

DHRUVARYAN CHUGH

AI/ML Engineer | Software Engineer | Data Engineer

New Delhi, India | +91 8178018929 | chughdhruvaryan@gmail.com

linkedin.com/in/dhruvaryan-chugh | github.com/Dhruvaryan-10 | portfoliodac.vercel.app

SUMMARY

Computer Science Engineering student (GPA 9.34/10) specializing in Artificial Intelligence and Machine Learning. Builds multi-agent LLM systems, RAG pipelines with measured retrieval evaluation, and applied ML models on real datasets in Python, LangChain, PyTorch, Scikit-learn, pandas, and NumPy, deployed over FastAPI and PostgreSQL.

EDUCATION

SRM Institute of Science and Technology

Chennai, Tamil Nadu

Bachelor of Technology, Computer Science and Engineering

Expected May 2027

GPA: 9.34 / 10.0 | Specialization: Artificial Intelligence, Machine Learning, Data Science

Coursework: Machine Learning, Statistics, Probability, Data Analytics, Data Structures, Database Management Systems

PROJECTS

SDLC Automation Platform

Python, LangChain, Google Gemini API, Pydantic, Streamlit

github.com/Dhruvaryan-10

- Designed a six-agent LLM workflow that plans and executes multi-step generation of requirement, design, and QA artifacts from a Statement of Work, using Google Gemini and LangChain with one externalized prompt template per agent.
- Implemented append-only artifact versioning with draft, final, and locked states and cross-artifact staleness tracking, so downstream documents are flagged as outdated rather than silently regenerated.
- Deployed a live instance on Streamlit Cloud with a seven-step guided workspace, DOCX export, and natural-language refinement that patches only the affected sections.

SmartCA – AI Personal Finance and Tax Workspace

Python, FastAPI, PostgreSQL, Drizzle ORM, React.js, RAG, LLM Tool Calling

github.com/Dhruvaryan-10

- Built a RAG tax-law retrieval system over a chunked corpus of authoritative Indian tax sources, scoped by assessment year, statutory section, and authority tier, returning a typed refusal when evidence is absent.
- Measured retrieval quality with an evaluation harness reporting recall@1/@3/@5, MRR, evidence support, distractor intrusion, and refusal accuracy across gold and provisional test cases.
- Architected an LLM orchestrator over six server-executed tools (tax-law search, transaction query, financial summary, tax calculation, regime comparison, simulation) backed by a deterministic versioned tax engine, so computed figures never originate from the model.
- Hardened the model boundary with tool allowlists, server-side user scoping, and an answer-validation layer that withholds unsupported figures, verified against a synthetic-model test environment.

Multi-Omics Drug Sensitivity Prediction

Python, LightGBM, PyTorch, Scikit-learn, pandas, NumPy, RDKit, SHAP

github.com/Dhruvaryan-10

- Achieved R-squared of 0.51 and RMSE of 1.96 on IC50 drug-sensitivity regression using LightGBM, outperforming Random Forest, Neural Network, and TabNet baselines.
- Built a feature pipeline over the LINC1275 and GDSC datasets covering data cleaning, exploratory data analysis, autoencoder dimensionality reduction, feature engineering, hyperparameter tuning, and SHAP attribution.

EXPERIENCE

GIZ (Deutsche Gesellschaft für Internationale Zusammenarbeit), India

New Delhi, India

Data Analyst Intern

May 2025 – July 2025

- Reduced report generation time by 65%, cutting turnaround from three days to same-day, by building automated reporting workflows in Microsoft Lists that replaced manual status collection.
- Performed data extraction, cleaning, validation, and reconciliation across datasets from 4 regional program teams, and built Tableau dashboards and PowerPoint summaries for non-technical stakeholders.

BrainMint Development

Chennai, Tamil Nadu

Frontend Development Engineering Intern

Feb 2025 – July 2025

- Built React.js data visualization interfaces surfacing mission telemetry into a single analytics dashboard, cutting page load times 30% through rendering and asset-delivery optimization.
- Integrated telemetry pipelines alongside backend engineers, standardizing API contracts to eliminate cross-service data mismatches.

SKILLS

Languages: Python, SQL, JavaScript, Java, C++

GenAI and LLMs: LangChain, Google Gemini API, OpenAI API, Retrieval-Augmented Generation (RAG), Agentic AI, Multi-Agent Systems, Prompt Engineering, Tool/Function Calling, LLM Evaluation, Retrieval Evaluation, Guardrails, Structured Outputs, Pydantic

Machine Learning: PyTorch, Scikit-learn, LightGBM, Autoencoders, SHAP, Explainable AI, Feature Engineering, Model Evaluation, Regression, Statistical Modeling

Data: pandas, NumPy, PostgreSQL, Drizzle ORM, Exploratory Data Analysis, Data Cleaning, Data Validation, Document Extraction, Data Visualization, Tableau, Power BI, Jupyter Notebook

Engineering: FastAPI, REST API Design, JWT Authentication, Backend Architecture, System Design, Git, GitHub, Streamlit, React.js, AWS (certified)

CERTIFICATIONS

AWS Certified Machine Learning Engineer – Associate (2026) | AWS Certified AI Practitioner (2026) | AWS Certified Cloud Practitioner (2026) | GitHub Copilot Certification, Microsoft Learn (2026)