

Learning Humanoid Locomotion with Perceptive Internal Model

Junfeng Long*, Junli Ren*, Moji Shi*, Zirui Wang, Tao Huang, Ping Luo, Jiangmiao Pang

Abstract—In contrast to quadruped robots that can navigate diverse terrains using a “blind” policy, humanoid robots require accurate perception for stable locomotion due to their high degrees of freedom and inherently unstable morphology. However, incorporating perceptual signals often introduces additional disturbances to the system, potentially reducing its robustness, generalizability, and efficiency. This paper presents the Perceptive Internal Model (PIM), which relies on onboard, continuously updated elevation maps centered around the robot to perceive its surroundings. We train the policy using ground-truth obstacle heights surrounding the robot in simulation, optimizing it based on the Hybrid Internal Model (HIM), and perform inference with heights sampled from the constructed elevation map. Unlike previous methods that directly encode depth maps or raw point clouds, our approach allows the robot to perceive the terrain beneath its feet clearly and is less affected by camera movement or noise. Furthermore, since depth map rendering is not required in simulation, our method introduces minimal additional computational costs and can train the policy in 3 hours on an RTX 4090 GPU. We verify the effectiveness of our method across various humanoid robots, various indoor and outdoor terrains, stairs, and various sensor configurations. Our method can enable a humanoid robot to continuously climb stairs and has the potential to serve as a foundational algorithm for the development of future humanoid control methods.

I. INTRODUCTION

The control methods for legged robots have advanced rapidly in recent years [1], [2], driven by robot learning techniques such as reinforcement learning and simulation-to-real transfer. Among these methods, “blind” policies, which rely solely on proprioceptive information as observations, have proven effective in enabling quadruped robots to achieve robust and agile locomotion across various terrains [3], [4], [5].

Due to the higher degrees of freedom and inherently unstable morphology of humanoid robots, it is significantly more difficult to achieve stable locomotion across diverse terrains using a single policy, making the integration of perceptual information essential. However, considering that incorporating perceptual signals may introduce additional disturbances and costs to the system, potentially compromising its robustness, generalizability, and efficiency, most recent efforts in humanoid locomotion have followed the “blind” policy [6], [7], [8] or a two-phase training paradigm [9], [10]. They may rely on prior motion trajectories [11], [12], [13] or multi-stage, computationally expensive training processes [9] to execute human-like locomotion actions that can be deployed

on real-world robots. While these approaches successfully enable humanoid robots to walk on flat surfaces and mildly uneven terrains such as low stairs, they often end up falling and can not handle more challenging scenarios. Humanoid Parkour Learning [10] exploits perceptive information but is not able to overcome terrains requiring fine-grained footholds such as continuous stairs.


This paper presents the Perceptive Internal Model (PIM), which relies on onboard, continuously updated elevation maps centered around the robot to perceive its surroundings. It is built upon the Hybrid Internal Model (HIM) [3], which uses batch-level contrastive learning to optimize the simulated robot’s response, incorporating both explicit velocity and implicit latent, to the robot’s successor state. To incorporate perceptive information, PIM directly uses the ground-truth obstacle height maps surrounding the robot as additional observations to train the policy in simulation. During inference, we construct the elevation map using a LiDAR or RGB-D camera and sample points from it to align with the policy’s observations.

In contrast to previous methods that directly encode depth maps or raw point clouds, PIM advances in: (a) It benefits from HIM, which supports batch-level learning, single-stage training, and high training efficiency. (b) The elevation map accounts for the robot’s odometry and continuously maintains a larger local map, allowing the robot to clearly perceive the terrain beneath its feet. Compared to raw sensory data, it is more robust to sensor movement and noise. (c) Rendering depth images in simulation incurs additional costs in both computational memory and efficiency, whereas directly querying terrain heights does not. Furthermore, mitigating the domain gap between simulated depth maps and real camera data is challenging. As a result, our method introduces minimal additional computational costs and completes policy training within 3 hours on an RTX 4090 GPU.


