

Curriculum Proximal Policy Optimization with Stage-Decaying Clipping for Self-Driving at Unsignalized Intersections

Zengqi Peng, Xiao Zhou, Yubin Wang, Lei Zheng, Ming Liu, *Senior Member, IEEE*, and Jun Ma

Abstract—Unsignalized intersections are typically considered as one of the most representative and challenging scenarios for self-driving vehicles. To tackle autonomous driving problems in such scenarios, this paper proposes a curriculum proximal policy optimization (CPPO) framework with stage-decaying clipping. By adjusting the clipping parameter during different stages of training through proximal policy optimization (PPO), the vehicle can first rapidly search for an approximate optimal policy or its neighborhood with a large parameter, and then converges to the optimal policy with a small one. Particularly, the stage-based curriculum learning technology is incorporated into the proposed framework to improve the generalization performance and further accelerate the training process. Moreover, the reward function is specially designed in view of different curriculum settings. A series of comparative experiments are conducted in intersection-crossing scenarios with bi-lane carriageways to verify the effectiveness of the proposed CPPO method. The results show that the proposed approach demonstrates better adaptiveness to different dynamic and complex environments, as well as faster training speed over baseline methods.

I. INTRODUCTION

In the past few decades, both academia and industry have witnessed the rapid development of autonomous driving technology [1]–[3]. However, ensuring safe and efficient passage at intersections with high vehicle density and frequent vehicle interactions remains a challenging task for autonomous driving [4], particularly in the presence of numerous human-driven vehicles exhibiting unpredictable behaviors. Inaccurate prediction of surrounding vehicle behavior can certainly influence the decision-making of the autonomous vehicle and even pose a threat to its safety. The situation becomes even more complicated when it comes to the unsignalized intersection, where the autonomous vehicle could interact with surrounding vehicles from multiple different directions simultaneously. In this sense, the increasing number of surrounding vehicles and their mutual influences lead to more complex behavior modes that are challenging to forecast,

thus severely affecting the safety and travel efficiency of the autonomous vehicle.

Currently, substantial research efforts focus on autonomous driving development, including rule-based methods, optimization-based methods, and learning-based methods. As a representative of the approach, rule-based methods show the promising effectiveness due to their transparency and comprehensibility. A set of rules are proposed to clarify the sequence of vehicles traversing the unsignalized intersection scenario. With planned rules, each vehicle makes the decision for preempting or yielding the surrounding vehicles to guarantee the road traffic safety [5]. Generally, such a rule-based strategy is designed to prioritize road traffic safety and avoid potential collisions with other social vehicles at any cost [6]. Moreover, optimization-based methods, such as model predictive control (MPC), are also widely utilized owing to their effectiveness in generating the control strategy while dealing with various constraints [7], [8]. An autonomous and safe intersection-crossing strategy is developed in [9] where the optimized trajectories of a team of autonomous vehicles approaching the intersection area are generated by centralized MPC. An effective intersection-crossing algorithm for autonomous vehicles based on vehicle-to-infrastructure communication capability is proposed in [10], where all vehicles are navigated by decentralized MPC with the sharing setting of the expected time of entering a critical zone. However, rule-based and optimization-based methods suffer from adaptiveness to the time-varying traffic situation due to the high complexity and dynamicity in real-world driving scenarios.

On the other hand, learning-based methods have recently been developed in the field of robotics and autonomous driving. Particularly, imitation learning-based methods leverage expert experience data to train the agent to generate trajectory or control commands [11], [12]. However, as supervised learning methods, the quality of the expert demonstration dataset will influence the actual performance of the agent significantly, which makes the training process rather challenging. Reinforcement learning (RL) is a promising direction to handle self-driving tasks, such as deep Q-learning, soft actor-critic, and proximal policy optimization (PPO) [3], [13], [14]. The target of RL-based methods is to train the agent to obtain the policy that maximizes the future cumulative reward by exploring the environment. A deep RL framework is proposed for navigation at occluded intersections by combining the deep Q-network and time-to-go representation [15], which demonstrates higher travel efficiency and lower collision rate than the time-to-collision

This work was supported in part by the Guangzhou-HKUST(GZ) Joint Funding Scheme under Grant 2023A03J0148; and in part by the Project of Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone under Grant HZQB-KCZYB-2020083.

Zengqi Peng, Xiao Zhou, Yubin Wang, and Lei Zheng are with the Robotics and Autonomous Systems Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China.

