

Peter Henderson

✉ phend@cs.stanford.edu

🏠 www.peterhenderson.co | [in phende](https://www.linkedin.com/in/phende) | [f Breakend](https://www.facebook.com/Breakend) | [G dy_JBs0AAAAJ](https://github.com/dy_JBs0AAAAJ) (list of publications)

EDUCATION

J.D./Ph.D. Computer Science

September 2018 -

Stanford University, Stanford, California

Ph.D. Advisor: Dan Jurafsky

Activities: Natural Language Processing Group, Graduate Fellow at the Regulation, Evaluation, and Governance Lab, Stanford Domestic Violence Pro Bono Project (leadership), Stanford Native Law Pro Bono Project, Three Strikes Project, OpenPhilanthropy AI Fellow

M.Sc. Computer Science

2016 - 2018

McGill University, Montréal, Québec

Thesis: Reproducibility and Reusability in Deep Reinforcement Learning

Advisors: David Meger, Joelle Pineau

Activities: Institut des algorithmes d'apprentissage de Montréal (MILA), Centre for Intelligent Machines (CIM), Mobile Robotics Lab (MRL)

B.Eng. Software Engineering

2011 - 2015

McGill University, Montréal, Québec

Thesis: Autonomous Swarm Behaviour in Mesh Networked Agents

Advisor: Mark Coates

HONORS

Carole Hafner Best Paper Award, Int'l Conference on Artificial Intelligence and Law 2021

Gerald Gunther Prize for Outstanding Performance in Property Law 2021

OpenPhilanthropy AI Fellowship 2020 - Present

Stanford Computer Science Departmental 1st Year Fellowship 2018

James McGill Scholarship 2011

Andrew Fayne Memorial Scholarship 2011

INDUSTRY

Applied Scientist

October 2017 - July 2018

EXPERIENCE

[Amazon](#) (Alexa Conversational AI)

Cambridge, MA

- As part of the Conversational AI group led by Alborz Geramifard, developed deep reinforcement learning dialogue management policies which learned to ask clarifying questions under uncertainty in ASR/NLU processing.

Software Development Engineer

June 2015 - August 2016

[Amazon Web Services](#) (New Initiatives, Storage Gateway)

Cambridge, MA

- Developed an internal tool to analyze high severity events and provide recommended actions to drive quicker response times.
- Developed tools for migrating live production database systems safely and with no downtime.
- Developed and architected the core code of [AWS Storage Gateway for Files](#) along with a small agile team.
- Participated in maintaining the live production Storage Gateway system and developed features to address customer issues.

Software Engineering Intern

May–August 2014

[A Thinking Ape](#) (Data Analytics, Realtime Bidding)

Vancouver, BC

- Worked on Realtime Bidding System for buying mobile ad space using Mopub and OpenX.
- Lead developer of data management, visualization, and analysis tools for development of bidding models.
- Developed features for and saw Realtime Bidding system to a soft launch with a bid round trip time of < 100 ms for 10,000+ requests per second.

Software Engineering Intern

January–August 2013

[Ericsson](#) (Web Communication Gateway)

Montréal, QC

- Worked on Web Communication Gateway, a Java backend application for exposing Real Time Communication services (RCS) to clients over a REST interface that can handle

- millions of users per rack and hundreds of thousands of users per instance. Saw product redesign to soft launch.
- Developed and designed proprietary TTCN-3 Framework for SIP/HTTP/MSRP signaling and automated with Jenkins

LAW+POLICY
EXPERIENCE

- California Supreme Court** June 2021 - August 2021
Summer Extern with Justice Cuéllar
- Evaluated and drafted memoranda on petitions for review. Wrote long-form memoranda on topics related to criminal, civil, and administrative law to assist with draft opinions.
- Internal Revenue Service** July 2020 - Present
Office of Research, Applied Analytics and Statistics
- Working as a volunteer researcher on the National Research Program Redesign effort.
 - Developing machine learning algorithms for more robust and fair audit selection, while simultaneously retaining accurate population and sub-population estimates of the tax gap.
- AI and Crisis Stability Working Group** November 2020 - June 2021
Stanford University Center for International Security and Cooperation
- Wrote memos on AI-related risks in national security contexts and took part in related working group discussions led by Under Secretary of Defense for Policy Colin Kahl.
- Institute for Security + Technology** March 2020 - Present
Adjunct Technical Advisor NC3
- Advising on analysis of risks from machine learning use in nuclear command, control, and communications (NC3).

