Jonathan Peter Lorraine

Homepage: www.jonlorraine.com
Email: lorraine@cs.toronto.edu
Location: Toronto, Ontario, Canada
Citizenship: Canadian/U.S.

Professional Experience

2022-Current **Research Scientist, NVIDIA**
Working on generative modelling with Sanja Fidler's group at the Toronto AI Lab.

2017-Current **Graduate Researcher, Vector Institute**
Papers accepted to multiple top machine learning conferences. Researching hyperparameter optimization, learning in games, and nested optimization.

2021-2022 **Research Scientist Intern, Google**
Worked on research and applied engineering for an AutoML platform used across Google for production-ready models. Designed a method used by our team to select parameters for production performance with $\sim 10\times$ less compute. Led a paper on the method. Also, mentored senior engineers on best hyperparameter optimization methods, and promoted cutting-edge methods.

2021 **Research Scientist Intern, Facebook AI Research (Now Meta AI)**
Led a spotlight paper acceptance at AAMAS. Worked with Jakob Foerster (homepage) on machine learning in multi-agent systems. Advised product teams on state-of-the-art hyperparameter optimization techniques, leveraging personal research to enhance model performance and efficiency across various projects.

2016-2022 **(Lead) Teaching Assistant, University of Toronto**
Designed content, conducted lectures, and organized grading for 10+ Computer Science courses, including four graduate-level courses on ML topics such as Deep Learning and Natural Language Computing. Courses include:

- CSC2547: Automated Reasoning with Machine Learning (Winter 2023)
- CSC2626: Imitation Learning for Robotics (Fall 2022)
- CSC2547: Learning to Search (Fall 2019)
- CSC413/2516: Neural Networks and Deep Learning (Winter 2020, Winter 2021)
- CSC412/2506: Probabilistic Learning and Reasoning (Winter 2019)
- CSC411/2515: Machine Learning (Fall 2018, Winter 2021)
- CSC401/2511: Natural Language Computing (Winter 2022)
- CSC311: Introduction to Machine Learning (Fall 2020)
- CSC165: Mathematical Expressions and Reasoning for Computer Science (Fall 2016)

2016-2019 **Machine Learning Engineer / Data Scientist, Electronica AI**
Collaborated with internal and client C-suite leaders to design, test, validate, and optimize automated financial trading strategies, leveraging machine learning to improve predictive accuracy and trade execution. Led multidisciplinary teams of more than five internal stakeholders across Software Engineering, Finance, and Legal to combine quantitative methods with AI, resulting in market-making strategies with daily order flow of 200M+.

2016-2018 **Software Developer, Alkemi**
Conducted stress testing on financial exchanges to identify bottlenecks and vulnerabilities in real-world scenarios. Also, performed optimization and validation of automated trading strategies based on intra/inter-exchange price discrepancies.

2013-2016 **Research Assistant, University of Toronto**
Worked with Statistics and Operations Research Professor Dmitry Krass (homepage) to generate computational results resulting in 3 publications.

2012-2014 **Mobile Application Developer, First Class Education Software**
Developed applications for teaching university level biology on Android and iOS. (app. 1, 2)

Education

2018-2024 **Ph.D. in Computer Science, University of Toronto**
Machine learning group, advised by Professor David Duvenaud (homepage).
Thesis: Scalable Nested Optimization for Deep Learning

2016-2018 **M.Sc. in Applied Computing, University of Toronto**
Specialist in Data Science, advised by Professor David Duvenaud.
Thesis: Black-Box Optimizers for Robust and Efficient Optimization of Trading Strategies.

2012-2016 **H.B.Sc., University of Toronto**
Specialist in Computer Science, major in Mathematics, and a minor in Economics. Graduated with high distinction.

Publications

2024 Mehta, B., Lorraine, J., Masson, S., Arunachalam, R., Pervaiz Bhat, Z., Lucas, J., George Zachariah, A. **Improving Hyperparameter Optimization with Checkpointed Model Weights**
We input checkpointed model weights to hyperparameter optimization surrogate functions to featurize the architecture, and more. (link)

2024 Xie, K., Lorraine, J., Cao, T., Gao, J., Lucas, J., Torralba, A., Fidler, S., Zeng, X. **LATTE3D: Large-scale Amortized Text-To-Enhanced3D Synthesis**
We generate high-quality meshes from text robustly in 400ms by combining 3D priors, amortized optimization, and a surface rendering stage. Accepted at ECCV 2024. (link)

2024 Bae, J., Lin, W., Lorraine, J., Grosse, R. **Training Data Attribution via Approximate Unrolled Differentiation**
We introduce a training data atttribution method that works in non-converged models with multi-stage training. (link)

