Simons Institute for the Theory of Computing

7th Annual Industry Day

Wednesday, October 20, 2021

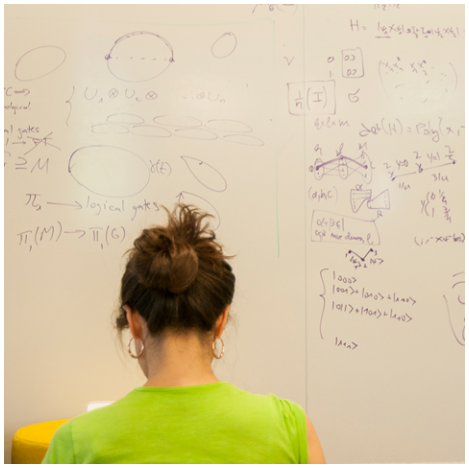

Table of Contents

Welcome

Dear friends,

We warmly welcome you to the seventh annual Simons Institute Industry Day.

Despite the ongoing global challenges, our community of scientists and scholars — hailing from industry, academia, and government — continues to engage actively in our distinctive research platform, whether virtually or in person.

Today's event is once again a chance for our research fellows and industry scientists to take the stage, to share their work, and to uncover new connections and synergies in research.

Our industry partnerships are vital to the relevance of our research. They represent a variety of industries, including computing software and hardware, social media, computing security, financial services, risk and reinsurance. They often provide a sounding board for thoughts and reveal new perspectives and directions for research. The theory of computing is critical and ever changing as the demand grows for technology at scale. We applaud our industry partners for supporting and being involved closely with core science. We see our shared pursuits in research as responsible actions. As the world becomes more complex and in some cases more fragile, together we can improve the human condition through joint discovery.

As the Simons Institute enters its 10th year of discovery, we invite you to remain closely connected as we announce anniversary activities, which will include a three-day, forward-looking symposium and celebratory gala at the end of May 2022.

Thank you all for your visionary partnership, and your dedication to this important work.

Yours,

Shafi Goldwasser, Director
Peter Bartlett, Associate Director
Prasad Raghavendra, Senior Scientist

Return to Top

Industry Day Agenda

8:45 a.m. Host asks participants to add their affiliations next to their names in Zoom

9:30 a.m. Introduction - (Prasad Raghavendra and Peter Bartlett) 10 mins

9:40 a.m. Lightning talks industry partners (9 @ 6 mins/each =) 54-60 mins

Algorand*, Tal Rabin*
Apple*, Audra McMillian*
Google*, Rina Panigrahy*
J.P. Morgan*, Alberto Pozanco Lancho*
Microsoft Research, *Jerry Li*
NTT Research*, Vipul Goyal*
PayPal, *Ignacio De Loizaga*
Roc360, *Elizabeth Barton*
VMware*, Udi Wieder*

10:40 a.m. Brief tutorial for Gather.Town and then networking break-out - 30 mins (via Gather.Town)

11:10 a.m. Lightning talks research fellows (10 @ 5 mins/each =) 50-60 mins

Learnability of hamiltonians from quantum many-body Gibbs states
Anurag Anshu
Confidence Intervals for Deep Learning Predictions
Stephen Bates
Perfect Sampling for Better Samples
Siddharth Bhandari
Proxy Convexity: A Unified Framework for the Analysis of Neural Networks Trained by Gradient Descent
Spencer Frei
An optimal algorithm for smooth strongly convex decentralized optimization
Adil Salim
On Maximal Advantage for Quantum Query Algorithms
Makrand Sinha
Robust Statistics and Fast Semidefinite Programming
Kevin Tian
Constrained optimization on Riemannian manifolds
Melanie Weber
Understanding Statistical-vs-Computational Tradeoffs via Low-Degree Polynomials
Alex Wein
Stateful Offline Contextual Policy Evaluation and Learning
Angela Zhou

12:10 p.m Program presentations (6 presentations @ 5 mins/each =) 30 mins

Computational Complexity of Statistical Inferencel: *Aug. 18–Dec. 17, 2021*
Guy Bresler
Geometric Methods in Optimization and Sampling: *Aug. 18–Dec. 17, 2021*

Philippe Rigollet
Causality: *Jan. 11–May 13, 2022*
Frederick Eberhardt
Learning and Games: *Jan. 11–May 13, 2022*
Vasilis Syrgkanis
Extended Reunion Lattices: Algorithms, Complexity, and Cryptography: *May 23–Jun. 24, 2022*
Vinod Vaikuntanathan
Extended Reunion: The Quantum Wave in Computing: *May 23–Jun. 24, 2022*
Umesh Vazirani

12:40 p.m. Final networking break-out - 50 mins (via Gather.Town)

Return to Top

Preparation for Industry Day 2021

Connecting

Zoom information: https://berkeley.zoom.us/j/95311601854
Gather.Town information: https://gather.town/app/cwMUdjOH6tyXMCuf/SimonsInstitute?spawnToken=znL1DlPepm92fp7m

Tech Support - day of event: Quelani Penland, Associate Event Coordinator
email: qpenland@berkeley.edu

Conference Format

The main event will be via Zoom meeting to allow participants to engage fully in the activities. Participants are asked to remain muted with their videos on during talks. The chat feature will be unrestricted between speakers and participants.

The networking portion of the event will be as follows.

1) *Informal Networking* (two sessions) via Gather.Town. Gather is a virtual room, which in this case will be set up as a large conferencing space with tables logoed for each industry partner, as shown in the screenshot below. Participants will be able to navigate the space as avatars, moving next to people with whom they wish to speak.

Introduction

Peter Bartlett, Simons Institute Associate Director.

Emcee

Prasad Raghavendra, Simons Institute Senior Scientist, will run the show and manage the talks by the industry partners, research fellows and program organizers.

