GAPONET: NONLINEAR OPERATOR LEARNING FOR BRIDGING THE HUMANOID SIM-TO-REAL GAP

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

033

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

The sim-to-real gap, arising from imperfect actuator modeling, contact dynamics, and environmental uncertainty, poses fundamental challenges for deploying simulated policies on physical robots. In humanoids, object manipulation further amplifies this gap: end-effector payloads alter joint inertia, gravity torques, and transmission efficiency, introducing state- and payload-dependent nonlinearities. Yet existing approaches lack both systematic analysis and a generalizable representation of this payload-induced degradation. To address this limitation, we propose **GapONet**, a payload-conditioned nonlinear operator that maps simulation context functions to residual actions for hardware. We then introduce a payloadaware (collect-analyze-solve) framework to learn this operator **GapONet**. First, we curate a sim-real paired dataset **TWINS** spanning multiple payloads, robots, motions, actuation rates, and simulators, comprising more than 11,298 motion sequences. Second, we perform payload-aware system identification to isolate payload-related effects and quantify their contributions, and analyze sim-to-real gaps across different simulators. Third, we train the operator **GapONet** to predict delta action for real-time, generalized, payload-conditioned compensation. We further introduce actuation functions and sensor predictors, which enable parallel RL training of **GapONet** with substantially reduced energy consumption. While tracking unseen motions, **GapONet** keeps the incidence of large sim-to-real gaps below 0.09%, whereas competing methods remain near 10%. By correcting upperbody gaps, GapONet also stabilizes lower-body locomotion tracking, laying the foundation for improved performance in humanoid loco-manipulation tasks.

1 Introduction

Policies trained in simulation benefit from GPU acceleration and massively parallel sampling, enabling fast and scalable optimization under approximate physics such as mass, friction, and damping (Makoviychuk et al., 2021; Tan et al., 2018). However, object interactions in the real world often diverge from these idealizations due to unmodeled or state-dependent effects, most notably in friction, inertia, and contact—leading to a persistent model—plant mismatch (Tobin et al., 2017; Zhao et al., 2020). This sim-to-real gap is further exacerbated in humanoids that manipulate objects of different masses. Variations in end-effector payload induce coupled drifts in equivalent joint inertia, gravity—torque amplitudes via center-of-mass and lever-arm shifts, transmission friction and efficiency, thereby altering closed-loop dynamics (Spong et al., 2006). Yet during policy training, these payload-dependent adjustments are typically simplified or held fixed, which leaves the gap largely unaddressed. The sim-to-real gap can grow in complex, nonpredictive ways, posing a substantial obstacle to robust policy transfer and reliable real-world deployment (Zhang et al., 2023).

Prevailing approaches either calibrate simulators via system identification to tune masses, frictions, and damping (Ljung, 1998; Åström & Eykhoff, 1971; Nelles, 2002); broaden training distributions through domain randomization and observation noise to reduce overfitting (Mehta et al., 2020; Tobin et al., 2017; Chen et al., 2021; Laskey et al., 2017; Zhang et al., 2020; Matas et al., 2018); or stage learning with curricula or progressively harder terrains to harden policies over time (Luo et al., 2020; Wang et al., 2021; Peng et al., 2020; Heess et al., 2017) to bridge the sim-to-real gap. However, the interacted object (payload) is a structured operating condition, not mere noise (Slotine & Li, 1987): it deterministically alters gravity loading, effective inertia, dissipation, and hence the closed-loop gain/phase under PD control. Single-point identification cannot capture behavior across payloads,

and domain randomization or curricula largely treat the payload as unstructured uncertainty. Thus, while these strategies can improve robustness, they hinge on manual design (randomization ranges, noise schedules, curriculum pacing) and provide limited diagnostic attribution. Critically, they do not yield a generalizable representation of the sim-to-real gap for humanoid interaction.

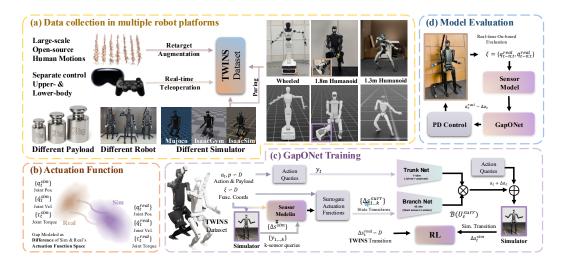


Figure 1: **The overall architecture of both data collection and GapONet training.** (a) **TWINS**, a paired sim—real dataset via motion retargeting and real-time teleoperation across diverse payloads, robots, and simulators. (b) The sim—real gap is formulated as a discrepancy between *actuation function* spaces, providing functional coordinates. (c) **GapONet** learns a payload-conditioned nonlinear operator that maps simulation context to residual actions, and training uses parallel RL. (d) Online evaluation on unseen hardware with PD control and sensor modeling to quantify sim—real alignment.

A complementary line of work learns dynamics directly from real data, either as state-transition models or action-to-effect maps (Shi et al., 2019; Xiao et al., 2024; He et al., 2025). From a control standpoint, however, identifying payload-dependent dynamics from passive logs requires persistence of excitation and explicit treatment of operating conditions. In practice, motion patterns, contact regimes, and payload values co-vary, so a single black-box model fit to mixed data tends to entangle payload effects with task-specific artifacts, yielding spurious correlations. As a result, such models often need large volumes of paired sim—real trajectories to cover the space and still exhibit poor cross-payload and unseen-motion generalization. The missing ingredient is a representation that disentangles exogenous operating parameters from state evolution, rather than collapsing them into a single dynamics model. Such a formulation enables a more faithful mapping between the simulator and real-world domains.

We present a \(collect\)-analyze-solve\) framework to learn this representation for bridging the simto-real gap in humanoids. We first curate **TWINS**, a time-synchronized sim-real corpus with a structured factorial design. Unlike prior collections (Wu et al., 2024; Mao et al., 2024; AgiBot-World-Contributors et al., 2025), our dataset design over diverse payload levels, humanoid platforms, actuation rates, simulations, and motion families, enabling further controlled analyses. To clarify the **GapONet** 's learning target, we first perform gray-box, block-wise system identification atop a PD control model, attributing error reductions to specific payload-related terms and quantifying their contributions. We then analyze identical motions across payloads and simulators, showing structured residuals dominated by actuator nonlinearities, which motivates a more generalizable nonlinear operator rather than a pointwise approximation function.

We then propose **GapONet**, a payload-conditioned nonlinear operator that maps simulation context functions to a residual actions for hardware. Our operator is parameterized with a branch—trunk decomposition (Lu et al., 2019): The branch net encodes the local dynamics of the physical world in which our robot resides as a function, and the trunk network encodes the input variables to that function, including payload weight and target pose. This separation provides a strong structural inductive bias, disentangling the conditioning context from the queried response, thereby enhancing the operator's generalization capacity. We also propose the sensor predictor, enabling parallel RL training

of **GapONet** with lower energy cost while preserving generalization beyond pointwise regression. While tracking unseen motions, **GapONet** keeps the incidence of large sim-to-real gaps below 0.09%, whereas competing methods remain near 10%. By correcting upper-body gaps, **GapONet** also stabilizes lower-body locomotion tracking, laying the foundation for improved performance in humanoid loco-manipulation tasks.

This paper makes three primary contributions:

- We develop a sim-real data collection pipeline and we curate the first dataset TWINS focusing on payload-induced sim-real gap across multiple payloads, robots, motions, and simulators.
- We reproduced over 30 hours of real data across four simulators and conducted controlled, ceteris paribus comparisons, yielding quantitative evidence that sim-to-sim evaluation improves the deployability of humanoid controllers.
- We introduce **GapONet**, a payload-conditioned nonlinear operator that maps simulation context functions to residual actions for hardware, and demonstrate its training via RL.

2 RELATED WORK

108

109

110

111

112

113

114 115

116

117

118

119

120

121

122

123 124 125

126 127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143 144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

Sim-to-Real Gap Sim-to-real research has largely moved from system identification—calibrating masses, frictions, and control gains to align simulation with measurements (Sobanbabu et al., 2025; Gu et al., 2024; Zhang et al., 2024)—to domain randomization, which perturbs dynamics and observations to harden policies (Peng et al., 2018; Xie et al., 2021; Mehta et al., 2020; Chen et al., 2021). The former can deliver high fidelity but typically demands accurate structural assumptions and extensive hardware time (Ljung, 1998; Miller et al., 2025); the latter proved influential for legged and humanoid control (Xie et al., 2020; Margolis et al., 2024; Li et al., 2023) yet can bias policies toward conservatism (He et al., 2024). In practice, both families often require substantial manual retuning across agents, tasks, and operating regimes, motivating data-driven directions that learn from collected data. One line models actuator nonlinearities with fine granularity to capture motorlevel effects (Hwangbo et al., 2019); another emphasizes residual correction, learning delta actions for online compensation with lighter overhead (He et al., 2025). In parallel, simulation-real fusion seeks coverage and speed from simulators while retaining real-world grounding (Fey et al., 2025; Zhang et al., 2023; Bjelonic et al., 2025; Xu et al., 2025; Ouyang & Cui), and new benchmarks standardize evaluation (Bjelonic & Hutter, 2025). Despite these advances, both simulator-centric and data-centric pipelines still struggle with broad generalization under real-world variability (Muratore et al., 2022), which limits general gap-bridging in complex systems, such as humanoids.

