James **Lucas**

Research Scientist

UK

jlucas@cs.toronto.edu | www.cs.toronto.edu/jlucas/ | AtheMathmo | jlucas-94

Education

University of Toronto — *Toronto, Canada*

PhD - Machine Learning — *Jan. 2017 - Oct 2022*

Advised by Roger Grosse and Richard Zemel

Thesis title: Optimization and loss landscape geometry of deep learning

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 — *Remote, UK*

Senior Research Scientist — *Jun 2023 -*

- Leading and executing on research projects to push basic ML capabilities
- Research focus on 3D deep learning, graphics, and generative modeling
- Designed algorithmic and architecture-based acceleration of generative AI

NVIDIA — *Remote, UK*

Research Scientist — *Jan 2022 - Jun 2023*

- Solving research problems for video games, autonomous driving, and content creation
- Designed, developed, and shipped material super-resolution tools as part of RTX Remix

NVIDIA — *Toronto, ON*

Research Intern — *May 2021 - Dec 2022*

- 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

Conference Publications

2024	**SpaceMesh: A continuous representation for learning manifold surface meshes**, Tianchang Shen, Zhaoshuo Li, Marc Law, Matan Atzmon, Sanja Fidler, James Lucas, Jun Gao, Nicholas Sharp	*SIGGRAPH*
2024	**LATTE3D: Large-scale amortized text-to-enhanced 3D synthesis**, Kevin Xie, Jonathan Lorraine, Tianshi Cao, Jun Gao, James Lucas, Antonio Torralba, Sanja Fidler, Xiaohui Zeng	*ECCV*
2024	**Graph metanetworks for processing diverse neural architectures**, Derek Lim, Haggai Maron, Marc Law, Jonathan Lorraine, James Lucas	*ICLR*
2023	**Spacetime Representation Learning**, Marc Law, James Lucas	*ICLR*
2023	**Bridging the sim2real gap with CARE**, Viraj Prabhu, David Acuna, Andrew Liao, Rafid Mahmood, Marc T Law, Judy Hoffman, Sanja Fidler, James Lucas	*TMLR*
2023	**ATT3D: Amortized text-to-3D object synthesis**, Jonathan Lorraine, Kevin Xie, Xiaohui Zeng, Chen-Hsuan Lin, Towaki Takikawa, Nicholas Sharp, Tsung-Yi Lin, Ming-Yu Liu, Sanja Fidler, James Lucas	*ICCV*
2022	**Optimizing data collection for Machine Learning**, Rafid Mahmood, James Lucas, Jose M. Alvarez, Sanja Fidler, Marc Law	*NeurIPS*
2022	**How much more data do I need? Estimating requirements for downstream tasks**, Rafid Mahmood, James Lucas, David Acuna, Daiqing Li, Jonah Philion, Jose M. Alvarez, Zhiding Yu, Sanja Fidler, Marc Law	*CVPR*
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

2022	**The Calibration Generalization Gap**, Annabelle Carrell, Neil Mallinar, James Lucas, Preetum Nakkiran	*arXiv preprint*
2021	**Causal Scene BERT: Improving object detection by searching for challenging groups**, Cinjon Resnick, Or Litany, Amlan Kar, Karsten Kreis, James Lucas, Kyunghyun Cho, Sanja Fidler	*AVVision ICCV*
2020	**Probing few-shot generalization with attributes**, Mengye Ren*, Eleni Triantafillou*, Kuan-Chieh Wang*, James Lucas*, Jake Snell, Xaq Pitkow, Andreas S. Tolias, Richard Zemel	*arXiv preprint*
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, Goerge Tucker, Roger Grosse, Mohammad Norouzi	*DeepGenStruct ICLR*

*Equal contribution

Awards

2023	**Real-Time Live! 1st Place Winner**, SIGGRAPH
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

Teaching

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

Academic services

2023	**Reviewer** ICML, ICLR
2022	**Reviewer** JMLR, CVPR, ICLR, NeurIPS, ICML
2021	**Reviewer** CVPR, ICML
2019	**Reviewer** NeurIPS 2019 Workshop on Machine Learning with Guarantees
2018,2019	**Reviewer** NeurIPS
2018,2019	**Reviewer** ICLR

Skills

Programming	Python, CUDA, C#, C++, MATLAB, R, Java, Javascript, Rust
ML Frameworks	Tensorflow, Pytorch, etc.

Other things I do

Parent	Proud father of two wonderful, noisy, small humans.
Mentor	Many of the above publications were completed with undergraduate mentees and graduate interns
Game Developer	I make PC games for fun!