

DAG-Plan: Generating Directed Acyclic Dependency Graphs for Dual-Arm Cooperative Planning

Zeyu Gao^{1*}, Yao Mu^{2*}, Jinye Qu¹, Mengkang Hu⁴, Shijia Peng⁵, Chengkai Hou³, Lingyue Guo¹,
Ping Luo^{4,5}, Shanghang Zhang^{3†} and Yanfeng Lu^{1†}

Abstract—Dual-arm robots promise greater efficiency but require planning for complex tasks with nonlinear sub-task dependencies. Current methods using Large Language Models (LLMs) suffer from a fundamental trade-off: generating linear sequences is efficient but fails to model parallelism and adapt to changes, while iterative querying is adaptive but too slow and costly. To bridge this gap, we introduce DAG-Plan, a novel task planning framework that for the first time employs a Directed Acyclic Graph (DAG) as the central representation for dual-arm coordination. The key insight is that a DAG natively captures complex sub-task dependencies and explicitly reveals opportunities for parallel execution. Within this framework, an LLM is used only once as a powerful semantic parser to translate a natural language instruction into a structured DAG. During execution, our system dynamically assigns candidate nodes to the suitable arm based on real-time environmental observations, enabling truly adaptive and parallel operation. Extensive evaluation on a dual-arm kitchen benchmark shows that DAG-Plan’s structured approach fundamentally outperforms existing paradigms. It achieves a 48% higher success rate than single-query linear sequence methods with dual arm by robustly managing dependencies, and an 84.1% higher execution efficiency than iterative querying methods by eliminating the latency of repeated LLM calls. Our work demonstrates that a principled, graph-based representation is the key to unlocking efficient and reliable LLM-based planning for complex robotic systems. More demos and code are available on <https://sites.google.com/view/dag-plan>.

I. INTRODUCTION

Achieving effective bimanual coordination in robotics is challenging due to the complexities and long-horizon of dual-arm operations, requiring precise spatial and temporal coordination [1], [2]. While humans effortlessly coordinate

Zeyu Gao* and Yao Mu* are co-first authors. Yanfeng Lu[†] and Shanghang Zhang[†] are the corresponding authors. This work is supported by the Strategic Priority Research Program of the Chinese Academy of Sciences (XDA0450200, XDA0450202), and Central Government Guidance Funds for Local Science and Technology Development (YDZX2025124), the National Natural Science Foundation of China (62476011), and the Beijing Natural Science Foundation (L252060).

¹ Zeyu Gao, Jinye Qu, Lingyue Guo and Yanfeng Lu are with the State Key Laboratory of Multimodal Artificial Intelligence Systems, Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing 100190, China, and also with the School of Artificial Intelligence, University of Chinese Academy of Sciences (UCAS), Beijing 100049, China (e-mail: gaوزهy2023@ia.ac.cn; yanfeng.lv@ia.ac.cn).

² Yao Mu is with the School of Computer Science, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: muyao@sjtu.edu.cn)

³ Chengkai Hou and Shanghang Zhang are with State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University, Beijing 100871, China (e-mail: shanghang@pku.edu.cn).

⁴ Mengkang Hu and Ping Luo are with the Department of Computer Science, The University of Hong Kong, Hong Kong 999077, China

⁵ Shijia Peng and Ping Luo are with OpenGVLab, Shanghai AI Laboratory, Shanghai 200232, China

their hands in daily long-horizon tasks, replicating such coordination in robots presents significant challenges. Traditional method employed hand-designed primitives to manage the movements of dual robotic arms [3], [4]. Whereas these methods often fall short in long-horizon task as they lack the flexibility required for adaptive task execution. Large language models (LLMs) have emerged as powerful tools endowed with extensive knowledge and sophisticated reasoning abilities. By systematically breaking down tasks into actionable sub-tasks and leveraging their commonsense knowledge and implicit reasoning capabilities, LLMs empower robots to effectively adapt to long-horizon tasks in the wild [5], [6].

The prevalence of linear task sequences, however, introduces a fundamental bottleneck for dual-arm coordination. While LLM-based planners have demonstrated remarkable reasoning capabilities, we observe that their performance in bimanual scenarios is constrained by the inherent limitations of the linear representation paradigm itself. This is true even for advanced multi-agent frameworks [7]–[9] that decompose tasks for multiple robots: their output remains a set of linear steps, which fails to explicitly model the complex, non-linear dependencies and parallel opportunities critical for efficient dual-arm execution. Consequently, planners are forced to choose between the rigidity of a single, pre-generated linear sequence and the prohibitive inefficiency of iterative LLM querying to navigate execution.

We identify that these challenges are not independent issues but rather symptomatic manifestations of a single root cause: the reliance on linear task representations. This representation is fundamentally mismatched to the problem for three key reasons. First, its sequential nature is incapable of explicitly representing the parallel execution opportunities that are essential for efficiency. Second, a pre-generated linear sequence is inherently rigid, decoupled from the real-time state of the environment and unable to adapt the execution order or arm assignment dynamically. Third, while iterative linear planning attempts to recover adaptability by frequently querying the LLM, it does so at the prohibitive cost of latency and computational expense, merely treating the symptom rather than curing the disease. Thus, the field is currently trapped in a false choice between inflexibility and inefficiency.

