Laboratory for Information and Decision Systems

The Laboratory for Information and Decision Systems (LIDS) is an interdepartmental laboratory staffed by faculty, research scientists, and graduate students from several departments and centers across MIT. LIDS provides an intellectually cohesive and collaborative environment that fosters high-quality, forward-looking research and instills in our students the disciplinary depth and interdisciplinary understanding required of research and engineering leaders of today and tomorrow.

Participants and Collaborations

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 (AeroAstro). However, LIDS has long been interdisciplinary, and recent research foci and the pervasiveness of the analytical methodologies advanced by LIDS researchers have broadened our collaborative scope. Some of the many entities at MIT with which LIDS has a strong relationship include the Computer Science and Artificial Intelligence Laboratory (CSAIL), the Research Laboratory of Electronics, the Operations Research Center (ORC), the Department of Brain and Cognitive Sciences, the Department of Civil and Environmental Engineering (CEE), the Department of Mathematics, the Department of Mechanical Engineering, the Department of Economics, and the Sloan School of Management.

LIDS is also a key component of the Institute for Data, Systems, and Society (IDSS), which advances disciplinary methodologies in statistics, data science, and information and decision systems and nurtures cross-cutting connections, especially with the social sciences, to address complex societal challenges. 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), designing new academic programs (such as those in statistics and social and engineering systems), and continuing to be heavily involved in the search for new faculty in the fields of networks and statistics. In particular, LIDS faculty have played a central role in launching the IDSS MicroMasters in Statistics and Data Science program.

LIDS researchers continue to have great success in obtaining funding for our broad and deep research agenda, and we continue to develop relationships with industrial organizations and national laboratories including Draper Laboratory, Lincoln Laboratory, the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL), the MIT-IBM Watson AI Lab, IBM, SES, and Accenture. 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.

LIDS has also been strengthened significantly this academic year with the hiring of a new faculty member, Cathy Wu (CEE). Professor Wu focuses on understanding and shaping the impact of autonomy on society from a dynamical systems perspective. Research directions include sample-efficient reinforcement learning, coping with distribution shifts, bridging machine learning and automation science, and applying automation science in the context of mobility.

Additionally, the lab has recently redefined the role of LIDS affiliate members, fostering stronger ties with a number of MIT faculty who, although not LIDS principal investigators (PIs), are intellectually aligned with the lab's research agenda and involved in substantial interactions. The current affiliate members are Professors Saurabh Amin (CEE), Hamsa Balakrishnan (AeroAstro), Tamara Broderick (EECS), David Gamarnik (Sloan), Song Han (EECS), Stefanie Jegelka (EECS), Youssef Marzouk (AeroAstro), and Elchanan Mossel (Mathematics).

The already strong core of existing LIDS PIs, together with the extended community of affiliate members, has transformed LIDS into a preeminent entity—both within MIT and more broadly in the academic world—in the fields of data science and foundations of machine learning. At the same time, traditional LIDS core areas (communications, information theory, networks, optimization, control, and autonomy) remain active and strong.

Intellectual Vision

The mission of LIDS is to develop and apply rigorous approaches and tools for modeling, analysis, design, and optimization of physical or artificial systems that process information and rely on information for decision making. Research in LIDS encompasses the development of new analytical methodologies as well as the adaptation and application of advanced methods to specific contexts and application domains.

Many of the important recent technological advances involve systems that collect, exchange, and process data and information. Information is then exploited to make decisions, including statistical decisions, resource allocation decisions, and real-time control decisions. This schema captures much of what is happening in an impressive range of fields such as robotics, autonomous systems, intelligent systems, machine learning, life science informatics, computer networks, societal infrastructures, electric power systems, and more. Advances in all of these domains are made in several labs and departments, and domain expertise is typically critical. At the same time, the "information to decisions" viewpoint rests on an intellectual core and on fundamental methodologies that can be applied across disciplines and domains. The objective of LIDS is to serve as a focal point for this intellectual core while advancing work in selected application domains that—in a virtuous cycle—also provide inspiration for further methodological research.

To achieve these aims, LIDS research is underpinned by:

- 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, statistical data processing, data science, and machine learning; transmission of information; networks; and systems and control
- A set of broad challenges in traditional and emerging applications of critical societal importance

The simultaneous efforts along each of these dimensions within the same lab lead to strong synergies: work in the mathematical disciplines leads to new methodologies that advance core engineering disciplines and interdisciplinary applied investigations; conversely, work on new interdisciplinary challenges provides the inspiration and direction for fundamental disciplinary research as well as 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.

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.

Networks: this area includes communications, information theory, and networking, with applications to wireless systems, optical networks, and data centers. Research in this area includes development of fundamental limits on communications systems, design of optimal resource allocation schemes for wireless networks, and design of optimal architectures and control algorithms for data centers and cloud networks. 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.

Moreover, 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 ubiquitous themes such as data fusion, distributed learning and decision making, and issues of scalability, robustness, and performance limits.

Particular areas of significant recent activity are as follows:

- Biological systems and biomedical data analysis
- High-performance unmanned autonomous systems
- Energy systems analysis, economics, and design
- Human-level perception for robotics platforms
- Machine learning for recommendation systems and social media
- Network scheduling and routing
- High-dimensional inference in graphical models
- Networking and information transmission in the context of the Internet of Things
- Social network analysis and characterization
- Network navigation and localization
- Transportation network analysis, control, and design
- Ultra-wideband and other emerging communications technologies

Furthermore, the recognition that research within traditional boundaries in information and decision sciences is not adequate to address many of the emerging societal challenges has motivated LIDS to branch out to areas at the intersection of several disciplines. As a result, LIDS is now 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:

- Foundations of 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
- 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

- Development of scalable and efficient inference algorithms for problems involving "big data," including basic research on graphical models
- Development of causal inference methods for gene regulation and early disease diagnostics

Research Highlights

Detailed information on the research carried out by each LIDS PI is provided below. Some recent highlights include the following.

Professor Guy Bresler and his collaborators have been developing a rich and powerful theory on the (average case) computational difficulty of high-dimensional statistical problems; the theory involves a web of reductions between seemingly disparate problems, in the same spirit as classical worst-case complexity theory for deterministic problems.

Professor Munther Dahleh, in collaboration with LIDS principal research scientist Mardavij Roozbehani and other colleagues, began a new line of research to create a data-driven platform for risk sharing in rural African farming. Using economic models, this work has the potential to boost stakeholders across the farming ecosystem and aims to lift African farmers out of poverty by empowering them with advanced predictive technologies.

Professor Jonathan How and his group developed a framework to perform distributed loop closure detection under constraints on communication bandwidth and computation capacity. The approach, which was presented at the 2018 International Symposium on Experimental Robotics, produces near-optimal performance on real-world data sets and shows promise for use in scenarios such as wilderness search and rescue.

A team led by Professor Ali Jadbabaie was awarded an Army Research Office Multidisciplinary University Research Initiative grant for the project Foundations of Decision Making with Behavioral and Computational Constraints, which proposes a new interdisciplinary research program that works toward a unified theory of human decision making. The six-person team includes LIDS affiliate member Elchanan Mossel.

Professor Luca Carlone and his group developed the first provably correct perception algorithms for localization and mapping with outliers and the first polynomial-time method for object detection with extreme outlier rates. These advances enable robots to identify objects hidden in a three-dimensional (3D) cloud of data quickly and accurately.

In addition, LIDS faculty members play leadership roles in major new initiatives with applications of critical societal importance, including projects 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 [MITEI]), and social, economic, and financial networks in close collaboration with Economics and Sloan.

Faculty Activities

The 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

McAfee Professor of Electrical Engineering Dimitri Bertsekas 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 approximate value and policy iteration methods and their applications. This year he published a new monograph, *Reinforcement Learning and Optimal Control* (Athena Scientific), that aims to clarify the connections between control theory and artificial intelligence (AI) and incorporates Professor Bertsekas's latest research. 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 continuing work may lead to fruitful cross fertilization between the two fields.

Robert C. Berwick

Professor Berwick's research during the past year proved several new results concerning the formal power of "deep learning" systems and human (natural) language, along with new formal characterizations of modern linguistic theory.

First, Berwick's group established the formal power of recurrent neural networks (RNNs), long short-term memory (LSTM), and convolutional neural networks (CNNs) in relationship to long-known results about the formal descriptive power of human language. While all of these networks have been used to model human languages, a long-standing result is that not all human language can be represented via finite-state machines, that is, as if items were lined up like beads on a string. At least the power of context-free grammars, to mirror the hierarchical structure of sentences, is required. However, this past year Berwick's research demonstrated that RNNs, LSTM, and CNNs all have at most the power of finite-state "beads on a string" models, and thus they cannot formally represent all human language hierarchical structures. In order to do so, the weights used by these neural networks all must be able to store arbitrarily precise numbers. This result calls into question the nature of these models as true "neural" representations of human language.

Second, Berwick's research developed a formal logical model, an axiom system, for the current linguistic theory developed at MIT, known as minimalism. With this set of axioms in hand, one can build an automatic inference system for minimalist grammars as follows: given a set of sentences associated with a group of simple "who did what to whom" labels (so-called predicate-argument structures), modern-day theorem provers (in particular, the Z3 system) can "solve" for what is unknown, namely the features associated with the words, outputting a lexicon (a dictionary) that can generate all of the example sentences with their associated correct syntax and predicate-argument structure. This appears to be the first working axiom system that has been converted into a feasible learning system for minimalist linguistic theory.

