

Hierarchical Reinforcement Learning for Multi-agent MOBA Game

Zhijian Zhang, Haozheng Li, Luo Zhang, Tianyin Zheng, Ting Zhang,
Xiong Hao, Xiaoxin Chen, Min Chen, Fangxu Xiao, Wei Zhou

vivo AI Lab

{zhijian.zhang, haozheng.li, luo.zhang, zhengtianyin, haoxiong}@vivo.com

Abstract

Real Time Strategy (RTS) games require macro strategies as well as micro strategies to obtain satisfactory performance, since it has large state space, action space, and hidden information. This paper presents a novel hierarchical reinforcement learning model for mastering Multiplayer Online Battle Arena (MOBA) games, a sub-genre of RTS games. The contributions are: (1) proposing a hierarchical framework, where agents execute macro strategies by imitation learning and carry out micromanipulations through reinforcement learning, (2) developing a simple self-learning method to get better sample efficiency for training, and (3) designing a dense reward function for multi-agent cooperation in the absence of game engine or Application Programming Interface (API). Finally, various experiments have been performed to validate the superior performance of the proposed method over other state-of-the-art reinforcement learning algorithms. Agents successfully learn to combat and defeat bronze-level built-in AI with 100% win rate, and experiments show that our method can create a competitive multi-agent for a kind of mobile MOBA game *King of Glory* in 5v5 mode.

1 Introduction

Deep reinforcement learning (DRL) has become a promising tool for game AI since its success in playing game Atari [Mnih *et al.*, 2015], AlphaGo [Silver *et al.*, 2017], Dota 2 [OpenAI, 2018], and so on. Researchers verify algorithms by conducting experiments in games quickly, and transfer this ability to real world applications such as robotics control, recommend services. Despite of initial success, there are still many challenges in practice. More and more researchers have started to conquer more complex Real Time Strategy (RTS) games such as StarCraft and Defense of the Ancients (Dota) recently. Dota is a kind of Multiplayer Online Battle Arena (MOBA) game which includes 5v5 and 1v1 modes. To achieve victory in a MOBA game, players need to control their only one agent to destroy enemies' crystal.

MOBA games take up more than 30% of the online gameplays all over the world, including League of Legends, Dota,



(a) 5v5 map

(b) 1v1 map

Figure 1: (a) Screenshot from 5v5 map of *KOG*. Players can get the position of allies, towers, enemies in view and know whether jungles alive or not from mini-map. From the screen, players can observe surrounding information including what kind of skills are released and releasing. (b) Screenshot from 1v1 map of *KOG*, known as solo mode.

King of Glory (*KOG*), and others [Murphy, 2015]. Fig. 1a shows a 5v5 map of *KOG*, where players control the motion of heroes by controlling the left bottom steer button, while using skills by controlling right bottom set of buttons. The upper-left corner shows mini-map, with the blue markers indicating own towers and the red markers indicating the enemies' towers. Each player can obtain gold and experience by killing enemies, jungling and destroying towers. The ultimate goal of this game is to destroy enemies' crystal. As shown in Fig. 1b, there are two players in 1v1 map.

Compared with Atari, the main challenges of MOBA games for us are: (1) **Game engine or Application Programming Interface (API) is not available for us.** We need to extract features by multi-target detection, and run the game on mobile phones, which is restricted by low computational power. However, the computational complexity can be up to $10^{20,000}$, while AlphaGo is about 10^{250} [OpenAI, 2018]. (2) **Rewards are severely delayed and sparse.** The ultimate goal of the game is to destroy the enemies' crystal, which means that rewards are seriously delayed. Meanwhile, the rewards are really sparse if we set the rewards of $-1/1$ according to the final result loss/win. (3) **Multi-agents' communication and cooperation are challenging.** Communication and cooperation are crucially important for RTS games especially in 5v5 mode.

