거대 언어 모델을 이용한 뉴로-심볼릭 작업 재계획법

Neuro-Symbolic Task Replanning using Large Language Models

권 민 서[1] · 김 영 준[†]
Minseo Kwon[1], Young J. Kim[†]

Abstract: We introduce a novel task replanning algorithm that combines a symbolic task planner with a multimodal Large Language Model (LLM). Our algorithm starts by describing the scene by extracting the semantic and spatial relationships of objects in the environment through a multimodal LLM and an open-vocabulary object detection model. Then, the LLM formulates a planning problem in symbolic form based on the scene description and the user's goal description, which are then processed by the symbolic planner to create task plans. These plans are converted into low-level executable codes for the robot, with the LLM performing syntax and semantic checks to ensure validity and facilitate replanning if necessary. We demonstrate the application of our replanning pipeline using dual UR5e manipulators in various benchmark tasks, including pick-and-place operations, block-stacking, and block rearrangement.

Keywords: Task Planning, Large Language Models, Replanning

1. Introduction

Integrating Large Language Models (LLMs) into robotic task planning presents promising results due to their enhanced commonsense reasoning and generalization abilities. LLMs can interpret unstructured input, providing a detailed understanding of tasks based on natural language and enabling more intuitive interactions between robots and humans.

Previous research leveraging LLMs in robotics has largely focused on extensive prompting strategies to directly obtain action sequences from models. However, due to the frequent generation of hallucinatory outputs, LLMs excel in aiding in solving planning tasks, such as using an LLM output to guide a search-based planner rather than solving tasks on their own[1-3].

Consequently, significant research has focused on integrating LLMs with symbolic planning formulations like Planning Domain Definition Language (PDDL)[4] to improve reliability. Some approaches leverage the commonsense knowledge of LLMs to aid in formulating planning problems[5-8], while others use LLMs to supplement the results of symbolic planners[9].

Furthermore, methods such as iteratively reprompting LLMs with feedback from external evaluators have also been employed to refine outputs, improving both task relevance and execution accuracy[10,11]. This approach helps resolve one of the main challenges in LLM-based task planning—errors during the planning process, such as syntax or semantic issues in the problem specification[7], which can result in a failed plan. Without replanning, these errors would cause the system to halt, requiring manual intervention. In real-world scenarios, where continuous adaptation by the robot is critical, replanning is essential for automatically detecting and solving these issues.

In this paper, we propose a novel task replanning pipeline that takes advantage of LLMs – their commonsense reasoning and capacity to process natural language queries - while mitigating

1. M.S. Student, Dept. of Computer Science and Engineering, Ewha Womans University, Seoul, Korea (minseo.kwon@ewha.ac.kr)
† Professor, Corresponding author: Dept. of Computer Science and Engineering, Ewha Womans University, Seoul, Korea (kimy@ewha.ac.kr)

their limitations by combining a symbolic planner with a multimodal LLM and employing a replanning method when planner failures occurred. Our approach uses a multimodal LLM to analyze semantic relationships among target objects in the scene and a 2D open-vocabulary object detection model to obtain the geometric information. Based on the scene description and user-provided goal description, the LLM encodes the planning problem using PDDL formulation. Given this problem PDDL and domain PDDL, a symbolic PDDL planner finds a plan for the task, which is later translated into low-level code such as Python code with action parameter selection to invoke motion planning for robot execution.

Throughout this whole process, replanning occurs when a failure is detected. When a symbolic PDDL planner fails to find a valid plan, LLM validates syntax and semantic errors in the problem PDDL, and we replan the task by reprompting the error messages to LLM as error feedback. If an exception occurs when executing the Python code, we replan the task by reprompting the exception messages to LLM.

We also conducted experiments using a dual robot manipulator setup and an LLM across three task planning domains to demonstrate the effectiveness of our pipeline on diverse and complex robotic tasks.

In summary, the main contributions of our work are:

- We propose a novel neuro-symbolic task replanning algorithm for executing various robotics tasks by combining symbolic planner and multimodal LLM and enabling LLM as a syntax and semantic checker.
- We demonstrated the utility of our algorithmic pipeline on real-world robotic tasks. We also highlighted cases where planning failures occurred and showed how the LLM refined these failures.
- We also demonstrated the effectiveness of replanning, showing that success rates increased by up to 31.92% compared to cases without replanning while observing how the occurrence of each failure cause decreases.

The rest of the paper is structured as follows. In Section 2, we overview relevant work to robot task planning using traditional methods and LLM-based methods. In Section 3, we outline the overall task replanning pipeline in four steps. In Section 4, we present the task planning results in real-world scenarios. Lastly, we conclude the paper and discuss the future work in Section 5.

