

André Barreto

Curriculum Vitae

Google DeepMind
6 Pancras Square
London N1C 4AG UK

Phone: +44 7756 356402
E-mail: andrebarreto@google.com

Professional experience

- 2016– Staff Research Scientist at Google DeepMind, London, UK
- 2014–2016 Adjunct professor in the School of Computer Science at McGill University, QC, Canada
- 2013–2016 Assistant researcher in the Department of Applied and Computational Mathematics at the National Laboratory for Scientific Computing, RJ, Brazil
- 2011–2013 Postdoctoral fellow in the Reasoning and Learning Laboratory, School of Computer Science, McGill University, QC, Canada
- 2008–2011 Postdoctoral fellow in the National Laboratory for Scientific Computing, RJ, Brazil

Education

- 2003–2008 D.Sc. degree in Civil Engineering (Major: Computational Systems)
Federal University of Rio de Janeiro (UFRJ), RJ, Brazil
Partial time in the Colorado State University, CO, USA
Title: *Approximate Solutions for Sequential Decision Problems under Uncertainty*
- 2001–2003 M.Sc. degree in Civil Engineering (Major: Computational Systems)
Federal University of Rio de Janeiro (UFRJ), RJ, Brazil
Title: *Genetic Orthogonal Least Squares Algorithm for Training Radial Basis Function Networks*
- 1996–2000 B.S. degree in Computer Science
Federal University of Juiz de Fora (UFJF), MG, Brazil
Title: *Graph Layout Using a Genetic Algorithm*

Publications

Journal Articles

Matheus R. F. Mendonça, Artur Ziviani, and André Barreto. Graph-based skill acquisition for reinforcement learning. *ACM Computing Surveys*, 52(1):6:1–6:26, February 2019.

André Barreto, Doina Precup, and Joelle Pineau. Practical Kernel-Based Reinforcement Learning. *Journal of Machine Learning Research*, 17(67):1–70, 2016.

Amir-massoud Farahmand, Doina Precup, André Barreto, and Mohammad Ghavamzadeh. Classification-based Approximate Policy Iteration. *IEEE Transactions on Automatic Control*, 60(11):2989–2993, 2015.

André Barreto, Joelle Pineau, and Doina Precup. Policy Iteration Based on Stochastic Factorization. *Journal of Artificial Intelligence Research*, 50:763–803, 2014.

Susan Higashi, [André Barreto](#), Maurício E. Cantão, and Ana Tereza R. Vasconcelos. Analysis of composition-based metagenomics classification. *BMC Genomics*, 13:1–11, 2012.

[André Barreto](#) and Marcelo D. Fragoso. Computing the Stationary Distribution of a Finite Markov Chain Through Stochastic Factorization. *SIAM Journal on Matrix Analysis and Applications*, 32:1513–1523, 2011.

[André Barreto](#) and Charles Anderson. Restricted Gradient-Descent Algorithm for Value-Function Approximation in Reinforcement Learning. *Artificial Intelligence*, 172(4-5):454–482, 2008.

Artem Sokolov, Darrell Whitley, and [André Barreto](#). A Note on the Variance of Rank-Based Selection Strategies for Genetic Algorithms and Genetic Programming. *Genetic Programming and Evolvable Machines*, 8(3):221–237, 2007.

[André Barreto](#), Helio J. C. Barbosa, and Nelson F. F. Ebecken. GOLS—Genetic Orthogonal Least Squares Algorithm for Training RBF Networks. *Neurocomputing*, 69(16-18):2041–2064, 2006.

Referred Conference Papers

[André Barreto](#), Diana Borsa, Shaobo Hou, Gheorghe Comanici, Eser Aygün, Philippe Hamel, Daniel Toyama, Jonathan Hunt, Shibl Mourad, David Silver, and Doina Precup. The option keyboard: Combining skills in reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. 2019. (Acceptance rate: **21.2%**).

