







RODRIGO TORO ICARTE

Contact Information

December 14th, 2020

	Full Name	Rodrigo Andrés Toro Icarte
	Institution	Department of Computer Science, University of Toronto
	Email	rntoro@cs.toronto.edu
	Website	www.cs.toronto.edu/~rntoro/
	Code	bitbucket.org/RToroIcarte/
	Mailing Address	1603-15 Dundonald Street, Toronto, ON M4Y 1K4, Canada

AFFILIATIONS

University of Toronto

Department of Computer Science.

Vector Institute for Artificial Intelligence

Vector Institute graduate student.

EDUCATION

Ph.D. in Computer Science

Aug., 2021

Thesis: *On learning, search, and knowledge: RL agents with reasoning capabilities.*

Supervisor: Sheila McIlraith

Department of Computer Science

University of Toronto (UofT)

MSc. in Computer Science

Dec., 2017

Thesis: *Teaching multiple tasks to an RL agent using LTL.*

Supervisor: Sheila McIlraith

Department of Computer Science

University of Toronto (UofT)

MSc. in Computer Science

Dec., 2015

Thesis: *How to exploit commonsense knowledge in learning-based image retrieval.*

Supervisors: Álvaro Soto and Jorge Baier

Department of Computer Science

Pontificia Universidad Católica de Chile (PUC)

Licentiate in the Engineering Sciences

Dec., 2013

Department of Computer Science

Pontificia Universidad Católica de Chile (PUC)

TEACHING EXPERIENCE

Course instructor

2014 - 2015

Introduction to Programming Languages

Department of Computer Science

Pontificia Universidad Católica de Chile

Student supervision

2020

Phillip Christoffersen, Final year project course

Fourth year computer science undergraduate student

University of Toronto

Student supervision 2019
Ethan Waldie, BAsC Thesis
Fourth year engineering science undergraduate student
University of Toronto

Student supervision 2019
Xi Yan, Final year project course
Fourth year computer science undergraduate student
University of Toronto

INTERNSHIPS

Samsung Research America – AI Center Toronto 2020
Knowledge Representation Intern
Computer Vision and Language Group
Supervisor: Dr. Allan Jepson

CONFERENCE PUBLICATIONS

M. Shvo, A. Li, R. Toro Icarte, S. McIlraith (2020). Interpretable Sequence Classification via Discrete Optimization. *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. **(to appear)**

L. Illanes, X. Yan, R. Toro Icarte, S. McIlraith (2020). Symbolic Plans as High-Level Instructions for Reinforcement Learning. In: *Proceedings of the 30th International Conference on Automated Planning and Scheduling (ICAPS)*, 540–550. **(Talk)**

R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, S. McIlraith (2019). Learning Reward Machines for Partially Observable Reinforcement Learning. In: *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, 15497–15508. **(Spotlight)**

R. Toro Icarte, L. Illanes, M. Castro, A. Cire, S. McIlraith, J.C. Beck (2019). Training Binarized Neural Networks using MIP and CP. In: *Proceedings of the 25th International Conference on Principles and Practice of Constraint Programming (CP)*, 401–417. **(Talk)**

A. Camacho, R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2019). LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 6065–6073. **(Talk)**

R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, S. McIlraith (2019). Searching for Markovian Subproblems to Address Partially Observable RL. In: *Proceedings of the 4th Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 22–26. **(Poster)**

L. Illanes, X. Yan, R. Toro Icarte, & S. McIlraith (2019). Symbolic Planning and Model-Free Reinforcement Learning: Training Taskable Agents. In: *Proceedings of the 4th Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 191–195. **(Poster)**

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. In: *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2112–2121. **(Long Talk)**

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Teaching Multiple Tasks to an RL Agent using LTL. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 452–461. **(Talk)**

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Advice-Based Exploration in Model-Based Reinforcement Learning. In: *Proceedings of the 31st Canadian Conference in Artificial Intelligence (Canadian AI)*, 72–83. **(Talk)**

R. Toro Icarte, J. Baier, C. Ruz, A. Soto (2017). How a General-Purpose Commonsense Ontology can Improve Performance of Learning-Based Image Retrieval. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 1283–1289. **(Talk)**

