

Integrated Exploration and Sequential Manipulation on Scene Graph with LLM-based Situated Replanning

Heqing Yang^{1,2} Ziyuan Jiao^{1,2†} Shu Wang³ Yida Niu^{2,4} Si Liu^{1†} Hangxin Liu²

Abstract—In partially known environments, robots must combine exploration to gather information with task planning for efficient execution. To address this challenge, we propose EPoG, an Exploration-based sequential manipulation Planning framework on Scene Graphs. EPoG integrates a graph-based global planner with a Large Language Model (LLM)-based situated local planner, continuously updating a belief graph using observations and LLM predictions to represent known and unknown objects. Action sequences are generated by computing graph edit operations between the goal and belief graphs, ordered by temporal dependencies and movement costs. This approach seamlessly combines exploration and sequential manipulation planning. In ablation studies across 46 realistic household scenes and 5 long-horizon daily object transportation tasks, EPoG achieved a success rate of 91.3%, reducing travel distance by 36.1% on average. Furthermore, a physical mobile manipulator successfully executed complex tasks in unknown and dynamic environments, demonstrating EPoG’s potential for real-world applications.

I. INTRODUCTION

To autonomously perform complex tasks, robots require an environment representation that is both expressive and readily queryable for planning. Recently, graph-based scene representations have emerged as a unifying paradigm [1], modeling scenes from 3D perception [2–4] and augmenting them with predicate-like attributes [5–7]. These representations facilitate action-centric reasoning [8] by grounding symbolic structure in perception, reducing the burden of manual domain engineering, and supporting complex manipulation tasks that are difficult to specify or solve with traditional pipelines [9–12].

Existing planning methods that leverage graph-based scene representations largely target two classes of problems: exploration and sequential manipulation. In exploration, robots navigate unknown or partially observed environments for objectives such as mapping [3] or object search [11]. These settings typically involve limited physical interaction with the scene, which constrains the range of tasks that can be executed. Sequential manipulation methods, in contrast, explicitly model interaction [6, 10] but often assume a fully known, static environment.

Real-world deployment violates these assumptions. A robot may have incomplete prior knowledge and may not immediately observe changes induced by human activity [13].

This work was supported in part by National Natural Science Foundation of China (Grant No. U25B6011, 52305007, U23B2010), by Beijing Natural Science Foundation (Grant No. 4262046), and by NSFC-RGC Project(N.CUHK498/24).

[†] Corresponding authors. This work was conducted during Heqing Yang’s internship at the Beijing Institute for General Artificial Intelligence (BIGAI).
¹ Beihang University. ² State Key Laboratory of General Artificial Intelligence, BIGAI, Beijing, China. ³ University of California, Los Angeles.
⁴ Institute for Artificial Intelligence, Peking University.

Instruction: Bring a banana[unknown] to the desk[known].



Fig. 1: An example illustrating the challenges of integrating exploration and sequential manipulation: (a) Robots must prioritize potential exploration locations and balance exploration with manipulation tasks to execute efficiently. (b) The robot needs to engage in situated planning to handle unexpected situations.

Consequently, it must interleave information gathering with task planning, which introduces three coupled challenges: 1) locating task-relevant objects under partial observability, requiring prioritized exploration; 2) trading off exploration and manipulation to reduce overall execution cost; and 3) reasoning under uncertainty to produce situated, executable plans. Fig. 1 illustrates these challenges. Although prior work addresses subsets of them [14–16], existing approaches often depend on hand-engineered heuristics, which limit robustness and scalability in long-horizon tasks.

To address these challenges, we propose EPoG, a framework that integrates Exploration and sequential manipulation Planning on Scene Graphs. EPoG adopts a bilevel planning architecture for partially observed environments, leveraging pretrained LLMs for informed exploration and situated replanning. As shown in Fig. 2, the global planner incrementally constructs a belief graph from onboard observations and LLM-based predictions, computes graph edit operations between the belief and goal graphs, and generates a candidate action sequence via topological sorting. The robot executes this sequence while continuously updating the belief graph with new observations and LLM predictions, thereby interleaving exploration and manipulation in a closed loop. When execution deviates from the nominal plan, the local planner invokes LLMs for situated replanning to resolve exceptions. We evaluate EPoG on five long-horizon object-transport tasks across 46 household scenes from the ProcThor-10k dataset [17]. Quantitative results demonstrate

