

Paula Nicoleta Gradu

CONTACT	1441 Walnut Street, Berkeley, CA 94709	pgradu@berkeley.edu https://paula-gradu.github.io/
---------	---	--

EDUCATION	PhD in Computer Science	Aug 2021 — Aug 2026
-----------	--------------------------------	---------------------

University of California, Berkeley, Berkeley, CA.

Advisors Benjamin Recht, Michael I. Jordan

GPA 4.0 / 4.0

Bachelor of Arts

Sep 2017 — May 2021

Princeton University, Princeton, NJ.

GPA 3.885 / 4.0

Major Mathematics

Major GPA 3.954 / 4.0

Certificates Statistics and Machine Learning (*equivalent to minor program*)
Applications of Computing (*equivalent to CS minor*)
Applied and Computational Mathematics

Awards The Middleton Miller '29 Prize 2021
(*awarded for the best independent work in mathematics*)

The Andrew H. Brown Prize 2020
(*awarded to the top 2-4 juniors in the Mathematics department*)

The Shapiro Prize for Academic Excellence 2018
(*awarded for outstanding academic achievement in first year of study*)

PUBLICATIONS (α - β denotes alphabetical ordering)

- 1. From Individual Experience to Collective Evidence: An Incident-Based Framework for Identifying Systemic Discrimination.** *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.
(α - β): Jessica Dai, Paula Gradu, Inioluwa Deborah Raji, Benjamin Recht.
- 2. Online Control For Adaptive Tapering Of Medications.** *Proceedings of the 62nd IEEE Conference on Decision and Control (CDC)*, 2023. *Nominated for the Robert Tempo Best CDC Paper Award.*
Paula Gradu, Benjamin Recht.
- 3. Valid Inference After Causal Discovery.** *Journal of the American Statistical Association (JASA)*, 2024.
Paula Gradu*, Tijana Zrnic*, Yixin Wang, Michael I. Jordan.
- 4. CLIP-OGD: An Experimental Design for Adaptive Neyman Allocation in Sequential Experiments.** *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, 2023. *Spotlight Presentation (awarded to 10% of accepted papers).*
(α - β): Jessica Dai, Paula Gradu, Chris Harshaw.

5. **Adaptive Regret for Control of Time-Varying Dynamics.** *Proceedings of the 5th Conference on Learning for Dynamics and Control (L4DC)*, 2023.
(α - β): Paula Gradu, Elad Hazan, Edgar Minasyan.
6. **Projection-free Adaptive Regret with Membership Oracles.** *Proceedings of the 34th Conference on Algorithmic Learning Theory (ALT)*, 2023.
Zhou Lu, Nataly Brukhim, Paula Gradu, Elad Hazan.
7. **Lyapunov Density Models: Constraining Distribution Shift in Learning-Based Control.** *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.
Katie Kang, Paula Gradu, Jason Choi, Michael Janner, Claire Tomlin, Sergey Levine.
8. **Online Control of Unknown Time-Varying Dynamical Systems.** *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
Edgar Minasyan, Paula Gradu, Max Simchowitz, Elad Hazan.
9. **Non-Stochastic Control with Bandit Feedback.** *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
(α - β): Paula Gradu, John Hallman, Elad Hazan.
10. **SyntaxNet for Romanian: Results and Potential.** *Proceedings of the 12th International Conference “Linguistic Resources and Tools for Processing the Romanian Language” (ConsILR)*, 2016.
Paula Gradu, Radu Ion.

PREPRINTS

1. **Large language model-based extraction of cycle-level chemotherapy-induced peripheral neuropathy phenotypes from clinical notes.** In submission to *NPJ Precision Oncology*.
Jenna Fields, Paula Gradu, Ahmed Alaa, Travis Zack.
2. **A Practical Control-Theoretic Framework for Taxonomizing Equipose in Chemotherapy Dosing.** In preparation for submission to *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*.
Jenna Fields, Paula Gradu, Benjamin Recht, Ahmed Alaa, Travis Zack.

WORKSHOP PAPERS

1. **Machine Learning for Mechanical Ventilation Control.** *Poster at the Machine Learning for Health (ML4H) Symposium*, 2021.
Daniel Suo, Naman Agarwal, Wenhan Xia, Xinyi Chen, Udaya Ghai, Alexander Yu, Paula Gradu, Karan Singh, Cyril Zhang, Edgar Minasyan, Julianne LaChance, Tom Zajdel, Manuel Schottdorf, Daniel Cohen, Elad Hazan.
 2. **DELUCA — A Differentiable Control Library: Environments, Methods, and Benchmarking.** *Poster at the Differentiable Computer Vision, Graphics, and Physics in Machine Learning Workshop, Neural Information Processing Systems (NeurIPS)*, 2020.
Paula Gradu*, John Hallman*, Daniel Suo*, Alex Yu, Naman Agarwal, Udaya Ghai, Karan Singh, Cyril Zhang, Anirudha Majumdar, Elad Hazan.
-

EMPLOYMENT	Student Researcher, Google AI Princeton	July 2019 - June 2021
	Research Assistant, Princeton University	
	<i>Geophysical Fluid Dynamics Laboratory</i>	Feb 2019 - July 2019
	<i>Socio-Cognitive Processes Lab</i>	Oct 2017 - July 2018
	Research Intern, Romanian Academy	Feb 2016 - Mar 2017
	<i>Research Institute for Artificial Intelligence “Mihai Drăgănescu”</i>	
SERVICE	Women in Computer Science and Electrical Engineering (WiCSE)	
	<i>Co-President</i>	2023-2024
	<i>Social Chair</i>	2022-2023
	Women in Machine Learning (WiML)	
	<i>General Chair, The 3rd WiML Un-Workshop</i>	2022
PEER REVIEW	International Conference on Neural Information Processing Systems (NeurIPS)	2022 (top 8%), 2023, 2024, 2025
	International Conference on Machine Learning (ICML)	2022, 2023, 2024, 2025
	International Conference on Artificial Intelligence and Statistics (AISTATS)	2025
	Learning for Dynamics and Control Conference (L4DC)	2023, 2024
	Journal of Machine Learning Research (JMLR)	2022
TEACHING	Graduate Student Instructor, Statistical Learning Theory. University of California, Berkeley, Fall 2023.	
	Teaching Assistant, Computational Control Theory. Princeton University, Fall 2020.	
GRADUATE COURSEWORK	University of California, Berkeley	
	Machine Learning Systems (LLM Edition)	
	Machine Learning and Econ	
	Causal Inference	
	Linear Systems Theory	
	A Guided Tour of Key Statistical Methods and Analysis Tools	
	The Philosophy and History of Automated Decision Making	
	Princeton University	
	Theoretical Machine Learning	
	Optimization for Machine Learning	
	Probability Theory	
	Information Theory and Applications	
	Deep Learning for Natural Language Processing	
	Optimization on Smooth Manifolds	
	Theory of Deep Weakly Supervised Learning	
	Foundations of Reinforcement Learning	
	Mathematical Understanding of Deep Learning	