

# Expert-Free Online Transfer Learning in Multi-Agent Reinforcement Learning

Alberto Castagna<sup>a,\*</sup> and Ivana Dusparic<sup>a</sup>

<sup>a</sup>School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland

**Abstract.** Transfer learning in Reinforcement Learning (RL) has been widely studied to overcome training challenges in Deep-RL, i.e., exploration cost, data availability and convergence time, by bootstrapping external knowledge to enhance learning phase. While this overcomes the training issues on a novice agent, a good understanding of the task by the expert agent is required for such a transfer to be effective.

As an alternative, in this paper we propose *Expert-Free Online Transfer Learning (EF-OnTL)*, an algorithm that enables expert-free real-time dynamic transfer learning in multi-agent system. No dedicated expert agent exists, and transfer source agent and knowledge to be transferred are dynamically selected at each transfer step based on agents' performance and level of uncertainty. To improve uncertainty estimation, we also propose *State Action Reward Next-State Random Network Distillation (sars-RND)*, an extension of RND that estimates uncertainty from RL agent-environment interaction.

We demonstrate *EF-OnTL* effectiveness against a no-transfer scenario and state-of-the-art advice-based baselines, with and without expert agents, in three benchmark tasks: Cart-Pole, a grid-based Multi-Team Predator-Prey (MT-PP) and Half Field Offense (HFO).

Our results show that *EF-OnTL* achieves overall comparable performance to that of advice-based approaches, while not requiring expert agents, external input, nor threshold tuning. *EF-OnTL* outperforms no-transfer with an improvement related to the complexity of the task addressed.

## 1 Introduction

Transfer learning (TL) in reinforcement learning (RL) has been introduced to address two main shortcomings of RL: data availability, as acquiring sufficient interaction samples can be prohibitive due to large state and action space, and lowering the exploration cost due to the partial observability, sparse feedback and safety concerns that may incur in real world environments [6].

Most TL in RL solutions utilize an expert, either agent or human, which supervises a novice agent through the Teacher-Student framework [22]. This can be exploited in two ways, by supporting other agents during exploration, i.e., [13, 14, 17, 28], or by training a novice agent to emulate an expert behaviour, i.e., [10, 20]. While this solves the issue for "student" agents, by lowering the amount of required training data and time for a specific task, it does not solve the broader problem, as an expert agent is required for each task. Furthermore, there may be task where an expert is unavailable to support student agents or does not yet exist.

This paper addresses the above RL limitations by introducing *Expert-Free Online Transfer Learning (EF-OnTL)*. *EF-OnTL* removes the need for a fixed expert in the system. Instead, it takes advantage of different knowledge gained in different parts of a multi-agent system. At each step, a current source agent is dynamically selected based on system performance and/or epistemic confidence. The source agent's experience is partially transferred to target agents. Roles are assigned at each transfer step, while knowledge shared are agent-environment transitions  $(s^t, a^t, r^t, s^{t+1})$  labelled with agent's epistemic uncertainty  $(u^t)$ .

To estimate uncertainty, we propose an extension of Random Network Distillation RND [1]: *State Action Reward Next-State Random Network Distillation (sars-RND)*. sars-RND estimates agent's epistemic uncertainty from the full agent-environment interaction. Thus, while RND approximates a state-visit counter, sars-RND provides a more fine-grained estimation based on the action taken and on its outcome.

Source selection relies on two criteria, *Average Uncertainty* ( $\bar{U}$ ) and *Best Performance* (BP).  $\bar{U}$  evaluates uncertainties over collected samples while BP analyses agent's performance in a given interval. Once transfer roles are defined for the time step, target agents receive a personalised batch of experience, selected based on sender-receiver uncertainties and expected surprise. Expected surprise is estimated by computing the expected loss on target's side [25]. We compare our proposed method against three baselines: (i) a *no-transfer* scenario; (ii) baseline in which, instead of transferring experience, an advice is transferred, i.e. a recommended action (in order to evaluate the impact of sharing advice versus experience with *Online Confidence-Moderated Advice Sharing (OCMAS)*), as work in [11] and (iii) state-of-the-art advice-based approach *Requesting Confidence-Moderated Policy advice* [5] (*RCMP*), in which previously trained agents provide action-advice on demand, when an agent is in a state with high uncertainty.

We evaluate *EF-OnTL* in three environments, (1) Cart-Pole, (2) Multi-Team Predator-Prey (MT-PP) and (3) Half Field Offense (HFO) [8].

Therefore, contribution of this paper is threefold:

- We propose *EF-OnTL*, an algorithm that enables online transfer learning by experience-sharing without a need for an expert agent and is applicable to a range of RL-methods;
- We propose sars-RND, an extension of RND, that estimates uncertainty based on a full RL environment transition tuple  $(s^t, a^t, r^t, s^{t+1})$ ;
- We assess several criteria to enable dynamic selection for source of transfer and shared knowledge. As result, at each step, each

\* Corresponding Author. Email: acastagn@tcd.ie.

target agent receives a tailored batch of experience, specific to its current belief.

The rest of this paper is organized as follows. Related work on agent-to-agent TL and overview of uncertainty estimators is presented in Section 2. *EF-OnTL*, sars-RND and proposed criteria are detailed in Section 3. Simulation setup is in Section 4 while evaluation results are presented in Section 5. Lastly, Section 6 discusses the framework’s limitations and future work directions.

## 2 Related Work

This section provides an overview of the existing work on agent-to-agent TL and methods used to approximate agent’s epistemic uncertainty.

### 2.1 Transfer Learning in RL

Most of agent-to-agent TL work is based on the teacher-student framework [22] where the teacher is an expert and the student is a novice agent. The student can request the teacher’s supervision, and due to resource constraints, their interaction is limited by a budget.

Preferred form of advice is action-based, i.e., [4, 5, 11, 13, 15, 17, 19, 21, 22, 28]. This lowers the overall exploration cost for the novice by asking for an action to follow in certain states. Other form of advice is Q-values [12, 13, 29] used to influence the action-selection process of a target agent. Finally, advice can be provided as policy to be followed for a certain number of steps [26] or a batch of RL experience, similar to a demonstration, as in [2, 3, 7, 14, 23].

