SACHIN KAHAWALA

Melbourne, AU

S.Kahawala@latrobe.edu.au LinkedIn GitHub Google Scholar

SUMMARY

Data Scientist and AI Engineer with 6+ years of experience and a PhD in Artificial Intelligence. Experienced in designing, building, and shipping AI and machine learning systems across research and industry settings, spanning generative AI, deep learning, forecasting, and computer vision. Comfortable working end-to-end from rapid prototyping and experimentation to production deployment on cloud platforms, with a track record of delivering reliable, measurable outcomes on complex, data-intensive problems in real-world environments. Brings a strong algorithmic and problem-solving foundation built through years of competitive programming.

SKILLS

LLM	LangChain, LangGraph, LlamaIndex, DeepEval, Open AI GPTx, Azure Open AI, Gemini, Open Source LLMs, Microsoft semantic kernel, Hugging Face, DSPy
Programming	Python, C++, R
Machine Learning & DL	PyTorch, Scikit Learn, Tensorflow
Databases	SQL DBs, NoSQL, Redis
Web, API & Frontend	Flask, FastAPI, React
Data Viz & BI platform	PowerBI, Tableau
DevOps/MLOps/CI/CD	MLFlow, GitHub Actions, Azure DevOps, Git
Cloud Technologies	Azure, GCP, AWS
AI and Data Platforms	Azure ML, Microsoft Foundry, Vertex AI, DataBricks
Microservices	Docker

EDUCATION

PhD in Artificial Intelligence, La Trobe University, Australia 2021 - 2024

B.Sc. (Hons) Computer Science and Engineering, University of Moratuwa, Sri Lanka 2015 - 2019

GCE Advanced Level, Dharmaraja College, Sri Lanka 2014
Ranked 18th (Top 0.03%) out of approximately 60,000 candidates in the Physical Science stream in Sri Lanka

EXPERIENCE

Senior Data Scientist and AI Architect
Jan 2026 - Present
PS Hummingbird *Melbourne, AU*

Data Scientist and AI Architect
Oct 2024 - Dec 2025
PS Hummingbird *Melbourne, AU*

- Led the design and deployment of production-grade AI systems for enterprise clients, delivering solutions across multiple domains during the rapid adoption of LLM technologies.
- Built end-to-end LLM pipelines to transform unstructured data into structured outputs for downstream applications.
- Developed systems to incorporate relevant contextual information into model outputs, improving accuracy and robustness across applications.
- Designed multi-agent workflows that break down complex tasks into structured stages, mimicking human decision-making processes and handling dependencies between tasks.
- Implemented scalable orchestration and processing systems to support concurrent workloads, ensuring reliability, fault tolerance, and consistent output quality.

- Established evaluation approaches combining rule-based checks and model-based assessment to monitor performance and iteratively improve system outputs.
- Deployed solutions on Microsoft Azure, integrating AI pipelines into production environments with monitoring, logging, and continuous improvement practices.

Data Scientist Sep 2020 - Sep 2024
Centre for Data Analytics and Cognition, La Trobe University *Melbourne, AU*

- Developed and implemented AI, ML, and CI/CD solutions, seamlessly integrating them into production environments.
- Collaborated with technical and business stakeholders to offer consultancy services for building AI-driven solutions.
- Analyzed complex, data-intensive problems to extract actionable insights and create detailed visualizations.
- Conducted thorough research and feasibility studies to identify the most suitable technologies and best practices for effectively addressing business challenges.

Doctoral Candidate May 2021 - July 2024
La Trobe University *Melbourne, AU*

- Developed state-of-the-art energy-efficient algorithms using Vector Function Architecture (VFA)
  - Utilizing VFA for few-shot manifold learning and enhanced integration with neuromorphic hardware.
  - Developed a general framework for the representation of graphs using VFA: revolutionizing GNNs by eliminating learning, leading to significant gains in both performance and efficiency compared to conventional GNNs.
  - The research is geared towards practical advancements in AI, emphasizing real-world applicability, performance, and efficiency.

Research Collaborator Sep 2020 - June 2024
Luleå University of Technology *Luleå, Sweden*

- Energy efficient algorithm development for Neuromorphic hardware
  - Developing unsupervised learning algorithms for the Intel Neuromorphic chip (Loihi) in collaboration with Intel

Postgraduate Tutor June 2022 - Jan 2024
La Trobe University *Melbourne, AU*

- Tutored Master of Business Analytics courses
  - Artificial Intelligence and Hyperautomation - Topics include ML & AI concepts, NLP, BI Tools, Agile
  - Principles of Business Analytics - Topics include Analytical tools (SAP, Power BI) and SQL
  - Cloud Platforms and Analytics - Topics include managing and leveraging cloud platforms for automation, security, data integration, and AI applications.

Research Assistant Jan 2020 - Aug 2020
IntelliSense Lab, University of Moratuwa *Sri Lanka*

- Designed and developed a deep learning–based apparel defect classifier for automated visual quality inspection in garment manufacturing.

Research Assistant (internship) July 2018 - Jan 2019
Zeptolytics (Pvt) Ltd *Sri Lanka*

- Design and development of a recommendation system

SELECTED PUBLICATIONS

- **Sachin Kahawala** et al. "Graph Vector Function Architecture." In: *Neural Networks* (2025).
- **Sachin Kahawala** et al. "Hypervector Approximation of Complex Manifolds in Smart City Internet of Things Infrastructure." In: *Smart Cities* 7.6 (2024), 131.
- **Sachin Kahawala** et al. "Robust Multi-Step Predictor for Electricity Markets with Real-Time Pricing." In: *Energies* 14.14 (2021).
- Evgeny Osipov, **Sachin Kahawala** et al. "Hyperseed: Unsupervised Learning With Vector Symbolic Architectures." In: *IEEE Transactions on Neural Networks and Learning Systems* (2022), pp. 1–15.
- Rashmika Nawaratne, **Sachin Kahawala** et al. "A Generative Latent Space Approach for Real-Time Road Surveillance in Smart Cities." In: *IEEE Transactions on Industrial Informatics* 17.7 (2021), pp. 4872–4881.

