# 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

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

## EXPERIENCES

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

- Analyzed training signals for **GRPO** algorithm in **reinforcement learning** for **LLMs/SLMs**, developed unified formulation modeling probability distributions, advantages, importance ratio, and clipping to improve training stability.
- Developed token probability variance-based adaptive clip-range scheduling method in **VeRL** framework, achieving **3%** stable performance improvement across multiple benchmarks, model sizes, and off-policy levels.
- Designed a distribution-aware compression path search algorithm with **supervised fine-tuning (SFT)** to improve LLM reasoning efficiency, reducing response length by **30%** with **<2%** accuracy loss and enabling **1.5×** faster inference.

**Research Intern – Ohio State University** *Apr 2024 – Nov 2024*

- Developed **UGround**, a universal pixel-level visual grounding model to improve GUI agents, enabling robust **UI grounding** across diverse applications for **Computer Use Agents(CUA)**, Mobile Agents and OS Agents.
- 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**, applying **Chain-of-Thought** for multimodal reasoning 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%**.

## PUBLICATION

[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%)

## SELECTED PROJECTS

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

- Developed an **AI-native computer use agent** for desktop 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 a modular backend architecture supporting **multi-tool use and orchestration** for scalable task automation.
- Built a **memory-augmented conversation system** with context management and user preference learning.

**Life Buddy – AI Lifestyle Assistant** *2nd Winner @Baidu AGI HACKATHON*

- Led team to develop an **AI agent** for personalized restaurant, entertainment, and trip planning recommendations.
- Designed a **context-aware recommendation pipeline** integrating function calling, SQL queries, and **vector search**, enabling real-time, preference-based suggestions.

## SKILLS

| | |
|---|---|
| **Languages & Tools** | Python, C/C++, SQL, Shell, Docker, Git, AWS, Linux, Hadoop, VectorDB (FAISS, Pinecone) |
| **ML/AI Frameworks** | PyTorch, Transformers, Hugging Face, Scikit-learn, LangChain, Ray, PySpark, DeepSpeed |
| **Specialization** | Natural Language Processing, Large Language Models, PEFT, RAG, Distributed Training |
| **Agent Systems** | Computer Use/Code/GUI Agents, Reinforcement Learning (GRPO/VeRL), UI Grounding |