

---

# Variational Delayed Policy Optimization

---

**Qingyuan Wu** <sup>\*</sup>  
University of Southampton

**Simon Sinong Zhan** <sup>\*</sup>  
Northwestern University

**Yixuan Wang**  
Northwestern University

**Yuhui Wang**  
King Abdullah University of Science and Technology

**Chung-Wei Lin**  
National Taiwan University

**Chen Lv**  
Nanyang Technological University

**Qi Zhu**  
Northwestern University

**Chao Huang** <sup>†</sup>  
University of Southampton

## Abstract

In environments with delayed observation, state augmentation by including actions within the delay window is adopted to retrieve Markovian property to enable reinforcement learning (RL). However, state-of-the-art (SOTA) RL techniques with Temporal-Difference (TD) learning frameworks often suffer from learning inefficiency, due to the significant expansion of the augmented state space with the delay. To improve learning efficiency without sacrificing performance, this work introduces a novel framework called Variational Delayed Policy Optimization (VDPO), which reformulates delayed RL as a variational inference problem. This problem is further modelled as a two-step iterative optimization problem, where the first step is TD learning in the delay-free environment with a small state space, and the second step is behaviour cloning which can be addressed much more efficiently than TD learning. We not only provide a theoretical analysis of VDPO in terms of sample complexity and performance, but also empirically demonstrate that VDPO can achieve consistent performance with SOTA methods, with a significant enhancement of sample efficiency (approximately 50% less amount of samples) in the MuJoCo benchmark. Code is available at <https://github.com/QingyuanWuNothing/VDPO>.

## 1 Introduction

Reinforcement learning (RL) has achieved considerable success across various domains, including board game [32], video game [27], cyber-physical systems [40, 41, 43]. Most of these achievements lack stringent timing constraints, and, therefore, overlook delays in agent-environment interaction. However, delays are prevalent in many real-world applications stemming from various sensors, computation, etc, and significantly affect learning efficiency [17], performance [6], and safety [26]. While observation-delay, action-delay, and reward-delay [10] are all crucial, observation-delay receives the most attention [7, 33, 42]. Unlike reward-delay, observation-delay, which is proved to be a superset of action-delay [19, 29], disrupts the Markovian property inherent to the environments. In this work, we focus on the reinforcement learning with a constant observation-delay  $\Delta$ : at any time step  $t$ , the agent can only observe the state  $s_{t-\Delta}$ , without access to states from  $t-\Delta+1$  to  $t$ .

---

<sup>\*</sup>Equal Contribution

<sup>†</sup>Correspondence to: Chao Huang, [chao.huang@soton.ac.uk](mailto:chao.huang@soton.ac.uk)

Augmentation-based approach is one of the promising methodologies [4, 19]. It retrieves the Markovian property by augmenting the state along with the actions within the window of delays to a new state  $x_t$ , i.e.,  $x_t = \{s_{t-\Delta}, a_{t-\Delta}, \dots, a_{t-1}\}$ , yielding a delayed MDP. However, the underlying sample complexity issue remains a central challenge. Pioneering works [5, 29] directly conduct classical temporal-difference (TD) learning methods, e.g., Deep Q Network [28] and Soft Actor-Critic [13], over the delayed MDP. However, due to the significant growth of the dimensionality, the sample complexity of these techniques increases tremendously. State-of-the-art (SOTA) methods [20, 39, 42] mitigate this issue by introducing an auxiliary delayed task with shorter delays to help learning the original longer delayed task (e.g., improving the long-delayed policy based on the short-delayed value function). However, sample inefficiency is not addressed sufficiently due to the TD learning paradigm still being affected significantly by the increased delays. The memory-less approach [33] improves the learning efficiency by ignoring the absence of the Markovian property of observation-delay RL and learning over the original state space with a cost of serious performance drop. Therefore, the critical challenge still remains: how to **improve learning efficiency without compromising performance** in the delayed setting.

To overcome such a challenge, we propose Variational Delayed Policy Optimization (VDPO), a novel delayed RL framework. Inspired by existing variational RL methods [1, 2, 25], VDPO can utilize extensive optimization tools to resolve the sample complexity issue effectively via formulating the delayed RL problem as a variational inference problem. Specifically, VDPO operates alternatively: (1) learning a reference policy over the delay-free MDP via TD learning and (2) imitating the behaviour of the learned reference policy over the delayed MDP via behaviour cloning. In the high dimensional delayed MDP, VDPO replaces the TD learning paradigm with the behaviour cloning paradigm, which considerably reduces the sample complexity. Furthermore, we demonstrate that VDPO not only effectively improves the sample complexity, but also achieves consistent theoretical performance with SOTAs. Empirical results show that compared to the SOTA approach [42], our VDPO has significant improvement in sample efficiency (*approximately 50% less amount of samples*) along with comparable performance at most MuJoCo benchmarks.

This paper first introduces notations related to delayed RL and variational RL (Sec. 2). In Sec. 3, we present how to formulate the delayed RL problem as a variational inference problem followed by our approach VDPO. Through theoretical analysis, we show that VDPO can effectively reduce the sample complexity without degrading the performance in Sec. 3.2. Practical implementation of VDPO is presented in Sec. 3.3. In Sec. 4, the experimental results over various MuJoCo benchmarks under diverse delay settings validate our theoretical observations. Overall, our contributions are summarized as follows:

- We propose Variational Delayed Policy Optimization (VDPO), a novel framework of delayed RL algorithms emerging from the perspective of variational RL.
- We demonstrate that VDPO enhances sample efficiency, by minimizing the KL divergence between the reference delay-free policy and delayed policy in a behaviour cloning fashion.
- We illustrate that VDPO shares the same theoretical performance as SOTA techniques, by showing that VDPO converges to the same fixed point.
- We empirically show that VDPO not only exhibit superior sample efficiency but also achieves competitive performance comparable to SOTAs across various MuJoCo benchmarks.

## 2 Preliminaries

**MDP.** A delay-free RL problem can be formalized as a Markov Decision Process (MDP), denoted by a tuple  $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \rho \rangle$ , where  $\mathcal{S}, \mathcal{A}$  represent state space and action space respectively,  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  represents the transition function; the reward function is denoted as  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ ;  $\gamma \in (0, 1)$  is the discount factor, and  $p(s_0)$  is the initial state distribution. At each time step  $t$ , the agent takes the action  $a_t \sim \pi(\cdot|s_t)$  based on the current observed state  $s_t$  and the policy  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ , and then observes the next state  $s_{t+1} \sim \mathcal{P}(\cdot|s_t, a_t)$  and a reward signal  $r_t = \mathcal{R}(s_t, a_t)$ . The objective of an RL problem is to find the policy  $\pi$  which can maximize the expected return  $\mathbb{E}_{\tau \sim p_\pi(\tau)} [\mathcal{J}(\tau)] := \mathbb{E}_{\tau \sim p_\pi(\tau)} [\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t)]$  where  $p_\pi(\tau)$  is the trajectory distribution induced by policy  $\pi$ . We use  $d^\pi(s_t)$  to denote the visited state distribution of policy  $\pi$ .

**Delayed MDP.** A delayed RL problem with a constant delay is originally not an MDP, but can be reformulated as a delayed MDP with Markov property based on the augmentation approaches [4, 19]. Assuming the constant delay being  $\Delta$ , the delayed MDP is denoted as a tuple  $\langle \mathcal{X}, \mathcal{A}, \mathcal{P}_\Delta, \mathcal{R}_\Delta, \gamma, \rho_\Delta \rangle$ , where the augmented state space is defined as  $\mathcal{X} := \mathcal{S} \times \mathcal{A}^\Delta$  (e.g., an augmented state  $x_t = \{s_{t-\Delta}, a_{t-\Delta}, \dots, a_{t-1}\} \in \mathcal{X}$ ),  $\mathcal{A}$  is the action space, the delayed transition function is defined as  $\mathcal{P}_\Delta(x_{t+1}|x_t, a_t) := \mathcal{P}(s_{t-\Delta+1}|s_{t-\Delta}, a_{t-\Delta})\delta_{a_t}(a'_t)\prod_{i=1}^{\Delta-1}\delta_{a_{t-i}}(a'_{t-i})$  where  $\delta$  is the Dirac distribution, the delayed reward function is defined as  $\mathcal{R}_\Delta(x_t, a_t) := \mathbb{E}_{s_t \sim b(\cdot|x_t)}[\mathcal{R}(s_t, a_t)]$  where  $b$  is the belief function defined as  $b(s_t|x_t) := \int_{\mathcal{S}^\Delta} \prod_{i=0}^{\Delta-1} \mathcal{P}(s_{t-\Delta+i+1}|s_{t-\Delta+i}, a_{t-\Delta+i}) ds_{t-\Delta+i+1}$ , the initial augmented state distribution is defined as  $\rho_\Delta = \rho \prod_{i=1}^{\Delta} \delta_{a_{-i}}$ .

**Variational RL.** Formulating the RL problem as a probabilistic inference problem [21] allows us to use extensive optimization tools in solving the RL problem. From the existing variational RL literature [30, 36], we usually define  $O = 1$  as the optimality of the task (e.g., the trajectory  $\tau$  obtains the maximum return). Then the probability of trajectory optimality can be represented as  $p(O = 1|\tau)$ . Then, the objective of variational RL becomes finding policy  $\pi$  with highest log evidence:  $\max_\pi \log p_\pi(O = 1)$ . Then, we can derive the lower bound of  $\log p_\pi(O = 1)$  by introducing a prior knowledge of trajectory distribution  $q(\tau)$ .

$$\log p_\pi(O = 1) \geq \mathbb{E}_{\tau \sim q(\tau)}[\log p(O = 1|\tau)] - \text{KL}(q(\tau)||p_\pi(\tau)) = \text{ELBO}(\pi, q), \quad (1)$$

where  $\text{KL}$  is the Kullback-Leibler (KL) divergence and  $\text{ELBO}(\pi, q)$  is the evidence lower bound (ELBO) [2, 30]. The objective of variational RL is maximizing the ELBO, which can be achieved by various optimization techniques [1, 2, 9, 30].