2. Related Works

2.1 Symbolic Task Planning

Symbolic task planning is based on a classical approach in AI planning, where tasks are represented using symbolic languages, such as PDDL[4]. This approach defines the world as a set of logical states and actions, where each action has preconditions and effects that change the state. Symbolic planners search these states using heuristics to generate a sequence of actions that achieves a given goal.

Additionally, advancements in hierarchical planning and the integration of symbolic task planning with motion planning (TAMP) have further improved performance. A key factor for generalizing TAMP approaches lies between discrete *high-level* task planning and continuous *low-level* motion planning. This involves selecting hybrid action parameters that satisfy constraints, such as how to grasp or where to place an object. The downward refinement property assumes that every solution to the high-level task plan has a corresponding low-level motion solution. However, this assumption often does not hold in real-world scenarios, limiting the practical application of TAMP in dynamic, real-world environments[12].

2.2 LLM-based Robot Task Planning

LLM-based robot task planning utilizes LLMs to enhance understanding of the real world and generate action sequences that lead to the goal. For instance, [13] demonstrated that LLMs can effectively combine language understanding with action grounding, allowing robots to execute tasks based on their capabilities and real-world affordances. [14] also employed LLMs, prompting them with environmental state feedback to generate task plans in Python code incorporating robot action primitives. Additionally, research has focused on frameworks that utilize LLMs to create spatial relationships between objects and provide motion planning feedback, addressing TAMP problems[15,16]. However, despite these advantages, LLMs often encounter difficulties with large-scale tasks and may yield unreliable outputs in complex scenarios[2]. There has also been a study that divided a given goal into multiple subgoals to prevent a significant increase in planning time for large-scale tasks, solving each subgoal using either symbolic or LLM-based planners[17].
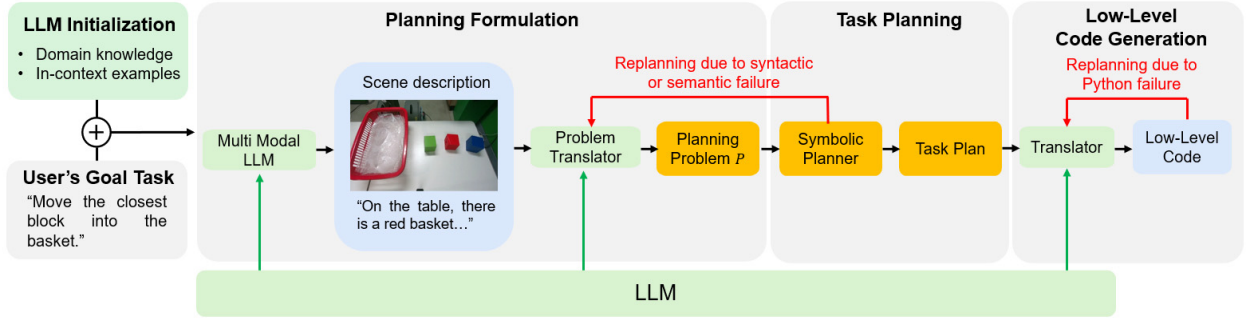
[Fig. 1] Our neuro-symbolic task replanning pipeline. The green blocks represent the use of LLM, and the orange blocks represent symbolic planning using symbolic languages. Red arrows show two cases of replanning, where the first arrow indicates syntax/semantic errors in problem formulation, and the second one contains Python exceptions in low-level codes

Integrating LLM with symbolic planners has also been a significant research area. [6] and [18] utilized LLMs as translators, converting natural language problem descriptions into PDDL problems through few-shot prompting[19]. [7] addressed TAMP problems by translating natural language problem descriptions into STL problem specifications and correcting their syntax and semantic errors through reprompting. However, these methods are not directly applicable to robot manipulation tasks in real-world scenes where planning problems are not provided in natural language.

Building on this, [5] combined LLMs with object detection and image captioning models to generate a problem specification, then solved by a symbolic PDDL planner. [9] presented a framework leveraging LLM's commonsense knowledge in household environments to reduce plan length generated by a symbolic PDDL planner. However, these approaches are limited as they primarily address task planning and do not extend to the low-level details required for robot execution, such as action parameter selection.

2.3 Corrective Reprompting with LLM

The concept of corrective reprompting [20] with LLMs has been explored to address planning errors caused by the LLM's hallucinatory outputs. [21] focused on detecting unmet action precondition errors in LLM-generated plans and reprompting the LLM to adjust actions accordingly. [11] incorporated the PDDL plan validator VAL[22] to identify errors in LLM-generated plans and refine them through interactive debugging. Additionally, [7] employed a rule-based STL syntax checker for syntax errors alongside an LLM correction module for semantic errors.

