

M^4 Diffuser: Multi-View Diffusion Policy with Manipulability-Aware Control for Robust Mobile Manipulation

Ju Dong^{1,2,3}, Lei Zhang^{1,3†}, Liding Zhang², Yao Ling², Yu Fu², Kaixin Bai^{1,3},
 Zoltán-Csaba Márton³, Zhenshan Bing², Zhaopeng Chen³, Alois Christian Knoll², Jianwei Zhang¹

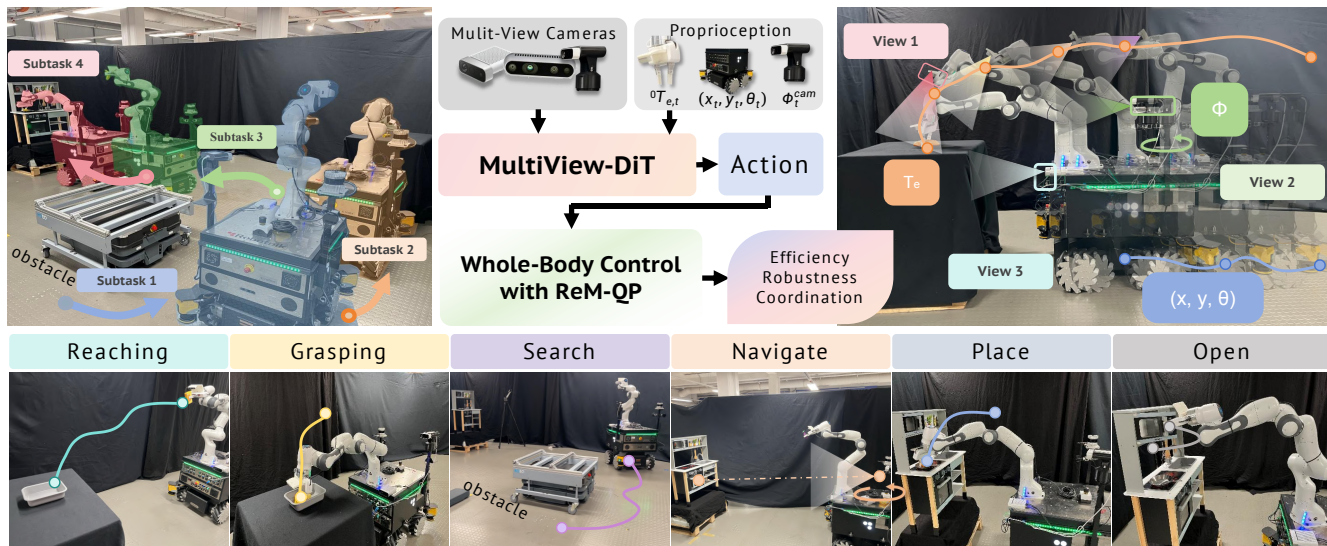


Fig. 1: M^4 Diffuser: Multi-View Diffusion Policy and ReM-QP controller for robust whole-body mobile manipulation.

Abstract—Mobile manipulation requires the coordinated control of a mobile base and a robotic arm while simultaneously perceiving both global scene context and fine-grained object details. Existing single-view approaches often fail in unstructured environments due to limited fields of view, exploration, and generalization abilities. Moreover, classical controllers, although stable, struggle with efficiency and manipulability near singularities. To address these challenges, we propose M^4 Diffuser, a hybrid framework that integrates a Multi-View Diffusion Policy with a novel Reduced and Manipulability-aware Quadratic Programming (ReM-QP) controller for mobile manipulation. The diffusion policy leverages proprioceptive states and complementary camera perspectives with both close-range object details and global scene context to generate task-relevant end-effector goals in the world frame. These high-level goals are then executed by the ReM-QP controller, which eliminates slack variables for computational efficiency and incorporates manipulability-aware preferences for robustness near singularities. Comprehensive experiments in simulation and real-world environments show that M^4 Diffuser achieves 7%–56% higher success rates and reduces collisions by 3%–31% over baselines. Our approach demonstrates robust performance for smooth whole-body coordination, and strong generalization to unseen tasks, paving the way for reliable mobile manipulation in unstructured environments. Details of the demo and supplemental

material are available on our project website <https://sites.google.com/view/m4diffuser>.

I. INTRODUCTION

Mobile manipulation, which integrates the coordinated operation of a mobile base and a robotic arm, is a crucial capability for enabling robots to autonomously perform long-horizon tasks in real-world, unstructured environments. In recent years, notable progress has been made in this field through the development of whole-body control frameworks [1, 2], motion planning algorithms [3], and imitation learning strategies [1], allowing robots to accomplish tasks such as table-to-table transfer, pick-and-place, and manipulation in kitchen scenarios. Nevertheless, several key limitations remain.

On the one hand, classical optimization-based controllers (e.g., whole-body quadratic programming controllers) provide stability and real-time performance, but they often rely on the introduction of slack variables to ensure task feasibility [4]. This design substantially increases computational overhead, reduces trajectory smoothness, and severely degrades manipulability near Jacobian singularities, frequently leading to execution failures.

On the other hand, learning-driven methods [5] demonstrate stronger adaptability and generalization across tasks. However, their deployment in practice often lacks sufficient stability. For instance, policies tend to fail when the visual input is occluded or out of view, making it difficult for the robot to maintain

†Corresponding author. lei.zhang-1@studium.uni-hamburg.de

¹TAMS (Technical Aspects of Multimodal Systems), Department of Informatics, University of Hamburg, Hamburg, Germany.