### 3 Our Approach: Variational Delayed Policy Optimization

In this section, we present a new delayed RL approach, Variational Delayed Policy Optimization (VDPO) from the perspective of variational inference. By viewing the delayed RL problem as a variational inference problem, VDPO can utilize extensive optimization tools to address sample complexity and performance issues properly. We first illustrate how to formulate delayed RL as the probabilistic inference problem with an elaborated optimization objective. Subsequently, we theoretically show that the inference problem is equivalent to a two-step iterative optimization problem. Then, we present the framework of VDPO along with the practical implementation.

#### 3.1 Delayed RL as Variational Inference

Delayed RL can be treated as an inference problem: given the desired goal  $O$ , and starting from a prior distribution over trajectory  $\tau$ , the objective is to estimate a posterior distribution over  $\tau$  consistent with  $O$ . The posterior can be formulated by a Boltzman like distribution  $p(O = 1|\tau) \propto \exp\left(\frac{\mathcal{J}(\tau)}{\alpha}\right)$  [2, 31] where  $\alpha$  is the temperature factor. Based on the above definition, the optimization objective of delayed RL can be defined as follows.

$$\max_{\pi_\Delta} \log p_{\pi_\Delta}(O = 1) = \max_{\pi_\Delta} \log \int p(O = 1|\tau)p_{\pi_\Delta}(\tau)d\tau, \quad (2)$$

where  $p_{\pi_\Delta}(O = 1)$  is the probability of the optimality of the delayed policy  $\pi_\Delta$ , and  $p_{\pi_\Delta}(\tau)$  is the trajectory distribution induced by  $\pi_\Delta$ . Based on Eq. (1) and Eq. (2), we can also show that the ELBO for optimization purpose is as follows (derivation of Eq. (3) can be found in Appendix B).

$$\log p_{\pi_\Delta}(O = 1) \geq \underbrace{\mathbb{E}_{\tau \sim p_\pi(\tau)}[\log p(O = 1|\tau)]}_{A \uparrow} - \underbrace{\text{KL}(p_\pi(\tau)||p_{\pi_\Delta}(\tau))}_{B \downarrow} = \text{ELBO}(\pi, \pi_\Delta), \quad (3)$$

where  $p_\pi(\tau)$  is the trajectory distribution induced by a newly-introduced *reference policy*  $\pi$ . As shown in Eq. (3), we transform the original optimization problem as a two-step iterative optimization problem: maximizing term  $A$  while minimizing term  $B$ . Next, we detail how our VDPO optimizes objectives  $A$  and  $B$  separately.

### 3.1.1 Maximizing the performance of reference policy by TD Learning

In this section, we discuss the treatment of term  $A$  in Eq. (3) and investigate the performance and sample complexity of reference policy  $\pi$  under different MDP settings. Maximizing term  $A$  in Eq. (3) is equivalent to maximizing the performance of  $\pi$  as follows.

$$\max_{\pi} \mathbb{E}_{\tau \sim p_{\pi}(\tau)} [\log p(O = 1|\tau)] = \max_{\pi} \mathbb{E}_{\tau \sim p_{\pi}(\tau)} [\mathcal{J}(\tau)]. \quad (4)$$

For Eq. (4), we can train the reference policy  $\pi$  in various MDPs with different delays or even delay-free settings. We show that the performance (Lem. 3.1) and sample complexity (Lem. 3.2) of reference policy  $\pi$  are correlated to the specific MDP setting. Based on existing literature [12, 22] and motivated by existing works [42], VDPO chooses training the reference policy in the delay-free MDP for gaining the edge in terms of **performance** and **sample complexity**.

**Performance:** Lem. 3.1 indicates that the performance of the optimal policy is likely decreased by increasing delays. This motivates us to learn the reference policy in the delay-free MDP for proper performance.

**Lemma 3.1** (Performance in delayed MDP, Theorem 4.3.1 in [22]). *Let  $\mathcal{M}_1, \mathcal{M}_2$  be two constant delayed MDPs with respective delays  $\Delta_1, \Delta_2$  ( $\Delta_1 < \Delta_2$ ). For the optimal policies in  $\mathcal{M}_1, \mathcal{M}_2$ , we have  $\mathcal{J}_1^* \geq \mathcal{J}_2^*$ .*

**Sample Complexity:** Furthermore, for a specific TD-based delayed RL method (e.g., model-based policy iteration), delays also affect its sample efficiency as stated in Lem. 3.2 that stronger delays will lead to much higher sample complexity, resulting in relative learning inefficiency. Therefore, learning the delay-free reference policy makes VDPO superior in sample complexity compared to learning under delay settings.

**Lemma 3.2** (Sample complexity of model-based policy iteration, Theorem 2 in [12]). *Let  $\mathcal{M}$  be the constant delayed MDP with delays  $\Delta$ . Model-based policy iteration finds an  $\epsilon$ -optimal policy with probability  $1 - \sigma$  using sample size  $\mathcal{O}\left(\frac{|\mathcal{X}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma}\right)$ , where  $|\mathcal{X}| = |\mathcal{S}||\mathcal{A}|^\Delta$ .*

Based on the above analysis and inspired by the existing work [23, 42], VDPO adopts a delay-free policy as the reference policy. More rigorous analyses are presented in Sec. 3.2, and we will detail the practical implementation in Sec. 3.3.

### 3.1.2 Minimizing the behaviour difference by Behaviour Cloning

With a fixed reference policy  $\pi$ , minimizing term  $B$  in Eq. (3) can be treated as behaviour cloning at the trajectory level. However, behaviour cloning at the trajectory level is relatively inefficient compared with training at the state level as we have to collect an entire trajectory before training. We next show that we can directly minimize the state-level KL divergence  $\text{KL}(\pi(a_t|s_t)||\pi_{\Delta}(a_t|x_t))$  as presented in Proposition 3.3.

**Proposition 3.3** (State-level KL divergence, proof in Proposition C.1). *For a fixed reference policy  $\pi$ , the trajectory-level KL divergence can be reformulated to state-level KL divergence as follows.*

$$\text{KL}(p_{\pi}(\tau)||p_{\pi_{\Delta}}(\tau)) = \underbrace{\sum_{t=0}^{\infty} \int d^{\pi}(s_t) \text{KL}(\pi(a_t|s_t)||\pi_{\Delta}(a_t|x_t)) ds_t}_{\text{State-level KL divergence}} + \text{Const.}, \quad (5)$$

where  $\text{Const.} = \text{KL}(\rho(s_0)||\rho_{\Delta}(x_0))$

$$+ \sum_{t=0}^{\infty} \int d^{\pi}(s_t) \int \pi(a_t|s_t) \text{KL}(\mathcal{P}(s_{t+1}|s_t, a_t)||b(s_t|x_t)\mathcal{P}_{\Delta}(x_{t+1}|x_t, a_t)) da_t ds_t.$$

Since transition dynamics, initial state distributions, and reference policy are all fixed at this point, we can minimize the state-level KL divergence instead of the trajectory-level KL divergence for efficient training, and then the optimization objective becomes as follows.

$$\min_{\pi_{\Delta}} \text{KL}(p_{\pi}(\tau)||p_{\pi_{\Delta}}(\tau)) \Rightarrow \min_{\pi_{\Delta}} \text{KL}(\pi(a_t|s_t)||\pi_{\Delta}(a_t|x_t)). \quad (6)$$

In this way, VDPO divides the delayed RL problem into two separate optimization problems including Eq. (4) and Eq. (6). How to practically implement VDPO to solve these optimization problems will be presented in Sec. 3.3.

### 3.2 Theoretical Property Analysis

Next, we explain why our VDPO achieves better sample efficiency compared with conventional delayed RL methods, followed by performance analysis of VDPO.

**Sample Complexity Analysis.** In fact, VDPO can use any delay-free RL method to improve the performance of the reference policy (maximizing  $A$ ). Here, we assume that VDPO maximizes  $A$  by the model-based policy iteration, and the sample complexity of maximizing  $A$  is  $\mathcal{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2}\ln\frac{1}{\sigma}\right)$  as described in [Lem. 3.2](#). And minimizing  $B$  in VDPO is equivalent to state-level behaviour cloning which has the sample complexity of  $\mathcal{O}\left(\frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)$  as stated in [Lem. 3.4](#).

**Lemma 3.4** (Sample complexity of behaviour cloning, Theorem 15.3 in [\[3\]](#)). *Given the demonstration from the optimal policy, behaviour cloning finds an  $\epsilon$ -optimal policy with probability  $1 - \sigma$  using sample size  $\mathcal{O}\left(\frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)$ .*

Based on [Lem. 3.2](#) and [Lem. 3.4](#), we can drive the sample complexity of VDPO ([Lem. 3.5](#)).

**Lemma 3.5** (Sample complexity of VDPO, proof in [Lem. C.2](#)). *Assumed that maximizing  $A$  in [Eq. \(3\)](#) by model-based policy iteration while minimizing  $B$  in [Eq. \(3\)](#) by behaviour cloning, VDPO finds an  $\epsilon$ -optimal policy with probability  $1 - \sigma$  using sample size*

$$\mathcal{O}\left(\max\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2}\ln\frac{1}{\sigma}, \frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)\right).$$

Then, based on [Lem. 3.5](#), we show that our VDPO has better sample complexity than most TD-only methods (e.g., model-based policy iteration [\[12\]](#), Soft Actor-Critic [\[5, 20, 42\]](#)) as follows.

**Proposition 3.6** (Sample complexity comparison, proof in [Proposition C.3](#)). *In the delayed MDP, as  $\sigma \rightarrow 0$ , the sample complexity of VDPO ([Lem. 3.5](#)) is less or equal to the sample complexity of model-based policy iteration ([Lem. 3.2](#)):*

$$\mathcal{O}\left(\max\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2}\ln\frac{1}{\sigma}, \frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)\right) \leq \mathcal{O}\left(\frac{|\mathcal{X}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2}\ln\frac{1}{\sigma}\right).$$

[Proposition 3.6](#) tells us that VDPO can reduce the sample complexity effectively, reaching the same performance but requiring fewer samples compared to model-based policy iteration.

