

Ashish Gaurav

ashish.gaurav@uwaterloo.ca
scholar.google.ca/citations?user=5CqEn6YAAAAJ

RESEARCH SUMMARY

In recent years, **reinforcement learning** methods have achieved impressive results across a variety of domains. Reward design and specification is crucial to the success of these methods in practical settings. Improperly specified rewards can induce unintended behaviour that can, in turn, have disastrous consequences. To avoid the laborious & manual process of reward design, several **inverse reinforcement learning** methods have been proposed to learn a suitable reward from some given expert demonstrations, or from some human feedback. In practice, these methods are usually expensive for real world use cases. Moreover, in many cases, the agent might not just be maximizing some notion of reward, but also satisfying certain **constraints**. In domains like robotics, constraints are usually easier and more natural to work with. Therefore, my current research focuses on **understanding the viability of constraint based behaviour specification and proposing efficient and practically applicable algorithms for learning such constraints from expert demonstrations**. These constraints formulations are grounded in realistic assumptions (for example, expected or probabilistic constraints) compared to prior methods, and can recover arbitrarily complex specification from given expert behaviour. Moreover, these algorithms are applicable to several domains (robotics, autonomous driving).

Skills: Python, Pytorch, Keras, Tensorflow

EDUCATION

University of Waterloo, Canada Ph.D. Candidate, Computer Science <i>Supervisor:</i> Pascal Poupart (Professor, University of Waterloo)	<i>Sep 2020 - current</i> Cumulative GPA: 93.5%
University of Waterloo, Canada M. Math., Computer Science <i>Supervisor:</i> Krzysztof Czarnecki (Professor, University of Waterloo)	<i>Sep 2017 - Jan 2020</i> Cumulative GPA: 89.75%
Birla Institute of Technology, Mesra, India B. Engg., Computer Science First Class with Distinction	<i>Sep 2013 - Apr 2017</i> Absolute GPA: 8.95/10.0
Chiranjiv Bharati School, Palam Vihar, Gurugram, India Primary, Middle and High School	<i>Apr 2001 - Mar 2013</i>

EXPERIENCE

Consulting ML Researcher @ Kodem Law (<i>part-time, remote</i>) Developed and applied state-of-the-art information extraction techniques for document understanding.	<i>Nov 2023 - current</i>
Research Assistant @ Waterloo AI Group, University of Waterloo <i>Topics:</i> inverse reinforcement learning, constraint learning, causality Developed efficient and practically applicable techniques for constraint learning from expert demonstrations, including the ICL [8] (<i>patent under review</i>) & IPCL [9] algorithms and contributed towards the development of VICRL [7] algorithm and constraint learning benchmarks. These algorithms have been demonstrated on	<i>Sep 2020 - current</i>

challenging real world scenarios in autonomous driving (lane change and maintaining a gap in highway driving).

Research Intern @ Vector Institute

Jul 2020 - Aug 2020

Visiting Researcher @ Waterloo AI Group, University of Waterloo

March 2020 - Jun 2020

Topics: reinforcement learning, transformer networks, temporal credit assignment

Benchmarked methods for temporal credit assignment problem, which are used to understand critical decisions in expert data. For example, we applied these methods to ice-hockey games, where it may be desired to know the timesteps/actions that were necessary to the outcome of win/loss.

Research Assistant @ W.I.S.E. Lab, University of Waterloo

Sep 2017 - Feb 2020

Topics: reinforcement learning, continual learning, autonomous driving

Developed benchmarks and methods for safe reinforcement learning (Linear Temporal Logic based reward formulation) [3], transfer RL [6], planning for autonomous driving [2], continual learning [5, T1] and out-of-distribution detection [1, 4]. Also designed the “high-level behavior planning” module (written in ROS/C++) for autonomoose, Waterloo’s self driving car project, which led to two successful [autonomous driving demos](#).

Teaching Assistant & Instructional Apprentice @ University of Waterloo

Sep 2017 - current

Performed exam & assignment marking, proctoring and conducted tutorials:

CS105: Introduction to Computer Programming (Fall 2023)

CS116: Introduction to Computer Science 2 (Winter 2024)

CS135: Designing Functional Programs (Fall 2020)

CS136: Elementary Algorithm Design and Data Abstraction (Fall 2017, Winter 2019)

CS145: Designing Functional Programs, Advanced Version (Fall 2019)

CS486: Introduction to Artificial Intelligence (Winter 2018)

CS885: Reinforcement Learning (Fall 2021, Winter 2022, Fall 2022)

Subject Matter Expert @ University of Toronto (*part-time, remote*)

Jul 2018 - Oct 2018

Designed content, including presentations, structure and scripts for teaching modules in the [Coursera’s Self-Driving Cars Specialization](#), offered by Steven Waslander, Professor at University of Toronto (UTIAS).

