

# Compose by Focus: Scene Graph-based Atomic Skills

Han Qi<sup>1</sup>, Changhe Chen<sup>2</sup>, and Heng Yang<sup>1</sup>

<sup>1</sup>School of Engineering and Applied Sciences, Harvard University

<sup>2</sup>Robotics Department, University of Michigan

**Abstract**—A key requirement for generalist robots is compositional generalization—the ability to combine atomic skills to solve complex, long-horizon tasks. While prior work has primarily focused on synthesizing a planner that sequences pre-learned skills, robust execution of the individual skills themselves remains challenging, as visuomotor policies often fail under distribution shifts induced by scene composition. To address this, we introduce a *scene graph-based* representation that *focuses on* task-relevant objects and relations, thereby mitigating sensitivity to irrelevant variation. Building on this idea, we develop a scene-graph skill learning framework that integrates graph neural networks with diffusion-based imitation learning, and further combine “focused” scene-graph skills with a vision-language model (VLM) based task planner. Experiments in both simulation and real-world manipulation tasks demonstrate substantially higher success rates than state-of-the-art baselines, highlighting improved robustness and compositional generalization in long-horizon tasks.

## I. INTRODUCTION

Solving long-horizon manipulation tasks requires breaking the whole task into multiple sub-tasks and executing individual skills (i.e., manipulation primitives). While much prior work has focused on how to compose already learned skills (e.g., by learning a high-level planner) [1], [2], [3], in this paper we address an orthogonal problem: how should the individual skills themselves be constructed so that they can be effectively composed?

**Motivating example.** Suppose the task is to pick up all the vegetables and put them into a basket. When presented with a cluttered scene containing many objects (including vegetables and distractors, as shown in Fig. 1), a high-level planner—such as a Vision-Language Model (VLM) or Task and Motion Planner (TAMP) [4]—can decompose the task into sub-goals: 1) pick up the carrot, 2) pick up the eggplant, 3) pick up the corn, etc. However, if each skill has only been trained in a plain environment (e.g., imitation learning from demonstrations with a single object on a clean table), this policy often fails in cluttered settings (see evaluation results in §IV). The failure is not due to the planner but arises because the visuomotor policies lack robustness to distribution shifts. Prior studies [5] also find that state-of-the-art imitation learning methods struggle on compositional tasks, largely due to brittle visual processing.

We argue that for skills to be composable, they must be *focused*—attending only to scene elements relevant to the skill at hand while ignoring “distractors”.

Towards this goal, we introduce a scene graph-based representation of visual data for efficient skill composition. Whereas most state-of-the-art approaches represent observations as raw RGB images [6], [7], [8] or 3D point




Fig. 1: **Real world vegetable picking.** A policy trained to pick up a single vegetable on a clean table is evaluated by placing all vegetables into the basket in a cluttered scenario.

clouds [9], [10], we instead transform visual input into dynamic semantic 3D scene graphs [11]. These graphs are *relevant-object centric*—constructed to include only objects and relations pertinent to the current skill, thus filtering out irrelevant visual noise. Nodes encode 3D geometry and semantic features of objects, while edges dynamically capture inter-object relations inferred from multimodal cues. We build these graphs by leveraging visual foundation models (e.g., Grounded-SAM [12]) for object segmentation and vision-language models (e.g., ChatGPT) for relation inference, yielding structured inputs that bypass raw image processing. We then apply Graph Neural Networks (GNNs) [13] to extract graph features, which are then used to condition diffusion-based visuomotor policies. This design substantially mitigates distribution shift in skill execution and enables robust composition of many atomic skills into long-horizon behaviors. In short, our method proposes using *scene graph* as an effective and interpretable way of data representation for *skill composition*, improving the *compositional generality* for solving long-horizon manipulation tasks.

**Contributions.** Our main contributions are:

- 1) We propose encoding structural scene graphs as general and interpretable inputs for vision-based policy learning using behavior cloning, where graphs capturing objects and relations are constructed with the help of VLMs and visual foundation models.
- 2) We integrate this representation with diffusion-based imitation learning and evaluate it on new skill composition manipulation benchmarks, showing substantial improvements over state-of-the-art baselines and strong robustness to visual perturbations in both simulation and real-world settings.

## II. RELATED WORK

### A. Compositional Generalization

Skill composition is essential for generalization in robotics, as real-world tasks are long-horizon and multi-skilled. Decomposing tasks into sub-goals enables flexible recombination, and recent advances in language models [14], [15] have supported structured planning. Works combining LLMs with robotics [2], [3] focus on high-level planning via meta-skill libraries [16], [17], but low-level execution remains underexplored. Task and motion planning (TAMP) [18], [19], [20] integrates discrete and continuous planning but depends on symbolic or geometric models, limiting generalization to complex or dynamic scenes. More recently, object-centric methods [21], [22] use VLMs for segmentation [23], [24] to build 3D features, improving over raw inputs but lacking explicit structure and broad generalization. In contrast, we propose scene graphs that encode objects and relations, enabling end-to-end skill composition and robust recomposition of atomic skills.

