

Semi-Supervised Imitation Learning with Mixed Qualities of Demonstrations for Autonomous Driving

Gunmin Lee¹, Wooseok Oh^{1%}, Jeongwoo Oh^{1%}, Seungyeon Shin², Dohyeong Kim¹, Jaeyeon Jeong¹,
Sungjoon Choi³, and Songhwa Oh^{1*}

¹Department of Electrical and Computer Engineering and ASRI, Seoul National University,
Seoul 08826, Korea ({first name}.{last name}@rllab.snu.ac.kr, songhwa@snu.ac.kr)

²School of Computer Science and Engineering Dongguk University, Seoul, 04620, Korea (2018112005@dongguk.edu)

³Department of Artificial Intelligence, Korea University, Seoul 02841, Korea, (sungjoon-choi@korea.ac.kr)

* Corresponding author % Equal contribution

Abstract: In this paper, we consider the problem of autonomous driving using imitation learning in a semi-supervised manner. In particular, both labeled and unlabeled demonstrations are leveraged during training by estimating the quality of each unlabeled demonstration. If the provided demonstrations are corrupted and have a low signal-to-noise ratio, the performance of the imitation learning agent can be degraded significantly. To mitigate this problem, we propose a method called semi-supervised imitation learning (SSIL). SSIL first learns how to discriminate and evaluate each state-action pair's reliability in unlabeled demonstrations by assigning higher reliability values to demonstrations similar to labeled expert demonstrations. This reliability value is called leverage. After this discrimination process, labeled and unlabeled demonstrations with estimated leverage values are utilized while training the policy in a semi-supervised manner. The experimental results demonstrate the validity of the proposed algorithm using unlabeled trajectories with mixed qualities. Moreover, the hardware experiments using an RC car are conducted to show that the proposed method can be applied to real-world applications.

Keywords: Autonomous driving, Deep learning, Imitation learning

1. INTRODUCTION

Designing a real-time controller for robots with high-speed motion is a challenging problem. Conventional approaches of designing a rule-based feedback controller [1], [2] have shortcomings in covering all the exception cases. Recent approaches include a method that can autonomously cover all the exception cases, such as training the controller via deep reinforcement learning (RL) methods [3], [4]. However, it is challenging to use RL methods for physical robots due to their high sample complexity, safety concerns, and difficulty in designing feasible reward functions [5].

To address the above issues, imitation learning (IL) has been a trending topic of designing a controller. IL methods do not need to design a reward function and tend to show better performance due to their simplicity of following an ideal trajectory. In general, the main objective of imitation learning is to find a controller that successfully mimics the given expert demonstrations.

Standard approaches to imitation learning (IL) [6], [7], [8] assume that only the expert data are given to the agent. However, this assumption may not always be valid. One of this assumption's main limitations is that collecting pure expert data is challenging due to the limited budget and time. The data may be noisy, not collected from experts, or contain incomplete demonstrations. This causes the learning quality of imitation learning to degrade. On the other hand, collecting diverse quality data is cheaper. However, it is not straightforward to evaluate and discriminate such data for training due to the variability in quality.

To resolve this problem, Choi et al. [9], [10] proposed

a method that can utilize demonstrations with mixed qualities by estimating the quality of each data through the lens of correlation. The cornerstone of the method is to model multiple policies with correlated Gaussian processes by assuming that the expert's policy is the dominant pattern in the dataset. This method was extended to neural networks in [11] to address the scalability issue of Gaussian process regression. Mixed demonstration generative adversarial imitation learning (MixGAIL), proposed in [12], is an imitation learning method that can be trained with both good and bad demonstrations. But it is assumed that all demonstrations are labeled.

In this paper, we introduce a novel imitation learning method named *semi-supervised imitation learning* (SSIL) which can evaluate the quality of unlabeled data using the predictive uncertainty. First, we train a variational autoencoder (VAE) from a small number of demonstrations from experts. From the VAE we trained, we can measure the reconstruction losses on unlabeled demonstrations. Then, we normalize the reconstruction loss and convert them into the confidence value of each state-action pair in each unlabeled demonstration. This confidence value is called the leverage value of the given demonstration. We then provide a state-action-leverage value triplet set to the SSIL framework.

