Will **Grathwohl**

wgrathwohl@cs.toronto.edu www.cs.toronto.edu/ wgrathwohl github.com/wgrathwohl

Education

University of Toronto
Toronto, ON

PhD in Computer Science — *December 2018 - Current*

- Advised by Richard Zemel and David Duvenaud
- Research focuses on generative modeling, variational inference, unsupervised learning.

University of Toronto
Toronto, ON

M.S. in Computer Science — *July 2017 - December 2018*

- Advised by Richard Zemel and David Duvenaud
- Master's Thesis: "Backpropagation Through Discrete Distributions."

Massachusetts Institute of Technology
Cambridge, MA

B.S. in Mathematics with Computer Science — *September 2010 - June 2014*

- Conducted research in the Anyscale Learning for All Group under Una-May O'Reilly.
- Conducted research in the Learning and Intelligent Systems group under Leslie Kaelbling.

Professional Experience

Google Brain
Toronto, ON

Student Researcher — *Febuary 2019 - Current*

- Re-investigating Energy-Based Models within the context of modern deep learning.
- Our work led to the publication "Your Classifier is Secretly and Energy-Based Model and You Should Treat it Like One" which received an oral presentation at ICLR 2020
- Our work led to "No MCMC for Me" which has been accepted to ICLR 2021.
- Our work led to the recent work "Oops I took a Gradient" which makes training Deep EBMs on discrete data possible.

OpenAI
San Francisco, CA

Research Intern — *May 2018 - September 2018*

- Worked with Ilya Sutskever and Durk Kingma on efficient generative models for high-dimensional data.
- Created FFJORD, the current state-of-the-art invertible generative model for high-dimensional data which received an oral presentation at ICLR 2019.

Lawrence Livermore National Labs
Livermore, CA

Computer Vision Specialist — *March 2016 - June 2017*

- Developed variational methods for improving transfer learning of deep neural network features for image prediction tasks.
- Developed novel variational autoencoder models for unsupervised representation learning for videos.
- Used deep reinforcement learning techniques with biological simulations to learn optimal drug-delivery policies for treating sepsis.

Zinc.io
San Francisco, CA

Independent Contractor — *December 2015*

- Built a system to automatically solve image captchas in real time.
- Designed and trained multiple deep neural networks for different captcha types.
- Automated system solves captchas with greater accuracy (99.7%) than human labelers from a professional captcha solving service.
- System runs in production today solving approximately 1,000,000 captchas a day.

Blackbird Technologies
Menlo Park, CA

Artificial Intelligence Engineer — *May 2015 - March 2016*

- Designed algorithms for automatic product indexing based on visual and text characteristics (mainly neural network based).
- Conducted original research on image classification and localization.
- Worked on boosting image classification generalization in data-constrained environments.
- Designed and built a fully automated data collection system which has accurately labeled millions of images to date.

Daqri *Mountain View, CA*

Computer Vision Engineer *May 2014 - May 2015*

- Designed and implemented algorithms for SLAM and visual-inertial odometry.
- Designed and implemented algorithms for 2D template tracking with accuracy and speed rivaling QUALCOMM's vuforia.
- Developed a novel method to combine keypoint optical flow with patch tracking for smooth, drift-free tracking.
- Wrote a performant SIMD linear algebra library in neon assembly code for ARM processors.

Mirador Technologies *Menlo Park, CA*

Co-Founder *May 2013 - May 2014*

- Founded a startup that created ML-powered tools to automatically detect and filter offensive content from social media platforms.
- Personally built the ML back-end powering Mirador's computer vision services.
- Collected and curated multiple labeled image datasets containing millions of images.
- Designed and trained multiple deep convolutional neural networks for pornography detection.
- Built the most accurate image moderation system on the market at the time.

Publications

Oops I Took a Gradient: Scalable Sampling for Discrete Distributions *ICML 20201*

Will Grathwohl, Kevin Swerskey, Milad Hashemi, David Duvenaud, Chris J. Maddison

Long Oral Presentation (Top 3% of submissions).