To break this deadlock, we propose DAG-Plan, a novel framework that replaces the linear representation with a structured Directed Acyclic Graph (DAG). DAG explicitly captures sub-tasks as nodes and, crucially, their temporal dependencies as directed edges, thereby natively encoding both the execution constraints and the parallel opportunities

that are opaque to a linear sequence. Within this framework, the role of the LLM is refined to a single, initial query where it acts as a powerful semantic parser, translating the natural language instruction into this structured DAG. The subsequent planning and execution intelligence resides in our Task Planning Inference module. This module continuously monitors the environmental state and the DAG’s structure—dynamically assigning the feasible and efficient pair of executable sub-tasks to the dual arms. This enables true parallel and adaptive execution without the latency of repeated LLM queries. Our main contributions are summarized as follows:

- We identify that the bottleneck in LLM-based dual-arm planning is the reliance on linear task representations, which are ill-suited for expressing the complex, parallelizable nature of bimanual tasks and necessitate a trade-off between efficiency and adaptability.
- We propose DAG-Plan, a novel framework that introduces a Directed Acyclic Graph (DAG) as the structured task representation for dual-arm planning. This representation explicitly models sub-task dependencies and parallelism, enabling efficient and adaptive execution with a single LLM query.
- We demonstrate that our approach outperforms existing paradigms. It achieves a 48% higher success rate than single-query linear sequence methods with dual arm by robustly managing dependencies, and an 84.1% higher execution efficiency than iterative querying methods by eliminating the latency of repeated LLM calls.

II. RELATED WORKS

A. Task Planning with LLMs

Large Language Models have revolutionized task planning by translating natural language goals into executable action sequences for embodied agents [10], [11]. Their success stems from an unparalleled ability to leverage commonsense knowledge and in-context learning. However, a prevalent limitation across this field is the reliance on linear action sequences as the primary output representation. This limitation becomes acutely apparent in multi-entity systems, such as multi-robot teams [7], [8] or dual-arm robots [9], [12], [13]. While these works demonstrate advanced decomposition or communication strategies (e.g., independent task decomposition in Twostep [7] or inter-robot discussion in RoCo [8]), their final output remains a linear sequence of steps for each agent. This representation obscures the temporal dependencies and parallel opportunities between sub-tasks, forcing a trade-off between plan quality and execution efficiency.

B. Dual-arm Robot Manipulation

Research on dual-arm manipulation has made significant strides, primarily at the skill level, across industrial [1], [2], agricultural [14], [15], and domestic settings [16], [17]. This progress is driven by advances in motion planning [17], foundation models [18], and learning-based methods [19], [20], enabling robots to perform human-like bimanual skills such as lifting or assembling. However, autonomously orchestrating a long-horizon task composed of multiple such skills

remains an open challenge. While LLMs have been applied to dual-arm systems, they often only control one arm at a time [12] or rely on pre-defined bimanual skills [13], thus failing to fully exploit the system’s potential for parallel and adaptive operation at the task level. Our work addresses this gap by proposing a framework for task-level planning that dynamically composes and coordinates skills for both arms.

C. Structured Task Decomposition (STD)

STD aims to break down complex tasks into structured representations like graphs, explicitly modeling dependencies and parallelism. Traditional STD methods were heavily constrained by data availability, relying on crowd-sourcing [21], [22] or complex query systems [23], [24]. The advent of LLMs, with their vast repositories of commonsense knowledge, has emerged as a powerful new paradigm for STD. For instance, TaskLAMA [25] demonstrated that LLMs can effectively decompose tasks into graphs with temporal relations. This line of work provides a critical foundation, confirming the feasibility of using LLMs for structured output. DAG-Plan pushes this frontier further. We explore how to leverage LLM-generated DAG not merely as a structured output, but as the central representation for real-time, adaptive, and parallel execution on a dual-arm robotic system.

D. Comparison with Symbolic Planners

Our approach is also related to symbolic planners like PDDL [26] and Scene Graph [27], which offer formal guarantees but require hand-crafted world models, creating a significant scalability bottleneck. PDDL and Scene Graph offers high formalism and precision, but low generalizability, and require significant human effort to define the domain. Our DAG-Plan method offers high generalizability and requires minimal human effort only a skill and task nature language description but provides less formal guarantee and relies on the LLM’s implicit knowledge. This makes DAG-Plan suited for open-world environments where formal models are impractical, while still producing a more explicit and parallelizable plan than a linear sequence.

III. METHODS

The DAG-Plan framework is designed to overcome the fundamental limitations of linear representations by introducing a structured DAG as the core planning representation. This representation natively captures complex temporal dependencies and explicit parallelism, enabling efficient and adaptive dual-arm coordination with a single LLM query. An overview of DAG-Plan pipeline is illustrated in Fig. 1.

A. Sub-task Dependency Directed Acyclic Graph Generation

We propose a novel approach that uses an LLM as a semantic parser to decompose a complex natural language instruction into a structured Directed Acyclic Graph (DAG). This graph representation is fundamental to our framework, as it natively captures the temporal dependencies and parallel opportunities inherent in bimanual tasks, which are opaque to linear sequences.

