

# Decision Making in Monopoly Using a Hybrid Deep Reinforcement Learning Approach

Trevor Bonjour<sup>1</sup>, Marina Haliem<sup>1</sup>, Aala Alsalem<sup>1</sup>, Shilpa Thomas, Hongyu Li,  
Vaneet Aggarwal<sup>1</sup>, *Senior Member, IEEE*, Mayank Kejriwal<sup>2</sup>, *Member, IEEE*,  
and Bharat Bhargava<sup>2</sup>, *Life Fellow, IEEE*

**Abstract**—Learning to adapt and make real-time informed decisions in a dynamic and complex environment is a challenging problem. Monopoly is a popular strategic board game that requires players to make multiple decisions during the game. Decision-making in Monopoly involves many real-world elements such as strategizing, luck, and modeling of opponent's policies. In this paper, we present novel representations for the state and action space for the full version of Monopoly and define an improved reward function. Using these, we show that our deep reinforcement learning agent can learn winning strategies for Monopoly against different fixed-policy agents. In Monopoly, players can take multiple actions even if it is not their turn to roll the dice. Some of these actions occur more frequently than others, resulting in a skewed distribution that adversely affects the performance of the learning agent. To tackle the non-uniform distribution of actions, we propose a hybrid approach that combines deep reinforcement learning (for frequent but complex decisions) with a fixed-policy approach (for infrequent but straightforward decisions). We develop learning agents using proximal policy optimization (PPO) and double deep Q-learning (DDQN) algorithms and compare the standard approach to our proposed hybrid approach. Experimental results show that our hybrid agents outperform standard agents by 20% in the number of games won against fixed-policy agents. The hybrid PPO agent performs the best with a win rate of 91% against fixed-policy agents.

**Index Terms**—Monopoly, deep reinforcement learning, decision making, proximal policy optimization, double deep Q-learning.

## I. INTRODUCTION

**D**ESPITE numerous advances in deep reinforcement learning (DRL), the majority of successes have been in two-player, zero-sum games, where it is guaranteed to converge to an optimal policy [1], such as Chess and Go [2]. Rare (and relatively recent) exceptions include Blade & Soul [3], no-press

Manuscript received 25 January 2022; accepted 14 March 2022. Date of publication 16 May 2022; date of current version 28 November 2022. This work was supported in part by Defense Advanced Research Projects Agency (DARPA) and in part by Air Force Research Laboratory (AFRL) under the Contract No. W911NF2020003. (Trevor Bonjour and Marina Haliem contributed equally to this work.) (Corresponding author: Trevor Bonjour.)

Trevor Bonjour, Marina Haliem, Aala Alsalem, Vaneet Aggarwal, and Bharat Bhargava are with Purdue University, West Lafayette, IN 47907 USA (e-mail: tbonjour@purdue.edu; mwadea@purdue.edu; alsalema@purdue.edu; vaneet@purdue.edu; bbshail@purdue.edu).

Shilpa Thomas, Hongyu Li, and Mayank Kejriwal are with the University of Southern California, Los Angeles, CA 90007 USA (e-mail: shilpath@usc.edu; lihongyu@isi.edu; kejriwal@isi.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TETCI.2022.3166555>, provided by the authors.

Digital Object Identifier 10.1109/TETCI.2022.3166555

diplomacy [4], Poker<sup>1</sup> [6], and StarCraft [7], [8]. In particular, there has been little work on agent development for the full 4-player game of Monopoly, despite it being one of the most popular strategic board games in the last 85 years.

Monopoly is a turn-based real-estate game in which the goal is to remain financially solvent. The objective is to force the opponents into bankruptcy by buying, selling, trading, and improving (building a house or a hotel) pieces of property. A player is only allowed to improve property when they achieve a monopoly. A monopoly is when a player owns all the properties that are part of the same color group. The game resembles the real-life business practice of cornering the market to achieve a real-estate monopoly.

During the game, a player can take multiple actions even when it is not their turn to roll the dice. Imagine you are in the middle of playing Monopoly with friends. It is not your turn to roll the dice, but one of your friends just acquired a property that will give you a monopoly. You know you will need that property if you want to have a chance at winning the game. You initiate a trade request, but you need to make an offer that they will probably accept. You need to think about an amount of money you could offer, or if you have a property that might be of interest to them to offer as an exchange for the property of interest. Maybe you need to mortgage or sell a property to generate cash for the trade - would it even be worth it in the long run. This scenario is a snapshot in time of how many different decisions one needs to make during Monopoly. This complexity makes it a fascinating but challenging problem to tackle.

Previous attempts [9], [10] at Monopoly overlook these complexities and consider a simplified version of the game. In both, the authors model Monopoly as a Markov Decision Process (MDP). [9] gives a novel representation for the state space. [10] find that a higher-dimensional representation of the state improves the learning agent's performance. However, both these attempts consider a very limited set of actions: *buy*, *sell*, *do nothing* in case of [9] and only *buy*, *do nothing* in case of [10]. Unlike previous attempts, we do not simplify the action space in Monopoly. Instead, we consider all possible actions (Table I), including trades, to make the game as realistic as possible. This consideration makes the task more challenging since we now need to deal with a high-dimensional action space.

<sup>1</sup>We note that, even in this case, a two-player version of Texas Hold 'em was initially assumed [5] but later superseded by a multi-player system.

