Computer Life
ISSN: 1819-4818 | Vol. 12, No. 1, 2024

Research Status of End-or-End Autonomous Driving Technology

Shixin Zhao [1,2], Feng Pan [1,2,*]

[1]Beijing Key Laboratory of Information Service Engineering, Beijing Union University, Beijing 100101, China

[2]College of Robotics, Beijing Union University, Beijing 100101, China

Abstract: With the continuous development of autonomous driving technology, it holds significant potential in reducing the risk of traffic accidents, alleviating traffic congestion, and im-proving traffic efficiency. Traditional autonomous driving systems employ a modular deployment strategy, dividing the development into separate modules for perception, decision-making, planning, and control, which are then integrated into the vehicle. Currently, end-to-end autonomous driving methods have emerged as a research trend in the field of autonomous driving. This approach directly maps input data from the perception stage to driving behavior, simplifying the overall architecture of the autonomous driving system to reduce complexity. This paper provides a summary of research progress in end-to-end methods in the field of autonomous driving control and concludes by discussing future research directions and challenges in end-to-end autonomous driving technology.

Keywords: Autonomous vehicles; end-or-end; reinforcement learning; vehicle control.

1. Introduction

With the continuous promotion of economic globalization and technological advancement, the number of motor vehicles worldwide is showing a rapid growth trend. With the continuous increase in the number of motor vehicles, the risk of road traffic accidents also increases accordingly, requiring more measures to gradually improve traffic safety conditions and reduce the occurrence of traffic accidents.

The continuous development of autonomous driving technology can effectively reduce vehicle traffic accidents and, at the same time, has tremendous potential in alleviating traffic congestion, improving traffic efficiency, and reducing energy consumption[1]. However, autonomous vehicles face extremely complex and highly uncertain situations on the road. This uncertainty arises from various factors, including the diversity of traffic environments, constantly changing road conditions, and the possibility of drastic changes in weather conditions in a short period of time. These factors may affect the performance of sensors, thereby increasing the uncertainty of autonomous driving systems.

Currently, end-to-end autonomous driving methods have become an important trend in the research of autonomous driving. End-to-end autonomous driving methods directly map from raw sensor data to trajectory points or control signals, simplifying the architecture of the entire autonomous driving system and integrating multiple modules into a single neural network model, reducing system complexity. Compared with traditional methods, there are no external gaps between perception modules and control modules, and few human-written algorithms are embedded. Therefore, end-to-end autonomous driving methods handle the interaction between vehicles and the environment more efficiently and are expected to achieve expert-level performance in the field of autonomous driving[2]. This paper will summarize the research progress of end-to-end methods in the field of autonomous driving control and discuss the research directions and challenges of future end-to-end autonomous driving technologies.

2. End-to-End Technology

End-to-end autonomous driving technology is a method that trains and executes the entire autonomous driving system as a single model. The goal of this approach is to enable vehicles to directly learn driving behavior from raw sensor data (such as cameras and LiDAR) without the need to decompose it into traditional modular subsystems (perception, planning, and control). As shown in Figure 1, compared to traditional autonomous driving frameworks, end-to-end autonomous driving takes raw sensor data (such as cameras, LiDAR, and GPS) as input and generates planning or low-level control actions (steering angle, throttle, and brake) as output.

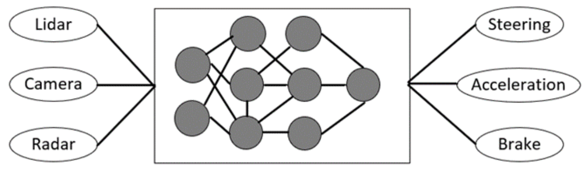

Figure 1. End-to-End Autonomous Driving Framework

The advantage of the end-to-end approach lies in its simple structure, eliminating the need for complex feature engineering and selection processes. Additionally, this method relies on learning from large-scale data to optimize the objective function, thus demonstrating stronger robustness in addressing unpredictable changes in the real world. End-to-end autonomous driving technology reduces the complexity of the entire autonomous driving system and eliminates the integration requirements between multiple independent modules and sensors. End-to-end models can adapt to various scenarios, including complex traffic situations and environmental changes, by leveraging the data accumulated during the training process.

3. Based on End-to-End Autonomous Driving Technology

End-to-end autonomous driving technology refers to a single neural network model that directly maps from sensor data input (such as cameras, radar, LiDAR, etc.) to vehicle control outputs (such as steering, acceleration, and braking). End-to-end autonomous driving can be categorized into two different types: imitation learning and reinforcement learning.