²Technical University of Munich, Germany.

³Agile Robots SE, Munich, Germany.

This work is supported by National Key Research and Development Program of China (2025YFE0217000).

continuous tracking and complete the task. Single-view perception is particularly limited, as it cannot simultaneously capture both the global scene and fine-grained object details, thereby constraining the exploration ability and generalization of the policy. To address these challenges, we propose M^4 Diffuser, a hybrid control framework with two key innovations:

- **Multi-View Diffusion Transformer Policy:** By combining complementary perspectives, the policy can jointly capture local object details and global scene context, enabling the generation of robust end-effector goals in the world frame. Experimental results show that multi-view input substantially improves robustness and enhances generalization, achieving 23%–56% higher success rates and reducing collisions by 14%–31% compared to different diffusion policy variants.
- **Reduced and Manipulability-aware Quadratic Programming Controller (ReM-QP):** At the low-control level, we remove the slack variables of traditional QP formulations to reduce computational complexity, and introduce an inverse condition number (ICN)-based manipulability preference to ensure stability and smoothness near singularities. Simulation results demonstrate that ReM-QP reduces task execution time by 28% and lowers end-effector jerk by 35%, achieving a favorable balance between efficiency and smoothness.

In summary, our method achieves robust mobile manipulation in unstructured environments through the tight integration of multi-view perception and whole-body control. Across different tasks in both simulation and real-world environments, M^4 Diffuser significantly outperforms both traditional planning methods and purely learning-based approaches, achieving an average **+28.4% improvement in success rate** and a **69% reduction in collisions** compared to the OMPL baseline. Furthermore, it surpasses state-of-the-art (SOTA) methods such as HoMeR [1], with **+10.0% higher success rates** and **5.2% fewer collisions**, while also demonstrating strong generalization to unseen objects and novel task configurations.

II. RELATED WORK

A. Mobile Manipulation

In recent years, the field of mobile manipulation has made significant progress in addressing several fundamental challenges. Robots are now capable of integrating navigation and manipulation.

In the field of traditional methods, research on whole-body control frameworks [6, 2, 4] and grasp detection algorithms [7] has enabled a certain degree of coordination between the mobile base and the manipulator. In addition, several works have tackled issues such as collision avoidance and adaptation to dynamic environments [8, 9, 10]. However, in real-world unstructured scenarios, these traditional methods still suffer from limited generalization ability, making it difficult to cope with complex and highly variable conditions.

Enhancing the searching ability is also of critical importance for mobile manipulation. Visual perception can significantly

improve a robot’s searching capability [11]; however, in practice, it is frequently hindered by the out-of-view problem. Traditional methods provide partial solutions to this issue, yet their effectiveness is often constrained by poor generalization performance. For example, EHC-MM [3] propose a monitor-position-based servoing (MPBS) function to prevent the robot from losing track of the target during its operation.

Learning-based approaches have demonstrated stronger task and scene generalization in structured and semi-structured settings [12, 13]. Recent studies have also explored learning-based coordination between the base and manipulator, for instance with diffusion policies [12] and reinforcement learning [14]. Nevertheless, purely learning-driven methods typically require large-scale datasets to achieve stable performance, and the stability and reliability of the learned policies remain insufficient for industrial applications.

B. Controller for Mobile Manipulation

Classical controllers for mobile manipulation, such as hierarchical task and motion planning [6] and whole-body QP controllers [15], provide strong stability and real-time performance but suffer from critical drawbacks. Holistic QP formulations often introduce slack variables to relax task constraints, which increases computational overhead and reduces trajectory smoothness in long-horizon tasks. Moreover, these optimization-based controllers are sensitive to Jacobian singularities, degrading manipulability and causing failures in cluttered environments [16]. Reactive whole-body controllers [2, 8] improve adaptability in dynamic settings but frequently produce jerky motions and lack global task consistency. To address these limitations, we propose the Reduced and Manipulability-aware QP (ReM-QP) controller, which eliminates slack variables for efficiency and incorporates manipulability-aware preferences for robustness.

C. Hybrid Control for Mobile Manipulation

Traditional controllers ensure stability and real-time performance but lack adaptability, while learning-based approaches offer flexibility and generalization at the cost of stability and safety. To address this, recent research has proposed hybrid control architectures that combine the stability of classical planning with the flexibility of data-driven strategies [1]. These methods have achieved promising levels of success and robustness in laboratory environments and benchmark tasks. However, our experimental results indicate that even state-of-the-art controllers still face performance bottlenecks in base–arm coordination, which ultimately limits the stability and applicability of hybrid approaches. Moreover, current mobile manipulation systems exhibit limited exploration ability. To overcome this limitation, we propose a multi-view policy learning framework with a manipulability-aware whole-body QP controller.

D. Imitation Learning for Robotic Manipulation

Imitation learning and generative models have been widely adopted in robotic manipulation due to its data efficiency and




Fig. 2: Diffusion transformer policy architecture. Multi-view RGB observations and proprioceptive states are encoded into latent features, which condition a denoising diffusion process implemented with a Transformer. The policy outputs a desired end-effector goal in the world frame, which is converted into a twist and executed by the ReM-QP controller.

ease of training [17, 18, 19, 20, 21]. Early approaches relied on behavior cloning and offline expert demonstrations. More recent advances, such as Diffusion Policy [5] and ACT [22], have introduced denoising diffusion models and transformer architectures to achieve superior generalization and success rates. To collect training data for these policies, we develop a teleoperation platform for mobile manipulation [23].

