

# Aldo Pacchiano

## Curriculum vitae – September 2022

### Personal Information

Citizenship Mexico  
Email [apacchiano@microsoft.com](mailto:apacchiano@microsoft.com)  
www [www.aldopacchiano.ai](http://www.aldopacchiano.ai)

### Research Interests

Online Learning, Reinforcement Learning, Active Learning and Algorithmic Fairness.

### Current Position

7/2021 – present Postdoctoral Researcher  
Microsoft Research, NYC

### Education

- 1/2016 – 6/2021 Ph.D. Student in Computer Science  
Advised by Peter Bartlett and Michael Jordan  
University of California, Berkeley
- 9/2015-1/2016 Ph.D. Student in Operations Research and Financial Engineering  
Princeton University
- 8/2013-6/2014 Masters of Engineering in Electrical Engineering and Computer Science  
Advised by Constantinos Daskalakis  
Massachusetts Institute of Technology
- 10/2012-7/2013 Masters of Advanced Study in Theoretical Mathematics  
University of Cambridge
- 8/2008-6/2012 Bachelors of Science Computer Science and Mathematics  
Massachusetts Institute of Technology

Ph.D. Thesis Model Selection for Contextual Bandits and Reinforcement Learning

### Publications

1. T. Lin\*, A. Pacchiano\*, Y. Yu\*, M. Jordan (2022). Online Nonsubmodular Minimization with Delayed Costs: From Full Information to Bandit Feedback. In: International Conference on Machine Learning (ICML) (Acceptance rate: 21.9%.)
2. R Müller, A Pacchiano (2022). Meta Learning MDPs with linear transition models. In: International Conference on Artificial Intelligence and Statistics (AISTATS)

3. T. Moskovitz, M. Arbel, J. Parker-Holder, A. Pacchiano (2022). Towards an Understanding of Default Policies in Multitask Policy Optimization. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
4. A. Pacchiano\*, S. Singh\*, E. Chou, A. Berg, J. Foerster (2021). Neural Pseudo-Label Optimism for the Bank Loan Problem. In: Advances in Neural Information Processing Systems (NeurIPS) (Acceptance rate: 26%.)
5. M. Papini, A. Trinzioni, A. Pacchiano, M. Restelli, A. Lazaric, M. Pirotta (2021). Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection. In: Advances in Neural Information Processing Systems (NeurIPS) (Acceptance rate: 26%.)
6. N. Chatterji\*, A. Pacchiano\*, P. Bartlett, M. Jordan (2021). On the Theory of Reinforcement Learning with Once-per-episode Feedback. In: Advances in Neural Information Processing Systems (NeurIPS) (Acceptance rate: 26%.)
7. A. Pacchiano\*, J. Lee, P. Bartlett, O. Nachum (2021). Near Optimal Policy Optimization via REPS. In: Advances in Neural Information Processing Systems (NeurIPS) (Acceptance rate: 26%.)
8. T. Moskovitz, J. Parker-Holder, A. Pacchiano, M. Arbel, M. Jordan (2021). Tactical Optimism and Pessimism for Deep Reinforcement Learning. In: Advances in Neural Information Processing Systems (NeurIPS) (Acceptance rate: 26%.)
9. A. Pacchiano, P. Ball, J. Parker-Holder, K. Choromanski, S. Roberts. Towards Tractable Optimism in Model-Based Reinforcement Learning. In: Uncertainty in Artificial Intelligence (UAI).
10. A. Cutkosky\*, C. Dann\*, A. Das\*, C. Gentile\*, A. Pacchiano\*, M. Purohit\* (2021). Dynamic Balancing for Model Selection in Bandits and RL. In: International Conference on Machine Learning (ICML) (Acceptance rate: 21.5%.)
11. D. Malik, A. Pacchiano, V. Srinivasan, Y. Li (2021). Sample Efficient Reinforcement Learning in Continuous State Spaces: A Perspective Beyond Linearity. In: International Conference on Machine Learning (ICML) (Acceptance rate: 21.5%.)
12. A. Pacchiano, H. Jiang, M. Jordan (2021). Robustness Guarantees for Mode Estimation with an Application to Bandits. In: Conference on Artificial Intelligence (AAAI).
13. J. Lee, A. Pacchiano, V. Muthukumar, W. Kong, E. Brunskill (2021). In: International Conference on Artificial Intelligence and Statistics (AISTATS)
14. A. Pacchiano, M. Ghavamzadeh, P. Bartlett, H. Jiang (2021). Stochastic Bandits with Linear Constraints. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
15. H. Jiang\*, Q. Jiang\*, A. Pacchiano\* (2021). Learning the Truth from Only One Side of the Story. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
16. J. Parker-Holder, L. Metz, C. Resnick, H. Hu, A. Lerer, A. Letcher, A. Peysakhovich, A. Pacchiano, J. Foerster (2020). Ridge Rider: Finding Diverse Solutions by Following Eigenvalues of the Hessian. In: Advances in Neural Information Processing Systems (NeurIPS.)
17. A. Pacchiano, M. Phan, Y. Abassi-Yadkori, A. Rao, J. Zimmert, T. Lattimore, C. Szepesvari (2020). Model Selection in Contextual Stochastic Bandit Problems. In: Advances in Neural Information Processing Systems (NeurIPS.)
18. A. Pacchiano, J. Parker-Holder, K. Choromanski, S. Roberts (2020). Effective Diversity in Population Based Reinforcement Learning. In: Advances in Neural Information Processing Systems (NeurIPS.)
19. X. Song, W. Gao, Y. Yang, K. Choromanski, A. Pacchiano, Y. Tang (2020). ES-MAML: Simple Hessian-free Meta Learning. In: International Conference on Learning Representations (ICLR.)
20. K. Choromanski\*, A. Pacchiano\*, J. Parker-Holder\*, Y. Tang, D. Jain, Y. Yang, A. Iscen, J. Hsu, V. Sindwani (2020). Probably Robust Blackbox Optimization for Reinforcement Learning. In: Conference on Robot Learning (CoRL).
21. E. Mazumdar\*, A. Pacchiano\*, Y. Ma, M. Jordan, P. Bartlett (2020). On Approximate Thompson Sampling with Langevin Algorithms. In: International Conference on Machine Learning (ICML)
22. J. Lee, A. Pacchiano, P. Bartlett, M. Jordan (2020). Accelerated Message Passing for Entropy-regularized MAP inference. In: International Conference on Machine Learning (ICML)
23. A. Pacchiano\*, J. Parker-Holder\*, Y. Tang\*, K. Choromanski, A. Choromanska, M. Jordan (2020). Learning to Score Behaviors for Guided Policy Optimization. In: International Conference on Machine Learning (ICML)
24. K. Choromanski, D. Cheikh, J. Davis, V. Likhoshesterov, A. Nazaret, A. Bahamou, X. Song, M. Akarte, J. Parker-Holder, J. Bergquist, Y. Gao, A. Pacchiano, T. Sarlos, A. Weller, V. Sindhwani (2020). Stochastic Flows

