

---

# The Emergence of Individuality

---

Jiechuan Jiang<sup>1</sup> Zongqing Lu<sup>1</sup>

## Abstract

Individuality is essential in human society. It induces the division of labor and thus improves the efficiency and productivity. Similarly, it should also be a key to multi-agent cooperation. Inspired by that individuality is of being an individual separate from others, we propose a simple yet efficient method for the emergence of individuality (EOI) in multi-agent reinforcement learning (MARL). EOI learns a probabilistic classifier that predicts a probability distribution over agents given their observation and gives each agent an intrinsic reward of being correctly predicted by the classifier. The intrinsic reward encourages the agents to visit their own familiar observations, and learning the classifier by such observations makes the intrinsic reward signals stronger and in turn makes the agents more identifiable. To further enhance the intrinsic reward and promote the emergence of individuality, two regularizers are proposed to increase the discriminability of the classifier. We implement EOI on top of popular MARL algorithms. Empirically, we show that EOI outperforms existing methods in a variety of multi-agent cooperative scenarios.

## 1. Introduction

Humans develop into distinct individuals due to both genes and environments (Freund et al., 2013). Individuality induces the division of labor (Gordon, 1996), which improves the productivity and efficiency of human society. Analogically, the emergence of individuality should also be essential for multi-agent cooperation.

Although multi-agent reinforcement learning (MARL) has been applied to multi-agent cooperation, it is widely observed that agents usually learn similar behaviors, especially when the agents are homogeneous with shared global reward and co-trained (McKee et al., 2020). For example, in

multi-camera multi-object tracking (Liu et al., 2017), where camera agents learn to cooperatively track multiple objects, the camera agents all tend to track the easy object. However, such similar behaviors can easily make the learned policies fall into local optimum. If the agents can respectively track different objects, they are more likely to solve the task optimally. Many studies formulate such a problem as task allocation or role assignment (Sander et al., 2002; Dastani et al., 2003; Sims et al., 2008). However, they require that the agent roles are rule-based and the tasks are pre-defined, and thus are not general methods. Some studies intentionally pursue difference in agent policies by diversity (Lee et al., 2020; Yang et al., 2020a) or by emergent roles (Wang et al., 2020a), however, the induced difference is not appropriately linked to the success of task. On the contrary, the emergence of individuality along with learning cooperation can automatically drive agents to behave differently and take a variety of roles, if needed, to successfully complete tasks.

Biologically, the emergence of individuality is attributed to innate characteristics and experiences. However, as in practice RL agents are mostly homogeneous, we mainly focus on enabling agents to develop individuality through interactions with the environment during policy learning. Intuitively, in multi-agent environments where agents respectively explore and interact with the environment, individuality should emerge from what they experience. In this paper, we propose a novel method for the emergence of individuality (EOI) in MARL. EOI learns a probabilistic classifier that predicts a probability distribution over agents given their observation and gives each agent an intrinsic reward of being correctly predicted probability by the classifier. Encouraged by the intrinsic reward, agents tend to visit their own familiar observations. Learning the probabilistic classifier by such observations makes the intrinsic reward signals stronger and in turn makes the agents more identifiable. In this closed loop with positive feedback, agent individuality emerges gradually. However, at early learning stage, the observations visited by different agents cannot be easily distinguished by the classifier, meaning the intrinsic reward signals are not strong enough to induce agent characteristics. Therefore, we propose two regularizers for learning the classifier to increase the discriminability, enhance the feedback, and thus promote the emergence of

---

<sup>1</sup>Peking University. Correspondence to: Zongqing Lu <zongqing.lu@pku.edu.cn>.

individuality.

EOI is compatible with centralized training and decentralized execution (CTDE) methods. We realize EOI on top of two popular MARL methods, MAAC (Iqbal & Sha, 2019) and QMIX (Rashid et al., 2018). For MAAC, as each agent has its own critic, it is convenient to shape the reward for each agent. For QMIX, we introduce an auxiliary gradient and update the individual value function by both minimizing the TD error of the joint action-value function and maximizing its cumulative intrinsic rewards. In experiments, we verify the effectiveness of the intrinsic reward and confirm that the proposed regularizers indeed improve the emergence of individuality even if agents have the same innate characteristics by ablation studies. And we empirically demonstrate that EOI outperforms existing methods in both grid-world and large-scale environments. Finally, we discuss and numerically show the similarity and difference between EOI and DIAYN (Eysenbach et al., 2019).

