

BeamDojo: Learning Agile Humanoid Locomotion on Sparse Footholds

Huayi Wang^{1,2} Zirui Wang^{1,3} Junli Ren^{1,4} Qingwei Ben^{1,5} Tao Huang^{1,2}

Weinan Zhang^{1,2,†} Jiangmiao Pang^{1,†}

¹Shanghai AI Laboratory ²Shanghai Jiao Tong University ³Zhejiang University

⁴The University of Hong Kong ⁵The Chinese University of Hong Kong

[†]Corresponding Authors.

Website: <https://why618188.github.io/beamdojo>




Fig. 1: Our proposed framework, **BEAMDOJO**, enables agile and robust humanoid locomotion across challenging sparse footholds. **Top Row:** Utilizing global mapping information from LiDAR, the humanoid demonstrates the ability to precisely traverse both forward and backward on stepping stones, with arrows indicating the direction of movement. **Bottom Left:** The humanoid skillfully traverses a narrow balance beam. **Bottom Middle:** Despite being trained without exposure to gaps and balancing beams, the humanoid achieves zero-shot generalization to various sparse foothold terrains. **Bottom Right:** Humanoid exhibits remarkable robustness, maintaining stable locomotion under external disturbances and additional payloads.

Abstract—Traversing risky terrains with sparse footholds poses a significant challenge for humanoid robots, requiring precise foot placements and stable locomotion. Existing learning-based approaches often struggle on such complex terrains due to sparse foothold rewards and inefficient learning processes. To address these challenges, we introduce **BEAMDOJO**, a reinforcement learning (RL) framework designed for enabling agile humanoid locomotion on sparse footholds. **BEAMDOJO** begins by introducing a sampling-based foothold reward tailored for polygonal feet, along with a double critic to balancing the learning process between dense locomotion rewards and sparse foothold rewards. To encourage sufficient trial-and-error exploration, **BEAMDOJO** incorporates a two-stage RL approach: the first stage relaxes the terrain dynamics by training the humanoid on flat terrain while providing it with task-terrain perceptive observations, and the second stage fine-tunes the policy on the actual task terrain. Moreover, we implement an onboard LiDAR-based elevation map to enable real-world deployment. Extensive simulation and real-world experiments demonstrate that **BEAMDOJO** achieves

efficient learning in simulation and enables agile locomotion with precise foot placement on sparse footholds in the real world, maintaining a high success rate even under significant external disturbances.

I. INTRODUCTION

Traversing risky terrains with sparse footholds, such as stepping stones and balancing beams, presents a significant challenge for legged locomotion. Achieving agile and safe locomotion on such environment requires robots to accurately process perceptive information, make precise footstep placement within safe areas, and maintain base stability throughout the process [60, 63].

Existing works have effectively addressed this complex task for quadrupedal robots [15, 16, 27, 35, 55, 60, 61, 63, 64]. However, these methods encounter great challenges when applied to humanoid robots, primarily due to a key difference

in foot geometry. Although the foot of most quadrupedal robots and some simplified bipedal robots [9, 30] can be modeled as a point, the foot of humanoid robots is often represented as a polygon [5, 17, 24, 49]. For traditional model-based methods, this requires additional half-space constraints defined by linear inequalities, which impose a significant computational burden for online planning [7, 8, 17, 38, 49]. In the case of reinforcement learning (RL) methods, foothold rewards designed for point-shaped feet are not suitable for evaluating foot placement of polygon-shaped feet [60]. Hybrid methods, which combine RL with model-based controllers, face similar challenges in online planning for humanoid feet [15, 27, 55]. Furthermore, the higher degrees of freedom and the inherently unstable morphology of humanoid robots make it even more difficult to achieve agile and stable locomotion over risky terrains.

On the other hand, recent advancements in learning-based humanoid robot locomotion have demonstrated impressive robustness across various tasks, including walking [2, 45, 46, 52], stair climbing [6, 18, 34], parkour [66], and whole-body control [3, 13, 19, 20, 22, 28, 29], etc. However, these methods still struggle with complex terrains and agile locomotion on fine-grained footholds. Enabling agile movement on risky terrains for humanoid robots presents several challenges. First, the reward signal for evaluating foot placement is sparse, typically provided only after completing a full sub-process (e.g., lifting and landing a foot), which makes it difficult to assign credit to specific states and actions [50]. Second, the learning process is highly inefficient, as a single misstep often leads to early termination during training, hindering sufficient exploration. Additionally, obtaining reliable perceptual information is challenging due to sensory limitations and environmental noise [66].

In this work, we introduce BEAMDOJO, a novel reinforcement learning-based framework for controlling humanoid robots traversing risky terrains with sparse footholds. The name BEAMDOJO combines the words “beam” (referring to sparse footholds such as beams) and “dojo” (a place of training or learning), reflecting the goal of training agile locomotion on such challenging terrains. We begin by defining a sampling-based foothold reward, designed to evaluate the foot placement of a polygonal foot model. To address the challenge of sparse foothold reward learning, we propose using double critic architecture to separately learn the dense locomotion rewards and the sparse foothold reward. Unlike typical end-to-end RL methods [60, 63], BEAMDOJO further incorporates a two-stage approach to encourage fully trial-and-error exploration. In the first stage, terrain dynamics constraints are relaxed, allowing the humanoid robot to practice walking on flat terrain while receiving perceptive information of the target task terrain (e.g., sparse beams), where missteps will incur a penalty but do not terminate the episode. In the second stage, the policy is fine-tuned on the true task terrain. To enable deployment in real-world scenarios, we further implement a LiDAR-based, robot-centric elevation map with carefully designed domain randomization in simulation training.

As shown in Fig. 1, BEAMDOJO skillfully enables humanoid robots to traverse risky terrains with sparse footholds, such as stepping stones and balancing beams. Through extensive simulations and real-world experiments, we demonstrate the efficient learning process of BEAMDOJO and its ability to achieve agile locomotion with precise foot placements in real-world scenarios.

The contributions of our work are summarized as follows:

- We propose BEAMDOJO, a two-stage RL framework that combines a newly designed foothold reward for the polygonal foot model and a double critic, enabling humanoid locomotion on sparse footholds.
- We implement a LiDAR-based elevation map for real-world deployment, incorporating carefully designed domain randomization in simulation training.
- We conduct extensive experiments both in simulation and on Unitree G1 Humanoid, demonstrating agile and robust locomotion on sparse footholds, with a high zero-shot sim-to-real transfer success rate of 80%. To the best of our knowledge, BEAMDOJO is the first learning-based method to achieve fine-grained foothold control on risky terrains with sparse footholds.

II. RELATED WORKS

A. Locomotion on Sparse Footholds

Walking on sparse footholds has been a long-standing application of perceptive legged locomotion. Existing works often employ model-based hierarchical controllers, which decompose this complex task into separate stages of perception, planning, and control [16, 17, 26, 40, 41, 43, 53]. However, model-based controllers react sensitively to violation of model assumptions, which hinders applications in real-world scenarios. Recent studies have explored combining RL with model-based controllers, such as using RL to generate trajectories that are then tracked by model-based controllers [15, 61, 55], or employing RL policies to track trajectories generated by model-based planners [27]. While demonstrating remarkable performance, these decoupled architectures can constrain the adaptability and coordination of each module.

Subsequent works have explored end-to-end learning frameworks that train robots to walk on sparse footholds using perceptive locomotion controllers [1, 4, 59, 60, 63]. Despite their focus being limited to quadrupeds, a majority of these works rely on depth cameras for exteroceptive observations, which are limited by the camera’s narrow field of view and restrict the robot to moving backward [1, 4, 59, 60]. Additionally, an image processing module is often necessary to bridge the sim-to-real gap between the captured depth images and the terrain heightmap used during training [1, 4, 59, 60].

In contrast to the aforementioned literature, this work achieves agile humanoid locomotion over risky terrains that addressing unique challenges specific to humanoid systems, such as foot geometry. Additionally, we implement a LiDAR-based elevation map to enhance the task, demonstrating that the robot can move smoothly both forward and backward using the robotics-centric elevation map as the perception module.




Fig. 2: Foothold Reward. We sample n points under the foot. Green points indicate contact with the surface within the safe region, while red points represent those not in contact with the surface.

