

ADVERSARIAL ADVANTAGE ACTOR-CRITIC MODEL FOR TASK-COMPLETION DIALOGUE POLICY LEARNING

Baolin Peng^{*} Xiujun Li[†] Jianfeng Gao[†] Jingjing Liu[†] Yun-Nung Chen[‡] Kam-Fai Wong^{*}

[†]Microsoft Research, Redmond, WA, USA

^{*}The Chinese University of Hong Kong, Hong Kong

[‡]National Taiwan University, Taipei, Taiwan

ABSTRACT

This paper presents a new method — adversarial advantage actor-critic (Adversarial A2C), which significantly improves the efficiency of dialogue policy learning in task-completion dialogue systems. Inspired by generative adversarial networks (GAN), we train a discriminator to differentiate responses/actions generated by dialogue agents from responses/actions by experts. Then, we incorporate the discriminator as another critic into the advantage actor-critic (A2C) framework, to encourage the dialogue agent to explore state-action within the regions where the agent takes actions similar to those of the experts. Experimental results in a movie-ticket booking domain show that the proposed Adversarial A2C can accelerate policy exploration efficiently.

Index Terms— task-completion dialogue, reward function, adversarial learning, policy learning, reinforcement learning

1. INTRODUCTION

There has been growing interest in exploiting reinforcement learning (RL) for policy learning in task-oriented dialogue systems [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. One of the biggest challenges in these approaches is the *reward sparsity* issue. Dialogue policy learning for complex tasks, such as movie-ticket booking and travel planning, requires exploration in a large state-action space, and it often takes many conversation turns between the user and the agent to fulfill a task, leading to a long trajectory. Thus, the reward signals (usually provided by users at the end of a conversation) are often delayed and sparse.

To deal with reward sparsity, different approaches have been proposed recently, with promising empirical results. One approach is to leverage prior knowledge learned from expert-generated (or human-human) dialogue. For example, instead of learning a dialogue policy from scratch, we construct an initial policy learned from human-human dialogues, via imitation learning or hand-crafted rules. Prior work showed that a pre-trained supervised policy or a weak rule-based policy can significantly improve the efficiency

of exploration [4, 11]. Another approach is to introduce heuristics, often in the form of the intrinsic reward to guide the exploration [12, 13, 14, 15]. While the extrinsic reward (e.g., feedback provided by users at the end of a conversation) could be sparse, it is possible to get intrinsic reward after each action in order to guide the agent to explore the region more effectively. For example, VIME maximizes information gain about the agents belief of environment dynamics [14]. It adds an intrinsic reward bonus to the reward function, which quantifies the agents *surprise* to encourage the agent to explore the regions that are relatively unexplored. BBQN encourages the agent to explore those state-action regions where the agent is relatively uncertain in action selection [11]. UNREAL converts the training signals from three auxiliary tasks as intrinsic rewards, which significantly improved the learning speed and the robustness of the agent [15].

In this paper, we present a new method that combines the strength of the two approaches mentioned above. Similar to the first approach, we also leverage expert-generated dialogues as prior knowledge. However, instead of constructing an initial dialogue policy using prior knowledge, we, inspired by generative adversarial networks (GAN) [16], train a discriminator to differentiate the responses (or actions) generated by dialogue agents from those by human experts. Then, we use the output of the discriminator as intrinsic reward to encourage the dialogue agent to explore state-action regions in which the agent takes actions similar to what human experts do. Specifically, we incorporate the discriminator as another critic into the advantage actor-critic (A2C) framework, resulting in a new model, called adversarial advantage actor-critic (Adversarial A2C). The modeling assumption behind our method is that the expert policies (embedded in the expert-generated dialogues) are reasonably good, thus the agent-selected actions, which are more similar to expert-selected ones, lead more often to successful dialogues with positive rewards. In a word, we remedy the reward sparse problem on two fronts, by leveraging human-human dialogues as prior knowledge and by introducing intrinsic rewards. Experiments in a movie-ticket booking domain show that the proposed Adversarial A2C model can




Fig. 1. Illustration of a task-completion dialogue system.

significantly improve dialogue policy learning in terms of both effectiveness and efficiency.

2. METHODOLOGY

Figure 1 illustrates a typical task-completion dialogue system that contains three main components: language understanding (LU) that converts natural language to system-readable semantic frames, natural language generation (NLG) that converts system actions to natural language, and a dialogue manager (DM). The dialogue manager controls state tracking and policy learning, where dialogue policy learning can be regarded as a sequential decision process. The system will learn to select the best response action at each step, by maximizing the long-term objective associated with a reward function.

