

# Zhenhao Zhang

New York, NY | (585) 470-3118 | zz3530@columbia.edu | LinkedIn | GitHub  
AI Agents · Multimodal Learning · Reinforcement Learning · Computer-Use Agents

## EDUCATION

### Columbia University

New York, NY

*M.S. in Artificial Intelligence; Concentration: AI and Advanced Computing*

*Aug. 2026 – Dec. 2028 (Expected)*

### University of Rochester

Rochester, NY

*B.S. in Computer Science; B.S. in Business: Information Systems*

*Aug. 2022 – Dec. 2025*

- GPA: **3.85/4.00**; Cum Laude (Computer Science & Business Information Systems); Highest Distinction in Computer Science; Dean's List (2023–2025)

## SELECTED PUBLICATIONS

Mitigating Conversational Inertia in Multi-Turn Agents

**ICML 2026**

SeekJudge: A Practical Reward Framework for Reinforcement Learning in Computer-Use Agents

**arXiv, 2026**

Political Symbols and Urban Amenities: The Spatial Logic of Flag Placement in Istanbul

**SSRN, 2026**

## RESEARCH EXPERIENCE

### Zhejiang University, College of Computer Science and Technology

Hangzhou, China / Hybrid

*Researcher; Supervisor: Dr. Linchao Zhu*

*Feb. 2025 – Present*

#### **Computer-Use CAPTCHA Agents** | *Lead Researcher*

*Jan. 2026 – Present*

- Led development of an interactive multimodal benchmark and training environment with **50K CAPTCHA instances** across **20 task categories** and 5 interaction patterns, using mouse/keyboard actions and executable browser-verified ground truth.
- Designed an answer-conditioned reasoning distillation pipeline for **46K multi-turn trajectories**, integrating step-level reasoning supervision, symbol-to-pixel grounding, persistent PLAN/STATE memory, and independent VLM quality filtering.
- Trained a unified **Qwen3.5-9B** multimodal agent across all 20 task categories on a 37K+ per-turn SFT corpus, experimenting with LoRA, full-parameter fine-tuning, frozen vision encoders, 32K contexts, and multi-task data scheduling.
- Engineered large-scale multimodal SFT/GRPO training infrastructure with DeepSpeed ZeRO-3, Liger FLCE, distributed checkpointing, executable environment rewards, and custom memory optimizations for long-context training with a 248K-token vocabulary.

#### **SeekJudge** | *Computer-Use Agent Reward Modeling* | *arXiv 2026*

*Jan. 2026 – Aug. 2026*

- Built and manually annotated a **300+ trajectory** cross-platform benchmark for long-horizon computer-use agents, with trajectory-level verdicts and dense step-level supervision across web, desktop, and mobile environments.
- Designed fine-grained reward criteria covering milestones, correct decisions, error correction, off-path behavior, and harmful actions; conducted human evaluation and model-vs.-human disagreement analysis at the step level.
- Performed systematic error analysis of rule-based and model-based reward signals, identifying rule-human disagreement across multiple CUA environments and motivating learned reward models for long-horizon agent evaluation.
- Contributed to distillation and SFT training of **SeekJudge-9B**, analyzing reasoning, formatting, and termination failure modes and iterating on training data and supervision strategies.

#### **Agent Memory / Conversational Inertia** | *ICML 2026*

*Feb. 2025 – Aug. 2026*

- Implemented task-specific trajectory summarization for long-horizon LLM agents across Maze, WebShop, Wordle, SciWorld, and TextCraft, compressing multi-turn observation-action histories while preserving task-critical state and progress.
- Designed environment-aware memory representations and summarization prompts, selectively retaining navigation state, completed subgoals, constraints, prior actions, and unresolved information for downstream agent decision-making.
- Benchmarked summarization against alternative context-management strategies across heterogeneous **AgentGym** environments, analyzing the trade-off between history retention and conversational inertia; training-free summary prompts improved Maze success from 85.5% to **90.0%**.
- Conducted failure and behavioral analysis of long-context agents, building attention diagnostics that link system-token ratio → system attention → task success, and studying action repetition, exploration, and error propagation.

