

# ResMimic: From General Motion Tracking to Humanoid Whole-body Loco-Manipulation via Residual Learning

Siheng Zhao<sup>12§</sup> Yanjie Ze<sup>13§</sup> Yue Wang<sup>2</sup> C. Karen Liu<sup>13†</sup>  
Pieter Abbeel<sup>14†</sup> Guanya Shi<sup>15†</sup> Rocky Duan<sup>1†</sup>

<sup>1</sup>Amazon FAR (Frontier AI & Robotics) <sup>2</sup>USC <sup>3</sup>Stanford University <sup>4</sup>UC Berkeley <sup>5</sup>CMU

<https://resmimic.github.io/>




Fig. 1: We deploy ResMimic on a Unitree G1 humanoid to demonstrate diverse whole-body loco-manipulation capabilities. (a) Carrying 4.5kg heavy payloads via whole-body contact; (b, c) expressive motion while carrying box; (d) general humanoid object interaction beyond manipulation; (e, f) carrying irregularly shaped heavy objects with instance-level generalization.

**Abstract—** Humanoid whole-body loco-manipulation promises transformative capabilities for daily service and warehouse tasks. While recent advances in general motion tracking (GMT) have enabled humanoids to reproduce diverse human motions, these policies lack the precision and object awareness required for loco-manipulation. To this end, we introduce ResMimic, a two-stage residual learning framework for precise and expressive humanoid control from human motion data. First, a GMT policy, trained on large-scale human-only motion, serves as a task-agnostic base for generating human-like whole-body movements. An efficient but precise residual policy is then learned to refine

the GMT outputs to improve locomotion and incorporate object interaction. To further facilitate efficient training, we design (i) a point-cloud-based object tracking reward for smoother optimization, (ii) a contact reward that encourages accurate humanoid body-object interactions, and (iii) a curriculum-based virtual object controller to stabilize early training. We evaluate ResMimic in both simulation and on a real Unitree G1 humanoid. Results show substantial gains in task success, training efficiency, and robustness over strong baselines. Videos are available at [resmimic.github.io](https://resmimic.github.io).

## I. INTRODUCTION

Humanoid robots have been showing promising applications thanks to their flexibility. Unlike quadrupeds or wheeled

<sup>§</sup>Work done during the internship of Siheng Zhao and Yanjie Ze at Amazon FAR (Frontier AI & Robotics) <sup>†</sup>FAR Co-Lead

manipulators, they can perform expressive whole-body loco-manipulation, combining locomotion and manipulation in a coordinated, human-like manner. This unique capability opens the door to transformative applications in daily services and industrial operations, while leveraging existing human-centric infrastructure.

However, realizing precise and expressive humanoid loco-manipulation remains a fundamental challenge to humanoids. Compared to locomotion [1], [2] or tabletop manipulation [3], [4], loco-manipulation demands higher precision, rich whole-body contacts, and data that is not easily accessible at scale. While direct imitation of human motions [5] is attractive—since humans naturally perform coordinated full-body control—retargeting human demonstrations to humanoid robots introduces significant embodiment gap: contact locations and relative object pose in human demonstrations often fail to translate, leading to floating contacts or penetrations, as shown in Figure 2.

Recent progress on general motion tracking (GMT) policies [6], [7], trained on large-scale human-only datasets [8], [9], shows that humanoids can reproduce diverse human motions with high fidelity. However, these policies are unaware of manipulated objects. On the other hand, existing humanoid loco-manipulation approaches rely on highly task-specific designs, such as stage-wise controllers [10] or handcrafted data pipelines [11], which limit scalability and generality. As a result, there is still no unified, efficient, and precise framework for humanoid loco-manipulation yet.

In parallel, breakthroughs in foundation models [12] have demonstrated the power of pre-training on large-scale data followed by post-training. In robotics manipulation, models like  $\pi_0$  [13] and LBM [14] highlight that this paradigm resolves data scarcity while improving robustness and generalization. In graphics, pretrained latent spaces are similarly leveraged for downstream generation and editing [15], [16]. However, in humanoid whole-body control, such a powerful pretrain–finetune paradigm remains largely unexplored.

In this paper, our key insight is that while diverse human motions can be effectively mimicked by a pre-trained GMT policy, object-centric loco-manipulation requires task-specific corrections. Importantly, many whole-body motions—such as balance, stepping, or reaching—are shared across tasks, while only the fine-grained object interaction requires adaptation. This motivates a residual learning paradigm, where a stable motion prior is augmented with lightweight task-specific adjustments.

To this end, we propose *ResMimic*, a two-stage residual learning framework for humanoid loco-manipulation. First, a GMT policy is trained on large-scale motion capture data to reproduce diverse human motions, serving as a robust prior for human-like whole-body behavior. Second, a task-specific residual policy is trained efficiently to condition on object reference trajectory, outputting corrective actions that refine the GMT policy and enable precise object manipulation. This decoupling alleviates the need for per-task reward engineering, improves data efficiency, and yields a general framework applicable not only to loco-manipulation but also

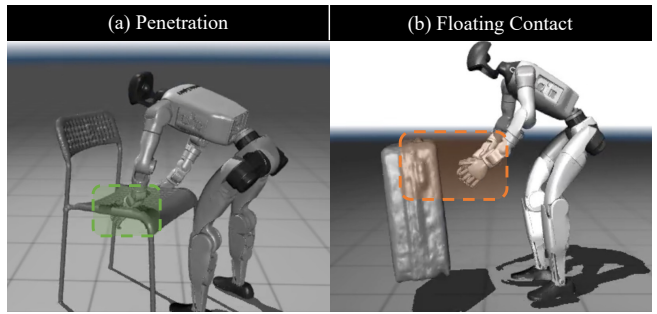


Fig. 2: Visualization of imperfect humanoid–object interaction data caused by the embodiment gap during retargeting: (a) hand–chair penetration; (b) hand–box floating contact.

to locomotion enhancement.

