

Model-free Deep Reinforcement Learning for Urban Autonomous Driving

Jianyu Chen*, Bodi Yuan* and Masayoshi Tomizuka

Abstract—Urban autonomous driving decision making is challenging due to complex road geometry and multi-agent interactions. Current decision making methods are mostly manually designing the driving policy, which might result in sub-optimal solutions and is expensive to develop, generalize and maintain at scale. On the other hand, with reinforcement learning (RL), a policy can be learned and improved automatically without any manual designs. However, current RL methods generally do not work well on complex urban scenarios. In this paper, we propose a framework to enable model-free deep reinforcement learning in challenging urban autonomous driving scenarios. We design a specific input representation and use visual encoding to capture the low-dimensional latent states. Several state-of-the-art model-free deep RL algorithms are implemented into our framework, with several tricks to improve their performance. We evaluate our method in a challenging roundabout task with dense surrounding vehicles in a high-definition driving simulator. The result shows that our method can solve the task well and is significantly better than the baseline.

I. INTRODUCTION

A highly intelligent decision making system is crucial for urban autonomous driving with dense surrounding dynamic objects. It must be able to handle the complex road geometry and topology, complex multi-agent interactions, and accurately follow the high level command such as routing information.

Current autonomous driving decision making systems are focusing on the non-learning model-based approaches, which often requires to manually design a driving policy [12], [21]. Although it might not be difficult to design a qualified driving policy with the help of human prior knowledge, the manually designed policy could suffer from several weaknesses: 1) Accuracy: the predefined motion heuristics and models will lead to bias, especially for highly interactive environments. 2) Generality: for different scenarios and tasks, the model might need to be redesigned manually.

Recent advances in machine learning enables the possibility for learning based approaches for autonomous driving decision making. The most popular approach is imitation learning, which can learn a driving policy automatically from expert driving data [3], [7], [2], [6]. However, there are some shortcomings for imitation learning: 1) It needs to collect a huge amount of expert driving data in real-world and in real-time, which can be costly and time consuming. 2) It can only

learn driving skills that are demonstrated in the dataset. This might lead to serious safety issues because expert drivers generally do not provide dangerous demonstrations so the autonomous vehicle cannot learn how to deal with those dangerous cases. 3) Since the human driver experts act as the supervision for learning, it is impossible for an imitation learning policy to exceed human-level performance.

Combined with deep learning techniques, reinforcement learning (RL) has brought a series of breakthroughs in recent years. Agents trained with deep reinforcement learning achieves super-human-level performance in Atari game playing [19], [20], go playing [25], [27], and complex strategic games [30], [1]. With reinforcement learning, a policy can be learned automatically without any expert data. It can simulate various kinds of different cases, including some dangerous ones. It is also possible to achieve better performance than a human expert. Thus a good alternative to imitation learning for autonomous driving decision making is to use deep reinforcement learning.

However, there are not many successful applications for deep reinforcement learning in autonomous driving, especially in complex urban driving scenarios. This is due to: 1) Most of the methods directly use front view image as the input and learn the policy end-to-end. The extremely complex high dimensional visual features dramatically enlarge the sample complexity for learning. 2) Deep reinforcement learning is a fast evolving research area, but its application to autonomous driving has lag behind. Most researchers are still using basic deep RL algorithms such as deep Q network, which is not able to solve some complex problems. Much more powerful deep RL algorithms were developed in recent years but few of them have been applied to autonomous driving tasks.

In this paper, we design a specific input representation to reduce the sample complexity instead of directly using the front view image. We then use visual encoding to capture the low-dimensional latent states for urban autonomous driving tasks, which makes the problem more tractable for reinforcement learning. Several state-of-the-art model-free deep RL algorithms are implemented to learn a driving policy in a complex roundabout scenario with multiple surrounding vehicles running. Several tricks are developed to improve the performance of the algorithms, including modified exploration strategies, frame skip, network architectures and reward designs. Final result shows that our method can robustly learn a driving policy that is able to navigate through the complex urban driving scenario.

