

Yuning Chen

✉ cyn10221017@gmail.com • [ycucm.github.io](https://github.com/ycucm) • [in LinkedIn](#) • [Google Scholar](#)

RESEARCH INTERESTS

LLM Post-training, Agentic RL, Memory-aware reward modeling, Generative recommendation and user simulation.

PROFESSIONAL EXPERIENCE

Research Scientist @ Meta, Bellevue, WA

May. 2025 - Now

Project: Semantic Personalization for Lead Ads Quality

- Initiated and led the development and deployment of a Qwen3-4B-based pipeline to analyze Q&A in user-submitted lead forms and infer commerce intent. Leveraged RL post-training (GRPO) to enhance cross-answer analysis and attribution, combining rubric-based rewards with LLM-as-a-judge to mitigate reward hacking. Implemented a hierarchical KV cache to boost throughput by 545.9%.
- Drove production adoption by partnering with product teams to integrate the output as proxy labels for multi-task learning in the ads reranking teacher model. Improved normalized entropy by 1.96% and achieved a statistically significant production A/B-test win, increasing sanitized ads value by 11.6%.

Project: Self-evolving Oncall Automation Agent with Episodic and Procedural Memory

- Initiated and built a Qwen3-8B-based long-horizon tool-calling agent for root-cause analysis and automated remediation across 40+ alert types, each requiring 10+ diagnostic steps; reduced response time by over 60% and saved 520 engineering hours annually.
- Designed a two-tier experience-memory architecture: episodic memory captured incident context, tool calls, diagnostic decisions, outcomes, and supporting work-chat evidence; procedural memory consolidated recurring successful trajectories into reusable and retrievable debugging skills.
- Developed memory compression and consolidation policies to classify new incidents, deduplicate related experiences, identify novel failure patterns, and trigger skill induction for previously unseen alert categories; improved long-term task Pass@1 by 17.6%.
- Designed the post-training and evaluation framework, combining trajectory-level SFT with step-level Agentic RL (HGPO with stepwise reward assignment) over 42K roll-out trajectories, reducing reward hacking and improving Pass@1 by 5.9%.

Project: Generative Recommender as User Simulator for Ads Ranking Training Data Synthesis

- Defined the scientific formulation of generative recommenders as high-fidelity user behavior simulators, using SFT + GSPO-based RL post-training for stable long-sequence generation and new-user synthesis to expand cold-start, long-tail, and OOD coverage.
- Designed the training and evaluation framework for synthetic traces, demonstrating transferable training signals for downstream sequential recommenders with 8.4%-13.1% Pass@32 improvements while supporting offline simulator-based RL policy learning.
- Submitted a research paper to NeurIPS'26 and received high review ratings.

Software Engineer Intern @ Bytedance, San Jose, CA

May. 2024 - Aug. 2024

Machine Learning for Systems, Network Engineering.

- Designed and developed a proactive ingress behavior probing tool for TikTok DNS and CDN users on IP anycast. After the successful deployment on 10+ edges globally, an ingress report for more than 240k DNS resolver IPs can be generated within 10 minutes.

Research Intern @ Alibaba Cloud U.S., Sunnyvale, CA

May. 2023 - Aug. 2023

Code Generation, Large Language Model Evaluation, Cloud Computing, IP Anycast.

- Proposed a coding benchmark with 1011 problems for cloud-native configuration generation and a scalable evaluation platform for LLM assessment, and analyzed the performance of 13 LLMs, offering strategies for quality and efficiency improvement.
- Built an IP anycast optimization system that formulates the AS path prepending optimization as a weighted MAX-SAT problem and proposed an efficient scanning method to obtain the potential connections for global IP anycast clients.

EDUCATION

Ph.D. in Electrical Engineering and Computer Science, **University of California, Merced**

Aug. 2021 - May. 2025

Advisor: Prof. [Wan Du](#). Research Topic: Machine Learning; Generative AI; Networked Systems.

B.Eng. in Computer Software Engineering, **Tsinghua University**

Aug. 2017 - May. 2021

Advisor: Prof. [Mingsheng Long](#). Honor: Outstanding Bachelor's Graduate (10%).