### B. Visual Imitation Learning

Imitation learning is an efficient approach for training manipulation policies, with visual imitation learning becoming increasingly popular due to challenges in state estimation and reward engineering. Most works [6], [25], [26], [27], [28] rely on 2D images as policy input, and robot foundation models similarly consume 2D data [7], [8]. While effective, 2D inputs lack spatial richness; 3D point-cloud policies [9] and 4D temporal extensions [29] add structure but still lack explicit reasoning of objects and relations.

Beyond raw 2D/3D features, zero-shot generalization has been explored via disentangled latent spaces [30], [31], but these often sacrifice interpretability. In contrast, we propose an explicit structural representation: scene graphs that encode objects and relations while embedding 3D information. This principled design supports robust skill composition by focusing on task-relevant context across varied environments.

### C. Scene Graphs in Robotics

Scene graphs have been widely explored as structured representations for perception and reasoning. Early work introduced 3D scene graphs as a way to unify semantics, geometry, and spatial relationships in visual scenes [11]. More recently, scene graphs have been used in robotics for high-level reasoning, such as task planning and navigation in cluttered or embodied environments [32], [33]. However, these approaches typically employ scene graphs as auxiliary structures for planning or symbolic reasoning, rather than as direct inputs to low-level visuomotor policies. In contrast, our work explicitly encodes scene graphs into a behavior cloning framework, enabling robust skill composition under visual variation.

## III. METHOD

### A. Problem Setup

We address long-horizon manipulation tasks by composing atomic skills. Training data consist of expert demonstrations

of individual skills, each manipulating objects toward a short-term goal (e.g., placing an apple in a bowl or pulling a cube with a tool). Unlike fixed primitives such as “pick” or “place” [34], [35], our formulation defines atomic skills by task context. Importantly, demonstrations cover skills only in isolation, never their compositions in cluttered scenes.

Formally, we transform observations  $O$  into scene graphs  $G$  and have skill descriptions as  $L$ . A single visuomotor policy  $\pi : (G, L) \mapsto A$  is trained to execute all atomic skills ( $A$  denotes the action space). The pipeline is illustrated in Fig. 2. We next describe (i) scene graph construction (§III-B), (ii) policy learning on graph features (§III-C), and (iii) test-time skill composition (§III-D).

### B. Scene Graph Construction

We obtain RGB and depth images from expert demonstrations in both simulation and real-world settings. Depth is converted to point clouds using camera intrinsics and extrinsics [9], and the robot gripper’s point cloud is assumed available. Each atomic skill is paired with a language description specifying the relevant objects  $B$ .

To extract object-level information, we employ a vision foundation model (Grounded SAM [12]) to segment masks of the *task-relevant* objects from RGB images and obtain their corresponding point clouds. These point clouds are then downsampled using farthest point sampling [36] and encoded into compact vector representations via a lightweight MLP network, DP3 Encoder [9]. Each object’s embeddings serve as the nodes of the scene graph.

Edges capture dynamic inter-object relations (e.g., ‘grasp’, ‘next’, ‘inside’), inferred from RGB with a VLM (e.g., ChatGPT). The resulting subgraph focuses only on task-relevant entities—robot gripper, objects, target, and optional obstacles (Fig. 3). Training on such concise subgraphs enables robust composition of learned skills in complex scenarios.

### C. Multi-skill Policy Training

After transforming every frame of observation data into a task relevant sub-scene-graph, we implement a two-layer Graph Attention Network (GAT) [37] to transform the scene graph into feature embedding.

Each node  $i \in \mathcal{V}$  is initialized with an input feature  $\mathbf{h}_i^{(0)} \in \mathbb{R}^{d_{in}}$ . At each GAT layer  $\ell$ , node features are updated as

$$\mathbf{h}_i^{(\ell+1)} = \sigma_\ell \left( \left\| \sum_{m=1}^{H_\ell} \sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(\ell,m)} \mathbf{W}^{(\ell,m)} \mathbf{h}_j^{(\ell)} \right\| \right),$$

where  $\alpha_{ij}^{(\ell,m)}$  denotes the attention coefficient between node  $i$  and  $j$  under head  $m$ , and  $\mathbf{W}^{(\ell,m)}$  is the learnable weight matrix of head  $m$ . Here,  $\left\| \sum_{m=1}^{H_\ell} \right\|$  denotes the concatenation of the outputs from all  $H_\ell$  attention heads. Formally, if each head produces  $\mathbf{z}_i^{(\ell,m)} \in \mathbb{R}^d$ , then

$$\left\| \sum_{m=1}^{H_\ell} \mathbf{z}_i^{(\ell,m)} \right\| = [\mathbf{z}_i^{(\ell,1)}; \mathbf{z}_i^{(\ell,2)}; \dots; \mathbf{z}_i^{(\ell,H_\ell)}] \in \mathbb{R}^{H_\ell d}.$$