Our contributions are four-fold:

- We propose a two-stage residual learning framework that combines pre-trained GMT with task-specific corrections, enabling efficient, precise humanoid loco-manipulation.
- To improve training efficiency and sim-to-real transfer, we propose (i) a point-cloud–based object tracking reward for smoother optimization, (ii) a contact reward that explicitly guides humanoid–object contacts, and (iii) a virtual object controller that provides a curriculum-based warm start.
- We conduct extensive evaluations in simulation and real world. Results demonstrate significant improvements in human motion tracking, object motion tracking, task success rate, training efficiency, robustness, and generalization on challenging loco-manipulation tasks.
- To accelerate research on humanoid loco-manipulation, we will release our GPU-accelerated simulation infrastructure, sim-to-sim evaluation prototype, and motion data.

## II. RELATED WORK

### A. Learning-Based Humanoid Control

Reinforcement learning (RL) enables real-time whole-body control of humanoids by training directly via interactions with the environment. However, such training typically suffers from low data efficiency and requires substantial effort in task-specific reward design [17]. As a result, most prior work has focused primarily on locomotion [18], [19] rather than versatile whole-body control, or is limited in specific task such as getting up [20] and keeping balance [21].

To enable more general control, learning from human motions has emerged as a promising direction [5]. These methods typically rely on kinematic retargeting to map human motions to humanoids, addressing the large embodiment gap. It enables accurate tracking of individual motions [22]–[24] or aims for versatile general motion tracking [6], [7], [25], [26]. While these methods demonstrate whole-body control ability, they remain limited in their ability to interact with physical world. VideoMimic [27] makes notable progress by not only tracking human motion but also reconstructing the environment, enabling contextual motions such as sitting on a chair. However, it still interacts only with static environments and does not extend to dynamic object interactions—an

essential step toward real-world utility. Different from these works, our work focuses on dynamic loco-manipulation.

### B. Humanoid Loco-Manipulation

Humanoid loco-manipulation remains a particularly challenging learning problem. Recent work has demonstrated promising results by leveraging teleoperation [6], [28]–[30]. However, these approaches lack explicit object awareness and require a human operator. Building on teleoperation, some works train autonomous imitation learning policies from collected data [4], [31]. Nonetheless, these efforts remain restricted to tabletop manipulation with limited whole-body expressiveness—tasks that could often be achieved more effectively by dual-arm mobile manipulators. The most related works to ours are [10], [11]. Dao *et al.* [10] propose a modular sim-to-real RL pipeline for box loco-manipulation, decomposing the task into distinct phases (e.g., walking, box-picking) with separate policies for each. Liu *et al.* [11] introduce an end-to-end learning pipeline using reference motions generated by task-specific trajectory optimization. However, all these approaches demonstrate loco-manipulation with limited whole-body contact (e.g., using only hands) and expressiveness, as well as rely on highly task-specific design. In contrast, our method leverages a GMT policy as prior, enabling more expressive whole-body loco-manipulation under a unified framework.

### C. Residual Learning for Robotics

Residual learning on top of predefined or learned base models has been widely adopted in robotics. Early works introduced residual policies to refine hand-designed policies or model predictive controllers, achieving more precise and contact-rich manipulation [32], [33]. Building on this idea, later approaches extended residual learning to policies initialized from demonstrations [34], [35]. In dexterous hand manipulation, residual policies have been used to adapt human hand motions for task-oriented control [36], [37]. Notably in humanoids, ASAP [22] leverages residual learning to compensate for dynamics mismatch between simulation and reality, enabling agile whole-body skills. Distinct from these directions, our method leverages a pre-trained general motion tracking (GMT) policy as a foundation and learns a residual policy to enable expressive whole-body loco-manipulation.

## III. METHOD

We formulate our whole-body loco-manipulation task as a goal-conditioned RL problem within a Markov Decision Process (MDP)  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  the action space,  $\mathcal{T}$  the transition dynamics,  $\mathcal{R}$  the reward function, and  $\gamma$  the discount factor. At time  $t$ , the state  $s_t \in \mathcal{S}$  includes: (i) robot proprioception  $s_t^r$ , (ii) object state  $s_t^o$ , (iii) motion goal state  $\hat{s}_t^r$ , and (iv) object goal state  $\hat{s}_t^o$ . The action  $a_t$  specifies target joint angles, which are executed on the robot through a PD controller. The reward is defined as  $r_t = \mathcal{R}(s_t, a_t)$ , while the training objective is to maximize the expected cumulative discounted reward:  $\mathbb{E}[\sum_{t=1}^T \gamma^{t-1} r_t]$ .

### A. Two-Stage Residual Learning

To transfer human–object interaction data into humanoid whole-body loco-manipulation policies in a task-agnostic manner, our objective is to avoid task-specific reward engineering, which may benefit individual tasks but limits generality. Instead, we propose a two-stage residual learning framework, as shown in Figure 3. **Stage I: General Motion Tracking (GMT).** We first train a general human motion tracking policy  $\pi_{\text{GMT}}$  as the backbone controller. Given robot proprioception  $s_t^r$  and reference motion  $\hat{s}_t^r$ , the policy outputs a coarse action  $a_t^{\text{gmt}} = \pi_{\text{GMT}}(s_t^r, \hat{s}_t^r)$ , optimized to maximize the motion tracking reward  $\mathbb{E}[\sum_{t=1}^T \gamma^{t-1} r_t^m]$ . **Stage II: Residual Refinement.** Building on the pretrained GMT policy, we train an efficient and precise residual policy  $\pi_{\text{Res}}$  per-task that refines the coarse action using both robot and object information:  $\pi_{\text{Res}}(s_t^r, s_t^o, \hat{s}_t^r, \hat{s}_t^o) = \Delta a_t^{\text{res}}$ . The final action is computed as  $a_t = a_t^{\text{gmt}} + \Delta a_t^{\text{res}}$ , and the residual policy is optimized to maximize the combined motion and object rewards  $\mathbb{E}[\sum_{t=1}^T \gamma^{t-1} (r_t^m + r_t^o)]$ . Both stages are trained using PPO [38].