Despite the good result shown by the teacher-student framework, roles are kept fixed and transfer impact is limited by the teacher knowledge. Therefore, insufficiently skilled expert might worsen novice’s performance. To overcome this limitation, [15, 26] introduce a real-time training of a centralised super-entity based exclusively on a subset of collected demonstration. This new entity provides on-demand advice when needed. While that improves advice quality over time, introducing a new super agent results in an additional cost to gather the experiences and to train the underneath model. [4, 11, 12, 19, 29] improved the base teacher-student framework by relying on confidence-based and importance-based methods to dynamically select one or more agents as teachers. In these models, an agent can ask for and provide advice simultaneously.

Baselines used in this paper are based on the above state-of-the-art methods: *OCMAS* is based on [11] and *RCMP* on [5]. In [11], an uncertain agent asks for advice and more confident agents recommend an action to follow. Confidence is measured through sars-RND and advice is given only when the advice seeker is the most uncertain agent across the team. Final action is taken by majority voting. In *RCMP* [5], an agent asks for advice during exploration based on its epistemic uncertainty. Advice, in form of an action, is given by a fixed expert-demonstrator. In our implementation, we replaced the single advisor with a jury of multiple trained agents to ward off any bias that might arise from using a single trained agent as teacher and action is selected by majority voting.

Table 1 summarises relevant recent work to this research. For each method, we indicate whether a prior external expert is needed or not. Note that we do not distinguish teacher expertise, i.e., optimal, sub-optimal, etc. When expert is not required, we specify whether the algorithm enables dynamic roles or not. Lastly, we report type of advice transferred and the frequency of transfer. The latter does not take into account when nor how often the advice is used.

Table 1: Summary of related work.

Reference	Expert required	Dynamic roles	Transferred advice	Transfer Frequency
[5, 17, 21] [22, 28]	✓	✗	action	dynamic
[4, 11, 12] [19, 29]	✗	✓	action Q-values	dynamic dynamic
[15] [26]	✗	✗	action policy	dynamic dynamic
[2, 3, 7] [14, 23]	✓	✗	RL-tuple	fixed
[13]	✓	✗	action & Q-values	fixed
<b><i>EF-OnTL</i></b>	✗	✓	<b>RL-tuple</b>	<b>fixed</b>

Our work differs from the above as it removes the need for presence of an expert agent by dynamically selecting a learning agent as a teacher, to exploit different knowledge gained in different parts of a multi-agent system. Furthermore, while most of the related work continuously override target policy by advising actions, Q-values and policy, *EF-OnTL* combines local and received knowledge by transferring, less frequently, an experience batch specifically selected based on a gap between source and target beliefs. Lastly, *EF-OnTL* does not require the tuning of additional parameters, i.e., uncertainty threshold, to improve the quality of transfer.

### 2.2 Uncertainty Estimation in TL

In RL there are two types of uncertainties, *aleatoric* and *epistemic*. *Aleatoric* comes from the environment and is generated by stochasticity in observation, reward and actions. *Epistemic* uncertainty comes from the learning model and indicates whether the agent has adequately explored a certain state. Most recent TL frameworks rely on epistemic uncertainty to determine if an agent requires guidance.

A possible approach to approximating agent’s epistemic uncertainty in a particular task is to determine the frequency of visits to each state. In fact, [4, 15, 19] estimate uncertainty by relying on function defined over state visits counter. Similarly, [29] relies on number of visits over a state-action pair. However, state space might be continuous or very large making the counting unfeasible. Thus, state visit counts could be approximated by RND [1]. RND has been proposed to encourage exploration within agents but has already been exploited as uncertainty estimator for TL in RL in [2, 11]. RND consists of two networks, a target with an unoptimised and randomly initialised parameters and a main predictor network. Throughout time, the former is distilled within the latter and uncertainty is defined as prediction error between the two outputs.

Other sophisticated models can be used to estimate the epistemic uncertainty. [23] proposes uncertainty estimation through neural network, decision tree and Gaussian process. Although, neural networks seem to be preferred overall as are also used in [5], where the agent’s learning model is expanded by replacing control-layer with an ensemble to estimate agent’s uncertainty. However, despite the different underlying technique used to estimate uncertainty, all the presented methods rely uniquely on visited state.

To overcome this limitation, this paper introduces sars-RND as an extension of RND. sars-RND computes epistemic uncertainty over a full RL interaction  $(s^t, a^t, r^t, s^{t+1})$ .

### 3 Expert-Free Online Transfer Learning

This section introduces the three main contributions of this research: (1) sars-RND: Section 3.1 addresses the uncertainty estimator model used within this work; (2) *EF-OnTL*: Section 3.2 provides the details of the core of our research work towards an autonomous framework for online transfer learning; (3) Selection criteria: Section 3.3 presents two criteria used to dynamically select transfer source and Section 3.4 introduces criteria used to identify RL experience worth to be shared, by evaluating source and transfer beliefs. These metrics enable the batch to be personalised based on specific target knowledge shortcomings.

#### 3.1 sars-RND

We propose *State Action Reward Next-State Random Network Distillation* (sars-RND) as an extension of RND to improve the estimation of epistemic uncertainty. When using RND as an uncertainty estimator, agent might lose important information that should be taken into account when computing uncertainty. For instance, in a sparse-reward environment and/or continuous control space, i.e., HFO and our Multi-Team Predator-Prey implementation described at Section 4, by not considering other details, such as action, agent might become confident within a state where the goal fulfilment is close but not yet achieved.

To overcome this potential limitation we propose sars-RND, an uncertainty estimator that takes into account a full RL interaction  $(s^t, a^t, r^t, s^{t+1})$ , rather than just state, to compute epistemic uncertainty of an agent at a specific time.

