

# Triple-GAIL: A Multi-Modal Imitation Learning Framework with Generative Adversarial Nets

Cong Fei<sup>1,2†</sup>, Bin Wang<sup>1†\*</sup>, Yuzheng Zhuang<sup>1†</sup>, Zongzhang Zhang<sup>3</sup>, Jianye Hao<sup>1</sup>,  
Hongbo Zhang<sup>1</sup>, Xuewu Ji<sup>2</sup> and Wulong Liu<sup>1</sup>

<sup>1</sup>Huawei Noah's Ark Lab

<sup>2</sup>Tsinghua University

<sup>3</sup>Nanjing University

feic18@mails.tsinghua.edu.cn, {wangbin158, zhuangyuzheng}@huawei.com,  
zhangzongzhang@gmail.com, {haojianye, zhanghongbo888}@huawei.com, jixw@mail.tsinghua.edu.cn,  
liuwulong@huawei.com

## Abstract

Generative adversarial imitation learning (GAIL) has shown promising results by taking advantage of generative adversarial nets, especially in the field of robot learning. However, the requirement of isolated single modal demonstrations limits the scalability of the approach to real world scenarios such as autonomous vehicles' demand for a proper understanding of human drivers' behavior. In this paper, we propose a novel multi-modal GAIL framework, named Triple-GAIL, that is able to learn skill selection and imitation jointly from both expert demonstrations and continuously generated experiences with data augmentation purpose by introducing an auxiliary skill selector. We provide theoretical guarantees on the convergence to optima for both of the generator and the selector respectively. Experiments on real driver trajectories and real-time strategy game datasets demonstrate that Triple-GAIL can better fit multi-modal behaviors close to the demonstrators and outperforms state-of-the-art methods.

## 1 Introduction

Imitation learning aims to mimic expert behavior directly from human demonstrations, without designing explicit reward signal as reinforcement learning (RL) [Abbeel and Ng, 2004; Fang *et al.*, 2019], and has made achievements in a variety of tasks. Recent work in imitation learning, especially generative adversarial imitation learning (GAIL) [Ho and Ermon, 2016], optimizes a policy directly from expert demonstrations without estimating the corresponding reward function, and overcomes compounding errors caused by behavioral cloning (BC) [Reddy *et al.*, 2019] as well as reduces the computational burden of inverse reinforcement learning (IRL) [Wulfmeier *et al.*, 2015; Pflueger *et al.*, 2019]. Existing imitation learning methods, including GAIL, mostly focus on

reconstructing expert behavior based on the assumption of single modality. However, most of real world demonstrations have multiple modalities with various skills and habits. For example, there are three distinct intentions in a driving task: lane-change left, lane keeping and lane-change right. In that case, imitation learning algorithms like GAIL will cause the mode collapse problem due to the inability of discovering and distinguishing mode variation in expert demonstrations. Besides, much of real world tasks like aforementioned driving task need to select behavior mode based on current situation for decision-making adaptively (i.e., human drivers will determine whether to change lanes based on traffic conditions) instead of specified manually.

Some extensions of GAIL have been proposed to deal with multi-modal tasks. [Li *et al.*, 2017; Kuefler and Kochenderfer, 2018; Wang *et al.*, 2017] learn latent codes in an unsupervised manner and recover multi-modal policies from unlabeled demonstrations, which need random sampling of latent codes. [Merel *et al.*, 2017; Lin and Zhang, 2018; Merel *et al.*, 2017] reconstruct modal information directly from expert demonstration labels or add an auxiliary classifier to assist the adversary, involving a supervised learning process. However, most of the extensions only focus on learning to distinguish different skills with random sampling of skill labels, thus they are not able to deal with those real world scenarios which require adaptive skill selection conditioned on environmental situations. This motivates our research.

In this paper we propose a new approach learning to select skill labels and imitate multi-modal policy simultaneously. The algorithm, called Triple-GAIL, is an extension of GAIL for distinguishing multiple modalities accurately and efficiently enhancing the performance on label-conditional imitation learning tasks. In particular, the contributions of this paper are as follows: (i) Similar to Triple-GAN [Chongxuan *et al.*, 2017], we propose a novel adversarial game framework which extends the original GAIL with an auxiliary selector. The selector and the generator in Triple-GAIL characterize the conditional distribution given the state-action pairs and state-label pairs while the discriminator distinguishes whether a state-action-label pair comes from expert demonstrations or not. (ii) Both of the generator and the selector

\*Corresponding author

†The first three authors contribute equally.

have been proved to converge to their own optima respectively with compatible utilities, which means Triple-GAIL can learn a good skill selector and a conditional generator simultaneously. (iii) We apply our algorithm in a driving task and a real-time strategy (RTS) game. Experimental results demonstrate that Triple-GAIL can distinguish multiple modalities clearly as well as enhance the performance on label-conditional imitation learning tasks.

