Research Statement

Ke Guo

The rapid advancement of autonomous driving and intelligent transportation systems has revolutionized the field of transportation. To ensure transportation safety and efficiency, it is crucial to understand and accurately model human behavior in traffic systems. Specifically, a human-like autonomous driving model can enable surrounding drivers to better anticipate and avoid potential conflicts with autonomous vehicles. Moreover, an accurate human driving model can be employed to construct realistic traffic simulators, aiding transportation planners in analyzing and optimizing traffic flow through methods such as analyzing the impact of routing choices or learning a traffic-aware autonomous driving policy in the simulator. However, accurately modeling human behaviors is challenging due to the intricate decision-making processes influenced by various factors, including road conditions, weather, and traffic density. **I dedicate my research to accurately modeling and replicating human behaviors, particularly in complex and highly interactive urban traffic environments**. The current rule-based traffic model lacks accuracy due to its oversimplification, while learning-based models struggle to achieve long-term stability. **My goal is to apply control theory in developing learning-based traffic models that equip long-term, robust performance**.

1 Past Research

As an initial step towards my research goal, I **developed a rule-based crowd motion model** called VR-ORCA [1], which can navigate a crowd to their respective goals without collisions. I improve navigation performance by allowing a pair of agents to adjust their collision avoidance responsibilities based on the information from neighboring agents. Our key insight is that if an agent's motion is constrained by its neighbors, it should assume more collision avoidance responsibility, thereby enabling the other agent to utilize the potential free space (video). However, these heuristic rule-based models are oversimplified with limited parameters and may not capture the full complexity of human behavior in realistic environments. To overcome these limitations, my subsequent research focused on data-driven approaches that leverage machine learning techniques to accurately predict and model human behavior in real-world traffic environment.

Learning a human motion model is indeed a challenging task due to the multi-modal uncertainty arising from human intentions and external factors. In my research, I focused on addressing this challenge by **learning a model that can predict a multi-modal distribution of human future trajectories** [2]. Many previous approaches have utilized Gaussian mixture models to capture the multi-modal nature of trajectory distributions. However, optimizing such models can be difficult, and they are prone to overfitting. To overcome these limitations, I explore the potential of generative models in capturing complex distributions and generating diverse samples. Specifically, I propose a novel end-to-end conditional variational autoencoder (CVAE) model using a coarse future plan within a neighboring grid as a latent variable for interpretability. The plan generation process is achieved through the application of a value iteration network to achieve the end-to-end learning process. Furthermore, to ensure both diversity and precision in the predicted trajectory distribution, I employ an approximate symmetric cross-entropy loss during model training. Through extensive experimentation on two real-world datasets, our method outperforms all previous works in terms of prediction performance (video).

Building upon the predictive human driving model, I considered **developing a human-like autonomous driving model** [3]. However, I encountered a surprising result *a model with excellent open-loop prediction performance can have unsatisfactory closed-loop performance*. This discrepancy can be attributed to the "covariate shift" issue, where the state induced by the model's policy gradually deviates from the distribution exhibited by expert drivers. Existing solutions to this problem, such as DAgger or GAIL, require risky online interactions with real-world traffic systems during training. Instead, I propose a novel offline approach called context-conditioned imitation learning (CCIL). In CCIL, the model generates the ego vehicle's target trajectory based on its context state (the states of all other observed objects and ego goal positions), rather than producing its action relying on both ego and context states as in the classical BC method. The insight is that while the ego state is susceptible to learned policy errors, static elements in the context, such as lanes or crosswalks, remain unaffected by the ego vehicle and dynamic elements like human drivers attempt to recover from the perturbations from the ego vehicle. Based on an assumption of context stability, I prove a robustness assurance of the closed-loop system using control theoretic knowledge. In practice, simply removing historical ego information from the policy inputs is insufficient, as the choices made regarding the coordinate system could still inadvertently leak ego information and affect observation distributions. To mitigate this issue, I introduce an ego-perturbed goal-oriented coordinate system that minimizes the leaks caused by ego motion. In this coordinate system, the origin is the current position of the ego vehicle plus a zero mean Gaussian perturbation, while the x-axis points towards the ego vehicle's goal position. I validate our method using two real-world large-scale urban driving datasets, Lyft and nuPlan, and achieve state-of-the-art closed-loop performance in log-replay simulator (video).

