

# Opportunities and Challenges from Using Animal Videos in Reinforcement Learning for Navigation <sup>★</sup>

Vittorio Giammarino <sup>\*</sup> James Queeney <sup>\*</sup>  
Lucas C. Carstensen <sup>\*\*</sup> Michael E. Hasselmo <sup>\*\*</sup>  
Ioannis Ch. Paschalidis <sup>\*\*\*</sup>

<sup>\*</sup> Division of Systems Engineering, Boston University, Boston, MA 02215, USA, E-mail: {vgiammar, jqueeney}@bu.edu.

<sup>\*\*</sup> Center for Systems Neuroscience, Department of Psychological and Brain Sciences and Graduate Program for Neuroscience, Boston University, Boston, MA 02215 USA, E-mail: {lucas, hasselmo}@bu.edu.

<sup>\*\*\*</sup> College of Engineering and Hariri Institute for Computing & Computational Science & Engineering, Boston University, Boston, MA 02215, USA, E-mail: yannisp@bu.edu.

---

**Abstract:** We investigate the use of animal videos (observations) to improve Reinforcement Learning (RL) efficiency and performance in navigation tasks with sparse rewards. Motivated by theoretical considerations, we make use of weighted policy optimization for off-policy RL and describe the main challenges when learning from animal videos. We propose solutions and test our ideas on a 2D navigation task. We show how the use of animal videos improves performance over RL algorithms that do not leverage such observations.

**Keywords:** Learning for Control; Reinforcement Learning; Deep Learning; RL with demonstrations; Imitation Learning; Bio-inspired Control.

---

## 1. INTRODUCTION

Reinforcement Learning (RL) is a powerful tool to learn control policies from experience where the learning agent does not have explicit information about the task but is only allowed to interact with the environment and observe rewards (Sutton and Barto, 2018). Unfortunately, the trial-and-error nature of RL is often an obstacle for its actual deployment in real-world scenarios (Dulac-Arnold et al., 2019). A solution to this problem is offered by the offline RL framework (Fu et al., 2020), where control policies are learned entirely from offline datasets without further interaction with the environment.

In realistic scenarios, leveraging offline data becomes a much more difficult task for the following reasons. First, datasets are not always collected in a demonstration format, that is, sequences of states, actions, and reward tuples; but in an observation format, i.e., sequences of states only, preventing the use of pure offline algorithms. Even when states and actions are collectable, the presence of domain differences between the source, i.e., the agent from which we collect data, and the target agent (the learner) limits the usefulness of the data. Finally, no information about the data collection policy is available and therefore techniques

for off-policy corrections, such as Importance Sampling (IS) (refer to Queeney et al. 2021), become unusable.

In this paper, we leverage videos of animals navigating an environment as offline datasets. We propose a hybrid approach that uses learner interactions to extract ready-to-use demonstrations from videos and uses the demonstrations to improve RL efficiency and performance. In order to accomplish our goal, we address three main problems throughout the paper: (i) dealing with Domain Adaptation (DA) issues, (ii) dealing with videos as offline datasets, and (iii) theoretically motivating off-policy learning algorithms.

**Contributions:** We start by formulating an adversarial DA algorithm (Tzeng et al., 2017) that maps both the offline dataset and the online interactions of the learning agent into a common space, allowing for the use of data from different domains. Further, we outline a simple method based on supervised learning to recover demonstrations (states, actions, and rewards) from observations (states only). Finally, we propose an IS-free off-policy algorithm called *Advantage Weighted Policy Optimization (AWPO)* and discuss its connection to weighted regression RL algorithms such as Relative Entropy Policy Search (REPS) (Peters et al., 2010), Maximum a posteriori Policy Optimization (MPO) (Abdolmaleki et al., 2018), Advantage Weighted Regression (AWR) (Peng et al., 2019), and Advantage Weighted Actor Critic (AWAC) (Nair et al., 2020) in the light of a recent work on principled off-policy learning (Queeney et al., 2021).

---

<sup>★</sup> Research partially supported by the NSF under grants CCF-2200052, IIS-1914792, and DMS-1664644, by the ONR under grants N00014-19-1-2571 and N00014-21-1-2844, by the NIH under grants R01 GM135930 and UL54 TR004130, and by the Boston University Kilachand Fund for Integrated Life Science and Engineering.

**Notation:** Unless indicated otherwise, we use uppercase letters (e.g.,  $S_t$ ) for random variables, lowercase letters (e.g.,  $s_t$ ) for values of random variables, script letters (e.g.,  $\mathcal{S}$ ) for sets, and bold lowercase letters (e.g.,  $\theta$ ) for vectors. Let  $[t_1 : t_2]$  be the set of integers  $t$  such that  $t_1 \leq t \leq t_2$ ; we write  $S_t$  such that  $t_1 \leq t \leq t_2$  as  $S_{t_1:t_2}$ . We denote with  $\mathbb{E}[\cdot]$  expectation, with  $\mathbb{P}(\cdot)$  probability, and with  $\mathbb{1}_S(\cdot)$  the indicator function which is equal to 1 when  $S = s$  and 0 otherwise.  $\mathcal{N}(\mu, \sigma)$  denotes a normal distribution with mean equal to  $\mu$  and standard deviation equal to  $\sigma$ .  $\mathbb{D}_{\text{KL}}(\cdot || \cdot)$  denotes the Kullback-Leibler (KL) divergence and  $\mathbb{D}_{\text{TV}}(\cdot, \cdot)$  the total variation (TV) distance.

