

Reinforcement Learning With Dual-Observation for General Video Game Playing

Chengpeng Hu[✉], Ziqi Wang[✉], Tianye Shu[✉], Hao Tong[✉], *Graduate Student Member, IEEE*, Julian Togelius[✉], *Senior Member, IEEE*, Xin Yao[✉], and Jialin Liu[✉], *Senior Member, IEEE*

Abstract—Reinforcement learning (RL) algorithms have performed well in playing challenging board and video games. More and more studies focus on improving the generalization ability of RL algorithms. The General Video Game AI (GVGAI) Learning Competition aims to develop agents capable of learning to play different game levels that were unseen during training. This article summarizes the five years' GVGAI Learning Competition editions. At each edition, three new games were designed. The training and test levels were designed separately in the first three editions. Since 2020, three test levels of each game were generated by perturbing or combining two training levels. Then, we present a novel RL technique with dual-observation for general video game playing, assuming that it is more likely to observe similar local information in different levels rather than global information. Instead of directly inputting a single, raw pixel-based screenshot of the current game screen, our proposed general technique takes the encoded, transformed global, and local observations (LOs) of the game screen as two simultaneous inputs, aiming at learning local information for playing new levels. Our proposed technique is implemented with three state-of-the-art RL algorithms and tested on the game set of the 2020 GVGAI Learning Competition. Ablation studies show the outstanding performance of using encoded, transformed global, and LOs as input.

Manuscript received 9 November 2020; revised 6 November 2021 and 25 February 2022; accepted 21 March 2022. Date of publication 1 April 2022; date of current version 16 June 2023. This work was supported by the Research Institute of Trustworthy Autonomous Systems (RITAS), the Guangdong Provincial Key Laboratory under Grant 2020B121201001, in part by the Program for Guangdong Introducing Innovative and Entrepreneurial Teams under Grant 2017ZT07X386, in part by the Shenzhen Science and Technology Program under Grant KQTD2016112514355531, and in part by the National Natural Science Foundation of China under Grant 61906083. (Corresponding author: Jialin Liu.)

Chengpeng Hu, Ziqi Wang, and Tianye Shu are with the Research Institute of Trustworthy Autonomous System (RITAS), Southern University of Science and Technology, Shenzhen 518055, China, and also with the Guangdong Provincial Key Laboratory of Brain-Inspired Intelligent Computation, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China (e-mail: 12132333@mail.sustech.edu.cn; 12132362@mail.sustech.edu.cn; 12132356@mail.sustech.edu.cn).

Jialin Liu is with the Guangdong Provincial Key Laboratory of Brain-Inspired Intelligent Computation, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China (e-mail: liujl@sustech.edu.cn).

Hao Tong and Xin Yao are with the Research Institute of Trustworthy Autonomous System (RITAS), Southern University of Science and Technology, Shenzhen 518055, China, also with the Guangdong Provincial Key Laboratory of Brain-Inspired Intelligent Computation, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China, and also with the University of Birmingham, B15 2TT Birmingham, U.K. (e-mail: htong6@outlook.com; xiny@sustech.edu.cn).

Julian Togelius is with New York University, New York, NY 11201 USA (e-mail: julian@togelius.com).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TG.2022.3164242>.

Digital Object Identifier 10.1109/TG.2022.3164242

Index Terms—Artificial intelligence, Atari, general video game artificial intelligence (GVGAI), general video game playing (GVGP), reinforcement learning (RL).

I. INTRODUCTION

GAMES have always been popular benchmarks for testing artificial intelligence methods. In recent years, video games have been widely used due to multiple favorable characteristics, such as continuous dynamics, large action/state space, partial observability, and speed, which make them useful proxies of real-world problems. Although reinforcement learning (RL) has shown considerable success in learning to play the game of Go [1], [2], and various video games [3]–[5], the trained models were only used to play the games and levels on which they have been trained. Therefore, researchers have more recently also addressed the challenge of making RL *generalize* [6]–[8].

The word *generalize* in the context of general video game playing (GVGP) has several meanings: An agent can perform well on unseen levels after training on given levels of an identical game; or an agent trained on some games can work well on new games. Many studies on GVGP focus on the former aspect. Braylan *et al.* [9] used transfer learning to reuse the neural modules of agents in new agents for playing new games. Similarly, Tutum *et al.* [10] evolved a context module to recognize temporal variations in games and autonomous driving simulations. Techniques widely used in deep learning, such as regularization, data augmentation, and batch normalization, have also been used to enhance the performance of RL algorithms or to increase their generalisation performance [11]–[13]. Justesen *et al.* [14] enlarged the training set with procedural content generation methods to improve the generality of RL agents. It was shown that for some games, with a relatively large amount of generated levels and training steps, the agent can work well on a variety of levels within the same distribution [14]. Partly inspired by this, the Procgen benchmark [15] includes 16 game-like environments with procedurally generated levels to evaluate the generalization of RL agents.

The General Video Game AI (GVGAI) Learning Competition [16], [17], first organized in 2017, requires agents to learn to play games without game simulators and achieve good performance on levels that are unseen during training. These unseen levels follow the original game rules but have a different layout and even new elements. The submitted agents are ranked according to their average win rate and the score they achieved while playing the unseen levels.

Despite the facilitation of using the GVGAI learning platform, only a few papers describing research using the platform have been published [6], [14], [18]–[21] and no significantly well-performed agents were received in the 2017, 2018, or 2019 editions of the competition [17]. After observing the behaviors of these entries, we assume that their poor performance can probably be blamed on 1) the significant differences between the training and test level maps (including the map size), even when the game rules remain unchanged, and 2) the direct use of game screenshots as inputs.

Motivated by verifying the above guess and developing a better GVGAI learning agent, we investigate the following questions in this article: 1) Can the specification of game states be reduced for better generality? 2) Can generating more game states as training data improve the performance of an RL agent when playing unseen levels?

In this article, we seek to answer the above research questions in the context of RL for GVGP. The main contributions of our work are as follows. 1) We propose a novel RL technique that simultaneously takes a transformed global observation (GO) and a transformed local observation (LO) (referred to as *dual-observation* from now on) as input. Moreover, we use a novel *tile-vector encoding* method for rapidly encoding observations and handling new tiles that may appear in unseen levels. 2) The avatar is randomly placed at the beginning of each game during training as a simple technique to increase variety in training data.

The aforementioned dual-observation, tile-vector encoding and training with random initial positions are applied to three state-of-the-art RL algorithms from stable-baselines3 [22]. Ablation studies on the 2020 GVGAI Competition games with dense, periodic, and sparse rewards show that dual-observation and tile-vector encoding significantly improves the tested algorithms' performance in playing most of the unseen levels.¹ Moreover, the results of five years' GVGAI Learning Competition editions, particularly the 2019–2021 editions, are reported and summarized in this article. The methodology of designing the 2020 and 2021 competition games are also described.

In the remainder of this article, Section II presents the GVGAI Learning Competition and the competition results of all its five editions. The details of our proposed RL technique with dual-observation and tile-vector encoding, and corresponding experimental studies are described in Section III. Section IV compares the proposed RL technique with dual-observation implemented with three state-of-the-art RL algorithms on GVGAI Learning Competition games and discusses them. Finally, Section V concludes this article.

II. GVGAI LEARNING COMPETITION

A. GVGAI Learning Platform

The GVGAI competition [16], [17] has offered an easy-to-use, open-sourced, common platform for researchers, and practitioners to test and compare AI methods since 2014. During the last few years, the framework has been expanded into several tracks to meet the demand of different research

directions [16], [17]. The GVGAI learning platform [6] particularly focuses on GVGP using RL methods. The GVGAI learning platform not only contains more than 100 arcade-style single-player games and dozens of two-player games, but also offers the potential of adding an unlimited number of games and levels thanks to the video game definition language (VGDL) [23], [24]. This functionality facilitates the enlargement of training and test datasets by allowing for games or levels to be designed either manually or with procedural content generation methods [25]–[29].

B. Competition Tasks and Rules

Three different games, each with five levels, are provided in the learning competition. For each game, two levels are released as training levels, the other three remain unknown to participants for testing. At every game tick, an agent receives a screenshot of the current game screen, a list of legal actions, its game score as well as the game termination state (PLAYER_WINS or PLAYER_LOSES if the game is terminated, otherwise NO_WINNER), and is expected to return an action to play within a certain time. The agent is expected to be able to play levels that were unseen during training. Readers are referred to [16], [17] and the competition website² for more detailed competition task and rules.

The rules of the GVGAI Learning Competition basically follow the ones of the GVGAI Single-player Planning Competition [30]. In the training phase, agent entries can make the best use of the two given training levels of each game. In the testing phase, all submitted entries are executed using an identical Python script. Each entry is evaluated by its win rate and average score over 10 independent runs on each test level of each game, where each game has a maximum length of 2000 game ticks unless a different maximum length is explicitly defined in the game rules. For each level of each game, the entries are ranked first by their average scores, and then their win rates in the decreasing order. For entries that receive the same number of wins and the same average score, the game length is considered to break the tie. Finally, points are assigned to entries according to their rankings per level per game (25 for 1st, 18 for 2nd, 15 for 3rd, 12 for 4th, 10 for 5th, 0 for the others). The final scores of the individual entries are accumulated across the game levels. The one with the most final points wins the first prize in the competition.

C. Five Years' Experience of the GVGAI Learning Competition

The GVGAI Learning Competition has been organized every year since 2017. Disappointingly, no particularly high performing agent has been received between 2017 and 2019 [16], [17]. First, in the past editions of competition, some of the test level maps were of a different size to the training maps, resulting in different sizes of screenshots and failure of some implementations of RL algorithms [17]. Second, the training and testing level maps were designed separately and dissimilar. It is not

¹Code available on GitHub: <https://github.com/SUSTechGameAI/DORL>

²[Online]. Available: <http://aigames.cn/gvgai/ppscog2020>

surprising that an RL agent plays randomly when it encounters a game state that is very different from the one used in training. The former reason poses technical challenges, while the latter makes it necessary to reflect the design of general learning agents. Therefore, the methodology for designing competition games and levels was changed from 2020.

D. Design of Competition Game Sets in 2020 and 2021

The games used in 2020 and 2021 competition editions have been carefully designed with the same screen size for all levels of a given game and different goals from each other.

1) *Game Sets*: The variety in reward frequency is a challenge to the generality of learning algorithms. The goal of the games varied from resource-collection to maze navigation to survival. This followed the tradition of designing GVGAI competition games since 2017, where designers have aimed for diversity. Meanwhile, each game has a unique time limit corresponding to the task difficulty.

a) *Game Set in 2020*: The 2020 competition edition provided three games, namely *GoldDigger*, *TreasureKeeper*, and *WaterPuzzle*, with dense, periodic, and sparse rewards, respectively.

- 1) *GoldDigger* is a resource-collection game. The avatar, thus the game character controlled by a player/agent, is expected to avoid monsters and collect all the jewels in a level to win the game. When a collision of the avatar with any monster occurs, it suddenly loses the game. It is also considered as a failure if a maximum of 2000 game ticks is passed. Once a jewel is collected, a score will be rewarded to the avatar. The avatar can also kill monsters to gain a score. Hence, *GoldDigger* is a game with dense rewards because of the large number of jewels and monsters in levels.
- 2) *TreasureKeeper* is a Sokoban-like game aiming at keeping the avatar and treasure chests away from monsters. The game is designed with a periodic reward given every 100 game ticks. Either a collision of monster and avatar or a monster and treasure chest will terminate the game with failure. If the avatar and all treasure chests survive for more than 600 game ticks, it wins.
- 3) *WaterPuzzle* is a maze game with sparse rewards. There are only three sprites in each level: the avatar, an immobile key, and an immobile door. The agent is rewarded for collecting the key and for touching the door after collecting the key. The agent wins the game only if it collects the key and then touches the door in 1500 game ticks.

b) *Game Set in 2021*: Similarly, the 2021 competition edition³ provides three games, namely *TrappedHero*, *BraveKeeper*, and *GreedyMouse*, with dense, periodic, and sparse rewards, respectively. Description of the games can be found on the competition website.

