Perceptive Humanoid Parkour: Chaining Dynamic Human Skills via Motion Matching

Anonymous CVPR submission

Paper ID *****

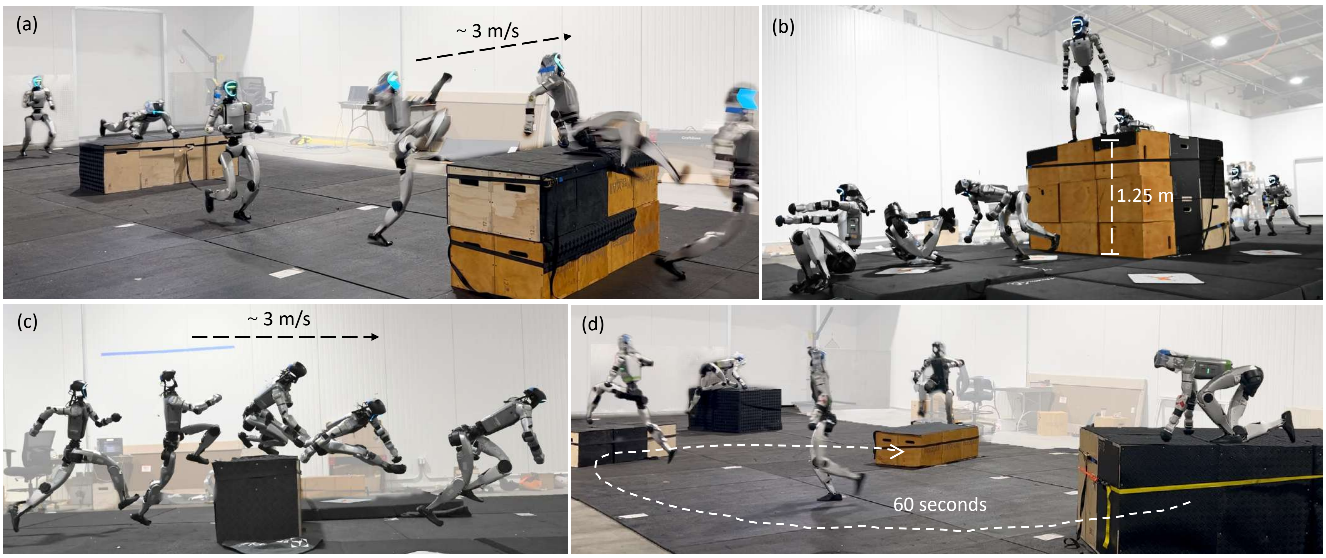

Figure 1. **Perceptive Humanoid Parkour (PHP)** enables a Unitree G1 humanoid robot to execute highly dynamic, long-horizon parkour behaviors using onboard perception. By composing various agile human skills via motion matching and a teacher-student training pipeline, we train a single multi-skill visuomotor policy capable of complex contact-rich maneuvers.

Abstract

*We present Perceptive Humanoid Parkour (*PHP*), a modular framework for autonomous, vision-based humanoid parkour. Our approach uses motion matching to compose retargeted atomic human skills into long-horizon kinematic trajectories with smooth transitions and natural dynamics. We then train motion-tracking RL experts and distill them into a single depth-based multi-skill policy via DAgger and RL. Using only onboard depth sensing and a discrete 2D velocity command, the robot autonomously selects and executes skills such as stepping over, climbing, vaulting, and rolling off obstacles of varying geometries and heights. Real-world experiments on a Unitree G1 humanoid demonstrate highly dynamic parkour behaviors, including climbing obstacles up to 1.25,m and long-horizon multi-obstacle traversal with closed-loop adaptation to obstacle perturbations.*

1. Introduction

Achieving agile and adaptive humanoid locomotion over complex terrains remains a central challenge in robotics. Humans traverse diverse environments by rapidly selecting and chaining dynamic whole-body skills based on visual context. In this work, we study parkour as a testbed for this capability. Parkour requires highly dynamic and contact-rich behaviors, such as climbing body-height obstacles and vaulting within fractions of a second, while tightly coupling control with perception for adaptation to terrain variation and perturbations. Moreover, solving long-horizon obstacle courses requires consolidating many dynamic skills into a single visuomotor policy.

