Frontiers of AI/ML: MIT's Quest for Intelligence

July 14, 2020 - July 16, 2020

Part 1: Tuesday, July 14, 2020

11:00 am - 11:40am

Challenges and Opportunities for Robotics and AI in the Face of the Pandemic Daniela Rus Director, MIT Computer Science and Artificial Intelligence Laboratory (CSAIL)



Daniela Rus Director MIT Computer Science and Artificial Intelligence Laboratory (CSAIL)

Daniela Rus is the Andrew (1956) and Erna Viterbi Professor of Electrical Engineering and Computer Science, director of MIT's Computer Science and Artificial Intelligence Laboratory. She brings deep expertise in robotics, artificial intelligence, data science and computation. She is a member of the National Academy of Engineering and the American Academy of Arts and Sciences, and a fellow of the Association for the Advancement of Artificial Intelligence, the Institute of Electrical and Electronics Engineer, and the Association for Computing Machinery. She is also a recipient of a MacArthur Fellowship, a National Science Foundation Career award, and an Alfred P. Sloan Foundation fellowship. Rus earned her PhD in computer science from Cornell University.

View full bio

The digitization of practically everything coupled with the mobile Internet, the automation of knowledge work, and advanced robotics promises a future with democratized use of computation and wide-spread use of customization and data-driven decision making. Recognizing this extraordinary transformation, MIT launched the Schwarzman College of Computing (SCC). In this talk I will introduce the SCC and discuss recent advances in robotics, machine learning, and artificial intelligence. I will also present results from three projects that are addressing current societal needs in the war against Covid-19.

11:40am - 12:20pm

Predictive and Prescriptive Analytics at Scale Rahul Mazumder Robert G. James Career Development Professor Associate Professor, Operations Research and Statistics MIT Sloan School of Management and Operations Research Center Rahul Mazumder Robert G. James Career Development Professor Associate Professor, Operations Research and Statistics MIT Sloan School of Management and Operations Research Center

Rahul Mazumder is the Robert G. James Career Development Professor and associate professor of operations research and statistics at the MIT Sloan School of Management. He is also affiliated with MIT's Operations Research Center and Center for Statistics and Data Science. His research interests are in statistical machine learning and large scale mathematical optimization; and their applications in recommender systems, computational biology, computational social science and finance. Before joining MIT, Mazumder was an assistant professor at Columbia University. He earned a PhD in statistics from Stanford University.

View full bio

Two key ingredients in an analytics pipeline are the predictive and prescriptive components. Usually, the former involves obtaining high-quality predictions from a machine learning algorithm, and the latter takes the resulting predictions as inputs to an optimization problem for downstream decision-making. Key challenges remain in addressing scalability concerns for each of these tasks. I will discuss some recent developments for creating large-scale algorithms for Gradient Boosting Machines and relatives, that also allow for interpretability and model compactness. For the prescriptive part, I will present new algorithms to solve extreme scale linear programs involving trillions of decision variables, that arise in several web-applications such as email volume optimization, matching, promotion optimization, ranking and recommender systems.

This presents joint research with members of my research group, and collaborators at Google and LinkedIn.

12:20pm - 1:00pm The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks Michael Carbin Assistant Professor, Department of Electrical Engineering and Computer Science Lead, Programming Systems Group Michael Carbin Assistant Professor, Department of Electrical Engineering and Computer Science Lead, Programming Systems Group

Michael Carbin is an assistant professor in MIT's Department of Electrical Engineering and Computer Science and a principal investigator at the Computer Science and Artificial Intelligence Laboratory, where he leads the Programming Systems Group. His group investigates the semantics, design, and implementation of systems that operate in the presence of uncertainty in their environment (perception), implementation (neural networks or approximate transformations), or execution (unreliable hardware). Carbin has received a Sloan Research Fellowship, a Facebook Research Award, a Google Faculty Research Award and an NSF Career Award. He earned a BS in computer science from Stanford University, and an MS and PhD in electrical engineering and computer science from MIT.

View full bio

Neural network pruning techniques can reduce the parameter counts of trained networks by over 90%, decreasing storage requirements and improving computational performance of inference without compromising accuracy. However, contemporary experience is that the sparse architectures produced by pruning are difficult to train from the start and, instead, training must first begin with large, overparameterized networks.

In this talk, I'll present our work on The Lottery Ticket Hypothesis, showing that at a standard pruning technique, iterative magnitude pruning, naturally uncovers subnetworks that are capable of training effectively from early in training. These subnetworks hold out the promise of more efficient machine learning methods, including inference, fine-tuning of pre-trained networks, and sparse training.