2) *Creation of Competition Level Sets*: The competition level set for each game is composed of two training levels and three test levels. Instead of using very distinct test levels that were




Fig. 1. Training levels in 2020. Red circle indicates the avatar. (a) *GoldDigger-0*. (b) *GoldDigger-1*. (c) *TreasureKeeper-0*. (d) *TreasureKeeper-1*. (e) *WaterPuzzle-0*. (f) *WaterPuzzle-1*.


dissimilar to the ones for training as in the 2017–2019 competition editions, all the test levels in the 2020 and 2021 editions were generated from training levels by making *changing one or several tiles* or performing *map combination*, described later in Section II-D2b.

a) *Training Level Set*: Screenshots of training levels are illustrated in Fig. 1. The training levels for each game were designed to have different skill-depth. *GoldDigger-1* is harder than *GoldDigger-0*. On the contrary, more surrounding obstacles in *TreasureKeeper-0* make it easier to protect the treasure chests from monsters. In *WaterPuzzle*, the distances between sprites affect the difficulty of levels. Consequently, *WaterPuzzle* levels not only use different mazes but also differ in sprite locations.

b) *Test Level Set*: For each game, a human designer generated three test levels based on training levels with one of the following operators separately.

- 1) *Single-tile change*: *Level-3* is designed by changing one single tile from *Level-1*. For instance, the single-tile change occurs at the top right corner of *GoldDigger-1*, where an obstacle is replaced by a jewel (see Fig. 2).
- 2) *Multi-tile change*: *Level-2* is designed by changing multiple tiles from *Level-0*. Taking *GoldDigger* as an example, most of the tiles in the middle row of *GoldDigger-0* are changed to generate *GoldDigger-2*, a more challenging level due to the obstacles.
- 3) *Map combination*: To validate if the knowledge of all training levels has been transferred, we designed the final test level of each game by combining the two training levels. The last test level is simply a combination of the top half of *level-0* and the bottom half of *Level-1*. Notably, the avatar remains unique, and its location is changed at each level.

³[Online]. Available: <http://aingames.cn/gvgai/cog2021/>

Fig. 2. Design of *GoldDigger* test levels.TABLE I
COMPETITION RANKINGS OF THE 2019 EDITIONS

(a) Competition rankings at IEEE CEC2019.

Rank	Entry	Points	# Wins
1	mballa	174	10/180
2	Random	158	1/180
3	Ben	151	14/180

(b) Competition rankings at IEEE CoG2019.

Rank	Entry	Points	# Wins
1	Random	158	9/180
2	mballa	153	20/180
3	TNTBot	124	1/180
4	UjiAgent	120	0/180

For example, Fig. 2 illustrates how the test levels are generated for *GoldDigger* using the above operators.

E. Competition Entries and Rankings

Competition entries and rankings of 2017 and 2018 editions have been reported in the work of [17]. In this article, we report the entries and rankings of the 2019–2021 editions.

1) *The 2019 Competition Editions*: Two editions of the GVGAI Learning Competition were held at the 2019 IEEE Congress on Evolutionary Computation (IEEE CEC2019)⁴ and the 2019 IEEE Conference on Games (IEEE CoG2019)⁵, which received two and three entries, respectively.

The final results and rankings of those editions according to the rules described in Section II-B are shown in Table I.⁶ “Random” refers to a baseline agent that plays actions uniformly at random.

⁴[Online]. Available: <http://aingames.cn/gvgai/cec2019/>

⁵[Online]. Available: <http://aingames.cn/gvgai/cog2019/>

⁶Unfortunately, description of the entries, including algorithms and training methods, was not mandatory when submitting an entry to the 2019 competition editions.

TABLE II
RANKINGS OF THE 2020 COMPETITION EDITION

Rank	Entry	Points	# Wins
1	Xybb	189	10/180
2	Random	149	4/180
3	Visionpack2	129	2/180
4	Eliaentity	109	0/180
5	Visionpack3	102	1/180
6	Visionpack	101	2/180

2) *The 2020 Competition Edition*: The 2020 GVGAI Learning Competition was held at the Sixteenth International Conference on Parallel Problem Solving from Nature (PPSN-XVI) and the 2020 IEEE Conference on Games (IEEE CoG2020). The following five entries were received and ranked together (cf. Table II).

- *Xybb*: Three DQN models were trained for 200 000 timesteps on each game separately using the two training levels and some new levels manually generated by the authors from the given training levels.
- *Eliaentity*: A DQN model from stable-baselines was trained for 2 million timesteps on *GoldDigger-0* and used in test levels of *GoldDigger*. Random actions were played in the other two games.
- *Visionpack*: For each of the three games, a PPO2 model from stable-baselines3 [22] was trained on level 0 of each game for 1 million timesteps. While testing on different games, the corresponding model, identified by the size of the game screen, was used to play the test levels.
- *Visionpack2*: Two separate DQN models from stable-baselines3 were trained on *GoldDigger-0* and *GoldDigger-1*. An image classifier was created to help select a suitable model for *GoldDigger*’s test levels. Random actions were played in the other two games.
- *Visionpack3*: A PPO2 model from stable-baselines3 was trained on *GoldDigger*’s training levels for 2 million timesteps using the concept of vectorized environments and an ACER model [31] was trained on *Waterpuzzle-0* only for 2 million timesteps. *Treasurekeeper* was played randomly.

The detailed results on each game and the running log are presented on the competition website. According to Table II, *Xybb* performs significantly better than Random and all the other entries, which could not even beat Random.

3) *The 2021 Competition Edition*: The 2021 GVGAI Learning Competition was held at the 2021 IEEE Conference on Games (IEEE CoG2021). No entry was received in this edition. A PPO agent with dual-observation, named *Arcane*, described later in Section III, was used as a baseline agent. Its performance in playing training and test levels of 2021 competition games is illustrated in Table VIII. Besides *Arcane*, the performance of several planning agents, assuming the availability of forward model, are also reported for comparison. Open Loop Expectimax Tree Search (OLETS), the winner of the first GVGAI Single-Player Planning Competition [30], performs the best on the three games.

TABLE III
FOUR GROUPS OF AGENTS CATEGORISED BY INPUT

Notation	Input			
	Pixel	One-hot	GO	LO
G1	✓		✓	
G2	✓		✓	✓
G3		✓	✓	
DORL		✓	✓	✓

TABLE IV
EXPERIMENTAL SETTING AND NETWORK CONFIGURATION

Hyper-parameter		Value
All	Total training steps	1,000,000
	Discount factor	0.99
	σ of scaled softmax	3
DQN	Initial exploration	1
	Final exploration	0.1
	Replay memory size	40,000
	Learning starts	0
	Mini-batch size	32
A2C	Learning rate	0.001
	Rollout length	5
	Learning rate	0.0007
PPO	Number of environments	8
	Rollout length	2,048
	Learning rate	0.0003
	Number of environments	8

More detailed description of parameters can be found in [22].

Despite the our best attempts at making GVGAI learning platform easy to use, only a few papers have been published describing research done with the platform [6], [14], [18]–[21]. Notably, no agents performing much better than the baselines were received in the past competition editions [17].

Another video game-like testbed for testing and enabling generalization in RL is Obstacle Tower, which is a single game but where levels are procedurally generated and multiple aspects of visual input, including lighting and art style, are systematically varied [32].

III. RL WITH DUAL-OBSERVATION

As previously reported in Section II-E, no significantly effective entry has been reported in the GVGAI Learning Competition editions. The only well-performing ones are the baseline agents, *xybb* and *Arcane*, in the 2020 and 2021 editions, respectively. Those two agents are instantiations of our proposed RL agent with dual-observation and tile-vector encoding detailed in this section. The agents submitted to the 2017–2020 competition editions used the raw screenshots of the whole game screen as a single input [17]. Under the assumption that local information has a higher chance of remaining invariant across different levels, we design a dual-observation as the input to RL agents (“DORL” for short) that makes use of the local information during training and test. DORL is illustrated in Fig. 3. DORL processes the screenshot of the game screen into two observations (cf. Section III-A), a global one and a local one, and then encodes them with tile-vector encoding and one-hot

matrices (cf. Section III-B), which are later used as inputs of the learning model (cf. Section III-C). DORL can be implemented with various learning algorithms. When using DQN, either the deterministic policy or the stochastic one can be used by DORL for determining the action to play (cf. Section III-D). Section III-E clarifies the novelty of this work.

A. Transformed Game Observations

Different areas of the game observation are of different importance. The avatar’s surrounding area usually has a more immediate and greater impact than the distant areas on the action selection. Therefore, DORL is proposed to transform the received screenshot of the game screen into a GO and a local one. The transformation procedure is described as follows and illustrated on the left part of Fig. 3.

- 1) A *GO* is not the direct use of the screenshot of the game screen. Instead, it is a $h_g \times w_g$ RGB image in which the avatar is placed at centre (cf. Fig. 4), where $h_g = 2h - h_{\text{tile}}$ and $w_g = 2w - w_{\text{tile}}$, w and h refers to the width and height of the original screenshot, w_{tile} and h_{tile} refer to the width and height of one single tile, respectively. This guarantees that the original game screen is always contained in the GO wherever the avatar moves.
- 2) A *LO* is a $h_l \times w_l$ RGB image ($h_l < w$, $h_l < h$) in which the avatar is centred (cf. Fig. 4). In our experiments, h_l and w_l are set as 50, which means 5×5 surrounding tiles are used.

In the context of VGDL [23], [24], the avatar can be easily located with the tile dictionary described in the following section.

B. Tile-Vector Encoded Inputs

The GVGAI platform represents levels with tile-based maps [23], [24]. Those tiles, either for different types of sprites or for accessible areas, have a predefined, fixed size of 10×10 [17]. Hence, we convert the RGB images into tile-based matrices as input, on which one-hot encoding can be easily applied.

All distinct tiles appeared in the training levels have been collected to build a *tile dictionary*, which is a set of $\langle \text{code}, \text{reference vector} \rangle$ tuples. The *code* $\in \mathbf{N}$ identifies a unique 10×10 -pixel *tile image*. For each tile, 5 pixels are selected from its 4th row with a stride of 2, and another 5 are selected from its 5th column with a stride of 2. For each selected 10 pixels, its averaged RGB value is calculated as $\frac{1}{3}(v_R + v_G + v_B)$, where v_R , v_G , and v_B refer to the red, green, and blue values of the pixel, respectively. As a result, a *reference vector* composed of 10 averaged RGB values is obtained.

At every game tick during training or testing, the two RGB images of transformed global and LOs are considered as $\frac{h_g}{10} \times \frac{w_g}{10}$ and $\frac{h_l}{10} \times \frac{w_l}{10}$ grids of tiles, respectively. The blank spaces added during transformation, which are not part of the game screen, are replaced by zeros. For both grids, we can easily calculate the corresponding RGB-value vector for every tile inside and find its closest reference vector in the dictionary, measured by the Manhattan distance. Assuming that some unexpected tiles may