#### 3.2 Transfer Framework

This section outlines the core contribution of this work *Expert-Free Online Transfer Learning (EF-OnTL)*. *EF-OnTL* is a novel online transfer learning algorithm that overcomes the need of an unique expert agent by dynamically selecting a temporary expert within each transfer iteration. Selected agent is used as source of transfer and some of its collected experience is made available to others. Target agent can then filter and sample a batch of experience to be integrated into its learning process and finally update its belief. Transferred batch contains five-elements tuples,  $(s^t, a^t, r^t, s^{t+1}, u^t)$ , source agent's RL interaction at time  $t$  and  $u^t$ , an additional value that identifies source's uncertainty over the RL interaction updated at time  $t$  of visit.

*EF-OnTL* is independent of an underlying RL algorithm used. Thus, it can be exploited on a range of RL methods, both tabular and neural network-based ones.

To define *EF-OnTL* framework, we specify the following:

- $N$  – number of RL-based agents available during a simulation;
- a set of *Agents*  $A = \{A_1, \dots, A_N\}$ ;
- a set of *Learning Processes*  $LP = \{LP_1, \dots, LP_N\}$  with the following relation  $f_1 = (A_i, LP_j) \in A \times LP \leftrightarrow i = j$ ;
- a set of *Uncertainty Estimators*  $UE = \{UE_1, \dots, UE_N\}$  and relation  $f_2 = (A_i, UE_j) \in A \times UE \leftrightarrow i = j$ ;
- a set of *Transfer Buffers*  $TB = \{TB_1, \dots, TB_N\}$  related to *agents* by  $f_3 = (A_i, TB_j) \in A \times TB \leftrightarrow i = j$ ;
- $B$  – constrained *Transfer Budget* to be used within a single transfer interaction;
- $TF$  – *Transfer Frequency* defined as number of episodes that occur between two consecutive transfer steps;
- $SS$  – *Source Selection* technique used to select source of transfer;

- $TM$  – *Transfer Methods* used to filter relevant knowledge to be transferred.

---

#### Algorithm 1 Expert-Free Online Transfer Learning

---

```

1: Given:  $A, LP, UE, TB, B, TF, SS, TM, f_1, f_2, f_3$ 
2: for  $ep$  in Episodes do
3:   for all  $A_i \in A$  do ▷ follow normal policy
4:     get state  $s_i^t$  for  $A_i$  at time  $t$ 
5:     sample an action  $a_i^t$  based on  $LP_i$ 
6:     perform a step and observe  $o_i^t = (s_i^t, a_i^t, r_i^t, s_i^{t+1})$ 
7:     estimate uncertainty  $u_i^t = UE_i(o_i^t)$ 
8:     push  $(o_i^t, u_i^t)$  to  $TB_i$  ▷ FIFO queue
9:     optimize  $UE_i$  on  $o_i^t$ 
10:    if time to update  $LP_i$  then
11:      optimize  $LP_i$ 
12:    end if
13:  end for
14:  if  $ep \% TF$  is 0 then ▷ start transfer-step
15:    select source agent  $A_s$  by  $SS$ 
16:    for all  $A_t \in (A \setminus A_s)$  do ▷ transfer from  $A_s$  to  $A_t$ 
17:      apply  $TM$  over  $TB_s$  and sample  $B$  tuples
18:      update  $LP_t$  with the sampled tuples
19:    end for
20:  end if
21: end for

```

---

Algorithm 1 introduces high level procedure followed by agents for sharing experience one to another throughout their simultaneous exploration processes.

First, line 1 defines *EF-OnTL* parameters. Then, at lines 3 – 6 agents retrieve observation from an environment, sample an action based on their learning process and finally take a step. Afterward, each agent  $A_i$  estimates its epistemic uncertainty  $u_i^t$  over the interaction  $(s_i^t, a_i^t, r_i^t, s_i^{t+1})$ , updates its  $UE_i$  model based on new sampled evidence and publishes its uncertainty-labelled collected tuple to the associated transfer buffer  $TB_i$  (7 – 9). When  $TB_i$  is at full capacity, new labelled interactions replace the oldest tuples following a FIFO scheme. Learning process  $LP_i$  is then updated based on the underlying model used, lines 10 – 12. Lines 14 – 20 show a transfer-step. At line 15, source is selected among the agents w.r.t. some fixed criteria  $SS$ . Secondly, remaining agents, i.e., targets, apply a filtering  $TM$  over source transfer buffer and then sample a batch composed by a fixed number  $B$  of tuples (17). Finally, at line 18, each target agent updates its learning process based on the obtained batch.

While the above provides a full description of the algorithm, Fig. 1 sketches *EF-OnTL* workflow for the simplest scenario with only 2 agents.

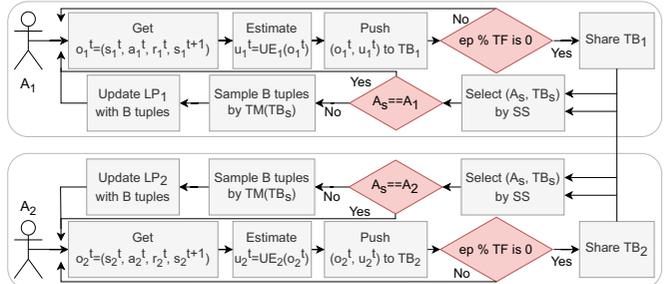


Figure 1: *EF-OnTL* workflow with a 2 agents scenario.

### 3.3 Source Selection Criteria

One of the most crucial parts of *EF-OnTL* is accurately selecting the source from which knowledge is extracted. In this section, we propose *Source Selection Metrics (SS)* to select a source agent among several candidates. Roles are dynamically assigned at the beginning of each transfer iteration and *SS* goal is to identify an experienced candidate whose knowledge will benefit others.

We propose two different methods to identify transfer source in *EF-OnTL*: *Average Uncertainty ( $\bar{U}$ )* and *Best Performance (BP)*. *Average Uncertainty ( $\bar{U}$ )* relies on average uncertainty  $\bar{u}_i$  over transfer buffer  $TB_i$ . Therefore, source agent  $A_s \leftarrow \min_{i=0}^N(\bar{u}_i)$  where,  $\bar{u}_i = \frac{\sum_{t=0}^{|TB_i|} u_t^i \in TB_i}{|TB_i|}$ . On the other hand, BP relies on performance achieved by agents during latest episodes. As performance measure we use average cumulated reward over episodes' finite-horizon undiscounted return ( $\bar{R}_i$ ).  $A_s \leftarrow \max_{i=0}^N(\bar{R}_i)$  where,  $\bar{R}_i = \frac{\sum_{e=0}^E \sum_{t=0}^{\tau_i^e} r_i^{e,t}}{E}$ .  $E$  is the number of evaluated episodes and  $\tau_i^e$  represents the length of the  $e$ -th episode for the  $i$ -th agent. Hence, average sum of rewards returned by the environment over episodes from initial to goal state with a finite number of steps.