Diana Borsa, [André Barreto](#), John Quan, Daniel Mankowitz, Hado van Hasselt, Rémi Munos, David Silver, and Tom Schaul. Universal successor features approximators. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

Jonathan Hunt, [André Barreto](#), Timothy Lillicrap, and Nicolas Heess. Composing entropic policies using divergence correction. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 97, pages 2911–2920, Long Beach, California, USA, 2019. PMLR.

Carlos Riquelme, Hugo Penedones, Damien Vincent, Hartmut Maennel, Sylvain Gelly, Timothy A. Mann, [André Barreto](#), and Gergely Neu. Adaptive temporal-difference learning for policy evaluation with per-state uncertainty estimates. In *Advances in Neural Information Processing Systems (NeurIPS)*. 2019. (Acceptance rate: **21.2%**).

[André Barreto](#), Diana Borsa, John Quan, Tom Schaul, David Silver, Matteo Hessel, Daniel Mankowitz, Augustin Židek, and Rémi Munos. Transfer in Deep Reinforcement Learning Using Successor Features and Generalised Policy Improvement. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 80, pages 501–510, 2018. (Acceptance rate: **25.1%**).

Steven Hansen, Alexander Pritzel, Pablo Sprechmann, [André Barreto](#), and Charles Blundell. Fast deep reinforcement learning using online adjustments from the past. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 10590–10600, 2018. (Acceptance rate: **20.8%**).

Rafael Beirigo, Marco Todorov, and [André Barreto](#). Online TD(λ) for discrete-time Markov jump linear systems. In *IEEE Conference on Decision and Control (CDC)*, pages 2229–2234, 2018.

Matheus R. F. de Mendonça, Artut Ziviani, and [André Barreto](#). Abstract state transition graphs for model-based reinforcement learning. In *Brazilian Conference on Intelligent Systems (BRACIS)*, pages 115–120, 2018.

André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor Features for Transfer in Reinforcement Learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4055–4065. 2017. **Among 3.5% of submissions selected as a spotlight** (Acceptance rate: 20.9%) .

David Silver, Hado van Hasselt, Matteo Hessel, Tom Schaul, Arthur Guez, Tim Harley, Gabriel Dulac-Arnold, David P. Reichert, Neil C. Rabinowitz, André Barreto, and Thomas Degris. The predictron: End-to-end learning and planning. In *Proceedings of the International Conference on Machine Learning, (ICML)*, pages 3191–3199, 2017. (Acceptance rate: 25.9%).

Zhongwen Xu, Joseph Modayil, Hado P van Hasselt, André Barreto, David Silver, and Tom Schaul. Natural value approximators: Learning when to trust past estimates. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2120–2128. 2017. **Among the 3.5% of submissions selected as a spotlight** (Acceptance rate: 20.9%) .

Amir-Massoud Farahmand, André Barreto, and Daniel Nikovski. Value-Aware Loss Function for Model-based Reinforcement Learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 1486–1494, 2017.

Rafael Beirigo, Marco Todorov, and André Barreto. Count-based quadratic control of Markov jump linear systems with unknown transition probabilities. In *IEEE Conference on Decision and Control (CDC)*, pages 4315–4320, 2017.

Rafael Beirigo, Marco Todorov, and André Barreto. Transfer on count-based quadratic control of Markov jump linear systems with unknown transition probabilities. In *Proceedings of Conferência Brasileira de Dinâmica, Controle e Aplicações (DINCON)*, 2017.

André Barreto, Rafael L. Beirigo, Joelle Pineau, and Doina Precup. Incremental Stochastic Factorization for Online Reinforcement Learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 1468–1475, 2016. (Acceptance rate: 26%).

André Barreto, Rafael L. Beirigo, Joelle Pineau, and Doina Precup. An Expectation-Maximization Algorithm to Compute a Stochastic Factorization From Data. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3329–3336, Buenos Aires, Argentina, 2015.

André Barreto. Tree-Based On-line Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2417–2423, Quebec City, QC, Canada, 2014. (Acceptance rate: 28%).

André Barreto, Doina Precup, and Joelle Pineau. On-line reinforcement learning using incremental kernel-based stochastic factorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1493–1501, Lake Tahoe, USA, 2012. (Acceptance rate: 24%).