We show the strength of the proposed method by comparing it against MixGAIL [12] and GAIL [6]. The experiments are performed in TORCS [13] and a real-world RC car driving scenario. SSIL shows superior performance compared to GAIL and MixGAIL when only more than 0.61% of demonstrations are given as labeled expert demonstrations. Moreover, in a real-world experiment with an RC car, SSIL shows maximum score improvements of 19.0% and 12.1% over GAIL and MixGAIL, respectively. In addition, SSIL shows a similar standard deviation of the performance score on three different trials with different seeds as MixGAIL, which is 58.1% of

This work was supported by the Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2019-0-01309, Development of AI Technology for Guidance of a Mobile Robot to its Goal with Uncertain Maps in Indoor/Outdoor Environments).

the standard deviation of the performance score of GAIL. Note that to show the robustness of our algorithm, the domain for collecting data and the domain used for testing our algorithm are different.

2. BACKGROUND

2.1 Mixed Demonstration Generative Adversarial Imitation Learning

Generative adversarial imitation learning (GAIL) is a powerful imitation learning method. In this algorithm, the agent imitates the given expert demonstration. While GAIL [6] uses only expert demonstrations, mixed demonstration generative adversarial imitation learning (MixGAIL) [12] uses demonstrations with poor qualities (negative demonstrations) along with expert demonstrations.

The objective function of MixGAIL is given as:

$$\min_{\pi \in \Pi} \max_D E_{\pi} [\log(D(s, a)) - \eta \text{CoR}(s, s_E, s_N)] + E_{\pi_E} [\log(1 - D(s, a))] - \lambda H(\pi), \quad (1)$$

where π_E and π are stochastic policies of expert trajectory and generator, respectively. E_{π} and E_{π_E} are the expected returns of policy π and π_E , respectively. The state is $s \in S$, and action is $a \in A$, where S and A are the state and action spaces. $D : S \times A \mapsto (0, 1)$ is a discriminator, and the causal entropy of π is represented as $H(\pi)$. s_E is a set of states of expert demonstrations, while s_N is a set of states in negative demonstrations. CoR is a term called the *constraint reward*. η and λ control the proportion of CoR and H , respectively. The constraint reward, CoR, is given as:

$$\text{CoR}(s, s_E, s_N) = \frac{(1 + \frac{d_{s_E}}{\alpha})^{-\frac{\alpha+1}{2}}}{(1 + \frac{d_{s_E}}{\alpha})^{-\frac{\alpha+1}{2}} + (1 + \frac{d_{s_N}}{\alpha})^{-\frac{\alpha+1}{2}}}, \quad (2)$$

$$d_{s_E} = \sqrt{\frac{1}{|s_E|} \sum_{s_e \in s_E} (\|s - s_e\|_2)^2},$$

$$d_{s_N} = \sqrt{\frac{1}{|s_N|} \sum_{s_n \in s_N} (\|s - s_n\|_2)^2},$$

where $\|\cdot\|_2$ is l_2 norm and α is a hyperparameter used to control the sharpness of CoR. CoR estimates how much the current state is close to expert states or negative states. This value is used as an estimated reward function to the policy network before the discriminator is moderately trained.

3. SEMI-SUPERVISED IMITATION LEARNING

One of the shortcomings of GAIL [6] is that, without expert demonstrations, GAIL may not be adequately trained. Hence, the performance of using GAIL is limited by the quality of expert demonstrations. Also, if the quality of demonstrations given to the algorithm is not consistent with the policy, the learning process becomes slower, the trained agent may not perform as well, or worse, the agent may not learn from them. If we have a set of state-action pairs from expert demonstrations, we can evaluate and give corresponding reliability values to other unlabeled state-action pairs and use them as a sub-optimal dataset. Since each unlabeled demonstration has a different degree of closeness to expert demonstrations, the reliability value for unlabeled state-action pairs has to be continuous, not binary. Therefore, an algorithm using

a continuous reliability value is necessary. We call this estimated reliability value a leverage value in this paper.

We introduce a novel method called semi-supervised imitation learning (SSIL). This method uses both labeled expert demonstrations and unlabeled demonstrations with unknown qualities. Our main contributions are twofold. First, we propose a novel method to evaluate each unlabeled state-action pair based on known expert state-action pairs. Second, we propose a method using continuous leverage values for training an autonomous agent.

