HumanoidPF: Collision-Free Humanoid Traversal via Geometric Guidance

Fig. 1: Using a single generalist policy, our humanoid robot achieves collision-free traversal in cluttered indoor environments, including (a) detouring through narrow passages, (b) crouching under low-hanging obstacles, (c) squeezing through tight indoor spaces, and (d) hurdling over objects scattered on the floor.

***Abstract*— We study the problem of collision-free humanoid traversal in cluttered indoor scenes, such as hurdling over objects scattered on the floor, crouching under low-hanging obstacles, or squeezing through narrow passages. To achieve this goal, the humanoid needs to map its perception of surrounding obstacles with diverse spatial layouts and geometries to the corresponding traversal skills. In this work, we propose Humanoid Potential Field (HumanoidPF), a geometric representation that encodes humanoid–obstacle relationships as collision-free motion directions, significantly facilitating RL-based traversal skill learning. We further observe that HumanoidPF exhibits a surprisingly small sim-to-real gap as a perceptual representation. To enable generalizable traversal skills, the policy is trained on diverse and challenging cluttered indoor scenes via a teacher-to-student distillation pipeline. We successfully transfer the learned policy to the real world, where users can command the humanoid to traverse cluttered indoor environments with a single click. Demos and code can be found in our website: https://axian12138.github.io/CAT/.**

I. Introduction

Humanoid robots operating in indoor environments frequently need to traverse cluttered spaces while avoiding collisions with surrounding objects. For instance, a domestic robot moving between rooms may need to hurdle over objects scattered on the floor, crouch under low-hanging obstacles, or squeeze through narrow passages. Such tasks require the robot to perceive obstacles with diverse spatial layouts and geometries and map them to appropriate traversal behaviors.

While legged locomotion in complex environments has seen rapid progress for quadrupeds [1]–[17] and humanoids [18]–[30], existing approaches are often limited in cluttered indoor scenes with full-spatial obstacle layouts and intricate geometries. A key challenge lies in the lack of an effective representation for humanoid–obstacle relationships during collision avoidance. Most prior works [9]–[17], [22], [29] obtain penalty signals only after collisions occur, resulting in sparse supervision for reinforcement learning (RL). Meanwhile, policies are typically trained on raw, high-dimensional environmental measurements without explicitly modeling humanoid–obstacle spatial relationships.

We identify that existing humanoid traversal learning suffers from lack of structured humanoid–obstacle representation. To address this, we propose HumanoidPF, a geometric representation that: (1) provides directional guidance before collision occurs, (2) enables dense reward shaping, (3) significantly improves generalization.

Inspired by Artificial Potential Fields (APF) [31], HumanoidPF represents the influence of surrounding obstacles as a continuous gradient field that indicates collision-free motion directions. First, the field is queried at multiple key body parts and used as policy observations, providing