We verify the effectiveness of our proposed method across various humanoid robots, including Unitree H1 and Fourier GR-1, on a range of challenging indoor and outdoor terrains, such as stairs, gaps, and high platforms, and with different sensor configurations, including Mid-360 LiDAR and RealSense D435+T265 cameras. Our method enables humanoid robots to navigate a variety of difficult terrains with a natural gait and a high success rate. It achieves a success rate of over 90% when continuously climbing stairs. We hope that PIM can serve as a foundational algorithm for the development of future humanoid control algorithms.

All authors are with Shanghai AI Laboratory. * denotes equal contribution. Junli Ren and Ping Luo are with the University of Hong Kong. Zirui Wang is with Zhejiang University. Tao Huang is with Shanghai Jiao Tong University. Corresponding author: Jiangmiao Pang (pangjiangmiao@gmail.com). Project Page: <https://junfeng-long.github.io/PIM/>.

(a) One-stage Simulation Training



(b) Humanoid-centered Elevation map



(c) Zero-shot Sim2real Deployments




Fig. 1: We propose a perceptive humanoid locomotion policy capable of mastering various challenging terrains. This policy requires only a few hours of single-stage training and can be deployed on hardware in a zero-shot manner. By integrating a LiDAR-based elevation map, we achieve accurate perception without the need for historical vision embedding.

II. RELATED WORK

A. Learning Humanoid Locomotion

The rise of the humanoid robot industry has motivated extensive research of humanoid locomotion in the past year [11], [9], [13], [14], [6], [15], [16], [12], [17], [8], [18], [7], [19], [20], [21], [10], a majority of these researches follows the “training in simulation & Sim2Real transfer” procedure, akin to the methods employed in quadruped research. Indeed, when it comes to humanoid robots, there is an emphasis on performing human-like movements [12], [22]. To achieve this, [11], [7], [13] utilize prior motion trajectories to learn natural whole-body motion while walking on diverse flat grounds. Despite the expressive motions demonstrated

in these studies, subsequent research has indicated that learning from scratch can also yield coordinated whole-body movement [20], [10], especially when the task is focused on traversing challenging terrains.

On the other hand, learning a robust legged locomotion policy requires the prediction of accurate robot states from accessible observations. Given that humanoid robots typically have higher dimensions of action and observation space, recent works have proposed different training diagrams to address this issue: [8] incorporates linear velocity into observations from a separate state estimator, [10], [9] encompasses multi-stage training process to narrow the sim2real gap, [6] predicts the states from observation as part of a denoising process.

Unlike the previous works, we introduce a novel locomotion policy that accomplishes state prediction utilizing perceptive information, thereby eliminating the necessity for additional training stages or explicit decoding of the predicted states.

B. Perceptive Legged Locomotion

Previous research on quadrupeds has formulated several successful systems that introduce exteroceptive sensors into the locomotion policy [23], [24], [25], [26]. These policies, which either receive heightmaps or depth images, have demonstrated that integrating a more comprehensive vision system and vision observations into the policy results in better performance on challenging terrains [27]. In the context of humanoid robots, while the aforementioned blind policies have demonstrated robust performance on various surfaces [7], against disturbances [11], [12], and on simple stairs [6], a comprehensive vision-aided policy is necessary to achieve perceptive locomotion. A significant challenge in perceptive locomotion is reducing the perception gap during sim2real transfer. Existing research [10] employs an extra phase of training to simulate the extensive hardware noise. However, this approach compromises perception accuracy and prevents the policy from accomplishing tasks that demand more precise perception, such as stair climbing.

Compared to existing methods, we propose a more comprehensive perceptive locomotion system that achieves more precise hardware perception by incorporating a LiDAR-based elevation map. The integration of a robot-centered elevation map eliminates the need for introducing vision into the historical observation sequence and achieves a more comprehensive state prediction. Extensive experiments on terrains involving stairs, gaps, and platforms underscore the superiority of the proposed method compared to existing studies.