## 2 Background and Related Work

### 2.1 Generative Adversarial Imitation Learning

GAIL is a promising imitation learning method based on generative adversarial nets (GANs) [Goodfellow *et al.*, 2014]. In GAIL, the generator serves as a policy to imitate expert behavior by matching the state-action  $(s, a)$  distribution of demonstrations, while the discriminator plays a role of surrogate reward to measure the similarity between the generated data and demonstration data. GAIL directly optimizes the policy without solving the reward function in IRL. The objective of GAIL is formulated as the min-max form:

$$\min_{\pi} \max_{D \in (0,1)} \mathbb{E}_{\pi} [\log D(s, a)] + \mathbb{E}_{\pi_E} [\log (1 - D(s, a))] - \lambda H(\pi) \quad (1)$$

where  $\pi$  and  $D$  are the generator (policy) and the discriminator respectively. The casual entropy  $H(\pi)$  serves as a regularization term of policy together with hyper-parameter  $\lambda$ . In practice, trust region policy optimization (TRPO) [Schulman *et al.*, 2015] is used to update the policy  $\pi$  with the surrogate reward function:  $r = -\log D(s, a)$ .

### 2.2 Multi-modal Imitation Learning Algorithms

There have been some extensions of GAIL to address multi-modal tasks. One typical way is to distinguish modal information in an unsupervised manner. InfoGAIL [Li *et al.*, 2017] infers latent codes by maximizing the mutual information between latent variables and observed state-action pairs. Burn-InfoGAIL [Kuefler and Kochenderfer, 2018] uses the maximum mutual information from the perspective of Bayesian inference to draw modal variables from burn-in demonstrations. VAE-GAIL [Wang *et al.*, 2017] introduces a variational autoencoder to infer modal variable, which allows for smoothing policy interpolation. The above algorithms can learn multi-modal policies from unlabeled demonstrations. However, due to lack of labels in demonstrations, these algorithms tend to distinguish latent labels without considering semantic information or task context.

Another direction focuses on labeled data in expert demonstrations. In CGAIL [Merel *et al.*, 2017], the modal labels are directly sent to the generator and the discriminator, which are conditioned onto the label itself. ACGAIL [Lin and Zhang, 2018] introduces an auxiliary classifier to reconstruct the modal information specially while the discriminator is only responsible for distinguishing whether the input sample  $(s, a)$  is from the demonstrations. In ACGAIL the classifier cooperates with the discriminator by sharing parameters, both of which provide adversarial loss to the generator.




Figure 1: The overall structure of Triple-GAIL. Label  $c$  is drawn from expert data (blue line) in the training process, while inferred from state-action pairs in application (red line), which characterizes the conditional distribution  $p_{C_\alpha}$  (red block). The generator characterizes the conditional distribution  $p_{\pi_\theta}$  (blue block). All these data sequences are sent to the discriminator together with expert data  $p_{\pi_E}$ .  $t-k:t$  means time sequences.  $E$  denotes an encoder.

Note that that above methods mainly leverage random sampling of latent labels from known prior distribution to distinguish multiple modalities. Once trained, model outputs the corresponding actions based on the manually specified labels. However, in this paper we focus on dealing with those tasks which require adaptive skill mode selection according to environmental situations. Moreover, we are interested in labeled expert demonstrations with multiple modalities. Different from existing works, the proposed Triple-GAIL is able to learn skill selection and imitation jointly from both expert demonstrations and continuously generated experiences.

## 3 Method

Suppose we can get a mixed set of labeled demonstrations with multiple expert modalities. In this paper we propose to learn one policy simultaneously from multiple expert demonstrations. More specifically, the expert policy including multiple skill labels is presented as  $\pi_E = \{\pi_{E_1}, \dots, \pi_{E_k}\}$ , which is determined by  $p(\pi|c)$ , where  $c$  is the skill label. In order to select skill labels from current environmental observations adaptively instead of specifying manually and then reconstruct multi-modal policy simultaneously, a novel adversarial imitation framework is introduced as follows.