## 2. PRELIMINARIES

We consider an infinite-horizon discounted Markov Decision Process (MDP) defined by the tuple  $(\mathcal{S}, \mathcal{A}, P, r, d_0, \gamma)$  where  $\mathcal{S}$  is the set of states and  $\mathcal{A}$  is the set of actions, both possibly infinite.  $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$  is the transition probability function and  $\Delta_{\mathcal{S}}$  denotes the space of probability distributions over  $\mathcal{S}$ . The function  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  maps state-action pairs to rewards.  $d_0 \in \Delta_{\mathcal{S}}$  is the initial state distribution and  $\gamma \in [0, 1]$  the discount factor. We model the decision agent as a stationary policy  $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$ , where  $\pi(a|s)$  is the probability of taking action  $a$  in state  $s$ . The goal is to choose a policy that maximizes the expected total discounted reward  $J(\pi) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$ , where  $\tau = (s_0, a_0, s_1, a_1, \dots)$  is a trajectory sampled according to  $s_0 \sim d_0$ ,  $a_t \sim \pi(\cdot|s_t)$  and  $s_{t+1} \sim P(\cdot|s_t, a_t)$ . A policy  $\pi$  induces a normalized discounted state visitation distribution  $d^\pi$ , where  $d^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(s_t = s | d_0, \pi, P)$ . We define the corresponding normalized discounted state-action visitation distribution as  $\mu^\pi(s, a) = d^\pi(s) \pi(a|s)$ . We denote the state value function of  $\pi$  as  $V^\pi(s) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | S_0 = s]$ , the state-action value function as  $Q^\pi(s, a) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | S_0 = s, A_0 = a]$  and the advantage function as  $A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s)$ . Finally, when a policy is parameterized with parameters  $\theta \in \Theta \subset \mathbb{R}^k$  we write  $\pi_\theta$ .

## 3. DOMAIN ADAPTATION AND ESTIMATION OF ACTIONS-REWARDS FROM VIDEOS

In this section we formalize the concepts of domain adaptation and actions-rewards estimation (Higgins et al., 2017; Xing et al., 2021; Schmeckpeper et al., 2020); we present ways to address them within the RL framework.

We denote the source and the target domains as  $D_S$  and  $D_T$ , respectively. Each domain corresponds to an MDP:  $D_S \equiv (\mathcal{S}_S, \mathcal{A}_S, P_S, r_S, d_0^S, \gamma)$  and  $D_T \equiv (\mathcal{S}_T, \mathcal{A}_T, P_T, r_T, d_0^T, \gamma)$  where  $\gamma$  is a common discount factor.  $\mathcal{S}$  and  $d_0$  are different between the source and target domains, while  $\mathcal{A}$  is shared and  $P$  and  $r$  have structural similarities. Refer to Fig. 1 for an example. Fig. 1a shows our target  $D_T$  and Fig. 1b and 1c show two possible sources  $D_S$ . The state spaces, made of RGB pixels, show significant differences between  $D_S$  and  $D_T$  ( $\mathcal{S}_S \neq \mathcal{S}_T$ ). On the other hand, all the domains share the same action spaces ( $\mathcal{A}_S = \mathcal{A}_T$ ), since the control action can be modeled using the same motion primitives: north, south, east, west. Finally, transition probabilities and rewards have similarities ( $P_S \approx P_T$  and  $r_S \approx r_T$ ) since all the domains are governed by the same deterministic transition dynamics and the performance depends on the

position of the agent with respect to a point in the environment. In every source domain  $D_S$ , we can record a video of a quasi-optimal agent while performing the task. We refer to this video as the offline dataset  $(s_S, s'_S)_{1:M}$  made only of observations. Our goal is leveraging  $(s_S, s'_S)_{1:M}$  to achieve faster policy learning in  $D_T$ .

Typically, in deep RL from pixels, agents learn an end-to-end mapping from states  $s_T \in \mathcal{S}_T$  to actions  $a_T \in \mathcal{A}_T$  (Mnih et al., 2013, 2015). In the process, a function  $f_T : \mathcal{S}_T \rightarrow \mathcal{H}$  that maps the high-dimensional pixel space into a lower dimensional embedding space  $\mathcal{H}$  is implicitly learnt and a policy  $\pi_\theta : \mathcal{H} \rightarrow \Delta_{\mathcal{A}}$  learns the mapping from  $\mathcal{H}$  to  $\Delta_{\mathcal{A}}$ . In order to leverage  $(s_S, s'_S)_{1:M}$  in  $D_T$ , we define, in addition to  $f_T$ , another encoding function  $f_S : \mathcal{S}_S \rightarrow \mathcal{H}$ . Our domain adaptation step will focus on training  $f_S$  such that the features extracted from  $\mathcal{S}_S$  by  $f_S$  will match as much as possible those extracted from  $\mathcal{S}_T$  by  $f_T$ . When this is accomplished, the policy  $\pi_\theta$  can be directly trained on the embedding space  $\mathcal{H}$  leveraging data collected from both  $D_T$  and  $D_S$ .

