Kuan-Chieh Lo

(+1)310-994-7190 | lo.311@osu.edu | linkedin/kuan-chieh-lo | kclo0908.github.io | Columbus, OH

Education

The Ohio State University — Columbus, Ohio
Ph.D. in Computer Science and Engineering (Advisor: Srinivasan Parthasarathy) — Aug 2022 – Aug 2027 (Expected)

National Taipei University of Technology — Taiwan
M.S. in Electrical Engineering — Sep 2012 – Aug 2014

National Changhua University of Education — Taiwan
B.S. in Mechatronics Engineering — Sep 2008 – Jun 2012

Technical Skills

LLM & Agentic AI: LangChain, LangGraph, RAG, Multi-agent Orchestration, MCP, Agent Skills
Machine Learning / Data Science: PyTorch, TensorFlow, Scikit-learn, NLTK, Pandas, NumPy, Matplotlib
Data & Vector Databases: ChromaDB, FAISS, MongoDB, MySQL
Programming Languages: Python, Java, JavaScript, C/C++
Web Development & Tools: Flask, Django, Node.js, Vue, Docker, Git, Nginx

Research Experience

The Ohio State University — Columbus, Ohio
Graduate Student Researcher — Aug 2022 – Present

- **Multi-Agent Adversarial Claim Verification** — Proposed a multi-agent LLM framework organizing heterogeneous agents across multiple foundation models, incorporating claim decomposition, multi-hop knowledge retrieval, and adversarial verification. Achieved significant accuracy gains over single-agent and homogeneous multi-agent baselines on fact-verification benchmarks. (*ICWSM 2026*)
- **Agentic AI for Crisis Response** — Developed Crisis Observatory, a multi-agent LLM system for crisis response that extracts credible signals from social media by integrating topic modeling, geolocation extraction, and a RAG pipeline grounded in authoritative external knowledge bases for credibility assessment and misinformation filtering. (*ICDM 2025*)
- **Fairness in Federated Graph Learning** — Proposed FairWAG, a fairness-aware federated learning framework that applies Shapley Values to quantify client contributions, enabling adaptive aggregation weights that optimize the performance-fairness trade-off under demographic skew across clients. (*EAAMO 2025*)
- **LLM Safety: Jailbreak Attacks on Large Reasoning Models** — Uncovered a prompt injection vulnerability in Large Reasoning Models (LRMs), where adversarially crafted prompts inject a spoofed chain-of-thought block to bypass safety alignment; evaluated across six open-source LRMs on adversarial benchmarks, achieving high attack success rates against strong baselines. (*Under Review*)
- **Compositionality Evaluation of VLMs** — Assessed compositional reasoning capabilities of state-of-the-art vision-language models (VLMs), exposing critical skill gaps across object detection, relational extraction, and attribute binding tasks.

Academia Sinica — Taipei, Taiwan
NLP Research Scientist (Mentor: Dr. Lun-Wei Ku) — Jan 2019 – Aug 2022

- **Misinformation Mitigation via Recommendation Systems** — Developed VICTOR, a reinforcement learning-based module that implicitly re-ranks news recommendations to surface verified articles, reducing users' exposure to misinformation without explicit content removal. (*WWW 2022, WSDM 2021*)
- **Echo Chamber Reduction** — Built a news-analysis platform that applies a stance classification model to present multi-source perspectives on events, visualizing stance distributions and political ideology of news sources, reducing filter bubble effects. (*WWW 2021*)
- **Visual Storytelling Evaluation** — Developed Vrank, a reference-free automatic evaluation metric for visual storytelling (VIST) trained on the VHED human evaluation dataset, achieving ˜30% higher accuracy than existing metrics when ranking story pairs and demonstrating generalizability to pure textual stories. (*ACL 2022*)

Work Experience

ASUSTeK Computer Inc. Taiwan

Software and Firmware Engineer 2014 – 2018

- Developed system firmware for ASUS products, including gaming desktops, business and consumer laptops.

Projects

Creativity in Large Language and Vision-Language Models

- Investigated the compositional factors influencing creative output in LLMs and vision-language models (VLMs), analyzing how individual components (e.g., concept, style, structure) interact to drive creativity and exploring controllability of creative generation across modalities.

LLM-Powered Agentic Chatbot for Bridal Consultation

- Designed and developed a multi-agent AI system for a bridal company integrating a conversational LLM chatbot, a personalized recommender system agent, and a meeting scheduling agent into a unified pipeline, enabling end-to-end automated client consultation and appointment booking.

LLM-Assisted Legal Document Processing System for Law Firms

- Built an NLP pipeline for law firms leveraging LLMs and retrieval-augmented generation (RAG) to automatically retrieve semantically relevant past legal cases from a document corpus and populate structured legal documents, reducing manual effort and accelerating attorney workflows.

Selected Peer-Reviewed Publications

- **Lo, K.-C.**, Shalin, V., Garrett, K., Parthasarathy, S. *Claim Verification with Adversarial Reasoning and Planning.* International AAAI Conference on Web and Social Media (ICWSM), May 2026
- **Lo, K.-C.**, He, Y., Jiang, Y., Parthasarathy, S. *FairWAG: Fairness-aware Weighted Aggregation for Graph Learning in a Federated Setting.* ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), November 2025
- **Lo, K.-C.**, Maneriker, P., Ganesh, S., et al. *Crisis Observatory: Extracting Credible Signals During a Crisis in the Age of LLMs.* IEEE International Conference on Data Mining (ICDM), November 2025
- **Lo, K.-C.**, Dai, S.-C., Xiong, A., Jiang, J., & Ku, L.-W. *VICTOR: An Implicit Approach to Mitigate Misinformation via Continuous Verification Reading.* ACM Web Conference (WWW), April 2022
- **Lo, K.-C.**, Dai, S.-C., Xiong, A., Jiang, J., & Ku, L.-W. *Escape from An Echo Chamber.* ACM Web Conference (WWW), June 2021
- **Lo, K.-C.**, Dai, S.-C., Xiong, A., Jiang, J., & Ku, L.-W. *All the Wiser: Fake News Intervention Using User Reading Preferences.* ACM International Conference on Web Search and Data Mining (WSDM), March 2021
- Hsu, C.-Y., Chu Y.-W., Chen V., **Lo, K.-C.**, Chen C., Huang T.-H., & Ku, L.-W. *Learning to Rank Visual Stories From Human Ranking Data.* Association for Computational Linguistics (ACL), May 2022

Under review:

- **Lo, K.-C.**, et al. *Jailbreak Vulnerabilities in Large Reasoning Models via Chain-of-Thought Injection.*

Academic Services and Teaching Experience

Conference Reviewer

- ICML (2026), ICWSM (2026), NeurIPS (2024), ACL (2024, 2022), AAAI (2022), IAAI (2022, 2021, 2020), NLPCC (2021), EMNLP (2020).

Graduate Teaching Associate (The Ohio State University)

- Database Systems (Autumn 2023, Spring 2024)
- Data Mining (Autumn 2024, Autumn 2025, Spring 2026)
- Network Science (Spring 2025)