TABLE V
GOLDDIGGER: AVERAGE SCORE, STANDARD DEVIATION AND NUMBER OF WINS (IN THE BRACKETS) OVER 20 INDEPENDENT TRIALS

Agent	Level-0 (167 / -20)			Level-1 (98 / -20)			Level-2 (167 / -20)			Level-3 (98 / -20)			Level-4 (196 / -30)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
DQN	G1	-4.8 $\pm$ 1.1	0.0/-5.0(0)	-0.8 $\pm$ 1.8	0.0/-5.0(0)		-2.0 $\pm$ 2.4	0.0/-5.0(0)		-0.2 $\pm$ 1.1	0.0/-5.0(0)		-1.5 $\pm$ 3.2	4.0/-5.0(0)	
	G2	40.6 $\pm$ 41.3	78.0/-5.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		-1.5 $\pm$ 2.3	0.0/-5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		88.0 $\pm$ 64.7	144.0/-5.0(0)	
	G3	22.8 $\pm$ 2.5	29.0/17.0(0)	4.7 $\pm$ 1.5	6.0/-1.0(0)		-2.5 $\pm$ 2.5	0.0/-5.0(0)		5.3 $\pm$ 0.6	7.0/4.0(0)		-0.5 $\pm$ 3.7	5.0/-5.0(0)	
	DORL	129.1 $\pm$ 44.3	159.0/23.0(0)	41.3 $\pm$ 20.1	81.0/19.0(0)		-1.2 $\pm$ 2.2	0.0/-5.0(0)		27.9 $\pm$ 1.9	31.0/23.0(0)		99.4 $\pm$ 26.1	116.0/26.0(0)	
DQN _s	G1	-3.2 $\pm$ 2.4	1.0/-5.0(0)	0.2 $\pm$ 5.8	18.0/-5.0(0)		-0.2 $\pm$ 1.1	0.0/-5.0(0)		3.2 $\pm$ 9.1	23.0/-5.0(0)		15.4 $\pm$ 22.9	65.0/-5.0(0)	
	G2	99.8 $\pm$ 47.0	136.0/-5.0(0)	23.0 $\pm$ 13.3	41.0/1.0(0)		84.2 $\pm$ 52.2	135.0/-5.0(0)		24.2 $\pm$ 12.3	48.0/4.0(0)		97.4 $\pm$ 56.7	146.0/-2.0(0)	
	G3	111.4 $\pm$ 50.9	157.0/11.0(0)	26.8 $\pm$ 8.1	42.0/13.0(0)		95.6 $\pm$ 62.8	158.0/10.0(0)		26.1 $\pm$ 12.6	58.0/7.0(0)		111.7 $\pm$ 48.3	163.0/-5.0(0)	
	DORL	110.8 $\pm$ 61.0	157.0/-5.0(0)	53.4 $\pm$ 30.6	97.0/5.0(2)		105.8 $\pm$ 54.0	159.0/-5.0(0)		59.4 $\pm$ 33.6	97.0/-5.0(1)		92.4 $\pm$ 73.6	189.0/-5.0(0)	
PPO	G1	48.6 $\pm$ 49.1	147.0/-5.0(0)	23.7 $\pm$ 19.9	50.0/-5.0(0)		33.0 $\pm$ 46.8	117.0/-5.0(0)		31.1 $\pm$ 15.2	59.0/7.0(0)		7.2 $\pm$ 19.1	61.0/-5.0(0)	
	G2	96.5 $\pm$ 65.2	153.0/-5.0(0)	51.4 $\pm$ 22.4	86.0/-2.0(0)		51.0 $\pm$ 62.8	152.0/0.0(0)		42.6 $\pm$ 18.6	73.0/14.0(0)		61.6 $\pm$ 60.0	160.0/-5.0(0)	
	G3	100.7 $\pm$ 65.6	157.0/-5.0(0)	44.0 $\pm$ 25.9	84.0/1.0(0)		89.0 $\pm$ 64.5	149.0/-5.0(0)		37.3 $\pm$ 27.6	75.0/-2.0(0)		46.9 $\pm$ 63.9	164.0/-5.0(0)	
	DORL	90.1 $\pm$ 67.8	152.0/-5.0(0)	57.7 $\pm$ 28.0	87.0/7.0(0)		95.6 $\pm$ 66.3	157.0/-5.0(0)		50.8 $\pm$ 26.8	83.0/7.0(0)		128.4 $\pm$ 59.8	173.0/-5.0(0)	
PPO _D	G1	95.0 $\pm$ 57.2	144.0/-5.0(0)	9.7 $\pm$ 12.2	37.0/-5.0(0)		54.1 $\pm$ 58.7	135.0/-5.0(0)		10.1 $\pm$ 13.2	39.0/-5.0(0)		38.8 $\pm$ 50.1	130.0/-5.0(0)	
	G2	110.0 $\pm$ 59.3	157.0/-5.0(0)	45.4 $\pm$ 25.0	88.0/1.0(0)		102.0 $\pm$ 47.2	153.0/-5.0(0)		54.9 $\pm$ 23.9	84.0/11.0(0)		112.4 $\pm$ 68.0	179.0/-1.0(0)	
	G3	108.4 $\pm$ 60.0	155.0/-5.0(0)	33.6 $\pm$ 22.2	78.0/-1.0(0)		84.7 $\pm$ 66.4	153.0/-5.0(0)		34.3 $\pm$ 19.9	67.0/-2.0(0)		77.8 $\pm$ 70.4	162.0/-5.0(0)	
	DORL	73.6 $\pm$ 76.4	163.0/-5.0(5)	51.5 $\pm$ 26.5	82.0/2.0(0)		58.5 $\pm$ 71.7	162.0/-5.0(1)		50.2 $\pm$ 27.5	79.0/1.0(0)		91.2 $\pm$ 81.7	182.0/-5.0(0)	
A2C	G1	-1.6 $\pm$ 3.5	4.0/-5.0(0)	5.6 $\pm$ 4.7	13.0/-4.0(0)		-1.2 $\pm$ 3.6	7.0/-5.0(0)		2.4 $\pm$ 5.1	13.0/-5.0(0)		0.1 $\pm$ 4.2	9.0/-5.0(0)	
	G2	0.2 $\pm$ 4.0	11.0/-5.0(0)	29.2 $\pm$ 11.0	48.0/2.0(0)		17.7 $\pm$ 43.4	147.0/0.0(0)		27.6 $\pm$ 11.4	55.0/5.0(0)		3.8 $\pm$ 9.8	37.0/-5.0(0)	
	G3	54.0 $\pm$ 42.1	152.0/-2.0(0)	62.7 $\pm$ 25.4	81.0/4.0(0)		23.4 $\pm$ 31.4	64.0/-5.0(0)		54.6 $\pm$ 25.4	80.0/7.0(0)		102.4 $\pm$ 55.1	183.0/28.0(0)	
	DORL	149.6 $\pm$ 3.6	159.0/145.0(0)	55.1 $\pm$ 38.6	93.0/-1.0(1)		-1.0 $\pm$ 2.0	0.0/-5.0(0)		46.6 $\pm$ 37.2	90.0/-1.0(0)		167.7 $\pm$ 30.9	188.0/70.0(0)	
Ye <i>et al.</i> [21]		22.8 $\pm$ 19.3	64/-5(0)	1.5 $\pm$ 4.9	12/-5(0)		11.2 $\pm$ 11.4	30/-5(0)		3.5 $\pm$ 5.0	13/-5(0)		4.2 $\pm$ 12.4	48/-5(0)	
Random		3.8 $\pm$ 11.9	42 / -5 (0)	9.4 $\pm$ 12.5	35 / -5 (0)		9.3 $\pm$ 24.4	86 / -5 (0)		7.6 $\pm$ 11.8	38 / -5 (0)		2.0 $\pm$ 16.5	69 / -5 (0)	
MCTS		155.1 $\pm$ 8.2	162 / 126 (8)	67.4 $\pm$ 12.2	90 / 36 (0)		129.4 $\pm$ 36.6	161 / 6 (1)		67.7 $\pm$ 11.4	88 / 43 (0)		156.5 $\pm$ 31.1	191 / 75 (1)	
OLETS		164.2 $\pm$ 0.9	166 / 162 (20)	93.1 $\pm$ 6.2	98 / 75 (9)		165.0 $\pm$ 1.6	167 / 161 (20)		86.8 $\pm$ 11.9	98 / 63 (8)		192.5 $\pm$ 3.5	196 / 184 (13)	
RHEA		146.2 $\pm$ 25.4	166 / 89 (4)	61.8 $\pm$ 18.0	87 / 19 (0)		124.3 $\pm$ 51.0	167 / 1 (3)		57.9 $\pm$ 15.3	77 / 8 (0)		156.5 $\pm$ 30.3	181 / 58 (0)	
RS		155.0 $\pm$ 14.3	167 / 110 (6)	61.0 $\pm$ 19.8	84 / 22 (0)		142.6 $\pm$ 30.8	166 / 56 (1)		64.3 $\pm$ 14.5	85 / 36 (0)		167.1 $\pm$ 22.2	193 / 97 (0)	

The bold numbers denotes the highest average scores among four input types or planning agents.

TABLE VI
TREASUREKEEPER: AVERAGE SCORE, STANDARD DEVIATION AND NUMBER OF WINS (IN THE BRACKETS) OVER 20 INDEPENDENT TRIALS

Agent	Level-0 (35 / 5)			Level-1 (35 / 5)			Level-2 (35 / 5)			Level-3 (35 / 5)			Level-4 (35 / 5)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
DQN	G1	27.2 $\pm$ 9.9	35.0/10.0(11)	11.5 $\pm$ 10.0	35.0/5.0(2)		20.2 $\pm$ 10.8	35.0/5.0(5)		6.0 $\pm$ 2.0	10.0/5.0(0)		20.2 $\pm$ 11.1	35.0/5.0(5)	
	G2	26.0 $\pm$ 11.1	35.0/10.0(12)	11.2 $\pm$ 8.0	35.0/5.0(1)		26.0 $\pm$ 10.2	35.0/10.0(10)		13.5 $\pm$ 10.4	35.0/5.0(2)		25.0 $\pm$ 10.1	35.0/10.0(8)	
	G3	23.5 $\pm$ 10.1	35.0/10.0(7)	9.2 $\pm$ 6.4	30.0/5.0(0)		24.0 $\pm$ 9.3	35.0/10.0(6)		6.2 $\pm$ 2.7	15.0/5.0(0)		22.8 $\pm$ 11.0	35.0/5.0(8)	
	DORL	28.2 $\pm$ 9.8	35.0/10.0(13)	9.2 $\pm$ 7.1	35.0/5.0(1)		25.2 $\pm$ 10.9	35.0/5.0(9)		10.8 $\pm$ 9.1	35.0/5.0(2)		23.2 $\pm$ 12.5	35.0/5.0(10)	
DQN _s	G1	15.5 $\pm$ 8.2	30.0/5.0(0)	9.8 $\pm$ 8.7	35.0/5.0(1)		16.2 $\pm$ 10.1	35.0/5.0(3)		7.8 $\pm$ 6.8	35.0/5.0(1)		16.5 $\pm$ 8.8	35.0/5.0(3)	
	G2	27.0 $\pm$ 8.4	35.0/10.0(9)	12.8 $\pm$ 8.9	35.0/5.0(2)		24.5 $\pm$ 9.2	35.0/5.0(6)		10.5 $\pm$ 6.7	25.0/5.0(0)		19.8 $\pm$ 11.1	35.0/5.0(4)	
	G3	22.5 $\pm$ 10.9	35.0/5.0(7)	10.2 $\pm$ 9.1	35.0/5.0(1)		20.8 $\pm$ 12.7	35.0/5.0(7)		7.0 $\pm$ 4.8	25.0/5.0(0)		13.5 $\pm$ 6.1	25.0/5.0(0)	
	DORL	24.0 $\pm$ 11.9	35.0/5.0(9)	14.5 $\pm$ 10.2	35.0/5.0(3)		25.0 $\pm$ 11.4	35.0/5.0(9)		10.8 $\pm$ 8.8	35.0/5.0(1)		22.5 $\pm$ 12.4	35.0/5.0(9)	
PPO	G1	16.0 $\pm$ 11.7	35.0/5.0(5)	14.0 $\pm$ 10.6	35.0/5.0(3)		14.8 $\pm$ 9.9	35.0/5.0(2)		9.8 $\pm$ 7.3	35.0/5.0(1)		11.5 $\pm$ 9.2	35.0/5.0(2)	
	G2	20.0 $\pm$ 11.3	35.0/5.0(5)	12.2 $\pm$ 7.0	25.0/5.0(0)		17.5 $\pm$ 10.2	35.0/5.0(3)		13.2 $\pm$ 9.4	35.0/5.0(2)		14.8 $\pm$ 8.4	35.0/5.0(2)	
	G3	21.0 $\pm$ 12.0	35.0/5.0(7)	11.5 $\pm$ 9.0	35.0/5.0(2)		17.5 $\pm$ 9.8	35.0/5.0(3)		10.5 $\pm$ 8.4	35.0/5.0(1)		18.0 $\pm$ 11.2	35.0/5.0(5)	
	DORL	20.0 $\pm$ 10.7	35.0/5.0(5)	12.2 $\pm$ 10.3	35.0/5.0(3)		16.2 $\pm$ 11.3	35.0/5.0(3)		9.2 $\pm$ 5.8	25.0/5.0(0)		11.8 $\pm$ 8.1	35.0/5.0(1)	
PPO _D	G1	15.5 $\pm$ 11.6	35.0/5.0(3)	9.5 $\pm$ 8.0	35.0/5.0(1)		15.5 $\pm$ 10.6	35.0/5.0(3)		9.8 $\pm$ 8.9	35.0/5.0(2)		14.0 $\pm$ 10.7	35.0/5.0(3)	
	G2	20.2 $\pm$ 11.5	35.0/5.0(5)	13.8 $\pm$ 9.9	35.0/5.0(1)		19.5 $\pm$ 9.1	35.0/5.0(3)		7.8 $\pm$ 7.2	35.0/5.0(1)		17.5 $\pm$ 11.6	35.0/5.0(4)	
	G3	12.0 $\pm$ 7.3	35.0/5.0(1)	17.8 $\pm$ 12.3	35.0/5.0(6)		20.2 $\pm$ 10.4	35.0/5.0(4)		7.0 $\pm$ 3.7	15.0/5.0(0)		16.0 $\pm$ 9.7	35.0/5.0(2)	
	DORL	22.8 $\pm$ 11.2	35.0/5.0(6)	6.2 $\pm$ 3.1	15.0/5.0(0)		22.8 $\pm$ 12.2	35.0/5.0(9)		8.8 $\pm$ 6.9	30.0/5.0(0)		15.5 $\pm$ 12.1	35.0/5.0(4)	
A2C	G1	25.2 $\pm$ 9.3	35.0/5.0(6)	6.0 $\pm$ 2.0	10.0/5.0(0)		16.8 $\pm$ 9.3	35.0/5.0(2)		5.8 $\pm$ 1.8	10.0/5.0(0)		16.8 $\pm$ 9.5	35.0/5.0(2)	
	G2	23.8 $\pm$ 10.5	35.0/5.0(8)	6.8 $\pm$ 4.0	20.0/5.0(0)		21.8 $\pm$ 11.4	35.0/5.0(7)		6.0 $\pm$ 2.0	10.0/5.0(0)		19.0 $\pm$ 11.6	35.0/5.0(6)	
	G3	31.2 $\pm$ 8.0	35.0/5.0(15)	15.2 $\pm$ 10.8	35.0/5.0(3)		25.5 $\pm$ 11.4	35.0/5.0(11)		14.2 $\pm$ 9.4	35.0/5.0(1)		27.0 $\pm$ 10.3	35.0/5.0(10)	
	DORL	26.5 $\pm$ 10.4	35.0/5.0(11)	6.5 $\pm$ 2.3	10.0/5.0(0)		23.2 $\pm$ 10.8	35.0/5.0(7)		6.0 $\pm$ 2.5	15.0/5.0(0)		20.0 $\pm$ 9.5	35.0/5.0(4)	
Ye <i>et al.</i> [21]		7.5 $\pm$ 5.1	30/5(0)	8.2 $\pm$ 5.1	25/5(0)		7.3 $\pm$ 3.1	15/5(0)		7.3 $\pm$ 5.4	30/5(0)		6.8 $\pm$ 3.0	15/5(0)	
Random		10.0 $\pm$ 3.5	20 / 5 (0)	5.3 $\pm$ 1.1	10 / 5 (0)		9.8 $\pm$ 6.8	30 / 5 (0)		5.0 $\pm$ 0.0	5 / 5 (0)		7.3 $\pm$ 4.0	20 / 5 (0)	
MCTS		32.3 $\pm$ 6.6	35 / 15 (17)	12.5 $\pm$ 7.5	20 / 5 (0)		33.3 $\pm$ 4.6	35 / 15 (16)		15.5 $\pm$ 8.4	35 / 5 (1)		32.8 $\pm$ 2.5	35 / 30 (11)	
OLETS		24.3 $\pm$ 9.5	35 / 10 (8)	5.0 $\pm$ 0.0	5 / 5 (0)		24.8 $\pm$ 7.5	35 / 15 (5)		5.0 $\pm$ 0.0	5 / 5 (0)		26.0 $\pm$ 6.6	35 / 10 (4)	
RHEA		25.0 $\pm$ 11.1	35 / 10 (10)	5.0 $\pm$ 0.0	5 / 5 (0)		27.0 $\pm$ 9.8	35 / 10 (9)		5.0 $\pm$ 0.0	5 / 5 (0)		25.5 $\pm$ 9.2	35 / 10 (9)	
RS		33.3 $\pm$ 5.3	35 / 15 (18)	11.8 $\pm$ 10.4	30 / 5 (0)		34.8 $\pm$ 1.1	35 / 30 (19)		10.0 $\pm$ 10.3	35 / 5 (2)		27.0 $\pm$ 10.4	35 / 5 (12)	