### 3.1 Triple-GAIL Framework

Triple-GAIL consists of three main components represented by neural networks as shown in Figure 1: a selector  $C_\alpha$  parameterized by  $\alpha$ , which produces skill labels given states and actions; a generator  $\pi_\theta$  parameterized by  $\theta$ , which acts as a corresponding policy conditioned on states and skill labels; and a discriminator  $D_\psi$  parameterized by  $\psi$ , which distinguishes whether the state-action-label pairs come from expert demonstrations or not. The joint distribution of state-action-label pairs in the setting of Triple-GAIL can be defined in two directions:  $\pi_\theta$  approximately characterizes the conditional distribution  $p_{\pi_\theta}(a|s, c)$  given  $s$  and  $c$ , while  $C_\alpha$  characterizes the conditional distribution  $p_{C_\alpha}(c|s, a)$  given  $s$  and  $a$ . We make the mild assumption that  $p(s, c)$  and  $p(s, a)$

can be obtained from the demonstrations and generated data respectively, then the two distributions are defined as follows:

$$p_{\pi_\theta}(s, a, c) = p(s, c)p_{\pi_\theta}(a|s, c) \quad (2)$$

$$p_{C_\alpha}(s, a, c) = p(s, a)p_{C_\alpha}(c|s, a) \quad (3)$$

where  $p_{\pi_\theta}(s, a, c)$  and  $p_{C_\alpha}(s, a, c)$  are the joint distributions defined by  $\pi_\theta$  and  $C_\alpha$  respectively, and  $p_{\pi_E}(s, a, c)$  denotes the expert distribution. We expect to achieve an equilibrium that both  $p_{\pi_\theta}(s, a, c)$  and  $p_{C_\alpha}(s, a, c)$  converge to the expert distribution.

In this game, we can draw skill labels  $c$  from expert demonstrations and then produces actions  $a$  conditioned on the current states and labels following Eqn. (2), which generates pseudo joint pairs  $(s^g, a^g, c^g)$ . Similarly, the selector provides pseudo skill labels given the current states and last actions generated from interaction, hence pseudo joint pairs  $(s^c, a^c, c^c)$  can also be sampled following Eqn. (3). Besides, real joint pairs  $(s^e, a^e, c^e)$  from expert demonstrations are provided as positive samples. Then, the joint pairs sampled from  $p_{\pi_\theta}(s, a, c)$ ,  $p_{C_\alpha}(s, a, c)$  and  $p_{\pi_E}(s, a, c)$  are all sent to the discriminator  $D_\psi$  for judgement. Note that different from Triple-GAN, which introduces a classifier for label classification in semi-supervised learning, the selector in Triple-GAIL in essence plays the role of skill selection under current circumstance, not just for distinguishing skill labels.

In such a three-player game, the generator and the selector work cooperatively against the discriminator. The adversarial loss, in turn, guides the generator as well as refines the selector, resulting in a multi-modal policy. In analogy with GAIL, the framework of Triple-GAIL is first formulated as a game to minimize  $\alpha$ ,  $\theta$  and maximize  $\psi$ :

$$\begin{aligned} \min_{\alpha, \theta} \max_{\psi} & \mathbb{E}_{\pi_E} [\log(1 - D_\psi(s, a, c))] \\ & + \omega \mathbb{E}_{\pi_\theta} [\log D_\psi(s, a, c)] \\ & + (1 - \omega) \mathbb{E}_{C_\alpha} [\log D_\psi(s, a, c)] - \lambda_H H(\pi_\theta) \end{aligned} \quad (4)$$

where  $\mathbb{E}_{\pi_E}$ ,  $\mathbb{E}_{\pi_\theta}$  and  $\mathbb{E}_{C_\alpha}$  denote  $\mathbb{E}_{(s, c, a) \sim p_{\pi_E}(s, c, a)}$ ,  $\mathbb{E}_{(s, c, a) \sim p_{\pi_\theta}(s, c, a)}$ ,  $\mathbb{E}_{(s, c, a) \sim p_{C_\alpha}(s, c, a)}$  respectively,  $\omega \in (0, 1)$  is a hyper-parameter that balances the weights of policy generation and skill selection, and  $H(\pi_\theta)$  is the policy casual entropy defined as  $\mathbb{E}_{\pi_\theta} [-\log \pi_\theta(a|s, a)]$  with hyper-parameter  $\lambda_H > 0$ .

Since the optimal solution should be obtained if and only if the pseudo distributions defined by the generator and the selector both converge to the distribution of expert, we introduce two cross-entropy terms  $R_E$  and  $R_G$  (See in Sec. 3.2), and define the objective function of this three-player game as:

$$\begin{aligned} \min_{\alpha, \theta} \max_{\psi} & \mathbb{E}_{\pi_E} [\log(1 - D_\psi(s, a, c))] \\ & + \omega \mathbb{E}_{\pi_\theta} [\log D_\psi(s, a, c)] \\ & + (1 - \omega) \mathbb{E}_{C_\alpha} [\log D_\psi(s, a, c)] \\ & + \lambda_E R_E + \lambda_G R_G - \lambda_H H(\pi_\theta) \end{aligned} \quad (5)$$

where  $\lambda_E$  and  $\lambda_G$  weigh the relative importance of two supervised loss.