## 2. Related Work

### 2.1. MARL

We consider the formulation of Decentralized Partially Observable Markov Decision Process (Dec-POMDP) (Oliehoek et al., 2016). There are  $n$  agents in an environment. At each timestep  $t$  each agent  $i$  receives a local observation  $o_i^t$ , takes an action  $a_i^t$ , and gets a shared global reward  $r^t$ . Agents together aim to maximize the expected return  $\mathbb{E} \sum_{t=0}^T \gamma^t r^t$ , where  $\gamma$  is a discount factor and  $T$  is the time horizon. Many methods have been proposed for Dec-POMDP, most of which follow the paradigm of centralized training and decentralized execution. Value function factorization methods decompose the joint value function into individual value functions. VDN (Sunehag et al., 2018) and QMIX (Rashid et al., 2018) respectively propose additivity and monotonicity for factorization structures. Qatten (Yang et al., 2020b) is a variant of VDN, which uses a multi-head attention structure to utilize global information. QPLEX (Wang et al., 2021a) uses a duplex dueling network architecture for factorization and theoretically analyzes the full representation expressiveness. Some methods extend policy gradient into multi-agent cases, which contain a centralized critic with global information and decentralized actors which only have access to local information. COMA (Foerster et al., 2018) proposes a counterfactual baseline for multi-agent credit assignment. MADDPG (Lowe et al., 2017) is an extension of DDPG algorithm (Lillicrap et al., 2016). DOP (Wang et al., 2021c) replaces the conventional critic with a value decomposed critic. Communication methods (Das et al., 2019; Jiang et al., 2020; Ding et al., 2020) share information between agents for implicit coordination, and DCG (Böhmer et al., 2020) adopts coordination graphs for explicit coordination.

### 2.2. Behavior Diversification

Many cooperative multi-agent applications require agents to take different behaviors to complete the task successfully. Behavior diversification can be handcrafted or emerge through agents' learning. Handcrafted diversification is widely studied as task allocation or role assignment. Heuristics (Sander et al., 2002; Dastani et al., 2003; Sims et al., 2008; Macarthur et al., 2011) assign specific tasks or pre-defined roles to each agent based on goal, capability, visibility, or by search. M<sup>3</sup>RL (Shu & Tian, 2019) learns a manager to assign suitable sub-tasks to rule-based workers with different preferences and skills. These methods require that the sub-tasks and roles are pre-defined, and the worker agents are rule-based. However, in general, the task cannot be easily decomposed even with domain knowledge and workers are learning agents.

The emergent diversification for single agent has been studied in DIAYN (Eysenbach et al., 2019), which learns reusable diverse skills in complex and transferable tasks without any reward signal by maximizing the mutual information between states and skill embeddings as well as entropy. In multi-agent learning, SVO (McKee et al., 2020) introduces diversity into heterogeneous agents for more generalized and high-performing policies in social dilemmas. Some methods are proposed for behavior diversification in multi-agent cooperation. ROMA (Wang et al., 2020a) learns a role encoder to generate role embedding, and learns a role decoder to generate neural network parameters from embedding. Two regularizers are introduced for learning identifiable and specialized roles. However, due to large parameter space, generating various parameters for the emergent roles is inefficient. And mode collapse would happen in the role decoder even with different role embeddings. Learning low-level skills for each agent using DIAYN is considered in (Lee et al., 2020; Yang et al., 2020a), where agents' diverse low-level skills are coordinated by the high-level policy. However, the diversity is not considered in the high-level policy.

## 3. Method

Individuality is of being an individual separate from others. Motivated by this, we propose EOI, where agents are intrinsically rewarded in terms of being correctly predicted by a probabilistic classifier that is learned based on agents' observations. If the classifier learns to accurately distinguish agents, agents should behave differently and thus individuality emerges. Two regularizers are introduced for learning the classifier to enhance intrinsic reward signal and promote individuality. EOI directly correlates individuality with the task by intrinsic reward, and thus individuality emerges naturally during agents' learning. EOI can be applied to Dec-POMDP tasks and trained along with CTDE algorithms.




Figure 1. Multi-camera multi-target capturing

We design practical techniques to implement EOI on top of two popular MARL methods, MAAC and QMIX.

### 3.1. Intrinsic Reward

As illustrated in Figure 1, two camera agents learn to capture two targets, where the closer and slower target 1 is the easier one. Capturing each target, they will get a global reward of +1. Sequentially capturing the two targets by one agent or together is sub-optimal, sometimes impossible when the task has a limited time horizon. The optimal solution is that the two agents go capturing different targets simultaneously. It is easy for both agents to learn to capture the easier target 1. However, after that, target 2 becomes even harder to be explored, and the learned policies easily fall at local optimum. Nevertheless, the emergence of individuality can address the problem, *e.g.*, agents prefer to capture different targets.

