

A Model-Based Approach for Improving Reinforcement Learning Efficiency Leveraging Expert Observations

Erhan Can Ozcan¹, Vittorio Giammarino¹, James Queeney², and Ioannis Ch. Paschalidis³

Abstract— This paper investigates how to incorporate expert observations (without explicit information on expert actions) into a deep reinforcement learning setting to improve sample efficiency. First, we formulate an augmented policy loss combining a maximum entropy reinforcement learning objective with a behavioral cloning loss that leverages a forward dynamics model. Then, we propose an algorithm that automatically adjusts the weights of each component in the augmented loss function. Experiments on a variety of continuous control tasks demonstrate that the proposed algorithm outperforms various benchmarks by effectively utilizing available expert observations.

I. INTRODUCTION

Deep Reinforcement Learning (RL) has shown its potential to solve complex sequential decision making problems with its impressive success in learning to play video games [1], autonomous driving [2], and robotics applications [3]. Among RL algorithms, model-free approaches are appealing since they can serve as general-purpose tools for learning complex tasks [4], [5]. However, these approaches are generally not sample efficient, and often require a significant number of interactions with the environment during learning. Since this process can be costly, time-consuming, and unsafe, the use of these algorithms can be limited in real-world problems [6], [7].

The efficiency issue of RL algorithms can be mitigated by leveraging available expert data. Unfortunately, expert data contain limited information, possibly having only access to observations of the expert (e.g., states), but not to their private decision making (e.g., actions). This mandates estimating expert actions or rewards before employing an RL algorithm to learn a policy [8], [9]. Since the performance of these types of algorithms can be susceptible to estimation errors, leveraging expert observations remains as a challenging task [10], [11].

In this study, we focus on improving the performance of an RL algorithm by leveraging expert observations. First,

ECO, VG, and ICP were partially supported by the DOE under grants DE-AC02-05CH11231 and DE-EE0009696, the NSF under grants CCF-2200052, DEB-2433726, EECs-2317079, and IIS-1914792, and by the ONR under grants N00014-19-1-2571 and N00014-21-1-2844. JQ was exclusively supported by Mitsubishi Electric Research Laboratories (MERL).

¹Erhan Can Ozcan and Vittorio Giammarino are with Division of Systems Engineering, Boston University, Boston, MA 02215, USA {cozcan, vgiammar}@bu.edu

²James Queeney is with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA queeney@merl.com

³Ioannis Ch. Paschalidis is with the Department of Electrical and Computer Engineering, Division of Systems Engineering, and Department of Biomedical Engineering, Boston University, Boston, MA 02215, USA yannisp@bu.edu

we define a policy loss function by combining a maximum entropy RL objective with a behavioral cloning loss to imitate the state transitions of the expert. This behavioral cloning loss leverages a forward dynamics model, which is learned in an online manner, to address the absence of expert actions. Next, to mitigate the detrimental effect that an inaccurate forward model can have on policy learning, we propose a practical algorithm that automatically adjusts the weight of the behavioral cloning loss during policy learning based on the accuracy of the model. Finally, by conducting experiments in the DeepMind Control Suite [12], we demonstrate that utilizing available expert data can significantly accelerate the training of RL algorithms.

II. RELATED WORK

One possible way to simplify policy learning in RL is to consider a set of expert data and utilize offline RL algorithms [13], [14], [15], [16]. However, all these studies assume that the expert dataset is collected in a demonstration format, i.e., sequences of states, actions, and reward tuples, which is unrealistic in many scenarios. Therefore, it is important to consider alternative algorithms that can work with expert datasets containing limited information.

Imitation Learning (IL) is another field that investigates how to incorporate expert data into policy learning in the absence of reward information. The goal in *Inverse Reinforcement Learning* (IRL), a popular sub-field of IL, is to retrieve a reward signal based on the expert data and utilize this information to learn a policy [11]. Various IRL approaches utilize the expert's state-action pairs [17], [18], [19]. However, data showing both expert states and actions on a specific task can be either limited or unavailable, which hinders the deployment of these types of algorithms in real-world problems. On the contrary, expert observations can be easily accessible in today's world, thus attention has shifted towards studies that focus on learning a policy by using only expert states. In [20], Generative Adversarial Imitation Learning proposed in [18] has been extended to a state-only setting. In [21], an algorithm that relies only on expert states has been developed to learn a non-stationary policy for discrete actions. In [22], a model-based framework, which promotes exploration in the face of uncertainty, is proposed to learn a policy solely from expert states. Finally, two recent works explore the possibility of adversarial imitation from expert's visual observations [23], [24]. However, all these approaches optimize an adversarial min-max problem during policy learning since they lack a reward signal, making optimization difficult [11], [25], [26].

