Pawan Parida

Final-Year Data Analytics Student | AI Engineering Intern | Building Production-Ready ML Systems

13zero7two005@gmail.com | +91 8700579954 | New Delhi, Delhi, India | GitHub

Professional Summary

Final-year Data Analytics student who reduced AI inference latency by 18x (450ms→25ms) through algorithmic optimization, achieving 616 FPS on resource-constrained systems. Built 8+ end-to-end ML applications: real-time interview coach, medical image classifier (91.3% accuracy), resume screener (98.45% across 25+ domains), and RAG systems (88.3% retrieval accuracy, <2% hallucination rate). Seeking AI Engineering internship to apply model optimization, deployment, and performance tuning to production challenges.

Education

DSEU Rajokri Campus - B.Sc in Information Technology - Data Analytics

2023-06 – 2025-06

Undergraduate program focusing on data analysis, statistics, machine learning, predictive modeling, and data visualization. Completed capstone project in real-time AI systems.

Work Experience

Project Intern - Edunet Foundation

2025-04 – 2025-05 | New Delhi, India

Built agriculture recommendation system achieving 99.77% prediction accuracy across 2,200 crop samples and 100% accuracy on fertilizer recommendations. Processed soil parameter datasets (NPK ratios, pH, rainfall) and deployed decision tree models via Streamlit interface for farmer accessibility. Reduced recommendation time from manual research (hours) to instant predictions, enabling data-driven farming decisions.

Projects

AIVOX - AI Interview Coach | GitHub

Real-time interview platform with speech recognition, emotion analysis, and AI-driven feedback | 2025-09

Built real-time interview coaching platform processing facial emotion analysis and speech transcription simultaneously. Reduced inference latency 18x (450ms→25ms) by replacing CNN with MediaPipe geometric landmark detection, maintaining 30+ FPS face tracking with ~15% CPU usage. Integrated Groq Whisper API for speech-to-text and Gemini

API for context-aware interview responses. Deployed responsive frontend (Tailwind CSS) enabling seamless user interaction. Tech: FastAPI, MediaPipe, Groq Whisper API, Gemini API, Python.

News Article Research Chatbot | GitHub

RAG-based chatbot for multi-document news analysis with citation tracking | 2025-06

Built RAG-based news research agent synthesizing insights from 50+ concurrent URLs with 88.3% retrieval accuracy (NDCG@5) and <2.1s query latency using Gemini 2.0 Flash. Reduced hallucination rates to <2% through context-grounding and source citation pipelines. Implemented scalable LangChain workflow automating content extraction and semantic summarization. Enables researchers to cross-reference multiple sources through natural language queries with verifiable citations. Tech: LangChain, ChromaDB, Gemini 2.0 Flash, Streamlit.

AI Grammar Tutor | GitHub

Grammar tutoring application using LLM with prompt optimization | 2025-07

Developed GenAI-powered grammar assistant achieving 92.3% detection accuracy for syntax and tense errors across 10 error categories (subject-verb agreement, tense, pronouns). Fine-tuned prompt engineering (Temperature 0.2) reducing model hallucinations for consistent pedagogical feedback. Built high-concurrency FastAPI backend handling multiple requests with <1.8s average response time. Tech: FastAPI, LangChain, Gemini API, JavaScript.

MySQL Natural Language Query Engine | GitHub

Convert natural language to SQL queries for e-commerce databases | 2025-06

Implemented Natural Language to SQL (NL2SQL) engine achieving 91.2% query generation accuracy on e-commerce database schema. Designed few-shot semantic prompting strategy using ChromaDB enabling complex joins and aggregations. Reduced data retrieval time from minutes (manual SQL) to ~0.8s for non-technical stakeholders. Eliminates SQL knowledge barrier through plain English query interface. Tech: LangChain, MySQL, ChromaDB, Streamlit.

Resume Classification System | GitHub

Automated resume screening across 25+ job domains | 2025-05

Developed automated resume classification system achieving 98.45% accuracy across 25+ job domains using OneVsRest classifier with TF-IDF vectorization. Outperformed baseline models by 12% through feature engineering and hyperparameter tuning. Streamlined recruitment funnel by automating candidate categorization, reducing manual screening time from hours to seconds. Handles multi-label classification for candidates with diverse skill sets. Tech: Scikit-learn, TF-IDF, Pandas, Streamlit.

Skills

Languages: Python, JavaScript, SQL, HTML/CSS

Machine Learning & AI: CNN & Transfer Learning, NLP (BERT, TF-IDF, Word Embeddings), Computer Vision (MediaPipe, OpenCV), RAG Systems, LLM Integration (Gemini, Whisper), Few-Shot Prompting

Data Science: Pandas, NumPy, Scikit-learn, Feature Engineering, Data Preprocessing, Model Evaluation

Frameworks & Tools: TensorFlow/Keras, FastAPI, Streamlit, LangChain, Google Gemini API, Groq Whisper API

Databases & Vector Search: MySQL, ChromaDB, Vector Databases, Schema Design

Frontend & Deployment: React.js, Tailwind CSS, REST API Design, Responsive Web Design

Core Concepts: RAG Systems, Prompt Engineering, Object-Oriented Programming, Git/GitHub

Open Source Contributions

Refactored KL Divergence Function - Optax - google-deepmind/optax | GitHub PR #1513
2025-11

Successfully contributed documentation clarification to Google DeepMind's Optax library. Updated docstrings to distinguish between standard KL divergence and Generalized KL Divergence for unnormalized distributions. Collaborated with maintainers through review process, demonstrating professional OSS practices including clear documentation, API design understanding, and responsiveness to feedback. Merged into production.

Certifications

Data Science Fundamentals - TestDome (2025-10) | Credential

Big Data Foundations - Level 1 - IBM (2025-08) | Credential

Quantium - Software Engineering Job Simulation - Forage (2025-09) | Credential

AI and Data Scientist - OneRoadmap (2025-08) | Credential

Languages

English (Fluent), Hindi (Native), Odia (Native)