

Ken/Yuangang Li



yuanganl@uci.edu | (+1) 213-421-9602 | yuangang.me | Irvine, CA

SUMMARY

Yuangang Li is a Ph.D. student at UCI with research and engineering capabilities in AI, specializing in **trustworthy LLMs**, **efficient ML systems**, and **AI for Software Engineering**. He has conducted research at Stanford, USC, and UVA, and previously worked at SenseTime developing production-grade **cloud-native infrastructure**. He brings a combination of deep research expertise and hands-on production experience. He also serves on the program committee for top-tier conferences, including KDD 2025 (Outstanding Reviewer) and 2026 (Senior Area Chair), ICLR 2026, ICML 2025, AAAI 2026, MLSys 2026, IEEE MLSP 2025, ICASSP 2025, and WWW 2026.

PUBLICATIONS

- [1] *Mitigating Hallucinations in Large Language Models via Causal Reasoning* ([AAAI 2026](#)). Authors: **Yuangang Li**, et.al.
- [2] *S2D-Align: Shallow-to-Deep Auxiliary Learning for Anatomically-Grounded Radiology Report Generation* ([AAAI 2026](#)). Authors: Jiechao Gao, Chang Liu, **Yuangang Li**
- [3] *LLM-Guided Semantic Bootstrapping for Interpretable Text Classification with Tsetlin Machines* (Submit to [ACL 2026](#)). Authors: Jiechao Gao, Rohan Kumar Yadav, **Yuangang Li**, Jie Wang
- [4] *NLP-ADBench: NLP Anomaly Detection Benchmark* ([EMNLP 2025](#)). Authors: **Yuangang Li**, et.al.
- [5] *AD-LLM: Benchmarking Large Language Models for Anomaly Detection* ([ACL 2025](#)). Authors: Tiankai Yang, Yi Nian, Shawn Li, Ruiyao Xu, **Yuangang Li**, et.al.
- [6] *A Large-Scale Simulation on Large Language Models for Decision-Making in Political Science* ([arxiv](#)). Authors: Chenxiao Yu, Jinyi Ye, **Yuangang Li**, Zheng Li, et.al.
- [7] *H-FedSN: Personalized Sparse Networks for Efficient and Accurate Hierarchical Federated Learning for IoT Application* ([arxiv](#)). Authors: **Yuangang Li***, et.al.
- [8] *FedMetaMed: Federated Meta-Learning for Personalized Medication in Distributed Healthcare Systems* ([BIBM 2024](#)). Authors: Jiechao Gao, **Yuangang Li**
- [9] *FedLDR: Federated Local Data-infused Graph Creation with Node-centric Model Refinement* ([ICDM 2024](#)). Authors: Jiechao Gao, **Yuangang Li**, et.al.
- [10] *Artificial Intelligence-Aided Digital Twin Design: A Systematic Review* ([Preprints 2024](#)). Authors: Nan Hao*, **Yuangang Li***, et.al.
- [11] *FedBCGD: Communication-Efficient Accelerated Block Coordinate Gradient Descent for Federated Learning* ([ACM MM 2024](#)). Authors: Junkang Liu, Fanhua Shang, Yuanyuan Liu, Hongying Liu, **Yuangang Li**, et.al.

RESEARCH EXPERIENCES

University of California, Irvine | Research Assistant

09.2025-Present

Research in Trustworthy AI and Software Engineering

- Evaluation of LLMs reasoning quality, focusing on coherence and causal consistency in step-by-step reasoning
- Explainability for code generation, to ensure the safety of generated code by interpreting the model's internal mechanisms

