The Duke Humanoid: Design and Control for Energy-Efficient Bipedal Locomotion Using Passive Dynamics

Boxi Xia, Bokuan Li, Jacob Lee, Michael Scutari, Boyuan Chen
Duke University

Duke | PRATT SCHOOL of ENGINEERING

General Robotics Lab

website

Overview

The Duke Humanoid, a 10-DOF **open-source** platform for locomotion research, mimics human physiology with frontal-plane symmetry to maintain static balance with **straight knees**. We deployed a zero-shot end-to-end reinforcement learning (RL) policy for velocity-tracking walking and a **passive RL policy** that **selectively disables joint motors** to exploit passive dynamics. Experiments show the passive policy **reduces cost of transport** (COT) by up to 50% in simulation and **31%** in real-world tests.

Hardware

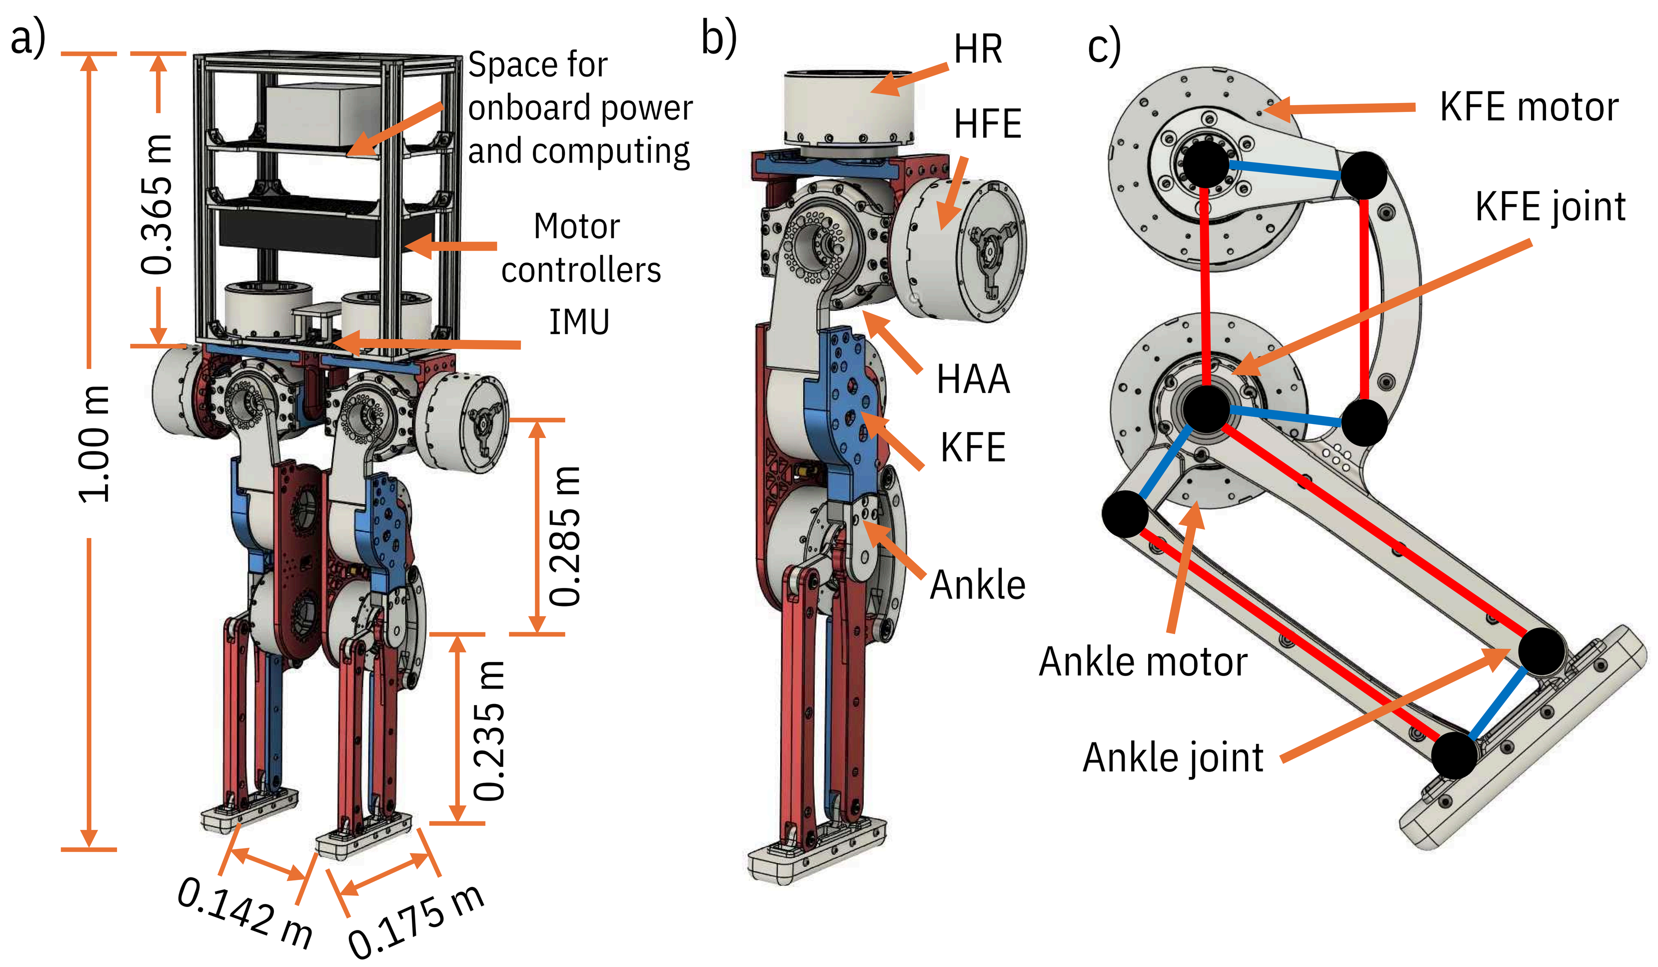

Mechanical Design Overview: a) Major dimensions and the extensible body design. b) All joints in the left leg. c) Two parallel linkages in the knee and ankle.

10 DoFs, 5 per leg: 3 at the hip (HR, HFE, HAA), 1 for knee flexion/extension (KFE), and 1 for ankle plantar/dorsiflexion. Hip center of mass aligned in the frontal plane, enabling balance and straight-knee walking

Joint	HR	HAA	HFE	KFE	Ankle
Human [deg]	[-50,40]	[-40,20]	[-30,110]	[0,150]	[-20,50]
Ours [deg]	[-90,60]	[-40,40]	[-90,90]	[0,110]	[-45,45]
Coverage	100%	100%	85%	73%	95%

Our range of motion largely matches that of human leg joints.

Robot	Symmetry*	Leg length [m]	Leg DoF	Mass [kg]	Max HFE [N·m]	Max KFE [N·m]
MIT[23]		0.55	5	24	68	136
Berkeley[5]		0.4	6	16	63	81
UnitreeG1[3]		0.6	6	35	88	120
HECTOR[6]		0.44	5	16	33.5	51.9
iCub[9]		0.4	6	24	40	40
COMAN[22]		0.44	6	55	55	40
Ours	✓	0.5	5	30	264	238

Symmetry*: hip arrangement symmetry across the frontal plane

Comparison of key parameters for mid-sized academics and industrial humanoids.

