2025 IEEE International Conference on Robotics and Automation (ICRA)
May 19-23, 2025. Atlanta, USA

SELP: Generating Safe and Efficient Task Plans for Robot Agents with Large Language Models

Yi Wu, Zikang Xiong, Yiran Hu, Shreyash S. Iyengar, Nan Jiang,
Aniket Bera, Lin Tan, and Suresh Jagannathan

***Abstract*— Despite significant advancements in large language models (LLMs) that enhance robot agents' understanding and execution of natural language (NL) commands, ensuring the agents adhere to user-specified constraints remains challenging, particularly for complex commands and long-horizon tasks. To address this challenge, we present three key insights, *equivalence voting*, *constrained decoding*, and *domain-specific fine-tuning*, which significantly enhance LLM planners' capability in handling complex tasks. *Equivalence voting* ensures consistency by generating and sampling multiple Linear Temporal Logic (LTL) formulas from NL commands, grouping equivalent LTL formulas, and selecting the majority group of formulas as the final LTL formula. *Constrained decoding* then uses the generated LTL formula to enforce the autoregressive inference of plans, ensuring the generated plans conform to the LTL. *Domain-specific fine-tuning* customizes LLMs to produce safe and efficient plans within specific task domains. Our approach, Safe Efficient LLM Planner (SELP), combines these insights to create LLM planners to generate plans adhering to user commands with high confidence. We demonstrate the effectiveness and generalizability of SELP across different robot agents and tasks, including drone navigation and robot manipulation. For drone navigation tasks, SELP outperforms state-of-the-art planners by 10.8% in safety rate (i.e., finishing tasks conforming to NL commands) and by 19.8% in plan efficiency. For robot manipulation tasks, SELP achieves 20.4% improvement in safety rate. Our datasets for evaluating NL-to-LTL and robot task planning will be released in github.com/lt-asset/selp.**

I. INTRODUCTION

Recent advancements in large language models have sig-

multiple pre- and post-conditions, or tasks may span long time horizons, requiring flawless execution at each step. Typically, to evaluate a planner's ability to handle complex tasks, two critical metrics are considered: *safety*, defined as the planner's compliance with given commands, and *efficiency*, measured as the time at which a robot completes a task. With increasingly more complex commands, existing LLM planners produce more unsafe and inefficient plans, preventing them from being applied to complex or real-world domains. Fig. 1 shows an example where the user requires the drone to visit rooms with some constraints on the visiting order. GPT-4 generates an unsafe plan (shown in the purple block) that disobeys the constraints.

Another challenge appears when fine-tuning LLM planners with safety and efficiency objectives. These objectives can sometimes conflict, making it difficult for a model to learn to balance them. Safety often requires conservative planning, incorporating redundancies, and thorough checks to avoid errors, which can lead to longer execution times and lower efficiency. On the other hand, optimizing for efficiency typically involves minimizing the number of steps and the time taken to complete a task, which can increase the risk of unsafe plans.

SELP effectively addresses these limitations. Similar to the existing technique [11], SELP starts with translating NL into a set of LTL specifications as an intermediate representation. However, SELP provides confidence in the correctness of