Speakers

- **For connectivity tech support, Quelani Penland - email: qpenland@berkeley.edu**
- Slides and other visuals:
    - Program Organizers and Industry Partners will **manage their own visuals** and share their screens when they are speaking.
    - Research Fellows' visuals will be managed by the administrator, whom they will instruct to progress slides.
- Speakers will be allocated a precise amount of time:
    - Program Organizers - 6 minutes
    - Research Fellows - 4 minutes
    - Industry Partners - 10 minutes *(except Novi+Facebook: two speakers @ 5 mins each)*
- We suggest speakers **leave the final minute** for questions from the floor.
- **Questions will be fielded** by Prasad Raghavendra.
- If your **connection is faulty**, we may move down the queue of speakers until you are ready.
- Speakers may **sign up for a technical run-through** with Omeid Far (omiedf@berkeley.edu) on Wednesday, October 20 between 8:15 – 8:45 a.m. Please contact him directly to choose your time.

Recording Talks

- Speakers who have signed release forms will be recorded, and their talks will be available after the event on the Simons Institute YouTube channel. If you would like to be recorded but have not yet received a release form, please contact Amy Ambrose amyambrose@berkeley.edu

Managing Q & A

- Participants are asked to **remain muted during the talks**, and when called on, to unmute to ask their question.
- Participants may use the **raise hand function** or pose questions directly in the **chat**.
- At the end of each talk, Peter or Prasad will either call on one raised hand or select a question from the chat and ask the participant to **unmute to speak**.
- Participants are encouraged to continue asking questions in the networking sessions.

Points of Contact – Before & During the Event

Atiya Rashid Associate Visitor Services Coordinator arashid123@berkeley.edu	Drew Mason Information Systems Analyst dmason@berkeley.edu	Quelani Penland Associate Event Coordinator simonsevents@berkeley.edu	Amy Ambrose Sr. Development Director amyambrose@berkeley.edu +1.510.944.6674

Return to Top

Industry Partners

Algorand in a Nutshell

The Algorand Foundation is dedicated to fulfilling the global promise of blockchain technology. With a commitment to ensuring strong, lasting support for blockchain innovation, the Algorand Foundation fosters an open and transparent system where anyone can participate, develop, and innovate on the permissionless Algorand protocol guided by the vision of an inclusive economy.

Tal Rabin is a Professor of Computer Science at University of Pennsylvania and a member of the Algorand Foundation research team. Prior to that she had been at IBM Research for 23 years as a Distinguished Research Staff Member and the manager of the Cryptographic Research Group. She received her PhD from the Hebrew University in 1995.

Tal's research focuses on cryptography and, more specifically, on secure multiparty computation, threshold cryptography, and proactive security and recently adapting these technologies to the blockchain environment. She has initiated and organizes the Women in Theory Workshop, a biennial event for graduate students in Theory of Computer Science. Tal is currently the chair of the SIGACT Executive Board.

Tal is an ACM Fellow, an IACR (International Association of Cryptologic Research) Fellow and member of the American Academy of Arts and Sciences. Tal's work won the 30 year test of time award at STOC. She is the 2019 recipient of the RSA Award for Excellence in the Field of Mathematics. She was named by Forbes in 2018 as one of the Top 50 Women in Tech in the world. In 2014 Tal won the Anita Borg Women of Vision Award winner for Innovation and was ranked by Business Insider as the #4 on the 22 Most Powerful Women Engineers.

Return to Top

Apple

Privacy Amplification by Shuffling

Differential privacy (DP) is a model of privacy-preserving machine learning that has garnered significant interest in recent years due to its rigorous privacy guarantees. An algorithm differentially private if the output is stable under small changes in the input database. While DP has been adopted in a variety of applications, most applications of DP in industry actually satisfy a stronger notion called local differential privacy. In local differential privacy data subjects perturb their data before it reaches the data analyst. While this requires less trust, it comes a substantial cost to accuracy. Recent work of Erlingsson, Feldman, Mironov, Raghunathan, Talwar, and Thakurta [EFMRTT19] demonstrated that random shuffling amplifies differential privacy guarantees of locally randomized data. Such amplification implies substantially stronger privacy guarantees for systems in which data is contributed anonymously [BEMMRLRKTS17] and has led to significant interest in the shuffle model of privacy [CSUZZ19, EFMRTT19]. In this talk, we will discuss a new result on privacy amplification by shuffling, which achieves the asymptotically optimal dependence in the local privacy parameter. Our result is based on a new proof strategy which is simpler than previous approaches, and extends to a lightly weaker notion known as approximate differential privacy with nearly the same guarantees. Based on joint work with Vitaly Feldman and Kunal Talwar (https://arxiv.org/abs/2012.12803)

Audra McMillan is a research scientist at Apple. Her research focuses on privacy-preserving data analysis. Specifically, she is interested in performing fundamental statistical tasks under differential privacy, and it's variations. Her work has applications not only to privacy but also to robust data analysis and controlling false discovery rates. She received her PhD from the Department of Mathematics at the University of Michigan and prior to joining Apple was a postdoc split between Boston University and Northeastern University.

Return to Top

What could be the data structures of the Mind?

What is a reasonable architecture for an algorithmic view of the mind? Is it akin to a single giant deep network or is it more like several small modules connected by some graph? How is memory captured -- is it some lookup table? Take a simple event like meeting someone over coffee -- how would your mind remember who the person was, what was discussed? Such information needs to be organized and indexed in a way so that it can be quickly accessed in future if say I met the same person again.

Rina Panigrahy is a research scientist at Google specializing in applied and theoretical algorithms in areas such as deep learning, high dimensional search, hashing, sketching, streaming, prediction and graph analysis with engineering and research impact covering over 75 publications and 50 patents. His masters thesis work at MIT was used in founding Akamai Technologies. Before Google he has held research and engineering positions at Microsoft (principal researcher) and Cisco Systems. He obtained his Ph.D. in Algorithms from Stanford, and did his undergrad from IIT Mumbai after securing the top rank at the IIT-JEE entrance examination all over India. He is a recipient of a Gold medal at the International Math Olympiad and a winner of several best paper awards.

Return to Top

Proving Security of Cryptographic Protocols using Automated Planning

Lately, financial institutions have started implementing cryptographic protocols for various privacy-related use cases. One of the key aspects in applying these protocols consists of proving that they are secure. This talk presents ongoing work on using Automated Planning for that task.

Alberto Pozanco is a Research Scientist working at the JP Morgan's AI Research team lead by Manuela Veloso. Before that, he completed his PhD on "goal reasoning for autonomous agents" at Universidad Carlos III de Madrid. The main focus of his current work includes the use of automated planning and optimization techniques to solve real world problems at JP Morgan.