**Performance Analysis.** We investigate the convergence of the delayed policy in VDPO ([Lem. 3.7](#)) and show that VDPO can also achieve the same performance as existing SOTAs ([Proposition 3.8](#)). As mentioned above, [Eq. \(4\)](#) in VDPO can be solved by existing delay-free RL method (e.g., model-based policy iteration) to learn an optimal reference policy  $\pi^*$ . Then, we can get the convergence of the delayed policy  $\pi_\Delta^*$  via [Eq. \(6\)](#).

**Lemma 3.7** (Convergence of delayed policy in VDPO, proof in [Lem. C.4](#)). *Let  $\pi^*$  be the optimal reference policy which is trained by a delay-free RL algorithm. The delayed policy  $\pi_\Delta$  converges to  $\pi_\Delta^*$  satisfying that*

$$\pi_\Delta^*(a_t|x_t) = \mathbb{E}_{s_t \sim b(\cdot|x_t)} [\pi^*(a_t|s_t)], \forall x_t \in \mathcal{X}. \quad (7)$$

Based on [Lem. 3.7](#), we show that the convergence of VDPO is consistent with that of existing SOTA methods ([Proposition 3.8](#)).

**Proposition 3.8** (Consistent fixed point, proof in [Proposition C.5](#)). *VDPO shares the same fixed point ([Eq. \(7\)](#)) with DIDA [\[23\]](#), BPQL [\[20\]](#) and AD-SAC [\[42\]](#) for the same delayed MDP.*

[Proposition 3.6](#) and [Proposition 3.8](#) together illustrate that VDPO can effectively improve the sample efficiency while guaranteeing consistent performance with SOTAs [\[20, 23, 42\]](#).

### 3.3 VDPO Implementation

In this section, we detail the implementations of VDPO, specifically the maximization [Eq. \(4\)](#) and the minimization [Eq. \(6\)](#) respectively. The pseudocode of VDPO is summarized in [Alg. 1](#), and the training pipeline of VDPO is presented in [Fig. 1](#).

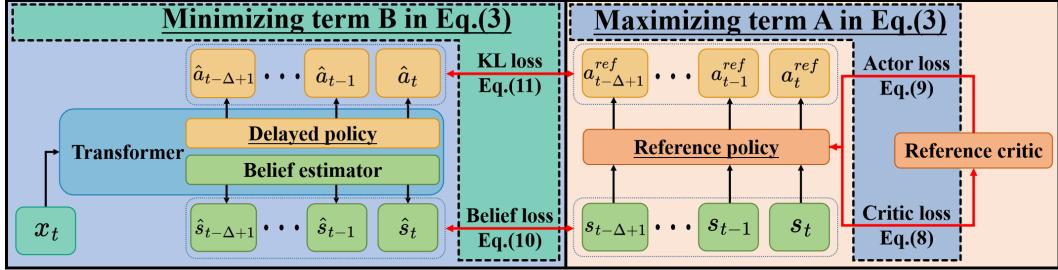


Figure 1: The training pipeline of VDPO.

**Algorithm 1** Variational Delayed Policy Optimization

**Input:** the reference policy  $\pi_\psi$  and critic  $Q_\theta$ ; transformer with belief head  $b_\phi$  and policy head  $\pi_\varphi$ ;  
**for** each update step **do**  
    # A of Eq. (3): Maximizing the performance of the reference policy  $\pi$   
        Updating critic  $Q_\theta$  via Eq. (8) # Soft policy evaluation  
        Updating policy  $\pi_\psi$  via Eq. (9) # Soft policy improvement  
    # B of Eq. (3): Minimizing the state-level KL between  $\pi$  and  $\pi_\Delta$   
        Updating belief head  $b_\phi$  via Eq. (10) # Belief representation learning  
        Updating policy head  $\pi_\varphi$  via Eq. (11) # Behaviour cloning  
**end for**  
**Output:**  $b_\phi$  and  $\pi_\varphi$

**Eq. (4)** aims to maximize the performance of the reference policy  $\pi$  in the delay-free setting, which VDPO addresses using Soft Actor-Critic [13]. Specifically, given transition data  $(s_t, a_t, r_t, s_{t+1})$ , SAC updates the critic  $Q_\theta$  parameterized by  $\theta$  via minimizing the soft TD error:

$$\nabla_\theta \left[ \frac{1}{2} (Q_\theta(s_t, a_t) - \mathbb{Y})^2 \right], \quad (8)$$

where  $\mathbb{Y} = r_t + \gamma \mathbb{E}_{a_{t+1} \sim \pi_\psi(\cdot | s_{t+1})} [Q_\theta(s_{t+1}, a_{t+1}) - \log \pi_\psi(a_{t+1} | s_{t+1})]$  where  $\pi_\psi$  is the reference policy parameterized by  $\psi$ . And the reference policy  $\pi_\psi$  is optimized by the gradient update:

$$\nabla_\psi \mathbb{E}_{\hat{a} \sim \pi_\psi(\cdot | s_t)} [\log \pi_\psi(\hat{a} | s_t) - Q_\theta(s_t, \hat{a})], \quad (9)$$

**Eq. (6)** aims to minimize the state-level KL divergence between the reference policy  $\pi$  and delayed policy  $\pi_\Delta$ . Note that the true state  $s_t$  under the delayed environment is inaccessible. Thus VDPO adopts a two-head transformer [37] to approximate not only the delayed policy  $\pi_\Delta$ , but also the belief estimator  $b$  that predicts the state  $\hat{s}_t$ , as transformer shows a superior representation performance in behaviour cloning [8, 18]. We also discuss how different neural representations influence the RL performance later in Sec. 4.2.3. A similar transformer architecture proposed in [24] is adopted, which serializes the augmented state  $x_t = \{s_{t-\Delta}, a_{t-\Delta}, \dots, a_{t-1}\}$  to  $\{(s_{t-\Delta}, a_{t-i})\}_{i=1}^{\Delta}$  as the input. Based on the information bottleneck principle [34], the encoder needs to encode the input as embeddings with sufficient information related to the true states. Thus, the belief decoder and the policy decoder share a common encoder which is only trained while training the belief decoder, and we freeze the gradient backward of the encoder in training the policy decoder.

Specifically, for a given augmented state  $x_t$  and true states  $\{s_{t-\Delta+i}\}_{i=1}^{\Delta}$ , the belief decoder  $b_\phi$  parameterized by  $\phi$  aims to reconstruct the states  $\{s_{t-\Delta+i}\}_{i=1}^{\Delta}$  based on the  $x_t$ . Therefore, the belief decoder  $b_\phi$  is optimized by the reconstruction loss:

$$\nabla_\phi \sum_{i=1}^{\Delta} \left[ \text{MSE}(b_\phi^{(i)}(x_t), s_{t-\Delta+i}) \right], \quad (10)$$

where  $b_\phi^{(i)}(x_t)$  is the  $i$ -th reconstructed state of the belief decoder  $b_\phi$  and MSE is the mean square error loss. Given the reference policy  $\pi_\psi$  and the pair of augmented state and states  $(x_t, \{s_{t-\Delta+i}\}_{i=1}^{\Delta})$ , the policy decoder  $\pi_\varphi$  parameterized by  $\varphi$  is optimized by minimizing the KL loss:

$$\nabla_\varphi \sum_{i=1}^{\Delta} \left[ \text{KL}(\pi_\varphi^{(i)}(\cdot | x_t) || \pi_\psi(\cdot | s_{t-\Delta+i})) \right], \quad (11)$$

where  $\pi_\varphi^{(i)}(\cdot | x_t)$  is the  $i$ -th output of the policy decoder  $\pi_\varphi$ .

## 4 Experimental Results

### 4.1 Experiment Settings

We evaluate our VDPO in the MuJoCo benchmark [35]. For the selection of baselines, we choose the existing SOTAs including Augmented SAC (A-SAC) [13], DC/AC [5], DIDA [23], BPQL [20] and AD-SAC [42]. The setting of hyper-parameters is presented in [Appendix A](#). We investigate the sample efficiency ([Sec. 4.2.1](#)) followed by performance comparison under different settings of delays ([Sec. 4.2.2](#)). We also conduct the ablation study on the representation of VDPO ([Sec. 4.2.3](#)). Each method was run over 10 random seeds. The training curves can be found in the [Appendix E](#).

### 4.2 Experimental Results

#### 4.2.1 Sample Efficiency

We first evaluate the sample efficiency in the MuJoCo with 5 constant delays. Using the performance of a delay-free policy trained by SAC,  $Ret_{df}$ , as the threshold, we report the required steps to reach this threshold within 1M global steps in [Table 1](#). From the results, we can tell that VDPO shows strong superiority in terms of sample efficiency, successfully reaching the threshold in all tasks and achieving the best sample efficiency in 7 out of 9 tasks. Specifically, VDPO only requires 0.42M and 0.67M steps to reach the threshold in *Ant-v4* and *Humanoid-v4* respectively, while none of the baselines can reach the threshold within 1M steps. In *Halfcheetah-v4*, *Hopper-v4*, *Pusher-v4*, *Swimmer-v4* and *Walker2d-v4*, the steps taken by our VDPO is around 51% (ranging from 25% to 78%) of that required by AD-SAC, SOTA baseline. Based on these results, we can conclude that VDPO shows a significant advantage in sample complexity compared to other baselines. Additional experimental results of 25 and 50 constant delays are presented in [Appendix D](#).

Table 1: Amount of steps required to reach the threshold  $Ret_{df}$  in MuJoCo tasks with 5 constant delays within 1M global steps, where  $\times$  denotes that failed to hit the threshold within 1M global steps. The best result is in blue.

