


Overview & Contributions

Plain RL with reward shaping yields lower-body-dominated humanoid gaits; imitation-based RL transfers coordinated skills but only *replays* references without perceptive adaptation. We couple a **diffusion-based motion generator** (terrain-aware online references) with an **RL-based whole-body tracker**, and **fine-tune the tracker in closed loop with the frozen generator**:


- Framework combining **diffusion-based motion generation** with **RL-based motion tracking** for perceptive whole-body humanoid locomotion.
- **Closed-loop RL fine-tuning** with a frozen 2-step diffusion generator mitigates the distribution mismatch of online-generated references; the tracker acts as a motion filter that suppresses unsafe behaviors.
- **Fully onboard** on Unitree G1 (LiDAR-based elevation mapping, TensorRT-accelerated diffusion ~ 0.02 s, 0.5 s receding-horizon control), supporting directional goal-reaching.
- Traverses boxes, hurdles, stairs, and **mixed-terrain sequences unseen during training** — without heavily engineered multi-teacher distillation.



Website

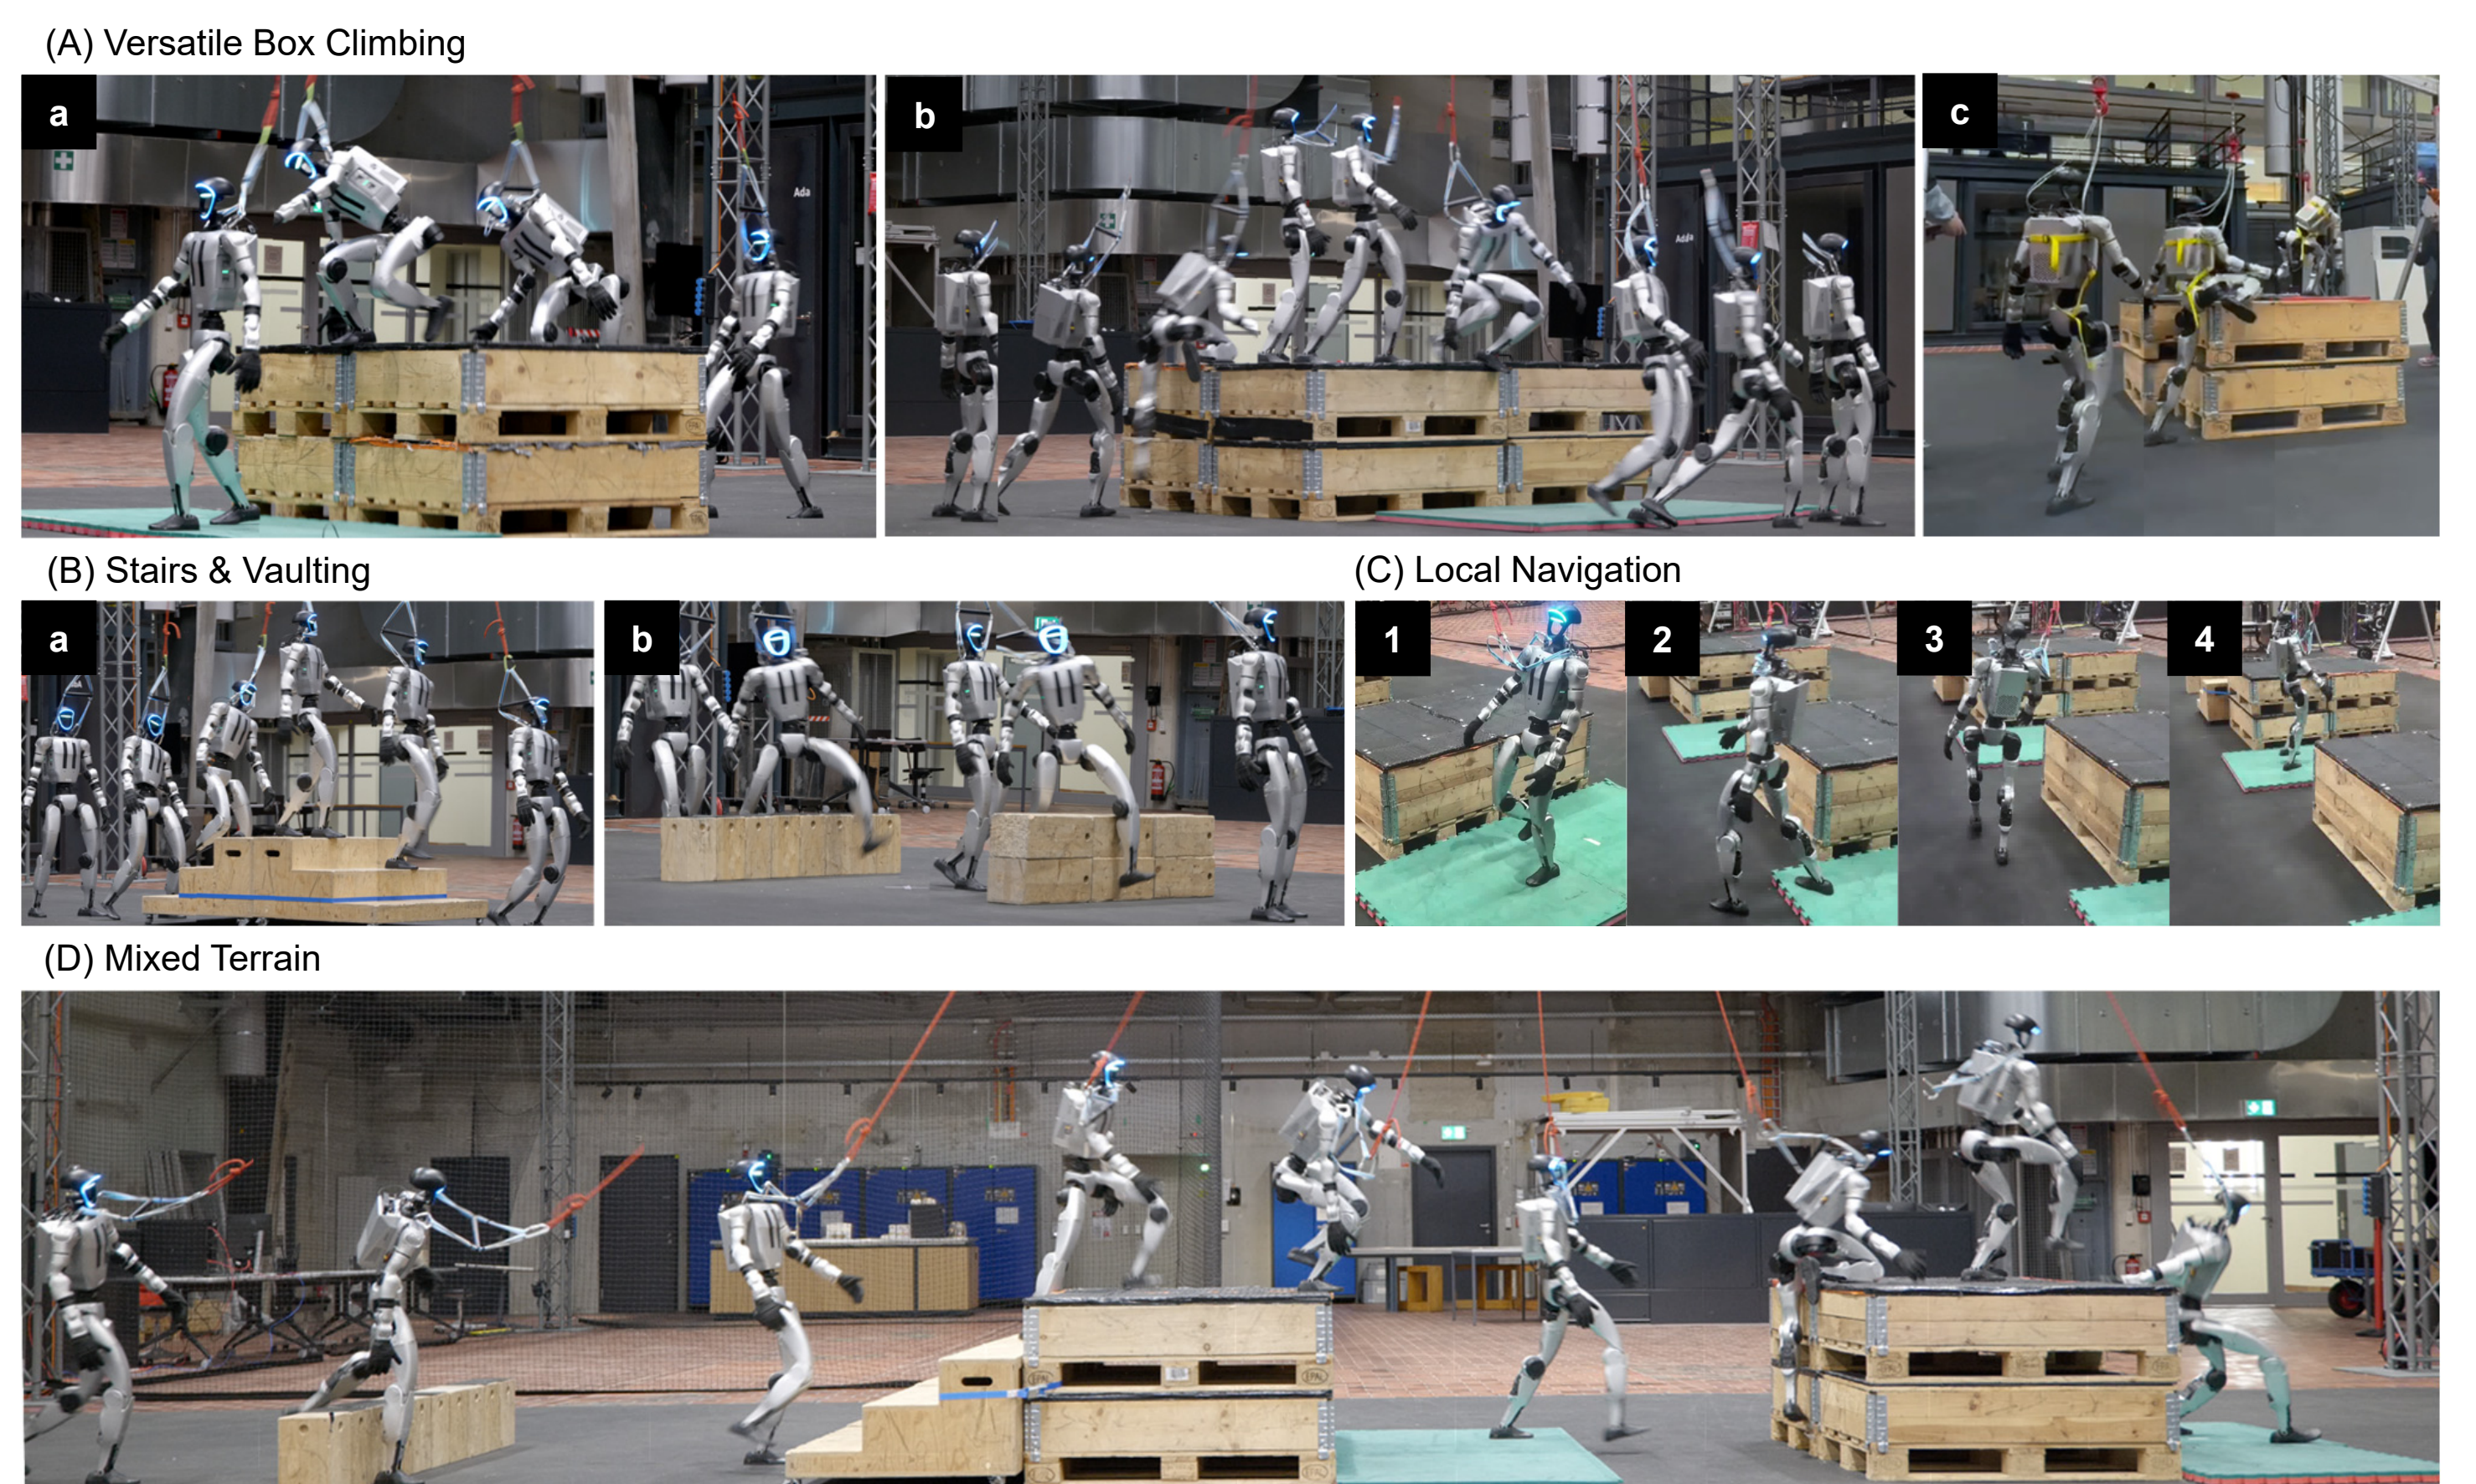


Paper



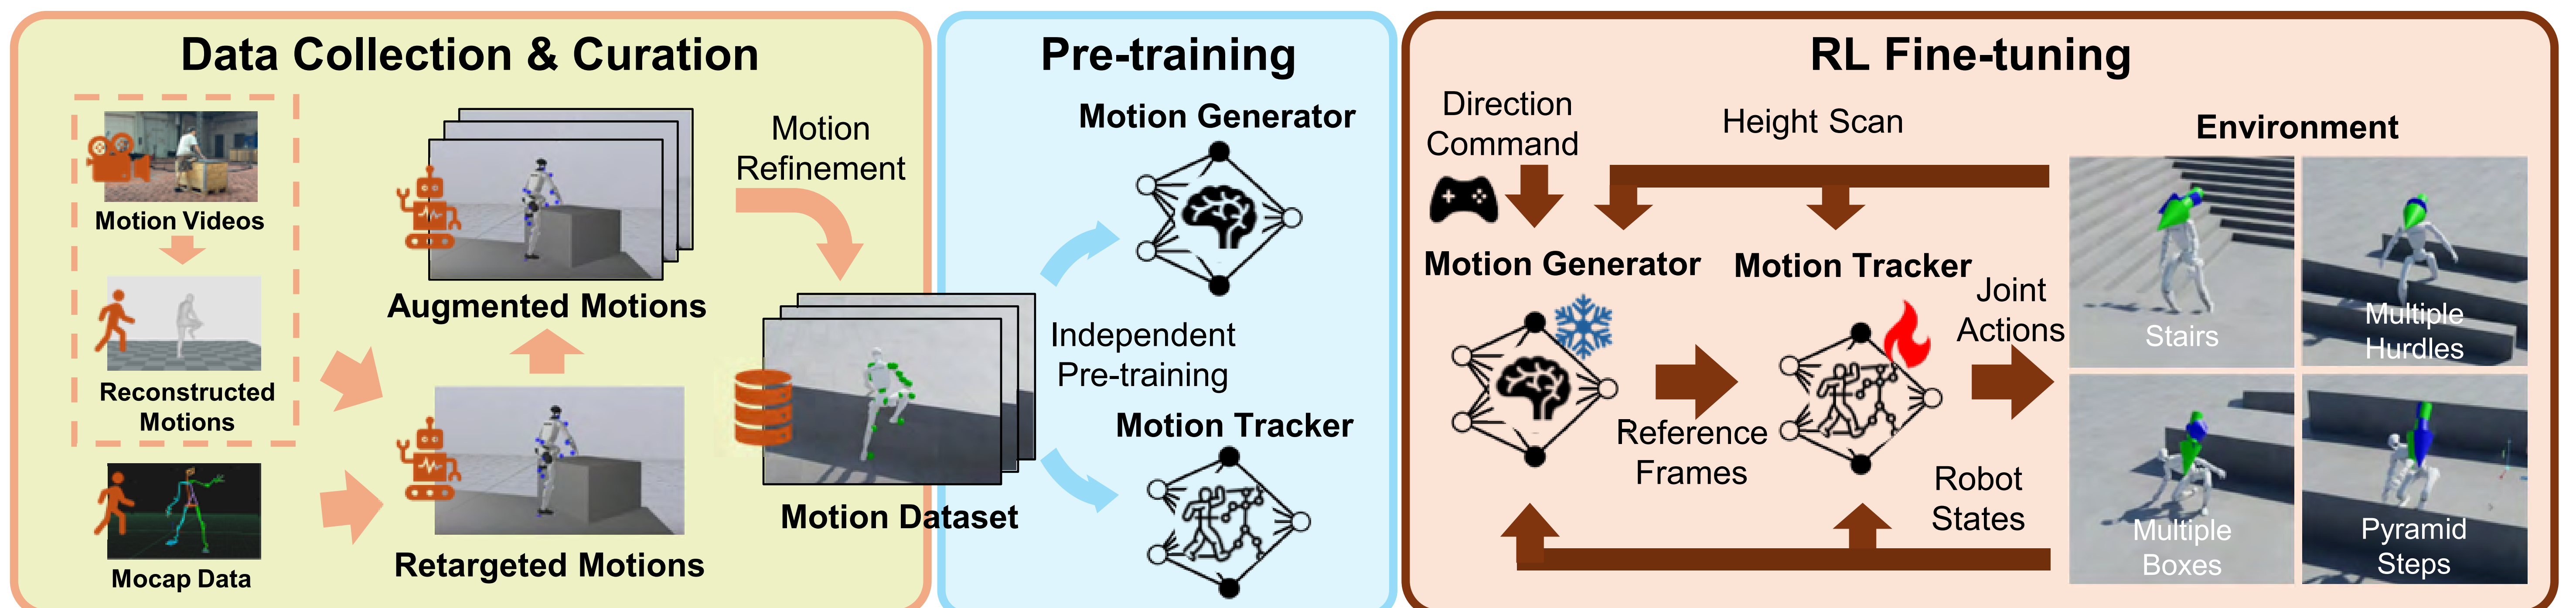
Video

Hardware Experiments — Unitree G1, fully onboard