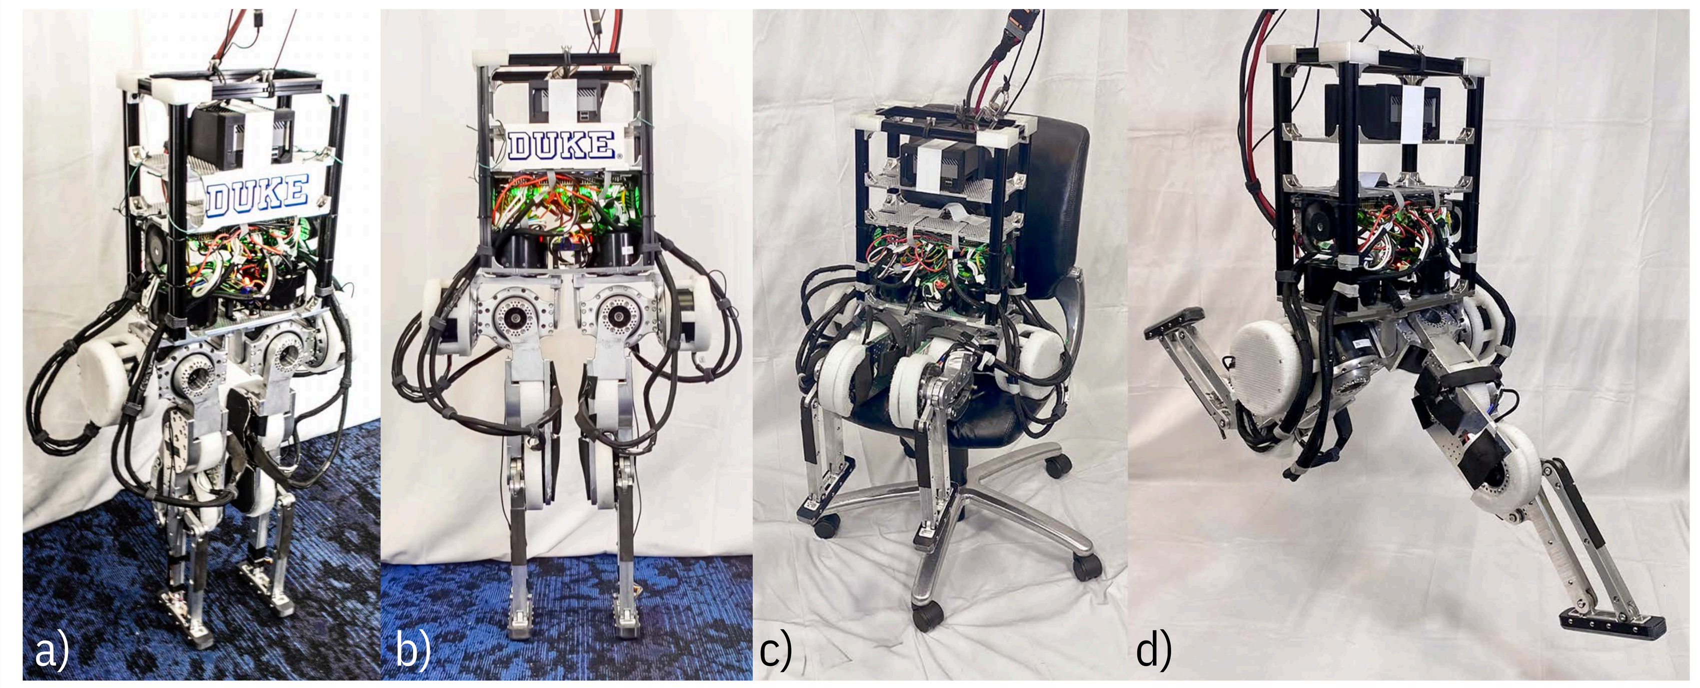

Duke Humanoid V1: a) and b) The frontal plane symmetry of the hip enables static standing with straight knees. b) and c) Additional poses demonstrating the robot's range of motion.

Leveraging passive dynamics in RL locomotion control

To improve locomotion efficiency, we learn a per-joint scaling factor [0,1] that modulates torque, letting joints go passive instead of always being actively controlled.

$$\tau = \underbrace{\alpha}_{\text{scaling factor}} \underbrace{\left(k_p(\mathbf{q}^* - \mathbf{q}) + k_d(\dot{\mathbf{q}}^* - \dot{\mathbf{q}})\right)}_{\text{baseline position PD control}}$$

τ : joint torque
q^* : learned desired joint position
$\dot{q}^*$: desired joint velocity
k_p : stiffness parameters
α : learned scaling factor
q : joint position
$\dot{q}$: joint velocity
k_d : damping parameters

At 0.3 m/s, three real-world trials showed the passive policy reduced cost of transport (COT) by 31% compared to the baseline.

