Beyond Decoupled Control: Rethinking Whole-Body Navigation for Humanoid Robots

Haihong Hao[1,2] Hui Zhu[2] Tingda Zhuang[2] Wenwen Luo[3,2] Chen Liang[4,2]
Lei Chen[1] Mingfei Han[1] Xiaojun Chang[1] Zhihui Li[1]
[1] University of Science and Technology of China [2] Honor
[3] Beihang University [4] University of Science and Technology Beijing

Abstract

Humanoid robots hold promise for general-purpose embodied intelligence, yet their navigation systems largely inherit decoupled paradigms from wheeled and quadrupedal platforms. While such modular designs enable scalable optimization, they underexploit a key advantage of humanoids: the tight coupling of perception, locomotion, and physical interaction. Realizing this potential requires rethinking the conventional separation between navigation and manipulation. In this work, we revisit whole-body humanoid navigation and present HumiAct*, a unified framework that couples a high-level reasoning planner with a shared vision-language-action policy and whole-body controller for coordinated planning, locomotion, and manipulation. We establish a training and evaluation ecosystem spanning heterogeneous datasets and physics-based simulation with paired-route protocols. Paired runs on the same routes reveal a* vision dependence dilemma*. Camera feedback helps the humanoid stay on course, yet changing image appearance alone degrades performance due to reliance on scene-specific cues. We introduce* latent action guidance *to distill action-conditioned motion priors from cross-embodiment navigation and manipulation trajectories into the shared policy, reducing appearance sensitivity without inference overhead. Built on GR00T, HumiAct's single shared policy outperforms fine-tuned baselines and separately trained specialists on direct navigation, planner-guided navigation, manipulation, and tasks that combine navigation with physical interaction. We further validate the system on a real humanoid across navigation and interaction tasks. Together, HumiAct and its data-to-robot ecosystem advance humanoid navigation toward unified whole-body intelligence.*

1. Introduction

Humanoid robots combine whole-body mobility with the ability to physically interact with their surroundings, a capability brought increasingly within reach by recent advances in humanoid control and generalist robot policies [20, 64]. Yet many humanoid systems still treat navigation and manipulation as separate problems. While this separation simplifies individual skill development, it underexploits a defining advantage of humanoid embodiment: coordinating perception, locomotion, and physical interaction through the same body. A humanoid approaching an object must coordinate where it moves, what it sees, and how it configures its body for the interaction that follows. These demands motivate a broader question: *How can a humanoid perceive, move, and physically interact with its environment as a unified embodied agent?*

Existing deployments typically pair separate navigation and manipulation policies with partial-body controllers [56], imposing boundaries between behaviors executed by the same body (Fig. 1a). We present *HumiAct*, a three-level humanoid agent that unifies navigation and manipulation within a shared policy–controller stack under high-level task planning (Fig. 1b). A language-model planner interprets instructions, decomposes long-horizon tasks into short language subtasks, and monitors progress. Rather than switching between separate skill modules across subtasks, one shared vision-language-action (VLA) policy maps the active subtask, visual observations, and proprioception to whole-body motion commands, while one low-level whole-body controller tracks these commands and maintains balance. This design eliminates policy–controller handoffs between navigation and manipulation and supports coupled loco-manipulation such as stepping while pushing.

Unifying navigation and manipulation in a mobile VLA policy, however, exposes a *vision dependence dilemma*. While whole-body traversal and interaction span wide viewpoint and background shifts, humanoid demonstrations remain scarce and confined to few scenes. Removing vision causes accumulated drift and trajectory deviation, whereas incorporating vision enables closed-loop correction relative to the environment. Yet when trained on limited humanoid scenes, the policy readily latches onto static appearance shortcuts correlated with training actions, degrading sharply under unfamiliar scenes or image corruptions. The central challenge is to preserve closed-loop visual feedback without overfitting to scene-specific appearance. We ad-

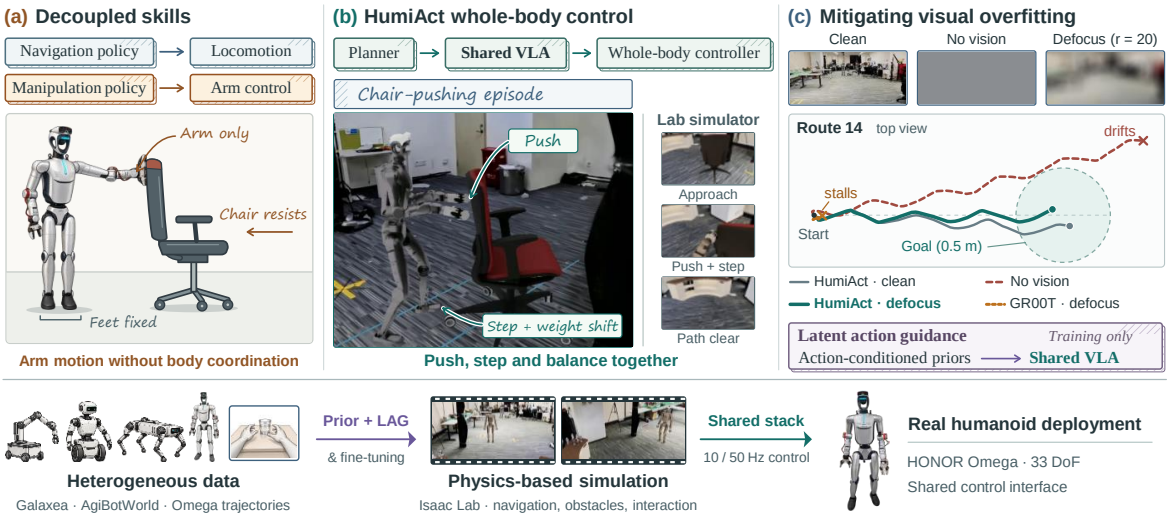


Figure 1. Motivation and overview of HumiAct. (a) Decoupled skills: separate navigation and manipulation stacks limit whole-body coordination, illustrated by arm-only chair pushing with fixed feet. **(b) Whole-body control:** HumiAct connects a task planner, a shared VLA policy, and a whole-body controller to coordinate pushing, stepping, and balance (chair-pushing episode). **(c) Mitigating visual overfitting:** Under strong defocus, HumiAct follows nearly the same path and reaches nearly the same endpoint as with clean images.

dress this challenge by grounding visual representations in *action-centric predictive priors* learned from diverse cross-embodiment trajectories. Predicting how actions change visual observations pushes features to capture task-relevant geometry and ego-motion rather than static background textures, while cross-embodiment navigation and manipulation datasets provide the scene diversity missing from humanoid-only data. Although different embodiments use distinct control interfaces, their motions can be abstracted into shared patterns of how actions transform camera views. We therefore introduce *latent action guidance* (LAG) to transfer these motion priors into the shared VLA policy. Mapping heterogeneous robot trajectories onto a compact vocabulary of shared latent actions, an action-conditioned predictor first learns across embodiments how each behavior alters what the camera sees; its latent current and predicted future features are then distilled into the VLA backbone prior to whole-body fine-tuning, shaping action-aware visual representations offline without adding any inference-time module (Fig. 1c). Evaluated in a physics-based reconstruction of an indoor workspace with paired-route protocols, HumiAct outperforms fine-tuned GR00T by 20 percentage points in success rate on obstacle routes and by 9.4 points on combined navigation–interaction tasks, while averaging 84% success across five image corruptions on goal-facing routes. We further deploy HumiAct on a physical humanoid robot, demonstrating unified navigation and whole-body interaction in the real world.

Our contributions are threefold:

- **A unified humanoid agent for navigation and interaction.** HumiAct couples a language-model planner with one shared VLA policy and one whole-body controller, so that navigation, manipulation, and the transition between them are produced by the same embodied agent.
- **Action-centric guidance for visual control.** We identify the vision dependence dilemma and introduce LAG, which transfers motion priors learned from heterogeneous embodied trajectories into the shared VLA policy to reduce reliance on scene-specific appearance.
- **A data-to-robot ecosystem.** We build a training and evaluation pipeline spanning heterogeneous embodied data, a physics-based reconstruction of an indoor workspace, and real-humanoid deployment, with paired-route protocols that isolate policy, planner, and visual-input effects.

2. Related Work

Navigation and mobile manipulation. Vision-and-language navigation has progressed from graphs to continuous spaces [4, 48] via large-scale simulation and RL [74, 84], advancing alignment, history reasoning, and multilingual grounding [12, 19, 28, 31, 38, 49, 83], alongside cross-embodiment goal-conditioned navigation [76–78]. Language planners ground instructions into skills via affordances, feedback, code, or value maps [2, 16, 39, 40, 53], while mobile manipulators pair navigation with grasping [21, 56]. Unlike wheeled or modular stacks, HumiAct unifies navigation and manipulation in one visuomotor policy and whole-body controller.

Whole-body humanoid control and VLA policies. Motion priors and legged controllers [41, 61, 67–69, 73] underpin humanoid sim-to-real locomotion [23], teleoperation and imitation [20, 33, 34], and motion tracking [35, 54, 57]. Meanwhile, generalist VLA policies [8, 9, 22,

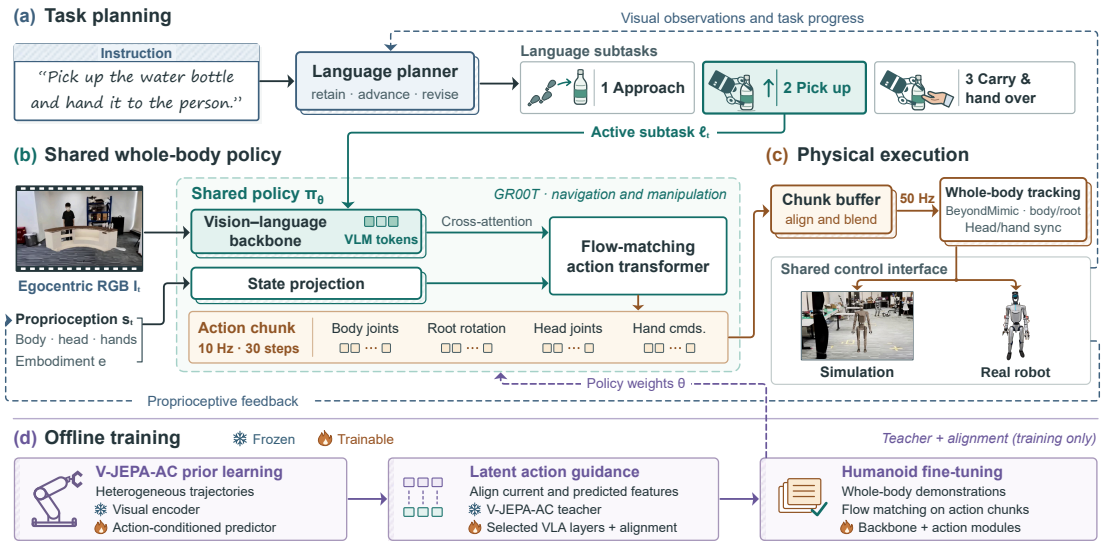


Figure 2. Overview of HumiAct. (a) A task planner updates the active language subtask using visual progress. (b) On hardware, the shared VLA predicts 30-step whole-body action chunks at 10 Hz. (c) Buffered and blended references are streamed at 50 Hz for whole-body tracking and synchronized head and hand execution. (d) Predictive prior learning, LAG transfer, and humanoid fine-tuning provide the policy weights; the teacher and alignment modules are used only during training.

46, 64, 92] leverage large cross-robot datasets [15, 45, 65, 81]. Recent humanoid systems explore video-learned loco-manipulation [43], cluttered whole-body navigation [52], and embodiment-routed stereo [86]; HumiAct instead unifies planner-guided navigation, manipulation, and their composition in one policy while explicitly regulating its visual dependence.

Latent actions, predictive learning, and visual generalization. Prior methods compress multimodal behavior or infer visual latent actions [10, 51, 58, 75, 82, 87], learn latent world models [24–27, 30], or guide policies with predictive video, state, denoising, and view-invariant representations [6, 7, 29, 42, 79, 91]. Visual overfitting is typically countered via randomization, augmentation, or pretrained encoders [14, 47, 50, 59, 62, 80, 89] alongside self-, language-, and masked-supervised pretraining [5, 11, 32, 66, 70, 71]. LAG is complementary: conditioning prediction on five shared cross-embodiment behaviors avoids equating raw control vectors while distilling current and future features into the VLA without pixel prediction or test-time modules.

3. Method

3.1. Unified Whole-Body Control

For a humanoid, approaching a target couples locomotion, viewpoint, body orientation, and the posture required for the interaction that follows. HumiAct captures this coupling with the three levels of Fig. 2; navigation and manipulation share both the policy and the execution stack, so their transition is part of one continuous perception–action process.

Task planning with progress feedback. Given a long-horizon instruction g and an initial observation I_{t_0}, a language-model planner produces an ordered sequence of short language subtasks,

$$\mathcal{P}_0 = (\ell_1, \ldots, \ell_M) = \Pi_{\text{plan}}(g, I_{t_0}), \tag{1}$$

such as moving forward, turning, or interacting with an object. It periodically evaluates progress from incoming observations and the execution context, and retains the current subtask, advances to the next one, or revises the remaining plan. The subtask ℓ_t active at time t is the language condition of the shared policy.

