

HAIC: Humanoid Agile Object Interaction Control via Dynamics-Aware World Model

Dongting Li^{1,4*}, Xingyu Chen^{2,4*}, Qianyang Wu⁴, Bo Chen⁴, Sikai Wu⁴, Hanyu Wu³, Guoyao Zhang⁴

Liang Li⁴, Mingliang Zhou⁴, Diyun Xiang⁴, Jianzhu Ma^{1†}, Qiang Zhang^{2†}, Renjing Xu^{2†}

¹Tsinghua University ²HKUST(Guangzhou) ³ETH Zurich ⁴Xiaomi Robotics Lab

*Equal Contribution [†]Equal Advising

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




Fig. 1: Our proposed framework HAIC enables a robot to perform complex interactions, including (a) **Underactuated Object Interaction**. The robot learns interaction skills such as skateboarding, cart pushing, and cart pulling. (b) **Sequential Interaction**. HAIC supports multi-object interaction, enabling a whole-body controller to pick up the box, load it onto the cart, then drive them forward in one policy. (c) **Multi-terrain Interaction**. HAIC enables robots to maintain interactions with both an object and the scene, such as carrying a box across various types of terrains, such as platforms, slopes, and stairs.

Abstract—Humanoid robots exhibit significant potential for executing complex whole-body interaction tasks in unstructured environments. While recent advancements in Human-Object Interaction (HOI) have been substantial, prevailing methodologies predominantly address the manipulation of fully actuated objects, where the target is rigidly coupled to the robot’s end-effector and its state is strictly constrained by the robot’s kinematics. This paradigm neglects the pervasive class of underactuated objects characterized by independent dynamics and non-holonomic constraints, which pose significant control challenges due to complex coupling forces and frequent visual occlusions. To bridge this gap, we propose HAIC, a unified framework designed to enable robust interaction across a spectrum of object dynamics without reliance on external state estimation. Central to our approach is a novel dynamics predictor that infers high-order object states, specifically velocity and acceleration, solely from proprioceptive history. These predictions are explicitly projected onto static

geometric priors to construct a spatially grounded representation of dynamic occupancy, allowing the policy to internalize collision boundaries and contact affordances in visual blind spots. We employ an asymmetric fine-tuning strategy where the world model continuously adapts to the student policy’s exploration, ensuring robust state estimation under distribution shifts. We evaluate our framework on a Unitree G1 humanoid robot. Empirical results demonstrate that HAIC achieves high success rates in agile object interactions, including skateboarding, cart pushing, and cart pulling under various weight load conditions, by compensating for inertial physical perturbations. Furthermore, by accurately predicting multi-object dynamics, it seamlessly masters sequential tasks and carries a box across composed terrains.

I. INTRODUCTION

Humanoid robots, characterized by their anthropomorphic morphology, possess immense potential for executing complex




Fig. 2: HAIC excels at complex interactions, particularly with underactuated objects, and significantly outperforms the baseline.

whole-body interaction tasks in unstructured environments. From transporting payloads to interacting dynamically with tools, these tasks require not only robust locomotion stability but also the capacity to perform precise interactions with the environment. Recent data-driven approaches in Human-Object Interaction have achieved significant milestones, enabling robots to acquire diverse interaction skills [16, 6, 37, 59].

However, prevailing research predominantly focuses on the Interaction of **fully actuated objects**, such as carrying a rigidly grasped box [6, 37]. In these scenarios, the object’s state is strictly constrained by the robot’s end-effector, resulting in relatively simple and fully controlled dynamics. In contrast, the real world is populated with **underactuated objects** possessing independent dynamics and non-holonomic constraints, such as skateboards, carts, or luggage. Manipulating such objects presents a dual challenge. First, **visual occlusion** is frequent. Large objects often obstruct onboard sensors, depriving the robot of direct observation of both the terrain and the object itself [26]. Second, **dynamics coupling** becomes critical. Unlike static payloads, underactuated objects generate independent acceleration and inertial forces that propagate to the robot. Traditional control methods often rely on external high-precision state estimation to handle these disturbances, which is impractical for real-world deployment [51]. Meanwhile, existing reinforcement learning policies, lacking explicit modeling of high-order dynamics, often suffer from response lag, leading to instability during high-speed interactions [18, 33, 47].

Humans exhibit remarkable adaptability when facing similar constraints. Even when vision is obstructed by a carried object or when balancing on an accelerating skateboard, humans rely on proprioception and an internal mental model to infer the object’s state. This capability stems from a fusion of geometric reasoning (knowing the object’s boundaries) and dynamic sensing (feeling the forces to anticipate acceleration), enabling proactive center-of-mass adjustments.

Inspired by this cognitive mechanism, we introduce HAIC, a unified learning framework designed to enable robust and agile object interaction in visual blind spots. Unlike prior

models focused on visual prediction [10, 11, 13, 9, 39], HAIC is proprioception-centric and addresses the control of both fully actuated and underactuated objects through two key mechanisms in a dynamics-aware world model (DWM):

- **High-Order Dynamics Prediction:** By analyzing proprioceptive history, HAIC explicitly predicts the object’s linear and angular accelerations. This allows the robot to capture the independent motion trends of underactuated objects, facilitating proactive inertia compensation.
- **Explicit Geometric Projection:** HAIC explicitly project these predicted dynamic states onto a static geometric prior (point cloud) of the object. This constructs a spatially grounded representation of dynamic occupancy, enabling the policy to perceive collision boundaries and contact affordances even without visual feedback.

To bridge the gap between simulation and reality, we employ an asymmetric fine-tuning strategy where the world model continuously adapts to the student policy’s exploration online, effectively mitigating distribution shifts [46, 57].

The main contributions of this work are summarized as follows:

- **Dynamics-Aware World Model:** We propose a novel dynamics predictor that integrates high-order acceleration inference with explicit geometric projection, unifying state estimation for both static payloads and dynamic underactuated objects.
- **Asymmetric Adaptive Distillation:** We design a robust two-stage training pipeline that transfers privileged geometric and dynamic reasoning into a sensor-limited student policy via continuous world model adaptation.
- **Agile Real-World Performance:** We demonstrate SOTA performance on a Unitree G1 humanoid robot. Our framework achieves a 100% success rate in blind multi-terrain box carrying and, for the first time without external sensing, enables the robot to master agile underactuated tasks, such as dynamic skateboarding.

II. RELATED WORK

A. Humanoid Whole-Body Control and Tracking

Deep reinforcement learning (RL) [49] has achieved significant milestones in tracking large-scale human motion capture (MoCap) datasets like AMASS [41]. Recent General Motion Trackers (GMT) [5, 71] and retargeting frameworks [3, 34, 16] allow humanoids to reproduce diverse kinematic behaviors. To support these data-driven approaches, scalable collection systems like TWIST [74] and TWIST2 [75] have been developed to provide high-quality whole-body demonstrations.

Early studies demonstrated that motion imitation and adversarial motion priors enable policies to track diverse human motion data, producing natural and robust whole-body control [44, 45, 65, 34, 14, 17]. These works establish motion tracking as a scalable paradigm for acquiring rich motor skills while substantially reducing manual reward engineering.

Building upon this foundation, subsequent research [75, 12, 73, 76] investigates generalized whole-body control through

large-scale motion datasets and unified latent representations, enabling a single policy to compose heterogeneous skills and adapt to diverse task conditions. Policies are pretrained via large-scale data and unsupervised reinforcement learning, and conditioned on language, vision, or high-level task descriptors to enable controllable and zero-shot skill adaptation [31, 40, 22, 50]. These models demonstrate strong generalization across diverse motion patterns, highlighting the potential of foundation training for whole-body control and advanced interactions with environments, especially dynamic objects.