Inspired by the impressive closed-loop autonomous driving performance of the CCIL method, I tried to scale it up to **develop a realistic learning-based traffic simulator** [4]. However, directly applying CCIL to the multi-agent imitation learning problem by sharing the same policy among all agents proved ineffective. This is because overlooking agents' historical trajectories in the policy's input will limit the policy's performance and realism. To address this challenge, I propose a novel approach called Learner-Aware Supervised Imitation Learning (LSAIL) method. By leveraging a variational autoencoder simultaneously modeling the expert and learner state distribution, our approach augments expert states such that the augmented state is aware of learner state distribution. By application of the learner-aware data augmentation method, I build the first traffic simulator based on imitation learning that can realistically replicate long-term microscopic urban traffic scenarios lasting more than 10 minutes. I demonstrate the effectiveness of our method on a real-world large-scale datasets, pNEUMA, achieving better short-term microscopic and long-term macroscopic similarity to real-world data than state-of-the-art baselines including SUMO (video).

In conclusion, my research is focused on finding principled solutions to address complex real-world traffic problems. I firmly believe that identifying and leveraging the unique structure and properties of the traffic system and developing mathematically-sound algorithms is crucial in tackling these challenges. One key aspect I have explored extensively is the closed-loop robustness of the traffic system. I have grounded this belief and dedicated my efforts to various aspects of traffic research, including traffic motion prediction, autonomous driving, and traffic simulation.

2 Future Research Directions

Motivated by the significant progress I have made in traffic modeling, I am excited to expand my research and construct a more sophisticated traffic model using recent advancements in machine learning. In pursuit of this goal, I plan to explore the following research directions:

Mixed Autonomy: The concept of mixed autonomy refers to the coexistence of autonomous vehicles (AVs) and human-driven vehicles on the road. As AV technology continues to advance, it is crucial to explore the implications and challenges associated with mixed autonomy. Future research can focus on studying human behavior difference in the presence of AVs. Developing effective traffic model in the mixed autonomy will be helpful to foster optimization algorithms and control strategies to maximize the benefits of mixed autonomy. This involves developing traffic-aware AVs considering factors like cooperative maneuvers, platooning, and coordination at intersections. Developing efficient traffic signal control systems that adapt to mixed autonomy scenarios can also enhance traffic flow and minimize congestion.

Interpretability and Controllability: Developing interpretable and explainable learning-based models is crucial for traffic prediction and simulation. Future research can focus on designing learning-based algorithms that provide causal reasoning processes, allowing users to understand the factors influencing traffic behaviors and outcomes. One potential approach is to leverage large language models, such as GPT-4, for labeling and explaining traffic behaviors. Moreover, incorporating human driving styles into traffic models can also enhance their interpretability and controllability. By integrating psychological and sociological factors, the models can capture realistic human interactions and responses to varying traffic conditions. This will greatly contribute to more interpretable and controllable traffic simulations.

Transfer Learning and Generalization: To ensure the scalability and applicability of developed models and methodologies, future research can investigate transfer learning techniques. By leveraging knowledge learned from several datasets, models can be applied to different scenarios or domains with limited labeled data. One potential avenue is the development of a large world model learned through self-supervised learning. This model would be trained on a diverse range of traffic datasets, capturing various traffic participant behaviors and scenarios. By learning from this rich and varied dataset, the model can acquire a holistic understanding of traffic dynamics and be able to predict the behavior of different types of traffic participants in new and unseen environments.

References

[1] K. Guo, D. Wang, T. Fan, and J. Pan, "Vr-orca: Variable responsibility optimal reciprocal collision avoidance," *IEEE Robotics and Automation Letters*, vol. 6, pp. 4520–4527, 2021.

[2] K. Guo, W. Liu, and J. Pan, "End-to-end trajectory distribution prediction based on occupancy grid maps," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2232–2241, 2022.

[3] K. Guo, W. Jing, J. Chen, and J. Pan, "CCIL: Context-conditioned imitation learning for urban driving," in *Robotics: Science and Systems (RSS)*, 2023.

[4] K. Guo, Z. Miao, W. Jing, W. Liu, W. Li, D. Hao, and J. Pan, "Lasil: Learner-aware supervised imitation learning for long-term microscopic traffic simulation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR, Under Review)*, 2024.