

Task-Driven Co-Design of Mobile Manipulators

Raphael Schneider*, Daniel Honerkamp*, Tim Welschehold, and Abhinav Valada

Abstract—Recent interest in mobile manipulation has resulted in a wide range of new robot designs. A large family of these designs focuses on modular platforms that combine existing mobile bases with static manipulator arms. They combine these modules by mounting the arm in a tabletop configuration. However, the operating workspaces and heights for common mobile manipulation tasks, such as opening articulated objects, significantly differ from tabletop manipulation tasks. As a result, these standard arm mounting configurations can result in kinematics with restricted joint ranges and motions. To address these problems, we present the first Concurrent Design approach for mobile manipulators to optimize key arm-mounting parameters. Our approach directly targets task performance across representative household tasks by training a powerful multitask-capable reinforcement learning policy in an inner loop while optimizing over a distribution of design configurations guided by Bayesian Optimization and HyperBand (BOHB) in an outer loop. This results in novel designs that significantly improve performance across both seen and unseen test tasks, and outperform designs generated by heuristic-based performance indices that are cheaper to evaluate but only weakly correlated with the motions of interest. We evaluate the physical feasibility of the resulting designs and show that they are practical and remain modular, affordable, and compatible with existing commercial components. We open-source the approach and generated designs to facilitate further improvements of these platforms.

Index Terms—Concurrent Design, Mobile Manipulation, Mechanism Design

I. INTRODUCTION

IN the last few years, we have seen an increasing adoption of mobile manipulator robots, leading to a wide range of new robot designs. Most commonly, mobile manipulators are constructed as modular systems based on existing commercially available mobile bases and manipulator arms, as shown in Fig. 2. This is a simple and cost-efficient strategy. However, we find that these platforms overwhelmingly follow the simplest possible mounting strategy, placing the manipulator arm on the base in the same configuration as for tabletop manipulation. In this work, we evaluate this design choice across a wide range of mobile manipulation tasks and find it to often lead to suboptimal kinematics, with restrictive joint limits that make it difficult to perform mobile manipulation tasks that often require different motions and workspaces than tabletop task, such as sideways grasping of drawers and handles at various heights, as illustrated in Fig. 1.

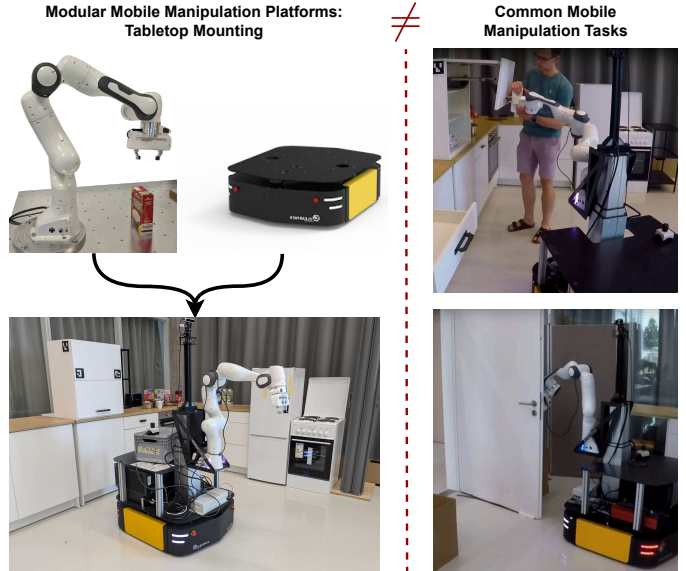


Fig. 1. Modular platforms have become a popular, cost-efficient design for mobile manipulators. However, they typically mount the manipulator arms identically to tabletop settings. We investigate this design choice and find that it is suboptimal for common mobile manipulation tasks that require different sets of motions. We then present a concurrent design approach to optimize the mounting parameters directly based on task success.

Concurrent Design (co-design) [1], [2] simultaneously refines the robot’s morphology and control policy, yielding more holistic and high-performing solutions. These methods have been applied to systems ranging from locomotion [3], [4] to manipulator arms [5], [6]. In this work, we identify the mounting parameters of mobile manipulators as a crucial design element and propose the first approach for the co-design of modular mobile manipulators. While our approach is not limited to these parameters, by focusing on arm mounting, the generated designs retain modularity and cost-efficiency, making them easily realizable in the real world.

Having identified the relevant parameters, we propose to directly optimize task performance by using the success rate on representative mobile manipulation tasks as a scoring function. This is made possible by the use of a powerful reinforcement learning policy, N^2M^2 [7], that trains a mobile manipulator to enable arbitrary end-effector motions. Importantly, this approach was shown to generate effective whole-body motions across a wide range of robot kinematics and tasks [8]–[11] and can be trained without an expensive physics simulator. In an outer loop, we maintain a distribution over possible designs and use Bayesian Optimization and HyperBand (BOHB) [12] to efficiently sample promising regions of the high-dimensional design space. In contrast to existing work that largely focuses on heuristic scoring functions or optimizes for very specific tasks [4], [6], [13], we hypothesize that this direct optimization for task performance with a general policy will result in designs and policies that generalize to the diverse tasks and

* Equal contribution. All authors are with the Department of Computer Science, University of Freiburg, Germany.

Manuscript received: December 22, 2024; Revised April 1, 2025; Accepted May 12, 2025.

This paper was recommended for publication by Editor Markus Vincze upon evaluation of the Associate Editor and Reviewers’ comments.

This work was partially funded by the German Research Foundation (DFG): 417962828, an academic grant from NVIDIA, and the BrainLinks-BrainTools center of the University of Freiburg.

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

environments that mobile manipulators are confronted with in human-centered environments. We evaluate the performance of the resulting designs on unseen test tasks and compare our approach to heuristic-based performance indices that allow the evaluation of a much larger number of designs but are not directly task-aligned. We find that the designs generated by all proposed methods yield significantly better performance than the default tabletop-style arm mounting. Moreover, the proposed direct task-based optimization consistently yields the best performance across different base drives and arms. We find that the designs generated (or optimized) by all the proposed methods yield significant better performance than the default tabletop-style mounting.

