DHRUV JAIN

+91-9198700175 | Indian Institute of Technology (BHU), Varanasi
Email | LinkedIn | GitHub | Website

PROFESSIONAL SUMMARY

Research Engineer at Ola Krutrim with experience building, training, and evaluating large-scale language and speech models across production and research environments. Worked on end-to-end model training pipelines, data curation, and deployment for multilingual and Indic AI systems. First-author publications at AAAI 2026 and under-review work at ICLR 2026 focused on improving the reliability and reasoning capabilities of open-source LLMs and SpeechLMs.

WORK EXPERIENCE

Research Engineer - Ola Krutrim Bangalore, India
LLM & Speech Modeling Team *June 2025 - Present*

- Post-trained open-source LLMs for the AI assistant platform Kruti.ai using SFT, DPO, and GRPO, improving safety (moderation, jailbreak and identity), search, and tool orchestration. Reduced production latency by **40%** and lowered operational costs by consolidating complex LLM chains into unified models for real-time traffic.
- Owned the end-to-end training cycle, including open-source and synthetic data curation, training and evaluation, and iterative improvement driven by failure analysis. Built multi-node distributed training and inference pipelines with TRL, DeepSpeed ZeRO and vLLM for efficient training and production deployment.
- Developed VoiceAgentBench, the first multilingual speech-agent benchmark with 6k+ spoken queries spanning English, Hindi, and five Indic languages, enabling systematic evaluation of SpeechLMs on agentic tasks such as tool calling, multi-turn planning, and adversarial safety across diverse linguistic and cultural contexts.

PUBLICATIONS

VoiceAgentBench: Are Voice Assistants ready for agentic tasks? arXiv | GitHub
Dhruv Jain*, *Harshit Shukla**, *et al.*, Under Review at ICLR 2026

Improving Physics Reasoning in Large Language Models Using Mixture of Refinement Agents arXiv
*Raj Jaiswal**, ***Dhruv Jain****, *et al.*, AAAI 2026 TrustAgent Workshop

SwissNYF: Tool Grounded LLM Agents for Black Box Setting arXiv
Somnath Sendhil Kumar, ***Dhruv Jain****, et al.*

EDUCATION

Indian Institute of Technology (BHU), Varanasi *Dec 2021 - May 2025*
Bachelor of Technology - Electronics Engineering; **CGPA : 8.36**

ACADEMIC RESEARCH

Research Intern - MIDAS Lab, IIIT Delhi *Jan 2025 - May 2025*
Supervisor: Prof. Rajiv Ratn Shah GitHub

- Implemented SCoRe (Self-Correcting RL) for high-school physics reasoning: a two-attempt pipeline where small LLMs generate a solution, receive step-level error feedback from a fine-tuned verifier model, and produce a corrected solution, trained end-to-end via on-policy RL with KL-regularized reward maximization.
- Curated PhysicsQA, a diagnostic benchmark of 370 high-school physics problems with verified chain-of-thought solutions. Developed a novel error taxonomy, categorizing failures into miscomprehension, conceptual error, and calculation mistakes to enable fine-grained analysis of LLM reasoning.
- Built a memory-efficient training pipeline using DeepSpeed Engine with ZeRO, LoRA adapters, and KV-cache-based incremental log-probability computation, enabling multi-turn RL training.

Research Intern - NTU Speech Lab *May 2024 - Nov 2024*
Supervisor: Prof. Chng Eng Siong GitHub

- Developed a joint-task framework for ASR post-processing using Llama-2-7B, integrating Generative Error Correction (GER) and Inverse Text Normalization (ITN) to produce high-fidelity ASR transcripts.
- Curated HyPost, a 400K-example training dataset derived from six diverse corpora, specifically designed to eliminate alignment issues in ASR post-processing through joint multi-task training.
- Implemented a domain adaptation strategy using Mixture of LoRAs (X-LoRA), closing up to 80% of the performance gap on out-of-distribution data while maintaining denormalization accuracy.

SKILLS

Programming Languages	C++, Python
Tools & Frameworks	PyTorch, HuggingFace (Transformers, TRL, PEFT), DeepSpeed, verl, vLLM, LangChain, LlamaIndex, Autogen, Docker
Areas of Interest	LLM Reasoning, AI Agents, Speech Modeling, Post-Training, AI Safety

PROJECTS

Dynamic Multi-Agent RAG System GitHub | Blog
Inter-IIT Tech Meet 13.0

- Designed a dynamic agentic RAG system for long legal and financial documents, featuring a interleaving reasoning-retrieval approach that enables LLMs to dynamically toggle between multi-hop reasoning and context retrieval.
- Engineered a fault-tolerant tool-calling architecture using code-driven reasoning and reflection to mitigate API failures, along with a Retrieval Memory module to cache context and optimize multi-turn query efficiency.
- Benchmarked retrieval performance against LongRAG and hierarchical baselines on CUAD and FinanceBench datasets, achieving superior Mean Reciprocal Rank (MRR) across complex multimodal documents.

LLM Agent for Black Box Tool Access GitHub | Paper
Inter-IIT Tech Meet 12.0

- Developed TOPGUN, a novel black-box tool planning method that leverages program synthesis to strategize tool usage without requiring direct access to API calls or their implementations and internal logic.
- Architected SwissNYF, a comprehensive end-to-end pipeline integrating a Tool Inventory, Retriever, and Tool Planner to enable seamless tool integration and efficient execution of complex, multi-step queries.

LEADERSHIP, HONOURS AND ACHIEVEMENTS

- **Institute Blue, IIT (BHU)** - Awarded the institute's highest distinction for outstanding technical achievements, research excellence, and impactful leadership contributions.
- **General Secretary 2024-25** - *Science and Technology Council, IIT (BHU)* - Led institute-wide technical initiatives, coordinated large-scale technical events, and drove research and innovation activities across 8 clubs.
- **Joint Secretary 2023-24** - *Robotics Club, IIT (BHU)* - Oversaw event execution, cross-team coordination, and mentored student-led robotics and ML projects.
- **Contingent Leader** - *Inter-IIT Tech Meet 13.0* - Led the institute contingent (100+ students) across 14 technical competitions among 23 IITs.
- **Inter-IIT Participation (ML)** - *Inter-IIT Tech Meet 11.0 (2023), 12.0 (2024), 13.0 (2025)* - Represented the institute in the Machine Learning Hackathons among 23 IITs.