

# Teaching on a Budget in Multi-Agent Deep Reinforcement Learning

Ercüment İlhan, Jeremy Gow and Diego Perez-Liebana  
School of Electronic Engineering and Computer Science  
Queen Mary University of London  
London, United Kingdom  
{e.ilhan, jeremy.gow, diego.perez}@qmul.ac.uk

**Abstract**—Deep Reinforcement Learning (RL) algorithms can solve complex sequential decision tasks successfully. However, they have a major drawback of having poor sample efficiency which can often be tackled by knowledge reuse. In Multi-Agent Reinforcement Learning (MARL) this drawback becomes worse, but at the same time, a new set of opportunities to leverage knowledge are also presented through agent interactions. One promising approach among these is peer-to-peer action advising through a teacher-student framework. Despite being introduced for single-agent RL originally, recent studies show that it can also be applied to multi-agent scenarios with promising empirical results. However, studies in this line of research are currently very limited. In this paper, we propose heuristics-based action advising techniques in cooperative decentralised MARL, using a nonlinear function approximation based task-level policy. By adopting Random Network Distillation technique, we devise a measurement for agents to assess their knowledge in any given state and be able to initiate the teacher-student dynamics with no prior role assumptions. Experimental results in a gridworld environment show that such an approach may indeed be useful and needs to be further investigated.

**Index Terms**—multi-agent, reinforcement learning, deep q-networks, action advising, teacher-student

## I. INTRODUCTION

Reinforcement Learning (RL) [1] is a prominent framework for solving sequential decision-making tasks. In recent years, the success of deep learning has led to the emergence of a new set of techniques called deep RL that achieved important breakthroughs, reaching super-human level of play in many complex domains, including Atari games [2], the game of Go [3] and StarCraft II [4]. However, modern RL methods suffer from sample inefficiency and long training times, resulting in limited scalability and applicability in practical situations. Furthermore, in the cases of having multiple RL agents interacting in a shared environment, which is referred to as Multi-Agent Reinforcement Learning (MARL), these problems are further intensified due to the domain inherent challenges such as non-stationarity and curse of dimensionality introduced by multi-agent dynamics.

One natural solution to accelerate learning in deep RL through overcoming sample inefficiency is *knowledge reuse*. A considerable amount of methods to leverage past knowledge are proposed to this date [5]. On the one hand, they are mostly limited to single-agent RL due to the difficulties of MARL.

On the other hand, agent interactions in MARL present unique opportunities of knowledge reuse.

As shown in [6], various classes of methods are proposed to take advantage of multiple agents for this purpose, such as imitation learning, inverse RL, and learning from demonstrations. *Action advising*, which involves agents exchanging advice in the form of actions between each other to drive their exploration is one of the most flexible approaches among these. The only requirement it has for the agents is to have a common action space and be able to access their local observations through observing or communicating at the time they need an advice. It has a strong track record in single-agent RL and seems to be a promising direction to investigate in MARL.

As Section II shows, the research on applications of action advising techniques in MARL is currently in its early stages and the methods are very limited. In this study, we aim to investigate the performance of heuristic-based teacher-student framework in cooperative MARL with agents employing nonlinear function approximation representations in their task-level policies. The problem setting we address does not hold any assumptions of student-teacher roles for agents, is completely decentralised in training and execution stages, and assumes that the communication between agents is limited; which we believe is realistic. To the best of our knowledge, this is the first study to utilise classical action advising approach in conjunction with deep MARL, and we hope to shed a light on advantages and shortcomings to determine further research directions.

The paper is structured as follows: Section II provides a review of related work. Afterwards, the background information regarding the studied techniques are provided in Section III. Section IV describes the game environment used in this project and Section V details the proposed approach. Section VI presents the experimental setup and Section VII the results obtained. Finally, Section VIII concludes the paper and outlines directions for future work.

## II. LITERATURE REVIEW

### A. Advising Methods in Single-Agent Reinforcement Learning

The idea of peer-to-peer knowledge sharing via advising has its roots in single-agent RL. [7] is one of the first studies to adopt this method, in a form of student-initiated advising, in which the learning agent is assisted by an expert agent whenever it asks for advice to make a decision. In [8], it

was shown that it is possible for one of the peers to be a human and provide feedback to accelerate the learning of the agent. Later on, [9] proposed the *teacher-student framework*. According to this method, an expert teacher constantly monitors a student agent’s learning process and provides it with action advices at appropriate times, limited with a budget, considering the practical concerns regarding attention and communication. Addressing the challenge of *when* to advise became more crucial with the introduction of a budget in this line of research. Following this work, [10] treated the optimal way of spending this budget in terms of student’s learning performance as an RL problem and attempted to learn how to teach. In [11], a new approach called jointly-initiated is presented. As opposed to the previous methods in which advising occurs with student [7] or teacher [9] initiation, they claim that their approach discards the need for student agent to be constantly monitored, making these techniques more feasible for human-agent settings. [12] extended teacher-student framework to take advantage of getting advice from multiple teachers by combining the advices using a voting based selection. Moreover, the case of getting suboptimal advices when student surpasses teacher’s performance is also addressed. Even though this improvement relaxed the requirement of expert teachers, it still assumes them to follow fixed policies that are good enough to provide advice. Finally, [13] made further investigations in learning to teach concept based on [10], and distinguished the qualities of being an expert and being a good teacher, claiming that best performers are not always the best teachers.

All of these previously mentioned studies are based on agents that operate in isolated single-agent environments. In MARL, multiple agents simultaneously learn in the same environment while affecting each others’ policies, rendering it non-stationary. Due to this, even if we have expert agents, they can no longer be guaranteed to follow a fixed policy. Moreover, the assumption of teacher and student roles within agents are not applicable in MARL, especially with large number of agents. Despite being challenging to adapt on, these properties reflect a more natural and realistic way of peer-to-peer knowledge sharing.

