KAI-CHUN SU

Tainan, Taiwan | +886-98981238 | pipichun2@gmail.com | github.com/vanillaSky00

PROFESSIONAL SUMMARY

CS student at NCKU focused on LLM agent systems, RAG pipelines, and AI evaluation frameworks. Passionate about building production-grade agentic backends using LangGraph with function calling, pgvector-based retrieval, and LLM-as-judge evaluation. Competition experience applying RAG to financial Q&A (AI CUP — E.SUN Bank) and deep learning to time-series forecasting. Grounded in application security through CTF competitions and penetration testing.

EDUCATION

National Cheng Kung University — B.S. in Computer Science and Information Engineering **Sep 2023 - Present**

- Outstanding Academic Achievement Award (2024).

TECHNICAL SKILLS

Languages: Python, Java, C, C#, JavaScript/TypeScript, Bash
Backend & Infra: FastAPI, Spring Boot, Docker, PostgreSQL, Redis, Git, GitHub Actions
Cloud & DevOps: AWS (SQS, ECS, S3, CloudWatch), OIDC, CI/CD pipelines, Linux
AI/LLM: LangGraph, RAG pipelines, pgvector, Function Calling, Context Engineering, LLM-as-Judge
Security: Burp Suite, SQL/prompt injection, XSS, privilege escalation, network traffic analysis (Wireshark), file upload attacks
Other: Blockchain (Merkle trees, SHA-256), WebSocket, RESTful APIs, SQL/NoSQL, Embedded Linux (QEMU/Buildroot)

SELECTED PROJECTS & COMPETITION

Human-AI Co-Worker Cooking Game **2025 - 2026**

LangGraph, Function Calling, Context Engineering, LangSmith, AWS (S3, OIDC), FastAPI, WebSocket, Docker, Unity/C#

- Architected a self-improving LLM agent for a human-AI co-worker cooking game, inspired by Voyager — proposing tasks, retrieving learned procedures, executing actions via function calling, and accumulating reusable skills via an 8-node LangGraph FSM backed by pgvector and a decoupled Unity/FastAPI pipeline.
- Resolved LLM state ambiguity through SayCan-inspired affordance-based context engineering — structuring perception into authoritative blocks (held item, reachable objects, assembly state). A two-level retry harness — function-level tool recovery and a critic-retry-curriculum node-based feedback loop — ensures stable end-to-end burger completion within 5 minutes.
- Set up CI/CD via GitHub Actions with AWS OIDC federation, running LLM-as-judge evaluation on EC2 with results stored on S3. Generated synthetic test datasets using LLM-as-user and integrated LangSmith tracing for end-to-end agent behavior monitoring, driving evaluation and context engineering iteration loops. Benchmarked at ~5–10s per decision cycle.

picoCTF 2025 — CMU Cybersecurity Competition **2025**

Prompt Injection, SQL Injection, XSS, Session Hijacking, Linux Privilege Escalation, Digital Forensics

- Placed 262 / 10,460 globally, applied techniques from Hack The Box training — Burp Suite for web exploitation, ffuf for endpoint fuzzing, Wireshark for network packet analysis, and LLM prompt injection and jailbreaking.

AI CUP 2024 — E.SUN Commercial Bank Financial RAG Competition **2024**

Python, RAG, bge-m3, Hybrid Retrieval, Reranking, Domain-Specific Chunking, OCR (Tesseract)

- Ranked top 25% (487 teams) building a RAG system for financial Q&A across three domains (bank FAQ, insurance policies, financial reports), with domain-specific chunking strategies tailored to each document structure.
- Achieved 96% accuracy on FAQ, 94% on insurance, and 84% on finance retrieval using bge-m3 embeddings with hybrid retrieval and keyword-based reranking. Built an OCR sub-pipeline (Tesseract) to extract text from financial report images.

AI CUP 2024 — Microclimate Solar Power Generation Prediction **2024**

Python, PyTorch, LSTM, Time-Series Forecasting, Deep Learning, pandas, matplotlib

- Ranked top 30% (934 teams) predicting next-day solar power output across 17 microclimate locations from sensor data (irradiance, temperature, humidity, wind).
- Built a PyTorch LSTM regression pipeline with per-location normalization and a custom 9am-to-9am windowing scheme mapping today's sensor features to tomorrow's output.

LEADERSHIP

Technical Project Lead — NCKU CSIE (Paprika / Mewi) **Mar 2025 - Present**

- Coordinated two concurrent AI agent projects via GitHub Projects, managing task assignment, code reviews, and PRs across cross-functional teams; authored Architecture Decision Records to align technical direction.

NCKU Rollerblading Club — Club Leader **Mar 2025 - Present**

- Participated in a 100+ person overnight New Year's Eve street-skating event, managing club registration and safety coordination. Organized 50+ participant cross-university skating initiatives, overseeing route planning and risk-management protocols.

NCKU CTF Club — Design Manager **Sep 2024 - Present**

- Scaled community engagement to 1,000+ members across Instagram, Facebook, and Discord; campaign materials featured at HITCON national cybersecurity conference. Hack The Box Academy: Top 5% ranking, 88 targets compromised.