

MoManipVLA: Transferring Vision-language-action Models for General Mobile Manipulation

Zhenyu Wu¹ Yuheng Zhou² Xiuwei Xu³ Ziwei Wang² Haibin Yan^{1*}

¹Beijing University of Posts and Telecommunications, School of Intelligent Engineering and Automation

²Nanyang Technological University, School of Electrical and Electronic Engineering

³Tsinghua University, Department of Automation

{wuzhenyu, eyanhaibin}@bupt.edu.cn zhou0484@e.ntu.edu.sg

xxw21@mails.tsinghua.edu.cn ziwei.wang@ntu.edu.sg

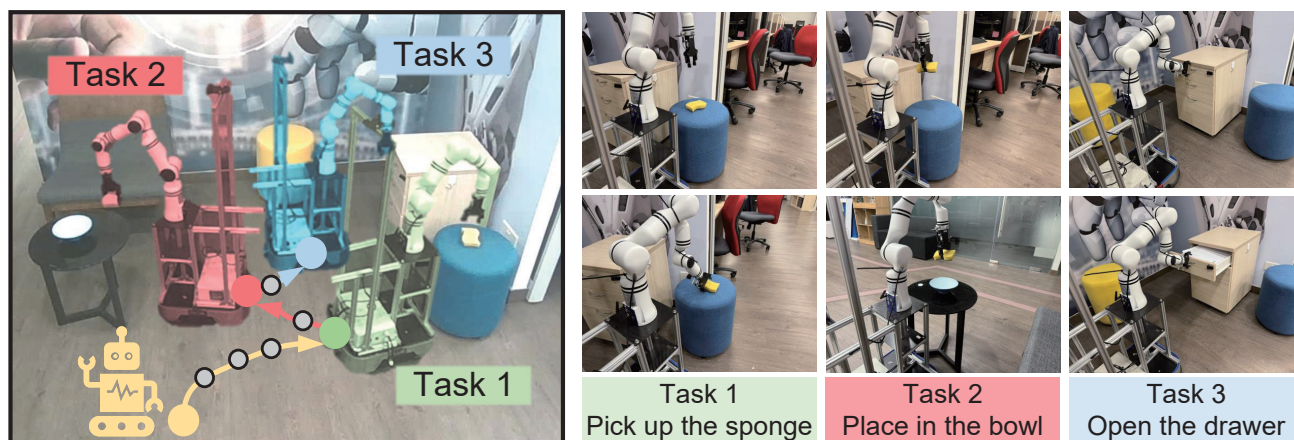


Figure 1. Transferring pre-trained VLA models to mobile manipulation significantly enhances the generalization ability of the policy across tasks and environments. Our MoManipVLA can complete diverse household tasks such as object picking, object delivery and drawer opening in large working area. Guided by the waypoints from VLA models, the motions of the mobile base and the robot arm are jointly generated with physical feasibility constraints.

Abstract

Mobile manipulation is the fundamental challenge for robotics to assist humans with diverse tasks and environments in everyday life. However, conventional mobile manipulation approaches often struggle to generalize across different tasks and environments because of the lack of large-scale training. In contrast, recent advances in vision-language-action (VLA) models have shown impressive generalization capabilities, but these foundation models are developed for fixed-base manipulation tasks. Therefore, we propose an efficient policy adaptation framework named MoManipVLA to transfer pre-trained VLA models of fixed-base manipulation to mobile manipulation, so that high generalization ability across tasks and environments can be achieved in mobile manipulation policy. Specifically,

we utilize pre-trained VLA models to generate waypoints of the end-effector with high generalization ability. We design motion planning objectives for the mobile base and the robot arm, which aim at maximizing the physical feasibility of the trajectory. Finally, we present an efficient bi-level objective optimization framework for trajectory generation, where the upper-level optimization predicts waypoints for base movement to enhance the manipulator policy space, and the lower-level optimization selects the optimal end-effector trajectory to complete the manipulation task. Extensive experimental results on OVMM and the real world demonstrate that MoManipVLA achieves a 4.2% higher success rate than the state-of-the-art mobile manipulation, and only requires 50 training cost for real world deployment due to the strong generalization ability in the pre-trained VLA models. Our project page can be found [here](#).

* Corresponding author

1. Introduction

Mobile manipulation empowers robotics with the ability to perform complex manipulation tasks across large spaces, which requires the whole-body control for the mobile base and the arm [37, 39]. With the increasing popularity of intelligent robotic systems, several fields such as household service [36], manufacturing [28], and logistics [31] are in urgent requirement of mobile manipulation capabilities since robotics are needed to perform cross-space manipulations autonomously. However, the requirement to perform diverse tasks in unstructured environments (e.g., assist humans in their daily lives) presents significant challenges.

Conventional mobile manipulation frameworks separately train navigation and fixed-base manipulation modules, leading to compounding errors [5, 7, 9, 41, 43]. Recent end-to-end methods jointly optimize navigation and manipulation actions, but high collection costs for demonstrations limit dataset scale and generalization [16, 42]. Meanwhile, although VLA models show strong generalization in diverse manipulation tasks [19, 21], their focus on fixed-base tasks prevents them from generating cooperative actions between the mobile base and robot arm for mobile manipulation.

