

Regularized Soft Actor-Critic for Behavior Transfer Learning

1st Mingxi Tan
Ubisoft La Forge
Ubisoft
Chengdu, China
ming-xi.tan@ubisoft.com

2nd Andong Tian
Ubisoft La Forge
Ubisoft
Chengdu, China
an-dong.tian@ubisoft.com

3rd Ludovic Denoyer
Ubisoft La Forge
Ubisoft
Bordeaux, France
ludovic.denoyer@ubisoft.com

Abstract—Existing imitation learning methods mainly focus on making an agent effectively mimic a demonstrated behavior, but do not address the potential contradiction between the behavior style and the objective of a task. There is a general lack of efficient methods that allow an agent to partially imitate a demonstrated behavior to varying degrees, while completing the main objective of a task. In this paper we propose a method called Regularized Soft Actor-Critic which formulates the main task and the imitation task under the Constrained Markov Decision Process framework (CMDP). The main task is defined as the maximum entropy objective used in Soft Actor-Critic (SAC) and the imitation task is defined as a constraint. We evaluate our method on continuous control tasks relevant to video games applications.

Index Terms—CMDP, behavior style, video game

I. INTRODUCTION

In video games, non-playable characters (NPCs) and bots are usually expected to complete a mission while adopting specific behavior styles. Traditional ways require to build complex behavior trees, which can be a challenging and time-consuming work. We therefore propose to treat this problem as a behavior imitation learning problem for Reinforcement Learning. In recent years, many behavior imitation learning algorithms have been proposed [1]–[6]. These algorithms, such as behavioral cloning [6], inverse reinforcement learning [5], general adversarial imitation learning [4] and soft Q imitation learning [3], mainly focus on how to make an agent mimic the entire demonstrated behaviors rather than partially imitating the demonstrated behaviors to varying degrees. However, it is a common need in video games to have NPCs behaving in a specific style while still accomplishing their main tasks.

Although this need can be addressed with a multi-objective framework, where the first objective is to accomplish the main task and the second objective is to imitate the demonstrated behavior. An effective multi-objective framework requires explicit constant rewards for both the main task and the imitation task, which can be time-consuming to design as it usually requires a reward-shaping loop to find suitable values [7]–[9].

In this work we focus on making the agent partially imitate the demonstrated behavior to varying degrees while completing a main task by formulating the imitation learning under the Constrained Markov Decision Processes (CMDPs) framework. We propose to extend the Soft Actor-Critic (SAC) [15] to an

efficient partial imitation learning algorithm named Regularized Soft Actor-Critic (RSAC). We evaluate our algorithm on a Rabbids¹ video game, which provides a relevant context to the application of reinforcement learning to the design of NPCs behavior styles.

II. BACKGROUND

A. Markov Decision Process (MDPs)

The Markov decision processes (MDPs) defined by the tuple $(S; A; p_s; r)$ where S denotes the state space, A the action space, $p_s(s'|s, a)$ the transition distribution, $r(s_t, a_t)$ the reward function. The agent's policy π is a state-conditional distribution over actions, where $\pi(a_t|s_t)$ denotes the probability of taking action a_t in state s_t . We will use $\rho_\pi(s_t)$ and $\rho_\pi(s_t, a_t)$ to denote the state and state-action marginals of the trajectory distribution induced by $\pi(a_t|s_t)$. Giving any scalar function of actions and states $f : S \times A \rightarrow \mathbb{R}$, the expected discounted sum of f is defined as

$$J(\pi) := E_{(s_t, a_t) \sim \rho_\pi} \left[\sum_{t=0}^T \gamma^t f(s_t, a_t) \right], \quad (1)$$

where $\gamma \in [0, 1]$ is a discount factor used to reduce the relative importance of future rewards. In Reinforcement learning, the goal is to learn a policy $\pi(a_t|s_t)$ that maximizes the reward function over the whole trajectory

$$\pi^* = \underset{\pi \in \Pi}{\operatorname{argmax}} \left[E_{(s_t, a_t) \sim \rho_\pi} \sum_{t=0}^T r(s_t, a_t) \right], \quad (2)$$

where Π is the set of possible policies.

