




Harsh Poonia

Curriculum Vitae

 <https://cse.iitb.ac.in/~harshpoonia> |  harshpoonia@cse.iitb.ac.in |  +91 81264 71862

EDUCATION

Indian Institute of Technology Bombay

2021-2025

B.Tech. (with **Honors**) in Computer Science and Engineering (Grade : **9.42/10**)

RESEARCH INTERESTS

Natural Language Processing, Causal Machine Learning, Foundation Models, Machine Unlearning, Explainable AI

PUBLICATIONS

1. χ SPN: Characteristic Interventional Sum-Product Networks for Causal Inference in Hybrid Domains
[Paper] (UAI 2024)
Harsh Poonia, Moritz Willig, Zhongjie Yu, Matej Zečević, Kristian Kersting, Devendra Singh Dhami

RESEARCH EXPERIENCE

Equivariant Representations for Speech

July 2024 - ongoing

Bachelor's Thesis Project / Prof. Preethi Jyothi

IIT Bombay

- Developing an **equivariant-by-design** network for a suite of speech tasks including classification, diarization, recognition on **accented** speech with the aim of learning representations despite dearth of data and compute
- Studied the concepts of transformations, lie groups and **invariant** subspaces for **G-equivariant** neural networks
- Adopting **canonization** methods for invariance to compute canonical representations for perturbed, accented speech
- Working on inherently **shift-equivariant** vision transformer models, and reviewing equivariant **position embeddings**
- Adapting **capsule networks** (for interpretability) and **lie-group** invariance to sequence models like transformers

Neuro-Causal Models for Tractable Inference

May 2023 - Febraury 2024

Research / Prof. Devendra Singh Dhami

Hessian Center for Artificial Intelligence & TU Darmstadt

- Developed a novel **characteristic interventional** sum-product network for causal inference in **hybrid domains** aimed at learning robust, generalizable causal representations; resulting in a **first-author publication at UAI 2024**
- Reviewed causal approaches to fairness, counterfactual data augmentation and **invariant** feature/mechanism learning
- Employed probabilistic circuits to model **interventions** in data-generating process, maintaining **tractability** of inference

Neural Granger Causality using xLSTMs

July 2024 - ongoing

Research / Prof. Devendra Singh Dhami

Hessian Center for Artificial Intelligence & TU Eindhoven

- Working on a novel adaptation of Extended Long Short-Term Memory (**xLSTM**) to discover **granger causal** relationships between **multivariate** time series, with the goal of learning **sparsest** relations that forecast accurately
- Explored granger causality literature based on component-wise MLP and recurrent models, with associated sparsity constraints
- Performing ablation studies with **scalar** and **matrix** xLSTM memories for modeling of **long-range** dependencies

Improving Text-to-SQL using In-Context Learning

Spring 2024

Prof. Sunita Sarawagi / RnD Project

IIT Bombay

- Researched ideas on teaching sequences for **meta learning**, and bayesian perspective on choosing good demonstrations from viewing LLMs as **latent variable** models to build a better theoretical understanding of in-context learning
- Compared the ICL ability of LLMs with explicit **finetuning** and related it to **gradient descent** and meta optimisation
- Worked on **question decomposition** methods and **fine-grained** sharing to improve performance on BIRD bench

INTERNSHIPS

Automated Desk Mandates

Software Development Intern, Amsterdam

May 2024 - July 2024

Optiver

- Completed a week-long trading training program, covering financial instruments such as futures, ETFs, options, and their **risk measures** using the **Greeks** and **Black-Scholes** model for option pricing across various expirations
- Developed a **daily execution** component for **T+1 postprocessing** of trade logs, helping with regulatory compliance
- Published **automated desk-level mandates** using the Atlassian API and jinja templating engine, after reconciling various instrument data sources with a focus on **scalability**, **speed** and minimal dependency risks for **robustness**

Smart Contract Security Audit

Research and Development / Internship

December 2022 - January 2023

Marsh McLennan

- Researched smart contract **security vulnerabilities** like re-entrancy, denial of service, integer overflow, timestamp dependance, common hacks and exploits like frontrunning, phishing and griefing and **safeguards** against them
- Tested security analysis tools **slither**, **mythril**, **manticore**, **securify** and built APIs for their use in the platform
- Created API wrappers around **AWS S3** buckets and **DynamoDB** for the backend using **FastAPI** and **Nodejs**

KEY PROJECTS

Adversarially Robust Generalisation

Advanced Machine Learning, Prof. Sunita Sarawagi

Spring 2024

Course Project

- Employed **robust self supervised** training with OOD unlabeled samples to mitigate adversarial attacks on the domain
- Studied improved **generalisation error** bounds for abundant unlabeled data supplementing scarce labeled data
- Infused dynamic **open-set** noise labels into RSS training to further reduce **overfitting** on inherently noisy labels

Building C Compiler from Scratch

Implementation of Programming Languages Lab, Prof. Uday Khedker

Spring 2024

Course Project

- Built a compiler to incrementally support a set of features in C including conditionals, **control flow** and **recursion**
- Generated **intermediate representations** - Abstract Syntax Trees, Three Address Codes, Register Transfer Language, MIPS assembly code with **register allocation**, using tools like lex and yacc for scanning and bottom-up parsing

Geodesic Graph Neural Networks

Learning with Graphs, Prof. Abir De

Autumn 2023

Course Project

- Perused the expressive power and lower computational complexity of GNNs based on **geodesic** (representation of shortest path between nodes) embeddings and task-specific pooling for link prediction, node and graph classification
- Extended the implementation by adapting embeddings to **heterogenous** graphs and **attention-weighted** pooling