TABLE I  
ACTIONS IN MONOPOLY

Action Type	Associated Properties	Game Phase	Action Parameters	Dimensions
Make Trade Offer (Exchange)	All	Pre-roll, out-of-turn	To player, property offered, property requested, cash offered, cash requested	2268
Make Trade Offer (Sell)	All	Pre-roll, out-of-turn	To player, property offered, cash requested	252
Make Trade Offer (Buy)	All	Pre-roll, out-of-turn	To player, property requested, cash offered	252
Improve Property	Color-group	Pre-roll, out-of-turn	Property, flag for house/hotel	44
Sell House or Hotel	Color-group	Pre-roll, post-roll, out-of-turn	Property, flag for house/hotel	44
Sell Property	All	Pre-roll, post-roll, out-of-turn	Property	28
Mortgage Property	All	Pre-roll, post-roll, out-of-turn	Property	28
Free Mortgage	All	Pre-roll, post-roll, out-of-turn	Property	28
Skip Turn	None	Pre-roll, post-roll, out-of-turn	None	1
Conclude Actions	None	Pre-roll, post-roll, out-of-turn	None	1
Use <i>get out of jail</i> card	None	Pre-roll	None	1
Pay Jail Fine	None	Pre-roll	None	1
Accept Trade Offer	None	Pre-roll, out-of-turn	None	1
Buy Property	All	Post-roll	Property	1

We observe that neither of the earlier state representations contains enough information for the agent to learn winning strategies for Monopoly when considering all the actions. To deal with the high-dimensional action space, we develop an enhanced state space that provides a higher representation power and helps the agent consistently get high win rates against other fixed-policy baseline agents. [10] use a sparse reward function where the agent receives a reward at the end of each game. Our experiments show that a sparse reward function is not ideal and cannot handle the complexities accompanying the full version of Monopoly. [9] use a dense reward function where the agent receives a reward within a game after taking any action. We formulate a dense reward function that performs better than the one given by [9].

In Monopoly, some actions occur more frequently than others resulting in a skewed distribution. For instance, a player is allowed to *trade* with other players at any point in the game, but a player can only *buy* an unowned property when they land on the property square. This rare occurrence of a particular state-action pair increases the computational complexity for a standard DRL agent. There is already some evidence emerging that a pure DRL approach may not always be the only (or even best) solution for solving a complex decision-making task. Recently hybrid DRL approaches have surfaced that result in faster convergence, sometimes to a better policy, in other domains such as operations [11], robotics [12], [13], and autonomous vehicles [14]–[16]. To deal with the non-uniform distribution of actions, we propose a hybrid DRL approach for Monopoly. Specifically, we use a fixed-policy approach for infrequent but straightforward decisions and use DRL for frequent but complex decisions. We show that our hybrid agents have a faster convergence rate and higher win rates against baseline agents when compared to the standard DRL agents.

We summarize the key contributions of the paper as follows:

- We consider all decisions that a player may need to make during Monopoly and develop a novel and comprehensive action space representation (Section IV-B).

- We design an enhanced state space representation (Section IV-A) and an improved reward function (Section IV-C) for Monopoly, using which the learning agents converge sooner and to a better policy in contrast to previous attempts (Section III).
- We develop standard DRL-based agents (Section V-A) that learn winning strategies for Monopoly against different fixed-policy agents. The standard DRL agents win at least 25% more games than the best fixed-policy agent.
- We devise a hybrid approach (Section V-B) to solve the complex decision-making task using DRL for a subset of decisions in conjunction with fixed-policy for infrequent actions. During training, the hybrid agents converge sooner and to a better policy as compared to the standard DRL agents. Our experiments (Section VII) show that the hybrid agent outperforms the standard learning agent by more than 20% in the number of games won against the fixed-policy agents.
- We develop a complete four-player open-sourced simulator for Monopoly (Section VI-A) together with three different fixed-policy baseline agents. The baseline agents (Section VI-B) are implemented based on common winning strategies used by human players in Monopoly tournaments.

## II. BACKGROUND

### A. Monopoly

Monopoly is a board game where players take turns rolling a pair of unbiased dice and make decisions based on their position on the board. Fig. 1 shows the conventional Monopoly game board that consists of 40 square locations. These include 28 property locations, distributed among eight color groups (22 real-estate properties), four railroads, and two utility, that players can buy, sell, and trade. Additionally, there are two tax locations that charge players a tax upon landing on them, six card locations that require players to pick a card from either

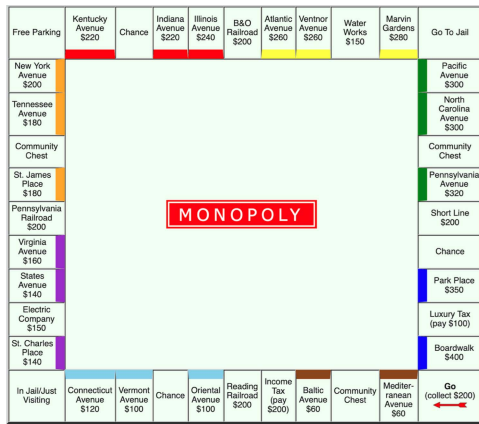


Fig. 1. Monopoly game board.

the *community chest* card deck or the *chance* card deck, the *jail* location, the *go to jail* location, the *go* location, and the *free parking* location. Our game schema also specifies all assets, their corresponding purchase prices, rents, and color. Each square shows the purchase prices that correspond to an asset in Fig. 1. In Monopoly, players act as property owners who seek to buy, sell, improve or trade these properties. The winner is the one who forces every other player into bankruptcy.

### B. Markov Decision Process

An MDP is defined by the tuple  $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$  where  $\mathcal{S}$  is the set of all possible states and  $\mathcal{A}$  is the set of all possible actions. The transition function  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the probability that an action  $a \in \mathcal{A}$  in state  $s \in \mathcal{S}$  will lead to a transition to state  $s' \in \mathcal{S}$ . The reward function  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  defines the immediate reward that an agent would receive after executing action  $a$  resulting in a transition from state  $s$  to  $s'$ .

### C. Reinforcement Learning

Solving an MDP yields a policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$ , which is a mapping from states to actions. An optimal policy  $\pi^*$  maximizes the expected sum of rewards. Reinforcement Learning (RL) is a popular approach to solve an MDP [17] without explicit specification of the transition probabilities. In RL, an agent interacts with the environment in discrete time steps in order to learn the optimal policy through trial and error.

DRL [18] makes use of deep neural networks to approximate the optimal policy or the value function to deal with the limitations of traditional RL methods. The use of deep neural networks as function approximators enables powerful generalization but requires critical decisions about representations. Poor design choices can result in estimates that diverge from the optimal policy [19]–[21]. Existing model-free DRL methods are broadly characterized into *policy gradient* and *value-based* methods.