Return to Top

Does your quantum computer have quantum memory?

Jerry is a Researcher at Microsoft Research Redmond. Prior to that, he finished his PhD at MIT under the supervision of Ankur Moitra, and completed a postdoc at the Simons Institute as the VMWare Simons Fellow. His work spans a number of areas in machine learning theory, with an emphasis on robust machine learning, and high dimensional statistics. For his PhD work on the former he was awarded the George Sprowls award for the best thesis in CS at MIT.

Return to Top

Post-Quantum Secure Multi-Party Computation

Secure cryptographic multi-party computation (MPC) deals with the setting where there are multiple parties each holding a private input. The parties wish to compute a joint function of these inputs without revealing their individual inputs to each other. MPC is a very general tool for computing on private data. Introduced over 3 decades ago, MPC has led to an incredibly successful line of research.

In this work, I will talk about quantum and post-quantum MPC. With the recent boom in quantum cryptography, studying quantum MPC is natural. However, quantum MPC comes with a series of beautiful novel technical challenges which I will talk about.

Vipul Goyal is a Senior Scientist at NTT Research, and an Associate Professor of Computer Science at CMU. Previously, he was a researcher in the Cryptography and Complexity group at Microsoft Research. He received his PhD in Computer Science from University of California, Los Angeles in Dec 2009.

Dr. Goyal is a winner of several honors including a 2016 ACM CCS test of time award, a Microsoft Research graduate fellowship, and, a Google outstanding graduate student award. He was named to the Forbes magazine 30 under 30 list of people changing science and healthcare in 2013. His research has received media coverage at popular science publications such as MIT technology reviews, Slashdot, and, Nature news. He has served on program committees of conferences such as Crypto, Eurocrypt, STOC, and ACM CCS.

Return to Top

PayPal Computing Research Areas of Interest

High level description of PayPal areas of interest around the theory of computing, with a focus on the ongoing research around artificial intelligence, quantum computing, and crypto currency technology.

Ignacio is a MSc in Telecommunications Engineering and an MBA, with more than twenty years of experience in private technology companies. He has played many different roles, ranging from software development in the beginnings, all the way chief product officer, with episodes of sales and business development. His areas of expertise include blockchain technologies, data internetworking. and information security. Among other hats, he is now leading technological partnerships with the academia for PayPal's blockchain, crypto and digital currencies business unit.

Return to Top

Refreshing the US Housing Stock: A Technology-Based Approach

Homes in the US are aging: 80% are more than 20 years old and almost 40% are more than 50 years old. Green technology has trickled into only a tiny fraction of US houses. At Roc 360, we are exploring the use of technology to bring modernization and efficiency to the heavily fragmented industry of home renovation. I will survey some of the issues and applications related to the use of AI and optimization.

Betsy Barton is a former Astronomer and Associate Professor at the University of California, Irvine. She spent 7 years in quantitative finance before leading a data science team as a Senior Director at Walmart. She has recently been consulting with Roc 360 to shape their data science efforts.

Return to Top

Omnipredictors

Loss minimization is a dominant paradigm in machine learning, where a predictor is trained to minimize some loss function that depends on an uncertain event (e.g., "will it rain tomorrow?"). Different loss functions imply different learning algorithms and, at times, very different predictors. While widespread and appealing, a clear drawback of this approach is that the loss function may not be known at the time of learning, requiring the algorithm to use a best-guess loss function. We introduce the notion of an (L,C)-omnipredictor, which could be used to optimize any loss in a family L. Once the loss function is set, the outputs of the predictor can be post-processed (a simple univariate data-independent transformation of individual predictions) to do well compared with any hypothesis from the class C. The post processing is essentially what one would perform if the outputs of the predictor were true probabilities of the uncertain events. In a sense, omnipredictors extract all the predictive power from the class C, irrespective of the loss function in L. We show that such "loss-oblivious" learning is feasible through a connection to multi- calibration, a notion introduced in the context of algorithmic fairness.

Udi Wieder is a Senior Researcher at VMware. His research interests are in algorithms for machine learning, statistical analysis of big data and visualizations. He has done considerable work on Data Structures, Randomized Algorithms, Scheduling and Load Balancing. On the latter topic, he produced a survey. Udi holds a PhD from the Weizmann Institute of Science in Israel. His PhD advisor was Moni Naor and his M.Sc. advisor was Uri Feige. Prior to joining VMware in 2014, Udi was a researcher at MSR Silicon Valley for nine years.

Return to Top

Research Fellows

Learnability of hamiltonians from quantum many-body Gibbs states

We consider the problem of learning the Hamiltonian of a quantum many-body system given samples from its Gibbs (thermal) state. The classical analog of this problem, known as learning graphical models or Boltzmann machines, is a well-studied question in machine learning and statistics. This lightning talk will describe a sample-efficient algorithm for the quantum Hamiltonian learning problem at all constant temperatures. In particular, we prove that polynomially many samples in the number of particles (qudits) are necessary and sufficient for learning the parameters of a spatially local Hamiltonian in l_2-norm. Our main contribution is in establishing the strong convexity of the log-partition function of quantum many-body systems.

Anurag Anshu

Quantum Postdoctoral Fellow, UC Berkeley

Dates of Visit: Jan. 1, 2021 – Aug. 31, 2022

Anurag is a postdoctoral researcher at the University of California, Berkeley. Prior to this, he was a joint postdoctoral researcher at the Institute for Quantum Computing and the Perimeter Institute for Theoretical Physics, Waterloo. He obtained his PhD from the Centre for Quantum Technologies, National University of Singapore. He is interested in quantum complexity theory, quantum many-body physics, quantum communication and quantum learning theory. A unifying theme of his research in these areas is the interplay between computation and physics.

Return to Top

Confidence Intervals for Deep Learning Predictions

We introduce Learn then Test, a framework for calibrating machine learning models so that their predictions satisfy explicit, finite-sample statistical guarantees regardless of the underlying model and (unknown) data-generating distribution. The framework addresses confidence sets for classification/regression, false discovery rate control in multi-label classification, intersection-over-union control in instance segmentation, and the simultaneous control of the type-1 error of outlier detection and confidence set coverage in classification or regression. We demonstrate these techniques in computer vision examples.

