The 19th LIDS Student Conference

is held on Jan 30-31, 2014 in Room 32-155.

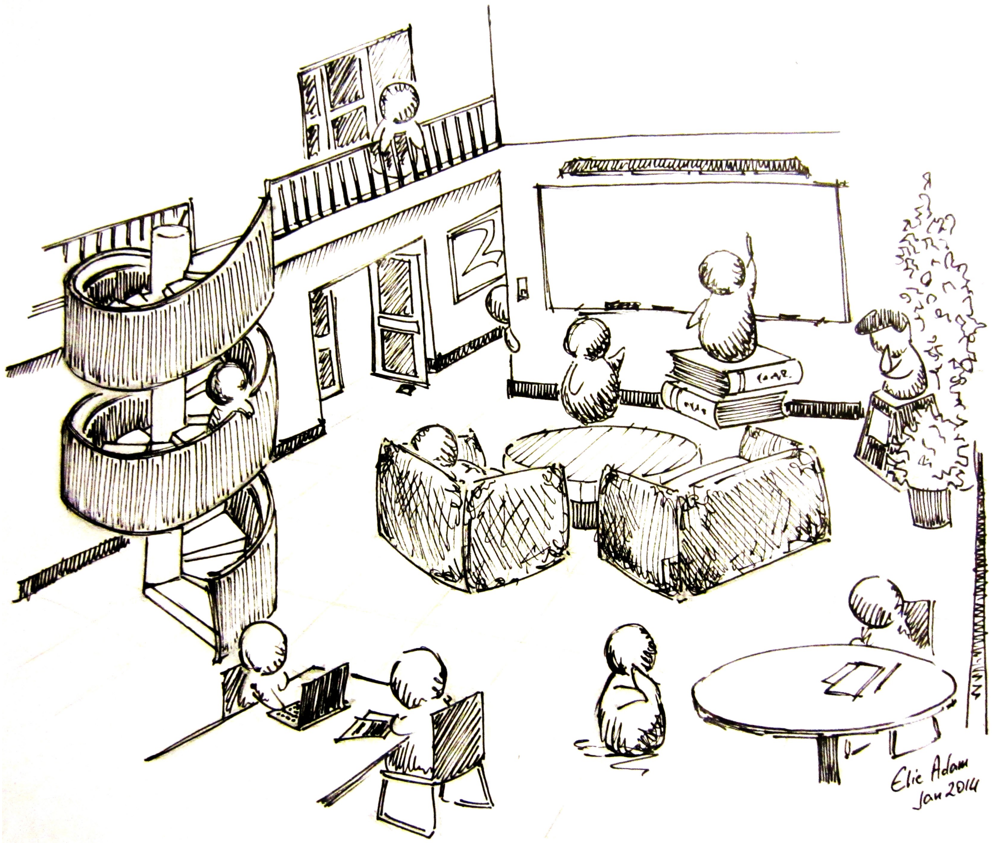

The Conference Program

more information on `lidsconf.mit.edu`

This event is sponsored by

DRAPER LABORATORY

Welcome to the 19th LIDS Student Conference.

On behalf of everyone at the Laboratory for Information and Decision Systems (LIDS) at MIT, it is our great pleasure to welcome you to the 2014 LIDS Student Conference. The two days of the conference feature student presentations, four plenary talks, and a panel discussion, showcasing many of the research themes and fields of study pursued at LIDS.

The conference is an excellent opportunity for students to present their work to both their peers and the wider community, gaining valuable feedback and presentation experience. The conference also provides a overview of the work done at LIDS, serving to foster communication and collaboration between research groups within LIDS and across MIT.

We are fortunate to have with four distinguished plenary speakers:

Ali Jadbabaie
Alfred Fitler Moore Professor of Network Science, University of Pennsylvania.

Radhika Nagpal
Fred Kavli Professor of Computer Science, Harvard University.

Anant Sahai
Associate Professor of Electrical Engineering and Computer Science, UC Berkeley.

Gerald Jay Sussman
Panasonic Professor of Electrical Engineering, MIT

The entire conference will take place in the Stata Center (building 32) at MIT. All plenary and student presentations will be in room 32-155, the panel discussion will be in room 32-124, the student social will be held in the 6th floor lounge in LIDS, and the banquet will be held in the R&D Common area on the 4th floor of the Stata Center.

We would like to thank Draper Laboratory for sponsoring the conference, and Prof. Peter Falb for the additional support.

Finally, we would like to acknowledge those whose help has been invaluable in planning this conference and making it happen: Jennifer Donovan, Lynne Dell, and Brian Jones at LIDS headquarters, and our student and plenary speakers and our panelists for their willingness to tell us about their work and ideas! We would like to thank Draper Laboratory for their ongoing sponsorship of this event over many years, and Peter Falb for providing additional support. Finally we would like to thank the attendees for making this event a success.

