Curriculum Vitae

Jinhu Qi

Email: jinhuqi@usc.edu / qijinhu1218@gmail.com

Address: 138 Trumpet Flower, Irvine, CA, USA

Nationality: China

Contact mobile phone: China - +86-13910807337, US - +1-4582101212

Research Interest and Plans

For my PhD research, I focus on utilizing Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) techniques, combined with knowledge graphs and vector databases, to overcome communication barriers across different languages and cultural backgrounds. By integrating language models with knowledge graphs and vector databases, I aim to perform reasoning in a seemingly "white box" environment, thereby addressing the hallucination problem of LLMs. By reducing communication barriers caused by language and cultural differences, my research aims to promote effective communication in social media and everyday life, thereby enhancing societal communication efficiency and understanding, and fostering social harmony and progress.

Education

- **University of Southern California** — Los Angeles, CA
  Master of Science in Analytics — *Jan 2022 - Dec 2023*

- **University of Oregon** — Eugene, OR
  Bachelor of Science in Computer and Information Science — *Sep 2017 - June 2021*
  Minor in Korean

Work Experience

- **Research Assistant,** YiSa Lab**,** School of Computer Science and Software Engineering, Sichuan University Jin Cheng College
  Advisor: Ke Wang — *May 2024 - Present*

  Research and Application of Intelligent Tourism Service Systemin Tibet Based on LLM：

  - Developed a large-scale project in the lab focused on smart cultural tourism in Tibet, leveraging Large Language Models (LLMs).
  - Conducted research on fine-tuning LLMs, using Vector-RAG and Graph-RAG techniques to optimize hallucination issues in Tibet tourism recommendation systems.
  - Published two papers accepted by ICWOC 2024 and AIPR 2024:
    - Paper 1: *Research on Tibetan Tourism Viewpoints information generation system based on LLM*
    - Paper 2: *RAG-Optimized Tibetan Tourism LLMs: Enhancing Accuracy and Personalization*
  - Currently writing a third paper on integrating Graph-RAG with LLMs to efficiently solve the hallucination problem in tourism recommender systems.

  Research on Tibetan Tourism Viewpoints information generation system based on LLM：

  - Tackled the hallucination problem in LLMs when providing recommendations for unfamiliar tourist viewpoints.
  - Collected and fine-tuned LLMs with data on Tibetan tourist viewpoints, improving hotel information extraction accuracy from **0.47** to **0.98**.
  - Achieved an average accuracy score of **80** through comparative experiments with fine-tuning methods like SFT and ORPO.
  - Fine-tuned models significantly reduced hallucinations, with results published in ICWOC 2024.

  RAG-Optimized Tibetan Tourism LLMs: Enhancing Accuracy and Personalization

  - Explored LLM hallucinations in tourism recommendation systems by applying Vector-RAG techniques.

- Collected data from the web, converted it into vectors using a pre-trained BERT model and a TF-IDF model and stored them in an FAISS vector database.
- Achieved **65%** accuracy in retrieving the top three most relevant results by testing various vector retrieval methods.
- Integrated retrieved results into the LLM, improving the Llama 3 model's quality score from **0.7429** to **0.9225**, with results published in AIPR 2024.

- **Research Assistant,** USC HUMANS Lab, Thomas Lord Department of Computer Science and Information Sciences Institute, Viterbi School of Engineering, University of Southern California

  Advisor: Emilio Ferrara *Aug 2023 - Present*

  The 2024 Election Integrity Initiative：
  - Developed and optimized Python scripts to collect extensive YouTube video and 4chan text data related to the 2024 U.S. election.
  - Participated in monitoring social media platforms and fringe communities to track election-related content and potential misinformation.
  - Executed comprehensive data cleaning processes to ensure high-quality datasets. Analyzed collected data and generated detailed reports to support research findings.
  - Collaborated with the research team to provide insights and support for the Election Integrity Initiative, contributing to the observatory's mission to uphold election integrity.

  Detection of Misinformation on Social Media:
  - Worked with Eun Cheol Choi, Ongoing Communication PhD, on a project to detect misinformation on social media.
  - Designed and implemented the code and logical framework for fine-tuning a large language model using LLaMA-Factory, tailoring it for effective misinformation detection.
  - Conducted experiments and iterative improvements to achieve a **75%** accuracy rate in identifying misinformation, enhancing the model's robustness.
  - Evaluated the model's performance through comprehensive testing and analysis. Developed reports detailing the model's accuracy and reliability in detecting false information.

