

RoboDuet: Learning a Cooperative Policy for Whole-body Legged Loco-Manipulation

Guoping Pan, Qingwei Ben, Zhecheng Yuan, Guangqi Jiang, Yandong Ji,
Shoujie Li, Jiangmiao Pang, Houde Liu, *Member, IEEE*, Huazhe Xu

Abstract—Fully leveraging the loco-manipulation capabilities of a quadruped robot equipped with a robotic arm is non-trivial, as it requires controlling all degrees of freedom (DoFs) of the quadruped robot to achieve effective whole-body coordination. In this letter, we propose a novel framework RoboDuet, which employs two collaborative policies to realize locomotion and manipulation simultaneously, achieving whole-body control through mutual interactions. Beyond enabling large-range 6D pose tracking for manipulation, we find that the two-policy framework supports zero-shot transfer across quadruped robots with similar morphology and physical dimensions in the real world. Our experiments demonstrate that RoboDuet achieves a 23% improvement in success rate over the baseline in challenging loco-manipulation tasks employing whole-body control. To support further research, we provide open-source code and additional videos on our website: locomanip-duet.github.io.

Index Terms—Legged robot, whole-body control, loco-manipulation, reinforcement learning.

I. INTRODUCTION

Mobile robots have increasingly been deployed to assist humans and demonstrated remarkable capabilities [1], [2]. Typically, these robots are equipped with wheeled mobile bases, making them less adaptable to diverse terrains and unable to adjust their base postures. This limitation has sparked interest in developing legged robots to undertake manipulation tasks, offering enhanced versatility and adaptability in diverse environments. By employing whole-body control in legged robots and robotic arms, it is possible to effectively overcome terrain constraints and significantly expand the manipulation workspace of the arms [3], [4], [5]. However, training a legged loco-manipulation policy to achieve whole-body control, along

This work was supported by Institute for Interdisciplinary Information Sciences, Tsinghua University, Shanghai Qi Zhi Institute and Shanghai Artificial Intelligence Laboratory. (Guoping Pan and Qingwei Ben are co-first authors.) (Corresponding author: Huazhe Xu.)

Guoping Pan, Qingwei Ben, Zhecheng Yuan and Shoujie Li are with Tsinghua University, Beijing 100084, China. (email: ppg23@mails.tsinghua.edu.cn, elgceben@gmail.com, yuanzc23@mails.tsinghua.edu.cn, lsj20@mails.tsinghua.edu.cn)

Guangqi Jiang and Yandong Ji are with University of San Diego, California 92093, USA. (email: gqjiang@ucsd.edu, ydji1024@gmail.com)

Jiangmiao Pang is with Shanghai Artificial Intelligence Laboratory, Shanghai 200032, China. (email: pangjiangmiao@gmail.com)

Houde Liu is with Tsinghua Shenzhen International Graduate School, Shenzhen 518055, China, and also with Jianghuai Advance Technology Center, Hefei 230000. (email: liu.hd@sz.tsinghua.edu.cn)

Huazhe Xu is with the Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China, and also with Shanghai Qi Zhi Institute, Shanghai 200030, China, as well as with Shanghai Artificial Intelligence Laboratory, Shanghai 200032, China. (email: huazhe_xu@mails.tsinghua.edu.cn)




Fig. 1. RoboDuet is a framework that affords loco-manipulation and zero-shot transfer on morphologically and dimensionally similar quadruped robots. Top row: zero-shot transfer. This feature enables control across six configurations generated by the permutation of two robotic arms and three quadruped robots, allowing for the replacement of quadruped robots without the need to retrain the entire system. Middle row: loco-manipulation. From left to right, the robot walks to pick up a small ball on grass, grasps a doll from a high table in a café, grabs a bottle on lower stairs, and picks up a cup from an office desk. Bottom row: whole-body control. Given multiple target end-effector poses, the robot can adjust its entire body posture while moving to reach the desired targets and maintain stability.

with accurate 6D pose tracking capabilities, presents a substantial challenge to researchers.

As a pioneering effort in this domain, Fu *et al.* [3] has utilized a unified control policy to accomplish coordinated manipulation and locomotion. Despite the implementation of a whole-body control framework, it cannot tackle accurate 6D end-effector pose tracking, a capability that is crucial for manipulation tasks. On the other hand, while GAMMA [4] and GeFF [5] are capable of grasping objects based on 6-DoF end-effector control, their operation strategies separate the arm and the quadruped systems, thus falling short of achieving whole-body control. This distinction restricts the workspace of the arm. Consequently, accomplishing large-range manipulation tasks across the entire workspace necessitates a novel training paradigm. This approach must ensure enhanced coordination between the quadruped and the manipulator arm while also improving training efficiency and generalization capabilities.

In awareness of these challenges, we introduce the **RoboDuet**: an integrated legged loco-manipulation framework tailored for large-range 6D pose tracking and whole-body control. As shown in Fig. 1, RoboDuet endows the policies with the capabilities of zero-shot transfer across quadruped

robots with similar morphology and physical dimensions, locomanipulation for diverse tasks, and robust whole-body control. These capabilities greatly enhance the system’s adaptability to different environments and reduce the need for retraining.

To simplify the whole-body control problem, we adopt a cooperative mechanism, achieved through the coordinated collaboration of a locomotion policy and an arm policy. The interaction between the locomotion policy and the arm’s actions exhibits a duet-like performance, where the locomotion policy utilizes the actions of the arm as guidance to adjust its posture, while the arm is complemented by the actions of the locomotion policy aiming to expand the robot’s workspace. The training process for RoboDuet is structured in two stages, as illustrated in Fig. 2. In stage 1, we develop the locomotion policy to endow the legged robot with essential mobility capabilities. Following stage 1, stage 2 involves training the arm policy that can coordinate with the locomotion policy. We argue that employing a two-stage training strategy enhances the stability of the training process, resulting in the acquisition of agile and large-range 6D pose tracking agents.

