

Zhangqi (Gavin) Duan

✉ zduan@cs.umass.edu | 🌐 [Website](#) | [LinkedIn](#) | [Google Scholar](#)

Amherst, MA

Summary Ph.D. student in Computer Science at UMass Amherst specializing in LLM training, alignment, and evaluation, with a focus on education technology and human-centered AI systems. Experienced in LLM fine-tuning (SFT), RL-based optimization (GRPO, DPO), reward design, and multitask training for interpretable model behavior. Published in leading venues (EMNLP, LAK, AIED, and WebSci) with a strong record of translating research prototypes into robust ML systems for real-world NLP and reasoning tasks.

Education **University of Massachusetts Amherst**

- PhD in Computer Science (Advisor: [Prof. Andrew Lan](#)), GPA: 3.83/4.00 2024.9 – Present
- MS in Computer Science, GPA: 3.82/4.00 2022.9 – 2024.5
- BS in Computer Science & BS in Mathematics, GPA: 3.85/4.00 2018.5 – 2022.5

Work **UMass Amherst Center for Data Science**

Amherst, MA

Experience Machine Learning Engineer Intern

2023.5 – 2023.8

- Analyzed 6.5M tweets using scalable data processing methods to detect key patterns and outliers
- Applied Python Numpy, Pandas modules to perform advanced data cleaning and feature engineering resulted in 23% decrease in model runtime
- Built a Topic Model to categorize diverse topics, uncovering unknown events within Non-English textual data
- Implemented a User Model pipeline to target primary contributors to ongoing conversations with unsupervised learning algorithms (Spherical K-means clustering, hierarchical clustering)

Siemens Ltd

Beijing, China

Software Engineer Intern

2019.5 – 2019.8

- Conducted database creation for 10,000 clients using SQL Subqueries/Merge Statement
- Resolved technical issues reported by clients or detected through internal monitoring
- Performed code reviews and provided feedback to team members

Publications **Knowledge-Aligned Student Error Simulator for Open-Ended Coding Tasks**

Zhangqi Duan, Nigel Fernandez, Andrew Lan

ACL, 2026 (Under Review)

Automated Knowledge Component Generation for Interpretable Knowledge Tracing in Coding Problems

Zhangqi Duan, Nigel Fernandez, Arun Narayanan, Mohammad Hassany, Rafaella Alencar, Peter Brusilovsky, Bitu Akram, Andrew Lan

ACL, 2026 (Under Review)

Using LLMs for Knowledge Component-level Correctness Labeling in Open-ended Coding Problems

Zhangqi Duan, Arnav Kankaria, Dhruv Kartik, Andrew Lan

AIED, Short Paper, 2026

Test Case-Informed Knowledge Tracing for Open-ended Coding Tasks

Zhangqi Duan, Nigel Fernandez, Alexander Hicks, Andrew Lan
LAK, 2025

Multilingualism, Transnationality, and K-pop in the Online #StopAsianHate Movement
Tessa Masis, **Zhangqi Duan**, Weiai Wayne Xu, Ethan Zuckerman, Jane Yeahin Pyo, Brendan O'Connor
WebSci, 2025

README: Bridging Medical Jargon and Lay Understanding for Patient Education through Data-Centric NLP
Zonghai Yao, Nandyala Siddharth Kantu, Guanghao Wei, Hieu Tran, **Zhangqi Duan**, Sunjae Kwon, Zhichao Yang, README annotation team, Hong Yu
EMNLP Findings, 2024

Agent Performing Autonomous Stock Trading under Good and Bad Situations
Yunfei Luo, **Zhangqi Duan**
AI4ABM (ICLR Workshop), 2023

Research Experience **LLM Training with RL for Code Generation** 2025.10 – 2026.1
AI for Education

- Developed **RL-based** methodologies (**GRPO**) for student code simulation, improving controllability and diversity of generated codes
- Designed a custom **RL reward** function combining code similarity metrics and error level IoU to better align generation with student error distributions
- Evaluated SFT-only vs. SFT+RL training at scale, yielding a **13.1%** improvement in error coverage
- Built an end-to-end LLM training and evaluation pipeline in PyTorch, supporting reproducible experimentation across reward variants and datasets

Multitask LLM Training for Interpretable Prediction and Generation 2024.9 – 2025.1
AI for Education

- Created the first LLM-based pipeline for knowledge component generation, validated against expert-level quality
- Designed a **soft-token conversion mechanism** to support end-to-end gradient flow through multiple models
- Developed KC-Gen KT, a **multitask learning** framework jointly optimizing code generation and correctness prediction, integrating interpretable student representations
- Achieved a **5%** AUC improvement on prediction tasks and a **4%** CodeBLEU gain on code generation tasks over state-of-the-art models, demonstrating the effectiveness of KC-Gen KT across diverse tasks

Multitask LLM Training for Interpretable Prediction and Generation 2024.9 – 2025.1
AI for Education

- Created the first LLM-based pipeline for knowledge component generation, validated against expert-level quality
- Designed a **soft-token conversion mechanism** to support end-to-end gradient flow through multiple models
- Developed KC-Gen KT, a **multitask learning** framework jointly optimizing code generation and correctness prediction, integrating interpretable student representations
- Achieved a **5%** AUC improvement on prediction tasks and a **4%** CodeBLEU gain on code generation tasks over state-of-the-art models, demonstrating the effectiveness of KC-Gen KT across diverse tasks

Retrieval-Augmented Lay Definition Generation in BioNLP

Healthcare NLP 2021.9 – 2022.1

- Improved patient comprehension of clinical notes using natural language generation
- Applied pretrained sequence-to-sequence models (BART) for controlled definition generation
- Built data pipelines to handle linguistic ambiguity and domain shift in medical text
- Deployed a full-stack NLP application using React and Django to deliver real-time explanations to end users

Teaching Experience **Teaching assistant, College of Info. and Computer Science** 2022.9, 2023.1
CS 240: Reasoning Under Uncertainty

- Assigned different tasks such as checking assignment and quiz correctness to other TAs and UCAs as head TA
- Assisted in preparation of coursework and taught discussion sections, improving the learning experience for 200+ undergraduate students

Service **Mentoring** 2025.9 – Present
Two undergraduate students, resulting in a publication at AIED 2026

Program Chair at AIED 2026

Honors & Awards Bay State Fellowship (Full Scholarship Sponsored by UMass CICS) 2022
Graduated Magna Cum Laude In MS from University of Massachusetts Amherst
Graduated Magna Cum Laude In BS from University of Massachusetts Amherst
Dean's List in all semesters

Skills **Machine Learning:** deep neural network architecture design, reinforcement learning algorithm implementation and reward/loss function design, preference alignment
Large Language Models: prompt engineering and RAG, LLM-driven annotation and evaluation, fine-tuning and LoRA using SFT, DPO, GRPO
Software Engineering: full-stack web development
Programming Languages: Python, Java, JavaScript, HTML, CSS, SQL
Frameworks: PyTorch, HuggingFace, TRL, vLLM, React, Node.js