To summarize, this work makes the following contributions:

- We identify the arm mounting parameters as a crucial design element in common modular mobile manipulators.
- We develop the first co-design method for mobile manipulators and demonstrate significant design improvements that generalize across a wide range of tasks.
- We develop a directly task-driven objective for mobile manipulation and show that its direct optimization outperforms heuristic-based design metrics for manipulators.
- We demonstrate the methodology’s versatility across different robot models and show that the approach can improve the designs of a wide range of modular mobile manipulator platforms.
- We make the code publicly available at <https://moma-codesign.cs.uni-freiburg.de>.

II. RELATED WORK

Mobile Manipulator Designs: Mobile manipulators can take a wide range of forms. Commercial platforms often follow human-inspired designs with shoulder-mounted arms, such as the PR2 or TIAGo. However, these designs created from scratch come at high costs and details on their design decisions are not publicly available. While Wu *et al.* [14] open-source the base design of a mobile manipulator, they do not study the arm design. As a result, modular mobile manipulator designs consisting of widely available manipulator arms on existing mobile bases have gained large adoption, as shown in Fig. 2. However, these robots overwhelmingly mount manipulator arms in a tabletop configuration, not taking into account the differences between mobile manipulation and static manipulation tasks. While some different arm mounts exist, their design is driven to enable human teleoperation [15], [16] or directly given by the arm, such as the shoulder-mounting inspired ABB YuMi arms [17] without further considerations. *Concurrent Design* jointly optimizes both the physical design and control strategies to enhance the robot’s performance. A wide range of methods have been proposed to tackle this problem in various settings. Derivative-based methods have been used to optimize design and motion policies in a single nonlinear program [23]. Gradients obtained via sensitivity analysis have enabled the optimization along the manifold of the optimal motions [24] and the use of hard constraints [4]. In parallel, sampling-based methods such as Bayesian Optimization [25] and evolutionary algorithms [5], [26], [27] have been developed to improve the fitness of




Fig. 2. A wide range of mobile manipulator designs have been developed. Most commonly, they follow a modular design and mount existing manipulator arms in a tabletop configuration on existing mobile bases. References from top left to bottom right [8], [18]–[22].

functions over time. Reinforcement learning has been used to learn control policies in the inner loop [28], design parameters in the outer loop [29], or for more efficient overall co-design by learning value functions over design parameters [30], a single policy across designs [31], [32] or optimizing both design and policy parameters jointly by differentiating through the design parameters [13], [28], [33]. While these methods have been applied to complex problems, such as quadruped locomotion with high numbers of degrees of freedom, these methods have overwhelmingly focused on the optimization for single tasks such as forward locomotion [23], [26], [27], [30], [33], jumping [4], moving an object forward [24] or dexterous hand grasping [28]. Stochastic programming has incorporated optimization over an uncertain distribution of scenarios [3], [34]. Dong *et al.* [35] combine Bayesian Optimization over task and design pairs with reinforcement learning.

Within manipulation, global performance indices such as manipulability [36] are commonly used to measure the fitness of robotic arms [6], [37]–[39]. However, these indices do not directly capture task performance and can only serve as heuristics. Vaish *et al.* [40] use environmental constraints to efficiently optimize a 3 degree of freedom (DoF) arm for a writing task as measured by vertical force and horizontal velocity resolution, however only investigate fixed directional constraints. Xu *et al.* [13] individually optimize a manipulator for tasks such as flipping a box or rotating a cube. Kulz *et al.* [5] optimize complex arm designs for collision-free trajectories across a range of synthetic goal-reaching tasks. In contrast, we investigate designs of redundant mobile manipulators for tasks that require complex environment interactions such as opening articulated objects for which finding optimal policies for a single given design is already a challenging problem [7], [41]. We, therefore, conceptually follow a black-box approach that uses a decoupled inner policy optimization loop with a sampling-based outer loop to maintain a distribution over designs [27], [30]–[32], and directly optimize for task success across representative tasks instead of weaker correlated heuristics. To the best of our knowledge, this work is the first to systematically optimize mobile manipulator designs beyond the robot base [14]. AVATRINA [42] evaluate arm and gripper mounting candidates for a mobile avatar, however they focus on a small set of candidates and do not include the base mobility in the evaluation metrics.

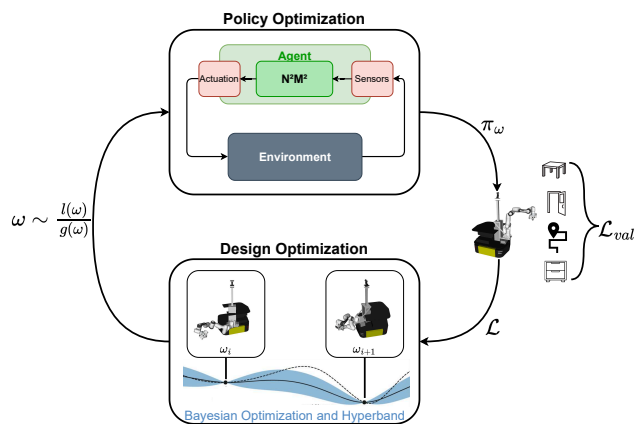


Fig. 3. MoMa-CoDesign maintains a distribution over design optimality based on Bayesian Optimization and Hyperband to propose new design parameters ω for evaluation. In an inner loop, we use a powerful reinforcement learning agent to learn a policy for this design. We then propose to directly use task-based success rates as a scoring function \mathcal{L} to evaluate the design.

III. DESIGN OPTIMIZATION FOR MOBILE MANIPULATORS

Co-design introduces a complex dependency between the design evaluation and policy optimization. Formally, our aim is to jointly optimize a morphological design $\omega \in \Omega$ over the design space Ω and a control policy $\pi_\omega(a|s)$ for this design to perform optimally as measured by the performance over a distribution of target tasks \mathcal{T} . We can define this as the optimization of

$$\omega^*, \pi_{\omega^*}^* = \arg \max_{\omega, \pi_\omega} \mathbb{E}_{\pi_\omega, \eta \sim \mathcal{T}} \mathcal{L}_\eta(\pi_\omega), \quad (1)$$

where \mathcal{L}_η is a task specific scoring function.