Our contributions are summarized as follows:

- We propose a novel RL-based whole-body control framework for quadrupedal loco-manipulation that enables 6D end-effector pose tracking during locomotion.
- We introduce a cooperative mechanism that decouples multi-task learning, allowing the two policies to collaborate while maintaining focus and resulting in enhanced loco-manipulation capabilities compared to the unified policy.
- Our approach achieves at least a 23% improvement in success rate for challenging mobile manipulation tasks over model-based controllers and demonstrates zero-shot transfer capabilities across quadruped robots with a 2.7 kg increase in weight.

II. RELATED WORKS

In recent years, the capabilities of legged robot locomotion have advanced rapidly, particularly in traversing complex terrains [6], [7], thereby facilitating their integration into diverse and challenging application domains. Furthermore, given the redundant degrees of freedom of quadruped robots compared to wheeled or tracked platforms, there are growing research efforts focused on whole-body control for legged robots with manipulators to achieve loco-manipulation capabilities. The current technological landscape features three primary strategies. The first approach leverages control-based techniques such as model predictive control (MPC), which explicitly consider the system’s kinematic and dynamic properties, making them well-suited for scenarios involving dense contact forces. However, these methods often entail substantial engineering effort and exhibit inherent limitations in adaptability and robustness when applied to unknown environments [8], [9], [10]. The second strategy adopts learning algorithms to generate high-level commands for legged robot or robotic arm, which are then translated into low-level joint control instructions based on built-in controllers or inverse kinematics (IK) [4], [5], [11], [12]. However, these approaches lack

effective coordination between the legged platform and the arm or suffer from infeasible IK solutions, thus failing to maximize the potential for coordinating the pose of the legged robot to extend the operational workspace of the arm. The third approach leverages deep reinforcement learning (DRL) to realize whole-body control. In earlier research, a pipeline was developed to integrate a model-based manipulator with an RL-based locomotion policy, demonstrating robustness against external force disturbances induced by the manipulator [13]. However, the locomotion policy focused solely on resisting external disturbances, lacking the ability to achieve whole-body coordination. Later study proposed training a unified policy to control the entire system for achieving whole-body control, but this approach were limited to achieving only end-effector position tracking [3]. A follow-up study extends this approach to 6D task-space pose tracking [14], but it is limited by the constraints of data collection. The most recent research leverages keypoint tracking to train a whole-body end-effector pose tracking policy, which requires an additional policy to provide locomotion capabilities [15]. This indicates that the method lacks the ability to perform simultaneous mobility and manipulation, which is crucial for executing dynamic tasks such as door opening. Consequently, there is a clear need for innovative frameworks that can fully harness the locomotive advantages of quadruped robots while ensuring seamless coordination between the upper arm and the lower legged platform.

III. METHODS

A. Cooperative Policy for Whole-body Control

RoboDuet consists of a loco policy for locomotion and an arm policy for manipulation. The two policies are harmonized as a whole-body controller. Specifically, the loco policy adjusts its actions by following the instructions of the arm policy.

1) **Loco policy:** The goal of the loco policy π_{loco} is to follow a target command $\mathbf{c}_t = (v_x^{cmd}, v_y^{cmd}, \omega_{yaw}^{cmd}, \phi_{pitch}^{cmd}, \phi_{roll}^{cmd})$, where $v_{x,y}^{cmd} = (v_x^{cmd}, v_y^{cmd})$ is the desired linear velocity in base frame along x- and y- axes, ω_{yaw}^{cmd} is the desired angular velocity in yaw axis, $\phi^{cmd} = (\phi_{pitch}^{cmd}, \phi_{roll}^{cmd})$ denotes the desired pitch and roll angles of the base which plays an important role in achieving whole-body control and many downstream tasks. The observation of loco policy o_t^{loco} contains leg states $s_t^{leg} \in \mathbb{R}^{24}$ (leg joint positions and velocities), base states ϕ_t (roll and pitch angles), target commands \mathbf{c}_t , clock time t_t , and last leg action $a_{t-1}^{leg} \in \mathbb{R}^{12}$. The leg action a_t^{leg} represents joint position offsets that are added to the default joint positions to specify the target positions for the twelve leg joint motors.

2) **Arm policy:** The goal of the arm policy π_{arm} is to accurately track the 6D end-effector pose. The observation of arm policy o_t^{arm} is composed of arm states $s_t^{arm} \in \mathbb{R}^{12}$ (arm joint positions and velocities), target end-effector pose $\chi_t \in \mathbb{R}^9$, base states ϕ_t , and last arm action $a_{t-1}^{arm} \in \mathbb{R}^8$. The actions of the arm policy consist of two parts: the first six actions $a_t^{arm,j} \in \mathbb{R}^6$ represent the target joint position offsets corresponding to six arm joint actuators. Then, it will be concatenated with the output of the loco policy to achieve synchronous control of the overall system. Notably, the position targets are

