

HumanoidExo: Scalable Whole-Body Humanoid Manipulation via Wearable Exoskeleton

Rui Zhong^{*1}, Yizhe Sun⁴, Junjie Wen³, Jinming Li⁵, Chuang Cheng^{†1},
Wei Dai¹, Zhiwen Zeng¹, Huimin Lu^{†1}, Yi Xu⁶, Yichen Zhu^{*†2}

<https://humanoid-exo.github.io/>

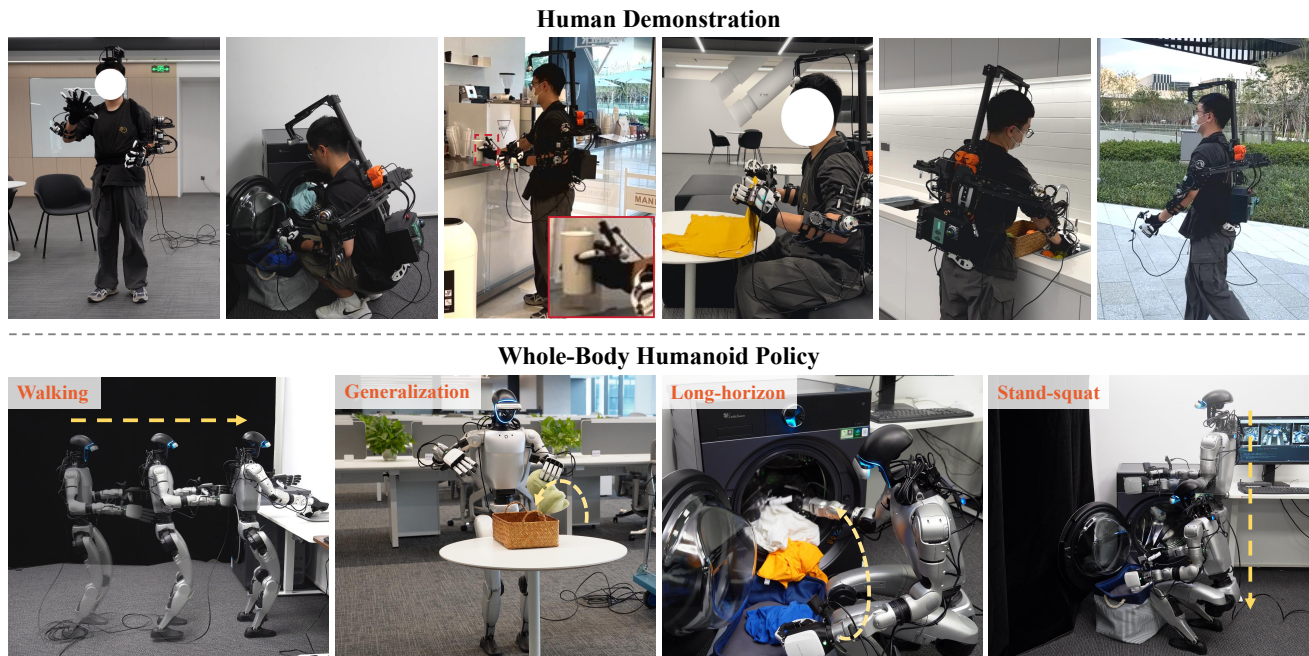


Fig. 1: **HumanoidExo**, a wearable exoskeleton system that transfers human motion to **whole-body** humanoid data. HumanoidExo provides an efficient solution that bridges the embodiment gap between humans and robots, facilitating the collection of diverse datasets. We tested it on three real-world tasks, and results show that it helps humanoid robots generalize to new environments, learn complex tasks from limited data, and acquire new skills like walking.

Abstract—A significant bottleneck in humanoid policy learning is the acquisition of large-scale, diverse datasets, as collecting reliable real-world data remains both difficult and cost-prohibitive. To address this limitation, we introduce *HumanoidExo*, a novel system that *transfers human motion to whole-body humanoid data*. HumanoidExo offers a high-efficiency solution that minimizes the embodiment gap between the human demonstrator and the robot, thereby tackling the scarcity of whole-body humanoid data. By facilitating the collection of more voluminous and diverse datasets, our approach significantly enhances the performance of humanoid robots in dynamic, real-world scenarios. We evaluated our method across three challenging real-world tasks: table-top manipulation, manipulation integrated with stand-squat motions, and whole-body manipulation. Our results empirically demonstrate that HumanoidExo is a crucial addition to real-

robot data, as it enables the humanoid policy to generalize to novel environments, learn complex whole-body control from only five real-robot demonstrations, and even acquire new skills (i.e., walking) solely from HumanoidExo data.

I. INTRODUCTION

Humanoid policy learning is a rapidly advancing field, now spanning locomotion, manipulation, and language-conditioned tasking. This progress is largely propelled by foundational model initiatives for general-purpose humanoids, such as Nvidia’s GR00T [1] and Figure AI. To mitigate the high cost of real-robot demonstrations, researchers have introduced several data-efficient pipelines. These include sim-to-real transfers, learning from web-scale human videos (e.g., EgoMimic [2]), and the development of diverse teleoperation systems for more effective data collection.

Despite these efforts, scaling humanoid data collection remains a significant challenge for two primary reasons. First, simulation and human video data both suffer from severe embodiment gaps. Simulated robot dynamics inevitably mis-

¹National University of Defense Technology, ²University of Toronto, ³East China Normal University, ⁴Tianjin University of Science and Technology, ⁵Shanghai University, ⁶Midea Group.