In this paper, we propose an efficient policy transfer framework named MoManipVLA to generalize the fixed-base VLA models to mobile manipulation tasks. Unlike existing mobile manipulation methods which suffer from low generalization ability across tasks and environments, MoManipVLA efficiently transfers generalizable fix-base manipulation policy from pre-trained VLA models to mobile manipulation. More specifically, we employ pre-trained VLA models to predict end-effector waypoints with a high generalization ability to guide the generation of mobile manipulation trajectories. We then design motion planning objectives of scene constraints for the mobile base and the robot arm including end-effector reachability, trajectory smoothness and collision avoidance, which aim to maximize the physical feasibility of the trajectories. To efficiently plan the whole-body motion, we propose a bi-level trajectory optimization framework of the objectives, where the upper-level optimization predicts the waypoints for base movement to strengthen the manipulator policy space, and the lower-level optimization selects the optimal end-effector trajectory for task completion. Figure 1 illustrates the comparison between the conventional and the proposed approaches, MoManipVLA efficiently transfers the pre-trained VLA models to diverse mobile manipulation tasks, where the mobile base and the robot arm coordinately perform actions with physically feasible trajectories. Extensive experimental results on OVMM [40] and real world demonstrate that our method achieves a 4.2% higher success rate than state-of-the-art mobile manipulation techniques. Due to the strong generalization of our pre-trained VLA models, MoManipVLA only requires 50

expert episodes for real-world deployment. Our contributions include:

- We propose a policy adaptation framework that transfers VLA models for mobile manipulation, enabling high generalization across tasks and environments.
- We introduce motion planning objectives and a bi-level trajectory optimization framework to enhance the physical feasibility and efficiency of trajectories.
- Extensive experiments in OVMM and the real world validate the method’s generalization and efficiency.

2. Related Work

Mobile Manipulation Framework: Mobile manipulation requires agents to have the ability to interact with objects in large spaces based on human instructions. Existing mobile manipulation frameworks can be categorized into two types: end-to-end and modular. End-to-end approaches [2, 8, 10] employ imitation learning to directly predict mobile manipulation actions based on visual observations. Qiu *et al.* [29] propose to learn a unified scene information representation for navigation and manipulation, which leverages both geometric and semantic information to improve the success rate of the manipulation. Yan *et al.* [38] utilized diffusion policy to generate mobile manipulation whole-body control trajectories, which control task-specific diffusion policy generation through diverse energy terms. However, imitation learning with expert trajectories leads to expensive training costs [24]. Therefore, modular mobile manipulation frameworks Home-Robot [40] and OK-Robot [22], which contain foundation model planners and reinforcement learning-based controllers are proposed to efficiently handle long-horizon mobile manipulation tasks. SPIN [34] proposes to utilize a reactive mobile manipulation framework to achieve proactive scene perception which follows the full-body and hand-eye coordination capabilities of humans. However, existing mobile manipulation methods struggle to generalize to diverse real-world tasks and environments due to insufficient large-scale pre-training.

Vision-language-action Models: The current state-of-the-art VLA models [12, 19–21] directly output end-effector 7-DoF actions based on RGB vision observations without relying on predicted object categories and poses. Early VLA frameworks following the vision-language model(VLM) [33] architecture represented actions with autoregressive discretization. ManipLLM [21] constructed the Chain-of-Thought(CoT) to stimulate manipulation reasoning capabilities in the foundation model and further introduced an active impedance adaptive strategy to plan the next waypoints. OpenVLA [19] explored the impact of visual encoders on the performance of the VLA models, which combines different foundation model components to achieve satisfactory performance. To fur-

ther enhance the VLA model to mine the association between visual input and action trajectories, TinyVLA [35] proposes to utilize the foundation model feature prior to guide the action decoder diffusion process, which significantly enhances the generalization of the VLA model on changes in viewpoints, objects, etc. Recent works also explore the ability of VLA models on embodiments with high degrees of freedom. RDT-1B [23] generalized diffusion policy to a bimanual manipulation by introducing a physically interpretable unified action space to handling more complex tasks. $\pi 0$ [1] proposed a novel flow-matching architecture based on VLM to learn internet-scale knowledge, which empowered cross-embodied manipulation planning for VLA models. Although existing VLA models demonstrate impressive generalization across tasks and environments, they fail to generate the cooperative actions of the mobile base and the robot arm.

Trajectory Optimization: In robotic manipulation tasks, trajectory optimization plays a pivotal role in achieving efficient and precise robot actions. Early research has focused on the multi-objective and constraint characteristics of trajectory optimization [3, 11, 17], aiming to address the demands for accuracy and low latency within dynamic environments. Handcrafted optimization objectives confine conventional methods to specific tasks, hindering their deployment in dynamic and complex environments. Consequently, data-driven approaches are now used to tailor trajectory optimization to the dynamic challenges of real-world scenes. Contact-GraspNet [32] and O2O-Afford [26] guide trajectory generation in densely cluttered scenes by predicting grasp poses and object affordances, respectively. Michael *et al.* [6] developed a collision model that predicts 6DOF object constraints, effectively tackling challenges associated with occlusions in point cloud data. To leverage the knowledge embedded in the demo episodes, some works employ imitation learning to mine trajectory policies directly from expert demonstrations. Chi *et al.* [4] and Huang *et al.* [13] train neural networks to generate collision-free trajectories through imitation learning from human or expert demonstrations, eliminating the need for manually defined constraints or explicit environmental modeling. Benefit from the development of foundation models, recent methods such as VoxPoser [14] and ReKep [15] have demonstrated that leveraging foundational models to infer scene affordances and constraints can significantly enhance the generalization of trajectory optimization. Inspired by ReKep’s success in manipulation tasks, we developed a motion planning objective that efficiently adapts pre-trained fixed-base VLA models for mobile manipulation.

3. Approach

In this section, we first introduce the overall framework of policy transfer from the fixed-base VLA models to mo-

bile manipulation tasks. Then we present the motion planning objectives for evaluation of the base and the arm trajectory to maximize the physical feasibility. Finally, we propose an efficient bi-level objective optimization framework for trajectory generation.

