

Guandao Yang

gy46@cornell.edu | (607) 379-8872 | github.com/stevenygd | www.guandaoyang.com

EDUCATION **Cornell University** Sep 2018 - May 2023
Computer Science PhD

Cornell University Sep 2014 - May 2018
Bachelor of Arts, double major in CS and Math GPA: 3.933

RESEARCH **SWALP : Stochastic Weight Averaging for Low Precision Training**

Guandao Yang, Tianyi Zhang, Haonan Chen, Christopher De Sa (Aiming for SysML)
Proposed a novel combination between stochastic weight averaging (SWA) and low precision training. Evaluate the advantage of SWALP both theoretically and empirically. Theoretically we prove that SWALP has strong bounds in convex problems. Empirically we show that SWALP consistently outperforms low precision SGD in all experiment settings including deep learnings.

A Unified Framework for Single-View 3D Reconstruction with Limited Pose Supervision

Guandao Yang, Yin Cui, Serge Belongie, Bharath Hariharan **(ECCV 2018)**
Proposed a unifying model combining autoencoder with GAN to evaluate the effect of reducing pose supervision on single-image 3D reconstruction task. The model achieves competitive performance with state-of-the-arts models with substantially less amount of pose supervisions.

Learning to Evaluate Image Captioning [[link](#)]

Yin Cui, Guandao Yang, Andreas Veit, Xun Huang, Serge Belongie **(CVPR 2018)**
Developed a learning-based evaluation metric for image captioning task with state-of-the-art human correlations. Designed and conducted experiments to investigate the roles of data augmentation; the results inspired the decision of using pathological transformations as data augmentation methods. Achieved state-of-the-art caption level human correlations in Flickr8k.

Deep Fundamental Matrices Estimation [[link](#)] **(GMDL 2018, workshop)**

Guandao Yang, Omid Poursaeed, Hanqing Jiang, Qiuren Fang, Bharath Hariharan, Serge Belongie
Estimating the F-Matrices from stereo images using end-to-end trainable deep learning models. Designed and implemented a novel Siamese network architecture trained with regression loss without losing mathematical properties of F-matrices or using key-points correspondences.

Fast Reading Comprehension with ConvNets [[arxiv](#)]

Felix Wu, Ni Lao, John Blitzer, Guandao Yang, Kilian Weinberger
Speeding up reading comprehension models using dilated CNN. Conducted experiments on open-domain question answering setting. Reproduced Document Reader's open-domain experiments. Achieved comparable or better results using dilated CNN in several datasets.

EXPERIENCE **Software Engineer Intern, Google** May 2017 - Aug 2017

Performers Extraction: designed a generalizable algorithm to extract performers from public events data; designed human interactive task to evaluate performance; improved coverage and recall by 1000x from baseline; achieve > 40% recall and > 89% precision.
Data Analysis/Visualization: designed an extensible architecture to build dashboard to visualize and analyze large volume of data; implemented it for public events dataset.

Research Assistant, Cornell SonicMEMS Lab Sep 2017 - Dec 2017

Used computer vision to extract nail-to-nail fingerprints from a video using inverse slit-cam scanning. Entered the IARPA N2N challenge final round as the only college undergraduate team; only 8 teams are qualified in the final round. Working to reconstruct 3D fingerprint models from video and combine signals from ultrasonic sensors to detect fake fingers.

Teaching Assistant, Cornell CS Dept. Sep 2015 - May 2018

Machine Learning for Intelligence System (CS 4780) Implemented & tested homeworks in Julia language; managed grade in-class kaggle competition for image classification.
Introductory and Intermediate iOS (CS 2048, CS 2049) Designed and held labs for Swift, iOS debugging & testing; assisted lecturer with live coding session of 50+ students.

Functional Programming Courses (CS 3110) Designed the first consensus protocol assignment; organized TA training & grading session; held course recitations twice a week.

Cofounder, Tech Lead, Hyphen Connect [[App Store](#)]

Apr 2015 - May 2016

Led a team of 5 to build the iOS app; managed weekly iteration and code review; designed the iOS architecture based on MVVM; analyzed data from Mixpanel to improve user experience; implemented server side logic using NodeJS and Promise; implemented iOS business logics synchronizing data between Push Notification, local cache in SQLite, and remote database.

Full-stack Web Developer, Cornell Linguistics Dept.

Sep 2014 - May 2016

Solo projects include [SpeechTerrors](#) used by hundreds of students from different courses; developed and maintained experiment website [[link](#)] using PHP for Cornell Phonetics Lab.

Software Engineer Intern, dxTechnology, Guangzhou.

Jul 2014 - Aug 2014

Built UI pages with JS framework Famo.us; reduced layout overhead to improve animation performance from 20 FPS to > 40 FPS; implemented more than one-third of front-end pages.

PROJECTS

Learning the Game of Renju using Neural Networks [[poster](#)]

Initiated the project to combine tree search and heuristics learned by neural networks to play the game of Renju. Lead a group of six students to architect, build, test, and optimize the system. Our current AI is competitive with top-10 AIs in GomokuCup.

Deep Learning Library in Julia [[codes](#)]

Implementing basic structure of NN.jl, a lightweight deep learning framework in Julia language; currently developing asynchronous lock-free updates following HOGWILD!.

Textmatters [[slide](#)]

Initiated the project to evaluate how in-image text will affect state-of-the-arts captioning system by conducting ablation studies using annotations from COCO-Text dataset.

Investigating GAN Domain Adaptation

Reproduce several GAN domain adaptation paper; explore the possibility to combine generative approach with discriminative approaches; implemented WGAN versions.

Dialogue Disentangling

Using unsupervised learning method to disentangle dialogues; identified key issue from the dialogue disentangling dataset and designed a new AMT pipeline to collect better labels.

Search Engines in OCaml

Built two general purpose search engines using vector space model and machine learning algorithms; implemented TF-IDF, Pagerank, RT Rank algorithm and MapReduce.

Compiler in OCaml

Built a compiler for a Java-like language with additional features such as pattern matching; implemented IR translation, optimizations, calling convention and OOP feature.

PATENT

Sono-Optical Nail-to-nail Fingerprint Extraction (Under application for provisional patent)

Use fingerprint scanned from an ultrasonic imaging devices to cross validate the fingerprint reconstructed from the video using inverse slit-cam. This combination allows efficient extraction of both the optical fingerprint features and the textural informations, adding the ability to detect fake fingerprint made of wrong materials, thus making nail-to-nail fingerprint more secure.

SKILLS

Mobile Proficient in Objective-C; familiar with Swift, Auto-layout, and Interface Builder.

Web Proficient in PHP, JS, SQL; familiar with Nodejs, AngularJS, Codelgniter, Laravel.

ML Proficient in Tensorflow and PyTorch.

Others Proficient in Java, Python, OCaml.