Anurag Koul

New York, USA | +1-541-908-5875 | hire@koulanurag.dev
Github.com/koulanurag| Linkedin.com/in/koulanurag| koulanurag.dev

Education:

Ph.D. in Computer Science (advised by Prof. Alan Fern) -- Sep 2016 - Sept 2022
Oregon State University, Oregon, USA -- 3.7/4.0
Bachelor of Computer Science -- Aug 2010 - Jul 2014
University of Mumbai, Maharashtra, India -- 70% (1st Class)

Work Experience :

PostDoctoral Researcher, **Microsoft Research, New York** -- Oct 2022 - working
- Reinforcement Learning research under John Langford

Research Intern, **Microsoft Research, New York** -- Jun 2022 - Sep 2022
- Safe model-based reinforcement Learning for system control

Research Intern, **Intel AI Labs, California (Part-time)** -- Oct 2020 - Mar 2021

Research Intern, **Intel AI Labs, California (Full-time)** -- Jun 2020 - Sep 2020
- Model-Based Reinforcement Learning for control tasks

Research Intern, **SAS, North Carolina, USA** -- May 2019 - Aug 2019
- Collaborative Multi-Agent Reinforcement Learning for Inventory Optimization

Graduate Research Assistant, **Oregon State University, USA** -- Jan 2017 - Jun 2022
- Working on Deep Reinforcement Learning, Planning, and Explainable AI

Senior Software Engineer, **Capgemini, India** -- Sep 2015 - Aug 2016

Software Engineer, **Capgemini, India** -- Aug 2014 - Sep 2015
- Full-stack web development of products in the telecommunication sector.
- Projects: Information Delivery Portal, Off-net Automation, Real-time Price engine

Research and Developer Intern, **QuantumVentura, India** -- Jan 2014 - Jul 2014
- Developed real-time twitter-stream analytics for sentiment and influence analysis

---

Research Papers:

- "PcLast: Discovering Plannable Latent States" --ICML, 2024
  Anurag Koul*, ShivaKanth*, Alex Lamb, Shaoru Chen, Ben Evans, Lili Wu, Byron Xu, Lekan Molu, Rajan Chari, Yonathan Efroni, Miro Dunik, John Langford
- "Offline Policy Comparison with Confidence: Benchmarks and Baselines" -- Offline RL workshop, NeurIPS 2022
  Anurag Koul, Mariano Phielipp**,** Alan Fern
- "Dream and Search to Control: Latent Space Planning for Continuous Control" -- DRL workshop, ICLR 2020
  Anurag Koul, Varun Kumar Vijay, Somdeb Majumdar, Alan Fern
- "Understanding Finite State Representations of Recurrent Policy Networks" -- ICML 2020
  Mohamad Danesh, **Anurag Koul**, Alan Fern, Saeed Khorram
- "Learning Finite State Representations of Recurrent Policy Networks" -- ICLR, 2019
  Anurag Koul, Sam Greydanus, Alan Fern
- "Explaining Reinforcement Learning via Reward Decomposition" -- XAI workshop, IJCAI, 2019
  Zoe Juozapaitis, **Anurag Koul**, Alan Fern, Martin Erwig, Finale Doshi-Velez
- "Visualizing and Understanding Atari Agents" -- ICML, 2018
  Sam Greydanus, **Anurag Koul**, Alan Fern
- "Explaining Deep Adaptive Programs via Reward Decomposition" -- XAI workshop, IJCAI, 2018
  Martin Erwig, Alan Fern, Magesh Murali, **Anurag Koul**

Active Research Directions:

- Large Language Models (LLM) for Planning:
  - Learn to search for policy improvement via ***"future as input"*** **to transformer** policies.
  - Understanding **Decision Transformers** for Offline RL.
- Efficient Planning:
  - Backward Planning for Credit Assignment
  - Planning in Combinatorial State Space Problems
  - **Hierarchical Abstraction** for planning in control
  - Learning **Factored Markov Decision Processes** for efficient planning

---

Past Projects:

- TD3 with Variable-Step Frequency: Learning the step frequency for control tasks.
- Model-Based Reinforcement Learning with Discrete State Space.
- Understanding uncertainty estimation and safe-policy improvement in model-based offline reinforcement learning
- Investigating Offline Reinforcement Learning for Real-World **Cassie Robot**
- Planning in Abstract Space: Learn policy over a fine-step model with planning from option model
- Adaptive PID controller: Investigating learning of adaptive PID attributes for control policy.
- Learning Finite Space Gated Recurrent Neural Networks.
- Policy Gradient with Reward Decomposition: variation of policy gradient to utilize factored rewards.
- Performance comparison of Deep Reinforcement Learning algorithms: DQN, DDQN, Dueling Architecture, and A3C tested over Atari.
- Reducing Ambiguity in Deep Neural Nets using Conformal Prediction: Minimization of output-set size of conformal prediction in deep neural networks leading to less uncertainty.
- Direct Perception for Autonomous Driving: Enhanced existing approach by capturing temporal features from the observation.

Technical Skills:

- *Programming Languages:* Python, Java, Ruby, Haskell, Rust, C++
- *Machine Learning Libraries:* Pytorch, Tensorflow, Keras
- *Web Development:* Html5, CSS3, JavaScript, Jquery, AngularJs, NodeJs

---

Academic Activities:

- Invited to give a talk on "Explaining RL Agents from the lens of perception, memory, and uncertainties." at XAI Workshop, AAAI,2024 --27th Feb 2024
- **Reviewer** in L4DC 2024, ICAPS 2023, ICML (2022, 2021, 2020), NeurIPS (2022, 2021), ICLR 2021, JAIR 2021, , AAAI 2021, IJCAI 2020, JMLR 2020
- Program Committee Member in "Explainable Agency in Artificial Intelligence" Workshop - AAAI 2022
- Attended "Deep Learning + Reinforcement Learning Summer School" 2021 by CIFAR -- Jul 26, 2021 - Jul 31, 2021
- Participated in Neuro-Symbolic Artificial Intelligence reading group -- Mar 2021 - May 2021
- Invited for a virtual talk on Explainability in Reinforcement Learning @ Queen's University, Canada --25th Feb 2021
- Co-President of Artificial Intelligence Graduate Student Association(AIGSA) --Jan 2020 - Apr 2022
- Organizer of Reinforcement Learning reading group -- Dec 2019 - Jan 2021

Ph.D. Coursework*:*

- Fall 2016 - Machine Learning, Algorithms
- Winter 2017 - Deep Learning, Computer Vision I
- Spring 2017 - Reinforcement Learning, Computer Vision II
- Fall 2017- Convex Optimization, Functional Programming
- Winter 2018 - Artificial Intelligence
- Spring 2018 - Theory of Computation
- Winter 2019 - Probabilistic Graphical Models
- Fall 2019 - Analysis of Algorithms
- Fall 2020 - Estimation, Filtering, and Detection
