# Methodical Advice Collection and Reuse in Deep Reinforcement Learning

Sahir

Alberta Machine Intelligence Institute (Amii) & University of Alberta Edmonton, Alberta, Canada sahir1@ualberta.ca

> Srijita Das University of Alberta Edmonton, Alberta, Canada srijita1@ualberta.ca

# ABSTRACT

Reinforcement learning (RL) has shown great success in solving many challenging tasks via use of deep neural networks. Although using deep learning for RL brings immense representational power, it also causes a well-known sample-inefficiency problem. This means that the algorithms are data-hungry and require millions of training samples to converge to an adequate policy. One way to combat this issue is to use action advising in a teacher-student framework, where a knowledgeable teacher provides action advice to help the student. This work considers how to better leverage uncertainties about when a student should ask for advice and if the student can model the teacher to ask for less advice. The student could decide to ask for advice when it is uncertain or when both it and its model of the teacher are uncertain. In addition to this investigation, this paper introduces a new method to compute uncertainty for a deep RL agent using a secondary neural network. Our empirical results show that using dual uncertainties to drive advice collection and reuse may improve learning performance across several Atari games.

# **KEYWORDS**

Reinforcement Learning, Action Advising, Teacher-Student Framework

# **1 INTRODUCTION**

Deep RL has shown promising results in many challenging problems, ranging from the game of Go [22] to Atari video games [16]. While these problems have opened up avenues for applying RL research to a wide variety of important problems like drug discovery [10], molecular optimization [26] and healthcare [15], current Deep RL algorithms are *sample-inefficient* and require millions of interactions with the environment to converge. This restricts deploying Deep RL systems in real-world applications where acquiring training samples incur cost.

A number of well-known techniques like reward shaping [17], policy shaping [11] and imitation learning [20] have been studied to address the sample-efficiency problem by using expert knowledge Ercüment İlhan Queen Mary University of London London, England, United Kingdom e.ilhan@qmul.ac.uk

Matthew E. Taylor University of Alberta & Alberta Machine Intelligence Institute (Amii) Edmonton, Alberta, Canada matthew.e.taylor@ualberta.ca

to better guide the agent. A popular paradigm for providing expert advice is a teacher-student framework [6] where a teacher is an experienced agent, trained on the same or related task, who provides advice to a naive student agent to help the student learn quickly. We assume the teacher cannot directly transfer its knowledge into the student, but instead must use a limited advice budget to tell the student what action to take in a given state. Our work takes a step in this direction where we propose a dual uncertainty based framework to drive advice collection and reuse. On the one hand, the proposed framework allows the student to seek advice from the teacher only when required. On the other hand, the framework reuses teacher's advice wisely by building a model of the teacher and querying it whenever applicable based on the model's confidence. Much of the existing work in action-advising has focused on either student initiated [8] or teacher initiated [7] criteria to seek advice effectively. Ilhan et al. [12, 13] more effectively use an advice budget by having the student build a model of the teacher and use it when possible once the budget has been exhausted. Inspired by these directions, we propose a flexible and systematic uncertainty based framework of effectively seeking and providing advice by using the model of the teacher. Our framework is motivated by real-world teacher-student interaction, where the student directly asks for advice only if he/she is uncertain about a problem. The teacher's prior advice is used by the student to build a mental model of the teacher. Furthermore, when the teacher is unavailable (e.g., the budget is exhausted), the student uses this model of the teacher to solve a problem if it seems familiar or tries to do it on their own otherwise. Our proposed approach is based on such interactions, where the student seeks advice in uncertain regions and the model of the teacher provides advice when its uncertainty is low (states that seem familiar).