Behavioral Cloning (BC) is another sub-field of IL, where the goal is to learn a direct mapping from states to actions in a supervised manner. In this field, one group of work assumes that there exists an expert providing the optimal action for any given state, and propose BC approaches based on dataset aggregation [27], [28]. However, having this type of an expert can be either extremely costly or unrealistic in many cases. In [11], [29], BC algorithms that can be trained on offline expert demonstrations are outlined, where the agent does not need to interact with the environment during policy learning, but requiring access to expert actions along with the states. Similar to IRL, studies that focus on learning a policy by using only expert states are of interest in BC. In [30], a BC approach is introduced where an expert can be imitated by learning a latent forward dynamics model. Finally, in [31], a BC approach is proposed that utilizes only expert states. Unfortunately, pure BC approaches may suffer from compounding errors as a result of distribution shift [27], [32]. Therefore, the performance of BC can be poor in domains where consecutive actions are required for success [20].

In order to address the shortcomings of both IL and RL approaches, combining these two disciplines is an appealing strategy when possible, i.e., a reward signal can be received from the environment after the interaction. In [33], a first-order oracle provides policy update information for any given state during the policy imitation stage, which is followed by traditional RL steps. In [34], policy learning begins with IL on available expert demonstrations, and then the learned policy is refined via RL to achieve a human level performance. Finally, in [35], an actor-critic algorithm is proposed by optimizing an RL objective combined with a behavioral cloning loss defined on expert demonstrations. This is similar to the augmented loss considered in this work, but the approach in [35] utilizes expert actions during policy training that we assume are not available. Finally, the approaches that combine IL and RL introduce a pre-training phase before policy learning [33], [34], [35]. However, in our case, it is not possible to perform these pre-training techniques since expert actions are unknown.

State-only expert data can be incorporated into policy optimization by learning a world model, which reflects the transition dynamics. Neural networks are powerful function approximators that can model complex non-linear dynamics even in high-dimensional problems; thus, they can be good candidates for modeling the dynamics in real-world problems [36], [37]. Most of the existing model-based approaches utilize the learned model as a simulator to generate huge amounts of data, which is how they achieve sample efficiency during policy optimization [38], [39]. One common drawback of this strategy is the challenge of determining the length of the counterfactual trajectories generated by the model since the model errors accumulate as the model horizon increases. Therefore, contrary to existing work, we utilize the learned predictive model solely to facilitate the use of expert observations in our work.

III. PRELIMINARIES

Consider an infinite horizon *Markov Decision Process* (MDP) characterized by the tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma, \rho_0)$. Following the standard notation, $\mathcal{S} \subseteq \mathbb{R}^n$ and $\mathcal{A} \subseteq \mathbb{R}^m$ denote the continuous state and action spaces, respectively, $p : \mathcal{S} \times \mathcal{A} \rightarrow P(\mathcal{S})$ is the transition model where $P(\mathcal{S})$ denotes the space of probability measures over \mathcal{S} , $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $\gamma \in (0, 1)$ is the discount rate, and ρ_0 is the initial state distribution.

The agent interacts with the environment by following a stochastic policy parameterized by θ , $\pi_\theta : \mathcal{S} \rightarrow P(\mathcal{A})$, and the expected sum of discounted rewards $\eta(\pi)$ can be expressed as follows:

$$\eta(\pi) = \mathbb{E}_{(s_t, a_t) \sim \rho_\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right],$$

where $(s_t, a_t) \sim \rho_\pi$ denotes a trajectory sampled according to $s_0 \sim \rho_0$, $a_t \sim \pi_\theta(\cdot | s_t)$, and $s_{t+1} \sim p(\cdot | s_t, a_t)$.