III. PROBLEM STATEMENT AND METHODS

A. Problem Formulation

We address the problem of coordinated mobile manipulation in unstructured environments. The goal is to control a mobile manipulator, comprising a holonomic mobile base and a 7-DoF robotic arm, using multi-modal observations to perform long-horizon tasks that require both global navigation and fine-grained manipulation.

At each timestep t , the observation is defined as:

$$o_t = \{RGB_t^{1:3}, \mathbf{T}_{e,t}, \phi_t^{\text{cam}}, (x_t, y_t, \theta_t)\} \quad (1)$$

where: $RGB_t^{1:3}$ denotes synchronized images from three RGB cameras. $\mathbf{T}_{e,t} \in SE(3)$ represents current end-effector pose. ϕ_t^{cam} denotes rotating rear-mounted camera angle. (x_t, y_t, θ_t) represents mobile base pose in world coordinates. Given this input, the goal is to compute a joint action for the full robot body that achieves robust task execution. Our solution decomposes this into two tightly coupled modules: Multi-View Diffusion Policy and ReM-QP Controller.

The high-level Multi-View Diffusion Policy is designed to capture global scene context and local object geometry through multi-view visual inputs and proprioceptive feedback. It learns a conditional diffusion process that predicts a desired end-effector pose in the world frame:

$$\mathbf{T}_e^* \sim \pi_\theta(a_t | o_t) \quad (2)$$

The learning objective of this module is to train a conditional diffusion model that denoises perturbed actions and reconstructs the desired end-effector pose given multi-view observations and proprioceptive states. The network is introduced in Sec. III-B and shown in Fig. 2.

Whole-body controller then takes the output of the policy as input and computes the corresponding joint and base commands for execution. Formally, the process can be expressed as:

$$\dot{\mathbf{q}}^* = \text{ReM-QP}\left(f_{\text{servo}}\left(\mathbf{T}_{e,t}, \pi_\theta(o_t)\right)\right), \quad (3)$$

This formulation shows that the whole-body velocity command $\dot{\mathbf{q}}^*$ is obtained by passing the end-effector goal $\pi_\theta(o_t)$ predicted by the policy through a position-based servoing law $f_{\text{servo}}(\cdot)$ and the proposed ReM-QP controller $\text{ReM-QP}(\cdot)$.

The hybrid control structure is introduced in Sec. III-C and proposed ReM-QP controller is detailed in Sec. III-D.




Fig. 3: ReM-QP controller. The reduced QP formulation eliminates slack variables for faster optimization, while ICN-based preferences improve robustness near singularities. This ensures efficient and stable whole-body execution of the high-level policy outputs.

B. Multi-View Diffusion Transformer Policy

Our policy network integrates a Transformer-based encoder with a conditional denoising diffusion process to generate robust task-space actions from raw sensory inputs[24, 25, 26, 27]. Multi-view RGB observations are first processed by convolutional backbones and concatenated with proprioceptive encodings. The fused features are passed through a Transformer encoder, which performs cross-view fusion and temporal reasoning via self-attention, producing a compact scene-state representation $\mathbf{h} = \text{Transformer}(\mathbf{x})$.

Conditioned on \mathbf{h} , the diffusion process iteratively denoises Gaussian noise into feasible action distributions. Given an expert action \mathbf{a}_0 and diffusion time step k , the forward process is defined as

$$q(\mathbf{a}_k | \mathbf{a}_0) = \mathcal{N}(\mathbf{a}_k; \sqrt{\bar{\alpha}_k} \mathbf{a}_0, (1 - \bar{\alpha}_k) \mathbf{I}), \quad (4)$$

where $\bar{\alpha}_k = \prod_{s=1}^k (1 - \beta_s)$ is determined by the noise schedule. The reverse process is modeled as

$$p_\theta(\mathbf{a}_{k-1} | \mathbf{a}_k, \mathbf{h}) = \mathcal{N}(\mathbf{a}_{k-1}; \mu_\theta(\mathbf{a}_k, k, \mathbf{h}), \Sigma_\theta(\mathbf{a}_k, k, \mathbf{h})), \quad (5)$$

where \mathbf{h} conditions the denoising network through cross-attention. The training objective reduces to noise prediction:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{\mathbf{a}_0, \epsilon, k} [\|\epsilon - \epsilon_\theta(\mathbf{a}_k, k, \mathbf{h})\|^2]. \quad (6)$$

At inference, the final denoised latent is decoded into task-space commands, which are executed by a downstream whole-body controller.

