Reinforcement Learning for the Beginning of Starcraft II Game

Yukang Chen
1155154501
Department of Computer Science
The Chinese University of Hong Kong
yukangchen@cse.cuhk.edu.hk

Chu Ruihang
1155156766
Department of Computer Science
The Chinese University of Hong Kong
rhchu@cse.cuhk.edu.hk

Abstract

Starcraft II is a popular real-time strategy game that is welcomed by many young people. This game is really complicated to play due to the long game timeline, various society/buildings/units, large amount of actions or selections, different constraints (*e.g.*, population limit) and the partially observed environments. In this project, we plan to develop a reinforcement learning model for the beginning of Starcraft II game, instead of the full-length game. The beginning of the game is essential for the further economy, population increase, and technology development. Our project is based on the SC2LE (StarCraft II Learning Environment) platform. We build a feasible pipeline for training reinforcement learning models and design random, scripted and our actor critic based agents. Experiments show that our actor critic based agents are able to learn valuable knowledge in this task. The video has been publicly available. [1]

1 Introduction

In recent years, reinforcement learning (RL) has been widely used in computer vision, natural language processing, robotics Levine et al. [2018] and so on. In terms of games, RL achieves remarkable success in Atari games Mnih et al. [2015b] and the Go game Silver et al. [2016]. Based on these achievements, researchers gradually pay attentions on the more complicated games, *e.g.* the real-time strategy (RTS) video games Starcraft II. In RTS games, players make their actions simultaneously, instead of acting in turns. This kind of games are more difficult to play well. It costs common teenagers several years to become a master in the Starcraft II. Therefore, it is also uneasy to train an agent to play this game. We list the difficulties of this problem as the following:

- **Large action space**: If the agent uses the point-and-click interface just like human players, there would be hundreds of millions of possible actions Vinyals et al. [2017] in each frame, for the high game resolution. In addition, with the game time going, players will possess different units, buildings and technologies, which involves additional unique actions.
- **Indistinct reward**: The most clear reward is the results of the game (wining or losing). However, as each game time involves thousands of frames and actions in the sequence, it is too sparse to use the final reward as instruction. What's more, apart from the micro-level actions, macro-level plannings are important for the results of the game. It is a hard job to design an approximate reward function that is neither too sparse or shortsighted.
- **Partially observation**: The environment is partially observed for two reasons. First, the map camera is locally posed. Players need to move the view to get more information. Second, there is a "fog-of-war" to cover unexplored area in the map. Players need to assign units to find out the positions and activities of their enemies.

[1] https://drive.google.com/file/d/11S6t3rNKjM1CJEkFSN1lkZiyLTdqHhzQ/view?usp=sharing

Figure 1: Illustration of units and buildings involved in the begining of this game.

In this project, we focus on the game beginning, which is essential for the full game and is also a complicated task. The beginning of the game is to set up the necessary units and buildings in a reasonable order and do their right jobs. It has less difficulty compared to the full game but is the essential task. Figure 1 1 shows an example layout of the beginning of the games with units and buildings noted. We design actor critic agents to learn how to collect economy and build marines and compare them to random and scripted agents. Our experiments are based on SC2LE, a StarScraft II Learning Environment platform, which is developed by the collaboration of DeepMind and Blizzard. It provides rich environments and interfaces for our agents to interact with StarCraft II. In Section 2, we introduce related works. In Section 3, we introduce our baseline agents and our RL based agents. We show the experiments in details in Section 4.2 and make a conclusion in Section 5.

2 Related Work

Reinforcement learning (RL) is studied to be suitable for playing real-time strategy games, achieving great advancements in classic board games Campbell et al. [2002], Silver et al. [2016] and video games Mnih et al. [2015a], Pathak et al. [2017], Jaderberg et al. [2019]. Containing much theoretic and domain complexities, RL's playing Starcraft II gains increasing research attention. To facilitate it, DeepMind developed SC2LE Vinyals et al. [2017], a RL environment for StarCraft II. It standardizes critical RL components such as the observation, action, and reward for Starcraft II domain, based on which various approaches could be fairly verified.