TABLE VII
WATERPUZZLE: AVERAGE SCORE, STANDARD DEVIATION AND NUMBER OF WINS (IN THE BRACKETS) OVER 20 INDEPENDENT TRIALS

Agent	Level-0 (15 / 0)			Level-1 (15 / 0)			Level-2 (15 / 0)			Level-3 (15 / 0)			Level-4 (15 / 0)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
DQN	G1	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G2	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G3	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	DORL	0.0 $\pm$ 0.0	0.0/0.0(0)	5.0 $\pm$ 0.0	5.0/5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
DQN _s	G1	0.0 $\pm$ 0.0	0.0/0.0(0)	2.0 $\pm$ 2.4	5.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G2	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.5 $\pm$ 1.5	5.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G3	15.0 $\pm$ 0.0	15.0/15.0(20)	4.2 $\pm$ 1.8	5.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		1.5 $\pm$ 4.5	15.0/0.0(2)	
	DORL	7.5 $\pm$ 4.3	15.0/5.0(5)	5.0 $\pm$ 0.0	5.0/5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.2 $\pm$ 1.1	5.0/0.0(0)	
PPO	G1	0.0 $\pm$ 0.0	0.0/0.0(0)	15.0 $\pm$ 0.0	15.0/15.0(20)		1.5 $\pm$ 2.3	5.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G2	4.8 $\pm$ 1.1	5.0/0.0(0)	15.0 $\pm$ 0.0	15.0/15.0(20)		8.5 $\pm$ 4.8	15.0/5.0(7)		0.0 $\pm$ 0.0	0.0/0.0(0)		4.5 $\pm$ 1.5	5.0/0.0(0)	
	G3	15.0 $\pm$ 0.0	15.0/15.0(20)	15.0 $\pm$ 0.0	15.0/15.0(20)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	DORL	15.0 $\pm$ 0.0	15.0/15.0(20)	7.5 $\pm$ 4.3	15.0/5.0(5)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
PPO _D	G1	0.0 $\pm$ 0.0	0.0/0.0(0)	5.0 $\pm$ 0.0	5.0/5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G2	15.0 $\pm$ 0.0	15.0/15.0(20)	12.5 $\pm$ 4.3	15.0/5.0(15)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G3	15.0 $\pm$ 0.0	15.0/15.0(20)	15.0 $\pm$ 0.0	15.0/15.0(20)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	DORL	15.0 $\pm$ 0.0	15.0/15.0(20)	15.0 $\pm$ 0.0	15.0/15.0(20)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
A2C	G1	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
	G2	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		5.0 $\pm$ 0.0	5.0/5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		4.8 $\pm$ 1.1	5.0/0.0(0)	
	G3	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		5.0 $\pm$ 0.0	5.0/5.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		5.0 $\pm$ 0.0	5.0/5.0(0)	
	DORL	0.0 $\pm$ 0.0	0.0/0.0(0)	0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
Ye <i>et al.</i> [21]	1.0 $\pm$ 3.0	15/0(1)	0.7 $\pm$ 1.7	5/0(0)			2.7 $\pm$ 3.3	15/0(1)		0.0 $\pm$ 0.0	0/0(0)		1.3 $\pm$ 2.2	5/0(0)	
Random	4.8 $\pm$ 6.2	15 / 0 (5)	1.0 $\pm$ 3.4	15 / 0 (1)			2.8 $\pm$ 3.7	15 / 0 (1)		0.3 $\pm$ 1.1	5 / 0 (0)		3.0 $\pm$ 4.6	15 / 0 (2)	
MCTS	9.0 $\pm$ 7.4	15 / 0 (12)	5.3 $\pm$ 6.0	15 / 0 (5)			6.5 $\pm$ 5.3	15 / 0 (5)		1.8 $\pm$ 3.6	15 / 0 (1)		3.5 $\pm$ 2.3	5 / 0 (0)	
OLETS	15.0 $\pm$ 0.0	15 / 15 (20)	11.3 $\pm$ 5.2	15 / 0 (13)			15.0 $\pm$ 0.0	15 / 15 (20)		9.8 $\pm$ 5.4	15 / 0 (10)		15.0 $\pm$ 0.0	15 / 15 (20)	
RHEA	7.0 $\pm$ 7.3	15 / 0 (9)	2.8 $\pm$ 4.6	15 / 0 (2)			7.8 $\pm$ 6.2	15 / 0 (8)		1.0 $\pm$ 2.0	5 / 0 (0)		3.5 $\pm$ 4.5	15 / 0 (2)	
RS	9.3 $\pm$ 7.1	15 / 0 (12)	2.5 $\pm$ 2.5	5 / 0 (0)			5.5 $\pm$ 5.9	15 / 0 (5)		0.3 $\pm$ 1.1	5 / 0 (0)		2.0 $\pm$ 2.5	5 / 0 (0)	

The bold numbers denotes the highest average scores among four input types or planning agents.




Fig. 3. Illustration of RL with dual-observation.

appear in the test levels, we find the closest tile with distance between reference vectors instead of simply finding the code with the same reference vector. Then, the tile-based grids are converted to two matrices of tile codes, i.e., integers. Finally, one-hot encoding is applied. Since there are no more than 16 tile types in each game, the size of final matrix is $\frac{h_a}{10} \times \frac{w_a}{10} \times 16$ for the transformed GO and $\frac{h_l}{10} \times \frac{w_l}{10} \times 16$ for the local one.

C. Network Architecture

We use a convolutional neural network (CNN). The CNN consists of three modules, namely Conv-G, Conv-L, and FC. The Conv-G and the Conv-L modules are used to process the encoded GO and LO, respectively.

For agents using pixel input, the Conv-G module is composed of three convolutional layers with $\langle \text{channel}, \text{kernel size},$

TABLE VIII
COMPARING DORL WITH PLANNING AGENTS ON THE 2021 COMPETITION GAMES

(a) *GreedyMouse*

Agent	Level-0 (98 / -40)			Level-1 (67 / -60)			Level-2 (98 / -50)			Level-3 (67 / -60)			Level-4 (77 / -70)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
<i>Arcane</i>	6.0 $\pm$ 2.6	8.0/1.0(0)		17.1 $\pm$ 6.0	33.0/9.0(0)		-0.6 $\pm$ 8.9	35.0/-10.0(0)		-0.8 $\pm$ 4.5	16.0/-5.0(0)		4.7 $\pm$ 12.6	40.0/-5.0(0)	
Random	0.8 $\pm$ 1.8	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		2.8 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		0.3 $\pm$ 1.1	5 / 0 (0)	
MCTS	2.3 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		7.8 $\pm$ 4.9	15 / 0 (6)		0.0 $\pm$ 0.0	0 / 0 (0)		0.5 $\pm$ 1.5	5 / 0 (0)	
OLETS	12.3 $\pm$ 5.6	15 / 0 (16)		11.0 $\pm$ 6.2	15 / 0 (14)		11.3 $\pm$ 5.9	15 / 0 (14)		10.0 $\pm$ 6.9	15 / 0 (13)		9.5 $\pm$ 5.0	15 / 5 (9)	
RHEA	2.0 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		5.3 $\pm$ 4.6	15 / 0 (3)		0.0 $\pm$ 0.0	0 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)	
RS	2.3 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		6.8 $\pm$ 4.3	15 / 0 (4)		0.0 $\pm$ 0.0	0 / 0 (0)		0.5 $\pm$ 1.5	5 / 0 (0)	

(b) *BraveKeeper*

Agent	Level-0 (100 / -40)			Level-1 (90 / -30)			Level-2 (100 / -40)			Level-3 (90 / -30)			Level-4 (90 / -30)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
<i>Arcane</i>	21.8 $\pm$ 17.2	50.0/0.0(4)		32.0 $\pm$ 27.36	75.0/5.0(3)		12.3 $\pm$ 9.4	35.0/0.0(1)		41.3 $\pm$ 26.4	75.0/5.0(3)		46.0 $\pm$ 23.9	75.0/5.0(12)	
Random	7.0 $\pm$ 5.6	20 / 0 (0)		5.0 $\pm$ 0.0	5 / 5 (0)		8.5 $\pm$ 6.9	30 / 0 (0)		5.3 $\pm$ 1.1	10 / 5 (0)		33.5 $\pm$ 16.7	55 / 10 (7)	
MCTS	49.0 $\pm$ 4.4	55 / 40 (20)		64.5 $\pm$ 10.7	85 / 45 (20)		48.5 $\pm$ 3.6	55 / 45 (20)		36.3 $\pm$ 20.1	70 / 15 (11)		62.0 $\pm$ 4.0	70 / 55 (20)	
OLETS	56.5 $\pm$ 5.3	65 / 45 (20)		59.8 $\pm$ 9.7	80 / 45 (20)		46.3 $\pm$ 8.9	65 / 35 (20)		49.3 $\pm$ 8.8	65 / 40 (20)		66.3 $\pm$ 6.5	80 / 55 (20)	
RHEA	34.0 $\pm$ 10.6	55 / 0 (17)		33.5 $\pm$ 17.5	55 / 5 (10)		43.5 $\pm$ 5.7	55 / 30 (19)		27.5 $\pm$ 18.8	55 / 5 (7)		56.8 $\pm$ 12.3	75 / 10 (18)	
RS	44.5 $\pm$ 5.9	60 / 35 (20)		28.5 $\pm$ 15.7	55 / 5 (9)		33.0 $\pm$ 14.7	45 / -5 (15)		25.5 $\pm$ 22.3	65 / 5 (6)		61.0 $\pm$ 7.0	75 / 50 (18)	

