

Bhrij Patel

bbp13@umd.edu | www.linkedin.com/in/bhrijpatel | [bridge00.github.io](https://github.com/bridge00) | [Google Scholar](#)

EDUCATION

University of Maryland, College Park

College Park, MD

Ph.D. in Computer Science, co-advised by Dinesh Manocha and Amrit Singh Bedi Aug. 2022 – May 2026 (Expected)

Duke University

Durham, NC

B.S. in Computer Science and Math, Creative Writing Minor, advised by Cynthia Rudin Aug. 2018 – May 2022

RESEARCH INTERESTS AND TECHNICAL SKILLS

Interests: Efficient Optimization, Alignment, Reinforcement Learning, LLMs

Languages/Frameworks: Python (Pytorch, Numpy, SciPy, Pandas), Java, HTML

RESEARCH LABS

GAMMA Lab, University of Maryland, College Park

Aug. 2022 – Present

- **Thesis Research: *Policy Optimization for Machine Learning-based Agents and Systems***
- Improving the sample-efficiency of Markovian optimization for Reinforcement Learning methods
- Exploring the use of LLMs for Object Navigation of non-stationary items and for Embodied Question-Answering
- Improving the optimization of AI Systems via multiple LLM-based evaluators

Interpretable Machine Learning Lab, Duke University

Jan 2019 – Mar. 2022

- Worked on an algorithm to generate ultra hi-res portraits of a person given low-res images examples through unsupervised representation learning with AEs
- Cleaned and explored criminal history from a public dataset of Broward County, FL, with ~ 150,000 records, and private sets from Kentucky with ~3,200,000 records

SELECTED PUBLICATIONS AND PREPRINTS

Sample Efficient Optimization for Reinforcement Learning

1. [ICML 2024] **Bhrij Patel**, et al., *Towards Global Optimality for Practical Average Reward Reinforcement Learning without Mixing Time Oracles*. (Link to Paper, Poster)
2. **Bhrij Patel**, et al., *CCE: Sample Efficient Sparse Reward Policy Learning for Robotic Navigation via Confidence-Controlled Exploration*. (2023) (Link to Paper)

LLM/AI System Optimization

1. [ONGOING PROJECT] **Bhrij Patel**, et al., *Provable Sample Efficient LLM Finetuning without Impractical Mixing Time Oracles*. (2024)
2. **Bhrij Patel**, et al., *AIME: AI System Optimization via Multiple LLM Evaluators*. (2024) (Link to Paper)

Embodied Exploration with LLM-based Agents

1. **Bhrij Patel**, et al., *Central Answer Modeling for an Embodied Multi-LLM System*. (2024) (Link to Paper)
2. Vishnu Sashank Dorbala*, **Bhrij Patel***, et al., *Right Place, Right Time! Towards ObjectNav for Non-Stationary Goals*. (2024) (Link to Paper)

WORK EXPERIENCE

REIN | Data Science Intern

Mar. 2020 – May 2020

- Set up automated web extraction of truck accident records from 1975-2018 with Python and SQL
- Cleaned and integrated trucking data into database for development of risk models