B. Physics-Based Human-Object Interaction

Humanoid-object interaction (HOI) is a longstanding challenge bridging computer vision and robotics. Large-scale HOI datasets provide diverse demonstrations to facilitate data-driven learning [29, 55, 4, 21, 77, 66]. Based on these data, advanced frameworks [60, 56, 61, 43, 67, 35] have been developed to train policies that master versatile whole-body interactions in simulation. Similarly, video-based imitation methods extract interactive skills directly from human videos [63, 2, 72]. Prior studies on loco-manipulation, which focus only on the interaction between the hand and objects, show that integrating force adaptation [78], hierarchical control [32, 23, 28], or residual action [54] enables precise manipulation behaviors with stable locomotion. And visual sim-to-real transfer methods substantially improve real-world generalization for continuous manipulation tasks [69, 15, 20]. Some specialized systems demonstrate high-speed adaptation for specific sports but lack generality [52, 68, 36, 62]. As for full-body interactions, Omniretarget [70] improves the quality of retargeted motion by constructing the interaction mesh. HDMI [63] and ResMimic [80] incorporate residual action modules to adapt policies to virtual object states. Other methods typically rely on explicit object tracking to estimate object states, such as PhysHSI (LiDAR) [58], HDMI (markers) [63], VisualMimic (depth) [72], and DemoHLM (RGB) [8]. However, all these methods are restricted to interactions with static environments or fully actuated objects. In contrast, HAIC can learn coupled humanoid-object dynamics through a specialized world model, enabling agile interactions with underactuated objects or multi-terrain environments.

C. World Models for Humanoids

World models [10, 11] act as internal simulators that enable agents to anticipate future outcomes, and have been widely adopted in robotic manipulation to predict tactile feedback and object dynamics [7, 19, 51]. Recent work leverages world models to address high-dimensional humanoid dynamics and complex contact sequences. DWL [9] reconstructs complete physical states from noisy observations via denoising learning, enabling robust whole-body control, while RWM [27] mitigates long-horizon drift through a dual-autoregressive prediction scheme. WMR [53] decouples dynamics learning from policy gradients and enforces state reconstruction to prevent representation collapse. World models are also used as self-supervised objectives for perceptual robustness, as




Fig. 3: Overview of our Dynamics-aware World Model. It predicts object dynamics from proprioception and explicitly projects them onto a static geometric prior for reconstructing the privileged information.

demonstrated by WMP [25], and extended to active interaction in Ego-VCP [38] for ego-centric contact planning. Any2Track [79] learns dynamics-aware embeddings from observation histories, enabling implicit inference of environmental properties and robust sim-to-real transfer. However, these methods primarily model proprioceptive dynamics for whole-body control, often neglecting the coupled dynamics inherent in object interaction.

III. METHODOLOGY

A. Problem Formulation

The purpose of humanoid-object interaction is to train a whole-body control policy that simultaneously tracks the reference trajectory of both the robot and the object. We formulate this proprioception-based interaction task as a partially observable Markov decision process (POMDP) [26]. At each timestep t , the full state $s_t \in \mathcal{S}$ includes the robot's proprioception s_t^{prop} , the reference motion s_t^{ref} , the point cloud of objects $\mathcal{P}$ and the privileged state s_t^{priv} .

Since the object state s_t^{obj} cannot be acquired directly, we adapt the teacher-student framework [24] to distill the privileged information. In training, the teacher has access to the full state s_t . During deployment, the student only observes:

- Onboard proprioception s_t^{prop} (joint data, IMU data).
- Reference motion s_t^{ref} .
- Nominal geometric template $\mathcal{P} \in \mathbb{R}^{N \times 3}$: a standard-sized point cloud acting as a semantic prior.

The goal is to explore a policy $\pi(a_t | s_t^{prop}, s_t^{ref}, a_{t-1}, \mathcal{P})$ in the simulation that enables the humanoid robot to acquire robust interaction skills with objects in the real world.

B. Dynamics-Aware World Model

Our motivation is to use the robot's state s_t^{prop} to predict the inaccessible state of objects s_t^{obj} [51]. Therefore, we design a predictive framework to internalize the object's dynamics relative to the robot. And we propose a pipeline that first estimates object dynamics, explicitly projects these dynamics onto the geometric prior, and subsequently feeds them into the downstream network, as shown in Fig. 3.

1) *Object Adapter*: The core component of the dynamics-aware world model is the object adapter, which functions as a real-time dynamics predictor. It takes the history of predicted object state $\hat{s}_{t-1}^{obj}$, robot observations $\mathcal{H}_t = \{s_{t-n}^{prop}, \dots, s_t^{prop}\}$ and the future of reference motion $\mathcal{F}_t = \{s_t^{ref}, \dots, s_{t+m}^{ref}\}$ as input and predicts the object's relative spatial and temporal state $\hat{s}_t^{obj}$:

$$\hat{s}_t^{obj} = f_{OA}(\hat{s}_{t-1}^{obj}, \mathcal{H}_t, \mathcal{F}_t) \quad (1)$$




Fig. 4: Framework overview. We train policies in the simulation from scratch. The framework includes a privileged teacher and a dynamics-aware student. The student policy utilizes the learned world model to perform robust interaction tasks such as skateboarding on a real humanoid robot.

where $\hat{s}_t^{obj}$ includes the predicted relative position $\hat{p}_t$, orientation $\hat{R}_t$, linear/angular velocities ($\hat{v}_t^{lin}, \hat{v}_t^{ang}$), and linear/angular accelerations ($\hat{a}_t^{lin}, \hat{a}_t^{ang}$).

2) *Explicit Geometric Projection*: While dynamics-aware approaches using signed distance fields (SDFs) or point clouds have been explored for stationary arms [30, 19] and HOI animation [1, 64], the application to the high-dimensional dynamics of humanoid whole-body control remains limited. By conditioning our predictor on a point-cloud-based geometric prior, the robot can mentally track an invisible object, effectively transforming a partially observable Markov decision process into a manageable control task without the need for the real-time privileged states.

To bridge the gap between abstract state estimation and physical interaction, we introduce an explicit geometric projection step. We transform the canonical point cloud $\mathcal{P}$ using the predicted $\hat{s}_t^{obj}$ into the robot’s root coordinate frame:

$$\mathcal{P}_t^{dyn} = \hat{R}_t \cdot \mathcal{P} + \hat{p}_t \quad (2)$$

By explicitly generating the dynamic object point cloud $\mathcal{P}_t^{dyn}$, we provide the privilege adapter with a spatially grounded representation of the object’s occupancy. Note that $\mathcal{P}$ serves as a geometric prior rather than a ground-truth measurement.

3) *Privilege Adapter*: Finally, the estimated dynamic geometry is fused with proprioception via the privilege adapter. This module takes the dynamic point cloud $\mathcal{P}_t^{dyn}$, current proprioception s_t^{prop} , and the reference motion s_t^{ref} to produce a predicted privileged feature $\hat{z}_t^{priv}$. This feature serves as the input to the student actor, effectively mimicking the information flow of a privileged actor.

C. Two-Stage Training Framework

Fig. 4 illustrates the overall architecture of HAIC. To streamline our formulation, we define the unified world model loss ($\mathcal{L}_{WM}$) as the weighted sum of dynamics prediction error and latent feature alignment:

$$\mathcal{L}_{WM} = \lambda_{obj} \|\hat{s}_t^{obj} - s_t^{obj}\|^2 + \lambda_{priv} \|\hat{z}_t^{priv} - z_t^{priv}\|^2 \quad (3)$$

where λ_{obj} and λ_{priv} balance the supervision strength of physical state regression and latent space distillation, respectively. We denote the standard PPO [49] objective (including policy gradient, value loss, and entropy) as $\mathcal{L}_{PPO}(\pi, V)$.

1) *Stage 1: Feature Alignment & Warm-up*: In the first stage, we train the privileged Teacher actor π_T via reinforcement learning. Simultaneously, we train the world model and the student actor π_S to mirror the teacher’s behavior and internal representation [48]. This “warm-up” phase ensures that the student starts Stage 2 with a learned sensorimotor primitive closer to the expert distribution. The objective is:

$$\mathcal{L}_{Stage1} = \mathcal{L}_{PPO}(\pi_T, V_T) + \mathcal{L}_{WM} + \lambda_{\text{distill}} D_{KL}(\pi_S || \pi_T) \quad (4)$$

Here, the action distillation term λ_{distill} forces the student to produce actions consistent with the teacher.

2) *Stage 2: Asymmetric Fine-tuning*: In the second stage, we fine-tune the student actor π_S using an asymmetric actor-critic setup [46]. The critic V_T retains access to privileged states for variance reduction.

To address the distribution shift caused by the student’s independent exploration, we continue to optimize the world model alongside the policy. The total objective is:

$$\mathcal{L}_{Stage2} = \mathcal{L}_{PPO}(\pi_S, V_T) + \mathcal{L}_{WM} \quad (5)$$

By consistently minimizing $\mathcal{L}_{WM}$ in both stages, we ensure the world model continuously adapts to the evolving trajectory distribution of the student policy, preventing the “blind” policy from collapsing due to inaccurate state estimation.

