COLT19

Conference on Learning Theory

June 25–28, 2019

Part of ACM FCRC
Phoenix, AZ, USA

Thanks to Our Sponsors

We are extremely grateful to our sponsors for being so very generous in supporting COLT, critically ensuring the conference's continued success, and making attendance more affordable for students through subsidized registration fees and travel awards which were provided to more than 35 student and postdocs authors. Thank you for all your support and generosity.

Platinum sponsors

Western Digital®

Gold sponsors

Silver sponsors

Foreword

Welcome to the 32nd Conference on Learning Theory!

This year we will enjoy 2 invited talks, 115 long talks (10 min plus a poster each). We thank the plenary speakers, Emma Brunskill and Moritz Hardt, for their participation and we also thank and congratulate the different authors for their contributions.

Alina Beygelzimer and Daniel Hsu
Program Chairs

Peter Grunwald and Yishay Mansour
Local Arrangements Chairs

Organizers

Program Chairs

Alina Beygelzimer (Yahoo! Research)
Daniel Hsu (Columbia University)

Local Arrangements Chairs

Peter Grunwald (Centrum Wiskunde & Informatica)
Yishay Mansour (Tel Aviv University & Google)

Program Committee

Jacob Abernethy (Georgia Tech)
Alekh Agarwal (Microsoft Research)
Shivani Agarwal (University of Pennsylvania)
Shipra Agrawal (Columbia University)
Alexandr Andoni (Columbia University)
Sivaraman Balakrishnan (Carnegie Mellon University)
Peter Bartlett (UC Berkeley)
Mihail Belkin (Ohio State University)
Shai Ben-David (University of Waterloo)
Guy Bresler (MIT)
Emma Brunskill (Stanford University)
Sebastien Bubeck (Microsoft Research)
Nicolo Cesa-Bianchi (University of Milan)
Moses Charikar (Stanford University)
Arnak Dalalyan (ENSAE)
Sanjoy Dasgupta (UC San Diego)
Ilias Diakonikolas (University of Southern California)
John Duchi (Stanford University)
Tim van Erven (Universiteit Leiden)
Vitaly Feldman (Google)
David Gamarnik (MIT)
Rong Ge (Duke University)
Claudio Gentile (Universita degli Studi dellInsubria)
Peter Grunwald (Centrum Wiskunde & Informatica)
Moritz Hardt (UC Berkeley)
Elad Hazan (Princeton University and Google Brain)
Prateek Jain (Microsoft Research)
Satyen Kale (Google)
Varun Kanade (University of Oxford)
Daniel Kane (UC San Diego)
Adam Klivans (UT Austin)
Wouter Koolen (Centrum Wiskunde & Informatica)
Tomer Koren (Google)
Akshay Krishnamurthy (Microsoft Research)
Edo Liberty (Amazon AI Labs)
Roi Livni (Tel Aviv University)
Gabor Lugosi (ICREA and Pompeu Fabra University)
Haipeng Luo (University of Southern California)
Raghu Meka (UC Los Angeles)
Ankur Moitra (MIT)
Andrea Montanari (Stanford University)
Praneeth Netrapalli (Microsoft Research)

Gergely Neu (Universitat Pompeu Fabra)
Behnam Neyshabur (New York University)
Francesco Orabona (Boston University)
Vianney Perchet (ENS Paris-Saclay and Criteo)
Luis Rademacher (UC Davis)
Maxim Raginsky (University of Illinois)
Sasha Rakhlin (MIT)
Aaditya Ramdas (Carnegie Mellon University)
Philippe Rigollet (MIT)
Robert Schapire (Microsoft Research)
Ohad Shamir (Weizman Institute)
Aaron Sidford (Stanford University)
Karthik Sridharan (Cornell University)
Xiaorui Sun (University of Illinois at Chicago)
Vasilis Syrgkanis (Microsoft Research)
Csaba Szepesvari (University of Alberta)
Matus Telgarsky (University of Illinois)
Ambuj Tewari (University of Michigan)
Gregory Valiant (Stanford University)
Michal Valko (INRIA)
Santosh Vempala (Georgia Tech)
Roman Vershynin (UC Irvine)
Aravindan Vijayaraghavan (Northwestern University)
Nisheeth Vishnoi (cole Polytechnique Fdrale de Lausanne)
Mengdi Wang (Princeton University)
Manfred Warmuth (UC Santa Cruz)
Rebecca Willett (University of Chicago)
Bob Williamson (Data61 and Australian National University)
David Woodruff (Carnegie Mellon University)

Publications Chair

Mark D. Reid (The Australian National University & NICTA)

Webmaster

Philippe Rigollet (MIT)

Women in ML Theory Lunch

Ruth Urner (York University)

Open Problems Chair

Jake Abernethy (Georgia Tech)

Sponsorship chairs

Satyen Kale (Google)
Robert Schapire (Microsoft Research)

Acknowledgements

Sponsor lists

Platinum:

AWS
American Express
Google
Jump trading
Western Digital

Gold:

Baidu
Microsoft
Oracle

Silver:

DE Shaw & Co
Mobileye
Two Sigma
Yahoo!