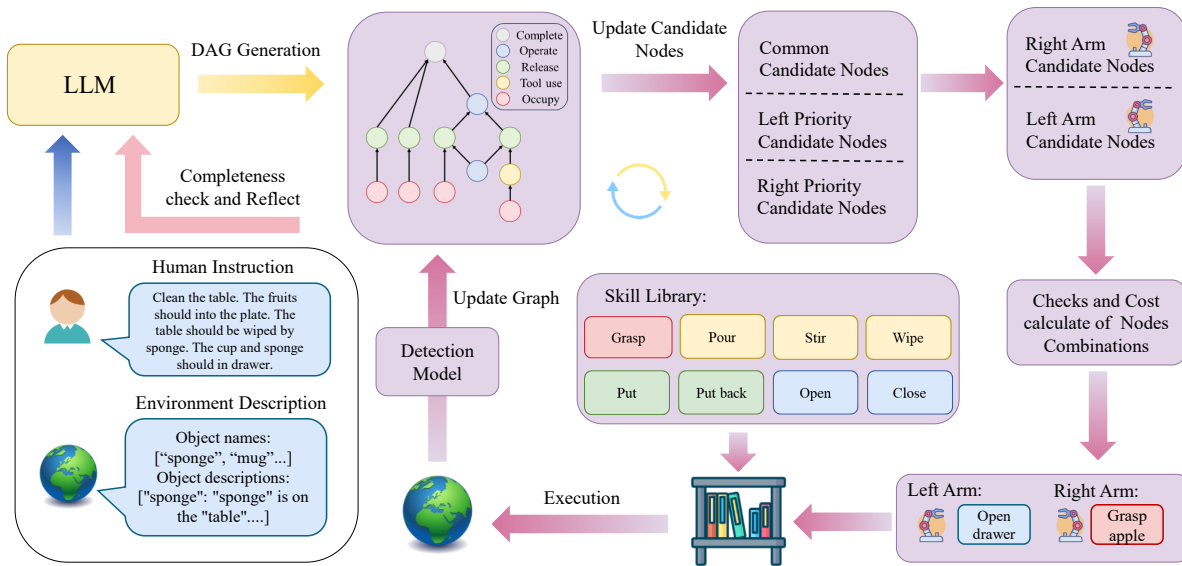


Fig. 1. An overview of DAG-Plan. The DAG-Plan generates a DAG based on human instruction and environmental description. It checks the graph's completeness and reflects the LLM to regenerate if incomplete. Once a valid DAG is obtained, DAG-Plan performs task inference to identify executable candidate nodes. The occupied arm and free arm are assigned priority candidate nodes and common candidate nodes respectively. The framework then evaluates all candidate combinations for feasibility and cost. DAG-Plan selects the nodes with the lowest cost and employs skill in library for execution. DAG-Plan updates the graph, iterating inference until the DAG is fully executed.

We formally define the dual-arm task graph as a tuple:

$$G = (V, E, T, N), \quad (1)$$

where:

- $V = \{v_1, v_2, \dots, v_n\}$ is the set of sub-task nodes.
- $E \subseteq V \times V$ is the set of directed edges representing temporal dependencies. An edge $e_{ij} = (v_i, v_j) \in E$ indicates that sub-task v_i must be completed before v_j can begin.
- $T : V \rightarrow \mathcal{T}$ maps each node to a task type from the set \mathcal{T} .
- $N : V \rightarrow \{1, 2\}$ specifies the number of arms required for the execution of each node.

The task type set \mathcal{T} is defined as $\{Occupy, Tool\ use, Release, Operate, Complete\}$. This typology is critical for modeling arm state transitions and ensuring logical soundness:

- *Occupy*: Engages a gripper, leaving the arm *occupied* (e.g., grasp).
- *Tool use*: Requires a tool, using it and keeping it in the gripper (e.g., wipe).
- *Release*: Releases an object, freeing the arm (e.g., place).
- *Operate*: A general task that occupies the arm only during execution (e.g., open).
- *Complete*: The terminal node means the task is finished when DAG only remain this node.

The task type $T(v)$ dictates the state transition of the robotic arms. We define an *occupy-release* pair as a critical subgraph $G' \subseteq G$ that manages object manipulation:

$$G' = (V', E'), \text{ s.t. } V' = \{v_o, v_{t1}, \dots, v_{tk}, v_r\}$$

$$T(v_o) = Occupy, \quad T(v_r) = Release, \quad T(v_t) = Tool\ use, \quad (2)$$

where a path $v_o \rightsquigarrow v_r$ exists via E' . This structure enforces the fundamental constraint that for any release operation,

there must exist a corresponding, preceding occupy operation, ensuring an object is grasped before it is released.

The LLM is prompted to generate a DAG. The output is validated for structural integrity. Graphs with incomplete *occupy-release* pairs or disconnected components are rejected and regenerated, ensuring a complete plan for execution.

B. Task Planning Inference with Directed Acyclic Graph

The generated DAG serves as a static plan. The Task Planning Inference module is the dynamic execution engine that traverses this graph in real-time, enabling adaptive and parallel operation. Its core function is to continuously select the optimal pair of executable sub-tasks for the two arms based on the current graph state and environmental observations.

1) *Stateful Node Selection*: The executor maintains the state of the task graph. Each node $v_i \in V$ has an execution status $s(v_i) \in \{Pending, Ready, Done\}$. A node becomes *Ready* when all its prerequisite nodes are *Done*.

The set of executable nodes is dynamically prioritized based on the semantic context of the graph:

- *Common Candidates* (C_c): The set of all *Ready* nodes.
- *Priority Candidates* (C_p): For any *Ready* node v_r where $T(v_r) = Release$, its entire antecedent path within the same *occupy-release* pair is given priority. This ensures that arms committed to an ongoing manipulation sequence complete it before starting new ones, which is critical for preventing logical errors and deadlocks. Each left and right arm has a Priority Candidates

2) *Execution Loop*: The executor runs a continuous loop: First, update candidate nodes: Recompute C_c and identify C_p based on the new graph state. Second, checks and cost calculate of nodes combinations: For each arm, if a priority