Nonlinear Operator Nonlinear operator learning. Rather than learning pointwise mappings, operator learning targets mappings between function spaces, where both inputs and outputs are functions (Kovachki et al., 2023). Within this paradigm, Unstacked Deep Operator Network (DeepONet) offers a principled route to learn nonlinear operators via an operator-level universal approximation result (Lu et al., 2019). Its branch-trunk decomposition encodes input functions in the branch network and query locations in the trunk, combining them (e.g., via inner products) to produce function values; the construction connects to low-rank approximations and RKHS viewpoints, lending theoretical footing to the architecture (Hornik et al., 1989; Lu et al., 2021). Building on these foundations, recent studies have pushed operator learning toward control and engineering settings: formulations grounded in Hamilton–Jacobi policy iteration suggest a pathway to control-theoretic operators (Lee & Kim, 2025); physics-informed treatments extend the approach to optimal control (Na & Lee, 2024); and model-predictive control has been instantiated with deep operator networks to handle online decision-making under dynamical constraints (de Jong et al., 2025). Beyond control, multiphysics applications demonstrate operator surrogates for solution fields in materials processing and additive manufacturing, highlighting scalability to complex PDE-governed phenomena (Kushwaha et al., 2024). Despite this progress, most deployments remain either theory-centric or domain-specific, with limited attention to robotics sim-to-real—in particular, to humanoid systems subject to shifting operating conditions such as payload changes. This gap motivates operator-based formulations that explicitly encode conditioning on task and environment variations while preserving sample efficiency and real-time viability.

3 DATA COLLECTION AND GAP ANALYSIS

End-effector payloads reshape joint dynamics and closed-loop behavior—raising reflected inertia, shifting gravity torques, and coupling with actuator and contact nonlinearities. Divergent simulator treatments of these effects produce a persistent, multi-factor sim-to-real gap. This section provides a structured diagnosis: Section 3.1 isolates payload-induced terms via gray-box system identification; Section 3.3 compares simulators on identical payload-bearing motions under matched controllers; Section 3.2 details **TWINS** and its collection pipeline.

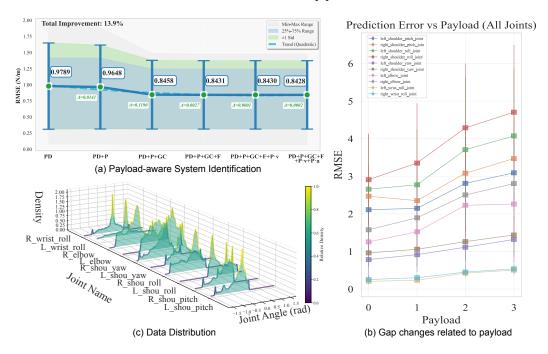


Figure 2: **System identification and data distribution** (a) Prediction residuals after adding payload-related parameters; notably, adding gravity compensation yields a clear improvement. (b) The vertical axis shows the change in the joint-wise gap as the payload increases. (c) Data distribution of **TWINS**; the z-axis indicates the probability density of each joint action.

3.1 PAYLOAD-AWARE SYSTEM IDENTIFICATION

Using bipedal humanoids that demand precise control as exemplars (Unitree H1-2 and G1), both operate under joint-space PD control tailored to locomotion (details in Section A.3.1). With added end-effector payloads P, we adopt a gray-box identification scheme: start from a rigid PD baseline and progressively augment the torque model with physically grounded terms salient in manipulation. For each joint, we fit a linear in parameters regression that attributes the sim-to-real discrepancy to gravity scaling, reflected inertia, actuator and transmission nonlinearities, and contact compliance, and we quantify their marginal contributions:

$$\tau = K_{p} (q_{\text{cmd}} - q) + K_{d} (\dot{q}_{\text{cmd}} - \dot{q}) + K_{v} \dot{q} + K_{c} \tanh\left(\frac{\dot{q}}{\varepsilon}\right)$$

$$+ K_{\text{payload}} P$$

$$+ K_{P \sin} P \sin q + K_{P \cos} P \cos q$$

$$+ K_{P\dot{q}} P \dot{q} + K_{P\ddot{q}} P \ddot{q}$$

$$+ \tau_{0}.$$

$$(1)$$

Here, K_p and K_d are proportional and derivative gains; K_v and K_c model viscous and Coulomb friction with ε smoothing the latter; K_{payload} scales the main payload P; $K_{P \sin}$ and $K_{P \cos}$ capture gravity and posture coupling under payload; $K_{P\dot{q}}$ and $K_{P\ddot{q}}$ model interactions between payload and joint velocity or acceleration; τ_0 is a constant bias. The remaining symbols are τ for joint torque;

 q, \dot{q}, \ddot{q} for joint position, velocity, and acceleration; $q_{\rm cmd}, \dot{q}_{\rm cmd}$ for commanded references; and P for payload magnitude interpreted as mass or equivalent inertia at the end effector. All K coefficients are identified per joint. This compact form separates baseline PD, friction, and payload dependent effects and enables clear attribution of simulation to real error.

Using over 2,000 data collected from real robots, we fit Equation (1) by minimizing RMSE between its torque and measurements. Adding payload-dependent terms reduces error Figure 2(a), with gravity compensation giving an early gain, but at higher payloads Equation (1) no longer captures the closed loop response Figure 2(b). The equation is not a replica of the simulator; it is a control equivalent surrogate that covers dominant channels under matched controllers. Identification on synchronized inputs with persistently exciting motions enables term level attribution, and the residual exposes nonlinear dynamics not captured by compact models. Learning a nonlinear operator, rather than a pointwise nonlinear function, better supports generalization across trajectories, payload schedules, actuation rates, and robots by mapping context functions to control signals.

3.2 **TWINS** COLLECTION

Section 3.1 shows with block-wise identification that the prediction to measurement gap is nonlinear and uncertain. Given the lack of suitable data, to validate this conclusion on genuine sim to real pairs, we present **TWINS**, the first dataset focused on payload induced sim to real gaps across multiple robots, standardized payload levels, and motion classes. **TWINS** records humanoid dynamics hierarchically, from single joints to full upper body motions with 3 different low-body gaits, using four Unitree H1-2 units with end effector masses from 0 to 3 kg (standard calibration weights) and actuation rates of 50 Hz and 100 Hz. The real data totals 30.17 hours, 11,298 sequences, and 307,273 synchronized frames. The distribution appears in Figure 2(c).

Each sequence is time synchronized with a matched high fidelity simulation replica in three widely used humanoid training simulators (MuJoCo, Isaac Gym, Isaac Sim), enabling comparison of real and simulated executions at the frame level and yielding a fourfold paired corpus of 120.68 (one real trace plus three simulated replicas). For every frame we record joint positions $q_{\rm sim}$, $q_{\rm real}$, velocities $\dot{q}_{\rm sim}$, $\dot{q}_{\rm real}$, accelerations $\ddot{q}_{\rm sim}$, $\ddot{q}_{\rm real}$, torques $\tau_{\rm sim}$, $\tau_{\rm real}$, payload P, and motor temperature $T_{\rm real}$. Further details of our collection pipeline and dataset on different robots are in Section A.2.

3.3 SIM-TO-REAL GAP ANALYSIS

After post-processing the paired data **TWINS**, we conduct a targeted analysis of the sim-to-real gap to guide operator design for payload-induced nonlinearities. The analysis tests concordance with the block wise identification in Section 3.1, determines whether the effect is concentrated in the upper body or extends to the whole body, and quantifies differences across simulators when reproducing the same motion under matched control.

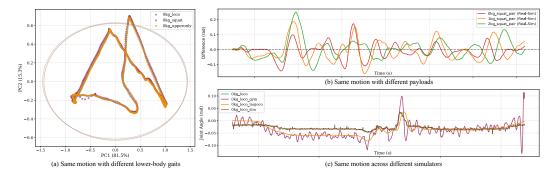


Figure 3: **Gap Analysis.** (a) The outer ellipse marks a shared kinematic envelope across gaits, while the central PCA trajectory of a single motion shows only minor variations with overall consistency. (b) Heavier payloads yield larger state gaps and longer delays. (c) Simulator comparison: Isaac Gym oscillates more, while MuJoCo and Isaac Sim track more stably.

