

Marlos Cholodovskis Machado

Google Brain
Montréal, QC, Canada

✉ marlosm@google.com
🌐 <http://mcmachado.info>
👤 [mcmachado](#)

RESEARCH INTERESTS

My research interests lie broadly in *Artificial Intelligence* and particularly focus on *Reinforcement Learning*.

EDUCATION

<i>Ph.D. in Computing Science</i> University of Alberta, Canada Supervisors: Michael Bowling and Marc G. Bellemare	2013 – 2019
<i>M.Sc. in Computer Science</i> Universidade Federal de Minas Gerais, Brazil Supervisors: Luiz Chaimowicz and Gisele L. Pappa	2011 – 2013
<i>B.Sc. in Computer Science with First Class Honors</i> Universidade Federal de Minas Gerais, Brazil	2006 – 2010

PUBLICATIONS

Journal Articles

- [1] **M. C. Machado**, M. G. Bellemare, E. Talvitie, M. J. Hausknecht, and M. Bowling (2018): “Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents”, in *Journal of Artificial Intelligence Research (JAIR)* 61:523–562.
- [2] H. van Seijen, A. R. Mahmood, P. M. Pilarski, **M. C. Machado** and R. S. Sutton (2016): “True Online Temporal-Difference Learning”, in *Journal of Machine Learning Research (JMLR)* 17(145):1–40.
- [3] R. L. F. Cunha, **M. C. Machado** and L. Chaimowicz (2014): “RTSmate: Towards an Advice System for RTS Games”, in *ACM Computers in Entertainment*, 11(4):1–20.

Refereed Conferences

- [4] **M. C. Machado**, M. G. Bellemare, and M. Bowling (2020): “Count-Based Exploration with the Successor Representation”, in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. [Acceptance rate: ~20.6%]
- [5] Y. Jinnai, J. W. Park, **M. C. Machado**, and G. Konidaris (2020): “Exploration in Reinforcement Learning with Deep Covering Options”, in *Proceedings of the International Conference on Learning Representations (ICLR)*. [Acceptance Rate: ~26.5%].
- [6] A. A. Taiga, W. Fedus, **M. C. Machado**, A. Courville, M. G. Bellemare (2020): “On Bonus Based Exploration Methods In The Arcade Learning Environment”, in *Proceedings of the International Conference on Learning Representations (ICLR)*. [Acceptance Rate: ~26.5%].
- [7] **M. C. Machado**, C. Rosenbaum, X. Guo, M. Liu, G. Tesauro, and M. Campbell (2018): “Eigenoption Discovery through the Deep Successor Representation”, in *Proceedings of the International Conference on Learning Representations (ICLR)*. [Acceptance Rate: ~36%].
- [8] C. Sherstan, **M. C. Machado**, P. Pilarski (2018): “Accelerating Learning in Constructive Predictive Frameworks with the Successor Representation”, in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. [Acceptance Rate: ~46.7%].
- [9] **M. C. Machado**, M. G. Bellemare, and M. Bowling (2017): “A Laplacian Framework for Option Discovery in Reinforcement Learning”, in *Proceedings of the International Conference on Machine Learning (ICML)*. [Acceptance Rate: ~25%].

- [10] Y. Liang, **M. C. Machado**, E. Talvitie, and M. Bowling (2016): “State of the Art Control of Atari Games Using Shallow Reinforcement Learning”, in *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. [Acceptance Rate: $\sim 25\%$].
- [11] C. Sherstan, A. White, **M. C. Machado**, and P. Pilarski (2016): “Introspective Agents: Confidence Measures for General Value Functions”, in *Proceedings of the Conference on Artificial General Intelligence (AGI)*. [Acceptance Rate: $\sim 39\%$].
- [12] **M. C. Machado**, G. L. Pappa, and L. Chaimowicz (2012): “A Binary Classification Approach for Automatic Preference Modeling of Virtual Agents in Civilization IV”, in *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG)*. [Acceptance Rate: $\sim 51\%$].
- [13] **M. C. Machado**, G. L. Pappa, and L. Chaimowicz (2012): “Characterizing and Modeling Agents in Digital Games”, in *Proceedings of the Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*. [Acceptance Rate: $\sim 54\%$].
- [14] **M. C. Machado**, E. P. C. Fantini, and L. Chaimowicz (2011): “Player Modeling: Towards a Common Taxonomy”, in *Proceedings of the International Conference on Computer Games (CGames)*. [Acceptance Rate: $\sim 75\%$].
- [15] **M. C. Machado**, B. S. L. Rocha, and L. Chaimowicz (2011): “Agents Behavior and Preferences Characterization in Civilization IV”, in *Proceedings of the Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*. [Acceptance Rate: $\sim 49\%$].
- [16] **M. C. Machado**, and L. Chaimowicz (2011): “Combining Metaheuristics and CSP Algorithms to solve Sudoku”, in *Proceedings of the Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*. [Acceptance Rate: $\sim 49\%$].