III. METHODOLOGY

A. Framework Overview

Compared to quadruped robots, humanoid robots have significantly lower stability due to their morphology. The perceptive information is necessary for the control system. It needs to determine the movements of its actuators for a desired velocity given its proprioceptive information and perceptive information. We model the humanoid locomotion task as a sequential decision problem. The entire framework is depicted in Fig. 2, we construct a perceptive internal model that utilizes both proprioceptive information and perceptive information for state estimation and optimize the policy with PPO [28].

B. Observations

The policy observations \mathbf{o}_t consist of commands, proprioceptive information, perceptive information, and action of last timestep \mathbf{a}_{t-1} . The commands are the desired velocity $\mathbf{c}_t = [v_x^c, v_y^c, \omega_{yaw}^c]$ indicates the linear velocity in longitudinal and lateral directions, and the angular velocity in the horizontal direction, respectively. The proprioceptive

information includes its joint position θ_t , joint velocity $\dot{\theta}_t$, base angular velocity ω_t and gravity direction in robot frame \mathbf{g}_t . The perceptive information \mathbf{p}_t is an elevation map around the robot. We divide these observations into two parts, non-perceptive part \mathbf{o}_t^n and perceptive part \mathbf{o}_t^p , i.e. $\mathbf{o}_t^n = [\mathbf{c}_t, \omega_t, \mathbf{g}_t, \theta_t, \dot{\theta}_t, \mathbf{a}_{t-1}]$ and $\mathbf{o}_t^p = [\mathbf{p}_t]$. Unlike the policy, the critic is allowed to access privileged information such as the linear velocity v_t of the robot at the training stage and provide a more accurate estimation of state values.

C. Preliminary: Hybrid Internal Model

HIM (Hybrid Internal Model) [3] provides accurate state estimation and sim2real ability. Basically, HIM estimates the robot's linear velocity in the next timestep and proprioceptive information in a latent space given its proprioceptive observation history. The linear velocity estimation is trained with simple regression to the ground truth linear velocity obtained from the simulator. The next step proprioceptive information prediction is trained through contrastive learning. HIM considers a pair of \mathbf{o}_{t+1}^n and $\mathbf{o}_{t-H:t}^n$ from the same trajectory as a positive pair and others as negative pairs. The pairs are optimized by swapping assignments tasks similar to SwAV[29].

D. Perceptive Internal Model

a) *Elevation Maps*: In the real-world experiments, we employ the elevation map module [27] to capture perceptive observations as demonstrated in Fig. 3. It requires both odometry and the ground point cloud to generate a grid-based representation of ground heights. Unlike applications involving quadrupeds, it is critical to ensure that the z-axis of the map frame aligns with the direction of gravity. Thus, we implement an initial rotation of the map frame to ensure alignment with the average acceleration direction (i.e. gravity direction). Additionally, the point cloud is processed to filter out points based on height, preserving only the ground points. The filtered point cloud and odometry data are then inputted into the elevation map module to produce the grid-based height map in the map frame. We test elevation maps with two sensor setups: the first setup utilized a single Mid-360 LiDAR to provide both odometry, integrated via FAST-LIO [30], [31], and point cloud data. The second setup employed a Realsense T265 camera for odometry and a Realsense D435 camera for point cloud acquisition. While both setups provide accurate elevation maps during stable locomotion, the LiDAR-based setup demonstrated enhanced robustness in scenarios involving rapid or irregular movements.

b) *Elevation Sampling*: We sample 96 points in the frame of the robot aligning its z-axis to the negative direction of gravity. The points are distributed in a $0.8m \times 1.2m$ square with the robot as the center. These points' relative z coordinates to the base link are used as our perceptive input.

c) *Terrain-based Perceptive Internal Model*: To make the most of perceptive information, we not only let the policy use it for foothold planning but also use it for state

