

---

CONTACT INFORMATION	University of Alberta Edmonton, AB Canada T5J 3B1	✉ <a href="mailto:marlos.cholodovskis@gmail.com">marlos.cholodovskis@gmail.com</a> 🌐 <a href="http://mcmachado.info">http://mcmachado.info</a> 📄 <a href="#">mcmachado</a>
---------------------	---	--

---

RESEARCH INTERESTS     **Machine learning.** Specifically, reinforcement learning, representation learning, and real-world applications.

RESEARCH POSITIONS     **Fellow in Residence**     2023 – Present  
Alberta Machine Intelligence Institute (Amii)

**Adjunct Professor**     2021 – Present  
*Canada CIFAR AI Chair*  
Alberta Machine Intelligence Institute (Amii) Fellow  
Department of Computing Science, University of Alberta

**Senior Research Scientist**     2021 – 2023  
DeepMind

**Research Scientist**     2019 – 2021  
Google Research, Brain Team

EDUCATION     Doctor of Philosophy in Computing Science     2013 – 2019  
*University of Alberta, Canada*  
Advisors: Michael Bowling and Marc G. Bellemare

Master of Science in Computing Science     2011 – 2013  
*Universidade Federal de Minas Gerais, Brazil*  
Advisors: Luiz Chaimowicz and Gisele L. Pappa

Bachelor of Science in Computer Science with First Class Honors     2006 – 2010  
*Universidade Federal de Minas Gerais, Brazil*

PUBLICATIONS     **Preprints**

[CoRR-3] Z. Abbas, R. Zhao, J. Modayil, A. White, M. C. Machado. Loss of Plasticity in Continual Deep Reinforcement Learning. *CoRR abs 2303.07507*, 2023. [Under review at CoLLAs].

[CoRR-2] H. Wang, E. Miah, M. White, M. C. Machado, Z. Abbas, R. Kumaraswamy, V. Liu, A. White. Investigating the Properties of Neural Network Representations in Reinforcement Learning. *CoRR abs 2203.15955*, 2022. [Under review at AIJ – Submitted on Feb/23]

[CoRR-1] R. S. Sutton, M. C. Machado, G. Z. Holland, D. Szepesvari, F. Timbers, B. Tanner, A. White. Reward-Respecting Subtasks for Model-Based Reinforcement Learning. *CoRR abs 2202.03466*, 2022. [Under review at AIJ – Submitted on Nov/22. Current status: “Accept after minor revision”].

### Journal Articles

[TMLR-1] R. Y. Tao, A. White, M. C. Machado. Agent-State Construction with Auxiliary Inputs. *Transactions on Machine Learning Research (TMLR)*, 2023.

[JMLR-2] M. C. Machado, A. Barreto, D. Precup, M. Bowling. Temporal Abstraction in Reinforcement Learning with the Successor Representation. *Journal of Machine Learning Research (JMLR)*, 24(80):1–69, 2023.

[Nature-1] [Alphabetical order] M. G. Bellemare, S. Candido, P. S. Castro, J. Gong, M. C. Machado, S. Moitra, S. Ponda, Z. Wang. Autonomous Navigation of Stratospheric Balloons using Reinforcement Learning. *Nature* 588:77–82, 2020.

[JAIR-1] M. C. Machado, M. G. Bellemare, E. Talvitie, M. J. Hausknecht, M. Bowling. Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents. *Journal of Artificial Intelligence Research (JAIR)* 61:523–562, 2018.

[JMLR-1] H. van Seijen, A. R. Mahmood, P. M. Pilarski, M. C. Machado, R. S. Sutton. True Online Temporal-Difference Learning. *Journal of Machine Learning Research (JMLR)* 17(145):1–40, 2016.

[CiE-1] R. L. F. Cunha, M. C. Machado, L. Chaimowicz. RTSmate: Towards an Advice System for RTS Games. *ACM Computers in Entertainment (ACM CIE)*, 12(1):1–20, 2014.

