

Learning Humanoid End-Effector Control for Open-Vocabulary Visual Loco-Manipulation

Runpei Dong[†] Ziyuan Li[†] Xialin He Saurabh Gupta

University of Illinois Urbana-Champaign
hero-humanoid.github.io

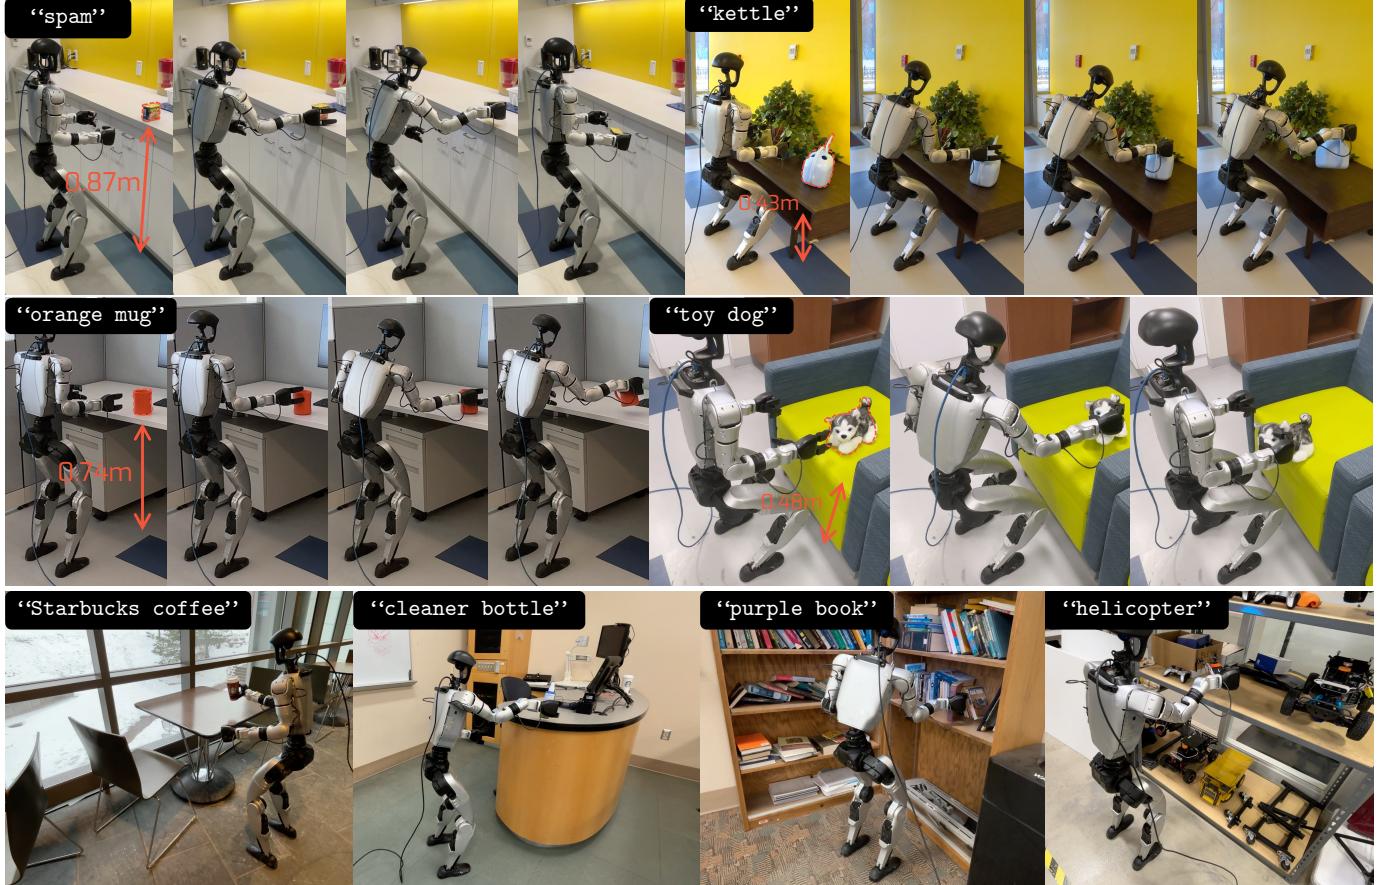


Fig. 1: We build capability for a humanoid to *autonomously* loco-manipulate *novel* objects in *novel* scenes using onboard sensors. We achieve this through a modular system powered by large vision models for visual generalization and an accurate end-effector tracking policy. Our system achieves an 83.8% average success rate at reaching and picking up novel objects in novel scenes in the real world in challenging scenarios that involve whole-body control via bending, squatting, and twisting.

Abstract—Visual loco-manipulation of arbitrary objects in the wild with humanoid robots requires accurate end-effector (EE) control and a generalizable understanding of the scene via visual inputs (e.g., RGB-D images). Existing approaches are based on real-world imitation learning and exhibit limited generalization due to the difficulty in collecting large-scale training datasets. This paper presents a new paradigm, HERO, for object loco-manipulation with humanoid robots that combines the strong generalization and open-vocabulary understanding of large vision models with strong control performance from simulated training. We achieve this by designing an accurate residual-aware EE tracking policy. This EE tracking policy combines classical robotics with machine learning. It uses a) inverse kinematics to convert

residual end-effector targets into reference trajectories, b) a learned neural forward model for accurate forward kinematics, c) goal adjustment, and d) replanning. Together, these innovations help us cut down the end-effector tracking error by 3.2 \times . We use this accurate end-effector tracker to build a modular system for loco-manipulation, where we use open-vocabulary large vision models for strong visual generalization. Our system is able to operate in diverse real-world environments, from offices to coffee shops, where the robot is able to reliably manipulate various everyday objects (e.g., mugs, apples, toys) on surfaces ranging from 43cm to 92cm in height. Systematic modular and end-to-end tests in simulation and the real world demonstrate the effectiveness of our proposed design. We believe the advances in this paper can open up new ways of training humanoid robots to interact with daily objects.

[†] Equal contributions.

I. INTRODUCTION

Think about reaching to pick up the objects placed on the various table tops in Fig. 1. As humans we can reliably and robustly use our whole bodies to execute such pick ups. We can use our back to reach over and across the table, or rotate our torso to pick up objects kept on the side, and use our legs to squat down to pick up objects kept on coffee tables, all while maintaining balance on our two legs. We can pick up seen objects on seen tables, but equally effortlessly also pick up novel objects on novel tables in novel scenes. Once we have glanced at the object and scene, we can even do this with our eyes closed if we want. In this paper, we develop an autonomous system that equips a humanoid robot with this fundamental capability: *reach over to pick up novel objects in novel everyday environments around us*.

Humanoids are doing backflips [24, 27, 29, 44, 86], so why would we be writing about such a mundane and seemingly unimpressive task? There are two key differences that make our problem of manipulating novel objects harder: the need for processing the high-dimensional RGB-D image observations to infer object locations and scene collision geometry, and the precision necessary for object manipulation. Most existing humanoid control works aren't conditioned on perceptual input from RGB-D camera: either they don't need any scene information or scene information is provided as input using environment sensors (e.g. MOCAP). Second, while back flips are impressive, it doesn't really matter *where* you land, as long as you land safely. In contrast, if a robot needs to pick up an object, it needs to get its hand *where the object actually is*. Object manipulation requires precise goal-directed behavior. Operation in novel environments, sensing using on-board RGB-D sensors, precise EE control, and the complexity of maintaining balance while moving around, make this problem challenging.

State-of-the-art for training humanoids for such tasks is end-to-end imitation learning in the real world [2, 4, 18, 73]. However, difficulty in collecting large datasets for learning limits generalization capabilities of learned policies. This causes them to fall short of the goal of manipulating novel objects in novel environments, which requires broad generalization. In this paper, we pursue an alternate approach. We take inspiration from strong results with modular systems for table-top object manipulation problems [12, 47, 62]. Vision foundation models translate high-level instructions into actionable plans (e.g., grasp ‘‘the red coke can’’), grasp synthesis models (e.g., AnyGrasp [13]) convert this into grasps, and a low-level control module conveys the robot EE to the grasp location. Being able to use large pre-trained models enables broad generalization and even open-world reasoning. In many ways, this is the more direct, obvious, and performant way to build such a system. So, why isn't such a modular method the go-to method for building a humanoid object manipulation system?

While it is easy to get a Franka Emika robot to where you want, it turns out it is extremely difficult for current methods to accurately control a humanoid hand. Leading

policies for this task achieve approximately 8–13cm error for end-effector tracking, an error rate that is simply too large for object manipulation. Our key technical contribution is to develop an accurate end-effector tracking policy that enables object manipulation applications. This unlocks the possibility of developing modular humanoid systems for object manipulation that generalize without large-scale real-world imitation demos.

So what are the ingredients of building a highly accurate end-effector tracker? Our accurate end-effector control algorithm is based on multiple innovations. First, rather than just directly trying to get the end-effector to the target location, we use a motion planner to generate an upper-body reference motion that gets the end-effector to the desired target. Second, the policy receives as input not just the current and target joint angles (output using the motion planner), but also the current and target end-effector position. Third, it is important to obtain a high-quality estimate of the current end-effector position, as we found that analytical forward kinematics and odometry on a low-cost humanoid robot like Unitree G1 are not accurate. We mitigate this issue by training neural forward models. We train two models. A neural forward kinematics model maps that provides an accurate end-effector pose relative to the base. And a neural odometry model that provides an accurate base pose relative to the stationary feet. Even after improvements, the tracker would make systematic errors. We further mitigate these errors by adjusting the desired target passed into the tracker based on the current EE pose tracking error, to encourage the tracker to get where we want it to be. Together, these innovations improve the state-of-the-art for end-effector. Our proposed tracker HERO achieves a tracking error of 2.5cm compared to 8–13 cm for previous state-of-the-art [39, 94]. In real-world testing in a MOCAP room, our full system achieves an average end-effector tracking error of 2.44cm.

Using this performant end-effector tracking policy, we develop a modular system for picking up open-vocabulary novel objects in novel everyday environments. This modular system leverages an open-vocabulary perception module to detect and segment the target object using large pre-trained vision models (Grounding DINO 1.5 [48] and SAM-3 [7]). We next use the AnyGrasp model [13] to produce parallel jaw grasps on the candidate object. We retarget them to the Dex3 hand on the Unitree robot. Finally, we use our tracker as a low-level controller to achieve the grasp pose. In real-world testing for grasping open-vocabulary object queries in novel environments, our system achieves a average success rate of 83.8% across more than 25 daily objects, 10 daily and cluttered scenes kept on diverse table heights.

II. RELATED WORKS

A. Legged Loco-Manipulation

1) *Loco-manipulation via motion tracking*: Motion tracking is vital for loco-manipulation, where teleoperation has emerged as one of the primary paradigms for facilitating imitation learning [2, 73, 92]. Inspired by DeepMimic-style tracking [57], He et al. [22, 23], Fu et al. [18], and Cheng et al. [10] pioneer this direction via training motion tracking policy with reinforcement

learning (RL) and sim2real transferring, where remarkable loco-manipulation results have been obtained on an H1 humanoid platform [79]. Abundant works have been proposed to improve whole-body tracking accuracy [26, 32, 39, 54, 60], agility [15, 24, 30], generalization [9, 42, 51], robustness [41, 94, 96], object interaction [84, 97], and reachability [95]. With generated reference motion, several works have demonstrated success in humanoid loco-manipulation [40, 55]. In the meanwhile, teleoperation systems have been rapidly developed, with which imitation learning has demonstrated remarkable loco-manipulation results [5, 45, 50, 59, 67, 83, 91, 93].

2) *Visual loco-manipulation*: One line of research focuses on imitation learning, where the visual loco-manipulation data is collected via human teleoperation, and great progress has been made [2, 18, 21, 22, 73]. Another line of research relies on reference state-based policy learning. For example, Liu et al. [46] utilize depth inputs that learn to predict high-level manipulation commands that control low-level policies. Yin et al. [87] propose to generate motions from visual depth inputs, where the robot interacts with objects via tracking generated motions. Our work lies in this category, where our system takes the visual input and outputs the EE pose target for loco-manipulation. More recently, He et al. [25] explores an end-to-end RL-based loco-manipulation visual policy learning and demonstrated great success on cylinder-style objects (*e.g.*, coke can). However, such methods cannot generalize to open-vocabulary queries as training assets are rather limited.

B. System Identification

Real-world robots have intrinsic errors due to hardware inaccuracy. One assumption is that the robot's joint has elasticity [16, 71, 72], which makes the joint positions encoded by motors unreliable. Researcher have proposed to conduct a system identification to mitigate this sim2real gap [3, 34]. To do this, two mainstream methods are used: online and offline system identification. Online methods aim at learning to compensate for the hardware and dynamic error during inference, where an adaptation model is trained to mitigate possible errors [15, 17, 36, 38, 43, 56, 58, 81, 88–90]. The other line of research, instead, utilizes the offline collected data for system identification. For example, ASAP [24] utilizes the offline collected data for a dynamic adaptation model training, which enables extremely agile motion tracking. Focusing on the same direction of optimizing motor dynamics, several works have been proposed and achieved promising system identification results [37, 70]. In this work, we propose two residual models that correct the forward kinematics and robot base odometry trained on collected MOCAP data, which can be categorized as an offline system identification method.

III. HERO: HUMANOID END-EFFECTOR CONTROL

Given a desired end-effector pose in the robot frame, the end-effector control policy outputs motor commands for all 29-DOFs of a G1 humanoid robot [80] to convey the arm to reach the desired end-effector pose. Note that even the robot doesn't need to take any steps, reaching far away targets

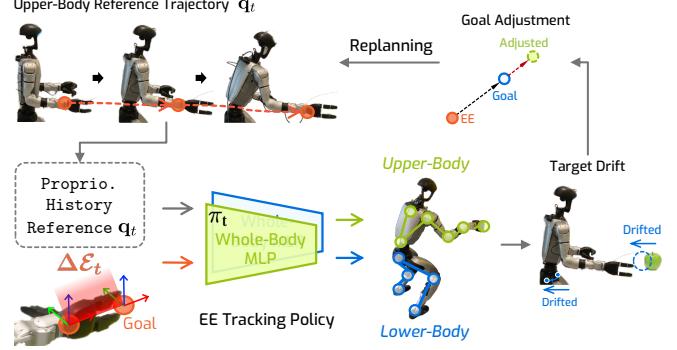


Fig. 2: **HERO is an accurate end-effector control framework.** Given an EE goal pose, HERO first uses IK to convert it into an upper-body goal. It then uses motion planning to generate an upper-body reference trajectory that is tracked via a learned tracking policy π_t (Sec. III). In addition to reference joints, π_t also takes accurate estimates of the residual EE error (obtained via a learned neural forward model (Sec. III-B, Sec. III-B2). HERO also employs periodic replanning (Sec. III-C) to adapt to drifts and goal adjustment (Sec. III-D) to mitigate systematic tracking errors. Accurate tracking enables building modular object manipulation systems (Sec. IV, Fig. 3).

requires whole-body coordination and balancing (and thus the control of all DOFs): bending at the waist, twisting the torso, or bending down the legs. Our innovation lies in the design of the policy architecture and input/output representations, where we combine classical robotics components with learned ones in novel ways. As our experiments will illustrate, a monolithic learning solution, that directly attempts to learn this mapping, fails. The overall design of our tracker is shown in Fig. 2.

