

David Abel

Senior Research Scientist, DeepMind

📍 Edinburgh, UK

🏠 david-abel.github.io

✉️ dmabel@google.com

EDUCATION

- 2020 **Ph.D in Computer Science**, *Brown University*, Providence, RI.
Advisor: Michael L. Littman.
Committee: George Konidaris, Stefanie Tellex, Peter Stone, Will Dabney.
Thesis: A Theory of Abstraction in Reinforcement Learning.
- 2019 **M.A. in Philosophy**, *Brown University*, Providence, RI.
Advisor: Joshua Schechter.
Thesis: Concepts in Bounded Rationality: Perspectives from Reinforcement Learning.
- 2015 **M.S. in Computer Science**, *Brown University*, Providence, RI.
Advisor: Stefanie Tellex.
Thesis: Learning to Plan in Complex Stochastic Domains.
- 2013 **B.A. in Computer Science & Philosophy**, *Carleton College*, Northfield, MN.
Advisors: David Liben-Nowell (CS), Anna Moltchanova (Philosophy).
Thesis: Toward the Defense of Hypercomputation.

SELECTED AWARDS

Outstanding Paper Award, "On the Expressivity of Markov Reward", NeurIPS 2021.
Presidential Award for Excellence in Teaching, Brown University.
Runner-up, AAAI/ACM SIGAI Doctoral Dissertation Award, 2020.
7x Top Reviewer, ICML / NeurIPS / AISTATS.

PUBLICATIONS

Journal Papers

People Construct Simplified Mental Representations to Plan.
Mark K. Ho, David Abel, Carlos G. Correa, Michael L. Littman, Jonathan D. Cohen, Thomas L. Griffiths.
Nature 2022.

The Value of Abstraction.
Mark K. Ho, David Abel, Thomas L. Griffiths, Michael L. Littman.
Current Opinions in Behavioral Sciences 2019.

Conference Papers

A Definition of Continual Reinforcement Learning.
David Abel, André Barreto, Benjamin Van Roy, Hado van Hasselt, Doina Precup, Satinder Singh.
NeurIPS 2023.

Settling the Reward Hypothesis.
Michael Bowling, John D. Martin, David Abel, Will Dabney.
ICML 2023.

Meta-Gradients in Non-Stationary Environments.

Jelena Luketina, Sebastian Flennerhag, Yannick Schroecker, David Abel, Tom Zahavy, Satinder Singh.

CoLLAs 2022.

On the Expressivity of Markov Reward

David Abel, Will Dabney, Anna Harutyunyan, Mark K. Ho, Michael L. Littman, Doina Precup, Satinder Singh.

NeurIPS 2021 (**Outstanding Paper Award**).

Revisiting Peng's $Q(\lambda)$ for Modern Reinforcement Learning

Tadashi Kozuno, Yunhao Tang, Mark Rowland, Rémi Munos, Steven Kapturowski, Will Dabney, Michal Valko, David Abel.

ICML 2021.

Lipschitz Lifelong Reinforcement Learning

Erwan Lecarpentier, David Abel, Kavosh Asadi, Yuu Jinnai, Emmanuel Rachelson, Michael L. Littman.

AAAI 2021.

What can I do here? A Theory of Affordances in Reinforcement Learning

Khimya Khetarpal, Zafarali Ahmed, Gheorghe Comanici, David Abel, Doina Precup.

ICML 2020.

Value Preserving State-Action Abstractions.

David Abel, Nathan Umbanhowar, Khimya Khetarpal, Dilip Arumugam, Doina Precup, Michael L. Littman.

AISTATS 2020.

The Efficiency of Human Cognition Reflects Planned Use of Information Processing.

Mark K. Ho, David Abel, Jonathan D. Cohen, Michael L. Littman, Thomas L. Griffiths.

AAAI 2020.

The Expected-Length Model of Options.

David Abel*, John Winder*, Marie DesJardins, Michael L. Littman

IJCAI 2019.

Finding Options that Minimize Planning Time.

