Huan Zhang

University of California, Los Angeles
http://www.huan-zhang.com

☎ 530-601-3829
✉ huanzhang@ucla.edu
404 Westwood Plaza
Los Angeles, CA 90095-1596
http://www.huan-zhang.com

Education

Ph.D. in Computer Science (in progress), UCLA (2018-), UC Davis (2014-2018).
• *Advisor:* Prof. Cho-Jui Hsieh (2015-present) and Prof. Venkatesh Akella (2014-2018)
• *Area:* deep learning, large-scale machine learning and optimization. Some recent works include exploring the robustness properties and adversarial examples in deep neural networks (especially on formal robustness verification), accelerating gradient boosted decision tree (GBDT) training on GPUs, and designing large-scale asynchronous parallel stochastic gradient descent and coordinate descent optimizers in multi-core and distributed settings, especially for deep learning training.

M.S. in Computer Engineering, UC Davis (2012-2014).
• *Advisor:* Prof. Venkatesh Akella (akella@ucdavis.edu)
• *Area:* computer architecture and parallel computing, especially on improving scalability of parallel workloads on heterogeneous multi-core processors and distributed systems.

Bachelor of Engineering, Zhejiang University (2008 - 2012), China.
• *Major:* Information Engineering

Work Experience

Internship in Microsoft Research AI, Redmond, WA, Jun 2018 - Sep 2018.
• *Mentor: Pengchuan Zhang (penzhan@microsoft.com), Po-Sen Huang and Lin Xiao*
I worked in the deep learning group and explored several ideas on improving the training stability of generative adversarial networks (GAN), as well as proposed a new algorithm for efficiently bounding the Jacobian matrix of a neural network and applied it to neural network robustness verification.

Internship in Amazon A9.com Search Lab, Palo Alto, CA, Nov 2017 - Mar 2018.
• *Mentor: Inderjit Dhillon (isd@a9.com)*
I worked on the design and implementation of a deep-learning based product search suggestion system for amazon.com production website. Given an incomplete search-bar query, our LSTM-based model provides candidates to complete the query in real-time.

Internship in IBM T.J. Watson Research Center, Yorktown Heights, NY, Jun 2017 - Nov 2017; Apr 2018 - Jun 2018, • *Mentor: Jinfeng Yi (jinfengy@us.ibm.com), Pin-Yu Chen (Pin-Yu.Chen@ibm.com).*
I worked on exploring the robustness issues of deep neural networks by crafting adversarial examples in a black-box setting using zeroth order optimization (ZOO) based attack, achieving very high success rates. I also investigated the robustness of image captioning system with recurrent neurons and attention mechanism and proposed the Show-and-Fool attack.

Selected Publications (* indicates equal contribution)

▪ *Neural Network Robustness Verification*

Efficient Neural Network Robustness Certification with General Activation Functions, Advances in Neural Information Processing Systems (NIPS), 2018, Huan Zhang*, Tsui-Wei Weng*, Pin-Yu Chen, Cho-Jui Hsieh, Luca Daniel. (* Equal contribution).

RecurJac: An Efficient Recursive Algorithm for Bounding Jacobian Matrix of Neural Networks and Its Applications, AAAI Conference on Artificial Intelligence (AAAI), 2019, Huan Zhang, Pengchuan Zhang, Cho-Jui Hsieh.

Towards Fast Computation of Certified Robustness for ReLU Networks, International Conference on Machine Learning (ICML), 2018, Tsui-Wei Weng*, Huan Zhang*, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Duane Boning, Inderjit S. Dhillon, Luca Daniel. (* Equal contribution).

Evaluating the Robustness of Neural Networks: An Extreme Value Theory Approach,

International Conference on Learning Representations (ICLR), 2018, Tsui-Wei Weng*, Huan Zhang*, Pin-Yu Chen, Jinfeng Yi, Dong Su, Yupeng Gao, Cho-Jui Hsieh, Luca Daniel (* Equal contribution).