This paper focuses on dialogue policy learning (the bottom-right part of Figure 1), where the input of the policy learner is the dialogue state representation s that consists of the latest user action (e.g., `request_moviename(genre=action, date=today)`), the last agent action (e.g. `request_location`), history dialogue turns, and available database results. The learned dialogue policy then helps the agent decide what action a to take in each turn of the conversation, in order to maximize the future cumulative reward.

As aforementioned, dialogue policy optimization can be formulated as a sequential decision problem to maximize the long term objective associated with a reward function. The advantage actor-critic (A2C) method has achieved superior performance on solving sequential decision problems [17, 18, 19]. Su et al. applied the actor-critic model to dialogue policy optimization and proved its superiority on convergence to other methods such as deep Q-networks [20]. Similarly, we employ an actor-critic approach to learn dialogue policy in our model. In addition, inspired by GAN [16] (using a discriminator to guide the training of generative models), we form a minimax game between a generator (an actor that selects actions in our scenario) and a discriminator, to judge whether an action is performed by the expert or the actor. The discriminator can be regarded as another critic and serves

as a heuristic intrinsic reward function to guide the actor towards expert-like regions. Another related topic is inverse reinforcement learning [21], which is to recover the reward function from expert demonstrations, samples of the trajectories executed by experts [22]. Ho and Ermon also drew a connection between inverse reinforcement learning and generative adversarial networks to learn the reward function in the GAN framework [23]. Compared to their work that focused on learning the extrinsic reward, in this paper, we use intrinsic reward to speech up the training.

2.1. Advantage Actor-Critic for Dialogue Policy Learning

The training objective of policy-based approaches is to find a policy π that maximizes the expected reward R (minimizes the loss J) over all possible dialogue trajectories. The expected reward is defined as $R = \sum_{t=0}^{T-1} \gamma^t r_t$ over a dialogue with the length T , where r_t is the reward at time stamp t , and γ is the discount factor. The policy π is a parametrized probabilistic mapping function between the state space and the action space:

$$\pi_{\theta}(a | s) = P(A_t = a | S_t = s; \theta), \quad (1)$$

where θ represents the parameters learned by policy gradient algorithms [24]. Given the objective function, the gradients of the parameters are computed as

$$\nabla_{\theta} J(\theta) = \mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(a | s) Q^{\pi_{\theta}}(s, a)] \quad (2)$$

where $Q^{\pi_{\theta}}(s, a)$ is the long-term reward value. However, the gradients usually have high variance, which makes the learning task more challenging. A baseline function $B(s)$ is usually employed to reduce the variance, while keeping the estimated gradient unchanged [25]. Here we can simply choose the state value function as a baseline $B(s) = V^{\pi_{\theta}}(s)$. With this strategy, we can rewrite (2) using the advantage function $A^{\pi_{\theta}}(s, a)$:

$$\nabla_{\theta} J(\theta) = \mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(a | s) A^{\pi_{\theta}}(s, a)], \quad (3)$$

$$A^{\pi_{\theta}}(s, a) = Q^{\pi_{\theta}}(s, a) - V^{\pi_{\theta}}(s). \quad (4)$$

However, in this setting, there are two functions and parameters that need to be learned. In order to reduce the number of required parameters and improve stability, temporal difference (TD) error is employed as an unbiased estimate of the advantage function,

$$\delta^{\pi_{\theta}} = r + \gamma V^{\pi_{\theta}}(s') - V^{\pi_{\theta}}(s), \quad (5)$$

In this way, the policy gradient with the TD error can be computed as

$$\nabla_{\theta} J(\theta) = \mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(a | s) \delta^{\pi_{\theta}}]. \quad (6)$$

The policy network π_{θ} is termed as the *actor* to yield a dialogue system action, and the advantage function $A^{\pi_{\theta}}$ is the *critic*, indicating “good” or “bad” for executing an action given a state. The classic A2C architecture is shown at the bottom part of Figure 2 without discriminator.