Human motion priors have become essential for learning dynamic humanoid behaviors [15, 30], but highly dynamic motion data is inherently scarce, often containing only a few short demonstrations per skill. To address this limitation, we leverage motion matching [4, 7] to compose retargeted atomic skills into diverse long-horizon trajectories through nearest-neighbor retrieval in a feature space, densifying sparse motion libraries while preserving realistic transitions.

Learning a visuomotor policy for dozens of highly dynamic skills further requires perceptive inputs that can be efficiently simulated and transferred to the real world. Prior work typically trains privileged state-based experts and distills them into vision-based students using DAgger [23], but pure imitation is insufficient for humanoid parkour, where compounding errors can quickly derail dynamic maneuvers. To address this, we augment distillation with an RL objective that provides task-level corrective feedback, improving traversal robustness across diverse skills.

To this end, we present Perceptive Humanoid Parkour (PHP), a modular framework that combines human motion priors, long-horizon skill composition, and perceptive control. We first retarget human motions into robot-compatible atomic skills, then compose them into long-horizon trajectories via motion matching. Finally, we train motion-tracking experts and distill them into a single depth-conditioned multi-skill policy using DAgger and RL, enabling autonomous transitions between skills such as stepping, climbing, and vaulting using onboard depth sensing.

2. Related Works

2.1. Perceptive Terrain Traversal for Legged Robots

Perception has enabled legged robots to traverse challenging terrains such as sparse footholds, gaps, and tall obstacles [1, 20, 31, 33, 34]. Building on these advances, prior work demonstrated parkour-style terrain traversal on quadrupeds through agile jumping and climbing behaviors [6, 11, 17, 24]. While these quadrupedal skills can often be learned from scratch via reward shaping, this approach scales poorly to humanoids due to the difficulty of high-dimensional whole-body control.

As a result, perceptive humanoid locomotion has primarily focused on lower-dynamic tasks such as stair climbing, sparse terrain traversal, and stepping onto low platforms [2, 10, 16, 25, 32, 37]. To improve training efficiency, many works adopt a teacher-student framework where privileged experts are distilled into vision-based policies via DAgger [6, 24]. We follow this paradigm, but augment distillation with RL to improve robustness for highly dynamic humanoid skills.

2.2. Humanoid Skill Chaining with Human Motion Data

Human motion priors enable agile and natural humanoid behaviors [5, 15, 21, 27, 30, 36], but long-horizon skill chaining remains challenging due to the heterogeneous nature of motion data. Prior work addressed this either by learning unified skill distributions with RL [22, 26, 29], or by generating intermediate kinematic trajectories with learned motion models [8, 13, 18, 28, 35]. However, these approaches often struggle in the low-data regimes common in highly dynamic parkour.

Instead, we leverage motion matching [4, 12], a simple yet effective nearest-neighbor retrieval approach widely used in animation and games [3, 9]. While previously explored mainly for simpler quadruped behaviors [14], we show that motion matching provides high-quality long-horizon references for chaining dynamic humanoid parkour skills over complex terrains.

3. Adaptive and Agile Long-Horizon Parkour

3.1. Overview

We first synthesize long-horizon kinematic references by composing locomotion and atomic parkour skills via motion matching. We then train motion-tracking expert policies with privileged observations and distill them into a single depth-based student policy using DAgger and PPO for zero-shot sim-to-real transfer (Fig. 2).

3.2. Skill Composition via Motion Matching

We adopt motion matching [4, 7] as an offline motion synthesis module for composing scarce atomic parkour skills into long-horizon trajectories. Each database frame contains a pose $\boldsymbol{q}_i$ and feature vector $\boldsymbol{x}_i$ encoding future trajectory, foot states, and root velocity [12]. Given the current state and a desired velocity command, we construct a query feature $\hat{\boldsymbol{x}}_t$ and retrieve the nearest frame

$$i_t^\star = \arg\min_{i \in \mathcal{C}_t} \|\hat{\boldsymbol{x}}_t - \boldsymbol{x}_i\|^2, \tag{1}$$

where $\mathcal{C}_t$ is the search window. Playback then transitions to the matched frame with short blending to smooth transitions.