Task (Delays=5)	A-SAC	DC/AC	DIDA	BPQL	AD-SAC	VDPO (ours)
Ant-v4	$\times$	$\times$	$\times$	$\times$	$\times$	<b>0.42M</b>
HalfCheetah-v4	$\times$	$\times$	$\times$	0.99M	0.56M	<b>0.44M</b>
Hopper-v4	0.83M	0.35M	$\times$	0.29M	0.12M	<b>0.07M</b>
Humanoid-v4	$\times$	$\times$	$\times$	$\times$	$\times$	<b>0.67M</b>
HumanoidStandup-v4	0.64M	0.35M	0.10M	<b>0.09M</b>	0.14M	0.14M
Pusher-v4	0.17M	0.02M	0.10M	0.27M	0.04M	<b>0.01M</b>
Reacher-v4	$\times$	0.61M	<b>0.10M</b>	0.90M	0.44M	0.77M
Swimmer-v4	$\times$	0.94M	0.10M	$\times$	0.13M	<b>0.07M</b>
Walker2d-v4	$\times$	$\times$	$\times$	0.52M	0.67M	<b>0.25M</b>

#### 4.2.2 Performance Comparison

The performance of VDPO and baselines are evaluated on MuJoCo with various settings and a normalized indicator [39, 42]  $Ret_{nor} = \frac{Ret_{alg} - Ret_{rand}}{Ret_{df} - Ret_{rand}}$ , where  $Ret_{alg}$  and  $Ret_{rand}$  are the performance of the algorithm and random policy, respectively. The results of MuJoCo benchmarks with 5, 25, and 50 constant delays are shown in the [Table 2](#), showing that VDPO and AD-SAC outperform other baselines significantly in most tasks. Overall, VDPO and AD-SAC (SOTA) show a comparable performance, which is consistent with the theoretical observation in [Sec. 3.2](#).

#### 4.2.3 Additional Discussions

In this section, we conduct the ablation study to investigate the performance of VDPO using different neural representations. Furthermore, we explore whether VDPO is robust under stochastic delays.

**Ablation Study on Representations.** As mentioned in [Sec. 3.3](#), we investigate how the choice of neural representations for belief and policy influences the performance of VDPO. Baselines include multiple-layer perceptron (MLP) and Transformer without belief decoder. The results presented in [Table 3](#) show that the two-head transformer used by our approach yields the best performance

Table 2: Normalized Performance  $Ret_{nor}$  in MuJoCo tasks with 5, 25, and 50 constant delays for 1M global steps, where  $\pm$  denotes the standard deviation. The best performance is in blue.

Task	Delays	A-SAC	DC/AC	DIDA	BPQL	AD-SAC	VDPO (ours)
Ant-v4	5	0.18 $\pm$ 0.01	0.25 $\pm$ 0.05	0.89 $\pm$ 0.03	0.96 $\pm$ 0.03	0.72 $\pm$ 0.25	<b>1.11<math>\pm</math>0.04</b>
	25	0.07 $\pm$ 0.07	0.19 $\pm$ 0.02	0.29 $\pm$ 0.07	0.57 $\pm$ 0.11	<b>0.66<math>\pm</math>0.04</b>	0.56 $\pm$ 0.06
	50	0.02 $\pm$ 0.04	0.19 $\pm$ 0.02	0.19 $\pm$ 0.05	0.38 $\pm$ 0.07	<b>0.48<math>\pm</math>0.06</b>	0.46 $\pm$ 0.07
HalfCheetah-v4	5	0.35 $\pm$ 0.15	0.40 $\pm$ 0.23	0.90 $\pm$ 0.01	1.00 $\pm$ 0.06	<b>1.07<math>\pm</math>0.06</b>	1.03 $\pm$ 0.08
	25	0.04 $\pm$ 0.01	0.16 $\pm$ 0.07	0.12 $\pm$ 0.03	<b>0.87<math>\pm</math>0.04</b>	0.71 $\pm$ 0.12	0.70 $\pm$ 0.17
	50	0.12 $\pm$ 0.17	0.12 $\pm$ 0.13	0.15 $\pm$ 0.03	0.73 $\pm$ 0.17	<b>0.74<math>\pm</math>0.10</b>	0.72 $\pm$ 0.21
Hopper-v4	5	1.02 $\pm$ 0.28	1.16 $\pm$ 0.25	0.40 $\pm$ 0.40	<b>1.25<math>\pm</math>0.09</b>	1.07 $\pm$ 0.30	1.22 $\pm$ 0.08
	25	0.13 $\pm$ 0.04	0.19 $\pm$ 0.04	0.27 $\pm$ 0.08	<b>1.21<math>\pm</math>0.18</b>	0.86 $\pm$ 0.25	0.82 $\pm$ 0.40
	50	0.04 $\pm$ 0.01	0.04 $\pm$ 0.01	0.09 $\pm$ 0.01	0.71 $\pm$ 0.13	<b>0.72<math>\pm</math>0.03</b>	0.22 $\pm$ 0.04
Humanoid-v4	5	0.13 $\pm$ 0.02	0.59 $\pm$ 0.17	0.08 $\pm$ 0.04	0.96 $\pm$ 0.05	0.98 $\pm$ 0.07	<b>1.15<math>\pm</math>0.07</b>
	25	0.05 $\pm$ 0.01	0.04 $\pm$ 0.01	0.07 $\pm$ 0.00	0.12 $\pm$ 0.01	<b>0.25<math>\pm</math>0.16</b>	0.12 $\pm$ 0.02
	50	0.04 $\pm$ 0.01	0.03 $\pm$ 0.01	0.07 $\pm$ 0.00	0.08 $\pm$ 0.01	0.10 $\pm$ 0.01	<b>0.12<math>\pm</math>0.00</b>
HumanoidStandup-v4	5	1.02 $\pm$ 0.08	1.16 $\pm$ 0.12	1.00 $\pm$ 0.00	1.13 $\pm$ 0.07	1.22 $\pm$ 0.03	<b>1.29<math>\pm</math>0.02</b>
	25	0.97 $\pm$ 0.09	1.03 $\pm$ 0.03	0.97 $\pm$ 0.02	1.09 $\pm$ 0.05	<b>1.15<math>\pm</math>0.08</b>	1.13 $\pm$ 0.12
	50	0.90 $\pm$ 0.02	1.02 $\pm$ 0.07	0.89 $\pm$ 0.06	1.06 $\pm$ 0.04	<b>1.12<math>\pm</math>0.02</b>	1.04 $\pm$ 0.16
Pusher-v4	5	1.11 $\pm$ 0.02	1.29 $\pm$ 0.05	1.01 $\pm$ 0.01	1.06 $\pm$ 0.08	<b>1.36<math>\pm</math>0.01</b>	1.17 $\pm$ 0.06
	25	0.49 $\pm$ 0.32	1.12 $\pm$ 0.02	1.04 $\pm$ 0.01	1.07 $\pm$ 0.06	1.29 $\pm$ 0.03	<b>1.31<math>\pm</math>0.07</b>
	50	0.00 $\pm$ 0.05	1.13 $\pm$ 0.01	1.04 $\pm$ 0.02	1.09 $\pm$ 0.05	1.23 $\pm$ 0.02	<b>1.33<math>\pm</math>0.05</b>
Reacher-v4	5	0.97 $\pm$ 0.01	1.02 $\pm$ 0.00	<b>1.03<math>\pm</math>0.00</b>	1.00 $\pm$ 0.01	1.03 $\pm$ 0.01	1.02 $\pm$ 0.03
	25	0.96 $\pm$ 0.02	1.00 $\pm$ 0.00	0.98 $\pm$ 0.01	0.87 $\pm$ 0.05	0.98 $\pm$ 0.02	<b>1.02<math>\pm</math>0.03</b>
	50	0.86 $\pm$ 0.02	0.89 $\pm$ 0.01	0.93 $\pm$ 0.02	0.90 $\pm$ 0.02	0.91 $\pm$ 0.03	<b>1.02<math>\pm</math>0.03</b>
Swimmer-v4	5	0.88 $\pm$ 0.09	1.11 $\pm$ 0.30	1.05 $\pm$ 0.01	0.97 $\pm$ 0.02	1.82 $\pm$ 0.78	<b>2.30<math>\pm</math>0.36</b>
	25	0.72 $\pm$ 0.02	0.78 $\pm$ 0.12	0.93 $\pm$ 0.09	1.36 $\pm$ 0.56	<b>2.52<math>\pm</math>0.40</b>	2.35 $\pm$ 0.27
	50	0.69 $\pm$ 0.04	0.68 $\pm$ 0.06	0.87 $\pm$ 0.03	2.23 $\pm$ 0.55	<b>2.71<math>\pm</math>0.14</b>	2.42 $\pm$ 0.22
Walker2d-v4	5	0.76 $\pm$ 0.21	0.85 $\pm$ 0.12	0.61 $\pm$ 0.07	1.20 $\pm$ 0.11	1.12 $\pm$ 0.09	<b>1.27<math>\pm</math>0.04</b>
	25	0.12 $\pm$ 0.02	0.26 $\pm$ 0.08	0.10 $\pm$ 0.02	0.59 $\pm$ 0.30	<b>0.72<math>\pm</math>0.11</b>	0.27 $\pm$ 0.11
	50	0.11 $\pm$ 0.02	0.11 $\pm$ 0.02	0.08 $\pm$ 0.01	<b>0.23<math>\pm</math>0.10</b>	0.23 $\pm$ 0.11	0.11 $\pm$ 0.03

compared to other candidates. The results also confirm that an explicit belief estimator implemented by a belief decoder can effectively improve performance.

Table 3: Normalized Performance  $Ret_{nor}$  of VDPO using different representations in MuJoCo tasks with 5 constant delays, where  $\pm$  denotes the standard deviation. The best result is in blue.

