Laboratory for Information and Decision Systems

The Laboratory for Information and Decision Systems (LIDS) is an interdepartmental laboratory devoted to research and education in systems, networks, control, optimization, communications, and statistical inference. These disciplines, which span the domain of the analytical information and decision sciences, have broad applications and play a critical role in addressing important systems-focused research questions in science, engineering, and society. LIDS provides an intellectually cohesive and collaborative environment that fosters forward-looking research and instills in our students the disciplinary depth and interdisciplinary understanding that will position them to be the research and engineering leaders of today and tomorrow.

LIDS continues to be at the forefront of foundational research in the traditional core disciplines of the field. Over the past few years, the scope of LIDS research has broadened to include new and emerging areas such as the following:

- Modeling, analysis, and optimization of multi-agent networked systems, including infrastructure systems (e.g., communications networks, traffic networks, and the power grid) and social, economic, and financial systems
- Efficient large-scale inference (using optimization, statistics, and algorithms)
- The theory and practice of autonomous systems

The faculty members within LIDS are principally drawn from the Department of Electrical Engineering and Computer Science (EECS) and the Department of Aeronautics and Astronautics. Due to the increasingly interdisciplinary nature of the lab's work, LIDS has also built strong relationships with many other entities at MIT, including the Computer Science and Artificial Intelligence Laboratory, the Research Laboratory of Electronics, the Operations Research Center (ORC), the Media Laboratory, the Harvard-MIT Division of Health Sciences and Technology, the Department of Civil and Environmental Engineering, the Department of Mechanical Engineering, the Department of Economics, and the Sloan School of Management. LIDS faculty members have leadership roles in major new initiatives at the Institute with applications of critical societal importance. These initiatives include projects focusing on transportation systems (e.g., the Future Urban Mobility initiative of the Singapore-MIT Alliance for Research and Technology [SMART]), energy systems (e.g., programs in collaboration with the MIT Energy Initiative), and social, economic, and financial networks in close collaboration with Economics and Sloan.

The lab is part of the Institute for Data, Systems, and Society (IDSS), whose mission is to address complex societal challenges by advancing education and research at the intersection of statistics, data science, information and decision systems, and social sciences. LIDS became part of IDSS in 2015 and continues to be a key partner in its growth and development. LIDS faculty play a pivotal role in defining the IDSS intellectual agenda, leading efforts in statistics and flagship projects (e.g., in finance, autonomy, and smart cities), and designing new academic programs in statistics and social and engineering systems. Also, they are heavily involved in the search for new faculty in networks and statistics.

LIDS has continued its growth across many dimensions this year and its community is excited to welcome Professor Philippe Rigollet, a recently tenured Mathematics Department faculty member who works at the intersection of statistics, machine learning, and optimization, focusing primarily on the design and analysis of statistical methods for high-dimensional problems. Professor Rigollet's presence is a significant addition to the critical mass of LIDS researchers in the area of inference and data science. In addition, Audun Botterud has joined LIDS as a principal research scientist. Botterud, who is also a member of the research staff at the Argonne National Laboratory, adds to the lab's critical mass in the area of energy systems.

LIDS researchers continue to have great success in obtaining funding for our broad and deep research agenda, and we continue to develop our relationships with industrial organizations and national laboratories including Draper Laboratory, Lincoln Laboratory, Aurora Flight Sciences, the Coca-Cola Company, Accenture, Microsoft Research, Anheuser-Busch InBev, Rockwell Collins Inc., and NVIDIA. Also, thanks to a rich history of research excellence and leadership, LIDS remains a magnet for the very best, attracting not only outstanding students but a continuous stream of world-leading researchers as visitors and collaborators.

Intellectual Vision

The mission of LIDS is to develop and apply rigorous approaches and tools for system modeling, analysis, design, and optimization. This mission encompasses the development of novel analytical methodologies as well as the adaptation and application of advanced methods to specific contexts and application domains. LIDS research addresses physical and man-made systems, their dynamics, and the associated information processing.

LIDS research spans the range from theory to practice and, as such, comprises several different dimensions:

- A set of core mathematical disciplines, including probability and statistics, dynamical systems, and optimization and decision theory
- A set of core engineering disciplines, including inference and statistical data processing, transmission of information, networks, and systems and control
- A set of broad challenges in traditional and emerging applications of critical societal importance

Research at LIDS involves activities within and across all of these dimensions, which creates strong synergies: work in each of the mathematical disciplines leads to new methodologies that enable advances in core engineering disciplines and in interdisciplinary applied investigations; conversely, work on new interdisciplinary challenges provides the inspiration and direction for fundamental disciplinary research, including the charting of emerging new disciplines.

Research Areas

The lab's multiple research strands are usually cross cutting and cannot be neatly organized into categories. Nevertheless, they can be broadly classified in terms of the following core areas.

Networks: this area includes communications, information theory, and networking, with applications to wireless and optical systems and data centers. Additional recent directions include analyses of social networks and of agent interactions in networked systems, with applications ranging from analysis of data generated by large-scale social networks to the study of dynamics and risk in large, interconnected financial, transportation, and power systems.

Statistical inference and machine learning: this area deals with complex systems, phenomena, and data that are subject to uncertainty and statistical variability. It also includes creation of large-scale data-processing software systems. Research ranges from development of basic theory, methodologies, algorithms, and computational infrastructures to adaptations of this work for challenging applications in a broad array of fields. Typical applications involve causal inference in experimental design, social data processing, and e-commerce, as well as image processing, computer vision, and automation of data engineering. Other current topics include reinforcement learning and online optimization, recommendation systems, graphical models, large-scale software systems for data engineering, medical image processing, causal inference in genetics, and high-dimensional statistics.

Optimization: the aim of this area is to develop analytical and computational methods for solving optimization problems in engineering, data science, and operations research, with applications in communication networks, control theory, power systems, machine learning, and computer-aided manufacturing. In addition to linear, nonlinear, dynamic, convex, and network programming, research focuses on methods that exploit the algebraic structure of large-scale problems and simulation-based methods.

Systems theory, control, and autonomy: this area deals with all aspects of system identification, inference, estimation, control, and learning for feedback systems. Theoretical research includes quantification of fundamental capabilities and limitations of feedback systems, development of practical methods and algorithms for decision making under uncertainty, robot sensing and perception, and inference and control over networks, as well as architecting and coordinating autonomy-enabled infrastructures for transportation, energy, and beyond.

The availability of increasingly capable sensing, communication, and computation enables the collection and transfer of large amounts of data pertaining to complex and heterogeneous interconnected systems. This opens up many new avenues for methodological research in all of the above areas, with some common themes such as data fusion, distributed learning and decision making, and issues of scalability, robustness, and performance limits.

Notable areas of current active research include the following:

- Biological systems and biomedical data analysis
- Coordination of unmanned autonomous systems
- Data center architectures and algorithms

- Energy services and information systems
- General, scalable tools for machine learning
- Human language learning
- Information transmission theory for networks
- Learning in social networks
- Localization and navigation
- Machine learning for recommendation systems and social media
- Network scheduling and routing
- Online learning and optimization
- Optimization and inference using algebraic methods
- Renewable energy integration
- Social network analysis and characterization
- Transportation network analysis, control, and design
- Ultra-wideband and other emerging communications technologies

Furthermore, it is well understood that research within traditional boundaries in information and decision sciences is not adequate to address many of the emerging societal challenges. This has motivated LIDS to become engaged in several research thrusts that cut across disciplinary boundaries and involve considerable interaction and collaboration with colleagues in other MIT units and in other disciplines. Examples include:

- Foundational research in network science, including network dynamics, control, and efficient algorithms
- Foundational research in game theory and mechanism design involving the study of new equilibrium notions and dynamics in games, as well as the design of efficiently computable incentive methods for large-scale networked, dynamic environments
- Development of new frameworks for modeling and understanding systemic risk
- Fundamental issues in cyber-physical systems, including architectural design, security and privacy, cross-layer algorithms, and tools for analysis, verification, and performance guarantees
- Foundational theory for multi-resolution modeling
- Development of scalable and efficient inference algorithms for problems involving “big data,” including basic research on graphical models

Faculty Activities

The research activities described in the sections to follow are organized in terms of individual faculty members. Nevertheless, many of the activities not only cut across the disciplines, applications, and emerging areas mentioned above but are also collaborative with others within LIDS and elsewhere at MIT.

Dimitri Bertsekas

Professor Bertsekas (McAfee Professor of Electrical Engineering) performs research on problems of sequential decision making under uncertainty, which are pervasive in communication networks, manufacturing systems, logistics, and control of nonlinear dynamical systems. In theory, such problems can be addressed with dynamic programming techniques. In practice, only problems with a moderately sized state space can be handled. This research effort deals with the application of neural networks and other approximation and interpolation methodologies to overcome the curse of dimensionality of real-world stochastic control problems. Recent efforts have focused on analysis of properties of value and policy iteration methods and their applications in deterministic optimal control. In addition, extensions of the fundamental concept of regular policies in abstract dynamic programming were developed as a follow-up to a research monograph on the subject published in 2013. A second edition of this monograph, which will incorporate Professor Bertsekas's latest research, is currently being written and is planned for publication in late 2017.

Professor Bertsekas is also interested in deterministic optimization problems and the role of convexity in solving them, possibly through the use of duality. In recent work, he has been able to establish a connection between the proximal algorithm, which is central in convex optimization, and temporal difference methods, which are prominent in approximate dynamic programming. This may lead to fruitful cross fertilization between the two fields.

Robert C. Berwick

Professor Berwick's research during the past year led to new advances regarding the biological and computational nature of human language and cognition, especially how human knowledge of language differs from purely computer-driven approaches. Surprisingly, he found that human language learning is easier with "small data" than with "big data." One can conclude that children do not construct rules the way a Google engine figures out what a cat looks like; thus, while making great strides in artificial intelligence, we remain far from an explanation of human intelligence. Children can construct rules using one or several examples but flounder in the face of thousands. This has been tested not only for language rules but also for learning of numbers.

In other work, a systematic comparison of "deep neural net" models for language processing against linguistically based descriptions revealed that, in surprising ways, deep neural nets fail to match the knowledge of language possessed by five-year-old children. For example, a sentence such as "This doll is easy to see" is understood by children to mean that it is easy for someone to see the doll; however, the analysis is not done correctly by the most advanced Google-type parsers, no matter how many examples they have been trained on. Apparently, these "more modern" parsers simply use the wrong representations, in the sense that they do not correspond to what we know about human language. Professor Berwick has constructed a parsing system that aligns more closely with the way children acquire language and people parse language. Professor Berwick used novel brain imaging analysis (magnetoencephalography) to test this system, recording subjects while they listened to a story being read aloud, and determined that the time course of brain activity more closely matched the linguistically correct model.

