

BHAVESH GURNANI

✉ cs5210594@iitd.ac.in 📞 (+91) 8949778677 📧 bhavesh-gurnani-410a68217 🌐 8have5h 🌐 <https://8have5h.github.io/>

EDUCATION

- **Indian Institute of Technology Delhi (IIT Delhi)** *Expected May 2026*
 - B.Tech and M.Tech in Computer Science & Engineering (Dual Degree) | CGPA: 8.52/10.0
- **Lord Buddha Public School, Kota** *Apr 2021 (CBSE XII), Apr 2019 (CBSE X)*
 - Class XII (PCM): 84% | Class X: 94%

SCHOLASTIC ACHIEVEMENTS

- **JEE Advanced 2021:** All India Rank **116** (out of approx. 200,000 candidates)
- **JEE Mains 2021:** All India Rank **493** (out of >1,000,000 candidates); **100% Percentile** in Mathematics
- **KVPY Fellowship 2020 (IISc):** All India Rank **305** (SA Stream)
- **NTSE Scholar 2019 (NCERT):** Ranked among Top **2000** students nationwide
- **INMO Qualifier:** Ranked among Top **300** students nationwide

EXPERIENCE

- **Research Intern** | *Microsoft Research* *Feb 2026 – April 2026*
 - Developing conversational agents for medical diagnosis via multi-turn dialogue modeling and clinical reasoning.
- **AI Security Intern** | *National University of Singapore — Prof. Prateek Saxena, KISP Lab* *Dec 2025 – April 2026*
 - Extended *OverThink* (Kumar et al., 2025; arXiv:2502.02542) with a filtering-evasive slowdown attack on LLMs.
 - Showed it increases reasoning-token generation without changing outputs, revealing compute-amplification risks.
- **Post-Training Research Intern** | *Neosigma / Remote* *Dec 2025 – Jan 2026*
 - Improved 128k-context agentic LLM post-training by tool-call truncation during SFT, improving context retention.
 - Set up AWS multi-node training infrastructure for scalable SFT, DPO, and GRPO with `ver1` at 128k context.
- **Quantitative Researcher Intern** | *Ebullient Securities / Gurugram* *May 2025 – July 2025*
 - Built and backtested alpha signals for short-term equity directionality from equities data.
 - Designed and live-tested a custom trading strategy in C++ using top alphas (40% annualised, Calmar 12).
 - Automated alpha-generation pipeline using LLMs with custom evaluation loops for iterative refinement.
- **Software/Research Engineer Intern** | *CompilerAI Labs Private Limited / New Delhi* *May 2024 – July 2024*
Project under the guidance of Prof. Sorav Bansal
 - Implemented MISRA C static analysis rules within Clang compiler across Preprocessor, AST, and LLVM stages.
 - Enhanced Equivalence Checker tool for formal verification of compiled executable code against C source code.

TEACHING ASSISTANTSHIP

- **TA, COL772: Natural Language Processing** | *IIT Delhi / Prof. Mausam* *Jan 2026 – May 2026*
 - Designed assignments and quizzes; built autograder to evaluate student submissions; resolved Piazza queries.
- **TA, COL333: Artificial Intelligence** | *IIT Delhi / Prof. Mausam* *Jul 2025 – Dec 2025*
 - Created quizzes and assignments; evaluated student scripts.

PUBLICATIONS

- **Hierarchical Masked Diffusion Language Models (In Review)**
Authors: Parul Gupta, Avni Aggarwal, **Bhavesh Gurnani**, Arnab Kumar Mondal, Dinesh Khandelwal, Dinesh Garg, Parag Singla
 - Proposed **HMDLM**, a two-stage diffusion model that predicts token-level generation priorities before decoding.
 - Derived novel ELBO with per-level noise schedules for controllable non-autoregressive generation.
 - Achieved 0.91 PPL improvement on LM1B, 1.4% gain on GSM8K, and 3% on code generation.
- **ProbMedTOD: A Probabilistic Task-Oriented Dialogue System for Patient History Taking (In Review)**
Authors: Vishal Vivek Saley, **Bhavesh Gurnani**, Dinesh Raghu, Mausam
 - Introduced **ProbMedTOD**, integrating LLMs with Bayesian Networks for patient history-taking.
 - Designed patient simulator and dialogue-level metric; achieved 21-pt MRR gain over baselines.

