Ziang Zheng

(+86) 189-5046-7116 | ziang_zheng@foxmail.com | github.com/ZiangZheng | ziangzheng.github.io

Education

Tsinghua University | ME, School of Vehicle and Mobility | Master 2024.09 – Present

- **Full-Stack Humanoid Researcher: scalable RL for embodied agents**, **predictive world models**, **foundation behavior priors** (Transformer / diffusion / generative); **closed-loop sim-to-real on Unitree G1**; end-to-end pipelines across IsaacLab / IsaacSim / mjlab / Genesis.
- Founder, **Renforce Dynamics** open-source org (20+ repos, **1,100+ stars**); gait group lead, THMOS humanoid soccer team. **Keywords:** Adaptive Whole-Body Policies · Predictive World Models · Self-Improving RL · Cross-Embodiment Sim-to-Real.

Central South University | CS & IoT, School of Computer Science | Bachelor 2021.09 – 2024.06

GPA: 92.63/100 (Ranked 1st) **Skills:** C/C++, Python, ROS, embedded / IoT systems, computer architecture, Linux system-level development.

Research Experience

THU iDlab | *Humanoid Whole-Body Control & Foundation Behaviors* 2025.03 – Present

- [**Master's Thesis Platform**] *Whole-Body Learning for Dynamic Humanoid Table Tennis on Unitree G1* — closed-loop Nokov mocap perception, model-based ball-trajectory prediction, and a two-stage **Hit-and-Move** whole-body skill policy trained in mjlab (**MuJoCo-Warp, 116k env-steps/s**) with FSM-coordinated hardware deployment.
- [**Adversarial Imitation — Theory**] *TCMP: Task-Centric Policy Optimization from Misaligned Motion Priors.* Reformulates AMP-style imitation as a **first-order task-priority constrained optimization**, treating style as a **conditional regularizer** that selects within the task-optimal policy set Π^*_{task}. **ICRA** (**Co-First Author**, Under Review).
- [**Adversarial Imitation — System**] *PBMP: Pretrain-Bias Motion Prior.* Replaces AMP's jointly trained discriminator with a **tracker-distilled, training-free style reward** under a **constrained-MDP** formulation; **∼10× wall-clock speedup** over AMP at matched task reward with **zero-shot sim-to-real** on Unitree G1. **IEEE RA-L** (**First Author**, Under Review).
- [**Diffusion Motion Prior**] *SMP: Score-Distilled Diffusion Motion Priors for Stylized Whole-Body Humanoid Control.* Replaces the AMP discriminator with a **TinyMDM diffusion motion prior** and an **ESM-SDS** (energy-score-matching SDS) style reward, with **command-gated style modulation** for stylistically diverse, command-conditioned whole-body behaviors. **CoRL 2026** (**First Author**, Under Review).
- [**Motion Tracking**] *M3imic: Versatile Whole-Body Controller for Multimodal Motion Mimicking.* A single whole-body tracker unifying retargeted mocap, raw mocap, and monocular video through a **shared latent command space**. **IEEE RA-L** (**Co-First Author**, Under Review, 161+ stars).
- [**Open-Source Infra.**] **beyondAMP** (AMP framework) · **TrackerLab** (multi-mode tracking) · **GenesisLab** (personal-PC RL) · **MuGS** (MuJoCo Gaussian Splatting) · **SoccerLab**/**IsaacNPC** (multi-agent humanoid soccer, Demo).

THU iDlab | *Legged Robot Locomotion* 2025.01 – 2025.09

- [**Cross-Embodiment**] *Learning a Shared Latent-to-Latent Locomotion Policy for Diverse Legged Robots.* A **morphology-agnostic** locomotion policy that aligns heterogeneous embodiments through a **shared latent-to-latent representation**. **IROS 2025** (**First Author**, Accepted).
- [**Self-Improving Locomotion**] *Jump-Start RL with Self-Evolving Priors for Extreme Monopedal Locomotion.* **Self-evolving motion priors** are iteratively re-distilled from on-policy rollouts and reused as jump-start initializations, escaping sparse-reward local optima on extreme monopedal hopping. **IEEE RA-L** (**First Author**, Under Review).

THU iDlab | *Offline & Model-Based RL (RL foundations)* 2023.09 – Present

- Co-first author on *Safe & Efficient MPC with Lagrange Multiplier Neural Network* (**IEEE TCST**, Accepted) · *Canonical Form of Datatic Description in Control Systems* (**ACC 2025**, Accepted) · *In-Distribution Feasible Policy Optimization* (Under Review).

Industrial Projects

XVirobotics | *Humanoid Foundation Behavior Models / Part-time Researcher* 2025.06 – Present

- Research collaborator under **Flood Sung** (ex post-training RL lead, Kimi / Moonshot AI) on humanoid foundation behavior models: a **predictive video world model** produces egocentric rollouts and a **VLM-based inverse dynamics module** infers action chains, addressing the **egocentric POMDP** of first-person observation (*arXiv 2026*). Lead a small undergraduate research team.

Efort Robotics | *Full-Scale Humanoid Sim-to-Real / iDlab Collaboration* 2025.06 – 2025.12

- Co-led iDlab whole-body sim-to-real on Efort's full-scale humanoid (retargeting, sim training, deployment); **first real-hardware hurdle-jumping** on the platform.