Software Engineering Intern @ Google

May 2016 - Jul 2016

Location: Mountain View, United States

Host: Andrew de los Reyes, Software Engineer, Google

Developed sandboxing and container software for Chrome OS.

PUBLICATIONS

[9] Inverse Probabilistic Constraint Learning

Under review, 2024

[8] Learning Soft Constraints From Constrained Expert Demonstrations

A Gaurav, K Rez  e, G Liu, P Poupart

International Conference on Learning Representations, 2023 (spotlight)

Patent filed: [Systems and methods to learn constraints from expert demonstrations](#)

Also featured at: [Vector Institute’s Endless Summer School, 2023](#)

[7] Benchmarking Constraint Inference in Inverse Reinforcement Learning

G Liu, Y Luo, A Gaurav, K Rez  e, P Poupart

International Conference on Learning Representations, 2023 (poster)

[T1] Safety-Oriented Stability Biases for Continual Learning

A Gaurav
Master's Thesis, 2020

[6] Transfer RL for Autonomous Driving: From WiseMove to WiseSim

A Balakrishnan, J Lee, A Gaurav, K Czarnecki, S Sedwards
ACM Transactions on Modeling and Computer Simulation, 2021, Vol. 31, Issue 3

[5] Simple Continual Learning Strategies for Safer Classifiers

A Gaurav, S Vernekar, J Lee, V Abdelzad, K Czarnecki, S Sedwards
Workshop on AI Safety (SafeAI), AAAI, 2020

[4] Out-of-distribution Detection in Classifiers via Generation

S Vernekar, A Gaurav, V Abdelzad, T Denouden, R Salay, K Czarnecki
Safety & Robustness in Decision Making Workshop, NeurIPS, 2019

[3] WiseMove: A Framework to Investigate Safe Deep RL for Autonomous Driving

J Lee*, A Balakrishnan*, A Gaurav*, K Czarnecki, S Sedwards*
International Conference on Quantitative Evaluation of Systems, 2019

[2] Design Space of Behaviour Planning for Autonomous Driving

M Ilievski, S Sedwards, A Gaurav, A Balakrishnan, A Sarkar, J Lee, F Bouchard, R Iaco, K Czarnecki
ArXiv, 2019

[1] Analysis of Confident-Classifiers for Out-of-Distribution Detection

S Vernekar*, A Gaurav*, T Denouden, B Phan, V Abdelzad, R Salay, K Czarnecki
Safe Machine Learning Workshop, International Conference on Learning Representations, 2019

SELECTED COURSE REPORTS

Spotify-viz, an interactive visualization system (CS 889)

Bayesian neural networks, for adversarial attack and defense (CS698)

Causal discovery for small graphs, using probabilistic programming (CS 842)

A survey of causal discovery algorithms (CS 886)

Bootstrapped DQN with UCB on Flappy Bird game (STAT 946)

Capsule networks, reproducing Sabour et al. (2017) and critique (STAT 841)

SELECTED PROJECTS

Cite-with-JS, a simple citation setup in Javascript for writing scientific articles

WiseMove, a hierarchical RL framework to investigate safety

GridDriving, a Gym simulator for RL experiments, based on CarRacing-v0

SSH-Scan, an SSH-based vulnerability scanner, *part of*: Mozilla Winter of Security 2016 - 2017

Porting Perl for Native Client architecture, porting Perl, *part of*: Google Summer of Code 2015

OTHER ACTIVITIES

Reviewer

2023 - present

Served as reviewer for several machine learning conferences including NeurIPS 2023, ICLR 2024, AISTATS 2024 and ICML 2024 (*ongoing*).

Vice President, Education @ Data Science Club, University of Waterloo

Jan 2019 - Dec 2019

Conducted several talks and tutorials on machine learning frameworks, reinforcement learning, autonomous driving. Also conducted a reading group on reinforcement learning.

AWARDS, SCHOLARSHIPS & ACHIEVEMENTS

Student Research Grant, Vector Institute

2021, 2022, 2023

CS Graduate Excellence Award, University of Waterloo

Sep 2020 - Dec 2020

International Doctoral Student Award

Sep 2020 - current

NSERC CREATE for Cyber-Physical Systems, Student Grant

May 2018 - Dec 2019

University of Waterloo Graduate Scholarship

Sep 2017 - Aug 2018

International Masters Student Award

Sep 2017 - Dec 2019

Academic Scholar, Chiranjiv Bharati School, Palam Vihar, Gurugram, India

2009 - 2012