

**MIT Industrial Liaison Program Faculty Knowledgebase Report**

---

AI and Autonomy @MIT

---

---

February 23, 2021 10:00 am - 12:10  
pm

---

10:00am – 10:15am

Reverse Engineering Visual Intelligence  
James DiCarlo  
Peter de Florez Professor of Neuroscience  
Co-Director, MIT Quest for Intelligence  
Head, Brain and Cognitive Sciences  
Investigator, McGovern Institute



James DiCarlo  
Peter de Florez Professor of Neuroscience  
Co-Director, MIT Quest for Intelligence  
Head, Brain and Cognitive Sciences  
Investigator, McGovern Institute

James DiCarlo is the Peter de Florez Professor of Neuroscience at MIT and co-director of the MIT Quest for Intelligence. He also heads the Department of Brain and Cognitive Science and is a principal investigator at the McGovern Institute for Brain Research. His research focuses on using computational methods to understand the brain's visual system, and with this knowledge, developing brain-machine interfaces to restore or augment lost senses. DiCarlo has received an Alfred P. Sloan fellowship, a Pew Scholar Award, and a McKnight Scholar Award. He earned a PhD in biomedical engineering, and an MD, from Johns Hopkins University.

[View full bio](#)

The brain and cognitive sciences are hard at work on a great scientific quest — to reverse engineer the human mind and its intelligent behavior. Yet these field are still in their infancy. Not surprisingly, forward engineering approaches that aim to emulate human intelligence (HI) in artificial systems (AI) are also still in their infancy. Yet the intelligence and cognitive flexibility apparent in human behavior are an existence proof that machines can be constructed to emulate and work alongside the human mind. I believe that these challenges of reverse engineering human intelligence will be solved by tightly combining the efforts of brain and cognitive scientists (hypothesis generation and data acquisition), and forward engineering aiming to emulate intelligent behavior (hypothesis instantiation and data prediction). As this approach discovers the correct neural network models, those models will not only encapsulate our understanding of complex brain systems, they will be the basis of next-generation computing and novel brain interfaces for therapeutic and augmentation goals (e.g., brain disorders). In this session, I will focus on one aspect of human intelligence — visual object categorization and detection — and I will tell the story of how work in brain science, cognitive science and computer science converged to create deep neural networks that can support such tasks. These networks not only reach human performance for many images, but their internal workings are modeled after— and largely explain and predict — the internal workings of the primate visual system. Yet, the primate visual system (HI) still outperforms current generation artificial deep neural networks (AI), and I will show some new clues that the brain and cognitive sciences can offer. These recent successes and related work suggest that the brain and cognitive sciences community is poised to embrace a powerful new research paradigm. More broadly, our species is the beginning of its most important science quest — the quest to understand human intelligence — and I hope to motivate others to engage that frontier alongside us.

10:15am – 10:30am

Towards AI that Learns to Write Code  
Armando Solar-Lezama  
Professor, Electrical Engineering and Computer Science  
Associate Director and COO, CSAIL  
Armando Solar-Lezama  
Professor, Electrical Engineering and Computer Science  
Associate Director and COO, CSAIL

Armando Solar-Lezama is a Professor in the department of Electrical Engineering and Computer Science at MIT and is also Associate Director and COO of the Computer Science and Artificial Intelligence lab. He also leads the NSF Funded Expeditions project "Understanding the World Through Code", a large multi-institution effort that works on applying neurosymbolic reasoning techniques to support scientific discovery.

[View full bio](#)

In this talk, I describe some recent work on neurosymbolic program synthesis, a new approach to program synthesis that combines machine learning and symbolic reasoning about programs in order to build new kinds of program synthesizers.

10:30am – 10:45am

Generative Models as Data++  
Phillip Isola  
Assistant Professor, Electrical Engineering and Computer Science  
Phillip Isola  
Assistant Professor, Electrical Engineering and Computer Science

Phillip Isola is the Bonnie & Marty Tenenbaum Career Development Chair Assistant Professor in EECS at MIT. He completed his Ph.D. in Brain & Cognitive Sciences at MIT, advised by Edward Adelson, then spent two years as a postdoc at UC Berkeley, under Alexei Efros, and one year as a visiting research scientist at OpenAI. He was a recipient of an NSF graduate fellowship and an NSF postdoctoral fellowship. His work focuses on why we represent the world the way we do, and how we can replicate these abilities in machines.

