Keyu He

(213) 713-2973 | keyuhe@cmu.edu | keyu-he.github.io | www.linkedin.com/in/keyuhe/

Education

Carnegie Mellon University, Pittsburgh, PA **May 2027**
Master of Science in Intelligent Information Systems @ School of Computer Science **GPA: 4.08/4.33**

University of Southern California, Los Angeles, CA **May 2025**
Bachelor of Science in Computer Science & Bachelor of Art in Applied Mathematics **GPA: 3.98/4.00**
Minor in Artificial Intelligence Applications

Skills

Programming: C++, C, Python, Java, MySQL, HTML, CSS, JS, x86-64 Assembly
Frameworks/Tools/Software: PyTorch, Pandas, NumPy, LangChain, FAISS, Git, AWS, LaTeX
Areas of Expertise: Machine Learning, Natural Language Processing (NLP), Large Language Models (LLMs), Vision Language Models (VLMs), Data Science / Data Engineering

Publications

1. **Keyu He**, Tejas Srinivasan, Brihi Joshi, Xiang Ren, Jesse Thomason, Swabha Swayamdipta. *Believing without Seeing: Quality Scores for Contextualizing Vision-Language Model Explanations*. Submitted to ACL 2026. Under Review. 10.48550/arXiv.2509.25844
   [VLM metrics design; calibration; human-subject studies; statistical analyses]
2. Brihi Joshi*, **Keyu He***, Sahana Ramnath, Sadra Sabouri, Kaitlyn Zhou, Souti Chattopadhyay, Swabha Swayamdipta, Xiang Ren. *ELI-Why: Evaluating the Pedagogical Utility of LLM Explanations*. Findings of ACL 2025. 10.48550/arXiv.2506.14200
   [Benchmark/Dataset design; human-subject studies; web-app implementation; multi-metric evaluation of LLM explanations]
3. Huihan Li*, Arnav Goel*, **Keyu He**, and Xiang Ren. *Attributing Culture-Conditioned Generations to Pretraining Corpora*. ICLR 2025. 10.48550/arXiv.2412.20760.
   [large-scale corpus analysis; Python data engineering and visualization; survey server infrastructure design]

Experience

CMind, Remote May 2025 – Aug. 2025
Data Engineer Intern
- Built a LangChain-based RAG prototype over an internal analytics DB (chunking, embeddings, FAISS indexing/retrieval, LLM answer synthesis) with metadata filters and notebook/API utilities.
- Designed CEO-speech trait metrics and an auto-evaluation pipeline with custom rubrics, achieving Fleiss' κ of 0.89–0.91 against human annotations.

University of Southern California, Los Angeles, CA Sep. 2022 – May 2025
Teaching & Grading Assistant
- Coordinated logistics and grading for two CS courses (CSCI-102, CSCI-360), serving ∼300 students per term.
- Led weekly OHs/discussions for ∼20 students, clarifying concepts and providing feedback on assignments.

Projects

AI-Based Career Advisor — University of Southern California Nov. 2024 – Dec. 2024
- Built an AI career advisor matching user skills/interests to roles via sentence-embedding similarity over 1.3M+ job-skill entries (JobSkills, LinkedIn Jobs).
- Integrated Bing AI to surface resources and job application links, plus consistency checks on recommendations.
- Evaluated recommendations with a T5-based entailment verification model to filter low-relevance matches.

LLM Prompt Recovery Project — University of Southern California Mar. 2024 – Apr. 2024
- Built a system to reconstruct user prompts from original text and Gemma-modified outputs.
- Fine-tuned Mixtral-8x7B with custom similarity metrics (sentence-T5-base + sharpened cosine), achieving a 0.65 leaderboard score.
- Earned a Kaggle silver medal (ranked 75/2175, **top 3.4%**) and released the final model on Kaggle.