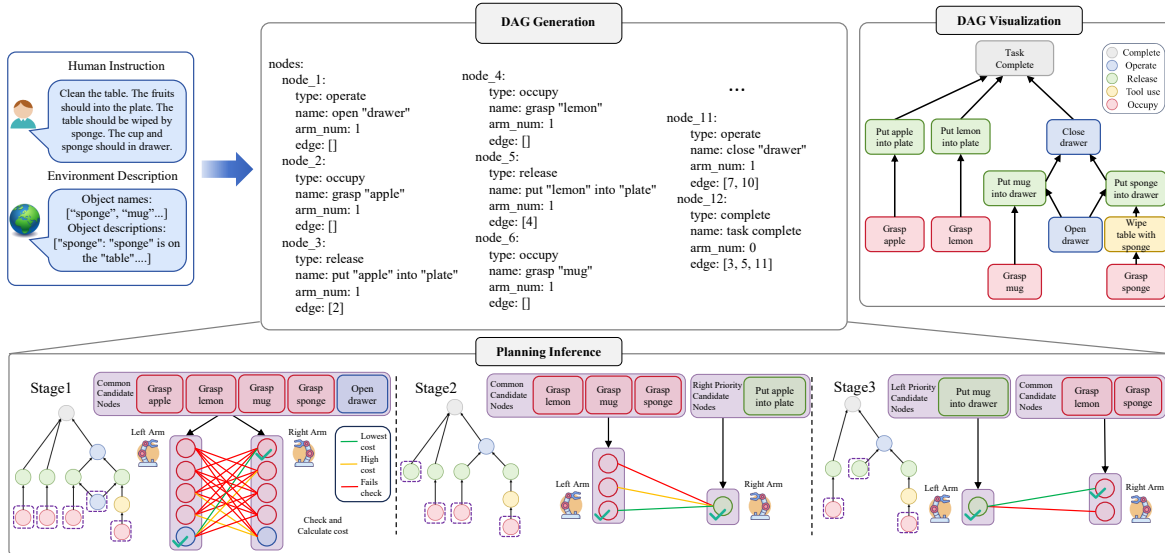


Fig. 2. The process of Task Planning Inference. In task 2 “clean the table (Hard)”, DAG-Plan initializes common candidate nodes based on the DAG. It evaluates node combinations, checks feasibility, and calculates costs. In stage 1, the right arm is selected to grasp apple and the left to open drawer. After execution, the task graph and nodes are updated, adding subsequent *release* nodes to the priority candidate nodes for right arm. In stage 2, right arm is assigned corresponding priority candidate nodes, checked and left arm still selected node in common candidate nodes with empty priority candidate nodes. The task graph and nodes are updated again. In stage 3, left arm put mug into drawer and right arm grasp lemon.

candidate exists, nodes are selected from them. Otherwise, nodes are selected from the common candidates. Third, update graph: Upon a node’s completion, it is marked *Done* and removed from G . This updates the *Ready* status of its successors.

3) *Feasibility Checking and Cost Calculation*: The system evaluates all possible pairs of selected nodes. Each candidate assignment is validated through a series of geometric feasibility checks:

- *Dependency Check*: A logical check rejects a candidate if it is an *Occupy* node whose subsequent *Release* node is not yet *Ready*, preventing an arm from being indefinitely occupied.
- *Reachable Check*: Calculates the Euclidean distance between target objects for both arms. A pair is rejected if distance between left target and right target is greater than $d_{\text{Reachable}}$, ensuring both targets are within a feasible simultaneous workspace.
- *Across Avoidance*: The distance between the two target object must greater than a safety threshold d_{Across} , indicating a high risk of inter-arm collision. This prevents scenarios where the arms would need to cross each other or operate in dangerously close proximity.

Object position for these checks are acquired using Grounding DINO [28] and SAM2 [29] for segmentation, and SAR-Net [30] for 6D pose estimation of specific categories. After completing the checks, we calculate the distance to left and right hand of the candidate nodes based on the environment state as the cost. We aim for the target objects to be close to the robotic hands facilitating dual-arm operations. The pair with the lowest combined cost is selected for execution. Finally, we briefly introduce the task planning inference process in Alg. 1.

Algorithm 1 Task Planning Inference with DAG

```

1: Input: Task Graph  $G = (V, E, T, N)$ 
2: Initialize:  $s(v_i) \leftarrow \text{Pending}, \forall v_i \in V$ 
3: while  $\exists v_i \in V \mid s(v_i) \neq \text{Done} \wedge T(v_i) \neq \text{Complete}$  do
  // Update candidate nodes
4:  $C_c \leftarrow \{v_i \in V \mid s(v_i) = \text{Ready}\}$ 
5:  $C_{pl} \leftarrow \{v \in \text{left occupy-release paths} \mid v \in C_c\}$ 
6:  $C_{pr} \leftarrow \{v \in \text{right occupy-release paths} \mid v \in C_c\}$ 
7:  $C_l \leftarrow (C_{pl} \neq \emptyset) ? C_{pl} : C_c$ 
8:  $C_r \leftarrow (C_{pr} \neq \emptyset) ? C_{pr} : C_c$ 
  // Checks and cost calculate of nodes combinations
9: Filter feasible pairs  $(v_l, v_r) \in C_l \times C_r$  through
  dependency, reachability, and across-arm checks
10: Calculate cost for each feasible pair based on distance
  to hands  $(v_l^*, v_r^*) \leftarrow \arg \min \text{cost}(v_l, v_r)$ 
11: Execute  $(v_l^*, v_r^*)$ 
  // Update graph
12:  $s(v_l^*) \leftarrow \text{Done}, s(v_r^*) \leftarrow \text{Done}$ 
13: for  $v_i \in V \mid s(v_i) = \text{Pending}$  do
14:   if  $\forall (v_j, v_i) \in E, s(v_j) = \text{Done}$  then
15:      $s(v_i) \leftarrow \text{Ready}$ 
16:   end if
17: end for
18: end while

```

C. Sub-tasks Execution with Foundation Model

Once the sub-tasks to be executed are determined, the robot needs to perform the corresponding actions to bridge the gap between textual instructions and the physical environment. The collective perceptual inferences specific to objects, along with physical insights and parameters for manipulation, are methodically organized and converted into a structured format



Clean the table (Easy) Clean the table (Hard) Stack bowls Make cup of coffee Boil vegetables

Fig. 3. Snapshots of 5 Tasks of Dual-arm Kitchen Benchmark.