With a long-term measurable goal of beating human players, recent works Usunier et al. [2016], Shao et al. [2018], Campbell et al. [2002] explores RL algorithms and architectures to improve win rates, with hand-crafted sub-systems Churchill et al. [2017], high-level actionsSun et al. [2018], rule-based systems Gehring et al. [2018] and so on. Many works Sun et al. [2018], Lee et al. [2018], Pang et al. [2019] also described the hierarchical reinforcement learning paradigm Kulkarni et al. [2016], Vezhnevets et al. [2017], Nachum et al. [2018]. Specifically, they generally operate on multi-level abstraction, each handled by learnt strategies and controllers. Owing to its hierarchical architecture, the long-horizon task is decomposed into a series of sub-policies, along with it the huge action could be substantially reduced. As a milestone, AlphaStar Vinyals et al. [2019] exploited a model-free learning method that uses data from human and agent game, as well as avoiding the imperfect models, which is typical of search-based method. Finally, it reached the grandmaster level in this game.

Due to the severe difficulty of the full task of Starcraft II, *i.e.,* win the game, Blizzard and DeepMind designed seven minigames Vinyals et al. [2017], Hu and Kaneko [2019] as well. Each can be regarded

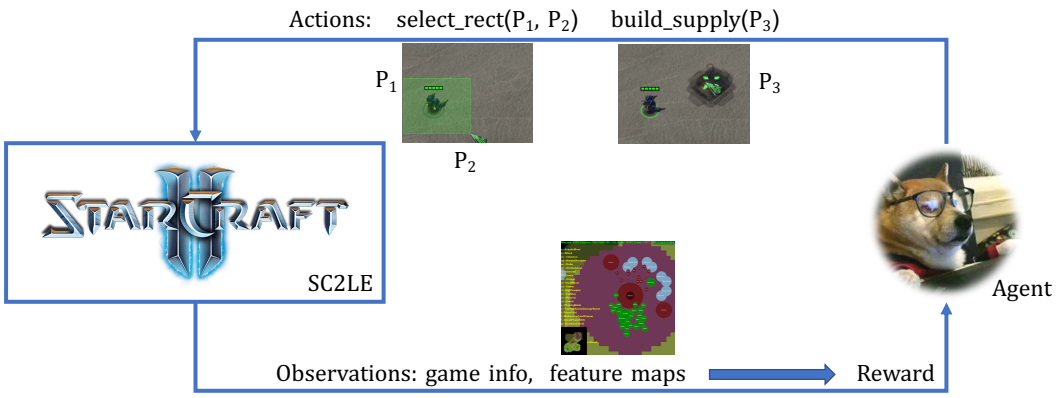

Figure 2: The interaction of the Starcraft II learning environment (SC2LE) platform and the agent.

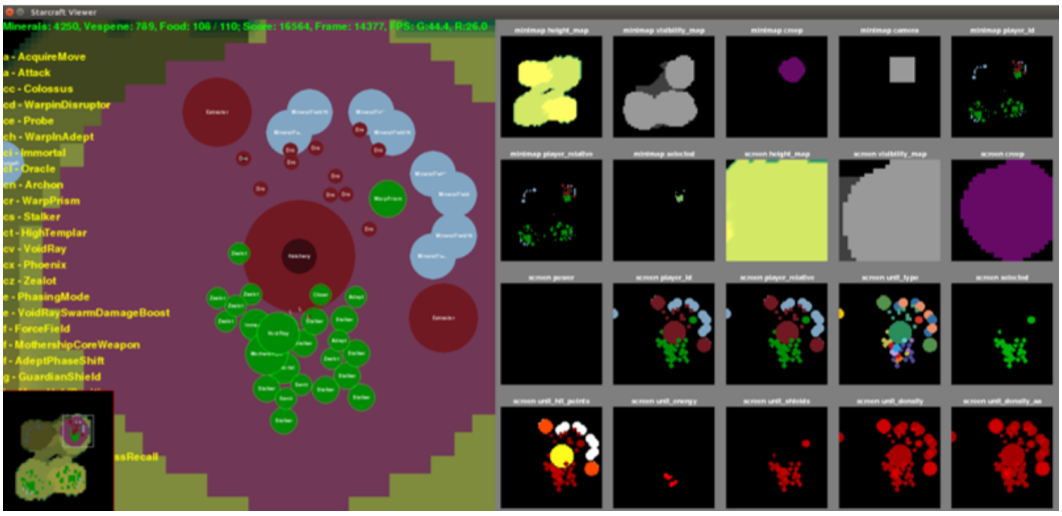

Figure 3: Processed input feature map in SC2LE platform.

as a sub-task and focuses on different perspective of the game, such as building base and collecting resource. BuildMarines, as one of mini games, encourages the agent to build as many Marines as possible within the fixed amount of time. Some earlier works Ring et al. [2018], Vinyals et al. [2017] searched low-level actions brutally with many meaningless attempt, thereby causing unsatisfactory results. By contrast, Liu. *et.al* Liu et al. [2019] tackled this problem by splitting it several into basic skills. Then the DQN reinforcement learning algorithm and a ruled-based method were proposed to schedule and execute these skills, respectively. In this proposal, we aim to improve the performance from the different aspects, *i.e.,* model-free agent construction.