Tasks (Delays=5)	MLP	Transformer w/o belief	Transformer w/ belief (ours)
Ant-v4	0.86 $\pm$ 0.20	1.09 $\pm$ 0.05	<b>1.11<math>\pm</math>0.04</b>
HalfCheetah-v4	0.95 $\pm$ 0.08	<b>1.37<math>\pm</math>0.11</b>	1.03 $\pm$ 0.08
Hopper-v4	1.11 $\pm$ 0.19	1.13 $\pm$ 0.29	<b>1.22<math>\pm</math>0.08</b>
Humanoid-v4	0.78 $\pm$ 0.11	0.89 $\pm$ 0.43	<b>1.15<math>\pm</math>0.07</b>
HumanoidStandup-v4	1.28 $\pm$ 0.05	1.28 $\pm$ 0.08	<b>1.29<math>\pm</math>0.02</b>
Pusher-v4	<b>1.35<math>\pm</math>0.04</b>	1.34 $\pm$ 0.05	1.17 $\pm$ 0.06
Reacher-v4	1.02 $\pm$ 0.05	1.02 $\pm$ 0.04	<b>1.02<math>\pm</math>0.03</b>
Swimmer-v4	2.29 $\pm$ 0.37	2.11 $\pm$ 0.08	<b>2.30<math>\pm</math>0.36</b>
Walker2d-v4	1.13 $\pm$ 0.20	1.15 $\pm$ 0.14	<b>1.27<math>\pm</math>0.04</b>

**Stochastic Delays.** We compare the performance in the MuJoCo with 5 stochastic delays where  $\Delta = 5$  is with a probability of 0.9 and  $\Delta \in [1, 5]$  is with a probability of 0.1. The results in the Table 4 demonstrate that VDPO outperforms other baselines at most tasks under stochastic delays. Especially in the *Ant-v4* and *Walker2d-v4*, VDPO performs approximately 62% and 49% better than the second best approach, respectively. In the *Reacher-v4* and *Swimmer-v4*, VDPO achieves a comparative performance with the best baseline. We will conduct a theoretical analysis of VDPO under stochastic delays in the future.

**Limitations and Future Works.** We mainly consider deterministic benchmarks in this paper, which are commonly adopted in the SOTAs [5, 23, 39]. However, the recent work ADRL [42] illustrates that existing approaches may have performance degeneration in stochastic environments, which can be mitigated by learning an auxiliary delayed task concomitantly. We will investigate in the future to integrate VDPO with ADRL to address stochastic applications.

Table 4: Normalized Performance  $Ret_{nor}$  in MuJoCo tasks with 5 stochastic delays for 1M global steps, where  $\pm$  denotes the standard deviation. The best result is in blue.

Tasks	A-SAC	DC/AC	DIDA	BPQL	AD-SAC	VDPO (ours)
Ant-v4	0.18 $\pm$ 0.01	0.27 $\pm$ 0.02	0.55 $\pm$ 0.08	0.58 $\pm$ 0.12	0.69 $\pm$ 0.17	1.12 $\pm$ 0.04
HalfCheetah-v4	0.36 $\pm$ 0.12	0.36 $\pm$ 0.18	0.75 $\pm$ 0.02	0.76 $\pm$ 0.16	1.03 $\pm$ 0.06	1.07 $\pm$ 0.09
Hopper-v4	0.85 $\pm$ 0.22	0.94 $\pm$ 0.29	0.31 $\pm$ 0.08	0.68 $\pm$ 0.34	1.05 $\pm$ 0.22	1.35 $\pm$ 0.11
Humanoid-v4	0.15 $\pm$ 0.06	0.67 $\pm$ 0.18	0.07 $\pm$ 0.01	0.40 $\pm$ 0.42	0.97 $\pm$ 0.07	1.06 $\pm$ 0.00
HumanoidStandup-v4	1.03 $\pm$ 0.05	1.20 $\pm$ 0.08	1.00 $\pm$ 0.00	1.10 $\pm$ 0.07	1.26 $\pm$ 0.07	1.27 $\pm$ 0.01
Pusher-v4	1.11 $\pm$ 0.02	1.17 $\pm$ 0.02	1.02 $\pm$ 0.01	1.07 $\pm$ 0.05	1.22 $\pm$ 0.01	1.34 $\pm$ 0.05
Reacher-v4	0.98 $\pm$ 0.01	1.02 $\pm$ 0.01	1.02 $\pm$ 0.00	0.85 $\pm$ 0.11	1.05 $\pm$ 0.01	1.01 $\pm$ 0.04
Swimmer-v4	0.82 $\pm$ 0.10	1.47 $\pm$ 0.58	1.03 $\pm$ 0.02	1.53 $\pm$ 0.52	2.36 $\pm$ 0.64	2.13 $\pm$ 0.18
Walker2d-v4	0.68 $\pm$ 0.28	0.89 $\pm$ 0.08	0.54 $\pm$ 0.09	0.59 $\pm$ 0.30	0.63 $\pm$ 0.39	1.33 $\pm$ 0.11

## 5 Related Works

Compared to the common delay-free setting, delayed RL with disrupted Markovian property [4, 19] is closer to real-world complex applications, such as robotics [17, 26], transportation systems [6] and financial market trading [14]. Existing delayed RL techniques conduct learning over either original state space (referred to as direct approach) or augmented state space (referred to as augmentation-based approach). Direct approaches enjoy high learning efficiency by learning in the original small state space. However, early approaches simply ignore the absence of Markovian property caused by delay and directly conduct classical RL techniques based on delayed observations, which distinctly suffer from serious performance drops. The subsequent improvement is to train based on unobserved instant observations, which are predicted by various generative models, e.g., deterministic generative models [38], Gaussian distributions [7], and transformers [24]. However, the inherent approximation errors in these learned models introduce prediction inaccuracy and result in sub-optimal performance issues [24]. To summarize, direct approaches achieve high learning efficiency, but commonly with a cost of performance degeneration.

The augmentation-based approach is notably more promising as it retrieves Markovian property via augmenting the state with the actions related to delays and thus legitimately enables RL techniques over the yielded delayed MDP [4, 19]. However, the augmentation-based approach works in a significantly larger state space, which is thus plagued by the curse of dimensionality, resulting in learning inefficiency. To mitigate this issue, DC/AC [5] leverages the multi-step off-policy technique to develop a partial trajectory resampling operator to accelerate the learning process. Based on the dataset aggregation technique, DIDA [23] generalizes the pre-trained delay-free policy into an augmented policy. Recent attempts [20, 39] evaluate the augmented policy by a non-augmented Q-function for improving learning efficiency. ADRL [42] suggests introducing an auxiliary delayed task with changeable auxiliary delays for the trade-off between the learning efficiency and performance degeneration in the stochastic MDP. However, these approaches still suffer from the sample complexity issue due to the fundamental challenge of TD learning in high dimensional state space.

The conceptualization of RL as an inference problem has gained attraction recently, allowing the adaption of various optimization tricks to enhance RL efficiency [11, 15, 21, 31]. For instance, VIP [30] integrates different projection techniques into the policy search approach based on variational inference. Virel [9] introduces a variational inference framework that reduces the actor-critic method to the Expectation-Maximization (EM) algorithm. MPO [1, 2] is a family of off-policy entropy-regularized methods in the EM fashion. CVPO [25] extends MPO to the safety-critical settings. The novel trial in this work of viewing the delayed RL as a variational inference problem allows us to use extensive optimization tools to address the sample complexity issue in delayed RL.

## 6 Conclusion

This work explores the challenges of RL problems in environments with inherent delays between agent interactions and their consequences. Existing delayed RL methods often suffer from learning inefficiency as temporal-difference learning in the delayed MDP with high dimensional augmented state space demands an increased sample size. To address this limitation, we present VDPO, a new delayed RL approach rooted in the variational inference principle. VDPO redefines the delayed RL problem into a two-step iterative optimization problem. It alternates between (1) maximizing the

performance of the reference policy by temporal-difference learning in the delay-free setting and (2) minimizing the KL divergence between the reference and delayed policies by behaviour cloning. Furthermore, our theoretical analysis and the empirical results in the MuJoCo benchmark validate that VDPO not only effectively improves the sample efficiency but also maintains a robust performance level.

### **Acknowledgments and Disclosure of Funding**

We sincerely acknowledge the support by the grant EP/Y002644/1 under the EPSRC ECR International Collaboration Grants program, funded by the International Science Partnerships Fund (ISPF) and the UK Research and Innovation, and the grant 2324936 by the US National Science Foundation. This work is also supported by Taiwan NSTC under Grant Number NSTC-112-2221-E-002-168-MY3.

## References

- [1] A. Abdolmaleki, J. T. Springenberg, J. Degrave, S. Bohez, Y. Tassa, D. Belov, N. Heess, and M. Riedmiller. Relative entropy regularized policy iteration. *arXiv preprint arXiv:1812.02256*, 2018.
- [2] A. Abdolmaleki, J. T. Springenberg, Y. Tassa, R. Munos, N. Heess, and M. Riedmiller. Maximum a posteriori policy optimisation. *arXiv preprint arXiv:1806.06920*, 2018.
- [3] A. Agarwal, N. Jiang, and S. M. Kakade. Reinforcement learning: Theory and algorithms. 2019.
- [4] E. Altman and P. Nain. Closed-loop control with delayed information. *ACM sigmetrics performance evaluation review*, 20(1):193–204, 1992.
- [5] Y. Bouteiller, S. Ramstedt, G. Beltrame, C. Pal, and J. Binas. Reinforcement learning with random delays. In *International conference on learning representations*, 2020.
- [6] Z. Cao, H. Guo, W. Song, K. Gao, Z. Chen, L. Zhang, and X. Zhang. Using reinforcement learning to minimize the probability of delay occurrence in transportation. *IEEE transactions on vehicular technology*, 69(3):2424–2436, 2020.
- [7] B. Chen, M. Xu, L. Li, and D. Zhao. Delay-aware model-based reinforcement learning for continuous control. *Neurocomputing*, 450:119–128, 2021.
- [8] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [9] M. Fellows, A. Mahajan, T. G. Rudner, and S. Whiteson. Virel: A variational inference framework for reinforcement learning. *Advances in neural information processing systems*, 32, 2019.
- [10] V. Firoiu, T. Ju, and J. Tenenbaum. At human speed: Deep reinforcement learning with action delay. *arXiv preprint arXiv:1810.07286*, 2018.
- [11] J. Fu, K. Luo, and S. Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [12] M. Gheshlaghi Azar, R. Munos, and H. J. Kappen. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91:325–349, 2013.
- [13] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- [14] J. Hasbrouck and G. Saar. Low-latency trading. *Journal of Financial Markets*, 16(4):646–679, 2013.
- [15] J. Ho and S. Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- [16] S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. G. Araújo. Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
- [17] J. Hwangbo, I. Sa, R. Siegwart, and M. Hutter. Control of a quadrotor with reinforcement learning. *IEEE Robotics and Automation Letters*, 2(4):2096–2103, 2017.
- [18] M. Janner, Q. Li, and S. Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- [19] K. V. Katsikopoulos and S. E. Engelbrecht. Markov decision processes with delays and asynchronous cost collection. *IEEE transactions on automatic control*, 48(4):568–574, 2003.
- [20] J. Kim, H. Kim, J. Kang, J. Baek, and S. Han. Belief projection-based reinforcement learning for environments with delayed feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [21] S. Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.

