

Yang Yuan

FIT 4-6007
Tsinghua University
Haidian, Beijing, 100084

Phone: +86-13581688475
E-mail: yuanyang@tsinghua.edu.cn
<http://people.iis.tsinghua.edu.cn/~yuanyang/>

Research Interests AI+healthcare, AI theory and applied category theory.

Work Experience Associate Professor at TSINGHUA UNIVERSITY July, 2025 – Now
Institute for Interdisciplinary Information Sciences (IIIS)
College AI

Assistant Professor at TSINGHUA UNIVERSITY July, 2019 – June, 2025
Institute for Interdisciplinary Information Sciences (IIIS)
College AI (Since April, 2024)

Postdoctoral Associate at MIT July, 2018 – June, 2019
Institute for Foundations of Data Science
Advisor: Prof. Piotr Indyk and Prof. Aleksander Madry

Education CORNELL UNIVERSITY Aug, 2012 – May, 2018
Ph.D. in Computer Science
Advisor: Prof. Robert Kleinberg
Visiting student at Princeton University Sep, 2016 – Dec, 2016
Visiting student at Microsoft Research New England Sep, 2014 – June, 2015

PEKING UNIVERSITY Sep, 2008 – Jun, 2012
B.S. *summa cum laude* in Computer Science (honors track) GPA: 3.77/4.0, Rank: 2/138
Thesis advisor: Prof. Yao Guo. With Distinguished Dissertation Award (top 3%)

Publications (α - β) denotes alphabetical ordering. * denotes equal contribution.

Tensor Product Attention Is All You Need,
Yifan Zhang, Yifeng Liu, Huizhuo Yuan, Zhen Qin, Yang Yuan, Quanquan Gu, Andrew
Chi-Chih Yao,
Arxiv.

Cumulative Reasoning with Large Language Models,
Yifan Zhang, Jingqin Yang, Yang Yuan, Andrew C Yao,
TMLR 2025.

Autonomous Data Selection with Zero-shot Generative Classifiers for Mathematical Texts,
Yifan Zhang, Yifan Luo, Yang Yuan, Andrew C Yao,
ACL 2025 Findings.

Hierarchical Attention Generates Better Proofs,
Jianlong Chen, Chao Li, Yang Yuan, Andrew C Yao,
ACL 2025 Main.

CatCode: A Comprehensive Evaluation Framework for LLMs On the Mixture of Code
and Text,
Zhenru Lin, Yiqun Yao, Yang Yuan,
COLM 2024.

Matrix Information Theory for Self-Supervised Learning,
Yifan Zhang, Zhiqian Tan, Jingqin Yang, Weiran Huang, Yang Yuan,

ICML 2024.

Information Flow in Self-Supervised Learning,
Zhiquan Tan, Jingqin Yang, Weiran Huang, Yang Yuan, Yifan Zhang,
ICML 2024.

Contrastive Learning is Spectral Clustering on Similarity Graph,
Zhiquan Tan, Yifan Zhang, Jingqin Yang, Yang Yuan,
ICLR 2024.

Trade-off Between Efficiency and Consistency for Removal-based Explanations,
Yifan Zhang, Haowei He, Zhiquan Tan, Yang Yuan,
NeurIPS 2023.

Towards Data-Algorithm Dependent Generalization: a Case Study on Overparameterized
Linear Regression,
Jing Xu, Jiaye Teng, Yang Yuan, Andrew C. Yao,
NeurIPS 2023.

On the Power of Foundation Models,
Yang Yuan,
ICML 2023.

On Uni-Modal Feature Learning in Supervised Multi-Modal Learning,
Chenzhuang Du, Jiaye Teng, Tingle Li, Yichen Liu, Tianyuan Yuan, Yue Wang, Yang
Yuan, Hang Zhao,
ICML 2023.

Finding Generalization Measures by Contrasting Signal and Noise,
Jiaye Teng, Bohang Zhang, Ruichen Li, Haowei He, Yequan Wang, Yan Tian, Yang Yuan,
ICML 2023.

Predictive Inference with Feature Conformal Prediction,
Jiaye Teng*, Chuan Wen*, Dinghuai Zhang*, Yoshua Bengio, Yang Gao, Yang Yuan,
ICLR 2023.

Towards Understanding Generalization via Decomposing Excess Risk Dynamics,
Jiaye Teng, Jianhao Ma, Yang Yuan,
ICLR 2022.

T-SCI: A Two-Stage Conformal Inference Algorithm with Guaranteed Coverage for Cox-
MLP,
Jiaye Teng, Zeren Tan, Yang Yuan,
ICML 2021.

A Stratified Approach to Robustness for Randomly Smoothed Classifiers,
Guang-He Lee, Yang Yuan, Shiyu Chang, Tommi S. Jaakkola,
NeurIPS 2019.