Audun Botterud

Audun Botterud has been a principal research scientist in LIDS since September 2016. He is also a research affiliate in MIT's Center for Energy and Environmental Policy Research, and he holds a co-appointment at the Argonne National Laboratory. The main goal of Botterud's 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. Towards 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.

Grid Integration of Renewable Energy

The rapid expansion of wind and solar energy resources in the power grid gives rise to challenges in the operation and planning of the grid. In the past year, Botterud and colleagues have been investigating electricity market design questions that arise from more renewables on the grid along with the so-called capacity value of energy storage (i.e., how much capacity can be relied upon during periods of scarcity in the grid). In particular, using mathematical programming and stochastic simulations, Botterud and student Conleigh Byers have analyzed interactions between the capacity values of energy storage and variable renewable energy such as wind and solar. In addition, Botterud and visiting student Ian Scott have proposed ways to speed up the solution of large-scale capacity expansion problems through optimal week selection algorithms. Furthermore, Botterud, Visiting Professor Magnus Korpås (NTNU, Norway), and Professor Hans Auer (TU Wien, Austria) are investigating electricity market designs with increasing shares of renewable energy and energy storage.

Botterud and his group have also analyzed the role of distributed energy resources in the energy system. In particular, they have investigated energy sharing and the pricing of solar photovoltaics and energy storage among local consumers, comparing Stackelberg games and social surplus optimization. Botterud and visiting student Andreas Fleischhacker have expanded their work in this domain to consider investment decisions using cooperative game theory. Also, Botterud, students Scott Burger and Ian Schneider, and Visiting Professor Ignacio Perez-Arriaga have engaged in work on tariff design to provide efficient incentives for distributed energy resources; Botterud and postdoc Mehdi Jafari have completed a project on the potential transition to a low-carbon energy system in Italy; and Botterud and student Ivan Rudnick are looking at the roles of renewable energy and natural gas in India through large-scale simulations of the Indian power system. Finally, Botterud and visiting students Jingwei Yang, Yibao Jiang, and Espen Flo Bødal have started work on optimization across different energy infrastructure systems (electricity, district heating, natural gas, hydrogen) through development of novel formulations and decomposition schemes for computational tractability of complex scheduling and planning problems.

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. A recent direction for this energy storage analytics work is factoring in empirical laboratory test results in model development and parameter tuning for power grid optimization models. In turn, insights from the power grid are used to guide early-stage laboratory work on promising battery chemistries by providing laboratory testing schemes for more realistic grid conditions. These directions are being investigated in an MITEI-funded project (with Professor Fikile Brushett of the Department of Chemical Engineering and postdoc Mehdi Jafari) with a specific focus on flow batteries. Also, Botterud and research affiliate Jiachen Mao are looking at the role of energy storage in distributed and isolated power systems; specifically, they are investigating how different battery technologies might contribute to meeting cost, emissions, and reliability targets considering the impact of battery degradation. In addition, Botterud, Apurba Sakti of MITEI, and postdoc Mehdi Jafari are analyzing the optimal sizing and location of batteries for off-shore wind developments, and Botterud, postdoc Bolun Xu, and Professor Magnus Korpås (NTNU, Norway) are developing improved analytical and computational algorithms to solve the multi-period stochastic battery scheduling problem.

Guy Bresler

Bonnie and Marty Tenenbaum Assistant Professor Guy Bresler 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.

In joint work with student Matthew Brennan, Professor Bresler has continued to investigate the interplay between computational and statistical limits in statistical inference problems. Building on their prior work in average-case complexity of statistical problems, they have produced a pair of new results on sparse principal component analysis (SPCA). First, they have completed the phase diagram for SPCA, characterizing for which parameter values the problem is information-theoretically impossible, feasible both statistically and computationally, or computationally difficult. Second, they have introduced the notion of strong comparison between high-dimensional statistics problems and shown that SPCA is more difficult than planted clique or planted dense subgraph problems in a precise and strong sense. Ongoing research seeks to show the first strong equivalence result between two high-dimensional statistics problems and to address many other problems.

In joint work with Matthew Brennan and Wasim Huleihel, Professor Bresler investigated another aspect of the computational complexity of statistical problems: how sensitive are the observed phenomena to specifics of the problems, and what behaviors are universal? Focusing on submatrix detection, they produced the first universality principle for computational lower bounds obtained through average-case reductions. In the general submatrix detection problem, the task is to detect the presence of a small k by k submatrix with entries sampled from a distribution P in an n by n matrix of samples from Q. This formulation includes a number of well-studied problems, such as biclustering when P and Q are Gaussians and the planted dense subgraph formulation of community detection when the submatrix is a principal minor and P and Q are Bernoulli random variables. These problems all seem to exhibit a universal phenomenon: there is a statistical-computational gap depending on P and Q between the minimum k at which the task can be solved and the minimum k at which it can be solved in polynomial time. The group's main result is tightly characterizing this computational barrier as a tradeoff of k with the Kullback-Leibler divergences between P and Q through average-case reductions from the planted clique conjecture. These computational lower bounds hold given mild assumptions on P and Q arising naturally from classical binary hypothesis testing.

In joint work with students Enric Boix-Adsera and Matthew Brennan, Professor Bresler studied the computational complexity of counting cliques of a given size k in Erdos-Renyi random graphs (where each edge is included independently with some probability p). They gave a worst-case to average-case reduction, showing that counting cliques in a random graph with constant edge probability p is essentially of the same complexity as in a worst-case graph. This is counterintuitive since it *a priori* seems much more difficult to count cliques in an adversarial graph than in a random one. Their work also shows a reduction to counting cliques in graphs with shrinking p, and despite the problem being much easier as p becomes small, the optimal tradeoff in complexity is achieved starting from the single worst-case problem. The computational lower bounds they derived, under a complexity assumption known as the exponential time hypothesis, are tight against what is achievable algorithmically, and indeed they provided new algorithms showing this. There are many interesting directions for further work in this area, such as a better understanding of the complexity of other combinatorial problems in random graphs (for instance, finding cliques of a given size instead of counting them).

Luca Carlone

Luca Carlone started as an assistant professor in the Department of Aeronautics and Astronautics and as a principal investigator in LIDS in 2017 after being a postdoctoral associate and a research scientist in LIDS. The main goal of his research is to enable safe, reliable, and efficient perception and understanding of autonomous systems (e.g., micro aerial vehicles, self-driving vehicles, ground robots) operating in complex and unstructured environments. Machine perception involves the problem of using sensor data and prior knowledge to infer an internal model of the surrounding environment that a machine can use for situational awareness and to support control and decision making.

Robust and efficient perception is crucial to enable safety-critical applications of robotics and autonomous vehicles and provides opportunities to enhance efficiency and effectiveness in service robotics and consumer applications. Toward these aims, Professor Carlone's work involves a combination of rigorous theory and practical algorithms and implementations. In particular, his research interests include nonlinear estimation, numerical and distributed optimization, and probabilistic inference applied to sensing, perception, and decision making in single- and multi-robot systems.

Verifiable and Provably Correct Robot Perception

Perception algorithms have been increasingly used in safety-critical applications, including intelligent transportation and military endeavors wherein algorithmic failures may place human lives at risk. During the last year, Professor Carlone's group focused on the design of perception algorithms that work in the presence of large amounts of outliers (e.g., due to false detections, sensor malfunction, or adversarial attacks) and provide formal performance guarantees. In particular, Carlone developed the first provably correct solution for localization and mapping with outliers and the first polynomial-time method for object detection with extreme outlier rates. These algorithms compute a solution (e.g., an estimate of the location of the robot), an estimated set of inliers, and a certificate regarding the quality of the set of inliers. Provably correct localization, mapping, and object detection are crucial to (1) reduce the amount of human supervision in off-line applications (e.g., 3D reconstruction), (2) reduce failure rates during online operations (e.g., pedestrian misdetection in self-driving cars), and (3) increase trustworthiness in online operation of autonomous vehicles. This research is partially funded by the Army Research Laboratory (ARL), Lincoln Laboratory, the Defense Advanced Research Projects Agency (DARPA), and Google; it has been published in *IEEE Robotics and Automation Letters* and presented at the International Conference on Robotics and Automation, the International Conference on Robots and Intelligent Systems, and the Robotics: Science and Systems conference.

Efficient Task-Driven Perception

Pervasive applications of perception in virtual and augmented reality (AR/VR) and miniaturized aerial platforms require perception algorithms to run in real time on computationally constrained platforms. This issue is exacerbated by the fact that perception has to process high-rate, high-dimensional sensor data (e.g., camera images). This year Professor Carlone's group started developing a framework for task-driven perception that selects the most relevant subset of sensor data to complete a control task under given constraints on performance and computation/power budgets. Moreover, in collaboration with Professor Sertac Karaman (AeroAstro) and Professor Vivienne Sze (EECS), the group has designed the first chip for visual-inertial navigation on miniaturized drones. This work, at the boundary between algorithm and hardware design, has already been featured on the front page of the MIT website. Efficient low-power perception can revolutionize a number of application fields, including nano drones, AR/VR, and self-driving cars (where power consumption is an issue). This research is partially funded by the Office of Naval Research (ONR), ARL, and Lincoln Laboratory.