PEER-REVIEWED
JOURNAL
PUBLICATIONS

- [1] Mark Krass, **Peter Henderson**, Michelle M. Mello, David M. Studdert, and Daniel E. Ho. “Artificial Intelligence, Covid-19 and Law: The Need for Evaluation in the Age of Many Models.” *BMJ*. 2021.
- [2] **Peter Henderson**, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. “Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning.” *Journal of Machine Learning Research*. 2020.
- [3] Vincent Francois-Lavet, **Peter Henderson**, Riashat Islam, Marc G. Bellemare, and Joelle Pineau. “An Introduction to Deep Reinforcement Learning.” *Foundations and Trends in Machine Learning*. 2018.
- [4] Iulian Vlad Serban, Ryan Lowe, **Peter Henderson**, Laurent Charlin, and Joelle Pineau. “A survey of available corpora for building data-driven dialogue systems.” *Dialogue and Discourse*. 2018.

PEER-REVIEWED
CONFERENCE
PUBLICATIONS

- [5] Lucia Zheng*, Neel Guha*, Brandon R. Anderson, **Peter Henderson**, Daniel E. Ho. When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset. *International Conference on Artificial Intelligence and Law (ICAIL)*, 2021. (**Carole Hafner Best Paper Award**)
- [6] Joshua Romoff, **Peter Henderson**, David Kanaa, Emmanuel Bengio, Ahmed Touati, Pierre-Luc Bacon, Joelle Pineau. TDprop: Does Jacobi Preconditioning Help Temporal Difference Learning? *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2021.
- [7] Dallas Card, **Peter Henderson**, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, Dan Jurafsky. With Little Power Comes Great Responsibility. *Empirical Methods in Natural Language Processing (EMNLP)*. 2020.
- [8] Joshua Romoff*, **Peter Henderson***, Ahmed Touati, Yann Ollivier, Joelle Pineau, Emma Brunskill. Separating value functions across time-scales. *International Conference on*

Machine Learning (ICML) and *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*. 2019.