Audun Botterud

Audun Botterud started as a principal research scientist in LIDS in September 2016 after more than 10 years on the research staff at the Argonne National Laboratory, where he still holds a co-appointment. The main goal of his research is to improve understanding of the complex interactions among engineering, economics, and policy in electricity markets and ultimately enable the transition to a cost-efficient and reliable low-carbon energy system. Toward this end, he uses analytical methods from operations research and decision sciences combined with fundamental principles of electrical power engineering and energy economics. At a more general level, his research focuses on decision making under uncertainty in complex systems. Botterud collaborates extensively with faculty, students, and researchers at MIT, Argonne, and many other institutions within the United States and abroad.

Grid Integration of Renewable Energy

The rapid expansion of wind and solar energy resources in the electric power grid gives rise to challenges in the operation and planning of the power grid. In the past year, Botterud's group has explored several different optimization methods for how to best address uncertainty and variability in wind power, completing work on interval optimization, chance-constrained goal programming, and fuzzy programming applied to power system operational decision problems. On the computational side, he has explored the use of decomposition methods and high-performance computing to speed up solutions to the classical transmission-constrained unit commitment problem, a notoriously hard problem to solve for large-scale power systems.

Another key challenge resulting from higher shares of renewables is to increase flexibility in the power grid. Toward this end, Botterud's group has carried out research showing significant advantages of having wind power provide ramping capacity to the grid. Moreover, he has collaborated with nuclear engineers to develop novel models of nuclear power plants that consider flexible operation of this generation technology within its physical constraints. Extensive simulations indicate substantial economic benefits to the power system as well as nuclear power plant owners from switching to a flexible operating scheme. Botterud is also looking at electricity market design questions that arise from more renewables on the grid.

A new direction in the past year, enabled by a grant from the Siebel Energy Institute, is exploration of decision making within distributed urban energy systems through bi-level optimization and game theoretical algorithms for optimal sharing of local energy resources within multi-resident buildings. This work is being conducted jointly with Professor Asu Ozdaglar.

Energy Storage Analytics

Energy storage holds the promise of solving multiple problems in the power grid but is still considered an expensive technology relative to competing solutions. In the past year, Botterud has worked with multiple collaborators on the potential use of energy storage in current and future electricity markets. This has involved modeling of several grid applications of energy storage, from provision of short-term balancing and load leveling services to the grid to the impact on long-term investment planning for a de-carbonized energy system. His research group has also developed stochastic programming models for optimal trading of energy storage under price uncertainty in day-ahead and real-time energy markets and demonstrated significant improvements in trading performance relative to traditional deterministic benchmark strategies. Another key focus area has been the development of more realistic battery representations for use in power system optimization models, factoring in that battery losses and power limits are a function of battery state. Moreover, Botterud and his group have looked at long-term battery degradation and lifetime effects. Their results emphasize the importance of considering degradation in optimal operational battery schemes to extend battery lifetimes and discounted economic lifetime benefits.

A new direction within energy storage analytics, enabled by a seed grant from the MIT Energy Initiative, is to extend research on improved modeling of batteries in the power grid by directly incorporating empirical laboratory test results in model development and parameter turning, as well as guiding early-stage laboratory work on promising battery chemistries by providing testing schemes for more realistic grid conditions. This work is being conducted jointly with Professor Fikile Brushett of the Department of Chemical Engineering.

Guy Bresler

Professor Bresler (Bonnie and Marty Tenenbaum Assistant Professor of Electrical Engineering and Computer Science) works in information theory, statistics, and applied probability. Specifically, his research aims to understand the relationship between the combinatorial structure and computational tractability of high-dimensional inference in graphical models and other statistical models.

With PhD student Mina Karzand, Professor Bresler has been working on the problem of recommendation system design. The goal is to understand the optimal algorithms in an idealized setting. To this end, they are assessing a latent variable model of user preferences with a finite number of user types and item types. All users of a given type either like or dislike all items from a given type. At each time step, each user in the system is recommended an item and provides noiseless feedback to the system. The performance metric is the amount of accumulated “reward” up to any given time (i.e., the number of liked recommendations). Their main results are derived from analyses of three algorithms along with corresponding information-theoretic lower bounds showing their optimality in certain system operating regimes. Bresler and Karzand proved that an algorithm based on classical user-user collaborative filtering is optimal in a setting with a huge number of item types and relatively few user types. If there are a huge number of user types and few item types, a variant of classical item-item collaborative filtering is optimal. Finally, in all system operating regimes, a new algorithm that incorporates the best of both user-user and item-item collaborative filtering is optimal. This is a step in a broader research program focused on understanding recommendation system design. Next steps include solving more realistic models for user preferences and conducting experiments on real data.

With PhD student Dheeraj Nagaraj, Professor Bresler has been developing the theory of graphical model sparsification. This work will have implications for learning graphical models from data as well as for speedup of inference algorithms such as Markov chain Monte Carlo (MCMC) and message-passing algorithms. The goal is to understand the structure of the space of probability distributions described by graphical models in terms of the complexity of the models. Bresler and Nagaraj initiated their study by proving both negative and positive results regarding approximation of the Curie-Weiss model, an Ising model on the complete graph with uniform edge interactions. They showed that it is always possible to approximate the Curie-Weiss model using a model with only linearly many edges under a certain local total variation metric that captures the accuracy of predictions made on the basis of partial observations. The results are encouraging, and their next goal will be to sparsify general Ising models.

Munther Dahleh

Professor Dahleh, director of the Institute for Data, Systems, and Society and the William A. Coolidge Professor of Electrical Engineering and Computer Science, is researching robustness scaling in large networks, transportation problems, and virtual storage.

Robustness Scaling in Large Networks

Professor Dahleh and SM student Tuhin Serkar have proposed new asymptotic notions that capture the robustness of large-scale networks as the dimension of the system grows. Motivated by applications in transportation, energy, and economic networks, they are studying network topologies that are robust to scaling and network characteristics that result in poor scaling. In addition, they are analyzing how perturbation of links affects network robustness and relating nodal degrees to network robustness.

Capacity Control for Traffic Management

Professor Dahleh and postdoc Yasin Yazicioglu are focusing on transportation problems. Their efforts are building on their past work in this area in which such problems were modeled as dynamic flows on a network, with each flow being split at the nodes according to a local decision strategy and exiting a link according to a nonlinear outflow function. Their past work showed that, in general, the fragility of such networks for a broad class of decision strategies is captured by a min-cut-type property of the graph. These results led to analyses of situations wherein cascaded failures can occur in the network for which more elaborate fragility conditions are derived. They considered the question of how to dynamically adjust link capacities to minimize congestion and prevent cascaded failures. In practice, capacities can be adjusted by controlling speed limits on a link as well as reallocating part of the capacity through lane shifting. Recognizing that fragility measures are not monotone in link capacities, Dahleh and Yazicioglu obtained the unintuitive result that reduction of capacity can increase network robustness. This is particularly important when fast links with high capacities feed into slower links with smaller capacities (e.g., exits from highways). With this understanding, they proposed a distributed real-time strategy for adjusting capacities that can minimize traffic congestion. This provides a direct mechanism for reducing congestion that can be employed by a central authority as opposed to relying on creation of incentives for drivers to alter their behavior.

Virtual Storage

Professor Dahleh and postdoc Daria Madjidian are investigating the ability of a homogeneous collection of deferrable energy consumers to behave as a battery, that is, to absorb and release energy in a controllable fashion up to fixed and predetermined limits on volume, charge rate, and discharge rate. They have derived bounds on the batteries that can be emulated and shown that there is a fundamental tradeoff between the ability to absorb and release energy at high rates. Also, they have introduced a new class of dynamic priority-driven feedback policies that balance these abilities and characterized the batteries that they can emulate. They believe that this approach can be very effective in characterizing the "virtual" capacity of urban transportation systems. Work on this project continues.

Jonathan How

Professor How leads research efforts focused on the control of multiple autonomous agents, with an emphasis on distributed decision making under uncertainty; path planning, activity, and task assignment; mission planning for unmanned aerial vehicles and unmanned ground vehicles; sensor network design; and robust, adaptive, and nonlinear control. Professor How is also the principal investigator for the Aerospace Control Laboratory. Recent research includes the following.

Decentralized Multi-Task Learning

Many real-world tasks involve multiple agents with partial observability and limited communication. While the previous efforts of Professor How's group improved multi-agent planning scalability, this work assumed knowledge of a high-level domain model (e.g., probabilistic models over sensors and actions). Learning without knowledge of the underlying model is challenging in these settings due to the local viewpoints of agents, in which the world is perceived as non-stationary as a result of concurrently exploring teammates. Approaches that learn specialized policies for individual tasks face major problems when applied to the real world: not only do agents have to learn and store a distinct policy for each task, but in practice the identity of the task is often non-observable, making these algorithms inapplicable. This work formalizes and addresses the problem of multi-task, multi-agent reinforcement learning under partial observability. The group's first contribution is a decentralized single-task learning approach that is robust to the concurrent interactions of teammates. Its second contribution is an approach for distilling single-task policies into a unified policy that performs well across multiple related tasks without explicit provision of task identity. Both the single-task and multi-task phases of the algorithm have been demonstrated to achieve good performance on a set of multi-agent domains. The group's approach makes no assumptions about communication capabilities and is fully decentralized during both learning and execution. According to Professor How and his group, this is the first formalization of decentralized multi-agent, multi-task learning under partial observability.

Truncated Bayesian Nonparametric Models

Bayesian nonparametric models have infinitely many latent parameters, which has the effect of allowing the number of parameters with an influence on observed data to grow as the amount of observed data increases. In order to exploit conditional independence properties in the model for efficient inference, it is necessary to condition on the latent parameters; however, this is not possible when there are infinitely many parameters. Professor How and his group addressed this issue by developing principled finite approximations of priors from Bayesian nonparametric models. These approximations involved truncating a sequential representation of the countable infinity of latent parameters. They developed error bounds in terms of L1 norms on the marginal density of observed data or, alternatively, the probability of the truncated model making a "mistake" when any data point selects a parameter from the truncated tail. This work paves the way for tractable, efficient truncated vibrational and MCMC inference schemes for nonparametric Bayesian models.