### 3.2 Objective Design and Theoretical Analysis

This section provides the formal theoretical analysis of Triple-GAIL. Following the similar proof framework of

Triple-GAN, first the theoretical prediction form of the optimal discriminator can be derived as in Lemma 1:

**Lemma 1.** *For any fixed generator and selector, the optimal form of the discriminator is denoted as:*

$$D_{\psi^*} = \frac{\omega p_{\pi_\theta} + (1 - \omega) p_{C_\alpha}}{p_{\pi_E} + \omega p_{\pi_\theta} + (1 - \omega) p_{C_\alpha}} = \frac{p_\omega}{p_{\pi_E} + p_\omega} \quad (6)$$

where  $p_{\pi_E}$ ,  $p_{\pi_\theta}$  and  $p_{C_\alpha}$  denote  $p_{\pi_E}(s, a, c)$ ,  $p_{\pi_\theta}(s, a, c)$  and  $p_{C_\alpha}(s, a, c)$  respectively, and  $p_\omega$  is defined as  $\omega p_{\pi_\theta} + (1 - \omega) p_{C_\alpha}$ .

Given the optimal  $D_{\psi^*}$  and omit policy entropy term, we can derive the equilibrium conditions and properties.

**Lemma 2.** *The min-max game in Eqn. (4) can achieve the multiple equilibrium that  $p_{\pi_E} = p_\omega$ , where variable  $\omega$  is a mixing coefficient between  $p_{\pi_\theta}(s, a, c)$  and  $p_{C_\alpha}(s, a, c)$ .*

Since  $\omega$  is a variable, Lemma 2 only proves a mixed distribution of  $p_{\pi_\theta}(s, a, c)$  and  $p_{C_\alpha}(s, a, c)$  converges to the true distribution of expert but fails to guarantee each of them converges to  $p_{\pi_E}(s, a, c)$ , that is  $p_{\pi_\theta}(s, a, c) = p_{C_\alpha}(s, a, c) = p_{\pi_E}(s, a, c)$ . To address this problem, we introduce two cross-entropy terms  $R_E$  and  $R_G$  as follows:

$$\begin{aligned} R_E &= \mathbb{E}_{\pi_E} [-\log p_{C_\alpha}(c|s, a)] \\ &\approx -\frac{1}{N} \sum_{i=0}^N \frac{1}{T} \sum_{t=1}^T c_{i,t}^e \log p_{C_\alpha}(c_{i,t}^e | s_{i,t}^e, a_{i,t-1}^e) \end{aligned} \quad (7)$$

$$\begin{aligned} R_G &= \mathbb{E}_{\pi_\theta} [-\log p_{C_\alpha}(c|s, a)] \\ &\approx -\frac{1}{N} \sum_{i=0}^N \frac{1}{T} \sum_{t=1}^T c_{i,t}^g \log p_{C_\alpha}(c_{i,t}^g | s_{i,t}^g, a_{i,t-1}^g) \end{aligned} \quad (8)$$

where the superscripts  $e$ ,  $g$  and  $c$  indicate the samples provided by the discriminator, the generator and the selector respectively. The subscript indicates the timestep. Consider that labels are drawn from expert demonstrations in the training phase,  $c_t^e = c_t^g$ .  $R_E$  is the standard supervised loss ensuring that the selector converges to expert distribution.  $R_G$  is essentially the divergence between the pseudo distribution  $p_{C_\alpha}(s, c, a)$  and generated distribution  $p_{\pi_\theta}(s, c, a)$ . And this optimizes the selector using the generated data from interaction, which can be viewed as data augmentation for the selector. A learning rate schedule is also introduced for  $R_G$  to boost training performance. The advantage of  $R_G$  is shown in Sec. 4. By combining Eqn. (7) and Eqn. (8) to the initial objective Eqn. (4), we can obtain the final form of objective given in Eqn. (5):

**Theorem 1.** *Eqn. (5) ensures the existence and uniqueness of the global equilibrium, which is achieved if and only if  $p_{\pi_\theta}(s, a, c) = p_{C_\alpha}(s, a, c) = p_{\pi_E}(s, a, c)$ .*

Following Theorem 1, we can guarantee both the generator and the selector can converge to their optima respectively.

