

Yuanqing Wang

Email: yuanqing.wang@utdallas.edu

Mobile: +1 (945) 713-1976

<https://wang-yuanqing.github.io/>

EDUCATION

- **The University of Texas at Dallas** Richardson, TX
Ph.D. Student in Mathematics; Advisors: Prof. Baris Coskunuzer Aug. 2024 – Present
 - GPA: 3.77/4.00.
 - Selected Coursework: Numerical Linear Algebra (A), Deep Learning (A).
- **Capital Normal University** Beijing, China
M.S. in Mathematics; GPA: 3.86/4.00 Sep. 2021 – Jun. 2024
 - Thesis: Calabi Conjecture and Kähler-Einstein Metric. Graduate training in differential geometry, topology, and analysis.
- **Xiamen University** Xiamen, China
B.S. in Information and Computational Science; Elite Students' Program Sep. 2016 – Jun. 2021
 - Elite Students' Scholarship. China Undergraduate Mathematical Contest in Modeling, 1st Prize in Fujian Province (2018).

RESEARCH INTERESTS

My research focuses on self-improving LLM agents, particularly credit assignment, long-horizon agentic reinforcement learning, and the evaluation of agent-proposed improvements.

SELECTED RESEARCH

- **dwDPO: Training Multi-Turn LLM Agents via Divergence-Weighted Direct Preference Optimization** *EMNLP 2026 Main*
Jinghao Lin^{*}, Yuhang Wu^{*}, **Yuanqing Wang^{*}**, Yuchen Li, Kangtianxingjian, Baris Coskunuzer[†], Xiawu Zheng[†]
 - Proposed divergence-weighted DPO for multi-turn credit assignment; outperformed standard DPO on all 6 long-horizon benchmarks, improving average accuracy from 35.9 to 37.0 at no added inference cost.
 - **Contribution:** Co-first author; conceived the entire method, developed all mathematical formulations and theoretical analysis, and led and coordinated collaborators in executing the experiments.
- **Skill-CDPO: Evolving Agent Tool-Use via Critical Step Preference Optimization** *EMNLP 2026 Findings*
Yuchen Li, Jinghao Lin, **Yuanqing Wang**
 - Introduced Critical Step DPO (CDPO), which identifies capability bottleneck steps via rollout divergence and constructs distributional preference pairs weighted by step-level criticality and pair-level score gaps. 8B model matched or exceeded GPT-5.2 on medical agent benchmarks.
 - **Contribution:** Designed the CDPO method, mathematical formulations, and theoretical analysis.
- **Heat Field Signatures: Density-Aware Geometry for Point Cloud Classification** *In Submission*
Yuanqing Wang, Yapeng Tian[†], Baris Coskunuzer[†]
 - Uses differential geometry and training-free heat-field features for point-cloud classification. HFS variants lead seven main benchmarks; HFS-full achieves 59.6% versus 35.7% mean class accuracy for Point Transformer on SCOP, with an 8.5× median end-to-end speedup over DGCNN across four datasets.
 - **Contribution:** Implemented the proposed method, selected the evaluation datasets, discovered its rotation-invariance property, ran all experiments, analyzed the results, and drafted the experimental section of the paper.
- **Multimodal Learning and Evaluation** *In Submission, ACL ARR*
First author
 - Research on multimodal understanding and model evaluation.

EARLIER RESEARCH

- **Solving Differential Equations Using Artificial Neural Networks** Xiamen University
National Undergraduate Innovation Project (Lead); Advisor: Prof. Chuanju Xu Mar. 2019 – Apr. 2020
 - Led a national-level undergraduate innovation project. Implemented MLP-based neural network solvers for PDEs with prescribed boundary conditions, among the early efforts applying neural networks to PDE problems.

INDUSTRY EXPERIENCE

- **Gathssen Investment Co., Ltd.** Shenzhen, China
Quantitative Strategy Developer Intern *Jun. 2021 – Aug. 2021*
 - Built a backtesting framework for dollar-cost averaging strategies and benchmarked multiple variants. Contributed to the development of a cash-and-carry arbitrage strategy.
 - Independently developed a cointegration-based cross-asset hedging strategy with full backtesting pipeline, from research to simulation.

TECHNICAL SKILLS

- **LLM Post-Training:** verl (GRPO, DPO), preference optimization, rollout and trajectory generation, benchmark design and evaluation
- **LLM Systems:** PyTorch, Hugging Face Transformers, vLLM, multimodal model deployment, Gemini API
- **Systems & Infrastructure:** Linux, Git, Slurm, Apptainer, H100/H200/GH200 GPU clusters
- **Programming & Tools:** Python, NumPy, pandas, Gradio

HONORS & AWARDS

- **NSM Summer Research Fellowship, UT Dallas** Summer 2026
- **NSM McDermott PhD Admission Fellowship, UT Dallas** Sep. 2024
- **Excellent Academic Scholarship, Capital Normal University** Dec. 2022 – Dec. 2023
- **China Undergraduate Mathematical Contest in Modeling, 1st Prize (Fujian Province)** Sep. 2018
- **Elite Students' Scholarship, Xiamen University** Sep. 2018 – Sep. 2019

TEACHING

- **Teaching Assistant, UT Dallas — Linear Algebra, Differential Calculus** 2025 – 2026
- **Teaching Assistant, Capital Normal University — Real Analysis, Complex Analysis** 2023 – 2024