B. Constrained Markov Decision Process (CMDPs)

Constrained Markov decision processes (CMDPs) restrict the policies set Π to a typically smaller set Π_C by introducing the set of constraints constructed by cost functions $C_k : S \times A \rightarrow \mathbb{R}$ and their corresponding thresholds $d_k \in \mathbb{R}$, with $k = 1, \dots, K$. The goal of Reinforcement Learning under

¹https://en.wikipedia.org/wiki/Raving_Rabbids

CMDPs is to find a plausible policy by solving the constrained optimization problem

$$\begin{aligned} \pi^* &= \underset{\pi \in \Pi_C}{\operatorname{argmax}} \left[E_{(s_t, a_t) \sim \rho_\pi} \sum_{t=0}^T r(s_t, a_t) \right] \\ \text{s.t. } & J_{C_k}(\pi) \leq d_k, k = 1, \dots, K, \end{aligned} \quad (3)$$

where Π_C is the set of possible policies that satisfy all constraints.

III. OUR APPROACH

A. Problem Setting

In this paper, we focus on finding the policy that allows the agent to complete the task while having a similar behavior as the behavior shown in a demonstration clip. The similarity between the behavior of the agent and the behavior in the demo clip can be defined by the cross-entropy between the policy of the agent π and the policy of the demo clip π^d . π^d is trained by following the imitation learning algorithm of [3]. Then our problem can be formulated as a policy search in CMDPs. The objective can be defined as a maximum entropy objective following the prior work of [15] and the constraint can be defined by the cross-entropy between the policy of the agent and the policy of the behavior in the demo clip

$$\begin{aligned} \max_{\pi_{0:T}} & E_{(s_t, a_t) \sim \rho_\pi} \left[\sum_{t=0}^T [r(s_t, a_t) + \alpha H(\pi(\cdot|s_t))] \right] \\ \text{s.t. } & E_{(s_t, a_t) \sim \rho_\pi} [\log \pi^d(a_t|s_t)] \geq \overline{CE}, \end{aligned} \quad (4)$$

where $H(\pi(\cdot|s_t))$ is entropy of π , α is temperature parameter of entropy term, \overline{CE} is the required cross-entropy between the policy of the agent and the policy of the demo clip. For every time step, the (4) can be presented as a dual problem

$$\begin{aligned} \max_{\pi_{0:T}} & E_{(s_t, a_t) \sim \rho_\pi} [r(s_t, a_t) + \alpha H(\pi(\cdot|s_t))] = \\ \min_{\beta_t \geq 0} & \max_{\pi_{0:T}} E_{(s_t, a_t) \sim \rho_\pi} [r(s_t, a_t) + \alpha H(\pi(\cdot|s_t))] \\ & - \beta_t CE(\pi(\cdot|s_t), \pi^d(\cdot|s_t)) - \beta_t \overline{CE}, \end{aligned} \quad (5)$$

where β is a Lagrangian multiplier.

We use the strong duality, which holds because the objective function is convex, and the constraint (Cross-Entropy) is also a convex function in π_t . This dual objective can be considered as the extended version of the maximum entropy objective of [15], denoted as regularized maximum entropy objective (RMEO) with respect to the policy as

$$\begin{aligned} \pi_t^* &= \underset{\pi \in \Pi}{\operatorname{argmax}} E_{(s_t, a_t) \sim \rho_\pi} [r(s_t, a_t) + \alpha H(\pi(\cdot|s_t))] \\ & - \beta_t CE(\pi(\cdot|s_t), \pi^d(\cdot|s_t)) \end{aligned} \quad (6)$$

Note the optimal policy at time t is a function of the dual variable β_t . For every time step, we can firstly solve the optimal policy π_t at time t and then solve the optimal variable β_t as

$$\beta_t^* = \underset{\beta_t \geq 0}{\operatorname{argmax}} E_{(s_t, a_t) \sim \pi_t^*} [\beta_t \log \pi_t^d(a_t|s_t)] - \beta_t \overline{CE}. \quad (7)$$

To optimize the RMEO in (6), we first derive a tabular Q-iteration method (RMEO-Q), then present RMEO actor-critic (RMEO-AC), a practical deep reinforcement learning algorithm.