### B. Advising Methods in Multi-Agent Reinforcement Learning

The application of action advising methods in MARL is a challenging and a fairly new subject. [14] was the first to propose a teacher-student framework that is suitable with multi-agent settings. They extended the heuristics from [9] by introducing several metrics based on the number of state visits to measure confidence in a given state, in order to overcome the challenge of having no fixed roles of student and teacher. They tested their methods in a cooperative team of agents utilizing SARSA( $\lambda$ ) with linear function approximation as task-level policies. In [15], teaching in MARL is approached as advising-level meta-learning problem. They proposed a centralised training and decentralised execution procedure using multi-agent actor-critic for teaching-level policies and tabular Q-learning for task-level policies. Following this work, [16] extended the idea of teaching to agents that use deep neural networks in their task-level policies and act in environment with

continuous state and action spaces. Despite having promising results, both of these learning based approaches are currently limited with 2 agents only.

## III. BACKGROUND

### A. Multi-Agent RL

We are interested in a cooperative multi-agent setting where agents get local observations and act in a decentralised fashion. *Decentralised partially observable Markov decision process* (Dec-POMDP)[17], which is a generalisation of *Markov decision process* (MDP) for problems with multiple decentralised agents with local observations, is a suitable formalisation for this purpose. A Dec-POMDP is defined by a tuple  $\langle \mathcal{I}, \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \Omega, \mathcal{O}, \gamma \rangle$  where  $\mathcal{I}$  is the set of  $n$  agents,  $\mathcal{S}$  is the set of states,  $\mathcal{A} = \times_{i \in \mathcal{I}} \mathcal{A}^i$  is the set of joint actions,  $\mathcal{T}$  is the state transition probabilities,  $\mathcal{R}$  is the reward function,  $\Omega = \times_{i \in \mathcal{I}} \Omega^i$  is the finite set of joint observations,  $\mathcal{O}$  is the set of conditional observation probabilities, and  $\gamma \in [0, 1)$  is the discount factor. At each timestep  $t$  in a state  $s'$ , each agent  $i$  perceives observation  $o_t^i$  determined by  $\mathcal{O}(o_t^i | s', \mathbf{a})$ , where  $\mathbf{a} = \langle a^1, \dots, a^n \rangle$  is the joint action that caused the state transition from  $s$  to  $s'$  according to  $\mathcal{T}(s' | s, \mathbf{a})$ , and receives reward  $r^i$  determined by  $\mathcal{R}(s, \mathbf{a}_t)$ . In fully-cooperative case, which is the setup we are interested in, every agent gets the same shared reward as  $r^i = \dots = r^n$ . In our problem framework, agents are able to observe each other at any time and infer their local observations. However, they still receive local observations individually from environment.

### B. Deep Q-Networks

Deep Q-Networks (DQN) [18] is a deep RL algorithm which employs deep neural networks to represent environment states and learn approximations of state-action values similar to Q-learning in an end-to-end fashion. It overcomes the limitations present in tabular and linear function approximation methods (like tile coding) which struggle to generalise across states and deal with large state spaces. At each timestep  $t$ , an action  $a_t$  is selected randomly (depending on the exploration technique) or greedily according to  $\max_a Q(o_t, a; \omega)$  based on state observation  $o_t$ , where  $\omega$  are the parameters of the neural network. The transition information  $o_t, a_t, r_{t+1}, o_{t+1}$  is then stored in a buffer  $D$ , called replay memory. Using minibatches of samples drawn from  $D$ , the network parameters  $\omega$  are trained periodically with gradient descent to minimise the loss  $(r_{t+1} + \gamma \max_{a'} Q(o_{t+1}, a') - Q(o_t, a_t))^2$ .

There has been many improvements on DQN over the years [19], and some of them became essential parts of the algorithm as they greatly improve convergence with no significant drawbacks. In this paper, we utilize double Q-learning, prioritised replay, dueling networks and noisy nets enhancements among these.

### C. Random Network Distillation

Random Network Distillation (RND) [20] is a technique built to provide an intrinsic curiosity reward for RL agents to enhance their exploration capabilities in complex environments.

It involves usage of two neural networks alongside the actual task-level RL algorithm, namely *target* and *predictor* networks, denoted by differentiable functions  $G$  and  $\hat{G}$  respectively. These networks are identical in structure, which is defined arbitrarily, and are able to map observations to embeddings as in  $G: \mathcal{O} \rightarrow \mathbb{R}^k$  and  $\hat{G}: \mathcal{O} \rightarrow \mathbb{R}^k$ . They are initialised with different random weights, and they produce different embeddings for identical inputs in their initial state. Over the course of training, samples used in task-level RL algorithm are also used to train predictor network  $\hat{G}$  using gradient descent to minimise the mean squared error  $\|G(x) - \hat{G}(x; \theta)\|^2$ . By doing so, the predictor network becomes more accurate at matching target network’s outputs for the samples it observes more, which is referred to as distilling a randomly initialised network. The error between target and predictor are expected to be higher for type of states that are seen less frequently and it acts as a natural indicator of state novelty. The authors of [20] state that, due to its simplicity, this method does not have to deal with error types generated by environment stochasticity, model misspecification and learning dynamics, unlike its counterparts that try to predict state transitions.

#### D. Teaching on a Budget

The teacher-student framework [9] is a peer-to-peer knowledge transfer procedure originally developed for two agents, namely the teacher and the student. While the student is learning to perform a single-agent task, the teacher provides action suggestions to assist its exploration whenever it is appropriate. These advices are expected to accelerate the student’s learning. Teacher agents also have a form of budget  $b_{give}$  which is the number of advices they can provide to limit the amount they can exchange to reflect realistic scenarios. The challenge of this framework comes from distributing advices in the most efficient way to accelerate the student’s learning progression.