In contrast, our approach utilizes a symbolic planner to ensure plan success without relying on external syntax checkers or verifiers. Moreover, we address both syntax and semantic errors in the planning problem formulation using a zero-shot approach with our automatic LLM reprompting.

3. Task Replanning Pipeline

Our task planning algorithm consists of four parts: planning formulation, task planning with symbolic planner, low-level code generation, and replanning with syntax and semantic checking. The overall pipeline is shown in [Fig. 1]. We will go through each step in the following subsections.
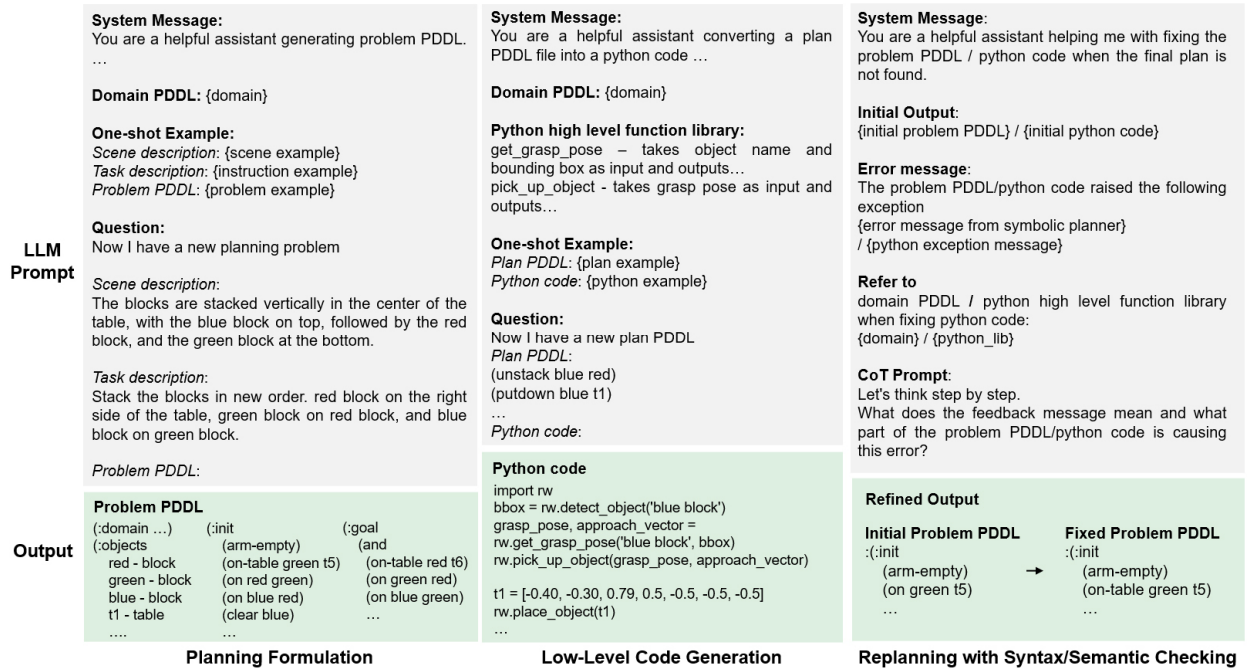
3.1 Planning Formulation

Our objective is to generate a problem PDDL, which can be formulated as below:

$$P \equiv\ <S, O, A, T, s_0, S^* >\ \tag{1}$$

where S is a finite set of all possible fully observable states, O is a set of environment objects, A is a finite set of all actions, $T: S \times A \to S$ is a deterministic state transition function, $s_0 \in S$ is an initial state, and $S^* \subset S$ is a set of possible goal states.

To enable the robot to interpret the initial scene and encode it into PDDL, we need to gather information about the types of objects in the environment and their spatial relationships. We use GPT-4-Turbo as multimodal LLM to simultaneously understand image and text prompts. By providing a color image from the

System Message:
You are a helpful assistant generating problem PDDL.
…

Domain PDDL: {domain}

One-shot Example:
Scene description: {scene example}
Task description: {instruction example}
Problem PDDL: {problem example}

Question:
Now I have a new planning problem

Scene description:
The blocks are stacked vertically in the center of the table, with the blue block on top, followed by the red block, and the green block at the bottom.

Task description:
Stack the blocks in new order. red block on the right side of the table, green block on red block, and blue block on green block.