B. Regularized Maximum Entropy Reinforcement Learning

Following on a similar logic as [15], the regularized soft Q-function and regularized value function can be defined as

$$\begin{aligned} Q_{rs}^*(s_t, a_t) &= r_t + E_{(s_{t+1}, \dots) \sim \rho_\pi} \left[\sum_{l=1}^{\infty} \gamma^l (r_{t+l} + \right. \\ & \left. \alpha H(\pi^*(\cdot|s_{t+l})) - \beta CE(\pi(\cdot|s_{t+l}), \pi^d(\cdot|s_{t+l}))) \right]. \end{aligned} \quad (8)$$

$$V_{rs}^*(s_t) = \alpha \log \sum_{a'_t} \left(\pi^d(a'_t|s_t)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) \right) \right). \quad (9)$$

The optimal policy of (6) is given by

$$\begin{aligned} \pi^* &= \frac{\left(\pi^d(a'_t|s_t) \right)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) \right)}{\sum_{a'_t} \left(\pi^d(a'_t|s_t) \right)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) \right)} \\ &= \left(\pi^d(a'_t|s_t) \right)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) - \frac{1}{\alpha} V_{rs}^*(s_t) \right). \end{aligned} \quad (10)$$

Then we define the regularized soft Bellman equation for the regularized soft state-action value function Q as

$$Q_{rs}^*(s_t, a_t) = r_t + \gamma E_{(s_{t+1}) \sim p_s} [V_{rs}(s_{t+1})]. \quad (11)$$

Let $Q_{rs}^*(\cdot, \cdot)$ and $V_{rs}^*(\cdot)$ be bounded and assume that $\sum_{a'_t} \left(\pi^d(a'_t|s_t)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) \right) \right) \leq \infty$ exists, we can find a solution to (10) with a fixed-point iteration, which we call regularized soft Q-iteration as

$$\begin{aligned} Q_{rs}(s_t, a_t) &\leftarrow r_t + \gamma E_{(s_{t+1}) \sim p_s} [V_{rs}(s_{t+1})], \forall s_t, a_t, \\ V_{rs}(s_t) &\leftarrow \alpha \log \sum_{a'_t} \left(\pi^d(a'_t|s_t)^{\frac{\beta}{\alpha}} \exp \left(\frac{1}{\alpha} Q_{rs}^\pi(s_t, a'_t) \right) \right), \forall s_t, \end{aligned} \quad (12)$$

converge to the optimal Q_{rs}^* and V_{rs}^* .

To make RMEO-Q effectively imitate the behavior of the demo clip, we use two replay buffers D^d and D^{new} to store the demonstration experiences and new experiences and use them to update $Q_{rs}(s_t, a_t)$ separately. We firstly use the new experiences to update Q_{rs} , allowing the agent to learn to complete the task by setting β to zero. Secondly, we continually use the same new experiences to solve β by (7). If the cross-entropy of π and π^d is bigger than the required value (\overline{CE}), β will be augmented. Otherwise, β will be gradually decreased to zero. Then we use this β and the demonstration experiences to update Q_{rs} , allowing the agent to learn to imitate the demonstrated behavior. A large β encourages the agent to learn to mimic the demonstrated behavior, and a small β encourages the agent to learn to complete the task. Thus β works as an automatic regulator that can help find the policy to satisfy both the task completion and the behavior imitation.

C. Regularized Maximum Entropy Objective Actor-Critic

To practically implement our method, we propose the RMEO actor-critic (RMEO-AC), which uses neural networks as function approximators for both the regularized Q-function and policy and optimizes both networks with stochastic gradient descent. We parameterize the regularized Q-function and policy by $Q_\theta(s_t, a_t)$ and $\pi_\phi(a_t|s_t)$.

