

Themistoklis Haris

Boston University
Center for Computing and Data Sciences,
665 Commonwealth Ave, Boston,
MA 02215

Email: tharis@bu.edu
URL: <https://themisharis.com>
Github: <https://github.com/tkhar>

Current positions

Computer Science PhD Student
Research Intern, Google

Advisors: Krzysztof Onak, Venkatesh Saligrama
Host: Maryam Karimzadehgan

Research Interests

Foundation Models • Theoretical Computer Science
Sublinear Algorithms • Efficient Transformers • LLM Interpretability • 3D Generation • Differential Privacy • Diffusion Models • Adversarial Robustness • Discrete Optimization • Graph Algorithms • Theoretical Deep Learning

Professional Experience

- 2025 **Research Internship** at the National Institute of Informatics, Tokyo, Japan: Work on Transformer Model Sensitivity with Professor Yuichi Yoshida
- 2021-2023 **CoreOS Kernel Software Engineer, Microsoft**: Hypervisor Memory Management, OS Containerization and Kernel Object Management.

Education

- 2027 PhD in Computer Science, Boston University (Expected January 2027)
- 2024 MSc in Computer Science, Boston University (GPA: 4.0 / 4.0)
- 2021 BA in Computer Science and Mathematics, Dartmouth College (GPA: 3.94 / 4.0)

Publications & Talks

CONFERENCE PAPERS

- 2026 Fabian Spaeh, *Themistoklis Haris*, Huy Nguyen, Alina Ene, “**Efficient Algorithms for Influence Maximization in General Models and Observed Cascades**” (IJCAI 2026)
- 2026 *Themistoklis Haris*, Diptaksho Palit, “**Testing k -submodularity**” (RANDOM 2026)
- 2026 *Themistoklis Haris*, Fabian Spaeh, Nithin Varma, Yuichi Yoshida, “**Stable Max Coverage Under a Cardinality Constraint**” (in submission)
- 2026 Patrick Lutz, *Themistoklis Haris*, Arjun Chandra, Aditya Gangrade, Venkatesh Saligrama, “**Transformers Discover Efficient Preconditioners**” (in submission)
- 2026 Patrick Lutz, *Themistoklis Haris*, Arjun Chandra, Aditya Gangrade, Venkatesh Saligrama, “**Layerwise Dynamics for In-Context Classification in Transformers**” (ICML 2026)

- 2026 *Themistoklis Haris*, Zihan Zhang, Yuichi Yoshida “**Noise Stability of Transformer Models**” (ICLR 2026)
- 2026 *Themistoklis Haris*, Steve Choise Mutiraj Laksanawisit “**Fast-MWEM: Private Data Release in Sublinear Time**” (TPDP 2026)
- 2025 *Themistoklis Haris*, Fabian Spaeh, Spyros Dragazis, Charalampos Tsourakakis “**Estimating Hitting Times Locally At Scale**” (NeurIPS 2025)
- 2025 Alexandr Andoni, *Themistoklis Haris*, Esty Kelman, Krzysztof Onak, “**From Search to Decision: A Framework for Adversarially Robust Approximate Nearest Neighbor Search**” (NeurIPS 2025 Workshop on Adversarial Examples in Machine Learning)
- 2025 *Themistoklis Haris*, Krzysztof Onak, “**Compression Barriers for Autoregressive Transformers**” ([Paper](#)) (COLT 2025)
- 2024 *Themistoklis Haris*, “**Demystifying kNN Attention: A Theoretical Foundation**” ([Paper](#)) (ICLR 2025)
- 2022 *Haris*, *Themistoklis* and Chakrabarti, Amit, “**Counting Simplices in Hypergraph Streams**,” Proceedings of the European Symposium on Algorithms (ESA), Sep 5-7, 2022, LIPIcs.ESA.2022.32, pp. 32:1-32:19, [URL](#)

