Dami Choi

choidami@cs.toronto.edu • https://www.cs.toronto.edu/~choidami/

EDUCATION

University of Toronto 2019 – Present
PhD in Computer Science advised by Chris J. Maddison and David Duvenaud

University of Toronto 2013 – 2018
BASc in Engineering Science (Major in Electrical and Computer Engineering)
- Cumulative GPA: 3.91/4.0

RELEVANT EXPERIENCE

Transluce Aug 2024- Present
Member of Technical Staff

Constellation, Astra Fellowship Jan 2024- Jun 2024

Google, Google Research, Translate Team Sep 2022- Feb 2023
Student Researcher
- Studied the optimization dynamics of multi-task learning in the data-imbalanced setting.

Google, Google Research, Brain Team Jun 2018- Jun 2019
AI Resident
- Studied *data echoing*, a method to train neural networks faster in the presence an bottleneck in the input pipeline.
- Studied common deep learning optimizers, their relationship with each other, and their empirical performance.

PUBLICATIONS

J. Treutlein, **D. Choi**, J. Betley, S. Marks, C. Anil, R. Grosse, O. Evans "Connecting the Dots: LLMs can Infer and Verbalize Latent Structure from Disparate Training Data."

D. Choi, Y. G. Shavit, D. Duvenaud "Tools for Verifying Proofs-of-Training-Data," In *Advances in Neural Information Processing Systems 36*, 2023

D. Choi, D. Xin, J. Gilmer, H. Dadkhahi, A. Garg, O. Firat, C. Yeh, A. M. Dai, B. Ghorbani "Order Matters in the Presence of Dataset Imbalance for Multilingual Learning," In *Advances in Neural Information Processing Systems 36*, 2023

M. B. Paulus, **D. Choi**, D. Tarlow, A. Krause, C. J. Maddison "Gradient Estimation with Stochastic Softmax Tricks," In *Advances in Neural Information Processing Systems 33*, 2020 (**Oral**).

R. T. Chen, **D. Choi**, L. Balles, D. Duvenaud, P. Hennig "Self-Tuning Stochastic Optimization with Curvature-Aware Gradient Filtering," In *Workshop on "I Can't Believe It's Not Better!", NeurIPS*, 2020.

D. Choi, C. J. Shallue, Z. Nado, J. Lee, C. J. Maddison, G. E. Dahl "On Empirical Comparisons of Optimizers for Deep Learning," arXiv preprint arXiv:1910.05446.

D. Choi, A. Passos, C. J. Shallue, G. E. Dahl "Faster Neural Network Training with Data Echoing," arXiv preprint arXiv:1907.05550.

N. Maheswaranathan, L. Metz, G. Tucker, **D. Choi**, J. Sohl-Dickstein "Guided evolutionary strategies: escaping the curse of dimensionality in random search," In *Proceedings of 36th International Conference on Machine Learning*, Long Beach, California, USA, 2019.

W. Grathwohl, **D. Choi**, Y. Wu, G. Roeder, D. Duvenaud "Backpropagation through the Void: Optimizing control variates for black-box gradient estimation.," In *Proceedings of 6th International Conference on Learning Representations*, Vancouver, British Columbia, Canada, 2018.

AWARDS & SCHOLARSHIPS

- Open Phil AI Fellowship 2020 – Present
- NSERC CGS D 2020 – 2023
- Top 10% Reviewer at NeurIPS 2020
- NSERC Undergraduate Student Research Award May 2017 – Aug 2017
- Dean's Honours List 2013 – 2018
  University of Toronto, Faculty of Applied Science and Engineering

INVITED TALKS

- Tools for Verifying Neural Models' Training Data, AI Safety Hub Edinburgh. August 2023
- Gradient Estimation with Stochastic Softmax Tricks, Differentiable Almost Everything: Differentiable Relaxations, Algorithms, Operators, and Simulators at ICML. July 2023

- Backpropagation through the Void: Optimizing control variates for black-box gradient estimation, Endless Summer School session at Vector Institute, Toronto, Canada. March 2018