3.1 Evaluating Unlabeled Demonstrations

We propose an algorithm that evaluates unlabeled state-action pairs. The variational autoencoder (VAE) is used to evaluate the quality of a state by calculating the reconstruction loss of unlabeled data. If the reconstruction loss is higher than a predetermined threshold, the data pair corresponding to the reconstruction loss can be classified as harmful data and vice versa [14]. However, binary classification cannot accurately evaluate the data because the data quality is not binary, expert or non-expert, but continuous. Therefore, evaluating the quality of each data is essential for higher performance, especially for autonomous driving. We obtain a leverage value by normalizing the reconstruction loss as follows:

$$L_{Recon} = \{r_j | r_j = \|s_{in_j} - s_{out_j}\|_2, 0 \leq j \leq |s_T|\},$$

$$r_{max} = \max_{r_j \in L_{Recon}} r_j, r_{min} = \min_{r_j \in L_{Recon}} r_j,$$

$$l_{T_j} = \frac{r_j - r_{min}}{r_{max} - r_{min}},$$

$$l_T = \{l_{T_j} | \forall j \in \mathbb{N}, \mathcal{K} \leq \mathcal{J} \leq |\sim_{\mathbb{T}}|\}, \quad (3)$$

where L_{Recon} is the reconstruction error set, r_j is the reconstruction error of j th state, s_{in_j} and s_{out_j} are the input state and result of VAE, respectively, l_{T_j} is the leverage value of the j th state, l_T is the set of leverage values, and $|s_T|$ is the number of states in all demonstrations.

While the mixture density network (MDN) [15] and GPR [16] use both state and action data for training, VAE only uses state data as its input, which is a loss of information. To compensate for this issue, we also propose to use two consecutive state data as input to VAE, which is called windowed VAE (windowVAE). By encoding two consecutive state data, the VAE can be seen as encoding the state data and the environment-action interaction. This encoding method allows VAE to encode not only the state data but also the effect of action. The results in Table 2 shows that the leverage values are evaluated as we intended.

3.2 Leverage Constraint Reward

The main problem of the *constraint reward* (CoR) in MixGAIL [12] is that this term can only be used on binary classified trajectories or state-action pairs. If each state-action pair's leverage value is not binary, the given data's potential is not fully utilized. Therefore, we propose a *leverage constraint reward* (LCoR), which can fully exploit the given leverage values with varying degrees for training an agent. The objective function for SSIL can be written as:

$$\min_{\pi \in \Pi} \max_D E_{\pi} [\log(D(s, a)) - \eta \text{LCoR}(s, s_T, l_T)] + E_{\pi_E} [\log(1 - D(s, a))] - \lambda H(\pi), \quad (4)$$

where the notations are the same as that of Section 2.1, except for s_T being a set of states of both labeled and unlabeled demonstrations, and l_T being the matching set of leverage value for s_T calculated from Section 3. LCoR is defined as follows:

$$\text{LCoR}(s, s_T, l_T) = \frac{\sum_{j=1}^{|l_T|} l_{T_j} \left(1 + \frac{\|s - s_{T_j}\|_1}{\alpha}\right)^{-\frac{\alpha+1}{2}}}{\sum_{j=1}^{|l_T|} \left(1 + \frac{\|s - s_{T_j}\|_1}{\alpha}\right)^{-\frac{\alpha+1}{2}}} \quad (5)$$

where $|l_T|$ is the size of a leverage value set, s_{T_j} and l_{T_j} are the state value and its corresponding leverage value of the j th demonstration, and the value of α is larger than -1 . The LCoR value adjusts the reward function inside the GAIL algorithm by giving a reward based on leverage values which are weighted by the distance to the state s . If states with high leverage values are close, LCoR becomes larger. This allows LCoR to operate as a positive reward incentive, as states with high leverage value are closer to states in the labeled expert demonstrations. On the other hand, if neighboring states have low leverage values, LCoR becomes smaller, causing LCoR to operate as a penalty to the overall reward function. Note that this value only impacts the policy update part, not affecting the discriminator update part.

If l_{T_j} is closer to 1, the corresponding state is closer to expert demonstrations while if l_{T_j} is close to 0, the state is far from expert demonstrations. The modified reward function for the policy is as follows:

$$r(s, a) = -\log(D(s, a)) + \eta \text{LCoR}(s, s_T, l_T). \quad (6)$$

Note that only the labeled state-action pairs are used to calculate $-\log(D(s, a))$ in (6), and the labeled and unlabeled dataset are used to calculate the value of LCoR.