We synthesize long-horizon parkour trajectories by composing locomotion with atomic parkour skills:

$$\texttt{Locomotion} \rightarrow \texttt{Skill} \rightarrow \texttt{Locomotion}.$$

Locomotion serves as a shared transition manifold, enabling scalable composition without requiring pairwise transitions between all skills.

For each skill clip, we annotate start/end frames (s_k, e_k) and define a pre-skill entry window

$$\mathcal{E}_k := [s_k - H_k, s_k], \tag{2}$$

corresponding to the approach phase before the maneuver. During locomotion, motion matching searches the locomotion database. For locomotion-to-skill transitions, search is restricted to $\mathcal{E}_k$, after which the skill clip is replayed sequentially until e_k. Motion matching is then resumed to return to locomotion.

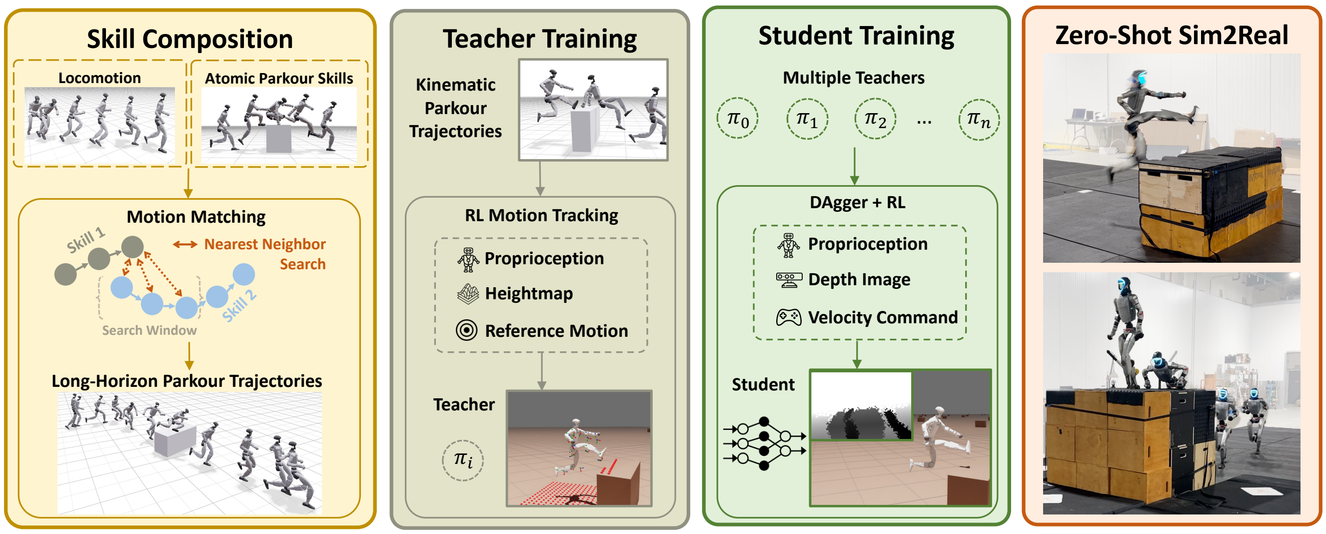

Figure 2. Perceptive Humanoid Parkour system overview.

To generate training trajectories, we roll out this composition process using randomized velocity commands, locomotion durations, and terrain configurations. Motion matching naturally produces diverse skill entrances with different stride phases and approach distances, substantially densifying sparse motion data. We additionally randomize obstacle geometry, pose, and distractors to improve robustness.

3.3. Learning a Highly-Dynamic Visuomotor Policy

We first train skill-specific motion-tracking experts and then distill them into a unified student policy.