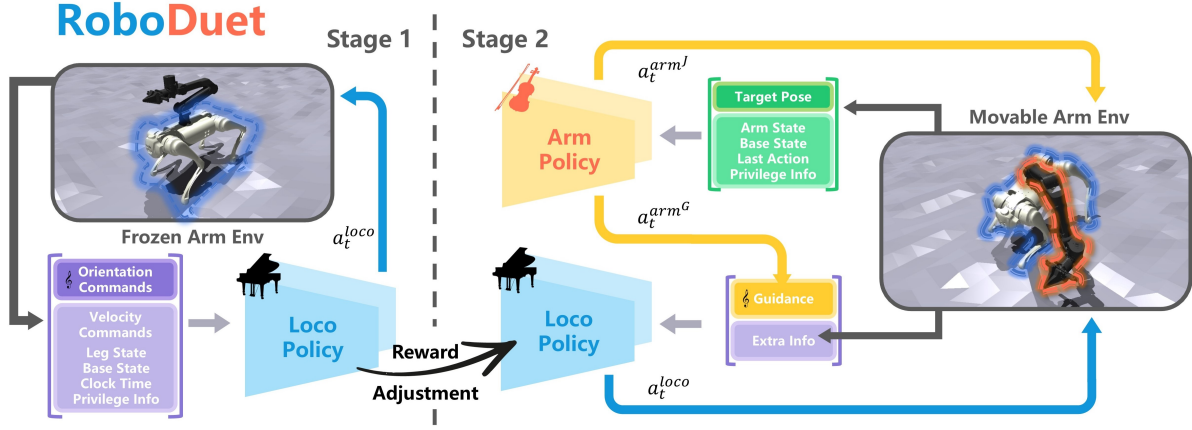


Fig. 2. Overview of RoboDuet. In Stage 1, the loco policy is trained with fixed robotic arm, enabling the quadruped robot to achieve robust locomotion. In Stage 2, the loco and arm policies are trained simultaneously in a cooperative manner. The loco policy from stage 1 is reused, but the original body orientation commands are replaced by guidance signals $a_t^{arm^G}$ generated by the arm policy, enabling coordinated whole-body control. During the transition between the two stages, reward adjustment is introduced to effectively leverage locomotion priors, facilitating seamless integration for whole-body loco-manipulation.

tracked using a proportional-derivative controller. To expand the manipulation workspace with whole-body control, the rest part of the arm policy $a_t^{arm^G} = (a_t^{arm^G_p}, a_t^{arm^G_r}) \in \mathbb{R}^2$ is used to replace ϕ^{cmd} , providing additional degrees of freedom for pose tracking to cooperate with the loco policy. Due to the distinct focuses of the two policies, only essential information is shared to maintain a favorable cooperative relationship while mitigating disruptions caused by factors such as differences in task learning efficiency.

B. Two-stage Training

To ensure robust loco-manipulation, we employ a two-stage training strategy. Our method enables a seamless transition between these stages by maintaining consistent input and output dimensions for the policies throughout the entire process. Since the primary tasks differ across the stages—stage 1 focuses on locomotion, while stage 2 emphasizes manipulation—we introduce reward adjustment to integrate the latter without compromising the performance of the former.

1) **Stage 1:** This stage focuses on obtaining the robust locomotion capability. To ensure that the loco policy adapt to the center of mass and the inertia offset of the whole robot throughout the entire training process, we keep all the arm joints fixed at their default positions $[0, 0.8, 0.8, 0, 0, 0]$. In this stage, the arm policy is inactive, and the target end-effector pose χ_t is set to zero. Inspired by the powerful blind locomotion algorithm [16], we similarly apply a vector of behavior parameters \mathbf{b}_t to represent a heuristic gait reward, which is expressed as follows:

$$\mathbf{b}_t = [\theta_1^{cmd}, \theta_2^{cmd}, \theta_3^{cmd}, \mathbf{f}^{cmd}, \mathbf{h}_z^{cmd}, \phi^{cmd}, \mathbf{s}^{cmd}, \mathbf{h}_z^{f,cmd}] \quad (1)$$

where $\theta^{cmd} = (\theta_1^{cmd}, \theta_2^{cmd}, \theta_3^{cmd})$ are the timing offsets between pairs of feet. Since our goal is to achieve pose tracking rather than diverse locomotion behaviors, we fix some gait parameters to speed up the convergence of training. In the following description, we specifically highlight the modified parts. We set the timing offsets θ^{cmd} to $[0.5, 0, 0]$ to achieve a

stable trotting gait. To enable the loco policy to recognize the rhythm of stepping, the clock time t_t is computed from the offset timings of each foot, with the mathematical definitions as follows:

$$\begin{aligned} \mathbf{t}_t &= [\sin(2\pi t^{FR}), \sin(2\pi t^{FL}), \sin(2\pi t^{RR}), \sin(2\pi t^{RL})] \quad (2) \\ [t^{FR}, t^{FL}, t^{RR}, t^{RL}] &= [t + \theta_2^{cmd} + \theta_3^{cmd}, t + \theta_1^{cmd} + \theta_3^{cmd}, \\ &\quad t + \theta_1^{cmd}, t + \theta_2^{cmd}] \quad (3) \end{aligned}$$

where t is a counter variable that advances from 0 to 1 during each gait cycle and FR, FL, RR, RL are the four feet respectively. When the base velocity is zero, the jitter caused by marching on the spot will reduce the precision of manipulation. In this situation, we set clock time \mathbf{t}_t to $[1, 1, 1, 1]$ to force all feet to maintain a stationary position. \mathbf{f}^{cmd} is the stepping frequency which is set to 3 Hz. \mathbf{h}_z^{cmd} is the body height command which is not used. $\mathbf{s}^{cmd} = (s_x^{cmd}, s_y^{cmd})$ is the foot clearance which is set to $[0.45, 0.3]$. $\mathbf{h}_z^{f,cmd}$ is the footswing height command which is set to 0.06 m.

During stage one, most components of the reward r^{loco} utilized by the loco policy are identical to those described in [16], which can be referred to for more details. However, we remove the body height tracking component $r_{\mathbf{h}_z^{cmd}}^{loco}$ to allow the body height to adjust adaptively according to the task. The additional payload introduced by the arm increases the torque demands on the leg motors, which will result in motor overheating [14]. To mitigate this, we incorporate an energy reward r_{energy}^{loco} for the quadruped robot to ensure the stability of instantaneous power [17], as formulated below:

$$r_{energy}^{loco} = -0.00004 \cdot \sum_{i \in \text{leg joints}} |\tau_i \dot{q}_i|^2 \quad (4)$$

where τ_i and \dot{q}_i are the torque and velocity of the i^{th} leg joint.

