Building General Humanoid Robots in the Physical World

Baoxiong Jia
BIGAI

Figure generated by Nano Banana Pro

Mobile-base policies work like magic

ACT-1, Sunday 2026

GENE-26.5, Genesis 2026

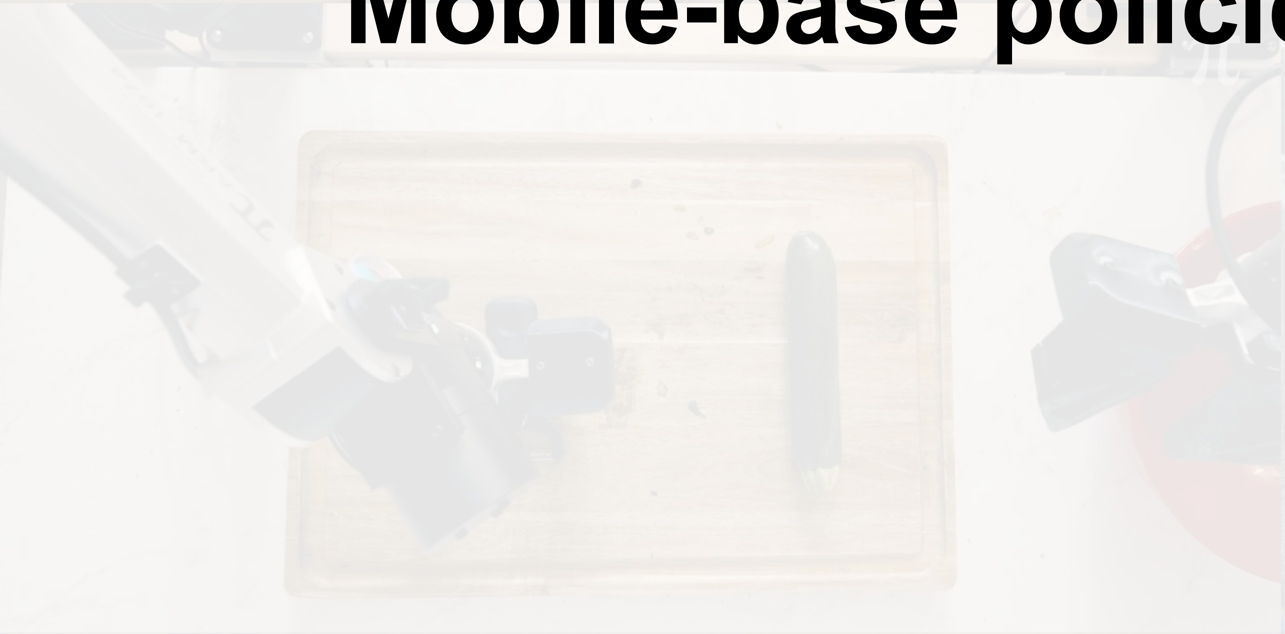

π-0.7, Physical Intelligence 2026


GEN-1, Generalist 2026

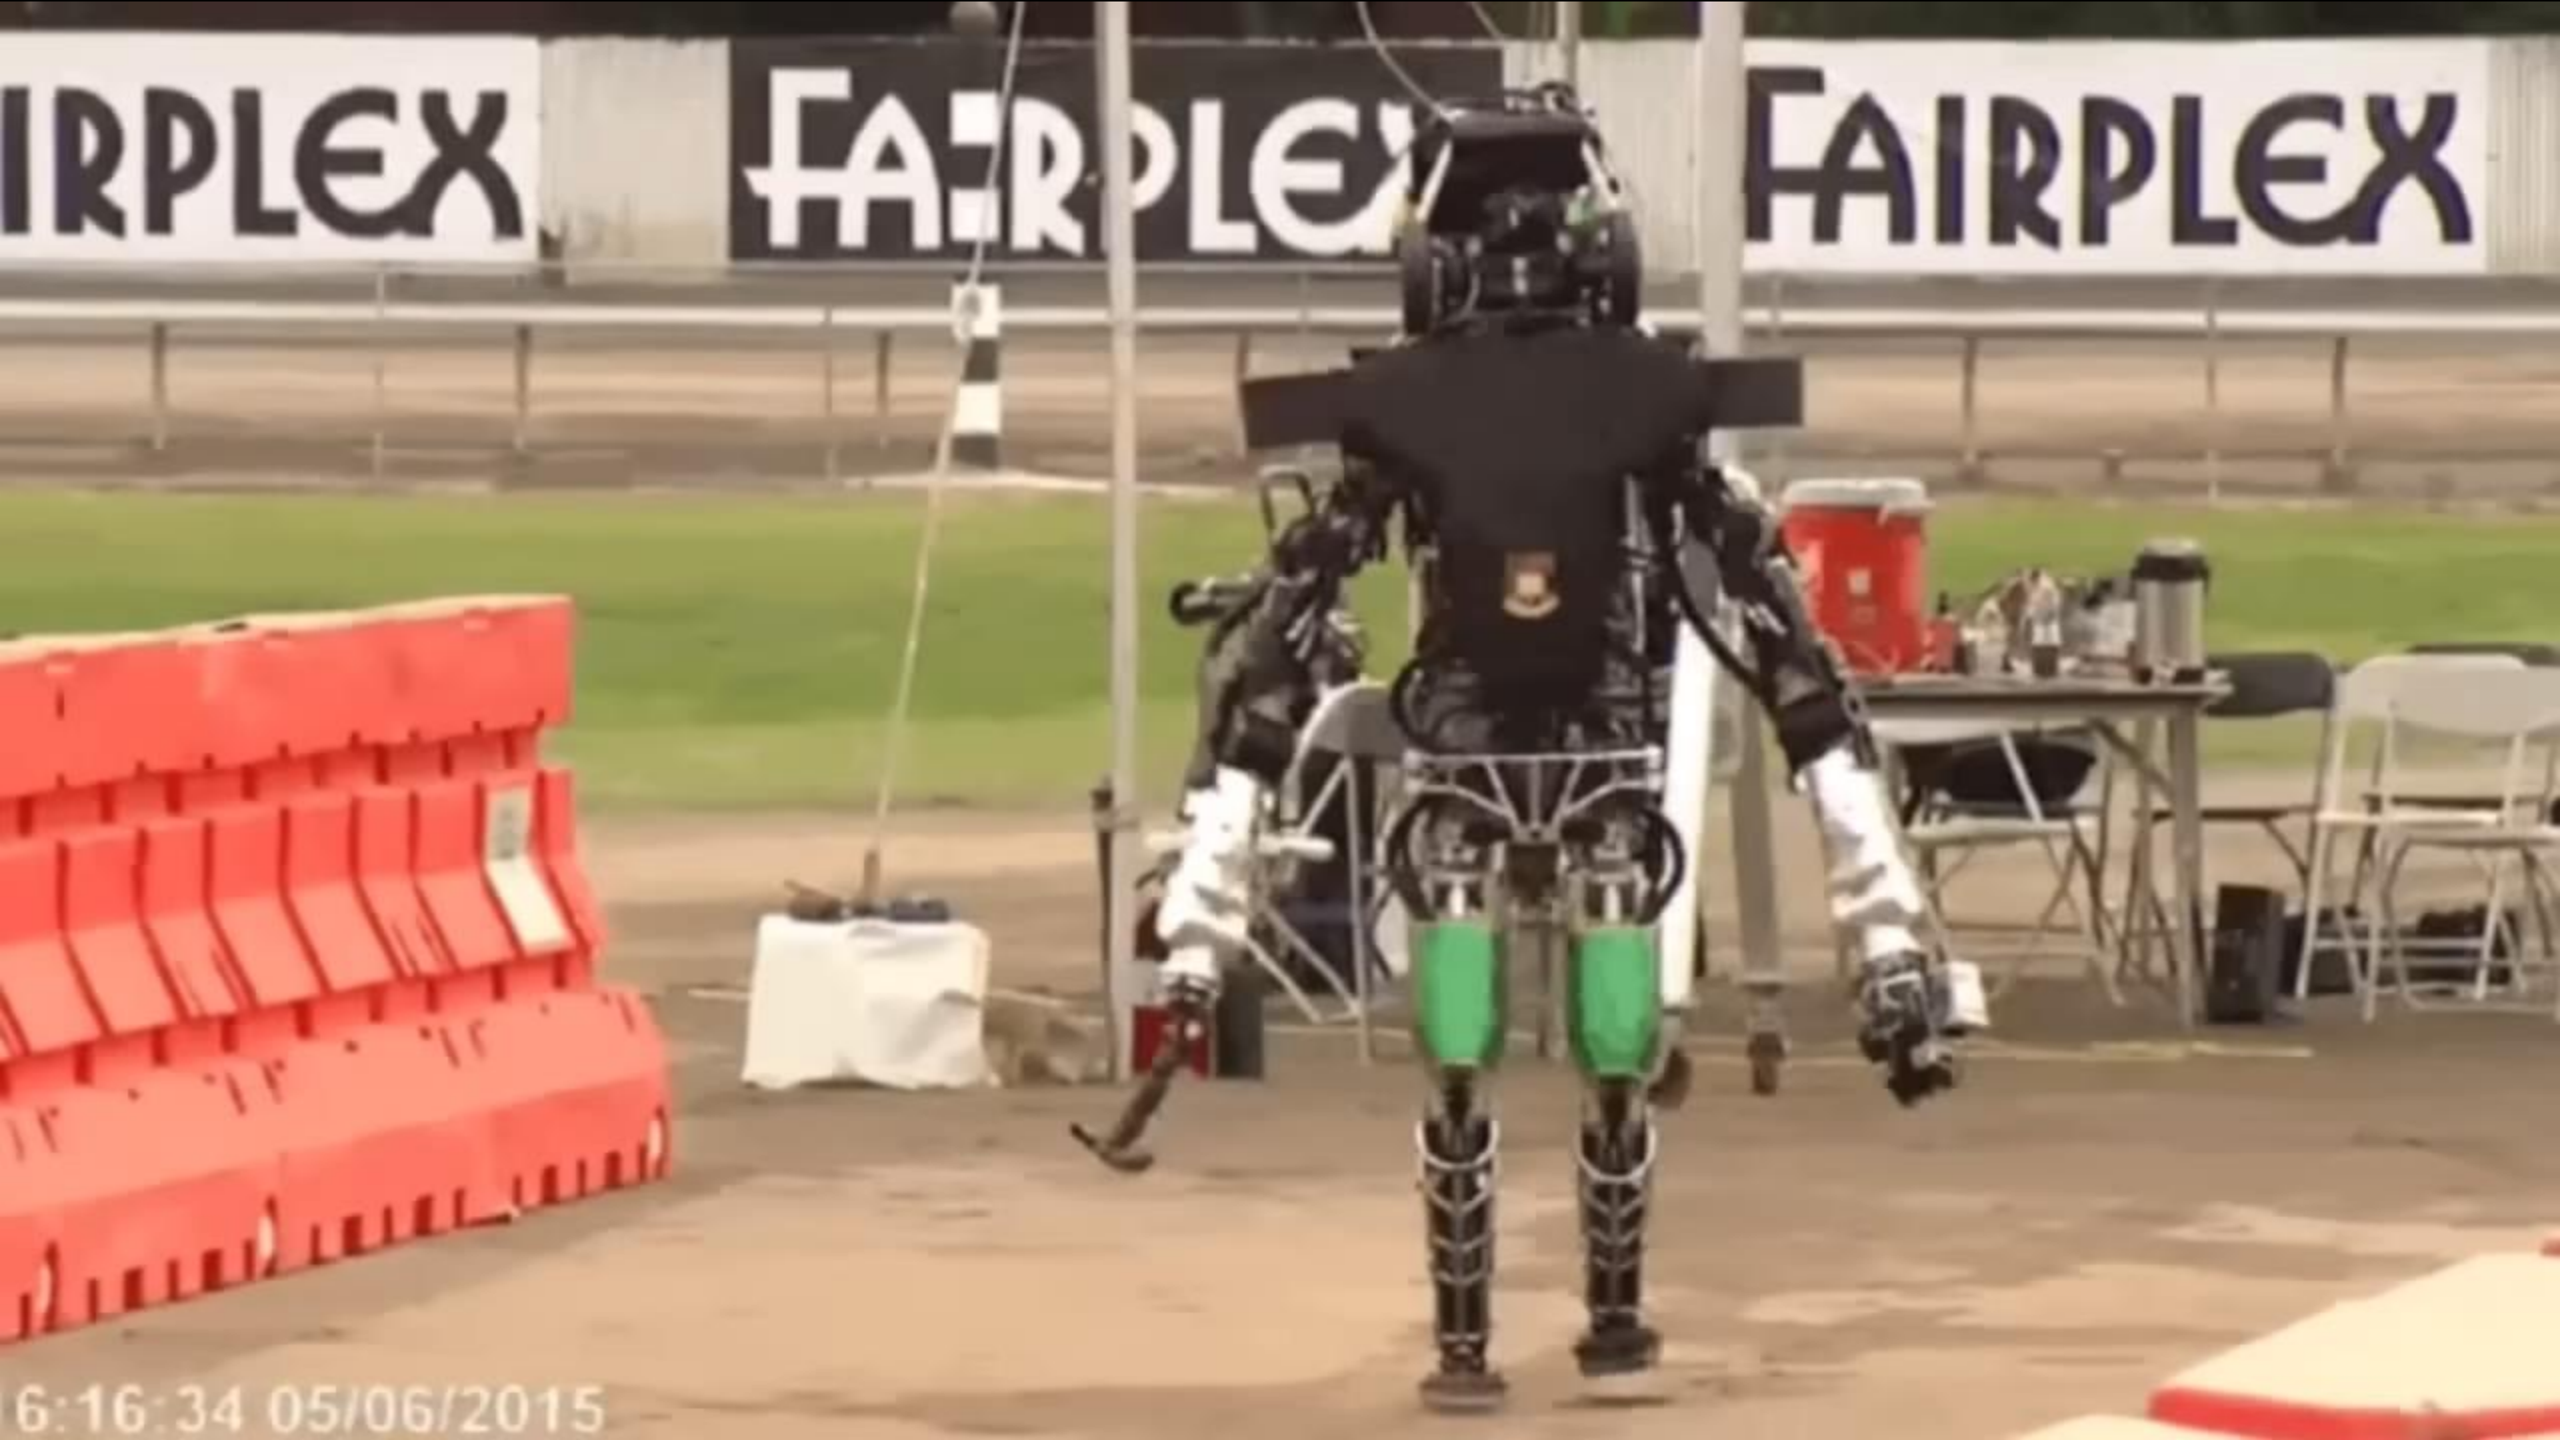

IRPLEX
FAIRPLEX
FAIRPLEX
6:16:34 05/06/2015

dns_di @Denis_13_1982 · May 8

Why not give them wheels instead of feet, they will be much faster and more energy efficient?

UniTree Kungfu Kid 6.0, UniTree 2025

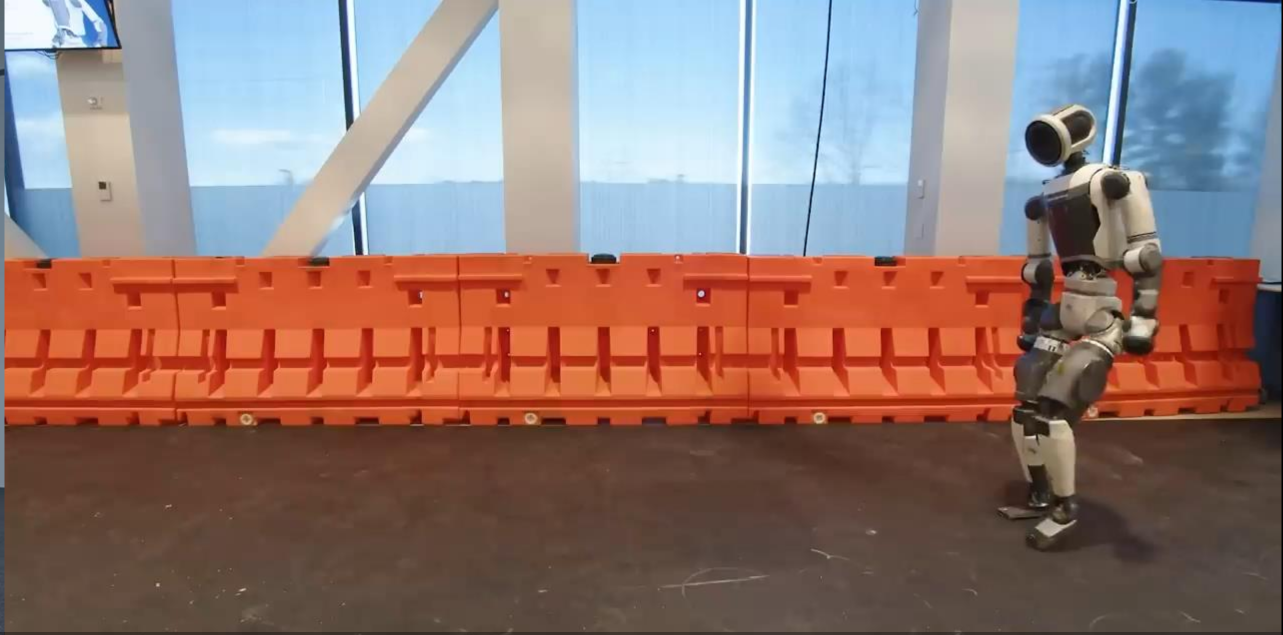

Walk, Run, Crawl, RL Fun, Boston Dynamics 2024

HONOR Robotics D1, HONOR 2026

Booster Robotics’ New Striker, Booster 2026

Why we don’t see that many “working” humanoids

Specific Scenes & Tasks

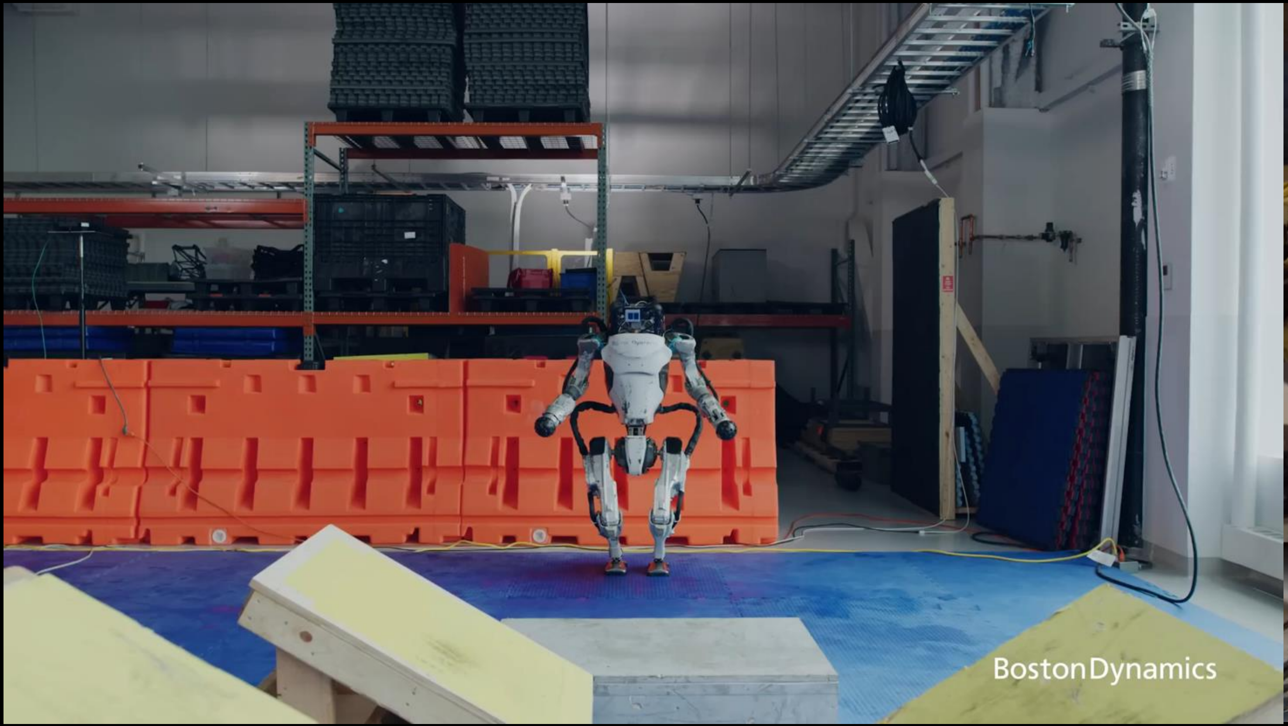

Boston Dynamics, Atlas | Partners in Parkour, 2022
https://www.youtube.com/watch?v=tF4DML7FIWk

General Tasks & Intearctions

Damen et al., Scaling Egocentric Vision: The Epic-Kitchens Dataset, 2018

Language

1.2B Hours

What’s the reason?

Manipulation

0.5M Hours

Whole-Body

1K? Hours

autonomous

Human motion for humanoid whole-body control

MoCap

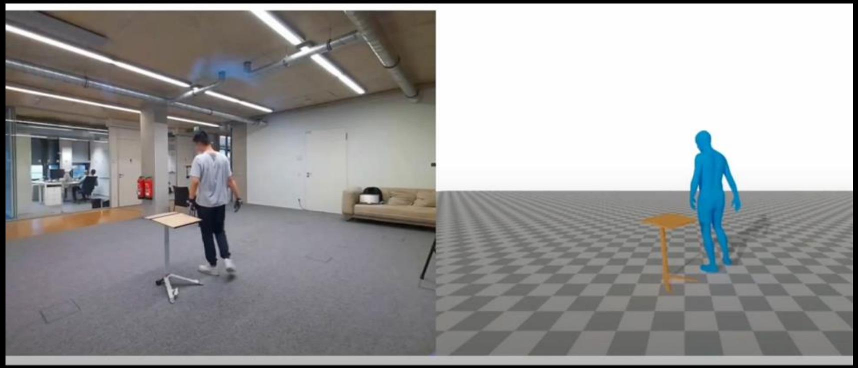

Motion Generation

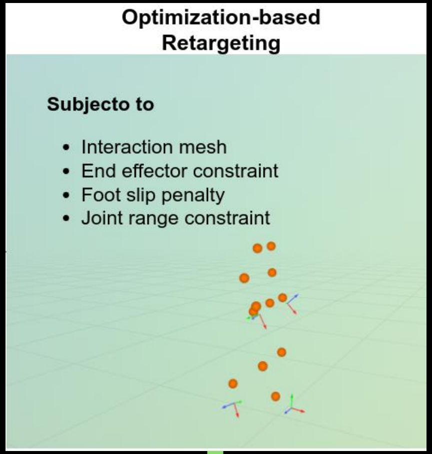

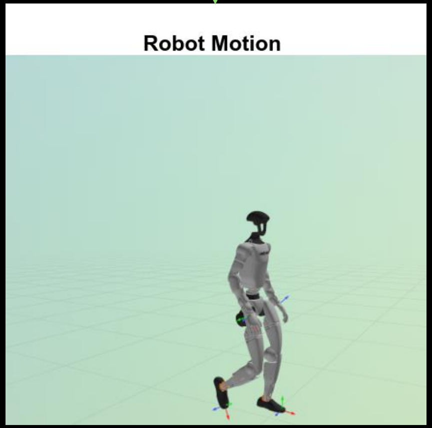

Humanoid Controller

MLP
MoE
Transformer
PID
MPC
…

Human Motion → ***Motion Retargeting*** → ***Controller*** → ***Robot Execution***

CLONE: Humanoid Whole-Body Teleoperation

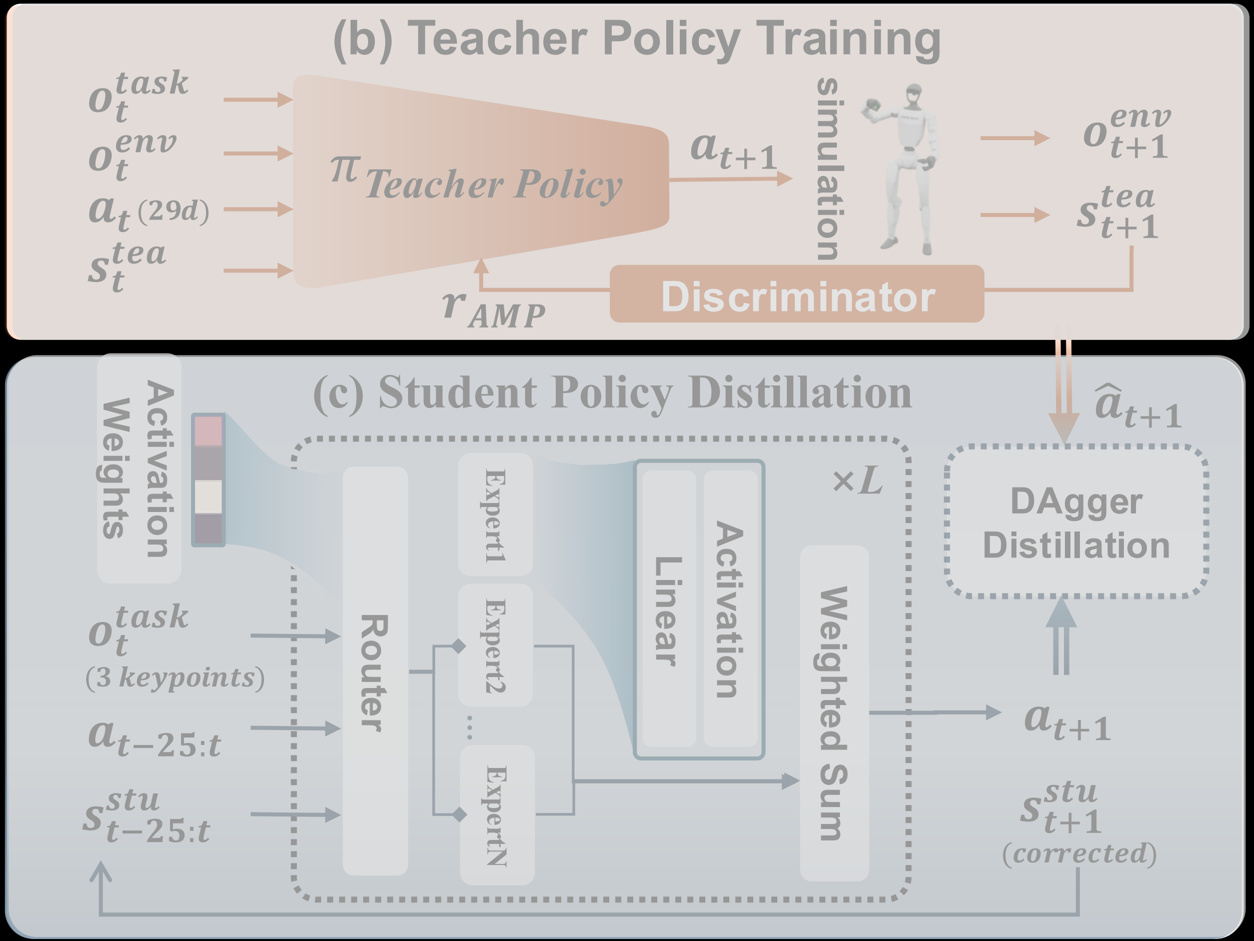

Learning a teacher policy with privileged information for human motion tracking

Distilling a MoE-based student policy with Behavior Cloning (Dagger)