While this can be solved within a single optimization problem, the optimization of policies for redundant mobile manipulators is difficult and prone to local optima, as is defining unbiased evaluators that capture general task performance. In this work, we claim that, for expensive platforms such as mobile manipulators, the cost invested in the optimization and evaluation of direct task performance of strong policies is a valuable tradeoff over the use of cheaper but biased heuristics.

As such, we propose to solve this as a hierarchical decision-making problem, in which we optimize over the design in an outer loop and train a design-specific reinforcement learning policy in the inner loop. This results in an iterative algorithm that solves an MDP for a single sampled design to generate a robust performance score for the design optimization, which we term MoMa-CoDesign. An overview is shown in Fig. 3.

A. Policy Optimization

Given a design ω , our aim is to optimize a policy and produce an indicative estimate of its performance over a wide range of mobile manipulation tasks $\eta \in \mathcal{T}$. The policy is acting in a design-specific MDP $\mathcal{M}_{\omega, \eta} = \text{MDP}(\mathcal{S}_{\omega, \eta}, \mathcal{A}_\omega, \mathcal{P}_{\omega, \eta}, \mathcal{R}_\eta, \gamma)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, \mathcal{P} is the transition function, \mathcal{R}_η is the reward function, and $\gamma \in [0, 1]$ is the discount factor. As we are investigating arm mounting decisions of modular mobile manipulators acting in the same physical world, all tasks share the same underlying state space




Fig. 4. Overview of the N^2M^2 mobile manipulation policy [7]. Given an end-effector motion the robot state and a local occupancy map, a reinforcement learning agent generates controls for the base, torso and the speed of the end-effector motions with the aim to ensure their kinematic feasibility while avoiding any obstacles.

\mathcal{S} and the same physical dynamics \mathcal{P} . But different tasks may act on different objects or motions, resulting in each task acting in a different subspace of \mathcal{S} (in turn implying different dynamics as the state space differs) [43]. The design parameters furthermore impact the state space and transition functions through the robot kinematics, and potentially the action space. The objective of the reinforcement learning agent is to find the policy $\pi_\omega(a_t|s_t)$ that maximizes the expected discounted returns over the distribution of tasks

$$\mathbb{E}_{\pi_\omega, \eta \sim \mathcal{T}} \left[\sum_t \gamma^t r_{\eta, t}(s_t, a_t) \right]. \quad (2)$$

We propose to use a recent state-of-the-art multitask reinforcement learning agent, N^2M^2 [7], [9] that has been demonstrated to outperform optimization-based methods and to perform well across a wide range of tasks in the real world [8], [10], [11]. This enables us to train a single agent to perform across the set of tasks, removing the need for an additional optimization loop over tasks.

The N^2M^2 agent [7], shown in Fig. 4, receives a desired end-effector motion in the base frame of the robot. This motion consists of translational and rotational velocities \vec{v}_{ee} to the next desired pose as well as a more distant end-effector subgoal g in the form of a 6-DoF pose to indicate the longer-term plan. The agent then converts these end-effector motions to whole-body motions by generating velocities \vec{v}_b and v_{torso} for the base and torso of the robot. Finally, an inverse kinematics solver completes the motions for the remaining arm joints. The agent is also given control over the speed at which the end-effector motions are executed through a learned scaling factor $\|\vec{v}_{ee}\|$. The agent further observes the current joint positions, its previous actions and a local occupancy map to learn how to avoid obstacles. At test time, the end-effector plan can be replaced with any arbitrary system to produce motions for new tasks. We train this agent on a general random goal-reaching task in a procedurally generated obstacle environment that was shown to produce strong generalization [7]. This environment analytically integrates the velocities of the robot over time and checks for collisions of the robot base and end-effector, enabling very fast simulation without an expensive physics simulator. We then return the optimized policy π_ω^* for the sampled design to the outer loop.

B. Task-Driven Design Generation

In the outer loop, we maintain a distribution over the unknown scoring function $\mathcal{L} : \Omega \rightarrow \mathbb{R}$ that describes the

optimality of a given design ω . We model the distribution $p(\mathcal{L}|D)$ for observed data points $D = \{\omega_i, \mathcal{L}_i\}$ with the Bayesian Optimization and HyperBand (BOHB) algorithm. BOHB combines the advantages of Bayesian optimization’s model-based search with HyperBand’s efficient resource allocation strategy, making it particularly effective for hyperparameter tuning in high-dimensional design spaces with constrained computational budgets [12] while achieving strong final performance. Instead of directly modeling $p(\mathcal{L}|D)$, BOHB maintains a parzen kernel density estimator over $l(\omega) = p(\mathcal{L} < \alpha|\omega)$ and $g(\omega) = p(\mathcal{L} > \alpha|\omega)$ to calculate the expected improvement. Hyperband then decides on the resource allocation per design by selecting configurations for Successive Halving [44]. It thereby evaluates many designs with a small budget and increases the budget for the best designs [12].

The BOHB algorithm iteratively updates its distribution and resamples the agent’s design parameters to optimize the scoring function. Once new design parameters are generated, they are written to the robot’s definition files. The reinforcement learning agent is then relaunched with the updated robot design, entering a new cycle of evaluation and optimization. This continuous feedback loop ensures progressive enhancement of the agent’s capabilities in performing the mobile manipulation tasks. We propose to directly use task performance as measured by success rates over a set of representative tasks as a scoring function to measure the general performance of design and policy. Given a set of representative mobile manipulation tasks \mathcal{T} , we split it into training tasks \mathcal{T}_{train} , validation tasks \mathcal{T}_{val} and held-out test tasks \mathcal{T}_{test} . The validation tasks are used to approximate the generalization performance during design optimization, while the test tasks are only used for the evaluation of the final design and policy. This results in a scoring function

$$\mathcal{L}_{val} = \frac{1}{N} \sum_{\eta \in \mathcal{T}_{val}} \nu_{\eta}, \quad (3)$$

where ν_{η} is the average success rate on task η . We follow N²M² and define the success rate as completing the full end-effector motion without violating a maximal translational or orientational tracking error and while avoiding any collisions with the environments.

