Tianqi Chen

CONTACT INFORMATION

Paul G. Allen School of
Computer Science & Engineering
University of Washington

E-mail: tqchen@cs.washington.edu
Phone: (+1) 206-669-7495
http://homes.cs.washington.edu/~tqchen/

RESEARCH

Machine Learning, Systems.

EDUCATION

PhD. Paul G. Allen School of Computer Science & Engineering, University of Washington — 2013-present

M.S. Department of Computer Science, Shanghai Jiao Tong University — 2010-2013

B.S. Department of Computer Science, Shanghai Jiao Tong University — 2006-2010

AWARDS

Google Ph.D. Fellowship
Champion of KDD Cup 2012.
3rd place of KDD Cup 2011.

SELECTED PUBLICATIONS

Learning to Optimize Tensor Programs.
T. Chen, L. Zheng, E. Yan, Z. Jiang, T. Moreau, L. Ceze, C. Guestrin, A. Krishnamurthy.
NeurIPS 2018

TVM: An Automated End-to-End Optimizing Compiler for Deep Learning.
T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, M. Cowan, H. Shen, L. Wang, Y. Hu, L. Ceze, C. Guestrin, A. Krishnamurthy.
OSDI 2018

XGBoost: A Scalable Tree Boosting System.
T. Chen, C. Guestrin.
KDD 2016

Net2Net: Accelerating Learning via Knowledge Transfer.
T. Chen, I. Goodfellow, J. Shlens.
ICLR 2016

A Complete Recipe for Stochastic Gradient MCMC
Y. Ma, T. Chen, E. Fox.
NeurIPS 2015

Efficient Second-Order Gradient Boosting for Conditional Random Fields.
T. Chen, S. Singh, B. Taskar, C. Guestrin.
AISTATS 2015

Stochastic Gradient Hamiltonian Monte Carlo.
T. Chen, E. Fox, C. Guestrin.
ICML 2014

OPEN SOURCE LEARNING SYSTEMS

TVM tvm.ai
Automated, learning based, end-to-end deep learning compiler.

MXNet mxnet.io
Scalable deep learning framework, adoped by Amazon.

XGBoost xgboost.ai
One of the de-facto tools used daily by data scientists, adopted in production pipelines of the major companies such as Uber, Airbnb, Amazon and Google Cloud.

SERVICES

Organizer of the LearningSys workshop at NeurIPS 2015.
Reviewer of ICML, IJCAI, NeurIPS, ICLR, JMLR.

SELECTED TALKS

TVM: An Automated End-to-End Optimizing Compiler for Deep Learning.
2018, OSDI Conference.

TVM: End-to-End Compilation Stack for Deep Learning.
2017, SOSP AISys Workshop.

XGBoost: A Scalable Tree Boosting System.
2016, KDD Conference.

Net2Net: Accelerating Learning via Knowledge Transfer.
2016, ICLR Conference.

Stochastic Gradient Hamiltonian Monte Carlo.
2014, ICML Conference.

ADVISING

I am advising multiple graduate students in the UW SAMPL lab, which I helped form.

- Eddie Yan. Scaling up learning-based automated program optimizations.
- Meghan Cowan. Optimizing deep learning models with low bits models.
- Ziheng Jiang. Machine learning on the edge, on-device training.
- Pratyush Patel. Deep learning systems on bare metal hardware.

TEACHING

I created and lectured two new courses:

- Systems for Machine Learning, University of Washington. Spring 2017, 2018
- Machine Learning and Data Mining in Practice, SJTU. Spring 2012

I also taught the following courses:

- Lecturer, Operation System Design and Implementation, SJTU. Fall 2010
- Teaching Assistant, Database Design and Implementation, SJTU. Spring 2010
- Teaching Assistant, Operation System Design and Implementation, SJTU. Fall 2009
- Teaching Assistant, Compiler Design and Implementation, SJTU. Spring 2009

INDUSTRY EXPERIENCE

Google Brain. Summer 2015
Research intern with Ian Goodfellow.
I developed Net2Net, a new method to accelerating learning via knowledge transfer. This work led to an ICLR paper and an US patent.

Turi Inc. Summer 2014
Software engineering intern with Jay Gu.
I developed the boosted tree and neural net toolkit in Turi Create.

Huawei Noah's Ark Lab, Hong Kong. Summer 2012
Research Intern with with Hang Li and Qiang Yang.
I developed a new functional gradient boosting algorithm for recommender systems.

Microsoft Research Asia, Beijing. Summer 2009
Research Intern, with Jun Yan.
I studied transfer learning for behavioral targeting.