Li et al., CLONE: Closed-Loop Whole-Body Humanoid Teleoperation for Long-Horizon, CoRL 2025

Enabling Humanoid Training with Diverse Data

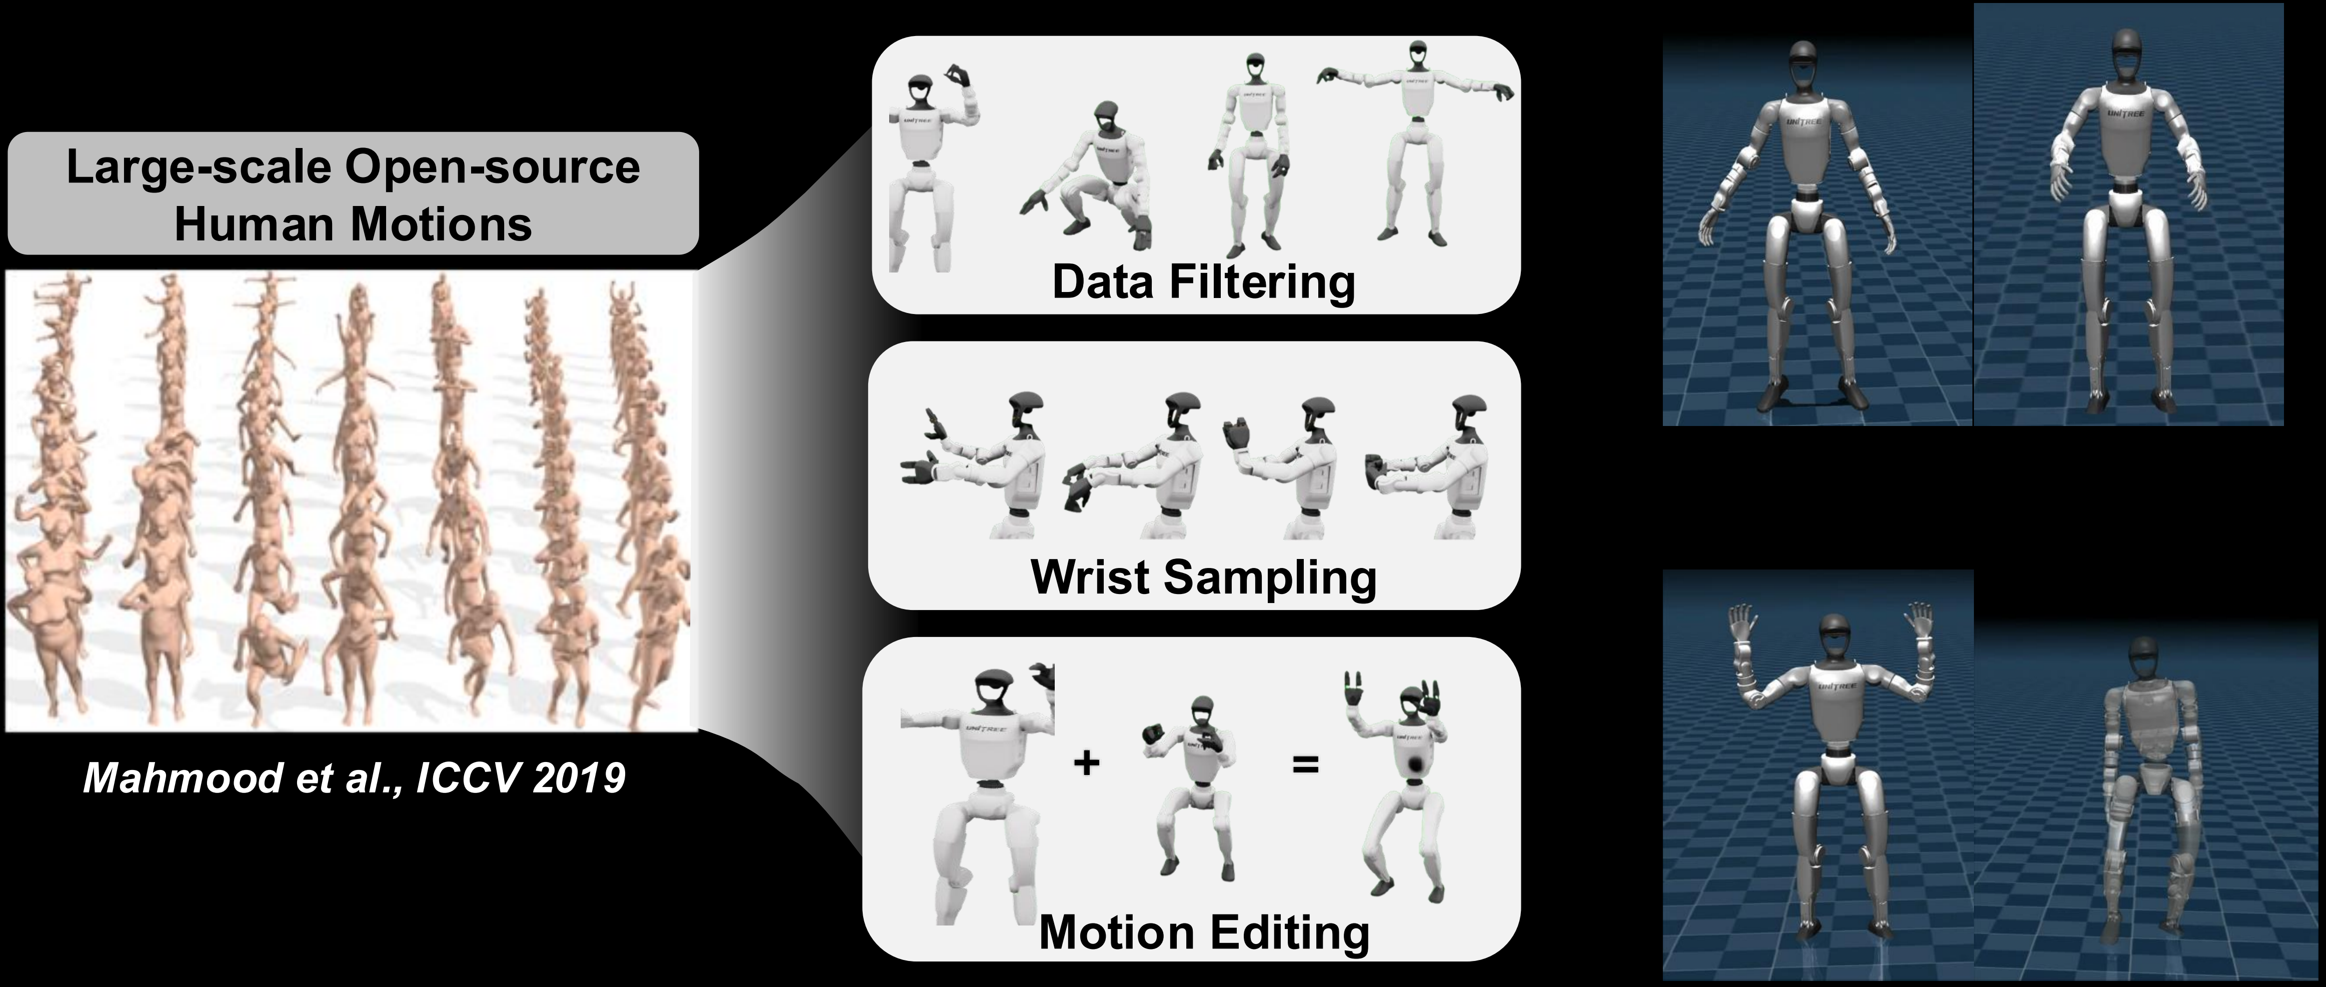

Mahmood et al., ICCV 2019

Augmenting existing motion datasets for wide pose coverage


Tracking for 15 meters and returning to the starting position (On an outdoor balcony)
Long-horizon Control in Outdoor Environments 1.5 x speed

BIGAI@World Humanoid Robot Game (WHRG) Aug. 2025

OmniBench for evaluating skill imbalances

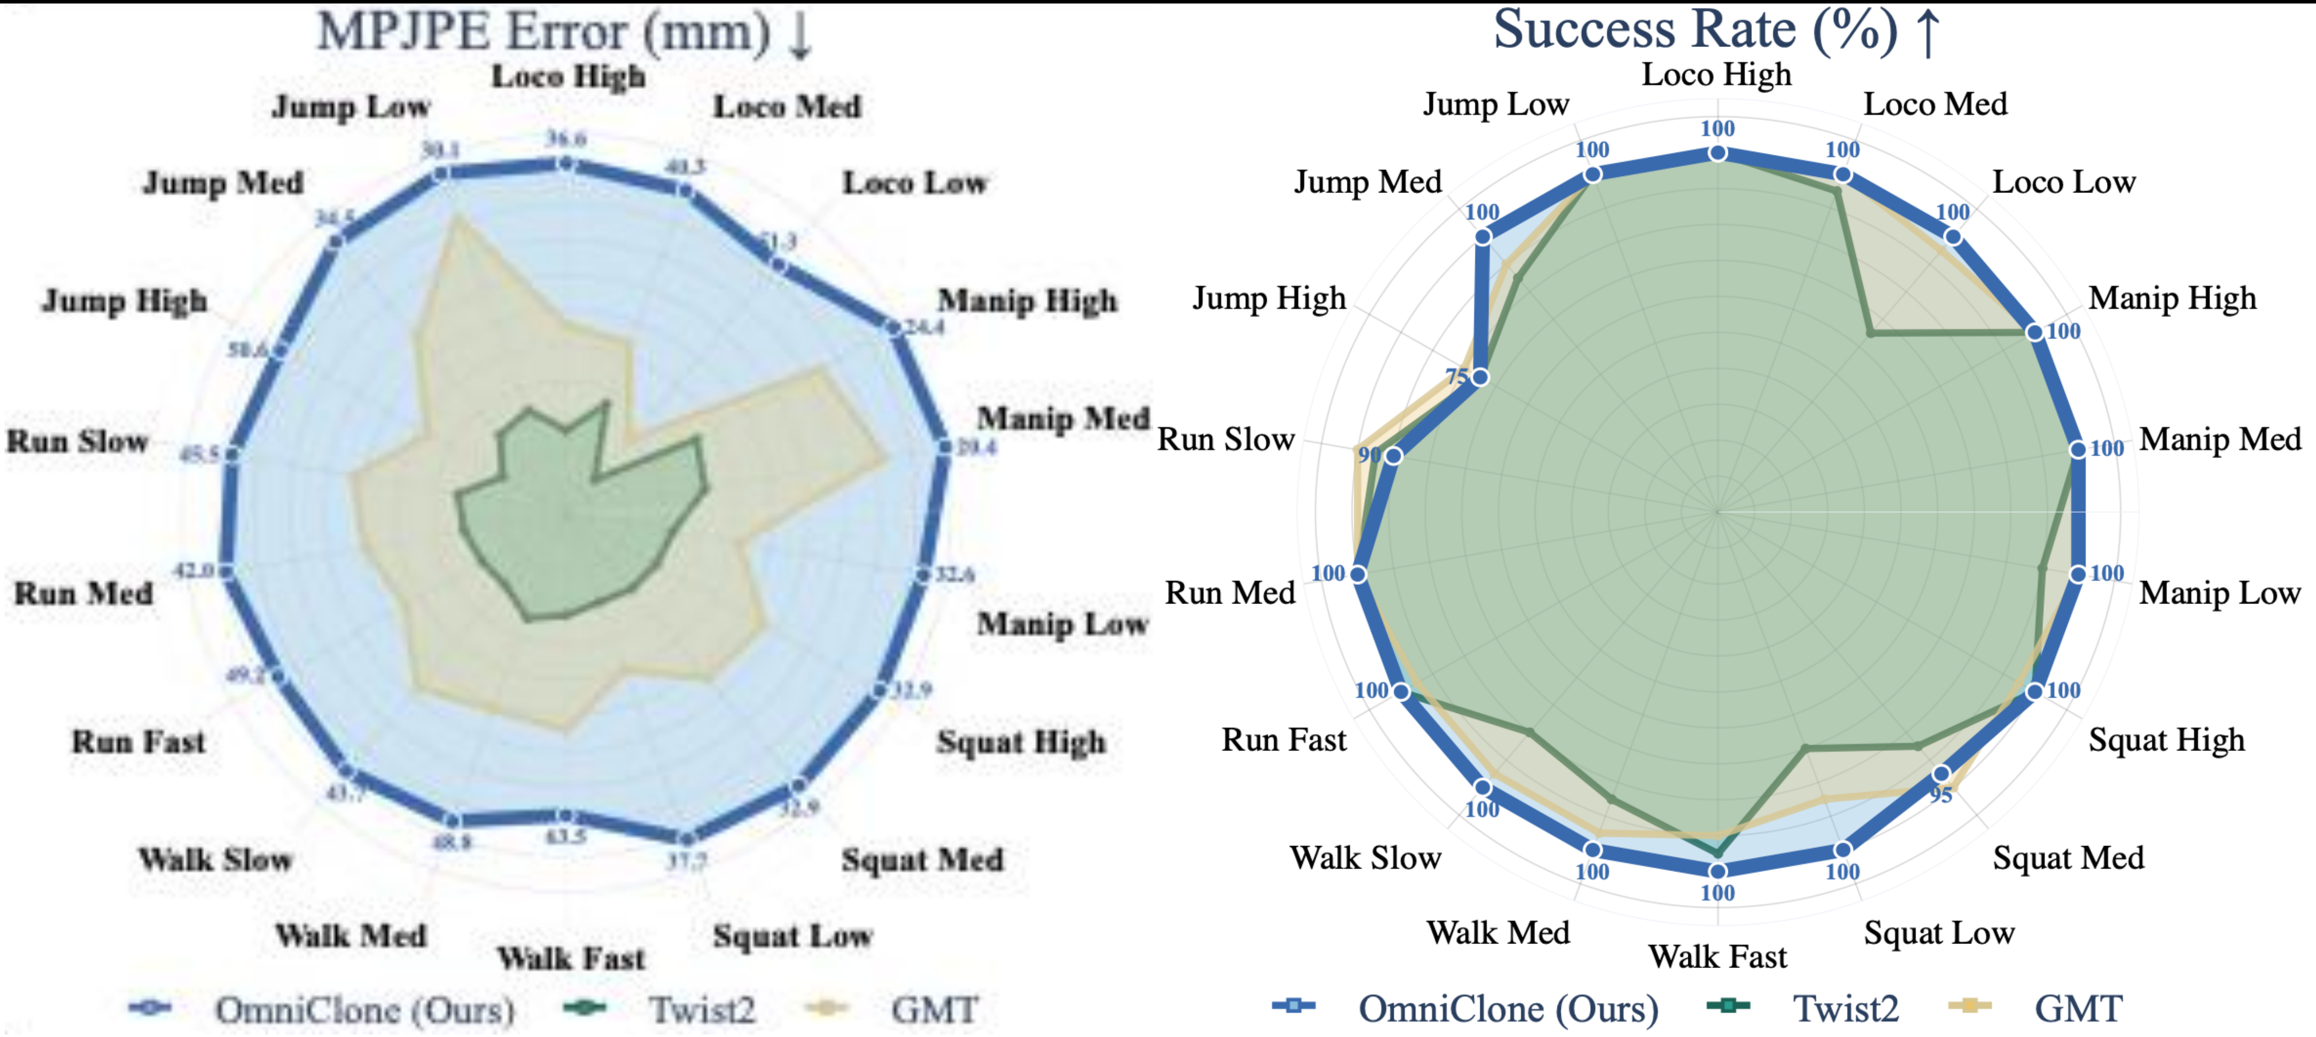

Li et al., OmniCLONE: Engineering a Robust, All-Rounder Whole-Body Humanoid Teleoperation System, ArXiv’26

