

Binjie Zhang

+65 80273505 / +86 15216801797 | binjie97@u.nus.edu | [Personal Website](#) | [Google Scholar](#)

EDUCATION

National University of Singapore (NUS) Singapore
Ph.D. in Computer Science – Advisor: Asst. Prof. Mike Zheng Shou Jan 2023 – Jan 2027 (expected)

Tsinghua University (THU) Beijing, China
M.Eng. in Computer Science and Technology – Advisor: Prof. Chun Yuan Aug 2019 – Jul 2022

East China University of Science and Technology (ECUST) Shanghai, China
B.Eng. in Information Engineering – GPA 3.77/4.00, Rank 1/92 Aug 2015 – Jul 2019

Research Interests: LLM agents (harness, memory, tool-use, self-evolution); vision-language and video understanding; world models; lifelong / compatible representation learning.

SELECTED PUBLICATIONS

First-author works at top-tier venues

- [1] “ReGRPO: Reflection-Augmented Policy Optimization for Tool-Using Agents.” *ECCV* 2026 (under review).
- [2] “RCFC: Lifelong Imitation Learning via Prototype Replay and Coarse-to-Fine Compatibility.” *ICML* 2026 (under review, review scores: 4, 4, 4, 2).
- [3] “Ego-centric Predictive Model Conditioned on Hand Trajectories.” **ICML 2026 (under review, review scores: 4, 4, 4, 5.) (top 15%),**
- [4] “Training-Free Optimization of Agent Memory for Long-Horizon Tool-Use.” *NeurIPS* 2026 (in preparation, 1st author).
- [5] “TaCA: Upgrading Your Visual Foundation Model with a Task-Agnostic Compatible Adapter.” arXiv, 2023.
- [6] “Darwinian Model Upgrades: Model Evolving with Selective Compatibility.” *AAAI*, 2023.
- [7] “Towards Universal Backward-Compatible Representation Learning.” *IJCAI* (long oral), 2022.
- [8] “Hot-Refresh Model Upgrades with Regression-Alleviating Compatible Training in Image Retrieval.” *ICLR*, 2022.

RESEARCH & WORK EXPERIENCE

AI Research Intern, TikTok (ByteDance) – GNE / IPR Agent Platform Singapore
Sole owner of three production agent systems: harness, evolution, memory Jan – Apr 2026 (3 mo.)

- o Pivoted from a vision / world-model PhD background to a large-scale LLM-agent stack in **2 weeks**, and drove all three projects from design to production deployment inside the 3-month window.

#1 • FIRE-Core – Skill-Based Agent Harness (Agent Application) *Fire Core 2.0 Fire Core 2.1*

- o **My role:** Core designer & technical lead – architected the V2/V2.1 harness end-to-end and drove cross-team integration, review, and rollout.
- o **V2 redesign** – migrated a centralized MAS to a **declarative Skill system** with LLM semantic routing + 3-layer Progressive Disclosure, cutting routing token cost by **97.8%** and shortening per-scenario dev cycle from **weeks** → **hours**.
- o **V2.1 hardening** (17 sub-tasks, 5 waves) – UnifiedToolRegistry w/ per-tool metadata, IoC PermissionInterceptor, Hook registry, concurrency controller, XML-tag multi-skill merging, Code-Interpreter-style sandbox, and **Skill Studio**
- o **Production impact** – unified Aicolate workflows and OneService SQL into declarative SKILL.md Tools; scaled coverage **3** → **27 skills** live across **9 country markets**.

#2 • EvoA – Auto-Evolution Agent (Agent-Native Self-Evolving ML) *Lark doc: []*

- o **My role:** Project lead & sole architect – designed the Agent-Native loop, implemented the framework, and delivered the Gandalf MVP from zero to production.
- o **Agent-Native closed loop** (Diagnose → Hypothesize → Experiment → Evaluate → Reflect) with the LLM as the primary decision maker – SOUL.md + **7 knowledge playbooks** + **30+** @tool wrappers + **12 skills**, unifying Merlin (training), Dorado (data), Aeolus (monitoring) and ModelHub (deploy) behind one agent.
- o **Knowledge & efficiency layer** – append-only experiment.log with Mandatory Reflection (lesson / dont_try_again / failure_mode) turns failed runs into future GPU savings; Pilot-Gated Pipeline (15% data → gate → full train) lifted GPU utilization from ~50% to **>80%** and enabled 3–4 parallel hypotheses per round.

- **Production impact on Gandalf** (9-country video-review model) – Balanced Accuracy **+18.4%**, leak rate **-65%** (1.87%→1.4%), TRUE AUC **0.862** (first trustworthy per-country metric), single-iteration wall-clock **8h** → **4h**, human operator time **-94%**; **L3 autonomous** zero-intervention end-to-end run validated.

AI Research Intern, Tencent ARC Lab

Shenzhen, China

Compatible Representation Learning – large-scale image / video retrieval

2020 – 2023

- **My role:** Core research contributor – led four first-author papers (ICLR'22, IJCAI'22 long oral, AACL'23, arXiv'23) that defined the team's backward-compatible retrieval line.
- Built multi-GPU CLIP-style / contrastive training pipelines and proposed regression-alleviating methods enabling **hot-refresh model upgrades without re-encoding** billion-scale galleries; winner of the Tencent Technology Breakthrough Award (2022).

AI Research Intern, Tencent ARC Lab

Shenzhen, China

Cross-Modality Video Understanding

2019 – 2020

- **My role:** Core developer – built the PyTorch data / training / evaluation codebase for multi-modal research, and shipped baselines for video–text retrieval and temporal grounding.

SELECTED RESEARCH PROJECTS

Reflection-Augmented RL for Tool-Using Agents^[1]: Proposed ReGRPO, a group-relative policy-optimization variant that injects structured self-reflection into multi-modal CoT, improving tool-selection reward and sample efficiency for long-horizon tool-use agents.

Lifelong Robot Imitation Learning^[2]: Designed RCFC, a lightweight framework for lifelong imitation learning using compact prototype replay + coarse-to-fine compatibility regularization; reduces forgetting and improves cross-task transfer on LIBERO.

Ego-Centric World / Predictive Model^[3]: Two-stage model: predicts future hand trajectories and uses them to condition a Latent Diffusion Model for egocentric future-video generation; SOTA on Ego4D, BridgeData, and RL Bench.

Task-Agnostic Compatible Adapter^[5]: Proposed TaCA, a parameter-efficient adapter for seamless upgrades between visual foundation models (e.g. CLIP variants) without retraining downstream tasks; validated on large-scale video–language benchmarks.

HONORS AND AWARDS

- Tencent Technology Breakthrough Award – Hot-Refresh Model Upgrades 2022
- SZCCF Science and Technology Award – Efficient Model Upgrades 2022
- Excellent Master Degree Graduate in Beijing & Outstanding Master's Graduation Thesis 2022
- Annual College Personage Award (highest student honor at ECUST) 2018
- National Scholarship for Undergraduates (twice), Ministry of Education of China 2017, 2016

TECHNICAL SKILLS

Research Areas: LLM agents (harness / memory / tool-use / RL), VLMs, world & predictive models, video understanding, lifelong and compatible representation learning

Production & Infra (via TikTok intern): Merlin training platform, Hive, Dorado, Aeolus dashboards, ModelHub, distributed multi-GPU training, sandboxed code execution