### 3.4 Transfer Filtering Criteria

Another crucial part of *EF-OnTL* is accurately selecting the knowledge to be transferred in each source and target transfer, as shown in line 17 of Algorithm 1. In this section we introduce *Transfer Filtering Criteria (TM)* to prioritise certain tuples based on their score over a set of measures.

To simplify the readability of this section we describe the simplest case with two agents, which, at each transfer step, are referred to as source  $A_s$  and target  $A_t$ . We rely on two criteria to identify experiences to be transferred, *uncertainty* and *expected surprise*. *Expected surprise* [25], is defined over target agent and is approximated through *Temporal Difference-error* (TD-error). For instance, given a DQN with predictor and target networks, TD-error is defined as Mean Squared Error (MSE) between the two outputs.

Assuming that  $UE_s$  and  $UE_t$  use a standard method with fixed parameters, we use their estimated uncertainties to select relevant tuples to be transferred. In detail, given  $(o_s^i, u_s^i)$ , an observation with its epistemic uncertainty explored by  $A_s$  at time  $i$ , then  $A_t$  estimates current uncertainty  $u_t^{o_s^i}$  over the interaction  $o_s^i$  sampled by  $A_s$ . Hence,  $u_t^{o_s^i} = UE_t(o_s^i)$ . We then define  $\Delta$ -conf as discrepancy between the two estimations:  $\Delta$ -conf =  $u_t^{o_s^i} - u_s^i$ . Note that  $u_t^{o_s^i}$  is an epistemic estimation at the time of transfer and changes over time.

As a result,  $A_t$  receives a personalised batch of experience that aims to fill shortcomings in its belief. Based on these two criteria we can define multiple filtering functions for incoming knowledge. In this paper, we introduced the following criteria to prioritise transferred experience:

- *rnd  $\Delta$ -conf* – *Random transfer from Delta Confidence with threshold*. By choosing this technique, an agent randomly samples  $B$  interactions with an associated  $\Delta$ -conf above the median computed across all samples.
- *high  $\Delta$ -conf* – *transfer Higher Delta Confidence*. Agent selects the top  $B$  entries with higher  $\Delta$ -conf.
- *loss & conf* – *higher Loss and Confidence mixed*. While previous filters are defined only over uncertainty, this one also considers expected surprise. To balance the different scales, values are nor-

malised within a  $[0, 1]$  interval and then summed up. Finally, target agent pulls  $B$  tuples from the buffer with the highest values.

The criteria presented above are explicitly or implicitly used in other TL frameworks. Nonetheless, we are the first to explicitly apply them to online transfer learning based on experience sharing, and evaluate *EF-OnTL* performance while exploring different combinations of *SS* and *TM* to assess their impact.

## 4 Evaluation Setup

We demonstrate *EF-OnTL* by using two different off-policy Deep-RL methods: the Dueling DQN [24] in Cart-Pole and MT-PP and Deep Deterministic Policy Gradient with parametrised action space (PA-DDPG) [9] for HFO.

**Table 2:** *Expert-Free Online Transfer Learning* parameters.

Parameter	Cart-Pole	MT-PP	HFO
N	5	4	3
TF	200	300	400
TB Capacity	10,000	100,000	25,000
SS:BP Eval. Interval	200	400	$\times$
Ep. Start Transfer	600	2,500	2,400
Max Timestep	400	200	500
Max Episode	1,800	8,000	20,000

Table 2 presents the parameters used to assess *EF-OnTL* across the evaluated tasks. In addition, for Cart-Pole and MT-PP we also vary  $B : \{500, 1500, 5000\}$ ,  $SS : \{\bar{U}, BP\}$ , and  $TM : \{rnd \Delta$ -conf, high  $\Delta$ -conf, loss & conf $\}$ , obtaining 18 transfer settings defined on  $B \times SS \times TM$ . In HFO, we have constrained the evaluation to a single transfer setting which has proved the best trade-off between positive transfer and cost of transfer over simpler environments, with the following parameters:  $B = 100$ ,  $SS = \bar{U}$  and  $TM = high \Delta$ -conf.

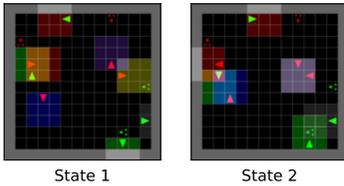
While MT-PP and HFO are naturally multi-agent environments, for Cart-Pole we used multiple parallel and independent instances. We collected 20 runs for Cart-Pole and MT-PP to study the impact of different transfer settings and 10 runs for HFO.

We compare *EF-OnTL* against *no-transfer*, where transfer is disabled, and two different advice-based baselines with and without external expertise:

- *OCMAS – Online Confidence-Moderated Advice Sharing*, based on [11], to evaluate the impact of sharing action-advice versus experience. For each state visited by an agent  $A_i$ , uncertainties are estimated by all agents and whether  $A_i$  is the most uncertain then it asks for advice. Other agents provide their best estimated action and  $A_i$  uses majority voting to take the final action. Allocated budget matches the number of our total interaction transferred per agent with  $B = 5,000$ , hence, 30,000 for Cart-Pole and 85,000 for MT-PP.
- *RCMP – Requesting Confidence-Moderated Policy advice* [5] to benchmark our expert-free transfer framework against the state-of-the-art expert-based teacher-student framework. We provide a jury, composed by 3 trained agents, to ward off any bias induced by using a single trained agent as teacher. An agent asks for an advice whenever its uncertainty is higher than a threshold and advice is selected as the most frequent action returned from the jury. Advice is then given until episode terminates.

