Statistics Winter Workshop: 2023

Modern Computational Statistics

Friday, January 13th
9:00AM—5:00PM

Saturday, January 14th
9:00AM—12:30PM

J. Wayne Rietz Union
Gainesville, FL

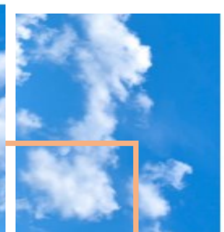

Venue

J. Wayne Reitz Union

655 Reitz Union Drive
University of Florida Campus
Gainesville, FL 32611

Rooms

Presentations: Auditorium
Refreshments: 2365
Posters: 2335

Parking Information

The Visitor Welcome Center and Bookstore parking garage is located at the Reitz Union, at the corner of Museum Road and Reitz Union Drive. The garage is an unrestricted pay facility available to all members of the university community. This garage can accommodate 300 vehicles. There are 45 short term parking spaces located in the garage.

The garage hours of operation are Monday through Friday, 7:30 am to 4:30 pm. Short-term and daily fees apply during this time. The garage may be used during non-operating hours for short-term parking, free of charge.

Program

Friday, January 13, 2023

8:30 AM: Breakfast (Rm 2365)

9:00 AM: Welcome

Dr. Michael Daniels, University of Florida
Department Chair, Department Statistics

9:15 AM: Session 1

Chair: Dr. Daniels

Dr. Martin Wainwright, Massachusetts Institute of Technology

"Computational-statistical Interplay in Reinforcement Learning"

Dr. Yuting Wei, (virtual) University of Pennsylvania
"A Non-asymptotic Framework for Approximate Message Passing"

10:45 AM: Morning Break (Rm 2365)

11:00 AM: Session 2

Chair: Dr. Aaron Molstad

Dr. Rajesh Ranganath, New York University
"Out of Distribution Generalization"

11:45 AM: Lunch (Rm 2365)

1:00PM: Session 3 (virtual)

Chair: Dr. Kshitij Khare
Dr. Hui Zou, University of Minnesota
"High-dimensional Clustering via Latent Semiparametric Mixture Models"

Dr. Ambuj Tewari ,University of Michigan
"The Tortoise and the Hare: How Learning Theory is Catching Up to Applied Machine Learning "

2:30 PM: Afternoon Break (Rm 2365)

2:45 PM: Session 4 (virtual)

Chair: Dr. Hani Doss

Dr. Matt Hoffman, GOOGLE
"Running Many-Chain MCMC on Cheap GPUs"

3:30 PM: Poster Session (Rm 2335)

Program

Saturday, January 14, 2023

8:30 AM: Breakfast (Rm 2365)

9:00 AM : Session 5 (virtual)

Chair: Dr. Zhihua Su

Dr. Yuejie Chi, Carnegie Mellon University
"Offline Reinforcement Learning: Towards Optimal Sample Complexity and Distributional Robustness"

Dr. Hua Zhou, UCLA, Los Angeles
"Orthogonal Trace-sum Maximization: Applications, Local Algorithms, and Global Optimality"

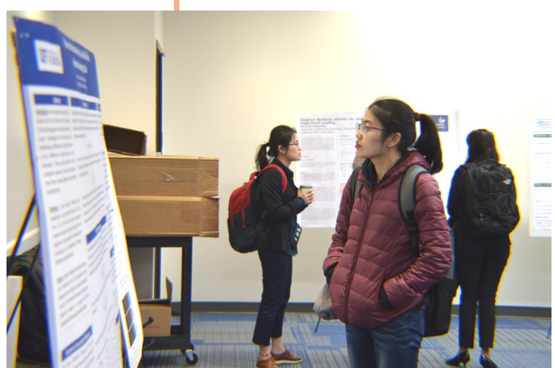

10:30 AM: Morning Break

10:45 AM: Session 6

Chair: Dr. Bikram Karmakar

Dr. Yufeng Liu, (virtual) University of North Carolina
"Distribution-Free Contextual Dynamic Pricing"

Dr. Max Simchowitz, Massachusetts Institute of Technology
"Randomized Smoothing, Online Learning, and Planning Through Contact"

"Offline Reinforcement Learning: Towards Optimal Sample Complexity and Distributional Robustness"

