Prediction & Reaction in Motion: Humanoid Control for Non-Inertial Worlds and Athletic AI

Yan Gu

Associate Professor

School of Mechanical Engineering

Purdue University

Where humanoid robotics are heading…

	Classical 1960-	**Learning-based** 2015-	**Emerging** 2023-
The controller considers	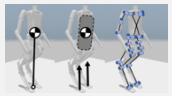		
	Simplified and complex dynamics models	Physics-based simulation Specific limited real-world data	Rich real-world dynamics Internet-scale data
Tools	Reduced-order modeling Control theory Convex optimization	GPU-parallel simulation Reinforcement learning Imitation learning	Foundation models World models Diffusion models
Drawing inspiration from	Control Optimization Neuro/biomechanics	Machine learning Computer graphics	Natural language processing Computer vision
Tasks			
	Stable locomotion on restrictive terrains	Agile locomotion on various terrains	Multi-task loco-manipulation in the open world

Gu, et al, **Evolution of Humanoid Locomotion Control**, *Science Robotics*, 2026.

Two Extremes on the Path to Humanoids in a Non-stationary Human World

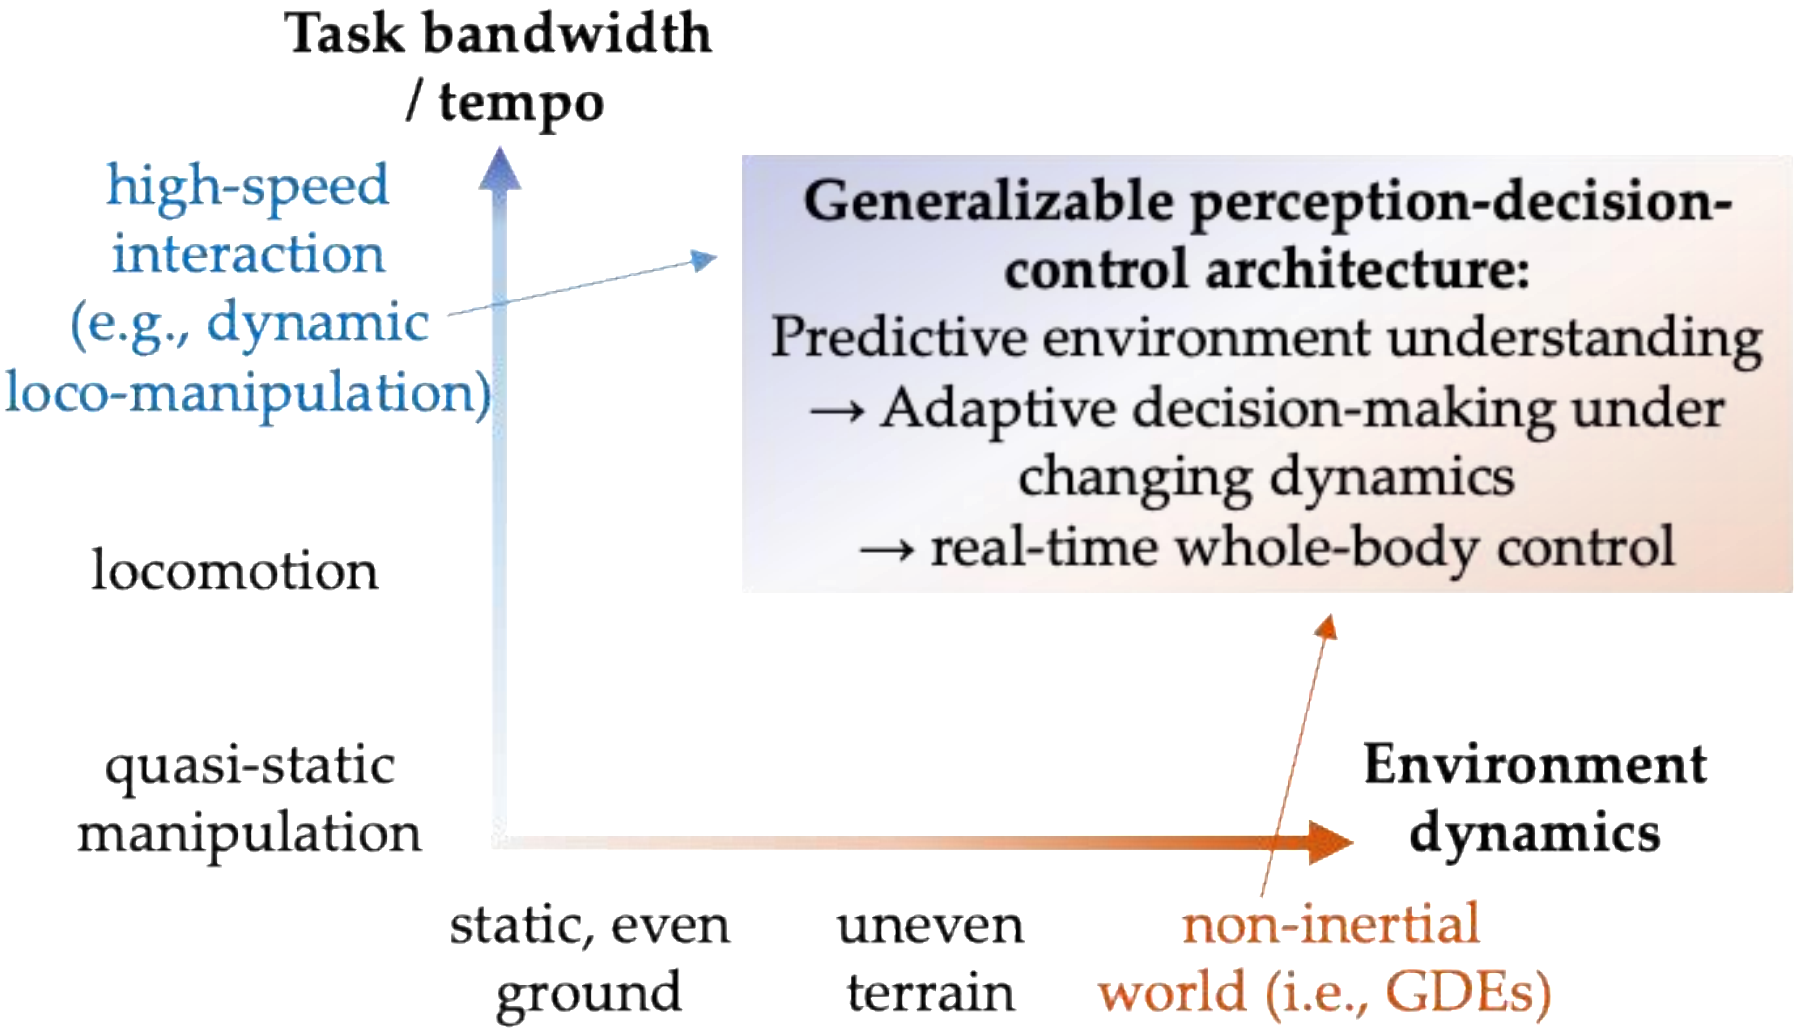

Research goal

Develop perception and decision-making approaches that enable legged robots to perceive, adapt, and act in complex dynamic environments.

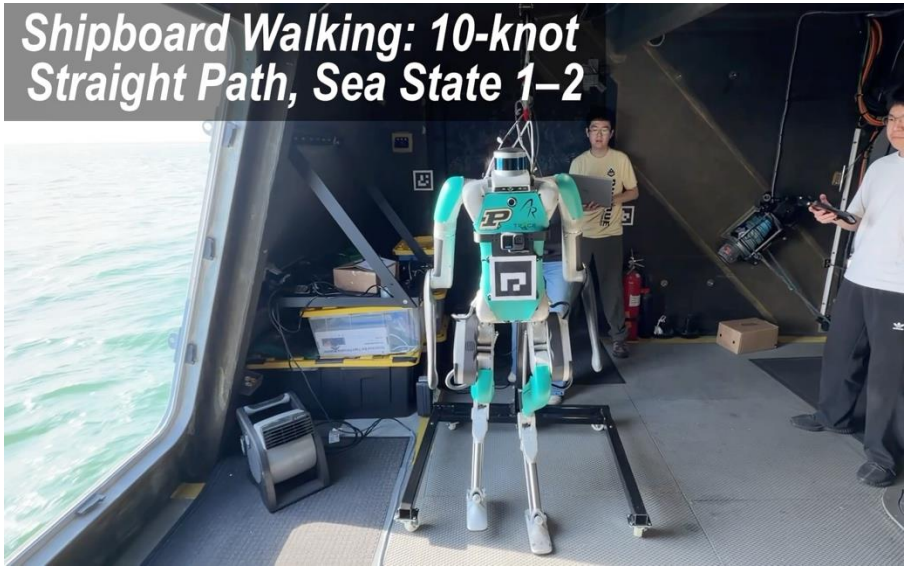

Non-inertial world

Humanoids walk and handle boxes aboard a moving Navy vessel, with stability guarantees from our control theory.

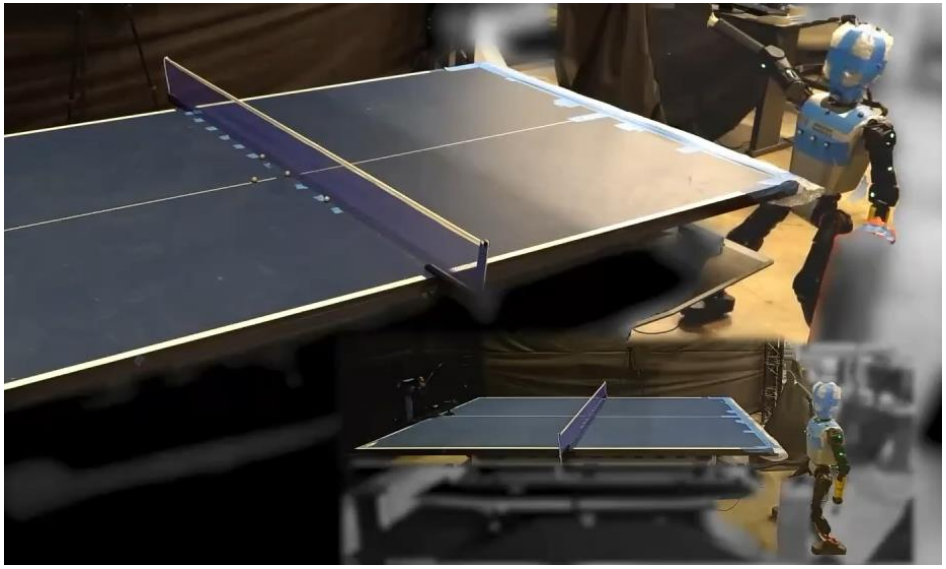

Athletic humanoids

PACE, a humanoid that plays table tennis. Physics-based prediction feeds learned control.

Humanoids in a Non-inertial World

Performance on a stationary ground



vs.

Performance gaps of commercial systems in GDEs

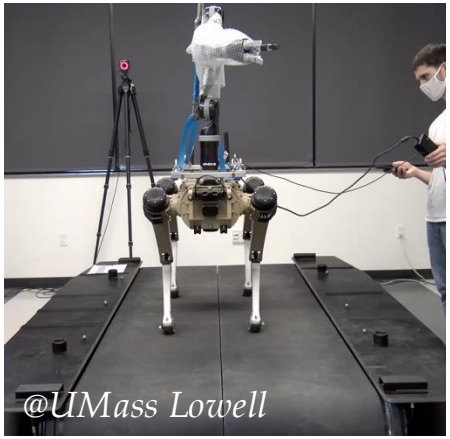

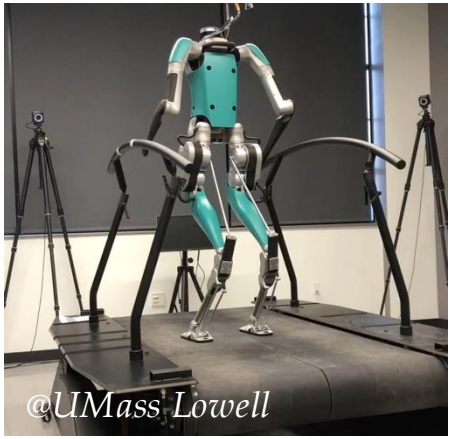





***Why** do robot dynamics **differ** between static and **globally moving (i.e., non-inertial)** environments?*

***How to** ensure **reliable** robot operations under ground motions?*

(S. Misenti, B. Hertel, B. Weng, R. Donald, A. P. Jawaji, M. Kosoko-Thoroddsen, J. G. Trafton, A. Norton, R. Azadeh, Y. Gu, *IJIRA*, 2025)

Challenges

- Unknown, unpredictable, and persistently and rapidly changing ground motions
- Complex (hybrid, nonlinear, high-dimensional, time-varying) robot dynamics
- Limited sensing (e.g., lack of GPS and inertial landmarks)

Humanoids in a Non-inertial World

Proprietary controller

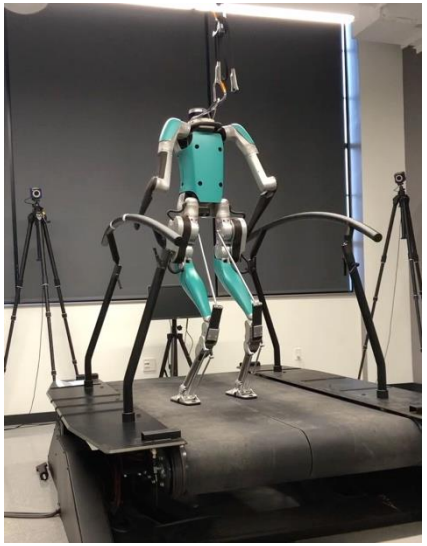

Failure to sustain stability (assumes *static* ground)

vs.