Same motion with different lower-body gaits We execute 17 upper-body motion sequences under three lower-body conditions: bipedal locomotion, static squat, and stance support only. As shown in Figure 3(a), the outer ellipse marks the shared kinematic envelope, while the center trajectory is the PCA trace of a single motion; across gaits, this trace is nearly retraced with only small phase/offset shifts. With envelopes matched, the upper-body sim-to-real gap is therefore largely insensitive to the lower-body condition, and residual differences are dominated by payload-amplified channels. We quantify this via joint-wise normalized RMSE, commanded-measured phase lag, and torque-saturation incidence. Note that, unlike fixed-base dual-arm platforms, upper-body actions in humanoids couple back to locomotion and can stress the gait controller; full experiments and analysis are in Section 5.2.

Same motion with different payloads As shown in Figure 3(b), each colored trajectory plots the joint-wise sim-real residual over time. Increasing payload amplifies both residual magnitude and phase lag, yielding larger state gaps and longer delays. Across **TWINS**, payload consistently widens the gap, and the residual grows nonlinearly with payload mass, in line with the block-wise identification trends reported in Section 3.1.

Same motion across different simulators Current methods always apply sim-to-sim evaluation as the cross-validation before hardware deployment (He et al., 2025; Liu et al., 2024). To characterize simulator-specific differences and their dependence on payload, we compare identical motions across MuJoCo, Isaac Gym, and Isaac Sim under matched controllers and simulator-adapted generic parameters over a standardized payload grid. Experiments Figure 3(c) show that MuJoCo yields smoother trajectories but larger peaks in high-acceleration segments; Isaac Gym exhibits occasional joint-level jitter; Isaac Sim achieves the most stable alignment in our evaluations, but still leaves a nonlinear gap during interaction. To stay aligned with prevailing practice and minimize simulator-induced confounds, we adopt Isaac Sim for subsequent experiments, as it exhibits the smallest sim-to-real gap in our analysis. We also release paired data for MuJoCo and Isaac Gym to enable cross-simulator comparisons and support future research. More results in Section A.3.

In summary, across payload levels, all simulators show a nonlinear increase in error relative to real hardware, with simulator-specific modes. This pattern persists across lower-body gaits: when kinematic envelopes are matched, the distributions of upper-body error and phase metrics remain closely aligned. The discrepancy arises from coupled channels—gravity, friction, Coriolis and inertial coupling, actuator limits and efficiency drift, sensing noise, and delays—that a pointwise function mapping cannot capture or generalize. A nonlinear operator is better suited: **GapONet** provides a compact, transferable representation by mapping context functions to corrective control signals across trajectories, payload schedules, actuation rates, and robot morphologies.

4 METHOD

We propose **GapONet**, a payload-conditioned nonlinear operator that maps simulation context functions to a residual actions for hardware. **GapONet** learns a functional correspondence from simulator space to real dynamics and introduces **actuation functions** that encode command and feedback histories. We then propose the **sensor predictor**, which enables parallel RL training of **GapONet**, overcoming the high energy consumption of the original approach while maintaining generalization beyond pointwise regression.

4.1 PROBLEM FORMULATION

Previous methods lack an explicit model of both the simulator and the real world (Mehta et al., 2020; Tobin et al., 2017; Matas et al., 2018; Shi et al., 2019; Xiao et al., 2024; He et al., 2025), which reduces distributional diversity and constrains generalization. We therefore propose **actuation functions**, which bridge the gap between the simulator and the real world by learning a surrogate mapping in the function space. These functions characterize the mapping from actions (together with task-specific parameters) to state transitions, under different joint configurations and dynamics, both in simulation and on the real robot.

Formally, bridging the sim-to-real gap can be posed as learning an operator that maps \mathcal{U}^{sim} to $\mathcal{U}^{\text{real}}$ rather than approximating multiple collected dynamics, where \mathcal{U} denotes the underlying function

space. Each element of \mathcal{U} , i.e., an actuation function U, is associated with a natural coordinate representation $\xi \triangleq (q,\dot{q})$, corresponding to joint positions and velocities. With this, an actuation function is written as $U_{\xi}: A \to Q \times V$. The goal of **GapONet** is to learn an operator \mathcal{G} such that $\mathcal{G}(U_{\xi}^{\mathrm{sim}}) \approx U_{\xi}^{\mathrm{real}}$.

4.2 Network Structure

To learn the operator effectively, we adopt a DeepONet (Lu et al., 2019)-style architecture in which a branch network encodes the output of actuation functions as the conditioning function and a trunk network encodes the variables over which generalization is supposed to occur as queries. Our rationale for selecting DeepONet is detailed in Section A.5.

Inspired by dynamic modeling (He et al., 2025), our model (**GapONet**) predicts delta actions for each joint, compensating for discrepancies between simulated and real-world dynamics (Craig, 2009). Here, we define the input of actuation functions as the sensor $x_t = \{(q_t, \dot{q}_t)\}$. Specifically, given k fixed sensor locations x_1, \ldots, x_k , we first query the simulated actuation function:

$$S_i(U_{\mathcal{E}}^{\text{sim}}) = \Delta f^{\text{sim}}(s_{\text{sim}}^{\xi}, x_i) = U_{\mathcal{E}}^{\text{sim}}(x_i), \quad i = 1, \dots, k.$$
 (2)

These outputs are then embedded into a latent representation via the Branch Net:

$$\mathcal{B}(U_{\xi}^{\text{sim}}) = [\mathcal{B}_1(U_{\xi}^{\text{sim}}(x_1)), \dots, \mathcal{B}_k(U_{\xi}^{\text{sim}}(x_k))], \tag{3}$$

where each component \mathcal{B}_i captures a distinct feature of the actuation state, allowing the network to decompose complex dynamics into interpretable subcomponents.

The Trunk Net encodes query signals consisting of both the payload and the current action:

$$y \in Y = P \times A, \quad \mathcal{T}(y) = [\mathcal{T}_1(y), \dots, \mathcal{T}_k(y)],$$
 (4)

where A and P denote actions and payloads, respectively. This serves to condition the latent space, aligning actuator dynamics with task objectives.

Finally, Branch and Trunk features are fused:

$$G_{\theta}(\xi, y) = \mathcal{B}(U_{\xi}^{\text{sim}}) \cdot \mathcal{T}(y) = \sum_{i=1}^{k} \mathcal{B}_{i}(U_{\xi}^{\text{sim}}(x_{i})) \cdot \mathcal{T}_{i}(y), \tag{5}$$

yielding the delta action $\Delta a^j = G^j$ for each joint j. This correction augments the simulator's nominal command, bridging the sim-to-real gap. The overall operator is then defined as:

$$\mathcal{G}(U_{\xi}^{\text{sim}})(y_t) = \Delta f^{\text{sim}}\left(s_{\text{sim}}^{\xi}, a_t + G_{\theta}(U_{\xi}^{\text{sim}}(y_t))\right). \tag{6}$$

4.3 GPU-PARALLEL OPERATOR LEARNING

A key challenge arises when applying this network in parallel Reinforcement Learning (RL) environments: computing sensor values for every ξ is computationally prohibitive. To address this, we introduce a **sensor model** S_{ϕ} , which predicts sensor readings directly from the actuation coordinates ξ :

$$\mathcal{L}_{\text{sensor}} = \mathbb{E}_{\xi} \left[\sum_{i} ||\Delta f^{\text{sim}}(s_{\text{sim}}^{\xi}, x_i) - (S_{\phi}(\xi))_i||_2^2 \right]. \tag{7}$$

Optimizing ϕ yields a surrogate function space $\mathcal{U}^{\text{surr}} \triangleq S_{\phi}(\mathcal{U}^{\text{sim}})$. By interpreting S_{ϕ} as an operator mapping U_{ξ}^{sim} to U_{ξ}^{surr} with approximately equal sensor output, our framework reduces to learning an operator from $\mathcal{U}^{\text{surr}}$ to $\mathcal{U}^{\text{real}}$, with the training objective:

minimize dist
$$(\mathcal{G}(S_{\phi}(U_{\varepsilon}^{\text{sim}})), U_{\varepsilon}^{\text{real}}).$$
 (8)

If the distance metric is defined analogously to L^2 distance in the function space, this can be reformulated as a RL problem with reward:

$$r_t = -w||(s_{\text{real}}^{t+1} - s_{\text{real}}^t) - \mathcal{G}(U_{\mathcal{E}}^{\text{surr}})(y_t)||_2^2,$$
 (9)

where s_{real} , y are sampled from **TWINS**.

5 EXPERIMENT

Our experimental evaluation comprises two parts: Section 5.1 evaluates **GapONet** 's zero-shot generalization to unseen robots and motions; Section 5.2 measures improvements in humanoid locomotion stability through online residual compensation on hardware.