* indicates equal contribution

J. Chen, B. Yuan and M. Tomizuka are with Department of Mechanical Engineering, University of California, Berkeley, CA94720, USA. Corresponding to jianyuchen@berkeley.edu

This work was supported by Denso International at America

II. RELATED WORKS

With recent progresses in deep learning, learning-based approaches find their applications in autonomous driving, both in real world and in simulation. NVIDIA used a deep convolutional neural network to learn a lane following policy end-to-end from the front view image [3], [4]. Waymo also used imitation learning to learn a urban driving policy from a huge amount of human driver data [2]. Based on CARLA, an open-source simulator for autonomous driving research, [7], [6] applied deep imitation learning to learn a policy to navigate through a complex virtual urban environment.

Researchers have also tried reinforcement learning approaches for autonomous driving. Wolf et al. [32] used a Deep Q Network to learn to steer an autonomous car in simulation. The action space is discrete and only allows coarse steering angles. Lillicrap et al. [17] developed a continuous control deep reinforcement learning algorithm which is able to learn a deep neural policy to drive the car on a simulated racing track. Sallab et al. [22] proposed a deep RL framework for autonomous driving and applied it to the lane following task in simulation. Chen et al. [5] proposed a hierarchical deep RL approach which is able to solve some complex temporal delayed problems such as traffic light passing. Nonetheless, the above approaches were developed either for simple scenarios without complex road geometry and multi-agent interaction, or used manually designed feature representations or policy models.

Bird-view representation for autonomous driving decision making is proposed recently to reduce the complexity of visual features [8], [2], [6]. They generally compress information of the map and objects to a rasterized image. Direct RGB images are very high dimensional, one can use variational auto-encoder (VAE) [16], [9] for dimension reduction, which learns a low dimensional latent representation of the image via an unsupervised style.

Model-free deep reinforcement learning is a rich research area and is evolving rapidly. Related algorithms range from the Q learning based approaches such as DQN [19], [20] and double DQN [29], the actor-critic approaches such as A3C [18], DDPG [17] and TD3 [11], the policy optimization approaches such as TRPO [23] and PPO [24], and the maximum entropy approaches such as soft actor-critic [13], [14].

III. PROBLEM FORMULATION

Directly applying reinforcement learning algorithms on the raw sensor input can hardly work well on the urban autonomous driving problem. We thus proposed a framework to reduce the complexity of the problem, making it possible to be solved well by the current model-free deep reinforcement learning techniques.

Fig.1 shows the proposed framework. Our intelligent driving agent receives perception and routing information from the driving environment. It then processes the information into a bird-view image as the input representation. The image is then encoded into low dimensional latent states. Reinforcement learning algorithms are then adopted to train

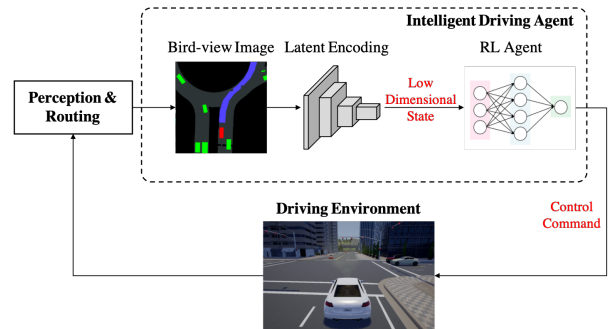


Fig. 1: Framework overview of the our system. The agent takes information from the perception and routing modules, generates a bird-view image and encodes it to low-dimensional latent states. Reinforcement learning is then applied to learn a policy to generate the correct control command.

a deep neural policy which takes the encoded states as input and generates the control command such as acceleration and steering angle.

A. Input Representation

The input to the autonomous driving agent needs to cover enough information for urban driving scenarios. A straight forward input representation is to directly use the raw sensor data such as front view image. However, the raw sensor data contains extremely high dimensional information such as appearances and textures of the roads and objects, weather conditions, and light conditions. In order to obtain good generalization, the dataset must cover enough data for each dimension of the raw sensor information.