Furthermore, given  $(s_S, s'_S)_{1:M}$  and the two encoders  $f_T$  and  $f_S$ , our actions-rewards estimation will consist in training two inverse models  $f_{\chi_a}^{\text{im}} : \mathcal{H} \times \mathcal{H} \rightarrow \Delta_{\mathcal{A}}$  and  $f_{\chi_r}^{\text{im}} : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$  using  $(h_T, a_T, r_T, h'_T) \equiv (f_T(s_T), a_T, r_T, f_T(s'_T))$  in order to estimate actions and rewards for  $(h_S, h'_S)_{1:M} \equiv (f_S(s_S), f_S(s'_S))_{1:M}$ . The full set  $(h_T, a_T, r_T, h'_T) \cup (h_S, \hat{a}_S, \hat{r}_S, h'_S)$  will be then used to train  $\pi_\theta$ .

### 3.1 Domain adaptation

We define two different encoding functions  $f_{\delta_T}^{\text{enc}} : \mathcal{S}_T \rightarrow \mathcal{H}$  and  $f_{\delta_S}^{\text{enc}} : \mathcal{S}_S \rightarrow \mathcal{H}$ . Since we are interested in learning from videos, both  $f_{\delta_T}^{\text{enc}}$  and  $f_{\delta_S}^{\text{enc}}$  are parameterized with Convolutional Neural Networks (CNN) which take RGB images as input and respectively output  $h_T \in \mathcal{H}$  and  $h_S \in \mathcal{H}$ . We train both  $f_{\delta_T}^{\text{enc}}$  and  $f_{\delta_S}^{\text{enc}}$  end-to-end with a value function  $V_{\theta_V} : \mathcal{H} \rightarrow \mathbb{R}$  or a state-action value function  $Q_{\theta_Q} : \mathcal{H} \times \mathcal{A} \rightarrow \mathbb{R}$  (see Algorithm 1). Additionally, we train  $f_{\delta_S}^{\text{enc}}$  using adversarial DA as in Tzeng et al. (2017). Our adversarial DA step is formulated as the following min-max game

$$\min_{\delta_S} \max_{\phi} \frac{1}{m} \sum_{i=1}^m [f_{\phi}^{\text{disc}}(f_{\delta_T}^{\text{enc}}(s_T^{(i)})) - f_{\phi}^{\text{disc}}(f_{\delta_S}^{\text{enc}}(s_S^{(i)}))], \quad (1)$$

where  $\{s_S^{(i)}\}_{i=1}^m \sim (s_S, s'_S)_{1:M}$ ,  $\{s_T^{(i)}\}_{i=1}^m \sim d^{\pi_{\theta_k}}$  and  $f_{\phi}^{\text{disc}} : \mathcal{H} \rightarrow \mathbb{R}$  is a discriminator function which recognizes from which domain ( $D_S$  or  $D_T$ ) the embedding has been generated.  $f_{\delta_S}^{\text{enc}}$  is trained in order to make the features  $h_S$  as close as possible to  $h_T$ . At optimality, the discriminator is no longer able to distinguish from which MDP the embedding  $h \in \mathcal{H}$  has been generated, meaning that both  $\mathcal{S}_S$  and  $\mathcal{S}_T$  have been mapped into the same common space  $\mathcal{H}$ . Note that in (1) we use the Wasserstein loss as introduced in Arjovsky et al. (2017). We prefer this loss to the original binary cross-entropy in Goodfellow et al. (2020) for stability reasons.

### 3.2 Actions and rewards estimation

Recall  $h_S \equiv f_{\delta_S}^{\text{enc}}(s_S)$ ,  $h_T \equiv f_{\delta_T}^{\text{enc}}(s_T)$  and the source data  $(s_S, s'_S)_{1:M}$  which do not contain actions  $(a_S)_{1:M}$

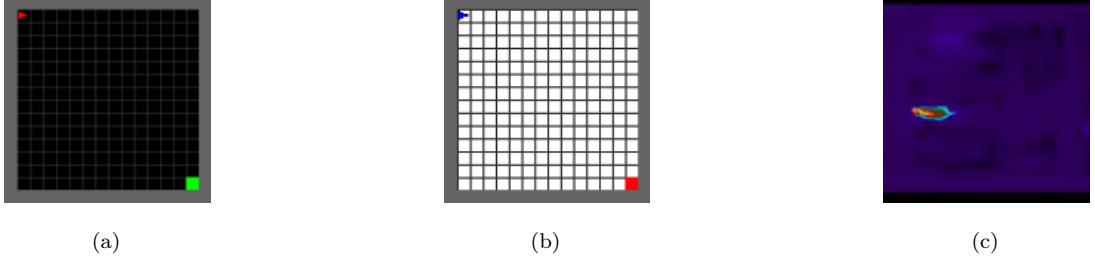


Fig. 1. MDPs used in the RL experiments. For all the MDPs the state space is made of  $128 \times 128$  RGB images showing the top view of the environment. The action spaces are the navigation primitives north, south, east and west and the reward is +1 when the bottom right corner is reached and a 0.001 penalty for each decision step. Fig. 1a shows the target MDP. Fig. 1b and Fig. 1c show two different source MDPs from which we collect videos used as offline datasets in Algorithm 1. A more extensive description of the rodent video in Fig. 1c is provided in Fig. 2. For both the environments in Fig. 1a and Fig. 1b, we rely on the minigrid library (Chevalier-Boisvert et al., 2018).