For this purpose, several heuristics are proposed to determine when to advise in teacher-initiated strategy:

- **Early advising:** The intuition is that the earlier states are generally more important in learning, and advice is better taken early on. Following this idea, student is given advice by teacher in the first  $b_{give}$  states.
- **Importance advising:** Even though earlier states are usually more important, not all of them may worth to be given advice in. Therefore, a budget is better spent on the states that are determined to be more important. In this method, importance of a state  $s$  is computed as:

$$I(s) = \max_a Q^{teacher}(s, a) - \min_a Q^{teacher}(s, a) \quad (1)$$

Advice is provided in a state  $s$  if  $b_{give} > 0$  and  $I(s)$  is greater than the importance threshold  $\tau_{imp}$ .

#### IV. THE GAME ENVIRONMENT

Our experiment domain is a gridworld game with discrete space and time, in which the agents must spread and cover the landmark positions cooperatively to maximise their shared rewards. The grid is sized  $10 \times 10$  and consists of 3 agents

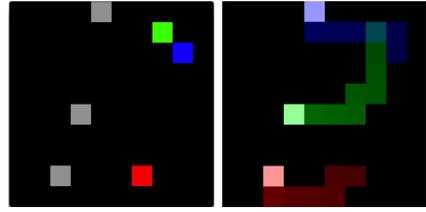


Fig. 1. Two frames of the game used for this study. On the left, the initial game state, with three landmarks (grey) and three agents (green, blue and red). On the right, the terminal state with all agents over the landmarks. Agents’ trajectories are shown as a shaded colour.

and 3 landmarks. The objective of the agents is to move as quickly as possible to cover all the landmarks of the level. While the landmarks are stationary, the agents can navigate around the grid by using a discrete set of actions  $\mathcal{A} = \{ Do\ Nothing, Move\ Up, Move\ Down, Move\ Left, Move\ Right \}$ . Movement actions executed by an agent change its position on the grid by 1 tile in the corresponding direction. At each timestep, the agents observe the relative positions of the other agents and landmarks in form of x-axis and y-axis values, plus the current timestep. All values are normalised to  $[-1, 1]$  using the maximum distance and game duration, respectively. Agents also receive a common reward in  $[0, 1]$ , determined by how many of the landmarks have an agent covering it. The game’s reward is calculated as:

$$\frac{1}{|\mathcal{L}|} \sum_{l=1}^{|\mathcal{L}|} \mathbb{1}_{(\sum_{i \in \mathcal{I}} \mathbb{1}_{\|d_{li}\|=0}) \geq 1} \geq 1 \quad (2)$$

where  $\mathcal{I}$  is the set of agents,  $\mathcal{L}$  is the set of landmarks,  $\mathbb{1}$  is the indicator function, and  $d_{li}$  is the Manhattan distance between landmark  $l$  and agent  $i$ . In other words, in order to act optimally and maximise the total reward, the agents must determine the lowest cost (in terms of distance) of the agent-landmark pair set and move to the appropriate landmark following the shortest path. Since the agents have no access to a forward model, they are expected to learn these strategies through interactions.

Despite this environment’s representational and mechanical simplicity, it still is capable of presenting complex behavioural challenges for MARL. Therefore, we decided to conduct experiments in this setting to focus on behavioural learning while keeping the computational expense at minimum.

We tuned the environment complexity in terms of size and reward sparsity, in order to ensure that the environment is solvable at least sub-optimally by independent DQN agents, and there remains to be some exploration challenge to keep the knowledge of the peers valuable to the agent.

#### V. PROPOSED METHOD

Our approach aims to accelerate learning of a cooperative team of agents in a multi-agent environment with a teacher-student action advising approach. The problem setting we take into consideration is fully observable, completely decentralised in every stage, and most importantly, assumes that

the environment is too complex for tabular or linear function approximation methods to be successful, unlike the majority of related work. In this settings, every agent operates with a local observation, however they are able to observe each other and infer any other agent’s local observations. Additionally, they can exchange action advising requests and responses in forms of actions. We take the budget constraints in inter-agent interactions into consideration as introduced in [9], and we want to achieve acceleration while keeping the agent interactions at minimum, since we believe that such behaviour will be the most useful in practical applications.

The general structure of our proposal can be expressed in two parts: a task-level policy and the teacher-student framework, which we explain in the following sections.

### A. Task-Level Policy

At task-level learning and decision making, agents are designed to employ a nonlinear function approximation according to the target problem structure. Our agents employ DQN along with its well-established improvements of double Q-learning, prioritised replay, dueling networks and noisy nets. In our framework, maybe the most important of these is prioritised replay as it helps with sampling transitions produced by advices which are likely to have larger temporal-difference error more frequently.

The network structure is a multilayer perceptron (MLP), which is the archetypal form of the deep learning models [21], with a single hidden layer formed by 256 units. All layer connections are dense and noisy. This addition removes the need to follow an explicit exploration policy such as  $\epsilon$ -greedy while learning, which we think is more suitable for an action advising approach.

As the multi-agent learning strategy, we follow independent learners approach in which each agent treats other agents as part of a non-stationary environment and behaves by taking only its own actions into account. This is the simplest approach in MARL without requiring the algorithm to have any specialisations for multiple agents. Therefore, we keep DQN as it is without any further modifications in this regard, which is referred to as independent DQNs. Despite having no theoretical convergence guarantees in multi-agent environments, independent DQNs were able to exhibit promising empirical results in previous studies and are often used as a baseline for further improvements [22][23].

### B. Teaching on a Budget

We adopt action advising in the form of the teacher-student framework to perform inter-agent knowledge transfer. This method requires only a common action set and minimal similarity between agents. In addition, considering communication costs, exchanging actions instead of episodes or policy parameters is more preferable especially when task-level policy employs a complex model like DQNs. In our approach, agents can broadcast requests for advices when they need it, and can also respond to requests for advices from other agents. They have separate budgets of asking  $b_{ask}$  and giving  $b_{give}$  advices,

which determine the total number of times they can perform these interactions following either early advising or importance advising heuristics.