2) **Stage 2:** This stage aims to coordinate locomotion and manipulation to achieve whole-body large-range loco-manipulation. The arm policy is activated simultaneously with all the robotic arm joints. We adopt 6D target pose of end-effector as policy input. To eliminate the influence brought

by body rotation [3], we similarly use a posture-independent spherical coordinate to define the target end-effector pose χ_t . The target position of end-effector is represented by radius l , latitude p , and longitude y . The target orientation is represented using a 6-D vector [18]. To improve the accuracy of end-effector orientation tracking, we use euler angles $[roll, pitch, yaw]$ in Z-Y-X order for sampling, which can intuitively exclude many illegal postures, and convert them to included angle along each axis of the coordinate. The mathematical form of sampling can be expressed as follows:

$$R = R_{yaw} \cdot R_{pitch} \cdot R_{roll} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \quad (5)$$

$$[\alpha, \beta, \gamma] = [\tan^{-1}(\frac{r_{21}}{r_{11}}), \tan^{-1}(\frac{r_{32}}{r_{22}}), \tan^{-1}(\frac{r_{13}}{r_{33}})] \quad (6)$$

Here, R is the composite rotation obtained by sequentially rotating around the z-axis, y-axis, and x-axis. γ, β, α represent the included angles with corresponding axes. To simultaneously minimize both position and orientation errors of the end-effector, the target pose tracking reward r_{χ_t} is constructed in exponential form:

$$r_{\chi_t} = e^{-w \cdot \Delta(l,p,y)} \cdot e^{-\Delta(\alpha,\beta,\gamma)} \quad (7)$$

$$\begin{cases} \Delta(l, p, y) = \sum_{u \in (l,p,y)} k_u \cdot |u_t - u^{\text{cmd}}| \\ \Delta(\alpha, \beta, \gamma) = \sum_{u \in (\alpha,\beta,\gamma)} k_u \cdot |u_t - u^{\text{cmd}}| \end{cases} \quad (8)$$

where weight coefficient w is used to balance the priority of the two components. The function $\Delta(\cdot)$ is defined as the sum of absolute errors between each variable and its respective target value within the set, and k_i represents the reciprocal of the sampling range for each variable, which is used to rescale the errors. The target pose tracking reward r_{χ_t} is utilized for both the loco and arm policies in stage 2.

TABLE I
KEY REWARD TERMS, EQUATION, AND
WEIGHTS USED IN STAGE 1 AND STAGE 2

Term	Equation	Weight	
		stage 1	stage 2
$r_{s^{\text{cmd}}}^{\text{loco}}$	$(\mathbf{p}_{x,y,\text{foot}}^f - \mathbf{p}_{x,y,\text{foot}}^{f,\text{cmd}}(s^{\text{cmd}}))^2$	-10.0	0.0
r_{hip}^{loco}	$\sum_{hip} a_t ^2$	0.0	-0.05
$r_{v_{x,y}}^{\text{cmd}}$	$\exp\left\{- v_{x,y} - v_{x,y}^{\text{cmd}} ^2 / \sigma_{v_{x,y}}\right\}$	1.0	0.5
$r_{\omega_{yaw}}^{\text{cmd}}$	$\exp\left\{- \omega_{yaw} - \omega_{yaw}^{\text{cmd}} ^2 / \sigma_{\omega_{yaw}}\right\}$	0.5	0.25
r_{smooth}^{arm}	$ a_{t-1}^{\text{arm}} - a_t^{\text{arm}} ^2$	0.0	-0.1
r_{guide}^{arm}	$ \phi_t - a_t^{\text{arm}G} $	0.0	-10.0

3) **Reward Adjustment:** While the rewards in stage 1 drive the quadruped robot to achieve robust locomotion, the fixed foot placements are not conducive to the whole-body control required in stage 2, when the robot needs to adapt its body posture by altering foot placements. To address this,

we replace the raibert heuristic footswing tracking reward $r_{s^{\text{cmd}}}^{\text{loco}}$ with the hip joint constraint reward r_{hip}^{loco} . Additionally, the velocity tracking rewards, $r_{v_{x,y}}^{\text{cmd}}$ and $r_{\omega_{yaw}}^{\text{cmd}}$, are scaled to half of their original values to accommodate different tasks. The reward adjustments are detailed in Table I. To achieve more coordinated whole-body control, we introduce r_{smooth}^{arm} and r_{guide}^{arm} to promote smoother arm motion and enhance the control of the body orientation. After the reward adjustments, the arm policy and locomotion policy will be synchronized to start stage 2 training.

IV. EXPERIMENT DESIGN

A. Robot System




Fig. 3. The robot system, consisting of a quadruped robot (Unitree Go1 Edu) with a mounted robotic arm (ARX5) and a camera (RealSense D435i).

kg and a maximum reach of 620 mm. The control frequency is set to 50 Hz for both training and deployment.