Problem PDDL:

Problem PDDL
```
(:domain …)        (:init              (:goal
(:objects          (arm-empty)         (and
  red - block      (on-table green t5)     (on-table red t6)
  green - block    (on red green)          (on green red)
  blue - block     (on blue red)           (on blue green))
  t1 - table       (clear blue)            …
  ….               …
```

Planning Formulation

System Message:
You are a helpful assistant converting a plan PDDL file into a python code …

Domain PDDL: {domain}

Python high level function library:
get_grasp_pose – takes object name and bounding box as input and outputs…
pick_up_object - takes grasp pose as input and outputs…

One-shot Example:
Plan PDDL: {plan example}
Python code: {python example}

Question:
Now I have a new plan PDDL
Plan PDDL:
(unstack blue red)
(putdown blue t1)
…
Python code:

Python code
```
import rw
bbox = rw.detect_object('blue block')
grasp_pose, approach_vector =
rw.get_grasp_pose('blue block', bbox)
rw.pick_up_object(grasp_pose, approach_vector)

t1 = [-0.40, -0.30, 0.79, 0.5, -0.5, -0.5]
rw.place_object(t1)
…
```

Low-Level Code Generation

System Message:
You are a helpful assistant helping me with fixing the problem PDDL / python code when the final plan is not found.

Initial Output:
{initial problem PDDL} / {initial python code}

Error message:
The problem PDDL/python code raised the following exception
{error message from symbolic planner}
/ {python exception message}

Refer to
domain PDDL / python high level function library when fixing python code:
{domain} / {python_lib}

CoT Prompt:
Let's think step by step.
What does the feedback message mean and what part of the problem PDDL/python code is causing this error?

Refined Output

Initial Problem PDDL   **Fixed Problem PDDL**
:(:init   :(:init
  (arm-empty)   →   (arm-empty)
  (on green t5)   (on-table green t5)
  …   …

Replanning with Syntax/Semantic Checking

LLM Prompt

Output

[Fig. 2] An example of LLM inputs and outputs of each planning step. Using these prompt configurations, the LLM produces the problem PDDL, low-level code, and refined problem PDDL or low-level code in cases of detected failures

robot's perspective together with the prompt, "*What objects are on the table? Tell me each of their appearance and spatial relationships.*", the LLM can generate a scene description about the objects on the table, including their relative position, and spatial relationships.

Using this scene description, along with the user-provided goal task, domain PDDL, and an in-context example, the LLM formulates the planning problem P. Moreover, instead of relying on multiple in-context examples, we use one-shot prompting [19] to improve the LLM's output[3]. Furthermore, LLM can obtain information about PDDL predicates from the domain PDDL. A detailed example of the prompt is provided in [Fig. 2].

The objects identified in the scene description become the set O (e.g., *(:objects red - block green - block blue - block)*), later serving as parameters for PDDL actions and predicates. From the spatial relationships between objects along with the positions of the objects (e.g., "*with the blue block on top, followed by the red block, and the green block at the bottom.*"), we translate it into predicates (e.g., *(on red green) (on blue red)*), and this set of predicates forms the initial state s_0. Additionally, based on the user-provided goal task, the LLM translates the natural language goal into a PDDL goal description (e.g. *(on-table red t6) (on green red) (on blue green)*), forming S^*. The remaining components S, A and T are derived from the domain PDDL.

3.2 Task Planning with Symbolic Planner

Once the planning problem P is formulated, the objective is to use a symbolic task planner to find a policy $\pi = \{a_1, \cdots, a_n | \forall a_i \in A\}$ for P. The generated problem PDDL and domain PDDL are put into a search-based symbolic planner to produce a plan PDDL. We utilize the Fast Downward planner[23], specifically employing the "*seq-opt-fdss-1*" configuration.

Planning is successful if the Fast Downward planner generates a PDDL plan starting from the initial state s_0 and reaching one of the possible goal states $s_g \in S^*$ within the given search time limit. The planning attempt is considered unsuccessful if the planner fails to generate such a plan within the time limit.

3.3 Low-Level Code Generation

As mentioned in Section 2.1, to execute the high-level plan obtained from task planning, it is necessary to search for hybrid action parameters that satisfy constraints and then call the motion planner[12]. Similarly, in our pipeline, the plan PDDL generated by the symbolic planner is converted into Python code. We

prompt the LLM to translate each action a_i from the plan PDDL into predefined Python action primitives, such as '*pick_up _object*' and '*place_object*'[24]. A detailed example of the prompt for LLM translator is shown in [Fig. 2]. By calling these primitives, the motion planner [25] is invoked, and the robot will execute the path. While the PDDL actions are high-level, requiring discrete and semantic parameters such as object names, these action primitives are low-level, requiring continuous and real-valued parameters, such as the grasp or place poses for each object. Therefore, the intermediate process of selecting continuous action parameters for Python action primitives is necessary.