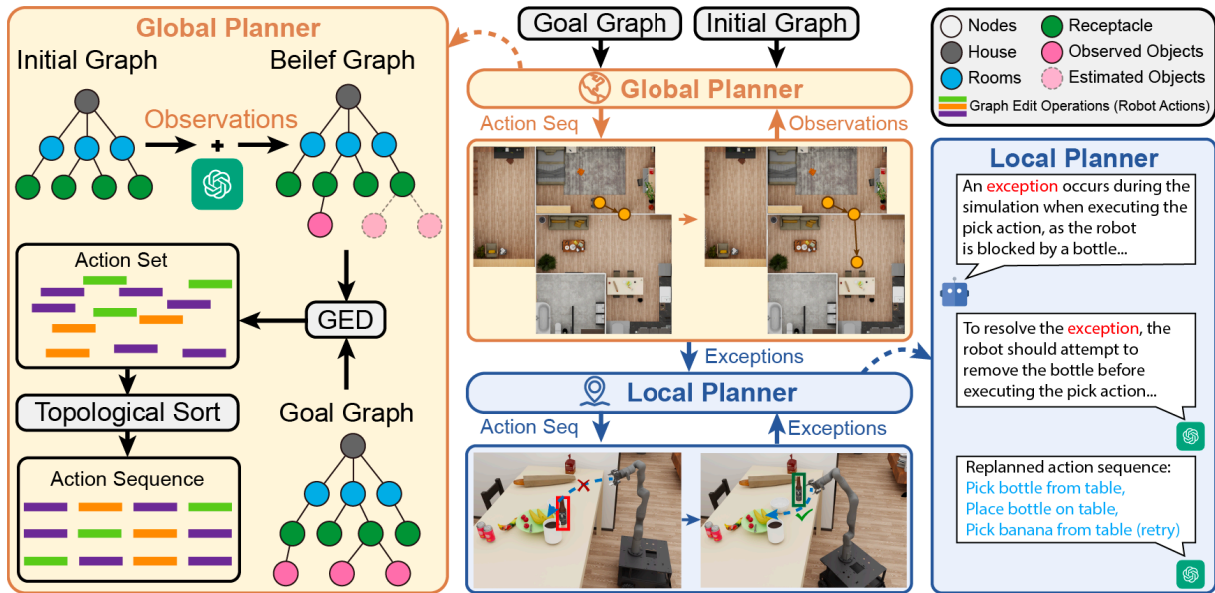


Fig. 2: **Overview of the proposed EPoG framework.** In the global planner, the belief graph is updated by estimating the locations of target objects present in the goal graph but missing in the initial graph, using new observations and LLM predictions. The action plan is then obtained by performing the topological sort on the graph edit operations between the belief and the goal graph. In the local planner, the LLM generates a situated action sequence to handle exceptions encountered during execution.

the effectiveness of EPoG, and ablations show that: i) the proposed formulation naturally couples exploration and manipulation, reducing total execution effort; ii) EPoG outperforms purely LLM-based planners on long-horizon tasks, achieving substantially higher success rates; and iii) LLM-guided heuristics and the low-level planner reduce manual design effort while improving execution efficiency. Experiments on a mobile robot further validate real-world applicability.

A. Related Work

Task planning with LLMs is increasingly popular in robotics for interpreting natural-language instructions and leveraging broad commonsense priors [18–20]. They can translate high-level commands into executable action sequences [21], enabling flexible task specifications across diverse environments. However, LLMs remain brittle on long-horizon problems due to limited spatial grounding, imperfect handling of temporal dependencies, and difficulty reasoning under environmental uncertainty [22–24]. To mitigate these limitations, recent work couples LLMs with classical planning. For example, Zhou *et al.* [25] use LLMs to generate Planning Domain Definition Language (PDDL) domain descriptions, while other approaches use LLMs as heuristics to guide search or exploration [19, 26]. Such hybrid systems combine the contextual flexibility of LLMs with the reliability and verifiability of symbolic planners. Nevertheless, most of these methods operate over predicate-based symbolic states, leaving task planning over graph-structured scene representations comparatively underexplored. Graph-based representations [3, 6, 9–12] offer a complementary substrate by integrating perception with structured spatial and semantic relations, often providing richer geometric context than flat predicates. This motivates studying how LLMs can be integrated with graph-based planning: LLMs can supply context-dependent priors and action proposals, while graph structure provides grounded relational context for reasoning, constraint checking, and replanning. Together, these ingredients may

yield more robust and efficient planning in complex, partially observed environments.

Task planning using scene graphs primarily falls into two distinct categories: methods based on known scene graphs that leverage predefined information for task planning [11, 27]; and methods based on unknown scene graphs, which involve robot exploration and the creation of new graphs for dynamic task execution [8]. However, these approaches face three key limitations: 1) Current methodologies predominantly utilize 2D topological relationships while neglecting essential 3D spatial attributes (*e.g.*, volumetric occupancy and geometric constraints), which can compromise task feasibility; 2) the inherent context length limits of transformer-based models, such as LLMs, impose restrictions on processing large-scale graphs, which is common for real-world environments; 3) current planning schemes for graph-based scene representations inadequately capture spatial-temporal relationships between scene entities that evolve over time or by human activities. These challenges motivate the investigation of new frameworks that can effectively harness both geometric and semantic information from graph-based scene representations while addressing computational scalability requirements.

Task planning in unknown environments requires robots to dynamically adapt plans based on new observations and continuously update their understanding of the environment [3, 28, 29]. Effective planning in such settings necessitates integrating exploration with task execution, balancing the acquisition of missing information with goal achievement. Existing approaches either explore the entire environment before planning [30] or rely on manually defined strategies [31], resulting in inefficiency and limited adaptability. In contrast, our work integrates exploration and planning using a graph-based representation that accounts for movement costs and employs LLMs for local replanning to handle exceptions, improving execution efficiency and versatility in dynamic environments.

II. PRELIMINARY

In this section, we first introduce the graph-based representation used in this work and then elaborate on how the planning problem is formulated using this representation.

A. Graph-Based Scene Representation

Following [2, 5, 10], the graph-based scene representation is structured as a tree, providing a hierarchical semantic abstraction of the environment. A scene graph $G = (V, E, A)$ consists of scene nodes V , edges E , and task-dependent attributes A associated with the nodes, which define the potential interactions with each node.

Scene Nodes $v_i \in V$ represent the entities in the scene, where $v_i = \langle o_i, c_i, M_i, B_i \rangle$. Here, o_i is the object's unique identifier in the graph, c_i is its semantic label, $M_i = \{m_i^j \mid j = 1, \dots, |M_i|\}$ is the set of geometric primitives describing its full geometry, and B_i is its 3D bounding box.