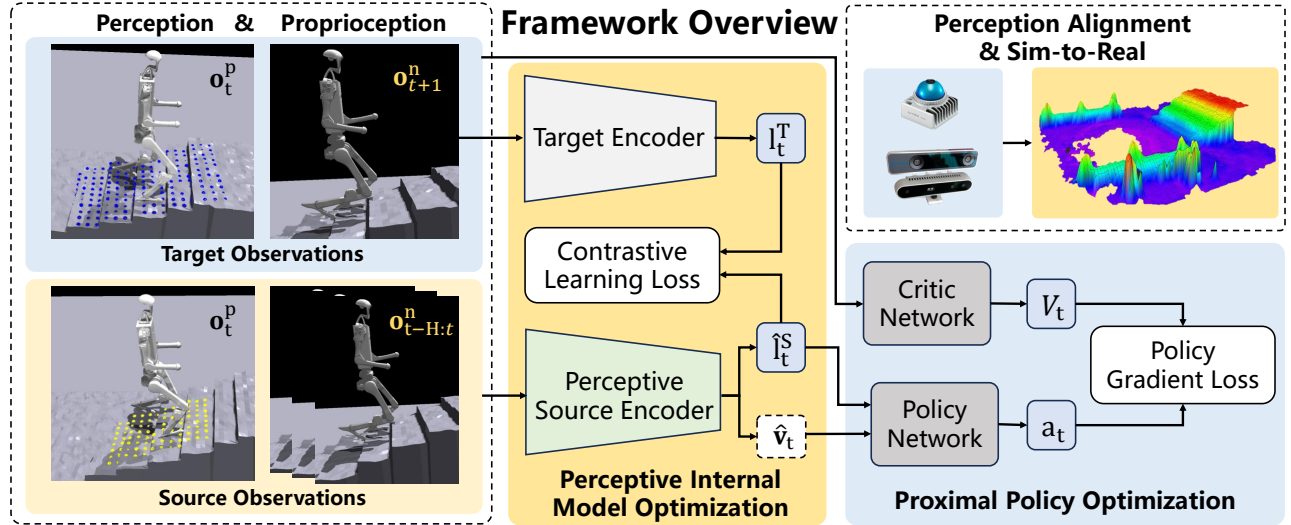


Fig. 2: Overview of our framework. Within PIM, we integrate perceptive information into the state predictor to achieve more comprehensive and accurate state prediction. A LiDAR-based elevation map serves as the perception model, enabling more precise perception alignment between simulation and real-world environments.




Fig. 3: Terrian Perception module implemented by a single LiDAR. The map frame overlaps the initial position of the torso link and is fixed with the world. The colored grid map demonstrates the height of the terrain in the map frame.

estimation, which we concluded as PIM (Perceptive Internal Model). PIM is similar to HIM, however, instead of using proprioceptive observation history to estimate the robot's next state, current perceptive observation \mathbf{o}_t^p is concatenated with the proprioceptive observation history $\mathbf{o}_{t-H:t}^n$ to provide more information for more accurate estimation of next-step state since terrain plays an important role in robot state transfer. The effectiveness is shown in Section IV-A.

E. Action Curriculum

We introduce a curriculum on action space to simplify the training. To be more specific, we choose some of the joints that are less important in the locomotion task, i.e., joints of

the robot's arm, and joints of the robot's waist. We set the range of these joints to zero and increase their range as the training goes on.