[22] P. Liotet. Delays in reinforcement learning. *arXiv preprint arXiv:2309.11096*, 2023.

[23] P. Liotet, D. Maran, L. Bisi, and M. Restelli. Delayed reinforcement learning by imitation. In *International Conference on Machine Learning*, pages 13528–13556. PMLR, 2022.

[24] P. Liotet, E. Venneri, and M. Restelli. Learning a belief representation for delayed reinforcement learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.

[25] Z. Liu, Z. Cen, V. Isenbaev, W. Liu, S. Wu, B. Li, and D. Zhao. Constrained variational policy optimization for safe reinforcement learning. In *International Conference on Machine Learning*, pages 13644–13668. PMLR, 2022.

[26] A. R. Mahmood, D. Korenkevych, B. J. Komer, and J. Bergstra. Setting up a reinforcement learning task with a real-world robot. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4635–4640. IEEE, 2018.

[27] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

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

[29] S. Nath, M. Baranwal, and H. Khadilkar. Revisiting state augmentation methods for reinforcement learning with stochastic delays. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1346–1355, 2021.

[30] G. Neumann. Variational inference for policy search in changing situations. 2011.

[31] D. Ramachandran and E. Amir. Bayesian inverse reinforcement learning. In *IJCAI*, volume 7, pages 2586–2591, 2007.

[32] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

[33] E. Schuitema, L. Bušoni, R. Babuška, and P. Jonker. Control delay in reinforcement learning for real-time dynamic systems: A memoryless approach. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3226–3231. IEEE, 2010.

[34] N. Tishby and N. Zaslavsky. Deep learning and the information bottleneck principle. In *2015 ieee information theory workshop (itw)*, pages 1–5. IEEE, 2015.

[35] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.

[36] M. Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056, 2009.

[37] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[38] T. J. Walsh, A. Nouri, L. Li, and M. L. Littman. Learning and planning in environments with delayed feedback. *Autonomous Agents and Multi-Agent Systems*, 18:83–105, 2009.

[39] W. Wang, D. Han, X. Luo, and D. Li. Addressing signal delay in deep reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2023.

[40] Y. Wang, S. Zhan, Z. Wang, C. Huang, Z. Wang, Z. Yang, and Q. Zhu. Joint differentiable optimization and verification for certified reinforcement learning. In *Proceedings of the ACM/IEEE 14th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2023)*, pages 132–141, 2023.

[41] Y. Wang, S. S. Zhan, R. Jiao, Z. Wang, W. Jin, Z. Yang, Z. Wang, C. Huang, and Q. Zhu. Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments. In *International Conference on Machine Learning*, pages 36593–36604. PMLR, 2023.

- [42] Q. Wu, S. S. Zhan, Y. Wang, C.-W. Lin, C. Lv, Q. Zhu, and C. Huang. Boosting long-delayed reinforcement learning with auxiliary short-delayed task. *arXiv preprint arXiv:2402.03141*, 2024.
- [43] S. S. Zhan, Y. Wang, Q. Wu, R. Jiao, C. Huang, and Q. Zhu. State-wise safe reinforcement learning with pixel observations. *arXiv preprint arXiv:2311.02227*, 2023.

## A Implementation Details

The hyper-parameters setting of VDPO is presented in [Table 5](#). For baselines, we adopt similar hyper-parameters settings as suggested by original works, including A-SAC, DC/AC [\[5\]](#), DIDA [\[23\]](#), BPQL [\[20\]](#) and AD-SAC [\[42\]](#). The implementation of VDPO is based on CleanRL [\[16\]](#), and we also provide the code and guidelines to reproduce our results in the supplemental material. Each run of VDPO takes approximately 6 hours on 1 NVIDIA A100 GPU and 8 Intel Xeon CPUs.

Table 5: Hyper-parameters table of VDPO.

Hyper-parameter	Value
Buffer Size	1,000,000
Batch Size	256
Global Timesteps	1,000,000
Discount Factor	0.99
Reference Policy	
Learning Rate for Actor	3e-4
Learning Rate for Critic	1e-3
Network Layers	3
Network Neurons	[256, 256]
Activation	ReLU
Optimizer	Adam
Initial Entropy	$-\lvert \mathcal{A} \rvert$
Learning Rate for Entropy	1e-3
Train Frequency for Actor	2
Train Frequency for Critic	1
Soft Update Factor for Critic	5e-3
Delayed Policy (Transformer)	
Sequence Length	$\Delta$
Embedding Dimension	256
Attention Heads	1
Layers Num	3
Attention Dropout Rate	0.1
Residual Dropout Rate	0.1
Embedding Dropout Rate	0.1
Learning Rate	3e-4
Optimizer	Adam
Train Frequency for Belief Decoder	1
Train Frequency for Policy Decoder	1,000

## B Evidence Lower BOund (ELBO) of Delayed RL

Following the similar derivation sketch with [2, 25], we provide the derivation of Eq. (2) as follows.

$$\begin{aligned}
\log p_{\pi_\Delta}(O = 1) &= \log \int p(O = 1|\tau)p_{\pi_\Delta}(\tau)d\tau \\
&= \log \int p_\pi(\tau) \frac{p(O = 1|\tau)p_{\pi_\Delta}(\tau)}{p_\pi(\tau)} d\tau \\
&\geq \int p_\pi(\tau) \log \left[ \frac{p(O = 1|\tau)p_{\pi_\Delta}(\tau)}{p_\pi(\tau)} \right] d\tau \\
&= \mathbb{E}_{\tau \sim p_\pi(\tau)} [\log p(O = 1|\tau)] - \text{KL}(p_\pi(\tau) || p_{\pi_\Delta}(\tau)).
\end{aligned}$$

## C Theoretical Analysis

**Proposition C.1** (State-level KL divergence). *For a fixed reference policy  $\pi$ , the trajectory-level KL divergence can transform into the formulation of state-level KL divergence as follows*

$$\text{KL}(p_\pi(\tau) || p_{\pi_\Delta}(\tau)) = \underbrace{\sum_{t=0}^{\infty} \int d^\pi(s_t) \text{KL}(\pi(a_t|s_t) || \pi_\Delta(a_t|x_t)) ds_t}_{\text{State-level KL divergence}} + \text{Const.} \quad (12)$$

$$\begin{aligned}
\text{where Const.} &= \text{KL}(\rho(s_0) || \rho_\Delta(x_0)) \\
&+ \sum_{t=0}^{\infty} \int d^\pi(s_t) \int \pi(a_t|s_t) \text{KL}(\mathcal{P}(s_{t+1}|s_t, a_t) || b(s_t|x_t) \mathcal{P}_\Delta(x_{t+1}|x_t, a_t)) da_t ds_t.
\end{aligned}$$

*Proof.* The trajectory distribution  $p_\pi(\tau)$  induced by  $\pi$  is given by:

$$p_\pi(\tau) = \rho(s_0) \prod_{t=0}^{\infty} P(s_{t+1}|s_t, a_t) \pi(a_t|s_t).$$

Similarly, the trajectory distribution  $p_{\pi_\Delta}(\tau)$  induced by  $\pi_\Delta$  is given by:

$$p_{\pi_\Delta}(\tau) = \rho_\Delta(x_0) b(s_0|x_0) \prod_{t=0}^{\infty} b(s_{t+1}|x_{t+1}) \mathcal{P}_\Delta(x_{t+1}|x_t, a_t) \pi_\Delta(a_t|x_t).$$