Ming Liu and Jun Ma are with the Robotics and Autonomous Systems Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, also with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China, and also with the HKUST Shenzhen-Hong Kong Collaborative Innovation Research Institute, Futian, Shenzhen, China.

All correspondence should be sent to Jun Ma (e-mail: jun.ma@ust.hk).

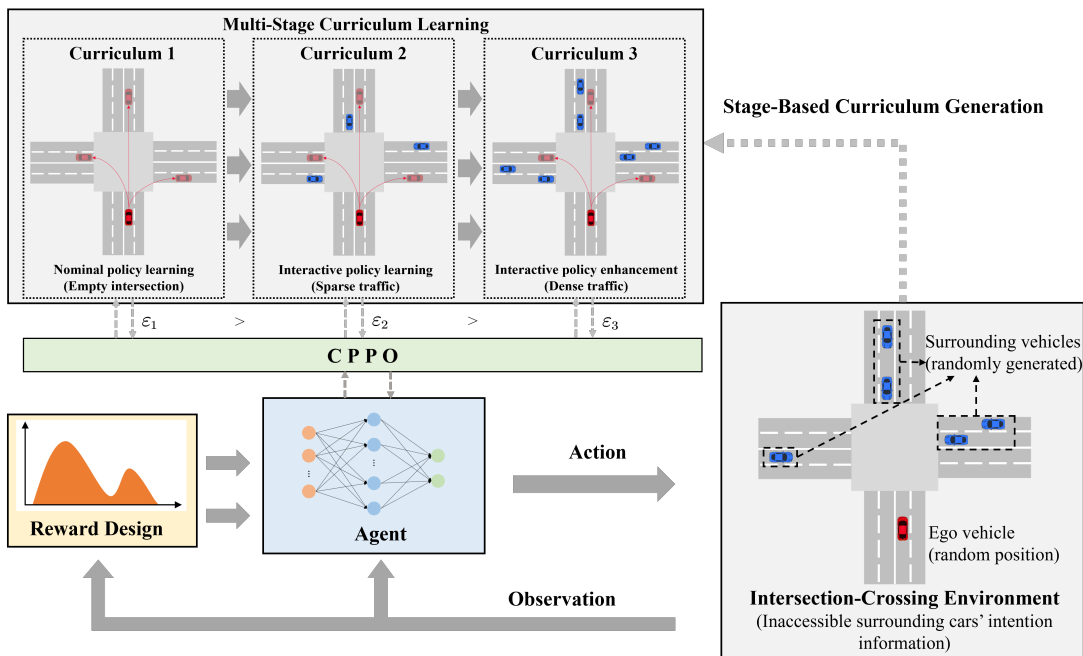


Fig. 1. Overview of the proposed CPPO framework for autonomous driving at unsignalized intersection. In the four-way intersection scenario, the ego vehicle is depicted in red and the surrounding vehicles are in blue. The solid vehicle represents the start point, and the semi-transparent vehicle represents the goal point.

method. A hierarchical decision algorithm is proposed for self-driving at intersections by integrating an RL-based decision module and an MPC-based planner [16]. However, this work only considers the situation of vehicles traveling straight, yet considering the situation where the vehicle turns left or right will significantly increase the complexity of the problem. In [17], an RL-based car following model is proposed for connected and automated vehicles at signalized intersections. The arrival time prediction is introduced into the reward function to train the agent.

Nevertheless, a major drawback of the RL-based method is that it always requires a long training process to obtain acceptable driving policies for complex self-driving tasks, such as unsignalized intersection-crossing. Because the training environment is unknown to the agent, which leads to that the agent needs to spend plenty of time interacting with the environment to understand its characteristics before learning a satisfying strategy. To solve this problem, a model-accelerated PPO algorithm is proposed in [18], where a prior model is incorporated into the RL framework to accelerate the training process. However, due to the black-box nature of neural networks, the safety of policies cannot be strictly guaranteed. Transfer learning is a class of techniques to leverage external expert knowledge before the learning process of target tasks, which helps to speed up the training procedure [19]. Moreover, curriculum learning is an alternative solution to expedite network training, which initiates the training process from easier tasks [20], [21]. Specifically, in curriculum learning, a series of course tasks with increasing difficulties are designed to enable agents to learn optimal strategies faster and more efficiently. In

[22], the curriculum learning technology is introduced into the soft actor-critic algorithm for autonomous overtaking tasks, which leads to faster convergence speed compared to the vanilla RL method. Besides, an automated curriculum mechanism is proposed in [23] to train agents for traversing the unsignalized intersection. The agent can obtain a fine-tuned policy in the final phase by dropping the future state information during the training process. However, in this work, surrounding vehicles are assumed not to interact with other vehicles, and future trajectories of opposing vehicles are accessible. Besides, the total number of surrounding vehicles is fixed. These assumptions and settings could possibly limit the generalization of the trained policy in more dynamic intersection scenarios.