Best student paper sponsors

Machine Leaning Journal
Mark Fulk Foundation

WiML lunch sponsor

WiML.

Overview: (all talks in 231ABC – all poster sessions in 301 Foyer)

Monday June 24
5.30 PM Opening Reception (until 7.30 PM, room 103AB)

Tuesday June 25
8.15 AM Continental Breakfast (until 9:00 AM)
8.50 AM Session 1 (STOC Sister Session)
11.00 AM Coffee Break - FCRC Keynote (11:20 AM) / Lunch (12:30 PM)
2.00 PM Session 2 (Online Learning)
3.30 PM Afternoon Break
4.00 PM Session 3 (Testing and Distribution Learning)
6.00 PM Sponsor's Talks
6.00 PM Poster Session 1 (until 7.30 PM)

Wednesday June 26
8.15 AM Continental Breakfast (until 9:00 AM)
9.00 AM Session 4 (Inference and Estimation)
11.00 AM Coffee Break - FCRC Keynote (11:20)/ Lunch/WiML Lunch (12.30 PM)
2.00 PM Session 5 (Neural Networks)
3.30 PM Afternoon Break
4.00 PM Session 6 (Bandits)
6.00 PM Poster Session 2 (until 7.30 PM)
7.30 PM Conference Dinner (in 120A) (until 9.30 PM)

Thursday June 27
8.15 AM Continental Breakfast (until 9:00 AM)
9.00 AM Session 7 (Active Learning, Experimental Design, and Exploration)
10.00 AM Keynote Talk: Moritz Hardt
11.00 AM Coffee Break - FCRC Keynote (11:20 AM) / Lunch (12:30 PM)
1.30 PM Open Problems Session
2.10 PM Session 8 (Privacy and Robustness)
3.30 PM Afternoon Break
4.00 PM Session 9 (Optimization)
6.00 PM Poster Session 3 (until 7.30 PM)

Friday June 28
8.15 AM Continental Breakfast (until 9:00 AM)
9.00 AM Session 10 (Reinforcement Learning and Control)
10.00 AM Keynote Talk: Emma Brunskill
11.00 AM Coffee Break - FCRC Keynote (11:20 AM) / Lunch (12:30 PM)
1.30 PM Business Meeting
2.10 PM Session 11 (Sampling)
3.30 PM Afternoon Break
4.00 PM Session 12 (Statistical Learning Theory and Algorithms)
6.00 PM Impromptu talks

Monday June 24

Opening reception 5:30-7:30pm, 103AB

Tuesday June 25

Continental Breakfast (8:15 AM - 9:00AM)

Session 1 (STOC Sister Session)

8:50 AM	**Opening remarks**
9:00 AM	Jerry Li, Aleksandar Nikolov, Ilya Razenshteyn, Erik Waingarten **On Mean Estimation for General Norms with Statistical Queries**
9:10 AM	Samuel B. Hopkins, Jerry Li **How Hard is Robust Mean Estimation?**
9:20 AM	Yeshwanth Cherapanamjeri, Nicolas Flammarion, Peter Bartlett **Fast Mean Estimation with Sub-Gaussian Rates**
9:30 AM	Daniel Alabi, Adam Tauman Kalai, Katrina Ligett, Cameron Musco, Christos Tzamos, Ellen Vitercik **Learning to Prune: Speeding up Repeated Computations**
9:40 AM	Michal Derezinski **Fast determinantal point processes via distortion-free intermediate sampling**
9:50 AM	Yair Carmon, John C. Duchi, Aaron Sidford, Kevin Tian **A Rank-1 Sketch for Matrix Multiplicative Weights**
10:00 AM	Zohar Karnin, Edo Liberty **Discrepancy, Coresets, and Sketches in Machine Learning**
10:10 AM	Matthew Brennan, Guy Bresler **Optimal Average-Case Reductions to Sparse PCA: From Weak Assumptions to Strong Hardness**
10:20 AM	Matthew Brennan, Guy Bresler, Wasim Huleihel **Universality of Computational Lower Bounds for Submatrix Detection**
10:30 AM	Jan Hązła Jadbabaie, Elchanan Mossel, M. Amin Rahimian **Reasoning in Bayesian Opinion Exchange Networks Is PSPACE-Hard**
10:40 AM	Dylan Foster, Andrej Risteski **Sum-of-squares meets square loss: Fast rates for agnostic tensor completion**
10:50 AM	Samuel B. Hopkins, Tselil Schramm, Jonathan Shi **A Robust Spectral Algorithm for Overcomplete Tensor Decomposition**

