Open Problem 1

Mon, Aug 16, 09:30 – 09:40
MAIN TRACK

- **Open Problem: Can Single-Shuffle SGD be Better than Reshuffling SGD and GD?**
  Chulhee Yun (MIT) ; Suvrit Sra (Massachusetts Institute of Technology, USA) ;
  Ali Jadbabaie (Massachusetts Institute of Technology)

Minimally Supervised Learning (A)

Mon, Aug 16, 09:50 – 10:15
TRACK 1

- **Exponential savings in agnostic active learning through abstention**
  Nikita Puchkin (National Research University Higher School of Economics) ; Nikita Zhivotovskiy (ETH)
- **Learning sparse mixtures of permutations from noisy information**
  Anindya De (-) ; Ryan O'Donnell ; Rocco Servedio (Columbia University)
- **The Bethe and Sinkhorn Permanents of Low Rank Matrices and Implications for Profile Maximum Likelihood**
  Nima Anari (Stanford) ; Moses Charikar (Stanford University, California) ;
  Kirankumar Shiragur (Stanford University) ; Aaron Sidford (Stanford)
- **Bounded Memory Active Learning through Enriched Queries**
  Max Hopkins (University of California San Diego) ; Daniel Kane (University of California San Diego) ;
  Shachar Lovett (University of California San Diego) ; Michal Moshkovitz (UC San Diego)
- **Adaptivity in Adaptive Submodularity**
  Hossein Esfandiari (Google Research) ; Amin Karbasi (Yale) ; Vahab Mirrokni (Google)

Optimization (A)

Mon, Aug 16, 09:50 – 10:10
TRACK 2

- **Optimizing Optimizers: Regret-optimal gradient descent algorithms**
  Philippe Casgrain (ETH Zürich) ; Anastasis Kratsios (ETH Zürich)
- **Frank-Wolfe with a Nearest Extreme Point Oracle**
  Dan Garber (Technion) ; Noam Wolf (Technion)
- **(Nearly) Dimension Independent Private ERM with AdaGrad Rates\\ via Publicly Estimated Subspaces**
  Peter Kairouz (Google) ; Monica Ribero Diaz (University of Texas at Austin) ; Keith Rush (Google Research) ;
  Abhradeep Thakurta (Google)
- **Thinking Inside the Ball: Near-Optimal Minimization of the Maximal Loss**
  Yair Carmon (Tel Aviv University) ; Arun Jambulapati (Stanford University) ; Yujia Jin (Stanford University) ;
  Aaron Sidford (Stanford)

Minimally Supervised Learning (B)

Mon, Aug 16, 10:20 – 10:40
TRACK 1

- **Learning and testing junta distributions with sub cube conditioning**
  Xi Chen (Columbia University) ; Rajesh Jayaram (Carnegie Mellon University) ; Amit Levi (University of Waterloo) ; Erik Waingarten
- **Outlier-Robust Learning of Ising Models Under Dobrushin's Condition**
  Ilias Diakonikolas (UW Madison) ; Daniel M. Kane (UCSD) ; Alistair Stewart (Web 3 Foundation) ; Yuxin Sun (University of Wisconsin - Madison)
- **Breaking The Dimension Dependence in Sparse Distribution Estimation under Communication Constraints**
  Wei-Ning Chen (Stanford University) ; Peter Kairouz (Google) ; Ayfer Ozgur (Stanford University)
- **The Sample Complexity of Robust Covariance Testing**
  Ilias Diakonikolas (UW Madison) ; Daniel Kane (UCSD)

Optimization (B)

Mon, Aug 16, 10:20 – 10:40
TRACK 2

- **Projected Stochastic Gradient Langevin Algorithms for Constrained Sampling and Non-Convex Learning**
  Andrew Lamperski (University of Minnesota)
- **The Last-Iterate Convergence Rate of Optimistic Mirror Descent in Stochastic Variational Inequalities**
  Waïss Azizian (Ecole Normale Supérieure de Paris) ; Franck Iutzeler (Univ. Grenoble Alpes) ; Jérôme Malick (CNRS) ; Panayotis Mertikopoulos (CNRS and Criteo AI Lab)
- **The Min-Max Complexity of Distributed Stochastic Convex Optimization with Intermittent Communication**
  Blake Woodworth (TTI-Chicago) ; Brian Bullins (TTI Chicago) ; Ohad Shamir (Weizmann Institute of Science) ; Nathan Srebro (Toyota Technical Institute of Chicago)
- **A Local Convergence Theory for Mildly Over-Parameterized Two-Layer Neural Network**
  Mo Zhou (Duke University) ; Rong Ge (Duke University) ; Chi Jin (Princeton University)