Stephen Bates
Machine Learning Postdoctoral Fellow, UC Berkeley
Dates of Visit: Aug. 18, 2021 – Aug. 17, 2022

Stephen Bates is a postdoctoral researcher with Michael I. Jordan in the Statistics and EECS departments at UC Berkeley. He works on developing methods to analyze modern scientific data sets, leveraging sophisticated black box models while providing rigorous statistical guarantees. Specifically, Stephen works on problems in high-dimensional statistics (especially false discovery rate control), statistical machine learning, conformal prediction and causal inference.

Previously, Stephen completed his PhD in the Stanford Department of Statistics advised by Emmanuel Candès. His thesis introduced methods for conditional independence testing and false discovery rate control in genomics, and he received the Ric Weiland Graduate Fellowship and the Theodore W. Anderson Theory of Statistics Dissertation Award for this work. Before his PhD, he studied statistics and mathematics at Harvard University.

Return to Top

Perfect Sampling for Better Samples

"In this talk we will look at the basic ideas/challenges of perfect sampling and how it might be helpful in sampling tasks where ""mixing time"" bounds are not known.

Consider a Markov Chain such as Glauber Dynamics or some other mechanism for sampling from a given target distribution. Usually, approximate sampling algorithms run a Markov Chain for a given number of steps (or until a certain statistic is met) and require proving mixing time bounds to ensure the quality of the output sample which is a hard task.
In absence of such guarantees the output distribution may be far away from the
target distribution, thereby rendering the output unreliable for statistical applications.
In contrast, perfect sampling algorithms are typically designed in such a way that the output produced when the algorithm stops is distributed exactly according to the target distribution. One might be unable to formally guarantee that the expected running time is small; yet the quality of the output is never in question.

Siddharth Bhandari
Simons-Berkeley Postdoctoral Fellow, UC Berkeley
Dates of Visit: Aug. 1, 2021 – Jul. 31, 2023

Siddharth Sandipkumar Bhandari completed his PhD at the School of Technology and Computer Science at the Tata Institute of Fundamental Research, Mumbai, under the guidance of Prahladh Harsha. He obtained a BTech degree in Mechanical Engineering from IIT, Kharagpur and a MSc degree in Computer Science from CMI.

Siddharth's research interests lie broadly in the field of randomized computation. He has dabbled in zero-error information theory, coding theory, sampling algorithms and theoretical guarantees for generative networks.

Return to Top

Proxy Convexity: A Unified Framework for the Analysis of Neural Networks Trained by Gradient Descent

In this work, we propose a unified non-convex optimization framework for the analysis of neural network training. We introduce the notions of proxy convexity and proxy Polyak-Lojasiewicz (PL) inequalities, which are satisfied if the original objective function induces a proxy objective function that is implicitly minimized when using gradient methods. We show that stochastic gradient descent (SGD) on objectives satisfying proxy convexity or the proxy PL inequality leads to efficient guarantees for proxy objective functions. We further show that many existing guarantees for neural networks trained by gradient descent can be unified through proxy convexity and proxy PL inequalities. This is joint work with Quanquan Gu (UCLA) and will appear at NeurIPS 2021.

Spencer Frei
Postdoctoral Researcher, UCLA
Machine Learning Postdoctoral Fellow, UC Berkeley
Dates of Visit: Aug. 15, 2021 – Aug. 16, 2022

Spencer Frei is currently finishing his PhD in Statistics at UCLA under the supervision of Quanquan Gu and Ying Nian Wu. He received his MSc in Mathematics from the University of British Columbia, Vancouver, and his BSc in Mathematics from McGill University. He is interested in the theory of deep learning and statistical learning theory.

For the 2021-2022 academic year, he will be a postdoctoral fellow at UC Berkeley, working under the mentorship of Peter Bartlett and Bin Yu as a part of the NSF/Simons program Collaboration on the Theoretical Foundations of Deep Learning.

Return to Top

An optimal algorithm for smooth strongly convex decentralized optimization

We consider the task of decentralized minimization of the sum of smooth strongly convex functions stored across the nodes of a network. For this problem, lower bounds on the number of gradient computations and the number of communication rounds required to achieve a given accuracy have recently been proven. We propose the first optimal algorithm for this problem: its computation and communication complexities match the lower bounds. Our approach relies on viewing the proposed algorithm as an accelerated primal dual algorithm.

Adil Salim
Postdoctoral Researcher, KAUST
Google Research Fellow
Program: Geometric Methods in Optimization and Sampling
Dates of Visit: Sep. 20 – Dec. 31, 2021

Adil Salim earned his Master degrees in Data Science from ENSAE ParisTech and in Probability Theory from Paris XI University in France in 2015. He received his Ph.D. in Machine Learning from Télécom ParisTech in France in 2018. Adil Salim is mainly interested in machine learning, optimization and sampling algorithms. From 2019 to 2021, he was a Postdoctoral Research Fellow working with Professor Peter Richtarik at KAUST. In Fall 2021, Adil will be a Simons-Berkeley Research Fellow, and he is expected to join Microsoft Research as a Senior Researcher in early 2022.

Return to Top

On Maximal Advantage for Quantum Query Algorithms

The blackbox or the query model has turned out to be a powerful abstraction to give rigorous guarantees towards understanding the relative power of quantum vs classical algorithms. In this model, the runtime of the algorithm is modeled by the number of queries to the input that are needed to solve the task as hand.

We show that a task called k-fold Forrelation gives the maximal advantage for quantum algorithms in this model, thus confirming a conjecture made by Aaronson and Ambainis. In particular, we show that this task can be solved with only constant number of quantum queries, but classically requires almost linearly many queries, which is essentially the maximum runtime in this model.

Makrand Sinha
Simons-Berkeley Postdoctoral Fellow, UC Berkeley
Dates of Visit: Jul. 1, 2021 – Jun. 30, 2023

Makrand Sinha is a Simons-Berkeley postdoctoral fellow at the Simons Institute at UC Berkeley. He received a PhD in 2018 from the University of Washington and subsequently completed an appointment as a postdoctoral researcher at CWI Amsterdam. His research interests lie in the foundations of quantum and classical computation and optimization, and specifically in understanding the relative power of quantum vs classical algorithms and communication protocols, understanding limitations of various approaches in optimization such as Linear or Semidefinite Programs, and designing algorithms for various optimization problems.