B. Reinforcement Learning in Locomotion Control

Reinforcement learning has been widely applied in legged locomotion control [4, 21, 31, 32, 33, 39, 42, 44, 65], benefiting from the policy update stability and high data efficiency provided by Proximal Policy Optimization (PPO) [48]. To adapt learned policies to diverse target tasks and ensure hardware deployability, previous works have designed two-stage training frameworks that aim to bridge the sim-to-real gap in the observation space [31, 32]. In contrast, this work introduces a novel two-stage training approach specifically aimed at improving sample efficiency, particularly addressing the challenge of early termination when walking on sparse terrains. This design not only enhances performance but also ensures more efficient learning in complex, real-world environments.

III. PROBLEM FORMULATION

This work aims to develop an terrain-aware humanoid locomotion policy, where controllers are trained via reinforcement learning (RL). The RL problem is formulated as a Markov Decision Process (MDP) $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, T, \mathcal{O}, r, \gamma \rangle$, where $\mathcal{S}$ and $\mathcal{A}$ denote the state and action spaces, respectively. The transition dynamics are represented by $T(s' | s, a)$, the reward function by $r(s, a)$, and the discount factor by $\gamma \in [0, 1]$. The primary objective is to optimize the policy $\pi(a_t | s_t)$ to maximize the discounted cumulative rewards:

$$\max_{\pi} J(\mathcal{M}, \pi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]. \quad (1)$$

In this work, the agent only has access to partial observations $\mathbf{o} \in \mathcal{O}$ due to sensory limitations and environmental noise, which provide incomplete information about the true state. Consequently, the agent functions within the framework of a Partially Observable Markov Decision Process (POMDP).

IV. METHODS

A. Foothold Reward

To accommodate the polygonal foot model of the humanoid robot, we introduce a sampling-based foothold reward that evaluates foot placement on sparse footholds. This evaluation is determined by the overlap between the foot's placement and designated safe areas, such as stones and beams. Specifically,

we sample n points on the soles of the robot's feet, as illustrated in Fig. 2. For each j -th sample on foot i , let d_{ij} denotes the global terrain height at the corresponding position. The penalty foothold reward r_{foothold} is defined as:

$$r_{\text{foothold}} = - \sum_{i=1}^2 \mathbb{C}_i \sum_{j=1}^n \mathbb{1}\{d_{ij} < \epsilon\}, \quad (2)$$

where $\mathbb{C}_i$ is an indicator function that specifies whether foot i is in contact with the terrain surface, and $\mathbb{1}$ is the indicator function for a condition. The threshold ϵ is a predefined depth tolerance, and when $d_{ij} < \epsilon$, it indicates that the terrain height at this sample point is significantly low, implying improper foot placement outside of a safe area. This reward function encourages the humanoid robot to maximize the overlap between its foot placement and the safe footholds, thereby improving its terrain-awareness capabilities.

B. Double Critic for Sparse Reward Learning

The task-specific foothold reward r_{foothold} is a sparse reward. To effectively optimize the policy, it is crucial to carefully balance this sparse reward with dense locomotion rewards which are crucial for gait regularization [62]. Inspired by [25, 56, 62], we adopt a double critic framework based on PPO, which effectively decouples the mixture of dense and sparse rewards.

In this framework, we train two separate critic networks, $\{V_{\phi_1}, V_{\phi_2}\}$, to independently estimate value functions for two distinct reward groups: (i) the regular locomotion reward group (dense rewards), $R_1 = \{r_i\}_{i=0}^n$, which have been studied in quadruped locomotion tasks [37] and humanoid locomotion tasks [34], and (ii) the task-specific foothold reward group (sparse reward), $R_2 = \{r_{\text{foothold}}\}$.

The double critic process is illustrated in Fig. 3. Specifically, each value network V_{ϕ_i} is updated independently for its corresponding reward group R_i with temporal difference loss (TD-loss):

$$\mathcal{L}(\phi_i) = \mathbb{E} \left[\|R_{i,t} + \gamma V_{\phi_i}(s_{t+1}) - V_{\phi_i}(s_t)\|^2 \right], \quad (3)$$

where γ is the discount factor. Then the respective advantages $\{\hat{A}_{i,t}\}$ are calculated using Generalized Advantage Estimation (GAE) [47]:

$$\delta_{i,t} = R_{i,t} + \gamma V_{\phi_i}(s_{t+1}) - V_{\phi_i}(s_t), \quad (4)$$

$$\hat{A}_{i,t} = \sum_{l=0}^{\infty} (\gamma \lambda)^l \delta_{i,t+l}, \quad (5)$$

where λ is the balancing parameter. These advantages are then individually normalized and synthesized into an overall advantage:

$$\hat{A}_t = w_1 \cdot \frac{\hat{A}_{1,t} - \mu_{\hat{A}_{1,t}}}{\sigma_{\hat{A}_{1,t}}} + w_2 \cdot \frac{\hat{A}_{2,t} - \mu_{\hat{A}_{2,t}}}{\sigma_{\hat{A}_{2,t}}}, \quad (6)$$

where w_i is the weight for each advantage component, and $\mu_{\hat{A}_{i,t}}$ and $\sigma_{\hat{A}_{i,t}}$ are the batch mean and standard deviation of




Fig. 3: Overview of BEAMDOJO. (a) **Training in Simulation:** In stage 1, proprioceptive and perceptive information, locomotion rewards and the foothold reward are decoupled respectively, with the former obtained from flat terrain and the latter from task terrain. The double critic module separately learns two reward groups. In stage 2, the policy is fine-tuned on the task terrain, utilizing the full set of observations and rewards. (b) **Real-world deployment:** The robot-centric elevation map, reconstructed using LiDAR data, is combined with proprioceptive information to serve as the input for the actor.

each component. This overall advantage is then used to update the policy:

$$\mathcal{L}(\theta) = \mathbb{E} \left[\min \left(\alpha_t(\theta) \hat{A}_t, \text{clip}(\alpha_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right], \quad (7)$$

where $\alpha_t(\theta)$ is the probability ratio, and ϵ is the clipping hyperparameter.

This double critic design provides a modular, plug-and-play solution for handling specialized tasks with sparse rewards, while effectively addressing the disparity in reward feedback frequencies within a mixed dense-sparse environment [62]. The detailed reward terms are provided in Appendix VI-A.

C. Learning Terrain-Aware Locomotion via Two-Stage RL

To address the early termination problem in complex terrain dynamics and encourage full trial-and-error exploration, we design a novel two-stage reinforcement learning (RL) approach for terrain-aware locomotion in simulation. As illustrated in Fig. 3, in the first stage, termed the “soft terrain dynamics constraints” phase, the humanoid robot is trained on flat terrain while being provided with a corresponding height map of the true task terrains (e.g., stepping stones). This setup encourages broad exploration without the risk of early termination from missteps. Missteps are penalized but do not lead to termination, allowing the humanoid robot to develop foundational skills for terrain-aware locomotion. In the second stage, termed the “hard terrain dynamics constraints” phase, we continue training the humanoid on the real terrains in simulation, where missteps result in termination. This stage fine-tunes the robot’s ability to step on challenging terrains accurately.

1) Stage 1: Soft Terrain Dynamics Constraints Learning: In this stage, we first map each task terrain (denoted as $\mathcal{T}$) to a flat terrain (denoted as $\mathcal{F}$) of the same size. Both terrains share the same terrain noise, with points are one-to-one corresponding. The only difference is that the flat terrain $\mathcal{F}$ fills the gaps in the real terrain $\mathcal{T}$.

We let the humanoid robot traverse the terrain $\mathcal{F}$, receiving proprioceptive observations, while providing perceptual feedback in the form of the elevation map of terrain $\mathcal{T}$ at the corresponding humanoid’s base position. This setup allows the robot to “imagine” walking on the true task terrain while actually traversing the safer flat terrain, where missteps do not lead to termination. To expose the robot to real terrain dynamics, we use the foothold reward (introduced in Section IV-A). In this phase, this reward is provided by the terrain $\mathcal{T}$, where $d_{i,j}$ is the height of the true terrain at the sampling point, while locomotion rewards are provided by the terrain $\mathcal{F}$.

This design successfully decouples the standard locomotion task and the task of traversing sparse footholds: flat terrain, $\mathcal{F}$, provides proprioceptive information and locomotion rewards to learn regular gaits, while risky task terrain, $\mathcal{T}$, offers perceptive information and the foothold reward to develop terrain-awareness skills. We train these two reward components separately using a double critic framework, as described in Section IV-B.