Given the end-effector target location \mathcal{T}^{EE} in the robot frame, we first use inverse kinematics to transform it into robot base height $h \in \mathbb{R}$, and upper body joint angles $\mathbf{q}^* \in \mathbb{R}^{17}$ (3 DoF waist plus $2 \times 7 = 14$ DoF left and right arm joints). We then use a collision-free motion planner [68] that takes the egocentric depth image and robot configurations to compute a joint trajectory $\{\mathbf{q}_t\}_{t=1}^T$ that conveys the robot arm from the starting configuration \mathbf{q}_0 to the target configuration \mathbf{q}^* .

We then use our proposed whole-body end-effector tracking policy π_t to execute this whole-body reaching motion. π_t is a learned neural network policy (described in Sec. III-A) that outputs joint angles commands for position control, where torque is obtained via a PD controller operating at 50Hz. π_t itself relies on a learned forward model η (described in Sec. III-B) and a learned base odometry model ξ (described in Sec. III-B2) to map robot configurations to end-effector poses (rather than forward kinematics, which is significantly less accurate). There is still a systematic offset in where the end-effector ends up after π_t is executed. We account for this by systematically shifting the desired end-effector pose goal in the opposite direction to further improve tracking accuracy (Sec. III-D). Finally, because π_t requires reference

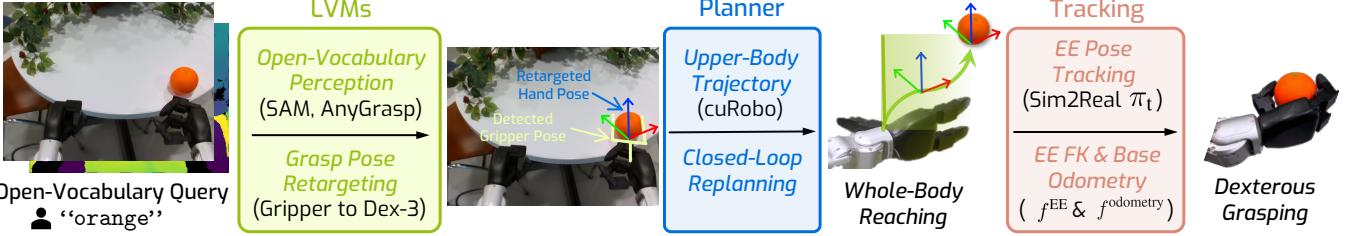


Fig. 3: **Overall architecture for our proposed modular system for open-vocabulary object grasping.** Given a free-form natural language text query indicating which object needs to be picked, we use open-vocabulary large vision models (LVMs like Grounding DINO [48] and SAM [7]) to segment out the object of interest and predict parallel jaw grasps (using the AnyGrasp model [13]). We retarget the predicted grasp to a Dex-3 hand. We use our proposed whole-body end-effector tracker to convey the robot arm to the predicted grasp before picking up the object. By decomposing *action planning* (*i.e.* identifying which object to pick and using what grasp) from *action execution* (*i.e.* actual control of the robot), we inherit the strong visual generalization from pre-trained models as well as strong control capabilities for simulated training of the tracking policy.

poses $\{\mathbf{q}_t\}_{t=0}^T$ as input, we replan to recompute the reference poses every k steps during execution (Sec. III-C).

A. Whole-body End-Effector Tracking Policy, π_t

To track the target end-effector pose $\mathcal{T}^{\text{EE}} \in \text{SE}(3)$ defined in the robot frame, our whole-body EE tracking policy π_t first obtains reference trajectory $\{\mathbf{q}_t\}_{t=1}^T$ and the corresponding reference EE pose trajectory $\{ee_t\}_{t=1}^T$ from a motion planner. Given the trajectory, the current proprioceptive state \mathbf{s}_t , and other commands, the tracking policy predicts the 29-DoF joint angles that are passed to per-joint PD controllers.

1) *Residual-aware end-effector tracking*: π_t output actions \mathbf{a}_t at time t as follows:

$$\mathbf{a}_t = \pi_t(\mathbf{s}_t, h_t, \mathbf{q}_t, \Delta\mathcal{E}_t, \mathbf{v}_t, \mathbf{s}_{t-5:t-1}, \mathbf{a}_{t-5:t-1}),$$

where \mathbf{s}_t is the current proprioception, h_t is the reference base height, \mathbf{q}_t are the reference upper-body joint angles, \mathbf{v}_t are the linear and angular velocity locomotion commands, and $\mathbf{s}_{t-5:t-1}, \mathbf{a}_{t-5:t-1}$ are five time steps of proprioception and action history. The proprioception include the robot's joint angles, joint velocities, angular velocity, projected gravity, and roll and pitch encoded from the IMU. We don't use the IMU yaw as it is inaccurate [22]. $\Delta\mathcal{E}_t$ represents the residual pose error between the current and target end-effector pose in the robot frame, *i.e.*,

$$\Delta\mathcal{E}_t = f^{\text{EE}}(\mathbf{x}_t) \ominus ee_t, \quad (1)$$

where $f^{\text{EE}}(\mathbf{x}_t)$ maps the arm states $\mathbf{x}_t \in \mathbb{R}^{17}$ to the end-effector pose $T_t^{\text{EE}} \in \text{SE}(3)$, and \ominus is the inverse pose composition operator.¹

2) *Policy architecture*: π_t is implemented as two three-hidden-layer MLPs that decouple whole-body control into upper- and lower-body actions. Both networks take the same observation as input and together predict 29-DoF whole-body actions, obtained by combining the upper-body components from one MLP with the lower-body components from the other.

¹We use \oplus to denote the pose composition operator: $\mathbf{T}_1 \oplus \mathbf{T}_2 = \mathbf{T}_2 \cdot \mathbf{T}_1$ and \ominus is the inverse pose composition operator: $\mathbf{T}_1 \ominus \mathbf{T}_2 = \mathbf{T}_2^{-1} \cdot \mathbf{T}_1$.

3) *Training*: π_t is trained in simulation via reinforcement learning with PPO [65]. We utilize both the AMASS dataset [52] ($\sim 8K$ motion sequences) and a curated set of typical end-effector targets ($\sim 8K$) that are encountered in everyday object reaching. Specifically, we randomly sample reaching targets with $[x, y, z]$ coordinates defined in robot frame (z axis is computed from the ground better understanding) ranging from $[0.1\text{m}, -0.5\text{m}, 0.65\text{m}]$ to $[0.5\text{m}, 0.5\text{m}, 1.15\text{m}]$, with the yaw orientation randomly sampled from -60° to 60° . The reference trajectories are obtained via a motion planner that outputs the upper-body and end-effector tracking goals.

B. Learned Residual Neural Forward Models

1) *Residual neural FK*: Our residual neural forward kinematics function, η , learns a correction to the output of the analytical forward kinematic function, FK, to output accurate end-effector poses. Specifically, given the current proprioceptive state of one operating arm and waist $\mathbf{x}_t \in \mathbb{R}^{10}$ at timestep t and output from analytical FK, $\text{FK}(\mathbf{x}_t)$, the final corrected end-effector pose $f^{\text{EE}}(\mathbf{x}_t)$ is obtained via pose composition:

$$f^{\text{EE}}(\mathbf{x}_t) = \text{FK}(\mathbf{x}_t) \oplus \eta(\mathbf{x}_t, \text{FK}(\mathbf{x}_t)).$$

Note that the analytical forward kinematics function $\text{FK}(\mathbf{x}_t)$ uses the robot geometry and coordinate transformations to compute the 6-DoF end-effector pose in the robot base frame. For precise robots, FK is itself quite accurate, however as our experiments will show, FK is inaccurate for our humanoid, necessitating the need for learning a correction.

2) *Residual Neural Leg Odometry*: Different from fixed-base object manipulation tasks, where the robot's base is fixed and stable, humanoid robots' base *floats* in the air for whole-body balancing during reaching. This movement makes the original reaching target inaccurate, as the reaching target defined in the robot frame is no longer the same place where the object lies. One might consider using the egocentric visual information for replanning or motion adjustment. However, as shown in Fig. 4, the egocentric view of the robot is too narrow for the robot to see the object when the robot's arm and waist movements are

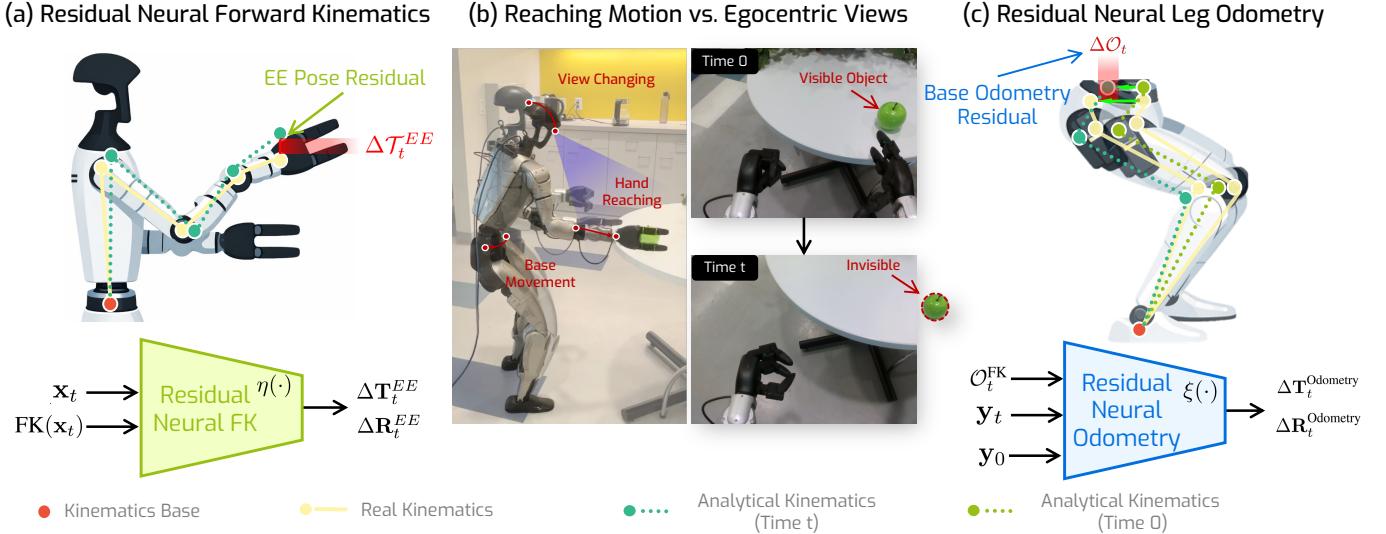


Fig. 4: **Learned neural forward kinematics model and odometry model.** (a) To correct for inaccurate analytical forward kinematics (FK) that maps joint angles and robot geometry to end-effector poses, we design a *residual neural forward kinematics model* (Sec. III-B), η , that predicts corrections $\Delta\mathbf{T}_t^{EE}$ & $\Delta\mathbf{R}_t^{EE}$ to the analytical forward kinematics output. (b) As the humanoid’s base moves around while reaching the object, the object can go out of view due to large whole-body motions, making closed-loop adjustment from vision infeasible. Thus, it is necessary to accurately estimate base motion. (c) Our residual neural odometry model accurately estimates base odometry from lower body joint states and by assuming that feet remain fixed (Sec. III-B2).

large. As a result, knowing the robot’s odometry, which helps in adjusting the reaching goal, is critical. We assume the feet to be static on the ground and use the lower body joint angles to predict the base pose. By assuming the robot ankle joint as the root joint and the robot base as the end-effector, we can compute the base pose using forward kinematics.

However, similar to the error in analytical FK for EE, analytical FK to compute the base pose is also inaccurate (see analysis in Sec. V). Similar to our solution for EE, we adopt a residual model to mitigate these inaccuracies. The only difference here is that we reason about base pose transformation relative to time step 0, rather than the absolute base pose.

Concretely, let $\mathbf{y}_t \in \mathbb{R}^6$ be the 6DOF state of the left (or right) leg motors. We can get analytical base odometry, *i.e.* base pose relative to base pose at time step 0, $\mathcal{O}^{FK}(\mathbf{y}_t, \mathbf{y}_0) \in SE(3)$, using analytical forward kinematics and $SE(3)$ difference:

$$\mathcal{O}^{FK}(\mathbf{y}_t, \mathbf{y}_0) = FK(\mathbf{y}_0) \ominus FK(\mathbf{y}_t).$$

Our residual neural leg odometry model ξ learns the residual:

$$f^{\text{odometry}}(\mathbf{y}_t, \mathbf{y}_0) = \mathcal{O}^{FK}(\mathbf{y}_t, \mathbf{y}_0) \oplus \xi(\mathbf{y}_t, \mathbf{y}_0, \mathcal{O}^{FK}(\mathbf{y}_t, \mathbf{y}_0)).$$

3) *Policy Architecture*: η is realized using a 3-layer multi-layer perceptron and outputs the residual transformation comprising a residual rotation and a residual translation. One head predicts the residual translation in \mathbb{R}^3 and another head predicts the first two columns of the residual rotation as a vector in \mathbb{R}^6 following [98]. Similar to the η , ξ is also a 3-layer MLP that predicts a translation and rotation for the residual transformation between ground truth and analytical FK result $\mathcal{O}_t^{FK}(\mathbf{y}_t, \mathbf{y}_0)$.

4) *Training data and Loss Function*: η is trained by collecting data in an MOCAP room using the Optitrack system [1]. Specifically, we deploy a tracking policy (it doesn’t have to be an accurate tracking policy) to move the robot end-effector across the workspace. We collect the robot state (joint angles) using the motor encoders and the true end-effector pose using the MOCAP system. We employ the Kabsch-Umeyama algorithm [33, 78] to transform the captured MOCAP marker coordinates to the robot’s end-effector and base pose, which achieves an accurate estimation with $< 1.5\text{mm}$ RMSE error (more details can be found in the supplementary). In this fashion, we collect a total of 3 hours of data $\mathcal{D}^{EE} = \{\mathcal{T}_t^{EE}, \mathbf{x}_t\}_{t=1}^N$ that spans the workspace we want the robot to operate in, we use the first 2 hours of data for training and the last 1 hour of data for offline validation. ξ is also trained on data captured in a MOCAP room, following the same protocol: 2 hours of training data and 1 hour of validation data. As ξ needs temporal pairs, here we create training data by sampling pairs of time steps from temporal trajectories.