### Refereed Conference Articles

[ICML-4] B. Daley, M. White, C. Amato, M. C. Machado. Trajectory-Aware Eligibility Traces for Off-Policy Reinforcement Learning. *International Conference on Machine Learning*, 2023. [27.9% accept. rate]

[ICML-3] M. Klissarov, M. C. Machado. Deep Laplacian-based Options for Temporally-Extended Exploration. *International Conference on Machine Learning*, 2023. [27.9% accept. rate]

[UAI-1] A. Erraqabi, M. C. Machado, M. Zhao, S. Sukhbaatar, A. Lazaric, L. Denoyer, Y. Bengio: Temporal Abstractions-Augmented Temporally Contrastive Learning: An Alternative to the Laplacian in RL. *Conference on Uncertainty in Artificial Intelligence*, 2022. [32.3% accept. rate]

[AISTATS-1] S. Vaswani, O. Bachem, S. Totaro, R. Müller, S. Garg, M. Geist, M. C. Machado, P. S. Castro, N. Le Roux. A General Class of Surrogate Functions for Stable and Efficient Reinforcement Learning. *International Conference on Artificial Intelligence and Statistics*, Oral, 2022. [29.2% overall accept. rate, 2.6% spotlight accept. rate, 0.2% best paper nominee]

[ICML-2] [Double 1st author] W. Chung, V. Thomas, M. C. Machado, N. Le Roux. Beyond Variance Reduction: Understanding the True Impact of Baselines on Policy Optimization. *International Conference on Machine Learning*, 2021. [21.5% accept. rate]

[ICLR-4] R. Agarwal, M. C. Machado, P. S. Castro, M. G. Bellemare. Contrastive Behavioral Similarity Embeddings for Generalization in Reinforcement Learning. *International Conference on Learning Representations*, Spotlight, 2021. [28.7% overall accept. rate, 5.6% spotlight accept. rate]

[NeurIPS-1] D. Ghosh, M. C. Machado, N. Le Roux. An Operator View of Policy Gradient Methods. *Neural Information Processing Systems*, 2020. [20.1% accept. rate]

[AAAI-1] M. C. Machado, M. G. Bellemare, and M. Bowling. Count-Based Exploration with the Successor Representation. *AAAI Conference on Artificial Intelligence*, 2020. [20.6% accept. rate]

[ICLR-3] Y. Jinnai, J. W. Park, M. C. Machado, and G. Konidaris. Exploration in Reinforcement Learning with Deep Covering Options. *International Conference on Learning Representations*, 2020. [26.5% accept. rate]

[ICLR-2] A. A. Taiga, W. Fedus, M. C. Machado, A. Courville, M. G. Bellemare. On Bonus Based Exploration Methods In The Arcade Learning Environment. *International Conference on Learning Representations*, 2020. [26.5% accept. rate]

[ICLR-1] M. C. Machado, C. Rosenbaum, X. Guo, M. Liu, G. Tesauro, and M. Campbell. Eigenoption Discovery through the Deep Successor Representation. *International Conference on Learning Representations*, 2018. [36.0% accept. rate]

[IROS-1] C. Sherstan, M. C. Machado, P. Pilarski. Accelerating Learning in Constructive Predictive Frameworks with the Successor Representation. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018. [46.7% accept. rate]

[ICML-1] M. C. Machado, M. G. Bellemare, M. Bowling. A Laplacian Framework for Option Discovery in Reinforcement Learning. *International Conference on Machine Learning*, 2017. [25.4% accept. rate]

[AAMAS-1] Y. Liang, M. C. Machado, E. Talvitie, M. Bowling. State of the Art Control of Atari Games Using Shallow Reinforcement Learning. *International Conference on Autonomous Agents and Multiagent Systems*, 2016. [24.9% accept. rate, 0.7% best paper nominee]

[AGI-1] C. Sherstan, A. White, M. C. Machado, P. Pilarski. Introspective Agents: Confidence Measures for General Value Functions. *Conference on Artificial General Intelligence*, 2016. [38.8% accept. rate]

[CIG-1] M. C. Machado, G. L. Pappa, L. Chaimowicz. A Binary Classification Approach for Automatic Preference Modeling of Virtual Agents in Civilization IV. *IEEE Conference on Computational Intelligence and Games*, 2012. [51.3% accept. rate]

