

KUSHAGRA GUPTA

+91 9269972395 ✉ kkgupta.work@gmail.com [LinkedIn](#) [Github](#)

Education

JK Lakshmipat University, Jaipur

Expected 2026

B.Tech in Computer Science & Engineering

7.6/10 CGPA

- Semester Exchange Student — **IIIT, Hyderabad** — Monsoon 2025 — 7/10 CGPA
- Semester Exchange Student — **IIT, Gandhinagar** — Dec 2023

Technical Skills

Languages: Python 3.10+, JavaScript (ES6+), C++17, SQL (PostgreSQL, MySQL), Bash

Machine Learning: Supervised/Unsupervised Learning, SVM, Random Forests, Naive Bayes, PCA, Scikit-learn

Deep Learning: PyTorch 2.x, CNN (ResNet, YOLO), RNN (LSTM, GRU), Transformers (Self-Attention), ViT, LoRA

GenAI & Agents: RAG, LangChain, CrewAI, Prompt Engineering, RLHF, Vector DBs (FAISS), Multimodal Pre-training

Cloud & Tools: Google Cloud Platform (Vertex AI, Cloud Run, CloudSQL), Docker, MS Excel (VLOOKUP, Pivot), Git

Work Experience

Full Stack Development Intern

May 2025 – July 2025

Dobby Ads

Remote

- Automated 30+ Cypress test cases, reducing post-release bugs by 25% and cutting regression time by 40%.
- Architected and deployed two MERN-stack web applications using MongoDB, optimizing scalable architecture.
- Integrated Prometheus and Grafana dashboards for real-time API monitoring, improving resolution time by 20%.

Android Development Intern

May 2024 – July 2024

Vedic Bodhi Pvt. Ltd.

Remote

- Optimized app performance by refining API call structures and implementing responsive UI designs in Android.
- Engineered secure user authentication using Firebase and integrated REST APIs for real-time data sync.

Projects

Subject Matter Expert (SME) AI Agent

Nov 2026

Python 3.12, FastAPI, RAG, LangChain, Elasticsearch, FAISS

- Developed a production-grade RAG system using FastAPI, integrating FAISS for dense retrieval and BM25 for sparse search.
- Streamlined document ingestion for 50k+ chunks using hierarchical chunking, cutting irrelevant retrievals by 40%.
- Architected an agentic workflow engine for multi-step queries, including automated PDF generation.

Legal Redline Sandbox (Hackathon - Global top 80 teams)

Nov 2026

Python, Google Gemini API, GCP Vertex AI, CloudSQL

- Spearheaded GCP infrastructure and data persistence layer, designing relational schemas to manage chat sessions.
- Implemented an asynchronous notification system for real-time status updates during complex contract analysis.

Multilingual Language Model (English–Hindi–Bengali)

Sept 2026

Python, PyTorch 2.1, Transformers, LoRA, SentencePiece

- Architected an 18.5M parameter Transformer using PyTorch, implementing LoRA to optimize Gemma-270M.
- Implemented SentencePiece tokenizer (50k vocab), achieving a <0.0002% OOV rate for multilingual text.

Automatic License Plate Recognition (ALPR)

Sept 2025

PyTorch, OpenCV, EasyOCR, CNN

- Engineered a 5-layer CNN for bounding-box regression, achieving 71.48% mean IoU on 9,600 images.
- Increased inference stability through Test-Time Augmentation (TTA), integrating the model into a deployable pipeline.

Relevant Coursework

Machine Learning & AI: Markov Decision Processes, Alpha-Beta Pruning, GDA, SVM Kernels, Constraint Satisfaction.

Deep Learning & NLP: Backpropagation (Adam/RMSProp), Batch Norm, Seq2Seq, Word2Vec, Attention, RLHF.

Leadership and Extracurriculars

- **Google Cloud GenAI Skill Badges:** Validated expertise in Vertex AI, Gemini API, RAG, and Prompt Design (2025).
- **IBM AI Fundamentals:** Certified in ML principles and real-world neural network applications (2025).