Scene Edges $e_{i,j} \in E$ are directed edges representing the supporting relationship between nodes. Each edge $e_{i,j} = \langle v_i, v_j, t_{i,j}, c_{i,j} \rangle$ is defined by a spatial transformation $t_{i,j}$ from the parent node v_i to the child node v_j and a semantic description $c_{i,j}$ of the supporting relationship, such as $c_{i,j} \in \{\text{on, contain}\}$.

Attributes $A = \{(A_i^s, A_i^c) \mid i \leq |V|\}$ represent the attributes assigned to each node v_i , where A_i^s is a supporting attribute indicating whether an object v_i can support another. The attribute A_i^c is optional and only assigned to container objects. Both attributes are used to verify whether a supporting relationship $c_{i,j}$ holds, which helps determine the feasibility of an action during task planning.

Based on the above definition, in this paper, an indoor scene is hierarchically structured into four levels: *Level 0*: a House as the root; *Level 1*: Rooms within the House; *Level 2*: Receptacles like containers and surfaces within the rooms, which serve as support structures; and *Level 3+*: Objects, which are the primary task targets. This hierarchical mapping supports downstream tasks and has demonstrated good results. Fig. 3 provides an example illustration of an indoor scene with this structure.

B. Problem Definition

In this work, we address the problem of sequential manipulation planning in a setting with unknown object states, a finite number of actions, and deterministic transitions. This problem \mathcal{P} is represented by a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{O}, s_{init}, s_{goal} \rangle$. Given an environment state $s \in \mathcal{S}$, an action $a \in \mathcal{A}$ can be selected from the set of applicable actions. The transition function $T : s \xrightarrow{a_i} s'$ defines the dynamics of the environment, indicating that executing action a_i in the state s results in a new state s' . The details of these actions will be introduced in Section III. The observation $obs = \{V_{obs}, E_{obs}\} \in \mathcal{O}$ is what the robot perceives upon reaching a new state. A solution to the planning problem \mathcal{P} is a sequence of actions $\pi = (a_1, \dots, a_k)$ that transforms the initial state s_{init} to the goal state s_{goal} .

In this paper, the state s is defined as G . The initial state s_{init} typically includes only House, Room, and Receptacle nodes, as well as the edges between them, excluding Object nodes, assuming the robot does not know object locations. This assumption is valid, as rooms and large




Fig. 3: An example illustration of an indoor scene.

objects are less likely to be frequently moved by humans, whereas smaller objects may be. The goal state s_{goal} is also represented as a graph, including all task-relevant nodes and their relationships. Object relationships are simplified for clarity into semantic descriptions, e.g., apple on table.

III. EPOG FRAMEWORK

The proposed EPOG framework addresses task planning for robots in partially known environments using scene graphs as the planning representation. It integrates exploration and sequential manipulation to generate efficient action sequences, enabling robots to explore unknown environments while completing tasks. A pre-trained LLM facilitates situated replanning to handle unexpected situations, improving robustness and reducing design effort. The following sections detail the framework and its components.

A. EPOG Overview

Algorithm 1 outlines the EPOG framework, which utilizes a bi-level planning scheme. The global planner begins with an initial belief graph representing the environment, combining known and unknown objects and their relationships. This graph is built from prior knowledge, including object locations, relationships, and spatial constraints, while accounting for uncertainty in areas that the robot has yet to explore, using commonsense knowledge from LLMs. As the robot executes tasks and gathers new observations, the belief graph is dynamically updated, supporting long-horizon planning through Graph Edit Distance (GED) and topological sorting. The process iterates until task completion. If exceptions occur during execution, an LLM-based local planner recursively resolves the issues until the action succeeds. The following sections detail the global and local planners.

B. Global Planner

`EstimateBeliefGraph(\cdot)` estimates the distribution of task-relevant objects and their relationships before planning. Initially, the belief graph may miss nodes and edges for task-relevant objects. For each missing node, a pre-trained LLM infers probable object locations, which are added to the belief graph. Two prompts are generated per object: (1) estimate the Room where the object is likely located, and (2) predict the Receptacle where it might

Algorithm 1: EPoG

```

Input : Initial Graph  $G_{init}$ ; Goal Graph  $G_g$ 
1 // Global Planner
2 // Initial estimation of the belief graph
3  $G_b \leftarrow \text{EstimateBeliefGraph}(G_{init});$ 
4 // Planning on scene graph; see Algorithm 2
5  $P_g \leftarrow \text{GraphBasedPlanner}(G_b, G_g);$ 
6 while  $P_g$  is not empty do
7    $action \leftarrow \text{Pop}(P_g);$ 
8    $obs, exception \leftarrow \text{RollOut}(action);$ 
9    $G_b, replan \leftarrow \text{UpdateGraph}(G_b, obs);$ 
10  if  $replan$  is not none then
11     $P_g \leftarrow \text{GraphBasedPlanner}(G_b, G_g);$ 
12  end
13  else if  $exception$  is not none then
14     $\text{LocalPlanner}(exception);$ 
15  end
16 end
17 // Local Planner
18 Function  $\text{LocalPlanner}(exception)$ :
19   // LLM-based planner; see Section III-C
20    $actions \leftarrow \text{LLMPlanner}(exception);$ 
21   foreach  $action$  in  $actions$  do
22      $obs, exception \leftarrow \text{RollOut}(action);$ 
23     // Recursively resolving exceptions
24     if  $exception$  is not none then
25        $\text{LocalPlanner}(exception);$ 
26     end
27   end

```

be placed. To complete the belief graph G_b , we add task-relevant nodes V_{task} and attributes A_{task} from G_g , and edges E_{est} based on the predicted relationships. For each edge $e_i = \langle v_i, v_j, t_{i,j}, c_{i,j} \rangle$, we do not predict the transformation $t_{i,j}$, as the object pose remains unknown until detection. In addition, we assume that the information for missing nodes $v_i = \langle o_i, c_i, M_i, B_i \rangle$ is given in the G_g . The final description of the updated belief graph is thus $G_b = (V_{init} \cup V_{task}, E_{init} \cup E_{est}, A_{init} \cup A_{task})$.