C. Hybrid Control Strategy

In summary, our framework tightly couples a high-level diffusion transformer policy with a reduced whole-body QP (ReM-QP) controller. The diffusion policy leverages multi-view RGB observations and proprioceptive inputs to generate robust end-effector trajectories in unstructured environments. The ReM-QP controller then executes these high-level goals as real-time whole-body commands, achieving both computational efficiency (via slack elimination) and robustness (via ICN-based manipulability preferences), while ensuring safety through standard inequality constraints. The overall hybrid control pipeline can be expressed as:

$$\begin{aligned} {}^0T_e^* &\sim \pi_\theta(a_t | o_t), \\ \nu_{ee}^* &= f_{\text{servo}}({}^0T_{e,t}, {}^0T_e^*), \\ \dot{q}^* &= \text{ReM-QP}(\nu_{ee}^*). \end{aligned} \quad (7)$$

where the diffusion transformer policy π_θ outputs a desired end-effector pose ${}^0T_e^*$, the servoing law converts this pose into a desired twist ν_{ee}^* , and the ReM-QP controller maps the twist to coordinated joint and base velocities \dot{q}^* .

D. Reduced and Manipulability-aware Quadratic Programming (ReM-QP) Controller

The ReM-QP controller serves as the low-level whole-body solver, as shown in Fig. 3. Building upon the holistic QP formulation of Jesse Haviland et al. [15], but introduces two key improvements: (i) elimination of slack variables for improved computational efficiency, and (ii) an ICN-based preference term for robust manipulability. These modifications yield a reduced strictly-convex QP that is both faster to solve and more robust near singular configurations, enabling safe execution of the high-level policy outputs.

1) Whole-Body Kinematics

We define four coordinate frames: the world \mathcal{F}_0 , the mobile base \mathcal{F}_b , the arm base \mathcal{F}_a , and the end-effector \mathcal{F}_e . The whole-body forward kinematics is

$${}^0T_e = {}^0T_b(x, y, \theta) {}^bT_a {}^aT_e(q_a), \quad (8)$$

where (x, y, θ) parameterizes the holonomic base, and $q_a \in \mathbb{R}^m$ are the arm joint angles. The end-effector twist follows

$$\nu_{ee} = J(q) \dot{q}, \quad \nu_{ee} \in \mathbb{R}^6, \quad (9)$$

where $\dot{q} \in \mathbb{R}^n$ concatenates base and arm velocities ($n = m + 3$), and $J(q) \in \mathbb{R}^{6 \times n}$ is the whole-body Jacobian. A desired twist ν_{ee}^* is obtained from a closed-loop servoing law [28].

2) Reduced QP Formulation

The original holistic controller introduces slack variables δ to soften task constraints [3, 6]:

$$\begin{aligned} \min_{\dot{q}, \delta} \quad & \frac{1}{2} \begin{bmatrix} \dot{q} \\ \delta \end{bmatrix}^\top \begin{bmatrix} Q_{qq} & 0 \\ 0 & Q_{\delta\delta} \end{bmatrix} \begin{bmatrix} \dot{q} \\ \delta \end{bmatrix} + c_q^\top \dot{q}, \\ \text{s.t.} \quad & J\dot{q} + \delta = \nu_{ee}^*, \\ & A\dot{q} \leq b, \end{aligned} \quad (10)$$

where $Q_{qq} \succ 0$ regularizes generalized velocities, $Q_{\delta\delta} \succ 0$ penalizes the slack, c_q encodes preference terms, and $A\dot{q} \leq b$ collects velocity bounds and dampers. We eliminate the slack by substituting $\delta = \nu_{ee}^* - J\dot{q}$, reducing the QP to a problem in \dot{q} only:

$$\min_{\dot{q}} \quad \frac{1}{2} \dot{q}^\top Q_{\text{red}} \dot{q} + c_{\text{red}}^\top \dot{q}, \quad \text{s.t.} \quad A\dot{q} \leq b, \quad (11)$$

with

$$Q_{\text{red}} = Q_{qq} + J^\top Q_{\delta\delta} J, \quad (12)$$

$$c_{\text{red}} = c_q - J^\top Q_{\delta\delta} \nu_{ee}^*. \quad (13)$$

Here, Q_{red} is the reduced Hessian, which combines velocity regularization and task-tracking penalties, while c_{red} is the reduced linear term. The latter aggregates both task-consistent preferences (e.g., manipulability) and the residual cost of tracking ν_{ee}^* .

3) ICN-Based Manipulability Preference

To enhance robustness near singular configurations, we introduce a preference term based on the *inverse condition number (ICN)* of the Jacobian [16, 29]:

$$\text{ICN}(q) = \frac{\sigma_{\min}(J)}{\sigma_{\max}(J)} \in [0, 1], \quad (14)$$

where σ_{\min} and σ_{\max} are the smallest and largest singular values of J , respectively. The gradient of ICN is approximated via centered finite differences:

$$\frac{\partial \text{ICN}(q)}{\partial q_i} \approx \frac{\text{ICN}(q + \delta e_i) - \text{ICN}(q - \delta e_i)}{2\delta}, \quad i = 1, \dots, n. \quad (15)$$

The negative gradient $-\nabla \text{ICN}(q)$ is incorporated into the linear coefficient c_q , biasing the solution toward configurations with better conditioning and away from singularities.

4) Safety Constraints

Based on previous method [15], we retain standard inequality constraints $A\dot{q} \leq b$ to encode joint velocity limits and velocity dampers near joint limits.

Since the reduced Hessian Q_{red} is strictly positive definite, the resulting QP remains strictly convex and admits

a unique global optimum at each control cycle. The ICN-based preference only modifies the linear term c_{red} and does not affect feasibility or convexity. Overall, the ReM-QP controller improves over the holistic baseline by eliminating slack variables for faster computation and introducing an ICN-based manipulability preference for robustness. Together with standard safety constraints, this yields an efficient, safe, and real-time whole-body controller that faithfully executes the diffusion policy outputs.