3 Method

3.1 Problem Formulation

In Figure. 2, we illustrate the interaction process of the SC2LE platform and the agent. In the following, we introduce the details of each parts in this system and the approach we adopt.

- **Environment**: The environment is the Starcraft II API in SC2LE. It receives the actions from the player or agent, interact the game, and returns the corresponding observations.
- **Action**: The agent actions in this platform directly follow the form of human actions. For example, to control a unit, it first select the action 'select_rect', and then picks two points (P_1,

Algorithm 1: Random Agent

```
Initialise s;
for each step do
    Randomly sample a from A_s, available action set in s;
    Get transition s', a, s, r;
    s ← s';
end
```

Algorithm 2: Scripted Agent

```
BUILD-SUPPLY ← False;
BUILD-BARRACK ← False;
Initialise s;
for each step do
    Command all available workers to mine;
    Do the following judgements based on s;
    if not BUILD-SUPPLY then
        Choose a worker to build a supply;
        BUILD-SUPPLY ← True;
    end
    else if not BUILD-BARRACK then
        Choose a worker to build a barracks;
        BUILD-BARRACK ← True;
    end
    else
        Command the barrack to build marines
    end
    if Population meet the limit then
        BUILD-SUPPLY ← False;
    end
    Summary the current observation as s';
    s ← s';
end
```

P_2) of a rectangle, which should cover the unit. To further order the unit to do something, *e.g.*, to build a supply, it first select this action 'build_supply' and then choose the point P_3.

- **Observations** The observations involve two aspects, the game information and feature maps. The game information includes resources, available actions, building queues and other visible information. Feature maps is the processed segmentation maps of screen maps. This kind of feature maps are more easier for the agent to learn than the original screen. We have many different kinds of feature maps for use as shown in Figure 3.

3.2 Baseline Models

Before conducting our training process, we set up two baseline models, random agent and scripted agent. We introduce them in details in this section.

Random Agent Random agent is a trivial naive model, which randomly samples action from the available action choices in each frames. Its outputs are meaningless. Thus, it gets unsatisfied results. We mark the random agent as a brief algorithm in Algorithm 1.

Scripted Agent Actually, we have enough experience and know the rules about this RTS game. Thus, in addition to the scripted agent, we also include a well-designed scripted agent as baseline. It totally follows a series of if-else rules and case judgement. It checks the current situation and response an appropriate action from the current valid action set. Specially, because the purpose of this game is to build marines, which is trained from the barrack, this agent checks if the barrack has

Algorithm 3: Actor Critic Algorithm

Initialise s, θ;
Sample $a \sim \pi_\theta$;
for *each step* **do**
  Perform a_t according to policy $\pi(a_t|s_t;\theta)$;
  Receive reward $r = R_s^a$; sample transition $s' \sim P_s^a$;
  Sample action $a' \sim \pi_\theta(s', a')$;
  $\delta = r + \gamma Q_w(s', a') - Q_w(s, a)$;
  $\theta = \theta + \alpha \nabla_\theta log \pi_\theta(s, a) Q_w(s, a)$;
  $w \leftarrow w + \beta \delta \phi(s, a)$;
  $a \leftarrow a', s \leftarrow s'$;
end

been built. If not, an worker will be selected and build a barrack. At the same time, other workers are mining (collecting money). After that, it builds marines continuously until it counters the population limit. Otherwise, a worker will be selected to build a supply, which can enhance the population limit. This process is repeated until the end. We formulate this process in Algorithm 2. Noted that we does not include very detailed actions in the algorithm, like how to select a worker, for clear illustration.

3.3 Training Strategy

We design our training strategies based on Actor Critic and Asynchronous Advantage Actor Critic (A3C) algorithms. We first introduce the reward function. Then, we introduce the detailed design as follow.

Reward: For the game beginning task with regard to building Marines, here we directly correspond the reward value with the marine number, *i.e.,* +1 if increasing a marine. It is different from the full game where the reward function is just the winning or losing (+1/-1).

