

Aldo Pacchiano

Microsoft Research
300 Lafayette Street
New York, NY, 10012
✉ apacchiano@microsoft.com
🌐 www.aldopacchiano.ai

Research Interests

Online learning, reinforcement learning and algorithmic fairness.

Education

- 2016-2021 **University of California Berkeley**, *Ph.D in Computer Science*.
○ Thesis: [Model Selection for Contextual Bandits and Reinforcement Learning](#).
○ Advisors: Peter Bartlett and Michael Jordan.
- 2015 **Princeton University**, *Ph.D Student in Operations Research and Financial Engineering*.
Transferred to UC Berkeley.
- 2013-2014 **Massachusetts Institute of Technology**, *MEng in Electrical Engineering and Computer Science*.
○ Thesis: [Trace Reconstruction Problem](#).
○ Advisor: Constantinos Daskalakis.
- 2012-2013 **University of Cambridge**, *MASt in Mathematics*.
○ Essay: Mixing Time of Random Transpositions.
○ Examiner: Nathanael Berestycki.
- 2008-2012 **Massachusetts Institute of Technology**, *B.S. in Mathematics and in Computer Science and Engineering*.

Research Work Experience

- 2021- **Postdoctoral Researcher**, *Microsoft Research, New York, NY, USA*.
- 2019-2021 **Visiting Researcher**, *Facebook AI Research, Menlo Park, CA, USA*.
○ Mentors: Jakob Foerster and Mohammad Ghavamzadeh.
- 05-08/2020 **Research Intern**, *Microsoft Research, New York, NY, USA*.
○ Mentors: Miroslav Dudik and Robert Schapire.
- 05-08/2019 **Research Intern**, *Google Research, New York, NY, USA*.
○ Mentor: Ariel Kleiner.
- 2018-2019 **Research Intern**, *Deepmind, London, UK*.
○ Mentor: William Dabney.
- 05-08/2019 **Research Intern**, *Google Brain, New York, NY, USA*.
○ Mentor: Krzysztof Choromanski.

Fellowships and Awards

- 2021 **Funded BAIR Commons Partnership**, Facebook AI Research and Berkeley Artificial Intelligence Research Lab (BAIR).
○ Project: “*Mitigating Emergent Biases in Online Learning*”.
- 2020 **Funded BAIR Commons Partnership**, Google and Berkeley Artificial Intelligence Research Lab (BAIR).
○ Project: “*LP-based Algorithms for Reinforcement Learning*”.
- 2019-2020 **Funded BAIR Commons Partnership**, Facebook AI Research and Berkeley Artificial Intelligence Research Lab (BAIR).
○ Project: “*Stochastic Bandits with Linear Constraints*”.
- 2016-2019 **BAIR Huawei Fellowship**, Berkeley Artificial Intelligence Research Lab.

Press

- 2020 **Venture Beat**, “*Researchers develop technique to increase sample efficiency in reinforcement learning*”.

Submissions and Preprints

* *Indicates co-first authorship.*

2. A. Gupta, D. Alvarez-Melis, **A. Pacchiano**, T. Moskovitz. (2022). Task-agnostic Imitation Learning with Trajectory Based Optimal Transport. *Working paper*.
3. J. Lee, W. Kong, **A. Pacchiano**, V. Muthukumar, E. Brunskill. (2022). Estimating Optimal Policy Value in General Linear Contextual Bandits. *Working paper*.
1. A. Wagenmaker, **A. Pacchiano** (2022). Leveraging Offline Data in Online Reinforcement Learning. *arXiv preprint*.
4. **A. Pacchiano**, O. Nachum, N. Tripuraneni, P. Bartlett (2022). Joint Representation Training in Sequential Tasks with Shared Structure. *arXiv preprint*. Under review at *Transactions of Machine Learning Research (TMLR)*.
5. **A. Pacchiano**, A. Saha, J. Lee (2022). Dueling RL: Reinforcement Learning with Trajectory Preferences. *arXiv preprint*. Under review at *International Conference on Artificial Intelligence and Statistics (AISTATS) 2023*.
6. **A. Pacchiano**, P. Bartlett, M. Jordan (2022). An Instance-Dependent Analysis for the Cooperative Multi-Player Multi-Armed Bandit. *arXiv preprint*. Under review at *Algorithmic Learning Theory (ALT) 2023*.
7. **A. Pacchiano**, D. Wulsin, R. Barton, L. Voloch (2022). Neural Design for Genetic Perturbation Experiments. *arXiv preprint*. Under review at *International Conference on Learning Representations (ICLR) 2023*.

8. E. Gal, S. Singh, **A. Pacchiano**, B. Walker, T. Lyons, J. Foerster (2022). Unbiased Decisions Reduce Regret: Adversarial Optimism for the Bank Loan Problem. Under review at *International Conference on Learning Representations (ICLR) 2023*.
9. J. Chan*, **A. Pacchiano***, N. Tripuraneni*, Y. Song, P. Bartlett, M. Jordan (2021). Parallelizing Contextual Linear Bandits. *arXiv preprint*.
10. **A. Pacchiano**, C. Dann, C. Gentile, P. Bartlett (2021). Regret Bound Balancing and Elimination for Model Selection in Bandits and RL. *arXiv preprint*.
11. S. Chiappa*, **A. Pacchiano*** (2020). Fairness with Continuous Optimal Transport. *arXiv preprint*. Journal submission under review at *Machine Learning*.