Policy gradient methods use deep networks to optimize the policy directly. Such methods are useful for physical control where the action space is continuous. Some popular policy gradient methods are Deep deterministic policy gradient (DDPG) [22], asynchronous advantage actor-critic (A3C) [23],

trust region policy optimization (TRPO) [24], proximal policy optimization (PPO) [25].

Value based methods, on the other hand, are based on estimating the value of being in a given state. Deep Q-Network (DQN) [26] is a well-known value-based DRL method. There have been many extensions of the DQN algorithm over the past few years, including double DQN (DDQN) [27], distributed DQN [28], prioritized DQN [29], dueling DQN [30], asynchronous DQN [23], and rainbow DQN [31].

In this paper, we implement both a policy gradient method (PPO) and a value based method (DDQN) to train our standard DRL and hybrid agents (Section V).

## III. RELATED WORK

Despite the popularity of Monopoly, a learning-based approach for decision-making for the full game has not been studied previously. There are older attempts to model Monopoly as a Markov Process such as [32]. [9] and more recently [10] propose modeling Monopoly as an MDP. However, both attempts consider a simplified version of the game. Specifically, both consider a limited set of actions (*buy, sell, do nothing*) with neither work considering trades between players. In [9], an RL agent is trained and tested against a random and a fixed-policy agent. [9] employs a Q-learning strategy along with a neural network. In recent work [10], authors apply a feed-forward neural network with the concept of experience replay to learn to play the game. Their approach supports the idea that no single strategy can maintain high win rates against all other strategies.

Settlers of Catan is a similar board game that involves trades between players. In both, the action distribution is not uniform: certain action types (making trades) are more frequently valid than others. In Monopoly, a player is allowed to trade with other players at any point in the game. However, a player can only buy an unowned property (currently owned by the bank) when they land on the property square. [33] use a model-based approach, Monte Carlo Tree Search (MCTS), for Settlers of Catan. The authors in [33] address the skewed action space by first sampling from a distribution over the types of legal actions followed by sampling individual actions from the chosen action type.

There is evidence emerging in other domains that hybrid DRL techniques reduce the computational complexity of the decision-making task and may provide a better alternative to a pure DRL approach. In [15], authors combine DQN (for high-level lateral decision-making) with the rule-based constraints (to ensure safety) for autonomous driving to achieve a safe and efficient lane change behavior. The authors in [15] use the rules to ensure safety after the DRL has taken a decision. The rule-based method does not select an action instead it just returns whether it is safe to take the action returned by the DQN agent. [12] presents a framework for robots to pick up the objects in clutter by combining DRL and rule-based methods. They divide the task of picking up objects into two parts - pushing and grasping. They make use of the DDPG algorithm with a continuous output for pushing other objects out of the way and use a rule-based method to decide whether to grasp or not. [11] propose an algorithm

for the power-increase operation that uses an A3C agent for the continuous control module and a rule-based system for the discrete control components.

In this work, we use DRL to solve decision-making in Monopoly. Like [9], [10], we represent Monopoly using an MDP, but unlike previous attempts, we do not simplify the game. To make the game as realistic as possible, we consider all possible actions (Table I), including trades. The inclusion of all actions makes the decision-making task more challenging since we need to deal with a high-dimensional action space. We also provide an improved state space representation and reward function when compared to the previous attempts. To handle the non-uniform action space, we propose a hybrid agent that combines a fixed-policy (or rule-based) approach for decisions involving rare actions with DRL for decisions involving remaining actions. Though the idea of hybrid DRL exists in literature, the separation of actions between DRL and fixed-policy is domain-specific. In [11], [12] this distinction is made on the basis that a subset of actions is continuous. In our case, the distinction is not obvious since all our actions are discrete. We are the first to propose the frequency of actions as a criterion for the selection of actions for rule-based and DRL respectively. In this work, we explicitly show that using a hybrid approach is better than using a standard one when all other variables are kept constant. In addition, we also show a comparison of using a policy gradient and value based method for our hybrid agent.

#### IV. MDP MODEL FOR MONOPOLY

We design novel state and action space representations and utilize a combination of dense and sparse reward functions to model the full 4-player game of Monopoly as an MDP.

##### A. State Space

We represent the state as a combination of player and property representation. For the player representation, we consider the current location, amount of cash with the player, a flag denoting if the player is currently in jail and another flag for whether the player has a *get out of jail free* card. Since all other cards force a player to take an action immediately and are not part of the decision-making process, we do not consider them in the state space. For the property representation, we include the 28 property locations. These constitute 22 real-estate properties, four railroad properties, and two utility properties. The property representation consists of owner representation, a flag for a mortgaged property, a flag denoting whether the property is part of a monopoly, and the fraction of the number of houses and hotels built on the property to the total allowed number. We represent the owner as a 4-dimensional one-hot-encoded vector with one index for each player with all zeros indicating the bank. In Monopoly, one can only build a house or a hotel on properties that belong to a color group. Thus for the non-real-estate properties, these values are always zero. We do not include the other locations from the board (Fig. 1) as they do not warrant a decision to be taken by the agent. Overall, the state space representation is a 240-dimensional vector: 16 dimensions for

the player representation and 224 dimensions for the property representation.

##### B. Action Space

We consider all actions that require a decision to be made by the agent. We do not include compulsory actions, like paying tax, moving to a specific location because of a chance card, or paying rent when you land on a property owned by another player. An exhaustive list of actions considered can be found in Table I.