Speech Seminar Series

Automatic Speech Recognition, Prof. Preethi Jyothi

Spring 2024

Course Project

- Delivered a talk on a **multimodal** pretraining model that encodes **acoustic** and **textual** information into a unified representation, allowing pretraining from either modalities to improve speech to text translation
- Wrote a blog on Meta's **Voicebox**: a multilingual, **non-autoregressive** speech synthesis model based on flow matching
- Created and presented **poster** for Align with Purpose: a general framework to improve a given property in CTC models

Text Style Transfer

Artificial Intelligence and Machine Learning, Prof. Preethi Jyothi

Autumn 2023

Course Project

- Implemented SOTA methods to manipulate text style attributes while preserving content without explicit disentanglement, including **style transformer** and latent representation **steering** on an autoencoder by nudging towards desired style
- Improved performance on Yelp and Amazon datasets by adding cyclic and self reconstruction objectives in training

OTHER PROJECTS

Railway Journey Planner

Data Structures and Algorithms Lab, Prof. Supratik Chakraborty

Autumn 2022

Course Project

- Structured a railway planner in C++ using data structures such as **BSTs**, **Dictionaries**, **Heaps**, **Tries** and **AVLs**
- Optimized searching for station name using **Trie** data structure and searching for reviews using **KMP** algorithm
- Adapted quicksort on linked list to list departures sorted by time and **priority queue** to fetch list of top-rated trains

Operating Systems

Operating Systems Lab, Prof. Purushottam Kulkarni

Autumn 2023

Course Project

- Extended memory management in xv6 by adding on-demand paging and **Copy-on-Write** optimisation on shared pages
- Developed various synchronization primitives, such as **semaphores** and **barriers**, utilising atomic instructions in xv6
- Added support for **user space threads**, a custom priority process scheduler and low level filesystem interaction

Cache and Memory Hierarchy Optimizations for Graph Analytics

Computer Architecture, Prof. Biswabandan Panda

Spring 2023

Course Project

- Simulated combinations of various LLC cache associativities, replacement policies (LRU/LFU/FIFO), inclusive/exclusive hierarchies and maintained coherence using **MSI** in **Champsim**, optimised IPC for BFS, Dijkstra, PageRank traces
- Coded a **best offset** data prefetcher based on recent requests, with throttling and MSHR threshold to conserve bandwidth

First Order Model for Image Animation

Computer Vision, Prof. Sharat Chandran

Spring 2023

Course Project

- Analysed a self-supervised **object-agnostic** model to decouple appearance and motion information for image animation
- Generated a dense optical flow from sparsely learnt **keypoints** and jacobians combined with local affine transformations through **taylor** approximations, and an occlusion mask for better aligned source-driver input to a generation module

SCHOLASTIC ACHIEVEMENTS

- Topped a class of **70** students in the graduate-level course on **Computer Vision** ('23)
- Secured **All India Rank 61** in **JEE (Advanced)** among **150,000** candidates ('21)
- Achieved **99.956 percentile** and **All India Rank 616** in **JEE (Main)** among **~ 1 million+** candidates ('21)
- Received the prestigious **KVPY Fellowship** with an All India Rank of **178** from the **Govt. of India** ('20)
- Qualified for the **Indian National Mathematical Olympiad** and awarded the **NTSE Scholarship** ('19)

POSITIONS OF RESPONSIBILITY

Chief Editor | Department Editorial Board

(2023 - 2024)

- Directed the creation and publication of **BitStream**, the CSE department magazine, leading an editorial team of **14**
- Conceptualised and co-authored articles on **multiple genres** - insightful interviews, personal pieces and geeky analyses

Teaching Assistant, Calculus I | Department of Mathematics

Autumn 2023

- Conducted weekly problem solving sessions for a tutorial batch of **50** students, handling doubts and explaining concepts

Teaching Assistant, Foundations of Machine Learning | Department of CS

Autumn 2024

- Responsible (in-part) for setting assignments, and grading exams for a **graduate-level** course with **200** students

Department Academic Mentor & Institute Student Mentor | Mentorship Program

May 2023 - ongoing

- Mentored a group of **6** sophomores in navigating academic and personal difficulties, working toward their **holistic** development, contributing in dept. endeavours such as writing **blogs**, organising **sessions** and compiling **resources**
- Helping **12** freshmen acclimatise to college life and make informed decisions on academic and extracurricular avenues

TECHNICAL SKILLS

Programming

C/C++, Python, Java, Javascript, Bash, Solidity, x86 Assembly

Data Science

PyTorch, Tensorflow, Keras, MATLAB, SciPy, NumPy, OpenCV, Pandas

RELEVANT COURSES

Machine Learning
(Graduate Courses)

Advanced Machine Learning, Foundations of Intelligent and Learning Agents, Learning with Graphs, Automatic Speech Recognition, Computer Vision, Formal Methods in Machine Learning, Learning and Inference in High Dimensions

Theoretical and
Systems

Data Structures and Algorithms, Discrete Structures, Design and Analysis of Algorithms, Computer Architecture, Computer Networks, Automata Theory, Optimisation Models, Operating Systems and Database & Information Systems, Programming Languages

EXTRACURRICULAR ACHIEVEMENTS

- Secured **2nd place** in the state of Gujarat in the **Indian National Cartographic Association** Map Quiz ('17)
- Secured **5th place** in SciComp GC, an algorithmic contest with problems in domains of mathematics, physics and computational astronomy, and reached the **top 8** in **CodeWars**, a python based battle game bot contest ('22)
- Recreated the classic **Snake Game** in **Vanilla Javascript** for DevCom IIT Bombay's Game Dev Hackathon ('22)