This work-in-progress paper has three contributions, introducing and evaluating: (1) An action-advising framework that allows the student to have flexibility in following either its own policy, reusing advice from the model of the teacher, or asking the teacher directly for advice; (2) a methodical uncertainty-based advice reuse approach that leverages the uncertainty of both the student and/or the student's model of the teacher; and (3) a new method of computing the uncertainty for the student agent using a secondary neural network. To the authors' surprise, while our new methods perform

Proc. of the Adaptive and Learning Agents Workshop (ALA 2022), Cruz, Hayes, da Silva, Santos (eds.), May 9-10, 2022, Online, https://ala2022.github.io/. 2022.

well, they do not outperform state-of-the-art advice reuse methods. Our hope is that this paper will encourage additional work in this area, as we believe there are still many opportunities for significant improvement in how to ask for and reuse, advice from a teacher. Additional details such as pseudocode, experimental, and hyperparameter settings, are available in the appendix of the archived version of this paper.

## 2 RELATED WORK

Learning from Demonstration (LfD) [2] and Action Advising [25] are two widely accepted solutions to the sample inefficiency problem in RL algorithms. The sample inefficiency refers to the need of numerous amount of samples to learn a policy for solving a task at hand. This problem is more evident in deep RL where neural networks are used as function approximators and require a large amount of data to converge. LfD is a technique that uses expert demonstrations to bootstrap the learning of agents. LfD is usually an offline process where the expert demonstrations, spanning as long as full trajectories, are provided to the learning agent prior to the start of actual training. Action advising however, is an online technique where an RL agent is provided with actions as advice via another teacher agent or human expert who is optimal or suboptimal.

Action advising. Action advising paradigm is used extensively in student-teacher framework [1, 6, 8, 12, 13, 24, 25]. Typically, the student agent is considered a novice that could potentially perform better in the presence of a teacher agent that is pre-trained and known to perform well in the given task. The teacher is available for a limited number of interactions which is usually referred to as the advising budget. Torrey et al. [25] introduced action advising along with multiple heuristics for deciding how to provide advice. They mainly introduced teacher-driven advising methods such as Early Advising and Importance Advising, among others. Studentdriven advising is also explored in the literature where heuristics like epistemic uncertainty of the student agent [8, 18] and advice novelty [14] are used to decide when to ask for advice. To reduce the overhead of the teacher, jointly-initiated action advising was investigated by Amir et al. [1], where a student could seek advice and the teacher could provide advice to the student's queries. Moreover, the student-teacher framework has also been extended by Da Silva et al. [7] to accommodate multiple agents, without having to fix roles of being a student or teacher, and training simultaneously using advice from each other. Lastly, Omidshafiei et al. [19] used a similar extension of the framework to propose multiple objective functions to learn when and what to advise instead of relying on heuristics.

Advice reuse. Given the vast literature on action advising, there exists limited work that aims to collect advice for later reuse. This is particularly useful because it allows the student to spend the teacher's budget wisely by not asking redundant or similar queries. The idea of reusing collected advice was introduced by Zhu et al. [27] in a student-teacher framework where advice were reused based on different heuristics such as *QChange, Budget reuse*, and *Decay Reusing Probability* in tabular RL algorithms. Due to the tabular nature, they store the action advice in a fixed size buffer

for reuse. Our work is heavily inspired by the deep RL compatible work of Ilhan et al. [12, 13], where supervised learning is used to train a model of the teacher from previously collected advice. This model would then predict an action (similar to the teacher) which could be reused if the predetermined probability threshold is met. The limitation of their work includes the student's dependency on the model of the teacher for advice collection. As the student would continue to experience more states in the environment, its measure of uncertainty would account for a larger subset of the state space. We propose a principled framework to address these issues.

#### 3 BACKGROUND

Reinforcement learning. We follow the standard RL framework [23] which is modelled as a Markov Decision Process (MDP). It is represented by the tuple  $(S, \mathcal{A}, R, P, \gamma)$  where S is the state space,  $\mathcal{A}$  is the action space, R is the reward function, P is the state-transition probability and  $\gamma \in [0, 1)$  is the discount factor for infinite horizon problems. At time-step t, an RL agent starts in a state  $s_t$ , takes an action  $a_t$ , receives a reward  $r_{t+1}$  after interacting with the environment and transitions to the next state  $s_{t+1}$ . At each time-step, the agent tries maximizing the expected return  $\mathcal{R}_t = \mathbb{E}[\sum_{k=0}^T \gamma^k r_{t+k+1}]$ , which is the discounted sum of rewards that an agent receives when starting from time-step t and following a policy  $\pi$  until time step T.

