

From Cooking Recipes to Robot Task Trees – Improving Planning Correctness and Task Efficiency by Leveraging LLMs with a Knowledge Network

Md. Sadman Sakib, and Yu Sun

Abstract—Task planning for robotic cooking involves generating a sequence of actions for a robot to prepare a meal successfully. This paper introduces a novel task tree generation pipeline producing correct planning and efficient execution for cooking tasks. Our method first uses a large language model (LLM) to retrieve recipe instructions and then utilizes a fine-tuned GPT-3 to convert them into a task tree, capturing sequential and parallel dependencies among subtasks. The pipeline then mitigates the uncertainty and unreliable features of LLM outputs using task tree retrieval. We combine multiple LLM task tree outputs into a graph and perform a task tree retrieval to avoid questionable nodes and high-cost nodes to improve planning correctness and execution efficiency. Our evaluation results show its superior performance in task planning accuracy and efficiency compared to previous works.

I. INTRODUCTION

Robotic cooking has emerged as a highly promising domain within robotics, presenting notable advantages such as convenience and the potential for enhanced efficiency and precision in meal preparation. To effectively automate cooking tasks, the key component is efficient task planning. This entails generating a series of actions guiding the robot in accomplishing a specific goal. However, this is an intricate field of research because cooking tasks typically involve lengthy sequences of actions encompassing various ingredients and tools. Moreover, they necessitate attaining numerous crucial ingredient states throughout the process. Additionally, the cooking conditions, processes, and requirements are exceptionally diverse. Approaches like state-space planning, learning from demonstration, and even knowledge network retrieval encounter difficulties when confronted with unseen starting conditions and requests.

In cooking tasks, ingredients or objects can vary in form, shape, and size, and there are multiple states to consider during recipe execution. The manipulation of an ingredient depends on its specific state, and certain ingredients may not be readily available in the desired state. Additionally, robots have varying capabilities, making some actions easier for them to perform than others. A task planning method should consider these factors and propose the most suitable plan for the robot to execute efficiently. Previous work has created a knowledge network consisting of 140 cooking recipes called

the Functional Object-Oriented Network (FOON) [1], [2]. However, generating plans in novel scenarios where FOON lacks knowledge about the recipe or an ingredient proved challenging. Furthermore, expanding the knowledge base was difficult due to the reliance on manual annotation.

Recently, the emergence of Large Language Models (LLMs) [3], [4], [5] has addressed the limitation of limited knowledge. These LLMs possess the ability to generate “likely” viable solutions for different scenarios and requests. While their results may not always be correct or optimal, their notable capacity for generalization can help overcome the limitations of search-based task tree retrieval methods. The search-based retrieval approach with a comprehensive knowledge network, on the other hand, can detect, remove, and replace the wrong elements in the LLM outputs.

The primary focus of this research paper is to tackle the task planning challenge in robotic cooking through the introduction of an innovative task tree generation approach (Figure 1). We aim to generate a task plan that is both error-free and cost-effective. To enhance the accuracy of the task plan, we employ a method that involves detecting incorrect components within the task trees generated by GPT-3 and searching for alternative options either within other task trees or within the FOON knowledge graph. This approach allows us to improve the overall quality and reliability of the generated task plan. The pipeline ensures effective resource utilization while achieving the desired objectives by carefully selecting the most optimized plan from these alternatives. The effectiveness of the task tree generation pipeline is evaluated through a comparative analysis with a previous approach. The results demonstrate the superiority of the proposed method, showcasing enhanced task planning accuracy and improved cost efficiency.




Fig. 1: Overview of our approach. Given a meal preparation instruction, the model generates a list of tasks specifically designed for the robot.

Yu Sun and Md. Sadman Sakib are members of the Robot Perception and Action Lab (RPAL), which is part of the Department of Computer Science & Engineering at the University of South Florida, Tampa, FL, USA.
Email: {madsadman, yusun}@usf.edu

Our contributions in this paper are as follows: (i) We

propose a novel task tree generation approach that accepts any dish of the user’s choice and produces a robot task tree with state-of-the-art accuracy and efficiency; (ii) We fine-tune GPT-3 to convert natural language instructions into a task tree structure; (iii) We improve the accuracy of the task plan by detecting incorrect components in the GPT generated task trees and finding alternatives either in other task trees or FOON; (iv) We optimize execution costs by performing weighted retrieval in a mini-FOON combined from multiple GPT outputs or FOON. We demonstrate the superiority of our model through a comparison with a previous approach.