Coffee Break (11:00 AM) / FCRC Keynote (11:20 AM) / Lunch (12:30 PM)

Session 2 (Online Learning)

2:00 PM — Dan Garber
On the Regret Minimization of Nonconvex Online Gradient Ascent for Online PCA

2:10 PM — Naman Agarwal, Alon Gonen, Elad Hazan
Learning in Non-convex Games with an Optimization Oracle

2:20 PM — Ashok Cutkosky
Combining Online Learning Guarantees

2:30 PM — Ashok Cutkosky
Artificial Constraints and Hints for Unbounded Online Learning

2:40 PM — Zakaria Mhammedi, Wouter M. Koolen, Tim van Erven
Lipschitz Adaptivity with Multiple Learning Rates in Online Learning

2:50 PM — Christian Coester, James R. Lee
Pure Entropic Regularization for MTS

3:00 PM — Yun Kuen Cheung, Georgios Piliouras
Vortices Instead of Equilibria in MinMax Optimization: Chaos and Butterfly Effects of Online Learning in Zero-Sum Games

3:10 PM — Mingda Qiao, Gregory Valiant
A Theory of Selective Prediction

3:20 PM — Vaggos Chatziafratis, Tim Roughgarden, Joshua R. Wang
On the Computational Power of Online Gradient Descent

Afternoon Break (3:30 PM)

Session 3 (Testing and Distribution Learning)

4:00 PM — Damian Straszak, Nisheeth K. Vishnoi
Maximum Entropy Distributions: Bit Complexity and Stability

4:10 PM — Jonathan Weed, Quentin Berthet
Estimation of smooth densities in Wasserstein distance

4:20 PM — Olivier Bousquet, Daniel Kane, Shay Moran
The Optimal Approximation Factor in Density Estimation

4:30 PM — Ilias Diakonikolas, Themis Gouleakis, Daniel M. Kane, Sankeerth Rao
Communication and Memory Efficient Testing of Discrete Distributions

4:40 PM — Jayadev Acharya, Clément L. Canonne, Himanshu Tyagi
Inference under Local Constraints: Lower Bounds from Chi-Square Contractions

4:50 PM — Maryam Aliakbarpour, Themis Gouleakis, John Peebles, Ronitt Rubinfeld, Anak Yodpinyanee
Towards Testing Monotonicity of Distributions Over General Posets

5:00 PM — Maryam Aliakbarpour, Ravi Kumar, Ronitt Rubinfeld
Testing Mixtures of Discrete Distributions

5:10 PM	Ilias Diakonikolas, Daniel M. Kane, John Peebles **Testing Identity of Multidimensional Histograms**
5:20 PM	Meimei Liu, Zuofeng Shang, Guang Cheng **Sharp Theoretical Analysis for Nonparametric Testing under Random Projection**
5:30 PM	Ivona Bezakova, Antonio Blanca, Zongchen Chen, Daniel Stefankovic, Eric Vigoda **Lower bounds for testing graphical models: colorings and antiferromagnetic Ising models**
5:40 PM	Yeshwanth Cherapanamjeri, Peter Bartlett **Testing Markov Chains Without Hitting**
5:50 PM	Anindya De, Elchanan Mossel, Joe Neeman **Is your function low dimensional?**

