

# Appendix

## A Implementation Details

This section outlines the procedures for generating diverse robot morphologies, randomizing environment dynamics, designing rewards, and training the LocoFormer policy.

### A.1 Procedural Robot Generation

We generate four categories of morphologies: quadrupeds, bipeds, and their respective wheeled counterparts. Random samples are illustrated in Fig. 1.

**Quadrupeds** Quadrupeds are constructed with hip pitch, hip roll, and knee joints, where the hip roll joints are randomly fixed. We randomize link length, mass, joint positions, and joint limits. A global scaling factor  $\mu$  uniformly scales sizes and quadratically scales masses. Two baseline configurations are included: A1-style (all knees bending backward) and AnyMalC-style (rear knees bending forward). The resulting quadrupeds vary from small robots ( $0.19 \times 0.14 \times 0.27$  m, 3 kg) to large structures ( $1.2 \times 0.66 \times 0.96$  m, 64 kg).

**Bipeds** Bipeds feature hip roll/pitch/yaw, knee, and ankle roll/pitch joints, with hip roll/yaw and ankle joints randomly fixed. We similarly randomize kinematic and dynamic parameters, including a variably sized trunk. Foot designs include human-like flat feet and point feet (TRON1-PF style). Leg configurations can bend forwards or backwards. Generated bipeds range from  $0.08 \times 0.12 \times 0.75$  m (5 kg) to  $0.32 \times 0.92 \times 3.20$  m (196 kg).

**Wheeled Quadrupeds and Bipeds** Wheeled quadrupeds and bipeds are derived by replacing feet with wheels of radius 5–15 cm, scaled by  $\mu$ .

While our procedural generation covers a wide morphological space, it is not exhaustive. Our aim is to train a policy capable of generalizing to out-of-distribution robots. In practice, unseen robots often deviate from our design assumptions. For instance, Berkeley-Humanoid exhibits anedral angles in both hip roll and yaw joints—our bipeds include at most one. Additionally, robots like G1 have joint axes offset along the  $x$ -axis, a feature not captured in our parameterization.


### A.2 Domain Randomization

To enhance task diversity for adaptation, we apply extensive domain randomization across physical and environmental parameters. Table 1 summarizes the maximum ranges used.


To accommodate robots with diverse masses, we define a base stiffness and damping per morphology, scaled proportionally to each robot’s mass. Motor strength further randomizes joint-level stiffness and damping, producing varied MDPs.

Parameters	Range
Center of mass (quadruped trunk)	$[-0.2, 0.2]$
Center of mass (bipeds trunk)	$[-0.05, 0.05]$
Center of mass (other links)	$[-0.03, 0.03]$
Motor strength <sup>†</sup>	$[0.5, 2.0]$
Payload <sup>†</sup>	$[0, 10]$
friction <sup>†</sup>	$[0.4, 1.6]$
Action latency	$\{0, 0.02\}$
Action EMA	$[0.5, 1.0]$


Table 1: Maximum ranges for domain randomization. Parameters marked with <sup>†</sup> are randomized with half intensity during training and full intensity at evaluation to assess out-of-distribution performance.




(a)



(b)



(c)



(d)

Figure 1: Visualization of sample procedurally generated robots. (a) Quadrupeds. (b) Quadrupeds with wheels. (c) Bipeds. (d) Bipeds with wheels.

### 32 A.3 Reward Design

33 We design a unified reward structure combining morphology-agnostic and morphology-specific terms.  
 34 For clarity, the time subscript  $t$  is omitted in some expressions.

#### 35 Shared Reward Components

- 36 • **Linear velocity command tracking**  $\exp(\|v_{xy} - v_{xy}^{cmd}\|^2/s_1)$
- 37 • **Angular velocity command tracking**  $\exp(\|w_z - w_z^{cmd}\|^2/s_2)$
- 38 • **Base linear velocity penalty**  $\|v_z\|^2$
- 39 • **Base angular velocity penalty**  $\|w_{xy}\|^2$
- 40 • **Base orientation penalty**  $\|R \cdot [0, 0, 1]\|_{xy}^2$
- 41 • **Base height penalty**  $\|x_z - x_z^{nominal}\|^2$
- 42 • **Joint acceleration penalty**  $\|\ddot{q}\|^2$
- 43 • **Torque penalty**  $-\|\tau\|^2$
- 44 • **Torque change penalty**  $-\|\dot{\tau}\|^2$
- 45 • **Penalty for movement of selected joints**  $\|q_s - q_s^{nominal}\|^2$
- 46 • **Penalty for contact on undesired links**  $\|f_c\|^2$
- 47 • **Penalty for actions in unused joints**  $\|a_{unused}\|^2$
- 48 • **Alive reward:** 1

49 These terms are scaled by 1.0, 0.5, -1.0, -0.05, -1.0, -10.0, -2.5e-7, -1e-4, -1e-5, -1e-7, -1.0, -1.0, -0.01,  
 50 0.4. Specifically, *penalty for movement of selected joints* is applied to hip roll joints for quadrupeds,  
 51 hip roll/yaw joints for bipeds (abandon this for bipeds with anhedral angles in hip), and all non-wheel  
 52 joints in wheeled robots.

#### 53 Morphology-Specific Terms

- 54 • **Slipping foot penalty**  $\|v_{xy}\|^2 \times [f > 0]$  (weighted by -0.2 for bipeds, -0.1 for quadrupeds)
- 55 • **Foot air time reward:** When a foot makes ground contact after being airborne, we reward  
 56  $t_{\text{on-the-air}} - 0.5$  (quadrupeds: 0.5 weight)
- 57 • **Foot orientation penalty:** Penalizes misalignment between foot and trunk orientation  
 58 (bipeds: -0.1 weight)

### 59 A.4 Training Details

60 We train our policy using 128 GPUs, each simulating 2048 parallel environments. The GPUs are  
 61 allocated in a 1:1:3:3 ratio across quadrupeds, wheeled quadrupeds, bipeds, and wheeled bipeds,  
 62 respectively. The checkpoint used for evaluation was trained on rough terrain across all morphologies.  
 63 However, we observed significant ground penetration issues in biped robots under these conditions,  
 64 resulting in unnatural ankle pose. To address this, future versions of the policy will use flat terrain  
 65 specifically for bipeds.

## 66 B Experiment Details

67 Fig. 2 shows the 10 simulated robots used for comparison against baseline methods. Among them,  
 68 we observe a notable sign of cross-trial adaptation in the TRON1 robot, which achieves a 10%  
 69 performance improvement through 5-seconds pre-evaluation adaptation. We find the robot initially  
 70 fails in early trials but gradually learns to stabilize and control itself effectively. An illustrative  
 71 example of this adaptation process is shown in Fig. 3.

72 We provide videos of real world evaluation on [generalist-locomotion.github.io](https://generalist-locomotion.github.io).




Figure 2: Robots where we compare with the baseline methods in simulation. First row: Unitree G1, H1, Fourier GR1, LimX Dynamics TRON1, Berkeley Humanoid. Second row: Unitree A1, Boston Dynamics Spot, ETH AnyMal C, LimX Dynamics TRON1-WF, Unitree Go2-W.




Figure 3: Cross-trial adaptation of TRON1 in simulation. The robot fails in early trials but progressively learns to stabilize and locomote effectively by Trial 4.