With the help of a perception module, the above complexity can be reduced largely. For example, we can use object detection to get the position, heading and velocity of each surrounding object; we can obtain lane related features such as longitudinal and lateral distance to a specific lane marking; we can also represent high level routing information as a set of trajectory points. However, although this kind of representation has been widely used with model-based policies, it brings problems for learning-based policies: 1) The number of surrounding objects is varying; 2) The information about road geometry and topology is highly structured. These information is extremely difficult to be transformed to a tensor with fixed shape, which can be used as the input to a learning based neural network policy.

To alleviate this problem, we proposed to convert the output of perception module (object detection and localization), as well as the routing information, to a bird-view representation that applied as the input to our policy. This bird-view representation contains enough information for decision making, and is a big simplification of the complex visual and spatial information of raw sensor data. As shown in Fig.2, our bird-view input representation is composed of the following four parts:

1) **Map:** The map contains information of road geometry. Here we render the drivable roads as gray polygons and other undrivable area part to be black.




Fig. 2: *Input Representation of our framework. The bird-view observation combines information of map, routing, historical detected objects and historical ego states.*

2) **Routing:** The routing information is calculated by a route planner. It contains information of a sequence of way points that the ego vehicle should follow. It is rendered as a thick blue polyline.

3) **Historical Detected Objects:** The historical bounding boxes of detected surrounding objects (e.g, vehicles, bicycles, pedestrians) in a past sliding window are rendered as green boxes, with reduced level of brightness meaning earlier time-steps.

4) **Historical Ego States:** Similar to the detected objects, the historical ego states are represented as boxes with reduced brightness. The color of the boxes are set red.

The final bird-view image is rendered with pixel size 256×256 and resized to 64×64 , and is always aligned with the ego vehicle view. The actual size of the field of view is $(40m, 40m)$, where the ego vehicle is positioned at $(20m, 8m)$.

B. Latent State Encoding

Even after we replace the complex raw sensor information with processed information, the input is still a high dimensional bird-view image. It will be extremely hard to learn a good policy with this high dimensional input. Meanwhile, the complex input can make it easy to get over-fitted during reinforcement learning process.

In order to further reduce the input complexity, we propose to use variational auto-encoder (VAE) [16] to learn a low dimensional latent representation. VAE is composed of two parts, an encoding network $q_\phi(\mathbf{x}|\mathbf{o})$ which encodes the original high dimensional observation \mathbf{o} to a low dimensional latent state \mathbf{x} , and a decoding network $p_\theta(\mathbf{o}|\mathbf{x})$ which decodes \mathbf{x} to \mathbf{o} . Here ϕ and θ denote the parameters of the encoder and the decoder, respectively. To obtain the parameters, the following objective needs to be maximized:

$$\mathcal{L}(\phi, \theta, \mathbf{o}) = D_{KL}(q_\phi(s|\mathbf{o})||p_\theta(s)) - \mathbb{E}_{q_\phi(s|\mathbf{o})}(\log p_\theta(\mathbf{o}|s)) \quad (1)$$

where D_{KL} represents kullback-leibler (KL) divergence. The prior distribution is usually a multivariate Gaussian $p_\theta(s) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. Some samples of the encoding results are shown in Fig.3 where the first row shows the original input images, and the second row shows the reconstructed images. We can see the reconstructed images are pretty close to the




Fig. 3: *Examples of the latent state encoding. The first row shows the original input images. The second row shows the reconstructed images based on the encoded latent states.*

original images, both for the road geometry and objects. This indicates that the encoded low dimensional latent state successfully preserves the core information of the original image input.

C. Reinforcement Learning

With access to the state s and reward function r , the objective of reinforcement learning is to find the optimal policy π^* which optimize the expected future total rewards:

$$R(\pi, r) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right] \quad (2)$$

where γ is the discount factor of the reward. The solution $\pi^*(a_t|s_t)$ can then be used as the controller of our agent, which takes the current input state s_t and output the control command a_t to be applied to the vehicle. The next section will introduce how we can obtain the policy.