Helio J. C. Barbosa, Heder S. Bernardino, and André Barreto. Exploring performance profiles for analyzing benchmark experiments. In *Proceedings of the Metaheuristics International Conference (MIC)*, Udine, Italy, 2011.

André Barreto, Doina Precup, and Joelle Pineau. Reinforcement learning using kernel-based stochastic factorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 720–728, Granada, Spain, 2011. (Acceptance rate: 24%).

André Barreto and Marcelo D. Fragoso. Lumping the states of a finite Markov chain through stochastic factorization. In *Proceedings of the 18th World Congress of the International Federation of Automatic Control (IFAC)*, pages 4206–4211, Milano, Italy, 2011.

Douglas A. Augusto, Helio J. C. Barbosa, André Barreto, and Heder S. Bernardino. Evolving numerical constants in grammatical evolution with the ephemeral constant method. In *Progress in Artificial Intelligence (Proceedings of the Portuguese Conference on Artificial Intelligence)*, volume 7026 of *Lecture Notes in Computer Science*, pages 110–124, Lisbon, Portugal, 2011. Springer Berlin Heidelberg.

André Barreto, Heder S. Bernardino, and Helio J. C. Barbosa. Probabilistic performance profiles for the experimental evaluation of stochastic algorithms. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, pages 751–758, New York, NY, USA, 2010. ACM Press.

Helio J. C. Barbosa, Heder S. Bernardino, and André M. S. Barreto. Using performance profiles to analyze the results of the 2006 CEC constrained optimization competition. In *Proceedings of the IEEE World Congress on Computational Intelligence (WCCI)*, pages 4586–4593, USA, 2010. IEEE Computer Society.

André Barreto, Douglas A. Augusto, and Helio J. C. Barbosa. On the characteristics of sequential decision problems and their impact on evolutionary computation and reinforcement learning. In *Artificial Evolution (Proceedings of the International Conference on Artificial Evolution)*, volume 5975 of *Lecture Notes in Computer Science*, pages 194–205, Strasbourg, France, 2010. Springer Berlin Heidelberg.

Darrel Whitley, Marc Richards, Ross Beveridge, and André Barreto. Alternative evolutionary algorithms for evolving programs: evolution strategies and steady-state GP. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, volume 1, pages 919–926, Seattle, Washington, USA, 2006. ACM Press. **Winner best Genetic Programming paper.**

André Barreto, Helio J. C. Barbosa, and Nelson F. F. Ebecken. Growing compact RBF networks using a genetic algorithm. In *Proceedings of the VII Brazilian Symposium on Neural Networks (SBRN)*, pages 61–66, Porto de Galinhas, PE, Brazil, 2002. IEEE Computer Society.

Helio J.C. Barbosa and André Barreto. An interactive genetic algorithm with coevolution of weights for multiobjective problems. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation (GECCO)*, pages 203–210, San Francisco, CA, USA, 2001. Morgan Kaufmann.

André Barreto and Helio J.C. Barbosa. Graph layout using a genetic algorithm. In *Proceedings of the VI Brazilian Symposium on Neural Networks (SBRN)*, pages 179–184, Rio de Janeiro, RJ, Brazil, 2000. IEEE Computer Society.

Technical Reports, Workshop Papers and Extended Abstracts

Steven Hansen, Will Dabney, André Barreto, Tom Van de Wiele, David Warde-Farley, and Volodymyr Mnih. Fast task inference with variational intrinsic successor features, 2019.

Hado van Hasselt, John Quan, Matteo Hessel, Zhongwen Xu, Diana Borsa, and André Barreto. General non-linear Bellman equations. In *Proceedings of the Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

Hugo Penedones, Damien Vincent, Hartmut Maennel, Sylvain Gelly, Timothy Arthur Mann, and André Barreto. Temporal difference learning with neural networks - study of the leakage propagation problem. *CoRR*, abs/1807.03064, 2018.