Shared whole-body action generation. Let I_t be the egocentric RGB observation, s_t the proprioceptive state of body, head, and hands, and e the embodiment identity. Following action-sequence policy designs [13, 90], one policy predicts an action chunk of horizon H for every behavior:

$$\hat{\mathbf{A}}_t = \hat{\mathbf{a}}_{t:t+H-1} = \pi_\theta(c_t), \qquad c_t = (I_t, s_t, \ell_t, e). \tag{2}$$

We build on the GR00T architecture [64]: a vision-language backbone, embodiment-conditioned state and action projections, and an action transformer f_θ trained with flow matching [55]. For a demonstration chunk $\mathbf{A}$, noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$, and $\tau \in [0, 1]$, the noised action and its target velocity are

$$\mathbf{X}_\tau = (1-\tau)\boldsymbol{\epsilon} + \tau\mathbf{A}, \qquad \mathbf{V}^\star = \mathbf{A} - \boldsymbol{\epsilon}, \tag{3}$$

and the policy is trained with

$$\mathcal{L}_{\text{FM}} = \mathbb{E}\left[\frac{\|\mathbf{m} \odot (f_\theta(\mathbf{X}_\tau, \tau, c_t) - \mathbf{V}^\star)\|_F^2}{\|\mathbf{m}\|_1}\right], \tag{4}$$

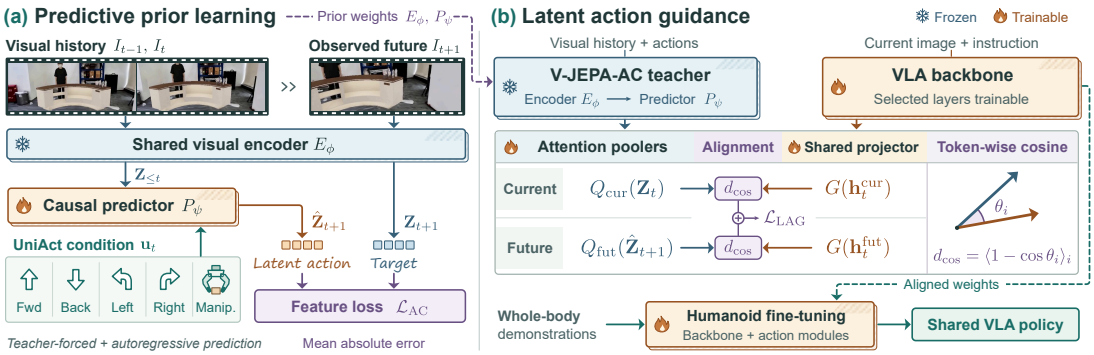


Figure 3. Learning and transferring action-conditioned motion priors. (a) Predictive prior learning: heterogeneous transitions mapped to five shared behaviors supervise causal predictor P_ψ over frozen encoder E_ϕ via teacher-forced and autoregressive rollouts. **(b) Latent action guidance:** frozen V-JEPA-AC features pooled by Q_{cur} and Q_{fut} supervise projected VLA auxiliary representations (G) via cosine alignment. **Bottom:** the aligned backbone and action modules are fine-tuned on humanoid demonstrations; teacher, auxiliary tokens, and alignment modules are discarded at inference.

where the mask $\mathbf{m}$ selects valid action dimensions and steps; at inference, integrating f_θ from $\mathbf{X}_0 = \boldsymbol{\epsilon}$ over $\tau \in [0, 1]$ yields $\hat{\mathbf{A}}_t$. Because legs, heading, head, waist, and arms are predicted within one trajectory, an interaction subtask continues from the body configuration and viewpoint that the approach produced.

Whole-body references and multi-rate execution. Each predicted step is decoded into a synchronized reference

$$\mathbf{r}_{t+j} = \left(\hat{\mathbf{q}}_{t+j}^{\mathrm{body}}, \hat{\mathbf{R}}_{t+j}^{\mathrm{root}}, \hat{\mathbf{q}}_{t+j}^{\mathrm{head}}, \hat{\mathbf{g}}_{t+j}\right), \quad j = 0, \ldots, H-1, \tag{5}$$

of body-joint positions, root orientation, head-joint positions, and hand commands. On hardware, inference and execution are asynchronous: chunks predicted at 10 Hz are buffered, aligned to the execution timeline, and blended at transitions, and references are streamed at 50 Hz to a whole-body controller that tracks body and root targets from proprioceptive feedback and maintains balance, with head and hand commands executed in sync.

3.2. Learning Latent Actions from Heterogeneous Trajectories

Embodiments differ in their control interfaces, yet their trajectories share one learning signal: how actions shape visual dynamics (Fig. 3a). **Unified action semantics (UniAct).** From the motion records and task annotations of open robot datasets [1, 44] and our humanoid trajectories, we map every transition to one of five behaviors: moving forward, moving backward, turning left, turning right, and manipulation. Base translation and rotation identify the navigation classes, arm and hand motion identify manipulation, and idle or ambiguous intervals are discarded. The resulting one-hot condition $\mathbf{u}_t$ gives a visual transition the same action meaning on every embodiment. It is used only as the condition for predictive learning; physical execution keeps the continuous whole-body commands of Sec. 3.1.

Action-conditioned future prediction. We build V-JEPA-AC on the predictive architecture of V-JEPA 2 [6]. A frozen pretrained encoder E_ϕ defines the representation space and a causal predictor P_ψ models its evolution under the semantic condition:

$$\mathbf{Z}_t = E_\phi(I_t), \qquad \hat{\mathbf{Z}}_{t+1} = P_\psi\left(\mathbf{Z}_{\le t}, \mathbf{u}_{\le t}\right), \tag{6}$$

with normalized patch-token representations $\mathbf{Z}_t$. Predictions are supervised by encoded future observations,

$$\mathcal{L}_{\mathrm{AC}} = \mathbb{E}\left[d\left(\hat{\mathbf{Z}}_{1:T-1}^{\mathrm{TF}}, \mathbf{Z}_{1:T-1}\right) + d\left(\hat{\mathbf{Z}}_{1:K}^{\mathrm{AR}}, \mathbf{Z}_{1:K}\right)\right], \tag{7}$$

where d is the mean absolute feature difference and the two terms cover teacher-forced prediction and K-step autoregressive rollouts. The predictor's final-layer tokens, projected into the encoder feature space, are the latent actions: they encode what the scene will look like if the given action is executed.

3.3. Latent Action Guidance

LAG transfers these action-conditioned dynamics into the shared VLA (Fig. 3b). Following knowledge and intermediate-feature distillation [37, 72, 88], the frozen teacher provides current features $\mathbf{Z}_t$ and predicted future features $\hat{\mathbf{Z}}_{t+1}$. Given the current image and instruction, the VLA produces auxiliary-token representations $\mathbf{h}_t^{\mathrm{cur}}$ and $\mathbf{h}_t^{\mathrm{fut}}$ at an intermediate backbone layer. A shared projector G maps them into the teacher space, learned attention poolers Q_{cur} and Q_{fut} form compact teacher targets, and we minimize

$$\begin{aligned}\mathcal{L}_{\mathrm{LAG}} = \mathbb{E}\big[&d_{\cos}(G(\mathbf{h}_t^{\mathrm{cur}}), Q_{\mathrm{cur}}(\mathbf{Z}_t)) \\ &+ d_{\cos}\Big(G(\mathbf{h}_t^{\mathrm{fut}}), Q_{\mathrm{fut}}(\hat{\mathbf{Z}}_{t+1})\Big)\big],\end{aligned} \tag{8}$$

with $d_{\cos}$ the mean token-wise cosine distance. The backbone thus learns to represent the current scene together with

Table 1. Training data. Stage-specific episode counts. All sources cover navigation and manipulation; the whole-body demonstrations additionally cover navigation–interaction.

Source	Embodiment	Episodes	Stage
Galaxea [44]	R1 Lite / Pro	34,353	Prior, LAG
AgiBotWorld [1]	A2D	142,069	Prior, LAG
Humanoid trajectories	Target humanoid	1,962	Prior, LAG
Whole-body demos	Target humanoid	5,421	Fine-tuning

its expected evolution under the instructed behavior, a target that is defined by motion and not by the appearance of a particular environment. The aligned backbone then initializes the shared VLA, which is trained for whole-body control with $\mathcal{L}_{\mathrm{FM}}$. Teacher, auxiliary tokens, and alignment modules are discarded afterwards, so LAG adds no inference cost and leaves the policy interface of Eq. 2 unchanged.

3.4. Data-to-Robot Ecosystem

Our ecosystem connects heterogeneous motion learning, humanoid adaptation, and closed-loop evaluation through a common observation–action interface.

Multi-source data. We combine navigation and manipulation trajectories from open robot datasets with whole-body humanoid demonstrations (Table 1). The heterogeneous trajectories, about 178k episodes from three embodiments, provide broad coverage of action-dependent visual dynamics for prior learning and transfer. The humanoid demonstrations then connect these priors to coordinated locomotion, object handling, and navigation–interaction on the target platform. **Training pipeline.** Training has three stages. (i) The predictor P_ψ is trained with $\mathcal{L}_{\mathrm{AC}}$ on the heterogeneous corpus, with the pretrained V-JEPA 2 encoder frozen. (ii) The frozen teacher transfers motion priors into a GR00T-initialized VLA through $\mathcal{L}_{\mathrm{LAG}}$ on the same corpus, updating selected backbone layers and the alignment modules. (iii) Backbone and action modules are jointly fine-tuned on the whole-body demonstrations with $\mathcal{L}_{\mathrm{FM}}$. All stages train in BF16 on 16 NVIDIA H20 GPUs: 5k prior steps (24-layer P_ψ, $K{=}2$), 32-token LAG alignment, and 200k fine-tuning steps (batch 96) under video-disjoint splits (Appendix). **Physics-based simulation and task evaluation.** We build the simulation environment on Isaac Lab [63] with PhysX, combining a reconstructed indoor workspace with collision-enabled obstacles and manipulable objects. We vary robot initialization, object placement, illumination, and object mass to evaluate behavior across visual and physical conditions, and seeded resets with automatic task checks give repeatable measurements of navigation progress, collisions, grasping, and object transport. An adapter connects the unified motion references to a BeyondMimic tracker [54] with synchronized head and hand execution; physics runs at 200 Hz and whole-body tracking at 50 Hz. The simulator exposes the policy interface used on hardware, enabling closed-loop evaluation and paired reruns of the same route under different policies or visual inputs. **Real-humanoid deployment.** We deploy HumiAct on the same 33-DoF humanoid platform used in EATR-Stereo [86]. The robot has synchronized head-mounted cameras and a 37-dimensional proprioceptive state comprising joint positions and base orientation. The shared VLA runs on an NVIDIA RTX 5000 Ada at 10 Hz and predicts 30-step action chunks; BF16 inference with four flow-integration steps keeps action generation within the control budget. A ROS 2 client exchanges visual and proprioceptive observations with the planner–policy service through ZeroMQ and streams whole-body references at 50 Hz, while asynchronous buffering and overlap blending maintain continuous execution.

4. Experiments

4.1. Experimental Setup

Tasks and metrics. We evaluate the humanoid in the reconstructed indoor workspace (Sec. 3.4). *Direct navigation* uses obstacle-free routes with a goal-facing start, a forward-motion instruction, a 10 s episode, and a 0.5 m goal radius. *Planner navigation* uses obstacle routes on which the planner selects and updates the instruction. *Manipulation* starts at the work surface, and *navigation–interaction* uses farther starts that require an approach followed by the object task. We report success rate (SR), success weighted by path length (SPL), and the fraction of trials with an obstacle contact above 10 N (Coll.). In simulation every policy executes the first 30 steps of each 40-step chunk at 20 Hz. GPT-5.6 Terra is the default planner.

Compared policies. All policies are fine-tuned on our humanoid demonstrations and share the whole-body controller. *GR00T* [64] is fine-tuned from the public base model and represents a vision-only whole-body VLA. GR00T uses the same 200k-step fine-tuning (batch 96). The visual-input study also evaluates *GR00T*†, a second GR00T checkpoint fine-tuned with HumiAct's downstream recipe from the public base model. The *navigation-only* specialist shares HumiAct's LAG initialization but sees no manipulation data, the *manipulation-only* specialist sees no navigation data, and the *prior-initialized reference* is an earlier LAG-initialized policy without the final joint adaptation.

4.2. Whole-Body Task Performance

One policy for navigation and manipulation. Table 2 compares policies on four whole-body tasks. On direct routes HumiAct completes 13 of 15 (SPL 0.705) against 10 of 15 (0.517) for GR00T, gaining three routes and losing none (route-paired ΔSR $+20.0$ points, 95% bootstrap interval $[0.0, 40.0]$). With the same planner it completes 8 obstacle routes against 5, with SPL rising from 0.258 to

Table 2. Whole-body task performance in simulation. Direct navigation uses the standard reset without an extra camera render. All planner-navigation rows use GPT-5.6 Terra. SR and collision rate (Coll.) in %; 15 routes per navigation setting and 32 trials for navigation–interaction, where *Reach* counts arrival at the interaction region and *Complete* additionally requires the object task. –: setting not applicable to, or not run for, the method.

