

Praneeth Dathu

+31 6 16911077 | praneethdathu46@gmail.com | [linkedin.com/in/praneethdathu46](https://www.linkedin.com/in/praneethdathu46) | github.com/praneeth-496 | [portfolio](#) | Amsterdam, Netherlands

PROFESSIONAL SUMMARY

MSc Artificial Intelligence (Leiden University) with hands-on machine learning research and applied LLM engineering. Built and benchmarked 20+ deep learning models on HPC GPU clusters, outperforming published state-of-the-art on a public BLE localization dataset (1.269 m MAE vs 1.973 m baseline; 87.3% live room accuracy). Side-projects span LLM integration with RAG, MCP-based tool-use, and cost-optimized inference (Claude Haiku, Groq). Strong in PyTorch, Transformers, generative models, and reproducible experimentation. Targeting AI / Machine Learning Engineer roles in the Netherlands. Available full-time from August 2026.

CORE COMPETENCIES

Machine Learning

Deep Learning

PyTorch

Transformers

LLM Integration

RAG / Vector Retrieval

MCP (Model Context Protocol)

Generative Models (VAE / GAN)

Reinforcement Learning

Recommender Systems

Signal Processing

Feature Engineering

Python

SQL

HPC / SLURM

Linux / Git

Model Benchmarking

Data Pipelines

WORK EXPERIENCE

Almende B.V. / Crownstone Feb 2026 - Jul 2026

Machine Learning Research Intern (MSc Thesis) - Rotterdam, Netherlands

- Designed and shipped a full indoor localization pipeline using BLE RSSI signals from 15 Crownstone anchors across 4 room environments, owning the work end to end (data collection, preprocessing, modeling, evaluation, reporting).
- Benchmarked 20+ machine learning models - CatBoost, TCN, Transformer, CNN-LSTM, KNN, MLP, PINN - on the ALICE HPC GPU cluster using SLURM job scheduling and reproducible experiment configs.
- Achieved **87.3% live room-classification accuracy** with CatBoost and **99.81% offline accuracy** with TCN; outperformed published state-of-the-art on a public BLE benchmark - KNN reached 1.269 m MAE vs the best published 1.973 m, with 6 of 11 models beating the prior baseline.
- Built a production-ready Python pipeline with automated preprocessing, Hampel filtering, z-score normalization, and a 75-dimensional engineered feature set; integrated InfluxDB / MQTT for real-time data flow.
- Communicated results to supervisors via structured bi-weekly reports, decks, and LaTeX-documented methodology.

Digital Ellanky Aug 2023 - Jul 2024

AI Researcher - Hyderabad, India

- Analyzed large-scale customer and campaign datasets with Python and SQL to extract KPIs and drive data-informed marketing decisions.
- Built predictive and clustering models for audience targeting; developed automated preprocessing and metric-reporting pipelines that reduced manual work for the analytics team.
- Designed performance dashboards tracking ROI, customer acquisition cost, and behavioral trends across multiple campaigns.
- Stack: Python, Pandas, NumPy, scikit-learn, SQL, Google Analytics.

SELECTED PROJECTS

Almende Internal Tooling Dashboard - UI & Chatbot Integration

Developed an interactive internal dashboard for the Almende team. Built the UI in Lovable and integrated an in-dashboard chatbot via the Groq API for low-latency LLM responses. Wired up Model Context Protocol (MCP) server connections so the assistant can act over internal tools and data. System concept led by a colleague; UI and chatbot integration owned by me.

Vault AI - Local-First Personal Finance Assistant

Built a privacy-first desktop app that ingests monthly bank transactions, classifies spending categories, and uses an LLM to surface anomalies, recurring leaks, and budgeting recommendations in plain language. Transaction data lives in a local SQLite store; Google Drive integration provides user-owned backup and cross-device sync. Stack: Python, SQLite, LLM API, Google Drive API.

Game AI - Pacman Capture the Flag with MCTS Rank 4 / 37 teams

Built competitive multi-agent game-AI bots using Monte Carlo Tree Search with RAVE; tuned hyperparameters via 100+ automated match simulations to reach a top-4 tournament finish.

Persona-Conditioned Conversational Agent with Long-Term Memory

Built a multi-turn conversational LLM agent with custom persona prompts and retrieval-augmented memory for coherence across long sessions. Deployed on a Hetzner VPS for always-on availability independent of the local machine. Optimized inference cost by routing routine turns to Claude Haiku and applying prompt-caching strategies. Stack: OpenClaw agent platform, Claude Haiku API, RAG / vector retrieval, Hetzner cloud.

Reinforcement Learning - DQN, A2C, PPO on CartPole

Implemented and compared DQN, A2C, and PPO in PyTorch; built an evaluation pipeline tracking convergence speed, cumulative reward, and training stability with structured logging and visualizations.

Neural Collaborative Filtering - MovieLens 1M

Implemented a Neural Collaborative Filtering recommender in PyTorch on MovieLens 1M; achieved NDCG@10 = 0.2117 on a held-out test set, trained on an RTX 3090.

EDUCATION

MSc Artificial Intelligence - Leiden University, Netherlands

Sept 2023 - Aug 2026

GPA 7.5/10. Specialization in Machine Learning, Deep Learning, and Data Analysis. Coursework: ML, Deep Learning, Reinforcement Learning, Recommender Systems, Data Mining, Natural Language Processing.

B.Tech Computer Science Engineering - Malla Reddy Engineering College, India

Jan 2019 - Jan 2023

GPA 7.18/10. Coursework: Data Structures & Algorithms, Database Systems, Operating Systems, Web Technologies.

TECHNICAL SKILLS

Machine Learning & AI: PyTorch, scikit-learn, CatBoost, XGBoost, Transformers, CNNs, RNNs, TCN, PINN, VAEs, GANs, WGAN-GP, Reinforcement Learning (DQN, A2C, PPO), Recommender Systems, Monte Carlo Tree Search.

LLM & GenAI: LLM API integration (Anthropic Claude, Groq), Retrieval-Augmented Generation (RAG), prompt engineering & prompt caching, Model Context Protocol (MCP), persona / agent design, cost-aware inference routing.

Programming & Data: Python (Pandas, NumPy, Matplotlib, Seaborn), SQL, SQLite, Jupyter, Bash.

Infrastructure & MLOps: SLURM (HPC scheduling), ALICE GPU cluster, Hetzner VPS deployment, Linux, Git, SSH, VS Code, reproducible experiment pipelines.

Data Engineering: InfluxDB, MQTT, signal processing (Hampel filtering, z-score normalization), feature engineering, Google Drive API integration.

Reporting: LaTeX / Overleaf, structured technical writing, stakeholder presentations.

Languages: English (fluent), Hindi (fluent), Telugu (native), Dutch (A2).