IV. ALGORITHMS

Three state-of-the-art model-free deep reinforcement learning algorithms are applied in our framework to learn the driving policy. We will introduce them briefly in this chapter.

A. Double Deep Q-Network (DDQN)

The most representative model-free reinforcement learning is Q learning [31], which is based on an estimation of the Q value $Q_\pi(s, a)$, defined as the expected future total rewards when taking action a at state s and then follow the policy π thereafter:

$$Q_\pi(s, a) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \middle| s_0 = s, a_0 = a \right] \quad (3)$$

The optimal Q value is $Q^*(s, a) = \max_\pi Q_\pi(s, a)$ and the optimal policy is the highest valued action $\pi^*(s) = \arg \max_a Q^*(s, a)$. Note that the action set for Q learning must be finite discrete. The optimal Q value can be learned using temporal difference (TD) learning [28].

When using a deep neural network to parametrize the Q value as $Q(s, a; \theta)$, where θ are parameters of the neural network, the state can be extended to high dimensional and continuous. The corresponding algorithm is called deep Q network (DQN) [19], [20]. DQN builds a replay buffer \mathcal{D} , and defines an additional target Q network with parameters θ' , besides the online Q network with parameters θ . During learning, the transition pairs (s, a, r, s') are stored into the

replay buffer and uniformly sampled a mini-batch at each step. Then the learning targets are defined as:

$$y = r + \gamma \max_{a'} Q(s', a'; \theta') \quad (4)$$

Supervised learning is applied to minimize the objective $\min_{\theta} \sum (y - Q(s, a; \theta))$ and θ is copied to θ' every τ steps.

The original DQN algorithm has the problem of overestimating the Q values. Thus Double DQN (DDQN) [29] is proposed to avoid the overoptimistic value estimates. The main idea is to separate action selection from the value estimation. DDQN thus obtains the action using the online Q network, but estimates the Q value using the target Q network. The algorithm is mostly same with DQN but replaces the learning target (4) with:

$$y = r + \gamma Q\left(s', \arg \max_{a'}(s', a'; \theta); \theta'\right) \quad (5)$$

B. Twin Delayed Deep Deterministic Policy Gradient (TD3)

DQN and DDQN can only solve problems with finite discrete action spaces. However most environments need continuous actions. For continuous control, a policy network π_{ϕ} is introduced, and it can be optimized through the deterministic policy gradient algorithm (DPG) [26], known as actor-critic. The policy can be optimized by taking gradient steps with respect to the expected future total rewards:

$$\nabla_{\phi} J(\phi) = \mathbb{E}_{\pi} \left[\nabla_a Q_{\pi}(s, a)|_{a=\pi(s)} \nabla_{\phi} \pi_{\phi}(s) \right] \quad (6)$$

Combined with deep neural networks, the deep deterministic policy gradient (DDPG) algorithm [17] is proposed. Similar to DQN algorithm, DDPG also builds both an online and a target Q network with parameters θ and θ' , as well as an online and a target policy network with parameters ϕ and ϕ' . Transitions are stored into a replay buffer and mini-batches are sampled at each step. To update the Q network, the Q value targets are set as:

$$y = r + \gamma Q(s', \pi_{\phi'}(s'); \theta') \quad (7)$$

The policy network is updated using the sampled policy gradient:

$$\nabla_{\phi} J(\phi) = \frac{1}{N} \sum_i \nabla_a Q(s, a; \theta)|_{s=s_i, a=\pi_{\phi}(s_i)} \nabla_{\phi} \pi_{\phi}(s)|_{s_i} \quad (8)$$

The target Q and policy networks are updated using temporal difference with respect to the online networks:

$$\begin{aligned} \theta' &\leftarrow \tau\theta + (1 - \tau)\theta' \\ \phi' &\leftarrow \tau\phi + (1 - \tau)\phi' \end{aligned} \quad (9)$$

DDPG can suffer from the function approximation errors which lead to overestimated values and suboptimal policies. Thus the twin delayed deep deterministic policy gradient (TD3) algorithm [11] is proposed to address the problem. It borrows the idea of double Q learning to build an additional Q network, and take the minimum value between the pair of Q networks when setting the target Q value. It also suggests to delay the target network update steps.