We broadly classify the actions in Monopoly into three groups, those associated with all 28 properties, 22 color-group properties, or no properties. We represent all actions that are not associated with any properties as binary variables. Note that we only have one action that pertains to using a card: Use *get out of jail* card. We do not consider any other card since all the other cards require the player to take the action mentioned on the card immediately. Only a *get out of jail* card can be used at a later time, requiring the player to decide on when to use it. Since improvements (building a house or a hotel) in Monopoly are only allowed for properties belonging to a color group, we represent both *improve property* and *sell house or hotel* as a 44-dimensional vector where 22 dimensions indicate building a house with the other 22 indicating building a hotel on a given property. Actions that are associated with all properties, except for *buy property* and *make trade offer*, are represented using a 28-dimensional one-hot-encoded vector with one index for each property. A player is only allowed to buy an unowned property when they directly land on the property square. Hence, though the action is associated with all properties, the decision to buy or not can be represented using a binary variable.

Trades are possibly the most complex part of the game. A player is allowed to trade with other players anytime during the game. A trade offer has multiple parameters associated with it: it needs to specify the player to whom the trade is being offered. It may further include an offered property, a requested property, the amount of cash offered, and the amount of cash requested. We divide the trade offers into three sub-actions: *sell property trade offers*, *buy property trade offers*, and *exchange property trade offers*. For the buy/sell trade offers, we discretize the cash into three parts: below market price ( $0.75 \times$  purchase price), at market price ( $1 \times$  purchase price) and, above market price ( $1.25 \times$  purchase price). Since we have three other players, 28 properties, and three cash amounts, we represent these using a 252-dimensional ( $3 \times 28 \times 3$ ) vector. To keep the dimensions in check for exchange trade offers, we use the market price for both assets. Thus, we only need to account for the properties and the player. We represent the exchange trade offers using a 2268-dimensional ( $3 \times 28 \times 27$ ) vector. Altogether, the action space has 2922 dimensions.

One thing to note here is that not all actions are valid all the time. Depending on the phase (Section VI-A) of the game, only a subset of possible actions is allowed (Table I).

### C. Reward Function

We use a combination of a dense and a sparse reward function ((1)). In order to reward or penalize a player for the overall policy at the end of each game, we use a constant value of  $\pm c$  for a win/loss respectively.

$$r = \begin{cases} \pm c & \text{for a win/loss} \\ r_x & \text{if the game is not over} \end{cases} \quad (1)$$

where  $r_x$  is the in-game reward for player  $x$  and  $c$  is a constant scalar. In order to choose the value of  $c$ , we ran an experiment with different values and chose the one that gives us the best results in terms of convergence time and the number of wins. More details are presented in Section VII.

During a single game, we use a reward function ((4)) defined as the ratio of the current players' net-worth ((2)) to the sum of the net worth of other active players. We update the net worth of each active player after they take any action.

$$nw_x = c_x + \sum_{a \in A_x} p_a \quad (2)$$

where  $nw_x$  is the net worth of player  $x$ ,  $c_x$  is the current cash with player  $x$ ,  $p_a$  is the price of asset  $a$ , and  $A_x$  is the set of assets owned by player  $x$ . The value for  $p_a$  is calculated as:

$$p_a = (b_p - m_v) * b + n_h * p_h + n_H + p_H \quad (3)$$

where  $b_p$  is the base price (as shown in Fig. 1) of an asset,  $m_v$  is the mortgage value,  $b$  denotes a bonus constant,  $n_h$  and  $p_h$  denote the number of houses and the price of each house on a given property and,  $n_H$  and  $p_H$  denote the number of hotels and price of each hotel built on the property. If an asset is mortgaged,  $m_v$  equals the mortgaged owed on the property. If the asset is not mortgaged, then this value is zero. Since, in Monopoly, it is often better to have property than have cash of the same value, we give a bonus constant to encourage the agent to buy more properties.  $b = 1.5$  if the asset is not part of a Monopoly and  $b = 2$  if the asset is part of a Monopoly. For the in-game reward we have:

$$r_x = \frac{nw_x}{\sum_{y \in X_i \setminus x} nw_y} \quad (4)$$

where  $r_x$  is the in-game reward for player  $x$  and  $X$  is the set of all active players. This reward value is bounded between  $[0,1]$  and helps distinguish the relative value of each state-action pair within a game.

## V. APPROACH

We approach Monopoly from a single agent perspective and treat the other players as part of the environment. We adopt a DRL approach to tackle the decision-making task. As we saw in the previous section (Section IV-B), there are multiple actions an agent can take at any given stage of the game resulting in a complex learning problem. We propose two learning-based approaches, a policy gradient method (PPO) and a value based method (DDQN). In both cases, we train a standard agent that uses a model-free DRL paradigm for all decisions and a hybrid

agent that uses DRL for a subset of actions in conjunction with a fixed-policy for remaining actions.

### A. Standard DRL Agents

1) *Actor-Critic PPO Agent*: We implement the actor-critic PPO algorithm with a clipped surrogate objective function [25]. To estimate the advantage, we make use of the truncated version of generalized advantage estimation used in [23]. The actor-critic PPO implementation makes use of two independent networks - the actor network and the critic network. PPO is an on-policy algorithm in which the agent follows a policy dictated by the actor network for a fixed number of time-steps (much less than the episode length). The actor network parameters are initialized randomly, and the agent *explores* by sampling actions according to the latest version of its stochastic policy. As the training proceeds, the policy typically becomes less random since the update rule encourages it to *exploit* rewards. The critic network is responsible for calculating the value of a given state which is used in the advantage estimation when optimizing the objective function. More details on the approach are presented in Appendix A-A (Supplementary material).

2) *DDQN Agent*: A common issue with using vanilla DQN is that it tends to overestimate the expected return. Double Q-learning [34] overcomes this problem by making use of a double estimator. To avoid over-estimation of the Q-values, we implement the DDQN [27] algorithm to train our agent. Similar to the standard DQN approach, DDQN makes use of an experience replay [35] and a target network. Similar to [26], we make use of the  $\epsilon$ -greedy exploration policy to select actions. Initially, the agent *explores* the environment by randomly sampling from allowed actions. As the learning proceeds and the agent learns which actions are more successful, its exploration rate decreases in favor of more *exploitation* of what it has learned. We mask the output of the network to only the allowed actions to speed up the training process. The action masking ensures that the learning agent selects a valid action at any given time. More details on the approach are presented in Appendix A-B (Supplementary material).

