FLoRA: A Framework for Learning Scoring Rules in Autonomous Driving Planning Systems

Zikang Xiong, Joe Eappen, and Suresh Jagannathan

***Abstract*—In autonomous driving systems, motion planning is commonly implemented as a two-stage process: first, a trajectory proposer generates multiple candidate trajectories, then a scoring mechanism selects the most suitable trajectory for execution. For this critical selection stage, rule-based scoring mechanisms are particularly appealing as they can explicitly encode driving preferences, safety constraints, and traffic regulations in a formalized, human-understandable format. However, manually crafting these scoring rules presents significant challenges: the rules often contain complex interdependencies, require careful parameter tuning, and may not fully capture the nuances present in real-world driving data. This work introduces FLoRA, a novel framework that bridges this gap by learning interpretable scoring rules represented in temporal logic. Our method features a learnable logic structure that captures nuanced relationships across diverse driving scenarios, optimizing both rules and parameters directly from real-world driving demonstrations collected in NuPlan. Our approach effectively learns to evaluate driving behavior even though the training data only contains positive examples (successful driving demonstrations). Evaluations in closed-loop planning simulations demonstrate that our learned scoring rules outperform existing techniques, including expert designed rules and neural network scoring models, while maintaining interpretability. This work introduces a data-driven approach to enhance the scoring mechanism in autonomous driving systems, designed as a plug-in module to seamlessly integrate with various trajectory proposers. Our video and code are available on xiong.zikang.me/FLoRA/.**

***Index Terms*—Autonomous vehicle navigation, motion and path planning, machine learning for robotics, temporal logic.**

I. Introduction

large, intricate neural networks that incorporate elements of randomness, such as dropout or sampling from probability distributions. While powerful, such characteristics can make it challenging to predict the system's behavior consistently [6]. By implementing interpretable scoring rules as a final evaluation layer, we can assess these generated plans against clear, understandable criteria, thereby introducing much-needed predictability and reliability to these complex systems. In essence, scoring techniques serve as a critical bridge, connecting the raw outputs of planning algorithms to the final, executable plans. This additional evaluation step helps mitigate the uncertainties inherent in complex planning systems, significantly enhancing the safety, efficiency, and overall performance of autonomous vehicles as they navigate through our highways and cities.

With real-world driving data available in datasets like NuPlan [1], most current learning methods focus on directly learning motion plan proposers rather than learning interpretable scoring mechanisms to evaluate these plans. We instead focus on learning scoring rules represented in temporal logic for evaluating autonomous driving plans, which assess and rank plans generated by motion plan proposers. These scoring rules capture the latent relationships between various driving rules and constraints; for example, if a vehicle has a safe time-to-collision with surrounding vehicles, it should always be subject to all comfort constraints. By applying these rules to the output of a motion planner, we can score and select desirable plans, ensuring that the planned paths adhere to safety standards and traffic regu-