C. Soft Actor Critic (SAC)

Although the above methods work well in a range of challenging decision making tasks, they usually suffer from high sample complexity and brittle convergence property, which leads to extraordinary hyper-parameter tuning. The soft actor critic (SAC) algorithm [13], [14] is proposed to alleviate this problem based on the theory of maximum entropy reinforcement learning. In this framework, the policy needs to maximize both the expected rewards and the entropy:

$$\max_{\pi} J(\pi) = \sum_{t=0} \mathbb{E}_{\pi} [r(s_t, a_t) + \alpha \mathcal{H}(\pi(\cdot|s_t))] \quad (10)$$

where α is the weight of the entropy term. This modified reward function defines a so-called soft Q function and the corresponding soft Bellman backup operator:

$$\mathcal{T}^{\pi} Q(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} [V(s_{t+1})] \quad (11)$$

where

$$V(s_{t+1}) = \mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t) - \log \pi(a_t|s_t)] \quad (12)$$

is the soft value function. It can be shown that applying the soft bellman backup will converge to the optimal soft Q value. SAC defines a soft Q network Q_{θ} , a soft value network Q_{θ} , and a policy network π_{ϕ} , as well as an additional target value network $V_{\psi'}$. Similar to DQN and DDPG, the Q network and value network can be updated using supervised learning by defining the following targets according to (11) and (12):

$$\hat{V}(s_t) = E_{a_t \sim \pi_{\phi}} [Q_{\theta}(s_t, a_t) - \log \pi_{\phi}(a_t|s_t)] \quad (13)$$

$$\hat{Q}(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} [V_{\psi'}(s_{t+1})] \quad (14)$$

The policy is a stochastic neural network $f_{\phi}(\varepsilon_t; s_t)$ where ε_t is sampled from some fixed distribution such as a multivariate normal distribution. Then policy network can be optimized by applying policy gradient to the expected future rewards:

$$J_{\pi}(\phi) = E_{s_t, \varepsilon_t} [\log \pi_{\phi}(f_{\phi}(\varepsilon_t; s_t)|s_t) - Q_{\theta}(s_t, f_{\phi}(\varepsilon_t; s_t))] \quad (15)$$

V. EXPERIMENTS

A. Simulation Environment and scenario

We train and evaluate our method on the CARLA simulator [10], which is a high-definition open-source simulator for autonomous driving research. Fig.4 (a) shows a sample view of the driving simulation environment we use. It includes various urban scenarios such as intersection and roundabout. The range of the map is $400m \times 400m$, containing around $6km$ total length of roads.

We choose the most challenging scenario, the central roundabout in this map for training and testing our method. As shown in Fig.4 (b), the task is to start from an entrance of the roundabout, safely and efficiently enter the roundabout, pass through the first two exits, drive out to the desired exit and reach the final goal point. There are 100 vehicles initially sampled in the whole map, and nearly half of them are sampled around the roundabout, making the traffic quite busy. The sampled vehicles will randomly choose a direction when encountering with multiple choices of routes, then follow the route, and slow down if there are front vehicles.




Fig. 4: The simulation environment and experiment scenario. Left is a sample view of the simulation. Right is a bird-view of our roundabout task scenario.

B. Network Architectures

We train our variational auto-encoder (VAE) on 50k bird-view images, which is generated by using a simple controller to drive around the roundabout. Note that we do not need any labels in the training dataset, such as the action commands or ground truth positions of vehicles, but only the raw bird-view images. We also do not require the controller to drive well, it does not matter if the vehicle drives out of lane or collides with other objects. The VAE model has 4 conv-layers of 3×3 kernel size, with 32, 64, 128, and 256 channels separately. The stride is set to 2. Each conv-layer is followed by ReLU activation. The latent space layer of size 64 is then fully connected to the last conv-layer. The VAE model is trained from scratch using Adam optimizer [15], with learning rate 10^{-4} for 100 epochs. Some examples of trained VAE results are given in Figure 3.