With this, the independence of the inner and outer optimization loop enables us to decouple the scoring function \mathcal{L} from the dense reward function \mathcal{R} used for RL training. While the dense reward function for RL training is designed to provide an optimal learning signal, it may not be perfectly correlated with the sparse, final task performance.

C. Manipulability Heuristic

While our proposed direct task performance metric is highly correlated with the desired outcome, the training of a reinforcement learning agent for each configuration remains costly. Performance indices have been widely used as simple to compute metrics to summarize the quality of manipulator designs [6], [37]–[39]. While less task-specific, this allows the evaluation of a larger set of designs at the same time. As we focus on arm mounting parameters and do not optimize




Fig. 5. Arm mounting parameters of the design space Ω , cf. Sec. IV, with the rotations α and β applied in sequence.

actuators or link lengths, we create an alternative method that uses whole-body manipulability as a scoring function. Yoshikawa’s manipulability ellipsoid represents how easily the end-effector can move in any direction, with the volume of the ellipsoid serving as a scalar measure of manipulability [36]. It is defined as

$$m = \sqrt{\det(JJ^T)}, \quad (4)$$

where J is the Jacobian Matrix and J^T its transpose. We can then define global manipulability over the workspace W as

$$\mu = \frac{\int m dW}{\int dW}. \quad (5)$$

We approximate this global manipulability by calculating the manipulability across end-effector poses on a discretized 3D grid centered around the robot and use it as a scoring function in the same optimization loop. We extend the manipulability commonly used for manipulator arms to include the based links of the robot. We base the grid size on the identified task heights in Sec. V and the forward reach of an average human of roughly 0.7 m [45]. Including the size of the base, this results in a grid with x, y, z corner coordinates $(-0.2, -0.8, 0.1)$ and $(0.2, 0.8, 1.7)$ relative to the base link of the robot and six orientations per point. For each point, we find an inverse kinematics solution starting from a random configuration and then calculate the manipulability for this configuration. If no inverse kinematics solution exists, we set the manipulability to zero.

To evaluate the efficacy of this heuristic-based optimization, we subsequently train an RL agent using the N²M² algorithm for the top three designs with the highest global manipulability. This allows us to directly compare the performance of the manipulability-optimized designs with those obtained through the task-based optimization method.

IV. DESIGN PARAMETERS

We identify the mounting parameters of the existing manipulator arms on mobile robot bases as an important, but often overlooked design aspect of common modular mobile manipulator robots. These parameters have a core impact on the overall kinematic structure. Importantly, they are applicable to a wide range of robots without requiring any change of the individual components, maintaining the cost-efficiency of using existing commercial arms and bases. We identify a set of

TABLE I
TIPOVER PARAMETERS

Parameter	Weight m	Position Vector \vec{p}
Franka link 0	2.40 kg	(430 mm, 330 mm, 907 mm)
Franka link 1	2.79 kg	(489 mm, 389 mm, 907 mm)
Franka link 2	2.54 kg	(681 mm, 581 mm, 879 mm)
Franka link 3	2.25 kg	(821 mm, 721 mm, 907 mm)
Franka link 4	2.20 kg	(837 mm, 737 mm, 884 mm)
Franka link 5	2.29 kg	(1005 mm, 965 mm, 880 mm)
Franka link 6	1.35 kg	(1105 mm, 1005 mm, 880 mm)
Franka link 7	0.36 kg	(1113 mm, 1013 mm, 959 mm)
End-effector	0.71 kg	(1260 mm, 1160 mm, 928 mm)
Max. arm payload	3.00 kg	(1260 mm, 1160 mm, 928 mm)
Ewellix TLT Tower	22.50 kg	(300 mm, 200 mm, 280 mm)
Ridgeback Base	135.00 kg	(0 mm, 0 mm, 140 mm)

The position vector is the x,y,z coordinate of the COM of the individual component relative to the center of the base, for the configuration with the furthest possible distance from the center (cf. Sec. IV).

main mounting parameters based on expert knowledge. These parameters are shown in Fig. 5 and consist of the following:

Arm Pitch α : forward tilt of the arm base link, ranging from upright (0°) to straight forward (90°).

Arm Yaw β : yaw angle of the arm base link, ranging from -90° to 90° . Applied after rotation α , thereby allowing, e.g., mounting of the arm on the side of the tower.

End-Effector Pitch ρ : mounting angle of the end effector at the end of the manipulator's arm, ranging from forward (0°) to downward facing (90°).

Forward translation x : mounting point of the tower lift link along the x-axis of the robot. We simultaneously shift the position of the arm base link by the opposite value, keeping its absolute position on the robot base the same. This impacts self-collisions between the arm and tower by changing the spacing between them. Limited to -5 cm to 15 cm.

Lateral translation y : translation of the tower lift link along the y-axis of the robot. Constrained to the robot base width of -20 cm to 20 cm.

Tower Yaw ϕ : yaw angle of the tower lift link, ranging from -90° to 90° .

This results in a six-dimensional optimization space $\Omega \in \mathbb{R}^6$. Note that while in this work we focus on arm mounting parameters, the approach is agnostic to the nature of these parameters and applicable to a much wider range of parameters, such as joint structure or link lengths.

Physical Feasibility

As we change the weight distribution of the platforms, we have to ensure the physical feasibility of the resulting designs. The main risk is tipping over if the arm is mounted with a much larger forward reach. We calculate the risk of tipping over both statically and dynamically. We calculate this for the design with the most extreme weight distribution, where the arm is mounted at the edge of the base, with tower at maximum height, the base rotated 45° and fully stretched out forward, and carrying the maximum possible payload. The parameters of each component are shown in Tab. I.

Statically, we approximate the stability of the arm configuration by assuming a point mass at the center of each component.