Our previous discrete footstep controller (2023-2025)

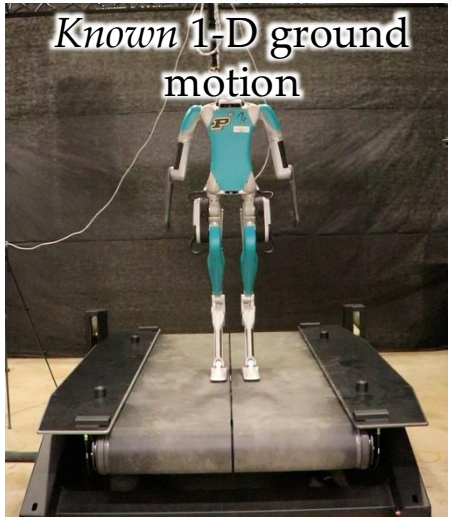

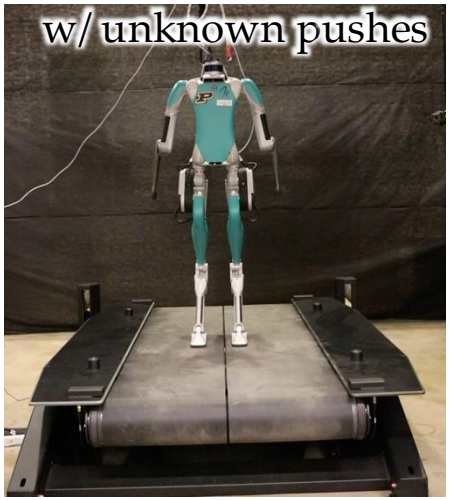

(Y. Gao, V. Paredes, Y. Gong, Z. He, A. Hereid, Y. Gu, IEEE TRO 2025)

Adaptive ankle control (2025-2026)

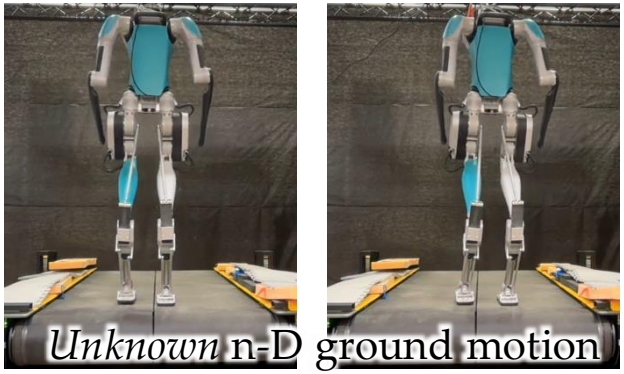

(2025)

- Extension to general system dynamics
- Real-world shipboard tests

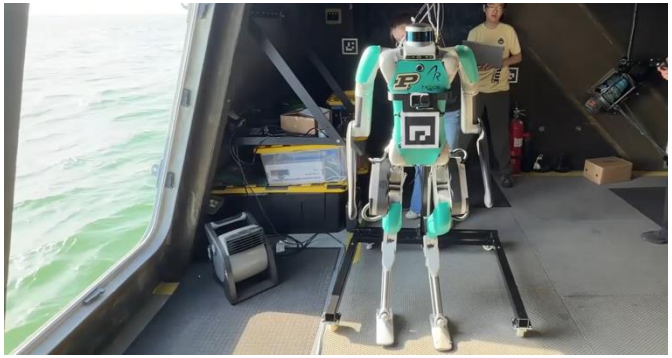

(2026)

Humanoids in a Non-inertial World

Adaptive ankle control under unknown, complex ground motion

Challenges:

- *Unknown, unpredictable,* and *persistently and rapidly changing* ground motions
- *Hybrid, nonlinear, high-dimensional, time-varying* robot dynamics

Key novelty:

- Extension of *adaptive ankle* torque control to *general hybrid systems*

Approach:

- Formulation of a *hybrid*, *time-varying* robot model with *ankle control*

$$\begin{cases} \begin{bmatrix} \dot{x}_{sc} \\ \ddot{x}_{sc} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \dfrac{g + \ddot{z}_{ws}(t)}{z_{sc}} & 0 \end{bmatrix} \begin{bmatrix} x_{sc} \\ \dot{x}_{sc} \end{bmatrix} - \begin{bmatrix} 0 \\ \ddot{x}_{ws}(t) \end{bmatrix} - \begin{bmatrix} 0 \\ \dfrac{\tau_a}{m z_{sc}} \end{bmatrix}, \quad t \neq t_k \\ \begin{bmatrix} \Delta x_{sc} \\ \Delta \dot{x}_{sc} \end{bmatrix} = \begin{bmatrix} -u_x \\ 0 \end{bmatrix}, \quad t = t_k \end{cases}$$

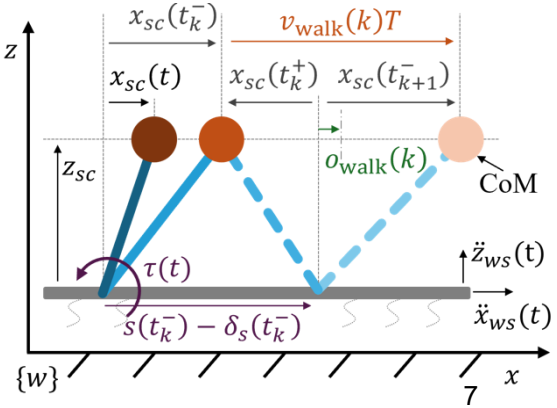

(I.C. Chang, J. Stewart, P.A. Ioannou, Y. Gu, ACC 2025 & manuscript under review)

Humanoids in a Non-inertial World

Adaptive ankle control under unknown, complex ground motion

Approach:

- Adaptive control law for *general hybrid linear* systems with *hybrid, unknown* sinusoidal disturbances

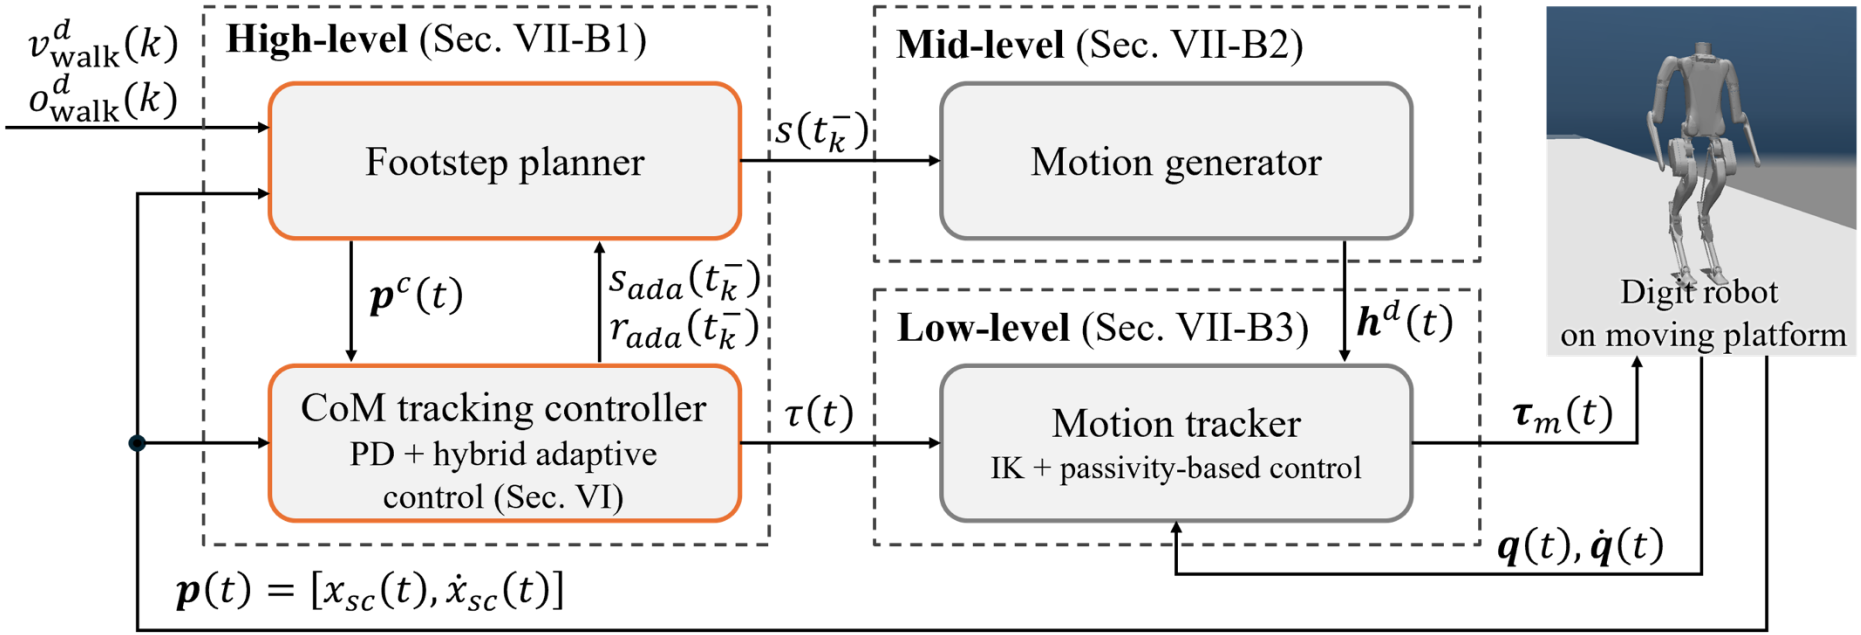

(I.-C. Chang, J. Stewart, P.A. Ioannou, Y. Gu, ACC 2025& manuscript under review)