The 19th LIDS Student Conference Committee,

Co-chairs: Elie Adam and James Saunderson

Kimon Drakopoulos, Hamza Fawzi, Ali Kazerani, Christina Lee, Ying Liu, Marzieh Parandehgheibi, Dhruv Parthasarathy, Frank Permenter, Luis Reyes-Castro, Shreya Saxena, Omer Tanovic.

Thursday, Jan 30.

08:30 - 09:30 **Breakfast and Registration**

09:30 - 09:45 **Opening remarks**

09:45 - 10:45 **Plenary Talk**

Prof. Anant Sahai. Glimpses into an Information Theory for Control

10:45 - 11:00 **Coffee break**

11:00 - 12:00 **Student Session on Communication Systems**

Matt Johnston. Opportunistic Scheduling with Limited Channel State Information: An Application of Causal Rate Distortion

Tianheng Wang. On the Stability of Scheduling Algorithms for Network Navigation

Austin Collins. What To Transmit From Your Many Antennas

Abhishek Sinha. Throughput Optimal Algorithms for Multicast in a Wireless Network

12:00 - 01:30 **Lunch break**

01:30 - 02:30 **Student Session on Inference and Learning I**

Matt Johnson. Analyzing Hogwild Parallel Gaussian Gibbs Sampling

Ying Liu. Learning Gaussian Graphical Models with Small Feedback Vertex Sets: Observed or Latent

Dan Levine. Focused Active Inference

02:30 - 02:45 **Coffee break**

02:45 - 03:45 **Student Session on Inference and Learning II**

Mark Cutler. Reinforcement Learning with Multi-Fidelity Simulators

Ammar Ammar. Learning the Mixed Multinomial Logit Model

Trevor Campbell. Batch-Sequential Clustering via Low-Variance Bayesian Nonparametrics

Hajir Roozbehani. Classification of Conditional Independence Models: a Topological Perspective

03:45 - 04:00 **Coffee break**

04:00 - 05:00 **Plenary Talk**

Prof. Gerald Jay Sussman. We Really Don't Know How to Compute!

05:15 **LIDS Social**

Friday, Jan 31.

08:30 - 09:30 **Breakfast**

09:30 - 10:30 **Plenary Talk**

Prof. Radhika Nagpal. Taming the Swarm

10:30 - 10:45 **Coffee break**

10:45 - 12:15 **Student Session Biological Systems + Optimization**

Noele Norris. Exploring the Optimality of Various Bacterial Motility Strategies

Narmada Herath. Trade-offs Between Retroactivity and Noise in Connected Transcriptional Components

Andras Gyorgy. From Single Modules to Multi-Module Systems in Gene Transcription Networks

Hamza Fawzi. Convex Lower Bounds For Atomic Cone Ranks with Applications to Nonnegative rank and cp-rank

Andrew Mastin. Randomized Minmax Regret for Combinatorial Optimization Under Uncertainty

Frank Permenter. Partial Facial Reduction: Simplified, Equivalent SDPs via Approximations of the PSD Cone

12:15 - 01:30 **Lunch break**

01:30 - 02:30 **Panel Discussion**

02:30 - 02:45 **Coffee break**

02:45 - 03:45 **Student Session on Systemic Efficiency and Resilience**

Alex Teytelboym. Stability and Competitive Equilibrium in Networks with Bilateral Contracts

Diego Feijer. Credit Markets, Systemic Risk and Financial Stability

Kimon Drakopoulos. Network Resilience Against Epidemic Spread

Kuang Xu. Queueing with Future Information

03:45 - 04:00 **Coffee break**

04:00 - 05:00 **Plenary Talk**

Prof. Ali Jadbabaie. Learning and Coordination in Social Networks

05:45 **Banquet (by invitation only)**

Plenary Talk I.

Thursday, Jan 30

ANANT SAHAI,

Associate Professor of Electrical Engineering and Computer Science,
UC Berkeley

Glimpses into an Information Theory for Control

Abstract. Control is an intellectual sibling to communication. Both are about removing uncertainty with limited resources — communication by sharing something about the world and control by shaping the world itself. While information theory has for decades been providing insights into problems of communication, traditional approaches to control did not use information-theoretic techniques or ideas. Recently, we have found some surprising connections between wireless information theory and some central (and long-open) problems in decentralized control.

Just as most problems in multiterminal information theory are open, so are most problems in decentralized control. These problems resist all standard control-theoretic attacks based on linear systems theory and optimization because they seem to be hopelessly nonconvex in all but corner cases and nonlinear control strategies are needed to achieve optimal performance. It turns out that the machinery of linear deterministic models that has been so helpful in understanding problems of relaying and interference in communication can be brought to shed light on the fundamental limits of performance in decentralized control.

