

WILLIAM B. FEDUS

EDUCATION

- University of Montreal* September 2020 (Expected)
Doctor of Philosophy in Computer Science
Montreal, QC
Advisors: Yoshua Bengio and Hugo Larochelle
- University of California, San Diego* June 2016
Masters of Science in Physics
La Jolla, CA
Advisors: David Meyer (Math) and Garrison Cottrell (CSE)
- Massachusetts Institute of Technology* June 2010
Bachelors of Science, Physics
Cambridge, MA
- University of Cambridge, Robinson College* June 2009
Physics (Junior Year)
Cambridge, England

RESEARCH AND EXPERIENCE

- Google Brain* April 2018 - Present
Student Researcher Scholar Program, Montreal, QC and Mountain View CA
- Led three projects on fundamental reinforcement learning (RL) research. The first characterized time-preferences of reinforcement learning (RL) agents and designed new agents with more flexible discount functions. The second empirically established the issue of catastrophic forgetting for RL agents within a single environment. The third formed the connection between n-step returns and the experience replay capacity.
 - Helped build a python API and benchmark the first neural network architectures as part of a hybrid research-engineering team designing a large-scale distributed infrastructure for massive, asynchronous, shared models.
 - Co-leading an extension of mixture-of-expert neural (MoE) Transformer networks two orders of magnitude (1T+ parameters) on self-supervised natural language fill-in-the-blank tasks.
- Montreal Institute of Learning Algorithms (Mila) at UdeM* September 2017 – Present
Graduate Student Researcher, Montreal, QC
- Co-led a branch of reinforcement learning (RL) research investigating intrinsically motivated agents that learn by discovering controllable and independently variable factors of the environment.
 - Co-advised in research in Generative Adversarial Networks (GANs) for natural language which established new quality-diversity metrics analogous to a receiver-operating curve to assess progress.
- Google Brain* December 2016 - August 2017
Software Engineering and Research Intern, Mountain View, CA
- Extended Generative Adversarial Networks (GANs) to natural language processing tasks by using a discriminator to set rewards in order to train a generative model in a reinforcement learning environment – proposed fill-in-the-blank natural language task for bidirectional sequence models later extended in BERT, XLNet and ELECTRA.
 - Co-led an investigation into dynamics of GAN-training which empirically questioned the viewpoint that divergence minimization was governing learning dynamics.

Google Research, Accelerated Sciences June 2016 - December 2016
Software Engineering and Research Intern, Mountain View, CA

- Designed, in collaboration with Calico, a spatiotemporal 3D convolutional model with temporal dilation to perform semantic segmentation in microscopy video to identify Yeast cell-types.
- Benchmarked scaling properties of a convolutional neural network (CNN) that predicts the fluorescent labels using only the transmitted light of unlabeled fixed of biological samples.

Zillow Group Summer 2015
Data Scientist Intern, Seattle, WA

- Co-led the first use of deep convolutional neural networks (CNN) for home value assessment based on images and developed a recurrent neural architecture employing Long Short-Term Memory (LSTM) cells to process sets of images.
- This approach improved the home valuation algorithm, the Zestimate, and has been deployed since 2019.

Garrison Cottrell Group at UCSD October 2015 – June 2016
Graduate Student Researcher, La Jolla, CA

- Deep learning research on recurrent neural network (RNN) architectures with attention mechanisms.
- Teaching assistant for three courses in machine learning and artificial intelligence and served as head-TA for the graduate neural networks course.

David Meyer Group at UCSD June 2014 – June 2016
Graduate Student Researcher, La Jolla, CA

- Applied research in persistent homology, a topological data analysis technique, for modeling the topological structure of data.

Compact Muon Solenoid (CMS) Experiment at CERN April 2013– June 2014
Graduate Student Researcher, La Jolla, CA

- Redesigned and improved late-iteration seeding algorithms for particle track reconstruction for the CMS collaboration.
- This demonstrated the feasibility and efficacy of a more efficient seeding algorithm: an algorithm which generates candidate particle tracks. This sped-up the entire track reconstruction process without reducing efficiency of finding charged particle tracks.

Fidelity Management and Research Company August 2010 – March 2013
Global Equity Research Associate, Boston, MA

- Equity research analyst for \$200M-20B companies using a combination of fundamental and quantitative techniques in two fields: oil and gas and technology and media – top performance in a synthetic technology and media model portfolio.
- Developed a model of explorer and producer oil and gas companies which used Monte-Carlo simulations to better assess risk profiles across drilling portfolio.