We use  $H_0 = 4$  heads with concatenation and  $\sigma_0 = \text{ELU}$  (Exponential Linear Unit) in layer 1,  $H_1 = 1$  head

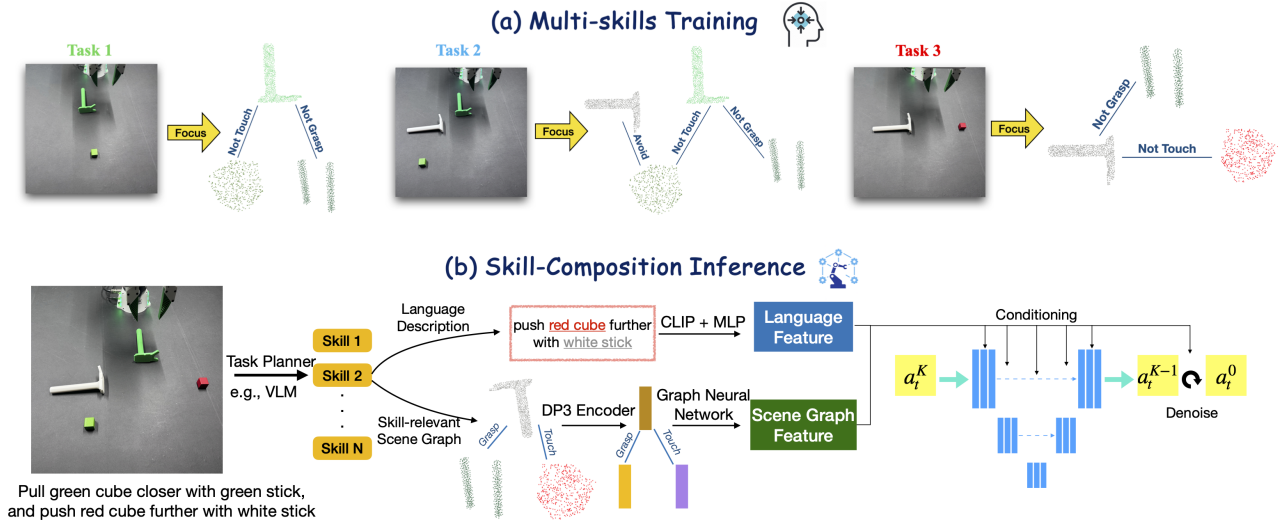


Fig. 2: **Compose by focus.** (a) A single policy is trained on focused scene graphs across all sub-skills. (b) During inference, a task planner (e.g., VLM) decomposes a long-horizon task into  $N$  sub-skills. For each sub-skill, CLIP encodes the description, Grounded SAM segments the relevant objects and extracts point clouds as graph nodes, and edges represent inter-object relations. DP3 Encoder [9] embeds the nodes, a GNN encodes the scene graph, and a diffusion policy conditioned on both description and graph features iteratively denoises actions.

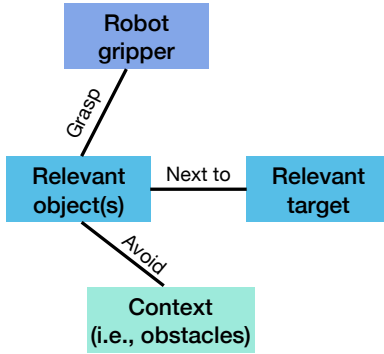


Fig. 3: **Illustration of a scene graph.**

without concatenation and  $\sigma_1 = \text{Id}$  in layer 2, yielding node embeddings  $\{\mathbf{h}_i^{(2)}\}$ . Graph representation (global mean pool) is then obtained as

$$\mathbf{F} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \mathbf{h}_i^{(2)}.$$

With the embedded feature vector, this structured representation is able to be integrated into the state-of-the-art imitation learning framework, Diffusion Policy [6].

For the language description of every atomic skill, we encode the text with CLIP encoder [38], getting the skill description feature  $P$ .

We implement a visuomotor policy with conditional denoising diffusion model [6] to learn these atomic skills, conditioning on scene graph features  $\mathbf{F}$ , skill description features  $P$  and robot poses  $Q$  during  $T_o$  observation horizons and then denoises random Gaussian noises into actions  $A_t$ , containing  $T_p$  action steps starting from time  $t$ . Concretely, starting from the Gaussian noise  $A_t^K$ , the denoising network  $\epsilon_\theta$  performs  $K$  iterations to denoise the random noise  $A_t^K$

into the noise-free action  $A_t^0$ ,

$$A_t^{k-1} = \alpha_k (A_t^k - \gamma_k \epsilon_\theta (A_t^k, k, \mathbf{F}, P, Q)) + \sigma_k \mathcal{N}(0, \mathbf{I}), \quad (1)$$

where  $\mathcal{N}(0, \mathbf{I})$  is Gaussian noise,  $\alpha_k$ ,  $\gamma_k$ , and  $\sigma_k$  depend on the noise scheduler.