### B. Hybrid Agents

Standard DRL techniques have a high sample complexity. DRL requires each state-action pair to be visited infinitely often, the main reason we use  $\epsilon$ -greedy in the DDQN agent. If some states are rare, we do not want to force the agent to explore them - especially if the related decisions are straightforward and we have an idea of what actions might be good in the given state. When playing Monopoly, a player can only buy an unowned property (property still owned by the bank) when they exactly land on the property square. During our simulations, we observed that the *buy property* action is seldom allowed. Similarly, *accept trade offer* is only valid when there is an outstanding trade offer from another player. The resulting rare-occurring action-state pairs further increase the sample and computational complexity of the learning task. We hypothesize that by using a rule-based approach for the rare occurring but simple decisions

and a learning-based approach for the more frequent but complex decisions, we can improve the overall performance.

We design hybrid agents that integrate the DRL approaches presented earlier (Section V-A) with a fixed-policy approach. We use a fixed-policy to make *buy property* and *accept trade offer* decisions. For all other decisions, we use DRL. During training, if there is an outstanding trade offer, the execution flow shifts from the learning-based agent to a fixed-policy agent to decide whether to accept the trade offer or not. The agent accepts an offer if the trade increases the number of monopolies. If the number of monopolies remains unchanged, the agent only accepts if the net worth of the offer is positive. The net worth of an offer is calculated using:

$$nw_o = (p_o + c_o) - (p_r + c_r) \quad (5)$$

where  $nw_o$  denotes the net worth of the trade offer,  $p_o$  is the price of the property offered,  $c_o$  is the amount of cash offered,  $p_r$  is the price of the property requested, and  $c_r$  is the amount of cash requested.

Similarly, whenever the agent lands on a property owned by the bank, the fixed-policy agent decides whether or not to buy the property. The agent buys the property if it results in a monopoly as long as it can afford it. For all other properties, if the agent has \$200 more than the property price, it decides to buy. Our experiments show that in both cases, the hybrid agent converges faster and significantly outperforms the standard DRL agent when playing against other fixed-policy agents.

## VI. EXPERIMENTAL SETUP

### A. Monopoly Simulator

We develop an open-sourced, complete simulator for a four-player game of Monopoly using Python, available on GitHub.<sup>2</sup> The simulator implements the conventional Monopoly board with 40 locations shown in Fig. 1 and enforces rules similar to the US version of the game,<sup>3</sup> barring some modifications. We do not consider the game rules associated with rolling doubles (for example, a double can get a player out of jail). We treat them as any other dice roll. Trading is an integral part of Monopoly. Players can use trades to exchange properties with or without cash with one or more players. We enforce the following rules for trading:

- Players can trade only unimproved (no houses or hotels) and unmortgaged properties.
- Players can make trade offers simultaneously to multiple players. The player who receives a trade offer is free to either accept or reject it. The trade transaction gets processed only when a player accepts an offer. Once a trade transaction is processed, we terminate all other simultaneous trade offers for the same property.
- A player can have only one outstanding trade offer at a time. A player needs to accept or reject a pending offer before another player can make a different trade offer.

In the conventional setting, players can take certain actions like mortgaging or improving their property even when it is not their turn to roll dice. If multiple players take simultaneous actions, the game could become unstable. To avoid this and to be able to keep track of all the dynamic changes involved in the game, we divide the gameplay into three phases:

- Pre-roll: The player whose turn it is to roll the dice is allowed to take certain actions before the dice roll in this phase. To end the phase, the player needs to *conclude actions*.
- Out-of-turn: Once the pre-roll phase ends for a player, the other players can make some decisions before this player rolls the dice. Every player is allowed to take actions in a round-robin manner in this phase until all players decide to *skip turn* or a predefined number of out of turn rounds are complete.
- Post-roll: Once the player rolls dice, their position is updated based on the sum of the number on the dice. This player then enters the post-roll phase. If the player lands on a property that is owned by the bank, they need to decide whether or not to buy during this phase.

Table I shows the game phases associated with each action. If a player has a negative cash balance at the end of their post-roll phase, they get a chance to amend it. If they are unsuccessful in restoring the cash balance, bankruptcy procedure begins following which the player loses the game.

### B. Baseline Agents

We develop baseline agents that, in addition to buying or selling properties, can make trades. We base the policies of these agents on successful tournament-level strategies adopted by human players. Several informal sources on the Web have documented these strategies though they do not always agree<sup>4</sup> A complete academic study on which strategies yield the highest probabilities of winning has been lacking. Perhaps the complex rules of the game have made it difficult to formalize analytically.

We develop three fixed-policy agents: FP-A, FP-B, and FP-C. All three agents can make one-way (buy/sell) or two-way (exchange) trades with or without cash involvement. They are also capable of rolling out trade offers simultaneously to multiple players. By doing so, the agent increases the probability of a successful trade, so it can acquire properties that lead to monopolies of a specific color group more easily. The fixed-policy agents try to offer properties that hold a low value (for example, a solitary property) to the agent itself but may be of value (gives the other player a monopoly) to the other player and vice versa when making trade requests. To yield a higher cash balance, the agents seek to improve their monopolized properties by building houses and hotels.

All three agents place the highest priority in acquiring a monopoly but differ on the priority they base on each property. FP-A gives equal priority to all the properties, FP-B and FP-C give a high priority to the four railroad properties. Additionally,

<sup>2</sup>[Online]. Available: <https://github.com/mayankkejriwal/GNOME-p3>

<sup>3</sup>[Online]. Available: <https://www.hasbro.com/common/instruct/monies.pdf>

<sup>4</sup>Two resources include <http://www.amnesta.net/monopoly/> and <https://www.vice.com/en/article/mgbzq/10-essential-tips-from-a-monopoly-world-champion>.




Fig. 2. Comparison of win rate (number of wins every 100 games) for different values of  $c$  in the reward function (1) for hybrid PPO agent.




Fig. 3. Comparison of win rate (number of wins every 100 games) for standard PPO and hybrid PPO agent during training. The hybrid agent converges sooner and to a better policy as compared to the standard agent.