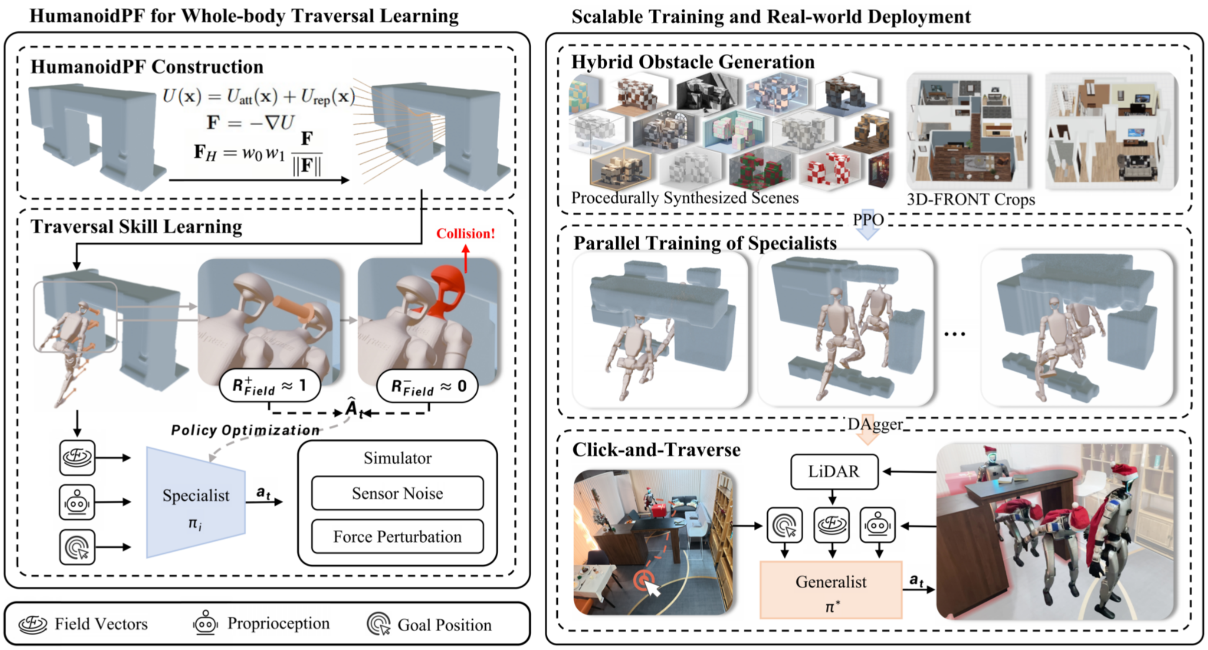

Fig. 2: **Pipeline overview.** Left: HumanoidPF for whole-body traversal learning, providing perceptual representation and collision-avoidance reward. Right: Scalable training and real-world deployment.

directional cues for collision-free motion. Second, the field naturally induces a reward structure that encourages the policy to align its motion with safe traversal directions. This representation provides informative and anticipatory guidance for RL and improves generalization across scenes.

To further improve robustness in diverse cluttered environments, we adopt a hybrid scene generation strategy that combines crops of realistic 3D indoor scenes with procedurally generated obstacles. This exposes the robot to a wide range of challenging obstacle configurations during training. We successfully transfer the learned policy to the real world, where a user can simply click a target location to command the humanoid to safely traverse cluttered indoor environments.

Our contributions are summarized as follows:

- We propose **HumanoidPF**, a geometric representation that explicitly encodes humanoid–obstacle relationships and facilitates RL-based traversal skill learning.
- We introduce a hybrid scene generation strategy that improves robustness and generalization across complex indoor environments.
- We deploy the learned policy in a real-world teleoperation system that enables intuitive humanoid traversal in cluttered indoor scenes.

II. Method

We study collision-free humanoid traversal in cluttered indoor scenes. Given a target position $\mathbf{g} \in \mathbb{R}^3$ and a set of indoor obstacles $\mathcal{O} = \{O_i\}_{i=1}^N$, the humanoid needs to reach $\mathbf{g}$ without colliding with $\mathcal{O}$. Our method consists of two parts: (i) **HumanoidPF**, which encodes humanoid–obstacle relationships to facilitate traversal learning, and (ii) a hybrid scene generation strategy that improves generalization to diverse and challenging indoor scenes.

A. *HumanoidPF for whole-body traversal learning*

We extend classical Artificial Potential Fields (APF) to **HumanoidPF** that explicitly considers whole-body coordination and supports learning-based whole-body traversal.

1) HumanoidPF construction

We construct a guidance field by combining an attractive component that pulls the humanoid toward the goal and a repulsive component that pushes it away from nearby obstacles. The resulting field provides a directional signal indicating safe, collision-free motion.

Directly applying such a field to a humanoid can lead to inconsistent guidance across different body parts. To address this, we introduce a priority-weighting scheme to coordinate whole-body behavior.

Specifically, we assign higher importance to the root body, which dominates global motion and balance, while other body parts receive lower base weights. In addition, we incorporate a collision-urgency mechanism that increases the influence of body parts that are closer to obstacles or moving toward them.

The final guidance field is obtained by normalizing the combined field and modulating it with these weights. This design amplifies informative signals from critical body parts while suppressing conflicting directions, resulting in a coherent guidance field for coordinated whole-body traversal.

2) Traversal skill learning with HumanoidPF

Observation. We construct a compact observation by querying the HumanoidPF at a set of key body parts. Each queried vector represents the locally safe motion direction induced by surrounding obstacles and the goal. This allows the policy to directly reason about traversal decisions, without relying on raw geometric inputs.