Stabilization via EMA: To further mitigate instability arising from the coupled training of the policy and the world model, we maintain an exponential moving average (EMA) [42] of the world model weights for the rollout policy. The policy interacts with the environment using the EMA-smoothed predictions, while gradients are backpropagated through the online network.




Fig. 5: Multiple Objects Contact Guidance Strategy.




Fig. 6: Real-world performance comparison with the baseline across various tasks. With the specifically designed dynamics-aware world model, HAIC maintains robust stability throughout the interaction, whereas the baseline suffers from balance failures and trajectory drift.

D. Multiple Objects Contact Reward

To enable robust manipulation of multiple objects, we introduce a contact reward r_{contact} that unifies geometric alignment and force regulation. As illustrated in Fig. 5, this reward dynamically guides the robot through two phases:

- Guidance (Fig. 5a): Before contact, the reward minimizes the distance between end-effectors and semantic targets.
- Execution (Fig. 5b-c): Upon contact, it applies force appropriately while maintaining geometric constraints.

The contact reward is calculated as the average reward across all active objects $o \in \mathcal{O}$ and their corresponding end-effectors $e \in \mathcal{E}_o$:

$$r_{\text{contact}} = \frac{1}{|\mathcal{O}|} \sum_{o \in \mathcal{O}} \left(\frac{1}{|\mathcal{E}_o|} \sum_{e \in \mathcal{E}_o} \mathbb{I}_{o,e} \cdot r_{\text{pos}}^{o,e} \cdot r_{\text{force}}^{o,e} \right). \quad (6)$$

Here, $\mathbb{I}_{o,e}$ is a binary mask from the reference motion indicating if contact is active. The individual terms account for physical tolerances to prevent optimization instability. Position term $r_{\text{pos}}^{o,e}$ encourages the end-effector position $\mathbf{p}_e$ to match the target $\mathbf{p}_{\text{tgt}}$ within a tolerance ϵ_{tol} :

$$r_{\text{pos}}^{o,e} = \exp(-\max(0, \|\mathbf{p}_e - \mathbf{p}_{\text{tgt}}\| - \epsilon_{\text{tol}}) / \sigma_p). \quad (7)$$

Force term $r_{\text{force}}^{o,e}$ ensures the contact force $\mathbf{F}_e$ reaches a required threshold F_{thr} . We penalize the force deficit only when it falls below this threshold:

$$r_{\text{force}}^{o,e} = \exp(-\max(0, F_{\text{thr}} - \|\mathbf{F}_e\|) / \sigma_f). \quad (8)$$

More detailed reward definitions and parameters are provided in Appendix C2.

IV. EXPERIMENTS

A. Experiment Setup

1) *Task Description:* To validate our method, we designed three types of interaction tasks:

Underactuated Object Interaction. This task requires the robot to interact with underactuated wheeled objects, which introduces unstructured, time-varying external physical disturbances upon contact. Unlike interactions with fully actuated objects, the motion of underactuated objects actively affects the robot's balance, requiring the controller to predict and compensate for these disturbances in real time.

Sequential interaction. In addition to interacting with a single object, we also design composite sequential tasks involving multiple objects, such as loading a payload onto a mobile cart and then manipulating the cart. These scenarios challenge the policy's capacity to handle discrete mode switching and maintain stable control over extended horizons while coordinating interactions with multiple dynamic bodies.

Multi-terrain interaction. The final task requires the robot to carry a box across composite terrains, including platforms, slopes, and stairs. This task places high demands on state estimation because the box obscures the robot's view of the ground. As a result, the policy relies solely on self-perception and predictions from the world model. The controller must simultaneously handle the dynamic load on the arms and the complex interactions with the terrain.

2) *Implementation Details:* We describe the specific configurations for data processing, training infrastructure, and deployment setups below.

Data Collection and Processing. To capture high-fidelity interaction behaviors, we utilized an optical motion capture system tracking both the human motion and interactive objects via synchronized markers. The raw kinematic data were retargeted to the Unitree G1 and refined through a physics-based simulation pipeline to generate physically valid states with contact information (see Appendix B for details).

Training Setup. We trained the base policy and world model using PPO in the Isaac Sim simulator. The two-stage training of a model was conducted on a single NVIDIA RTX 4090 GPU and took approximately 8 hours.

Sim-to-Real Transfer. We employed extensive domain randomization to bridge the reality gap. For real-world deployment, the policy ran on an external PC, transmitting joint position targets to the robot’s onboard control board via Ethernet. The PD controller operated at 200 Hz in simulation and 500 Hz on the real robot to ensure accurate tracking of the PD targets, while the control policies ran at 50 Hz.

3) *Evaluation Metrics:* To rigorously quantify the robot’s kinematic tracking performance, we define the following metrics based on onboard proprioception. We report the Mean and Standard Deviation for error metrics, and the Success Rate (SR) for task completion.

Robot State Metrics:

- $E_{\text{mpbpe}} / E_{\text{mpboe}}$: Root-relative Mean Per Body Position (mm) and Orientation (10^{-3} rad) Error.
- $E_{\text{g-mpbpe}}$: Global Mean Per Body Position Error (mm).
- $E_{\text{mpjpe}} / E_{\text{mpjve}}$: Mean Per Joint Position (10^{-3} rad) and Velocity (10^{-3} rad/frame) Error.
- $E_{\text{mpbve}} / E_{\text{mpbae}}$: Mean Per Body Velocity (mm/frame) and Acceleration (mm/frame²) Error.

Object State Metrics:

- $E_{\text{mpope}} / E_{\text{mpoae}}$: Mean Per Object Position (mm) and Orientation (10^{-3} rad) Error.
- $E_{\text{mpove}} / E_{\text{mpoae}}$: Mean Per Object Velocity (mm/frame) and Acceleration (mm/frame²) Error.

We evaluate the performance of our proposed framework on the Unitree G1 humanoid robot in real-world scenarios. The robot operates autonomously using onboard computation, without external motion capture systems, relying solely on proprioceptive feedback and our dynamics-aware World Model.

4) *Baseline:* We compare our method with HDMI in simulation and the real world. We utilize **HDMI***, a proprioception-only variant of the original HDMI framework [63] as its lower bound. As illustrated in Fig. 2b, we also present the sim-to-sim upper bound performance of HDMI for comparison with the real-world success rates of HAIC and HDMI*. To ensure a fair comparison, this baseline was retrained from scratch using the same setup as our method. Instead of privileged object states, HDMI* receives a history of proprioceptive observations to implicitly infer the environmental context.

B. Real-world Deployment and Comparison

In real-world experiments, our dynamics-aware world model functions as a core estimator across three distinct scenarios:

(a) **Underactuated Object Interaction.** For underactuated objects like skateboards and carts, the world model infers the object’s full dynamics, including translation, rotation, linear/angular velocity, and acceleration. These high-order predictions are explicitly applied to the object’s canonical point cloud. (b) **Sequential Interaction.** In sequential tasks involving multiple objects, HAIC predicts the dynamics for both the box and the cart. It applies the respective predicted rotations and translations to align each object’s geometric prior independently. (c) **Multi-terrain Interaction.** HAIC extends its prediction to environmental features. It estimates the dynamics of the carried box while simultaneously predicting the translation and orientation of terrain structures. By projecting these estimates onto the geometry and feeding the aligned features into the downstream network, the robot successfully navigates uneven ground. We compared HAIC with the baseline in these scenarios. HAIC exhibits superior robustness due to its structured representation incorporating both geometric and dynamic information. The original HDMI performed well with fully-driven objects, but its performance declined on underactuated objects, since it does not account for geometry and higher-order physical quantities. As demonstrated by the case study results below, HAIC shows significant advantages in both success rate and motion quality.

TABLE I: Real-world results for **Skateboarding**. We report the success rate for both Gliding and Completing. HAIC achieves stable gliding and significantly higher completion rates compared to the baseline.