F. Symmetry Regularization

We also introduce a symmetry regularization technique similar to [32] to improve the harmony of the gait. In detail, we involve three operators: G_o^n which flips the robot's proprioceptive observation with respect to the x-z plane, G_o^p which flips the robot's perceptive observation with respect to the x-z plane, and G_a which flips the robot's action with respect to the x-z plane. Then during optimization of PIM, we augment the data with G_o . During policy optimization, we involve two additional losses:

$$\begin{aligned} \mathcal{L}_{symmetry}^{policy} &= MSE(G_a(\pi(\mathbf{o}_t^n, \mathbf{o}_t^p, PIM(\mathbf{o}_{t-H:t}^n, \mathbf{o}_t^p)), \\ &\quad \pi(G_o^n(\mathbf{o}_t^n), G_o^p(\mathbf{o}_t^p), PIM(G_o^n(\mathbf{o}_{t-H:t}^n), G_o^p(\mathbf{o}_t^p)))) \\ \mathcal{L}_{symmetry}^{value} &= MSE(V(\mathbf{o}_t^n, \mathbf{o}_t^p), V(G_o^n(\mathbf{o}_t^n), G_o^p(\mathbf{o}_t^p))) \end{aligned} \quad (1)$$

G. Training Process

At each time, the proprioception observation history $\mathbf{o}_{t-H:t}^n$ and perceptive observation \mathbf{o}_t^p are fed into the perceptive internal model to obtain an estimation of the robot's linear velocity $\hat{\mathbf{v}}_{t+1}$ and a latent variable l_t predicting the robot's next proprioception information. Then they are fed into the policy network with current observation \mathbf{o}_t to obtain action \mathbf{a}_t . During the policy optimization process, the perceptive internal model is frozen and we only update the policy network and value network using PPO. After policy optimization, we will use the collected trajectories to optimize the perceptive internal model.

H. Reward Functions

Our reward functions guide the robot to follow certain velocity commands and maintain soft contact with the ground. Besides some reward functions borrowed from quadrupedal locomotion, we designed some new rewards for humanoid locomotion. For example, we wish that the distance of two feet is not too close, then we involve the reward called **feet lateral distance**, which computes the lateral distance of two feet in the robot's frame and penalizes distance closer than d_{min} . We also wish that the feet of the humanoid robot are parallel to the ground, so we use a reward function called **feet ground parallel** by adding five sample points on each foot of the humanoid robot, indicating the front, middle, hind, left and right, then penalize the variance of distances of these sample points to the ground. Detailed reward functions are listed in Table I.

TABLE I: Rewards

Reward	Equation (r_i)	Weight (w_i)
Lin. velocity tracking	$\exp\left\{-\frac{\ \mathbf{v}_{xy}^{cmd} - \mathbf{v}_{xy}\ _2^2}{\sigma}\right\}$	1.0
Ang. velocity tracking	$\exp\left\{-\frac{(\omega_{yaw}^{cmd} - \omega_{yaw})^2}{\sigma}\right\}$	1.0
Linear velocity (z)	v_z^2	-0.5
Angular velocity (xy)	$\ \omega_{xy}\ _2^2$	-0.025
Orientation	$\ \mathbf{g}_x\ _2^2 + \ \mathbf{g}_y\ _2^2$	-1.25
Joint accelerations	$\ \dot{\theta}\ _2^2$	-2.5×10^{-7}
Joint power	$\frac{\ \tau\ \dot{\theta}^T}{\ \mathbf{v}\ _2^2 + 0.2\ \omega\ _2^2}$	-2.5×10^{-5}
Body height w.r.t. feet	$(h_{target} - h)^2$	0.1
Feet clearance	$\sum_{feet} (p_z^{target} - p_z^i)^2 \cdot v_{xy}^i$	-0.25
Action rate	$\ \mathbf{a}_t - \mathbf{a}_{t-1}\ _2^2$	-0.01
Smoothness	$\ \mathbf{a}_t - 2\mathbf{a}_{t-1} + \mathbf{a}_{t-2}\ _2^2$	-0.01
Feet stumble	$\mathbf{1}\{\exists i, \mathbf{F}_i^{xy} > 3 F_i^z \}$	-3.0
Torques	$\sum_{all\ joints} \left \frac{\tau_i}{k p_i}\right ^2$	-2.5×10^{-6}
Joint velocity	$\sum_{all\ joints} \dot{\theta}_i^2$	-1×10^{-4}
Joint tracking error	$\sum_{all\ joints} \theta_i - \theta_i^{target} ^2$	-0.25
Arm joint deviation	$\sum_{arm\ joints} \theta_i - \theta_i^{default} ^2$	-0.1
Hip joint deviation	$\sum_{hip\ joints} \theta_i - \theta_i^{default} ^2$	-0.5
Waist joint deviation	$\sum_{waist\ joints} \theta_i - \theta_i^{default} ^2$	-0.25
Joint pos limits	$\sum_{all\ joints} out_i$	-2.0
Joint vel limits	$\sum_{all\ joints} RELU(\hat{\theta}_i - \hat{\theta}_i^{max})$	-0.1
Torque limits	$\sum_{all\ joints} RELU(\hat{\tau}_i - \hat{\tau}_i^{max})$	-0.1
No fly	$\mathbf{1}\{\text{only one feet on ground}\}$	0.25
Feet lateral distance	$ y_{left\ foot}^B - y_{right\ foot}^B - d_{min}$	2.5
Feet slip	$\sum_{feet} \mathbf{v}_i^{toot} * \sim \mathbf{1}_{new\ contact}$	-0.25
Feet ground parallel	$\sum_{feet} Var(H_i)$	-2.0
Feet contact force	$\sum_{feet} RELU(F_i^z - F_{th})$	-2.5×10^{-4}
Feet parallel	$Var(D)$	-2.5
Contact momentum	$\sum_{feet} v_i^z * F_i^z $	-2.5×10^{-4}