2024 Lim, D., Marron, H., Law, M., Lorraine, J., Lucas, J. **Graph Metanetworks for Processing Diverse Neural Architectures**
We propose a way to process diverse network architectures with equivariance to parameter permutation symmetries. Spotlight at ICLR 2024 (link)

2023 Zhang, M., Desai, N., Bae, J., Lorraine, J., Ba, J. **Using Large Language Models for Hyperparameter Optimization**

We use LLMs for optimizing hyperparameters, allowing new capabilities like treating our code as a hyperparameter. Accepted at NeurIPS 2023 FMDM Workshop. (link)

2023 Lorraine, J., Xie, K., Zeng, X., Lin, C., Takikawa, T., Sharp, N., Lin, T., Liu, M., Fidler, S., Lucas, J. **ATT3D: Amortized Text-To-3D Object Synthesis**

We train a single, amortized model to output objects for various text prompts, allowing generalization, interpolations, and reduced training cost. Accepted at ICCV 2023. (link)

2022 Lorraine, J., Anderson, N., Lee, C., De Laroussilhe, Q., Hassen, M. **Task Selection for AutoML System Evaluation**

We select relevant development tasks to accurately assess the impact of AutoML system changes on holdout tasks with different distributions, as in production. (link)

2022 Vicol, P., Lorraine, J., Duvenaud, D., Grosse, R. **On Implicit Bias in Overparameterized Bilevel Optimization**

We characterize implicit regularization from various bilevel optimization methods. Oral at ICML 2022. (link)

2022 Lorraine, J., Acuna, D., Vicol, P., Duvenaud, D. **Complex Momentum for Optimization in Games**

We generalize gradient descent with momentum, which improves convergence in adversarial games with near identical compute cost. Accepted at AISTATS 2022. (link)

2022 Lorraine, J., Vicol, P., Parker-Holder, J., Kachman, T., Metz, L., Foerster, J. **Lyapunov Exponents for Diversity in Differentiable Games**

We find multiple solutions in games by branching optimization near bifurcations, found with Lyapunov exponent based objectives. Oral at AAMAS 2022. (link)

2021 Richter-Powell, J., Lorraine, J., Amos, B. **Input Convex Gradient Networks**

We model convex gradients by integrating a JVP parameterized by a neural network. Spotlight at the Optimal Transport and ML Workshop at NeurIPS 2021. (link)

2021 Raghu, A., Lorraine, J., Kornblith, S., McDermott, M., Duvenaud, D. **Meta-learning to Improve Pre-training**

We meta-learn pre-training hyperparameters with gradients by combining iterative and implicit differentiation. Accepted at NeurIPS 2021. (link)

2020 Lorraine, J., Vicol, P., Duvenaud, D. **Optimizing Millions of Hyperparameters by Implicit Differentiation**

We jointly tune as many hyperparameters as parameters, while being only few times more costly compute than standard training. Accepted at AISTATS 2020. (link)

2019 Lorraine, J., Hossein, S. **JacNet: Learning Functions with Structured Jacobians**

We learn functions by parameterizing their derivative, allowing us to easily enforce constraints on higher-order information. Accepted at the ICML 2019 INNF Workshop. (link)

2018 MacKay, M., Vicol, P., Lorraine, J., Duvenaud, D., Grosse, R. **Self-Tuning Networks: Bilevel Optimization of Hyperparameters using Structured Best-Response Functions**

Our hypernetwork architecture scales hyperparameter optimization to modern networks. Accepted at ICLR 2019. (link)

2018 Adam, G., Lorraine, J. **Understanding Neural Architecture Search Techniques**
We investigate failure modes of neural architecture search methods, and propose solutions by modulating hidden state interpretability. (link)

2017 Lorraine, J., Duvenaud, D. **Stochastic Hyperparameter Optimization through Hypernets**
We learn a differentiable loss for hyperparameters, which can scale to thousands of dimensions. Accepted at the NIPS 2017 Meta-learning Workshop. (link)

Other Research Experience

2016 **Optimizing Facility Location and Design**
Developed a method to solve a non-linear concave knapsack problem. Accepted at the European Journal of Operational Research. (link)

2016 **Budgetary Effects on Pricing Equilibrium in Online Markets**
Extended prior work to account for vendor cost affect on equilibria. (link)

2015 **On Covering Location Problems on Networks with Edge Demand**
Made a method for finding numerical solutions to a max covering problem on a network with edge-based demand. Accepted at the journal of Computers & Operations Research. (link)