Method	Success Rate		Robot State Metrics					
	Glide	Comp.	E_{mpbpe}	E_{mpboe}	E_{mpjpe}	E_{mpjve}	E_{mpbae}	E_{mpjve}
HDMI*	20%	0%	132.1 $\pm$ 35.7	666.7 $\pm$ 213.6	188.7 $\pm$ 12.6	5.39 $\pm$ 0.62	4.06 $\pm$ 0.04	10.40 $\pm$ 0.19
HAIC	100%	60%	81.5 $\pm$ 15.4	453.5 $\pm$ 88.2	127.1 $\pm$ 11.5	4.72 $\pm$ 0.75	4.48 $\pm$ 0.50	12.15 $\pm$ 1.87

TABLE II: Real-world results for **Cart Manipulation**. HAIC achieves 100% success in both tasks, whereas the baseline fails in “Push”.

Task	Method	SR $\uparrow$	$E_{\text{mpbpe}} \downarrow$	$E_{\text{mpboe}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{mpjve}} \downarrow$	$E_{\text{mpbae}} \downarrow$	$E_{\text{mpjve}} \downarrow$
Pull	HDMI*	40%	118.0 $\pm$ 5.8	556.7 $\pm$ 43.0	165.0 $\pm$ 2.3	6.11 $\pm$ 0.13	3.73 $\pm$ 0.14	10.32 $\pm$ 1.00
Cart	HAIC	100%	76.2 $\pm$ 6.8	425.0 $\pm$ 41.7	132.3 $\pm$ 6.4	5.08 $\pm$ 0.53	3.66 $\pm$ 0.21	9.50 $\pm$ 1.10
Push	HDMI*	0%	136.2 $\pm$ 23.3	778.5 $\pm$ 218.6	158.9 $\pm$ 4.7	7.36 $\pm$ 0.48	5.32 $\pm$ 0.21	15.62 $\pm$ 0.74
Cart	HAIC	100%	82.0 $\pm$ 6.3	418.7 $\pm$ 36.6	150.0 $\pm$ 9.3	5.29 $\pm$ 0.17	4.47 $\pm$ 0.09	11.61 $\pm$ 0.80

TABLE III: Real-world results for **Sequential Interaction**. The baseline is omitted as it failed to complete the prerequisite cart tasks. HAIC demonstrates capability in sequential object transport.

Task	Method	SR $\uparrow$	$E_{\text{mpbpe}} \downarrow$	$E_{\text{mpboe}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{mpjve}} \downarrow$	$E_{\text{mpbae}} \downarrow$	$E_{\text{mpjve}} \downarrow$
Pull Cart	HDMI*	0%						
w/ Box	HAIC	40%	88.2 $\pm$ 8.2	420.7 $\pm$ 16.7	157.2 $\pm$ 8.2	5.91 $\pm$ 0.07	5.02 $\pm$ 0.09	12.20 $\pm$ 0.31
Push Cart	HDMI*	0%						
w/ Box	HAIC	100%	70.5 $\pm$ 5.7	381.7 $\pm$ 24.7	143.6 $\pm$ 4.1	5.27 $\pm$ 0.26	3.92 $\pm$ 0.18	11.10 $\pm$ 0.85

TABLE IV: Real-world results for **Multi-terrain Interaction**. HAIC succeeds on the complex slope-stair terrain where baseline fails.

Task	Method	SR $\uparrow$	$E_{\text{mpbpe}} \downarrow$	$E_{\text{mpboe}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{mpjve}} \downarrow$	$E_{\text{mpbae}} \downarrow$	$E_{\text{mpjve}} \downarrow$
Carry	HDMI*	100%	66.6 $\pm$ 2.0	344.0 $\pm$ 19.8	171.7 $\pm$ 1.7	4.69 $\pm$ 0.10	3.14 $\pm$ 0.01	8.10 $\pm$ 0.13
Box	HAIC	100%	58.8 $\pm$ 6.1	305.6 $\pm$ 20.8	162.0 $\pm$ 6.8	4.21 $\pm$ 0.13	3.03 $\pm$ 0.05	7.75 $\pm$ 0.18
w/ Stair	HDMI*	100%	66.2 $\pm$ 4.1	316.2 $\pm$ 2.8	158.9 $\pm$ 13.8	4.79 $\pm$ 0.36	4.35 $\pm$ 0.24	10.10 $\pm$ 0.98
	HAIC	100%	55.5 $\pm$ 3.9	331.2 $\pm$ 23.1	118.4 $\pm$ 0.6	4.06 $\pm$ 0.02	4.05 $\pm$ 0.01	8.66 $\pm$ 0.01
w/ Slope	HDMI*	100%	53.5 $\pm$ 0.7	289.0 $\pm$ 1.7	130.6 $\pm$ 1.9	3.68 $\pm$ 0.03	3.39 $\pm$ 0.01	7.41 $\pm$ 0.02
	HAIC	100%	60.5 $\pm$ 4.3	324.7 $\pm$ 28.0	120.2 $\pm$ 1.6	8.77 $\pm$ 0.11	4.60 $\pm$ 0.13	3.66 $\pm$ 0.05
w/ Stair + Slope	HDMI*	0%	126.8 $\pm$ 4.4	701.0 $\pm$ 32.1	162.3 $\pm$ 2.0	5.98 $\pm$ 0.07	4.03 $\pm$ 0.10	10.87 $\pm$ 0.27
	HAIC	100%	60.6 $\pm$ 13.4	286.5 $\pm$ 59.8	128.5 $\pm$ 16.2	4.74 $\pm$ 0.17	4.05 $\pm$ 0.04	10.67 $\pm$ 0.52

1) *Underactuated Object Interaction:* This task involves high-dynamic interactions, where object interactions significantly affect the robot’s balance: skateboarding and manipulating wheeled carts. Fig. 6 demonstrates that HAIC maintains

TABLE V: Ablation results for **Skateboarding**. Vec-Pose can balance but fails to dismount due to a lack of acceleration prediction. HAIC achieves a high completion rate with the lowest robot tracking errors and the most stable object acceleration control.

Method	SR $\uparrow$		Robot State Metrics							Object State Metrics			
	Glide	Comp.	$E_{mpbpe}\downarrow$	$E_{g-mpbpe}\downarrow$	$E_{mpboe}\downarrow$	$E_{mpjpe}\downarrow$	$E_{mpbve}\downarrow$	$E_{mpbae}\downarrow$	$E_{mpjve}\downarrow$	$E_{mpope}\downarrow$	$E_{mpoe}\downarrow$	$E_{mpove}\downarrow$	$E_{mpoae}\downarrow$
Proprio	40%	0%	91.1 $\pm$ 9.6	548.1 $\pm$ 186.4	490.6 $\pm$ 58.0	163.3 $\pm$ 3.4	5.08 $\pm$ 0.28	4.18 $\pm$ 0.10	12.45 $\pm$ 0.55	549.9 $\pm$ 167.5	1219 $\pm$ 179	9.63 $\pm$ 0.64	11.02 $\pm$ 0.25
Vec-Pose	100%	0%	77.1 $\pm$ 12.7	430.3 $\pm$ 248.2	394.0 $\pm$ 55.0	167.0 $\pm$ 4.7	4.25$\pm$0.20	3.94$\pm$0.09	10.14$\pm$0.10	434.8 $\pm$ 259.9	1192 $\pm$ 147	8.78 $\pm$ 1.25	10.75 $\pm$ 0.10
Vec-Dyn	100%	100%	74.6 $\pm$ 0.8	444.7 $\pm$ 13.3	416.7 $\pm$ 6.8	147.7 $\pm$ 2.3	4.87 $\pm$ 0.16	4.11 $\pm$ 0.16	12.20 $\pm$ 0.56	432.1$\pm$9.1	1038 $\pm$ 137	8.64 $\pm$ 0.36	10.90 $\pm$ 0.28
Geo-Pose	100%	100%	60.2 $\pm$ 2.5	513.0 $\pm$ 86.2	354.5 $\pm$ 18.3	121.6 $\pm$ 1.8	4.55 $\pm$ 0.08	4.01 $\pm$ 0.04	12.42 $\pm$ 0.22	565.8 $\pm$ 83.2	1071 $\pm$ 72	8.77 $\pm$ 0.98	10.74 $\pm$ 0.11
HAIC	100%	100%	57.8$\pm$2.8	417.8$\pm$36.4	335.5$\pm$19.7	117.6$\pm$0.7	4.54 $\pm$ 0.11	4.04 $\pm$ 0.05	12.30 $\pm$ 0.31	473.8 $\pm$ 43.1	1035$\pm$57	8.08$\pm$0.32	10.66$\pm$0.01

TABLE VI: Ablation results for **Cart Manipulation**. HAIC shows consistently high performance, achieving either the best or second-best results across most metrics.