- [9] **Peter Henderson**, Matthew Vertescher, David Meger, and Mark Coates. “Cost Adaptation for Robust Decentralized Swarm Behaviour.” *Proceedings of The IEEE International Conference on Intelligent Robots and Systems (IROS)*. 2018.
- [10] Joshua Romoff*, **Peter Henderson***, Alexandre Piché, Vincent Francois-Lavet, and Joelle Pineau. “Reward Estimation for Variance Reduction in Deep Reinforcement Learning.” *Proceedings of the Conference on Robot Learning (CoRL)*. 2018.
- [11] **Peter Henderson***, Riashat Islam*, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. “Deep Reinforcement Learning that Matters.” *Proceedings of 32nd AAAI Conference on Artificial Intelligence (AAAI)*. 2018.
- [12] **Peter Henderson**, Wei-Di Chang, Pierre-Luc Bacon, David Meger, Joelle Pineau, and Doina Precup. “OptionGAN: Learning Joint Reward-Policy Options using Generative Adversarial Inverse Reinforcement Learning.” *Proceedings of 32nd AAAI Conference on Artificial Intelligence (AAAI)*. 2018.
- [13] **Peter Henderson**, Koustuv Sinha, Nicolas Angelard-Gontier, Nan Rosemary Ke, Genevieve Fried, Ryan Lowe, and Joelle Pineau. “Ethical Challenges in Data-Driven Dialogue Systems.” *Proceedings of the First AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*. 2018.
- [14] Florian Shkurti, Wei-Di Chang, **Peter Henderson**, Md Jahidul Islam, Juan Camilo Gamboa Higuera, Jimmy Li, Travis Manderson, Anqi Xu, Gregory Dudek, and Junaed Sattar. “Underwater Multi-Robot Convoying using Visual Tracking by Detection.” *Proceedings of The IEEE International Conference on Intelligent Robots and Systems (IROS)*. 2017.
- PEER-REVIEWED [15] Dilip Arumugam, **Peter Henderson**, Pierre-Luc Bacon. An Information-Theoretic Perspective on Credit Assignment in Reinforcement Learning. *Biological and Artificial Reinforcement Learning Workshop (NeurIPS)*. 2020.
- WORKSHOP [16] Joshua Romoff, **Peter Henderson**, David Kanaa, Emmanuel Bengio, Ahmed Touati, Pierre-Luc Bacon, Joelle Pineau. TDprop: Does Jacobi Preconditioning Help Temporal Difference Learning? *Theoretical Foundations of RL Workshop (ICML)* and *Beyond First order methods in ML Systems Workshop (ICML)*. 2020.
- PUBLICATIONS [17] **Peter Henderson** and Emma Brunskill. “Distilling Information from a Flood: A Possibility for the Use of Meta-Analysis and Systematic Review in Machine Learning Research.” *Critiquing and Correcting Trends in Machine Learning Workshop (NeurIPS)*. 2018.
- [18] **Peter Henderson**, Joshua Romoff, and Joelle Pineau. “Where Did My Optimum Go?: An Empirical Analysis of Gradient Descent Optimization in Policy Gradient Methods.” *European Workshop on Reinforcement Learning (EWRL)*. 2018.
- [19] Christos Tsirigotis, Xavier Bouthillier, François Corneau-Tremblay, **Peter Henderson**, Reyhane Askari, Samuel Lavoie-Marchildon, Tristan Deleu, Dendi Suhubdy, Michael Noukhovitch, Frédéric Bastien, and Pascal Lamblin. “Orion: Experiment Version Control for Efficient Hyperparameter Optimization.” *Reproducibility in Machine Learning Workshop (ICML)*. 2018.
- [20] Joshua Romoff, Alexandre Piché, **Peter Henderson**, Vincent Francois-Lavet, and Joelle Pineau. “Reward Estimation for Variance Reduction in Deep Reinforcement Learning.” *International Conference on Learning Representations – Workshop Track (ICLR)*. 2018.
- [21] Riashat Islam*, **Peter Henderson***, Maziar Gomrokchi, and Doina Precup. “Reproducibility of Benchmarked Deep Reinforcement Learning Tasks for Continuous Control.” *Reproducibility in Machine Learning Workshop (ICML)*. 2017.

- [22] **Peter Henderson**, Wei-Di Chang, Florian Shkurti, Johanna Hansen, David Meger, and Gregory Dudek. “Benchmark Environments for Multitask Learning in Continuous Domains.” *Lifelong Learning: A Reinforcement Learning Approach Workshop (ICML)*. 2017.
- [23] **Peter Henderson***, Thang Doan*, Riashat Islam, and David Meger. “Bayesian Policy Gradients via Alpha Divergence Dropout Inference.” *Bayesian Deep Learning Workshop (NeurIPS)*. 2017.
- [24] Maryam Fazel-Zarandi, Shang-Wen Li, Jin Cao, Jared Casale, **Peter Henderson**, David Whitney, and Alborz Geramifard. “Learning Robust Dialog Policies in Noisy Environments.” *Conversational AI Workshop, (NeurIPS)*. 2017.

POLICY REPORTS
AND ARTICLES

- [25] On the Opportunities and Risks of Foundation Models
Report in collaboration with the Stanford Center for Research on Foundation Models
- [26] Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies
Report submitted to the Administrative Conference of the United States (2020)
- [27] Prioritizing Public Health Resources for COVID-19 Investigations: How Administrative Data Can Protect Vulnerable Populations
Health Affairs Blog (2020)
- [28] Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims
Report in collaboration with OpenAI, PAI, and CSER, among other organizations (2020).

