

# ANDREW LI

✉ Email [andrewli@cs.toronto.edu](mailto:andrewli@cs.toronto.edu)  
🌐 Webpage [www.cs.toronto.edu/~andrewli/](http://www.cs.toronto.edu/~andrewli/)  
☎ Phone Number +1 (519) 992-6886  
🇨🇦 Citizenship Canada 🍁

## RESEARCH DIRECTION

---

My current research focuses on the intersection of **machine learning** (particularly **deep reinforcement learning**) and **symbolic reasoning** with the goal of building more reliable and transparent AI agents. My research to date covers the following topics:

- \* Formal logics as a zero-shot specification language for RL. (**C4, C6**)
- \* Formal languages in RL under uncertain semantic interpretations. (**C8**)
- \* Learning belief state representations from high-dimensional data for partially observable RL. (**C7**)
- \* Learning automata representations from high-dimensional data for classification. (**C3, W1**)
- \* Learning Bayesian networks from data. (**C1**)

## SKILLS

---

**Software Engineering:** Through multiple internships at top technology companies, I gained extensive experience in software development and algorithm design. I was previously recognized as one of Canada's leading competitors in informatics competitions.

**Machine Learning:** I have over 5 years of academic experience publishing in top ML venues and training ML models in frameworks such as PyTorch, TensorFlow, and, Theano.

**Reinforcement Learning:** I have real-world industry experience training large-scale RL models for autonomous driving, in addition to years of academic experience developing RL algorithms. I am familiar with modern RL best practices, including multi-process and GPU computing, and JAX optimization.

## EDUCATION

---

**University of Toronto** *2021 - 2025 (anticipated)*  
Candidate for *PhD* in Computer Science  
Advisor: Sheila A. McIlraith  
(GPA: 3.96/4.0)

**University of Toronto** *2019 - 2021*  
*MSc* in Computer Science  
Advisor: Sheila A. McIlraith  
(GPA: 4.0/4.0)

**University of Waterloo** *2015 - 2019*  
*BMath* in Computer Science and Combinatorics & Optimization (Double Major)  
(GPA: 94.5/100)

## CONFERENCE PUBLICATIONS (REFEREED)

---

(\*) denotes equal first-author contribution.

- C8** Andrew C. Li, Zizhao Chen, Toryn Q. Klassen, Pashootan Vaezipoor, Rodrigo Toro Icarte, Sheila A. McIlraith. *Reward Machines for Deep RL in Noisy and Uncertain Environments*. 38th Conference on Neural Information Processing Systems (**NeurIPS**). Vancouver, British Columbia Canada, December 2024.
- C7** Andrew Wang\*, Andrew C. Li\*, Rodrigo Toro Icarte, Toryn Q. Klassen, Sheila A. McIlraith. *Learning Belief Representations for Partially Observable Deep RL*. 40th International Conference on Machine Learning (**ICML**). Honolulu, Hawaii USA, July 2023.
- C6** Mathieu Tuli, Andrew C. Li, Pashootan Vaezipoor, Toryn Q. Klassen, Scott Sanner, Sheila A. McIlraith. *Learning to Follow Instructions in Text-based Games*. 36th Conference on Neural Information Processing Systems (**NeurIPS**). New Orleans, Louisiana USA, November 2022.
- C5** Pashootan Vaezipoor\*, Andrew C. Li\*, Rodrigo Toro Icarte, Sheila A. McIlraith. *Achieving Zero-Shot Task Generalization with Formal Language Instructions*. 5th Multi-disciplinary Conference on Reinforcement Learning and Decision Making (**RLDM**). Providence, Rhode Island USA, June 2022.
- C4** Pashootan Vaezipoor\*, Andrew C. Li\*, Rodrigo Toro Icarte, Sheila A. McIlraith. *LTL2Action: Generalizing LTL Instructions for Multi-Task RL*. 38th International Conference on Machine Learning (**ICML**). Virtual conference, July 2021.
- C3** Maayan Shvo, Andrew C. Li, Rodrigo Toro Icarte, Sheila A. McIlraith. *Interpretable Sequence Classification via Discrete Optimization*. 35th AAAI Conference on Artificial Intelligence (**AAAI**). Virtual conference, Feb. 2021.
- C2** Neda Paryab, Alexander Sachs, Andrew Li, Meiyappan Nagappan and Jesse Hoey. *Relating Values and Social Network Structure*. 5th International Conference on Computational Social Science (**IC2S2**). Amsterdam, The Netherlands, July 2019.
- C1** Andrew C. Li and Peter van Beek. *Bayesian Network Structure Learning with Side Constraints*. 9th International Conference on Probabilistic Graphical Models (**PGM**). Prague, Czech Republic, Sept. 2018.

## WORKSHOP PRESENTATIONS (REFEREED)

---

(\*) denotes equal first-author contribution.

