

MAXWELL FISHELSON

maxkfish.com

maxfish@mit.edu ◊ (917) 434-7425

EDUCATION

Ph.D. in Theoretical Computer Science **Massachusetts Institute of Technology (MIT)**
Supervisor: Prof. Constantinos Daskalakis

B.S. in Mathematics **Massachusetts Institute of Technology (MIT)**
GPA: 4.9/5

RESEARCH

From External to Swap Regret 2.0: An Efficient Reduction for Large Action Spaces **STOC 2024**

- Devised the Tree-Swap algorithm: the first swap regret minimization algorithm with rate depending logarithmically on the number of actions N , resolving an open problem of Blum and Mansour from 2007
- Proved an equivalent swap regret rate only dependent on the Littlestone dimension of a concept class, establishing the first no-swap-regret algorithm for potentially infinite concept classes
- Established matching lower bounds, showing even oblivious, ℓ_1 -constrained adversaries can force a swap regret learner to expend exponentially many rounds
- Improved the best known dependence on N for swap regret learning in the setting of bandit feedback
- Resolved open problems about the query and communication complexity of computing ϵ -approximate correlated equilibrium in games
- Proved ϵ -approximate correlated equilibria in extensive-form games can be computed efficiently for constant ϵ , advancing a long-standing open problem for extensive-form games
- <https://arxiv.org/pdf/2310.19786.pdf>

Online Learning and Solving Infinite Games with an ERM Oracle **COLT 2023**

- Devised the first algorithm for online learning of arbitrary potentially-infinite hypothesis classes using only ERM oracle calls
- Provided a theoretical backing to the prevalent “double-oracle” algorithms that are the standard for solving large games in practice, unlike previous online learning algorithms that relied on the SOA oracle and had no practical utility
- Obtained similar results for non-parametric games, providing learning algorithms that only rely on best response oracles and converge to approximate-minimax equilibria in two-player zero-sum games and approximate coarse correlated equilibria in multi-player general-sum games
- <https://arxiv.org/pdf/2307.01689.pdf>

Near-Optimal No-Regret Learning for Correlated Equilibria in Multi-Player General-Sum Games **STOC 2022**

- Constructed algorithmic extensions of Optimistic Hedge attaining $poly(\log T)$ internal-regret and swap-regret in multi-player general-sum games respectively
- Established a means of simulating the Stoltz-Lugosi Optimistic-Hedge algorithm as an instance of Optimistic Hedge on a combinatorial space, demonstrating that the no-external-regret of Optimistic Hedge implies no-internal-regret for Stoltz-Lugosi

- Introduced novel techniques for the analysis of Taylor expansions of multinomial functions arising from no-regret algorithms, enabling a proof that Blum-Mansour Optimistic Hedge achieves no-swap-regret
- A corollary of our bound is that both Stoltz-Lugosi and Blum-Mansour Optimistic Hedge converge to correlated equilibrium in general games at a rate of $\tilde{O}(1/T)$.
- <https://arxiv.org/pdf/2111.06008.pdf>

Near-Optimal No-Regret Learning in General Games

NeurIPS 2021
Oral Presentation

- Established that Optimistic Hedge – a common variant of multiplicative-weights-updates with recency bias – attains $poly(\log T)$ regret in multi-player general-sum games
- Exponentially improved on the best known regret attainable by no-regret learners in general games
- Introduced many novel techniques for the analysis of the performance of regret minimization algorithms, including Fourier analysis
- A corollary of our bound is that Optimistic Hedge converges to coarse correlated equilibrium in general games at a rate of $\tilde{O}(1/T)$.
- <https://arxiv.org/pdf/2108.06924.pdf>

Multi-item Non-truthful Auctions Achieve Good Revenue

SICOMP 2022

(a.k.a. Simple, Credible, and Approximately Optimal Multi-item Auctions) EC 2020

- Established that first-price-type auctions can achieve a constant factor of the optimal revenue in the multi-item auction setting, resolving an open question
- Established the first credible and static multi-item auction that is approximately revenue optimal
- Obtained approximately-revenue-optimal multi-item mechanisms with fixed entry fees that are amenable to tuning via online learning techniques
- Proved a geometric lemma that enabled analysis of the utility of a first price auction, showing that welfare loss in a first price auction is at most 4 times the revenue of the posted price mechanism
- <https://arxiv.org/pdf/2002.06702.pdf>

Pattern Avoidance Over a Hypergraph

Electronic Journal
of Combinatorics

- Achieved a generalization of the Stanley-Wilf theorem, bounding the number of n -permutations avoiding a fixed sub-permutation at indices corresponding to the edges of a hypergraph
- Achieved bounds for both random and deterministic avoidance hypergraphs
- In deterministic case, devised a hypergraph formulation of pattern-avoidance, enabling the use of the hypergraph containers method
- <https://arxiv.org/pdf/1906.09659.pdf>

Szemerédi-Trotter: Polynomials and Incidences

Mathematical Reflections 2016

- For a set of reals A , proved a lower bound on the magnitude of either the set $A + A = \{x + y | x, y \in A\}$ or the set $f(A) + g(A) = \{f(x) + g(y) | x, y \in A\}$
- Initially a submission to the Intel Science Talent Search, published as an abridged version later
- One of seven math research papers awarded semifinalist in Intel STS
- <https://bit.ly/2rlmeUF>

WORK EXPERIENCE

EconCS Research Extern

Microsoft Research, New England 2020

- Goal of research to devise ways of compressing the information in matching problems via agent classifications while maintaining welfare guarantees for the derived matchings

- Position intended for graduate students; was employed during undergrad; only 1 opening

Tutor

MIT Math Learning Center 2019

- Help undergraduates with coursework from classes spanning MIT's mathematics curriculum
- Provide individual attention to students struggling on a specific topic, helping them gain intuition
- Offer generalized lectures to groups of students in the same classes

Quantitative Research and Trading Intern

Optiver US LLC 2018

- Created a machine learning model to predict volatility of the S&P500 following a day with abnormally high realized volatility
- Invented a strategy to adjust predicted volatility in response to a market input using unbiased historical data
- Traded S&P futures and options in a simulated environment using real market data
- Coded automated trading algorithms to compete against fellow interns in market making games

Teaching Assistant

Awesome Math Summer Program 2015-16

- 4 three-week sessions at Cornell U. (twice), UC Berkeley, and Univ. of Puget Sound, WA
- Worked alongside instructor explaining high-level olympiad math contest techniques from geometry, combinatorics, and number theory
- Worked with 50 students each camp (ages 12-16) helping them work through problems and gain problem solving intuition during classes and office hours

AWARDS AND HONORS

Honorable Mention, USA Junior Math Olympiad (#12 nationally; 1 point from winning) 2014
Honorable Mention, USA Junior Math Olympiad (#15 nationally; 2 points from winning) 2013
INTEL National Semi-finalist, Math Research 2016
#1 Team, PClassic Computer Programming Competition 2015
Harvard-MIT Math Tournament HMMT individual round (#19 internationally) 2015
AMC 10 perfect score (150) 3-time AMC Distinguished Honor Roll 2012-16
#1 Individual Scorer, NY State Math Tournament (Curt Boddie Award) 2014
#1 Individual Scorer, NYC Math Tournament (NYCIML), 3 years in a row 2014-16