II. OVERVIEW OF THE FUNCTIONAL OBJECT-ORIENTED NETWORK

A. Definitions

FOON and related knowledge graphs have been used in many tasks for robots, such as robotic cooking [6] and furniture assembly [7], [8], [9]. The one used here is a knowledge graph constructed through manual annotation of video demonstrations. It consists of two types of nodes: object nodes and motion nodes. These nodes are connected by directed edges, which depict the preconditions and effects of actions. The functional unit is the fundamental building block of FOON, representing a single action observed in the video demonstration. It consists of one or more input nodes, one or more output nodes, and a single motion node. The input nodes specify the required state of objects before the action, while the output nodes describe the resulting state after the action is executed. The motion node represents the action itself. Functional units provide a detailed and vivid representation of the actions observed in the video demonstrations. Figure 2 shows two functional units of slicing an onion and placing onion to cooking pan. The current FOON dataset (available in [10]) consists of 140 annotated recipes sourced from platforms such as YouTube, Activity-Net [11], and EPIC-KITCHENS [12].

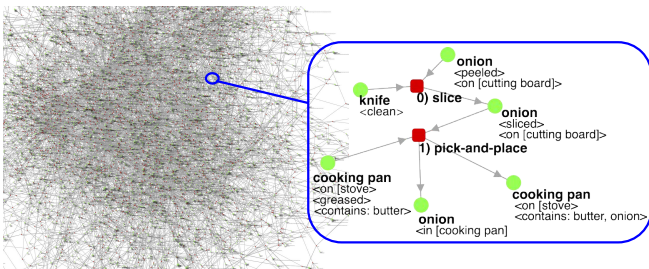


Fig. 2: Two functional units from FOON depicting slicing an onion and placing it to the cooking pan. Objects and motions nodes are denoted by green circle and red square respectively.

1) *Task Planning with FOON*: The utilization of FOON (or similar knowledge representation structures [13], [14]) as a knowledge base for task planning offers several advantages, including the ability to provide recipe variations. Task planning with FOON involves searching the network to find a goal node and retrieving a path, referred to as a

task tree, that leads to achieving the desired objective. A task tree consists of a sequence of functional units that need to be executed to prepare the dish. To illustrate, consider the task tree associated with boiling water, which comprises actions such as placing a pot on the stove, pouring water, turning on the stove, and turning off the stove. Each of these procedural steps is represented by input object nodes, signifying the prerequisites for executing the action; a motion node, denoting the action itself; and output object nodes, denoting the effect of executing the action. The task tree retrieval algorithm proposed in [1], [15] focuses on finding a path that utilizes only the ingredients available in the kitchen. On the other hand, [16], [17] retrieves a plan that can be executed with human-robot collaboration at home, while [17] learns assembly task plans from non-expert demonstrations using FOON. Nevertheless, these approaches have a limitation when it comes to generating a plan for a task that is not explicitly available in FOON. For example, if a user asks for a plan to prepare a mango milkshake, but there is no dedicated recipe for it in FOON, the system may be unable to provide a plan, even if there is a recipe for a banana milkshake. To address this limitation, a novel task tree retrieval method [18] was introduced that can learn from similar recipes in FOON and make necessary modifications to match the user’s requirements. While this approach introduces some level of generalization, the quality of the generated plan heavily relies on the availability of closely matched recipes in FOON. In this work, we leverage LLMs to overcome this dependency on closely matched recipes and generate high-quality task trees for any recipe, thereby enhancing the flexibility and effectiveness of the task planning process.

B. Related Works

In the domain of robotic cooking and task planning, several strategies have been proposed to tackle the challenges associated with generating effective action sequences for executing user instructions. One prominent approach involves the use of knowledge graphs to address this challenge. Notably, the KNOWROB framework [19], [20] has made significant contributions in this area by leveraging a knowledge base constructed from data collected in sensor-equipped environments. [21] introduced a task generalization scheme that relaxes the requirement of having multiple task demonstrations to perform tasks in unknown environments. This scheme integrates the task plan with a knowledge graph derived from observations in a virtual simulator. The impact of knowledge graphs on a robot’s decision-making process was further investigated in [22]. However, these approaches heavily rely on the limited information contained in their respective knowledge bases. In contrast, our approach harnesses the power of Language Models (LLMs) to alleviate the burden of creating a knowledge base, offering a more comprehensive and flexible solution.

