TTT-Parkour: Rapid Test-Time Training for Perceptive Legged Humanoid Parkour

Shaoting Zhu[12*] Baijun Ye[12*] Jiaxuan Wang[1†] Jiakang Chen[1†] Ziwen Zhuang[12] Linzhan Mou[3] Runhan Huang[1] Hang Zhao[12‡]

[1] Tsinghua University [2] Shanghai Qi Zhi Institute [3] Princeton University

* Equal contribution † Equal contribution ‡ Corresponding author

ttt-parkour.github.io

CoRL 2026, Austin

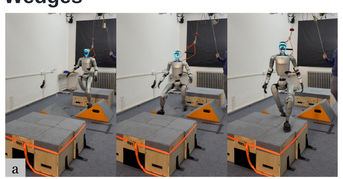
Research problem

How can a humanoid cross difficult terrain unseen during pre-training?

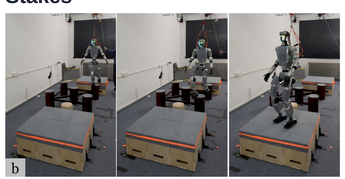
Reconstruct the target scene, fine-tune a pre-trained policy in simulation, then deploy the adapted policy on the robot.

Real-world parkour

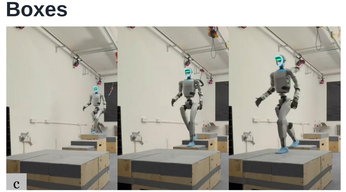
Wedges



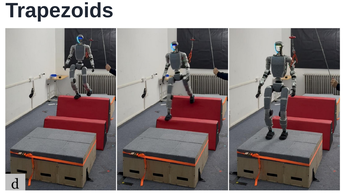
Stakes



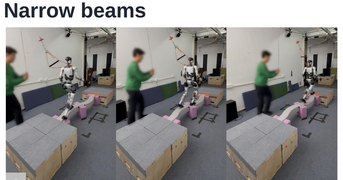
Boxes



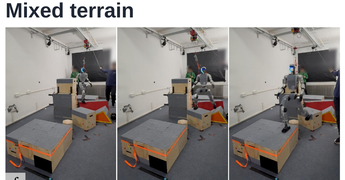
Trapezoids



Narrow beams



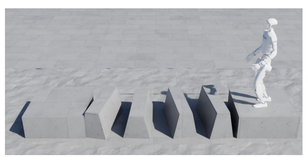
Mixed terrain



Legged traversal using foot contacts only. Ground contact or falling counts as failure.

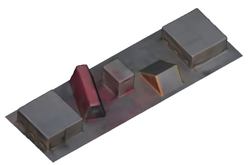
Real-to-sim-to-real adaptation

1 Pre-train



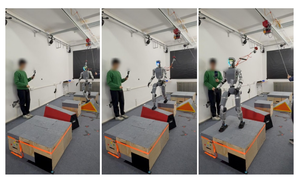
100k PPO iterations on 200 procedural terrains across five terrain families.

2 Adapt in simulation



Reconstruct the target terrain. Fully fine-tune the pre-trained policy on its mesh.

3 Deploy



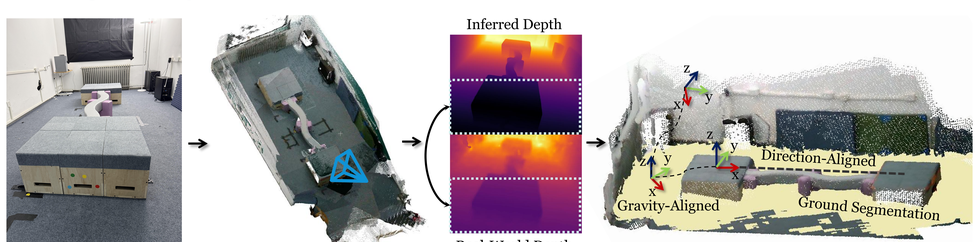
Deploy at 50 Hz with no further policy updates on the robot.

Perceptive policy

Depth history ➡ **CNN + proprioception** ➡ **Actor** ➡ **29 joint targets**

Asymmetric actor-critic with task, regularization, safety and AMP-style rewards.

Metric geometry for reliable transfer

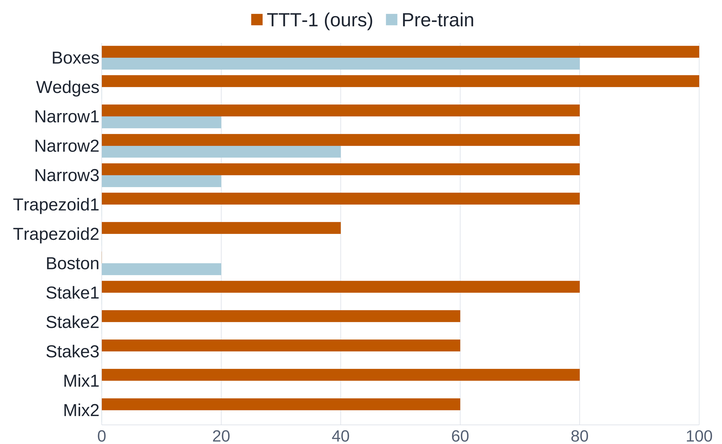


1. Real-World Capture 2. Feed-forward Reconstruction 3. Scale Recovery 4. Physically-consistent Frame Alignment

Sensor depth recovers metric scale. Ground and start/goal semantics align gravity and traversal direction.

Experimental results

Real-world traversal success (%)



5 real-world trials per terrain. Boston remains unsuccessful after adaptation.

Rapid specialization

Terrain / group	TTT-1 iterations
10 of 13 terrains	**≤120**
Stake2	**200**
Trapezoid2	**240**
Boston	**560**

Iterations to 97% simulation success. TTT-13 converges at 1,500 iterations with comparable final success: 99.56% vs. 99.83% for TTT-1.

Scope: manual capture and workstation fine-tuning. Transfer to unrelated terrains is limited. Benefits over stronger generalist policies remain untested.

TL;DR **TTT-Parkour** adapts a pre-trained humanoid policy to unseen terrain through **RGB-D mesh reconstruction** and **test-time training in simulation**, enabling direct deployment on real robots for parkour in **about 10 minutes** end to end on most tested terrains.