Despite its simplicity, adapting the teacher-student framework to MARL with deep RL task-level policies is not straightforward. First of all, in our problem setting, every agent is learning simultaneously and no longer have fixed policies. Moreover, they now have different levels of knowledge about the task, and have no information on each other’s expertise since they can not have any access to internal information and they hold on assumptions about roles of being teacher or student. In order to overcome this, we follow a jointly-initiated advising strategy [11]. Every time an agent requests an advice from another agent, the agent in the position of teacher executes this interaction if it thinks that it can take this role. Furthermore, as agents have no fixed roles or knowledge levels that are set previously, they must be capable of determining if they need an advice based on the state they are, or if they are experienced enough to give advice. In [14], they rely on number of visits to measure an agent’s certainty in a given state. However, this is not applicable when state space is large and nonlinear function approximation is used to represent states. We propose to use RND technique as a metric to measure agents’ uncertainties as an alternative for nonlinear function approximation. Every time an agent uses a sample of state observation to train its internal task-level policy, this sample is also used to *distill* its predictor network  $\hat{G}$ . Consequently, it will be able to measure how uncertain it is in a state with observation  $o$  by measuring the error  $\|\hat{G}(o; \theta) - G(o)\|^2$  when determining if it needs advice or is capable of giving it. This approach can be treated as number of visits in nonlinear function approximation regime.

At each timestep, each agent  $i$  chooses its action according to teacher-student augmented action selection procedure defined in Algorithm 1. If it has a remaining asking budget  $b_{ask}$ , it measures its uncertainty  $\mu$  using its internal models  $G$  and  $\hat{G}$ ; if  $\mu$  is higher than a predefined asking threshold  $\tau_{ask}$  the agent broadcasts its advice requests to other agents as described in Algorithm 2. Agents who have any remaining advice giving budget  $b_{give}$  attempt to respond to this request. If agent receives any advice, it determines which one to follow by using majority voting (ties broken at random), to then execute such action. Otherwise, if no advice is received, it continues exploring by following its own policy.

Responding to an advice request happens as described in Algorithm 3. Upon getting a request, the agent first checks if it has any remaining giving budget  $b_{give}$ . If so, it then determines its expertise by using its internal  $G$  and  $\hat{G}$ , computing  $\|G(o) - \hat{G}(o)\|^2$ . It then compares this value with threshold  $\tau_{give}$  to decide to proceed or not. Finally, if it is set to follow the Early Advising heuristic, it broadcasts an action advice according to its policy; if the heuristic is set to Importance Advising, it also checks if the state is important enough to give advice for by using Equation 1.

The general learning flow of an agent with DQN policy and action advising mechanism is shown in Algorithm 4. As can be seen, the task-level policy can easily be isolated from these

---

**Algorithm 1** Teacher-student augmented action selection

---

**Require:** state observation  $o$ 

```
1: procedure SELECTACTION
2:    $a \leftarrow \text{None}$ 
3:   if  $b_{ask} > 0$  then
4:      $\mu \leftarrow \|G(o) - \hat{G}(o)\|^2$   $\triangleright$  Determine uncertainty
5:     if  $\mu > \tau_{ask}$  then
6:        $a \leftarrow \text{ASKFORADVICE}(o)$ 
7:   if  $a = \text{None}$  then
8:      $a \leftarrow \pi(o)$ 
9:   else
10:     $b_{ask} \leftarrow b_{ask} - 1$ 
11:  return  $a$ 
```

---

**Algorithm 2** Advice seeking

---

**Require:** state observation  $o$ 

```
1: procedure ASKFORADVICE
2:    $a \leftarrow \text{None}$ 
3:    $A \leftarrow \emptyset$ 
4:   for every other agent  $i$  do
5:      $a_{advice} \leftarrow i.\text{ADVISE}$ 
6:     if  $a_{advice} \neq \text{None}$  then
7:       add  $a_{advice}$  to  $A$ 
8:   if  $A \neq \emptyset$  then
9:      $a \leftarrow$  perform majority voting in  $A$ 
10:  return  $a$ 
```

---

**Algorithm 3** Respond to advice request

---

```
1: procedure ADVISE
2:    $a \leftarrow \text{None}$ 
3:   if  $b_{give} > 0$  then
4:      $o_s \leftarrow$  observe student's state.
5:      $\mu \leftarrow \|G(o_s) - \hat{G}(o_s)\|^2$   $\triangleright$  Determine uncertainty
6:     if  $\mu < \tau_{give}$  then
7:       switch teaching method do
8:         case Early Advising
9:            $a \leftarrow \pi(o)$ 
10:           $b_{give} \leftarrow b_{give} - 1$ 
11:        case Importance Advising
12:           $i_o \leftarrow \max_a Q(o_s, a) - \min_a Q(o_s, a)$ 
13:          if  $i_o > \tau_{importance}$  then
14:             $a \leftarrow \pi(o)$ 
15:             $b_{give} \leftarrow b_{give} - 1$ 
16:  return  $a$ 
```

---

enhancements except for its action selection policy. Therefore, it can be changed for any other policy as long as it satisfies the assumption of computing state-action values, for importance advising to be applicable.

## VI. EXPERIMENTAL SETUP

The objective of the evaluation is to understand if and how the proposed modifications can enhance the learning performance. We conducted experiments through multiple

---

**Algorithm 4** Training of task-level policy (DQN) with teacher-student framework

---

```
1: procedure TRAIN
2:   Initialise DQN model
3:   Initialise replay memory
4:   Initialise  $G$  and  $\hat{G}$ 
5:   for all training episodes do
6:      $o \leftarrow$  initial observation
7:     for all episode steps do
8:        $a \leftarrow \text{SELECTACTION}(o)$ 
9:       Execute  $a$  and get reward  $r$  and observation  $o'$ 
10:      Store transition  $(o, a, r, o')$  in replay memory
11:       $o \leftarrow o'$ 
12:      Sample a minibatch from replay memory
13:      Train DQN with minibatch
14:      Train  $\hat{G}$  with minibatch
```

---

learning sessions, each one consisting of a set fixed number of different game episodes which are initialised with random, yet non-overlapping, agent and landmark positions.