Recently, task planning with LLMs has become a prominent area of research, capitalizing on LLMs’ impressive language understanding and generation capabilities. Various

studies have explored using LLMs to generate step-by-step plans for long-horizon tasks. For instance, Erra et al. [23] proposed an approach that employs LLMs to generate plans for complex tasks. [24], [25], [26] have also utilized LLMs for plan generation in different domains. However, these works often do not explicitly consider the robot’s capability to perform specific actions. One limitation of relying solely on LLMs is the lack of interaction with the environment. To address this limitation, SayCan [27] introduced a framework that combines the high-level knowledge of LLMs with low-level skills, enabling the execution of plans in the real world. By grounding LLM-generated plans with the robot’s capabilities and environmental constraints, SayCan bridges the gap between language-based planning and physical execution. In addition, recent research efforts such as Text2Motion [28], ProgPrompt [29] have integrated LLMs with learned skill policies. They exhibit trust in the LLM-generated plan and proceed to execute it, whereas we focus on enhancing LLM’s accuracy to generate an optimal task plan.

III. PROPOSED METHOD

Our objective is to develop a robust pipeline that generates highly accurate and executable task trees for robotic operations. To achieve this, we employ a multi-step approach that leverages the capabilities of LLMs and FOON. Initially, we utilize ChatGPT [30] to generate a recipe based on the user’s meal specifications. However, the output is in natural language, which may pose challenges for direct robot execution. To address this, we employ a fine-tuned GPT-3 model to convert the recipe instructions into a task tree format. Due to the uncertainty of the generative model, the task plan may not be always correct or most efficient. To enhance reliability and efficiency, we look for alternative options in other task trees generated by GPT-3 or in FOON. From these alternatives, the selected task tree is expected to be accurate and easier for the robot to execute. A visual representation of our pipeline and its key components are presented in Figure 3. In the following subsections, we will provide detailed explanations of each component.

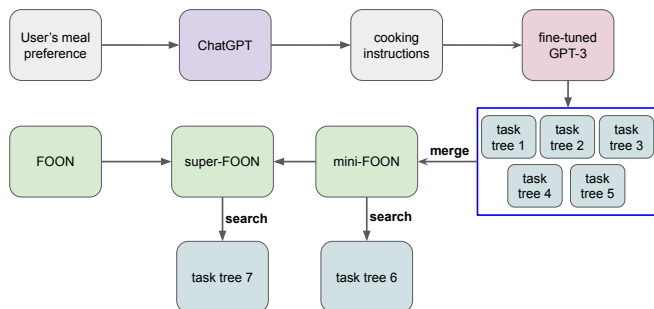


Fig. 3: Overview of our task tree generation procedure. Starting with a meal specification as the query, our pipeline generates a task plan represented as task tree 7.

A. Prompt engineering for recipe generation

Our system is designed to receive user specifications such as dish, a list of desired ingredients or exclude certain ingre-

dients, gluten-free, vegetable-based, or non-dairy options. We engine a prompt using the specifications and retrieve a recipe from ChatGPT. For easier parsing, we have designed the prompt to include numbered instructions within the response of ChatGPT.

B. Converting instructions to a task tree

When a robot performs an action, several factors need to be considered, such as preconditions, effects, and the objects involved. Additionally, understanding the state of these objects is crucial for determining the appropriate grasp or manipulation technique [31]. However, extracting all this information from a recipe written in natural language poses significant challenges. To address it, we propose translating the instructions into structured functional units that encapsulate all the necessary details. By organizing these functional units into a task tree, we provide a step-by-step guide for the robot to execute the task effectively. To accomplish this, we created a dedicated dataset for fine-tuning a GPT-3 Davinci model. This model takes a recipe as input and translates it into a task tree. The dataset comprises 180 recipe examples sourced from FOON, each consisting of natural language instructions and a corresponding FOON task tree. Due to the limitation of the maximum token count (at the time), some recipes had to be divided into multiple parts, resulting in multiple task trees for a single recipe.

C. Creating a mini-FOON

To address the potential presence of errors in the task plans generated by the fine-tuned model, we adopt a strategy of generating multiple task trees for the same recipe. Our aim is to search for a task tree that is error-free and one that is efficient for the robot to execute from the combined graph mini-FOON. FOON has revealed that merging recipes in a graph structure can lead to the emergence of novel cooking methods. This merging process allows recipes to share information and learn diverse approaches for accomplishing subtasks. Inspired by this idea of exploring new paths, we employ a similar graph structure to merge the five task trees generated by GPT. This merged structure is referred to as a mini-FOON.