The last few years have seen an explosion of powerful generative models -- models that can synthesize fake faces, landscapes, text, audio and more. The results are fascinatingly realistic, but it's not immediately clear what they are useful for. We already have billions of images of faces, why do we need a model to make more? I will argue that the real power of these models is not their ability to make random fake data but that they make a new kind of data: data that comes bundled up with controllable latent variables. I will focus on deep generative models of images, which synthesize a photo given an input vector of latent variables. The latent variables are knobs that control what the output will look like: a user can tune them to change the lighting conditions in a photo, rotate objects, add or remove elements of a scene, and much more. I will show applications in image editing and scientific data visualization, and I will suggest that this new kind of data, sampled from deep generative models, can be thought of as data++: it looks just like regular data, but comes with extra functionality.

10:45am – 11:00am

Engineering and reverse-engineering intelligence with the MIT Inference Stack  
Vikash Mansinghka  
Principal Research Scientist, Brain and Cognitive Sciences



Vikash Mansinghka  
Principal Research Scientist, Brain and Cognitive Sciences

Vikash Mansinghka is a Principal Research Scientist at MIT, where he leads the Probabilistic Computing Project. Vikash holds S.B. degrees in Mathematics and in Computer Science from MIT, as well as an M.Eng. in Computer Science and a PhD in Computation. He also held graduate fellowships from the National Science Foundation and MIT's Lincoln Laboratory. His PhD dissertation on natively probabilistic computation won the MIT George M. Sprowls dissertation award in computer science, and his research on the Picture probabilistic programming language won an award at CVPR. He co-founded three VC-backed startups: Prior Knowledge (acquired by Salesforce in 2012), Empirical Systems (acquired by Tableau in 2018), and Common Sense Machines (co-founded in 2020). He served on DARPA's Information Science and Technology advisory board from 2010-2012, currently serves on the editorial boards for the Journal of Machine Learning Research and the journal Statistics and Computation, and co-founded the International Conference on Probabilistic Programming.

[View full bio](#)

Humans see, think, and learn far more robustly, flexibly, and efficiently than current AI systems. Can we achieve human-level performance, using AI architectures that people can understand and trust?

We have been developing a new AI programming model that narrows the gaps between human and machine intelligence by unifying probabilistic, symbolic, and neural approaches. This talk will focus on three emerging AI capabilities, developed in partnership with industry: (i) inferring 3D objects from 2D images, using models of human common sense; (ii) deduplicating and cleaning dirty, denormalized databases with millions of records, using models of human domain expertise; (iii) enabling people without statistics training to solve data analysis problems, by emulating judgment calls made by human statisticians. It will highlight the common AI engineering principles and computing abstractions underlying these diverse capabilities, as well as ongoing opportunities for MIT-industry partnership.

11: 00am – 11:15am

Electrochemical artificial synapses for brain-inspired computing  
Bilge Yildiz  
Professor, Nuclear Science and Engineering  
Professor, Materials Science and Engineering  
Bilge Yildiz  
Professor, Nuclear Science and Engineering  
Professor, Materials Science and Engineering