and rewards  $(r_S)_{1:M}$ . A popular approach in the literature is to estimate  $(a_S)_{1:M}$  and  $(r_S)_{1:M}$  using inverse models, where the online interactions  $(h_T, a_T, r_T, h'_T)$  collected by the learner in  $D_T$  are used for supervised learning (Torabi et al., 2018; Schmeckpeper et al., 2020; Pathak et al., 2017). We define two inverse models  $f_{\chi_a}^{\text{im}} : \mathcal{H} \times \mathcal{H} \rightarrow \Delta_{\mathcal{A}}$  and  $f_{\chi_r}^{\text{im}} : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$  for actions and rewards, respectively. We collect  $(h_T, a_T, r_T, h'_T)$  in  $D_T$  and train the two inverse models to minimize the following losses

$$\begin{aligned} \mathcal{L}_a^{\text{im}}(\chi_a) &= \|p(a_T) - f_{\chi_a}^{\text{im}}(h_T, h'_T)\|^2, \\ \mathcal{L}_r^{\text{im}}(\chi_r) &= (r_T - f_{\chi_r}^{\text{im}}(h_T, h'_T))^2, \end{aligned} \quad (2)$$

where  $p(a_T)$  is a distribution over  $\mathcal{A}$ . In practice, when  $\mathcal{A}$  is a discrete action space  $p(a_T) = \mathbb{1}_{A_T}(a_T)$ , and when  $\mathcal{A}$  is continuous  $p(a_T) = \mathcal{N}(a_T, \sigma)$  with  $\sigma$  arbitrary. We then sample  $\hat{a}_S \sim f_{\chi_a}^{\text{im}}(h_S, h'_S)$  and compute  $\hat{r}_S = f_{\chi_r}^{\text{im}}(h_S, h'_S)$  obtaining the tuples  $(h_S, \hat{a}_S, \hat{r}_S, h'_S)_{1:M}$  which are combined with  $(h_T, a_T, r_T, h'_T)$  for policy learning in  $D_T$ . Finally, because we assume that the demonstrating agent in  $D_S$  is quasi-optimal, we add a small constant bonus  $r_i$  to  $\hat{r}_S$  to encourage the learner to follow the demonstrator and compensate for reward sparsity.

#### 4. ADVANTAGE WEIGHTED POLICY OPTIMIZATION

The final challenge in order to successfully leverage videos of animals in RL is determining the proper algorithm for off-policy learning. Recall  $\pi_{\theta}$  is defined as  $\pi_{\theta} : \mathcal{H} \rightarrow \Delta_{\mathcal{A}}$ ; we therefore consider  $\mathcal{H}$  as our state space of interest. In the following sections we state instrumental results for our derivations.

**Policy improvement lower bound:** Our starting point is the policy improvement lower bound originally developed in Kakade and Langford (2002), refined in Achiam et al. (2017) and further generalized for the off-policy setting in Queeney et al. (2021).

*Theorem 1.* (Queeney et al. 2021). Consider the current policy  $\pi_k$  and the behavioral policy  $\beta$ . For any future policy  $\pi$  we have

$$\begin{aligned} J(\pi) - J(\pi_k) &\geq \frac{1}{1-\gamma} \mathbb{E}_{(h,a) \sim \mu^{\beta}} \left[ \frac{\pi(a|h)}{\beta(a|h)} A^{\pi_k}(h, a) \right] \\ &\quad - \frac{2\gamma C^{\pi, \pi_k}}{(1-\gamma)^2} \mathbb{E}_{h \sim d^{\beta}} [\mathbb{D}_{\text{TV}}(\beta(\cdot|h), \pi(\cdot|h))], \end{aligned} \quad (3)$$

where  $C^{\pi, \pi_k} = \max_{h \in \mathcal{H}} |\mathbb{E}_{a \sim \pi(\cdot|h)} [A^{\pi_k}(h, a)]|$  and  $\mathbb{D}_{\text{TV}}(\beta(\cdot|h), \pi(\cdot|h)) = \frac{1}{2} \int_{a \in \mathcal{A}} |\beta(a|h) - \pi(a|h)| da$  is the total variation distance between the distributions  $\beta(\cdot|h)$  and  $\pi(\cdot|h)$ .

Note that Theorem 1 considers the general formulation in Queeney et al. (2021). By setting  $\beta \equiv \pi_k$  we recover the bound in Achiam et al. (2017).