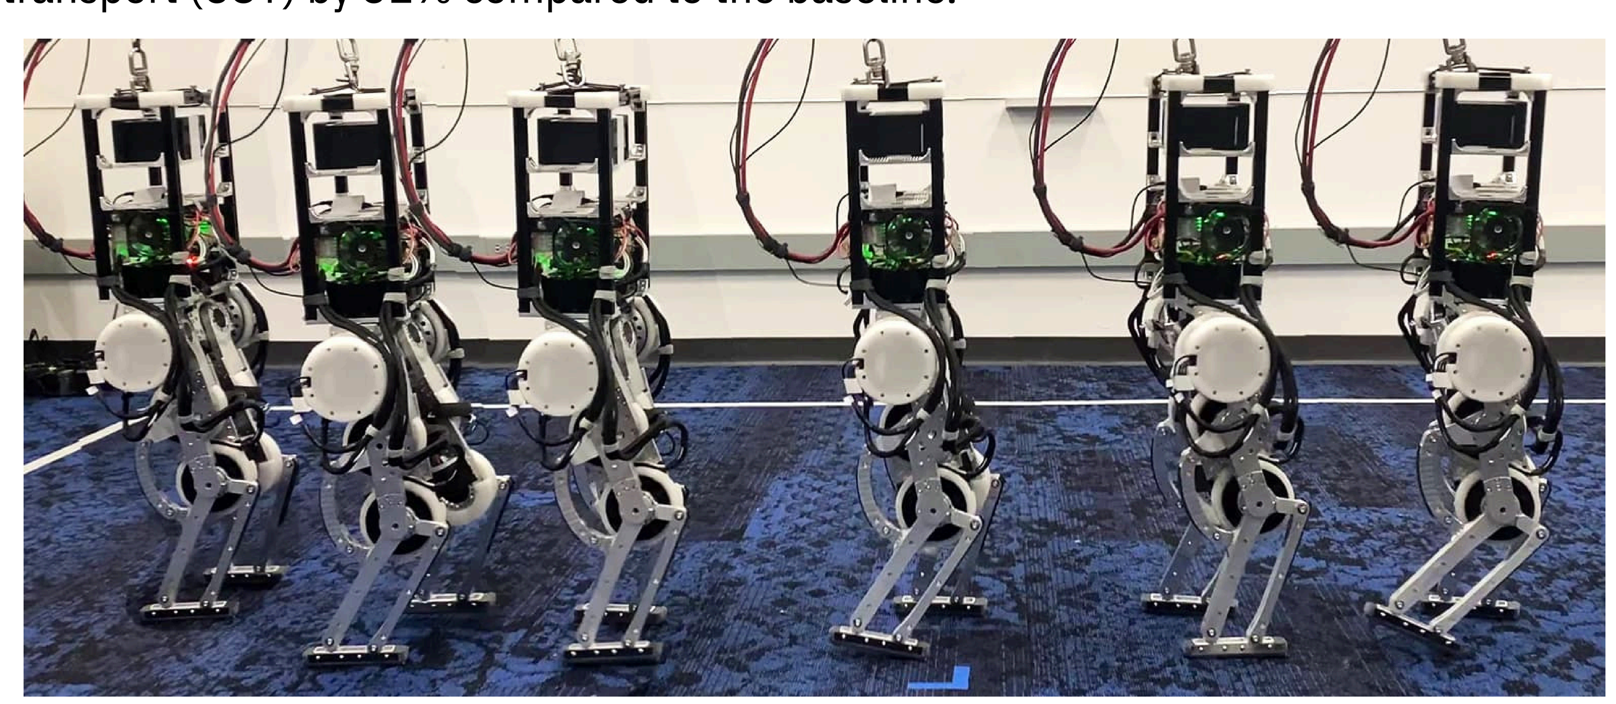

walking using the **baseline** RL policy. CoT 1.13 ± 0.1

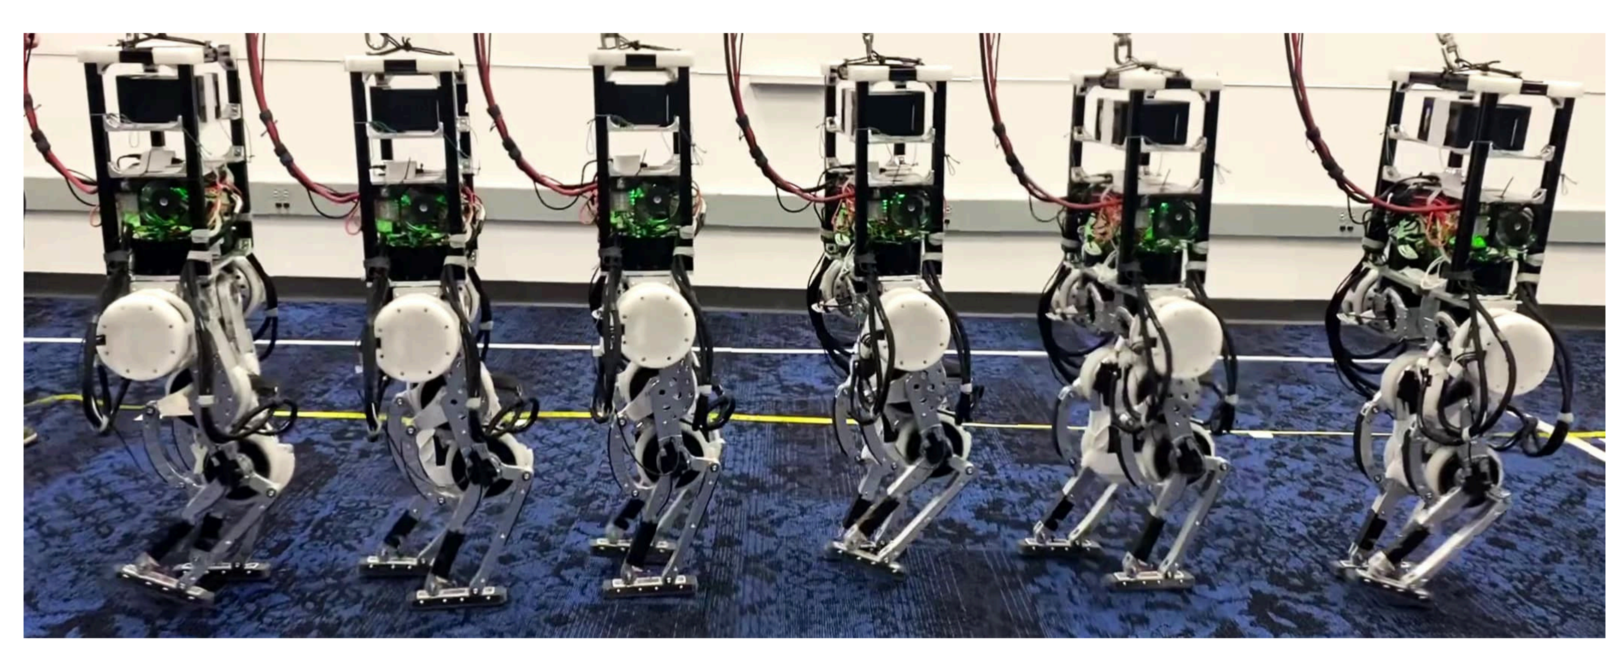