Furthermore, by allowing the humanoid robot to traverse the flat terrain while applying penalties for missteps instead of terminating the episode, the robot can continuously attempt foothold placements, making it much easier to obtain successful positive samples. In contrast, conventional early termination disrupts entire trajectories, making it extremely difficult to acquire safe foothold samples when learning from

scratch. This approach significantly improves sampling efficiency and alleviates the challenges of exploring terrains with sparse footholds.

2) *Stage 2: Hard Terrain Dynamics Constraints Learning*: In the second stage, we fine-tune the policy directly on the task terrain $\mathcal{T}$. Unlike in Stage 1, missteps on $\mathcal{T}$ now result in immediate termination. This enforces strict adherence to the true terrain constraints, requiring the robot to develop precise and safe locomotion strategies.

To maintain a smooth gait and accurate foot placements, we continue leveraging the double-critic framework to optimize both locomotion rewards and the foothold reward r_{foothold} . Here, d_{ij} again represents the height of terrain $\mathcal{T}$ at the given sampling point.

D. Training in Simulation

1) *Observation Space and Action Space*: The policy observations, denoted as $\mathbf{o}_t$, consist of four components:

$$\mathbf{o}_t = [\mathbf{c}_t, \mathbf{o}_t^{\text{proprio}}, \mathbf{o}_t^{\text{percept}}, \mathbf{a}_{t-1}]. \quad (8)$$

The commands $\mathbf{c}_t \in \mathbb{R}^3$ specify the desired velocity, represented as $[\mathbf{v}_x^c, \mathbf{v}_y^c, \omega_{\text{yaw}}^c]$. These denote the linear velocities in the longitudinal and lateral directions, and the angular velocity in the horizontal plane, respectively. The proprioceptive observations $\mathbf{o}_t^{\text{proprio}} \in \mathbb{R}^{64}$ include the base angular velocity $\omega_t \in \mathbb{R}^3$, gravity direction in the robot's frame $\mathbf{g}_t \in \mathbb{R}^3$, joint positions $\theta_t \in \mathbb{R}^{29}$, and joint velocities $\dot{\theta}_t \in \mathbb{R}^{29}$. The perceptive observations $\mathbf{o}_t^{\text{percept}} \in \mathbb{R}^{15 \times 15}$ correspond to an egocentric elevation map centered around the robot. This map samples 15×15 points within a 0.1 m grid in both the longitudinal and lateral directions. The action of last timestep $\mathbf{a}_{t-1} \in \mathbb{R}^{12}$ is also included to provide temporal context.

The action $\mathbf{a}_t \in \mathbb{R}^{12}$ represents the target joint positions for the 12 lower-body joints of the humanoid robot, which are directly output by the actor network. For the upper body joints, the default position is used for simplicity. A proportional-derivative (PD) controller converts these joint targets into torques to track the desired positions.

2) *Terrain and Curriculum Design*: Inspired by [23, 60, 63], we design five types of sparse foothold terrains for the two-stage training and evaluation:

- *Stones Everywhere*: This is a general sparse foothold terrain where stones are scattered across the entire terrain. The center of the terrain is a platform surrounded by stones, as shown in Fig. 4(a). The stones are uniformly distributed within sub-square grids. As the curriculum progresses, the stone size decreases and the sparsity increases.
- *Stepping Stones*: This terrain consists of two lines of stepping stones in the longitudinal direction, connected by two platforms at each end, as shown in Fig. 4(b). Each stone is uniformly distributed in two sub-square grids, with the same curriculum effect as in *Stones Everywhere*.
- *Balancing Beams*: In the initial curriculum level, this terrain has two lines of separate stones in the longitudinal




Fig. 4: Terrain Setting in Simulation. (a) is used for stage 1 training, while (b) and (c) are used for stage 2 training. The training terrain progression is listed from simple to difficult. (b)-(e) are used for evaluation.

direction. As the curriculum progresses, the size of the stones decreases and their lateral distance reduces, eventually forming a single line of balancing beams, as shown in Fig. 4(c). This terrain is challenging for the robot as it must learn to keep its feet together on the beams without colliding with each other, while maintaining the center of mass. This requires a distinct gait compared to regular locomotion tasks.

- *Stepping Beams*: This terrain consists of a sequence of beams to step on, randomly distributed along the longitudinal direction, with two platforms at either end, as illustrated in Fig. 4(d). This terrain, along with the *Stones Everywhere* and *Stepping Stones* terrains, requires the robot to place its footholds with high precision.
- *Gaps*: This terrain consists of several gaps with random distances between them, as shown in Fig. 4(e). This terrain requires the robot to make large steps to cross the gaps.

We begin by training the robot on the *Stones Everywhere* terrain in Stage 1 with soft terrain constraints to develop a generalizable policy. In Stage 2, the policy is fine-tuned on the *Stepping Stones* and *Balancing Beams* terrains with hard terrain constraints. The commands used in these two stages

TABLE I: Commands Sampled in Two Stage RL Training

Term	Value (stage 1)	Value (stage 2)
$\mathbf{v}_x^c$	$\mathcal{U}(-1.0, 1.0)$ m/s	$\mathcal{U}(-1.0, 1.0)$ m/s
$\mathbf{v}_y^c$	$\mathcal{U}(-1.0, 1.0)$ m/s	$\mathcal{U}(0.0, 0.0)$ m/s
ω_{yaw}^c	$\mathcal{U}(-1.0, 1.0)$ rad/s	$\mathcal{U}(0.0, 0.0)$ m/s

are detailed in Table I. Note that in Stage 2, only a single x-direction command is given, with no yaw command provided. This means that if the robot deviates from facing forward, no correction command is applied. We aim for the robot to learn to consistently face forward from preceptive observation, rather than relying on continuous yaw corrections.

For evaluation, the *Stepping Stones*, *Balancing Beams*, *Stepping Beams*, and *Gaps* terrains are employed. Remarkably, our method demonstrates strong zero-shot transfer capabilities on the latter two terrains, despite the robot being trained exclusively on the first three terrains.

The curriculum is designed as follows: the robot progresses to the next terrain level when it successfully traverses the current terrain level three times in a row. Furthermore, the robot will not be sent back to an easier terrain level before it pass all levels, as training on higher-difficulty terrains is challenging at first. The detailed settings of the terrain curriculum are presented in the Appendix VI-B.

3) *Sim-to-Real Transfer*: To enhance robustness and facilitate sim-to-real transfer, we employ extensive domain randomization [51, 54] on key dynamic parameters. Noise is injected into observations, humanoid physical properties, and terrain dynamics. Additionally, to address the large sim-to-real gap between the ground-truth elevation map in simulation and the LiDAR-generated map in reality—caused by factors such as odometry inaccuracies, noise, and jitter—we introduce four types of elevation map measurement noise during height sampling in the simulator:

- *Vertical Measurement*: Random vertical offsets are applied to the heights for an episode, along with uniformly sampled vertical noise added to each height sample at every time step, simulating the noisy vertical measurement of the LiDAR.
- *Map Rotation*: To simulate odometry inaccuracies, we rotate the map in roll, pitch, and yaw. For yaw rotation, we first sample a random yaw noise. The elevation map, initially aligned with the robot’s current orientation, is then resampled by adding the yaw noise, resulting in a new elevation map corresponding to the updated orientation. For roll and pitch rotations, we randomly sample the biases $[h_x, h_y]$ and perform linear interpolation from $-h_x$ to h_x along the x -direction and from $-h_y$ to h_y along the y -direction. The resulting vertical height map noise is then added to the original elevation map.
- *Foothold Extension*: Random foothold points adjacent to valid footholds are extended, turning them into valid footholds. This simulates the smoothing effect that occurs during processing of LiDAR elevation data.

- *Map Repeat*: To simulate delays in elevation map updates, we randomly repeat the map from the previous timestep.

The detailed domain randomization settings are provided in Appendix VI-C.

E. Real-world Deployment

1) *Hardware Setup*: We use Unitree G1 humanoid robot for our experiments in this work. The robot weighs 35 kg, stands 1.32 meters tall, and features 23 actuated degrees of freedom: 6 in each leg, 5 in each arm, and 1 in the waist. It is equipped with a Jetson Orin NX for onboard computation and a Livox Mid-360 LiDAR, which provides both IMU data and feature points for perception.

