# Peter Henderson

www.peterhenderson.co | phende | Breakend | dy_JBs0AAAAJ

## Education

**Ph.D. Computer Science** — September 2018 -
*Stanford University*, Stanford, California
Advisor: Dan Jurafsky
Natural Language Processing Group, Graduate Fellow at the Regulation, Evaluation, and Governance Lab

**M.Sc. Computer Science** — 2016 - 2018
*McGill University*, Montréal, Québec
Thesis: Reproducibility and Reusability in Deep Reinforcement Learning
Advisors: David Meger, Joelle Pineau
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
Advisors: Mark Coates

## Industry Experience

**Applied Scientist** — October 2017 - July 2018
Amazon (Alexa Conversational AI) — Cambridge, MA

**Software Development Engineer** — June 2015 - August 2016
Amazon Web Services (New Initiatives, Storage Gateway) — Cambridge, MA

**Software Engineering Intern** — May–August 2014
A Thinking Ape (Data Analytics, Realtime Bidding) — Vancouver, BC

**Software Engineering Intern** — January–August 2013
Ericsson (Web Communication Gateway) — Montréal, QC

## Fellowships, Scholarships, and Grants

| | |
|---|---|
| OpenPhilanthropy AI Fellowship | 2020 |
| Stanford Computer Science Departmental 1st Year Fellowship | 2018 |
| James McGill Scholarship | 2011 |
| Andrew Fayne Memorial Scholarship | 2011 |

## Journal Publications

[1] 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.

[2] 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.

## Refereed Conference Publications

[3] 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.

[4] **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.

[5] 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.

[6] **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.

[7] **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.

[8] **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.

[9] 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.

## Refereed Workshop Publications

[10] **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.

[11] **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.

[12] 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. "Oríon: Experiment Version Control for Efficient Hyperparameter Optimization." *Reproducibility in Machine Learning Workshop (ICML)*. 2018.

[13] 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.

[14] 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.

[15] **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.

[16] **Peter Henderson***, Thang Doan*, Riashat Islam, and David Meger. "Bayesian Policy Gradients via Alpha Divergence Dropout Inference." *Bayesian Deep Learning Workshop (NeurIPS)*. 2017.

[17] 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.

## Pre-prints

[18] **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." *ArXiv PrePrint*. 2020.

[19] 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.

[20] **Peter Henderson**, Koustuv Sinha, Rosemary Nan Ke, and Joelle Pineau. "Adversarial Gain" *ArXiv PrePrint*. 2018.

[21] **Peter Henderson** and Matthew Vertescher. "An Analysis of Parallelized Motion Masking Using Dual-Mode Single Gaussian Models." *ArXiv PrePrint*. 2017.

[22] **Peter Henderson** and Muthucumaru Maheswaran. "Chaotic Memory Randomization for Securing Embedded Systems." *ArXiv PrePrint*. 2016.

[23] **Peter Henderson.** "Implanted intracortical electrodes as chronic neural interfaces to the central nervous system." *PeerJ PrePrint*. 2015.

## Policy Work

Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies 2020
*Report submitted to the Administrative Conference of the United States*

Prioritizing Public Health Resources for COVID-19 Investigations: How Administrative Data Can Protect Vulnerable Populations 2020
*Health Affairs Blog*

Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims 2020
*Policy Report*

## Invited Talks and Panels

**Separating Value Functions Across Time-scales** 2019
*Center for Human Compatible Aritifial 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 | AAAI AI Magazine Special Issue on Conversational AI | 2019 |

| | |
|---|---|
| Organizer | Machine Learning Retrospectives Workshop (NeurIPS 2019), Workshop on Reproducibility in Machine Learning (ICML 2018) |
| Program Committee | Machine Learning Retrospectives 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, IROS 2017/2018, CoRL 2018, ICLR 2018/2020, AAAI Reproducibility in Artificial Intelligence Workshop 2019, ICLR Task Agnostic Reinforcement Learning Workshop 2019, NeurIPS ML Retrospectives Workshop 2019, NeurIPS Reproducibility Challenge 2019, IEEE Robotics and Automation Letters 2020 |
| Graduate-level Coursework | *Stanford University*<br>Foundations of Statistical and Scientific Inference, Cognitive Neuroscience, Statistical Mechanics, Human Neuroimaging Methods, Machine Learning and Causal Inference, Foreign Policy Decision Making in Comparative Perspective, Regulating Artificial Intelligence, Experimental and Behavioral Economics, Matching and Market Design, Econometrics, Meta-research and Meta-analysis, Computational Complexity, Advanced Immigration Policy Reform<br>*McGill University*<br>Machine Learning, Reinforcement Learning, Applied Machine Learning, Natural Language Processing, Network Analysis, Intelligent Robotics, Computer Graphics, Fundamentals of Computer Vision, Micro- and Nano-Bioengineering |
| Languages | English (Native), Russian (Heritage/Native), Spanish (Intermediate), French (Beginner) |
| Certifications | First Aid/CPR, EMT-B (Massachusetts, NREMT), CITI IRB BioMed/GCP Research |
| More Information | More information and auxiliary documents can be found at http://www.peterhenderson.co/. |