The robot system used in our experiments consists of a 12-DoF legged robot Unitree Go1 Edu and a 6-DoF robotic arm ARX5 with a parallel gripper, as shown in Fig. 3. The Go1 weighs 12 kg, and the ARX5, mounted on its back, weighs 3.35 kg. A RealSense D435i camera is positioned above the gripper to maintain a fixed relative pose. The ARX5 offers a rated load capacity of 1.5

B. Algorithm Comparison

In order to examine the performance of our proposed method on whole-body control, 6D pose tracking, cross-embodiment, and loco-manipulation, we set up comparisons of the following algorithms:

- **Floating Base (FB)+IK:** A floating base policy without whole-body control for the quadruped robot, combined with an IK solver [19] for arm.
- **Adaptive Floating Base (AFB)+ID:** An adaptive floating base policy coupled with an inverse dynamics based arm control strategy, corresponding to the built-in controllers of the Unitree Go1 and ARX5.
- **Unified:** A unified policy with two output heads corresponding to the control of quadruped robot and arm. Training starts directly from stage 2.
- **Two-Stage:** The same policy settings as the Unified are used, but with a two-stage training approach.
- **Cooperated:** This setting uses the cooperative policy. Training starts directly from stage 2.
- **RoboDuet:** The same policy settings as the Cooperated are used, but with a two-stage training approach.

The first setting serves as a baseline to evaluate the solution stability of our method. The second setting represents a more powerful real-world baseline, where the legged robot actively adjusts its posture in response to the interaction forces induced

by the arm’s motion. The primary distinctions among the remaining four settings lie in whether they use a cooperative policy and a two-stage training approach, which are the two key components of our method. This comparison allows us to assess the effectiveness of each component.

C. Training Details

We train 4096 parallel agents using the IsaacGym simulator [20] and the Proximal Policy Optimization (PPO) algorithm [21]. We train all algorithms for 50,000 iterations across 3 seeds, with the two-stage training comprising 10,000 iterations for stage 1 and 40,000 for stage 2, employing the asymmetric actor-critic framework used in [22]. All neural networks are designed as the Multilayer Perceptron (MLPs) with ELU [23] activations for hidden layers. For the cooperative policy, the network architecture consists of three layers with 512, 256, and 128 neurons, respectively. In contrast to the cooperative policy, the unified policy merges feature extractors into one, so we double the number of neurons in the hidden layers. All training is performed on NVIDIA RTX 4090 GPUs.

D. Sample space

The sample ranges utilized for both training and evaluation are detailed in Table II. The evaluation sample space is slightly larger than the training one to demonstrate the generalization of algorithms, which already covers all 6D poses in the front hemisphere space. During the evaluation, we uniformly sample 200,000 target commands at random from the ranges to facilitate a comparative analysis of various algorithms.

TABLE II
RANGES OF COMMANDS USED IN TRAINING AND EVALUATION

Parameter	Range	
	Training	Evaluation
v_x (m/s)	[-1.00, 1.00]	[-1.50, 1.50]
ω_z (rad/s)	[-0.60, 0.60]	[-1.00, 1.00]
l (m)	[0.30, 0.70]	[0.20, 0.80]
p (rad)	[-0.45 π , 0.45 π]	[-0.50 π , 0.50 π]
y (rad)	[-0.50 π , 0.50 π]	[-0.50 π , 0.50 π]
α (rad)	[-0.45 π , 0.45 π]	[-0.50 π , 0.50 π]
β (rad)	[-0.33 π , 0.33 π]	[-0.50 π , 0.50 π]
γ (rad)	[-0.42 π , 0.42 π]	[-0.50 π , 0.50 π]

E. Metrics

To quantify the performance of the algorithms, we define several metrics: (1) velocity tracking error, measured by the difference between the target commands and actual states for v_x , v_y , and ω_{yaw} ; (2) position tracking error, determined by the error for l , p , y , and the Euclidean distance D ; (3) orientation tracking error, calculated based on the angles α , β , γ , and the quaternion geodesic distance ζ ; (4) survival rate, assessed by randomly applying forces of 10 to 20 newtons to the base for 2 seconds and calculating the proportion of robots that maintain a base height above 0.26 m across all samples; (5) solvability,

defined as the ratio of collision-free poses to the total number of samples; and (6) workspace, defined as the area of the convex hull formed by all target commands within the tracking error threshold, ensuring no self-collisions throughout the tracking process. A sample is considered successful if the tracking of the end-effector pose meets the thresholds of $D \leq 0.03$ m for distance and $\theta \leq \pi/36$ for orientation, where θ is the angle between the two vectors obtained by applying the target orientation and end-effector orientation to the same unit vector. The maximum time allowed to reach the target command is 4 seconds, after which the average error over the following 2 seconds is computed.

V. RESULTS

A. Simulation Experiments

1) *Whole-body Control*: To validate the effectiveness of our proposed whole-body control training framework RoboDuet, we make a comparison with the FB+IK. To be fair, we evaluate the solvability and workspace under the same error threshold by applying the average Euclidean distance error and the quaternion geodesic distance error of RoboDuet as the pose error thresholds for IK solver. For both, the initial joint positions of the arm are set to the same as default positions mentioned in Section III-B1. According to Fig. 4, RoboDuet demonstrates a 14.83% improvement in solvability compared to the FB+IK configuration. Furthermore, the workspace comparison suggests that whole-body control can significantly enhance the operational capabilities of the robotic arm. It is worth reiterating that solvability is defined as reaching the target pose from the current pose without any collisions throughout the entire process. Additionally, many samples are inherently unreachable due to the limitations imposed by the structure of the robotic arm itself.




Fig. 4. Comparison of Solvability and Workspace between FB+IK and RoboDuet.

2) *Ablation*: To validate the significance of the two-stage training and the cooperative mechanism, which are the key components of RoboDuet, we conduct sufficient ablation experiments by comparing with Unified, Two-Stage and Co-operated. The results are shown in Table III, which demonstrate that all configurations meet the requirements for stable standing. Although two-stage training does not significantly improve end-effector pose tracking, it demonstrates a notable

TABLE III
 ABLATION RESULTS (SCALED BY 10^{-2}) OF DIFFERENT METHODS IN STILL AND MOVE MODES
 "STILL" MEANS ALL VELOCITY COMMANDS ARE ZERO, WHILE "MOVE" INVOLVES RANDOM VELOCITY SAMPLING DURING EVALUATION.