The effect of data mixing

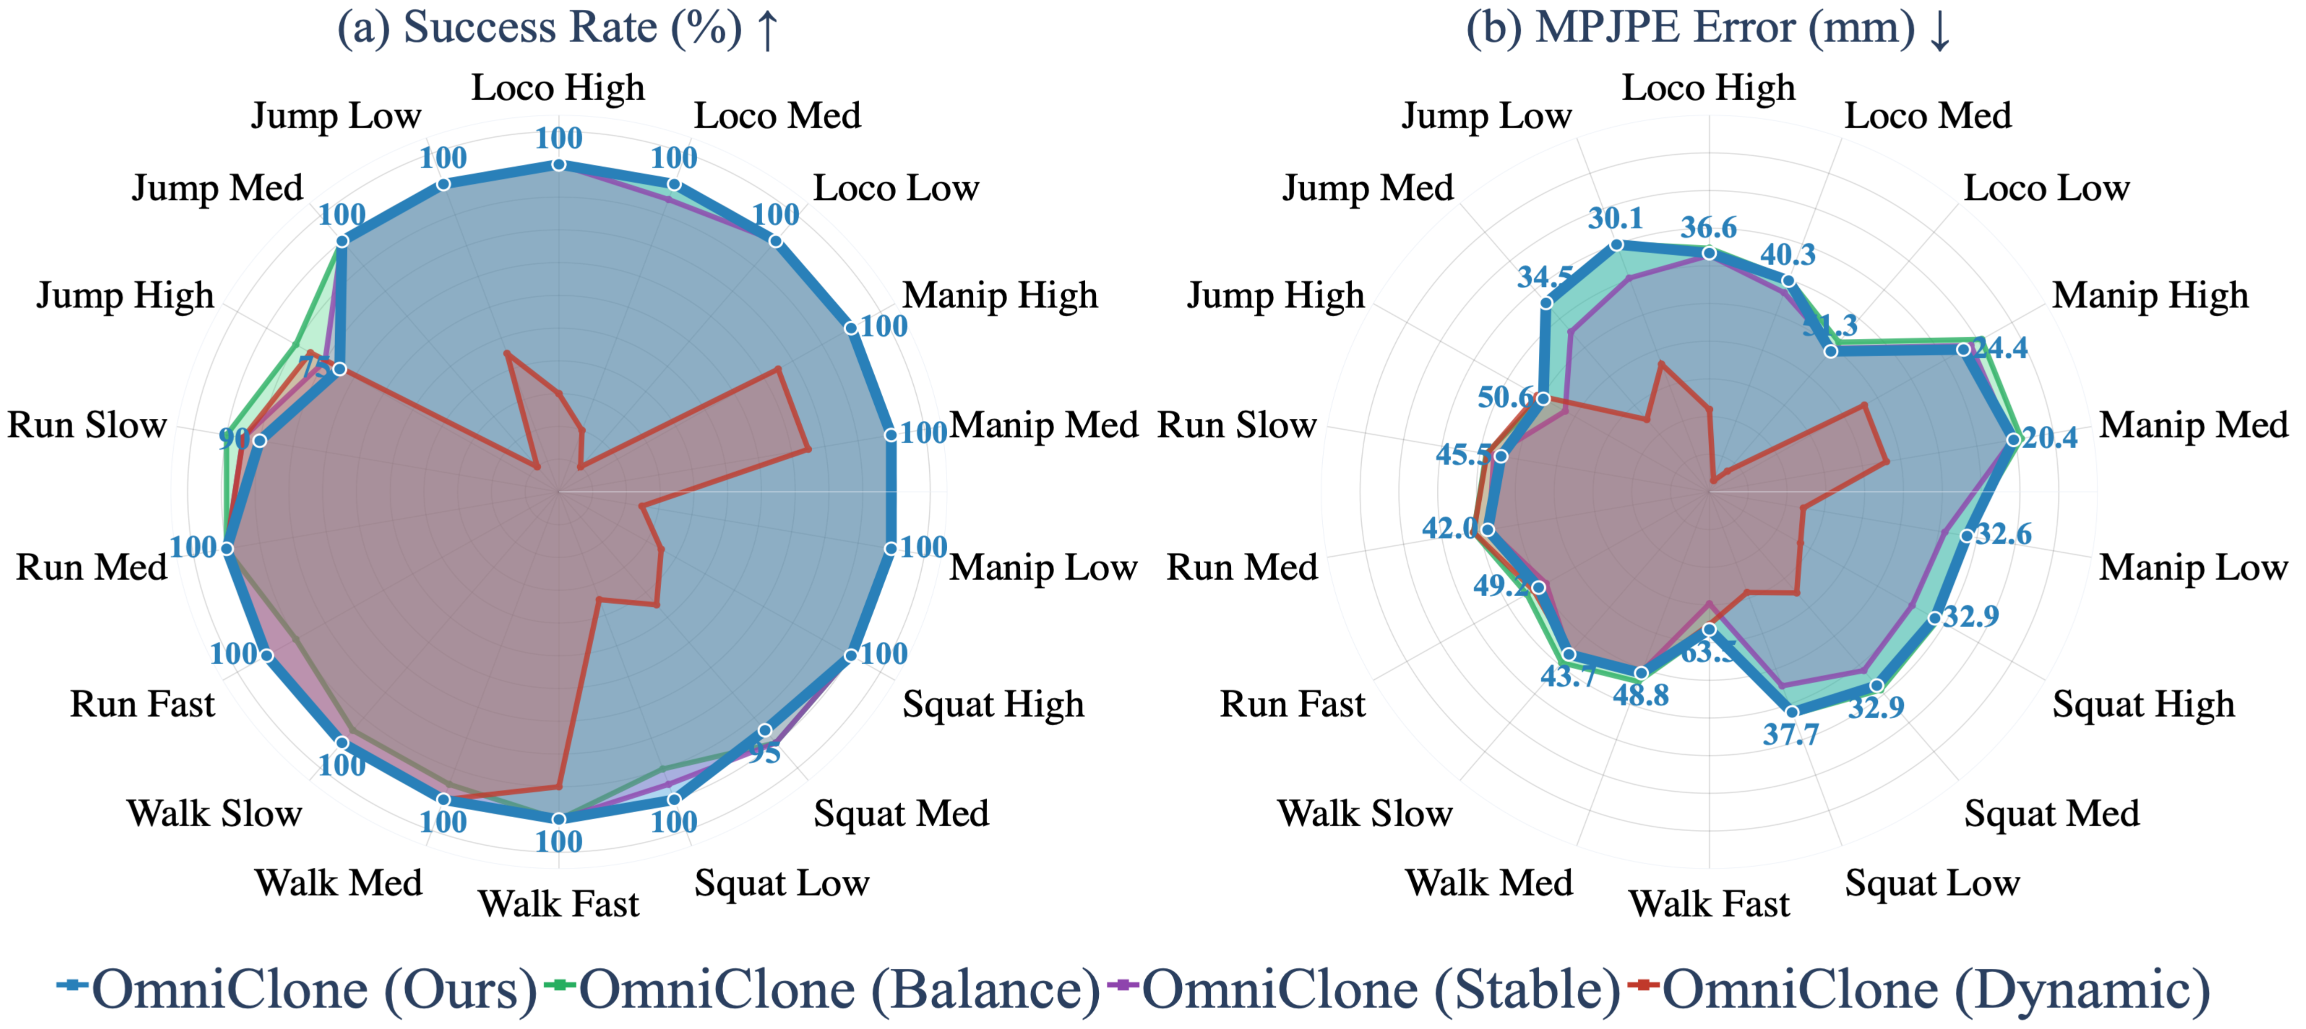

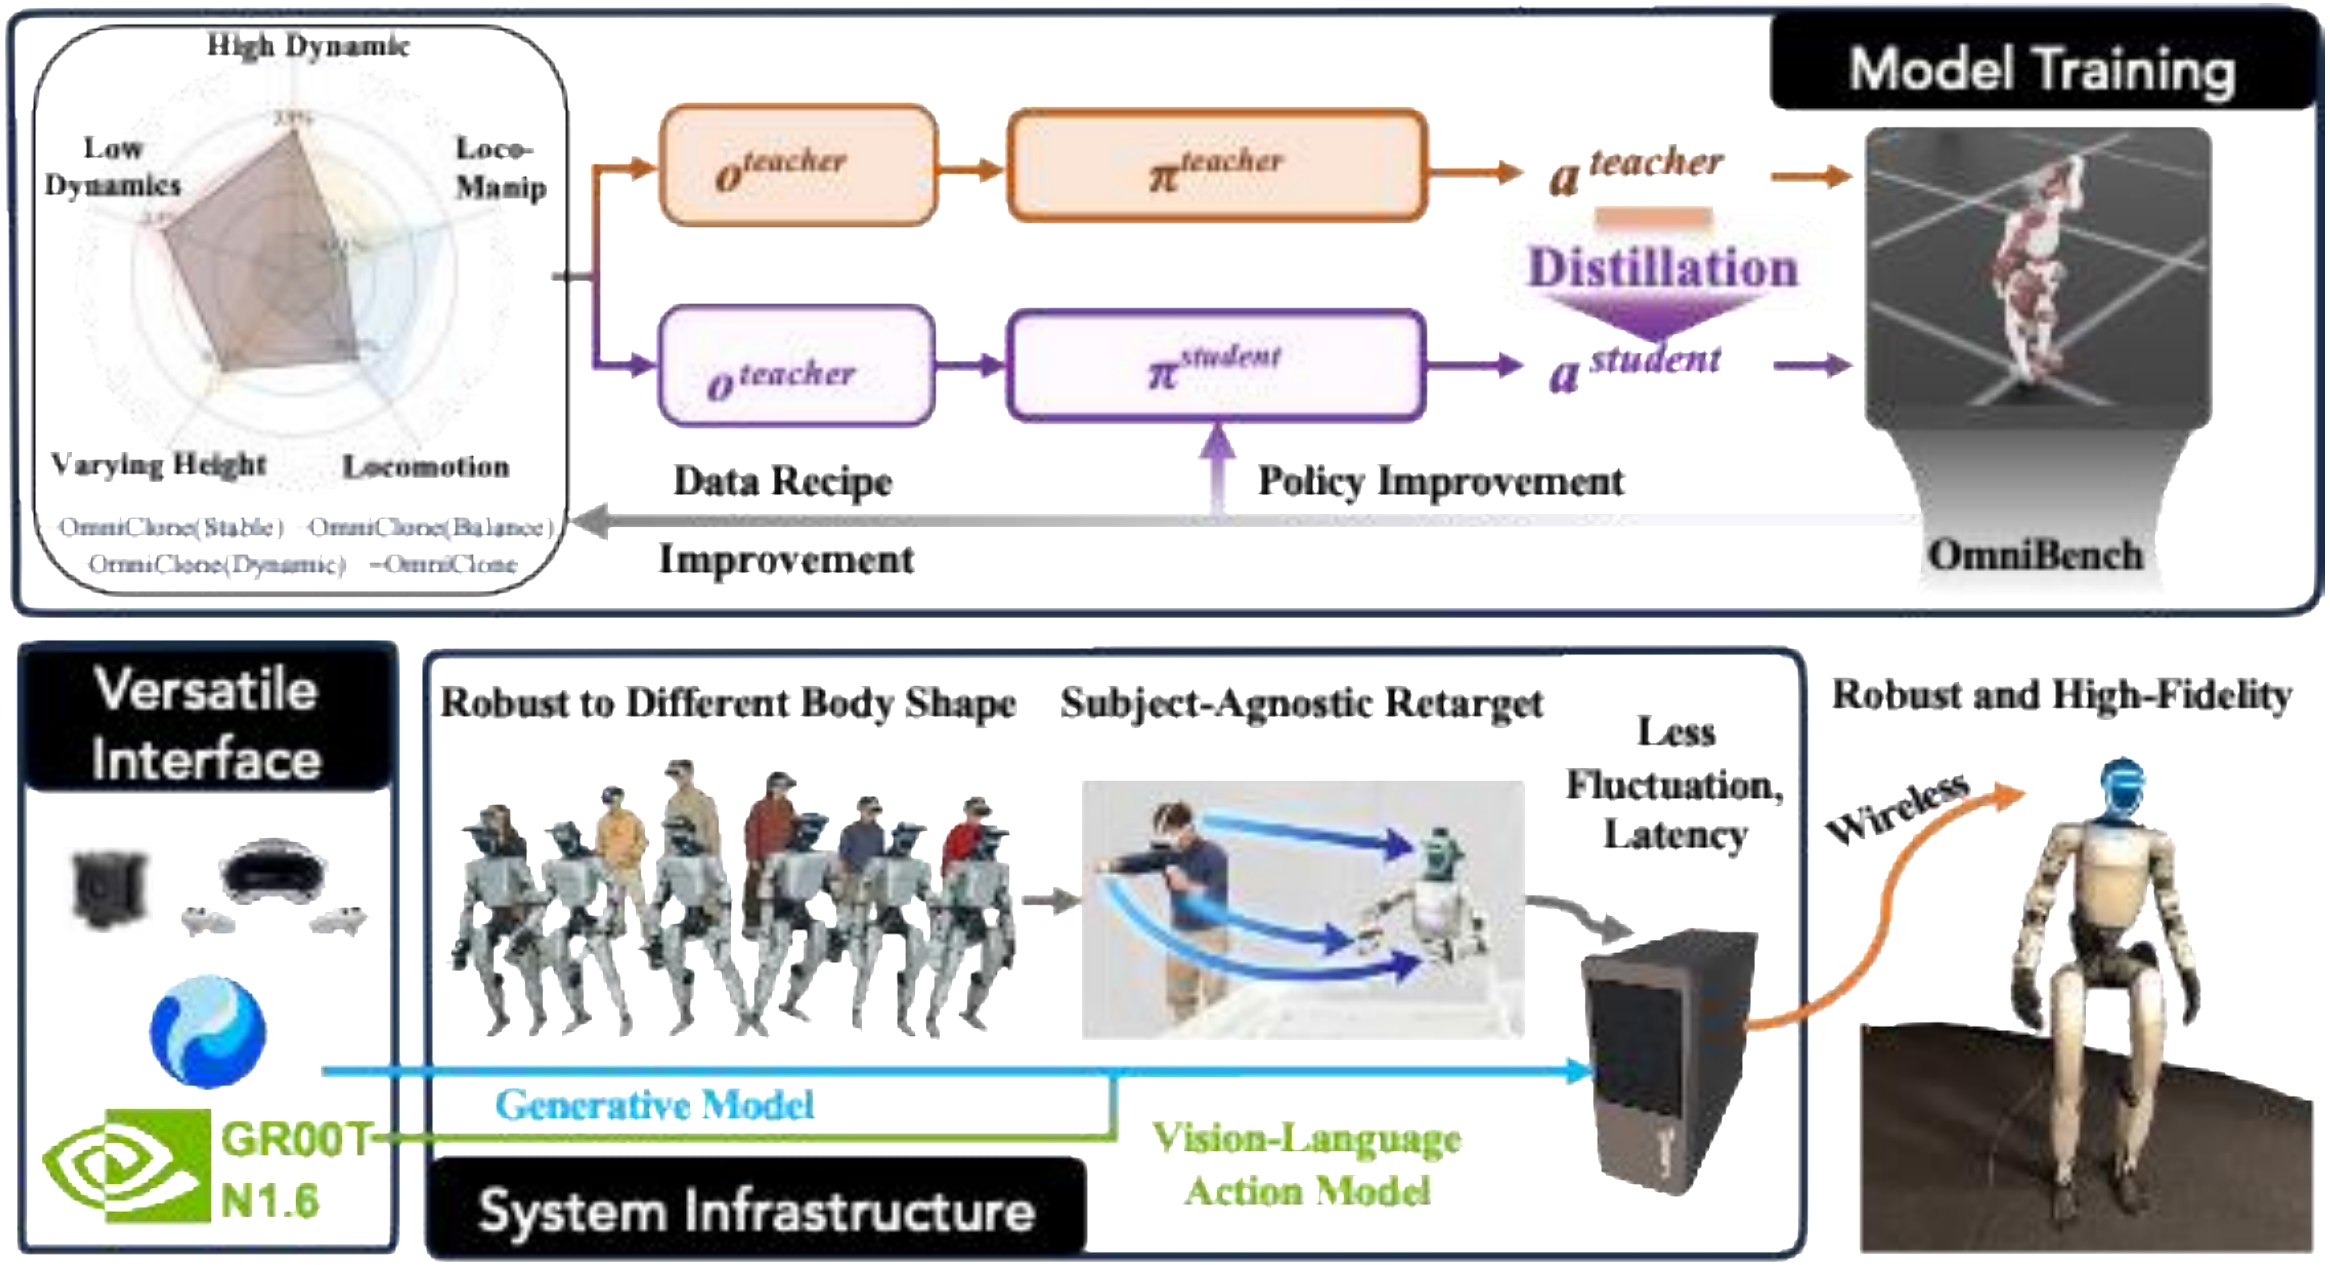
Model Training
High Dynamic
Low Dynamics
Loco-Manip
Varying Height
Locomotion
OmniClone(Stable)
OmniClone(Balance)
OmniClone(Dynamic)
OmniClone
$o^{teacher}$
$\pi^{teacher}$
$a^{teacher}$
Distillation
$o^{teacher}$
$\pi^{student}$
$a^{student}$
Data Recipe
Policy Improvement
Improvement
OmniBench
Versatile Interface
GR00T N1.6
Robust to Different Body Shape
Subject-Agnostic Retarget
Less Fluctuation, Latency
Robust and High-Fidelity
Wireless
Generative Model
Vision-Language Action Model
System Infrastructure

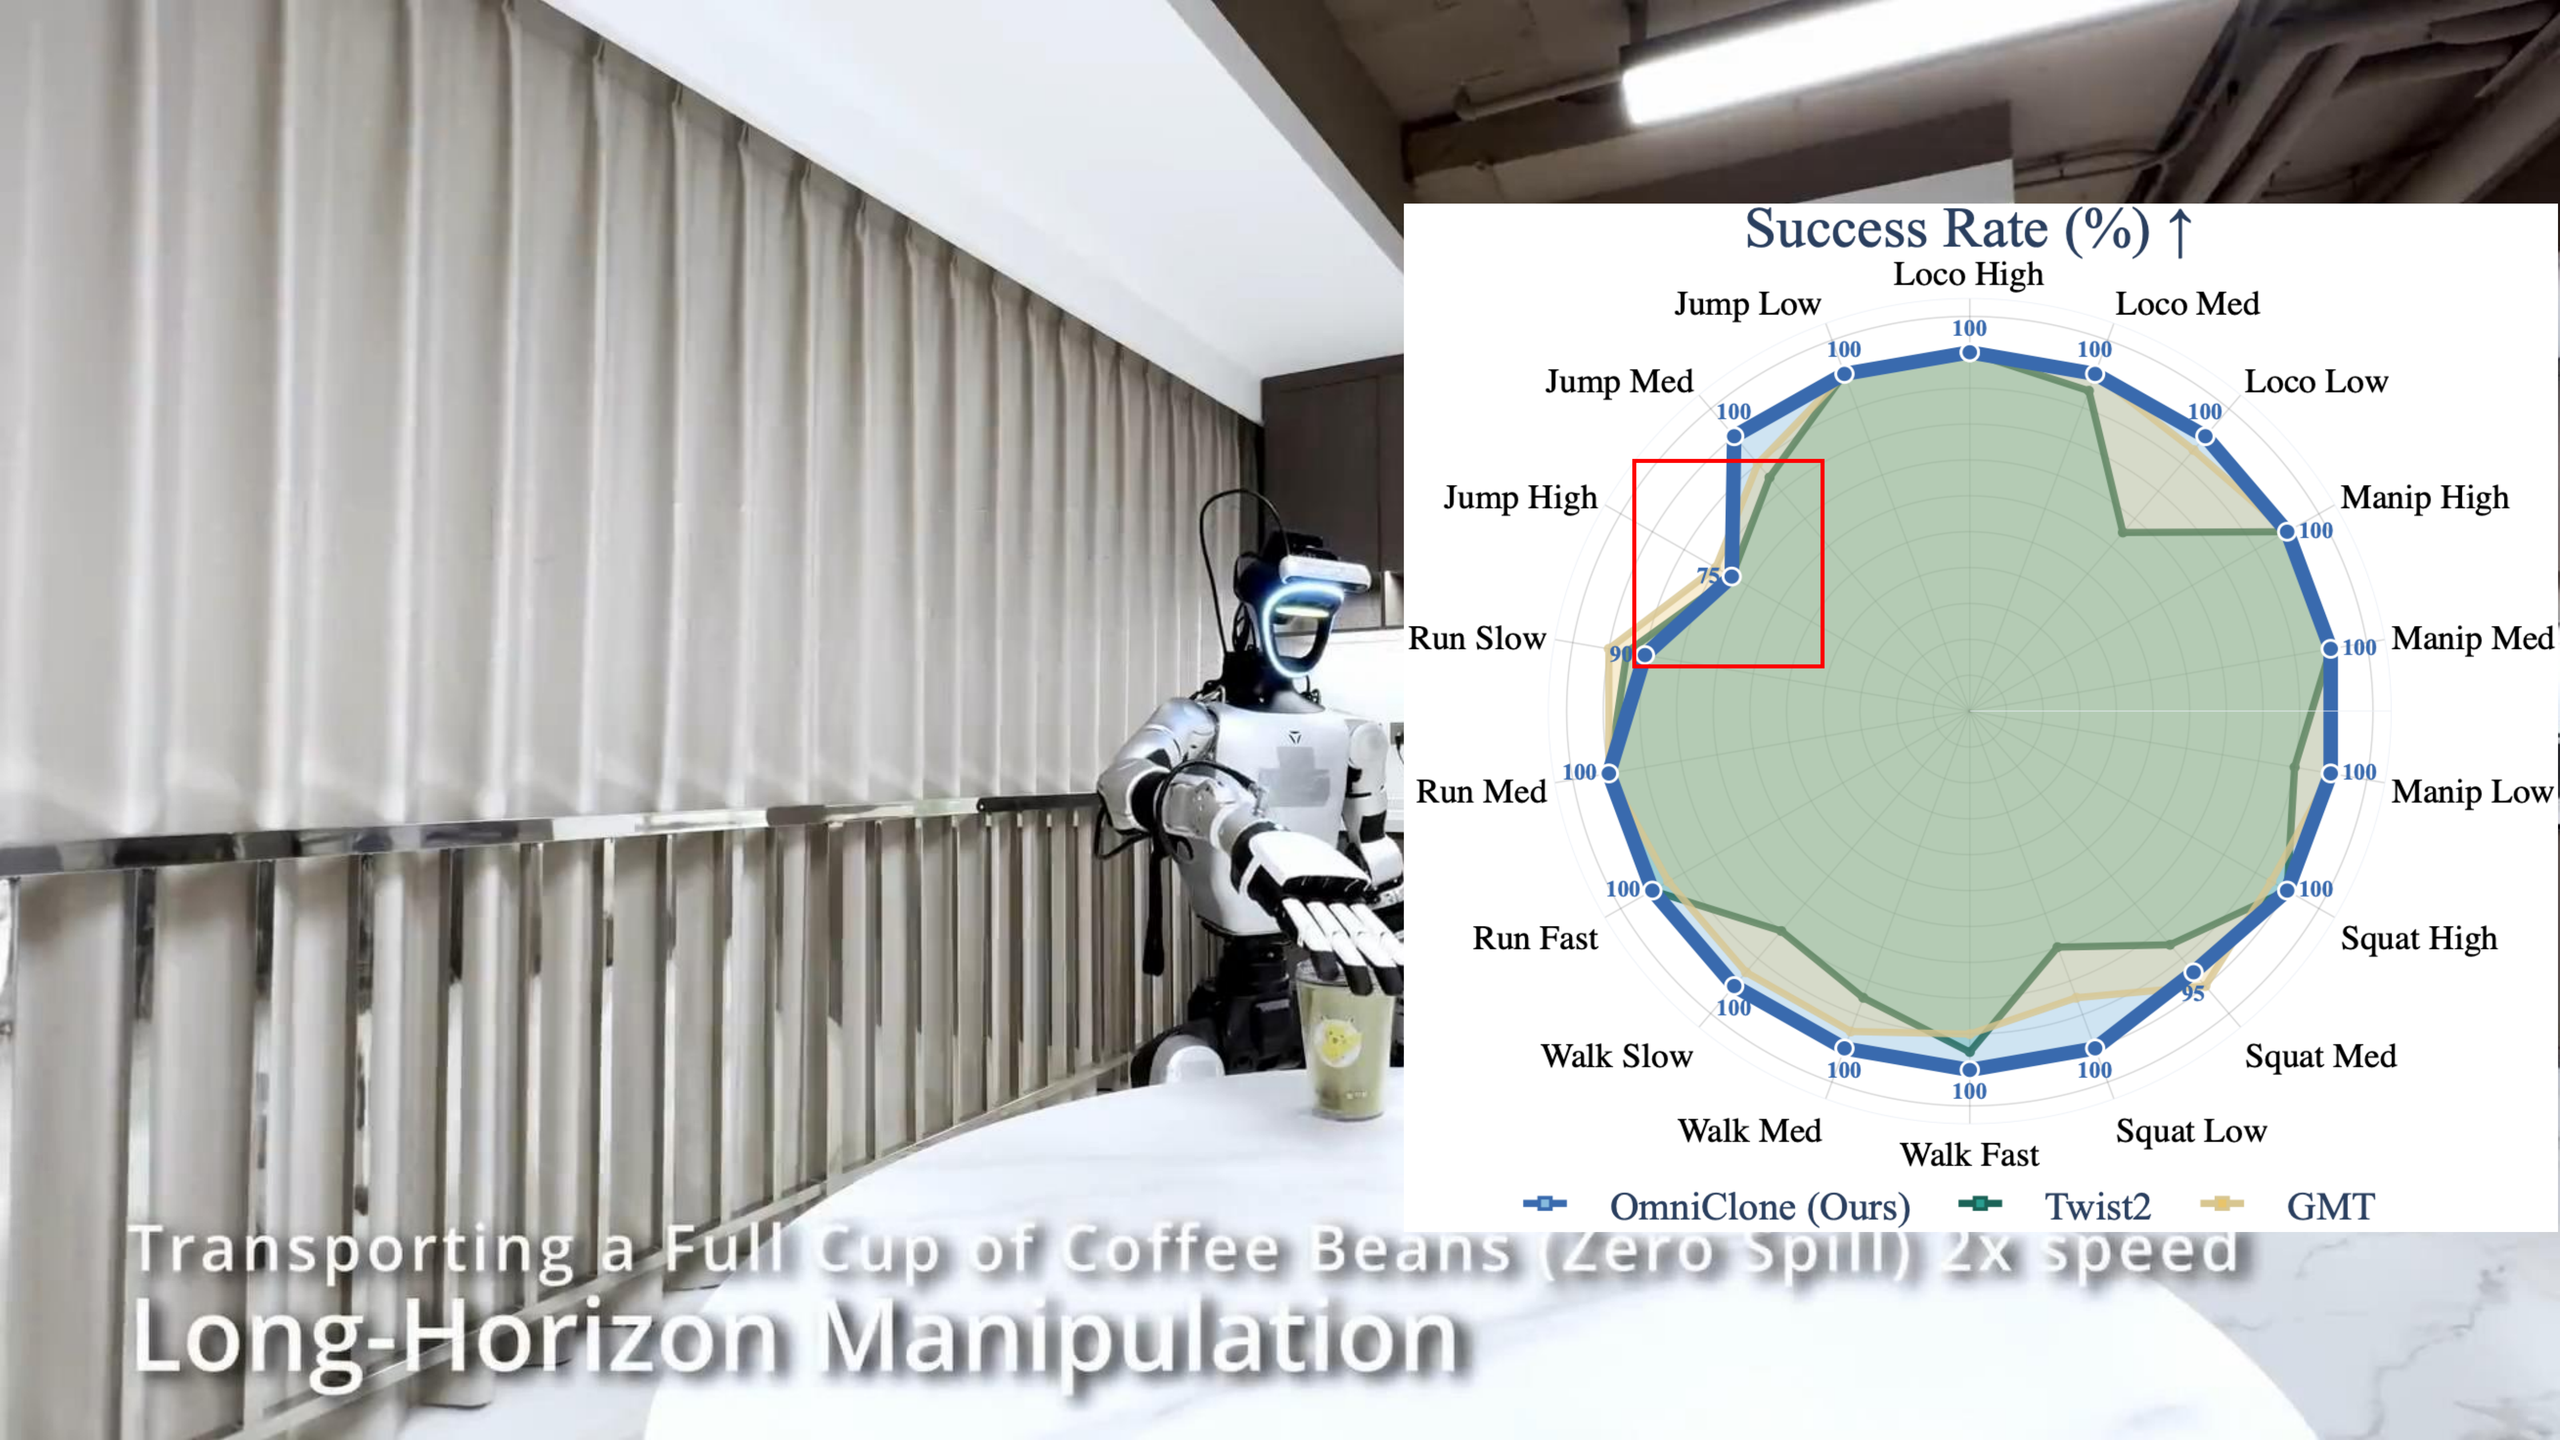
Success Rate (%) ↑
Loco High
Loco Med
Loco Low
Manip High
Manip Med
Manip Low
Squat High
Squat Med
Squat Low
Walk Fast
Walk Med
Walk Slow
Run Fast
Run Med
Run Slow
Jump High
Jump Med
Jump Low
100
100
100
100
100
100
100
95
100
100
100
100
100
100
90
75
100
100
OmniClone (Ours)
Twist2
GMT
Transporting a Full Cup of Coffee Beans (Zero Spill) 2x speed
Long-Horizon Manipulation

The problem with reference motion

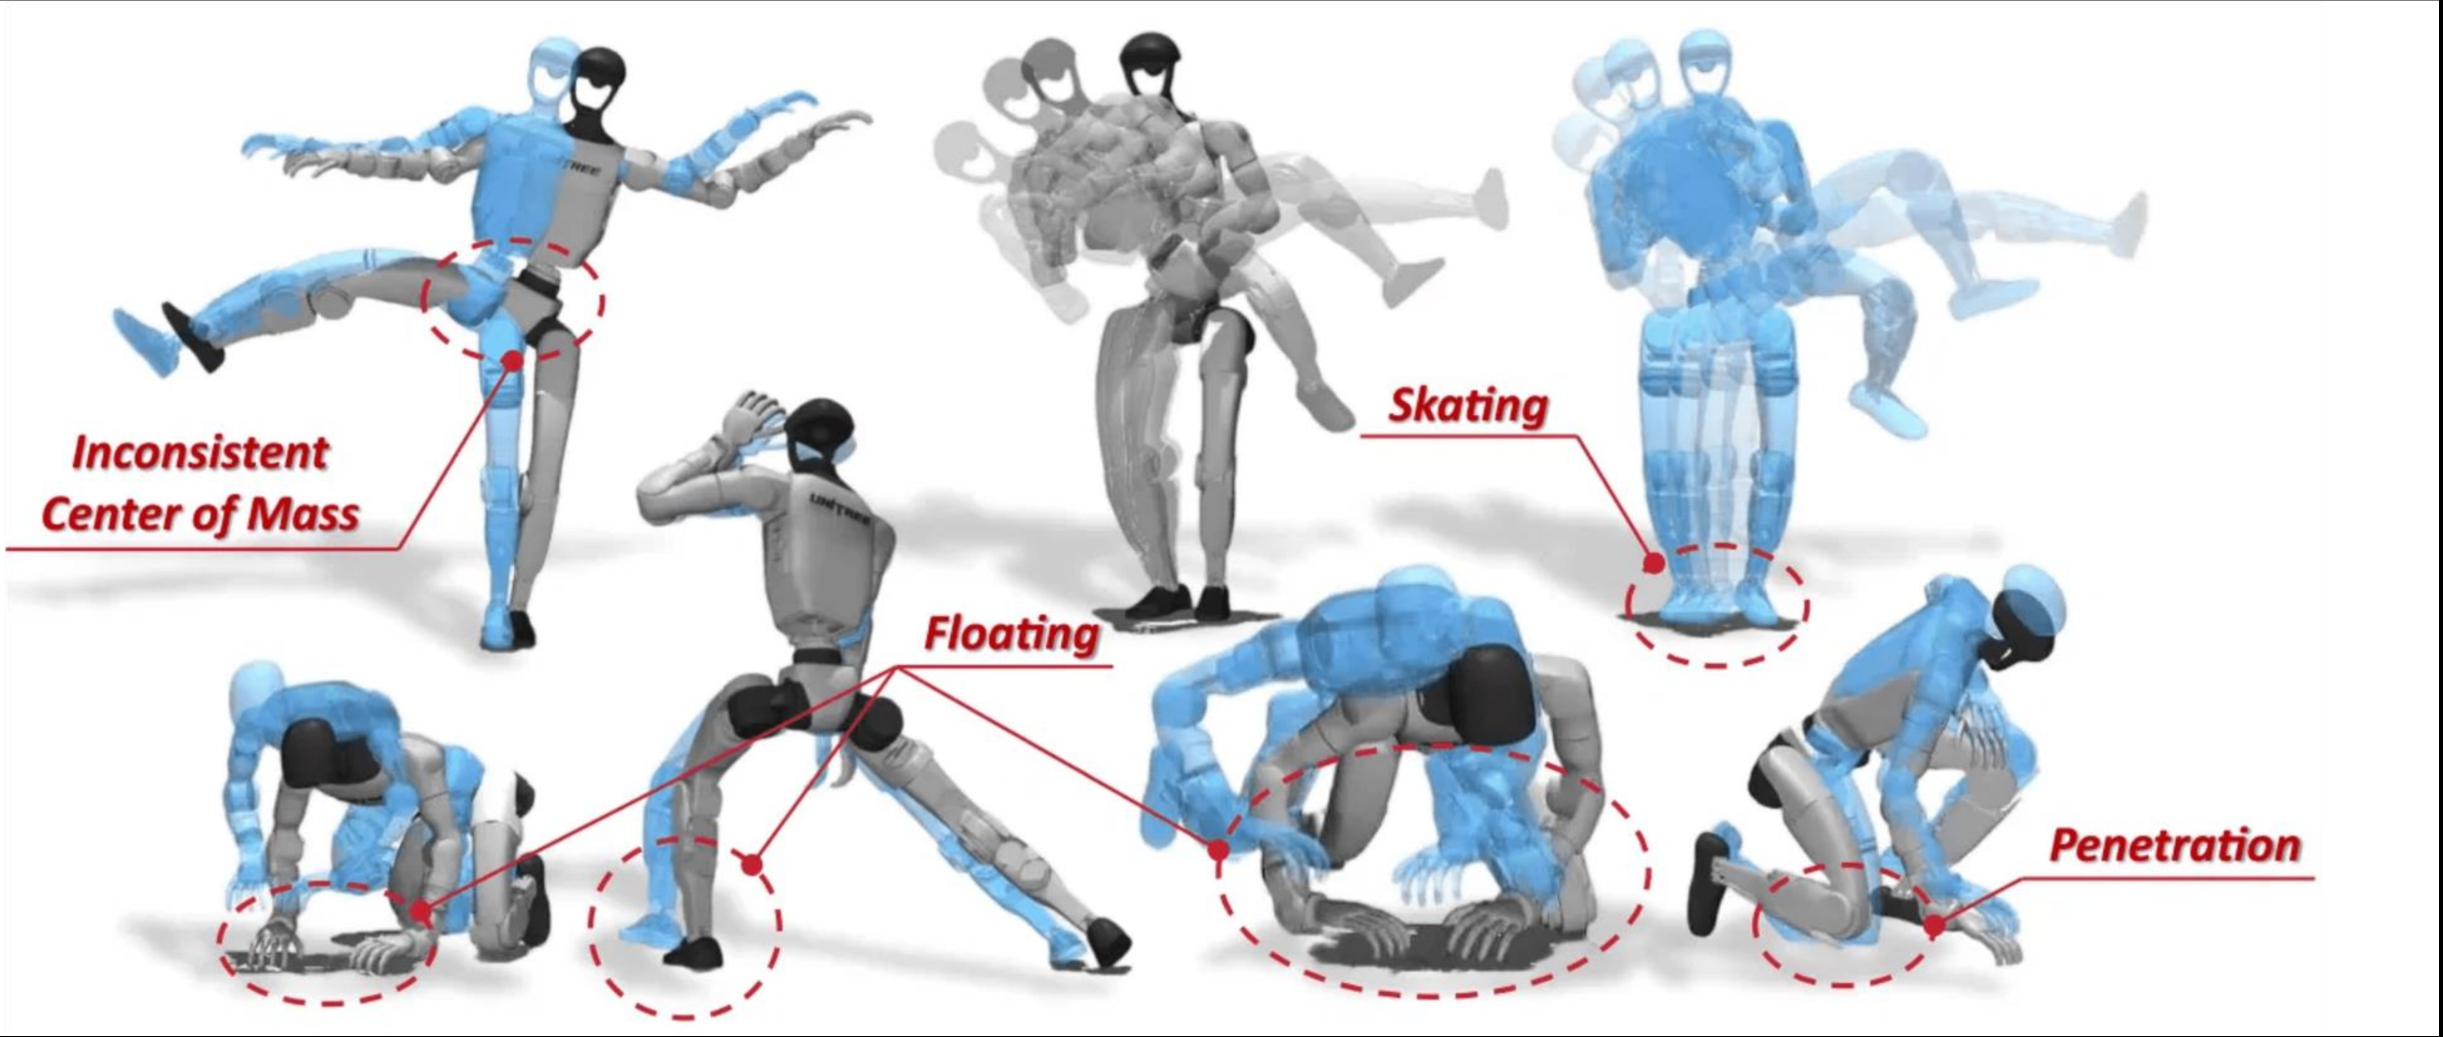

Li et al., OmniTrack: General Motion Tracking via Physics-Consistent Reference, ArXiv'26

