# Matthew Faw

+1 (336)-262-1938 • mfaw3@gatech.edu • matthewfaw.github.io

## Research Interests

Stochastic Optimization, Online Decision-Making, Reinforcement Learning and Control, Time-series Forecasting

## Education

**Georgia Institute of Technology** **Atlanta, GA**
*ARC Postdoctoral Fellow (Focus: Model-free/online control, Online resource allocation)* *2024–present*
Mentors: Siva Theja Maguluri, Sahil Singla

**The University of Texas at Austin** **Austin, TX**
*Ph.D. Machine Learning (Thesis: Adaptive Algorithms for Stochastic Optimization and Bandit Learning),* *2018–2024*
Advisors: Sanjay Shakkottai, Constantine Caramanis.

**Duke University** **Durham, NC**
*B.S.E. Electrical & Computer Engineering, B.S. Computer Science, A.B. Math (cum laude),* *2013–2017*
Advisors: Nick Buchler, Richard Fair, Benjamin C. Lee

## Industry Experience

**Google Research** **Remote**
*Visiting Researcher* *March 2025–September 2025*
Designed, trained, and evaluated novel discrete tokenization approaches to time-series forecasting with transformers.

**Google Research** **Mountain View, CA**
*Ph.D. Student Researcher* *June-September 2024*
Designed, trained, and evaluated an in-context learning approach for time-series forecasting using a decoder-only transformer architecture.

**Verato** **McLean, VA**
*Software Engineer (Full-time)* *June 2017-July 2018*
Co-designed and built a custom CI system using Kubernetes. Created software for entity search, updates, and analysis.

**MathWorks** **Natick, MA**
*Software Engineering Intern, Stateflow Semantics* *May-August 2016*
Co-designed + implemented proof-of-concept architectural change to code generation process to improve extensibility and performance.

## Selected Publications (Full List on Google Scholar)

### In Submission

**2026**: "Online Control with Multiple Sensors", **F**, S. T. Maguluri (*Georgia Tech*)
*Proposed a generalization of the online control problem, and proved near-optimal regret upper-bounds for a novel algorithm.*

### Conference Papers

**ICML 2025**: "In-Context Fine-Tuning for Time-Series Foundation Models", **F**, R. Sen, Y. Zhou, A. Das (*Georgia Tech, Google Research*)
*Empirical evaluation of a new approach to time-series forecasting using in-context learning and a decoder-only transformer model.*

**ICML 2025**: "On Mitigating Affinity Bias through Bandits with Evolving Biased Feedback", **F**, C. Caramanis, J. Hoffmann (*Georgia Tech, UT Austin, Google DeepMind*)
*Algorithms and new lower bounds for regret under an affinity-biased feedback model (preliminary results appeared at NeurIPS 2023 workshop: Algorithmic Fairness through the Lens of Time)*

**COLT 2023**: "Beyond Uniform Smoothness: A Stopped Analysis of Adaptive SGD", **F**$^{=}$, L. Rout$^{=}$, C. Caramanis, S. Shakkottai (*UT Austin*)
*First algorithm which converges at order-optimal rate under a generalized smoothness assumption in standard noise regime.*

**COLT 2022**: "The Power of Adaptivity in SGD: Self-Tuning Step Sizes with Unbounded Gradients and Affine Variance", **F**$^{=}$, I. Tziotis$^{=}$, C. Caramanis, A. Mokhtari, S. Shakkottai, R. Ward (*UT Austin*)
*Resolved an open problem on convergence of AdaGrad-Norm for smooth stochastic non-convex optimization.*

**SIGMETRICS 2022/POMACS**: "Learning To Maximize Welfare with a Reusable Resource", **F**$^{=}$, O. Papadigenopoulos$^{=}$, C. Caramanis, S. Shakkottai (*UT Austin*)

*Optimal prophet inequalities, learning variants, and lower bounds for prophet inequalities with dynamic constraints.*

**SODA 2022, Theory of Computing 2026**: "Single-Sample Prophet Inequalities via Greedy-Ordered Selection", C. Caramanis, P. Dütting, **F**, P. Lazos, S. Leonardi, O. Papadigenopoulos, E. Pountourakis, R. Reiffenhäuser (alphabetical order, *UT Austin, Google Switzerland, Sapienza University of Rome, Drexel*)
*Improved single-sample prophet inequalities for nearly all combinatorial settings considered in prior work.*

**NeurIPS 2020**: "Mix and Match: An Optimistic Tree-Search Approach for Learning Models from Mixture Distributions", **F**, R. Sen, K. Shanmugam, C. Caramanis, S. Shakkottai (*UT Austin, IBM Research*)
*Optimistic bandit tree-search for multi-source domain adaptation.*