Given data \mathcal{D}^{EE} , and the residual transformation for η defined as $\Delta\mathbf{T}_t^{EE} = \mathcal{T}_t^{EE} \ominus FK(\mathbf{x}_t)$, we use a decoupled representation for $\Delta\mathbf{T}_t^{EE}$, where the translation is $\Delta\mathbf{T}_t^{EE} \in \mathbb{R}^3$, and the orientation $\Delta\mathbf{R}_t^{EE} \in \mathbb{R}^6$ is the first two columns of the orientation transformation matrix [98]. The residual neural FK model, η is trained using the MSE loss: $\mathcal{L}_{\text{MSE}}(\eta(\mathbf{x}_t, FK(\mathbf{x}_t)), \Delta\mathbf{T}_t^{EE})$. Similarly, ξ utilize the same decoupled continuous representation and MSE loss $\mathcal{L}_{\text{MSE}}(\xi(\mathbf{y}_m, \mathbf{y}_n, \mathcal{O}_{n-m}^{FK}), \Delta\mathcal{O}_{n-m})$, where $0 \leq m \leq n$ are two timestamps randomly sampled from collected temporal sequence.

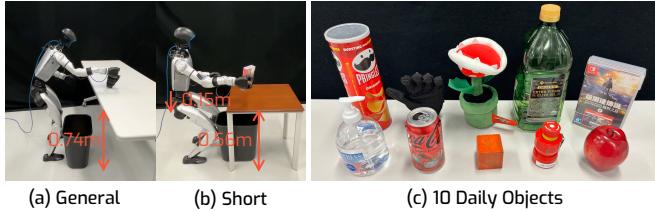


Fig. 5: **Novel test environments and novel test objects** used for end-to-end testing of our proposed humanoid open-vocabulary object grasping system. (a-b) Standard table (0.74m) and short table (0.56m) setups; note that for short table the robot would first squat down by 15cm the motion planner suggested. (c) 10 daily objects with different shapes, physical properties, appearances, *etc* (details can be found in Sec. B.7).

C. Replanning

Over the course of execution, the robot may drift sufficiently far from the reference joint trajectory that the tracking policy π_t is being asked to track. This could lead to poor tracking performance because of the input being out-of-distribution or the tracking target simply being too far away. We mitigate this by replanning the reference motion every $k = 300$ time steps (6 seconds). We replan using the same motion planner [68] as before (Sec. III). A replanning at time t updates the remaining $\{\mathbf{q}_t\}_t^T$ and $\{\mathbf{ee}_t\}_t^T$ and takes ≈ 20 ms.

D. Goal Adjustment

Finally, due to the sim2real mismatch there may still be a systematic gap between where the policy is commanded to go to and where the policy ends up at. To eliminate these systematic errors, we adjust the target in the direction opposite to the current error to encourage the policy to get to where we want it to get to. We only modify the $\Delta\mathcal{E}_t$ that is being input into the policy. We simply scale up the current error by a factor of $\alpha = 1.6$. We only scale up the translation error vector as we found no benefit from also scaling up the rotation error. The goal adjustment starts when the EE tracking error $\Delta\mathcal{E}_t \leq 0.15$ m, and stops when the $\Delta\mathcal{E}_t \leq 0.02$ m, which helps stabilize the robot when sufficiently close.

IV. A MODULAR SYSTEM FOR HUMANOID OBJECT GRASPING

Our overall task of picking up novel open-vocabulary objects in novel environments. Specifically, given a humanoid robot that is standing at a table, the goal is for it to pick up objects described by a free-form natural language query only using its onboard egocentric sensors.

We design a modular system for this task by building on top of our HERO tracker from Sec. III. Concretely, we first segment out the object corresponding to the query using Grounding DINO [48]. We next use an AnyGrasp model [13] to produce parallel jaw grasps on the candidate object. We filter the predicted grasps based on how parallel to the table they are. We retarget the selected grasp to the Dex3 hand on the Unitree

TABLE I: **Success rate** for the end-to-end open-vocabulary grasping task on novel objects on general and short tables in the real world. Our proposed system achieves a 90% success rate at this task across objects placed at two different heights.

Language Query	Success Rate	
	General Table Height: 0.74 m	Short Table Height: 0.56m
red coke can	3/3	3/3
emergency stop button	3/3	3/3
red piranha plant	3/3	3/3
orange cube	3/3	3/3
olive oil bottle	2/3	2/3
game cartridge	2/3	3/3
chip can	2/3	3/3
hand soap bottle	3/3	3/3
robot hand	3/3	2/3
red apple	3/3	2/3
Total	27/30	27/30

robot. Finally, we use our proposed HERO tracker to convey the robot end-effector to the predicted location.

Specifically, to map parallel gripper grasps to the Dex3 hand, we rotate the AnyGrasp output pose by 45° around the z-axis, enabling the thumb to form one jaw while the other two fingers form the opposing jaw. This configuration provides larger contact areas and better force closure than straight finger extension, improving grasp robustness and pose error tolerance. We also clip the end-effector orientation within 70° to ensure natural full-body motion. Excessive rotations cause twisted upper-body postures from inverse kinematics, degrading EE tracking precision.

V. EXPERIMENTS

We design experiments to test the performance of our overall system as well as the effectiveness of each individual module. Specifically, we seek answers to the following questions: a) what matters for building an accurate end-effector tracking system for humanoids? b) how does the end-effector tracking performance vary depending on the target reaching location? c) what is the extent of error in forward kinematics and can a learnt model successfully mitigate it? d) to what extent can a modular system enable a humanoid to manipulate novel objects in novel environments in the real world?

A. Experimental Setup

All our experiments use an *unmodified* Unitree G1 Humanoid robot with Dex-3 hands. We use the inputs from the head-mounted RGB-D camera (Intel D435i) and proprioception (dof angles, velocities), and base-mounted IMU for our system.

Experiments involving visual sensing are all conducted in novel environments on novel objects in the real world. Experiments that assess the quality of tracking or the accuracy in forward kinematics are done in a MOCAP room equipped with 13 Optitrack cameras (more details can be found in

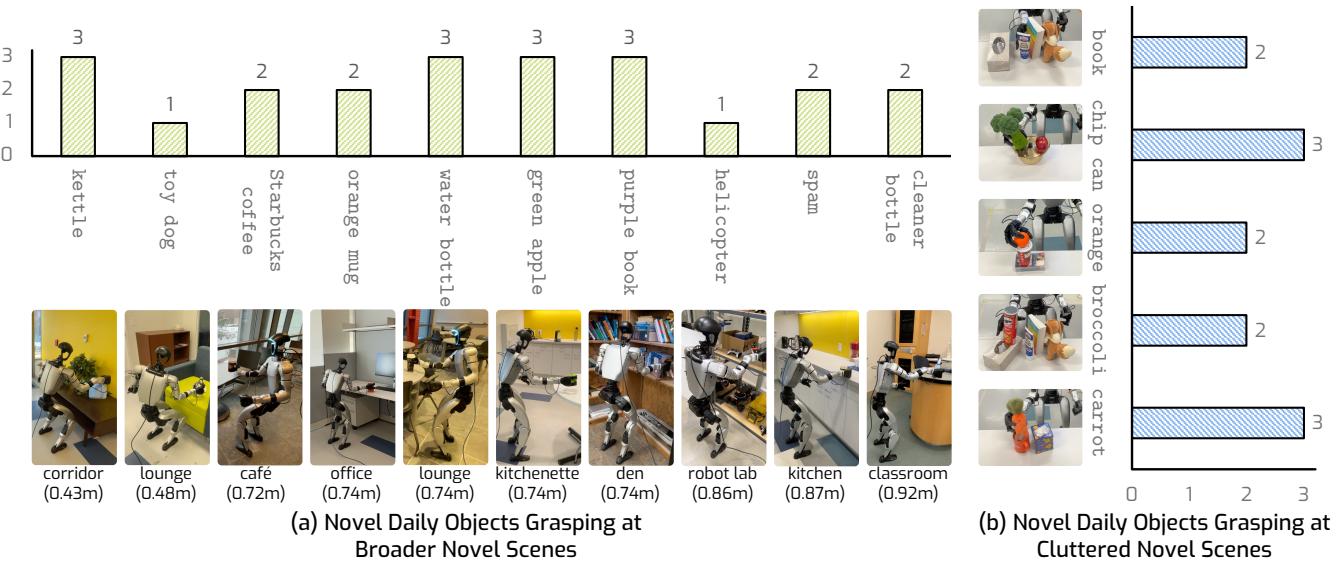


Fig. 6: **Success rate** for the end-to-end open-vocabulary grasping task on novel objects in (a) broader and (b) cluttered novel scenes in the real world. (a) We test HERO in 10 daily scenes on 10 new daily objects, such as office lounge and coffee shops. HERO achieves an overall 22/30 (73.3%) success rate, demonstrating strong scene generalization capability. Details of tested locations can be found in Sec. B.8. (b) We also test HERO in 5 random cluttered scenes with different layouts. HERO achieves an overall 12/15 (80%) success rate, demonstrating the generalization capability in using language as an accurate proxy for manipulating objects in cluttered scenes.

Sec. B.1). Finally, some design choices are validated in simulation using Issac Gym [53] and MoJuCo [75].

B. End-to-end System Testing

Our end-to-end system test starts with the humanoid robot standing 10-20cm in front of tables of varying heights (43cm – 92cm) and the goal for the robot is to pick up an object described via natural language. Fig. 5 shows the test objects, and Tab. I lists the open-vocabulary text queries we experimented with. The robot needs to operate entirely using its onboard sensors: the head-mounted RGB-D camera, proprioception, and a base-mounted IMU sensor. During trials, objects are placed randomly in the area reachable by the robot. A trial is deemed successful if the robot is able to pick up the correct object and lift it off the table for more than 2 seconds. We conduct 3 trials per object per table height.

10 Daily Objects. Tab. I reports the success rate at this end-to-end test. Across the different open world queries, two table heights, and objects, our system achieves a success rate of 90%, demonstrating the effectiveness of our modular system design. This establishes the effectiveness of a modular system for open-vocabulary object grasping with a humanoid robot, previously thought impossible for a humanoid robot. It also serves as the first demonstration of open-vocabulary grasping capability on humanoid robots operating under whole-body control.

10 Daily Scenes. Fig. 6(a) plots the success rate on broader novel daily scenes and objects for a generalization evaluation. We test our system across various scenes like robot lab and

classroom, and the results show that HERO achieves an overall 73.3% success rate, demonstrating remarkable scene and object generalization capability.

5 Cluttered Layouts. To further evaluate the open-vocabulary capability and robustness of our system, we test HERO on grasping objects from cluttered layouts, shown in Fig. 6(b). HERO achieves an 80% success rate, which demonstrates that our open-vocabulary system is language-sensitive and can reliably grasp objects in such challenging cluttered scenes without any human demonstration.

C. Forward Kinematics vs. Learned Forward Kinematics

We assess the accuracy of analytical forward kinematics on the G1 humanoid robot. We conduct this test in a MOCAP room. We affix tracking markers to the robot base and the robot end-effector, and use the relative transformation between the robot base and the robot end-effector as the ground truth. We execute a reaching policy and record the end-effector positions (from MOCAP, to use as ground truth) and the corresponding joint angles. We use the joint angles either with the *analytical forward kinematics* or with our learned forward kinematics model to obtain predicted end-effector poses.

End-Effector Pose. Tab. II reports the translation and orientation errors between the predicted end-effector poses and the ground truth poses measured using MOCAP. We average over 60 samples captured in diverse poses in the robot’s workspace. As we can see, analytical forward kinematics is off by as much as 1.76cm. At the same time, our learned model is much better and achieves an error of 0.27cm. The fact that a learned

TABLE II: Inaccuracies in analytical forward kinematics and how our learned model mitigates them. a) We report mean translation and rotation error in the estimate of end-effector position via different methods. Ground truth comes from MOCAP measurements. Analytical FK (*i.e.* using robot joint angles along with robot geometry) is inaccurate for our humanoid robot causing a 1.76cm translation error in the end-effector pose. Our proposed learned residual forward kinematics model is able to mitigate these errors with a residual design being more effective. b) Bottom part of the table reports metrics for base odometry and exhibits the same trends.

Method	Translation Error (cm)	Rotation Error (deg)
a) End-effector Pose		
Analytical FK	1.76	5.87
Learned FK (ours)	0.27	2.30
Learned FK, no residual (ours)	0.35	2.98
b) Base Odometry		
Analytical FK	1.10	0.49
Learned FK (ours)	0.33	0.36
Learned FK, no residual (ours)	0.37	0.42

model reduces the error by $6 \times$ indicates that analytical forward kinematics model errors are systematic. Tab. II also reports an ablation where we directly try to predict the end-effector pose *without* using the estimate from analytical FK and note that it does worse.

Base Odometry. Tab. II also reports base odometry results. Once again, we see that the residual neural model is better than both analytical FK and a non-residual neural model.

Execution Curves. We compare analytical and neural forward-model errors over time in Fig. 7. While analytical FK exhibits a persistent bias above 1.75 cm during reaching, our neural model η remains below 0.25 cm throughout. Similarly, as whole-body balancing causes analytical odometry drift to grow, our neural odometry model ξ reduces this drift by about $3 \times$ relative to the analytical baseline.

D. End-effector Tracking Accuracy Evaluation

We evaluate the performance of different tracking methods in simulation and in the real world under MOCAP. We evaluate different tracking methods on a fixed set of 180 reaching goals. The goal distribution is designed to reflect realistic grasping scenarios: we sample from three different table heights, with z -coordinates ranging from 5-15cm above each surface and x - y positions uniformly distributed across the table workspace. This yields 60 poses per table height (180 total), capturing the typical distribution of manipulation targets in everyday environments.