During the reinforcement learning process, all networks have the same conv-layers and first dense layer copied from the pretrained VAE model, denoted as visual encoding layers here. To keep the visual encoding stable, the weights of the visual encoding layers are fixed without updating. Now we introduce the remaining layers of the networks used in the three reinforcement learning algorithms as follows:

1) *Double Deep Q-Network (DDQN)*: DDQN has one Q network and a target Q network with the same architecture. After the visual encoding layers, the network is then followed with 5 dense layers, with hidden layers ranging from 256 to 32 nodes. The output size is equal to the number of the possible actions, each representing the Q value of the corresponding action. The networks are trained using Adam optimizer with learning rate of 10^{-3} .

2) *Twin Delayed Deep Deterministic Policy Gradient (TD3)*: TD3 has a policy network and two Q networks. After the visual encoding layers, all of these networks have 4 dense layers, with 64, 200 and 20 hidden units separately, followed by batch normalization and leaky ReLU for each layer. The negative slope of leaky ReLU is set to 0.01. The last dense layer of the policy network, corresponding to the action output, has 2 units, with a tanh activation function to limit the action range. The Q networks are a little different, their latent states from the visual encoding are first concatenated with the corresponding action, and the last layer has only one unit,

denoting the Q value. Also, there is no batch normalization and no tanh activation used in Q networks. All the networks are trained using Adam optimizer. The policy network and Q networks are trained with learning rate of 10^{-4} and 10^{-3} respectively.

3) *Soft Actor Critic (SAC)*: The soft actor critic has two Q networks, one value network, and a policy network. The Q and value networks has the same architecture with one output unit. The visual encoding layers are followed by 5 dense layers with hidden units ranging from 256 to 32. The policy network's architecture is the same except that for the last layer, it splits to two branches. This is because the policy network of SAC is stochastic. The first branch represents the mean of the action and the second branch represents its variance. All networks are trained using Adam optimizer with learning rate of 3×10^{-4} .

C. Implementation Details

To improve the performance on the proposed task, we designed specific reward functions and added several tricks such as frame skip and new exploration strategies.

1) *Rewards Design*: After testing a variety of reward combinations, we came up with a five-term reward function which works well. The reward function is given below:

$$r = r_v + r_\alpha + r_c + r_o + c. \quad (16)$$

where r_v is the term encouraging moving forward to make progress. It is set equal to the ego vehicle's speed. Moreover, we set $r_v \leftarrow 10 - r_v$ if $r_v > 5$ to penalize exceeding speed limit. r_α is the term penalizing the magnitude of steering angle α to improve driving smoothness, where $r_\alpha = 0.5 * \alpha^2$. r_c is the term penalizing collision with other surrounding vehicles, where $r_c = -10$ if there is a collision, otherwise $r_c = 0$. r_o is the term penalizing running out of the lane, where $r_o = -1$ if the distance d between the ego vehicle the routing baseline is larger than $2m$, otherwise $r_o = 0$ (note that we also tried to set r_o continuously, such as linear to the distance d or d^2 , but it turns out the discrete one works better). The last term c is a small constant set to -0.1 , which was used to penalize the ego vehicle for stopping still.

2) *Frame Skip*: We use the frame skip trick during training, where we keep the action unchanged for k consecutive frames. In other words, during the training, each action made by the ego vehicle will last k frames until the new action starts to be effective. This technique heavily reduces the training complexity, one can regard it as reducing the search depth by a factor of k .

However, k shouldn't be set too large. If k is too large, the correct action space might be too small or even not existed. For example, when we set $k = 10$, it is difficult for the ego vehicle to make turns successfully. In our experiments, we use $k = 4$.

3) *Exploration Strategies*: We proposed different exploration strategies in our implementation. For DDQN, we use a different kind of epsilon greedy strategy. Generally, epsilon greedy chooses the random actions from uniform distribution. However, the uniform distribution contains no heuristic information, which is inefficient. We thus propose

to sample the random actions according to each action's Q values, where actions with bigger Q values will be sampled more frequently. This strategy introduces the information of expected future rewards and is thus more efficient.