^{*}Equal contributions. [†]Corresponding authors. This work was done while Rui Zhong, Yizhe Sun, Junjie Wen, Jinming Li and Yichen Zhu were at Midea Group.

match their real-world counterparts, while the morphological and kinematic differences between humans and robots make direct video-to-policy transfer notoriously difficult. Second, direct teleoperation is difficult to scale. This approach typically requires a one-to-one, human-to-robot setup, which is expensive and resource-intensive. Furthermore, the process is physically and mentally demanding, leading to operator fatigue that limits the duration of collection sessions and requires highly skilled personnel. This reliance on expert operators and the inherent risk of damaging costly hardware during live sessions create a major bottleneck for generating large-scale datasets.

In this work, we introduce **HumanoidExo**, an integrated system that advances both data collection and policy learning for humanoid robotics. We utilize a custom-designed, lightweight, and flexible wearable exoskeleton [3] to capture human motion without impeding the operator’s natural movements. This design enables the comfortable performance of diverse daily tasks, while our system records and translates the operator’s actions into structured data for robot learning. To capture comprehensive, whole-body motion, a back-mounted LiDAR sensor tracks the operator’s torso, providing a 6D pose to record base movements such as walking, squatting, and bending. By fusing data from the exoskeleton and LiDAR, our system generates kinematically feasible, whole-body trajectories ready for large-scale policy learning.

To leverage this data, we present **HumanoidExo-VLA (HE-VLA)** in short, a refined, vision-language-action model for whole-body humanoid policy learning from exoskeleton data. This hybrid approach uses imitation learning as a foundation and incorporates reinforcement learning to ensure the robot maintains balance and stability during movement and manipulation. The synergy between our hardware and software enables stable, efficient policy learning from exoskeleton data, leading to policies that are directly deployable on a physical humanoid robot.

To validate the effectiveness of our method, we conduct a comprehensive study across three challenging real-world tasks. These tasks include table-top manipulation, dexterous manipulation involving stand-squat motions, and whole-body manipulation that requires walking to a table. Our experimental results highlight the critical role of **HumanoidExo** data in enhancing policy performance: 1) **Generalization**: It enables the learned policy to generalize effectively to novel scenes and environments. 2) **Data Efficiency**: It allows an end-to-end model to learn complex tasks with as few as five real-robot demonstrations. 3) **Skill Acquisition**: It empowers the humanoid robot to acquire entirely new skills (e.g., walking) using only data from the exoskeleton, without any real-robot demonstrations. We believe that **HumanoidExo** represents a significant step toward achieving scalable whole-body humanoid policy learning.

II. RELATED WORK

A. Data Collection with In-The-Wild Equipment

To enable scalable robot data collection, several in-the-wild systems have recently been proposed to make the

process more affordable and accessible. These systems have shown great promise in specific domains. For instance, **Human Policy** [4] focuses on transferring human action primitives to bridge embodiment differences. **DexCap** [5] uses a wearable glove to capture precise wrist and fingertip poses for dexterous tasks. **AirExo** [6], [7] leverages low-cost hardware with direct kinematic mapping for arm manipulation. The **Universal Manipulation Interface (UMI)** [8] introduced a simple handheld controller for collecting bimanual data at scale, which **DexUMI** [9] later extended to dexterous hands. However, these pioneering works share two common limitations. First, they are primarily designed for conventional robotic arms. Second, their focus is typically restricted to upper-body manipulation. In contrast, our work, **HumanoidExo**, addresses these gaps directly. To the best of our knowledge, this is the first system designed for in-the-wild, whole-body policy learning on a humanoid robot. We introduce multiple techniques to bridge the embodiment gap between humans and humanoids, effectively translating natural human motion into executable robot policies. For the first time, we demonstrate that this paradigm can successfully scale up the acquisition of whole-body training data, paving the way for more capable and generalist humanoid robots.

B. Humanoid Whole-Body Manipulation

Reinforcement learning algorithms for whole-body locomotion of humanoid robots have been extensively studied [10], [11], and whole-body teleoperation of humanoid robots can be achieved through various methods [12]–[19]. In robotic manipulation, imitation learning [8], [20]–[23] with rich visual representations, especially from robot foundation models like **Vision-Language-Action (VLA)** models [24]–[27], has become a central paradigm. This line of work typically focuses on humanoid manipulation [28]; for instance, **iDP3** [29] trains egocentric 3D diffusion policies, and **Dexmimicgen** [30] learns bimanual skills via a sim-to-real strategy. These methods mainly rely on teleoperated data collected directly from the physical humanoid robot. In contrast, our approach, **HumanoidExo**, focuses on learning humanoid control through a novel mixture of limited real robot data and a larger dataset collected from a human operator wearing an exoskeleton. More importantly, our method is not confined to manipulation but extends to complex whole-body control tasks such as standing, squatting, and walking.

III. METHODOLOGY

This section provides a comprehensive technical overview of our proposed **HumanoidExo** system. We begin by elucidating the foundational concept of our approach, which is designed to bridge the significant embodiment gap between a human demonstrator and a humanoid robot. Then, we detail the specific implementation for data acquisition, breaking it down into our methodologies for capturing and refining upper-body motion and for tracking lower-body dynamics. Subsequently, we describe the motion retargeting pipeline that integrates these data streams and adjusts the complete whole-body trajectories for effective robot data. To conclude,