Reward. A key advantage of HumanoidPF is that it provides dense and anticipatory supervision. Instead of only

	Side-Hurdle		Side-Crouch		Multi-Hurdle		Hurdle-Crouch		Side-Hurdle-Crouch	
	SR(%) ↑	*DE(m)* ↓	*SR(%)* ↑	*DE(m)* ↓	*SR(%)* ↑	*DE(m)* ↓	*SR(%)* ↑	*DE(m)* ↓	*SR(%)* ↑	*DE(m)* ↓
ASTraversal	37.1±3.1	0.54±0.32	56.0±9.9	0.48±0.05	82.1±8.7	0.26±0.43	28.1±10.4	1.11±0.78	0.5±0.5	1.06±0.39
Humanoid Parkour	45.1±4.5	0.62±0.39	64.4±19.3	0.56±0.16	88.7±2.6	0.23±0.35	33.3±6.1	1.16±0.63	0.4±0.3	1.49±0.04
Ours w/o OBS_{Field}	60.4±9.6	0.53±0.68	90.1±5.3	0.19±0.35	90.5±3.5	0.09±0.11	77.8±5.4	0.33±0.23	53.7±9.9	0.59±0.08
Ours w/o R_{Field}	71.4±9.9	0.5±0.34	80.3±15.4	0.23±0.06	88.3±14.1	0.23±0.5	21.9±15.8	1.27±0.71	0.0±0.0	1.57±0.003
Ours	**95.4**±3.9	**0.06**±0.34	**96.9**±2.1	**0.05**±0.09	**95.0**±4.9	**0.06**±0.1	**93.9**±2.7	**0.08**±0.16	**86.6**±5.2	**0.2**±0.32

TABLE I: **Validation of HumanoidPF for skill learning.** Performance across five obstacle traversal scenarios.

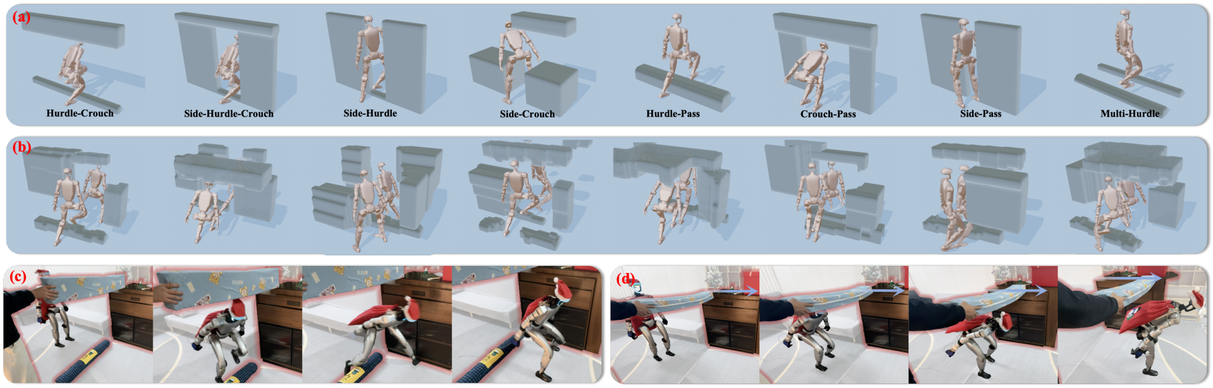

Fig. 3: **Collision-free humanoid traversal in both simulation and the real world.**

penalizing collisions after they occur, the field encourages the humanoid to align its motion with safe directions before collisions happen.

Concretely, for each body part, we encourage its motion direction to align with the corresponding HumanoidPF vector. The strength of this guidance is adaptively scaled based on the magnitude of the field, so that body parts under higher collision risk receive stronger supervision.

The overall reward aggregates these alignment signals across all body parts. This design provides dense guidance for policy learning, improves sample efficiency, and promotes coordinated whole-body obstacle avoidance.