Return to Top

Robust Statistics and Fast Semidefinite Programming

Algorithmic robust statistics, the design of statistical estimators or learners tolerant to arbitrary adversarially-corrupted points, is a rapidly-developing area with many applications in classical settings as well as modern machine learning tasks, e.g. defenses against data poisoning. However, while many recent exciting results in algorithmic robust statistics have yielded polynomial-time methods, the resulting runtimes are oftentimes still somewhat impractical. In this talk, we briefly overview a number of advances which combine insights from the robust statistics literature with lightweight, custom-designed optimization machinery to yield robust estimators and learners computable in nearly-linear time; applications include mean estimation, stochastic optimization tasks, PCA, and clustering.

Kevin Tian
PhD candidate, Stanford University
VMware Research Fellow
Program: Geometric Methods in Optimization and Sampling
Dates of Visit: Aug. 18 – Dec. 31, 2022

Kevin is a final-year Ph.D. student with the Theory Group in Computer Science at Stanford, advised by Aaron Sidford. He completed undergraduate studies in Computer Science and Math at MIT from 2012-2015. He is broadly interested in fundamental algorithmic problems in modern data science, often at the intersection of continuous optimization and high-dimensional statistics.

Return to Top

Constrained optimization on Riemannian manifolds

Many applications involve non-Euclidean data, such as graphs, strings or matrices. In such cases, exploiting Riemannian geometry can deliver algorithms that are computationally superior to standard nonlinear programming approaches. This observation has resulted in an increasing interest in Riemannian methods in the optimization and machine learning community. In this talk, we consider the problem of optimizing a function on a Riemannian manifold subject to convex constraints. We will introduce a class of algorithms (Riemannian Frank-Wolfe methods) for solving such problems and see examples, where a Riemannian approach is superior to its Euclidean counterpart.

Melanie Weber
PhD candidate, Princeton University
Research Fellow
Program: Geometric Methods in Optimization and Sampling
Dates of Visit: Aug. 30 – Dec. 3, 2021

Melanie is a PhD student at Princeton University, where she is very fortunate to be advised by Charles Fefferman. Her research focuses on understanding the geometric features of data mathematically and on developing machine learning methods that utilize this knowledge.

Return to Top

Understanding Statistical-vs-Computational Tradeoffs via Low-Degree Polynomials

What is the best way to extract a hidden signal from a large noisy dataset? This question has been well studied in various statistical models such as sparse PCA in spiked matrix models, community detection in the stochastic block model, tensor PCA in the spiked tensor model, and more. Many of these models exhibit a striking gap between what can be achieved statistically (via a brute force algorithm) and what can be achieved with known polynomial-time algorithms. I will give an overview of the "low-degree polynomial framework", which is a method for predicting the limitations of poly-time algorithms and giving rigorous evidence for average-case computational hardness of statistical problems.

Alex Wein
Postdoctoral Researcher, New York University
Research Fellow
Program: Computational Complexity of Statistical Inference
Dates of Visit: Aug. 19 – Dec. 31, 2022

Alex Wein is a Courant Instructor at the Math Department in NYU's Courant Institute. My research interests include: statistical and computational limits of high-dimensional inference problems; connections to statistical physics; group actions and invariant theory; low-degree polynomials as a proxy for computational tractability.

Return to Top

Stateful Offline Contextual Policy Evaluation and Learning

We study off-policy evaluation and learning from sequential data in a structured class of Markov decision processes that arise from repeated interactions with an exogenous sequence of arrivals with contexts, which generate unknown individual-level responses to agent actions that induce known transitions. This is a relevant model, for example, for dynamic personalized pricing and other operations management problems in the presence of potentially high-dimensional user types. The individual-level response is not causally affected by the state variable. In this setting, we adapt doubly-robust estimation in the single-timestep setting to the sequential setting so that a state-dependent policy can be learned even from a single timestep's worth of data. We introduce a \textit{marginal MDP} model and study an algorithm for off-policy learning, which can be viewed as fitted value iteration in the marginal MDP. We also provide structural results on when errors in the response model leads to the persistence, rather than attenuation, of error over time. In simulations, we show that the advantages of doubly-robust estimation in the single time-step setting, via unbiased and lower-variance estimation, can directly translate to improved out-of-sample policy performance. This structure-specific analysis sheds light on the underlying structure on a class of problems, operations research/management problems, often heralded as a real-world domain for offline RL, which are in fact qualitatively easier.

Angela Zhou
Machine Learning Postdoctoral Fellow, UC Berkeley
Dates of Visit: Jul. 1, 2021 – Jun. 30, 2022

Angela Zhou is interested in developing prescriptive analytics with theoretical guarantees for data-driven decision-making. Currently, she is working on leveraging causal inference and machine learning as a language for prescriptive analytics, making robust recommendations for action in view of fundamental practical challenges in observational/operational data. Her work emphasizes credibility as a form of reliability, developing robust inferential procedures subject to analyst-tunable violations of assumptions. More broadly, Angela is interested in the interplay of statistics and optimization for decision-making, with applications to e-commerce, healthcare and policy.

Return to Top

Program Presentations

Computational Complexity of Statistical Inference

Aug. 18 – Dec. 17, 2021

https://simons.berkeley.edu/programs/si2021

The two basic lines of inquiry in statistical inference have long been: (i) to determine fundamental statistical (i.e., information-theoretic) limits; and (ii) to find efficient algorithms achieving these limits. However, for many structured inference problems, it is not clear if statistical optimality is compatible with efficient computation. Statistically optimal estimators often entail an infeasible exhaustive search. Conversely, for many settings the computationally efficient algorithms we know are statistically suboptimal, requiring higher signal strength or more data than is information-theoretically necessary. This phenomenon is both fascinating and unsettling. It suggests that the information-theoretic limit on the signal-to-noise ratio (or the amount of data) for these problems, as studied since the beginning of mathematical statistics, is not the practically relevant benchmark for modem high-dimensional settings. Instead, the practically relevant benchmark is the fundamental statistical limit for *computationally efficient* algorithms.