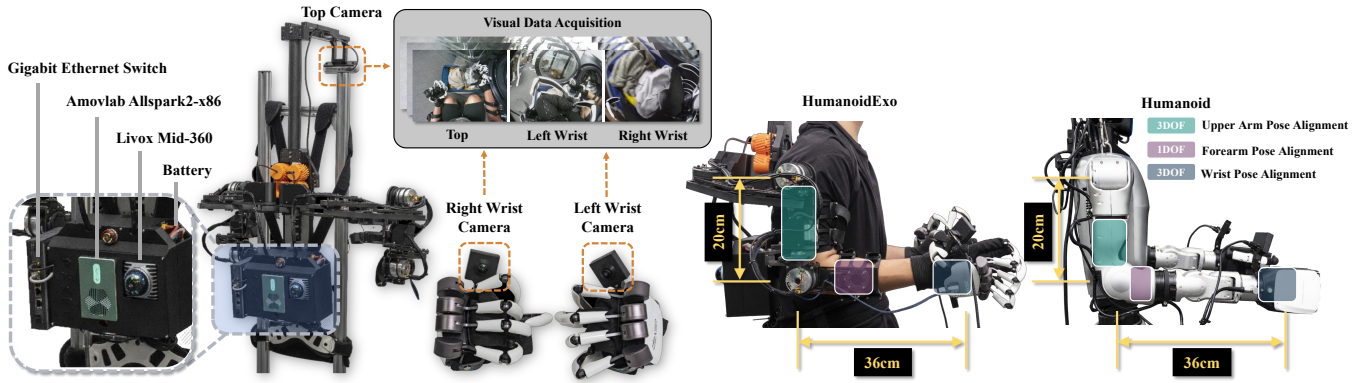


Fig. 2: **Hardware overview for HumaniodExo.** We integrated a Mid-360 LiDAR for acquiring exoskeleton motion odometry. For visual information acquisition, we added two wrist cameras to capture new operational perspectives and enrich environmental perception. These cameras, installed on the Dexmo force-feedback gloves, were mounted at angles identical to those of the robot’s cameras. Since the HumanoidExo system adopts a joint space control method with angle remapping, we redesigned the exoskeleton’s key parameters to match the arm length of the Unitree G1 robot. In addition, we recruited data collectors with upper-body dimensions similar to the G1’s anthropometric parameters for data collection and teleoperation, minimizing end-effector errors arising from dimensional mismatches.

we specify the hardware configuration of the target humanoid platform used in our experiments.

A. Bridging the Embodiment Gap Between Human and Humanoid

One shortcoming that preventing scaling humanoid data with exoskeleton or universal interface control is let in the control mapping methods. Specifically, robot arm control typically involves two primary methods: Cartesian space control (end-effector control) and joint space control. For devices such as the Meta Quest or Apple Vision Pro, which are capable of capturing the end-effector pose, Cartesian control is often an appropriate choice. However, in scenarios where precise arm joint angles are required—such as when placing clothes in a washing machine—solely tracking the end-effector pose without considering the other arm joints can easily result in collisions with the environment. Moreover, since a typical humanoid robot arm has seven degrees of freedom (DoF), its inverse kinematics has multiple solutions, leading to a null space. End-effector control in such cases necessitates additional constraints, making real-time and accurate joint angle computation particularly challenging in small, confined spaces. Therefore, it is more practical to directly map the joint space, aligning human joint movements with those of the robot.

B. Upper-Body Alignment between HumanoidExo and Robot

HumanoidExo is specifically designed to read all seven joints of the human arm. The rotational axes of its exoskeleton arm are precisely aligned with the corresponding axes of the human joints, making the exoskeleton isomorphic to the human arm. Additionally, we have introduced two extra DOF at the Glenohumeral (GH) joint, enhancing the system’s ergonomics and making it more suitable for daily wear. Due to these extra DoF and the incorporation of a new link system, HumanoidExo’s joint configuration differs

significantly from that of a typical humanoid robot arm. The motion retargeting algorithm used for this process is outlined below.

Upper Arm Pose Alignment: The exoskeleton’s base is used as the reference coordinate system, fixed to the wearer’s torso. To align the upper arm pose, we first extract all rotational joints leading to the upper arm attachment point. This includes both active joints driven by motors and passive joints within the linkages and timing belts. By analyzing the relative positions of these joints, we construct a Denavit-Hartenberg (DH) parameter table [3]. Using forward kinematics, we then compute the relative pose between the upper arm attachment point and the base, expressed as a quaternion $q^{upper\ arm}$. Subsequently, an iterative inverse kinematics method is employed to map the upper arm’s pose onto the first three joints of the robot’s arm, thereby achieving upper arm pose alignment.

Forearm Pose Alignment: The elbow is the only moving joint between the human upper arm and forearm. HumanoidExo is equipped with a motor at the elbow, enabling direct joint space control. This motor allows us to map the human elbow’s bending angle directly to the robot’s forearm, achieving forearm pose alignment.

Wrist Pose Alignment: The human wrist can be treated as a spherical joint [31]. To simplify the structural design and reduce the wrist’s weight while still enabling the exoskeleton to capture wrist pose, we utilize Inertial Measurement Units (IMUs) as the joint data source. To capture the wrist’s three rotational DoF relative to the forearm, we place one IMU on the forearm and another on the Dexmo force feedback glove. First, we record the quaternions of the wrist IMU, q^w , and the forearm IMU, q^f . Upon initializing the exoskeleton, the operator returns the wrist to the home position. Now record both IMU readings at this time q^{w_0} and q^{f_0} . The rotation quaternion q^{wrist} representing the wrist’s orientation relative

