

Siqi Chen

chentingjia1209@163.com | github.com/tj-chen-1209 | tj-chen-1209.github.io

Availability: Visiting research internship in 2026 (full year, flexible).

EDUCATION

Harbin Institute of Technology (HIT)

M.Eng. in Control Theory and Control Engineering (School of Astronautics)

Harbin, China

Sep 2025 – Present

Harbin Institute of Technology (HIT)

B.Eng. (School of Astronautics)

Harbin, China

Sep 2021 – Jun 2025

- GPA: 91.21/100 (overall rank: top 15%).
- Coursework: Calculus (98), Linear Algebra (96), Probability & Statistics (94), Automatic Control Theory (95.8).
- English: CET-6 507.

PUBLICATIONS

ResMerge

ICML 2026 submission (under review)

Fourth author

- Conducted small-model ablation experiments on LIBERO-Object under a single RTX 4090 setup; implemented evaluation scripts for result aggregation and plotting.
- Created core paper figures (teaser and method diagram) to improve clarity of the submission.

EXPERIENCE

PKU–PsiBot Joint Lab

Beijing, China

Research Intern

Nov 2025 – Feb 2026

- Delivered project-oriented engineering support for embodied robotics research, including system integration and experiment tooling.
- Implemented and maintained training/evaluation workflows; improved reproducibility and iteration speed for benchmark experiments.

Standard Robots

Shenzhen, China

Embodied AI Intern

Jul 2025 – Sep 2025

- Built robot teleoperation and data collection workflows; produced structured logs for downstream learning or debugging.
- Participated in debugging DARWIN-01 electrical bus issues and contributed to NVIDIA Jetson-based development; gained hands-on understanding of the full robot system stack.

SELECTED PROJECTS

GR00T-N1.6 Fine-tuning on LIBERO | *DeepSpeed ZeRO-2, BF16, LIBERO, VLA*

2026

- Reuse an end-to-end fine-tuning and evaluation pipeline for NVIDIA GR00T-N1.6 (3B VLA) on LIBERO (4 suites, 40 tasks).
- Applied a parameter-efficient recipe: froze vision encoder and LLM; trained the projector and diffusion decoder (~40% trainable parameters).
- Achieved 97.8% average success (782/800 episodes) and released 4 benchmark-validated checkpoints on HuggingFace.

RDT Fine-tuning on LIBERO (Full FT vs LoRA) | *PyTorch, DeepSpeed, SigLIP, T5-XXL*

2025

- Implemented a complete supervised fine-tuning pipeline for Robotics Diffusion Transformer (RDT) on LIBERO, including HDF5 parsing and state/action remapping.
- Added T5-XXL instruction embedding caching to remove redundant encoding and improved training throughput by ~30%.
- Built automated checkpoint evaluation with profiling; found LoRA (rank 8–32) saturates at ~20% while full fine-tuning reaches 76–94% (suite-dependent).

RDT Integration into RLinf | *RL framework integration, LIBERO, diffusion policy*

2026

- Integrated a diffusion-based manipulation policy (RDT) into the RLinf framework and enabled checkpoint-based rollout evaluation on LIBERO.
- Aligned observation interfaces by adding 9-DoF joint proprioception (arm + gripper) to env wrappers and distributed I/O.
- Identified the core RL blocker: diffusion sampling does not directly provide action log-probabilities required by PPO-style methods (RL training in progress).

OptiTrack → ROS2 → MoveIt2 Teleop Data Capture | *ROS2, TF, PoseStamped, MoveIt2* 2025

- Implemented a motion-capture streaming module that converts OptiTrack Motive (NatNet) rigid-body tracking into ROS2 PoseStamped/TF.
- Mapped the tracked pose into RViz/MoveIt2 for teleoperation demonstration capture, logging, and simulation-side verification.

TECHNICAL SKILLS

Learning & Embodied AI: Reinforcement Learning, Imitation Learning, Diffusion Policies, VLA fine-tuning, Benchmarking & Ablations

Robotics: ROS2, teleoperation & data logging, simulation (LIBERO)

ML/Systems: PyTorch, DeepSpeed, CUDA, distributed training/evaluation, wandb)

Programming: Python, C++, Linux, Git

AWARDS

Silver Medal, 2025 World Humanoid Robot Games (Material Sorting Track)

Operator & Algorithm Developer

Aug 2025