In Theorem 1 the first term on the right-hand side in (3) is a surrogate objective, while the second can be seen as a penalty term or a soft constraint. Theorem 1 shows that by enforcing any future policy  $\pi$  to stay close to data generating policy  $\pi_k$  or  $\beta$ , a monotonic policy improvement can be guaranteed. As a result, the following surrogate optimization problem for Theorem 1 with  $\beta \equiv \pi_k$  is derived in Schulman et al. (2015a):

$$\begin{aligned} \max_{\pi} \quad & \mathbb{E}_{(h,a) \sim \mu^{\pi_k}} \left[ \frac{\pi(a|h)}{\pi_k(a|h)} A^{\pi_k}(h, a) \right] \\ \text{s.t.} \quad & \mathbb{E}_{h \sim d^{\pi_k}} [\mathbb{D}_{\text{KL}}(\pi_k(\cdot|h) \parallel \pi(\cdot|h))] \leq \epsilon. \end{aligned} \quad (4)$$

Based on (4), the popular on-policy Proximal Policy Optimization (PPO) (Schulman et al., 2017) and its off-policy version Generalized PPO (GePPO) (Queeney et al., 2021) are formalized.

**AWPO derivations:** In the following, we focus on deriving from Theorem 1 an off-policy algorithm that we call Advantage Weighted Policy Optimization (AWPO). This belongs to the family of weighted regression algorithms and has strong similarities with REPS, MPO, AWR and AWAC. The main source of difference lies in the need for an off-policy advantage estimation and in the way the algorithm is derived. We start from the lower bound in Theorem 1 and proceed analogously to Achiam et al. (2017). Note that the total variation distance is a metric (Tsybakov, 1997), hence  $\mathbb{D}_{\text{TV}}(\beta(\cdot|h), \pi(\cdot|h)) = \mathbb{D}_{\text{TV}}(\pi(\cdot|h), \beta(\cdot|h))$ . By applying Pinsker's inequality followed by Jensen's inequality, we see that

$$\begin{aligned} \mathbb{E}_{h \sim d^{\beta}} [\mathbb{D}_{\text{TV}}(\pi(\cdot|h), \beta(\cdot|h))] &\leq \mathbb{E}_{h \sim d^{\beta}} \left[ \sqrt{\frac{1}{2} \mathbb{D}_{\text{KL}}(\pi(\cdot|h) \parallel \beta(\cdot|h))} \right] \\ &\leq \sqrt{\frac{1}{2} \mathbb{E}_{h \sim d^{\beta}} [\mathbb{D}_{\text{KL}}(\pi(\cdot|h) \parallel \beta(\cdot|h))].} \end{aligned}$$

The expression above leads to the surrogate problem

$$\begin{aligned} \max_{\pi} \quad & \mathbb{E}_{h \sim d^\beta} \left[ \mathbb{E}_{a \sim \pi(\cdot|h)} [A^{\pi_k}(h, a)] \right] \\ \text{s.t.} \quad & \mathbb{E}_{h \sim d^\beta} [\mathbb{D}_{\text{KL}}(\pi(\cdot|h) || \beta(\cdot|h))] \leq \epsilon, \end{aligned} \quad (5)$$

where  $\beta$  is any behavioral policy. By converting the KL divergence in (5) as a soft constraint and solving the Lagrangian problem in closed form, we obtain the following optimal solution

$$\pi^*(a|h) = \beta(a|h) \exp \left( \frac{A^{\pi_k}(h, a)}{\lambda} \right) \frac{1}{Z(h)},$$

where  $Z(h)$  is a normalizing factor ensuring  $\pi^*(\cdot|h)$  is a probability distribution and  $\lambda$  is a Lagrange multiplier. Further, we project  $\pi^*$  on the manifold of function approximations parameterized by  $\theta$  (Peng et al., 2019). Given  $\pi_{\theta_k}$  the policy parameterized by  $\theta_k$  at step  $k$ , for  $k+1$  we obtain

$$\theta_{k+1} = \arg \max_{\theta} \mathbb{E}_{(h, a) \sim \mu^\beta} \left[ \exp \left( \frac{A^{\pi_{\theta_k}}(h, a)}{\lambda} \right) \log \pi_{\theta}(a|h) \right]. \quad (6)$$

Note that the Lagrange multiplier  $\lambda$  in (6) is provided as hyperparameter.

*Remark 1.* (Novelty). We show direct connection with the generalized policy improvement in Theorem 1 which is a missing piece in all the aforementioned literature (Peters et al., 2010; Abdolmaleki et al., 2018; Peng et al., 2019; Nair et al., 2020). Moreover, a novel small nuance of our formulation is the need for off-policy policy evaluation since we make use of samples from  $\beta$  to compute  $A^{\pi_{\theta_k}}$  for the current policy  $\pi_{\theta_k}$ .

*Remark 2.* (Off-policy policy evaluation). Off-policy policy evaluation is defined as leveraging data generated by a behavioral policy  $\beta$  to evaluate the performance metric of the evaluation policy  $\pi_{\theta_k}$  (refer to Chandak et al. 2021). Well-known algorithms for off-policy policy evaluation are Retrace in Munos et al. (2016), V-trace in Espeholt et al. (2018) and UnO in Chandak et al. (2021). However, all of these require IS correction, which implies knowledge of  $\beta$ . When learning from videos of animals,  $\beta$  is unknown and cannot be estimated using Imitation Learning as in Kostrikov et al. (2021); Giammarino et al. (2022) given the absence of actions. We consider two main algorithms for off-policy evaluation which do not perform IS correction: Peng Q-Lambda (PQL) (Peng and Williams, 1994; Kozuno et al., 2021) and Generalized Advantage Estimation (GAE) (Schulman et al., 2015b). PQL makes use of an explicit correction mechanism for off-policy policy evaluation, while GAE is originally conceived as an on-policy algorithm which, despite not accounting for off-policy data, works well in practice. We refer to Chandak et al. (2021); Munos et al. (2016); Kozuno et al. (2021) for additional details on policy evaluation with off-policy correction.