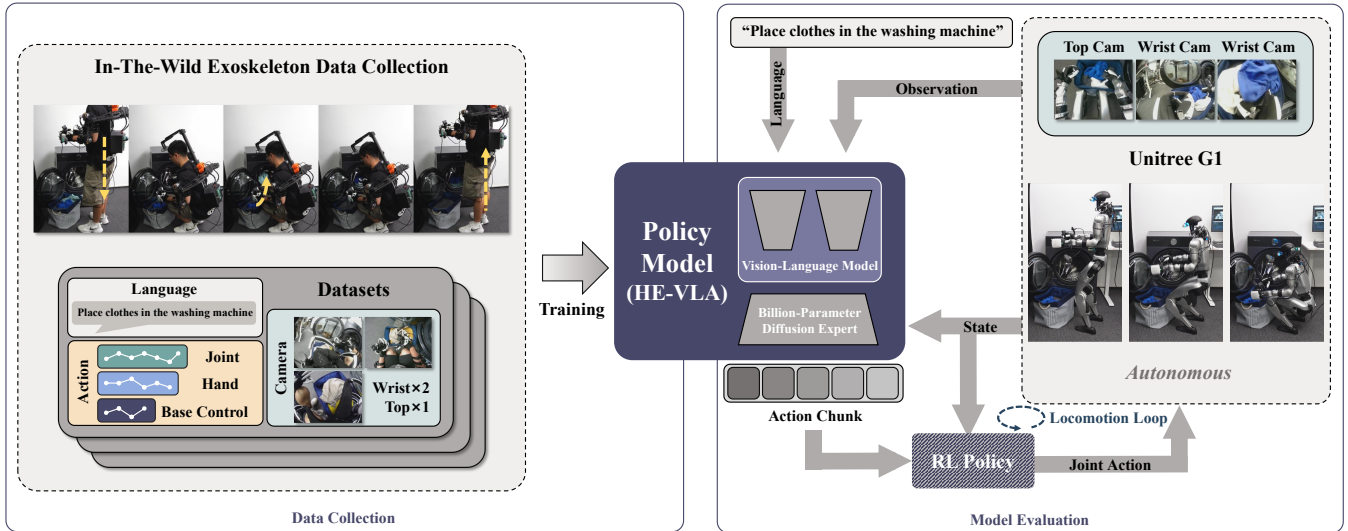


Fig. 3: **The overview of the HE-VLA.** The left side of the figure illustrates our data collection process and the composition of the dataset, with the primary source of training data being in-the-wild data gathered by the HumanoidExo independently of the controlled robot. The right side of the figure presents the model deployment and inference pipeline. We employ two systems to control the robot: a Vision-Language-Action (VLA) model and a Reinforcement Learning (RL) model. The VLA model generates high-level control commands and transmits them to the lower-level RL model. The RL model is then responsible for maintaining balance control and executing the joint movement commands.

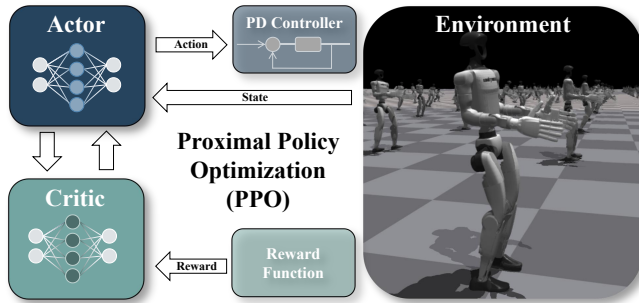


Fig. 4: **The actor-critic reinforcement learning in HE-VLA.** This module works in conjunction with the primary VLA model, guaranteeing the humanoid can reliably stand, squat, and walk during policy inference.

to the forearm at any given moment:

$$q^{wrist} = q^t q^{t_0} = q^{f*} q^w q^{w_0*} q^{f_0}$$

where q^* represents the conjugate of the quaternion. This quaternion serves as the target pose for the wrist, and an iterative inverse kinematics method is employed once more to map this pose to the final three joints of the robot’s arm, achieving wrist pose alignment.

C. Lower-Body Alignment between HumanoidExo and Robot

For humans, lower-limb movements such as walking, squatting, and standing are fundamental to daily activities, as they directly determine mobility, flexibility, and the ability to interact with complex environments. Motivated by this, our system adopts base translational velocity, base rotational velocity, and height as the command for the RL-based balance

control algorithm. These variables capture the essential DoF of human lower-body motion, enabling intuitive and effective control while also providing a foundation for high-level task planning.

Beyond supporting whole-body teleoperation of humanoid robots, our HumanoidExo system is designed with an independent data collection capability. Specifically, it can perceive and record its own motion state in real time—including displacement, orientation, and variations in body height—without depending on external operator input.

To realize this, we integrated a Mid-360 LiDAR on the back of the exoskeleton and developed a LiDAR-odometry module based on FAST-LIO [32] for real-time estimation of the system’s spatial position and orientation. Compared with approaches relying on vision or external localization beacons, LiDAR odometry provides superior environmental adaptability and robustness to lighting conditions. It maintains high localization accuracy even under drastic illumination changes or in texture-sparse environments. Moreover, it eliminates the need for external infrastructure, thereby enabling truly independent operation and reliable data collection in open, unstructured, or entirely unknown in-the-wild scenarios.

The specific mounting position and configuration of the LiDAR are illustrated in Fig. 2 (left). The placement was carefully designed to balance field of view, resilience to interference, and overall motion stability of the system.

D. HE-VLA: A Whole-Body Humanoid Policy Learning Method

This section introduces our method for learning a whole-body humanoid policy from our collected exoskeleton data. Our approach, namely **HumanoidExo-VLA (HE-VLA in**

short), consists of two key components: a pre-trained Vision-Language-Action (VLA) model that learns foundational whole-body motion control, and a reinforcement learning method that ensures robust whole-body balance.

Vision-Language-Action Model. Given a set of expert demonstrations that contain complex humanoid skill trajectories, we want to learn a visuomotor policy $\pi : \mathcal{O} \mapsto \mathcal{A}$ that maps the visual observations $o \in \mathcal{O}$ to actions $a \in \mathcal{A}$, such that our robots not only *reproduce* the skill but also *generalize* beyond the training data. To tackle the challenge of manipulating complex humanoid skills, we leverage DexVLA [24], a pre-trained vision-language-action model, for the tasks described in our experiments. Notably, the pre-training dataset for this foundation model does not include data from the humanoid robot used in our experiments, the Unitree G1. By leveraging a pre-trained robot foundation model, we ensure its parameters are effectively warmed up, facilitating an easier acquisition of complex real-world skills. However, we observe that finetuning the model alone does not guarantee stable whole-body control. In fact, the fine-tuned model often causes the robot to fall, failing to maintain whole-body stability due to instability induced by dynamic upper-body motions. Therefore, we introduce an improvement to address this critical problem.