Approximately-optimal strategies can be found and the control-theoretic counterparts to ideas like generalized degrees-of-freedom and cut-set bounds can be discovered. There are control/estimation counterparts to ideas like non-coherent communication channels. And even some classical results in decentralized LTI control theory can be reinterpreted as really being about network coding in the complex field. All this suggests that there is an entire parallel realm of information theory that connects to control problems — just waiting to be explored. This talk will give some glimpses into this.

Biography. Anant Sahai received his B.S. in 1994, from the University of California, Berkeley, and his S.M. and Ph.D. from MIT in 1996 and 2001, respectively. He is an associate professor in the EECS Department at Berkeley, where he joined as an assistant professor in 2002. Prior to that, he spent a year as a research scientist at the wireless startup Enuvis (a company founded by John Tsitsiklis and Ben Van Roy, another LIDS alum) in South San Francisco, developing software-radio signal-processing algorithms to enable very sensitive GPS receivers for indoor operation. From 2007 to 2009, he was the treasurer for the IEEE Information Theory Society. His current research interests are at the intersection of information theory and decentralized control, as well as in wireless communication, particularly dynamic spectrum sharing and its regulatory dimensions. He enjoys working very closely with his small group of graduate students on fun and deep problems. He usually teaches small intimate courses but this semester, is teaching a giant intro course with hundreds of students.

Plenary Talk II.

Thursday, Jan 30

Gerald Jay Sussman,

Panasonic Professor of Electrical Engineering,
Massachusetts Institute of Technology

We Really Don't Know How to Compute!

Abstract. We have been building and programming computing machines for about 60 years (over 50 years for me!), and we have seen lots of progress. We can now compute absurdly fast on enormous data sets. We have learned a great deal about composition and abstraction, so we can conquer great complexity. But we have just begun to scratch the surface.

A quick look at the natural world shows that we have a lot to learn. A mammalian neuron takes about ten milliseconds to respond to a stimulus. A human driver can respond to a visual stimulus in a few hundred milliseconds, and decide an action, such as making a sharp turn or activating the brake. So the computational depth of this behavior is only a few tens of steps. Imagine the power if we could make such a machine out of electronics, which is vastly faster than neurons! But, we don't know how to make such a machine, and we don't know how to program one if we had it.

As a further embarrassment, note that the human genome is about 1GB of information. So the information required to build a human from a single, undifferentiated eukariotic cell is about 1GB – a bit smaller than the information required to build a modern operating system. We see that the instructions to build a mammal are written in very dense code. But more importantly, the program is very flexible. Only small patches to the human genome are required to build a cow or a dog rather than a human. Bigger patches result in a frog or a snake. At this point we don't have the foggiest notions how to make a description of such a complex machine as a mammal that is both dense and flexible, with lots of code reuse from earlier versions.

Our idea of system security is similarly archaic. We are continually attacked by mutating parasites, but we can survive for about four score and ten years. How can this possibly work, with only 1GB in the code base?

Our languages, and the principles that they are organized around, are woefully inadequate for capturing processes like these. We will need new design principles and new linguistic support for expressing our designs. I will address this issue and show some ideas that can perhaps get us to the next phase of engineering design.

Biography. Gerald Jay Sussman is the Panasonic (formerly Matsushita) Professor of Electrical Engineering at MIT. He received the S.B. and the Ph.D. degrees in mathematics from the MIT in 1968 and 1973, respectively. He has been involved in artificial intelligence research at M.I.T. since 1964. His research has centered on understanding the problem-solving strategies used by scientists and engineers, with the goals of automating parts of the process and formalizing it to provide more effective methods of science and engineering education. Sussman has also worked in computer languages, in computer architecture and in VLSI design.

Among his contributions, Sussman is a coauthor (with Hal Abelson and Julie Sussman) of "Structure and Interpretation of Computer Programs", the introductory computer science textbook used at M.I.T. As a result of this and other contributions to computer-science education, Sussman received the ACM's Karl Karlstrom Outstanding Educator Award in 1990, and the Amar G. Bose award for teaching in 1991. Over the past decade Sussman and Wisdom have developed a subject that uses computational techniques to communicate a deeper understanding of advanced Classical Mechanics. Sussman and Jack Wisdom, with Meinhard Mayer, have produced a textbook, "Structure and Interpretation of Classical Mechanics," to capture these ideas. Sussman and Wisdom have also presented a functional programmer's view of differential geometry in their monograph "Functional Differential Geometry."