The whole training procedure of Triple-GAIL is summarized in Algorithm 1. Triple-GAIL has three models consisting of four neural networks: The generator  $\pi_\theta$  consists of a policy network and a value network. The selector and the discriminator are characterized by the selector network  $C_\alpha$

---

**Algorithm 1** The Training Procedure of Triple-GAIL
 

---

**Input:** The multi-intention trajectories of expert  $\tau_E$ ; **Parameter:** The initial parameters  $\theta_0, \alpha_0$  and  $\psi_0$

```

1: for  $i = 0, 1, 2, \dots$  do
2:   for  $j = 0, 1, 2, \dots, N$  do
3:     Reset environments by the demonstration episodes with fixed label  $c_j$ ;
4:     Run policy  $\pi_\theta(\cdot|c_j)$  to sample trajectories:  $\tau_{c_j} = (s_0, a_0, s_1, a_1, \dots, s_{T_j}, a_{T_j}|c_j)$ 
5:   end for
6:   Update the parameters of  $\pi_\theta$  via TRPO with rewards:  $r_{t_j} = -\log D_\psi(s_{t_j}, a_{t_j}, c_j)$ 
7:   Update the parameters of  $D_\psi$  by gradient ascending with respect to:
    
```

$$\nabla_\psi \frac{1}{N_e} \sum_{n=1}^{N_e} \log(1 - D_\psi(s_n^e, a_n^e, c_n^e)) + \frac{1}{N} \sum_{j=1}^N \left[ \frac{\omega}{T_j} \sum_{t=1}^{T_j} \log D_\psi(s_t^g, a_t^g, c_j^g) + \frac{1-\omega}{T_j} \sum_{t=1}^{T_j} \log D_\psi(s_t^c, a_t^c, c_j^c) \right] \quad (9)$$

```

8:   Update the parameters of  $C_\alpha$  by gradient descending with respect to:
    
```

$$\nabla_\alpha \frac{1}{N} \sum_{j=1}^N \left[ \frac{1-\omega}{T_j} \sum_{t=1}^{T_j} \log D_\psi(s_t^c, a_t^c, c_j^c) - \frac{\lambda_E}{T_j} \sum_{t=1}^{T_j} c_j^e \log p_{C_\alpha}(c_t^e | s_t^e, a_{t-1}^e) - \frac{\lambda_G}{T_j} \sum_{t=1}^{T_j} c_j^e \log p_{C_\alpha}(c_t^c | s_t^g, a_{t-1}^g) \right] \quad (10)$$

```

9: end for
    
```

---

and the discriminator network  $D_\psi$ , respectively. Firstly, we reset the environment by the labeled episodes, namely, each episode has a fixed true label. The agent runs the policy with these labels and gathers the generated data. The generated data are then sent to the selector to produce skill labels. The data generated by generator and selector are all sent to the discriminator as the pseudo data while the demonstration data is served as true data. Then the discriminator network is updated by ascending the gradient with above sampled data with (9) while the selector network is updated by descending the gradient with (10). Our policy network and value network are updated by TRPO. Once trained, the selector adaptively generates skill label based on state-action pairs, and the skill label is input to the generator to produce corresponding actions.

## 4 Experiments

Note that we tend to solve real world tasks which need to adaptively select skill mode and guide decision-making based on current environmental situations. So we demonstrate the performance of our method on two typical tasks, both of which drive their policies with multiple explicit skills under specific circumstances. We first apply it to a driving task where the agent adaptively selects whether to change lanes based on the highway traffic conditions and imitates corresponding driving behaviors. Then, we extend it to an RTS game, where the agent needs to choose its skills considering the enemy's tactical intentions and take corresponding strategies. For both tasks, our algorithm is evaluated against three baselines: BC, GAIL and CGAIL. Note that InfoGAIL and ACGAIL need to manually specify skill labels by experts, which beyond our comparison domain. We modify original CGAIL by adding a classifier with the same structure of selector in Triple-GAIL. This classifier is trained by supervised learning and the parameters are fixed.

### 4.1 Experimental Setup

For the driving task, we follow the state and action representations as in [Henaff *et al.*, 2019]. For the generator network, the images are input to a 3-layer convolutional network with  $64 \times 128 \times 256$  feature maps, while the vectors are run through fully connected network with  $256 \times 256$  hidden units with a final layer expands its size the same as the output of the convolutional network. The skill label is run through fully connected network with  $128 \times 128$  hidden units with corresponding expansion to the size of input states. Then the input states and the skill label are now the same size and concatenated together, and are run through fully connected network with  $256 \times 128$  hidden units to output actions and value. The selector network has the same architecture of the generator with input states and actions then outputs skill label. The discriminator network is similar to the selector, and we just adjust all hidden units to 64 for the input states, actions and skill label, then run through fully connected network with  $128 \times 64$  hidden units to output reward value. In the RTS task, the state information includes current resource of players and the images with  $20 \times 20$  dimensions.