The performance of the agents is assessed as a team through a learning session. They are evaluated every 100 episodes in a predefined set of 50 evaluation levels. During evaluation, learning and teaching procedures are disabled, and the levels used are fixed across all learning sessions. An evaluation score is calculated by normalising the average episode rewards obtained across 50 levels with the maximum possible total reward (determined by a set of hand-crafted expert agents), giving a score in  $[0, 1]$ , where 1 indicates the optimal performance.

The performance of the proposed methods can be assessed by looking at the evaluation scores across a learning session, according to the following two metrics:

- **Asymptotic performance:** This is measured directly by looking at the evaluation scores values, and represents how good the agents are at solving the game.
- **Learning speed:** This is measured by looking at area under the curve of evaluation scores against the number of training episodes graph.

For agents to be able to benefit from knowledge transfer, there must be some form of knowledge heterogeneity within the team. In MARL, such heterogeneity tends to arise when the agents explore different parts of the state space, use different task-level policies in terms of complexity and representation, and are in different stages of training. Since our environment is fully observable, and we use identical agents, only the latter is applicable in our setting.

One objective of this study is to determine how the proposed methods work in different types of knowledge heterogeneity. Therefore, we design the following 2 scenarios:

- Scenario I: we train a team of agents in levels from a single distribution of levels; then, we take agents from different stages of pre-training to form a team to be evaluated.
- Scenario II: we train 3 sets of agents in 3 different level distributions, in which landmarks and initial agents

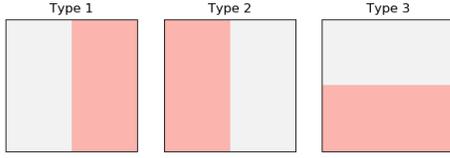


Fig. 2. Level structure types in terms of possible regions for initial positions of agents (gray) and landmarks (red).

locations are strictly limited to predefined regions (see Figure 2). Then, we take one agent with moderate performance arbitrarily from each level type to form a team to evaluate. Note that the learning sessions used to pre-train agents are generated with different seeds than the ones where we run the final evaluation.

The following is a list of the different types of agents used in this study. These agents are trained in the levels first to acquire some knowledge, and then picked to form teams of 3 agents as explained above. These agents are:

- **A0:** Agent with no prior knowledge.
- **A10:** Agent taken from a team of agents trained for 10K episodes. This agent obtained an evaluation score of 0.45 in a level distribution identical to the evaluation sessions.
- **A20:** Agent taken from a team of agents trained for 20K episodes. The evaluation score is 0.91 in a level distribution identical to the learning evaluation sessions.
- **A10-1:** Agent taken from a team of agents trained for 10K episodes, with 0.51 evaluation score in a restricted level distribution of type 1.
- **A10-2:** Agent taken from a team of agents trained for 10K episodes, with 0.41 evaluation score in a restricted level distribution of type 2.
- **A10-3:** Agent taken from a team of agents trained for 10K episodes, with 0.43 evaluation score in a restricted level distribution of type 3.

The methods we evaluate along with the agent teams used in them are as follows:

- **XP:** The team to be evaluated is formed by agents A0, A10, A20 and no advising is used. This method serves as a baseline for the first scenario.
- **XP-EA:** The team to be evaluated is formed by agents A0, A10, A20, and `early advising` is enabled.
- **XP-IA:** The team to be evaluated is formed by agents with different knowledge levels A0, A10, A20, and `importance advising` is enabled.
- **XP-L:** The team to be evaluated is formed by agents A10-1, A10-2, A10-3 and no advising is used. This method serves as a baseline for the second scenario.
- **XP-L-EA:** The team to be evaluated is formed by agents A10-1, A10-2, A10-3, using `early advising`.
- **XP-L-IA:** The team to be evaluated is formed by agents with different knowledge levels A10-1, A10-2, A10-3, using `importance advising`.

The hyperparameters of the learning algorithms are determined empirically, as shown in Table I. They are kept the same

TABLE I  
HYPERPARAMETERS.

Parameter name	Value
Minimum replay memory size to start learning	10000
Replay memory capacity	25000
Prioritisation type	proportional
Prioritisation exponent $\alpha$	0.6
Prioritisation importance sampling $\beta$	0.5 $\rightarrow$ 1
Target network update period	10000
Minibatch size	64
Learning rate	0.001
Train period	2
Discount factor $\gamma$	0.99
Adam $\epsilon$	$1.5 \times 10^{-4}$
Huber loss $\delta$	1
Asking threshold $\tau_{ask}$	10
Giving threshold $\tau_{give}$	3
Importance threshold $\tau_{imp}$	1

across different scenarios and configurations.

We run the tests for the methods with teaching enabled, namely XP-EA and XP-IA, with different advice budgets of 1K, 2K, 5K, 10K, 20K and unlimited ( $\infty$ ). For Scenario I (XP, XP-EA, XP-IA), learning is run for 10K episodes, while 20K episodes are run for Scenario II (XP-L, XP-L-EA, XP-L-IA). The budgets are separate for each individual agent and are set the same for asking and giving advice initially. Moreover, due to stability concerns with deep RL methods, especially in multi-agent settings, every experiment is repeated 10 times with different random seeds.