	Direct navigation		Planner navigation			Manipulation	Navigation + interaction	
Method	SR ↑	SPL ↑	SR ↑	SPL ↑	Coll. ↓	SR ↑	Reach ↑	Complete ↑
Manipulation-only specialist	–	–	–	–	–	32.0	–	–
Navigation-only specialist + LAG	73.3	0.574	46.7	0.330	40.0	–	–	–
GR00T [64]	66.7	0.517	33.3	0.258	**26.7**	33.0	43.8 (14/32)	6.3 (2/32)
Prior-initialized reference	73.3	0.540	**53.3**	0.385	46.7	–	–	–
HumiAct	**86.7**	**0.705**	**53.3**	**0.410**	46.7	**36.0**	**56.3 (18/32)**	**15.6 (5/32)**

Table 3. Vision dependence on the 15 direct-navigation routes after a camera refresh at reset. Methods share route geometry and seeds within each comparison. **(a)** SR (%) and SPL under each input condition (seed 42). **(b)** Successful clean routes (of 15) under three inference seeds. **(c)** Mean maximum lateral deviation from the start–goal line (Lat., m) and number of stalled routes (St., progress < 0.5 m). GR00T† denotes the additional GR00T checkpoint and Nav.+LAG the navigation specialist.

(a) Corrupted and removed visual input

		GR00T		GR00T†		Nav.+LAG		HumiAct	
	Condition	SR	SPL	SR	SPL	SR	SPL	SR	SPL
	Clean	80.0	.606	86.7	.669	86.7	.674	**93.3**	**.719**
Appearance change	Defocus (r=6)	80.0	.634	**86.7**	**.687**	73.3	.516	80.0	.629
	Noise (σ=40)	**86.7**	**.683**	**86.7**	.681	80.0	.586	80.0	.590
	Brightness ($\times$1.8)	73.3	.550	66.7	.502	73.3	.529	**86.7**	**.614**
	Motion blur (23 px)	73.3	.556	80.0	.618	80.0	.577	**86.7**	**.633**
	JPEG (q=10)	73.3	.560	80.0	.626	80.0	.562	**86.7**	**.649**
	Mean of five	77.3	.597	80.0	**.623**	77.3	.554	**84.0**	**.623**
	Strong defocus (r=20)	46.7	.347	**73.3**	**.570**	**73.3**	.535	**73.3**	.546
	Strong noise (σ=100)	60.0	.443	73.3	**.559**	73.3	.521	**80.0**	.553
Rem.	Gray image	**73.3**	.495	66.7	**.504**	66.7	.418	46.7	.305
	Central occlusion (25%)	**80.0**	**.574**	66.7	.492	73.3	.479	60.0	.443

(b) Clean routes completed per inference seed

Seed	GR00T	GR00T†	Nav.+LAG	HumiAct
42	12	13	13	14
43	11	10	11	12
44	11	12	13	12
Mean	11.3	11.7	12.3	**12.7**

(c) Trajectory diagnostics

	GR00T		GR00T†		Nav.+LAG		HumiAct	
Condition	Lat.	St.	Lat.	St.	Lat.	St.	Lat.	St.
Clean	.28	0	.36	0	.25	0	.21	0
Gray image	.50	1	.37	0	.27	0	.46	1
Occlusion	.30	0	.34	0	.35	0	.29	0
Strong defocus	.25	6	.33	0	.28	0	.34	0
Strong noise	.31	0	.29	0	.33	0	.32	0

0.410; Fig. 4a shows a recorded direct-navigation episode for both policies. Sharing the policy does not trade one skill for the other: the navigation specialist with the same initialization reaches 73.3% and 46.7% where the joint policy reaches 86.7% and 53.3%, and manipulation is retained (36.0% against 32.0% for the manipulation specialist and 33.0% for GR00T).

Composed navigation–interaction. The composed task is where a unified body matters most, since the approach determines the viewpoint and posture from which the interaction starts. HumiAct reaches the interaction region in 18/32 trials and completes the task in 5/32, against 14/32 and 2/32 for GR00T. The navigation specialist shares HumiAct's LAG initialization, so its gap (13.4/6.6 points) isolates joint training from LAG (ablated in Table 7).

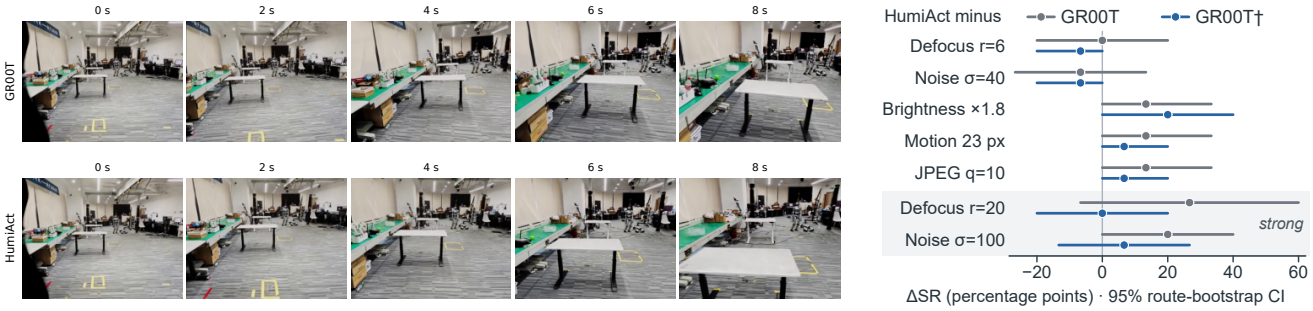
4.3. Vision Dependence Dilemma

Following the use of appearance corruptions to probe visual robustness [36], Table 3 intervenes on the policy image alone; language, proprioception, and controller inputs are untouched. The visual-input study renders the camera after each reset and before the first policy action. With this protocol, HumiAct completes 14/15 clean routes, GR00T† 13/15, and GR00T 12/15, compared with 13/15 for HumiAct and 10/15 for GR00T under the standard reset in Table 2.

Table 4. Direct-navigation calibration (protocol and metrics follow Table 3). With clean RGB, HumiAct reaches more goals with tighter lateral tracking than non-learned gait-replay references, including privileged-pose steering.

Controller	Input	SR ↑	SPL ↑	Lat. ↓
Non-learned references				
Open-loop gait	none	80.0	.616	.39
Heading hold	root yaw	80.0	.617	.39
Oracle steering	true pose, goal	86.7	.807	.23
Learned policies, camera image				
GR00T [64]	RGB	80.0	.606	.28
GR00T†	RGB	86.7	.669	.36
Nav.+LAG	RGB	86.7	.674	.25
HumiAct	RGB	93.3	.719	.21
Learned policies, uniform gray image				
GR00T [64]	gray	73.3	.495	.50
HumiAct	gray	46.7	.305	.46

Calibrating direct navigation. We calibrate the goal-



(a) Egocentric observations on the same direct-navigation route

(b) Route-paired ΔSR of HumiAct

Figure 4. Execution and paired robustness. (a) Recorded simulator frames at 0, 2, 4, 6, and 8 s for GR00T (top) and HumiAct (bottom) on route 04. **(b)** HumiAct minus GR00T (gray) and minus GR00T† (blue) in SR points for each corruption; bars are 95% intervals from resampling the 15 route pairs.

facing routes against three non-learned gait-replay references (Table 4). Blind replay already completes most routes, making SPL and lateral drift more informative than SR alone, while yaw-only heading hold leaves drift unchanged because orientation cannot observe lateral displacement. With camera feedback, HumiAct improves goal arrival and nearly halves lateral drift, matching privileged-pose oracle steering.

Visual feedback during walking. Replacing the image with uniform gray or central occlusion substantially increases lateral drift across learned policies and drops HumiAct below blind gait replay (Table 3a,c). Rather than reverting to a nominal open-loop gait, the policy actively relies on camera feedback to stay on course and is pulled off line when visual information is removed.

Appearance corruptions. HumiAct leads across the five standard corruptions (84.0% mean SR vs. 80.0% for GR00T† and 77.3% for GR00T; Table 3a, Fig. 4b), avoids the stalling that afflicts GR00T under strong defocus (Table 3c), and maintains the highest clean completion across inference seeds (Table 3b). Fig. 5 shows the five inputs.

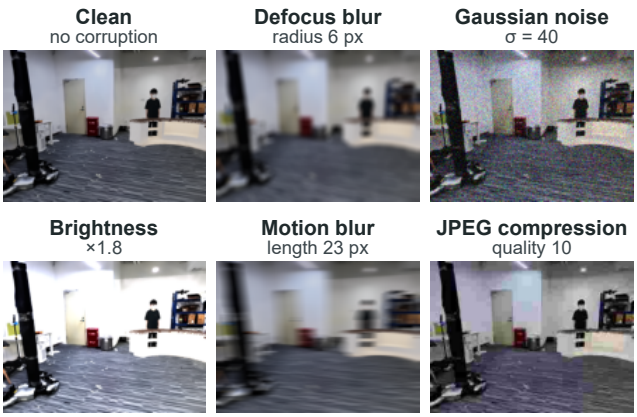


Figure 5. Five standard visual corruptions. Each panel shows the first policy input under the indicated corruption.

Table 5. Latent action-sequence discrimination (%) on 500 balanced held-out clips. *Exact* / *Step*: full-sequence / per-step accuracy (UniAct 95% CI half-widths ≤ 4.3).

	$H{=}1$	$H{=}2$		$H{=}3$	
Condition	Acc.	Exact	Step	Exact	Step
Random chance	20.0	4.0	20.0	0.8	20.0
Train-marginal majority	20.0	20.0	31.8	20.0	39.8
Permuted action labels	3.0	0.7	4.6	0.3	5.2
UniAct	**88.2**	**67.8**	**81.5**	**53.8**	**79.3**

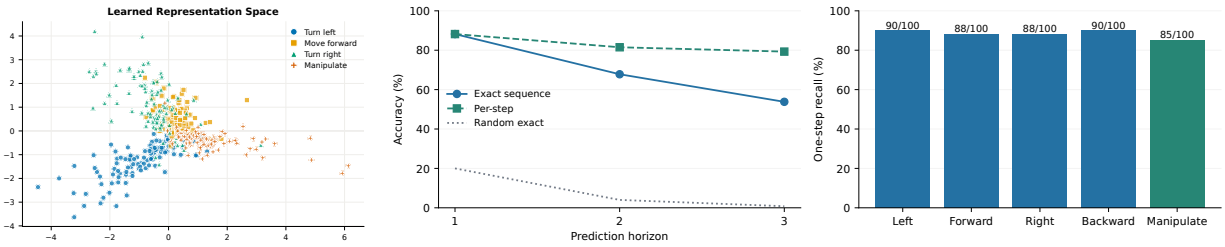
4.4. Learned Representation Space

LAG can only help if the teacher encodes how behaviors change the scene. We test this directly on held-out clips whose source videos and action records are disjoint from training, with the predictor frozen.

Predicted futures identify the action. For each clip, we roll out all candidate action sequences and select the predicted future closest to the observed one. Across the 500 balanced clips, accuracy reaches 88.2% at $H{=}1$ and 67.8%/53.8% on exact two-/three-step sequences (vs. 4.0%/0.8% chance; Table 5), with three-step per-step accuracy near 80% (Fig. 6b). One-step recall spans 85–90% across all five behaviors (Fig. 6c), confirming that minority navigation actions are represented as well as manipulation.

The evidence comes from the action condition. Permuting the action labels at inference, with the predictor unchanged, lowers one-step accuracy to 3.0%, far below the 20% of chance: the predictor commits to the future of the action it is given, so the wrong label systematically selects the wrong future. The representation is therefore organized by action-conditioned dynamics and not by visual statistics of the clips.

Structure of the space. Fig. 6a visualizes the same evidence on the 400 clips of the four behaviors other than backward motion. Candidate argmax selects the true behavior for 359/400 clips (89.75%). After removing the component shared by all candidates of a clip, two principal components



(a) Relative predictive evidence (b) Sequence accuracy over horizons (c) One-step recall per behavior

Figure 6. Learned representation space. (a) Each point is a held-out clip. For every candidate behavior the predictor imagines a future, which is scored by cosine similarity to the observed future; the scores are centered within the clip, standardized, and projected by PCA (two components, 76.2% of the variance). Clips group by their true behavior although labels are used only for coloring. **(b)** Exact-sequence and per-step accuracy against analytic chance. **(c)** One-step recall for each of the five behaviors (100 clips each).

Table 6. Planner choice with the HumiAct policy fixed, on the 15 obstacle routes. SR and collision rate in %.

Planner	SR ↑	SPL ↑	Coll. ↓
GPT-6 Astra	**93.3**	**0.716**	40.0
GPT-6.1 Sol	73.3	0.569	46.7
GPT-5.6 Terra (default)	53.3	0.410	46.7
GPT-5.6 Sol	46.7	0.379	40.0
Qwen3.7-plus	20.0	0.147	13.3
GPT-5.6 Luna	6.7	0.041	**6.7**

explain 76.2% of the variance of the candidate scores, and clips arrange by behavior: left and right turns extend in opposite directions along the second component, manipulation extends along the first, and forward motion lies between them. LAG distills these action-conditioned predictive features into the policy backbone.

4.5. Task Planning

With the HumiAct policy fixed, Table 6 varies only the planner on the obstacle routes. SR ranges from 6.7% to 93.3%: GPT-6 Astra completes 14/15 routes with SPL 0.716, well above the 8/15 of the default planner used in Table 2. The same motor policy therefore completes far more obstacle routes when it is given better instructions, and the planner-navigation numbers in Table 2 are conservative with respect to the planner. Collision rate has to be read together with SR: GPT-5.6 Luna has the fewest contacts because it rarely advances past the obstacle (1/15 routes).