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2017). Using Advice in Model-Based Reinforcement Learning. In: *Proceedings of the 3rd Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 199–203. **(Poster)**

SELECTED WORKSHOP PUBLICATIONS

L. Illanes, X. Yan, R. Toro Icarte, S. McIlraith (2020). Symbolic Plans as High-Level Instructions for Reinforcement Learning (Abridged). In: *Proceedings of the Knowledge Representation & Reasoning Meets Machine Learning Workshop at NeurIPS 2020*. **(Best paper award)**

R. Toro Icarte, R. Valenzano, T. Klassen, P. Christoffersen, A. Farahmand, S. McIlraith (2020). The Act of Remembering: A Study in Partially Observable RL – RWRL Report. In: *Proceedings of the Challenges of Real World Reinforcement Learning Workshop at NeurIPS 2020*.

R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, S. McIlraith (2019). Learning Reward Machines for Partially Observable Reinforcement Learning. In: *Proceedings of the Optimization Foundations of Reinforcement Learning Workshop at NeurIPS 2019*.

R. Toro Icarte, L. Illanes, M. Castro, A. Cire, S. McIlraith, J.C. Beck (2019). Training Binarized Neural Networks using MIP and CP (Abridged Report). In: *Proceedings of the Workshop on Machine Learning with Guarantees at NeurIPS 2019*.

L. Illanes, X. Yan, R. Toro Icarte, S. McIlraith (2019). Leveraging Symbolic Planning Models in Hierarchical Reinforcement Learning. In: *Proceedings of the Knowledge Representation & Reasoning Meets Machine Learning workshop at NeurIPS 2019*.

A. Camacho, R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2019). LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In: *Proceedings of the Knowledge Representation & Reasoning Meets Machine Learning workshop at NeurIPS 2019*

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Teaching Multiple Tasks to an RL Agent using LTL. In: *Proceedings of the Learning by Instruction Workshop at NeurIPS 2018*.

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Advice-Based Exploration in Model-Based RL. In: *Proceedings of the Learning by Instruction Workshop at NeurIPS 2018*.

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2018). Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. In: *Proceedings of the 1st Workshop on Goal Specifications for Reinforcement Learning at ICML 2018*.

UNDER-REVIEW OR IN PREPARATION

R. Toro Icarte, T. Klassen, R. Valenzano, S. McIlraith (2020). Reward Machines: Exploiting Reward Function Structure in Reinforcement Learning. *ArXiv preprint arXiv:2010.03950*. Under review at the Journal of Artificial Intelligence Research (JAIR). **(Q1 Journal/under-review)**

R. Toro Icarte, R. Valenzano, T. Klassen, P. Christoffersen, A. Farahmand, S. McIlraith (2020). The Act of Remembering: A Study in Partially Observable Reinforcement Learning. *ArXiv preprint arXiv:2010.01753*. Under review at the 9th International Conference on Learning Representations (ICLR). **(AI Conference/under-review)**

R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, S. McIlraith. How to Learn Reward Machines in Fully-Observable and Partially-Observable Environments. To be submitted to the Journal of Artificial Intelligence (AIJ). **(Q1 Journal/in preparation)**

P. Sanabria, T. Tapia, R. Toro Icarte, A. Neyem. A Deep Reinforcement Learning Approach for Job Scheduling in Dew Computing. To be submitted to the Engineering Applications of Artificial Intelligence Journal. **(Q1 Journal/in preparation)**

GRANTS, HONORS & AWARDS

Best paper award at the KR2ML workshop at NeurIPS20 (sponsored by Amazon Science)	2020
Spotlight presentation at NeurIPS19 (Top 12% of accepted papers)	2019
Long-talk presentation at ICML18 (Top 32% of accepted papers)	2018
2018/19 Vector Institute Postgraduate Affiliate Program.	2018
Becas Chile, international Ph.D. scholarship	2017
Becas Chile, international M.Sc. scholarship	2015
Becas Chile, national M.Sc. scholarship	2014
Highest ranking among my graduating class (558 students)	2013
Engineering honor scholarship at PUC	2012

SELECTED INVITED TALKS

R. Toro Icarte. Towards Reinforcement Learning Agents with Reasoning Capabilities. Invited talk at the *Big AI Lab*, Brown University, 2020.