Bilge Yildiz is a professor in the Nuclear Science and Engineering and the Materials Science and Engineering Departments at Massachusetts Institute of Technology (MIT), where she leads the Laboratory for Electrochemical Interfaces. She received her PhD degree at MIT in 2003 and her BSc degree from Hacettepe University in Turkey in 1999. After working at Argonne National Laboratory as research staff, she returned to MIT as an assistant professor in 2007. Her leadership responsibilities at MIT include the Low Carbon Energy Center on Materials in Energy and Extreme Environments, and one of the Integrated Research Groups of MIT's NSF sponsored Materials Research Science and Engineering Center. Yildiz's research focuses on laying the scientific groundwork to enable next generation electrochemical devices for energy conversion and information processing. The scientific insights derived from her research guide the design of novel materials and interfaces for efficient and durable solid oxide fuel cells, electrolytic water splitting, brain-inspired computing, and solid state batteries. Her approach combines computational and experimental analyses of electronic structure, defect mobility and composition, using *in situ* scanning tunneling and X-ray spectroscopy together with first-principles calculations and novel atomistic simulations. Yildiz's work has made significant contributions to advancing the molecular-level understanding of oxygen reduction and oxidation kinetics on solid surfaces, and of ion and electron transport, under electro-chemo-mechanical conditions. The scientific insights derived from her research guide the design of novel material chemistries for efficient and durable solid oxide fuel cells, thermo-/electro-chemical splitting of H<sub>2</sub>O and CO<sub>2</sub>, high energy density solid state batteries, and red-ox based memristive analog information processing. Her teaching and research efforts have been recognized by the Argonne Pace Setter (2016), ANS Outstanding Teaching (2008), NSF CAREER (2011), IU-MRS Somiya (2012), the ECS Charles Tobias Young Investigator (2012), and the ACerS Ross Coffin Purdy (2018) awards.

[View full bio](#)

Deep learning is a hugely successful and powerful algorithm for machine learning applications such as computer vision and natural language processing. However, the training of these neural networks is limited by the traditional von Neumann architecture of our current CPUs and GPUs. Shuttling data back and forth between the separate memory and computation units in such architecture results in significant energy consumption; many orders of magnitude greater than the energy consumption in human brain. Our research focuses on designing materials and hardware that can instead perform data storage and computation in a single architecture using ions, inspired by the human brain. In the project that I will present as an example, we have designed a protonicelectrochemical synapse that changes conductivity deterministically by current-controlled shuffling of dopant protons across the active device layer; resulting in energy consumption on par with biological synapses in the brain. Through these strategies, we exhibit a path towards neuromorphic hardware that has high yield and consistency, performs data storage and computation in a single device, and uses significantly lesser energy as compared to current systems.

11:15am – 11:25am

MIT Startup Exchange Lightning Talks

Hosta Labs - Automated digital structural assessment

Henriette Fleischmann  
Co-founder & COO  
[Hosta Labs](#)

Leela - Automatic scene understanding for workplace safety & compliance

Cyrus Shaoul  
CEO  
[Leela](#)

Farmwise Labs - Autonomous weeding robot for vegetable farms

Pauline Canteneur  
Business Strategist  
[Farmwise Labs](#)

OnSpecta - Unique Virtualization Technology for Best Inference Hardware Performance

Victor Jakubiuk  
Co-founder and CTO  
[OnSpecta](#)

Nara Logics - Brain-like AI platform for digital advisor

Jana Eggers  
CEO  
[Nara Logics](#)

11:25am – 11:40am

Building Dependable and Verifiable Autonomous Systems  
Chuchu Fan  
Wilson Assistant Professor, Aeronautics and Astronautics  
Chuchu Fan  
Wilson Assistant Professor, Aeronautics and Astronautics

Chuchu Fan is the Wilson assistant professor of Aeronautics and Astronautics at MIT. Before joining MIT in 2020, she was a postdoctoral researcher in the Department of Computing and Mathematical Sciences at the California Institute of Technology. She received her Ph.D. in 2019 from the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. She received her Bachelor's degree from Tsinghua University, Department of Automation in 2013. Her research interests are in the areas of formal methods, machine learning, and control for safe autonomy.

[View full bio](#)

The introduction of machine learning (ML) and artificial intelligence (AI) creates unprecedented opportunities for achieving full autonomy. However, learning-based methods in building autonomous systems can and do fail, due to poor quality data, modeling errors, the coupling with other agents, and the complex interaction with human and computer systems in modern operational environments. In this talk, I will present several of our recent efforts that address this challenge and advance the use of AI and ML techniques to enable the design of provably dependable and safe autonomous systems. The topics I am going to cover are 1. How to generate safety certificates for complex autonomous systems; 2 How to learn certified safe decision and control; and 3. How to build certified correct simulators.