5.1 ZERO-SHOT MOTION TRACKING

GapoNet can generalize to unseen target joint-position sequence (motion) under the branch-trunk architecture. To test this capability beyond our dataset **TWINS**, we collected an unseen-motion test set of 100 sim-real pairs: 35 sequences at 0 kg, 23 at 1 kg, 22 at 2 kg, and 20 at 3 kg. The test set also spans three lower-body gaits in a 6:3:1 ratio for static stance, squat, and locomotion. For quantitative assessment, we report **Large Gap Ratio** (the percentage of frames whose error exceeds a predefined threshold), **IQR** (the interquartile range of the gap over all motions), and **Gap Range** (the framewise gap range from minimum to maximum).

We benchmark **GapONet** with four baselines: (i) an MLP learned dynamics model (He et al., 2025), (ii) a Transformer learned dynamics model that exploits temporal context better, (iii) system identification, a classical approach to bridging the sim-to-real gap, and (iv) PD control with official gains. Each experiment is repeated multiple times, and we report the mean and standard deviation in the table. As shown in Table 1, **GapONet** attains the best or tied-best scores on nearly all metrics, with a pronounced improvement in LGR. These results indicate smoother, more controllable zero-shot gap bridging than the learned dynamics baselines and consistent gains over system identification across motions from multiple robots.

Table 1: Zero-shot sim-to-real gap on unseen-motion test set across four payloads.

Method	0 kg			1 kg			
	LGR(%)	IQR (↓)	Range (↓)	LGR(%)	IQR (↓)	Range (↓)	
PD control	$12.7^{\pm 3.3}$	$0.138^{\pm0.007}$	$0.538^{\pm0.019}$	$10.6^{\pm0.1}$	$0.139^{\pm0.028}$	$0.667^{\pm0.011}$	
MLP	$10.0^{\pm0.8}$	$0.108^{\pm0.012}$	$0.480^{\pm0.088}$	$10.8^{\pm0.1}$	$0.125^{\pm0.002}$	$0.589^{\pm0.029}$	
Transformer	$9.55^{\pm0.3}$	$0.127^{\pm0.014}$	$0.465^{\pm0.067}$	$5.60^{\pm0.4}$	$0.140^{\pm0.005}$	$0.525^{\pm 0.041}$	
System Identification	$12.4^{\pm0.3}$	$0.141^{\pm0.015}$	$0.505^{\pm0.032}$	$9.01^{\pm 1.0}$	$0.140^{\pm0.029}$	$0.609^{\pm0.122}$	
GapONet (Ours)	$0.09^{\pm0.03}$	$0.093^{\pm0.016}$	$0.449^{\pm0.117}$	$0.22^{\pm0.11}$	$0.115^{\pm0.013}$	$0.537^{\pm0.148}$	
Method	2 kg			3 kg			
	LGR(%)	IQR (↓)	Range (↓)	LGR(%)	IQR (↓)	Range (↓)	
PD control	$11.2^{\pm0.1}$	$0.205^{\pm0.001}$	$0.625^{\pm0.038}$	$12.8^{\pm0.1}$	$0.499^{\pm0.008}$	$0.642^{\pm0.060}$	
MLP	$10.8^{\pm0.1}$	$0.252^{\pm0.003}$	$0.621^{\pm0.023}$	$12.2^{\pm0.9}$	$0.460^{\pm0.013}$	$0.668^{\pm0.060}$	
Transformer	$0.44^{\pm0.3}$	$0.140^{\pm0.002}$	$0.606^{\pm0.040}$	$9.82^{\pm0.1}$	$0.416^{\pm0.002}$	$0.573^{\pm0.178}$	
System Identification	$9.53^{\pm0.7}$	$0.193^{\pm0.102}$	$0.601^{\pm0.031}$	$12.1^{\pm0.5}$	$0.494^{\pm0.003}$	$0.611^{\pm0.127}$	
GapONet (Ours)	$0.39^{\pm0.10}$	$0.161^{\pm0.004}$	$0.578^{\pm0.112}$	$0.84^{\pm0.23}$	$0.317^{\pm0.005}$	$0.498^{\pm0.157}$	

5.2 LOCOMOTION TRAJECTORY TRACKING

Section 5.1 demonstrates the generalization and gap-solving capabilities of **GapONet**, but improving upper-body tracking alone is insufficient to prove system-level benefits. For broader humanoid applications, lower-body motion must also be considered. As shown in Section 3.3, lower-body gaits have minimal impact on upper-body motion distributions, while upper-body compensation affects the lower-body dynamics through coupled torques and contact forces, influencing the center of mass trajectory (Zhang et al., 2025). To further validate **GapONet**'s ability to address the upper-body gap, we introduce an online residual compensation method that adapts to varying lower-body states.

We provide both qualitative and quantitative results to evaluate the performance of **GapONet**. We conducted tests on 14 motion sequences (7 at 0 kg and 7 at 1 kg payloads) using a previously unseen Unitree H1-2 robot. For quantitative assessment, we report **Trajectory Consistency** (velocity dis-

Table 2: Sim-to-real gap in locomotion trajectory tracking on an unseen humanoid robot.

Method	Trajectory Consistency (↓)		Smoothness (↓)		Robustness (↓)	
	0 kg	1 kg	0 kg	1 kg	0 kg	1 kg
PD control	$20.33^{\pm 1.982}$	$27.49^{\pm 1.057}$	$53.76^{\pm0.257}$	$25.76^{\pm0.277}$	$10.16^{\pm0.007}$	10.14 ^{±0.026}
MLP	$19.18^{\pm0.919}$	$28.82^{\pm 1.560}$	$53.48^{\pm0.343}$	$25.55^{\pm0.361}$	$10.15^{\pm0.027}$	$10.14^{\pm0.024}$
Transformer	$19.13^{\pm0.689}$	$29.05^{\pm 1.576}$	$53.57^{\pm0.290}$	$26.56^{\pm0.385}$	$10.14^{\pm0.007}$	$10.16^{\pm0.012}$
System Identification	$19.16^{\pm0.489}$	$28.59^{\pm1.343}$	24.99 ^{±0.298}	$25.16^{\pm0.378}$	$10.14^{\pm0.011}$	$10.17^{\pm0.008}$
GapONet (Ours)	$18.78^{\pm 1.147}$	$23.23^{\pm 5.245}$	$53.36^{\pm0.486}$	$25.08^{\pm0.181}$	$10.13^{\pm0.167}$	$10.14^{\pm0.017}$

Values are reported as mean with superscript \pm standard deviation (three decimals). The best result in each column is highlighted in light green and bold.

crepancy between simulation and real data), **Smoothness** (mean acceleration gap), and **Robustness** (per-joint gap with added noise). Each experiment was repeated multiple times, and the results are presented as mean and standard deviation to ensure validity. Detailed metric calculations can be found in Section A.6.

Results in Table 2 show that **GapONet** outperforms other methods in trajectory tracking, maintaining excellent performance even with payloads, and exhibiting the smallest error growth. In qualitative analysis, as shown in Figure 4, when a humanoid robot follows the same trajectory from the same starting point with identical commands, the real execution trajectory (depicted by the white lines) exhibits significant deviations. Robots without the residual model show frequent tilting and large trajectory shifts, while the policy with **GapONet** follows better. Full video demonstrations and more details can be found in Section A.6 and the supplementary material.



Figure 4: **Locomotion trajectory tracking.** (a) shows trajectory tracking using PD control, where the path (white line) deviates significantly, and the robot's torso tilts drastically, indicating instability. (b) shows the full-body motion after upper-body correction with **GapONet**. Although there is still some rightward deviation, the trajectory is much more stable, and the robot's torso remains upright.

These results collectively demonstrate the generalization and gap-solving capabilities of **GapONet**. It not only outperforms current baselines on unseen motions under different payloads but also achieves higher stability in lower-body locomotion on an unseen robot, laying the foundation for improved performance in humanoid loco-manipulation tasks.

6 Conclusion

We present an end-to-end data-collection pipeline and curate 120+ hours of paired sim-real data across multiple robots. We characterize payload-related parameters, compare sim-to-real gaps across simulators, and assess the impact of lower-body actions on whole-body behavior. We then learn a payload-conditioned nonlinear operator **GapONet** mapping simulation context functions to residual actions for hardware. On zero-shot motion tracking, the large-gap ratio is 0.09%, with improved robustness and smoothness in locomotion trajectory tracking, strengthening the basis for humanoid loco-manipulation. Future work and limitations are discussed in Section A.8.

ETHICS STATEMENT

The dataset used and planned for release in this work has been fully anonymized and does not contain any personal or individually identifiable information, but rather consists of a collection of publicly accessible content. The paper does not include any analysis, reporting, or disclosure of private user details, and care has been taken to ensure that all data handling aligns with privacy regulations and ethical guidelines.

REPRODUCIBILITY STATEMENT

We include real-world experimental footage to substantiate the reported results and release a subset of sim-real paired data for cross-validation; both are provided in the supplementary materials. Key implementation details and experimental settings are described in the main paper (Section 4, Section 5) and supplementary materials Section A.7.