Publications

* Indicates co-first authorship. The fields I publish are conference-driven.

1. **A. Pacchiano**, C. Dann, C. Gentile (2022). Best of Both Worlds Model Selection. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (Acceptance rate: 25.6%.)
2. J. Parker-Holder*, Y. Xu*, P. Ball*, **A. Pacchiano***, O. Rybkin, S. Roberts, T. Rocktäschel, E. Grefenstette (2022). Learning World Models in a Handful of Reward Free Deployments. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (Acceptance rate: 25.6%.)
3. A. Gupta*, **A. Pacchiano***, Y. Zhai, S. Kakade, S. Levine (2022). Unpacking Reward Shaping: Understanding the Benefits of Reward Engineering on Sample Complexity. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (Acceptance rate: 25.6%.)
4. 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%. *Spotlight paper, acceptance rate 2%*.)
5. R Müller, **A Pacchiano** (2022). Meta Learning MDPs with linear transition models. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. (Acceptance rate: 29%.)
6. 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)*. (Acceptance rate: 29%. *Oral presentation. One of 4 papers out of 1685 submissions nominated for best paper award.*)

7. **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%.)
8. M. Papini, A. Trinzioni, **A. Pacchiano**, M. Restelli, A. Lazaric, M. Pirootta (2021). Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (Acceptance rate: 26%.)
9. 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%.)
10. **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%.)
11. 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%.)
12. **A. Pacchiano**, P. Ball, J. Parker-Holder, K. Choromanski, S. Roberts (2021). Towards Tractable Optimism in Model-Based Reinforcement Learning. In: *Uncertainty in Artificial Intelligence (UAI)*. (Acceptance rate: 26.3%.)
13. A. Cutkosky*, C. Dann*, A. Das*, C. Gentile*, **A. Pacchiano***, M. Purohit* (2021). Dyanamic Balancing for Model Selection in Bandits and RL. In: *International Conference on Machine Learning (ICML)*. (Acceptance rate: 21.5%. *Spotlight paper, acceptance rate 3%.*)
14. 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%. *Spotlight paper, acceptance rate 3%.*)
15. **A. Pacchiano**, H. Jiang, M. Jordan (2021). Robustness Guarantees for Mode Estimation with an Application to Bandits. In: *Conference on Artificial Intelligence (AAAI)*. (Acceptance rate: 21%.)
16. J. Lee, **A. Pacchiano**, V. Muthukumar, W. Kong, E. Brunskill (2021). Online Model Selection for Reinforcement Learning with Function Approximation. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. (Acceptance rate: 29.8%.)
17. **A. Pacchiano**, M. Ghavamzadeh, P. Bartlett, H. Jiang (2021). Stochastic Bandits with Linear Constraints. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. (Acceptance rate 29.8%.)

18. 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)*. (Acceptance rate 29.8%.)
19. 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)*. (Acceptance rate 20%.)
20. **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)*. (Acceptance rate 20%.)
21. **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)*. (Acceptance rate 20%. *Spotlight presentation, acceptance rate 2.9%.*)
22. 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)*. (Acceptance rate 26.5%.)
23. 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)*. (Acceptance rate 34.7%.)
24. 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)*. (Acceptance rate 21.8%.)
25. J. Lee, **A. Pacchiano**, P. Bartlett, M. Jordan (2020). Accelerated Message Passing for Entropy-regularized MAP inference. In: *International Conference on Machine Learning (ICML)*. (Acceptance rate 21.8%.)
26. **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)*. (Acceptance rate 21.8%.)
27. K. Choromanski, D. Cheikhi, 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)*. (Acceptance rate 21.8%.)

28. 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)*. (Acceptance rate 21.8%.)
29. 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)*. (Acceptance rate 30%.)
30. 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)*. (Acceptance rate 30%.)
31. 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)*. (Acceptance rate 20.6%.)
32. 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)*. (Acceptance rate 21.1%.)
33. R. Jiang*, **A. Pacchiano***, T. Stepleton, H. Jiang, S. Chiappa (2019). Wasserstein Fair Classification. In: *Uncertainty in Artificial Intelligence (UAI)*. (Acceptance rate 26%. *Oral presentation, acceptance rate 7.7%*.)
34. N. Chatterji*, **A. Pacchiano***, P. Bartlett (2019). Online Learning with Kernel Losses. In: *International Conference on Machine Learning (ICML)*. (Acceptance rate 22.6%. *Oral long talk.*.)
35. **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)*. (Acceptance rate 24%.)
36. 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)*. (Acceptance rate 32.4%.)
37. 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)*. (Acceptance rate 20.8%. *Spotlight presentation, acceptance rate 3.5%*.)
38. 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)*. (Acceptance rate 20.8%.)