During training, the discount factor, η , is decayed. However, each decay causes the target value of the value function to fluctuate. From (6), the target value consists of the discriminator value and the LCoR value. Both values are estimated using neural networks. To mitigate the fluctuation of the value function, η is decayed if the difference of the discriminator value function before and after the RL algorithm update is larger than the difference of LCoR value function before and after the RL algorithm update when both functions are given initial state. The discriminator value function and the LCoR value function are shown below:

$$\begin{aligned} V_\pi(s) &= V_{D_\pi}(s) + \eta V_{C_\pi}(s), \\ V_{D_\pi}(s) &= E_\pi \left[\sum_{k=t}^T -\gamma^{k-1} \log D(s_k, a_k) \middle| s_t = s \right], \\ V_{C_\pi}(s) &= E_\pi \left[\sum_{k=t}^T \gamma^{k-1} \text{LCoR}(s_k, s_E, s_N) \middle| s_t = s \right], \\ \Delta V_C &= V_{C_{\pi'}}(s_0) - V_{C_\pi}(s_0), \\ \Delta V_D &= V_{D_{\pi'}}(s_0) - V_{D_\pi}(s_0), \end{aligned} \quad (7)$$

where V_D is the value function of the discriminator, V_C is the value function of LCoR, s_0 is an initial state, and γ is the decay factor. The η is decayed by multiplying γ when $\Delta V_D > \Delta V_C$.

During training, we estimate discriminator D and policy π using D_ω and π_θ , where ω and θ are weights of the neural networks for D_ω and π_θ , respectively. The training process is explained in Algorithm 1.

Algorithm 1 SSIL

Require: Expert demonstrations $\tau_E \sim \pi_E$, unlabeled demonstrations $\tau_U \sim \pi_U$, number of policies n , expert states set $s_E = \{s_e | \forall (s_e, a_e) \in \tau_E\}$, unlabeled states set $s_U = \{s_u | \forall (s_u, a_u) \in \tau_U\}$, total state set $s_T = s_E \cup s_U$, initial policy and discriminator parameters θ_0, ω_0 , weighting parameter of LCoR $\eta \geq 0$, η 's decay rate ϵ

- 1: Train variational autoencoder by expert demonstrations τ_E
- 2: Feed unlabeled demonstrations to trained network and collect reconstruction loss set U for each demonstration
- 3: Normalize the reconstruction loss set U into l_T , smallest value to 1 and largest value to 0
- 4: **for** $i = 0, 1, 2, \dots$ **do**
- 5: Sample demonstrations $\tau_i \sim \pi_{\theta_i}$
- 6: Update the discriminator parameters from ω_i to ω_{i+1} with the gradient

$$\hat{E}_{\tau_i}[\nabla_\omega \log(D_\omega(s, a))] + \hat{E}_{\tau_E}[\nabla_\omega \log(1 - D_\omega(s, a))]$$

- 7: Take a policy step from θ_i to θ_{i+1} , using the TRPO rule with cost function $\log(D_{\omega_{i+1}}(s, a)) - \eta \text{LCoR}(s, s_T, l_T)$. Specifically, take a KL-constrained natural gradient step with

$$\hat{E}_{\tau_i}[\nabla_\theta \log \pi_\theta(a|s) A_{\pi_\theta}(s, a)] - \lambda \nabla_\theta H(\pi_\theta)$$

where

$$\begin{aligned} A_{\pi_\theta}(s, a) &= -\log D_{\omega_{i+1}}(s, a) + \gamma V_{D_{\pi_\theta}}(s') \\ &\quad - V_{D_{\pi_\theta}}(s) + \eta (\text{LCoR}(s, s_T, l_T) \\ &\quad + V_{C_{\pi_\theta}}(s') - V_{C_{\pi_\theta}}(s)), \end{aligned}$$

- and s' is the next state after state s .
- 8: update $\eta \leftarrow \epsilon \eta$, when $\Delta V_D > \Delta V_C$
 - 9: **end for**
-

4. EXPERIMENTAL RESULTS

4.1 Simulation Experiment

We evaluate the proposed algorithm on a race car simulator called TORCS [13]. In this experiment, a total of 49,200 state-action pairs are collected as demonstrations. The first 1,200 expert state-action pairs are used as the labeled expert demonstrations. The other 48,000 state-action pairs, which consist of 24,000 expert demonstrations and 24,000 negative demonstrations, are used as unlabeled demonstrations. GAIL and MixGAIL are used as baselines for the evaluation of the proposed method. In addition, SSIL using leverage values derived from different algorithms, GPR and MDN, are compared with the proposed algorithm.

4.1.1 Collecting Expert and Negative Demonstrations

The expert demonstrations are collected from a trained PPO [4] network. The negative demonstrations are collected from an under-trained PPO network. The unlabeled data of this part consists of part of expert demonstrations and all of the negative demonstrations. A half of the unlabeled data is expert demonstrations, and the other half consists of negative demonstrations. We changed the amount of labeled data to show the efficiency of SSIL. On average, one expert trajectory consists of 2,400 state-action pairs, and one state-action pair is composed of 29 states and two actions.