Deep Q-Networks. Deep Q-Networks (DQN) [16] is a state-ofthe-art off-policy deep RL algorithm to approximate the Q-function for high dimensional tasks with continuous state-space and discrete actions. An underlying neural network is used as a function approximator and the loss function for the iteration *i* is  $L_i(\theta_i) = \mathbb{E}_{s,a\sim P}[y_i - Q(s, a; \theta_i)]^2$  which is the squared error between the target  $y_i$  and the Q-value output by the current model parameterized by  $\theta_i$ . The Bellman equation is used to approximate the target value  $y_i = r + \gamma \max_{a'} Q(s', a'; \theta_{i-1})$  where s' and a' refers to the next state-action pair, respectively. An experience replay buffer is used to store all the agent interactions. At every training step, a mini-batch of experience is sampled from this buffer to update the parameters of the Q-function.

Advice reuse using behavior cloning. Imitation learning [21] aims at learning the policy of an expert by collecting demonstrations. Behavior Cloning [20] is a specific approach under Imitation learning that uses supervised learning to approximate the conditional distributions of actions given the state. Deep neural networks have been used for behavior cloning where the objective is to minimize the negative log-likelihood,  $\mathcal{L}(\eta) = \sum_{(s,a) \in \mathcal{T}} -\log M(a|s, \eta)$  where  $\mathcal{L}(\eta)$  is the loss-function,  $\mathcal{T}$  refers to the demonstration data and  $\eta$  is the parameter of the supervised learning model *M*. Ilhan et al. [12] used a behavior cloning module to train a neural network using the advice collected from a teacher agent which was referred to as Advice Imitation. It was done to enable the student to replicate the teacher's action decisions without needing to query it. Furthermore, epistemic uncertainty via employing dropout regularization in this imitation model lets the student to avoid using the advising budget for similar states for which prior advice had already been taken. We use a similar behavior cloning module for advice reuse in our proposed framework.

# 4 UNCERTAINTY-DRIVEN STUDENT-TEACHER FRAMEWORK FOR ADVICE COLLECTION AND REUSE

This section discusses the problem formulation, introduces the proposed algorithms, and introduces the new way of estimating epistemic uncertainty.

#### 4.1 **Problem formulation**

We propose a methodical student-teacher framework for action advising that aims to wisely reuse advice and minimize redundant advice requests made to the teacher. The focus of this work is to compute and use uncertainty as a metric to decide whether to ask for advice from the teacher, reuse advice from the model of the teacher, or simply let the student continue following its own policy. However, the framework is meant to be general, and any valid metric like *value of information* [4] can be used to drive this decision making To formalize the proposed problem:

**Given:** A student agent with its own policy  $\pi_S$ , a skilled teacher agent  $\pi_T$ , advice budget *b* which specifies the number of times the teacher *T* can be queried by the student for action-advice

**To-do:** learn optimal student policy  $\pi_S^*$  by leveraging advice from the teacher following  $\pi_T$ , reusing the model of the teacher build from prior advice  $M_\eta$  or following its own policy  $\pi_S$  subject to the budget *b* 

## 4.2 Student Uncertainty driven action-advice and reuse

We propose two algorithms that use uncertainties to drive decision making. In all algorithms, there are two ways to leverage the computed uncertainty. One way is to use a fixed uncertainty threshold to decide whether the student or model of the teacher is uncertain. The other way is to have a dynamically changing threshold based on the previously observed uncertainty values. We keep dynamic uncertainty thresholds for our experiments. The proposed algorithms are as below:

- (1) Student's Uncertainty-driven Advising (SUA)
- (2) Student's Uncertainty-driven Advising with Advice Imitation & Reuse (SUA-AIR)

SUA and SUA-AIR, use the student's uncertainty estimates to drive the advice collection process. Student agent's uncertainty to drive action advising has been explored previously in literature, as mentioned in Section 2. To the best of our knowledge, student uncertainty driven advice collection in an advice-reuse framework hasn't been explored previously. In addition, the process we use to compute the student's uncertainty is novel, which has its advantages, discussed in more detail in Section 4.3.