(c) *TrappedHero*

Agent	Level-0 (15 / 0)			Level-1 (15 / 0)			Level-2 (15 / 0)			Level-3 (15 / 0)			Level-4 (15 / 0)		
	Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)		Avg. $\pm$ Std	Max/Min(Wins)	
<i>Arcane</i>	15.0 $\pm$ 0.0	15.0/15.0(20)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)		0.0 $\pm$ 0.0	0.0/0.0(0)	
Random	0.8 $\pm$ 1.8	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		2.8 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		0.3 $\pm$ 1.1	5 / 0 (0)	
MCTS	2.3 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		7.8 $\pm$ 4.9	15 / 0 (6)		0.0 $\pm$ 0.0	0 / 0 (0)		0.5 $\pm$ 1.5	5 / 0 (0)	
OLETS	12.3 $\pm$ 5.6	15 / 0 (16)		11.0 $\pm$ 6.2	15 / 0 (14)		11.3 $\pm$ 5.9	15 / 0 (14)		10.0 $\pm$ 6.9	15 / 0 (13)		9.5 $\pm$ 5.0	15 / 5 (9)	
RHEA	2.0 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		5.3 $\pm$ 4.6	15 / 0 (3)		0.0 $\pm$ 0.0	0 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)	
RS	2.3 $\pm$ 2.5	5 / 0 (0)		0.0 $\pm$ 0.0	0 / 0 (0)		6.8 $\pm$ 4.3	15 / 0 (4)		0.0 $\pm$ 0.0	0 / 0 (0)		0.5 $\pm$ 1.5	5 / 0 (0)	

Average score and standard deviation over 20 independent runs on testing and training levels are reported, as well as the number of wins shown in the brackets. The bold numbers denotes the highest average scores among planning agents.




Fig. 4. Illustration of transformed game observations.

stride of $\langle 32, 8, 4 \rangle$, $\langle 64, 4, 3 \rangle$, and $\langle 64, 3, 1 \rangle$. The Conv-L module is composed of two convolutional layers with $\langle \text{channel}, \text{kernel size}, \text{stride} \rangle$ of $\langle 32, 6, 4 \rangle$, $\langle 64, 4, 2 \rangle$. We consider agents using one-hot input and pixel input separately.

For agents using one-hot input, the Conv-G module is composed of two convolutional layers with $\langle \text{channel}, \text{kernel size}, \text{stride} \rangle$ of $\langle 32, 4, 3 \rangle$ and $\langle 64, 3, 2 \rangle$. The Conv-L module has only one convolutional layer with $\langle \text{channel}, \text{kernel size}, \text{stride} \rangle$ of $\langle 32, 3, 1 \rangle$.

The Conv-G module and the Conv-L module of agents using pixel input and one-hot input are different because the dimensionality of pixel input is larger than the one-hot input, so it

requires much more down-sampling for the model of agents using pixel input.

All of those convolutional layers use ReLU function as activation. Both Conv-G's and Conv-L's outputs are reshaped to a vector and then fully connected to a linear layer with 256 outputs. The FC module has two $512 \times 64 \rightarrow 64 \times \#actions$ linear layers with ReLU activation function and no activation function respectively, where $\#actions$ denotes the number of actions. For other agents with GO only in this article, the Conv-L module is removed and the first layer of FC just uses a 256×64 linear layer.

D. Decision Policy

Action selection strategy is a core component of RL algorithms. It directly affects the experience transitions got by agents. When training DQN agents, we use the ϵ -greedy strategy with linearly decreasing ϵ value. while testing, the following two decision policies have been considered.

The first one uses a greedy approach, i.e., a DQN agent takes action with the highest Q-value. It is referred to as the *deterministic policy* in our paper. A greedy decision policy can clearly show the training effect of agents. However, sometimes it may be too strict and rigid for an agent.

To encourage exploration, we also design a *stochastic policy* to make DQN agents not always select the action with the highest

Q-value but also the “suboptimal” ones. However, the ordinary epsilon greedy that controls exploration probability proportional to Q-value is not suitable in our case since playing an action shown to be insufficient during training is more likely to be an impaired decision on test levels. Therefore, we design a “scaled softmax” that decreases the probability of selecting actions with very low Q-values, described as follows. First, a probability distribution is computed over actions $\mathcal{A}$ based on their Q-values $\mathbf{Q}^*$. For any action $a \in \mathcal{A}$, its probability of being selected is defined as

$$p_a = \frac{\exp(k \cdot q'_a)}{\sum_{i \in \mathcal{A}} \exp(k \cdot q'_i)},$$

where q' denotes the Q-value shifted by $-(\max \mathbf{Q}^* + \min \mathbf{Q}^*)/2$, and $k = \sigma/(\max \mathbf{Q}^* - \min \mathbf{Q}^*)$ to avoid overflow. σ is a control parameter. Then, an action randomly selected according to the resulting probability distribution is played. The parameter σ is introduced to control how greedy the strategy is. The higher the σ is, the higher selection probability the action with maximum Q-value has.

E. Novelty of This Work

It is worth mentioning that the work of [21], which has some similarities with ours, only used the raw image of either the global or the LO as a single input. DORL not only uses both inputs simultaneously, but also adapts the structure of the neural network to take tile-vector encoded observations as inputs. Moreover, only A2C was considered in [21], while DORL is a more general technique, and its instantiations using several different state-of-the-art RL algorithms, including A2C, DQN, and PPO, are implemented and compared in Section IV. The work of [11] and [13] used image transformation as a data augmentation technique during training, but their agents were trained and tested on identical game levels thus the test levels are also training levels. However, in our case, the test levels are unseen during training. Moreover, our proposed technique pre-processes the inputs with the novel tile-vector encoding method aiming at accelerating the process of encoding observations and handling the possible appearance of new tiles in unseen test levels.

IV. EXPERIMENTAL STUDY

DORL is implemented using three different state-of-the-art RL algorithms (PPO, A2C, and DQN) to validate its performance. For the purpose of examining the benefit brought by different components of DORL, for each RL algorithm, four groups of agents categorized by input format (one-hot or pixel input; GO alone or dual-observation), as summarized in Table III, have been considered and compared on the 2020 GVGAI Learning Competition games (see Sections IV-B and IV-C). The PPO agent with DORL is used as a baseline agent in the 2021 GVGAI Learning Competition (see Section IV-D). Experimental details are provided in Section IV-A.




Fig. 5. Average scores obtained by PPO agents on *GoldDigger*'s training levels over 20 independent trials each. (a) *GoldDigger-0*. (b) *GoldDigger-1*.

A. Experimental Setting and Baseline

These four groups of agents (cf. Table III) were trained separately with DQN, PPO, and A2C, adapted from stable-baselines3 [22]. All DQN agents were tested with the deterministic policy and the stochastic policy, denoted as “DQN” and “DQN_s,” respectively.

As a baseline, the agent of Ye *et al.* [21] is applied with all its three operations, *rotation*, *transformation*, and *crop*, and trained with its default parameters, network structure, and pixel input [21]. Additionally, the performance of four planning agents given forward models of games, including Monte Carlo Tree Search (MCTS) [33], Open Loop Expectimax Tree Search (OLETS) [33], Rolling Horizon Evolution (RHEA) [34], and Random Search (RS), are also reported.

All learning agents were trained with the same procedure on the same machine with an Intel Xeon Gold 6240 CPU and four TITAN RTX GPUs. A learning agent is trained for 1 000 000 timesteps for each game on the two randomly alternated training levels. Three mono-agents are trained to build a meta-agent. When testing, the meta-agent first determines the game by its screen size and then deploys the corresponding mono-agent. All test experiments are repeated independently 20 times. Experimental setting and network configuration are summarized in Table IV.

B. Comparing Agents on the 2020 Competition Games

The average score over 20 independent runs, the highest score, the lowest score, and the number of wins obtained by each trained agent on the training levels and another 20 trials after training on the test levels of the 2020 competition games, *GoldDigger*, *TreasureKeeper*, and *WaterPuzzle*, are provided in Tables V–VII, respectively. Wilcoxon rank-sum test has been performed on the game scores reported in Tables V–VII, and the results are provided in *Appendix*. As an illustrative example, Fig. 5 shows the training curves of the four PPO agents with different input types (G1, G2, G3, and DORL as indicated in Table III).

By comparing the performance of instantiations of G1, G2, G3, and DORL, implemented with PPO, A2C, and DQN, we observe outstanding performance of DORL when playing all the test levels of *GoldDigger* and *TreasureKeeper*. Highlights are as follows.

- 1) Any algorithm with dual-observation and one-hot encoding (i.e., agents shown on rows or columns entitled with “DORL”) obtains significantly higher score than or similar score to its version using other observation or pixel-based encoding (referred to as G1, G2, and G3) when playing all test levels of *GoldDigger* and *TreasureKeeper*. When playing *WaterPuzzle*, using dual-observation is not always helpful.
- 2) DQN_s-DORL and PPO agents perform significantly better than or similar to the agent of Ye *et al.* [21] in terms of game score and number of wins in all the test levels of *GoldDigger* and *TreasureKeeper*, while this is not always the case for *WaterPuzzle* (cf. Table XI in Appendix).
- 3) Using a stochastic policy for DQN does not always lead to an overall better performed agent. On one hand, DQN_s obtains statistically higher or similar game score compared to DQN with the same input observation in 7 out of 9 test levels (i.e., 3 test levels for each game). On the other hand, DQN agent wins *TreasureKeeper* more times than DQN_s with the same input observation.

More discussions on the agents’ performance on individual games are provided as follows.

1) *GoldDigger With Risky Reward*: According to Table V and the results of statistical tests provided in Appendix, all agents with dual-observation (G2 and DORL) perform significantly better than or similar to those using GO only (G1 and G3) in playing training levels and test levels. Among those four types of input, DORL agents implemented with DQN and PPO usually holds the first place. For example, DQN agent with DORL achieves the highest score among all five levels. The performance of A2C agents is not stable enough due to the high variance of gradients. Any agent with DORL obtains significantly higher or similar game score compared to its version using other observation or pixel-based encoding.

A notable observation of DQN agents is the disparate results of using the deterministic policy and the stochastic policy, especially in *Level-2*. The former obtains -1.2 and the latter obtains 105.8 as an average score. *Level-2* is adapted from *Level-0* by adding a wall (cf. Fig. 2), making the avatar trap at a corner. The stochastic policy helps the avatar to explore the space and escape from the corner. All planning agents perform well in this game, but the random agent plays badly.

Despite the high average scores of DORLs in playing *GoldDigger*, they sometimes get deficient scores which can be determined by the high variance shown in Table V. PPO agent with DORL obtains 67.8 in *Level-0* and 66.3 in *Level-2* as standard deviations of scores. It may be explained by the fact that the agents sometimes attempt to approach a monster to attack it. However, due to the random behaviors of the monsters, the agent could be killed before attacking. DORLs rarely win since collecting all the jewels is hard. On the other hand, it may be blamed on the drawback of traditional RL algorithms [35]. RL agents usually pay more attention to getting a higher long-term expected reward, but sometimes ignore risks caused by some uncertainties, e.g., agents bias more to gain score by attacking monsters but ignore the risk of being killed due to the direct use

of the game score as the reward. Designing a suitable reward function for efficiently training a general agent is crucial.

2) *Deceptive TreasureKeeper*: As shown in Table VI, most of the agents can obtain the maximum score (35) from *Level-0* to *Level-4*. According to the results of statistical tests provided in Appendix, DQN agents with dual-observation (G2 and DORL) obtain significantly higher or similar scores than the ones with the GO only (G1 and G3) when playing test levels. When using A2C as the base algorithm, the agent using tile-vector GO only, thus G3, beats all the other agents. In particular, among all the agents, only A2C-G3 can win *Level-1* and *Level-3* with the maximum score at *TreasureKeeper*.

