

# PRABHAT NAGARAJAN

PhD Student | University of Alberta | Edmonton, AB

Homepage: [prabhatnagarajan.com](http://prabhatnagarajan.com)

Google Scholar  $\diamond$  Github

## RESEARCH INTERESTS

---

Reinforcement Learning & Sequential Decision-making

## EDUCATION

---

**University of Alberta** (Edmonton, AB) 2021 - present  
Pursuing *Doctor of Philosophy* in **Computing Science** 2021 - present

- Advisors: Marlos C. Machado & Martha White
- GPA: 4.0

**The University of Texas at Austin** (Austin, TX) 2014 - 2018  
*Master of Science* in **Computer Science** 2018

- Thesis: *Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning*
- Committee: Peter Stone (advisor), Scott Niekum

*Bachelor of Science* in **Computer Science** (*high honors* - top 10%) 2018

*Bachelor of Science* in **Mathematics** (*high honors* - top 10%) 2017

- Specialization: Pure Mathematics

## PROFESSIONAL EXPERIENCE

---

**Microsoft Research** Feb. 2025 - May 2025  
*Research Intern* New York, NY

I will work with Alex Lamb & John Langford on language models.

**Sony AI America** Jan. 2023 - Aug. 2023  
*AI Research Intern* Remote

I worked with Dustin Morrill on algorithms for reward shaping.

**Preferred Networks** Sept. 2018 - May 2021  
*Engineer* Tokyo, Japan

Reinforcement Learning - Open source software and RL applications

**Facebook** May 2017 - Aug. 2017  
*Software Engineer Intern* New York, NY

Messenger Infrastructure.

**Microsoft** May 2016 - Aug. 2016  
*Software Engineer Intern* Redmond, WA

Visual Studio Team Services.

**Yahoo!** May 2015 - Aug. 2015  
*Technical Intern* Sunnyvale, CA

Ads Targeting.

## CONSULTING & PART-TIME EXPERIENCE

---

**Alberta Machine Intelligence Institute**

*ML Foundations Facilitator*

Deliver monthly educational courses on machine learning basics to industry students.

Edmonton, AB, Canada

July 2022 - Dec. 2022

**Awenyx**

Consulted on applications of deep RL to black-box optimization problems

November 2021

## PUBLICATIONS

---

### Journal Publications

- [1] Yasuhiro Fujita, **Prabhat Nagarajan**, Toshiki Kataoka, and Takahiro Ishikawa.  
*ChainerRL: A Deep Reinforcement Learning Library.*  
Journal of Machine Learning Research (JMLR). 22(77):1-14, April 2021.

### Conference Publications

- [1] Shin-ichi Maeda, Hayato Watahiki, Yi Ouyang, Shintaro Okada, Masanori Koyama, and **Prabhat Nagarajan**.  
*Reconnaissance for Reinforcement Learning with Safety Constraints.*  
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD) 2021.
- [2] Zhang-Wei Hong, **Prabhat Nagarajan**, and Guilherme J. Maeda  
*Periodic Intra-Ensemble Knowledge Distillation for Reinforcement Learning.*  
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD) 2021.
- [3] Yasuhiro Fujita, Kota Uenishi, Avinash Ummadisingu, **Prabhat Nagarajan**, Shimpei Masuda, and Mario Castro.  
*Distributed Reinforcement Learning of Targeted Grasping with Active Vision for Mobile Manipulators.*  
Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2020.
- [4] Daniel Brown, Wonjoon Goo, **Prabhat Nagarajan**, and Scott Niekum.  
*Extrapolating Beyond Suboptimal Demonstrations via Inverse Reinforcement Learning from Observations.*  
Proceedings of the 36th International Conference on Machine Learning (ICML), June 2019.

### Workshop Publications

- [1] Aaron Havens, Yi Ouyang, **Prabhat Nagarajan**, and Yasuhiro Fujita.  
*Learning Latent State Spaces for Planning through Reward Prediction.*  
In Workshop on Deep Reinforcement Learning at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), December 2019.
- [2] **Prabhat Nagarajan**, Garrett Warnell, and Peter Stone.  
*Deterministic Implementations for Reproducibility in Deep Reinforcement Learning.*  
In AAAI 2019 Workshop on Reproducible AI, January 2019.

- [3] **Prabhat Nagarajan**, Garrett Warnell, and Peter Stone.  
*The Impact of Nondeterminism on Reproducibility in Deep Reinforcement Learning.*  
In ICML Reproducibility in Machine Learning Workshop, July 2018.

## Theses

- [1] **Prabhat Nagarajan**.  
Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning.  
*Master's Thesis*, The University of Texas at Austin, August 2018

## SCHOLARSHIPS & AWARDS

---

- Alberta Innovates Graduate Student Scholarship (2024) 2024  
(Value: \$31000 CAD for one year)
- University of Alberta Graduate Recruitment Scholarship 2021  
(Value: \$5000 CAD)
- NDSEG Fellowship (*declined*) 2021  
(Value: \$43,200 USD per year for 3 years)

## TEACHING

---

### **The University of Alberta**, Edmonton, AB

- Teaching Assistant* Intro to Reinforcement Learning (CMPUT 365) Fall 2024  
*Teaching Assistant* Basics of Machine Learning (CMPUT 267) Fall 2022  
*Teaching Assistant* Intermediate Machine Learning (CMPUT 396) Fall 2021

### **The University of Texas at Austin**, Austin, TX

- Tutor* Discrete Mathematics for Computer Science (CS 311) Fall 2015  
*Tutor* Data Structures (CS 314) Fall 2015

## COURSE DEVELOPMENT

---

### **The University of Alberta**, Edmonton, AB

- Deep Reinforcement Learning (CMPUT 628) Winter 2025

## REVIEWING

---

### **Journal Reviewing**

- Journal of Machine Learning Research (JMLR) - 2021, 2022
- Neural Networks - 2024

### **Conference Reviewing**

- Conference on Neural Information Processing Systems (NeurIPS) - 2023
- International Conference on Machine Learning (ICML) - 2022, 2023, 2024
- International Conference on Learning Representations (ICLR) - 2024, 2025
- Reinforcement Learning Conference (RLC) - 2024
- AAAI Conference on Artificial Intelligence (AAAI) - 2025
- International Conference on Robotics and Automation (ICRA) - 2021

## Workshop Reviewing

- RSS 2020 Workshop on *Closing the Academia to Real-World Gap in Service Robotics*

## SERVICE

---

### Workshop Organization

*Saving the Phenomena of Minds* workshop at RLDM 2025  
*Finding the Frame* workshop at the Reinforcement Learning Conference (RLC) 2024

### The University of Texas at Austin

*MS Admissions Committee* 2018

## ADVISING

---

### Preferred Networks (Tokyo, Japan)

#### *Intern Advising*

- *Aaron Havens* - Model-based reinforcement learning from latent reward models
- *Zhang-Wei Hong* - Improving ensemble reinforcement learning with knowledge distillation. Resulted in an ECML 2021 paper.

## TALKS

---

*Revisiting Overestimation in Value-based Deep Reinforcement Learning*

*AI Seminar*

Edmonton, AB, Canada. April, 2024.

*Deterministic Implementations for Reproducibility in Deep Reinforcement Learning*

*AAAI 2019 Workshop on Reproducibility in AI*

Honolulu, HI, January, 2018.

## MEMBERSHIPS

---

- AAAI - Member (2018–2023)

## PERSONAL DETAILS

---

Citizenship: USA

Languages: English (Native), Japanese (JLPT N4)