OmniTrack: Physical motion generation for motion tracking

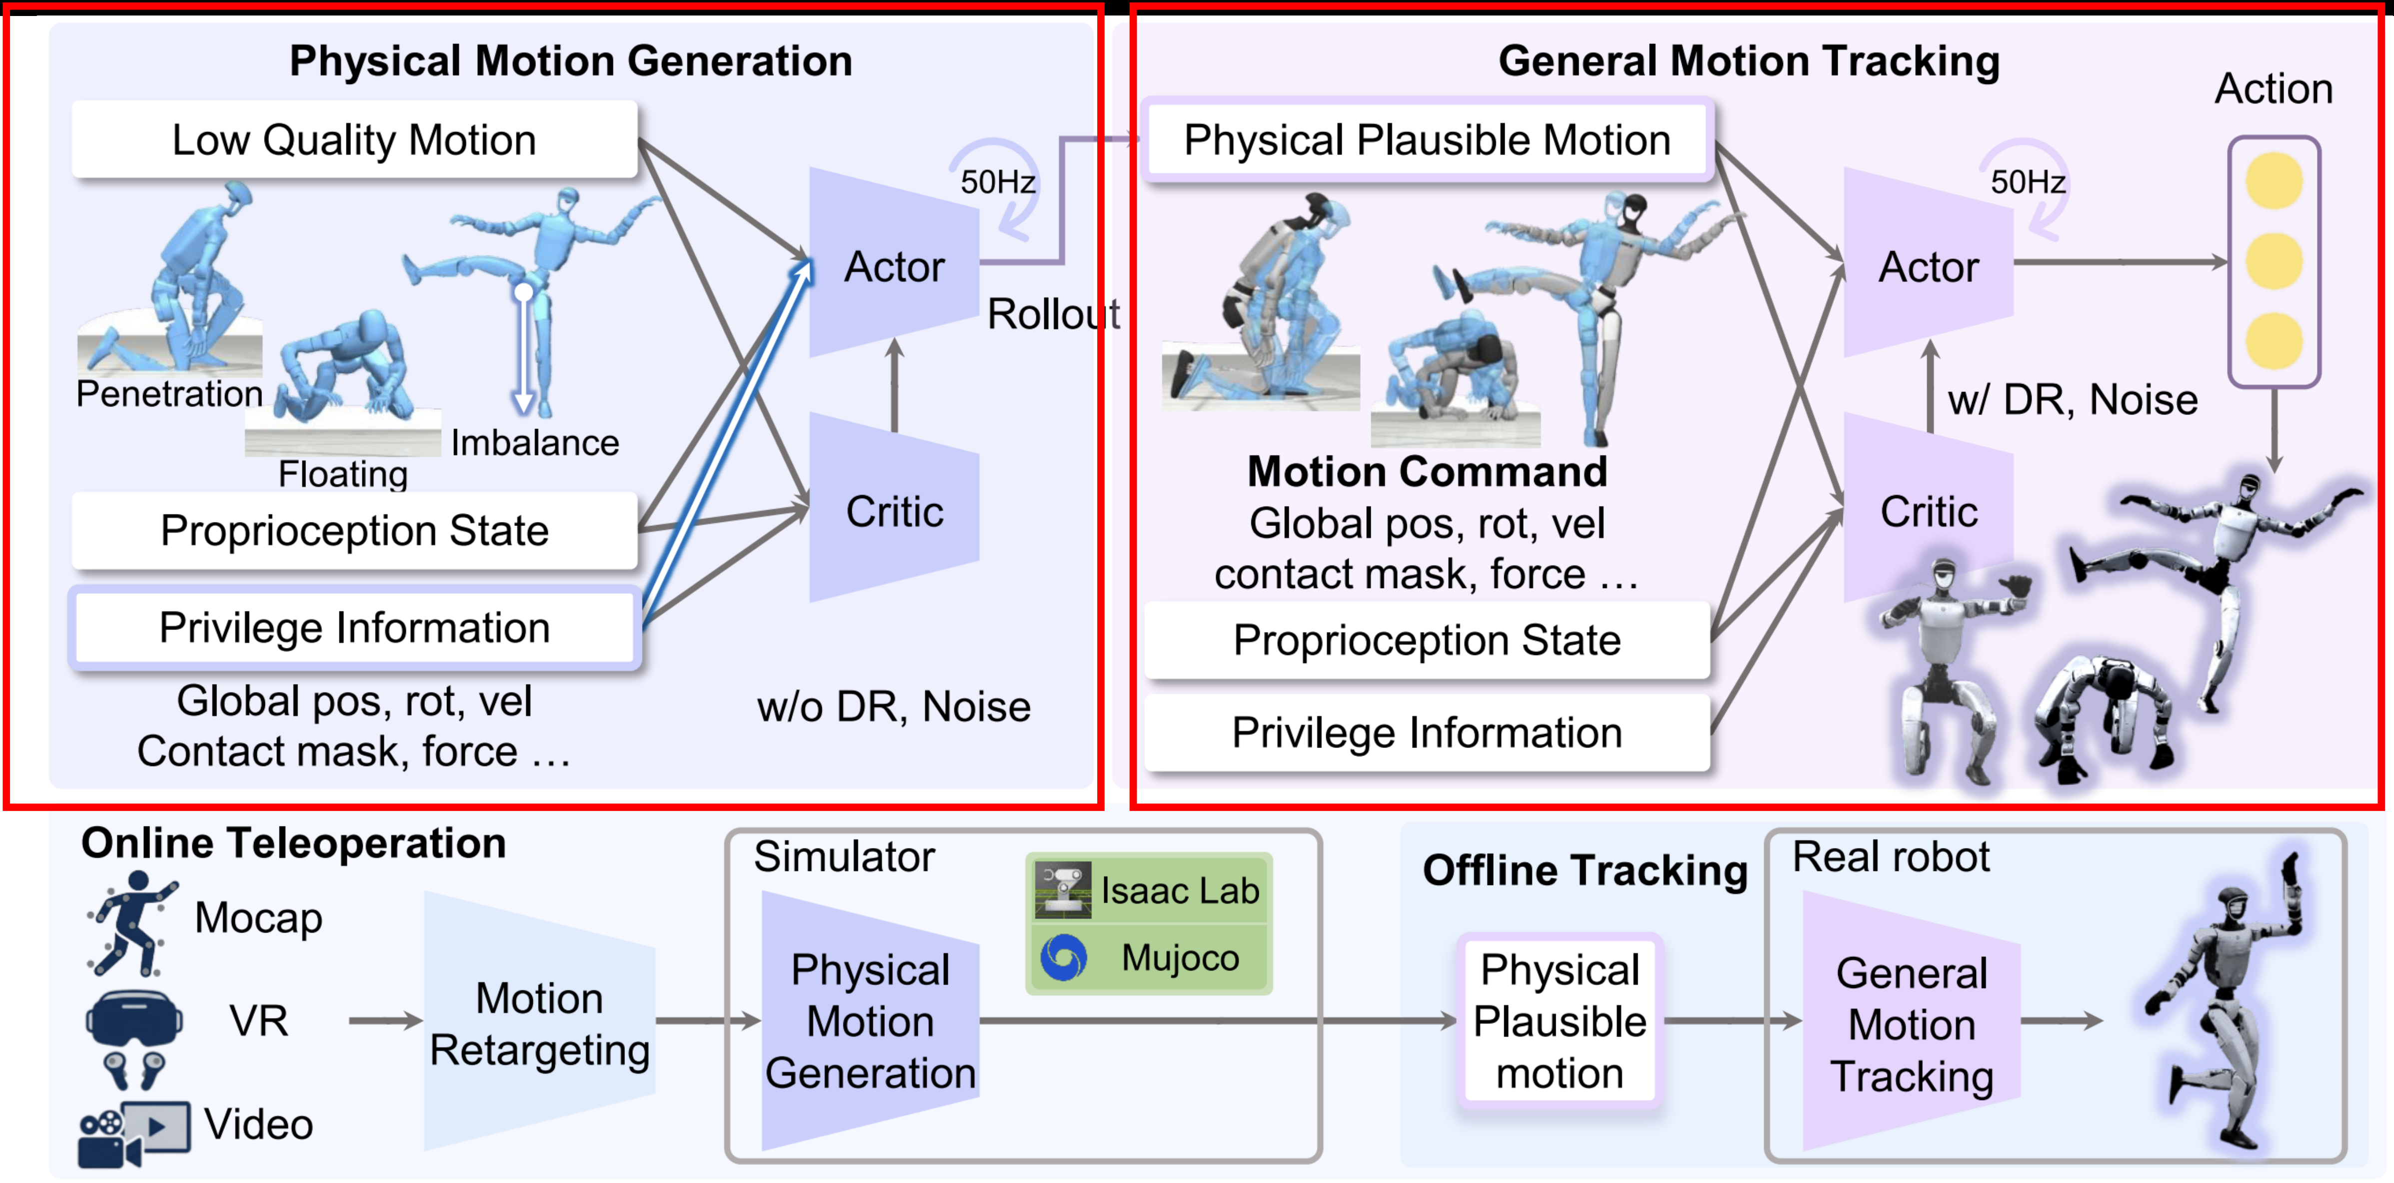

Constructing physical consistent reference motions

Dataset	Reference	Penetration	Floating	Smoothness	MPJPE
LAFAN1	Raw Ref.	20.3%	2.52%	33.7	0.000
	Physical Ref.	**0.0%**	**0.0%**	**31.8**	21.0
AMASS	Raw Ref.	67.7%	3.29%	19.5	0.000
	Physical Ref.	**0.0%**	**0.0%**	**15.8**	16.0

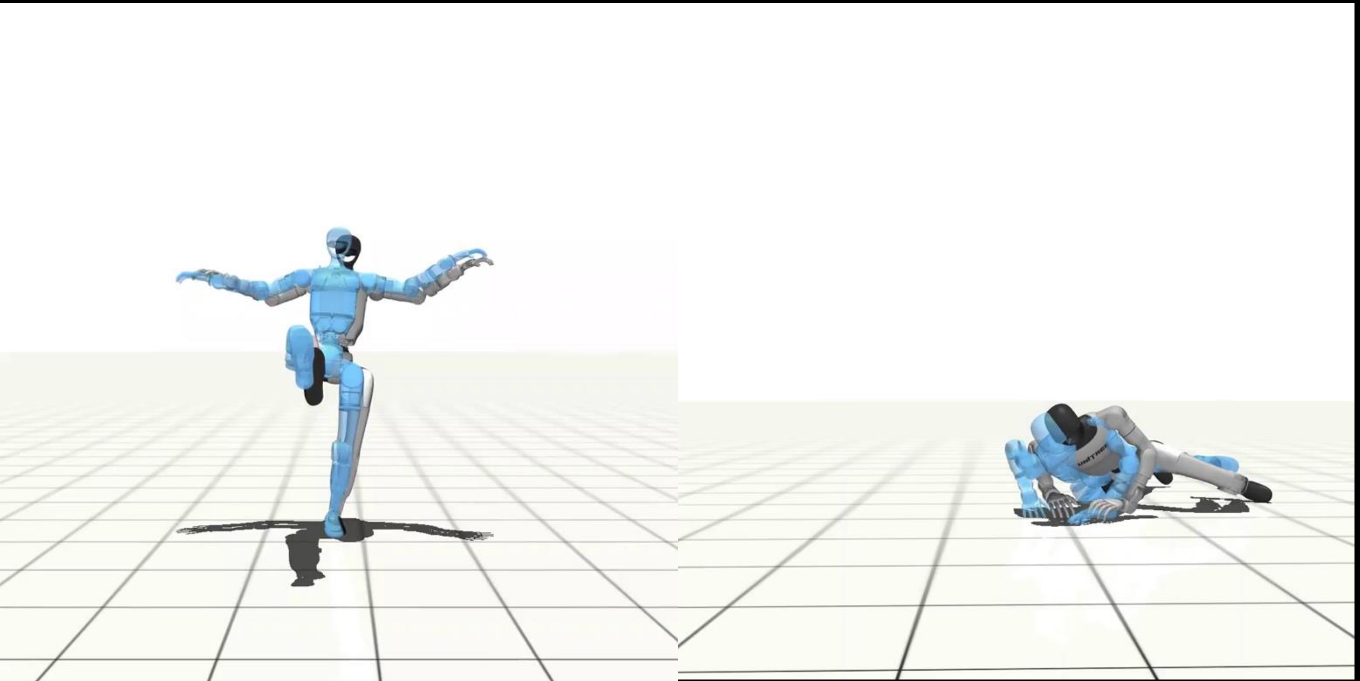

- **Retargeting causes physical inconsistencies**
- **Roll out with a zero DR teacher policy leads to motion quality improvement**

Physical motion for scaling general motion tracking

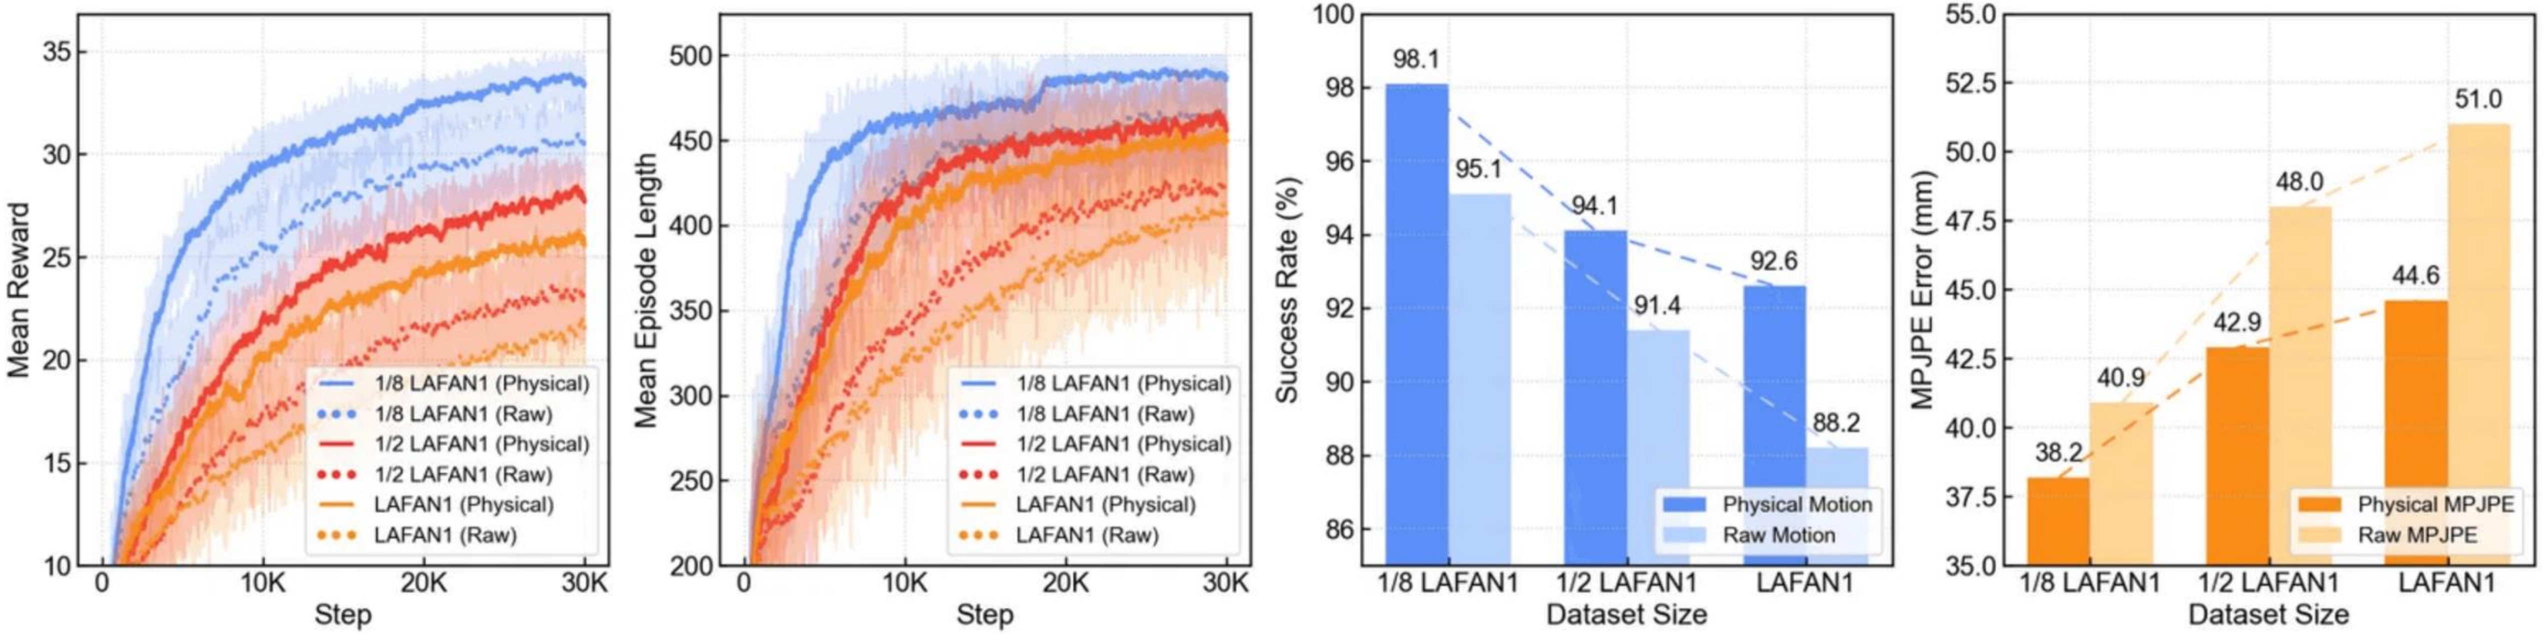

Physical reference motions improve both the success rate and the tracking quality of the raw motion with scaling effects

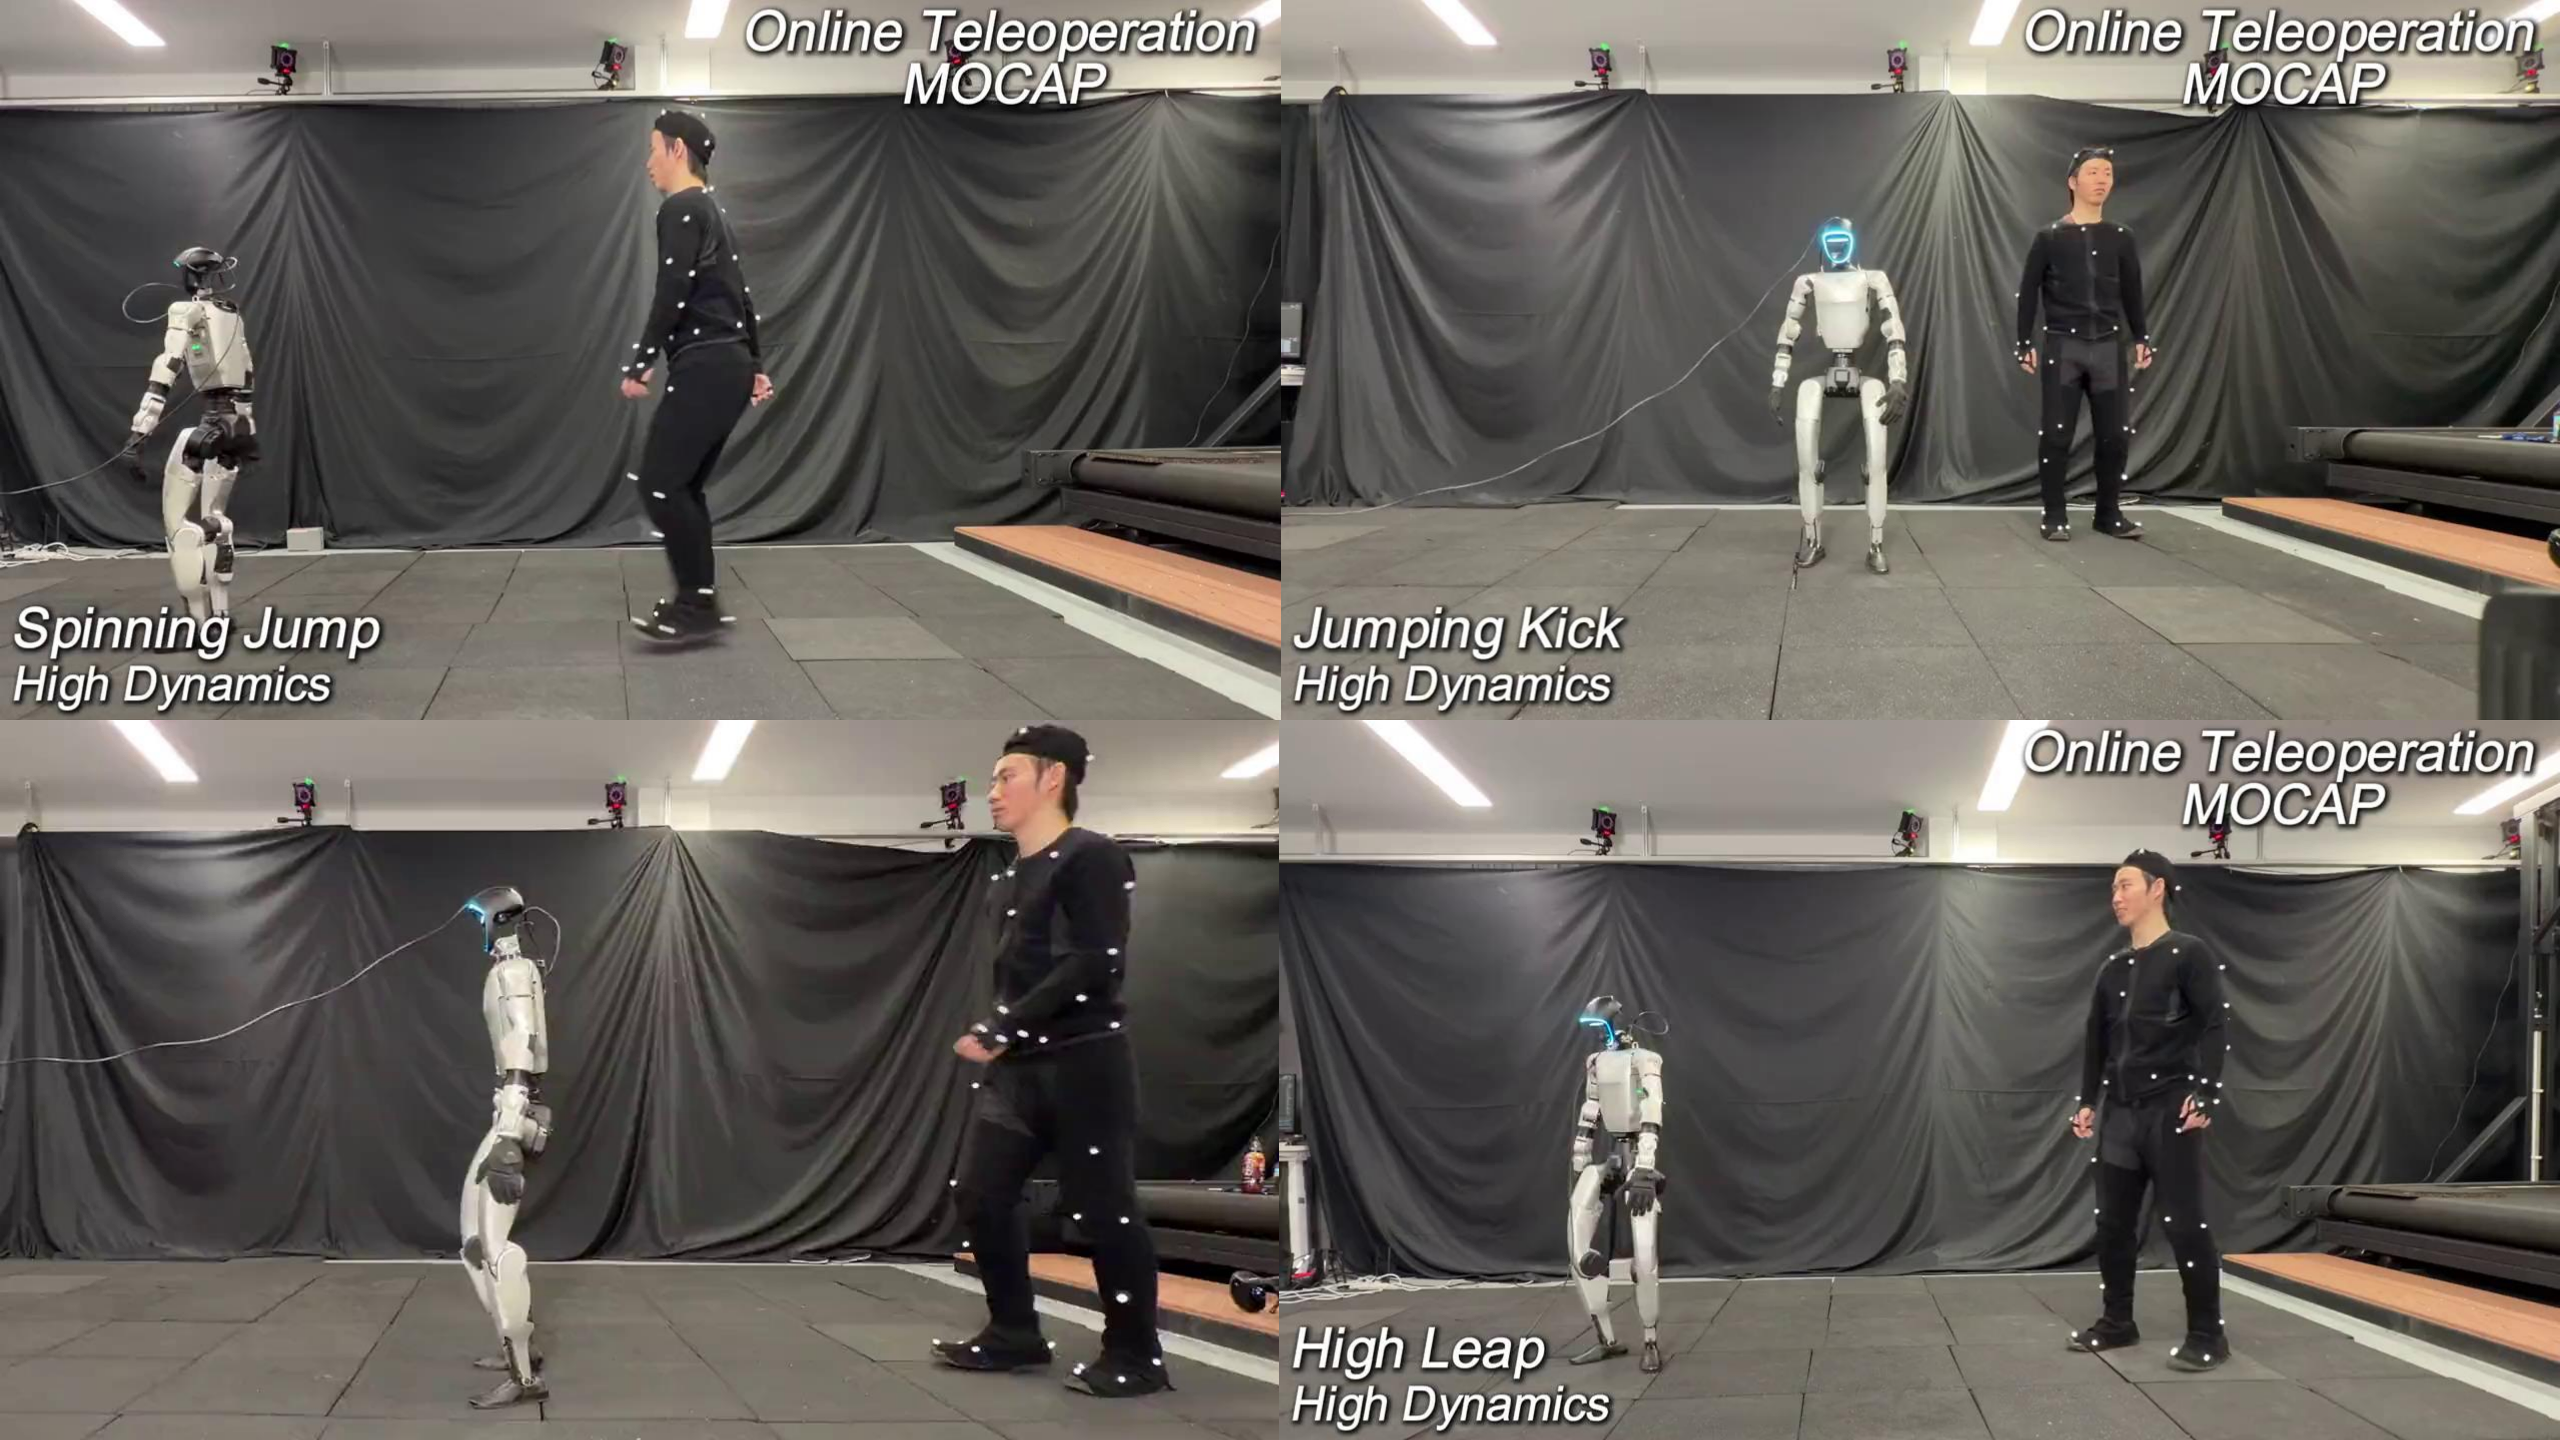
Online Teleoperation
MOCAP
Spinning Jump
High Dynamics
Online Teleoperation
MOCAP
Jumping Kick
High Dynamics
Online Teleoperation
MOCAP
High Leap
High Dynamics