Physical robot. The same planner–policy service and whole-body interface run on the physical humanoid (Sec. 3.4). In planner-free trials, the navigation-adapted policy achieved a higher forward-walking success rate than GR00T. The appendix reports the trial setup and success rates for each command.

4.6. LAG Ablation

With both feature targets, direct-navigation SR is 86.7% with aligned teacher actions, 73.3% without action conditioning, and 46.7% with shuffled actions (Table 7). The current-only variant matches HumiAct on Terra-guided navigation (53.3%) and averages 78.7% across the five corruptions, compared with 84.0% for HumiAct.

Table 7. LAG ablation. Success rate (%). Direct and Terra-guided navigation follow Table 2. The last column averages SR across defocus, Gaussian noise, brightness change, motion blur, and JPEG compression (Table 3a).

Distilled targets	Teacher action	Direct	Terra	Mean
GR00T (none)	–	66.7	33.3	77.3
Current	Aligned	73.3	**53.3**	78.7
Future	None	66.7	46.7	76.0
Current + future	None	73.3	46.7	80.0
Current + future	Shuffled	46.7	33.3	74.6
HumiAct (both)	Aligned	**86.7**	**53.3**	**84.0**

5. Conclusion

HumiAct takes a step toward unified humanoid navigation and manipulation by using one whole-body policy throughout the approach to an object and the ensuing interaction. During training, latent action guidance distills action-conditioned visual predictions from cross-embodiment trajectories; at inference, the policy runs without the predictor. In simulation, HumiAct reached the interaction region more often and completed more navigation–interaction trials than GR00T. On the real humanoid, HumiAct demonstrated coordinated stepping and pushing during chair repositioning.

Whole-body navigation and manipulation is challenging because the approach determines the stance, heading, and viewpoint for the subsequent interaction, so locomotion errors compound into manipulation failures. This coupling motivates a single policy that handles both phases without switching controllers. Future work includes controlled comparisons with modular expert pipelines, evaluation in unseen scenes, and broader real-robot tasks requiring concurrent locomotion and manipulation.

References

[1] AgiBot World Contributors, Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shenyuan Gao, Xindong He, Xuan Hu, Xu Huang, Shu Jiang, Yuxin Jiang, Cheng Jing, Hongyang Li, Jialu Li, Chiming Liu, Yi Liu, Yuxiang Lu, Jianlan Luo, Ping Luo, Yao Mu, Yuehan Niu, Yixuan Pan, Jiangmiao Pang, Yu Qiao, Guanghui Ren, Cheng Ruan, Jiaqi Shan, Yongjian Shen, Chengshi Shi, Mingkang Shi, Modi Shi, Chonghao Sima, Jianheng Song, Huijie Wang, Wenhao Wang, Dafeng Wei, Chengen Xie, Guo Xu, Junchi Yan, Cunbiao Yang, Lei Yang, Shukai Yang, Maoqing Yao, Jia Zeng, Chi Zhang, Qinglin Zhang, Bin Zhao, Chengyue Zhao, Jiaqi Zhao, and Jianchao Zhu. AgiBot World Colosseo: A Large-scale Manipulation Platform for Scalable and Intelligent Embodied Systems. arXiv preprint arXiv:2503.06669, 2025. 4, 5

[2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. arXiv preprint arXiv:2204.01691, 2022. 2

[3] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir R. Zamir. On Evaluation of Embodied Navigation Agents. arXiv preprint arXiv:1807.06757, 2018. 14

[4] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683, 2018. 2

[5] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023. 3

[6] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Mojtaba Komeili, Matthew J. Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, Sergio Arnaud, Abha Gejji, Ada Martin, Francois Robert Hogan, Daniel Dugas, Piotr Bojanowski, Vasil Khalidov, Patrick Labatut, Francisco Massa, Marc Szafraniec, Kapil Krishnakumar, Yong Li, Xiaodong Ma, Sarath Chandar, Franziska Meier, Yann LeCun, Michael Rabbat, and Nicolas Ballas. V-JEPA 2: Self-Supervised Video Models Enable Understanding, Prediction and Planning. arXiv preprint arXiv:2506.09985, 2025. 3, 4, 14

[7] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. Revisiting Feature Prediction for Learning Visual Representations from Video. *Transactions on Machine Learning Research*, 2024. 3

[8] Kevin Black, Noah Brown, Danny Driess, et al. π_0: A Vision-Language-Action Flow Model for General Robot Control. arXiv preprint arXiv:2410.24164, 2024. 2

[9] Anthony Brohan, Noah Brown, Justice Carbajal, et al. RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems*, 2023. 2

[10] Jake Bruce, Michael D. Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Maria Elisabeth Bechtle, Feryal Behbahani, Stephanie C. Y. Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando de Freitas, Satinder Singh, and Tim Rocktäschel. Genie: Generative Interactive Environments. In *Proceedings of the 41st International Conference on Machine Learning*, pages 4603–4623, 2024. 3

[11] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *IEEE/CVF International Conference on Computer Vision*, pages 9630–9640, 2021. 3

[12] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History Aware Multimodal Transformer for Vision-and-Language Navigation. In *Advances in Neural Information Processing Systems*, pages 5834–5847, 2021. 2

[13] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. In *Robotics: Science and Systems*, 2023. 3

[14] Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying Generalization in Reinforcement Learning. In *International Conference on Machine Learning*, pages 1282–1289, 2019. 3

[15] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. RoboNet: Large-Scale Multi-Robot Learning. In *Proceedings of the Conference on Robot Learning*, pages 885–897. PMLR, 2020. 3

[16] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. PaLM-E: An Embodied Multimodal Language Model. In *International Conference on Machine Learning*, pages 8469–8488, 2023. 2

[17] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall/CRC, 1993. 14

[18] R. A. Fisher. On the Interpretation of χ^2 from Contingency

Tables, and the Calculation of P. *Journal of the Royal Statistical Society*, 85(1):87–94, 1922. 14

[19] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-Follower Models for Vision-and-Language Navigation. In *Advances in Neural Information Processing Systems*, pages 3318–3329, 2018. 2

[20] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. HumanPlus: Humanoid Shadowing and Imitation from Humans. In *Conference on Robot Learning*, pages 2828–2844, 2024. 1, 2

[21] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation. In *Conference on Robot Learning*, pages 4066–4083, 2024. 2

[22] Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Quan Vuong, Ted Xiao, Pannag R. Sanketi, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An Open-Source Generalist Robot Policy. In *Robotics: Science and Systems*, 2024. 2

[23] Xinyang Gu, Yen-Jen Wang, and Jianyu Chen. Humanoid-Gym: Reinforcement Learning for Humanoid Robot with Zero-Shot Sim2Real Transfer. arXiv preprint arXiv:2404.05695, 2024. 2

[24] David Ha and Jürgen Schmidhuber. World Models. arXiv preprint arXiv:1803.10122, 2018. 3

[25] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning Latent Dynamics for Planning from Pixels. In *International Conference on Machine Learning*, pages 2555–2565, 2019.

[26] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to Control: Learning Behaviors by Latent Imagination. In *International Conference on Learning Representations*, 2020.

[27] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering Atari with Discrete World Models. In *International Conference on Learning Representations*, 2021. 3

[28] Mingfei Han, Haihong Hao, Liang Ma, Kamila Zhumakhanova, Ekaterina Radionova, Jingyi Zhang, Xiaojun Chang, Xiaodan Liang, and Ivan Laptev. Implicit Geometry Representations for Vision-and-Language Navigation from Web Videos. arXiv preprint arXiv:2603.09259, 2026. 2

[29] Haihong Hao, Lei Chen, Mingfei Han, Changlin Li, Dong An, Yuqiang Yang, Zhihui Li, and Xiaojun Chang. LatentPilot: Scene-Aware Vision-and-Language Navigation by Dreaming Ahead with Latent Visual Reasoning. In *European Conference on Computer Vision*, pages 75–93, 2026. 3

[30] Haihong Hao, Lei Chen, Mingfei Han, Chenglong Yan, Dong An, Yuqiang Yang, Changlin Li, Minghao Guo, Salman Khan, and Xiaojun Chang. Seeing Together, Acting Apart: Shared Environmental Understanding for Multi-Robot Navigation. In *Advances in Neural Information Processing Systems*, 2026. 3

[31] Haihong Hao, Mingfei Han, Changlin Li, Zhihui Li, and Xiaojun Chang. CoNav: Collaborative Cross-Modal Reasoning for Embodied Navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(11):14239–14253, 2026. 2

[32] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders Are Scalable Vision Learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15979–15988, 2022. 3

[33] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. OmniH2O: Universal and Dexterous Human-to-Humanoid Whole-Body Teleoperation and Learning. In *Conference on Robot Learning*, pages 1516–1540, 2024. 2

[34] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning Human-to-Humanoid Real-Time Whole-Body Teleoperation. arXiv preprint arXiv:2403.04436, 2024. 2

[35] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Jan Kautz, Changliu Liu, Guanya Shi, Xiaolong Wang, Linxi Fan, and Yuke Zhu. HOVER: Versatile Neural Whole-Body Controller for Humanoid Robots. In *IEEE International Conference on Robotics and Automation*, pages 9989–9996, 2025. 2

[36] Dan Hendrycks and Thomas Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. In *International Conference on Learning Representations*, 2019. 6, 18

[37] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. arXiv preprint arXiv:1503.02531, 2015. 4

[38] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez Opazo, and Stephen Gould. VLN BERT: A Recurrent Vision-and-Language BERT for Navigation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1643–1653, 2021. 2

[39] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Noah Brown, Tomas Jackson, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. Inner Monologue: Embodied Reasoning through Planning with Language Models. In *Conference on Robot Learning*, 2022. 2

[40] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. VoxPoser: Composable 3D Value Maps for Robotic Manipulation with Language Models. In *Conference on Robot Learning*, pages 540–562, 2023. 2

[41] 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*, 4(26):eaau5872, 2019. 2

[42] Youngjoon Jeong, Junha Chun, and Taesup Kim. Learning to Act Robustly with View-Invariant Latent Actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6781–6790, 2026. 3

[43] Haoran Jiang, Jin Chen, Qingwen Bu, Li Chen, Modi Shi, Yanjie Zhang, Delong Li, Chuanzhe Suo, Chuang Wang,

Zhihui Peng, and Hongyang Li. WholeBodyVLA: Towards Unified Latent VLA for Whole-Body Loco-Manipulation Control. In *International Conference on Learning Representations*, 2026. 3

[44] Tao Jiang, Tianyuan Yuan, Yicheng Liu, Chenhao Lu, Jianning Cui, Xiao Liu, Shuiqi Cheng, Jiyang Gao, Huazhe Xu, and Hang Zhao. Galaxea Open-World Dataset and G0 Dual-System VLA Model. arXiv preprint arXiv:2509.00576, 2025. 4, 5

[45] Alexander Khazatsky, Karl Pertsch, Suraj Nair, et al. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. arXiv preprint arXiv:2403.12945, 2024. 3

[46] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An Open-Source Vision-Language-Action Model. In *Conference on Robot Learning*, pages 2679–2713, 2024. 3

[47] Ilya Kostrikov, Denis Yarats, and Rob Fergus. Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels. In *International Conference on Learning Representations*, 2021. 3

[48] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments. In *European Conference on Computer Vision*, pages 104–120, 2020. 2

[49] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-Across-Room: Multilingual Vision-and-Language Navigation with Dense Spatiotemporal Grounding. In *Conference on Empirical Methods in Natural Language Processing*, pages 4392–4412, 2020. 2

[50] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement Learning with Augmented Data. In *Advances in Neural Information Processing Systems*, 2020. 3

[51] Seungjae Lee, Yibin Wang, Haritheja Etukuru, H. Jin Kim, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Behavior Generation with Latent Actions. In *Proceedings of the 41st International Conference on Machine Learning*, pages 26991–27008, 2024. 3

[52] Anqi Li, Yuxin Chen, Zhaobo Li, Zhuo Cao, Junli Ren, Masayoshi Tomizuka, and Dhruv Shah. TANGO: Humanoid Navigation in Cluttered Environments with a Whole-Body Vision-Language-Action Model. arXiv preprint arXiv:2609.09158, 2026. 3

[53] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as Policies: Language Model Programs for Embodied Control. In *IEEE International Conference on Robotics and Automation*, pages 9493–9500, 2023. 2

[54] 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. *Science Robotics*, 11(117): eadx8924, 2026. 2, 5, 14

[55] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling. In *International Conference on Learning Representations*, 2023. 3, 14

[56] Peiqi Liu, Yaswanth Orru, Jay Vakil, Chris Paxton, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. OK-Robot: What Really Matters in Integrating Open-Knowledge Models for Robotics. arXiv preprint arXiv:2401.12202, 2024. 1, 2

[57] Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Fernando Castañeda, Sirui Chen, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, Jinhyung Park, David Sami, Zi Wang, Xingye Da, Runyu Ding, Cyrus Hogg, Lina Song, Edy Lim, Eugene Jeong, Tairan He, Haoru Xue, Wenli Xiao, Simon Yuen, Jan Kautz, Yan Chang, Umar Iqbal, Linxi Jim Fan, and Yuke Zhu. SONIC: Supersizing Motion Tracking for Natural Humanoid Whole-Body Control. *Science Robotics*, 11(117):eaed4592, 2026. 2

[58] Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning Latent Plans from Play. In *Proceedings of the Conference on Robot Learning*, pages 1113–1132, 2020. 3

[59] Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Yecheng Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Tingfan Wu, Jay Vakil, Pieter Abbeel, Jitendra Malik, Dhruv Batra, Yixin Lin, Oleksandr Maksymets, Aravind Rajeswaran, and Franziska Meier. Where are we in the search for an Artificial Visual Cortex for Embodied Intelligence? In *Advances in Neural Information Processing Systems*, 2023. 3

[60] Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947. 14

[61] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7(62), 2022. 2

[62] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3M: A Universal Visual Representation for Robot Manipulation. In *Proceedings of the 6th Conference on Robot Learning*, pages 892–909, 2023. 3

[63] NVIDIA. Isaac Lab: A GPU-Accelerated Simulation Framework for Multi-Modal Robot Learning. arXiv preprint arXiv:2511.04831, 2025. 5

[64] NVIDIA, Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Jim Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang, Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu, Edith Llontop, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed, You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie, Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao, Ruijie Zheng, and Yuke Zhu. GR00T N1: An Open Foundation Model for Generalist Humanoid Robots. arXiv preprint arXiv:2503.14734, 2025. 1, 3, 5, 6, 14

[65] Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. arXiv preprint arXiv:2310.08864, 2023. 3

[66] Maxime Oquab, Timothée Darcet, Théo Moutakanni, et al. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*, 2024. 3

[67] 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*, 37(4):143, 2018. 2

[68] 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*, 40(4):144:1–144:20, 2021.