Reinforcement Learning for Motion Balance. Relying solely on imitation learning to directly output joint positions for whole-body control introduces significant stability risks. Minor deviations from the learned trajectories can result in falls, posing a threat of catastrophic damage to the robot and its environment. To overcome this limitation, we leverage reinforcement learning to train a robust whole-body locomotion controller. This controller is responsible for maintaining dynamic balance while executing commands for base speed, yaw rate, and a target torso height.

Specifically, at each step, the policy receives $\mathbf{o}_t = [c_t, \boldsymbol{\omega}_t^b, \mathbf{g}_t^b, \mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{a}_{t-1}]$, and $c_t = [v_x^*, \omega_z^*, h^*]$, where $\boldsymbol{\omega}_t^b$ and \mathbf{g}_t^b are base angular velocity and gravity in the torso frame, and $(\mathbf{q}_t, \dot{\mathbf{q}}_t)$ are joint states. The action \mathbf{a}_t specifies desired lower-body joint targets $\mathbf{q}^{\text{des}} = \mathbf{q}^0 + \mathbf{a}_t$. Joint torques are applied by a PD law

$$\tau_i = k_{p,i}(q_i^{\text{des}} - q_i) - k_{d,i}\dot{q}_i,$$

driving locomotion while the upper-body joints follow the operator directly.

To achieve stable walking with dynamic upper-body movement, we use a curriculum that scales the admissible joint range by a ratio $r \in [0, 1]$. This ratio increases as the agent succeeds, smoothly expanding motion from static to fully expressive. Commanded squats are realized through a reward function that tracks a target base height and shapes knee flexion accordingly. We train this by dedicating one-third of our parallel environments to squatting and the rest to walking, periodically switching their roles to learn seamless transitions. Each transition is mirrored across the robot’s x - z plane (swap left/right states and actions and flip the yaw command), and both original and mirrored samples are stored. Auxiliary actor/critic symmetry losses encourage

consistent predictions on mirrored pairs, improving sample efficiency and reducing unintended left-right bias. As a result, the learned policy walks, turns, and squats to commanded heights while remaining stable under continuously changing operator-driven upper-body poses.

E. Robot Configuration

Camera View. The visual data acquisition system is composed of a primary head-mounted camera (Realsense D455) and two supplementary wrist-mounted fisheye cameras, which offer expanded perspectives during manipulation. To mitigate the visual embodiment gap, we standardized the hardware by using the same camera models on both the exoskeleton and the robot. Furthermore, the field of view (FoV) of the wrist cameras on the exoskeleton is carefully calibrated to match that of the robot’s cameras.

We identified two primary challenges related to the head camera. First, the Unitree G1 robot features a non-actuated head, making it impossible to translate the human operator’s head movements into robot camera motion. Second, a perfect positional correspondence between the exoskeleton’s and the robot’s head cameras is not achievable. Nevertheless, our empirical results demonstrate that these discrepancies do not hinder policy learning; the model proves capable of extracting salient visual features even when trained with the non-stationary head camera view from the human demonstrator.

Robot Setup. Our experimental platform is the Unitree G1, a humanoid robot with a height of 1.3m and 29 DoF in its body. The robot is equipped with two Inspire-Hand end-effectors. Each hand is an underactuated gripper with 12 DoF (six active and six passive). The thumb contains two active and two passive DoFs, while each of the remaining fingers has one active and one passive DoF. As a result, the entire robotic system has a total of 41 active degrees of freedom. Our policy model outputs target joint angles to control the robot’s entire body.

IV. EXPERIMENT

In the experiments section, we aim to answer the following questions to demonstrate the effectiveness of HumanoidExo:

- How does the HumanoidExo system compare to state-of-the-art policy learning methods?
- Does training with HumanoidExo data improve the policy’s robustness to environmental variations and physical disturbances?
- Can HumanoidExo enable the learning of novel and complex skills for humanoid robots?
- To what extent does exoskeleton data facilitate learning skills that involve mobility and locomotion?

A. Task Descriptions & Implementation Details

As shown in Fig. 5, HumanoidExo performs on challenging whole-body humanoid manipulation tasks. We conducted experiments on three tasks that represent different categories of humanoid skills.

- **PlaceToy:** This is a tabletop manipulation task. The robot is required to pick up a toy, whose position is