### Journal Papers

**Theory of Computing 2026**: "Single-Sample Prophet Inequalities via Greedy-Ordered Selection", C. Caramanis, P. Dütting, **F**, P. Lazos, S. Leonardi, O. Papadigenopoulos, E. Pountourakis, R. Reiffenhäuser (alphabetical order, *UT Austin, Google Switzerland, Sapienza University of Rome, Drexel*)
*Journal version of SODA'22*

**POMACS 2022**: "Learning to Maximize Welfare with a Reusable Resource", **F**$^{=}$, O. Papadigenopoulos$^{=}$, C. Caramanis, S. Shakkottai (*UT Austin*)
*Journal version of SIGMETRICS'22*

**TOCS 2017**: "Computational Sprinting: Architecture, Dynamics, and Strategies", S. Zahedi, S. Fan, **F**, E. Cole, B. Lee (*Duke University*)
*Evaluated performance and economic viability of several strategies for system-level computational sprinting for Spark applications.*

## Awards + Honors

**2023**: Dr. Brooks Carlton Fowler Endowed Presidential Graduate Fellowship in ECE, 2023-2024 academic year

**2022**: Top 10% reviewer for NeurIPS'22 and AISTATS'22, Highlighted reviewer for ICLR 2022

**2020**: NXP Foundation Fellowship, 2020-2021 academic year

**2017**: Cum Laude Graduation Honors, Duke University

**2016**: Member, Tau Beta Pi and Eta Kappa Nu Honor Societies, Duke University

**2014**: Gold medal, International Genetically Engineered Machine Competition

## Talks and Poster Presentations

### Invited Talks

**October 2025**: INFORMS Job Market Showcase, Atlanta, GA: "Optimal Control with Multiple Sensors"

**July 2025**: INFORMS APS Conference, Atlanta, GA: "Fundamental Limits of Regret Minimization in Stochastic Bandits with Evolving Biased Feedback"

**February 2025**: Google Research, (Virtual): "Order-Optimal Convergence Rates with Adaptive SGD"

**October 2024**: INFORMS Annual Meeting, Seattle, WA: "Order-Optimal Convergence Rates with Adaptive SGD"

**March 2024**: Georgia Tech ARC Colloquium, Atlanta, GA: "The Power of Adaptivity in SGD"

### Talks

**July 2023**: COLT 2023, Bangalore, India: "Beyond Uniform Smoothness: A Stopped Analysis of Adaptive SGD"

**April 2023**: IFML Workshop, UW "Beyond Uniform Smoothness: A Stopped Analysis of Adaptive SGD"

**July 2022**: COLT 2022, London, UK: "The Power of Adaptivity in SGD: Self-Tuning Step Sizes with Unbounded Gradients and Affine Variance"

**June 2022**: SIGMETRICS 2022, IIT Bombay, Mumbai, IN: "Learning To Maximize Welfare with a Reusable Resource"

**April 2022**: Machine Learning Lab Research Symposium, UT Austin: "The Power of Adaptivity in SGD: Self-Tuning Step Sizes with Unbounded Gradients and Affine Variance"

**January 2022**: SODA 2022, Virtual: "Single Sample Prophet Inequalities via Greedy-Ordered Selection"

### Poster Presentations

**October 2022**: Joint IFML/Data-Driven Decision Processes Workshop, Simons Institute, UC Berkeley, "The Power of Adaptivity in SGD: Self-Tuning Step Sizes with Unbounded Gradients and Affine Variance"

**December 2020**: NeurIPS 2020, Virtual, "Mix and Match: An Optimistic Tree-Search Approach for Learning Models from

Mixture Distributions"

**November 2019**: Texas Wireless Summit, UT Austin, "Mix and Match: An Optimistic Tree-Search Approach for Learning Models from Mixture Distributions"

## Conference Reviewing

**2021-Present**: AISTATS, ALT, ICLR, ICML, JMLR, NeurIPS

## Teaching Experience

**Georgia Tech**: ISyE 3770 Statistics and Applications *Instructor*

**UT Austin**: EE 460J Data Science Lab *TA*

**Duke**: CS 308 Software Design and Implementation *TA*, ECE 280 Signals & Systems *TA*, Synthetic Biology House Course *Co-Instructor*

## Technical Skills

**Programming**: Python, Java, C/C++, JavaScript

**Infrastructure**: Kubernetes, AWS, Google Cloud, Mongo, Solr

## References

**Sanjay Shakkottai** (Ph.D. Advisor) sanjay.shakkottai@utexas.edu

**Constantine Caramanis** (Ph.D. Advisor) constantine@utexas.edu

**Siva Theja Maguluri** (Postdoc Mentor) siva.theja@gatech.edu