3.1. Problem Statement

The mobile manipulation task requires the robot to interact with objects in different locations, cooperatively controlling the mobile base and the robot arm to accomplish the interaction actions in physically feasible trajectories. The observation at the t_{th} step consists of the visual input and the proprioception state. The visual input contains RGB images, depth images, and the camera poses, and the proprioception state includes locations p_b^t and rotations r_b^t for the base in the world frame, and locations p_e^t , rotations r_e^t , and gripper openness g_e^t of the end-effector in the base frame. We apply the transformation function Γ to obtain the location \hat{p}_e^t and rotation \hat{r}_e^t of the end-effector in the world frame. Based on observation and language instructions, the agent is required to generate the optimal action to achieve the next-best proprioception state of the mobile base and the robot arm for task completion.

Existing mobile manipulation frameworks usually learn the mobile manipulation policy by imitating expert demonstrations. The scale of expert demonstrations is limited because of the high cost of data collection in mobile manipulators with high DoFs, which leads to low generalization ability across tasks and environments. Our goal is to efficiently transfer pre-trained VLA model policies to mobile manipulation, which jointly generates trajectories of both the mobile base and the robot arm to achieve high generalization ability.

3.2. Policy Transfer Framework

VLA models pre-trained on web-scale data [27] demonstrate impressive generalization capabilities across tasks and environments. Nevertheless, existing efforts are limited to fixed-base manipulation and cannot deal with mobile manipulation that requires whole-body policy generation for both the base and the arm. To this end, we propose to transfer the fixed-base manipulation policy predicted by pre-trained VLA models by whole-body motion planning with physical feasibility. Figure 2 demonstrates the overall pipeline of MoManipVLA. We first employ pre-trained VLA models to generate the optimal end-effector waypoints based on observation and human instruction. To enable the end-effector to achieve the goal waypoints, we jointly generate the trajectories for the base and the arm with the highest physical feasibility. We design motion planning objectives to maximize the physical feasibility including safety, smoothness and reachability. To efficiently solve the trajectory, we propose a bi-level objective optimization frame-

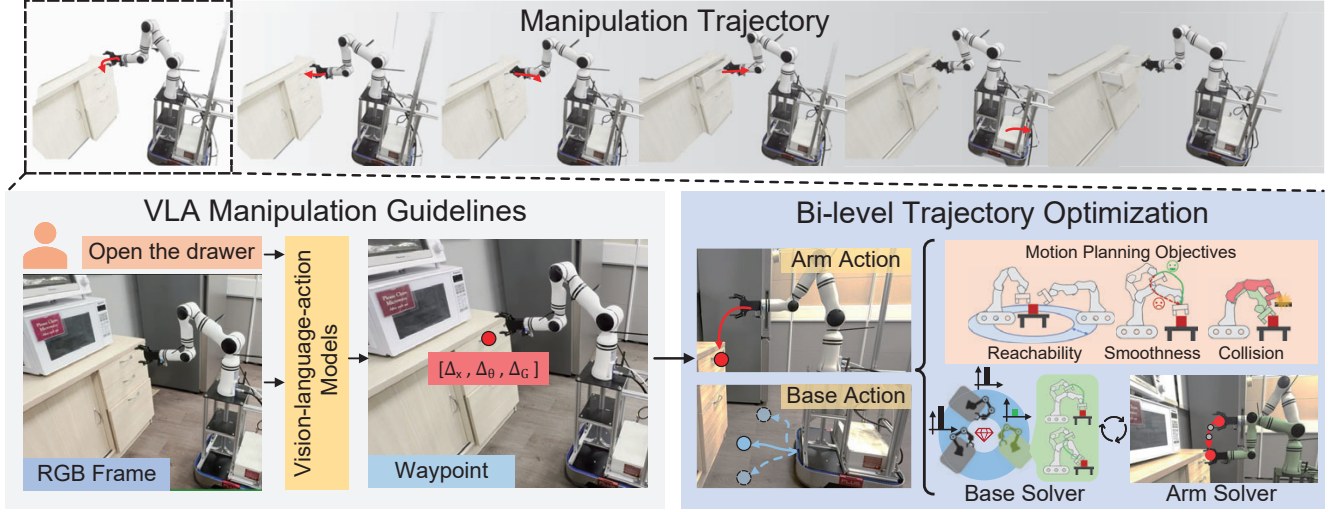


Figure 2. The pipeline of MoManipVLA. The pre-trained VLA models predict highly generalized end-effector waypoints to guide the mobile manipulation task, through which the trajectory of the mobile base and the robot arm can be generated with objectives of physical feasibility. The objectives consider the reachability, smoothness and collision, and the trajectory is acquired via bi-level optimization.

work, where the upper-level optimization predicts the base movement waypoints to enhance the subsequent manipulation policy space, and the lower-level optimization selects the optimal end-effector trajectory to achieve the waypoints generated by VLA models. Therefore, the motion planning framework enables pre-trained fixed-base policies to be efficiently adapted to mobile manipulation with negligible training cost.

3.3. Mobile Motion Planning Objectives

Motion planning aims to generate feasible trajectories for the mobile base and the robot arm between given the consecutive waypoints of end-effectors, which should achieve pose reachability, trajectory smoothness, and safety during the entire mobile manipulation process. To enhance the generalization ability across tasks and environments, the i_{th} waypoint of the end-effector \mathbf{Q}_i in the base frame during the entire mobile manipulation process can be generated via a pre-trained fixed-base VLA model, where RGB images of the current scene and the proprioception state of the robot arm are considered following the VLA pipeline. For trajectory planning, we first transform the waypoint \mathbf{Q}_i to that the world frame denoted as $\hat{\mathbf{Q}}_i$ in the following:

$$\hat{\mathbf{Q}}_i = \Gamma(\mathbf{P}_i^0, \mathbf{Q}_i) \quad (1)$$

where \mathbf{P}_i^0 demonstrates the proprioception state of the base when predicting the waypoint \mathbf{Q}_i with the VLA model. Therefore, the optimization objective can be represented as follows to generate the trajectory given the consecutive waypoints $\hat{\mathbf{Q}}_i$ and $\hat{\mathbf{Q}}_{i+1}$:

$$\begin{aligned} \min \sum_{t=0}^{T_i} \mathcal{O}(\mathbf{x}_{i,b}^t, \mathbf{x}_{i,e}^t) \\ \text{s.t. } \text{quad}\Gamma(\mathbf{x}_{i,b}^0, \mathbf{x}_{i,e}^0) = \hat{\mathbf{Q}}_i, \quad \Gamma(\mathbf{x}_{i,b}^{T_i}, \mathbf{x}_{i,e}^{T_i}) = \hat{\mathbf{Q}}_{i+1} \end{aligned} \quad (2)$$

where T_i denotes the number of steps in the planned trajectory between above consecutive waypoints, and \mathcal{O} is the objective function evaluating the physical feasibility of trajectories. For trajectory planning, $\mathbf{x}_{i,b}^t$ means the planned base pose in the world frame at the t_{th} step including location $p_{i,b}^t$ and rotation $r_{i,b}^t$, and $\mathbf{x}_{i,e}^t$ demonstrates the planned end-effector pose in the base frame at the t_{th} step including location $p_{i,e}^t$, rotation $r_{i,e}^t$ and gripper openness $g_{i,e}^t$. Our goal is to minimize the cost function to enhance the physical feasibility of the entire trajectory with waypoints from VLA models, and the constraints indicate that the starting pose and the ending pose should be consistent with the waypoints for trajectory smoothness. Given the poses of the base and the arm in the trajectory, the joint angles of the arm are iteratively solved by Pinocchio IK solver, and the base is driven with translation and rotations. The objective functions are introduced in the following.

Reachability cost: Since the mobile manipulation requires the robot to interact with objects in large area, improper pose of the mobile base may lead to an unreachable pose of the arm. Therefore, we evaluate the reachability of each candidate trajectory. Assuming the solution of joint angles can be acquired via N_{IK} iterations, the reachability cost \mathcal{F}_r can be represented as follows:

$$\mathcal{F}_r = \begin{cases} N_{IK}/N_{max} & \text{if } N_{IK} \leq N_{max}, \\ C_0 & \text{if } N_{IK} > N_{max}, \end{cases} \quad (3)$$

where N_{max} is a hyperparameter representing the maxi-

imum iteration time for the IK solver, and C_0 is a large constant. Acquiring the joint angle solution within the maximum iterations indicates reachable trajectories. Slower IK solving with more iterations means that the joint angle is closer to the range limit, which is more vulnerable to noise in the execution the mobile manipulation actions. This shows lower reachability for the robot arm. When the iteration number is larger than the budget, the pose in the candidate trajectory is unreachable without feasible joint angle solutions. Therefore, the cost will be assigned to an extremely high value.

Smoothness cost: The smoothness constraint keeps continuous and smooth changes for joint angles of the robot arm and translation and rotation of the base, where sudden changes are avoided to keep the safety of the motor and the embodiment. We define trajectory smoothness as the difference of joint angles θ^t of the arm and proprioception of the base between consecutive poses in the candidate trajectory:

$$\mathcal{F}_s = \sum_{t=0}^{T_i} \|\theta^{t+1} - \theta^t\|_2 + \|\mathbf{x}_b^{t+1} - \mathbf{x}_b^t\|_2 \quad (4)$$

where subscript i for the base pose in the trajectory is omitted for simplicity. We leverage the joint angle solved via IK for smoothness constraint instead of the proprioception of the arm because small changes in the arm pose cannot guarantee slight differences in joint angles.

Collision cost: The robot is required to avoid any collision among the robot arm, the mobile base and objects in the environment to keep safety during the mobile manipulation process. We leverage nvblox [25] to compute the ESDF for object surface in the environment according to the RGB-D images, and randomly sample N_q query points on the robot surface to evaluate the collision cost:

$$\mathcal{F}_c = \sum_{t=0}^T \sum_{j=1}^{N_q} \max(0, \epsilon_0 - \mathcal{D}(q_j^t, \Omega)) \quad (5)$$

where q_j^t means the j_{th} query point at the t_{th} pose in the trajectory, and ϵ_0 is a hyperparameter that controls the safety margin for collision avoidance. $\mathcal{D}(q_j^t, \Omega)$ means the distance between q_j^t and the surface Ω , which is predicted by the ESDF model. We expect the distance between the robot and the object surface can be maximized if it is within the safety margin. Otherwise, the trajectory will be regarded as collision-free with no contribution to the cost.

We acquire the overall objectives by combining the reachability, smoothness and collision costs with the hyperparameter $\{\lambda_i\}_i$:

$$\mathcal{O} = \lambda_1 \mathcal{F}_r + \lambda_2 \mathcal{F}_s + \lambda_3 \mathcal{F}_c \quad (6)$$

By generating physically feasible trajectories, the agent can fully leverage the generalization ability across tasks and en-