University of Southern California | Research Assistant

07.2023-08.2025

Research in Trustworthy AI

- Proposed CDCR-SFT, a novel **fine-tuning method** that internalizes causal DAG-based reasoning, boosting **causal reasoning to 95.3% (from 72.9%)** on the CLADDER benchmark, and **reducing hallucinations by 11%** on HaluEval^[1]
- Proposed NLP-ADBench, the **first comprehensive NLP anomaly detection benchmark**, establishing a standardized foundation by showing the strong effectiveness of LLM-based features across 19 methods and 8 datasets^[4]
- Developed AD-LLM to evaluate LLMs for NLP-AD, showing strong generalization in zero-shot detection, data augmentation, and model selection, and positioning LLMs as a robust backbone for future NLP-AD systems^[5]
- Designed a multi-step reasoning framework for LLMs to simulate voter decisions, cutting election prediction error by 77% and exposing bias and overfitting issues in models like GPT-4o and LLaMA 3.1^[6]

Stanford University | Research Assistant

05.2025-08.2025

Research in Trustworthy AI and Explainable AI

- Fine tune (LoRA, DPO, etc.) open-source LLMs** on the first self-constructed dataset with causal structures, improving causal reasoning and mitigating hallucinations, offering the first empirical evidence linking causality to hallucination^[1]
- Proposed S2D-ALIGN, a progressive fine-tuning framework that learns anatomically grounded alignment from coarse to fine using reference reports and anatomical key phrases, achieving SOTA results on MIMIC-CXR and IU X-RAY^[2]
- Developed an LLM-guided semantic bootstrapping method that distills LLM knowledge into interpretable Tsetlin Machines, achieving BERT-level accuracy without embeddings or inference, while preserving clause-level transparency^[3]

University of Virginia | Research Collaboration

12.2023-12.2024

Research in Effective Machine Learning Systems

- Designed an automated method to generate optimally compressed LLMs, focusing on enhancing model accuracy and deployment efficiency through innovative quantization and calibration techniques
- Proposed H-FedSN, a personalized sparse network method for hierarchical federated learning in IoT, reducing communication cost by up to 238× while preserving accuracy via structured masking and Bayesian aggregation^[7]
- Developed the FedMetaMed, integrating federated learning and meta-learning to enhance personalized medication

strategies across distributed healthcare systems, improving model adaptability and privacy preservation^[8]

- Single-handedly developed FedLDR, a brand new federated learning algorithm that employs GCN to enhance spatio-temporal data analysis through local data integration and node-centric optimization^[9]
- FedMetaMed^[8] and FedLDR^[9] were accepted by BIBM and ICDM respectively, and I gave talks on these research

Chinese Academy of Sciences | Research Assistant

05.2021-09.2021

Research in Software Engineering

- Contributed to the R&D of a web-based automated testing tool using Record and Playback technology, significantly enhancing test case management by enabling streamlined recording, editing, execution, analysis, and result generation. This implementation boosted end-to-end testing efficiency by 300% and saved over 15 hours per week
- Independently developed a script parser using Node.js that converts user actions recorded in JSON format into executable Robot Framework and Selenium scripts, enabling the replay and repeated execution of these user actions
- Single-handedly created innovative UI components using Vue.js and AceEditor, orchestrated the optimal containerization of the program with Docker, and automated the DevOps pipeline to maximize development efficiency

WORK EXPERIENCES

University of Southern California | Research Engineer

03.2023-07.2023

Developing a system for ML distributed execution

- Contributed to "[Ablator](#)", an open-source Deep Learning framework used by 40+ USC researchers for horizontal scaling of ablation experiments and hyperparameter tuning, encompassing 70 pull requests
- Implemented distributed experiment execution with **Ray**, managed open-source projects, set up **CI pipelines** via GitHub Actions, oversaw release management and version control, and authored **pytest** unit tests with **97%** coverage
- Solely launched 'python-rclone' on PyPI, a Python API for RClone that streamlines cloud data synchronization for 'Ablator', removing pre-installation requirements and enabling automatic binary selection ([python-rclone](#))

