Suraj Ranganath

suraj.ranganath@gmail.com | GitHub | Google Scholar | LinkedIn | San Diego, CA

Positioning

AI systems builder and researcher working across agentic AI, identity security, graph machine learning, AI safety, and world models. I turn research ideas into usable systems, published work, patents, and production products.

Education

University of California San Diego Sept 2025 – June 2027
M.S. Data Science, GPA 4.0/4.0

Vellore Institute of Technology Aug 2020 – May 2024
B.Tech. Computer Science and Engineering, GPA 3.98/4.0, Top 5 in class

Experience

BalkanID May 2023 – Jun 2025
Machine Learning Engineer

- Built BalkanID Copilot, an agentic identity-security system using LLMs and knowledge graphs; adopted by 6+ enterprise customers and associated with $200K+ revenue.
- Designed text-to-Cypher over a multi-tenant Identity Intelligence Graph with 5M+ entities and relationships.
- Built MCP server and connector integrations for actions across 20+ enterprise services.
- Built ETL pipelines ingesting about 50GB per tenant from 100+ enterprise sources.
- Owned model reliability, evals, monitoring, and safe rollout loops.
- Co-inventor on 4 U.S. patents in agentic AI and graph machine learning for identity security.

IIT Bombay, EdTech Lab Jan 2023 – Jun 2023
Research Intern

- Managed EEG, eye-tracking, GSR, video, interview, and spatial problem-solving data for 300+ hours of multimodal learning analytics data.
- Contributed to a Nature Scientific Data publication through data collection, anonymization, validation, and analysis.

MySivi / Speakify, YC W22 Jul 2022 – Aug 2022
Machine Learning Intern

- Shipped a T5 grammar correction system for Indian English to 50,000+ users.

Selected Research and Systems

StealthRL: GRPO-based adversarial paraphrasing framework for stress-testing AI-text detectors.
KV Cache Quantization: 33-method empirical study for self-forcing video generation efficiency.
PIPE-RDF: LLM-assisted pipeline for schema-specific enterprise RDF benchmark generation.
TrailKarma: Offline-first trail intelligence app using Android, BLE relay, Databricks, Solana, and on-device ML.
Marin NAMO/NAMO-D: Merged open-source optimizer support in marin-community/marin.

Skills

LLM agents, RAG, MCP, text-to-Cypher, knowledge graphs, graph ML, AI safety, red teaming, GRPO, LoRA, evaluation harnesses, model monitoring, world models, video diffusion, KV-cache quantization, multimodal data, learning analytics, PyTorch, JAX, FastAPI, Docker, Kubernetes, Neo4j, Databricks, Android, BLE.