Algorithm 1: Bi-Level Trajectory Optimization

Require: Solution tolerance: μ_s

Require: Upper and lower iterations: N_{max}^{up} , N_{min}^{low}

- 1: Get i_{th} waypoint $\hat{\mathbf{Q}}_i \leftarrow \Gamma(x_b^0, \mathbf{Q}_i)$
 - 2: Get initial solution $\{x_e^0, x_b^0\}$ by linear interpolation
 - 3: Initialize current stage step $t \leftarrow 0$
 - 4: **repeat**
 - 5: Update robotic state $\{p_b^t, r_b^t\}, \{p_e^t, r_e^t\} \leftarrow x_b^t, x_e^t$
 - 6: **for** $k = 1$ to N_{max}^{up} **do** ▷ Update Base
 - 7: Transform end-effector position $\hat{p}_e^t \leftarrow \Gamma(p_b^t, p_e^t)$
 - 8: Sampling end-effector trajectories L_s via \hat{p}_e^t
 - 9: Update robotic base x_b^{t+1} via Ep.6 and L_s
 - 10: Update current stage step $t \leftarrow t + 1$
 - 11: **end for**
 - 12: **for** $k = 1$ to N_{max}^{low} **do** ▷ Update End Effector
 - 13: Transform end-effector position $\hat{p}_e^t \leftarrow \Gamma(p_b^t, p_e^t)$
 - 14: Update robotic base position p_e^{t+1} via Ep.6
 - 15: Update robotic end-effector $x_b^{t+1} \leftarrow x_b^t$
 - 16: Update current stage step $t \leftarrow t + 1$
 - 17: **end for**
 - 18: **until** $\|\Gamma(p_b^t, p_e^t) - \hat{\mathbf{Q}}_i\|_2 < \mu_s$
-

vironments in VLA models for the challenging mobile manipulation task.

3.4. Bi-Level Trajectory Optimization

Directly searching the optimal solution to the objective (6) is highly complex, because the pose search space of the mobile base and the robot arm is extremely large with high DoFs. To address this problem, we propose a bi-level trajectory optimization framework to improve the efficiency of trajectory generation, where the upper level optimizes the waypoints for the base to enhance the manipulator policy space, while the lower level further optimizes the waypoints for the end-effector to follow the guidance of the pre-trained VLA models for completing manipulation tasks. Although the bi-level optimization framework is greedy, but mobile manipulation requires searching a 10-DoF space (7 for the arm and 3 for the base), resulting in a large search space that cannot be proven convex, where direct searching may easily lead to local optima. Notably, decomposed greedy search is commonly used in other robotics fields, such as humanoid control [44], which splits the 39-DoF humanoid actions into separate 27-DoF upper and 12-DoF lower body actions. We employ Dual Annealing search algorithms for the objective optimization, where gradient-based local optimizer SLSQP is used to refine the solution. The details of bi-level trajectory optimization are demonstrated in Algorithm 1.

Initialization: According to the constraints in the objective, the first and the last poses of the trajectory are set to the current waypoint and the next one predicted by VLA models. The intermediate poses in the trajectory are initialized

as the interpolation states with equal intervals between the consecutive waypoints.

Upper-level optimization: In the upper-level, we search the optimal base poses to form the search space of arm poses. In the iterative updating stage of the Dual Annealing search algorithms, we only update the base pose to generate new candidates of whole-body pose in the trajectory. For the j th pose in the trajectory, we randomly sample different arm poses given the base pose, and evaluate the quality of the search space with the expected objective and top-tier objective:

$$J_{up} = \sum_{x_e \in \mathcal{M}} \mathcal{O}(x_b^j, x_e) + \alpha \sum_{x_e \in \mathcal{N}} \mathcal{O}(x_b^j, x_e) \quad (7)$$

where x_e means the candidate arm pose. \mathcal{M} and \mathcal{N} respectively mean the set of sampled candidates and those with the top-k lowest objectives. The search space with low expected objectives across candidates indicates high quality, and lower top-tier objective means the higher upper bound for the candidate performance.

Lower-level optimization: After the base poses in the trajectory achieve optimal, the low-level optimization searches for the best arm poses in the trajectory. Given the search space, the quality of each candidate can be estimated by the objective (6). We leverage the optimal candidate in the given search space as the whole-body pose in the trajectory. With efficient policy transfer, our agent achieves high generalization ability across tasks and environments with guidance from pre-trained VLA models.

4. Experiment

In this section, we conduct comprehensive experiments in OVMM and the real world to demonstrate the effectiveness of the proposed method. We first describe the implementation details and compare them with state-of-the-art methods. Then we validate the effectiveness of each cost term in the motion planning objective through ablation experiments. Finally, we test the method performance in the real world to verify generalization.

4.1. Datasets and Implementation Details

The Open Vocabulary Mobile Manipulation (OVMM) benchmark [40] contains 60 scene models approximating the layout of a real house, along with over 18k 3D models of everyday objects. The mobile manipulation task in the OVMM benchmark is formally defined as ‘‘Move a target object from container A to container B’’, where the target object is a small object that can be grasped by Hello Robot Stretch [18]. The robot is initialized in an unknown environment and needs to perform ‘‘Nav to Recp-A, Gaze, Pick Object, Nav to Recp-B, and Place’’ stages sequentially to complete the mobile manipulation, where errors in any of these

stages will result in a failed manipulation. We collect mobile manipulation expert trajectories with the heuristic baseline provided by OVMM to fine-tune the off-the-shelf VLA models to bridge the cross-embodied gaps. We employ OpenVLA-7B to generate fine-grained interaction trajectories. Specifically, each expert trajectory consists of a series of tuples containing visual perception, robotic states, and execution actions. We collect 200 pick-and-place demonstration episodes and efficiently fine-tune 10K epochs with LORA on 4 RTX 3090 GPUs. For the trajectory optimization framework, we employ double annealing to search for physically feasible trajectories between waypoints. The number of intermediate steps is calculated based on the position and rotation changes between waypoints with the specified step size (0.05). The scene ESDF is constructed with 3D object models. We follow the OVMM benchmark setting and utilize the open vocabulary instance segmentation model as the visual perception module. The cost hyperparameters λ_1 , λ_2 and λ_3 in the optimization objective are set to 10.0, 1.0 and 0.6, respectively. The safety threshold ϵ_0 in the collision cost is set to 0.1.