Task	Method	SR $\uparrow$	Robot State Metrics							Object State Metrics			
			$E_{mpbpe}\downarrow$	$E_{g-mpbpe}\downarrow$	$E_{mpboe}\downarrow$	$E_{mpjpe}\downarrow$	$E_{mpbve}\downarrow$	$E_{mpbae}\downarrow$	$E_{mpjve}\downarrow$	$E_{mpope}\downarrow$	$E_{mpoe}\downarrow$	$E_{mpove}\downarrow$	$E_{mpoae}\downarrow$
Pull Cart	Proprio	60%	97.9 $\pm$ 34.7	341.9 $\pm$ 114.6	575.4 $\pm$ 188	137.4 $\pm$ 6.5	5.58 $\pm$ 0.73	3.75 $\pm$ 0.22	11.39 $\pm$ 1.69	448.8 $\pm$ 190	1551 $\pm$ 149	5.90 $\pm$ 2.41	3.59 $\pm$ 0.06
	Vec-Pose	60%	100.0 $\pm$ 26.6	388.0 $\pm$ 221.0	589.6 $\pm$ 141	145.3 $\pm$ 13.5	5.86 $\pm$ 0.74	3.83 $\pm$ 0.31	11.96 $\pm$ 2.06	437.7 $\pm$ 223	1530 $\pm$ 212	5.78 $\pm$ 1.96	3.55 $\pm$ 0.03
	Vec-Dyn	100%	71.3 $\pm$ 0.5	244.1 $\pm$ 22.3	406.9 $\pm$ 9.3	133.7 $\pm$ 1.0	5.12 $\pm$ 0.04	3.54 $\pm$ 0.01	9.88 $\pm$ 0.18	230.1 $\pm$ 16.0	1640 $\pm$ 108	4.05$\pm$0.02	3.52$\pm$0.01
	Geo-Pose	100%	66.1 $\pm$ 0.3	261.5 $\pm$ 32.7	384.1 $\pm$ 21.7	127.3 $\pm$ 4.4	5.29 $\pm$ 0.66	3.65 $\pm$ 0.19	10.85 $\pm$ 1.52	320.3 $\pm$ 89.5	1210$\pm$11.6	4.09 $\pm$ 0.11	3.57 $\pm$ 0.04
	HAIC	100%	64.8$\pm$2.8	225.7$\pm$3.4	372.7$\pm$34.1	123.1$\pm$0.3	4.75$\pm$0.01	3.48$\pm$0.01	9.39$\pm$0.09	219.5$\pm$6.6	1241 $\pm$ 247	4.19 $\pm$ 0.02	3.55 $\pm$ 0.01
Push Cart	Proprio	0%	112.9 $\pm$ 9.2	276.6 $\pm$ 80.5	538.4 $\pm$ 129.2	172.1 $\pm$ 14.2	6.50 $\pm$ 0.41	4.78 $\pm$ 0.14	15.38 $\pm$ 0.85	460.7 $\pm$ 286.2	1406 $\pm$ 16	4.61 $\pm$ 0.91	3.75 $\pm$ 0.04
	Vec-Pose	100%	61.5 $\pm$ 1.8	451.8 $\pm$ 82.9	287.7$\pm$3.1	125.3$\pm$2.0	4.65 $\pm$ 0.17	4.16 $\pm$ 0.10	11.20 $\pm$ 0.91	338.7 $\pm$ 63.9	1152$\pm$58	4.25 $\pm$ 0.43	3.71 $\pm$ 0.01
	Vec-Dyn	100%	62.3 $\pm$ 2.5	439.0 $\pm$ 75.7	299.9 $\pm$ 16.3	126.4 $\pm$ 2.2	4.58$\pm$0.05	4.12$\pm$0.04	10.86$\pm$0.39	358.4 $\pm$ 73.1	1172 $\pm$ 26	4.63 $\pm$ 1.03	3.70$\pm$0.01
	Geo-Pose	100%	64.0 $\pm$ 4.8	249.1 $\pm$ 43.3	346.8 $\pm$ 28.5	141.0 $\pm$ 3.1	4.93 $\pm$ 0.06	4.28 $\pm$ 0.03	12.70 $\pm$ 0.16	259.8$\pm$66.8	1153 $\pm$ 107	3.65$\pm$0.41	3.70$\pm$0.01
	HAIC	100%	61.3$\pm$2.1	221.0$\pm$74.3	319.9 $\pm$ 13.0	145.9 $\pm$ 5.5	5.00 $\pm$ 0.10	4.23 $\pm$ 0.08	12.36 $\pm$ 0.39	267.9 $\pm$ 85.6	1262 $\pm$ 96	4.04 $\pm$ 1.07	3.70$\pm$0.01

robust stability in skateboarding and cart tasks, whereas the baseline suffers from frequent failures and trajectory drift.

Skateboarding: Table I reports two success metrics: glide success rate (maintaining balance on the skateboard $>1s$) and complete success rate. While the baseline struggles to maintain basic balance, HAIC achieves a high glide success rate, demonstrating robust balance control. Although the complex dismount phase remains challenging, our method significantly outperforms the baseline, which fails to complete the task.

Cart Manipulation: Table II highlights the robustness of HAIC. Notably, in the “Push Cart” task, the baseline fails as it cannot account for the cart’s external dynamics and non-holonomic constraints. HAIC predicts the object’s forward velocity and acceleration, allowing the robot to proactively lean into the motion, resulting in a high success rate for both “Pull” and “Push” tasks.

2) *Sequential Interaction:* We evaluate sequential tasks requiring mode switching: “Pull Cart w/ Box” (load the box, then pull the cart) and “Push Cart w/ Box” (load the box, then push the cart). The baseline was unable to attempt these sequential composite tasks, as it cannot reliably complete a single cart manipulation task. Therefore, Table III reports only the performance of HAIC. Despite the compounded complexity and error accumulation over a longer horizon, our framework enables the robot to transition smoothly, achieving a high success rate in the “Push Cart w/ Box” task.

3) *Multi-terrain Interaction:* We evaluate the robot’s ability to carry a box across varying terrains, including flat ground, 15-degree slopes, stairs, and mixed slope-stair terrains. In these scenarios, HAIC successfully picks and places a box across various types of terrains, where the policy completely occludes the robot’s visual sensors. As shown in Table IV, HAIC demonstrates superior robustness on complex terrains. While the baseline handles simple slopes, it suffers catastrophic failures on the “w/ Slope + Stair” task due to error accumulation

in state estimation. In contrast, HAIC maintains a high success rate with lower orientation errors, validating that our geometric prior effectively compensates for the lack of visual feedback on uneven terrain.

C. Ablation Study on HAIC

To rigorously analyze the contribution of each component in our framework—specifically, the benefits of higher-order dynamics prediction and explicit geometric projection—we define five model variants for our ablation study:

- **Proprio:** A standard proprioception-only policy that observes only the robot’s joint states and reference motion, without any explicit object representation.
- **Vec-Pose:** Extends the Proprio by adding an auxiliary module that predicts the object’s position and orientation ($\hat{p}, \hat{R}$). These are fed into the privilege adapter together.
- **Vec-Dyn:** Extends Vec-Pose by predicting the comprehensive dynamics, including velocity and acceleration ($\hat{v}, \hat{a}$), which are also fed as flattened vector inputs.
- **Geo-Pose:** Predicts the same state ($\hat{p}, \hat{R}$) as Vec-Pose. The predicted pose is used to transform the canonical point cloud into the robot frame. The explicit geometric feature and predicted pose are fed to the privilege adapter.
- **HAIC:** The proposed full method. It predicts the comprehensive dynamics ($\hat{p}, \hat{R}, \hat{v}, \hat{a}$), projects the pose onto the geometric feature, and then feeds all the object’s representations into the privilege adapter.

1) *Analysis of Training Efficiency:* Fig. 8 illustrates the learning curves of all methods on the challenging “Push Cart w/ Box” task over 3,800 training steps. HAIC demonstrates superior sample efficiency and achieves the highest asymptotic return compared to all baselines. We further evaluate the variants in the MuJoCo simulator within the same setup for quantitative comparison.




Fig. 7: Sim-to-real performance across various HOI tasks, including pushing, carrying, and terrain traversal. HAIC achieves robust interactions with diverse terrains and underactuated objects, and demonstrates strong generalization across object size, terrain orientation, and load weight.




