

Aldo Pacchiano

Boston University
665 Commonwealth Avenue
Boston, MA, 02215
✉ pacchian@bu.edu
🌐 www.aldopacchiano.ai

Research Interests

Reinforcement learning, online learning, language models and bandits.

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*.
- 2012-2013 **University of Cambridge**, *MASt in Mathematics*.
- 2008-2012 **Massachusetts Institute of Technology**, *B.S. in Mathematics and in Computer Science and Engineering*.

Research Work Experience

- 2024-Present **Assistant Professor**, *Boston University, Boston, MA, USA*.
- 2024-Present **Visiting Scientist**, *Broad Institute of MIT and Harvard, Cambridge, MA, USA*.
- 2023-2024 **Eric and Wendy Schmidt Postdoctoral Fellow**, *Broad Institute of MIT and Harvard, Cambridge, MA, USA*.
- 2021-2023 **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.

Selected Submissions and Preprints

* *Indicates co-first authorship.*

1. A. Russo, R. Welch, **A. Pacchiano** (2025). Learning to Explore: An In-Context Learning Approach for Pure Exploration. *arXiv preprint*.

2. I. Shenfeld, F. Faltings, P. Agrawal, **A. Pacchiano** (2025). Language Model Personalization via Reward Factorization. *arXiv preprint*.
3. D. Misra*, **A. Pacchiano***, Ta-Chung Chi, Ge Gao (2025). Principled Fine-tuning of LLMs from User-Edits: A Medley of Preference, Supervision and Reward.
4. Y. Chen, **A. Pacchiano***, I. Paschalidis (2025). Self-Reinforce or Imitate: Hybrid Learning for Reasoning Distillation.
5. **A. Pacchiano**, C. Dann, C. Gentile, P. Bartlett (2021). Regret Bound Balancing and Elimination for Model Selection in Bandits and RL. *arXiv preprint*.

Selected Publications

* Indicates co-first authorship.

1. N. Brukhim*, **A. Pacchiano***, M. Dudik, R. Schapire. On the Hardness of Bandit Learning. In: *Conference on Learning Theory (COLT) 2025*.
2. Y. Chen, **A. Pacchiano**, I. Paschalidis. Multiple-policy Evaluation via Density Estimation. In: *International Conference on Machine Learning (ICML) 2025*.
3. A. Russo, **A. Pacchiano**. Adaptive Exploration for Multi-Reward Multi Policy Evaluation. In: *International Conference on Machine Learning (ICML) 2025*.
4. **A. Pacchiano**. Second Order Bounds for Contextual Bandits with Function Approximation. In: *International Conference on Learning Representations (ICLR) 2025*.
5. C. Kausik, M. Mutti, **A. Pacchiano**, A. Tewari. A Framework for Partially Observed Reward-States in RLHF. In: *International Conference on Learning Representations (ICLR) 2025*.
6. C. Zhang, Z. Hong, **A. Pacchiano**, P. Agrawal. ORSO: Accelerating Reward Design via Online Reward Selection and Policy Optimization. In: *International Conference on Learning Representations (ICLR) 2025*.
7. N. Das, S. Chakraborty, **A. Pacchiano**, S. Ray Chowdhury. Provably Sample Efficient RLHF via Active Preference Optimization. In: *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) 2025*.
8. D. Misra*, **A. Pacchiano***, R. Schapire* (2024). Provable Interactive Learning with Hindsight Instruction Feedback. In: *International Conference on Machine Learning (ICML) 2024*.
9. **A. Pacchiano**, C. Dann, C. Gentile. Data-Driven Regret Balancing for Online Model Selection in Bandits. In: *International Conference on Artificial Intelligence and Statistics (AISTATS) 2024*.
10. J. Lee*, A. Xie*, **A. Pacchiano**, Y. Chandak, C. Finn, O. Nachum, E. Brunskill. Supervised Pretraining Can Learn In-Context Reinforcement Learning. In: *Advances in Neural Information Processing Systems, (NeurIPS) 2023*. (*Spotlight paper.*)
11. A. Wagenmaker, **A. Pacchiano**. Leveraging Offline Data in Online Reinforcement Learning. In: *International Conference on Machine Learning (ICML) 2023*.

