# ANTHONY FAUSTINE

## Principal AI Engineer | AI Systems Architect & Production ML Leader

sambaiga@gmail.com +353-899-896-099 sambaiga sambaiga.github.io Dublin, Ireland

## PROFESSIONAL SUMMARY

*Principal AI Engineer and PhD-level researcher with 10+ years designing, building, and operating production-grade AI systems across energy, industrial, and semiconductor domains. Deep expertise in **transformer-based and attention architectures**, **sequence-to-sequence models**, **CNNs**, **graph neural networks**, and **probabilistic deep learning**, from composable backbone-and-head design through scalable inference and lifecycle governance. Author of **10+ peer-reviewed publications** in IEEE Transactions on Power Systems, Smart Grid, and Industrial Informatics, and creator of **Twiga**, an open-source production forecasting library built on these architectures. As a Principal Investigator, secured **€1.6M** in competitive funding and delivered measurable outcomes including a **25%** reduction in forecasting error and enterprise-wide MLOps standardisation. Awarded the **Eaton STAR Leadership Award**. Known for raising the engineering bar through architecture ownership, cross-functional technical leadership, and translating deep research into resilient, production-scale AI systems.*

## KEY ACHIEVEMENTS

- **AI Systems Architecture**: Designed composable AI architectures combining **sequence-to-sequence models**, **graph attention networks**, **gated attention fusion**, and **CNNs** with interchangeable probabilistic heads (Normal, Laplace, quantile, conformal), deployed at scale across power, manufacturing, and semiconductor domains.
- **Strategic Funding & Leadership**: Appointed Principal Investigator and secured **€1.6M** in competitive research funding, driving AI innovation across intelligent energy systems and data-centre infrastructure.
- **Deep Learning Research to Production**: Delivered a **25%** reduction in forecasting error by designing multi-horizon models including **FPSeq2Q** (fully parameterised seq2seq with uncertainty), **N-HiTS**, and **conformal MLP**; productised as the open-source **Twiga** library (REST-serving, MLOps, HPO built in).
- **Cloud-Native MLOps at Scale**: Standardised enterprise MLOps workflows on **AWS**/**Azure** using **Docker**, **Airflow**, **MLflow**, and **CI/CD pipelines**, enabling automated retraining, governance, drift-triggered promotion, and scalable model serving.
- **Model Reliability & Observability**: Implemented production drift-detection and monitoring frameworks using **Evidently AI** and **Weights & Biases**, improving operational stability across multi-model deployments.
- **Recognition**: Recipient of the **Eaton STAR Leadership Award** for engineering innovation, data-driven decision support, and cross-functional technical leadership.
- **Thought Leadership**: **10+ peer-reviewed publications** in *IEEE Transactions on Power Systems*, *Smart Grid*, and *Industrial Informatics* spanning deep learning architectures, probabilistic forecasting, and neural energy disaggregation; creator of open-source **Twiga** and technical speaker at international IEEE events.

## EXPERIENCE

2026–Present **Analog Devices**

**Principal Engineer, Machine Learning** Limerick, Ireland

**Summary:** Building production-ready ML systems for supply chain management, leveraging demand forecasting and intelligent planning to ensure ADI's Intelligent Edge technology reaches the right place at the right time across a complex global network.

- **MLOps Platform Architecture:** Architect of five-layer platform (Governance, Contracts, Execution, Pipelines, SDK) with formal data/model contracts, schema validation, and ModelCard-gated promotion with MLflow lineage — delivered in 23 PRs within 2 months (Level 0 → Level 1 maturity).
- **ML Algorithms & Systems:** Unified 6 quantile forecasting algorithms under composable BaseRegressor interface with single registry and Optuna HPO; engineered Numba-compiled inference engine maintaining legacy parity while enabling memory-safe prediction at scale.
- **Production Optimization:** Eliminated OOM failures by replacing unbounded CROSS JOINs with range JOINs and lazy DuckDB scans; reduced peak memory 8–15 GB → <2 GB, enabling full forecast cycles in standard containers.
- **Technical Leadership:** Established CI/CD quality baseline (1,275+ pytest tests, Ruff, ty, GitHub Actions) and mentored engineers on platform architecture and production ML best practices as the team scales globally.

2022–2026 **Eaton**

**Lead Data Scientist & Principal Investigator (NTI)**, Centre for Intelligent Power (CIP) Dublin, Ireland

**Summary:** Lead Data Scientist responsible for architecting production AI systems, advancing forecasting capabilities, and delivering strategic analytics across power distribution, data centres, and EV infrastructure.