### B. General Motion Tracking Policy

In this stage, our goal is to train a real-world deployable general motion tracking policy  $\pi_{\text{GMT}}$  that takes only humanoid proprioception  $s_t^r$  and human reference motion  $\hat{s}_t^r$  as input, and outputs actions  $a_t^{\text{gmt}}$  for the humanoid robot to mimic the reference motion. While ResMimic is a general framework that can incorporate any GMT policy as the base policy, we present one specific implementation as example.

1) *Dataset:* An important motivation of our two-stage training pipeline is to decouple human motion tracking from object interaction. The general motion tracking policy relies solely on human motion capture data, avoiding the need for costly and hard-to-obtain manipulation data.

We leverage several publicly available MoCap datasets, including AMASS [8] and OMOMO [9], which together contain over 15,000 clips (approximately 42 hours). Motions that are impractical for our setting, such as stair climbing, are filtered out. After curating the human motion dataset, we apply kinematics-based motion retargeting (e.g. GMR [39]) to transfer human motions into a humanoid reference motion dataset  $\{\hat{S}_i^r = \{\hat{s}_t^r\}_{t=1}^T\}_{i=1}^D$ .

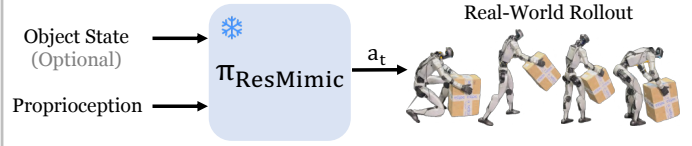
2) *Training Strategy:* To train a real-world deployable general motion tracking policy  $\pi_{\text{GMT}}$ , we adopt a single-stage RL framework in simulation without access to privileged information. The proprioceptive observation  $s_t^r$  is defined as  $[\theta_t, \omega_t, q_t, \dot{q}_t, a_t^{\text{hist}}]_{t-10:t}$ , where  $\theta_t$  is the root orientation,  $\omega_t$  is the root angular velocity,  $q_t \in \mathbb{R}^{29}$  is the joint position,  $\dot{q}_t$  is the joint velocity, and  $a_t^{\text{hist}}$  is the recent action history. The reference motion input  $\hat{s}_t^r$  is defined as  $[\hat{p}_t, \hat{\theta}_t, \hat{q}_t]_{t-10:t+10}$ , where  $\hat{p}_t$  is the reference root translation,  $\hat{\theta}_t$  is the reference root orientation, and  $\hat{q}_t$  is the reference joint position. To improve tracking quality, we also incorporate future reference motion into the input, enabling the policy to anticipate and plan for upcoming targets, which yields smoother tracking.



### Stage 1: Base Policy Training



### Stage 3: Real-World Deployment



### Stage 2: Residual Policy Training

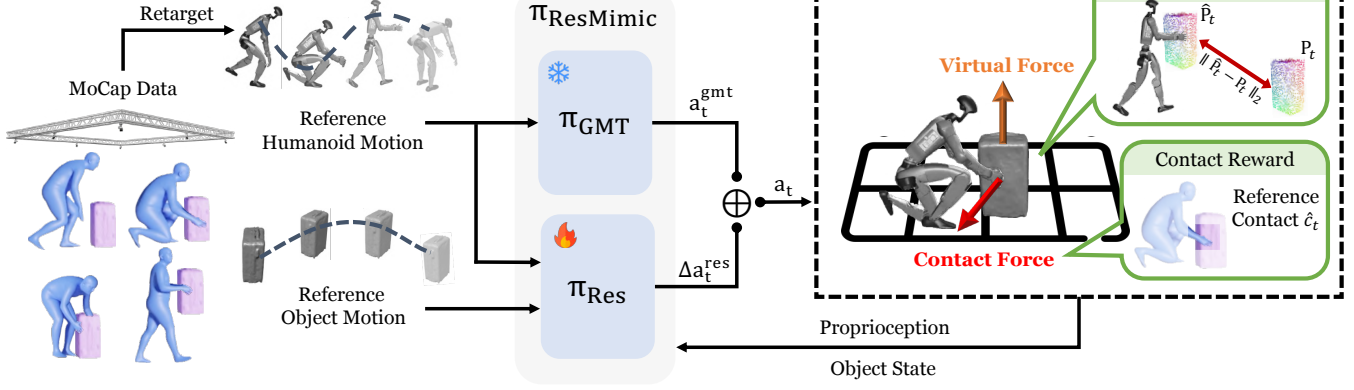


Fig. 3: Overview of ResMimic : (1) A general motion tracking policy is trained on large-scale human motion data to serve as base policy. (2) A task-specific residual policy is efficiently trained with **virtual force**, **object** and **contact reward**, to refine the base policy outputs. (3) During real-world deployment, the combined policy is employed for loco-manipulation control.

3) *Reward and Domain Randomization*: Following TWIST [6], the motion tracking reward  $r_t^m$  is formulated as the sum of three components: (i) task rewards, (ii) penalty terms, and (iii) regularization terms. To promote robust and generalizable sim-to-real transfer, we further apply domain randomization during training. See [6] for more details.

#### C. Residual Refinement Policy

Building on the pre-trained  $\pi_{\text{GMT}}$ , we introduce a residual policy  $\pi_{\text{Res}}$  to refine the coarse actions predicted by the base policy and thereby complete the desired task.