By now dozens of fundamental high-dimensional statistical estimation problems are conjectured to have different computational and statistical limits. These problems (for example, sparse linear regression or sparse phase retrieval) are ubiquitous in practice and well-studied theoretically, yet the central mysteries remain: What are the fundamental data limits for computationally efficient algorithms? How do we find optimal efficient algorithms? At a more basic level, are these statistical-computational gaps in various problems appearing for a common reason? Is there hope of building a widely applicable theory describing and explaining statistical-computational trade-offs?

The objective of the program is to advance the methodology for reasoning about the computational complexity of statistical estimation. Over the last decade several disparate communities and lines of work have begun to make progress on these questions. This program aims to stimulate work towards developing a deeper understanding and building a coherent theory by forming new collaborations between researchers in complexity theory, algorithms, statistics, learning theory, probability, and information theory.

Guy Bresler is associate professor in the Department of Electrical Engineering and Computer Science at Massachusetts Institute of Technology, and a member of LIDS, Center for Statistics, and IDSS. Previously, he was a postdoc at MIT. Before that Guy received his PhD from the EECS Department at the University of California Berkeley. His undergraduate degree is from the University of Illinois at Urbana-Champaign. In the last several years, his research has focused on the interface between computation and statistics with the aim of understanding the relationship between combinatorial structure and computational tractability of high-dimensional inference.

Return to Top

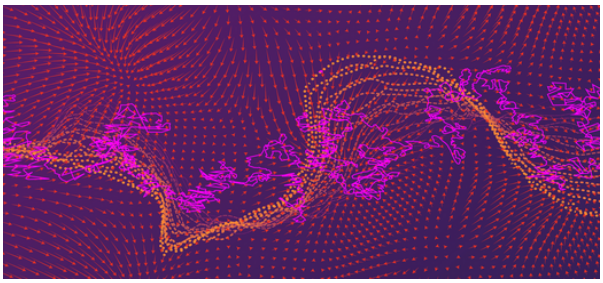

Geometric Methods in Optimization and Sampling
Aug. 18 – Dec 17, 2021
https://simons.berkeley.edu/programs/gmos2021

Optimization and sampling are two of the most important mathematical topics at the interface of data science and computation. The two questions are, in fact, connected mathematically through a powerful framework articulated around the geometry of probability distributions. The geometric toolbox that underlies optimization and sampling was initiated in the study of partial differential equations (PDEs) and has evolved into different mathematical disciplines: probability, calculus of variations, analysis and geometry. While connections are slowly beginning to percolate across disciplines, this program is aimed to be a catalyst for new and interdisciplinary ideas using a principled and unified approach to optimization and sampling.

A central goal of this program is to develop and promote a geometric approach to various computational problems in sampling, optimization, and PDEs. For example, the geometry of Optimal Transport has been instrumental to establish fruitful connections between diffusion processes, gradient flows, and diffusive PDEs by eliciting hidden convexity. This success calls for a versatile toolbox to tackle algorithmic questions arising in sampling, optimization, and particle methods for solving PDEs by leveraging the hidden geometric structure of each problem in a systematic way. Moreover, in a large class of problems this geometric structure is supplemented by additional symmetries or other algebraic structures that can be exploited to design better algorithms.

These recent connections between sampling, optimization, and PDEs have placed the fields in a unique position for mutual impact. This program aims at bringing together researchers from various backgrounds to tackle these challenging problems using a unified approach by focusing on the following aspects:

- Sampling as an optimization problem
- Geometry and optimal transport
- The PDE perspective on sampling and optimization
- Eliciting convexity via geometry in sampling and optimization
- The interplay of algebra and geometry in optimization

Philippe Rigollet is a professor of statistics at Massachusetts Institute of Technology. He works at the intersection of statistics, machine learning, and optimization, focusing primarily on the design and analysis of statistical methods for high-dimensional problems. His recent research focuses on the statistical limitations of learning under computational constraints. At the University of Paris VI, Rigollet earned a BS in statistics in 2001, a BS in applied mathematics in 2002, and a PhD in mathematical statistics in 2006 under the supervision of Alexandre Tsybakov. Prior to joining MIT in 2015, he held positions as a visiting assistant professor at the Georgia Institute of Technology, and assistant professor at Princeton University.

Return to Top

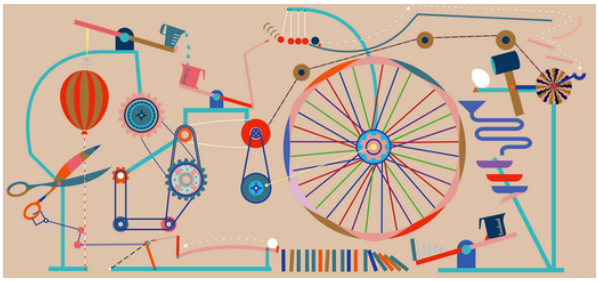

Causality

Jan. 11 – May 13, 2022

https://simons.berkeley.edu/programs/Causality2022

This program aims to integrate advances and techniques from theoretical computer science into methods for causal inference and discovery.

Although attempts to characterize causal relations can be found in some of the oldest written records, the history of the usage of causal concepts within scientific discussions over the past 100 years has been rocky, varying from the outright denial of any role of causality in mature scientific theories to a disingenuous usage of ambiguous terms that obscure the role of cause and effect (e.g., "link," "connection," etc.).

A substantive development of new formal approaches to causality in the 1970s and 1980s precipitated a change in attitude toward the scientific investigation of causal questions. The change was led by the development of two largely intertranslatable mathematical frameworks: the potential outcome framework and the causal graphical models framework. These frameworks integrated three concepts central to the notion of causation: (i) the connection between the underlying causal relations and observed data; (ii) the difference that interventions can make to a causal system; and (iii) counterfactual statements about a system. All of these aspects of causality play a central role in scientific testing and explanation, often constituting the goal of scientific inquiry itself.

The mathematization of questions of causality has resulted in the development of inference techniques and learning methods to infer causal relations from data. These formal approaches are now starting to spread throughout the applied sciences, where just about any field of study is seeing a renewed and explicit interest in tackling causality.

Broad application of these theoretical frameworks in scientific domains requires not only conceptual clarity and "in principle" methods, but a detailed understanding of how the methods behave in practice, how to scale and approximate the ideally desired computations, and how to optimize methods for the particular constraints present in a domain.