REFERENCES

- 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, 2025.
- Karl Johan Åström and Peter Eykhoff. System identification—a survey. *Automatica*, 7(2):123–162, 1971.
- Filip Bjelonic and Marco Hutter. Pace dataset for sim-to-real transfer in legged robots: Joint dynamics identification and locomotion experiments across multiple robot platforms. 2025.
- Filip Bjelonic, Fabian Tischhauser, and Marco Hutter. Towards bridging the gap: Systematic simto-real transfer for diverse legged robots. *arXiv preprint arXiv:2509.06342*, 2025.
- Xiaoyu Chen, Jiachen Hu, Chi Jin, Lihong Li, and Liwei Wang. Understanding domain randomization for sim-to-real transfer. *arXiv preprint arXiv:2110.03239*, 2021.
- John J Craig. Introduction to robotics: mechanics and control, 3/E. Pearson Education India, 2009.
- Thomas Oliver de Jong, Khemraj Shukla, and Mircea Lazar. Deep operator neural network model predictive control. *arXiv preprint arXiv:2505.18008*, 2025.
- Nolan Fey, Gabriel B Margolis, Martin Peticco, and Pulkit Agrawal. Bridging the sim-to-real gap for athletic loco-manipulation. *arXiv preprint arXiv:2502.10894*, 2025.
- Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. *arXiv preprint arXiv:2408.14472*, 2024.
- 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.
- Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbab, Chaoyi Pan, et al. Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. *arXiv preprint arXiv:2502.01143*, 2025.
- Nicolas Heess, Dhruva Tb, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.

- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
 - 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.
 - Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
 - Shashank Kushwaha, Jaewan Park, Seid Koric, Junyan He, Iwona Jasiuk, and Diab Abueidda. Advanced deep operator networks to predict multiphysics solution fields in materials processing and additive manufacturing. *Additive Manufacturing*, 88:104266, 2024.
 - Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. Dart: Noise injection for robust imitation learning. In *Conference on robot learning*, pp. 143–156. PMLR, 2017.
 - Jae Yong Lee and Yeoneung Kim. Hamilton–jacobi based policy-iteration via deep operator learning. *Neurocomputing*, pp. 130515, 2025.
 - Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Robust and versatile bipedal jumping control through reinforcement learning. *arXiv* preprint *arXiv*:2302.09450, 2023.
 - Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv* preprint arXiv:2010.08895, 2020.
 - Yun Liu, Bowen Yang, Licheng Zhong, He Wang, and Li Yi. Mimicking-bench: A benchmark for generalizable humanoid-scene interaction learning via human mimicking. *arXiv* preprint *arXiv*:2412.17730, 2024.
 - Lennart Ljung. System identification. In *Signal analysis and prediction*, pp. 163–173. Springer, 1998.
 - Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv* preprint arXiv:1910.03193, 2019.
 - Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
 - Sha Luo, Hamidreza Kasaei, and Lambert Schomaker. Accelerating reinforcement learning for reaching using continuous curriculum learning. In 2020 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE, 2020.
 - Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
 - Jiageng Mao, Siheng Zhao, Siqi Song, Tianheng Shi, Junjie Ye, Mingtong Zhang, Haoran Geng, Jitendra Malik, Vitor Guizilini, and Yue Wang. Learning from massive human videos for universal humanoid pose control, 2024.
 - Gabriel B Margolis, Ge Yang, Kartik Paigwar, Tao Chen, and Pulkit Agrawal. Rapid locomotion via reinforcement learning. *The International Journal of Robotics Research*, 43(4):572–587, 2024.
 - Jan Matas, Stephen James, and Andrew J Davison. Sim-to-real reinforcement learning for deformable object manipulation. In Conference on Robot Learning, pp. 734–743. PMLR, 2018.
 - Bhairav Mehta, Manfred Diaz, Florian Golemo, Christopher J Pal, and Liam Paull. Active domain randomization. In *Conference on Robot Learning*, pp. 1162–1176. PMLR, 2020.

- AJ Miller, Fangzhou Yu, Michael Brauckmann, and Farbod Farshidian. High-performance reinforcement learning on spot: Optimizing simulation parameters with distributional measures. arXiv preprint arXiv:2504.17857, 2025.
 - Fabio Muratore, Fabio Ramos, Greg Turk, Wenhao Yu, Michael Gienger, and Jan Peters. Robot learning from randomized simulations: A review. *Frontiers in Robotics and AI*, 9:799893, 2022.
 - Kyung-Mi Na and Chang-Hun Lee. Physics-informed deep learning approach to solve optimal control problem. In *AIAA SCITECH 2024 Forum*, pp. 0945, 2024.
 - Oliver Nelles. Nonlinear system identification. *Measurement Science and Technology*, 13(4):646–646, 2002.
 - Yutao Ouyang and Jingzhi Cui. Bridging the sim-to-real gap for efficient and robust robotic skill acquisition. In *Tsinghua University Course: Advanced Machine Learning*.
 - Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
 - Xue Bin Peng, Erwin Coumans, Tingnan Zhang, Tsang-Wei Lee, Jie Tan, and Sergey Levine. Learning agile robotic locomotion skills by imitating animals. *arXiv preprint arXiv:2004.00784*, 2020.
 - Guanya Shi, Xichen Shi, Michael O'Connell, Rose Yu, Kamyar Azizzadenesheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural lander: Stable drone landing control using learned dynamics. In 2019 international conference on robotics and automation (icra), pp. 9784–9790. IEEE, 2019.
 - Jean-Jacques E Slotine and Weiping Li. On the adaptive control of robot manipulators. *The international journal of robotics research*, 6(3):49–59, 1987.
 - Nikhil Sobanbabu, Guanqi He, Tairan He, Yuxiang Yang, and Guanya Shi. Sampling-based system identification with active exploration for legged robot sim2real learning. *arXiv* preprint *arXiv*:2505.14266, 2025.
 - Mark W Spong, Seth Hutchinson, Mathukumalli Vidyasagar, et al. *Robot modeling and control*, volume 3. Wiley New York, 2006.
 - Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *Robotics: Science and Systems (RSS)*, 2018.
 - 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 (IROS)*, 2017.
 - Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4555–4576, 2021.
 - Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation, 2024.
 - Wenli Xiao, Haoru Xue, Tony Tao, Dvij Kalaria, John M Dolan, and Guanya Shi. Anycar to anywhere: Learning universal dynamics model for agile and adaptive mobility. *arXiv preprint* arXiv:2409.15783, 2024.
 - Zhaoming Xie, Patrick Clary, Jeremy Dao, Pedro Morais, Jonanthan Hurst, and Michiel Panne. Learning locomotion skills for cassie: Iterative design and sim-to-real. In *Conference on Robot Learning*, pp. 317–329. PMLR, 2020.
 - Zhaoming Xie, Xingye Da, Michiel Van de Panne, Buck Babich, and Animesh Garg. Dynamics randomization revisited: A case study for quadrupedal locomotion. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 4955–4961. IEEE, 2021.

- Jie Xu, Eric Heiden, Iretiayo Akinola, Dieter Fox, Miles Macklin, and Yashraj Narang. Neural robot dynamics. *arXiv preprint arXiv:2508.15755*, 2025.
- Bohao Zhang, Daniel Haugk, and Ram Vasudevan. System identification for constrained robots. *arXiv preprint arXiv:2408.08830*, 2024.
- Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Mingyan Liu, Duane Boning, and Cho-Jui Hsieh. Robust deep reinforcement learning against adversarial perturbations on state observations. *Advances in neural information processing systems*, 33:21024–21037, 2020.
- Xiang Zhang, Changhao Wang, Lingfeng Sun, Zheng Wu, Xinghao Zhu, and Masayoshi Tomizuka. Efficient sim-to-real transfer of contact-rich manipulation skills with online admittance residual learning. In *Conference on Robot Learning*, pp. 1621–1639. PMLR, 2023.
- Yuanhang Zhang, Yifu Yuan, Prajwal Gurunath, Tairan He, Shayegan Omidshafiei, Ali-akbar Aghamohammadi, Marcell Vazquez-Chanlatte, Liam Pedersen, and Guanya Shi. Falcon: Learning force-adaptive humanoid loco-manipulation. *arXiv preprint arXiv:2505.06776*, 2025.
- Wenshuai Zhao, Jorge Peña Queralta, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In 2020 IEEE symposium series on computational intelligence (SSCI), pp. 737–744. IEEE, 2020.

A APPENDIX

A.1 THE USE OF LARGE LANGUAGE MODELS (LLMS)

We employed LLMs only for grammar/style rewrites and equation/notation formatting corrections. We appreciate the steadily improving reasoning capabilities of LLMs, which helped us identify linguistic issues more quickly and maintain a more consistent scholarly style. However, all research ideation, theoretical development and formula derivations, methodological choices, and experimental design and execution were performed exclusively by the authors. Accordingly, the LLM did not play a significant role in research ideation or writing and should not be regarded as a contributor.