We can then calculate the center of mass as follows:

$$\text{COM} = \frac{\sum m_i \times \vec{p}_i}{\sum m_i} = (132 \text{ mm}, 109 \text{ mm}). \quad (6)$$

With the four wheels located at $(\pm 319 \text{ mm}, \pm 276 \text{ mm})$, the center of mass lies clearly within the area spanned by the wheels, remaining statically stable. Dynamically, we approximate the case when the robot drives at full speed and brakes abruptly. The critical torque is

$$\tau_{critical} = m_{base} \times g \cdot d = 231.76 \text{ N m}, \quad (7)$$

where m_{base} is the mass of the base, g is the weight force and $d = 0.319 \text{ m} - 0.144 \text{ m}$ is the distance between the rotational axis located at the front wheels of the robot and the center of mass of the robot, as calculated above¹. The maximum torque exhibited by the system, $\tau_{max} = \tau_{acc} + \tau_{grav}$, consists of the acceleration torque and the gravitational torque:

$$\tau_{acc} = \sum_l m_l \cdot a \cdot z_l \cdot \cos(\beta_l) = 21.6 \text{ N m}, \quad (8)$$

$$\tau_{grav} = \sum_l m_l \cdot g \cdot r_l \cdot \sin(\alpha_l) = 136.8 \text{ N m}, \quad (9)$$

where l are the different components given by the arm links and the payload. Respectively, m_l are their masses, r_l the distances from the pivot point to their centers of mass, z_l their heights, α_l the angle from the pivot point to the center of each component, $\beta = 90^\circ - \alpha$, and $a = \frac{\Delta v}{\Delta t} = \frac{1.1 \text{ m s}^{-1}}{0.5 \text{ s}} = 2.2 \text{ m s}^{-2}$ is the negative acceleration during the braking process. We find that the maximum torque applied is therefore lower than the critical torque, $\tau_{max} < \tau_{critical}$, meaning the system remains stable even under maximal braking. Additionally, the arm pulling on a fixed object can exert an additional torque. We find that the system remains stable even in the case of exerting torques up to the safety limits of 30 N m of the arm. However, this must be carefully considered for arms with higher safety limits.

V. EXPERIMENTS

We aim to empirically answer the following questions:

- What is the impact of the identified arm mounting parameters for mobile manipulation tasks?
- How does the use of direct task metrics compare to the optimization of cheaper-to-evaluate but less correlated performance indices?
- What is the impact of the task-distribution \mathcal{T}_{val} used during optimization?

A. Tasks

We parameterize a set of representative mobile manipulation tasks to match common household settings. The tasks are designed to collectively test the robot's navigation, object manipulation, and interaction capabilities in scenarios that mimic real-world household activities. Additional details in [7].

¹Assuming the worst case design described above would place the rotational axis even further away diagonally at only one of the front wheels. To remain conservative, we continue with the lower estimate of the critical torque.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Random Goal: The robot starts at a random position and has to follow a linear end-effector motion from its initial end-effector pose to a random goal pose on an empty map. We draw goals in the range of the human vertical accessibility [46] at the height of 0.1 m to 1.7 m (defined for the wrist frame of the robot, resulting in gripper tip poses reaching down to the floor). At the edges of the workspace, we restrict the goal orientations to a pitch matching a top or bottom grasp, respectively.

Random Obstacle (RO): The robot is placed in a procedurally generated obstacle map, requiring navigation in narrow, cluttered environments while following end-effector motions interpolated from an A*-planner towards similar random goals. This task was developed to generate highly diverse scenarios, and to learn general behaviors.

Pick and Place: The robot has to pick up an object from one table and place it on another table. The pick and place heights are set to the same values as for the random goals, with orientation constraints for top, and bottom grasps at the edges of the workspace height.

Open Door: The robot has to follow end-effector motions inferred from human demonstrations [47] to navigate to a door, press the handle down, and open it while passing through the narrow door frame.

Open Drawer: The robot has to open a drawer with a prismatic joint placed at typical kitchen heights of 0.4 m to 1.2 m [46], following motions from human demonstrations.

Open Cabinet: Open a cabinet door with a revolute joint, located at typical kitchen heights of 0.4 m to 1.7 m [46], following motions from human demonstrations.

For the Pick and Place, Open Door, Open Drawer and Open Cabinet tasks we also randomly place an obstacle in the path between the robot and the task object, requiring the robot to complete the motions from different directions. The motions from human demonstrations adapt to different object poses through their parametrization in an object-centric frame [47]. The tasks are shown in the video.

B. Robots and Optimization Setup

To test the effectiveness across mobile manipulator platforms, we evaluate all approaches for two drive types, omnidirectional and differential, and two arms, the 7 Degree of Freedom (DoF) Franka and the 6 DoF UR5 arm. To evaluate the impact of these different modules, we substitute them into the same base design of the FMM robot, shown in Fig. 1. The default configuration of this robot consists of an omnidirectional base, a height-adjustable Ewellix column to change the height, and a Franka arm.

All hyperparameters of the optimization process are reported in Tab. II. We set the budget for Hyperband between 300,000 and 1M training steps. These values are chosen based on the average success rate of the original N²M² agent [7]. Approximately after 200,000 training steps, the curve starts to rise, converging at around 700,000 steps. We set the boundaries optimistically to make sure we find the best performing agents. The parameters of the N²M² agent are set according to [7], which were shown to generalize across robot platforms.

TABLE II
BOHB ALGORITHM PARAMETERS.

Parameter	Value
Iterations	20
Sampled Designs	60
Reduction Factor η	3
Random Fraction ρ	1/3
Minimum Budget b_{min} (steps)	300,000
Maximum Budget b_{max} (steps)	1,000,000

All other parameters are set to the default values of the HpBandSter framework [12].

C. Models

Tabletop Mount: The default tabletop-based configuration of the arm without any design optimization.

Manipulability: Co-design based on the global manipulability as described in Sec. III-C.