Magazine Articles

- [17] 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. T. 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 (2015): “Reports from the 2015 AAAI Workshop Program”, in *AI Magazine* 36(2): 90-101.

Additional Extended Abstracts, Workshops, Symposia, and Tutorials

- [18] A.A. Taiga, W. Fedus, **M. C. Machado**, A. Courville, M. G. Bellemeare (2019): “Benchmarking Bonus-Based Exploration Methods on the Arcade Learning Environment”, in *ICML Workshop on Exploration in Reinforcement Learning*.
- [19] J. Farebrother, **M. C. Machado**, and M. Bowling (2018): “Generalization and Regularization in DQN”, in *NeurIPS Deep Reinforcement Learning Workshop*. Also presented at the *4th Multidisciplinary Conference on Reinforcement Learning and Decision Making*.
- [20] **M. C. Machado**, M. G. Bellemare, and M. Bowling (2018): “Count-Based Exploration with the Successor Representation”, in *ICML Workshop on Exploration in Reinforcement Learning*. Also presented at the *4th Multidisciplinary Conference on Reinforcement Learning and Decision Making*.
- [21] **M. C. Machado**, M. G. Bellemare, E. Talvitie, M. J. Hausknecht, and M. Bowling (2018): “Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents (Extended Abstract)”, in *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [22] M. Liu, **M. C. Machado**, G. Tesauro, and M. Campbell (2017): “The Eigenoption-Critic Framework”, in *NIPS Workshop on Hierarchical Reinforcement Learning*.
- [23] **M. C. Machado**, and M. Bowling (2016): “Learning Purposeful Behaviour in the Absence of Rewards”, in *ICML Workshop on Abstraction in Reinforcement Learning*.

[24] **M. C. Machado**, S. Srinivasan, and M. Bowling (2015): “Domain-Independent Optimistic Initialization for Reinforcement Learning”, in *AAAI Workshop on Learning for General Competency in Video Games*.

[25] **M. C. Machado**, E. P. C. Fantini, and L. Chaimowicz (2011): “Player Modeling: What is it? How to do it?”, in *Tutorials of the X Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*.

Theses

[26] “Efficient Exploration in Reinforcement Learning through Time-Based Representations“, *Ph.D.*, 2019.

[27] “A Methodology for Player Modeling based on Machine Learning”, *M.Sc.*, 2013.

[28] “Classification of Documents Modeled as Graphs” (*in Portuguese*), *B.Sc.*, 2010.

AWARDS

- Best Paper Award at the ICML Workshop on Exploration in Reinforcement Learning 2019
- NIPS Top 10% Highest-Scoring Reviewer 2019
- Best Paper Award at the ICML Workshop on Exploration in Reinforcement Learning 2018
- NIPS Top 10% Highest-Scoring Reviewer 2018
- ICML Top 10 Reviewer 2018
- Best Paper Award Nominee at the International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2016
- IJCAI Outstanding PC Member 2016
- Provincial Alberta Innovates Technology Futures Graduate Student Scholarship 2013 – 2018
- M.Sc. Early Achievement Award 2012
- IEEE CIS Outstanding Student Paper Travel Grant 2012
- B.Sc. First Class Honors 2010
- Brazil. Council for Scient. and Tech. Development Undergraduate Research Scholarship 2007 – 2008

RESEARCH EXPERIENCE

Google Brain – Montréal 2019

I am currently a Research Scientist at Google Brain.

DeepMind – Alberta/London 2018

I interned in the Deep Learning group with Vlad Mnih.

IBM Research – T.J. Watson Research Center 2017

I interned in the AI Foundations group with Gerald Tesauro and Murray Campbell.

Microsoft Research – New York Lab 2016

I interned in the ML group with Alekh Agarwal, Fernando Diaz, Miro Dudik, and Robert Schapire.

