ARNOL P S

Senior Engineer - Data Science | AI/ML Researcher

Kerala, India | +91 9048357172 | arnolps172@gmail.com

Professional Summary

Data Scientist and AI/ML Engineer with 5+ years of experience in designing and deploying production-grade machine learning systems. Specialized in **Computer Vision**, **Deep Learning**, **Vision-Language Models**, and **Natural Language Processing**. Led patent-pending research on biometric identification combining wavelet-based image enhancement with DINOv2 feature extraction and vector similarity search. Demonstrated expertise in building end-to-end AI pipelines—from research prototyping to scalable deployment. Seeking to pursue doctoral research in **medical image analysis**, leveraging multimodal learning and transfer learning to develop clinically impactful diagnostic tools.

Education

Master of Science in Computer Science *(Specialization: Data Analytics)* Indian Institute of Information Technology and Management – Kerala (IIITM-K) *June 2017 – May 2019*

Bachelor of Science in Computer Science, Statistics and Mathematics Kristu Jayanti College, Bengaluru *June 2014 – March 2017*

Professional Experience

Senior Engineer - Data Science

Reflections Info Systems Pvt. Ltd. | *August 2025 – Present*

Leading research and development of AI-powered enterprise solutions with focus on multi-agent architectures and conversational AI systems.

- Architected **Agentic AI systems** using LangGraph for multi-agent orchestration, achieving autonomous task decomposition and execution
- Designed **hybrid RAG pipelines** with two-stage retrieval (vector search + cross-encoder reranking), achieving 94.7% query classification accuracy
- Developed **Vision-Language Model pipelines** for document understanding, supporting 13+ file formats with automatic VLM fallback mechanisms
- Built **real-time conversational AI** integrating voice synthesis, WebRTC streaming, and LLM-powered sentiment analysis with sub-256ms latency
- Implemented production **compliance verification systems** processing 149+ regulatory criteria using multi-agent LLM orchestration

Senior Data Scientist - Consultant

Digital University Kerala | *December 2024 – July 2025*

- Led development of **cow identification system** using DINOv2 (self-supervised vision transformers), achieving **92% accuracy**
- Designed **semantic document search engine** using dense retrieval and vector databases, significantly improving search relevancy

- Implemented **LLM-based document Q&A systems** using AWS Bedrock and LangChain for automated summarization and information extraction
- Guided team in developing advanced NLP and ETL solutions, delivering project milestones ahead of schedule

Senior Software Engineer - AI/ML

Techversant Infotech | *June 2024 – November 2024*

- Developed **RAG applications with memory** capabilities for enhanced contextual accuracy in conversational systems
- Built **face recognition systems** leveraging state-of-the-art deep learning models for precise identification
- Designed **AI-powered proctoring tools** using YOLO and real-time object detection for monitoring applications
- Integrated AI capabilities into existing enterprise systems with cross-functional teams

Senior Engineer - Data Science

Digital University Kerala | *September 2023 – June 2024*

- Engineered **ETL pipelines** for document extraction and indexing in Elasticsearch, optimizing data retrieval efficiency
- Spearheaded development of **ML-based document search engine** with semantic understanding capabilities
- Led team of three engineers in developing search infrastructure and data processing modules
- Created backend infrastructure for "**Fun With AI**" web application showcased at Global Science Fest Kerala, utilizing AWS for scalability

Data Analyst

Digital University Kerala | *June 2021 – August 2023*

- Optimized database performance through efficient data modeling and indexing strategies
- Created **automated data pipelines** reducing processing time and enabling faster analytics
- Developed interactive data visualizations for financial and operational reporting

Research Fellow

ICFOSS (International Centre for Free and Open Source Software) | *September 2019 – September 2020*

- Developed **Morphological Analyzer** for Malayalam language using computational linguistics approaches
- Built **sentiment analysis systems** for Malayalam, improving language analysis efficiency by 10%
- Conducted sentiment analysis research on YouTube comment data
- Managed and analyzed large-scale datasets for NLP research

Research Projects

Biometric Cattle Identification System *(Patent Pending)*

Digital University Kerala | December 2024 – July 2025

Computer vision system for individual cattle identification using muzzle (nose) patterns as biometric markers—analogous to human fingerprint recognition. Developed novel methodology combining wavelet-based image enhancement with self-supervised vision transformers.

Problem: Traditional RFID-based livestock identification is susceptible to tampering, loss, and requires time-consuming manual verification, creating challenges for traceability and fraud prevention.