2) *Elevation Map and System Design*: The raw point cloud data obtained directly from the LiDAR suffers from significant occlusion and noise, making it challenging to use directly. To address this, we followed [34] to construct a robot-centric, complete, and robust elevation map. Specifically, we employed Fast LiDAR-Inertial Odometry (FAST-LIO) [57, 58] to fuse LiDAR feature points with IMU data provided by the LiDAR. This fusion generates precise odometry outputs, which are further processed using robot-centric elevation mapping methods [10, 11] to produce a grid-based representation of ground heights.

During deployment, the elevation map publishes information at a frequency of 10 Hz, while the learned policy operates at 50 Hz. The policy’s action outputs are subsequently sent to a PD controller, which runs at 500 Hz, ensuring smooth and precise actuation.

V. EXPERIMENTS

A. Experimental Setup


We compare our proposed framework BEAMDOJO, which integrates two-stage RL training and a double critic, with the following baselines:

- BL 1) **PIM** [34]: This one-stage method is designed for humanoid locomotion tasks, such as walking up stairs and traversing uneven terrains. We additionally add our foothold reward r_{foothold} to encourage the humanoid to step accurately on the foothold areas.
- BL 2) **Naive**: This method neither include the two-stage RL nor the double critic. The only addition is the foothold reward. This is an naive implementation to solve this task.
- BL 3) **Ours w/o Soft Dyn**: This is an ablation which removing the first stage of training with soft terrain dynamics constraints.
- BL 4) **Ours w/o Double Critic**: This is an ablation which uses a single critic to handle both locomotion rewards and foothold reward, instead of using a double critic. This follows the traditional design in most locomotion tasks.

The training and simulation environments are implemented in IsaacGym [36]. To ensure fairness, we adapt all methods to two stages. For stage 1, we train the humanoid on the *Stones*

TABLE II: Benchmarked Comparison in Simulation.

	Stepping Stones		Balancing Beams		Stepping Beams		Gaps	
	R_{succ} (%), $\uparrow$	R_{trav} (%), $\uparrow$	R_{succ} (%), $\uparrow$	R_{trav} (%), $\uparrow$	R_{succ} (%), $\uparrow$	R_{trav} (%), $\uparrow$	R_{succ} (%), $\uparrow$	R_{trav} (%), $\uparrow$
Medium Terrain Difficulty								
PIM	71.00 (± 1.53)	78.29 (± 2.49)	74.67 (± 2.08)	82.19 (± 4.96)	88.33 (± 3.61)	93.16 (± 4.78)	98.00 (± 0.57)	99.16 (± 0.75)
Naive	48.33 (± 6.11)	47.79 (± 5.76)	57.00 (± 7.81)	71.59 (± 8.14)	92.00 (± 2.52)	92.67 (± 3.62)	95.33 (± 1.53)	98.41 (± 0.67)
Ours w/o Soft Dyn	65.33 (± 2.08)	74.62 (± 1.37)	79.00 (± 2.64)	82.67 (± 2.92)	98.67 (± 2.31)	99.64 (± 0.62)	96.33 (± 1.53)	98.60 (± 1.15)
Ours w/o Double Critic	83.00 (± 2.00)	86.64 (± 1.96)	88.67 (± 2.65)	90.21 (± 1.95)	96.33 (± 1.15)	98.88 (± 1.21)	98.00 (± 1.00)	99.33 (± 0.38)
BEAMDOJO	95.67 (± 1.53)	96.11 (± 1.22)	98.00 (± 2.00)	99.91 (± 0.07)	98.33 (± 1.15)	99.28 (± 0.65)	98.00 (± 2.65)	99.21 (± 1.24)
Hard Terrain Difficulty								
PIM	46.67 (± 2.31)	52.88 (± 2.86)	33.00 (± 2.31)	45.28 (± 3.64)	82.67 (± 2.31)	90.68 (± 1.79)	96.00 (± 1.00)	98.27 (± 3.96)
Naive	00.33 (± 0.57)	21.17 (± 1.71)	00.67 (± 1.15)	36.25 (± 7.85)	82.00 (± 3.61)	88.91 (± 3.75)	31.00 (± 3.61)	62.70 (± 4.08)
Ours w/o Soft Dyn	42.00 (± 6.56)	47.09 (± 6.97)	51.00 (± 4.58)	72.93 (± 4.38)	87.33 (± 2.08)	89.41 (± 1.75)	93.00 (± 1.00)	95.62 (± 2.50)
Ours w/o Double Critic	55.67 (± 3.61)	60.95 (± 2.67)	70.33 (± 3.06)	85.64 (± 3.24)	94.67 (± 1.53)	96.57 (± 1.42)	94.33 (± 3.06)	95.62 (± 2.50)
BEAMDOJO	91.67 (± 1.33)	94.26 (± 2.08)	94.33 (± 1.53)	95.15 (± 1.82)	97.67 (± 2.08)	98.54 (± 1.43)	94.33 (± 1.15)	97.00 (± 1.30)


Fig. 5: Foothold Error Comparison. The foothold error benchmarks of all methods are evaluated in (a) stepping stones and (b) balancing beams, both tested under medium terrain difficulty.

Everywhere with curriculum learning. In this stage, our method and BL 4) use soft terrain dynamics constraints, while all other baselines use hard terrain dynamics constraints. For stage 2, we conduct fine-tuning on the *Stepping Stones* and *Balancing Beams* terrains with curriculum learning.


For evaluation, we test all methods on the *Stepping Stones*, *Balancing Beams*, *Stepping Beams* and *Gaps* terrains. We evaluate performance using three metrics:

- **Success Rate** R_{succ} : The percentage of successful attempts to cross the entire terrain.
- **Traverse Rate** R_{trav} : The ratio of the distance traveled before falling to the total terrain length (8 m).
- **Foothold Error** E_{foot} : The average proportion of foot samples landing outside the intended foothold areas.

B. Simulation Experiments

1) *Quantitative results*: We report the success rate (R_{succ}) and traverse rate (R_{trav}) for four terrains at medium and hard difficulty levels (terrain level 6 and level 8, respectively) in Table II. For each setting, the mean and standard deviation are calculated over three policies trained with different random seeds, each evaluated across 100 random episodes. Our key observations are as follows:

- Leveraging the efficient two-stage RL framework and the double critic, BEAMDOJO consistently outperforms


Fig. 6: Learning Efficiency. The learning curves show the maximum terrain levels achieved in two training stages of all methods. Faster attainment of terrain level 8 indicates more efficient learning.

single-stage approaches and ablation designs, achieving high success rates and low foothold errors across all challenging terrains. Notably, the naive implementation struggles significantly and is almost incapable of traversing stepping stones and balancing beams at hard difficulty levels.

- Existing humanoid controllers [34] face difficulties when adapting to risky terrains with fine-grained footholds, primarily due to sparse foothold rewards and low learning efficiency.
- Despite the our method not being explicitly trained on *Stepping Beams* and *Gaps*, it demonstrates impressive zero-shot generalization capabilities on these terrains.

2) *Detailed Ablation Analysis*: We conduct additional ablation studies by comparing BEAMDOJO with BL 2), BL 3), and BL 4).

Foot Placement Accuracy: As shown in Fig. 5, BEAMDOJO achieves highly accurate foot placement with low foothold error values, largely due to the contribution of the double critic. In comparison, the naive implementation shows higher error rates, with a substantial proportion of foot placements landing outside the safe foothold areas. This demonstrates the precision and effectiveness of our method in

TABLE III: Gait Regularization. We conduct experiments on stepping stones and evaluate three representative gait regularization reward metrics: smoothness, feet air time, and feet clearance. Detailed definitions of the reward functions can be found in Table VII.

Designs	Smoothness ($\downarrow$)	Feet Air Time ($\uparrow$)
Naive	1.7591 (± 0.1316)	-0.0319 (± 0.0028)
Ours w/o Soft Dyn	0.9633 (± 0.0526)	-0.0169 (± 0.0014)
Ours w/o Double Critic	1.2705 (± 0.1168)	-0.0229 (± 0.0033)
BEAMDOJO	0.7603 (± 0.0315)	-0.0182 (± 0.0027)




Fig. 7: Foot Placement Planning Visualization. We illustrate two trajectories for the foot placement process: the yellow line represents BEAMDOJO, while the red line corresponds to *Ours w/o Double Critic*. Points along the trajectories are marked at equal time intervals. From A to C, the method without the double critic exhibits significant adjustments only when approaching the target stone (at point B).

challenging terrains.