We implemented *RCMP* with 5 heads on advantage branch for dueling DQN. Each head provides  $|A|$  estimation of advantages and we approximate uncertainty by normalising head’s logit within a  $[0, 1]$  interval, as such, uncertainty ranges between  $[0, 0.3]$ . We set the threshold for Cart-Pole to 0.25 and allocated a budget of 300 episodes. For MT-PP, after several trials, threshold is set to 0.02 and maximum budget to 6,000 episodes. For consistency with *OCMAS*, in Section 5 we report used budget in terms of overall action-advice suggested.

For MT-PP, different versions exist and for our configuration we rely on a grid-based implementation [2]. A depiction of this problem is reported in Figure 2. 8 predators and 4 prey are fairly divided into two colour-based teams and randomly spread in a  $12 \times 12$  grid.



**Figure 2:** Multi-Team Predator-Prey (MT-PP) environment. 8 predators (filled oriented triangles), and 4 prey (dotted oriented triangles), are fairly split into two colour-based teams.

We control each predator with a RL agent while prey follow a random policy. An episode terminates when one or both teams have no more prey left to capture. Each agent perceives a  $3 \times 3$  grid centered around the first cell it is facing, as highlighted in Figure 2. Thus, observation is composed of a 3-dimensional  $3 \times 3$  matrix. First channel describes object type, i.e., *void*, *wall*, *predator*, *prey*. Second channel identifies the team-membership, i.e., *none* when not applicable, *red* and *green*. Third channel provides rotation of an object, i.e., *none* when non applicable, *up*, *down*, *left*, *right*. Predators have 5 possible actions: *rotate left*, *rotate right*, *forward*, *pick* and *hold*. First two actions rotates a predator left or right while remaining in the same cell. *Forward* moves an agent to the next cell it is facing, if it is empty. When facing a prey, *pick* is used to catch the prey.

Finally, while all predators are controlled by RL agents, only a single team is enabled to share knowledge, in order to better isolate performance impact of sharing. Results are therefore presented as aggregated performance across the *sharing-enabled* team.

**Table 3:** Dueling Q-Network parameters.

Parameter	Cart-Pole	MT-PP
Input Layer	FC(4, 128)	Conv1d(3, 7, $k = 1$ )
Hidden Layer(s)	FC(128, 64)	Conv1d(7, 15, $k = 1$ ) FC(135, 256)
Value Branch	FC(64, 1)	FC(256, 1)
Adv. Branch	FC(64, 2)	FC(256, 5)
Activation	ReLU	ReLU
Loss	MSE	MSE
Optimiser	Adam	Adam
Learning Rate	$1e - 5$	$1e - 5$
Target Replacement	1,000	10,000
Exploration	Softmax	$\epsilon$ -decay
Replay Buffer Size	10,000	10,000
Batch Size	32	32
Ep. Start Training	100	100

Dueling DQN architectures are specified in Table 3 while approximation of Q-values ( $Q$ ) is given by combining the two output streams

of the network as follow  $Q = value + adv - \overline{adv}$ . In Cart-Pole we balance exploration and exploitation through *Softmax* function [18] while in MT-PP we follow an  $\epsilon$ -decay probability that eventually anneals to 0 at episode 7,450. Model is optimised at each time-step after the 100<sup>th</sup> episode. As replay buffer, we used the proportional version of prioritised experience replay with suggested parameters [16].



**Figure 3:** Instance of Half Field Offense 3v2. Offense is depicted in yellow, defense in red and goalie in violet.

Lastly, we compare *EF-OnTL* against a no-transfer baseline and *RCMP* in HFO. We omitted *OCMAS* as synchronizing all the player agents, in a complex scenario, for each time-step to evaluate their confidence is unfeasible. Our HFO setup is a 3v2 where offense players are controlled by our agents with noise free perception and defense by *HELIOS* baseline [8]. Figure 3 depicts a snapshot of our configuration.

For *RCMP*, we empirically found it more advantageous to provide single-step advice rather than full trajectory needed to accomplish an episode. Thus, we allocated a maximum budget of 10,000 per agent and we fixed the uncertainty threshold to  $1e - 4$ .

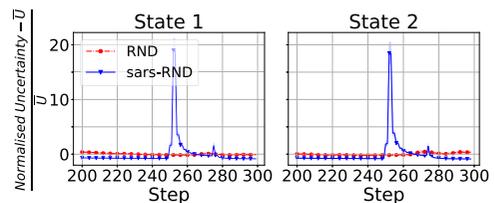
As PA-DDPG architecture for HFO, we use an existing configuration publicly available [27].

## 5 Results

This section presents the results and analysis of *EF-OnTL* performance. As a preliminary study, we also first evaluate the performance of our proposed sars-RND uncertainty estimator against RND, to compare the sensitiveness of the two models.

### 5.1 sars-RND Evaluation

To assess sars-RND capability we use the MT-PP benchmark. Each estimator visits 2 states depicted by Fig. 2 while sampling different actions. We fixed architectures of both target and predictor across the two estimators, but their input layer dimensions differ as sars-RND requires extra neurons to accommodate action, reward and next state. Agents perform action  $a_0$  for 250 steps to strengthen the prediction of both uncertainty estimators. Afterward, action sampled is changed to  $a_1$  for 25 steps. Finally, agents resample  $a_0$  for 25 steps.



**Figure 4:** RND and sars-RND normalised estimated uncertainty with their 95% confidence interval over states depicted by Figure 2 with different actions sampled over time.