1) *Reference Motions*: We obtain reference trajectories using a MoCap system that simultaneously records human motion  $\{\hat{h}_t\}_{t=1}^T$  and object motion  $\{\hat{o}_t\}_{t=1}^T$ . The human motion is retargeted to the humanoid robot using GMR [39], yielding humanoid reference trajectories  $\{\hat{s}_t^r = \text{GMR}(\hat{h}_t)\}_{t=1}^T$ . The object motion is directly used as the reference  $\hat{s}_t^o$ . Together, these form the complete reference trajectory  $\{(\hat{s}_t^r, \hat{s}_t^o)\}_{t=1}^T$  for training the residual policy.

2) *Training Strategy*: We adopt single-stage RL with PPO for residual learning. The residual policy  $\pi_{\text{Res}}$  takes  $\langle s_t^r, s_t^o, \hat{s}_t^r, \hat{s}_t^o \rangle$  as input and outputs a residual action  $\Delta a_t^{\text{res}} \in \mathbb{R}^{29}$ . The object state is represented as  $s_t^o = [p_t^o, \theta_t^o, v_t^o, \omega_t^o]$ , and the reference object trajectory as  $\hat{s}_t^o = [\hat{p}_t^o, \hat{\theta}_t^o, \hat{v}_t^o, \hat{\omega}_t^o]_{t-10:t+10}$ , where  $p_t^o$  denotes the object root translation,  $\theta_t^o$  the root orientation,  $v_t^o$  the root velocity, and  $\omega_t^o$  the root angular velocity.

**Network Initialization.** At the start of training, the humanoid already closely mimics the reference human motion, so the residual policy should ideally output values near zero. To enforce this, we initialize the final layer of the PPO actor using Xavier uniform initialization with a small gain

factor (i.e., a scalar that scales the variance of the initialized weights) so that the initial outputs are close to zero [40].

**Virtual Object Force Curriculum.** While the two-stage residual framework performs well on simple tasks with reliable reference trajectories and light objects, it often fails early when reference motions are noisy or objects are heavy. Failures arise mainly from (i) penetration introduced by kinematic retargeting, which causes the humanoid to push or knock over the object when naively imitating the reference—leading the initial policy to retreat rather than engage—and (ii) instability when handling large object masses.

To address this, inspired by [41], we introduce a virtual object controller curriculum that stabilizes training by driving the object toward its reference trajectory. At each timestep, PD controllers apply virtual forces and torques:

$$\mathcal{F}_t = k_p(\hat{p}_t^o - p_t^o) - k_d v_t^o, \quad \mathcal{T}_t = k_p(\hat{\theta}_t^o \ominus \theta_t^o) - k_d \omega_t^o,$$

where  $\mathcal{F}_t$  and  $\mathcal{T}_t$  are the control force and torque, and  $\ominus$  denotes the rotation difference. The controller gains  $(k_p, k_d)$  are gradually decayed, so that early training is stabilized by strong virtual assistance, while later training forces the policy to take over and complete the task autonomously.

3) *Reward and Early Termination*: Decoupling motion tracking from object interaction brings an additional advantage: we avoid carefully tuning the relative weights between motion and object rewards. Instead, we directly reuse the motion reward  $r_t^m$  and domain randomization from GMT training, and introduce two additional terms: the object tracking reward  $r_t^o$ , which encourages task completion, and the contact tracking reward  $r_t^c$ , which provides explicit guidance on body-object contact, improving real-world deployability.




Fig. 4: We deploy ResMimic on Unitree G1 with MoCap-based object states. (a) Lifting a box from random object initial poses across 11 trials; (b) Autonomous consecutive kneeling and box lifting; (c) Reactive behavior to external perturbations.

**Object Tracking Reward.** Prior work [11], [42] typically measures object tracking using pose differences between the simulated and reference objects, e.g.,  $r_t^o = \exp(-\lambda_p \|p_t^o - \hat{p}_t^o\|_2) + w \cdot \exp(-\lambda_\theta \|\theta_t^o - \hat{\theta}_t^o\|_2)$ . Instead, we propose an alternative with smoother reward landscape: sample  $N$  points from the object mesh surface and compute the point-cloud difference between the current and reference states,

$$r_t^o = \exp(-\lambda_o \sum_{i=1}^N \|\mathbf{P}[i]_t - \hat{\mathbf{P}}[i]_t\|_2),$$

where  $\mathbf{P}_t \in \mathbb{R}^{N \times 3}$  denotes the sampled 3D points. This approach naturally accounts for both translation and rotation, eliminating the need for task-specific weight tuning.

**Contact Reward.** To encourage correct physical interactions during whole-body manipulation while remaining efficient, we discretize contact locations into meaningful links, e.g., torso, hip, and arms, excluding feet since they mainly contact the ground. Oracle contact information is obtained from the reference human-object interaction trajectory:

$$\hat{c}_t[i] = \mathbf{1}(\|\hat{d}_t[i]\| < \sigma_c),$$

where  $i$  indexes the links,  $\mathbf{1}(\cdot)$  is the indicator function, and  $\|\hat{d}_t[i]\|$  is the distance between link  $i$  and the object surface. The contact tracking reward is then defined as

$$r_t^c = \sum_i \hat{c}_t[i] \cdot \exp\left(-\frac{\lambda}{f_t[i]}\right),$$

where  $f_t[i]$  is the contact force at link  $i$ .

**Early Termination.** Commonly used in motion tracking, early termination [5] ends an episode if a body part makes unintended ground contact or deviates substantially from the reference, preventing the policy from overvaluing invalid states. For humanoid whole-body loco-manipulation, we introduce additional conditions: (i) the object mesh deviates from its reference beyond a threshold,  $\|\mathbf{P}_t - \hat{\mathbf{P}}_t\|_2 > \sigma_o$ , or (ii) any required body-object contact is lost for more than 10 consecutive frames.

