Knowledge Distillation from Single-Task Teachers to Multi-Task Student for End-to-End Autonomous Driving

Pedram Agand

Computer science, Simon Fraser University, Burnaby, Canada.
pagand@sfu.ca

Abstract

In the domain of end-to-end autonomous driving, conventional sensor fusion techniques exhibit inadequacies, particularly when facing challenging scenarios with numerous dynamic agents. Imitation learning hampers the performance by the expert and encounters issues with out-of-distribution challenges. To overcome these limitations, we propose a transformer-based algorithm designed to fuse diverse representations from RGB-D cameras through knowledge distillation. This approach leverages insights from multi-task teachers to enhance the learning capabilities of single-task students, particularly in a Reinforcement Learning (RL) setting. Our model consists of two primary modules: the perception module, responsible for encoding observation data acquired from RGB-D cameras and performing tasks such as semantic segmentation, semantic depth cloud mapping (SDC), ego vehicle speed estimation, and traffic light state recognition. Subsequently, the control module decodes these features, incorporating additional data, including a rough simulator for static and dynamic environments, to anticipate waypoints within a latent feature space. Vehicular controls (e.g., steering, throttle, and brake) are obtained directly from measurement features and environmental states using the RL agent and are further refined by a PID algorithm that dynamically follows waypoints. The model undergoes rigorous evaluation and comparative analysis on the CARLA simulator across various scenarios, encompassing normal to adversarial conditions. Our code is available at https://github.com/pagand/e2etransfuser/ to facilitate future studies.

Introduction and Research Questions

The objective of this research is to develop an end-to-end autonomous driving system that ensures safe navigation without violating traffic rules while avoiding collisions with dynamic obstacles. The research aims to address the challenges of perception and control in autonomous driving through the utilization of deep learning techniques and transformer-based sensor fusion. The key research questions include:

Can the proposed model outperform our previous approach, LeTFuser (Agand et al. 2023), in terms of route completion and driving score? How can the fusion of multiple RGB-D camera representations with the SDC map improve scene understanding and enhance the model's ability to handle complex driving scenarios? What benefits can be obtained from incorporating additional metadata, such as traffic light and ego vehicle speed information, into the perception module? How can the model be further improved through distillation modules and modified gradient normalization (MGN) for different driving tasks, as well as 2D attention-based fusion for knowledge augmentation from various sources? Can we utilize the transformer's capability to learn the environment and optimal policy as a sequence modeling problem instead of relying solely on imitation learning?

Related Work

The research builds upon the existing work of using deep learning and imitation learning for autonomous driving. The baseline method proposed in (Natan and Miura 2022a) serves as a foundation, employing EfficientNet for perception and GRU-based control modules. The LeTFuser model advances this baseline by incorporating transformer-based sensor fusion, utilizing convolutions vision transformer (CVT) (Wu et al. 2021) for local and global feature extraction, and introducing multi-task learning for simultaneous perception and control tasks (Agand et al. 2023). Additionally, the study draws inspiration from recent advancements such as the distillation model from (Jacob, Agarwal, and Stenger 2023), trajectory guided control (Wu et al. 2022), and the modified gradient normalization (MGN) (Natan and Miura 2022b) for improved task-based learning. Chen et al. (2021) introduced an innovative approach, transforming the RL problem into a sequential modeling problem by employing the reward-to-go technique. Nevertheless, it encounters challenges due to the inherent distinction between policy and environment Markov Decision Process (MDP), leading to a decline in performance in intricate scenarios such as autonomous driving. Concurrently, Janner, Li, and Levine (2021) exclusively concentrate on enhancing the maximum log-likelihood of all tokens rather than minimizing the Mean Squared Error (MSE) between predicted and ground-truth actions from the trajectory. They argue that learning to predict states and returns-to-go is crucial for optimal performance.