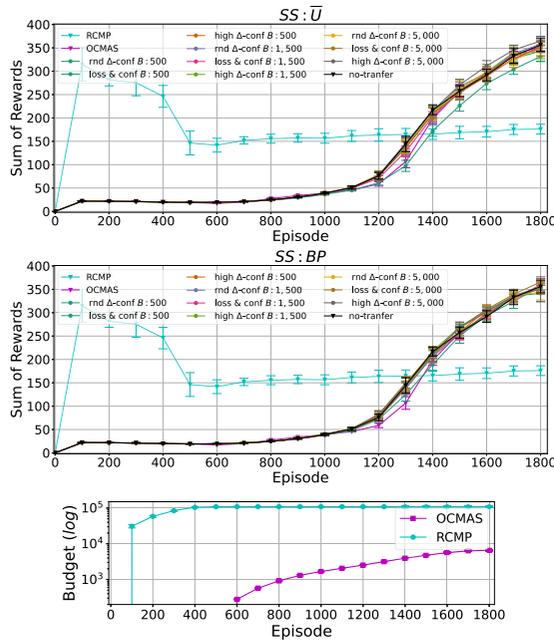
Overall, we observed higher uncertainty within the first steps while curves decrease asymptotically to 0 for both estimators as they increasingly visit the states. However, by evaluating uncertainty in a narrowed interval as in Figure 4, we notice that while RND keeps a flat trend, sars-RND registers a spike in uncertainty on the change of action, firstly, at step 250 when action  $a_0$  is replaced by  $a_1$ , and secondly, at 275 when  $a_0$  is resampled. Second spike is smaller compared to first as sars-RND recognises the previous seen interaction.

This confirms that sars-RND provides a more reliable uncertainty estimation by considering executed action alongside visited state while RND is not sensitive to action variation. Furthermore, given the observed decreasing trend, sars-RND generalises well enough to recognise familiar states even while sampling a different action. Thus, uncertainty estimated for previously un-encountered state is higher than previously seen state while exploring different actions.

## 5.2 EF-OnTL Evaluation

**Cart-Pole results** Cart-Pole is a simple environment with a continuous observation space and a binary decision for the control policy. We evaluate the performance of our TL technique by evaluating the learning curve over the episodes.

Given the minimalistic environment, generally we have observed very similar performance across all the techniques: *no-transfer*, action-advice based baseline and *EF-OnTL*. Figure 5 shows the



**Figure 5:** Learning curves in Cart-Pole and budget used by advice-based baselines.

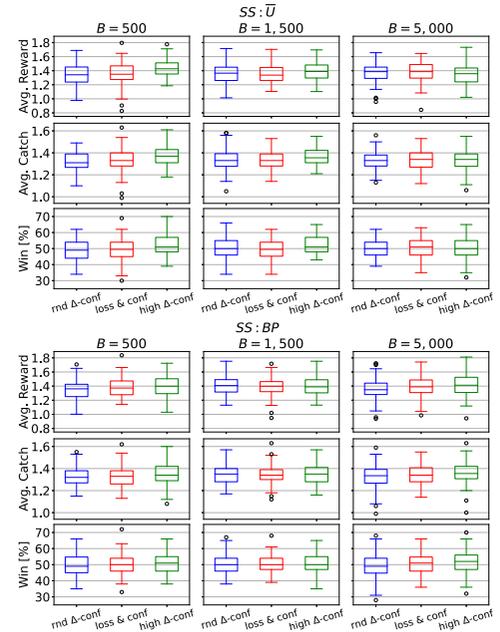
learning curves with their 95% confidence interval.

We observe a fluctuating transfer impact for *EF-OnTL* achieving both positive and negative transfer. *EF-OnTL* positive transfer starts around episode 1,400, and its benefit is generally linear to quantity of shared information. Despite the mild improvement, we cannot establish a configuration that prevails overall but we observe that different transfer settings impact differently the performance throughout agent’s life. We conclude that in this scenario, *EF-OnTL* has room for improvement by enabling dynamic selection of techniques used to select both the source and experience to transfer.

Based on our evaluation, *RCMP* shows jump-start within first episodes (0 – 1,200) while later performance is kept to a similar value until the end. Thus, despite the initial improvement, at the end it caps the target’s performance.

**Multi-Team Predator-Prey results** While in Cart-Pole we observed low improvements by our approach, MT-PP shows a more interesting outcome. We measured the performance of a sharing-enabled team in MT-PP through *Average Reward*, *Average Catch* and *Win Probability*. We do not report the distribution of miscaught prey since all *EF-OnTL* agents learn not to catch other team’s prey.

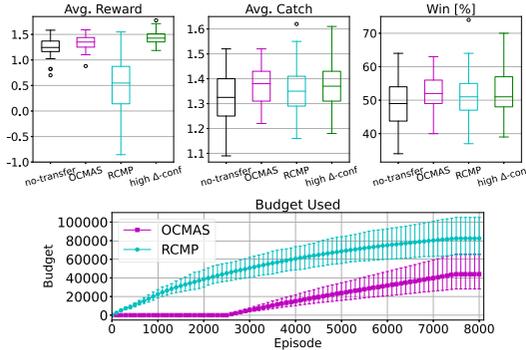
First, we report performance across the 18 transfer settings defined for *EF-OnTL* in Figure 6. All predators learnt a similar behaviour to fulfil their goal. Generally, all variations have a common median value across all performance metrics and the difference between one *EF-OnTL* configuration to the other is very small. This common trend is opposed to what we previously observed in Cart-Pole where the benefit of transfer was generally correlated to budget used. This could suggest that for more challenging environment a lower budget may be the best trade-off to keep low transfer overhead cost and to improve agent’s performance.



**Figure 6:** Performance achieved by *EF-OnTL* transfer-enabled team in MT-PP across 500 test episodes.

Furthermore, while fixing budget  $B$  and  $SS$ , generally using *rnd*  $\Delta$ -*conf* performs the worst, *high*  $\Delta$ -*conf* performs the best, while *loss & conf* usually lies in between. Thus, overall from start of sharing to end, filtering experience to share by discrepancy between source and target uncertainty (i.e., *high*  $\Delta$ -*conf*), is the most effective in this type of environment.

Figure 7 shows the comparison between *no-transfer*, *OCMAS* and *RCMP* baselines, and the best *EF-OnTL* configuration:  $B = 500$ ,  $SS = \bar{U}$  and  $TM = \text{high } \Delta\text{-conf}$ . Contrary to what we expected, *RCMP* does not significantly outperform the others and actually decreases the average reward. With *RCMP* predator has learnt to perform more frequently expensive actions, i.e., catch. Despite this different behaviour leads to similar win and positive catch distributions, overtime leads also to lose reward and to miscatch preys.