walking using the **passive** RL policy. CoT 0.77 ± 0.1

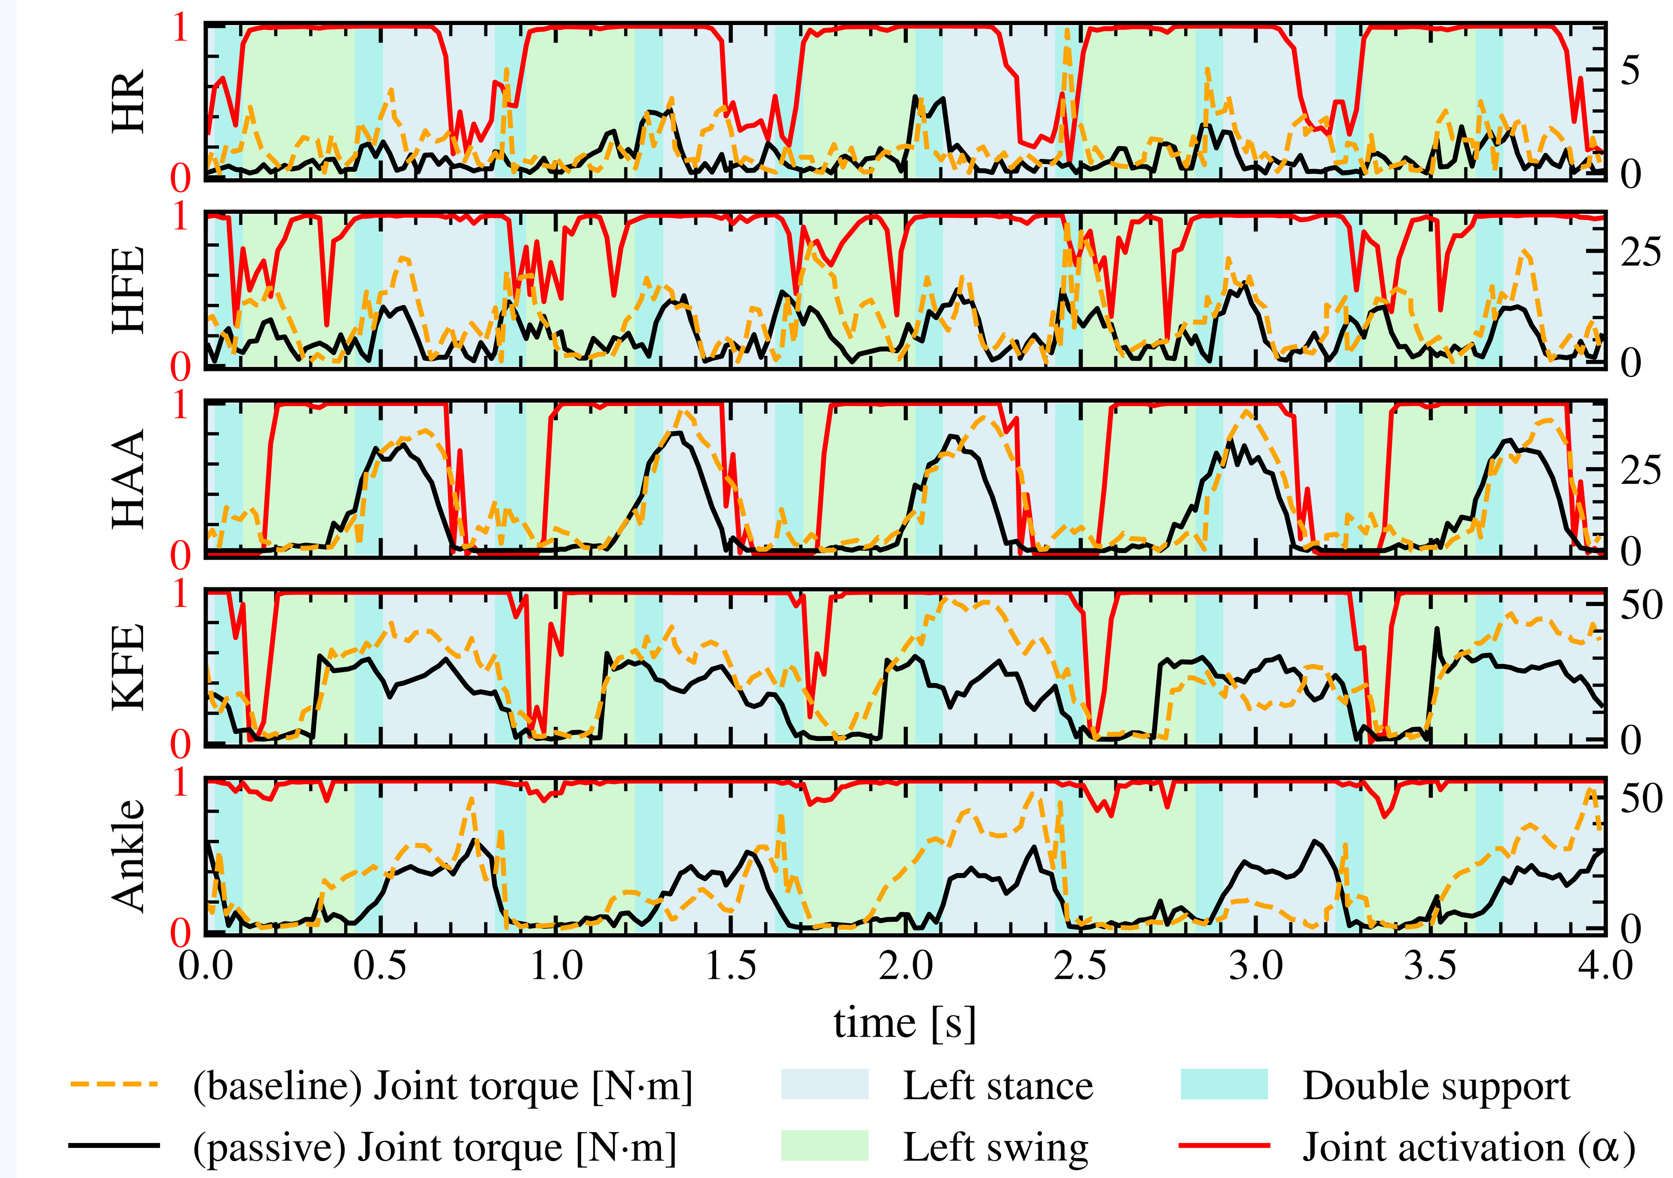

Comparison of left leg absolute joint torque for real-world baseline vs. passive walking: KFE deactivation during swing in passive walking demonstrates the utilization of passive dynamics, contrasting with the generally higher torque of the baseline policy