Metric-Semantic Perception for Spatial Artificial Intelligence

Situational awareness in complex, unstructured, and dynamic environments requires autonomous agents to build and maintain a multifaceted model of the environment, including both a 3D geometric model (useful for navigation and coordination) and semantic labels (useful for characterizing areas of interest and providing more succinct information to human operators). Over the last year, Professor Carlone's group has begun investigating metric-semantic perception with the goal of developing practical algorithms for joint metric-semantic understanding. This research has provided two main contributions: a fast solver for Markov random fields based on semi-definite programming and a decoupled approach for metric-semantic 3D reconstruction that detects planar surfaces in the environment. Joint metric-semantic understanding has the potential to (1) increase robustness of perception by leveraging the redundancy between semantics and geometry, (2) provide opportunities for data compression, and (3) enhance high-level human-robot spatial interactions. This research, partially funded by ONR, ARL, and Lincoln Laboratory, has been published in *IEEE Robotics and Automation Letters* and presented at the International Conference on Robotics and Automation.

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, focuses on machine learning with applications in neural learning and control.

Robustness Scaling in Large Networks

This work examines the dependence of performance measures on network size with a focus on large networks. Professor Dahleh and colleagues developed a framework where it is shown that poor performance can be attributed to dimension-dependent scaling of network energy. Drawing from previous work, they showed that such scaling in undirected networks can be attributed to the proximity of the network spectrum to unity, or distance to instability. However, such a simple characterization does not exist in the case of directed networks. In fact, they showed that any arbitrary scaling can be achieved for a fixed distance to instability. Their results demonstrate that it is always favorable, in terms of performance scaling, to balance a large network. This justifies a popular observation that undirected networks generally perform better. They showed the relationship between network energy and performance measures, such as output shortfall probability or centrality, that are used in economic or financial networks. The strong connection between them explains why a network topology behaves in a qualitatively similar manner under different performance measures. These results suggest that there is a need to study performance degradation in large networks that focus on topological dependence and transcend application-specific analyses.

Learning Noisy Linear Time-Invariant Systems

Professor Dahleh and colleagues addressed the problem of learning the parameters of a stable linear time-invariant (LTI) system with an unknown latent space dimension (or order) from its noisy input-output data. In particular, they focused on learning the parameters of the best lower-order approximation allowed by the finite data. This was achieved by constructing a Hankel-like representation of the underlying system using ordinary least squares. Such a representation circumvents the non-convexities that typically arise in system identification and allows an accurate estimation of the underlying LTI system. Their results relied on a careful analysis of a self-normalized martingale difference term that helps bound identification errors up to logarithmic factors of the lower bound. They provided a data-dependent scheme for order selection and obtained a realization of system parameters corresponding to that order through an approach closely related to the Kalman-Ho subspace algorithm. They showed that this realization is a good approximation of the underlying LTI system with high probability.

Learning Noisy Stochastic Jump Linear Time-Invariant Systems

Professor Dahleh and colleagues addressed the problem of learning the parameters of a mean square stable switched linear system (SLS) with an unknown latent space dimension (or order) from its noisy input-output data. In particular, they focused on learning a good lower-order approximation of the underlying model allowed by finite data. This was achieved by constructing Hankel-like matrices from data and obtaining suitable approximations via singular value decomposition truncation wherein the truncation threshold is purely data dependent. By exploiting tools from a theory of model reduction for SLS, they found that the system parameter estimates were close to a balanced truncated realization of the underlying system with high probability.

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 (ACL). A key aspect of ACL is RAVEN (Real-time indoor Autonomous Vehicle test ENvironment), a unique experimental facility that uses a motion capture system to enable rapid prototyping of aerobatic flight controllers for helicopters and aircraft, robust coordination algorithms for multiple vehicles, and a ground projection system to allow real-time animation of the planning environment, beliefs, uncertainties, vehicle intentions, and predicted behaviors (e.g., trajectories).

Resource-Aware Distributed Loop Closure Detection

Distributed loop closure detection is a fundamental task for popular applications in robotics and computer vision, such as collaborative simultaneous localization and mapping and distributed image retrieval. In real-world scenarios, this process is resource intensive, as it involves exchanging many observations over a wireless network and verifying a large number of potential matches. ACL researchers proposed a principled framework to perform distributed loop closure detection under constraints on communication bandwidth and computation capacity. Agents first exchange compact queries to identify a set of potential loop closures and then select a subset of resource-feasible potential loop closures for verification that maximize a task-driven performance metric. While solving this problem optimally is challenging, ACL researchers proposed efficient approximation algorithms with provable performance guarantees. The proposed approach produces near-optimal performance on real-world data sets and shows promise for practical use in resource-constrained applications such as wilderness search and rescue.

Design of Reliable Graphs for Simultaneous Localization and Mapping

Estimation-over-graphs is a class of estimation problems that admit a natural graphical representation. Several key problems in robotics and sensor networks, including those related to sensor network localization, synchronization over a group, and simultaneous localization and mapping, fall into this category. In this work, ACL researchers were able to characterize the impact of the graphical structure of these problems on estimation reliability. Specifically, several connections were drawn between different notions of graph connectivity and various properties of the underlying estimation problem. These results enable agents to evaluate estimation reliability solely on the basis of the graphical representation of the problem. In addition, ACL researchers established a number of new theoretical results to facilitate the synthesis of sparse but well-connected graphs for estimation problems. These structures were exploited to design a pair of fast approximation algorithms with provable guarantees. The proposed framework was applied to various forms of the measurement selection problem in resource-constrained localization and mapping. The resulting algorithms and theoretical findings were validated using random graphs, existing and new synthetic benchmarks, and publicly available real robotic data sets.

Alignment and Fusion of Representations

A fundamental challenge in many robotics and computer vision applications is correctly aligning and fusing observations of multiple sensors and agents to form a collective model of the environment. To address this problem, ACL researchers developed a framework that enables simultaneous and fast alignment of semantic labels and representations. This development was motivated by two shortcomings of the existing work: high computational complexity, which prohibits real-time applications, and inconsistent alignment solutions, which drastically reduce fusion accuracy. The proposed approach is based on reformulating the problem in a graph-theoretic framework wherein a novel spectral graph clustering technique is used to align the observations. This approach provides accurate and consistent alignment solutions while having a considerably smaller run time than the state-of-the-art methods. The resulting framework can significantly improve the accuracy and efficiency of existing discrete assignment problems, which traditionally use pairwise (but potentially inconsistent) correspondences.

Certifiably Correct Distributed Synchronization

Pose-graph optimization (estimating a robot's trajectory based on relative measurements) admits a semidefinite relaxation that produces exact solutions in low-noise regimes. Solving the relaxed problem, however, is computationally challenging in large problems. The standard approach is to exploit low-rank and geometric structures via fast local Riemannian optimization techniques that can provide globally optimal solutions in low-noise regimes. Inspired by these results, ACL researchers have designed distributed optimization algorithms on Riemannian manifolds for certifiably correct collaborative pose-graph optimization and collaborative simultaneous localization and mapping. Specifically, ACL has developed a fast block-coordinate minimization technique for distributed optimization on the Stiefel manifold and established global convergence rate estimates. Furthermore, ACL has developed a parallel execution framework by exploiting the sparsity structure of typical pose-graph optimization problems. It has been shown that parallelization significantly accelerates convergence by eliminating the dependence of the convergence rate on the size of the pose graph. Ongoing efforts are focused on developing distributed global optimality verification schemes for pose-graph optimization in collaborative localization and mapping.

Marija Ilic

Marija Ilic is a permanent senior staff member at Lincoln Laboratory and a senior research scientist at LIDS. She is also affiliated with MIT's Institute for Data, Systems, and Society. She has continued to proactively work toward building a collaboration among Lincoln Laboratory Group 73 (Energy Systems), LIDS, IDSS, and MITEI. 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. In addition, she has played a key role in the growth of MIT's Electric Power Systems Center (an MITEI low-carbon energy center). Ilic is leading projects in major research areas such as modeling and control for changing electric energy systems, resiliency, and electricity markets. In addition to advising three EECS doctoral students and one EECS master's student, she has served as a co-advisor to two MITEI postdocs. Two master's students graduated under her supervision. She is the PI on one large multi-year MITEI project with industry and a technical lead on several Lincoln Laboratory projects; some of these projects are carried out in collaboration with LIDS. The talks, publications, and provisional patent filings described below evolve around the novel modeling for simulation and control of electric energy systems pursued under her leadership in the Electric Energy Systems Group (EESG@MIT).