A.2 DATA COLLECTION

A.2.1 LEGGED HUMANOID ROBOT

We collect paired sim—real data on two humanoids: the 1.8 m Unitree H1-2 and the 1.3 m Unitree G1. Joint naming and kinematic locations are shown in Figure 5. In our setup, we log the full upper body and locomotion-relevant joints (27-DoF configuration in code), along with IMU and actuator telemetry.

ROS setup and topics Data acquisition is implemented as a ROS 2 Python node (rclpy, node name deploy_node). The node subscribes to low-level robot state messages and publishes torque/position commands:

- Subscriptions: LowState (joint positions/velocities/currents, IMU, wireless remote), used to buffer sensor streams and teleop events.
- Publications: LowCmd on topic lowcmd_buffer at 50 Hz (control period $\Delta t \approx 20 \, \mathrm{ms}$). Commands include per-joint PD terms and optional feedforward residuals (CRC is appended before transmission).

Teleoperation triggers (e.g., start/stop, emergency stop) are parsed from the wireless controller and gate recording and command streaming.

What is recorded For each trial, we write files (per-trial timestamped) with the following datasets, matching the code:

- command_time_list (s): wall-clock times when commands are produced.
- command_val_list: commanded action vectors (per 20 ms tick).
- robot/joint_time_list (s): time stamps associated with the sensed robot state.
- robot/joint_angle_list, robot/joint_velocity_list, robot/joint_current_list, robot/joint_temperature_list: actuator telemetry.
- robot/imu_list, robot/ang_vel_list: IMU linear orientation proxies and angular rates.
- motion_name, current_time: metadata for the retargeted/teleop motion and file creation time.

Spatiotemporal synchronization We use a single monotonic clock started at node initialization to time-stamp both the command loop and the sensor callback buffers. During acquisition, the node executes a fixed-rate control loop (50 Hz) and performs rclpy.spin_once with a short timeout each tick; the current monotonic time is appended to both command_time_list and robot/joint_time_list. This yields frame-accurate alignment between the actuation stream and the sensed state at the controller cadence. Since logging and control are co-located on the same machine, no cross-machine NTP is required; residual jitter is bounded by the loop period and handled in post-processing by resampling to a common time base when needed.

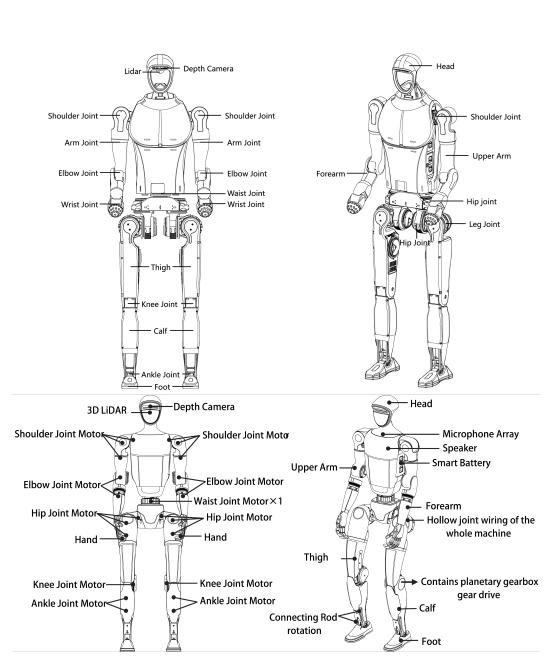


Figure 5: Joint names and positions on Unitree H1-2 and G1 robots

Libraries The implementation relies on rclpy (ROS 2), numpy, torch (policy inference/logging utilities), mujoco (simulation), h5py (file I/O), and transforms3d (frame utilities). All topics and message types (LowState, LowCmd, MotorState, IMUState) come from the unitree_hg.msg package.

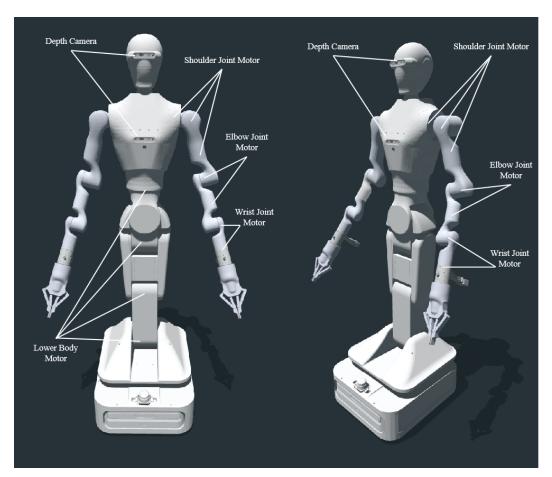


Figure 6: Joint names and positions on RealMan WR75S robot

A.2.2 WHEELED HUMANOID ROBOT

We also collect motion execution data on dual-arm wheeled robots (RealMan). Our setup logs the full arm joint configurations along with actuator telemetry through UDP communication using the official RealMan API.

Communication Setup Data acquisition uses the RealMan official API with UDP communication. Position commands are sent to each arm at dedicated ports (8080, 8576), while real-time state data is received through UDP callbacks on separate ports (8089, 8090). The system registers callback functions to process telemetry streams containing joint positions, velocities, currents, and temperatures.

Data Recording Structure For each trial, we save timestamped datasets in HDF5 format with the following structure matching our dual-arm configuration:

- command_time_list (s): wall-clock timestamps when commands are issued.
- command_val_list: commanded action vectors for both arms concatenated (14-dimensional for dual 7-DoF arms).

- 864
- 866 867
- 869
- 870 871
- 872 873 874
- 875 876
- 877 878 879
- 882
- 883 885
- 887
- 888
- 889 890

894 895

897 899

900

901

902

903 904 905

906

907

908 909

910 911

912 913

914 915

916 917 robot1/joint_time_list, robot2/joint_time_list (s): sensor timestamps

for left and right arms respectively.

• robot1/joint_angle_list, robot2/joint_angle_list: joint positions in radians for each arm.

• robot1/joint_velocity_list, robot2/joint_velocity_list: joint velocities in rad/s for each arm.

- robot1/joint_current_list, robot2/joint_current_list: motor currents for each arm.
- robot1/joint_temperature_list, robot2/joint_temperature_list: actuator temperatures for each arm.
- motion_name, slowdown_factor, current_time: metadata for trial identification.

Spatiotemporal synchronization We employ a unified monotonic clock initialized at data collection start to timestamp both command transmission and sensor reception. During execution, commands are sent via rm_movej_canfd API calls while the monotonic timestamp is recorded for both command and sensor streams. Since both command generation and sensor processing occur on the same machine with shared timing, cross-machine synchronization is unnecessary. The UDP callback mechanism ensures frame-accurate alignment between actuation commands and sensed states at the controller frequency. Residual timing jitter is bounded by the loop period and handled through post-processing resampling when temporal alignment is required for analysis. The system continuously monitors joint enable flags and error codes, with joint disable events prioritized as critical errors and other malfunctions classified as general errors, triggering immediate data cleanup and graceful termination.

A.3 GAP ANALYSIS

A.3.1PD CONTROL

We use a basic joint-space proportional-derivative controller to track commanded trajectories with low latency. The proportional term corrects position error (stiffness), and the derivative term provides damping to reduce overshoot:

$$\tau = K_p \left(q_{\text{cmd}} - q \right) + K_d \left(\dot{q}_{\text{cmd}} - \dot{q} \right). \tag{10}$$

Here $q_{\rm cmd}$ and $\dot{q}_{\rm cmd}$ are the desired joint position/velocity, q and \dot{q} are the measured states, and K_p, K_d (typically diagonal, positive) set tracking stiffness and damping. Optional gravity/feedforward terms can be added when needed, but the above is the minimal PD law.

In equation 1, $K_p(q_{\rm cmd}-q)+K_d(\dot{q}_{\rm cmd}-\dot{q})$ is the standard joint-space PD action (typically diagonal gains). The extra linear terms $K_v \dot{q}$ and $K_c \tanh(\dot{q}/\varepsilon)$ model viscous damping and smoothed Coulomb friction, respectively; $\varepsilon > 0$ regularizes the sign function to avoid chattering. The scalar (or diagonal) P denotes the payload descriptor (e.g., mass/COM proxy). The bias $K_{payload}$ P provides a load-dependent offset, while $K_{P\sin}P\sin q$ and $K_{P\cos}P\cos q$ capture load-scaled gravity/-COM components in joint coordinates. Velocity/acceleration couplings $K_{P\dot{q}}P\dot{q}$ and $K_{P\ddot{q}}P\ddot{q}$ address payload-amplified damping/inertial effects. The constant τ_0 compensates residual biases (e.g., calibration offsets).