Imitation learning refers to the learning strategy of an agent based on expert trajectories, typically provided with expert decisions and control information [3]. The specific objective of the agent is to evaluate the best match between states and actions, aiming to replicate the expert's trajectory as closely as possible. In 2016, Bojarski et al. proposed a groundbreaking approach where a convolutional neural network was trained to predict steering commands solely from a forward-facing monocular camera, focusing primarily on lateral control without considering longitudinal control signals, which exhibited good performance in a limited number of less complex scenarios [4]. In 2018, Codevilla et al. introduced a conditional end-to-end imitation learning model for autonomous driving control, incorporating both lateral and longitudinal control [5]. This approach builds upon classical imitation learning by incorporating expert instruction vectors and jointly determining the current action based on the state space.

Conditional imitation learning is a milestone in the field of imitation learning for autonomous driving, demonstrating that convolutional neural networks (CNNs) can learn to autonomously perform lane and road tracking tasks. However, end-to-end imitation learning models lack interpretability, and many researchers have attempted to address this challenge by introducing intermediate representation layers. Chen et al. proposed a novel paradigm called direct perception, specifically designed for prediction in urban autonomous driving scenarios [6]. This approach explicitly captures features of the surrounding environment, which are then fed to lower-level controllers to generate steering and acceleration commands for the vehicle. Sauer et al. further developed an advanced direct perception model [7] that leverages both video and high-level commands for intermediate representation, which is then used to compute control signals as output. Urtasun and her team also introduced two interpretable end-to-end planners [8][9], both utilizing raw LiDAR data and high-definition maps to predict safe trajectories and intermediate representations that play a crucial role in strategic decision-making in response to surrounding environmental scenes.

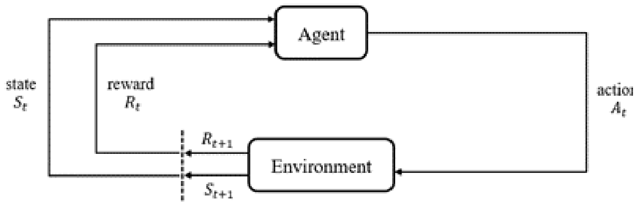

Figure 2. End-to-End Autonomous Driving Framework

End-to-end imitation learning methods require a large amount of manually labeled data, and different drivers may make completely different decisions in the same situation, leading to significant uncertainty during training. Value-based and policy-based reinforcement learning methods are effective approaches for tackling more complex problems and adapting better to real-world driving scenarios. Training autonomous vehicles using reinforcement learning methods has become a growing trend in end-to-end autonomous driving research [2]. Mnih et al. employed a deep learning approach based on Q-Learning to directly learn control signals from screen captures [10]. Wolf et al. introduced the Q-Learning method to the field of intelligent vehicles [11], where the vehicle selects corresponding actions based on image information in the Gazebo simulator [12]. To alleviate the issue of poor stability with high-dimensional perceptual inputs, the Conditional DQN approach was proposed, which utilizes deblurring algorithms to enhance the prediction stability of different motion commands. The proposed model achieved performance comparable to human driving in specific scenarios. Saxena et al. utilized the Proximal Policy Optimization (PPO) algorithm [14] for learning control policies in the continuous motion planning space [13]. Their model simulated interactions with other vehicles to avoid collisions and improve passenger comfort. Building upon this work, Ye et al. utilized PPO to learn automatic lane-changing strategies in real highway scenarios [15]. With the self-vehicle and surrounding vehicle states as inputs, the agent learned to avoid collisions and drive smoothly. This demonstrated the effectiveness of PPO-based reinforcement learning algorithms in end-to-end autonomous driving. The safety of end-to-end autonomous driving has also raised significant concerns. Constrained Policy Optimization (CPO) is a pioneering general policy utilization algorithm for constrained reinforcement learning, ensuring proximity to constraint satisfaction at each iteration. Li et al. incorporated risk perception algorithms into a deep reinforcement learning framework to learn risk-aware driving decision policies for lane-changing tasks with minimum expected risk. Mo et al. utilized Monte Carlo tree search to mitigate unsafe behaviors in overtaking tasks in highway scenarios.Hu and his team[16] conducted an analysis and research on the key components of perception and prediction, organizing them in a certain priority to ensure that all subtasks contribute to planning. They proposed a Unified Autonomous Driving framework (UniAD), which is the first framework to integrate full-stack driving tasks into a single deep neural network. It leverages the advantages of each subtask and module to execute safe vehicle planning. On the other hand, Zheng[17] and his team combined generative artificial intelligence with end-to-end autonomous driving technology, introducing Generative End-to-End Autonomous Driving (GenAD). GenAD proposes an instance-centric scene representation, initially transforming the surrounding scenes into maps and perceptual instances. Subsequently, it utilizes a variational autoencoder to learn future trajectory distributions in a structured latent space for trajectory prior modeling, achieving state-of-the-art performance in end-to-end autonomous driving.