Humanoids in a Non-inertial World

In-lab Experimental Validation during Treadmill Pitch and Sway

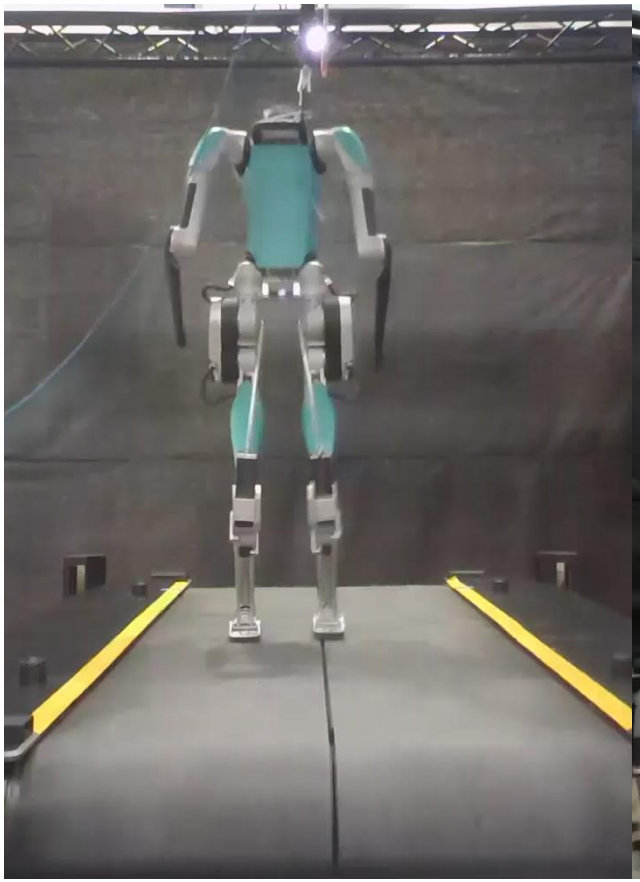

Pitch = $4° \sin(\frac{2\pi}{8} t)$

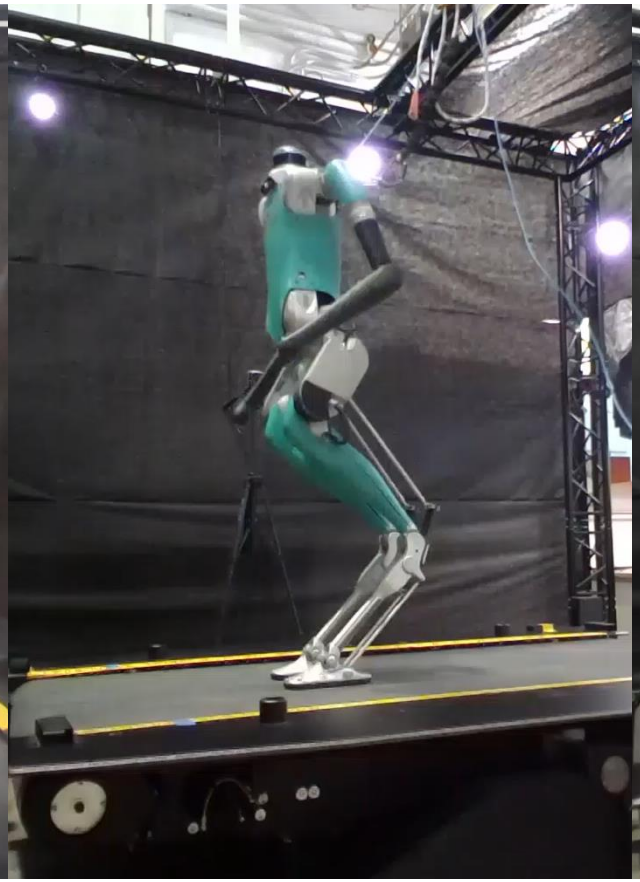

Sway =$5 \sin(\frac{2\pi}{8} t)\ cm$

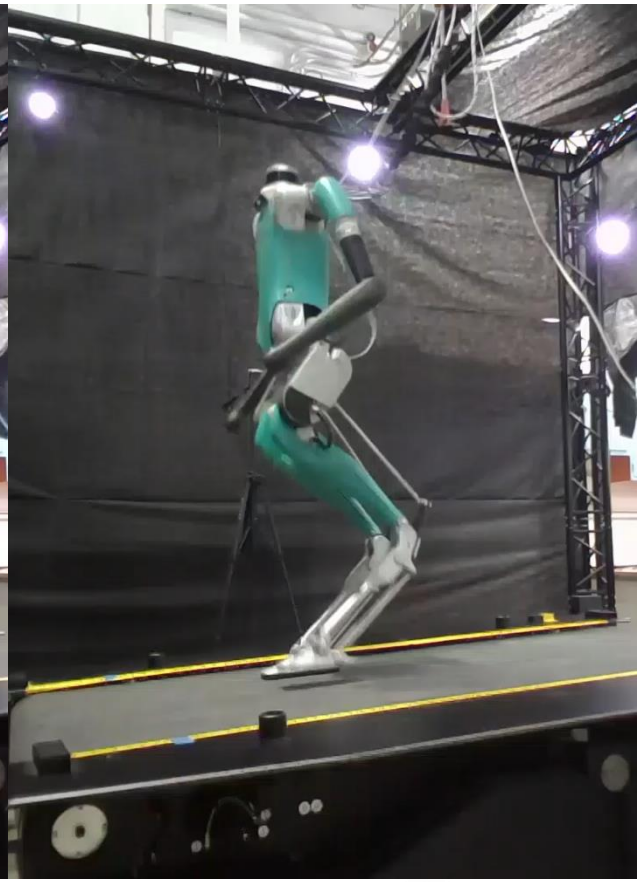

Pitch = $3° \sin(\frac{2\pi}{8} t)$

Sway =$3 \sin(\frac{2\pi}{10} t)\ cm$

***Robust** walking under various **unknown**, **multi-directional** ground motions*

Humanoids in a Non-inertial World

Shipboard Field Test

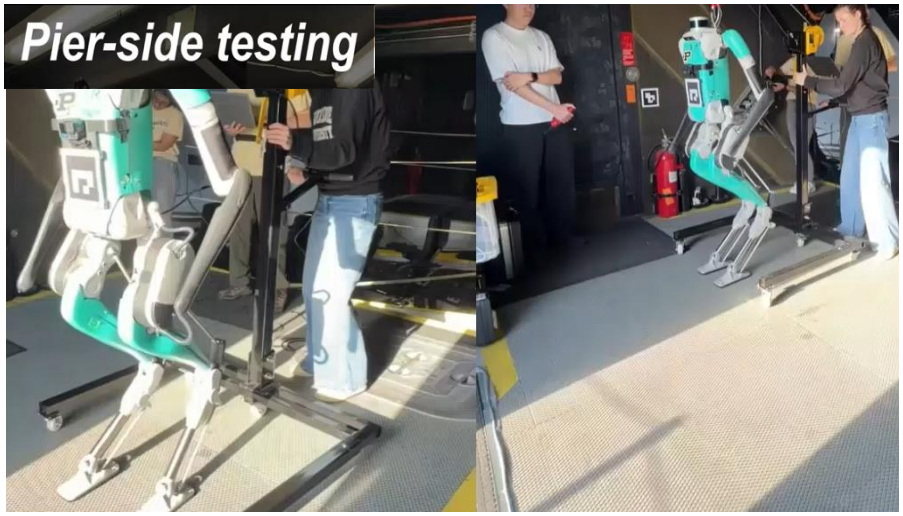

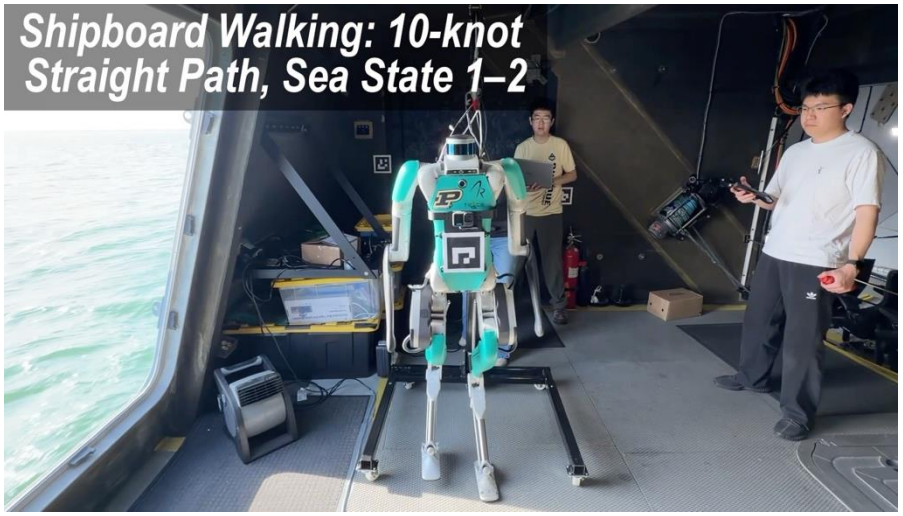



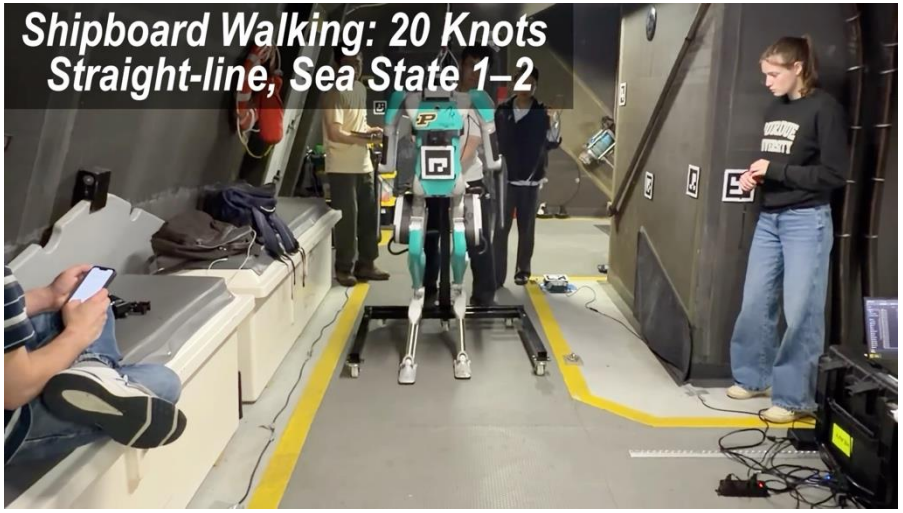

(The gantry is used only for fall protection and cable management; it does not provide balance or weight support)

