DEMI RUOHAN WANG

(+1) 123-456-7890 | demiw@cs.cmu.edu | dem1tasse.github.io | Linkedin | GitHub

EDUCATION

Carnegie Mellon University *Aug 2025 - Dec 2026 (Expected)*
Master of Intelligent Information Systems, School of Computer Science GPA: 4.17 / 4.0

Tongji University *Sept 2020 - Jun 2025*
Bachelor of Software Engineering GPA: 3.95 / 4.0

EXPERIENCES

Graduate Research Assistant – CMU (Advisor: Prof. Graham Neubig) *Sep 2025 – Present*

- Built **code-based web agent** system using **OpenHands**, **Browser Use**, **BrowserGym**, implementing executable Playwright-based, domain-separated skill libraries for automated skill extraction, reuse, and composition.
- Achieved **7% absolute improvement** over action-based agents on **WebArena-Verified**, demonstrating that code generation can effectively serve as an **augmentation layer** for robust and transferable web automation.

Research Intern – Microsoft Research, Asia *Mar 2025 – Jul 2025*

- Investigated **GRPO** gradient dynamics in **RLVR** for **LLM reasoning**, revealing that off-policy degree fundamentally reshapes token-level update contributions via IS-ratio and clipping, reconciling conflicting conclusions in prior studies.
- Developed ACPO, a variance-aware clipping method on **VeRL**, improving training stability and showing robust performance across 3 tasks, 2 off-policy regimes, and 3B/7B model scales, achieving **12/18 first-place** over DAPO/CISPO.

Research Intern – Ohio State University (Advisor: Prof. Yu Su) *Apr 2024 – Nov 2024*

- Developed **UGround (200k+ Hugging Face downloads)**, a universal pixel-level visual grounding model for **GUI agents**, supporting robust **UI grounding** across Web, Mobile, and OS environments.
- Created a dataset of **9M** element examples from **773K** real-world website screenshots by designing an efficient synthetic data pipeline, combining web crawling and large language model annotation.
- Led model evaluations across multiple benchmarks (e.g. Mind2Web, AndroidControl, OmniAct), achieving **state-of-the-art** results with up to **36%** improvement in grounding accuracy over previous models.

Machine Learning Engineer Intern – ByteDance *Oct 2023 – Feb 2024*

- Fine-tuned LLaVA-based **vision-language models** with **LoRA**, designing a structured multi-step reasoning workflow on **100K** examples, boosting precision in detecting *off-platform traffic diversion violations* from **62.3%** to **90.2%**.
- Designed a self-supervised example selection pipeline for **in-context learning**, improving F1-Score on *livestream interaction violation* detection by **5.2%** and reducing manual review workload by **40%**.

PUBLICATIONS

[1] Navigating the Digital World as Humans Do: Universal Visual Grounding for GUI Agents
Gou B., **Wang R.**, Zheng B., Xie Y., Chang C., Shu Y., Sun H., Su Y. ICLR 2025 Oral (1.8%)

[2] What are Key Factors for Updates in RL for LLM Reasoning?
Wang P., **Wang R.**, Luo X., Xu J., Yang X., Feng S., Yang Y., Li D. Under submission.

SELECTED PROJECT

Miko – AI-Native Desktop Companion [Demo] *2nd Winner @AdventureX 2025 Kimi Track*

- Developed an **AI-native desktop agent** for productivity, capable of executing system-level and application tasks (e.g., app control, Gmail, Python execution, file operations, web search) through a conversational interface.
- Designed **tool orchestration** and **memory-augmented** context management for personalized task execution.

SKILLS

Languages & Tools	Python, C/C++, SQL, Shell, Docker, Git, AWS, Linux, Hadoop, VectorDB(FAISS)
ML/AI	PyTorch, Transformers, vLLM, RL (GRPO/VeRL), PEFT/LoRA, RAG, Distributed Training
Agent Systems	GUI/Computer-Use/Code Agents, Agent Skills, Tool Use, Benchmarking, UI Grounding