[SBGames-3] M. C. Machado, G. L. Pappa, L. Chaimowicz. Characterizing and Modeling Agents in Digital Games. *Brazilian Symposium on Computer Games and Digital Entertainment*, 2012. [54% accept. rate]

[CGames-1] M. C. Machado, E. P. C. Fantini, L. Chaimowicz. Player Modeling: Towards a Common Taxonomy. *International Conference on Computer Games*, 2011. [75% accept. rate]

[SBGames-2] M. C. Machado, B. S. L. Rocha, L. Chaimowicz. Agents Behavior and Preferences Characterization in Civilization IV. *Brazilian Symposium on Computer Games and Digital Entertainment*, 2011. [49.2% accept. rate]

[SBGames-1] M. C. Machado, L. Chaimowicz. Combining Metaheuristics and CSP Algorithms to solve Sudoku. *Brazilian Symposium on Computer Games and Digital Entertainment*, 2011. [49.2% accept. rate]

### **Selected Magazine Articles, Extended Abstracts, and Workshop Papers**

[Earlier versions of conference papers that were also presented at workshops are omitted]

[WS-4] J. Farebrother, M. C. Machado, M. Bowling. Generalization and Regularization in DQN. *NeurIPS Deep Reinforcement Learning Workshop & 4th Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2018.

[EA-1] M. C. Machado, M. G. Bellemare, E. Talvitie, M. J. Hausknecht, M. Bowling. Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents (Extended Abstract). *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018. [Invited paper]

[WS-3] M. Liu, M. C. Machado, G. Tesauro, M. Campbell. The Eigenoption-Critic Framework. *NeurIPS Workshop on Hierarchical Reinforcement Learning*, 2016.

[WS-2] M. C. Machado, M. Bowling. Learning Purposeful Behaviour in the Absence of Rewards. *ICML Workshop on Abstraction in Reinforcement Learning*, 2016.

[M-1] S. V. Albrecht, J. Christopher L., D. L. Buckeridge, A. Botea, C. Caragea, C. H. Chi, T. Damoulas, B. N. Dilkina, E. Eaton, P. Fazli, S. Ganzfried, M. Lindauer, M. C. Machado, Y. Malitsky, G. Marcus, S. Meijer, F. Rossi, A. Shaban-Nejad, S. Thiébaux, M. M. Veloso, T. Walsh, C. Wang, J. Zhang, Y. Zheng. Reports from the 2015 AAI Workshop Program. *AI Magazine* 36(2): 90-101, 2015.

[WS-1] M. C. Machado, S. Srinivasan, M. Bowling. Domain-Independent Optimistic Initialization for Reinforcement Learning. *AAAI Workshop on Learning for General Competency in Video Games*, 2015.

### **Patents**

[P-2] R. Agarwal, M. C. Machado, P. S. Castro, M. G. Bellemare. Contrastive Behavioral Similarity Embeddings for Generalization in Reinforcement Learning. US Patent App. 17/487,769, 2023.

[P-1] S. Candido, J. Gong, M. G. Bellemare, M. C. Machado. Systems and Methods for Navigating Aerial Vehicles Using Deep Reinforcement Learning. US Patent App. 16/667,424, 2021.

### Theses

[T-2] M. C. Machado. Efficient Exploration in Reinforcement Learning through Time-Based Representations. Ph.D. thesis, University of Alberta, 2019.

[T-1] M. C. Machado. A Methodology for Player Modeling based on Machine Learning. M.Sc. thesis, Universidade Federal de Minas Gerais, 2013.

RESEARCH GRANTS	<b>NSERC Discovery Grant</b>	2023 – 2028
	\$205,000 over five years. Sole PI: “Integrated Architectures for State and Temporal Abstraction in Reinforcement Learning”.	
	<b>NSERC Discovery Launch Supplement</b>	2023 – 2024
	\$12,500 associated with “Integrated Architectures for State and Temporal Abstraction in Reinforcement Learning” (sole PI).	
	<b>Canada CIFAR AI Catalyst Grant</b>	2023 – 2025
\$100,000 over two years. Co-PI (w/ Samira E. Kahou and Ulrich Aïvodji): “Hiccups on the Road to Explainable Reinforcement Learning (XRL)”.		
	<b>Canada CIFAR AI Recruitment Chair</b>	2021 – 2025
\$500,000 over five years. Sole PI: “Discovering Temporal and Spatial Abstractions in Reinforcement Learning”.		