12. **A. Pacchiano***, A. Saha*, J. Lee. Dueling RL: Reinforcement Learning with Trajectory Preferences. In: *International Conference on Artificial Intelligence and Statistics (AISTATS) 2023*.
13. **A. Pacchiano**, D. Wulsin, R. Barton, L. Voloch. Neural Design for Genetic Perturbation Experiments. In: *International Conference on Learning Representations (ICLR) 2023*. (*Spotlight paper*.)
14. A. Gupta*, **A. Pacchiano***, Y. Zhai, S. Kakade, S. Levine. Unpacking Reward Shaping: Understanding the Benefits of Reward Engineering on Sample Complexity. In: *Advances in Neural Information Processing Systems (NeurIPS) 2022*.
15. N. Chatterji*, **A. Pacchiano***, P. Bartlett, M. Jordan. On the Theory of Reinforcement Learning with Once-per-episode Feedback. In: *Advances in Neural Information Processing Systems (NeurIPS) 2021*.
16. T. Moskovitz, J. Parker-Holder, **A. Pacchiano**, M. Arbel, M. Jordan. Tactical Optimism and Pessimism for Deep Reinforcement Learning. In: *Advances in Neural Information Processing Systems (NeurIPS) 2021*.
17. **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) 2021*.
18. **A. Pacchiano**, M. Ghavamzadeh, P. Bartlett, H. Jiang. Stochastic Bandits with Linear Constraints. In: *International Conference on Artificial Intelligence and Statistics (AISTATS) 2021*.
19. **A. Pacchiano***, M. Phan*, Y. Abassi-Yadkori, A. Rao, J. Zimmert, T. Lattimore, C. Szepesvari. Model Selection in Contextual Stochastic Bandit Problems. In: *Advances in Neural Information Processing Systems (NeurIPS) 2020*.
20. X. Song, W. Gao, Y. Yang, K. Choromanski, **A. Pacchiano**, Y. Tang. ES-MAML: Simple Hessian-free Meta Learning. In: *International Conference on Learning Representations (ICLR) 2020*.
21. E. Mazumdar*, **A. Pacchiano***, Y. Ma, M. Jordan, P. Bartlett. On Approximate Thompson Sampling with Langevin Algorithms. In: *International Conference on Machine Learning (ICML) 2020*.
22. **A. Pacchiano***, J. Parker-Holder*, Y. Tang*, K. Choromanski, A. Choromanska, M. Jordan. Learning to Score Behaviors for Guided Policy Optimization. In: *International Conference on Machine Learning (ICML) 2020*.
23. R. Jiang*, **A. Pacchiano***, T. Stepleton, H. Jiang, S. Chiappa. Wasserstein Fair Classification. In: *Uncertainty in Artificial Intelligence (UAI) 2019*. (*Oral presentation*.)
24. P. Etingof, S. Gong, **A. Pacchiano**, Q. Ren, T. Schedler. Computational Approaches to Poisson Traces Associated to Finite Subgroups of $\mathrm{Sp}_{2n}(\mathbb{C})$. In: *Experimental Mathematics 2012*.

Selected Refereed Workshops without Published Proceedings

* Indicates co-first authorship.

1. A. Gupta*, T. Moskovitz, D. Alvarez-Melis, **A. Pacchiano**. (2022). Transfer Reinforcement Learning via the Undo Maps Formalism. In: *New Frontiers in Learning, Control, and Dynamical Systems Workshop, ICML 2023*.

Patents

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

Selected Talks

How to Learn Sequential Decision-Making Algorithms from Data.

King Abdullah University of Science and Technology (KAUST), Rising Stars in AI Workshop, April 2025.

A Tale of Two Algorithmic Principles: Optimism and Posterior Sampling.

Columbia University, IEOR/DRO Seminar, December 2024.

Experiment Planning with Function Approximation.

ICML, LatinX Workshop, July 2024.

Brown University, BigAI Talk, April 2024.

Information Theory and Applications Workshop (ITA), San Diego, February 2024.

The Dissimilarity Dimension: Sharper Bounds for Optimistic Algorithms.

University of California San Diego, TILOS Seminar, 2024.

Professional Activities and Service

- Reviewer NeurIPS (2019 - 2021), NeurIPS 2022 (*top reviewer*), ICML (2019,2022, 2024), L4DC 2020, ICLR (2021), AISTATS (2020-2022), TMLR, TPAMI.
- Area Chair AISTATS 2023, NeurIPS 2023, AISTATS 2024, NeurIPS 2024, AISTATS 2025, ICLR 2025, ICML 2025, NeurIPS 2025.
- Organizer **TTIC Summer Workshop 2024 "Adaptive Learning in Complex Environments"**.
- Organizer **Deployable Reinforcement Learning Workshop @ RLC 2024**.
- Program Co-Chair **International Symposium on Artificial Intelligence and Mathematics. ISAIM 2024**.
- Organizer **TTIC Summer Workshop 2022 "New Models in Online Decision Making for Real-World Applications"** .

Other Relevant Work Experience

2014-2015 **Associate**, *Goldman Sachs, London, UK*.

Personal Information

- Citizenship Mexico, Italy, Green Card Holder.
- Languages Spanish, English.
- Full Name Aldo Pacchiano Camacho (Mexico), Aldo Pacchiano (Italy).