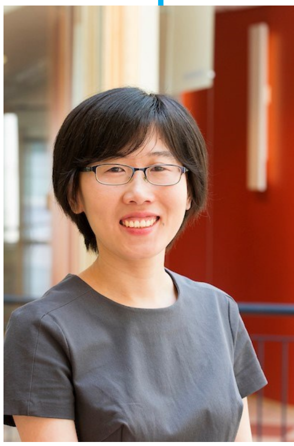

Dr. Yuejie Chi *Professor, Dept. of Electrical & Computer Engineering*

Carnegie Mellon University

Abstract

Offline or batch reinforcement learning seeks to learn a near-optimal policy using history data without active exploration of the environment. To counter the insufficient coverage and sample scarcity of many offline datasets, the principle of pessimism has been recently introduced to mitigate high bias of the estimated values. However, prior algorithms or analyses either suffer from suboptimal sample complexities or incur high burn-in cost to reach sample optimality, thus posing an impediment to efficient offline RL in sample-starved applications. In this talk, we demonstrate that the model-based (or "plug-in") approach achieves minimaxoptimal sample complexity without burn-in cost for tabular Markov decision processes (MDPs). Our algorithms are "pessimistic" variants of value iteration with Bernstein-style penalties, and do not require sophisticated variance reduction. We further consider a distributionally robust formulation of offline RL, focusing on tabular robust MDPs with an uncertainty set specified by the Kullback-Leibler divergence. Again, a model-based algorithm that combines distributionally robust value iteration with the principle of pessimism achieves a near-optimal sample complexity up to a polynomial factor of the effective horizon length.

Biography

Dr. Yuejie Chi is a Professor in the department of Electrical and Computer Engineering, and a faculty affiliate with the Machine Learning department and CyLab at Carnegie Mellon University. She received her Ph.D. and M.A. from Princeton University, and B. Eng. (Hon.) from Tsinghua University, all in Electrical Engineering.
Her research interests lie in the theoretical and algorithmic foundations of data science, signal processing, machine learning and inverse problems, with applications in sensing, imaging, and societal systems, broadly defined. Among others, Dr. Chi received the Presidential Early Career Award for Scientists and Engineers (PECASE), the inaugural IEEE Signal Processing Society Early Career Technical Achievement Award for contributions to high-dimensional structured signal processing, and held the inaugural Robert E. Doherty Early Career Development Professorship. She was named a Goldsmith Lecturer by IEEE Information Theory Society and a Distinguished Lecturer by IEEE Signal Processing Society.

“Running Many-Chain MCMC on Cheap GPUs”

Dr. Matt Hoffman

Research Scientist

Google

Abstract

Traditional Markov chain Monte Carlo (MCMC) workflows run a small number of parallel chains for many steps. This workflow made sense in the days when parallel computation was limited to the number of cores on one’s CPU, since it lets one amortize the cost of burn-in/warmup (i.e., discarding early samples to eliminate bias) over as many useful samples as possible. But modern commodity GPUs (<$500 retail, or a few dollars per hour to rent in the cloud), in conjunction with modern software frameworks like TensorFlow Probability, make it possible to run 50–100 chains in parallel without paying much of a price in wallclock time. For many Bayesian inference problems, this means that we can get reasonably low-variance estimates of posterior expectations using as little as one sample per chain, dramatically reducing the time we spend waiting for results. In this talk, I will present some of our work on new MCMC algorithms and workflows that can take full advantage of GPUs .

Biography

Dr. Matthew D. Hoffman is a Research Scientist at Google. His main research interests are in the computational issues that arise in probabilistic modeling and inference, and applications that arise in various areas of statistics, generative modeling, and deep learning.

Dr. Yufeng Liu

Professor, Depts. of Statistics & Operations Research, Genetics, & Biostatistics

University of North Carolina

"Distribution-Free Contextual Dynamic Pricing"

Abstract

Contextual dynamic pricing aims to set personalized prices based on sequential interactions with customers. At each time period, a customer who is interested in purchasing a product comes to the platform. The customer's valuation for the product is a linear function of contexts, including product and customer features, plus some random market noise. The seller does not observe the customer's true valuation, but instead needs to learn the valuation by leveraging contextual information and historical binary purchase feedbacks. Existing models typically assume full or partial knowledge of the random noise distribution. In this talk, we consider contextual dynamic pricing with unknown random noise. Our distribution-free pricing policy learns both the contextual function and the market noise simultaneously. A key ingredient of our method is a novel perturbed linear bandit framework, where a modified linear upper confidence bound algorithm is proposed to balance the exploration of market noise and the exploitation of the current knowledge for better pricing.