AWARDS AND HONORS	<b>Organizations</b>	
	<i>Canada CIFAR AI Chair</i>	2021
	CIFAR Canada	
	<i>Amii Fellow</i>	2021
Alberta Machine Intelligence Institute (Amii)		

### Paper Distinctions

– Conferences

*Best Paper Honorable Mention: AISTATS* 2022

A General Class of Surrogate Functions for Stable and Efficient Reinforcement Learning

*Best Paper Honorable Mention: AAMAS* 2016

State of the Art Control of Atari Games Using Shallow Reinforcement Learning

– Workshops

*Best Paper: ICML Workshop on Exploration in Reinforcement Learning* 2019

Benchmarking Bonus-Based Exploration Methods on the Arcade Learning Environment [Preliminary version of the work ICLR-2]

*Best Paper: ICML Workshop on Exploration in Reinforcement Learning* 2018

Count-Based Exploration with the Successor Representation [Preliminary version of the work AAAI-1]

**Reviewing / Area Chairing**

<i>Notable Area Chair</i>	2023
International Conference on Learning Representations (ICLR)	
<i>Top 9% Highest-Scoring Reviewer</i>	2022
International Conference on Learning Representations (ICLR)	
<i>Top 8% Highest-Scoring Reviewer</i>	2021
Conference on Neural Information Processing Systems (NeurIPS)	
<i>Top 10% Highest-Scoring Reviewer</i>	2020
Conference on Neural Information Processing Systems (NeurIPS)	
<i>Top 33% Highest-Scoring Reviewer</i>	2020
International Conference on Machine Learning (ICML)	
<i>Top 10% Highest-Scoring Reviewer</i>	2019
Conference on Neural Information Processing Systems (NeurIPS)	
<i>Top 10% Highest-Scoring Reviewer</i>	2018
Conference on Neural Information Processing Systems (NeurIPS)	
<i>Top 10 Reviewer Award</i>	2018
International Conference on Machine Learning (ICML)	
<i>Outstanding PC Member</i>	2016
International Joint Conference on Artificial Intelligence (IJCAI)	

**University**

<i>Nomination for Ph.D. Outstanding Thesis Award</i>	2019
University of Alberta	
<i>M.Sc. Early Achievement Award</i>	2012
Universidade Federal de Minas Gerais (UFMG)	
<i>B.Sc. First Class Honors</i>	2010
Universidade Federal de Minas Gerais (UFMG)	

**Scholarships**

<i>Provincial Alberta Innovates Technology Futures Scholarship</i>	2013 – 2018
\$126,000 over four years in Ph.D..	
<i>Brazilian Research Scholarship (CNPq)</i>	2007 – 2008
5,400 BRL over eighteen months in B.Sc..	

SUPERVISION

**Graduated Students (University of Alberta)**

– M.Sc. students

Erfan Miah (w/ Martha White)	2021 – 2022
<i>Feature Generalization in Deep RL: An Investigation into Representation Properties</i>	
Currently a researcher at the Alberta Machine Intelligence Institute (Amii).	
Ruo Yu (David) Tao (w/ Adam White)	2021 – 2022
<i>Agent-State Construction with Auxiliary Inputs</i>	
<b>Nominated for M.Sc. Outstanding Thesis Award</b>	
Currently a Ph.D. student at Brown University.	

**Current Students (University of Alberta)**

*Because of my adjunct status, my students are required to have another professor as co-supervisor.*

– Ph.D. students

Brett Daley (w/ Martha White) 2022 – present  
 Alex Lewandowski (w/ Dale Schuurmans) 2022 – present

– M.Sc. students

Edan Meyer (w/ Adam White) 2022 – present  
 Subhojeet Pramanik (w/ Adam White) 2022 – present  
 Diego Fernando Gomez Noriega (w/ Michael Bowling) 2023 – present

**Interns (DeepMind)**

Martin Klissarov, Ph.D., McGill University 2022

**Interns (Google Brain)**

Taylor W. Killian, Ph.D., University of Toronto 2020  
 Valentin Thomas (w/ Nicolas Le Roux), Ph.D., Université de Montréal 2019