SELECTED PROJECTS

- **Hierarchical Masked Diffusion Language Models** | *MTP under Prof. Parag Singla* *June 2025 – Present*
 - Implemented hierarchical masked-diffusion with staged generation: descriptions → comments → code tokens.
 - Achieved empirical improvements on LLADA for MBPP and HumanEval code-generation benchmarks.
- **Improving Dialogue Systems for Medical Diagnosis** | *BTP under Prof. Mausam* *July 2024 – May 2025*
 - Modeled large-scale medical text data as Bayesian Networks to capture causal disease-symptom relationships
 - Designed novel inference techniques for Bayesian Networks, improving performance over standard RAG systems

COURSE PROJECTS

- **Object-Centric Learning for Scene Decomposition | *DL Course Final Project*** *Jan 2024 – Apr 2024*
 - Implemented Slot Attention on CLEVRText and designed slot-conditioned diffusion model with custom U-Net
 - Benchmarked performance using ARI metric and achieved interpretable object-centric scene representations
- **Intelligent Bot for Alternating Markov Games | *AI Course Final Project*** *Aug 2023 – Dec 2023*
 - Developed AI agents for strategic games using depth-limited Minimax with domain heuristics optimized via RL
 - Implemented DQN enhanced with MCTS inspired by Expert Iteration, achieving superior gameplay performance
- **Lexer, Parser, Interpreter for Language with Rationals | *Prof. S. Arun Kumar*** *Jan 2023 – Apr 2023*
 - Built complete compiler with custom BigInt/Rational modules, lexer using ML-Lex, and parser with ML-Yacc
 - Implemented interpreter with static scoping supporting variables, conditionals, loops, and recursion
- **Custom MultiThreading Library | *Prof. Abhilash Jindal*** *Feb 2023 – Mar 2023*
 - Designed a user-level multithreading library in C with round-robin scheduling and synchronization primitives
 - Benchmarked against POSIX pthread, demonstrating efficient scheduling with reduced overhead
- **Table Cell Classification for Question Answering | *Prof. Mausam*** *Feb 2024 – Mar 2024*
 - Designed a model to identify answer cells in custom tables for complex natural language queries
 - Built a two-stage transformer architecture: column predictor followed by row predictor conditioned on query and predicted columns
 - Prompt-tuned Gemma, improving exact match accuracy from 43% (two-model pipeline) to 67%
- **Text-to-Math Program Synthesis | *Prof. Parag Singla*** *Jan 2024 – Mar 2024*
 - Developed Seq2Seq models to generate executable math formulas from text problem descriptions, with an interpreter for answer computation
 - Implemented Bi-LSTM encoder-decoder with GloVe embeddings and Bahdanau attention; achieved 67% exact match and 73% execution accuracy using beam search
 - Fine-tuned BERT as encoder, improving to 78% exact match and 81% execution accuracy
- **Lay Summarization of Biomedical Research Articles | *Prof. Mausam*** *Apr 2024 – Jun 2024*
 - Fine-tuned PLMs on HPC clusters to generate layperson summaries of biomedical research articles
 - Applied PEFT techniques (Adapters, LoRA) with mixed-precision training for efficient large-scale fine-tuning
 - Evaluated Flan-T5 and BioGPT on relevance, readability, and factual consistency metrics

EXTRACURRICULAR ACTIVITIES

- **Tower Research Limestone Data Challenge:** Achieved **Top 50** selection **twice** across offline sessions
- **Inter IIT Tech Meet 2024:** Led team on Adobe Research problem; built a detection and explainability pipeline for AI-generated images, identifying visual artifacts and providing reasoning for model predictions
- **ML in Astronomy Workshop (2024):** Conducted a session on ML–Astronomy integration; built LSTM-Attention models on light curves for object classification [\[Code\]](#)
- **ARIES Research Team (2024):** Conducted NLP sessions and contributed to multiple industry research projects
- **Technical Executive, Economics Club, CAIC (2022–23):** Developed a React.js website to manage club competitions and events
- **Technical Engineer, Infinity Hyperloop, CAIC (2022–23):** Designed a PyQt5 interface for pod control; implemented sensor processing and STM communication over CAN protocols
- **Academic Mentor (2023):** Mentored 150+ freshers for MTL100 (Calculus) via hybrid doubt sessions