To train the denoising network  $\epsilon_\theta$ , we add the noise  $\epsilon^k$  at the  $k$ -th iteration to action  $A_t^0$  [6], and the training objective is to predict the noise added to the original action:

$$\mathcal{L} = \text{MSE}(\epsilon^k, \epsilon_\theta(\bar{\alpha}_k A_t^0 + \bar{\beta}_k \epsilon^k, k, \mathbf{F}, P, Q)). \quad (2)$$

where  $\bar{\alpha}_k$  and  $\bar{\beta}_k$  are the noise scheduler. In implementation, we end-to-end train the point cloud encoder, graph encoder and diffusion model.

#### D. Testing Time Skill Composition

To solve a long-horizon task, we can leverage VLM (i.e., ChatGPT-4V) for high-level planning. For each sub-goal  $S$ , the VLM is used to parse the instruction and identify the relevant objects. Given each observation, Grounded SAM segments the point clouds of these objects in visually complex scenes, while VLM infers their semantic relationships. Based on this information, a dynamic sub-scene graph corresponding to the sub-goal is constructed, and its features are encoded using a GNN. As illustrated in Fig. 2 (b), the trained policy then predicts actions conditioned on the sub-scene graph features and the sub-goal description. We assume that all atomic skills required at test time have been learned during training.

## IV. SIMULATION EXPERIMENT

In this section, we evaluate our approach in simulated multi-skill environments. The primary goal of this work is to demonstrate the effectiveness of using scene graphs as inputs for behavior cloning, with a particular emphasis on

skill composition. Our experiments are designed to highlight two main points. First, we show that scene graph-based representations enable *zero-shot generalization* in *compositional scenes*, whereas policies trained on raw 2D or 3D visual inputs suffer from severe distribution shifts. Second, we conduct an ablation study to assess the importance of the scene graph representation in our framework, demonstrating that every component plays a critical role in enabling robust skill composition. Together, these experiments validate the effectiveness and necessity of our design choices for solving long-horizon manipulation tasks.

### A. Task Description

Since most existing simulation benchmarks (e.g., [39], [40]) focus primarily on single-skill tasks, we design five sets of multi-skill, long-horizon tasks built on ManiSkill2 [39], with illustrations provided in Fig. 4. These tasks cover 13 atomic skills spanning a range of manipulation primitives (e.g., pick, place, push, pull). Each multi-skill task is designed to require visual reasoning, generalization, and action understanding in order to compose atomic skills within visually complex scenes. To further increase difficulty, some decisions are conditioned on the scene—for example, whether the robot arm must avoid an obstacle when pulling a cube (see Fig. 5). We also provide motion-planning scripts for all atomic skills, which are used to generate expert demonstrations. For each atomic skill, we collect 100 demonstrations. The tasks are described below:

- **Cube out and in:** The atomic skills are *removing the red cube from the bin* and *putting the blue cube into the bin*. The evaluation scenario begins with a red cube inside the bin and a blue cube outside. The robot is asked to “*put the blue cube into the bin so that only the blue cube remains inside.*”
- **Sort by color:** The atomic skills are to *place an object (apple, banana, or cube) onto a target elliptical place*. In evaluation, three elliptical places with different colors are presented, and the robot is asked to “*put the three objects onto the ellipses of same color.*”
- **Blocks stacking game:** The atomic skills involve logical operations on cubes, inspired by how humans naturally play with them: (a) if two cubes are stacked, push them to the center area together; (b) if the red cube is behind the green cube, lift the red cube before moving it to the center; (c) stack the purple cube on top of the red cube if it is empty on top. In evaluation, all three cubes appear together, and the robot is asked to “*move the red cube to the center while avoiding the green cube, and then stack the purple cube on top of the red cube if there is no other cube on it.*” The policy must understand the logic underlying these atomic skills to successfully compose them.
- **Tools usage:** The atomic skills are *using the L-shaped tool to pull a cube back* and *using the stick to push a cube away*. In evaluation, both tools and two cubes are present, and the robot is asked to “*pull back the blue*

*cube with red tool and push away the green cube with yellow tool*” (or vice versa).

- **Obstacle avoidance:** This is a more challenging variant of tool usage. If pulling a cube straight back would result in a collision, the robot must instead pull it diagonally to avoid obstacles (see Fig. 5). The evaluation task is to “*pull one cube back with red tool and push another cube away with yellow tool, while avoiding the obstacles.*”

### B. Baselines

For baselines, we compare against diffusion policies with different visual representations, including the 2D Diffusion Policy [6] and the 3D Diffusion Policy [9]. We also evaluate against  $\pi_0$  [7], a foundation model trained on large-scale data, as a baseline for large-scale pretraining, and compare its performance with our scene graph-based approach on skill composition tasks. These comparisons motivate our hypothesis that a structured and focused representation of observations is essential for robust skill composition. It is important to note that our work emphasizes *end-to-end* policy learning with scene graph representations. Although we leverage VLMs for high-level planning, they are only used to decompose abstract goals into concrete sub-steps. For a fair comparison, we also allow baselines to do the high-level planning at the beginning.