Then, the goal of standard RL is to find a policy maximizing the expected sum of discounted rewards. By augmenting the standard reinforcement learning objective with an entropy term, the maximum entropy objective can be obtained [40]. Accordingly, the optimal policy π^* of this problem maximizes the entropy visited at each state as well as the discounted reward:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{(s_t, a_t) \sim \rho_\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) + \alpha H(\pi(\cdot | s_t)) \right],$$

where H represents entropy and α is the temperature parameter adjusting the relative weight of the entropy term with respect to the discounted reward. *Soft Actor-Critic* (SAC) [41] is an effective off-policy algorithm within this maximum entropy RL framework, and it alternates between estimating a soft Q-value function $Q_\psi(s, a)$ using the soft Bellman backup operator and a policy improvement step to improve the current policy π_θ . Using a replay buffer \mathcal{D} that stores previously sampled states and actions, the soft Q-value parameters, ψ , can be obtained by optimizing the following soft Bellman residual error:

$$J_Q(\psi, \mathcal{D}) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[\frac{1}{2} (Q_\psi(s_t, a_t) - (r_t + \gamma V_{\bar{\psi}}(s_{t+1})))^2 \right],$$

where

$$V_{\bar{\psi}}(s_{t+1}) = \mathbb{E}_{a_{t+1} \sim \pi_\theta} [Q_{\bar{\psi}}(s_{t+1}, a_{t+1}) - \alpha \log \pi_\theta(a_{t+1} | s_{t+1})], \quad (1)$$

and $\bar{\psi}$ are delayed target parameters that are maintained via exponential smoothing with coefficient τ . Then, SAC updates the policy towards the exponential of soft Q-value based on the following maximum entropy policy loss:

$$J_\pi(\theta, \mathcal{D}) = \mathbb{E}_{s_t \sim \mathcal{D}, a_t \sim \pi_\theta} [\alpha \log \pi_\theta(a_t | s_t) - Q_\psi(s_t, a_t)]. \quad (2)$$

The final ingredient in SAC is to devise a mechanism to automate the temperature selection for each task based on

a target entropy value \bar{H} . Therefore, α is updated based on the gradients of the following function:

$$J(\alpha, \mathcal{D}) = \mathbb{E}_{s_t \sim \mathcal{D}, a_t \sim \pi_\theta} [-\alpha \log \pi_\theta(a_t | s_t) - \alpha \bar{H}]. \quad (3)$$

The algorithm proposed in this paper builds on top of this maximum entropy RL framework. While we incorporate the same Q-value estimation step and temperature update rule into our algorithm, we modify the policy loss given in (2) to utilize expert observations during policy updates, thereby significantly accelerating training.

IV. MAXIMUM ENTROPY POLICY LEARNING WITH EXPERT OBSERVATIONS

We now present our algorithm *Soft Actor-Critic with Expert Observations* (SAC-EO), which learns a policy with maximum entropy RL while also minimizing deviations from the expert’s state trajectory. The algorithm has two essential components: (1) We learn a forward dynamics model that allows us to utilize expert observations to improve policy learning without requiring expert actions, and (2) we propose an augmented policy objective that combines the maximum entropy RL objective with a behavioral cloning loss, leveraging expert observations.

Incorporation of Expert Data: Although we only observe the expert states, it is still possible to utilize this information during policy optimization by learning a forward dynamics model. Learning a global model that can accurately predict the next state and the reward for any given state-action pair can be impractical, as it requires exploring all parts of the state space. Furthermore, learning a globally accurate model can be impossible when the system dynamics are complex [38]. However, this is not necessary for the success of our algorithm since we employ our model solely to incorporate expert data into the policy improvement.

Our goal is to modify the policy π_θ by utilizing a learned model so that the policy can generate transitions similar to the expert. We model the transition dynamics using a probabilistic feed-forward neural network parameterized by ϕ , whose outputs represent the mean vector μ and diagonal covariance matrix Σ of a Gaussian distribution. Then, for a state-action pair (s_t, a_t) at time t , the trained model \hat{m}_ϕ estimates the state at time $t + 1$:

$$\hat{s}_{t+1} \sim \hat{m}_\phi(s_t, a_t) = \mathcal{N}(\mu_\phi(s_t, a_t), \Sigma_\phi(s_t, a_t)).$$

Similar to model-based RL algorithms, we refine the learned model gradually as the policy improves to overcome the state *distribution shift* issue [38], [39]; more details related to model learning can be found in [42].