Online Teleoperation
MOCAP
Walk and Run
Human Style
Online Teleoperation
MOCAP
Kneel and Sit
Contact Rich

General Motion Tracking

All results are generated by a single general policy

OmniXtreme: Breaking the Generality Barrier in High-Dynamic Humanoid Control

Yunshen Wang*, Shaohang Zhu*, Peiyuan Zhi, Yuhan Li, Jiaxin Li,
Yong-Lu Li, Yuchen Xiao, Xingxing Wang, Baoxiong Jia, Siyuan Huang

Problem 1: Sim2Real Gap

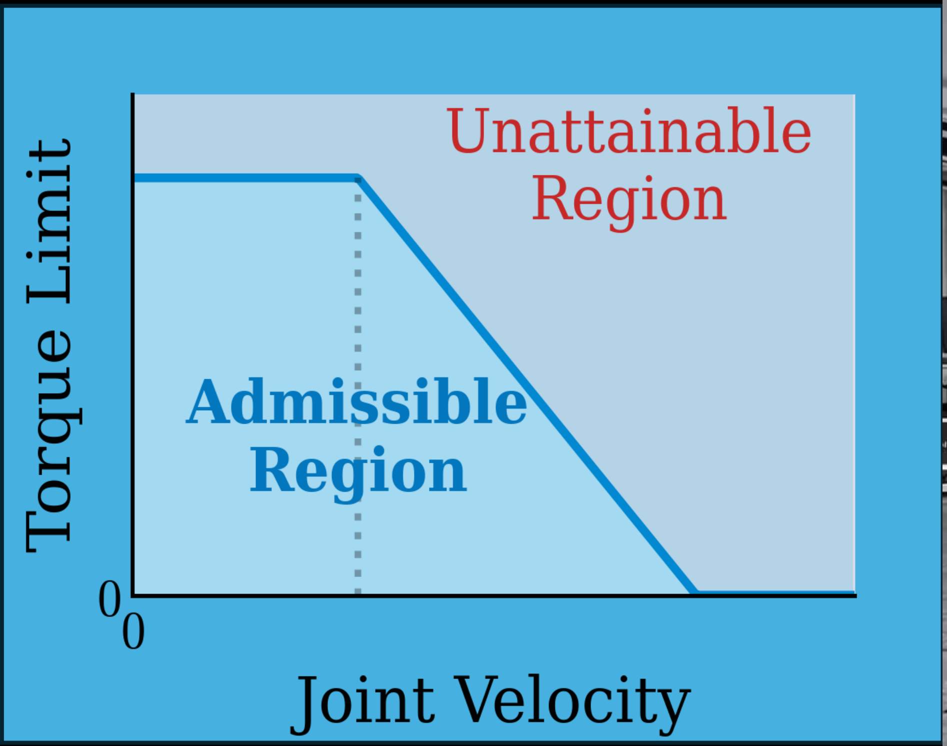

Real motors cannot produce arbitrary torque

Problem 1: Sim2Real Gap

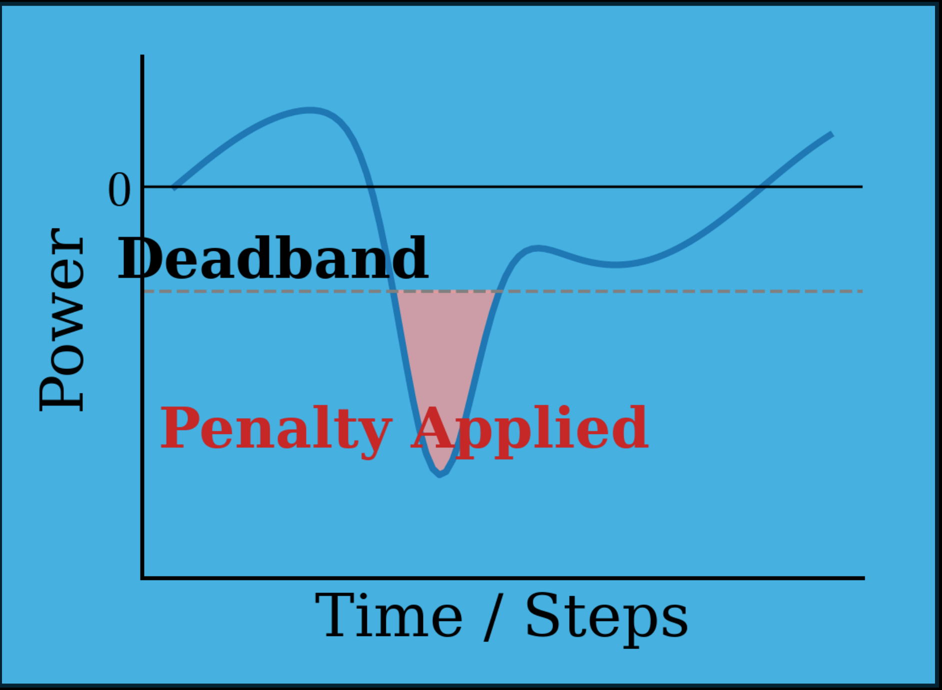

Energy-aware control improves both stability and hardware safety

Scaling on extreme motions

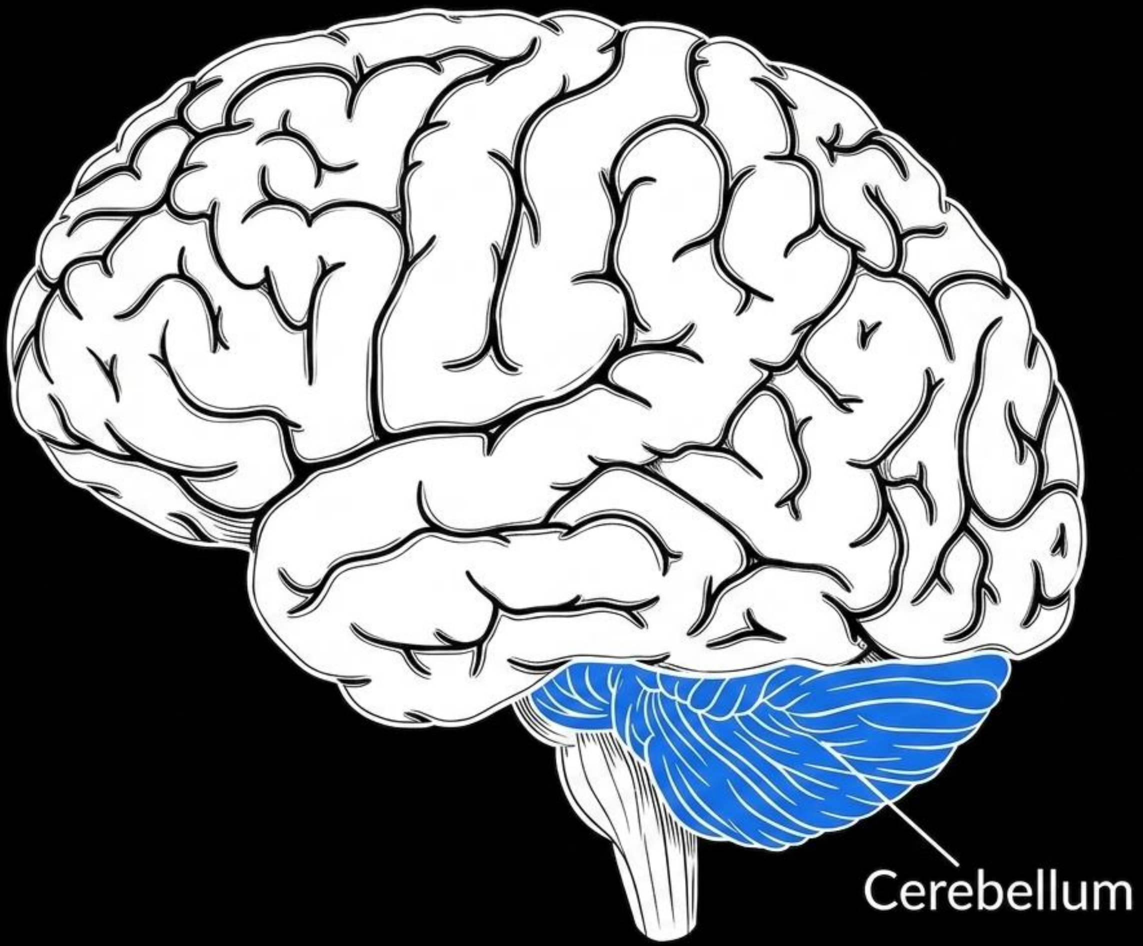

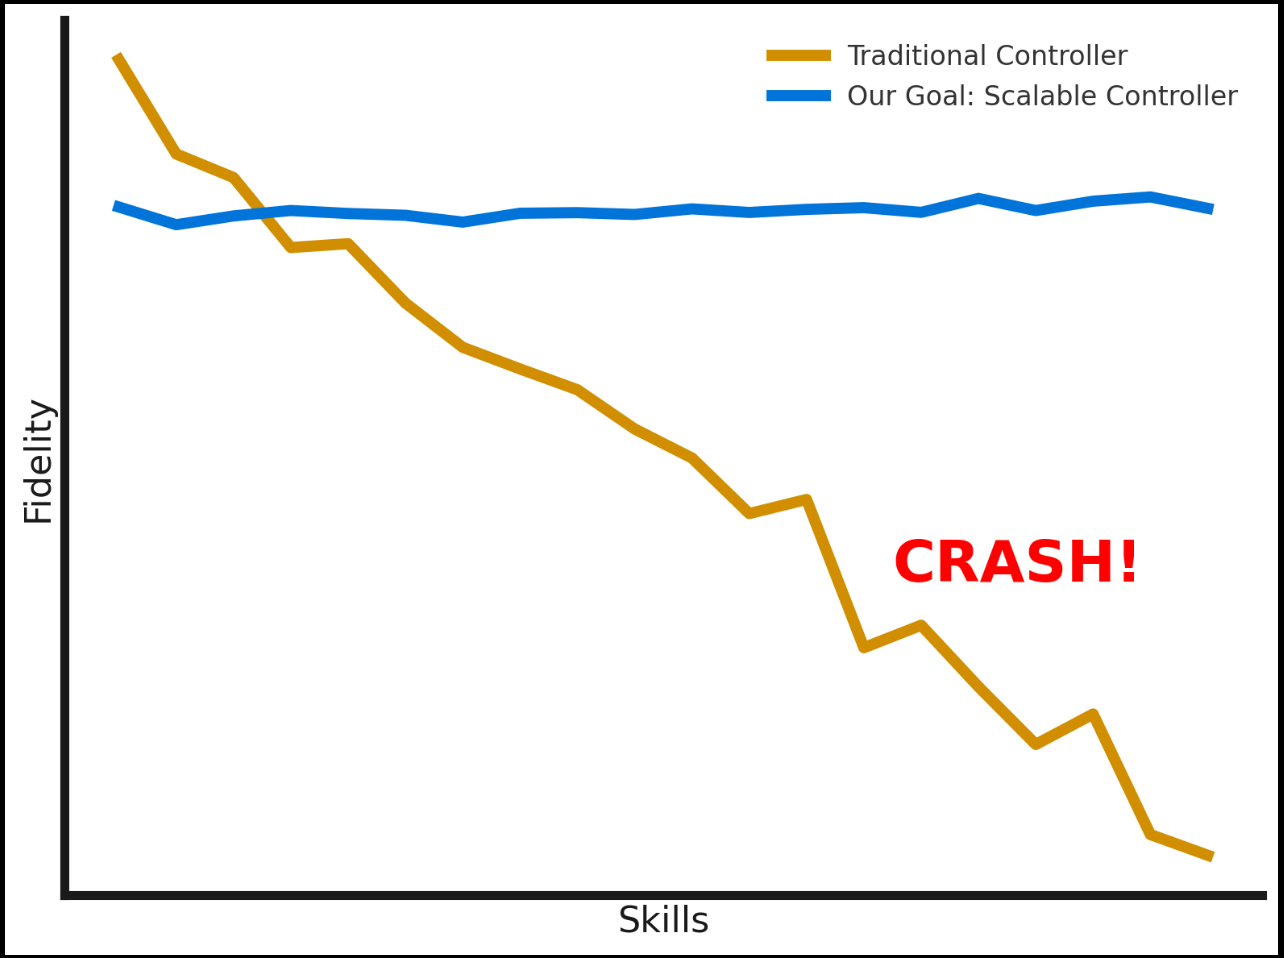

Cerebellum contains roughly 80% of the brain's neurons

Problem 2: Scaling effects on the policy

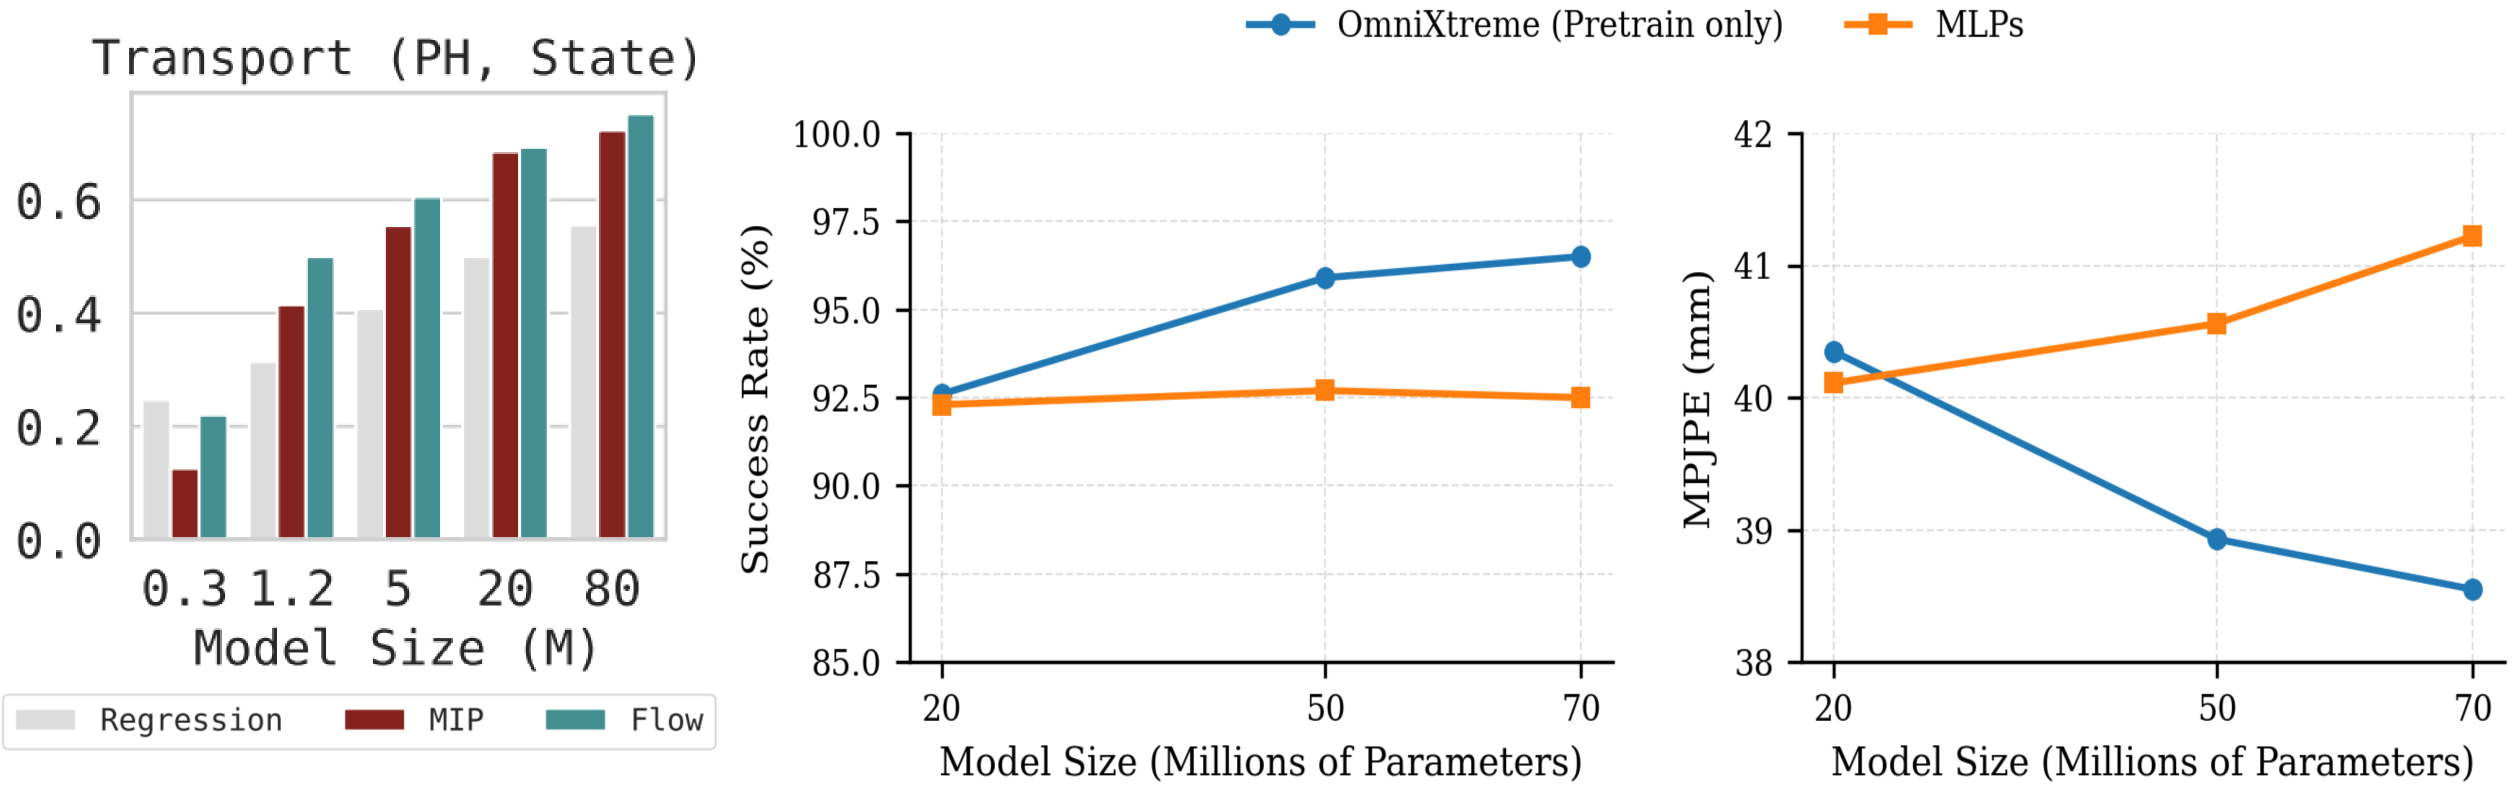

Naïve scaling of policy network size (MLPs) is not very effective
Flow-matching as a better supervision for scaling and representation learning

Pan et al., Much Ado About Noising: Dispelling the Myths of Generative Robotic Control, ArXiv'25

