

Senior Data Scientist with 8+ years of experience designing and deploying end-to-end AI systems across computer vision, document intelligence, and Agentic AI systems. Strong background in machine vision and data science, with proven ability to translate complex real-world problems into scalable, production-grade AI solutions. Passionate about building reliable, high-impact AI systems that drive measurable business outcomes.

SKILLS

Tools and Languages	Python, Git, SQL, Linux environments, Azure cloud services, RESTful API development (FastAPI), Distributed & production ML workflows, Strong debugging, code review, and system design skills
Machine Learning & AI	Supervised & Unsupervised Learning, Deep Learning (PyTorch, TensorFlow), Probabilistic Modeling & Uncertainty Estimation, Model Evaluation, Experiment Design, and Error Analysis
AI System Design & MLOps	End-to-end ML system design, Model deployment & monitoring, Docker & Kubernetes for ML services, Pipeline orchestration & scalable inference
Agentic AI & LLM Systems	Agentic AI architectures (Plan-Act-Observe loops), LLM tool orchestration & function calling, Multi-step task planning & execution, Prompt engineering & reasoning control strategies, LangChain and LangGraph for agent workflow orchestration
Data Engineering	Data preprocessing & feature engineering, PDF/text data pipelines, Structured & unstructured data handling, Batch and stream-oriented ML workflows
Software Engineering	Clean architecture principles, Modular code design, Test-driven development, Code maintainability, Collaborative development and code reviews

TECHNICAL EXPERIENCE

Senior Data Scientist **Dec 2023 — Present**
CUBE Content Governance Global Limited – Reg-Transform Team *Melbourne, Australia*

- Led end-to-end ML and AI development for large-scale regulatory document processing systems.
- Designed Agentic AI document intelligence architectures combining OCR, NLP, and vision models.
- Architected a high-performance title extraction pipeline optimised for latency, accuracy, and scalability.
- Built and optimised table extraction models and pipelines for reliable structured data reconstruction.
- Enhanced document QA and metadata extraction using prompt engineering and LLM-based workflows.
- Designed and deployed FastAPI-based ML microservices with automated zero-downtime updates.
- Managed production ML workflows, monitoring, and optimisation in cloud environments.

Senior Machine Learning Engineer **Apr 2023 — Nov 2023**
CUBE Content Governance Global Limited – Reg-Transform Team *Colombo, Sri Lanka (Remote)*

- Developed and optimised machine learning models for document table extraction and information retrieval.
- Built scalable pipelines to convert unstructured PDFs into structured, machine-readable data.
- Deployed serverless ML workflows on Azure, integrating high-performance REST APIs.
- Implemented monitoring, alerting, and optimisation strategies for production ML systems.
- Created robust unit and integration testing frameworks using PyTest.
- Stress-tested ML services using Locust to validate system performance under production workloads.

Lead Machine Learning Engineer **Oct 2021 — Apr 2023**
Future Insight Technologies – Machine Learning Team *Colombo, Sri Lanka*

- Led the design and development of end-to-end machine learning pipelines for document intelligence solutions.
- Researched and implemented advanced document processing and extraction techniques.
- Architected scalable, maintainable ML systems for long-term production use.
- Deployed and managed machine learning systems on managed servers with performance and reliability focus.
- Established a self-sustaining ML lifecycle with automated data ingestion, training, deployment, and monitoring.
- Provided technical leadership and guidance across the machine learning development lifecycle.

Associate Tech Lead – Cognitive Machine Learning **Sep 2021 — Nov 2021**
Zone24x7 (Pvt) Ltd *Colombo, Sri Lanka*

- Led the Computer Vision team responsible for delivering production-grade ML solutions.
- Designed end-to-end machine learning system architectures, from data ingestion to deployment.
- Provided technical leadership in model selection, system design, and performance optimisation.
- Oversaw integration of deep learning models into scalable, production environments.
- Mentored junior engineers and conducted technical reviews to ensure code quality and robustness.
- Collaborated with cross-functional teams to translate business requirements into AI-driven solutions.

Senior Machine Learning Engineer

Zone24x7 (Pvt) Ltd

2019 — 2021
Colombo, Sri Lanka

- Core member of the Machine Learning team working on multiple computer vision products.
- Developed and trained deep learning models using PyTorch and TensorFlow for real-world applications.
- Designed and implemented image classification, object detection, and image segmentation pipelines.
- Optimised model inference using OpenVINO for low-latency, real-time systems.
- Built reliable data processing and feature extraction pipelines using Python, OpenCV, Pandas, and Scikit-learn.
- Contributed to production deployment, monitoring, and maintenance of ML systems.
- Worked closely with product and engineering teams to iterate on model performance and system reliability.

Machine Learning Engineer

Zone24x7 (Pvt) Ltd

2017 — 2019
Colombo, Sri Lanka

- Developed and evaluated machine vision algorithms for industrial and enterprise use cases.
- Implemented classical and deep learning-based image processing techniques using OpenCV.
- Supported end-to-end ML workflows including data preparation, model training, evaluation, and deployment.
- Assisted in building proof-of-concept solutions that evolved into production ML products.
- Gained strong exposure to real-world constraints such as latency, reliability, and deployment limitations.

EDUCATION

Master of Science in Computer Science and Engineering

University of Moratuwa

Jan 2020 — Jun 2022
Sri Lanka

- GPA: 3.9 / 4.2
- Master's thesis focused on computer vision and applied AI techniques.

Bachelor of Science in Electronic and Telecommunication Engineering

University of Moratuwa

Jan 2013 — Apr 2017
Sri Lanka

- GPA: 3.24 / 4.2

RECENT PROJECTS

Agentic Document Insight Assistant

2025 — Ongoing

- Designing an agentic AI system for document understanding using goal-driven planning and tool orchestration.
- Implemented explicit plan-act-observe loops with reflective reasoning for structured insight extraction.
- Focused on reliability, failure handling, and governance of LLM-driven workflows.

Effective Date Extraction (RegBooks)

2025 — Ongoing

- Building an automated effective date extraction pipeline for regulatory documents from RegBooks.
- Combined text processing with LLM-assisted interpretation to handle diverse date expressions.
- Implemented asynchronous workflows for efficient processing and improved system scalability.

Document Title Extraction

2024 — 2025

- Developed a scalable document title extraction pipeline using LLM-based reasoning.
- Implemented controlled extraction workflows with validation and fallback mechanisms.

RegInsight Document Summarisation with Guardrails

2024 — 2025

- Built a production-grade document summarisation service with guardrails for safe, reliable outputs.
- Architected a multi-stage pipeline integrating Azure OpenAI, Google Gemini, and internal processing components.
- Implemented routing, validation, and re-summarisation logic to ensure concise and structured summaries.

AI-Driven Document Scraper

2023 — 2024

- Designed an intelligent document scraper combining traditional scraping, OCR, and AI-assisted understanding.
- Implemented adaptive extraction logic based on document quality and structural characteristics.

Document Structure Detection

2021 — 2023

- Built a scalable document structure detection system using computer vision and OCR.
- Implemented layout analysis and rule-based decision logic for structured output generation.

Common Engineering & Deployment Practices

- Deployed containerised ML/AI services using Kubernetes with scalable, production-ready architectures.
- Implemented CI/CD pipelines with automated testing, validation, and rollback mechanisms.
- Designed systems with monitoring, logging, and performance optimisation for large-scale workloads.
- Ensured reliability and maintainability across distributed, production ML systems.

ACTIVITIES

Toastmasters International – Competent Leader Award

2018

Inter-University Rugby Champions

2014 — 2016

All Island Chess Competition – 5th Place

2008