Two Extremes on the Path to Humanoids in a Non-stationary Human World

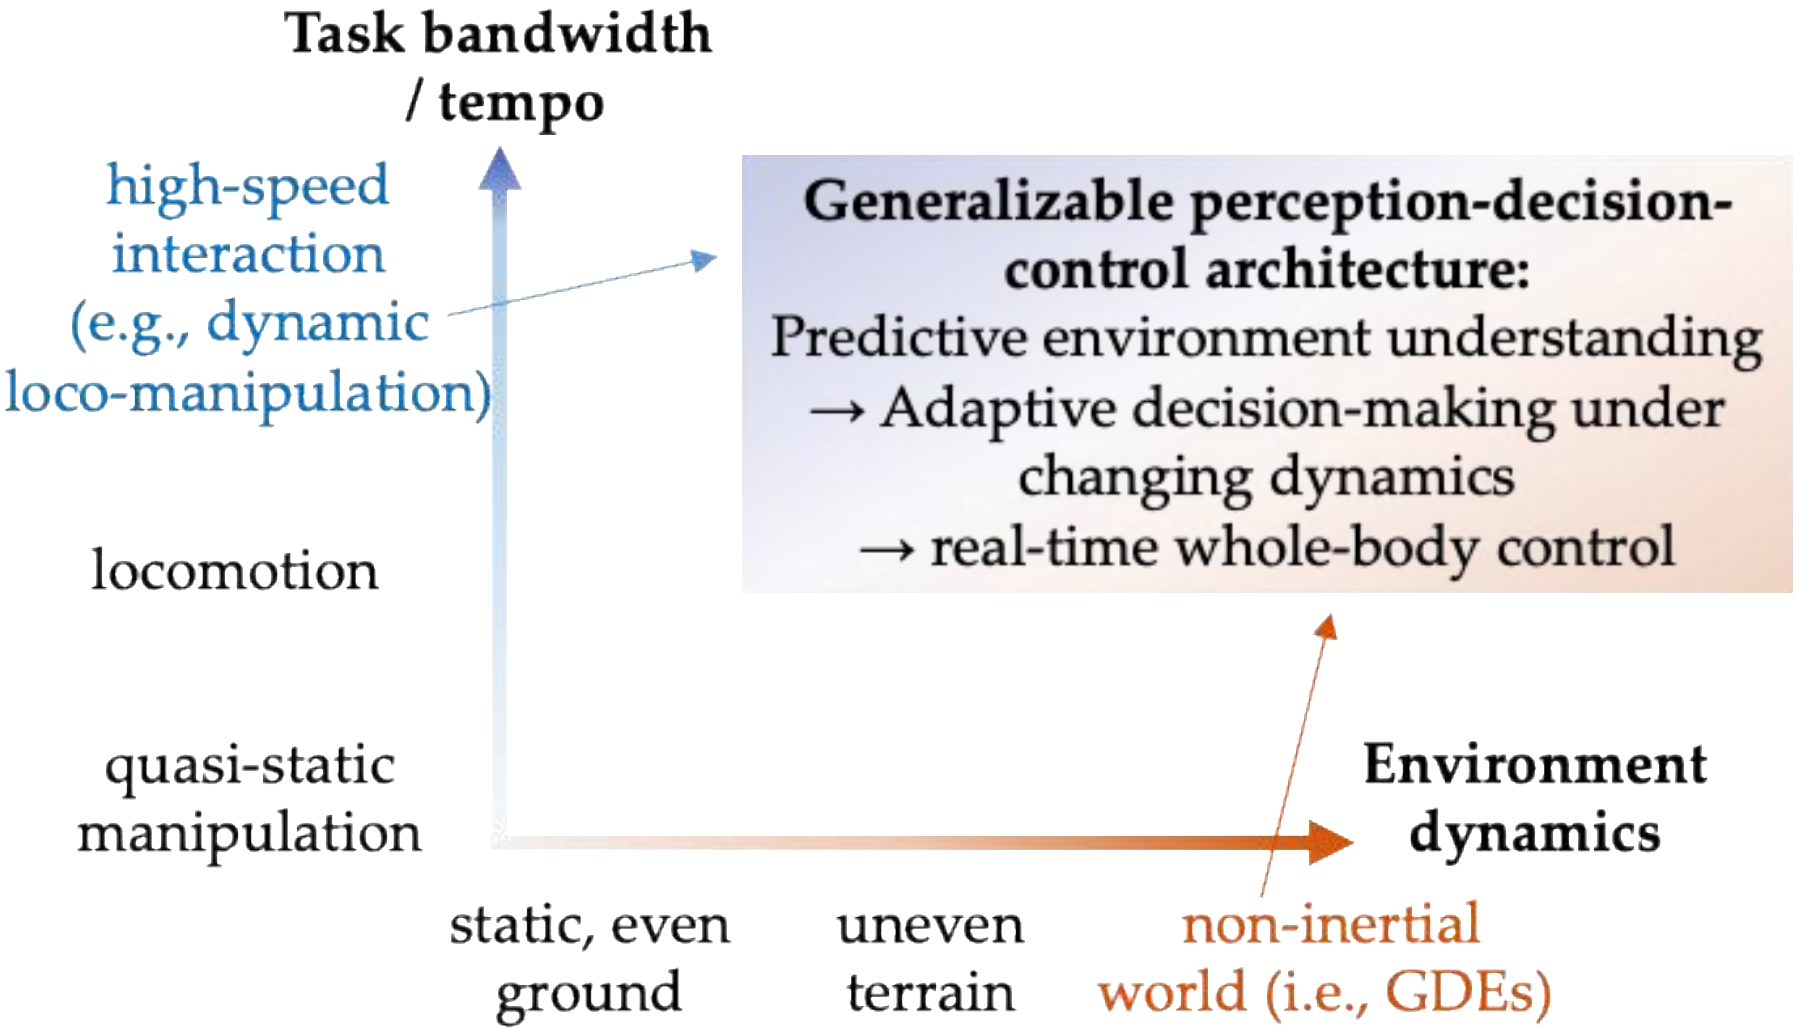

Athletic Humanoids

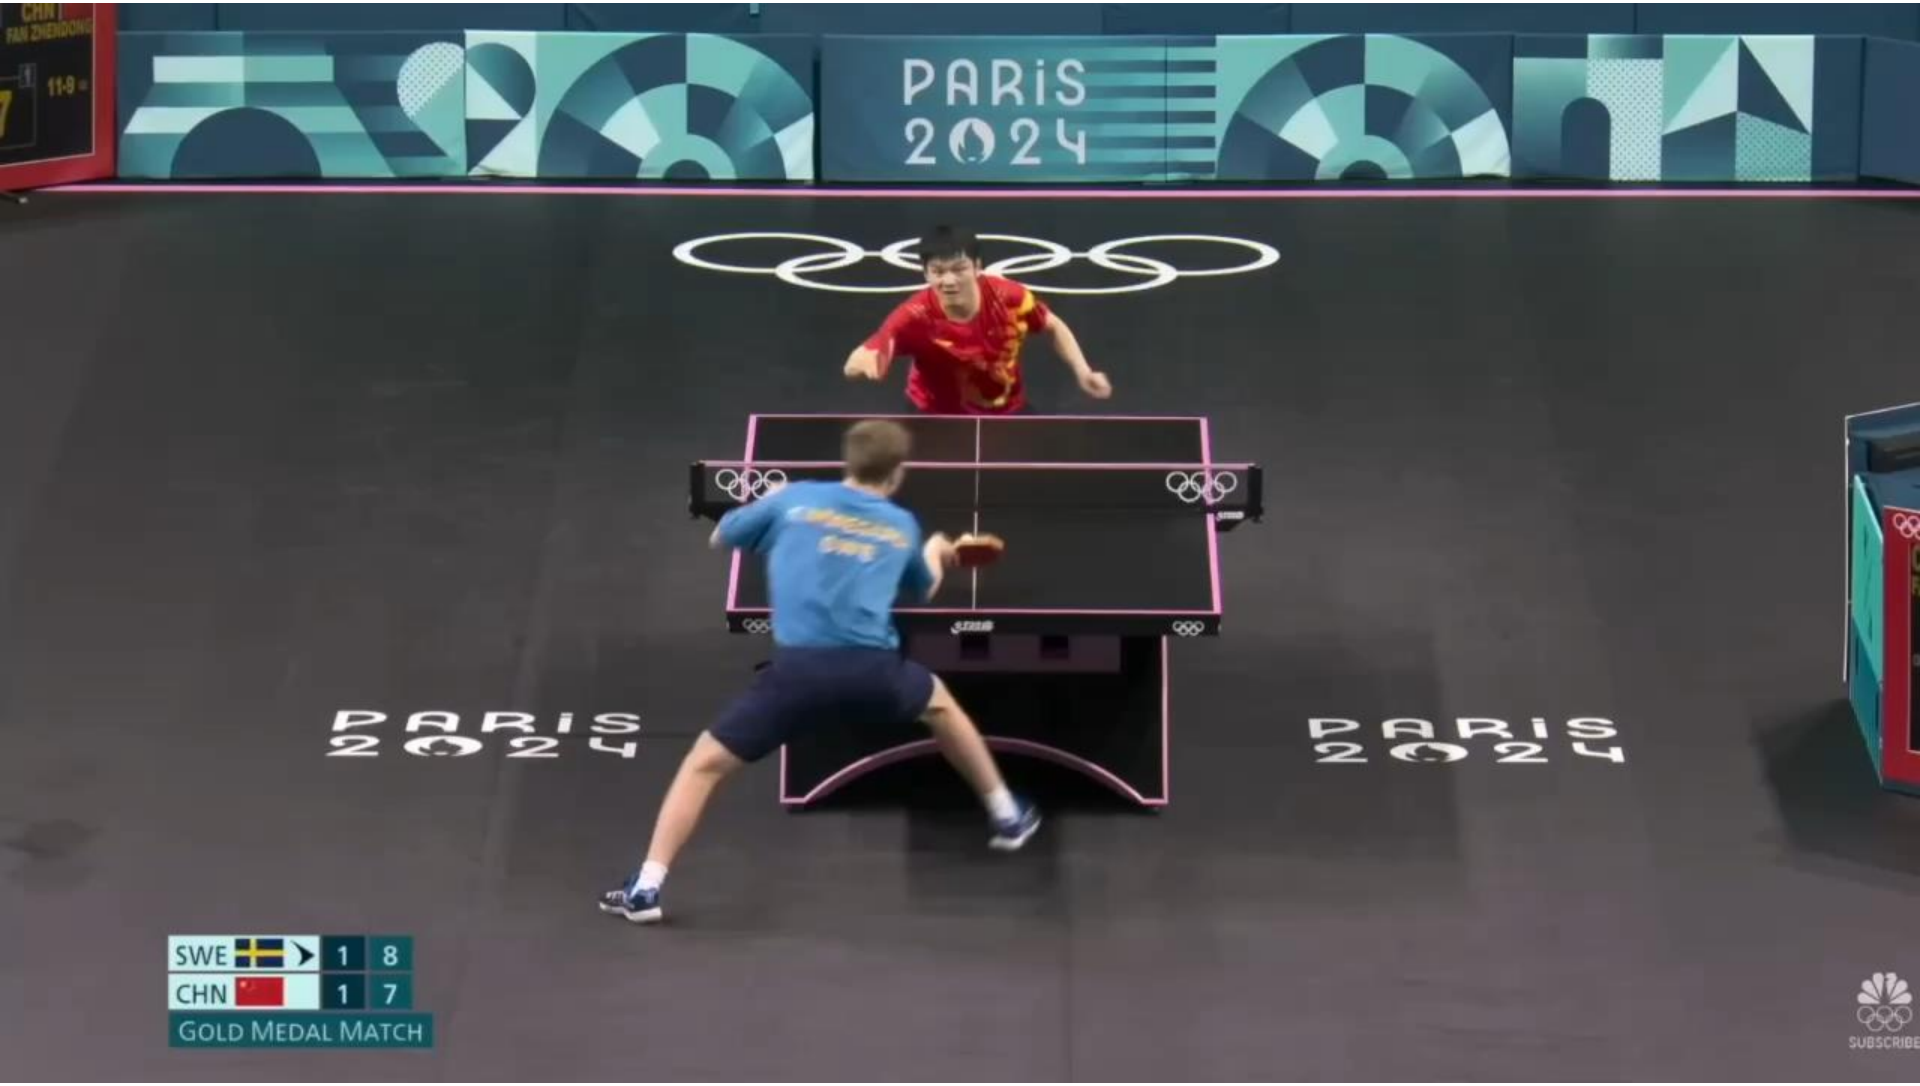

(https://www.youtube.com/watch?v=cN14JgUV2ag)

Table tennis is fast "physical" chess

involving fast, contact-rich , whole-body sports skills

Proactively making contact with ball and ground

Fast, precise perception of **fast-moving** ball and opponent

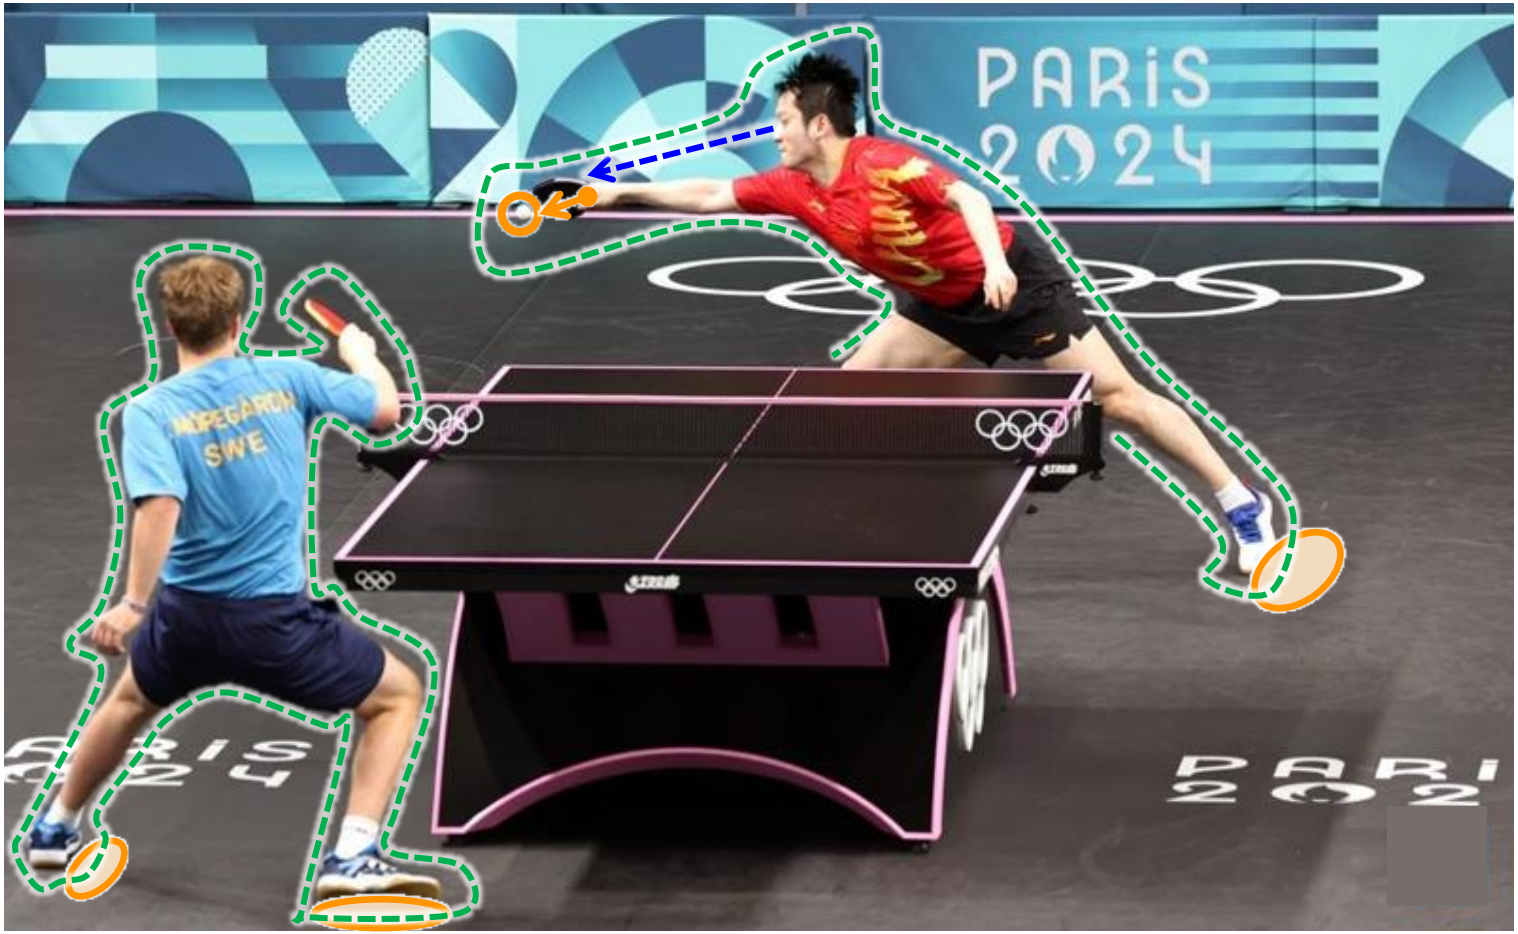

Highly coordinated and dynamic **whole-body** movement (**arm strike + foot work**)

Challenges in mastering table tennis

- **Demanding performance goals:**
  - **Fast (~200 ms)** and **precise** perception and decision making
- **Complex physics involving humans, objects, and environment:**
  - **Rich contact** (foot-ground, ball-paddle, ball-table)
  - **Fast ball** motion→ **highly dynamic, whole-body player** movement
  - **High dimensional** action space, **inherently unstable** locomotion

We introduce an end-to-end RL framework for whole-body humanoid table tennis.

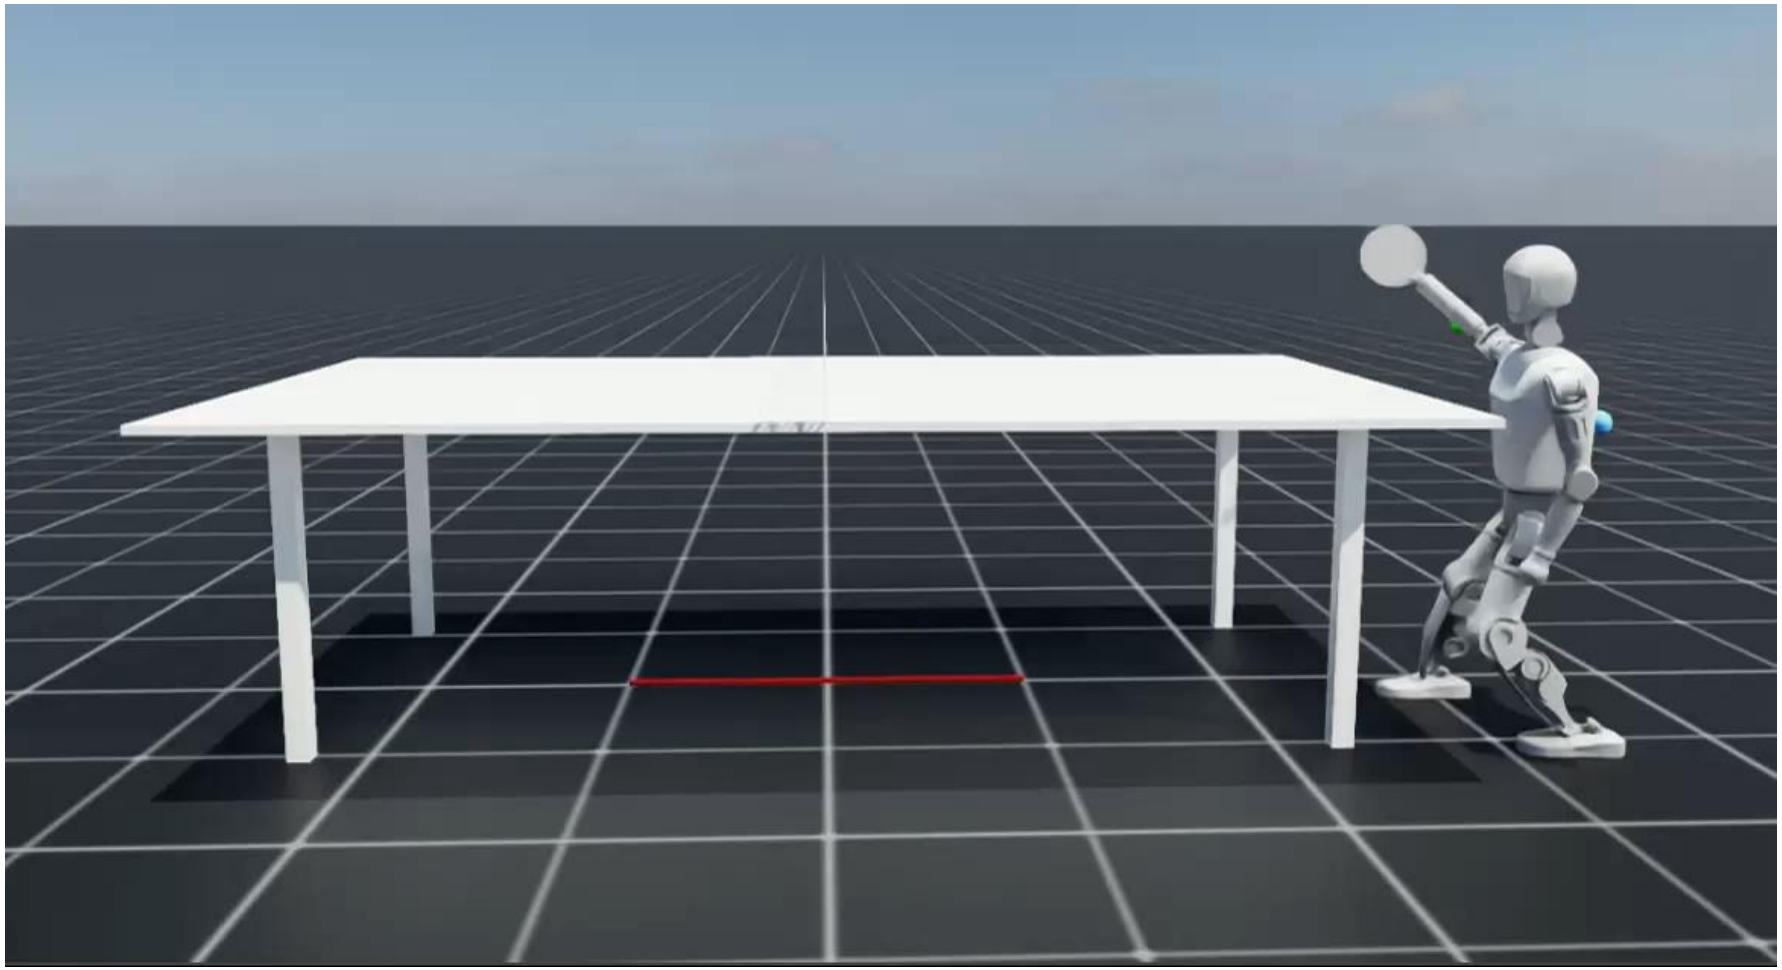

M. Hu*, W. Chen*, W. Li*, F. Mandali**, Z. He**, R. Zhang**, P. Krisna**, K. Christian, L. Benaharon, D. Ma, K. Ramani, Y. Gu, ICRA 2026