- **Data Scientist**, Machine Learning Intern, DerbySoft Ltd

  Advisor: Chao Yang *June 2023 - Dec 2023*

  Based on ChatGPT to customize and recommend hotels that are most relevant to customers:
  - Developed ChatGPT plugins, showcasing proficiency in NLP and demonstrating the potential in LLMs by optimizing OpenAI's specialized workflows, enhancing customer service query efficiency by **50%.**
  - Executed advanced predictive modeling and utilized A/B testing to refine ad targeting and customer engagement strategies, mirroring the methodical approach needed for empirical research in AI.
  - Enhanced the company's recommendation algorithms by integrating vector databases with existing APIs, using Python to achieve an **80%** accuracy in hotel recommendations, showcasing the ability to apply machine learning techniques in real-world applications.

- **Teaching Assistant,** Industrial and Systems Engineering, Viterbi School of Engineering, University of Southern California

  Advisor: Carl F Kesselman & Bruce Wilcox *Aug 2022 - May 2023*

  DSCI/ISE 559-Introduction to Data Management & ISE535-Data Mining:
  - Facilitated coursework in Data Management and Data Mining, directly relevant to managing and analyzing large datasets, a fundamental skill in LLM research.
  - Led tutorial sessions in R programming and Exploratory Data Analysis (EDA), enriching students' analytical
  - capabilities and mirroring the deep technical understanding required for LLM research.
  - Managed TA teams to ensure efficient task completion, demonstrating leadership and collaborative skills vital for multidisciplinary AI research projects.

- **Software Engineer Intern,** Hotel Development Department, Tongcheng International Travel Service Co. Ltd.
  Advisor: Frank Wang *Jul 2018 - Aug 2018*

  Research and develop keyword hotel search through left and right entropy algorithm:
  - Modified the program to optimize the hotel search system, extracted and analyzed the left and right entropy of keywords to improve the accuracy of predicting customer input keywords by **20%** using Python and SQL.
  - Increased the user's loyalty to the hotel's booking website by **15%** and assembled the data visualization to provide support to the stakeholders to make decisions.

Publication

Jinhu Qi, Shuai Yan, Wentao Zhang, Yibo Zhang, Zirui Liu, Ke Wang, (I am the First Author), "*Research on Tibetan Tourism Viewpoints information generation system based on LLM*" Paper accepted by 2024 12th International Conference on Intelligent Computing and Wireless Optical Communications (ICWOC 2024) June 21-23, 2024. *The patent is in an application, forthcoming in 2025.* ***(https://arxiv.org/abs/2407.13561)***

Jinhu Qi, Shuai Yan, Yibo Zhang, Wentao Zhang, Rong Jin, Yuwei Hu, Ke wang, (I am the First Author), "*RAG-Optimized Tibetan Tourism LLMs: Enhancing Accuracy and Personalization*" Paper accepted by 2024 7th International Conference on Artificial Intelligence and Pattern Recognition (AIPR 2024) September 20-22, 2024. *The patent is in an application, forthcoming in 2025.(**https://arxiv.org/abs/2408.12003**)*

Research Project

- **My ongoing project that will apply for a patent, forthcoming 2025:**

  Enhancing Cross-Cultural Communication on Social Media through Offensive Content Detection with LLMs: *May 2024 - Present*
  - Develop a system to detect and categorize culturally offensive content on social media, focusing on differences between Chinese and American cultural sensitivities.
  - Utilize the BERT model to analyze and classify social media content (e.g., tweets, videos) into specific cultural offense categories. Implements tagging of content and stores results in a vector database for further analysis.
  - Use a vector database to identify if new content matches existing cultural offense categories.
  - Fine-tune a large language model (LLM) by SFT/ORPO/KTO with data from BERT classifications to classify new content.
  - Compare the performance of methods using evaluation metrics such as BERTScore and accuracy.
  - Aims to create a reliable system for identifying and mitigating culturally offensive content, enhancing cross-cultural understanding and communication on social media.