Biography

Dr. Yufeng Liu is professor in Department of Statistics and Operations Research, Department of Genetics, and Department of Biostatistics at University of North Carolina at Chapel Hill. His research interests include statistical machine learning, high dimensional data, precision medicine, and applications in bioinformatics and e-commerce. He received the CAREER Award from National Science Foundation in 2008, Leo Breiman Junior Award from American Statistical Association in 2017. He is an elected fellow at American Statistical Association (ASA) and Institute of Mathematical Statistics (IMS).

"Out of Distribution Generalization"

Dr. Rajesh Ranganath
Assistant Professor, Courant Institute of Mathematical Scines
New York University

Abstract

Supervised learning typically requires test data being drawn from the same distribution as the training data. This assumption does not hold in many real-world settings. As an example setting consider a model trained in one hospital and used in another. In such a setting, the test distribution is related to the training distribution, but they are not the same. In this talk, I will discuss our work on representation learning for out of distribution generalization. I will construct a family of representations that generalize under changing spurious correlations with applications to images and chest X-rays.

References:

https://arxiv.org/pdf/2107.00520.pdf

https://arxiv.org/pdf/2210.01302.pdf

Biography

Dr. Rajesh Ranganath is an assistant professor at NYU's Courant Institute of Mathematical Sciences and the Center for Data Science. He is also affiliate faculty at the Department of Population Health at NYUMC. His research focuses on approximate inference, causal inference, probabilistic models, and machine learning for healthcare. Rajesh completed his PhD at Princeton and BS and MS from Stanford University. Rajesh has won several awards including the NDSEG graduate fellowship, the Porter Ogden Jacobus Fellowship, given to the top four doctoral students at Princeton University, the Savage Award in Theory and Methods, and an NSF Career Award.

“Randomized Smoothing, Online Learning, and Planning Through Contact”

Dr. Max Simchowitz
Postdoctoral Researcher, CSAIL
Massachusetts Institute of Technology

Abstract

Trajectory planning through discontinuous and “stiff” (large Lipschitz) dynamics is a fundamental problem in robotic manipulation. In this talk, we explore the statistical and computational challenges of stiff trajectory optimization through the lens of randomized smoothing, where the dynamics are intentionally corrupted with a small bit of continuous noise. We begin with empirical evidence and a heuristic justification of randomized smoothing for trajectory optimization in comparison to alternatives such as automated differentiation. We then introduce the framework of smoothed online learning and turn our attention to piecewise affine systems, a class of models that exhibit the stiffness property encountered in contact. Finally, we propose novel oracle-efficient, low-regret algorithms for smooth online learning that lead to algorithms for online prediction, online simulation, and robust planning under piecewise affine dynamics, as well as considerable generalizations. Joint work with Adam Block, Terry Suh, Tao Peng, Kaiqing Zhang, and Russ Tedrake at MIT.

Biography

Dr. Max Simchowitz is a postdoctoral researcher under Russ Tedrake in the Robot Locomotion group, part of CSAIL at MIT. He received his PhD in the EECS department at UC Berkeley under Michael I. Jordan and Benjamin Recht, generously supported by Open Philanthropy, NSF GRFP, and Berkeley Fellowships, and was fortunate enough to receive two ICML outstanding paper awards. His work focuses broadly on machine learning theory, with a recent focus on the plentiful intersections between statistical learning theory, online learning, non-convex optimization, and control theory.

Dr. Ambuj Tewari
Professor, Dept. of Statistics
University of Michigan

"The Tortoise and the Hare: How Learning Theory is Catching Up to Applied Machine Learning"

Abstract

Efron and Hastie wrote Computer Age Statistical Inference giving us a masterful survey of the major developments in statistics and machine learning that occurred during the 20th century. They compared the race between algorithms and their underlying theory to the classic hare and tortoise story. With machine learning deployed widely, the gap between practice and theory only seems to be widening in the first quarter of the 21st century. However, theoreticians have been working hard to make the theory tortoise move a little faster and catch up to the algorithmic hare. A lot of recent theoretical research in machine learning has been driven by concerns arising in applications. How do we build predictors while keeping users' data private? How do we tackle non-stationarity? How do we protect deployed systems from adversarial attacks? How do we account for the fact that predictions are often fed into optimization solvers to make decisions further downstream? In my talk I will highlight some recent theoretical advances in addressing such concerns with a bias towards topics that my own research group has worked on recently.