4.2. Comparison with State-of-the-art Methods

Because our proposed MoManipVLA relies on the VLA models for waypoint guidance, we generate fine-grained base and arm trajectories for the OVMM task exclusively at the end of the agents’ navigation phase within MoManipVLA. Table 1 demonstrates the performance of MoManipVLA on OVMM compared to state-of-the-art methods. Benefiting from the generalization ability of the pre-trained VLA models, our method achieves 4.2% overall success rate and 11.2% partial success rate gain, respectively, which efficiently verify that the proposed approach efficiently transfers the pre-trained VLA model policies to the mobile manipulation task. We demonstrate that our method can coordinate the motion of the robotic base and arm in order to keep the end-effector in a reasonable spatial relationship with the target object. Table 2 demonstrates the efficiency of our approach compared to the baseline, Latency denotes the inference latency for generating the interaction policy. Notably, the step numbers of our method are close to the heuristic method while the latency is close to the RL-based method, further demonstrating the efficiency of the proposed policy transfer and hierarchical optimization framework. We further counted the failure cases of mobile manipulation as illustrated in Table 3. ‘‘Find_recep’’ means that the grasp receptacle is not found, ‘‘Nav_to_place’’ means that the placement receptacle is not found, and ‘‘Orient_to_place’’ means that the robot’s navigation strategy fails (it is not aligned to the target), further validating the importance of the base waypoints in mobile manipulation.

Table 1. Comparison results on the OVMM benchmark, where partial success rates indicate the execution of each stage. We follow the OVMM setting employing the baseline Nav and Gaze model to find the target object, replacing the Pick and Place policy to validate the effectiveness of the proposed method on mobile manipulation.

| Method | Partial Success Rates | | | Overall SR | Partial SR | Step |
|------------------|-----------------------|-------|---------|------------|------------|--------|
| | FindObj | Pick | FindRex | | | |
| UniTeam | 49.2% | 42.8% | 19.6% | 9.2% | 30.8% | 1006.6 |
| OVMM (RL) | 32.4% | 15.6% | 9.2% | 1.2% | 14.6% | 1132.5 |
| OVMM (Heuristic) | 30.8% | 14.4% | 3.6% | 0.8% | 12.4% | 1009.8 |
| RoboAI | 41.2% | 21.2% | 6.4% | 0.0% | 17.2% | 906.2 |
| KUZHUM | 55.7% | 50.2% | 35.2% | 11.6% | 38.2% | 1153.3 |
| MoManipVLA | 66.1% | 62.6% | 53.1% | 15.8% | 49.4% | 1240.5 |

Table 2. Ablation experiments for optimization cost terms, where the base method denotes the full setting of our approach. We verify its effectiveness by gradually eliminating the cost term. We further explored the effectiveness of bi-level optimization.

| Method | Search Policy | | Partial Success Rates | | | Overall SR | Partial SR | Step | Latency |
|---------------------|---------------|--------|-----------------------|-------|---------|------------|------------|--------|---------|
| | Bi-level | Direct | FindObj | Pick | FindRex | | | | |
| Base Method | ✓ | | 66.1% | 62.6% | 53.1% | 15.8% | 49.4% | 1240.5 | 693.1 |
| Base Method | | ✓ | 65.9% | 61.1% | 52.5% | 14.1% | 48.5% | 1373.3 | 742.9 |
| Base Method w/o GT. | ✓ | | 23.7% | 12.7% | 7.1% | 1.7% | 11.3% | 1187.8 | 737.6 |
| w/o Reachability | ✓ | | 65.8% | 61.3% | 52.6% | 13.1% | 48.2% | 1174.7 | 682.7 |
| w/o Smoothness | ✓ | | 66.1% | 61.7% | 53.1% | 13.8% | 48.7% | 1185.3 | 688.2 |
| w/o Collision | ✓ | | 66.3% | 61.4% | 52.7% | 15.3% | 48.9% | 1174.2 | 692.1 |

4.3. Ablation Study

To fully verify the importance of each cost term in trajectory optimization, we conduct ablation experiments on the baseline method. We respectively remove the terms of reachability, smoothness, and collision and evaluate the success rate and efficiency. The results of the ablation experiments are illustrated in Table 2, where “Base Method” denotes our proposal method, and “w/o GT” means using the instance segmentation. Each term improves the overall success rate with the obvious contribution. Among all three terms, reachability increases the success rate most significantly because mobile manipulation requires target interaction in large regions, where it is challenging for the agent to reach the target due to the base location. In the mobile manipulation task, the cooperative control of the mobile base and the arm to ensure end-effector reachability is a key bottleneck in performance, as well as the transfer of the pre-trained VLA model to mobile manipulation.

Meanwhile, we also leverage different pose search methods for trajectory generation. Besides the bi-level objective optimization in our method, we also use the ordinary Dual Annealing search algorithm without search space decomposition. Despite a higher success rate, our bi-level objective optimization method also reduces the latency significantly. This indicates the effectiveness of our algorithm to avoid search deficiency. We further leverage the groundtruth ob-

Table 3. Failure cases in OVMM simulation environment

| Find_recep | Nav_to_place | Orient_to_place | Other |
|------------|--------------|-----------------|--------|
| 14.62% | 0.33% | 72.09% | 12.96% |

ject masks as the visual input for our system to fairly evaluate our mobile manipulation policy. We also evaluate our method by utilizing vision foundation models such as Detic for object mask generation. As the objects are usually highly cluttered in the household environments, the quality of visual perception degrades significantly. Therefore, the overall success rate is far below methods using groundtruth masks.

4.4. Real World Experiments

Table 4 demonstrates the performance of our approach on real-world robotics. We employ the hexman echo plus base and RM65 robotic arm components offline mobile platforms and utilize nvblox to reconstruct the scene ESDF and Grounding SAM [30] to obtain masks for the robotic arm and target object, using the arm mask as the query point and excluding the target object mask in ESDF construction. We follow the ORB-SLAM setting and use the Intel Realsense T265 tracking camera to acquire the real-time camera pose and base pose. Benefiting from the generalization of the pre-trained VLA models, our approach requires only




Figure 3. Real-world mobile manipulation visualization results. We demonstrate a real-world mobile robotic Pick-and-Place action sequence, which generates mobile manipulation trajectories through iterative bi-layer search of the base and arm.