11:40am – 11:55am

Autonomous Flight in Urban environments: Challenges for Perception and Planning  
Nicholas Roy  
Bisplinghoff Professor, Aeronautics & Astronautics  
Director of Quest Systems Engineering, MIT Quest for Intelligence  
Nicholas Roy  
Bisplinghoff Professor, Aeronautics & Astronautics  
Director of Quest Systems Engineering, MIT Quest for Intelligence

Nicholas Roy is the Bisplinghoff Professor of Aeronautics & Astronautics and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology. He has a B.Sc. in Physics and Cognitive Science and an M.Sc. in Computer Science, both from McGill University. He received his Ph. D. in Robotics from Carnegie Mellon University in 2003. He has made research contributions to planning under uncertainty, machine learning, human-computer interaction and aerial robotics. He founded and led Project Wing at Google [X] from 2012-2014. He is currently the Director of Quest Systems Engineering in MIT's Quest for Intelligence.

[View full bio](#)

Small UAS have tremendous promise for providing many different services in urban environments, such as inspection or delivery. But, autonomous flight in the urban environment also brings substantial challenges in terms of sensing, perception and decision making. Small UAS need to be able to understand where they are and what is around them to a much greater degree than ever before. I will talk about recent progress in perception and planning for small UAS, and what the next generation of onboard autonomy may look like.

11:55am – 12:10pm

The MIT MOOS-IvP Open Source Marine Autonomy Project

Michael Benjamin  
Research Scientist, Mechanical Engineering  
Michael Benjamin  
Research Scientist, Mechanical Engineering

Michael Benjamin is a research scientist in the Center for Ocean Engineering, a part of the Department of Mechanical Engineering at MIT. He is also a member of the Laboratory for Autonomous Marine Sensing Systems and the Marine Robotics Group in the Computer Science and Artificial Intelligence Laboratory. Until December 2010, he was with the Naval Undersea Warfare Center in Newport Rhode Island.

Benjamin's work is focussed on algorithms and software for autonomous marine vehicles, some of which are shown to the right. In 2007 he founded moos-ivp.org at MIT, hosting the MOOS-IvP open source project in marine autonomy software. A key part of this project is the use of a behavior based architecture for autonomous decision-making using multi-objective optimization with interval programming for reconciling competing behaviors. This work is driven by the belief that multi-objective optimization is a fundamental component of robust decision-making. Formulating a decision-making problem into distinct specialized components also promotes the development of an autonomous system with contributions from varied developers and organizations. It also allows for a system comprised of public open source general-purpose code alongside non-public specialized code.

[View full bio](#)

Unmanned underwater and surface vessels hold enormous potential in understanding our ocean as remote ocean monitoring and sensing systems. Autonomous surface vessels may also delivery other platforms or act as communication and navigation aids for remotely deployed underwater vehicles. The same autonomy technology may one day soon be used to deploy lightly-crewed or completely unmanned surfaces vessels for transportation.

In any of these applications, reasoning about collision avoidance with other surface vessels is a key aspect of ensuring safe operation. Typically an autonomy system reasoning about collision avoidance in marine surface vehicles includes consideration of the COLREGS or the Coast Guard Collision Regulations. However, the COLREGS were written for humans and prescribe actions to be taken to avoid collisions with a single other vessel. It is assumed that humans will apply common sense to extenuating circumstances, and generalize reasonably when multiple vehicles need to be avoided simultaneously. Humans are resilient in this manner, routinely handling arbitrarily complex and unique situations.

To enable this resiliency in an automated system requires an autonomy architecture that also extends to an arbitrary number of simultaneous vessels and mission considerations. Over the last 20 years, we have designed such an autonomy system from the ground up, based on our developed mathematical model for multi-objective optimization called Interval Programming (IvP).

This architecture is known as MOOS-IvP and is has been distributed at MIT under an open source license since 2006 at [www.moos-ivp.org](http://www.moos-ivp.org). The public code-base now represents roughly 40 work years of development effort over many dozens of autonomy and support modules. The IvP mathematical model supports a behavior-based architecture extendible by users for their own missions and platforms, allowing for commercial or proprietary extensions layered on top of the publicly available code-base. The first version of the COLREGS collision avoidance modules was included in the 2017 release. MOOS-IvP has been used around the world on dozens of unmanned marine platforms in academia, industry and defense.