TABLE I: Sim-to-Sim evaluation of ResMimic against baseline approaches in MuJoCo across all four tasks.

Method	Task	SR $\uparrow$	Iter. $\downarrow$	$E_o \downarrow$	$E_m \downarrow$	$E_j \downarrow$
Base Policy	Kneel	0%	—	$0.76 \pm 0.01$	$3.30 \pm 0.53$	$0.28 \pm 0.01$
	Carry	0%	—	$0.29 \pm 0.02$	$2.47 \pm 0.26$	$1.19 \pm 0.30$
	Squat	40%	—	$0.19 \pm 0.01$	$0.93 \pm 0.07$	$0.90 \pm 0.08$
	Chair	0%	—	$1.19 \pm 0.48$	$30.18 \pm 33.45$	$1.20 \pm 0.23$
	Mean	10%	—	0.61	9.22	<b>0.89</b>
Train from Scratch	Kneel	0%	$\times$	$0.69 \pm 0.00$	$5.20 \pm 0.62$	$3.41 \pm 0.07$
	Carry	0%	6500	$0.70 \pm 0.03$	$5.39 \pm 0.38$	$2.33 \pm 0.06$
	Squat	0%	5000	$0.68 \pm 0.05$	$7.56 \pm 2.31$	$4.28 \pm 0.70$
	Chair	0%	2000	$0.97 \pm 0.08$	$10.01 \pm 1.28$	$13.36 \pm 0.92$
	Mean	0%	4500	0.76	7.04	5.84
Finetune	Kneel	0%	$\times$	$0.87 \pm 0.01$	$5.92 \pm 0.81$	$3.02 \pm 0.18$
	Carry	30%	4500	$0.33 \pm 0.01$	$2.49 \pm 0.18$	$2.39 \pm 0.06$
	Squat	0%	2000	$0.47 \pm 0.05$	$5.07 \pm 1.06$	$2.53 \pm 0.13$
	Chair	0%	700	$0.15 \pm 0.01$	$0.28 \pm 0.05$	$1.26 \pm 0.09$
	Mean	7.5%	2400	0.46	3.44	2.30
ResMimic (Ours)	Kneel	90%	2000	$0.14 \pm 0.00$	$0.23 \pm 0.06$	$2.17 \pm 0.06$
	Carry	100%	1000	$0.11 \pm 0.00$	$0.08 \pm 0.00$	$1.24 \pm 0.03$
	Squat	80%	1500	$0.07 \pm 0.01$	$0.07 \pm 0.03$	$1.18 \pm 0.03$
	Chair	100%	700	$0.16 \pm 0.01$	$0.13 \pm 0.02$	$0.55 \pm 0.01$
	Mean	<b>92.5%</b>	<b>1300</b>	<b>0.12</b>	<b>0.13</b>	1.29

## IV. EXPERIMENTS

We assess the effectiveness of ResMimic through a combination of large-scale simulation studies and real-world deployment on a Unitree G1 humanoid robot (29-DoF, 1.3 m tall).

The evaluation is designed to investigate both algorithmic efficiency and deployment robustness. In particular, we focus on the following research questions:

**Q1:** Can a general motion tracking (GMT) policy, without task-specific retraining, accomplish diverse loco-manipulation tasks?

**Q2:** Does initializing from a pre-trained GMT policy improve training efficiency and final performance compared to training from scratch?

**Q3:** When adapting GMT policies to loco-manipulation tasks, is residual learning more effective than fine-tuning?

**Q4:** Beyond simulation, can ResMimic achieve precise,




Fig. 5: Comparison between IsaacGym and MuJoCo results for task *Chair* (left) and *Carry* (right). Corresponding curves quantify object tracking error for **Train from Scratch**, **Finetune**, and **ResMimic**.

expressive, and robust control in the real-world?

#### A. Experiment Setup

1) **Tasks**: We design four challenging whole-body locomanipulation tasks that stress different aspects of humanoid control and generalization. Task (i) *Kneel on one knee and lift a box* — requires expressive, large-amplitude motion and precise lower-body coordination. (ii) *Carry a box onto the back* — demands whole-body expressiveness while maintaining balance under shifting load distribution. (iii) *Squat and lift a box with arms and torso* — highlights the challenge of whole-body contact-rich manipulation. (iv) *Lift up a chair* — involves manipulating a heavy, irregularly shaped object. For clarity, we will refer to these tasks as *Kneel*, *Carry*, *Squat*, and *Chair*, respectively. The human-object interaction reference motions for these tasks are collected using OptiTrack motion capture system.

2) **Evaluation Metrics**: We evaluate ResMimic in terms of training efficiency, motion fidelity, manipulation accuracy, and overall task completion. The following metrics are used throughout our experiments: (i) **Training Iterations (Iter.)**:

To eliminate discrepancies due to hardware differences, we report convergence speed in terms of training iterations rather than wall-clock time. Here we determine convergence when reward stops increasing approximately. (ii) **Object Tracking Error**:  $E_o = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N \|\mathbf{P}[i]_t - \hat{\mathbf{P}}[i]_t\|_2$ , where  $\mathbf{P}_t$  and  $\hat{\mathbf{P}}_t$  denote sampled point clouds on the manipulated object mesh and its reference at time  $t$ . (iii) **Motion Tracking Error**:  $E_m = \frac{1}{T} \sum_{t=1}^T \sum_i \|p_t[i] - \hat{p}_t[i]\|_2$ , where  $p_t[i]$  and  $\hat{p}_t[i]$  are the global positions of link  $i$  at time  $t$ . (iv) **Joint Tracking Error**:  $E_j = \frac{1}{T} \sum_{t=1}^T \|q_t - \hat{q}_t\|_2$ , where  $q_t, \hat{q}_t \in \mathbb{R}^{29}$  denote the robot’s joint angles and their reference values. This evaluates joint-level precision. (v) **Task Success Rate (SR)**: A rollout is considered successful if  $E_o$  is below a predefined threshold and the robot remains balanced.