**Advantage estimation:** As mentioned in Remark 2, we perform advantage estimation using either GAE or PQL. GAE makes use of a parameterized value function  $V_{\theta_V}$ , which is trained by collecting rollouts of length  $N$  in  $D_T$  and length  $M$  in  $D_S$  and minimizing  $\mathcal{L}(\theta_V, \delta_T, \delta_S)$  in (7)

$$\begin{aligned} \mathcal{L}_T(\theta_V, \delta_T) &= \frac{1}{N} \sum_{t=1}^N \|\hat{V}_T^{(t)} - V_{\theta_V}(f_{\delta_T}^{\text{enc}}(s_T^{(t)}))\|^2, \\ \mathcal{L}_S(\theta_V, \delta_S) &= \frac{1}{M} \sum_{t=1}^M \|\hat{V}_S^{(t)} - V_{\theta_V}(f_{\delta_S}^{\text{enc}}(s_S^{(t)}))\|^2, \\ \mathcal{L}(\theta_V, \delta_T, \delta_S) &= \frac{1}{2} (\mathcal{L}_T(\theta_V, \delta_T) + \mathcal{L}_S(\theta_V, \delta_S)), \end{aligned} \quad (7)$$

where  $\hat{V}_T^{(t)} = \sum_{l=0}^N \gamma^l r_T^{(t+l)}$  and  $\hat{V}_S^{(t)} = \sum_{l=0}^M \gamma^l r_S^{(t+l)}$  are discounted sums of rewards. This approach for estimating the value function is called TD(1) or Monte Carlo (Sutton and Barto, 2018). Note that  $f_{\delta_T}^{\text{enc}}$  and  $f_{\delta_S}^{\text{enc}}$  are also trained to minimize  $\mathcal{L}(\theta_V, \delta_T, \delta_S)$  in (7). On the other hand, PQL makes use of a parameterized state-action value function  $Q_{\theta_Q}$  and follows the update rule in Peng and Williams (1994); Kozuno et al. (2021). As for GAE, we update  $f_{\delta_T}^{\text{enc}}$  and  $f_{\delta_S}^{\text{enc}}$  together with  $Q_{\theta_Q}$ .

The full algorithm is summarized in Algorithm 1, which combines domain adaptation, source actions-rewards estimation and AWPO.

---

**Algorithm 1:** AWPO from Videos with Adversarial Domain Adaptation

---

Input:  $(s_S, s'_S)_{1:M}, \pi_{\theta}, V_{\theta_V}(Q_{\theta_Q}), f_{\delta_S}^{\text{enc}}, f_{\delta_T}^{\text{enc}}, f_{\phi}^{\text{disc}}, f_{\chi_a}^{\text{im}}, f_{\chi_r}^{\text{im}}, r_i, \alpha, \lambda, K, N$ .

for  $k = 0, \dots, K-1$  do

$$\begin{aligned} & \text{Generate } (s_T, a_T, r_T, s'_T)_{1:N} \text{ from interaction of } \pi_{\theta_k} \text{ with the environment.} \\ & \text{Perform adversarial DA as in (1), and} \\ & \hat{a}_S \text{ and } \hat{r}_S \text{ estimation as in (2).} \\ & \text{Set } \hat{r}_S = f_{\chi_r}^{\text{im}}(h_S, h'_S) + r_i. \\ & \text{Set } \mathcal{D} \equiv (h_S, \hat{a}_S, \hat{r}_S, h'_S)_{1:M} \cup (h_T, a_T, r_T, h'_T)_{1:N}. \\ & \text{Perform } A^{\pi_{\theta_k}}(h, a) \text{ estimation.} \\ & \text{Update } V_{\theta_V}(Q_{\theta_Q}), f_{\delta_S}^{\text{enc}}, f_{\delta_T}^{\text{enc}} \text{ using (7) (or Peng and Williams 1994).} \\ & \theta_{k+1} = \theta_k + \alpha \nabla_{\theta} \mathbb{E}_{\mathcal{D}} \left[ \exp \left( \frac{A^{\pi_{\theta_k}}(h, a)}{\lambda} \right) \log \pi_{\theta}(a|h) \right]. \\ & \theta_k \leftarrow \theta_{k+1}. \\ & \theta \leftarrow \theta_K. \\ & \text{return } \pi_{\theta} \end{aligned}$$


---

## 5. EXPERIMENTS

We conduct a series of experiments in order to evaluate all components of our algorithm. We focus on 2D navigation tasks with sparse rewards from visual observations. We test Algorithm 1 in RL with offline observations without DA, with simplified DA (environment in Fig. 1a as target and Fig. 1b as source for both these experiments), and finally with hard DA where the rodent video in Fig. 2 is used as an offline dataset (Fig. 1a for target and Fig. 1c for source). Videos of rats foraging in an open-field arena were collected at  $1024 \times 1024$  spatial resolution and 14-bit thermal resolution per pixel using thermal infrared cameras (FLIR SC8000, FLIR Systems, Inc.). The dataset used is a subset utilizing only the top-down view of the Rodent3D dataset, in which a full description is available at Patel et al. (2022). Finally, we parameterize the encoders using a CNN with 3 layers and the discriminator, the policy network, the value function, and the inverse models using multilayer perceptrons (MLP) with a single hidden layer.