To enable agents to develop individuality, EOI learns a probabilistic classifier  $P(I|O)$  to predict a probability distribution over agents given on their observation, and each agent takes the correctly predicted probability as the intrinsic reward at each timestep. Thus, the reward function for agent  $i$  is modified as

$$r + \alpha p(i|o_i),$$

where  $r$  is the global environmental reward,  $p(i|o_i)$  is the predicted probability of agent  $i$  given its observation  $o_i$ , and  $\alpha$  is a tuning hyperparameter to weight the intrinsic reward. With the reward shaping, EOI works as follows. If there is initial difference between agent policies in terms of visited observations, the difference is captured by  $P(I|O)$  as it is fitted using agents' experiences. The difference is then fed back to each agent as an intrinsic reward. As agents maximize the expected return, the difference in agents' policies is exacerbated together with optimizing the environmental return. Therefore, the learning process is a closed loop with positive feedback. As agents progressively behave more identifiably, the classifier can distinguish agents more

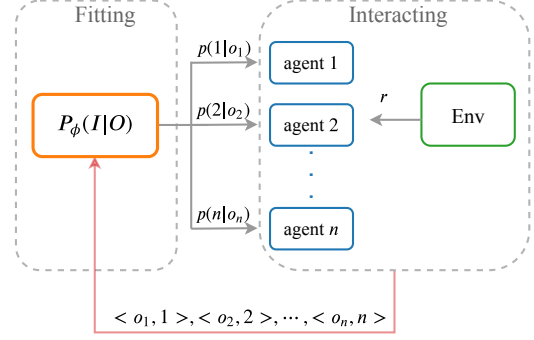


Figure 2. EOI

accurately, and thus individuality emerges gradually.

The classifier  $P_\phi(I|O)$  is parameterized by a neural network  $\phi$  and learned in a supervised way. At each timestep, we take each agent  $i$ 's observation  $o_i$  as input and the agent index  $i$  as the label and store the pair  $\langle o_i, i \rangle$  into a buffer  $\mathcal{B}$ .  $\phi$  is updated by minimizing the cross-entropy loss (CE), which is computed based on the uniformly sampled batches from  $\mathcal{B}$ . The learning process of EOI is illustrated in Figure 2.

### 3.2. Regularizers of $P_\phi(I|O)$

In the previous section, we assume that there is some difference between agents' policies. However, in general, the difference between initial policies is small (even no differences if agents' policies are initially by the same network weights), and the policies will quickly learn similar behaviors as in the example in Figure 1. Therefore, the intrinsic rewards are nearly the same for each agent, which means no feedback in the closed loop. To generate the feedback in the closed loop, the observation needs to be identifiable and thus the agent can be distinguished in terms of observations by  $P_\phi(I|O)$ . To address this, we propose two regularizers: positive distance (PD) and mutual information (MI) for learning  $P_\phi(I|O)$ .

**Positive Distance.** The positive distance is inspired from the triplet loss (Schroff et al., 2015) in contrastive learning, which is proposed to learn identifiable embeddings. Since  $o_i^t$  and its previous observations  $\{o_i^{t-\Delta t}, o_i^{t-\Delta t+1}, \dots, o_i^{t-1}\}$  are distributed on the trajectory generated by agent  $i$ , the previous observations in the  $\Delta t$ -length window could be seen as the positives of  $o_i^t$ . To make the probability distribution on the anchor  $o_i^t$  close to that on the positives, we sample an observation  $o_i^{t'}$  from  $\{o_i^{t-\Delta t}, o_i^{t-\Delta t+1}, \dots, o_i^{t-1}\}$  and minimize the cross-entropy loss

$$\text{CE} \left( p_\phi(\cdot | o_i^t), p(\cdot | o_i^{t'}) \right).$$

The positive distance minimizes the intra-distance between the observations with the same "identity", which hence enlarges the margin between different "identities". As a

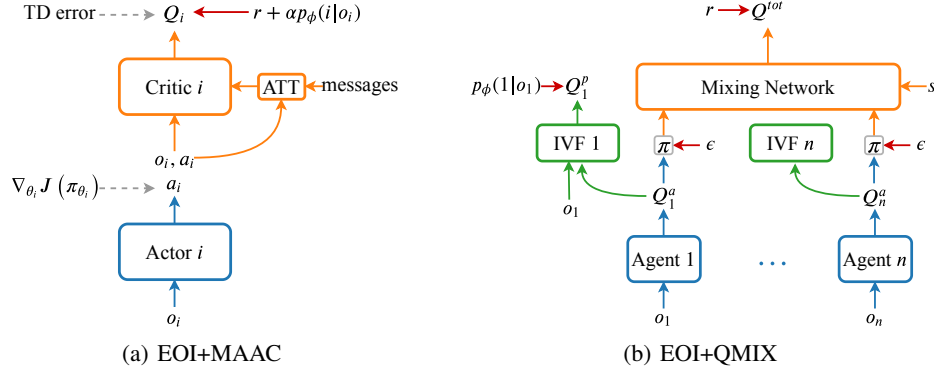


Figure 3. Illustration of EOI with MAAC and QMIX.

result, the observations become more identifiable. Since the positives are naturally defined on the trajectory, the identifiability generated by the positive distance is actually induced by the agent policy. Defining the negatives is hard but we find that just using the positive distance works well in practice.