Metrics. We measure the translation and rotation errors in the end-effector position as the primary metrics. We report the mean and std of the errors. We use end-effector position as measured under MOCAP to compute these metrics. We also report the joint tracking error for the upper body joints. Here

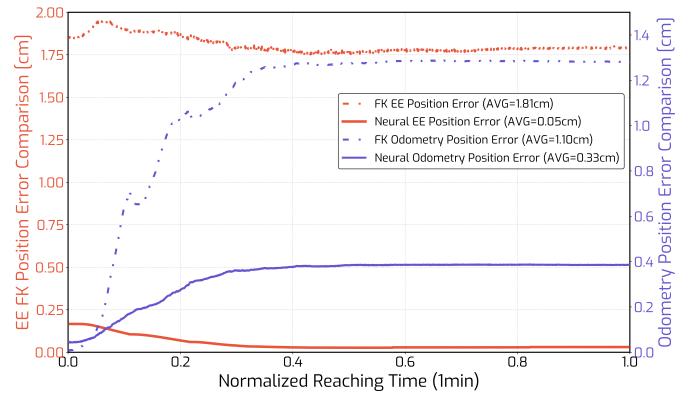


Fig. 7: Impact of neural forward models in the real world. We plot the neural and analytical forward kinematics and base odometry as a function of execution time steps. The plot shows 1 minute of execution at 50Hz.

TABLE III: End-effector tracking evaluation against prior methods in simulation (Sec. V-D1). We report end-effector translation and rotation errors and the upper body joint tracking error. Across all three table heights, our proposed HERO tracker achieves the lowest end-effector tracking error as compared to recent tracking methods FALCON [94] and AMO [39]. On average, our translation errors are $3.2 \times$ lower than the best baseline.

Method	Translation Error (cm)	Orientation Error (deg)	Joint Tracking Error (rad)
General Table (H: 0.74m)			
FALCON [94]	11.12 ± 3.58	16.07 ± 4.00	0.02 ± 0.00
AMO [39]	8.32 ± 3.41	14.38 ± 5.86	0.02 ± 0.00
HERO (ours)	2.21 ± 0.58	10.85 ± 3.95	0.16 ± 0.03
Tall Table (H: 0.88m)			
FALCON [94]	6.80 ± 2.51	13.30 ± 3.02	0.03 ± 0.01
AMO [39]	8.46 ± 4.17	14.04 ± 5.16	0.02 ± 0.00
HERO (ours)	3.30 ± 1.91	8.93 ± 3.06	0.20 ± 0.04
Low Table (H: 0.5m)			
FALCON [94]	22.80 ± 7.15	28.63 ± 8.87	0.02 ± 0.00
AMO [39]	8.10 ± 3.87	13.12 ± 6.70	0.02 ± 0.00
HERO (ours)	1.92 ± 0.96	13.92 ± 4.91	0.16 ± 0.03

we use measurements from the motor encoders.

Comparisons. We compare against AMO [39] and FALCON [94], two recent tracking methods that report impressive tracking performance. We retrain both policies on the same end-effector targets that our policy is trained on. This mitigates any bias due to a mismatch in training distributions between the different policies. In addition, we also compare to ablated versions of our method. Specifically, we ablate the effect of forward model quality (Sec. III-B and Sec. III-B2), re-planning (Sec. III-C), and goal adjustment (Sec. III-D).

Results. We discuss the results and our key takeaways.

1) *Comparisons against state-of-the-art.*: Tab. III reports the tracking accuracy of HERO against AMO [39] and

TABLE IV: Learned forward models improve end-effector tracking in the real world (Sec. V-D2). We study the impact of using different estimates, MOCAP vs. Analytical Forward Kinematics (FK) vs. our neural model, for end-effector (EE) pose and base pose. Use of our proposed neural forward models improves end-effector tracking accuracy over using analytical forward kinematics and achieves comparable errors to an oracle version that uses estimates from MOCAP (last row).

EE Pose	Base Pose	Trans. Error (cm)	Orient. Error (deg)	Joint Error (rad)
FK	FK	4.67 ± 1.30	14.59 ± 3.99	0.20 ± 0.03
Ours	FK	3.35 ± 0.70	14.07 ± 3.93	0.19 ± 0.03
FK	Ours	3.89 ± 1.06	14.28 ± 4.75	0.20 ± 0.04
Ours	Ours	2.56 ± 1.23	12.06 ± 4.38	0.18 ± 0.03
<i>MoCap</i>	<i>MoCap</i>	2.44 ± 0.86	14.29 ± 4.55	0.21 ± 0.05

TABLE V: Importance of Replanning (Sec. III-C) and Goal Adjustment (Sec. III-D) in the real world. Both goal adjustment and replanning lead to improvement in end-effector tracking accuracy, with replanning being more important. All methods in this table use accurate end-effector and base pose from MOCAP.

Method	Trans. Error (cm)	Orient. Error (deg)	Joint Error (rad)
w/o Replan	5.17 ± 2.21	16.13 ± 4.66	0.21 ± 0.03
w/o Goal Adjustment	2.71 ± 0.87	9.38 ± 2.72	0.20 ± 0.03
HERO (full)	2.44 ± 0.86	8.22 ± 3.52	0.21 ± 0.05

FALCON [94] in simulation across 3 different table heights: 0.5m, 0.74m, and 0.88m. HERO obtains a much lower end-effector translation and rotation error. Average translation error for HERO is 2.48cm compared to 8.29cm for AMO [39] and 13.57cm for FALCON [94]. Interestingly, HERO has a larger joint tracking error but a much lower end-effector tracking error. Unlike baseline methods that optimize purely in joint space, HERO leverages end-effector shift observations to directly improve task-space accuracy, demonstrating that minimizing joint error does *not* guarantee optimal end-effector positioning.

2) *Importance of Accurate Forward Kinematics*: Tab. IV reports control experiments where we replace the end-effector and base pose to be from analytical forward kinematics or MOCAP rather than our learned model. Using accurate end-effector and base pose from our learned models (denoted by ours) leads to the lowest errors. These errors are to those of an oracle that uses ground truth estimates from MOCAP (last row). Using FK instead of our learned model for either the end-effector or the base leads to worse performance.

3) *Importance of Replanning and Goal Adjustment*: We next assess the importance of replanning and goal adjustment by removing each of these components one at a time. To maximally isolate the effect of these components, we conduct this experiment where all methods have access to end-effector and base pose from MOCAP. As reported in Tab. V, both

components are important with the replanning component being more important than the goal adjustment component.

Execution Curves. We further analyze how the error evolves as the execution proceeds with and without replanning in Fig. 8. The cyan line that doesn’t use replanning plateaus off at a much higher error, whereas the blue line with replanning achieves a much lower overall error. In addition, the close performance of HERO using MOCAP observations and learned neural forward models demonstrates that our learned forward models are feasible and accurate. This ensures the possibility of transferring HERO to anywhere outside the ideal experimental space using MOCAP, making HERO a useful real-world system.

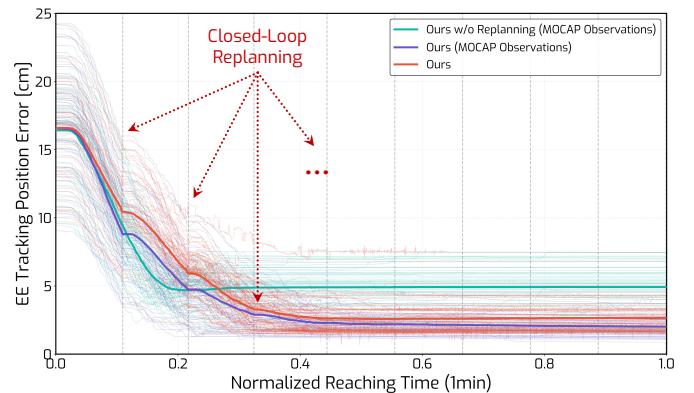


Fig. 8: Impact of replanning on end-effector tracking error in the real world. We plot the end-effector translation error as a function of execution time steps. The plot shows 1 minute of execution at 50Hz. The transparent lines are individual 60 real-world rollouts, and the corresponding solid line indicates the average value. The gray vertical dashlines indicate the replanning every 6 seconds (0.15Hz). Cyan line shows HERO without replanning and while purple line shows HERO with replanning. Re-planning leads to more accurate tracking. Orange line uses end-effector estimates from our neural model which leads to tracking performance very close to the oracle purple line that uses end-effector estimates from MOCAP.

E. Failure Mode Analysis

Fig. 9 categorizes the observed failures into two distinct modes, both emphasizing the need for precise end-effector control. The first mode, *object slipping* (Fig. 9(a)), arises when the limited dexterity of the Dex-3 hand fails to secure large, irregularly shaped objects. The second mode, *object knocked over* (Fig. 9(b)), reveals the importance of feasible grasping orientation and the limitation of the Dex-3 hand that has simply too large fingers. When interacting with unstable items like standing books, the margin for error is minimal; an insufficient hand orientation retargeting during reaching can trigger a collision. This highlights a stringent accuracy requirement for an appropriate retargeting strategy that helps stable and accurate grasping with the clumsy Dex-3 hand.

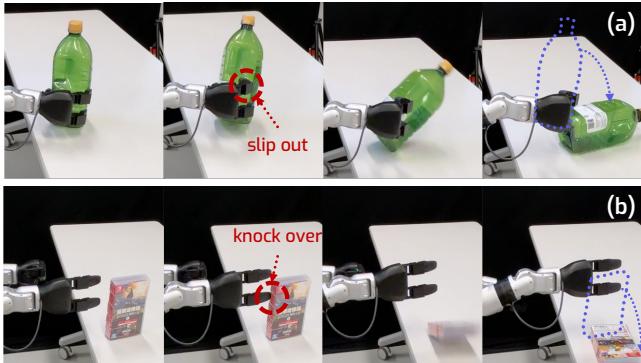


Fig. 9: **Failure mode examples.** We identify two main failure modes: a) *object slipping out* during grasping, which happens when grasping large objects with irregular shapes that are challenging for a Dex-3 hand endowed with limited dexterity; b) *object knocked over* during reaching, which happens when the retargeted hand orientation is not sufficiently large while the object stands unstably (e.g., a game cartridge or a book that has a thin ridge).

TABLE VI: **Reachable workspace volume** across configurations.

Configuration	Single Arm (m ³)	Combined (LUR) (m ³)
Arms-only (14 DoFs)	0.166	0.248
Arms+Waist (17 DoFs)	0.426	0.523

F. Whole Body Reaching Space Analysis

1) *Importance of Waist Bending:* To quantify how enabling torso motion via the waist bending affects end-effector reachability, we estimate the EE workspace under two kinematic settings: *i) arms-only*, where IK optimizes the 14-DoF arm joints, and *ii) arms+waist*, where IK additionally optimizes the waist DoFs (17 DoFs total). We define an axis-aligned 3D candidate region in the robot base frame:

$$x \in [0, 1.0] \text{ m}, \quad y \in [-1.0, 1.0] \text{ m}, \quad z \in [-0.5, 1.0] \text{ m} \quad (2)$$

and uniformly discretize it with a grid resolution of 0.02 m. For each sampled point \mathbf{p} , we test feasibility by solving inverse kinematics with cuRobo [68] under joint-limit constraints. A point is marked *reachable* if the IK solver converges within a fixed iteration budget and achieves an EE position residual below a preset tolerance; otherwise, it is marked *unreachable*. The reachable workspace volume is approximated via voxel counting:

$$V \approx N_{\text{reach}} \cdot (0.02)^3, \quad (3)$$

where N_{reach} denotes the number of reachable grid points.

The results are shown in Tab. VI. Quantitatively, enabling waist DoFs substantially increases workspace volume. The combined workspace (LUR) grows from 0.248 m³ in the arms-only setting to 0.523 m³ with arms+waist, corresponding to a $\sim 2.1 \times$ increase. Similarly, the single-arm workspace increases from 0.166 m³ to 0.426 m³. This gain is primarily attributed to

the waist’s control to bend and twist the torso, which effectively repositions the shoulder frame and allows the EE to cover farther-forward and lower-height targets that are infeasible with a non-actuated waist.

2) *Workspace Showcasing:* Fig. 10 illustrates HERO retrieving various objects distributed across an expansive tabletop workspace. Since every object is positioned beyond 0.4m to the robot base, the task requires whole-body coordination to maximize reachability and precision. As shown, HERO enables the robot to coordinate expressive whole-body motion while maintaining the necessary precision for successful grasping.

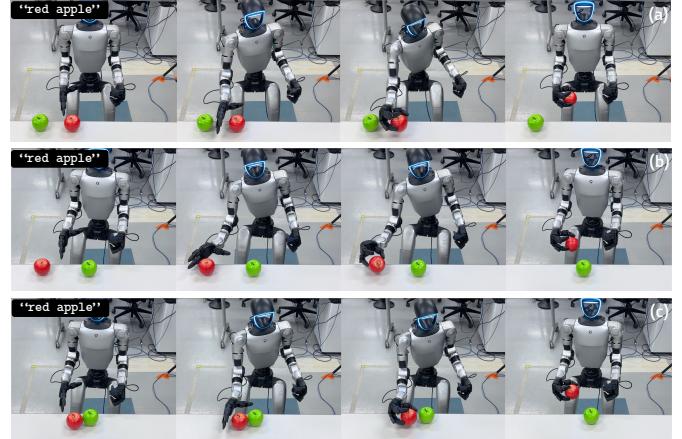


Fig. 10: **HERO enables a humanoid picking up objects from a standard table (0.74m) across a large workspace with open-vocabulary queries.** (a-c) The robot can reach and pick up a red apple placed at different heights, poses, and locations.

VI. LIMITATIONS

This work proposes HERO, which learns to control humanoid end-effectors, with which open-vocabulary visual loco-manipulation task is targeted. While HERO achieves remarkable results, we identify the following potential limitations. i) The egocentric field of view (FoV) under current hardware setup is rather limited, making it difficult for the humanoid to see objects farther than 1m or higher than 0.9m. Besides, as stated in the paper, the humanoid may twist the whole body a lot, which makes the object disappear from egocentric vision. One potential is to enable neck DoFs that achieves active vision and action [61, 93], which enables the robot to achieve improved closed-loop reasoning from vision. ii) HERO relies on a classical motion planner [68], which may lead to extremely twisted motion that is not an optimal or energy-efficient choice for loco-manipulation. One potential is to leverage a learning-based prior for trajectory optimization, which may help unleash more energy-efficient locomotion. iii) As HERO is a modular system, the limitations of these systems also apply. For example, the failure may come from submodules like the LVMs when the operating environment is too complicated [20, 47]. iv) The current system is limited in dexterity as we use a Dex-3 hand. With the rapid progress

in tackling the “hand dexterity challenge”, there exists a great potential to leverage better embodiment for a more dexterous whole-body dexterous manipulation, which remains a critical challenge in robotics research.