Since *TreasureKeeper* is a survival game, the game score increases along with the survival time of the avatar. Some core actions that make an avatar survive longer (e.g., pushing a box on purpose to block monsters), do not directly lead to any immediate reward. Random decisions can also result in such actions. Nevertheless, the game score occasionally happens to increase when the avatar is pushing a box. Such coincidences easily confuse a learning agent and encourage it to push boxes more frequently, which does not usually help guard the treasure and earn the reward. In such a situation, even if an agent has learnt how to avoid monsters, the treasure might still be attacked by monsters, which leads to a failure of the game. Planning agents can hardly win *Level-1* and its variation *Level-3*.

3) *Simple Yet Challenging WaterPuzzle*: All PPO agents can achieve acceptable average scores, especially G3 and DORL which get average scores of 15, over 20 trials of playing training levels (cf. Table VII). However, their performance on test levels is disappointing. A2C and DQN agents rarely win the games. No agent can get any score in *Level-3*, and a few agents, including PPO agents with G1/G2, and A2C agents with G2/G3 can get some points in *Level-2* and *Level-4*. *WaterPuzzle* is challenging for planning agents. Most of the planning agents, except OLETS, rarely win the game.

The puzzle game *WaterPuzzle* is extremely easy for human players but hard for RL agents due to the sparse reward. Some agents learned how to win the game at training levels but failed in test levels. It is probably because what agents actually learned is the winning path in a specific map.

C. Further Discussion

Agents with dual-observation show promising results for playing 2020 competition games. Some remarkable phenomena are discussed as follows.

1) *Dual-Observation Improves Generalization*: Agents enhanced by dual-observation have an advantage over other agents, especially in playing *GoldDigger*. Comparing their performance in playing training levels and unseen test levels, agents with dual-observation usually achieve superior results to others in all three games. Although those unseen levels have different layouts, they share the same game rules. Taking *GoldDigger* as an example, an agent gains reward by approaching and digging diamonds. Agents need to pay more attention to the close elements, such as monsters, diamonds, and treasures. This playing strategy remains the same in training levels and unseen

test levels. With the help of dual-observation, knowledge gained in playing training levels enhances the agents' ability to play unseen test levels.

The main benefit brought by using dual-observation is that an agent can receive the whole global information and at the same time collect possibly more important local information around it. Using the GO alone treats all elements equally, while using the LO alone is myopic and ignores the information in further area. An agent can consider different areas with different weights when using this dual-observation. In other words, it focuses on the core information that may affect the gained reward since the actions that an agent takes usually have limited effect on the GO but are more meaningful from local view without losing the reception of remote areas that produce possible future rewards from global view. It can be attributed to the coordination of the global and LOs that the agents with dual-observation usually get better results than others at a new level. Even the layout changes in new levels, there is still some familiar local information in the view of the agents.

2) *Tile-Vector Encoding Helps*: Agents using pixel input (G1 and G2) usually perform inferior to those using tile-vector encoded input (cf. Tables V–VII). Taking *GoldDigger* as an example, PPO agent with G3 beats G1 for all levels. A2C and DQN_s meet the same situation. Using tile-vector encoding, each RGB image is converted into $\frac{h}{10} \times \frac{w}{10} \times 16$ one-hot encoded matrices using a precalculated tile dictionary, as described in Section III-B. It reduces the state space in representation. The compressed state space in representation helps our agents learn better within the same training budget.

At the same time, we have to admit that one-hot has a restricted scope of use. Since the three games we use here all act with a speed of one tile per game tick, it is easy for us to convert the image to one-hot matrix. Nevertheless, if it is not the case, one-hot matrix may not be converted easily.

3) *Deterministic Versus Stochastic Decision-Making*: Agents with the stochastic policy and the deterministic policy perform very differently, particularly in *GoldDigger* and *WaterPuzzle*. DQN agents with the deterministic policy play badly in test levels, while a much higher score is obtained with the stochastic policy for the same levels. The reason here is that the deterministic policy traps the agent in one state, such as moving forward to a wall, which makes no sense. With the stochastic policy, the agent can get away from those states. However, the stochastic policy works worse than the deterministic one in *TreasureKeeper*. Possibly, the stochastic policy makes agents move forward to the monsters mistakenly.

The deterministic policy and stochastic policy for deciding the action to play have their advantages. The stochastic policy can avoid the agent being trapped in one state, but if the penalty due to mistakes is too high, the deterministic one may be a better choice. For example, in *GoldDigger-2*, the agent with the greedy strategy may trap in the wall if it does not choose an “attack” action due to a relatively low Q-value. By adding some randomness during action selection, the agent can possibly get rid of this situation.

4) *Training With Random Starts*: To improve the generalization ability of learning agents, we apply a simple data addition technique to PPO agent, denoted as PPO_D. We keep the layout of

training level maps unchanged during training, but a new initial position is given to the avatar when launching a new episode. Thus, the agent keeps playing the same training level maps but with different initial positions.

According to Tables V–VII, generating more levels using this simple technique for training a PPO agent does not lead to a better performed agent. PPO_D obtains statistically higher game score than PPO with the same input observation in playing two test levels of three games and statistically lower score than PPO with the same input observation in playing five test levels of three games. PPO-DORL performs worse with the aforementioned random start technique, but dual-observation still helps. When playing test levels of *GoldDigger*, PPO_D with dual-observation still obtains better results than using GO only, comparing G2 with G1, and DORL with G3.

The PPO_D agents show stable performance in playing training levels, especially for *WaterPuzzle*, where they can get an average reward of 15. But those agents can not get any points in playing test levels.

D. Arcane: PPO-DORL Playing 2021 Competition Games

PPO-DORL is used as a baseline agent in the 2021 competition edition, named as *Arcane* and trained on the competition games, *GreedyMouse*, *BraveKeeper*, and *TrappedHero*, as described in Section IV-A. As no competition submission was received, *Arcane* is also compared to planning agents in Table VIII.

Arcane achieves the higher max score in four levels of *GreedyMouse* compared with OLETS and beats other planning agents regarding average score in *Level-0*, *Level-1*, and *Level-4*. When playing *BraveKeeper*, *Arcane* is not stable enough, although its max scores are close to OLETS' scores in *Level-1*, *Level-3*, and *Level-4*. *Arcane* achieves a win rate of 100% in playing *TrappedHero-0*, higher than all the planning agents. However, it failed to earn any score in the other levels.

V. CONCLUSION

In this article, we first summarize the five years' GVGAI Learning Competition editions, in particular, the latest ones in 2019, 2020, and 2021. Viewing the poor performance of competition entries, we design a novel RL technique with dual-observation (DORL) for GVGAI.

DORL simultaneously takes the tile-vector encoded, transformed GO, and LO of the game screen as input, aiming at learning local information, which may exist in unseen games or levels during training. As a general technique, DORL is applied to DQN, A2C, and PPO separately and compared in playing the games of the 2020 GVGAI Learning Competition. Two versions of DQN-DORL with the stochastic policy and deterministic policy for deciding actions to play during test, respectively, are performed. DORL shows its superior performance on the game set of the 2020 GVGAI Learning Competition, thanks to the introduction of LO. Additionally, the training dataset is enlarged with new levels generated by hanging the avatar's initial position. We observe that applying this simple data addition technique is

TABLE IX
RESULTS OF WILCOXON RANK-SUM TEST OF AGENTS' PERFORMANCE ON *GOLDDIGGER* REPORTED IN TABLE V

Agent	DQN				DQN _s				Ye <i>et al.</i> [21]	
	G1	G2	G3	DORL	G1	G2	G3	DORL		
DQN	G1	--	0/2/1 (0\0)	0/2/1 (0\0)	0/1/2 (0\0)	0/2/1 (0\0)	0/0/3 (0\0)	0/0/3 (0\0)	0/0/3 (0\1)	0/1/2 (0\0)
	G2	1/2/0 (0\0)	--	1/1/1 (0\0)	0/2/1 (0\0)	1/1/1 (0\0)	0/1/2 (0\0)	0/1/2 (0\0)	0/1/2 (0\1)	1/0/2 (0\0)
	G3	1/2/0 (0\0)	1/1/1 (0\0)	--	0/1/2 (0\0)	0/2/1 (0\0)	0/0/3 (0\0)	0/0/3 (0\0)	0/0/3 (0\1)	1/1/1 (0\0)
	DORL	2/1/0 (0\0)	1/2/0 (0\0)	2/1/0 (0\0)	--	2/1/0 (0\0)	0/2/1 (0\0)	0/1/2 (0\0)	0/1/2 (0\1)	2/0/1 (0\0)
DQN _s	G1	1/2/0 (0\0)	1/1/1 (0\0)	1/2/0 (0\0)	0/1/2 (0\0)	--	0/0/3 (0\0)	0/0/3 (0\0)	0/0/3 (0\1)	1/1/1 (0\0)
	G2	3/0/0 (0\0)	2/1/0 (0\0)	3/0/0 (0\0)	1/2/0 (0\0)	3/0/0 (0\0)	--	0/3/0 (0\0)	0/2/1 (0\1)	3/0/0 (0\0)
	G3	3/0/0 (0\0)	2/1/0 (0\0)	3/0/0 (0\0)	2/1/0 (0\0)	3/0/0 (0\0)	0/3/0 (0\0)	--	0/2/1 (0\1)	3/0/0 (0\0)
	DORL	3/0/0 (1\0)	2/1/0 (1\0)	3/0/0 (1\0)	2/1/0 (1\0)	3/0/0 (1\0)	1/2/0 (1\0)	1/2/0 (1\0)	--	3/0/0 (1\0)
Ye <i>et al.</i> [21]	2/1/0 (0\0)	2/0/1 (0\0)	1/1/1 (0\0)	1/0/2 (0\0)	1/1/1 (0\0)	0/0/3 (0\0)	0/0/3 (0\0)	0/0/3 (0\1)	--	

Agent	PPO				PPO _D				Ye <i>et al.</i> [21]	
	G1	G2	G3	DORL	G1	G2	G3	DORL		
PPO	G1	--	0/1/2 (0\0)	0/2/1 (0\0)	0/0/3 (0\0)	1/2/0 (0\0)	0/0/3 (0\0)	0/1/2 (0\0)	0/1/2 (0\1)	1/2/0 (0\0)
	G2	2/1/0 (0\0)	--	0/3/0 (0\0)	0/2/1 (0\0)	1/2/0 (0\0)	0/1/2 (0\0)	0/3/0 (0\0)	0/3/0 (0\1)	2/1/0 (0\0)
	G3	1/2/0 (0\0)	0/3/0 (0\0)	--	0/2/1 (0\0)	2/1/0 (0\0)	0/1/2 (0\0)	0/3/0 (0\0)	0/3/0 (0\1)	2/1/0 (0\0)
	DORL	3/0/0 (0\0)	1/2/0 (0\0)	1/2/0 (0\0)	--	3/0/0 (0\0)	0/3/0 (0\0)	2/1/0 (0\0)	0/3/0 (0\1)	3/0/0 (0\0)
PPO _D	G1	0/2/1 (0\0)	0/2/1 (0\0)	0/1/2 (0\0)	0/0/3 (0\0)	--	0/0/3 (0\0)	0/2/1 (0\0)	0/2/1 (0\1)	1/2/0 (0\0)
	G2	3/0/0 (0\0)	2/1/0 (0\0)	2/1/0 (0\0)	0/3/0 (0\0)	3/0/0 (0\0)	--	2/1/0 (0\0)	0/3/0 (0\1)	3/0/0 (0\0)
	G3	2/1/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/1/2 (0\0)	1/2/0 (0\0)	0/1/2 (0\0)	--	0/3/0 (0\1)	3/0/0 (0\0)
	DORL	2/1/0 (1\0)	0/3/0 (1\0)	0/3/0 (1\0)	0/3/0 (1\0)	1/2/0 (1\0)	0/3/0 (1\0)	0/3/0 (1\0)	--	2/1/0 (1\0)
Ye <i>et al.</i> [21]	0/2/1 (0\0)	0/1/2 (0\0)	0/1/2 (0\0)	0/0/3 (0\0)	0/2/1 (0\0)	0/0/3 (0\0)	0/0/3 (0\0)	0/1/2 (0\1)	--	

Agent	A2C				Ye <i>et al.</i> [21]
	G1	G2	G3	DORL	
A2C	G1	--	0/1/2 (0\0)	0/1/2 (0\0)	0/2/1 (0\0)
	G2	2/1/0 (0\0)	--	0/1/2 (0\0)	1/2/0 (0\0)
	G3	2/1/0 (0\0)	2/1/0 (0\0)	--	2/1/0 (0\0)
	DORL	2/1/0 (0\0)	1/1/1 (0\0)	1/2/0 (0\0)	2/0/1 (0\0)
Ye <i>et al.</i> [21]	1/2/0 (0\0)	0/2/1 (0\0)	0/1/2 (0\0)	1/0/2 (0\0)	--

“+/-/” indicates the number of times that the row agent performs statistically better/similar/worse to the column agent, respectively. The number of wins of the row agent and the column agent for three levels are shown in the brackets.

not always helpful. DORL implemented with PPO (named as *Arcane*), is used as a baseline agent in the 2021 competition.