**Mutual Information.** If the observations are more identifiable, it is easier to infer the agent that visits the given observation most, which indicates the higher mutual information between the agent index and observation. Therefore, to further increase the discriminability of the classifier, we maximize their mutual information,

$$\begin{aligned} \text{MI}(I; O) &= \mathcal{H}(I) - \mathcal{H}(I|O) \\ &= \mathcal{H}(I) - \mathbb{E}_{o \sim p(o)} \left[ \sum_i -p(i|o) \log p(i|o) \right]. \end{aligned}$$

Since we store  $\langle o_i, i \rangle$  of every agent in  $\mathcal{B}$ , the number of samples for each agent is equal. Fitting  $P_\phi(I|O)$  using batches from  $\mathcal{B}$  ensures  $\mathcal{H}(I)$  is a constant. To maximize  $\text{MI}(I; O)$  is to minimize  $\mathcal{H}(I|O)$ . Therefore, equivalently, we sample batches from  $\mathcal{B}$  and minimize

$$\text{CE}(p_\phi(\cdot|o_i^t), p_\phi(\cdot|o_i^{t'})).$$

Therefore, the optimization objective of  $P_\phi(I|O)$  is to minimize

$$\begin{aligned} &\text{CE}(p_\phi(\cdot|o_i^t), \text{one\_hot}(i)) + \beta_1 \text{CE}(p_\phi(\cdot|o_i^t), p_\phi(\cdot|o_i^{t'})) \\ &+ \beta_2 \text{CE}(p_\phi(\cdot|o_i^t), p_\phi(\cdot|o_i^{t'})), \end{aligned}$$

where  $\beta_1$  and  $\beta_2$  are hyperparameters. The regularizers increase the discriminability of  $P_\phi(I|O)$ , make the intrinsic reward signals stronger to stimulate the agents to be more distinguishable, and eventually promote the emergence of individuality. In this sense,  $P_\phi(I|O)$  not only is the posterior statistics, but also serves as the inductive bias of agents' learning.

### 3.3. Implementation with MAAC and QMIX

Existing methods for reward shaping in MARL focus on independent learning agents (McKee et al., 2020; Wang et al., 2020b; Du et al., 2019). How to shape the reward in Dec-POMDP for centralized training has not been deeply studied. Since the intrinsic reward can be exactly assigned to the specific agent, individually maximizing the intrinsic reward is more efficient than jointly maximizing the sum of all agents' intrinsic rewards (Hughes et al., 2018). Adopting this idea, we respectively present the implementation with MAAC and QMIX for realizing EOI. MAAC is an off-policy actor-critic algorithm, where each agent learns its own critic, thus it is convenient to directly give the shaped reward  $r + \alpha p_\phi(i|o_i)$  to the critic of each agent  $i$ , without modifying other components, as illustrated in Figure 3(a). The TD error of the critic and the policy gradient are the same as in MAAC (Iqbal & Sha, 2019).

In QMIX, each agent  $i$  has an individual action-value function  $Q_i^a$ . All the individual action-value functions are monotonically mixed into a joint action-value  $Q^{\text{tot}}$  by a mixing network. Each agent selects the action with the highest individual value, but the individual value has neither actual meaning nor constraints (Rashid et al., 2018). Therefore, we can safely introduce an auxiliary gradient of the intrinsic reward to the individual action-value function  $Q_i^a$ . Each agent  $i$  learns an intrinsic value function (IVF)  $Q_i^p$ , which takes as input the observation  $o_i$  and the individual action-value vector  $Q_i^a(o_i)$  and approximates  $\mathbb{E} \sum_{t=0}^T \gamma^t p(i|o_i^t)$  by minimizing the TD error,

$$\begin{aligned} &\mathbb{E}_{\langle o_i, Q_i^a(o_i), o_i' \rangle \sim \mathcal{D}} [(Q_i^p(o_i, Q_i^a(o_i)) - y)^2], \\ &\text{where } y = p_\phi(i|o_i) + \gamma \bar{Q}_i^p(o_i', \bar{Q}_i^a(o_i')). \end{aligned}$$

$\bar{Q}_i^a$  and  $\bar{Q}_i^p$  are the target value functions and  $\mathcal{D}$  is the replay buffer. In order to improve both the global reward and intrinsic reward, we update  $Q_i^a$ , parameterized by  $\theta_i$ , towards maximizing  $\mathbb{E}[Q_i^p(o_i, Q_i^a(o_i; \theta_i))]$  along with minimizing the TD error of  $Q^{\text{tot}}$  (denoted as  $\delta^{\text{tot}}$ ), as illustrated in Fig-






Figure 4. Illustration of Pac-Men and Windy Maze.