IV. EXPERIMENT

A. Real-World and Simulation Experiment Setup

We employ the *DARKO* robot for experiments, as shown in Fig. 4. For policy training, we use teleoperation to collect 350 trajectories for each task. Implementation details of teleoperation are introduced in Supplementary Materials. We simulate the *DARKO* mobile manipulator in the MuJoCo physics engine. The scene consists of two tables (Table A and Table B) and a single graspable object initially placed on Table A.




Fig. 4: The *DARKO* robot platform. It is built on an omnidirectional mobile base (RB-Kairos) with mecanum wheels, and equipped with a Franka-Emika-Panda robotic arm (1) carrying a wrist-mounted RealSense D435i camera (2). Additional sensing equipment includes two Sick MicroScan 2D lidars (3), a front-facing Azure Kinect RGB-D camera (4), and an adjustable Azure Kinect RGB-D camera (5)

B. Experiment Setup of Whole-Body Control in Simulation

We evaluate proposed whole-body control (WBC) in the simulation environment. The robot is controlled to move the object from Table A to Table B. As shown in Fig. 5, the task is decomposed into the following four phases: (i) **Nav A** — navigate to Table A; (ii) **Desk A** — grasp and lift the object; (iii) **Nav B** — navigate to Table B while holding the object; (iv) **Desk B** — place the object and lift the arm. This structured setup enables phase-wise comparison across controllers and allows us to jointly evaluate omnidirectional navigation and dexterous manipulation under a unified whole-body control framework.

Variants: For ablation study and comparison experiment, we evaluate the following controller variants:

- **Baseline:** Holistic whole-body controller with slack variables [15].
- **Elim. Slacks:** Our reduced QP formulation that eliminates slack variables to improve computational efficiency.




Fig. 5: Task phases for the mobile manipulation benchmark in the MuJoCo simulator. The four phases consist of two navigation segments (Nav A, Nav B) and two dexterous manipulation segments (Desk A, Desk B).

TABLE I: Overall performance across controllers. ↓ lower is better, ↑ higher is better.

Method	EE Acc. RMS (m/s ²)↓	EE Jerk RMS (m/s ³)↓	Total Time (s)↓
Baseline [15]	0.151	4.61	146.67
Elim. Slacks	0.184	9.47	83.86
Baseline+ICN	0.055	1.58	128.57
Ours (ReM-QP)	0.087	2.98	105.87




Fig. 6: Visualization of ReM-QP whole-body control. The controller ensures coordinated base–arm behavior by automatically coupling base motion to end-effector trajectories.

- **Baseline+ICN:** Holistic baseline controller augmented with ICN-based grasp preference.
- **Ours (ReM-QP):** Our method with slack elimination and ICN-based preferences.

Evaluation Metrics: We measure phase-wise execution time and global motion quality metrics, including RMS end-effector acceleration, RMS jerk, and total task time. These metrics jointly reflect both efficiency (execution time) and motion quality (trajectory smoothness).

C. Qualitative Results of Whole-Body Control with ReM-QP

To complement the quantitative metrics, we visualize the trajectories of both the end-effector and the mobile base. Fig. 6 illustrates how the ReM-QP controller couples end-effector goals with coordinated base movements: when the end-effector advances toward the target, the base follows with forward motions (green arrows), while backward adjustments (red arrows) emerge when the controller repositions the base to maintain reachability and manipulability. This visualization highlights the key property of our controller: **end-effector goals actively guide base motion**, enabling smooth whole-body coordination across both navigation and manipulation.

TABLE II: Per-phase execution time (s). Task is split into navigation (Nav A/B) and dexterous phases (Desk A/B).

Phase	Baseline	Elim. Slacks	Baseline+ICN	Ours (ReM-QP)
Nav A (to Table A)	58.59	23.06	36.21	30.15
Desk A (grasp+lift)	9.15	6.60	8.65	7.33
Nav B (to Table B)	59.06	41.85	67.64	47.00
Desk B (place+lift)	19.87	12.17	16.07	21.40
Total	146.67	83.86	128.57	105.87

D. Quantitative Results of Whole-Body Control with ReM-QP

The quantitative results of whole-body control are summarized in Tab. I and Tab. II. The four controllers are compared in terms of overall motion quality (end-effector acceleration/jerk RMS and total time) and phase-wise execution time.

As shown in Tab. I, the baseline controller suffers from high jerk (4.61 m/s^3) and long completion time (147 s). Eliminating slack variables (Elim. Slacks) drastically reduces the total execution time to 83.9 s, highlighting the efficiency gain of the reduced QP formulation. However, this speed-up comes at the cost of higher acceleration and jerk, reflecting less smooth motions. In contrast, augmenting the baseline with ICN-based preferences (Baseline+ICN) significantly improves smoothness with the lowest acceleration (0.055 m/s^2) and jerk (1.58 m/s^3). However, the overall time remains long (128.8 s) due to inefficiencies in navigation.

Our proposed method (ReM-QP) combines the benefits of both approaches. By jointly eliminating slacks and incorporating ICN-based preferences, it achieves a balanced performance: the total task time is reduced to 105.87 s, while maintaining substantially lower jerk (2.98 m/s^3) than Elim. Slacks. This indicates that ReM-QP preserves smooth and stable grasping while still being considerably faster than the baseline controller.