Table 4. Real-world results. We tested each task 10 times.

| Task | SR | Step | Latency(ms) |
|-------------|-------|-------|-------------|
| Stack Block | 30.0% | 67.0 | 580.0 |
| Open Drawer | 10.0% | 89.0 | 592.0 |
| Put in Bowl | 40.0% | 102.0 | 585.0 |

50 samples to complete the fine-tuning and achieves 40% success rate on the mobile manipulation task. The drawer opening task is a challenging hinge object interaction in mobile manipulation, requiring the robotic trajectory to meet physical constraints and avoid collision with hinge object movement, which leads to 10% success rate.

We demonstrate the qualitative results in Figure 3, which illustrates the sequence of the Pick-and-Place task of the mobile manipulation. The robot initializes the surrounding of the target object to ensure the effectiveness of the waypoints predicted by the VLA models. After the robot approaches the target object by navigation, the motion of the mobile base and arm is generated based on the waypoints produced by the pre-trained VLA models, ensuring that the end-effector follows the VLA model guidelines. Initially, target blocks are positioned within the robotic arm’s native workspace, enabling task completion through direct execution of VLA models predicted waypoints via generated arm trajectories. When the agent needs to move to place the red box, the motion planning at this point is costly and the base needs to be moved to augment the motion space of the arm to indeed reach the waypoints. We initialize base optimization with heuristic navigation, then refine the base waypoint using bi-level optimization, so that the red box reenters the workspace to complete placement. Our proposed MoManipVLA tightly integrates navigation and manipulation, en-

sure efficient execution of mobile manipulation tasks, and further leverages the generalization ability of the pre-trained VLA models.

5. Conclusion

In this paper, we have presented an efficient policy transfer framework to generalize pre-trained VLA models to mobile manipulation tasks. We first predict the waypoints for the end-effectors with the fixed-base VLA models, and generate the trajectories for the mobile base and the robot arm with maximized physical feasibility. By designing objectives considering the end-effector reachability, trajectory smoothness and collision avoidance, we leverage an efficient bi-level objective optimization framework to jointly search for the optimal trajectory pose for the base and the arm. Extensive experiments demonstrate the powerful generalization of our approach in both simulation and the real-world tasks. The limitations of our method are two-fold. First, our proposed method is limited to the performance of the VLA models and can only be deployed in highly constrained movement space. Meanwhile, for other moving platforms extra motion planning constraints need to be designed to ensure feasibility. Second, our system cannot generate effective motions for long-horizon tasks due to the lack of task planning modules. Designing an efficient trajectory generation framework and integrating task planning from foundation models will be the future works.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 62376032, in part by the NTU Start up Grant 024303-00001.

References

- [1] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024. 3
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 2
- [3] Michael Chang, Tomer Ullman, Antonio Torralba, and Joshua Tenenbaum. A compositional object-based approach to learning physical dynamics. In *International Conference on Learning Representations*, 2017. 3
- [4] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023. 3
- [5] Daniel Choi, Angus Fung, Haitong Wang, and Aaron Hao Tan. Find everything: A general vision language model approach to multi-object search. *arXiv preprint arXiv:2410.00388*, 2024. 2
- [6] Michael Danielczuk, Arsalan Mousavian, Clemens Eppner, and Dieter Fox. Object rearrangement using learned implicit collision functions. In *IEEE International Conference on Robotics and Automation*, pages 6010–6017, 2021. 3
- [7] Hao-Shu Fang, Chenxi Wang, Hongjie Fang, Minghao Gou, Jirong Liu, Hengxu Yan, Wenhai Liu, Yichen Xie, and Cewu Lu. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics*, 2023. 2
- [8] Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation using low-cost whole-body teleoperation. In *Conference on Robot Learning*, 2024. 2
- [9] Jiayuan Gu, Devendra Singh Chaplot, Hao Su, and Jitendra Malik. Multi-skill mobile manipulation for object rearrangement. *arXiv preprint arXiv:2209.02778*, 2022. 2
- [10] Abhinav Gupta, Adithyavairavan Murali, Dhiraj Prakashchand Gandhi, and Lerrel Pinto. Robot learning in homes: Improving generalization and reducing dataset bias. *Advances in Neural Information Processing Systems*, 31, 2018. 2
- [11] Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):269–296, 2020. 3
- [12] Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. An embodied generalist agent in 3d world. In *International Conference on Machine Learning*, pages 20413–20451, 2024. 2
- [13] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3d scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16750–16761, 2023. 3
- [14] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *Conference on Robot Learning*, pages 540–562. PMLR, 2023. 3
- [15] Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. *arXiv preprint arXiv:2409.01652*, 2024. 3
- [16] Xiaoyu Huang, Dhruv Batra, Akshara Rai, and Andrew Szot. Skill transformer: A monolithic policy for mobile manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10852–10862, 2023. 2
- [17] Matthew Kelly. An introduction to trajectory optimization: How to do your own direct collocation. *SIAM Review*, 59(4): 849–904, 2017. 3
- [18] Charles C Kemp, Aaron Edsinger, Henry M Clever, and Blaine Matulevich. The design of stretch: A compact, lightweight mobile manipulator for indoor human environments. In *ICRA*, pages 3150–3157, 2022. 6
- [19] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024. 2
- [20] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, et al. Vision-language foundation models as effective robot imitators. *arXiv preprint arXiv:2311.01378*, 2023.
- [21] Xiaoqi Li, Mingxu Zhang, Yiran Geng, Haoran Geng, Yuxing Long, Yan Shen, Renrui Zhang, Jiaming Liu, and Hao Dong. Manipllm: Embodied multimodal large language model for object-centric robotic manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18061–18070, 2024. 2
- [22] Peiqi Liu, Yaswanth Orru, Jay Vakil, Chris Paxton, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Ok-robot: What really matters in integrating open-knowledge models for robotics. *arXiv preprint arXiv:2401.12202*, 2024. 2
- [23] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint arXiv:2410.07864*, 2024. 3
- [24] Russell Mendonca, Emmanuel Panov, Bernadette Bucher, Jiguang Wang, and Deepak Pathak. Continuously improving mobile manipulation with autonomous real-world rl. *arXiv preprint arXiv:2409.20568*, 2024. 2
- [25] Alexander Millane, Helen Oleynikova, Emilie Wirbel, Remo Steiner, Vikram Ramasamy, David Tingdahl, and Roland Siegwart. nvblox: Gpu-accelerated incremental signed distance field mapping. In *IEEE International Conference on Robotics and Automation*, pages 2698–2705, 2024. 5
- [26] Kaichun Mo, Yuzhe Qin, Fanbo Xiang, Hao Su, and Leonidas Guibas. O2o-afford: Annotation-free large-scale