Online Learning, Game Theory 1 (A)

Mon, Aug 16, 11:10 – 11:30
TRACK 1

- **Lazy OCO: Online Convex Optimization on a Switching Budget**
  Uri Sherman (Tel Aviv University) ; Tomer Koren (Tel Aviv University and Google)
- **Deterministic Finite-Memory Bias Estimation**
  Tomer Berg (TAU) ; Or Ordentlich (HUJI) ; Ofer Shayevitz (Tel-Aviv University)
- **Online Learning from Optimal Actions**
  Yuri Fonseca (Columbia University)
- **Online Learning with Simple Predictors and a Combinatorial Characterization of Minimax in 0/1 Games**
  Steve Hanneke (Toyota Technological Institute at Chicago) ; Roi Livni (Tel Aviv University) ; Shay Moran (Technion)

Approximation and Complexity

Mon, Aug 16, 11:10 – 11:35
TRACK 2

- **Reconstructing weighted voting schemes from partial information about their power indices**
  Emmanouil Vlatakis-Gkaragkounis (Columbia University) ; Huck Bennett (University of Michigan) ; Anindya De (-) ; Rocco Servedio (Columbia University)
- **Weak learning convex sets under normal distributions**
  Anindya De (-) ; Rocco Servedio (Columbia University)
- **Statistical Query Algorithms and Low Degree Tests Are Almost Equivalent**
  Matthew Brennan (MIT) ; Guy Bresler (MIT) ; Sam Hopkins ; Jerry Li (Microsoft) ; Tselil Schramm (Stanford University)
- **A Dimension-free Computational Upper-bound for Smooth Optimal Transport Estimation**
  Adrien Vacher (LIGM) ; Boris Muzellec (INRIA, Ecole Normale Superieure) ; Alessandro Rudi (INRIA, Ecole Normale Superieure) ; Francis Bach (INRIA - Ecole Normale Supérieure) ; Francois-Xavier Vialard (University Paris-Est)
- **Spectral Planting and the Hardness of Refuting Cuts, Colorability, and Communities in Random Graphs**
  Afonso Bandeira (ETH) ; Jess Banks (UC Berkeley) ; Dmitriy Kunisky (New York University) ; Christopher Moore (Santa Fe) ; Alex Wein (NYU)

Online Learning, Game Theory 1 (B)

Mon, Aug 16, 11:35 – 11:55
TRACK 1

- **Online Markov Decision Processes with Aggregate Bandit Feedback**
  Alon Cohen (Technion and Google Inc.) ; Haim Kaplan (Tel Aviv University and Google) ; Tomer Koren (Tel Aviv University and Google) ; Yishay Mansour (Tel Aviv University)
- **Optimal Dynamic Regret in Exp-Concave Online Learning**
  Dheeraj Baby (University of California Santa Barbara) ; Yu-Xiang Wang (UC Santa Barbara)
- **Exponential Weights Algorithms for Selective Learning**
  Mingda Qiao (Stanford University) ; Gregory Valiant (Stanford University)
- **Impossible Tuning Made Possible: A New Expert Algorithm and Its Applications**
  Liyu Chen (USC) ; Haipeng Luo (USC) ; Chen-Yu Wei (University of Southern California)

Randomized Linear Algebra

Mon, Aug 16, 11:40 – 11:55
TRACK 2

- **Sparse sketches with small inversion bias**
  Michal Derezinski (UC Berkeley) ; Zhenyu Liao (University of California, Berkeley) ; Edgar Dobriban (University of Pennsylvania) ; Michael Mahoney ("University of California, Berkeley")
- **Exponentially Improved Dimensionality Reduction for l1: Subspace Embeddings and Independence Testing**

Taisuke Yasuda (Carnegie Mellon University) ; David Woodruff (Carnegie Mellon University) ; Yi Li (Nanyang Technological University)

- **Average-Case Communication Complexity of Statistical Problems**
  Cyrus Rashtchian (UCSD) ; David Woodruff (Carnegie Mellon University) ; Peng Ye (Tsinghua University) ; Hanlin Zhu (Institute for Interdisciplinary Information Sciences, Tsinghua University)

Keynote: Eva Tardos (Cornell University)

Mon, Aug 16, 12:05 – 13:00
MAIN TRACK

Open Problem 2

Tue, Aug 17, 09:35 – 09:45
MAIN TRACK

- **Open Problem: Is There an Online Learning Algorithm That Learns Whenever Online Learning Is Possible?**
  Steve Hanneke (Toyota Technological Institute at Chicago)

Neural Networks/Deep Learning (A)