$\text{RollOut}(\cdot)$ is performed using a motion planner to generate the motion plan [20] before executing an action. If the motion planner raises an exception, the action will not be executed. Otherwise, it is carried out in the environment. After execution, the robot receives a new observation $obs = \{V_{obs}, E_{obs}\}$, where V_{obs} represents the visible objects and E_{obs} denotes the relationships between them. In this work, we assume that whenever the robot enters a Room, it can observe all objects and their relationships, except those contained within a closed Receptacle.

$\text{UpdateGraph}(\cdot)$ maintains the belief graph in sync with current observations, $\mathcal{O} = (V_{obs}, E_{obs})$, where V_{obs} represents visible objects (nodes) and E_{obs} captures their spatial and semantic relationships (edges). If a target object's initial position estimate e_{err} is absent from V_{obs} , the global planner updates the belief graph by replacing (v_i, e_{err}) with a new estimate (v_i, e_{new}) based on G_b . For instance, if an apple is expected on a table but is not found there, the planner re-estimates its location.

$\text{GraphBasedPlanner}(\cdot)$, as detailed in Algorithm 2, is invoked after updating the belief graph to generate a task plan with topological sort. Given an initial and goal graph,

Algorithm 2: Optimized Graph-Based Task Planner

```

Input : Belief graph  $G_b$ , Goal graph  $G_g$ 
Output : Optimal plan  $\pi^*$ 
1 // Step 1: Generate primitive actions
2  $K \leftarrow \text{GED}(G_b, G_g);$  // min-cost edit
   operations
3  $C \leftarrow \text{GetConstraints}(K);$  // constraints
   between actions
4 // Step 2: Initialize search frontier
5  $\text{Stack } stack \leftarrow \emptyset$ 
6  $min\_cost \leftarrow \infty$ 
7  $\pi^* \leftarrow \emptyset$ 
8 foreach  $a \in \text{UnconstrainedActions}((K, C))$  do
9    $stack.push([a, K \setminus \{a\}, C \setminus \{a\}]);$  // Search
   Node
10 end
11 // Step 3: Depth-first search
12 while  $\neg stack.empty()$  do
13    $(\pi, K_r, C_r) \leftarrow stack.pop()$ 
14    $cost \leftarrow \text{HeuristicCost}(\pi)$ 
15   // Prune suboptimal branches
16   if  $cost \geq min\_cost$  then
17     continue
18   else
19     // Goal check
20     if  $\text{IsGoalReached}(\pi)$  then
21        $\pi' \leftarrow \text{InsertWalkActions}(\pi)$ 
22       if  $cost < min\_cost$  then
23          $min\_cost \leftarrow cost$ 
24          $\pi^* \leftarrow \pi'$ 
25       end
26     continue
27   end
28   // Expand successors
29   foreach  $a \in \text{NextActions}(K_r, C_r)$  do
30      $\pi_{new} \leftarrow \pi \oplus a;$  // append action
31      $K_{new} \leftarrow K_r \setminus \{a\};$  // remove action
32      $C_{new} \leftarrow C_r \setminus \{a\};$  // remove
   constraints
33      $stack.push((\pi_{new}, K_{new}, C_{new}))$ 
34   end
35 end
36 end
37 return  $\pi^*$ 

```

GED provides a set of graph edit operations that transform the initial graph into the goal graph with minimal action cost. Formally,

$$\text{GED}(G_1, G_2) = \min_{\{a_1, \dots, a_k\} \in \mathcal{K}(G_1, G_2)} \sum_{i=1}^k \text{cost}(a_i) \quad (1)$$

where $\{a_1, \dots, a_k\} \in \mathcal{K}(G_1, G_2)$ denotes a set of edit operations transforming G_1 into G_2 , and $\text{cost}(a) \geq 0$ is the cost of each graph edit operation a . We consider four types of edit operations, corresponding to robot actions:

- $\text{delete}(e_{i,j}) \rightarrow \text{Pick}(v_i, v_j)$: Pick object v_j from v_i .
- $\text{insert}(e_{i,j}) \rightarrow \text{Place}(v_i, v_j)$: Place object v_j on v_i .
- $\text{substitute}(A_i^c, \text{opened}) \rightarrow \text{Open}(v_i)$: Open door v_i to make contained objects accessible and observable.
- $\text{substitute}(A_i^c, \text{closed}) \rightarrow \text{Close}(v_i)$: Close door v_i to make contained objects inaccessible and unobservable.