MoMa-CoDesign (RO Task): Our approach with $\mathcal{T}_{train} = \mathcal{T}_{val} = \{\text{Random Obstacle}\}$. To ensure the designs learn to pass narrow passages, we reduce the spacing between obstacles during training to 1.7 m.

MoMa-CoDesign (All Tasks): Our approach with $\mathcal{T}_{train} = \{\text{Random Obstacle}\}$ and $\mathcal{T}_{val} = \mathcal{T}_{test} = \{\text{All tasks}\}$, equally weighted.

We run all methods with the same time budget of roughly eight days on a server with a Nvidia A40 GPU and AMD EPYC 7513 processor, resulting in the evaluation of 60 designs for MoMa-CoDesign and 2,500 designs for the cheaper manipulability approach. Note that these compute costs remain negligible compared to the hardware cost of the platforms.

D. Experimental Results

We evaluate the resulting policies and designs over 100 episodes per test task and report the success rates in Tab. III. All approaches achieve improvements over the original tabletop mounting design, confirming the importance of the arm mounting parameters. MoMa-CoDesign achieves the best results across platforms, more than doubling success rates for the Franka arm. In contrast, while able to sample many more designs in the same time, the designs generated by the manipulability heuristic do not reach the same level of improvement. We report the manipulability for all final designs on the right of Tab. III. The numbers confirm our hypothesis that the performance index is only weakly correlated with task performance, with a Pearson correlation coefficient of $\rho = 0.58$. Moreover, the definition of the workspace over which to evaluate the metric further biases the designs, e.g., a symmetric workspace induces a prior to cover all directions. While a control policy could move the base to account for one-sided designs, the manipulability can only take into account the local impact of the base on the Jacobian, not larger repositionings.

Comparing the evaluation of MoMa-CoDesign validated on all tasks versus the general random obstacle task, we find that the random obstacle task results in significantly better designs for the two Franka-based platforms, with small advantages of incorporating all tasks for the UR5 arm. We hypothesize that this effect is partially caused by negative correlations between

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE III
EVALUATION OF THE GENERATED DESIGNS ACROSS TASKS WITH STANDARD ERROR.

Drive Arm	Model	Random Obstacle	Random Goal	Pick & Place	Door	Drawer	Cabinet	Average Success	Manipulability	
Omni	Franka	Tabletop Mount	30 ± 4.6	54 ± 5.0	39 ± 4.9	6 ± 2.4	36 ± 4.8	26 ± 4.4	31.8 ± 1.9	1.15
		Manipulability	54 ± 5.0	73 ± 4.4	33 ± 4.7	18 ± 3.8	54 ± 5.0	53 ± 5.0	47.5 ± 2.0	1.45
		MoMa-CoDesign (RO task)	66 ± 4.7	91 ± 2.9	51 ± 5.0	57 ± 5.0	83 ± 3.8	88 ± 3.2	72.7 ± 1.8	1.18
		MoMa-CoDesign (all tasks)	44 ± 5.0	74 ± 4.4	30 ± 4.6	66 ± 4.7	74 ± 4.4	66 ± 4.7	59.0 ± 2.0	0.88
Differential	Franka	Tabletop Mount	16 ± 3.7	38 ± 4.9	6 ± 2.4	1 ± 1.0	15 ± 3.6	5 ± 2.2	13.5 ± 1.4	0.64
		Manipulability	41 ± 4.9	84 ± 3.7	36 ± 4.8	24 ± 4.3	47 ± 5.0	10 ± 3.0	40.3 ± 2.0	1.45
		MoMa-CoDesign (RO task)	57 ± 5.0	92 ± 2.7	37 ± 4.8	31 ± 4.6	59 ± 4.9	49 ± 5.0	54.2 ± 2.0	1.33
		MoMa-CoDesign (all tasks)	23 ± 4.2	75 ± 4.3	23 ± 4.2	55 ± 5.0	25 ± 4.3	25 ± 4.3	37.6 ± 2.0	1.19
Omni	UR5	Tabletop Mount	26 ± 4.4	46 ± 5.0	9 ± 2.9	4 ± 2.0	26 ± 4.4	34 ± 4.7	24.2 ± 1.7	0.05
		Manipulability	28 ± 4.5	57 ± 5.0	9 ± 2.9	15 ± 3.6	32 ± 4.7	33 ± 4.7	29.0 ± 1.9	0.53
		MoMa-CoDesign (RO task)	35 ± 4.8	47 ± 5.0	17 ± 3.8	16 ± 3.7	28 ± 4.5	17 ± 3.8	26.7 ± 1.8	0.18
		MoMa-CoDesign (all tasks)	24 ± 4.3	44 ± 5.0	22 ± 4.1	19 ± 3.9	32 ± 4.7	34 ± 4.7	29.2 ± 1.9	0.11

Left: success rates in %, best highlighted in bold. Right: average manipulability of the designs as defined in Sec. III-C. RO task: optimization on the random obstacle task. ± indicates estimated standard errors.




Fig. 6. Designs generated by the different methods across the different robots. Note that we do not render connecting links between tower and arm and between arm and end-effector.

certain tasks, leading to a more fragmented design optimality function that would require more data to discover global optima. It furthermore relates to the generality of the random obstacle task, which was designed for diverse experience, in contrast to the application-driven downstream tasks. Nevertheless, it also points to further possible improvements in the design sampling approach. Improvements are generally smaller for the UR5 arm. This hints at the configuration space already being better aligned with the mobile manipulation tasks than the Franka arm, which was specifically designed for tabletop tasks. At the same time, overall performance on the tasks also remains clearly below the Franka arm, and so does the average manipulability. This demonstrates the importance of the additional DoF of the Franka arm. We also evaluate the share of steps in which the arm is in a near singular configuration. We find no significant increases, except for the cabinet task of the RO trained agents with the Franka arm. However, the success rates remain by far the highest, indicating that the agent is successfully using the mobility of the base to compensate for limitations of the arm.

