Caleb Koch

Contact Information

(434)-857-6440
ckoch@stanford.edu
cs.stanford.edu/~ckoch/

Research Interests

learning theory, algorithms, complexity theory, explainable ML

Education

Stanford University — Sept 2020 –
PhD in Computer Science
Advisor: Li-Yang Tan

Cornell University — May 2020
Bachelor of Arts in Computer Science (*summa cum laude*) and Mathematics (*magna cum laude*)

University of Oxford, St. Catherine's College — 2018 – 2019
Visiting Student

Awards

2024	FOCS 2024 Special Issue Invitation
2024	CCC 2024 Special Issue Invitation
2020–2023	NDSEG Fellowship
2020	NSF Fellowship (declined)
2020	Jack Kent Cooke Graduate Scholarship (2 years funding, deferred)
2020	Phi Beta Kappa, Cornell University
2019	Computer Science Book Prize. St. Catherine's College, Oxford. (awarded for top performance in computer science)
2016–2020	Jack Kent Cooke Undergraduate Scholarship (full tuition, 4 years)

Publications

Fast decision tree learning solves hard coding-theoretic problems, with Carmen Strassle and Li-Yang Tan. *IEEE Symposium on Foundations of Computer Science (FOCS) 2024*.

The sample complexity of smooth boosting and the tightness of the hardcore theorem, with Guy Blanc, Alexandre Hayderi, and Li-Yang Tan. *IEEE Symposium on Foundations of Computer Science (FOCS) 2024*. Invited to the special issue of FOCS.

A strong direct sum theorem for distributional query complexity , with Guy Blanc, Carmen Strassle, and Li-Yang Tan. *Computational Complexity Conference (CCC) 2024*. Invited to the special issue of CCC.

Superconstant inapproximability of decision tree learning, with Carmen Strassle and Li-Yang Tan. *Conference on Learning Theory (COLT) 2024*.

A strong composition theorem for junta complexity and the boosting of property testers, with Guy Blanc, Carmen Strassle, and Li-Yang Tan. *IEEE Symposium on Foundations of Computer Science (FOCS) 2023*.

Properly learning decision trees with queries is NP-hard, with Carmen Strassle and Li-Yang Tan. *IEEE Symposium on Foundations of Computer Science (FOCS) 2023*.

Certification with an NP oracle, with Guy Blanc, Jane Lange, Carmen Strassle, and Li-Yang Tan. *Innovations in Theoretical Computer Science (ITCS) 2023*.

Automata learning with an incomplete teacher, with Mark Moeller, Thomas Wiener,

Alaia Solko-Breslin, Nate Foster, Alexandra Silva. *European Conference on Object-Oriented Programming (ECOOP) 2023*.

Superpolynomial lower bounds for decision tree learning and testing, with Carmen Strassle, and Li-Yang Tan. *ACM-SIAM Symposium on Discrete Algorithms (SODA) 2023*.

A query-optimal algorithm for finding counterfactuals, with Guy Blanc, Jane Lange, and Li-Yang Tan. *International Conference on Machine Learning (ICML) 2022*.

The query complexity of certification, with Guy Blanc, Jane Lange, and Li-Yang Tan. *ACM Symposium on Theory of Computing (STOC) 2022*.

Hyperprofile-based computation offloading for mobile edge networks, with Andrew Crutcher, Kyle Coleman, Jon Patman, Flavio Esposito, and Prasad Calyam. *IEEE Conference on Mobile Ad Hoc and Sensor Systems (MASS) 2017*.

Internships

Simons Institute for the Theory of Computing — Summer 2024
Visiting Student Researcher
Sublinear algorithms program

Summer 2019

Cornell University, Computer Science Department
Undergraduate Researcher
Identity-Suppressed Decision Diagrams for NetKAT
Advisor: Nate Foster

JPMorgan Chase — Summer 2018
Software development intern

University of Missouri – Columbia — Summer 2017
CS REU in Networks, Undergraduate Researcher
Computation Offloading in Mobile Edge Networks
Advisor: Prasad Calyam

US Army Corps of Engineers, Geospatial Research Laboratory — Summer 2016
Research Intern
Interpolation Methods for Mining Trajectory Data
Advisor: Crystal Chen

Teaching Experience

Course Assistant, CS 254B at Stanford (Complexity Theory II)	Spring 2024
Course Assistant, CS 254 at Stanford (Complexity Theory I)	Winter 2024
Teaching Assistant, CS 4810 (theory of computing) at Cornell	Fall 2019
Course Assistant, Math 1120 at Cornell (Calculus II)	Fall 2016, Spring 2017, Fall 2017

Software libraries

Identity-suppressed decision diagrams, with Steffen Smolka, Eamon Woods, and Nate Foster. *OCaml Package Manager*. `opam.ocaml.org/packages/idds/`. 2019.

NetKAT, with Steffen Smolka, and Nate Foster. *OCaml Package Manager*. `opam.ocaml.org/packages/netkat/`. 2019.

Talks

A strong direct sum theorem for distributional query complexity, July 2024.
CCC Conference Presentation.

Superconstant inapproximability of decision tree learning, June 2024.
COLT Conference Presentation.

Fast decision tree learning solves hard coding-theoretic problems, March 2024.
Cornell Theory Seminar.
Columbia CS Theory Student Seminar.
Penn Theory Seminar.

A strong composition theorem for junta complexity and the boosting of property testers, November 2023.
FOCS Conference Presentation.

The complexity of DNF minimization, June 2023.
PhD Qualifying Exam Presentation.

Block sensitivity and Boolean function composition, October 2022.
Stanford Teach Us Anything Seminar.

A Query-Optimal Algorithm for Finding Counterfactuals, July 2022.
ICML Conference Presentation.

The Query Complexity of Certification, June 2022.
STOC Conference Presentation.

The Query Complexity of Certification, December 2021.
Stanford Theory Lunch.

PARITY $\notin AC^0$, September 2019.
Cornell Math Club.

Hyperprofile-Based Computation Offloading for Mobile Edge Networks. November 2017.
IEEE Mobile Adhoc and Sensor Systems.

Professional activities

Research Blog Posts
DNF Minimization, Part I. Theory Dish. `https://theorydish.blog/2023/08/28/dnf-minimization-part-i/`.
DNF Minimization, Part II. Theory Dish. `https://theorydish.blog/2023/08/28/dnf-minimization-part-ii/`.

Workshops
Swiss Winter School on Theoretical Computer Science, January 2023, PhD student participant

Journal Reviewing
ACM Transactions on Algorithms (TALG). 2021

Conference Reviewing
ACM Symposium on Theory of Computing (STOC), 2023
Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2023
International Colloquium on Automata, Languages, and Programming (ICALP), 2023
ACM-SIAM Symposium on Discrete Algorithms (SODA), 2023, 2024