Marlos Cholodovskis Machado

Department of Computing Science
University of Alberta
3-05 RLAI - Comp. Science Centre
Edmonton, AB, Canada

☎ +1 (587) 598-2265
✉ machado@ualberta.ca
www.cs.ualberta.ca/~machado
mcmachado

RESEARCH INTERESTS

My research interests lie broadly in *Artificial Intelligence* and particularly focus on **Reinforcement Learning (RL)**. Recently, I have been working on problems such as exploration, option discovery, representation learning, transfer learning in RL. I am concerned with general competency in complex domains such as the Arcade Learning Environment.

EDUCATION

Ph.D. in Computing Science 2019
University of Alberta, Canada
Supervisors: Michael Bowling and Marc G. Bellemare

M.Sc. in Computer Science 2013
Universidade Federal de Minas Gerais, Brazil
Supervisors: Luiz Chaimowicz and Gisele L. Pappa

B.Sc. in Computer Science with First Class Honors 2010
Universidade Federal de Minas Gerais, Brazil

PUBLICATIONS

Preprint

[1] **M. C. Machado**, M. G. Bellemare, and M. Bowling: "Count-Based Exploration with the Successor Representation", in *CoRR abs/1807.11622 (arXiv)*.

Journal Articles

[2] **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.

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

[4] 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

[5] **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%].

[6] 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%].

[7] **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%].

[8] 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: 25%].

[9] 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: 39%].

[10] **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: 51%].

[11] **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: 54%].

[12] **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: 75%].

[13] **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: 49%].

[14] **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: 49%].

Magazine Articles

[15] 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

[16] **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*.

[17] **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)*.

[18] M. Liu, **M. C. Machado**, G. Tesauro, and M. Campbell (2017): "The Eigenoption-Critic Framework", in *NIPS Workshop on Hierarchical Reinforcement Learning*.

[19] **M. C. Machado**, and M. Bowling (2016): "Learning Purposeful Behaviour in the Absence of Rewards", in *ICML Workshop on Abstraction in Reinforcement Learning*.

[20] **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*.

[21] **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

[22] **M. C. Machado** (2013): "A Methodology for Player Modeling based on Machine Learning", *M.Sc. thesis*.

[23] **M. C. Machado** (2010): "Classification of Documents Modeled as Graphs" (*in Portuguese*), *B.Sc. thesis*.

AWARDS

- Best Paper Award at the ICML Workshop on Exploration in Reinforcement Learning 2018
  3 out of 51 submitted papers (30 accepted) at the ICML Workshop on Exploration in RL, with a shorter version of the paper "Count-Based Exploration with the Successor Representation".
- ICML Top 10 Reviewer 2018
  Each area chair was asked to identify the most informative review(s) for each paper they've handled. The ten reviewers who received the largest number of nominations were given this award.
- AAMAS Best Paper Nominee 2016
  4 out of 137 full papers at the International Conference on Autonomous Agents and Multi-Agent Systems, with the paper "State of the Art Control of Atari Games Using Shallow Reinforcement Learning".
- IJCAI Outstanding PC Member 2016
  134 out of 1357 PC (Main Track) members were recognized as such.
- Provincial Alberta Innovates Technology Futures Graduate Student Scholarship 2013
  Value: C$ 126,000 over four years.
- M.Sc. Early Achievement Award 2012
  Value: R$3,000
- IEEE CIS Outstanding Student Paper Travel Grant 2012
  Value: $800
- Gold Medal Honour 2010
  Awarded for finishing the undergraduate course with the highest grades
- Brazilian National Council for Scientific and Technological Development (CNPq) Undergraduate Student Research Scholarship 2007
  Value: R$5,400 over eighteen months

RESEARCH INDUSTRY EXPERIENCE

DeepMind – London 2018

I am currently interning at DeepMind (summer and fall of 2018). I am working in the Deep Learning group with Vlad Mnih on a deep reinforcement learning project. More details will be made available in the future.

IBM Research – T.J. Watson Research Center 2017

I was an intern at IBM Research during the summer of 2017. I worked in the AI Foundations group with Gerald Tesauro and Murray Campbell. I developed a new algorithm that discovers options while also learning a non-linear representation of states from raw pixels. The algorithm is based on the concept of the successor representation and on recent successes in the deep reinforcement learning literature. The work developed during the 12-week long internship has lead to two papers [5, 18].