For additional implementation details we refer to the code publicly available at our GitHub repository<sup>1</sup>.

### 5.1 Results

We focus on the 2D navigation task in Fig. 1a, where the state  $s_T \in \mathcal{S}_T$  is the  $128 \times 128$  RGB image showing the full grid and the goal is to reach the bottom right corner labeled by the green square. The agent receives +1 when reaching the goal and receives a 0.001 penalty for each step. The episode starts with the agent in the top-left corner and the maximum episode length is set to 512 steps. We run all the experiments for  $400k$  training steps and evaluate the policy every  $2k$  steps for 10 episodes. This procedure is repeated for 10 different random seeds. Fig. 3 summarizes the final results where we report the reward obtained at the end of the training process averaged over episodes and seeds. In Fig. 3, AWPO+GAE and AWPO+PQL refer to pure RL algorithms where the agent learns only by its own experience. The other labels refer instead to Algorithm 1 where for  $D_S = 1b$  we are leveraging a video collected in Fig. 1b as an offline dataset, while for  $D_S = 1c$  we are leveraging the rodent video in Fig. 2. Fig. 3 shows that both versions of AWPO benefit from the introduction of videos as offline datasets and outperform their RL counterpart for both  $D_S = 1b$  and  $D_S = 1c$ . We notice that our DA step is particularly useful when learning from the rodent (Fig. 3b), while for  $D_S = 1b$  this is not the case and our algorithm works better when the DA step is turned off. For  $D_S = 1b$ , we notice  $f_\phi^{\text{disc}}$  is able to discriminate well between Fig. 1a and Fig. 1b due to the background differences, making the DA problem in (1) much more difficult for  $f_{\delta_S}^{\text{enc}}$ . Improving the DA step in the context of policy learning will be the main focus of future works.

## 6. CONCLUSIONS

This paper analyzes the challenges of using videos in RL as offline datasets in order to increase learning efficiency and performance. The main problems we identify consist in dealing with domain discrepancies between source and target domains, learning from offline observations rather than demonstrations, and leveraging off-policy data for policy updates. We allow for structural differences in  $\mathcal{S}$  between source and target domains, which we account for by applying domain adaptation techniques in order to make the offline dataset useful to the learning agent. The importance of this step is evident in Fig. 3 for the  $D_S = 1c$  experiments. Furthermore, learning from offline observations deals with the problem of estimating actions and rewards from an offline dataset consisting of observations only. Our solution is based on inverse models trained using supervised learning and learner’s interactions. Finally, considering recent works on principled off-policy learning, we re-derive the AWPO algorithm and test it with and without offline videos on a 2D navigation task. Future works will focus on improving each individual step with particular attention to DA which shows the greatest margin for improvement. Moreover, we plan to increase the scope of our experiments, focusing not only on the

<sup>1</sup> <https://github.com/VittorioGiammarino/Opportunities-and-Challenges-of-Using-Animal-Videos-in-Reinforcement-Learning-for-Navigation.git>

navigation setup, but also on other learning tasks such as locomotion and manipulation.

## REFERENCES

Abdolmaleki, A., Springenberg, J.T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. (2018). Maximum a posteriori policy optimisation. *arXiv preprint arXiv:1806.06920*.

Achiam, J., Held, D., Tamar, A., and Abbeel, P. (2017). Constrained policy optimization. In *International Conference on Machine Learning*, 22–31. PMLR.

Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, 214–223. PMLR.

Chandak, Y., Niekum, S., da Silva, B., Learned-Miller, E., Brunskill, E., and Thomas, P.S. (2021). Universal off-policy evaluation. *Advances in Neural Information Processing Systems*, 34.

Chevalier-Boisvert, M., Willems, L., and Pal, S. (2018). Minimalistic gridworld environment for gymnasium. <https://github.com/Farama-Foundation/MiniGrid>.

Dulac-Arnold, G., Mankowitz, D., and Hester, T. (2019). Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*.

Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., et al. (2018). IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, 1407–1416. PMLR.

Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. (2020). D4RL: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.

Giammarino, V., Dunne, M.F., Moore, K.N., Hasselmo, M.E., Stern, C.E., and Paschalidis, I.C. (2022). Learning from humans: Combining imitation and deep reinforcement learning to accomplish human-level performance on a virtual foraging task. *arXiv preprint arXiv:2203.06250*.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.

Higgins, I., Pal, A., Rusu, A., Matthey, L., Burgess, C., Pritzel, A., Botvinick, M., Blundell, C., and Lerchner, A. (2017). DARLA: Improving zero-shot transfer in reinforcement learning. In *International Conference on Machine Learning*, 1480–1490. PMLR.

Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer.

Kostrikov, I., Fergus, R., Tompson, J., and Nachum, O. (2021). Offline reinforcement learning with Fisher divergence critic regularization. In *International Conference on Machine Learning*, 5774–5783. PMLR.