Learning Efficiency: Although we train for 10,000 iterations in both stages to ensure convergence across all designs, BEAMDOJO converges significantly faster, as shown in Fig. 6. Both the two-stage training setup and the double critic improve learning efficiency, with the two-stage setup contributing the most. In contrast, the naive implementation struggles to reach higher terrain levels in both stages.

The advantage of two-stage learning lies in its ability to allow the agent to continuously attempt foot placements, even in the presence of missteps, making it easier to accumulate a substantial number of successful foot placement samples. Meanwhile, the double-critic setup separates the foothold reward from the locomotion rewards, ensuring that its updates remain unaffected by the noise of unstable locomotion signals, particularly in the early training phase. Both strategies play a crucial role in enhancing learning efficiency.

Gait Regularization: The combination of small-scale gait regularization rewards with sparse foothold reward can hinder gait performance, as shown in Table III, where the naive design and the ablation without the double critic exhibit poor performance in both smoothness and feet air time. In contrast, our method and the ablation with double critic demonstrates superior motion smoothness and improved feet clearance. This improvement arises because, in the double-critic framework, the advantage estimates for the dense and sparse reward groups are normalized independently, preventing the sparse rewards from introducing noise that could disrupt the learning of regularization rewards.

Foot Placement Planning: As illustrated in Fig. 7, we

TABLE IV: Agility Test. We evaluate the agility of the humanoid robot on stepping stones with a total length of 2.8m.

v_x^c (m/s)	Time Cost (s)	Average Speed (m/s)	Error Rate (% , $\downarrow$)
0.5	6.33 (± 0.15)	0.45 (± 0.05)	10.67 (± 4.54)
0.75	4.33 (± 0.29)	0.65 (± 0.05)	13.53 (± 6.52)
1.0	3.17 (± 0.58)	0.88 (± 0.04)	11.83 (± 8.08)
1.25	2.91 (± 0.63)	0.96 (± 0.03)	22.74 (± 5.32)
1.5	2.69 (± 0.42)	1.04 (± 0.05)	30.68 (± 6.17)

observe that the double critic also benefits foot placement planning of the entire sub-process of foot lifting and landing. Our method, BEAMDOJO, enables smoother planning, allowing the foot to precisely reach the next foothold. In comparison, the baseline excluding double critic demonstrates reactive stepping, where adjustments are largely made when the foot is close to the target stone. This behavior indicates that the double critic, by learning the sparse foothold reward separately, helps the policy adjust its motion over a longer horizon.

C. Real-world Experiments

1) *Result:* As demonstrated in Fig.1, our framework achieves zero-shot transfer, successfully generalizing to real-world dynamics. To showcase the effect of height map domain randomization (introduced in Section IV-D3) in sim-to-real transfer, we compare our proposed method with an ablation that excludes height map randomization (denoted as “ours w/o HR”). We conduct five trials on each terrain and report the success and traversal rates in Fig. 8, with following conclusions:

- BEAMDOJO achieves a high success rate in real-world deployments, demonstrating excellent precise foot placement capabilities. Similar to simulation results, it also exhibits impressive generalization performance on *Stepping Beams* and *Gaps*, even though these terrains were not part of the training set.
- The ablation, lacking height map domain randomization, results in a significantly lower success rate, highlighting the importance of this design.
- It is also worth mentioning that BEAMDOJO enables backward movement in risky terrains, as shown in Fig. 1(b). This advantage is achieved by leveraging LiDAR to its full potential, whereas a single depth camera cannot handle such scenarios.

2) *Agility Test:* To assess the agility of our method, we provide the humanoid robot with five commanded longitudinal velocities, v_x^c : 0.5, 0.75, 1.0, 1.25, and 1.5 m/s, and check the tracking ability. Each test was conducted over three trials, and the results are reported in Table IV. The results show minimal tracking error up to the highest training command velocity of 1.0 m/s, where the robot achieves an average speed of 0.88 m/s, demonstrating the agility of our policy. However, performance degrades significantly beyond 1.25 m/s, as maintaining such speeds becomes increasingly difficult on these highly challenging terrains.

Stepping Stones	R_{succ}	R_{trav}	Balancing Beams	R_{succ}	R_{trav}	Stepping Beams	R_{succ}	R_{trav}	Gaps	R_{succ}	R_{trav}
Ours w/o HR	1/5	38.20%	Ours w/o HR	0/5	12.37%	Ours w/o HR	1/5	30.00%	Ours w/o HR	3/5	60.00%
BEAMDOJO	4/5	92.18%	BEAMDOJO	4/5	88.16%	BEAMDOJO	3/5	70.00%	BEAMDOJO	5/5	100.00%

Fig. 8: Real-world Experiments. We build terrains in the real world similar to those in simulation. (a) **Stepping Stones:** stones with a size of 20 cm, a maximum distance of 45 cm between stones, and a sparsity level of 72.5%. (b) **Balancing Beams:** beams with a width of 20 cm. (c) **Stepping Beams:** beams with a size of 20 cm, a maximum distance of 45 cm between beams, and a sparsity level of 66.6%. (d) **Gaps:** gaps with a maximum distance of 50 cm.




Fig. 9: Robustness Test. We evaluate the robustness of the humanoid robot in real-world scenarios with: (a) heavy payload, (b) external forces, and (c) recovering from missteps.

3) *Robustness Test:* To evaluate the robustness of our precise foothold controller, we conducted the following experiments on real-world experiment terrains:

- **Heavy Payload:** As shown in Fig. 9(a), the robot carried a 10 kg payload—approximately 1.5 times the weight of its torso—causing a significant shift in its center of mass. Despite this challenge, the robot effectively maintained agile locomotion and precise foot placements, demonstrating its robustness under increased payload conditions.
- **External Force:** As shown in Fig. 9(b), the robot was subjected to external forces from various directions. Starting from a stationary pose, the robot experienced external pushes, transitioned to single-leg support, and finally recovered to a stable standing position with two-leg support.
- **Misstep Recovery:** As shown in Fig. 9(c), the robot traverse terrain without prior scanning of terrain dynamics. Due to occlusions, the robot lacked information about the terrain underfoot, causing initial missteps. Nevertheless, it demonstrated robust recovery capabilities.

D. Extensive Studies and Analysis

1) *Design of Foothold Reward:* As discussed in Section IV-A, our sampling-based foothold reward is proportional to the number of safe points, making it a relatively continuous reward: the larger the overlap between the foot placement and the safe footholds, the higher the reward the agent receives.

TABLE V: Comparison of Different Foothold Reward Designs. The success rate and foothold error for each foothold reward design are evaluated on stepping stones with medium terrain difficulty.

Designs	R_{succ} (%), $\uparrow$	E_{foot} (%), $\downarrow$
foothold-30%	93.67 (± 1.96)	11.43 (± 0.81)
foothold-50%	92.71 (± 1.06)	10.78 (± 1.94)
foothold-70%	91.94 (± 2.08)	14.35 (± 2.61)
BEAMDOJO	95.67 (± 1.53)	7.79 (± 1.33)

We compare this approach with other binary and coarse reward designs: when $p\%$ of the sampled points fall outside the safe area, a full penalty is applied; otherwise, no penalty is given. This can be defined as:

$$r_{\text{foothold}-p\%} = - \sum_{i=1}^2 \mathbb{C}_i \cdot \mathbb{1} \left\{ \left(\sum_{j=1}^n \mathbb{1}\{d_{ij} < \epsilon\} \right) \geq p\% \cdot n \right\}. \quad (9)$$

We compare our continuous foothold reward design with three variants of the coarse-grained approach, where $p = 30, 50$, and 70 (denoted as foothold-30%, foothold-50%, and foothold-70% respectively). The success rate R_{succ} and the foothold error E_{foot} on stepping stones are reported in Table V.

It is clear that our fine-grained design enables the robot to make more accurate foot placements compared to the other designs, as this continuous approach gradually encourages maximizing the overlap. Among the coarse-grained

TABLE VI: Comparison of Different Curriculum Designs. The success rate and traverse rate for each curriculum design are evaluated on stepping stones with medium and hard terrain difficulty respectively.

Designs	Medium Difficulty		Hard Difficulty	
	R_{succ}	R_{trav}	R_{succ}	R_{trav}
w/o curriculum-medium	88.33	90.76	2.00	18.36
w/o curriculum-hard	40.00	52.49	23.67	39.94
BEAMDOJO	95.67	96.11	82.33	86.87

approaches, foothold-50% performs better than foothold-30% and foothold-70%, as a 30% threshold is too strict to learn effectively, while 70% is overly loose.