All methods are trained on the same expert demonstrations (RGB and depth images) and evaluated under the same maximum step limits. Each task is tested with 50 randomized seeds (i.e., initial positions), and we report the average success rate across seeds. Each trial has a maximum success rate of 1.0, defined as the proportion of completed sub-skills.

### C. Results

**Evaluation on atomic skills:** Table I(a) shows that scene graph inputs enable near-perfect success ( $\approx 1.0$ ) across all atomic skills, confirming their effectiveness for behavior cloning and the suitability of GNNs for encoding them. Additionally, the baselines also achieve good performance on single-skill tasks, confirming correct implementation and their ability to learn atomic skills. In contrast, as shown in the skill composition experiments of next part, the baselines fail to generalize when multiple skills must be composed.

**Evaluation on skill composition:** Next, we evaluate the policy on skill composition tasks. As shown in Table I(b), our composing-by-focus method yields only a small performance gap between atomic-skill tasks and compositional tasks. The policy consistently achieves high success rates across all compositional tasks. In contrast, the baseline methods experience substantial degradation: on average, their evaluation scores drop below 50%, and in some cases they fail completely. We want to emphasize a few key findings based on the results:

- 1) **Sensitivity of behavior cloning to visual perturbations.** When additional objects are introduced into the scene or when the background changes, both 2D- and




Fig. 4: **Simulation tasks.** [Left] Visualization of atomic skills in each task, which is the training data. [Right] Evaluation scenarios, involving multiple objects to be operated with possibly changed background.

Algorithm \Task	Cube Out and In	Sort by Color	Blocks Stacking Game	Tools Usage	Obstacle Avoidance
Diffusion Policy	<b>0.98</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>
DP3	0.94	<b>1.0</b>	0.96	0.92	<b>1.0</b>
$\pi_0$	0.7	<b>1.0</b>	0.94	0.26	<b>1.0</b>
Scene Graph	<b>0.98</b>	<b>1.0</b>	0.94	<b>1.0</b>	0.92

(a) Single Skill

Algorithm \Task	Cube Out and In	Sort by Color	Blocks Stacking Game	Tools Usage	Obstacle Avoidance
Diffusion Policy	0.0	0.04	0.47	0.0	0.52
DP3	0.27	0.07	0.48	0.13	0.44
$\pi_0$	0.15	0.02	0.77	0.07	0.49
Scene Graph	<b>0.78</b>	<b>0.79</b>	<b>0.93</b>	<b>0.88</b>	<b>0.90</b>

(b) Skill Composition

TABLE I: **Evaluation results for all simulation tasks.** We evaluate in-domain performance on single-skill tasks (Table (a)), where all methods achieve high scores. On composed tasks (Table (b)), baseline performance drops by 50% – 70% due to limited visual adaptation, while our method consistently outperforms them with higher and more stable success rates.

3D-based policies often generate erroneous behaviors, failing to complete multi-step tasks.

- 2) **Limitations of data scaling for skill composition.** The baseline  $\pi_0$ , despite being pre-trained on large-scale robot manipulation datasets and further fine-tuned on our atomic-skill dataset, performs poorly on our skill composition experiments, indicating that it does not generalize well in these settings.
- 3) **Challenges of domain adaptation.** For behavior cloning baselines, achieving skill composition would require demonstrations covering all permutations of atomic skills—an exponential growth relative to the number of skills, not mentioning the labor for collecting long-horizon demonstrations. In contrast, our approach enables composition directly from policies

trained on atomic skills. The key insight is that VLMs/LLMs already demonstrate strong zero-shot reasoning, and scene graph representations provide a natural interface to integrate these capabilities with diffusion policies.

- 4) **Advantages of scene graph representations.** Scene graphs offer flexibility in the number of nodes and the information contained within them, while explicitly encoding inter-object relationships. This structure helps disambiguate subtle action differences (e.g., obstacle avoidance).

#### D. Ablation Study

We conduct three types of ablation studies to evaluate the importance of our 3D scene graph representation in training




Fig. 5: **Obstacle avoidance.** The sub-scene graph includes the robot gripper, relevant objects, and an *obstacle node* (yellow stick), enabling the policy to learn trajectories that depend on the cube–obstacle relation.

a policy with atomic skills for skill composition. While keeping high-level planning and object-centric representations fixed, we examine the contribution of specific design choices in our visual data processing pipeline. The ablations are performed on the “Sort by Color”, “Blocks Stacking Game”, and “Tools Usage” tasks.