3.3.1. Training Expert Policies

We follow prior motion-tracking frameworks [15, 30]. Experts observe reference motion targets, proprioception, and privileged terrain information including a local height scan. We additionally provide global root states to improve recovery from drift during terrain-coupled motions. Training uses tracking rewards, adaptive sampling, and standard domain randomization.

3.3.2. Distilling a Unified Student Policy

Prior work commonly applies DAgger [23] to distill multiple experts into a single visuomotor policy. However, we find pure imitation insufficient for highly dynamic skills such as climbing and vaulting, where successful execution depends on brief high-magnitude control bursts. Per-step imitation losses do not account for rollout outcomes and therefore may fail to encourage actions required for successful traversal.

To address this, we jointly optimize DAgger and PPO:

$$\mathcal{L} = \lambda_{\text{PPO}}\mathcal{L}_{\text{PPO}} + \lambda_D\mathcal{L}_D, \tag{3}$$

with a curriculum that gradually shifts from imitation to RL.

Table 1. Baseline success rate comparison.

Commanded Velocity	1.0 m/s			2.0 m/s		
	36 cm	58 cm	76 cm	36 cm	58 cm	76 cm
Velocity Tracking	**1.00**	0.00	0.00	**1.00**	0.00	0.00
Uncomposed Data	0.06	0.02	0.00	0.37	0.27	0.07
End-to-end Depth	0.95	0.07	0.08	0.78	0.19	0.14
Ours	**1.00**	**0.99**	**0.95**	**1.00**	**0.99**	**0.95**

PPO provides task-level corrective feedback that improves robustness across diverse dynamic skills.

The student policy observes proprioception, depth images rendered with Nvidia WARP [19], and velocity commands. We apply camera randomization, depth noise, latency randomization, and standard sim-to-real domain randomization. To improve stability, we gradually relax termination thresholds and enable PPO-specific exploration control only after sufficient imitation pretraining.

4. Experiments

4.1. Real-World Results

Human-Level Agility. The robot successfully executes highly dynamic parkour maneuvers, including climbing a 1.25 m wall (96% robot height) in 2.67 s, cat vaults, and drop landings (Fig. 3(a,b)). Furthermore, despite training only on single-obstacle trajectories, the policy generalizes to long-horizon multi-obstacle courses by autonomously composing stepping and climbing behaviors while adapting online to obstacle perturbations (Fig. 3(c)).

4.2. Simulation Results

4.2.1. Baseline Comparison

We compare against: (1) reward-shaping velocity tracking without human motion priors, (2) training on uncomposed atomic skills without motion matching, and (3) end-