Word from Our Sponsors

6:00 PM **Sponsors' Talks**

6:00 PM **Poster Session 1**

1. On Mean Estimation for General Norms with Statistical Queries
2. How Hard is Robust Mean Estimation?
3. Fast Mean Estimation with Sub-Gaussian Rates
4. Learning to Prune: Speeding up Repeated Computations
5. Fast determinantal point processes via distortion-free intermediate sampling
6. A Rank-1 Sketch for Matrix Multiplicative Weights
7. Discrepancy, Coresets, and Sketches in Machine Learning
8. Optimal Average-Case Reductions to Sparse PCA: From Weak Assumptions to Strong Hardness
9. Universality of Computational Lower Bounds for Submatrix Detection
10. Reasoning in Bayesian Opinion Exchange Networks Is PSPACE-Hard
11. Sum-of-squares meets square loss: Fast rates for agnostic tensor completion
12. A Robust Spectral Algorithm for Overcomplete Tensor Decomposition
13. On the Regret Minimization of Nonconvex Online Gradient Ascent for Online PCA
14. Learning in Non-convex Games with an Optimization Oracle
15. Combining Online Learning Guarantees
16. Artificial Constraints and Hints for Unbounded Online Learning
17. Lipschitz Adaptivity with Multiple Learning Rates in Online Learning
18. Pure Entropic Regularization for MTS
19. Vortices Instead of Equilibria in MinMax Optimization: Chaos and Butterfly Effects of Online Learning in Zero-Sum Games
20. A Theory of Selective Prediction
21. On the Computational Power of Online Gradient Descent
22. Maximum Entropy Distributions: Bit Complexity and Stability

23. Estimation of smooth densities in Wasserstein distance
24. The Optimal Approximation Factor in Density Estimation
25. Communication and Memory Efficient Testing of Discrete Distributions
26. Inference under Local Constraints: Lower Bounds from Chi-Square Contractions
27. Towards Testing Monotonicity of Distributions Over General Posets
28. Testing Mixtures of Discrete Distributions
29. Testing Identity of Multidimensional Histograms
30. Sharp Theoretical Analysis for Nonparametric Testing under Random Projection
31. Lower bounds for testing graphical models: colorings and antiferromagnetic Ising models
32. Testing Markov Chains Without Hitting
33. Is your function low dimensional?
34. The Gap Between Model-Based and Model-Free Methods on the Linear Quadratic Regulator: An Asymptotic Viewpoint
35. Finite-Time Error Bounds For Linear Stochastic Approximation and TD Learning
36. Model-based RL in Contextual Decision Processes: PAC bounds and Exponential Improvements over Model-free Approaches
37. Non-asymptotic Analysis of Biased Stochastic Approximation Scheme
38. Learning Linear Dynamical Systems with Semi-Parametric Least Squares

Wednesday June 26

Continental Breakfast (8:15 AM -- 9:00AM)

Session 4 (Inference and Estimation)

9:00 AM	Laurent Massoulié, Ludovic Stephan, Don Towsley **Planting trees in graphs, and finding them back**
9:10 AM	Sami Davies, Miklos Racz, Cyrus Rashtchian **Reconstructing Trees from Traces**
9:20 AM	Ludovic Stephan, Laurent Massoulié **Robustness of spectral methods for community detection**
9:30 AM	Yingjie Fei, Yudong Chen **Achieving the Bayes Error Rate in Stochastic Block Model by SDP, Robustly**
9:40 AM	Robert Busa-Fekete, Dimitris Fotakis, Balazs Szorenyi, Manolis Zampetakis **Optimal Learning for Mallows Block Model**
9:50 AM	Vishesh Jain, Frederic Koehler, Jingbo Liu, Elchanan Mossel **Accuracy-Memory Tradeoffs and Phase Transitions in Belief Propagation**
10:00 AM	Surbhi Goel, Daniel M. Kane, Adam R. Klivans **Learning Ising Models with Independent Failures**
10:10 AM	Victor-Emmanuel Brunel **Learning rates for Gaussian mixtures under group invariance**

10:20 AM — Jeongyeol Kwon, Wei Qian, Constantine Caramanis, Yudong Chen, Damek Davis
Global Convergence of the EM Algorithm for Mixtures of Two Component Linear Regression

10:30 AM — Arun Sai Suggala, Kush Bhatia, Pradeep Ravikumar, Prateek Jain
Adaptive Hard Thresholding for Near-optimal Consistent Robust Regression

10:40 AM — Constantinos Daskalakis, Themis Gouleakis, Christos Tzamos, Emmanouil Zampetakis
Computationally and Statistically Efficient Truncated Regression

10:50 AM — Galen Reeves, Jiaming Xu, Ilias Zadik
The All-or-Nothing Phenomenon in Sparse Linear Regression

Coffee Break (11:00 AM) / FCRC Keynote (11:20 AM) / WiML Lunch (12:30)/ Lunch (12:30 PM)

Session 5 (Neural Networks)

2:00 PM — Ziwei Ji, Matus Telgarsky
The implicit bias of gradient descent on nonseparable data

2:10 PM — Pedro Savarese, Itay Evron, Daniel Soudry, Nathan Srebro
How do infinite width bounded norm networks look in function space?

