

Grigoris Velegkas

CONTACT INFORMATION

Email: gvelegkas41@gmail.com
Phone: +1-475-434-6748

RESEARCH INTERESTS

Statistical Learning Theory, Algorithmic Game Theory, Reinforcement Learning

EDUCATION

Yale University, New Haven, CT, USA 2019 - 2025

PhD in Computer Science, Advisor: Amin Karbasi

National Technical University of Athens (NTUA), Athens, Greece 2013 - 2018

BSc & MSc in Electrical Engineering and Computer Science, Honors

EMPLOYMENT

Google Research, Research Scientist (US-MTV) August 2025 - current
Algorithms and Optimization Teams

Google Research, Research Intern (US-NYC) May 2024 - October 2024
Host: William Kong

Google Research, Student Researcher in Market Algorithms (US-NYC) October 2023 - May 2024
Host: Yuan Deng

Google Research, Research Intern in Market Algorithms (US-MTV) May 2023 - August 2023
Hosts: Andres Perlroth, Gagan Aggarwal

HONORS & AWARDS

Leventis Foundation Scholarship for Academic Excellence 2024-2025

Onassis Foundation Scholarship for Academic Excellence 2020-2024

Bodossaki Foundation Scholarship for Academic Excellence 2020-2024

Gerondelis Foundation Scholarship for Academic Excellence 2019-2020

Eurobank “A great moment for education” Award 2013

PUBLICATIONS

On Agnostic PAC Learning in the Small Error Regime

Julian Asilis, Mikael Møller Høgsgaardm, Grigoris Velegkas (alphabetical order)
39th Annual Conference on Neural Information Processing Systems, 2025 (**NeurIPS 2025**)
Spotlight

On Union-Closedness of Language Generation

Steve Hanneke, Amin Karbasi, Anay Mehrotra, Grigoris Velegkas (alphabetical order)
39th Annual Conference on Neural Information Processing Systems, 2025 (**NeurIPS 2025**)

(Im)possibility of Automated Hallucination Detection in Large Language Models

Amin Karbasi, Omar Montasser, John Sous, Grigoris Velegkas (alphabetical order)
2nd Conference on Language Modeling (**COLM 2025**)

Procurement Auctions via Approximately Optimal Submodular Optimization

Yuan Deng, Amin Karbasi, Vahab Mirrokni, Renato Paes Leme, Grigoris Velegkas, Song Zuo (alphabetical order)
42nd International Conference on Machine Learning, 2025 (**ICML 2025**)
Spotlight

On the Limits of Language Generation: Trade-Offs Between Hallucination and Mode-Collapse

Alkis Kalavasis, Anay Mehrotra, Grigoris Velegkas (alphabetical order)
57th Annual ACM Symposium on Theory of Computing (**STOC 2025**)

Understanding Aggregations of Proper Learners in Multiclass Classification

Julian Asilis, Mikael Møller Høgsgaardm, Grigoris Velegkas (alphabetical order)
36th International Conference on Algorithmic Learning Theory (**ALT 2025**)

Randomized Truthful Auctions with Learning Agents

Gagan Aggarwal, Anupam Gupta, Andres Perloth, Grigoris Velegkas (alphabetical order)
38th Annual Conference on Neural Information Processing Systems, 2024 (**NeurIPS 2024**)

Universal Rates for Active Learning

Steve Hanneke, Amin Karbasi, Shay Moran, Grigoris Velegkas (alphabetical order)
38th Annual Conference on Neural Information Processing Systems, 2024 (**NeurIPS 2024**)

On the Computational Landscape of Replicable Learning

Alkis Kalavasis, Amin Karbasi, Grigoris Velegkas, Felix Zhou (alphabetical order)
38th Annual Conference on Neural Information Processing Systems, 2024 (**NeurIPS 2024**)

Injecting Undetectable Backdoors in Deep Learning and Language Models

Alkis Kalavasis, Amin Karbasi, Argyris Oikonomou, Katerina Sotiraki, Grigoris Velegkas, Manolis Zampetakis (alphabetical order)
38th Annual Conference on Neural Information Processing Systems, 2024 (**NeurIPS 2024**)

Universal Rates for Regression: Separations between Cut-Off and Absolute Loss

Idan Attias, Steve Hanneke, Alkis Kalavasis, Amin Karbasi, Grigoris Velegkas (alphabetical order)
37th Annual Conference on Learning Theory (**COLT 2024**)

Replicable Learning of Large-Margin Halfspaces

Alkis Kalavasis, Amin Karbasi, Kasper Green Larsen, Grigoris Velegkas, Felix Zhou (alphabetical order)
41st International Conference on Machine Learning, 2024 (**ICML 2024**)