Suppose we are provided a set of expert observations (states) $\mathcal{D}^e \equiv \{s_t^e\}_{t=0}^N$. Then, for any given expert state s_t^e , we can estimate the next state of our policy by utilizing the learned model: $\hat{s}_{t+1} \sim \hat{m}_\phi(s_t^e, \pi_\theta(s_t^e))$. Note that \hat{s}_{t+1} is a function of the policy parameters θ , and our goal is to imitate the state transitions of the expert. Therefore, for a given model, policy and expert dataset, we can define a

term measuring the deviation from the state trajectory of the expert by using our model:

$$\text{MSE}(\phi, \theta, \mathcal{D}^e) = \frac{\sum_{t=0}^{N-1} \|\hat{s}_{t+1} - s_{t+1}^e\|^2}{N}, \quad (4)$$

where $\|\cdot\|$ denotes the ℓ_2 norm of a vector throughout the paper. Since we learn a parametric policy and model, we can calculate the gradient of the MSE term with respect to θ , and utilize this information during policy optimization. Finally, note that our learned model is only used to calculate the MSE term in (4), unlike model-based RL methods that utilize the learned model as a simulator to generate additional data for training.

Augmented Policy Objective: Similar to [41], our algorithm optimizes an RL objective, namely the maximum entropy RL objective given in (2), based on data available in the replay buffer \mathcal{D} . In addition to this, our algorithm takes into account the expert data with the help of the learned model, and minimizes the deviations from the state trajectory of the expert defined in (4). The augmented function minimized by our algorithm at the policy improvement stage is as follows:

$$J_\pi(\theta, \phi, \mathcal{D}, \mathcal{D}^e, \epsilon) = (1 - \epsilon) J_\pi(\theta, \mathcal{D}) + \epsilon \text{MSE}(\phi, \theta, \mathcal{D}^e), \quad (5)$$

where ϵ is the expert state matching parameter controlling the relative importance of the MSE term compared to the soft policy improvement term. Next, we describe how to automatically adjust ϵ throughout training.

A. Automatic Adjustment of the Expert State Matching Coefficient

The ingredients above describe how we can incorporate expert observations into a maximum entropy RL algorithm. Nevertheless, the success of our algorithm heavily depends on the quality of the learned model, which is used to calculate the MSE term in (5). In the early stages of training, the learned model may not be able to accurately represent transition dynamics in states visited by the expert. In addition, it may be difficult to learn an accurate model in environments with complex dynamics. In these scenarios, we want to place more emphasis on the RL objective.

To integrate this concept into our framework, it is important to devise a strategy that determines the importance of each term based on the accuracy of the model on the expert dataset, and automatically adjusts the value of the expert state matching parameter. However, evaluating the performance of the model on the expert trajectories requires knowledge of both expert actions and states. Since we only have access to expert states, evaluating the performance of the model is not trivial.

In model-based RL frameworks, it is common to learn a model ensemble and to define a discrepancy measure between models to assess their accuracy. For example, [43] utilizes a discrepancy measure to prevent policies from visiting states where transition dynamics are uncertain, while [22] utilizes a discrepancy measure to incentivize the policy to explore unknown parts of the state space. In our study, we want to assess the reliability of our models on the states

visited by the expert. Therefore, we design a rule based on this measure.

Suppose that we have a pair of dynamics models, \hat{m}_{ϕ_1} and \hat{m}_{ϕ_2} , that are initialized with the parameters ϕ_1 and ϕ_2 , respectively, but trained on the same data. We can define the maximum discrepancy δ^{max} between models on expert dataset \mathcal{D}^e as follows:

$$\delta^{max}(\mathcal{D}^e) = \max_{s \in \mathcal{D}^e} \|\hat{m}_{\phi_1}(s, \pi_\theta(s)) - \hat{m}_{\phi_2}(s, \pi_\theta(s))\|. \quad (6)$$

Using this discrepancy measure, we can define an adaptive expert state matching coefficient by

$$\epsilon_k = \frac{1}{1 + \beta \delta^{max}(\mathcal{D}^e)}, \quad (7)$$

where $\beta \geq 0$ is a tunable parameter to scale the discrepancy. Then, the policy loss with an adaptive epsilon ϵ_k can be expressed as follows:

$$J_\pi(\theta, \phi, \mathcal{D}, \mathcal{D}^e, \epsilon_k) = (1 - \epsilon_k) J_\pi(\theta, \mathcal{D}) + \epsilon_k \frac{\sum_{i=1}^2 \text{MSE}(\phi_i, \theta, \mathcal{D}^e)}{2}. \quad (8)$$

Finally, we assume that our algorithm is trained by collecting K trajectories, each of length E , and we update the models and adaptive state matching coefficient at the end of each trajectory. The pseudocode showing the steps of SAC-EO is given in Algorithm 1.

Algorithm 1 SAC-EO

Input: Scale parameter $\beta \geq 0$, Expert observations \mathcal{D}^e .