- **Research Leadership & Funding:** Secured **€1.6M** as Principal Investigator for AI-driven data-centre power and cooling management; awarded the **Eaton STAR Leadership Award** for engineering innovation and cross-functional impact.
- **AI Architecture & Deep Learning:** Designed and implemented production forecasting systems using **transformer-based**, **seq2seq**, **CNN**, and **graph attention** architectures (N-HiTS, FPSeq2Q, GANF, gated-attention fusion) with composable probabilistic heads, delivering a **25%** improvement in long-horizon accuracy across PV, load, and EV pipelines.
- **Cloud-Native AI Infrastructure:** Built and operated scalable AI services on **AWS** using containerised microservices (Docker, FastAPI), **MLflow** model registry, and **Airflow**-orchestrated retraining pipelines with automated drift-triggered promotion.
- **EV Infrastructure Planning:** Architected long-horizon demand forecasting for EV charging infrastructure using statistical, Bayesian, and gradient-boosting approaches to support capacity expansion and investment decisions.
- **Substation Intelligence:** Built a real-time AI analytics platform covering thousands of substations with anomaly detection, equipment-risk scoring, and maintenance prioritisation; productised as a customer-facing SaaS offering.
- **Battery Analytics:** Engineered a predictive framework for BMS integrating validation and labelling protocols to enable high-accuracy State of Health (SoH) and Remaining Useful Life (RUL) estimation.
- **Model Reliability & Inference:** Implemented production monitoring, drift detection, and model-serving infrastructure (FastAPI, Uvicorn) with online performance tracking and inference optimisation for low-latency endpoints.
- **Technical Leadership:** Mentored senior engineers and data scientists, advised executives on AI strategy, and represented Eaton at international technical events; established AI engineering standards adopted across the NTI organisation.

2021–2022 **Irish Manufacturing Research**

**Senior Industrial Analytics Researcher**, Industrial Analytics Dublin, Ireland

**Summary:** Delivered predictive maintenance and computer vision solutions that improved reliability, efficiency, and operational visibility across advanced manufacturing environments.

- **Predictive Maintenance Leadership:** Directed the Data-Driven Maintenance Service (DDMS) program, reducing unplanned downtime by **30%** and increasing OEE by **25%** through IoT-driven monitoring and predictive modelling.
- **CNC Health Analytics:** Engineered spindle health diagnostics using unsupervised learning and **Hidden Markov Models (HMMs)**, enabling early-stage fault detection and informed maintenance decisions.
- **Computer Vision Automation:** Developed AI-driven anomaly detection systems for visual inspection, cutting defect detection time by **40%** and improving operational efficiency by **20%**.
- **MLOps & Deployment Practices:** Implemented model versioning, automated monitoring, and scalable deployment workflows, improving development efficiency by **25%**.
- **Cross-Functional Delivery:** Collaborated with manufacturing engineers, operations teams, and solution integrators to translate ML outputs into actionable improvements on industrial production lines.

2020–2021 **CeADAR**

**Data Scientist & ML Researcher**, Applied Analytics Dublin, Ireland

**Summary:** Delivered large-scale load analytics, Bayesian forecasting, and geospatial automation systems supporting energy intelligence and computer vision research programs.

- **Large-Scale Load Profiling:** Built a scalable load-profiling system using variational auto-encoders and clustering to segment **10,000+** building load profiles for downstream analytics.
- **Probabilistic Forecasting:** Implemented cluster-specific forecasting using **Bayesian state-space models**, improving accuracy by **22%** and reducing peak-load planning errors by **18%**.
- **Geospatial Labelling Automation:** Developed an automated pipeline for satellite imagery labelling using Sentinel-2 data and **UNet**-based models, generating **1,000+** labelled samples with **95%** accuracy.
- **Deployment & Workflow Optimization:** Delivered a Python/GDAL/Docker workflow that reduced manual labelling efforts by **70%**, improving dataset throughput and model development velocity.

2017–2019 **IDLab, imec, Ghent University**

**Machine Learning Researcher**, Smart Energy & IoT — Ghent, Belgium

**Summary:** Conducted applied machine learning research in Non-Intrusive Load Monitoring(NILM), multimodal sensing, and smart-home analytics, advancing academic and industry-funded innovation programs.

