

James Lucas

MACHINE LEARNING PHD CANDIDATE

Charles Street West, Toronto, ON, M4Y 1R5

+1 647-570-5772 | ✉ jlucas@cs.toronto.edu | 🏠 www.cs.toronto.edu/jlucas/ | 📺 [AtheMathmo](#) | 📺 [jlucas-94](#)

Education

University of Toronto

Toronto, Canada

PHD CANDIDATE - MACHINE LEARNING

Jan. 2017 -

Advised by Roger Grosse and Richard Zemel

Thesis title: Loss landscapes and optimization in deep learning [In progress]

Research focus: Deep learning practice/theory, optimization, few-shot learning

University of Cambridge

Cambridge, UK

MA 1ST CLASS HONS - MATHEMATICS

Oct. 2012 - Jul. 2015

Focus on: Statistics/probability theory, measure theory, optimization

Experience

NVIDIA

Toronto, ON

RESEARCH INTERN

May 2021 -

- Intern in research simulation technologies group, led by Sanja Fidler
- Developing methods for robust learning in realistic at-scale applications
- Developing generative models for 3D data

Google Brain

Toronto, ON

STUDENT RESEARCH COLLABORATOR

Apr. 2019 - Oct. 2019

- Invited to continue work within Geoffrey Hinton's group as a student research collaborator
- Researching practical improvements for deep learning optimization and deep generative models

Google Brain

Toronto, ON

RESEARCH INTERN

Jan. 2019 - Apr. 2019

- Working with Mohammad Norouzi and George Tucker.
- Developed internal framework for training and evaluating deep generative models
- Identified and corrected theoretical shortcomings of existing methods and published corresponding research paper as first-author (NeurIPS 2019)

University of Toronto

Toronto, ON

SESSIONAL INSTRUCTOR

Sep. 2017 - Dec. 2017

- Taught fourth-year Intro to Machine Learning course (CSC411) to over 400 students
- Prepared materials (lecture slides, demos, and tutorials); organized TAs and handled admin

LoopUp

San Francisco, CA

SOFTWARE ENGINEER

Sep. 2015 - Dec. 2016

- Designed and implemented in-house ML-based anomaly detection system to combat fraud
- Developed tools for load testing and data-focused profiling
- Designed and produced WebRTC screen sharing platform

Microsoft Research

Cambridge, UK

RESEARCH INTERN

Jun. 2015 - Aug. 2015

- Second undergraduate internship awarded through Bright Minds Competition
- Developed statistical models combining Gaussian processes and TrueSkill
- Evaluated models for ranking Olympic athletes over time

Microsoft Research

Cambridge, UK

RESEARCH INTERN

Jun. 2014 - Aug. 2014

- First undergraduate internship awarded through Bright Minds Competition
- Developed PGM approaches for multidimensional ranking systems using data from Dota2
- Investigated influence of data tournament-graph structure on feasibility of inference

Awards

- 2019 **Outstanding Reviewer**, ICLR
2017 **E. F. Burton And F. W. Burton Graduate Scholarship**, University of Toronto
2015 **Foundation Scholarship**, Queens' College, Cambridge
2014 / 2015 **Bright Minds competition winner**, Microsoft Research, Cambridge

Conference Publications

- 2021 **Analyzing Monotonic Linear Interpolation in Neural Network Loss Landscapes**, James Lucas, Juhan Bae, Michael R. Zhang, Stanislav Fort, Richard Zemel, Roger Grosse *ICML*
- 2021 **Theoretical bounds on estimation error for meta-learning**, James Lucas, Mengye Ren, Irene Kameni, Toniann Pitassi, Richard Zemel *ICLR*
- 2020 **Regularized linear autoencoders recover the principal components, eventually**, Xuchan Bao, James Lucas, Sushant Sachdeva, Roger Grosse *NeurIPS*
- 2019 **Don't Blame the ELBO: A linear VAE perspective on Posterior Collapse**, James Lucas, George Tucker, Roger Grosse, Mohammad Norouzi *NeurIPS*
- 2019 **Preventing Gradient attenuation in Lipschitz constrained convolutional networks**, Qiyang Li*, Saminul Haque*, Cem Anil, James Lucas, Roger Grosse, Joern-Henrik Jacobsen *NeurIPS*
- 2019 **Lookahead Optimizer: k steps forward, 1 step back**, Michael Zhang, James Lucas, Geoff Hinton, Jimmy Ba *NeurIPS*
- 2019 **Sorting out Lipschitz function approximation**, Cem Anil*, James Lucas*, Roger Grosse *ICML*
- 2019 **Aggregated Momentum: Stability Through Passive Damping**, James Lucas, Shengyang Sun, Richard Zemel, Roger Grosse *ICLR*
- 2018 **Adversarial Distillation of Bayesian Neural Network Posteriors**, Kuan-Chieh Wang, Paul Vicol, James Lucas, Li Gu, Roger Grosse, Richard Zemel *ICML*

Workshop and pre-prints

- 2020 **Flexible Few-Shot Learning of Contextual Similarity**, Mengye Ren*, Eleni Triantafillou*, Kuan-Chieh Wang*, James Lucas*, Jake Snell, Xaq Pitkow, Andreas S. Tolias, Richard Zemel *MetaLearn Neurips*
- 2020 **On Monotonic Linear Interpolation of Neural Network Parameters**, James Lucas, Juhan Bae, Michael Zhang, Jimmy Ba, Richard Zemel, Roger Grosse *OptML Neurips*
- 2019 **Information-theoretic limitations on novel task generalization**, James Lucas, Mengye Ren, Irene Kameni, Toniann Pitassi, Richard Zemel *MLWG Neurips Contributed Oral*
- 2019 **Understanding posterior collapse in generative latent variable models**, James Lucas, George Tucker, Roger Grosse, Mohammad Norouzi *DeepGenStruct ICLR*

*Equal contribution

Teaching

- 2019 **Deep Learning I - TA**, Vector Institute
2018 **CSC411 - Head TA**, University of Toronto
2017 **CSC411 - Instructor**, University of Toronto

Skills

- Programming** Python, C#, C++, MATLAB, R, Java, Javascript, Rust
ML Frameworks Tensorflow, Pytorch

Other things I do

- Parent** Proud father of a wonderful, noisy, small human.
Mentor Several of the above publications were completed with undergraduate mentees.
Game Developer I make PC games. Hoping to publish a game on steam soon!