The designs generated by the different approaches are shown in Fig. 6. We find that all designs prefer to tilt the arm forward and place the tower towards a corner of the base, increasing its




Fig. 7. Manipulability across the workspace of the design for the omni-directional Franka robot. Left: design generated by optimizing the manipulability, right: design generated by MoMa-CoDesign (SO task). Green indicates high manipulability, and red is low. Points with zero manipulability are omitted. Each point depicts the average over six orientations.

reach. The resulting designs differ strongly from the tabletop mounting configuration, again highlighting the importance of this design factor. At the same time, the orientations of the arm significantly differ between the manipulability-optimized designs and MoMa-CoDesign. Furthermore, we find that MoMa-CoDesign converged to similar designs across the four Franka arm configurations, showing that the design optimization process moved towards a similar optimum. We show the designs performing the different tasks in the accompanying video.

Fig. 7 shows the manipulability for the design generated by the manipulability method. As we can see, the resulting design covers large areas of the workspace. However, as we showed previously, this does not directly correlate with the required workspace and motions for diverse mobile manipulation tasks.

VI. CONCLUSION

In this work, we identify arm mounting parameters as an important design factor for modular mobile manipulation platforms. We showed that default tabletop-like mounting configurations are suboptimal for a wide range of mobile manipulation tasks that often require different workspaces and motions at different heights. To address these problems, we presented a co-design approach that optimizes a powerful reinforcement learning agent directly on task metrics. We then empirically demonstrated the approach across different base and arm modules. In comparison to performance heuristics, we find that the designs generated by our method lead to significantly higher performance, as heuristics are only weakly correlated

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

with the objective of interest, and as a result, even with the evaluation of a much larger number of designs, converge to worse optima. Finally, we open-sourced the resulting designs and approach to enable the community to test and generate alternative designs for different base and arm modules. In future work, we plan to investigate the implementation of the resulting designs in the real world, as well as the use of heuristic-based methods to initialize the search of the presented co-design approach through computationally cheaper methods.