To the best of our knowledge, this paper is the first attempt to propose reinforcement learning in MOBA game, which does not obtain information from API, but captures information from game video directly. We cope with the curse of computational complexity through imitation learning of the

macro strategies, and it is hard to train for agents because the rewards in this part are severely delayed and sparse. Meanwhile, we develop a distributed platform for sampling to accelerate the training process and combine the A-Star path planning algorithm to do navigation. To test the performance of our method, we take the game *KOG*, a popular mobile MOBA game, as our experiment environment, and systematic experiments have been performed.

The main contributions of this work include: (1) proposing a novel hierarchical reinforcement learning framework for a kind of mobile MOBA game *KOG*, which combines imitation learning and reinforcement learning. Imitation learning is responsible for macro strategies such as where to go to, when to offend and defend, while reinforcement learning is in charge of micromanipulations such as which skill to release and how to move in battle; (2) developing a simple self-learning method which learns to compete with agent's past good decisions and come up with an optimal policy to accelerate the training process; (3) developing a multi-target detection method to extract global features composing the state of reinforcement learning; (4) designing a dense reward function and using real-time data to communicate with each other. Experiments show that our agents learn better policy than other reinforcement learning methods.

2 Related Work

2.1 RTS Games

There has been a history of studies on RTS games such as StarCraft [Ontanón *et al.*, 2013] and Dota [OpenAI, 2018]. One practical way using rule-based method by bot SAIDA achieved championship on SSCAIT recently. Based on the experience of the game, rule-based bots can only choose pre-defined action and policy at the beginning of a game, which is insufficient to deal with large and real time state space throughout the game, and it is difficult to keep learning and evolving. Dota2 AI created by OpenAI, named OpenAI Five, has made great success by using Proximal Policy Optimization (PPO) algorithm together with well-designed rewards. However, OpenAI Five has used huge computing resources due to lacking of macro strategy.

Related work has also been done in macro strategies by Tencent AI Lab in game *KOG* [Wu *et al.*, 2018], and their 5-AI team achieved 48% win rate against human player teams which are ranked top 1% in the player ranking system. However, 5-AI team used supervised learning and the training data can be obtained from game replays processed by game engine and API, which run on server. This method is not possible for us because we don't have access to the game engine or API, and we need to run on mobile phones.

2.2 Hierarchical Reinforcement Learning

Traditional reinforcement learning methods such as Q-learning or Deep Q Network (DQN) is difficult to manage due to large state space in environment, Hierarchical reinforcement learning [Barto and Mahadevan, 2003] tackles this kind of problem by decomposing a high dimensional target into several sub-target which is easier to cope with.

Hierarchical reinforcement learning has been explored in different scenarios. As for games, somewhat related to our hierarchical architecture exists [Sun *et al.*, 2018], which takes advantage of prior knowledge of a game to design macro strategies, and there is neither imitation learning nor experienced experts' guidance. There have been many novel hierarchical reinforcement learning algorithms proposed recently. One approach of combining meta-learning with a hierarchical learning is Meta Learning Shared Hierarchies (MLSH) [Frans *et al.*, 2017], which is mainly used in multi-task learning and transfer learning. Hierarchically Guided Imitation Learning/Reinforcement Learning have been showed effective in speeding up learning [Le *et al.*, 2018], but it needs high-level expert guidance in micro strategies which is hard to design for RTS games.

2.3 Multi-agent Reinforcement Learning in Games

Multi-agent reinforcement learning(MARL) has certain advantages over single agent learning. Different agents can complete tasks faster and better through knowledge sharing, and there are some challenges as well. For example, the computational complexity increases due to larger state space and action space compared with single agent learning. Because of the above challenges, MARL mainly focuses on stability and adaption.

Simple applications of reinforcement learning to MARL is limited, such as no communication and cooperation among agents [Sukhbaatar *et al.*, 2016], lacking of global rewards [Rashid *et al.*, 2018], and failure to consider enemies' strategies when learning policy. Some recent studies relevant to this challenge have been investigated. [Foerster *et al.*, 2017] introduced a concentrated criticism of the cooperative settings with shared rewards. The approach interprets experiences in the replay memory as off-environment data and marginalize the action of a single agent while keeping others unchanged. These methods enable successful combination of experience replay with multi-agent. Similarly, [Jiang and Lu, 2018] proposed an attentional communication model based on actor-critic algorithm for MARL, which learns to communicate and share information when making decision. Therefore, this approach can be a complement for the proposed research. Parameter sharing multi-agent gradient descent Sarsa (PS-MASGDS) algorithm [Shao *et al.*, 2018] used a neural network to estimate the value function and proposed a reward function to balance the units move and attack in the game of StarCraft, which can be learned from. However, these methods require a lot of computing resources.