B. Scalable training in diverse scenes

To improve robustness and generalization, we train the policy in a diverse set of cluttered indoor environments constructed from both realistic and synthetic sources. Specifically, we combine crops of real 3D indoor scenes with procedurally generated obstacles. The realistic scenes provide natural layouts and object geometries, while the synthetic obstacles introduce challenging and highly constrained configurations that are rare in existing datasets.

This hybrid setup exposes the humanoid to a wide range of traversal scenarios. Importantly, our HumanoidPF-based formulation generalizes across diverse environments without requiring any task-specific reward tuning.

III. Experiment

We evaluate our method in the MuJoCo [32] simulator to validate the effectiveness of **HumanoidPF** for humanoid traversal in cluttered environments.

A. Validation of HumanoidPF for skill learning

We compare our method with existing approaches on a set of cluttered traversal scenarios.

Experiment Setting. To analyze performance under different obstacle layouts and geometries, we design five types of cluttered scenes for evaluation. All scenes are manually generated and include diverse obstacle configurations.

Metrics. We evaluate performance using Success Rate (SR, %), defined as reaching within 0.1 m of the goal in 5 s without collision, and Distance Error (DE, m), the closest horizontal distance to the goal during traversal.

Baselines. We compare against ASTraversal [14] (elevation map-based), Humanoid Parkour [22] (collision penalty-based), and two ablations: replacing HumanoidPF observations with elevation maps (w/o OBS_{Field}) and removing directional rewards in favor of collision penalties (w/o R_{Field}).

Results. As shown in Table I, our method consistently outperforms all baselines across all scenarios. Notably, prior methods that rely on elevation maps or post-hoc collision penalties struggle in cluttered environments, where safe traversal requires anticipating collisions and coordinating whole-body motion.

In contrast, HumanoidPF provides explicit directional guidance before collisions occur, enabling the policy to make informed traversal decisions. The ablation results further confirm that both the observation (OBS_{Field}) and reward (R_{Field}) components are critical for performance.

IV. Conclusion

We introduce HumanoidPF, a representation that encodes humanoid–obstacle relationships for RL-based collision-free traversal, and demonstrate its robust whole-body traversal capability across diverse obstacle configurations. However, the current framework does not yet exploit contact-rich interactions (e.g., leaning on or stepping onto support surfaces), which remains an important direction for future work.

References

[1] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, "Perceptive locomotion through nonlinear model-predictive control," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3402–3421, 2023.

[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] G. Gagné-Labelle, V. Atanassov, and I. Havoutis, "Acrobotics: A generalist approach to quadrupedal robots' parkour," in *Annual Conference Towards Autonomous Robotic Systems*. Springer, 2025, pp. 124–137.

[4] C. Zhang, N. Rudin, D. Hoeller, and M. Hutter, "Learning agile locomotion on risky terrains," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 11 864–11 871.

[5] Q. Liao, Z. Li, A. Thirugnanam, J. Zeng, and K. Sreenath, "Walking in narrow spaces: Safety-critical locomotion control for quadrupedal robots with duality-based optimization," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 2723–2730.

[6] J. Lee, M. Bjelonic, A. Reske, L. Wellhausen, T. Miki, and M. Hutter, "Learning robust autonomous navigation and locomotion for wheeled-legged robots," *Science Robotics*, vol. 9, no. 89, p. eadi9641, 2024.

[7] R. Yang, G. Yang, and X. Wang, "Neural volumetric memory for visual locomotion control," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 1430–1440.

[8] N. Rudin, D. Hoeller, M. Bjelonic, and M. Hutter, "Advanced skills by learning locomotion and local navigation end-to-end," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 2497–2503.

[9] Z. Wang, T. Ma, Y. Jia, X. Yang, J. Zhou, W. Ouyang, Q. Zhang, and J. Liang, "Omni-perception: Omnidirectional collision avoidance for legged locomotion in dynamic environments," *arXiv preprint arXiv:2505.19214*, 2025.

[10] 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.

[11] T. Miki, J. Lee, L. Wellhausen, and M. Hutter, "Learning to walk in confined spaces using 3d representation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8649–8656.

[12] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, "Robot parkour learning," *arXiv preprint arXiv:2309.05665*, 2023.