VII. DISCUSSIONS

We designed a system that enables a humanoid robot to grasp open-vocabulary novel objects in novel environments. This is a challenging task as it requires strong visual generalization (grounding open-vocabulary queries into robot’s egocentric perception) as well as strong control performance (squatting, bending, twisting to reach for objects while maintaining stability). Our system achieves a 90% success rate in the real world on this challenging task. This was made possible by adopting a decomposed design that separated *action planning* (ingesting RGB-D images to select specified objects and predict grasps for them) from *action execution* (actually conveying the robot to the predicted grasps). The separation enabled the use of large pre-trained models for action planning, leading to strong generalization and open-vocabulary capabilities. At the same time, *action execution* benefited from large-scale simulated training. Overall, we demonstrated that there is a more scalable alternative to the current practice of real world imitation learning for getting humanoid robots to rearrange objects.

Along the way, we encountered a number of surprising findings. It turned out, forward kinematics on a humanoid robot isn’t accurate and leads to the end-effector pose estimates that are off by 1.76cm on average. However, the error are systematic. This enabled the design of a residual neural forward model that reduced the error to 0.27cm.

Second, even though many recent works produce whole-body humanoid trackers, it turns out they incur large (8 – 13cm) end-effector tracking errors. This precludes the use of such trackers for manipulation applications. However, as we demonstrated, by careful design, particularly one that combines classical robotics (forward kinematics, inverse kinematics, and motion planning) with learning, it is possible to build accurate end-effector tracking policies. Our tracking policies achieve an end-effector tracking error rate of 2.5cm in the real world.

Finally, we believe that this successful demonstration of a modular system that enables a humanoid robot to manipulate objects will unlock the unification of the work on manipulation and humanoid control. Being able to accurately control humanoid end-effectors would mean that manipulation policies trained for other platforms (*e.g.*, via UMI [11, 21]) can be readily deployed on versatile humanoids.

ACKNOWLEDGMENTS

We are grateful to the Coordinated Science Laboratory and the Center for Autonomy for their help with experimental space and the MOCAP room, and for lending us the G1 humanoid robot used in this work. This work was supported by an NSF CAREER Award (IIS-2143873).

REFERENCES

- [1] Optitrack. URL <https://optitrack.com/>. 5, 19
- [2] Helix: A vision-language-action model for generalist humanoid control, 2025. URL <https://www.figure.ai/news/helix>. 2, 3
- [3] Karl Johan Åström and Peter Eykhoff. System identification—a survey. *Automatica*, 7(2):123–162, 1971. 3
- [4] Jose Barreiros, Andrew Beaulieu, Aditya Bhat, Rick Cory, Eric Cousineau, Hongkai Dai, Ching-Hsin Fang, Kunitatsu Hashimoto, Muhammad Zubair Irshad, Masha Itkina, et al. A careful examination of large behavior models for multitask dexterous manipulation. *arXiv preprint arXiv:2507.05331*, 2025. 2
- [5] Qingwei Ben, Feiyu Jia, Jia Zeng, Junting Dong, Dahu Lin, and Jiangmiao Pang. HOMIE: Humanoid Loco-Manipulation with Isomorphic Exoskeleton Cockpit. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025. 3
- [6] Gary Bradski. The opencv library. *Dr. Dobb’s Journal: Software Tools for the Professional Programmer*, 25(11):120–123, 2000. 19
- [7] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris Coll-Vinent, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, RISHI HAZRA, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollar, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer. SAM 3: Segment anything with concepts. In *The Fourteenth International Conference on Learning Representations*, 2026. 2, 4, 18, 20
- [8] Yiming Chen, Zekun Qi, Wenyao Zhang, Xin Jin, Li Zhang, and Peidong Liu. Reasoning in space via grounding in the world. In *The Fourteenth International Conference on Learning Representations*, 2026. 17
- [9] Zixuan Chen, Mazeyu Ji, Xuxin Cheng, Xuanbin Peng, Xue Bin Peng, and Xiaolong Wang. Gmt: General motion tracking for humanoid whole-body control. *arXiv preprint arXiv:2506.14770*, 2025. 3
- [10] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive Whole-Body Control for Humanoid Robots. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. 2
- [11] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. 11
- [12] Murtaza Dalal, Min Liu, Walter Talbott, Chen Chen, Deepak Pathak, Jian Zhang, and Ruslan Salakhutdinov.

Local policies enable zero-shot long-horizon manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13875–13882. IEEE, 2025. 2

[13] Hao-Shu Fang, Chenxi Wang, Hongjie Fang, Minghao Gou, Jirong Liu, Hengxu Yan, Wenhui Liu, Yichen Xie, and Cewu Lu. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics*, 39(5):3929–3945, 2023. 2, 4, 6, 18

[14] Hao-Shu Fang, Hengxu Yan, Zhenyu Tang, Hongjie Fang, Chenxi Wang, and Cewu Lu. Anydexgrasp: General dexterous grasping for different hands with human-level learning efficiency. *arXiv preprint arXiv:2502.16420*, 2025. 18

[15] Nolan Fey, Gabriel B. Margolis, Martin Peticco, and Pulkit Agrawal. Bridging the Sim-to-Real Gap for Athletic Loco-Manipulation. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 3

[16] Mirjana Filipović, Veljko Potkonjak, and Miomir Vukobratović. Elasticity in humanoid robotics. *Scientific Technical Review, Military Technical Institute, Belgrade*, 1:24–33, 2007. 3

[17] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: Learning a unified policy for manipulation and locomotion. In Karen Liu and Dana Kulic andD Jeffrey Ichnowski, editors, *Conference on Robot Learning, CoRL 2022, 14–18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 138–149. PMLR, 2022. 3

[18] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *8th Annual Conference on Robot Learning*, 2024. 2, 3

[19] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014. 19

[20] Arjun Gupta, Michelle Zhang, Rishik Sathua, and Saurabh Gupta. Demonstrating MOSART: Opening Articulated Structures in the Real World. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 10, 17

[21] Huy Ha, Yihuai Gao, Zipeng Fu, Jie Tan, and Shuran Song. Umi-on-legs: Making manipulation policies mobile with manipulation-centric whole-body controllers. In *Conference on Robot Learning*, pages 5254–5270. PMLR, 2025. 3, 11

[22] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris M. Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *8th Annual Conference on Robot Learning*, 2024. 2, 3, 4, 20

[23] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2024, Abu Dhabi, United Arab Emirates, October 14–18, 2024*, pages 8944–8951. IEEE, 2024. 2

[24] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbabu, Chaoyi Pan, Zeji Yi, Guannan Qu, Kris Kitani, Jessica K. Hodgins, Linxi Fan, Yuke Zhu, Changliu Liu, and Guanya Shi. ASAP: Aligning Simulation and Real-World Physics for Learning Agile Humanoid Whole-Body Skills. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 2, 3, 20

[25] Tairan He, Zi Wang, Haoru Xue, Qingwei Ben, Zhengyi Luo, Wenli Xiao, Ye Yuan, Xingye Da, Fernando Castañeda, Shankar Sastry, et al. Viral: Visual sim-to-real at scale for humanoid loco-manipulation. *arXiv preprint arXiv:2511.15200*, 2025. 3

[26] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Jan Kautz, Changliu Liu, Guanya Shi, Xiaolong Wang, et al. Hover: Versatile neural whole-body controller for humanoid robots. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9989–9996. IEEE, 2025. 3

[27] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Learning Getting-Up Policies for Real-World Humanoid Robots. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 2, 20

[28] Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic manipulation through spatial constraints of parts with foundation models. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9488–9495. IEEE, 2024. 17

[29] Tao Huang, Junli Ren, Huayi Wang, Zirui Wang, Qingwei Ben, Muning Wen, Xiao Chen, Jianan Li, and Jiangmiao Pang. Learning Humanoid Standing-up Control across Diverse Postures. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 2

[30] Tao Huang, Huayi Wang, Junli Ren, Kangning Yin, Zirui Wang, Xiao Chen, Feiyu Jia, Wentao Zhang, Junfeng Long, Jingbo Wang, and Jiangmiao Pang. Towards adaptable humanoid control via adaptive motion tracking. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, 2026. 3

[31] Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In *Conference on Robot Learning*, pages 4573–4602. PMLR, 2025. 17

[32] Mazeyu Ji, Xuanbin Peng, Fangchen Liu, Jialong Li, Ge Yang, Xuxin Cheng, and Xiaolong Wang. Exbody2: Advanced expressive humanoid whole-body control. In *RSS 2025 Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond*. 3

[33] Wolfgang Kabsch. A solution for the best rotation to

relate two sets of vectors. *Foundations of Crystallography*, 32(5):922–923, 1976. 5, 19

[34] F Kozin and HG Natke. System identification techniques. *Structural safety*, 3(3-4):269–316, 1986. 3

[35] Eric Krotkov, Douglas Hackett, Larry Jackel, Michael Perschbacher, James Pippine, Jesse Strauss, Gill Pratt, and Christopher Orlowski. The darpa robotics challenge finals: Results and perspectives. In *The DARPA robotics challenge finals: Humanoid robots to the rescue*, pages 1–26. Springer, 2018. 16

[36] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. RMA: rapid motor adaptation for legged robots. In *Robotics: Science and Systems XVII, Virtual Event, July 12-16, 2021*, 2021. 3

[37] Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Sci. Robotics*, 4(26), 2019. 3

[38] Kuang-Huei Lee, Ofir Nachum, Tingnan Zhang, Sergio Guadarrama, Jie Tan, and Wenhao Yu. PI-ARS: accelerating evolution-learned visual-locomotion with predictive information representations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2022, Kyoto, Japan, October 23-27, 2022*, pages 1447–1454. IEEE, 2022. 3

[39] Jialong Li, Xuxin Cheng, Tianshu Huang, Shiqi Yang, Ri-Zhao Qiu, and Xiaolong Wang. AMO: Adaptive Motion Optimization for Hyper-Dexterous Humanoid Whole-Body Control. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. 2, 3, 8, 9, 16

[40] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyo Seo, Georgios Pavlakos, and Yuke Zhu. Okami: Teaching humanoid robots manipulation skills through single video imitation. In *Conference on Robot Learning*, pages 299–317. PMLR, 2025. 3

[41] Yitang Li, Yuanhang Zhang, Wenli Xiao, Chaoyi Pan, Haoyang Weng, Guanqi He, Tairan He, and Guanya Shi. Hold my beer: Learning gentle humanoid locomotion and end-effector stabilization control. In *Conference on Robot Learning*, pages 4506–4523. PMLR, 2025. 3

[42] Yitang Li, Zhengyi Luo, Tonghe Zhang, Cunxi Dai, Anssi Kanervisto, Andrea Tirinzoni, Haoyang Weng, Kris Kitani, Mateusz Guzek, Ahmed Touati, Alessandro Lazaric, Matteo Pirotta, and Guanya Shi. BFM-zero: A promptable behavioral foundation model for humanoid control using unsupervised reinforcement learning. In *The Fourteenth International Conference on Learning Representations*, 2026. 3

[43] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. *Int. J. Robotics Res.*, 44(5):840–888, 2025. 3

[44] Qiayuan Liao, Takara E Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C Karen Liu. Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion. *arXiv preprint arXiv:2508.08241*, 2025. 2

[45] Fukang Liu, Zhaoyuan Gu, Yilin Cai, Ziyi Zhou, Hyunyoung Jung, Jaehwi Jang, Shijie Zhao, Sehoon Ha, Yue Chen, Danfei Xu, et al. Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation. *IEEE Robotics and Automation Letters*, 2025. 3

[46] Minghuan Liu, Zixuan Chen, Xuxin Cheng, Yandong Ji, Ri-Zhao Qiu, Ruihan Yang, and Xiaolong Wang. Visual whole-body control for legged loco-manipulation. In *Conference on Robot Learning*, pages 234–257. PMLR, 2025. 3

[47] Peiqi Liu, Yaswanth Orru, Jay Vakil, Chris Paxton, Nur Muhammad (Mahi) Shafullah, and Lerrel Pinto. Demonstrating ok-robot: What really matters in integrating open-knowledge models for robotics. In Dana Kulic, Gentiane Venture, Kostas E. Bekris, and Enrique Coronado, editors, *Robotics: Science and Systems XX, Delft, The Netherlands, July 15-19, 2024*, 2024. 2, 10, 17

[48] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer, 2024. 2, 4, 6, 20

[49] Ilya Loschilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 20

[50] Chenhao Lu, Xuxin Cheng, Jialong Li, Shiqi Yang, Mazeyu Ji, Chengjing Yuan, Ge Yang, Sha Yi, and Xiaolong Wang. Mobile-television: Predictive motion priors for humanoid whole-body control. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5364–5371. IEEE, 2025. 3

[51] Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Sirui Chen, Fernando Castañeda, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, et al. Sonic: Supersizing motion tracking for natural humanoid whole-body control. *arXiv preprint arXiv:2511.07820*, 2025. 3

[52] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 4

[53] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance GPU based physics simulation for robot learning. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021. 7, 20

[54] Jiageng Mao, Siheng Zhao, Siqi Song, Chuye Hong, Tianheng Shi, Junjie Ye, Mingtong Zhang, Haoran Geng,

Jitendra Malik, Vitor Guizilini, and Yue Wang. Universal humanoid robot pose learning from internet human videos. In *2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)*, pages 1–8, 2025. 3

[55] James Ni, Zekai Wang, Wei Lin, Amir Bar, Yann LeCun, Trevor Darrell, Jitendra Malik, and Roei Herzig. From generated human videos to physically plausible robot trajectories. *arXiv preprint arXiv:2512.05094*, 2025. 3

[56] Michael O’Connell, Guanya Shi, Xichen Shi, Kamyar Azizzadenesheli, Anima Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural-fly enables rapid learning for agile flight in strong winds. *Sci. Robotics*, 7(66), 2022. 3

[57] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4):143, 2018. 2

[58] Xue Bin Peng, Erwin Coumans, Tingnan Zhang, Tsang-Wei Edward Lee, Jie Tan, and Sergey Levine. Learning agile robotic locomotion skills by imitating animals. In Marc Toussaint, Antonio Bicchi, and Tucker Hermans, editors, *Robotics: Science and Systems XVI, Virtual Event / Corvalis, Oregon, USA, July 12-16, 2020*, 2020. 3

[59] Oliver Porges, Mathilde Connan, Bernd Henze, Andrea Gigli, Claudio Castellini, and Máximo Alejandro Roa Garzon. A wearable, ultralight interface for bimanual teleoperation of a compliant, whole-body-controlled humanoid robot. In *2019 International Conference on Robotics and Automation, ICRA 2019*. IEEE, 2019. 3

[60] Tifanny Portela, Andrei Cramariuc, Mayank Mittal, and Marco Hutter. Whole-body end-effector pose tracking. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11205–11211. IEEE, 2025. 3