Sussman is a fellow of the Institute of Electrical and Electronics Engineers (IEEE). He is a member of the National Academy of Engineering (NAE), a fellow of the American Association for the Advancement of Science (AAAS), a fellow of the Association for the Advancement of Artificial Intelligence (AAAI), a fellow of the Association for Computing Machinery (ACM), a fellow of the American Academy of Arts and Sciences, and a fellow of the New York Academy of Sciences (NYAS). Sussman is a founding director of the Free Software Foundation. He has been a bonded locksmith. He is a life member of the American Watchmakers-Clockmakers Institute (AWI), a member of the Massachusetts Watchmakers-Clockmakers Association, a member of the Amateur Telescope Makers of Boston (ATMOB), and a member of the American Radio Relay League (ARRL).

Plenary Talk III.

Friday, Jan 31

Radhika Nagpal,

Fred Kavli Professor of Computer Science,
Harvard University

Taming the Swarm

Abstract. Biological systems, from cells to social insects, get tremendous mileage from the cooperation of vast numbers of cheap, unreliable, and limited individuals. What would it take to create our own artificial collectives of the scale and complexity that nature achieves? In this talk, I will discuss one of our recent and ongoing endeavors - the Kilobot project - a 1024 ("kilo") robot swarm testbed for studying collective intelligence. I will describe some of the challenges for building and programming robot swarms at this scale, and I will discuss how we have used the Kilobot swarm to study collective algorithms inspired by both engineering (e.g. coordinate system formation) and nature (collective transport, flocking, self-assembly). A central theme for our work is understanding the global-to-local relationship: how complex and robust collective behavior can be systematically achieved from large numbers of simple agents.

Biography. Radhika Nagpal is the Kavli Professor of Computer Science at Harvard University and a core faculty member of the Wyss Institute for Biologically Inspired Engineering. She received her PhD degree in Computer Science from MIT, and spent a year as a research fellow at the Harvard Medical School. At Harvard she leads the self-organizing systems research group and her research interests span computer science, robotics, and biology.

Plenary Talk IV.

Friday, Jan 31

ALI JADBABAIE,

Alfred Fitler Moore Professor of Network Science,
University of Pennsylvania

Learning and Coordination in Social Networks

Abstract. We present a model of social learning and information aggregation where agents in a social network would like to learn a true state of the world using a stream of private information and exchanges of opinions (represented by beliefs) with their neighbors. Motivated by the practical difficulties of Bayesian updating of beliefs in a network setting, we study a variant of the seminal model of DeGroot, according to which agents linearly combine their personal experiences with the views of their neighbors. We examine how the structure of a social network and the quality of information available to different agents determine the speed of social learning and show that a bound on the rate of learning has a simple analytical characterization in terms of the relative entropy of agents' signal structures and their eigenvector centralities, a measure of importance of nodes in the network. The proposed characterization establishes that the way information is dispersed throughout the social network has non-trivial implications for the rate of learning. Next, we will discuss and study a repeated game in which a group of players attempt to coordinate on a desired, but only partially known, outcome. Agents have an incentive to correctly estimate the state, while trying to coordinate with and learn from others, similar to a Keynesian Beauty Contest game. We will show that myopic but Bayesian agents who repeatedly play this game and observe the actions of their neighbors over a network (that satisfies some weak connectivity condition) eventually succeed in coordinating on a single action. We also show that if the agents' private observations are not functions of the history of the game, the private observations are optimally aggregated in the limit. (Joint work with Pooya Molavi and Alireza Tahbaz-Salehi)

Biography. Ali Jadbabaie received his Bachelors degree with High Honors in Electrical

Engineering (with a focus on Control Systems) from Sharif University of Technology, Tehran, Iran, in 1995. After working as a control engineer for a year, he moved to the US, where he received a Masters degree in Electrical and Computer Engineering from the University of New Mexico, Albuquerque in 1997 and a Ph.D. in Control and Dynamical Systems from California Institute of Technology (Caltech) in 2001. From July 2001-July 2002 he was a postdoctoral scholar at the department of Electrical Engineering at Yale University. Since July 2002 he has been at the University of Pennsylvania, Philadelphia, PA, where he is the Alfred Fitler Moore Professor of Network Science in the department of Electrical and Systems Engineering, with secondary appointments in departments of Computer and Information Sciences and Operations and Information Management (in the Wharton School). He is the director of the Raj and Neera Singh Program in Networked and Social Systems (NETS) at Penn Engineering, a new interdisciplinary undergraduate degree program that blends Network Science, Operations Research, Economics, and Computer Science with Information and Decision Systems. He is an IEEE Fellow, and the recipient of an NSF Career award, ONR Young Investigator Award, the O. Hugo Schuck Best Paper award of the American Automatic Control Council, and the George S. Axelby Outstanding Paper Award of the IEEE Control Systems Society. His current research interests are on multi-agent coordination and control, network science and optimization, with applications to social and economic networks and collective robotics.

Panel Discussion

We'll be holding a panel discussion on Friday Jan 31, 1:30 to 2:30 pm in Room 32-124.

Research, LEGOs and some glue, just in case.