Based on this correspondence, we use GED to obtain a

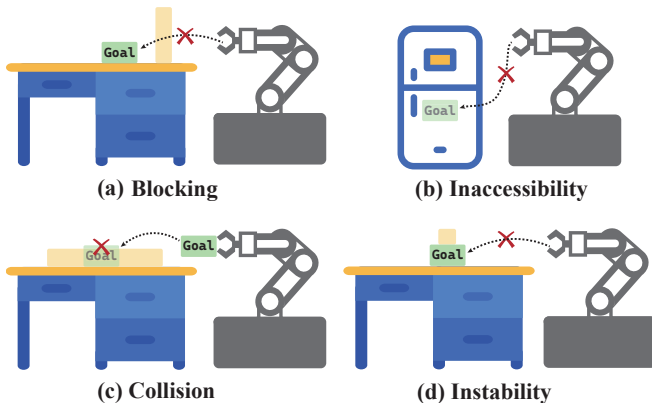


Fig. 4: **Examples of exceptions in motion planning.** (a) **Blocking:** The robot must avoid collisions with other objects while placing or picking up objects. (b) **Inaccessibility:** Successful object retrieval or placement within a container requires the container to be opened first. (c) **Collision:** The robot must ensure that placed objects do not collide with the environment. (d) **Instability:** The robot must maintain the stability of stacked objects during manipulation; for example, retrieving a book beneath a cup may cause instability.

set of necessary graph edit operations between the belief and goal graphs, resulting in a set of unsorted actions K . These actions have temporal dependencies, forming a partially ordered set. For example, for node v_i , the action $\text{pick}(v_i, v_j)$ must occur before $\text{place}(v_i, \cdot)$. These temporal dependencies create a constraint set $C = \{(a_i, a_j) \mid i \neq j\}$, where each action pair (a_i, a_j) must satisfy the condition $a_i < a_j$, meaning that a_i must occur before a_j . This ensures that actions are performed in the correct sequence, respecting the temporal order necessary to complete the task.

The task planning problem on G can then be formulated as a topological sorting problem on (K, C) , to generate a sorted sequence of actions π . To address this, we define a search node $\mathcal{N} = (K', C', \pi')$, where $K' \subseteq K$ represents the unexplored actions, $C' \subseteq C$ denotes the remaining temporal constraints, $\pi' \subseteq \pi$ is the subsequence of the action sequence. $K \setminus \{a\}$ is the action set except for the action $\{a\}$.

$\text{HeuristicCost}(\cdot)$ estimates the robot’s travel distance for the current action sequence. Since some object poses are unknown to the robot, a `Walk` action is inserted where the robot first moves near its parent `Receptacle` node. The simple yet effective A^* algorithm is used to estimate the travel distance between actions. As a result, `GraphBasedPlanner` accounts for the robot’s movement cost and seeks to minimize travel distance. Additionally, a pruning strategy is employed to continuously update the upper bound of the moving cost, effectively reducing the search space and optimizing the path planning.

C. Local Replanning

Robots often lack full awareness of environments with novel objects, leading to failures in global planners that account for positional uncertainties (e.g., object location errors) but overlook manipulative space constraints (e.g., joint limits, end-effector accessibility). Rule-based methods like those in [10] offer partial solutions via handcrafted heuristics but lack adaptability for long-horizon reasoning. As shown in Fig. 4, we introduce a local planner module to handle four motion planning exceptions. This taxonomy

System message: As an assistive robot, you must implement corrective actions to address errors that arise during task executions.
Task description: You will resolve errors during task executions using the following primitives: `Pick(x, y)` to pick item x from y , `Place(x, y)` to place item x on y , `Open(x)` to open x , and `Close(x)` to close x . Additionally, you may temporarily place objects in a designated parking area.

Example: The motion exception is: "Placing v_1 on v_2 fails because v_1 collides with v_3 . The parking place is v_0 ."

Analysis: Object v_1 collides with v_3 , so v_3 must be adjusted to allow v_1 to be placed on v_2 .

Robot Hand State: Since the failed action was a placement, my hand is occupied with v_1 .

Steps to Resolve:

1. Place v_1 in the parking area: `Place(1, 0)`.
 2. Remove the collision by picking v_3 : `Pick(3, 2)`.
 3. Move v_3 to the parking area: `Place(3, 0)`.
 4. Pick v_1 from the parking area: `Pick(1, 0)`.
 5. Retry the failed action by placing v_1 on v_2 : `Place(1, 2)`.
- I summarize the action sequence: [`Place(1, 0)`, `Pick(3, 2)`, `Place(3, 0)`, `Pick(1, 0)`, `Place(1, 2)`]

Question: You attempted to $\langle \text{failure_action} \rangle$, but it failed due to $\langle \text{exception} \rangle$. The parking place is $\langle \text{parking_place} \rangle$.

Fig. 5: **Prompt templates used by LLMPlanner.** Chain of thought (CoT) prompts, and action primitives prompts.

addresses two key limitations: 1) While global planning considers positional uncertainty, it neglects constraints from robotic embodiment (e.g., reachability, stability). Our classification systematically captures these interaction constraints; 2) The exception types align with standard motion planner diagnostics, allowing direct mapping to native planner outputs (e.g., Virtual Kinematic Chain (VKC)’s collision checking [32])—a critical feature for real-time exception handling in manipulation tasks. The prompts for guiding the LLM in inserting corrective actions for these exceptions during task execution are summarized in Fig. 5.

IV. SIMULATION AND EXPERIMENT

In simulations, we conduct an ablation study of the EPoG framework on five complex, long-horizon daily object transportation tasks across realistic scenes from the ProcThor-10k dataset [17]. Then, we illustrate the planned action sequences, highlighting EPoG’s efficiency in task execution. We finally validate the EPoG framework through experiments on a physical robotic mobile manipulator with perception, confirming its effectiveness in real-world scenarios.

A. Simulation Setup

We evaluate five complex, long-horizon daily object transportation tasks, summarized in Table I: 1) Breakfast Preparation, 2) Bedroom Work, 3) Movie and Snack Preparation, 4) Tea Making and Relaxation, and 5) Bath Preparation. A total of 46 scenes were filtered from the ProcThor-10k dataset, each containing task-relevant objects. Four types of exceptions were randomly generated near these objects in each scene, with two exceptions present per scene. In the study, each task is performed once for each scene, and the following metrics are evaluated: **%SR:** The percentage of successfully completed tasks out of the total scenes. **%EN:** The average percentage difference in newly explored nodes.

