Kartik Vats

+91 8700451204 | kartikvatsj@gmail.com | LinkedIn | Github | Leetcode| Codeforces | Codechef | Portfolio

Education

Year	Degree	Institute	Grade
2023 – 2027	B.Tech (CSE)	Delhi Technological University (DTU), Delhi	8.38 CGPA
2023	12th CBSE	The Vivekananda School	91.2%
2021	10th CBSE	DAV Public School, Delhi	97.6%

Experience

Deep Learning Research Intern May 2025 - Present
ML Research Lab, DTU under Prof. Anil Singh Parihar *Delhi, India*

- **Designed a novel statistical low-light preprocessing pipeline** achieving **state-of-the-art 38.42%** accuracy (scaling to **53.89%**), outperforming prior neural baselines by **+6.02%/+26.83%** on **ELLAR** (ELL/LL splits); adapted and extended **up-projection concepts**, implementing custom blocks improving **feature reconstruction and stability** over standard deconvolution in extreme low-light settings.
- Evaluated **ResNet, ViT, CLIP, BLIP, Grounded DINO** encoders for **low-light representation learning** with emphasis on **semantic supervision**, **cross-task generalization**, ablation analysis, reproducibility; identified **data scarcity** and **illumination bias** as key limitations of neural approaches; ongoing work toward research publication.
- **NeurIPS 2026 submission (under review):** *LegalBrain Indic Legal Corpus* — independently built a **20K-context** Indian legal QA benchmark; **BM25 (R@1=0.706)** outperforming 8 retrieval configs and 3 generation baselines (**250M–8B params**); dataset on Hugging Face (CC BY 4.0).

Projects

SafeGen Arena - RL Environment for Image-AI Safety | Github | HF Space |*GRPO, LoRA, CLIP, OpenEnv*

- **Built an OpenEnv-compliant RL gym** training Qwen2.5-1.5B + LoRA as a prompt-rewriting safety layer for diffusion models with 3-way action space (allow/transform/reject); GRPO training from $-0.05 \rightarrow +0.33$ reward over 1,300 steps with zero mode collapse.
- **Designed 4-term reward** combining Llama Guard 3, CLIP concept-arithmetic intent residual, NudeNet/Q16 visual judge, and over-refusal penalty; novel residual `embed(prompt)` $-\alpha\times$ `embed(unsafe_axis)` penalises verbatim copies.
- ***Top 100, Meta × OpenEnv Hackathon 2026 (5,000+ teams).***

BidNet - Real-Time Bidding Model | Github | *AutoEncoders, Contrastive Learning, Feature Engineering, NLP*

- Innovated a raw-to-dense feature pipeline, transforming raw bid data using **contrastive embeddings** and **autoencoders**, followed by an **ANN-based feature compression** for efficient representation.
- Devised a low-latency RTB model predicting bid price and bidding decisions within **5ms per request**, optimized via **Grid Search** and **Adaptive Learning Rate Scheduling**.
- Achieved **82% classification accuracy**, scaling to handle **100K+ bid requests/second** in large-scale ad exchanges.

LLM Safety - Red Teaming & Mitigation | Github | *PyTorch, Hugging Face, DistilBERT, CoT*

- Engineered a **DistilBERT safety classifier (87% F1, 3 classes)** with batch inference and confidence scoring; **end-to-end red-teaming pipeline** (500+ adversarial prompts, 3 harm categories, 3 mitigation agents) **reducing unsafe generations by 42%.**
- Built **observability suite** (confusion matrices, per-class metrics, JSON/visual exports) for **real-time monitoring and drift detection** in production LLM safety workflows.

InquireAI - AI Search Assistant | Github | *Flutter, Tavily API, WebSocket, Sentence Transformers, FastAPI*

- Created a **real-time cross-platform AI assistant** using Flutter and FastAPI, integrating **sentence transformers** with **WebSocket streaming**.
- Built custom **RAG pipeline** combining Gemini API with Tavily Search, achieving **85% accuracy** in source ranking.

Technical Skills

Languages: Python, C/C++, JavaScript, HTML/CSS, Dart, SQL
Web/App Frameworks and Libraries: Flutter, React.js, TailwindCSS, Node.js, WebSockets, AsyncIO
AI/ML Technologies: PyTorch, TensorFlow, Scikit-Learn, OpenCV, HuggingFace, Transformers, GANs, FastAPI
Specialized Techniques: Prompt Engineering, CoT, RAG, MCP, Evaluation Pipelines, RL Environments
Database Systems: MongoDB, Firebase, PostgreSQL, ChromaDB, VectorDB
Developer Tools: Git/Github, Linux, Anaconda, Jupyter Notebook, Postman, Azure, Vercel, Docker
Coursework: OOPs, Computer Networks, Operating Systems, DBMS, Linear Algebra, Probabilty and Statistics
Certifications: Machine Learning Specialization - Coursera, Deep Learning Specialization - Coursera, Google Gen AI study Jam , Google Cloud study Jam

Achievements

- **Specialist(1535)** on CodeForces, **Knight(1982)** on Leetcode, **3 star(1636)** on CodeChef.
- Ranked **1345** in Meta Hacker Cup, **1677** in CodeForces Round 1034 and **972** in LeetCode Weekly Contest 462, **535** & **515** in CodeChef Starters 171 and 172 respectively (among 25,000+) and solved over **900+ problems**.
- **Winner** of h4ck0n-CTF'25 **(1st place)**, IICON-CTF'25 **(1st place)** and Hack-IIT Kanpur **(3rd place)**.
- Secured **2nd position** at **VisionXAI Hackathon**, held during DTU's annual techfest, Invictus.
- **Grand Finalist** (Ranked 12 out of 10,000 teams) in **NCIIPC-AICTE Pentathon 2025**.
- Ranked **4th in university (723rd worldwide)** in **IEEE Xtreme 18.0**, a CP contest among 9,500+ teams.
- **Under 500** in Amazon ML Challenge 2025 among 20,000 teams.
- Core member of EHAX – DTU's Cybersecurity Society; co-organized EHAX CTF with 3,000+ global participants and contributed to achieving **India Rank 7** and **World Rank 33** on CTFtime.