ure 3(b). Since the intrinsic value function is differentiable with respect to the individual action-value vector  $Q_i^a(o_i; \theta_i)$ , and the action-value vector is continuous, we can establish the connection between  $Q_i^p$  and  $Q_i^a$  by the chain rule, like the policy update in DDPG (Lillicrap et al., 2016), and the gradient of  $\theta_i$  is,

$$\nabla_{\theta_i} J(\theta_i) = \frac{\partial \delta^{tot}}{\partial \theta_i} - \alpha \frac{\partial Q_i^p(o_i, Q_i^a(o_i; \theta_i))}{\partial \theta_i}.$$

MAAC and QMIX are off-policy algorithms, and the environmental rewards are stored in the replay buffer. However, the intrinsic rewards are recomputed in the sampled batches before each update, since  $P_\phi(I|O)$  is co-evolving with the learning agents and hence the previous intrinsic reward is outdated. The joint learning process of the classifier and agent policies can be mathematically formulated as a bi-level optimization, which is detailed in Appendix A. Note that in EOI, we can easily replace local observation with trajectory by taking the hidden state of RNN (it takes the trajectory as input) as the input of the classifier, which might be helpful in complex environments.

## 4. Experiments

In the evaluation, we first verify the effectiveness of the intrinsic reward and the two regularizers on both MAAC and QMIX by ablation studies. Then, we compare EOI against EDTI (Wang et al., 2020b) to verify the advantages of individuality with sparse reward, and against ROMA (Wang et al., 2020a) and Lee et al. (2020) (denoted as HC) to investigate the advantages of individuality over emergent roles and diversity. We also discuss the similarity and difference between EOI and DIAYN (Eysenbach et al., 2019), and provide the numerical results. In grid-world environments, the agents do not share the weights of neural network since parameter sharing causes similar agent behaviors. All the curves are plotted using mean and standard deviation. The details about the experimental settings and the hyperparameters are available in Appendix B.

### 4.1. Performance and Ablation

To clearly interpret the mechanism of EOI, we test EOI in the two scenarios: Pac-Men and Windy Maze, as illustrated




Figure 5. Learning curves in Pac-Men (top) and Windy Maze (bottom).

in Figure 4.

- *Pac-Men*. There are four agents initialized at the maze center, and some randomly initialized dots. At each timestep, each agent can move to one of four neighboring grids or eat a dot.
- *Windy Maze*. There are two agents initialized at the bottom of the T-shaped maze, and two dots initialized at the two ends. At each timestep, each agent could move to one of four neighboring grids or eat a dot. There is a wind running right from the dotted line. Shifted by the wind, forty percent of the time, the agent will move to the right grid whichever action it takes.

Each agent has a local observation that contains a square view with  $5 \times 5$  grids centered at the agent itself, and could only get a global reward, *i.e.*, the total eaten dots, at the final timestep.

In Figure 5 (top), MAAC and QMIX get the lowest environmental reward respectively, since some agents learn to go to the same room and compete for the dots. This can be verified by the position distribution of QMIX agents in Figure 6(a), where three agents move to the same room. MAAC agents behave similarly. At the early training, it is easy for the agents to explore the bottom room and eat dots there to improve the environmental reward. Once the agents learn such policies, it is hard to explore other rooms, so the agents learn similar behaviors and fall at the local optimum.

Driven by the intrinsic reward without both regularizers, EOI obtains better performance than MAAC and QMIX.




Figure 6. Distributions of agents' positions of QMIX (a) and EOI+QMIX (b). The darker color means the higher value.





Figure 7. Kernel PCA of agents' observations of EOI+QMIX in Pac-Men. Darker color means higher correctly predicted probability of  $P_\phi(I|O)$ .

But the improvement is not significant since the observations of different agents cannot be easily distinguished when there is little difference between initial policies. The regularizers of PD and MI can increase the discriminability of  $P_\phi(I|O)$ , providing stronger intrinsic signals. Guided by  $P_\phi(I|O)$  with PD or MI, the agents go to different rooms and eat more dots. MI theoretically increases the discriminability even the initial policies have no differences, while PD makes the observations distinguishable according to policies. Combining the advantages of the two regularizers leads to higher and steadier performance, as shown in Figure 5 (top). With both two regularizers, the agents respectively go to the four rooms and achieve the highest reward, which is indicated in Figure 6(b). We also visualize the observations of different agents by kernel PCA, as illustrated in Figure 7, where darker color means higher correctly predicted probability of  $P_\phi(I|O)$ . We can see  $P_\phi(I|O)$  can easily distinguish agents given their observations.


Similar results could be observed in Figure 5 (bottom). Under the effect of wind, it is easy to eat the right dot, even if the agent acts randomly. Limited by the time horizon, it is impossible that the agent first eats the right dot and then goes to the left end for the other dot. MAAC and QMIX only achieve the mean reward of 1. Due to the wind and the small state space, the trajectories of the two agents are similar, thus EOI without regularizers provides little help, where the intrinsic reward is not strong enough to induce individuality. With regularizers, the observations on the right path and the left path can be discriminated gradually. In the learning process, the agents will first both learn to eat



(a) EOI+MAAC

(b) EOI+QMIX

Figure 8. Learning curves of intrinsic reward and environmental reward.



