig. 1a shows the learning curves on ICG for different algorithms. The difference of rewards between ROMMEO-Q and FMQ-c10 may seem small because of the small reward margin between the global optimum and the local one. However, ROMMEO-Q actually outperforms all baselines significantly in terms of converging to the global optimum, which is shown

¹The experiment code and appendix are available at <https://github.com/rommeoijcai2019/rommeo>.




Figure 2: Experiment on Max of Two Quadratic Game. (a) Reward surface and learning path of agents. Scattered points are actions taken at each step; (b) Learning curve of ROMMEO and baselines. (c) Mean of agents’ policies π and opponent models ρ .

in Fig. 1b. To further analyze the opponent modeling described in Sec. 2.3, we visualize the probability of agent $-i$ taking the optimal action A estimated by agent i ’s opponent model ρ^i and its true policy π^{-i} in Fig. 1c. Agent i ’s opponent model “thinks ahead” of agent $-i$ and converges to agent $-i$ ’s optimal policy before agent $-i$ itself converges to the optimal policy. This helps agent i to respond to its opponent model optimally by choosing action A , which in turn leads to the improvement of agent $-i$ ’s opponent model and policy. Therefore, the game converges to the global optimum. To note, the big drop of $P(A)$ for both policies and opponent models at the beginning of the training comes from the severe punishment of miscoordination associated with action A .

5.2 Differential Games

We adopt the differential Max of Two Quadratic Game [Wei *et al.*, 2018] for continuous case. The agents have continuous action space of $[-10, 10]$. Each agent’s reward depends on the joint action following the equations: $r^1(a^1, a^2) = r^2(a^1, a^2) = \max(f_1, f_2)$, where $f_1 = 0.8 \times [-(\frac{a^1+5}{3})^2 - (\frac{a^2+5}{3})^2]$, $f_2 = 1.0 \times [-(\frac{a^1-5}{1})^2 - (\frac{a^2-5}{1})^2] + 10$. We compare the algorithm with a series of baselines including PR2 [Wen *et al.*, 2019], MASQL [Wei *et al.*, 2018; Grau-Moya *et al.*, 2018], MADDPG [Lowe *et al.*, 2017] and independent learner via DDPG [Lillicrap *et al.*, 2015]. To compare against traditional opponent modeling methods, similar to [Rabinowitz *et al.*, 2018; He *et al.*, 2016], we implement an additional baseline of DDPG with an opponent module that is trained online with supervision in order to capture the latest opponent behaviors, called DDPG-OM. We trained all agents for 200 episodes with 25 steps per episode.

This is a challenging task to most continuous gradient based RL algorithms because gradient update tends to direct the training agent to the sub-optimal point. The reward surface is provided in Fig. 2a; there is a local maximum 0 at $(-5, -5)$ and a global maximum 10 at $(5, 5)$, with a deep valley staying in the middle. If the agents’ policies are initialized to $(0, 0)$ (the red starred point) that lies within the basin of the left local maximum, the gradient-based methods would tend to fail to find the global maximum equilibrium point due to the valley blocking the upper right area.

A learning path of ROMMEO-AC is summarized in Fig. 2a and the solid bright circle on the right corner implies the convergence to the global optimum. The learning curve is

presented in Fig. 2b, ROMMEO-AC shows the capability of converging to the global optimum in a limited amount of steps, while most of the baselines can only reach the sub-optimal point. PR2-AC can also achieve the global optimum but requires many more steps to explore and learn. Additionally, fine tuning on the exploration noise or separate exploration stage is required for deterministic RL methods (MADDPG, DDPG, DDPG-OM, PR2-AC), and the learning outcomes of energy-based RL method (MASQL) is extremely sensitive to the annealing scheme for the temperature. In contrast, ROMMEO-AC employs a stochastic policy and controls the exploration level by the weighting factor α . It does not need a separate exploration stage at the beginning of the training or a delicately designed annealing scheme for α .