3 Methods

This section introduces hierarchical architecture, state representation and action definition. Then the network architecture and training algorithm are presented. The reward function design and self-learning method are discussed lastly.

3.1 Hierarchical Architecture

The hierarchical architecture is shown in Fig. 2. There are four types of macro actions including attack, move, purchase and adding skill points, which are selected by imitation learning. Then reinforcement learning algorithm chooses specific

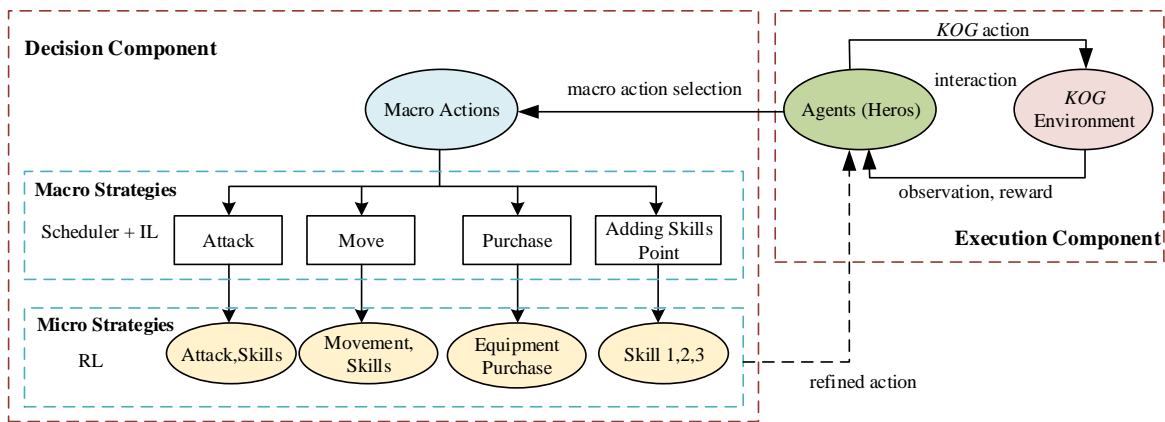


Figure 2: Hierarchical architecture

States	Dimensionality	Type
Extracted features	116	R
Mini-map information	$64 \times 64 \times 3$	R
Current view information	$84 \times 42 \times 3$	R
Action_A	7	one-hot
Action_M	9	one-hot

Table 1: The dimension and data type of our states and action

action a according to policy π for making micro strategies in state s . The encoded action is performed and the agent can get reward r and next observation s' from *KOG* environment. $R_\pi = \sum_{t=0}^T \gamma^t r_t$ is defined as the discounted return, where $\gamma \in [0,1]$ is a discount factor. The aim of agents is to learn a policy that maximizes the expected discounted returns, defined as:

$$J = E_\pi[R_\pi] \quad (1)$$

The Scheduler module designed by observation from game video information is responsible for switching between reinforcement and imitation learning. It is also possible to replace the imitation learning part with high-level expert system for the fact that the data in imitation learning model is produced by high-level expert guidance.

3.2 State Representation and Action Definition

State Representation

It is an open problem on how to represent the state of RTS games optimally. This paper constructs a state representation as inputs to neural network from features extracted by multi-target detection, mini-map information of the game, and current view of the agent, which have different dimensions and data types, as illustrated in Table 1. Current view information is RGB image in the view of the agent, and mini-map information is from RGB image in the upper-left corner of the screenshot.