- 1: Initialize policy π_θ , predictive models \hat{m}_{ϕ_1} and \hat{m}_{ϕ_2} , critics Q_{ψ_1} and Q_{ψ_2} , target critics \bar{Q}_{ψ_1} and \bar{Q}_{ψ_2} , replay buffer \mathcal{D} , model training buffer \mathcal{D}^m .
 - 2: **for** $k = 1 \dots K$ **do**
 - 3: Train models \hat{m}_{ϕ_1} and \hat{m}_{ϕ_2} on \mathcal{D}^m via maximum likelihood.
 - 4: $\delta^{max} = \max_{s \in \mathcal{D}^e} \|\hat{m}_{\phi_1}(s, \pi_\theta(s)) - \hat{m}_{\phi_2}(s, \pi_\theta(s))\|$.
 - 5: Set $\epsilon_k = \frac{1}{1 + \beta \delta^{max}}$.
 - 6: Sample a mini-batch of expert states from \mathcal{D}^e , and store them in \mathcal{D}_b^e .
 - 7: **for** E steps **do**
 - 8: $a_t \sim \pi_\theta(\cdot | s_t)$
 - 9: $s_{t+1}, r_t \sim \text{EnvironmentStep}(a_t)$
 - 10: $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, r_t, s_{t+1})$
 - 11: $\mathcal{D}^m \leftarrow \mathcal{D}^m \cup (s_t, a_t, s_{t+1})$
 - 12: Sample a mini-batch of environment data from \mathcal{D} , and store them in \mathcal{D}_b .
 - 13: Split \mathcal{D}_b^e into two parts randomly: $\mathcal{D}_{b,1}^e$ and $\mathcal{D}_{b,2}^e$.
 - 14: $\psi_i \leftarrow \psi_i - \lambda_Q \nabla_{\psi_i} J_Q(\psi_i, \mathcal{D}_b)$ for $i \in \{1, 2\}$
 - 15: $L_\theta = (1 - \epsilon_k) J_\pi(\theta, \mathcal{D}_b) + \epsilon_k \frac{\sum_{i=1}^2 \text{MSE}(\phi_i, \theta, \mathcal{D}_{b,i}^e)}{2}$
 - 16: $\theta \leftarrow \theta - \lambda_\pi \nabla_\theta L_\theta$
 - 17: $\alpha \leftarrow \alpha - \lambda \nabla_\alpha J(\alpha, \mathcal{D}_b)$
 - 18: $\bar{\psi}_i \leftarrow \tau \psi_i + (1 - \tau) \bar{\psi}_i$ for $i \in \{1, 2\}$
 - 19: **end for**
 - 20: **end for**
 - 21: **return** π_θ .
-

Note that setting $\beta = 0$ in Algorithm 1 results in a modified version of the algorithm *Behavioral Cloning From Observation* (BCO) proposed in [31], and the algorithm converges to classic SAC as $\beta \rightarrow \infty$.

V. EXPERIMENTS

We analyze the use of expert observations within a maximum entropy RL framework over 6 tasks available in the DeepMind Control Suite [12]. The tasks we consider include *cheetah-run*, *hopper-hop*, *hopper-stand*, *walker-run*, *walker-walk*, and *walker-stand*, which all have fixed-length horizons of 1000 steps with rewards in the unit interval, i.e., $r(s, a) \in [0, 1]$. We obtain expert observations by training a different model-free RL algorithm *Maximum a Posteriori Policy Optimization* (MPO) [5] for 5 million steps. The full details of the SAC-EO algorithm, including network architectures, model training, and all hyperparameters are summarized in [42]. Finally, the source code is available on GitHub to promote reproducibility¹.

A. Performance Comparison on DeepMind Control Suite

Although our strategy to incorporate expert data into policy improvement can be generalized into any model-free RL algorithm, we build our algorithm on top of SAC due to its flexibility to learn complex tasks. Therefore, the first benchmark we consider in our experiments is SAC. Another benchmark we consider in our experiments is MPO since expert data required by our algorithm are generated by employing an MPO agent. We consider a problem setting where only expert states are available during training, thus offline RL algorithms cannot be used as baselines. However, we include a modified version of BCO, obtained by setting $\beta = 0$ in Algorithm 1, among the benchmarks for a fair comparison. The reason to use modified-BCO is to ensure a one to one correspondence between methods, and one fundamental difference between modified-BCO and BCO is that modified-BCO utilizes a forward dynamics model instead of an inverse model during policy optimization. Figure 1 shows the average of total reward obtained by our algorithm across 5 seeds on different tasks, compared to SAC, MPO, and modified-BCO.