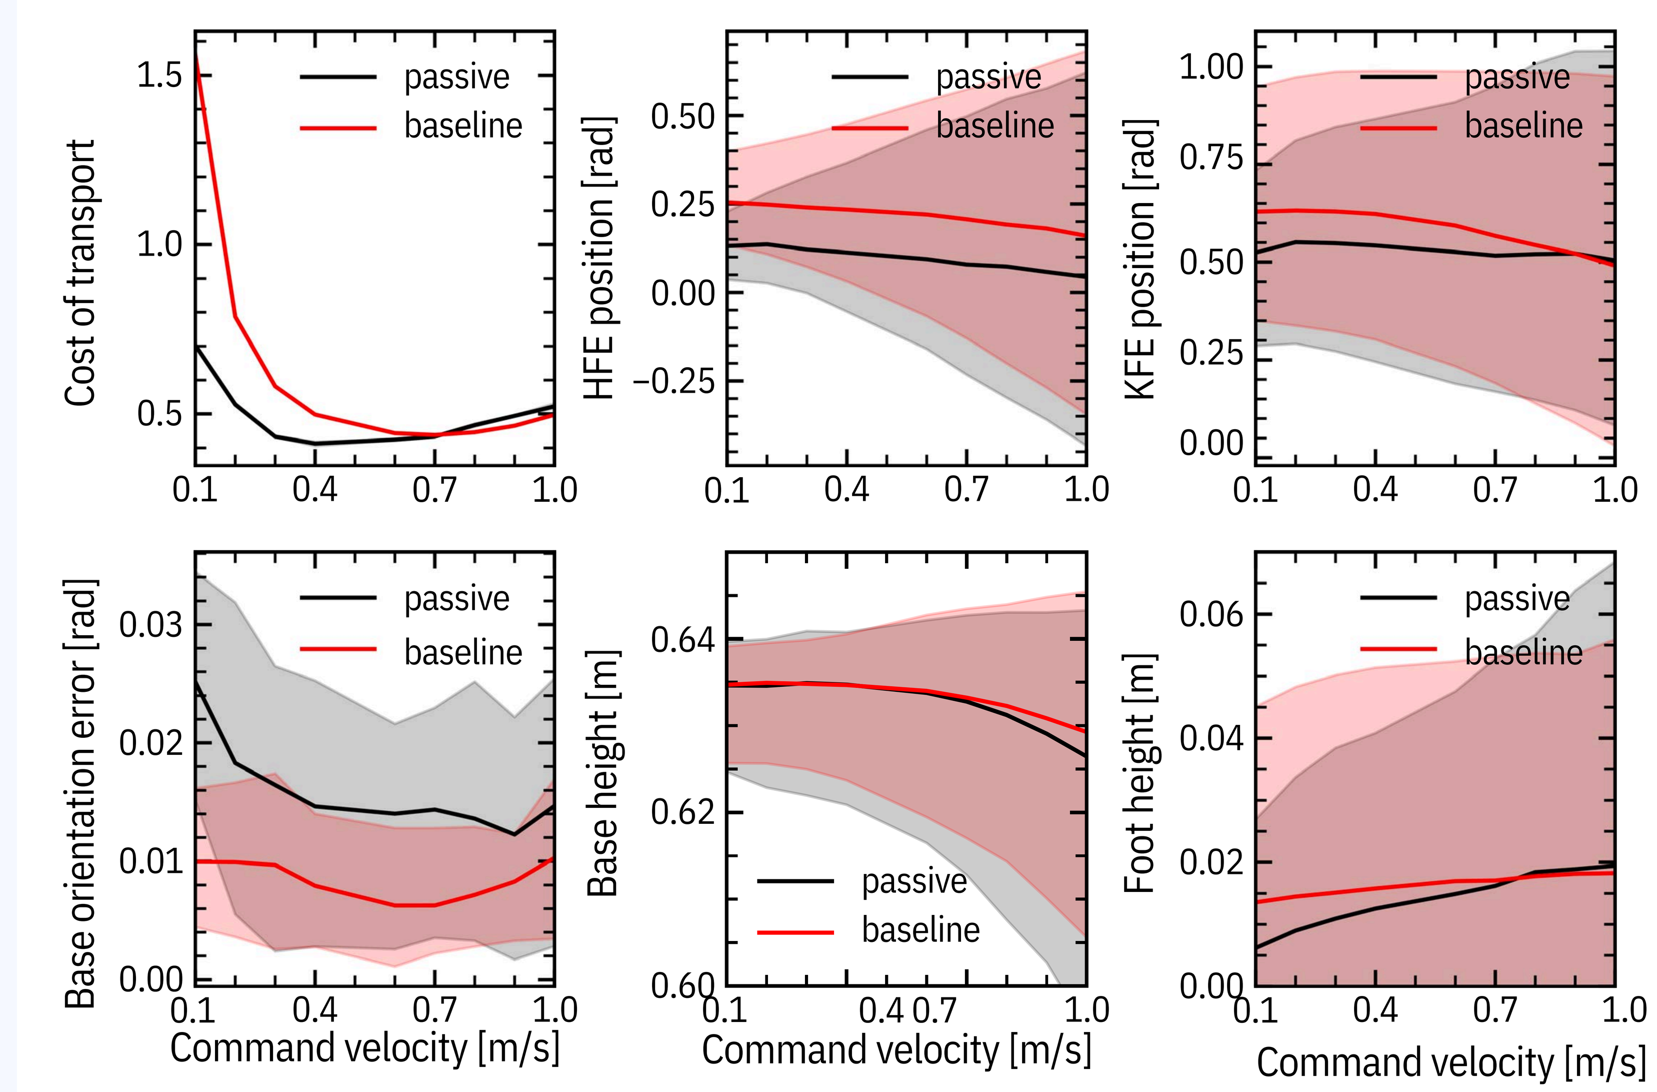

Baseline vs. passive walking policy in Simulation: Passive is up to 50% more energy-efficient at low speed, with reduced KFE/HFE positions (less action). Shaded regions show 95% confidence intervals.