**Interns (University of Alberta)**

Jesse Farebrother, B.Sc., University of Alberta 2018 – 2019  
 Nicolas Carion, M.Sc., École Normale Supérieure de Lyon 2015

**Ph.D. Supervisory Committee**

Abhishek Naik, University of Alberta 2021 – Present  
*(w/ R. Sutton, supervisor, and D. Schuurmans)*  
 Chen Ma, University of Alberta 2021 – Present  
*(w/ R. Sutton, supervisor, and M. White)*  
 Han Wang, University of Alberta 2022 – Present  
*(w/ M. White, and A. White, supervisors)*

**M.Sc. Thesis Examining Committee**

Archit Sakhadeo, University of Alberta 2021  
*No More Pesky Hyperparameters: Offline Hyperparameter Tuning For Reinforcement Learning*  
 Rohan Calum Nuttall, University of Alberta 2022  
*Uncertainty Methods in Active Reinforcement Learning*  
 Jiamin He, University of Alberta 2023  
*Consistent Emphatic Weightings for Off-Policy Reinforcement Learning*

RESEARCH  
 INTERNSHIPS

**DeepMind** 2018  
 Deep learning group w/ Vlad Mnih.  
**IBM Research – T.J. Watson Research Center** 2017  
 AI Foundations group w/ Gerald Tesauro and Murray Campbell.  
**Microsoft Research – New York Lab** 2016  
 ML group w/ Alekh Agarwal, Fernando Diaz, Miro Dudik, & Robert Schapire.  
**Vetta Labs LTDA** 2009 – 2010

TEACHING ASSISTANT EXPERIENCE	<b>CMPUT 366: Intelligent Systems</b>	2016
	<b>CMPUT 403: Practical Algorithmics</b>	2016
	<b>DCC 865: Design and Analysis of Algorithms</b>	2012
SOFTWARE ENGINEERING EXPERIENCE	<b>Avenue Code</b>	2013
	<b>Synergia: Engenharia de Software e Sistemas</b>	2011 – 2013
	<b>Ilusis Interactive Graphics</b>	2010 – 2011
SELECTED TALKS	<i>Representation-driven Option Discovery in Reinforcement Learning</i>	
	Microsoft Research – New York City, USA	Mar. 2023
	University of Alberta – Edmonton, Canada	Mar. 2023
	McGill University – Montréal, Canada	Feb. 2023
	<i>Temporal Abstraction in Reinf. Learning with the Successor Representation</i>	
	Keynote at the Brazilian Conference on Intelligent Systems	Nov. 2022
	Microsoft Workshop on Reinforcement Learning, Forwards and Backwards: Insights from Neuroscience	Oct. 2021
	Stanford University – Stanford, USA	Feb. 2020
	<i>Autonomous Navigation of Stratospheric Balloons using Reinforcement Learning</i>	
	Amii AI Meetup – Edmonton, Canada	Jun. 2022
	University of Alberta – Edmonton, Canada	Jan. 2021
	<i>How Atari Started the Golden Age of Reinforcement Learning</i>	
	• Amii’s AI Week (w/ Michael Bowling) – Edmonton, Canada	May 2022
	<i>An Operator View of Policy Gradient Methods</i>	
	University of Alberta – Edmonton, Canada	Nov. 2020
	DeepMind – London, UK	Oct. 2020
	<i>Purposeful Exploration in Reinforcement Learning</i>	
	Facebook AI Research – Montréal, Canada	Oct. 2018
	Google Brain – Montréal, Canada	Oct. 2018
	Microsoft Research – Montréal, Canada	Oct. 2018
	<i>Count-Based Exploration with the Successor Representation</i>	
	RLDM – Montréal, Canada	Jul. 2019
	ICML WS on Exploration in RL, Best paper – Stockholm, Sweden	Jul. 2018
	<i>Eigenoption Discovery through Diffusion Models of Information Flow</i>	
	McGill University – Montréal, Canada	Nov. 2017
	Microsoft Research – Montréal, Canada	Nov. 2017
	<i>Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents</i>	
	IJCAI, Journal track – Stockholm, Sweden	Jul. 2018