According to Figure 1, SAC-EO matches the performance of the expert, which was trained for five millions steps, in less than one million steps for 4 out of 6 tasks. Furthermore, utilizing expert observations during policy updates accelerates the training of SAC in all tasks. We further analyze this point in Table I, which reports the number of steps each algorithm needs to reach a percentage of the expert performance. Lastly, the poor performance of modified-BCO in all tasks reveals that the behavioral cloning loss defined on expert observations alone is insufficient to learn a good policy, since it is not possible to learn a model that reflects the transition dynamics on the expert states unless the agent explores this part of the state space. This highlights the importance of utilizing this type of a loss within an RL framework.

¹<https://github.com/noc-lab/sac-expert>.

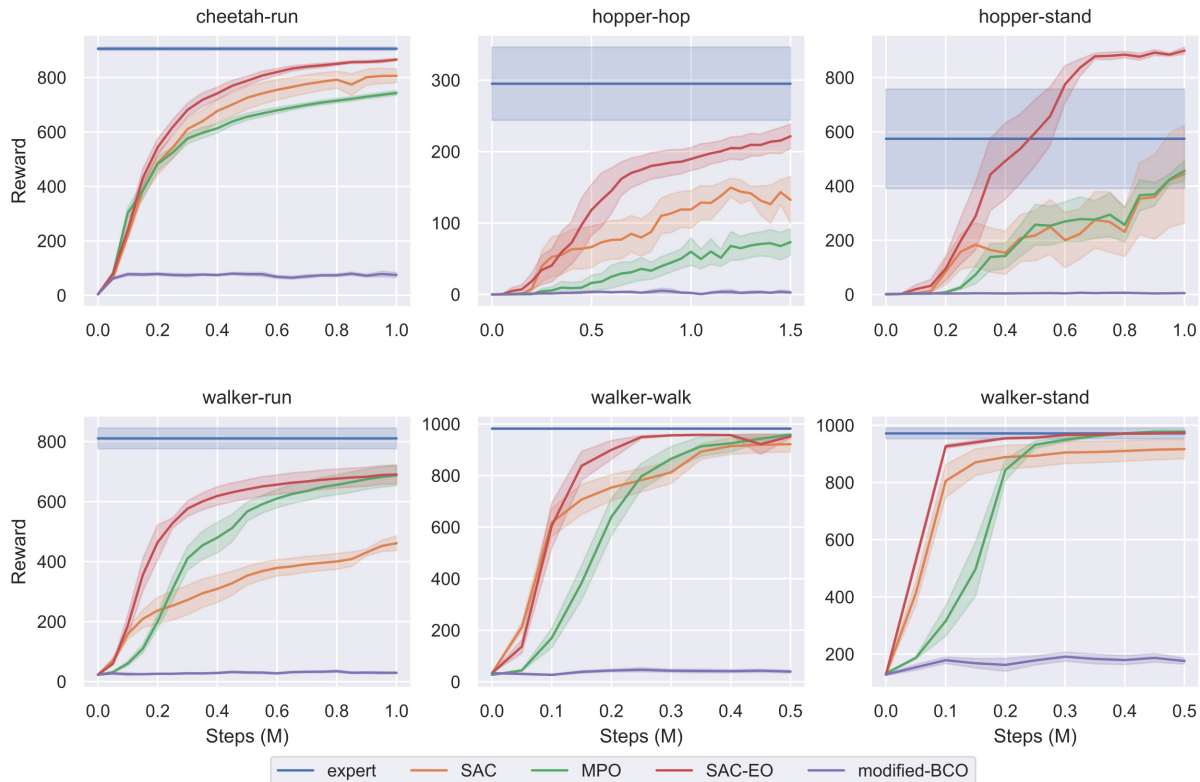


Fig. 1. Comparison of algorithms across tasks. Horizontal blue line represents the expert performance. SAC-EO and modified-BCO are supplied four expert trajectories during training. Shading denotes one standard error across policies.