JOURNAL PAPERS

- 2019 Qijia Shao, Amy Sniffen, Julien Blanchet, Megan Hillis, Xinyu Shi, *Themistoklis Haris* et al. 2020. “**Teaching American Sign Language in Mixed Reality**”. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 4, 4, Article 152 (December 2020), 27 pages. <https://doi.org/10.1145/3432211>

TALKS

- 2026 On efficient $2D \rightarrow 3D$ generation (Google)
- 2026 On Noise Stability for Transformer Models (BU AI Seminar)
- 2026 On Efficient Autoregressive Inference (Google)
- 2025 On Language Generation in the Limit (BU Oral Exam)
- 2025 On Efficient Attention Mechanisms: AI Research Seminar, Boston University
- 2025 On Robust Data Structures: Theory Seminar Talk, Boston University
- 2025 Theory Seminar Talk, National Institute of Informatics [Tokyo, Japan]
- 2025 Talk on space lower bounds for Transformers. [COLT, Lyon]
- 2025 Talk on Adversarially Robust Algorithms for Searching (NTU Singapore)
- 2025 Talk on space lower bounds for Transformers (EnCore LLM Theory Workshop, UCSD)
- 2025 Talk on k -NN Methods for Efficient Attention Computation (Piotr Indyk Reading Group, MIT)
- 2023 *CoreOS Learning Day*: Outlined the Windows Container Virtualization Stack to a cohort of around 100 Microsoft Engineers and Program Managers.
- 2022 *Counting Simplices in Hypergraph Streams*, ESA 2022, 9/7/2022, Potsdam, Germany

Grants, honours & awards

2025

2025 **Top Reviewer**, NeurIPS 2025, ICLR 2026, ICML 2026
 2025 **Outstanding Reviewer** award at “SCOPE” ICLR workshop 2025.
 2023 **Boston University Dean’s Fellowship**
 2022 **CoreOS Hackathon Winner**, Microsoft
 2021 **Lovelace Research Scholarship**, Dartmouth College
 2021 **Francis L. Town Scientific Prize in Computer Science, Dartmouth College**: The prize is offered annually to “one meritorious and deserving student in each department of scientific study at the College” at the end of the sophomore year.
 2020 **James O. Freedman Presidential Scholar**
 2017-2022 **13 Citations for Academic Excellence**, Dartmouth College

Service

2025 **Reviewer** for TMLR, ICML 2026, ICLR 2026, AISTATS 2026, NeurIPS 2026, ICLR 2025, SCOPE workshop ICLR 2025.
 2024 **Theory Group Seminar Organizer and Website Maintainer**, Boston University
 2023- **Social Activities Committee**: Organized social events for the CS PhD community at BU
 2019-2021 **President of Dartmouth CoderDojo**: Designed curricula and taught computer science to local high-school students.

Workshops and Conferences

2025 “Theory in LLMs” workshop, EnCore Institute, UCSD.
 2024 “Frontiers in Complexity Theory: A Graduate Workshop”, DIMACS Institute
 2024 “Workshop on Extroverted Sublinear Algorithms”, Simons Institute, UC Berkeley.
 2023 “Theory and Practice of Foundation Models” workshop, *Yale FDS*

Teaching

2025 **Teaching Assistant**: for *Algorithms*, Boston University
 2024 **Summer Challenge at BU**: Computer Science Instructor for a high school program
 2021 **Teaching Fellow** for *Theory of Computation*, Dartmouth College
 2020 **Teaching Assistant** for *Discrete Math, Machine Learning, Computer Graphics, Algorithms, Software Engineering*, Dartmouth College

Technical Skills

PROGRAMMING LANGUAGES

C • C++ • Python • Java • Ruby • JavaScript • Rust • R# • CUDA • Triton

DEEP LEARNING

LLM Pre-Training, PyTorch, Transformers, Diffusion Models, Computer Vision, 3D Generation, LLM Inference Optimization, Efficient Attention Mechanisms

ALGORITHMS

Discrete and Continuous Optimization, Graph Theory, Complexity Theory