Ilic presented many invited talks this year, including:

- "Future of Electric Power Grids" (North Carolina State University)
- "Dynamic Monitoring and Decision System (DyMonDS) Framework for Minimally Coordinated Efficient, Reliable and Resilient Electricity Services" (New York State Energy Research and Development Authority)
- "New Energy Space Modeling and Implications on Complexity of Decision Making and Control in Electric Energy Systems" (plenary talk, DIMACS/TRIPODS/MOPTA Conference, Lehigh University)
- "Future of Control and Communications: Challenges and Opportunities in Integrating Power Electronics in Electric Power Systems" (invited position talk, Future of Power Electronics Conversion conference, Svalbard, Norway)
- "Modeling of Energy/Power Flow Dynamics" (Exergy Analysis & Design for 21st Century Aerospace Systems, Control, & Optimization workshop, Dayton, OH)
- "Dynamic Monitoring and Decision Systems (DyMonDS) Framework for Data-Enabled Integration in Complex Electric Energy Systems" (Workshop on Mathematics for Energy Systems, Isaac Newton Institute, Cambridge, England)

In addition, Ilic produced a number of publications over the past year. For example, along with her group, she published NIST Transactive Energy Modeling and Simulation Challenge Phase II Final Report, which assessed and offered economic and technical solutions for grid operations with high solar power integration. Other co-authored publications included:

- "Recursive Algorithm for Resource Allocation in Radial Network Systems" (2019 Mediterranean Conference on Embedded Computing)
- "Secure Blockchain-Enabled DyMonDS Design" (ACM [Association for Computing Machinery] International Conference on Omni-Layer Intelligent Systems)

- "Fundamental Modeling and Conditions for Realizable and Efficient Energy Systems" (2018 IEEE [Institute of Electrical and Electronics Engineers] Conference on Decision and Control)
- "Enhanced Automatic Generation Control (E-AGC) for Electric Power Systems with Large Intermittent Renewable Energy Sources" (2019 IEEE Power & Energy Society General Meeting)

Ilic and her graduate students made three provisional patent filings, the first concerning stable and efficient operations of electric power grids, the second focusing on methods and systems for secure scheduling and dispatching of synthetic regulation reserves, and the third involving integrative modeling and control of difficult electric aircraft missions.

Ilic participated in the recent Advanced Research Projects Agency-Energy Performance-based Energy Resource Feedback, Optimization, and Risk Management workshop. She also helped co-organize the first Exergy Analysis & Design for 21st Century Aerospace Systems, Control, & Optimization workshop, held at the Wright Brothers Innovation Institute. She has been invited to organize the second workshop in 2020 on the same topic at MIT. She continued to serve as the editor-in-chief of *Foundations and Trends® in Electric Energy Systems* and as an associate editor for several journals. She has collaborated internationally with colleagues in Germany and Japan in response to their interest in the energy modeling approach and the Scalable Electric Power Systems Simulator, which is under development at LIDS. Notably, she served as the technical lead on a study funded by the Department of Homeland Security and RAND that addressed candidate power grid architectures for future Puerto Rico electricity services.

Ali Jadbabaie

Professor Jadbabaie led a recent multi-investigator project sponsored by DARPA on scalable non-convex optimization. The project, which involved four PIs from LIDS and CSAIL, focused on developing fundamental tools, methods, and algorithms for scaling up optimization algorithms for machine learning applications. In particular, work over the past year led to a new understanding of acceleration phenomena, according to which simple modifications of the well-known gradient descent algorithm lead to a speedup. While these modifications have been known to work for decades, the exact mechanism by which they lead to accelerated optimization remained a mystery. In other developments, project researchers developed new tools for understanding graph neural networks, as well as new insights that lead to a better understanding of how machine learning algorithms learn from examples and how the solutions generalize.

In other work Professor Jadbabaie, together with LIDS research scientist Amir Ajorlou and social and engineering systems PhD student Chin-Chia Hsu, developed a theory for the study of the spread of misinformation in social networks. This topic has recently received quite a bit of attention, and there is a pressing need to develop an understanding of empirical observations indicating that fake news travels wider, faster, and deeper than real news in social networks such as Twitter. To study the issue, Jadbabaie and colleagues developed a model in which an individual has to decide whether to share a piece of uncertain/inaccurate news with his or her followers on Twitter. Their findings showed that if the individual's main motive for sharing news is to persuade his or her followers to shift their opinion, the news is shared only if the followers can be convinced to shift. In certain situations, this could lead to a higher probability of news sharing cascades for inaccurate/fake news than for real news.

Patrick Jaillet

The research of Professor Jaillet (Dugald C. Jackson Professor of Electrical Engineering and Computer Science and ORC co-director) focuses on online optimization and learning, machine learning, and decision making under uncertainty. Examples include online and dynamic versions of assignment/matching problems, secretary problems, routing problems, and their generalizations. The 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), reinforcement learning, game-theoretic concepts (price of anarchy), and their integrations.

Motivating applications arise from modern transportation sharing systems and dynamic resource allocation problems in areas relating to 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 eight PhD students from ORC (Maximilien Burq, Phillip Chodrow, Samuel Gilmour, Jason Liang, Sebastien Martin, Konstantina Mellou, Julia Romanski, and Sohil Shah) and one SM student from CEE/EECS (Youssef Aboutaleh). His research group in Singapore included two postdocs (Supriyo Ghosh and Tien Mai from SMART) and six PhD students (Zhongxiang Dai, 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).

Funded research programs over the past academic year originated from ONR (Online Optimization and Learning in a Complex Environment) and SMART (Future Mobility, a large project involving several other MIT principal investigators).

Sertac Karaman

Professor Karaman, the Class of 1948 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 dynamic data-driven application systems.

Alexandre Megretski

Professor Megretski works on problems associated with the analysis and design of nonlinear dynamical systems. Specific areas of interest include optimization of signal processing systems used in digital communications; design and certification of dynamical strategies in multi-agent systems, particularly in the presence of vision-based feedback; and study of fundamental performance limitations for online learning algorithms in dynamical feedback loops.

Sanjoy Mitter

Professor Mitter continued his research on various aspects of information theory as they relate to control in environments where sensors, actuators, and controllers are linked by noisy communication channels.

Professor Mitter's main research has focused on using free energy as a unifying concept in Bayesian inference and equilibrium statistical mechanics. Free energy is defined as the sum of average energy and relative entropy. This leads to the Gibbs variational principle in statistical mechanics and to the characterization of conditional distribution as minimizing free energy in Bayesian inference.

Eytan Modiano

Professor Modiano (associate director of LIDS) 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 communication networks that are cost effective and scalable and meet emerging needs for high-datarate 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 20 journal and conference papers.

Application domains such as autonomous vehicles, command and control systems, virtual reality, and sensor networks rely heavily upon the distribution of time-critical information. Age of Information (AoI) is a recently proposed metric that captures the freshness of information from the perspective of an application. AoI measures the time that elapsed from the moment that the most recently received packet was generated to the present time. Over the past few years, CNRG has been developing network algorithms for optimizing AoI in wireless networks. This past year, graduate student Rajat Talak developed fundamental limits on AoI. In particular, it was shown that, in certain queueing systems, heavy-tailed service time distribution minimizes AoI. This is in sharp contrast to known results in queueing theory showing that heavy-tailed service time distribution maximizes delay. The CNRG results demonstrate that, indeed, AoI is fundamentally different from delay. Moreover, graduate student Vishrant Tripathi considered the problem of information dissemination over a mobile wireless network. In particular, this work considered a mobile agent that disseminates information to nodes on a graph and develops trajectories for the agent that minimizes AoI.

The group continues to work on 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, wherein failures in the cyber network lead to failures in the physical network, and vice versa. Over the past year, Professor Modiano and his student Jianan Zhang have studied the robustness of interdependent communications and computation networks. Such networks are increasingly prevalent, as many network tasks rely on some form of in-network computing (e.g., Internet search, video streaming). In this work, they have been able to characterize the robustness of a distributed computing network in terms of the so-called “minimum cut” of the network (i.e., the minimum number of nodes/links that need to be removed in order to disrupt the computation/communication task).

During the past year, the group started exploring the use of reinforcement learning in network control algorithms. In particular, they have been trying to develop reinforcement learning algorithms that can provide provable performance guarantees. To that end, they have been exploring various model-based reinforcement learning schemes that can be applied to network control and optimization. In particular, they were able to show that the upper-confidence reinforcement learning (UCRL) algorithm can be used in queueing networks with unbounded state spaces to obtain optimal performance. This is in contrast to known results, according to which UCRL can be used only in systems with bounded state spaces. In related work, graduate student Qingkai Liang considered the problem of network control in networks with adversarial dynamics and developed algorithms based on online learning and reinforcement learning to optimize performance.

In addition, with support from the Army Research Office, the group developed a wireless networking testbed using National Instruments programmable radios. The testbed consists of 30 radios with programmable physical and channel access layers that can be used to implement and validate algorithms created by the group.

The group continues to pursue industrial collaborations in order to increase the impact of their work on practical systems. Over the past year, the group collaborated with researchers at Nokia Bell Labs on the problem of resource allocation in distributed computing networks. In addition, the group continues to collaborate closely 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, queueing theory, graph theory, network protocols and algorithms, hardware design, and physical layer communications.

Asuman Ozdaglar

Professor Ozdaglar is the head of the Department of Electrical Engineering and Computer Science. Her research group focuses on modeling, analysis, and optimization of large-scale, dynamic multi-agent networked systems. The group’s 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.

Pablo A. Parrilo

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

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, information-theoretic methods in statistics and machine learning, fault-tolerant circuits, and probabilistic methods in combinatorics.