Fig. 8: Learning curves on the “Push Cart w/ Box” task.

2) *Skateboarding*: Table V presents the results for the dynamic skateboarding task, where Proprio struggles to maintain balance even during the glide phase, confirming the insufficiency of proprioception alone. Interestingly, while Vec-Pose achieves perfect gliding success, it fails during the dismount phase, indicating that position perception aids balance but lacks the foresight to handle impulsive forces upon stepping off. Comparing this with dynamics-aware methods highlights that predicting object acceleration is critical for managing transient disturbances. Consequently, HAIC achieves high success rates across both phases with the lowest tracking errors, indicating that our world model generates necessary feedforward compensation for stability.

3) *Cart Manipulation*: Table VI presents the results for cart manipulation. Proprio fails to resolve object states, while Vec-Pose and Vec-Dyn suffer from severe trajectory drift despite completing the task. In contrast, methods utilizing geometric projection (Geo-Pose and HAIC) achieve significantly tighter global tracking by strictly aligning with the cart’s steering axis. Ultimately, HAIC delivers the best overall performance by combining this geometric alignment with dynamics prediction for precise and stable manipulation.

4) *Other Complex Scenarios*: We further extend the evaluation to sequential and multi-terrain tasks. Detailed metrics are provided in Appendix Table XIV and Table XV. Experimental results highlight two key advantages of HAIC:

- **Sequential Stability**: In sequential tasks like “Pull Cart w/ Box”, baselines fail completely due to error accumulation. HAIC leverages its world model to mitigate drift, thereby ensuring high success rates and precise object steering across extended interaction sequences.

- **Terrain Adaptability**: On the challenging “Slope + Stair” terrain, Proprio and Vec-Pose suffer from stability issues. HAIC maintains a high success rate by predicting dynamic disturbances caused by uneven ground.

D. Robust interaction and Generalization

To evaluate the policy’s robustness and generalization capabilities when encountering unseen object attributes and environmental changes, we designed a series of rigorous variations, including object size, terrain rotation, and loaded weight. The results are summarized in Fig. 7.

1) *Object Size and Weight Generalization*: We tested beyond the training range (0.1m–0.2m, 1kg–2kg). The input geometric prior $\mathcal{P}$ remained fixed as the original standard template. The policy successfully transported a large box with a side length of 0.25m and a heavy box weighing 3kg. For heavy objects, the policy exhibited a lower center of gravity and a slower movement speed. For large-sized objects, the policy adjusted the grasping position and arm posture to avoid self-collisions. This indicates that the world model learned general relationships between mass, inertia, and motion dynamics.

2) *Terrain Orientation Perturbations*: We randomly rotated slopes and stairs. The policy successfully adapted to these changes, maintaining stable walking by adjusting foot landing angles and body orientation. This demonstrates that the policy does not memorize specific terrain geometries but understands general equilibrium principles on inclined support surfaces.

3) *External Perturbations*: We applied instantaneous lateral thrusts during robot carrying tasks. The strategy rapidly estimated changes in object and self-momentum caused by the disturbance using the world model, swiftly restoring balance through stepping and arm swinging, demonstrating the ability to adapt online to unknown perturbations.

4) *Extend experiment.*: We also extend the experiment to more underactuated interaction tasks on everyday furniture in Fig. 9, and HAIC achieves a 100% success rate on all tasks, including pushing a table, a suitcase, and chairs with an extra yaw freedom on the shaft.

These experiments demonstrate that by combining geometric priors and predictions of dynamics in end-to-end training, HAIC can acquire an understanding of the essence of physical interactions, and exhibit zero-shot generalization capabilities.



Fig. 9: Real-world extend deployment experiment on everyday furniture. HAIC perform 100% SR on table, suitcase and chairs.




Fig. 10: Comparison of the relative object states in robot's local frame from dynamics-aware world model, MuJoCo and reference motion.

E. Analysis and Explanation of world model

1) *Comparison on world model*: To evaluate the effectiveness of our explicit world modeling approach, we compare HAIC against Rapid Motor Adaptation (RMA) [24]. While RMA relies on an implicit environmental latent to adapt to dynamics, our DWM explicitly models the object's physical behaviors. As summarized in Table VII, HAIC achieves superior or competitive success rates and significantly lower tracking errors across most kinematic metrics in both cart manipulation tasks. These results validate that for handling complex underactuated interactions, explicit dynamics prediction is more robust and effective than implicit adaptation.

TABLE VII: Compared to the RMA baseline, HAIC achieves superior or competitive success rates and significantly lower tracking errors.

Task	Method	SR $\uparrow$	$E_{\text{mpbpe}} \downarrow$	$E_{\text{g-mpbpe}} \downarrow$	$E_{\text{mpboe}} \downarrow$	$E_{\text{mpipe}} \downarrow$	$E_{\text{mpope}} \downarrow$	$E_{\text{mpooe}} \downarrow$
Pull Cart	RMA	70%	99.1 $\pm$ 38.2	323.1 $\pm$ 110.8	530.6 $\pm$ 168.6	139.9 $\pm$ 5.6	530.6 $\pm$ 168.6	1474 $\pm$ 178.5
	HAIC	100%	64.8 $\pm$ 2.8	225.7 $\pm$ 3.4	372.7 $\pm$ 34.1	123.1 $\pm$ 0.3	219.5 $\pm$ 6.6	1241 $\pm$ 246.5
Push Cart	RMA	100%	74.1 $\pm$ 6.3	267.0 $\pm$ 52.8	340.8 $\pm$ 30.2	140.1 $\pm$ 6.0	519.9 $\pm$ 92.9	1120 $\pm$ 31.0
	HAIC	100%	61.3 $\pm$ 2.1	221.0 $\pm$ 74.3	319.9 $\pm$ 13.0	145.9 $\pm$ 5.5	267.9 $\pm$ 85.6	1262 $\pm$ 96.2

2) *DWM and Failure case analysis*: We visualize the robot frame object state from DWM prediction, MuJoCo, and reference motion in Fig. 10. We select several highly salient object state components, namely translation along x or y-axis, and rotation on z-axis. For box state, the prediction closely match the simulated results. However, for challenging tasks like pushing or pulling cart and skateboarding, the predicted

states exhibit drift from the simulated trajectories, tending to revert toward the reference motion state. This error eventually leads to an observation mismatch, preventing it from adapting effectively. When its prediction of the object state converges within a error threshold, it can ensure the policy success. When skateboard decelerates and lands, divergence from the actual state will lead to failure (60% SR). As illustrated in Fig. 10, transforming the static terrain into the robot's local frame equips the DWM with crucial spatial awareness. This enables the system to actively alleviate heading drifts and maintain the intended path, even under terrain rotations.

3) *Robustness to Hyperparameters and Annotations*: As shown in Table VIII, even with imprecise thresholds, HAIC maintains a high success rate with only minor kinematic degradation. Furthermore, our binary contact masks are automatically computed via a standardized script rather than manually annotated. This algorithmic approach is generalizable across tasks and avoids human annotation noise.

TABLE VIII: Ablation study on contact reward thresholds. We varied the tolerance (ϵ_{tol}) and force threshold (F_{thr}) parameters (Eq. 7 and 8).

ϵ_{tol}	F_{thr}	SR $\uparrow$	$E_{\text{mpbpe}} \downarrow$	$E_{\text{g-mpbpe}} \downarrow$	$E_{\text{mpboe}} \downarrow$	$E_{\text{mpipe}} \downarrow$	$E_{\text{mpope}} \downarrow$	$E_{\text{mpooe}} \downarrow$
0.1	0	100%	77.6 $\pm$ 15.1	511.7 $\pm$ 525.8	372.7 $\pm$ 50.3	157.3 $\pm$ 16.5	795.7 $\pm$ 722.9	1254 $\pm$ 59.4
0.05	5	100%	68.4 $\pm$ 2.4	388.4 $\pm$ 37.6	334.7 $\pm$ 16.9	137 $\pm$ 1.4	401.5 $\pm$ 37.9	1147 $\pm$ 227.4
0	10	100%	61.3 $\pm$ 2.1	221.0 $\pm$ 74.3	319.9 $\pm$ 13.0	145.9 $\pm$ 5.5	267.9 $\pm$ 85.6	1262 $\pm$ 96.2

V. CONCLUSION

This paper introduced HAIC, a novel framework that enables humanoid robots to perform robust whole-body interactions with underactuated objects under visual occlusion by integrating a dynamic-aware world model. By fusing static geometric priors with dynamic proprioceptive feedback, our approach effectively tracks unobserved object states and anticipates environmental changes through higher-order dynamics prediction. This capability was validated in various scenarios, including climbing stairs, dynamic skateboarding, and cart manipulation, demonstrating proactive gait adjustment and balance recovery. Real-world experiments on the Unitree G1 humanoid confirm that our distilled policy achieves high performance, successfully bridging the sim-to-real gap for deployment in sensor-limited environments.