The flow of decision-making involved in **SUA** is shown in Figure 1. At any state, the student agent can ask the teacher agent for advice if its uncertainty  $u_s$  is greater than the student's adaptive uncertainty threshold  $c_1$ . In other words, the student can ask the teacher for advice when it is uncertain. The teacher agent can then provide action advice  $a_t$  if the teaching budget is not consumed (b > 0). Lastly, the student agent would continue to follow its own policy if no advice is received  $(a_t \text{ is None})$ . This is possible if either

the student agent is certain in the given state  $u_s \leq c_1$ , or if the teaching budget is consumed (b = 0).

SUA-AIR, on the other hand, uses the uncertainty of the student for advice collection and the uncertainty of the model of the teacher for advice reuse. The extent of flexibility available in SUA-AIR and the conditions that trigger those choices are shown in Figure 2. At any state, the student agent can ask the teacher agent for advice if its uncertainty  $u_s$  is greater than the student's adaptive uncertainty threshold  $c_1$ . The teacher agent will then provide action advice  $a_t$ for the queried state if the teaching budget is not consumed (b > 0). The student agent can ask for advice reuse from the model of the teacher if action advice  $a_t$  is not received (budget exhausted). The model of the teacher would provide advice if the model uncertainty  $u_m$  is less than its uncertainty threshold  $c_2$ , or simply if the model is certain, and if advice reuse is enabled (reuse\_enabled = True). At any step, advice reuse would be enabled if reuse probability  $\rho$  is greater than a random probability. Lastly, the student agent would follow its own policy if action  $a_t$  is not determined.







Figure 2: Flow of SUA-AIR.

The pseudocode for SUA and SUA-AIR are shown in the appendix of the archived version of this paper. SUA serves as a baseline advising method for our experiments. SUA-AIR builds upon SUA and AIR to leverage advice reuse with our proposed method of student's uncertainty-driven advising. This allows the student agent to be more independent in asking the teacher for advice, instead of random advice collection or relying on the model of the teacher to get advice as in [12, 13]. Moreover, the uncertainty estimates of the student agent would account for a larger subset of the state-space, enabling more accurate estimates.

#### 4.3 Computing Student's Epistemic Uncertainty

The computation of uncertainty occurs at two distinct places, one within the student agent and the other within the model of the teacher. This is evident in Figure 2, where  $u_s$  denotes the uncertainty of the student and  $u_m$  denotes the uncertainty of the model of the teacher.

The way of computing uncertainty is different at each place. As done in [13], the model of the teacher that is trained in a supervised fashion includes dropout to compute the uncertainty over the action probabilities. However, using this strategy to compute student agent's uncertainty would come at a cost. Unless there is a generalisation goal, simply adding a dropout layer in an RL agent can potentially hurt the performance of the learning agent due to the introduction of extra variance that is already plentiful in RL [9]. Thus, we propose a new way of computing the RL model's uncertainty by using a secondary neural network with dropout that mimics its learning.

The secondary neural network is trained simultaneously using the samples and targets of the student agent's RL algorithm. Since this network is updated over the same data seen by the learning agent, it can be assumed that its uncertainty estimations can serve as a good proxy to the student agent's uncertainty. Another benefit of having this secondary network approach is that it is easy to decouple. For example, if the student does not require advising (No Advising) or if the uncertainty estimates are not needed, the secondary network can simply be turned off.