Vetta Labs LTDA 2009 – 2010

SOFTWARE ENGINEERING EXPERIENCE

Avenue Code 2013

Synergia: Engenharia de Software e Sistemas 2011 – 2013

Ilusis Interactive Graphics 2010 – 2011

TEACHING ASSISTANT EXPERIENCE

CMPUT 366: Intelligent Systems	2016
CMPUT 403: Practical Algorithmics	2016
DCC 865: Design and Analysis of Algorithms	2012

SUPERVISION EXPERIENCE**Students**

- Jesse Farebrother (B.Sc., University of Alberta) 2018 – 2019

Interns

- Valentin Thomas (Ph.D., Université de Montréal) 2019
- Nicolas Carion (M.Sc., École Normale Supérieure de Lyon) 2015

SELECTED TALKS*Temporal Abstraction in Reinforcement Learning with the Successor Representation*

- Stanford University – Stanford, USA Feb. 2020

Purposeful Exploration in Reinforcement Learning

- Facebook AI Research – Montréal, Canada Oct. 2018
- Google Brain – Montréal, Canada Oct. 2018
- Microsoft Research – Montréal, Canada Oct. 2018

Count-Based Exploration with the Successor Representation

- Multi-disciplinary Conference on RL and Decision Making Montréal, Canada Jul. 2019
- ICML Workshop on Exploration in Reinf. Learning (Best paper track) – Stockholm, Sweden Jul. 2018

Eigenoption Discovery through Diffusion Models of Information Flow

- McGill University – Montréal, Canada Nov. 2017
- Microsoft Research – Montréal, Canada Nov. 2017

Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents

- International Joint Conference on Artificial Intelligence (Journal track) – Stockholm, Sweden Jul. 2018
- IJCAI Workshop on Computer Games Workshop (Invited speaker) – Stockholm, Sweden Jul. 2018
- University of Alberta AI Seminar – Edmonton, Canada Oct. 2017

A Laplacian Framework for Option Discovery in Reinforcement Learning

- International Conference on Machine Learning – Sydney, Australia Aug. 2017
- ICML Workshop on Abstractions in Reinforcement Learning – Sydney, Australia Aug. 2017
- Multi-disciplinary Conference on RL and Decision Making – Ann Arbor, USA Jun. 2017
- University of Alberta AI Seminar – Edmonton, Canada May 2017

Exploration in Reinforcement Learning: The Quest for Purposeful Behavior

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

The Arcade Learning Environment: What comes next?

- IJCAI Wksp on General Intel. and Game-Playing Agents (Invited speaker) – New York, USA Jul. 2016

SERVICE AND OUTREACH**Journal Reviewer**

- Journal of Artificial Intelligence Research 2019 – 2020
- Journal of Machine Learning Research 2017 – 2018
- Adaptive Behavior 2017
- ACM Transactions on Autonomous and Adaptive Systems 2016
- IEEE Transactions on Computational Intelligence and AI in Games 2013 – 2017

Program Committee

- Neural Information Processing Systems (NeurIPS) 2018 – 2020
- International Conference on Machine Learning (ICML) 2018 – 2020
- International Conference on Learning Representations (ICLR) 2020
- AAAI Conference on Artificial Intelligence (AAAI) 2018 – 2020
- 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

- RLDM Workshop on Modeling Inductive Biases in Reinforcement Learning 2019
with Anne Collins (University of California Berkeley), Todd Gureckis (New York University),
Anna Harutyunyan (DeepMind), and Doina Precup (McGill University and DeepMind) as the moderator.
- ICML Workshop on Exploration in Reinforcement Learning 2018
with Ian Osband (DeepMind), Finale Doshi-Velez (Harvard University), Martha White
(University of Alberta), and Benjamin Van Roy (Stanford University) as the moderator.

Workshop Program Committee

- NeurIPS Reproducibility Challenge 2019
- NeurIPS Workshop on Optimization Foundations for Reinforcement Learning 2019
- Montreal AI Symposium 2019
- ICML Workshop on Lifelong Learning: A Reinforcement Learning Approach 2018 – 2019
- AAMAS Workshop on Adaptive Learning Agents (ALA) 2018
- NIPS Workshop on Hierarchical Reinforcement Learning 2017

University Service

- President, Computing Science Graduate Student Association 2015 – 2016