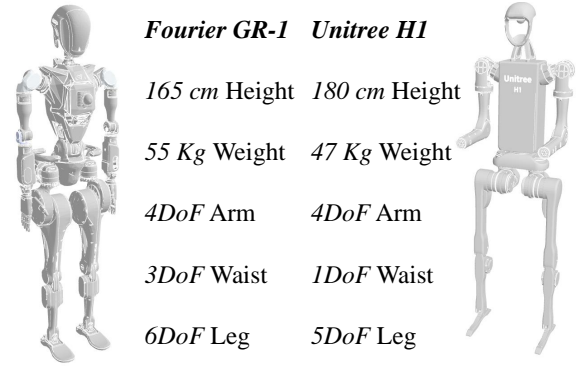


Fig. 4: Robot Hardware Setups in terms of Height, Weight, and Degree of Freedom (DoF)




Fig. 5: Estimation Loss and Terrain Level During Training

IV. RESULTS

In the experiments, we validate the efficiency and generalizability of the proposed humanoid locomotion on two humanoid robots: Unitree H1 and Fourier GR-1 (Fig. 4), the experiments emphasize:

- Effectiveness of Perception in estimating next-step state.
- Effectiveness of Perceptive Locomotion: We demonstrate the proposed method successfully navigates extreme terrains including stairs ($\geq 0.15m$), wooden platforms ($\geq 0.4m$), slopes (15°), and gaps.
- Effectiveness of the Perceptive Internal Model(PIM): The terrain-based internal model accurately predicts the robot's states, enhancing the stability and adaptability of the locomotion.
- Cross-Platform Validation: We validate the proposed method on two different humanoid robots in challenging environments, demonstrating its robustness in mastering disturbances and difficult terrains.

A. PIM vs HIM

We compare the estimation accuracy of PIM and HIM, here the difference is that PIM uses perception for estimating while HIM does not, both of their corresponding policies use perception for locomotion. As shown in Fig 5, where terrain level indicates the difficulty of terrain that the robot can traverse and estimation loss is the error between estimated velocity and real velocity. PIM can achieve higher training efficiency and provide higher estimation accuracy, as well as enable robots to traverse more difficult terrain.




Fig. 6: We conducted extensive hardware experiments to validate the effectiveness of the proposed method across different humanoid robots. (a) H1 traverses stairs with heights of 15 cm. (b) H1 consecutively steps onto a wooden platform and jumps over a gap between two platforms. (c) We successfully deploy the proposed policy on GR1 .