R. Toro Icarte. An Introduction to Reinforcement Learning and Reward Machines. Invited talk at the *Center for Mathematical Modeling*, Universidad de Chile, 2020.

R. Toro Icarte. Training Binarized Neural Networks using MIP and CP. Invited talk at the *Conference on Data Science and Optimization*, 2019.

R. Toro Icarte. Using Reward Machines for High-Level Task Specification and Decomposition in RL. Invited talk at the *Artificial Intelligence Lab*, Universidad Católica de Chile, 2018.

SELECTED CONTRIBUTED TALKS

R. Toro Icarte. Learning Reward Machines for Partially Observable Reinforcement Learning. Recent published research track at the *17th International Conference on Principles of Knowledge Representation and Reasoning*, KR 2020.

R. Toro Icarte. Learning Reward Machines for Partially Observable Reinforcement Learning. Spotlight presentation at the *33rd Conference on Neural Information Processing Systems*, NeurIPS 2019.

R. Toro Icarte. Learning Reward Machines for Partially Observable Reinforcement Learning. Spotlight presentation at the *Optimization Foundations of Reinforcement Learning Workshop*, NeurIPS 2019.

R. Toro Icarte. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. Long-talk presentation at the *7th Goal Reasoning Workshop*, ACS 2019.

R. Toro Icarte. Advising and Instructing Reinforcement Learning Agents with LTL and Automata. Long-talk presentation at the *Reasoning about Actions and Processes: Highlights of Recent Advances Workshop*, ICAPS 2019.

R. Toro Icarte. Training Binarized Neural Networks using MIP and CP. Long-talk presentation at the *25th International Conference on Principles and Practice of Constraint Programming*, CP 2019.

R. Toro Icarte. Teaching Multiple Tasks to an RL Agent using LTL. Long-talk presentation at the *Learning by Instruction Workshop*, NeurIPS 2018.

R. Toro Icarte. Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. Long-talk presentation at the *35th International Conference on Machine Learning*, ICML 2018.

R. Toro Icarte. Teaching Multiple Tasks to an RL Agent using LTL. Long-talk presentation at the *17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS 2018.

R. Toro Icarte. How a General-Purpose Commonsense Ontology can Improve Performance of Learning-Based Image Retrieval. Long-talk presentation at the *26th International Joint Conference on Artificial Intelligence*, IJCAI 2017.

OTHER

Computer skills

Programming languages: Python, Java, C#, C/C++, and Matlab.

Web and mobile: Ruby on Rails, Android, JavaScript, PostgreSQL, HTML, and PHP.

Machine learning: TensorFlow, PyTorch, Keras, and Scikit-Learn.

Optimization: Gurobi, Cplex, and CP Optimizer.

Open source code: bitbucket.org/RToroIcarte/ and github.com/RodrigoToroIcarte/

Service

PC member: NeurIPS 2020.

Sub-reviewer: ICAPS 2017, CDC 2019, and ICRA 2020.

Languages

Spanish (native) and English (fluent).

REFERENCES

Sheila A. McIlraith

Professor and Canada CIFAR AI Chair

Department of Computer Science, University of Toronto

Email address: sheila@cs.toronto.edu

Michael L. Littman

Professor

Department of Computer Science, Brown University

Email address: michael_littman@brown.edu

Jorge A. Baier

Associate Professor

Department of Computer Science, Pontificia Universidad Católica de Chile

Email address: jabaier@ing.puc.cl

Allan Jepson

Vice President and Chief Scientist

Samsung Research America – AI Center Toronto

Email address: allan.jepson@samsung.com

Amir-massoud Farahmand

Faculty Member and Canada CIFAR AI Chair

Vector Institute for Artificial Intelligence

Email address: farahmand@vectorinstitute.ai

Richard Valenzano

Applied Research Scientist

Element AI

Email address: rick.valenzano@elementai.com