

ZIHENG QIU

✉ theqiuZh@gmail.com · Ph.D. Student in Embodied AI ·  · 

EDUCATION

Fudan University, Ph.D. in Computer Science and Technology Sep. 2025 - Present

- Research: embodied AI and robot manipulation; Advisor: Prof. Wenchao Ding

University of Electronic Science and Technology of China, B.Eng. in Software Engineering Sep. 2021 - Jun. 2025

- Advisor: Prof. Yong Liao

PUBLICATIONS

AutoSpeed: Annotation-Free Stage-Adaptive Motion Speed Learning for Robot Manipulation *ECCV 2026*

Qingda Hu[†], **Ziheng Qiu**[†], *et al.* († Equal contribution)

 [arXiv](#)  [Project](#)

- Introduced cost-aware selection over speed-retimed trajectory targets, enabling end-to-end stage-adaptive speed learning without speed or stage annotations.
- Developed DCT-based trajectory retiming for smooth non-integer speed changes and Nonlinear Temporal Aggregation for deployment.
- Evaluated on 62 simulation tasks and four Agilex Piper bimanual tasks, achieving a $1.78\times$ average real-world speedup while maintaining or improving success rates.

Resolving State Ambiguity in Robot Manipulation via Adaptive Working Memory Recoding *RA-L 2026*

Qingda Hu, **Ziheng Qiu**, *et al.*

 [arXiv](#)  [Project](#)

- Proposed PAM, a visuomotor policy that separates motion primitives from temporal context and recodes multi-scale history into compact working memory.
- Supported a 300-frame history window with stable training and inference above 20 Hz, while exposing the history segments and modalities used by the policy.
- Trained one policy across seven real-world state-ambiguity tasks, reaching a 0.91 average success rate versus 0.61 for LongDP and 0.45 for MTIL.

HONORS

Graduate Academic Scholarship, Fudan University 2026

SKILLS

- **Programming and deep learning:** Python, PyTorch; model training, debugging, evaluation, and experimental analysis
- **Generative and multimodal models:** diffusion models, large language models, and vision-language models for visuomotor policies and robot manipulation
- **Robot learning:** imitation learning, reinforcement learning, action representations, world models, and long-horizon modeling
- **Simulation and deployment:** ROS, LIBERO, Meta-World, and RoboTwin; deployment and debugging on Franka and AgileX single- and dual-arm platforms
- **Research engineering:** multi-GPU distributed training, robot data collection and preprocessing, Linux, Bash, Git, and Docker