- and Geometric Optimization on the Orthogonal Group. In: International Conference on Machine Learning (ICML)
25. P. Ball, J. Parker-Holder, A. Pacchiano, K. Choromanski, S. Roberts (2020). Ready Policy One: World Building Through Active Learning. In: International Conference on Machine Learning (ICML)
  26. J. Lee\*, A. Pacchiano\*, M. Jordan (2020). Convergence Rates of Smooth Message Passing with Rounding in Entropy Regularized MAP Inference. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
  27. K. Choromanski\*, A. Pacchiano\*, J. Parker-Holder\*, Y. Tang\* (2020). Practical Nonisotropic Monte Carlo Sampling in High Dimensions via Determinantal Point Processes. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
  28. S. Chiappa, R. Jiang, T. Stepleton, A. Pacchiano, H. Jiang, J. Aslanides (2020). A General Approach to Fairness with Optimal Transport. In: Conference on Artificial Intelligence (AAAI).
  29. K. Choromanski\*, A. Pacchiano\*, J. Parker-Holder\*, Y. Tang\*, V. Sindhwani (2019). From Complexity to Simplicity: Adaptive ES-Active Subspaces for Blackbox Optimization. In: Advances in Neural Information Processing Systems (NeurIPS.)
  30. R. Jiang\*, A. Pacchiano\*, T. Stepleton, H. Jiang, S. Chiappa (2019). Wasserstein Fair Classification. In: Uncertainty in Artificial Intelligence (UAI).
  31. N. Chatterji\*, A. Pacchiano\*, P. Bartlett (2019). Online Learning with Kernel Losses. In: International Conference on Machine Learning (ICML.)
  32. A. Pacchiano. Y. Bachrach (2019). Computing Stable Solutions in Threshold Network Flow Games with Bounded Treewidth. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS.)
  33. K. Choromanski, A. Pacchiano, J. Pennington, Y. Tang (2019). KAMA-NNs: Low-dimensional rotation based neural networks. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
  34. M. Rowland, K. Choromanski, F. Chalus, A. Pacchiano. T. Sarlos, R. Turner, A. Weller (2018). Geometrically Coupled Monte Carlo Sampling. In: Advances in Neural Information Processing Systems (NeurIPS.)
  35. K. Bhatia\*, A. Pacchiano\*, N. Flammarion, P. Bartlett, M. Jordan (2018). Gen-Oja: A Two-time-scale Approach for Streaming CCA. In: Advances in Neural Information Processing Systems (NeurIPS.)
  36. M. Rowland, A. Pacchiano, A. Weller (2017). Conditions Beyond Treewidth for Tightness of Higher-Order LP Relaxations. In: International Conference on Artificial Intelligence and Statistics (AISTATS)
  37. A. Pacchiano. O Williams (2015). Real Time Clustering of Time Series Using Triangular Potentials. International Conference on Artificial Intelligence and Applications (AIFU).
  38. P. Etingof, S. Gong, A. Pacchiano, Q. Ren, T. Schedler (2012). Computational Approaches to Poisson Traces Associated to Finite Subgroups of  $Sp_{2n}(\mathbb{C})$ . Experimental Mathematics.