Then we can replace the Q-iteration with Q-learning and train θ to minimize

$$J_Q(\theta) = E_{(s_t, a_t) \sim D} \left[\frac{1}{2} (Q_\theta(s_t, a_t) - r(s_t, a_t) - \gamma E_{(s_{t+1}) \sim p_s} [\bar{V}_{rs}(s_{t+1})])^2 \right] \quad (13)$$

$$V_{rs}(s_{t+1}) = Q_{\bar{\theta}}(s_{t+1}, a_{t+1}) - \alpha \log(\pi_\phi(a_{t+1}|s_{t+1})) + \beta \log(\pi^d(a_{t+1}|s_{t+1})),$$

where $Q_{\bar{\theta}}$ is target function for providing relatively stable target values. We can also find β by minimizing

$$J(\beta) = E_{(s_t, a_t) \sim D^{new}} [\beta \pi^d(a_t|s_t) - \beta \overline{CE}], \quad (14)$$

The policy parameters can be learned by directly minimizing the KL-divergence

$$J_\pi(\phi) = E_{(s_t) \sim D} \left[D_{KL} \left(\pi_\phi(\cdot|s_t) \parallel \frac{(\pi^d(\cdot|s_t))^{\frac{\beta_t}{\alpha}} \exp(\frac{1}{\alpha} Q_{rs}^\pi(s_t, \cdot))}{Z_\theta(s_t)} \right) \right], \quad (15)$$

Since the partition function $Z_\theta(s_t)$ normalizes the distribution, it does not contribute to the gradient and can be ignored. By using the reparameterization trick $a_t = f_\phi(\epsilon_t; s_t)$, we can rewrite the objective above as

$$J_\pi(\phi) = E_{(s_t) \sim D, \epsilon \sim N} [\alpha \log(f_\phi(\epsilon_t; s_t)|s_t) - \beta \log(\pi^d(f_\phi(\epsilon_t; s_t)|s_t)) - Q_\theta(s_t, f_\phi(\epsilon_t; s_t))] \quad (16)$$

The gradient of (13), (14), (15) with respect to the corresponding parameters are

$$\nabla_\theta J_Q(\theta) = \nabla_\theta Q_\theta(s_t, a_t) (Q_\theta(s_t, a_t) - r(s_t, a_t) - \gamma \bar{V}(s_{t+1})), \quad (17)$$

$$\nabla_\beta J(\beta) = \log \pi^d(a_t|s_t) - \overline{CE}, \quad (18)$$

$$\begin{aligned} \nabla_\phi J_\pi(\phi) = & \nabla_\phi \alpha \log(\pi_\phi(a_t|s_t)) + \nabla_\phi f_\phi(\epsilon_t; s_t) (\nabla_{a_t} \alpha \log(\pi_\phi(a_t|s_t)) \\ & - \nabla_{a_t} \beta \log(\pi^d(a_t|s_t)) - \nabla_{a_t} Q_\theta(s_t, a_t)), \end{aligned} \quad (19)$$

The final algorithm is listed in Algorithm 1.

Algorithm 1 Regularized Soft Actor-Critic

Input: $\theta, \phi, \beta, \pi^d, D^d$

```

1:  $\bar{\theta} \leftarrow \theta$ 
2:  $D^{new} \leftarrow \emptyset$ 
3: for each iteration do
4:   for each environment step do
5:      $a_t \sim \pi_\phi(a_t|s_t)$ 
6:      $s_{t+1} \sim P(s_{t+1}|s_t, a_t)$ 
7:      $D^{new} \cup \{(s_t, a_t, r(s_t, a_t), s_{t+1})\}$ 
8:   end for
9:   for each gradient step do
10:    Sample a mini-batch from  $D^{new}$ 
11:    First set  $\beta = 0$ 
12:     $\theta \leftarrow \theta - \lambda_Q \nabla_\theta J_Q(\theta)$ 
13:     $\phi \leftarrow \phi - \lambda_\pi \nabla_\phi J_\pi(\phi)$ 
14:     $\bar{\theta} \leftarrow \tau \theta + (1 - \tau) \bar{\theta}$ 
15:    Then
16:     $\beta \leftarrow \beta - \nabla_\beta J(\beta)$ 
17:    Sample a mini-batch from  $D^d$ 
18:     $\theta \leftarrow \theta - \lambda_Q \nabla_\theta J_Q(\theta)$ 
19:     $\phi \leftarrow \phi - \lambda_\pi \nabla_\phi J_\pi(\phi)$ 
20:     $\bar{\theta} \leftarrow \tau \theta + (1 - \tau) \bar{\theta}$ 
21:  end for
22: end for
Output:  $\theta, \phi$ 