This work addresses the unsignalized intersection-crossing task where the ego vehicle interacts with varying numbers of surrounding vehicles. The main contributions of this work are summarized as follows: a curriculum proximal policy optimization (CPPO) framework with stage-decaying clipping is proposed for training the agent in the highly dynamic intersection-crossing self-driving task, where the reward function is particularly designed to balance the safety and travel efficiency in different traffic situations. The stage-based curriculum learning technology is introduced into the PPO method with a decaying clipping parameter to accelerate the training process and improve the generalization of the trained policy. By learning a series of courses with increasing difficulty levels, the agent can capture the uncertainties of surrounding vehicles implicitly and adapt to situations effectively with varying numbers of surrounding vehicles. A series of simulations in different intersection scenarios

are conducted to evaluate the performance of the proposed method and baseline method in *Highway_Env* [24]. The CPPO method demonstrates faster training speed and better generalization than the standard PPO method.

The rest of this paper is organized as follows. Section II gives an introduction to the formulation of the intersection problem addressed in this work. Section III illustrates the proposed methodology. Section IV presents the experimental results. The conclusion and future work are discussed in Section V.

II. PROBLEM DEFINITION

In this section, we first introduce the task scenarios to be solved in this work. Then, the formulation of the learning environment is illustrated.

A. Problem Statement

The problem to be solved in this work is to control the ego vehicle to traverse an unsignalized four-way intersection and reach the goal position. Furthermore, each road consists of two lanes. The task scenario is shown in Fig. 1.

We assume that the ego vehicle always starts from a random position (denoted by a solid red vehicle) in the lower zone of the intersection, and the goal position of the ego vehicle (denoted by a semi-transparent red vehicle) is also randomly generated within the left, upper, and right zones. There are several surrounding vehicles driving from other lanes towards different target lanes. They will react to the behavior of the ego vehicle. Without loss of generality, we assume that the position and velocity information of surrounding vehicles can be accessed by the ego vehicle. But the information about the driving intention of surrounding vehicles is unknown to the ego vehicle, which increases the difficulties of this task. The objective of this task is to generate a sequence of actions that enables the ego vehicle to expeditiously approach the target point while ensuring collision avoidance with surrounding vehicles and staying within the road boundaries.

B. Learning Environment

In this work, we frame the agent’s learning objective as the optimal control of a Markov Decision Process by defining state space, action space, state transition dynamics, reward function, and discount factor. Then the RL problem can be represented by a tuple $\mathcal{E} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$.

State space \mathcal{S} : In this scenario, the state space \mathcal{S} is explicitly defined by *Highway_Env* environment, which consists of state matrices \mathbf{S}_t . At each timestep t , the agent observes the kinematic features of vehicles in the intersection, so the state matrix can be defined as follows,

$$\mathbf{S}_t = [\mathbf{s}_t^0 \ \mathbf{s}_t^1 \ \dots \ \mathbf{s}_t^N]^T \quad (1)$$

where N denotes the total number of vehicles in the intersection. The first row of \mathbf{S}_t , i.e., $(\mathbf{s}_t^0)^T$, is a vector consisting of the kinematic features of the ego vehicle, while other rows of \mathbf{S}_t , i.e., $(\mathbf{s}_t^i)^T$ ($i = 1, 2, \dots, N$), represent vectors of the

kinematic features of the surrounding vehicles. The vector of the kinematic features is defined as follows,

$$\mathbf{s}_t^i = [x_t^i \ y_t^i \ v_{x,t}^i \ v_{y,t}^i \ \sin \psi_t^i \ \cos \psi_t^i]^T \quad (2)$$

where x_t^i and y_t^i are the vehicle’s current position in the world coordinate system, respectively; $v_{x,t}^i$ and $v_{y,t}^i$ are the speed of the vehicle along the X-axis and Y-axis, respectively; ψ_t^i is the heading angle of the vehicle at timestep t .