3) **Baselines**: To validate the effectiveness and efficiency of our pipeline, we compare ResMimic against three representative and strong baselines:

- (i) **Base Policy**: The pre-trained GMT policy is directly deployed to follow human reference motion, without access to object information.
- (ii) **Train from Scratch**: A single-stage RL policy is trained






Fig. 6: Real-world qualitative results comparing ResMimic against all other baselines.

from scratch to track both human motion and object trajectories, without leveraging the GMT policy. For fairness, we use the same reward terms as ResMimic across all tasks, without task-specific tuning.

- (iii) **Base Policy + Fine-tune:** The base GMT policy is fine-tuned to track both human motion and object trajectories. The reward terms are identical to those used in ResMimic. However, due to the limitations of direct fine-tuning, the policy cannot incorporate explicit object information as input.

### B. Sim-to-Sim Evaluation

All policies are trained in IsaacGym to leverage its massive parallelism for accelerated data collection. To evaluate generalization, however, we perform sim-to-sim transfer into MuJoCo, which is widely regarded as a better proxy for real-world physics and thus a closer benchmark for sim-to-real performance. We report quantitative results for all baselines across the four loco-manipulation tasks in Table I, and highlight key takeaways below.

**(Q1) GMT alone cannot complete loco-manipulation tasks, but provides a strong initialization.** We first compare ResMimic against directly deploying the pre-trained GMT policy. As shown in Table I, the base GMT policy achieves only a 10% success rate, compared to 92.5% for ResMimic. Although GMT alone yields slightly lower joint tracking error (due to its training objective), it performs poorly on object tracking and overall task completion, as it lacks access to object information. This indicates that while GMT captures joint-level precision, it is insufficient for manipulation without adaptation.

**(Q2) Using GMT as base policy significantly improves training efficiency and effectiveness.** Next, we compare ResMimic with policies trained entirely from scratch under identical settings (reward function, domain randomization, etc.). As shown in Table I, training from scratch fails to solve the tasks in MuJoCo and converges much more slowly. A side-by-side comparison between IsaacGym and MuJoCo as shown in Figure 5 reveals that trained-from-scratch policies sometimes show partial success in IsaacGym but collapse entirely under sim-to-sim transfer. In contrast, ResMimic maintains strong performance with minimal degradation.




Fig. 7: Ablation on virtual object controller.

This demonstrates the necessity of using GMT as a foundation: its large-scale pretraining imbues generalization and robustness against sim-to-sim gaps.

**(Q3) Residual learning outperforms direct fine-tuning.** Finally, we evaluate residual learning against fine-tuning the GMT policy. While fine-tuning yields slight improvements over training from scratch, it neither outperforms the base GMT policy nor approaches the performance of ResMimic as shown in Table I. A key limitation is that fine-tuning cannot incorporate additional object observations, since the GMT policy architecture is restricted to human motion inputs. Although object-tracking rewards provide some supervision, the lack of explicit object state prevents learning robust behaviors, particularly under randomized object poses. Moreover, fine-tuning tends to overwrite the generalization capability of the GMT policy, leading to instability across tasks. A comparison between IsaacGym and MuJoCo as shown in Figure 5 reveals that fine-tuned policies succeed on task *Chair-Lift* in IsaacGym but fail to transfer in MuJoCo, underscoring the superiority of residual learning as a more generalizable and extensible adaptation strategy.

### C. Real-world Evaluation

As shown in Figure 1, we deploy ResMimic on a Unitree G1 humanoid and demonstrate precise, expressive, robust whole-body loco-manipulation. ResMimic supports both blind (without object state input) and non-blind deployment (with MoCap-based object state input). For simplicity, all real-world results in Figure 1 are under blind deployment.

- **Expressive carrying motions:** the robot kneels on one knee to pick up a box, or carry the box on its back, highlighting expressive whole-body movement.
- **Humanoid-object interaction beyond manipulation:** the robot sits down on a chair and then stands up while maintaining balance and contact with the environment.
- **Heavy payload carrying with whole-body contact:** the robot successfully carries a 4.5kg box, while the G1’s wrist payload limit is around 2.5kg, demonstrating the necessity of leveraging whole-body contact.
- **Generalization to irregular heavy objects:** the robot lifts and carries chairs weighing 4.5kg and 5.5kg, showing instance-level generalization to novel, non-box geometries. We also conduct a qualitative comparison of ResMimic

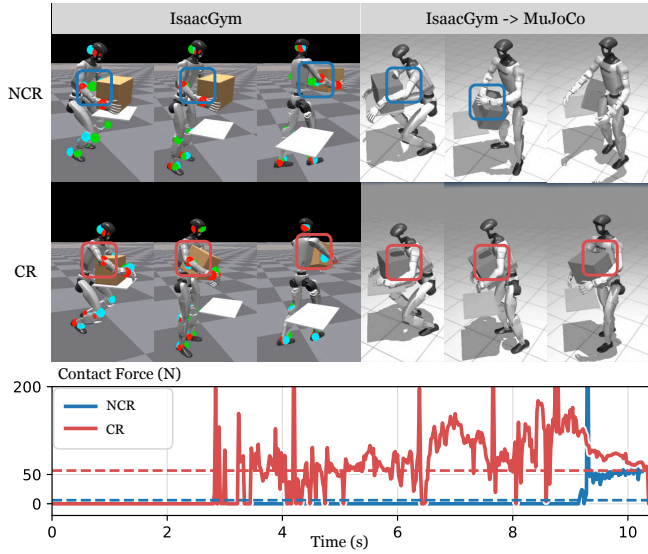


Fig. 8: **Ablation on contact reward.** Here NCR denotes “No Contact Reward”, and CR denotes “with Contact Reward”. Corresponding curves (bottom) quantify torso contact force.

against all baselines in the real world, as shown in Fig. 6. The results indicate that while the base policy can superficially mimic human motion, it lacks object awareness—an issue that becomes even more pronounced when the demonstration data is imperfect. Training from scratch and finetuning, on the other hand, fail entirely due to sim-to-real gap.