1) *Merging task trees*: During the merging process, our objective is to eliminate any incorrect functional units and remove duplicates. An incorrect functional unit can arise in two ways: (i) syntax error and (ii) an erroneous object-action relationship. Syntax verification involves checking whether the functional unit includes the necessary components such as input and output objects, as well as a motion node. Additionally, it verifies if each object has an assigned state. On the other hand, validating the object-action relationship poses the challenge of determining if the state transition for an action is correct. To tackle this challenge, we compiled a comprehensive list of all valid state transitions from FOON. Based on this list, we can assess the correctness of a transition. For instance, if a transition such as “sliced \rightarrow whole” is not present in FOON, it would be identified as incorrect.

Functional units that successfully pass the verification criteria are then added to the mini-FOON.

D. Creating a super-FOON

We integrate the mini-FOON with the original FOON, forming a combined network known as the super-FOON. During this merging process, our primary focus is on node consolidation, as the mini-FOON and FOON may use different names for the same object or motion node. To achieve consolidation, we follow a set of basic rules. For instance, we convert all object names to their singular form. We observed that GPT-3 often generates plural forms such as “strawberries” and “onions” while FOON represents them as “strawberry” and “onion” respectively. By applying these rules, we try to ensure consistency and compatibility between the node names in the mini-FOON and FOON within the super-FOON network.

E. Task tree retrieval

Taking the desired dish as the goal node, we employ a search procedure similar to [16] to retrieve all paths leading to the goal. We execute the same search algorithm in both the mini-FOON and super-FOON. This approach often yields multiple task plans, exceeding five in number, which may differ in the number of cooking steps involved. For instance, when preparing a banana milkshake, one plan may suggest adding the whole peeled banana to the blender, while another plan may propose cutting the banana in half before blending. Once the incorrect functional units have been filtered out, the task tree retrieval procedure does not select them. Instead, it prioritizes the available correct functional units to construct the task plan. For instance, if the functional unit for “slicing an apple” is found to be incorrect in the first generated tree but correct in the other four task trees, the search procedure will choose the functional unit of slicing an apple from those four task trees. From the generated plans, we must select the most feasible one for the robot to execute. The feasibility of executing an action depends on the robot’s configuration. For example, a robot with only one hand may find pouring easier than chopping. Consequently, the success rate of executing a task tree varies among different robots. Following the approach of [16], we assign a cost value ranging from 0 to 1 to each action. These values are determined by three factors: 1) the physical capabilities of the robot, 2) its past experiences and ability to perform actions, and 3) the tools or objects that the robot needs to manipulate. A higher cost value indicates a more challenging action to execute. Based on these costs, we select task tree 6 from the mini-FOON and task tree 7 from the super-FOON. Ideally, task tree 7 should never be worse than task tree 6 since the super-FOON encompasses all the task trees from the mini-FOON. Task tree 7 serves as the final output of this pipeline. Figure 4 illustrates an example of cost optimization using the super-FOON, where two pouring actions are preferred over scooping due to the significantly lower cost assigned to pouring (0.1) compared to scooping (0.4). We assigned a

low cost to pouring based on the successful pouring accuracy achieved by Huang et al. [32], [33] with a robot.

IV. EXPERIMENTS AND RESULTS

Our experiment aims to assess both the quality of the generated task trees and the associated execution costs. Simultaneously, we seek to compare the performance of our model in generating recipes across different dish categories. To accomplish this, we curated a dataset consisting of 60 randomly selected recipes from the Salad, Drink, and Muffin categories. These recipes were extracted from Recipe1M+ [34], a comprehensive collection of over one million recipes encompassing a wide range of dish types and ingredients. For more experiments, please refer to our recent article [35] on consolidating task trees.

A. Evaluation Metric

Validating the plan of a cooking task in an automated manner is challenging due to the absence of a fixed method for preparing a dish. Two task plans for the same dish can differ in their cooking approaches, yet both can be deemed correct. As a result, manual verification becomes necessary. However, the original format of a task tree can be difficult for humans to comprehend. To address this, we convert the task trees into progress lines as used in [18] to illustrate how the ingredients are manipulated and undergo changes throughout the cooking process. This simplified visualization facilitates the detection of errors in the task plan by humans. We consider a recipe correct if the progress lines for all ingredients used in the recipe are accurate. A progress line example for a Greek Salad recipe is provided in Figure 5.

