

[Home](#) > [People](#) >

Vahab S. Mirrokni

Vahab Mirrokni is a Google Fellow and VP at Google Research, leading [algorithm and optimization research groups](#) at Google. These research teams include: [market algorithms](#), [large-scale graph mining](#), and [large-scale optimization](#). Previously he was a distinguished scientist and senior research director at Google. He received his PhD from MIT in 2005 and his B.Sc. from Sharif University of Technology in 2001. He joined Google Research in 2008, after research positions at Microsoft Research, MIT and Amazon.com. He is the co-winner of best paper awards at KDD, ACM EC, and SODA. His research areas include algorithms, distributed and stochastic optimization, and computational economics. Recently he has been working on various algorithmic problems in machine learning, online optimization and mechanism design, and [large-scale graph-based learning](#). His full publication list by year can be found [here](#). He also has an out-dated [personal homepage](#).

Research Areas

[Algorithms and theory](#)[Networking](#)[Anti abuse](#)[Software engineering](#)

Research

Authored Publications

Sort By ▾

See Filters ≡

Google ✕

Synthetic Text Generation for Training Large Language Models (LLMs) via Gradient Matching

Dang Nguyen • Zeman Li • [Mohammad "Hossein" Bateni](#) • [Vahab Mirrokni](#) • Meisam Razaviyayn • Baharan Mirzasoleiman • International Conference on Machine Learning (ICML) (2025)

Titans: Learning to Memorize at Test Time

Ali Behrouz • Peilin Zhong • [Vahab Mirrokni](#) • The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS) (2025)

PiKE: Adaptive Data Mixing for Large-Scale Multi-Task Learning Under Low Gradient Conflicts

Zeman Li • [Yuan Deng](#) • Peilin Zhong • Meisam Razaviyayn • [Vahab Mirrokni](#) • 2025

TRELLIS: Learning to Compress Key-Value Memory in Attention Models

Mahdi Karami • Ali Behrouz • Praneeth Kacham • [Vahab Mirrokni](#) • 2025

Lattice: Learning to Compress the Cache in the Attention

Mahdi Karami • Razvan Pascanu • [Vahab Mirrokni](#) • 2025

Titans: Learning to Memorize at Test Time

[Vahab Mirrokni](#) • Ali Behrouz • Peilin Zhong • 2025

Procurement Auctions via Approximate Submodular Optimization

[Yuan Deng](#) • Amin Karbasi • [Vahab Mirrokni](#) • [Renato Paes Leme](#) • Grigoris Veleghkas • [Song Zuo](#) • Forty-second International Conference on Machine Learning (2025)

Addax: Utilizing Zeroth-Order Gradients to Improve Memory Efficiency and Performance of SGD for Fine-Tuning Language Models

Zeman Li • Xinwei Zhang • Peilin Zhong • [Yuan Deng](#) • Meisam Razaviyayn • [Vahab Mirrokni](#) • 2025

Trellis: Learning to Compress Key-Value Memory in Attention

Addax: Utilizing Zeroth-Order Gradients to Improve Memory Efficiency and Performance of SGD for Fine-Tuning Language Models

Zeman Li • Xinwei Zhang • Peilin Zhong • [Yuan Deng](#) • Meisam Razaviyayn • [Vahab Mirrokni](#) • ICLR (2025)

Join us

We're always looking for more talented, passionate people.

[See opportunities](#)



Follow us



Explore our other initiatives

Google AI

Discover how Google AI is committed to enriching knowledge and solving complex challenges

Google Cloud

High-performance infrastructure for cloud computing, data analytics & machine learning

Research

Products

Overview

Build

Solutions

Research

Products

Responsibility

Pricing

Societal Impact

Resources

About

Google DeepMind

Our mission is to build AI responsibly to benefit humanity

Models

Research

Science

About

Google Labs

Explore the future of AI responsibly with Google Labs

About

Experiments

Stay connected

Google

About Google

Google Products

Privacy

Terms