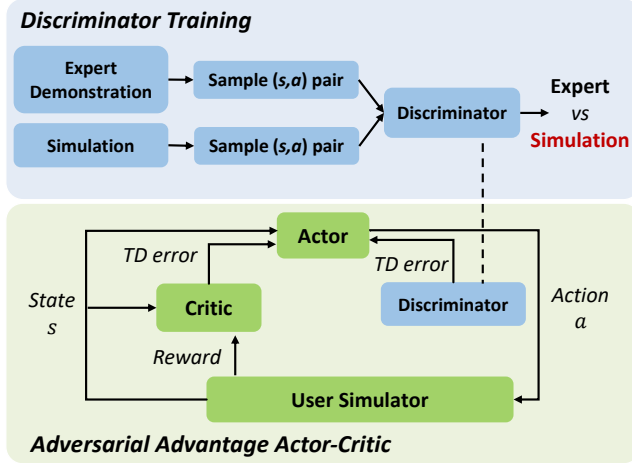


Fig. 2. Illustration of the proposed adversarial advantage actor-critic for dialogue policy learning.

2.2. Adversarial Model for Dialogue Policy Learning

GAN is a minimax competing game between a generator and a discriminator. In our scenario, the actor π can be viewed as a generator G , which aims to generate actions that can purposefully confuse a discriminator D . The discriminator D is expected to identify a state-action pair (s, a) as either an expert demonstration or a simulation experience. When D cannot distinguish actions generated from the actor π and those from the experts, we believe that π has been improved from the previous state. Moreover, D can be viewed as a reward function extracted from the experts' trajectories. Figure 2 shows the discriminator training procedure using adversarial learning.

The training objective is to find a saddle point (π, D) of

$$\mathbb{E}_{\pi}[\log(D(s, a))] + \mathbb{E}_{\text{Demo}}[\log(1 - D(s, a))]. \quad (7)$$

More specifically, let θ_D denote the parameters of the discriminator D . The training objective of D is simply to maximize the probability of classifying each state-action pair (s, a) correctly:

$$\begin{aligned} \min_{\theta_D} \mathcal{L}_D &= -\mathbb{E}_{(s,a) \sim \text{Simu}} \log D(s, a; \theta_D) \\ &\quad -\mathbb{E}_{(s,a) \sim \text{Demo}} \log(1 - D(s, a; \theta_D)) \end{aligned} \quad (8)$$

where *Simu* and *Demo* represent simulation experience and expert demonstration, respectively. As thus, the actor π can be improved using actor-critic, with $-\log(1 - D(s, a))$ as the reward function. The updated gradients can then be reformed as:

$$\nabla_{\theta} J(\theta) = \mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(a | s) A_{\text{GAN}}^{\pi_{\theta}}(s, a)]. \quad (9)$$

Similarly, we use TD error as an unbiased estimation of the advantage function:

$$\delta_{\text{GAN}}^{\pi_{\theta}} = r_{\text{GAN}} + \gamma V_{\text{GAN}}^{\pi_{\theta}}(s') - V_{\text{GAN}}^{\pi_{\theta}}(s). \quad (10)$$

Algorithm 1 Adversarial Advantage Actor-Critic Model

- 1: **Input:** Expert demonstrations *Demo*, initialize actor π , discriminator D and two value functions $V^{\pi_{\theta}}(s)$, $V_{\text{GAN}}^{\pi_{\theta}}(s)$
- 2: **for** $i=1:N$ **do**
- 3: Restart the dialogue simulator, get state representation s , initialize transition tuple $\text{buffer} = []$
- 4: **while** s is not a terminal state **do**
- 5: Perform the action a_t according to the actor $\pi(a_t | s; \theta)$
- 6: Receive the reward r_t and switch to a new state s'
- 7: Store (s, a_t, r_t, s') to the transition tuple buffer
- 8: $s := s'$
- 9: **end while**
- 10: Train the actor π with gradients (6)
- 11: Train value function $V^{\pi_{\theta}}(s)$ by minimizing the TD error (5)
- 12: Sample state action pairs (s, a) from expert demonstration *Demo*
- 13: Train the actor π with gradients (9)
- 14: Update the reward with $-\log(1 - D(s, a))$ and train value function $V_{\text{GAN}}^{\pi_{\theta}}(s)$ by minimizing the TD error (10)
- 15: Update the discriminator parameters (8)
- 16: **end for**

2.3. Adversarial Advantage Actor-Critic

Furthermore, training dialogue policy with a stand-alone adversarial model can be impractical, due to the high dimensionality of its state-and-action space. To address this issue, we propose the adversarial advantage actor-critic (Adversarial A2C) method as depicted in Figure 2, which combines A2C with a reward function learned from an adversarial model that serves as another additional critic for the actor π . There are several ways to combine two critics, such as linear combination of two reward functions or alternately optimizing with each reward function. In our experiments, we use alternating optimization. Algorithm 1 outlines the full procedure of training the Adversarial A2C model. The goal is to encourage the actor to select better actions guided by a discriminator, in order to improve the efficiency and effectiveness of the exploration.