2) *Design of Curriculum:* To validate the effectiveness of the terrain curriculum introduced in Section IV-D2, we introduce an ablation study without curriculum learning. In this design, we train using only medium and hard terrain difficulties in both stages (denoted as “w/o curriculum-medium” and “w/o curriculum-hard”). Similarly, we report the R_{succ} and R_{trav} for both ablation methods, along with our method, on stepping stones terrain at two different difficulty levels in Table VI. The results show that incorporating curriculum learning significantly improves both performance and generalization across terrains of varying difficulty. In contrast, without curriculum learning, the model struggles significantly with challenging terrain when learning from scratch (“ours w/o curriculum-hard”), and also faces difficulties on other terrain types, severely limiting its generalization ability (“ours w/o curriculum-medium”).

3) *Design of Commands:* As mentioned in Section IV-D2, in the second stage, no heading command is applied, and the robot is required to learn to consistently face forward through terrain dynamics. We compare this approach with one that includes a heading command (denoted as “ours w/ heading command”), where deviation from the forward direction results in a corrective yaw command based on the current directional error. In the deployment, we use the LiDAR odometry module to update the heading command in real time, based on the difference between the current orientation and the initial forward direction.

We conduct five trials on the stepping stones terrain in the real world, comparing our proposed method with the “ours w/ heading command” design. The success rates are 4/5 and 1/5, respectively. The poor performance of the heading command design is primarily due to two factors: (1). In the simulation, the model overfits the angular velocity of the heading command, making it difficult to handle noisy real-world odometry data; (2). In the real world, precise manual calibration of the initial position is required to determine the correct forward direction, making the heading command approach less robust. In contrast, BEAMDOJO, without heading correction, proves to be more effective.

4) *Generalization to Non-Flat Terrains:* We observe that BEAMDOJO also generalizes well to non-flat terrains such as stairs and slopes. The main adaptation involves calculating the




Fig. 10: Generalization Test on Non-Flat Terrains. We conduct real-world experiments on (a) stairs with a width of 25 cm and a height of 15 cm, and (b) slopes with a 15-degree incline.




Fig. 11: Failure Case Analysis. We evaluate the success rate on varying (a) stone sizes, and (b) step distances.

base height reward relative to the foot height rather than the ground height on these uneven surfaces. Additionally, Stage, 1 pre-training becomes unnecessary for stairs and slopes, as the footholds are no longer sparse. We validate our approach through hardware experiments on stairs and slopes, as shown in Fig. 10. The results demonstrate that BEAMDOJO enables the robot to successfully traverse stairs and slopes with success rates of 8/10 and 10/10, respectively.

5) *Failure Cases:* To investigate the framework’s performance limitations, we evaluate its performance across varying stone sizes and step distances, as shown in Fig. 11. We compare policies trained with different terrain parameters, with the smallest stone size (20 cm, 15 cm and 10 cm) or largest step distance (45 cm, 50 cm and 55 cm). The results indicate that while tougher training enhances adaptability, performance still drops sharply on 10 cm stones (approximately half the foot length) and 55 cm steps (roughly equal to the leg length), even under the most challenging training settings. In these cases, the difficulty shifts toward maintaining balance on very small footholds and executing larger strides—challenges that the current reward function does not adequately address.

6) *Limitations:* On the one hand, the performance of our method is significantly constrained by the limitations of the perception module. Inaccuracies in LiDAR odometry, along with issues such as jitter and map drift, present considerable challenges for real-world deployment. Furthermore, when processing LiDAR data, the trade-off between the confidence in noisy measurements and the dynamic changes in terrain—such as the jitter of stones, which is difficult to simulate in the simulation—makes it challenging to effectively handle sudden disturbances or variations. As a result, the system struggles to quickly and flexibly adapt to unexpected changes in the

environment.

On the other hand, our method has yet to fully leverage the information provided by the elevation map, and has not adequately addressed the challenges of terrains with significant foothold height variations. In future work, we aim to develop a more generalized controller that enables agile locomotion, extending to a broader range of terrains, including stairs and other complex surfaces that require footstep planning, as well as terrains with significant elevation changes.

VI. CONCLUSION

In this paper, we proposed a novel framework, BEAMDOJO, which enables humanoid robots to traverse with agility and robustness on sparse foothold terrains such as stepping stones and balance beams, and generalize to a wider range of challenging terrains (e.g., gaps, balancing beams). The key conclusions are summarized as follows:

- **Accuracy of Foot Placement:** We introduced a foothold reward for polygonal feet, which is proportional to the contact area between the footstep and the safe foothold region. This continuous reward effectively encourages precise foot placements.
- **Training Efficiency and Effectiveness:** By incorporating a two-stage reinforcement learning training process, BEAMDOJO enables full trial-and-error exploration. Additionally, the double-head critic significantly enhances the learning of sparse foothold rewards, regularizes gait patterns, and facilitates long-distance foot placement planning.
- **Agility and Robustness in the Real World:** Our experiments demonstrate that BEAMDOJO empowers humanoid robots to exhibit agility and achieve a high success rate in real-world scenarios. The robots maintain stable walking even under substantial external disturbances and the inevitable sway of beams in real world. Notably, by leveraging LiDAR-based mapping, we have achieved stable backward walking, a challenge typically encountered with depth cameras.

ACKNOWLEDGMENTS

This work is funded in part by the National Key R&D Program of China (2022ZD0160201), and Shanghai Artificial Intelligence Laboratory. We thank the authors of PIM [34] for their kind help with the deployment of the elevation map. We thank Unitree Robotics for their help with the Unitree G1 humanoid robot. We thank RSS reviewers for their careful and insightful feedback, which helped improve the quality of this work.

REFERENCES

- [1] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning (CoRL)*, pages 403–415. PMLR, 2023.
- [2] Zixuan Chen, Xialin He, Yen-Jen Wang, Qiayuan Liao, Yanjie Ze, Zhongyu Li, S Shankar Sastry, Jiajun Wu, Koushil Sreenath, Saurabh Gupta, et al. Learning smooth humanoid locomotion through lipschitz-constrained policies. *arXiv preprint arXiv:2410.11825*, 2024.
- [3] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024.
- [4] Xuxin Cheng, Kexin Shi, Ananye Agarwal, and Deepak Pathak. Extreme parkour with legged robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 11443–11450. IEEE, 2024.
- [5] Joel Chestnutt, Manfred Lau, German Cheung, James Kuffner, Jessica Hodgins, and Takeo Kanade. Footstep planning for the honda asimo humanoid. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pages 629–634. IEEE, 2005.
- [6] Wenhao Cui, Shengtao Li, Huaxing Huang, Bangyu Qin, Tianchu Zhang, Liang Zheng, Ziyang Tang, Chenxu Hu, NING Yan, Jiahao Chen, et al. Adapting humanoid locomotion over challenging terrain via two-phase training. In *Conference on Robot Learning (CoRL)*, 2024.
- [7] Hongkai Dai and Russ Tedrake. Planning robust walking motion on uneven terrain via convex optimization. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 579–586. IEEE, 2016.
- [8] Robin Deits and Russ Tedrake. Footstep planning on uneven terrain with mixed-integer convex optimization. In *2014 IEEE-RAS international conference on humanoid robots*, pages 279–286. IEEE, 2014.
- [9] Johannes Engelsberger, Christian Ott, Máximo A Roa, Alin Albu-Schäffer, and Gerhard Hirzinger. Bipedal walking control based on capture point dynamics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4420–4427. IEEE, 2011.
- [10] Péter Fankhauser, Michael Bloesch, Christian Gehring, Marco Hutter, and Roland Siegwart. Robot-centric elevation mapping with uncertainty estimates. In *International Conference on Climbing and Walking Robots (CLAWAR)*, 2014.
- [11] Péter Fankhauser, Michael Bloesch, and Marco Hutter. Probabilistic terrain mapping for mobile robots with uncertain localization. *IEEE Robotics and Automation Letters (RA-L)*, 3(4):3019–3026, 2018. doi: 10.1109/LRA.2018.2849506.
- [12] Zipeng Fu, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Minimizing energy consumption leads to the emergence of gaits in legged robots. In *Conference on Robot Learning (CoRL)*, 2021.
- [13] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *Conference on Robot Learning (CoRL)*, 2024.
- [14] Kunihiro Fukushima. Visual feature extraction by a multilayered network of analog threshold elements. *IEEE Transactions on Systems Science and Cybernetics*, 5(4):