Task-Driven Navigation and Mapping with Resource Constraints

A fundamental problem for autonomous systems is the ability to simultaneously map the environment and localize within, especially when there is no global reference. This problem is often referred to as simultaneous localization and mapping (SLAM). Professor How and his group study how to process fast-growing data with more advanced sensing technologies and use sparse models to account for robots' computation, communication, and memory constraints. In the first part of the work, a subset of landmarks and measurements of landmarks are selected from a pre-collected data set to build a sparse map, such that the robot still achieves good navigation performance (minimal collision) with this map. The group then extends the robot's capability to plan its own trajectories while autonomously exploring an unknown environment to build maps. A topological feature graph is developed to maintain the sparsity of the map but still enable collision checks for path planning. The approach uses a unified information metric to explicitly balance exploration of new environments and exploitation of mapped environments. Finally, The group uses deep neural networks to detect real-world objects as landmarks for map building. The algorithm explicitly takes into account false positives in object detection and performs object data associations and SLAM simultaneously. The proposed approaches are compared with existing methods using both detailed simulations and real-world experiments. The results show that the new approaches have good navigation and mapping performance with significantly less memory and fewer computation resources.

Planning with Limited Communication Networks

Communication plays a key role in multi-agent planning and estimation, serving as a critical component of the robot system. Many multi-robot coordination schemes require some sort of consensus across the team, such as in task allocation, consensus filtering, and coverage control. In all of these planning algorithms, robots are able to estimate and plan in a distributed fashion with distributed communication. Communication is especially important in highly dynamic environments, where information can change very quickly and unpredictably, and thus team-wide communication is necessary to maintain a consistent team-wide belief. In these dynamic scenarios, it is imperative to have fast and reliable communication that still maintains the decentralized nature of most of these algorithms. In this work, Professor How and his team are exploring the performance of a real ad hoc network using 802.11 Wi-Fi adapters and low-cost Raspberry Pi computers placed across the sixth floor of the Stata Center. They are interested in understanding the reliability and delays associated with using the 802.11 ad hoc mode for broadcasting information across a team of agents. One of the goals is to better understand the effects of hidden neighbors, weak link signals, and limited buffer sizes on the team's ability to transmit information. Also, by co-designing planning algorithms and communication realities, the team can both develop planning algorithms that are robust to communication failures and improve team-wide communication.

Online Threat Assessment for Highway Driving Scenarios

One of the main aims of research on autonomous driving is the quest to increase road safety and reduce accidents. Highway driving not only constitutes the majority of driving time in developed nations but also is most likely to lead to accidents because of high speeds. The objective of this project is to develop an online threat assessment system for highway driving scenarios, with a special focus on predicting lane change maneuvers of surrounding traffic based solely on information available through on-board perception sensors. Past work highlighted the importance of obtaining accurate lateral velocity estimates for lane change predictions through a hidden Markov model–based approach. Also, vehicle interactions, explicitly highlighted as crucial for robust threat assessments, are not addressed in state-of-the-art lane change maneuver prediction approaches. The aim is to better understand and model the impact of tracking noise and interactions not only between the "ego" vehicle and relevant traffic but also among the vehicles constituting the relevant traffic.

Marija Ilic

Marija Ilic, formerly a professor of electrical and computer engineering at Carnegie Mellon University, became a permanent senior staff member of MIT's Lincoln Laboratory in February 2016. She joined the Institute for Data, Systems, and Society as a visiting professor and LIDS as a principal investigator in June 2016. Professor Ilic has proactively worked toward building a collaboration among Lincoln Laboratory Group 73, LIDS, IDSS, and the MIT Energy Initiative. She has done so by organizing and participating in several meetings with current and potential industry sponsors and by presenting her work and vision at workshops on electric energy systems. She worked closely with Professor Munther Dahleh on organizing a major IDSS brainstorming workshop, "Solving Complex Energy Systems: Challenges and Opportunities," in March 2017. Professor Ilic is leading projects in major research areas such as modeling and control of changing electric energy systems, resiliency, and electricity markets. She advises three EECS graduate students, and a fourth student completed his master's degree in June 2016 and will continue with her as a doctoral student in the fall. She is the principal investigator on three LIDS projects and several Lincoln Laboratory projects.

Looking forward, Professor Ilic is very eager to build a scalable electric power system simulator with the main intent of using it for demonstrating the effects of new technologies, cyber solutions, and electricity markets on physical electric energy systems. The idea for the simulator is an outgrowth of a collaboration with the National Institute of Standards and Technology that she led at Carnegie Mellon before coming to MIT. At present, an existing version of the simulator is being used by her doctoral students and her Lincoln Laboratory research colleagues. Professor Ilic will begin looking for support and help to both obtain space and raise funds for transforming the simulator into an open access national and international facility. She already has an active collaboration with the Japan Science and Technology Agency, and a similar collaboration has been initiated with the Interdisciplinary Center for Scientific Computing at Heidelberg University in Germany.

Ali Jadbabaie

Ali Jadbabaie is the JR East Professor of Engineering in the Department of Civil and Environmental Engineering, the associate director of IDSS, and the director of MIT's Sociotechnical Systems Research Center. He is a recognized expert in the fields of network science, decision and control theory, and multi-agent coordination.

Among other areas, his current research focuses on the interplay of dynamic systems and networks, with a specific emphasis on multi-agent coordination and control, distributed optimization, network science, and network economics.

Fundamental Limits of Learning Project

The goal of this project, involving eight LIDS principal investigators and sponsored by the Defense Advanced Research Projects Agency (DARPA), is to develop a foundational theory for learning with a focus on fundamentals and rigorous, provably correct approaches. Over the past few years, there has been an explosion in the use of machine learning algorithms and techniques in a variety of applications ranging from advertising to image recognition and classification, natural language processing, and automatic translation. Despite these successes, many fundamental questions in machine learning are still open, and a theory that quantifies the tradeoffs and elucidates what can and cannot be done with a given amount of data is missing. The LIDS principal investigators are tackling this problem head on, studying fundamental issues related to estimation network structures such as measuring associations and noisy data and addressing large-scale optimization problems.

Basic Research Challenge Program on Decentralized and Online Optimization

This project, sponsored by the Office of Naval Research (ONR), involves three LIDS principal investigators (Professors Jadbabaie, Ozdaglar, and Jaillet) across two different teams. The goal of the project is to understand the basic research challenges in distributed and online optimization and resource allocation problems.

Evolution of Cultural Norms and Dynamics of Sociopolitical Change Initiative

This multidisciplinary, multi-investigator university research initiative funded by the Army Research Office brings together investigators from LIDS, IDSS, and other universities, including systems theorists, computer scientists, economists, and political scientists, to analyze and model patterns of communication and collective decision making in networked societies. The goal is to gain a better understanding of social and political change.

Dynamic Pricing and Word of Mouth

Principal investigator Jadbabaie and LIDS postdoctoral scholar Amir Ajorlou have developed a theory that provides new insights on how offering frequent zero-price sales can allow online retailers to maximize their profit by dynamically adjusting prices. In situations where information about a product travels only via word of mouth, offering frequent zero-price sales allows retailers to balance their immediate profit with growing their customer base and increasing sales. Through holding such sales, firms can use free offers to attract consumers who would not otherwise buy the product but have friends (or friends of friends) who might be willing to pay more. In essence, firms can use consumers who might not want to buy the product to reach others who might be willing to pay a great deal more.

Collective Behavior and Collective Decision Making

Principal investigator Jadbabaie and his colleagues have developed a theory of how individuals might aggregate disparate sources of information in their social networks and how the structure of these networks might impact whether they collectively reach optimal decisions when they do not have access to all of the information available to their peers. Given the well-known cognitive burden of keeping track of everyone's information at all times, they study how groups can reach correct decisions even when individuals cannot fully recall all pieces of information revealed to them in the past. They also identify how, under such deviations, individuals may double count information delivered from the same source but through multiple paths.

Patrick Jaillet

The research of Professor Jaillet (Dugald C. Jackson Professor and ORC co-director) and his group focuses on online learning and optimization problems. In particular, they address online and dynamic versions of assignment/matching and secretary problems as well as some of their generalizations. Their research deals with provable results (algorithmic design and analysis) on how to solve such problems under uncertainty, with or without explicit stochastic modeling of uncertainty. Methodological tools include those from online optimization (competitive analysis), stochastic optimization (robust analysis), online learning (min-max regret analysis, Bayesian updates), game theoretic concepts (price of anarchy), and their integrations.

Motivating applications include routing and location problems that arise from transportation and logistics networks, data communication and sensor networks, and/or autonomous multi-agent systems, as well as dynamic resource allocation problems in various applications arising from the digital economy (search engines and online auctions), health care (kidney exchange programs), and social interactions (job search and house exchanges).

Professor Jaillet's research group at MIT this past academic year included nine PhD students from ORC (Maximilien Burq, Arthur Flajolet, Virgile Galle, Chong Yang Goh, Swati Gupta, Nikita Korolko, Sebastien Martin, Konstantina Mellou, and Julia Romanski) and four SM students from Leaders for Global Operations (Scott Foster, Lila Fridley, Boyan Kelchev, and Jonathan Zanger). His research group in Singapore included one postdoc (Jie Chen from SMART) and five PhD students (Phong Nguyen and Haibin Yu from the National University of Singapore, Anatoliy Prokhorchuk from Nanyang Technological University, Meghna Lowalekar from Singapore Management University, and Gary Goh from the Singapore University of Technology and Design).

Current funded research programs originate from ONR ("Online Optimization and Learning under Uncertainty" and "Decentralized Online Optimization in Multi-Agent Systems in Dynamic and Uncertain Environments") and SMART ("Future Mobility," a large project involving several other MIT principal investigators).

Sertac Karaman

Professor Karaman, the Class of '48 Career Development Chair, carries out research in the areas of control theory, optimization, formal methods, stochastic processes, and applied probability, with applications to robotics, mobile sensor networks, cyber-physical systems, and data-driven application systems. His current work focuses on computing for complex, agile autonomous vehicles.

High-Dimensional Computing for Autonomous Systems

A number of problems in autonomy are plagued by the curse of dimensionality: the computational effort required to solve many problems of inference, estimation, control, and learning scales exponentially with increasing numbers of variables describing the state of the underlying system. In their recent work, Professor Karaman and his group have proposed novel algorithms that run in polynomial time for all problem instances admitting a certain low-rank structure. The key insight involves compressing a certain cost-to-go function (as in data compression) and working on the compressed version of the function iteratively, in order to improve it towards an optimal one. The group rigorously proved that the new compression-based algorithms run in polynomial time with increasing degrees of freedom. The new algorithms also scale polynomially with the "rank" of the value function. Hence, if the value function has a low-rank structure due to, for example, geometry and physics (in other words, if the value function is "compressible"), the new algorithms provide substantial computational savings. The group's work has applications in a range of areas, including inference, state estimation, control synthesis, differential games, planning under uncertainty, and reinforcement learning.