Tue, Aug 17, 09:50 – 10:15
TRACK 1

- **Modeling from Features: a Mean-field Framework for Over-parameterized Deep Neural Networks**
  Cong Fang (University of Pennsylvania) ; Jason Lee (Princeton) ; Pengkun Yang (Tsinghua University) ; Tong Zhang (The Hong Kong University of Science and Technology)
- **Provable Memorization via Deep Neural Networks using Sub-linear Parameters**
  Sejun Park (KAIST) ; Jaeho Lee (KAIST) ; Chulhee Yun (MIT) ; Jinwoo Shin (KAIST)
- **The Effects of Mild Over-parameterization on the Optimization Landscape of Shallow ReLU Neural Networks**
  Itay Safran (Princeton University) ; Gilad Yehudai (Weizmann Institute of Science) ; Ohad Shamir (Weizmann Institute of Science)
- **Non-asymptotic approximations of neural networks by Gaussian processes**
  Ronen Eldan (Weizmann) ; Dan Mikulincer (Weizmann Institute) ; Tselil Schramm (Harvard University)
- **Size and Depth Separation in Approximating Benign Functions with Neural Networks**
  Gal Vardi (Weizmann Institute of Science) ; Daniel Reichman (Worcester Polytechnic Institute) ; Toniann Pitassi (University of Toronto) ; Ohad Shamir (Weizmann Institute of Science)

Robustness, Privacy and Fairness (A)

Tue, Aug 17, 09:50 – 10:15
TRACK 2

- **Differentially Private Nonparametric Regression Under a Growth Condition**
  Noah Golowich (Massachusetts Institute of Technology)

- **On Avoiding the Union Bound When Answering Multiple Differentially Private Queries**
  Badih Ghazi (Google) ; Ravi Kumar (Google) ; Pasin Manurangsi (Google)
- **Non-Euclidean Differentially Private Stochastic Convex Optimization**
  Raef Bassily (The Ohio State University) ; Cristobal Guzman (PUC Chile) ;
  Anupama Nandi (The Ohio State University)
- **The Sparse Vector Technique, Revisited**
  Haim Kaplan (TAU, GOOGLE) ; Yishay Mansour (Tel Aviv University and Google Research) ;
  Uri Stemmer (Ben-Gurion University)
- **Machine Unlearning via Algorithmic Stability**
  Enayat Ullah (Johns Hopkins University) ; Tung Mai (Adobe Research) ; Anup Rao (Adobe Research) ;
  Ryan Rossi (Adobe Research) ; Raman Arora (Johns Hopkins University)

Neural Networks/Deep Learning (B)

Tue, Aug 17, 10:20 – 10:40
TRACK 1

- **The Connection Between Approximation, Depth Separation and Learnability in Neural Networks**
  Eran Malach (Hebrew University Jerusalem Israel) ; Gilad Yehudai (Weizmann Institute of Science) ;
  Shai Shalev-Schwartz (Hebrew University of Jerusalem) ; Ohad Shamir (Weizmann Institute of Science)
- **Implicit Regularization in ReLU Networks with the Square Loss**
  Gal Vardi (Weizmann Institute of Science) ; Ohad Shamir (Weizmann Institute of Science)
- **On the Approximation Power of Two-Layer Networks of Random ReLUs**
  Daniel Hsu (Columbia University) ; Clayton Sanford (Columbia) ; Rocco Servedio (Columbia University) ;
  Emmanouil Vlatakis-Gkaragkounis (Columbia University)
- **When does gradient descent with logistic loss interpolate using deep networks with smoothed ReLU activations?**
  Niladri Chatterji (UC Berkeley) ; Philip Long (Google) ; Peter Bartlett

Robustness, Privacy and Fairness (B))

Tue, Aug 17, 10:20 – 10:40
TRACK 2

- **A Law of Robustness for Two-Layers Neural Networks**
  Sebastien Bubeck (Microsoft Research) ; Yuanzhi Li (CMU) ;
  Dheeraj Nagaraj (Massachusetts Institute of Technology)
- **Adversarially Robust Low Dimensional Representations**
  Pranjal Awasthi (Rutgers University/Google) ; Vaggos Chatziafratis (Google) ;
  Xue Chen (George Mason University) ; Aravindan Vijayaraghavan
- **Adversarially Robust Learning with Unknown Perturbation Sets**
  Omar Montasser (TTI-Chicago) ; Steve Hanneke (Toyota Technological Institute at Chicago) ;
  Nathan Srebro (Toyota Technical Institute of Chicago)
- **Moment Multicalibration for Uncertainty Estimation**