TABLE I
TASK LIST OF DUAL-ARM KITCHEN BENCHMARK.

Index	Task Name	Human Instruction
Task 1	Clean the table (Easy)	Clean the table. Put objects into plate.
Task 2	Clean the table (Hard)	Clean the table. The fruits should into the plate. The table should be wiped by sponge. The mug and sponge should in drawer.
Task 3	Stack bowls	Stack the bowls onto the wooden tray with the green bowl, blue bowl and yellow bowl order.
Task 4	Make cup of coffee	Make a cup of coffee. You should add the coffee, water and milk in order. Finally, stir it with the spoon.
Task 5	Boil vegetables	Boil vegetables. Pour water and put vegetables into the pot.

of executable action code.

1) *Occupy*: We use the pre-trained AnyGrasp [18] to produce a variety of grasp pose proposals.

2) *Tool use*: We have specifically designed tool-usage skills based on the categories of tools employed, and have provided adjustable parameters to enable the robot to perform flexible tool-manipulating actions.

3) *Release*: Based on the current grasping position of the object and the 3D bounding box of the target object, we propose an appropriate release pose to ensure that the target object can be accurately placed at the desired location.

4) *Operate*: We employ the GPartNet [31] to forecast the physical characteristics of articulated objects. This model works by dividing the point cloud of an articulated object into its constituent rigid parts and subsequently calculating the articulation parameters.

IV. EXPERIMENTS

In the experiments, we are interested in answering the following questions about DAG-Plan:

- **Q1**: To what extent can the task-graph approach improve performance compared to the task-sequence method?
- **Q2**: How does the cost and efficiency of DAG-Plan compare to iterative methods?
- **Q3**: How does the performance of DAG-Plan compare to multi-agents methods?
- **Q4**: Can DAG-Plan be deployed in real-world setting?

A. Experimental Setup

To validate the correctness and execution efficiency of our method, we created a Dual-arm Kitchen Benchmark, including plan tests and physical simulation tests. This benchmark fills the gap in long-sequence operation physics simulation benchmarks for dual-arm robots. The goal of this benchmark is to validate the success rate and efficiency of dual-arm robot planning and executable ability in complex scenarios. The dual-arm robot should correctly plan and fully utilize both the left and right arms, completing tasks with as few execution stages as possible. The benchmark consists of 5 sequential

tasks are shown in Tab. I and Fig. 3, comprising a total of 44 sub-tasks. Our physical simulation scene is built on the Sapien [32]. The embodied platform is Agilex CobotMagic, each 6 DOF arm equipped with a two-finger gripper. The platform is equipped with an Intel RealSense L515 RGB-D camera, attached on the robot’s head.

1) *Evaluation of Planning Effectiveness and Conciseness*: In this experiment, we focus on testing the conciseness of the generated plans and the number of stages required. In a stage, the robot can execute a right arm node and a left arm node. This requires that the plans generated by the LLMs can achieve the task goals in terms of language logic and do not violate the preconditions for stage execution. We used GPT-4o to generate 10 plans for each task, evaluating their **Success Rate (SR)** and the minimum **Stage** of the passed plan required at language level. Due to cost considerations, we also include average **Tokens** usage as an important metric for evaluating cost of the algorithm. Finally, we calculated the average success rate and the average number of stages for all tasks. We defined **Stage Efficiency** as the ratio of single-arm plan stages to the stages required by each method. For failed dual-arm plans, we calculated the stage count based on the single-arm plan stages to ensure a fair comparison.

2) *Evaluation with Physical Simulation*: In this experiment, we will test the executability and execution efficiency of the plans in physical simulation scenarios. Compared to plan tests, physical simulation tests validate both the high-level planning and low-level execution capabilities. Then evaluate the **Success Rate (SR)** and minimum **Time** which consists of **query time** and **execution time** in the simulator with 10 trials. Finally, we calculated the average success rate and the time for all tasks. We defined **Execution Efficiency** as the ratio of single-arm plan time to the time required by each method. For failed dual-arm plans, we calculated the average time based on the single-arm time to ensure a fair comparison.

3) *Baseline Algorithms and Proposed Method*: We compare the planning algorithms baselines and our proposed method:

"Clean the table. The fruits should into the plate. The table should be wiped by sponge. The mug and sponge should in drawer."

T




Fig. 4. Simulation snapshots of the execution process of long-horizon task 2.

TABLE II

PERFORMANCE COMPARISON ON PLAN TESTS. WE REPORT THE SUCCESS RATE, MINIMUM STAGE AND AVERAGE TOKEN USAGE OF THE 10 PLANS GENERATED BY THE LLM FOR EACH TASK.

	Task1			Task2			Task3			Task4			Task5			Average		
	SR↑	Stage↓	Tokens↓	SR↑	Stage↓	Tokens↓	SR↑	Stage↓	Tokens↓	SR↑	Stage↓	Tokens↓	SR↑	Stage↓	Tokens↓	SR↑	Stage Efficiency↑	Tokens↓
TP-S	100%	8	2030.0	70%	11	2424.8	100%	6	2086.8	90%	12	2413.0	100%	7	2096.6	92%	100.0%	2210.2
TP-D	100%	4	2084.1	0%	Fail	2355.6	0%	Fail	2038.6	20%	8	2405.2	80%	5	2040.6	40%	129.4%	2180.8
TwoStep [7]	100%	4	3699.6	0%	Fail	4248.2	0%	Fail	3823.4	0%	Fail	4130.2	100%	5	3810.2	40%	115.8%	3942.3
RoCo [8]	100%	4	26615.0	0%	Fail	68921.4	100%	4	25410.2	40%	8	65242.0	100%	4	26985.0	68%	141.9%	42634.7
DAG-Plan	100%	4	2031.0	70%	7	2207.2	100%	4	1980.6	70%	7	2277.0	100%	4	1973.0	88%	169.2%	2093.8

TABLE III

PERFORMANCE COMPARISON ON PHYSICAL SIMULATION TESTS. WE REPORT THE SUCCESS RATE AND MINIMUM TIME FOR EACH TASK WITH 10 TRIALS. TIME CONSISTS OF QUERY TIME AND EXECUTION TIME.