Communication Complexity of Estimating Correlations

Professor Polyanskiy and colleagues characterized the communication complexity of the following distributed estimation problem: Alice and Bob observe infinitely many independent and identically distributed copies of ϱ-correlated unit-variance (Gaussian or ±1 binary) random variables with unknown $\varrho \in [-1, 1]$. By interactively exchanging k bits, Bob wants to produce an estimate $\hat{\varrho}$ of ϱ. Polyanskiy and his group showed that the best possible performance (optimized over interaction protocol Π and estimator $\hat{\varrho}$) satisfies $\inf_{\Pi, \hat{\varrho}} \sup_{\varrho} \mathrm{E}[|\varrho - \hat{\varrho}|^2] = 1/(k\, 2 \ln 2) + o(1/k)$. Curiously, the number of samples in their achievability scheme was exponential in k; by contrast, a naive scheme exchanging k samples achieved the same $\Omega(1/k)$ rate but with a suboptimal prefactor. Their protocol achieving optimal performance is one way (non-interactive). They also proved the $\Omega(1/k)$ bound even when ϱ is restricted to any small open sub-interval of $[-1, 1]$ (i.e., a local minimax lower bound). Their proof techniques relied on symmetric strong data-processing inequalities and various tensorization techniques from information-theoretic interactive common-randomness extraction. The results also imply an $\Omega(n)$ lower bound on the information complexity of the Gap-Hamming problem, for which they showed a direct information-theoretic proof.

Improved Inequalities, Hypercontractivity, and Uncertainty Principle on the Hypercube

Log-Sobolev inequalities (LSIs) upper bound entropy via a multiple of the Dirichlet form (i.e., norm of a gradient). In recent work, Professor Polyanskiy and colleagues proved a family of entropy-energy inequalities for the binary hypercube that provide a nonlinear comparison between the entropy and the Dirichlet form and improve on the usual LSIs for functions with small support. These nonlinear LSIs, in turn, imply a new version of the hypercontractivity for such functions. As another consequence, they derived a sharp form of the uncertainty principle for the hypercube: a function whose energy is concentrated on a set of small size and whose Fourier energy is concentrated on a small Hamming ball. The tradeoff between the sizes that they derived is asymptotically optimal. This new uncertainty principle implies a new estimate of the size of Fourier coefficients of sparse Boolean functions. The group observed that an analogous (asymptotically optimal) uncertainty principle in the Euclidean space follows from the sharp form of Young's inequality. This hints that nonlinear LSIs augment Young's inequality (which itself is sharp for finite groups).

Estimating Information Flow in Deep Neural Networks

Professor Polyanskiy and colleagues assessed estimations of the mutual information $I(X; T)$ between input X to a deep neural network (DNN) and the output vector T of its ith hidden layer (an "internal representation"). Focusing on feedforward networks with fixed weights and noisy internal representations, they developed a rigorous framework for accurate estimation of $I(X; T)$. By relating $I(X; T)$ to information transmission over additive white Gaussian noise channels, they revealed that compression—that is, reduction in $I(X; T)$ over the course of training—is driven by progressive geometric clustering of representations of samples from the same class. Experimental results verified this connection. They also explored purely deterministic DNNs wherein $I(X; T)$ is provably vacuous and showed that, nevertheless, these models also cluster inputs belonging to the same class. The binning-based approximation of $I(X; T)$ employed in past work to measure compression was identified as a measure of clustering, thus clarifying that these experiments were in fact tracking the same clustering phenomenon. Leveraging the clustering perspective, the group provided new evidence that compression and generalization may not be causally related and discussed potential future research ideas.

Comparison of Channels: Criteria for Domination by a Symmetric Channel

Professor Polyanskiy and his group studied the basic question of whether a given channel V can be dominated (in the precise sense of being more noisy) by a q-ary symmetric channel. The concept of a "less noisy" relationship between channels originated in network information theory (broadcast channels) and is defined in terms of mutual information or Kullback-Leibler divergence. The group provided an equivalent characterization in terms of chi-square divergence. Furthermore, they developed a simple criterion for domination by a q-ary symmetric channel in terms of the minimum entry of the stochastic matrix defining the channel V. The criterion is strengthened for the special case of additive noise channels over finite Abelian groups. The group also showed that domination by a symmetric channel implies (via comparison of Dirichlet forms) a logarithmic Sobolev inequality for the original channel.

List-Decodable Zero-Rate Codes

Professor Polyanskiy and colleagues considered list decoding in the zero-rate regime for two cases: the binary alphabet and the spherical codes in Euclidean space. Specifically, they studied the maximal $\tau \in [0, 1]$ for which there exists an arrangement of M balls of relative Hamming radius τ in the binary hypercube (of arbitrary dimension) with the property that no point of the latter is covered by L or more of them. As $M \to \infty$, the maximal τ decreases to a well-known critical value τ_L. In their work, the group proved several results regarding the rate of this convergence. For the binary case, they showed that the rate is $1/M$ when L is even, thus extending the classical results of Plotkin and Levenshtein for $L = 2$. For $L = 3$, the rate was shown to be $M^{2/3}$. For the similar question about spherical codes, they proved that the rate is between $1/M$ and $1/M^{2L/(L(L-1)+2)}$.

Quasi-Static Fading Multi-Access Communication

Consider a (multiple-access) wireless communication system wherein users are connected to a unique base station over shared-spectrum radio links. Each user has a fixed number of bits k to send to the base station, and users' signals are attenuated by random channel gain (quasi-static fading). In their work in this area, Professor Polyanskiy and colleagues considered many-user asymptotics (wherein the number of users grows linearly with block length) and adopted a per-user probability of error criterion (as opposed to a classical joint-error probability criterion). Under these two settings, they derived bounds on the optimal required energy per bit for reliable multi-access communication (MAC). They confirmed the curious behavior (previously observed for non-fading MAC) of the possibility of perfect multi-user interference cancellation for user densities below a critical threshold. Furthermore, they demonstrated the suboptimality of standard solutions such as orthogonalization and treatment of interference as noise.

Improved Bounds on Gaussian Multi-Access Communication

Professor Polyanskiy and his group considered the Gaussian multiple-access channel with two critical departures from classical asymptotics: the number of users is proportional to block length, and each user sends a fixed number of data bits. They provided improved bounds on the tradeoff between user density and energy per bit. Interestingly, in this information-theoretic problem, they relied on Gordon's lemma from Gaussian process theory. From an engineering standpoint, they discovered a surprising new effect: effective coded-access schemes can achieve perfect multi-user interference cancellation at low user density. In addition, through a similar method, they analyzed the limits of false discovery in binary sparse regression problems in the asymptotic regime of measurements moving to infinity at fixed ratios with problem dimension, sparsity, and noise level. Their rigorous bound matched the formal replica-method prediction for a range of parameters with imperceptible numerical precision.

Optimality of the Plug-In Estimator for Differential Entropy Estimation

Professor Polyanskiy and colleagues established the optimality of the plug-in estimator for the problem of differential entropy estimation under Gaussian convolutions. Specifically, they considered estimations of the differential entropy $h(X + Z)$, where X and Z are independent d-dimensional random variables with $Z \sim N(0, \sigma^2)$. The distribution of X is unknown and belongs to a nonparametric class, but n independently and identically distributed samples from it are available. The group first showed that, despite the regularizing effect of noise, any good estimator (within an additive gap) for this problem must have an exponential in d sample complexity. They then analyzed the absolute-error risk of the plug-in estimator and showed that it converges as $c^d/\sqrt{n}$, thus attaining the parametric estimation rate. This implies the optimality of the plug-in estimator for the considered problem. They provided numerical results comparing the performance of the plug-in estimator with general-purpose (unstructured) differential entropy estimators (based on kernel density estimation [KDE] or k nearest neighbors [kNN] techniques) applied to samples of $X + Z$. These results revealed significant empirical superiority of the plug-in to state-of-the-art KDE- and kNN-based methods.

Alexander Rakhlin

Professor Rakhlin investigates statistical properties of learning procedures that interpolate (memorize) training data. This research is motivated by the good out-of-sample performance of over-parametrized neural networks that achieve zero training error. Rakhlin has turned to the question of whether kernel ridge regression can be a consistent estimator when regularization is turned off. In joint work with Tengyuan Liang, he showed that generalization of such minimum-norm interpolating solutions can indeed arise because of implicit regularization. This new phenomenon occurs due to high dimensionality of the data, curvature of the kernel function, and favorable geometric properties of the data. Furthermore, Professor Rakhlin and Xiyu Zhai showed that the minimum-norm interpolant is inconsistent if data dimensionality is small.

In another line of work, Tuhin Sarkar and Professor Rakhlin provided a detailed finite-sample analysis for learning linear time-invariant systems from a single trajectory. This research was extended to partially observed linear time-invariant systems in joint work with Professor Dahleh.

Mardavij Roozbehani

Principal research scientist Roozbehani led several research efforts related to mathematical modeling, optimization, and control for cyber-physical networked systems. Application areas of focus included energy networks, transportation networks, financial networks, and digital or precision farming. The main themes in Roozbehani's research are concentrated around models and methodologies for quantifying robustness, efficiency, and risk in energy networks and understanding the tradeoffs that system architectures induce between these metrics. For instance, in energy networks with price-sensitive demand and renewable (uncertain) supply, research conducted by Roozbehani and colleagues revealed important tradeoffs between efficiency and risk and how they depend on the extent of cooperation versus competition among participating agents. Market mechanisms that a system operator can adopt to influence behavior and strike optimal tradeoffs were obtained as by-products of the tradeoff characterization. Development of abstract analytical models at both the agent and network levels for modeling the dynamic interaction between agents and markets, as well as strategic interactions among agents, has been instrumental in producing the above-mentioned findings and an important thrust of the research in general. These abstractions guided the group's research in cyber-physical systems and resulted in deeper insights on and a better understanding of the robustness and fragility of general flow networks.