2014 **Structural Properties of Voronoi Diagrams in Facility Location with Continuous Demand**
Designed an algorithm for finding a point to add a Voronoi diagram, with a Voronoi cell that has maximal area. Accepted in Discrete Applied Mathematics. (link)

Service

2023 Reviewer: The IEEE/CVF Conference on Computer Vision and Pattern Recognition
2022 Reviewer: International Conference on Autonomous Agents and Multi-agent Systems
2022 Reviewer: International Conference on Automated Machine Learning
2021,22,23 Reviewer: International Conference on Machine Learning
2021,22 Reviewer: International Conference on Artificial Intelligence and Statistics
2020,21 Reviewer: International Conference on Learning Representations
2020,21,24 Reviewer: Conference on Neural Information Processing Systems
2018 Reviewer: NIPS Smooth Games Optimization and Machine Learning Workshop

2019-2021 Organizer: Meta-learning reading group at the University of Toronto and Vector institute jointly with Mengye Ren and Eleni Triantafillou.

Grants & Awards

2024 Monica Ryckman Bursary 4000 CAD
2024 Doctoral Completion Award: 4100 CAD
2023 Vector Research Grant: 6000 CAD
2023 Doctoral Completion Award: 4000 CAD
2022 Vector Research Grant: 3000 CAD

2022 AAMAS Student Scholarship Award: 1200 USD
2021 Vector Research Grant: 6000 CAD
2021 Ray Reiter Graduate Award: 1000 CAD
2020 Vector Research Grant: 6000 CAD
2019 Vector Research Grant: 6000 CAD
2017 MITACS Accelerate Research Grant: 30 000 CAD
2014 NSERC Undergraduate Research Award: 4500 CAD
2013 Cheng Yu Tung Grant: 4800 CAD
2012 GE-STAR Award: 4500 USD

Talks

2023 ATT3D: Amortized Text-to-3D Object Synthesis. NVIDIA Computer Vision, Remote.
2023 ATT3D: Amortized Text-to-3D Object Synthesis. NVIDIA Omniverse, Remote.
2022 AutoML and Scalable Hyperparameter Optimization. NVIDIA, Remote.
2022 Nested-optimization-aware Hyperparameter Optimization. University of Toronto. Remote.
2022 Lyapunov Exponents for Diversity in Differentiable Games. AAMAS, 2022. Remote.
2022 Bilevel, Hypergradient Hyperparameter Optimization. University of Toronto. Remote.
2022 Complex Momentum for Optimization in Games. AISTATS, 2022. Remote.
2022 Lyapunov Exponents for Diversity in Differentiable Games. Google Brain. Remote.
2021 Gradient-Based Hyperparameter Optimization. Google. Remote.
2021 AutoTFX Reading Group. Google. Remote.
2021 Using Bifurcations for Diversity in Learning. BOFM Workshop ICML, 2021. Remote.
2021 Gradient-Based Hyperparameter Optimization. Facebook AI Research. Remote.
2020 Optimizing Millions of Hyperparameters by Implicit Differentiation. AISTATS, 2020. Remote.
2019 Meta-Learning Reading Group, Vector Institute. Toronto, Canada.
2019 JacNet: Learning Functions with Structured Jacobians, ICML 2019 Invertible networks and normalizing flows workshop. Long Beach, USA.
2018 Self-Tuning Networks: Bilevel Optimization of Hyperparameters using Structured Best-Response Functions, Vector Institute. Toronto, Canada.
2018 Applied Research in Action 2018, University of Toronto. Toronto, Canada.
2017 Complex Order Book Strategies and Intra/Inter-Exchange Arbitrage, Alkemi AI. Toronto, Canada.
2017 Hyperparameter Opt. with Hypernets, NIPS Meta-Learning Workshop 2017. Long Beach, USA.
2017 Bayesian Optimization for Trading Strategies, Electronic AI. Toronto, Canada.
2016 Maximizing the Trading Area of a new Facility, Rotman School of Management. Toronto, Canada.
2015 On Covering Location Problems on Networks with Edge Demand, Rotman School of Management. Toronto, Canada.

Student Supervision & Mentorship

2021-2022 Jack Richter-Powell (McGill Undergrad. → UofT Research → MIT PhD student → NVIDIA Intern)
2018-2019 Haoping Xu (UofT Undergrad. → UofT PhD student)

Patents

2023 Lorraine, J., Xie, K., Gao, J., Fidler, S., Lucas, J., Zeng, X. **Training Data Sampling for Neural Networks**
2023 Lorraine, J., Xie, K., Fidler, S., Lucas, J. **Neural Network-Based Digital Asset Generation**

Last updated: July 1, 2024