B. Stairs

We validate that the proposed method enables the robot to successfully traverse stairs with a height of 15 cm (Fig. 6-a), surpassing existing methods in the literature, such as [6], [8], which achieve a maximum height of only 10 cm. This demonstrates the enhanced capability of the proposed perceptive locomotion in handling more challenging terrains. In addition, we would like to highlight the challenge posed by the foot length, which equals the stair width (30cm). This constraint requires precise foot placement to avoid stepping too far forward, which would cause the foot to get stuck on the next step and prevent proper lifting, or stepping too far back, which would result in a loss of balance and the robot falling. This highlights the critical need for accurate state prediction and perception, further demonstrating the effectiveness of the proposed perceptive locomotion policy in handling challenging environments with precise control.

C. High platform & gap

Compared to stair climbing, jumping onto high platforms and between gaps demands less perception accuracy but requires the policy to respond more rapidly to changes in perception. As demonstrated in Fig. 6-b, H1 successfully performs consecutive jumps onto platforms and over gaps, which highlights the effectiveness of the proposed method in perceptive integrated state estimation and executing extreme parkour actions.

D. Cross-Platform Validation

In addition to H1, we applied the same training process of the proposed method to Fourier GR-1, unlike H1, whose ankle has only one tandem DOF, GR-1 has two DOF at the ankles powered by **parallel joints**. Our method can still enable it to traverse various terrains as demonstrated in Fig. 6-c. This highlights the method’s robustness to different robot configurations, such as variations in height, mass, joint configuration, and foot design. These results suggest the

potential of the proposed method to be developed into a generalized solution for humanoid locomotion across diverse platforms.

E. Whole-body Movement

As a whole-body control method, we observed that the robot performs coordinated upper-body movements to maintain balance without prior knowledge, particularly during extreme actions such as jumping. In general, the robot avoided walking with the same hand or foot during flat-ground tests. During jumping motions, as demonstrated in the fourth image of Fig. 6-b, the robot swings its left arm forward heavily as its right leg jumps, maintaining balance during these extreme actions. This outcome indicates that with appropriate reward design and model prediction as achieved by our method, the robot can generate simple, human-like movements without prior data or imitation learning.

V. CONCLUSION

In this work, we propose an effective framework for perceptive humanoid locomotion in complex terrain. The framework includes Perceptive Internal Model(PIM) for next-state prediction, which utilizes perceptive information in addition to HIM. We also present some novel techniques in humanoid locomotion including action curriculum, symmetry regularization, and reward function design. We demonstrate success in sim-to-real deployment on different humanoid robots with different heights, mass, joint configurations, and foot designs. To the best of our knowledge, this work is the first perceptive locomotion framework that works well across different robot platforms.

REFERENCES

- [1] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [2] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [3] J. Long, Z. Wang, Q. Li, L. Cao, J. Gao, and J. Pang, "Hybrid internal model: Learning agile legged locomotion with simulated robot response," in *The Twelfth International Conference on Learning Representations*, 2024.
- [4] J. Long, W. Yu, Q. Li, Z. Wang, D. Lin, and J. Pang, "Learning h-infinity locomotion control," in *8th Annual Conference on Robot Learning*, 2024.
- [5] G. B. Margolis and P. Agrawal, "Walk these ways: Tuning robot control for generalization with multiplicity of behavior," in *Conference on Robot Learning*. PMLR, 2023, pp. 22–31.
- [6] X. Gu, Y.-J. Wang, X. Zhu, C. Shi, Y. Guo, Y. Liu, and J. Chen, "Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning," *arXiv preprint arXiv:2408.14472*, 2024.
- [7] I. Radosavovic, B. Zhang, B. Shi, J. Rajasegaran, S. Kamat, T. Darrell, K. Sreenath, and J. Malik, "Humanoid locomotion as next token prediction," *arXiv preprint arXiv:2402.19469*, 2024.
- [8] Q. Liao, B. Zhang, X. Huang, X. Huang, Z. Li, and K. Sreenath, "Berkeley humanoid: A research platform for learning-based control," *arXiv preprint arXiv:2407.21781*, 2024.
- [9] W. Cui, S. Li, H. Huang, B. Qin, T. Zhang, L. Zheng, Z. Tang, C. Hu, N. Yan, J. Chen *et al.*, "Adapting humanoid locomotion over challenging terrain via two-phase training," in *8th Annual Conference on Robot Learning*.