The phase-wise breakdown in Tab. II provides further insights. Navigation phases (Nav A/B) dominate task time. Elim. Slacks yields the fastest navigation, but at the expense of motion quality. Baseline+ICN, on the other hand, improves dexterous phases (Desk A/B) with smoother end-effector motions, but shows slower navigation due to less effective base–arm coordination. Our method strikes a middle ground: navigation is faster than baseline and Baseline+ICN, while desk phases remain smooth and reliable. Notably, Desk A is completed in 7.3 s, close to the efficiency of Elim. Slacks, but without compromising grasp stability.

Overall, these results demonstrate that the proposed ReM-QP controller achieves a favorable trade-off: reducing task time by 28% compared to the baseline while lowering jerk by 35%, thus ensuring both efficiency and robustness in long-horizon mobile manipulation. These findings justify the use of ReM-QP as the low-level controller in the subsequent real-world M^4 Diffuser experiments.

E. Experiment Setup of Real-World Mobile Manipulation

We evaluate the M^4 Diffuser framework on the physical DARKO robot across a suite of mobile manipulation tasks that require both reliable navigation and dexterous manipulation.

Each task consists of two stages: (i) navigation through the scene while avoiding obstacles, and (ii) a tabletop manipulation task. We consider both simplified environments (“easy”) and realistic kitchen scenes with unstructured layouts.

- **Reaching (easy/ kitchen):** The robot grasps an object and navigates around obstacles to place it into a container on a table or on the kitchen stove.
- **Pick-and-Place (easy/ kitchen):** The robot navigates around obstacles, grasps an object from a table, and places it into a container on the same table or delivers it to a designated location in the kitchen setup.
- **Open Door (kitchen):** The robot navigates to a microwave in the kitchen scene and open the door.

Specifically, LiDAR-based SLAM method is utilized to obtain the location of mobile base.

Baselines for Comparison Experiments: We compare our approach against the following baselines:

- **OMPL baseline with AprilTags:** A traditional decoupled planning pipeline implemented using the OMPL library [30]. The mobile base is driven by LiDAR-based SLAM, while arm are planned separately in joint space.
- **Diffusion policy (single-view/ multi-view):** Diffusion policy is conditioned only on the single-view or multi-view information outputs joint-level actions for both the arm and the base.
- **Diffusion policy (multi-view) + ReM-QP:** The multi-view diffusion policy predicts high-level end-effector goals, which are then converted into whole-body joint commands by our proposed ReM-QP controller.
- **HoMeR [1](multi-view) + WBC:** combines RGB-based relative actions and point cloud-based keyposes, executed via a WBC.
- **Diffusion Transformer (multi-view) + ReM-QP (ours):** Multi-view diffusion transformer policy with the ReM-QP controller.

For all multi-view settings except HoMeR, the rear camera is mounted on an adjustable bracket, allowing flexible viewpoint adaptation; HoMeR relies on fixed viewpoints.

F. Qualitative Results of Real-World Experiments

Fig. 7 illustrate representative failures in real-world rollouts. The *OMPL baseline* frequently fails in kitchen scenes, mainly due to its reliance on AprilTags and the brittleness of decoupled planning. The *Diffusion Policy (multi-view)* variant, although marker-free, still suffers from unstable base–arm coordination, leading to jitter and object slips. These qualitative examples highlight two major limitations of existing approaches: strong dependence on artificial markers and unstable execution without optimization. Both issues are addressed by our framework. These cases highlight the brittleness of marker-dependent planning and the instability of direct policy execution, which motivate the need for our design.

G. Quantitative Results of Real-World Experiments

The quantitative results are detailed in Tab. III. On average, the **OMPL baseline** achieves a success rate of 54% and a




Fig. 7: Failure cases in real-world experiments. (Case 1) OMPL baseline fails to pick a banana: wrist camera loses AprilTag (blue arrow/circle). (Case 2) OMPL baseline fails to place a banana: trajectory error causes collision with the cabinet edge (red arrow/circle). (Case 3) OMPL baseline fails to place a croissant: arm collides with the cabinet door (red arrow/circle). (Case 4) OMPL baseline fails to grasp a croissant: arm strikes the table surface (red arrow/circle). (Case 5) DP (multi-view) fails to pick Vaseline: jitter leads to object slip (yellow arrow/box).

TABLE III: Task-wise performance of different methods. Each task was evaluated over 50 trials. Metrics: success rate (%) and collision rate (%).