In most TAMP systems where the downward refinement property is not satisfied, these parameters must satisfy constraints like collision avoidance and robot joint limits. If any constraints are violated, the system backtracks and tries alternative high-level plans[12]. But for simplicity, we assume that the downward refinement property holds, which means no such constraints exist in our real-world scene.

The process for determining action parameters is as follows. We use a 2D open-vocabulary object detection model to compute the bounding boxes of target objects. Leveraging the scene description provided by the multimodal LLM, the LLM assigns names to objects, and these names, along with the detection model, help generate 2D bounding boxes. These 2D bounding boxes are then expanded into 3D by integrating depth information from segmented object masks captured by an RGB-D camera from the robot's perspective. We employ Grounded-Segment-Anything [26] as the object detection model, which combines Grounding DINO [27] and Segment Anything [28]. This step corresponds to the Python function '*detect_object*' as in [Fig. 2]. Then, a grasp pose selection algorithm [29] is applied within the 3D bounding boxes, corresponding to the function '*get_grasp_pose*' as in [Fig. 2]. The resulting grasp pose serves as the continuous parameter for the '*pick_up_object*' action. For the '*place_object*' action, continuous parameters are set based on predefined positions for different table sections (e.g., *t1 = [-0.40, -0.30, 0.79, 0.5, -0.5, -0.5, -0.5], t2 = ⋯*).

3.4 Replanning with Syntax and Semantic Checking

Given that the LLM may produce erroneous outputs, we have integrated an automatic replanning module to prevent program interruptions caused by failures. This module detects planning failures and reprompts the LLM to resolve the issues. Failures within the pipeline typically stem from two main sources: errors in problem PDDL generation and low-level code generation.

Errors in problem PDDL generally fall into two categories: syntax and semantic errors. Syntax errors, such as misplaced parentheses or incorrect object names in the set O, cause the planner to terminate during the parsing stage due to invalid PDDL inputs. Semantic errors occur when the initial state s_0 does not match the actual scene or when the goal description S^* misaligns with the user's intended goal. As a result, the planner is unable to find a valid path from state s_0 to any goal state $s_g \in S^*$, leading to a dead end. In both cases, LLM is reprompted using the planner output and a zero-shot Chain of Thought (CoT) [30] prompt, guiding it to analyze the planner error message. If a syntax error is detected, the LLM corrects the problem by fixing the incorrect syntax. When a semantic error is identified, the LLM adjusts s_0 or S^* or both, and refines the planning problem P. A prompt example of LLM syntax and semantic checker is shown in [Fig. 2].

While Python code errors are less frequent than problem PDDL errors, they typically involve simple runtime issues, such as LLM using incorrect action names or neglecting to define action primitive parameters before using it. In such cases, the LLM is similarly reprompted with the Python exception message and a zero-shot CoT prompt to refine the code.

4. Experiments

4.1 Experimental Setup

The experiments used an Intel Core i9 CPU and NVIDIA RTX Ada 6000 GPU. For the physical robot setup, we used UR5e dual manipulators, each equipped with Robotiq 3F grippers and an Intel RealSense D455 RGBD camera mounted above the table for a top-down view. We utilized GPT-4-Turbo [31] as the multimodal LLM and Fast Downward [23] as the symbolic PDDL planner.

All experiments were based on a PDDL domain inspired by the well-known Blocksworld domain. Additionally, the table was divided into six sections, with the position of each section specified in the prompts.

[Fig. 3] Demonstration of a pick-and-place task. Initially, three RGB-colored blocks are placed in a row next to the basket (leftmost image). The goal is to identify the closest block to the basket and drop it into the basket (image 2). Actions in images 1 through 2 were derived using our task planning pipeline and executed successfully

[Fig. 4] Demonstration of a block stacking. In this case, the blocks are stacked in blue, red, and green from top to bottom (leftmost image). The goal is to restack the blocks in the order of green, red, and blue from top to bottom on the right side of the table (image 8). Actions in images 1 to 8 were executed successfully using our task planning pipeline

4.2 Robot Demonstration

We conducted physical robot demonstrations on three types of tasks: pick-and-place, block-stacking, and block rearrangement. These tasks showcased our framework's ability to perceive detailed attributes of the objects, such as letters on the blocks and abstract positions of blocks, and its capacity for task planning. The LLM initialization for the real robot demonstration was configured as shown in [Fig. 2].