Capacity-Constrained Real-Time Scheduling Algorithms

In this research, Roozbehani and colleagues studied the problem of scheduling sporadic preemptible tasks. These types of problems generally admit a large degree of freedom in determining how to allocate resources to tasks while completing each of them by their deadlines. The group developed a computationally efficient and compact characterization of all of the scheduling decisions at a given time that maintains the feasibility of the overall scheduling problem. This allowed optimizing the scheduling decision based on exogenous optimality criteria that become available only in real time. Potential applications include electric vehicle charging, cloud computing, and web advertisements.

Digital Farming and Quantifying the Value of Data

The long-term goal of this new line of research is to empower farmers, particularly in rural Africa and other developing areas, to improve their production through investments in technologies. Traditional investment strategies have the farmers take loans to upgrade their operations by mortgaging their land. As a result, in such settings, investors participate in the profits but not in the losses that result from external shocks (e.g., adverse weather and pests). This makes the farmers vulnerable to shocks, which serve as a deterrent from engaging in these upgrades. The idea is to establish mechanisms to leverage the potential increases (through technology) in productivity and profits to provide insurance for shocks. Roozbehani and colleagues' approach is based on a data-sharing platform that incentivizes investors and rewards farmers for participation. The value of data in this setting is related to the increase in aggregate welfare after adoption of technologies. Two economic models need to be contrasted to quantify the value of data: an economy without instrumentation and data collection and one with data collection. The idea is that the researchers can conceivably test all possible controls (interventions) on all different farms. Assuming that the farms have some similarity, one can leverage tests from other farms. This leads to a matrix completion problem in the following way: rows indicate M farms and columns indicate N treatments. Let the (i, j) cell indicate a success/failure metric of applying treatment j to farm i. This provides a sparse matrix, and the idea is to use matrix completion to determine success for untreated farms. Then the value of the treatment (and hence the data) is computed by looking at the best outcomes for each farm.

Optimal Routing in Web-Based Private Transportation Networks

Thanks to the growth of web-based transportation companies and the increasing availability of data, there is an opportunity for substantially improving the efficiency of transportation networks. Existing vehicle routing solutions that are deployed by ride providers are far from efficient. For instance, drivers make ad hoc or random decisions about their travel paths between rides. Using graph network models and principles of dynamic programming, Roozbehani and colleagues addressed this problem by developing an efficient algorithm to determine the optimal policy for drivers between rides in order to maximize driver profits. Their approach accounts for various factors, including the frequency of passenger ride requests at different locations, traffic conditions, and surge pricing. The approach was implemented, tested, and proven effective based on road network data from Boston and New York City.

Devavrat Shah

Professor Devavrat Shah and his research group are currently involved in developing theoretical foundations and algorithmic solutions for questions arising in the context of "social" data processing and decision making. Social data, the data generated through modern services including e-commerce portals, media, polls, and more, provides us with a granular lens on the inner workings of our society. On the positive side, they can help in the design of better policies, improve social living by enabling efficient labor markets, and help businesses operate efficiently. On the negative side, they can help manipulate social opinions or even election outcomes; this motivates research on methods for identifying the origins and mechanisms of such manipulation and possible preventive or punitive measures.

Intellectually, this requires development of robust statistical models that capture the universal aspects of social behavior, the ability to engage in causal inference with extremely limited information in the presence of a large number of potential causes, and algorithms that can scale with the amount of data while extracting meaningful information in high-dimensional settings. Specifically, addressing the challenges just described requires fundamental progress at the interface of statistics, machine learning, computation, and social sciences. Professor Shah's group has been working to address these challenges through their focus on canonical issues in social data processing, including learning choice, recommendations, ranking, crowdsourcing, and causal inference.

When evaluating the impact of a policy on a metric of interest, it may not be possible to conduct a randomized control trial. Over the past few years, Professor Shah and his group have introduced a robust variant of the synthetic control method, a popular data-driven approach that estimates a synthetic control by combining measurements of “similar” units (donors). This method, called robust synthetic control (RSC), is a generalization of synthetic control that overcomes the challenges of missing data in high levels of noise while removing the reliance on domain knowledge for selecting donors. However, synthetic control, RSC, and their variants suffer from poor estimation when the pre-intervention period is too short.

Professor Shah’s group and a collaborator at Columbia University have now developed a novel method, mRSC, as a generalization of unidimensional RSC to multidimensional RSC. This proposed mechanism incorporates multiple metrics to estimate a synthetic control, thus overcoming the challenge of poor inference from limited pre-intervention data. According to their approach, data are in effect “multiplied” because multiple metrics can be combined. This data amplification procedure allows for accurate synthetic control identification even in the presence of extremely limited data.

A key requirement of any sound scientific method is the ability to refute; thus, the group developed a diagnostic test that evaluates the utility of including additional metrics. As a byproduct, their method provides a way to forecast time series using related and auxiliary data. This has led to novel methods for sports analytics including causal inference–based rules for the game of cricket.

Suvrit Sra

Suvrit Sra joined the Department of Electrical Engineering and Computer Science as an assistant professor in 2018 and was promoted to associate professor in 2019; he continues as a principal investigator at LIDS, where he has been since January 2015. His research spans machine learning and optimization and pure and applied mathematics and statistics.

Professor Sra’s primary research is in large-scale optimization for machine learning. Over the last year, with his collaborators and students, he worked on developing deeper understandings of the theory of non-convex optimization as well as the theory of deep neural networks. In particular, his students Chengtao Li, Hongyi Zhang, and Zelda Mariet brought a line of research on optimization and probability theory to fruition while opening paths to new explorations. Together with his students, Professor Sra’s work on the theory of deep learning continues to expand; one of the group’s noteworthy findings was that typical neural networks can have infinitely many local minima, and an absence of “bad” local minima is not a property that explains the empirical success of neural networks. In optimization, a key result from Professor Sra and his group was that the version of stochastic gradient descent (SGD) that is actually implemented on computers can be theoretically analyzed and shown to be provably faster than the version of SGD that is almost invariably analyzed in the hundreds of published works on this important algorithm. Their work opens the path to bridging this notable gap between theory and practice. Finally, in another noteworthy result, the group obtained methods provably faster than gradient descent by discretizing ordinary differential equations (ODEs) using standard Runge-Kutta integrators; this work is the first to provide a non-reverse-engineered method faster than gradient descent via the ODE view.

John Tsitsiklis

Professor Tsitsiklis (LIDS director, IDSS associate director, and Clarence J. LeBel Professor of Electrical Engineering) has been primarily focused on LIDS administrative aspects and an online class offered through the IDSS MicroMasters in Statistics and Data Science program.

On the research side, he 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 are typically 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.

Replication Strategies in Data Centers and Other Multiserver Systems

Consider a processing system, consisting of several servers, meant to process a random stream of incoming tasks. A popular strategy that can reduce average delay—both in theory and in practice—involves replicating incoming jobs and sending different copies of the task to different servers. Existing analyses of such schemes point to performance improvements and provide some insights but under very restrictive and unrealistic assumptions. Together with student Martin Zubeldia, who is co-supervised by Professor David Gamarnik (MIT Sloan), Professor Tsitsiklis has continued the development of a complex model that describes in a more realistic manner the random slowdowns in different servers and the correlation of such slowdowns with job size. This work is close to completion, and their focus is now on some technical aspects of their conjectures and proofs.

Sensitivity Theory for a Class of Hybrid Dynamical Systems

There is a popular (and much used in practice) 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, Professor Tsitsiklis 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 to settle some problems that were left open in an earlier LIDS doctoral thesis. During the current year, some previously submitted manuscripts were refined, revised, resubmitted, and eventually accepted. In addition, some general structural results were developed on the relation between different qualitative properties of an important class of hybrid dynamical systems.

Privacy

Professor Tsitsiklis and Kuang Xu have initiated research on a general formulation of the problem of privacy-preserving decision making. This work is at an initial stage. The goal is to express different notions of privacy (namely differential privacy and privacy of an agent who is striving for a specific goal) within a unified framework and then develop a general theory.

Caroline Uhler

Henry L. and Grace Doherty Associate Professor Caroline Uhler carries out research in the areas of machine learning and statistics with applications to genomics; her particular focus is on algebraic statistics, graphical models, and causal inference.

Causal Inference and Gene Regulation

Determining the causal structure of a set of variables is critical for both scientific inquiry and decision making. However, this is often challenging in practice due to limited interventional data. Given that randomized experiments are usually expensive to perform, Professor Uhler, in collaboration with her PhD students Raj Agrawal and Karren Dai Yang, MEng student Chandler Squires, and Karthik Shanmugam from IBM, proposed a general framework and theory based on optimal Bayesian experimental design to select experiments for targeted causal discovery. Specifically, they assumed that the experimenter is interested in learning some function of the unknown graph (e.g., all descendants of a target node) subject to design constraints such as limits on the number of samples and rounds of experimentation. While it is in general computationally intractable to select an optimal experimental design strategy, the group provided a tractable implementation with provable guarantees on both approximation and optimization quality based on submodularity. They evaluated the efficacy of their method on both synthetic and real data sets and demonstrated that it realizes considerable performance gains over baseline strategies such as random sampling.