### 4.2 Learning to Drive in Dense Traffic

The experiment is conducted with the Next Generation Simulation program's Interstate 80 (NGSIM I-80) dataset [Halkias and Colyar, 2006]. NGSIM I-80 dataset includes various complex driver skills or behaviors such as lane changes, merges and sudden accelerations, and is recorded at different times of day with different traffic conditions, which contains uncongested and congested peak period. [Henaff *et al.*, 2019] provides an interactive simulation environment with this dataset, which is used in our driving task.

We first apply preprocessing of the dataset for learning skill selection. The expert trajectories are labeled manually

Algorithms	Success Rate (%)	Mean Distance (m)	KL Divergence		
			Lane-change Left	Lane-keeping	Lane-change Right
BC	$6.8 \pm 3.2$	$81.7 \pm 3.1$	$3827 \pm 358$	$4008 \pm 486$	$2581 \pm 371$
GAIL	$73.9 \pm 1.3$	$168.4 \pm 5.2$	$1764 \pm 279$	$1893 \pm 378$	$606 \pm 278$
CGAIL	$65.5 \pm 0.9$	$149.8 \pm 7.2$	$1297 \pm 255$	$1892 \pm 279$	$977 \pm 109$
Triple-GAIL	<b><math>80.9 \pm 1.2</math></b>	<b><math>179.6 \pm 3.6</math></b>	<b><math>447 \pm 122</math></b>	<b><math>685 \pm 214</math></b>	<b><math>392 \pm 127</math></b>
Expert	100	$210 \pm 2.1$	0	0	0

Table 1: The Success Rate, Mean Distance and KL Divergence of different algorithms.




Figure 2: The visualization of trajectories. The trajectories of lane-change right, lane keeping and lane-change left are represented by red, green and blue lines respectively. Dis. denotes displacement.

based on rules with three skills:  $[1, 0, 0]$  corresponds to lane-change left,  $[0, 1, 0]$  corresponds to lane keeping and  $[0, 0, 1]$  corresponds to lane-change right. 150 trajectories are sampled from demonstrations with the length of 13s (130 frames). Each of modalities has 50 trajectories.

### Performance of Skill Selection

We first estimate the performance of the selector, which acts on selecting driving skills. The rollouts are visualized in Figure 2, where the red, green and blue trajectories indicate various skills labeled by expert, corresponding to lane-change left, lane keeping and lane-change right respectively.

As shown in Figure 2(a), the BC model has the shortest trajectory length because of collision. And the trajectories generated by BC tend to deviate to left side due to cascade errors. GAIL has longer trajectories, but fails to distinguish skill labels as shown in the mixed color of Figure 2(b). We conclude that BC and GAIL have poor ability in distinguishing and selecting different skills from demonstrations due to the mode collapse problem.

Compared with BC and GAIL, CGAIL and Triple-GAIL have a more clear separation among rollout trajectories with different driving skills. However, CGAIL has relative short

Algorithms	Accuracy
CGAIL	83.2%
Triple-GAIL	<b>90.7%</b>
Triple-GAIL $\setminus R_E$	69.6%
Triple-GAIL $\setminus R_G$	81.3%

 Table 2: Selection accuracies of driving skills. Triple-GAIL  $\setminus R_E$  removes supervised loss  $R_E$  while Triple-GAIL  $\setminus R_G$  removes supervised loss  $R_G$ .

trajectories, also indicated in *Mean Distance* in Table 1. We guess that the pretrained classifier of CGAIL is unable to choose skill label adaptively guiding the generator for decision-making without joint optimization. In contrast, Triple-GAIL can successfully select driving skill as well as learn effective lane-change driving policies accordingly.

### Overall Result of Policy Imitation

We first evaluate our algorithm in two metrics: *Success Rate* and *Mean Distance*. The *Success Rate* indicates the probability of moving through the road segment without collision, while the *Mean Distance* indicates the distance travelled before the episode ends. Then, in order to further evaluate the statistical distance between learned policy and expert policy, we sample the trajectory positions  $(x, y)$  from demonstrations and rollouts separately, which can be regarded as the low-dimension projection of policies. The *KL Divergence* of position distribution between generated data and demonstrations is also calculated.

From Table 1 we conclude that Triple-GAIL outperforms all other three baselines in both *Success Rate* and *Mean Distance*, and is closest to expert demonstrations. Furthermore, for different skills in this task including lane-change left, lane keeping and lane-change right, Triple-GAIL indicates the smallest *KL Divergence* compared with other baselines, which means with appropriate skill selection, Triple-GAIL is able to imitate corresponding expert policies effectively.