High-Throughput Computing for Autonomous Systems

Being able to perceive its environment is of utmost importance to any robotic vehicle. However, most robot perception problems are high-dimensional large-scale inference problems (e.g., due to high-rate sensing data). Professor Karaman and his group designed task-driven perception algorithms that implement attention mechanisms to extract the most meaningful data set enabling efficient perception. They showed that such problems can be formulated as subset selection problems, and they proved that the resulting problems are NP (nondeterministic polynomial-time) hard. Yet, using submodularity properties, they were able to design efficient algorithms providing bounds on sub-optimality. They applied these algorithms to navigation using visual sensors. In experiments, they showed that these techniques are essential in enabling high-performance agile robotic vehicles such as fast-flying drones.

Low-Energy Computing for Autonomous Systems

Computers that can enable fingernail-sized flapping-wing robotic vehicles or control small insects must run under tight energy budgets. A state-of-the-art insect-sized flapping-wing robotic vehicle requires 100 milliwatts of power to lift itself. Unfortunately, state-of-the-art general-purpose computers need at least two orders of magnitude more power to run algorithms that make such robotic vehicles autonomous. In joint work with Professor Vivienne Sze and the Energy-Efficient Multimedia Systems Group, Karaman and his group are co-designing the algorithms and the computing hardware. The resulting systems save at least an order of magnitude in computing energy and may become the computing elements for next-generation miniature robotic vehicles.

Autonomy-Enabled Transportation Systems

Autonomous vehicles hold the potential to revolutionize transportation, particularly in urban centers. For example, drones may enable the next generation of high-throughput, low-delay urban logistics services. This research area aims to better understand the potential impact of autonomous vehicles by studying new autonomous vehicle systems. For instance, one potential technology allows trucks to follow each other closely, much closer than what human drivers can safely achieve. Thanks to this new technology, trucks can form platoons on highways, allowing them to save energy. However, forming long platoons requires some vehicles to wait for others, resulting in excessive transportation delays. In their most recent work, Karaman and his group developed a queuing model to understand this energy-delay tradeoff. In particular, they analyzed both open-loop (timetable) policies and feedback policies and showed that the addition of feedback benefits the system very little. These ideas are not limited to truck platooning. They can be applied to ride sharing to reduce transportation costs, formation flight for commercial airlines to save energy, and many other areas.

Sanjoy Mitter

Professor Mitter's research focuses on information theory and stochastic control along with issues related to learning manifold structures from noisy data. A variety of stochastic control problems involving minimization of directed information subject to distortion constraints have been considered. These problems are related to his earlier work on sequential rate distortion theory (conducted jointly with Sekhar Tatikonda, now a faculty member at Yale University). The main question that needs to be answered is how to quantify the amount of information needed to achieve a desired control objective. It is conjectured that information-theoretic optimal allocation of resources takes place between identification of parameters and control objectives in problems of stochastic adaptive control.

Eytan Modiano

Professor Modiano (associate director of LIDS and associate head of the Department of Aeronautics and Astronautics) leads the Communications and Networking Research Group (CNRG), which consists of eight graduate students and one postdoc. The primary goal of CNRG is the design of architectures for aerospace networks that are cost effective and scalable and meet emerging needs for high-data-rate and reliable communications. In recent years, the group has focused on robust network designs, wireless networks, data center networks, and interdependent cyber-physical networks.

During the past year, CNRG members authored more than a dozen journal and conference papers in areas such as robust network designs, dynamic resource allocations, and fundamental performance limits in wireless networks. Their contributions included the design of throughput optimal network routing and scheduling algorithms, the design of transmission scheduling algorithms for wireless networks with delay-sensitive traffic, the study of robustness in interdependence networks, and the design of ultra-high-capacity data center network architectures using optical switching technology. The most significant contributions are highlighted below.

Over the past 10 years, CNRG has developed a number of network control algorithms for wireless networks. However, these schemes were developed assuming that they can be applied universally, to all nodes in the network. The group is now focusing on the design of network control algorithms that can operate efficiently using legacy network devices. Their approach involves a novel overlay architecture for implementing optimal network control algorithms over a legacy network using a limited number of overlay nodes. This new paradigm allows optimal control algorithms to be incrementally deployed alongside existing schemes, thus providing a migration path for new control algorithms and the promise of dramatic improvements in network performance.

In a related area, Professor Modiano and graduate student Abhishek Sinha developed a new framework for throughput-optimal data dissemination in wireless networks in the presence of an arbitrary mix of unicast, broadcast, multicast, and anycast traffic. In particular, they developed an online dynamic policy, universal max-weight (UMW), that optimally routes packets and schedules transmissions. UMW is the first throughput-optimal algorithm in the context of generalized network flow problems.

The group continued work on its project addressing the robustness of interdependent networks. Many engineering systems involve interactions between two or more networked systems. Cyber-physical systems, for example, consist of networked computer systems that are used to control physical systems such as the power grid, water or gas distribution systems, and transportation networks. While this cyber-physical interaction is critical for the functionality of the overall system, it also introduces vulnerabilities in the form of interdependence failure cascades, with failures in the cyber network leading to failures in the physical network and vice versa. Over the past year, Professor Modiano and his student Jianan Zhang developed a novel model for interdependence networks and new methods for assessing reliability in such networks. Their novel approach will lead to new understandings of interdependence between networked systems such as the power grid and the communication networks used to control the grid.

In recent years, Professor Modiano and his group have been pursuing industrial collaborations to increase the impact of their work on practical systems. To that end, the group is conducting a joint project with BBN Technologies on resilient overlay networks and a project with Qualcomm on mission-critical communications. Also, it is involved in a close collaboration with researchers at Lincoln Laboratory on the design of network architectures and protocols for military communications.

CNRG's research crosses disciplinary boundaries by combining techniques from network optimization, queuing theory, graph theory, network protocols and algorithms, hardware design, and physical layer communications.

Asuman Ozdaglar

Professor Ozdaglar (EECS associate department head and Keithley Professor of Electrical Engineering and Computer Science) and her research group focus on modeling, analysis, and optimization of large-scale, dynamic multi-agent networked systems. Their research draws on advances in game theory, optimization theory, dynamical systems, and stochastic network analysis. It focuses on both investigating substantive problems in these areas and developing new mathematical tools and algorithms for the analysis and optimization of these systems and for processing large-scale data.

In a recent project, the group investigated Bayesian learning by consumers in online review systems. Despite the growing importance of online reviews, there has been relatively little work on whether they are effective in aggregating the information of millions of users. The group's earlier work, wherein individuals updated their beliefs on the basis of observation of others who had made similar decisions in the past, does not give a uniformly encouraging answer: individuals may herd on past decisions, preventing efficient aggregation of information. Review systems differ from pure observational learning problems because some individuals provide their own assessments of the quality of the product or seller in question, potentially providing more information to users but introducing another challenge: those purchasing a good are not a representative sample of users but, rather, are selected from among those more likely to have a positive assessment. Professor Ozdaglar and her group, in work conducted jointly with graduate student Ali Makhdoumi, Professor Daron Acemoglu from the Department of Economics, and Azarakhsh Malekian (former LIDS postdoc and now an assistant professor at the University of Toronto), showed how this problem can be analyzed in a general setting and provided general conditions for efficient aggregation of information. Extending previous work, they were also able to provide a tight characterization of the speed of learning. This latter result is particularly important, since many platforms seek to attract users by providing them with better information so as to improve their decisions. The group's characterization clarifies how platforms can best design review systems to achieve faster learning and attract more users.

Professor Ozdaglar's group also works on developing novel algorithms for processing large-scale data and distributed optimization over networks. Many statistical learning problems are based on empirical risk minimization, wherein a parameter of interest is estimated by optimizing an objective function defined as the sum of a loss function averaged over data and a penalty term related to model complexity. These problems motivate incremental algorithms that can exploit the additive problem structure and lead to efficient algorithms that are scalable (both with the dimension of the problem and the size of the data sets). The group's recent work, conducted jointly with former postdoc Mert Gurbuzbalaban (now an assistant professor at Rutgers), graduate student Denizcan Vanli, and Professor Pablo Parrilo, has made significant contributions to this area, filling in important gaps in theory and leading to improved algorithms. These projects are supported by DARPA and ONR.

In another recent collaboration, with former postdoc Insoon Yang (now an assistant professor at the University of Southern California) and Daron Acemoglu, Professor Ozdaglar is studying the efficiency implications of upgrading to clean energy technologies and the penetration of renewables in the energy-producing sector. Their model generalizes an industry equilibrium model of entry, exit, and firm dynamics proposed in the economics literature by adding a technology upgrade component. Given a time-varying carbon tax sequence, they first characterize the nonstationary dynamic equilibrium. They show that the optimal upgrading, entry, and exit strategies are characterized as thresholding rules on the firm's generation efficiency. They then determine a time-varying carbon tax that maximizes the utility of consumers subject to the equilibrium, which leads to a bilevel dynamic optimization problem. Their model will provide qualitative insight on the role of policy instruments in sustainable development of clean generating technologies to reduce greenhouse gases as well as an optimal strategy for firms to upgrade technologies and enter and exit the energy-producing sector. This work is supported by a National Science Foundation (NSF) Cyber-Physical Systems grant (with Berkeley, Michigan, and Vanderbilt), a joint Masdar Institute/MIT project, and a Siebel grant.

Pablo A. Parrilo

Professor Parrilo and his research group are focused on mathematical optimization, systems theory, and control, with emphasis on development and application of computational tools based on convex optimization and algorithmic algebra.

Parrilo and PhD student Diego Cifuentes developed and are currently investigating chordal networks, a novel representation of structured polynomial ideals. The sparsity structure of a polynomial ideal is often described through a graph that captures the interactions among decision variables. Chordal networks provide a computationally convenient decomposition of the polynomial ideal into simpler (triangular) polynomial sets while preserving the underlying graphical structure. Parrilo and Cifuentes showed that many interesting families of polynomial ideals involve compact chordal network representations, even though the number of components might be exponentially large. Chordal networks can be computed for arbitrary polynomial systems using a refinement of their earlier chordal elimination algorithm. They applied their methods to examples from algebraic statistics and vector addition systems and showed that algorithms based on chordal networks can outperform existing techniques by orders of magnitude.