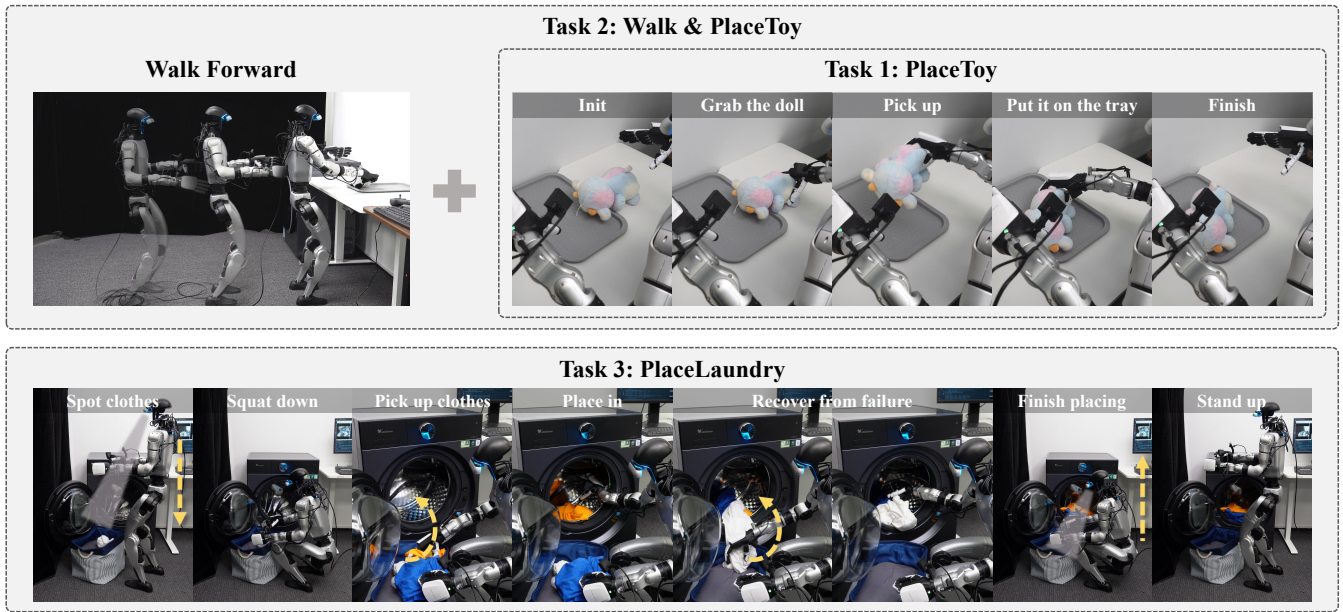


Fig. 5: Examples for PlaceToy (Task 1), Walk & PlaceToy (Task 2), and PlaceLaundry (Task 3). We designed three tasks to showcase the effectiveness of HumanoidExo in robot skill learning: Task 1 tests dexterity, Task 2 combines locomotion and manipulation, with the mobile-manipulation data entirely collected by HumanoidExo, and Task 3 involves whole-body manipulation of the humanoid robot.

randomized on its left or right side, and place it into a tray at the center. This task tests dexterity without requiring locomotion.

- **Walk & PlaceToy:** This task combines locomotion with manipulation. The humanoid is required to walk to a table, stop, and then place a toy into a tray at the center.
- **PlaceLaundry:** This is a whole-body manipulation task. The humanoid must squat down, pick up an article of clothing from a basket, place it into a washing machine on its right, and then return to a standing position. The primary challenge is maintaining balance during the dynamic squatting and standing motions.

We give examples for each task in Figure 5. For all experiments, we fine-tune the model for 50k iterations with a batch size of 128. We use a cosine learning rate scheduler with an initial learning rate of $2e-5$. All models and experiments use the same set of hyperparameters to ensure fair comparison.

B. Generalizable Manipulation with HumanoidExo

The experiments in this section aim to measure the efficiency of HumanoidExo data and the generalizability that such data bring.

Fig. 6 summarizes the average success rates for the tabletop manipulation task across three experimental setups: 1) training with 200 real teleoperated demonstrations; 2) mixed training with 5 teleoperated and 195 HumanoidExo demonstrations; 3) training with only 5 teleoperated demonstrations. This experiment aims to evaluate the model’s performance under an extreme data-scarce scenario, where real data constitutes just 2.5% of the mixed dataset.

The results reveal two key findings. First, HumanoidExo

data provides a significant performance boost when augmenting a small set of real demonstrations. For instance, while training with only five teleoperated demonstrations yields a mere 5% success rate, adding 195 HumanoidExo demonstrations boosts the success rate dramatically to about 80%. Second, for this in-domain task, the performance achieved with the mixed dataset suggests that HumanoidExo data has a utility comparable to that of real-robot data. This is particularly significant because collecting HumanoidExo data is substantially more cost-effective and scalable than robot teleoperation, highlighting the practical value and efficiency of our method.

Object & Environment Generalization. A key metric for evaluating a policy is its ability to generalize to novel objects and unseen environments. To assess this, we conduct two generalization tests: one involving four objects not present in the training data, and another in a completely new environment with different objects (a novel toy and tray).

Across all generalization settings, we observe that a mixed training approach—combining teleoperated data with HumanoidExo data—yields a higher average success rate than training solely on teleoperated data. These findings demonstrate the effectiveness of our proposed system in improving policy generalization.

C. Exoskeleton Data Brings New Skill

Given that HumanoidExo data can serve as an effective substitute for teleoperated data, a natural question arises: can it also be used to teach the robot entirely new skills?

To investigate this, we designed an experiment that builds upon the previous tabletop manipulation setup. We collected

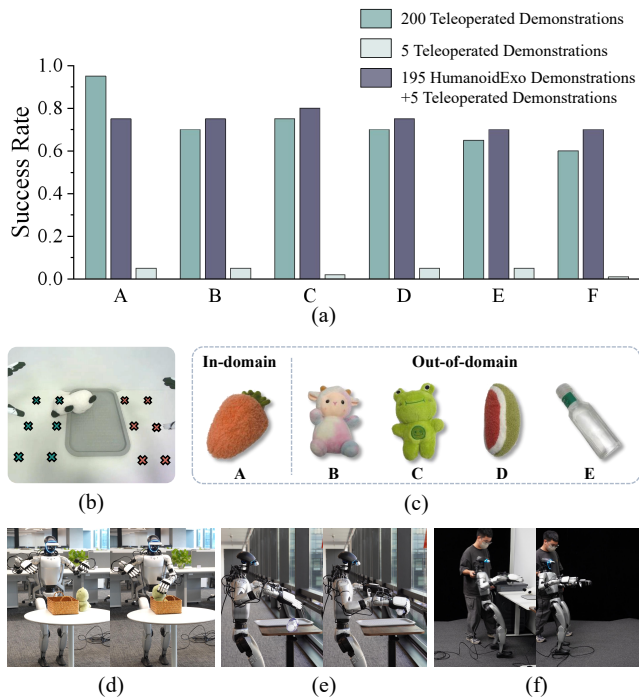


Fig. 6: **Model Generalization.** (a) **Model success rates.** Labels A-E correspond to the robot’s success rate for grasping items A-E shown in (c), and the number of trials for each experiment is 60. (b) **Item placement locations for model testing.** (c) **Training datasets.** Out-of-domain data represents items that did not appear in the Teleoperated Demonstrations but were present in the HumanoidExo Demonstrations. (d)&(e) **Tasks in new environment.** The robot’s success rate for completing the task is represented by the label F in (a). (f) **Robustness to disturbance.** The robot could autonomously walk back to the table and resume the tabletop task (Task 2) after being forcibly moved away.