### Ablation Study

To further estimate the performance of the selector and the joint optimization of the selector and the generator, the selection accuracies of driving skills is compared in Table 2. We show that both CGAIL and Triple-GAIL have high selection accuracies, while Triple-GAIL is slightly higher up to 90%. The comparison between CGAIL and Triple-GAIL illustrates that the joint optimization of the selector and the generator

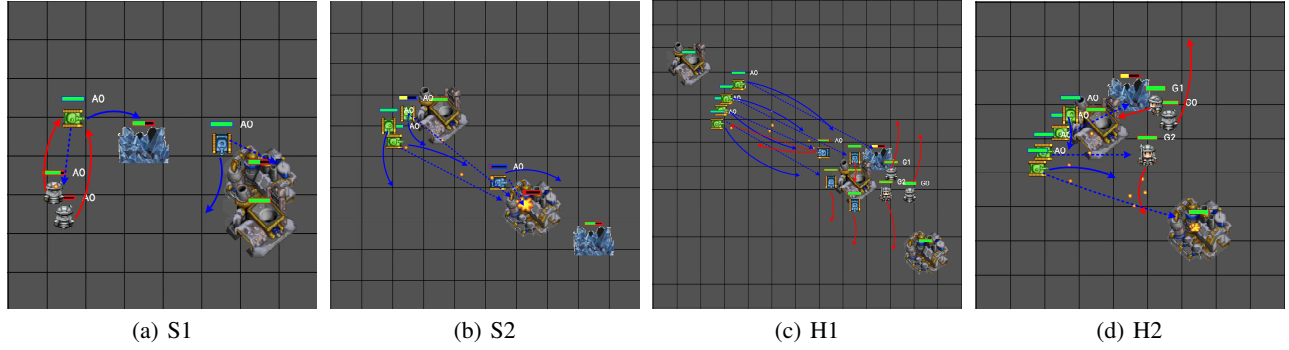


Figure 3: Game screenshots between trained agent (blue) and built-in agent (red). The agents’ colors are shown on the arrows and the boundary of hit point gauges. (a)~(b) is an example case against *SIMPLE* and (c)~(d) is an example case against *HIT-N-RUN*. The solid line indicates the movements of agents while the dotted line points to the attack target.

Algorithms	<i>SIMPLE</i>	<i>HIT-N-RUN</i>
BC	$42.7 \pm 5.6$	$31.5 \pm 3.8$
GAIL	$60.4 \pm 7.2$	$52.4 \pm 6.4$
CGAIL	$67.5 \pm 5.6$	$60.5 \pm 6.3$
Triple-GAIL	<b><math>73.9 \pm 3.2</math></b>	<b><math>68.3 \pm 5.6</math></b>
CGAIL+label	$69.5 \pm 6.9$	$62.7 \pm 7.3$
Triple-GAIL+label	$78.9 \pm 3.2$	$76.3 \pm 6.6$
Expert Matched	$95.2 \pm 1.7$	$90.6 \pm 4.9$
Expert Mismatched	$38.2 \pm 4.5$	$54.3 \pm 5.9$

Table 3: Win rates of different algorithms competing with built-in agents over 10k games. +label denotes that the inferred labels are replaced by true labels. Matched denotes the agents run against the targeted built-in agents while Mismatched denotes the agents run against the mismatched built-in agents.

in Triple-GAIL is superior to CGAIL which is pretrained and fixed. There is a significant difference in Triple-GAIL with and without the cross-entropy loss term  $R_E$ , which shows that the supervised signal from demonstrations plays an important role in training the selector. In addition, the comparison of Triple-GAIL with and without the  $R_G$  also validates the advantage of  $R_G$ .

### 4.3 Learning to Play RTS Game

We then verify the Triple-GAIL in a Mini-RTS game, which is a miniature version of StarCraft [Tian *et al.*, 2017]. In Mini-RTS game, players are required to gather resources, build troops and finally invade/defend the enemy until one player wins. There are two built-in agents in Mini-RTS: *SIMPLE* and *HIT-N-RUN*. *SIMPLE* is a conservative strategy, where all troops stay on the defensive until the number of melee tanks reaches up to 5. After that, all troops launch a counterattack. *HIT-N-RUN* usually builds 2 range tanks that move towards enemy base to harass the opponent, taking advantage of long attack range and high speed. We train the agent with frame-skip of 50 and history length of 20. Other settings and details refer to [Tian *et al.*, 2017].