Problem formulation of End-to-end RL

- **Objective:** To design end-to-end RL for whole-body humanoid table tennis with dynamic, coordinated footwork and arm strike

Observations

MoCap



Base position ($\mathbf{p}_{\text{base}} \in \mathbb{R}^3$)
Ball position ($\mathbf{p}_{\text{ball}} \in \mathbb{R}^3$)
Robot heading vector ($\mathbf{e}_{xy} \in \mathbb{R}^2$)

IMU/Gyro

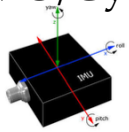

Base angular velocity ($\boldsymbol{\omega}_{\text{base}} \in \mathbb{R}^3$)
Projected gravity vector ($\mathbf{g}_{\text{base}} \in \mathbb{R}^3$)

Encoder

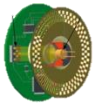

Joint positions ($\mathbf{q} \in \mathbb{R}^{21}$)
Joint velocities ($\dot{\mathbf{q}} \in \mathbb{R}^{21}$)

Buffer



Previous action ($\mathbf{a}_{\text{last}} \in \mathbb{R}^{21}$)

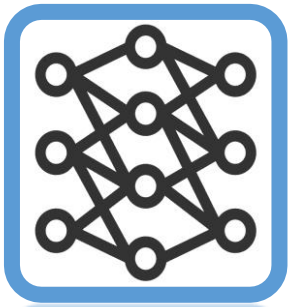

Actor Policy

21 joint targets

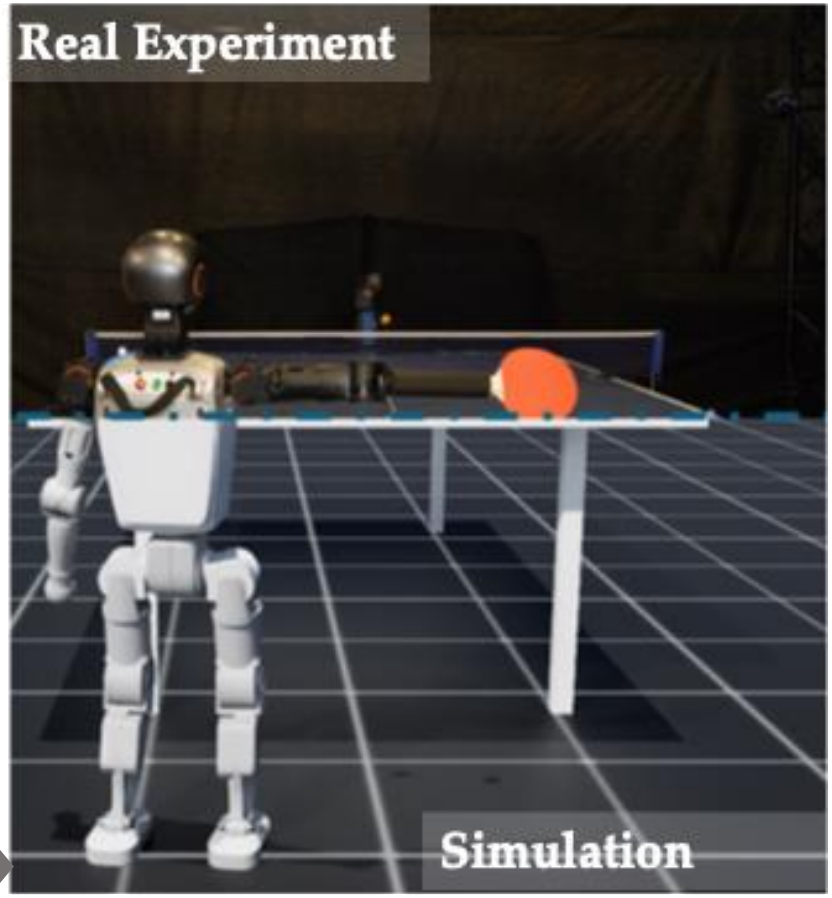

Challenges: Sparse, binary, delayed rewards related to rich contact

Ball hitting the paddle

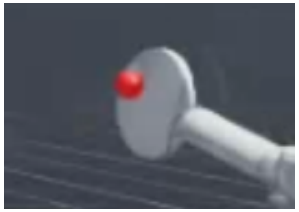

Sparse, binary reward

Successful ball return to other side

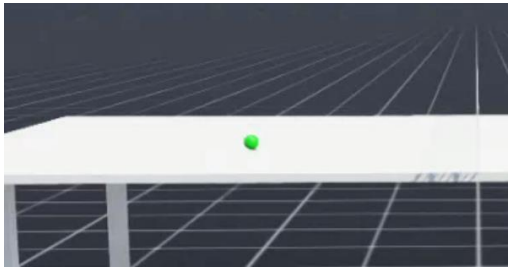

Sparse, binary, delayed reward

PPO fails to explore an effective hitting strategy with such rewards!

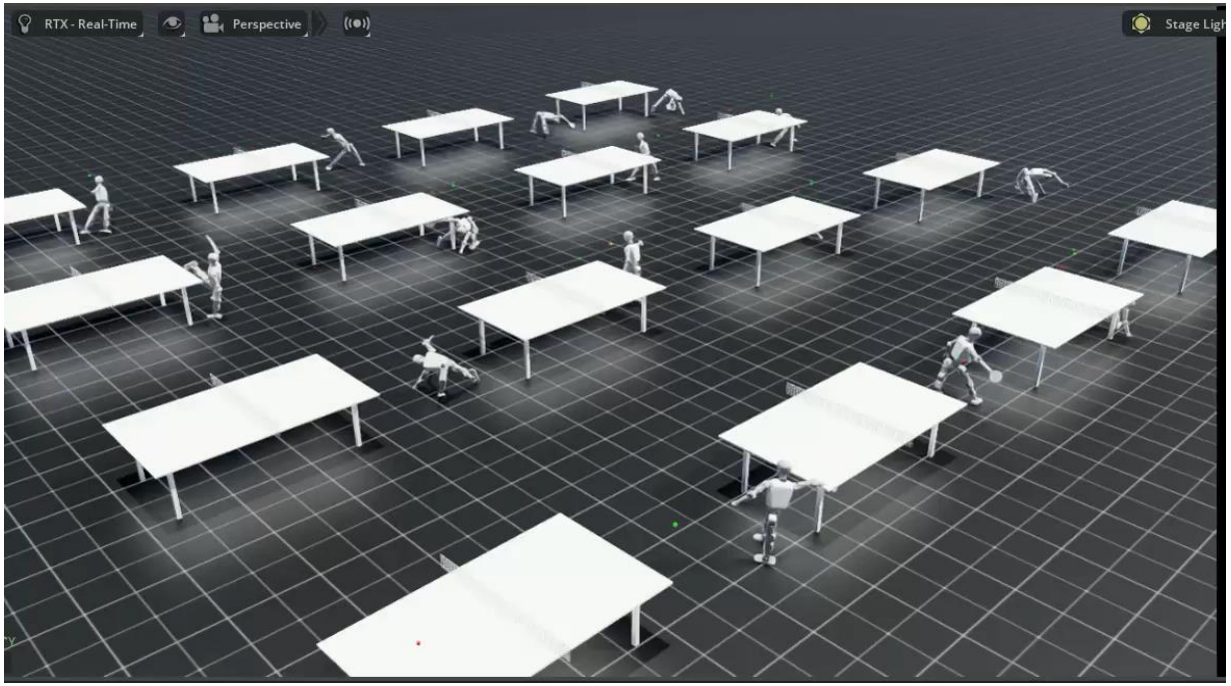

Methodology: Impact-aware RL with prediction augmentation

- **Key idea 1: Lightweight, learned prediction for augmenting actor**
  - Input: **Actual ball position**
  - Ground truth: **Physics-based prediction** (ball aerodynamics, ball-table impact)
  - Output: **Ball apex point after bounce** → **ball-paddle hitting point**
    - Relaxes virtual hitting-plane assumption → versatile motion
    - Included in actor policy's observation → **promote *proactive* action**

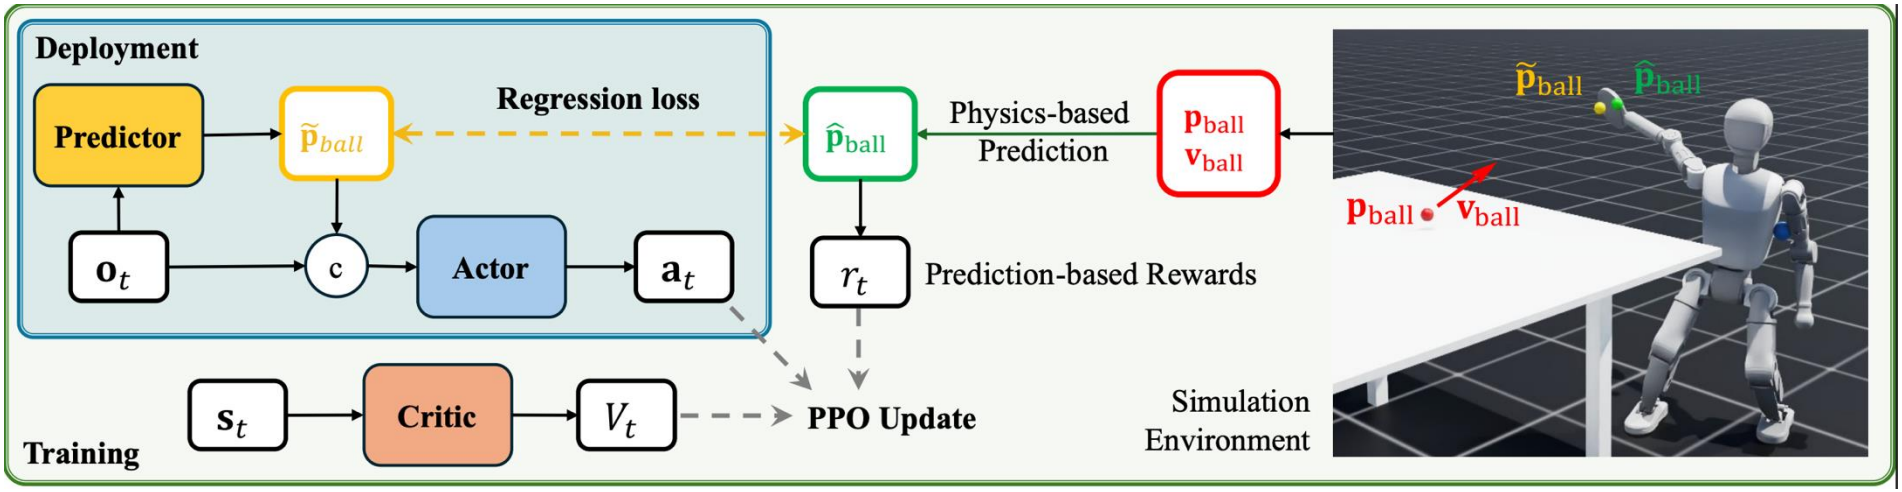

Methodology: Impact-aware RL with prediction augmentation

- **Key idea 2: Impact-aware, prediction-based reward shaping**
  - **Before ball-paddle impact:**
    - *Physics-based* prediction of ball-paddle hitting position

      → Converting **sparse, binary** to **dense, continuous reward**

      → *Proactive footwork* + *paddle preparation*

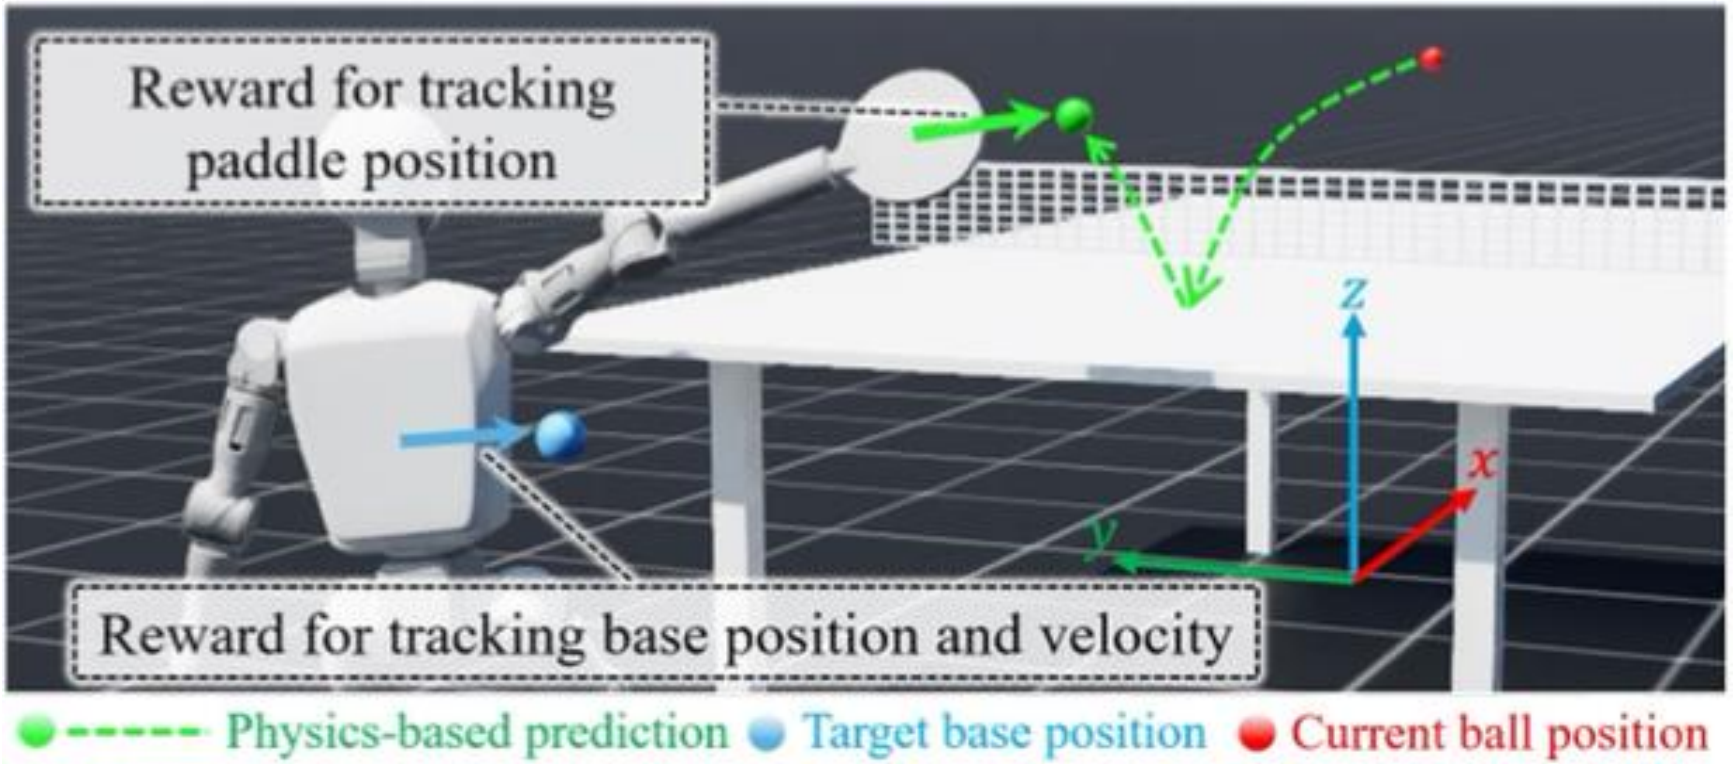

Methodology: Impact-aware RL with prediction augmentation

- **Key idea 2: Impact-aware, prediction-based reward shaping**
  - **After ball-paddle impact:**
    - ***Physics-based*** prediction of landing position and net-clear margin