Action space \mathcal{A} : Similar to the human driver’s operation in driving, the action space of the agent consists of five basic discrete actions as follows,

$$\mathcal{A} = \{ A^0, A^1, A^2, A^3, A^4 \} \quad (3)$$

where A^0 and A^2 are the left and right lane-changing action, respectively; A^1 is the motion keeping action; A^3 and A^4 represent the deceleration and acceleration action, respectively.

The vehicle is controlled by two low-level controllers, the longitudinal and lateral controllers, which convert discrete actions (3) to the continuous control input $[a, \delta]^T$ of the vehicle, where a and δ are the acceleration and steering angle, respectively. Considering the physical limitations, the maximum acceleration and steering angle are set as 8 m/s² and 45 degrees, respectively.

State transition dynamics $\mathcal{P}(\mathbf{S}_{t+1}|\mathbf{S}_t, a_t)$: The transition function \mathcal{P} defines the transition of environment state, which follows the Markov transition distribution. The next state generated by \mathcal{P} is related to the current state \mathbf{S}_t and the applied action $a_t \in \mathcal{A}$. The transition dynamics $\mathcal{P}(\mathbf{S}_{t+1}|\mathbf{S}_t, a_t)$ is implicitly defined by *Highway_Env* environment and unknown for the agent.

Reward function \mathcal{R} : The reward function assigns a positive reward for successfully completing an episode and for maintaining survival. It penalizes collisions, out-of-the-road, and lane-changing behaviors. When there are few vehicles on the road, the acceleration behaviors will be rewarded, and vice versa. The details about the structure of the reward function will be introduced in Section III-C.

Discount factor $\gamma \in (0, 1)$: The future reward is accumulated with a discount factor γ .

III. METHODOLOGY

In this section, we first illustrate the details of the proposed CPPO framework with stage-decaying clipping. Then, the curriculum setting is introduced to enhance the training process of the RL agent. Lastly, the multi-objective reward structure for the agent is presented.

A. Curriculum Proximal Policy Optimization with Stage-Decaying Clipping

PPO is a model-free RL framework to solve the sequential decision-making problem under uncertainties. It alternatively constructs a clipped surrogate objective to replace the original function, which generates a lower bound on the unclipped objective and avoids the incentive for an excessively large policy update. Therefore, PPO algorithm facilitates the learning of policies in a faster and more efficient way. Specifically,

the objective function in the PPO algorithm is shown as follows,

$$J_{clip}(\theta) = \mathbb{E}_t \left[\min \left(\rho_t(\theta) \hat{A}_t, \text{clip}(\rho_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t \right) \right] \quad (4)$$

where $\rho_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)}$ is the probability ratio between the new policy and old policy, \hat{A}_t is the estimated advantage function at timestep t , ε is a hyperparameter.

The intuitive idea is to change the value of the hyperparameter ε in the different training periods. Empirically, we set $\varepsilon = \{0.25, 0.2, 0.15\}$. During the beginning stage of training, a large parameter $\varepsilon_1 = 0.25$ is used for the rough exploration, which is then adjusted to the second largest parameter $\varepsilon_2 = 0.2$ during the middle stage, and finally adjusted to a smaller parameter $\varepsilon_3 = 0.15$ during the later stage. However, determining when to make adjustments to the parameters is a problem.

Remark 1: If $\hat{A} > 0$, the probability ratio will not exceed $1 + \varepsilon$. Otherwise, the ratio will be less than or equal to $1 - \varepsilon$. Therefore, the magnitude of the clipping parameter ε will influence the training speed and the performance of the trained policy. A larger clipping parameter allows for a more significant step size of the update, which means that the agent can have a faster training process, but the optimality of the trained policy cannot be ensured. On the contrary, a smaller clipping parameter may lead to slow update speed and drop into a local optimal policy. Therefore, stage-decaying clipping can capitalize on the strengths of both approaches while circumventing their weaknesses.

In intersection-crossing tasks, the ego vehicle needs to interact with a variable number of interactive surrounding vehicles with different driving behaviors from the other three directions. Therefore, these scenarios are rather complex. It is difficult to get a satisfactory driving policy by directly deploying the PPO algorithm in these high-dynamic scenarios. Here, we introduce stage-based curriculum learning technology to generate a task sequence with increasing complexity for training acceleration and better generalization. Besides, the clipping parameter can be adjusted when switching the curriculum, which addresses the issue of when to change the clipping parameter mentioned before.