Proposed Research Plan

As of August 2023, significant progress has been made in the development of the model. It has built on top of our previous lightweight model, LeTFuser, successfully incorporating CvT for sensor fusion and multi-task learning, and evaluating its performance using the CARLA simulator. The control module is functional, utilizing a PID algorithm and direct prediction for vehicular controls. Since then, the fusion module has been enhanced, and the distillation technique on one task has been implemented. Evaluation results have demonstrated the model's capability to achieve higher driving scores. However, several aspects require further investigation and improvement. The research plan for the future includes as follows:

Implement distillation modules to improve the extraction of single-task teachers for the multi-task student in the perception module. Employ transformers to convert a reinforcement learning problem into a sequential model learning approach to improve imitation learning downsides. Develop a fusion module that utilizes 2D attention to augment knowledge from different sources, including RGB, SDC map, and simulation measurements. Evaluate the enhanced model using various scenarios (different weather, adversarial situations) in the CARLA simulator, comparing its performance against the most recent models.

Timeline

March 2022: Worked on smart traffic light control and proposed EcoLight (Agand, Iskrov, and Chen 2023), a reward shaping scheme for RL algorithms to reduce CO2 emissions while maintaining competitive results in other metrics like travel time.
September 2022: Undertook the second project on monocular depth estimation for autonomous driving and proposed DMODE, a class-agnostic method that estimates object distance without requiring object class information (Agand, Chang, and Chen 2022). Employed differential approaches to combine changes in object size over time and camera motion for distance estimation.
March 2023: Made significant progress in developing the perception and control modules of the LeTFuser model for end-to-end autonomous driving with transformer-based sensor fusion and multi-task learning (Agand et al. 2023).
May 2023: Developed and integrated the fusion module with 2D attention in the model for improved feature representation and knowledge augmentation.
September 2023: Implemented distillation modules and modified gradient normalization (MGN) to enhance perception learning in the model.
February 2024: Implemented the sequential modeling of the decision-making problem with the help of pretrained transformers to optimize vehicular command.
August 2024: Completed the research plan, including the final evaluation and comparative analysis of the model using the CARLA simulator.

Conclusion

Throughout my PhD journey, I have engaged in various projects, each contributing to the field of AI and autonomous driving. From smart traffic light control (EcoLight) and monocular depth estimation (DMODE) to end-to-end autonomous driving, each project has provided valuable insights and experience in different aspects. These diverse projects have prepared me for the challenges in developing a foundational/pretrained model to optimize the driving score and route completion for different driving situation.

References

Agand, P.; Chang, M.; and Chen, M. 2022. DMODE: Differential Monocular Object Distance Estimation Module without Class Specific Information. *arXiv preprint arXiv:2210.12596*.

Agand, P.; Iskrov, A.; and Chen, M. 2023. Deep Reinforcement Learning-based Intelligent Traffic Signal Controls with Optimized CO2 emissions. In *2023 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE.

Agand, P.; Mahdavian, M.; Savva, M.; and Chen, M. 2023. LeTFuser: Light-weight End-to-end Transformer-Based Sensor Fusion for Autonomous Driving with Multi-Task Learning. *arXiv preprint arXiv:2310.13135*.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.

Jacob, G. M.; Agarwal, V.; and Stenger, B. 2023. Online Knowledge Distillation for Multi-task Learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2359–2368.

Janner, M.; Li, Q.; and Levine, S. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286.

Natan, O.; and Miura, J. 2022a. End-to-end autonomous driving with semantic depth cloud mapping and multi-agent. *IEEE Transactions on Intelligent Vehicles*.

Natan, O.; and Miura, J. 2022b. Towards compact autonomous driving perception with balanced learning and multi-sensor fusion. *IEEE Transactions on Intelligent Transportation Systems*, 23(9): 16249–16266.

Wu, H.; Xiao, B.; Codella, N.; Liu, M.; Dai, X.; Yuan, L.; and Zhang, L. 2021. Cvt: Introducing convolutions to vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22–31.

Wu, P.; Jia, X.; Chen, L.; Yan, J.; Li, H.; and Qiao, Y. 2022. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. *arXiv preprint arXiv:2206.08129*.