In this article, for each game, an agent is trained separately as a base agent. A meta-agent is then composed of those trained base agents without knowledge sharing for playing different games. As future work, we will investigate knowledge transfer across games.

APPENDIX

Tables IX–XI provide the results of Wilcoxon rank-sum test performed on the game scores reported in Tables V–VII for comparing agents' performance in playing test levels of *GoldDigger*, *TreasureKeeper*, and *WaterPuzzle*. “+/-/” indicates the number of times that the row agent obtains statistically higher/similar/lower game score compared to the column agent, respectively. The numbers of wins of the row agent and the column agent are repeated in the brackets. Highlights are as follows.

- 1) Any algorithm with dual-observation and one-hot encoding (i.e., agents shown on rows or columns entitled with “DORL”) obtains significantly higher score than or similar score to its version using other observation or pixel-based

encoding (referred to as G1, G2, and G3) when playing all test levels of *GoldDigger* and *TreasureKeeper*. When playing *WaterPuzzle*, using dual-observation is not always helpful.

- 2) DQN_s-DORL and PPO agents perform significantly better than or similar to the agent of Ye *et al.* [21] in terms of game score and number of wins in all the test levels of *GoldDigger* and *TreasureKeeper*, while this is not always the case for *WaterPuzzle* (cf. Table XI).
- 3) Using a stochastic policy for DQN does not always lead to an overall better performed agent. On one hand, DQN_s obtains statistically higher or similar game score compared to DQN with the same input observation in 7 out of 9 test levels (i.e., 3 test levels for each game). On the other hand, DQN agent wins *TreasureKeeper* more times than DQN_s with the same input observation.
- 4) Generating more levels for training a PPO agent does not lead to a better performed agent. PPO_D obtains statistically higher game score than PPO with the same input observation in playing two test levels of three games and statistically lower score than PPO with the same input observation in playing five test levels of three games.

TABLE X
RESULTS OF WILCOXON RANK-SUM TEST OF AGENTS' PERFORMANCE ON *TREASUREKEEPER* REPORTED IN TABLE VI

Agent	DQN				DQN _s				Ye <i>et al.</i> [21]	
	G1	G2	G3	DORL	G1	G2	G3	DORL		
DQN	G1	--	0/2/1 (10\20)	0/3/0 (10\14)	0/3/0 (10\21)	0/3/0 (10\7)	0/2/1 (10\10)	0/3/0 (10\7)	0/3/0 (10\19)	2/1/0 (10\0)
	G2	1/2/0 (20\10)	--	1/2/0 (20\14)	0/3/0 (20\21)	3/0/0 (20\7)	0/3/0 (20\10)	2/1/0 (20\7)	0/3/0 (20\19)	3/0/0 (20\0)
	G3	0/3/0 (14\10)	0/2/1 (14\20)	--	0/3/0 (14\21)	1/2/0 (14\7)	0/2/1 (14\10)	1/2/0 (14\7)	0/3/0 (14\19)	2/1/0 (14\0)
	DORL	0/3/0 (21\10)	0/3/0 (21\20)	0/3/0 (21\14)	--	1/2/0 (21\7)	0/3/0 (21\10)	1/2/0 (21\7)	0/3/0 (21\19)	2/1/0 (21\0)
DQN _s	G1	0/3/0 (7\10)	0/0/3 (7\20)	0/2/1 (7\14)	0/2/1 (7\21)	--	0/2/1 (7\10)	0/3/0 (7\7)	0/2/1 (7\19)	2/1/0 (7\0)
	G2	1/2/0 (10\10)	0/3/0 (10\20)	1/2/0 (10\14)	0/3/0 (10\21)	1/2/0 (10\7)	--	1/2/0 (10\7)	0/3/0 (10\19)	3/0/0 (10\0)
	G3	0/3/0 (7\10)	0/1/2 (7\20)	0/2/1 (7\14)	0/2/1 (7\21)	0/3/0 (7\7)	0/2/1 (7\10)	--	0/2/1 (7\19)	2/1/0 (7\0)
	DORL	0/3/0 (19\10)	0/3/0 (19\20)	0/3/0 (19\14)	0/3/0 (19\21)	1/2/0 (19\7)	0/3/0 (19\10)	1/2/0 (19\7)	--	2/1/0 (19\0)
Ye <i>et al.</i> [21]	0/1/2 (0\10)	0/0/3 (0\20)	0/1/2 (0\14)	0/1/2 (0\21)	0/1/2 (0\7)	0/0/3 (0\10)	0/1/2 (0\7)	0/1/2 (0\19)	--	

Agent	PPO				PPO _D				Ye <i>et al.</i> [21]	
	G1	G2	G3	DORL	G1	G2	G3	DORL		
PPO	G1	--	0/3/0 (5\7)	0/2/1 (5\9)	0/3/0 (5\4)	0/3/0 (5\8)	0/3/0 (5\8)	0/3/0 (5\6)	0/2/1 (5\13)	1/2/0 (5\0)
	G2	0/3/0 (7\5)	--	0/3/0 (7\9)	0/3/0 (7\4)	0/3/0 (7\8)	1/2/0 (7\8)	1/2/0 (7\6)	0/3/0 (7\13)	3/0/0 (7\0)
	G3	1/2/0 (9\5)	0/3/0 (9\7)	--	0/3/0 (9\4)	0/3/0 (9\8)	0/3/0 (9\8)	0/3/0 (9\6)	0/3/0 (9\13)	2/1/0 (9\0)
	DORL	0/3/0 (4\5)	0/3/0 (4\7)	0/3/0 (4\9)	--	0/3/0 (4\8)	0/3/0 (4\8)	0/3/0 (4\6)	0/3/0 (4\13)	2/1/0 (4\0)
PPO _D	G1	0/3/0 (8\5)	0/3/0 (8\7)	0/3/0 (8\9)	0/3/0 (8\4)	--	0/3/0 (8\8)	0/3/0 (8\6)	0/3/0 (8\13)	2/1/0 (8\0)
	G2	0/3/0 (8\5)	0/2/1 (8\7)	0/3/0 (8\9)	0/3/0 (8\4)	0/3/0 (8\8)	--	0/3/0 (8\6)	0/3/0 (8\13)	2/1/0 (8\0)
	G3	0/3/0 (6\5)	0/2/1 (6\7)	0/3/0 (6\9)	0/3/0 (6\4)	0/3/0 (6\8)	0/3/0 (6\8)	--	0/3/0 (6\13)	2/1/0 (6\0)
	DORL	1/2/0 (13\5)	0/3/0 (13\7)	0/3/0 (13\9)	0/3/0 (13\4)	0/3/0 (13\8)	0/3/0 (13\8)	0/3/0 (13\6)	--	2/1/0 (13\0)
Ye <i>et al.</i> [21]	0/2/1 (0\5)	0/0/3 (0\7)	0/1/2 (0\9)	0/1/2 (0\4)	0/1/2 (0\8)	0/1/2 (0\8)	0/1/2 (0\6)	0/1/2 (0\13)	--	

Agent	A2C				Ye <i>et al.</i> [21]
	G1	G2	G3	DORL	
A2C	G1	--	0/3/0 (4\13)	0/0/3 (4\22)	0/3/0 (4\11)
	G2	0/3/0 (13\4)	--	0/2/1 (13\22)	0/3/0 (13\11)
	G3	3/0/0 (22\4)	1/2/0 (22\13)	--	2/1/0 (22\11)
	DORL	0/3/0 (11\4)	0/3/0 (11\13)	0/1/2 (11\22)	--
Ye <i>et al.</i> [21]	0/1/2 (0\4)	0/1/2 (0\13)	0/0/3 (0\22)	0/1/2 (0\11)	--

"+/~/-" indicates the number of times that the row agent performs statistically better/similar/worse to the column agent, respectively. The number of wins of the row agent and the column agent for three levels are shown in the brackets

TABLE XI
RESULTS OF WILCOXON RANK-SUM TEST OF AGENTS' PERFORMANCE ON *WATERPUZZLE* REPORTED IN TABLE VII

Agent	DQN				DQN _s				Ye <i>et al.</i> [21]
	G1	G2	G3	DORL	G1	G2	G3	DORL	
DQN	G1	--	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
	G2	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
	G3	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
	DORL	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
DQN _s	G1	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
	G2	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\2)	0/3/0 (0\0)	0/1/2 (0\1)
	G3	0/3/0 (2\0)	0/3/0 (2\0)	0/3/0 (2\0)	0/3/0 (2\0)	0/3/0 (2\0)	--	0/3/0 (2\0)	0/2/1 (2\1)
	DORL	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\2)	--	0/2/1 (0\1)
Ye <i>et al.</i> [21]	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	1/2/0 (1\2)	2/1/0 (1\0)	--

Agent	PPO				PPO _D				Ye <i>et al.</i> [21]
	G1	G2	G3	DORL	G1	G2	G3	DORL	
PPO	G1	--	0/1/2 (0\7)	1/2/0 (0\0)	1/2/0 (0\0)	1/2/0 (0\0)	1/2/0 (0\0)	1/2/0 (0\0)	0/2/1 (0\1)
	G2	2/1/0 (7\0)	--	2/1/0 (7\0)	2/1/0 (7\0)	2/1/0 (7\0)	2/1/0 (7\0)	2/1/0 (7\0)	2/1/0 (7\1)
	G3	0/2/1 (0\0)	0/1/2 (0\7)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/1/2 (0\1)
	DORL	0/2/1 (0\0)	0/1/2 (0\7)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/1/2 (0\1)
PPO _D	G1	0/2/1 (0\0)	0/1/2 (0\7)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/1/2 (0\1)
	G2	0/2/1 (0\0)	0/1/2 (0\7)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/3/0 (0\0)	0/1/2 (0\1)
	G3	0/2/1 (0\0)	0/1/2 (0\7)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/3/0 (0\0)	0/1/2 (0\1)
	DORL	0/2/1 (0\0)	0/1/2 (0\7)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	0/3/0 (0\0)	--	0/1/2 (0\1)
Ye <i>et al.</i> [21]	1/2/0 (1\0)	0/1/2 (1\7)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	2/1/0 (1\0)	--

Agent	A2C				Ye <i>et al.</i> [21]
	G1	G2	G3	DORL	
A2C	G1	--	0/1/2 (0\0)	0/1/2 (0\0)	0/3/0 (0\0)
	G2	2/1/0 (0\0)	--	0/3/0 (0\0)	2/1/0 (0\0)
	G3	2/1/0 (0\0)	0/3/0 (0\0)	--	2/1/0 (0\0)
	DORL	0/3/0 (0\0)	0/1/2 (0\0)	0/1/2 (0\0)	--
Ye <i>et al.</i> [21]	2/1/0 (1\0)	0/1/2 (1\0)	0/1/2 (1\0)	2/1/0 (1\0)	--

"+/~/-" indicates the number of times that the row agent performs statistically better/similar/worse to the column agent, respectively. The number of wins of the row agent and the column agent for three levels are shown in the brackets.

ACKNOWLEDGMENT

The authors would like to thank Ye *et al.* [21] for helping with correctly executing their agent, the whole GVGAI team for implementing and maintaining the framework. And also like to thank Mr. Yang Tao for designing the 2020 competition games, and all the participants to the GVGAI Learning Competition editions for their participation.