Yuu Jinnai, David Abel, D. Ellis Hershkowitz, Michael L. Littman, George Konidaris.

ICML 2019.

Discovering Options for Exploration by Minimizing Cover Time.

Yuu Jinnai, Jee Won Park, David Abel, George Konidaris.

ICML 2019.

State Abstraction as Compression in Apprenticeship Learning.

David Abel, Dilip Arumugam, Kavosh Asadi, Yuu Jinnai, Michael L. Littman, Lawson L.S. Wong.

AAAI 2019.

State Abstractions for Lifelong Reinforcement Learning.

David Abel, Dilip Arumugam, Lucas Lehnert, Michael L. Littman.
ICML 2018.

Policy and Value Transfer in Lifelong Reinforcement Learning

David Abel*, Yuu Jinnai*, Yue Guo, George Konidaris, Michael L. Littman.
ICML 2018.

Bandit-Based Solar Panel Control.

David Abel, Edward C. Williams, Stephen Brawner, Michael L. Littman, Emily Reif.
IAAI 2018.

Near Optimal Behavior via Approximate State Abstraction.

David Abel*, D. Ellis Hershkowitz*, Michael L. Littman.
ICML 2016.

Goal-Based Action Priors.

David Abel, D. Ellis Hershkowitz, Gabriel Barth-Maron, Stephen Brawner, Kevin O'Farrell, James MacGlashan, Stefanie Tellex.
ICAPS 2015.

TEACHING EXPERIENCE

Primary Instructor A First Byte of Computer Science, *110 students*, Brown University.
Artificial Intelligence and Society, *20 students*, Summer at Brown.

Guest Lecturer Continual Reinforcement Learning, Stanford.
On Research (2022), MIT.
On Research (2021a), Harvard University.
On Research (2021b), Harvard University.
On Research (2020), Harvard University.
Reinforcement Learning, Brown University.
Data Science, Brown University.
A First Byte of Computer Science, Brown University.
Data Science Summer Program, Microsoft Research NYC Summer School.

Teaching Assistant A First Byte of Computer Science, Brown University.
Artificial Intelligence, Brown University.
Data Structures, Carleton College.
Logic, Carleton College.
Intro to Computer Science $\times 3$, Carleton College.

PROFESSIONAL EXPERIENCE

2022–Present **Senior Research Scientist**, *DeepMind London*.
2020–2022 **Research Scientist**, *DeepMind London*.
Summer 2019 **Research Intern**, *DeepMind London*, Hosted by Dr. Will Dabney.

Summer 2016 **Visiting Researcher**, *University of Oxford*, Hosted by Dr. Owain Evans.
Summer 2015 **Research Intern**, *Microsoft Research NYC*, Hosted by Dr. Fernando Diaz.

INVITED TALKS AND PANELS

- 2024 *A Careful Look at Agency*, Bath RL Workshop.
- 2023 *Three Dogmas of Reinforcement Learning*, ICML Workshop on Interactive Learning.
On Reinforcement Learning, AI Podden.
The Reward Hypothesis, Principles of Intelligence Podcast.
On Research, Panel at New in ML Workshop at NeurIPS.
AI and the Reward Hypothesis, Fidelity AI Lab.
- 2022 *On the Expressivity of Markov Reward*, Brown.
On the Expressivity of Markov Reward, Imperial College London.
On the Expressivity of Markov Reward, U. of Witwatersrand.
On the Expressivity of Markov Reward, Australian Reinforcement Learning Group.
On the Expressivity of Markov Reward, UCSB.
Agency as Reward Maximization, Comp. Bio. Conference at ICL.
- 2021 *On the Expressivity of Markov Reward*, Mila.
Abstraction in Reinforcement Learning, Purdue.
- 2020 *Abstraction in Reinforcement Learning*, Michigan.
Abstraction in Reinforcement Learning, MSR Redmond.
Abstraction in Reinforcement Learning, NYU.
Abstraction in Reinforcement Learning, Australian Reinforcement Learning Group.
- 2019 *Abstraction and Meta Reinforcement Learning*, NeurIPS Workshop on Meta-Learning.
Abstraction in Reinforcement Learning, UC Berkeley.
Abstraction in Reinforcement Learning, U. Mass..
Abstraction in Reinforcement Learning, UT Austin.
Abstraction in Reinforcement Learning, Oxford CS.
Abstraction in Reinforcement Learning, Oxford Philosophy.
- 2018 *State Abstraction in Reinforcement Learning*, UCLA.
State Abstraction in Reinforcement Learning, USC.
State Abstraction in Reinforcement Learning, UCSD.
State Abstraction in Reinforcement Learning, CU Boulder.
State Abstraction in Reinforcement Learning, Princeton.
State Abstraction in Reinforcement Learning, Oregon State.
State Abstraction in Reinforcement Learning, Baidu Sunnyvale.
Bandit-Based Solar Panel Control, Computational Sustainability Seminar.
- 2017 *How Artificial Intelligence Should Model the World*, Research Matters.
Abstraction and Lifelong Reinforcement Learning, Carnegie Mellon University.
Artificial Intelligence; How Intelligent, and How Soon?, Nordea Markets NYC.
Abstraction and Reinforcement Learning, MSR Redmond.