No MCMC for me: Amortized sampling for fast and stable training of energy-based models *ICLR 2021*

Will Grathwohl, Jacob Kelly, Milad Hashemi, Mohammad Norouzi, Kevin Swersky, David Duvenaud *Vienna, Austria*

Learning the Stein Discrepancy for Training and Evaluating Energy-Based Models without Sampling *ICML 2020*

Will Grathwohl, Kuan-Chieh Wang, Jorn-Henrik Jacobsen, David Duvenaud, Richard Zemel *Vienna, Austria*

Your Classifier is Secretly and Energy-Based Model and You Should Treat it Like One *ICLR 2020*

Will Grathwohl, Jackson Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, Kevin Swersky *Addis Abbaba, Ethiopia*

Oral Presentation (Top 1.5% of submissions).

Understanding the Limitations of Conditional Generative Models *ICLR 2020*

Jörn-Henrik Jacobsen, Ethan Fetaya, **Will Grathwohl** *Addis Abbaba, Ethiopia*

Invertible Residual Networks *ICML 2019*

Jens Behrmann*, **Will Grathwohl***, Ricky T. Q. Chen, David Duvenaud, Jörn-Henrik Jacobsen* *Long Beach, California*

Long Oral Presentation.

FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models *ICLR 2019*

Will Grathwohl*, Ricky T. Q. Chen*, Jesse Bettencourt, Ilya Sutskever, David Duvenaud *New Orleans, Louisiana*

Oral Presentation (Top 1.6% of submissions).

Backpropagation through the Void: Optimizing control variates for black-box gradient estimation *ICLR 2018*

Will Grathwohl, Dami Choi, Yuhuai Wu, Geoff Roeder, David Duvenaud *Vancouver, British Columbia*

Fellowships & Awards

2021 **Google PhD Fellowship in Machine Learning**, Declined due to graduation.
2020 **Didi Graduate student Award in Computer Science**, A merit-based fellowship worth $10,000.
2018 **Borealis AI Graduate Fellowship**, A $50,000, 2 year fellowship funding research in AI.
2018 **Best Paper Award**, Symposium on Advances in Approximate Bayesian Inference for FFJORD.
2017 **Huawei Prize**, A financial award based on academic and research performance.

Talks

ICLR 2021 Workshop on Energy-Based Models — *Online*

Using and Abusing Gradients for Discrete Sampling and Energy-Based Models — *May. 2021*

- Invited talk at the first ICLR workshop on Energy-Based Models

CMU Artificial Intelligence Seminar Series — *Pittsburgh, PA*

Using and Abusing Gradients for Discrete Sampling and Energy-Based Models — *Feb. 2021*

- Invited talk.

Generative Models and Uncertainty Quantification — *Copenhagen, Denmark*

Your Classifier is Secretly and Energy-Based Model and You Should Treat it Like One — *Oct. 2019*

- Invited talk at a workshop on generative models.

NIPS 2017 Deep RL Symposium — *Long Beach, California*

Backpropagation Through the Void — *Dec. 2017*

- Invited to give a contributed talk (1 of 12 selected submissions)

Blei Lab, Columbia University — *New York, New York*

Backpropagation Through the Void — *Nov. 2017*

- Presented my work from ICLR 2018 to David Blei and his research group

Courant Institute of Mathematical Sciences, New York University — *New York, New York*

Backpropagation Through the Void — *November 2017*

- Presented my work from ICLR 2018 to Rob Fergus and his research group

Borealis AI — *Toronto, ON*

Backpropagation Through the Void — *December 2017*

- Presented my work from ICLR 2018 to RBC's AI research group, Borealis.

Stanford University — *Palo Alto, CA*

Gradient Estimation for Discrete Objectives — *July 2018*

- Presented my work from ICLR 2018 to the Stanford University NLP group.

BAIR Lab — *Berkeley, CA*

Gradient Estimation for Discrete Objectives — *August 2017*

- Presented my work from ICLR 2018 to the BAIR lab at UC Berkeley.

Service

Neurips 2020

- Reviewed submissions for the conference.

Neurips 2019

- Reviewed submissions for the conference.

Neurips 2018

- Reviewed submissions for the conference.

NIPS 2017 Bayesian Deep Learning Workshop

- Member of the program committee and reviewed papers for the workshop.