

TIANYI ZHANG

tz58@stanford.edu \diamond github.com/Tiiiger \diamond tiiiger.github.io

EDUCATION

Stanford University, Computer Science Department 2020.9 – Present

Cornell University, College of Arts and Sciences 2016.8 – 2019.12

Major in Computer Science & College Scholar Program GPA: 3.99/4.3

Honors & Awards: 2019/2020 CRA Outstanding Undergraduate Researcher Award Runner-up; Lynne S. Abel Fund; Einhorn Discovery Grant; Cornell Undergraduate Research Fund;

RESEARCH (sorted by time)

SWALP: Stochastic Weight Averaging for Low Precision Training ICML 2019
advised by Prof. Christopher De Sa *Second Author*

We investigate the combination of stochastic weight averaging and low-precision training. This algorithm displays better convergence behaviors and performs well on common benchmarks.

Simplifying Graph Convolutional Networks ICML 2019
advised by Prof. Kilian Weinberger *Co-First Author*

Our work simplifies Graph Convolutional Networks to a linear model. The resulting model achieves competitive performance while being up to two orders of magnitude faster. Moreover, the simplified model leads to a better theoretical understanding of graph convolutional models.

QPyTorch: A Low-Precision Arithmetic Simulation Framework NeurIPS 2019 Workshop
advised by Prof. Christopher De Sa, *First Author*

We develop this open-source project to support the empirical research on low-precision training. QPyTorch can simulate various quantization strategies and integrates well with the cutting-edge research. Our software has been downloaded $9K+$ times through PyPI.

BERTScore: Evaluating Text Generation with BERT ICLR 2020
advised by Prof. Kilian Weinberger and Prof. Yoav Artzi *Co-First Author*

We leverage contextual embeddings to develop an automatic evaluation metric for text generation. The proposed metric correlates highly with human judgments and is more robust to adversarial inputs. Our software has been downloaded $37K+$ times through PyPI.

Identifying Mislabeled Data using the Area Under the Margin Ranking NeurIPS 2020
advised by Prof. Kilian Weinberger *Second Author*

We introduce a new method to discover mislabeled training samples and mitigate their impact on the training process of deep neural networks.

Revisiting Few-sample BERT Fine-tuning ICLR 2021
advised by Prof. Kilian Weinberger and Prof. Yoav Artzi *Co-First Author*

We identify and fix three suboptimal optimization practices in BERT Fine-tuning.

EXPERIENCE

Research Engineer, supervised by Kilian Weinberger, ASAPP Inc. 2019.03 – 2020.07

Reviewer, NeurIPS 2020, ICLR 2021

Student Volunteer, NeurIPS 2019 Conference 2019.12

Teaching Assistant, Introduction to Natural Language Processing 2018.08 – 2018.12