Well-designed engineering systems are typically built out of simpler subsystems designed with extensibility and modularity in mind. To what extent could (and should) we apply these principles to our research? After all, even theory-building for the purpose of analysis can be viewed as a problem of design.

What are the design principles for producing research that interfaces well with other people's work, and is more likely to be reused and built upon? At a more personal level, how might we organize our own thoughts and research processes in a way that also builds insight in a flexible, modular way? How might experienced researchers teach these things to research students? What approaches seem to work? What tends to fail spectacularly?

On the other hand what are the limitations of working (or aspiring to work) in this way? Does it disincentivise tearing up the foundations (either of an entire field or of one's own thinking) and starting again, or does it make it clearer when such action is necessary?

Panelists. The panel will be moderated by Prof. Pablo Parrilo, and composed of:

Prof. Munther Dahleh

Prof. Ali Jadbabaie

Prof. Radhika Nagpal

Prof. Anant Sahai

Prof. Gerald Jay Sussman

Session on Communication Systems.

Matthew Johnston

Opportunistic Scheduling with Limited Channel State Information: An Application of Causal Rate Distortion

Abstract. In this talk, we consider an opportunistic communication system in which a transmitter selects one of multiple channels over which to schedule a transmission, based on partial knowledge of the network state. We characterize a fundamental limit on the rate that channel state information must be conveyed to the transmitter in order to meet a constraint on expected throughput. This problem is modeled as a causal rate distortion optimization of a Markov source. We introduce a novel distortion metric capturing the impact of imperfect channel state information on throughput, and compute a closed-form expression for the causal information rate distortion function for the case of two channels. Additionally, we analyze the gap between the causal information rate distortion and the causal entropic rate distortion functions.

Austin Collins

What To Transmit From Your Many Antennas

Abstract. We analyze a model for wireless communication in the "delay constrained" setting, for example streaming video or a phone call. In this scenario, the transmitter has multiple antennas (MISO - multiple input single output), and the channel is modeled as an isotropic block fading channel with channel state information available at the receiver. We analyze the information theoretically best transmission scheme in this delay constrained setting, and discuss its relation to orthogonal designs used in space-time block coding.

Tianheng Wang

On the Stability of Scheduling Algorithms for Network Navigation

Abstract. Wireless navigation networks enable location-awareness in GPS-challenged environments. For such networks, scheduling algorithms are needed to improve the navigation accuracy through measurement pair selections under limited communication resource. In this paper, we develop an analytical framework to determine the location error evolution for different scheduling algorithms and network settings. Under this framework, we provide sufficient conditions for the stability of the location error evolution, and we quantify the time-averaged network location errors (NLEs) for scheduling algorithms with and without exploiting the network states. Furthermore, we show the optimality of the proposed scheduling algorithms in terms of the error scaling with respect to the agent density. These results provide fundamental insights into the effects of scheduling algorithms and network settings

on the location error evolution, leading to efficient scheduling algorithms for navigation networks.

Abhishek Sinha

Throughput Optimal Algorithms for Multicast in a Wireless Network

Abstract. In this talk we concentrate on the classical problem of throughput-optimal broadcast in both wired and wireless networks setting. A known throughput optimal solution from the literature is to pre-identify a set of disjoint spanning trees and route packets on them. However this approach is not scalable and not particularly convenient in the scenario of constantly changing link qualities and/or topologies in a dynamic wireless mobile environment. In this talk we give a provably optimal Max-Weight type distributed broadcast algorithm for a network whose underlying topology is a Directed Acyclic Graph (DAG). This algorithm does not maintain any global topology information like spanning trees yet achieves throughput optimality. In the subsequent part of the talk we present another broadcast algorithm which works for any network with arbitrary underlying topology. This algorithm has some interesting connections with classical link scheduling algorithms under adversarial arrival models.

Session on Inference and Learning I.

Matthew Johnson

Analyzing Hogwild Parallel Gaussian Gibbs Sampling

Abstract. Sampling inference methods often don't scale well because even weak global dependencies reduce opportunities for parallel computation. Gibbs sampling theory, for example, requires sample updates to be performed sequentially in the absence of strict conditional independence structure among variables. Yet empirical work has shown that some widely-used models with global dependencies can be sampled effectively by going "Hogwild" and simply running Gibbs updates in parallel with only periodic global communication. Little is understood about the successes and limitations of this strategy.

As a step towards such an understanding, in this talk we consider such Hogwild Gibbs sampling approaches for Gaussian distributions. We discuss a framework that provides sufficient conditions for convergence and error bounds along with simple proofs and connections to methods in numerical linear algebra.

Based on joint work with James Saunderson and Alan Willsky.