Metrics		Still				Move			
		Unified	Two-Stage	Cooperative	RoboDuet	Unified	Two-Stage	Cooperative	RoboDuet
Velocity Tracking ↓	v_x (m/s)	0.43±0.00	0.47±0.00	0.30±0.00	0.32±0.00	12.16±0.02	11.47±0.00	9.83±0.05	9.7±0.00
	v_y (m/s)	0.45±0.00	0.44±0.00	0.36±0.00	0.34±0.00	17.10±0.00	16.84±0.02	15.91±0.00	15.42±0.00
	ω_z (rad/s)	0.37±0.00	0.35±0.00	0.33±0.00	0.32±0.00	61.83±0.03	62.06±0.03	60.82±0.05	60.59±0.02
Position Tracking ↓	l (m)	4.25±0.04	4.13±0.00	1.90±0.02	1.97±0.00	4.01±0.01	4.25±0.01	1.88±0.00	1.91±0.02
	p (rad)	22.62±0.00	21.37±0.12	17.42±0.03	18.81±0.05	21.93±0.11	22.45±0.07	17.06±0.03	18.09±0.01
	y (rad)	12.01±0.01	11.52±0.04	7.94±0.00	7.43±0.00	12.15±0.00	11.43±0.00	7.70±0.00	7.22±0.00
	D (m)	14.47±0.01	13.84±0.01	10.79±0.01	11.08±0.00	13.56±0.02	13.05±0.01	10.58±0.00	10.75±0.00
Orientation Tracking ↓	α (rad)	44.15±0.03	47.76±0.03	40.46±0.05	39.17±0.02	42.50±0.08	47.63±0.07	43.22±0.04	39.89±0.01
	β (rad)	60.41±0.13	62.86±0.08	48.89±0.03	45.44±0.02	56.74±0.04	63.25±0.15	45.29±0.01	47.58±0.03
	γ (rad)	52.33±0.01	54.74±0.03	43.56±0.04	39.38±0.04	52.36±0.03	54.38±0.05	42.74±0.02	39.01±0.00
	ζ (-)	51.15±0.01	53.46±0.02	48.18±0.00	47.14±0.00	50.58±0.00	53.03±0.01	48.36±0.00	47.53±0.00
Survival Rate (%) ↑		87.35±0.78	95.66±0.06	94.49±0.03	98.20±0.01	93.03±1.34	99.09±0.01	98.56±0.02	99.96±0.00
Workspace (m^3) ↑		51.36±0.46	60.10±0.05	84.08±0.06	85.79±0.02	78.96±0.04	81.30±0.02	86.71±0.01	87.05±0.01

improvement in resisting external perturbations. Specifically, Two-Stage improves survival rates by 8.31% over the Unified in stationary standing, attributed to the focus of stage 1 on optimizing locomotion. On the other hand, Cooperated and RoboDuet outperform both the Unified and Two-Stage configurations across almost all metrics, indicating that the cooperative mechanism effectively decouples multi-task learning, enabling the two policies to collaborate while maintaining focus. Cooperated achieves the best performance in position tracking, largely due to its lack of the gait prior constraint from stage 1, allowing it to take a more aggressive approach in achieving higher tracking rewards which results a decrease on survival rates. In comparison, RoboDuet maintains comparable end-effector tracking performance while achieving robust locomotion motion ability and offers a larger operational workspace during motion. In summary, by effectively integrating cooperative policy and two-stage training, RoboDuet significantly enhances control performance in both tracking accuracy and gait stability, underscoring the essential role of these components.

TABLE IV
 WORKSPACES AND SURVIVAL RATES OF ZERO-SHOT POLICY TRANSFER
 FROM Go1 TO Go2 AND A1.

X+ARX	Workspace (m^3)		Survival Rate (%)	
	Still	Move	Still	Move
Go1	0.8579	0.8705	98.20	99.96
Go2	0.8552	0.8216	96.32	93.21
A1	0.8416	0.8121	97.07	94.34

3) *Zero-shot Transfer*: To assess the zero-shot transfer capability of RoboDuet, we introduce two additional quadruped robots, Unitree A1 and Unitree Go2, and mount the ARX5 on their backs. The policy trained on the Go1+ARX system is directly transferred to these newly integrated platforms. Subsequently, we evaluate their workspaces and survival rates

in the presence of external disturbances, with the results presented in Table IV. Notably, under stationary standing conditions, Go2+ARX achieves a workspace of 0.8552, while A1+ARX attains a survival rate of 97.07%, closely aligning with the performance of the original Go1+ARX system. Due to fundamental differences among the various embodiments, only Go1+ARX5 demonstrates an increase in stability and workspace during motion through its custom-tailored cooperative policy. In contrast, the other two systems exhibit slight performance degradation. Nevertheless, they still achieve a survival rate of nearly 93% under external perturbations.

B. Real-world Experiments

In terms of real-world experiments, we design three distinct types of tasks to evaluate the effectiveness of our policy in the real world and its proficiency in handling diverse locomotion challenges. We directly deploy FB+IK, AFB+ID, and RobotDuet on a real-world robotic system for testing.