      → Converting **binary, sparse, delayed** to **continuous, dense, current** reward

      → ***High hitting / success rate*** with ***fast*** ball speed (~6 m/s)

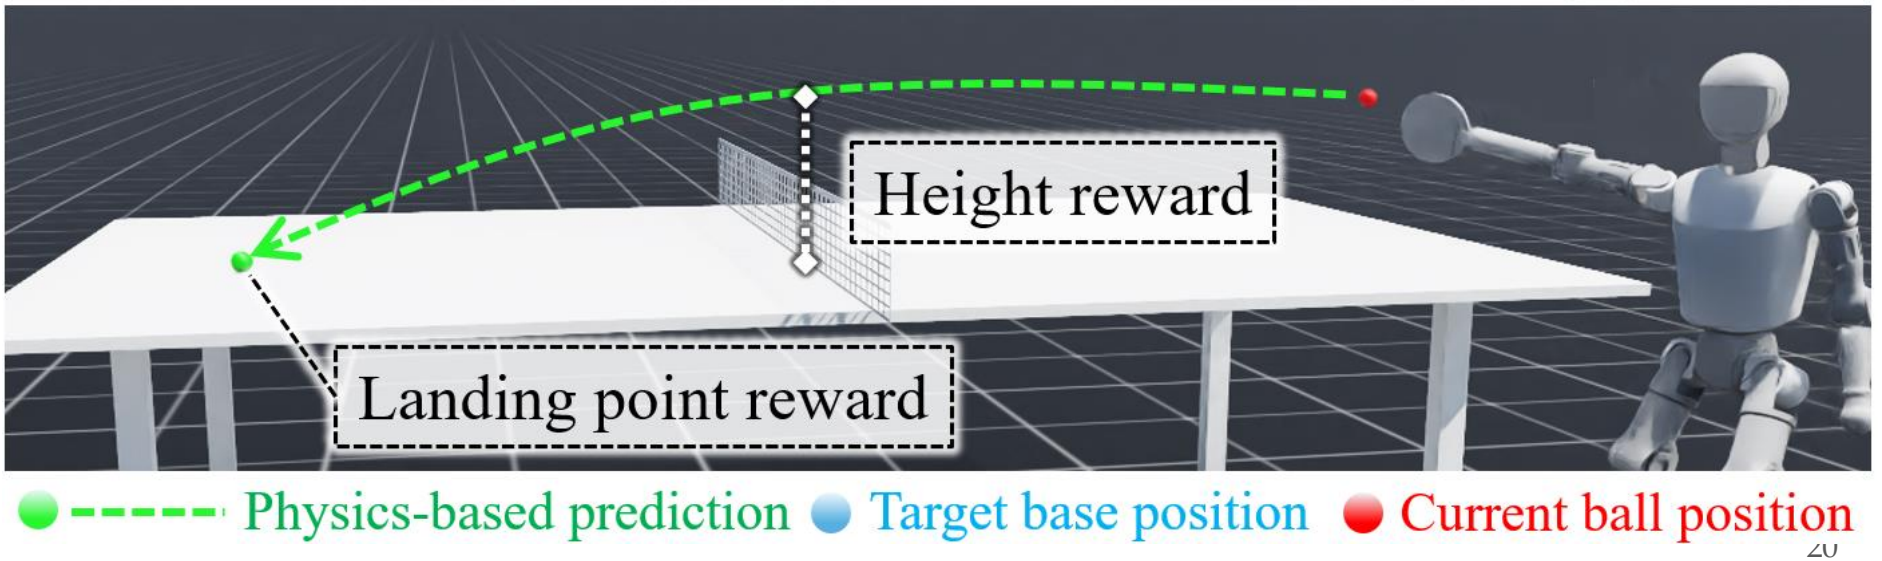

Training

Parallel Simulations in Issac Lab

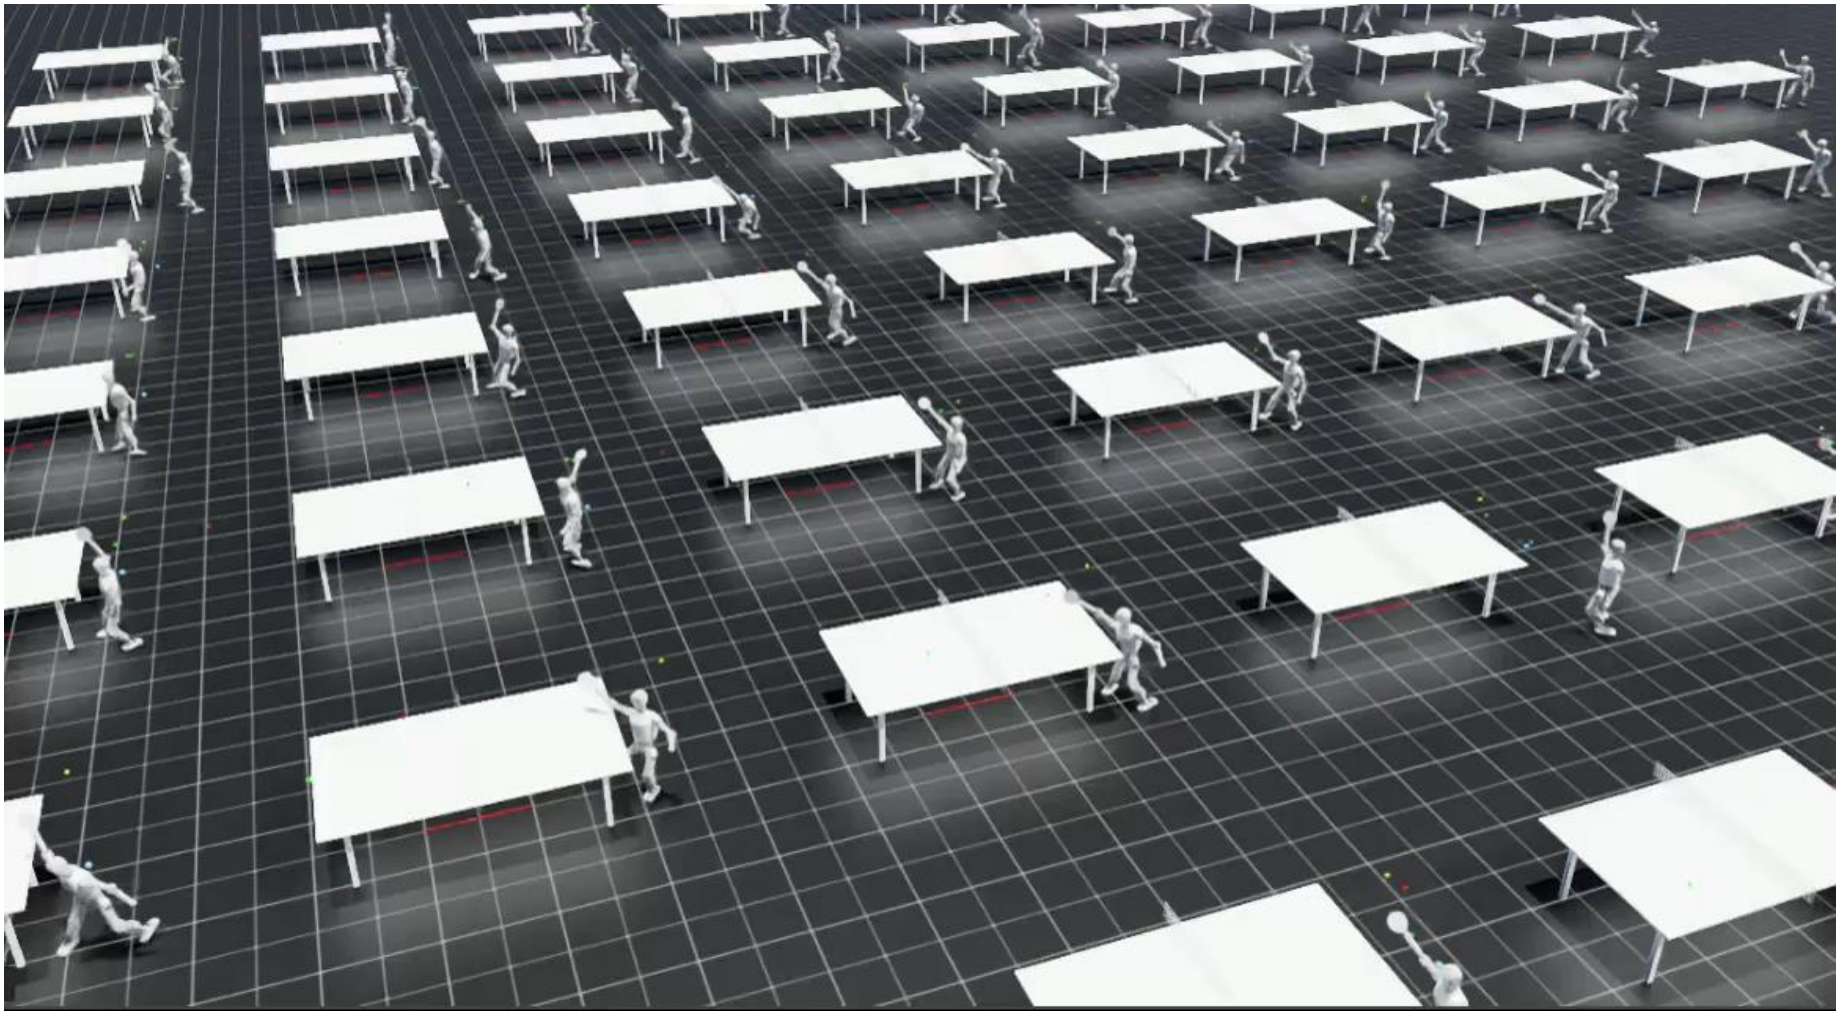

Training Results: Ablation study

- **Without the learned predictor**, PPO cannot learn proactive whole-body motion to return successfully.
- **Without the dense, continuous reward**, sparse and binary rewards are insufficient to help the robot learn an effective return strategy.

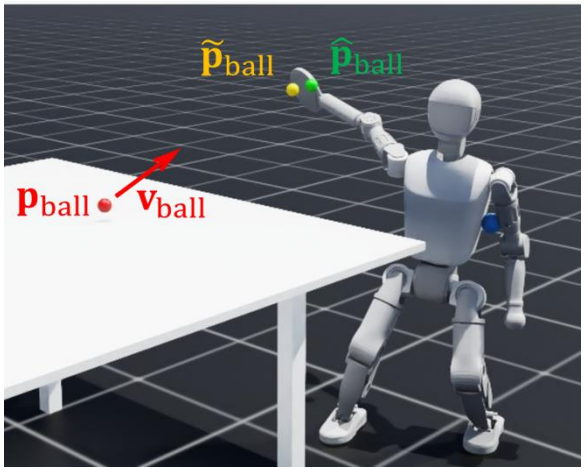

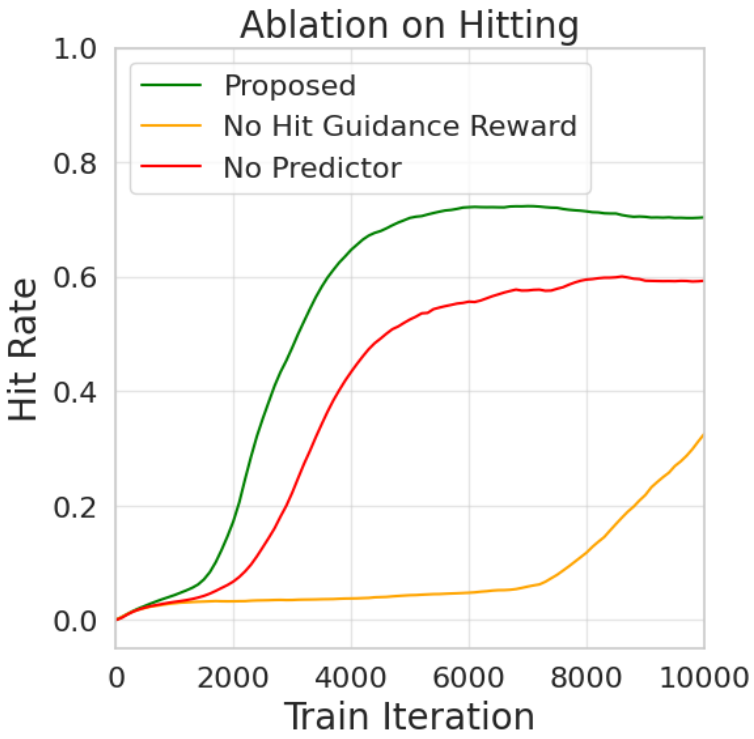

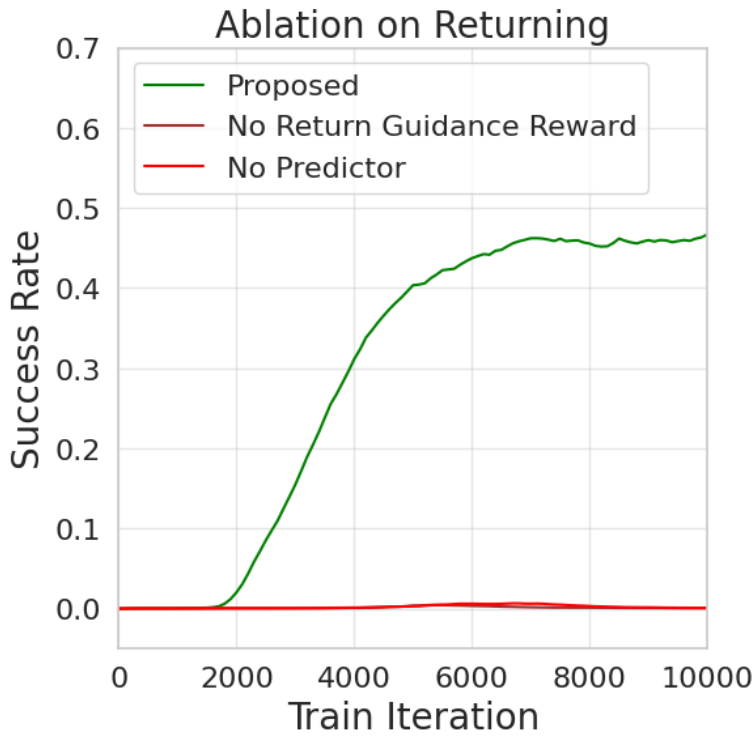

Training Results: Diverse striking points

- By eliminating the reliance on the **virtual hitting plane**, our work enables the emergence of diverse striking behaviors.