Start from PD only (K_p, K_d) , add K_v, K_c to reduce overshoot and stick-slip, then introduce $K_{\rm payload}, K_{P\sin}, K_{P\cos}$ for static/load gravity, and $K_{P\dot{q}}, K_{P\ddot{q}}$ for dynamic load effects; keep all gains bounded and ε small enough to smooth $\tanh(\cdot)$ without degrading response.

A.3.2 More analysis results

We present additional qualitative results here Figure 7 and Figure 8; further videos are provided in the supplementary materials.

A.4 NONLINEAR OPERATOR

What is an operator? In contrast to learning a finite-dimensional mapping $f: \mathbb{R}^n \to \mathbb{R}^m$, operator learning targets a mapping between function spaces, $\mathcal{G}:\mathcal{U}\to\mathcal{V}$, where the input $u\in\mathcal{U}$ is itself

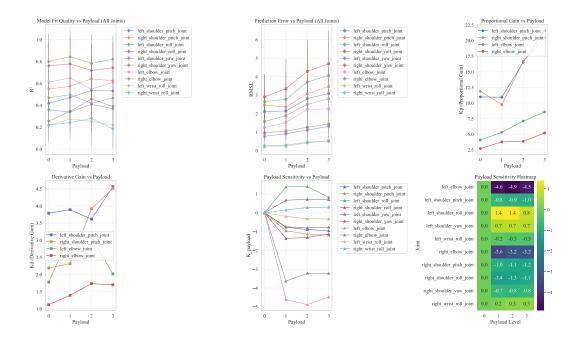


Figure 7: Data analysis on payload-related parameters

a function and the output $\mathcal{G}(u) \in \mathcal{V}$ is another function. Practically, we observe u via its sensor samples at locations $\{x_i\}_{i=1}^m$: $\{u(x_i)\}$, and we query the output at arbitrary y-locations to obtain values $\mathcal{G}(u)(y)$. This setup makes the learning objective function-to-function rather than pointwise regression, and enables generalization to unseen inputs u and query points y.;

Why not "learn a function" directly? Classical approximation fits (x,y) pairs for one target function. Operator learning instead aims to recover the rule that maps any admissible input function u to an output function $\mathcal{G}(u)$. To make this learnable from data, we draw a diverse family of input functions—e.g., samples from Gaussian Random Fields (SE/RBF kernels with tunable length-scales/variances) and orthogonal polynomial expansions (e.g., Chebyshev with random coefficients)—so the model is trained across a rich subset of $\mathcal U$ rather than around a single curve. This ensures the learned mapping reflects an operator over a function class, not merely a single function fit.

Low-rank/separable viewpoint Many learned operators can be written (or approximated) in a separable, low-rank form

$$\widehat{\mathcal{G}}(u)(y) = \sum_{k=1}^{p} b_k(u) t_k(y), \tag{11}$$

where $b_k(u)$ are functionals of the input function (computed from its samples) and $t_k(y)$ are basis functions over the query variable y. This mirrors RKHS/separable-kernel and POD/SVD intuitions and clarifies the roles of "encode the input function" versus "encode the query location.";

We adopt this operator perspective to learn **GapONet**, a mapping from simulation context functions to hardware-space responses, so that the model predicts an output function of state/time given an input function describing simulated context—setting the stage for the DeepONet factorization introduced next.

A.5 METHODS

A.5.1 WHY DO WE CHOOSE DEEPONET?

Our operator must (i) ingest simulation context functions with explicit payload conditioning, (ii) answer at arbitrary query points (current actions, payload) across heterogeneous robots and sim-

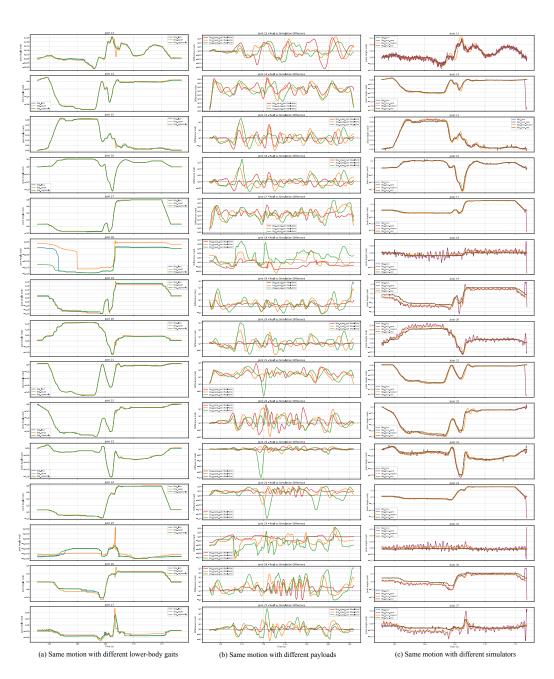


Figure 8: Results on all upper-body joints about the same motion with different payloads, simulations, and lower-body gaits.

ulators, (iii) train under a closed-loop RL objective without requiring paired function-to-function supervision at every query, and (iv) support low-latency on-board inference.

We have considered some alternatives and trade-offs, for example:

- Fourier/Neural Operators (FNO family) (Li et al., 2020; Kovachki et al., 2023): excel on fixed grids with spectral convolutions, but rely on discretization tied to resolution/geometry; cross-morphology deployment (different joint layouts) typically needs regridding or retraining, and spectral blocks add latency on embedded hardware.
- Graph/Galerkin/UNO-style operators (Kovachki et al., 2023): adapt to irregular meshes/graphs but require topology-aligned parameterization; when robots or sensor layouts change, weights/graphs must be remapped. Querying arbitrary state-time points is less natural than function-query separation. Capacity is high, but so are data and compute demands.
- Physics-informed neural operators (PINO): leverage known PDE residuals for sample efficiency, yet our residual field (sim—real actuation gap with delays/saturation) lacks a clean PDE form, making hard constraints difficult to specify and risking model-bias.

As for DeepONet's branch-trunk decomposition (Lu et al., 2019; 2021) aligns directly with our problem: the branch encodes context (multi-sensor histories, simulator traces, payload), and the trunk indexes continuous query variables (state/time/joint), producing residual action/torque values via a simple inner product. This yields (1) continuous space—time queries without grid lock-in, (2) clean conditioning on payload and robot-specific context without graph/topology rewiring, (3) RL-friendly training since supervision can be placed at arbitrary queried points along closed-loop rollouts, and (4) low-latency deployment because inference reduces to lightweight embeddings plus an inner product. Moreover, DeepONet comes with an operator-level universal approximation theorem that provides formal capacity guarantees for nonlinear operators (Lu et al., 2021), which we found attractive given the diversity of simulators, payloads, and hardware.

In summary, we choose DeepONet because its function—query factorization, theoretical operator approximation guarantees, and efficient, payload-conditioned querying match our requirements better than grid-bound spectral operators, topology-coupled graph variants, or physics-informed schemes that presume known PDE structure (Lu et al., 2019; 2021; Li et al., 2020; Kovachki et al., 2023). Our objective is to demonstrate that operator learning can achieve a mapping from simulation to reality, thereby aiding sim-to-real transfer. Determining the optimal operator architecture is outside the main scope of this work.

A.6 EXPERIMENT

A.6.1 METRICS

We report two metric families: (i) gap distribution (Table 1: large-gap ratio(LGR), interquartile range (IQR), and gap range) and (ii) kinematic quality of lower-body (Table 2: smoothness, trajectory consistency, and robustness). All metrics are computed per run and then aggregated by payload mass (the environment groups trials by mass buckets).

Let $q_t^{\mathrm{real}}, q_t^{\mathrm{sim}}$ be joint trajectories (or end-effector signals) sampled at uniform Δt . Define the gap $g_t = q_t^{\mathrm{real}} - q_t^{\mathrm{sim}}$ and its absolute value $|g_t|$. Central-difference operators approximate derivatives.

Large-gap ratio (**Table 1**) Fraction of samples with absolute joint error exceeding a threshold (0.5 rad by default):

Large-gap ratio =
$$\frac{\left|\{(t,i): |g_{t,i}| \ge \tau\}\right|}{\left|\{(t,i)\}\right|}, \quad \tau = 0.5 \text{ rad.}$$
 (12)

Captures the frequency of serious deviations.

Gap IQR (Table 1) Dispersion of absolute errors via the interquartile range:

$$\mathcal{G} = \{ |g_{t,i}| : t = 1, \dots, T, \ i = 1, \dots, J \}, \qquad \text{IQR} = Q_{0.75}(\mathcal{G}) - Q_{0.25}(\mathcal{G}). \tag{13}$$

Lower is a tighter error distribution.

Gap range (Table 1) Extreme-case spread of absolute errors:

$$Range = \max(|g|) - \min(|g|). \tag{14}$$

Highlights worst-case variability.