Task	OMPL baseline		DP (single-view)		DP (multi-view)		DP (multi-view) + ReM-QP		HoMeR [1]		M^4 Diffuser (ours)	
	Success	Coll.	Success	Coll.	Success	Coll.	Success	Coll.	Success	Coll.	Success	Coll.
Reaching (easy)	66.0	18.0	32.0	34.0	68.0	20.0	82.0	6.0	82.0	8.0	90.0	4.0
Reaching (kitchen)	54.0	22.0	28.0	38.0	60.0	22.0	76.0	8.0	74.0	14.0	84.0	6.0
Pick-and-Place (easy)	60.0	20.0	30.0	36.0	66.0	18.0	80.0	10.0	76.0	10.0	86.0	6.0
Pick-and-Place (kitchen)	52.0	24.0	24.0	40.0	58.0	24.0	76.0	12.0	72.0	12.0	82.0	8.0
Open Door (kitchen)	38.0	26.0	20.0	42.0	44.0	20.0	62.0	12.0	58.0	16.0	70.0	10.0
Average	54.0	22.0	26.8	38.0	59.2	20.8	75.2	9.6	72.4	12.0	82.4	6.8

collision rate of 22%. In the kitchen environment, the success rate decreased while the collision rate increased, indicating the higher difficulty of tasks in this setting. **DP (single-view)** performs worst overall, with success rates averaging only $\sim 27\%$ and collisions 38%, showing that single-view perception is insufficient. **DP (multi-view)** improves robustness compared to single-view and performs on par with OMPL, but collisions remain high ($\sim 21\%$). For **DP (multi-view) + ReM-QP**, optimization brings clear gains: success rises to $\sim 75\%$ and collisions drop below 10%. **HoMeR [1]** achieves intermediate performance, averaging 72.4% success and 12% collisions, slightly below DP (multi-view) + ReM-QP. Finally, our method achieves the best overall performance, averaging **82.4% success** with only **6.8% collisions**, and shows particularly strong improvements in challenging kitchen tasks.

These results confirm the necessity of our hybrid framework: OMPL is brittle and marker-dependent, DP alone is unstable, and even DP+ReM-QP is less reliable than our transformer-based hybrid. The ReM-QP controller is essential for translating high-level diffusion outputs into smooth, feasible whole-body motions, while the multi-view transformer policy provides robust spatial reasoning. In addition, the use of an adjustable rear-mounted camera ensures that target objects remain consistently within view, further improving robustness in real-world environments.

H. Generalization Experiments

To further evaluate robustness and generalization, we conducted additional experiments with unseen objects and altered task configurations. Although the policy was trained primarily on banana manipulation, we replaced the target with a croissant and an eggplant. The robot successfully grasped and placed both objects with high success rates of 80% in 50 trials, demonstrating strong object-level generalization. Moreover, when target placement positions were shifted to novel locations,

the hybrid method maintained reliable performance, showing robustness to position variations. These experiments highlight the adaptability of our multi-view diffusion transformer policy, reinforced by the ReM-QP controller, and underline its practical deployability in unstructured real-world environments.

V. CONCLUSION AND FUTURE WORK

This work presented a hybrid control framework M^4 Diffuser for mobile manipulation, integrating a **multi-view diffusion transformer policy** with a **reduced whole-body quadratic programming (ReM-QP) controller**. The framework addresses two long-standing challenges: the limited adaptability of classical model-based planners and the instability of purely learning-based policies. By leveraging multi-view visual observations and proprioceptive feedback, the diffusion transformer policy generates robust high-level end-effector goals. The ReM-QP controller then translates these goals into smooth and feasible whole-body motions, achieving real-time coordination between the mobile base and the manipulator arm. Extensive evaluations in simulation and real-world environments validate the framework’s effectiveness. In simulation, ReM-QP reduces execution time by 28% and lowers end-effector jerk by 35%. In real-world tasks, M^4 Diffuser achieves an average of 82% success rate with only 6.8% collisions, representing a relative gain of about 28% higher success and 69% fewer collisions compared to the OMPL baseline. Moreover, our method consistently outperforms strong learning-based baselines, including HoMeR [1], achieving about 10% higher success and over 5% fewer collisions. Generalization experiments further demonstrated robustness for unseen objects and novel scene configurations. Beyond these empirical gains, our framework advances the state of the art in three aspects: (1) it unifies navigation and manipulation through a single whole-body control pipeline, avoiding brittle decoupling strategies; (2) it removes reliance on fiducial markers, enabling fully marker-free

deployment; and (3) it achieves a favorable balance between the adaptability of learning-based policies and the precision of optimization-based control, ensuring scalability and deployability. In summary, hybridizing diffusion-based visuomotor policies with efficient whole-body optimization provides a robust and generalizable solution for mobile manipulation. Our framework achieves safe, marker-free operation in unstructured kitchen environments, while generalizing to unseen objects and task variations, paving the way toward practical deployment in homes, warehouses, and healthcare settings.

Future work will focus on language- and multimodal-guided mobile manipulation.