Christopher Jung (University of Pennsylvania) ; Changhwa Lee (University of Pennsylvania) ;
Mallesh Pai (Rice University) ; Aaron Roth (University of Pennsylvania) ; Rakesh Vohra (University of Pennsylvania)

Online Learning, Game Theory 2 (A)

Tue, Aug 17, 11:10 – 11:30
TRACK 1

- **Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information**
  Angeliki Giannou (National Technical University of Athens) ;
  Emmanouil Vlatakis-Gkaragkounis (Columbia University) ; Panayotis Mertikopoulos (CNRS and Criteo AI Lab)
- **Adaptive Learning in Continuous Games: Optimal Regret Bounds and Convergence to Nash Equilibrium**
  Yu-Guan Hsieh (Univ. Grenoble Alpes) ; Kimon Antonakopoulos (Inria) ;
  Panayotis Mertikopoulos (CNRS and Criteo AI Lab)
- **Learning in Matrix Games can be Arbitrarily Complex**
  Gabriel Andrade (University of Colorado Boulder) ; Rafael Frongillo (CU Boulder) ;
  Georgios Piliouras (Singapore University of Technology and Design)
- **Robust Online Convex Optimization in the Presence of Outliers**
  Tim van Erven (University of Amsterdam) ; Sarah Sachs (University of Amsterdam) ;
  Wouter Koolen (CWI Amsterdam) ; Wojciech Kotlowski (Poznan University of Technology)

Bandits, RL and Control 1 (A)

Tue, Aug 17, 11:10 – 11:30
TRACK 2

- **Improved Regret for Zeroth-Order Stochastic Convex Bandits**
  Tor Lattimore (DeepMind) ; Andras Gyorgy (DeepMind)
- **Minimax Regret for Stochastic Shortest Path with Adversarial Costs and Known Transition**
  Liyu Chen (USC) ; Haipeng Luo (USC) ; Chen-Yu Wei (University of Southern California)
- **Towards a Dimension-Free Understanding of Adaptive Linear Control**
  Juan Perdomo (University of California, Berkeley) ; Max Simchowitz (UC Berkeley) ; Alekh Agarwal (Microsoft) ;
  Peter Bartlett
- **Fine-Grained Gap-Dependent Bounds for Tabular MDPs via Adaptive Multi-Step Bootstrap**
  Haike Xu (Tsinghua University) ; Tengyu Ma (Stanford) ; Simon Du (University of Washington)

Online Learning, Game Theory 2 (B)

Tue, Aug 17, 11:35 – 11:55
TRACK 1

- **Black-Box Control for Linear Dynamical Systems**
  Xinyi Chen (Princeton University; Google) ; Elad Hazan (Princeton University)
- **Majorizing Measures, Sequential Complexities, and Online Learning**
  Adam Block (MIT) ; Yuval Dagan (MIT) ; Alexander Rakhlin (MIT)

- **Instance-Dependent Complexity of Contextual Bandits and Reinforcement Learning: A Disagreement-Based Perspective**
  Dylan Foster (MIT) ; Alexander Rakhlin (MIT) ; David Simchi-Levi (MIT) ; Yunzong Xu (MIT)
- **Sequential prediction under log-loss and misspecification**
  Meir Feder (TAU) ; Yury Polyanskiy (MIT)

Bandits, RL and Control 1 (B)

Tue, Aug 17, 11:35 – 11:55
TRACK 2

- **Is Reinforcement Learning More Difficult Than Bandits? A Near-optimal Algorithm Escaping the Curse of Horizon**
  Zihan Zhang (Tsinghua University) ; Xiangyang Ji (Tsinghua University) ; Simon Du (University of Washington)
- **Corruption-robust exploration in episodic reinforcement learning**
  Thodoris Lykouris (Microsoft Research NYC) ; Max Simchowitz (UC Berkeley) ; Alex Slivkins (Microsoft Research) ; Wen Sun (Cornell University)
- **On Query-efficient Planning in MDPs under Linear Realizability of the Optimal State-value Function**
  Gellert Weisz (DeepMind, UCL) ; Philip Amortila (University of Illinois at Urbana-Champaign) ; Barnabás Janzer (University of Cambridge) ; Yasin Abbasi-Yadkori (DeepMind) ; Nan Jiang (University of Illinois at Urbana-Champaign) ; Csaba Szepesvari (DeepMind/University of Alberta)
- **Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games**
  Chen-Yu Wei (University of Southern California) ; Chung-Wei Lee (University of Southern California) ; Mengxiao Zhang (University of Southern California) ; Haipeng Luo (USC)

Keynote: David Silver (DeepMind)

