

**Research Interests:** *Natural Language Processing, Foundation Models / Large Language Models (LLMs), Reinforcement Learning, LLM Reasoning, Context Management, Efficient Reasoning, Game AI*

## EDUCATION

- University of California, San Diego** September 2024—Expected: June, 2030
- PhD in Computer Science and Engineering
  - Advisor: **Dr. Prithviraj Ammanabrolu**
- Georgia Institute of Technology** August 2021—May 2024
- M.S. in **Interactive Intelligence**; GPA: 4.00/4.00
  - Advisors: **Dr. Mark O. Riedl, Dr. Thad Starner**
- University of North Carolina, Chapel Hill** August 2016—May 2020
- B.S in Computer Science; GPA: 3.72/4.00

## RESEARCH EXPERIENCE

- Research Intern** **Incoming: June 2026 — August 2026**  
*Microsoft Research, under Marc-Alexandre Côté & Xingdi “Eric” Yuan* *Montréal, Québec*
- Incoming summer research internship continuing collaboration on LLM agents and situated text environments.
- Graduate Research Assistant** **September 2024 — Current**  
*University of California, San Diego, under Dr. Prithviraj Ammanabrolu* *San Diego, California*
- First author of **Behavior Cue Reasoning: Monitorable Reasoning Improves Efficiency and Safety through Oversight** (Cui, Killian, Ammanabrolu — pre-print): led all implementation of a training method that anchors atomic LLM reasoning behaviors to special token sequences (*behavior cues*, e.g. [continue], [stop], [answer]), making chain-of-thought reasoning monitorable and externally controllable.
    - Reduced wasted reasoning tokens by up to **50%** under an efficiency objective.
    - **More than doubled** success rate under a safety-constrained objective by recovering safe actions from otherwise-unsafe reasoning traces **80%** of the time.
    - Trained and evaluated **Qwen3-8B** and **GLM-Z1-9B** on **AIME**, **TextWorld**, and **HazardWorld** (via a custom LLM-compatible text-wrapper that we built).
  - First author of **TALES: Text Adventure Learning Environment Suite** (Cui, Yuan, Xiao, Ammanabrolu, Côté — **Oral, REALM 2025**; non-archival): a progressive stress test of LLMs’ baseline agentic capabilities, spanning **122 tasks** across which **49 models** were benchmarked. Led paper writing and agent trace analysis; collaborated with Microsoft Research on codebase development.
  - Lead system installation and configuration for the lab’s local GPU compute clusters in partnership with UCSD campus IT.
- AI Engineer** **June 2025 — September 2025**  
*MBZUAI Institute of Foundation Models, under Dr. Taylor W. Killian* *Sunnyvale, California*
- Investigated context management strategies for multi-turn RL training of **LLM agents** — specifically the effect of retaining past reasoning traces in-context and combining them with self-critique reasoning. Found **up to 20% in-distribution performance gains** on agentic text-environment benchmarks.
  - Trained **Qwen3-8B** with **PPO** atop the **verl-agent** framework across four agentic text environments (TextWorld, TextWorld Express, ALFWorld, Lifegate) on H200 clusters.
  - Built a custom stochastic text environment (Lifegate) for stress-testing risk-sensitive RL training of LLM agents.
- Graduate Research Assistant** **January 2023 — July 2024**  
*Georgia Tech, under Dr. Thad Starner* *Atlanta, Georgia*
- Assisted in securing funding by presenting progress in development and deployment of Edtools to various faculty, including the past and present Deans of the College of Computing.
  - Research Lead of **Socratic Mind**: Using **Foundation Models** to create scalable, interactive, oral assessments.
    - Winner of the **2023-24 Tools Competition in Educational Technologies**
    - Garnered interest from 15+ professors across 7 universities
    - Successfully served over 1,000 college students in graduate and undergraduate level courses
  - Led research teams investigating **neural network** approaches to automatic plagiarism detection, anti-plagiarism tooling for online exams, and digital watermarks for coding assignments.
  - Led research teams on classroom uses of foundation models and on **Impasse**, an IDE extension for detecting student difficulty in coding assignments.
- Graduate Research Assistant** **March 2022 — May 2024**  
*Georgia Tech, under Dr. Mark O. Riedl* *Atlanta, Georgia*
- Explored the potential of combining **prompting with Foundation Models** with symbolic-based planning approaches to produce coherent story plot skeletons by using GPT-J with prompt-chaining to perform commonsense inference
  - Extended **LIGHT Textworld**’s map generation by streamlining map creation for user-generated layouts
  - Trained and tested **Reinforcement Learning Agents** in the **LIGHT Textworld** environment by developing a wrapper to integrate various agent architectures with the **LIGHT** game engine.
  - Explored few-shot training of **generalized Reinforcement Learning** agents via specialized prompt tuning and **attention-based expert guidance** for ensemble learning.