Formally, the uncertainty is computed from the secondary neural network by conducting the following steps. First, *N* forward passes are performed for a given state *s*. This would give us the matrix  $\mathbf{F} \in \mathbb{R}^{N \times A}$ , which contains Q-values for each action  $a \in [a_1, a_2, ..., a_A]$ , for each forward pass  $i \in [1, 2, ..., N]$ :

$$\mathbf{F} = \begin{bmatrix} Q_1(s, a_1) Q_1(s, a_2) \dots Q_1(s, a_A) \\ Q_2(s, a_1) Q_2(s, a_2) \dots Q_2(s, a_A) \\ \vdots \\ Q_N(s, a_1) Q_N(s, a_2) \dots Q_N(s, a_A) \end{bmatrix}$$

where *A* is the total number of possible actions. The matrix **F** is then used to compute the variance across each column or the Q-values  $Q_i(s, a)$  for each action *a*. This results in a row matrix  $\mathbf{Q} \in \mathbb{R}^{:1 \times A}$ :

$$\mathbf{Q} = \left[ var(Q(s, a_1)) var(Q(s, a_2)) \dots var(Q(s, a_A)) \right]$$

The average of the values in matrix **Q** then gives us the uncertainty  $u_s$ :

$$u_s = \frac{\sum_{j=1}^{A} var(Q(s, a_j))}{A}$$

Computing uncertainty in this manner does not restrict the architecture of the student RL agent to have built in modifications such as dropout or multiple heads. The secondary network, equipped with dropout, is separate from the student RL agent architecture. Moreover, we use adaptive uncertainty thresholds to avoid tuning domain-specific uncertainty thresholds to enable scalability of our methods.

## **5 EXPERIMENTS**

The proposed experiments were designed to answer the following research questions:

- **RQ1** Does adding uncertainties of both the student and the model of the teacher help improve student learning?
- **RQ2** How do the proposed algorithms perform against the baselines?
- **RQ3** Do the proposed algorithms effect the way teacher advice is requested or reused?

#### 5.1 Testing Environments

To evaluate our proposed algorithms against some baseline heuristics in action advising, we have selected five popular domains from the Arcade Learning Environment [3]: Pong, Freeway, Seaquest, Enduro, and Q\*bert.

Each environment operates in an RGB pixel space with an observation dimension of  $160 \times 120 \times 3$ . To reduce the complexities, these observations are pre-processed by converting them to grayscale and shrinking the dimensions to  $80 \times 80 \times 1$  via interpolation. Moreover, frames are skipped by repeating the agent actions for 4 consecutive frames to account for the high frame rate. Four resulting frames (or 16 in total) are then stacked upon each other to remove the effect of partial observability, making the final observation dimension to be  $80 \times 80 \times 4$ .

#### 5.2 Experimental Setup

All agents are trained for 5 million training frames and evaluated every 50,000 steps for 10 trials each. The following agents were tested on each domain for 10 independent runs:

- No Advising (NA): A student agent with no advising (no teacher). Early Advising (EA): A student agent with advising provided in early phase of training until the the teaching budget exhausts.
- **Random Advising (RA):** A student agent advised with 0.5 probability at every step until the budget runs out.
- **Advice Imitation & Reuse (AIR):** Previous state-of-the-art baseline [13] that uses the uncertainty of the model of the teacher to drive advice collection and reuse.
- **Student's Uncertainty-driven Advising (SUA):** A student agent that uses its adaptive uncertainty estimates to drive advice collection; no advice reuse.
- Student's Uncertainty-driven Advising with Advice Imitation & Reuse (SUA-AIR): A student agent that uses its adaptive uncertainty estimates to drive advice collection; paired with a teacher imitation model that uses its adaptive uncertainty thresholds for advice reuse.

All student agents use the same architecture, Double DQN with 3 convolutional layers followed by a fully-connected hidden layer and a dueling output. All agents use the  $\epsilon$ -greedy strategy for exploration where  $\epsilon$  is decayed over time. The secondary network's structure is similar to student DQN with two added differences. First, the secondary network is a supervised learning model, and second, has a dropout layer.<sup>1</sup> Dropout rate controls the percentage

<sup>&</sup>lt;sup>1</sup>Following the literature [5], the dropout rate is set to 0.2 so that it does not hamper the learning progress.