Extracted features include the position of all heroes, towers, and soldiers, blood volume, gold that the player have and skills released by heroes in the current view, as shown in Fig. 3. All the extracted features are embedded to a 116-dimensional vector. The inputs at current step are composed

of current state information, the last step information, and the last action which has been shown to be useful for the learning process in reinforcement learning. States with real value are normalized to $[0, 1]$.

Action Definition

In this game, we define the action into two parts including Action_M and Action_A. The motion movement Action_M includes Up, Down, Left, Right, Lower-right, Lower-left, Upper-right, Upper-left, and Stay still. When the selected action is attack Action_A, it can be Stay still, Skill-1, Skill-2, Skill-3, Attack, and summoned skills including Flash and Restore. Meanwhile, it is our first choice to attack the weakest enemy when action attack is available for each agent.

3.3 Network Architecture and Training Algorithm

Network Architecture

Table reinforcement learning such as Q-learning has limitations in large state space situations. To tackle this problem, the micro level algorithm design is similar to proximal policy optimization (PPO) algorithm [Schulman *et al.*, 2017]. Inputs of convolutional network are current view and mini-map information with a shape of $84 \times 42 \times 3$ and $64 \times 64 \times 3$ respectively. Meanwhile, the extracted features consist of a 116-dimensional vector. We use the rectified linear unit (ReLU) activation function in the hidden layer. The output layer's activation function is softmax function, which outputs the probability of each action. Our model in game *KOG*, including inputs and architecture of the network, and output of actions, is depicted in Fig. 3.

Training Algorithm

This paper proposes a Hierarchical Reinforcement Learning (HRL) algorithm for multi-agent learning, and the training process is presented in Algorithm 1. Firstly, we initialize our controller policy and global state. Then each agent takes a move and attack action pair $[a_t^m, a_t^a]$ and receive reward r_{t+1} and next state s_{t+1} . The agent can obtain both macro action through imitation learning and micro action from reinforcement learning from state s_{t+1} . The action probability likelihood is normalized to choose action $[a_{t+1}^m, a_{t+1}^a]$ from macro action A_{t+1} . At the end of each iteration, we use the ex-