39. 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)*. (Acceptance rate 31.7%.)
40. **A. Pacchiano**, O Williams (2015). Real Time Clustering of Time Series Using Triangular Potentials. International Conference on Artificial Intelligence and Applications (AIFU).
41. 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.

Refereed Workshops without Published Proceedings

* *Indicates co-first authorship.*

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: *Beyond First Order Methods Workshop ICML 2021*.
3. D. Malik, **A. Pacchiano**, V. Srinivasan, Y. Li (2021). Sample Efficient Reinforcement Learning in Continuous State Spaces: A Perspective Beyond Linearity. In: *Workshop on Reinforcement Learning Theory ICML 2021*.
4. M. Papini, A. Trinzioni, **A. Pacchiano**, M. Restelli, A. Lazaric, M. Pirotta (2021). Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection. In: *Workshop on Reinforcement Learning Theory ICML 2021*.
5. N. Chatterji*, **A. Pacchiano***, P. Bartlett, M. Jordan (2021). On the Theory of Reinforcement Learning with Once-per-episode Feedback. In: *Workshop on Reinforcement Learning Theory ICML 2021*. *Oral presentation*.
6. R Müller, **A Pacchiano** (2022). Meta Learning MDPs with linear transition models. In: *Workshop on Reinforcement Learning Theory ICML 2021*.
7. J. Lee, W. Kong, **A. Pacchiano**, V. Muthukumar, E. Brunskill (2021). Estimating Optimal Policy Value in Linear Contextual Bandits beyond Gaussianity. In: *Workshop on Reinforcement Learning Theory ICML 2021*.
8. 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*. *Spotlight presentation*.

9. R. Muller, J. Parker-Holder, **A. Pacchiano** (2020). Taming the Herd: Multi-Modal Meta-Learning with a Population of Agents. In: *Workshop LifelongML ICML 2020*.
10. **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*.
11. 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. *Workshop on Neural Architecture Search ICLR 2020*.
12. M. Abdullah*, **A. Pacchiano***, M. Draief (2018). Reinforcement Learning with Wasserstein Distance Regularisation, with Applications to Multipolicy Learning. In: *European Workshop on Reinforcement Learning*.

Public Repositories

PLOT - Neural bandits library with access to public datasets.

Patents

1. K. Choromanski, A. Pacchiano, V. Sindhvani. "Robust And Data-Efficient Blackbox Optimization". U.S. Patent 62/793,248. 2019.

Talks

On the Statistical Complexity of Batch Learning: Theory and Algorithms.

University of Oxford, AIMS Seminar, 2022.

Best of Both Worlds Model Selection.

Carnegie Mellon University, AI Seminar, 2022.

Online Model Selection: the principle of regret balancing.

Boston University, AIR Seminar, 2022.

University of Washington, 2022.

Unpacking Reward Shaping: Understanding the Benefits of Reward Engineering on Sample Complexity.

Princeton University, RL Seminar, 2022.

University College London, Gatsby Unit ML Seminar, 2022.

An Instance-Dependent Analysis for the Cooperative Multi-Player Multi-Armed Bandit.

Massachusetts Institute of Technology, 2022.

Model Selection in Stochastic Bandit Problems and RL, from CORRAL to Regret Balancing.

Microsoft Research, MSR NYC Seminar - Job Talk, 2021.

Regret Bound Balancing and Elimination for Model Selection.

Stanford University, 2020.

Learning to Score Behaviors for Guided Policy Optimization.

Massachusetts Institute of Technology, 2020 ([link](#)).

Teaching Assistantships

2016 **Introduction to Artificial Intelligence**, *University of California Berkeley.*

2016 **Introduction to Machine Learning**, *University of California Berkeley.*

2014 **Introduction to Inference**, *Massachusetts Institute of Technology.*

2013 **Theory of Computation**, *Massachusetts Institute of Technology.*

Professional Activities and Service

Reviewer NeurIPS (2019 - 2021), NeurIPS 2022 (*top reviewer*), ICML (2019-2022), L4DC 2020, ICLR (2021), AISTATS (2020-2022), TMLR, TPAMI.

Area Chair AISTATS 2023.

Organizer **TTIC Summer Workshop 2022 “New Models in Online Decision Making for Real-World Applications”.**

≈ 100 attendees from US and international institutions.

Organizer RIIAA 2018 (student-run AI conference in Mexico City).

Other Publications and Texts

2016 **La Unión de Morelos**, *Triángulos y voces, reflexiones sobre inteligencia Artificial, Newspaper Article.*

Newspaper article discussing the future of AI research in Spanish. Published in Mexican newspaper ‘La Unión de Morelos’.

2014 **Los Hijos de la Malinche**, *El Vandam, Nosotros y Ellos, La Otra Costa, Short Stories.*

Short stories published in Mexican literary magazine ‘Los Hijos de la Malinche’.

- **Copyrighted Literary Texts**, *Al otro lado del río, Cuestión de Ser, La cúpula de hielo, Novels*

- **Literary Website**, *Short stories and other texts.*

Personal Information

Citizenship Mexico.

Languages Spanish, Italian, English.