Therefore, the trajectory-level KL divergence can be written as

$$\begin{aligned}
&\text{KL}(p_\pi(\tau) || p_{\pi_\Delta}(\tau)) \\
&= \mathbb{E}_{\tau \sim p_\pi(\tau)} [\log p_\pi(\tau) - \log p_{\pi_\Delta}(\tau)] \\
&= \mathbb{E}_{\tau \sim p_\pi(\tau)} [\log p(s_0) + \sum_{t=0}^{\infty} [\log [\mathcal{P}(s_{t+1}|s_t, a_t) \pi(a_t|s_t)]] \\
&\quad - \log [p(x_0) b(s_0|x_0)] - \sum_{t=0}^{\infty} [\log b(s_{t+1}|x_{t+1}) + \log \mathcal{P}_\Delta(x_{t+1}|x_t, a_t) + \log \pi_\Delta(a_t|x_t)]] \\
&= \mathbb{E}_{\tau \sim p_\pi(\tau)} \left[ \sum_{t=0}^{\infty} [\log \mathcal{P}(s_{t+1}|s_t, a_t) + \log \pi(a_t|s_t)] - \sum_{t=0}^{\infty} [\log b(s_t|x_t) + \log \mathcal{P}_\Delta(x_{t+1}|x_t, a_t) + \log \pi_\Delta(a_t|x_t)] \right] \\
&= \mathbb{E}_{\tau \sim p_\pi(\tau)} \left[ \log \left[ \frac{\rho(s_0)}{\rho_\Delta(x_0)} \right] + \sum_{t=0}^{\infty} \left[ \log \frac{\mathcal{P}(s_{t+1}|s_t, a_t)}{b(s_t|x_t) \mathcal{P}_\Delta(x_{t+1}|x_t, a_t)} + \log \frac{\pi(a_t|s_t)}{\pi_\Delta(a_t|x_t)} \right] \right] \\
&= \underbrace{\mathbb{E}_{\tau \sim p_\pi(\tau)} \left[ \log \left[ \frac{\rho(s_0)}{\rho_\Delta(x_0)} \right] + \sum_{t=0}^{\infty} \left[ \log \frac{\mathcal{P}(s_{t+1}|s_t, a_t)}{b(s_t|x_t) \mathcal{P}_\Delta(x_{t+1}|x_t, a_t)} \right] \right]}_C + \underbrace{\mathbb{E}_{\tau \sim p_\pi(\tau)} \left[ \sum_{t=0}^{\infty} \left[ \log \frac{\pi(a_t|s_t)}{\pi_\Delta(a_t|x_t)} \right] \right]}_D.
\end{aligned}$$

For  $C$ , we have

$$\begin{aligned}
C &= \text{KL}(\rho(s_0) \parallel \rho_\Delta(x_0)) + \sum_{t=0}^{\infty} \int p_\pi(\tau) \log \frac{\mathcal{P}(s_{t+1}|s_t, a_t)}{b(s_t|x_t)\mathcal{P}_\Delta(x_{t+1}|x_t, a_t)} d\tau \\
&= \text{KL}(\rho(s_0) \parallel \rho_\Delta(x_0)) + \sum_{t=0}^{\infty} \int d^\pi(s_t) \int \pi(a_t|s_t) \int \mathcal{P}(s_{t+1}|s_t, a_t) \log \frac{\mathcal{P}(s_{t+1}|s_t, a_t)}{b(s_t|x_t)\mathcal{P}_\Delta(x_{t+1}|x_t, a_t)} ds_{t+1} da_t ds_t \\
&= \text{KL}(\rho(s_0) \parallel \rho_\Delta(x_0)) + \sum_{t=0}^{\infty} \int d^\pi(s_t) \int \pi(a_t|s_t) \text{KL}(\mathcal{P}(s_{t+1}|s_t, a_t) \parallel b(s_t|x_t)\mathcal{P}_\Delta(x_{t+1}|x_t, a_t)) da_t ds_t \\
&\geq 0.
\end{aligned}$$

Therefore,  $C$  is a constant determined by the transition function of the dynamics and the fixed reference policy  $\pi$ .

Then, for  $D$ , we have

$$\begin{aligned}
D &= \sum_{t=0}^{\infty} \int p_\pi(\tau) \log \frac{\pi(a_t|s_t)}{\pi_\Delta(a_t|x_t)} d\tau \\
&= \sum_{t=0}^{\infty} \int d^\pi(s_t) \int \pi(a_t|s_t) \log \frac{\pi(a_t|s_t)}{\pi_\Delta(a_t|x_t)} da_t ds_t \\
&= \sum_{t=0}^{\infty} \int d^\pi(s_t) \text{KL}(\pi(a_t|s_t) \parallel \pi_\Delta(a_t|x_t)) ds_t.
\end{aligned}$$

□

**Lemma C.2** (Sample complexity of VDPO). *Assumed that maximizing  $A$  in Eq. (3) by model-based policy iteration while minimizing  $B$  in Eq. (3) by behaviour cloning, VDPO finds an  $\epsilon$ -optimal policy with probability  $1 - \sigma$  using sample size*

$$\mathcal{O}\left(\max\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma}, \frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)\right).$$

*Proof.* Applying Lem. 3.2 and Lem. 3.4. □

**Proposition C.3** (Sample complexity comparison). *In the delayed MDP, as  $\sigma \rightarrow 0$ , the sample complexity of VDPO (Lem. 3.5) is less or equal to the sample complexity of model-based policy iteration (Lem. 3.2):*

$$\mathcal{O}\left(\max\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma}, \frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma\right)\right) \leq \mathcal{O}\left(\frac{|\mathcal{X}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma}\right).$$

*Proof.* It is obvious that

$$\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma} \leq \frac{|\mathcal{X}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma},$$

as we have  $|\mathcal{S}| \leq |\mathcal{X}| = |\mathcal{S}||\mathcal{A}|^\Delta$ .

Then, we show that

$$\frac{|\mathcal{X}|\ln|\mathcal{A}|}{(1-\gamma)^4\epsilon^2}\sigma \leq \frac{|\mathcal{X}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\sigma},$$

which is equivalent to

$$\frac{\ln|\mathcal{A}|}{|\mathcal{A}|} \leq -(1-\gamma)\frac{\ln\sigma}{\sigma}.$$

This inequality always holds when  $\sigma \rightarrow 0$  as

$$\lim_{\sigma \rightarrow 0} -(1-\gamma)\frac{\ln\sigma}{\sigma} = +\infty \gg \frac{1}{e} > \frac{\ln|\mathcal{A}|}{|\mathcal{A}|}.$$

□

**Lemma C.4** (Convergence of delayed policy in VDPO). *Let  $\pi^*$  be the optimal reference policy which is trained by a delay-free RL algorithm. The delayed policy  $\pi_\Delta$  converges to  $\pi_\Delta^*$  satisfying that*

$$\pi_\Delta^*(a_t|x_t) = \mathbb{E}_{s_t \sim b(\cdot|x_t)} [\pi^*(a_t|s_t)], \forall x_t \in \mathcal{X}. \quad (13)$$

*Proof.* We can derive the result from the solution of Eq. (6).  $\square$

**Proposition C.5** (Consistent fixed point). *VDPO shares the same fixed point (Eq. (7)) with DIDA [23], BPQL [20] and AD-SAC [42] for the same delayed MDP.*

*Proof.* The fixed points of DIDA (Eq. (3) in [23]), BPQL (Eq. (23) in [20]) and AD-SAC (Theorem 5.9 in [42]) all are

$$\pi_\Delta^*(a_t|x_t) = \mathbb{E}_{s_t \sim b(\cdot|x_t)} [\pi^*(a_t|s_t)], \forall x_t \in \mathcal{X}.$$

which is consistent with the fixed point of VDPO.  $\square$

## D Additional Experimental Results

In MuJoCo with 25 and 50 constant delays, We report the required steps to hit this threshold within 1M global steps in [Table 6](#) and [Table 7](#) respectively.

Table 6: Amount of steps required to hit the threshold  $Ret_{df}$  in MuJoCo tasks with 25 constant delays within 1M global steps, where  $\times$  denotes that failed to hit the threshold within 1M global steps. The best result is in blue.

Task (Delays=25)	A-SAC	DC/AC	DIDA	BPQL	AD-SAC	VDPO (ours)
Ant-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
HalfCheetah-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
Hopper-v4	$\times$	$\times$	$\times$	<b>0.69M</b>	$\times$	$\times$
Humanoid-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
HumanoidStandup-v4	$\times$	0.38M	$\times$	<b>0.09M</b>	<b>0.09M</b>	0.48M
Pusher-v4	$\times$	0.09M	0.10M	0.02M	0.03M	<b>0.01M</b>
Reacher-v4	$\times$	0.83M	$\times$	$\times$	$\times$	<b>0.22M</b>
Swimmer-v4	$\times$	$\times$	$\times$	0.39M	0.12M	<b>0.09M</b>
Walker2d-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$

Table 7: Amount of steps required to hit the threshold  $Ret_{df}$  in MuJoCo tasks with 50 constant delays within 1M global steps, where  $\times$  denotes that failed to hit the threshold within 1M global steps. The best result is in blue.

Task (Delays=50)	A-SAC	DC/AC	DIDA	BPQL	AD-SAC	VDPO (ours)
Ant-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
HalfCheetah-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
Hopper-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
Humanoid-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
HumanoidStandup-v4	$\times$	0.68M	$\times$	0.21M	<b>0.08M</b>	0.84M
Pusher-v4	$\times$	0.14M	0.10M	0.18M	0.02M	<b>0.01M</b>
Reacher-v4	$\times$	$\times$	$\times$	$\times$	$\times$	<b>0.39M</b>
Swimmer-v4	$\times$	$\times$	$\times$	0.29M	<b>0.11M</b>	0.15M
Walker2d-v4	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$