B. Task Planning Accuracy

We employed four different methods to generate task trees for the selected recipes. The quality of the generated trees was assessed using the progress line, and the corresponding accuracy results are shown in Figure 6. When relying solely on FOON, the task trees obtained for Salad and Drink recipes exhibited good quality. This was expected as FOON contained an ample number of recipes (10 each) for these categories. However, for Muffin recipes, the quality of the generated task trees suffered due to the scarcity of available examples in FOON (only one recipe). The FOON-search based approach heavily depends on finding a similar recipe in FOON as a reference for making necessary modifications to the task plan. Consequently, a high number of adjustments were required, leading to inaccuracies in the task plan. In the case of the fine-tuned GPT-3 model, errors in functional units frequently resulted in task plan failures. However, the introduction of the Mini-FOON helped mitigate these errors by providing a wider range of alternatives to achieve the desired objectives. Integrating FOON into our approach enabled us to choose a path from a broader set of options, resulting in higher accuracy. Compared to [18], our approach achieved a 4% higher accuracy for Salad, 6% higher accuracy for Drink, and a significant 45% higher accuracy for Muffin recipes. Notably, our fine-tuned model demonstrated good

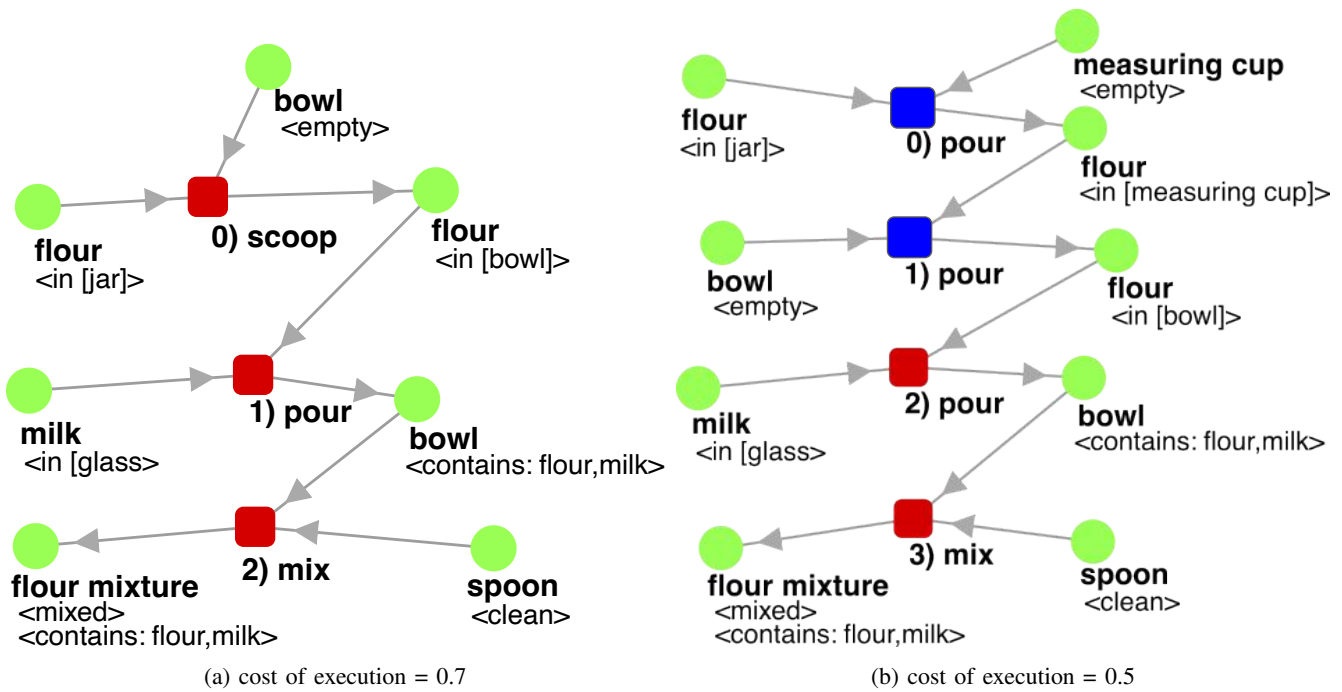


Fig. 4: Example of cost optimization: Comparison between task trees retrieved from the mini-FOON and super-FOON. The assigned costs for scooping, pouring, and mixing are 0.4, 0.1, and 0.1 respectively. (a) The task tree from the mini-FOON (b) The task tree from the super-FOON.