(a) Action distribution

(b) EOI+QMIX

Figure 9. Action distributions of the initial action-value functions in QMIX and learning curves with EOI+QMIX. The dotted lines are the version with the same initial action-value functions.

the right dot. Then one of them will change its policy and go left for higher intrinsic reward, and eventually the agents develop the distinct policies and get a higher mean reward.

To further investigate the effect of the regularizers, we show the learning curves of intrinsic reward and environmental reward of EOI with and without regularizers in Figure 8. EOI with the regularizers converges to higher intrinsic reward than that without regularizers, meaning agents behave more distinctly. With regularizers, the rapid increase of intrinsic reward occurs before that of environmental reward, which indicates the regularizers make  $P_\phi(I|O)$  also serve as the inductive bias for the emergence of individuality.

We also investigate the influence of the difference between initial policies. The action distributions (over visited observations) of the four initial action-value functions in QMIX are illustrated in Figure 9(a), where we find there is a large difference between them. The inherent difference makes agents distinguishable initially, which is the main reason EOI without the regularizers works well. We then initiate the four action-value functions with the same weights and re-run the experiments. The performance of EOI without the regularizers drops considerably, while EOI with the regularizers has almost no difference, as illustrated in Figure 9(b), indicating that PD and MI make the learning more robust to the initial policies. Agent individuality still emerges even with the same innate characteristics.

In Figure 10, we test EOI+MAAC with different  $\alpha$  to inves-





Figure 10. Learning curves of EOI+MAAC with different  $\alpha$  in Pac-Men.



(a) Firefighters (b) Learning curves


Figure 11. Firefighters (a) and learning curves (b).

tigate the effect of  $\alpha$  on the emergent individuality. When  $\alpha$  is too large, the agents will pay much attention to learn the individualized behaviors, which harms the optimization of the cumulative reward. The choice of  $\alpha$  is a trade-off between success and individuality.

#### 4.2. Comparison with the Existing Methods


To compare EOI with previous methods, we design a more challenging task, Firefighters. There are some burning grids in two areas and two rivers in other areas. Four firefighters (agents) are initiated at the same burning area, illustrated in Figure 11(a). The agent has a local observation that contains a square view with  $5 \times 5$  grids and can move to one of four neighboring grids, spray water, or pump water. They share a water tank, which is initiated with four units water. Once the agent sprays water on the burning grid, it puts out the fire and consumes one unit water. Once the agent pumps water in the river, the water in the tank increases by one unit. The agents only get a global reward, *i.e.*, the number of the extinguished burning grids, at the final timestep.

As illustrated by Figure 11(b), MAAC and QMIX fall into local optimum. The agents learn to put out the fire around the initial position until the water is exhausted, and get the mean reward 4. They do not take different roles of pumping water or putting out the fire farther. Benefited from the emergent individuality, EOI+MAAC achieves the best performance with a clear division of labor, where two agents go to the river for pumping water, and two agents go to different areas for fighting with fire. Since higher rewards are hard to explore, we compare EOI with EDTI, a multi-



(a) Battle (b) 10\_vs\_10

Figure 12. Illustration of Battle and 10\_vs\_10.



(a) Battle (b) 10\_vs\_10

Figure 13. Learning curves in Battle and 10\_vs\_10.

agent exploration method. EDTI gives the agent an intrinsic motivation, considering both curiosity and influence. It rewards the agent to explore the rarely visited states and to maximize the influence on the expected returns of other agents. EDTI could escape from the local optimum, but it converges to a lower mean reward than EOI. This is because encouraging curiosity and influence cannot help division of labor, and the computational complexity makes the learning inefficient especially in the environments with many agents.

ROMA uses a role encoder taking the local observation to generate role embedding, and decodes the role to generate the parameters of individual action-value function in QMIX. However, ROMA converges slower than EOI+QMIX, since generating various parameters for the emergent roles is less efficient than encouraging diverse behaviors by reward shaping. Moreover, since the agents are initiated with the same observations, the role encoder will generate the similar role embeddings, which means similar agent behaviors at the beginning, bringing difficulty to the emergent roles.

HC first provides each agent a set of diverse skills, which are learned by DIAYN (before  $6 \times 10^3$  episodes), then learns a high-level policy to select the skill for each agent. Since the skills are trained independently without centralized coordination, they could be diverse but might not be correlated with the success of task. So it is hard to explore and learn high-performing policies. Moreover, HC requires a centralized controller in execution. EOI encourages individuality with the coordination provided by centralized training, considering both of the global reward and other agents' policies.

To investigate the effectiveness of EOI in large-scale com-

plex environments, we test EOI in the two scenarios: the MAgent (Zheng et al., 2018) task Battle and the gfootball (Kurach et al., 2020) task 10\_vs\_10, as illustrated in Figure 12.