Tue, Aug 17, 12:05 – 13:00
MAIN TRACK

Open Problem 3

Wed, Aug 18, 09:30 – 09:40
MAIN TRACK

- **Are all VC classes learnable with computable learners?**
  Sushant Agarwal (University of Waterloo) ; Nivasini Ananthakrishnan (University of Waterloo) ; Shai Ben-David (University of Waterloo) ; Tosca Lechner (University of Waterloo) ; Ruth Urner (YorkUniversity)

Generalization and PAC-Learning 1 (A)

Wed, Aug 18, 09:50 – 10:15
TRACK 1

- **Generalizing Complex Hypotheses on Product Distributions: Auctions, Prophet Inequalities, and Pandora's Problem**
  Chenghao Guo (Massachusetts Institute of Technology) ; Zhiyi Huang (University of Hong Kong) ; Zhihao Gavin Tang (Shanghai University of Finance and Economics) ; Xinzhi Zhang (University of Washington)
- **PAC-Bayes, MAC-Bayes and Conditional Mutual Information: Fast rate bounds that handle general VC classes**
  Peter Grunwald (Centrum voor Wiskunde en Informatica) ; Thomas Steinke (IBM, Almaden) ; Lydia Zakynthinou (Northeastern University)
- **Agnostic Proper Learning of Halfspaces under Gaussian Marginals**
  Ilias Diakonikolas (UW Madison) ; Daniel Kane (UCSD) ; Vasilis Kontonis (University of Wisconsin-Madison) ; Christos Tzamos (UW-Madison) ; Nikos Zarifis (University of Wisconsin-Madison)
- **Benign Overfitting of Constant-Stepsize SGD for Linear Regression**
  Difan Zou (University of California, Los Angeles) ; Jingfeng Wu (Johns Hopkins University) ; Vladimir Braverman (Johns Hopkins University) ; Quanquan Gu (University of California, Los Angeles) ; Sham Kakade (University of Washington)
- **Boosting in the Presence of Massart Noise**
  Ilias Diakonikolas (UW Madison) ; Russell Impagliazzo (UC - San Diego) ; Daniel Kane (UCSD) ; Rex Lei (University of California - San Diego) ; Jessica Sorrell (University California, San Diego) ; Christos Tzamos (UW-Madison)

Regression and High-Dimensional Statistics (A)

Wed, Aug 18, 09:50 – 10:15
TRACK 2

- **Efficient Algorithms for Learning from Coarse Labels**
  Dimitris Fotakis (National Technical University of Athens) ; Alkis Kalavasis (National Technical University of Athens) ; Vasilis Kontonis (University of Wisconsin-Madison) ; Christos Tzamos (UW-Madison)
- **Impossibility of Partial Recovery in the Graph Alignment Problem**
  Luca Ganassali (INRIA (Paris)) ; Laurent Massoulie (Microsoft-Inria Joint Center) ; Marc Lelarge (INRIA-ENS)
- **Hypothesis testing with low-degree polynomials in the Morris class of exponential families**
  Dmitriy Kunisky (New York University)
- **Group testing and local search: is there a computational-statistical gap?**
  Fotis Iliopoulos (Princeton University) ; Ilias Zadik (NYU)
- **Johnson-Lindenstrauss Transforms with Best Confidence**
  Maciej Skorski (University of Luxembourg)

Generalization and PAC-Learning 1 (B)

Wed, Aug 18, 10:20 – 10:40
TRACK 1

- **Fast Rates for Structured Prediction**
  Vivien Cabannnes (INRIA) ; Francis Bach (INRIA - Ecole Normale Supérieure) ; Alessandro Rudi (INRIA, Ecole Normale Superieure)

- **A Priori Generalization Analysis of the Deep Ritz Method for Solving High Dimensional Elliptic Partial Differential Equations**
  Yulong Lu (University of Massachusetts Amherst) ; Jianfeng Lu (Duke University) ; Min Wang (Duke University)
- **On Empirical Bayes Variational Autoencoder: An Excess Risk Bound**
  Rong Tang (University of Illinois at Urbana Champaign) ; Yun Yang (Univeristy of Illinois at Urbana-Champaign)
- **Information-Theoretic Generalization Bounds for Stochastic Gradient Descent**
  Gergely Neu (Universitat Pompeu Fabra)

Regression and High-Dimensional Statistics (B)