Figure 3: Evaluation rewards for all agents tested on Enduro, Freeway, Pong, Q\*bert, and Seaquest.



Figure 4: Advice reuse (actions taken from the model of the teacher) and advice taken (actions taken directly from the teacher agent) in every 100 steps taken by the student in all domains.

of units that are dropped at every training step and the number of forward passes performed to compute the epistemic uncertainty is set to 100.

The supervised learning model of teacher (or imitation model), trained with the student-teacher interaction data, is equipped with the network structure identical to the student agent's secondary network. Here, the model of teacher predicts action probabilities (instead of Q-values in the secondary network). A dropout layer is also added to the model. Following AIR [13], the dropout rate is 0.35. Again, the number of forward passes to compute the epistemic uncertainty is set to 100. The list of all the hyperparamters for the student agents, and the model of the teacher (or imitation model) are shown in the appendix of the archived version of this paper.

For each game, a teacher agent is pre-trained and has a fixed policy. The teacher agents have the same network structure and algorithm as the student agent. The teacher agents can be considered competent, as compared to results in the literature, and obtain average evaluation scores of 1556 for Enduro, 28.8 for Freeway, 12 for Pong, 3705 for Q\*bert, and 8178 for Seaquest.

## **6 RESULTS**

The evaluation performance of all the agents in Pong, Freeway, Seaquest, Enduro, and Q\*bert is reported in Figure 3. All advised agents (except RA) performed better than no advising agent (NA) during the early training phases, with statistically significant differences as they are more than twice the standard error in the respective means (for 95% confidence), in all domains except Q\*bert, where the benefits of advising appear in later stages of training.

In Enduro, the performance of all advised agents during the early phases of training was better than NA agent with statistically significant differences. The benefits of advice reuse become slightly more apparent when we look at the performance of agents in Freeway. Students with advice reuse, SUA-AIR and AIR, showed a statistically significant boost in performance during early-to-mid training, as compared to other agents. Advising methods such as EA, RA, and SUA took more time to catch up to the policies of students using advice reuse in Freeway. The difference of performance between students with and without advice reuse becomes more evident in Pong. Students with advice reuse (SUA-AIR and AIR) showed a statistically significant boost in performance throughout the majority of training (other than the end of training). Advised agents without advice reuse (EA, SUA, and RA) failed to keep up with the performance of advice reuse students in Pong. In Q\*bert, the advice reuse agents (SUA-AIR and AIR) performed statistically better than all other agents towards the end of training. In Seaquest, most of the advised agents (except RA) took the lead over the NA agent throughout the majority of training (other than the end of training) with statistically significant differences.

In general, SUA-AIR and AIR performed in a similar fashion. The differences in performance between SUA-AIR and AIR are not statistically significant. Thus, the answer to RQ1 is inconclusive. However, SUA-AIR does perform statistically better in various phases of training (e.g. early, later phases) across different domains (e.g. Pong, Q\*bert), as compared to other agents (except AIR). A table reporting the evaluation scores of the student agents across all domains is shown in the archived version of this paper.

The advice taken and the reuse schedule are shown in Figure 4. The number of advice reused (top row) for AIR is higher than SUA and SUA-AIR in almost all steps. However, having a higher advice reuse rate does not necessarily correspond to a better evaluation performance. The amount of advice taken (bottom row) for SUA and SUA-AIR is, in general, spread in an erratic fashion across the environment steps. AIR seeks to consume most of the training budget early on due to the higher uncertainty of the model of the teacher. Whereas the proposed algorithms do a better job of asking the teacher as and when needed. Thus, to answer RQ3, the proposed algorithms do bring a change in the way advice is taken directly from the teacher.

#### **Evaluating Model Performance**

To further investigate the similar evaluation performance of SUA-AIR and AIR, we evaluate the accuracy of the model of teacher by comparing the actions of the teacher and the model for the states that the student visits. This evaluation is shown in Figure 5.