%TD: The average percentage difference in travel distance (in meters) required to complete the task in each scene. Both %EN and %TD are calculated relative to the results of EPoG.

TABLE I: Benchmark tasks and their goal predicates.

No.	Task	Scenes	Goal Conditions
1	Breakfast Prep.	10	{apple, bread, fork} → plate plate → diningtable
2	Bedroom Work	10	{alarmclock, CD, laptop, pencil} → desk
3	Movie & Snack	10	remote → sofa bread → plate; plate → diningtable
4	Tea & Relax	10	kettle → countertop; cup → diningtable remote → sofa
5	Bath Prep.	6	{soapbottle, cloth} → faucet

B. Ablation Study

To evaluate embodied task planning under environmental uncertainty systematically, we designed three baselines for the ablation study:

- The **LLM Planner** isolates data-driven planning by testing how LLMs overcome geometric infeasibility via iterative self-correction without explicit environment modeling—a scenario common in zero-shot deployments [21, 29].
- The **Exploration+LLM Planner** variant proactively explores the environment with optimized travel distance, establishing an upper bound for perception-planning pipelines in the absence of environmental priors. This configuration evaluates the extent to which exploration mitigates manipulative uncertainty.
- The **Exploration+PoG** variant implements a traditional rule-based planner with explicit exception handling from [10], serving as a reference for planners that rely on predefined spatiotemporal constraints and revealing differences in handling open-world task variance.

Our baseline design shows that planners integrating LLMs, including the **Exploration+PoG**, leverage the LLM’s open-world understanding to facilitate execution efficiency. However, LLMs may suffer from hallucinations and reasoning inconsistencies in complex or uncertain environments, whereas traditional symbol-based planners tend to achieve higher success rates due to their deterministic, rule-following nature.

Table II summarizes the results. Overall, the LLM planner exhibits lower success rates, particularly in long-horizon planning, as it often fails (with a maximum of 20 retries allowed per scene) to predict the next action correctly, such as walking to the wrong location for pick or place, or failing to find all task-relevant objects. This is because the LLM struggles to reason over complex, information-rich scene graphs and use them effectively for planning. Additionally, even in successful cases, the LLM planner generates less efficient plans for higher %TD. In contrast, EPoG achieves an overall success rate of 91.3%, with a reduction of 40.0% in explored nodes and 36.2% in travel distance compared

to the Exploration+PoG baseline. This is due to EPoG’s integrated manipulation planning and informed exploration strategy, which results in lower %EN and %TD. EPoG continuously updates the belief graph with LLM heuristics and observations, optimizing both exploration and task execution. However, EPoG still encounters 8.7% planning failures under compound exception conditions. Although EPoG’s global planning algorithm guarantees complete task accomplishment through graph editing operations, its local planning module manifests cognitive ambiguities when handling compound exceptions induced by combinatorial object configurations. Despite minor performance losses in the local planner, the LLM-based local planner effectively handles random exceptions without requiring manually defined rules.

C. Case Study

Fig. 6 provides a graphical illustration of the robot’s paths during two tasks. For clarity, only movements between Rooms or Receptacles are shown, omitting actions performed on Objects within the same Receptacle. In Fig. 6(a), EPoG demonstrated efficient task execution by first placing items on a plate in the kitchen (① → ② → ③), then transporting the plate to the living room (→ ④). In contrast, the LLM-based method did not effectively use the plate as a transport tool, instead opting to carry objects separately, leading to a longer action sequence and increased travel distance. In Fig. 6(b), EPoG initially used the LLM’s commonsense knowledge to incorrectly estimate the cloth’s location on the nightstand. However, EPoG enabled the robot to update the belief graph dynamically with new observations and regenerate a task plan. The robot then efficiently located the cloth in a cabinet and completed the task. Both tasks demonstrate EPoG’s ability to seamlessly integrate exploration and planning, resulting in efficient task execution by reducing overall travel distance. In contrast, exploration-first strategies led the robot to explore unnecessary locations. The LLM planner struggled with the complexity of large scene graphs, resulting in unnecessary travel between irrelevant locations and increased total travel distance. Although the PoG method optimized the plan, the separation of exploration and planning led to higher execution effort due to longer travel distances.

D. Experiment Setup

We conducted two real-world experiments on a physical mobile manipulator to evaluate EPoG’s ability to plan with unknown object states and resolve exceptions during task execution. Notably, our physical robot experiments focus on validating the task planning performance of the EPoG framework in real-world environments, while deliberately abstracting away implementation details of robot perception. The experiments utilized VKC-based mobile manipulation planning [32, 33] and Curobo [34] to generate real-time whole-body trajectories for object interactions.