Wed, Aug 18, 10:20 – 10:40
TRACK 2

- **Reduced-Rank Regression with Operator Norm Error**
  Praneeth Kacham (Carnegie Mellon University) ; David Woodruff (Carnegie Mellon University)
- **Rank-one matrix estimation: analytic time evolution of gradient descent dynamics**
  Antoine Bodin (EPFL) ; Nicolas Macris (Ecole Polytechnique Federale de Lausanne)
- **It was ``all'' for ``nothing'': sharp phase transitions for noiseless discrete channels**
  Ilias Zadik (NYU) ; Jonathan Niles-Weed (NYU)
- **Stochastic block model entropy and broadcasting on trees with survey**
  Emmanuel Abbe (EPFL) ; Elisabetta Cornacchia (EPFL) ; Yuzhou Gu (Massachusetts Institute of Technology) ; Yury Polyanskiy (MIT)

Networks and Graphs

Wed, Aug 18, 11:10 – 11:35
TRACK 1

- **Random Graph Matching with Improved Noise Robustness**
  Cheng Mao (Georgia Institute of Technology) ; Mark Rudelson ; Konstantin Tikhomirov (Georgia Institute of Technology)
- **Quantifying Variational Approximation for Log-Partition Function**
  Romain Cosson (MIT) ; Devavrat Shah (MIT)
- **Source Identification for Mixtures of Product Distributions**
  Spencer Gordon (California Institute of Technology) ; Bijan Mazaheri (California Institute of Technology) ; Yuval Rabani (Hebrew University of Jerusalem) ; Leonard Schulman (Caltech)
- **Learning to Sample from Censored Markov Random Fields**
  Ankur Moitra (MIT) ; Elchanan Mossel (MIT) ; Colin Sandon (MIT)
- **Learning from Censored and Dependent Data: The case of Linear Dynamics**
  Orestis Plevrakis (Princeton University)

Bandits, RL and Control 2 (A)

Wed, Aug 18, 11:10 – 11:30

TRACK 2

- **Cooperative and Stochastic Multi-Player Multi-Armed Bandit: Optimal Regret With Neither Communication Nor Collisions**
  Mark Sellke ; Sebastien Bubeck (Microsoft Research) ; Thomas Budzinski (ENS Lyon)
- **Softmax Policy Gradient Methods Can Take Exponential Time to Converge**
  Gen Li (Tsinghua University, China) ; Yuting Wei (University of Pennsylvania) ; Yuejie Chi (CMU) ; Yuantao Gu (Tsinghua University) ; Yuxin Chen (Princeton University)
- **Asymptotically Optimal Information-Directed Sampling**
  Johannes Kirschner (ETH Zurich) ; Tor Lattimore (DeepMind) ; Claire Vernade (DeepMind) ; Csaba Szepesvari (DeepMind/University of Alberta)
- **Fast Rates for the Regret of Offline Reinforcement Learning**
  Yichun Hu (Cornell University) ; Nathan Kallus (Cornell University) ; Masatoshi Uehara (Cornell University)

Bandits, RL and Control 2 (B)

Wed, Aug 18, 11:35 – 11:55
TRACK 2

- **Nearly Minimax Optimal Reinforcement Learning for Linear Mixture Markov Decision Processes**
  Dongruo Zhou (UCLA) ; Quanquan Gu (University of California, Los Angeles) ; Csaba Szepesvari (DeepMind/University of Alberta)
- **Parameter-Free Multi-Armed Bandit Algorithms with Hybrid Data-Dependent Regret Bounds**
  Shinji Ito (NEC Corporation)
- **Learning to Stop with Surprisingly Few Samples**
  Tianyi Zhang (Columbia University) ; Daniel Russo (Columbia University) ; Assaf Zeevi (Columbia University)
- **Improved Analysis of the Tsallis-INF Algorithm in Stochastically Constrained Adversarial Bandits and Stochastic Bandits with Adversarial Corruptions**
  Saeed Masoudian (University of Copenhagen) ; Yevgeny Seldin (University of Copenhagen)

Clustering

Wed, Aug 18, 11:40 – 11:55
TRACK 1

- **Towards a Query-Optimal and Time-Efficient Algorithm for Clustering with a Faulty Oracle**
  Pan Peng (University of Sheffield) ; Jiapeng Zhang (University of Southern California)
- **Exact Recovery of Clusters in Finite Metric Spaces Using Oracle Queries**
  Marco Bressan (University of Milan) ; Nicolò Cesa-Bianchi (University of Milan) ; Silvio Lattanzi (Google) ; Andrea Paudice (University of Milan & Istituto Italiano di Tecnologia)
- **Approximation Algorithms for Socially Fair Clustering**
  Yury Makarychev (TTIC) ; Ali Vakilian (Toyota Technological Institute at Chicago)

Keynote: Sara van de Geer (ETH Zurich)