Executing end-to-end reinforcement learning methods on real-world autonomous vehicles is a challenging task. In 2019, Kendall et al. implemented the Deep Deterministic Policy Gradient (DDPG) algorithm [18] on an intelligent vehicle, with a monocular image as the sole input. The agent learned a lane-following policy and achieved human-level performance in a 250-meter road test [17]. This work marked the first application of deep reinforcement learning on full-sized autonomous vehicles. To further enhance driving safety and comfort, Wang et al. introduced a strategy for lane-changing based on human expert knowledge. This method can be executed on single or multiple vehicles, facilitating smooth lane transitions.

4. Conclusion

Autonomous driving has achieved significant milestones, and its successful validation on urban semi-open roads is a strong testament to its feasibility. The end-to-end approach to autonomous driving aims to achieve comprehensive driving decisions from sensor inputs to vehicle control through a single neural network model, eliminating the need for complex segmented processing pipelines. While this approach has potential advantages such as reducing system complexity and higher levels of automation, it also faces a range of challenges.

1）Training an end-to-end autonomous driving model requires a large amount of data. This data needs to encompass various traffic scenarios, weather conditions, and road situations to ensure that the model possesses a wide-ranging generalization capability. The collection and annotation of such data incur significant costs.

2）Ensuring high levels of safety and reliability is crucial for the practical application of autonomous driving systems, as they are directly involved in driving operations. Relying solely on a single end-to-end model can lead to failures due to various reasons, such as the model's inability to comprehend extreme situations or perception failures. Ensuring the safety and reliability of the system is a highly challenging task.

Indeed, despite the challenges, end-to-end autonomous driving represents a new direction in autonomous driving technology and may receive further research and development in the future. It offers the potential for a simplified and more streamlined approach to autonomous vehicle control. As researchers continue to explore and refine this methodology, it may lead to advancements in the overall efficiency, safety, and scalability of autonomous driving systems.

References

[1] LIU Y F,HU X M,CHEN G W.Review of end-to-end motion planning for autonomous driving with visual perception[J].Journal of Image and Graphics, 2021,26(01):49-66.

[2] Teng S, Hu X, Deng P, et al. Motion planning for autonomous driving: The state of the art and future perspectives[J]. IEEE Transactions on Intelligent Vehicles, 2023.

[3] Zhang Z, Tian R, Sherony R, et al. Attention-based interrelation modeling for explainable automated driving[J]. IEEE Transactions on Intelligent Vehicles, 2022, 8(2): 1564-1573.

[4] Bojarski M, Del Testa D, Dworakowski D, et al. End to end learning for self-driving cars[J]. arXiv preprint arXiv: 1604.07316, 2016.

[5] Codevilla F, Müller M, López A, et al. End-to-end driving via conditional imitation learning[C]//2018 IEEE international conference on robotics and automation (ICRA). IEEE, 2018: 4693-4700.

[6] Chen C, Seff A, Kornhauser A, et al. Deepdriving: Learning affordance for direct perception in autonomous driving[C]//Proceedings of the IEEE international conference on computer vision. 2015: 2722-2730.

[7] Sauer A, Savinov N, Geiger A. Conditional affordance learning for driving in urban environments[C]//Conference on Robot Learning. PMLR, 2018: 237-252.

[8] Zeng W, Luo W, Suo S, et al. End-to-end interpretable neural motion planner[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019: 8660-8669.

[9] Sadat A, Casas S, Ren M, et al. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations[C]//Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16. Springer International Publishing, 2020: 414-430.

[10] Mnih V, Kavukcuoglu K, Silver D, et al. Playing atari with deep reinforcement learning[J]. arXiv preprint arXiv: 1312.5602, 2013.

[11] Wolf P, Hubschneider C, Weber M, et al. Learning how to drive in a real world simulation with deep q-networks[C]//2017 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2017: 244-250.

[12] Koenig N, Howard A. Design and use paradigms for gazebo, an open-source multi-robot simulator[C]//2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566). IEEE, 2004, 3: 2149-2154.

[13] Saxena D M, Bae S, Nakhaei A, et al. Driving in dense traffic with model-free reinforcement learning[C]//2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020: 5385-5392.

[14] Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms[J]. arXiv preprint arXiv:1707.06347, 2017.

[15] Ye F, Cheng X, Wang P, et al. Automated lane change strategy using proximal policy optimization-based deep reinforcement learning[C]//2020 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2020: 1746-1752.

[16] Hu Y, Yang J, Chen L, et al. Planning-oriented autonomous driving[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2023: 17853-17862.

[17] Zheng W, Song R, Guo X, et al. GenAD: Generative End-to-End Autonomous Driving[J]. arXiv preprint arXiv:2402.11502, 2024.

[18] Kendall A, Hawke J, Janz D, et al. Learning to drive in a day[C]//2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019: 8248-8254.

[19] Lillicrap T P, Hunt J J, Pritzel A, et al. Continuous control with deep reinforcement learning[J]. arXiv preprint arXiv: 1509.02971, 2015.