

Harsh Poonia

• harpoonix.github.io • hpoonia@cs.cmu.edu • linkedin.com/in/harshpoonia • github.com/harpoonix

EDUCATION

Carnegie Mellon University

Master of Science in Machine Learning (Grade: **4.19/4**)

Pittsburgh, PA
December 2026

- Courses: LLM Systems, Probabilistic Graphical Models, Convex Optimization, Deep RL

Indian Institute of Technology Bombay

B.Tech. (with **Honors**) in Computer Science and Engineering (Grade : **9.44/10**)

Mumbai, India
May 2025

- Courses: Graph + Advanced ML, Natural Language Processing, Speech Recognition, Formal Methods in ML, Reinforcement Learning, Statistical ML, Computer Vision, Linear Algebra, Optimization, Networks & OS, Algorithm Design, Databases
- **Awards:** Narotam Sekhsaria Scholar (Top 20/5,000+, **\$42K**), 2025; JEE Advanced AIR **61/150K (Top 0.04%)**, 2021

INTERNSHIPS

Mistral AI

AI Science Intern, Code Agents

Palo Alto, CA
May 2026 - August 2026

- Designed rewards, training environments, and evals for **RL post-training** of code agents, strengthening the agent's ability to **explore and understand codebases**, leading to downstream gains on challenging repository-level SWE benchmarks
- Drove a parallel initiative to diagnose failure modes in **synthetic RL environments** and improve data generation and filtering pipelines, enabling higher-quality, more cost-efficient scaling

Optiver

Software Engineer Intern

Amsterdam, Netherlands
May 2024 - July 2024

- Automated the **T+1 postprocessing** of order logs, replacing **6-month** regulatory compliance cycles with **daily** updates
- Published **automated desk-level mandates** after reconciling various instrument data sources and processing an **entire day's** trade data in **<1 min** with minimal dependency risks, using the Atlassian API and **PostgreSQL** databases

PUBLICATIONS

Exploring Neural Granger Causality with xLSTMs: Unveiling Temporal Dependencies in Complex Data [Paper]

Harsh Poonia, Felix Divo, Kristian Kersting, Devendra Singh Dhami

Published at **NeurIPS 2025**

- **Led** the adaptation of **xLSTM** to discover **sparse** Granger-causal relationships between noisy **multivariate** time series
- Created a stable, robust training paradigm by formulating a joint optimization algorithm with **proximal** gradient descent and dynamic **adaptive** group lasso penalty, consistently enabling **self-selection** of the best features across domains

χSPN: Characteristic Interventional Sum-Product Networks for Causal Inference in Hybrid Domains [Paper]

Harsh Poonia, Moritz Willig, Zhongjie Yu, Matej Zečević, Kristian Kersting, Devendra Singh Dhami

Published at **UAI 2024**

- Formulated **characteristic-function likelihoods** in the **spectral domain** to model mixed discrete-continuous causal data
- Employed probabilistic circuits to model **interventions** in data-generating process, maintaining **tractability** of inference

RESEARCH PROJECTS

Gradient Arbitration

Research Project | Two-Person Team | Prof. Max Simchowitz

Carnegie Mellon University
August 2025 - May 2026

- Investigated **angular and magnitude-based gradient interference** in multi-objective RL training under a **shared policy**
- Formulated a **drop-in gradient arbitration** method based on a **Hyperball** variant, boosting average performance across **10 robotic control tasks** from **82% to 94.5%** versus the best-tuned **AdamW** baseline

Curious LLMs: Guiding RL Exploration with a Curiosity Bonus [Report]

Course Project | Deep Reinforcement Learning

Carnegie Mellon University
October 2025 - December 2025

- Developed a **Semantic Random Network Distillation** framework, integrating embedding-based intrinsic rewards with GRPO to reward novel reasoning traces while fine-tuning **GPT-OSS-20B** on **8K** competition-level math problems
- Improved **DeepMath-500** by **~2.5%** over GRPO; raised **AIME 2024/25 Pass@16** from **60% to 70%**
- Showed **semantic curiosity** improves reasoning while lowering token entropy, distinct from standard entropy bonuses

Invariant Representations for Speech [Report]

Bachelor's Thesis Project | Prof. Preethi Jyothi

IIT Bombay
July 2024 - May 2025

- Developed an **invariant-by-design** network learning canonical representations of **accented** speech for zero-shot generalization
- Ideated and implemented a **flow matching** based generative model for **non-autoregressive** speech synthesis

SKILLS

Machine Learning Programming

PyTorch, CUDA, NumPy, Tensorflow, Keras, Pandas, SciPy, OpenCV, MATLAB
C/C++, Python, SQL, Javascript, Bash, Solidity, x86 Assembly