[69] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. ASE: Large-Scale Reusable Adversarial Skill Embeddings for Physically Simulated Characters. *ACM Transactions on Graphics*, 41(4):94:1–94:17, 2022. 2

[70] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*, pages 8748–8763, 2021. 3

[71] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-World Robot Learning with Masked Visual Pre-training. In *Proceedings of The 6th Conference on Robot Learning*, pages 416–426. PMLR, 2023. 3

[72] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for Thin Deep Nets. In *International Conference on Learning Representations*, 2015. 4

[73] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning. In *Proceedings of the 5th Conference on Robot Learning*, pages 91–100. PMLR, 2022. 2

[74] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9339–9347, 2019. 2

[75] Nur Muhammad Mahi Shafiullah, Zichen Jeff Cui, Ariuntuya Altanzaya, and Lerrel Pinto. Behavior Transformers: Cloning k modes with one stone. In *Advances in Neural Information Processing Systems*, 2022. 3

[76] Dhruv Shah, Ajay Sridhar, Arjun Bhorkar, Noriaki Hirose, and Sergey Levine. GNM: A General Navigation Model to Drive Any Robot. In *IEEE International Conference on Robotics and Automation*, pages 7226–7233, 2023. 2

[77] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. ViNT: A Foundation Model for Visual Navigation. In *Conference on Robot Learning*, pages 711–733, 2023.

[78] Ajay Sridhar, Dhruv Shah, Catherine Glossop, and Sergey Levine. NoMaD: Goal Masked Diffusion Policies for Navigation and Exploration. In *IEEE International Conference on Robotics and Automation*, pages 63–70, 2024. 2

[79] Yang Tian, Sizhe Yang, Jia Zeng, Ping Wang, Dahua Lin, Hao Dong, and Jiangmiao Pang. Predictive Inverse Dynamics Models are Scalable Learners for Robotic Manipulation. In *International Conference on Learning Representations*, 2025. 3

[80] 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 *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 23–30, 2017. 3

[81] Homer Rich Walke, Kevin Black, Tony Z. Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, Abraham Lee, Kuan Fang, Chelsea Finn, and Sergey Levine. BridgeData V2: A Dataset for Robot Learning at Scale. In *Proceedings of The 7th Conference on Robot Learning*, pages 1723–1736. PMLR, 2023. 3

[82] Chen Wang, Linxi Fan, Jiankai Sun, Ruohan Zhang, Li Fei-Fei, Danfei Xu, Yuke Zhu, and Anima Anandkumar. MimicPlay: Long-Horizon Imitation Learning by Watching Human Play. In *Proceedings of the 7th Conference on Robot Learning*, pages 201–221, 2023. 3

[83] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 2

[84] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: Learning Near-Perfect PointGoal Navigators from 2.5 Billion Frames. In *International Conference on Learning Representations*, 2020. 2

[85] Frank Wilcoxon. Individual Comparisons by Ranking Methods. *Biometrics Bulletin*, 1(6):80–83, 1945. 14

[86] Songwei Wu, Rui Zhao, Fan Yang, Zhongqiang Nie, Zhiduo Jiang, Wandong Sun, Yuwei Li, Jian Hu, Yang Liu, and Hong Liu. EATR-Stereo: Embodiment-Aware Token Routing of Paired Stereo Evidence for Humanoid Vision-Language-Action Control. arXiv preprint arXiv:2608.17453, 2026. 3, 5

[87] Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Sejune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Liden, Kimin Lee, Jianfeng Gao, Luke Zettlemoyer, Dieter Fox, and Minjoon Seo. Latent Action Pretraining from Videos. In *International Conference on Learning Representations*, 2025. 3

[88] Sergey Zagoruyko and Nikos Komodakis. Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *International Conference on Learning Representations*, 2017. 4

[89] Chiyuan Zhang, Oriol Vinyals, Remi Munos, and Samy Bengio. A Study on Overfitting in Deep Reinforcement Learning. arXiv preprint arXiv:1804.06893, 2018. 3

[90] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Robotics: Science and Systems*, 2023. 3

[91] Jia Zheng, Teli Ma, Yudong Fan, Zifan Wang, Shuo Yang, and Junwei Liang. MotionWAM: Towards Foundation World Action Models for Real-Time Humanoid Loco-Manipulation. arXiv preprint arXiv:2606.09215, 2026. 3

[92] Brianna Zitkovich, Tianhe Yu, Sichun Xu, et al. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. In *Proceedings of The 7th Conference on Robot Learning*, pages 2165–2183. PMLR, 2023. 3

Beyond Decoupled Control: Rethinking Whole-Body Navigation for Humanoid Robots

Supplementary Material

This appendix provides reproducibility details, route-paired statistical tests, root-cause analyses of anomalous metrics and failure modes, trajectory and corruption visualizations, UniAct/LAG ablation specifications, physical-robot trial details, and a summary of qualitative simulation and real-humanoid rollouts. Every metric, hyperparameter, and symbol follows the main paper.

A. Reproducibility, Compute, and Evaluation Protocols

Claim anchor. *Supports Main Sec. 3.4 and Sec. 4.1 (Tables 1–3). Takeaway:* All learned policies share the same humanoid corpus, downstream fine-tuning budget, and whole-body controller with zero inference-time overhead from LAG, while paired evaluation protocols maximize statistical resolution under hardware-aligned trial budgets.

Baseline Fairness and Tuning Budget. Table S1 summarizes architecture, data, optimization, compute, and seed configurations across all three training stages. All compared whole-body policies—GR00T [64], GR00T†, the navigation-only specialist (Nav.+LAG), the manipulation-only specialist, and HumiAct—share the same 84-directory whole-body demonstration corpus (5,421 episodes: 49 original navigation/manipulation directories totaling 2,205 episodes plus 35 converted directories totaling 3,216 episodes), the same 200k-step fine-tuning schedule (global batch size 96, learning rate 10^{-4}, weight decay 10^{-5}, 5% linear warmup, state dropout 0.2), the same flow-matching action transformer f_θ [55], and the same 50 Hz BeyondMimic whole-body tracker [54]. GR00T† additionally matches HumiAct's exact downstream data order and optimization recipe from the public GR00T base model. All hyperparameters are inherited from the base GR00T and V-JEPA 2 [6] recipes without per-method tuning on evaluation routes. Because the V-JEPA-AC teacher (E_ϕ, P_ψ), the 32 auxiliary alignment tokens, the poolers $(Q_{\text{cur}}, Q_{\text{fut}})$, and the projector G are discarded prior to deployment, HumiAct incurs zero inference latency or parameter overhead relative to GR00T.

Hardware-Aligned Trial Budgets and Paired Statistical Design. On a full-size 33-DoF physical humanoid, each closed-loop trial requires full workspace reset, battery and thermal monitoring, safety spotting, and balance stabilization. We intentionally designed the evaluation suite from the outset with a compact, hardware-feasible trial structure (N=15 routes for direct and planner navigation, N=32 starts for composed navigation–interaction, and N=3 repetitions per primitive on hardware) so that simulation and physical-robot evaluation remain structurally aligned under one shared control interface. To maintain statistical rigor under compact N, all simulation model comparisons use *route-paired evaluation*: compared checkpoints share the exact start pose, goal position, obstacle layout, environment seed, and per-request policy noise seed $42{+}10{,}000e{+}q$ for request q of episode e. Gaussian image noise draws from a separate deterministic stream indexed by episode, request, batch, and step. We quantify uncertainty via 10,000-draw route-paired percentile bootstrap 95% confidence intervals [17] (seed 20261111), exact two-sided McNemar tests [60] on discordant success pairs, paired Wilcoxon signed-rank tests [85] on continuous SPL [3], Fisher exact tests [18] on unpaired composed counts, and three-seed clean repeats (seeds 42–44, 45 rollouts/policy).

Metrics and Protocol Separation. Across N trials with binary goal indicators $S_i \in \{0,1\}$, positive reference start–goal distance $d_i > 0$ ($2.20 \le d_i \le 4.30$ m on direct navigation), and traveled path length L_i, we compute

$$\text{SR} = \frac{1}{N}\sum_{i=1}^{N} S_i, \quad \text{SPL} = \frac{1}{N}\sum_{i=1}^{N} S_i \frac{d_i}{\max(d_i, L_i)}. \quad (9)$$

On direct navigation (goal-aligned start, forward-motion instruction, 10 s horizon, 0.5 m goal radius, no obstacles), L_i spans the full 10 s episode and d_i is Euclidean start–goal distance; a route can therefore enter the goal radius and overshoot slightly afterward. On planner navigation (15 obstacle routes), L_i terminates at first goal arrival and obstacle collision rate (Coll.) records the marginal fraction of trials with any contact exceeding 10 N; a trial may incur a light contact and still reach the goal. On manipulation only, mean successful control steps are averaged over completed trials alongside SR. On composed navigation–interaction (N=32), letting $A_i, B_i \in \{0,1\}$ denote interaction-region arrival and object-task completion, we report Reach ($N^{-1}\sum_i A_i$) and Complete ($N^{-1}\sum_i A_i B_i$). In simulation, policies predict 40-step chunks at 20 Hz and execute the first 30 steps (1.5 s) with zero overlap.

Crucially, the whole-body comparison (Main Table 2) and the visual-intervention suite (Main Tables 3–4) use two distinct simulator reset entry points that are never cross-compared: Table 2 uses the standard environment reset, whereas Tables 3–4 execute two extra kinematic/camera

Table S1. Reproducibility, compute, and baseline-fairness summary across all three training stages and evaluation protocols. Every hyperparameter, episode count, and protocol parameter matches Main Sec. 3.4 and Sec. 4.1. Teacher weights, auxiliary tokens, poolers ($Q_{\text{cur}}, Q_{\text{fut}}$), and projector G are discarded after Stage 2, incurring zero inference overhead.

Component	Stage 1: V-JEPA-AC Prior	Stage 2: LAG Alignment	Stage 3: Whole-Body Fine-Tuning
Base / frozen modules	Frozen ViT-giant E_ϕ (384×384, patch 16, 576 tokens, width 1408)	Frozen teacher (E_ϕ, P_ψ); GR00T N1 base; 32 aux. tokens (16 cur + 16 fut)	LAG-aligned backbone + action modules (GR00T/GR00T† start from public base)
Trainable modules	24-layer causal P_ψ (width 1024, 16 heads, RoPE, 7-D zero-padded $\mathbf{u}_t$)	Selected LLM layers (3rd-from-last, width 2048), shared G, poolers $Q_{\text{cur}}, Q_{\text{fut}}$	Full VLA backbone, state/action projections, flow-matching action transformer f_θ
Training corpus	178,384 episodes: AgiBotWorld (142,069), Galaxea (34,353), Humanoid (1,962)	Same 178,384-episode heterogeneous corpus (video-/action-/parquet-disjoint split)	5,421 humanoid whole-body demos across 84 directories (49 nav./manip. + 35 converted)
Objective	$\mathcal{L}_{\text{AC}}$: normalized ℓ_1 (teacher-forced + K=2 AR steps; 4 frames @ 2 Hz)	$\mathcal{L}_{\text{LAG}}$: token-wise cosine distance (λ_{cur}=λ_{fut}=1.0; 576→16 pooled)	$\mathcal{L}_{\text{FM}}$: masked flow matching; 40-step chunk (30 exec. @ 20 Hz sim; 30 @ 10 Hz real)
Optimization	5,000 iters (10 ep. × 500), peak LR $1.5\times10^{-4} \to 10^{-6}$, 2-ep. warmup, wd 0.04	LR 10^{-4}, wd 3×10^{-5}, warmup 5%, color jitter; discarded before inference	200k steps, batch 96, LR 10^{-4}, wd 10^{-5}, warmup 5%, state dropout 0.2 (shared by GR00T/GR00T†)
Compute & precision	16 NVIDIA H20 GPUs, BF16; action weights $[2, 2, 2, 1, 2]$	16 NVIDIA H20 GPUs, BF16	16 NVIDIA H20 GPUs, BF16; deploy on RTX 5000 Ada (4 flow steps, 10/50 Hz)
Seeds & tuning budget	Seed 20260916; final ckpt. (5k iters); PCA seed 20260923; no test tuning	Inherited recipe; identical downstream budget for all LAG ablations (Table 7)	Eval seeds 42–44 ($42{+}10{,}000e{+}q$); bootstrap 10,000 draws (seed 20261111)

render passes immediately after reset (without advancing simulation time or controller history) so the first policy frame reflects the freshly reset pose before image corruption is injected.

