SANGAMESH GIRISH DANDIN

AI Engineer | Machine Learning | LLM Applications | AI Security

Dublin, Ireland | +353 89 255 9418 | sangameshs2003@gmail.com

LinkedIn | Github | Portfolio

PROFESSIONAL SUMMARY

AI Engineer and MSc Artificial Intelligence student at National College of Ireland who builds end-to-end AI applications, from model training and explainability to FastAPI services, React interfaces and Docker packaging. Project work includes a gradient-boosted malware classifier with SHAP explanations, a multi-LLM prompt injection defence layer, LLM-driven security automation, and ML-based vulnerability exploitation prediction. Brings software engineering and cybersecurity depth (CompTIA Network+, PNPT in progress) to ML work; published IEEE author.

TECHNICAL SKILLS

Languages: Python, SQL, JavaScript, Bash

Machine Learning & Deep Learning: Scikit-learn, XGBoost, LightGBM, Random Forest, PyTorch, TensorFlow/Keras, CNN-LSTM Models, SHAP Explainability, Feature Engineering, Model Evaluation

LLMs & NLP: LLM Application Development, Prompt Engineering, Multi-Agent LLM Pipelines, RAG, FAISS, Sentence Transformers, Semantic Search, Groq API (LLaMA 3), NLP

Data Science: Pandas, NumPy, Matplotlib, Jupyter, Data Preprocessing, Exploratory Data Analysis, Statistical Analysis, Predictive Modelling, Data Visualisation

Backend & Software Engineering: FastAPI, Flask, REST APIs, Server-Sent Events, Celery, Redis, SQLAlchemy, PostgreSQL, SQLite, React, Tailwind CSS, Git, GitHub Actions (CI)

Cloud & Deployment: AWS (EC2, S3, IAM, VPC), Docker, Docker Compose, Linux, Supabase, Vercel

AI Security & Cybersecurity: Prompt Injection Defence, LLM Security, Adversarial Machine Learning, Malware Classification, Threat Intelligence, CVE/CVSS Analysis, Vulnerability Assessment, OSINT, Nmap, GDPR-Aware Design

PROJECTS

ByteHunter | AI Malware Classification Engine GitHub

Tech: *Python, LightGBM, XGBoost, Scikit-learn, SHAP, LIEF, FastAPI, SQLite, React, Docker*

- Built a full-stack static malware analysis tool for Windows PE executables, using LIEF to extract the 2,381-dimension EMBER feature vector (header fields, section entropy, imports, byte histograms and strings).
- Developed a LightGBM-based malware detection model trained on the EMBER 2018 dataset, alongside an XGBoost classifier for malware family classification using Microsoft BIG-2015.
- Integrated SHAP TreeExplainer to show the features driving each verdict, and mapped imported Win32 APIs to suspicious behaviours such as code injection, persistence and anti-debugging.
- Served the pipeline through an async FastAPI service with batch scanning, SHA-256 result caching, SQLite scan history, PDF forensic reports and scheduled model hot-swapping, packaged with Docker Compose.

ZeroTrace AI | Threat Intelligence Platform

Tech: *Python, XGBoost, FAISS, LLMs (Groq API), FastAPI, Streamlit, PostgreSQL, Docker*

- Built a threat intelligence platform that consolidates vulnerability data from NVD, CISA KEV and ExploitDB.
- Trained an XGBoost model to predict the likelihood of a vulnerability being exploited from CVSS and CVE features.
- Implemented FAISS-based semantic search for vulnerability discovery and used LLMs to summarise CVEs into analyst-ready threat intelligence reports.

RedTeam Copilot | AI-Powered Penetration Testing Assistant GitHub

Tech: *Python, FastAPI, Groq LLM API, Nmap, NVD API, CISA KEV, Supabase, React, Docker*

- Built an agentic recon-to-report assistant for authorised testing that chains four async agents: OSINT, port scanning, CVE matching and LLM-generated reporting.
- Combined DNS resolution, subdomain brute-forcing and crt.sh certificate-transparency lookups for OSINT; ran Nmap service and version detection with timeouts and an asyncio TCP-connect fallback when Nmap is blocked.
- Matched discovered service versions to CVEs through the NVD API and flagged entries in the CISA Known Exploited Vulnerabilities catalogue to prioritise risk.
- Streamed live pipeline progress to a React dashboard with Server-Sent Events and generated PDF reports with a risk matrix, attack chain and remediation roadmap, stored in Supabase.