Professor Parrilo and PhD student Frank Permenter developed a new dimension reduction method for semidefinite programming through the use of Jordan algebras. Specifically, they considered orthogonal projections satisfying certain invariance conditions (the optimal set should be contained in its range). Adding this constraint to the original optimization problem yielded an equivalent primal-dual pair over a lower-dimensional symmetric cone—namely, the cone of squares of a Jordan subalgebra of symmetric matrices. They also proposed a simple lattice-theoretic algorithm for minimizing the rank of this projection and, hence, the dimension of the cone. Through the theory of Jordan algebras, the proposed method easily extends to linear programming, second-order cone programming, and, more generally, symmetric cone optimization. The method has been shown to be very effective with several examples from the literature.

In joint work with Hamza Fawzi and James Saunderson, Parrilo has developed a new methodology for handling the exponential/relative entropy cone using symmetric cone solvers. Their approach is based on highly accurate rational approximations of the logarithm function. The key to this approach is that their rational approximations, by construction, inherit the (operator) concavity of the logarithm. Importantly, their method extends to the matrix logarithm and other derived functions such as the matrix relative entropy, providing new semidefinite optimization-based tools for convex optimization involving these functions. They also produced and implemented their method for the MATLAB-based parser CVX. They compared their method with the existing successive approximation scheme in CVX and showed that it can be much faster, especially in the case of large problems.

Yury Polyanskiy

Professor Polyanskiy conducts research in the areas of mathematics of information (information theory), coding theory, and the theory of random processes. His current work focuses on non-asymptotic characterization of the performance limits of communication systems, non-Shannon information measures, redundant circuits, and probabilistic methods in combinatorics.

Joint Source-Channel Coding with Feedback

This work quantifies the fundamental limits of variable-length transmission of a general (possibly analog) source over a memoryless channel with noiseless feedback, under a distortion constraint. Polyanskiy and his group considered excess distortion, average distortion, and guaranteed distortion. In contrast to the asymptotic fundamental limit, a general conclusion is that allowing variable-length codes and feedback leads to a sizable improvement in the fundamental delay-distortion tradeoff. In addition, they investigated the minimum energy required to reproduce k source samples with a given fidelity after transmission over a memoryless Gaussian channel and showed that the required minimum energy is reduced with feedback and an average (rather than maximal) power constraint.

Minimum Energy to Send k Bits over Multiple-Antenna Fading Channels

Polyanskiy and his group investigated the minimum energy required to transmit k information bits with a given reliability over a multiple-antenna Rayleigh block-fading channel, with and without channel state information (CSI) at the receiver. No feedback was assumed. It is well known that the ratio between the minimum energy per bit and the noise level converges to -1.59 dB as k moves toward infinity, regardless of whether or not CSI is available at the receiver. The group showed that the lack of CSI at the receiver causes a slowdown in the speed of convergence to -1.59 dB as k → ∞ (relative to the case of perfect receiver CSI). Specifically, they showed that in the no-CSI case the gap to -1.59 dB is proportional to [(log k)/k]1/3, whereas when perfect CSI is available at the receiver this gap is proportional to 1/√k. In both cases, the gap to -1.59 dB is independent of the number of transmit antennas and channel coherence time. Numerically, they observed that when the receiver is equipped with a single antenna, one needs to transmit at least 7 × 107 information bits to achieve an energy per bit of -1.5 dB in the no-CSI case, whereas 6 × 104 bits suffice for the case of perfect CSI at the receiver.

Strong Data-Processing Inequalities for Channels and Bayesian Networks

The data-processing inequality—that is, $I(U; Y) \leq I(U; X)$—for a Markov chain $U \rightarrow X \rightarrow Y$ has been the method of choice for proving impossibility (converse) results in information theory and many other disciplines. Various channel-dependent improvements (labeled strong data-processing inequalities, or SDPIs) on this inequality have been proposed both classically and more recently. Polyanskiy and his group surveyed known results relating various notions of contraction for a single channel and considered the basic extension: given SDPI for each constituent channel in a Bayesian network, how can one produce an end-to-end SDPI?

Their approach was based on the Evans-Schulman method, which was demonstrated for three different kinds of SDPIs: the usual Ahlswede-Gács contraction coefficients (mutual information), Dobrushin's contraction coefficients (total variation), and finally the F_I curve (the best possible nonlinear SDPI for a given channel). Resulting bounds on the contraction coefficients were interpreted as probabilities of site percolation. As an example, they demonstrated how to obtain SDPIs for an n-letter memoryless channel with feedback given an SDPI for n = 1.

Finally, Polyanskiy and his group described a simple observation on the equivalence of a linear SDPI and a comparison with an erasure channel (in the sense of "less noisy" order). This led to a simple proof of a curious inequality of Samorodnitsky and shed light on how information spreads in the subsets of inputs of a memoryless channel.

Philippe Rigollet

Professor Philippe Rigollet's research focuses on understanding the fundamental limitations associated with high-dimensional statistical problems and developing methodology that performs optimally within these limits. Specifically, Rigollet's group produces methods with provable guarantees that blend tools from discrete and continuous optimization, statistical modeling, information theory, and computational complexity.

Learning Determinantal Point Processes

Determinantal point processes (DPPs) are a family of probabilistic models that arose from the study of quantum mechanics and random matrix theory. Discrete DPPs have found numerous applications in machine learning, including in document and timeline summarization, image search and segmentation, audio signal processing, bioinformatics, and neuroscience. What makes such models appealing is that they exhibit repulsive behavior and lend themselves naturally to tasks in which returning a diverse set of objects is important. While there are fast algorithms for sampling, marginalization, and conditioning, much less is known about learning the parameters of a DPP. Rigollet and his group have developed the first statistical results for learning DPPs, including:

- Characterizing the rates of convergence for the maximum likelihood estimator (MLE): by studying the local and global geometry of the expected log-likelihood function, they are able to establish rates of convergence for the MLE and provide a complete characterization of cases in which these rates are parametric.
- Determining optimal rates of convergence for this problem: these rates are achievable with the method of moments and are governed by a combinatorial parameter that is strongly linked to the independence structure of the process.
- Providing fast algorithms to implement the method of moments efficiently using tools from combinatorial optimization.

Current research efforts are focused on showing that the log-likelihood landscape does not have spurious local minima. Results in this direction would allow for provable guarantees for simple algorithms that implement the maximum likelihood estimator.

Learning under an Algebraic Constraint

The information era has had a transformative effect in virtually all aspects of science, with data analysis becoming a preponderant guide for scientific discovery. This is the case, for example, in cryo-electron microscopy (cryo-EM), which was selected by the journal *Nature Methods* as the 2015 method of the year "for its newfound ability to solve protein structures at near-atomic resolution." This promising imaging method comes with new challenges that resist known signal processing techniques. In cryo-EM, a molecule is imaged from different, unknown viewing directions that can be thought of as latent three-dimensional rotations acting on each molecule sample before the image is captured. The presence of the latent rotation has a singular effect on not only the computational complexity but also the statistical complexity of this problem. This feature is shared by various other examples of the same kind in areas such as robotics, structural biology, radar, crystalline simulations, and image registration problems in a number of important contexts, including geology, medicine, and paleontology. Unlike more traditional statistical problems in which a linear underlying structure is often built into the model, data-driven science generates problems with algebraic but often nonlinear structures that are typically dictated by physics.

Rigollet's group has developed a principled approach to analyze this family of problems together with a general procedure to construct computationally efficient algorithms using low-rank tensor decompositions. Importantly, these methods can be proved to be statistically optimal and therefore make the most efficient use of the collected data. These problems show new and surprising statistical limitations; for example, they require more observations than standard problems in signal processing to estimate at a certain precision.

Rigollet's group has focused on a specific model, multi-reference alignment. It is one of the simplest models able to capture fundamental aspects of this class of problems, rendering it ideal for theoretical study. Current efforts aim at extending these results and methods to cryo-EM. This work is supported by NSF.

Optimal Transportation for Statistics

The theory of optimal transportation has had a considerable impact in both pure and applied mathematics since its introduction by Gaspard Monge in the 18th century. Recently, it has found a new application in machine learning and specifically in computer imaging, where it has become mainstream to compute distances between images. This recent advent can be explained by a newfound use of entropic regularization that is pervasive to various scientific disciplines, including statistics. Rigollet's research efforts aim at shedding new light on the use of entropy in such problems. Current efforts are focusing on using entropy for statistical regularization, akin to its use in aggregation.

Causal Inference

Machine learning algorithms have been successfully employed in a wide range of applications such as biology, sociology, computer vision, marketing, and finance.

These algorithms seldom rely on modeling assumptions but instead focus on optimizing prediction performance. This agnostic approach has led to spectacular results, but at the cost of little to no understanding of the underlying prediction process. In various scientific disciplines such as biology, medicine, sociology, and neuroscience, understanding the mechanisms that lead to accurate predictions is paramount to advance knowledge and technology. A good representation of a complex system should enable predicting the state of one component given others, as is traditional in machine learning; however, the representation should also capture the causal relationships between system components. Current methods for causal inference can describe only local aspects of large, complex systems. At the same time, machine learning techniques can be scaled to large systems but are ill suited for the purpose of learning causal effects, as they are inherently designed to leverage observational data (i.e., data collected in a passive fashion). While some causal effects can be discovered in this manner, a full understanding of the causal relationships in a system requires interventional data obtained by deliberately and carefully altering one or more components. Rigollet's group is currently advancing on two fronts, as follows.

- Learning from observational data: by introducing new, average case statistical models for causal inference, Rigollet's group has made unprecedented advances in the context of learning causal models from observational data only. In particular, unlike the previous literature, this work allows for arbitrary causal cycles.
- Learning from interventional data: Rigollet's group has identified a set of necessary and sufficient conditions for interventions to uniquely determine a causal system, potentially with cycles. Based on such assumptions, Rigollet is in the process of deriving the first results characterizing the sample complexity associated with estimating such causal systems using the maximum likelihood approach.

Mardavij Roozbehani

Principal research scientist Roozbehani led three parallel research thrusts related to mathematical modeling, optimization, and control for networked systems. Each thrust focused on a single application area: energy networks, transportation networks, and financial networks. The common theme in these thrusts was their emphasis on robustness analysis, metrics for risk and fragility, and tradeoffs between efficiency and risk. This is aligned with the long-term goal of Roozbehani's research: to understand the sources of robustness, fragility, and systemic risk in large infrastructure networks. Funding for his research comes from the National Science Foundation, the MIT-Masdar Institute Cooperative Program, and the Kuwait-MIT Center for Natural Resources and the Environment.

