

Pranav R Bhat

Location: Mangalore / Bengaluru pranavrbm@gmail.com +91 8073714200 linkedin.com/in/pranavrbm
github.com/pranavrbm

Summary

AI and Data Science Specialist with hands-on experience building production-grade machine learning pipelines, RAG systems, LLM-powered automation, and full-stack applications. Published IEEE first-author researcher in AI-generated text detection. Experienced across the full AI stack — from local LLM inference and vector databases to cloud-native deployment — with a track record of shipping real products in agriculture, healthcare, and resource management.

Summary of Qualifications

- Practical experience in AI/ML model development, RAG pipelines, LLM orchestration (Ollama, LangChain, ChromaDB), and production deployment.
- Led and published peer-reviewed research in AI-generated text detection (IEEE DISCOVER 2024, first author).
- Strong command of Python, TensorFlow, Scikit-learn, XGBoost, CatBoost, SHAP, and Pandas for real-world predictive systems.
- Built end-to-end automation systems using Playwright, Google Business API, and local LLM inference at production scale.
- Capable of containerizing and serving ML models with Docker, FastAPI, and robust microservices architectures.
- Effective cross-functional collaborator — trained teammates on Claude Code, Git, and AI-assisted workflows.

Education

Bachelor of Technology (B.Tech) in AI and Data Science 2022 – 2026
NMAM Institute of Technology, Nitte, Karkala
Expected Graduation: 2026
Relevant Coursework: Machine Learning, Deep Learning, NLP, Data Structures, Algorithms, Big Data Analytics, Cloud Computing

Professional Experience

AI & Web Development Intern Mar 2026 – May 2026
Organic Mandya Bengaluru (Remote)

- **Google Reviews AI Automation:** Built an end-to-end review management system using Playwright and Google Business API. Deployed Llama 3.1:8b via Ollama with ChromaDB-backed RAG (crawled full website, structured into vector store) for brand-authentic AI reply generation across all store locations.
- Integrated Office365 SMTP email escalation for flagged reviews; deployed to Linux VPS (srv1322211) with a daily 8:30 PM IST cron job; configured Playwright browser auth via SSH X forwarding.
- **Know Your Food Portal:** Designed, built, and deployed know.organicmandya.com using Astro v6 and Tailwind CSS v4 — 15+ content categories, search, SEO, dark mode, UTM tracking, sitemap, and full accessibility pass.
- **Data Engineering:** Studied BigQuery schemas and ETL flows; completely migrated the organization's data pipeline (BigQuery, Shopify API, Odoo, all cron jobs) from an external server to the Organic Mandya Linux VPS; validated pipeline and set up Git version control.
- **Mobile App Design (Suggi):** Designed and prototyped a farmer/consumer mobile app using Figma MCP and React artifacts — Trust Score, crop planning calendar, plant scan, farm gate tourism, and a voice-first AI Guru assistant.

- Trained and onboarded two team members on Claude Code, Git workflows, prompting techniques, and shared brand guidelines.

Tech Stack: Python, Playwright, Ollama (Llama 3.1:8b), ChromaDB, RAG, Google Business API, Astro v6, Tailwind CSS v4, BigQuery, Shopify API, Odoo, Netlify, Linux Cron, Figma MCP, React

Data Science Intern

Jul 2025 – Sep 2025

Personify (via 1stop)

Remote

- Developed ML classification models for **Hate Speech Detection** using NLP preprocessing (tokenization, TF-IDF, embeddings), SVM, and Logistic Regression on real-world social media text.
- Built a **clinical ML model** predicting heart failure patient mortality using XGBoost, CatBoost, Random Forest, SHAP explainability, SMOTE class-imbalance handling, and rigorous cross-validation.
- Conducted deep-dive **EDA on the Home Credit Default Risk dataset** — uncovering statistical anomalies and key risk features, translating findings into actionable business insights.

Software Developer Intern (Backend & Image Processing)

Jun 2024 – Aug 2024

Krishitantra

Onsite

- Engineered backend compute pipelines and Python algorithms for agricultural image processing to estimate crop nitrogen levels.
- Built and deployed a scalable crop-analysis web app integrating a React frontend with a Django REST API backend for real-time model inference.
- Automated report generation for precision agriculture, improving usability and efficiency for end-users.