**Importance of 3D representation.** We ablate 3D object representations by cropping RGB images with Grounded SAM and extracting features using a pretrained DINOv2 encoder. Nodes are represented by these features while edges still encode inter-object relations. Trained with the same dataset and architecture, this 2D variant performs significantly worse, as shown in the left column of Fig. 6.

**Importance of graph representation.** A common alternative for 3D object-centric representations is to concatenate point clouds of related objects and encode them with a point cloud network (e.g., DP3 Encoder). While effective for simple single-skill tasks [21], this unstructured approach cannot handle varying object counts and weakens spatial relationships. Consequently, policies trained without scene graphs perform worse, as shown in Fig. 6 (middle).

**Importance of GNN to process the graph.** Processing node features with a GNN captures inter-object relations, handles variable node counts, and is permutation-invariant. In contrast, concatenating node features discards edge information, fixes input size, and is sensitive to object order. An ablation with shuffled inputs (Fig. 6, right) shows our scene graph method remains robust, while the concatenation-based approach fails.

## V. REAL WORLD EXPERIMENT

We evaluate the scene graph representation and skill composition in real-world across two multi-skill tasks, involving 6 atomic skills. We set up two RealSENSE L515 cameras to obtain real-world RGB images and depth images. Using intrinsics and extrinsics of the cameras, we are able to get the point clouds in world frame. We collect expert demonstrations by using a SpaceMouse, with 50 demonstrations for each atomic skill. For fair comparison, we give same number of maximum total steps for each evaluation among all the methods. We test 20 trials for each task. Each trial has a maximum score of 1.0, defined as the proportion of completed sub-skills (i.e., if there are three vegetables to pick and two are successfully picked, the score would be 0.67).

### A. Real World Vegetable Picking

Task	Diffusion Policy	DP3	$\pi 0$	Scene Graph
Single Skill (w. distractors)	0.6	0.7	<b>1.0</b>	<b>1.0</b>
Skill Composition	0.0	0.2	0.05	<b>0.97</b>

TABLE II: **Real-world vegetable picking experiments.** Success rate is defined as the proportion of successfully picked vegetables. Compared with all baselines, our scene graph-based method is more robust to scenario variations and skill composition.

**Single skill evaluation.** The atomic skill is to pick up a vegetable and place it in the basket. Training uses demonstrations of single vegetables (carrot/eggplant/corn) on a clean table, with scene graphs constructed via Grounded SAM [12] containing three nodes: robot hand, vegetable, and basket. The diffusion policy is trained on these graph features and atomic skill descriptions.

At evaluation, we test picking up a single vegetable in cluttered scenes (Fig. 1, row 1). As shown in Table II, our scene graph-based method achieves a perfect success rate (1.0), matching  $\pi 0$  but with far less pretraining. In contrast, 2D and 3D Diffusion Policies perform worse (0.6 and 0.7), highlighting our method’s robustness to distractors and data efficiency.

**Skill composition evaluation.** We then evaluate the policy’s ability to compose multiple atomic skills. In this setting, several vegetables (with possible distractors) are placed on the table (see the second row of Fig. 1), and the goal is described as “Pick up all vegetables and put them in the basket.” To execute this goal, we use ChatGPT to generate substeps (e.g., “1. Pick up the corn and put it in the basket. 2. Pick up the carrot and put it in the basket.”), and each sub-goal is passed sequentially to the policy. For each subtask, a sub-scene graph is constructed containing only the relevant objects, enabling the policy to focus on task-relevant context. An LLM is used to detect when a subtask finishes earlier than the maximum step limit, after which the next subtask begins.

As shown in Table II, our method achieves a success rate of 0.97 on skill composition, far outperforming the baselines: Diffusion Policy, DP3, and  $\pi 0$ . While the baselines can learn atomic skills in isolation, they fail to generalize when multiple skills must be composed in cluttered scenes. These results demonstrate that scene graph-based representations are crucial for enabling robust multi-skill composition.

### B. Real World Tool Usage

Task	Diffusion Policy	DP3	$\pi 0$	Scene Graph
Tool Usage	0.4	0.6	0.075	<b>0.9</b>

TABLE III: **Real-world tool usage experiments.** The score is the average score of 20 trials. Each trial has a maximum score of 1.0, defined as the proportion of completed sub-skills (i.e., pull cube or push cube).

This experiment is the real-world counterpart of the tool-usage task, where an L-shaped stick is used to pull a cube closer or push it away, sometimes requiring obstacle avoidance (Fig. 8). Atomic skills are defined as pulling or




Fig. 6: Ablation studies highlighting the importance of 3D scene graph representations for skill composition.




Fig. 7: Evaluation trajectories for a multi-skill task (pull then push with tools). The baseline (bottom) fails under visual complexity, while our method focuses on task-relevant objects via sub-scene-graphs and successfully composes the two skills. Additional trajectories are shown in the supplementary video.