From observational and interventional data, a directed acyclic graph (DAG) model can be determined only up to its interventional Markov equivalence class (I-MEC). Professor Uhler and colleagues investigated the size of MECs for random DAG models generated by uniformly sampling and ordering an Erdős-Rényi graph. For constant density, they showed that the expected log observational MEC size asymptotically (in the number of vertices) approaches a constant. Similarly, they characterized I-MEC size with high precision. In addition, they showed that the asymptotic expected number of interventions required to fully identify a DAG is a constant. These results were obtained by exploiting Meek rules and coupling arguments to provide sharp upper and lower bounds on the asymptotic quantities, which were then calculated numerically up to high precision. Their results have important consequences for the experimental design of interventions and the development of algorithms for causal inference.

Also, in collaboration with PhD students Yuhao Wang and Anastasiya Belyaeva and MEng student Chandler Squires, Professor Uhler considered the problem of estimating the differences between two causal DAG models with a shared topological order given independent and identically distributed samples from each model. This is of interest for example in genomics, where changes in the structure or edge weights of the underlying causal graphs reflect alterations in gene regulatory networks. They provided the first provably consistent method for directly estimating the differences in a pair of causal DAGs without separately learning two possibly large and dense DAG models and computing their difference. Their two-step algorithm first uses invariance tests between regression coefficients of the two data sets to estimate the skeleton of the difference graph and then orients some of the edges using invariance tests between regression residual variances. They demonstrated the properties of their method through a simulation study and applied it to the analysis of gene expression data from ovarian cancer and during T-cell activation.

Total Positivity

In a recent collaboration, Professor Uhler and colleagues 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. MTP2 was introduced because it is easier to verify than positive association. 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). Professor Uhler showed, in collaboration with Steffen Lauritzen and Piotr Zwiernik, that the maximum likelihood estimator (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.

Extending their results on Gaussian MTP2 models, they recently showed that any MTP2 exponential family is convex, and hence maximum likelihood estimations for this model family can be solved via a convex optimization problem. Moreover, they showed that MTP2 quadratic exponential families, which contain ferromagnetic Ising models and attractive Gaussian graphical models as special cases, are defined by intersecting the space of canonical parameters with a polyhedral cone whose faces correspond to conditional independence relations. Thus, MTP2 serves as an implicit regularizer for quadratic exponential families and leads to sparsity in the estimated graphical model. Furthermore, they showed that the MLE in an MTP2 binary exponential family exists if and only if the sign patterns $(1, -1)$ and $(-1, 1)$ are represented in the sample for every pair of vertices; in particular, this implies that the MLE may exist with $n = d$ samples, in stark contrast to unrestricted binary exponential families where 2^d samples are required. Finally, they developed a globally convergent algorithm for computing the MLE for MTP2 Ising models similar to iterative proportional scaling and applied it to an analysis of data related to two psychological disorders.

These results suggest that MTP2 constraints could be an interesting alternative to the standardly used graphical lasso for modeling in high-dimensional settings. It is therefore of interest to obtain sparsistency and non-parametric density estimation results under MTP2. Professor Uhler is currently working on such results together with her PhD student Yuhao Wang. In addition, they are exploring MTP2 constraints for applications to genomics, particularly identifying groups of genes that are spatially clustered and co-regulated.

Copula Models

Multivariate distributions are fundamental to modeling. Discrete copulas can be used to construct diverse multivariate joint distributions over random variables from estimated univariate marginals. The space of discrete copulas admits a representation as a convex polytope that can be exploited in entropy-copula methods relevant to hydrology and climatology. To allow for extensive use of such methods in a wide range of applied fields, it is important to have a geometric representation of discrete copulas with desirable stochastic properties. In collaboration with her former postdocs Elisa Perrone and Liam Solus, Professor Uhler showed that the families of ultramodular discrete copulas and their generalization to convex discrete quasi-copulas admit representations as polytopes. They drew connections to the prominent Birkhoff polytope, the alternating sign matrix polytope, and their most extensive generalizations in the discrete geometry literature. In doing so, they generalized some well-known results on these polytopes from both the statistics literature and the discrete geometry literature.

Kalyan Veeramachaneni

Kalyan Veeramachaneni joined LIDS in 2016 as a principal research scientist. His research group, the Data to AI Lab, develops automation technologies for data science, deriving insights from the huge amount of information produced by contemporary systems and building algorithms that emulate human thought.

Today, the demand for data scientists is growing at an unprecedented rate. To address that need, the research conducted by Veeramachaneni and his team focuses on creating algorithms that emulate thought processes. By studying with and working alongside data scientists as they analyze information in a variety of fields ranging from education to medicine, the Data to AI Lab works toward the broad goal of designing unsupervised computer programs that can perform the same tasks as human data scientists. Concurrent with this research, the group is pursuing the Human-Data Interaction Project, a research venture designed to enable seamless human interaction with these complex and still-developing technologies and to extend the capabilities of human researchers.

ATMSeer: Cracking Open the Black Box of Automated Machine Learning

As the demand for machine learning experts continues to grow, researchers have begun to construct tools that automate the process of creating machine learning models in order to allow domain experts to perform this normally specialized task themselves. However, practitioners often struggle to use these tools due to a lack of transparency and their inability to control the search process. Veeramachaneni and his team built an interactive visualization system that allows practitioners to visualize the search process, control its direction, and analyze the results. This system, called ATMSeer, is built on top of ATM (auto tune models), a search engine previously created by Veeramachaneni and his team.

When the team conducted user studies to understand how practitioners actually interact with ATMSeer, they found that the new system helped with several decisions, including whether or not to continue the search process, how to make a selection among equally competent models, and how to fine-tune the search space to allow for a more specifically directed search. Veeramachaneni and his team have open sourced their search engine as well as their new tool. They envision that this will lead to greatly expanded use of machine learning in multiple domains.

Improving Synthetic Data Generation

Another emerging area of machine learning involves synthetic data generation: creating a statistically useful data set by sampling new instances from a joint distribution learned on real data. Synthetic data have multiple uses; for example, they are useful for testing the scalability and robustness of new algorithms and can be safely shared openly, preserving the privacy and confidentiality of the actual data.

Generating synthetic data that are statistically similar to real data is difficult. In other work, Veeramachaneni, his student Yi Sun, and Professor Alfredo Cuesta-Infante from Universidad Rey Juan Carlos in Spain addressed this issue by employing a vine copula model. A vine model is a flexible high-dimensional dependence model that uses only bivariate building blocks. However, the number of possible configurations of a vine copula grows exponentially as the number of variables increases, making model selection a major challenge in development. In their work, the team formulated a vine structure learning problem with both vector and reinforcement learning representations. They used neural networks to find the embeddings for the best possible vine model and generate a structure. The team's model showed significant improvement over other generation methods on multiple benchmarking as well as real data sets.

Alan Willsky

Edwin Sibley Webster Professor (retired) Alan Willsky and his former doctoral student, James Saunderson, have continued their research into learning multiresolution/hierarchical graphical models for complex and large data sets with far fewer restrictions than previous methods. In particular, they are working on methods that provably learn correct models with multiple layers of hidden variables. A paper on this topic is in progress. In January, Professor Willsky gave an invited talk at KTH Stockholm at which he provided an overview of his work in this area over several decades.

Moe Win

Professor Moe Win leads the Wireless Information and Network Sciences Laboratory, a group involved in multidisciplinary research that encompasses developing fundamental theories, designing network-enabled 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. Professor Win has a strong track record of hosting students from both the Undergraduate Research Opportunities Program (UROP) and the MIT Summer Research Program. He 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 and the Graz University of Technology in Austria; the University of Lund in Sweden; the University of Oulu in Finland; 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; Tsinghua University in China; Draper Laboratory; the Jet Propulsion Laboratory; Mitsubishi Electric Research Laboratories; the Centre for Maritime Research and Experimentation in Italy; and the Defence Science and Technology Group in Australia.

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

Network Localization and Navigation

The group has made notable contributions to the field of network localization and navigation, in particular joint localization and synchronization and multi-object tracking (MOT). Of note, the group developed network localization and synchronization (NLS), a paradigm that jointly considers localization and synchronization to achieve potential performance gains in a completely asynchronous wireless network. Specifically, the group considers a network consisting of agents and anchors wherein the positions of agents are unknown. The network is completely asynchronous, such that the local clocks of both agents and anchors are affected by unknown skews and offsets. The group established a theoretical framework to investigate NLS performance. In particular, they formulated a computing problem for NLS under a non-Bayesian framework and investigated the fundamental performance limits of the problem through an equivalent Fisher information analysis in which they considered the waveforms received by the nodes as measurements. Also, they derived upper and lower bounds on the performance limits of individual nodes to obtain insights into the effects of network connectivity and scheduling schemes on performance limits. Analytical results were evaluated by means of simulations based on practical models that took the spatial consistency of wireless channels into account.