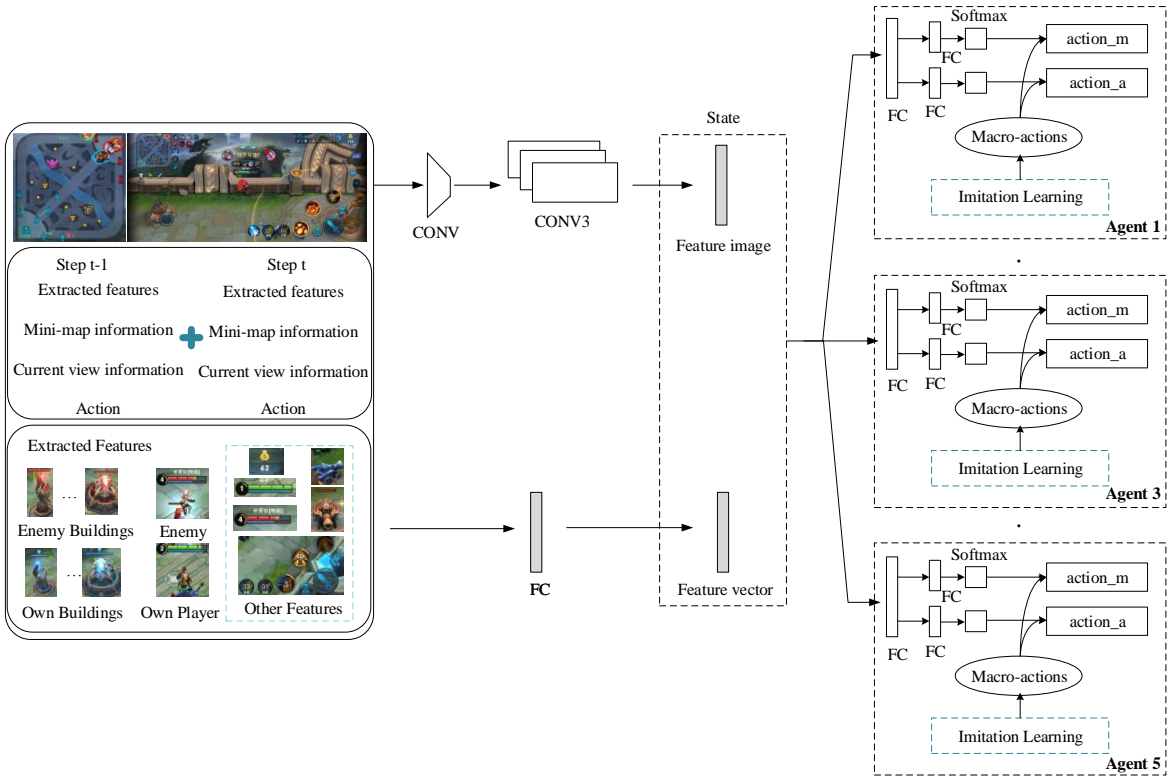


Figure 3: Network architecture of hierarchical reinforcement learning model

perience replay samples to update parameters of the policy network.

We take the loss of entropy and self-learning into account to encourage exploration in order to balance the trade-off between exploration and exploitation. Loss formula is defined as:

$$L_t^M(\theta) = E_t[w_1 L_t^v(\theta) + w_2 N_t^M(\pi, a_t) + L_t^{Mp}(\theta) + w_3 S_t^M(\pi, a_t)] \quad (2)$$

$$L_t^A(\theta) = E_t[w_1 L_t^v(\theta) + w_2 N_t^A(\pi, a_t) + L_t^{Ap}(\theta) + w_3 S_t^A(\pi, a_t)] \quad (3)$$

$$L_t(\theta) = L_t^M(\theta) + L_t^A(\theta) \quad (4)$$

where $L_t^M(\theta)$ is the loss of action move, $L_t^A(\theta)$ is the loss of action attack. w_1, w_2, w_3 are the weights of value loss, entropy loss and self-learning loss that we need to tune, N_t^M denotes the entropy loss of action move, N_t^A denotes the entropy loss of action attack, S_t^M means the self-learning loss of action move, and S_t^A means the self-learning loss of action attack. Total loss $L_t(\theta)$ is composed of the loss of move and attack for simply computation.

Value loss $L_t^v(\theta)$, policy loss $L_t^{Mp}(\theta)$, and policy loss $L_t^{Ap}(\theta)$ are defined as follows:

$$L_t^v(\theta) = E_t[(r(s_t, a_t) + V_t(s_t) - V_t(s_{t+1}))^2] \quad (5)$$

$$L_t^{Mp}(\theta) = E_t[\min(r_t(\theta)D_t^M, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon)D_t^M)] \quad (6)$$

$$L_t^{Ap}(\theta) = E_t[\min(r_t(\theta)D_t^A, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon)D_t^A)] \quad (7)$$

where $r_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{old}}(a_t|s_t)$, D_t^M and D_t^A are advantage of action move and action attack computed by the difference between return and value estimation.

3.4 Reward Design and Self-learning

Reward Design

Reward function plays a significant role in reinforcement learning, and good learning results of an agent are mainly depending on diverse rewards. The ultimate goal of the game is to destroy the enemies' crystal. If our reward is only based on the final result, it will be extremely sparse, and the seriously delayed reward leads to slow learning speed. Dense reward gives quick positive or negative feedback to the agent, and can help the agents to learn faster and better. Damage amount of an agent is not available for us since we don't have game engine or API. In our experiment, all agents can receive two parts of rewards including self-reward and global-reward. Self-reward consists of gold and Health Points (HP) loss/gain of the agent, while global-reward includes tower loss and death of allies and enemies.

$$\begin{aligned} r_t &= \rho_1 \times r_{self} + \rho_2 \times r_{global} \\ &= \rho_1((gold_t - gold_{t-1})f_m + (HP_t - HP_{t-1})f_H) \\ &\quad + \rho_2(tower_{loss_t} \times f_t + player_{death_t} \times f_d) \end{aligned} \quad (8)$$

where $tower_{loss_t}$ is positive when enemies' tower is destroyed, and is negative when own tower is destroyed, the

Algorithm 1 Hierarchical RL Training Algorithm

Input: Reward function R_n , max episodes M , function $IL(s)$ indicates imitation learning model.