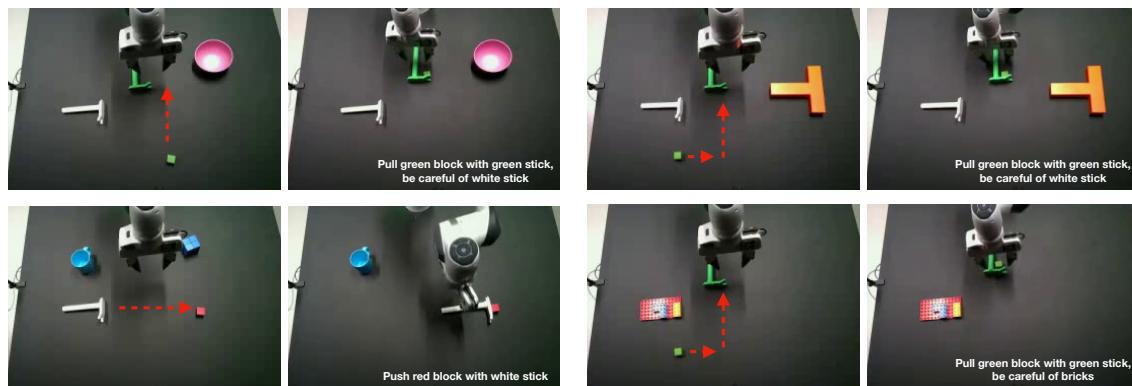


Fig. 8: Various evaluation settings for tool usage. We test the generality for our scene graph-based methods by putting different distractors and obstacles on the table. Our method can give correct trajectories (like how to avoid the different obstacles) without being distracted by unrelated objects on the table.

pushing a cube, with scene graphs containing nodes for ‘robot hand’, ‘stick’, ‘cube’, and optional ‘obstacle’.

For skill composition, two sticks and two cubes are placed on the table with the instruction: “Use the green stick to pull the green cube, be careful of the white stick; then use the white stick to push the red cube.” As in Table III, our method, which focuses on task-relevant objects, achieves a 0.9 success rate, far outperforming baselines. Example trajectories (Fig. 7) show baselines fail under visual variation and composition, while our approach composes skills reliably.

Beyond skill composition, our approach also improves robustness. As shown in Fig. 8, the policy (1) learns appropriate trajectories under different conditions (e.g., pulling straight versus pulling while avoiding an obstacle), and (2) remains resilient to scene variations such as background distractors or

unseen obstacles. For instance, when the obstacle is changed from a stick (seen during training) to bricks (unseen), our method still achieves successful execution.

## VI. CONCLUSIONS

We introduced a scene graph-based representation for robotic skill composition, where nodes capture 3D information of task-relevant objects and edges encode their spatial relationships. This structured input enables visuomotor policies to focus on essential context, improving robustness to visual variation. By supporting flexible composition of atomic skills, our method reduces the exponential data demands of long-horizon demonstrations and improves data efficiency. Moreover, scene graphs provide a natural interface between high-level planning and low-level execution, offering a unified framework for skill composition by focus.

## VII. LIMITATIONS

Our method has two limitations: (1) dynamic scene graph construction with VLMs adds computational overhead, though limited by focusing on a *small* sub-scene graph; (2) reliance on Grounded-SAM [12] may introduce errors due to imperfect segmentation. Improved VLMs and visual foundation models are expected to enhance both efficiency and accuracy.