INDUSTRY EXPERIENCE

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| <b>Merit Interactive Co., Ltd.</b><br><i>AI Engineer Intern</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Hangzhou, China<br><i>Jan. 2026 – Mar. 2026</i> |
| <ul style="list-style-type: none"><li>Designed an end-to-end multimodal retrieval system over <b>770K+ images</b>, supporting text-to-image and image-to-image search through FAISS ANN retrieval, reranking, and score fusion, with <b>&lt;50 ms</b> vector recall and ~100–500 ms end-to-end query latency.</li><li>Built a multilingual webpage classification pipeline across 8 topic domains; labeled ~6K pages, curated 7.5K+ samples, and benchmarked <b>13 approaches</b> spanning embedding classifiers, contrastive learning, LoRA, clustering, and retrieval/reranking.</li><li>Improved classification from 62.5% to <b>87.75% Top-1</b> (97.6% Top-2) and deployed the final lightweight 264 KB model as a containerized REST API.</li></ul> |                                                 |
| <b>Merit Interactive Co., Ltd.</b><br><i>Algorithm Intern</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Hangzhou, China<br><i>Jul. 2024 – Aug. 2024</i> |
| <ul style="list-style-type: none"><li>Built mobile-app analytics datasets from raw user logs with SQL; analyzed usage, installation, uninstallation, user distribution, and temporal trends and automated Top-100 ranking reports.</li><li>Developed app-trend forecasting models using Logistic Regression and LightGBM with behavioral and temporal features to predict applications likely to trend in the following week.</li></ul>                                                                                                                                                                                                                                                                                                                   |                                                 |

ADDITIONAL RESEARCH & TEACHING

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                               |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| <b>University of Rochester, Goergen Institute for Data Science and AI</b><br><i>DSCC 279 Independent Research; Supervisor: Prof. Cantay Çalışkan</i>                                                                                                                                                                                                                                                                                                                             | Rochester, NY<br><i>Aug. 2025 – Jan. 2026</i> |
| <ul style="list-style-type: none"><li>Built a geospatial vision pipeline over ~<b>748K Google Street View panoramas</b>, including image cleaning, coordinate alignment, spatial matching, multi-scale geographic aggregation, and district-level feature construction.</li><li>Applied nested geographic buffers, distance-weighted aggregation, and Earth Mover’s Distance to quantify structural differences in political-symbol visibility across urban districts.</li></ul> |                                               |
| <b>University of Rochester, Department of Computer Science</b><br><i>Undergraduate Teaching Assistant, CSC 280: Computer Models and Limitations</i>                                                                                                                                                                                                                                                                                                                              | Rochester, NY<br><i>Aug. 2025 – Dec. 2025</i> |

SELECTED PROJECTS & AWARDS

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>GIDS Biomedical Data Science Hackathon   Prize Winner</b>                                                                                                                                                                                                                                                                                                                                                                                                                 | Aug. 2026 |
| <ul style="list-style-type: none"><li><i>MERFISH cell-type classification</i>: led the modeling pipeline for 60-class cell-type prediction on mouse spinal-cord data (200 genes, 5K labelled cells): seed-bagged LightGBM ensembles over expression, PCA, and spatial/expression-neighbour label features with biological constraint masks and a leakage-free CV protocol; placed <b>3rd of 30 teams</b> with <b>80.7%</b> accuracy (majority-class baseline 14%).</li></ul> |           |
| <b>GIDS Biomedical Data Science Hackathon   Prize Winner</b>                                                                                                                                                                                                                                                                                                                                                                                                                 | Aug. 2025 |
| <ul style="list-style-type: none"><li><i>Perturb-seq gene expression prediction</i>: developed an enhanced Norman-method pipeline for single/double perturbation prediction using robust preprocessing, multi-model regression, and adaptive coefficient estimation; achieved <b>RMSD 0.40</b> on the test set.</li></ul>                                                                                                                                                    |           |
| <b>GIDS Biomedical Data Science Hackathon   Prize Winner</b>                                                                                                                                                                                                                                                                                                                                                                                                                 | Aug. 2024 |
| <ul style="list-style-type: none"><li><i>Single-cell multiome modeling</i>: built an MLP-based peak-to-gene prediction pipeline, achieving <b>75% accuracy</b> and outperforming baseline linear and random-forest models.</li></ul>                                                                                                                                                                                                                                         |           |
| <b>Finnov8 Hackathon   Second Best Overall Award</b>                                                                                                                                                                                                                                                                                                                                                                                                                         | Mar. 2024 |
| <ul style="list-style-type: none"><li>Achieved an individual score of <b>0.67</b> in a 4-member team and improved prediction from 0.60 to 0.67 using feature engineering, Multi-Output MLP, and LSTM models.</li></ul>                                                                                                                                                                                                                                                       |           |

LEADERSHIP

|                                                                             |                       |
|-----------------------------------------------------------------------------|-----------------------|
| <b>Minister of Publicity</b> , University of Rochester CSSA                 | Jun. 2024 – Jul. 2025 |
| <b>Member, External Relations Department</b> , University of Rochester CSSA | Jun. 2023 – May 2024  |

TECHNICAL SKILLS

|                                                                                                                  |  |
|------------------------------------------------------------------------------------------------------------------|--|
| <i>Languages</i> : Python, C/C++, Java, SQL, JavaScript/HTML/CSS                                                 |  |
| <i>ML/AI</i> : PyTorch, Hugging Face Transformers, Scikit-learn, FAISS, LoRA, SFT, GRPO, multimodal/VLM training |  |
| <i>Systems/Tools</i> : DeepSpeed ZeRO-3, Docker, Linux, Git, REST APIs, distributed training and checkpointing   |  |