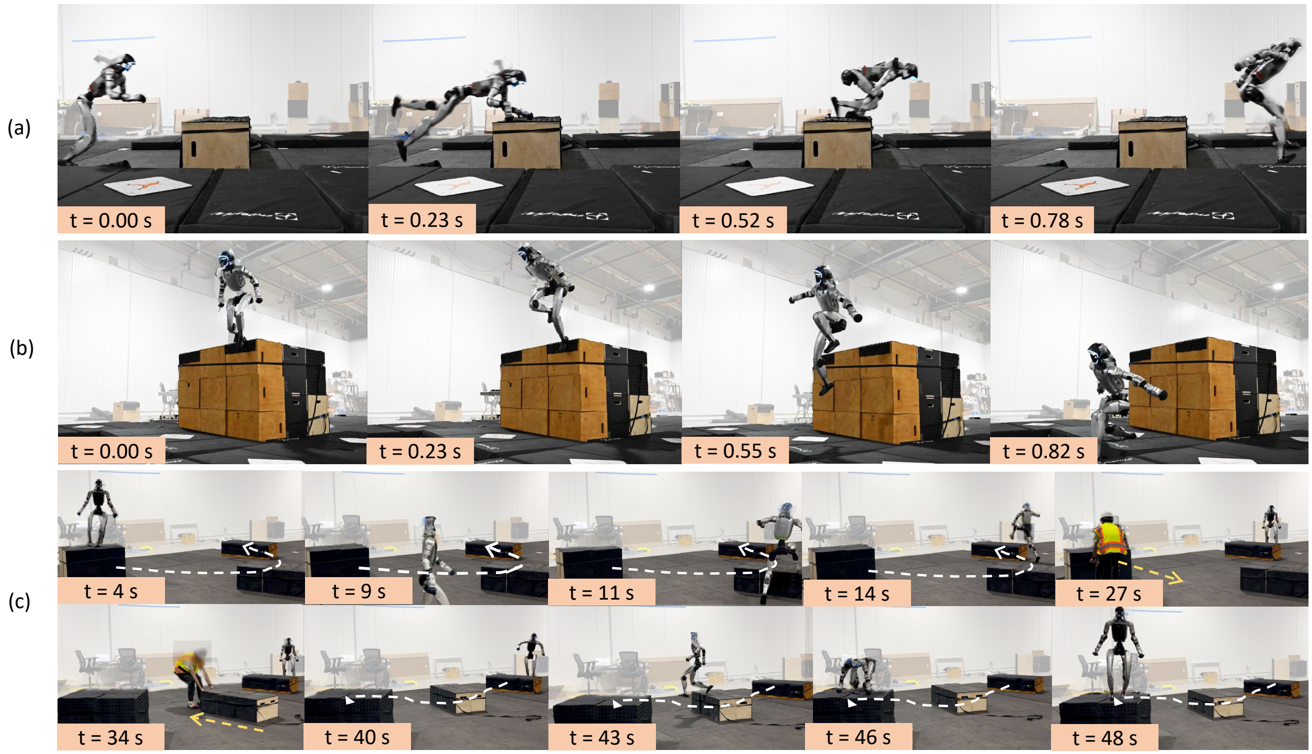

Figure 3. Hardware results including (a) cat vault, (b) drop landing, and (c) long-horizon terrain traversal.

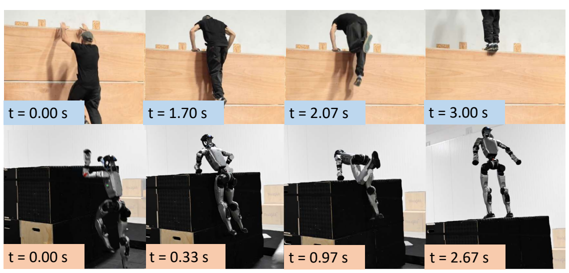

Figure 4. High-wall climb comparison.

Table 2. Success rate under different motion-matching densities and distillation strategies.

Method	1.0 m/s			2.0 m/s		
	58 cm	76 cm	94 cm	36 cm	58 cm	76 cm
Extreme Distances	0.99	0.62	0.64	0.98	0.60	0.58
Half Density	0.95	0.32	0.57	0.99	0.85	0.81
DAgger Only	0.16	0.03	0.12	0.63	0.09	0.10
Ours	0.99	**0.95**	**1.00**	**1.00**	**0.98**	**0.90**

to-end depth-policy training without expert distillation. As shown in Tab. 1, reward shaping fails on higher obstacles, uncomposed motion data lacks effective long-horizon skill transitions, and end-to-end depth training struggles on highly dynamic maneuvers due to exploration difficulty. In contrast, our framework achieves consistently high success rates across obstacle heights and speeds.

4.2.2. Ablation Study

We study motion-matching density and RL during distillation. Reducing motion-matching density significantly lowers success rates, especially on harder tasks, highlighting the importance of densereferences. We further find that DAgger-only training performs poorly on highly dynamic maneuvers, while our DAgger+RL distillation substantially improves robustness and traversal success.

5. Conclusion

We presented Perceptive Humanoid Parkour, a framework for autonomous long-horizon humanoid parkour using onboard perception. By combining motion-matching-based skill composition with a teacher-student RL pipeline, our approach enables agile and adaptive whole-body parkour behaviors on a Unitree G1 humanoid with zero-shot sim-to-real transfer. We show that motion matching provides effective long-horizon references, while augmenting DAgger with RL improves distillation of highly dynamic skills into a unified depth-based policy. Our current system remains limited by perception and hardware constraints, including short-range sensing and limited hand capabilities. Future work includes incorporating semantic scene understanding and richer conditioning signals for more flexible humanoid behaviors.