- W8** Andrew C. Li, Zizhao Chen, Toryn Q. Klassen, Pashootan Vaezipoor, Rodrigo Toro Icarte, Sheila A. McIlraith. *Reward Machines for Deep RL in Noisy and Uncertain Environments*. Robotic Tasks and How to Specify Them? — Task Specification for General-Purpose Intelligent Robots, at RSS. Delft, Netherlands, July 2024.
- W7** Pashootan Vaezipoor\*, Andrew C. Li\*, Rodrigo Toro Icarte, Sheila A. McIlraith. *LTL2Action: Generalizing LTL Instructions for Multi-Task RL*. AAAI Symposium Series (On the Effectiveness of Temporal Logics on Finite Traces in AI). San Francisco, California USA, March 2023. **Best Poster Award**.
- W6** Mathieu Tuli, Andrew C. Li, Pashootan Vaezipoor, Toryn Q. Klassen, Scott Sanner, Sheila A. McIlraith. *Learning to Follow Instructions in Text-based Games*. AAAI Symposium Series (On the Effectiveness of Temporal Logics on Finite Traces in AI). San Francisco, California USA, March 2023.
- W5** Andrew C. Li\*, Zizhao Chen\*, Pashootan Vaezipoor, Toryn Q. Klassen, Rodrigo Toro Icarte, Sheila A. McIlraith. *Noisy Symbolic Abstractions for Deep RL: A case study with Reward Machines*. Deep Reinforcement Learning Workshop, at NeurIPS. New Orleans, Louisiana USA, November 2022.
- W4** Andrew C. Li, Pashootan Vaezipoor, Rodrigo Toro Icarte, Sheila A. McIlraith. *Exploring Long-Horizon Reasoning with Deep RL in Combinatorially Hard Tasks*. Decision Awareness in Reinforcement Learning Workshop, at ICML. Baltimore, Maryland USA, July 2022.
- W3** Mathieu Tuli, Andrew C. Li, Pashootan Vaezipoor, Toryn Q. Klassen, Scott Sanner, Sheila A. McIlraith. *Instruction Following in Text-Based Games*. 3rd Wordplay: When Language Meets Games Workshop, at NAACL. Seattle, Washington USA, July 2022.

**W2** Maayan Shvo, **Andrew C. Li**, Rodrigo Toro Icarte, Sheila A. McIlraith. *Interpretable Sequence Classification via Discrete Optimization (Abridged Report)*. 4th Knowledge Representation and Reasoning Meets Machine Learning Workshop (KR2ML 2020), at NeurIPS. Virtual workshop, Dec. 2020.

**W1** Phillip J.K. Christofferson, **Andrew C. Li**, Rodrigo Toro Icarte, Sheila A. McIlraith. *Learning Symbolic Representations for Reinforcement Learning of Non-Markovian Behavior*. 4th Knowledge Representation and Reasoning Meets Machine Learning Workshop (KR2ML 2020), at NeurIPS. Virtual workshop, Dec. 2020.

## WORK EXPERIENCE

---

### Planning and Control PhD Intern

*Foster City, California*

Zoox Inc.

*Summer 2023*

Developed reinforcement learning algorithms for autonomous vehicle control and closed-loop metrics for their evaluation.

### Teaching Assistant

*Toronto, Ontario*

University of Toronto

\* CSC236 - Introduction to Theory of Computation

*Fall 2022*

\* CSC2515 - Introduction to Machine Learning

*Fall 2021*

\* CSC236 - Introduction to Theory of Computation

*Fall 2020*

\* CSC373 - Algorithm Design, Analysis & Complexity

*Winter 2020*

\* CSC236 - Introduction to Theory of Computation

*Fall 2019*

### Undergraduate Researcher

*Waterloo, Ontario*

University of Waterloo

\* Supervisor: Jesse Hoey

*Fall 2018*

\* Supervisor: Peter van Beek

*Winter 2018*

\* Supervisor: Justin Wan

*Fall 2017*

### Software Engineering Intern

*San Francisco, California*

Square Inc.

*Winter 2017*

Improved growth of the Appointments iOS user base by building a referrals feature in Swift and Objective-C. Reduced UI test running time by 10x by refactoring test infrastructure. Migrated large sections of the code base to Swift 3, supporting new releases of iOS.

### Software Engineering Intern

*San Francisco, California*

Square Inc.

*Summer 2016*

Improved user experience and increased revenue by building a Ruby-on-Rails feature enabling purchases directly in user onboarding. Added backend support to process credit card transactions without compromising sensitive user info. Condensed onboarding flow, generating a 5% absolute lift in new sign-ups, verified through A/B testing.

## PROFESSIONAL SERVICE

---

### Reviewing

\* NeurIPS 2022 (**Top 8% Reviewer**), 2024

\* ICML 2023

\* ICLR 2022

\* ML4OR Workshop at AAAI 2022

\* TaskSpec Workshop at RSS 2024.

## Mentorship

- \* Andrew Wang (Undergraduate) *Summer 2022 - Winter 2023*
- \* Zoe Chen (Undergraduate) *Summer 2022*
- \* Phillip Christofferson (Undergraduate) *Winter 2020 - Fall 2020*

## SCHOLARSHIPS AND AWARDS

---

- \* Ontario Graduate Scholarship (\$15,000) *Jan 2025*
- \* NSERC CGS-D (\$105,000) *2021-2024*
- \* NSERC CGS-M (\$17,500) *Sept. 2020*
- \* Vector Scholarship in Artificial Intelligence (\$17,500) *Sept. 2019*
- \* NSERC Undergraduate Research Award (\$4,500) *Jan. 2018*
- \* President's International Experience Award (\$1,500) *Jan. 2017*
- \* W.T. Tutte National Scholarship (\$20,000) *Sept. 2015*
  - Second largest of 15 national scholarships awarded to exceptional first-year math students.
- \* President's Scholarship of Distinction *Sept. 2015*
- \* Canadian Computing Olympiad *May 2014, 2015*
  - Invitational national-level algorithms competition for top 20 high school students in Canada.
  - *Bronze medal* (2014).
  - *Silver medal* (2015). 1st (out of 1472) in Canada on the qualifying round, 6th in Canada on the final round.