SenseTime | Infrastructure Engineer

01.2022-01.2023

Developing a SaaS Platform for AI and production-grade Application

- Built scalable, production-grade AI infrastructure with Kubernetes-native tooling, cloud-native messaging systems, and GPU-accelerated ML pipelines
- Developed "RocketMQ as a Service", akin to "RabbitMQ as a Service" in AWS Marketplace, offering fully managed SaaS-based RocketMQ clusters, increasing **300%** creation speed and saving **10+ hours/week** in manual operations
- Utilized Operator SDK to build a **Kubernetes**-based RocketMQ Operator and CRD automating lifecycle management
- Integrated **Prometheus & Grafana** for real-time monitoring, **cutting 35%** downtime and **boosting 40%** failure detection
- Optimized a machine learning training pipeline using **GPUs on Kubernetes** for enhanced computational efficiency
- Researched container runtimes (sysbox, crun, youki) to enhance **AI workload** isolation and resource management in K8S

Xiaoniu Translations (Beijing) Technology Co., Ltd. | FullStack Developer

01.2021-04.2021

- Developed an AI document translation system with Java/Spring/Maven, independently created a PDF/XML parsing module attracting 30,000 MAUs, used Nginx for reverse proxy, and managed version control with Git

PROJECT EXPERIENCES

SLinux OS (Similar to Linux 0.11) {Code} | C, File System, System Call Interface

06.2023-09.2023

- Developed a functional Unix-like Operating System, providing basic os functions and ensuring performance after tests
- Gained hands-on experience with kernel development, file systems, process scheduling, and other core OS concepts
- Developed modules for memory management, process scheduling, file system management, and process communication
- Implemented multi-threading and concurrency control to manage and execute multiple processes simultaneously

Game Development and Data Analysis | Unity, C#, Python, Firebase, Data Visualization, Game Design

01.2023-05.2023

- Contributed to 'Hexagon Adventure,' a 2D game, optimizing various aspects of game design and analytics
- Implemented automated data collection in Unity using C# to gather critical player behavior and game state data, and stored it in Firebase Realtime Database to support decision-making
- Analyzed complex game data using Python, focusing on key metrics like player progression, health point loss, and power-up usage, contributing to a 40% increase in player engagement

Job Posting Analytics System {Demo} | ETL, Data Mining, Database, Visualization, SpringBoot, Docker

05.2019-06.2020

- Developed a text mining system for collecting and analyzing recruitment data, which won the "National Level Innovative Excellence Project" award and was adopted by the school
- Extracted over **10 million+** recruitment records from job sites using **web scraping**, with Bloom Filters for deduplication
- Automated **ETL** processes using Python (numpy, pandas), and stored data in MongoDB and MySQL
- Created automated reports, dashboards, and data visualizations with Spring Boot, ReactJS, NodeJS, and AntV, enhancing data readability and visual impact, and integrated Docker with CI/CD pipelines for streamlined project automation

EDUCATION

University of California, Irvine | Ph.D. in Software Engineering

09.2025 - 06.2029

University of Southern California | M.S. in Computer Science (GPA: 3.81/4.0)

01.2022 - 05.2024

SPECIALIZATIONS & SKILLS

- **Languages & Tools:** Python, Java, Go, C, SQL, JavaScript, TypeScript, Shell, Bash, Git, Github, Github Action, GitLab
- **Cloud Native:** **Kubernetes**, **Docker**, Containerd, Operator SDK, CRD, Helm, Prometheus, Grafana, CI/CD, Github Action, AWS, GCP, GKE, Azure, Linux
- **Backend & Frontend:** Spring, Spring Boot, Spring MVC, Maven, Express, Flask, Gin, Redis, MongoDB, MySQL, Firebase, RocketMQ, JWT, Swagger, Postman, AngularJS, Vue.js, React.js, Bootstrap, LangChain, FastAPI

- Machine Learning & Data: **PyTorch**, **Ray**, JAX, **LLMs**, **vLLM**, DSPy, Transformer, Hugging Face, LoRA, DPO, PEFT, RLHF, PPO, Hadoop, Spark, NumPy, Pandas, Scikit-learn