B. Statistical Analysis, Trajectory Atlas, and Visual Corruptions

Claim anchor. *Supports Main Sec. 4.2–4.3 (Tables 2–4, Figs. 4–5). Takeaway:* Route-paired intervals confirm systematic SPL gains on direct navigation, statistically significant elimination of strong-defocus stalling vs. GR00T (p=0.031), tighter corridor tracking ($1/15$ vs. $4/15$ excursions > 0.5 m), and statistically significant degradation under gray-image removal (p=0.041), alongside exact pixel-space corruption definitions.

Paired Hypothesis Tests and Multi-Seed Variance. Table S2a reports paired ΔSR, ΔSPL, 95% bootstrap intervals, discordant pair counts (gained/lost), and exact non-parametric tests. On standard-reset direct navigation (Main Table 2), HumiAct gains 3 routes over GR00T and 2 over the prior-initialized reference while losing none, with strictly positive 95% bootstrap CIs on ΔSPL ($+0.188\,[+0.016, +0.391]$ vs. GR00T; $+0.165\,[+0.014, +0.349]$ vs. prior reference). Across the five standard corruptions (75 paired route–condition trials), HumiAct leads all baselines in mean SR (84.0% vs. 80.0% for GR00T† and 77.3% for GR00T and Nav.+LAG), with a strictly positive 95% ΔSPL interval over Nav.+LAG ($+0.069\,[+0.004, +0.142]$), isolating the benefit of joint navigation–manipulation training given identical LAG initialization. Across inference seeds 42–44 on clean reset-refreshed routes (Table S2b; 45 rollouts/policy), HumiAct achieves the highest mean completion ($12.7 \pm 1.2/15$, $84.4 \pm 7.7\%$) and mean SPL (0.645 ± 0.064). Figure S1 visualizes all four policies alongside the complete route-level SPL matrix across all seven corruption conditions.

Task Decomposition and Manipulation Retention. Table S3a decomposes the 32 composed navigation–interaction trials into mutually exclusive outcomes: HumiAct increases interaction-region arrival by $+12.5$ pp ($18/32$ vs. $14/32$) and full end-to-end completion by $+9.4$ pp ($5/32 - 2/32 = 3/32$). Among trials that reach the interaction region, HumiAct nearly doubles conditional object completion ($5/18 = 27.8\%$ vs. $2/14 = 14.3\%$), showing that whole-body approach quality carries directly into manipulation readiness. On manipulation-only evaluation (Table S3b), joint training retains manipulation accuracy (36.0% vs. 32.0% for the specialist and 33.0% for GR00T in Main Table 2) with only a 15-step difference in mean completion length (566 vs. 551 steps; on hardware, the separate manipulation-only baseline reaches 35.0%).

Root-Cause Analysis of Anomalous Metrics and Failure Modes. Several non-monotonic numbers in the main tables warrant explicit mechanistic explanation so they are not misread as contradictions:

1. *Larger SR drop of HumiAct under gray image and central occlusion (46.7% vs. 73.3% for GR00T; Table 3a):* Under a uniform gray image, HumiAct suffers a statistically significant drop relative to GR00T (ΔSR $= -26.7$ pp, 95% CI $[-46.7, -6.7]$; ΔSPL Wilcoxon p=0.041) while its lateral deviation more than doubles ($0.21 \to 0.46$ m). This is expected behavior

Table S2. Route-paired statistical analysis and multi-seed clean variance. **(a)** Paired SR (ΔSR, pp) and SPL (ΔSPL) differences with 95% route-bootstrap CIs (10,000 draws, seed 20261111), gained/lost route counts (G/L), exact two-sided McNemar p-values (pooled over 75 paired route–condition trials for the five-corruption mean), and paired Wilcoxon signed-rank p-values on SPL. Bold CIs exclude zero. Under strong defocus, GR00T stalls on 6/15 routes vs. 0/15 for HumiAct (exact McNemar p=0.031). **(b)** Clean direct-navigation completions (out of 15) and SPL across inference seeds 42–44 (45 rollouts/policy), matching Main Table 3b.

(a) Route-paired hypothesis tests and 95% bootstrap confidence intervals

Comparison	ΔSR	95% CI	G/L	p_{McN}	ΔSPL	95% CI	p_{Wilc}
Direct navigation, standard reset (Main Table 2; 15 paired routes)							
HumiAct − GR00T	+20.0	[0.0, 40.0]	3/0	0.250	+0.188	**[0.016, 0.391]**	0.136
HumiAct − Prior-init. ref.	+13.3	[0.0, 33.3]	2/0	0.500	+0.165	**[0.014, 0.349]**	0.221
Five-corruption mean, reset-refreshed (Main Table 3a; 15 routes × 5 corruptions = 75 paired trials)							
HumiAct − GR00T	+6.7	[−6.7, 24.0]	8/3	0.227	+0.026	[−0.094, 0.168]	0.594
HumiAct − GR00T†	+4.0	[−2.7, 10.7]	5/2	0.453	+0.000	[−0.057, 0.060]	0.917
HumiAct − Nav.+LAG	+6.7	[−1.3, 16.0]	7/2	0.180	+0.069	**[0.004, 0.142]**	0.087
Strong corruptions and visual-information removal, reset-refreshed (seed 42; 15 paired routes)							
Strong defocus (r=20) vs. GR00T	+26.7	[−6.7, 60.0]	6/2	0.289	+0.199	[−0.092, 0.475]	0.311
Strong defocus (r=20) vs. GR00T†	+0.0	[−20.0, 20.0]	1/1	1.000	−0.024	[−0.213, 0.165]	0.433
Strong noise (σ=100) vs. GR00T	+20.0	[0.0, 40.0]	3/0	0.250	+0.110	[−0.035, 0.282]	0.695
Gray image vs. GR00T	−26.7	**[−46.7, −6.7]**	0/4	0.125	−0.190	**[−0.351, −0.048]**	0.041
Central occlusion vs. GR00T	−20.0	[−46.7, 6.7]	1/4	0.375	−0.131	[−0.347, 0.084]	0.422

(b) Clean reset-refreshed direct navigation across inference seeds (completions / 15 and SPL)

Policy	Seed 42	Seed 43	Seed 44	Mean ± SD (/15)	Mean SR ± SD (%)	Mean SPL ± SD
GR00T	12 (0.606)	11 (0.583)	11 (0.549)	11.3 ± 0.6	75.6 ± 3.8	0.579 ± 0.029
GR00T†	13 (0.669)	10 (0.520)	12 (0.609)	11.7 ± 1.5	77.8 ± 10.2	0.599 ± 0.075
Nav.+LAG	13 (0.674)	11 (0.576)	13 (0.641)	12.3 ± 1.2	82.2 ± 7.7	0.630 ± 0.050
HumiAct	14 (0.719)	12 (0.612)	12 (0.605)	12.7 ± 1.2	84.4 ± 7.7	0.645 ± 0.064

Table S3. Task decomposition, manipulation retention, and corridor excursion (standard whole-body protocol, supporting Main Sec. 4.2 and Table 2). **(a)** Mutually exclusive outcomes on the 32 composed navigation–interaction trials; end-to-end gain is $3/32 = +9.4$ pp (two-sided Fisher exact: Reach p=0.454, Full p=0.426, Complete | Reach p=0.426). **(b)** Manipulation-only evaluation (mean steps over successful trials; GR00T achieves 33.0% SR in Main Table 2, and the separate real-robot manipulation baseline reaches 35.0%). **(c)** Pre-arrival lateral corridor excursion E_i (Eq. 11) across the 15 direct-navigation routes.

(a) Composed task (N=32)

Method	No arr.	Arr. only	Full	Full | Arr.
GR00T	18	12	2	14.3% (2/14)
HumiAct	14	13	5	27.8% (5/18)

(b) Manipulation

Policy	SR (%)	Steps
Manip.-only	32.0	551
HumiAct	36.0	566

(c) Direct-nav. excursion

Policy	Succ.	$E_i > 0.5$ m	Fall
GR00T	10/15	4/15	1/15
Prior ref.	11/15	0/15	1/15
HumiAct	13/15	1/15	1/15

of closed-loop visual control, not fragility: because the direct-navigation routes start goal-aligned, even a blind open-loop gait replay achieves 80.0% SR with 0.39 m drift (Main Table 4). A policy that falls back on proprioceptive open-loop walking when images vanish retains high SR on straight routes, whereas HumiAct actively steers via visual feedback (achieving 0.21 m clean drift, matching the 0.23 m privileged oracle) and is therefore pulled off line when visual information is stripped completely.

2. *Single-condition fluctuations under moderate noise (σ=40: 80.0% vs. 86.7%) and defocus (r=6: 80.0% vs. 86.7% for GR00T†):* On N=15 routes, one route equals 6.7 pp. Figure S1b confirms that both gaps stem from a single route flip (Route 9 under σ=40; Route 11 under r=6; exact McNemar p=1.0), whereas HumiAct wins on Brightness (+13.3/ + 20.0 pp), Motion blur (+13.3/ +6.7 pp), JPEG (+13.3/ +6.7 pp), Strong noise (+20.0/ + 6.7 pp), and Clean (+13.3/ + 6.7 pp). Under strong defocus (r=20), GR00T stalls (progress < 0.5 m) on 6/15 routes vs. 0/15 for HumiAct (exact McNemar p=0.031). Notably, GR00T† also avoids stalling (73.3%, 0 stalls) and shows a smaller relative SR drop on the two strong corruptions (15.4% vs. 17.9% for HumiAct and 33.3% for GR00T, computed via $D_{m,c} = (\text{SR}_{m,\text{clean}} - \text{SR}_{m,c})/\text{SR}_{m,\text{clean}}$). Thus, matching HumiAct's downstream training recipe largely eliminates catastrophic freezing, while LAG provides additional gains in clean lateral tracking (0.21 m vs. 0.36 m), five-corruption mean SR (+4.0 pp), and strong noise (+6.7 pp).
3. *Higher collision rate accompanying higher SR on planner navigation (46.7% vs. 26.7% for GR00T; 40.0% for GPT-6 Astra vs. 6.7% for GPT-5.6 Luna):* Collision rate is a marginal count of > 10 N contact events across 15 routes and is coupled to obstacle exposure. Policies or

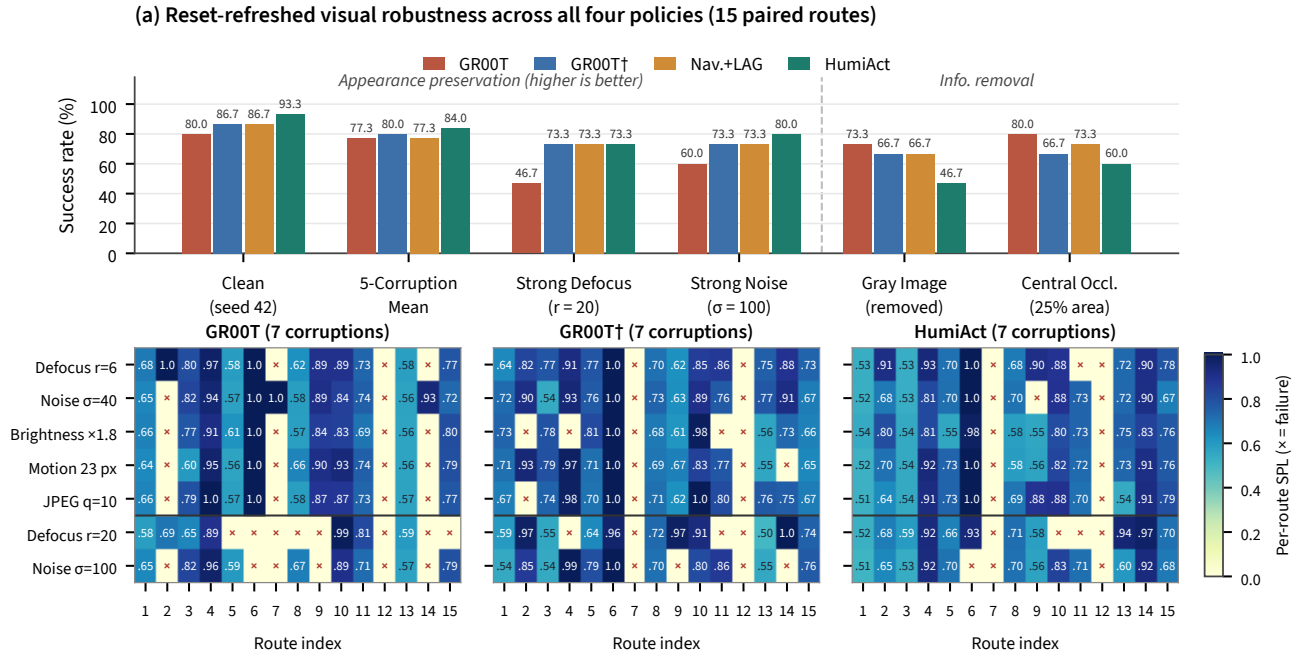


Figure S1. Four-policy visual robustness and complete per-route corruption breakdown (reset-refreshed protocol, seed 42, 15 paired routes). **(a)** Success rate across all four policies on clean input, the five-corruption mean, stronger corruptions, and visual-information removal (matching Main Table 3a). **(b)** Complete route-by-condition SPL heatmaps across all seven appearance corruptions for GR00T, GR00T†, and HumiAct ($\times$ denotes non-arrival). No route is omitted.

planners that stall or turn away before reaching the obstacle rarely trigger contact (e.g., GPT-5.6 Luna completes $1/15$ routes with $1/15$ collisions), whereas trajectories that actively navigate past narrow clearances accumulate brief contacts even while reaching the goal. Consequently, SR, SPL, and Coll. must always be interpreted jointly (Fig. S5a).