Output: Hierarchical reinforcement learning neural network.

```
1: Initialize controller policy  $\pi$ , global state  $s_g$  shared
   among our agents;
2: for  $episode = 1, 2, \dots, M$  do
3:   Initialize  $s_t, a_t^m, a_t^a$ ;
4:   repeat
5:     Take action  $[a_t^m, a_t^a]$ , receive reward  $r_{t+1}$ , next state
      $s_{t+1}$ , where  $a_t^m$  indicates a motion movement, and
      $a_t^a$  indicates a motion attack;
6:     Choose macro action  $\mathbf{A}_{t+1}$  from  $s_{t+1}$  according to
      $IL(s = s_{t+1})$ ;
7:     Choose micro action  $[a_{t+1}^m, a_{t+1}^a]$  from  $\mathbf{A}_{t+1}$  ac-
     cording to the output of RL in state  $s_{t+1}$ ;
8:     if  $a_{t+1}^i \notin \mathbf{A}_{t+1}$ , where  $i = 0, \dots, 16$  then
9:        $P(a_{t+1}^i | s_{t+1}) = 0$ ;
10:    else
11:       $P(a_{t+1}^i | s_{t+1}) = P(a_{t+1}^i | s_{t+1}) / \sum_j P(a_{t+1}^j | s_{t+1})$ ;
12:    end if
13:    Collect samples  $(s_t, a_t^m, a_t^a, r_{t+1})$ ;
14:    Update policy parameter  $\theta$  to maximize the expected
    returns;
15:  until  $s_t$  is terminal
16: end for
```

same as $player_{death}$, f_m is a coefficient of gold loss, the same as f_H , f_t and f_d , ρ_1 is the weight of self-reward and ρ_2 means the weight of global-reward. The reward function is effective for training, and the results are shown in the experiment section.

Self-learning

There are many kinds of self-learning methods for reinforcement learning such as Self-Imitation Learning (SIL) [Oh *et al.*, 2018] and Episodic Memory Deep Q-Networks (EMDQN) [Lin *et al.*, 2018]. SIL is applicable to actor-critic architecture, while EMDQN combines episodic memory with DQN. However, considering better sample efficiency and easier-to-tune of the system, the proposed method migrates EMDQN to reinforcement learning algorithm PPO [Schulman *et al.*, 2017]. Loss of self-learning part can be defined as follows:

$$S_t(\pi, a_t) = E_t[(V_{t+1} - V_H)^2] + E_t[\min(r_t(\theta)A_{H_t}, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon)A_{H_t})] \quad (9)$$

where the memory target V_H is the best value from memory buffer, and A_{H_t} means the best advantage from it.

$$V_H = \begin{cases} \max((\max_{R_i(s_t, a_t)}), R(s_t, a_t)), & \text{if } (s_t, a_t) \in \text{memory} \\ R(s_t, a_t), & \text{otherwise} \end{cases} \quad (10)$$

$$A_{H_t} = V_H - V_{t+1}(s_{t+1}) \quad (11)$$

Category	Training Set	Testing Set	Precision
Own Soldier	2677	382	0.6158
Enemy Solider	2433	380	0.6540
Own Tower	485	79	0.9062
Enemy Tower	442	76	0.9091
Own Crystal	95	17	0.9902
Enemy Crystal	152	32	0.8425

Table 2: The accuracy of multi-target detection

Scenarios	AI. 1	AI. 2	AI. 3	AI. 4
1v1 mode	80%	/	52%	58%
5v5 mode	82%	68%	66%	60%

Table 3: Win rates playing against AI. 1: AI without macro strategy, AI. 2: without multi-agent, AI. 3: without global reward and AI. 4: without self-learning method

where $i \in [1, 2, \dots, E]$, E represents the number of episodes in memory buffer that the agent has experienced.