4.1.2 Baseline Algorithms for Comparison

For a fair comparison among VAE, windowed VAE, GPR, and MDN, each evaluation method is tested by using it for training SSIL. For all the compared algorithms, the labeled data are used for training. After training, the unlabeled data are given to each algorithm to obtain leverage values using (3). VAE and windowed VAE use the only state for evaluating the unlabeled data.

On the other hand, GPR and MDN use state-action pairs, where the input to the algorithm is the state and the output of the algorithm is the action. In the case of GPR, the uncertainty estimation of each state-action pair can be obtained by calculating the standard deviation of the predicted output. In the case of MDN, the uncertainty estimation of each state-action pair can be calculated by adding the epistemic uncertainty and aleatoric uncertainty [17]. In the case of MDN and GPR, the reconstruction error in the equation is replaced with estimated uncertainty.

4.1.3 Simulation Results

We tested GAIL, MixGAIL, and SSIL three times, each using different random seeds.

The comparative results with SSIL (VAE) and SSIL (windowVAE) are shown in Table 1. We conduct the experiment based on the number of labeled data’s state-action pairs. 100, 300, 600, 1,200 state-action pairs indicates 8.3%, 25%, 50%, 100% of the labeled expert demonstrations, and 0.20%, 0.61%, 1.22%, 2.44% of all demonstrations, respectively.

SSIL vs GAIL First, to show the robustness of the proposed algorithm to the quality of data, SSIL is compared with GAIL using labeled demonstrations and unlabeled demonstrations as input to the discriminator. For GAIL, the unlabeled trajectories are given with the expert trajectories to the discriminator. For SSIL, the labeled trajectories are given to the discriminator and LCoR, while the unlabeled data is given to LCoR only. In addition, the comparison shows how much the labeled data should be given to SSIL in order to outperform GAIL.

100 Expert Demonstrations. The results of SSIL trained with leverage values from VAE (SSIL (VAE)) and windowed VAE (SSIL (windowVAE)) show that when the number of input expert state-action pairs is 100, the results are worse than using mixed input of expert and negative trajectories directly to GAIL. The performance of the experiment on 100 labeled state-action pairs and 48,000 unlabeled state-action pairs is only 9.75% and 57.6% of GAIL’s performance for SSIL (VAE) and SSIL (windowVAE), respectively, as shown in Table 1. The result can be interpreted as; VAE sees data that is not seen during training, calculates unseen data’s reconstruction high, causing a misallocation of the leverage value and worse training results than GAIL.

300 Expert Demonstrations. If 300 labeled state-action pairs and 48,000 unlabeled state-action pairs are used, the SSIL (windowVAE) shows 93.1% of the performance of GAIL, as shown in Table 1. On the other hand, the SSIL (VAE) shows an improvement of 7.27% over GAIL. The experimental result can be interpreted that in this domain, if 0.61% of demonstrations are known as expert demonstrations, SSIL (VAE) shows better performance than GAIL. In the case of SSIL (windowVAE), it can be inferred that more data is needed to encode two consecutive states.

600 & 1200 Expert Demonstrations. If the number of input data is more than 600, the SSIL (VAE) and SSIL (windowVAE) show higher performance and lower standard deviation than GAIL. In the case of 600 input data, SSIL (VAE) and SSIL (windowVAE) shows 41.0% and 35.3% improvements over GAIL, respectively, as shown in Table 1. In the case of 1200 input data, SSIL (VAE) and SSIL (windowVAE) shows 54.4% and 57.3% better performance than GAIL. The comparison

with GAIL with mixed quality demonstrations and SSIL shows that rather than using mixed qualities of data as input to GAIL, evaluating each data and exploiting them during training positively affects the performance.

SSIL vs GAIL with only labeled demonstrations SSIL (VAE) and SSIL (windowVAE) are compared with GAIL using only labeled demonstrations in the second experiment. In the case of 300, 600, and 1200 expert demonstrations, SSIL (VAE) shows 63.5%, 24.8%, and 13.4% improvements over GAIL, respectively, and SSIL (windowVAE) shows 41.9%, 19.8%, and 15.5% improvement over GAIL, respectively. However, in the case of 100 input data, SSIL (VAE) shows 20.9% of GAIL’s performance, while SSIL (windowVAE) shows 23.5% of improvement over GAIL. Therefore, it can be inferred that using unlabeled data and exploiting their leverage values during training positively affects the performance.