FP-B places a high priority on the high rent locations: *Park Place* and *Boardwalk* and assigns a low priority to utility locations. On the other hand, FP-C places high priority on properties in the orange color group (*St. James Place*, *Tennessee Avenue*, *New York Avenue*) or in the sky-blue color group (*Oriental Avenue*, *Vermont Avenue*, *Connecticut Avenue*). An agent tries to buy or trade properties of interest more aggressively, sometimes at the risk of having a low cash balance. It may also end up selling a lower priority property to generate cash for a property of interest.

## VII. EXPERIMENTS AND RESULTS

### A. Training of PPO Agents

We train both the standard PPO agent and the hybrid PPO agent using the actor critic PPO algorithm. We use the same architecture and parameters for both agents in order to draw a fair comparison. We use the state, action, and reward functions presented in Section IV. In the case of the hybrid agent, however, we permanently mask the actions that use a fixed-policy. During training, the agents play against the three fixed-policy agents. We train the learning agents for 2000 games each and randomize the turn order to remove any advantage one may get due to the player’s position. Details on the network architecture and parameters are provided in Appendix B-A (Supplementary material).

1) *Experiment 1 - Constant Value for Reward:* In order to see the effect of the constant value  $c$  on the performance of the agent, we train the hybrid PPO agent with different values of  $c \in \{0, 1, 10, 20, 50, 100\}$ . The win rate (wins per 100 games) is given in Fig. 2. Since PPO is an on-policy algorithm and relies more on the in-game reward, we see that the value of  $c$  (received during win/loss) does not have a major impact on the performance of the agent. We keep  $c = 0$  for the PPO agents for all the following experiments.

2) *Experiment 2 - Standard PPO vs. Hybrid PPO:* We compare the performance of the standard PPO agent to that of the hybrid PPO agent during training. The win rate (wins per 100 games) for each agent during training is shown in Fig. 5. As mentioned earlier, the difference between the two settings is that for the standard PPO agent, all the decisions are taken by the learning agent, whereas for the hybrid PPO agent, a subset of decisions is taken by following a fixed-policy approach. We see




Fig. 4. Comparison of win rate (number of wins every 100 games) for different values of  $c$  in the reward function (1) for hybrid DDQN agent.




Fig. 5. Comparison of win rate (number of wins every 100 games) for standard DDQN and hybrid DDQN agent during training. The hybrid agent converges sooner and to a better policy as compared to the standard DRL agent.

from the graph that the hybrid PPO agent converges sooner and to a better policy when compared to the standard PPO agent.

### B. Training of DDQN Agents

Previous attempts at Monopoly [9], [10] use Q-value based methods to solve the MDP task. In order to compare our approach to the previous approaches<sup>5</sup> and to see how a value based algorithm behaves, we train both the standard agent and the hybrid agent using the DDQN algorithm. For the two agents to be comparable, we use the same architecture and parameters for both. We use the state, action, and reward functions presented in Section IV. Like in the PPO case, we permanently mask

<sup>5</sup>We omit the comparison of action space representations since the agent does not win a single game against the fixed-policy agents when using a simpler action space presented in [9], [10].




Fig. 6. Comparison of win rates of the hybrid DDQN agent during training using our proposed state space representation (Section IV-A) to that previously given by Bailis *et al.* [9] and Arun *et al.* [10].




Fig. 7. Comparison of win rates of the hybrid DDQN agent during training using our proposed reward function (Section IV-C) to that previously given by Bailis *et al.* [9] and Arun *et al.* [10].

the actions that use a fixed-policy for the hybrid DDQN agent. Details on the network architecture and parameters are provided in Appendix B-B (Supplementary material).

1) *Experiment 1 - Constant Value for Reward*: In order to see the effect of the value of the constant  $c$  for the reward function ((1)), we run an experiment with different values of  $c \in \{0, 1, 10, 20, 50, 100\}$  when training the hybrid agent. The hybrid agent is trained for 6000 games against the three fixed-policy agents for each value of  $c$ . We present the results for the win rate (wins per 100 games) in Fig. 4. Since DDQN is an off-policy algorithm and uses a larger experience buffer, we see that, unlike the PPO agent, the value of  $c$  affects the performance of the DDQN agent. We get the fastest convergence and best win rates on average for  $c = 10$ . For all the following experiments we fix the value of  $c$  at 10.

2) *Experiment 2 - Standard DDQN vs. Hybrid DDQN*: In order to compare the standard DRL agent to the hybrid agent, we run an experiment to show how the two agents behave during training and show a comparison of the win rate while training. During training, the agents play against the three fixed-policy agents. We train the learning agents for 10000 games each and use an exponential decay function for the exploration rate. We randomize the turn order during training (and testing) to remove any advantage one may get due to the player’s position. The win rate (wins per 100 games) for each agent during training is shown in Fig. 5.

3) *Experiment 3 - Comparison of State Space Representations*: We perform an experiment to show the comparison of using different state space definitions for the DDQN agent. We compare our state space definition Section IV-A to the state space representation proposed by Bailis *et al.* [9] and Arun *et al.* [10]. Fig. 6 shows a comparison of win rates of the hybrid agent during training using the three different state representations. Please note, we use our action space representation (Section IV-B) and reward function (Section IV-C) for all three training runs and only change the state space representation.

3) *Experiment 4 - Comparison of Reward Functions*: We run an experiment to show how the choice of the reward function affects the performance of the learning agent. We compare our reward function Section IV-C to the reward functions proposed by Bailis *et al.* [9] and Arun *et al.* [10]. Fig. 7 shows a comparison of the win rates of the hybrid agent during training using the three different reward functions. For these training runs, we use

our state and action space representations and only change the reward functions.