4.2.1 Pick-and-Place

The first experiment involved a simple pick-and-place task. As in [Fig. 3] three colored blocks—red, green, and blue—and a basket were placed on the table. Given the user's goal task, "*Move the closest block into the basket*," the system identified the green block as the block nearest to the basket, generated actions (*pickup green_block t5)(putdown green_block basket*), and generated low-level code with grasp and place poses for the green block.

4.2.2 Block-Stacking

The second experiment focused on stacking blocks in a specified vertical order. As in [Fig. 4], three blocks were initially stacked at the center of the table. The user's goal was to "*Stack the blocks in a new order at the right side of the table. Red at the bottom, green in the middle, and blue at the top.*". The problem PDDL, plan PDDL and low-level code were generated, as shown in [Fig. 5] by the LLM and symbolic planner.

4.2.3 Block Rearrangement

The third experiment was rearranging blocks in alphabetical


```
        Problem PDDL
(define (problem stack)
 (:domain blocksworld)
 (:objects
  blue_block - block
  red_block - block
  green_block - block
  t1 t2 t3 t4 t5 t6 - table
 )
 (:init
  (arm-empty)
  (on-table green_block t5)
  (on red_block green_block)
  (on blue_block red_block)
  (clear blue_block)
  (clear-table t1)
  …
  (clear-table t6)
 )
 (:goal
  (and
   (on-table red_block t6)
   (on green_block red_block)
   (on blue_block green_block))))
```

```
         Plan PDDL
(unstack blue_block red_block)
(putdown blue_block t2)
(unstack red_block green_block)
(putdown red_block t6)
(pickup green_block t5)
(stack green_block red_block)
(pickup blue_block t2)
(stack blue_block green_block)
```

```
       Low-Level Code
t2 = [-0.40, -0.525, 0.79, 0.5, …]
…
bbox = rw.detect_object('blue block')
grasp_pose =
rw.get_grasp_pose('blue block', bbox)
rw.pick_up_object(grasp_pose)
rw.place_object(t2)
…
pose_blue_block =
pose_green_block.copy()
pose_blue_block[2] += 0.08
rw.place_object(pose_blue_block)
```

[Fig. 5] The problem PDDL, plan PDDL, and low-level code generated during the block-stacking demonstration

[Fig. 6] Demonstration of a block rearrangement. Initially, blue, red, and green blocks with the letters Y, Z, and X are arranged from left to right (leftmost image). The goal is to rearrange the blocks alphabetically from left to right: X, Y, and Z (image 8). Actions in images 1 to 8 were executed successfully using our task planning pipeline

[Fig. 7] The problem PDDL, plan PDDL, and low-level code generated during the block rearrangement demonstration

[Table 1] Success and failure rates (%) without replanning

Domain	Problem failure		Python failure	Success rates
	syntax	semantic		
Stack	6.6	16.7	3.3	73.3
Rearrange	3.3	3.3	0	93.3

[Table 2] Success and failure rates (%) with replanning

Domain	Problem failure		Python failure	Success rates
	syntax	semantic		
Stack	0	3.3	0	96.7
Rearrange	0	0	0	100

order based on letters marked on them. Given the user command "*Rearrange the blocks in alphabetic order from left to right, based on the letters on the top.*", the LLM recognized the letters on each block and their initial positions. Using its alphabet knowledge, the LLM generated a problem description to rearrange the blocks in the order X, Y, and Z from left to right, as shown in [Fig. 6]. The problem PDDL, plan PDDL and low-level code were generated, as shown in [Fig. 7] by the LLM and symbolic planner.

4.3 Experiment Results

[Table 1] and [Table 2] summarize the success rates and failure causes for two domains, block-stacking, and block rearrangement, comparing cases with and without replanning. For each domain, 30 problems were randomly generated, and we observed whether task planning and low-level code execution succeeded.

Without replanning, the block-stacking domain had a success rate of 73.3%. Among the failures, the most common cause was problem PDDL semantic errors, followed by problem PDDL syntax errors and failures caused by Python exceptions. In the block rearrangement domain, the success rate was higher at 93.3%, with problem PDDL semantic and syntax errors each accounting for 3.3% of all cases and no Python failures observed. The lower success rate in block-stacking is likely due to difficulties the multimodal LLM had in accurately capturing the spatial relationships between stacked blocks for the PDDL initial state.

In both domains, problem PDDL semantic errors were primarily caused by confusion between the *on-table* predicate (indicating a block on the table) and the *on* predicate (representing the relationship between blocks). Problem PDDL syntax errors were typically due to incorrect domain names or improper PDDL formatting. Python failures arose when the pose variable was used in the action primitive without being first retrieved by the '*get_grasp_pose*' function.