VI. ACKNOWLEDGMENT

We would like to thank Xiaomi Robotics for providing the experimental site and equipment, Qingfang Li for processing the teaser material, and Shimin Li and Dongfang Li for assisting with the optical motion capture.

REFERENCES

- [1] Rami Ali Al-Asqhar, Taku Komura, and Myung Geol Choi. Relationship descriptors for interactive motion adaptation. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 45–53, 2013.
- [2] Arthur Allshire, Hongsuk Choi, Junyi Zhang, David McAllister, Anthony Zhang, Chung Min Kim, Trevor Darrell, Pieter Abbeel, Jitendra Malik, and Angjoo Kanazawa. Visual imitation enables contextual humanoid control. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2025.
- [3] Joao Pedro Araujo, Yanjie Ze, Pei Xu, Jiajun Wu, and C. Karen Liu. Retargeting matters: General motion retargeting for humanoid motion tracking. *arXiv preprint arXiv:2510.02252*, 2025.
- [4] Bharat Lal Bhatnagar, Xianghui Xie, Ilya Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2022.
- [5] 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.
- [6] Jeremy Dao, Helei Duan, and Alan Fern. Sim-to-real learning for humanoid box loco-manipulation. In *International Conference on Robotics and Automation (ICRA)*, 2024.
- [7] Fan Feng, Phillip Lippe, and Sara Magliacane. Learning interactive world model for object-centric reinforcement learning. In *2511.02225, Thirty-ninth Conference on Neural Information Processing Systems (NeurIPS)*.
- [8] Yuhui Fu, Feiyang Xie, Chaoyi Xu, Jing Xiong, Haoqi Yuan, and Zongqing Lu. Demohlm: From one demonstration to generalizable humanoid loco-manipulation. *arXiv preprint arXiv:2510.11258*, 2025.
- [9] Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. *arXiv preprint arXiv:2408.14472*, 2024.
- [10] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3), 2018.
- [11] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- [12] Jinrui Han, Weiji Xie, Jiakun Zheng, Jiayuan Shi, Weinan Zhang, Ting Xiao, and Chenjia Bai. Kungfubot2: Learning versatile motion skills for humanoid whole-body control. *arXiv preprint arXiv:2509.16638*, 2025.
- [13] Nicklas Hansen, Jyothir SV, Vlad Sobal, Yann LeCun, Xiaolong Wang, and Hao Su. Hierarchical world models as visual whole-body humanoid controllers. *arXiv preprint arXiv:2405.18418*, 2024.
- [14] 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 Hodgins, Linxi “Jim” Fan, Yuke Zhu, Changliu Liu, and Guanya Shi. Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. *arXiv preprint arXiv:2502.01143*, 2025.
- [15] 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.
- [16] 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 *International Conference on Robotics and Automation (ICRA)*, 2025.
- [17] 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. *arXiv preprint arXiv:2502.08378*, 2025.
- [18] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 2019.
- [19] Youngjoon Jeong, Junha Chun, Soonwoo Cha, and Tae-sup Kim. Object-centric world model for language-guided manipulation. *arXiv preprint arXiv:2503.06170*, 2025.
- [20] Haoran Jiang, Jin Chen, Qingwen Bu, Li Chen, Modi Shi, Yanjie Zhang, Delong Li, Chuanzhe Suo, Chuang Wang, Zhihui Peng, et al. Wholebodyvla: Towards unified latent vla for whole-body loco-manipulation control. *arXiv preprint arXiv:2512.11047*, 2025.
- [21] Nan Jiang, Tengyu Liu, Zhexuan Cao, Jieming Cui, Zhiyuan Zhang, Yixin Chen, He Wang, Yixin Zhu, and Siyuan Huang. Full-body articulated human-object interaction. In *International Conference on Computer Vision (ICCV)*, 2023.
- [22] Nan Jiang, Zimo He, Wanhe Yu, Lexi Pang, Yunhao Li, Hongjie Li, Jieming Cui, Yuhao Li, Yizhou Wang, Yixin Zhu, et al. Uniact: Unified motion generation and action streaming for humanoid robots. *arXiv preprint arXiv:2512.24321*, 2025.
- [23] Dvij Kalaria, Sudarshan S Harithas, Pushkal Katara, Sangkyung Kwak, Sarthak Bhagat, Shankar Sastry, Srinath Sridhar, Sai Vemprala, Ashish Kapoor, and Jonathan Chung-Kuan Huang. Dreamcontrol: Human-inspired whole-body humanoid control for scene interaction via guided diffusion. *arXiv preprint arXiv:2509.14353*, 2025.
- [24] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. In *Robotics: Science and Systems (RSS)*, 2021.
- [25] Hang Lai, Jiahang Cao, Jiafeng Xu, Hongtao Wu, Yun-

- feng Lin, Tao Kong, Yong Yu, and Weinan Zhang. World model-based perception for visual legged locomotion. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11531–11537. IEEE, 2025.
- [26] Mikko Lauri, David Hsu, and Joni Pajarinen. Partially observable markov decision processes in robotics: A survey. *IEEE Transactions on Robotics*, 39(1):21–40, 2023.
- [27] Chenhao Li, Andreas Krause, and Marco Hutter. Robotic world model: A neural network simulator for robust policy optimization in robotics. *arXiv preprint arXiv:2501.10100*, 2025.
- [28] 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. *arXiv preprint arXiv:2505.03738*, 2025.
- [29] Jiaman Li, Jiajun Wu, and C Karen Liu. Object motion guided human motion synthesis. *ACM Trans. Graph.*, 42(6), 2023.
- [30] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyao Seo, Georgios Pavlakos, and Yuke Zhu. Okami: Teaching humanoid robots manipulation skills through single video imitation. In *Conference on Robot Learning (CoRL)*, 2024.
- [31] Yitang Li, Zhengyi Luo, Tonghe Zhang, Cunxi Dai, Anssi Kanervisto, Andrea Tirinzoni, Haoyang Weng, Kris Kitani, Mateusz Guzek, Ahmed Touati, et al. Bfm-zero: A promptable behavioral foundation model for humanoid control using unsupervised reinforcement learning. *arXiv preprint arXiv:2511.04131*, 2025.
- [32] Yitang Li, Yuanhang Zhang, Wenli Xiao, Chaoyi Pan, Haoyang Weng, Guanqi He, Tairan He, and Guanya Shi. Learning gentle humanoid locomotion and end-effector stabilization control. *arXiv preprint arXiv:2505.24198*, 2025.
- [33] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. *The International Journal of Robotics Research (IJRR)*, 2024.
- [34] 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.
- [35] Yuhang Lin, Yijia Xie, Jiahong Xie, Yuehao Huang, Ruoyu Wang, Jiajun Lv, Yukai Ma, and Xingxing Zuo. Simgenhoi: Physically realistic whole-body humanoid-object interaction via generative modeling and reinforcement learning. *arXiv preprint arXiv:2508.14120*, 2025.
- [36] Chenhao Liu, Leyun Jiang, Yibo Wang, Kairan Yao, Jinchun Fu, and Xiaoyu Ren. Humanoid whole-body badminton via multi-stage reinforcement learning. *arXiv preprint arXiv:2511.11218*, 2025.
- [37] Fukang Liu, Zhaoyuan Gu, Yilin Cai, Ziyi Zhou, Shijie Zhao, Hyunyoung Jung, Sehoon Ha, Yue Chen, Danfei Xu, and Ye Zhao. Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation. *arXiv preprint arXiv:2409.20514*, 2024.
- [38] Hang Liu, Yuman Gao, Sangli Teng, Yufeng Chi, Yakun Sophia Shao, Zhongyu Li, Maani Ghaffari, and Koushil Sreenath. Ego-vision world model for humanoid contact planning. *arXiv preprint arXiv:2510.11682*, 2025.
- [39] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. *arXiv preprint arXiv:2411.14386*, 2024.
- [40] 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.
- [41] 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.
- [42] Daniel Morales-Brotons, Thijs Vogels, and Hadrien Hendrikx. Exponential moving average of weights in deep learning: Dynamics and benefits. *Trans. Mach. Learn. Res.*, 2024.
- [43] Liang Pan, Zeshi Yang, Zhiyang Dou, Wenjia Wang, Buzhen Huang, Bo Dai, Taku Komura, and Jingbo Wang. Tokenhsi: Unified synthesis of physical human-scene interactions through task tokenization. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2025.
- [44] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.
- [45] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (ToG)*, 40(4):1–20, 2021.
- [46] Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor critic for image-based robot learning. *arXiv preprint arXiv:1710.06542*, 2017.
- [47] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 2024.
- [48] Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.
- [49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec

- Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [50] Yiyang Shao, Xiaoyu Huang, Bike Zhang, Qiayuan Liao, Yuman Gao, Yufeng Chi, Zhongyu Li, Sophia Shao, and Koushil Sreenath. Langwbc: Language-directed humanoid whole-body control via end-to-end learning. *arXiv preprint arXiv:2504.21738*, 2025.
- [51] Andrea Sipos and Nima Fazeli. Simultaneous contact location and object pose estimation using proprioception and tactile feedback. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [52] Zhi Su, Bike Zhang, Nima Rahmanian, Yuman Gao, Qiayuan Liao, Caitlin Regan, Koushil Sreenath, and S Shankar Sastry. Hitter: A humanoid table tennis robot via hierarchical planning and learning. *arXiv preprint arXiv:2508.21043*, 2025.
- [53] Wandong Sun, Long Chen, Yongbo Su, Baoshi Cao, Yang Liu, and Zongwu Xie. Learning humanoid locomotion with world model reconstruction. *arXiv preprint arXiv:2502.16230*, 2025.
- [54] Wandong Sun, Luying Feng, Baoshi Cao, Yang Liu, Yaochu Jin, and Zongwu Xie. Ulc: A unified and fine-grained controller for humanoid loco-manipulation. *arXiv preprint arXiv:2507.06905*, 2025.
- [55] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision (ECCV)*, 2020.
- [56] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 43(6):1–21, 2024.
- [57] 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)*, 2017.
- [58] Huayi Wang, Wentao Zhang, Runyi Yu, Tao Huang, Junli Ren, Feiyu Jia, Zirui Wang, Xiaojie Niu, Xiao Chen, Jiahe Chen, Qifeng Chen, Jingbo Wang, and Jiangmiao Pang. Physhsi: Towards a real-world generalizable and natural humanoid-scene interaction system. *arXiv preprint arXiv:2510.11072*, 2025.
- [59] Jin Wang, Rui Dai, Weijie Wang, Luca Rossini, Francesco Ruscelli, and Nikos Tsagarakis. Hypermotion: Learning hybrid behavior planning for autonomous loco-manipulation. In *Conference on Robot Learning (CoRL)*, 2024.
- [60] Yinhuai Wang, Jing Lin, Ailing Zeng, Zhengyi Luo, Jian Zhang, and Lei Zhang. Physshoi: Physics-based imitation of dynamic human-object interaction. *arXiv preprint arXiv:2312.04393*, 2023.
- [61] Yinhuai Wang, Qihan Zhao, Runyi Yu, Hok Wai Tsui, Ailing Zeng, Jing Lin, Zhengyi Luo, Jiwen Yu, Xiu Li, Qifeng Chen, Jian Zhang, Lei Zhang, and Ping Tan. Skillmimic: Learning basketball interaction skills from demonstrations. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 17540–17549, June 2025.
- [62] Yushi Wang, Changsheng Luo, Penghui Chen, Jianran Liu, Weijian Sun, Tong Guo, Kechang Yang, Biao Hu, Yangang Zhang, and Mingguo Zhao. Learning vision-driven reactive soccer skills for humanoid robots. *arXiv preprint arXiv:2511.03996*, 2025.
- [63] 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.
- [64] Qianyang Wu, Ye Shi, Xiaoshui Huang, Jingyi Yu, Lan Xu, and Jingya Wang. Thor: Text to human-object interaction diffusion via relation intervention. *arXiv preprint arXiv:2403.11208*, 2024.
- [65] Weiji Xie, Jinrui Han, Jiakun Zheng, Huanyu Li, Xinzhe Liu, Jiyuan Shi, Weinan Zhang, Chenjia Bai, and Xuelong Li. Kungfubot: Physics-based humanoid whole-body control for learning highly-dynamic skills. *arXiv preprint arXiv:2506.12851*, 2025.
- [66] Sirui Xu, Dongting Li, Yucheng Zhang, Xiyan Xu, Qi Long, Ziyin Wang, Yunzhi Lu, Shuchang Dong, Hezi Jiang, Akshat Gupta, Yu-Xiong Wang, and Liang-Yan Gui. Interact: Advancing large-scale versatile 3d human-object interaction generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2025.
- [67] Sirui Xu, Hung Yu Ling, Yu-Xiong Wang, and Liangyan Gui. Intermimic: Towards universal whole-body control for physics-based human-object interactions. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2025.
- [68] Zifan Xu, Myoungkyu Seo, Dongmyeong Lee, Hao Fu, Jiaheng Hu, Jiaxun Cui, Yuqian Jiang, Zhihan Wang, Anastasiia Brund, Joydeep Biswas, et al. Learning agile striker skills for humanoid soccer robots from noisy sensory input. *arXiv preprint arXiv:2512.06571*, 2025.
- [69] Haoru Xue, Xiaoyu Huang, Dantong Niu, Qiayuan Liao, Thomas Kragerud, Jan Tommy Gravdahl, Xue Bin Peng, Guanya Shi, Trevor Darrell, Koushil Sreenath, et al. Leverb: Humanoid whole-body control with latent vision-language instruction. *arXiv preprint arXiv:2506.13751*, 2025.
- [70] Lujie Yang, Xiaoyu Huang, Zhen Wu, Angjoo Kanazawa, Pieter Abbeel, Carmelo Sferrazza, C Karen Liu, Rocky Duan, and Guanya Shi. Omniretarget: Interaction-preserving data generation for humanoid whole-body loco-manipulation and scene interaction. *arXiv preprint arXiv:2509.26633*, 2025.
- [71] Kangning Yin, Weishuai Zeng, Ke Fan, Minyue Dai, Zirui Wang, Qiang Zhang, Zheng Tian, Jingbo Wang,

- Jiangmiao Pang, and Weinan Zhang. Unitracker: Learning universal whole-body motion tracker for humanoid robots. *arXiv preprint arXiv:2507.07356*, 2025.
- [72] 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.
 - [73] Mingqi Yuan, Tao Yu, Wenqi Ge, Xiuyong Yao, Dapeng Li, Huijiang Wang, Jiayu Chen, Xin Jin, Bo Li, Hua Chen, et al. Behavior foundation model: Towards next-generation whole-body control system of humanoid robots. *arXiv preprint arXiv:2506.20487*, 2025.
 - [74] Yanjie Ze, Zixuan Chen, João Pedro Araújo, Zi an Cao, Xue Bin Peng, Jiajun Wu, and C. Karen Liu. Twist: Teleoperated whole-body imitation system. *arXiv preprint arXiv:2505.02833*, 2025.
 - [75] 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. *arXiv preprint arXiv:2511.02832*, 2025.
 - [76] Weishuai Zeng, Shunlin Lu, Kangning Yin, Xiaojie Niu, Minyue Dai, Jingbo Wang, and Jiangmiao Pang. Behavior foundation model for humanoid robots. *arXiv preprint arXiv:2509.13780*, 2025.
 - [77] Juzhe Zhang, Haimin Luo, Hongdi Yang, Xinru Xu, Qianyang Wu, Ye Shi, Jingyi Yu, Lan Xu, and Jingya Wang. NeuralDome: A neural modeling pipeline on multi-view human-object interactions. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2023.
 - [78] Yuanhang Zhang, Yifu Yuan, Prajwal Gurunath, Ishita Gupta, Shayegan Omidshafiei, Ali-akbar Agha-mohammadi, Marcell Vazquez-Chanlatte, Liam Pedersen, Tairan He, and Guanya Shi. Falcon: Learning force-adaptive humanoid loco-manipulation. *arXiv preprint arXiv:2505.06776*, 2025.
 - [79] 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.
 - [80] 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.