TABLE I

STEPS TO REACH X PERCENTAGE OF THE EXPERT PERFORMANCE²

	75% of the Expert			95% of the Expert		
	MPO	SAC	SAC-EO	MPO	SAC	SAC-EO
cheetah-run	577K	368K	265K	NA	NA	770K
hopper-hop	NA	NA	NA	NA	NA	NA
hopper-stand	508K	250K	227K	799K	627K	287K
walker-run	571K	NA	330K	NA	NA	NA
walker-walk	190K	139K	83K	376K	357K	185K
walker-stand	160K	47K	29K	187K	268K	48K

TABLE II

THE BENEFIT OF UTILIZING ADAPTIVE ALGORITHM

Fixed vs Adaptive	$\epsilon_k = 0.1$	$\epsilon_k = 0.01$	$\epsilon_k = 0.001$	Adaptive
cheetah-run	614±93	808±33	843±34	867±11
hopper-hop	38±77	86±96	196±32	191±42
hopper-stand	910±24	564±435	602±419	919±25
walker-run	23±2	253±169	623±127	695±62
walker-stand	963±30	981±9	982±9	985±10
walker-walk	23±1	948±28	955±41	967±11

compared to using a fixed ϵ_k in most tasks.

B. The Benefit of Adaptive Epsilon

Next, we show the benefit of utilizing an adaptive ϵ_k , as we describe, over a fixed ϵ_k in our algorithm. Therefore, we train our algorithm by setting ϵ_k to 0.1, 0.01, and 0.001, respectively, and compare them with our adaptive algorithm. The average performance of each version after one million interactions with the environment is summarized in Table II.

According to Table II, it is promising to utilize our adaptive strategy, as the adaptive version performs better

²In Table I, NA means that the specified percentage of the expert performance cannot be achieved by the corresponding algorithm in one million steps.

VI. CONCLUSION

In this paper, we demonstrated how to increase the sample efficiency of an RL algorithm by utilizing expert observations without access to expert actions. We introduced the algorithm SAC-EO, which performs policy updates based on a loss function combining a maximum entropy RL objective with a behavioral cloning loss that utilizes a learned forward dynamics model. Based on our experiments, SAC-EO significantly accelerates training compared to standard model-free RL algorithms. Although our algorithm is built upon SAC, the strategy we present to incorporate expert observations into policy improvement is a general framework that can work with many RL algorithms. Therefore, it can

be interesting to explore the effect of employing our idea on different RL algorithms in future work. Furthermore, since it does not need explicit action information from experts, the proposed algorithm may have the potential to take advantage of large amounts of video available online during policy improvement. Hence, another interesting future direction can be to extend our work by incorporating visual observations.