Fig. 7 illustrates the task setup. In the first experiment (Fig. 7(a)), the robot is tasked with moving a basket onto a coffee table and placing a toy cabbage and a cup into the basket. Initially, the robot is aware of the cabbage and basket

TABLE II: Ablation Study

Method	Breakfast Prep.			Bedroom Work			Movie and Snack Prep.			Tea Making & Relax.			Bath Prep.			Total		
	%SR \uparrow	%EN \downarrow	%TD \downarrow	%SR \uparrow	%EN \downarrow	%TD \downarrow	%SR \uparrow	%EN \downarrow	%TD \downarrow	%SR \uparrow	%EN \downarrow	%TD \downarrow	%SR \uparrow	%EN \downarrow	%TD \downarrow	%SR \uparrow	%EN \downarrow	%TD \downarrow
LLM	10.0	+0.0	+110	10.0	+0.0	+48.7	10.0	+0.0	+179	40.0	+10.3	+115	16.7	+0.0	+16.7	17.4	+2.05	+93.9
Exp. + LLM	40.0	+77.9	+80.6	40.0	+75.9	+84.0	40.0	+50.4	+248	60.0	+75.0	+124	16.7	+168	+63.2	41.3	+89.4	+120
Exp. + PoG	100	+77.2	+59.0	100	+50.4	+25.0	100	+66.5	+112	100	+67.9	+122	100	+104	+60.5	100	+73.2	+75.6
EPoG (ours)	100	52.1	35.8	100	50.8	52.1	80.0	52.9	33.0	90.0	52.1	23.5	83.3	46.4	43.6	91.3	51.9	37.8




Fig. 6: Illustrations of the robot’s paths for two tasks, highlighting key locations and actions.

states, but the cup’s location is unknown. Using LLM, the robot infers that the cup might be on the coffee table, and the global planner of EPoG generates an optimized task plan to minimize travel distance. The robot first places the cabbage in the basket, moves the basket to the table, and then plans to retrieve the cup from the coffee table and place it into the basket. However, when the cup is not found on the table, the robot updates its belief graph and ultimately locates the cup on top of a cabinet, successfully completing the task. In the second experiment (Fig. 7(b)), the robot is instructed to place a tissue box on a coaster on the coffee table. The robot must handle exceptions like inaccessibility, blocking, and potential collisions during the task. The LLM correctly estimates the tissue box to be inside a closed closet, but the contents of the closet are unknown. After opening the closet, the robot discovers that a cup is blocking the tissue box, so it temporarily moves the cup to a shelf. Similarly, when a tea can obstructs the coaster, the robot sets down the tissue box, removes the tea can, and successfully places the box on the coaster. Throughout the process, EPoG primarily relies on its local planning module to manage these exceptions effectively. The robot’s execution details are documented in the supplementary video.

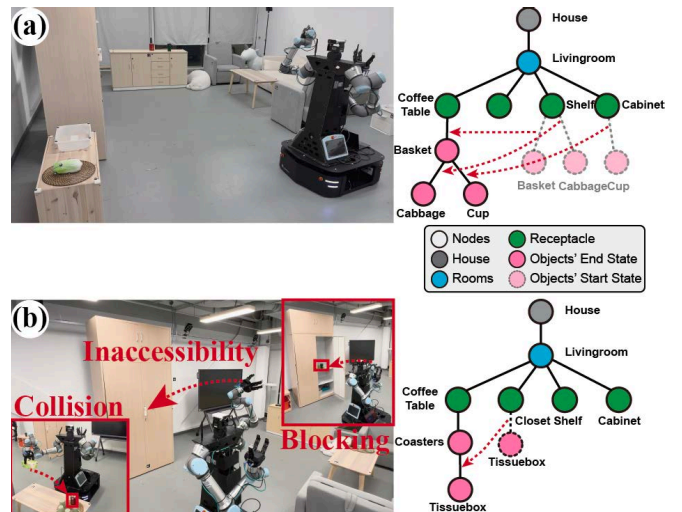


Fig. 7: **Real-world experiments.** (a) The robot must place three objects (with one unknown object) on a coffee table. (b) The robot must retrieve a tissue box from a closet and place it on a coaster, first removing a cup and a tea can that obstruct task execution. Object states in the closet and on the coffee table are initially unknown.

V. CONCLUSION

In conclusion, we presented EPoG, a planning framework that effectively integrates exploration and task planning within unknown environments using graph-based representations. By utilizing a bi-level planning scheme, EPoG combines graph-based sequential planning with LLM-based local replanning, enabling robust long-horizon operation and effective handling of unexpected situations. Studies across daily complex object transportation tasks demonstrated EPoG's efficiency in task execution and potential for real-world robotic applications. Future integration of a 3D scene graph perception module, along with the incorporation of Task and Motion Planning (TAMP) methods [35, 36], could further improve its practicality, advancing robots' capabilities in real-world scenarios.