Actor Critic To deal with the indistinct reward mentioned in Section 1, we adopt actor-critic architecture that directly parameterizes the policy $\pi_\theta(a|s)$ and updates the parameters θ by performing gradient ascent on $E[R_t]$. In this approach, the critic plays a role as evaluating the reward of actions, guiding the policy function used to generate actions, denoted as

$$Q_w(s, a) \approx Q^{\pi_\theta}(s, a) \tag{1}$$

where $Q(s, a)$ is the state-action value function, the critic and actor are responsible for updating the actor parameter w and policy parameter θ. In our project, the state s is the real-time game status as well as segmented screen map, action can be represented by vectors where each dimension is a action category. The pseudo-code is shown in Algorithm 3.

Under this strategy, the w and θ are updated by steps until convergence.

Asynchronous Advantage Actor Critic Asynchronous Advantage Actor Critic (A3C) is an updated version of Actor Critic algorithm. It accumulates the training process of traditional actor critic algorithm in an asynchronous way. In other words, the algorithm is running on multi threads in parallel. At the end of each epoch, the gradient is accumulated from multi threads and used to update the global parameter. Then the thread-specific parameters are synchronized with the global one. We illustrate this algorithm in Algorithm 4.

4 Experiments

4.1 Implmentation Details

We train our models and compare it with two agents: RandomAgent and ScriptAgent. RandomAgent randomly chooses an action at every time step during the game, while ScriptAgent is well-designed with hand-craft playing strategies. Specifically, We control ScriptAgent to focus on quickly mining resources and building Marines, without many irrelevant procedures. When the Marines number goes

Algorithm 4: Asynchronous Advantage Actor Critic Algorithm

Initialize global shared parameter θ; Assume thread-specific parameter θ';
Reset $d\theta \leftarrow 0$;
Initialize s, θ;
Sample $a \sim \pi_\theta$;
for *each step* **do**
  Perform a_t according to policy $\pi(a_t|s_t;\theta)$;
  Receive reward $r = R_s^a$; sample transition $s' \sim P_s^a$;
  Sample action $a' \sim \pi_\theta(s', a')$;
  Accumulate gradients wrt θ' $d\theta$;
  $\delta = r + \gamma Q_w(s', a') - Q_w(s, a)$;
  Perform asynchronous update of θ usirng $d\theta$.;
  $\theta = \theta + \alpha \nabla_\theta log \pi_\theta(s, a) Q_w(s, a)$;
  $w \leftarrow w + \beta \delta \phi(s, a)$;
  $a \leftarrow a', s \leftarrow s'$;
end

(a) (b) (c)

Figure 4: Comparison of (a) RandomAgent, (b) ScriptAgent, (c) Our Agent. (a) RandomAgent randomly sample actions at each frame and fail to build any marine to the end; (b) ScriptAgent builds marines in a normal way; (c) Our trained actor critic agent has learned to build marines but is stuck by the population limit.

up to an upper limit, ScriptAgent turns to build another Supply Station so as to accommodate more Marines. As for our agent, we train it on a MacBook Pro equipped with a 4GB Radeon Pro 555X GPU. For implementation details, we set the feature map resolution of both screen and minimap as 64×64. The feature extracting neural network is composed of 4 convolution layers and 6 fully connected layers. The maximum steps in one iteration is set as 2000 and the total iterations is 340. The learning rate is 5e-4, which is a carefully tuned value. The discount ratio 0.99.

Agent	Curriculum score
RandomAgent	0
ScriptAgent	27
Actor Critic Agent	11
A3C Agent	11

Table 1: Curriculum score comparison between different agents.

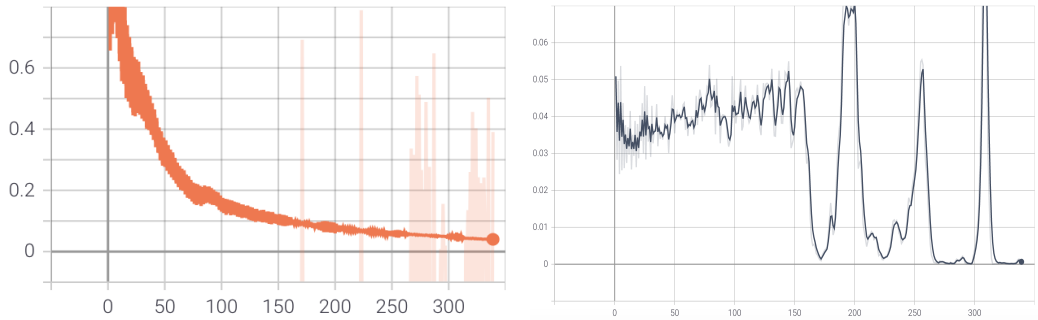

(a) (b)

Figure 5: (a) Policy loss and (b) Value loss over training iterations.

