Philip Amortila

philipa4@illinois.edu
www.philipamortila.com
Updated: Jan 2024

Education

2019–2024 **PhD Computer Science**, University of Illinois at Urbana-Champaign.
- Advisor: Nan Jiang [web page]
- Thesis Proposal: *Power and Limitations of Function Approximation for Efficient Reinforcement Learning* [pdf]
- Committee: Nan Jiang, Csaba Szepesvári, Maxim Raginsky, Arindam Banerjee.

2017–2019 **MSc Computer Science**, McGill University.
- Advisors: Prakash Panangaden, Marc G. Bellemare
- Thesis: *Couplings in Reinforcement Learning – Applications to State Abstraction and Algorithm Analysis* [pdf]
- GPA: 4.00/4.00

2013–2017 **BSc Honours Maths & Physics**, McGill University.
- Minor in Computer Science
- Distinctions: First Class Honours, Principal's Student-Athlete Honour Roll
- GPA: 3.75/4.00

Publications

Preprints

[1] **Scalable Online Exploration via Coverability**
Philip Amortila, Dylan J. Foster, Akshay Krishnamurthy
Preprint [pdf]

[2] **Mitigating Covariate Shift in Misspecified Regression with Applications to Reinforcement Learning**
Philip Amortila, Tongyi Cao, Akshay Krishnamurthy
Preprint [arXiv]

Conference Papers

[3] **Harnessing Density Ratios for Online Reinforcement Learning**
Philip Amortila, Dylan J. Foster, Nan Jiang, Ayush Sekhari, Tengyang Xie
ICLR 2024 (Spotlight) [arXiv]

[4] **The Optimal Approximation Ratios in Misspecified Off-Policy Value Function Estimation**
Philip Amortila, Nan Jiang, Csaba Szepesvári
ICML 2023 [arXiv]

[5] **A Few Expert Queries Suffices for Sample-Efficient RL with Resets and Linear Value Approximation**
Philip Amortila, Nan Jiang, Dhruv Madeka, Dean P. Foster
NeurIPS 2022 [arXiv, talk]

[6] **On Query-efficient Planning in MDPs under Linear Realizability of the Optimal State-value Function**
Gellert Weisz, Philip Amortila, Barnabás Janzer, Yasin Abbasi-Yadkori, Nan Jiang, Csaba Szepesvári
COLT 2021 [arXiv, talk]

[7] **Exponential Lower Bounds for Planning in MDPs With Linearly-Realizable Optimal Action-Value Functions**
Gellert Weisz, Philip Amortila, Csaba Szepesvári
ALT 2021 (Best Student Paper Award) [arXiv, talk]

[8] **Solving Constrained Markov Decision Processes via Backward Value Functions**
Harsh Satija, Philip Amortila, Joelle Pineau
ICML 2020 [arXiv, talk]

[9] **A Distributional Analysis of Sampling-Based Reinforcement Learning Algorithms**
Philip Amortila, Doina Precup, Prakash Panangaden, Marc G. Bellemare
AISTATS 2020 [arXiv, talk]
NeurIPS 2019 Optimization in RL Workshop (Spotlight talk)[talk]

[10] **Learning Graph Weighted Models on Pictures**
Philip Amortila and Guillaume Rabusseau
International Conference on Grammatical Inference 2018 [arXiv]

Technical Notes

[11] **A Variant of the Wang-Foster-Kakade Lower Bound for the Discounted Setting**
Philip Amortila, Nan Jiang, Tengyang Xie [arXiv]

Workshop Papers

[12] **Temporally Extended Metrics for Markov Decision Processes**
Philip Amortila, Marc G. Bellemare, Prakash Panangaden, Doina Precup
AAAI 2019 Safety in AI Workshop (Spotlight talk) [pdf]

Work/Research/Visiting Positions

Summer 2023 **Research Intern**, *Microsoft Research New England*, with Dylan Foster and Akshay Krishnamurthy.
Topic: Representation Learning & Modular Approaches to Rich Observation RL

Summer 2022 **Research Intern**, *Amazon NYC*, with Dean P. Foster.
Topic: Coordination & Communication in Partially Observed Cooperative Games

Fall 2021 **Research Intern**, *Amazon NYC*, with Dean P. Foster.
Topic: Optimal Algorithms for Expert-Assisted RL With Linear Features. *(See [5])*

Summer 2021 **Visiting Researcher**, *University of Alberta*, with Csaba Szepesvári.
Topic: Optimal Methods for Off-policy Evaluation With Misspecification. *(See [4])*

Fall 2020 **Visiting Graduate Student**, *Simons Institute @ UC Berkeley*.
Theory of Reinforcement Learning Program. [web page]

Summer 2020 **Visiting Researcher**, *University of Alberta*, with Csaba Szepesvári.
Topic: Limits of Sample-Efficient Learning With Linear Features (*See [6] and [7]*)

2019 - 2024 **Grad Research Assistant**, *UIUC*, with Nan Jiang.
Topic: Theory of Function Approximation for Reinforcement Learning

2017 - 2019 **Grad Research Assistant**, *McGill University*, with Prakash Pananagaden and Marc G. Bellemare.
Topic: Simple/Unified Convergence Results for Popular RL algorithms. (*See [9]*)

Summer 2016 **Undergrad CS Researcher**, *McGill University*, with Prakash Panangaden.
Topic: Minimization and learning of Weighted Automata. (*See [10]*)

Summer 2015 **Undergrad Physics Researcher**, *Simon Fraser University*, with Mike Hayden.
Topic: Algorithms/Hardware for Magnetic Particle Image Generation and Interpretation.

Fellowships & Awards

Academic **Finalist for Google PhD Fellowship 2023**.
1 of 3 selected amongst all applicants at UIUC. Not selected for national competition.

Academic **Finalist for Apple Scholar in AI/ML 2022**.
1 of 3 selected amongst all applicants at UIUC. Not selected for national competition.

Academic **Best Student Paper Award**, ALT 2021.

$63,000 **NSERC Postgraduate Doctoral Fellowship (PGS-D)**, *NSERC*, 2019.

$6,000 **Undergraduate Summer Research Award**, *NSERC*, 2016.

$6,000 **Undergraduate Summer Research Award**, *NSERC*, 2015.

$800 **Tomlinson Engagement Award for Mentoring**, *McGill University*, 2016.

$800 **Tomlinson Engagement Award for Mentoring**, *McGill University*, 2015.

Academic **Principal's Student-Athlete Honour Roll**, *McGill University*, 2017.

Service

Reviewer

Journal of Machine Learning Research (JMLR) 2022, 2023
Transactions of Machine Learning Research (TMLR) 2022, 2023
International Conference on Machine Learning (ICML) 2020, 2021, 2022, 2023, 2024
Neural Information Processing Systems (NeurIPS) 2020, 2021, 2022
International Conference on Learning Representations (ICLR) 2023

Teaching Experience

Teaching Assistant	**CS 542 Statistical Reinforcement Learning**	*UIUC*, Fall 2023
	CS 443 Reinforcement Learning	*UIUC*, Spring 2023
	CS 498 Reinforcement Learning	*UIUC*, Fall 2019
	CS 598 Foundations of Machine Learning	*McGill*, Fall 2018
	CS 551 Applied Machine Learning	*McGill*, Winter 2018
	CS 551 Applied Machine Learning	*McGill*, Fall 2017
	CS 302 Functional Programming (undergrad TA)	*McGill*, Winter 2017
Tutoring	**Mathematics Help Desk Tutor**	*McGill*, 2015–2017

Languages

English, French	*Fluent*
Arabic	*Mother Tongue*
Python, Mathematica	*Occasionally*