We generate a curriculum sequence with three stages, which is represented as $\Omega = \{\Omega_1, \Omega_2, \Omega_3\}$. The curriculum sequence is designed for different goals with increasing complexity.

Curriculum 1: Intersection without surrounding vehicles. In stage 1, which is denoted as Ω_1 , there is only the ego vehicle in the intersection. The objective of this curriculum is to learn a transferable nominal policy to obtain the nominal policy π_1 that can find an action sequence to the goal point. The ego agent is guided with empirically designed rewards in this curriculum stage, which can decrease the exploration time in the whole action space and avoid local optimum with poor generalization. In this curriculum, the hyperparameter of the clipped function is set as ε_1 .

Curriculum 2: Intersection with a few vehicles. In the

second stage, we load the policy trained in Curriculum 1 as the initial policy of this stage for the following training process. In this curriculum, there are a few surrounding vehicles in the intersection. We aim to train the nominal policy to obtain a new policy π_2 with the preliminary obstacle avoidance ability by maximizing the intersection-crossing reward in this stage. The hyperparameter is switched to ε_2 .

Curriculum 3: Intersection with numerous vehicles. In the third stage, we load the policy trained in Curriculum 2 as the initial policy. In this stage, there are numerous surrounding vehicles in the intersection. The objective of this curriculum is to train the previous policy to obtain the optimal policy π^* with better obstacle avoidance ability in a more complex environment. In the preceding episodes of this stage, the parameter ε remains at 0.2, and then transits to ε_3 to continue the course training.

B. Multi-Objective Reward Design

For the RL-based method, the design of the reward is essential for the success of policy training. An inappropriate reward function not only slows down the speed of training but also leads to a trained policy with poor performance. Therefore, it is challenging to guide the agent to obtain a satisfactory driving strategy in a complex scenario. In this work, the reward is smartly designed for intersection scenarios with different densities.

Considering the complexity of the target scenario, a comprehensive reward function is designed as follows,

$$r = r_{succ}(T, N_{car}) + r_{colli}(v, N_{car}) + r_{TO} + r_{OfR} + r_{LC} + r_l, \quad (5)$$

where r_{succ} and r_l are the reward for successfully completing tasks and surviving in the task, respectively; r_{colli} , r_{TO} , r_{OfR} , and r_{LC} are the penalty of collision with surrounding vehicles, time-out, out-of-the-road boundary, and lane-changing behavior, respectively.

Remark 2: Inspired by the idea of curriculum learning, several terms in the reward function are related to the setting of the scenario. The success reward term is related to both the time of finishing the task and the maximum number of surrounding vehicles in the intersection central area when the ego vehicle is crossing. It will encourage autonomous vehicles to expedite the intersection-crossing task, and greater rewards will be obtained when completing more complex tasks. Additionally, the collision penalty is related to the speed of the ego vehicle when the collision happens and the maximum number of surrounding vehicles. This reward term aims to encourage autonomous vehicles to maintain a lower speed when there are more vehicles in the central area of the intersection to ensure safety, and greater penalties will be imposed if collisions occur in more complex tasks. Thereby, the balance between the safety and travel efficiency is achieved through the particular designed reward function.

To sum up, the proposed CPPO framework with stage-decaying clipping is summarized in Algorithm 1.

Algorithm 1: Curriculum Proximal Policy Optimization with Stage-Decaying Clipping

Input: Environment state s_t , curriculum set Ω **Output:** $\pi^* = f(\theta^*)$

- 1 Initialize the policy network with parameter θ_0 ;
 - 2 **while** not terminated **do**
 - 3 Select curriculum Ω_i from curriculum set Ω ;
 - 4 Reset the environment according to the setting of curriculum Ω_i ;
 - 5 Select the clipping parameter $\varepsilon = \varepsilon_j$;
 - 6 **if** curriculum switched **then**
 - 7 Load policy π^* trained by Ω_{i-1} as the initial policy;
 - 8 **end**
 - 9 Update the policy network by maximizing the designed reward (5);
 - 10 Save the trained policy as π_i ;
 - 11 **end**
 - 12 The policy network obtained by the last curriculum is the final policy $\theta^* = \theta_3$.
-

IV. EXPERIMENTS

In this section, we implement the proposed algorithm in dynamic intersection-crossing scenarios. Then we compare the performance of CPPO with that of two baseline methods. The simulations are conducted in *Highway_Env* [24].