ZeroInject Shield | Prompt Injection Defence System GitHub

Tech: *Python, FastAPI, Groq API (LLaMA 3.3 70B, LLaMA 3.1 8B, Qwen3 32B), SQLAlchemy, React*

- Built zero-trust middleware that screens user prompts for injection, jailbreak and mixed-intent attacks before they reach an application's core LLM.
- Designed a five-stage pipeline (sanitise, verify, score, apply policy, enforce) in which three different LLM verifiers rate each prompt and a weighted consensus combines their scores into a single injection score.
- Implemented a policy engine that maps the score to BLOCK, SANITIZE or ALLOW, so ambiguous prompts are cleaned and the user's legitimate question still reaches the model instead of being rejected outright.
- Logged every decision to a database behind a React analytics dashboard, and built an e-commerce chatbot sandbox to test detection against benign and harmful prompt datasets.

SpearSim | AI Phishing Simulation Platform GitHub

Tech: *Python, FastAPI, Groq (LLaMA 3), React, Supabase (PostgreSQL), Celery, Redis, SendGrid, Docker, GitHub Actions*

- Built a platform for authorised phishing-awareness campaigns that uses LLaMA 3 to generate role-aware phishing emails across 5 scenario types and 3 difficulty levels.
- Designed for GDPR: pseudonymised employee data before every LLM call, recorded credential entry as a flag with zero credential storage, and added a consent gate, append-only audit log, erasure endpoint and 90-day retention.
- Implemented open and click tracking with per-employee risk scoring, campaign analytics, a post-simulation awareness quiz and LLM-generated remediation advice, exported as PDF and CSV reports.
- Added multi-organisation support with admin/analyst/viewer RBAC, Celery/Redis background email delivery via SendGrid, and GitHub Actions CI for linting, type checks and Docker builds.

PROFESSIONAL EXPERIENCE

Data Science Intern | CodeClause, Remote 2023

- Developed an Image Caption Generator on Flickr8k with an Xception CNN encoder and LSTM decoder in TensorFlow/Keras, supported by NLP text cleaning and tokenisation pipelines.
- Built a Loan Status Prediction system comparing 5 ML models (Logistic Regression, SVM, Decision Tree, Random Forest, Gradient Boosting) with automated preprocessing, feature engineering and a Tkinter GUI.
- Standardised reusable preprocessing and evaluation pipelines, reducing model experimentation and deployment iteration time across projects.

EDUCATION

MSc in Artificial Intelligence Jan 2026 - Jan 2027 (Expected)
National College of Ireland, Dublin, Ireland

B.E. in AI and Data Science Feb 2023 - Oct 2025
Nitte Meenakshi Institute of Technology, Bengaluru, India | CGPA: 7.49/10

Diploma in Computer Science Aug 2019 - Aug 2022
N.V. Polytechnic College, Kalaburagi, India | 80%

LEADERSHIP

Vice President | ACM NMIT Student Chapter, Bengaluru Jan 2024 - Jul 2025

- Led 10+ technical workshops and cybersecurity training sessions covering vulnerability research, offensive security and AI system design for 200+ participants.
- Managed cross-functional student teams and coordinated faculty partnerships to run semester-long technical programmes and security events.
- Organised AI security and adversarial ML research seminars while serving concurrently as Vice Secretary of IEEE CIS NMIT, contributing to 30% chapter membership growth.

CERTIFICATIONS

- PNPT: Practical Network Penetration Tester (In Progress) TCM Security
- CompTIA Network+ (N10-008) CompTIA
- AWS Academy Graduate: Cloud Architecting (EC2, S3, VPC, IAM) Amazon Web Services
- Cybersecurity Analyst Job Simulation Tata Group / Forage

PUBLICATION

AlertSphere: Alerting About Disasters and Ensuring Safety Using AI
Nayana B.P., K.S. Prakruthi, Ananya Rekha Ashok, **Sangamesh Girish Dandin**, Tanveer Akhlaque Ahmed.
IEEE ICAMIDA 2025, Aurangabad, India DOI: 10.1109/ICAMIDA64673.2025.11209472