References

[1] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on robot learning*, pages 403–415. PMLR, 2023. 2

[2] Qingwei Ben, Botian Xu, Kailin Li, Feiyu Jia, Wentao Zhang, Jingping Wang, Jingbo Wang, Dahua Lin, and Jiangmiao Pang. Gallant: Voxel grid-based humanoid locomotion and local-navigation across 3d constrained terrains, 2025. 2

[3] Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. Drecon: data-driven responsive control of physics-based characters. *ACM Transactions On Graphics (TOG)*, 38(6):1–11, 2019. 2

[4] Michael Büttner and Simon Clavet. Motion matching - the road to next gen animation. Proc. of Nucl.ai, 2015. 1, 2

[5] Zixuan Chen, Mazeyu Ji, Xuxin Cheng, Xuanbin Peng, Xue Bin Peng, and Xiaolong Wang. Gmt: General motion tracking for humanoid whole-body control. *arXiv preprint arXiv:2506.14770*, 2025. 2

[6] Xuxin Cheng, Kexin Shi, Ananye Agarwal, and Deepak Pathak. Extreme parkour with legged robots. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11443–11450. IEEE, 2024. 2

[7] Simon Clavet. Motion matching and the road to next-gen animation. Proc. of GDC, 2016. 1, 2

[8] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. *arXiv preprint arXiv:2406.10454*, 2024. 2

[9] Ruiyu Gou, Michiel van de Panne, and Daniel Holden. Control operators for interactive character animation. *ACM Transactions on Graphics (TOG)*, 2025. 2

[10] Junzhe He, Chong Zhang, Fabian Jenelten, Ruben Grandia, Moritz Bächer, and Marco Hutter. Attention-based map encoding for learning generalized legged locomotion. *Science Robotics*, 10(105):eadv3604, 2025. 2

[11] David Hoeller, Nikita Rudin, Dhionis Sako, and Marco Hutter. Anymal parkour: Learning agile navigation for quadrupedal robots. *Science Robotics*, 9(88):eadi7566, 2024. 2

[12] Daniel Holden, Anas Kanoun, Michiel Bŭttner, Sofien Bouaziz, Sebastian Thrun, and Aaron Hertzmann. Learned motion matching. *ACM Transactions on Graphics (TOG)*, 2020. 2

[13] Dvij Kalaria, Sudarshan S Harithas, Pushkal Katara, Sangkyung Kwak, Sarthak Bhagat, Shankar Sastry, Srinath Sridhar, Sai Vemprala, Ashish Kapoor, and Jonathan Chung-Kuan Huang. Dreamcontrol: Human-inspired whole-body humanoid control for scene interaction via guided diffusion. *arXiv preprint arXiv:2509.14353*, 2025. 2

[14] Dongho Kang, Simon Zimmermann, and Stelian Coros. Animal gaits on quadrupedal robots using motion matching and model-based control. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8500–8507. IEEE, 2021. 2

[15] Qiayuan Liao, Takara E Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C Karen Liu. Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion. *arXiv preprint arXiv:2508.08241*, 2025. 1, 2, 3

[16] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9997–10003. IEEE, 2025. 2

[17] Shixin Luo, Songbo Li, Ruiqi Yu, Zhicheng Wang, Jun Wu, and Qiuguo Zhu. Pie: Parkour with implicit-explicit learning framework for legged robots. *IEEE Robotics and Automation Letters*, 2024. 2

[18] Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Sirui Chen, Fernando Castañeda, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, et al. Sonic: Supersizing motion tracking for natural humanoid whole-body control. *arXiv preprint arXiv:2511.07820*, 2025. 2

[19] Miles Macklin. Warp: A high-performance python framework for gpu simulation and graphics. https://github.com/nvidia/warp, 2022. NVIDIA GPU Technology Conference (GTC). 3