195 new HumanoidExo demonstrations of a compound task: walking to the table, stopping, and then executing the ‘place toy’ action. The policy was then trained on a mixed dataset containing these 195 new HumanoidExo demonstrations and the same five teleoperated demonstrations from the previous experiment. Crucially, these five teleoperated demonstrations only contain the stationary manipulation portion of the task; they include no walking. Therefore, any walking ability exhibited by the final policy must be learned exclusively from the HumanoidExo data. The task is shown in Figure 5.

For this combined task, the policy achieved a 100% success rate on the walking portion in every trial. The robot consistently navigated to the table and stopped in the correct position before initiating the manipulation phase. Consequently, the overall task success rate is identical to that of the stationary pick-and-place experiment. This result demonstrates that our method successfully empowered the robot with the new skill of walking without degrading its previously learned manipulation capabilities. This promising phenomenon motivated us to conduct further tests on the

robustness and generalization of this newly acquired skill, which was learned exclusively from HumanoidExo data.

Robustness to Disturbance. An interesting observation from our experiments is the trained policy’s remarkable robustness to physical disturbances. As demonstrated in Figure 6(f), when we manually drag the Unitree G1 away from its workspace, the policy consistently recovers by walking the robot back to the table.

Notably, the model can recover even when displaced to distances far exceeding any seen in the HumanoidExo training data. It successfully navigates back to its station in front of the table before proceeding to complete the task. This behavior suggests that training with HumanoidExo data instills the policy with strong generalization capabilities against significant real-world perturbations.

Generalization to New Environment. To evaluate generalization, we tested the policy in a new environment by moving the table next to a window, as shown in Figure 6(d) and (e). The robot was then tasked with performing the same ‘walk and place’ routine. While we observed a slight decrease in the success rate for the manipulation phase, the walking skill transferred perfectly to this unseen setting. This result demonstrates the strong generalization ability of the locomotion capabilities learned through our approach, even when the visual context changes significantly.

D. HumanoidExo for Whole-Body Dexterous Manipulation

We conducted further experiments to test the whole-body dexterous manipulation capabilities of our humanoid robot. Specifically, we chose a ‘Place Laundry’ task, where the robot is required to squat, grasp clothes from a basket, and place them into a washing machine on its right. The robot repeats this process until the basket is empty, then stands up to signal task completion. This task presents several challenges: the clothes are deformable objects that are difficult for dexterous hands to grasp and place entirely inside the machine, which requires the model to exhibit recovery behaviors. Furthermore, the model must use robust visual observation while maintaining whole-body balance to avoid falling during the upper-body task execution.

Following the methodology of the previous section, we trained the HE-VLA model in three different configurations: one with 200 teleoperated demonstrations, a second with only 5 teleoperated demonstrations and 195 HumanoidExo data points, and a third with only 5 teleoperated demonstrations. The results, presented in Table I, show that the model trained on 200 teleoperated demonstrations performs similarly to our model, which used 195 HumanoidExo data points. These results demonstrate that even for complex whole-body dexterous manipulation, our method can be a strong replacement for teleoperated data and can even be more effective for task generalization.

V. CONCLUSIONS

In this work, we addressed the critical data bottleneck that hinders the development of capable, general-purpose humanoid robots. While existing methods like simulation,

TABLE I: **Experimental Results for PlaceLaundry Task.** The results demonstrate that HumanoidExo data is a viable substitute for teleoperated data, achieving comparable performance across both in-domain and out-of-domain scenarios.

Teleoperated Data	HumanoidExo Data	Seen Cloth Success Rate	Unseen Cloth Success Rate
200	0	80%	80%
5	195	80%	75%
5	0	5%	5%

human videos, and direct teleoperation have advanced the field, they suffer from significant limitations in scalability, cost, and embodiment mismatch. We introduced HumanoidExo, a lightweight, wearable exoskeleton system designed to provide a practical and effective solution for scalable, whole-body data collection. Our experiments confirm that this approach is highly effective. We’ve shown that data from HumanoidExo enables policies to generalize to new environments, achieve remarkable data efficiency by learning complex skills from as few as five real-robot demonstrations, and even acquire entirely new skills like walking without any prior robot data. These results validate our system as a powerful paradigm for generating large-scale, high-quality humanoid datasets.