In addition, the group has continued the development of a paradigm for scalable MOT that is based on factor graphs and the loopy sum-product algorithm (SPA). In particular, they have proposed Bayesian inference algorithms for localizing and tracking extended objects in the presence of data association uncertainty. These algorithms are based on factor graphs and SPA. The group has achieved a reduction in computing complexity in a principled manner by means of "stretching" factors in the graph. After stretching, new nodes have lower dimensions than the original nodes. This leads to reduced computing complexity of the resulting SPA. One of the introduced algorithms is based on an overcomplete description of data association uncertainty and has a computing complexity that scales only quadratically in number of objects and linearly in number of measurements. Without relying on suboptimal pre-processing steps, such as clustering of measurements, the designed algorithm can localize and track multiple objects that potentially generate a large number of measurements. Simulation results demonstrate the excellent accuracy-complexity tradeoff of the proposed algorithms. The group is also continuing its work on MOT by performing experiments using emerging millimeter wave radar technology.

Quantum Information Science

Quantum entanglement serves as a valuable resource for many important quantum operations. In this context, the group has devoted its effort to establishing entangled qubit pairs between agents that are far apart. To address the noise inherent in the system, recurrent quantum entanglement distillation (QED) algorithms have been proposed. However, the recursive structure requires storing a multitude of qubits, and quantum memory size is a critical limitation of near-term quantum devices. The group has characterized the memory size used in recurrent QED algorithms and reduced that size based on the insights gained from the characterization. In particular, they have developed a scheduling policy to achieve optimal memory efficiency. Since minimum memory size grows only linearly with respect to rounds of distillation, the work of the group removes the quantum memory size issue as a bottleneck for implementing QED.

Events and Communications

LIDS continues to organize its signature events: the broadly attended LIDS seminar series and the LIDS student conference, a student-run event that provides an interactive forum for students to discuss their research and features several distinguished plenary speakers each year. The conference celebrated its 24th year in spring 2019.

LIDS faculty also continue to be involved in the organization of major workshops and conferences. Two notable examples are as follows:

- The Learning for Dynamics and Control workshop, co-organized by Professors Ali Jadbabaie and Pablo Parrilo, which brought together experts in machine learning and AI along with researchers in control theory and robotics. The event welcomed nearly 400 attendees.
- The LIDS-hosted Control@MIT event, designed as a forum for MIT's control community to share current research, discuss challenges in the field, and identify opportunities for collaboration and future research. The event, organized by Professor Jonathan How, attracted 100-plus attendees including faculty, PIs, students, and postdocs from across the Institute.

In addition, LIDS faculty have been key contributors to the organization of various events under the umbrella of IDSS and the Statistics and Data Science Center, which sits within IDSS. These events include the annual Women in Data Science Conference (jointly presented by Harvard and MIT and co-organized by Professor Caroline Uhler); the annual Statistics and Data Science Day, a community-building event for those interested in statistics at MIT; and the weekly LIDS and Stats Teas, which provide students and postdocs an opportunity to give brief research presentations to the community in an informal setting.

The lab also continues to produce its annual community-oriented magazine, *LIDS|ALL*, which consistently receives great praise. The magazine includes articles on important events related to LIDS and profiles of individuals whose lives have been shaped by the lab.

Finally, we have made significant forays into reaching a wider audience for our work. This has included publishing substantial numbers of MIT News articles; coordinating news coverage with departments, labs, and centers throughout the Institute; and re-energizing our social media presence. As one measure of success, the LIDS Facebook page saw a 38% increase in followers this academic year.

Awards

Professors Dimitri Bertsekas and John Tsitsiklis were awarded the INFORMS (Institute for Operations Research and the Management Sciences) John von Neumann Theory Prize for their research monographs *Parallel and Distributed Computation* and *Neuro-Dynamic Programming*.

Principal research scientist Audun Botterud's paper "Combined Heat and Power Dispatch Using Simplified District Heat Flow Model" was selected for presentation in the Best Conference Papers on Power System Operations and Electricity Markets session at the 2019 IEEE Power and Energy Society General Meeting.

Professors Luca Carlone and Sertac Karaman, together with their collaborators (Amr Suleiman and Zhengdong Zhang), received the Best Student Paper Award at the 2019 Symposium on VLSI Circuits.

Alireza Fallah and Ian Schneider were named to the 2019 cohort of Siebel Scholars in recognition of their academic achievements, leadership, and commitment to addressing crucial global challenges.

Professor Jonathan How and his co-authors (Shayegan Omidshafiei, Dong-Ki Kim, Miao Liu, Gerald Tesauro, Matthew Riemer, Christopher Amato, and Murray Campbell) received an Outstanding Student Paper honorable mention at the 2019 AAAI (Association for the Advancement of Artificial Intelligence) Conference on Artificial Intelligence. Professor How was also a finalist for the 2018 International Conference on Robotics and Automation Best Multi-Robot Systems Paper Award.

Professor Ali Jadbabaie received a 2019 Multidisciplinary University Research Initiative Award from the Directorate of Basic Research at the Department of Defense. In addition, Professor Jadbabaie and Professor Asu Ozdaglar were co-authors of a paper that received a Best Student Paper Award at the IEEE 2018 International Conference on Acoustics, Speech and Signal Processing.

Professor Patrick Jaillet and collaborators were finalists in the Production and Operations Management Society's POMS-JD.com 2019 Best Data-Driven Research Paper Competition. Professor Jaillet, together with colleagues, also won the Best Applications Paper Award at the 2019 International Conference on Automated Planning and Scheduling.

Professor Thomas Magnanti received a Singapore National Day Award for his long-term work developing higher education in Singapore.

Professor Alexander Rakhlin received the Joseph A. Martore Award for Excellence in Teaching from MIT's Institute for Data, Systems, and Society.

Professor Devavrat Shah, along with LIDS alumni Shreevatsa Rajagopalan and Jinwoo Shin, received the 2019 ACM SIGMETRICS (Special Interest Group for the Computer Systems Performance Evaluation Community) Test of Time Paper Award; the award recognizes a SIGMETRICS paper that is still influential 10 to 12 years after its initial publication.

Professor Suvrit Sra received a National Science Foundation (NSF) CAREER Award as well as an NSF TRIPODS+X grant. He was also awarded an NSF BIGDATA grant for which he is the lead PI.

Omer Tanovic won the MIT School of Engineering Graduate Student Extraordinary Teaching and Mentoring Award.

Professor John Tsitsiklis received an honorary doctorate from Harokopio University in Athens, Greece.

Professor Caroline Uhler was named a 2019 Simons Investigator in Mathematics and Physical Sciences by the Simons Foundation.

Professor Alan Willsky received the 2019 IEEE Jack S. Kilby Signal Processing Medal for contributions to stochastic modeling, multi-resolution techniques, and control-signal processing synergies.

Professor Moe Win and his group received a 2018 R&D 100 Award for Peregrine, a network localization and navigation prototype.

Honors

Professor Guy Bresler was promoted to associate professor without tenure in the Department of Electrical Engineering and Computer Science effective July 1, 2019.

Professor Luca Carlone was elevated to an IEEE senior member. Also, Professor Carlone was named an AIAA senior member.

Professor Munther Dahleh was a plenary speaker at multiple workshops and other venues, including a presentation to the Vienna chamber of commerce.

Professor Sertac Karman was granted tenure by the Department of Aeronautics and Astronautics.

Professor Asu Ozdaglar was named the School of Engineering Distinguished Professor of Engineering.

Professor Pablo Parrilo was named the Joseph F. and Nancy P. Keithley Professor by the Department of Electrical Engineering and Computer Science.

Professor Suvrit Sra was promoted to associate professor without tenure in the Department of Electrical Engineering and Computer Science effective July 1, 2019.

Organization

Professor John Tsitsiklis continues to serve as LIDS director (having started in this role in April 2017). Professor Eytan Modiano joined as associate director effective July 1, 2017. On the occasion of this leadership change, a comprehensive review of all aspects of the lab was initiated, including systematic collection of opinions and input from the entire community of faculty, research and administrative staff, and students. This process led to a number of new initiatives aimed at improving intellectual and social life within the lab; examples include mentoring events, student-led events, and efforts to connect with the LIDS alumni community. These efforts will be furthered with a major celebration of the lab's 80th anniversary in fall 2019.

AY2019 Key Statistics

- Faculty PIs: 25
- Research staff PIs: 4
- Affiliate members: 8
- Administration, technical, and support staff: 10
- Postdocs and other research staff (non-PIs): 31
- Visitors and other affiliates: 40
- Graduate students: 122
- Visiting students: 17

Overall Outlook

LIDS is a world-leading center for fundamental research in the information and decision sciences. It occupies a unique niche at the interface of theory and applications in diverse areas and is a central component underlying many recent technological advances and challenges, including in the currently vibrant fields of data science, statistics, machine learning, and intelligent systems. There are of course many activities in these domains taking place outside LIDS, including prominent applications. Within this broad range of activities, LIDS serves as a focal point in the development of the underlying fundamental methodologies and as a meeting ground for like-minded researchers.

In addition to fundamental research, LIDS is engaged in furthering collaborative efforts that balance theory and practice for maximal impact. The umbrella provided by IDSS and the resulting opportunities for cross-cutting collaborations are very helpful in this respect.

Finally, while LIDS is a research-oriented entity, the lab's faculty maintain 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
Clarence J. Lebel Professor, Department of Electrical Engineering and Computer Science

Eytan Modiano
Associate Director
Professor, Department of Aeronautics and Astronautics