### C. Evaluation of Learning Agents

We run multiple experiments to evaluate the different learning agents. In the case of the standard PPO agent, all decisions are made using the pre-trained actor network. The actor network makes a subset of the decisions for the hybrid PPO agent and follows a fixed-policy for the remaining actions. We use the pre-trained policy network to take all the decisions in the case of the standard DDQN agent and a subset of decisions in the case of the hybrid DDQN agent. We set the exploration rate to zero for both the standard and the hybrid DDQN agents. For each experiment, we use different combinations of agents to play the four-player Monopoly game. Each experiment is run for five iterations of 2000 games each. The order of play is randomized for each game.

1) *Experiment 1 - PPO vs. Fixed-Policy*: For the first experiment, we evaluate the performance of the two PPO agents when playing against the three fixed-policy agents. We provide the results for the standard PPO agent and the hybrid PPO agent in Tables II and III respectively. We observe that the standard PPO agent has a win rate of **69.95%** while that of the hybrid PPO agent is **91.65%**. The result is in line with what we saw during training. The hybrid PPO agent outperforms the standard PPO agent by more than **20%** when playing against fixed-policy agents.

2) *Experiment 2 - DDQN vs. Fixed-Policy*: We perform a similar evaluation for the DDQN agents as we do for PPO agents in the first experiment. We run five iterations of 2000 games each against the three fixed-policy agents for both learning agents. The standard DRL agent achieves a win rate of **47.41%** as shown in Table IV. The hybrid agent significantly outperforms the standard agent and achieves a win rate of **76.91%** as shown in Table V.

3) *Experiment 3 - Hybrid vs. Standard vs. Fixed-Policy*: To test how the standard and the hybrid agents perform against each other, we evaluate a setting where the standard and hybrid agents play against each other as well against two fixed-policy agents. We present the results from one of those settings in Table VI for the PPO agents and in Table VII for the DDQN agents. In both cases, we observe that the hybrid agents have a higher win



TABLE II  
RESULTS FOR STANDARD PPO AGENT AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-A	FP-B	FP-C	Std PPO
1	172	177	278	1373
2	157	183	279	1381
3	125	174	275	1426
4	161	158	270	1411
5	154	187	255	1404
Win Rate	7.69%	8.79%	13.57%	<b>69.95%</b>

TABLE III  
RESULTS FOR HYBRID PPO AGENT AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-A	FP-B	FP-C	Hybrid PPO
1	35	53	68	1844
2	42	47	68	1843
3	48	51	69	1832
4	32	68	80	1820
5	50	69	55	1826
Win Rate	2.07%	2.88%	3.40%	<b>91.65%</b>

TABLE IV  
RESULTS FOR STANDARD DDQN AGENT AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-A	FP-B	FP-C	Std DDQN
1	307	307	424	962
2	307	298	455	940
3	286	332	430	952
4	267	333	451	949
5	329	331	402	938
Win Rate	14.96%	16.01%	21.62%	<b>47.41%</b>

TABLE V  
RESULTS FOR HYBRID DDQN AGENT AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-A	FP-B	FP-C	Hybrid DDQN
1	148	169	147	1536
2	143	172	142	1543
3	147	168	134	1551
4	161	176	124	1539
5	147	186	145	1522
Win Rate	7.46%	8.71%	6.92%	<b>76.91%</b>

TABLE VI  
RESULTS FOR HYBRID AND STANDARD PPO AGENTS AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-B	FP-C	Std PPO	Hybrid PPO
1	57	40	108	1795
2	53	51	135	1761
3	54	60	107	1779
4	46	52	113	1789
5	68	50	134	1748
Win Rate	2.78%	2.53%	5.97%	<b>88.72%</b>

TABLE VII  
RESULTS FOR HYBRID AND STANDARD DDQN AGENTS AGAINST FIXED-POLICY AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	FP-B	FP-C	Std DDQN	Hybrid DDQN
1	259	299	272	1170
2	326	293	268	1113
3	300	299	261	1140
4	287	269	293	1151
5	287	308	293	1112
Win Rate	14.59%	14.68%	13.87%	<b>56.86%</b>

TABLE VIII  
RESULTS FOR PPO AGENTS AGAINST DDQN AGENTS OVER FIVE RUNS OF 2000 GAMES EACH

Run	Std DDQN	Hybrid DDQN	Std PPO	Hybrid PPO
1	18	566	27	1389
2	15	594	24	1367
3	23	559	33	1385
4	17	592	32	1359
5	17	545	32	1406
Win Rate	0.9%	28.56%	1.48%	<b>69.06%</b>

rate (**88.72%** and **56.86%**) than the standard agents (**5.97%** and **13.87%**). We present results for additional settings in Appendix C (Supplementary material).

4) *Experiment 4 - PPO vs. DDQN*: For our final evaluation, we evaluate how the four learning agents perform when playing against each other. We see from Table VIII that the standard agents barely win any games. The hybrid PPO agent performs the best with a win rate of **69.06%** and the hybrid DDQN comes in second with a win rate of **28.56%**.

## D. DISCUSSION

From the results, we see that the learning agents outperform the fixed-policy agents by some margin. Although only two action choices separate the standard and the hybrid agents, we observe that the hybrid agents significantly outperform the standard agents. Evidently, instead of letting the agent explore the rare state-action pair it may be better suited if these are replaced by rule-based logic, especially if we know what actions might be good in the given state. Additionally, we see from Fig. 3 and Fig. 5, the hybrid agents converge sooner and to a better policy than the standard DRL agents. From Fig. 6 and Fig. 7 we see that our state representation and design of reward function improves the performance of the DDQN learning agent when compared to previous attempts. Though we see a similar pattern for both value based and policy gradient based methods, we note that the hybrid PPO agent outperforms the hybrid DDQN agent.

## VIII. CONCLUSION

We present the first attempt at modeling the full version of Monopoly as an MDP. Using novel state and action space representations and an improved reward function, we show that our DRL agent learns to win against different fixed-policy agents.