Furthermore, we analyze the learning path of policy π and modeled opponent policy ρ during the training, the results are shown in Fig. 2c. The red and orange lines are mean of modeled opponent policy ρ , which always learn to approach the optimum ahead of the policy π (in dashed blue and green lines). This helps the agents to establish the trust and converge to the optimum quickly, which further justifies the effectiveness and benefits of conducting a regularized opponent model proposed in Sec. 2.3.

6 Conclusion

In this paper, we use Bayesian inference to formulate MARL problem and derive a novel objective ROMMEO which gives rise to a new perspective on opponent modeling. We design an off-policy algorithm ROMMEO-Q with complete convergence proof for optimizing ROMMEO. For better generality, we also propose ROMMEO-AC, an actor critic algorithm powered by NNs to solve complex and continuous problems. We give an insightful analysis of the effect of the new learning process of the opponent modeling on agent’s performance in MARL. We evaluate our methods on the challenging matrix game and differential game and show that they can outperform a series of strong base lines. It is worthy of noting that Theorems 1,2 and 3 only guarantees the convergence to optimal solutions with respect to ROMMEO objective but not the optimum in the game. The achievement of the optimum in the game also relies on the opponent learning algorithm. In our work, we demonstrate that ROMMEO-Q/AC’s convergence to the optimum of the game in self-play setting. The convergence to optimum of games in non-self-play settings will be studied in our future work.