4 Experiments

The experiment setup will be introduced first. Then we evaluate the performance of our algorithms on two environments: (i) 1v1 map including entry-level, easy-level and medium-level built-in AI, and (ii) a challenging 5v5 map. We analyze the average rewards and win rates during training.

4.1 Setup

The experiment setup includes experiment platform and GPU cluster training platform. In order to increase the diversity and quantity of samples, we use 10 mobile phones for an agent to collect the distributed data. Meanwhile, we need to maintain the consistency of all the distributed mobile phones when training. We transmit the collected sample of all agents to the server and do a centralized training, and share the parameters of network among all agents. Each agent executes its policy based on its own states. As for the features obtained by multi-target detection, its accuracy and category are depicted in Table 2, which is adequate for our learning process. Moreover, the speed of taking an action is about 150 Actions Per Minute (APM), comparable to 180 APM of high level player. A-star path planning algorithm is applied when going to someplace. Parameters of w_1 , w_2 , and w_3 are set to 0.5, -0.01, and 0.1 respectively based on preliminary results. The training time is about seven days for agents on one Tesla P40 GPU.

4.2 1v1 mode of game KOG

As shown in Fig. 1(b), there are one agent and one enemy player in 1v1 map. We need to destroy the enemies' tower first and then destroy the crystal to get the final victory. We draw the win rates and average rewards when agent fights with different level of built-in AI.

Win Rates

The results when our AI plays against AI without macro-strategy, without multi-agent, without global reward and




Figure 4: The average rewards of our agent in 1v1 mode during training.

without self-learning method are listed in Table 3. 50 games are played against AI. 1, AI. 3 and AI. 4, and the win rates are 80%, 52% and 58% respectively. We have tested about 200 games for each level of built-in AI, and listed the win rates for algorithm PPO and HRL in 1v1 mode against different level of built-in AI, as shown in Table 4.

Average Rewards

Generally speaking, the target of our agent is to defeat the enemies as soon as possible. Fig. 4 illustrates the average rewards of our agent Angela in 1v1 mode when combatting with different enemies. In the beginning, the rewards are low because the agent is still a beginner and doesn't have enough learning experience. However, our agent is learning gradually and being more and more experienced. When the training episodes of our agent reach about 100, the rewards in each step become positive overall and our agent starts to have some advantages in battle. There are also some decreases in rewards when facing high level built-in AI because of the fact that the agent is unable to defeat the Warrior at first. To sum up, the average rewards are increasing obviously, and stay smooth after about 600 episodes.

4.3 5v5 mode of game KOG

As shown in Fig. 1(a), there are five agents and five enemy players in 5v5 map. What we need to do actually is to destroy the enemies' crystal. In this scenario, we train our agents with built-in AI, and each agent holds one model. In order to analyze the results during training, we illustrate the win rates in Fig. 5.

Win Rates

We have plotted the win rates in Fig. 5. there are three different levels of built-in AI that our agents combat with. When fighting with bronze-level built-in AI, agents learn fast and the win rates reach 100%. When training with gold-level built-in AI, the learning process is slow and agents can't win until 100 episodes. In this mode, the win rates are about 40% in the end. This is likely due to the fact that our agents can hardly obtain dense global rewards when playing against high level AI, which leads to hard cooperation in team battle. One way using supervised learning method from Tencent AI Lab obtains 100% win rate [Wu *et al.*, 2018]. However, the method used about 300 thousand game replays with the

Algorithm	Entry-level	Easy-level	Medium-level
PPO	62%	60%	55%
HRL	89%	83%	80%

Table 4: Win rates for HRL and PPO in 1v1 mode against different level of built-in AI.




Figure 5: Win rates of our agents in 5v5 mode against different level of built-in AI.

advantage of API. Another way is using PPO algorithm without macro strategy, which achieves about 22% win rate when combatting with gold-level built-in AI. Meanwhile, the results of our AI playing against AI without macro strategy, without multi-agent, without global reward and without self-learning method are listed in Table 3. These indicate the importance of each method in our hierarchical reinforcement learning algorithm.