Part 2: Thursday, July 16, 2020

11:00am - 11:40am

Why Do ML Models Fail? Aleksander Madry Professor of Computer Science Director of the Center for Deployable ML Computer Science and Artificial Intelligence Laboratory



Aleksander Madry Professor of Computer Science Director of the Center for Deployable ML Computer Science and Artificial Intelligence Laboratory

Aleksander Madry is a professor of computer science in MIT's Department of Electrical Engineering and Computer Science and a principal investigator at the Computer Science and Artificial Intelligence Laboratory. He is also director of the MIT Center for Deployable Machine Learning. His research interests span algorithms, continuous optimization, the science of deep learning, and developing reliable, trustworthy and secure machine learning systems. Before coming to MIT, he was a postdoc at Microsoft Research New England and on the faculty of EPFL in Switzerland. His honors include an NSF Career Award, an Alfred P. Sloan Research Fellowship and the European Association for Theoretical Computer Science's Presburger Award. Madry earned an undergraduate degree in theoretical physics and computer science from University of Wroclaw, and a PhD in computer science from MIT.

View full bio

Our current machine learning models achieve impressive performance on many benchmark tasks. Yet, these models remain remarkably brittle and susceptible to manipulation.

Why is this the case?

In this talk, we take a closer look at this question, and pinpoint some of the roots of this observed brittleness. Specifically, we discuss how the way current ML models "learn" and are evaluated gives rise to widespread vulnerabilities, and then outline possible approaches to alleviate these deficiencies.

BlockFLow: An Accountable and Privacy-Preserving Approach to Federated Learning Lalana Kagal Principal Research Scientist Computer Science & Artificial Intelligence Laboratory Lalana Kagal Principal Research Scientist Computer Science & Artificial Intelligence Laboratory

Lalana Kagal is a principal research scientist at MIT's Computer Science and Artificial Intelligence Lab and Internet Policy Research Initiative, where she co-directs the Decentralized Information Group. She is also a research fellow at the Web Science Research Institute and an editor-in-chief of the Journal of Web Semantics. Her research focuses on modeling how social norms and legal rules work in society in order to automate policy compliance in information systems. She is currently exploring various facets of information management and policy, such as the development of new paradigms for exploring and integrating distributed data, privacy aware analysis of big datasets, and accessibility of mobile apps. She earned PhD from the University of Maryland.

Federated learning enables collaborating agents to develop a shared model without requiring them to share their underlying data. However, naive implementations are susceptible to privacy and security threats. As sensitive data is not shared, the privacy risk is reduced. However, it is still possible to leak information about the training dataset from the model's weights or parameters. Also, malicious agents who train on random data, or worse, try to poison the model, can weaken the shared model and must be identified and held accountable.

In this talk, I will describe an initial implementation of an accountable federated learning system that is privacy-preserving. BlockFLow incorporates differential privacy to reduce information leakage, introduces a novel auditing mechanism for evaluating model contribution, and uses Ethereum smart contracts to incentivize good behavior. Its primary goal is to reward agents proportional to the quality of their contribution while protecting the privacy of the underlying datasets and being resilient to malicious adversaries.

12:20pm - 1:00pm

Watch and help: a platform for social perception and Human AI collaboration Antonio Torralba Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science Head AI+D (AI & Decision Making) faculty, EECS Computer Science and Artificial Intelligence Laboratory



Antonio Torralba Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science Head AI+D (AI & Decision Making) faculty, EECS Computer Science and Artificial Intelligence Laboratory

Antonio Torralba is the Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science at MIT. He also heads the faculty of artificial intelligence and decisionmaking in the MIT Schwarzman College of Computing. Previously, he led the MIT Quest for Intelligence as its inaugural director and the MIT–IBM Watson AI Lab as its MIT director. He researches computer vision, machine learning and human visual perception, with an interest in building systems that can perceive the world as humans do. He has received an NSF Career Award, the International Association for Pattern Recognition's JK Aggarwal Prize, a Frank Quick Faculty Research Innovation Fellowship and a Louis D. Smullin ('39) Award for Teaching Excellence. Torralba earned a BS from Telecom BCN, Spain, and a PhD from the Institut National Polytechnique de Grenoble, France.

View full bio

In this talk, I will describe Watch-And-Help, a platform for training and testing social intelligence in agents. In this platform, an AI agent needs to help a human-like agent perform a complex household task efficiently. To succeed, the AI agent needs to i) understand the underlying goal of the task by watching a single demonstration of the human-like agent performing the same task (social perception), and ii) coordinate with the human-like agent to solve the task in an unseen environment as fast as possible (human-AI collaboration). Experimental results demonstrate that in order to achieve success in the challenge, an AI agent has to accurately understand and predict the human-like agent's behaviors, and adapt its collaborative plan accordingly in novel environments.