Daniel J. Mankowitz, Augustin Zidek, [André Barreto](#), Dan Horgan, Matteo Hessel, John Quan, Junhyuk Oh, Hado van Hasselt, David Silver, and Tom Schaul. Unicorn: Continual learning with a universal, off-policy agent. In *Proceedings of the Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

A. M. Farahmand, [André Barreto](#), and Daniel Nikovski. Value-aware loss function for model learning in reinforcement learning. In *The European Workshop on Reinforcement Learning (EWRL)*, 2016.

[André Barreto](#), Borja Pigem, Joelle Pineau, and Doina Precup. Starting to uncover the relationship between stochastic factorization and hidden markov models. In *Advances in Neural Information Processing Systems (NIPS): Workshop on Novel Trends and Applications in Reinforcement Learning*, 2014.

Amir-massoud Farahmand, Doina Precup, [André Barreto](#), and Mohammad Ghavamzadeh. CAPI: Generalized Classification-Based Approximate Policy Iteration. In *Proceedings of the Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2013.

Anwarissa Asbah, [André Barreto](#), Clement Gehring, Joelle Pineau, and Doina Precup. Reinforcement Learning Competition 2013: Controllability and Kernel-Based Stochastic Factorization. In *Proceedings of the ICML Workshop on the Reinforcement Learning Competition*, 2013.

Douglas A. Augusto, Helio J. C. Barbosa, [André Barreto](#), and Heder S. Bernardino. A new approach for generating numerical constants in grammatical evolution. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, volume 1, pages 193–194, Dublin, Ireland, 2011. ACM Press.

[André Barreto](#) and Doina Precup. Kernel-based stochastic factorization for batch reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS), Learning and Planning from Batch Time Series Data Workshop*, 2010.

Raquel L. Costa and [André Barreto](#). Naive Bayes classifier applied to DNA barcoding. In *Proceedings of the 6th International Conference of the Brazilian Association for Bioinformatics and Computational Biology (X-Meeting)*, 2010.

Susan Higashi, [André Barreto](#), Maurício E. Cantão, and Ana Tereza R. Vasconcelos. Analysis of composition-based metagenomics classification. In *Proceedings of the 6th International Conference of the Brazilian Association for Bioinformatics and Computational Biology (X-Meeting)*, 2010.

[André Barreto](#), Douglas A. Augusto, and Helio J. C. Barbosa. On the characteristics of sequential decision problems and their impact on evolutionary computation. In *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, volume 1, pages 1767–1768, Montreal, Québec, Canada, 2009. ACM Press.

Book Chapters

Helio J. C. Barbosa, Heder S. Bernardino, and [André Barreto](#). Using performance profiles for the analysis and design of benchmark experiments. In *Advances in Metaheuristics*, pages 21–36. Springer, 1 edition, 2013.

Selected invited talks

- August, 2019 *Efficient Reinforcement Learning with Generalized Policy Updates*, MILA — Québec Artificial Intelligence Institute, Montreal, Québec, Canada.
- August, 2019 *Hierarchical Reinforcement Learning*, Deep Learning and Reinforcement Learning Summer School, Edmonton, Alberta, Canada.
- February, 2019 *The Option Keyboard: Combining Skills in Reinforcement Learning*, The 12th Barbados Workshop on Reinforcement Learning, Bridgetown, Barbados.
- May, 2018 *Recent Advances in Transfer for Reinforcement Learning*, Reasoning and Learning Laboratory, Montreal, Québec, Canada.
- April, 2018 *Recent Advances in Transfer for Reinforcement Learning*, Google Zürich, Zürich, Switzerland.
- August, 2014 *Tree-Based On-line Reinforcement Learning*, Reasoning and Learning Laboratory, School of Computer Science, McGill University
- April, 2012 *Incremental Kernel-Based Stochastic Factorization*, The 7th Barbados Workshop on Reinforcement Learning
- December, 2009 *Introduction to Reinforcement Learning*, Seminar on Data Mining, Department of Computer Science, Federal University of Juiz de Fora (UFJF)
- July, 2009 *Reducing the Dimension of Markov Models Through Stochastic Factorization*, Reasoning and Learning Laboratory, School of Computer Science, McGill University
- November, 2008 *Approximate Solutions for Sequential Decision Problems under Uncertainty*, Laboratory of Logic, Artificial Intelligence and Formal Methods, Institute of Mathematics and Statistics, University of São Paulo (USP)