[20] 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. 2

[21] Yixuan Pan, Ruoyi Qiao, Li Chen, Kashyap Chitta, Liang Pan, Haoguang Mai, Qingwen Bu, Hao Zhao, Cunyuan Zheng, Ping Luo, et al. Agility meets stability: Versatile humanoid control with heterogeneous data. *arXiv preprint arXiv:2511.17373*, 2025. 2

[22] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (TOG)*, 2021. 2

[23] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011. 2, 3

[24] Nikita Rudin, Junzhe He, Joshua Aurand, and Marco Hutter. Parkour in the wild: Learning a general and extensible agile locomotion policy using multi-expert distillation and rl finetuning. *arXiv preprint arXiv:2505.11164*, 2025. 2

[25] Huayi Wang, Zirui Wang, Junli Ren, Qingwei Ben, Tao Huang, Weinan Zhang, and Jiangmiao Pang. Beamdojo: Learning agile humanoid locomotion on sparse footholds. In *Robotics: Science and Systems (RSS)*, 2025. 2

[26] Jinze Wu, Guiyang Xin, Chenkun Qi, and Yufei Xue. Learning robust and agile legged locomotion using adversarial motion priors. *IEEE Robotics and Automation Letters*, 8(8): 4975–4982, 2023. 2

[27] Weiji Xie, Jinrui Han, Jiakun Zheng, Huanyu Li, Xinzhe Liu, Jiyuan Shi, Weinan Zhang, Chenjia Bai, and Xuelong Li. Kungfubot: Physics-based humanoid whole-body control for learning highly-dynamic skills. *arXiv preprint arXiv:2506.12851*, 2025. 2

[28] Michael Xu, Yi Shi, KangKang Yin, and Xue Bin Peng. Parc: Physics-based augmentation with reinforcement learning for character controllers. In *ACM SIGGRAPH*, 2025. 2

[29] Pei Xu, Zhen Wu, Ruocheng Wang, Vishnu Sarukkai, Kayvon Fatahalian, Ioannis Karamouzas, Victor Zordan, and C Karen Liu. Learning to ball: Composing policies for long-horizon basketball moves. *ACM Transactions on Graphics (TOG)*, 44(6):1–14, 2025. 2

[30] Lujie Yang, Xiaoyu Huang, Zhen Wu, Angjoo Kanazawa, Pieter Abbeel, Carmelo Sferrazza, C Karen Liu, Rocky Duan, and Guanya Shi. Omniretarget: Interaction-preserving data generation for humanoid whole-body loco-manipulation and scene interaction. *arXiv preprint arXiv:2509.26633*, 2025. 1, 2, 3

[31] 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. 2

[32] Tae Hoon Yang, Haochen Shi, Jiacheng Hu, Zhicong Zhang, Daniel Jiang, Weizhuo Wang, Yao He, Zhen Wu, Yuming Chen, Yifan Hou, et al. Locomotion beyond feet. *arXiv preprint arXiv:2601.03607*, 2026. 2

[33] Alan Yu, Ge Yang, Ran Choi, Yajvan Ravan, John Leonard, and Phillip Isola. Learning visual parkour from generated images. In *8th Annual Conference on Robot Learning*, 2024. 2

[34] 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. 2

[35] Yanjie Ze, Siheng Zhao, Weizhuo Wang, Angjoo Kanazawa, Rocky Duan, Pieter Abbeel, Guanya Shi, Jiajun Wu, and C Karen Liu. Twist2: Scalable, portable, and holistic humanoid data collection system. *arXiv preprint arXiv:2511.02832*, 2025. 2

[36] Tong Zhang, Boyuan Zheng, Ruiqian Nai, Yingdong Hu, Yen-Jen Wang, Geng Chen, Fanqi Lin, Jiongye Li, Chuye Hong, Koushil Sreenath, et al. Hub: Learning extreme humanoid balance. *CoRL*, 2025. 2

[37] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. *arXiv:2406.10759*, 2024. 2