REFERENCES

- [1] NVIDIA, “Gr00t n1: An open foundation model for generalist humanoid robots,” *arXiv preprint arXiv:2503.14734*, 2025.
- [2] S. Kareer, D. Patel, R. Punamiya, P. Mathur, S. Cheng, C. Wang, J. Hoffman, and D. Xu, “Egomimic: Scaling imitation learning via egocentric video,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 13 226–13 233.
- [3] R. Zhong, C. Cheng, J. Xu, Y. Wei, C. Guo, D. Zhang, W. Dai, and H. Lu, “Nuexo: A wearable exoskeleton covering all upper limb rom for outdoor data collection and teleoperation of humanoid robots,” *arXiv preprint arXiv:2503.10554*, 2025.
- [4] R. Qiu, S. Yang, X. Cheng, C. Chawla, J. Li, T. He, G. Yan, D. J. Yoon, R. Hoque, L. Paulsen, G. Yang, J. Zhang, S. Yi, G. Shi, and X. Wang, “Humanoid policy - human policy,” *arXiv preprint arXiv:2503.13441*, 2025.
- [5] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu, “Dexcap: Scalable and portable mocap data collection system for dexterous manipulation,” *arXiv preprint arXiv:2403.07788*, 2024.
- [6] H. Fang, H. Fang, Y. Wang, J. Ren, J. Chen, R. Zhang, W. Wang, and C. Lu, “Airexo: Low-cost exoskeletons for learning whole-arm manipulation in the wild,” *arXiv preprint arXiv:2309.14975*, 2024.
- [7] H. Fang, C. Wang, Y. Wang, J. Chen, S. Xia, J. Lv, Z. He, X. Yi, Y. Guo, X. Zhan, L. Yang, W. Wang, C. Lu, and H. Fang, “Airexo-2: Scaling up generalizable robotic imitation learning with low-cost exoskeletons,” *arXiv preprint arXiv:2503.03081*, 2025.
- [8] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” *arXiv preprint arXiv:2402.10329*, 2024.
- [9] M. Xu, H. Zhang, Y. Hou, Z. Xu, L. Fan, M. Veloso, and S. Song, “Dexumi: Using human hand as the universal manipulation interface for dexterous manipulation,” *arXiv preprint arXiv:2505.21864*, 2025.
- [10] Y. Zhang, Y. Yuan, P. Gurunath, T. He, S. Omidshafiei, A. akbar Aghamohammadi, M. Vazquez-Chanlatte, L. Pedersen, and G. Shi, “Falcon: Learning force-adaptive humanoid loco-manipulation,” *arXiv preprint arXiv:2505.06776*, 2025.
- [11] Q. Zhang, P. Cui, D. Yan, J. Sun, Y. Duan, G. Han, W. Zhao, W. Zhang, Y. Guo, A. Zhang, and R. Xu, “Whole-body humanoid robot locomotion with human reference,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 11 225–11 231.
- [12] 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.
- [13] Y. Li, Y. Lin, J. Cui, T. Liu, W. Liang, Y. Zhu, and S. Huang, “Clone: Closed-loop whole-body humanoid teleoperation for long-horizon tasks,” *arXiv preprint arXiv:2506.08931*, 2025.
- [14] 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.
- [15] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang, “Open-television: Teleoperation with immersive active visual feedback,” *arXiv preprint arXiv:2407.01512*, 2024.
- [16] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, “Humanplus: Humanoid shadowing and imitation from humans,” *arXiv preprint arXiv:2406.10454*, 2024.
- [17] Y. Jiang, R. Zhang, J. Wong, C. Wang, Y. Ze, H. Yin, C. Gokmen, S. Song, J. Wu, and L. Fei-Fei, “Behavior robot suite: Streamlining real-world whole-body manipulation for everyday household activities,” *arXiv preprint arXiv:2503.05652*, 2025.
- [18] 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.
- [19] Z. Luo, Y. Yuan, T. Wang, C. Li, S. Chen, F. Castañeda, Z.-A. Cao, J. Li, D. Minor, Q. Ben, X. Da, R. Ding, C. Hogg, L. Song, E. Lim, E. Jeong, T. He, H. Xue, W. Xiao, Z. Wang, S. Yuen, J. Kautz, Y. Chang, U. Iqbal, L. J. Fan, and Y. Zhu, “Sonic: Supersizing motion tracking for natural humanoid whole-body control,” *arXiv preprint arXiv:2511.07820*, 2025.
- [20] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *arXiv preprint arXiv:2304.13705*, 2023.
- [21] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *RSS*, 2023.
- [22] 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.
- [23] H. Zhang, S. Hu, Z. Yuan, and H. Xu, “Doglove: Dexterous manipulation with a low-cost open-source haptic force feedback glove,” *arXiv preprint arXiv:2502.07730*, 2025.
- [24] J. Wen, Y. Zhu, J. Li, Z. Tang, C. Shen, and F. Feng, “Dexvla: Vision-language model with plug-in diffusion expert for general robot control,” *arXiv preprint arXiv:2502.05855*, 2025.
- [25] K. Black, N. Brown, D. Driess *et al.*, “ π_0 : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [26] Physical Intelligence, “ $\pi_{0.5}$: a vision-language-action model with open-world generalization,” *arXiv preprint arXiv:2504.16054*, 2025.
- [27] J. Wen, Y. Zhu, J. Li, M. Zhu, K. Wu, Z. Xu, R. Cheng, C. Shen, Y. Peng, F. Feng *et al.*, “Tinyvla: Towards fast, data-efficient vision-language-action models for robotic manipulation,” *arXiv preprint arXiv:2409.12514*, 2024.
- [28] P. Ding, J. Ma, X. Tong, B. Zou, X. Luo, Y. Fan, T. Wang, H. Lu, P. Mo, J. Liu *et al.*, “Humanoid-vla: Towards universal humanoid control with visual integration,” *arXiv preprint arXiv:2502.14795*, 2025.
- [29] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, “3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations,” in *ICRA 2024 Workshop on 3D Visual Representations for Robot Manipulation*, 2024.
- [30] Z. Jiang, Y. Xie, K. Lin, Z. Xu, W. Wan, A. Mandlekar, L. J. Fan, and Y. Zhu, “Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 16 923–16 930.
- [31] Y. Zimmermann, M. Sommerhalder, P. Wolf, R. Riener, and M. Hutter, “Anyexo 2.0: A fully actuated upper-limb exoskeleton for manipulation and joint-oriented training in all stages of rehabilitation,” *IEEE Transactions on Robotics*, pp. 2131–2150, 2023.
- [32] W. Xu and F. Zhang, “Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.