Kozuno, T., Tang, Y., Rowland, M., Munos, R., Kapturowski, S., Dabney, W., Valko, M., and Abel, D. (2021). Revisiting Peng’s Q ( $\lambda$ ) for modern reinforcement learning. In *International Conference on Machine Learning*, 5794–5804. PMLR.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

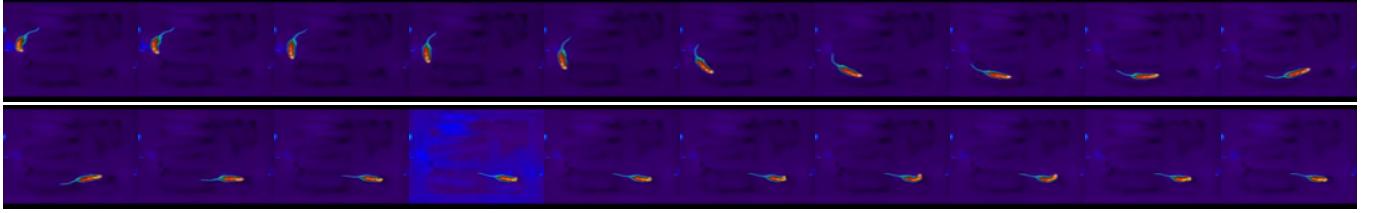
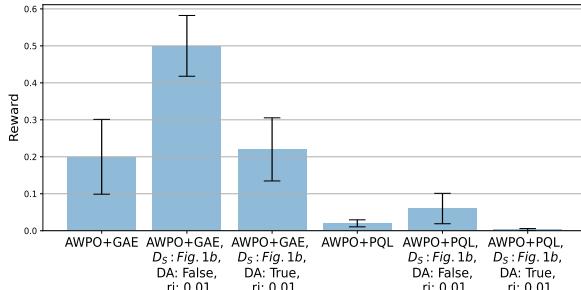
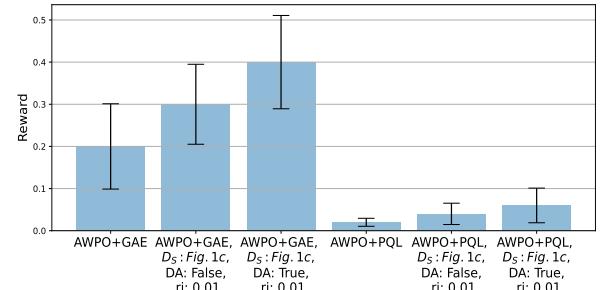


Fig. 2. Frames from the video used as offline dataset. The video consists of a top view of a rodent foraging in an empty open field maze. The video is recorded using a thermal camera.



(a)



(b)

Fig. 3. Experiment results. The reported reward is averaged over the final 10 evaluations and 10 seeds. Error bars show the standard error. Each experiment consists of 400k training steps. Both Fig. 3a and Fig. 3b contain the labels AWPO+GAE and AWPO+PQL which represent the RL-only benchmark where no use of offline datasets is made. The other labels indicate variations of Algorithm 1 where  $D_S$  is the source domain in which the offline dataset is collected (either Fig. 1b or Fig. 1c), DA indicates whether the DA step in (1) is used or not and  $r_i$  is the small bonus added to the estimated reward in the source domain.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidje land, A.K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.

Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. (2016). Safe and efficient off-policy reinforcement learning. *Advances in Neural Information Processing Systems*, 29.

Nair, A., Gupta, A., Dalal, M., and Levine, S. (2020). AWAC: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*.

Patel, M., Gu, Y., Carstensen, L., Hasselmo, M., and Betke, M. (2022). Animal pose tracking: 3d multimodal dataset and token-based pose optimization. *International Journal of Computer Vision*.

Pathak, D., Agrawal, P., Efros, A.A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, 2778–2787. PMLR.

Peng, J. and Williams, R.J. (1994). Incremental multi-step Q-learning. In *Machine Learning Proceedings 1994*, 226–232. Elsevier.

Peng, X.B., Kumar, A., Zhang, G., and Levine, S. (2019). Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*.

Peters, J., Mulling, K., and Altun, Y. (2010). Relative entropy policy search. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

Queeney, J., Paschalidis, I., and Cassandras, C. (2021). Generalized proximal policy optimization with sample reuse. *Advances in Neural Information Processing Systems*, 34.

Schmeckpeper, K., Rybkin, O., Daniilidis, K., Levine, S., and Finn, C. (2020). Reinforcement learning with videos: Combining offline observations with interaction. *arXiv preprint arXiv:2011.06507*.

Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015a). Trust region policy optimization. In *International Conference on Machine Learning*, 1889–1897. PMLR.

Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. (2015b). High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Sutton, R.S. and Barto, A.G. (2018). *Reinforcement learning: An introduction*. MIT press.

Torabi, F., Warnell, G., and Stone, P. (2018). Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*.

Tsybakov, A.B. (1997). On nonparametric estimation of density level sets. *The Annals of Statistics*, 25(3), 948–969.

Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. (2017). Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7167–7176.

Xing, J., Nagata, T., Chen, K., Zou, X., Neftci, E., and Krichmar, J.L. (2021). Domain adaptation in reinforcement learning via latent unified state representation. *arXiv preprint arXiv:2102.05714*.