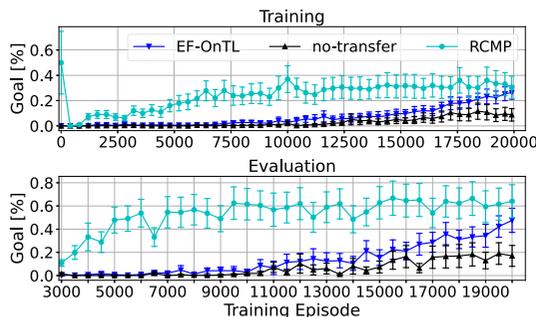
**Figure 7:** Team based *EF-OnTL* performance in MT-PP against the baselines with budget consumption.

**Half Field Offense results** HFO is a multi-agent environment where agents need to collaborate to score a goal. Collaborative interaction complicates the task that already has sparse reward function and continuous state and control space. Thus, in a 3 vs 2 situation, having a single poor performing agent within the team will likely result in capping the team performance.

To evaluate *EF-OnTL* against *no-transfer* and *RCMP* baseline, we measure probability of the team scoring a goal during training and testing phases. We estimate probabilities by counting the number of accomplished episodes within a sample. For training we use a sliding window of 100 episodes while on testing we run 500 episodes where agents use their up-to-date knowledge. These results are reported by Figure 8.

*EF-OnTL* shows good performance in both training and testing when compared against *no-transfer* and their difference increases over time. Within the latest 5,000 episodes, *EF-OnTL* powered agents, on average, doubled their goal probability when compared against *no-transfer* agents.

In HFO the exploitation of expert knowledge plays a crucial role. In fact, *RCMP* achieves outstanding results since an early stage of training. However, despite an initial jump-start, *RCMP* goal probability seems to fluctuate within the same range from episode 7,000 onwards resembling Cart-Pole trend.



**Figure 8:** HFO goal probability during training and testing.

***RCMP* additional notes** While implementing *RCMP* as baseline, we encountered some challenges to be addressed that are worth to be mentioned. Firstly, in Cart-Pole, given the normalisation to maintain a fixed scale for uncertainty and binary action decision, estimator shows low granularity in uncertainty estimation. In fact, uncertainty falls into 3 possibilities based on agreement across the 5

heads, 0, when all agree, 0.2 when a single head disagree, and 0.3 when two heads disagree with others. Secondly, in environment with multiple actions as MT-PP, the algorithm performance are very sensitive to the chosen threshold. Despite a lower threshold leads agent to be guided for longer, the final performance in learning curve and tracked metrics are significantly lower. We provide a deeper analysis on *RCMP* baseline in supplementary material [30]. Finally, in HFO expert teachers significantly improve student’s performance. Although, in PA-DDPG, the ensemble is on critic which outputs a single value and as such is impossible to predict a range and hence to normalise the estimated uncertainty. Consequently, the uncertainty curve decreases sharply within the first hundreds episodes and then ranges within a narrow interval making it again not trivial to find the right threshold given the tight uncertainty interval.

To conclude, while *RCMP* requires careful tuning on threshold, our proposed *EF-OnTL* does not require tuning or adaptation other than budget  $B$  and transfer frequency  $TF$ . Furthermore, *EF-OnTL* has shown comparable performance with both *OCMAS* and *RCMP* in Cart-Pole and MT-PP while significantly lowering the communication and synchronisation cost by only sharing every  $N$  episodes a personalised batch of experience tailored to target agent gaps.

In HFO, the use of *EF-OnTL* led the offense team to score more than twice as many goals as the *no-transfer* agents. Furthermore, in the latest 5,000 episodes, the gap between *RCMP* and our novel *Expert-Free Online Transfer Learning* faded away over time.

## 6 Conclusion

This paper introduced *Expert-Free Online Transfer Learning (EF-OnTL)*, a novel dynamic online transfer learning framework based on experience sharing and suitable for multi-agent implementation, and *State Action Reward Next-State Random Network Distillation (sars-RND)*, an extension of RND, to estimate agent epistemic uncertainty from a full RL interaction. *EF-OnTL* is evaluated varying budget used, teacher selection criteria and experience filtering criteria to improve shared batch quality across multiple agents.

We benchmark *EF-OnTL* against *no-transfer*, *RCMP* and *OCMAS* in three environments, Cart-Pole, Multi-Team Predator-Prey and Half Field Offense.

*EF-OnTL* has shown a significant improvement when compared against *no-transfer* baseline and overall similar, and in some conditions superior, performance when compared against advice-based baselines. Despite the minor improvement in performance, the communication cost is significantly reduced by sharing less frequently an experience batch from a temporary selected teacher to a target agent. The transferred batch is personalised upon the expertise of source and target shortcomings. Finally, this paper open up several directions for future work. Firstly, in our experiments on *EF-OnTL* we did not cover any empirical evaluation on transfer frequency nor optimal batch-size. Instead, we fixed this values to some pre-defined parameters as we preferred to study source selection and transfer filtering criteria, presented respectively in Section 3.3 and Section 3.4, used to design the transferred batch. Furthermore, experiments in Cart-Pole have highlighted how aforementioned criteria and budget affect TL performance throughout different stage of learning. This requires further study to assess *EF-OnTL* influence over target while enabling dynamic tuning of these parameters. Secondly, despite sars-RND can be dropped to minimize used resources when there is no-need of uncertainty, i.e., on deployment, sars-RND could be optimised by limiting the input to state action pairs rather than full agent-environment interaction. The proof of equivalence while reducing the input is yet

to be provided. Thirdly, we assumed that agents are homogeneously defined and all are engaged during a transfer step. Thus, their task is equally defined in observation space, action space, goal, reward model and transition function. This paper does not cover adaption of *EF-OnTL* and sars-RND across different tasks or goals. Lastly, uncertainty-labelled experience could be stored and used later to train a novice agent in an offline scenario.

## Acknowledgements

This work was sponsored, in part, by the Science Foundation Ireland under Grant No. 18/CRT/6223 (Centre for Research Training in Artificial Intelligence), and SFI Frontiers for the Future project 21/FFP-A/8957.