REFERENCES

- [1] P. Sundaresan, R. Malhotra, P. Miao, J. Yang, J. Wu, H. Hu, R. Antonova, F. Engelmann, D. Sadigh, and J. Bohg, "Homer: Learning in-the-wild mobile manipulation via hybrid imitation and whole-body control," *arXiv preprint arXiv:2506.01185*, 2025. 1, 2, 6, 7
- [2] A. Dietrich, T. Wimbock, A. Albu-Schaffer, and G. Hirzinger, "Reactive whole-body control: Dynamic mobile manipulation using a large number of actuated degrees of freedom," *IEEE Robotics Automation Magazine*, vol. 19, no. 2, pp. 20–33, 2012. 1, 2
- [3] J. Wang, Y. Jin, J. Shi, Y. A. D. Li, F. Sun, D. Luo, and B. Fang, "Ehc-mm: Embodied holistic control for mobile manipulation," 2025. 1, 2, 4
- [4] L. Zhang, K. Cai, Z. Sun, Z. Bing, C. Wang, L. Figueredo, S. Haddadin, and A. Knoll, "Motion planning for robotics: A review for sampling-based planners," *Biomimetic Intelligence and Robotics*, vol. 5, no. 1, p. 100207, 2025. 1, 2
- [5] 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*, 2024. 1, 3
- [6] A. Escande, N. Mansard, and P.-B. Wieber, "Hierarchical quadratic programming: Fast online humanoid-robot motion generation," *The International Journal of Robotics Research*, vol. 33, no. 7, pp. 1006–1028, 2014. 2, 4
- [7] S. Jauhri, S. Lueth, and G. Chalvatzaki, "Active-perceptive motion generation for mobile manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 1413–1419. 2
- [8] B. Burgess-Limerick, J. Haviland, C. Lehnert, and P. Corke, "Reactive base control for on-the-move mobile manipulation in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2048–2055, 2024. 2
- [9] B. Burgess-Limerick, C. Lehnert, J. Leitner, and P. Corke, "An architecture for reactive mobile manipulation on-the-move," 2022. 2
- [10] N. Marticorena, T. Fischer, J. Haviland, and N. Sunderhauf, "Rmmi: Reactive mobile manipulation using an implicit neural map," 2024. 2
- [11] M. Bajracharya, J. Borders, R. Cheng, D. Helmick, L. Kaul, D. Kruse, J. Leichthy, J. Ma, C. Matl, F. Michel *et al.*, "Demonstrating mobile manipulation in the wild: A metrics-driven approach," *arXiv preprint arXiv:2401.01474*, 2024. 2
- [12] S. Chen, J. Liu, S. Qian, H. Jiang, L. Li, R. Zhang, Z. Liu, C. Gu, C. Hou, P. Wang *et al.*, "Ac-dit: Adaptive coordination diffusion transformer for mobile manipulation," *arXiv preprint arXiv:2507.01961*, 2025. 2
- [13] S. Yan, Z. Zhang, M. Han, Z. Wang, Q. Xie, Z. Li, Z. Li, H. Liu, X. Wang, and S.-C. Zhu, "M2 diffuser: Diffusion-based trajectory optimization for mobile manipulation in 3d scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, p. 1–17, 2025. 2
- [14] W. Zhao, E. Rantala, S. Salimpour, Z. Li, J. Pajarinen, and J. P. Queralta, "Exploiting local observations for robust robot learning," 2025. 2
- [15] J. Haviland, N. Sunderhauf, and P. Corke, "A holistic approach to reactive mobile manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, p. 3122–3129, Apr. 2022. 2, 4, 5
- [16] N. Vahrenkamp, T. Asfour, G. Metta, G. Sandini, and R. Dillmann, "Manipulability Analysis," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2012, pp. 568–573. 2, 4
- [17] L. Zhang, K. Bai, G. Huang, Z. Bing, Z. Chen, A. Knoll, and J. Zhang, "Contactdextnet: Multi-fingered robotic hand grasping in cluttered environments through hand-object contact semantic mapping," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 8495–8502. 3
- [18] Y. Wang, L. Zhang, Y. Tu, H. Zhang, K. Bai, Z. Chen, and J. Zhang, "Tooleenet: Tool affordance 6d pose estimation," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 10 519–10 526. 3
- [19] H. Xu, L. Zhang, X. Hu, B. Zhong, K. Bai, Z.-C. Márton, Z. Bing, Z. Chen, A. C. Knoll, and J. Zhang, "Funcanon: Learning pose-aware action primitives via functional object canonicalization for generalizable robotic manipulation," *arXiv preprint arXiv:2509.19102*, 2025. 3
- [20] L. Zhang, J. Dong, K. Bai, M. Ni, Z.-C. Marton, Z. Chen, and J. Zhang, "Responsiblerobotbench: Benchmarking responsible robot manipulation using multi-modal large language models," *arXiv preprint arXiv:2512.04308*, 2025. 3
- [21] L. Zhang, D. Zheng, K. Bai, Z. Bing, Z.-C. Marton, Z. Chen, A. C. Knoll, and J. Zhang, "Omnidexvlg: Learning dexterous grasp generation from vision language model-guided grasp semantics, taxonomy and functional affordance," *arXiv preprint arXiv:2512.03874*, 2025. 3
- [22] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," 2023. 3
- [23] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation," in *Conference on Robot Learning (CoRL)*, 2024. 3
- [24] S. Dasari, O. Mees, S. Zhao, M. K. Srirama, and S. Levine, "The ingredients for robotic diffusion transformers," *arXiv preprint arXiv:2410.10088*, 2024. 3
- [25] A. Donat, X. Jia, X. Huang, A. Taranovic, D. Blessing, G. Li, H. Zhou, H. Zhang, R. Lioutikov, and G. Neumann, "Towards fusing point cloud and visual representations for imitation learning," 2025. 3
- [26] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, "Aloha unleashed: A simple recipe for robot dexterity," 2024. 3
- [27] H. Kim, Y. Ohmura, and Y. Kuniyoshi, "Transformer-based deep imitation learning for dual-arm robot manipulation," 2025. 3
- [28] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer International Publishing, 2009. 4
- [29] M. Togai, "An application of the singular value decomposition to manipulability and sensitivity of industrial robots," *SIAM J. Algebraic Discrete Methods*, vol. 7, no. 2, p. 315–320, Apr. 1986. 4
- [30] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, December 2012. 6