(A) 75 cm box climbing with three jump-down styles. (B) Stairs + vaulting hurdles. (C) Local navigation: bypassing the box to reach the goal. (D) Mixed terrain: vaulting + stairs + box climbing.

Method: Three-Stage Generation-and-Tracking Framework



(a) **Data.** ~ 5 min human motions (GVHMR, AMASS, OmniRetarget) \rightarrow retarget \rightarrow kinematic augmentation $\rightarrow \sim 1$ h.

(b) **Pre-training.** Tracker (PPO/IsaacLab) + MDM-style diffusion generator (0.5 s / 25-frame reference from heading, scan, 2-frame history).


(c) **RL Fine-tuning.** Closed-loop fine-tuning with frozen generator (2 denoising steps); tracker acts as a **motion filter**.

Result 1 — Online motion generation enables generalization

Success rate \pm SD across terrain height/yaw (500 spawns/setting).

Task	Tracker Only	Tracker + Gen.
Box climbing	0.859 ± 0.252	0.987 ± 0.014
Vaulting	0.805 ± 0.231	0.990 ± 0.026
Stair ascent	0.845 ± 0.300	0.997 ± 0.005

Result 2 — Fine-tuning the tracker is essential




Fine-tuning **consistently improves success across all five tasks, largest gains on the hardest settings.**

Qualitative Hardware Behaviors

- **Versatile box traversal** with on-top reorientation; knees/hands support, hands cushion on jump-down.
- **Continuous vaulting** over hurdles of varying heights, crossing directly rather than stepping onto them.
- **Stair** ascent and descent with terrain-aware foot placement.
- **Mixed-terrain transitions** unseen in training, e.g. stair ascent following a jump-down.
- **Local-navigation behavior:** tracker partially overrides the reference to *bypass* obstacles and still reach the goal.

Onboard Terrain Perception — LiDAR Elevation Mapping



Livox MID360 LiDAR + DLIO pose + Elevation Mapping CuPy \rightarrow real-time height-map reconstruction (top) of the real environment (bottom).