Tianyu (Anna) Yang

tyang4@nd.edu LinkedIn GitHub Personal Website Google Scholar

Overview

PhD candidate in Computer Science with 20 publications at top ML/NLP venues (ACL, EMNLP), including **11 first author**. Applied this research as a research intern at Zillow, Samsung Research, and Bosch AI Center, building Self-Evolving Agents, RL-based long-term memory consolidation, and agentic multimodal RAG systems. **Research focus: post-training LLMs/MLLMs for agents that perceive, reason, and act.**

Education

Ph.D., Computer Science Sep 2023 – 2027 (expected)
University of Notre Dame South Bend, IN
Advisor: Prof. Xiangliang Zhang; GPA: 3.90/4.0

M.S. (Research), Electrical Engineering Sep 2021 – May 2023
Columbia University New York, NY
Advisor: Prof. Shih-Fu Chang; GPA: 3.90/4.0

B.S., Automation Sep 2017 – May 2021
Beijing Jiaotong University Beijing, China
GPA: 3.90/4.0

Honors and Awards

Samsung Best Paper Award (3/9,000 company-wide) 2026
Research Fellowship, University of Notre Dame 2023
Research Fellowship, Columbia University (1 of 3 fellows selected program-wide) 2021
National Scholarship, China (0.2% selection rate) 2019

Work Experience

Zillow: Applied Scientist Intern Jun 2026 – Oct 2026
Reward-Driven Alignment Drift in Self-Evolving LLM Agents Seattle, WA

Investigated reward-driven alignment drift in self-evolving LLM agents through multi-round SFT, GRPO, and self-rewarded DPO across reasoning and agentic tasks, analyzing how repeated optimization affects model behavior and alignment robustness. Developed SEA, a neuron-level interpretability framework that identifies sparse refusal circuits through differentiable mask optimization and tracks their evolution across training checkpoints. Proposed Neuron-Targeted Alignment (NTA), using masked LoRA to constrain SFT/DPO updates to identified circuits. Across three LLMs, NTA reduced harmfulness by 63.21% and benign-query over-refusal by 58.44% while preserving 99.58% of original model utility.

Resulting publication: When Rewards Reshape Alignment: Drift in Self-Evolving Agents

LLM Post-Training *Self-Evolving Agents*

Samsung Research: Research Scientist Intern Jan 2026 – May 2026
TrustMem: RL-Based Long-Term Memory Consolidation for LLM Agents Mountain View, CA

Samsung Best Paper Award (3/9,000 company-wide)

Proposed TrustMem, an RL-based post-training framework that learns reliable memory-editing policies for long-term LLM agents. Developed a Memory Transition Verifier to assess information coverage, preservation, and faithfulness at each memory update. Introduced Transition-Ranked GRPO (TR-GRPO), combining multi-granularity rewards with verifier-guided preference ranking to optimize memory consolidation. Achieved SOTA across long-term memory benchmarks, improving MemoryAgentBench by 6.5 points and HaluMem by 12.14 F1 points while reducing transition-level memory corruption by 79.1%.

Resulting publication: TrustMem: Learning Trustworthy Memory Consolidation for LLM Agents with Long-Term Memory

Agent Memory *Reinforcement Learning* *long-horizon Agents* *LLM Post-Training*

Bosch AI Center: Research Scientist Intern May 2025 – Aug 2025
MM-R2: Reinforcement Learning for Agentic Multimodal Retrieval Pittsburgh, PA

Developed MM-R2, an agentic multimodal retrieval framework that learns evidence-selection and source-routing decisions through reinforcement learning. Designed Intent-Referent Binding and KnowledgeMap routing for grounded, relevance-aware retrieval, and constructed MM-R2-Traj, a large-scale dataset of 900K retrieval trajectories for agent post-training. Post-trained Qwen-VL-7B using SFT and GRPO with routing and answer-correctness rewards, achieving 54.3% InfoSeek accuracy (+7.8 points over the strongest baseline) and improving Route Recall@1 from 23.4% to 74.0%.

Resulting publication: Reason Before You Retrieve: Agentic Planning for Multi-modal RAG

RAG *Reinforcement Learning* *LLM Post-Training*

Google: Open-Source Contributor May 2023 – Aug 2023
G-API Algorithms for OpenCV Remote

- Developed Python tutorials and new machine learning algorithms for OpenCV's G-API module, extending official documentation and improving image-processing performance.

Academic Experience

AXIS^2: Two-Dimensional Latent Generation for Language Modeling Mar 2026 – Sep 2026

Proposed AXIS^2, a novel language generation paradigm that moves beyond conventional next-token prediction by reformulating generation as a two-dimensional process combining coarse-to-fine latent prediction and progressive within-resolution refinement. Designed a residual-quantized latent hierarchy and a lightweight refinement module that reuses cached backbone representations to model dependencies among latent variables without repeated full-model evaluations. Achieved up to 24.3% higher ROUGE-1 than a strong block-diffusion baseline (BD3-LM) under a matched training budget, with 10.4× faster inference than BD3-LM (0.55s vs. 5.72s/sample) and 15.5× faster inference than autoregressive decoding on a single H100.

Language Modeling *Latent Generation* *Efficient Inference*

Publications

** Co-first authors.*

[Preprint'26] **Of Errors and Echoes: Responsibility-Aware Credit Assignment for Tool-Use Agents**
Xingjian Diao*, **Tianyu Yang***, Wenjun Huang, Xiangchi Yuan, Chunhui Zhang, Xingchen Zhao, Weiyi Wu, Haotian Xu, Soroush Vosoughi, Xiangliang Zhang†, Jiang Gui†

[Preprint'26] **Reason Before You Retrieve: Agentic Planning for Multi-modal RAG**
Tianyu Yang, Minghao Chen, Zhenzhen Li, Simon Shir, Xiangliang Zhang

[Preprint'26] **When Rewards Reshape Alignment: Drift in Self-Evolving Agents**
Tianyu Yang, Austin Tian, Reid A. Johnson, Baoyuan Wang, Xiangliang Zhang

[Preprint'26] **AXIS^2: Two-Dimensional Latent Generation for Language Modeling**
Tianyu Yang et al.