This program will bring together theoretical and applied researchers from a broad variety of domains with the goal of understanding the complexity, optimizations, and possible approximation regimes required to turn the methods of causal inference into a broadly applicable scientific toolbox.

Frederick Eberhardt is professor of philosophy in the Division of the Humanities and Social Sciences at the California Institute of Technology. Before coming to Caltech he was assistant professor in the philosophy-neuroscience-psychology (PNP) program and the Department of Philosophy at Washington University in St. Louis. In 2011 he had a two-year research leave to work on causal discovery methods at Carnegie Mellon University with a grant from the James S. McDonnell Foundation. Before going to St. Louis, he was a McDonnell postdoc at the Institute of Cognitive and Brain Sciences at the University of California Berkeley. Frederick completed his PhD in the Philosophy Department at Carnegie Mellon University. His research interests lie at the formal end of the philosophy of science, the machine learning end of statistics and computer science, and the learning and modeling end of psychology and cognitive science. His work has focused primarily on methods for causal discovery from statistical data, the use of experiments in causal discovery, the integration of causal inferences from different data sets, and the philosophical issues at the foundations of causality and probability. He has done some work on computational models in cognitive science and some historical work on the philosophy of Hans Reichenbach, especially his frequentist interpretation of probability.

Return to Top

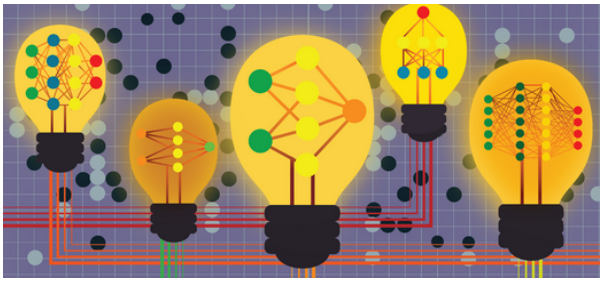

Learning and Games

Jan. 11 – May 13, 2022

https://simons.berkeley.edu/programs/games2022

The intersection of learning theory, game theory and mechanism design is becoming increasingly relevant: i) data input to machine learning algorithms are either owned or generated by self-interested parties, ii) machine learning is used to optimize economic systems (e.g. auction platforms) or to learn how to optimally act in strategic settings, iii) machine learning models used in critical systems are becoming prone to adversarial attacks, iv) several machine learning approaches can be framed as finding the equilibrium of a game, as opposed to the minimizer of an objective function. The theoretical foundations of these problems lie at the intersection of learning theory, game theory and mechanism design.

Already, online learning and game theory have played a key role in some landmark advances in Machine Learning. Online learning has provided some of the most successful optimization methods used in training large-scale deep neural networks. Game-theoretic modeling has enabled the design of Generative Adversarial Networks (GANS) and inspired approaches for training deep neural network classifiers that are robust to adversarial attacks. Finally, min-max tree search and regret minimization algorithms are central in solving Go and Texas Hold'em. More broadly, the world is moving towards the co-existence of multiple AIs that learn from their interaction, which might be collaborative, strategic or adversarial.

The objective of the program is to further advance the interaction between learning and games and rethink its mathematical foundations. The semester will study the foundations of: 1) Min-Max Optimization, 2) Multi-Agent Reinforcement Learning, 3) Dynamical Systems and Learning, 3) Behavioral Game Theory, 4) Econometrics and Learning, 5) Mechanism Design and Learning. Moreover, it will address practical challenges in application domains where these techniques seem most appropriate, such as: 1) generative adversarial networks, 2) adversarial robustness, 3) learning with humans in the loop, 4) learning as a

model of strategic behavior, 5) interactions of multiple learners.

The semester aims to bring together members of different communities, including machine learning, economics, operations research, theoretical computer science, and social computing.

Vasilis Syrgkanis is a Principal Researcher at Microsoft Research, New England, where he is also co-leading the project on Automated Learning and Intelligence for Causation and Economics (ALICE). His research lies at the intersection of theoretical computer science, machine learning and economics/econometrics. He received his Ph.D. in Computer Science from Cornell University, where he had the privilege to be advised by Eva Tardos and then spent two years as a postdoc researcher at Microsoft Research, NYC, as part of the Algorithmic Economics and the Machine Learning groups. He obtained his diploma in EECS at the National Technical University of Athens, Greece.

Return to Top

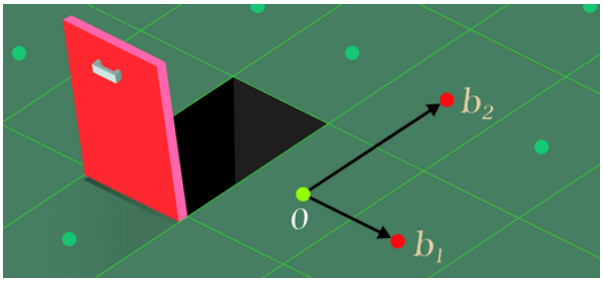

Extended Reunion — Lattices: Algorithms, Complexity, and Cryptography

May 23 – Jun. 24, 2022

https://simons.berkeley.edu/programs/extended-reunion-lattices2022

This extended reunion is for long-term participants in the program "Lattices: Algorithms, Complexity, and Cryptography" held in the Spring 2020 semester. It will provide an opportunity to meet old and new friends. Moreover, we hope that it will give everyone a chance to reflect on the progress made during the semester and since, and sketch which directions the field should go in the future. In an effort to keep things informal and to encourage open discussion, none of the activities will be recorded. Participation in the reunion is by invitation only.

The study of integer lattices serves as a bridge between number theory and geometry and has for centuries received the attention of illustrious mathematicians, including Lagrange, Gauss, Dirichlet, Hermite, and Minkowski. In computer science, lattices made a grand appearance in 1982 with the celebrated work of Lenstra, Lenstra, and Lovász, who developed the celebrated LLL algorithm to find short vectors in integer lattices. The role of lattices in cryptography has been equally, if not more, revolutionary and dramatic, with lattices first playing a destructive role as a potent tool for breaking cryptosystems and later serving as a new way to realize powerful and game-changing notions such as fully homomorphic encryption. These exciting developments over the last two decades have taken us on a journey through such diverse areas as quantum computation, learning theory, Fourier analysis, and algebraic number theory.