	Task1		Task2		Task3		Task4		Task5		Average	
	SR↑	Time (s)↓	SR↑	Time (s)↓	SR↑	Time (s)↓	SR↑	Time (s)↓	SR↑	Time (s)↓	SR↑	Execution Efficiency↑
TP-S	80%	9.5 + 53.5 = 63.0	40%	12.9 + 90.7 = 103.6	90%	8.3 + 35.6 = 43.9	80%	8.2 + 116.9 = 125.1	100%	8.2 + 55.8 = 64.0	78%	100.0%
TP-D	80%	8.5 + 26.1 = 34.6	0%	Fail	0%	Fail	0%	Fail	50%	8.3 + 43.2 = 51.4	26%	111.4%
TwoStep [7]	80%	11.4 + 27.3 = 38.7	0%	Fail	0%	Fail	0%	Fail	50%	11.3 + 43.2 = 54.5	26%	109.2%
RoCo [8]	80%	49.3 + 26.2 = 75.5	0%	Fail	90%	54.9 + 24.1 = 79.0	30%	155.3 + 78.9 = 234.2	100%	54.4 + 35.3 = 89.7	60%	68.7%
DAG-Plan	80%	8.5 + 27.6 = 36.1	50%	11.7 + 59.8 = 71.5	90%	8.3 + 24.0 = 32.3	50%	8.3 + 69.9 = 77.9	100%	8.3 + 35.4 = 43.7	74%	152.8%

- 1) *Task Planning for Single-arm (TP-S)*: TP-S directly uses LLMs to generate a full task sequence, with each stage involving a single arm to manipulate a single object.
- 2) *Task Planning for Dual-arm (TP-D)*: TP-D directly uses LLMs to generate a full task sequence, with each stage involving dual arms.
- 3) *Twostep [7]*: Twostep decomposes complex tasks into fully independent sub-tasks with dual agents, with the primary agent and the secondary agent each completing separate sub-tasks to accelerate task execution.
- 4) *RoCo [8]*: RoCo proposed multi-agents collaboration using LLMs for planning which is an iterative method. Robots query LLMs in each stages to collaboratively reason about task strategies. DABICO [9] is a dual-armed version of the RoCo, and is equivalent to the RoCo in our experiments settings.
- 5) *DAG-Plan (Proposed)*: Our method DAG-Plan generates a task graph, followed by task planning inference to iteratively generate nodes for each stage.

B. Evaluation of Planning Effectiveness and Conciseness

This tier evaluates the inherent planning capability of each method, isolating the quality of the high-level plan from low-level execution errors. Our analysis of the plan tests shown in Tab. II reveals the fundamental advantages of the structured DAG representation. It directly addresses Q1, Q2 and Q3.

1) *Q1 DAG Representation vs. Linear Sequences*: The DAG-Plan outperforms linear sequence methods. While TP-S is reliable yet inefficient, and TP-D is efficient yet unreliable, DAG-Plan delivers the best of both worlds. DAG-Plan matches the high success rate of TP-S (88% vs. 92%) while achieving the stage efficiency of the best dual-arm methods (169.2%). This shows that the DAG is a necessary representation for reliable and efficient dual-arm planning.

2) *Q2 Single-Query vs. Iterative Query Cost*: The cost advantage of DAG-Plan's single-query approach is definitive. RoCo's iterative paradigm consumes over 42k tokens per plan, making it prohibitively expensive. DAG-Plan reduces the token cost by nearly 20×, achieving a cost profile nearly identical to single-query linear methods but with vastly superior performance.

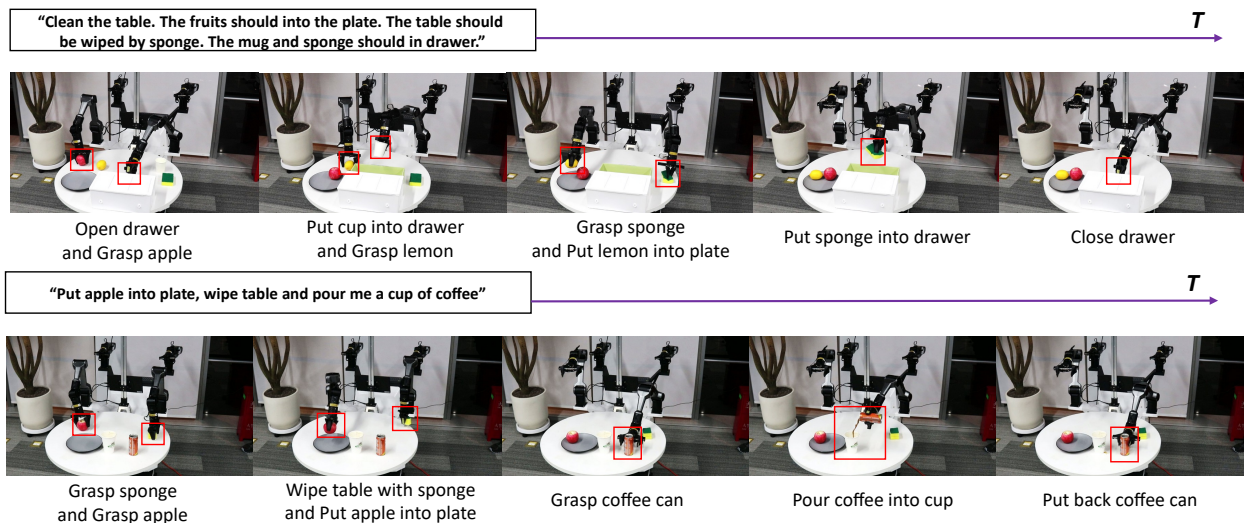


Fig. 5. Real-world snapshots of the execution process of long-horizon tasks.