- *Adversarial Examples in Deep Learning (Attacks and Defenses)*

The Limitations of Adversarial Training and the Blind-Spot Attack, International Conference on Learning Representations (ICLR), 2019, Huan Zhang*, Hongge Chen*, Zhao Song, Duane Boning, inderjit dhillon, Cho-Jui Hsieh (* Equal contribution).

Attacking Visual Language Grounding with Adversarial Examples: A Case Study on Neural Image Captioning, 56th Annual Meeting of the Association for Computational Linguistics (ACL), 2018, Hongge Chen*, Huan Zhang*, Pin-Yu Chen, Jinfeng Yi, Cho-Jui Hsieh (* Equal contribution).

Is Robustness the Cost of Accuracy? Lessons Learned from 18 Deep Image Classifiers, European Conference on Computer Vision (ECCV), 2018, Dong Su*, Huan Zhang*, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, Yupeng Gao. (* Equal contribution).

ZOO: Zeroth Order Optimization based Black-box Attacks to Deep Neural Networks without Training Substitute Models, 10th ACM Workshop on Artificial Intelligence and Security, 2017, Pin-Yu Chen*, Huan Zhang*, Yash Sharma, Jinfeng Yi and Cho-Jui Hsieh (* Equal contribution).

Towards Robust Neural Networks via Random Self-ensemble, European Conference on Computer Vision (ECCV), 2018, Xuanqing Liu, Minhao Cheng, Huan Zhang, Cho-Jui Hsieh.

EAD: Elastic-Net Attacks to Deep Neural Networks via Adversarial Examples, In AAAI Conference on Artificial Intelligence (AAAI), 2018, Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi and Cho-Jui Hsieh.

- *Large-scale Optimization, Machine Learning Acceleration and Computer Systems*

GPU-acceleration for Large-scale Tree Boosting, SysML Conference, 2018, Huan Zhang, Si Si and Cho-Jui Hsieh.

A Comprehensive Linear Speedup Analysis for Asynchronous Stochastic Parallel Optimization from Zeroth-Order to First-Order, Advances in Neural Information Processing Systems (NIPS), 2016, Xiangru Lian, Huan Zhang, Cho-Jui Hsieh, Yijun Huang and Ji Liu.

Can Decentralized Algorithms Outperform Centralized Algorithms? A Case Study for Decentralized Parallel Stochastic Gradient Descent, Advances in Neural Information Processing Systems (NIPS), 2017, Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang and Ji Liu.

HogWild++: A New Mechanism for Decentralized Asynchronous Stochastic Gradient Descent, IEEE International Conference on Data Mining (ICDM), 2016, Huan Zhang, Cho-Jui Hsieh and Venkatesh Akella.

Gradient Boosted Decision Trees for High Dimensional Sparse Output, 34th International Conference on Machine Learning (ICML), 2017, Si Si, Huan Zhang, Sathiya Keerthi, Dhruv Mahajan, Inderjit Dhillon and Cho-Jui Hsieh.

Fixing the Convergence Problems in Parallel Asynchronous Dual Coordinate Descent, IEEE International Conference on Data Mining (ICDM), 2016, Huan Zhang and Cho-Jui Hsieh.

Sublinear Time Orthogonal Tensor Decomposition, Advances in Neural Information Processing Systems (NIPS), 2016, Zhao Song, David P. Woodruff and Huan Zhang.

Open Source Projects

LightGBM on GPU, `https://github.com/huanzhang12/lightgbm-gpu`.
LightGBM is a popular tree boosting package with high efficiency on large-scale datasets. I accelerated its decision tree construction process on GPUs using a novel histogram based algorithm which efficiently exploits the GPU's capability to approximately find the best split. Experiment results show that the GPU algorithm is up to 7 to 8 times faster than LightGBM and 20 to 30 times faster than XGBoost on 28 CPU cores. My code (about 5,000 lines C++) reaches production quality and has been merged into the LightGBM official repository.