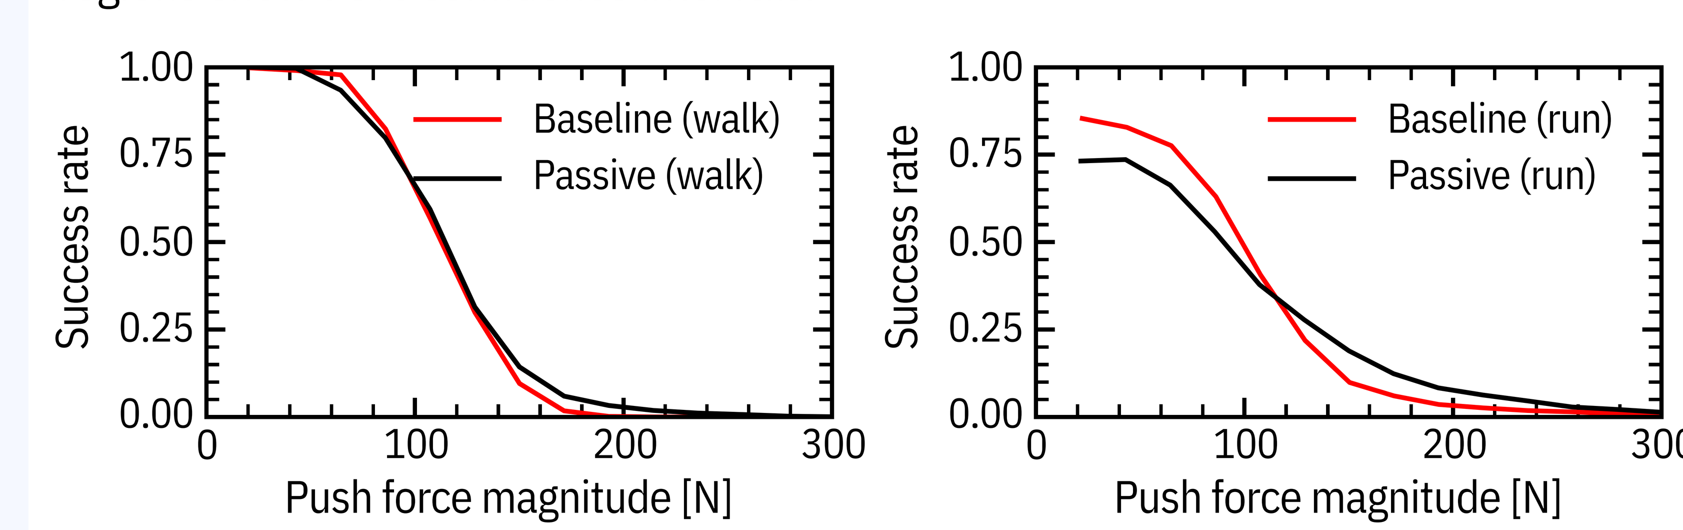

Push recovery in simulation: baseline vs. passive policy for walking (left) and running (right). Adding passive actions leaves robustness largely unchanged.