Biography

Dr. Ambuj Tewari is a professor in the Department of Statistics and the Department of EECS (by courtesy) at the University of Michigan, Ann Arbor. He is also affiliated with the Michigan Institute for Data Science (MIDAS). He obtained his PhD under the supervision of Peter Bartlett at the University of California at Berkeley. His research interests lie in machine learning including statistical learning theory, online learning, reinforcement learning and control theory, network analysis, and optimization for machine learning. He collaborates with scientists to seek novel applications of machine learning in behavioral sciences, psychiatry, and chemical sciences. His research has been recognized with paper awards at COLT 2005, COLT 2011, and AISTATS 2015. He received an NSF CAREER award in 2015, a Sloan Research Fellowship in 2017, and was named IMS Fellow in 2022.

Dr. Martin Wainwright

Cecil H. Green Professor, Laboratory ofr Information & Decision Systems

Massachusetts Institute of Technology

"Computational-statistical Interplay in Reinforcement Learning"

Abstract

Reinforcement learning and stochastic control involve statistical procedures for using data to determine near-optimal policies in sequential settings. The associated framework of Markov decision processes has applications in numerous areas (e.g., environmental control, robotics, supply chain management, competitive game-playing, industrial process control).

In this talk, we discuss some of computational and statistical challenges that arise in these domains, including fast Krylov-based procedures for solving non-parametric operator equations, projected fixed point methods for approximation, and Bellman orthogonality for off-policy problems.

Based on joint work with Yaqi Duan (MIT), Eric Xia (MIT) and Andrea Zanette (Berkeley).

Biography

Dr. Martin Wainwright is currently the Cecil H. Green Professor at MIT, affiliated with Laboratory for Information and Decision Systems and MIT Statistics and Data Science Center (EECS and Mathematics). He received a Bachelor's degree in Mathematics from University of Waterloo, Canada, a Ph.D. degree in EECS from Massachusetts Institute of Technology (MIT), and is currently on leave from Statistics/EECS at UC Berkeley. His research interests include high-dimensional statistics, stochastic control and reinforcement learning, causal inference, and information-theoretic methods in statistics. Among other awards, he has received the COPSS Presidents' Award (2014) from the Joint Statistical Societies; the David Blackwell Lectureship (2017) and Medallion Lectureship (2013) from the Institute of Mathematical Statistics; and Best Paper awards from the IEEE Signal Processing Society and IEEE Information Theory Society. He was a Section Lecturer at the In-

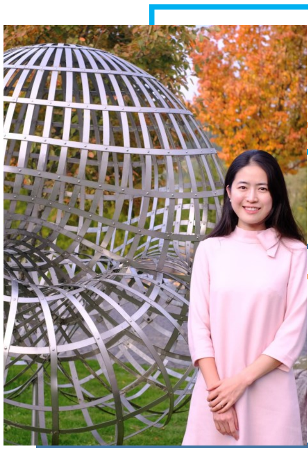

Dr. Yuting Wei
Assistant Professor of Statistics, Wharton School
University of Pennsylvania

"A Non-asymptotic Framework for Approximate Message Passing"

Abstract

Approximate message passing (AMP) emerges as an effective iterative paradigm for solving high-dimensional statistical problems. However, prior AMP theory --- which focused mostly on high-dimensional asymptotics --- fell short of predicting the AMP dynamics when the number of iterations surpasses $o(\log n/\log\log n)$ (with n the problem dimension). To address this inadequacy, this paper develops a non-asymptotic framework for understanding AMP in spiked matrix estimation. Built upon new decomposition of AMP updates and controllable residual terms, we lay out an analysis recipe to characterize the finite-sample behavior of AMP in the presence of an independent initialization, which is further generalized to allow for spectral initialization. As two concrete consequences of the proposed analysis recipe: (i) when solving Z_2 synchronization, we predict the behavior of spectrally initialized AMP for up to $O(n/\mathrm{poly}\log n)$ iterations, showing that the algorithm succeeds without the need of a subsequent refinement stage (as conjectured recently by Celentano et al.); (ii) we characterize the non-asymptotic behavior of AMP in sparse PCA (in the spiked Wigner model) for a broad range of signal-to-noise ratio.