SSIL vs MixGAIL In the third experiment, SSIL (VAE) and SSIL (windowVAE) are compared with MixGAIL using labeled and unlabeled demonstrations. For MixGAIL, the labeled data are used for s_E in CoR and discriminator input. On the other hand, we assume there is an oracle that can discriminate the unlabeled data to expert and negative data for the unlabeled data. Therefore, the expert data in the unlabeled data are used for s_E in CoR with the labeled data, and the negative data in the unlabeled data are used for s_N in CoR. This comparison empirically proves that SSIL has the ability to estimate the quality of each demonstration.

SSIL (VAE) results using 300 and 600 expert state-action pairs as input have 17.0%, and 3.11% improvements over MixGAIL, as shown in Table 1. In the case of SSIL (windowVAE), results using 100, 300, and 1200 expert state-action pairs as input show improvements of 204%, 1.55%, and 0.112% compared to MixGAIL, as shown in Table 1. SSIL (VAE) using 100 and 1200 expert state-action pairs, and SSIL (windowVAE) using 300 expert state-action pairs show 51.4%, 87.8% and 98.3% of the performance of corresponding results of MixGAIL, as shown in Table 1. From the results, it can be inferred that SSIL has a slight advantage over MixGAIL when each state-action data has different qualities. In addition, the comparison results show that the discrimination process using VAE is reliable.

SSIL (VAE), SSIL (windowVAE) vs SSIL (MDN), SSIL (GPR) In the fourth experiment, the different algorithms for evaluating the unlabeled data are compared. Leverage value evaluation methods, VAE and windowed VAE, are compared with GPR and MDN. The comparison is shown by the performance of SSIL trained with the leverage value calculated from each evaluation method. This experiment verifies the validity of the proposed leverage value quality assessments.

The performances of SSIL trained using leverage values from GPR (SSIL (GPR)) and MDN (SSIL (MDN)) were relatively poor compared to SSIL (VAE) and SSIL (windowVAE), except for the case of SSIL (MDN) using 300 expert state-action pairs for training. Especially, SSIL (GPR) showed worse performance than SSIL (windowVAE) in every case. The performance gain of SSIL (windowVAE) over SSIL (GPR) varies from 2.14% to 537%. SSIL (MDN) showed 0.661% to 85.4% better performance than SSIL (VAE) in most cases except for the case using 600 state-action pairs to input, as shown in Table 1. However, except for the case using 300 state-

labeled state-actions	GAIL	GAIL expert	MixGAIL	SSIL (GPR)	SSIL (MDN)	SSIL (VAE)	SSIL (windowVAE)
100	0.564/0.077	0.263/0.057	0.107/ 0.007	0.051/0.044	0.102/0.008	0.055/0.017	0.325/0.161
300	0.564/0.077	0.370/0.169	0.517/0.009	0.514/ 0.021	0.609 /0.085	0.605/0.125	0.525/0.071
600	0.564/0.077	0.637/0.129	0.771/0.054	0.643/0.118	0.614/0.102	0.795 /0.054	0.763/0.002
1200	0.564/0.077	0.768/0.095	0.886/ 0.001	0.880/0.020	0.873/0.014	0.871/0.013	0.887/0.022

Table 1.: The scaled performance mean (left) and standard deviation (right) of GAIL, MixGAIL, and SSIL on the TORCS domain. The categories written in bold are the proposed algorithm and the figure of the algorithm that has the best performance and the most negligible standard deviation in each comparison. GAIL expert indicates the GAIL trained with only labeled demonstrations. Each figure is scaled to the perfect score for each agent.

Algorithm	0-25%	25-50%	50-75%	100% (expert)
LGPR (expert)	0.2593	0.4258	0.2202	0.1738
LGPR (negative)	0.2341	0.4379	0.3890	0.3682
VAE	0.7185	0.7595	0.7639	0.8533
windowVAE	0.6467	0.6906	0.7273	0.8168

Table 2.: The average leverage values from different data quality calculated by different algorithms. The leverage value calculation algorithms are trained with four different qualities of the dataset. Each dataset consists of different qualities of data; 0 to 25% of maximum reward, 25% to 50% of maximum reward, 50% to 75% of maximum reward, and 100% of maximum reward (expert data). The parenthesis next to the LGPR algorithm indicates the quality of the data.

action pairs to input, SSIL (windowVAE) shows the superior result to SSIL (MDN), with performance improvement varying from 1.6% to 219%, as shown in Table 1. Considering the standard deviation results shown in Table 1, the best algorithm among the three evaluation methods is SSIL (VAE), as we think that SSIL (VAE) is the sweet spot between standard deviation and performance. Therefore, one can claim that the proposed leverage value quality assessment is adequate for the proposed algorithm.