- [10] Z. Zhuang, S. Yao, and H. Zhao, "Humanoid parkour learning," *arXiv preprint arXiv:2406.10759*, 2024.
- [11] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang, "Expressive whole-body control for humanoid robots," *arXiv preprint arXiv:2402.16796*, 2024.
- [12] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi, "Learning human-to-humanoid real-time whole-body teleoperation," *arXiv preprint arXiv:2403.04436*, 2024.
- [13] P. Dugar, A. Shrestha, F. Yu, B. van Marum, and A. Fern, "Learning multi-modal whole-body control for real-world humanoid robots," *arXiv preprint arXiv:2408.07295*, 2024.
- [14] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, "Humanplus: Humanoid shadowing and imitation from humans," *arXiv preprint arXiv:2406.10454*, 2024.
- [15] X. Gu, Y.-J. Wang, and J. Chen, "Humanoid-gym: Reinforcement learning for humanoid robot with zero-shot sim2real transfer," *arXiv preprint arXiv:2404.05695*, 2024.
- [16] N. Hansen, J. SV, V. Sobal, Y. LeCun, X. Wang, and H. Su, "Hierarchical world models as visual whole-body humanoid controllers," *arXiv preprint arXiv:2405.18418*, 2024.
- [17] Z. Jiang, Y. Xie, J. Li, Y. Yuan, Y. Zhu, and Y. Zhu, "Harmon: Whole-body motion generation of humanoid robots from language descriptions," in *8th Annual Conference on Robot Learning*.
- [18] J. Li, Y. Zhu, Y. Xie, Z. Jiang, M. Seo, G. Pavlakos, and Y. Zhu, "Okami: Teaching humanoid robots manipulation skills through single video imitation," in *8th Annual Conference on Robot Learning*.
- [19] B. van Marum, A. Shrestha, H. Duan, P. Dugar, J. Dao, and A. Fern, "Revisiting reward design and evaluation for robust humanoid standing and walking," *arXiv preprint arXiv:2404.19173*, 2024.
- [20] C. Zhang, W. Xiao, T. He, and G. Shi, "Wococo: Learning whole-body humanoid control with sequential contacts," *arXiv preprint arXiv:2406.06005*, 2024.
- [21] W. Zhao, Y. Zhao, J. Pajarinen, and M. Muehlebach, "Bi-level motion imitation for humanoid robots," in *8th Annual Conference on Robot Learning*.
- [22] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. Kitani, C. Liu, and G. Shi, "Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning," *arXiv preprint arXiv:2406.08858*, 2024.
- [23] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, "Extreme parkour with legged robots," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 443–11 450.
- [24] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Conference on robot learning*. PMLR, 2023, pp. 403–415.
- [25] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024.
- [26] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [27] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, "Elevation mapping for locomotion and navigation using gpu," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 2273–2280.
- [28] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [29] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in neural information processing systems*, 2020.
- [30] W. Xu and F. Zhang, "Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.
- [31] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lio2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [32] Z. Su, X. Huang, D. Ordoñez-Apaez, Y. Li, Z. Li, Q. Liao, G. Turrissi, M. Pontil, C. Semini, Y. Wu *et al.*, "Leveraging symmetry in rl-based legged locomotion control," *arXiv preprint arXiv:2403.17320*, 2024.