

Yash Bansal

yash_b@bt.iitr.ac.in | GitHub | LinkedIn

Summary

AI Engineer (-side ML researcher) building production-grade **agentic LLM systems**, RAG pipelines, and multi-modal systems with deep learning models across **computer vision** and **NLP**; published at **AAAI & ICVGIP**.

Education

Indian Institute of Technology, Roorkee	<i>August 2023 – Present</i>
B.Tech. in Biosciences and Bioengineering	CGPA: 8.27
Heera Lal Public School, Delhi (CBSE)	<i>2022</i>
Senior Secondary (Class XII)	93.20%
Delhi Public School, Sector 30, Noida (CBSE)	<i>2020</i>
Secondary (Class X)	96.80%

Research and Publications

- Yash Bansal, et al. “Zero-Shot Vision Language Reasoning via Dual-layer Scene Graph Chain of Thoughts.” AAAI 2026 Student Abstract. [Paper]
- Pranshu Danani, Yash Bansal, et al. “TWiST: Temporal Weakly-Supervised Triplets Recognition in Surgical Videos.” AAAI 2026 Student Abstract . [Paper]
- Yash Bansal, et al. “Weighted ColorMorphology Feature Fusion for Tuberculosis Bacilli Detection from Cytopathology Images.” ICVGIP 2025. [Paper]

Experience

Machine Learning Intern — Whyminds.ai	<i>June 2025 – July 2025</i>
• Engineered a multi-agentic pipeline using CrewAI automating LinkedIn content publishing: scraping AI news, NER-based topic tagging, LLM generation, and activity logging; publishing 5+ posts/week.	
• Built a document OCR pipeline using PaddleOCR + LayoutLMv3 for field extraction from noisy scanned PAN and tax documents with post-processing via regex normalization and confidence thresholding.	

Projects

DocVault: Zero-Hallucination RAG Pipeline for Policy Q&A	GitHub
• Built a production RAG pipeline with NLI verification (DeBERTa-v3) stripping contradicted claims at inference and Redis-backed session memory; achieved 0.995 faithfulness and 0.5% hallucination rate .	
• Engineered hybrid retrieval (Pinecone dense + FTS5 BM25) with RRF fusion , cross-encoder reranking, confidence gating and hierarchical chunking; reached context recall 0.84 (95% CI 0.78–0.90).	
• Designed LLM-as-judge harness measuring faithfulness, context recall/precision, and refusal rate; automated ingest-to-monitoring via Celery task chains , Prometheus metrics, per-query tracing, and hallucination alerts.	
Research-Swarm: Multi-Agent Research System	GitHub
• Architected a LangGraph multi-agent system (Planner → Researcher → Verifier → Analyst → Writer) with cost-aware model routing across open and closed LLMs, Qdrant-powered memory, and citation-grounded outputs.	
• Productionized with FastAPI, Docker, human-in-the-loop controls, and RAGAS-based eval pipelines ; live observability across agent traces and retrieval quality.	
Deep Learning Model for Detecting Acid-Fast Bacilli in Cytopathology Smears — CSE Dept.	
• Developed a TB bacilli classifier using a novel autoencoder-based reconstruction-loss weighted 6-channel input (lighting, saturation, hue + Sobel edge maps) with channel attention and false-positive focal loss .	
• Outperformed current SOTA, achieving 97.58% accuracy with Grad-CAM interpretability.	
Satellite-Driven Real Estate Valuation via Multimodal Learning — IIT Roorkee	GitHub
• Designed a multimodal system fusing satellite imagery (EfficientNet-V2-S + FPN-Lite + CBAM) with tabular data (MLP) via late-fusion for property price prediction; achieved R² = 0.8182 and RMSE = \$143K .	

Skills

LLM & Agents: LangGraph, LangChain, CrewAI, LiteLLM, OpenAI API, vLLM, Groq, mem0
Retrieval & RAG: FAISS, LanceDB, Qdrant, Pinecone, SQLite FTS5, RAGAS
ML / DL: PyTorch, Hugging Face Transformers, scikit-learn, OpenCV, spaCy, ONNX Runtime
MLOps & Infra: Docker, FastAPI, AWS, GCP, MLflow, W&B, Git/GitHub, Celery, Redis, Prometheus
Data & Databases: Pandas, NumPy, MySQL, MongoDB, SQL

Extra Curricular

Joint Secretary — iGEM, IIT Roorkee	<i>February 2024 – June 2025</i>
• Built a computational ML pipeline for in-silico prediction of antibacterial peptides; led dry lab ideation and mentored 10+ new recruits.	
Joint Secretary — Institute Alumni Relations Cell, IIT Roorkee	<i>May 2025 – June 2025</i>
• Led a team of 40+ students to conduct 200+ mock interviews for campus students and onboarded 100+ alumni.	
• Organized career talks attended by 400+ students and coordinated sponsorships for the events.	