Fig. 5: Progress lines for a Greek Salad recipe.

accuracy for Muffin recipes, despite not being specifically trained on this particular dish. This highlights the significant advantage of employing an LLM. Once the LLM is fine-tuned to comprehend the structure of a task tree, it can effectively generalize to various types of recipes.

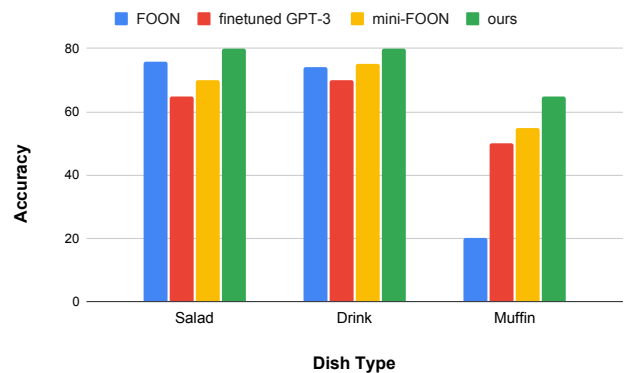


Fig. 6: Comparison of different approaches' accuracy on Salad, Drink, and Muffin dishes.

C. Execution cost

The objective of this experiment is to evaluate the extent to which our approach can optimize the execution cost of recipes. If a recipe cannot be optimized, it implies that there are no superior alternatives in FOON compared to the initial output generated by the fine-tuned model (task tree 1). In Figure 7, we present the number of optimized recipes by generating different numbers of task trees. When the number of task trees is 2, and we select the plan with the lower cost, it yields a better solution in 5% of the cases. Similarly, by gradually increasing the number of task trees up to 5 and selecting the one with the minimum cost, we obtain a better solution in 15% of the cases. More optimization occurs when we choose task tree 6 from the Mini-FOON, as it

combines subtasks from five different task trees, resulting in a lower cost. Ultimately, task tree 7, the final output from our pipeline, maximizes the advantages of FOON and minimizes the execution cost compared to task tree 1 in 40% cases.

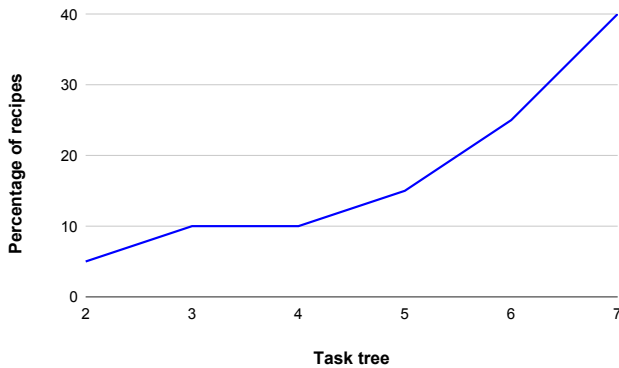


Fig. 7: Number of recipes that were optimized by generating varying numbers of task trees in comparison to Task Tree 1 (generated by the fine-tuned model).

V. DISCUSSION

A. Finetuning a GPT-3

We examined how the model’s understanding of the task tree structure improves with the addition of new training data (Figure 8). Initially, the training began with a dataset consisting of only 30 examples. Consequently, the model struggled to grasp the syntax of functional units, resulting in grammatical errors in the generated functional units. For instance, it would include multiple motion nodes within a single functional unit, whereas, according to the definition, a functional unit should contain only one motion node. As we increased the number of recipes, the model gradually reduced its syntactical errors. However, it still exhibited logical mistakes, such as incorrect state transitions or missing actions. Finally, after finetuning with 180 examples, the model achieved an accuracy of 67%.




Fig. 8: Impact of training dataset size on model accuracy.