- *Battle*. 20 agents (red) battle with 12 enemies (blue) which are controlled by built-in AI. The moving or attacking range of the agent is four neighbor grids. If an agent attacks an enemy, kills an enemy, and is killed by an enemy, it will respectively get a reward of +1, +10, and -2. The team reward is the sum of all agents' rewards.
- *10\_vs\_10*. 10 agents try to score facing 10 defenders which are controlled by built-in AI. If the agents shoot the ball into the goal, they will get a global reward of +1. If they lose control of the ball, they will get a global reward of -0.2.

Due to the large number of agents, we let the agents share the weights of neural network  $\theta$  and additionally feed a one-hot encoding of agent index into  $\theta$ . The classifier  $\phi$  does not use the index information to avoid overfitting. Parameter sharing might cause similar behaviors, which hinders the emergence of complex cooperative strategies. However, EOI helps a number of agents learn individualized policies even with shared parameters. As shown in Figure 13, EOI increases the behavior diversification and outperforms the baselines. Moreover, we find that the coefficient  $\alpha$  should be small in these environments where the success of the task is not always aligned with individuality. When individuality conflicts with success, focusing on individuality would negatively impact the maximization of environmental rewards. Hence we linearly decrease  $\alpha$  to zero during the training, which could accelerate the emergence of individuality at the early training and make the optimization objective unbiased at the later training. It is beneficial to explore the mechanism for adaptive  $\alpha$  to balance the success and individuality, *e.g.*, to adjust  $\alpha$  by meta gradient with the direction of the fastest increase of the environmental reward as in Lin et al. (2019).

#### 4.3. Similarity and Difference Between EOI and DIAYN

DIAYN is proposed to learn diverse skills with single-agent RL in the absence of any rewards. It trains the agent by maximizing the mutual information between skills ( $Z$ ) and states ( $S$ ), maximizing the policy entropy, and minimizing the mutual information between skills and actions ( $A$ ) given the state. The optimization objective is

$$\begin{aligned} & \text{MI}(S; Z) + \mathcal{H}(A|S) - \text{MI}(A; Z|S) \\ &= \mathcal{H}(Z) - \mathcal{H}(Z|S) + \mathcal{H}(A|S) - (\mathcal{H}(A|S) - \mathcal{H}(A|Z, S)) \\ &= \mathcal{H}(Z) - \mathcal{H}(Z|S) + \mathcal{H}(A|Z, S). \end{aligned}$$




Figure 14. Learning curves in sparse Pac-Men.

To maximize this objective, in practice DIAYN gives the learning agent an intrinsic reward  $\log q(z|s) - \log p(z)$ , where  $q(z|s)$  approximates  $p(z|s)$  and  $p(z)$  is a prior distribution. EOI gives each agent an intrinsic reward of  $p(i|o)$ . Let agents correspond to skills and observations correspond to states, then EOI has a similar intrinsic reward with DIAYN, though the motivations are distinct. Unlike from DIAYN, EOI employs two regularizers to capture the unique characteristics of MARL, strengthen the reward signals, and promote the emergence of individuality in MARL.

To empirically investigate the difference between EOI and the original DIAYN (making each skill as an agent), we test them in a sparse-reward version of Pac-Men, where the reward is defined as the minimum eaten dots of the four rooms. The results are shown in Figure 14. In this sparse-reward task, the agents could hardly obtain environmental reward before the individuality emerges. Without encouraging individuality, QMIX learns slowly. EOI+QMIX explores the environmental reward earlier and converges faster than DIAYN, indicating EOI has a stronger ability for individuality.

## 5. Conclusion and Discussion

We have proposed EOI, a novel method for the emergence of individuality in MARL. EOI learns a probabilistic classifier that predicts a probability distribution over agents given their observation and gives each agent an intrinsic reward of being correctly predicted by the classifier. Two regularizers are introduced to increase the discriminability of the classifier. We realized EOI on top of two popular MARL methods and empirically demonstrated that EOI outperforms existing methods in a variety of multi-agent cooperative tasks. We also discuss the similarity and difference between EOI and DIAYN.

However, EOI might be limited in some scenarios where the observation or trajectory cannot represent individuality. For example, if the agents could get the same full observation of the state, the probabilistic classifier cannot discriminate different agents based on the same observation. Moreover, if the local observation contains the identity information, *e.g.*, agent index, the probabilistic classifier would overfit the identity information and cannot help the emergence of individuality. We leave the limitations to future work.