## Revisions

1. S. Chiappa\*, A. Pacchiano\* (2020). Fairness with Continuous Optimal Transport. In: Major Revision by Machine Learning.

## Refereed Workshop Papers

1. A. Pacchiano, D. Wulsin, R. Barton, L. Voloch (2022). Neural Design for Genetic Perturbation Experiments. In: ReALML Workshop ICML 2022.
2. J. Lorraine, J. Parker-Holder, P. Vicol, A. Pacchiano, L. Metz, T. Kachman, J. Foerster (2021). Using Bifurcations for Diversity in Differentiable Games. In: ICML 2021 Beyond First Order Methods Workshop.
3. RL theory workshop
4. D. Malik, A. Pacchiano, V. Srinivasan, Y. Li (2021). Sample Efficient Reinforcement Learning in Continuous State Spaces: A Perspective Beyond Linearity. In: ICML 2021 Workshop on Reinforcement Learning Theory.

5. M. Papini, A. Trinzioni, A. Pacchiano, M. Restelli, A. Lazaric, M. Pirotta (2021). Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection. In: ICML 2021 Workshop on Reinforcement Learning Theory.
6. N. Chatterji\*, A. Pacchiano\*, P. Bartlett, M. Jordan (2021). On the Theory of Reinforcement Learning with Once-per-episode Feedback. In: ICML 2021 Workshop on Reinforcement Learning Theory.
7. R Müller, A Pacchiano (2022). Meta Learning MDPs with linear transition models. In: ICML 2021 Workshop on Reinforcement Learning Theory.
8. J. Lee, W. Kong, A. Pacchiano, V. Muthukumar, E. Brunskill (2021). Estimating Optimal Policy Value in Linear Contextual Bandits beyond Gaussianity. In: ICML 2021 Workshop on Reinforcement Learning Theory.
9. J. Parker-Holder, L. Metz, C. Resnick, H. Hu, A. Lerer, A. Letcher, A. Peysakhovich, A. Pacchiano, J. Foerster (2020). Ridge Rider: Finding Diverse Solutions by Following Eigenvalues of the Hessian. In: First Order Methods in ML Systems Workshop ICML 2020.
10. R. Muller, J. Parker-Holder, A. Pacchiano (2020). Taming the Herd: Multi-Modal Meta-Learning with a Population of Agents. In: ICML 2020 Workshop LifelongML.
11. A. Pacchiano, J. Parker-Holder, K. Choromanski, S. Roberts (2020). Effective Diversity in Population Based Reinforcement Learning. In: Fourth Lifelong Machine Learning Workshop ICML 2020.
12. X. Song, K. Choromanski, J. Parker-Holder, Y. Tang, W. Gao, A. Pacchiano, T. Sarlos, D. Jain, Y. Yang (2020). Reinforcement Learning with Chromatic Networks for Compact Architecture Search. ICLR Workshop on Neural Architecture Search 2020.
13. M. Abdullah\*, A. Pacchiano\*, M. Draief (2018). European Workshop on Reinforcement Learning.

## Preprints

1. A. Pacchiano, C. Dann, C. Gentile (2022). Best of Both Worlds Model Selection. arXiv preprint.
2. A. Pacchiano, O. Nachum, N. Tripuraneni, P. Bartlett (2022). Joint Representation Training in Sequential Tasks with Shared Structure. arXiv preprint.
3. A. Pacchiano, A. Saha, J. Lee (2021). Dueling RL: Reinforcement Learning with Trajectory Preferences.
4. A. Pacchiano, P. Bartlett, M. Jordan (2021). An Instance-Dependent Analysis for the Cooperative Multi-Player Multi-Armed Bandit.
5. S. Chiappa\*, A. Pacchiano\* (2020). Fairness with Continuous Optimal Transport. arXiv preprint.
6. J. Chan, A. Pacchiano, N. Tripuraneni, Y. Song, P. Bartlett, M. Jordan (2021). Parallelizing Contextual Linear Bandits. arXiv preprint.
7. A. Pacchiano, C. Dann, C. Gentile, P. Bartlett (2021). Regret Bound Balancing and Elimination for Model Selection in Bandits and RL. arXiv preprint.