[Preprint'26] **Agentic Multimodal RAG: Roles, Decisions, and Evaluation**
Tianyu Yang, Wenliang Guo, Xiujin Liu, Shijian Deng, Yilun Zhao, Yapeng Tian, Xiangliang Zhang

[Preprint'26] **DesignAgent: Interactive Object-Level Indoor 3D Scene Editing via Multimodal Agentic Reasoning**
Tianyu Yang*, Xiujin Liu*, Chen Zhao, Yilun Zhao, Xiangliang Zhang

[EMNLP'26] **TrustMem: Learning Trustworthy Memory Consolidation for LLM Agents with Long-Term Memory**
Tianyu Yang, Sudipta Paul, Vijay Srinivasan, Vivek Kulkarni, Srinivas Chappidi

[ACL'26] **Deconstructing Multimodal Mathematical Reasoning: Towards a Unified Perception-Alignment-Reasoning Paradigm**
Tianyu Yang, Sihong Wu, Yilun Zhao, Lisen Dai, Chen Zhao, Minhao Cheng, Arman Cohan, Xiangliang Zhang

[EMNLP'25] **Quest2DataAgent: Automating End-to-End Scientific Data Collection**
Tianyu Yang, Jason Rohr, Brett Scheffers, Nitesh Chawla, Xiangliang Zhang

[ACL'25] **CLIPErase: Efficient Unlearning of Visual-Textual Associations in CLIP**
Tianyu Yang, Lisen Dai, Xiangqi Wang, Minghao Chen, Yapeng Tian, Xiangliang Zhang

[EMNLP'24] **SaSR-Net: Source-Aware Semantic Representation Network for Enhancing Audio-Visual Question Answering**
Tianyu Yang, Yiyang Nan, Lisen Dai, Zhenwen Liang, Yapeng Tian, Xiangliang Zhang

[Preprint'26] **ADAPT: Audio-Depth Action Planning with Transition-Aware Rewards for Lights-Out Industrial Navigation**
Wenjun Huang, Shenghao Fu, Yian Jin, . . . , **Tianyu Yang**, et al.

[Preprint'26] **Seeing, Maintaining, and Learning: The Evolution of Multimodal Agent Memory**
Minghao Guo, Qingyue Jiao, Xiujin Liu, . . . , **Tianyu Yang**, et al.

[COLM'26] **Explainable AI-Generated Image Detection RewardBench**
Michael Yang, Shijian Deng, William T. Doan, Kai Wang, **Tianyu Yang**, Harsh Singh, Yapeng Tian

[WSDM'25] **WildlifeLookup: A Chatbot Facilitating Wildlife Management with Accessible Data and Insights**
Xiangqi Wang, **Tianyu Yang**, Jason Rohr, Brett Scheffers, Nitesh Chawla, Xiangliang Zhang

[EMNLP'25] **Self-Improvement in Multimodal Large Language Models: A Survey**
Shijian Deng, Kai Wang, **Tianyu Yang**, Harsh Singh, Yapeng Tian

[ACL'25] **Physics: Benchmarking Foundation Models for PhD-Qualifying Exam Physics Problem Solving**
Kaiyue Feng, Yilun Zhao, **Tianyu Yang**, Yixin Liu, Chen Zhao, John Sous, Arman Cohan

[ACL'25] **Beyond Single-Value Metrics: Evaluating and Enhancing LLM Unlearning with Cognitive Diagnosis**
Yicheng Lang, Kehan Guo, Yue Huang, Yujun Zhou, Haomin Zhuang, **Tianyu Yang**, Yao Su, Xiangliang Zhang

[ACL'24] **SceMQA: A Scientific College Entrance Level Multimodal Question Answering Benchmark**
Zhenwen Liang, Kehan Guo, Gang Liu, Taicheng Guo, Yujun Zhou, **Tianyu Yang**, Jiajun Jiao, Renjie Pi, Jipeng Zhang, Xiangliang Zhang

[EMNLP'23] **UniMath: A Foundational and Multimodal Mathematical Reasoner**
Zhenwen Liang, **Tianyu Yang**, Jipeng Zhang, Xiangliang Zhang

Service

- **Area Chair (2026)**: EMNLP, NAACL, ICASSP; **Reviewer (2024–2026)**: ACL, EMNLP, NAACL, NeurIPS, ICML, ICLR, KDD, CVPR, ECCV, ICCV, COLM.
- **Program Chair / Organizer (2025–2026)**: AAAI Workshop on AI for Science; ICCV Workshop on Knowledge-Intensive Multimodal Reasoning; **Tutorial Organizer (2025)**: CVPR Multimodal Mathematical Reasoning.

Skills

Research: LLM/MLLM Post-Training (SFT, DPO, PPO, GRPO), Reward Design & Preference Optimization, Agentic RL, Agentic RAG, Tool Use & Multi-Agent Systems, Long-Term Memory Agents, Multimodal Reasoning, Model Evaluation, Synthetic Data

Frameworks: PyTorch, Hugging Face (Transformers, TRL, PEFT/LoRA), verl, Ray, vLLM, SGLang, Accelerate, FAISS, Neo4j

Distributed Training & RL Infrastructure: Multi-Node Training, PyTorch DDP/FSDP, DeepSpeed ZeRO, Distributed RL Rollouts, Slurm, torchrun, NCCL, Kubernetes