2:20 PM — Song Mei, Theodor Misiakiewicz, Andrea Montanari
Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit

2:30 PM — Surbhi Goel, Adam R. Klivans
Learning Neural Networks with Two Nonlinear Layers in Polynomial Time

2:40 PM — Ainesh Bakshi, Rajesh Jayaram, David P. Woodruff
Learning Two Layer Rectified Neural Networks in Polynomial Time

2:50 PM — Santosh Vempala, John Wilmes
Gradient Descent for One-Hidden-Layer Neural Networks: Polynomial Convergence and SQ Lower Bounds

3:00 PM — Ohad Shamir
Exponential Convergence Time of Gradient Descent for One-Dimensional Deep Linear Neural Networks

3:10 PM — Itay Safran, Ronen Eldan, Ohad Shamir
Depth Separations in Neural Networks: What is Actually Being Separated?

3:20 PM — Samet Oymak
Stochastic Gradient Descent Learns State Equations with Nonlinear Activations

Afternoon Break (3:30 PM)

Session 6 (Bandits)

4:00 PM	Tor Lattimore, Csaba Szepesvari **An Information-Theoretic Approach to Minimax Regret in Partial Monitoring**
4:10 PM	Sandeep Juneja, Subhashini Krishnasamy **Sample complexity of partition identification using multi-armed bandits**
4:20 PM	Nadav Merlis, Shie Mannor **Batch-Size Independent Regret Bounds for the Combinatorial Multi-Armed Bandit Problem**
4:30 PM	Mark Braverman, Jieming Mao, Jon Schneider, S. Matthew Weinberg **Multi-armed Bandit Problems with Strategic Arms**
4:40 PM	Anupam Gupta, Tomer Koren, Kunal Talwar **Better Algorithms for Stochastic Bandits with Adversarial Corruptions**
4:50 PM	Shi Dong, Tengyu Ma, Benjamin Van Roy **On the Performance of Thompson Sampling on Logistic Bandits**
5:00 PM	Yingkai Li, Yining Wang, Yuan Zhou **Nearly Minimax-Optimal Regret for Linearly Parameterized Bandits**
5:10 PM	Daniele Calandriello, Luigi Carratino, Alessandro Lazaric, Michal Valko, Lorenzo Rosasco **Gaussian Process Optimization with Adaptive Sketching: Scalable and No Regret**
5:20 PM	Sébastien Bubeck, Yuanzhi Li, Haipeng Luo, Chen-Yu Wei **Improved Path-length Regret Bounds for Bandits**
5:30 PM	Akshay Krishnamurthy, John Langford, Aleksandrs Slivkins, Chicheng Zhang **Contextual Bandits with Continuous Actions: Smoothing, Zooming, and Adapting**
5:40 PM	Yifang Chen, Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei **A New Algorithm for Non-stationary Contextual Bandits: Efficient, Optimal, and Parameter-free** Peter Auer, Pratik Gajane, and Ronald Ortner **Adaptively Tracking the Best Bandit Arm with an Unknown Number of Distribution Changes**
5:50 PM	Wojciech Kotlowski, Gergely Neu **Bandit Principal Component Analysis**

6:00 PM **Poster session 2**

1. Planting trees in graphs, and finding them back
2. Reconstructing Trees from Traces
3. Robustness of spectral methods for community detection
4. Achieving the Bayes Error Rate in Stochastic Block Model by SDP, Robustly
5. Optimal Learning for Mallows Block Model
6. Accuracy-Memory Tradeoffs and Phase Transitions in Belief Propagation