The non-uniform action distribution in Monopoly makes the decision-making task more complex. To deal with the skewed action distribution we propose a hybrid DRL approach. The hybrid agents use DRL for the more frequent but complex decisions combined with a fixed-policy for the infrequent but simple decisions. Experimental results show that the hybrid agents significantly outperform standard DRL agents for both policy based (PPO) and value based methods (DDQN). We evaluate the learning agents in different settings and see that the hybrid agents consistently get a high win rate. The hybrid PPO agent performs the best in all scenarios with a win rate of 91% against the baseline agents and a win rate of 69.06% against other learning agents. In this work, we integrate a fixed-policy approach with a learning-based approach, but other hybrid approaches may be possible. For instance, instead of using a fixed-policy agent, the seldom occurring actions could be driven by a separate learning agent that could either be trained jointly or separately from the principal learning agent. In the future, we plan to explore other hybrid approaches, train multiple agents using Multi-Agent Reinforcement Learning (MARL) techniques, and extend the Monopoly simulator to support human opponents.

#### ACKNOWLEDGMENT

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, or the U.S. Government. The authors thank their team members on this project for all the discussions to develop this paper. Some of the ideas in this paper are based on our learning from the SAIL-ON meetings.

#### REFERENCES

- [1] A. Celli, A. Marchesi, T. Bianchi, and N. Gatti, "Learning to correlate in multi-player general-sum sequential games," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 13 076–13 086.
- [2] D. Silver *et al.*, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [3] I. Oh, S. Rho, S. Moon, S. Son, H. Lee, and J. Chung, "Creating pro-level ai for a real-time fighting game using deep reinforcement learning," *IEEE Trans. Games*, to be published, doi: [10.1109/TG.2021.3049539](https://doi.org/10.1109/TG.2021.3049539).
- [4] P. Paquette *et al.*, "No press diplomacy: Modeling multi agent gameplay," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 4474–4485.
- [5] M. Moravčík *et al.*, "DeepStack: Expert level artificial intelligence in heads up no limit poker," *Science*, vol. 356, no. 6337, pp. 508–513, 2017.
- [6] N. Brown and T. Sandholm, "Superhuman AI for multiplayer poker," *Science*, vol. 365, no. 6456, pp. 885–890, 2019.
- [7] O. Vinyals *et al.*, "Grandmaster level in starcraft II using multi agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [8] K. Shao, Y. Zhu, and D. Zhao, "Starcraft micromanagement with reinforcement learning and curriculum transfer learning," *IEEE Trans. Emerg. Top. Comput. Intell.*, vol. 3, no. 1, pp. 73–84, Feb. 2019.
- [9] P. Bailis, A. Fachantidis, and I. Vlahavas, "Learning to play monopoly: A reinforcement learning approach," in *Proc. 50th Anniversary Conv. Soc. Study Artif. Intell. Simul. Behav.*, 2014.
- [10] E. Arun, H. Rajesh, D. Chakrabarti, H. Cherala, and K. George, "Monopoly using reinforcement learning," in *Proc. IEEE Region 10 Conf.*, 2019, pp. 858–862.
- [11] D. Lee, A. M. Arigi, and J. Kim, "Algorithm for autonomous power-increase operation using deep reinforcement learning and a rule-based system," *IEEE Access*, vol. 8, pp. 196 727–196 746, 2020.
- [12] Y. Chen, Z. Ju, and C. Yang, "Combining reinforcement learning and rule-based method to manipulate objects in clutter," in *Proc. IEEE Int. Joint Conf. Neural Netw.*, 2020, pp. 1–6.
- [13] H. Xiong, T. Ma, L. Zhang, and X. Diao, "Comparison of end-to-end and hybrid deep reinforcement learning strategies for controlling cable-driven parallel robots," *Neurocomputing*, vol. 377, pp. 73–84, 2020.
- [14] A. Likmeta, A. M. Metelli, A. Tirinzoni, R. Giol, M. Restelli, and D. Romano, "Combining reinforcement learning with rule-based controllers for transparent and general decision-making in autonomous driving," *Robot. Auton. Syst.*, vol. 131, 2020, Art. no. 103568.
- [15] J. Wang, Q. Zhang, D. Zhao, and Y. Chen, "Lane change decision-making through deep reinforcement learning with rule-based constraints," in *Proc. IEEE Int. Joint Conf. Neural Netw.*, 2019, pp. 1–6.
- [16] Q. Guo, O. Angah, Z. Liu, and X. J. Ban, "Hybrid deep reinforcement learning based ECO-driving for low-level connected and automated vehicles along signalized corridors," *Transp. Res. Part C, Emerg. Technol.*, vol. 124, 2021, Art. no. 102980.
- [17] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [18] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 26–38, Nov. 2017.
- [19] L. Baird, "Residual algorithms: Reinforcement learning with function approximation," in *Machine Learning Proceedings*, New York, NY, USA: Elsevier, 1995, pp. 30–37.
- [20] S. Whiteson, "Evolutionary function approximation for reinforcement learning," *J. Mach. Learn. Res.*, vol. 7, pp. 877–917, 2006.
- [21] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1587–1596.
- [22] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," 2015, *arXiv:1509.02971*.
- [23] V. Mnih *et al.*, "Asynchronous methods for deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1928–1937.
- [24] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1889–1897.
- [25] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [26] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [27] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI Conf. Artif. Intell.*, 2016, pp. 2094–2100.
- [28] A. Nair *et al.*, "Massively parallel methods for deep reinforcement learning," 2015, *arXiv:1507.04296*.
- [29] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," 2015, *arXiv:1511.05952*.
- [30] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, "Dueling network architectures for deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1995–2003.
- [31] M. Hessel *et al.*, "Rainbow: Combining improvements in deep reinforcement learning," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018, pp. 3215–3222.
- [32] R. B. Ash and R. L. Bishop, "Monopoly as a Markov process," *Math. Mag.*, vol. 45, no. 1, pp. 26–29, 1972.
- [33] M. S. Dobre and A. Lascarides, "Exploiting action categories in learning complex games," in *Proc. IEEE Intell. Syst. Conf.*, 2017, pp. 729–737.
- [34] H. Hasselt, "Double Q-learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 2613–2621.
- [35] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Mach. Learn.*, vol. 8, no. 3/4, pp. 293–321, 1992.