Biography

Dr. Yuting Wei is currently an assistant professor in the Statistics and Data Science Department at the Wharton School, University of Pennsylvania. Prior to that, Yuting spent two years at Carnegie Mellon University as an assistant professor of statistics, and one year at Stanford University as a Stein Fellow. She received her Ph.D. in statistics at the University of California, Berkeley. She was the recipient of the 2022 NSF Career award, the Erich L. Lehmann Citation from the Berkeley statistics department in 2018. Her research interests include highdimensional and non-parametric statistics, statistical machine learning, and reinforcement learning

Dr. Hua Zhou

Professor, Dept. of Biostatistics

University of California, Los Angeles

Orthogonal Trace-sum Maximization: Applications, Local Algorithms, and Global Optimality”

Abstract

We study the problem of maximizing the sum of traces of matrix quadratic forms on a product of Stiefel manifolds. This orthogonal trace-sum maximization (OTSM) problem generalizes many interesting problems such as generalized canonical correlation analysis (CCA), Procrustes analysis, and cryo-electron microscopy of the Nobel prize fame. For these applications finding global solutions is highly desirable, but it has been unclear how to find even a stationary point, let alone test its global optimality. While the problem is nonconvex, the paper shows that its semidefinite programming relaxation solves the original nonconvex problems exactly with high probability, under an additive noise model with small noise in the order of O(m^{1/4}). In addition, it shows that the solution of a nonconvex algorithm considered in Won, Zhou, and Lange [SIAM J. Matrix Anal. Appl., 2 (2021), pp. 859-882] is also its global solution with high probability under similar conditions. These results can be considered as a generalization of existing results on phase synchronization.

Biography

Dr. Hua Zhou is a biostatistician interested in applying computational tools to solve problems arising from bioinformatics, genetics, neuroimaging, and electronic health records.

Dr. Hui Zou
Professor of Statistics
University of Minnesota

"High-dimensional Clustering via Latent Semiparametric Mixture Models"

Abstract

Cluster analysis is a fundamental task in machine learning. Several clustering algorithms have been extended to handle high-dimensional data by incorporating a sparsity constraint in the estimation of a mixture of Gaussian models. Though it makes some neat theoretical analysis possible, this type of approach is arguably restrictive for many applications. In this work we propose a novel latent variable transformation mixture model for clustering in which we assume that after some unknown monotone transformations the data follows a mixture of Gaussians. Under the assumption that the optimal clustering admits a sparsity structure, we develop a new clustering algorithm named CESME for high-dimensional clustering. The use of unspecified transformation makes the model far more flexible than the classical mixture of Gaussians. On the other hand, the transformation also brings quite a few technical challenges to the model estimation as well as the theoretical analysis of CESME. We offer a comprehensive analysis of CESME including identifiability, initialization, algorithmic convergence, and statistical guarantees on clustering. Extensive numerical study and real data analysis show that CESME outperforms the existing high-dimensional clustering algorithms.

Biography

Dr. Hui Zou received PhD in Statistics from Stanford University in 2005 and became a full professor of Statistics at University of Minnesota in 2014. His research interests include high-dimensional statistics, machine learning, and data science in general. Some of his well-known inventions include elastic net, adaptive lasso, sparse PCA, LLA for nonconvex regularization, composite quantile regression, rank-based inference of graphical model, and high-dimensional discriminant analysis. He is an elected Fellow of Institute of Mathematical Statistics and American Statistical Association. He has been Web of Science highly cited researcher in mathematics from 2014 to 2019. He received several other awards including NSF career award, IMS Tweedie award and the best paper award in 2019 ICCM.

Organizers

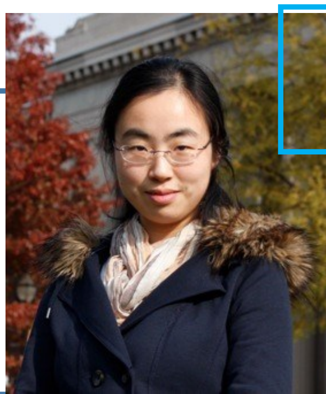

Dr. Zhihua Su
Associate Professor of Statistics
Department of Statistics
University of Florida

Dr. Aaron Molstad
Assistant Professor of Statistics
Department of Statistics
University of Florida

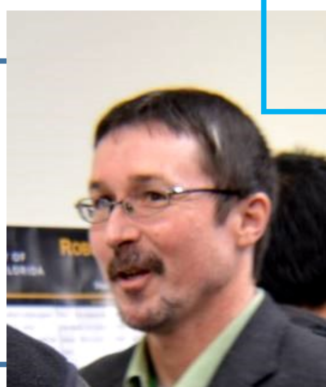

Dr. Michael Daniels
Professor, Chair and Andrew Banks Family Endowed Chair
Department of Statistics
University of Florida

Dr. George Michailidis
UFII Founding Director & Professor of Statistics
UF Informatics Institute
University of Florida

Thank you to our Sponsors!

UF | Department of Statistics

UF | Informatics Institute
UNIVERSITY *of* FLORIDA

UF Liberal Arts *and* Sciences

InfoTech®
| CONSULTING |