Wed, Aug 18, 12:05 – 13:00
MAIN TRACK

Open Problem 4

Thu, Aug 19, 09:30 – 09:40
MAIN TRACK

- **Open Problem: Tight Online Confidence Intervals for RKHS Elements**
  Sattar Vakili (MediaTek Research) ; Jonathan Scarlett (National University of Singapore) ; Tara Javidi (University of California San Diego)

Generalization and PAC-Learning 2 (A)

Thu, Aug 19, 09:50 – 10:15
TRACK 1

- **Concentration of Non-Isotropic Random Tensors with Applications to Learning and Empirical Risk Minimization**
  Mathieu Even (ENS - Inria) ; Laurent Massoulie (Microsoft-Inria Joint Center)
- **A Theory of Heuristic Learnability**
  Mikito Nanashima (Tokyo Institute of Technology)
- **From Local Pseudorandom Generators to Hardness of Learning**
  Amit Daniely (Hebrew University) ; Gal Vardi (Weizmann Institute of Science)
- **Functions with average smoothness: structure, algorithms, and learning**
  Yair Ashlagi (Ben Gurion University) ; Lee-Ad Gottlieb (Ariel University) ; Aryeh Kontorovich (Ben-Gurion University of the Negev)
- **Robust learning under clean-label attack**
  Avrim Blum (Toyota Technological Institute of Chicago) ; Steve Hanneke (Toyota Technological Institute at Chicago) ; Jian Qian (MIT) ; Han Shao (Toyota Technological Institute at Chicago)

Nonparametrics

Thu, Aug 19, 09:50 – 10:15
TRACK 2

- **Learning with invariances in random features and kernel models**
  Song Mei (UC Berkeley) ; Theodor Misiakiewicz (Stanford University) ; Andrea Montanari (Stanford University)
- **Kernel Thinning**
  Raaz Dwivedi (UNIVERSITY OF CALIFORNIA Berkeley) ; Lester Mackey (Microsoft Research New England)
- **On the Minimal Error of Empirical Risk Minimization**
  Gil Kur (MIT) ; Alexander Rakhlin (MIT)
- **Nonparametric Regression with Shallow Overparameterized Neural Networks Trained by GD with Early Stopping**
  Ilja Kuzborskij (DeepMind) ; Csaba Szepesvari (DeepMind/University of Alberta)

- **A Statistical Taylor Theorem and Extrapolation of Truncated Densities**
  Constantinos Daskalakis (MIT) ; Vasilis Kontonis (University of Wisconsin-Madison) ;
  Christos Tzamos (UW-Madison) ; Emmanouil Zampetakis (Massachusetts Institute of Technology)

Generalization and PAC-Learning 2 (B)

Thu, Aug 19, 10:20 – 10:40
TRACK 1

- **Improved Algorithms for Efficient Active Learning Halfspaces with Massart and Tsybakov Noise**
  Chicheng Zhang (University of Arizona) ; Yinan Li (University of Arizona)
- **The Optimality of Polynomial Regression for Agnostic Learning under Gaussian Marginals in the SQ Model**
  Ilias Diakonikolas (UW Madison) ; Daniel Kane (UCSD) ; Thanasis Pittas (University of Wisconsin-Madison) ;
  Nikos Zarifis (University of Wisconsin-Madison)
- **Near Optimal Distributed Learning of Halfspaces with Two Parties**
  Mark Braverman (Princeton University) ; Gillat Kol (Princeton University) ; Shay Moran (Technion) ;
  Raghuvansh R. Saxena (Princeton University)
- **Query complexity of least absolute deviation regression via robust uniform convergence**
  Xue Chen (George Mason University) ; Michal Derezinski (UC Berkeley)

Sampling Algorithms

Thu, Aug 19, 10:20 – 10:45
TRACK 2

- **Optimal dimension dependence of the Metropolis-Adjusted Langevin Algorithm**
  Sinho Chewi (Massachusetts Institute of Technology) ; Chen Lu (Massachusetts Institute of Technology) ;
  Kwangjun Ahn (MIT) ; Xiang Cheng (Postdoc at MIT (EECS)) ;
  Thibaut Le Gouic (Massachusetts Institute of Technology) ; Philippe Rigollet (MIT)
- **Random Coordinate Langevin Monte Carlo**
  Zhiyan Ding (University of Wisconsin-Madison) ; Qin Li (University of Wisconsin-Madison) ;
  Jianfeng Lu (Duke University) ; Stephen Wright (University of Wisconsin-Madison)
- **Near-Optimal Entrywise Sampling of Numerically Sparse Matrices**
  Vladimir Braverman (Johns Hopkins University) ; Robert Krauthgamer (Weizmann Institute of Science, Israel) ;
  Aditya Krishnan (Johns Hopkins University) ; Shay Sapir (Weizmann Institute of Science)
- **Structured Logconcave Sampling with a Restricted Gaussian Oracle**
  Yin Tat Lee (UW) ; Ruoqi Shen (University of Washington) ; Kevin Tian (Stanford University)
- **On the Convergence of Langevin Monte Carlo: The Interplay between Tail Growth and Smoothness**
  Murat Erdogdu (University of Toronto, Vector Institute) ; Rasa Hosseinzadeh (University of Toronto)

