Yutian Li

Email: hi@yutian.li Phone: +1 (650) 521-1357 Website: https://yutian.li

Work Experience

Aug 2018 - Current — Core Developer at Hudson River Trading
Working on large-scale cluster and data systems. Improving research job efficiency and user experience.

Jan - Jun 2018 — Teaching Assistant at Stanford University
Teaching assistant for CS 246: Mining Massive Data Sets and CS 240: Advanced Topics in Operating Systems.

Aug 2017 - Jan 2018 — Intern at Tesla
Autopilot
Designed and implemented fleet data engine for hard negative case mining. Neural networks training.

Jan 2016 - Jan 2017 — Engineer at Megvii, Beijing
Infrastructure Development
Cluster resource scheduling and containerized runtime. Large-scale neural networks training system. Support for cluster parallel training. Familiar with advanced C++ features and system design.

Jan - Feb 2015 — Assistant Developer Intern at Jane Street Asia Limited
Worked on projects both in horizontal and vertical scope in OCaml, a functional programming language. Landed features into production system.

Feb 2014 - Jan 2015 — Full-Time Intern at Microsoft Research Asia, Beijing
Systems Research Group
Distributed CUDA system by the name of Minerva. Designed for rapid training of deep neural networks. Devised a technique to speed up convolution by up to 50%. Contributed a major part of the code, consisting of interface design, memory and thread management, and scheduling.

Education

Sept 2015 - Jun 2018 — Master of Science, Computer Science
Stanford University
Systems track.

Aug 2010 - Jul 2015 — Bachelor of Science in Engineering, Computer Science and Technology
Tsinghua University
Major GPA 92/100, ranked 4th out of 123. Graduated with distinction.

Aug - Dec 2013 — Undergraduate Exchange Program, Computer Science
University of Texas at Austin
GPA: 4.0 with University Honors.

Projects

MinPy — *Pure NumPy with third party operators and gradients.*
Integrated NumPy exprience with MXNet. Offered Python interface native to NumPy with customizable operators and automatic gradient calculation.

VisualGenome — *Connecting structured image concepts to language.*
Worked under supervision of Dr. Fei-Fei Li, on the design and implementation of an automated data cleaning and storage systems. Aimed to be a dense knowledge base of images, descriptions, and concepts.

MXNet — *An efficient, flexible distributed framework for deep learning.*
Built upon experience from Minerva, CXXNET, and many more robust frameworks. Worked under DMLC to provide flexibility and portability for legacy and new code. Designed and implemented the engine to support efficient task scheduling. Published in Neural Information Processing Systems, Workshop on Machine Learning Systems, 2015.