3. EXPERIMENTS

To verify the performance of the proposed model, we evaluated it in a task-completion dialogue system for movie-ticket booking. In this system, the agent will gather information from users through conversations and eventually book the movie tickets for them. The environment then judges a binary outcome (success or failure) at the end of each conversation, based on: 1) whether a movie ticket is booked, and 2) whether the booked ticket satisfies the constraints requested by the user.

3.1. Experimental Setup

The dataset used in our experiment is raw conversational data collected via Amazon Mechanical Turk, annotated by domain

experts [5]. This single-domain movie-ticket booking dataset contains 11 dialogue acts and 29 slots, including *informable* slots (users can use these to narrow down the search), and *requestable* slots (where users can ask the agent for more information). There are in total 280 labeled dialogues, with an average length of 11 turns.

In order to perform end-to-end training for the dialogue system, a user simulator is required to interact with the system in a natural way. We adopted a publicly available, user-agenda-based simulator in our experiments [26]. In a task-completion dialogue setting, the user simulator first generates a user goal, and the dialogue agent tries to help the user accomplish that goal in the course of the conversation, without explicitly knowing the user goal. A user goal normally consists of two parts: *inform_slots* representing slot-value pairs that serve as constraints from the user, and *request_slots* representing slots whose value the user has no information about, but wants to get information from the agent through the conversation. In our experiment, the user goals were generated from labeled conversational data.

3.2. Implementation

In Figure 2, the expert demonstrations can be collected from either human or pre-trained agent. In our experiment, we collected 50 successful dialogues from a pre-trained agent. The discriminator is a binary classifier of a single-layer neural network with 80 hidden units. For the actor, we use a single-layer neural network with a hidden size of 80, pre-trained with rule-based examples in order to give acceptable initialization. During the Adversarial A2C model training, two critics (the critic and the discriminator in Figure 2) are applied alternatively, where their value functions are single-layer neural networks with 80 hidden units. All parameters are optimized with *RMSPprop*. During training, the model is updated at the end of each dialogue episode.

3.3. Evaluation Results

In the movie-ticket booking task, we benchmark the proposed Adversarial A2C model against three baseline models on three metrics: success rate, average rewards, and the average number of turns per dialogue session.

- *Rule Agent* is a handcrafted rule-based policy that informs and requests a hand-picked subset of necessary slots.
- *A2C Agent* is trained with a pre-defined reward function and a standard advantage actor-critic algorithm.
- *BBQN-Map Agent* is the best agent among a set of BBQN variants (including BBQN-VIME) and DQN variants, which has demonstrated great efficiency for policy exploration in task-completion dialogue systems [11].




Fig. 3. Learning curves of dialogue policies.

Agents	Success Rate	Reward	Turn
Rule	41.34	0.26	16.00
A2C	81.24	5.08	15.43
BBQN-Map	81.56	5.00	18.75
Adversarial A2C	87.52	5.93	13.52

Table 1. Final agent performance on 5K simulated dialogues.

Figure 3 shows the learning curves of all these dialogue agents mentioned above, and Table 1 shows the evaluation performance of each agent, averaged over 5 runs. The learning curves in Figure 3 shows that Adversarial A2C agent can learn much faster with better exploration capability. The learning curve is also more stable compared with others. Table 1 suggests that the Adversarial A2C agent can yield better dialogue policies than other approaches, in terms of success rate, average reward, and average number of turns per dialogue.

4. CONCLUSIONS

This paper presents an adversarial advantage actor-critic model, which can explore policy learning in task-completion dialogue systems with great efficiency. The proposed model learns a discriminator from expert demonstrations and online experience, and then the learned discriminator serves as an additional critic to guide policy learning. Our experiments in a movie-ticket booking domain demonstrate the superiority and efficiency of the proposed model in policy learning, compared with state-of-the-art approaches. The promising results suggest several interesting future directions: 1) employing variance-reducing methods to stabilize the gradient calculation, in order to address the high variance issue in policy gradient estimation, 2) applying the model to more complicated dialogue tasks, such as composite task-completion dialogues [8], and 3) extending this work to other deep reinforcement learning benchmark tasks and other domains.