4.2 Experimental Results

We evaluate the agent performance by the curriculum score, *i.e.,* the Marine number within certain time. We show the experimental results by screenshots as shown in Fig 4. It can be observed that RandomAgent is unable to obtain any score, while ScriptAgent benefits from the clear target and therefore wins the highest score 27. Our Agent learns the game strategy from random actions but wastes time on many meaningless steps. It has learned how to build marines to get positive rewards. However, it is stuck by the population limit. It does not know that building the Supply Station can enhance the population limit. So it underperforms compared to ScriptAgent and achieves score 11.

We try to solve this problem with the improved version, Asynchronous Advantage Actor Critic (A3C). We run 8 asynchronous threads at the same time in parallel. Other settings like learning rate and discount ratio are same to the original actor critic agent. However, we did not find any obvious advantage of A3C model over the actor critic agent. It also achieve the most 11 reward and fails to learn building supply to enhance the population limit. In addition, the training process sometimes becomes more unstable. Thus, the A3C agent requires further tuning.

We also record the loss curve of the training process for actor critic agent as in Figure 5. Specially, policy loss L_{policy} and value loss L_{value} are computed with the following equations.

$$A = v_{target} - v \tag{2}$$

$$L_{policy} = ReduceMean(log(P_a) * A) \tag{3}$$

$$L_{value} = ReduceMean(v * A) \tag{4}$$

where A means the difference between the target value and the predict value and P_a means the probability of taking actions. Figure 5 shows that the loss of our model converges to a low level but the stability needs improvements.

5 Conclusion

In this paper, we study the beginning of Starcraft II game. The pipeline for scripted agent and reinforcement learning based agents is built on the SC2LE platform. We design both random and scripted agents as baseline. Our agents are trained with Actor Critic and A3C algorithms with detailed analysis. Although their current performance is still inferior to the well-designed scripted agent, it learns much knowledge in this task and is much better than the random agent. This project presents a promising hope for the further research in this challenging task.

References

Murray Campbell, A Joseph Hoane Jr, and Feng-hsiung Hsu. Deep blue. *Artificial intelligence*, 134 (1-2):57–83, 2002.

David Churchill, Zeming Lin, and Gabriel Synnaeve. An analysis of model-based heuristic search techniques for starcraft combat scenarios. In *AIIDE Workshops*, 2017.

Jonas Gehring, Da Ju, Vegard Mella, Daniel Gant, Nicolas Usunier, and Gabriel Synnaeve. High-level strategy selection under partial observability in starcraft: Brood war. *arXiv preprint arXiv:1811.08568*, 2018.

ZheJie Hu and Tomoyuki Kaneko. Enhancing sample efficiency of deep reinforcement learning to master the mini-games of starcraft ii. 2019.

Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human-level performance in 3d multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.

Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*, pages 3675–3683, 2016.

Dennis Lee, Haoran Tang, Jeffrey O Zhang, Huazhe Xu, Trevor Darrell, and Pieter Abbeel. Modular architecture for starcraft ii with deep reinforcement learning. *arXiv preprint arXiv:1811.03555*, 2018.

Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *Int. J. Robotics Res.*, 37(4-5):421–436, 2018. doi: 10.1177/0278364917710318. URL `https://doi.org/10.1177/0278364917710318`.

Tianlin Liu, Xihong Wu, and Dingsheng Luo. A hierarchical model for starcraft ii mini-game. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 222–227. IEEE, 2019.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015a.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nat.*, 518(7540):529–533, 2015b. doi: 10.1038/nature14236. URL `https://doi.org/10.1038/nature14236`.

Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 3303–3313, 2018.

Zhen-Jia Pang, Ruo-Ze Liu, Zhou-Yu Meng, Yi Zhang, Yang Yu, and Tong Lu. On reinforcement learning for full-length game of starcraft. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4691–4698, 2019.

Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.

Roman Ring et al. Replicating deepmind starcraft ii reinforcement learning benchmark with actor-critic methods. 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*, 3(1):73–84, 2018.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

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.

Nicolas Usunier, Gabriel Synnaeve, Zeming Lin, and Soumith Chintala. Episodic exploration for deep deterministic policies: An application to starcraft micromanagement tasks. *arXiv preprint arXiv:1609.02993*, 2016.

Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1703.01161*, 2017.

Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575:350–354, 2019.