Robustness and Risk in Transportation Networks

Roozbehani's work with postdoctoral associate Yasin Yazicioglu on the robustness of transportation networks revealed new sources of fragility in these networks resulting from limited information exchange and decision making based on local observations. This research led to novel and unexpected results whereby locally optimal decisions or local infrastructure improvements have negative consequences and degrade global performance. These earlier results paved the way for the development of novel control algorithms aimed at establishing speed limits or alternative incentives for congestion pricing to increase the robustness of transportation networks to disturbances and link failures.

Scalable Control Architectures for Energy Networks

Roozbehani's work with MSc student Allison Fero of the Technology and Policy Program (TPP) led to new methodologies and distributed algorithms for maximizing system-wide reliability in energy networks. In collaboration with the Tata Center for Technology and Design and motivated by electrification problems in underprivileged and geographically remote areas in India, this research focused on the design of a computationally scalable communication and control architecture for the interconnection of small microgrids. From an application perspective, this work developed an optimization-based control architecture for resource sharing that offers optimal load scheduling and energy sharing decisions subject to system dynamics, power balance constraints, and congestion constraints while maximizing network-wide reliability. Such resource sharing technologies were shown to reduce costs significantly, with a huge impact on the target communities. The theory and methodological engine behind this technology is distributed optimal control for a network of linear quadratic regulators (LQRs) coupled through network flow decisions. The LQR solution is combined with network flow dual decomposition to generate a fully decomposed algorithm for determining the dynamic programming solution of the LQR subject to network flow couplings.

Coordination Games and Risk in Financial Networks

In collaboration with EECS graduate student Thiago Vieira, Roozbehani developed a new model of risk in financial networks arising from coordination games among investors with heterogeneous information structures. This model extends the debt global game models of coordination risk previously developed in the literature. The extended model enabled Roozbehani and Vieira to study coordination risk over a finite time horizon and introduce new information structures for investors. Their main results showed that, under certain conditions, positive public information for a financed project can cause investors to react negatively, decreasing the project's chances of success. The findings also revealed that when investors receive private and public noisy signals of past actions by other investors, a herding behavior occurs that can increase the fragility of the system.

Other work during the past year focused on electricity market designs for renewables (with EECS/TPP student Ian Schneider), methodologies for aggregation of distributed demand response resources and creation of virtual storage (with MIT postdoctoral associate Daria Madjidian), scheduling of flexible demand resources (with Donatello Materassi of the University of Tennessee), and development of a unifying framework for analysis and quantification of resilience and fragility in networks across several application domains (with EECS PhD candidate Tuhin Sarkar).

Devavrat Shah

Professor Shah and his research group are currently involved in developing theoretical foundations and algorithmic solutions for two classes of questions. The first set of questions arise in the context of "social data processing," that is, using data generated by people to make better decisions. At the highest level, this can help businesses operate efficiently, make better policy decisions, improve social living by enabling efficient labor markets, and so on. The big challenge here is that data are often highly unstructured and/or large in quantity, and thus fundamentally new statistical solutions and computationally efficient algorithms are needed. The second set of questions concern efficient operations of the substrate on which such large-scale data-processing systems are built: the data center. Specifically, the efficient operation of a network is crucial to the performance of distributed computation systems running in data centers. The goal is to schedule transfers of packets in data centers at the nanosecond time scale. Some details on the progress made in the past year in the context of these issues are provided below.

Estimating a Latent Variable Model Using Collaborative Filtering

To enable decision making using social data, it is essential to identify generic characteristics that can help in utilizing flexible models for making predictions. This is precisely what Shah and his group have done: social data are (or should be) anonymous. That is, from the data-processing perspective, it should not matter who has generated the data. To put it another way, the overall conclusion should remain invariant if one renames the individuals who have generated the data. For example, the results of a democratic election should not change even if the voters' names change, as long as the total number of votes for each candidate remains the same.

Anonymity, seemingly a weakness, somewhat surprisingly turns out to be the primary weapon in addressing the challenge of developing tractable and flexible models for social data. Mathematically, anonymity can be viewed as an underlying "probabilistic model" with a certain "exchangeability" property. A remarkable development in mathematical statistics provides a crisp, non-parametric characterization for such models: the latent variable model.

Shah and his group use the latent variable model to examine the question of designing personalization or recommendation systems such as those used by Netflix, YouTube, Amazon, and Spotify. Here the goal is using the history of an entire population's preferences to predict what movies, music, books, or other products individual consumers may like but have not already experienced. On one hand, the question is: what is the best algorithm to design for that end goal using the non-parametric model emerging from exchangeability? On the other hand, that question has been around since the dawn of the e-commerce era.

Shah and his group found a surprising answer: the popular algorithm called collaborative filtering (which in some form has been around since the early 1990s) is statistically optimal in estimating a latent variable model. This explains the success of such a vastly popular algorithm (e.g., YouTube and Amazon reportedly use it) that remained a mystery for decades.

In ongoing work, the group is exploring the application of the model and algorithm to social networks, time series analysis, and image de-noising, among other domains.

Centralized Scheduling in Data Centers

Data centers have become the backbone of the modern distributed computation environment. Such systems are fundamentally built using commodity machines, and hence the networks enabling communication between machines in data centers have become crucial in making computation efficient. Due to the legacy of Internet architecture and the ease of system design, modern data centers have distributed scheduling and congestion control algorithms. However, this leads to poor latency performance and loss of capacity.

Unlike the Internet, data centers are fundamentally well-controlled environments. As a consequence, the reasoning behind the distributed network design for the Internet is not valid. Based on this observation, Shah and his group have proposed a centralized network architecture to address this question and have built prototypes of the architecture.

Suvrit Sra

Principal research scientist Suvrit Sra has been with LIDS since January 2015. His research spans two broad categories: machine learning and optimization and pure and applied mathematics (including statistics).

Sra's primary research focuses on in large-scale optimization for machine learning and related areas. Over the past year, with his collaborators and students, he has continued to expand the repertoire of fast optimization algorithms. In particular, his work continues to address the theory of large-scale nonconvex optimization, which underlies the popular subarea called "deep learning." His work in this area, for the first time, provides an analysis of when and how fast we can expect certain large-scale optimization algorithms to reach a solution. The last year has witnessed new developments by Sra and his collaborators on this topic, as well as follow-ups by colleagues in the machine learning community.

Furthermore, together with collaborators at MIT, Sra is investigating theoretical properties that (broadly viewed) seek to explain when and why deep learning models work. Sra held a joint tutorial with Francis Bach at the 2016 Neural Information Processing Systems conference on the state of the art in large-scale nonconvex optimization algorithms; in June 2017, he held a series of invited lectures at the Machine Learning Summer School in Germany, again on the topic of optimization in machine learning. These tutorials and lectures helped inspire many new minds getting into the area.

Sra and his students and collaborators continued to build inroads into the topic of "geometric optimization," a new subarea of mathematical optimization that Sra has been helping to establish on a more solid computational complexity foundation. Over the past year, his work on this topic has generated faster algorithms that rely on differential geometry. In particular, this work is extending the scope of the well-established area of "convex optimization" to more general settings; in a few years, this material should be textbook knowledge.

Finally, Sra continues to expand his work in discrete probability theory (on determinantal point processes and the richer family of measures called strongly Rayleigh measures) in collaboration with other MIT faculty members.

John Tsitsiklis

Professor Tsitsiklis (LIDS director and Clarence J. LeBel Professor of Electrical Engineering) and his students work on system modeling, analysis, optimization, and control in possibly stochastic and dynamic environments and in the possible presence of multiple agents with conflicting interests. Their research activities have focused on developing methodologies, mathematical tools of broad applicability, and computational methods. Motivating applications for recent work have come from domains such as computer networks and social networks. Examples of the group's recent activities are provided below.

Private Search and Optimization

Motivated by the growing concern for the privacy of agents who interact with data providers, Tsitsiklis, student Zhi Xu, and Professor Kuang Xu, a LIDS alum, have initiated a new research direction. Specifically, they have introduced a new class of problems whereby an agent attempts to locate an object through queries to a database but wishes to do so without allowing an eavesdropper to learn the object's location. In a different variant of the problem, an agent wishes to optimize an objective in a setting where part of the objective is privately known by the agent but there is also dependence on information available from a data provider. The goal is to have the agent interact with the data provider and find an optimum solution but without revealing too much information (i.e., without allowing the data provider to also identify an optimal solution). Both of the above problems are stylized abstractions of all-too-common situations in which an agent needs information but does not wish the information provider to identify the conclusions to be deduced from that information. For both problems, Tsitsiklis and his colleagues have identified strategies under which the agent can achieve the given goal with the least possible obfuscation effort (i.e., redundant queries to the data provider).

Sensitivity Theory for a Class of Hybrid Dynamical Systems

There is a popular and much-used scheduling policy in computer networks known as the max-weight policy. This policy results in "hybrid" dynamics, whereby the derivative of the state vector is a piecewise constant function of the state. In the presence of stochastic fluctuations of the input to the scheduling system, one obtains perturbed trajectories. Much of the analysis of such policies depends on an understanding of the effect of such perturbations. Together with visiting student Arsalan Sharifnassab, Tsitsiklis and his group developed a theory on the effect of perturbations for a rather general class of hybrid systems. They then applied it to the special case of max-weight scheduling policies to derive theoretical results stronger than those previously available and also to settle some problems that were left open in an earlier LIDS doctoral thesis. There is much remaining work in progress that will generalize these results and allow their use in concrete settings.

Caroline Uhler

Henry L. and Grace Doherty Assistant Professor Caroline Uhler carries out research in the areas of mathematical statistics and optimization with applications to genomics and cell biology. Her current work focuses on causal inference and the use of total positivity for modeling.

Causal inference is a cornerstone of scientific discovery. It is of particular interest to determine causal or directional structures among variables based on observational data, since conducting randomized controlled trials is often impractical or prohibitively expensive.

Combinatorics of Markov Equivalence Classes

Unfortunately, observational data alone cannot uniquely identify a directed graphical model since different directed graphical models can satisfy the same conditional independence relations (such graphs are Markov equivalent). It is therefore important to understand the set of Markov equivalence classes and their sizes. Together with MEng student Adityanarayanan Radhakrishnan and former postdoc Liam Solus, Uhler initiated combinatorial enumeration of Markov equivalence classes on a fixed undirected graph, thereby recasting this important statistical problem into the language of combinatorial optimization.