A. Experimental Settings

The experiments are conducted on the Windows 11 system environment with a 3.90 GHz AMD Ryzen 5 5600G CPU. The task scenarios are constructed based on *Highway_Env* environment, where each road is a bi-lane carriageway. We use fully-connected networks with 1 hidden layer of 128 units (action network) and 64 units (critic network) to represent policies. The neural network is constructed in PyTorch [25] and trained with an Adam optimizer [26]. The MDP is solved using the proposed CPPO framework and the standard PPO method with two different fixed clipping parameters. The simulation frequency is set as 15 Hz. The hyperparameters for network training are listed in Table I. Here, we compare the proposed method, CPPO, and two baseline methods with different clipping parameters ($\varepsilon = 0.15, 0.25$). For the sake of fairness, other network parameters of these methods are the same.

TABLE I
HYPERPARAMETER SETTINGS.

Hyperparameter	Value
Learning rate for actor network	5×10^{-4}
Learning rate for critic network	1×10^{-3}
Discount factor	0.9
Number of epochs	20

Then we test all trained policies in intersection scenarios with different numbers N_{sv} of surrounding vehicles, whose

behaviors are characterized by the intelligent driver model (IDM) [27]. We record the success rate, collision rate, time-out rate, and out-of-road rate, respectively. By conducting these simulations, we can evaluate the generalization performance of these trained policies.

1) *No surrounding vehicles* ($N_{sv} = 0$): This scenario can be used to check whether the trained policy has the ability to find the nominal trajectory to achieve the goal point.

2) *Different number of surrounding vehicles* ($N_{sv} = 1, 2, \dots, 6$): These trained policies are tested in simple and complex scenarios to estimate their generalization performance and safety.

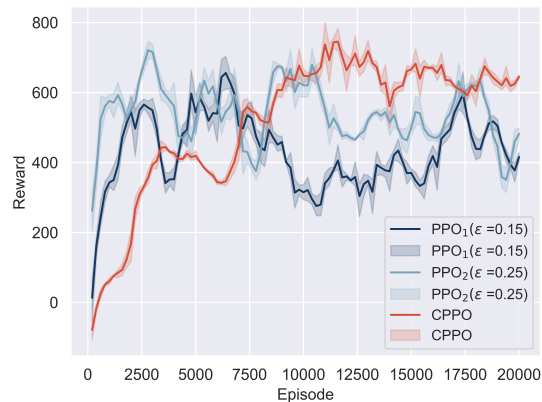


Fig. 2. Reward curve comparison among different methods. The training curves are smoothed by the Savitzky-Golay filter. The curriculum is switched at 2×10^3 th and 5×10^3 th episode, respectively.

TABLE II
TRAINING TIME OF DIFFERENT METHODS.

Methods	Training Time (hour:min.:sec.)
CPPO	3:34:28
PPO ₁ ($\varepsilon = 0.15$)	6:45:32
PPO ₂ ($\varepsilon = 0.25$)	5:58:28

B. Training Results

To illustrate the effectiveness of the proposed framework, the training time is listed in Table II for comparison. It is obvious that CPPO has a much faster training speed than the two baseline methods. Specifically, the training speed of CPPO is 47.2% faster than PPO₁, and 40.2% faster than PPO₂.

The change in reward during the training process is shown in Fig. 2. According to these three learning curves, we can find that the CPPO agent initially receives the least reward, which is because that the reward function is positively correlated with the number of vehicles in the environment upon successful task completion. However, as CPPO's policy network converges to the optimal policy in the final stage, its reward curve surpasses that of the other two baseline methods. As the baseline algorithm PPO₁, with the parameter $\varepsilon = 0.15$, is deployed directly in a complex environment




Fig. 3. Demonstration of the driving performance attained by the proposed CPPO method in an unprotected left-turn task at the unsignalized intersection. The green car and blue cars represent the ego vehicle and surrounding vehicles under normal driving conditions, respectively. The red cars represent the vehicles that have collided. (a)-(d) present four key snapshots during the performance evaluation demonstration.

TABLE III
COMPARISON OF SUCCESS RATE, COLLISION RATE, AND TIME-OUT RATE AMONG DIFFERENT METHODS.