4. *Non-learned gait-replay calibration (80.0% SR on goal-facing routes):* The three non-learned references in Main Table 4 loop one steady 33-step (1.65 s) forward-walking gait cycle resampled to 20 Hz through the same whole-body interface. Open-loop gait and IMU-style heading hold both reach $12/15$ goals and fail solely on the three longest routes (3.49, 4.25, and 4.30 m), which exceed their 10 s travel range; privileged oracle steering (replanning every 0.5 s from ground-truth pose) misses only the two longest (4.25 and 4.30 m). This calibration clarifies scope: goal-facing straight routes primarily test closed-loop lateral tracking (where HumiAct cuts drift nearly in half, $0.39 \rightarrow 0.21$ m) and visual robustness, whereas obstacle routes and composed tasks require closed-loop perception, planning, and loco-manipulation that gait replay cannot perform.
5. *Modest absolute completion on composed navigation–interaction (5/32 = 15.6%):* End-to-end loco-manipulation compounds errors across long horizons: small heading or foot-placement deviations at arrival shift the object out of the reliable grasp envelope (13/18 arrivals fail during interaction). This reflects genuine task difficulty rather than policy collapse, as confirmed by the nearly $2\times$ higher conditional completion rate over GR00T (27.8% vs. 14.3%), and identifies close-range stance alignment as a key direction for future work.

Route-Aligned Coordinates and Corridor Excursion. For route i with planar start $\mathbf{p}_i^{\text{start}} \in \mathbb{R}^2$ and goal $\mathbf{p}_i^{\text{goal}} \in \mathbb{R}^2$ ($d_i = \|\mathbf{p}_i^{\text{goal}} - \mathbf{p}_i^{\text{start}}\|_2 > 0$), let $\hat{\mathbf{e}}_i^{\parallel} = (\mathbf{p}_i^{\text{goal}} - \mathbf{p}_i^{\text{start}})/d_i = (e_{i,x}^{\parallel}, e_{i,y}^{\parallel})^{\top}$ and $\hat{\mathbf{e}}_i^{\perp} = (-e_{i,y}^{\parallel}, e_{i,x}^{\parallel})^{\top}$. Each world-frame root position $\mathbf{p}_i(t)$ is transformed into along-track and lateral coordinates

$$\begin{aligned} x_i(t) &= \left(\mathbf{p}_i(t) - \mathbf{p}_i^{\text{start}}\right)^{\top} \hat{\mathbf{e}}_i^{\parallel}, \\ y_i(t) &= \left(\mathbf{p}_i(t) - \mathbf{p}_i^{\text{start}}\right)^{\top} \hat{\mathbf{e}}_i^{\perp}. \end{aligned} \tag{10}$$

Letting $t_i^{\star}$ denote first goal entry (or the final episode step if unsuccessful), the pre-arrival maximum lateral corridor excursion is

$$E_i = \max_{t \le t_i^{\star}} |y_i(t)|. \tag{11}$$

Figure S2 displays all 15 routes under this canonical frame, and Table S3c reports corridor departures ($E_i > 0.5$ m): GR00T exceeds the 0.5 m corridor on $4/15$ routes, the prior-initialized reference on $0/15$, and HumiAct on $1/15$, with $1/15$ fall each. Notably, staying inside the lateral corridor

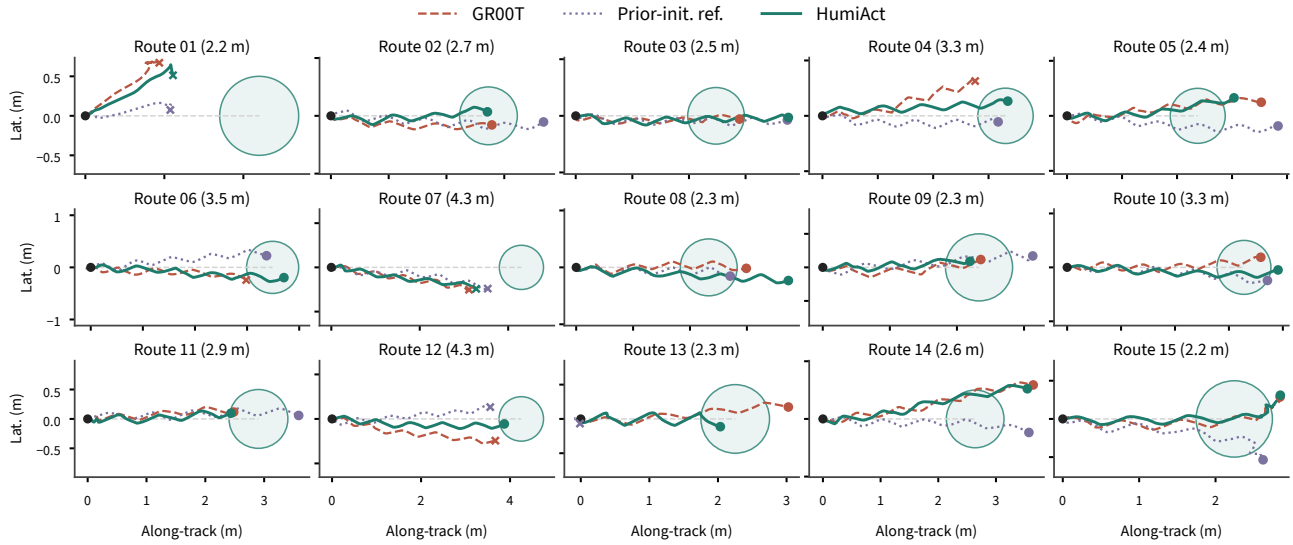


Figure S2. Complete 15-route direct-navigation trajectory atlas (standard whole-body protocol, Main Table 2). Each route i is translated to $\mathbf{p}_i^{\text{start}}$ and rotated so $\mathbf{p}_i^{\text{goal}}$ lies on the positive horizontal axis; shaded circles mark the 0.5 m success radius, filled dots mark completion, and $\times$ marks non-arrival.

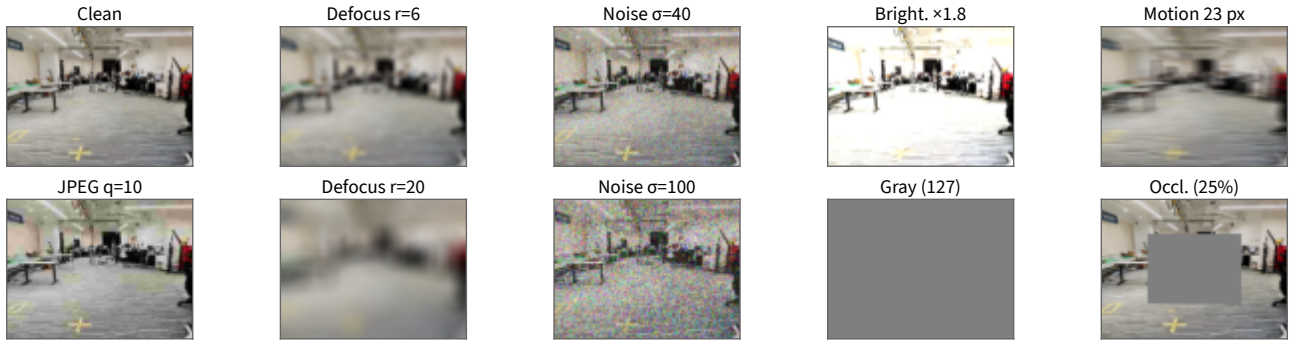


Figure S3. All ten visual-intervention inputs on Route 01 (240×320 left-eye RGB, applied immediately prior to policy inference while leaving proprioception, language, and controller inputs unchanged).

alone is insufficient without sustained forward progress (the prior reference has $0/15$ departures yet reaches only $11/15$ goals vs. $13/15$ for HumiAct).

Exact Visual Corruption Specifications. Corruptions [36] are applied to the resized 240×320 left-eye RGB image immediately before policy inference (Fig. S3): (i) *Defocus blur*: convolution with a normalized disk kernel of radius $r{=}6$ px (standard) or $r{=}20$ px (strong); (ii) *Gaussian noise*: additive zero-mean noise with $\sigma{=}40$ (standard) or $\sigma{=}100$ (strong) on the 0–255 RGB scale, clipped to $[0, 255]$ uint8; (iii) *Brightness*: multiplicative scaling by $1.8\times$ with clipping to 255; (iv) *Motion blur*: 23 px horizontal box kernel; (v) *JPEG compression*: encode–decode round trip at quality $q{=}10$ with explicit RGB/BGR conversion; (vi) *Gray image*: all pixels replaced with constant RGB 127; (vii) *Central occlusion*: central half-width by half-height box (25% area) filled with RGB 127.

C. UniAct Predictive Prior and LAG Ablation Details

Claim anchor. *Supports Main Sec. 3.2–3.3, Sec. 4.4, and Sec. 4.6 (Tables 5 and 7, Figs. 3 and 6). Takeaway:* Action-conditioned predictive accuracy is balanced across all five behaviors (85–90% one-step recall), and the LAG ablation confirms that downstream gains arise specifically from action-aligned future dynamics rather than generic feature regularization.

Unified Action Mapping and Split Isolation. Every heterogeneous trajectory transition is mapped to one of five shared behaviors—turn left, move forward, turn right, move backward, and manipulate—encoded in indices 0–4 of a 7-D zero-padded one-hot condition $\mathbf{u}_t \in \{0, 1\}^7$ (indices 5–6 remain zero). Base translation and rotation identify navigation classes, while arm and hand motion identify

Table S4. Complete held-out UniAct representation analysis (supporting Main Sec. 4.4, Table 5, and Fig. 6). **(a)** 500 balanced held-out clips (100 per first action; source mix: AgiBotWorld 348 [69.6%], Galaxea 127 [25.4%], Humanoid 25 [5.0%]) with 95% bootstrap CIs (derangement min–max range over all 44 label permutations for the permutation control). **(b)** Per-action $H{=}1$ scores on the balanced panel. **(c)** Separate natural-distribution panel (489/500 manipulation-first clips; constant majority gives 97.8% at $H{=}1$) and 4-class PCA projection statistics (seed 20260923).

(a) Sequence discrimination across horizons (%, balanced 500-clip held-out panel)

Method	$H{=}1$ Micro ↑	$H{=}1$ Macro ↑	$H{=}2$ Exact ↑	$H{=}2$ Step ↑	$H{=}3$ Exact ↑	$H{=}3$ Step ↑
Random chance (analytic)	20.0	20.0	4.0	20.0	0.8	20.0
Train-marginal majority	20.0 [16.6, 23.6]	20.0 [20.0, 20.0]	20.0 [16.6, 23.6]	31.8 [28.4, 35.3]	20.0 [16.6, 23.6]	39.8 [36.4, 43.3]
Permuted action labels (44 derang.)	3.0 [0.8, 5.0]	3.0 [0.8, 5.0]	0.7 [0.0, 1.6]	4.6 [3.1, 6.0]	0.3 [0.0, 0.6]	5.2 [3.4, 6.7]
UniAct (frozen predictor)	**88.2 [85.4, 91.0]**	**88.2 [85.4, 91.0]**	**67.8 [63.6, 71.8]**	**81.5 [79.0, 84.0]**	**53.8 [49.6, 58.2]**	**79.3 [77.1, 81.6]**

(b) One-step per-action discrimination (balanced)

Behavior	Clips	Recall	Prec.	F1
Turn left	100	90.0	93.8	91.8
Move forward	100	88.0	85.4	86.7
Turn right	100	88.0	86.3	87.1
Move backward	100	90.0	95.7	92.8
Manipulate	100	85.0	81.0	82.9

(c) Natural panel & 4-class PCA geometry

Quantity	Value
Natural $H{=}1$ Exact / Step	92.0 [89.6, 94.4] / 92.0 [89.6, 94.4]
Natural $H{=}2$ Exact / Step	75.8 [72.0, 79.4] / 87.6 [85.6, 89.5]
4-class argmax acc. (400 clips)	89.75% (359/400)
Mean true-action cosine	0.6381
PC1 / PC2 / Total var. (2D silh.)	41.62% / 34.62% / 76.24% (0.1731)

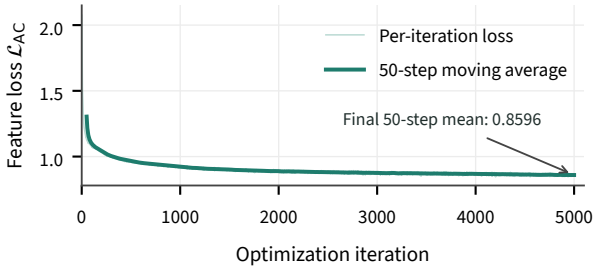


Figure S4. Predictive-prior optimization curve over 5,000 iterations (10 epochs × 500 steps; seed 20260916). Loss decreases smoothly from 2.0756 to a final 50-step mean of 0.8596.

manipulation; when both cues appear simultaneously, the largest normalized motion cue assigns the active class, and idle or ambiguous intervals are discarded so manipulation never acts as a catch-all. Training, validation, and test partitions strictly separate source video, action-array, and source-parquet identifiers to prevent clip leakage. Stage 1 trains P_ψ over frozen ViT-giant features (E_ϕ, 576 tokens of width 1408, 4 frames at 2 Hz) with first-action sampling weights $[2, 2, 2, 1, 2]$ (left, forward, right, backward, manipulate) to counter manipulation dominance (Fig. S4).

