

YASH VASUDEVA

Machine Learning Engineer

☎ +91 9318378900

✉ vasudevyash@gmail.com

🌐 yashvasudeva1

🌐 Yash Vasudeva

Education

B.Tech (Information Tech), Guru Gobind Singh Indraprastha University

8.8/10 2024 – 2028

Core Skills

Languages: Python, SQL, C++

Deep Learning: Tensorflow, Transformers, NLP

Data: Pandas, Polars, NumPy, MySQL, PostgreSQL

Monitoring: Prometheus, Grafana

Visualization: Matplotlib, Seaborn

ML: Scikit-Learn, PyTorch, XGBoost, LightGBM

GenAI: Hugging Face, LangChain, LangGraph

MLOps: FastAPI, Docker, MLflow, REST APIs, CI/CD

Tools: Git, GitHub, Jupyter Notebook

Experience

Business Analyst Intern – Machine Learning | Tata Power Delhi Distribution Limited

May 2026 – Jul 2026

Engineered an end-to-end portal for Non-Technical Loss (NTL) identification, analyzing electricity theft, meter tampering, and billing anomalies across 10,000 consumers using 12 meter-signal features.

- Benchmarked 15 classification algorithms and selected Logistic Regression for deployment, achieving 73.4% precision (target >65%), 97.2% recall, and 0.9696 PR-AUC
- Architected a FastAPI application spanning Data Explorer, Model Intelligence, and Live Inference, producing fraud-probability-ranked inspection lists while reconciling predictions with physical audit ground truth to reduce false-positive field visits costing Rs. 1,500–5,000 each

Technical Blog Writer | Namespace Community

Mar 2026 – Present

- Published 10+ tutorials for a 50,000+ member developer community, simplifying transformers, evaluation techniques, and applied artificial intelligence through practical explanations
- Audited contributor submissions for technical accuracy and established repository conventions covering READMEs, directory structures, and documentation, streamlining onboarding for new contributors

Projects

ResearchLens – Evidence-Grounded Research & Documentation Assistant

[\[GitHub\]](#)

- Created a RAG assistant for research papers and technical framework documentation, enabling developers to query complex references conversationally and reducing reliance on manual documentation search
- Combined dense retrieval, BM25, Reciprocal Rank Fusion, and cross-encoder reranking with context filtering and evidence gating to improve relevance and suppress unsupported responses
- Added claim-level citation verification and benchmarked retrieval, groundedness, and citation quality, achieving 1.000 Recall@5, 0.942 MRR, 94.2% groundedness, and 100% citation accuracy

ORCHESTRA – Multi-Agent Hackathon Evaluation Platform

[\[GitHub\]](#)

- Orchestrated five specialized GPT-4o judges in parallel, followed by bias auditing and chief-judge aggregation to score hackathon submissions and assign confidence tiers
- Implemented asynchronous evaluation workflows with Node.js, Express, PostgreSQL, Zod, and REST APIs, including GitHub/PPTX parsing, leaderboard generation, and automated feedback reports

Real-Time YouTube Sentiment Analyzer

[\[GitHub\]](#)

- Applied TF-IDF with LightGBM to 49,985 labeled YouTube comments, achieving 79.96% accuracy and 0.80 Macro F1 for real-time sentiment classification
- Integrated YouTube Data API ingestion with a Manifest V3 Chrome Extension and FastAPI backend, adding resilient pagination, trend visualization, and Dockerized inference reaching ~8,869 comments/sec

Achievements

MLH HackDay – Gemini Track Winner

- Won 1st place among 85+ teams and 250+ participants by developing an AI product with Google Gemini APIs within an 8-hour hackathon

National Hackathon Finalist – Top 5

- Placed among the top 5 teams across national hackathons hosted by The Economic Times, IIIT-Delhi, CGC University, and IIM Indore

Smt. Savitri Gupta Memorial Award – S.D. Public School

- Received the award for outstanding contributions to national and international school projects, co-curricular excellence, and leadership as Head Boy and Student Council Leader (2023–24)