3) Q3 Centralized Graph vs. Decentralized Multi-Agent:

The comparison with Twostep highlights a critical finding: decomposition into fully independent sub-tasks is ill-suited for the intrinsically collaborative nature of a dual-arm system. DAG-Plan’s centralized graph provides a global view, enabling it to explicitly model and enforce the necessary interdependencies between arms.

4) *Attribution of Success and Failure:* A deeper analysis of the planning results reveals the distinct failure modes of each baseline, highlighting the specific advantages of the DAG representation. TP-D fails on complex tasks (Task 2 & 4) due to the fundamental inability of linear sequences to model non-linear dependencies. Twostep fails as decomposing into independent sub-tasks is ill-suited for interdependent dual-arm coordination. RoCo and DAG-Plan share a common failure ceiling due to the LLM’s inherent limitation in temporal dependency prediction. DAG-Plan reaches this LLM-imposed ceiling, having eliminated failures caused by inferior representation TP-D or architecture Twostep.

C. Evaluation with Physical Simulation

This tier assesses the executability of the generated plans, testing whether high-level logical advantages translate into efficient and feasible physical execution. The simulation results (Tab. III and Fig. 4) strongly corroborate the planning-level findings and introduce critical insights into physical feasibility. It provides critical evidence that the performance gaps between methods are not just theoretical but manifest in simulation, thereby reinforcing the conclusion of Q1-3 and laying the groundwork for Q4.

1) *Bridging the Gap from Plan to Action:* A key insight is the stark difference between a logically sound plan and an executable one. TP-D’s performance collapses in simulation (26% SR vs. 40% in plan tests) because its linear sequences often generate physically infeasible actions (e.g., commanding both arms to grasp adjacent objects, causing collisions). DAG-Plan, equipped with its runtime feasibility checks, dynamically filters out such invalid actions, showcasing its critical role in bridging the sim-to-real gap.

2) *The Execution-Time Cost of Querying:* The results expose the crippling execution-time cost of the iterative paradigm. RoCo’s query time often exceeds its action execution time, leading to an abysmal overall execution efficiency of only 68.7%. This empirically confirms that query latency is a fundamental bottleneck. DAG-Plan’s execution efficiency reached 52.8% with the single-query approach completely avoids this penalty, achieving a time profile comparable to TP-D but with vastly higher reliability.

D. Real-world Experiments

We validate DAG-Plan in real-world setting to answer Q4, where a Agilex Cobot Magic dual-arm robot to complete tasks similar to simulation. When dealing with these tasks, the DAG-Plan framework can effectively assign the sub-task for the each arm. We show the execution process of DAG-Plan in real-world in Fig. 5 first row. We evaluate 10 runs for each task as shown in Fig. 6 for comparison of real-world and simulation success rates. The slightly lower success rate in real-world is primarily due to the inaccuracy of depth affect the foundation models for execution, whereas the depth in simulations is entirely precise.




Fig. 6. Comparison of real-world and simulation success rates of DAG-Plan.

In addition to the predefined tasks similar to simulation, we also conducted tests in the wild, as shown in Fig. 5 second row. By providing human instructions through audio input, the LLM automatically generates DAG, enabling the dual-arm robot to efficiently complete the tasks.

V. CONCLUSION

This work introduces DAG-Plan, which efficiently and accurately generates collaborative plans with LLMs for dual-arm robots. DAG-Plan decomposes complex tasks into directed acyclic graph (DAG) with clear temporal relationships and iteratively selects feasible sub-tasks based on environmental observations during execution. The main contribution is the replacement of task sequence with a DAG and the dynamic adjustment of the planning according to the current situation, allowing dual-arm robots to flexibly utilize both arms for sub-task execution. Compared to the single-arm approach, DAG-Plan demonstrates significantly higher execution efficiency. In contrast to dual-arm methods, DAG-Plan achieves a higher success rate in task planning. Compared to the iterative method, DAG-Plan maintains the ability to interact with environmental information while requiring shorter query time and reduced token usage.