Trajectory consistency (Table 2) Discrepancy in the *rate-of-change of velocity* (a curvature-like signal) between real and simulated motion:

$$v_{-}t^{\text{real}} = \nabla q_{-}t^{\text{real}}, \quad v_{-}t^{\text{sim}} = \nabla q_{-}t^{\text{sim}}, \quad \kappa_{-}t^{\text{real}} = \nabla v_{-}t^{\text{real}}, \quad \kappa_{-}t^{\text{sim}} = \nabla v_{-}t^{\text{sim}}, \quad (15)$$

TrajectoryConsistency =
$$\frac{1}{T} \sum_{t=1}^{T} \left| \kappa_t^{\text{real}} - \kappa_t^{\text{sim}} \right|$$
. (16)

Smaller values indicate that the simulator reproduces the evolution of motion patterns more faithfully.

Smoothness (Table 2) Discrepancy in *accelerations* between real and simulated trajectories:

Smoothness =
$$\frac{1}{T} \sum_{t=1}^{T} \left| a_t^{\text{real}} - a_t^{\text{sim}} \right|, \qquad a_t^{\text{real}} = \nabla^2 q_t^{\text{real}}, \quad a_t^{\text{sim}} = \nabla^2 q_t^{\text{sim}}.$$
 (17)

Lower scores mean closer kinematic smoothness to real motion.

Robustness (Table 2) Sensitivity of the sim–real gap to *measurement noise*. For noise levels $\sigma \in \{\sigma_1, \dots, \sigma_K\}$,

Robustness =
$$\frac{1}{K} \sum_{k=1}^{K} \left[\frac{1}{T} \sum_{t=1}^{T} \left| \left(q_t^{\text{real}} + \epsilon_t^{(k)} \right) - \left(q_t^{\text{sim}} + \tilde{\epsilon}_t^{(k)} \right) - g_t \right| \right], \tag{18}$$

$$g_t = q_t^{\text{real}} - q_t^{\text{sim}}, \quad \epsilon_t^{(k)}, \tilde{\epsilon}_t^{(k)} \sim \mathcal{N}(0, \sigma_k^2). \tag{19}$$

Smaller values indicate that the evaluation is stable under realistic perturbations.

Each motion is run at least six times. For each run, we compute every metric (optionally per joint and then averaged); otherwise, only real-stream statistics are used as specified by each metric. We then aggregate runs by payload/mass buckets and report means with standard errors. All three metrics are discrepancy-style measures; by construction, **smaller values indicate better performance**.

A.6.2 LOCOMOTION TRAJECTORY TRACKING

We generate locomotion commands using a phase-based trajectory: a normalized phase $\phi \in [0,1)$ advances at the control rate and indexes a trapezoidal base-velocity profile (accelerate-cruise-decelerate-pause). Forward and backward segments alternate automatically, while lateral velocity and yaw rate remain zero unless specified. The phase schedules lower-body gait timing and yields desired joint trajectories for the legs, tracked by a joint-space PD controller at 50 Hz with torque/rate limits and safety checks.

Fixed start pose and heading. Each real-robot run starts from the same world-frame pose—a fixed position and heading—followed by a short smooth interpolation into the nominal stand pose before the phase route is enabled. This ensures repeatable initial conditions, so the resulting base trajectory in SE(2) (odometry or motion-capture) can be compared across runs to assess tracking quality, drift, and sim—real alignment. Commands and sensor streams share a monotonic timestamp, keeping phase, velocity setpoints, and measured joint/IMU signals time-aligned for evaluation.

A.7 IMPLEMENTATION DETAILS

A.7.1 NETWORK STRUCTURE

Overview. The training pipeline with **GapONet** consists of three components: a Sensor Predictor to predict the sensor input of Branch Network, a Branch Network $\mathcal{B}(U_q(x))$ that encodes sensor-driven actuation functions and a Trunk Network $\mathcal{T}(y)$ that processes action queries. Both are implemented as multi-layer perceptrons (MLPs), fused via dot product to yield the operator output $\mathcal{G}(U_q(x))(y)$. These networks are trained end-to-end with Proximal Policy Optimization (PPO), and optimized using Adam.

1134 **Sensor Predictor** 1135 1136 • **Input:** For each time j at time step t, the Sensor Predictor receives a sequence of sensor 1137 states over a k-step history window: 1138 $\{q_j^{t-n}, \dot{q}_j^{t-n}, q_{j,d}^{t-n}\}_{n=0}^k,$ 1139 1140 where q_i , \dot{q}_i denote joint position and velocity, $q_{i,d}$ is the target position. 1141 • History Length: k=41142 • Input Dimension: 10 joint num \times (3 \times history length + 1 current position) = 130-dim 1143 1144 1145 • Output: $\Delta q \& \Delta \dot{q} \times 10$ joint = 20-dim vector 1146 • Sensor Number: 20 • Learning Rate: 1×10^{-4} 1148 1149 **Branch Net.** 1150 1151 • Input: 20-dim vector of sensor predictor output × 20 sensor num = 400-dim vector 1152 • Delta Action Duration: 1 step 1153 1154 • Architecture: 4-layer MLP with hidden sizes [256, 256, 256], each followed by ELU acti-1155 vation. 1156 • Output: p-dimensional latent representation (p = 160 by default) 1157 • Learning Rate: 1×10^{-4} 1158 1159 Trunk Net. 1160 1161 • Input: The Trunk Net receives the target query $y = q_{i,d}^{t+1}$ desired joint position + payload 1162 • Input Dimension: 11 1163 1164 • Architecture: 2-layer MLP with hidden sizes [128, 128, 128], ELU activations 1165 • Output: p-dimensional vector, same dimension as Branch output 1166 • Learning Rate: 1×10^{-4} 1167 1168 **Fusion.** The operator output is computed as the dot product: 1169 1170 $\mathcal{G}(U_q(x))(y) = \sum_{i=1}^{p} \mathcal{B}_i(x) \cdot \mathcal{T}_i(y).$ 1171 1172 1173 Training Details. 1174 1175 • PPO update with clipping ratio $\epsilon = 0.2$, batch size = 4096. 1176 • Reward defined as $r_t = -\|q^{t+1} - q_{\text{real}}^{t+1}\|^2$. 1177 1178 • Temporal smoothness penalty \mathcal{L}_{gap} with $\lambda = 0.01$. 1179 • Training duration: 1 hour on 1 RTX 3090Ti GPU. 1180 1181 1182 A.7.2 SIMULATIONS 1183 We evaluate on MuJoCo 3.2.3, Isaac Gym 1.0rc4, and Isaac Sim 4.5.0. To enhance reproducibility, 1184 each setting uses the simulator's official default parameters. The software environments are: 1185 1186 MuJoCo / Isaac Gym: Python 3.8.13, legged_gym 1.0.0, PyTorch 2.4.1, torchvision 0.19.1. 1187

• Isaac Sim: Python 3.10.4, isaaclab 0.40.21, PyTorch 2.5.1, torchvision 0.20.1.

Table 3: Hyperparameters for Branch Net.

A.8 LIMITATION AND FUTURE WORK

Our dataset and analysis primarily target the upper body, and although we include tests on locomotion trajectory tracking, the present system does not yet enable highly dynamic sim-real transfer for full humanoids. Going forward, we will (i) extend the current pipeline to high-dynamics, whole-

Hyper-Parameters Values History Length **Delta Action Duration** Sensor Number A, V, P, J U_q Input ΔS U_q Output [256, 256, 128]Layer Structure Output Number 0.1 Dropout Samples Per Update Iteration Policy/Value Function Minibatch Size Discriminators/Encoder Minibatch Size γ Discount 0.99 2×10^{-5} Learning Rate 0.95 $GAE(\lambda)$ $TD(\lambda)$ 0.95 PPO Clip Threshold 0.2 T Episode Length

Table 4: Hyperparameters for Trunk Net.

Hyper-Parameters	Values	
History Length	4	
Delta Action Duration	1	
Sensor Number	20	
y Input	a_d	
Layer Structure	[128, 128]	
Output Number	10	
Dropout	0.1	
Samples Per Update Iteration	131072	
Policy/Value Function Minibatch Size	16384	
Discriminators/Encoder Minibatch Size	4096	
γ Discount	0.99	
Learning Rate	2×10^{-5}	
$GAE(\lambda)$	0.95	
$TD(\lambda)$	0.95	
PPO Clip Threshold	0.2	
T Episode Length	300	

body loco-manipulation and to additional robot platforms, and (ii) address the strong dependence on a stable locomotion policy—even with relative metrics, unreliable gaits can cause catastrophic failures (cf. 'videos/failure.mp4') that preclude testing. A second focus is to train a robust full-body tracker for large-mass humanoids (e.g., H1-2), providing a stronger substrate for our operator-based sim—real mapping.