Methods	$N_{sv}=0$			$N_{sv}=2$			$N_{sv}=4$		
	succ.(%)	coll.(%)	time-out(%)	succ.(%)	coll.(%)	time-out(%)	succ.(%)	coll.(%)	time-out(%)
CPPO	100	0	0	90.5	9.5	0	78.5	21.5	0
PPO ₁ ($\epsilon = 0.15$)	55.5	0	44.5	86	14	0	76.5	23.5	0
PPO ₂ ($\epsilon = 0.25$)	69.5	0	30.5	83.5	16.5	0	72	28	0

for learning, its reward curve exhibits minor fluctuations in the later stages of training and is lower than that of the CPPO algorithm. For the agent trained by the baseline PPO₂ with $\epsilon = 0.25$, its reward curve exhibits a rapid increase initially. However, due to the large parameter ϵ , it oscillates significantly in later episodes. Although its reward curve ends up resembling that of PPO₁, the performance of its policy may be inferior to that of PPO₁. Above results illustrate that the introduction of stage learning allows for a more efficient sampling process, leading to a faster training speed and better convergence compared to baseline methods.

C. Performance Evaluation

To further demonstrate the superiority of CPPO in unsignalized intersection scenarios, comparative simulations are conducted. We test policies obtained by three methods in intersection scenarios with different numbers of surrounding vehicles. Each method is retested 200 times in each scenario. Among all testing results attained by the proposed CPPO method at the unsignalized intersection, we pick up the results from a left-turn task for demonstration, and the details are presented in Fig. 3. In this demonstration, the ego vehicle is generated on the right lane of the lower zone, and its goal lane is the left lane in the left zone. We can observe that the ego vehicle exhibits a safe interaction behavior of decelerating and steering left when encountering a surrounding vehicle approaching from its left side in Fig. 3. Specifically, in the first snapshot, the ego vehicle drives at a constant speed from the right lane in the lower zone towards the central area of the intersection, preparing for a left turn. Then, in the second snapshot, the ego vehicle perceives that it is getting close to a surrounding vehicle ahead, and the vehicle ahead shows no signs of slowing down to yield. As a result, the ego vehicle decelerates and steers left to avoid a collision. Afterwards, in the third snapshot,

the closest surrounding vehicle in the previous snapshot has moved away from the ego vehicle. The ego vehicle perceives that there are no other vehicles that could potentially collide with itself. Therefore, it adjusts its heading angle, accelerates towards the target lane, and continues driving until the completion of the unsignalized intersection task.

The success rates of the three methods in all evaluation scenarios are shown in Fig. 4. It is evident that the CPPO method achieves the highest overall task success rate. While its success rate decreases with the increasing complexity of the environment, it remains higher than that of the other two baseline methods. It indicates that the proposed method has better generalization performance. Furthermore, except for the scenario where there are no surrounding vehicles, PPO₁ achieves a slightly higher task success rate than PPO₂. This is because a smaller parameter ϵ enables a search for better policies. Besides, it is noted that both two baseline methods have a low task success rate in the scenario where there is only the ego vehicle.

In addition, we have summarized the results of the success rate, collision rate, and timeout rate of tests with $N_{sv} = 0, 2, 4$ surrounding vehicles in Table III. From this statistical result, we can observe that both baseline methods exhibit a large number of timeouts in testing, 44.5% and 30.5% for PPO₁ and PPO₂, where there is no surrounding vehicles. This suggests that directly deploying the agent in a complex environment for training may cause the policy to become stuck in a local optimum. For other task scenarios with surrounding vehicles, two baseline methods did not exhibit timeouts in testing results. Because of the integration of the curriculum sequence, there is no timeout case happening for the CPPO. Therefore, the introduction of the curriculum learning technologies enables the agent to converge to a better optimum compared to those agents trained directly.




Fig. 4. Success rate comparison of three algorithms in unsignalized intersection with a different number of surrounding vehicles.

V. CONCLUSION

In this paper, we proposed a novel CPPO framework with stage-decaying clipping for unsignalized intersection-crossing tasks. We formulate a curriculum sequence for guiding the agent to learn the driving policy in scenarios where task difficulty gradually heightens by increasing the number of surrounding vehicles, and the clipping parameter in PPO varies as the curriculum stage switches. Besides, the reward function is particularly designed to guide the agent to balance safety and travel efficiency in different situations. A series of experiments were conducted in *Highway_Env* environment to verify the effectiveness of the proposed method. We compared the performance of the proposed method and two baseline methods. The results show that the CPPO method has the fastest training speed and the highest task success rate among different settings, which demonstrates that the proposed method has better generalization performance than all baseline algorithms. In the future, we will consider incorporating game theoretic methods into the CPPO framework to enhance the effectiveness of our method.