References

- [Abdolmaleki *et al.*, 2018] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Rémi Munos, Nicolas Heess, and Martin A. Riedmiller. Maximum a posteriori policy optimisation. *CoRR*, abs/1806.06920, 2018.
- [Bowling and Veloso, 2001] Michael Bowling and Manuela Veloso. Rational and convergent learning in stochastic games. In *IJCAI*, San Francisco, CA, USA, 2001.
- [Brown, 1951] George W. Brown. Iterative solution of games by fictitious play. In *AAPA*. New York, 1951.
- [Chalkiadakis and Boutilier, 2003] Georgios Chalkiadakis and Craig Boutilier. Coordination in multiagent reinforcement learning: A bayesian approach. In *AAMAS*, pages 709–716, New York, NY, USA, 2003. ACM.
- [Claus and Boutilier, 1998] Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multi-agent systems. AAAI ’98/AAAI ’98, Menlo Park, CA, USA, 1998. American Association for Artificial Intelligence.
- [Florensa *et al.*, 2017] Carlos Florensa, Yan Duan, and Pieter Abbeel. Stochastic Neural Networks for Hierarchical Reinforcement Learning. *arXiv e-prints*, April 2017.
- [Foerster *et al.*, 2017] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual Multi-Agent Policy Gradients. *arXiv e-prints*, page arXiv:1705.08926, May 2017.
- [Fox *et al.*, 2015] Roy Fox, Ari Pakman, and Naftali Tishby. Taming the Noise in Reinforcement Learning via Soft Updates. *arXiv e-prints*, page arXiv:1512.08562, December 2015.
- [Grau-Moya *et al.*, 2018] Jordi Grau-Moya, Felix Leibfried, and Haitham Bou-Ammar. Balancing Two-Player Stochastic Games with Soft Q-Learning. *arXiv e-prints*, page arXiv:1802.03216, February 2018.
- [Haarnoja *et al.*, 2017] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. *CoRR*, abs/1702.08165, 2017.
- [Haarnoja *et al.*, 2018a] Tuomas Haarnoja, Vitchyr Pong, Aurick Zhou, Murtaza Dalal, Pieter Abbeel, and Sergey Levine. Composable Deep Reinforcement Learning for Robotic Manipulation. *arXiv e-prints*, page arXiv:1803.06773, March 2018.
- [Haarnoja *et al.*, 2018b] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *arXiv e-prints*, page arXiv:1801.01290, January 2018.
- [He *et al.*, 2016] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé III. Opponent modeling in deep reinforcement learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1804–1813, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [Heinrich and Silver, 2016] Johannes Heinrich and David Silver. Deep reinforcement learning from self-play in imperfect information games. *CoRR*, abs/1603.01121, 2016.
- [Kalman, 1960] Rudolf Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME - Journal of basic Engineering*, 82:35–45, 01 1960.
- [Kapetanakis and Kudenko, 2002] Spiros Kapetanakis and Daniel Kudenko. Reinforcement learning of coordination in cooperative multi-agent systems. In *Eighteenth National Conference on Artificial Intelligence*, Menlo Park, CA, USA, 2002.
- [Kappen, 2005] Hilbert J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11), 2005.
- [Levine and Koltun, 2013] Sergey Levine and Vladlen Koltun. Variational policy search via trajectory optimization. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *NIPS ’13*. 2013.
- [Levine, 2018] Sergey Levine. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. *arXiv e-prints*, page arXiv:1805.00909, May 2018.
- [Lillicrap *et al.*, 2015] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [Lowe *et al.*, 2017] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. *arXiv e-prints*, page arXiv:1706.02275, June 2017.
- [Nachum *et al.*, 2017] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the Gap Between Value and Policy Based Reinforcement Learning. *arXiv e-prints*, page arXiv:1702.08892, February 2017.
- [Rabinowitz *et al.*, 2018] Neil C Rabinowitz, Frank Perbet, H Francis Song, Chiyuan Zhang, SM Eslami, and Matthew Botvinick. Machine theory of mind. *arXiv preprint arXiv:1802.07740*, 2018.
- [Raileanu *et al.*, 2018] Roberta Raileanu, Emily Denton, Arthur Szlam, and Rob Fergus. Modeling others using oneself in multi-agent reinforcement learning. *CoRR*, abs/1802.09640, 2018.
- [Rashid *et al.*, 2018] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *arXiv e-prints*, page arXiv:1803.11485, March 2018.
- [Rawlik *et al.*, 2013] Konrad Rawlik, Marc Toussaint, and Sethu Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference (extended abstract). *IJCAI ’13*, pages 3052–3056, 2013.
- [Schulman *et al.*, 2017] John Schulman, Xi Chen, and Pieter Abbeel. Equivalence Between Policy Gradients and Soft Q-Learning. *arXiv e-prints*, page arXiv:1704.06440, April 2017.
- [Shapley, 1953] Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- [Theodorou *et al.*, 2010] Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *J. Mach. Learn. Res.*, 11, December 2010.
- [Tian *et al.*, 2018] Zheng Tian, Shihao Zou, Tim Warr, Lisheng Wu, and Jun Wang. Learning Multi-agent Implicit Communication Through Actions: A Case Study in Contract Bridge, a Collaborative Imperfect-Information Game. *arXiv e-prints*, page arXiv:1810.04444, October 2018.
- [Todorov, 2007] Emanuel Todorov. Linearly-solvable markov decision problems. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *NIPS ’07*. 2007.
- [Todorov, 2010] Emanuel Todorov. Policy gradients in linearly-solvable mdp. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *NIPS*. 2010.
- [Toussaint and Storkey, 2006] Marc Toussaint and Amos Storkey. Probabilistic inference for solving discrete and continuous state markov decision processes. *ICML ’06*, pages 945–952, New York, NY, USA, 2006. ACM.
- [Toussaint, 2009] Marc Toussaint. Robot trajectory optimization using approximate inference. *ICML ’09*, pages 1049–1056, New York, NY, USA, 2009. ACM.
- [Wei *et al.*, 2018] Ermo Wei, Drew Wicke, David Freelan, and Sean Luke. Multiagent soft q-learning. *AAAI*, 2018.
- [Wen *et al.*, 2019] Ying Wen, Yaodong Yang, Rui Luo, Jun Wang, and Wei Pan. Probabilistic recursive reasoning for multi-agent reinforcement learning. In *ICLR*, 2019.
- [Wilson *et al.*, 2010] Aaron Wilson, Alan Fern, and Prasad Tadepalli. Bayesian policy search for multi-agent role discovery. *AAAI ’10*, pages 624–629. AAAI Press, 2010.
- [Ziebart *et al.*, 2008] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI ’08*, 2008.