Held-Out Sequence Discrimination and PCA Geometry. Table S4 provides complete confidence intervals, per-action scores, and natural-distribution controls for Main Table 5 and Fig. 6. For five equiprobable actions, random exact-sequence accuracy is 5^{-H} ($20.0\%, 4.0\%, 0.8\%$ for $H{=}1, 2, 3$) while per-step chance stays 20.0%. On the balanced 500-clip panel (100 clips per first action), one-step recall spans 85.0–90.0% across all five behaviors, verifying that minority navigation behaviors are represented as faithfully as manipulation. At $H{=}3$, exact sequence accuracy reaches 53.8% while per-step accuracy stays at 79.3%, showing that multi-step rollouts preserve action-specific dynamics even when a single step errs. The separate natural-distribution panel (489/500 manipulation-first clips; Table S4c) reaches 92.0% at $H{=}1$ and 75.8% exact at $H{=}2$, but is kept separate because a constant manipulation baseline already achieves 97.8% first-action accuracy.

For the four-behavior geometry analysis (Main Fig. 6a; 400 clips excluding backward motion), given clip context and observed future representation $\mathbf{Z}_i^+$ and candidate action prediction $\hat{\mathbf{Z}}_i^+(u)$, we compute candidate cosine scores and centered relative evidence:

$$s_{i,u} = \cos\big(\hat{\mathbf{Z}}_i^+(u), \mathbf{Z}_i^+\big), \qquad e_{i,u} = s_{i,u} - \frac{1}{4}\sum_v s_{i,v}. \tag{12}$$

Argmax over $s_{i,u}$ (computed *before* PCA) selects the true behavior on $359/400$ clips (89.75%, mean true-action cosine 0.6381). Standardizing each coordinate of $e_{i,\cdot}$ across the 400 clips and applying 2-component PCA (seed 20260923) explains $41.62\% + 34.62\% = 76.24\%$ of variance (2D silhouette 0.1731).

LAG Distillation Interface and Ablation Design. In Stage 2, 32 auxiliary tokens (16 current, 16 future) are appended to the VLA backbone sequence, and student features $\mathbf{h}_t^{\text{cur}}, \mathbf{h}_t^{\text{fut}} \in \mathbb{R}^{16\times 2048}$ are extracted from the third-from-last language-backbone layer. A shared linear projector $G : \mathbb{R}^{2048} \to \mathbb{R}^{1408}$ maps them to the teacher width, while learned attention poolers $Q_{\text{cur}}, Q_{\text{fut}}$ compress the 576 teacher patch tokens into 16 target tokens, trained with equal unit weights ($\lambda_{\text{cur}}{=}\lambda_{\text{fut}}{=}1.0$, matching Main Eq. 8). All ablation variants in Main Table 7 undergo the exact same Stage 3 fine-tuning recipe:

- *Current + Aligned* sets $\lambda_{\text{fut}}{=}0$, keeping only current-frame teacher features; it matches HumiAct on Terra-guided navigation (53.3%) but trails on direct navigation (73.3% vs. 86.7%) and five-corruption mean (78.7% vs. 84.0%).

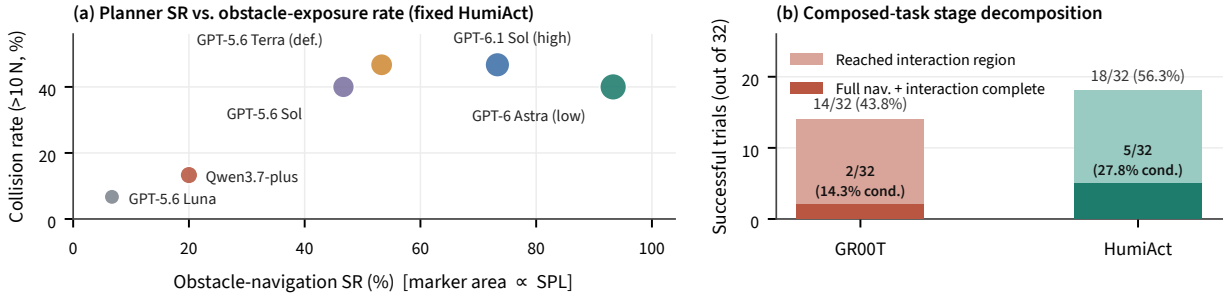


Figure S5. Planner exposure trade-off and composed-task funnel. **(a)** Obstacle-navigation SR vs. collision rate (> 10 N) across six planners with the HumiAct motor policy held fixed (marker area $\propto$ SPL; Main Table 6). **(b)** Composed-task arrival and conditional completion counts on 32 trials (Main Table 2).

- *Future* + *None* and *Current* + *future* + *None* remove the UniAct action condition $\mathbf{u}_t$ from P_ψ, testing unconditioned predictive distillation; both underperform HumiAct (66.7%/73.3% direct; 76.0%/80.0% corruption mean).
- *Current* + *future* + *Shuffled* permutes teacher action labels during distillation while keeping architecture and budget identical; performance drops sharply below even unguided GR00T (46.7% vs. 66.7% direct; 74.6% vs. 77.3% corruption mean), confirming that mis-aligned predictive targets induce destructive interference and that HumiAct's gains specifically require action-aligned dynamics.

D. Task Planner, Physical-Robot Evaluation, and Qualitative Rollouts

Claim anchor. *Supports Main Sec. 3.1, Sec. 3.4, and Sec. 4.5–5 (Table 6, Figs. 1 and 4). Takeaway:* Varying only the high-level planner reveals a strong SR–exposure trade-off (6.7% to 93.3% SR with the exact same motor policy), while physical-robot trials and ten representative rollouts substantiate closed-loop walking, bar-counter approach, grasping/delivery, and whole-body chair pushing.

Closed-Loop Planner Contract. After each executed action chunk, the planner receives the full episode observation history and outputs a structured JSON decision containing `action_id`, `objective`, `remaining_plan`, `decision` $\in$ {retain, advance, revise}, `reason`, and `visual_state`. All six planners in Main Table 6 drive the exact same HumiAct checkpoint and whole-body controller on the same 15 obstacle routes; GPT-6 Astra uses low reasoning effort, GPT-6.1 Sol uses high reasoning effort, and the remaining four planners use default settings. As shown in Fig. S5a, stronger visual-spatial planners (GPT-6 Astra: 14/15, SPL 0.716; GPT-6.1 Sol: 11/15, SPL 0.569) unlock substantially higher obstacle-navigation performance from the fixed motor policy than the default GPT-5.6 Terra planner (8/15, SPL 0.410), confirming that Main Table 2 is conservative with respect to planner quality.

Table S5. Physical-robot primitive motion trials (supporting Main Sec. 4.5). Planner disabled; 1 denotes successful execution and 0 denotes failure across three independent repetitions per command.

	GR00T				Nav.-adapted policy			
Command	T1	T2	T3	Succ.	T1	T2	T3	Succ.
Turn left	1	1	1	3/3	1	1	1	3/3
Turn right	1	1	1	3/3	1	1	1	3/3
Move forward	0	1	0	1/3	1	1	0	2/3

Physical-Robot Execution Stack and Primitive Trials. On the physical 33-DoF humanoid (37-D proprioceptive state, synchronized head-mounted cameras), the shared VLA runs on an NVIDIA RTX 5000 Ada at 10 Hz (BF16, 4 flow-integration steps, 30-step chunks) and communicates with the onboard ROS 2 client via ZeroMQ. Buffered references are aligned to the execution clock, smoothly blended across chunk transitions, and streamed at 50 Hz to the whole-body tracker with synchronized head and hand commands. Table S5 records all per-trial outcomes for the planner-free motion-command evaluation referenced in Main Sec. 4.5. Both GR00T and the navigation-adapted policy complete 3/3 left turns and 3/3 right turns; on forward walking, the adapted policy succeeds in 2/3 trials vs. 1/3 for GR00T (the single adapted-policy failure resulted from an initial heading drift before gait stabilization). As emphasized in Sec. A, these three-repetition primitive trials provide controlled hardware verification of motion execution and are distinct from full obstacle or composed navigation trials.

Qualitative Rollout Summary. Figure S6 presents keyframe sequences from both physics-based simulation and physical-robot rollouts, and Table S6 summarizes the ten representative qualitative demonstrations across simulation and real-humanoid evaluation (supporting Main Figs. 1 and 4, and Sec. 4.2–5).

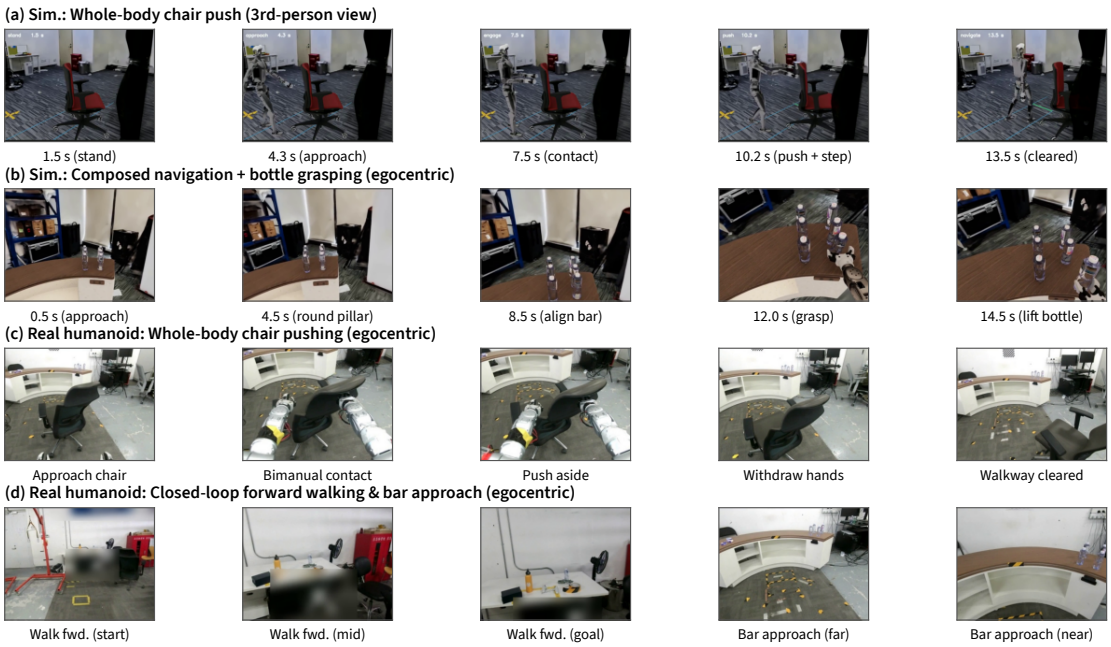


Figure S6. Qualitative filmstrips across simulation and real-humanoid deployment. **(a)** Simulated whole-body chair pushing (3rd-person): walking approach, bimanual contact, stepping while pushing, and clearing the walkway. **(b)** Simulated composed navigation and bottle grasping (32-trial benchmark). **(c)** Real 33-DoF humanoid bimanual chair pushing in front of the bar counter (egocentric). **(d)** Real 33-DoF humanoid closed-loop forward walking (Table S5, Trial 2) and bar-counter approach (egocentric).

Table S6. Summary of qualitative simulation and physical-humanoid rollouts. Representative closed-loop demonstrations across simulation and real-robot evaluation corresponding to the main text and Fig. S6.

Rollout scenario	Dom.	Anchor	Content summary
1. Direct-navigation comparison (Route 04)	Sim	Sec. 4.2	GR00T vs. HumiAct on direct-nav. Route 04 (Fig. 4a).
2. Whole-body chair push (dual-view)	Sim	Fig. 1b	Dual-view chair push: approach, bimanual push, traversal.
3. Chair push and bar approach (egocentric)	Sim	Sec. 3.1	Egocentric bimanual chair push + continuous bar approach.
4. Composed navigation and grasp (Trial 1)	Sim	Table 2	Composed 32-trial Trial 1: pillar approach + bottle grasp.
5. Bimanual chair push (cleanroom)	Sim	Sec. 3.4	Cleanroom physics check of bimanual push and arm swing.
6. Physical-humanoid chair pushing	Real	Sec. 5	Physical 33-DoF humanoid: two chair-push trials.
7. Physical-humanoid walking and bar approach	Real	Sec. 4.5	Physical humanoid: forward walk (Table S5, T2) + bar approach.
8. Composed navigation and grasp (Trial 2)	Sim	Table 2	Composed 32-trial Trial 2: pillar turn, alignment, bottle lift.
9. Closed-loop banana grasping	Sim	Sec. 3.4	Closed-loop banana top-pinch: alignment and 28 cm lift.
10. Closed-loop milk-carton delivery	Sim	Sec. 3.4	Full closed-loop milk-carton delivery (0 collisions, complete).