The OmniXtreme pipeline

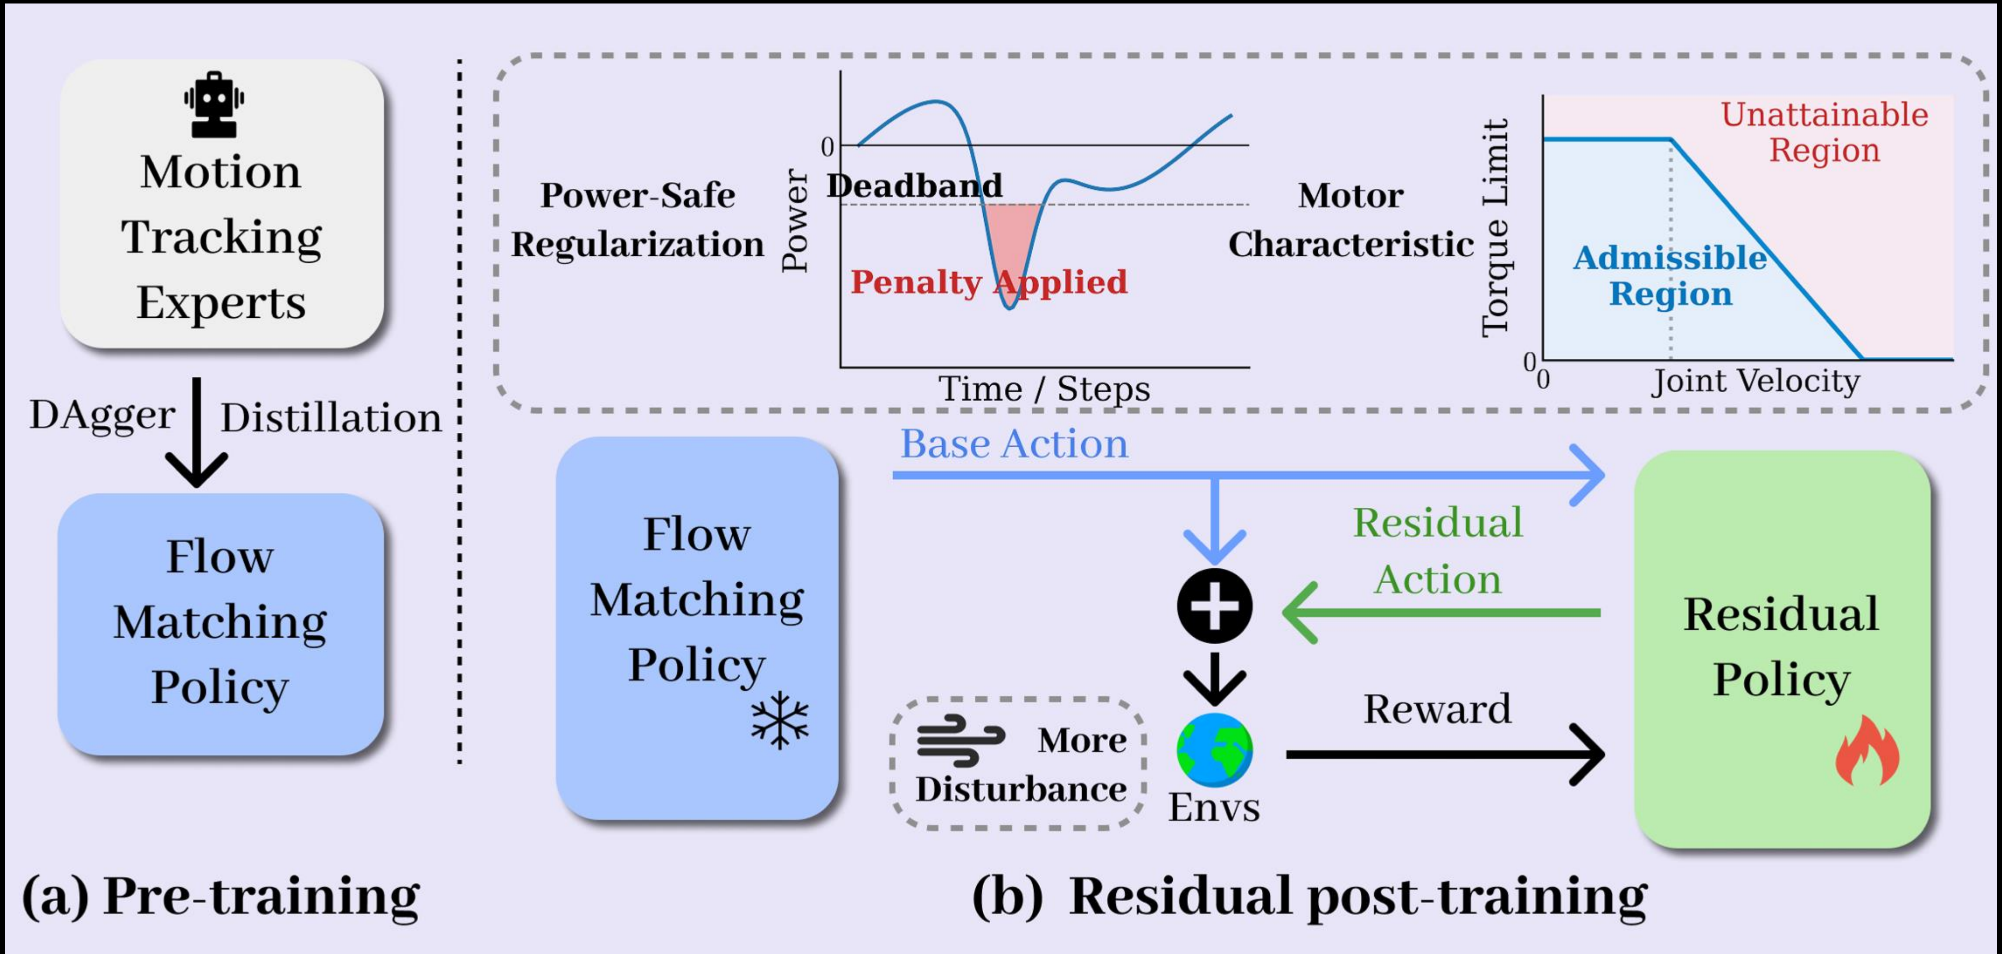

Two-stage DAgger with a flow-matching student policy with residual actions

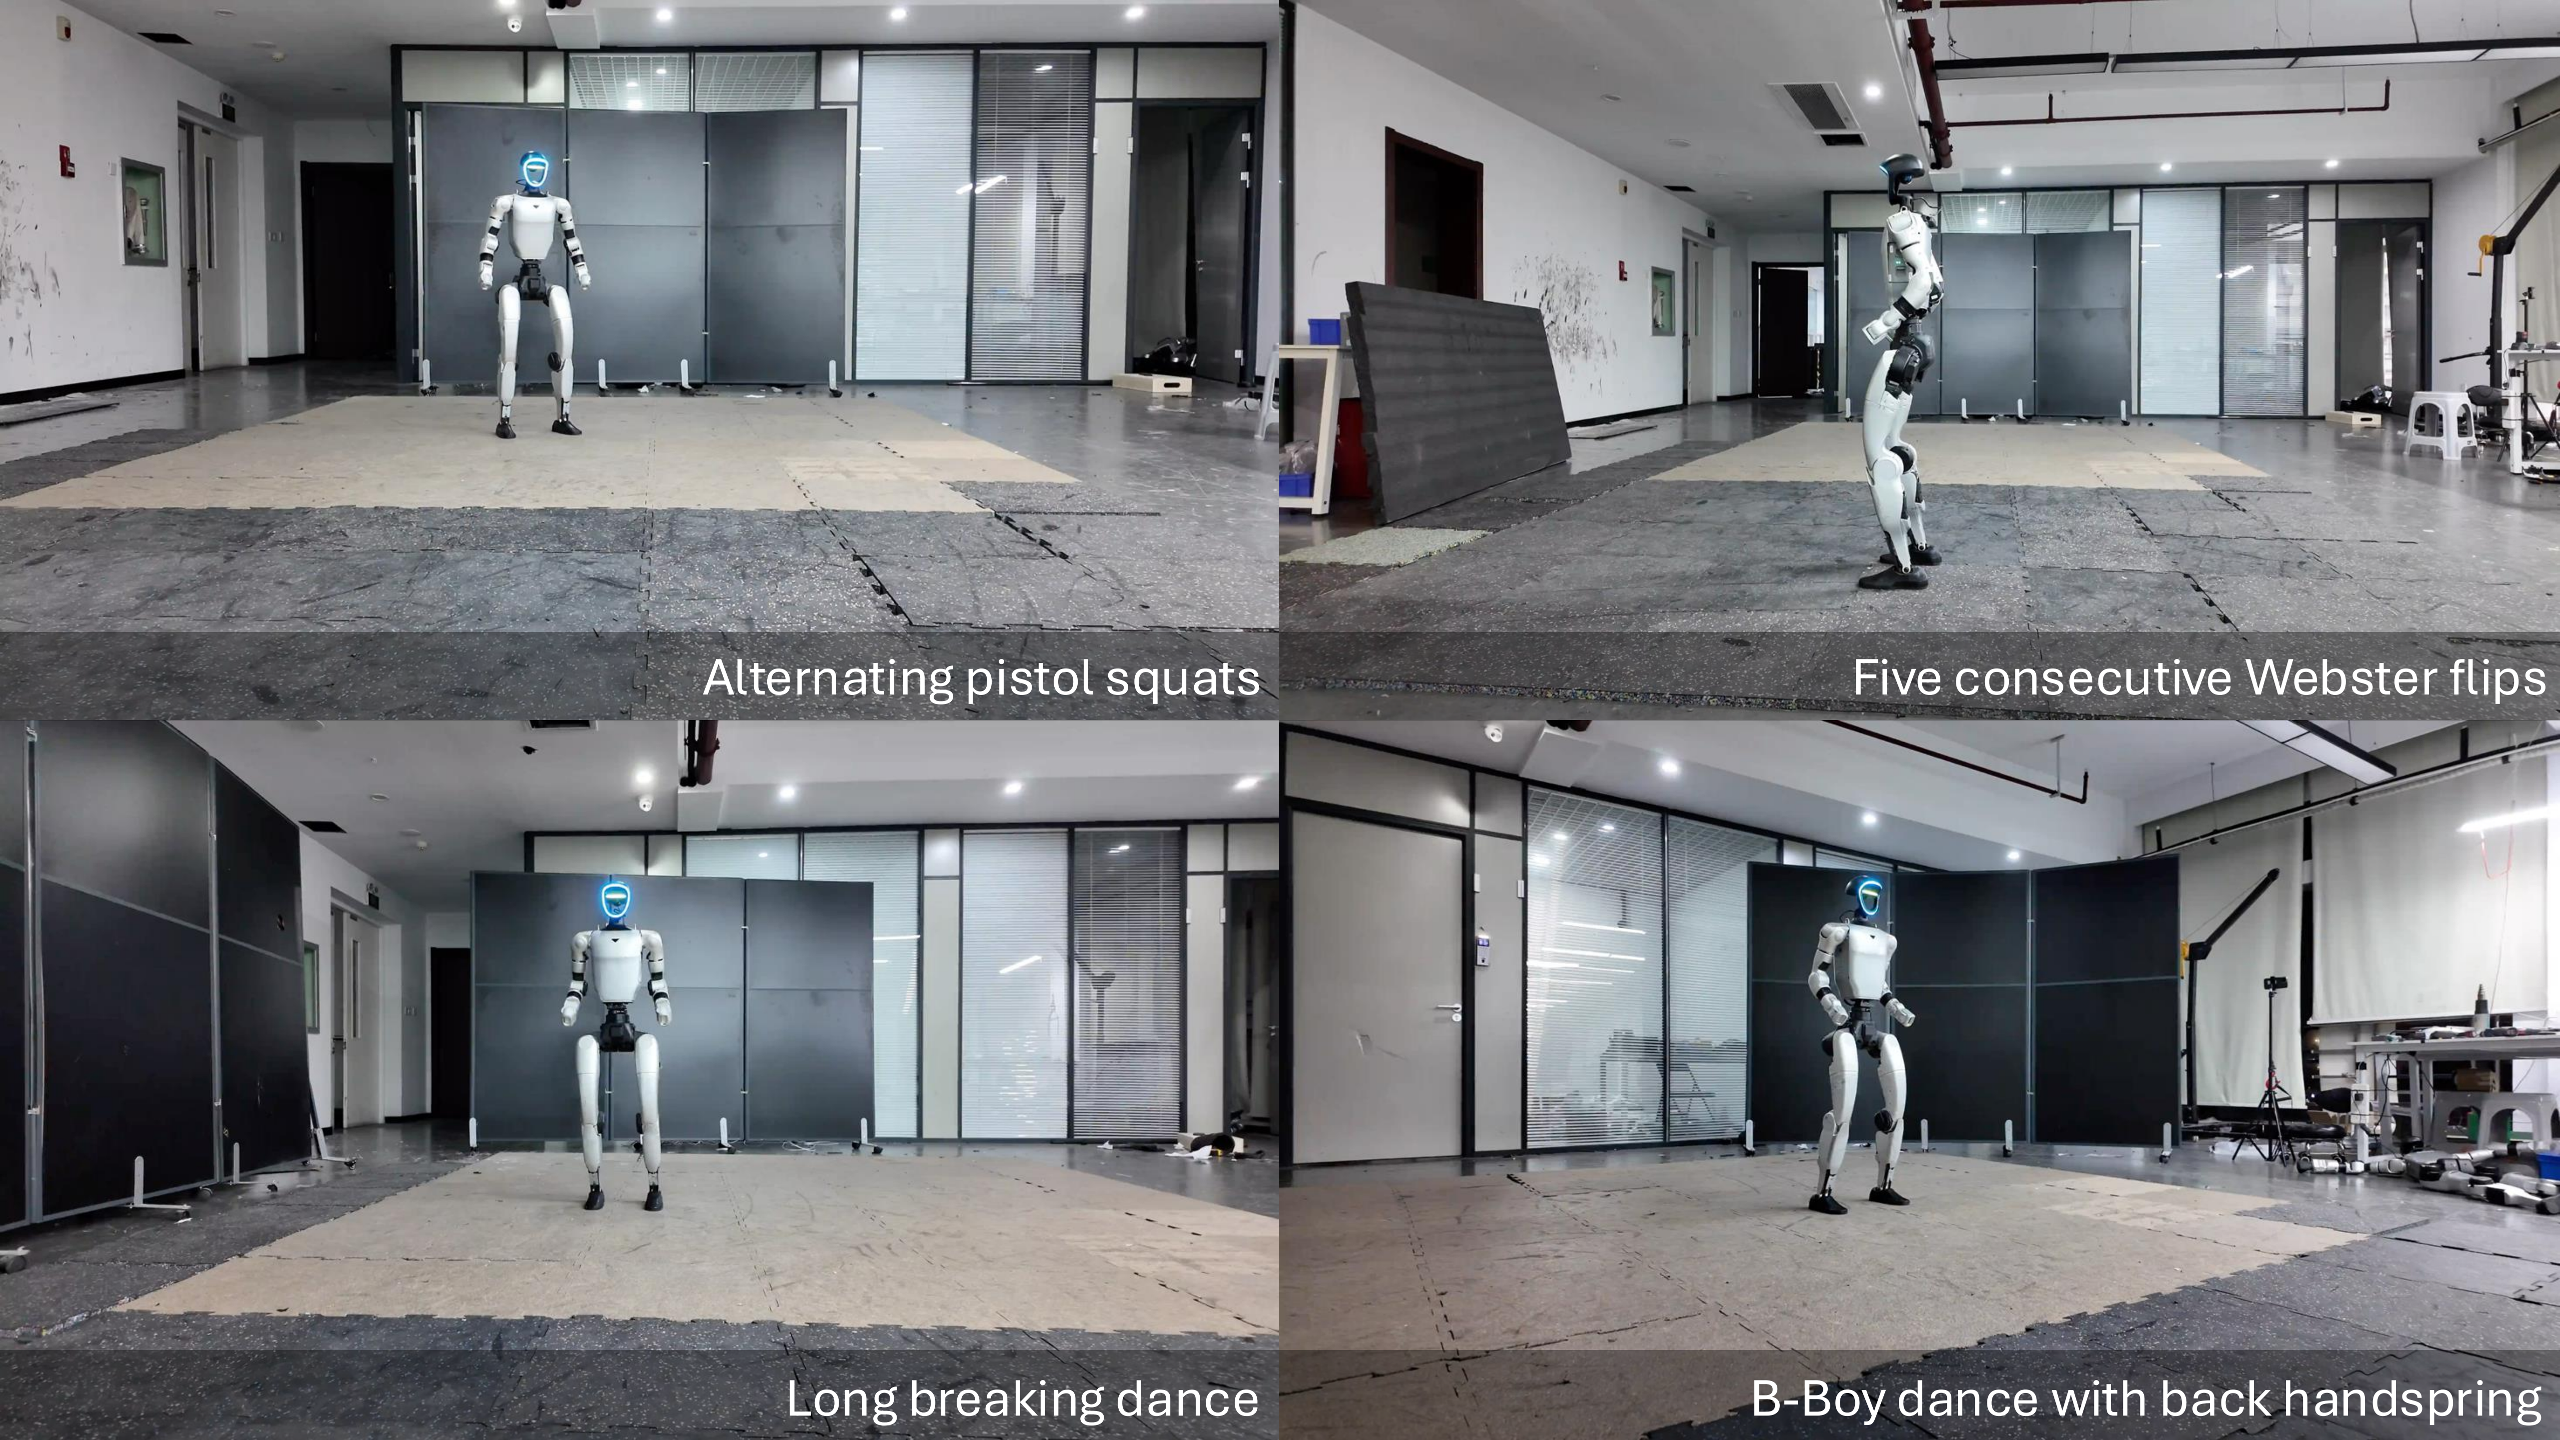
Alternating pistol squats
Five consecutive Webster flips
Long breaking dance
B-Boy dance with back handspring

Enabling autonomous humanoid policies

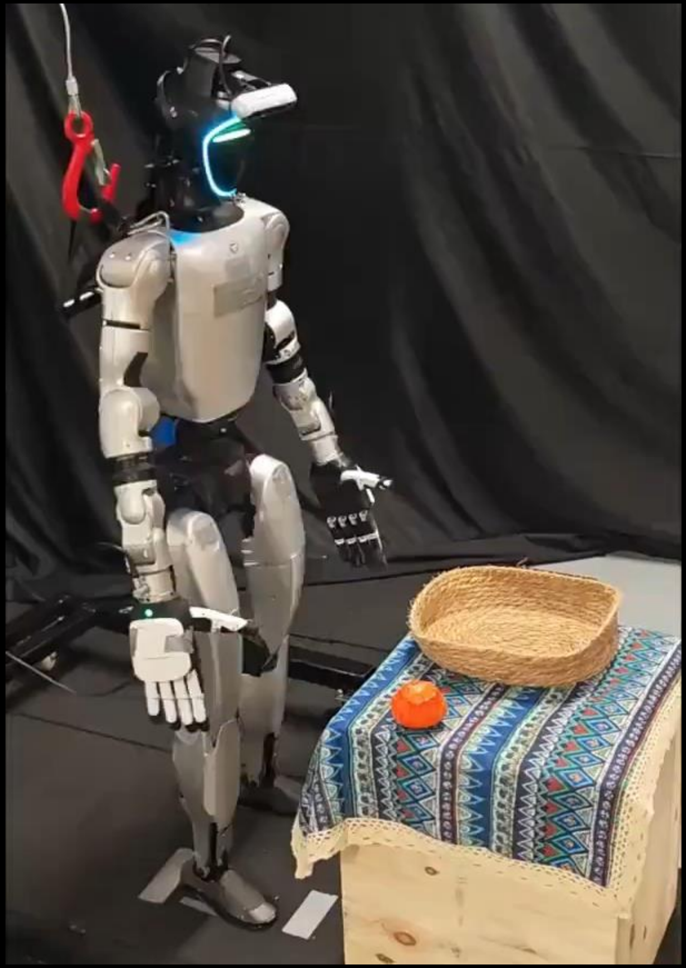 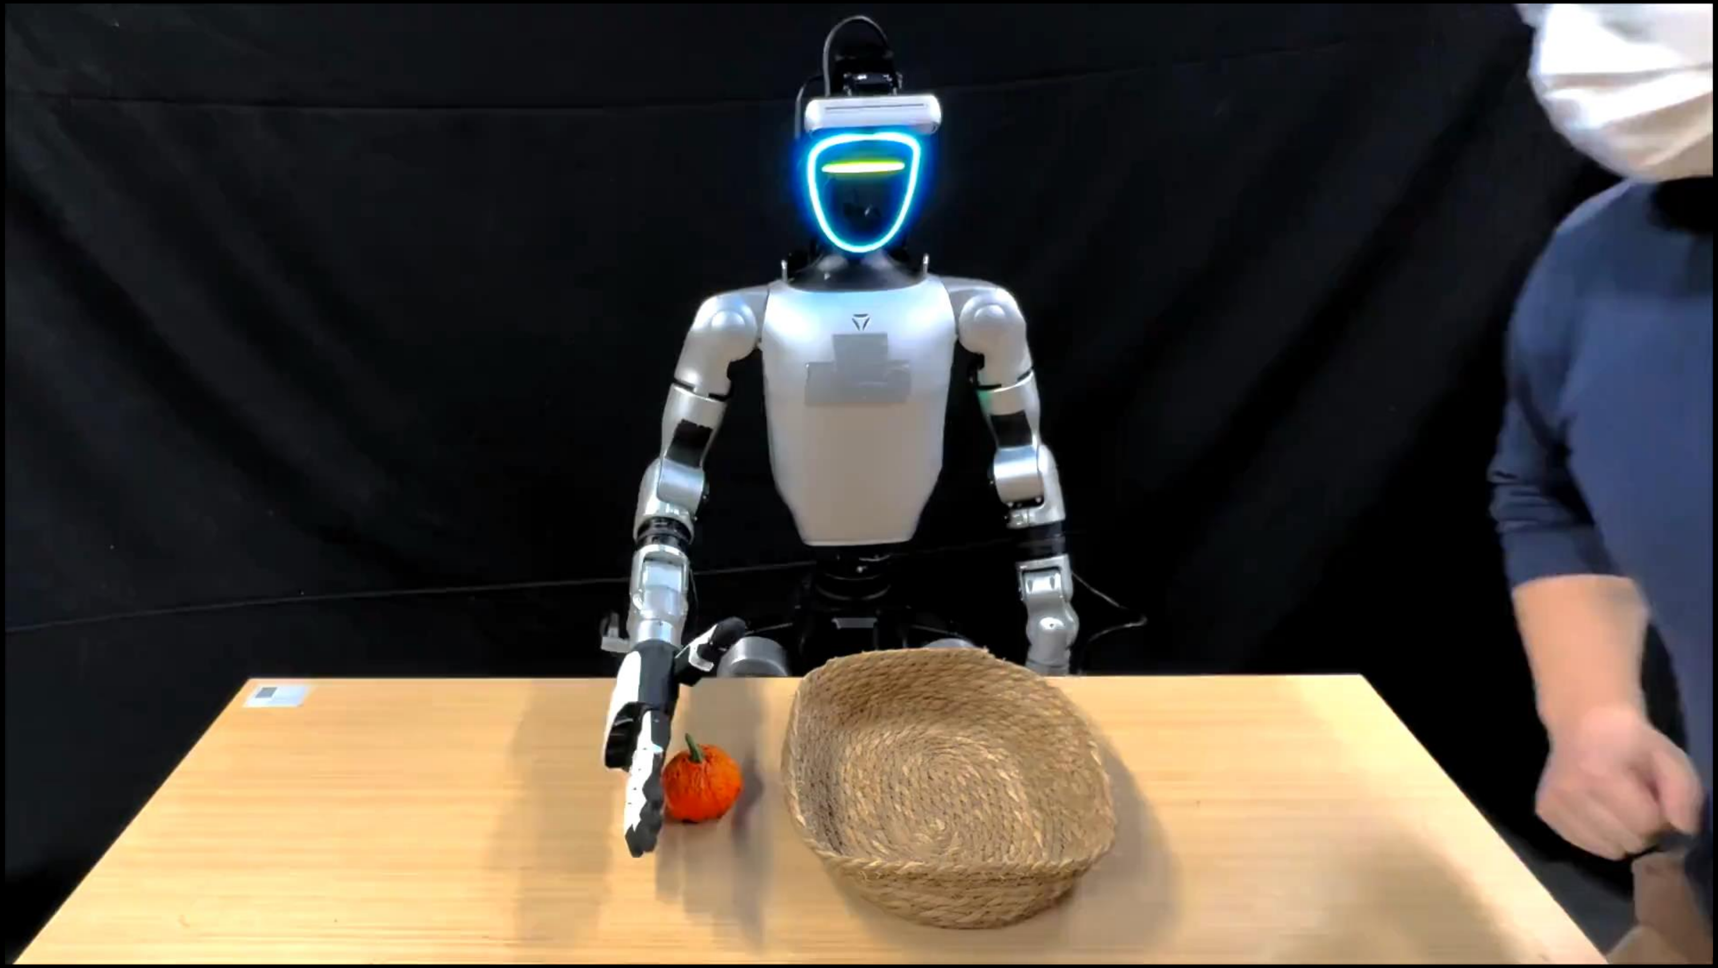

~50 demos real-robot teleoperation can unlock full-body manipulation capabilities

Task1: Pick and Place

"pick up the green tomato and put to the green basket"

OK, can we let the robot help us wipe the whiteboard first after meeting?

Let me collect the data and imitation learning will solve the rest ☺

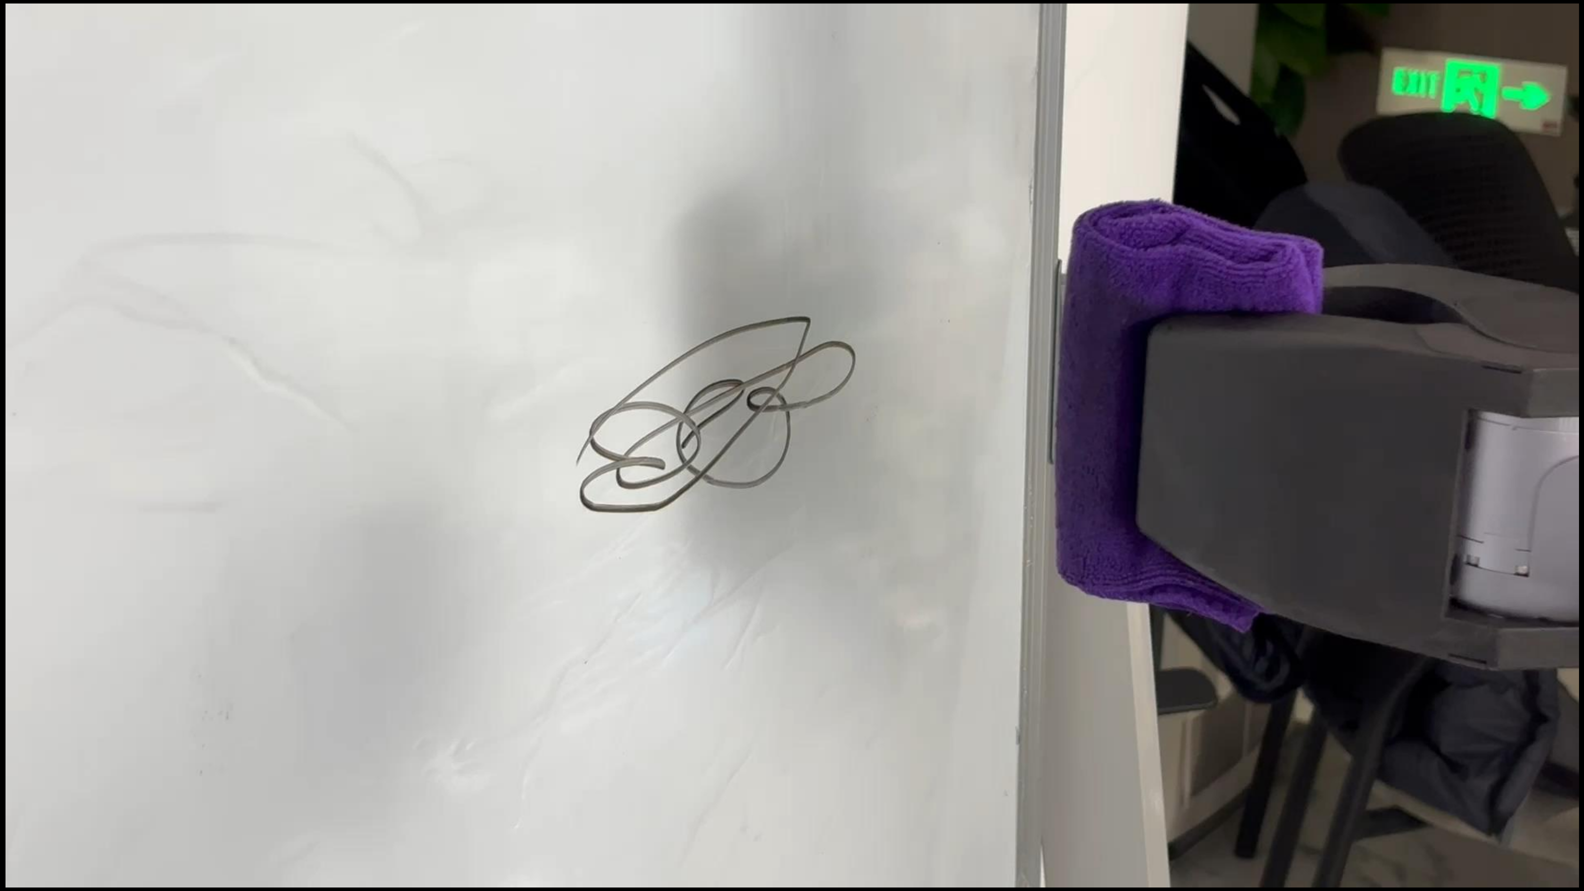

Data collection is a disaster

OK, can we let the robot help us wipe the whiteboard first after meeting?