Ying Liu

Learning Gaussian Graphical Models with Small Feedback Vertex Sets: Observed or Latent

Abstract. Gaussian Graphical Models (GGMs) or Gauss Markov random fields are widely used in many applications, and the trade-off between the modeling capacity and the efficiency of learning and inference has been an important research problem. In this paper, we study the family of GGMs with small feedback vertex sets (FVSs), where an FVS is a set of nodes whose removal breaks all the cycles. Exact inference such as computing the marginal distributions and the partition function has complexity $O(k^2 n)$ using message-passing algorithms, where k is the size of the FVS, and n is the total number of nodes. We propose efficient structure learning algorithms for two cases: 1) All nodes are observed, which is useful in modeling social or flight networks where the FVS nodes often correspond to a small number of high-degree nodes, or hubs, while the rest of the network is modeled by a tree. Regardless of the maximum degree, without knowing the full graph structure, we can exactly compute the maximum likelihood estimate in $O(kn^2 + n^2 \log n)$ if the FVS is known or in polynomial time if the FVS is unknown but has bounded size. 2) The FVS nodes are latent variables, where structure learning is equivalent to decomposing a inverse covariance matrix (exactly or approximately) into the sum of a tree-structured matrix and a low-rank matrix. By incorporating efficient inference into the learning steps, we can obtain a learning algorithm using alternating low-rank correction with complexity $O(kn^2 + n^2 \log n)$ per iteration. We also perform experiments using both synthetic data as well as real data of flight delays to demonstrate the modeling capacity with FVSs of various sizes.

Dan Levine

Focused Active Inference

Abstract. In many inferential settings, data are expensive to acquire or process and are heterogeneously informative about underlying processes of interest. We therefore consider observation selection (e.g., adaptive sampling, active inference, experimental design, etc.) as a way of simultaneously reducing sensing costs while improving inferential accuracy. Prior work has considered the problem of inferring all latent states by selecting locally maximally-informative observations. We consider the more general problem where only a subset of latent variables is of inferential interest, and interactions between observable and latent random variables are described by a high-dimensional, yet factorable, joint distribution containing nuisance variables.

We specifically consider the class of multivariate Gaussian distributions Markov to tree-shaped undirected graphs. For pairs of vertices connected by a unique path in the graph, we show that there exist decompositions of nonlocal mutual information into local information measures that can be computed efficiently from the output of message passing algorithms. We integrate these decompositions into a computationally efficient greedy selector where the computational expense of quantification can be distributed across nodes in the network. Experimental results demonstrate the comparative efficiency of our algorithms for sensor selection in high-dimensional distributions.

We additionally derive an online-computable performance bound based on augmentations of the relevant latent variable set that, when such a valid augmentation exists, is applicable for *any* distribution with nuisances.

Session on Inference and Learning II.

Mark Cutler

Reinforcement Learning with Multi-Fidelity Simulators

Abstract. We present a framework for reinforcement learning (RL) in a scenario where multiple simulators are available with decreasing amounts of fidelity to a real-world learning scenario. Our framework is designed to limit the number of samples used in each successively higher-fidelity/cost simulator by allowing the agent to choose to run trajectories at the lowest level that will still provide it with informative information. The approach transfers state-action Q-values from lower-fidelity models as heuristics for the "Knows What It Knows" Rmax family of RL algorithms, which is applicable over a wide range of possible dynamics and reward representations. Theoretical proofs of the framework's sample complexity are given and empirical results are demonstrated on a remote controlled car with multiple simulators. The approach allows RL algorithms to find near-optimal policies for the real world with fewer expensive real-world samples than previous transfer approaches or learning without simulators.

Ammar Ammar

Learning the Mixed Multinomial Logit Model

Abstract. Computing a ranking over choices using consumer data gathered from a heterogenous population has become an indispensable module for any modern consumer information system, e.g. Yelp, Netflix, Amazon and app-stores like Google play. The corresponding ranking or recommendation algorithm needs to extract meaningful information from noisy data *accurately* and in a *scalable* manner. A principled approach to resolve this challenge requires a *model* that connects observations to recommendation decisions and subsequently a tractable inference algorithm utilizing this model.

In this paper, we abstract all the preference data generated by consumers as noisy, partial realizations of their innate preferences, i.e. orderings or permutations over choices. Inspired by the seminal works of Samuelson (cf. *axiom of revealed preferences*) and that of McFadden (cf. discrete choice models for transportation), we model population's innate preferences as a Mixture of the so called Multinomial Logit model (MMNL). Now given a consumer's limited preferences, the recommendation problem boils down to (a) learning the MMNL model corresponding to the population data, (b) finding the mixture component of MMNL that closely represents the revealed preferences of the consumer at hand, and (c) recommending other choices to her/him that are ranked high according to thus found component.