Classical treatment

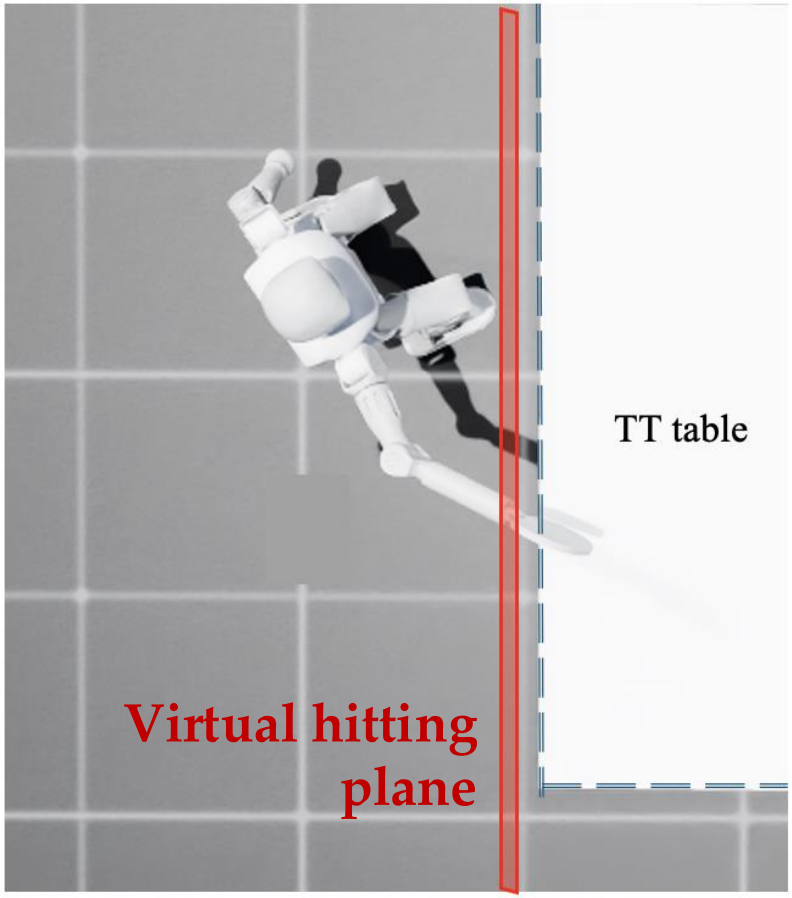

Our work

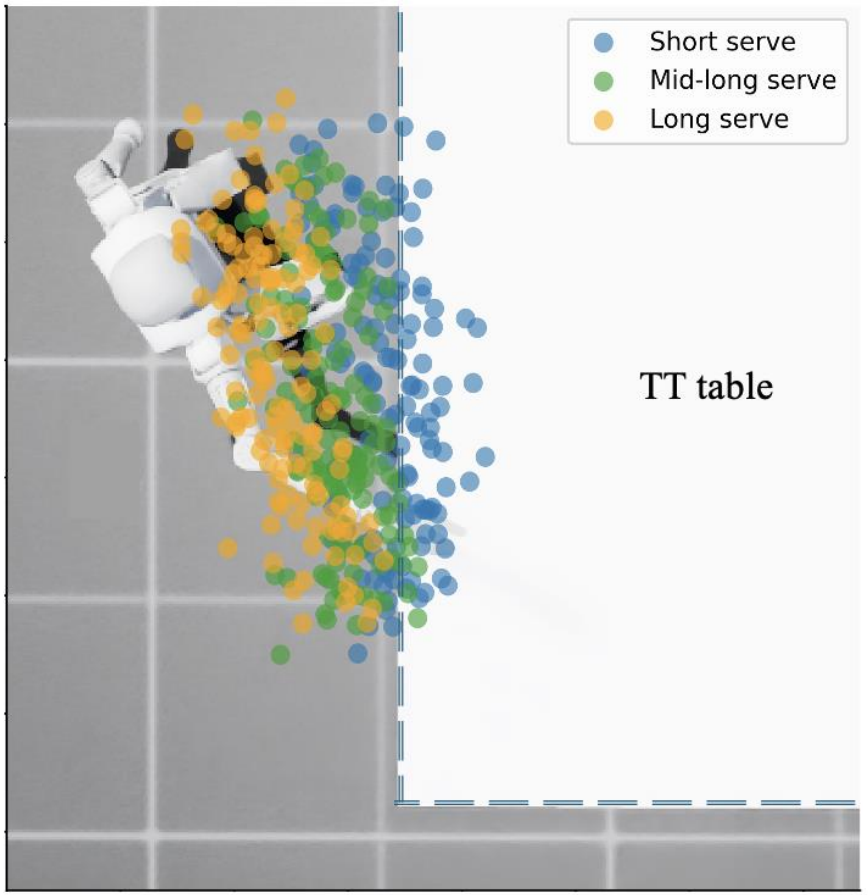

Deployment Results

- High **hit (93.5%)** and **success return (61.3%)** rate on the real hardware
- Competitive ball return **speed (6.9 m/s)**

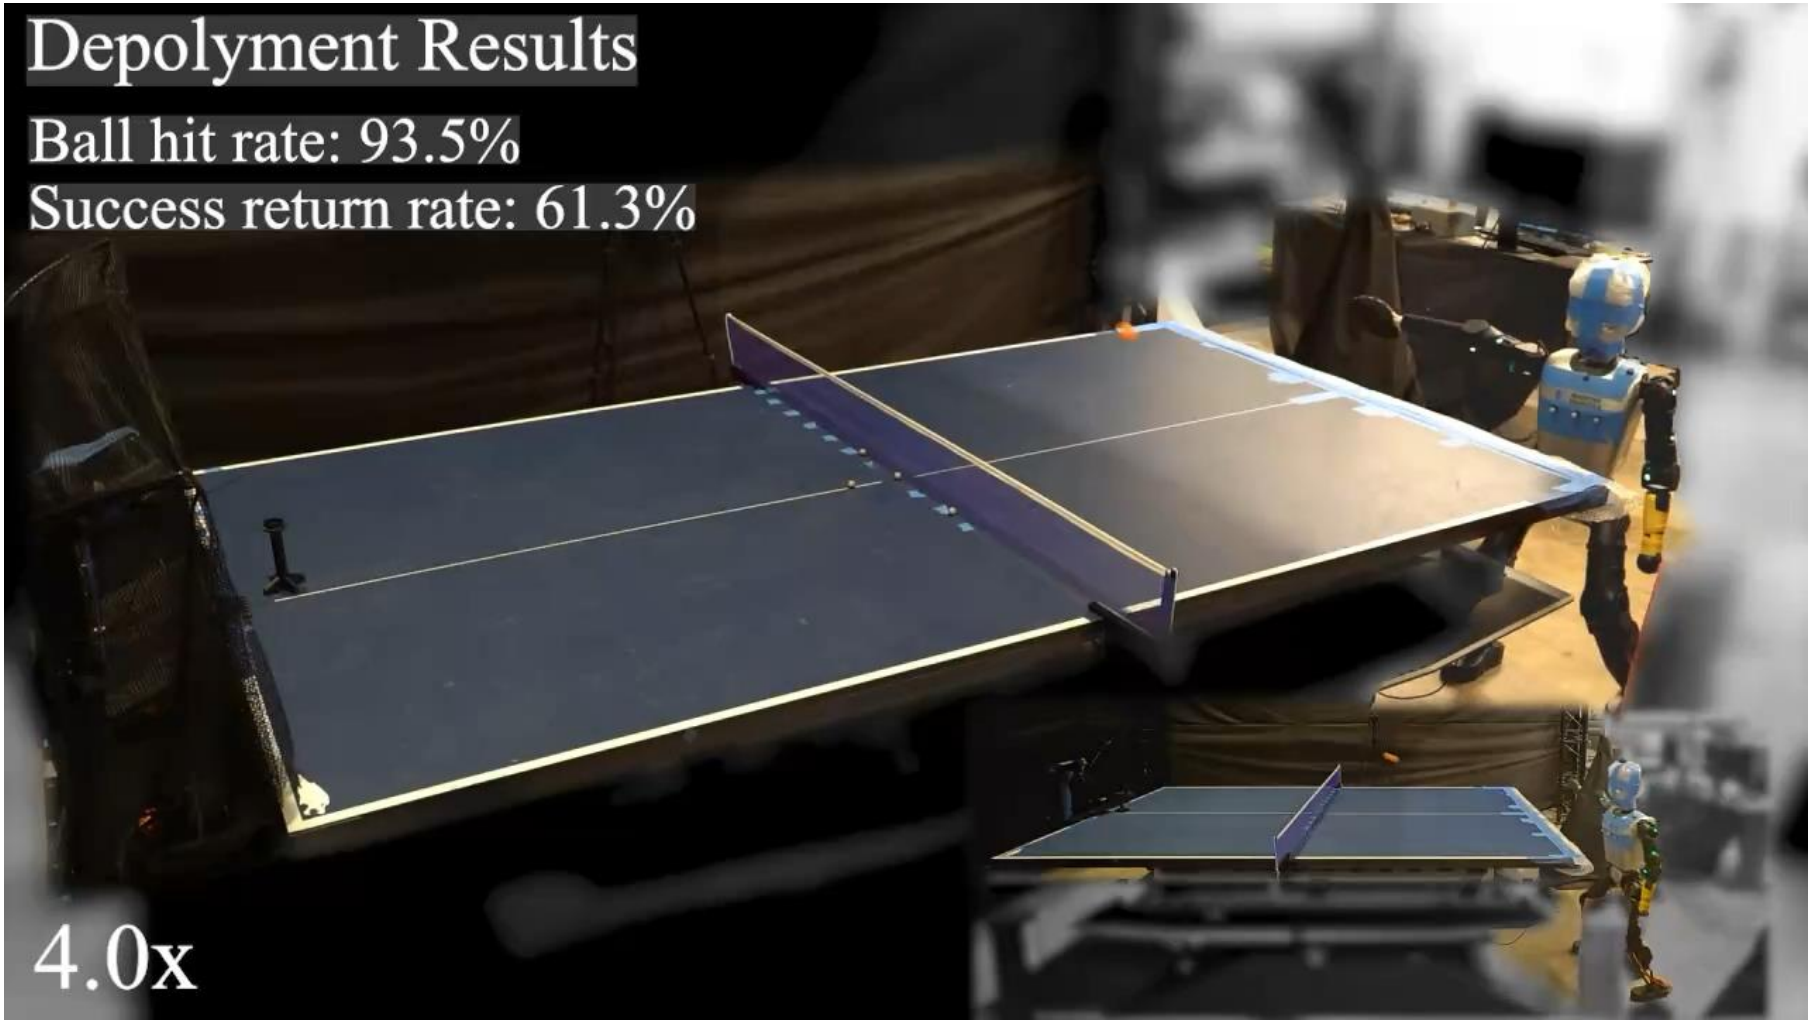

Deployment Results

- Dynamic **two-dimensional footwork** emerges from incoming ball with various positions and velocities.

