HARDIK JINDAL

Final Year Undergraduate
Department of Electrical Engineering
Indian Institute of Technology Kanpur

hardik1408
Hardik Jindal
+91-9582890169
hardikj22@iitk.ac.in

Degree/Certificate	Institute	Year	CPI/%
B.Tech	Indian Institute of Technology Kanpur	2022-2026	8.74/10
CBSE(XII)	Bharti Public School, Delhi	2022	94%
CBSE(X)	D.L.F Public School, Ghaziabad	2020	96%

SCHOLASTIC ACHIEVEMENTS

- Secured **All India Rank 889** in **JEE Advanced 2022**
- Secured **All India Rank 2257** in **JEE Mains 2022**
- Secured **All India Rank 894** in **KVPY SX Stream**
- Recipient of the **Dr. S.L. Batra Excellence Award** for exemplary academic performance in the academic year 2021-22

WORK EXPERIENCE

Oracle *May'25 - Jul'25*

Project Intern | Data Space Transactions Team

- Explored **Fuzzy Clustering Algorithm** as an alternative for Kmeans algorithm for **IVF Indexes**, enhancing soft assignment
- Developed hybrid pipelines (K-Means+fuzzy and full fuzzy) for dense embeddings using Oracle **Database 23ai** with IVF indexing
- Designed clustering workflows with membership matrix generation and **centroid refinement** for **high-dimensional data**
- Benchmarked using datasets like **MNIST, GloVe** (25,50,100) and **LastFM**, achieveing about **7**% **improvement** in accuracy

Google Summer of Code *May'24 - Aug'24*

Ceylon Computer Science Institute | Mentors: N.Weerasekara, Tingyuan Cui

- Developed a **RAG** based **RESTful** API in **Flask** for retrieval of cybersecurity news from **100+** articles by leading news outlets
- Transitioned the codebase to open LLMs like **LLaMA3, Gemma, MistralAI** by integrating a **HuggingFace-based** endpoint
- Integrated **Langchain** to establish a **LLM chain** and facilitate structured data processing and formatting for news handling
- Utilized **Pinecone** as a **vector database** in order to efficiently store and reduce retrieval times for **data embeddings**

MAJOR COMPETITIONS

IEEE IROS Robot Simulation Competition *Aug'23 - Dec'23*

Team Humanoid, IITK

- Integrated a **YOLOv5** model for recognizing the **position, depth** and **spatial orientation** of opponent in real-time
- Implemented **image segmentation** and **contour detection** techniques, including **Hough transform** for detection
- Achieved **2nd** position **worldwide**, competing against **30+ international** teams from various universities across the globe

Inter IIT Tech Meet 13.0 | *Oct'24 - Dec'24*

Dynamic Agentic RAG with Pathway | High Prep Event

- Designed a multi-agentic architecture using **Monte Carlo Tree Search and LATS**, provided by **LangGraph** for adaptive, autonomous task resolution
- Integrated **LiteLLM** for embeddings and indexing with **BM25** and **HNSW**, optimized via **Reciprocal Rank Fusion**
- Used **Topological Sorting** on a DAG to optimize data flow & enable **quasi parallel** execution and calibrate reasoning process
- Achieved an average cost of **USD 0.12** and inference time of **1.5 minutes** per query with **GPT-4o**

Inter IIT Tech Meet 12.0 *Oct'23 - Dec'23*

DevRev AI Agent 007: Tooling up for Success | High Prep Event

- Developed a Dynamic Tool Orchestration system for **LLMs** and finetuned the **LLaMA2** model using techniques like **QLoRA**
- Devised a custom **Retriever** using **FAISS** along with **L2 Metric** to retrieve tools & examples relevant to the query
- Achieved an **EM score** of **64%**, optimizing average inference cost to **USD 0.007** with **GPT 4 and 3.5 Turbo** LLM

KEY PROJECTS

Differentiable Discrete Sampling *Jan'25 - Present*

MIT Media Lab

- Analyzed limitations of **Gumbel-Softmax** and **StochasticAD** in hybrid **discrete-continuous systems** via variance diagnostics
- Conducted simulations across **Markovian** and n-ary random walks to quantify **gradient instability** in long-horizon settings
- Performed extensive experiments on **discrete control tasks**, **multi-agent coordination** and **stochastic simulations**
- Proposed **Adaptive Uncertainty Gating**, a scalable low-variance gradient estimator for discrete-continuous optimization

Optimization and LLMs *Aug'25 - Present*

MIT Sloan School of Management | Prof. Swati Gupta

- Exploring how we can use LLMs in order to solve **Multi-objective Reinforcement Learning** problems (MORL)
- Inspired by **OptiMUS-0.3** , designing an **Agentic AI** system leveraging LLM-driven reasoning to solve optimization problems

VLM Calibration *Jun'25 - Present*

Manifold Research Organization

- Ran **calibration** experiments on **VLMs** (Qwen, LLaMA, OLMo, Gemma, OpenAI) using Self-Ask, CoT, and confidence probing
- Trained **sparse auto-encoders** on CoT traces to identify calibrated vs mis-calibrated reasoning features in model outputs

High-Performance Parallel Analysis of 3D Data *Feb'25 - Apr'25*

CS633 Course Project | Prof. Preeti Malakar

- Developed **MPI-parallelized system** to compute local/global extrema using a 7-point stencil across a decomposed domain
- Optimized performance with **halo exchange** and **overlapped computation** to reduce latency and CPU idle time
- Implemented efficient **parallel I/O** for direct, non-contiguous data access for **scalable processing** of datasets up to 16 GB

Multimodal In-Context Learning with VLMs *Jul'25 - Present*

Undergraduate Project | Prof. Sayak Ray Chowdhury

- Analyzing the effect of **multimodal prompts** on generative fidelity and semantic consistency across **VLMs**
- Deployed multiple open-source VLMs **(LLaMA, Gemma, Qwen)** using **vLLM** for **high-throughput** and **low-latency serving**
- Experimenting how VLMs internalize **exemplar-conditioned** distributions and transfer **contextual patterns**

MRL for Multiresolution Image Synthesis *Jan'25 - May'25*

Undergraduate Project | Prof. Soumya Dutta

- Inspired by **'Matryoshka Representational Learning'**, devised a loss function that captures **latent features** at different resolutions
- Implemented Deep Learning architectures like **Autoencoders** and **Diffusion Models** and extensively tested on 3D volume data
- Improved the performance of Conditional Autoencoders by **3-5 PSNR** at lower resolutions with fewer training data

TECHNICAL SKILLS

- **Languages:** C | C++ | Python | Bash | Rust | Haskell | Ruby
- **Utilities:** Git | GitHub | Linux | HuggingFace | LaTeX
- **Libraries:** Transformers | OpenCV | Torch | LangChain | vLLM

POSITIONS OF RESPONSIBILITY

Coordinator, Programming Club IITK *Apr'24 - Apr'25*

- Lead a **two-tier** team comprising of **25 secretaries** to foster a culture of programming within the campus community

RELEVANT COURSES

(**:OUTSTANDING PERFORMANCE) (*:EXCELLENT PERFORMANCE) († ONGOING)

Data Structure and Algorithms	Fundamentals of Computing*	Probability and Statistics	Parallel Computing*
Convex Optimization for ML**	Introduction to ML	Linear Algebra and ODE	Statistical NLP†