

Bharat Prakash

Baltimore
MD, USA
☎ (+1) 970 368 0788
✉ bhp1@umbc.edu
📄 bharatprakash.github.io

Research Statement **The goal of my research is to build language guided hierarchical agents capable of solving complex long horizon tasks using RL and human feedback. More recently, I am interested in exploiting large language models to augment these agents..**

Education

- 2017–present **PhD, Computer Science**, *University of Maryland, Baltimore County (UMBC), Baltimore, MD.*
- 2012–2014 **MS, Computer Science**, *University of Maryland, Baltimore County (UMBC), Baltimore, MD.*
- 2006–2010 **BE, Information Technology**, *University of Pune, India.*

Experience

- 2023 **Visiting Researcher**, *Johns Hopkins University, Baltimore, MD.*
Hierarchical Reinforcement Learning. Large Language Model (LLM) Augmented Agents.
- Summer 2020 **Research Intern**, *ARL, Baltimore, MD.*
Unsupervised RL, Model-based RL and planning.
- Summer 2019 **Research Intern**, *Noodle AI, Palo Alto, CA.*
Decision making under uncertainty. Model-based deep reinforcement learning, MPC for manufacturing processes.
- 2017–present **Graduate Research Assistant**, *UMBC.*
Machine learning, deep reinforcement learning.
- 2014–2017 **Software Engineer**, *eBay/PayPal Inc., Baltimore, MD.*
Worked with PayPal Credit and PayPal Working Capital Team. Back-end web services and API infrastructure.
- Summer 2013 **Software Developer Intern**, *Comcast Labs, Washington DC.*
Xfinity Search Infrastructure Team
- Spring 2013 **Graduate Assistant**, *UMBC.*
Course: Discrete Structures. Grading and assisting the professor.
- 2010–2012 **Software Engineer**, *Persistent Systems, Pune, India.*

Publications and Academic Research Projects

- 2023 **LLM Augmented Hierarchical RL**, [*Reinforcement Learning, Hierarchical RL, Language conditioned RL, LLMs*].
Augmenting hierarchical reinforcement learning agents with LLMs for efficient exploration and generalization (Accepted at the FLAIRS 2024, LangRob workshop @ CoRL 2023 and FMDM workshop @ NeurIPS 2023)
- 2022 **Hierarchical Agents by Combining Language Generation and Semantic Goal Directed RL**, [*Reinforcement Learning, Hierarchical RL, Language conditioned RL*].
An interpretable hierarchical agent framework by combining sub-goal generation using language and semantic goal directed reinforcement learning (Accepted at the NeurIPS LaReL Workshop 2022)
- 2021 **Combining Learning from Human Feedback and Knowledge Engineering to Solve Hierarchical Tasks in Minecraft**, [*Reinforcement Learning, Imitation Learning, Learning from Human feedback*].
2021 NeurIPS MineRL BASALT Challenge Winners (AAAI 2022 Spring Symposium on Machine Learning and Knowledge Engineering for Hybrid Intelligence)

- 2021 **Interactive Hierarchical Guidance using Language**, [*Reinforcement Learning, Hierarchical RL, Language, Machine learning*].
A hierarchical agent framework where low-level sub-goals are specified using language (Accepted at AI-HRI at AAAI Fall Symposium 2021)
- 2021 **Automatic Goal Generation using Dynamical Distance Learning**, [*Reinforcement Learning, Self supervised Learning, Machine learning*].
A method for goal generation using dynamical distance functions thus automatically producing a curriculum.
- 2020 **Learning Robust World Models using Dynamical Distance Functions**, [*Self supervised Learning, Planning, Machine learning*].
A method uses dynamical distance functions as an intrinsic exploration strategy to learn robust world models without requiring access to any reward functions.
- 2020 **Guiding Safe Reinforcement Learning Policies Using Structured Language Constraints**, [*Reinforcement learning, Machine learning*].
A framework to train RL agents conditioned on constraints that are in the form of structured language (Accepted at SafeAI Workshop at AAAI 2020)
- 2019 **Improving Safety in Reinforcement Learning using Model-Based Architectures and Human Intervention**, [*Reinforcement learning, Machine learning*].
A hybrid architecture for reducing the human intervention time and improving safety by combining model-based and model-free methods. (Accepted at AAAI FLAIRS 2019)
- 2018 **Representation learning by solving auxiliary tasks on Xray images**, [*Computer vision, Machine learning*].
Learning image representations on unannotated Chest Xray images using the method described in Noroozi and Favaro. Achieved improvements in classification tasks compared to the baseline.
- 2017 **Colorizing monochrome photographs**, [*Computer vision, Machine learning*].
Using deep neural networks to colorize monochrome photos. Course project, Data Science.
- 2014 **Performance Analysis of LSA and PLSA**.
Latent semantic analysis and Probabilistic latent semantic analysis on social media data. (MS thesis)

Relevant Coursework

Undergraduate

Operating Systems, Databases, Data Structures, Programming Languages, Computer Networks, Distributed Systems, Theory of Computation, Calculus

Graduate

Advanced Operating Systems, Design and Analysis of Algorithms, Advanced Computer Architecture, Artificial Intelligence, Machine Learning, Introduction to Computer Vision, Introduction to Natural Language Processing, Data Science, Graphical and Statistical Models of Learning

Mentorship, Leadership Experience and Academic Service

- 2018-2023 Reviewer - AI-HRI 2021, AI-HRI 2022, IJCAI 2023, NeurIPS 2023, AAAI 2024, LangRob @ CoRL 2023, FMDM @ NeurIPS 2023
- Research Mentorship - Tejasvini Manjunath (UMBC), Haoran Ren (UMBC, Amazon), Mohit Khatwani (UMBC, Google)

Other Interests

Soccer (Football), Photography, Video Editing and Basic Animation, Urban Sketching, Volunteering at Soccer Without Borders