Highlights and Selected Awards

- 2018 Participated in the project *Training Action Selection Neural Networks Using a Differentiable Credit Function*, awarded a patent under number GP-202165-00-PCT.
- 2014 Author of the project “*Solutions of Sequential Decision-Making Problems Involving Big Data*” awarded a grant of BRL 30,000 on a nation-wide competition refereed by *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq/MCTI).
- 2014 Program co-chair of the *Sequential Decision Making with Big Data Workshop* at the AAAI Conference on Artificial Intelligence, Quebec City, Canada.
- 1996 First position in the Brazilian national admission exam to enter the School of Computer Science at *Universidade Federal de Juiz de Fora*.

Student Supervisions

- 2016– Matheus Mendonça. *Approximating Network Measures Using Machine Learning*, PhD candidate at the Graduate Program of the National Laboratory for Scientific Computing (co-supervision with Prof. Artur Ziviani, secondary supervisor)

- 2012–2017 Isabella Alvim Guedes. *Design of an Empirical Function to Estimate the Protein-Ligand Affinity Constant*, PhD candidate at the Graduate Program of the National Laboratory for Scientific Computing (co-supervision with Prof. Laurent Dardenne, secondary supervisor)
- 2011–2014 Igor Magalhães Ribeiro. *Kernel Methods for Dimensionality Reduction in Metagenomic Applications*, M.Sc at the Graduate Program of the National Laboratory for Scientific Computing (co-supervision with Prof. Ana Tereza Vasconcelos, main supervisor)
- 2012–2012 Alicia Bendz. *Kernel-Based Stochastic Factorization with Multiple Agents*, undergraduate internship in the School of Computer Science at McGill University (co-supervision with Prof. Joelle Pineau, secondary supervisor)
- 2012–2012 Ryan Primeau. *Application of Kernel-Based Stochastic Factorization to RC Helicopter Control*, undergraduate internship in the School of Computer Science at McGill University (co-supervision with Prof. Doina Precup, secondary supervisor)
- 2009–2011 Susan Higashi. *Analysis of Composition-Based Metagenomic Classification*, M.Sc at the Graduate Program of the National Laboratory for Scientific Computing (co-supervision with Prof. Ana Tereza Vasconcelos, secondary supervisor)

External thesis evaluator / Committee member

- August, 2016 Maurício Archanjo N. Coelho. *An Approach for Structured Data Prediction Using the Perceptron*, PhD at the Federal University of Juiz de Fora (UFJF)
- August, 2016 Tiene André Felisbino. *Statistical Learning Approaches for the Computation of Tensor Components for Face Recognition*, M. Sc. at the Graduate Program of the National Laboratory for Scientific Computing.
- March, 2014 Maria Fernanda Ribeiro Dias. *Empirical Analysis of Machine Learning Techniques for Classifying *Metarhizium anisopliae* Protein Sequences*, M. Sc. at the Graduate Program of the National Laboratory for Scientific Computing
- June, 2010 Maurício Archanjo N. Coelho. *Structured Data Prediction Using Perceptrons, With Applications on Path Planning*, M.Sc at the Federal University of Juiz de Fora (UFJF)

Teaching

- April–July, 2014 *Data-Driven Modeling*, Graduate Program of the National Laboratory for Scientific Computing
- August–October, 2010 *Machine Learning*, Graduate Program of the National Laboratory for Scientific Computing

Reviewing

- Journal of Machine Learning Research (JMLR)*
- Journal of Artificial Intelligence Research (JAIR)*
- Neural Information Processing Systems Conference (NIPS)*
- International Conference on Machine Learning (ICML)*

Personal information

Birthday June 6, 1977

Citizenship Brazilian

Languages Portuguese (mother tongue) and English

November 23, 2019