5 Conclusion

This paper proposed a novel hierarchical reinforcement learning framework for multi-agent MOBA game KOG, which learns macro strategies through imitation learning and micro actions by reinforcement learning. In order to obtain better sample efficiency, we presented a simple self-learning method, and extracted global features as a part of state input by multi-target detection. We performed systematic experiments both in 1v1 mode and 5v5 mode, and compared our method with PPO algorithm. Our results showed that this hierarchical reinforcement learning framework is encouraging for the MOBA game.

In the future, we will explore how to combine graph network with our method for multi-agent collaboration

Acknowledgments

We would like to thank two anonymous reviewers for their insightful comments and our colleagues, particularly Dr. Yang Wang, Dr. Hao Wang, Jingwei Zhao, and Guozhi Wang, for extensive discussion and suggestion. We are also very grateful for the support from vivo AI Lab.

References

- [Barto and Mahadevan, 2003] Andrew G Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1-2):41–77, 2003.
- [Foerster *et al.*, 2017] Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip HS Torr, Pushmeet Kohli, and Shimon Whiteson. Stabilising experience replay for deep multi-agent reinforcement learning. *arXiv preprint arXiv:1702.08887*, 2017.
- [Frans *et al.*, 2017] Kevin Frans, Jonathan Ho, Xi Chen, Pieter Abbeel, and John Schulman. Meta learning shared hierarchies. *arXiv preprint arXiv:1710.09767*, 2017.
- [Jiang and Lu, 2018] Jiechuan Jiang and Zongqing Lu. Learning attentional communication for multi-agent cooperation. *arXiv preprint arXiv:1805.07733*, 2018.
- [Le *et al.*, 2018] Hoang M Le, Nan Jiang, Alekh Agarwal, Miroslav Dudík, Yisong Yue, and Hal Daumé III. Hierarchical imitation and reinforcement learning. *arXiv preprint arXiv:1803.00590*, 2018.
- [Lin *et al.*, 2018] Zichuan Lin, Tianqi Zhao, Guangwen Yang, and Lintao Zhang. Episodic memory deep q-networks. *arXiv preprint arXiv:1805.07603*, 2018.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [Murphy, 2015] M Murphy. Most played games: November 2015–fallout 4 and black ops iii arise while starcraft ii shines, 2015.
- [Oh *et al.*, 2018] Junhyuk Oh, Yijie Guo, Satinder Singh, and Honglak Lee. Self-imitation learning. *arXiv preprint arXiv:1806.05635*, 2018.
- [Ontanón *et al.*, 2013] Santiago Ontanón, Gabriel Synnaeve, Alberto Uriarte, Florian Richoux, David Churchill, and Mike Preuss. A survey of real-time strategy game ai research and competition in starcraft. *IEEE Transactions on Computational Intelligence and AI in games*, 5(4):293–311, 2013.
- [OpenAI, 2018] OpenAI. Openai five, 2018. <https://blog.openai.com/openai-five/>, 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 preprint arXiv:1803.11485*, 2018.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Shao *et al.*, 2018] Kun Shao, Yuanheng Zhu, and Dongbin Zhao. Starcraft micromanagement with reinforcement learning and curriculum transfer learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2018.
- [Silver *et al.*, 2017] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [Sukhbaatar *et al.*, 2016] Sainbayar Sukhbaatar, Rob Fergus, et al. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, pages 2244–2252, 2016.
- [Sun *et al.*, 2018] Peng Sun, Xinghai Sun, Lei Han, Jiechao Xiong, Qing Wang, Bo Li, Yang Zheng, Ji Liu, Yongsheng Liu, Han Liu, et al. Tstarbots: Defeating the cheating level builtin ai in starcraft ii in the full game. *arXiv preprint arXiv:1809.07193*, 2018.
- [Wu *et al.*, 2018] Bin Wu, Qiang Fu, Jing Liang, Peng Qu, Xiaoqian Li, Liang Wang, Wei Liu, Wei Yang, and Yongsheng Liu. Hierarchical macro strategy model for moba game ai. *arXiv preprint arXiv:1812.07887*, 2018.