BRUCE (JIANING) FAN

(510) 945-7664 | U.S. Citizen | jf3774@columbia.edu | linkedin.com/in/brucefjn | github.com/brucefjn | New York, NY

EDUCATION

Columbia University — New York, NY
M.S. Data Science for Policy | *GPA: 3.91* — *Expected May 2027*
Coursework: Applying Machine Learning, Building AI Tools with LLMs, Python for Public Policy, Programming and Database Systems

University of California, Los Angeles (UCLA) — Los Angeles, CA
B.A. Political Science | *GPA: 3.91* — *Jun 2023*

EXPERIENCE

Columbia AI Alignment Club (CAIAC) — New York, NY
AI Safety Engineer (Internship) — *Sep 2026 – Nov 2026*

- Building a production **AI agent reliability platform** with an industry partner, using **LangGraph** to orchestrate multi-step planning, tool calling, retrieval, execution, and recovery workflows for real-world tasks.
- Developing a **React + FastAPI** agent monitoring interface that captures execution traces, tool calls, model outputs, human corrections, and task outcomes, enabling real-time debugging and human-in-the-loop intervention.
- Designing an automated **agent evaluation framework** measuring task success, tool-use accuracy, grounding, trajectory quality, failure recovery, latency, and cost; converting failure cases into structured evaluation and training datasets.
- Building an **RLHF/RLAIF-style feedback pipeline** that combines human preference labels, corrections, and model-based evaluations to generate preference data for iterative agent and model improvement, with deployment into the partner's production workflow.

Columbia University – School of International and Public Affairs (SIPA) — New York, NY
AI Software Engineer (Part-Time) — *Sep 2026 – Present*

- Benchmarking hosted and open-source LLMs for a **Canvas LMS AI tutor** across **model quality, latency, and inference cost** to support model selection.
- Building an **LLM evaluation pipeline** with scenario-based test cases and **regression testing** to identify failure modes across answer-seeking, misconception, insufficient-context, and off-topic student interactions.
- Designing **prompting strategies and guardrails** for guided reasoning, concept explanation, answer withholding, and topic boundaries.
- Building and evaluating **RAG pipelines** over course materials, testing **document chunking, retrieval, and context selection** to ground tutor responses in course-specific content.

Noah Medical — San Carlos, CA
AI/ML Engineer — *Dec 2024 – Jul 2025 & May 2026 – Sep 2026*

Agentic AI & Quality-Assurance Systems

- Architected a production **multi-agent workflow** for enterprise evidence retrieval, reasoning, synthesis, and human review, integrating Salesforce and Glean through **REST APIs**.
- Built an **LLM orchestration layer** with model routing, tool calling, RAG, structured outputs, confidence scoring, fallback logic, and human-in-the-loop escalation for regulated complaint analysis.
- Deployed containerized agent services using **Docker, Kubernetes, Redis, and CI/CD**, with automated testing, execution tracing, structured logging, and latency/error monitoring.
- Evaluated hosted and locally deployed LLMs across **output quality, latency, inference cost, and failure severity**, refining model selection and fallback behavior while reducing manual complaint-review time by ∼**95%**.

Complaint Analytics & Monitoring Platform

- Designed and deployed a **Flask backend on Microsoft Azure**, automating ∼**80%** of an FDA-regulated complaint-monitoring workflow with reusable analytics and alerting APIs.
- Built an end-to-end **ETL pipeline** using Salesforce REST APIs, **SOQL, and PostgreSQL**; added **23 pytest unit/integration tests** for regression detection and production reliability.

Sinopuncture Health Center | Website — Fremont, CA
Co-Founder & Full-Stack Web Developer — *Feb 2024 – Jul 2025*

- Designed and launched a production web application using **WordPress, JavaScript, and REST integrations**, implementing responsive service, provider, appointment, and customer-contact workflows.
- Built reusable frontend and booking workflows across desktop and mobile, owning development from requirements and implementation through deployment, debugging, and production maintenance.

U.S. House of Representatives, Office of Rep. Ted W. Lieu — Washington, DC
Legislative Research Intern — *Mar 2022 – Jun 2022*

- Built reproducible **NLP data pipelines** for legislative records and constituent feedback, covering preprocessing, labeling, dataset versioning, and train/test workflows.
- Fine-tuned a **DistilBERT transformer classifier** and benchmarked it against TF-IDF/logistic-regression baselines, achieving ∼**91% accuracy and macro-F1** on held-out data.
- Developed a scalable **LLM-assisted analysis pipeline** across **786K public comments, 6,145 EPA dockets, and 36 rulemakings**, combining semantic retrieval, LLM verification, and sentence-transformer classification.
- Built a **human-in-the-loop evaluation framework** with blinded annotation and inter-rater reliability analysis, identifying a classifier–human agreement failure of $\kappa = 0.137$ before downstream interpretation.

PUBLICATION

Jianing Fan, Yue Yao. "Who Gets Heeded? An Obligation-Level Audit of Responsiveness in EPA Rulemaking." **Accepted (Oral), ACM EAAMO 2026**, Aug. 2026. arXiv:2608.10329.

TECHNICAL SKILLS

Languages & Software: Python (PyTorch, TensorFlow), Java, JavaScript, HTML/CSS, SQL/SOQL, Flask, REST APIs, Git/GitLab, pytest
AI Agents & ML: LangGraph, RAG, tool calling, LLM orchestration, RLHF/RLAIF, Hugging Face Transformers, human-in-the-loop systems
Data & Infrastructure: PostgreSQL, Redis, ETL, data pipelines, Salesforce, AWS, GCP, Azure, Docker, Kubernetes, CI/CD