REFERENCES

- [1] 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.
- [2] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2021.
- [3] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 297–330, 2020.
- [4] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [5] 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.
- [6] J. Queeney, Y. Paschalidis, and C. G. Cassandras, "Generalized proximal policy optimization with sample reuse," in *Advances in Neural Information Processing Systems*, vol. 34, pp. 11 909–11 919, 2021.
- [7] J. Queeney, E. C. Ozcan, I. Paschalidis, and C. Cassandras, "Optimal transport perturbations for safe reinforcement learning with robustness guarantees," *Transactions on Machine Learning Research*, 2024. [Online]. Available: <https://openreview.net/forum?id=cgSXpAR4G1>
- [8] Y. Liu, A. Gupta, P. Abbeel, and S. Levine, "Imitation from observation: Learning to imitate behaviors from raw video via context translation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1118–1125.
- [9] V. Giammarino, J. Queeney, L. C. Carstensen, M. E. Hasselmo, and I. C. Paschalidis, "Opportunities and challenges from using animal videos in reinforcement learning for navigation," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 9056–9061, 2023.
- [10] R. F. Prudencio, M. R. Maximo, and E. L. Colombini, "A survey on offline reinforcement learning: Taxonomy, review, and open problems," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [11] B. Zheng, S. Verma, J. Zhou, I. Tsang, and F. Chen, "Imitation learning: Progress, taxonomies and challenges," *arXiv preprint arXiv:2106.12177*, 2021.
- [12] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, *et al.*, "Deepmind control suite," *arXiv preprint arXiv:1801.00690*, 2018.
- [13] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.
- [14] Y. Gao, H. Xu, J. Lin, F. Yu, S. Levine, and T. Darrell, "Reinforcement learning from imperfect demonstrations," *arXiv preprint arXiv:1802.05313*, 2018.
- [15] A. Nair, A. Gupta, M. Dalal, and S. Levine, "AWAC: Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020.
- [16] H. Xu, X. Zhan, J. Li, and H. Yin, "Offline reinforcement learning with soft behavior regularization," *arXiv preprint arXiv:2110.07395*, 2021.
- [17] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *Proceedings of the 33rd International Conference on Machine Learning*, pp. 49–58, 2016.
- [18] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [19] N. Baram, O. Anschel, and S. Mannor, "Model-based adversarial imitation learning," *arXiv preprint arXiv:1612.02179*, 2016.
- [20] F. Torabi, G. Warnell, and P. Stone, "Generative adversarial imitation from observation," *arXiv preprint arXiv:1807.06158*, 2018.
- [21] W. Sun, A. Vemula, B. Boots, and D. Bagnell, "Provably efficient imitation learning from observation alone," in *Proceedings of the 36th International Conference on Machine Learning*, pp. 6036–6045, 2019.
- [22] R. Kidambi, J. Chang, and W. Sun, "Mobile: Model-based imitation learning from observation alone," in *Advances in Neural Information Processing Systems*, vol. 34, pp. 28 598–28 611, 2021.
- [23] V. Giammarino, J. Queeney, and I. C. Paschalidis, "Adversarial imitation learning from visual observations using latent information," *arXiv preprint arXiv:2309.17371*, 2023.
- [24] M. Liu, T. He, W. Zhang, S. Yan, and Z. Xu, "Visual imitation learning with patch rewards," *arXiv preprint arXiv:2302.00965*, 2023.
- [25] A. Tucker, A. Gleave, and S. Russell, "Inverse reinforcement learning for video games," *arXiv preprint arXiv:1810.10593*, 2018.
- [26] D. Brown, W. Goo, P. Nagarajan, and S. Niekum, "Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations," in *Proceedings of the 36th International Conference on Machine Learning*, pp. 783–792, 2019.
- [27] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.
- [28] B. Kim and J. Pineau, "Maximum mean discrepancy imitation learning," in *Robotics: Science and Systems*, 2013.
- [29] M. K. Hanawal, H. Liu, H. Zhu, and I. C. Paschalidis, "Learning policies for Markov decision processes from data," *IEEE Transactions on Automatic Control*, vol. 64, no. 6, pp. 2298–2309, June 2019.
- [30] A. Edwards, H. Sahni, Y. Schroecker, and C. Isbell, "Imitating latent policies from observation," in *Proceedings of the 36th International Conference on Machine Learning*, pp. 1755–1763, 2019.
- [31] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," *arXiv preprint arXiv:1805.01954*, 2018.
- [32] M. Laskey, S. Staszak, W. Y.-S. Hsieh, J. Mahler, F. T. Pokorny, A. D. Dragan, and K. Goldberg, "SHIV: Reducing supervisor burden in dagger using support vectors for efficient learning from demonstrations in high dimensional state spaces," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 462–469, 2016.
- [33] C.-A. Cheng, X. Yan, N. Wagener, and B. Boots, "Fast policy learning through imitation and reinforcement," *arXiv preprint arXiv:1805.10413*, 2018.
- [34] V. Giammarino, M. F. Dunne, K. N. Moore, M. E. Hasselmo, C. E. Stern, and I. C. Paschalidis, "Combining imitation and deep reinforcement learning to human-level performance on a virtual foraging task," *Adaptive Behavior*, 2023.
- [35] V. G. Goecks, G. M. Gremillion, V. J. Lawhern, J. Valasek, and N. R. Waytowich, "Integrating behavior cloning and reinforcement learning for improved performance in dense and sparse reward environments," *arXiv preprint arXiv:1910.04281*, 2019.
- [36] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine, "Learning to poke by poking: Experiential learning of intuitive physics," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [37] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Proceedings of the 4th Conference on Robot Learning*, pp. 1101–1112, 2020.
- [38] A. Rajeswaran, I. Mordatch, and V. Kumar, "A game theoretic framework for model based reinforcement learning," in *Proceedings of the 37th International Conference on Machine Learning*, pp. 7953–7963, 2020.
- [39] M. Janner, J. Fu, M. Zhang, and S. Levine, "When to trust your model: Model-based policy optimization," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [40] B. D. Ziebart, *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University, 2010.
- [41] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, *et al.*, "Soft actor-critic algorithms and applications," *arXiv preprint arXiv:1812.05905*, 2018.
- [42] E. C. Ozcan, V. Giammarino, J. Queeney, and I. Paschalidis, "A model-based approach for improving reinforcement learning efficiency leveraging expert observations," *arXiv preprint arXiv:2402.18836*, 2024.
- [43] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims, "Morel: Model-based offline reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 810–21 823, 2020.