

Zuo Gou

Tsinghua University — gouz23@mails.tsinghua.edu.cn — tactino.github.io — github.com/tactino

EDUCATION

Tsinghua University

Sep 2023 – Expected Jun 2027

B.Eng., Automation (Xinya College & Dept. of Automation)

Beijing, China

· **GPA: 3.8/4.0** **Awards:** Tsinghua Comprehensive Scholarship (2024, 2025); Luo Yuehua Scholarship (2024)

MANUSCRIPTS

[1] **Z. Gou** et al. *Relative Noise Schedules in Joint Continuous–Discrete Diffusion*. In preparation; target ICML 2027.

[2] **Z. Gou** et al., W.-C. Ma. *Humanoid2Humanoid: Shared Sensing for Cooperative Loco-Manipulation*. Under review, CVPR 2027.

RESEARCH EXPERIENCE

Prof. Wei-Chiu Ma's Group (Dept. of Computer Science, Cornell University)

Jun 2026 – Nov 2026

Research Intern, Multi-Agent Humanoid Locomotion & Inter-Robot Communication

Ithaca, NY

- Studying how two humanoids can share what they see in real time, so that a robot walking backward climbs stairs hidden from its own cameras by relying on its partner's view.
- Established when such sharing matters: robots usually infer hidden terrain from gait and contact, and only in specific conditions does the partner's view become indispensable, recovering 20–40 success points.
- Built decentralized policies that act on their own sensing plus messages from the partner, halt safely when the link drops and resume when it returns; demonstrated a simulated pair climbing stairs together.

Prof. Ge Liu's Group (Dept. of Computer Science, University of Illinois Urbana-Champaign)

Mar 2026 – Present

Research Intern, Diffusion Models & Generative Modeling Theory

Remote

- Studying how a diffusion model that generates continuous and discrete quantities at once — atom positions with atom types and bonds — should pace noise across its two channels, a choice now hand-tuned per dataset.
- Characterized when the likelihood bound is indifferent to that pacing and what the best schedule looks like when it is not, and confirmed the predictions on public molecular models.
- Showed that reported likelihoods are dominated by a single coarse step of the sampler, and built an estimator reaching the same precision with roughly $2,600\times$ less computation.

Prof. Huazhe Xu's Group (IIIS, Tsinghua University)

Jul 2025 – Jan 2026

Research Assistant, Vision–Language–Action Models & RL

Beijing, China

- Architected *PlugRL* (project page), a distributed framework for reinforcement learning on vision–language–action policies.

Prof. Yebin Liu's Group (Dept. of Automation, Tsinghua University)

Sep 2024 – Jun 2025

Research Assistant, Dexterous Manipulation & Diffusion-based Control

Beijing, China

- Developed a diffusion model that generates physically plausible hand–object interaction for dexterous manipulation, with the simulation environments and demonstration pipelines it was trained on.

TECHNICAL SKILLS

Programming

Python (PyTorch), C++, CUDA, MATLAB, Shell

Robotics & simulation

IsaacGym, Isaac Lab, MuJoCo, Legged Gym, RSL-RL

Infrastructure

Linux, Slurm clusters, distributed multi-GPU training, Weights & Biases