Finally, we evaluate ResMimic under non-blind deployment with MoCap-based object state input as shown in Figure 4. In this setting, the robot demonstrates the ability to (i) manipulate objects from random initial poses, (ii) autonomously perform consecutive loco-manipulation tasks, and (iii) show reactive behavior to external perturbations.

#### D. Ablation Studies

1) *Effect of the Virtual Object Controller:* The virtual object controller stabilizes early-stage training by applying curriculum-based virtual forces that guide the object toward its reference trajectory. A qualitative example is shown in Figure 7. In this task, the reference motions contain imperfections, including penetrations between the humanoid’s hand and the object. The policy initially focuses only on motion tracking to reach the object, causing the object to be knocked over, which yields low object rewards and frequent early terminations. This quickly drives the policy into a local minimum where the robot retreats rather than engaging with the object. In contrast, with the virtual force curriculum, the object remains stabilized during early learning, enabling the policy to overcome motion-data imperfections and converge to precise manipulation strategies.

2) *Effect of the Contact Reward:* The contact reward provides explicit guidance on leveraging whole-body strategies. As illustrated in Figure 8, there are two possible ways to lift the box: (1) relying only on wrists and hands, or (2) engaging both torso and arm contact as demonstrated by humans. Without the contact reward, the policy converges to (1), which may succeed in IsaacGym but fails to transfer

to MuJoCo and the real world. With the contact reward, the humanoid instead adopts strategy (2), using coordinated torso and arm contact. This alignment with human demonstrations results in improved sim-to-sim and sim-to-real transfer, validating the importance of the contact reward.

## V. CONCLUSION

This work introduces ResMimic, a two-stage residual learning framework that enables precise, expressive, and robust humanoid loco-manipulation. We first pre-train a general motion-tracking policy on large-scale human motion data, and then refine it with a task-specific residual policy. Across extensive experiments, ResMimic delivers substantial gains in task success, motion fidelity, and training efficiency, and further demonstrates seamless deployment on a real Unitree G1 humanoid. These results highlight the transformative potential of pre-trained policies for humanoid control.

## REFERENCES

- [1] I. Radosavovic, B. Zhang, B. Shi, J. Rajasegaran, S. Kamat, T. Darrell, K. Sreenath, and J. Malik, “Humanoid locomotion as next token prediction,” *arXiv:2402.19469*, 2024.
- [2] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, “Learning humanoid locomotion with transformers,” *arXiv:2303.03381*, 2023.
- [3] J. Li, Y. Zhu, Y. Xie, Z. Jiang, M. Seo, G. Pavlakos, and Y. Zhu, “Okami: Teaching humanoid robots manipulation skills through single video imitation,” in *8th Annual Conference on Robot Learning*, 2024.
- [4] Y. Ze, Z. Chen, W. Wang, T. Chen, X. He, Y. Yuan, X. B. Peng, and J. Wu, “Generalizable humanoid manipulation with 3d diffusion policies,” *arXiv preprint arXiv:2410.10803*, 2024.
- [5] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, “Deepmimic: Example-guided deep reinforcement learning of physics-based character skills,” *ACM Trans. Graph.*, vol. 37, no. 4, pp. 143:1–143:14.
- [6] Y. Ze, Z. Chen, J. P. Araújo, Z. ang Cao, X. B. Peng, J. Wu, and C. K. Liu, “Twist: Teleoperated whole-body imitation system,” *arXiv preprint arXiv:2505.02833*, 2025.
- [7] Z. Chen, M. Ji, X. Cheng, X. Peng, X. B. Peng, and X. Wang, “Gmt: General motion tracking for humanoid whole-body control,” *arXiv:2506.14770*, 2025.
- [8] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black, “AMASS: Archive of motion capture as surface shapes,” in *International Conference on Computer Vision*, Oct. 2019.
- [9] J. Li, J. Wu, and C. K. Liu, “Object motion guided human motion synthesis,” 2023. [Online]. Available: <https://arxiv.org/abs/2309.16237>
- [10] J. Dao, H. Duan, and A. Fern, “Sim-to-real learning for humanoid box loco-manipulation,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.03191>
- [11] F. Liu, Z. Gu, Y. Cai, Z. Zhou, S. Zhao, H. Jung, S. Ha, Y. Chen, D. Xu, and Y. Zhao, “Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation,” *arXiv preprint arXiv:2409.20514*, 2024.
- [12] OpenAI, “Gpt-4 technical report,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [13] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mohtakuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky, “ $\pi_0$ : A vision-language-action flow model for general robot control,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.24164>
- [14] T. L. Team, “A careful examination of large behavior models for multitask dexterous manipulation,” 2025. [Online]. Available: <https://arxiv.org/abs/2507.05331>
- [15] D. Roich, R. Mokady, A. H. Bermano, and D. Cohen-Or, “Pivotal tuning for latent-based editing of real images,” 2021. [Online]. Available: <https://arxiv.org/abs/2106.05744>
- [16] R. Gal, O. Patashnik, H. Maron, G. Chechik, and D. Cohen-Or, “Stylegan-nada: Clip-guided domain adaptation of image generators,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.00946>