[61] Haozhi Qi, Yen-Jen Wang, Toru Lin, Brent Yi, Yi Ma, Koushil Sreenath, and Jitendra Malik. Coordinated humanoid manipulation with choice policies. *arXiv preprint arXiv:2512.25072*, 2025. 10

[62] Zekun Qi, Wenyao Zhang, Yufei Ding, Runpei Dong, XinQiang Yu, Jingwen Li, Lingyun Xu, Baoyu Li, Xialin He, Guofan Fan, Jiazhao Zhang, Jiawei He, Jiayuan Gu, Xin Jin, Kaisheng Ma, Zhizheng Zhang, He Wang, and Li Yi. Sofar: Language-grounded orientation bridges spatial reasoning and object manipulation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. 2, 17

[63] Tianhe Ren, Qing Jiang, Shilong Liu, Zhaoyang Zeng, Wenlong Liu, Han Gao, Hongjie Huang, Zhengyu Ma, Xiaoke Jiang, Yihao Chen, et al. Grounding dino 1.5: Advance the “edge” of open-set object detection. *arXiv preprint arXiv:2405.10300*, 2024. 18, 20

[64] Andreas J Schmid, Nicolas Gorges, Dirk Goger, and Heinz Worn. Opening a door with a humanoid robot using multi-sensory tactile feedback. In *2008 IEEE international conference on robotics and automation*, pages 285–291. IEEE, 2008. 16

[65] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 4

[66] Mili Shah, Roger D Eastman, and Tsai Hong. An overview of robot-sensor calibration methods for evaluation of perception systems. In *Proceedings of the Workshop on Performance Metrics for Intelligent Systems*, pages 15–20, 2012. 19

[67] Christopher Stanton, Anton Bogdanovych, and Edward Ratanasena. Teleoperation of a humanoid robot using full-body motion capture, example movements, and machine learning. In *Proc. Australasian Conference on Robotics and Automation*, volume 8, page 51, 2012. 3

[68] Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, et al. Curobo: Parallelized collision-free robot motion generation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8112–8119. IEEE, 2023. 3, 6, 10, 20

[69] Amy Tabb and Khalil M Ahmad Yousef. Solving the robot-world hand-eye (s) calibration problem with iterative methods. *Machine Vision and Applications*, 28(5):569–590, 2017. 19

[70] Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. In Hadas Kress-Gazit, Siddhartha S. Srinivasa, Tom Howard, and Nikolay Atanasov, editors, *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. 3

[71] Johannes Tenhumberg and Berthold Bäuml. Calibration of an elastic humanoid upper body and efficient compensation for motion planning. In *20th IEEE-RAS International Conference on Humanoid Robots, Humanoids 2021, Munich, Germany, July 19-21, 2021*, pages 98–103. IEEE, 2021. 3

[72] Johannes Tenhumberg, Dominik Winkelbauer, Darius Burschka, and Berthold Bäuml. Self-contained calibration of an elastic humanoid upper body using only a head-mounted RGB camera. In *21st IEEE-RAS International Conference on Humanoid Robots, Humanoids 2022, Ginowan, Japan, November 28-30, 2022*, pages 702–707. IEEE, 2022. 3

[73] Tesla. Artificial intelligence & autopilot, 2021. URL <https://www.tesla.com/AI>. 2, 3

[74] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017. 20

[75] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pages 5026–5033. IEEE, 2012. 7,

[76] Roger Y Tsai. Efficient and accurate camera calibration technique for 3d machine vision. In *IEEE conference on computer vision and pattern recognition*, 1985. 19

[77] Roger Y Tsai, Reimar K Lenz, et al. A new technique for fully autonomous and efficient 3 d robotics hand/eye calibration. *IEEE Transactions on robotics and automation*, 5(3):345–358, 1989. 19

[78] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 13(4):376–380, 2002. 5, 19

[79] Unitree. Unitree H1 / H1-2: Unitree’s first universal humanoid robot. 2023. URL <https://www.unitree.com/h1>. 3

[80] Unitree. Unitree G1: Humanoid Agent AI Avatar. 2024. URL <https://www.unitree.com/g1>. 3

[81] Hongxi Wang, Haoxiang Luo, Wei Zhang, and Hua Chen. CTS: concurrent teacher-student reinforcement learning for legged locomotion. *IEEE Robotics Autom. Lett.*, 9(11):9191–9198, 2024. 3

[82] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. 17

[83] Lai Wei, Xuanbin Peng, Ri-Zhao Qiu, Xuxin Cheng, and Xiaolong Wang. Hmc: Learning heterogeneous meta-control for contact-rich loco-manipulation. In *RSS 2025 Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond*. 3

[84] Haoyang Weng, Yitang Li, Nikhil Sobanbabu, Zihan Wang, Zhengyi Luo, Tairan He, Deva Ramanan, and Guanya Shi. Hdmi: Learning interactive humanoid whole-body control from human videos. *arXiv preprint arXiv:2509.16757*, 2025. 3

[85] Haoru Xue, Tairan He, Zi Wang, Qingwei Ben, Wenli Xiao, Zhengyi Luo, Xingye Da, Fernando Castañeda, Guanya Shi, Shankar Sastry, et al. Opening the sim-to-real door for humanoid pixel-to-action policy transfer. *arXiv preprint arXiv:2512.01061*, 2025. 16

[86] Lujie Yang, Xiaoyu Huang, Zhen Wu, Angjoo Kanazawa, Pieter Abbeel, Carmelo Sferrazza, C Karen Liu, Rocky Duan, and Guanya Shi. Omnidretarget: Interaction-preserving data generation for humanoid whole-body loco-manipulation and scene interaction. *arXiv preprint arXiv:2509.26633*, 2025. 2

[87] Shaofeng Yin, Yanjie Ze, Hong-Xing Yu, C Karen Liu, and Jiajun Wu. Visualmimic: Visual humanoid loco-manipulation via motion tracking and generation. *arXiv preprint arXiv:2509.20322*, 2025. 3

[88] Wenhao Yu, Jie Tan, C. Karen Liu, and Greg Turk. Preparing for the unknown: Learning a universal policy with online system identification. In Nancy M. Amato, Siddhartha S. Srinivasa, Nora Ayanian, and Scott Kuindersma, editors, *Robotics: Science and Systems XIII, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, July 12-16, 2017*, 2017. 3

[89] Wenhao Yu, C. Karen Liu, and Greg Turk. Policy transfer with strategy optimization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

[90] Wenhao Yu, Jie Tan, Yunfei Bai, Erwin Coumans, and Sehoon Ha. Learning fast adaptation with meta strategy optimization. *IEEE Robotics Autom. Lett.*, 5(2):2950–2957, 2020. 3

[91] Yanjie Ze, Zixuan Chen, Joao Pedro Araujo, Zi-ang Cao, Xue Bin Peng, Jiajun Wu, and Karen Liu. Twist: Teleoperated whole-body imitation system. In Joseph Lim, Shuran Song, and Hae-Won Park, editors, *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 2143–2154. PMLR, 27–30 Sep 2025. 3

[92] Yanjie Ze, Zixuan Chen, Wenhao Wang, Tianyi Chen, Xialin He, Ying Yuan, Xue Bin Peng, and Jiajun Wu. Generalizable humanoid manipulation with 3d diffusion policies. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2873–2880. IEEE, 2025. 2

[93] Yanjie Ze, Siheng Zhao, Weizhuo Wang, Angjoo Kanazawa, Rocky Duan, Pieter Abbeel, Guanya Shi, Jiajun Wu, and C Karen Liu. Twist2: Scalable, portable, and holistic humanoid data collection system. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, 2026. 3, 10

[94] Yuanhang Zhang, Yifu Yuan, Prajwal Gurunath, Ishita Gupta, Shayegan Omidshafiei, Ali-akbar Aghamohammadi, Marcell Vazquez-Chanlatte, Liam Pedersen, Tairan He, and Guanya Shi. Falcon: Learning force-adaptive humanoid loco-manipulation. In *8th Learning for Dynamics & Control Conference*, 2026. 2, 3, 8, 9, 16

[95] Zhikai Zhang, Chao Chen, Han Xue, Jilong Wang, Sikai Liang, Yun Liu, Zongzhang Zhang, He Wang, and Li Yi. Unleashing humanoid reaching potential via real-world-ready skill space. *IEEE Robotics and Automation Letters*, 11(2):2082–2089, 2025. 3

[96] Zhikai Zhang, Jun Guo, Chao Chen, Jilong Wang, Chenghuai Lin, Yunrui Lian, Han Xue, Zhenrong Wang, Maoqi Liu, Jiangran Lyu, et al. Track any motions under any disturbances. *arXiv preprint arXiv:2509.13833*, 2025. 3

[97] Siheng Zhao, Yanjie Ze, Yue Wang, C Karen Liu, Pieter Abbeel, Guanya Shi, and Rocky Duan. Resmimic: From general motion tracking to humanoid whole-body loco-manipulation via residual learning. *arXiv preprint arXiv:2510.05070*, 2025. 3

[98] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019. 5, 20

APPENDIX

A Additional Experimental Analysis	16
A.1 Language Sensitivity	16
A.2 Tracking Error Distribution Analysis	16
A.3 Moving Object Grasping with Visual Replanning	16
A.4 Extending HERO to Other Tasks Like Door Opening	16
A.5 Field of View Analysis	17
A.6 Visual Perception Illustration	17
A.7 Analytical FK Error Visualization	19
B Additional Implementation Details	19
B.1 MOCAP Setup	19
B.2 Onboard Egocentric RGB-D Camera	19
B.3 Hyper-parameters	19
B.4 Rewards	20
B.5 Policy Training	20
B.6 Deployment Hardware	20
B.7 Testing Assets Details	20
B.8 Testing Scenes Details	20

A ADDITIONAL EXPERIMENTAL ANALYSIS

A.1 Language Sensitivity

Fig. 11 shows that our system can correctly interpret language to pick up the correct object among relevant distractors. It picks up the red apple (and not the green one) when told to pick up the red apple in the top row, and vice versa in the bottom row.

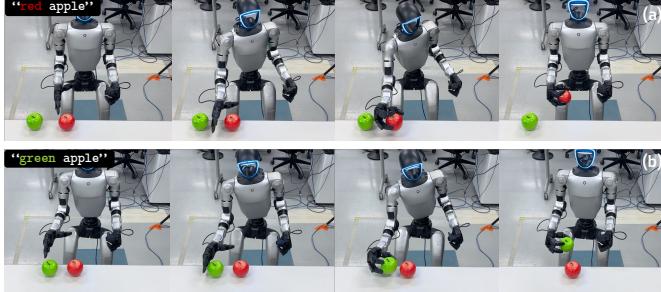


Fig. 11: **HERO is able to distinguish the target object via language queries.** (a) Picking up a red apple; (b) Picking up a green apple instead of red apple.

A.2 Tracking Error Distribution Analysis

We visualize the CDFs of EE tracking errors in Fig. 12. **Top row.** HERO dominates all baselines: at 80%, HERO achieves 3.9 cm / 6.7° (pos/rot), versus 20.9 cm / 25.3° for FALCON and 9.8 cm / 19.5° for AMO. At 90%, HERO remains below 4.6 cm and 8.2°, indicating strong tail robustness. **Bottom row.** Replanning (Sec. III-C) is critical: at 80%, HERO achieves 3.1 cm vs. 6.3 cm without replanning (2.1× worse), and at 90% 3.4 cm vs. 7.1 cm. Rotation gains are smaller but consistent (median 13.6° vs. 16.1°).

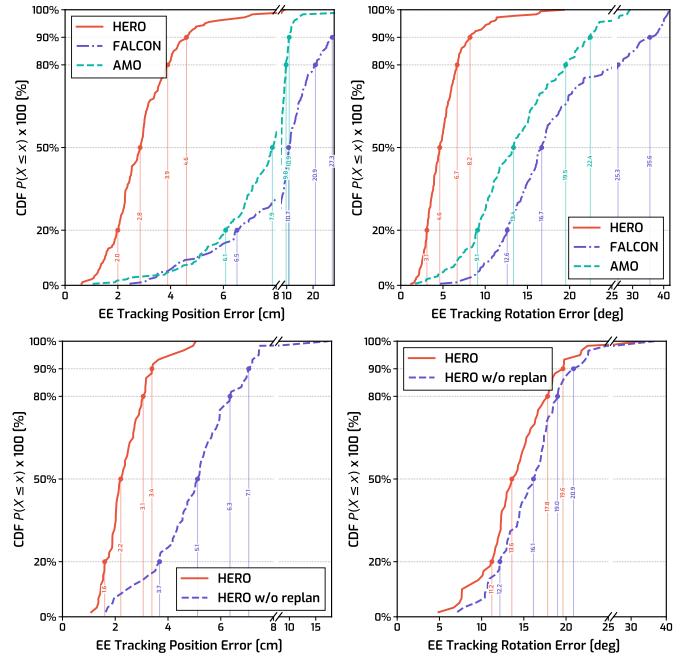


Fig. 12: **CDF analysis of end-effector tracking errors.** Top row: Comparison of translation (left) and rotation (right) error distributions across HERO, FALCON [94], and AMO [39] for all table heights. The steeper curves of HERO indicate consistently lower errors and tighter distributions. Bottom row: Ablation study showing translation (left) and rotation (right) error distributions with and without replanning (Sec. III-C). The steeper CDF curves with replanning demonstrate their significant contribution to tracking accuracy.

A.3 Moving Object Grasping with Visual Replanning

HERO derives target grasping poses from vision and language queries, enabling the robot to capture moving objects through closed-loop replanning. Fig. 13(a-b) illustrates this dynamic adaptation: while the system initially generates an EE trajectory based on the first vision perception, it re-estimates the pose as the object moves. This visual feedback triggers the update of the target grasp, allowing the robot to seamlessly adjust its trajectory and successfully secure the moving object. Note that, in these two trials, the robot successfully sees the object after moving, but the object can only be seen at the corner due to the rather limited field of view.