AWARDS AND ACHIEVEMENTS

- **Amazon ML Summer School [2024]**, Participant
- **IEEEXtreme v17.0 - Country Rank 5th [2023]**, Country rank – 5th, Global rank - 355th
- **National Patent [2022]** Method And Apparatus for Detecting Surface Defects
- **IEEEXtreme v11.0 - Country Rank 10th [2017]**, Country rank – 10th, Global rank - 201th
- **IEEEXtreme v10.0 - Country Rank 6th [2016]**, Country rank – 6th, Global rank - 181th

SELECTED INDUSTRY-ORIENTED PROJECT PORTFOLIO

Planning Compliance Automation Platform

Australia (Apr 2025 – Present)

- **Data:** Statutory planning schemes, council policies, developer submissions, prior approvals.
- **Objective:** Automate assessment of development proposals against multi-jurisdiction planning controls and generate draft compliance reports.
- **Outcomes:** Led the design and delivery of a production AI platform, deployed and used to process large planning and policy documents, identify compliance gaps, and generate structured, report-ready outputs with cited evidence. The system models human assessment workflows through multi-step pipelines, enabling scalable, consistent analysis and reducing manual review time by >70%. Developed an evaluation framework combining automated and human-in-the-loop assessment to ensure accuracy, citation grounding, and reliability in production. Instrumented the pipeline to track usage, cost, and latency for ongoing performance optimisation.
- **Tech:** Azure OpenAI, Azure AI Search, FastAPI, React, Azure Container Apps, Cosmos DB, Blob Storage, Azure Service Bus, Application Insights, LangGraph, DeepEval.

Enterprise Knowledge Retrieval Assistant (RAG)

Public Sector, Australia (Oct 2024 – Feb 2025)

- **Data:** Policy manuals, SOPs, FAQs, and internal documents (PDF, DOCX, HTML).
- **Objective:** Deploy a secure retrieval-based assistant to answer internal queries with traceable sources.
- **Outcomes:** Delivered a production RAG system with source-linked responses and feedback loops, achieving <1.8s latency and reducing Tier-2 escalations by ∼30%.
- **Tech:** Azure OpenAI, Azure AI Search, Python, LangChain, Azure Functions, Blob Storage.

Real-Time Voice Agent with CRM Integration

PS Hummingbird, Australia (Feb 2025 – March 2025)

- **Data:** Simulated voice conversations, sample customer records, and CRM data accessed via tool calls.
- **Objective:** Develop a working prototype of a real-time voice assistant to demonstrate agentic CRM automation capabilities to an enterprise client.
- **Outcomes:** Built an end-to-end voice agent demo using LiveKit for real-time media streaming and Gemini real-time models on Vertex AI as the con-

versational core. Integrated Azure Cosmos DB for contextual grounding and connected the agent to Microsoft Dynamics 365 through MCP tool servers, enabling it to query and update customer records mid-conversation. Delivered the prototype as a live client demonstration, showcasing low-latency turn-taking and agentic tool use to inform downstream solution design.

- **Tech:** Google Cloud (Vertex AI, Gemini real-time models), LiveKit, MCP, Python, FastAPI, Azure Cosmos DB, Microsoft Dynamics 365.

La Trobe Energy Analytics Platform (LEAP)

La Trobe University, Australia (Jan 2023 – Jan 2024)

- **Data:** Energy consumption, generation, and emissions data across multiple campuses.
- **Objective:** Forecast energy usage and optimise building operations in real time.
- **Outcomes:** Developed and deployed a cloud-based analytics platform for forecasting and scenario modelling, enabling optimisation of energy consumption and supporting sustainability initiatives.
- **Tech:** Azure, Databricks, XGBoost, PyTorch, MLFlow.

Waste Management Plan Automation System

Australia (Jan 2023 – June 2024)

- **Data:** Text documents, guidelines, construction plans, and waste management reports.
- **Objective:** Automate analysis and planning of waste management processes.
- **Outcomes:** Built a system combining computer vision and ML models to identify waste types, predict generation patterns, and generate structured waste management plans.
- **Tech:** Python, OpenCV, TensorFlow, Flask, Azure.

Clinical Note-taking from Doctor-Patient Conversations

Australia (June 2023 – June 2024)

- **Data:** De-identified clinical transcripts and medical records.
- **Objective:** Extract structured clinical information from conversations to support documentation.
- **Outcomes:** Developed an NLP/LLM-based system to extract key medical attributes and generate structured notes, supported by an evaluation pipeline combining rule-based and model-based methods.
- **Tech:** Python, LangChain, Hugging Face, Azure OpenAI, DeepEval.

Emotion Analysis of Prostate Cancer Forum Discussions

Australia (July 2022 – Dec 2022)

- **Data:** Online forum discussions.
- **Objective:** Analyse emotional trends in patient discussions.
- **Outcomes:** Built a classification model to identify emotional patterns, providing insights to support patient engagement strategies.
- **Tech:** Python, TensorFlow, Scikit-learn.

References available upon request