Limitations for Learning Bayesian Networks

There are two main classes of algorithms for learning causal graphs from observational data, namely constraint-based and score-based methods. Constraint-based methods treat causal inference as a constraint satisfaction problem in which the constraints are given by conditional independence relations. Score-based methods use greedy search algorithms to maximize a penalized likelihood score over the space of directed graphical models. Constraint-based methods have the advantage of being scalable to problems with thousands of nodes. However, such methods are very sensitive to errors in the constraints (i.e., the conditional independence relations), which need to be estimated from data. Uhler's research based on algebraic geometry showed that constraint-based methods require unrealistic sample sizes to estimate the correct causal structure, thereby implying severe limitations for such methods. Conversely, the greedy nature of score-based methods usually leads to better graph recovery rates for the same sample size but with a higher computational cost due to the large search space. Uhler's group is currently working on quantifying these statistical computational tradeoffs for causal inference and designing optimal algorithms for given sample sizes and computational budgets.

Permutation-Based Causal Inference from Observational Data

Greedy equivalence search, a score-based method, is one of the most widely used causal inference algorithms. It performs a clever greedy search on the space of Markov equivalence classes and comes with theoretical consistency guarantees. While it is an open question to find a combinatorial formula for the number of Markov equivalence classes on p nodes, it is conjectured from simulations that on average every Markov equivalence class is of size four. Since the space of all permutations is much smaller than the space of all Markov equivalence classes and since it is relatively easy to learn a directed graph given an ordering, it is natural to consider a greedy search on the space of permutations. Together with PhD student Yuhao Wang and former postdoc Liam Solus, Uhler proved that greedy permutation search is consistent, meaning that it outputs the correct equivalence class under the same conditions as greedy equivalence search but with a drastic reduction in the size of the search space. This algorithm can be seen as a simplex-type method in a particular class of polytopes, one that Uhler introduced in collaboration with Josephine Yu, Charles Wang, and former postdoc Fatemeh Mohammadi.

Permutation-Based Causal Inference from Interventional Data

Most excitingly, Uhler and her group recently showed that their greedy permutation search can be extended to deal with a mix of observational and interventional data. The interventional adaptation of the greedy permutation search is the first algorithm with provable consistency guarantees in the presence of observational and interventional data. This is now a fundamentally important problem due to recent technological developments in genomics that allow the generation of observational and interventional single-cell gene expression data at a very large scale (perturb-seq). In collaboration with Aviv Regev at the Broad Institute and PhD student Karren Dai Yang, Uhler is currently working on applications of her algorithm to the analysis of perturb-seq data to learn specific gene regulatory networks.

Totally Positive Gaussian Graphical Models

In a recent collaboration, Uhler and her group analyzed properties of distributions that are multivariate totally positive of order two (MTP2). This property, introduced in the 1970s, is a stronger version of positive association, an important notion in probability theory and statistical physics. A Gaussian distribution is MTP2 if and only if the inverse covariance matrix is an M matrix (i.e., all off-diagonal entries are non-positive). In collaboration with Steffen Lauritzen and Piotr Zwiernik, Uhler and her team recently showed that the MLE in Gaussian MTP2 distributions already exists for two observations. In addition, they showed that the MLE has intriguing properties with respect to conditional independence implications. In particular, it leads to sparse graphical models without the need of a tuning parameter. These results suggest that MTP2 constraints could be an interesting alternative to the standardly used graphical lasso for modeling in the high-dimensional setting, and the group is currently exploring this approach for applications to genomics.

Kalyan Veeramachaneni

Kalyan Veeramachaneni joined LIDS in 2016 as a principal research scientist. His research group develops automation technologies for data science, a burgeoning field that focuses on deriving insights from the huge amount of information produced by contemporary systems. To enable an unsupervised computer program to perform the same tasks as a data scientist, Veeramachaneni and his team spent years working alongside such scientists as they analyzed data in different domains, from education to medicine. They are now attempting to build algorithms that emulate these human thought processes. While this level of automation is necessary to keep up with the ever-increasing demand for data scientists, the group is also pursuing a concurrent effort, the "human-data interaction project." The aim of this project is to allow humans to interact seamlessly with these technologies, thus enhancing what each can achieve individually.

In pursuit of this goal, Veeramachaneni and his student James Max Kanter developed an end-to-end system called the Data Science Machine (DSM). After data scientists specify or formulate a prediction problem, the DSM transforms the raw data into a predictive model by generating variables, selecting a modeling method, and auto-tuning the model's hyperparameters.

Trane: A Formal Language to Describe Prediction Problems

Pushing the boundary even further, Veeramachaneni and his student Benjamin Schreck designed a formal language, called Trane, for describing prediction problems in relational data sets. In addition, they implemented a system that allows data scientists to specify problems in that language.

Trane was inspired by a limitation that experts commonly encounter early in the data science process. Coming up with a prediction problem is an iterative process, one carried out several times even before data scientists begin building a predictive model. This order of events makes it impossible for experts to deduce whether a good predictive model can be generated for a particular problem before they begin working on it. Although a slightly different problem formulation, or the selection of a closely related field, could result in more accurate models, data scientists have no way of knowing which of these improvements are possible. Moreover, when a new domain problem is encountered, the entire process must be repeated.

At the same time, the prediction problem formulations that data scientists generally come up with share many common elements. For example, most involve choosing a time-varying column, selecting a window, and applying a limited set of operations over that window in order to achieve the relevant outcome. The existence of these standards raises the possibility of automatically enumerating many or all of the prediction problems one would want to consider for a given data set and doing so without extensive prior knowledge or data manipulation.

Achieving this aim requires a common language that is descriptive enough to be general purpose but limited enough to be enumerable. This would allow data scientists to automatically formulate hundreds, if not thousands, of prediction problems for a given data set. With the DSM at their disposal, they could also solve these predictive problems automatically and present the results to domain experts and/or their peers.

Trane was designed for this purpose. When tested, the language was able to describe several prediction problems across many different domains, including 29 problems on the data science competition website KAGGLE. Veeramachaneni and Schreck then designed an interpreter that translates input from the user, specified in this language, into a series of transformation and aggregation operations. These operations can subsequently be applied to a data set in order to generate labels that can be used to train a supervised machine learning classifier. Using a smaller subset of the language, Veeramachaneni and Schreck tested this system on KAGGLE's Walmart Store Sales Forecasting data set by enumerating 1,077 prediction problems and then building models that attempted to solve them. Considering that just one out of those 1,077 problems was the focus of a 2.5-month-long competition on KAGGLE, this system is poised to enable a thousand-fold increase in data scientists' productivity.

This level of automation allows for the collapse of the iterative process, enabling data scientists and domain experts to collaboratively select problems to consider while having immediate access to an estimate of each problem's predictive accuracy (and therefore feasibility). It frees up data scientists' time for higher-level and more creative tasks and simultaneously increases the rate of exploration of the prediction problem space by an order of 1,000. It grants domain experts and data scientists the ability to know what types of prediction problems are possible, and whether or not they can achieve reasonable predictive accuracy, so that they can decide which are worth considering further. The system may even generate interesting problems that human actors would not have thought of on their own.

Synthetic Data Vault

An end-to-end data science endeavor requires human intuition as well as the ability to understand data and pose hypotheses and/or variables. To expand the pool of possible ideas, enterprises hire freelance data scientists as consultants and in some cases even crowdsource through the KAGGLE website. In conversations with numerous stakeholders, Veeramachaneni and his team found that the inability to share data due to privacy concerns often prevents enterprises from obtaining outside help. Even within an enterprise, development and testing can be impeded by factors that limit access to data.

Arguably, enterprises can sidestep these concerns and expand their pool of possible participants by generating synthetic data. To be effective, these synthetic data must meet two requirements. First, they must somewhat resemble the original data statistically to ensure realism and keep problems engaging for data scientists. Second, they must also resemble the original data formally and structurally so that any software written on top of them can be reused. In order to meet these requirements, the data must be statistically modeled in their original form so that one can sample from and re-create them. In most cases, that form is the database itself. Thus, modeling must occur before any transformations and aggregations are applied.

To generate synthetic data, Veeramachaneni and his student Neha Patki created the Synthetic Data Vault (SDV), a system that builds generative models of relational databases. The SDV relies on an algorithm that computes statistics at the intersection of related database tables and uses a state-of-the-art copula-based multivariate modeling approach to model these data. The SDV iterates through all possible relations, ultimately creating a model for the entire database. Once this model is computed, the same relational information allows the SDV to synthesize data by sampling from any part of the database.

Tests have shown that experts working with data generated by the SDV produce work whose quality is as high as or higher than that of those working with “real” data; also, they engage well with the data. After building the SDV, the team modeled five different relational data sets and created three versions of synthetic data, both with and without noise, for each of these sets. They hired 39 data scientists and divided them into groups to solve predictive problems defined over the five data sets. They presented different groups with different versions of the data, always giving one group the original data. For each data set, they compared the predictive accuracies of features generated from the original data with the accuracies of those generated by users who were given the synthetic data. Regardless of which group it came from, predictive accuracy for a feature was generated by executing that feature on the original data. For 11 out of 15 comparisons (more than 70%), data scientists using synthetic data performed the same as or better than those using the original data set. Data scientists working on synthetic data generated a total of 4,313 features on the five data sets.

Alan Willsky

Edwin Sibley Webster Professor (retired) Alan Willsky has a long-standing reputation as one of the leaders in the development of multi-resolution statistical models and the exploitation of these models to yield efficient and scalable processing algorithms. Methods he has developed are widely cited and have found application in domains ranging from computer vision to groundwater hydrology.

One very important area of research within this domain is the development of algorithms to learn the structure of multi-resolution models from data, an area in which Professor Willsky and his students and collaborators have made significant contributions over many years. These methods, however, all required a restricting assumption. In some cases, the placement of the data components on a multi-resolution tree with hidden layers is prescribed, as is the structure of the tree, but what is free to be learned are the nature and dimensions of the variables to be placed at the various hidden nodes. In other cases, the structure of the tree and the locations of the observed data elements are completely free, but the dimensionality of the hidden variables is fixed (and typically scalar).

Professor Willsky and former doctoral student James Saunderson are developing a method that removes both of these restrictions simultaneously. While on the surface it seems to be a very ill-posed problem to learn both the structure of the multi-resolution tree and the dimensions of the hidden variables, they have developed conditions under which it can be done and an algorithm to achieve the desired result.

Moe Win

The Wireless Communication and Network Sciences Laboratory, led by Professor Moe Win, is involved in multidisciplinary research that encompasses developing fundamental theories, designing cooperative algorithms, and conducting network experiments for a broad range of real-world problems.

