

CHIYUAN ZHANG

Research Scientist, Google Brain

✉ chiyuan@google.com
🌐 <http://pluskid.org>
🌐 <https://github.com/pluskid>

2012-2017 PhD in Electrical Engineering and Computer Science, MIT
Computer Science and Artificial Intelligence Laboratory (CSAIL)
Center for Brain, Minds & Machines (CBMM)

🎓 Tomaso Poggio

2014-2014 Summer Exchange Program in Japan
Graduate School of Informatics, Kyoto University
Computer Science Department, Tokyo Institute of Technology

🎓 Marco Cuturi
🎓 Masashi Sugiyama

2009-2012 MEng in Computer Science at Zhejiang University

🎓 Xiaofei He & Deng Cai

2005-2009 BEng in Chu Kochen Honors College & College of Computer Science at Zhejiang University

Selected Publications

- preprint **Chiyuan Zhang***, Vitaly Feldman*: *What Neural Networks Memorize and Why: Discovering the Long Tail via Influence Estimation*. preprint, 2020. *equal contribution.
- arXiv **Chiyuan Zhang***, Ziheng Jiang*, Kunal Talwar, Michael C. Mozer: *Characterizing Structural Regularities of Labeled Data in Overparameterized Models*. arXiv:2002.03206, 2020. *equal contribution.
- ICLR20 **Chiyuan Zhang**, Samy Bengio, Moritz Hardt, Michael C. Mozer, Yoram Singer: *Identity Crisis: Memorization and Generalization under Extreme Overparameterization*. The International Conference on Learning Representations (ICLR), 2020.
- arXiv **Chiyuan Zhang**, Samy Bengio, Yoram Singer: *Are All Layers Created Equal?*. arXiv:1902.01996, 2019.
- NeurIPS19 **Chiyuan Zhang***, Maithra Raghu*, Jon Kleinberg, Samy Bengio: *Transfusion: Understanding Transfer Learning with Applications to Medical Imaging*. Advances in Neural Information Processing Systems 33 (NeurIPS), 2019. *equal contribution.
- ICML18 Neil C. Rabinowitz, Frank Perbet, H. Francis Song, **Chiyuan Zhang**, S.M. Ali Eslami, Matthew Botvinick: *Machine Theory of Mind*. International Conference on Machine Learning (ICML), 2018.
- ICLR17 **Chiyuan Zhang**, Samy Bengio, Moritz Hardt, Benjamin Recht, Oriol Vinyals: *Understanding deep learning requires rethinking generalization*. 5th International Conference on Learning Representations (ICLR), **Best Paper Award**, 2017.
- NeurIPS15 **Chiyuan Zhang***, Charlie Frogner*, Hossein Mobahi, Mauricio Araya-Polo, Tomaso Poggio: *Learning with a Wasserstein Loss*. Advances in Neural Information Processing Systems 28 (NeurIPS), 2015. *equal contribution.
- LearningSys Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, **Chiyuan Zhang**, Zheng Zhang: *MXNet: A Distributed Deep Learning Framework for Efficiency and Flexibility*. NIPS Workshop on LearningSys, 2015.
- INTERSPEECH14 Stephen Voinea, **Chiyuan Zhang**, Georgios Evangelopoulos, Lorenzo Rosasco, Tomaso Poggio: *Word-level Invariant Representations from Acoustic Waveforms*. INTERSPEECH (Best Student Paper Award), 2014.
- JMLR13 Binbin Lin, Xiaofei He, **Chiyuan Zhang**, Ming Ji: *Parallel Vector Field Embedding*. Journal of Machine Learning Research (JMLR), 2013.

- NeurIPS12 Binbin Lin, Sen Yang, **Chiyuan Zhang**, Jieping Ye, Xiaofei He: *Multi-task Vector Field Learning*. Advances in Neural Information Processing Systems 25 (**NeurIPS**), 2012.
- NeurIPS11 Binbin Lin, **Chiyuan Zhang**, Xiaofei He: *Semi-supervised Regression via Parallel Field Regularization*. Advances in Neural Information Processing Systems 24 (**NeurIPS**), 2011.
- KDD10 Deng Cai, **Chiyuan Zhang**, Xiaofei He: *Unsupervised feature selection for multi-cluster data*. International Conference on Knowledge Discovery and Data Mining (**KDD**), 2010.

Projects & Experiences

- 2017 Summer Intern at DeepMind: Investigating and modeling the theory of mind (ToM) for reinforcement learning agents (with Neil Rabinowitz); Inspecting the notion of generalization and overfitting in deep reinforcement learning (with Oriol Vinyals).
- 2016 Summer Intern at Google Brain: Understanding the effect of regularization on large neural networks (Host: Moritz Hardt & Samy Bengio).
- 2015 **MXNet** and **MXNet.jl**: lightweight, portable, flexible distributed deep learning library with multiple language bindings; joint efforts by the **DMLC team**.
- 2014 **Mocha.jl**: efficient and extensible multi-backends (GPU & CPU) deep learning framework for Julia. The **3rd most star-ed** Julia package on github.
- 2014 Summer Visiting at Kyoto University and Tokyo Institute of Technology: Optimal Transport (Advisor: Marco Cuturi) & Density Ratio Estimation (Advisor: Masashi Sugiyama) and applications to Domain Adaptation Learning problems.
- 2013 Internship at **Shell Intl. E&P Inc.**: automatic geological feature detection from pre-stack seismic traces with Machine Learning models (Host: Mauricio Araya Polo).
- 2012 **SHOGUN** C++ Machine Learning Library: I surveyed and implemented various well-known multiclass learning algorithms for SHOGUN (Mentor: Cheng Soon Ong, Sören Sonnenburg).
- 2008 **YASnippet**: I created the most popular open source snippet automation plugin for GNU Emacs (now maintained by capitomorte). The **5th most star-ed** Emacs plugin on github.

Skills

- Programming C, C++, CUDA, Julia, Python, Ruby, Matlab, Javascript, with practical experiences.
- Language Chinese Mandarin (native), English (fluent), Japanese (intermediate proficiency, JLPT N2).

Professional Activities

- NeurIPS Reviewer (14-20) of *Advances in Neural Information Processing Systems* (NeurIPS).
- ICML Reviewer (16-20) of *International Conference on Machine Learning* (ICML).
- ICLR Reviewer (18-20) of *International Conference on Learning Representations* (ICLR).
- JMLR Reviewer (15,18,19) and webmaster (13-18) of *Journal of Machine Learning Research* (JMLR).