SERVICE

- Reviewing 2024: CoLLAs, ICML, JMLR, OpenMind, RLC
2023: AAAI, AISTATS, CoLLAs, ICML, JMLR, Nature, NeurIPS, TMLR, Workshops.
2022: AISTATS, CoLLAs, EWR, ICML, ICLR, NeurIPS, RLDM, JMLR, TMLR.
2021: AIJ, ICML, IEEE, JMLR, NeurIPS, Workshops.
2020: AAAI, AIJ, JAIR, JMLR, MLJ, ICML, ICLR, IEEE, NeurIPS, Workshops.
2019: AAAI, ACM, JMLR, ICLR, ICML, MLJ, NeurIPS, RLDM, Workshops.
2018: JMLR, ICML, IEEE, MLJ, NeurIPS, Workshops.
2017: ICAPS, IEEE.
- Service JMLR Editorial Board, 2020, 2021, 2022, 2023, 2024.
Area Chair, NeurIPS: 2023, 2024.
Area Chair, AAAI: 2023.
Senior PC Member, RL Conference: 2024.
Senior PC Member, CoLLAs: 2022, 2023, 2024.
PC for NeurIPS Workshop: 2021.
PC for IJCAI Workshop: 2021.
PC for NeurIPS Workshop: 2020.
PC for ICML Workshop: 2020.
- Events Co-Chair of RLC Workshop on Finding the Frame, 2024.
Co-Chair of RLDM Workshop on RL as a Model of Agency: 2022.
Co-Chair of ICLR Social on Philosophy of AGI: 2021.
Co-Chair of RLDM Workshop on Moral Decision Making: 2019.
- Brown CS Faculty-Grad Student Liaison: 2018-2019
Ph.D Recruiting Czar, 2018.
Grad Student Social Czar, 2016-2017.
Ph.D Mentorship Program Coordinator, 2017-2019.
Co-organizer, Reinforcement Learning Reading Group, 2017-2018.

AWARDS

- General **Outstanding Paper**, NeurIPS 2021 (*On the Expressivity of Markov Reward*).
Runner-up for the AAAI/ACM SIGAI Doctoral Dissertation Award.
Nominated for the Victor Lesser Distinguished Dissertation Award.
Top Reviewer, AISTATS 2022; ICML 2018, 2019, 2020, 2021; NeurIPS 2019, 2020.
- Brown **Presidential Award for Excellence in Teaching**, *University wide award presented to four graduate students per year for outstanding pedagogical achievement.*
University **Open Graduate Fellowship for Masters in Philosophy**, *University wide award granted to pursue a second masters during the Ph.D.*
Great Teaching Assistant Award, *Artificial Intelligence.*
Great Teaching Assistant Award, *A First Byte of Computer Science.*
- Carleton **Distinction in Major and Thesis**, *Philosophy.*
College **Distinction in Major and Capstone**, *Computer Science.*