Claire Kim

nkim48@asu.edu | https://clairekim59.github.io | https://www.linkedin.com/in/clairekim59

PhD candidate with AI/ML engineer experience, specializing in LLM inference systems, RAG, evaluation, and scalable production pipelines. Brings applied industry experience from Amazon, AMD, and federally funded research in trustworthy AI decision support.

EDUCATION

ARIZONA STATE UNIVERSITY **2021 – Aug 2026**
PhD, Computer Science Tempe, AZ
- Data Mining & Machine Learning Lab
- Advisors: Dr. Huan Liu and Dr. Mickey Mancenido
- Funded by DHS-CAOE

KOREA UNIVERSITY **2013 – 2019**
M.E., Computer Science & Engineering Seoul, South Korea
B.S., Computer Science & Engineering

WORK EXPERIENCE

AMAZON — Sep – Dec 2025
Applied Scientist Intern — **Bellevue, WA**
- Improved production sentiment forecasting system for community operations using advanced ML models, increasing prediction accuracy and robustness at scale
- Designed and implemented an end-to-end pipeline from data collection to production-ready Python package deployment leveraging AWS and MLflow for experiment tracking and CI/CD-driven model releases

AMD — Aug – Dec 2024, May – Aug 2025
AI/ML Intern — **Austin, TX**
- Reduced hallucination in LLM-based systems using retrieval-augmented generation and feedback-driven solution
- Designed and implemented LLM evaluation and inference pipelines for scalable model deployment
- Built synthetic QA generation pipeline and LLM-as-a-judge system enabling automated evaluation at scale
- Improved multi-agent LLM system by incorporating user feedback signals to increase production reliability
- Optimized data and inference pipelines for efficiency, improving throughput and reducing evaluation latency

DHS-CAOE (Center for Accelerating Operational Efficiency) — May 2022 – Aug 2024
Graduate Research Assistant — **Tempe, AZ**
- Designed NLP models for topic modeling and text summarization using BERT and Llama-2
- Collaborated with an interdisciplinary team to build a multi-agent RAG system for trustworthy AI-enabled intelligence analysis, validated via human-subject study using GPT-4
- Developed an interactive data analysis and visualization dashboard using NodeJS and Flask

ONR (Office of Naval Research Project) — Jan 2021 – Aug 2022
Graduate Research Assistant — **Tempe, AZ**
- Researched integration of online/offline COVID-19 datasets using topic modeling
- Analyzed 2M tweets for sentiment and stance detection in pandemic-related discussions

SELECTED PROJECTS

Reasoning-Level Fairness in LLMs (Ongoing)

- Investigating bias propagation in multi-step LLM inference and developing evaluation methods for fairness in reasoning processes
- Implementing scalable evaluation pipelines for analyzing fairness across large-scale generated outputs

Bayesian Learning based Uncertainty Measurement for Hallucination Mitigation (Ongoing)

- Design uncertainty-aware LLM system and developed evaluation workflows to demonstrate hallucination detection in practice

TECHNICAL SKILLS

- **Languages:** Python, Java, SQL, JavaScript, Bash
- **ML & DL Frameworks:** PyTorch, TensorFlow, scikit-learn, Hugging Face, Pandas, NumPy, LangChain
- **LLM & AI Systems:** RAG, vector DBs (Weaviate, Chroma, Faiss), LLM-as-a-judge, multi-agent LLM, hallucination, fine-tuning (SFT, RLHF, LoRA/PEFT), prompt engineering, inference-time scaling, responsible AI · Models: GPT-4o, Claude, Llama-3, Mistral, BERT
- **MLOps & Deployment:** CI/CD for ML, model registries, Docker, model serving (cloud & edge), inference optimization (quantization, precision tuning), PyTorch Profiler, TensorRT, MLflow, WandB
- **AI-assisted dev:** Claude Code, Cursor, GitHub Copilot, agentic coding workflows, MCP
- **Cloud & Infrastructure:** AWS, GCP, Docker, Node.js, Flask, Git, Linux
- **NLP & Evaluation:** text classification, forecasting, summarization, bias mitigation, benchmarking, evaluation metrics, data preprocessing & feature engineering

SELECTED PUBLICATIONS | Google Scholar

- M. Cohen, **N. Kim** et al., PADTHAI-MM: A Principled Approach for the Design of Trustworthy, Human-Centered AI systems using the MAST Methodology **AI Magazine'25**
- **N. Kim** et al., Robust Stance Detection: Understanding Public Perceptions in Social Media **ASONAM'24**
- P. Salehi, Y Ba, **N. Kim** et al., Evaluating Trustworthiness of AI-Enabled Decision Support Systems: Validation of the Multisource AI Scorecard Table (MAST) **JAIR'23**
- L. Cheng, **N. Kim**, H. Liu, Debiasing Word Embeddings with Nonlinear Geometry **COLING'22**
- **N. Kim**, D. Mosallanezhad, L. Cheng, B. Li, and H. Liu, Bridge the Gap: the Commonality and Differences Between Online and Offline COVID-19 Data **SBP-BRiMS'22**

ACADEMIC SERVICE & TEACHING

- Reviewer at AMLC 2025 (Amazon Gen AI Evaluation Workshop) 2025
- Invited Reviewer for NeurIPS 2025 (Reliable ML Workshop) 2025
- Program Committee (PC) member of ASONAM 2024 conference 2024
- Invited talk at Machine Learning Day 2024, Arizona State University 2024
- Invited talk at SCAI AI Day, Arizona State University 2023
- Program Committee (PC) member of ASONAM, SBP-BRiMS 2023 conferences 2023
- Invited Reviewer for EMNLP 2023 conference 2023
- Reviewer at ECML-PKDD, ACM MultiMedia, ASONAM, and AAAI conferences 2022
- Volunteer at WSDM 2022 conference 2022
- Graduate Teaching Assistant, Object-Oriented Programming & Data Structures (ASU) 2022