In order to gather demonstrations with multi-modal poli-

cies, two sets of targeted instructions adopted for the above two built-in agents respectively. The win rates of these two targeted instructions reach up to 90% while less than 60% if the agents run against the mismatched built-in agents, as shown in Table 3. Then the demonstrations with multiple policies are sampled by running the corresponding games and labeling the sampled state-action pairs. After that, Triple-GAIL and three baselines are trained.

We compare the performance of all four algorithms and the win rates are listed in Table 3. It is clear that Triple-GAIL provides better performance than all other baselines in both two built-in agents. We demonstrate that Triple-GAIL can effectively distinguish the enemy’s tactical intention and helps to adopt the corresponding policy. When specifying skill labels with expert demonstration (CGAIL+label) instead of inferring from the classifier of CGAIL, we confirm that the joint optimization of the selector and the generator in Triple-GAIL indeed improves the policy imitation performance.

Figure 3 provides typical game screenshots of Triple-GAIL. When the trained agents versus *SIMPLE*, the learned agent commonly builds one range tank and directly moves towards to enemy base and attack energy troops in range, continuing to harass the enemy, as shown in the top row of Figure 3(a). Once gaining an advantage, all melee and range tanks launch a general attack as shown in Figure 3(b); if the opponent is *HIT-N-RUN*, the trained agent will firstly build several range tanks against the harassment as shown in Figure 3(c), and immediately launch a counterattack if enemy is weakly guard as shown in the Figure 3(d).

## 5 Conclusion

In this paper, we propose Triple-GAIL, a novel multi-modal GAIL framework that is able to learn skill selection and imitation jointly from both expert demonstrations and continuously generated experiences by introducing an auxiliary selector. We provide theoretical guarantees on the convergence to optima for both of the generator and the selector respectively. Experiments on driving task and real-time strategy game demonstrate that Triple-GAIL can better fit multi-modal behaviors close to the demonstrators and outperforms state-of-the-art methods.



## References

- [Abbeel and Ng, 2004] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pages 1–8, 2004.
- [Chongxuan et al., 2017] LI Chongxuan, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *Advances in neural information processing systems*, pages 4088–4098, 2017.
- [Fang et al., 2019] Bin Fang, Shidong Jia, Di Guo, Muhua Xu, Shuhuan Wen, and Fuchun Sun. Survey of imitation learning for robotic manipulation. *International Journal of Intelligent Robotics and Applications*, pages 1–8, 2019.
- [Goodfellow et al., 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [Halkias and Colyar, 2006] John Halkias and James Colyar. Next generation simulation fact sheet. *US Department of Transportation: Federal Highway Administration*, 2006.
- [Henaff et al., 2019] Mikael Henaff, Alfredo Canziani, and Yann LeCun. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. *arXiv preprint arXiv:1901.02705*, 2019.
- [Ho and Ermon, 2016] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
- [Kuefler and Kochenderfer, 2018] Alex Kuefler and Mykel J Kochenderfer. Burn-in demonstrations for multi-modal imitation learning. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1071–1078, 2018.
- [Li et al., 2017] Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, pages 3812–3822, 2017.
- [Lin and Zhang, 2018] Jiahao Lin and Zongzhang Zhang. Acgail: Imitation learning about multiple intentions with auxiliary classifier gans. In *Pacific Rim International Conference on Artificial Intelligence*, pages 321–334. Springer, 2018.
- [Merel et al., 2017] Josh Merel, Yuval Tassa, Sriram Srinivasan, Jay Lemmon, Ziyu Wang, Greg Wayne, and Nicolas Heess. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.
- [Pflueger et al., 2019] Max Pflueger, Ali Agha, and Gaurav S Sukhatme. Rover-irl: Inverse reinforcement learning with soft value iteration networks for planetary rover path planning. *IEEE Robotics and Automation Letters*, 4(2):1387–1394, 2019.
- [Reddy et al., 2019] Siddharth Reddy, Anca D Dragan, and Sergey Levine. Sqil: imitation learning via regularized behavioral cloning. *arXiv preprint arXiv:1905.11108*, 2019.
- [Schulman et al., 2015] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897, 2015.
- [Tian et al., 2017] Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, and C Lawrence Zitnick. Elf: An extensive, lightweight and flexible research platform for real-time strategy games. In *Advances in Neural Information Processing Systems*, pages 2659–2669, 2017.
- [Wang et al., 2017] Ziyu Wang, Josh S Merel, Scott E Reed, Nando de Freitas, Gregory Wayne, and Nicolas Heess. Robust imitation of diverse behaviors. In *Advances in Neural Information Processing Systems*, pages 5320–5329, 2017.
- [Wulfmeier et al., 2015] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*, 2015.