REFERENCES

- [1] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [2] M. Han, Z. Zhang, Z. Jiao, X. Xie, Y. Zhu, S.-C. Zhu, and H. Liu, "Reconstructing interactive 3d scenes by panoptic mapping and cad model alignments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [3] A. Rosinol, A. Violette, M. Abate, N. Hughes, Y. Chang, J. Shi, A. Gupta, and L. Carlone, "Kimera: From slam to spatial perception with 3d dynamic scene graphs," *International Journal of Robotics Research (IJRR)*, 2021.
- [4] S.-C. Wu, K. Tateno, N. Navab, and F. Tombari, "Incremental 3d semantic scene graph prediction from rgb sequences," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [5] M. Han, Z. Zhang, Z. Jiao, X. Xie, Y. Zhu, S.-C. Zhu, and H. Liu, "Scene reconstruction with functional objects for robot autonomy," *International Journal of Computer Vision (IJCV)*, 2022.
- [6] C. Agia, K. M. Jatavallabhula, M. Khodeir, O. Miksik, V. Vineet, M. Mukadam, L. Paull, and F. Shkurti, "Taskography: Evaluating robot task planning over large 3d scene graphs," in *Conference on Robot Learning (CoRL)*, 2022.
- [7] A. Rosinol, A. Gupta, M. Abate, J. Shi, and L. Carlone, "3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans," *Robotics: Science and Systems (RSS)*, 2020.
- [8] K. Rana, J. Haviland, S. Garg, J. Abou-Chakra, I. Reid, and N. Sunderhauf, "Sayplan: Grounding large language models using 3d scene graphs for scalable task planning," *arXiv preprint*, 2023.
- [9] S. Chen, P.-L. Guhur, M. Tapaswi, C. Schmid, and I. Laptev, "Think global, act local: Dual-scale graph transformer for vision-and-language navigation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [10] Z. Jiao, Y. Niu, Z. Zhang, S.-C. Zhu, Y. Zhu, and H. Liu, "Sequential manipulation planning on scene graph," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [11] Q. Gu, A. Kuwajerwala, S. Morin, K. M. Jatavallabhula, B. Sen, A. Agarwal, C. Rivera, W. Paul, K. Ellis, R. Chellappa, *et al.*, "Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning," *arXiv preprint*, 2023.
- [12] R. Liu, X. Wang, W. Wang, and Y. Yang, "Bird's-eye-view scene graph for vision-language navigation," in *IEEE International Conference on Computer Vision (ICCV)*, 2023.
- [13] S. Hopko, J. Wang, and R. Mehta, "Human factors considerations and metrics in shared space human-robot collaboration: A systematic review," *Frontiers in Robotics and AI*, 2022.
- [14] Z. Yang, L. Ning, H. Wang, T. Jiang, S. Zhang, S. Cui, H. Jiang, C. Li, S. Wang, and Z. Wang, "Text2reaction: Enabling reactive task planning using large language models," *Robotics and Automation Letters (RA-L)*, 2024.
- [15] M. A. do Carmo Alves, A. Varma, Y. Elkhatib, and L. Soriano Marcolino, "Information-guided planning: an online approach for partially observable problems," *Advances in Neural Information Processing Systems (NIPS)*, vol. 36, pp. 69157–69177, 2023.
- [16] C. R. Garrett, C. Paxton, T. Lozano-Pérez, L. P. Kaelbling, and D. Fox, "Online replanning in belief space for partially observable task and motion problems," in *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020.
- [17] M. Deitke, E. VanderBilt, A. Herrasti, L. Weihs, K. Ehsani, J. Salvador, W. Han, E. Kolve, A. Kembhavi, and R. Mottaghi, "Proctor: Large-scale embodied ai using procedural generation," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, pp. 5982–5994, 2022.
- [18] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [19] Z. Zhao, W. S. Lee, and D. Hsu, "Large language models as common-sense knowledge for large-scale task planning," *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [20] 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," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [21] C. H. Song, J. Wu, C. Washington, B. M. Sadler, W.-L. Chao, and Y. Su, "Llm-planner: Few-shot grounded planning for embodied agents with large language models," in *IEEE International Conference on Computer Vision (ICCV)*, 2023.
- [22] 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 preprint arXiv:2402.01817*, 2024.
- [23] K. Valmeekam, A. Olmo, S. Sreedharan, and S. Kambhampati, "Large language models still can't plan (a benchmark for llms on planning and reasoning about change)," *arXiv preprint*, 2022.
- [24] Y. Yamada, Y. Bao, A. K. Lampinen, J. Kasai, and I. Yildirim, "Evaluating spatial understanding of large language models," *Transactions on Machine Learning Research*, 2024.
- [25] 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," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2081–2088, IEEE, 2024.
- [26] D. Shah, M. R. Equi, B. Osifski, F. Xia, B. Ichter, and S. Levine, "Navigation with large language models: Semantic guesswork as a heuristic for planning," in *Conference on Robot Learning (CoRL)*, pp. 2683–2699, PMLR, 2023.
- [27] Z. Ni, X. Deng, C. Tai, X. Zhu, Q. Xie, W. Huang, X. Wu, and L. Zeng, "Grid: Scene-graph-based instruction-driven robotic task planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 13765–13772, IEEE, 2024.
- [28] J. Rintanen, "Complexity of planning with partial observability," in *ICAPS*, vol. 4, pp. 345–354, 2004.
- [29] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, "React: Synergizing reasoning and acting in language models," *arXiv preprint*, 2022.
- [30] H. Jiang, B. Huang, R. Wu, Z. Li, S. Garg, H. Nayyeri, S. Wang, and Y. Li, "Roboexp: Action-conditioned scene graph via interactive exploration for robotic manipulation," *arXiv preprint arXiv:2402.15487*, 2024.
- [31] Y. Xu, Z. Zhang, J. Yu, Y. Shen, and Y. Wang, "A framework to co-optimize robot exploration and task planning in unknown environments," *Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 12283–12290, 2022.
- [32] Z. Jiao, Z. Zhang, W. Wang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu, "Efficient task planning for mobile manipulation: a virtual kinematic chain perspective," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [33] Z. Li, Y. Niu, Y. Su, H. Liu, and Z. Jiao, "Dynamic planning for sequential whole-body mobile manipulation," in *IEEE Conference on Industrial Electronics and Applications (ICIEA)*, IEEE, 2024.
- [34] B. Sundaralingam, S. K. S. Hari, A. Fishman, C. Garrett, K. Van Wyk, V. Blukis, A. Millane, H. Oleynikova, A. Handa, F. Ramos, *et al.*, "Curobo: Parallelized collision-free robot motion generation," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8112–8119, IEEE, 2023.
- [35] Z. Jiao, Z. Zhang, X. Jiang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu, "Consolidating kinematic models to promote coordinated mobile manipulations," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [36] Z. Jiao, Y. Niu, Z. Zhang, Y. Wu, Y. Su, Y. Zhu, H. Liu, and S.-C. Zhu, "Integration of robot and scene kinematics for sequential mobile manipulation planning," *IEEE Transactions on Robotics (T-RO)*, 2025.