## References

- [1] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov, ‘Exploration by random network distillation’, *arXiv preprint arXiv:1810.12894*, (2018).
- [2] Alberto Castagna and Ivana Dusparic, ‘Multi-agent transfer learning in reinforcement learning-based ride-sharing systems’, in *ICAART*, (2022).
- [3] Gabriel V Cruz Jr, Yunshu Du, and Matthew E Taylor, ‘Pre-training neural networks with human demonstrations for deep reinforcement learning’, *arXiv preprint arXiv:1709.04083*, (2017).
- [4] Felipe Leno Da Silva, Ruben Glatt, and Anna Helena Real Costa, ‘Simultaneously learning and advising in multiagent reinforcement learning’, in *Proceedings of the 16th conference on autonomous agents and multiagent systems*, pp. 1100–1108, (2017).
- [5] Felipe Leno Da Silva, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor, ‘Uncertainty-aware action advising for deep reinforcement learning agents’, in *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 5792–5799, (2020).
- [6] Gabriel Dulac-Arnold, Nir Levine, Daniel J Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester, ‘Challenges of real-world reinforcement learning: definitions, benchmarks and analysis’, *Machine Learning*, **110**(9), 2419–2468, (2021).
- [7] V Gabriel, Yunshu Du, and Matthew E Taylor, ‘Pre-training with non-expert human demonstration for deep reinforcement learning’, *The Knowledge Engineering Review*, **34**, (2019).
- [8] Matthew Hausknecht, Prannoy Mupparaju, Sandeep Subramanian, Shivaram Kalyanakrishnan, and Peter Stone, ‘Half field offense: An environment for multiagent learning and ad hoc teamwork’, in *AAMAS Adaptive Learning Agents (ALA) Workshop*, (May 2016).
- [9] Matthew Hausknecht and Peter Stone. Deep reinforcement learning in parameterized action space, 2015.
- [10] Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Andrew Sendonaris, Gabriel Dulac-Arnold, Ian Osband, John Agapiou, et al., ‘Learning from demonstrations for real world reinforcement learning’, (2017).
- [11] Ercument Ilhan, Jeremy Gow, and Diego Perez-Liebana, ‘Teaching on a budget in multi-agent deep reinforcement learning’, in *2019 IEEE Conference on Games (CoG)*, pp. 1–8. IEEE, (2019).
- [12] Yongyuan Liang and Bangwei Li, ‘Parallel knowledge transfer in multi-agent reinforcement learning’, *arXiv preprint arXiv:2003.13085*, (2020).
- [13] Wenzhang Liu, Lu Dong, Jian Liu, and Changyin Sun, ‘Knowledge transfer in multi-agent reinforcement learning with incremental number of agents’, *Journal of Systems Engineering and Electronics*, **33**(2), 447–460, (2022).
- [14] Ashvin Nair, Bob McGrew, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel, ‘Overcoming exploration in reinforcement learning with demonstrations’, in *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 6292–6299. IEEE, (2018).
- [15] Mojtaba Norouzi, Monireh Abdoos, and Ana LC Bazzan, ‘Experience classification for transfer learning in traffic signal control’, *The Journal of Supercomputing*, **77**(1), 780–795, (2021).
- [16] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver, ‘Prioritized experience replay’, *arXiv preprint arXiv:1511.05952*, (2015).
- [17] Sriram Ganapathi Subramanian, Matthew E Taylor, Kate Larson, and Mark Crowley, ‘Multi-agent advisor q-learning’, *Journal of Artificial Intelligence Research*, **74**, 1–74, (2022).
- [18] Richard S Sutton, Andrew G Barto, et al., ‘Introduction to reinforcement learning’, (1998).
- [19] Adam Taylor, Ivana Dusparic, Maxime Guériau, and Siobhán Clarke, ‘Parallel transfer learning in multi-agent systems: What, when and how to transfer?’, in *2019 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, (2019).
- [20] Matthew E Taylor and AI Borealis, ‘Improving reinforcement learning with human input’, in *IJCAI*, volume 328, pp. 5724–5728, (2018).
- [21] Matthew E Taylor, Nicholas Carboni, Anestis Fachantidis, Ioannis Vlahavas, and Lisa Torrey, ‘Reinforcement learning agents providing advice in complex video games’, *Connection Science*, **26**(1), 45–63, (2014).
- [22] Lisa Torrey and Matthew Taylor, ‘Teaching on a budget: Agents advising agents in reinforcement learning’, in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pp. 1053–1060, (2013).
- [23] Zhaodong Wang and Matthew E Taylor, ‘Improving reinforcement learning with confidence-based demonstrations.’, in *IJCAI*, pp. 3027–3033, (2017).
- [24] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas, ‘Dueling network architectures for deep reinforcement learning’, in *International conference on machine learning*, pp. 1995–2003. PMLR, (2016).
- [25] Adam White, Joseph Modayil, and Richard S Sutton, ‘Surprise and curiosity for big data robotics’, in *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, (2014).
- [26] Tianpei Yang, Weixun Wang, Hongyao Tang, Jianye Hao, Zhaopeng Meng, Wulong Liu, Yujing Hu, and Yingfeng Chen, ‘Learning when to transfer among agents: An efficient multiagent transfer learning framework’, *arXiv preprint ArXiv:2002.08030*, **46**, 55, (2020).
- [27] Longtao Zheng. Ltzheng/pddpg-hfo: Half field offense in robocup 2d soccer with reinforcement learning. <https://github.com/ltzheng/pddpg-hfo>, Accessed: 27 Jan. 2023.
- [28] Changxi Zhu, Yi Cai, Ho-fung Leung, and Shuyue Hu, ‘Learning by reusing previous advice in teacher-student paradigm’, in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1674–1682, (2020).
- [29] Changxi Zhu, Ho-Fung Leung, Shuyue Hu, and Yi Cai, ‘A q-values sharing framework for multi-agent reinforcement learning under budget constraint’, *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, **15**(2), 1–28, (2021).
- [30] Alberto Castagna and Ivana Dusparic, ‘Expert-Free Online Transfer Learning in Multi-Agent Reinforcement Learning’, *arXiv preprint arXiv:2303.01170*, (2023).