- 322–333, 1969.
- [15] Siddhant Gangapurwala, Mathieu Geisert, Romeo Orsolino, Maurice Fallon, and Ioannis Havoutis. Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control. *IEEE Transactions on Robotics*, 38(5):2908–2927, 2022.
 - [16] Ruben Grandia, Fabian Jenelten, Shaohui Yang, Farbod Farshidian, and Marco Hutter. Perceptive locomotion through nonlinear model-predictive control. *IEEE Transactions on Robotics*, 39(5):3402–3421, 2023.
 - [17] Robert J Griffin, Georg Wiedebach, Stephen McCrory, Sylvain Bertrand, Inho Lee, and Jerry Pratt. Footstep planning for autonomous walking over rough terrain. In *2019 IEEE-RAS 19th international conference on humanoid robots (humanoids)*, pages 9–16. IEEE, 2019.
 - [18] 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. In *Robotics: Science and Systems (RSS)*, 2024.
 - [19] 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 (CoRL)*, 2024.
 - [20] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8944–8951. IEEE, 2024.
 - [21] Tairan He, Chong Zhang, Wenli Xiao, Guanqi He, Changliu Liu, and Guanya Shi. Agile but safe: Learning collision-free high-speed legged locomotion. In *Robotics: Science and Systems (RSS)*, 2024.
 - [22] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Jan Kautz, Changliu Liu, Guanya Shi, Xiaolong Wang, et al. Hover: Versatile neural whole-body controller for humanoid robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
 - [23] 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.
 - [24] Armin Hornung and Maren Bennewitz. Adaptive level-of-detail planning for efficient humanoid navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 997–1002. IEEE, 2012.
 - [25] Changxin Huang, Guangrun Wang, Zhibo Zhou, Ronghui Zhang, and Liang Lin. Reward-adaptive reinforcement learning: Dynamic policy gradient optimization for bipedal locomotion. *IEEE transactions on pattern analysis and machine intelligence*, 45(6):7686–7695, 2022.
 - [26] Fabian Jenelten, Takahiro Miki, Aravind E Vijayan, Marko Bjelonic, and Marco Hutter. Perceptive locomotion in rough terrain—online foothold optimization. *IEEE Robotics and Automation Letters (RA-L)*, 5(4): 5370–5376, 2020.
 - [27] Fabian Jenelten, Junzhe He, Farbod Farshidian, and Marco Hutter. Dtc: Deep tracking control. *Science Robotics*, 9(86):eadh5401, 2024.
 - [28] Mazeyu Ji, Xuanbin Peng, Fangchen Liu, Jialong Li, Ge Yang, Xuxin Cheng, and Xiaolong Wang. Ex-body2: Advanced expressive humanoid whole-body control. *arXiv preprint arXiv:2412.13196*, 2024.
 - [29] Zhenyu Jiang, Yuqi Xie, Jinhan Li, Ye Yuan, Yifeng Zhu, and Yuke Zhu. Harmon: Whole-body motion generation of humanoid robots from language descriptions. In *Conference on Robot Learning (CoRL)*, 2024.
 - [30] Shuuji Kajita, Fumio Kanehiro, Kenji Kaneko, Kazuhito Yokoi, and Hirohisa Hirukawa. The 3d linear inverted pendulum mode: A simple modeling for a biped walking pattern generation. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, volume 1, pages 239–246. IEEE, 2001.
 - [31] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. In *Robotics: Science and Systems (RSS)*, 2021.
 - [32] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.
 - [33] Junfeng Long, Zirui Wang, Quanyi Li, Liu Cao, Jiawei Gao, and Jiangmiao Pang. Hybrid internal model: Learning agile legged locomotion with simulated robot response. In *International Conference on Learning Representations (ICLR)*, 2024.
 - [34] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
 - [35] Yidan Lu, Yinzhao Dong, Ji Ma, Jiahui Zhang, and Peng Lu. Learning an adaptive fall recovery controller for quadrupeds on complex terrains. *arXiv preprint arXiv:2412.16924*, 2024.
 - [36] 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. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
 - [37] Gabriel B Margolis and Pulkit Agrawal. Walk these ways: Tuning robot control for generalization with multiplicity of behavior. In *Conference on Robot Learning (CoRL)*, pages 22–31. PMLR, 2023.
 - [38] Gabriel B Margolis, Tao Chen, Kartik Paigwar, Xiang Fu, Donghyun Kim, Sangbae Kim, and Pulkit Agrawal.

- Learning to jump from pixels. In *Conference on Robot Learning (CoRL)*, 2021.
- [39] 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.
- [40] Carlos Mastalli, Ioannis Havoutis, Michele Focchi, Darwin G Caldwell, and Claudio Semini. Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping, and whole-body control. *IEEE Transactions on Robotics*, 36(6):1635–1648, 2020.
- [41] Oliwier Melon, Romeo Orsolino, David Surovik, Mathieu Geisert, Ioannis Havoutis, and Maurice Fallon. Receding-horizon perceptive trajectory optimization for dynamic legged locomotion with learned initialization. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 9805–9811. IEEE, 2021.
- [42] 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):eabk2822, 2022.
- [43] Bhavyansh Mishra, Duncan Calvert, Sylvain Bertrand, Hakki Erhan Sevil, and Robert Griffin. Efficient terrain map using planar regions for footstep planning on humanoid robots. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8044–8050. IEEE, 2024.
- [44] I Made Aswin Nahrendra, Byeongho Yu, and Hyun Myung. Dreamwaq: Learning robust quadrupedal locomotion with implicit terrain imagination via deep reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5078–5084. IEEE, 2023.
- [45] Ilija Radosavovic, Sarthak Kamat, Trevor Darrell, and Jitendra Malik. Learning humanoid locomotion over challenging terrain. *arXiv preprint arXiv:2410.03654*, 2024.
- [46] Ilija Radosavovic, Bike Zhang, Baifeng Shi, Jathushan Rajasegaran, Sarthak Kamat, Trevor Darrell, Koushil Sreenath, and Jitendra Malik. Humanoid locomotion as next token prediction. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [47] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations (ICLR)*, 2015.
- [48] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [49] Alexander Stumpf, Stefan Kohlbrecher, David C Conner, and Oskar von Stryk. Supervised footstep planning for humanoid robots in rough terrain tasks using a black box walking controller. In *2014 IEEE-RAS International Conference on Humanoid Robots*, pages 287–294. IEEE, 2014.
- [50] Richard Stuart Sutton. *Temporal credit assignment in reinforcement learning*. University of Massachusetts Amherst, 1984.
- [51] 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)*, pages 23–30. IEEE, 2017.
- [52] Bart van Marum, Aayam Shrestha, Helei Duan, Pranay Dugar, Jeremy Dao, and Alan Fern. Revisiting reward design and evaluation for robust humanoid standing and walking. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [53] Alexander W Winkler, C Dario Bellicoso, Marco Hutter, and Jonas Buchli. Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robotics and Automation Letters (RA-L)*, 3(3):1560–1567, 2018.
- [54] Zhaoming Xie, Xingye Da, Michiel Van de Panne, Buck Babich, and Animesh Garg. Dynamics randomization revisited: A case study for quadrupedal locomotion. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4955–4961. IEEE, 2021.
- [55] Zhaoming Xie, Xingye Da, Buck Babich, Animesh Garg, and Michiel van de Panne. Glide: Generalizable quadrupedal locomotion in diverse environments with a centroidal model. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 523–539. Springer, 2022.
- [56] Pei Xu, Xiumin Shang, Victor Zordan, and Ioannis Karamouzas. Composite motion learning with task control. *ACM Transactions on Graphics (TOG)*, 42(4): 1–16, 2023.
- [57] Wei Xu and Fu Zhang. Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters (RA-L)*, 6(2):3317–3324, 2021.
- [58] Wei Xu, Yixi Cai, Dongjiao He, Jiarong Lin, and Fu Zhang. Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics*, 38(4):2053–2073, 2022.
- [59] Ruihan Yang, Ge Yang, and Xiaolong Wang. Neural volumetric memory for visual locomotion control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1430–1440, 2023.
- [60] Ruiqi Yu, Qianshi Wang, Yizhen Wang, Zhicheng Wang, Jun Wu, and Qiuguo Zhu. Walking with terrain reconstruction: Learning to traverse risky sparse footholds. *arXiv preprint arXiv:2409.15692*, 2024.
- [61] Wenhao Yu, Deepali Jain, Alejandro Escontrela, Atil Iscen, Peng Xu, Erwin Coumans, Sehoon Ha, Jie Tan, and Tingnan Zhang. Visual-locomotion: Learning to walk on complex terrains with vision. In *Conference on Robot Learning (CoRL)*, 2021.
- [62] Fatemeh Zargarbashi, Jin Cheng, Dongho Kang, Robert

Sumner, and Stelian Coros. Robotkeyframing: Learning locomotion with high-level objectives via mixture of dense and sparse rewards. In *Conference on Robot Learning (CoRL)*, 2024.