- [17] C. Sferrazza, D.-M. Huang, X. Lin, Y. Lee, and P. Abbeel, "Humanoidbench: Simulated humanoid benchmark for whole-body locomotion and manipulation," 2024. [Online]. Available: <https://arxiv.org/abs/2403.10506>
- [18] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, "Real-world humanoid locomotion with reinforcement learning," *Science Robotics*, vol. 9, no. 89, 2024.
- [19] Z. Zhuang, S. Yao, and H. Zhao, "Humanoid parkour learning," in *8th Annual Conference on Robot Learning*, 2024.
- [20] T. Huang, J. Ren, H. Wang, Z. Wang, Q. Ben, M. Wen, X. Chen, J. Li, and J. Pang, "Learning humanoid standing-up control across diverse postures," *arXiv preprint arXiv:2502.08378*, 2025.
- [21] T. Zhang, B. Zheng, R. Nai, Y. Hu, Y.-J. Wang, G. Chen, F. Lin, J. Li, C. Hong, K. Sreenath, and Y. Gao, "Hub: Learning extreme humanoid balance," 2025. [Online]. Available: <https://arxiv.org/abs/2505.07294>
- [22] T. He, J. Gao, W. Xiao, Y. Zhang, Z. Wang, J. Wang, Z. Luo, G. He, N. Sobanbabu, C. Pan, Z. Yi, G. Qu, K. Kitani, J. Hodgins, L. J. Fan, Y. Zhu, C. Liu, and G. Shi, "Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills," *arXiv preprint arXiv:2502.01143*, 2025.
- [23] Q. Liao, T. E. Truong, X. Huang, G. Tevet, K. Sreenath, and C. K. Liu, "Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion," 2025. [Online]. Available: <https://arxiv.org/abs/2508.08241>
- [24] J. Mao, S. Zhao, S. Song, T. Shi, J. Ye, M. Zhang, H. Geng, J. Malik, V. Guizilini, and Y. Wang, "Learning from massive human videos for universal humanoid pose control," 2024. [Online]. Available: <https://arxiv.org/abs/2412.14172>
- [25] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang, "Expressive whole-body control for humanoid robots," *arXiv preprint arXiv:2402.16796*, 2024.
- [26] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. Kitani, C. Liu, and G. Shi, "Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning," *arXiv preprint arXiv:2406.08858*, 2024.
- [27] A. Allshire, H. Choi, J. Zhang, D. McAllister, A. Zhang, C. M. Kim, T. Darrell, P. Abbeel, J. Malik, and A. Kanazawa, "Visual imitation enables contextual humanoid control," *arXiv preprint arXiv:2505.03729*, 2025.
- [28] J. Li, X. Cheng, T. Huang, S. Yang, R. Qiu, and X. Wang, "Amo: Adaptive motion optimization for hyper-dexterous humanoid whole-body control," *Robotics: Science and Systems 2025*, 2025.
- [29] Q. Ben, F. Jia, J. Zeng, J. Dong, D. Lin, and J. Pang, "Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit," *arXiv preprint arXiv:2502.13013*, 2025.
- [30] Y. Zhang, Y. Yuan, P. Gurunath, T. He, S. Omidshafiei, A.-a. Aghamohammadi, M. Vazquez-Chanlatte, L. Pedersen, and G. Shi, "Falcon: Learning force-adaptive humanoid loco-manipulation," *arXiv preprint arXiv:2505.06776*, 2025.
- [31] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, "Humanplus: Humanoid shadowing and imitation from humans," 2024. [Online]. Available: <https://arxiv.org/abs/2406.10454>
- [32] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, and S. Levine, "Residual reinforcement learning for robot control," 2018. [Online]. Available: <https://arxiv.org/abs/1812.03201>
- [33] T. Silver, K. Allen, J. Tenenbaum, and L. Kaelbling, "Residual policy learning," 2019. [Online]. Available: <https://arxiv.org/abs/1812.06298>
- [34] M. Alakuijala, G. Dulac-Arnold, J. Mairal, J. Ponce, and C. Schmid, "Residual reinforcement learning from demonstrations," 2021. [Online]. Available: <https://arxiv.org/abs/2106.08050>
- [35] S. Haldar, J. Pari, A. Rai, and L. Pinto, "Teach a robot to fish: Versatile imitation from one minute of demonstrations," 2023. [Online]. Available: <https://arxiv.org/abs/2303.01497>
- [36] K. Li, P. Li, T. Liu, Y. Li, and S. Huang, "Maniptrans: Efficient dexterous bimanual manipulation transfer via residual learning," 2025. [Online]. Available: <https://arxiv.org/abs/2503.21860>
- [37] S. Zhao, X. Zhu, Y. Chen, C. Li, X. Zhang, M. Ding, and M. Tomizuka, "Dexh2r: Task-oriented dexterous manipulation from human to robots," 2024. [Online]. Available: <https://arxiv.org/abs/2411.04428>
- [38] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [39] Y. Ze, J. P. Araújo, J. Wu, and C. K. Liu, "Gmr: General motion retargeting," 2025, *gitHub repository*.
- [40] L. Ankile, A. Simeonov, I. Shenfeld, M. Torne, and P. Agrawal, "From imitation to refinement – residual rl for precise assembly," 2024. [Online]. Available: <https://arxiv.org/abs/2407.16677>
- [41] Z. Mandi, Y. Hou, D. Fox, Y. Narang, A. Mandlekar, and S. Song, "Dexmachina: Functional retargeting for bimanual dexterous manipulation," *arXiv preprint arXiv:2505.24853*, 2025.
- [42] S. Xu, H. Y. Ling, Y.-X. Wang, and L.-Y. Gui, "Intermimic: Towards universal whole-body control for physics-based human-object interactions," *arXiv preprint arXiv:2502.20390*, 2025.