[13] N. Rudin, J. He, J. Aurand, and M. Hutter, "Parkour in the wild: Learning a general and extensible agile locomotion policy using multi-expert distillation and rl fine-tuning," *arXiv preprint arXiv:2505.11164*, 2025.

[14] Y. Chen, J. Ma, Z. Luo, Y. Han, Y. Dong, B. Xu, and P. Lu, "Learning autonomous and safe quadruped traversal of complex terrains using multi-layer elevation maps," *IEEE Robotics and Automation Letters*, 2025.

[15] T. He, C. Zhang, W. Xiao, G. He, C. Liu, and G. Shi, "Agile but safe: Learning collision-free high-speed legged locomotion," *arXiv preprint arXiv:2401.17583*, 2024.

[16] 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.

[17] 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.

[18] Z. Li, J. Zeng, S. Chen, and K. Sreenath, "Autonomous navigation of underactuated bipedal robots in height-constrained environments," *The International Journal of Robotics Research*, vol. 42, no. 8, pp. 565–585, 2023.

[19] J. He, C. Zhang, F. Jenelten, R. Grandia, M. Bächer, and M. Hutter, "Attention-based map encoding for learning generalized legged locomotion," *Science Robotics*, vol. 10, no. 105, p. eadv3604, 2025.

[20] J. Sun, G. Han, P. Sun, W. Zhao, J. Cao, J. Wang, Y. Guo, and Q. Zhang, "Dpl: Depth-only perceptive humanoid locomotion via realistic depth synthesis and cross-attention terrain reconstruction," *arXiv preprint arXiv:2510.07152*, 2025.

[21] H. Duan, B. Pandit, M. S. Gadde, B. Van Marum, J. Dao, C. Kim, and A. Fern, "Learning vision-based bipedal locomotion for challenging terrain," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 56–62.

[22] Z. Zhuang, S. Yao, and H. Zhao, "Humanoid parkour learning," *arXiv preprint arXiv:2406.10759*, 2024.

[23] J. Long, J. Ren, M. Shi, Z. Wang, T. Huang, P. Luo, and J. Pang, "Learning humanoid locomotion with perceptive internal model," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 9997–10 003.

[24] J. Ren, T. Huang, H. Wang, Z. Wang, Q. Ben, J. Long, Y. Yang, J. Pang, and P. Luo, "Vb-com: Learning vision-blind composite humanoid locomotion against deficient perception," *arXiv preprint arXiv:2502.14814*, 2025.

[25] T. Huang, R. Xu, Y. Wang, W. Gao, and S. Zhang, "Traversing narrow paths: A two-stage reinforcement learning framework for robust and safe humanoid walking," *arXiv preprint arXiv:2508.20661*, 2025.

[26] H. Wang, Z. Wang, J. Ren, Q. Ben, T. Huang, W. Zhang, and J. Pang, "Beamdojo: Learning agile humanoid locomotion on sparse footholds," *arXiv preprint arXiv:2502.10363*, 2025.

[27] W. Sun, B. Cao, L. Chen, Y. Su, Y. Liu, Z. Xie, and H. Liu, "Learning perceptive humanoid locomotion over challenging terrain," *arXiv preprint arXiv:2503.00692*, 2025.

[28] A. Allshire, H. Choi, J. Zhang, D. McAllister, A. Zhang, C. M. Kim, T. Darrell, P. Abbeel, J. Malik, and A. Kanazawa, "Visual imitation enables contextual humanoid control," *arXiv preprint arXiv:2505.03729*, 2025.

[29] Q. Ben, B. Xu, K. Li, F. Jia, W. Zhang, J. Wang, J. Wang, D. Lin, and J. Pang, "Gallant: Voxel grid-based humanoid locomotion and local-navigation across 3d constrained terrains," *arXiv preprint arXiv:2511.14625*, 2025.

[30] Z. Zhang, J. Guo, C. Chen, J. Wang, C. Lin, Y. Lian, H. Xue, Z. Wang, M. Liu, J. Lyu, *et al.*, "Track any motions under any disturbances," *arXiv preprint arXiv:2509.13833*, 2025.

[31] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The international journal of robotics research*, vol. 5, no. 1, pp. 90–98, 1986.

[32] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.