1) *Extreme 6D Pose Tracking*: To evaluate the robustness and generalization of end-effector pose tracking in our method, we select five target poses to compare the self-collision and solvability between FB+IK and RoboDuet. The target poses are illustrated in Fig. 6, with four representing extreme 6D target configurations. Poses 2 and 3 are close to the ground and intersect with the robot's legs, while poses 4 and 5 are significantly elevated above the robot, beyond the training sampling range. Notably, pose 5 is located on the right side, nearing the robotic arm's maximum reach, and oriented towards the robot's head, representing an highly pathological configuration that is unsolvable. For each pose, we conduct eight repetitions, and the results are shown in Table V. The IK system fails to resolve poses 2, 4, and 5. However, RoboDuet maintains stability across all poses and approached optimal approximate solutions for physically unreachable poses. Although pose 3 can be solved using IK, the floating base policy lacks whole-

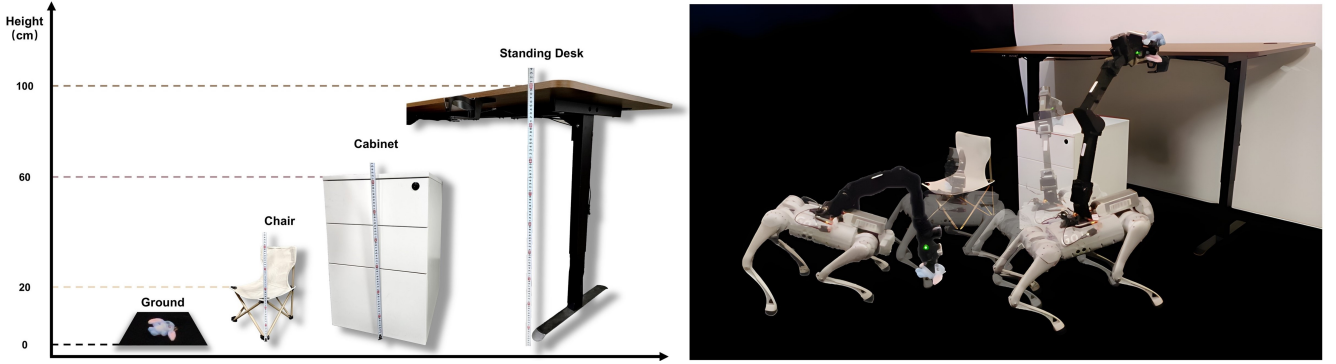


Fig. 5. Transferring objects between different heights, including the ground (0 cm), a camping chair (20 cm), a cabinet (60 cm), and the cup holder of a standing desk (100 cm). The right part illustrates how RoboDuet utilizes whole-body control to adapt its posture for varying grasp poses during transfer tasks.




Fig. 6. Tracking extreme 6D poses with RoboDuet. The target pose is represented as $[l, p, y, \text{roll}, \text{pitch}, \text{yaw}]$, with l in meters and the others in radians. For certain unreachable targets, RoboDuet achieves stable approximations to optimal solutions.

body control, leading to self-collision issues. Additionally, since the training does not include scenarios where the arm is in motion, tracking during poses 4 and 5 results in a higher risk of tipping, which we regard as self-collision. In contrast, due to the self-collision penalty integrated into the training process and its ability to maintain locomotion while the arm is in motion, RoboDuet experiences no collisions in nearly 95% of trials, with the arm stopping effectively near the legs in poses 2 and 3.

TABLE V
SELF-COLLISION AND IK FAILURE RATE OF
FB + IK VS. ROBODUET

Method	Self-collision Rate	IK Failure Rate
FB+IK	50%	60%
RoboDuet	5%	-

2) *Whole-body Control*: To further analyze the impact of whole-body control on loco-manipulation, we design a set of tasks involving transferring a doll from varying heights, as shown in Fig. 5. For each height, we calculate the average success rate from five repeated trials conducted by each of the three different operators. A trial is considered successful only if the object is picked up and transferred to the next height without being dropped at any moment during this process. The

system’s speed is controlled via the right VR handle, while the end-effector is operated with the left handle.

TABLE VI
SUCCESS RATES ACROSS DIFFERENT HEIGHTS FOR VARIOUS CONTROL
METHODS IN OBJECT TRANSFER TASKS.

Method	Change in Height (cm)				Avg
	0 → 20	20 → 60	60 → 100	100 → 0	
FB+IK	3/15	7/15	0/15	0/15	10/60
AFB+ID	9/15	13/15	6/15	4/15	32/60
RoboDuet	10/15	12/15	8/15	9/15	39/60

The experimental results are shown in Table VI. When grasping objects on the ground, it is crucial to avoid self-collisions. Lacking whole-body control capabilities, the FB+IK approach is prone to self-collisions and failure, particularly when picking up objects positioned in front of or below the head. In contrast, AFB+ID and RoboDuet can adjust the body posture to minimize collisions during grasping, resulting in an almost 40% increase in success rate. However, since AFB+ID can only adjust the rotation around the roll axis based on interaction forces, it still struggles to grasp objects positioned directly in front of the head. As the height increases, the need for body posture adjustments decreases, resulting in an improvement in success rates for all approaches. Placing or retrieving the doll from the cup holder on a standing desk presented a particularly challenging task. The robotic arm was required to maintain a horizontal or top-down pose while operating near its maximum reach, when IK often fails to find a valid solution. AFB+ID relies on its gravity compensation mechanism and inverse dynamics solution, achieving a placement success rate of 40%. However, due to joint limit constraints, it can only successfully retrieve the doll when its ears or head are positioned near the outer edge of the cup holder. RoboDuet demonstrates high solution stability and is able to reduce the vertical distance between the base and the standing desk by inclining the body, which successfully completed approximately 55% of the trials. Compared to AFB+ID, RoboDuet achieves a 23% improvement in success rate for more challenging loco-manipulation tasks.

3) *Zero-shot Transfer*: To evaluate the generalizability of our method across different robot embodiments, we directly deploy the policy trained on the Unitree Go1+ARX5 to the Unitree Go2+ARX5. The Unitree Go1 platform has a total mass of 18.35 kg, whereas the Unitree Go2 platform weighs 21.05 kg, representing a 14.7% increase. Despite the significant increase in mass, RoboDuet still demonstrated robust whole-body control across both platforms, as illustrated in Fig. 7. The figure depicts four key moments of the two systems tracking the same trajectory, demonstrating consistent 6D pose tracking and robust whole-body control capabilities across both configurations.