A.4 Extending HERO to Other Tasks Like Door Opening

As HERO constructed a modular system that coordinates high-level planning and low-level end-effector control, it reveals a possibility of extending HERO to broader tasks like door opening, which is a challenging loco-manipulation task [35, 64, 85]. In Fig. 14, we directly employ HERO to identify the target grasping pose for fridge door handle, followed by the same pipeline as object grasping, enabling the humanoid to grasp the fridge handle and finally open it when returning to the default position. Note that the door requires a large force to open because of magnetic attraction; we leave the

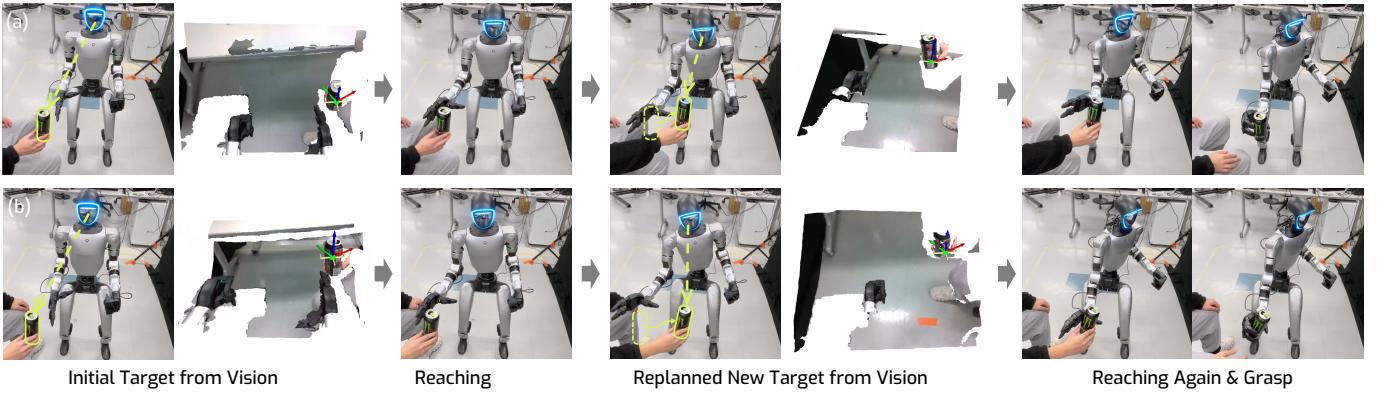


Fig. 13: **HERO enables a humanoid to grasp a moving object via visual closed-loop replanning.** (a-b) Two examples of visual closed-loop replanning. The goal is to grasp black can.

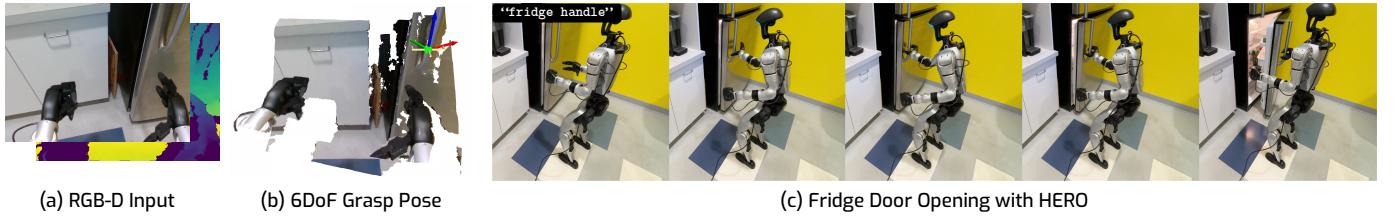


Fig. 14: **Door opening with HERO.** (a) The egocentric RGB-D visual inputs. (b) Given a language query for the door handle (e.g., ‘‘fridge handle’’ prompted here), our modular system obtains the grasping pose, same as the pipeline for picking up objects. (c) The robot executes the reaching trajectory and closes the hand when reaching the target poses closely, and then returns to the default pose with the door opened. Note that while the door is heavy, HERO successfully manages to open the door with a smooth and stable door motion.

door unlatched beforehand. This result shows *HERO*’s modular potential, and it is also possible to extend our system to broader loco-manipulation tasks by incorporating off-the-shelf trajectory generation frameworks [20, 28, 31, 62].

A.5 Field of View Analysis

Our system uses the onboard camera for visual perception, which, however, has a limited field of view. As shown in Fig. 15, the robot first stands randomly at a distance of 1.28m from the object, while the target object (stapler) is not visible at this distance; after walking forward under a consistent velocity command to a distance of 0.6m and continuously detecting the object, the egocentric view successfully captures the target object, which makes the robot stops at about 0.5m from the object. Then the robot coordinates the whole-body reaching motion and successfully grasps the object. This visualization indicates that the robot’s onboard egocentric view is limited, and the robot can only see the object within a short distance (< 0.6m), which makes the availability of the 3D spatial understanding beforehand [8, 47, 82] critical for searching objects, which could be a future exploration.

A.6 Visual Perception Illustration

Fig. 16 illustrates how *HERO* leverages LVMs to obtain the targeted EE grasping pose, following a modular perception-to-action design that is similar to prior systems [47, 62].

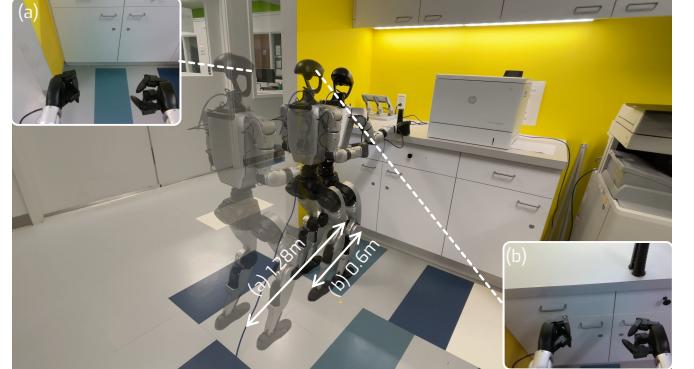


Fig. 15: **Field of view visualization using HERO.** We let the robot stand at a random distance from the object (e.g., 1.28m), and the robot keeps walking forward until the target object (stapler) is detected. After successfully detecting the target object, the robot stops walking and grasps the object via whole-body coordination. The robot can only see the object within 0.6m, which makes it hard for the robot to search for the object in a random room.

Object of Interest from Language. Given an ego-centric RGB-D observation and a natural-language query specifying the target object, *HERO* first applies GroundingDINO to produce

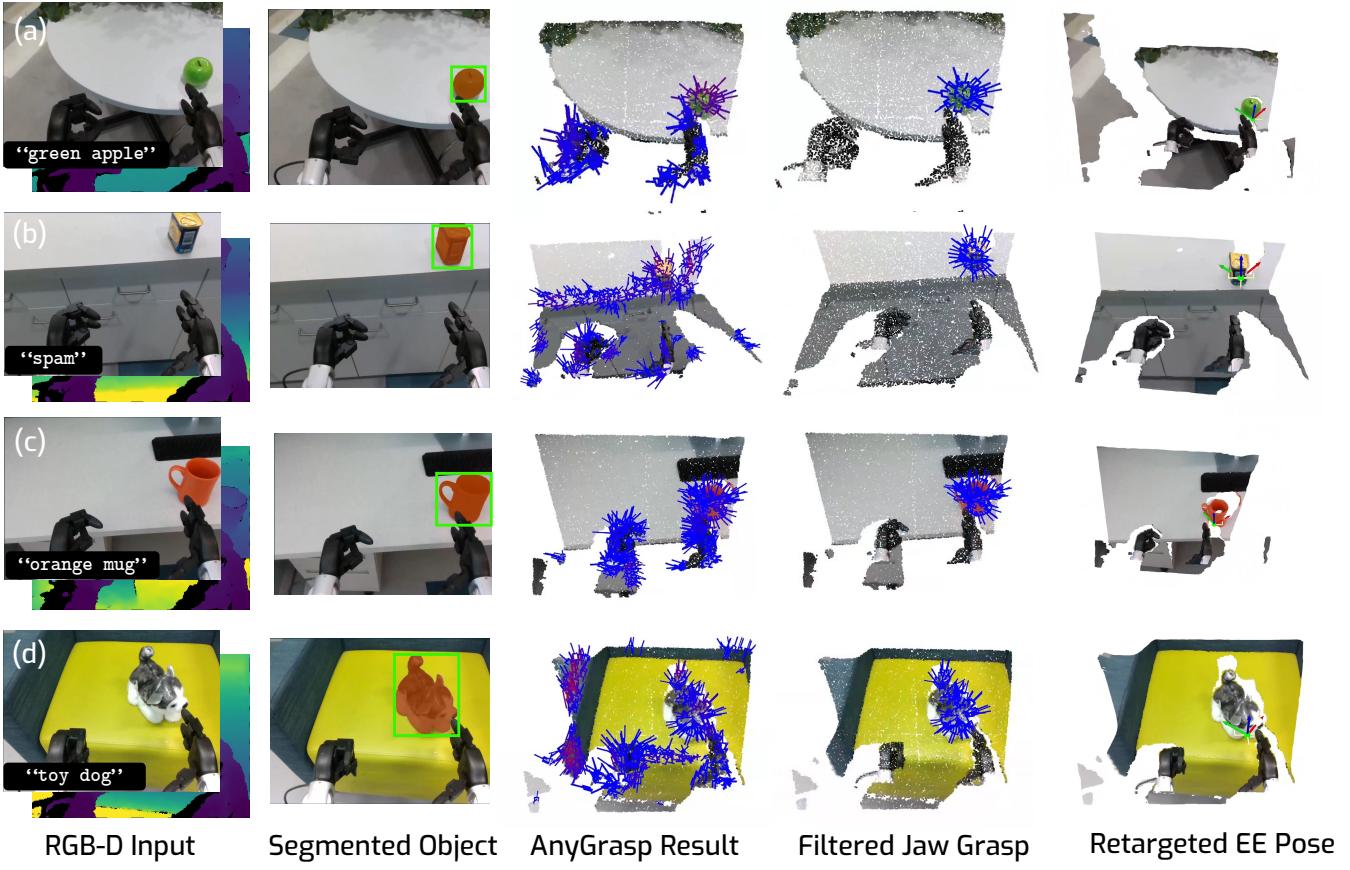


Fig. 16: **HERO visual perception pipeline illustration.** (a-d) Examples of ego-centric visual perception using LVMs, including GroundingDINO [63], SAM-3 [7], and AnyGrasp. Given the language query, GroundingDINO outputs the detection box, which is input to SAM for the segmentation mask. The mask is used to filter out jaw grasps predicted by AnyGrasp, which is finally retargeted to the 6-DoF end-effector pose for dexterous grasping with a Dex-3 hand.

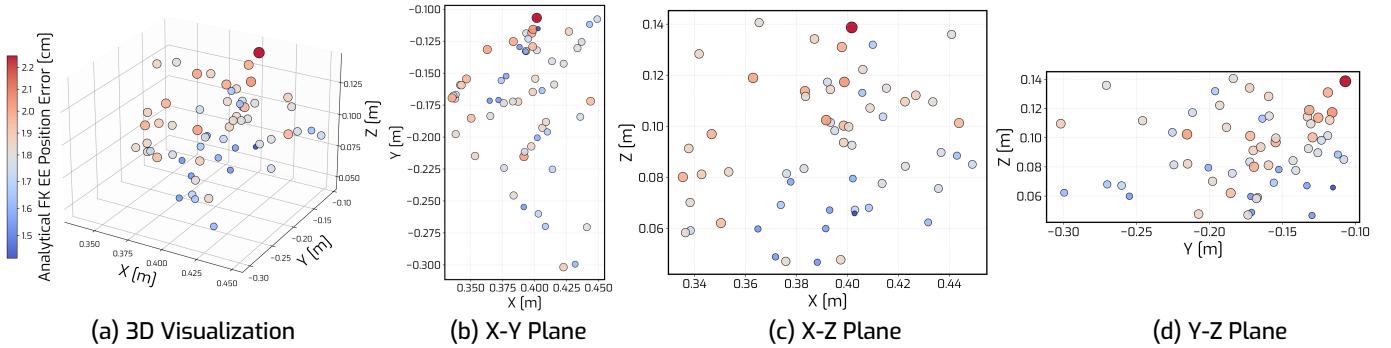


Fig. 17: **Visualization of analytical forward kinematics error.** We plot the 60 data points collected in the MOCAP room. The error is indicated in the color bar on the left side of the figure, and the size of the scatter also increases with the error.

a language-conditioned detection box [63]. The detected box is then used to prompt SAM-3 for segmentation of the object of interest [7].

Grasp Proposals. This mask serves as a spatial constraint for grasp proposal generation: HERO runs AnyGrasp [13] to produce a set of candidate grasps, and then filters out proposals outside the segmented object region. Note that

AnyDexGrasp [14] can also be used here, but we find that the Dex-3 hand lacks dexterity, and the difference between these two methods is limited.

Grasp Selection. To select the best jaw grasp, we first filter out the grasp poses that lie on the opposite side of the object relative to the robot’s hand (*e.g.*, for an object to the right side of the hand, the left approaching grasps are abandoned).

Then we filter out grasps that are too high or too low based on a gravity-aligned height estimation of objects using depth. Afterward, we select the grasp that lies most parallel to the ground with the highest confidence as the final grasp.

Grasp Retargeting. The selected grasp is retargeted to a 6-DoF end-effector pose for dexterous grasping with the Dex-3 hand. We first rotate the gripper pose by 45 degrees around the z -axis to improve the grasp robustness and pose error tolerance. After that, we clip the yaw angle within 70 degrees to ensure the orientation is not too large.

A.7 Analytical FK Error Visualization

In Fig. 17, we visualize the translation error of analytical forward kinematic results. We plot the error via the collected 60 samples in the MOCAP room, where the error is recorded when time is 1 minute. From the figure, we can observe that the error generally increases when the EE location becomes larger along the Y and Z axes, which may form a pattern that can be learned from a neural model.

B ADDITIONAL IMPLEMENTATION DETAILS

B.1 MOCAP Setup

MOCAP System. We use the modern MOCAP system Optitrack [1] with 13 cameras which provides $\leq 0.2\text{mm}$ measure accuracy.

Robot Link Pose. To obtain the end-effector pose in the robot frame, we put several markers onto both links, and we show markers on the hand in Fig. 18(a). Although the MOCAP system provides constructed asset poses via selected marker groups, there exists a misalignment between the MOCAP asset frame and the robot link frame. To address this, we carefully measure each marker’s relative offset to the link’s origin, followed by the Kabsch-Umeyama (KU) algorithm [33, 78] that transforms individual marker coordinates into 6-DoF link pose in the MOCAP frame within $< 1.5\text{mm}$ RMSE error. The relative transformation of EE and the robot base is thus obtained as they are all in the MOCAP frame. This approach ensures an accurate measurement of both the end-effector and the robot base, setting a solid ground for our evaluation and camera calibration, introduced next.

B.2 Onboard Egocentric RGB-D Camera

Setup. We use the onboard RGB-D camera D435i mounted on the humanoid’s head, as shown in Fig. 18(b). The humanoid’s neck features a pitch degree of freedom enabling head rotation within a limited range, necessitating precise camera calibration for accurate 3D perception.