REFERENCES

- [1] S. Doncieux, N. Bredeche, J.-B. Mouret, and A. E. Eiben, "Evolutionary robotics: what, why, and where to," *Frontiers in Robotics and AI*, vol. 2, 2015.
- [2] J. C. Bongard, "Evolutionary robotics," *Communications of the ACM*, vol. 56, no. 8, 2013.
- [3] G. Bravo-Palacios, A. Del Prete, and P. M. Wensing, "One robot for many tasks: Versatile co-design through stochastic programming," *IEEE Rob. and Auto. Letters*, vol. 5, no. 2, 2020.
- [4] T. Dinev, C. Mastalli, V. Ivan, S. Tonneau, and S. Vijayakumar, "A versatile co-design approach for dynamic legged robots," in *Int. Conf. on Intelligent Robots and Systems*, 2022.
- [5] J. Külz and M. Althoff, "Optimizing modular robot composition: A lexicographic genetic algorithm approach," in *Int. Conf. on Robotics & Automation*, 2024.
- [6] S. Hwang, H. Kim, Y. Choi, K. Shin, and C. Han, "Design optimization method for 7 dof robot manipulator using performance indices," *Int. Journal of Precision Engineering and Manufacturing*, vol. 18, 2017.
- [7] D. Honerkamp, T. Welschhold, and A. Valada, "N²m²: Learning navigation for arbitrary mobile manipulation motions in unseen and dynamic environments," *IEEE Trans. on Robotics*, vol. 39, no. 5, 2023.
- [8] D. Honerkamp, H. Maheshka, J. O. von Hartz, T. Welschhold, and A. Valada, "Whole-body teleoperation for mobile manipulation at zero added cost," *IEEE Rob. and Auto. Letters*, vol. 10, no. 4, pp. 3198–3205, 2025.
- [9] D. Honerkamp, T. Welschhold, and A. Valada, "Learning kinematic feasibility for mobile manipulation through deep reinforcement learning," *IEEE Rob. and Auto. Letters*, 2021.
- [10] D. Honerkamp, M. Buchner, F. Despinoy, T. Welschhold, and A. Valada, "Language-grounded dynamic scene graphs for interactive object search with mobile manipulation," *IEEE Rob. and Auto. Letters*, 2024.
- [11] F. Schmalstieg, D. Honerkamp, T. Welschhold, and A. Valada, "Learning hierarchical interactive multi-object search for mobile manipulation," *IEEE Rob. and Auto. Letters*, 2023.
- [12] S. Falkner, A. Klein, and F. Hutter, "Bohb: Robust and efficient hyperparameter optimization at scale," in *Int. Conf. on Mach. Learning*, 2018.
- [13] J. Xu, T. Chen, L. Zlokapa, M. Foshey, W. Matusik, S. Sueda, and P. Agrawal, "An end-to-end differentiable framework for contact-aware robot design," *Robotics: Science and Systems*, 2021.
- [14] J. Wu, W. Chong, R. Holmberg, A. Prasad, Y. Gao, O. Khatib, S. Song, S. Rusinkiewicz, and J. Bohg, "Tidybot++: An open-source holonomic mobile manipulator for robot learning," *Proc. of the Conf. on Rob. Learning*, 2024.
- [15] M. Schwarz, C. Lenz, R. Memmesheimer, B. Pätzold, A. Rochow, M. Schreiber, and S. Behnke, "Robust immersive telepresence and mobile telemanipulation: Nimbro wins ana avatar xprize finals," in *Int. Conf. on Humanoid Robots*, 2023.
- [16] K. Hauser, N. Watson, Eleanor, J. Bae, J. Bankston *et al.*, "Analysis and perspectives on the ana avatar xprize competition," *Int. Journal of Social Robotics*, pp. 1–32, 2024.
- [17] K. Blomqvist, M. Breyer, A. Cramariuc, J. Förster, M. Grinvald, F. Tschopp, J. J. Chung, L. Ott, J. Nieto, and R. Siegwart, "Go fetch: Mobile manipulation in unstructured environments," *arXiv preprint arXiv:2004.00899*, 2020.
- [18] KUKA AG. Kmr iiwa. [Online]. Available: <https://www.kuka.com/en-de/products/amr-autonomous-mobile-robotics/mobile-robots/kmr-iiwa>
- [19] H. Xiong, R. Mendonca, K. Shaw, and D. Pathak, "Adaptive mobile manipulation for articulated objects in the open world," *arXiv preprint arXiv:2401.14403*, 2024.
- [20] H. Raei, J. M. Gandarias, E. De Momi, P. Balatti, and A. Ajoudani, "A multipurpose interface for close- and far-proximity control of mobile collaborative robots," in *Int. onf. for Biomedical Robotics and Biomechanics*, 2024, pp. 457–464.
- [21] C.-C. Hsu, B. Abbatematteo, Z. Jiang, Y. Zhu, R. Martín-Martín, and J. Biswas, "Kinscene: Model-based mobile manipulation of articulated scenes," *arXiv preprint arXiv:2409.16473*, 2024.
- [22] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser, "Tidybot: Personalized robot assistance with large language models," *Autonomous Robots*, vol. 47, no. 8, 2023.
- [23] A. Spielberg, B. Araki, C. Sung, R. Tedrake, and D. Rus, "Functional co-optimization of articulated robots," in *Int. Conf. on Robotics & Automation*, 2017.
- [24] S. Ha, S. Coros, A. Alspach, J. Kim, and K. Yamane, "Joint optimization of robot design and motion parameters using the implicit function theorem," in *Robotics: Science and Systems*, vol. 13, 2017.
- [25] T. Liao, G. Wang, B. Yang, R. Lee, K. Pister, S. Levine, and R. Calandra, "Data-efficient learning of morphology and controller for a microrobot," in *Int. Conf. on Robotics & Automation*, 2019.
- [26] K. M. Digumarti, C. Gehring, S. Coros, J. Hwangbo, and R. Siegwart, "Concurrent optimization of mechanical design and locomotion control of a legged robot," in *Mobile Service Robotics*, 2014.
- [27] A. Gupta, S. Savarese, S. Ganguli, and L. Fei-Fei, "Embodied intelligence via learning and evolution," *Nature communications*, vol. 12, no. 1, 2021.
- [28] T. Chen, Z. He, and M. Ciocarlie, "Hardware as policy: Mechanical and computational co-optimization using deep reinforcement learning," *Proc. of the Conf. on Rob. Learning*, 2020.
- [29] A. Albright and J. Vaughan, "Selecting mechanical parameters of a monopode jumping system with reinforcement learning," *arXiv preprint arXiv:2212.01303*, 2022.
- [30] K. S. Luck, H. B. Amor, and R. Calandra, "Data-efficient co-adaptation of morphology and behaviour with deep reinforcement learning," in *Proc. of the Conf. on Rob. Learning*, 2020.
- [31] C. Schaff, D. Yunis, A. Chakrabarti, and M. R. Walter, "Jointly learning to construct and control agents using deep reinforcement learning," in *Int. Conf. on Robotics & Automation*, 2019.
- [32] C. Schaff, A. Sedal, and M. R. Walter, "Soft robots learn to crawl: Jointly optimizing design and control with sim-to-real transfer," *Robotics: Science and Systems*, 2022.
- [33] L. Jackson, C. Walters, S. Eckersley, P. Senior, and S. Hadfield, "Orchid: optimisation of robotic control and hardware in design using reinforcement learning," in *Int. Conf. on Intelligent Robots and Systems*, 2021.
- [34] G. Bravo-Palacios and P. M. Wensing, "Large-scale admm-based co-design of legged robots," in *Int. Conf. on Intelligent Robots and Systems*, 2022.
- [35] Y. Dong, S. Han, X. Cheng, W. Friedl, R. I. C. Muchacho, M. A. Roa, J. Tumova, and F. T. Pokorny, "Co-designing tools and control policies for robust manipulation," *arXiv preprint arXiv:2409.11113*, 2024.
- [36] T. Yoshikawa, "Manipulability and redundancy control of robotic mechanisms," in *Int. Conf. on Robotics & Automation*, vol. 2, 1985.
- [37] Y. Xiao, Z. Fan, W. Li, S. Chen, L. Zhao, and H. Xie, "A manipulator design optimization based on constrained multi-objective evolutionary algorithms," in *IEEE Conf. on industrial informatics*, 2016.
- [38] F. Molaei and S. Ghatrehsamani, "Kinematic-based multi-objective design optimization of a grapevine pruning robotic manipulator," *AgriEngineering*, vol. 4, no. 3, 2022.
- [39] Y. Feng, L. Fang, W. Bu, and J. Kang, "Multi-objective optimization for design of redundant serial robots," in *Chinese Automation Congress*, 2020.
- [40] A. Vaish and O. Brock, "Co-designing manipulation systems using task-relevant constraints," in *Int. Conf. on Robotics & Automation*, 2024.
- [41] J. Pankert and M. Hutter, "Perceptive model predictive control for continuous mobile manipulation," *IEEE Rob. and Auto. Letters*, vol. 5, no. 4, 2020.
- [42] J. M. Correia Marques, P. Naughton, J.-C. Peng *et al.*, "Immersive commodity telepresence with the avatrina robot avatar," *Int. Journal of Social Robotics*, pp. 1–29, 2024.
- [43] S. Sodhani, A. Zhang, and J. Pineau, "Multi-task reinforcement learning with context-based representations," in *Int. Conf. on Mach. Learning*, 2021.
- [44] K. Jamieson and A. Talwalkar, "Non-stochastic best arm identification and hyperparameter optimization," in *Artificial intelligence and statistics*, 2016.
- [45] J. Looker and T. Garvey, "Reaching for holograms," in *Proc. from International Design Congress*, 2015.
- [46] J. Hrovatin, S. Prekrat, L. Oblak, and D. Ravnik, "Ergonomic suitability of kitchen furniture regarding height accessibility," *Collegium antropologicum*, vol. 39, no. 1, 2015.
- [47] T. Welschhold, C. Dornhege, and W. Burgard, "Learning mobile manipulation actions from human demonstrations," in *Int. Conf. on Intelligent Robots and Systems*, 2017.