For TD3, we added a zero-mean Gaussian action noise with a specific variance σ :

$$\sigma = \delta * \lambda_t * \lambda_d * \lambda_p, \quad (17)$$

where δ is set to 0.5 and 0.1 for acceleration and steering, respectively. $\lambda_t = \max(0.5, 1 - t/T)$, where t is the current training step and T is the total decay time set to $100k$. $\lambda_d = \max(1 - t/T, 0.2 + t_p/T_p)$, where t_p is the current path step in a single exploration path, and T_p is the maximum path length which is set to 500. We also introduce a periodic coefficient, $\lambda_p = 1 + \sin(5\pi * t/T + \pi/2)$.

For SAC, the exploration strategy is actually included in the stochastic policy, which is learned and adapted during training, thus we do not need to add any specific exploration strategy for it.

VI. RESULTS

We evaluate our proposed framework using the metrics of average return, and success rate for entering and passing through the roundabout. The approaches are evaluated under two cases: without and with adding surrounding vehicles.

A. Roundabout without Vehicles

We first test our different RL approaches in the roundabout scenario without adding any surrounding vehicles. We also implemented a DDPG approach and trained directly on the front view images, with the same input sizes, network architectures, exploration strategies and hyper-parameters as in [17]. This DDPG approach is set as the baseline.




Fig. 5: Learning curves of different RL approaches applied in the roundabout scenario without adding vehicles. The shaded region represents half a standard deviation of returns with $1k$ evaluation steps. Curves are smoothed for visual clarity.

Fig. 5 shows the average returns for each approach. Rough estimations of returns for each checkpoint (entrance, first/second/desired exit, goal point) are drawn as horizontal dashed lines. As we can see, with the implementation under our proposed framework, SAC algorithm has the best performance, the trained ego vehicle can pass through the roundabout and reach the goal point (refer to Fig. 4) in

most of times after 100 epochs. Although DDQN and TD3 make slow progress after the ego vehicle gets around the first two exits (refer to Fig. 4), all of the three algorithms perform well in the roundabout before the second exit. Due to our limited computation sources, we stopped at 140 epochs, when our SAC policy can already reach the final goal point stably. We believe that given longer training time and more computational resource, these algorithms can perform even better. On the other hand, the baseline has very low average returns and can barely learn anything. Actually, the vehicle just keeps turning right and goes out of the lane. The policy gets stuck in a local optimum and is not able to make any progress.

B. Roundabout with Dense Surrounding Vehicles




Fig. 6: Learning curves of different RL approaches for the roundabout scenario with adding 100 vehicles. The shaded region represents half a standard deviation of returns with $1k$ evaluation steps. Curves are smoothed for visual clarity.

After adding dense surrounding vehicles to the environment, the task becomes much more complicated. In this scenario we only implement the three RL algorithms under our proposed framework but not the DDPG baseline since it can learn nothing even without any surrounding vehicles.

As shown in Fig. 6, SAC is still the best of all methods, and it can frequently reach the final goals after about 200 epochs. Model trained with SAC learns not only the lane following skills, but also how to interact with other vehicles appropriately in the busy roundabout. DDQN, TD3 can often successfully enter the roundabout or reach the first exit, but can barely reach the second exit. This might be due to the advantage of the adaptive exploration of SAC.

We then tested the three RL algorithms for 50 times each and summarized their success rates under our roundabout task with dense surrounding vehicles. The success rates are recorded at five checkpoints: the entrance, the first exit, the second exit, the desired exit, and the goal point (Fig.4). Results are shown in Table. I. We can see that the ego vehicle trained with all of the three approaches can successfully enter the roundabout with over 80% success rate, indicating they have learned correct behaviors when entering the roundabout such as yielding to upcoming vehicles. When the distance increases, the success rate for DDQN and TD3 decreases

dramatically. This is mainly because they are easy to suffer from insufficient exploration. Thanks to a good balance of exploration and exploitation, SAC performs the best and can reach the goal point quite often.