## PUBLICATIONS

### ***Behavior Cue Reasoning: Monitorable Reasoning Improves Efficiency and Safety through Oversight***

*Christopher Z. Cui, Taylor W. Killian, Prithviraj Ammanabrolu*  
Pre-print

### ***TALES: Text Adventure Learning Environment Suite***

*Christopher Z. Cui, Xingdi Yuan, Ziang Xiao, Prithviraj Ammanabrolu, Marc-Alexandre Côté*  
Oral Presentation, REALM 2025 (non-archival)

### ***Leveraging Past Assignments to Determine If Students Are Using ChatGPT for Their Essays***

*Yuhui Zhao, Chunhao Zao, Rohit Sridhar, Christopher Cui, Thad Starner*  
Published in *Proceedings of the Tenth ACM Conference on Learning@ Scale 2024*

### ***Examinator v4.0: Cheating Detection in Online Take-Home Exams***

*Christopher Cui\*, Jui-Tse Hung\*, Vaibhav Malhotra, Hardik Goel, Raghav Apoorv, Thad Starner (\*shared first author)*  
Published in *Proceedings of the Tenth ACM Conference on Learning@ Scale 2024*

### ***Socratic Mind: Scalable Oral Assessment Powered By AI***

*Jui-Tse Hung, Christopher Cui, Diana M. Popescu, Saurabh Chatterjee, Thad Starner*  
Published in *Proceedings of the Tenth ACM Conference on Learning@ Scale 2024*

### ***Answer Watermarking: Using Answer Generation Assistance Tools to Find Evidence of Cheating***

*Christopher Cui, Jui-Tse Hung, Pranav Sharma, Saurabh Chatterjee, Thad Starner*  
Published in *Proceedings of the Tenth ACM Conference on Learning@ Scale 2024*

### ***A Mixture-of-Experts Approach to Few-Shot Task Transfer in Open-Ended Text Worlds***

*Christopher Cui, Xiangyu Peng, Mark Riedl*  
Pre-print

### ***Thespian: Multi-Character Text Role-Playing Game Agents***

*Christopher Cui, Xiangyu Peng, Mark O. Riedl*  
Published in *Proceedings of the AIIDE Workshop on Experimental AI in Games (EXAG) 2023*

### ***Story Shaping: Teaching Agents Human-like Behavior with Stories***

*Xiangyu Peng\*, Christopher Cui\*, Wei Zhou, Renee Jia, Mark O. Riedl (\*shared first author)*  
Published in *Proceedings of the 19th AAI Conference on Artificial Intelligence in Interactive and Digital Entertainment (AIIDE) 2023*

### ***Examinator v3.0: Cheating Detection in Online Take-Home Exams***

*Jui-Tse Hung, Christopher Cui, Varun Agarwal, Saurabh Chatterjee, Raghav Apoorv, Rocko Graziano, Thad Starner*  
Published in *Proceedings of the Tenth ACM Conference on Learning@ Scale 2023*

### ***Neural Story Planning***

*Anbang Ye, Christopher Cui, Taiwei Shi, Mark O. Riedl*  
Published in *Proceedings of the AAI Workshop on Creative AI Across Modalities 2023*

## SERVICE & REVIEWING

- **Gold-tier Reviewer (top 25%)**, ICML 2026
- Reviewer, NeurIPS 2025
- Reviewer, ICLR 2026

## TEACHING EXPERIENCE

### **Graduate Teaching Assistant**

*University of California, San Diego*

**January 2026 — June 2026**

*San Diego, California*

- CSE 190, Introduction to Deep Reinforcement Learning (Spring 2026)
- CSE 291, AI Agents (Winter 2026)

### **(Head) Graduate Teaching Assistant**

*Georgia Institute of Technology*

**August 2022 — May 2024**

*Atlanta, Georgia*

- CS3600/6601, Artificial Intelligence — Head TA: Spring 2024 (*Dr. Thad Starner*), Fall 2023 (*Dr. Thomas Ploetz*); TA: Spring 2023 (*Dr. Thad Starner*), Fall 2022 (*Dr. Thomas Ploetz*)
- CS7650, Natural Language Processing — TA: Summer 2023 (*Dr. Mark Riedl*)

## SKILLS

- Technical: Large Language Models, Reinforcement Learning, RLHF / PPO / GRPO, Natural Language Processing, Distributed Training, LLM Reasoning, AI Agents
- Languages: Python, C/C++, Java, JavaScript
- Tools and Frameworks: PyTorch, Hugging Face Transformers, verl / verl-agent, OpenRLHF, vLLM, Ray, DeepSpeed