Stochastic Optimization (A)

Thu, Aug 19, 11:10 – 11:30
TRACK 1

- **SGD in the Large: Average-case Analysis, Asymptotics, and Stepsize Criticality**

Courtney Paquette (McGill University) ; Kiwon Lee (McGill University) ; Fabian Pedregosa (Google) ; Elliot Paquette (McGill University)

- **Stochastic Approximation for Online Tensorial Independent Component Analysis**
  Chris Junchi Li (UC Berkeley) ; Michael Jordan (UC Berkeley)
- **On the Stability of Random Matrix Product with Markovian Noise: Application to Linear Stochastic Approximation and TD Learning**
  Alain Durmus (ENS Paris Saclay) ; Eric Moulines (Ecole Polytechnique) ; Alexey Naumov (National Research University Higher School of Economics) ; Sergey Samsonov (National Research University Higher School of Economics) ; Hoi-To Wai (Chinese University of Hong Kong)
- **Shape Matters: Understanding the Implicit Bias of the Noise Covariance**
  Jeff Z. HaoChen (Stanford University) ; Colin Wei (Stanford University) ; Jason Lee (Princeton) ; Tengyu Ma (Stanford)

Bandits, RL and Control 3 (A)

Thu, Aug 19, 11:10 – 11:30
TRACK 2

- **Double Explore-then-Commit: Asymptotic Optimality and Beyond**
  Tianyuan Jin (National University of Singapore) ; Pan Xu (University of California, Los Angeles) ; Xiaokui Xiao (National University of Singapore) ; Quanquan Gu (University of California, Los Angeles)
- **Adaptive Discretization for Adversarial Lipschitz Bandits**
  Chara Podimata (Harvard University) ; Alex Slivkins (Microsoft Research)
- **Efficient Bandit Convex Optimization: Beyond Linear Losses**
  Arun Suggala (Carnegie Mellon University) ; Pradeep Ravikumar (Carnegie Mellon University) ; Praneeth Netrapalli (Microsoft Research)
- **Mirror Descent and the Information Ratio**
  Tor Lattimore (DeepMind) ; Andras Gyorgy (DeepMind)

Stochastic Optimization (B)

Thu, Aug 19, 11:35 – 11:55
TRACK 1

- **Almost sure convergence rates for Stochastic Gradient Descent and Stochastic Heavy Ball**
  Othmane Sebbouh (ENS Paris) ; Robert Gower (Telecom Paris Tech) ; Aaron Defazio (Facebook AI Research)
- **SGD Generalizes Better Than GD (And Regularization Doesn't Help)**
  Idan Amir (Tel-Aviv University) ; Tomer Koren (Tel Aviv University and Google) ; Roi Livni (Tel Aviv University)
- **Convergence rates and approximation results for SGD and its continuous-time counterpart**
  Xavier Fontaine (ENS Paris-Saclay) ; Valentin De Bortoli (ENS Paris-Saclay) ; Alain Durmus (ENS Paris Saclay)
- **Streaming k-PCA: Efficient guarantees for Oja's algorithm, beyond rank-one updates**
  De Huang (Caltech) ; Jonathan Niles-Weed (NYU) ; Rachel Ward (University of Texas)

Bandits, RL and Control 3 (B)

Thu, Aug 19, 11:35 – 11:55
TRACK 2

- **Multiplayer Bandit Learning, from Competition to Cooperation**
  Simina Branzei (Purdue University) ; Yuval Peres (N/A)
- **Non-stationary Reinforcement Learning without Prior Knowledge: an Optimal Black-box Approach**
  Chen-Yu Wei (University of Southern California) ; Haipeng Luo (USC)
- **Cautiously Optimistic Policy Optimization and Exploration with Linear Function Approximation**
  Andrea Zanette (Stanford University) ; Ching-An Cheng (Microsoft) ; Alekh Agarwal (Microsoft)
- **Regret Minimization in Heavy-Tailed Bandits**
  Shubhada Agrawal (TIFR, Mumbai) ; Sandeep Juneja (TIFR Mumbai) ; Wouter Koolen (CWI Amsterdam)

Keynote: Persi Diaconis (Standford University)

Thu, Aug 19, 12:05 – 13:00
MAIN TRACK