REFERENCES

- [1] K.-C. Ying, P. Pourhejazy, C.-Y. Cheng, and Z.-Y. Cai, "Deep learning-based optimization for motion planning of dual-arm assembly robots," *Computers & Industrial Engineering*, vol. 160, p. 107603, 2021.
- [2] J. Borrell, C. Perez-Vidal, and J. V. Segura, "Optimization of the pick-and-place sequence of a bimanual collaborative robot in an industrial production line," *The International Journal of Advanced Manufacturing Technology*, vol. 130, no. 9, pp. 4221–4234, 2024.
- [3] S. S. Mirrazavi Salehian, N. B. Figueroa Fernandez, and A. Billard, "Dynamical system-based motion planning for multi-arm systems: Reaching for moving objects," in *International Joint Conference on Artificial Intelligence*, 2017, pp. 4914–4918.
- [4] J. Grannen, Y. Wu, S. Belkhal, and D. Sadigh, "Learning bimanual scooping policies for food acquisition," *arXiv preprint arXiv:2211.14652*, 2022.
- [5] Y. Mu, Q. Zhang, M. Hu, W. Wang, M. Ding, J. Jin, B. Wang, J. Dai, Y. Qiao, and P. Luo, "Embodiedgpt: Vision-language pre-training via embodied chain of thought," in *Advances in Neural Information Processing Systems*, vol. 36, 2023, pp. 25 081–25 094.
- [6] Y. Mu, J. Chen, Q.-L. Zhang, S. Chen, Q. Yu, C. Ge, R. Chen, Z. Liang, M. Hu, C. Tao, P. Sun, H. Yu, C. Yang, W. Shao, W. Wang, J. Dai, Y. Qiao, M. Ding, and P. Luo, "Robocodex: Multimodal code generation for robotic behavior synthesis," in *International Conference on Machine Learning*, vol. 235. PMLR, 2024, pp. 36 434–36 454.
- [7] I. Singh, D. Traum, and J. Thomason, "Twostep: Multi-agent task planning using classical planners and large language models," *arXiv preprint arXiv:2403.17246*, 2024.
- [8] Z. Mandi, S. Jain, and S. Song, "Roco: Dialectic multi-robot collaboration with large language models," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 286–299.
- [9] Z. Zhao, X. Yue, J. Xie, C. Fang, Z. Shao, and S. Guo, "A dual-agent collaboration framework based on llms for nursing robots to perform bimanual coordination tasks," *IEEE Robotics and Automation Letters*, vol. 10, no. 3, pp. 2942–2949, 2025.
- [10] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," in *Conference on robot learning*. PMLR, 2023, pp. 287–318.
- [11] M. Hu, Y. Mu, X. Yu, M. Ding, S. Wu, W. Shao, Q. Chen, B. Wang, Y. Qiao, and P. Luo, "Tree-planner: Efficient close-loop task planning with large language models," *arXiv preprint arXiv:2310.08582*, 2023.
- [12] 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 preprint arXiv:2304.11477*, 2023.
- [13] F. Joubin, A. Ceravola, P. Smirnov, F. Ocker, J. Deigmoeller, A. Bernardino, C. Wang, S. Hasler, D. Tanneberg, and M. Gienger, "Copal: Corrective planning of robot actions with large language models," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 8664–8670.
- [14] D. Sepúlveda, R. Fernández, E. Navas, M. Armada, and P. González-De-Santos, "Robotic aubergine harvesting using dual-arm manipulation," *IEEE Access*, vol. 8, pp. 121 889–121 904, 2020.
- [15] T. Yoshida, Y. Onishi, T. Kawahara, and T. Fukao, "Automated harvesting by a dual-arm fruit harvesting robot," *Robomech Journal*, vol. 9, no. 1, p. 19, 2022.
- [16] P. Ögren, C. Smith, Y. Karayiannidis, and D. Kragic, "A multi objective control approach to online dual arm manipulation1," *IFAC Proceedings Volumes*, vol. 45, no. 22, pp. 747–752, 2012.
- [17] F. Ju, H. Jin, and J. Zhao, "A kinematic decoupling whole-body control method for a mobile humanoid upper body robot," in *2023 International Conference on Frontiers of Robotics and Software Engineering (FRSE)*. IEEE, 2023, pp. 85–90.
- [18] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3929–3945, 2023.
- [19] D. Jiang, H. Wang, and Y. Lu, "Mastering the complex assembly task with a dual-arm robot: A novel reinforcement learning method," *IEEE Robotics and Automation Magazine*, vol. 30, no. 2, pp. 57–66, 2023.
- [20] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation," *arXiv preprint arXiv:2401.02117*, 2024.
- [21] N. Kokkalis, T. Köhn, J. Huebner, M. Lee, F. Schulze, and S. R. Klemmer, "Taskgenies: Automatically providing action plans helps people complete tasks," *ACM Transactions on Computer-Human Interaction*, vol. 20, no. 5, pp. 1–25, 2013.
- [22] S. Zhou, L. Zhang, Y. Yang, Q. Lyu, P. Yin, C. Callison-Burch, and G. Neubig, "Show me more details: Discovering hierarchies of procedures from semi-structured web data," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 2998–3012.
- [23] A. Hassan Awadallah, R. W. White, P. Pantel, S. T. Dumais, and Y.-M. Wang, "Supporting complex search tasks," in *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*, 2014, pp. 829–838.
- [24] R. Mehrotra and E. Yilmaz, "Extracting hierarchies of search tasks & subtasks via a bayesian nonparametric approach," in *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*, 2017, pp. 285–294.
- [25] Q. Yuan, M. Kazemi, X. Xu, I. Noble, V. Imbrasaitė, and D. Ramachandran, "Tasklama: probing the complex task understanding of language models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, 2024, pp. 19 468–19 476.
- [26] C. Aeronautiques, A. Howe, C. Knoblock, I. D. McDermott, A. Ram, M. Veloso, D. Weld, D. W. Sri, A. Barrett, D. Christianson *et al.*, "Pddl—the planning domain definition language," *Technical Report, Tech. Rep.*, 1998.
- [27] I. Armeni, Z.-Y. He, J. Gwak, A. R. Zamir, M. Fischer, J. Malik, and S. Savarese, "3d scene graph: A structure for unified semantics, 3d space, and camera," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 5664–5673.
- [28] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *European Conference on Computer Vision*. Springer, 2024, pp. 38–55.
- [29] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson *et al.*, "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024.
- [30] H. Lin, Z. Liu, C. Cheang, Y. Fu, G. Guo, and X. Xue, "Sar-net: Shape alignment and recovery network for category-level 6d object pose and size estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 6707–6717.
- [31] H. Geng, H. Xu, C. Zhao, C. Xu, L. Yi, S. Huang, and H. Wang, "Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7081–7091.
- [32] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 097–11 107.