```

IV. EXPERIMENTS

To evaluate the proposed approach, we first define the main task and then define the specific behavior that should be imitated. The agent we train is car number 4 at the bottom-right corner of the map and its main task is to hit car number 1 at the top-left corner as shown in Fig. 1 (left). The demonstrated behavior is car number 4 navigating in circles and backwards in the lower half of the map as shown in Fig. 1 (right). The goal of the agent is to hit car number 1 while adopting the demonstrated behavior style. This task is challenging because if the agent mimics the behavior of the entire demonstration, it cannot hit car number 1. It requires car 4 to partially mimic the behavior shown while hitting car 1.

A. Baseline

We modify the Soft Q Imitation Learning algorithm of [3] by adding a constant main task reward +1.0. We perform a grid search on 11 imitation reward values: 0.0, +0.1, +0.2,



(a) Main task

(b) Behavior of demo clip

Fig. 1. Main task and behavior of demo clip

TABLE I
RESULTS OF BASELINE MODELS.

Imitation reward	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
CE	1.72	1.48	1.44	1.44	1.44	1.44	1.44	1.43	1.43	1.43	1.43
Mission reward	0.96	0.92	-0.24	-0.26	-0.31	-0.35	-0.38	-0.42	-0.41	-0.42	-0.43

+0.3, +0.5, +0.6, +0.7, +0.8, +0.9, +1.0 in the environment of a Rabbits video game. Since the reward for the main task is sparse, we use an auxiliary reward to guide the agent to complete the main task: when the agent is in the upper half of the map, if it is far from car 1, it will obtain a constant reward = -0.01. For each reward setting, we train the model for 1M steps using the same random seed, and then evaluate the model for 10 epochs. The cross-entropy and reward for each setting are shown in Table I. The first row is different imitation reward applied in experiment, the second row is the corresponding cross-entropy of π and π^d and the third row is the corresponding task rewards obtained. All values are the average over 10 episodes.

From the Table I, we make several important observations:

1) Except for the model trained with imitation reward=0.0 and reward=+0.1, the cross-entropy of all other models is in a very narrow range (second row), which means it is difficult to tune the constant reward of the imitation task to obtain an agent which behavior has a certain degree of similarity to the demonstrated behavior. 2) Out of 11 imitation reward settings, only one setting (imitation reward=+0.1) successfully trained a policy that can complete the main task in the demonstrated behavior style, which means the reward shaping is very challenging and time-consuming.

B. Experiments settings

We train three agents to complete the main task while partially imitating the demonstrated behavior to 3 different degrees: cross-entropy of π and π^d smaller than 1.5, 1.55 and 1.6. When the agent completes the main task, it will get a constant reward = +1.0. We apply the same auxiliary reward and train the same steps as before. For every setting, we train the model with 5 random seeds, we halt the training process every 20K steps and evaluate the model for 10 epochs. The results are shown in Fig. 2.

From Fig. 2, we can make several important observations.

1) Our approach can satisfy the requirements of different constraints (Fig. 2, a) while completing the main task (Fig. 2, b). 2) The tighter constraint given to the agent is, the lower is the total reward, as it influences the agent to imitate the demonstrated behavior, which will move it away from car 1 to get more auxiliary reward with the value of -0.01 (Fig. 2, b). 3) The tighter constraint needs more time to adjust the value of β (Fig. 2, c).

These results show that our approach is efficient because it can automatically trade off task completion and behavioral imitation to find a policy that satisfies both.

REFERENCES

[1] D. A. Pomerleau, "Efficient training of artificial neural networks for autonomous navigation," Neural Computation, vol. 3, pp. 88–97, 1991.