B. Executing a task tree

A task tree provides a high-level plan that lacks interaction with the environment. However, executing actions often requires additional information, such as the geometric position

of objects or the initial quantity of ingredients in a container. For instance, a task plan might involve adding ice to an empty glass, but the glass could be positioned upside down on a table. Therefore, before pouring the ice, the glass would need to be rotated back to its original position. This crucial step is missing in our high-level planning. Hence, there is a need for hierarchical planning, where the task tree can be converted into a low-level plan that can be executed in the real world. Paulius et al. [36] proposed a method to convert a task tree into a representation using Planning Domain Definition Language (PDDL) [37]. Each functional unit is treated as a planning operator, and a plan is generated based on the robot’s low-level motion primitives.

C. Limitations of our approach

(i) The generation of a task tree involves making 5 API calls to the fine-tuned model. Each API call takes approximately 5 seconds, resulting in a slow pipeline. The focus of this research was not on time complexity. In the future, if we aim to enhance the system’s speed, it may be necessary to explore fine-tuning locally installed LLMs. (ii) The generated plan sometimes introduces new names for ingredients, states or motions such as garnish. These unknown labels in functional units pose a challenge when attempting to find alternative options in FOON, as proper mapping to existing functional units becomes difficult. Furthermore, the detection of incorrect transitions is also hindered, as the possible transition list may not include these new labels. (iii) A fine-tuned GPT-3 model has a limitation where the combined query and completion cannot exceed 2048 tokens. Due to this constraint, generating a task tree becomes challenging when dealing with complex recipes that require a higher number of functional units.

VI. CONCLUSION

In this study, our objective was to propose a novel pipeline for task tree generation, leveraging the advantages offered by LLMs. We utilized ChatGPT to respond to user queries, and then fine-tuned a GPT-3 model to convert the response into a task tree representation. To enhance the accuracy and execution cost of the task tree, we integrated the output of the fine-tuned model with FOON, exploring multiple possibilities to achieve the desired objectives. Through our experiments, we demonstrated its superior performance, highlighting its remarkable generalization capabilities. In future, we intend to focus on addressing the challenges of task tree correction and re-planning in cases of planning or execution failures. It is worth noting that our pipeline exhibits a high degree of flexibility, allowing for the seamless substitution of GPT and FOON with more advanced Language Models or knowledge networks. We aim to incorporate image inputs into our system by utilizing the newly released GPT-4, which can handle both textual questions and accompanying images. This would allow users to upload images of dishes and inquire about their preparation methods.