- object-object affordance learning. In *Conference on Robot Learning*, pages 1666–1677, 2022. 3
- [27] Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration0. In *IEEE International Conference on Robotics and Automation*, pages 6892–6903, 2024. 3
- [28] Yanlong Peng, Zhigang Wang, Yisheng Zhang, Shengmin Zhang, Nan Cai, Fan Wu, and Ming Chen. Revolutionizing battery disassembly: The design and implementation of a battery disassembly autonomous mobile manipulator robot (beam-1). *arXiv preprint arXiv:2407.06590*, 2024. 2
- [29] Ri-Zhao Qiu, Yafei Hu, Ge Yang, Yuchen Song, Yang Fu, Jianglong Ye, Jiteng Mu, Ruihan Yang, Nikolay Atanasov, Sebastian Scherer, et al. Learning generalizable feature fields for mobile manipulation. *arXiv preprint arXiv:2403.07563*, 2024. 2
- [30] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 7
- [31] Petr Štibinger, George Broughton, Filip Majer, Zdeněk Rozsypálek, Anthony Wang, Kshitij Jindal, Alex Zhou, Dinesh Thakur, Giuseppe Loianno, Tomáš Krajník, et al. Mobile manipulator for autonomous localization, grasping and precise placement of construction material in a semi-structured environment. *IEEE Robotics and Automation Letters*, 6(2):2595–2602, 2021. 2
- [32] Martin Sundermeyer, Arsalan Mousavian, Rudolph Triebel, and Dieter Fox. Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 13438–13444, 2021. 3
- [33] Xiaoguang Tu, Zhi He, Yi Huang, Zhi-Hao Zhang, Ming Yang, and Jian Zhao. An overview of large ai models and their applications. *Visual Intelligence*, 2(1):1–22, 2024. 2
- [34] Shagun Uppal, Ananye Agarwal, Haoyu Xiong, Kenneth Shaw, and Deepak Pathak. Spin: Simultaneous perception interaction and navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18133–18142, 2024. 2
- [35] Junjie Wen, Yichen Zhu, Jinming Li, Minjie Zhu, Kun Wu, Zhiyuan Xu, Ran Cheng, Chaomin Shen, Yaxin Peng, Feifei Feng, et al. Tinyvla: Towards fast, data-efficient vision-language-action models for robotic manipulation. *arXiv preprint arXiv:2409.12514*, 2024. 3
- [36] Anxing Xiao, Nuwan Janaka, Tianrun Hu, Anshul Gupta, Kaixin Li, Cunjun Yu, and David Hsu. Robi butler: Remote multimodal interactions with household robot assistant. *arXiv preprint arXiv:2409.20548*, 2024. 2
- [37] Haoyu Xiong, Russell Mendonca, Kenneth Shaw, and Deepak Pathak. Adaptive mobile manipulation for articulated objects in the open world. *arXiv preprint arXiv:2401.14403*, 2024. 2
- [38] Sixu Yan, Zeyu Zhang, Muzhi Han, Zaijin Wang, Qi Xie, Zhitian Li, Zhehan Li, Hangxin Liu, Xinggang Wang, and Song-Chun Zhu. M2diffuser: Diffusion-based trajectory optimization for mobile manipulation in 3d scenes. *arXiv preprint arXiv:2410.11402*, 2024. 2
- [39] Ruihan Yang, Yejin Kim, Rose Hendrix, Aniruddha Kembhavi, Xiaolong Wang, and Kiana Ehsani. Harmonic mobile manipulation. *arXiv preprint arXiv:2312.06639*, 2023. 2
- [40] Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, et al. Homerobot: Open-vocabulary mobile manipulation. *arXiv preprint arXiv:2306.11565*, 2023. 2, 6
- [41] Hang Yin, Xiuwei Xu, Zhenyu Wu, Jie Zhou, and Jiwen Lu. Sg-nav: Online 3d scene graph prompting for llm-based zero-shot object navigation. *arXiv preprint arXiv:2410.08189*, 2024. 2
- [42] Naoki Yokoyama, Alex Clegg, Joanne Truong, Eric Underlander, Tsung-Yen Yang, Sergio Arnaud, Sehoon Ha, Dhruv Batra, and Akshara Rai. Asc: Adaptive skill coordination for robotic mobile manipulation. *IEEE Robotics and Automation Letters*, 9(1):779–786, 2023. 2
- [43] Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen, Li Erran Li, and Xiaolong Wang. Gnfactor: Multi-task real robot learning with generalizable neural feature fields. In *Conference on Robot Learning*, pages 284–301, 2023. 2