- [63] Chong Zhang, N. Rudin, David Hoeller, and Marco Hutter. Learning agile locomotion on risky terrains. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [64] Shaoting Zhu, Runhan Huang, Linzhan Mou, and Hang Zhao. Robust robot walker: Learning agile locomotion over tiny traps. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [65] Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher Atkeson, Sören Schwertfeger, Chelsea Finn, and Hang Zhao. Robot parkour learning. In *Conference on Robot Learning (CoRL)*, 2023.
- [66] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. In *Conference on Robot Learning (CoRL)*, 2024.

APPENDIX

A. Reward Functions

The reward functions we used during the training are shown in Table VII, which mainly comes from [1, 12, 31, 37, 34]. The corresponding symbols and their descriptions are provided in Table VIII.

TABLE VII: Reward Functions

Term	Equation	Weight
Group 1: Locomotion Reward Group		
xy velocity tracking	$\exp\{-\ \mathbf{v}_{xy} - \mathbf{v}_{xy}^c\ _2^2 / \sigma\}$	1.0
yaw velocity tracking	$\exp\{-(\omega_{\text{yaw}} - \omega_{\text{yaw}}^c)^2 / \sigma\}$	1.0
base height	$(h - h^{\text{target}})^2$	-10.0
orientation	$\ \mathbf{g}_x\ _2^2 + \ \mathbf{g}_y\ _2^2$	-2.0
z velocity	$\mathbf{v}_z^2$	-2.0
roll-pitch velocity	$\ \omega_{xy}\ _2^2$	-0.05
action rate	$\ \mathbf{a}_t - \mathbf{a}_{t-1}\ _2^2$	-0.01
smoothness	$\ \mathbf{a}_t - 2\mathbf{a}_{t-1} + \mathbf{a}_{t-2}\ _2^2$	-1e-3
stand still	$\ \theta\ _2^2 \cdot \mathbb{1}\{\ \mathbf{v}_{xy}^c\ _2^2 < \epsilon\}$	-0.05
joint velocities	$\ \dot{\theta}\ _2^2$	-1e-4
joint accelerations	$\ \ddot{\theta}\ _2^2$	-2.5e-8
joint position limits	$\text{ReLU}(\theta - \theta_{\max}) + \text{ReLU}(\theta_{\min} - \theta)$	-5.0
joint velocity limits	$\text{ReLU}(\dot{\theta} - \dot{\theta}_{\max})$	-1e-3
joint power	$ \tau \ \dot{\theta}\ ^T / (\ \mathbf{v}\ _2^2 + 0.2 * \ \omega\ _2^2)$	-2e-5
feet ground parallel	$\sum_{i=1}^2 \text{Var}(\mathbf{p}_{z,i})$	-0.02
feet distance	$\text{ReLU}(\mathbf{p}_{y,1} - \mathbf{p}_{y,2} - d_{\min})$	0.5
feet air time	$\sum_{i=1}^2 \left(t_{\text{air},i} - t_{\text{air}}^{\text{target}}\right) \cdot \mathbb{F}_i$	1.0
feet clearance	$\sum_{i=1}^2 \left(\mathbf{p}_{z,i} - \mathbf{p}_z^{\text{target}}\right)^2 \cdot \dot{\mathbf{p}}_{xy,i}$	-1.0
Group 2: Foothold Reward Group		
foothold	$-\sum_{i=1}^2 \mathbb{C}_i \sum_{j=1}^n \mathbb{1}\{d_{ij} < \epsilon\}$	1.0

TABLE VIII: Used Symbols

Symbols	Description
σ	Tracking shape scale, set to 0.25.
ϵ	Threshold for determining zero-command in stand still reward, set to 0.1.
τ	Computed joint torques.
h^{target}	Desired base height relative to the ground, set to 0.725.
$\text{ReLU}(\cdot)$	Function that clips negative values to zero [14].
$\mathbf{p}_i, \dot{\mathbf{p}}_i$	Spatial position and velocity of all sampled points on the i -th foot respectively.
$\mathbf{p}_z^{\text{target}}$	Target foot-lift height, set to 0.1.
$t_{\text{air},i}$	Air time of the i -th foot.
$t_{\text{air}}^{\text{target}}$	Desired feet air time, set to 0.5.
$\mathbb{F}_i$	Indicator specifying whether foot i makes first ground contact.
$d_{\min}$	Minimum allowable distance between two feet, set to 0.18.

B. Terrain Curriculum

The training terrains using curriculum comprises *Stones Everywhere*, *Stepping Stones*, and *Balancing Beams*. The *Stones Everywhere* terrain spans an area of $8m \times 8m$, while both *Stepping Stones* and *Balancing Beams* are $2m$ in width and $8m$ in length, with single-direction commands. The depth of gaps relative to the ground is set to $1.0m$, and all stones and beams exhibit height variations within $\pm 0.05m$. The depth tolerance threshold, ϵ , is set to $-0.1m$.

We define terrain difficulty levels ranging from 0 to 8, denoted as l . The specific terrain curriculum at each difficulty level are as follows:

- *Stones Everywhere*: The stone size is $\max\{0.25, 1.5(1 - 0.1l)\}$, and the stone distance is $0.05[l/2]$.
- *Stepping Stones*: The stone sizes follow the sequence $[0.8, 0.65, 0.5, 0.4, 0.35, 0.3, 0.25, 0.2, 0.2]$, with a maximum stone distance of $0.1 + 0.05l$.
- *Balancing Beams*: The stone size is $0.3 - 0.05[l/3]$, with the stone distance in x -direction $0.4 - 0.05l$, and in y -direction $[0.2, 0.2, 0.2, 0.25, 0.3, 0.35, 0.35, 0.4, 0.2]$. At the highest difficulty level, the terrain forms a single continuous balancing beam.

C. Domain Randomization

TABLE IX: Domain Randomization Setting

Term	Value
Observations	
angular velocity noise	$\mathcal{U}(-0.5, 0.5)$ rad/s
joint position noise	$\mathcal{U}(-0.05, 0.05)$ rad/s
joint velocity noise	$\mathcal{U}(-2.0, 2.0)$ rad/s
projected gravity noise	$\mathcal{U}(-0.05, 0.05)$ rad/s
Humanoid Physical Properties	
actuator offset	$\mathcal{U}(-0.05, 0.05)$ rad
motor strength noise	$\mathcal{U}(0.9, 1.1)$
payload mass	$\mathcal{U}(-2.0, 2.0)$ kg
center of mass displacement	$\mathcal{U}(-0.05, 0.05)$ m
Kp, Kd noise factor	$\mathcal{U}(0.85, 1.15)$
Terrain Dynamics	
friction factor	$\mathcal{U}(0.4, 1.0)$
restitution factor	$\mathcal{U}(0.0, 1.0)$
terrain height noise	$\mathcal{U}(-0.02, 0.02)$ m
Elevation Map	
vertical offset	$\mathcal{U}(-0.03, 0.03)$ m
vertical noise	$\mathcal{U}(-0.03, 0.03)$ m
map roll, pitch rotation noise	$\mathcal{U}(-0.03, 0.03)$ m
map yaw rotation noise	$\mathcal{U}(-0.2, 0.2)$ rad
foothold extension probability	0.6
map repeat probability	0.2

D. Hyperparameters

TABLE X: Hyperparameters

Hyperparameter	Value
General	
num of robots	4096
num of steps per iteration	100
num of epochs	5
gradient clipping	1.0
adam epsilon	$1e-8$
PPO	
clip range	0.2
entropy coefficient	0.01
discount factor γ	0.99
GAE balancing factor λ	0.95
desired KL-divergence	0.01
actor and double critic NN	MLP, hidden units [512, 216, 128]
BEAMDOJO	
w_1, w_2	1.0, 0.25