Limiting the number of replanning attempts to four, we observed a significant increase in success rates, approaching almost 100% in both domains. Notably, the rate of problem

PDDL syntax errors and Python failures dropped to zero with replanning, demonstrating the effectiveness of our LLM-based replanning method in ensuring plan correctness.

5. Conclusion and Future Work

In this paper, we proposed a neuro-symbolic task replanning pipeline that integrates multimodal LLMs and symbolic planners to address challenges in robot task planning. By leveraging LLMs' commonsense reasoning abilities, our system generates problem specifications, uses a symbolic planner to find a plan, and converts it into low-level code with action parameter selection. We also introduced an automatic replanning module to resolve failures during planning and demonstrated the system's effectiveness in real-world tasks, showing improved success rates with replanning.

While our approach assumes the downward refinement property for simple scenarios, it does not fully account for real-world complexities. Future work will focus on developing a complete TAMP algorithm that handles cases where downward refinement does not hold, addressing these real-world constraints.

References

[1] S. Kambhampati, K. Valmeekam, L. Guan, K. Stechly, M. Verma, S. Bhambri, L. Saldyt, and A. Murthy, "LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks," *arXiv:2402.01817*, 2024, DOI: 10.48550/arXiv.2402.01817.

[2] K. Valmeekam, S. Sreedharan, M. Marquez, A. Olmo, and S. Kambhampati, "On the planning abilities of large language models (a critical investigation with a proposed benchmark)," *arXiv:2302.06706*, 2023, DOI: 10.48550/arXiv.2302.06706.

[3] T. Silver, V. Hariprasad, R. S. Shuttleworth, N. Kumar, T. Lozano-Pérez, and L. P. Kaelbling, "PDDL planning with pretrained large language models," *NeurIPS 2022 foundation models for decision making workshop*, New Orleans, USA, 2022, [Online], https://openreview.net/forum?id=1QMMUB4zfl.

[4] M. Fox and D. Long, "PDDL2. 1: An extension to PDDL for expressing temporal planning domains," *Journal of Artificial Intelligence Research*, vol. 20, Dec., 2003, DOI: 10.1613/jair.1129.

[5] K. Shirai, C. C. Beltran-Hernandez, M. Hamaya, A. Hashimoto, S. Tanaka, K. Kawaharazuka, K. Tanaka, Y. Ushiku, and S. Mori, "Vision-language interpreter for robot task planning," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, pp. 2051-2058, 2024, DOI: 10.1109/ICRA57147.2024.10611112.

[6] B. Liu, Y. Jiang, X. Zhang, Q. Liu, S. Zhang, J. Biswas, and P. Stone, "Llm+ p: Empowering large language models with optimal planning proficiency," *arXiv:2304.11477*, 2023, DOI: 10.48550/arXiv.2304.11477.

[7] Y. Chen, J. Arkin, C. Dawson, Y. Zhang, N. Roy, and C. Fan, "Autotamp: Autoregressive task and motion planning with llms as translators and checkers," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, pp. 6695-6702, 2024, DOI: 10.1109/ICRA57147.2024.10611163.

[8] M. S. Sakib and Y. Sun, "Consolidating Trees of Robotic Plans Generated Using Large Language Models to Improve Reliability," *arXiv:2401.07868*, 2024, DOI: 10.48550/arXiv.2401.07868.

[9] R. Arora, S. Singh, K. Swaminathan, A. Datta, S. Banerjee, B. Bhowmick, K. M. Jatavallabhula, M. Sridharan, and M. Krishna, "Anticipate & Act: Integrating LLMs and Classical Planning for Efficient Task Execution in Household Environments," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, pp. 14038-14045, 2024, DOI: 10.1109/ICRA 57147.2024.10611164.

[10] Z. Zhou, J. Song, K. Yao, Z. Shu, and L. Ma, "Isr-llm: Iterative self-refined large language model for long-horizon sequential task planning," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, pp. 2081-2088, 2024, DOI: 10.1109/ICRA57147.2024.10610065.

[11] T. Silver, S. Dan, K. Srinivas, J. B. Tenenbaum, L. Kaelbling, and M. Katz, "Generalized planning in pddl domains with pretrained large language models," *The AAAI Conference on Artificial Intelligence*, Vancouver, Canada, pp. 20256-20264, 2024, DOI: 10.1609/aaai.v38i18.30006.

[12] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 265-293, May, 2021, DOI: 10.1146/annurev-control-091420-084139.

[13] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. M. J. Ruano, K. Jeffrey, S. Jesmonth, N. J. Joshi, R. C. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. M. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, and M. Yan, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv:2204.01691*, 2022, DOI: 10.48550/arXiv.2204.01691.