## Figure 5: Percentage of correct actions taken by the model of the teacher for SUA-AIR and AIR across training steps in Pong averaged over 3 independent runs.

The percentage of correct actions taken by the model (y-axis), correct actions are actions that are the same as teacher's, over the course of millions of environment steps (x-axis) taken by the student agent in Pong are reported. It is evident from Figure 5 that the model of the teacher for SUA-AIR and AIR show similar accuracies across training. Since AIR uses the model of the teacher to collect advice, it is natural to see it slightly more accurate than the model of SUA-AIR. However, this does not impact the evaluation performance of SUA-AIR. We expected SUA-AIR to be better than AIR, but it is not, and we intend to understand why in the future work. Due to the similar performance of these models of the teacher, SUA and SUA-AIR showed similar evaluation performance across all games.

### 7 CONCLUSION AND FUTURE WORK

We proposed two new methods, Student's Uncertainty-driven Advising (SUA) and Student's Uncertainty-driven Advising with Advice Imitation & Reuse (SUA-AIR). Both SUA and SUA-AIR use our proposed method for computing student's uncertainty to drive the advice collection process. This uncertainty of the student agent and the uncertainty of the model of teacher were then used for the advice collection (in SUA and SUA-AIR) and advice reuse processes (in SUA-AIR), respectively. The model of teacher was trained using advice interaction data between the student and teacher agents. Using this framework, the student agent can decide when to ask the teacher agent for direct advice, or the model of teacher for advice reuse, or when to follow the student's own policy. We hoped that SUA-AIR would outperform AIR, but we found that the differences were not statistically significant. Moreover, the results show that using advice reuse, in action advising RL agents, provides a significant boost in performance in different stages of training.

There are multiple avenues for future work. Currently, the student agent in SUA-AIR considerably leverages the model of the teacher for advice reuse after the teaching budget is consumed. This can be further extended to add more flexibility where the student agent could start asking the model of the teacher for reuse during the consumption of the teaching budget to use the teaching budget efficiently. Moreover, the decision to reuse advice in SUA-AIR precedes the decision of the student agent. This can be extended to account for the student agent's uncertainty before reusing advice from the model of the teacher. Furthermore, a more thorough study could be conducted to test different advice reuse schedules by modifying the initial reuse probability, final reuse probability, and the total decaying steps. For example, the final reuse probability could be set to 0 to allow the student agent to become independent from the model of the teacher towards the later stages of training. Lastly, we currently use fixed percentile values to compute the uncertainties for the student agent and the model of teacher. This could be better extended to follow a dynamic schedule where the percentile values start closer to 50 in early training and then restricted to values closer to 100 in the later stages of training. This change would be better suited for advice reuse since a lower percentile value for the student agent's uncertainty would not capture the states for which the student is genuinely uncertain.

#### ACKNOWLEDGEMENTS

This work has taken place in the Intelligent Robot Learning Lab at the University of Alberta, which is supported in part by research grants from the AI4Society; the Alberta Machine Intelligence Institute (Amii); a Canada CIFAR AI Chair, Amii; Compute Canada; and NSERC.