5. REFERENCES

- [1] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams, “POMDP-based statistical spoken dialog systems: A review,” *Proceedings of the IEEE*, vol. 101, no. 5, pp. 1160–1179, 2013.
- [2] Mehdi Fatemi, Layla El Asri, Hannes Schulz, Jing He, and Kaheer Suleman, “Policy networks with two-stage training for dialogue systems,” *arXiv preprint arXiv:1606.03152*, 2016.
- [3] Tiancheng Zhao and Maxine Eskenazi, “Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning,” in *Proceedings of SIGDIAL*, 2016.
- [4] Pei-Hao Su, Milica Gasic, Nikola Mrksic, Lina Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young, “Continuously learning neural dialogue management,” *arXiv preprint arXiv:1606.02689*, 2016.
- [5] Xuijun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz, “End-to-end task-completion neural dialogue systems,” in *Proceedings of IJCNLP*, 2017.
- [6] Jason D Williams, Kavosh Asadi, and Geoffrey Zweig, “Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning,” in *Proceedings of ACL*, 2017.
- [7] Bhuwan Dhingra, Lihong Li, Xuijun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng, “Towards end-to-end reinforcement learning of dialogue agents for information access,” in *Proceedings of ACL*, 2017, pp. 484–495.
- [8] Baolin Peng, Xuijun Li, Lihong Li, Jianfeng Gao, Asli Celikyilmaz, Sungjin Lee, and Kam-Fai Wong, “Composite task-completion dialogue policy learning via hierarchical deep reinforcement learning,” in *EMNLP*, 2017, pp. 2221–2230.
- [9] Bing Liu and Ian Lane, “Iterative policy learning in end-to-end trainable task-oriented neural dialog models,” *arXiv preprint arXiv:1709.06136*, 2017.
- [10] Baolin Peng, Xuijun Li, Jianfeng Gao, Jingjing Liu, and Kam-Fai Wong, “Integrating planning for task-completion dialogue policy learning,” *CoRR*, vol. abs/1801.06176, 2018.
- [11] Zachary C Lipton, Jianfeng Gao, Lihong Li, Xuijun Li, Faisal Ahmed, and Li Deng, “Efficient exploration for dialogue policy learning with BBQ networks & replay buffer spiking,” *arXiv preprint arXiv:1608.05081*, 2016.
- [12] Nuttapon Chentanez, Andrew G Barto, and Satinder P Singh, “Intrinsically motivated reinforcement learning,” in *NIPS*, 2005, pp. 1281–1288.
- [13] Shakir Mohamed and Danilo Jimenez Rezende, “Variational information maximisation for intrinsically motivated reinforcement learning,” in *NIPS*, 2015, pp. 2125–2133.
- [14] Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel, “Vime: Variational information maximizing exploration,” in *NIPS*, 2016, pp. 1109–1117.
- [15] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu, “Reinforcement learning with unsupervised auxiliary tasks,” *arXiv preprint arXiv:1611.05397*, 2016.
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *NIPS*, 2014, pp. 2672–2680.
- [17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [18] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al., “Mastering the game of go with deep neural networks and tree search,” *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [19] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al., “Mastering the game of Go without human knowledge,” *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [20] Pei-Hao Su, David Vandyke, Milica Gasic, Nikola Mrksic, Tsung-Hsien Wen, and Steve Young, “Reward shaping with recurrent neural networks for speeding up on-line policy learning in spoken dialogue systems,” in *Proceedings of SIGDIAL*, 2015.
- [21] Pieter Abbeel and Andrew Y Ng, “Apprenticeship learning via inverse reinforcement learning,” in *ICML*. ACM, 2004, p. 1.
- [22] Andrew Y Ng, Stuart J Russell, et al., “Algorithms for inverse reinforcement learning,” in *ICML*, 2000, pp. 663–670.
- [23] Jonathan Ho and Stefano Ermon, “Generative adversarial imitation learning,” in *NIPS*, pp. 4565–4573, 2016.
- [24] Ronald J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine Learning*, vol. 8, pp. 229–256, 1992.
- [25] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu, “Asynchronous methods for deep reinforcement learning,” in *ICML*, 2016, pp. 1928–1937.
- [26] Xuijun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen, “A user simulator for task-completion dialogues,” *arXiv preprint arXiv:1612.05688*, 2016.