  Graph-RAG: Integrating Knowledge Graphs for Improved Interpretability: *July 2024 - Present*
  - Continued research on combining knowledge graphs with LLMs to transform "black box" logic into "white box" interpretability. Generated knowledge graph triples for Tibetan tourist attractions using Microsoft's prompt engineering techniques.
  - Developed a function-calling LLM mechanism, improving triple keyword extraction accuracy by **7%** compared to fine-tuned LLMs.
  - Designed two knowledge graph retrieval engines: logical chain retrieval and multi-node brute-force retrieval.
  - Conducted a user survey where Graph-RAG LLM combined with logical chain retrieval achieved a **60%** higher user preference than baseline LLMs. Improved LLM transparency, accuracy, and relevance, with plans to refine knowledge graph methods further.

- **Researcher for the academic article written and prepared by Eun Cheol Choi, "Dissecting Deceptive Discourse by Using Artificial Social Media Feeds via LLMs" :** *Jan 2024 - Present*
  - Pre-trained the Mistral-7B-v0.2 model on Hugging Face to generate naturalistic human-like social media posts and tweets towards our goal of creating content with misinformation and toxicity.
  - Collected and labeled data by identifying accurate, neutral, and misinformation or deceptive information from social media platforms, such as X and Instagram, to provide training datasets.

- Enhanced the accuracy of identifying misinformation and toxic elements by pre-training and fine-tuning the model via SFT/ORPO and evaluated the accuracy performance through recall and F1 score.
- Suggested how generative LLMs can simulate and manipulate social media feeds under varying network structures, thereby impacting the susceptibility to misinformation.
- Applied our LLMs for generating realistic social media content to prepare for high-external-validity Randomized Controlled Trials (RCTs), aiming at dissecting the role of network structures in misinformation spread.

- **Researcher for the academic article written and prepared by Eun Cheol Choi, "Fact-Checking Augmentation via Claim Matching with LLMs"：** *Aug 2023 - Dec 2023*
  - Processed and prepared datasets for LLM fine-tuning, involving extensive preprocessing of tweets posted during the COVID-19 pandemic on Twitter (X), focusing on claims related to the pandemic.
  - Developed and implemented a classification system to categorize matched tweets and claims into coherent groups, distinguishing between identical, different, and ambiguous claims to support accurate information dissemination.
  - Fine-tuned and evaluated LLaMA 2-7b as the optimal LLM for the enhancing fact-checking mechanisms due to its superior predictive accuracy in classifying matched claims. Achieved **75%** accuracy rate in claim matching in fine-tuned LLaMA 2.

Teaching Performance

- **University of Southern California, Viterbi School of Engineering** Los Angeles, CA

Course Name	Year	Semester	Weeks	Hours per week	Count of Students	Instructor
ISE 535 Data Mining	2022	Fall	15	20	114	Bruce Wilcox
ISE/DSCI 559 Introduction to Data Management	2023	Spring	15	20	138	Bruce Wilcox / Carl Kesselman

Qualitative comments from students:

1. Anonymous Student from ISE 535: “Jinhu Qi consistently demonstrated exceptional teaching skills by explaining complex concepts clearly and making them accessible to me. His ability to simplify difficult topics was invaluable to the class.”
2. Anonymous Student from ISE/DSCI 559: “Jinhu Qi showed outstanding dedication in supporting students. He was approachable and responsive, always willing to go the extra mile to ensure me understood the material.”
3. Anonymous Student from ISE/DSCI 559: “Jinhu Qi made a significant contribution to the course by developing useful supplemental materials and organizing effective review sessions, greatly enhancing the overall learning experience.”

Honors and Award

- **Phi Beta Kappa Honor Society, University of Oregon** *2022*

- **Dean's List, University of Oregon** *Winter 2020, Fall 2020, Winter 2021*

Technical Skills

- **Programming Languages:** Python, LLMs Pre-training and Fine-tuning by SFT/ORPO/KTO, Vector Database, Machine Learning, Tableau, SQL, R, API Development, RAG, Graph RAG, LangChain, Neo4j with Knowledge Graph
- **Certification:** Datacamp - Data Analyst with Python, SQL, and Tableau

Referees

- **Emilio Ferrara** USC Annenberg School of Communication
  Professor of Computer Science and Communication
  emiliofe@usc.edu

- **Carl F Kesselman** USC Viterbi College of Engineering
  Professor of Industrial and Systems Engineering, Computer Science, Population and Public Health Sciences, and Biomedical Sciences
  carl@isi.edu

- **Ke Wang** Sichuan University Jin Cheng College
  Head of the Department of Big Data; Vice Dean of the Institute of Artificial Intelligence and Large Model Industry
  wangke@cdjcc.edu.cn