

Jiseung Hong

jiseungh@andrew.cmu.edu | (412) 638-3100 | [Website](#) | [LinkedIn](#)

Education

Carnegie Mellon University

Pittsburgh, PA

Master of Science in Intelligent Information Systems (MIIS) | GPA: 3.92/4.00

Dec 2025

Korea Advanced Institute of Science and Technology (KAIST)

Daejeon, South Korea

Bachelor's Degree – School of Computing

Feb 2024

Korea Student Aid Foundation (KOSAF) National Scholarship

Mar 2017 ~ Feb 2024

Academic Research and Projects

OpenHands PR-Arena: Platform for Evaluating and Benchmarking Agentic Coding Assistants [Website](#) | [GitHub](#)

Carnegie Mellon University

Aug 2024 ~ Dec 2025

- Designed and developed PR-Arena, a GitHub App that evaluates automated GitHub issue resolution by generating paired pull requests from state-of-the-art LLMs, enabling leaderboard-based comparison of human preferences across models
- Publicly released the leaderboard and performed comprehensive research analyzing discrepancies between the human preference data (i.e., PR-Arena evaluation) and the benchmark evaluation metrics (i.e., SWE-bench evaluation)
- Conducted extensive inference testing and performance evaluation of LLM coding agents on SWE-bench to assess issue resolution capabilities [Blog Post](#)

Tinker Tales: Interactive Storytelling Framework for Early Childhood Narrative Development and AI Literacy

Emory University

June 2025 ~ Dec 2025

- Built a conversational agent for children that scaffolds and evaluates narrative development and human-AI collaboration
- Developed conversation functions, backend database (NoSQL), and educator-facing web app on AWS

Measuring Sycophancy of Language Models in Multi-turn Dialogues [EMNLP 2025 Findings] [arXiv](#) | [GitHub](#)

Carnegie Mellon University & Emory University

Feb 2025 ~ May 2025

- Introduced *SYCON Bench*, a novel benchmark for evaluating and quantifying sycophantic behavior in real-world conversational settings (multi-turn, free-form, and open-ended)
- Evaluated 17 Large Language Models (LLMs) on *SYCON Bench*, analyzing how model scaling, reasoning-optimization, and alignment tuning influence sycophancy
- Experimented 5 mitigation strategies, demonstrating that third-person perspective prompting significantly reduces sycophancy

Measuring Social Biases in State-Spaces Models

Carnegie Mellon University

Feb 2025 ~ May 2025

- Evaluated comparable transformer-based and state-space model (SSM) on StereoSet and BBQ dataset; Showed that SSMs exhibit reduced bias and better trade-off between LM performance and stereotypes
- Explored mitigation strategies such as fine-tuning transformer-based model and SSMs on an augmented data

Korean Bio-Medical Corpus (KBMC) for Medical Named Entity Recognition [LREC-COLING 2024] [arXiv](#)

Seoul National University

May 2023 ~ Oct 2023

- Constructed the *Korean Bio-Medical Corpus (KBMC)*, the first open-source dataset for Korean medical named entity recognition
- Contributed primarily to evaluating six different language models on the KBMC dataset

Teaching Experience

Carnegie Mellon University

Pittsburgh, PA

Teaching Assistant (TA) - IITP Executive Education: Introductory NLP (Prof. David Mortensen)

Aug 2025 ~ Dec 2025

Work Experience

NCSOFT

Seongnam, South Korea

Intern, Natural Language Processing Narrative Laboratory

Mar 2021 ~ Aug 2021

- Developed a software prototype to organize clusters of news articles based on relevance to a given query by applying the *Learning-To-Rank (LTR)* technique to a Support Vector Machine (SVM)
- Outperformed the previous news ranking system in terms of precision, recall, and f1-score

Skills

DevOps & Tools – GitHub Actions (CI/CD), Docker, YAML, Firebase, Chroma DB, AWS