## Talks

1. Invited talk at Jan Peter's lab at TU Darmstadt. Jan 6, 2021. "Towards an Understanding of Transfer Learning in Reinforcement Learning Domains"
2. Talk at Ali Jadbabaie's lab at MIT. February 25, 2021. "An Instance-Dependent Analysis for the Cooperative Multi-Player Multi-Armed Bandit"
3. Invited talk at Microsoft Research NYC. January 21, 2021. "Model Selection in Stochastic Bandit Problems and RL, from CORRAL to Regret Balancing"
4. Talk at Pulkit Agrawal's lab at MIT. December 7, 2020. "Learning to Score Behaviors for Guided Policy Optimization"
5. Talk at Emma Brunskill's lab at Stanford University. November 19, 2020. "Regret Bound Balancing and Elimination for Model Selection"

## Teaching

1. I was a Teaching Assistant (TA) for the course “Introduction to Artificial Intelligence” in the computer science department at UC Berkeley. Summer and Fall 2016
2. I was a Teaching Assistant (TA) for the course “Introduction to Machine Learning” in the computer science department at UC Berkeley. Spring 2016.
3. I was a Teaching Assistant (TA) for the course “Introduction to Inference” in the computer science department at MIT. Spring 2014.
4. I was a Teaching Assistant (TA) for the course “Theory of Computation” in the mathematics department at MIT. Fall 2013.

## Academic Service

1. Conference Reviews: NeurIPS (2019, 2020, 2021, 2022), ICML (2019, 2020, 2021, 2022), L4DC 2020, ICLR (2021), AISTATS (2020, 2021, 2022)
2. Journal Reviews: Transactions on Machine Learning Research (TMLR), IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI).
3. Area Chair AISTATS 2023
4. Co-organizer of the TTIC Summer Workshop 2022 entitled “New Models in Online Decision Making for Real-World Applications”

## Graduate Coursework

EECS: Combinatorial Algorithms and Data Structures (Christos Papadimitriou), Convex Optimization (Laurent El Ghaoui), Randomness and Computation (Alistair Sinclair), Deep Reinforcement Learning (Sergey Levine), Convex Optimization and Applications (Martin Wainwright), Beyond Worst Case Analysis (Luca Trevisan), Information Theory (Thomas Courtrade), Human Intelligence Enterprise (Patrick Winston), The Society of Mind (Marvin Minsky), Theory of Computation (Michael Sipser), Convex Optimization (Dimitri Bertsekas), Linear and Nonlinear Optimization (Mengdi Wang), Quantum Computation (Cambridge) .

Statistics: Advanced Probability (Alan Sola), High Dimensional Statistics (Martin Wainwright), Statistical Learning Theory (Peter Bartlett), Statistical Theory and Methods (Samory Kpotufe), Probability Theory ( Patrick Cheridito)

Mathematics: Algebraic Combinatorics (Richard Stanley), Algebraic Number Theory (MIT), Additive Combinatorics (Ben Green), Representation Theory (Stuart Martin), Extremal Combinatorics (Imre Leader), Combinatorial Theory (MIT).

## Other Research Work Experience

10/2019 – 5/2020 and 12/2020 – 6/2021

Visiting Researcher  
Facebook AI Research  
Mentors: Jakob Foerster and Mohammad Ghavamzadeh

5/2020 – 8/2020

Research Intern  
Microsoft Research, NYC

Mentors: Miroslav Dudik and Robert Schapire

5/2019 – 8/2019

Research Intern  
Google Research  
Mentor: Ariel Kleiner

9/2018 – 1/2019

Research Intern  
DeepMind  
Mentor: William Dabney

5/2018 – 8/2018

Research Intern  
Google Brain  
Mentor: Krzysztof Choromanski

## Competitions

08/2014

Mexico City Youth Prize  
Mexico City Major  
Academic Achievement Recognition 2<sup>nd</sup> place

10/2012

National Youth Prize (Mexico)  
National Youth Institute Academic Achievement Recognition

07/2008, 07/2007, 07/2006

International Mathematics Olympiad  
Mexico Team  
Silver, Bronze and Honourable Mention