REFERENCES

- [1] D. Silver *et al.*, “Mastering the game of go with deep neural networks and tree search,” *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [2] D. Silver *et al.*, “Mastering the game of go without human knowledge,” *Nature*, vol. 550, pp. 354–359, Oct. 2017.
- [3] J. Schrittwieser *et al.*, “Mastering Atari, go, chess and shogi by planning with a learned model,” *Nature*, vol. 588, pp. 604–609, 2020.
- [4] L. Kaiser *et al.*, “Model-based reinforcement learning for Atari,” in *Proc. 8th Int. Conf. Learn. Representations*, 2020, pp. 1–28. [Online]. Available: <https://openreview.net/pdf?id=S1xCPJHtDB>
- [5] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, and W. Jaśkowski, “Vizdoom: A doom-based AI research platform for visual reinforcement learning,” in *Proc. IEEE Conf. Comput. Intell. Games*, 2016, pp. 1–8, doi: [10.1109/CIG.2016.7860433](https://doi.org/10.1109/CIG.2016.7860433).
- [6] R. R. Torrado, P. Bontrager, J. Togelius, J. Liu, and D. Perez-Liebana, “Deep reinforcement learning for general video game AI,” in *Proc. IEEE Conf. Comput. Intell. Games*, 2018, pp. 1–8, doi: [10.1109/CIG.2018.8490422](https://doi.org/10.1109/CIG.2018.8490422).
- [7] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song, “Assessing generalization in deep reinforcement learning,” 2018, *arXiv:1810.12282*.
- [8] J. Oh *et al.*, “Discovering Reinforcement Learning Algorithms,” in *Adv. Neural Inf. Process. Syst.*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Red Hook, NY, USA: Curran Associates, Inc., 2020, pp. 1060–1070.
- [9] A. Braylan, M. Hollenbeck, E. Meyerson, and R. Miikkulainen, “Reuse of neural modules for general video game playing,” in *Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 353–359.
- [10] C. C. Tutum, S. Abdulqudus, and R. Miikkulainen, “Generalization of agent behavior through explicit representation of context,” in *Proc. IEEE Conf. Games*, 2021, pp. 1–7, doi: [10.1109/CoG52621.2021.9619141](https://doi.org/10.1109/CoG52621.2021.9619141).
- [11] I. Kostrikov, D. Yarats, and R. Fergus, “Image augmentation is all you need: Regularizing deep reinforcement learning from pixels,” in *Proc. Int. Conf. Learn. Representations*, 2021, pp. 1–21. [Online]. Available: <https://openreview.net/pdf?id=GY6-6sTvGaf>
- [12] K. Cobbe, O. Klimov, C. Hesse, T. Kim, and J. Schulman, “Quantifying generalization in reinforcement learning,” in *Proc. 36th Int. Conf. Mach. Learn.*, vol. 97, Jun. 09–15, 2019, pp. 1282–1289.
- [13] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, “Reinforcement learning with augmented data,” *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 19884–19895, 2020.
- [14] N. Justesen, R. R. Torrado, P. Bontrager, A. Khalifa, J. Togelius, and S. Risi, “Illuminating generalization in deep reinforcement learning through procedural level generation,” *NeurIPS Deep RL Workshop*, 2018.
- [15] K. Cobbe, C. Hesse, J. Hilton, and J. Schulman, “Leveraging procedural generation to benchmark reinforcement learning,” in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 2048–2056.
- [16] D. Perez-Liebana, S. M. Lucas, R. D. Gaina, J. Togelius, A. Khalifa, and J. Liu, *General Video Game Artificial Intelligence*. San Rafael, CA, USA: Morgan & Claypool Publishers, 2020.
- [17] D. Perez-Liebana, J. Liu, A. Khalifa, R. D. Gaina, J. Togelius, and S. M. Lucas, “General video game AI: A multitrack framework for evaluating agents, games, and content generation algorithms,” *IEEE Trans. Games*, vol. 11, no. 3, pp. 195–214, Sep. 2019.
- [18] K. Kuanusont, S. M. Lucas, and D. Pérez-Liebana, “General video game AI: Learning from screen capture,” in *Proc. IEEE Congr. Evol. Comput.*, 2017, pp. 2078–2085.
- [19] D. Apeldoorn and G. Kern-Isberner, “An agent-based learning approach for finding and exploiting heuristics in unknown environments,” in *Proc. 13th Int. Symp. Commonsense Reasoning*, vol. 2052, 2017, pp. 1–8. [Online]. Available: <http://ceur-ws.org/Vol-2052/paper1.pdf>
- [20] A. Dockhorn and D. Apeldoorn, “Forward model approximation for general video game learning,” in *Proc. IEEE Conf. Comput. Intell. Games*, 2018, pp. 1–8, doi: [10.1109/CIG.2018.8490411](https://doi.org/10.1109/CIG.2018.8490411).
- [21] C. Ye, A. Khalifa, P. Bontrager, and J. Togelius, “Rotation, translation, and cropping for zero-shot generalization,” in *Proc. IEEE Conf. Games*, 2020, pp. 57–64.
- [22] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, “Stable-Baselines3: Reliable Reinforcement Learning Implementations,” *J. Mach. Learn. Res.*, vol. 22, no. 268, pp. 1–8, 2021.
- [23] T. Schaul, “A video game description language for model-based or interactive learning,” in *Proc. IEEE Conf. Comput. Intell. Games*, 2013, pp. 1–8, doi: [10.1109/CIG.2013.6633610](https://doi.org/10.1109/CIG.2013.6633610).
- [24] T. Schaul, “An extensible description language for video games,” *IEEE Trans. Comput. Intell. AI Games*, vol. 6, no. 4, pp. 325–331, Dec. 2014.
- [25] J. Togelius, G. N. Yannakakis, K. O. Stanley, and C. Browne, “Search-based procedural content generation: A taxonomy and survey,” *IEEE Trans. Comput. Intell. AI Games*, vol. 3, no. 3, pp. 172–186, Sep. 2011.
- [26] N. Shaker, J. Togelius, and M. J. Nelson, *Procedural Content Generation in Games*. Cham, Switzerland: Springer, 2016.
- [27] A. Summerville *et al.*, “Procedural content generation via machine learning (PCGML),” *IEEE Trans. Games*, vol. 10, no. 3, pp. 257–270, Sep. 2018.
- [28] S. Risi and J. Togelius, “Increasing generality in machine learning through procedural content generation,” *Nature Mach. Intell.*, vol. 2, no. 8, pp. 428–436, 2020.
- [29] J. Liu, S. Snodgrass, A. Khalifa, S. Risi, G. N. Yannakakis, and J. Togelius, “Deep learning for procedural content generation,” *Neural Comput. Appl.*, vol. 33, pp. 19–37, 2021.
- [30] D. Perez *et al.*, “The 2014 general video game playing competition,” *IEEE Trans. Comput. Intell. AI Games*, vol. 8, no. 3, pp. 229–243, Sep. 2016.
- [31] Z. Wang *et al.*, “Sample efficient actor-critic with experience replay,” in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–20. [Online]. Available: <https://static.aminer.cn/upload/pdf/239/1521/964/58d82fc8d649053542fd5854.pdf>
- [32] A. Juliani *et al.*, “Obstacle tower: A generalization challenge in vision, control, and planning,” in *Proc. 28th Int. Joint Conf. Artif. Intell.*, 2019, pp. 2684–2691.
- [33] D. P. Liébana, S. M. Lucas, R. D. Gaina, J. Togelius, A. Khalifa, and J. Liu, “General video game artificial intelligence,” *Synth. Lectures Games Comput. Intell.*, vol. 3, no. 2, pp. 1–191, 2019.
- [34] D. Perez, S. Samothrakakis, S. Lucas, and P. Rohlfshagen, “Rolling horizon evolution versus tree search for navigation in single-player real-time games,” in *Proc. 15th Annu. Conf. Genet. Evol. Comput.*, 2013, pp. 351–358.
- [35] J. Garcia and F. Fernández, “A comprehensive survey on safe reinforcement learning,” *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 1437–1480, 2015.



Chengpeng Hu received the B.E. degree in computer science from the Southern University of Science and Technology, Shenzhen, China, in 2021, where he is currently working toward the Postgraduate degree in computer science with the Department of Computer Science and Engineering.

His research interests include AI in games, evolutionary computation, and its applications in combinatorial optimisation.



Ziqi Wang received the B.E. degree in computer science from the Southern University of Science and Technology, Shenzhen, China, in 2021, where he is currently working toward the Postgraduate degree in computer science with the Department of Computer Science and Engineering.

He has coauthored several papers on procedural content generation and has previously published with the IEEE Conference on Games and the IEEE Congress on Evolutionary Computation. His research interests include procedural content generation through AI methods including deep reinforcement learning, generative adversarial networks, and evolutionary computation.



Tianye Shu received the B.E. degree in computer science from the Southern University of Science and Technology, Shenzhen, China, in 2021, where he is currently working toward the Postgraduate degree in computer science with the Department of Computer Science and Engineering.

His recent paper on Experience-driven PCG via reinforcement learning: A Super Mario Bros study has been published at the 2022 IEEE Conference on Games. His research interests include AI in games and evolutionary multiobjective optimization.



Hao Tong (Graduate Student Member) received the B.E. degree in electronic engineering from the China University of Mining and Technology, Xuzhou, China, in 2017, the M.Sc. degree in computer science from the Joint Program of Southern University of Science and Technology, Harbin Institute of Technology, Harbin, China, in 2019. He is currently working toward the Ph.D. degree with the University of Birmingham, Birmingham, U.K.

He has authored and coauthored various papers published in CEC, GECCO, PPSN, TEVC, etc. His

major research interests include evolutionary computation in vehicle routing problems and dynamic optimization. In particular, he was also very interested in AI in games. He has organized a series of competitions in CoG, GECCO, and PPSN about General Video Game AI.

Mr. Tong was a member of the Games Technical Committee of the IEEE Computational Intelligence Society from 2020 to 2021.




Julian Togelius (Senior Member, IEEE) received the B.A. degree from Lund University, Lund, Sweden, the M.Sc. degree from the University of Sussex, Brighton, U.K., and the Ph.D. degree from the University of Essex, Colchester, U.K., all in computer science.

He is currently an Associate Professor with the Department of Computer Science and Engineering, New York University, New York, NY, USA, and a co-founder of *modl.ai*. He works on artificial intelligence for games and on games for artificial intelligence. He

was with IDSIA, Lugano, Switzerland, and the IT University of Copenhagen, Copenhagen, Denmark. His research interests include procedural content generation in games, general video game playing, player modeling, and fair and relevant benchmarking of AI through game-based competitions. Additionally, he works on topics in evolutionary computation, quality-diversity algorithms, and reinforcement learning.

Dr. Togelius was the Editor-in-Chief of IEEE TRANSACTIONS ON GAMES from 2018 to 2021.



Xin Yao received the B.Sc. degree from the University of Science and Technology of China (USTC), Hefei, China, in 1982, the M.Sc. degree from the North China Institute of Computing Technologies, Beijing, China, in 1985, and the Ph.D. degree from the USTC, in 1990, all in computer science.

He is currently a Chair Professor and founding head of Department of Computer Science and Engineering, Southern University of Science and Technology (SUSTech), Shenzhen, China, and a part-time Professor of computer science with the University of

Birmingham, Birmingham, U.K. His major research interests include evolutionary computation and ensemble machine learning. In particular, he is very interested in iterated prisoner's dilemma games, and co-evolutionary learning of game-playing strategies.

Dr. Yao was a Distinguished Lecturer of the IEEE Computational Intelligence Society (CIS). From 2014 to 2015 and from 2003 to 2008, he was the President of IEEE CIS and Editor-in-Chief of IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION. He was the recipient of the 2010 IEEE Transactions on Evolutionary Computation Outstanding Paper Award for his 2008 paper on Measuring Generalization Performance in Co-evolutionary Learning, 2001 IEEE Donald G. Fink Prize Paper Award, 2016 and 2017 IEEE Transactions on Evolutionary Computation Outstanding Paper Awards, 2011 IEEE Transactions on Neural Networks Outstanding Paper Award, 2010 BT Gordon Radley Award for Best Author of Innovation (Finalist), and several best paper awards at conferences. He was also the recipient of the 2012 Royal Society Wolfson Research Merit Award, 2013 IEEE CIS Evolutionary Computation Pioneer Award, and 2020 IEEE Frank Rosenblatt Award.



Jialin Liu (Senior Member, IEEE) received the B.Sc. degree from the Huazhong University of Science and Technology (HUST), Wuhan, China, in 2010, the M.Sc. degree from the École Polytechnique and Université Paris-Sud, Paris, France, in 2013, the Diplôme d'Ingénieur from the Polytech'Paris-Sud, Paris, France, in 2012, and the Ph.D. degree from Université Paris-Saclay, Gif-sur-Yvette, Paris, France, in 2016.

She is currently a Tenure-Track Assistant Professor with the Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China. Her research interests include AI in games, evolutionary computation, portfolio methods, and their applications to planning, scheduling and learning in uncertain environments.

Dr. Liu is currently an Associate Editor for IEEE TRANSACTIONS ON GAMES and was the Chair of the Games Technical Committee of IEEE COMPUTATIONAL INTELLIGENCE SOCIETY from 2020 to 2021.