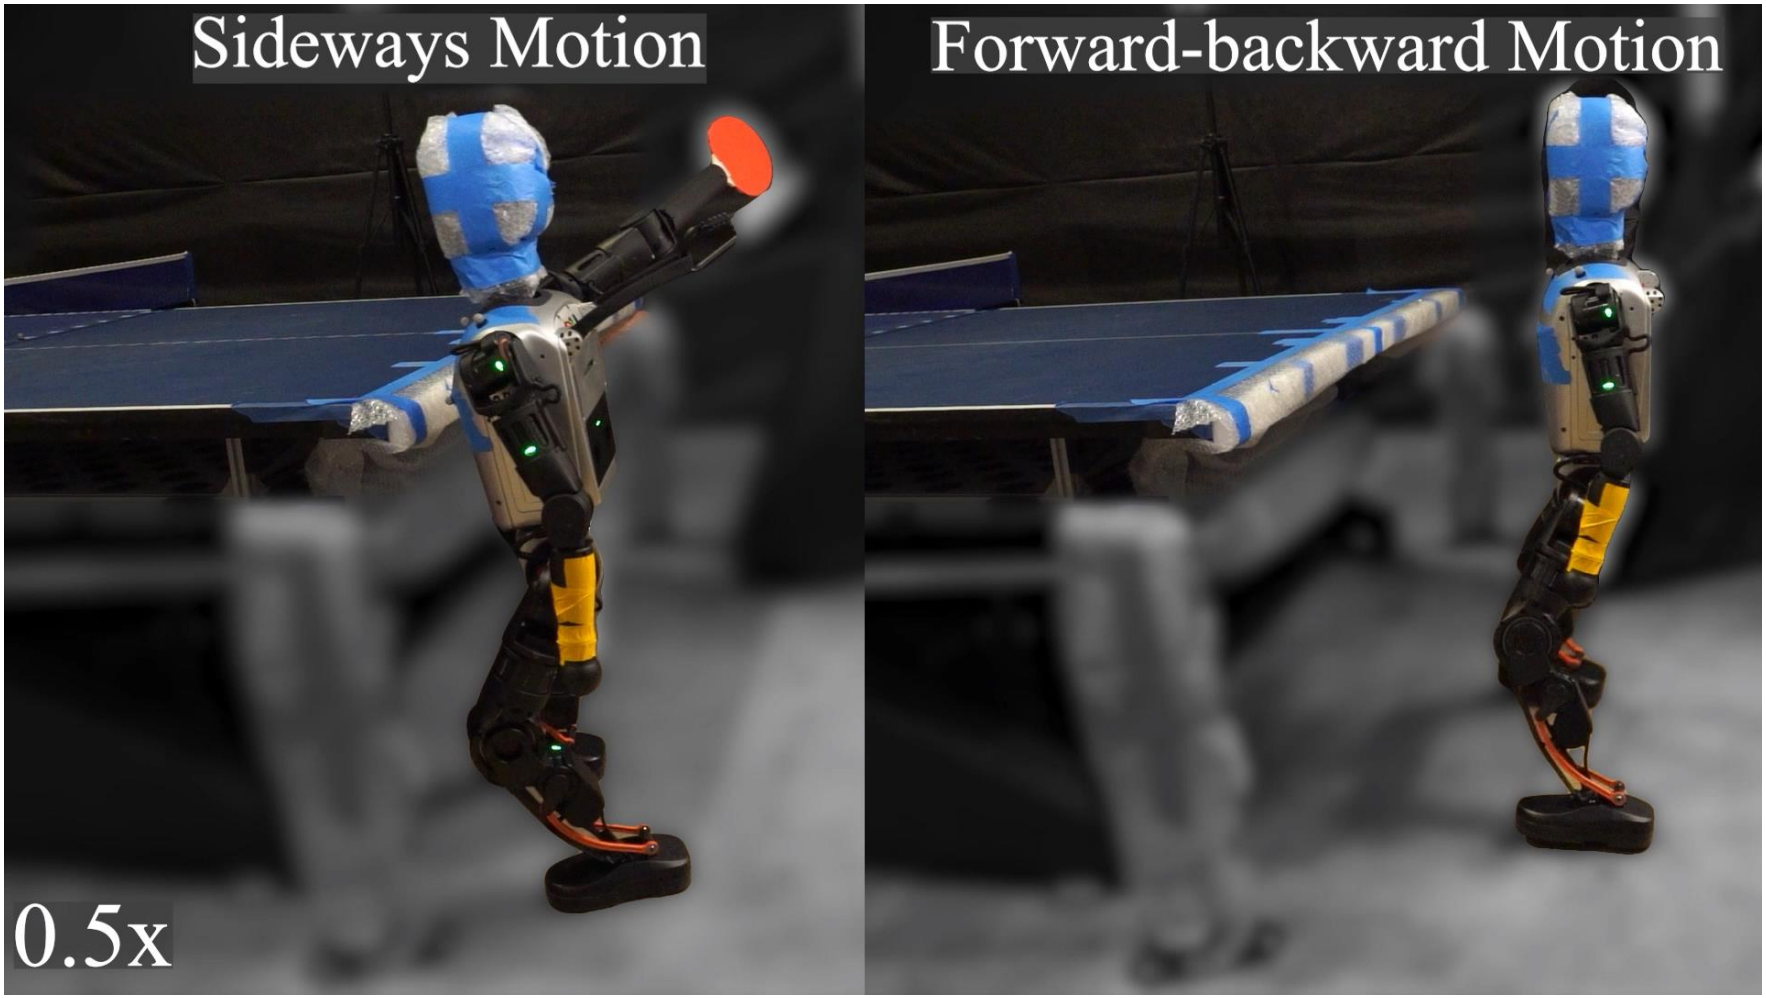

Key Takeaways

- **Reward design.** Coordinated, whole-body end-to-end RL is feasible, but it relies on ***carefully designed dense, continuous, physics-informed rewards*** to guide exploration and stabilize learning.
- **Prediction.** Adding a lightweight ball-trajectory predictor (world-model–like) enables anticipatory robot motion, such as ***proactive paddle reaching and footwork toward the incoming ball***.
- **Environment design & sim-to-real**. Training without an analytical planner demands:
  - ***Accurate physics simulations*** (restitution, drag, friction, contact)
  - ***Careful domain randomization***

  to produce reliable ball-hitting and shrink the transfer gap.



Yan Gu
Assoc. Prof, ME
yangu@purdue.edu

Adaptive and Robust Embodied Intelligence for Human-Centered Robotics

PURDUE UNIVERSITY

Robust Mobility in Dynamic Environments



Ours *Commercial*

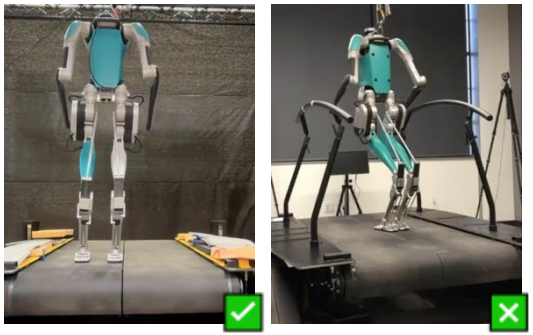

Safe, adaptive control on accelerating ground (for emergency response, inspection, and maintenance).

Trust in Human-Robot Interaction

with Shaoshuai Mou (AAE) and Raymond Yeh (CS)



New tasks or environments



Building trustworthy autonomy through uncertainty-aware decision making.

Physics-aware Learning for Contact-rich Loco-manipulation

with Karthik Ramani (ME)



Fast and precise perception-to-action learning at sub-second timescales.

Tactile-aware Loco-manipulation of Dynamic Objects

with Yu She (IE)

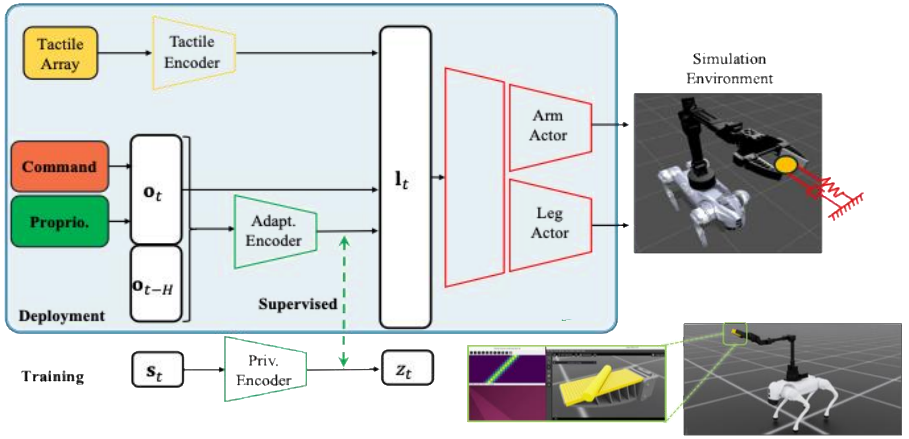



TRACE @ Purdue | Real-World Embodied AI

Physics-grounded robot learning, multimodal perception and decision-making, uncertainty-aware control and safe autonomy



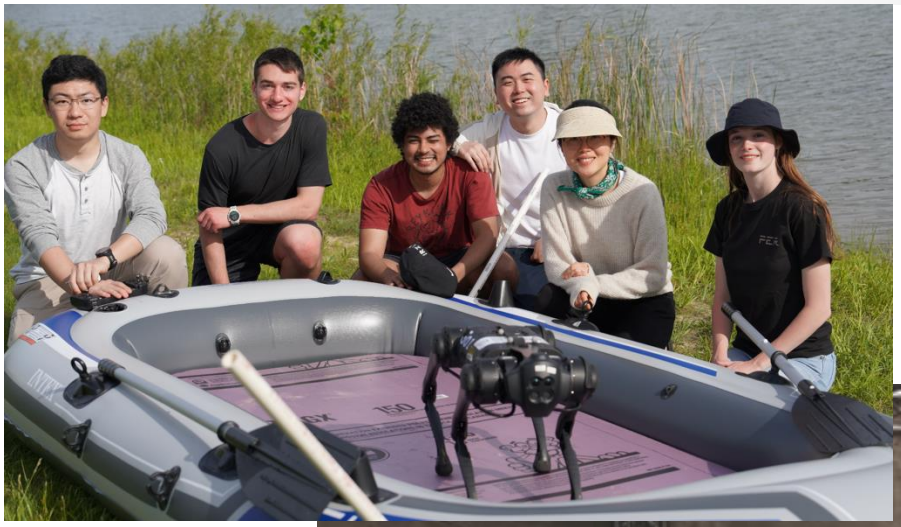

We are hiring and looking for collaboration opportunities!

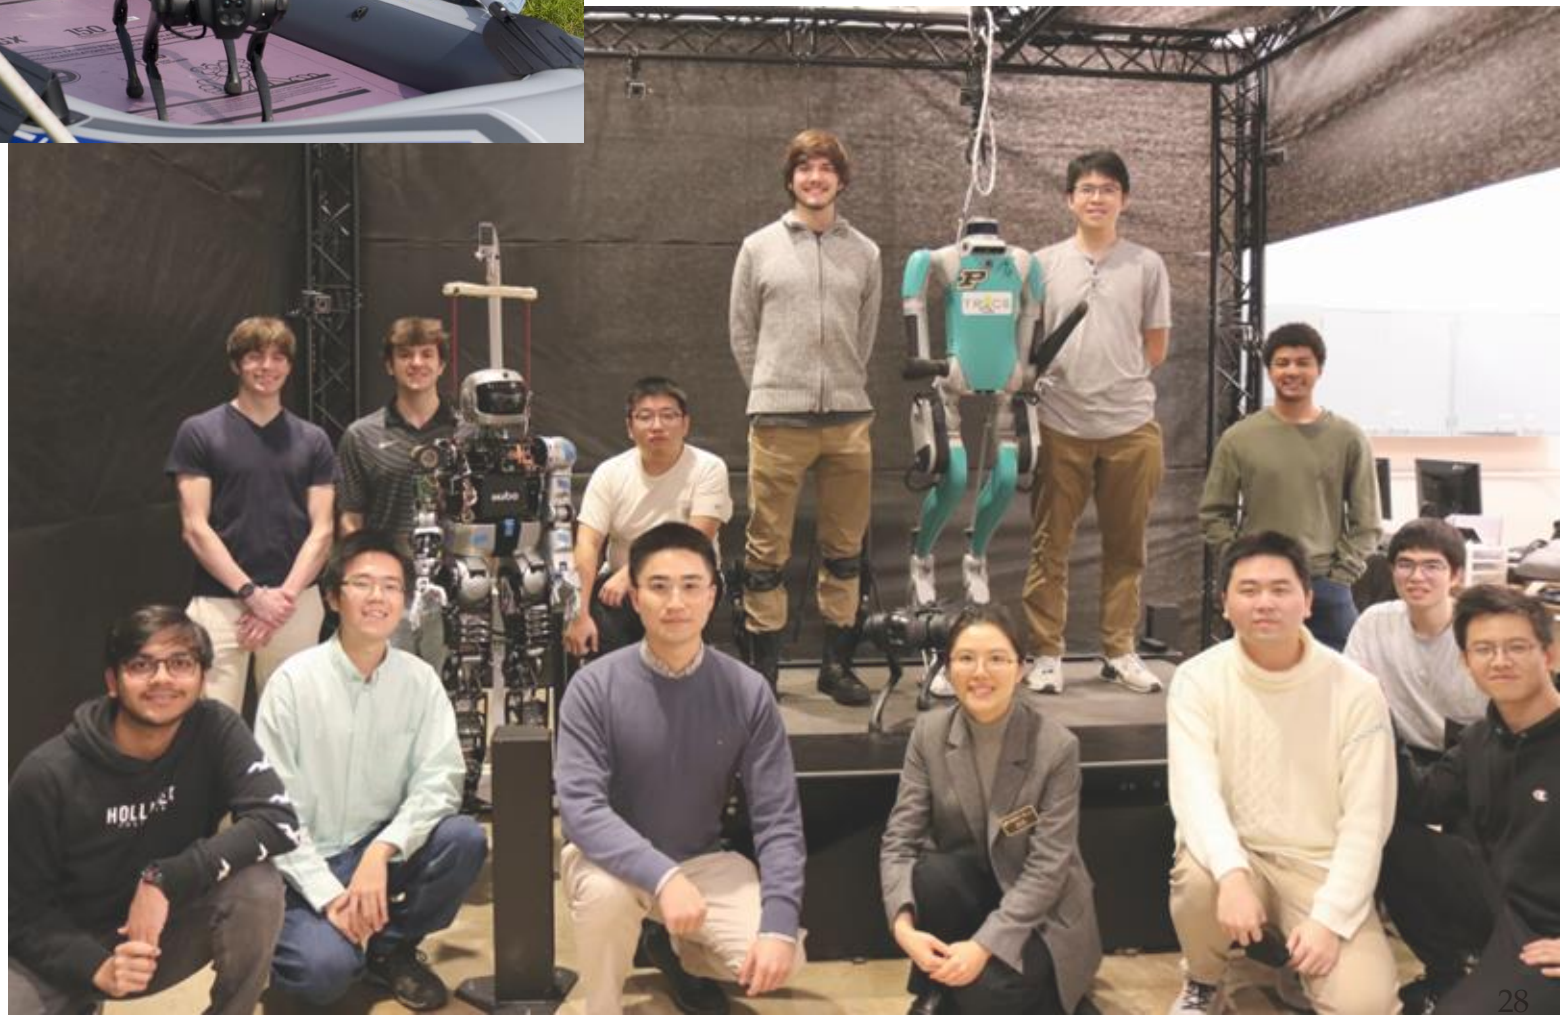

TRACE

PURDUE UNIVERSITY