Inspired by this, we address the problem of learning MMNL model from partial preferences. We identify fundamental limitations of *any* algorithm to learn such a model as well as provide conditions under which, a simple, data-driven (non-parametric) algorithm learns the model effectively. The proposed algorithm has a pleasant similarity to the standard

collaborative filtering for scalar (or star) ratings, but in the domain of permutations. This work advances the state-of-art in the domain of learning distribution over permutations.

Trevor Campbell

Batch-Sequential Clustering via Low-Variance Bayesian Nonparametrics

Abstract. Bayesian nonparametrics (BNPs) are a class of probabilistic models in machine learning in which the number of parameters is mutable and is learned alongside the values of those parameters during inference. This yields a richness and flexibility in describing uncertainty that cannot be achieved with traditional Bayesian models. However, the enhanced flexibility comes at a cost; inference procedures for BNPs (e.g. Gibbs sampling, variational Bayes, particle learning) are typically computationally expensive, limiting the applicability of BNPs to autonomous systems in which timely decision-making is crucial.

This talk will discuss BNP low-variance asymptotics, a recently developed methodology for finding hard clustering algorithms that share the flexibility of BNPs, but have computationally inexpensive learning algorithms. In particular, a low-variance analysis of the dependent Dirichlet process (DDP) Gaussian mixture model will be presented, along with two hard clustering algorithms that can be extracted from the analysis. Finally, results from experiments will be covered, demonstrating that these hard clustering algorithms require orders of magnitude less computational time than the DDP while still providing comparable or higher accuracy on both synthetic and real datasets.

Hajir Roozbehani

Classification of conditional independence models: a topological perspective

Abstract. Conditional independence (CI) relations have a natural structure as a real (or complex) manifold. As such, they look like a Euclidean space and have no interesting topology. In this talk we construct CI models as projective varieties and show that interesting topology emerges in this way.

To do this we work with CI models of regular Gaussian random variables. It is known that such models are the intersection of the zero set of polynomial equations with the positive definite cone. We show that relaxing the positive definite constraint adds only components with singular principal minors to the feasible set. This gives a convenient method to construct varieties from CI models and resolves a conjecture in the algebraic statistics literature. We then show how this construction gives rise to some simple invariants that can distinguish CI models. We conclude the talk by showing that topologists can tell independence from conditional independence provided that they study projective geometry.

Session on Biological Systems + Optimization.

NOELE NORRIS

Exploring the Optimality of Various Bacterial Motility Strategies

Abstract. While many species of bacteria are motile, they use various random strategies to determine where to swim. This chemotaxis allows bacterial populations to distribute themselves in accordance to distributions of nutrients found within an environment. We extend past work describing a chemotactic E. coli cell as an ergodic, stochastic hybrid system and use experimental data on bacterial motion in microfluidic environments to model other species of bacteria. Our focus is on understanding the differences between the run-and-tumble strategy of E. coli and the more complicated run-reverse-flick strategy of the marine bacterium Vibrio alginolyticus. We use stochastic stability theory to analyze the chemotaxis models in terms of their stationary distributions and also derive a diffusion approximation of the system that provides further insight into the performance of various strategies. By comparing general chemotactic strategies, we hypothesize why various strategies may be evolutionarily advantageous for particular environments. These results also provide intuition for designing minimalistic multi-agent robotic systems that can be used for various environmental monitoring and source-seeking tasks.

This presentation summarizes joint work with Filippo Menolascina, Roman Stocker, and Emilio Frazzoli.

NARMADA HERATH

Trade-offs between retroactivity and noise in connected transcriptional components

Abstract. At the interconnection of two gene transcriptional components in a biomolecular network, the noise in the downstream component can be reduced by increasing its gene copy number. However, this method of reducing noise increases the load applied to the upstream system, called ‘retroactivity, thereby causing a perturbation in the upstream system. In this work, we quantify the error in the system trajectories caused by perturbations due to retroactivity and noise, and analyze the trade-off between these two perturbations. We model the system as a set of nonlinear chemical Langevin equations and employ contraction theory for stochastic systems to find an upper bound for the error due to noise. We show that this upper bound can be minimized by increasing the gene copy number, but this leads to an increase in the magnitude of the error due to retroactivity.