Projects

Deceptive AI Review Detection System ↗

2024

- Designed an end-to-end NLP pipeline to classify deceptive AI-generated reviews in large unstructured text corpora, achieving **92% accuracy** using deep learning classifiers and TF-IDF feature engineering.
- Applied advanced preprocessing, data augmentation, and model evaluation to maximize generalization on imbalanced datasets.
- **Published as first author at IEEE DISCOVER 2024** — cited as a benchmark approach for LLM-generated fake review detection. [\[Paper\]](#) ↗

AI-Powered Predictive System for Water Crisis Management ↗

2025 – Present

- Built an urban water intelligence platform predicting shortage risk using climate, reservoir, groundwater, and demand datasets for Bengaluru.
- Implemented Python ETL and ML pipelines; developed an RBAC-enabled Next.js (T3) dashboard with Prisma and PostgreSQL for real-time prediction visualization.
- Containerized full ML ecosystem with **Docker Compose**; integrated secure FastAPI prediction service for automated model serving.

AI Health Microservices via Generative AI ↗

2024

- Integrated **Google Gemini Generative AI** with multi-turn prompt engineering to generate personalized health recommendations from user data.
- Designed a scalable microservices backend with Spring Boot, Docker, and AWS; secure auth via Keycloak and async messaging via RabbitMQ; React/Redux frontend.

LeafColor – Computer Vision Crop Analyzer ↗

2024

- Developed a Python/Cython application using image segmentation to analyze agricultural leaf images and estimate nitrogen content with 92% accuracy vs. manual expert ratings.
- Designed a user-friendly interface for image upload and automated scientific report generation.

- Built a scalable adoption platform with Next.js frontend, Prisma ORM, PostgreSQL, and NextAuth role-based authentication.

Publications

Identifying Deceptive AI Reviews: A Machine Learning Approach ↗

October 2024

Pranav R. Bhat, Nidhish Shettigar, Venugopala P. S., Amrith R. Naik, Vaishnavi

Conference: 2024 IEEE DISCOVER Conference on Distributed Computing, VLSI, Electrical Circuits, and Robotics

Location: St. Joseph Engineering College, Mangaluru, India

Conference Record: #62353

Link: [Access Publication on IEEE Xplore](#) ↗

Certifications

Machine Learning Specialization

Sep 2025

Stanford Online & DeepLearning.AI (Coursera)

Supervised/Unsupervised Learning, Advanced Algorithms, Reinforcement Learning

Verify: [B207U213G589](#) ↗

Professional Affiliations

- Student Member, IEEE Computer Society
- Member, ACM India Chapter

Technical Skills

AI / ML & Deep Learning:	TensorFlow, PyTorch, Scikit-learn, HuggingFace Transformers, Fine-tuning, Prompt Engineering, RAG, LangChain, ChromaDB, Ollama, TF-IDF, CNNs, RNNs, SHAP, XGBoost, CatBoost
Generative AI & LLMs:	Llama 3.1, Google Gemini API, OpenAI-compatible APIs, vector embeddings, semantic search, local LLM inference, multi-turn context management
Automation & Scraping:	Playwright, Google Business API, Linux Cron, SMTP email automation, web scraping pipelines
Data Science & Engineering:	Pandas, NumPy, EDA, ETL Pipelines, BigQuery, Shopify API, Odoo, feature engineering, model evaluation
Programming Languages:	Python, JavaScript, SQL, C++, Dart
Backend & APIs:	FastAPI, Django, Flask, Spring Boot, REST APIs, Node.js
Frontend & Frameworks:	React, Next.js, Astro v6, Tailwind CSS v4, HTML/CSS
Databases:	PostgreSQL, MySQL, MongoDB, Prisma, ChromaDB
DevOps & Infrastructure:	Docker, Docker Compose, Git, GitHub Actions, Linux VPS, AWS, GCP, Firebase, Netlify, Vercel
Design & Tooling:	Figma MCP, Claude Code, GitHub Copilot

Languages

- English (Fluent)
- Kannada (Native)
- Hindi (Intermediate)
- Japanese (Learning)