The promise of practical lattice-based cryptosystems together with their apparent quantum resistance is generating a tremendous amount of interest in deploying these schemes at Internet scale. However, before lattice cryptography goes live, we need major advances in understanding the hardness of lattice problems that underlie the security of these cryptosystems. Significant, groundbreaking progress on these questions requires a concerted effort by researchers from many areas: (algebraic) number theory, (quantum) algorithms, optimization, cryptography, and coding theory.

The goal of this extended reunion for the "Lattices: Algorithms, Complexity, and Cryptography" program is to bring back together experts in these areas to review progress, to continue to attack outstanding open questions, and to continue the exploration of connections among lattices, computer science, and mathematics. The need to thoroughly understand the computational landscape and cryptographic capabilities of lattice problems is greater now than ever given the possibility that secure communication on the Internet and secure collaboration on the cloud might soon be powered by lattices.

This program was supported in part by the Alfred P. Sloan Foundation.

Vinod Vaikuntanathan is a Steven and Renee Finn Career Development Assistant Professor of Computer Science at MIT. His main research interest is in the theory and practice of cryptography. He works on lattice-based cryptography, building advanced cryptographic primitives using integer lattices; leakage-resilient cryptography, defining and developing algorithms resilient against adversarial information leakage; and more recently, the theory and practice of computing on encrypted data, constructing powerful cryptographic objects such as fully homomorphic encryption and functional encryption. Vinod obtained his PhD from MIT, where he received a 2009 George M. Sprowls Award for the best MIT PhD thesis in computer science. He is also a recipient of the 2008 IBM Josef Raviv Postdoctoral Fellowship, the 2013 Alfred P. Sloan Research Fellowship, the 2014 Microsoft Faculty Fellowship, and a 2014 NSF CAREER award.

Return to Top

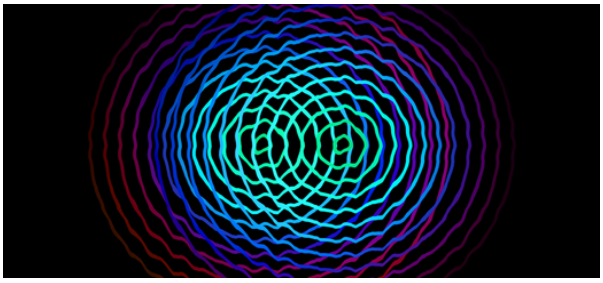

Extended Reunion: The Quantum Wave in Computing

May 23 – Jun. 24, 2022

https://simons.berkeley.edu/programs/extended-reunion-quantum2022

This extended reunion is for long-term participants in the program "The Quantum Wave in Computing" held in the Spring 2020 semester. It will provide an opportunity to meet old and new friends. Moreover, we hope that it will give everyone a chance to reflect on the progress made during the semester and since, and sketch which directions the field should go in the future. In an effort to keep things informal and to encourage open discussion, none of the activities will be recorded. Participation in the reunion is by invitation only.

Quantum computation is entering an exciting new period. Small- to medium-scale quantum computers are around the corner, and the biggest upcoming challenges are expected to be algorithmic. The first major challenge is identifying what kinds of computational tasks such computers will be useful for, given that for the foreseeable future, the scale issue will be compounded by minimal or nonexistent error correction. The second challenge is the testing of such devices, as direct simulation by classical computers is all but impossible and running a trace on the quantum program is ruled out by the basic laws of quantum physics.

Providing answers to these questions requires a collaboration between classical theoretical computer science and physics, chemistry, and mathematics. On the quantum algorithms front, there are great challenges in proposals for quantum machine learning and quantum annealing, with connections to classical machine learning, algorithms for low-rank matrix completion, and MCMC algorithms. The most promising algorithmic application for quantum computers in the long run, their "killer app," is expected to be the simulation of quantum systems and quantum chemistry.

On the theoretical computer science end, existing work on testing quantum devices has already led to exciting connections with the theory of interactive proof systems and theoretical cryptography. These connections will evolve into a beautiful and deep theory as the challenges in complexity theory, cryptography, and security raised by interactions with quantum devices are more systematically explored.

Umesh V. Vazirani is a Research Director for Quantum Computing Simons Institute, and the Roger A. Strauch Professor of Electrical Engineering and Computer Science at the University of California, Berkeley, and the director of the Berkeley Quantum Computation Center. He received his BTech in Computer Science from MIT in 1981 and his PhD in Computer Science from Berkeley in 1985. Vazirani's research interests include computational learning theory, combinatorial algorithms, computational complexity theory and quantum computing. He received the 2012 Fulkerson Prize (with Sanjeev Arora and Satish Rao) for his work on approximation algorithms for sparsest cut.

Return to Top

Thank you for your partnership!

Please join us for these coming programs:

Fall 2021

Computational Complexity of Statistical Inference	Aug. 18 – Dec. 17, 2021
Geometric Methods in Optimization and Sampling	Aug. 18 – Dec. 17, 2021

Spring 2022

Causality	Jan. 11 – May 13, 2022
Learning and Games	Jan. 11 – May 13, 2022

Summer 2022

Extended Reunion — Lattices: Algorithms, Complexity, and Cryptography	May 23 – Jun. 24, 2022
Extended Reunion: The Quantum Wave in Computing	May 23 – Jun. 24, 2022
Summer Cluster: Interpretable Machine Learning	Jun. 27 – Aug. 5, 2022
Computational Innovation and Data-Driven Biology	Jul. 5 – Aug. 5, 2022

Fall 2022

Data-Driven Decision Processes	Aug. 17 – Dec. 16, 2022
Graph Limits and Processes on Networks: From Epidemics to Misinformation	Aug. 17 – Dec. 16, 2022

Spring 2023

Meta-Complexity	Jan. 10 – May 12, 2023

Simons Institute contacts:

Shafi Goldwasser, Director, simonsdirector@berkeley.edu

Peter Bartlett, Associate Director, simonsassociatedirector@berkeley.edu

Christine Wang, Science and Outreach Administrator, chris24@berkeley.edu

Amy Ambrose, Sr. Development Director amyambrose@berkeley.edu