

Blake Woodworth

blakewoodworth@gmail.com
home.ttic.edu/~blake/

Education

INRIA - SIERRA Team

Postdoctoral Researcher
Advisor: Francis Bach

2021-

Toyota Technological Institute at Chicago

Ph.D. in Computer Science
M.S. in Computer Science
Advisor: Nathan Srebro

2015-2021

2015-2017

Yale University

B.S. in Computer Science with Distinction, Summa Cum Laude
Senior Thesis Advisor: Daniel Spielman

2011-2015

Publications

Two Losses Are Better Than One: Faster Optimization Using a Cheaper Proxy

Blake Woodworth, Konstantin Mishchenko, Francis Bach

Submitted, 2023

Asynchronous SGD Beats Minibatch SGD Under Arbitrary Delays

Konstantin Mishchenko, Francis Bach, Mathieu Even, *Blake Woodworth*

NeurIPS, 2022

Non-Convex Optimization with Certificates and Fast Rates Through Kernel Sums of Squares

Blake Woodworth, Francis Bach, Alessandro Rudi

COLT, 2022

Lower Bounds for Non-Convex Stochastic Optimization

Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Nathan Srebro, and *Blake Woodworth*.

Mathematical Programming 2022.

The Minimax Complexity of Distributed Optimization

Blake Woodworth

PhD Thesis, 2021.

A Stochastic Newton Algorithm for Distributed Convex Optimization

Brian Bullins, Kumar Kshitij Patel, Ohad Shamir, Nathan Srebro, *Blake Woodworth*

NeurIPS, 2021.

An Even More Optimal Stochastic Optimization Algorithm: Minibatching and Interpolation Learning

Blake Woodworth, and Nathan Srebro.

NeurIPS 2021.

On the Implicit Bias of Initialization Shape: Beyond Infinitesimal Mirror Descent

Shahar Azulay, Edward Moroshko, Mor Shpigel Nacson, *Blake Woodworth*, Nathan Srebro, Amir Globerson, and Daniel Soudry.

ICML 2021.

The Min-Max Complexity of Distributed Stochastic Convex Optimization with Intermittent Communication

Blake Woodworth, Brian Bullins, Ohad Shamir, and Nathan Srebro.
COLT 2021. Best Paper Award.

Mirrorless Mirror Descent: A More Natural Discretization of Riemannian Gradient Flow

Suriya Gunasekar, Blake Woodworth, and Nathan Srebro.
AISTATS 2021.

Implicit Bias in Deep Linear Classification: Initialization Scale vs Training Accuracy

Edward Moroshko, Suriya Gunasekar, Blake Woodworth, Jason D. Lee, Nathan Srebro, and Daniel Soudry.
NeurIPS 2020.

Minibatch vs Local SGD for Heterogeneous Distributed Learning

Blake Woodworth, Kumar Kshitij Patel, and Nathan Srebro.
NeurIPS 2020.

Is Local SGD Better than Minibatch SGD?

Blake Woodworth, Kumar Kshitij Patel, Sebastian U. Stich, Zhen Dai, Brian Bullins, H. Brendan McMahan, Ohad Shamir, and Nathan Srebro.
ICML 2020.

Kernel and Deep Regimes in Overparametrized Models

Blake Woodworth, Suriya Gunasekar, Jason D. Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro.
COLT 2020.

The Gradient Complexity of Linear Regression

Mark Braverman, Elad Hazan, Max Simchowitz, and Blake Woodworth.
COLT 2020.

Guaranteed Validity for Empirical Approaches to Adaptive Data Analysis

Ryan Rogers, Aaron Roth, Adam Smith, Nathan Srebro, Om Thakkar, and Blake Woodworth.
AISTATS 2020.

Open Problem: The Oracle Complexity of Convex Optimization with Limited Memory

Blake Woodworth and Nathan Srebro.
COLT 2019.

The Complexity of Making the Gradient Small in Stochastic Convex Optimization

Dylan J. Foster, Ayush Sekhari, Ohad Shamir, Nathan Srebro, Karthik Sridharan, and Blake Woodworth.
COLT 2019. Best Student Paper Award.

Graph Oracle Models, Lower Bounds, and Gaps for Parallel Stochastic Optimization

Blake Woodworth, Jialei Wang, Brendan McMahan, and Nathan Srebro.
NeurIPS 2018.

Training Well-Generalizing Classifiers for Fairness Metrics and Other Data-Dependent Constraints

Andrew Cotter, Maya Gupta, Heinrich Jiang, Nathan Srebro, Karthik Sridharan, Serena Wang, Blake Woodworth, and Seungil You.
FAT/ML 2018, ICML 2019.

The Everlasting Database: Statistical Validity at a Fair Price

Blake Woodworth, Vitaly Feldman, Saharon Rosset, and Nathan Srebro.
NeurIPS 2018.

Lower Bound for Randomized First Order Convex Optimization

Blake Woodworth and Nathan Srebro.

Technical Report 2017.

Implicit Regularization in Matrix Factorization

Suriya Gunasekar, Blake Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nathan Srebro.

NeurIPS 2017.

Learning Non-Discriminatory Predictors

Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro.

COLT 2017.

Tight Complexity Bounds for Optimizing Composite Objectives

Blake Woodworth and Nathan Srebro.

NeurIPS 2016.

Fellowships and Awards

COLT Best Paper Award 2021.

COLT Best Student Paper Award, 2019.

Google Research PhD Fellowship, 2019.

NSF Graduate Research Fellowship, 2017.

Best Poster Award, TTIC Student Workshop, 2016.

Phi Beta Kappa Society Junior Year Inductee, 2013.

Internships

Google Research, Princeton, NJ, Summer 2019.

Worked with Elad Hazan and Naman Agarwal on online learning and optimization algorithms.

Microsoft Research, Cambridge, UK, Summer 2017.

Collaborated with Ryota Tomioka and Alex Gaunt to implement graph neural networks in Tensorflow, optimize them for performance, and apply them to prediction tasks on protein data.

Service

Conference Reviewing *NeurIPS, COLT, ICML, ICLR, AISTATS, FAT*

Journal Reviewing *JMLR, SIMODS, Mathematics of Optimization Research*

Co-Organizer TTIC Student Workshop 2019

Co-Organizer TTIC Admitted Student Visit Days 2016-2019

Teaching

Teaching Assistant, Statistical and Computational Learning Theory, Professor Nathan Srebro
TTIC, Fall 2018

Teaching Assistant, Convex Optimization, Professor Nathan Srebro
TTIC, Winter 2018

Teaching Assistant, Mathematical Toolkit, Professor Madhur Tulsiani
TTIC, Fall 2016

Peer Tutor, Data Structures and Programming Techniques, Professor Stanley Eisenstat
Yale University, Spring 2014, 2015

Peer Tutor, Systems Programming & Computer Organization, Professor Stanley Eisenstat
Yale University, Fall 2014

Peer Tutor, Introduction to Computer Science, Professor Dana Angluin
Yale University, Spring 2013