LAW REVIEW
ARTICLES

- [29] *Vulnerabilities in Discovery Tech.* Harvard Journal of Law & Technology (forthcoming 2022) (with Neel Guha and Diego Zambrano)

PRE-PRINTS

- [30] Nicolas Gontier, Koustuv Sinha, **Peter Henderson**, Iulian Serban, Michael Noseworthy, Prasanna Parthasarathi, and Joelle Pineau. “The RLLChatbot: a solution to the ConvAI challenge.” *ArXiv PrePrint*. 2018.
- [31] **Peter Henderson**, Koustuv Sinha, Rosemary Nan Ke, and Joelle Pineau. “Adversarial Gain” *ArXiv PrePrint*. 2018.
- [32] **Peter Henderson** and Matthew Vertescher. “An Analysis of Parallelized Motion Masking Using Dual-Mode Single Gaussian Models.” *ArXiv PrePrint*. 2017.
- [33] **Peter Henderson** and Muthucumar Maheswaran. “Chaotic Memory Randomization for Securing Embedded Systems.” *ArXiv PrePrint*. 2016.
- [34] **Peter Henderson**. “Implanted intracortical electrodes as chronic neural interfaces to the central nervous system.” *PeerJ PrePrint*. 2015.

INVITED TALKS
AND PANELS

- How blockers can turn into a paper: A retrospective on “Towards The Systematic Reporting of the Energy and Carbon Footprints of Machine Learning”** 2020
ML Retrospectives Workshop at ICML
Online
- Separating Value Functions Across Time-scales** 2019
Center for Human Compatible Artificial Intelligence (CHAI) Seminar at UC Berkeley
Berkeley, USA
- Panel: What Are the Key Obstacles Preventing the Progression and Application of Deep RL in Industry?** 2019
Rework Deep Reinforcement Learning Summit
San Francisco, USA
- Benchmarking and Evaluation in Inverse Reinforcement Learning** 2018
New Benchmarks, Metrics, and Competitions for Robotic Learning Workshop at RSS

Pittsburgh, USA

**Reproducibility and Replicability in Deep Reinforcement Learning
(and Other Deep Learning Methods)** 2018
Statistical Society of Canada Annual Meeting
Montréal, Canada

Tutorial on Policy Gradients for Continuous Control 2017
Reinforcement Learning Summer School, Montréal, Canada
Montréal, Canada

**Show Me the Data! On the Reproducibility of Policy Gradient Methods
for Continuous Control** 2017
Reinforcement Learning Summer School
Montréal, Canada

TEACHING **Teaching Assistant**, Stanford University
CS 384: Ethical and Social Issues in Natural Language Processing Spring 2020
EMED 111A: Emergency Medical Technician Training Fall 2018
Teaching Assistant, McGill University
Comp 202: Introduction to Computing Winter 2017
Comp 303: Software Design Fall 2016
Math 303: Discrete Mathematics Winter 2015

EDITOR AAI AI Magazine Special Issue on Conversational AI 2019

ORGANIZER Machine Learning Retrospectives Workshop (NeurIPS 2020), Machine Learning Retrospectives
Workshop (NeurIPS 2019), Workshop on Reproducibility in Machine Learning (ICML 2018)

PROGRAM First Workshop on Causal Inference and NLP (EMNLP 2021), Machine Learning Retrospectives
COMMITTEE Workshop (NeurIPS 2019), Reproducibility in Artificial Intelligence Workshop (AAAI 2019),
Workshop on Reproducibility in Machine Learning (ICML 2018)

REVIEWER ICML Workshop on Reproducibility in Machine Learning 2018, JMLR 2018/2020, IROS 2017/2018,
CoRL 2018, ICLR 2018/2020, AAI Reproducibility in Artificial Intelligence Workshop 2019,
ICLR Task Agnostic Reinforcement Learning Workshop 2019, NeurIPS ML Retrospectives
Workshop 2019, NeurIPS Reproducibility Challenge 2019/2020, IEEE Robotics and Automation
Letters 2020, Georgetown Center for Security and Emerging Technology External Reviewer

LANGUAGES English (Fluent), Russian (Fluent), Spanish (Intermediate)

CERTIFICATIONS EMT-B (Massachusetts, NREMT)

MORE More information, a complete list of publications, and auxiliary documents can be found at
INFORMATION <http://www.peterhenderson.co/>.