IJCAI WS on Computer Games (Invited) – Stockholm, Sweden Jul. 2018  
University of Alberta – Edmonton, Canada Oct. 2017

*A Laplacian Framework for Option Discovery in Reinforcement Learning*

ICML – Sydney, Australia Aug. 2017  
ICML WS on Abstractions in RL – Sydney, Australia Aug. 2017  
RLDM – Ann Arbor, USA Jun. 2017  
University of Alberta – Edmonton, Canada May 2017

*Exploration in Reinforcement Learning: The Quest for Purposeful Behavior*

Univ. Federal de Minas Gerais (UFMG) – Belo Horizonte, Brazil Dec. 2016

*The Arcade Learning Environment: What comes next?*

IJCAI WS on General Intelligence and Game-Playing Agents (Invited) – New York, USA Jul. 2016

SERVICE AND  
OUTREACH

**Journal Reviewer**

ACM Transactions on Autonomous and Adaptive Systems  
Adaptive Behavior  
IEEE Transactions on Computational Intelligence and AI in Games  
Journal of Artificial Intelligence Research  
Journal of Machine Learning Research (JMLR)  
Machine Learning  
Nature Reviews Psychology  
Transactions on Machine Learning Research (TMLR)

**Area-Chair/Meta-Reviewer**

International Conference on Learning Representations (ICLR) 2021, 2023  
Montreal AI Symposium (MAIS) 2020

**Program Committee**

International Conference on Machine Learning (ICML) 2018 – 2022  
Neural Information Processing Systems (NeurIPS) 2018 – 2021  
AAAI Conference on Artificial Intelligence (AAAI) 2018 – 2020  
International Conference on Learning Representations (ICLR) 2020, 2022  
International Joint Conference on Artificial Intelligence (IJCAI) 2016 – 2019

**Workshops Organized**

AAAI Workshop on Learning for General Competency in Video Games 2015

**Participation in Panels**

Microsoft Summit Workshop on Reinforcement Learning, Forwards and Backwards: Insights from Neuroscience 2021

w/ Nathaniel Daw (Princeton University), Sam Gershman (Harvard University), Kimberly Stachenfeld (DeepMind), Geoff Gordon (Microsoft Research & Carnegie Mellon University), and Ida Momennejad (Microsoft Research) as moderator.

RLDM Workshop on Modeling Inductive Biases in Reinforcement Learning 2019

w/ Anne Collins (University of California Berkeley), Todd Gureckis (NYU), Anna Harutyunyan (DeepMind), and Doina Precup (McGill & DeepMind) as moderator.

ICML Workshop on Exploration in Reinforcement Learning 2018

w/ Ian Osband (DeepMind), Martha White (University of Alberta), Finale Doshi-Velez (Harvard), and Benjamin Van Roy (Stanford) as moderator.

### Workshop (WS) Program Committee

NeurIPS WS on Lifelong Learning Machine Learning 2021

ICML WS on Lifelong Learning 2020

NeurIPS Reproducibility Challenge 2019

NeurIPS WS on Optimization Foundations for Reinforcement Learning 2019

Montreal AI Symposium 2019

ICML WS on Lifelong Learning: A Reinf. Learning Approach 2019

ICML WS on Lifelong Learning: A Reinf. Learning Approach 2018

AAMAS WS on Adaptive Learning Agents (ALA) 2018

NeurIPS WS on Hierarchical Reinforcement Learning 2017

### University

President, Computing Science Graduate Student Association 2015 – 2016

CODE **Arcade Learning Environment versions 0.5 – 0.6** 2015 – 2017

RELEASED LANGUAGES: C++ AND PYTHON. Multiple versions of the ALE, including the introduction of modes and stochasticity in the environment, new functions, and a Python interface. This code was developed collaboratively.

**Source-code for multiple published papers** 2016 – 2019

LANGUAGES: C++ AND PYTHON. Source-code of several published papers, including True-Online Sarsa, Blob-PROST features, Eigenoptions,  $DQN_e^{MMC}$  + SR, and a gridworld library.

Last update: April 25, 2023.