

David Abel

Research Scientist, DeepMind

📍 London, UK

🏠 [david-abel.github.io](https://github.com/david-abel)

✉️ dmabel@google.com

EDUCATION

- 2020 **Ph.D in Computer Science**, *Brown University*, Providence, RI.
Advisor: Michael L. Littman.
Thesis: A Theory of Abstraction in Reinforcement Learning.
Committee: George Konidaris, Stefanie Tellex, Peter Stone, Will Dabney.
- 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.

PUBLICATIONS

Conference Papers

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 Umanhowar, 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.

Journal Papers

The Value of Abstraction.

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

Workshops, Symposia, and Extended Abstracts

Skill Discovery with Well-Defined Objectives.

Yuu Jinnai, David Abel, Jee Won Park, D. Ellis Hershkowitz, Michael L. Littman, and George Konidaris.
ICLR Workshop on Structure and Priors in Reinforcement Learning, 2019.

simple_rl: Reproducible Reinforcement Learning in Python.

David Abel.
ICLR Workshop on Reproducibility in Machine Learning, 2019.

Boundedly Rational Concept Formation.

David Abel, Mark K. Ho.
IACAP 2019.

A Theory of State Abstraction for Reinforcement Learning.

David Abel.

AAAI Doctoral Consortium 2019.

Toward Good Abstractions for Lifelong Learning.

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

NeurIPS Workshop on Hierarchical Reinforcement Learning, 2017.

Toward Improving Solar Panel Efficiency using Reinforcement Learning.

David Abel, Emily Reif, Edward C. Williams, Michael L. Littman.

EnvirolInfo 2017, early version at RLDM 2017.

Agent-Agnostic Human-in-the-Loop Reinforcement Learning.

David Abel, John Salvatier, Andreas Stuhlmüller, Owain Evans.

NeurIPS Workshop on the Future of Interactive Learning Machines, 2016.

Exploratory Gradient Boosting for Reinforcement Learning in Complex Domains.

David Abel, Alekh Agarwal, Fernando Diaz, Akshay Krishnamurthy, Robert Schapire.

ICML Workshop on Abstraction in Reinforcement Learning, 2016.

Reinforcement Learning as a Framework for Ethical Decision Making.

David Abel, James MacGlashan, Michael L. Littman.

AAAI Workshop on AI, Society, & Ethics, 2016.

Affordances as Transferrable Knowledge for Planning Agents.

Gabriel Barth-Maron, David Abel, James MacGlashan, Stefanie Tellex.

AAAI Symposium on Knowledge and Skill Transfer, 2014.

Toward Affordance-Aware Planning.

David Abel, Gabriel Barth-Maron, James MacGlashan, Stefanie Tellex.

RSS Workshop on Affordances in Vision for Cognitive Robotics, 2014.

TEACHING EXPERIENCE

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

Guest Lecturer **On Research**, 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 ×3, Carleton College.

Volunteer **Hour of Code**, Nathan Bishop Middle School.
Work **Math Tutor**, OES Elementary School.
Chess Coach, Vose Elementary School.

PROFESSIONAL EXPERIENCE

2020–Present **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

2020 *Abstraction in Reinforcement Learning*, Michigan, MSR Redmond, NYU.

2019 *Abstraction and Meta Reinforcement Learning*, NeurIPS Workshop on Meta-Learning.
Abstraction in Reinforcement Learning, UC Berkeley, U. Mass., UT Austin, Oxford

2018 *State Abstraction in Reinforcement Learning*, UCLA, USC, UCSD, CU Boulder, Princeton, Oregon State, 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 Melon University
Artificial Intelligence; How Intelligent, and How Soon?, Nordea Markets NYC
Abstraction and Reinforcement Learning, MSR Redmond

AWARDS

Brown University **Presidential Award for Excellence in Teaching**.
Open Graduate Fellowship for Masters in Philosophy.
Great Teaching Assistant Award, *Artificial Intelligence*.
Great Teaching Assistant Award, *A First Byte of Computer Science*.

Carleton College **Distinction in Major and Thesis**, *Philosophy*.
Distinction in Major and Capstone, *Computer Science*.

Travel Awards **RLDM, ICML, AAI, NeurIPS, Brown**.

Misc. **Top Reviewer Award**, ICML 2018, 2019, 2020; NeurIPS 2019, 2020.
The Automated Poet, Discover Art Show, Providence, RI.

SERVICE

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

Committees JMLR Editorial Board Reviewer.

RLDM Workshop on Moral Decision Making, Co-Chair w/ Dan Kasenberg, 2019

Reviewing 2020: AAAI, AIJ, JAIR, JMLR, MLJ, ICLR, ICML, IEEE, NeurIPS, Workshops

2019: AAAI, ACM, JMLR, ICLR, ICML, MLJ, NeurIPS, Workshops

2018: JMLR, ICML, IEEE, MLJ, NeurIPS, Workshops

2017: ICAPS, IEEE