

Education

- 2022–Present **PhD, Computer Science, *Stanford University***, Stanford, CA, GPA: 4.17
Thesis: Learning Representations Under Restricted Conditions.
Advised by Sebastian Thrun and Christopher Ré.
Graduated in 2.5 years (Department record).
- 2018–2018 **MEng, Computer Science, *Cornell University***, Ithaca, NY, GPA: 4.02
Research: Graph Neural Networks. Advised by Kilian Weinberger.
- 2014–2018 **BA, Computer Science, *Cornell University***, Ithaca, NY, GPA: 3.94
External Specialization: Mathematics.

Experience

- 2020–2022 **Research Engineer, *Google Brain***, Mountain View, CA
- Theoretical research in transfer learning, multi-task learning, and deep learning optimization.
 - Applied research in multi-modal (video, audio, text) modeling.
- 2019–2020 **Machine Learning Engineer, *YouTube***, Mountain View, Ca
- Machine learning algorithmic improvements to YouTube Algorithm(s).
 - 5-figure spot bonus from YouTube CTO for “outstanding, above-and-beyond contributions to YouTube”.
- 2018–2019 **Machine Learning Research Intern, *ASAPP***, Ithaca, NY
- Personally selected by Dr. Kilian Weinberger to be Employee no. 1 in ASAPP’s Ithaca Research Extension.
 - Research into NLP systems and intent disambiguation.

Projects

- 2019–2022 **imagineart.ai**
- I founded <https://imagineart.ai>, an eCommerce website that transforms any picture into a work of art using machine learning and gives our users the opportunity to buy this art as a canvas.
 - Invited to the final round of YC interviews as a solo founder (6% of all applicants).
- 2019–2022 **Machine Learning YouTube Content Creator**
- I created educational machine learning content about TensorFlow, hardware, and Natural Language Processing.
 - My channel has around 75 subscribers, over 10,000 views, and around 330 hours of watch time.

Talks

- April 2025 Oral presentation on Vector Quantization for generative models at ICLR 2025 in Singapore.
- Dec. 2021 Spotlight presentation on Task Affinity Groupings for multi-task learning at NeurIPS 2021 in Long Beach, CA.
- October 2021 I gave a talk to Level 5 (formerly Lyft’s self-driving division, now part of Toyota’s) regarding multi-task learning for autonomous vehicles.
- March 2021 I was invited by Waymo Research to give a talk on information transfer in multi-task learning systems.

Media

Personal Website, <https://cfifty.github.io>

- Oct. 2021 **Google AI Blog**, Deciding Which Tasks Should Train Together in Multi-Task Neural Networks.
<https://ai.googleblog.com/2021/10/deciding-which-tasks-should-train.html>

Publications: 4,610 citations (05/15/25)

- ICLR 2025 **Restructuring Vector-Quantization with the Rotation Trick**
Oral Award Christopher Fifty, Ronald G. Junkins, Dennis Duan, Aniketh Iyengar, Jerry W. Liu, Ehsan Amid, Sebastian Thrun, Christopher Ré.
- ICLR 2024 **Context-Aware Meta-Learning**
Christopher Fifty, Dennis Duan, Ronald Junkins, Ehsan Amid, Jure Leskovec, Christopher Ré, Sebastian Thrun.

- arXiv 2023 In-Context Learning for Few-Shot Molecular Property Prediction
Christopher Fifty, Jure Leskovec, Sebastian Thrun.
- NeurIPS 2023 Implicit Geometry and Interaction Embeddings Improve Few-Shot Molecular Property Prediction
Machine Learning for Structural Biology Workshop
Christopher Fifty, Joe Paggi, Ehsan Amid, Jure Leskovec, Ron Dror.
- TMLR 2022 Layerwise Bregman Representation Learning with Applications to Knowledge Distillation
Ehsan Amid, Rohan Anil, Christopher Fifty, Manfred K Marmuth.
- WWW 2022 Small towers make big differences
Yuyan Wang, Zhe Zhao, Bo Dai, Christopher Fifty, Dong Lin, Lichan Hong, Ed H Chi.
- arXiv 2022 N-Grammer: Augmenting Transformers with latent n-grams
Aurko Roy, Rohan Anil, ..., Christopher Fifty, ..., Yonghui Wu.
- NeurIPS 2021 Efficiently Identifying Task Groupings for Multi-Task Learning
Spotlight Christopher Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, Chelsea Finn.
- arXiv 2021 Co-training transformer with videos and images improves action recognition
Bowen Zhang, Jiahui Yu, Christopher Fifty, Wei Han, Andrew M Dai, Ruoming Pang, Fei Sha.
- ICML 2020 Step-Size Adaptation Using Exponentiated Gradient Updates
Spotlight *Beyond First Order Methods in Machine Learning Systems Workshop*
Ehsan Amid, Rohan Anil, Christopher Fifty, Manfred K Warmuth
- arXiv 2020 Measuring and Harnessing Transference in Multi-Task Learning
Christopher Fifty, Ehsan Amid, Zhe Zaho, Tianhe Yu, Rohan Anil, Chelsea Finn.
- ICML 2019 Simplifying Graph Convolutional Networks
Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, Kilian Weinberger

Community Contributions

- 2021-Present Peer reviewer for NeurIPS (**Outstanding Reviewer Award**), ICML, and ICLR conferences.
- 2021-Present Peer reviewer for *Transactions of Machine Learning Research* journal.
- 2024 Teaching assistant for CS 224n (Stanford), Natural Language Processing with Deep Learning.
- 2024 Teaching assistant for CS 299 (Stanford), Machine Learning.
- 2018-2022 Administrator of Cornell Computer Science Alumni Community.
- 2018-2022 Alumni Interviewer for Cornell University.
- 2018 Teaching assistant for CS 5740/4750 (Cornell), Natural Language Processing, taught by Claire Cardie.
- 2017 Teaching assistant for CS 4820 (Cornell), Introduction to Analysis of Algorithms, taught by Eva Tardos.

Open Source Software

- 29,500+** ☆ Google Research Open Source: <https://github.com/google-research>. Core Contributor
As a researcher at Google Brain, I contributed the source code for my research on Task Affinity Groupings in Multi-Task Learning.
- 2,700+** ☆ Lingvo: <https://github.com/tensorflow/lingvo>. Core Contributor
Lingvo is a framework for building neural networks in Tensorflow, particularly sequence models. My team at Google Brain is responsible for its internal and external development. This library powers Google Translate, the YouTube Algorithm, etc.
- 671** ☆ ActionML Recommender: <https://github.com/actionml/universal-recommender>. Core Contributor
ActionML's universal recommender is an open-source collaborative-filtering recommendation engine. My work was instrumental in porting the initial Scala implementation to Java.
- 101** ☆ The Rotation Trick: https://github.com/cifty/rotation_trick Primary Author
Open source implementation for Restructuring Vector Quantization with the Rotation Trick (ICLR 2025, Oral Award).
- 62** ☆ CAML: <https://github.com/cifty/CAML> Primary Author
Open source implementation for Context-Aware Meta-Learning (ICLR 2024).

26 ☆ NumpyGCN: <https://github.com/cifty/numpyGCN> Primary Author
NumpyGCN is a Numpy Implementation of Graph Convolutional Networks.

Selected Graduate Coursework

- Math 120 (Abstract Algebra)
- Math 143 (Differential Geometry)
- Math 144 (Topology and Geometry)
- Math 171 (Fundamental Concepts of Analysis)
- EE 276 (Information Theory)
- EE 364m (The Mathematics of Convexity)