## References

- Böhmer, W., Kurin, V., and Whiteson, S. Deep coordination graphs. In *International Conference on Machine Learning (ICML)*, 2020.
- Das, A., Gervet, T., Romoff, J., Batra, D., Parikh, D., Rabbat, M., and Pineau, J. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning (ICML)*, 2019.
- Dastani, M., Dignum, V., and Dignum, F. Role-assignment in open agent societies. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003.
- Ding, Z., Huang, T., and Lu, Z. Learning individually inferred communication for multi-agent cooperation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Du, Y., Han, L., Fang, M., Liu, J., Dai, T., and Tao, D. Liir: Learning individual intrinsic reward in multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations (ICLR)*, 2019.
- Foerster, J. N., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. Counterfactual multi-agent policy gradients. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- Freund, J., Brandmaier, A. M., Lewejohann, L., Kirste, I., Kritzler, M., Krüger, A., Sachser, N., Lindenberger, U., and Kempermann, G. Emergence of individuality in genetically identical mice. *Science*, 340(6133):756–759, 2013.
- Gordon, D. M. The organization of work in social insect colonies. *Nature*, 380(6570):121–124, 1996.
- Hughes, E., Leibo, J. Z., Phillips, M., Tuyls, K., Dueñez-Guzmán, E., Castañeda, A. G., Dunning, I., Zhu, T., McKee, K., Koster, R., et al. Inequity aversion improves cooperation in intertemporal social dilemmas. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Iqbal, S. and Sha, F. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2019.
- Jiang, J., Dun, C., Huang, T., and Lu, Z. Graph convolutional reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- Kurach, K., Raichuk, A., Stańczyk, P., Zajkac, M., Bachem, O., Espeholt, L., Riquelme, C., Vincent, D., Michalski, M., Bousquet, O., et al. Google research football: A novel reinforcement learning environment. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- Lee, Y., Yang, J., and Lim, J. J. Learning to coordinate manipulation skills via skill behavior diversification. In *International Conference on Learning Representations (ICLR)*, 2020.
- Li, C., WU, C., Wang, T., Yang, J., Zhao, Q., and Zhang, C. Celebrating diversity in shared multi-agent reinforcement learning. 2021a.
- Li, W., Wang, X., Jin, B., Sheng, J., Hua, Y., and Zha, H. Structured diversification emergence via reinforced organization control and hierarchical consensus learning. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2021b.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016.
- Lin, X., Baweja, H. S., Kantor, G., and Held, D. Adaptive auxiliary task weighting for reinforcement learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Liu, W., Camps, O., and Sznaier, M. Multi-camera multi-object tracking. *arXiv preprint arXiv:1709.07065*, 2017.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Macarthur, K. S., Stranders, R., Ramchurn, S., and Jennings, N. A distributed anytime algorithm for dynamic task allocation in multi-agent systems. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2011.
- McKee, K., Gemp, I., McWilliams, B., Duñez-Guzmán, E. A., Hughes, E., and Leibo, J. Z. Social diversity and social preferences in mixed-motive reinforcement learning. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020.
- Oliehoek, F. A., Amato, C., et al. *A concise introduction to decentralized POMDPs*, volume 1. Springer, 2016.
- Rashid, T., Samvelyan, M., De Witt, C. S., Farquhar, G., Foerster, J., and Whiteson, S. Qmix: monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2018.

- Sander, P. V., Peleshchuk, D., and Grosz, B. J. A scalable, distributed algorithm for efficient task allocation. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2002.
- Schroff, F., Kalenichenko, D., and Philbin, J. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- Shu, T. and Tian, Y. M<sup>3</sup>rl: Mind-aware multi-agent management reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2019.
- Sims, M., Corkill, D., and Lesser, V. Automated organization design for multi-agent systems. *Autonomous agents and multi-agent systems*, 16(2):151–185, 2008.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2018.
- Wang, J., Ren, Z., Liu, T., Yu, Y., and Zhang, C. Qplex: Duplex dueling multi-agent q-learning. In *International Conference on Learning Representations (ICLR)*, 2021a.
- Wang, T., Dong, H., Lesser, V., and Zhang, C. Multi-agent reinforcement learning with emergent roles. *arXiv preprint arXiv:2003.08039*, 2020a.
- Wang, T., Wang, J., Wu, Y., and Zhang, C. Influence-based multi-agent exploration. In *International Conference on Learning Representations (ICLR)*, 2020b.
- Wang, T., Gupta, T., Mahajan, A., Peng, B., Whiteson, S., and Zhang, C. Rode: Learning roles to decompose multi-agent tasks. 2021b.
- Wang, Y., Han, B., Wang, T., Dong, H., and Zhang, C. Dop: Off-policy multi-agent decomposed policy gradients. In *International Conference on Learning Representations (ICLR)*, 2021c.
- Yang, J., Borovikov, I., and Zha, H. Hierarchical cooperative multi-agent reinforcement learning with skill discovery. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020a.
- Yang, Y., Hao, J., Liao, B., Shao, K., Chen, G., Liu, W., and Tang, H. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939*, 2020b.
- Zheng, L., Yang, J., Cai, H., Zhou, M., Zhang, W., Wang, J., and Yu, Y. Magent: A many-agent reinforcement learning platform for artificial collective intelligence. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.