REFERENCES

- [1] 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*, vol. 1, no. 1, pp. 33–55, 2016.
- [2] S. Narayanan, E. Chaniotakis, and C. Antoniou, "Shared autonomous vehicle services: A comprehensive review," *Transportation Research Part C: Emerging Technologies*, vol. 111, pp. 255–293, 2020.
- [3] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2021.
- [4] L. Wei, Z. Li, J. Gong, C. Gong, and J. Li, "Autonomous driving strategies at intersections: Scenarios, state-of-the-art, and future outlooks," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 44–51, IEEE, 2021.
- [5] G. Lu, L. Li, Y. Wang, R. Zhang, Z. Bao, and H. Chen, "A rule based control algorithm of connected vehicles in uncontrolled intersection," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 115–120, IEEE, 2014.
- [6] A. Aksjonov and V. Kyrki, "Rule-based decision-making system for autonomous vehicles at intersections with mixed traffic environment," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 660–666, IEEE, 2021.

- [7] X. Qian, I. Navarro, A. de La Fortelle, and F. Moutarde, "Motion planning for urban autonomous driving using bézier curves and MPC," in *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pp. 826–833, IEEE, 2016.
- [8] Y. Wang, Y. Li, H. Ghazzai, Y. Massoud, and J. Ma, "Chance-aware lane change with high-level model predictive control through curriculum reinforcement learning," *arXiv preprint arXiv:2303.03723*, 2023.
- [9] L. Riegger, M. Carlander, N. Lidander, N. Murgovski, and J. Sjöberg, "Centralized MPC for autonomous intersection crossing," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 1372–1377, IEEE, 2016.
- [10] M. Kneissl, A. Molin, H. Esen, and S. Hirche, "A feasible MPC-based negotiation algorithm for automated intersection crossing," in *European Control Conference (ECC)*, pp. 1282–1288, IEEE, 2018.
- [11] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4693–4700, IEEE, 2018.
- [12] K. Menda, K. Driggs-Campbell, and M. J. Kochenderfer, "Ensemble-dagger: A Bayesian approach to safe imitation learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5041–5048, IEEE, 2019.
- [13] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," *arXiv preprint arXiv:1704.02532*, 2017.
- [14] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [15] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, "Navigating occluded intersections with autonomous vehicles using deep reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2034–2039, IEEE, 2018.
- [16] T. Tram, I. Batkovic, M. Ali, and J. Sjöberg, "Learning when to drive in intersections by combining reinforcement learning and model predictive control," in *IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 3263–3268, IEEE, 2019.
- [17] M. Zhou, Y. Yu, and X. Qu, "Development of an efficient driving strategy for connected and automated vehicles at signalized intersections: A reinforcement learning approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 1, pp. 433–443, 2019.
- [18] Y. Guan, Y. Ren, S. E. Li, Q. Sun, L. Luo, and K. Li, "Centralized cooperation for connected and automated vehicles at intersections by proximal policy optimization," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 11, pp. 12597–12608, 2020.
- [19] Z. Zhu, K. Lin, A. K. Jain, and J. Zhou, "Transfer learning in deep reinforcement learning: A survey," *arXiv preprint arXiv:2009.07888*, 2020.
- [20] X. Wang, Y. Chen, and W. Zhu, "A survey on curriculum learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 4555–4576, 2021.
- [21] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, "Curriculum learning for reinforcement learning domains: A framework and survey," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 7382–7431, 2020.
- [22] Y. Song, H. Lin, E. Kaufmann, P. Dürr, and D. Scaramuzza, "Autonomous overtaking in gran turismo sport using curriculum reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9403–9409, IEEE, 2021.
- [23] S. Khaitan and J. M. Dolan, "State dropout-based curriculum reinforcement learning for self-driving at unsignalized intersections," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 12219–12224, IEEE, 2022.
- [24] E. Leurent, "An environment for autonomous driving decision-making," <https://github.com/eleurent/highway-env>, 2018.
- [25] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [27] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, p. 1805, 2000.