Let me collect the data and imitation learning will solve the rest ☺

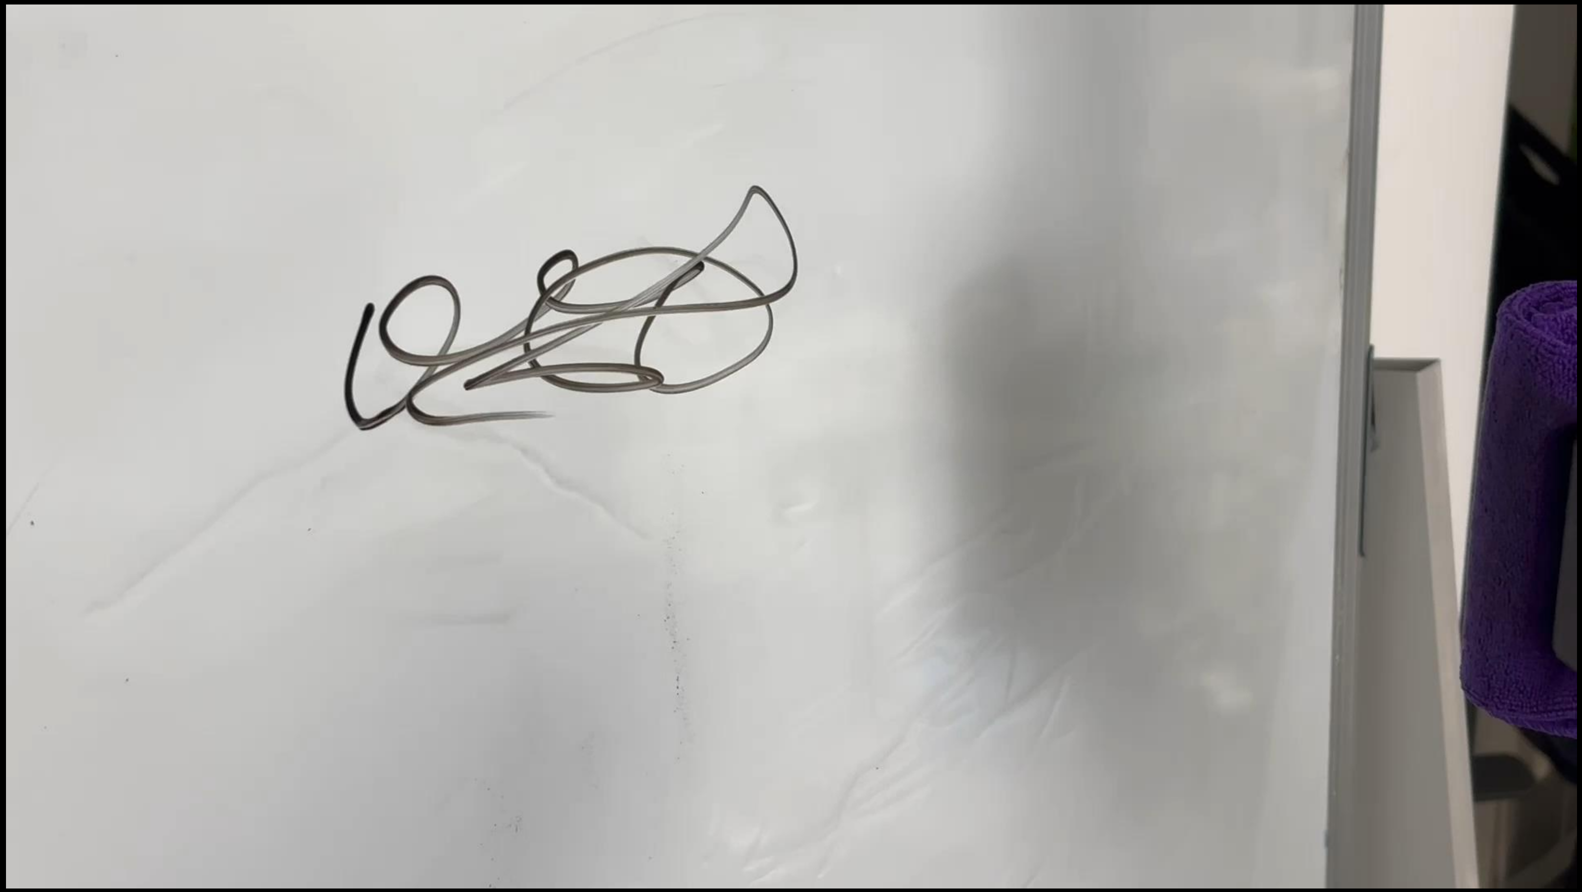

Of course, the learned policy failed no matter how much data used ☹

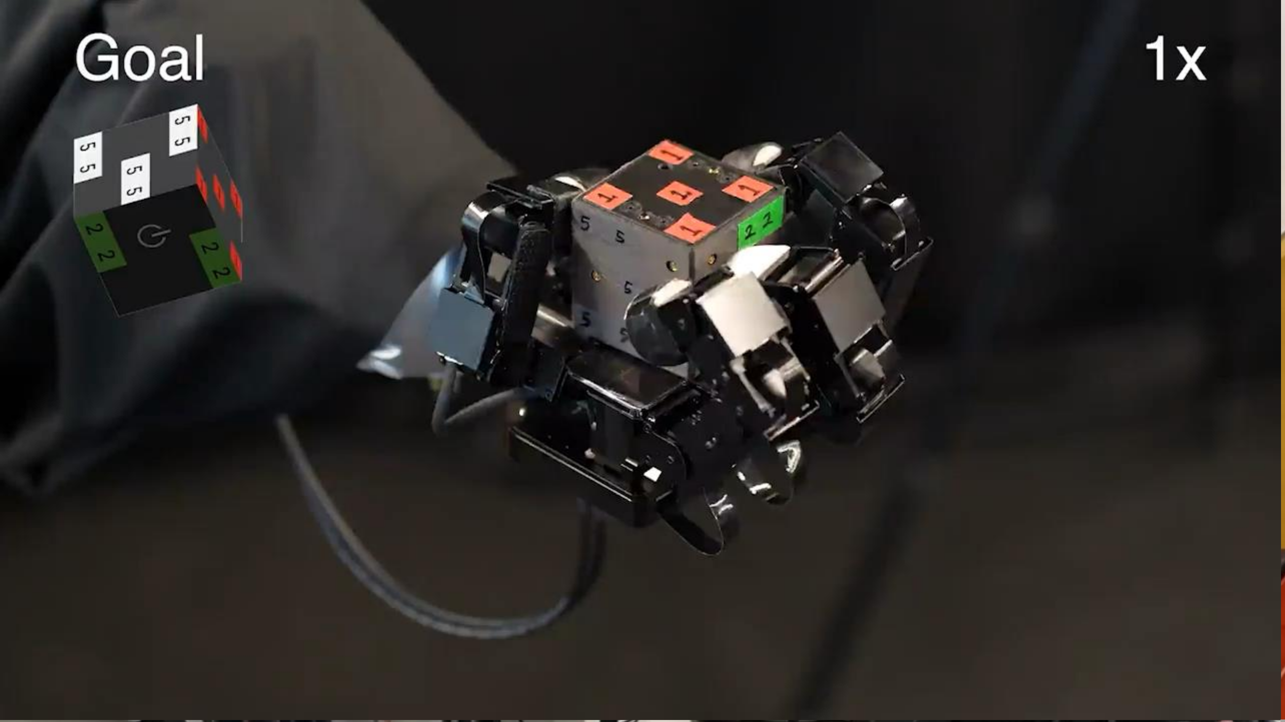
Goal
1x
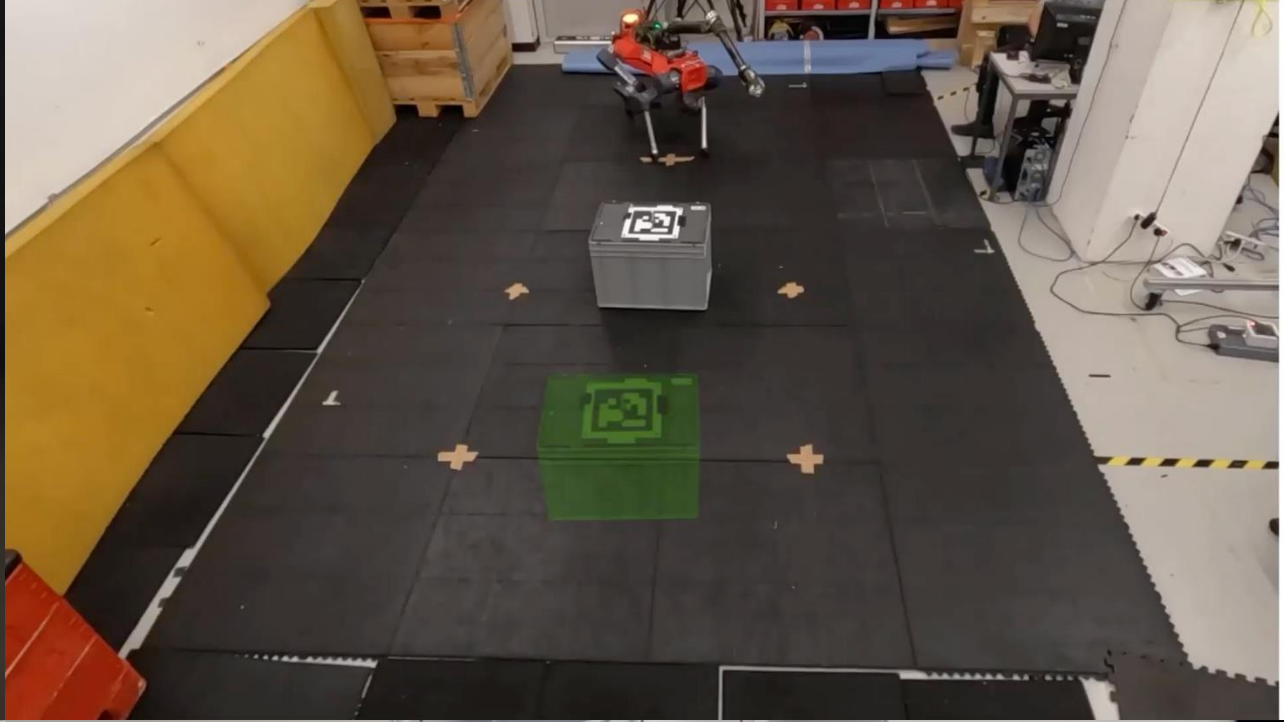


Task 1
Mustard Bottle
2x Speed

And when you don't have force-torque sensors...

mass-spring-damper system

$$\boldsymbol{F} = K\left(\mathrm{x} - \boxed{\mathrm{x}^{\mathrm{cmd}}}\right) + \cancel{D\left(\dot{\mathrm{x}} - \boxed{\dot{\mathrm{x}}^{\mathrm{cmd}}}\right)} + \cancel{M\left(\ddot{\mathrm{x}} - \boxed{\ddot{\mathrm{x}}^{\mathrm{cmd}}}\right)}$$

$$\mathrm{x} = \mathrm{x}^{\mathrm{cmd}} + \frac{F}{K}$$

And if the end effector moves really slowly...

Revisiting the control formulation

mass-spring-damper system

$$\boldsymbol{F} = K\left(\mathrm{x} - \mathrm{x}^{\mathrm{cmd}}\right) + \cancel{D\left(\dot{\mathrm{x}} - \dot{\mathrm{x}}^{\mathrm{cmd}}\right)} + \cancel{M\left(\ddot{\mathrm{x}} - \ddot{\mathrm{x}}^{\mathrm{cmd}}\right)}$$

$$\mathrm{x} = \mathrm{x}^{\mathrm{cmd}} + \frac{F}{K}$$

Force can be estimated via position offsets!
Tracking the force-adjusted position enables joint force-position control.

Formulating forces with positions

$$\boldsymbol{F} = K\left(\mathrm{x} - \mathrm{x}^{\mathrm{cmd}}\right)$$

Position control

$$\mathrm{x}^{\mathrm{target}} = \mathrm{x}^{\mathrm{cmd}}$$

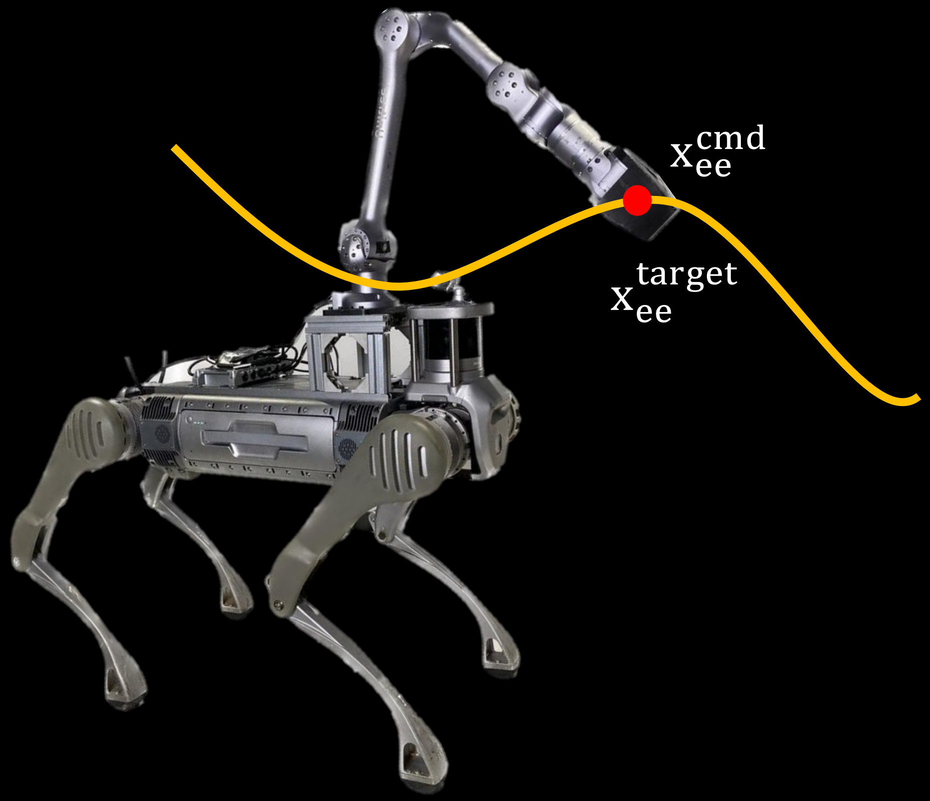

Formulating forces with positions

$$\boldsymbol{F} = K\left(\mathrm{x} - \mathrm{x}^{\mathrm{cmd}}\right)$$

When with external force

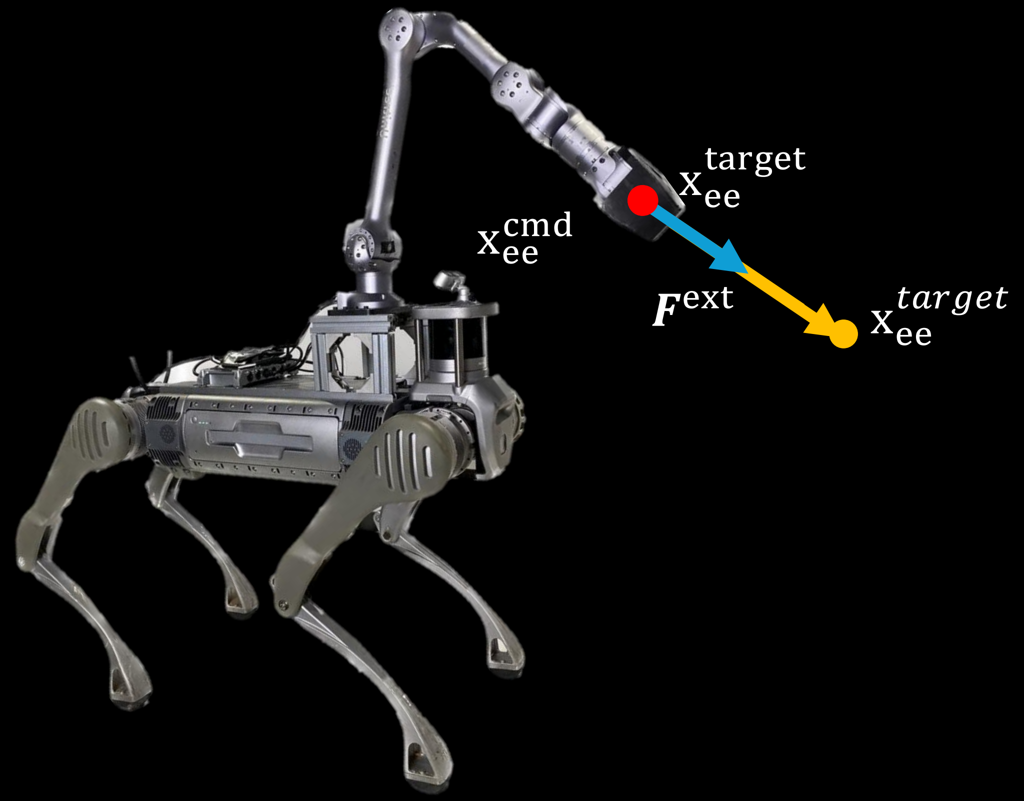

Formulating forces with positions

$$\boldsymbol{F} = K\left(\mathrm{x} - \mathrm{x}^{\mathrm{cmd}}\right)$$

Impedance control $$\mathrm{x}^{\mathrm{target}} = \mathrm{x}^{\mathrm{cmd}} + \frac{\boldsymbol{F}^{\mathrm{ext}}}{K}$$

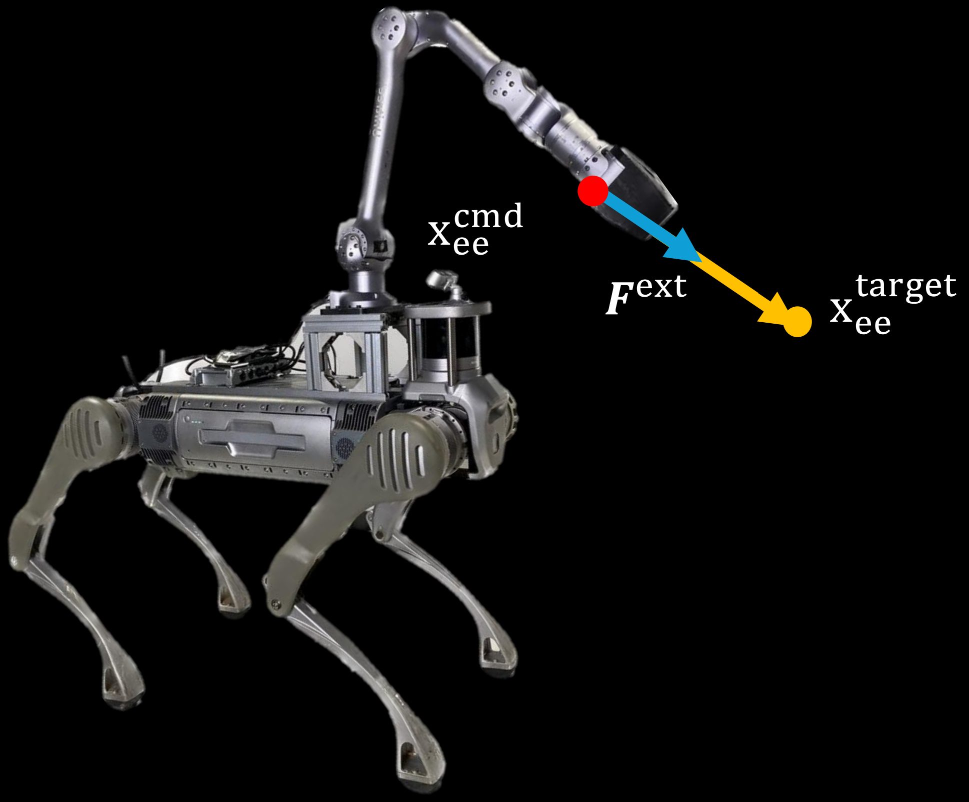

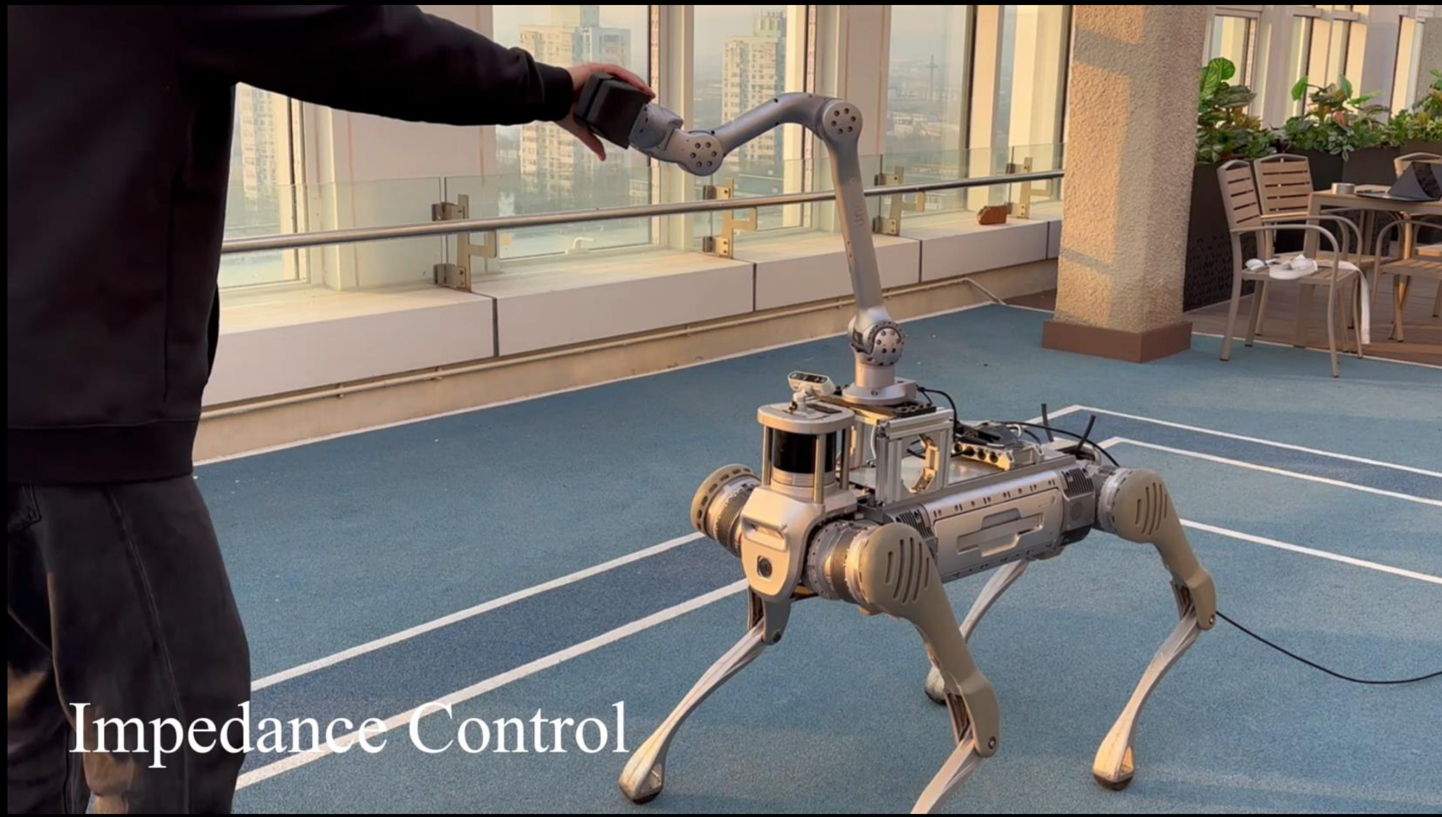

Revisiting the control formulation

mass-spring-damper system

$$\boldsymbol{F} = \cancel{K(\mathrm{x} - \mathrm{x}^{\mathrm{cmd}})} + D(\dot{\mathrm{x}} - \dot{\mathrm{x}}^{\mathrm{cmd}}) + \cancel{M(\ddot{\mathrm{x}} - \ddot{\mathrm{x}}^{\mathrm{cmd}})}$$

$$\dot{\mathrm{x}} = \dot{\mathrm{x}}^{\mathrm{cmd}} + \frac{F}{D}$$

Force-adjusted velocity enables compliant locomotion

Formulating forces with velocities

$$\boldsymbol{F} = D\left(\dot{\mathrm{x}} - \dot{\mathrm{x}}^{\mathrm{cmd}}\right)$$

Compliant locomotion

$$\dot{\mathrm{x}}^{\mathrm{target}} = \dot{\mathrm{x}}^{\mathrm{cmd}} + \frac{F^{\mathrm{ext}}}{D}$$

