

## Research interests

Machine learning, optimization, deep learning theory.

## Education

2016–2022 **Ph.D. in Computer Science**, *University of Illinois Urbana-Champaign*.

Advisor: Matus Telgarsky.

2012–2016 **B.Eng. in Computer Science and Technology (the ACM class)**, *Shanghai Jiao Tong University*.

## Work experience

2022- **Research Scientist**, *Google Research*.

Summer 2021 **Research Intern**, *Google Research*.

Summer 2019 **Research intern**, *Microsoft Research New York City*.

Summer 2017 **Software engineering intern**, *Google Inc.*

## Publications

All peer-reviewed work

- [1] Ziwei Ji, Kwangjun Ahn, Pranjal Awasthi, Satyen Kale, and Stefani Karp. Agnostic learnability of halfspaces via logistic loss. In *ICML*, 2022.
- [2] Yuzheng Hu, Ziwei Ji, and Matus Telgarsky. Actor-critic is implicitly biased towards high entropy optimal policies. In *ICLR*, 2022.
- [3] Ziwei Ji, Justin D Li, and Matus Telgarsky. Early-stopped neural networks are consistent. In *NeurIPS*, 2021.
- [4] Ziwei Ji, Nathan Srebro, and Matus Telgarsky. Fast margin maximization via dual acceleration. In *ICML*, 2021.
- [5] Daniel Hsu, Ziwei Ji, Matus Telgarsky, and Lan Wang. Generalization bounds via distillation. In *ICLR*, 2021.
- [6] Ziwei Ji and Matus Telgarsky. Characterizing the implicit bias via a primal-dual analysis. In *ALT*, 2021.
- [7] Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In *NeurIPS*, 2020.
- [8] Ziwei Ji, Miroslav Dudík, Robert E. Schapire, and Matus Telgarsky. Gradient descent follows the regularization path for general losses. In *COLT*, 2020.
- [9] Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. In *ICLR*, 2020.
- [10] Ziwei Ji, Matus Telgarsky, and Ruicheng Xian. Neural tangent kernels, transportation mappings, and universal approximation. In *ICLR*, 2020.
- [11] Ziwei Ji and Matus Telgarsky. Risk and parameter convergence of logistic regression. In *COLT*, 2019.
- [12] Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *ICLR*, 2019.
- [13] Ziwei Ji, Ruta Mehta, and Matus Telgarsky. Social welfare and profit maximization from revealed preferences. In *WINE*, 2018.

## Preprints

- [1] Kwangjun Ahn, Prateek Jain, Ziwei Ji, Satyen Kale, Praneeth Netrapalli, and Gil I Shamir. Reproducibility in optimization: theoretical framework and limits. *arXiv preprint arXiv:2202.04598*, 2022.

## Service

◦ Reviewer for *NeurIPS*, *ICML*, *ICLR*, *COLT*, *EC*, *ITCS*, *IEEE Transactions on Information Theory*.

## Teaching

Graduate teaching assistant, *UIUC*

- CS 598: Deep learning theory.
- CS 446: Machine learning.

Fall 2020, Fall 2021  
Spring 2019