## VII. RESULTS AND DISCUSSION

Tables II and III show the results for the scenarios described in this paper. These results show the asymptotic performance (score) and the area under the curve (AUC) values as indicators of the learning performance of the agents. Results are reported for different moments of the evaluated training: after 2.5K, 5K, 7.5K, 10K, 15K and 20K episodes, and they include the standard error of the measure (10 repetitions). Results that are significantly different than the first row (baseline) of each table, according to Welch's  $t$ -test ( $p < 0.05$ ) are denoted in bold. Additionally, evaluation scores of the methods with highest final AUC values from each of the scenarios are plotted against their respective baseline as shown in Figures 3 and 4.

In the first scenario, both of the algorithms provided slight accelerations in learning and achieved very similar final performances with the baseline method at 10Kth episode. The only significant difference from the baseline was seen at scores at 5Kth episode, by XP-IA (5K). The overall best performing agent in this scenario is XP-IA (10K). This can be seen as an indication that the importance metric was indeed a useful heuristic to distribute advice over more important states in this scenario. The budget seems to have more effect on XP-IA, achieving its best result at 10K, confirming the claim that too much advice may have negative effects on performance [9].

In the second scenario, XP-L-IA failed to show any significant advantage over the baseline XP-L except for the score at 5Kth episode with budget of 1K. This can be a result of

TABLE II

ASYMPTOTIC PERFORMANCE (SCORE) AND AREA UNDER THE CURVE (AUC) VALUES OF XP, XP-EA AND XP-IA AT LEARNING EPISODES 2.5K, 5K, 7.5K AND 10K WITH STANDARD ERRORS IN PARENTHESES. RESULTS OF XP-EA AND XP-IA THAT ARE SIGNIFICANTLY DIFFERENT THAN XP ALGORITHM ACCORDING TO WELCH'S  $t$ -TEST WITH  $p < 0.05$  ARE DENOTED IN BOLD. EACH ALGORITHM IS RUN 10 TIMES.

Algorithm	At 2500th		At 5000th		At 7500th		At 10000th	
	AUC	Score	AUC	Score	AUC	Score	AUC	Score
XP	13.94 (0.04)	0.58 (0.005)	28.72 (0.12)	0.62 (0.009)	47.15 (0.44)	0.83 (0.013)	68.26 (0.6)	0.87 (0.006)
XP-EA (1K)	13.9 (0.05)	0.58 (0.002)	28.81 (0.14)	0.63 (0.014)	47.18 (0.54)	0.81 (0.013)	68.27 (0.73)	0.86 (0.01)
XP-EA (2K)	13.85 (0.07)	0.57 (0.005)	28.58 (0.13)	0.61 (0.009)	46.39 (0.59)	0.79 (0.024)	67.16 (0.87)	0.86 (0.009)
XP-EA (5K)	13.88 (0.03)	0.58 (0.004)	28.75 (0.11)	0.63 (0.018)	47.02 (0.56)	0.81 (0.024)	67.86 (0.86)	0.87 (0.007)
XP-EA (10K)	13.85 (0.04)	0.57 (0.004)	28.57 (0.09)	0.61 (0.006)	46.3 (0.47)	0.8 (0.021)	67.28 (0.68)	0.87 (0.008)
XP-EA (20K)	13.76 (0.05)	0.57 (0.004)	28.56 (0.09)	0.64 (0.016)	47.38 (0.37)	0.84 (0.005)	68.87 (0.45)	0.88 (0.003)
XP-EA ( $\infty$ )	13.76 (0.05)	0.57 (0.003)	28.53 (0.06)	0.64 (0.011)	47.11 (0.24)	0.82 (0.006)	68.0 (0.31)	0.86 (0.003)
XP-IA (1K)	13.94 (0.04)	0.58 (0.003)	29.0 (0.15)	0.66 (0.016)	47.62 (0.59)	0.82 (0.012)	68.92 (0.71)	0.87 (0.006)
XP-IA (2K)	13.91 (0.08)	0.58 (0.005)	28.86 (0.21)	0.64 (0.02)	47.11 (0.61)	0.81 (0.015)	68.16 (0.84)	0.87 (0.009)
XP-IA (5K)	13.92 (0.03)	0.57 (0.004)	28.84 (0.12)	<b>0.65 (0.008)</b>	48.01 (0.38)	0.83 (0.006)	69.49 (0.41)	0.88 (0.004)
XP-IA (10K)	13.97 (0.04)	0.58 (0.003)	29.22 (0.22)	<b>0.68 (0.024)</b>	48.26 (0.72)	0.8 (0.019)	69.6 (0.98)	0.88 (0.006)
XP-IA (20K)	13.96 (0.04)	0.58 (0.003)	29.19 (0.2)	<b>0.67 (0.022)</b>	47.92 (0.58)	0.8 (0.012)	69.01 (0.75)	0.87 (0.004)
XP-IA ( $\infty$ )	13.96 (0.04)	0.58 (0.003)	29.19 (0.2)	<b>0.67 (0.022)</b>	47.92 (0.58)	0.8 (0.012)	69.01 (0.75)	0.87 (0.004)

TABLE III

ASYMPTOTIC PERFORMANCE (SCORE) AND AREA UNDER THE CURVE (AUC) VALUES OF XP-L, XP-L-EA AND XP-L-IA AT LEARNING EPISODES 5K, 10K, 15K AND 20K WITH STANDARD ERRORS IN PARENTHESES. RESULTS OF XP-L-EA AND XP-L-IA THAT ARE SIGNIFICANTLY DIFFERENT THAN XP-L ALGORITHM ACCORDING TO WELCH'S  $t$ -TEST WITH  $p < 0.05$  ARE DENOTED IN BOLD. EACH ALGORITHM IS RUN 10 TIMES.

