Marlos C. Machado

Curriculum vitae

CONTACT INFORMATION

DeepMind
Edmonton, AB
Canada T5J 3B1

✉ marlosm@deepmind.com
http://mcmachado.info
mcmachado

RESEARCH INTERESTS

Machine learning. Specifically, reinforcement learning, representation learning, optimization, and real-world applications.

FULL-TIME RESEARCH POSITIONS

Senior Research Scientist — Jul. 2021 – Present
Research Scientist — Jan. 2021 – Jul. 2021
DeepMind

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

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

[1] M. C. Machado, A. Barreto, D. Precup. Temporal Abstraction in Reinforcement Learning with the Successor Representation. *CoRR abs* 2110.05740, 2021.

Journal Articles

[2] [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.

[3] 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.

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

[5] R. L. F. Cunha, M. C. Machado, L. Chaimowicz. RTSmate: Towards an Advice System for RTS Games. *ACM Computers in Entertainment (ACM CIE)*, 11(4):1–20, 2014.

Refereed Conference Articles

[6] [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 (ICML)*, 2021. [21.5% accept. rate]

[7] 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 (ICLR)*, Spotlight, 2021. [28.7% overall accept. rate, 5.6% spotlight accept. rate]

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

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

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

[11] 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 (ICLR)*, 2020. [26.5% accept. rate]

[12] 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 (ICLR)*, 2018. [36.0% accept. rate]

[13] 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 (IROS)*, 2018. [46.7% accept. rate]

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

[15] 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 (AAMAS)*, 2016. [24.9% accept. rate]

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

[17] 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 (CIG)*, 2012. [51% accept. rate]

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

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

[20] 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 (SBGames)*, 2011. [49% accept. rate]

[21] M. C. Machado, L. Chaimowicz. Combining Metaheuristics and CSP Algorithms to solve Sudoku. *Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, 2011. [49% accept. rate]

Selected Magazine Articles, Extended Abstracts, and Workshop Papers
[Earlier versions of conference papers that were also presented at workshops are omitted]

[22] S. Vaswani, O. Bachem, S. Totaro, R. Müller, M. Geist, M. C. Machado, P. S. Castro, N. Le Roux. A Functional Mirror Ascent View of Policy Gradient Methods with Function Approximation. *ICML Workshop on Reinforcement Learning Theory*, 2021.

[23] A. Erraqabi, M. Zhao, M. C. Machado, Y. Bengio, S. Sukhbaatar, L. Denoyer, A. Lazaric. Exploration-Driven Representation Learning in Reinforcement Learning. *ICML Workshop on Unsupervised Reinforcement Learning*, 2021.

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

[25] 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]

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

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

[28] 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 AAAI Workshop Program. *AI Magazine* 36(2): 90-101, 2015.

[29] 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

[30] 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

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

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

Awards and Honors

Organizations

Canada CIFAR AI Chair 2021
CIFAR Canada

Amii Fellow 2021
Alberta Machine Intelligence Institute (Amii)

Paper Distinctions

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 by Taiga et al. (2020)]

Best Paper: ICML Workshop on Exploration in Reinforcement Learning 2018
Count-Based Exploration with the Successor Representation [Preliminary version of the work by Machado et al. (2020)]

Best Paper Nominee: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2016
State of the Art Control of Atari Games Using Shallow Reinforcement Learning

IEEE CIS Outstanding Student Paper Travel Grant 2012
A Binary Classification Approach for Automatic Preference Modeling of Virtual Agents in Civilization IV

Reviewing

Top 8% Highest-Scoring Reviewer 2021
Conference on Neural Information Processing Systems (NeurIPS)

Top 10% Highest-Scoring Reviewer 2020
Conference on Neural Information Processing Systems (NeurIPS)

Top 33% Highest-Scoring Reviewer 2020
International Conference on Machine Learning (ICML)

Top 10% Highest-Scoring Reviewer 2019
Conference on Neural Information Processing Systems (NeurIPS)

Top 10% Highest-Scoring Reviewer 2018
Conference on Neural Information Processing Systems (NeurIPS)

Top 10 Reviewer Award 2018
International Conference on Machine Learning (ICML)

Outstanding PC Member 2016
International Joint Conference on Artificial Intelligence (IJCAI)

University

Nomination for Ph.D. Outstanding Thesis Award 2019
University of Alberta

M.Sc. Early Achievement Award 2012
Universidade Federal de Minas Gerais (UFMG)

B.Sc. First Class Honors 2010
Universidade Federal de Minas Gerais (UFMG)

Scholarships

Provincial Alberta Innovates Technology Futures Scholarship 2013 – 2018
126,000 CAD over four years in Ph.D..

Brazilian Research Scholarship (CNPq) 2007 – 2008
5,400 BRL over eighteen months in B.Sc..

Research Grants

Canada CIFAR AI Recruitment Chair 2021 – 2025
$300,000 over five years. Sole PI: "Discovering Temporal and Spatial Abstractions in Reinforcement Learning".

Supervision

Current Students (University of Alberta)

Ruo Yu (David) Tao (w/ Adam White), M.Sc. 2021 – present

Erfan Miahi (w/ Martha White), M.Sc. 2021 – present

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)

Alex Lewandowski, University of Alberta 2021 – Present
(w/ D. Schuurmans, supervisor, and A. White)

M.Sc. Thesis Examining Committee

Archit Sakhadeo, University of Alberta 2021
No More Pesky Hyperparameters: Offline Hyperparameter Tuning For 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

CMPUT 366: Intelligent Systems 2016

CMPUT 403: Practical Algorithmics 2016

DCC 865: Design and Analysis of Algorithms 2012

Software Engineering Experience

Avenue Code 2013

Synergia: Engenharia de Software e Sistemas 2011 – 2013

Ilusis Interactive Graphics 2010 – 2011

Selected Talks

Temporal Abstraction in Reinf. Learning with the Successor Representation

Stanford University – Stanford, USA Feb. 2020
Microsoft Workshop on Reinforcement Learning, Forwards and Backwards: Insights from Neuroscience Oct. 2021

Autonomous Navigation of Stratospheric Balloons using Reinforcement Learning

University of Alberta – Edmonton, Canada Jan. 2021

An Operator View of Policy Gradient Methods

University of Alberta – Edmonton, Canada Nov. 2020
DeepMind – London, UK Oct. 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

RLDM – Montréal, Canada Jul. 2019
ICML WS on Exploration in RL, Best paper – 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

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

Area-Chair/Meta-Reviewer

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

Program Committee

Neural Information Processing Systems (NeurIPS)	2018 – 2021
International Conference on Machine Learning (ICML)	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 RELEASED

Arcade Learning Environment versions 0.5 – 0.6 2015 – 2017
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, $\text{DQN}_e^{\text{MMC}} + \text{SR}$, and a gridworld library.

Last update: October 21, 2021.