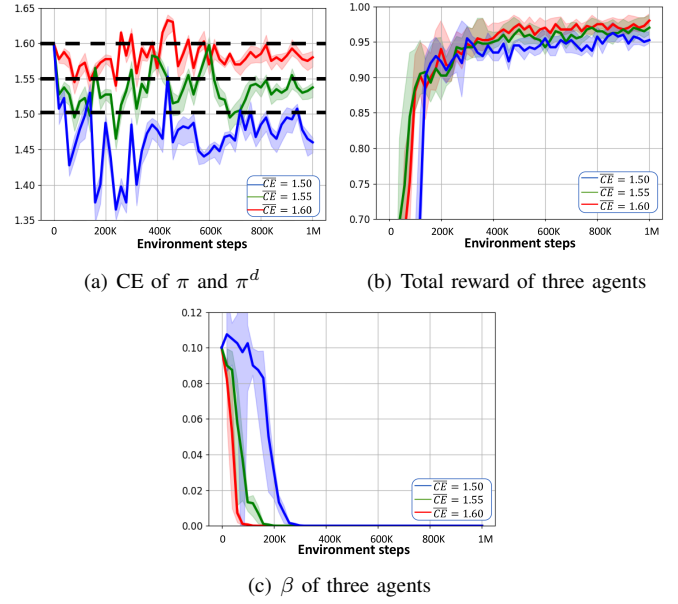


Fig. 2. Results of our approach under three different constraints (\overline{CE}). a) CE of π and π^d . b) Total reward of three agents. c) Three β values. All curves show the average over 5 random seeds and envelopes show the standard error around the mean.

[2] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," In the International Conference on Machine Learning (ICML), pp. 663–670, 2000.

[3] S. Reddy, A. D. Dragan and S. Levine, "SQIL: Imitation Learning via Reinforcement Learning with Sparse Rewards," In the International Conference on Learning Representations (ICLR), 2020.

[4] J. Fu, K. Luo and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," arXiv preprint arXiv:1710.11248, 2017.

[5] J. Ho and S. Ermon, "Generative adversarial imitation learning," Advances in Neural Information Processing Systems, pp. 4565–4573, 2016.

[6] S. Ross, G. J. Gordon and J. A. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," In Proceedings of the fourteenth international conference on artificial intelligence and statistics, pp. 627–635, 2011.

[7] H. Mossalam, Y. M. Assael, D. M. Roijers and S. Whiteson, "Multi-objective deep reinforcement learning," Computer Research Repository (CoRR), <https://arxiv.org/abs/1610.02707>, 2016.

[8] A. Abels, D. M. Roijers, T. Lenaerts, A. Nowé and D. Steckelmacher, "Dynamic weights in multi-objective deep reinforcement learning," In Proceedings of the 36th International Conference on Machine Learning (ICML), pp. TAB, 2019.

[9] R. Yang, X. Sun, and K. Narasimhan, "A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation," In Neural Information Processing Systems (NeurIPS), pp. 14610–14621, 2019.

[10] K. Kurach, M. Lucic, X. Zhai, M. Michalski and S. Gelly, "DThe gan landscape: Losses, architectures, regularization, and normalization," arXiv preprint arXiv:1807.04720, 2018.

[11] E. Altman, "Constrained Markov decision processes," CRC Press, 1999.

[12] T. Yang, J. Rosca, K. Narasimhan and P. J. Ramadge, "Projection-based constrained policy optimization," arXiv preprint arXiv:2010.03152, 2020.

[13] Y. Zhang, Q. Vuong and K. W. Ross, "First order constrained optimization in policy space," arXiv preprint arXiv:2002.06506, 2020.

[14] M. Turchetta, A. Kolobov, S. Shah, A. Krause and A. Agarwal, "Safe reinforcement learning via curriculum induction," arXiv preprint arXiv:2006.12136, 2020.

[15] T. Haarnoja, H. Tang, P. Abbeel and S. Levine, "Reinforcement learning with deep energy-based policies," Computer Research Repository (CoRR), <https://arxiv.org/abs/1702.08165>, 2017.