- **NILM Computer Vision:** Developed a CNN-based computer vision framework using V–I trajectories for non-intrusive load monitoring, improving appliance-level detection accuracy by **35%**.
- **Smart-Home Event Detection:** Designed and deployed a multimodal ML pipeline leveraging heterogeneous sensor data to achieve **92%** event-detection accuracy in smart-home environments.
- **Funding & Industry Collaboration:** Collaborated with industrial partners on research proposals, contributing to securing over **€500K** in funding to advance NILM and intelligent home analytics.

2010–2017 **University of Dodoma**

**Assistant Lecturer & Digital Innovation Researcher** — Dodoma, Tanzania

**Summary:** Led applied research and teaching initiatives in computer vision, wireless communication, mobile health, and precision agriculture, while advancing national capacity building in data science and machine learning across Tanzania.

- **AI for Agriculture:** Developed and deployed a computer vision model for banana disease detection, enabling early diagnosis and low-cost precision agriculture for smallholder farmers.
- **AI Capacity Building:** Designed and delivered a national data science and machine learning training program, including workshops and a summer school for university students across Tanzania.
- **Mobile Health Systems:** Built a mobile ADR reporting platform for the Tanzania Food and Drugs Authority, increasing national adverse drug reaction reporting.
- **IoT & Sensor Networks:** Developed wireless sensor and mobile-based systems for water quality monitoring in the Lake Victoria Basin.
- **Teaching & Leadership:** Taught Information Theory, Wireless Communication, Mobile Computing, and Data Management; organized seminars on IoT, AI, and emerging technologies.

## TECHNICAL SKILLS

- **Languages & Frameworks:** Python, SQL; PyTorch, TensorFlow, scikit-learn, Lightning, Hugging Face Transformers
- **AI Architectures:** Transformers (BERT/GPT-style, ViTs), seq2seq models, graph attention networks, N-HiTS, LSTM, CNNs; composable probabilistic heads, Bayesian deep learning, quantile regression, conformal prediction
- **Cloud & Infrastructure:** AWS (SageMaker, EC2, S3), Azure ML; Docker, Kubernetes, FastAPI, Uvicorn, CI/CD
- **MLOps & Governance:** MLflow, Airflow, Weights & Biases, Evidently AI; model versioning, drift detection, automated retraining, inference optimisation
- **Data & Systems:** Distributed data pipelines, streaming/batch processing, feature engineering, model-serving APIs

## EDUCATION

2025 **PhD in Computer Science & Engineering,** Instituto Superior Técnico — Lisbon, Portugal

2024 **MSc in Leadership, Technology & Innovation,** Technological University Dublin — Dublin, Ireland

2012 **MSc in Telecommunications Engineering,** University of Dodoma — Dodoma, Tanzania

2009 **BSc in Electronics and Communication Science,** The University of Dar es Salaam — Dar es Salaam, Tanzania

## SELECTED PUBLICATIONS

10+ total, full list at sambaiga.github.io

- Faustine, A., Jardim Nunes, N., Pereira, L. (2025). **Efficiency through Simplicity: MLP-based Approach for Net-Load Forecasting with Uncertainty Estimates in Low-Voltage Distribution Networks**. *IEEE Transactions on Power Systems*, 40(1), 46-57. DOI: 10.1109/TPWRS.2024.3400123.
- Faustine, A., Pereira, L. (2022). **FPSeq2Q: Fully Parameterized Sequence-to-Quantile Regression for Net-Load Forecasting With Uncertainty Estimates** [seq2seq deep learning architecture]. *IEEE Transactions on Smart Grid*, 13(3), 2440-2451. DOI: 10.1109/TSG.2022.3148699.
- Bousbiat, H., Faustine, A., Klemenjak, C., Pereira, L., Elmenreich, W. (2023). **Unlocking the Full Potential of Neural NILM: On Automation, Hyperparameters and Modular Pipelines** [neural network + MLOps]. *IEEE Transactions on Industrial Informatics*, 19(5), 7002-7010. DOI: 10.1109/TII.2022.3206322.

- Faustine, A., Pereira, L., Klemenjak, C. (2021). **Adaptive Weighted Recurrence Graphs for Appliance Recognition in Non-Intrusive Load Monitoring** [graph-based ML + CNNs]. *IEEE Transactions on Smart Grid*, 12(1), 398-406. DOI: 10.1109/TSG.2020.3010621.
- Faustine, A., Pereira, L. (2024). **Conformal Multilayer Perceptron-Based Probabilistic Net-Load Forecasting for Low-Voltage Distribution Systems with Photovoltaic Generation**. *IEEE SmartGridComm 2024*. 10.1109/SmartGridComm60555.2024.1073810.