ANDRAS GYORGY

From Single Modules to Multi-Module Systems in Gene Transcription Networks

Abstract. Predicting the dynamic behavior of a complex network from that of the composing modules is a central problem in systems and synthetic biology. Unfortunately, modules display context-dependent behavior. As a result, our current ability of predicting the emergent behavior of a network from that of the composing modules remains limited. One cause of context-dependence is retroactivity. This phenomenon is similar to loading in electrical networks and influences the dynamic performance of a module upon connection to other modules. Here, we establish an analysis framework for gene transcription networks that explicitly accounts for retroactivity to reliably predict how a module's behavior will change once it is connected to other systems. This framework carries substantial conceptual analogy with the electrical circuit theory based on equivalent impedances established by Thevenin. Specifically, relying on model reduction techniques for nonlinear systems, we demonstrate that a module's key interconnection properties are encoded by three retroactivity matrices: internal, scaling, and mixing retroactivity. All of them have a physical interpretation and can be computed from measurable biochemical parameters and from the modules' topology, similar to how one would compute the equivalent impedance of a network of interconnected electrical components. The internal retroactivity quantifies the effect of intramodular connections on an isolated module's dynamics. The scaling and mixing retroactivity establish how intermodular connections change the dynamics of connected modules. Based on these matrices and on the dynamics of modules in isolation, we can accurately predict the dynamic behavior of an arbitrary interconnection of modules. Further, using contraction theory we provide a quantitative metric that determines how robust the dynamic behavior of a module is to interconnection with other modules. Our metric can be employed both to evaluate the extent of modularity in natural networks and to establish concrete design guidelines to minimize retroactivity between modules in synthetic systems. Finally, we illustrate how our framework predicts and explains surprising and counter-intuitive dynamic properties of naturally occurring network structures, which could not be captured by existing models of the same dimension.

Hamza Fawzi

Convex lower bounds for atomic cone ranks with applications to nonnegative rank and cp-rank

Abstract. We study a certain class of atomic rank functions defined on a convex cone, which generalize several notions of "positive" ranks such as nonnegative rank or cp-rank (for completely positive matrices). We propose a new method to obtain lower bounds for such ranks using semidefinite programming and sum-of-squares techniques. For the case of nonnegative rank and cp-rank our lower bound has interesting connections with existing combinatorial bounds and satisfy important structural properties such as invariance under diagonal scaling or subadditivity.

Andrew Mastin

Randomized Minmax Regret for Combinatorial Optimization Under Uncertainty

Abstract. The minmax regret problem for combinatorial optimization under uncertainty can be viewed as a zero-sum game played between an optimizing player and an adversary, where the optimizing player selects a solution and the adversary selects costs with the intention of maximizing the regret of the player. Existing minmax regret models consider only deterministic solutions/strategies, and minmax regret versions of most polynomial solvable problems are NP-hard. In this paper, we consider a randomized model where the optimizing player selects a probability distribution (corresponding to a mixed strategy) over solutions and the adversary selects costs with knowledge of the player's distribution, but not its realization. We show that under this randomized model, the minmax regret version of any polynomial solvable combinatorial problem becomes polynomial solvable. This holds true for both the interval and discrete scenario representations of uncertainty. We also show that the maximum expected regret value under the randomized model is upper bounded by the regret under the deterministic model.

Frank Permenter
Partial facial reduction: simplified, equivalent SDPs via approximations of the PSD cone

Abstract. We develop practical semidefinite programming (SDP) facial reduction procedures that utilize computationally efficient approximations of the positive semidefinite cone. The proposed methods simplify SDPs with no positive definite solution by solving a sequence of easier optimization problems. We demonstrate effectiveness of our techniques on SDPs arising in practice and describe our publicly-available software implementation.

Session on Systemic Efficiency and Resilience.

Alex Teytelboym
Stability and competitive equilibrium in networks with bilateral contracts

Abstract. We consider general networks of bilateral contracts that include supply chains. We define a new stability concept called walk stability and show that any network of bilateral contracts has a walk-stable contract allocation whenever agents' preferences satisfy full substitutability. Walk-stable contract allocations may not be immune to group deviations, efficient, or in the core. We also show that competitive equilibrium exists in networked markets even in the absence of transferrable utility. The competitive equilibrium contract allocation is walk-stable.

Diego Feijer
Credit markets, systemic risk and financial stability

Abstract. We develop a model that speaks of banking crises, financial regulation and the unstable nature of credit. Feedback between financial constraints and credit market prices result in an externality that induces banks to borrow in amounts that are not Pareto optimal leaving the system fragile and susceptible to systemic runs. Regulatory policies that affect the cost of debt can restore constrained efficiency and reduce the probability and severity of financial crises.

Kimon Drakopoulos
Network Resilience Against Epidemic Spread

Abstract. We study the problem of minimizing the expected extinction time of an epidemic that evolves on a graph according to an SIS model. We consider curing policies for the nodes which exploit the knowledge of the current state of the epidemic. We characterize the family of graphs for which it is possible to achieve subpolynomial extinction time and provide a policy that achieves it. Moreover we prove that for all graphs outside this family the expected extinction time is polynomial.

Kuang Xu
Queueing with Future Information

Abstract. Can forecasts significantly improve the performance of queueing systems? We study an admissions control problem: to redirect away incoming jobs up to a fixed rate, in

order to minimize the resulting average queue length, Q. We show that Q diverges to infinity in the heavy-traffic regime under any online policy, but can be bounded given (even limited) future forecasts. Based on joint work with Joel Spencer and Madhu Sudan.