REFERENCES

- [1] D. Paulius, Y. Huang, R. Milton, W. D. Buchanan, J. Sam, and Y. Sun. Functional Object-Oriented Network for Manipulation Learning. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 2655–2662. IEEE, 2016.
- [2] Md Sadman Sakib, Hailey Baez, David Paulius, and Yu Sun. Evaluating Recipes Generated from Functional Object-Oriented Network. *arXiv preprint arXiv:2106.00728*, 2021. (Featured in *18th International Conference on Ubiquitous Robots (UR 2021)*).
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [5] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. *ArXiv*, abs/2109.01652, 2021.
- [6] Kota Takata, Takuya Kiyokawa, Ichel G. Ramirez-Alpizar, Natsuki Yamanobe, Weiwei Wan, and Kensuke Harada. Efficient task/motion planning for a dual-arm robot from language instructions and cooking images. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12058–12065, 2022.
- [7] Ross A. Knepper, Todd Layton, John Romanishin, and Daniela Rus. IkeaBot: An autonomous multi-robot coordinated furniture assembly system. In *2013 IEEE International Conference on Robotics and Automation*, pages 855–862, 2013.
- [8] Yizhak Ben-Shabat, Xin Yu, Fatemeh Sadat Saleh, Dylan Campbell, Cristian Rodriguez-Opazo, Hongdong Li, and Stephen Gould. The IKEA ASM Dataset: Understanding People Assembling Furniture through Actions, Objects and Pose. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 846–858, 2021.
- [9] Alexandros Vassiliades, Nikos Zarkadas, Nick Bassiliades, and Theodore Patkos. Onto-IKEA: A Knowledge Retrieval Framework Based on IKEA Ontology. In *JOWO*, 2021.
- [10] FOON Website: Graph Viewer and Videos. <http://www.foonets.com>, 2021. Accessed: 2023-06-06.
- [11] Bernard Ghanem Fabian Caba Heilbron, Victor Escorcía and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015.
- [12] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In *European Conference on Computer Vision (ECCV)*, 2018.
- [13] David Paulius and Yu Sun. A survey of knowledge representation in service robotics. *Robotics and Autonomous Systems*, 118:13–30, 2019.
- [14] Yu Sun, Shaogang Ren, and Yun Lin. Object–object interaction affordance learning. *Robotics and Autonomous Systems*, 62(4):487–496, 2014.
- [15] David Paulius, Ahmad B Jelodar, and Yu Sun. Functional Object-Oriented Network: Construction & Expansion. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5935–5941. IEEE, 2018.
- [16] D. Paulius, K. S. P. Dong, and Y. Sun. Task Planning with a Weighted Functional Object-Oriented Network. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [17] Yi Chen, David Paulius, Yu Sun, and Yunyi Jia. Robot learning of assembly tasks from non-expert demonstrations using functional object-oriented network. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 2012–2019. IEEE, 2022.
- [18] Md. Sadman Sakib, David Paulius, and Yu Sun. Approximate task tree retrieval in a knowledge network for robotic cooking. *IEEE Robotics and Automation Letters*, 7:11492–11499, 2022.
- [19] M. Beetz, D. Beßler, Andrei Haidu, M. Pomarlan, A. Bozcuoğlu, and Georg Bartels. KnowRob 2.0 — A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 512–519, 2018.
- [20] Michael Beetz, Ulrich Klank, Ingo Kresse, Alexis Maldonado, Lorenz Mösenlechner, Dejan Pangercic, Thomas Rühr, and Moritz Tenorth. Robotic roommates making pancakes. In *2011 11th IEEE-RAS International Conference on Humanoid Robots*, pages 529–536, 2011.
- [21] Angel Andres Daruna, Lakshmi Nair, Weiyu Liu, and S. Chernova. Towards robust one-shot task execution using knowledge graph embeddings. *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11118–11124, 2021.
- [22] Angel Andres Daruna, Devleena Das, and S. Chernova. Explainable knowledge graph embedding: Inference reconciliation for knowledge inferences supporting robot actions. *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1008–1015, 2022.
- [23] Chao Zhao, Shuai Yuan, Chunli Jiang, Junhao Cai, Hongyu Yu, Michael Wang, and Qifeng Chen. Erra: An embodied representation and reasoning architecture for long-horizon language-conditioned manipulation tasks. *IEEE Robotics and Automation Letters*, PP:1–8, 06 2023.
- [24] Wenlong Huang, P. Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *ArXiv*, abs/2201.07207, 2022.
- [25] Dhruv Shah, Blazej Osinski, Brian Ichter, and Sergey Levine. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*, 2022.
- [26] Andy Zeng, Adrian S. Wong, Stefan Welker, Krzysztof Choromanski, Federico Tombari, Aavek Purohit, Michael S. Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, and Peter R. Florence. Socratic models: Composing zero-shot multimodal reasoning with language. *ArXiv*, abs/2204.00598, 2022.
- [27] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Jayant Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego M Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, 2022.
- [28] Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. Text2motion: From natural language instructions to feasible plans. *arXiv preprint arXiv:2303.12153*, 2023.
- [29] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. ProgPrompt: Generating situated robot task plans using large language models. In *International Conference on Robotics and Automation (ICRA)*, 2023.
- [30] OpenAI. Chatgpt. <https://openai.com/>, 2021. Accessed on 8th June 2023.
- [31] Ahmad B Jelodar and Yu Sun. Joint Object and State Recognition using Language Knowledge. In *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019.
- [32] Yongqiang Huang and Yu Sun. Learning to pour. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7005–7010. IEEE, 2017.
- [33] Yongqiang Huang, Juan Wilches, and Yu Sun. Robot gaining accurate pouring skills through self-supervised learning and generalization. *Robotics and Autonomous Systems*, 136:103692, 2021.
- [34] Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1M+: A Dataset for Learning Cross-Modal Embeddings for Cooking Recipes and Food Images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [35] Md Sadman Sakib and Yu Sun. Consolidating trees of robotic plans generated using large language models to improve reliability. *International Journal of Artificial Intelligence and Robotics Research*, February 2024.
- [36] David Paulius, Alejandro Agostini, and Dongheui Lee. Long-horizon planning and execution with functional object-oriented networks. 2022.
- [37] Drew McDermott, Malik Ghallab, Adele E. Howe, Craig A. Knoblock,

Ashwin Ram, Manuela M. Veloso, Daniel S. Weld, and David E. Wilkins. Pddl-the planning domain definition language. 1998.