When we look at the failure cases, we found that almost all the failures come from rear-end to the front vehicle. This is because we do not explicitly incorporate velocity of vehicles in our input, which is essential for keeping safe distance with the front vehicle. Instead, we draw historical vehicle states using boxes with fading color, which only implicitly contain the velocity information. This problem is even more serious when the input is passed through an encoding network. As in Fig.3, we can barely see the fading color of some vehicles in the reconstructed images.

TABLE I: Success rate for the roundabout scenarios evaluated on our three models DDQN, SAC and TD3. The value represent percentage of success trials.

Approach	DDQN	TD3	SAC
Entrance	80%	88%	86%
First exit	52%	74%	80%
Second exit	38%	2%	74%
Desired exit	8%	0%	64%
Goal point	0%	0%	58%

VII. CONCLUSIONS

In this paper, we proposed a framework to enable model-free deep reinforcement learning in challenging urban autonomous driving scenarios. We designed a bird-view input representation to reduce the sample complexity, and used visual encoding to capture the low-dimensional latent states. We then applied three state-of-the-art model free deep RL algorithms (DDQN, TD3, SAC) into our framework, with several tricks to improve the performance. We evaluate our method in a challenging roundabout scenario with dense surrounding vehicles in a high-definition driving simulator. The results showed that that our method has the power to solve the tasks well. Although our method is significantly better than the baseline, it doesn't solve task perfectly. In the future, we will use more computation resources to dig the limit of our method, and also improve the learning efficiency. Furthermore, We will explore the generalization ability to other scenarios.

REFERENCES

- [1] K. Arulkumaran, A. Cully, and J. Togelius. Alphastar: An evolutionary computation perspective. *arXiv preprint arXiv:1902.01724*, 2019.
- [2] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.
- [3] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [4] M. Bojarski, P. Yeres, A. Choromanska, K. Choromanski, B. Firner, L. Jackel, and U. Muller. Explaining how a deep neural network trained with end-to-end learning steers a car. *arXiv preprint arXiv:1704.07911*, 2017.
- [5] J. Chen, Z. Wang, and M. Tomizuka. Deep hierarchical reinforcement learning for autonomous driving with distinct behaviors. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1239–1244. IEEE, 2018.
- [6] J. Chen, B. Yuan, and M. Tomizuka. Deep imitation learning for autonomous driving in generic urban scenarios with enhanced safety. *arXiv preprint arXiv:1903.00640*, 2019.
- [7] F. Codevilla, M. Miiller, A. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018.
- [8] N. Djuric, V. Radosavljevic, H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, and J. Schneider. Motion prediction of traffic actors for autonomous driving using deep convolutional networks. *arXiv preprint arXiv:1808.05819*, 2018.
- [9] C. Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.
- [10] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. Carla: An open urban driving simulator. *arXiv preprint arXiv:1711.03938*, 2017.
- [11] S. Fujimoto, H. van Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- [12] D. González, J. Pérez, V. Milanés, and F. Nashashibi. A review of motion planning techniques for automated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 17(4):1135–1145, 2016.
- [13] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [14] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- [15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [17] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [18] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- [19] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [20] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [21] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on intelligent vehicles*, 1(1):33–55, 2016.
- [22] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani. Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, 2017(19):70–76, 2017.
- [23] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [24] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [25] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- [26] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller. Deterministic policy gradient algorithms. In *ICML*, 2014.
- [27] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [28] R. S. Sutton, A. G. Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [29] H. Van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [30] O. Vinyals, I. Babuschkin, J. Chung, M. Mathieu, M. Jaderberg, W. Czarnecki, A. Dudzik, A. Huang, P. Georgiev, R. Powell, et al. Alphastar: Mastering the real-time strategy game starcraft ii, 2019.
- [31] C. J. C. H. Watkins. Learning from delayed rewards. 1989.
- [32] P. Wolf, C. Hubschneider, M. Weber, A. Bauer, J. Härtl, F. Dürr, and J. M. Zöllner. Learning how to drive in a real world simulation with deep q-networks. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 244–250. IEEE, 2017.