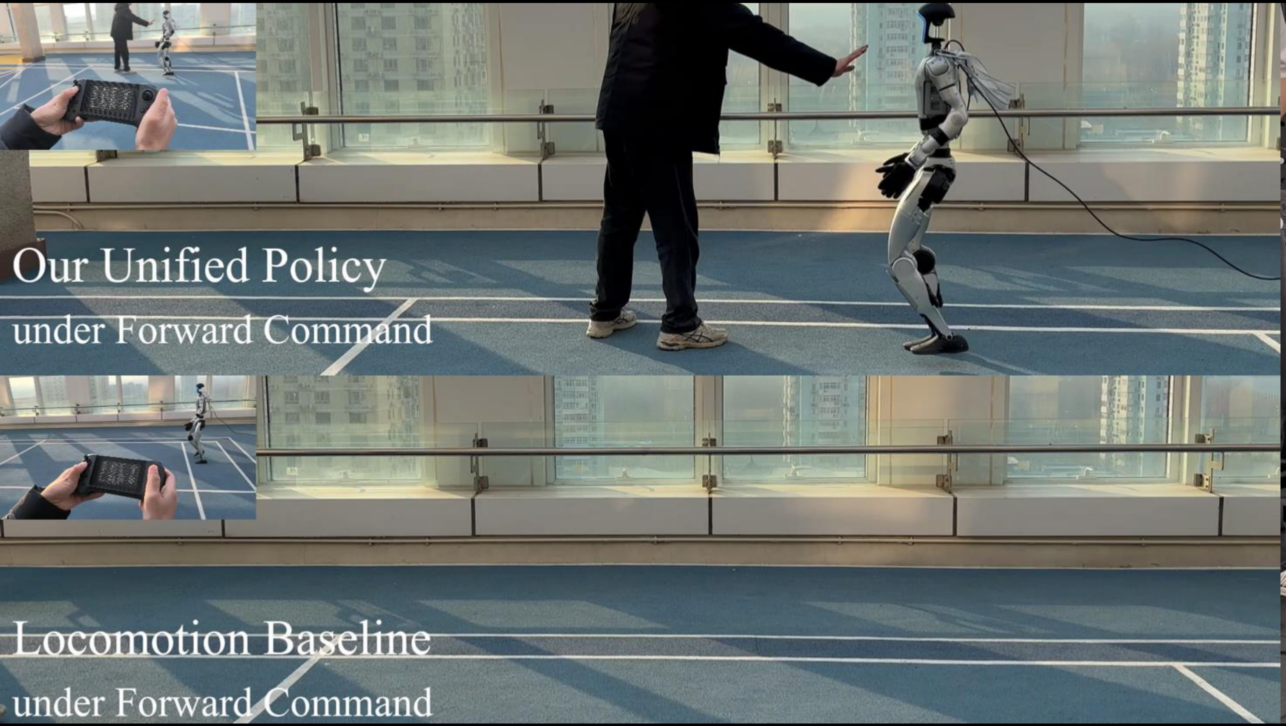

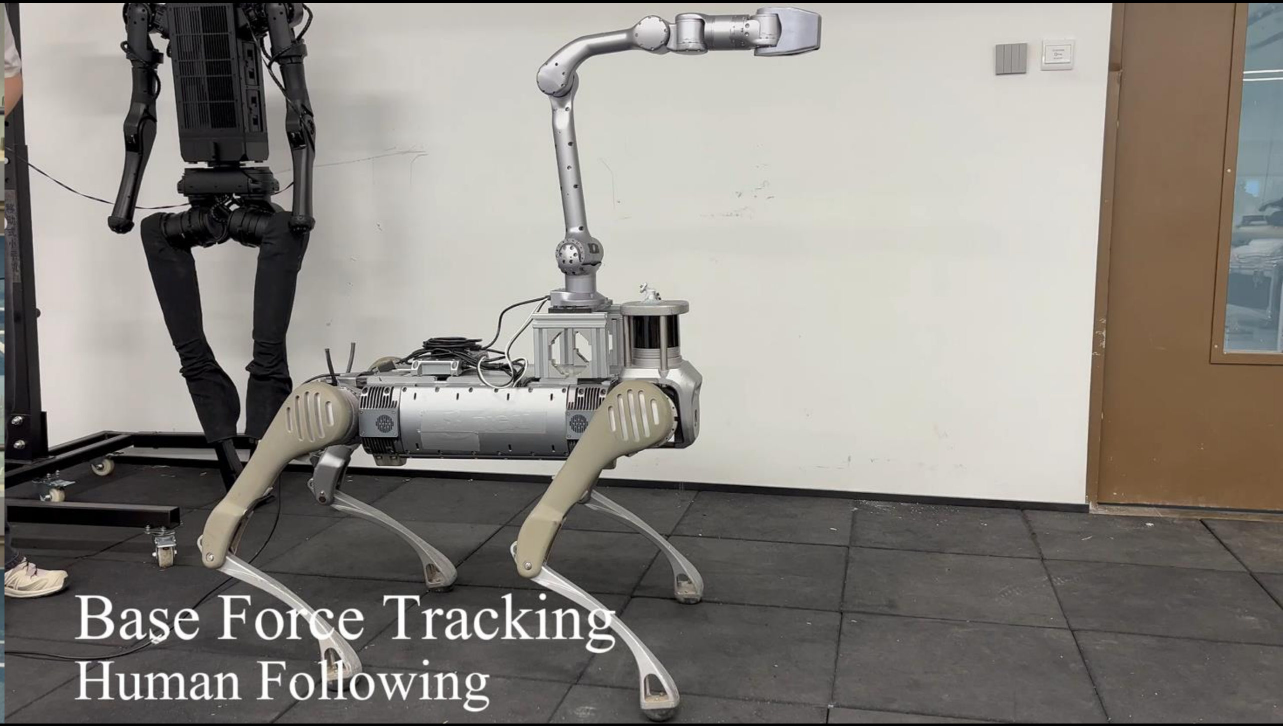

UniFP via RL with force-position sampling in simulator

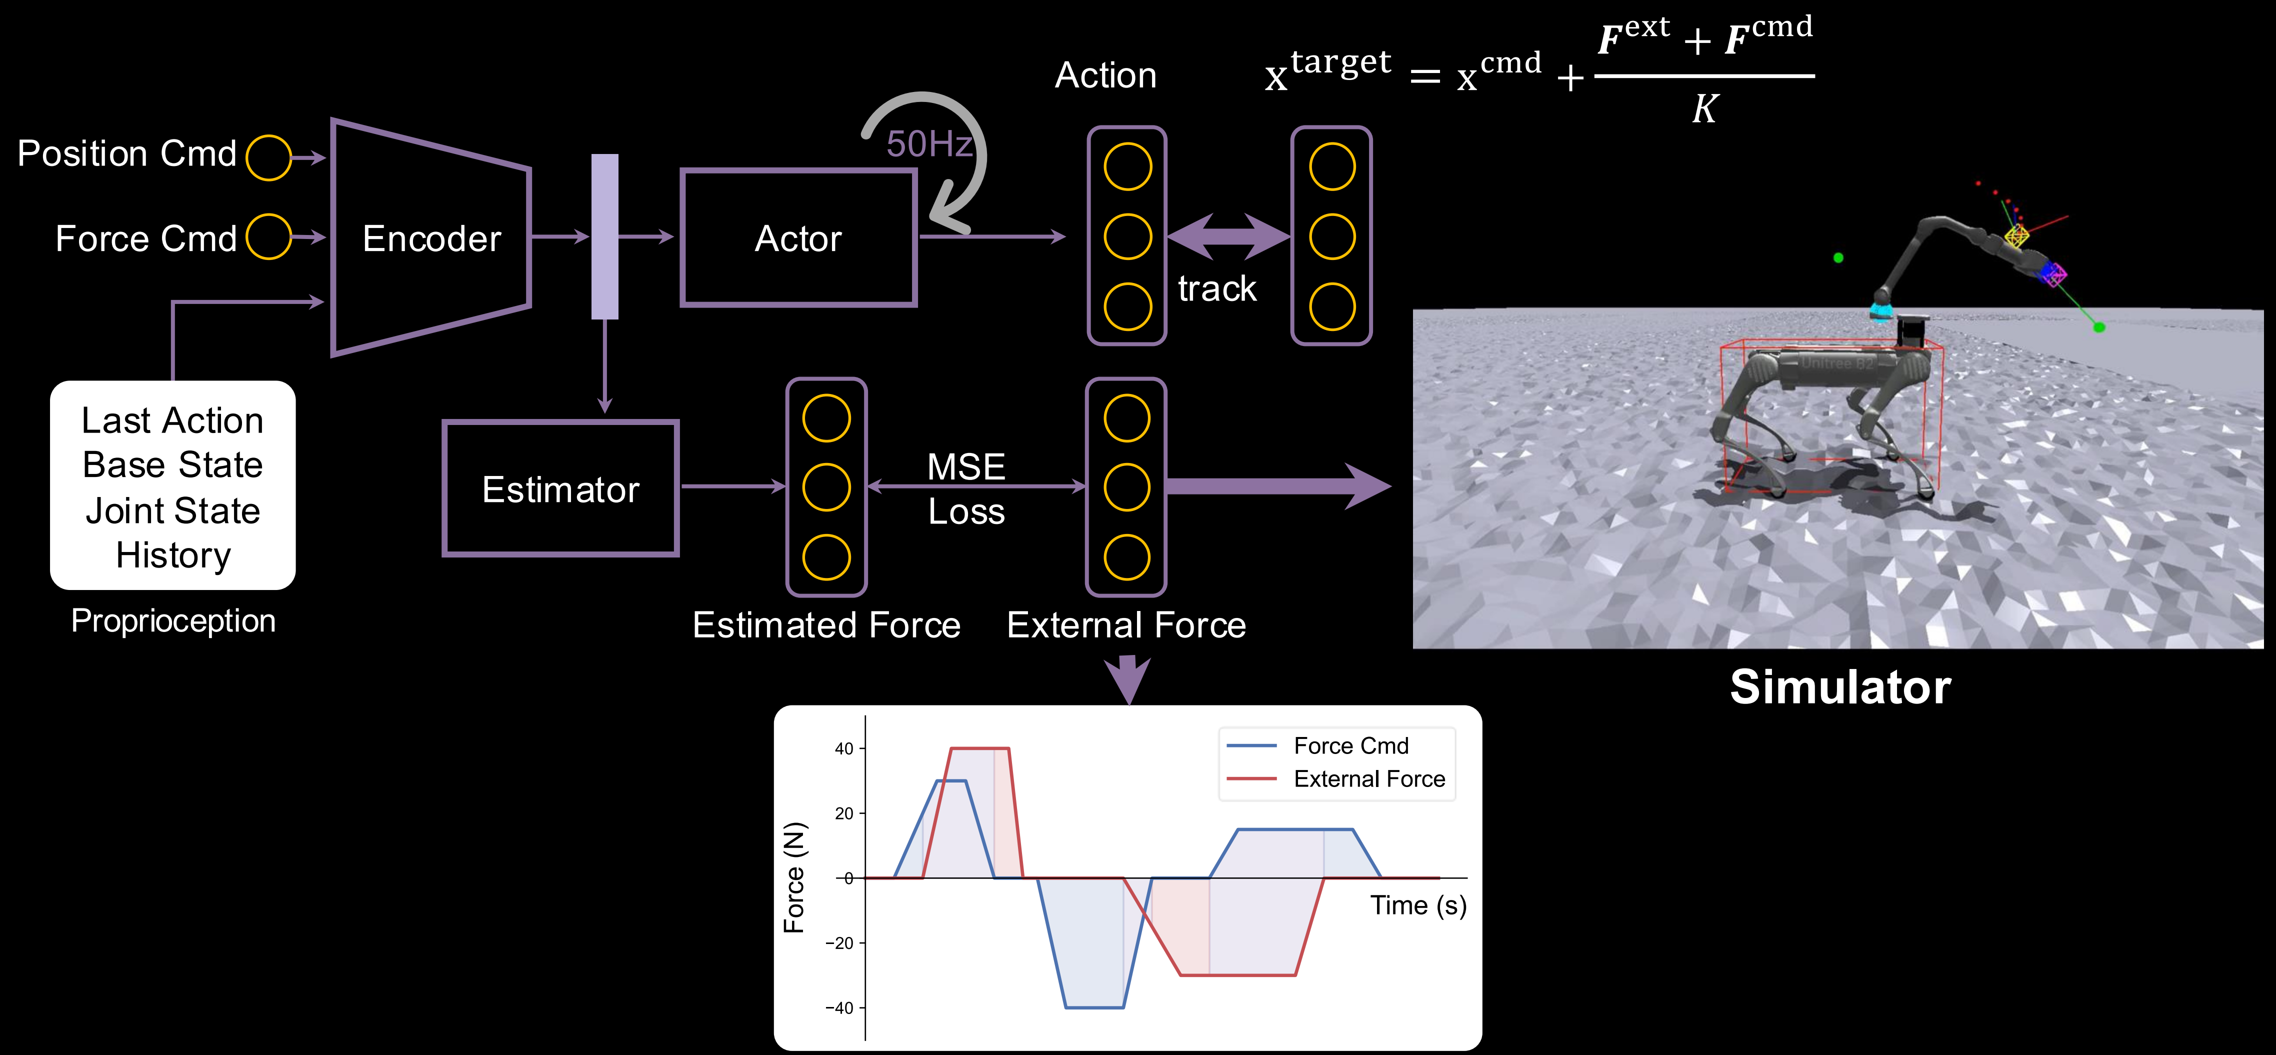

Zhi et al., UniFP: Learning a Unified Policy for Position and Force Control in Legged Loco-Manipulation, CoRL'25 Best Paper

UniFP for force-aware real-world imitation learning

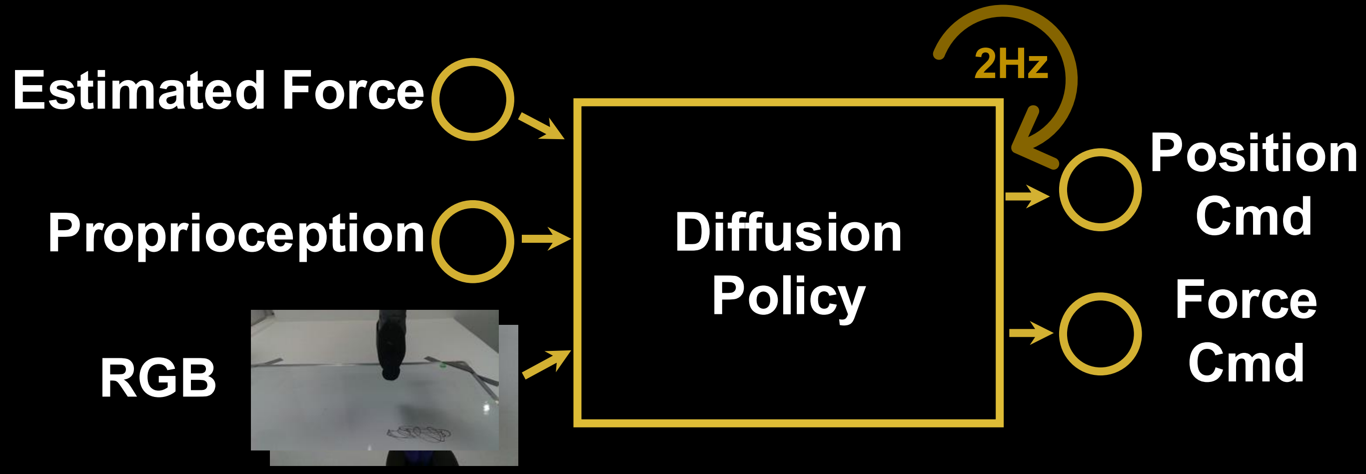

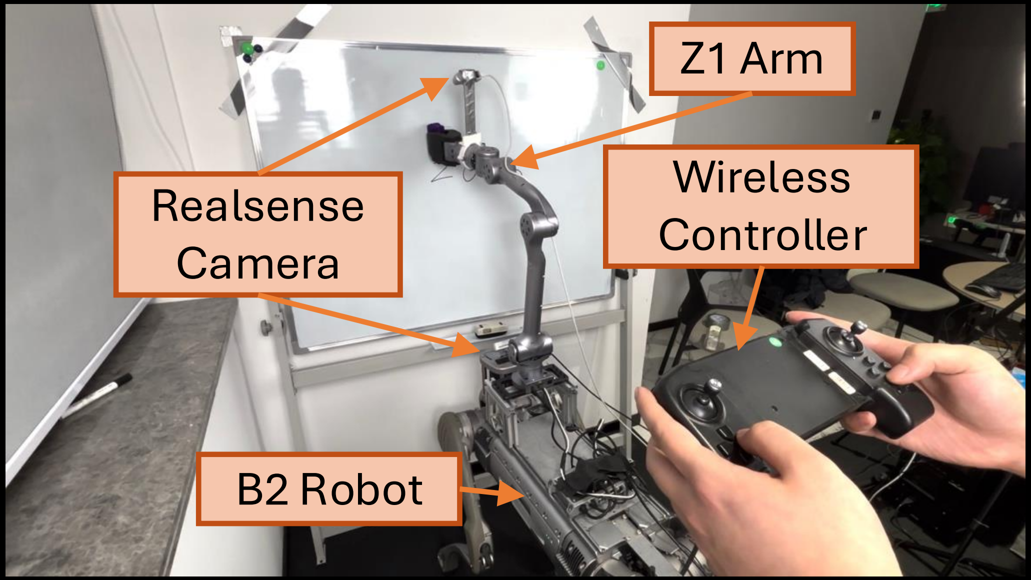

- Data collection with estimated forces
- Imitation learning with position and force command targets
- Inference with UniFP

UniFP for force-aware real-world imitation learning

Tested on 4 tasks with each task taking 50 demonstrations

UniFP for force-aware real-world imitation learning

Table A.3: Imitation learning results (**50 trials per task**)

Task	wipe-blackboard	open-cabinet	close-cabinet	open-drawer-occlusion
w/o Force	0.22	0.36	0.30	0.30
w/ Force	0.58	0.70	0.72	0.76

Achieves ~39.5% higher success rate than the vanilla DP policy


Robot Exercising in Gym
Kettlebell Pulling
Tug-of-war
Throw the trash

So how is this important…

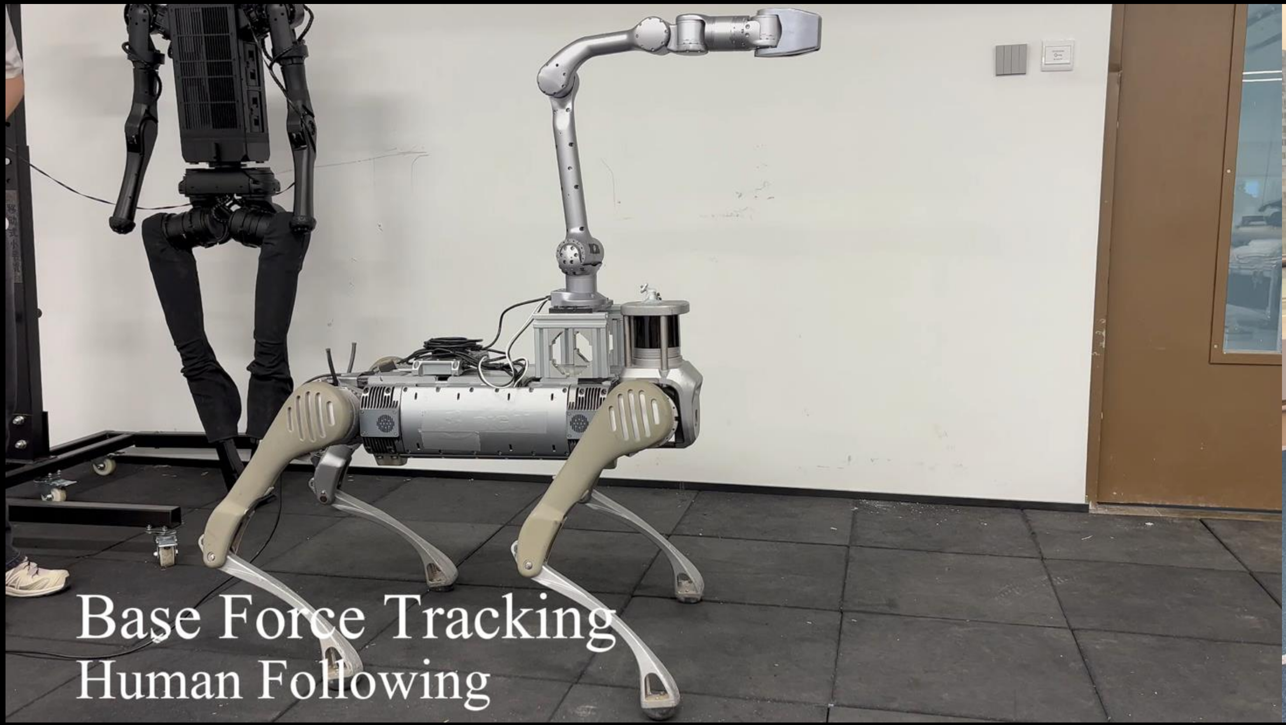

Movement Tracking

Compliant Holding

The composition of these behaviors works for human-robot collaboration


EXIT

COLA for collaborative object carrying

External Forces

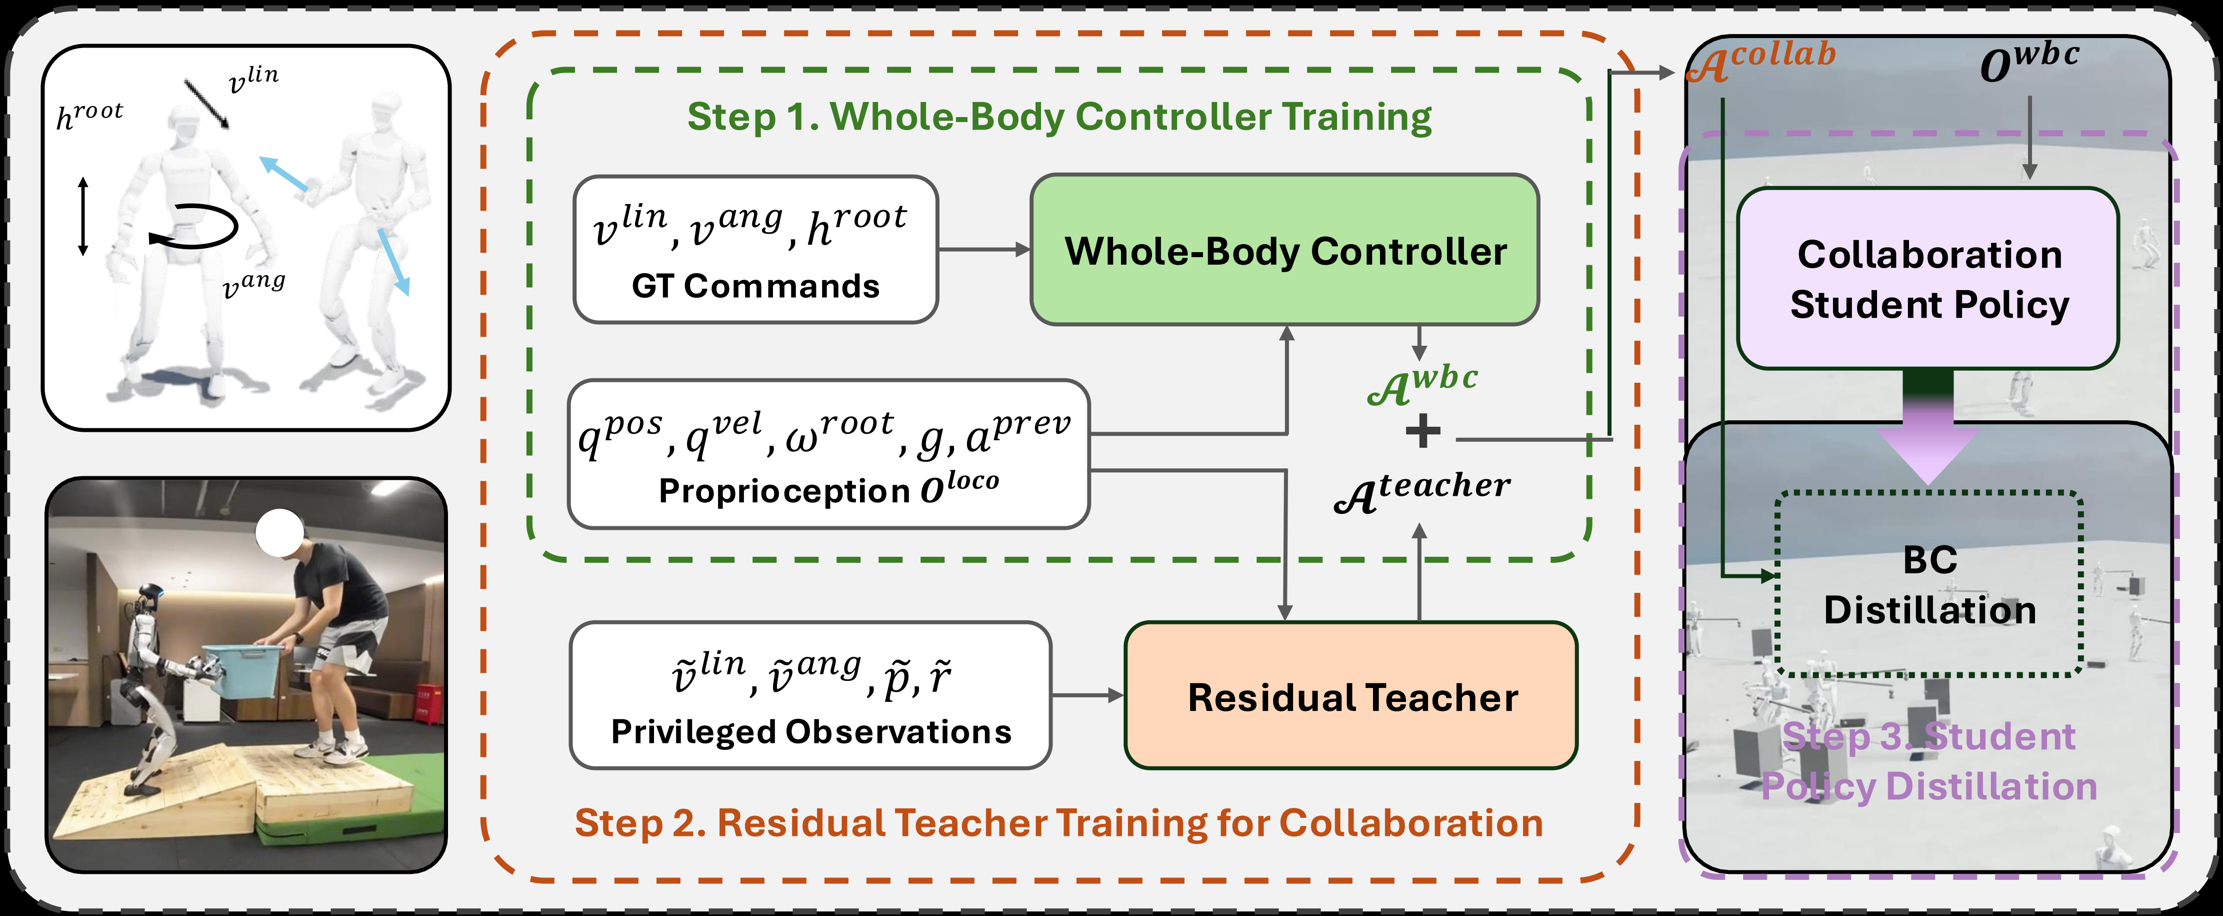

Du et al., Learning Human-Humanoid Coordination for Collaborative Object Carrying, ICRA’26


Speed Tracking, 1.0x Speed
Height Tracking, 1.0x Speed
Speed Tracking (Terrains), 1.0x Speed
Height Tracking, 1.25x Speed


Long Distance Testing (102.4m Total), 1.0x Speed

Speed Tracking, 1.0x Speed

Takeaways

- **Scaling controller learning itself as an important problem**
  - Weigh it at least the same as VLA/WAM learning
  - High-quality data with careful distribution selection
  - Clear hardware specs for sim2real gap solving
  - Use of HOI data for robot learning with perception considered?

Takeaways

- Scaling controller learning itself as an important problem
  - Weigh it at least the same as VLA/WAM learning
  - High-quality data with careful distribution selection
  - Clear hardware specs for sim2real gap solving
  - Use of HOI data for robot learning with perception considered?
- Compliance/force is a must for real-world tasks
  - Force modeling is not the same as domain randomization
  - Compliance for safety and collaboration, force for task execution
  - Use of tactile sensing for dexterous manipulation?

Takeaways

Thank you !

- Scaling controller learning itself as an important problem
  - ❖ Weigh it at least the same as VLA/WAM learning
  - ❖ High-quality data with careful distribution selection
  - ❖ Clear hardware specs for sim2real gap solving
  - ❖ Use of HOI data for robot learning with perception considered?
- Compliance/force is a must for real-world tasks
  - ❖ Force modeling is not the same as domain randomization
  - ❖ Compliance for safety and collaboration, force for task execution
  - ❖ Use of tactile sensing for dexterous manipulation?

At the end, it will still be about autonomous and generalizability

At present, controller+data > model/hardware design