#### REFERENCES

- Ofra Amir, Ece Kamar, Andrey Kolobov, and Barbara Grosz. 2016. Interactive Teaching Strategies for Agent Training. In In Proceedings of IJCAI 2016.
- [2] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57, 5 (2009), 469–483.
- [3] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal* of Artificial Intelligence Research 47 (2013), 253–279.
- [4] Georgios Chalkiadakis and Craig Boutilier. 2003. Coordination in Multiagent Reinforcement Learning: A Bayesian Approach. In Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (Melbourne, Australia) (AAMAS '03). Association for Computing Machinery, New York, NY, USA, 709–716. https://doi.org/10.1145/860575.860689
- [5] Lu Chen, Xiang Zhou, Cheng Chang, Runzhe Yang, and Kai Yu. 2017. Agentaware dropout dqn for safe and efficient on-line dialogue policy learning. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. 2454-2464.
- [6] Jeffery Allen Clouse. 1996. On integrating apprentice learning and reinforcement learning. University of Massachusetts Amherst.
- [7] Felipe Leno Da Silva, Ruben Glatt, and Anna Helena Reali Costa. 2017. Simultaneously learning and advising in multiagent reinforcement learning. In Proceedings of the 16th conference on autonomous agents and multiagent systems. 1100–1108.
- [8] Felipe Leno Da Silva, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. 2020. Uncertainty-aware action advising for deep reinforcement learning agents. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34. 5792–5799.
- [9] Jesse Farebrother, Marlos C. Machado, and Michael Bowling. 2018. Generalization and Regularization in DQN. CoRR abs/1810.00123 (2018).
- [10] Sai Krishna Gottipati, Yashaswi Pathak, Boris Sattarov, Sahir, Rohan Nuttall, Mohammad Amini, Matthew E.Taylor, and Sarath Chandar. 2021. Towered Actor Critic For Handling Multiple Action Types In Reinforcement Learning For Drug Discovery. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 35. 142–150.
- [11] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles Isbell, and Andrea L Thomaz. 2013. Policy shaping: Integrating human feedback with reinforcement learning. In NIPS.
- [12] Ercument Ilhan, Jeremy Gow, and Diego Perez Liebana. 2021. Action Advising with Advice Imitation in Deep Reinforcement Learning. In AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021, Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé (Eds.). ACM, 629–637. https://dl.acm.org/doi/10.5555/ 3463952.3464029

- [13] Ercument Ilhan, Jeremy Gow, and Diego Perez Liebana. 2021. Learning on a Budget via Teacher Imitation. CoRR abs/2104.08440 (2021). arXiv:2104.08440 https://arxiv.org/abs/2104.08440
- [14] Ercument Ilhan and Diego Perez Liebana. 2020. Student-Initiated Action Advising via Advice Novelty. CoRR abs/2010.00381 (2020). arXiv:2010.00381 https://arxiv. org/abs/2010.00381
- [15] Ning Liu, Ying Liu, Brent Logan, Zhiyuan Xu, Jian Tang, and Yanzhi Wang. 2019. Learning the dynamic treatment regimes from medical registry data through deep Q-network. *Scientific reports* 9, 1 (2019), 1–10.
- [16] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [17] A Ng, D Harada, and S Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*.
- [18] Phillip Odom and Sriraam Natarajan. 2016. Active Advice Seeking for Inverse Reinforcement Learning. In Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016. 512–520.
- [19] Shayegan Omidshafiei, Dong-Ki Kim, Miao Liu, Gerald Tesauro, Matthew Riemer, Christopher Amato, Murray Campbell, and Jonathan P How. 2019. Learning to teach in cooperative multiagent reinforcement learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33. 6128–6136.
- [20] Dean A Pomerleau. 1991. Efficient training of artificial neural networks for autonomous navigation. Neural computation 3, 1 (1991), 88–97.
- [21] Stefan Schaal. 1999. Is imitation learning the route to humanoid robots? Trends in cognitive sciences (1999).
- [22] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [23] Richard S Sutton and Andrew G Barto. 2018. Reinforcement learning: An introduction. MIT press.
- [24] Matthew E Taylor, Nicholas Carboni, Anestis Fachantidis, Ioannis Vlahavas, and Lisa Torrey. 2014. Reinforcement learning agents providing advice in complex video games. Connection Science 26, 1 (2014), 45-63.
- [25] Lisa Torrey and Matthew Taylor. 2013. Teaching on a budget: Agents advising agents in reinforcement learning. In Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems. 1053-1060.
- [26] Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. 2019. Optimization of molecules via deep reinforcement learning. *Scientific reports* 9, 1 (2019), 1–10.
- [27] Changxi Zhu, Yi Cai, Ho-fung Leung, and Shuyue Hu. 2020. Learning by reusing previous advice in teacher-student paradigm. In Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. 1674–1682.