## E Learning Curves

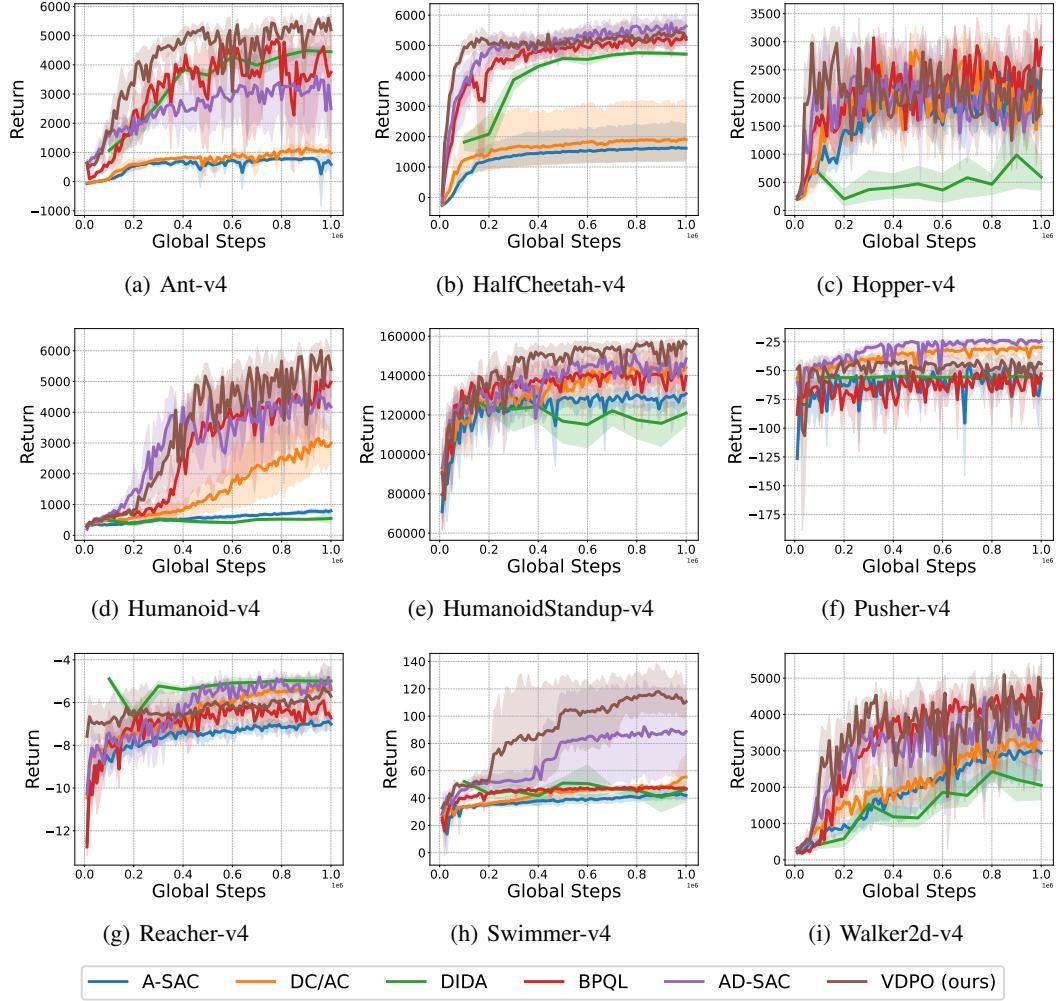


Figure 2: Learning curves in MuJoCo tasks with 5 constant delays where the shaded areas represented the standard deviation.

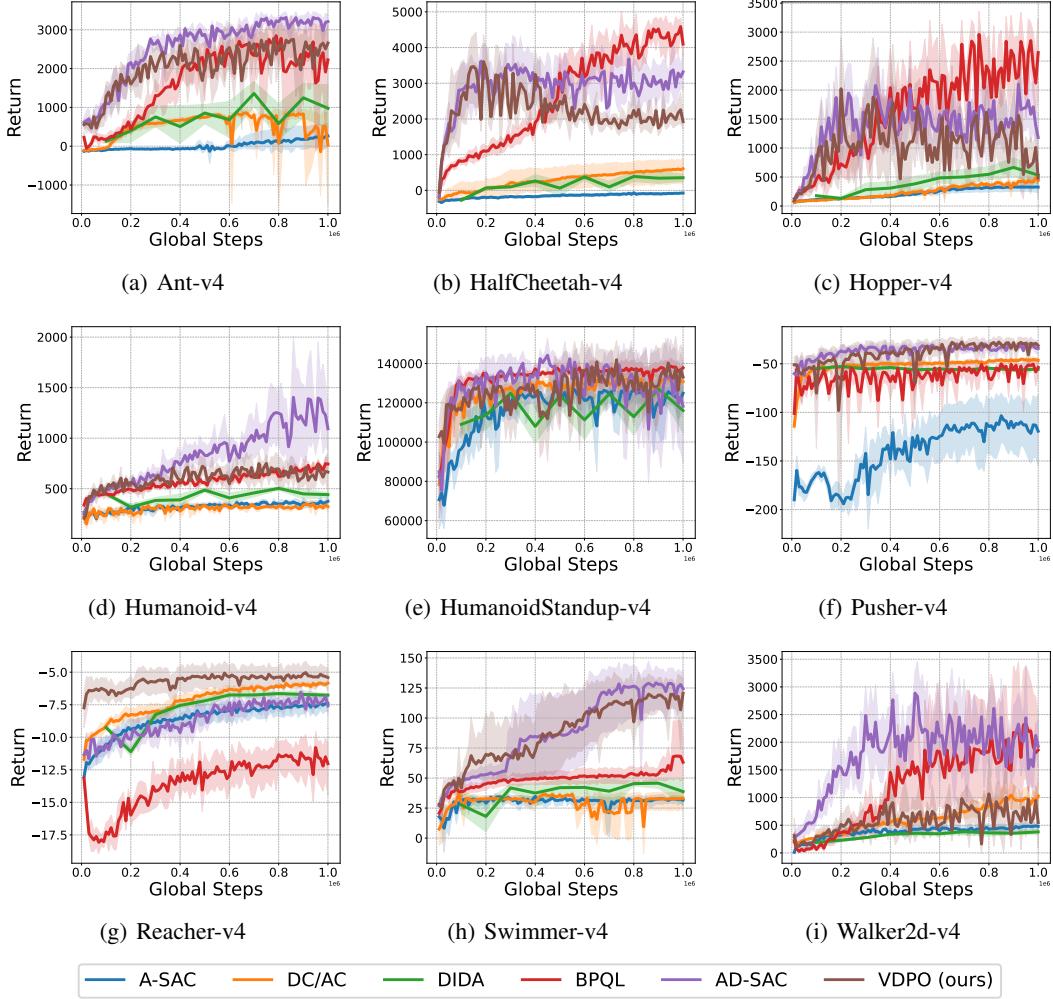


Figure 3: Learning curves in MuJoCo tasks with 25 constant delays where the shaded areas represented the standard deviation.

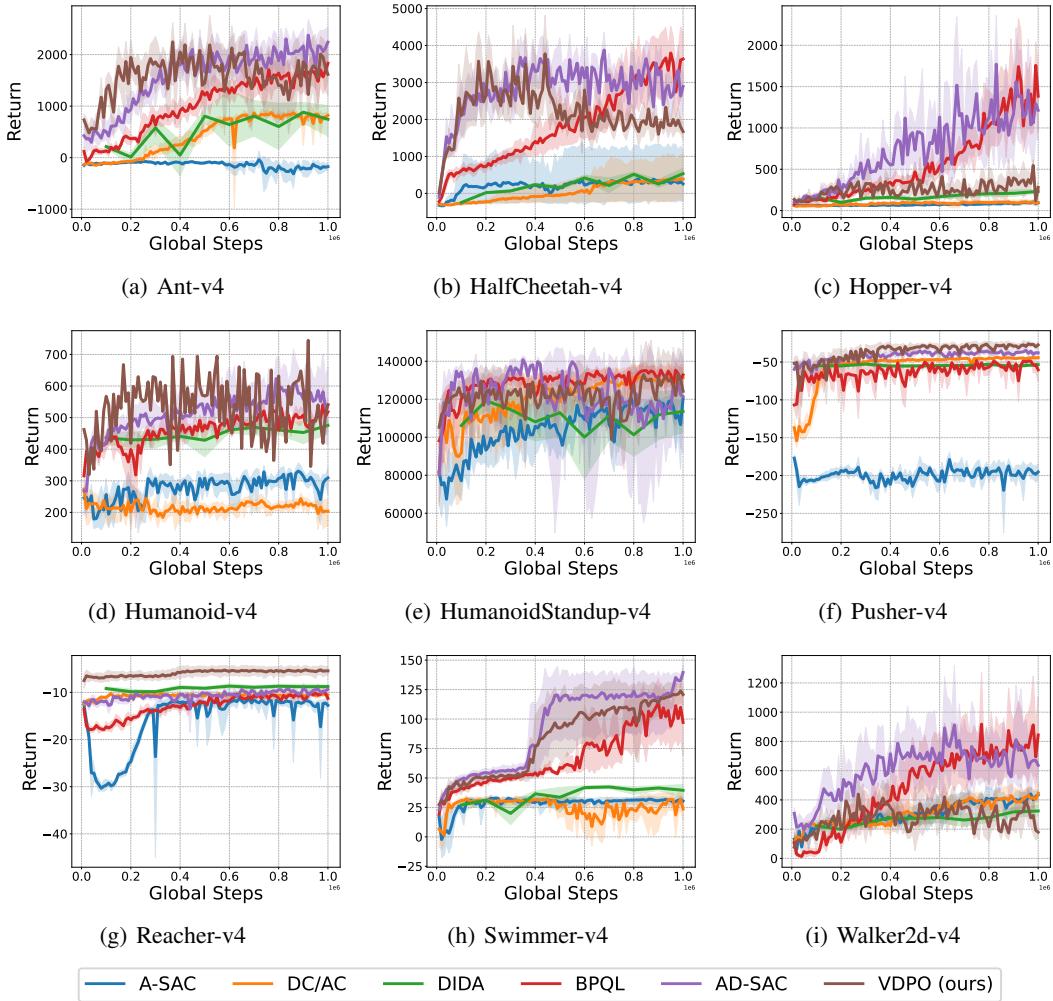


Figure 4: Learning curves in MuJoCo tasks with 50 constant delays where the shaded areas represented the standard deviation.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: We have clearly claimed our contribution and scope in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: In Section 4.2.3, we have discussed the limitation that this work suffers from performance degeneration issues in stochastic environments.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We have provided the full set of assumptions and complete and correct proof in the main text and appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We have provided all the information needed to reproduce our experimental results, including implementation details in the main text and appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We have provided our code to reproduce the experimental results in the supplemental material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We have provided the implementation details in the main text and appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: We have reported the experimental results averaged over multiple random seeds.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We have provided the computer resource in the implementation detail of the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: The research conducted in this paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: This paper aims to advance the field of Machine Learning. We believe there are no potential societal impacts of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: NA

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: NA

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: NA

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: NA

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: NA

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.