Spotlight

User Response in Ad Auctions: An MDP Formulation of Long-term Revenue Optimization

Yang Cai, Zhe Feng, Christopher Liaw, Aranyak Mehta, Grigoris Velegkas (alphabetical order)
Web Conference on World Wide Web, 2024 (**WWW 2024**)

Optimal Learners for Realizable Regression: PAC Learning and Online Learning

Idan Attias, Steve Hanneke, Alkis Kalavasis, Amin Karbasi, Grigoris Velegkas (alphabetical order)
37th Annual Conference on Neural Information Processing Systems, 2023 (**NeurIPS 2023**)

Oral Presentation

Replicable Clustering

Hossein Esfandiari, Amin Karbasi, Vahab Mirrokni, Grigoris Velegkas, Felix Zhou (alphabetical order)
37th Annual Conference on Neural Information Processing Systems, 2023 (**NeurIPS 2023**)

Replicability in Reinforcement Learning

Amin Karbasi, Grigoris Velegkas, Lin F. Yang, Felix Zhou (alphabetical order)
37th Annual Conference on Neural Information Processing Systems, 2023 (**NeurIPS 2023**)

Statistical Indistinguishability of Learning Algorithms

Alkis Kalavasis, Amin Karbasi, Shay Moran, Grigoris Velegkas (alphabetical order)
40th International Conference on Machine Learning, 2023 (**ICML 2023**)

Replicable Bandits

Hossein Esfandiari, Alkis Kalavasis, Amin Karbasi, Andreas Krause, Vahab Mirrokni, Grigoris Velegkas (alphabetical order)
11th International Conference on Learning Representations, 2023 (**ICLR 2023**)

Universal Rates for Interactive Learning

Steve Hanneke, Amin Karbasi, Shay Moran, Grigoris Velegkas (alphabetical order)
36th Annual Conference on Neural Information Processing Systems, 2022 (**NeurIPS 2022**)

Oral Presentation

Multiclass Learnability Beyond the PAC Framework: Universal Rates and Partial Concept Classes

Alkis Kalavasis*, Grigoris Velegkas*, Amin Karbasi

36th Annual Conference on Neural Information Processing Systems, 2022 (**NeurIPS 2022**)

Reinforcement Learning with Logarithmic Regret and Policy Switches

Grigoris Velegkas, Zhuoran Yang, Amin Karbasi

36th Annual Conference on Neural Information Processing Systems, 2022 (**NeurIPS 2022**)

Is Selling Complete Information (Approximately) Optimal?

Dirk Bergemann, Yang Cai, Grigoris Velegkas, Mingfei Zhao (alphabetical order)

23rd ACM Conference on Economics and Computation, 2022 (**EC 2022**)

An Efficient ϵ -BIC to BIC Transformation and Its Application to Black-Box Reduction in Revenue Maximization

Yang Cai, Argyris Oikonomou, Grigoris Velegkas, Mingfei Zhao (alphabetical order)

32nd ACM-SIAM Symposium on Discrete Algorithms, 2021 (**SODA 2021**)

How to Sell Information Optimally: an Algorithmic Study

Yang Cai, Grigoris Velegkas (alphabetical order)

12th Innovations in Theoretical Computer Science, 2021 (**ITCS 2021**)

(* Denotes equal contribution)

PREPRINTS

Characterizations of Language Generation With Breadth

Alkis Kalavasis, Anay Mehrotra, Grigoris Velegkas (alphabetical order)

(**under submission**)

Pointwise Lipschitz Continuous Graph Algorithms via Proximal Gradient Analysis

Quanquan C. Liu, Grigoris Velegkas, Yuichi Yoshida, Felix Zhou (alphabetical order)

(**under submission**)

Efficient Gradient Clipping Methods in DP-SGD for Convolution Models

William Kong, Monica Ribero, Grigoris Velegkas (alphabetical order)

(**under submission**)

ACADEMIC SERVICE **Program Committee Member:** Conference on Web and Internet Economics (WINE), Conference on Learning Theory (COLT), Conference on Economics and Computation (EC)

Journal Reviewer: Transactions on Economics and Computation (TEAC), Journal of Machine Learning Research (JMLR)

Conference Reviewer / Subreviewer: Conference on Neural Information Processing Systems (NeurIPS), International Conference on Machine Learning (ICML), International Conference on Learning Representations (ICLR), Symposium on Foundations of Computer Science (FOCS), Symposium on Theory of Computing (STOC), Symposium on Discrete Algorithms (SODA), Conference on Web and Internet Economics (WINE)

Tutorial: [Language Generation in the Limit](#), COLT 2025