7. Learning Ising Models with Independent Failures
8. Learning rates for Gaussian mixtures under group invariance
9. Global Convergence of the EM Algorithm for Mixtures of Two Component Linear Regression
10. Adaptive Hard Thresholding for Near-optimal Consistent Robust Regression
11. Computationally and Statistically Efficient Truncated Regression
12. The All-or-Nothing Phenomenon in Sparse Linear Regression
13. The implicit bias of gradient descent on nonseparable data
14. How do infinite width bounded norm networks look in function space?
15. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit
16. Learning Neural Networks with Two Nonlinear Layers in Polynomial Time
17. Learning Two Layer Rectified Neural Networks in Polynomial Time
18. Gradient Descent for One-Hidden-Layer Neural Networks: Polynomial Convergence and SQ Lower Bounds
19. Exponential Convergence Time of Gradient Descent for One-Dimensional Deep Linear Neural Networks
20. Depth Separations in Neural Networks: What is Actually Being Separated?
21. Stochastic Gradient Descent Learns State Equations with Nonlinear Activations
22. An Information-Theoretic Approach to Minimax Regret in Partial Monitoring
23. Sample complexity of partition identification using multi-armed bandits
24. Batch-Size Independent Regret Bounds for the Combinatorial Multi-Armed Bandit Problem
25. Multi-armed Bandit Problems with Strategic Arms
26. Better Algorithms for Stochastic Bandits with Adversarial Corruptions
27. On the Performance of Thompson Sampling on Logistic Bandits
28. Nearly Minimax-Optimal Regret for Linearly Parameterized Bandits
29. Gaussian Process Optimization with Adaptive Sketching: Scalable and No Regre
30. Improved Path-length Regret Bounds for Bandits
31. Contextual Bandits with Continuous Actions: Smoothing, Zooming, and Adapting
32. A New Algorithm for Non-stationary Contextual Bandits: Efficient, Optimal, and Parameter-free
33. Adaptively Tracking the Best Bandit Arm with an Unknown Number of Distribution Changes
34. Bandit Principal Component Analysis
35. Estimating the Mixing Time of Ergodic Markov Chains
36. Distribution-Dependent Analysis of Gibbs-ERM Principle
37. Theoretical guarantees for sampling and inference in generative models with latent diffusions
38. Sampling and Optimization on Convex Sets in Riemannian Manifolds of Non-Negative Curvature
39. Nonconvex sampling with the Metropolis-adjusted Langevin algorithm
40. Normal Approximation for Stochastic Gradient Descent via Non-Asymptotic Rates of Martingale CLT

Conference dinner in 120A (7.30 PM – 9.30 PM)

Thursday June 27

Continental Breakfast (8:15 AM - 9:00AM)

Session 7 (Active Learning, Experimental Design, and Exploration)

9:00 AM Tongyi Cao, Akshay Krishnamurthy
Disagreement-Based Combinatorial Pure Exploration: Sample Complexity Bounds and an Efficient Algorithm

9:10 AM Xue Chen, Eric Price
Active Regression via Linear-Sample Sparsification

9:20 AM Vivek Madan, Mohit Singh, Uthaipon Tantipongpipat, Weijun Xie
Combinatorial Algorithms for Optimal Design

9:30 AM Michal Derezinski, Kenneth L. Clarkson, Michael W. Mahoney, Manfred K. Warmuth
Minimax experimental design: Bridging the gap between statistical and worst-case approaches to least squares regression

9:40 AM Mark Braverman, Jieming Mao, Yuval Peres
Sorted Top-k in Rounds

Keynote Talk 1

10:00 Moritz Hardt, TBD

Coffee Break (11:00 AM) / FCRC Keynote (11:20 AM) / Lunch (12:30 PM)

Open Problems Session

1:30 PM **Open problems**

Session 8 (Privacy and Robustness)

2:10 PM Akshay Degwekar, Vinod Vaikuntanathan
Computational Limitations in Robust Classification and Win-Win Results

2:20 PM Omar Montasser, Steve Hanneke, Nathan Srebro
VC Classes are Adversarially Robustly Learnable, but Only Improperly

2:30 PM Yu Cheng, Ilias Diakonikolas, Rong Ge, David P. Woodruff
Faster Algorithms for High-Dimensional Robust Covariance Estimation

2:40 PM Dylan Foster, Vasilis Syrgkanis
Statistical Learning with a Nuisance Component

2:50 PM Gautam Kamath, Jerry Li, Vikrant Singhal, Jonathan Ullman
Privately Learning High-Dimensional Distributions

3:00 PM John Duchi, Ryan Rogers
Lower Bounds for Locally Private Estimation via Communication Complexity

3:10 PM Amos Beimel, Shay Moran, Kobbi NIssim, Uri Stemmer
Private Center Points and Learning of Halfspaces

3:20 PM Kwang-Sung Jun, Francesco Orabona
Parameter-free Online Convex Optimization with Sub-Exponential Noise

Afternoon Break (3:30 PM)

Session 9 (Optimization)

4:00 PM Nicholas J. A. Harvey, Christopher Liaw, Yaniv Plan, Sikander Randhawa
Tight analyses for non smooth stochastic gradient descent

4:10 PM Prateek Jain, Dheeraj Nagaraj, Praneeth Netrapalli
Making the Last Iterate of SGD Information Theoretically Optimal