MIT Dark Matter Detection (DMTPC Group) June 2007 – June 2010
Undergraduate Student Researcher, Cambridge, MA

- Designed and optimized our working prototype by analyzing calibration data and devising new mechanical systems for gas flow.
- Researched physically motivated classification algorithms to improve DMTPC's sensitivity to incident dark matter directionality.

Fidelity Management and Research Company Summer 2008 & 2009
Global Equity Research Associate Intern, Boston, MA

- Analyzed the impact of global telecom non-SMS mobile data adoption and completed an investment breakdown of the water industry. Correctly identified top performing and worst performing equities within my assigned sector.

- Developed & tested nine atmospheric science experiments now used in the curricula at MIT and five other universities.

PUBLICATIONS AND PRESENTATIONS

- [1] W. Fedus, C. Gelada, Y. Bengio, M. G. Bellemare, H. Larochelle. Learning Multiple Time Horizons for Improved Performance and Generalized Discounting. (*In review at ICML*), 2020.
- [2] W. Fedus, P. Ramachandran, R. Agarwal, Y. Bengio, M. Rowland, W. Dabney. Revisiting Fundamentals of Experience Replay. (*In review at ICML*), 2020.
- [3] W. Fedus*, D. Ghosh*, J. D. Martin, Y. Bengio, M. G. Bellemare, H. Larochelle. On Catastrophic Forgetting in Atari 2600 Games. (*In review at ICML*), 2020.
- [4] W. Fedus*, D. Ghosh*, J. Martin, M. G. Bellemare, H. Larochelle. MEMENTO: Further Progress Through Forgetting. *NeurIPS BARL Workshop* (Oral Presentation), 2019.
- [5] V. Jain, W. Fedus, H. Larochelle, D. Precup, M. G. Bellemare. Algorithmic Improvements for Deep Reinforcement Learning applied to Interactive Fiction. *AAAI*, 2020 (Oral Presentation).
- [6] A. Ali Taiga, W. Fedus, M. Machado, A. Courville, M. G. Bellemare. Benchmarking Bonus-Based Exploration Methods on the Arcade Learning Environment. *ICML Exploration Workshop*, 2019 (Oral Presentation; Best Paper Award) and *ICLR*, 2020.
- [7] W. Fedus, C. Gelada, Y. Bengio, M. G. Bellemare, H. Larochelle. Hyperbolic Discounting and Learning Over Multiple Horizons. *RLDM*, 2019 (Oral presentation; Best Paper Award).
- [8] P. Veličković, W. Fedus, W.L. Hamilton, P. Liò, Y. Bengio, R.D. Hjelm. Deep Graph Infomax. *NeurIPS Workshop*, 2018 and *ICLR 2019*.
- [9] M. Caccia*, L. Caccia*, W. Fedus, H. Larochelle, J. Pineau, L. Charlin. Language GANs Falling Short. *NeurIPS Critiquing Machine Learning Trends Workshop*, 2018; *ICLR*, 2020.
- [10] A. Goyal, P. Brakel, W. Fedus, T. Lillicrap, S. Levine, H. Larochelle, Y. Bengio. Recall Traces: Backtracking Models for Efficient Reinforcement Learning. *ICLR*, 2019.
- [11] V. Thomas*, E. Bengio*, W. Fedus*, J. Pondard, P. Beaudoin, H. Larochelle, J. Pineau, D. Precup, Y. Bengio. Disentangling the independently controllable factors of variation by interacting with the world. *NIPS Workshop*, 2017 (Spotlight presentation).
- [12] W. Fedus*, M. Rosca*, B. Lakshminarayanan, A. M. Dai, S. Mohamed, I. Goodfellow. Many Paths to Equilibrium: GANs Do Not Need to Decrease a Divergence At Every Step. *ICLR*, 2018.
- [13] W. Fedus, I. Goodfellow, A. M. Dai. MaskGAN: Better Text Generation via Filling in the _____. *ICLR*, 2018.
- [14] E. Christiansen, S. Yang, D. Ando#, A. Javaherian#, G. Skibinski#, S. Lipnick#, E. Mount†, A. O’Neil†, K. Shah†, A. K. Lee†, P. Goyal†, W. Fedus†, R. Poplin†, A. Esteva, L. Rubin, P. Nelson*, S. Finkbeiner*. In Silico Labeling. *Cell Journal*, 2018.
- [15] W. Fedus, M. Gartner, A. Georges, D. A. Meyer, D. Rideout. Persistent Homology for Mobile Phone Data Analysis. *Netmob 2015 Conference at MIT*, 2015.
- [16] J.P. Lopez, S. Ahlen, J. Battat, T. Caldwell, M. Chernicoff, C. Deaconu, D. Dujmic, A. Dushkin, W. Fedus, P. Fisher, F. Golub, S. Henderson, A. Inglis, A. Kaboth, G. Kohse, L. Kirsch, R. Lanza, A. Lee, J. Monroe, H. Ouyang, T. Sahin, G. Sciolla, N. Skvorodnev, H. Tomita, H. Wellenstein, I. Wolfe, R. Yamamoto, H. Yegoryan. Background Rejection in the DMTPC Dark Matter Search Using Charge Signals. *Proceedings of the DPF-2011 Conference*, 2011.
- [17] S. Ahlen, J. B. R. Battat, T. Caldwell, C. Deaconu, D. Dujmic, W. Fedus, P. Fisher, F. Golub, S. Henderson, A. Inglis, A. Kaboth, G. Kohse, R. Lanza, A. Lee, J. Lopez, J. Monroe, T. Sahin, G. Sciolla, N. Skvorodnev, H. Tomita, H. Wellenstein, I. Wolfe, R. Yamamoto, H. Yegoryan. First Dark Matter Search Results from a Surface Run of the 10-L DMTPC Directional Dark Matter Detector. *Physics Letters B*, 695 (124), 2011.
- [18] J.B.R. Battat, S. Ahlen, T. Caldwell, C. Deaconu, D. Dujmic, W. Fedus, P. Fisher, F. Golub, S. Henderson, A. Inglis, A. Kaboth, G. Kohse, R. Lanza, A. Lee, J. Lopez, J. Monroe, T. Sahin, G. Sciolla, N. Skvorodnev, H. Tomita, H. Wellenstein, I. Wolfe, R. Yamamoto, H. Yegoryan. DMTPC:

Dark matter detection with directional sensitivity. *International Journal of Modern Physics A*, 25:1-51,2010.

[19] W. Fedus. Reconstructing Nuclear Recoil Tracks in the Dark Matter Time Projection Chamber. *Senior Undergraduate Thesis*, 2010.

[20] S. Ahlen et al. The case for a directional dark matter detector and the status of current experimental efforts. *International Journal of Modern Physics A*, 25(1), 2010.

ACHIEVEMENTS AND HONORS

- Student Prize at NeurIPS 2019 BARL Workshop from UNIQUE Research Center
- Best Paper Award at ICML 2019 Exploration Workshop
- Best Paper Award at RLDM 2019
- PBEER fellowship. Doctoral research scholarships program for foreign students for \$35,000
- Amazon Web Services Research Grant for \$7,500
- Temporal Dynamics of Learning Center (TDLC) Small Grant for \$2,000
- Frontiers of Innovation Scholars Program (FISP) Fellowship for \$25,000
- UCSD Physics Excellence Grant for \$15,000
- UCSD SHORE Recruiting Award
- SMART Grant, Maine State Scholarship, ACT Grant at MIT
- 5th-Place Quantum Quandaries Event at National Science Olympiad
- National Society of Scholars

TEACHING AND SERVICE

- Reviewer for ICML
- Reviewer for ICLR
- Reviewer for NeurIPS
- Reviewer for AAAI
- Reviewer for Nature
- Reviewer for ICML Exploration workshop
- UCSD CSE 253, Head TA of Graduate Neural Networks. Professor Garrison Cottrell, Winter 2015.
- UCSD CSE 190, Undergraduate Neural Networks. Professor Garrison Cottrell, Fall 2015.
- UCSD CSE 150, Artificial Intelligence. Professor Garrison Cottrell, Spring 2015.
- UCSD CSE 150, Artificial Intelligence. Professor Lawrence Saul, Winter 2014.
- UCSD PHYS 2B, Electricity and Magnetism. Professor Ivan Schuller, Fall 2014.
- High School and Undergraduate tutor for Physics, Math and Computer Science, 2014-Present.
- UCSD Young Physicists Program (YPP), outreach program to junior high and high school students.