Mownetharan A K S
linkedin.com/in/mowne
mowne.co.in

mownetharan@gmail.com
+91 9840895746
Applied AI Engineer,
Bengaluru

EDUCATION

Indian Institute of Technology Madras **Chennai**
B.Tech in Aerospace Engineering (Elective work in Data Science) *May 2022*

- Coursework: Data Analytics, Applied Statistics, Probabilistic Systems Analysis & Applied Probability, Linear Algebra, Differential Equations, Complex Variables, Python for Data Science

PROFESSIONAL EXPERIENCE

ITO Health **Remote**
Applied AI Engineer *(From May 2026)*

- End-to-End development of HealthEligible, an AI-powered public assistant that screens Rhode Island households for SNAP, Medicaid, RI Works (TANF) and Child Care Assistance (CCAP) in a single text or voice conversation, built with FastAPI and LangGraph.
- Built the conversation flow that runs the screening: the assistant picks out household details from what people type or say, checks them against a fixed set of rules, and decides what to ask next.
- Implemented a vectorless, PageIndex-style RAG over state regulation documents that answers policy questions with section-level citations.
- Built an AI rules extraction pipeline that turns Rhode Island's SNAP, Medicaid, RI Works and CCAP regulations into version controlled JSON eligibility rules and applicant schemas, merged into one master schema. A regulation scraper detects upstream changes and suggests rule edits for human review.
- Added voice support through the OpenAI Realtime API, supporting English and Spanish. Deployed it on AWS Elastic Beanstalk with Langfuse tracing, emailed summaries (SES) and S3 share links.

Genpact **Bengaluru**
Assistant Manager – AI/ML Engineer *(Dec 2024 – April 2026)*

- Developed a Mortgage AI-based validation tool which is powered by LLMs and aims to automate validation of Full Mortgage Applications with supporting documents hosted in AWS Lambda.
- Developed a Data Extraction/ Query answering application for credit and loan agreements utilizing Llama3 in a RAG pipeline with deployment in AWS using S3 and Lambda services
- Boosted the capability of Dashboards using Explainable AI techniques which utilizes LLMs to explain the visuals of a dashboard with insights, trends and recommendations along with a Chabot which can be used to create custom plots

IQVIA **Bengaluru**
Analyst (Advanced Analytics) *(July 2022 – Dec 2024)*

- Implemented Advanced NLP techniques to perform Theme Generation, Sentiment Analysis, text-based Clustering and classification by utilizing Generative AI, Vector Embeddings & Semantic Search to provide actionable insights from large-scale survey feedback data with 20,000+ data points.
- Conducted Robustness Testing on Price Elasticity CI/CD pipelines in Cloudera CDSW along with comparative inflation analysis to evaluate pricing strategies on pharmaceuticals, ensuring strategic pricing decisions were data-driven.
- Developed an AI-ML based Survey Automation tool that utilizes parameter-tuned Boosting ML algorithms which results in a reduction in time taken to reach out and complete the surveys by 8%
- Enhanced various projects such as Market Segmentation, Persona Enrichment to improve target identification & market strategies with modelling and extrapolation techniques

PERSONAL PROJECTS

ML-Agent

- Built an AutoML tool for tabular data where an LLM suggests experiments, but regular code checks each plan, runs cross-validation and picks the best scikit-learn pipeline.
- The LLM only sees column names, summary statistics and earlier results, never the actual rows. It works with the OpenAI API or a Codex or Claude Code login, and each run saves the model, a leaderboard and a short report.

PyInterpret

- Wrote and published an open-source Python library that brings SHAP, LIME, permutation importance and partial dependence together under one API for scikit-learn models.
- Every method returns the same kind of result, switching between explanations and comparing them becomes easy without rewriting code.

ERP Assistant

- Built a chat assistant that adds and removes users, customers and offices from plain-English requests, using a LangGraph workflow to work out what the user wants.
- It reads CSV, Excel, PDF, text and JSON uploads and uses an LLM to fit the records to the right format. It can also turn a written description of who reports to whom into an interactive org chart.

SKILLS & TOOLS

LLMs: GPT models, Anthropic Claude, Gemini, AWS Bedrock, Openrouter
Agents: LangGraph, LangChain, tool calling, structured outputs
RAG: vector based AWS OpenSearch, vectorless PageIndex
Voice AI: speech-to-text and text-to-speech with ElevenLabs and AWS Polly/Transcribe, OpenAI Realtime
LLM evaluation and observability: Langfuse, Giskard, custom eval harnesses

Python: FastAPI, Pandas, scikit-learn, Pydantic, pytest
Machine learning: boosting models, AutoML, model interpretability (SHAP, LIME)
NLP: text classification, sentiment analysis, embeddings, semantic search
Cloud and deployment: AWS (Lambda, S3, Elastic Beanstalk, SES), Docker, CI/CD with GitHub Actions
Data: SQL, PySpark, EDA, Streamlit