4:20 PM Lijun Zhang, Zhi-Hua Zhou
Stochastic Approximation of Smooth and Strongly Convex Functions: Beyond the $O(1/T)$ Convergence Rate

Bo Jiang, Haoyue Wang, Shuzhong Zhang
An Optimal High-Order Tensor Method for Convex Optimization

4:30 PM Sebastien Bubeck, Qijia Jiang, Yin Tat Lee, Yuanzhi Li, Aaron Sidford
Near-optimal method for highly smooth convex optimization

Alexander Gasnikov, Pavel Dvurechensky, Eduard Gorbunov, Evgeniya Vorontsova, Daniil Selikhanovych, Cesar A. Uribe
Optimal Tensor Methods in Smooth Convex and Uniformly Convex Optimization

4:40 PM Adrien Taylor, Francis Bach
Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions

4:50 PM Ulysse Marteau-Ferey, Dmitrii M. Ostrovskii, Francis Bach, Alessandro Rudi
Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance

5:00 PM Yin Tat Lee, Zhao Song, Qiuyi Zhang
Solving Empirical Risk Minimization in the Current Matrix Multiplication Time

5:10 PM Jelena Diakonikolas, Cristóbal Guzmán
Lower Bounds for Parallel and Randomized Convex Optimization

5:20 PM Dylan Foster, Ayush Sekhari, Ohad Shamir, Nathan Srebro, Karthik Sridharan, Blake Woodworth
The Complexity of Making the Gradient Small in Stochastic Convex Optimization

5:30 PM Francis Bach, Kfir Y. Levy
A Universal Algorithm for Variational Inequalities Adaptive to Smoothness and Noise

5:40 PM Rong Ge, Zhize Li, Weiyao Wang, Xiang Wang
Stabilized SVRG: Simple Variance Reduction for Nonconvex Optimization

5:50 PM Cong Fang, Zhouchen Lin, Tong Zhang
Sharp Analysis for Nonconvex SGD Escaping from Saddle Points

6:00 PM **Poster session 3**

1. The Relative Complexity of Maximum Likelihood Estimation, MAP Estimation, and Sampling
2. Disagreement-Based Combinatorial Pure Exploration: Sample Complexity Bounds and an Efficient Algorithm
3. Active Regression via Linear-Sample Sparsification
4. Combinatorial Algorithms for Optimal Design
5. Minimax experimental design: Bridging the gap between statistical and worst-case approaches to least squares regression
6. Sorted Top-k in Rounds
7. Computational Limitations in Robust Classification and Win-Win Results
8. VC Classes are Adversarially Robustly Learnable, but Only Improperly
9. Faster Algorithms for High-Dimensional Robust Covariance Estimation
10. Statistical Learning with a Nuisance Component
11. Privately Learning High-Dimensional Distributions
12. Lower Bounds for Locally Private Estimation via Communication Complexity
13. Private Center Points and Learning of Halfspaces
14. Parameter-free Online Convex Optimization with Sub-Exponential Noise
15. Tight analyses for non smooth stochastic gradient descent
16. Making the Last Iterate of SGD Information Theoretically Optimal
17. Stochastic Approximation of Smooth and Strongly Convex Functions: Beyond the $O(1/T)$ Convergence Rate
18. An Optimal High-Order Tensor Method for Convex Optimization
19. Near-optimal method for highly smooth convex optimization
20. Optimal Tensor Methods in Smooth Convex and Uniformly Convex Optimization
21. Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions
22. Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance
23. Solving Empirical Risk Minimization in the Current Matrix Multiplication Time
24. Lower Bounds for Parallel and Randomized Convex Optimization
25. The Complexity of Making the Gradient Small in Stochastic Convex Optimization
26. A Universal Algorithm for Variational Inequalities Adaptive to Smoothness and Noise
27. Stabilized SVRG: Simple Variance Reduction for Nonconvex Optimization
28. Sharp Analysis for Nonconvex SGD Escaping from Saddle Points
29. Uniform concentration and symmetrization for weak interactions
30. Learning from Weakly Dependent Data under Dobrushin's Condition
31. High probability generalization bounds for uniformly stable algorithms with nearly optimal rate
32. When can unlabeled data improve the learning rate?
33. Classification with unknown class conditional label noise on non-compact feature spaces

34. Consistency of Interpolation with Laplace Kernels is a High-Dimensional Phenomenon
35. On Communication Complexity of Classification Problems
36. Space lower bounds for linear prediction
37. Affine Invariant Covariance Estimation for Heavy-Tailed Distributions
38. Approximate Guarantees for Dictionary Learning
39. Sample-Optimal Low-Rank Approximation of Distance Matrices
40. A near-optimal algorithm for approximating the John Ellipsoid