To advocate outreach and diversity, the group is committed to attracting graduate and undergraduate students from underrepresented and minority groups and to giving them exposure to theoretical and experimental research at all levels. The group has a strong track record for hosting students from both the Undergraduate Research Opportunities Program (UROP) and the MIT Summer Research Program. This year the group hosted three women: a graduate student, a postdoctoral fellow, and a recipient of the competitive Ibn Khaldun Fellowship for Saudi Arabian Women. Professor Win maintains dynamic collaborations and partnerships with academia and industry, including the University of Southern California; the University of California, Santa Barbara; Arizona State University; the University of Bologna and the University of Ferrara in Italy; the Vienna University of Technology in Austria; the University of Lund in Sweden; the King Abdullah University of Science and Technology in Saudi Arabia; the Singapore University of Technology and Design and Nanyang Technological University in Singapore; Kyung Hee University in Korea; Draper Laboratory; the Jet Propulsion Laboratory; and the Centre for Maritime Research and Experimentation in Italy.

Professor Win received funding from two new grants that began on June 1, 2017. The first grant (totaling $799,000), for his project "Situational Awareness for Emergencies Through Network-Enabled Technologies (SafeT-Net)," was provided by the National Institute of Standards and Technology. The second (totaling $127,103), for "A System for Efficient and Accurate Network Navigation," was provided by ONR.

Current research topics being investigated by Professor Win and his group include network localization and navigation, network interference exploitation, intrinsic wireless secrecy, adaptive diversity techniques, ultra-wide-bandwidth systems, and quantum information science. Details of a few specific projects are provided below.

Network Localization and Navigation

The group has made notable contributions to the field of network localization and navigation, in particular developing node deployment strategies for localization networks and designing and implementing network localization and navigation algorithms. The group investigated the node deployment problem in a general setting where the positions of the deployed nodes were assumed to be imperfectly known. Specifically, the group considered a navigation network consisting of a target node and a few assisting nodes, where the objective is to minimize the localization error of the target node by deploying additional assisting nodes. An optimization program has been formulated for such a setting and a near-optimal solution, with low implementation complexity and fast execution time, has been obtained. The group is also developing a localization system called Peregrine, which incorporates resource management, network scheduling, and network localization and navigation, including exploitation of spatio-temporal cooperation on commercially affordable hardware. In the meantime, the group is designing a real-time belief propagation algorithm to fuse map information with inter- and intra-user measurements. This algorithm is being implemented (with smartphones only) in a system called Mercury. Experimental results show that the system provides reliable location information and that spatial cooperation remarkably reduces the location uncertainty of users.

Intrinsic Network Secrecy

With the ever-increasing penetration of wireless communications into our lives, securing wireless transmissions has become a demanding necessity. In this context, the group has devoted effort to studying how to apply physical-layer security to strengthen wireless communication secrecy through exploiting the intrinsic properties of the communication channels along with engineering network interference. Based on the theoretical framework developed over the past few years, the group has created a cooperative interference engineering technique—generalized interference alignment—to enhance wireless network secrecy. The feasibility of this technique has been characterized, and insights have been gained on how the technique benefits wireless network secrecy. To determine how network topology affects wireless network secrecy, the group has also analyzed the secrecy performance of stochastic networks with heterogonous density.

To generalize the group's work for practical scenarios, studies on the intrinsic secrecy of heterogeneous networks have been conducted. The insights gained are particularly useful in securing communication within the Internet of Things, vehicular networks, and cyber-physical systems. It has been shown that the performance of local secrecy varies highly with the node distribution and that local densification is a key enabler for increasing local secrecy. Techniques for optimal positioning and optimal routing have been pointed to as viable and interesting solutions.

Highlights, Awards, and Events

The laboratory continues to organize the broadly attended LIDS seminar series and the LIDS student conference, an event organized entirely by LIDS students. Marking its 22nd year in 2017, the conference features distinguished plenary speakers and provides an interactive forum for students to present and discuss their research. LIDS produces a community-oriented magazine, *LIDS|ALL*, which continues to receive great praise. The magazine includes articles on important events related to LIDS and profiles of individuals whose lives have been shaped by LIDS.

LIDS faculty continue to be involved in the organization of major workshops and conferences. Within MIT, we wish to highlight two events: the LIDS Smart Urban Infrastructures Workshop (co-organized by Professors Karaman, Ozdaglar, and Parrilo), which brought together a diverse group of thought leaders, technologists, and theoreticians, and SDSCon 2017 (organized by Professor Shah), a celebration of the new Statistics and Data Science Center and a community-building event for those interested in statistics at MIT.

Finally, LIDS faculty, students, and alumni continue to receive substantial recognition for their contributions, with numerous national and international awards and honors. Some notable examples are listed below.

Awards

Postdoctoral fellow Stefania Bartoletti was named a 2016 Paul Baran Young Scholar by the Marconi Society.

Principal research scientist Audun Botterud, along with Professor Asu Ozdaglar, received a 2017 Siebel Energy Institute research grant for the project "Improved Analytics for Urban Energy Distribution Grids with Smart Buildings."

Noam Buckman was awarded a 2017 National Defense Science and Engineering Graduate Fellowship.

Zied Ben Chaouch received the EECS 2017 Frederick C. Hennie II Teaching Award.

Wenhan Dai, Zhenyu Liu, and Bryan Teague, working with their project advisors (Professor Moe Win, postdoctoral fellow Stefania Bartoletti, and Professor Andrea Conti of the University of Ferrara), won first prize in the 2016 IEEE (Institute of Electrical and Electronics Engineers) Communications Society Communication Technology Changing the World student competition.

Professor Ali Jadbabaie was awarded a 2016 Vannevar Bush Faculty Fellowship from the Office of the Secretary of Defense. This fellowship provides extensive, long-term financial support to distinguished university faculty and staff scientists and engineers. Professor Jadbabaie's project focuses on strategic decision making in large-scale networks.

Professor Patrick Jaillet received the EECS 2017 Ruth and Joel Spira Award for Excellence in Teaching.

System administrator Brian Jones won a 2017 MIT Excellence Award for Serving the Client, one of the highest honors awarded to MIT staff members.

Professor Sertac Karaman received the 2017 Young Investigator Award from the Office of Naval Research. Professor Karaman also received the 2017 IEEE Robotics & Automation Society Early Career Award for "contributions to motion planning and control algorithms for robots and autonomous vehicles."

Christina Lee was awarded a 2016 Claude E. Shannon Research Assistantship. She is supervised by Professors Asu Ozdaglar and Devavrat Shah.

Professor Eytan Modiano, together with his student Abhishek Sinha and Professor Leandros Tassiulas (Yale), received the best paper award at ACM (Association for Computing Machinery) MobiHoc 2016.

Shreya Saxena was named a 2017 Graduate Woman of Excellence by MIT'S Office of the Dean for Graduate Education. She is supervised by Professor Munther Dahleh.

Omer Tanovic won the EECS 2017 Carlton E. Tucker Teaching Award.

Professor John Tsitsiklis received the National HELORS (Hellenic Operational Research Society) Award.

Professor Caroline Uhler received a 2017 NSF Career Award for her project "Gaussian Graphical Models: Theory, Computation, and Applications." Professor Uhler was also awarded a 2017 Sloan Research Fellowship.

Professor Moe Win received the 2016 IEEE Communications Society Edwin Howard Armstrong Achievement Award "for contributions to the foundation of wireless communication and localization networks."

Honors

Professor Robert Berwick delivered the Foundation for Science and Faith's Distinguished Lecture at the Vatican in July 2016. Professor Berwick also delivered the American Society of Catholic Scientists Distinguished Lecture in April 2017.

Professor Munther Dahleh was a plenary speaker at the 2016 European Control Conference and the 2016 Optimization Days conference. Professor Dahleh also spoke at the ETH Global Lectures in June 2017.

Professor Marija Ilic presented a talk on cyber security for energy systems at the 2017 Governors' Summit on Energy Security and Infrastructure in Washington, DC.

Professor Sanjoy Mitter was the August-Wilhelm-Scheer Visiting Professor at the Technical University of Munich in Germany from March to May 2017.

Professor Eytan Modiano was named associate director of LIDS effective July 1, 2017. Professor Modiano was also appointed editor in chief of *IEEE/ACM Transactions on Networking*, the premier journal in the networking field.

Professor Asu Ozdaglar was named associate department head of EECS beginning in January 2017. Professor Ozdaglar also gave invited seminars, tutorials, and plenary talks on incremental/stochastic optimization and incentives and dynamics in networked systems at the annual meeting of the Society for Industrial and Applied Mathematics, the University of Toronto, the Stanford Marketplace Innovation Workshop, and the INFORMS (Institute for Operations Research and the Management Sciences) Applied Probability Society conference.

Professor Devavrat Shah became the inaugural director of the Statistics and Data Science Center effective July 1, 2016.

Professor John Tsitsiklis was appointed director of LIDS and associate director of the Institute for Data, Systems, and Society beginning in April 2017.

Principal research scientist Kalyan Veeramachaneni was named one of the 100 most creative people in business by Fast Company.

Organization

Professor John Tsitsiklis is serving as the new LIDS director, and Professor Eytan Modiano will become the associate director effective July 1. On the occasion of this leadership change, a comprehensive review of all aspects of the laboratory has been initiated, including systematic collection of opinions and input from the entire community of faculty, research and administrative staff, and students. This deliberative process will continue with the aim of enhancing the operational aspects of the lab and promoting cohesive intellectual efforts.

AY2017 Key Statistics

Faculty principal investigators: 22
Research staff (LIDS principal investigators): 6
Administration, technical, and support staff: 10
Postdocs: 23
Visitors: 20
Graduate students: 106

Future Outlook

LIDS continues to be a world-leading center for fundamental research in the information and decision sciences, occupying a unique niche at the interface of theory and applications in diverse areas. Some of our main upcoming efforts will be geared toward solidifying the critical mass in certain key areas and highlighting the particular flavor of LIDS research, its impact, and its brand.

Another of the immediate objectives for LIDS is to further develop collaborative efforts and nurture the balance of theory and practice for maximal impact. The umbrella provided by IDSS and the resulting potential for new cross-cutting collaborations are very helpful in this respect.

Finally, while the laboratory is a research-oriented entity, a constant goal is to have LIDS faculty play a leading role in curriculum innovation, thus bridging research and the classroom in areas such as data science, control and autonomy, and networks.

John Tsitsiklis
Director (as of April 1, 2017)
Clarence J. Lebel Professor, Department of Electrical Engineering and Computer Science

Munther Dahleh
Interim Director (January–March 2017),
Director, Institute for Data, Systems, and Society
William A. Coolidge Professor, Department of Electrical Engineering and Computer Science

Asu Ozdaglar
Director (July–December 2016)
Interim Head and Professor, Department of Electrical Engineering and Computer Science

Pablo Parrilo
Associate Director (July–December 2016)
Professor, Department of Electrical Engineering and Computer Science