Technical Approach:

- Designed **end-to-end deep learning pipeline**: Image → YOLOv11 Detection → Ridge Enhancement → DINOv2 Embedding → Vector Search
- Developed **novel wavelet-based ridge enhancement** algorithm adapting fingerprint recognition techniques for cattle muzzle patterns (BayesShrink denoising, biorthogonal wavelet contrast enhancement, morphological skeletonization)
- Leveraged **DINOv2 (ViT-L14) self-supervised learning** for 1024-dimensional feature extraction without task-specific fine-tuning
- Implemented **Pinecone vector database** for scalable real-time similarity search across registered animals

Results:

Metric	Value
Identification Accuracy	92%
Model Inference Time	49ms
End-to-End Latency	533ms
Training Dataset	420 annotated images

Technologies: PyTorch, YOLOv11, DINOv2, PyWavelets, OpenCV, Pinecone, ONNX Runtime, DVC

Contribution: Led end-to-end research and implementation including literature review (27 academic references), novel methodology development, dataset creation (300+ images), model training, and system architecture.

Industry Projects

Reflections Info Systems Pvt. Ltd. | August 2025 – Present

Vision-Based Document Extraction Pipeline

Multi-modal AI for structured data extraction from financial documents

- Created **adapter pattern architecture** supporting 13+ document formats (PDF, DOCX, XLSX, images, HTML)
- Implemented **VLM fallback strategy** with automatic escalation when OCR quality falls below threshold
- Designed **multi-provider AI routing** with priority-based failover (OpenRouter → OpenAI → Anthropic → Google)
- **Technologies**: GPT-4 Vision, Claude Vision, Gemini, Docling, PyTorch, vLLM

QFMA Compliance Verification System

AI-powered regulatory compliance using RAG and multi-agent orchestration

- Designed **multi-agent architecture** using LangGraph with specialized agents for extraction, verification, and reasoning
- Implemented **hybrid RAG pipeline** with ChromaDB vector search and confidence-based classification
- **Impact**: Automated verification of 39 articles and 149+ compliance criteria with bilingual (EN/AR) support

Production RAG System with Cross-Encoder Reranking

Semantic search with two-stage retrieval and content guardrails

- Developed **two-stage retrieval**: vector similarity search followed by cross-encoder reranking
- Built **query classifier** achieving 94.7% accuracy for routing to specialized handlers
- Implemented **6 chunking strategies** (fixed, sentence, paragraph, semantic, recursive, token-based)

Agentic AI for Roadside Assistance

Conversational AI platform with real-time analysis

- Built **multi-model LLM pipeline** (GPT-4o/Claude) for sentiment analysis, issue extraction, and intelligent routing
- Integrated **ElevenLabs conversational AI** with WebSocket-based live transcription (sub-256ms latency)

Technical Skills

Machine Learning & AI PyTorch, TensorFlow, Transformers (Hugging Face), Scikit-learn, ONNX Runtime, vLLM

Large Language Models GPT-4/4o, Claude (Anthropic), Gemini, LangChain, LangGraph, AWS Bedrock, OpenRouter

Natural Language Processing RAG Systems, Semantic Search, Sentence-Transformers, Cross-Encoders, Named Entity Recognition, Sentiment Analysis, Text Classification, Morphological Analysis

Computer Vision & Image Processing YOLO, DINOv2, Vision Transformers, OpenCV, Wavelet Analysis (PyWavelets), Morphological Operations, OCR (Tesseract), Image Enhancement

Vector Databases Qdrant, ChromaDB, Elasticsearch, FAISS

Voice & Conversational AI ElevenLabs, WebRTC, Silero VAD, Real-time Speech Processing

Backend Development FastAPI, Python (AsyncIO), WebSockets, PostgreSQL, Redis, REST APIs

Frontend Development Next.js, React, TypeScript, Tailwind CSS

Cloud & Infrastructure AWS (Bedrock, EC2, S3), Docker, NVIDIA CUDA, GPU Computing (RTX 4090)

Programming Languages Python, TypeScript, JavaScript, SQL

Certifications

- **Google Data Analytics Professional Certificate** – Google
- **Building Real-Time Video AI Applications** – NVIDIA Deep Learning Institute
- **Getting Started with Deep Learning** – NVIDIA Deep Learning Institute

Languages

- **English** – Professional proficiency
- **Malayalam** – Native
- **Swedish** – Beginner (Duolingo)

Research Interests

- Vision-Language Models for Medical Image Analysis
- Self-Supervised Learning and Transfer Learning for Computer Vision
- Biometric Pattern Recognition and Image Enhancement Techniques
- Weakly Supervised Learning and Explainable AI
- Multi-Modal Deep Learning Architectures

References

Available upon request.