Friday June 28

Continental Breakfast (8:15 AM - 9:00AM)

Session 10 (Reinforcement Learning and Control)

9:00 AM Stephen Tu, Benjamin Recht
The Gap Between Model-Based and Model-Free Methods on the Linear Quadratic Regulator: An Asymptotic Viewpoint

9:10 AM R. Srikant, Lei Ying
Finite-Time Error Bounds For Linear Stochastic Approximation and TD Learning

9:20 AM Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford
Model-based RL in Contextual Decision Processes: PAC bounds and Exponential Improvements over Model-free Approaches

9:30 AM Belhal Karimi, Blazej Miasojedow, Eric Moulines, Hoi-To Wai
Non-asymptotic Analysis of Biased Stochastic Approximation Scheme

9:40 AM Max Simchowitz, Ross Boczar, Benjamin Recht
Learning Linear Dynamical Systems with Semi-Parametric Least Squares

Keynote Talk 2

10:00 AM Emma Brunskill
Towards Efficient Effective Reinforcement Learning Algorithms That Interact With People

Coffee Break (11:00 AM) / FCRC Keynote (11:20 AM) / Lunch (12:30 PM)

1:30 PM **Business meeting**

Session 11 (Sampling)

2:20 PM Geoffrey Wolfer, Aryeh Kontorovich
Estimating the Mixing Time of Ergodic Markov Chains

2:30 PM Ilja Kuzborskij, Nicolò Cesa-Bianchi, Csaba Szepesvari
Distribution-Dependent Analysis of Gibbs-ERM Principle

2:40 PM Belinda Tzen, Maxim Raginsky
Theoretical guarantees for sampling and inference in generative models with latent diffusions

2:50 PM Navin Goyal, Abhishek Shetty
Sampling and Optimization on Convex Sets in Riemannian Manifolds of Non-Negative Curvature

3:00 PM Oren Mangoubi, Nisheeth K. Vishnoi
Nonconvex sampling with the Metropolis-adjusted Langevin algorithm

3:10 PM Andreas Anastasiou, Krishnakumar Balasubramanian, Murat Erdogdu
Normal Approximation for Stochastic Gradient Descent via Non-Asymptotic Rates of Martingale CLT

3:20 PM Christopher Tosh, Sanjoy Dasgupta
The Relative Complexity of Maximum Likelihood Estimation, MAP Estimation, and Sampling

Afternoon Break (3:30 PM)

Session 12 (Statistical Learning Theory and Algorithms)

4:00 PM Andreas Maurer, Massimiliano Pontil
Uniform concentration and symmetrization for weak interactions

4:10 PM Yuval Dagan, Constantinos Daskalakis, Nishanth Dikkala, Siddhartha Jayanti
Learning from Weakly Dependent Data under Dobrushin's Condition

4:20 PM Vitaly Feldman, Jan Vondrak
High probability generalization bounds for uniformly stable algorithms with nearly optimal rate

4:30 PM Christina Göpfert, Shai Ben-David, Olivier Bousquet, Sylvain Gelly, Ilya Tolstikhin, Ruth Urner
When can unlabeled data improve the learning rate?

4:40 PM Henry Reeve, Ata Kaban
Classification with unknown class conditional label noise on non-compact feature spaces

4:50 PM Alexander Rakhlin, Xiyu Zhai
Consistency of Interpolation with Laplace Kernels is a High-Dimensional Phenomenon

5:00 PM Daniel Kane, Roi Livni, Shay Moran, Amir Yehudayoff
On Communication Complexity of Classification Problems

5:10 PM Yuval Dagan, Gil Kur, Ohad Shamir
Space lower bounds for linear prediction

5:20 PM Dmitrii M. Ostrovskii, Alessandro Rudi
Affine Invariant Covariance Estimation for Heavy-Tailed Distributions

5:30 PM Aditya Bhaskara, Wai Ming Tai
Approximate Guarantees for Dictionary Learning

5:40 PM Piotr Indyk, Ali Vakilian, Tal Wagner, David Woodruff
Sample-Optimal Low-Rank Approximation of Distance Matrices

5:50 PM Michael B. Cohen, Ben Cousins, Yin Tat Lee, Xin Yang
A near-optimal algorithm for approximating the John Ellipsoid

6:00 PM **Impromptu talks**