Calibration with MOCAP. Standard hand-eye calibration [66, 69] typically relies on analytical forward kinematics to obtain end-effector poses. However, as demonstrated in the main paper (Sec. V-C), analytical forward kinematics exhibits systematic errors of approximately 1.8cm due to hardware inaccuracies—unsuitable for precise camera calibration.

We instead leverage the MOCAP system for ground-truth pose measurement. Following the marker-based approach described previously, we attach reflective markers to an ArUco

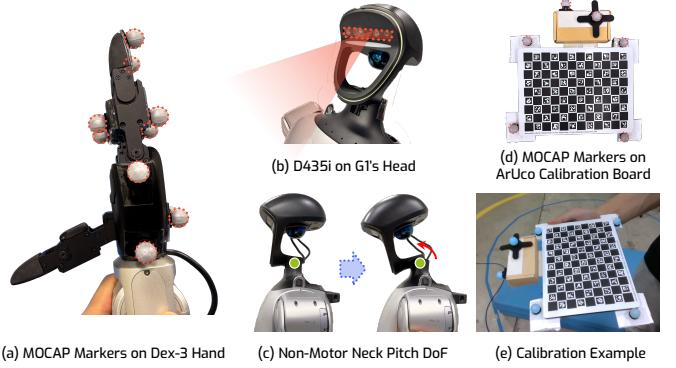


Fig. 18: **MOCAP markers, camera, and calibration setups.** (a) We put several MOCAP markers on both the robot’s end-effector and robot’s base (similar to the EE), and each marker’s relative location to the link’s base is measured carefully. By employing the Kabsch-Umeyama algorithm [33, 78], we are able to accurately obtain the robot’s link’s coordinate in the MOCAP frame from each marker’s individual coordinates in the MOCAP frame with $< 1.5\text{mm}$ RMSE error. (b) The onboard D435i camera mounted on the Unitree G1 humanoid robot’s head. (c) While no motor is set, there is a neck pitch DoF that allows the head to rotate along the y axis via external physical force, making the manufacturer-provided camera parameters far from the real setup. (d) Similar to EE and base, we put several MOCAP markers on a standard ArUco calibration board [19] to obtain an accurate relative transformation of the calibration board to the robot base. (e) Our calibration requires one person to hold the board in front of the camera to collect different board poses in the robot frame.

calibration board [19] and apply the KU algorithm [33, 78] for 6-DoF pose estimation. During data collection, we manually move the board through 60-70 diverse poses in front of the camera. For each pose i , we record: 1) the robot base pose in MOCAP frame $\mathcal{T}_{\text{MOCAP}}^{\text{base}}$, 2) the board pose in MOCAP frame $\mathcal{T}_{\text{MOCAP}}^{\text{board},i}$, and 3) the board pose in camera frame $\mathcal{T}_{\text{camera}}^{\text{board},i}$ via ArUco detection using the OpenCV library [6].

To compute the camera-to-base transformation $\mathcal{T}_{\text{base}}^{\text{camera}}$, we solve the eye-to-hand calibration problem:

$$\mathcal{T}_{\text{base}}^{\text{camera}} \oplus \mathcal{T}_{\text{camera}}^{\text{board},i} = \mathcal{T}_{\text{MOCAP}}^{\text{base}} \ominus \mathcal{T}_{\text{MOCAP}}^{\text{board},i} \quad (4)$$

using the Tsai-Lenz method [76, 77]. This MOCAP-assisted calibration achieves a reprojection error within 2.5mm, ensuring accurate egocentric 3D perception.

Image Resolution & FPS. We use the RGB-D images with a resolution of 640×480 in a 60Hz FPS.

B.3 Hyper-parameters

Motion planning For motion planning, we use cuRobo and set the planning dt to 7.25e-6.

Grasping Threshold When the robot approaches the object, it autonomously close the hand when the hand distance to the target grasp $\Delta\mathcal{E}_t \leq \delta$ where $\delta > 0$ is a threshold. At the moment when this threshold is reached, we pass the same

local waypoint of the planned motion trajectory to the policy to ensure stability, and the hand is immediately closed for grasping. In this paper, we utilize a threshold of $\delta = 1.5\text{cm}$, which we find most effective across tested objects.

B.4 Rewards

Tab. VII summarizes reward components and weights used for RL training of π_t , which is structured into four categories: *tracking task*, *penalties*, *regularization*, and *locomotion task*. To ensure precise manipulation, the tracking rewards weigh the alignment of the end-effector based on our newly proposed residual $\Delta\mathcal{E}_t$. Note that EE orientation is represented with the continuous 6D parameterization (first two columns of the rotation matrix) [98]. To encourage the planned upper-body posture (e.g., waist bending or torso twisting), we also add a joint-space tracking term. Penalties strictly enforce safety constraints (e.g., joint limits, termination), while regularization terms—such as costs on torque, acceleration, and stance symmetry—are essential for generating smooth, stable motions capable of robust and natural Sim2Real transfer. To train the robot to follow locomotion commands, we also use a flag variable to control the standing and waking mode switching.

B.5 Policy Training

Simulation & Training Setup. We train our end-effector tracking policy π_t with the IsaacGym simulator [53], and transfer this policy to the MuJoCo simulator [75] for Sim2Sim evaluations before deploying it in the real world. We train our policy with 4,096 environments for overall 20K iterations in parallel, with a learning rate of 1e-4 for both the actor and critic models. AdamW optimizer [49] is used with a weight decay of 1e-2. We use a high simulation frequency of 500Hz, with the low-level PD controller running at 50Hz. All the policy training is conducted on a single NVIDIA RTX 4090 or an L40S GPU.

Sim2Real Domain Randomization. Following previous works [22, 24, 27], we employ standard domain and dynamics randomization to facilitate Sim2Real transfer [74], including variations in link center of mass (CoM) and control delay. Notably, we identify that randomizing the *end-effector mass* is essential; without this specific randomization, the policy exhibits end-effector instability, leading to high-frequency hand oscillations that compromise tracking accuracy.

B.6 Deployment Hardware

We run all modules (e.g., π_t and SAM-3 [7]) off-the-shelf on a 32-GB RAM laptop equipped with NVIDIA RTX 5070Ti GPU and Intel Core Ultra 9 275HX CPU processor (24 CPU cores / 24 threads). We run cuRobo with CUDA graph acceleration, which largely improve the efficiency on the edge [68]. For the detection module, we have tested both Grounding DINO base [48] and Grounding DINO 1.5 [63], where the base version can be deployed on the laptop, and Grounding DINO 1.5 only provides access through online APIs. However, we find that Grounding DINO base is sufficient for most scenes and objects.

TABLE VII: **Reward components and weights.** Penalty rewards prevent unreasonable behaviors for sim2real transfer, regularization helps improve motion smoothness and stability, and task rewards ensure successful and precise end-effector and upper-body tracking.

TERM	EXPRESSION	WEIGHT
Tracking Task Rewards:		
End-effector exp	$\exp(-\ \Delta\mathcal{E}_t\ ^2)$	2.0
Upper-body DoF exp	$\exp(-\ \mathbf{q}_t^{\text{upper}(\text{Ref})} - \mathbf{q}_t^{\text{upper}}\ ^2)$	4.0
Base height exp	$\exp(-\ h^{\text{base}} - h^{\text{base}(\text{Ref})}\ ^2)$	4.0
Penalty:		
DoF position limits	$\mathbb{1}(\mathbf{q}_t \notin [\mathbf{q}_{\min}, \mathbf{q}_{\max}])$	-5.0
DoF velocity limits	$\mathbb{1}(\dot{\mathbf{q}}_t \notin [\mathbf{q}_{\min}, \mathbf{q}_{\max}])$	-5.0
Termination	$\mathbb{1}_{\text{termination}}$	-250
Regularization:		
End-effector linear velocity	$\ \mathbf{v}_{\text{EE}}^2\ $	-0.2
End-effector angular velocity	$\ \omega_{\text{EE}}^2\ $	-0.02
DoF acceleration	$\ \ddot{\mathbf{q}}_t\ _2$	-2.5e-7
DoF velocity	$\ \dot{\mathbf{q}}_t\ _2^2$	-1e-3
Action rate	$\ \mathbf{a}_t\ _2^2$	-0.1
Torque	$\ \tau_t\ $	-1e-5
Angular velocity	$\ \omega^2\ $	-0.05
Base velocity	$\ \mathbf{v}^2\ $	-2.0
Base orientation	$1 - \cos \theta_{\text{base}}$ $\cos \theta_{\text{base}} = \frac{\mathbf{g}_{\text{base}} \cdot \mathbf{g}_{\text{target}}}{\ \mathbf{g}_{\text{base}}\ \ \mathbf{g}_{\text{target}}\ }$	-1.5
Torso orientation	$1 - \cos \theta_{\text{torso}}$ $\cos \theta_{\text{torso}} = \frac{\mathbf{g}_{\text{torso}} \cdot \mathbf{g}_{\text{target}}}{\ \mathbf{g}_{\text{torso}}\ \ \mathbf{g}_{\text{target}}\ }$	-1.0
Stance symmetry	$\sum (q_{\text{left}}^{\text{s}} - q_{\text{right}}^{\text{s}} + q_{\text{left}}^{\text{a}} + q_{\text{right}}^{\text{a}})$ s: sagittal joints, a: anti-sagittal joints	-0.5
Ankle roll	$\sum (q_{\text{ankle,roll}} + q_{\text{ankle,roll}}^{\text{right}})$	-2.0
Feet contact	$\mathbb{1}(n_{\text{contact}} < 2)$ $+ \mathbb{1}(n_{\text{contact}} = 2 \vee n_{\text{contact}} = 0)$	-4.0
Feet orientation	$\ \mathbf{g}_{\text{xy}}^{\text{left foot}}\ + \ \mathbf{g}_{\text{xy}}^{\text{right foot}}\ $	-2.0
Negative knee DoFs	$\sum \mathbb{1}(q_{\text{knee}} < q_{\text{knee,min}})$	-1.0
Feet spread distance	$\mathbb{1}(\ \mathbf{p}_{\text{left foot}}^{\text{right foot}} - \mathbf{p}_{\text{xy}}^{\text{right foot}}\ < d_{\text{thresh}})$	-10.0
Walkings Task Rewards:		
Linear Velocity v_x	$\exp(-(v_x^{\text{cmd}} - v_x^{\text{base}})^2 / \sigma)$	2.0
Linear Velocity v_y	$\exp(-(v_y^{\text{cmd}} - v_y^{\text{base}})^2 / \sigma)$	1.5
Angular Velocity	$\exp(-(\omega_z^{\text{cmd}} - \omega_z^{\text{base}})^2 / \sigma)$	4.0

B.7 Testing Assets Details

In the paper, we have tested HERO with 20 daily objects; these objects have different sizes and weights, while being made with different materials, making it challenging to grasp with a Dex-3 hand. We list the detailed sizes, weights, materials, and language queries of all objects tested in Tab. VIII. Note that the size is roughly measured as the shape is irregular and cannot be easily described.

B.8 Testing Scenes Details

Tab. IX lists the details of the novel scenes tested in this paper, which are mainly chosen from the Coordinated Science Laboratory Studio (CSL Studio) and the Thomas M. Siebel Center for Computer Science at the University of Illinois Urbana-Champaign, Urbana, IL. The snapshot of these testing scenes can be found in Fig. 6.

TABLE VIII: **Testing objects, sizes, weights, materials, and language queries.** Sizes are roughly measured due to irregular shapes. Weights are measured with an accurate food scale.

Object	Size	Weight	Material	Language Query
<i>10 daily object evaluation.</i>				
	4.9×4.9×4.9 cm	58.06 g	Wood	orange cube
	12×6×6 cm	14.97 g	Aluminum	coke can
	8.5×7.5×7.5 cm	15.88 g	Plastic	red apple
	9×5×5 cm	137.89 g	Plastic & Metal	emergency stop button
	15×6×5 cm	239.95 g	Plastic & Metal	robot hand
	16.5×10.5×3.3 cm	185.07 g	Plastic	game cartridge
	27.5×9.5×9.5 cm	79.83 g	Plastic	olive oil bottle
	16×8.5×8.5 cm	392.81 g	Plastic & Liquid	hand soap
	23×7.9×7.9 cm	43.09 g	Paperboard & Plastic	chip can
	17×11×11 cm	73.94 g	Plush	red piranha plant
<i>10 daily objects used in 10 daily scenes evaluation.</i>				
	21×12×43 cm	215.91 g	Plastic	kettle
	16.5×11×21 cm	213.19 g	Plush	toy dog
	17.2×9×9 cm	24.95 g	Paperboard & Plastic	Starbucks coffee
	11.5×9.4×9.4 cm	526.17 g	Ceramic	orange mug
	18.5×6.3×6.3 cm	286.22 g	Plastic & Metal	water bottle
	7×7.7×7.7 cm	14.97 g	Plastic	green apple

Continued on next page

Object	Size	Weight	Material	Language Query
	18.8×2.4×13.3 cm	301.19 g	Paper	purple book
	14×8×39 cm	234.96 g	Plastic & Metal	helicopter
	8×5.6×10 cm	367.41 g	Metal & Spam	spam
	19×5×5 cm	86.18 g	Plastic & Liquid	cleaner bottle

Additional objects.

	24.5×8×6.5 cm	135.17 g	Plush	carrot
	23×8×8 cm	307.08 g	Plush	broccoli
	7.5×7.2×7.5 cm	19.05 g	Plastic	orange
	21×2.3×13.8 cm	376.03 g	Paper	book
	15.3×6.5×6.5 cm	18.14 g	Aluminum	black can

TABLE IX: **Testing scenes, table heights, language queries.** Here we list the novel scenes chosen in this paper for evaluation, and the corresponding table height. The snapshot of these scenes can be found in Fig. 6.

Scene	Location	Table Height	Language Query
corridor	CSL Studio @ UIUC	0.43m	kettle
office lounge	CSL Studio @ UIUC	0.48m	toy dog
building café	Thomas M. Siebel Center for Computer Science @ UIUC	0.72m	Starbucks coffee
office	CSL Studio @ UIUC	0.74m	orange mug
building lounge	Thomas M. Siebel Center for Computer Science @ UIUC	0.74m	water bottle
office	CSL Studio @ UIUC	0.74m	green apple
kitchenette			
building den	Thomas M. Siebel Center for Computer Science (RM 3333) @ UIUC	0.74m	purple book
robotics lab	CSL Studio @ UIUC	0.86m	helicopter
office kitchen	CSL Studio @ UIUC	0.87m	spam
classroom	Thomas M. Siebel Center for Computer Science (RM 1302) @ UIUC	0.92m	cleaner bottle