SELECTED PUBLICATIONS

* denotes equal contribution

- Yifei Xu*, Yuning Chen*, Xumiao Zhang*, Xianshang Lin, Pan Hu, Yunfei Ma, Songwu Lu, Wan Du, Z. Morley Mao, Dennis Cai, and Ennan Zhai. **CloudEval-YAML: A Practical Benchmark for Cloud Configuration Generation**. MLSys 2024. [Code](#)
- Yuning Chen, Kang Yang, Zhiyu An, Brady Holder, Luke Paloutzian, Khaled M. Bali, Wan Du. **MARLP: Time-series Forecasting Control for Agricultural Managed Aquifer Recharge**. ACM KDD 2024. [Code](#)
- Kang Yang, Yuning Chen, Wan Du. **Generalizable Radio-Frequency Radiance Fields for Spatial Spectrum Synthesis**. IEEE/CVF CVPR 2026. [Code](#)

- Minyuan Zhou*, Yuning Chen*, Yifei Xu, Jiaqi Zheng, Guihai Chen, Wanchun Dou, Pan Hu, Yongping Tang, Wendong Yin, Jie Lin, Qingyan Yu, Yuanchao Su, Songwu Lu, Wan Du. **AnyPro: Preference-Preserving Anycast Optimization based on Strategic AS-Path Prepending**. USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2026. [Code](#)
- Kang Yang*, Yuanlin Yang*, Yuning Chen*, Sikai Yang, Wan Du. **SoilX: Calibration-Free Comprehensive Soil Sensing Through Contrastive Cross-Component Learning**. ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys) 2026. [Best Paper Award](#). [Code](#)
- Kang Yang, Yuning Chen, and Wan Du. **OrchLoc: In-Orchard Localization via a Single LoRa Gateway and Generative Diffusion Model-based Fingerprinting**. ACM MobiSys 2024 (Extended version published in ACM TMC). [Code](#)
- Yifei Xu*, Yuning Chen*, Xumiao Zhang*, Xianshang Lin, Pan Hu, Yunfei Ma, Songwu Lu, Wan Du, Z. Morley Mao, Ennan Zhai, and Dennis Cai. **CloudEval-YAML: A Realistic and Scalable Benchmark for Cloud Configuration Generation**. Machine Learning for Systems Workshop, NeurIPS 2023. [Code](#)
- (1st author, under submission) **MEMOIR: Disentangling Memory Deficiency from Execution Error in Agentic Reinforcement Learning**.
- (1st author, under submission) **Generative Recommenders are High Fidelity User Simulators**.
- (1st author, under submission) **From Lead-Form Q&A to Commercial Intent: RL Post-Training for Calibrated Intent Scoring**.
- (1st author, under submission) **GenRTC: Enabling Generative Visual Captions in Live Video Streaming**.

TECHNICAL SKILLS

- **Programming:** Python, C++, Go, JavaScript, SQL, MATLAB.
- **ML & AI:** LLM post-training, supervised fine-tuning, reinforcement learning, reward modeling, agentic systems, agentic memory, large embedding models, PyTorch, vLLM, Hugging Face, TensorFlow.
- **Systems & Infrastructure:** Linux, Kubernetes, Git, WebRTC, OpenCV, FFMPEG, MySQL.
- **LLM and RL Systems:** VeRL, VeRL-Agent, Miles, Harbor, SLIME, OpenEnv.

ACADEMIC SERVICES

- **Conference Area Chair:** KDD'25-26 (Applied AI).
- **Conference Program Committee:** MLSys'26 (Agentic Systems).
- **Conference Reviewer:** NeurIPS'24-25, KDD'25-26, ICML'25, ICLR'25-27, AISTATS'25-26, MM'25-26, IMWUT'26, CIKM'24 TRAIL.
- **Conference Artifact Evaluation Committee:** MobiCom'24-25, MobiSys'24-25.
- **Journal Reviewer:** IEEE TPAMI, IEEE TMC, IEEE TNSE, ACM TOSN, IEEE IoTJ, ACM TIOT.