Algorithm	At 5000th		At 10000th		At 15000th		At 20000th	
	AUC	Score	AUC	Score	AUC	Score	AUC	Score
XP-L	15.38 (0.12)	0.36 (0.005)	38.34 (0.43)	0.52 (0.006)	66.97 (0.71)	0.66 (0.02)	105.46 (1.37)	0.86 (0.011)
XP-L-EA (1K)	15.58 (0.1)	<b>0.39 (0.008)</b>	<b>39.51 (0.33)</b>	<b>0.54 (0.005)</b>	<b>69.85 (0.85)</b>	0.73 (0.024)	<b>110.8 (1.55)</b>	0.88 (0.018)
XP-L-EA (2K)	15.4 (0.12)	0.37 (0.009)	38.28 (0.66)	0.53 (0.012)	68.55 (1.49)	0.7 (0.022)	108.4 (2.28)	0.86 (0.02)
XP-L-EA (5K)	15.51 (0.14)	<b>0.39 (0.007)</b>	<b>39.98 (0.37)</b>	<b>0.54 (0.006)</b>	<b>71.18 (0.68)</b>	<b>0.75 (0.022)</b>	<b>112.69 (1.27)</b>	<b>0.89 (0.005)</b>
XP-L-EA (10K)	15.53 (0.1)	<b>0.38 (0.008)</b>	<b>39.53 (0.35)</b>	<b>0.55 (0.005)</b>	<b>71.3 (1.02)</b>	<b>0.74 (0.026)</b>	<b>112.14 (1.77)</b>	0.87 (0.014)
XP-L-EA (20K)	15.53 (0.1)	<b>0.38 (0.008)</b>	<b>39.53 (0.35)</b>	<b>0.55 (0.005)</b>	<b>71.3 (1.02)</b>	<b>0.74 (0.026)</b>	<b>112.14 (1.77)</b>	0.87 (0.014)
XP-L-EA ( $\infty$ )	15.53 (0.1)	<b>0.38 (0.008)</b>	<b>39.53 (0.35)</b>	<b>0.55 (0.005)</b>	<b>71.3 (1.02)</b>	<b>0.74 (0.026)</b>	<b>112.14 (1.77)</b>	0.87 (0.014)
XP-L-IA (1K)	15.39 (0.13)	<b>0.38 (0.006)</b>	38.54 (0.5)	0.52 (0.008)	67.95 (0.97)	0.68 (0.025)	107.6 (1.59)	0.87 (0.01)
XP-L-IA (2K)	15.27 (0.14)	0.37 (0.008)	38.4 (0.43)	0.53 (0.007)	67.73 (0.69)	0.67 (0.017)	106.47 (1.42)	0.85 (0.016)
XP-L-IA (5K)	15.37 (0.09)	0.38 (0.009)	38.67 (0.32)	0.52 (0.005)	68.13 (0.59)	0.69 (0.025)	107.32 (1.46)	0.86 (0.019)
XP-L-IA (10K)	15.37 (0.1)	0.37 (0.009)	38.58 (0.34)	0.52 (0.006)	68.06 (0.65)	0.68 (0.027)	107.04 (1.6)	0.85 (0.02)
XP-L-IA (20K)	15.37 (0.09)	0.38 (0.009)	38.67 (0.32)	0.52 (0.005)	68.13 (0.59)	0.69 (0.025)	107.32 (1.46)	0.86 (0.019)
XP-L-IA ( $\infty$ )	15.37 (0.09)	0.38 (0.009)	38.67 (0.32)	0.52 (0.005)	68.13 (0.59)	0.69 (0.025)	107.32 (1.46)	0.86 (0.019)

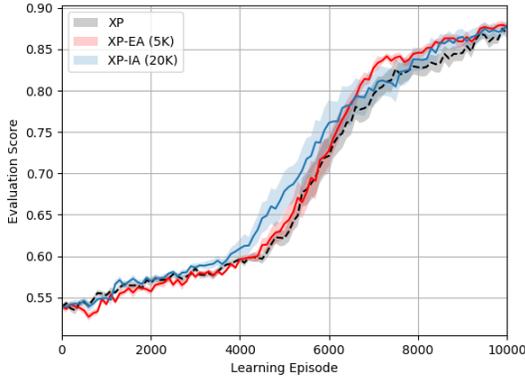


Fig. 3. Evaluation scores versus number of learning episodes of XP, XP-EA (5K) and XP-IA (20K). Shaded areas indicate 95% confidence intervals.

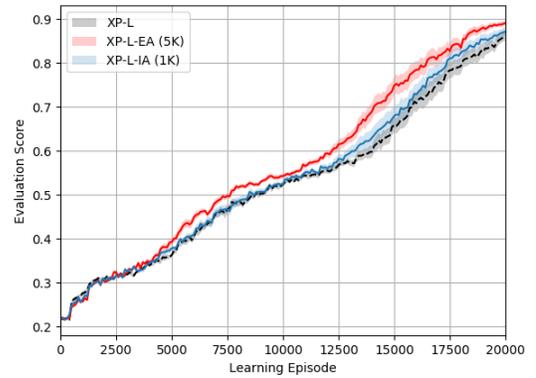


Fig. 4. Evaluation scores versus number of learning episodes of XP-L, XP-L-EA (5K) and XP-L-IA (1K). Shaded areas indicate 95% confidence intervals.

the importance metric not being accurate at reflecting actual relevance of states in this kind of agent knowledge setting. On the other hand, XP-L-EA performed very well with significant improvements in terms of asymptotic performance and learning speed at multiple stages of learning. Moreover, it even managed to achieve a significantly better final performance. This may be caused by the agents starting with similar (and moderate) amount of knowledge, so the early advices are likely to be useful for any of them without having a need for additional importance assessment. The identical results of XP-L with 10K or a higher budget is caused by not having the need to make use of it beyond some point, once the agent is certain of the decision it is making on its own (as controlled with the advice budgets). This can be considered as another benefit of using this uncertainty measurement technique if it is tuned well.