Fig. 7. Zero-shot deployment of RoboDuet, trained on the Unitree Go1+ARX5 (left), onto the Unitree Go2+ARX5 (right). The bottom-left image illustrates consistent agile 6D pose tracking, while the others highlight coordinated whole-body control.

VI. DISCUSSION AND LIMITATIONS

In this letter, we propose RoboDuet, a whole-body legged loco-manipulation framework that integrates two collaborative policies for velocity tracking and 6D end-effector pose control, enabling agile and robust performance across diverse tasks. Leveraging a two-stage training approach, RoboDuet effectively utilizes robust locomotion priors to enhance the system’s resistance to external disturbances. Additionally, our method supports zero-shot transfer deployment, allowing seamless hardware replacement without retraining, and it achieves real-world performance comparable to simulation.

While RoboDuet focuses on low-level whole-body control, achieving fully autonomous task execution will require integration with high-level planners and advanced collision detection. Moreover, our method can only handle terrains such as slopes and gravel while maintaining a fixed target pose. However, it has yet to achieve loco-manipulation on more complex terrains, such as stairs. These are key areas for future enhancement.

REFERENCES

- [1] 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.
- [2] 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,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023.
- [3] Z. Fu, X. Cheng, and D. Pathak, “Deep whole-body control: learning a unified policy for manipulation and locomotion,” in *2023 Conference on Robot Learning (CoRL)*. PMLR, 2023, pp. 138–149.
- [4] J. Zhang, N. Gireesh, J. Wang, X. Fang, C. Xu, W. Chen, L. Dai, and H. Wang, “Gamma: Graspability-aware mobile manipulation policy learning based on online grasping pose fusion,” in *2024 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2024.
- [5] R.-Z. Qiu, Y. Hu, G. Yang, Y. Song, Y. Fu, J. Ye, J. Mu, R. Yang, N. Atanasov, S. Scherer et al., “Learning generalizable feature fields for mobile manipulation,” *arXiv preprint arXiv:2403.07563*, 2024.
- [6] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, “Perceptive locomotion through nonlinear model-predictive control,” *IEEE Transactions on Robotics (TRO)*, 2023.
- [7] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, “Robot parkour learning,” in *2023 Conference on Robot Learning (CoRL)*. PMLR, 2023, pp. 73–92.
- [8] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, “A unified mpc framework for whole-body dynamic locomotion and manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4688–4695, 2021.
- [9] H. Ferrolho, V. Ivan, W. Merkt, I. Havoutis, and S. Vijayakumar, “Roloma: robust loco-manipulation for quadruped robots with arms,” *Auton. Robots*, vol. 47, no. 8, p. 1463–1481, Oct. 2023. [Online]. Available: <https://doi.org/10.1007/s10514-023-10146-0>
- [10] J.-P. Sleiman, F. Farshidian, and M. Hutter, “Versatile multicontact planning and control for legged loco-manipulation,” *Science Robotics*, p. eadg5014, 2023.
- [11] N. Yokoyama, A. W. Clegg, E. Undersander, S. Ha, D. Batra, and A. Rai, “Adaptive skill coordination for robotic mobile manipulation,” in *IEEE Robotics Autom. Lett. (RA-L)*. IEEE, 2024, pp. 779–786.
- [12] M. Liu, Z. Chen, X. Cheng, Y. Ji, R. Qiu, R. Yang, and X. Wang, “Visual whole-body control for legged loco-manipulation,” *CoRL*, 2024.
- [13] Y. Ma, F. Farshidian, T. Miki, J. Lee, and M. Hutter, “Combining learning-based locomotion policy with model-based manipulation for legged mobile manipulators,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2377–2384, 2022.
- [14] H. Ha, Y. Gao, Z. Fu, J. Tan, and S. Song, “UMI-on-legs: Making manipulation policies mobile with a manipulation-centric whole-body controller,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=3i7j8ZPnbm>
- [15] T. Portela, A. Cramariuc, M. Mittal, and M. Hutter, “Whole-body end-effector pose tracking,” *arXiv preprint arXiv:2409.16048*, 2024.
- [16] G. B. Margolis and P. Agrawal, “Walk these ways: Tuning robot control for generalization with multiplicity of behavior,” in *2023 Conference on Robot Learning (CoRL)*. PMLR, 2023, pp. 22–31.
- [17] Z. Fu, A. Kumar, J. Malik, and D. Pathak, “Minimizing energy consumption leads to the emergence of gaits in legged robots,” in *Conference on Robot Learning*, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:237372201>
- [18] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, “On the continuity of rotation representations in neural networks,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 5738–5746.
- [19] B. Sundaralingam, S. K. S. Hari, A. Fishman, C. Garrett, K. V. Wyk, V. Blukis, A. Millane, H. Oleynikova, A. Handa, F. Ramos, N. Ratliff, and D. Fox, “curobo: Parallelized collision-free minimum-jerk robot motion generation,” 2023.
- [20] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa et al., “Isaac gym: High performance gpu-based physics simulation for robot learning,” *2021 Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks (NeurIPS)*, 2021.
- [21] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [22] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, “Asymmetric actor critic for image-based robot learning,” in *Robotics: Science and Systems XIV*, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018, H. Kress-Gazit, S. S. Srinivasa, T. Howard, and N. Atanasov, Eds., 2018. [Online]. Available: <http://www.roboticsproceedings.org/rss14/p08.html>
- [23] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, “Fast and accurate deep network learning by exponential linear units (elus),” *arXiv: Learning*, 2015. [Online]. Available: <https://api.semanticscholar.org/CorpusID:5273326>