Leverage Value Comparison Finally, the correlations of the leverage value and the quality of demonstrations are shown using different qualities of negative data. The qualities of negative data are calculated based on the reward function of the environment. We compare the leverage value of different qualities of negative demonstrations to show that the leverage value reflects the quality of data. The number of labeled demonstrations and negative demonstrations is 600 and 2500, respectively. Note that the number of negative demonstrations is larger than the labeled demonstrations. We also compare the leverage value calculated from LGPR [9] with the leverage values calculated from the proposed method in order to show that the LGPR method cannot calculate the leverage value accurately when the negative demonstrations are dominant. In contrast, the proposed method can discriminate when the negative data is not dominant. The quality of the negative data is split into three parts: 0 to 25% of the maximum score, 25% of the maximum score to 50% of the maximum score, and 50% of the maximum score to 75% of the maximum score. The expert data is also used to compare this algorithm for the baseline. This comparison verifies the positive correlation between quality and leverage value.

The results show that the leverage value and the quality of demonstrations have a positive correlation in both VAE and windowed VAE, as shown in Table 2. Both results of the evaluation algorithms show that the evaluation methods (VAE, windowVAE) validly evaluate the reliability of unlabeled demonstrations. On the other hand, the LGPR assigns higher leverage value for negative demonstrations, as the negative demonstrations are dominant.

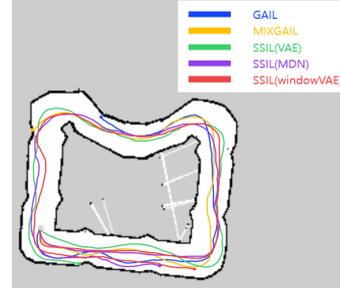


Fig. 1.: Bird-eye view of the track used for the experiment and its experimental results.

From this experiment, we can infer that while the LGPR cannot evaluate the negative demonstration and miscalculates the reliability of the data, the proposed method can well discriminate and evaluate the negative value when the negative demonstrations are dominant.

4.2 Real-World RC Car Experiment

We apply the proposed method to learn an RC Car controller in the real world to empirically prove that not only the reconstruction loss is an effective method of discriminating and labeling unlabeled data, but also the performance of the proposed method stands out among the compared methods. This experiment’s reward function is proportional to how much the RC car has traveled along the track. Since we give the RC car a constant speed, we used the timestep as the performance measure.

4.2.1 Collecting Expert and Negative Demonstrations

Expert demonstrations are collected by controlling the RC car manually. The expert demonstrations consist of 4,010 state-action pairs, which are ten trajectories. We use 2,005 state-action pairs, which is a half of the expert demonstrations, as labeled state-action pairs. The other 2,005 state-action pairs are given as unlabeled data with other 1,170 negative demonstrations by selecting 250-300 state-action pairs each from four different qualities of under-trained policies. The negative demonstrations were collected from under-trained GAIL. These demonstrations include episodes where RC car collides with walls. Note that to show the robustness of our algorithms in the real world, the track used to collect data is different from the track used to train the algorithm.

4.2.2 RC Car Experimental Results

GAIL, MixGAIL, and SSIL are trained three times, each using different random seeds. The plotted graph is the average rewards per training step and standard deviations.

The numerical results are shown in Table 3. The results of SSIL with leverage values calculated from VAE are 9.02% higher than MixGAIL and 15.8% higher than the original GAIL, but its average performance was worse than that of SSIL trained with MDN leverage values. As shown in Table 3, the SSIL using VAE has a lower average standard deviation than GAIL and SSIL trained with

evaluation metric	GAIL	MixGAIL	SSIL (MDN)	SSIL (VAE)	SSIL (windowVAE)
maximum score	0.679	0.721	0.813	0.786	0.808
standard deviation	0.086	0.051	0.089	0.050	0.051

Table 3.: The scaled performance of GAIL, MixGAIL and SSIL(MDN, VAE, windowed VAE) on RC car domain. The expert score of the environment is 400, and each algorithm’s scores are normalized by provided expert score.

MDN leverage values and shows a similar average standard deviation to MixGAIL. The windowed VAE version of SSIL also shows higher performance than GAIL and MixGAIL and similar performance to SSIL trained with MDN leverage value. The windowed VAE version of SSIL shows performance gains of 19% and 12.1% compared to GAIL and MixGAIL, respectively, and has a lower standard deviation than GAIL and SSIL trained with MDN leverage values, as shown in Table 3.