## REFERENCES

- [1] D. Shah, P. Xu, Y. Lu, T. Xiao, A. Toshev, S. Levine, and B. Ichter, "Value function spaces: Skill-centric state abstractions for long-horizon reasoning," *arXiv preprint arXiv:2111.03189*, 2021.
- [2] W. Mao, W. Zhong, Z. Jiang, D. Fang, Z. Zhang, Z. Lan, H. Li, F. Jia, T. Wang, H. Fan, *et al.*, "Robomatrix: A skill-centric hierarchical framework for scalable robot task planning and execution in open-world," *arXiv preprint arXiv:2412.00171*, 2024.
- [3] A. Mei, G.-N. Zhu, H. Zhang, and Z. Gan, "Replanvlm: Replanning robotic tasks with visual language models," *IEEE Robotics and Automation Letters*, 2024.
- [4] 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, no. 1, pp. 265–293, 2021.
- [5] S. Haresh, D. Dijkman, A. Bhattacharyya, and R. Memisevic, "Clevrskills: Compositional language and visual reasoning in robotics," *Advances in Neural Information Processing Systems*, vol. 37, pp. 38 235–38 266, 2024.
- [6] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [7] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, *et al.*, " $\pi_0$ : A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.
- [8] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [9] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations," *arXiv preprint arXiv:2403.03954*, 2024.
- [10] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki, "3d diffuser actor: Policy diffusion with 3d scene representations," *arXiv preprint arXiv:2402.10885*, 2024.
- [11] 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.
- [12] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, *et al.*, "Grounded sam: Assembling open-world models for diverse visual tasks," *arXiv preprint arXiv:2401.14159*, 2024.
- [13] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [14] Y. Du, S. Li, A. Torralba, J. B. Tenenbaum, and I. Mordatch, "Improving factuality and reasoning in language models through multiagent debate," in *Forty-first International Conference on Machine Learning*, 2023.
- [15] V. Subramaniam, Y. Du, J. B. Tenenbaum, A. Torralba, S. Li, and I. Mordatch, "Multiagent finetuning: Self improvement with diverse reasoning chains," *arXiv preprint arXiv:2501.05707*, 2025.
- [16] J. Mao, T. Lozano-Pérez, J. B. Tenenbaum, and L. P. Kaelbling, "Learning reusable manipulation strategies," in *Conference on Robot Learning*. PMLR, 2023, pp. 1467–1483.
- [17] Y. Zhu, P. Stone, and Y. Zhu, "Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4126–4133, 2022.
- [18] 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, no. 1, pp. 265–293, 2021.
- [19] T. Silver, R. Chitnis, J. Tenenbaum, L. P. Kaelbling, and T. Lozano-Pérez, "Learning symbolic operators for task and motion planning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3182–3189.
- [20] A. Curtis, T. Silver, J. B. Tenenbaum, T. Lozano-Pérez, and L. Kaelbling, "Discovering state and action abstractions for generalized task and motion planning," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 36, no. 5, 2022, pp. 5377–5384.
- [21] Y. Zhu, Z. Jiang, P. Stone, and Y. Zhu, "Learning generalizable manipulation policies with object-centric 3d representations," *arXiv preprint arXiv:2310.14386*, 2023.
- [22] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, "Viola: Imitation learning for vision-based manipulation with object proposal priors," in *Conference on Robot Learning*. PMLR, 2023, pp. 1199–1210.
- [23] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.
- [24] H. K. Cheng and A. G. Schwing, "Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model," in *European conference on computer vision*. Springer, 2022, pp. 640–658.
- [25] S. Haldar, J. Pari, A. Rai, and L. Pinto, "Teach a robot to fish: Versatile imitation from one minute of demonstrations," *arXiv preprint arXiv:2303.01497*, 2023.
- [26] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, "Mimicplay: Long-horizon imitation learning by watching human play," *arXiv preprint arXiv:2302.12422*, 2023.
- [27] H. Ha, P. Florence, and S. Song, "Scaling up and distilling down: Language-guided robot skill acquisition," in *Conference on Robot Learning*. PMLR, 2023, pp. 3766–3777.
- [28] H. Qi, H. Yin, and H. Yang, "Control-oriented clustering of visual latent representation," in *International Conference on Learning Representations (ICLR)*, 2025.
- [29] D. Niu, Y. Sharma, H. Xue, G. Biamby, J. Zhang, Z. Ji, T. Darrell, and R. Herzig, "Pre-training auto-regressive robotic models with 4d representations," *arXiv preprint arXiv:2502.13142*, 2025.
- [30] S. Batra and G. Sukhatme, "Zero-shot visual generalization in robot manipulation," *arXiv preprint arXiv:2505.11719*, 2025.
- [31] K. Hsu, J. I. Hamid, K. Burns, C. Finn, and J. Wu, "Tripod: Three complementary inductive biases for disentangled representation learning," *arXiv preprint arXiv:2404.10282*, 2024.
- [32] K. P. Singh, J. Salvador, L. Weihs, and A. Kembhavi, "Scene graph contrastive learning for embodied navigation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 10 884–10 894.
- [33] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, "Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation," in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [34] U. A. Mishra, S. Xue, Y. Chen, and D. Xu, "Generative skill chaining: Long-horizon skill planning with diffusion models," in *Conference on Robot Learning*. PMLR, 2023, pp. 2905–2925.
- [35] I. Mishani, Y. Shaoul, and M. Likhachev, "Mosaic: A skill-centric algorithmic framework for long-horizon manipulation planning," *arXiv preprint arXiv:2504.16738*, 2025.
- [36] C. Moenning and N. A. Dodgson, "Fast marching farthest point sampling," University of Cambridge, Computer Laboratory, Tech. Rep., 2003.
- [37] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *arXiv preprint arXiv:1710.10903*, 2017.
- [38] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PmlR, 2021, pp. 8748–8763.
- [39] J. Gu, F. Xiang, X. Li, Z. Ling, X. Liu, T. Mu, Y. Tang, S. Tao, X. Wei, Y. Yao, *et al.*, "Maniskill2: A unified benchmark for generalizable manipulation skills," *arXiv preprint arXiv:2302.04659*, 2023.
- [40] NVIDIA, "Isaac Sim." [Online]. Available: <https://github.com/isaac-sim/IsaacSim>