Learning-Based Low-Rank Approximations,
Piotr Indyk, Ali Vakilian, Yang Yuan, (α - β)
NeurIPS 2019.

Asymmetric Valleys: Beyond Sharp and Flat Local Minima,
Haowei He, Gao Huang, Yang Yuan,
NeurIPS 2019 (spotlight).

Expanding Holographic Embeddings for Knowledge Completion,

Yexiang Xue, Yang Yuan, Ashish Sabharwal,
NIPS 2018.

An Alternative View: When Does SGD Escape Local Minima?
Robert Kleinberg, Yuanzhi Li and Yang Yuan, (α - β)
ICML 2018.

YATES: Rapid Prototyping for Traffic Engineering Systems,
Praveen Kumar, Chris Yu, Yang Yuan, Nate Foster, Robert Kleinberg, Robert Soulé,
SOSR 2018.

Semi-Oblivious Traffic Engineering: The Road Not Taken,
Praveen Kumar, Yang Yuan, Chris Yu, Nate Foster, Robert Kleinberg, Petr Lapukhov,
Chiun Lin Lim and Robert Soulé,
NSDI 2018.

Hyperparameter Optimization: A Spectral Approach,
Elad Hazan, Adam Klivans and Yang Yuan, (α - β)
ICLR 2018. Previously appeared in NIPS DLTP Workshop 2017 (oral).

Convergence Analysis of Two-layer Neural Networks with ReLU Activation,
Yuanzhi Li and Yang Yuan, (α - β)
NIPS 2017.

Exploiting the Structure: Stochastic Gradient Methods Using Raw Clusters,
Zeyuan Allen-Zhu*, Yang Yuan* and Karthik Sridharan,
NIPS 2016.

Even Faster Accelerated Coordinate Descent Using Non-Uniform Sampling,
Zeyuan Allen-Zhu, Zheng Qu, Peter Richtárik and Yang Yuan, (α - β)
ICML 2016.

Improved SVRG for Non-Strongly-Convex or Sum-of-Non-Convex Objectives,
Zeyuan Allen-Zhu and Yang Yuan, (α - β)
ICML 2016.

Optimization Algorithms for Computational Geometry,
Zeyuan Allen-Zhu, Zhenyu Liao and Yang Yuan, (α - β)
ICALP 2016.

Simultaneous Nearest Neighbor Search,
Piotr Indyk, Robert Kleinberg, Sepideh Mahabadi and Yang Yuan, (α - β)
SoCG 2016.

Combinatorial Multi-Armed Bandit and Its Extension to Probabilistically Triggered Arms,
Wei Chen, Yajun Wang, Yang Yuan and Qinshi Wang, (α - β)
JMLR 2016.

(Conference version:)

Combinatorial Multi-Armed Bandit: General Framework, Results and Applications,
Wei Chen, Yajun Wang and Yang Yuan, (α - β)
ICML 2013.

Escaping From Saddle Points – Online Stochastic Gradient for Tensor Decomposition,
Rong Ge, Furong Huang, Chi Jin and Yang Yuan, (α - β)
COLT 2015.

Optimal Auctions vs. Anonymous Pricing,
Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis and Yang Yuan,

$(\alpha\text{-}\beta)$
Games and Economic Behavior, Volume 118, November 2019, Pages 494-510.
Conference version: FOCS 2015.

Simple and Near-Optimal Mechanisms for Market Intermediation,
Rad Niazadeh, Yang Yuan and Robert Kleinberg,
WINE 2014.

On the Ratio of Revenue to Welfare in Single-Parameter Mechanism Design,
Robert Kleinberg and Yang Yuan, $(\alpha\text{-}\beta)$
EC 2013.

Boreas: An Accurate and Scalable Token-based Approach to Code Clone Detection,
Yang Yuan and Yao Guo,
ASE 2012.

A Fast Parallel Branch and Bound Algorithm for Treewidth,
Yang Yuan,
ICTAI 2011.

Teaching

Machine Learning

Fall 2019 – 2024

Type-safe Modern System Practice

Summer 2022 – 2024

Introduction to Artificial Intelligence

Fall 2019 – 2020

AI+X

Fall 2021

AI Research Practice

Fall 2022

Awards

Tsinghua Outstanding Worker	2024
Outstanding Headteacher (Second Price)	2021
Beijing Zhiyuan Scholar	2020
Forbes China 30 Under 30 (Science)	2019
Amazon AWS Research Award	2017
Microsoft Azure Research Award	2017
Jeff Hawkins & Janet Strauss Fellowship	2012
Pivot of Merit Student	2011
Tung OOCL Scholarship	2010
Kwang-Hua Scholarship	2009
Huaying Scholarship	2007 – 2008