We would like to emphasize that the standard deviation of GAIL is the second-highest since the algorithm is tested on a similar but different track from the track used to collect the training data. MixGAIL showed worse performance than all SSIL methods because the unlabeled data were not labeled correctly, causing it to be seen as negative, although it is actually expert data. This process causes the *constraint reward* to be contaminated with wrong information. On the other hand, SSIL trained with VAE and windowed VAE safely discriminates and labels expert data, causing the standard deviation during training to lower and show better performance. This labeling process allows the proposed algorithm to gain valid information about each unlabeled data from the *leverage constraint reward*, allowing our agent to explore the environment safely. Also, SSIL trained with MDN shows a similar maximum score to windowed VAE but has a much higher standard deviation, from which we can infer that MDN leverage value is less stable than windowed VAE value.

5. CONCLUSION

This paper has considered handling unlabeled, perhaps sub-optimal, demonstrations in imitation learning. We have proposed a method to discriminate the demonstration at the state-action pair level. Moreover, we have proposed a new reward incentive called *leverage constraint reward* that is suited for continuous evaluation of demonstrations, which is called leverage. As shown in Table 2, the estimated leverage values correspond to the quality of data. The SSIL has been able to show better performance than the baseline algorithms in TORCS simulations, when the labeled expert demonstrations are less than 1% of the total demonstrations. In the RC car experiment, SSIL has shown a significant performance improvement over GAIL and MixGAIL in terms of the travel distance when given similar amounts of data.

REFERENCES

- [1] W. Zhao, M. A. Reinthal, D. D. Espy, and X. Luo, “Rule-based human motion tracking for rehabilitation exercises: Realtime assessment, feedback, and guidance,” *IEEE Access*, vol. 5, pp. 21382–21394, 2017.
- [2] W. Xiao, N. Mehdipour, A. Collin, A. Y. Bin-Nun, E. Frazzoli, R. D. Tebbens, and C. Belta, “Rule-based optimal control for autonomous driving,” in *Proc. of the ACM/IEEE 12th International Conference on Cyber-Physical Systems*, May 19-21, 2021.
- [3] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz, “Trust region policy optimization,” in

- Proc. of the 32nd International Conference on Machine Learning*, July 6-11, 2015.
- [4] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *CoRR*, vol. abs/1707.06347, 2017.
- [5] G. Dulac-Arnold, D. J. Mankowitz, and T. Hester, “Challenges of real-world reinforcement learning,” *CoRR*, vol. abs/1904.12901, 2019.
- [6] J. Ho and S. Ermon, “Generative adversarial imitation learning,” in *Proc. of the Advances in Neural Information Processing Systems 29*, December 5-10, 2016.
- [7] K. Lee, S. Choi, and S. Oh, “Maximum causal Tsallis entropy imitation learning,” in *Proc. of the Advances in Neural Information Processing Systems 31*, December 3-8, 2018.
- [8] Y. Li, J. Song, and S. Ermon, “InfoGAIL: Interpretable imitation learning from visual demonstrations,” in *Proc. of the Advances in Neural Information Processing Systems 30*, December 4-9, 2017.
- [9] S. Choi, K. Lee, and S. Oh, “Robust learning from demonstration using leveraged Gaussian processes and sparse-constrained optimization,” in *Proc. of the IEEE International Conference on Robotics and Automation*, May 16-21, 2016.
- [10] S. Choi, K. Lee, and S. Oh, “Robust learning from demonstrations with mixed qualities using leveraged Gaussian processes,” *IEEE Transactions on Robotics*, vol. 35, no. 3, pp. 564–576, 2019.
- [11] S. Choi, K. Lee, and S. Oh, “Scalable robust learning from demonstration with leveraged deep neural networks,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 24-28, 2017.
- [12] G. Lee, D. Kim, W. Oh, K. Lee, and S. Oh, “MixGAIL: Autonomous driving using demonstrations with mixed qualities,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 24, 2020 - January 24, 2021.
- [13] D. Loiacono, L. Cardamone, and P. L. Lanzi, “Simulated car racing championship: Competition software manual,” *CoRR*, vol. abs/1304.1672, 2013.
- [14] J. An and S. Cho, “Variational autoencoder based anomaly detection using reconstruction probability,” *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.
- [15] C. Bishop, “Mixture Density Networks,” Aston University, Birmingham, U.K., Tech. Rep. NCRG/94/004, February, 1994.
- [16] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [17] S. Choi, K. Lee, S. Lim, and S. Oh, “Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling,” in *Proc. of the IEEE International Conference on Robotics and Automation*, May 21-25, 2018.