[14] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, "Progprompt: Generating situated robot task plans using large language models," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, London, United Kingdom, pp. 11523-11530, 2023, DOI: 10.1109/ICRA48891.2023.10161317.

[15] Y. Ding, X. Zhang, C. Paxton, and S. Zhang, "Task and motion planning with large language models for object rearrangement," *2023 IEEE/RSJ International Conference on Intelligent Robots*

and Systems (IROS), Detroit, USA, pp. 2086-2092, 2023, DOI: 10.1109/IROS55552.2023.10342169.

[16] S. Wang, M. Han, Z. Jiao, Z. Zhang, Y. N. Wu, S.-C. Zhu, and H. Liu, "LLM3: Large Language Model-based Task and Motion Planning with Motion Failure Reasoning," *arXiv:2403.11552*, 2024, DOI: 10.48550/arXiv.2403.11552.

[17] M. Kwon, Y. Kim, and Y. J. Kim, "Fast and Accurate Task Planning using Neuro-Symbolic Language Models and Multi-level Goal Decomposition," *arXiv:2409.19250*, 2024, DOI: 10.48550/arXiv.2409.19250.

[18] Y. Xie, C. Yu, T. Zhu, J. Bai, Z. Gong, and H. Soh, "Translating natural language to planning goals with large-language models," *arXiv:2302.05128*, 2023, DOI: 10.48550/arXiv.2302.05128.

[19] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, and A. Askell, "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, pp. 1877-1901, 2020, [Online], https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

[20] M. Skreta, N. Yoshikawa, S. Arellano-Rubach, Z. Ji, L. B. Kristensen, K. Darvish, A. Aspuru-Guzik, F. Shkurti, and A. Garg, "Errors are useful prompts: Instruction guided task programming with verifier-assisted iterative prompting," *arXiv:2303.14100*, 2023, DOI: 10.48550/arXiv.2303.14100.

[21] S. S. Raman, V. Cohen, I. Idrees, E. Rosen, R. Mooney, S. Tellex, and D. Paulius, "CAPE: Corrective Actions from Precondition Errors using Large Language Models," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, pp. 14070-14077, 2024, DOI: 10.1109/ICRA57147.2024.10611376.

[22] R. Howey, D. Long, and M. Fox, "VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL," *16th IEEE International Conference on Tools with Artificial Intelligence*, Boca Raton, USA, pp. 294-301, 2004, DOI: 10.1109/ICTAI.2004.120.

[23] M. Helmert, "The fast downward planning system," *Journal of Artificial Intelligence Research*, vol. 26, pp. 191-246, Jul., 2006, DOI: 10.1613/jair.1705.

[24] S. H. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor, "Chatgpt for robotics: Design principles and model abilities," *IEEE Access*, vol. 12, pp. 55682-55696, Apr., 2024, DOI: 10.1109/ACCESS.2024.3387941.

[25] S. Chitta, I. Sucan, and S. Cousins, "Moveit![ros topics]," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 18-19, Mar., 2012, DOI: 10.1109/MRA.2011.2181749.

[26] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang, "Grounded sam: Assembling open-world models for diverse visual tasks," *arXiv:2401.14159*, 2024, DOI: 10.48550/arXiv.2401.14159.

[27] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, and J. Zhu, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," *arXiv:2303.05499*, 2023, DOI: 10.48550/arXiv.2303.05499.

[28] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, and P. Dollár, "Segment anything," *2023 IEEE/CVF International Conference on Computer Vision*, Paris, France, pp. 4015-4026, 2023, DOI: 10.1109/ICCV51070.2023.00371.

[29] A. ten Pas and R. Platt, "Using geometry to detect grasp poses in 3d point clouds," *Robotics Research*, vol. 2, pp. 307-324, Jul., 2017, DOI: 10.1007/978-3-319-51532-8_19.

[30] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," *arXiv:2205.11916*, 2022, DOI: 10.48550/arXiv.2205.11916.

[31] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, et al., "Gpt-4 technical report," *arXiv:2303.08774*, 2023, DOI: 10.48550/arXiv.2303.08774.

권 민 서

2024 이화여자대학교 컴퓨터공학과 학사

2024~현재 이화여자대학교 컴퓨터공학과 석사 과정

관심분야: Robot Task Planning

김 영 준

1993 서울대학교 계산통계학과 학사

1996 서울대학교 계산통계학과 석사

2000 Purdue University, Computer Science 박사

2003 University of North Carolina at Chapel Hill 박사후 연구원

2003~현재 이화여자대학교 컴퓨터공학과 교수

관심분야: Robot Motion Planning, Haptic Rendering, Physically-based Animation, Geometric Modeling