### VIII. CONCLUSIONS

This paper describes the application, for the first time, of action advising via heuristic-based teacher-student framework on Multi-Agent Reinforcement Learning (MARL) agents, employing policies with nonlinear function approximation. The environment used for training agent advising is a grid-based game in which three agents need to coordinate to place themselves in three different landmarks. The work described here shows that using off-policy learning can provide significant improvements in the speed and performance of agents that learn via advice, particularly when the team is composed of agents with heterogeneous knowledge. Moreover, Random Network Distillation (RND) can be a reliable metric to be utilised as a state visit counter through nonlinear function approximation when state space complexity is high.

Another interesting finding is that the state importance metric may be inefficient in some cases of knowledge distribution amongst agents, for example if they are all experts but on different state distributions. Additionally, it is worthwhile highlighting that, even if it is possible to determine the experience of the agents for their roles in knowledge exchange relationships, this experience is importantly biased by the other agents that were present at the time they built their knowledge. Further investigation on how to adapt the importance metric for agent advising is an interesting line of future work.

Off-policy learning through replay memory may be a slowing down factor in action advising, as it takes given advice into consideration in a delayed way and reduces the rate they influence the agent's current policy. Therefore, in addition to the enhancements like prioritised experience replay, it may be useful to implement more specific techniques like multi-step advice and continual monitoring of agents for fixed periods of time after advice exchange. This is similar to previous work in the field [11][16].

Finally, another interesting line of future work is to expand the problem to more than 3 agents, which brings interesting aspects to the discussion such as defining a more accurate peer selection and decision making beyond majority voting.

### REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.
- [2] V. Mnih, K. Kavukcuoglu, D. Silver, *et al.*, "Human-level Control Through Deep Reinforcement Learning," *Nature*, vol. 518:7540, pp. 529–533, 2015.
- [3] D. Silver, A. Huang, C. J. Maddison, *et al.*, "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [4] O. Vinyals, I. Babuschkin, J. Chung, *et al.*, *AlphaStar: Mastering the Real-Time Strategy Game StarCraft II*, <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>, 2019.
- [5] M. E. Taylor and P. Stone, "Transfer Learning for Reinforcement Learning Domains: A Survey," *Journal of Machine Learning Research*, vol. 10, pp. 1633–1685, 2009.
- [6] F. L. Da Silva and A. H. R. Costa, "A Survey on Transfer Learning for Multiagent Reinforcement Learning Systems," *Journal of Artificial Intelligence Research*, vol. 64, pp. 645–703, 2019.
- [7] J. A. Clouse, *On Integrating Apprentice Learning and Reinforcement Learning*. University of Massachusetts Amherst, 1996.
- [8] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy Shaping: Integrating Human Feedback with Reinforcement Learning," in *Neural Information Processing Systems*, 2013, pp. 2625–2633.
- [9] L. Torrey and M. E. Taylor, "Teaching on a Budget: Agents Advising Agents in Reinforcement Learning," in *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13, Saint Paul, MN, USA, May 6-10, 2013*, 2013, pp. 1053–1060.
- [10] M. Zimmer, P. Viappiani, and P. Weng, "Teacher-Student Framework: A Reinforcement Learning Approach," in *AAMAS Workshop Autonomous Robots and Multirobot Systems*, 2014.
- [11] O. Amir, E. Kamar, A. Kolobov, and B. J. Grosz, "Interactive Teaching Strategies for Agent Training," in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 2016, pp. 804–811.
- [12] Y. Zhan, H. Bou-Ammar, and M. E. Taylor, "Theoretically-Grounded Policy Advice from Multiple Teachers in Reinforcement Learning Settings with Applications to Negative Transfer," *CoRR*, vol. abs/1604.03986, 2016.
- [13] A. Fachantidis, M. E. Taylor, and I. P. Vlahavas, "Learning to Teach Reinforcement Learning Agents," *CoRR*, vol. abs/1707.09079, 2017.
- [14] F. L. da Silva, R. Glatt, and A. H. R. Costa, "Simultaneously Learning and Advising in Multiagent Reinforcement Learning," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '17, Sao Paulo, Brazil, 2017, pp. 1100–1108.
- [15] S. Omidshafiei, D. Kim, M. Liu, G. Tesauro, M. Riemer, C. Amato, M. Campbell, and J. P. How, "Learning to Teach in Cooperative Multiagent Reinforcement Learning," *CoRR*, vol. abs/1805.07830, 2018.
- [16] D. K. Kim, M. Liu, S. Omidshafiei, S. Lopez-Cot, M. Riemer, G. Habibi, G. Tesauro, S. Mourad, M. Campbell, and J. P. How, "Learning Hierarchical Teaching in Cooperative Multiagent Reinforcement Learning," *arXiv preprint arXiv:1903.03216*, 2019.
- [17] F. A. Oliehoek and C. Amato, *A Concise Introduction to Decentralized POMDPs*, ser. Springer Briefs in Intelligent Systems. Springer, 2016.
- [18] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, "Playing Atari with Deep Reinforcement Learning," *CoRR*, vol. abs/1312.5602, 2013.
- [19] M. Hessel, J. Modayil, H. van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. G. Azar, and D. Silver, "Rainbow: Combining Improvements in Deep Reinforcement Learning," in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, 2018, pp. 3215–3222.
- [20] Y. Burda, H. Edwards, A. J. Storkey, and O. Klimov, "Exploration by Random Network Distillation," *CoRR*, vol. abs/1810.12894, 2018.
- [21] I. J. Goodfellow, Y. Bengio, and A. C. Courville, *Deep Learning*, ser. Adaptive Computation and Machine Learning. MIT Press, 2016.
- [22] J. Foerster, I. A. Assael, N. de Freitas, and S. Whiteson, "Learning to Communicate with Deep Multi-Agent Reinforcement Learning," in *Neural Information Processing Systems*, 2016, pp. 2137–2145.
- [23] A. Tampuu, T. Maitinen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru, and R. Vicente, "Multiagent Cooperation and Competition with Deep Reinforcement Learning," *CoRR*, vol. abs/1511.08779, 2015.