Microsoft Research – New York Lab 2016

I was an intern at Microsoft Research during the summer of 2016. I worked in the Machine Learning group with Alekh Agarwal, Fernando Diaz, Miro Dudik, and Robert Schapire. I focused on discovering options for exploration in reinforcement learning. The developed algorithm draws ideas from contextual bandits and it was evaluated in Minecraft (Malmö project). I also designed practical linear representations for Minecraft and a reusable platform for future research in the Malmö project.

Vetta Labs LTDA 2009 - 2010

As a research intern, I was the main responsible for the project IDATA, which aimed at increasing newspapers' revenue by reducing the number of daily stranded papers. Using supervised learning techniques we were able to increase newspapers' profit in 6% by correctly predicting the number of newspapers to be delivered in each selling point, as described in this report (*in Portuguese*).

SOFTWARE ENGINEERING EXPERIENCE

Avenue Code 2013
Synergia: Engenharia de Software e Sistemas 2011-2013
Ilusis Interactive Graphics 2010-2011

TEACHING EXPERIENCE

CMPUT 366: Intelligent Systems 2016
Teaching assistant in a computer science course focused on reinf. learning. I was responsible for grading and proposing assignments, as well as for holding office hours.

CMPUT 403: Practical Algorithmics 2016
Teaching assistant in a computer science course focused on programming competitions. I was responsible for grading and I substituted the instructor when necessary.

DCC 865: Design and Analysis of Algorithms 2012
Teaching assistant in a graduate-level computer science course. I was responsible for grading and proposing assignments, as well as for holding office hours.

SUPERVISION EXPERIENCE

B.Sc. Students

· Jesse Farebrother (University of Alberta) 2018
Research assistant. We are working on a project where we investigate transfer learning and regularization in deep reinforcement learning using the Arcade Learning Environment as testbed.

· Nicolas Carion (École Normale Supérieure de Lyon) 2015
Summer internship. We worked on a project where we proposed a new transfer learning approach to reinforcement learning based on linear representations and incremental singular value decomposition. We used the Arcade Learning Environment as testbed.

SELECTED TALKS

Count-Based Exploration with the Successor Representation

· 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 Intelligenc (Journal track) – Stockholm, Sweden Jul. 2017
· IJCAI Workshop on Computer Games Workshop (Invited speaker) – Stockholm, Sweden Jul. 2017
· University of Alberta AI Seminar – Edmonton, Canada Oct. 2017

Option Discovery for Continual Learning in Reinforcement Learning

· IBM Research Thomas J. Watson Research Center – Yorktown Heights, NY, USA Aug. 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, MI, 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

Option Discovery for Global Exploration
· Microsoft Research – New York City, NY, USA Aug. 2016

The Arcade Learning Environment: What comes next?
· IJCAI Workshop on General Intelligence and Game-Playing Agents – New York, NY, USA Jul. 2016

Domain-Independent Optimistic Initialization for Reinforcement Learning
· AAAI Workshop on Learn. for General Competency in Video Games – Austin, TX, USA Jan. 2015

A Binary Classification Approach for Automatic Preference Modeling of Virtual Agents in Civilization IV
· IEEE Conference on Computational Intelligence and Games (CIG) – Granada, Spain Sep. 2012

Player Modeling: Towards a Common Taxonomy
· International Conference on Computer Games – Louisville, KY, USA Jul. 2011

SERVICE AND OUTREACH

Journal Reviewer
· 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
· AAAI Conference on Artificial Intelligence (AAAI) 2018-2019
· Neural Information Processing Systems (NIPS) 2018
· International Conference on Machine Learning (ICML) 2018
· International Joint Conference on Artificial Intelligence (IJCAI) 2016-2017

Conference Reviewer
· AAAI Conference on Artificial Intelligence (AAAI) 2015-2016
· IEEE Conference on Computational Intelligence and Games (CIG) 2012
· Brazilian Symposium on Computer Games and Digital Entertainment (SBGames) 2011-2012

Workshops Organized
· AAAI Workshop on Learning for General Competency in Video Games 2015

Participation in Panels
· 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
· ICML Workshop on Lifelong Learning: A Reinforcement Learning Approach 2018
· 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

Last update: August 22, 2018.